# WRESTLING WITH COSMIC SCALES

by

Victor Carl Chan

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy

David A. Dunlap Department of Astronomy and Astrophysics
University of Toronto

Wrestling with Cosmic Scales

Victor Carl Chan
Doctor of Philosophy

David A. Dunlap Department of Astronomy and Astrophysics
University of Toronto
2023

# Abstract

The ΛCDM model of cosmology has had resounding success in describing the evolution of the Universe as a whole throughout cosmic history. This has been corroborated by a wide variety of experiments and observations across all cosmic scales, and we have entered an era of precision experiments that allow us to scrutinize the model's predictions at smaller scales than ever. This thesis presents a collection of studies on the impact of existing and upcoming observational data on constraints of extensions to the standard cosmological model through thorough analyses of experiment systematics, and the development of new methodology. We begin with a study of the *Gaia* mission's parallax zero point, and we show that *Gaia* data can independently secure uncertainties of geometric anchors in future determinations of the local distance ladder. This will be useful for verifying or disputing tensions between early- and late-time estimates of the Hubble constant, which describes the present rate of the Universe's expansion. We continue with a study of the impact of detector crosstalk systematics on upcoming searches for primordial gravitational wave signatures in the cosmic microwave background, which require unprecedented levels of sensitivity and precision. The results of our work directly inform the development and design of future experiments. We then present the development of the Small Correlated Against Large Estimator for directly quantifying the small-scale statistics of gravitational lensing in the cosmic microwave background without the need for traditional reconstruction of the lensing deflection field. Our method can be used to improve our understanding of dark matter and matter clustering at the scales of galaxy clusters, and it can help to provide a detection of the minimum neutrino mass in upcoming cosmic microwave background experiments.

爸爸，媽媽，姊姊，多謝你們永遠支持我。

# Acknowledgements

My deepest gratitude goes to my supervisors Prof. Renée Hložek and Prof. Jo Bovy. Without either of you, I could not have become even a shadow of the scientist that I am today. As your student, I feel as if I am prepared to take on pretty much any scientific problem that I am interested in. To Renée, you worked extremely hard to tailor your advisory style to one that I was receptive to, and it has definitely been a deciding factor in the completion of my thesis. Your expertise, patience and guidance in my studies, research, and even sometimes personal struggles, have been invaluable to my growth as a scientist and as a person. To Jo, you took the steps to work with me even though my scientific interests were not comfortably within your *many* areas of expertise. You always kept me grounded, and were always keen to help solve the problems in our research. I am **SO** lucky to have both of you as role-models and mentors! I hope you are as proud of me as I am proud to call both of you my advisors.

It has also been an absolute pleasure to have a wonderful advisory committee with Prof. Keith Vanderlinde and Prof. Marten van Kerkwijk. Keith's AST121 course inspired to pursue astronomy (specifically cosmology) during my first year of undergraduate studies. Keith and Marten were also great mentors for my undergraduate thesis, and I am grateful to have had the opportunity to be under their guidance for my graduate studies. Thank you for always providing pragmatic advice for my careers and goals, while also being supportive of my personal circumstances.

I would like to extend extra thanks to Prof. Alexander van Engelen and Prof. Joel Meyers for their invaluable collaboration and mentorship. I felt as if I had four advisors at many times during my PhD, and I am truly honoured to have had the opportunity to do such exciting work with you.

I would like to additionally thank Prof. Dick Bond, Prof. Gwendolyn Eadie, Prof. Ting Li, Prof. Chris Matzner, Prof. Maria Drout, Prof. Josh Speagle, Dr. Christina Peters, Dr. Mubdi Rahman, Dr. Yilun Guan, and Dr. Keir Rogers for thoughtful discussions about the thesis content.

Graduate school would have been a much larger burden if I were to navigate it alone. The friendships forged within my cohort will likely last a lifetime. Thank you to Ariel, Jen, Thierry, Alex L., Abhinav, Georg, and Alex R. for being amazing friends (and especially for sharing memes). Emily, being GASA co-president with you may have been difficult at times, but I will always look back at that era fondly when I remember how hard we worked together. Taylor, you weren't part of my cohort, but struggling through our final year together was very fun. It was inspiring to see your growth this year, and that definitely helped to drive me through what I needed to do. Thank you for not only lending an ear whenever I asked, but also for giving your earnest takes. Tomás, what is your problem gordito?

Special thanks go out to the past and present members of the Hnauts and the Bovy Research Group (BANTER?) for commiserating in our research and grad-school troubles. Thank you to Anita, Jason, Morgan, Margaret, James, Aarya, Harrison, Martine, Nathaniel, Henry, Simran, and Bolin. To Adiv, Anna, and Ayushi, thanks for being good friends and for giving great advice. Natalie, you were never 'officially' my graduate student mentor, but you took it upon yourself to be one for me anyway. I am extremely grateful for your guidance and friendship.

I would be terribly remiss to omit the friends I've come to grow close with in the past year. To Ayush, Sam, Steffani, Phil, Fraser, Mairead, Anika, Mark, Shannon, Ziggy, Sarik, Alex, Dang, Jacob M., and Michael: having you around the department was monumental in regaining my joy in academics, to the point that I've backtracked on my original plans to leave astronomy. Witnessing your hard work and energy gives me the inspiration and motivation to keep going. I will miss the regular crossword alerts, but I will **not** stop doing the crossword. Special shoutout to Will Shortz, and whoever decided to invent the rebus.

I'd like to thank my friends outside of school for providing me with plenty of emotional energy to get through this degree. To Noel, Dan, Lillian, Chanakya, Hilton, Tim, Tony, Kyle, Ryan, Hughie: thanks for being there for me during the highest highs and the lowest lows. Your collective support means more to me than you could know. Thank you as well to Vismay, Visal, Bruce, and Sylar for staying in touch (and also for memes). A shoutout to my therapist Julieta Strugo for teaching me how to understand myself better, which in turn has helped me grow into a better person.

To Clarence, my closest confidant. We've been there for each other through thick and thin. The reason why I'll always be on your side is because I know you'll do the same for me. You always know exactly when and how to take my mind off of stressful situations, and you've saved me from myself more than you could know. Thank you for being my best friend.

We're nearing the end here! I give thanks to my extended family for always being supportive of my academic endeavours. This thesis was written in loving memory of grandpa Foo Kan. Thank you to grandma Sui Lan for always making sure I've had enough to eat. To my cousins Marrick, Alvin, and Carmen: thank you for being there for me and for listening to my troubles. To Bonnie: I know how much you look up to me, and I hope that if you see this, you'll be inspired to achieve your dreams as well.

To my sister Heather. You may be on the other side of the planet, but that never prevents you from providing me with limitless encouragement. You have always been there for me, especially when no one else could be, and definitely at times that were inconvenient for you. I am forever grateful to have the best big sister to look up to.

To my father Pak Hon, and to my mother Wendy. You have done everything in your power to be the most supportive parents in existence. You have gone above and beyond to provide me with the best you could. I may occasionally push the limits of what normal parents would tolerate, but I understand that you would offer me the world. I am extremely proud to call you my parents.

# Contents

# List of Tables

# List of Figures

# Chapter 1

# Introduction

## 1.1 Timeline and History of the Universe

CONTEMPORARY models of physical cosmology are generally constructed within the paradigm of the theory of general relativity (Einstein, 1916, 1917). Detailed courses in modern cosmology are provided by textbooks from which much of the introduction content is adapted in this chapter (Dodelson, 2003; Dodelson & Schmidt, 2020; Baumann, 2022). Models of cosmology follow spacetime metrics that describe homogeneous and isotropic geometries to be in line with observations. There exists a set of maximally-symmetric spacetime metrics that satisfy these properties, and within it, three subclasses can be classified with a parameter $k$ which controls the curvature of the spacetime geometry, as well as a parameter $a$ which scales the spatial coordinates. The scale parameter is usually allowed to vary with time: $a(t)$. This is commonly referred to as the FLRW metric (derived independently by its namesakes: Friedmann 1922; Lemaître 1933; Robertson 1935; Walker 1935), and it can be expressed in spherical polar coordinates as follows

$$\mathrm{d}s^2 = \mathrm{d}t^2 - a^2(t)\left[\frac{\mathrm{d}r^2}{1 - kr^2} + r^2\mathrm{d}\Omega^2\right], \tag{1.1}$$

where $\mathrm{d}\Omega^2 = \mathrm{d}\theta^2 + \sin^2\theta\mathrm{d}\phi^2$, and in units of the speed of light $c = 1$. The spatial curvature $k$ defines the specific subclass of maximally-symmetric geometries that exist within the FLRW metric. A spatially flat geometry can be defined with $k = 0$, which corresponds to Minkowski space in the absence of the scale parameter. Contemporary experiments report results that are consistent with a spatially flat Universe (Planck Collaboration et al., 2020a).

The scale parameter $a(t)$ is typically defined such that at present time $t_0$, $a_0 = a(t_0) = 1$, and solutions to the Einstein equations using the FLRW metric indicate that the evolution of the scale parameter are dependent on the energy density $\rho$ and pressure $P$ of all the contents of the Universe. In particular, this set of solutions is typically referred to as the Friedmann equations, shown below

$$H^2 \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{k}{a^2}, \tag{1.2}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3P). \tag{1.3}$$

One can define the critical density of the Universe today $\rho_{\mathrm{crit},0}$ as the density of today's Universe required for a flat geometry (setting $k = 0$ in Equation (1.2))

$$\rho_{\mathrm{crit},0} = \frac{3H_0^2}{8\pi G}. \tag{1.4}$$

It is simplest to consider the evolution of the scale parameter by considering the components of the Universe's energy density individually. The reaction of each component of the energy density then reacts uniquely to the expansion through the continuity equation. A simple way to obtain it is to apply a time derivative to both sides of Equation (1.2) for a flat Universe ($k = 0$), and substitute in the original version of Equation (1.2) (also flat) as well as Equation (1.3):

$$\frac{8\pi G}{3}\dot{\rho} = 2\frac{\dot{a}}{a}\left(\frac{\ddot{a}}{a} - \frac{\dot{a}^2}{a^2}\right) \tag{1.5}$$

$$= -\frac{8\pi G}{3}H(\rho + 2P + 2\rho) \tag{1.6}$$

$$\dot{\rho} = -3H(\rho + P)\,. \tag{1.7}$$

The continuity equation Equation (1.7) describes the way energy is conserved throughout the Universe's expansion, and combining it with the equation of state $w = P/\rho$ and solving for $\rho$ with Equation (1.7) yields $\rho \propto a^{-3(1+w)}$. For each major component of the energy density, this corresponds to:

$$\rho \propto \begin{cases} a^{-4} & \text{radiation}\,, \\ a^{-3} & \text{matter}\,, \\ a^{0} & \text{vacuum (dark) energy}\,. \end{cases} \tag{1.8}$$

By defining dimension-less density parameters $\Omega = \rho/\rho_{\mathrm{crit}}$, one can rework Equation (1.2) for a Universe containing radiation (subscript $r$), matter (subscript $m$), and vacuum energy (subscript $\Lambda$; also referred to as dark energy; see Einstein 1917; Zel'dovich 1968; Riess et al. 1998; Perlmutter et al. 1999 for derivation from field equations and observational evidence; see Li et al. 2013 for a review):

$$\frac{H^2(a)}{H_0^2} = \Omega_{r,0}\left(\frac{a}{a_0}\right)^{-4} + \Omega_{m,0}\left(\frac{a}{a_0}\right)^{-3} + \Omega_{k,0}\left(\frac{a}{a_0}\right)^{-2} + \Omega_{\Lambda,0}. \tag{1.9}$$

For a Universe with $k \neq 0$, the density parameter associated with spatial curvature (subscript $k$) is defined with respect to the sum of all other density parameters such that $\Omega_k = 1 - \sum \Omega_X$ for $X \in \{r, m, \Lambda\}$. The primary implication of Equation (1.9) is that the evolution of the Universe's scale factor $a(t)$ is uniquely affected according to the abundance, distribution and behaviour of each respective component of the Universe's contents. Observational evidence (such as the cosmic

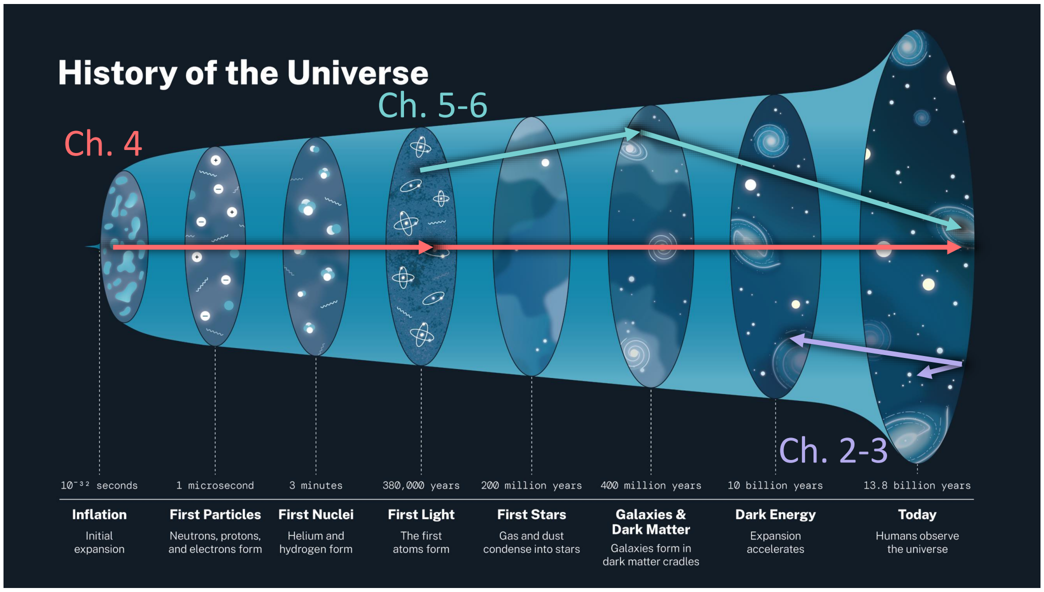

Figure 1.1: A simplified timeline of the Universe's history adapted from NASA (2023). The topics of this thesis span the entire range of the Universe's history, from the early Universe to the present day. Chapters 2-3 focus on improving local distance ladder measurements with *Gaia* data. Chapter 4 focuses on characterizing detector systematics in upcoming searches for inflationary signals in the polarized cosmic microwave background. Chapters 5-6 focus on the development and application of a novel estimator for the small-scale statistics of gravitational lensing in the cosmic microwave background.

microwave background described in §1.2) support the expanding Universe model, and imply the Universe's origin from a hot Big Bang.

An important effect of cosmic expansion is that it induces cosmological redshift to photons travelling extremely large distances, defined as a function of its observed wavelength $\lambda_0$ and emitted wavelength $\lambda$:

$$z = \frac{\lambda_0 - \lambda}{\lambda}. \tag{1.10}$$

The redshift can be related to the scale factor through the relation $1 + z = a^{-1}$. The expansion of space in our Universe is monotonic, so cosmic time can be described with any of $t$, $a$, or $z$. It is also useful to define the comoving distance $\chi$ and conformal time $\eta$. The comoving distance is a non-observable parameterization of cosmic distance that takes into account its expansion. In units of the speed of light $c = 1$, the comoving distance to a source at redshift $z$ and emitting light at time $t_1$ which we observe at $t_0$ is defined as:

$$\chi = \int_{t_1}^{t_0} \frac{\mathrm{d}t}{a(t)} = \int_0^z \frac{\mathrm{d}z}{H(z)}\,. \tag{1.11}$$

Similarly, conformal time is defined with respect to the time since the Big Bang $t = 0$:

$$\eta = \int_0^t \frac{\mathrm{d}t}{a(t)}\,. \tag{1.12}$$

In summary, the combination of Equation (1.9) with modern observations suggests that:

1. The Universe expands over time. That is to say that the distance between two points in space are expected to increase over time following $a(t)$,
2. The Universe's contents, which can be described by energy densities $\rho$ and $\Omega$, are diluted due to its expansion,
3. The Universe's expansion rate $H(t)$ is intricately tied to its contents at time $t$, and
4. The Universe had a hot, thermal origin followed by a period of expansion during which all cosmic structures such as stars, galaxies, galaxy clusters eventually formed. Refer to Figure 1.1 for a general timeline of the Universe's history.

The prevailing model of cosmology at present is spatially flat $\Lambda$CDM, in which dark energy ($\Lambda$) and cold dark matter (CDM) are the dominant components of the Universe's energy density today. It particularly imposes that CDM makes up the entirety of the dark matter in the Universe, and a majority of all matter. We further discuss dark matter models in §1.4.1. One of the attractive features of $\Lambda$CDM is that it can be described by a total of six free parameters, with all remaining physical cosmological parameters derived from some combination of those few (Planck Collaboration et al., 2020a). We provide a glimpse into the Universe's history in the context of $\Lambda$CDM in the following sections of the chapter. We begin with an extensive review of the cosmic microwave background in §1.2, with a particular focus on the effects of gravitational lensing from large-scale massive structures in the Universe. §1.3 provides a brief overview of cosmic distance measurements in the context of the Hubble constant $H_0 = H(z = 0)$. In §1.4, we bring into focus some of the questions that cosmologists are working to answer today, and we connect the subjects in §1.2-1.3 to the main content of this dissertation. Finally, we highlight some key statistical methods and tools used in this thesis in §1.5, and provide an outline of its contents in §1.6.

## 1.2 The Cosmic Microwave Background

The cosmic microwave background (CMB) is a radiation signal originating from the early Universe. This is when radiation was finally allowed to free-stream through the Universe after neutral hydrogen was first formed (commonly, but inaccurately, referred to as recombination). This occurred at a redshift of approximately $z \approx 1100$. The original form of the CMB was a mostly uniform blackbody of $\sim 3000$K which has since been cosmologically redshifted to $\sim 2.73$K (Mather et al., 1994). As a continuous signal observed in all directions, the CMB's signals have been studied by comparing correlation statistics across various angular scales on the sky. The CMB temperature power spectrum $C_\ell^{TT}$ (defined in §1.2.2) quantifies anisotropies in the CMB temperature signal at various angular scales. The observed anisotropies in the CMB temperature field today are on the order of $\sim 100\mu$K, and they have been studied with increasing ferocity over the last few decades. We begin this Section with an overview of CMB observations in §1.2.1, followed by a review of the physics contained within the CMB's features.

The signals present in the observed CMB can be broadly divided into two categories: primary and secondary anisotropies. Features of the CMB that were present as the photons were first free-streaming away from the surface of last scattering are considered primary anisotropies. The primary anisotropies in the CMB are sourced by density (also referred to as scalar) perturbations. Primary anisotropies reflect the characteristics of the CMB's uncontaminated form, and the power/cross spectra (defined as Equation (1.19) in §1.2.2) of the original CMB are heavily dependent on the geometry and contents of the Universe at recombination. The primary features of the CMB have been measured many times, and they have been shown to be consistent with ΛCDM (Bennett et al., 2013; Planck Collaboration et al., 2020a; Aiola et al., 2020; Sobrin et al., 2022). Further studies of the CMB's primary fluctuations seek uncover minor inconsistencies with ΛCDM, which would allow for extensions to the standard model that include novel physics and/or resolve tensions with other observations. Many contemporary experiments are designed with polarization sensitivity to characterize the CMB's polarized features, usually decomposed into $E$- and $B$-modes, which we define in §1.2.2. One example of an undetected feature of the primary CMB is the presence of $B$-mode polarization generated due to primordial gravitational waves sourced by inflation. ΛCDM predicts only $E$-mode polarization signals in the primary CMB, which are sourced from the Thomson scattering of photons in inhomogeneous media which made up the environment at recombination. Inflationary models naturally predict the existence of $B$-mode signals due to the gravitational waves (also referred to as tensor perturbations) created by the rapid expansion of space, but the predicted amplitudes of primary $B$-mode fluctuations are significantly lower than the $E$-mode signal, or even secondary sources of $B$-mode polarization signals. A deeper discussion of the primary anisotropies of the CMB is provided in §1.2.2.

Features imprinted on CMB photons *after* recombination as they travel towards us are considered secondary anisotropies. These are induced by the large-scale structures (LSS) in between us and the surface of last scattering through a variety of mechanisms. The physical processes that CMB photons can experience also change as the LSS evolves over cosmic time. While secondary anisotropies have historically been considered nuisances for the purposes of studying the primary CMB, deeper study of CMB's secondary anisotropies provide a unique probe of the development and evolution of LSS in our Universe. We briefly describe a range of CMB secondary anisotropies in §1.2.3. The weak gravitational lensing of CMB photons from LSS is an example of a CMB secondary feature, and it is a key focus of this dissertation. A review of CMB lensing is provided in §1.2.4. We finish this Section with a detailed description of a full observation and analysis procedure in §1.2.5.

### 1.2.1 History of Cosmic Microwave Background Observations

Evidence for the cosmic microwave background was indirectly observed as early as 1940, when McKellar (1940) observed CN absorption lines from a stellar spectrum to infer a 'maximum "effective" temperature of interstellar space' to be 2.7 K. The cosmic microwave background was first directly detected by Penzias & Wilson (1965), who jointly reported an excess radio blackbody temperature of $T_{\mathrm{CMB}} = 3.5 \pm 1.0$ K that could be observed in all directions across the sky. They eventually realized that their radio signature closely matched a radiative blackbody predicted by Dicke et al. (1965) and Peebles (1965) from the principles of evolutionary cosmology described in §1.1. Their pioneering work confirmed the existence of a smooth, isotropic radiation field predicted by contemporary cosmological models. A major subsequent breakthrough came from the **Co**smic **B**ackground **E**xplorer (COBE; Smoot et al. 1991; Efstathiou et al. 1992; Mather et al. 1994), which was a satellite mission able to detect anisotropic structure in the CMB at a temperature contrast of approximately $\Delta T/T \gtrsim 10^{-4}$ with their **D**ifferential **M**icrowave **R**adiometer (DMR; Smoot et al. 1991; Bennett et al. 1996). COBE also measured the CMB temperature to be

$T_{\mathrm{CMB}} = 2.726 \pm 0.010$ K with the **F**ar-**I**nfra**R**ed **A**bsolute **S**pectrophotometer (FIRAS; Mather et al. 1994) instrument. The subsequent **B**alloon **O**bservations **O**f **M**illimetric **E**xtragalactic **R**adiation **an**d **G**eophysics (BOOMERanG; Netterfield et al. 2002) mission followed up on COBE's discovery, reporting the spatial flatness of the Universe by measuring the angular size of the first acoustic peak of the power spectrum. The **W**ilkinson **M**icrowave **a**nisotropy **P**robe (WMAP) was another major leap forward in CMB observations. This satellite mission mapped out the CMB across the entire sky at 13 arcminute resolution across 5 frequencies along with polarization sensitivity (Dunkley et al., 2009; Bennett et al., 2011, 2013). WMAP's observations allowed cosmologists to precisely measure cosmological parameters which predict the angular power spectra that were observed. The most recent satellite mission was *Planck*, which further cemented our understanding of cosmological models (Planck Collaboration et al., 2014a,b, 2016a,b, 2020b,a). *Planck*'s ΛCDM results are typically used in the standard cosmological models for contemporary cosmological analyses, especially in the context of early-Universe studies. There is a focus shift in contemporary CMB experiments, which are now designed to study the CMB at higher resolution and sensitivity with polarization. One exception is the search for a primordial $B$-mode polarization signal, which remains a major goal to validate cosmic inflationary theories.

Current experiments include the **A**tacama **C**osmology **T**elescope (ACT; Fowler et al. 2010; Dunkley et al. 2011). ACT is a ground-based telescope whose team recently published their sixth data release (Madhavacheril et al., 2023; Qu et al., 2023). The **S**outh **P**ole **T**elescope (SPT; Schaffer et al. 2011; Story et al. 2015; Henning et al. 2018; Chown et al. 2018) is another ground-based telescope which has also recently upgraded their telescope to their SPT-3G configuration, and is continuing observations with $\sim 16\,000$ detectors with polarization sensitivity (Sobrin et al., 2022). Both ACT and SPT have mapped the CMB at high resolution, furthering our understanding of galaxy clustering related CMB secondaries including the Sunyaev-Zel'dovich effect and weak gravitational lensing (both discussed later in this Section). The **C**osmology **L**arge **A**ngular **S**cale **S**urveyor (CLASS; Essinger-Hileman et al. 2014; Eimer et al. 2023) is a current ground-based survey covering 70% of the sky to constrain parameters sensitive to large scales like $\tau$ or $r$. These experiments pave the way for the next generation of high-resolution CMB observations. The **S**imons **O**bservatory (SO; Ade et al. 2019) is approaching first light. It is a ground-based telescope which will observe the CMB at six frequencies with a total of 60 000 detectors targeting cosmological parameters that are sensitive to both large and small angular scales. CMB-Stage IV (CMB-S4; Abazajian et al. 2016) is a proposed ground-based experiment which will further the ambitions of the Simons Observatory at even higher sensitivity. A report of the collaboration's science goals can be found here: Abazajian et al. (2019). Balloon-borne experiments are an effective way of avoiding atmospheric noise while being cheaper than a space satellite mission. SPIDER (SPIDER Collaboration et al., 2021) is a good example of such an experiment. The **Lite** (Light) satellite for the studies of **B**-mode polarization and **I**nflation from cosmic background **R**adiation **D**etection (LiteBIRD; LiteBIRD Collaboration et al. 2023) satellite is a Japanese-led mission which will observe the CMB at 15 frequency bands with polarization sensitivity to constrain inflationary models.

The natural structure of the CMB angular power spectrum in $T$, $E$, and $B$ decreases with increasing multipole $\ell$ (or decreasing angular scale of resolution), and the power spectrum of the noise from the detectors (after deconvolution of the beam) is expected to increase with $\ell$. The desired signal in the $B$ fields is also expected to be significantly weaker than those observed from the other two. It is for this reason that each subsequent generation of CMB experiments shares the recurring theme of improving detector sensitivity, and also increasing the number of detectors.

### 1.2.2 Primary Anisotropies

A majority of the signal in the cosmic microwave background originates from the physical structure of the Universe's contents at the time of recombination. The major components of the Universe in the early Universe were the photons (radiation), dark matter, and baryonic matter. One way to understand the structure of these constituents is to work forwards in time beginning from the Big Bang. Current models of cosmology rely on inflationary theories in the earliest moments of the Universe to seed a random distribution of over/under-densities of matter and radiation. It is worth noting here that ΛCDM does *not* include cosmic inflation in its model, and further discussion on inflation is provided in §1.4.3. Invoking an epoch of cosmic inflation can generate the necessary starting point for structure evolution by dragging quantum curvature fluctuations into cosmic scales (Baumann & McAllister, 2015), generating the initial density field which evolves into the structures we see today. The distribution of matter at position $\mathbf{x}$ and conformal time $\eta$ can be expressed through the density $\rho(\mathbf{x})$, or the density contrast:

$$\delta(\mathbf{x},\eta) = \frac{\rho(\mathbf{x},\eta) - \bar{\rho}(\eta)}{\bar{\rho}(\eta)}\,, \tag{1.13}$$

where $\bar{\rho}(\eta)$ is the mean matter density of the Universe at conformal time $\eta$. The density contrast quantifies the relative density of matter at a given position and time. The statistics of the distribution of matter are then contained in the matter power spectrum, defined through the Fourier transform of the density contrast (with $\mathbf{k}$ as the Fourier analogue of $\mathbf{x}$). The Fourier transform $f(\mathbf{k})$ of a general function $f(\mathbf{x})$ is defined as:

$$f(\mathbf{k}) = \int \mathrm{d}^3\mathbf{x} f(\mathbf{x}) e^{-i\mathbf{k}\cdot\mathbf{x}}\,, \tag{1.14}$$

with inverse Fourier transform:

$$f(\mathbf{x}) = \int \frac{\mathrm{d}^3\mathbf{k}}{(2\pi)^3} f(\mathbf{k}) e^{i\mathbf{k}\cdot\mathbf{x}}\,. \tag{1.15}$$

The matter power spectrum $P(k,\eta)$ is defined in Fourier space with respect to the density contrast $\delta(\mathbf{k},\eta)$:

$$\langle \delta(\mathbf{k},\eta)\delta^*(\mathbf{k}',\eta)\rangle = (2\pi)^3 P(k,\eta)\delta^3(\mathbf{k}-\mathbf{k}')\,. \tag{1.16}$$

The initial seeds of density (sometimes referred to as scalar) perturbations can be described by an initial power spectrum, which are predicted differently depending on the flavour of inflationary theory. Cosmic inflation is also predicted to produce gravitational wave (or tensor) perturbations, which we cover briefly later in this section. Simple models of inflation predict an initial scalar perturbation amplitude $A_s$, as well as spectral tilt $n_s$ which describe the height and shape of the initial power spectrum respectively (Equations (1.28)–(1.29) in Baumann & McAllister 2015). One can then work forward in time, taking into account all the interactions of matter fluctuations, to predict the spectrum of fluctuations that is expected at the time of recombination.

While the Universe was seeded all over with countless perturbations of various amplitudes and sizes, it is informative to first consider the evolution of a single overdense perturbation in an otherwise uniform field. Both dark matter and baryonic matter interact gravitationally. They move towards regions of overdensities (or away from underdensities). Dark matter responds only to gravity, so we expect the dark matter component of our perturbation to grow over time. Meanwhile, photons interact directly with the (ionized) baryons through Thomson scattering off of free electrons, but not with the dark matter. Protons scatter photons much less efficiently than electrons due to their mass, and their nature as composites; however, they are strongly coupled to the free electrons through Coulomb scattering. Another component of the energy density that one may consider are cosmic neutrinos. These would have interacted gravitationally with all matter, but they were also relativistic in the early Universe. This means that they simply free-stream away from the initial perturbation. In summary, the baryons and photons were 'coupled', but the dark matter and photons were not. Our overdense perturbation corresponds to a pressure gradient for the photons (and by extension the baryons). The outward pressure of photons acts against gravity and carries the photon/baryon fluid away from the initial overdensity at the sound speed. This continues until the Universe has eventually expanded and cooled enough for the baryons to no longer be ionized ($z \sim 1100$). By this time, the acoustic wave will have propagated out to a distance depending on the sound speed of the fluid $c_s$. This is usually referred to as the sound horizon, which is parameterized as follows:

$$r_{s,\mathrm{CMB}} = \int_0^{t_{\mathrm{CMB}}} \frac{\mathrm{d}t}{a(t)} c_s(t) = \int_{z_{\mathrm{CMB}}}^{\infty} \frac{\mathrm{d}z}{H(z)} c_s(z)\,, \tag{1.17}$$

where we have parameterized the integrals to describe the sound horizon at recombination, or the surface of last scattering. See Eisenstein (2020) for a visual animation of the single perturbation's evolution, which produces a density peak at the sound horizon.

In reality, the Universe was seeded with many initial over/underdensities after inflation. These initial perturbations were of quite small amplitude on an otherwise smooth background energy density. The cosmological information of relevance is contained in the collective evolution of all these perturbations (commonly referred to as baryon acoustic oscillations, or BAOs), resulting in a statistical correlation, or preference, for structures at the scale of the sound horizon. The presence of these acoustic oscillations in the density field at recombination is encoded within correlations between CMB photons. The structural information of photons (and baryons) at the last scattering surface is contained within the primary anisotropies of the CMB; hot spots indicate regions that were relatively underdense, and cold spots indicate regions that were relatively overdense at the surface of last scattering. This is explained by the Sachs-Wolfe effect, which postulates that photons from overdense regions are redshifted more than photons from underdense regions due to the relative differences in depth of each region's respective gravitational potential wells. A full-sky CMB temperature map can be typically transformed into a spherical harmonic representation, which is analogous to the Fourier transform of a flat, 2D image. The spherical harmonic transform of a field described on the unit sphere $X(\hat{\mathbf{n}}) = X(\theta, \phi)$ is defined through a set of coefficients $a^X_{lm}$ for predefined spherical harmonics functions $Y_{\ell m}(\theta, \phi)$ of multipole (or degree) $\ell$ and order $m$:

$$X(\theta, \phi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a^X_{\ell m} Y_{\ell m}(\theta, \phi)\,. \tag{1.18}$$

In spherical harmonics, the angular multipole can be thought of as analogous to a Fourier mode, and it corresponds to an inverse of angular size $\alpha = 2\pi/(\ell+1)$. The distribution of temperature contrasts as well as their angular sizes and separations is quantitatively described by the angular power spectrum, which satisfies the following:

$$\langle a^X_{\ell m} a^Y_{\ell' m'} \rangle = C^{XY}_\ell \delta_{\ell\ell'} \delta_{mm'}, \tag{1.19}$$

where $a^X_{\ell m}$ are the set of spherical harmonics coefficients for the field $X$ at multipole $\ell$ and order $m$. It is possible to take the cross spectrum of two different fields $X$ and $Y$. The angular power spectrum is a measure of the variance of the CMB field at a given angular scale $\alpha = 2\pi/(\ell+1)$. The CMB temperature power spectrum $C^{TT}_\ell$ is shown in Figure 1.2. The first peak at $\ell \sim 200$ corresponds to an angular scale of 1°, which is the angular size subtended by the BAO sound horizon at the surface of last scattering $r_{s,\mathrm{CMB}}$ (also referred to as the fundamental mode). Subsequent peaks of the CMB temperature power spectrum are overtones of the fundamental mode, and still deeply tied to the BAO physics. The amplitude of the primary CMB temperature power spectrum is quickly suppressed at smaller angular scales, as the diffusion damping of photons erases the temperature contrast information at smaller scales $\ell \gtrsim 2000$ (Silk, 1968). In short, the photon/baryon fluid was not perfectly coupled, and photons were able to diffuse from hot regions to colder regions through the random motions in between scattering events.

The primary CMB is also expected to be linearly polarized due to quadrupolar temperature structure at the surface of last scattering. Observed polarization information from the CMB is typically decomposed into $E$- and $B$-mode components, which is a basis that is independent of observer coordinates (Seljak & Zaldarriaga, 1996; Zaldarriaga & Seljak, 1997; Kamionkowski et al., 1997; Zaldarriaga, 2001). The $E$- and $B$-mode transformation is defined such that both fields are invariant to translations and rotations. The $E$-modes are additionally defined to be invariant to reflections, and the $B$-modes are defined to change sign with a reflection. They are named as such because the electric and magnetic fields in electromagnetism follow the same behaviour under those transformations. CMB polarization is observed as linearly polarized light, which is described by the Stokes parameters $Q$ and $U$. The decomposition from linear Stokes parameters $Q(\boldsymbol{\ell})$ and $U(\boldsymbol{\ell})$ in Fourier space (where $\theta$ is defined as the angle between a reference axis and $\hat{\ell}$) is defined as follows (Dodelson, 2003):

$$E(\boldsymbol{\ell}) \equiv Q(\boldsymbol{\ell})\cos(2\theta) + U(\boldsymbol{\ell})\sin(2\theta) \tag{1.20}$$

$$B(\boldsymbol{\ell}) \equiv -Q(\boldsymbol{\ell})\sin(2\theta) + U(\boldsymbol{\ell})\cos(2\theta)\,. \tag{1.21}$$

The primary $E$-mode polarization is expected to be sourced by the quadrupolar temperature structure which is present from the existing distribution of density fluctuations at recombination. The epoch of inflation which we invoked earlier to source the scalar perturbations is also expected to generate gravitational waves, which are tensor perturbations that are expected to generate primary $B$-mode signals. In the absence of tensor perturbations, the primary $BB$ power spectrum signal is expected to be exactly zero. The $B$-mode polarization signal from tensor perturbations is expected to be significantly weaker than either the scalar $T$ or $E$-mode signals (refer back to Figure 1.2 for a visual comparison), and a detection of the primary $B$-mode signal is desired to quantify the

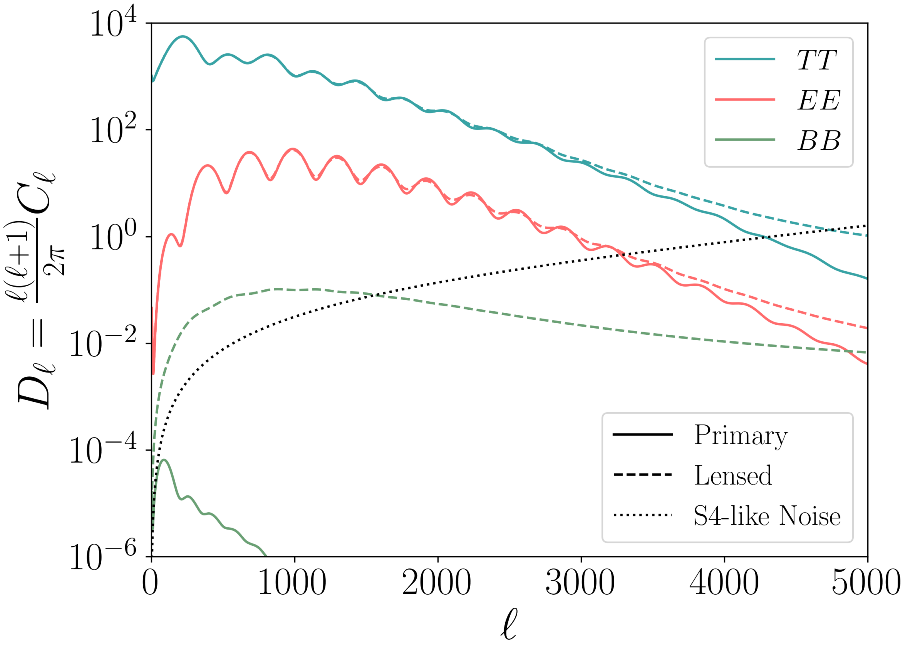

Figure 1.2: Theoretical CMB primary $TT$, $EE$, and $BB$ power spectra compared to their lensed counterparts. A tensor-to-scalar ratio $r = 10^{-3}$ is assumed. The noise spectrum $N_\ell^{TT}$ at proposed CMB-S4 levels is also shown.

ratio between the amplitudes of tensor and scalar perturbations. This parameter is the tensor-to-scalar ratio $r$, and different models of cosmic inflation predict their own value of $r$ (see Baumann & McAllister 2015 for a nice review of inflationary theories). The current upper limit on $r$ is $r \lesssim 0.06$ (Planck Collaboration et al., 2020b), and the next generation of CMB experiments are aiming to constrain $r \gtrsim 10^{-3}$ (Ade et al., 2019; Abazajian et al., 2019; LiteBIRD Collaboration et al., 2023).

### 1.2.3 Secondary Anisotropies

A majority of the features in the observed CMB at scales $\ell \lesssim 2000$ are considered to originate from the primary anisotropies discussed in the previous section, but there are also a number of secondary anisotropies that are expected to be present in the CMB. These are sourced by the large-scale structures (LSS) in between us and the surface of last scattering. One may take the perspective that secondary anisotropies contaminate the primary signal in the observed CMB, but we show in this Section that these secondaries impart valuable information about the distribution and evolution of the structures from which they are generated. This is especially the case as experiments reach higher sensitivities to probe the angular scales at which secondary anisotropies begin to dominate the observed signal. We also show that secondary anisotropies maybe characterized and cleaned from the observed CMB signal in many cases. Some examples of secondary anisotropies include the thermal and kinetic Sunyaev-Zel'dovich effects (tSZ and kSZ), the integrated Sachs-Wolfe effect, weak gravitational lensing, and others. In a similar sense to the secondary anisotropies, foreground signals from dust emission may also contaminate the observed CMB. We reserve a more thorough review of the weak gravitational lensing of the CMB for § 1.2.4, as it is the direct focus of Chapters 5-6 of this dissertation.

A notable secondary feature of the CMB is the Thomson scattering of CMB photons *after* recombination. The Universe remains neutral until the epoch of reionization, so this effect is typically parameterized as the optical depth to reionization $\tau$. The optical depth serves to generally suppress the observed CMB temperature fluctuations by a factor of $e^{-\tau}$, and a measurement of $\tau$ serves to constrain the time at which reionization occurred. Keating & Miller (2006) and Reichardt (2016) provide good reviews with further discussion of the CMB optical depth.

As photons stream away from the surface of last scattering, they may encounter massive structures that form over cosmic time. One physical interaction that may occur at late times is the inverse Compton scattering off of free (hot) electrons that surround galaxies and galaxy clusters. This is the Sunyaev-Zel'dovich effect, and there are a few different versions. The thermal Sunyaev-Zel'dovich effect (tSZ; Zeldovich & Sunyaev 1969; Sunyaev & Zeldovich 1970, 1972) describes how thermal electrons can impart a general boost to the photons' energies relative to the CMB blackbody as they pass through. This phenomenon is also produced by the newly freed electrons at the epoch of reionization as well as the hot intergalactic medium (IGM) at late cosmic times. The tSZ's spectral distortion effect can be characterized along the lines of sight of galaxy clusters in a CMB experiment if it is sensitive to multiple frequencies $\nu$, and it is typically characterized through maps of the Compton-$y$ parameter (Zeldovich & Sunyaev, 1969; Sunyaev & Zeldovich, 1970). The observed temperature of the CMB is affected through the tSZ effect as follows:

$$\Delta T_{\rm tSZ}(\nu) = y T_{\rm CMB}\Big(\frac{h\nu}{k_B T_{\rm CMB}}\frac{e^{h\nu/k_B T_{\rm CMB}}+1}{e^{h\nu/k_b T_{\rm CMB}}-1} - 4\Big). \tag{1.22}$$

The Compton-$y$ parameter is interpreted as a measure of the total integrated electron gas pressure along the line of sight. Compton-$y$ maps can be created from *Planck* data combined with ACT (Madhavacheril et al., 2020a) as well as SPT (Bleem et al., 2022), and they can be combined with observations of galaxy cluster density to probe the ionized content of the late Universe (Yan et al., 2021; Lokken et al., 2022).

The kinetic Sunyaev-Zel'dovich effect (kSZ; Sunyaev & Zeldovich 1980) is a Doppler shift in the energies of scattered photons due to the bulk motion of free electrons surrounding galaxy clusters and at the epoch of reionization. Consider the example of free electrons around a galaxy cluster. The kSZ effect is sourced by the peculiar velocities of the electrons, which move according to the gravitational potential of each cluster, as well as the general gravitational potential of the local large scale structure. As a Doppler shift, the relevant components of these motions are the radial velocities, and we expect the kSZ effect to manifest as small-angular scale anisotropies of the observed CMB temperature field such that $\Delta T/T_{CMB} = -\tau v_r/c$, where $\tau$ is the optical depth of the cluster. The kSZ is also expected to be sourced by the free electrons at the epoch of reionization, and Smith & Ferraro (2017) present an estimator to detect the reionization signal at small angular scales of the CMB. Such a detection is possible with upcoming experiments such as CMB-S4, and can help to constrain the time and duration of reionization (Alvarez et al., 2021). Alternatively, Foreman et al. (2023) show that it is possible to clean kSZ signals to better observe other CMB features at small-scales.

The Universe at late times ($z \to 0$) begins to become dominated by dark energy, which causes the expansion of the Universe to accelerate. This causes the gravitational potentials of the largest structures to decay as the expansion pushes them apart. A CMB photon passing through such structures at late-times experiences an initial gravitational blueshift as it travels into the gravitational potential, and a gravitational redshift as it travels out of the potential. The gravitational potential slowly decreases as the photon passes through, so it experiences a slightly stronger blueshift than redshift. This is often referred to as the integrated Sachs–Wolfe effect (see Nishizawa (2014) for a review), and it can be detected by combining observations of the CMB with maps of large-scale galaxy clustering data to place constraints on models of dark energy (Fosalba et al., 2003; Ho et al., 2008; Giannantonio et al., 2008; Ferraro et al., 2015; Dong et al., 2021).

Secondary anisotropies in the CMB can also be present in the polarization signal. If the existing linear polarization signal in the primary CMB experiences a net rotation, then the effect may be observed through correlations between the observed $E$- and $B$-mode fields of the CMB (i.e., the $EB$ cross spectrum). This phenomenon is referred to as cosmic birefringence (refer to Komatsu 2022 for a review). Cosmic birefringence is predicted to be sourced by some models of dark matter and dark energy such as the axion or ultra-light axion (Marsh, 2016; Ferreira, 2021), and a $\sim 3\sigma$ hint of its existence has been shown in WMAP and *Planck* data (Diego-Palazuelos et al., 2022; Eskilt & Komatsu, 2022).

### 1.2.4 The Lensed Cosmic Microwave Background

Gravitational lensing of the CMB by cosmological structures along the line of sight has become a standard observational tool to probe the content and evolution of the universe (see Lewis & Challinor (2006) for an extensive review). The CMB lensing deflections are sourced by the 3D gravitational potentials $\Phi(\mathbf{r}, \eta)$ of massive LSS. The gravitational potentials are defined at position $\mathbf{r} = \chi\hat{\mathbf{n}}$ that is at some comoving distance $\chi$ away from us in the direction $\hat{\mathbf{n}}$. The growth of structure over time causes the potentials to evolve, so the 3D gravitational potential is also parameterized by the conformal time $\eta$. The lensing potential $\phi(\hat{\mathbf{n}})$ at line of sight $\hat{\mathbf{n}}$ is typically defined such that the total deflection experienced by a CMB photon is the gradient of the lensing potential perpendicular to the line of sight (described with the operator $\boldsymbol{\nabla}_{\hat{\mathbf{n}}}$). The lensing potential is the projection of the 3D gravitational potential (which does the lensing) along the line of sight $\hat{\mathbf{n}}$ between the observer (us at comoving distance $\chi = 0$) and the surface of last scattering (at comoving distance $\chi = \chi_{\mathrm{CMB}}$) in a flat Universe:

$$\phi(\hat{\mathbf{n}}) \equiv -2 \int_0^{\chi_{\mathrm{CMB}}} \mathrm{d}\chi \Phi(\chi\hat{\mathbf{n}}, \eta_0 - \chi) W_\phi(\chi)\,, \tag{1.23}$$

where $\eta_0$ is the conformal time today, and $\eta_0 - \chi$ is then the conformal time at which the photon was located at $\chi\hat{\mathbf{n}}$. Note that there is an extra convolution with a lensing kernel $W_\phi(\chi)$ which controls the strength of the lensing as a function of the source potential $\Phi$'s distance away from us. Most of the lensing deflections are 'weak' lensing, meaning that the deflections are of order $\sim$ arcmin (refer to §1 of Lewis & Challinor 2006 for the reasoning), and massive structures at $z \approx 2$ are the most efficient lenses for the CMB due to the geometry of the source, lens, observer system. As defined earlier, the total deflection experienced by a CMB photon is then:

$$\boldsymbol{\alpha} = \boldsymbol{\nabla}_{\hat{\mathbf{n}}} \phi(\hat{\mathbf{n}}) = -2 \int_0^{\chi_{\mathrm{CMB}}} \mathrm{d}\chi \boldsymbol{\nabla}_{\hat{\mathbf{n}}} \Phi(\chi\hat{\mathbf{n}}, \eta_0 - \chi) W_\phi(\chi)\,. \tag{1.24}$$

The lensing convergence is also typically defined such that $\kappa(\hat{\mathbf{n}}) = -\boldsymbol{\nabla}_{\hat{\mathbf{n}}} \cdot \boldsymbol{\nabla}_{\hat{\mathbf{n}}} \phi(\hat{\mathbf{n}})/2$. The lensing potential $\phi$, deflection $\boldsymbol{\alpha}$, and convergence $\kappa$ are different physical quantities, but they all contain the same information relevant to the lensing physics: the distribution and nature of the gravitational potentials $\Phi$ of massive LSS. The lensing deflection power spectrum and lensing convergence power spectrum are typically defined in terms of the lensing potential power spectrum as $C_L^{\alpha\alpha} = L^2(L+1)^2 C_L^{\phi\phi}/2\pi$, and $C_L^{\kappa\kappa} = L^2(L+1)^2 C_L^{\phi\phi}/4$. Lewis & Challinor (2006) show in their Equations (3.4)–(3.11), and subsequently (3.14) that the angular power spectrum of the lensing potential $C_L^{\phi\phi}$ is directly tied to the physical power spectrum of the gravitational potentials $P_\Phi(k,\eta)$. This is directly related to the power spectrum of matter density perturbations $P(k,\eta)$ (Equation (3.15) of Lewis & Challinor 2006):

$$P_\Phi(k,\eta) = \frac{9\Omega_m^2(\eta)H^4(\eta)}{8\pi^2}\frac{P(k,\eta)}{k}\,. \tag{1.25}$$

Gravitational lensing imparts several effects on the observed CMB. At the map-level, the observed CMB temperature $\tilde{T}$ is simply shifted by the deflection angle:

$$\tilde{T}(\hat{\mathbf{n}}) = T(\hat{\mathbf{n}} + \boldsymbol{\alpha}) = T(\hat{\mathbf{n}}) + \boldsymbol{\nabla}_{\hat{\mathbf{n}}} T(\hat{\mathbf{n}}) \cdot \boldsymbol{\nabla}_{\hat{\mathbf{n}}}\phi(\hat{\mathbf{n}}) + \mathcal{O}(\phi^2)\,. \tag{1.26}$$

This has the general effect of magnifying and demagnifying the features of the observed CMB temperature field. In power spectra, this manifests as a transfer of power between angular multipoles. The CMB temperature field is equally likely to be magnified or demagnified, while preserving the total observed power (i.e, the integrated $TT$ power spectrum is conserved). This has the qualitative effect of smoothing out the observed CMB $TT$ power spectral features; the primary $TT$ peaks are suppressed, and the troughs are raised up. The damping tail at $\ell \gg 3000$ also receives power from the larger-scales, so we see a general increase in lensed $TT$ power to the point that it is the dominant signal at small angular scales. Both of these phenomena are featured in Figure 1.3. In polarization, the effects are similar, but one must remember that we observe the linear polarization, and transform it into $E$- and $B$-modes. The redistribution of linear polarization in map-space translates into a mixture of power between the observed $E$- and $B$-modes in combination with the transfer of power between multipoles described earlier. The $EE$ power spectrum experiences very similar qualitative changes as the $TT$ power spectrum, as shown in Figure 1.2. This can be explained with the same reasoning as the peak smearing phenomenon; there is significantly more primary $EE$ power than primary $BB$ power (if any), so it is more likely to leak $E$-modes into $B$-modes than the other way around. This is also seen in Figure 1.2, where the lensed $BB$ power is significantly higher in amplitude than its primary counterpart due to the added power from the original $E$ field. As a final note on the direct effects of lensing on the observed CMB field(s), we see in Equation (1.26) that the extra lensing contributions depend directly on the gradient (and higher order derivatives) of the original CMB field itself. We show later that this is taken advantage of in estimators for the lensing potential field and its statistics.

Estimates of the lensing potential angular power spectra directly inform us of the statistics of matter distribution in the Universe, which makes CMB lensing a powerful cosmological probe. Improvements in measurements of the CMB with telescopes like the Atacama Cosmology Telescope (ACT, Aiola et al. 2020), the South Pole Telescope (SPT, Henning et al. 2018) and the *Planck* satellite (Planck Collaboration et al., 2020b) have unveiled the fluctuations in the temperature and polarization signal of this primordial light down to arcminute scales.

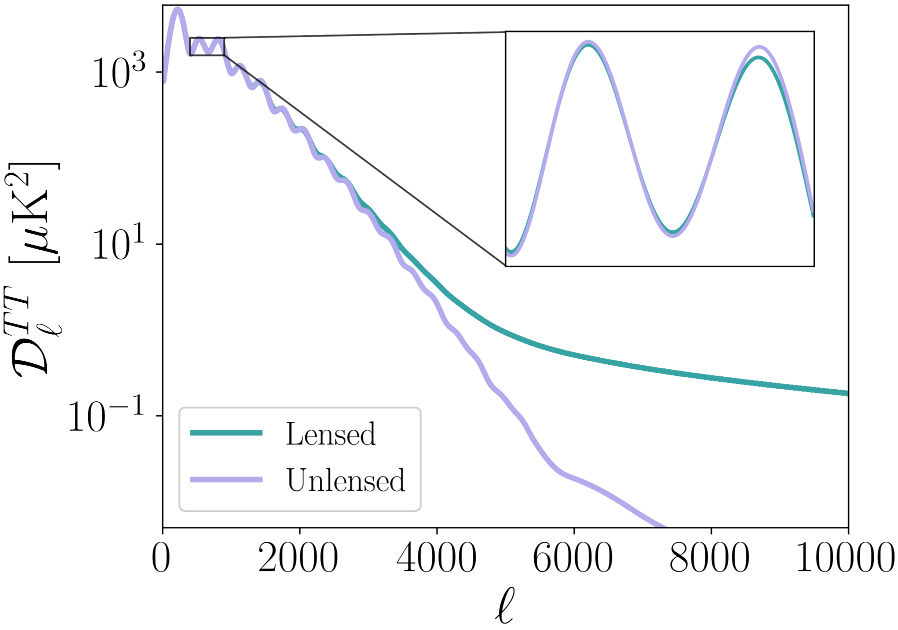

Figure 1.3: A comparison of the CMB $TT$ power spectrum before and after lensing. The peak smoothing can be observed in the inset zoomed in at $400 < \ell < 900$, and the excess power at $\ell \gg 3000$ is apparent.

Extracting (or 'reconstructing') the lensing signal from CMB temperature $T$ and/or polarization $E$ and $B$ maps can proceed via a number of approaches. The pioneering works of Hu & Okamoto (2002) and Okamoto & Hu (2003) developed the concept of the quadratic estimator (QE) for the lensing signal, which combines pairs of observed maps $TT$, $EE$, $BB$, $TE$, $TB$, and $EB$, to 'reconstruct' the lensing potential field. This is made possible by taking advantage of the correlations, or covariances, induced by the redistribution of power across multipoles and/or polarization. The Hu & Okamoto (2002); Hu et al. (2007) quadratic estimator for the deflection field $\alpha(\mathbf{L})$ constructed with the quadratic pair $XY \in \{TT, EE, BB, TE, TB, EB\}$ has the following form:

$$\alpha_{XY}(\mathbf{L}) = \frac{A_{XY}(L)}{L} \int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} X(\mathbf{l}_1) Y(\mathbf{L}-\mathbf{l}_1) F_{XY}(\mathbf{l}_1, \mathbf{L}-\mathbf{l}_1)\,, \tag{1.27}$$

where $A_{XY}(L)$ is a normalization factor dependent on the choice of filtering functions $F_{XY}$ (Hu & Okamoto, 2002, Equations (13-15)) or (Hu et al., 2007, Equations (19-22)) and $f_{XY}$ (Hu & Okamoto, 2002, Table 1) or (Hu et al., 2007, Equation (17)):

$$A_{XY}(L) = L^2 \left[ \int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} f_{XY}(\mathbf{l}_1, \mathbf{L}-\mathbf{l}_1) F_{XY}(\mathbf{l}_1, \mathbf{L}-\mathbf{l}_1) \right]^{-1}. \tag{1.28}$$

The filtering functions presented in Hu et al. (2007) differ from those in Hu & Okamoto (2002) through a slightly different treatment of the quadratic pair. In this case, the $X$ field is filtered to optimize for information from CMB gradient contribution in Equation (1.26), and the $Y$ field is filtered to optimize for information from the lensing contribution. The gradient filter in particular is constructed such that the $X$ field receives no contributions from small-scales $\ell \gtrsim 2000$ to mitigate

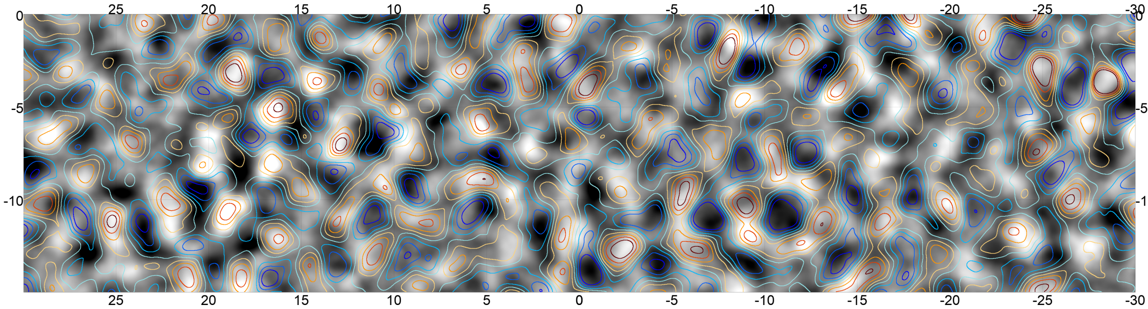

Figure 1.4: Part of the reconstructed lensing potential map from Madhavacheril et al. (2023, Figure 4, included with permission from the authors). Overlaid contours are from the *Planck* observations of the CIB. The reconstructed lensing potential visibly correlates well with the CIB, which can be used as a tracer for the massive Large Scale Structures that source the lensing potential.

bias (Hu et al., 2007). Note that this now includes the possibility of a $XY = ET$ estimator ($BT$ and $BE$ are neglected due to the low amplitude of the primary $B$ field). The filtering in Hu et al. (2007) is in contrast to the original Hu & Okamoto (2002) estimator which employs filters optimized to minimize variance in the reconstructed deflection field.

The iterative $EB$ estimator (Hirata & Seljak, 2003a; Smith et al., 2012) is expected to be especially effective at large angular scales ($L \lesssim 1000$) due to the transfer of power from $E$ modes into $B$ modes, the latter of which contains only a meager signal in the primary CMB caused by a possible epoch of cosmic inflation. Application of the QE on *Planck* data has allowed for a $40\sigma$ detection of gravitational lensing (Aghanim et al., 2020), and there has also been recent success from ACT (Qu et al., 2023; Madhavacheril et al., 2023). The reconstructed lensing potential map from Madhavacheril et al. (2023) is shown in Figure 1.4 (their Figure 4), which visually matches the structure of the observed cosmic infrared background (CIB) observed by *Planck* (used as a reasonable proxy for the distribution of galaxies sourcing the lensing field). However, the effectiveness of the QE may soon be limited as experiments push to smaller scales and lower noise. Indeed, the recent analysis of a set of very deep SPTpol data showed improved results compared with the more standard QE approach (Millea et al., 2021). The QE formalism, which only crudely approximates the full maximum likelihood estimate of the signal, is statistically suboptimal on small angular scales and in low-noise regimes (Hirata & Seljak, 2003b,a; Smith et al., 2012; Carron & Lewis, 2017; Horowitz et al., 2019; Hadzhiyska et al., 2019).

The derivation of the QE procedure relies on the assumption that lensing is a weak effect, in the sense that it has only a small effect on the statistics of the CMB sky. This is an appropriate approximation for most of the regimes in which the QE has historically been applied, but at lower noise levels, this approximation quickly begins to break down. On small angular scales, at $\ell \gg 2000$, the power resulting from gravitational lensing dominates over the primordial unlensed power spectrum, as can be seen in Figure 1.2,1.3. In this lensing-dominated regime, QE techniques are suboptimal in the limit of low noise levels. Figure 1.5 shows that the additive bias $N_L^{(0)}$ from the $TT$ reconstruction noise becomes lower than that of the iterative $EB$ estimator at small angular scales ($L \gtrsim 2000$) due to increasing noise from the $EB$ estimator. The zero-th order reconstruction noise, $N_L^{(0)}$, stems from the inherent variance in the primary CMB that is propagated through the reconstruction of the lensing potential. The reconstruction information is estimated from the

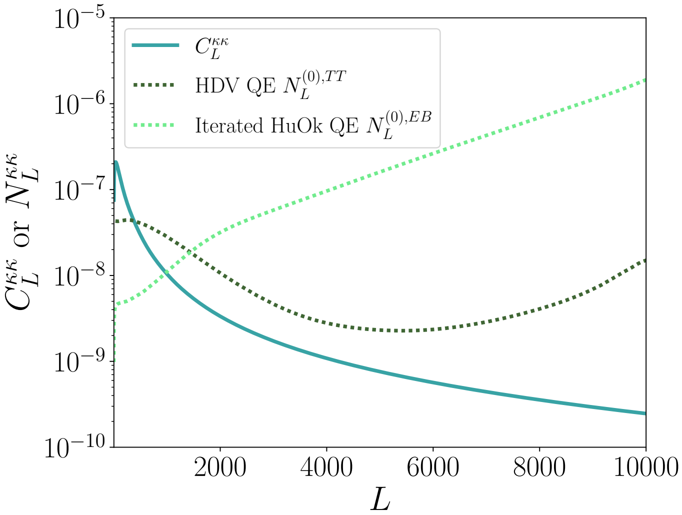

Figure 1.5: A comparison of the CMB lensing convergence power spectrum with the additive bias from reconstruction noise experienced by traditional quadratic estimator techniques with CMB-S4-like noise levels. The $EB$ estimator's effectiveness due to its low reconstruction noise at large scales ($L < 1000$) is apparent, but the reconstruction noise quickly grows at small-scales. At small-scales, the $TT$ reconstruction techniques will be more effective.

observed realization of the CMB, and the statistical limit to the amount of information one can obtain from a single realization sets the uncertainty. The noise power is greater than the signal power at small-scales; however, this bias can be removed with well-established methods (Dvorkin & Smith, 2009; Hanson et al., 2011). Bias from higher-order reconstruction noise $N_L^{(1)}$ also need to be estimated and removed at small angular scales; these stem from additional permutations of the connected four-point function as discussed in Kesden et al. (2003). Recent advances in computation and statistical methodology allow for the computation of the full maximum likelihood lensing map (Hirata & Seljak, 2003b; Carron & Lewis, 2017; Millea et al., 2020) thereby surpassing the performance of QE techniques in simulated data.

Quadratic estimators can also be shown to be suboptimal in the small-scale, low-noise limit because they are weighted by the sky-averaged variance of the large-scale modes, despite precise measurement of the large-scale modes. Due to cosmic variance on large scales, this weighting contributes to excess variance in the lensing reconstruction. This limitation can be circumvented with the so-called 'Gradient Inversion' approach to reconstruction (Seljak & Zaldarriaga, 2000; Horowitz et al., 2019; Hadzhiyska et al., 2019) which, unlike the QE technique, is not limited by cosmic variance exhibited by the large scale temperature fluctuations.

Quadratic estimators and other estimators like the gradient inversion estimator aim to reconstruct a map of the underlying lensing potential *explicitly*. A reconstructed map of the lensing field is valuable for delensing (Kesden et al., 2002; Knox & Song, 2002; Seljak & Hirata, 2004; Green et al., 2017; Hotinli et al., 2022) and for cross-correlation with other maps of large scale structure (e.g. (Sherwin et al., 2012; Bianchini et al., 2015; Liu & Hill, 2015; Schmittfull & Seljak, 2018; Robertson et al., 2021; Darwish et al., 2021b; Baxter et al., 2022; Chang et al., 2022; Lin et al., 2020; Piccirilli et al., 2022)), but the lensing power spectrum carries valuable information even without an associated map-level reconstruction. Phenomena which impact matter clustering can be constrained using measurements of the matter power spectrum, without requiring a map of overdensities. Examples include the effects of neutrino mass and related quantities (Kaplinghat et al., 2003; Lesgourgues & Pastor, 2012; Green & Meyers, 2021; Abazajian et al., 2022), dark matter interactions (Tulin & Yu, 2018; Gluscevic et al., 2019; Buen-Abad et al., 2022), ultralight dark matter (Hui et al., 2017; Ferreira, 2021), warm dark matter (Drewes et al., 2017), and baryonic feedback (Chisari et al., 2019).

Chapter 5 of this dissertation covers the development of the Small Correlated Against Large Estimator (SCALE), a novel estimator for the CMB lensing power spectrum at small angular scales ($\ell \gg 3000$). SCALE takes advantage of the correlations between the CMB temperature modes between large- and small-scales. We continue in Chapter 6 with an application of SCALE in the parameter estimation of cosmological models.

### 1.2.5 Observing the Cosmic Microwave Background

The process of designing, constructing, and operating a CMB experiment is increasingly long and complex. Upcoming experiments are being proposed and designed over 10 years before they are expected to operate (Abazajian et al., 2016, 2019; LiteBIRD Collaboration et al., 2023). This includes a wide range of experiments from ground-based, balloon-borne, and satellite missions. Although there are many experiments designed and optimized for different science cases, they have and continue to share similarities in their overall observation procedure.

Current and upcoming experiments employ the use of either superconducting transition edge sensors (TES; Suzuki et al. 2020) or microwave kinetic inductance detectors (MKIDs; Ulbricht et al. 2021; Baselmans et al. 2022). The detectors are designed to be sensitive to specific frequencies as well as linear polarization, and are typically arranged in large arrays. Two important properties of the detectors to consider are their noise equivalent temperature ($[\mathrm{NET}] = \mu\mathrm{K}/\sqrt{\mathrm{Hz}}$, also sometimes called sensitivity), and their beam shape/size (Knox, 1995). The noise equivalent temperature of a detector $\mathrm{NET}_{\mathrm{det}}$ describes the equivalent white noise that it experiences during observations over time, and it translates into a flat line in a power spectral decomposition: $P(f) = \mathrm{NET}^2$ (in the absence of detector motions). While this is an intrinsic property of the detector itself, an experiment is typically characterized by the noise equivalent temperature of the entire array: $\mathrm{NET}_{\mathrm{arr}} = \mathrm{NET}_{\mathrm{det}}/\sqrt{N_{\mathrm{det}}}$, where $N_{\mathrm{det}}$ is the total number of detectors in the array. The beam of a detector refers to the angular shape and size of its sensitivity region when pointed at a source, and it is analogous to the point spread function of an optimal instrument. While CMB detector beams can have non-trivial shapes and sidelobes, they are typically modelled by a 2-dimensional Gaussian with size $\theta_{\mathrm{FWHM}}$ such that their observed solid angle is approximately $\Omega_{\mathrm{det}} = \theta^2_{\mathrm{FWHM}}$.

An array of detectors is assembled onto a ‘focal plane’, and there is freedom in choosing the physical layout of the detectors and where they are pointing. Configurations are typically chosen such to maximize the number of detectors on the focal plane while minimizing the amount of overlap between the detector beams. The latter is done to avoid any leakage of signals between detectors, and is often referred to as optical crosstalk. There is also freedom to either focus on a few specific observing frequencies or to diversify into a wide range of frequency coverage. Choosing fewer observing frequencies allows for more detectors on an array to be dedicated to each frequency, lowering the overall noise equivalent temperature for the array at each given frequency. Choosing to cover more observing frequencies means taking a hit in the array’s overall noise equivalent temperature for each frequency, but can be informative for characterizing spectral features in observations. This is particularly useful for characterizing the foreground contamination in the CMB signal (van Engelen et al., 2014; Osborne et al., 2014; Flauger et al., 2014; Madhavacheril & Hill, 2018; Beck et al., 2020; Abylkairov et al., 2021; Darwish et al., 2021a; LiteBIRD Collaboration et al., 2023).

After a focal plane has been established, a survey area must be chosen for the survey. The overall noise levels of CMB observations are dependent on the total observation time: $\sigma = \mathrm{NET}/\sqrt{t_{\mathrm{pix}}}$, where $t_{\mathrm{pix}}$ is the total amount of time spent observing each pixel in the map. It is useful for later analysis to define a noise level such that $w \equiv (\sigma^2\Omega)^{-1}$. One may choose to focus observations on a small region of the sky, and frequently repeat observations to beat down noise. This would optimize for a cleaner picture of the observing region, but angular power spectrum analysis would be limited to multipoles higher than the angular size of the sky area. One may alternatively choose to survey a larger area of the sky, resulting in overall noisier maps. A positive trade-off for these maps is that they would then contain both information at lower multipoles, as well as more sky area to limit the effects of cosmic variance at higher multipoles. There are also several other considerations such as the galactic foregrounds that are present in the sky, and the presence of other astrophysical sources that may contaminate the CMB signal. Experiments may choose to avoid certain parts of the sky that are known to be heavily contaminated. The choice of survey area is typically a compromise between these considerations, and is often a major factor in the design of a CMB experiment.

The focal plane of an experiment does not cover the entirety of its survey area, so a scanning strategy must be chosen to map out the entire area. This is usually done in ground-based experiments by sweeping across the survey area, and slowly filling up the map in the perpendicular direction with this motion. This scanning motion is typically repeated iteratively for the duration of observations. A satellite experiment has access to the entire sky, and typically rotates along an axis to sweep out rings of the sky. The satellite's rotation axis is also made to precess, and the combination of the satellite's rotation, precession, and its orbit around the Earth/Sun allows it to eventually map out the full sky. It is important to consider here that the sweeping motions of the focal plane are expected to induce extra components of noise into the observations, regardless of the scanning strategy. It was mentioned previously that the detectors' intrinsic noise properties are captured by their noise equivalent temperature, but the repeated motions of the focal plane can cause the (originally) white noise to become correlated such that the power spectral decomposition of the detector timestreams, or time ordered data (TODs), is no longer flat. This is normally referred to as $1/f$ noise, and the power spectral decompositions (PSD) of a noise TOD can be modelled as follows:

$$P(f) = \mathrm{NET}^2_{\mathrm{det}}\left(\frac{f^\alpha + f^\alpha_{\mathrm{knee}}}{f^\alpha + f^\alpha_{\mathrm{min}}}\right), \tag{1.29}$$

where $\alpha$ is a spectral index that is typically $-1$ (hence $1/f$), $f_{\mathrm{min}}$ is the minimum frequency considered, and $f_{\mathrm{knee}}$ is a frequency at which the modelled PSD typically has a 'knee' in its shape. This shape is dependent on the choice of on the scanning strategy of the experiment. In the case of a satellite mission, it is dependent on the rotation and precession speed of the satellite.

As CMB experiment scans its survey area, groups of detectors in the focal plane are continuously connected to a set of readout electronics. Consider a group of detectors connected to a single wire that leads to a readout system. Each detector is typically assigned a unique frequency, and the readout system is designed to monitor the voltage of the circuit relative to a predetermined voltage bias for all readout frequencies at all times. Fluctuations in the voltage at a particular detector's frequency over time are registered as that detector's 'signal', or TOD. This is known as a frequency multiplexing (fMUX) readout system (Dobbs et al., 2009; de Haan et al., 2020). It is known that detectors sharing the same readout electronics can potentially leak signals into each other's TODs at this stage, and it is typically referred to as electrical crosstalk. Electrical crosstalk is expected to be strongest for pairs of detectors that are close together in the frequency schedule (i.e., their assigned readout frequency), and it is an effect that is expected to be at the level of less than 1% (Montgomery et al., 2022). Previous and current experiments have estimated their levels of crosstalk through calibration with cosmic rays, and found that their calibrations were sufficient to remove most effects of crosstalk from their data (Planck Collaboration et al., 2016c; Ade et al., 2022).

A set of detector TODs, $\mathbf{d}(t)$, must first be transformed into CMB map(s), $\mathbf{m}$, for use in cosmological analyses. A pointing matrix $\mathbf{P}$ can be constructed containing information about the direction of the sky at which each detector is pointing at a given time $t$. This can be used to transform the TODs into a set of sky maps using a mapmaking equation. If a set of detector TODs can be expressed as follows:

$$\mathbf{d} = \mathbf{P}\mathbf{m} + \mathbf{n}, \tag{1.30}$$

where the first term is the signal from the sky, and the second term is noise. The noise $\mathbf{n}$ can be further decomposed into correlated and non-correlated components, where the correlated component is due to the $1/f$ noise described in Equation (1.29). A simple mapmaking equation can be obtained by solving for $\mathbf{m}$ while minimizing its variance:

$$\mathbf{m} = (\mathbf{P}^T\mathbf{N}^{-1}\mathbf{P})^{-1}\mathbf{P}^T\mathbf{N}^{-1}\mathbf{d}, \tag{1.31}$$

where $\mathbf{N}$ is the covariance of all noise TODs $\mathbf{n}$. This includes the correlations between noise TODs from the effects of $1/f$ noise described in Equation (1.29), and the mapmaking equation is a simple, but effective way to transform observed CMB data from TOD space into map space (Keihänen et al., 2005; Kurki-Suonio et al., 2009; Keihänen et al., 2010).

A given experiment may produce a set of observed maps $\mathbf{m}$ in terms of Stokes parameters $I$ (which is readily transformed into relative temperature), as well as $Q$ and $U$ (which are linear polarization components). The observed linear polarization is dependent on our frame of reference, so it is useful for cosmological analyses to transform the polarization data into a basis that is independent of reference frame. This is typically done by decomposing the CMB polarization data into curl-free $E$-modes, and divergence-free $B$-modes (Zaldarriaga, 2001). CMB maps are typically projected onto **H**ierarchical **E**qual **A**rea iso**L**atitude **Pix**elation (`HEALPix`; Górski et al. 2005) projections of the sphere, which recursively divide the sky into 'rhombus'-like pixels with equal solid angle.

CMB maps provide a depiction of the experiment's data that is conceptually easy to understand, but it is not a form that is readily applicable for cosmological analyses. The Gaussian statistics of the maps can be determined through their angular power spectra, which can be computed through spherical harmonic decompositions, see Equation (1.19). The information in the power spectrum is considered Gaussian because it captures the variance between any two points on the field. In other words, it is a two-point function. Higher order statistics such as the skewness or kurtosis of the observed field can be quantified through the angular bispectrum (or a three-point function comparing any three points of the field at once) or the angular trispectrum (or a four-point function comparing any four points of the field at once) respectively. It is also possible at this stage to use information obtained from the different observing frequencies to characterize the foreground signals present in the CMB map, which do not follow the same spectral patterns as the CMB blackbody, and 'clean' them from the maps (Tegmark et al., 2003; Madhavacheril & Hill, 2018; Beck et al., 2020; Abylkairov et al., 2021; Darwish et al., 2021a). Assuming that the correlated noise contributions to the detector TODs was cleaned properly, then there should remain residual white noise in the CMB maps, which should manifest as a combination of white noise in the angular power spectrum. In terms of temperature power, this manifests as

$$N_\ell^{TT} = w^{-1} \exp\left(\ell(\ell+1)\sigma_b^2\right), \tag{1.32}$$

where $\sigma_b = \theta_{\mathrm{FWHM}}/2\sqrt{2\ln 2}$ is the width of the instrumental beam, and $w$ is the same noise level defined earlier (Knox, 1995). Typical CMB experiment noise levels are reported as $w^{-1/2}$ in units of $\mu$K-arcmin. The effective noise power in polarization power is simply twice that of the temperature noise power, as the intensity information is approximately split equally into the two linear polarization pieces.

The power spectra of the observed CMB maps are typically compared to theoretical models of the CMB, which are dependent on the cosmological parameters of interest. Parameter estimation is typically performed through a Markov Chain Monte Carlo (MCMC) analysis, which is a Bayesian statistical technique that is further described in § 1.5. To do this, a full set of covariances between the observed CMB band-powers ought to be determined. The primary CMB is expected to be a Gaussian random field, meaning that CMB power between disjoint multipoles is not expected to correlate. In practice, small correlations are expected due to the effects of foregrounds and gravitational lensing, among others. In the limit of these small correlations, the uncertainty of estimated band-powers is well described by

$$\Delta C_\ell = \sqrt{\frac{2}{(2\ell+1)\Delta\ell f_{\mathrm{sky}}}}(C_\ell + N_\ell), \tag{1.33}$$

where $\Delta\ell$ is the multipole width of each band-power, and $f_{\mathrm{sky}}$ is the fraction of the sky that is observed (Knox, 1995).

There exist many software packages to facilitate and/or simulate the entire process outlined in this section. The latter kind in particular are useful to predict the performance of upcoming experiments, which is informative for their designs. The **T**ime **O**rdered **A**strophysics **S**calable **T**ools (`TOAST`; Kisner et al. 2023)[1], which is a software based in `C++` with a `Python` wrapper built to simulate the CMB observation process. It takes in parameters describing an experiment's focal plane configuration, detector properties, and scanning strategy, and it creates 'sky signal' TODs based on the particular pixels of an input map that each detector is pointing at. It can also generate the noise TODs expected from Equation (1.29). The `TOAST` framework is designed to be scalable such that simulations can be modified to optimize for speed or memory depending on the machine running it. With a set of detector TODs, a common mapmaking software is called **MA**pmaking through **D**estriping for **A**nisotropy **M**easurements (`MADAM`; Keihänen et al. 2005, 2010)[2]. There exists an option within `TOAST` to automatically feed the detector TODs into `MADAM` to output

---

[1] https://github.com/hpc4cmb/toast

[2] https://github.com/hpc4cmb/libmadam

simulated observed CMB maps. There also exist a set of software called Boltzmann codes that predict theoretical CMB power spectra given a set of cosmological parameters. These work by seeding a set of initial conditions for the perturbations predicted by a given set of cosmological parameters, usually within a standard model like ΛCDM, and evolve them into observables in the CMB and LSS. Examples include the **C**ode for **A**nisotropies in the **M**icrowave **B**ackground (`CAMB`; Lewis & Challinor 2011)[3] and the **C**osmic **L**inear **A**nisotropy **S**olving **S**ystem (`CLASS`; Blas et al. 2011)[4]. Several common software packages have been developed to manipulate the CMB maps themselves. As mentioned previously, full sky CMB maps are typically represented with `HEALPix` pixelization (Górski et al., 2005)[5]. Standard `HEALPix` software includes routines for map-space and spherical harmonic operations. Another useful package is `pixell`[6], which is maintained by the Simons Observatory team, and contains many utility functions for CMB map analysis.

#### LiteBIRD: Multi-frequency probe of CMB Polarization

The Lite satellite for the study of B-mode polarization and Inflation from cosmic background Radiation Detection (LiteBIRD; LiteBIRD Collaboration et al. 2023) is a proposed satellite mission led by the **J**apanese **A**erospace E**x**ploration **A**gency (JAXA) with the primary goal of detecting an inflationary signature in the low multipole $\ell < 30$ moments of the $B$-mode polarization field of the CMB. Its main strategy to improve on existing observations is to include three separate focal planes in a Low Frequency Telescope (LFT), a Medium Frequency Telescope (MFT), and a High Frequency Telescope (HFT). By combining a large number of detectors in 15 frequency bands 34 to 448 GHz. For comparison, the *Planck* mission observed over 9 different frequency bands (Planck Collaboration et al., 2020b). Observing in more frequency bands allows for better characterization of foregrounds, which are known to be a significant source of contamination in the $B$-mode signal (BICEP2 Collaboration et al., 2014; Flauger et al., 2014; Bonaldi et al., 2014). We also saw before that adding more detectors to an experiment helps to beat down noise. As a satellite mission, LiteBIRD also has access to the entire sky to observe at lower multipoles at lower cosmic variance (corresponding to lower errors in Equation (1.33)) as well as avoiding most of the atmospheric contamination that would exist in ground-based observations (Errard et al., 2015). The LiteBIRD target sensitivity is $r \gtrsim 10^{-3}$, which is approximately a factor of 10 improvement over the current upper limit of $r \lesssim 0.06$ (Planck Collaboration et al., 2020b), so its design will need to overcome a plethora of noise, systematics, and contaminants to constrain such a weak signal.

## 1.3 Measures of Cosmic Distance

### 1.3.1 The Cosmic Distance Ladder

Measurements of astronomical distance are a challenge relevant to all fields of astronomy. One major application for cosmology is its application for quantifying the Hubble law at late times. Accurate measurements of distances to “nearby” galaxies in the Hubble flow can be combined with spectroscopic redshift observations to estimate the Hubble constant $H_0$ (Wong et al., 2020; Riess

---

[3] https://camb.info/

[4] https://github.com/lesgourg/class_public

[5] https://healpix.sourceforge.io/

[6] https://github.com/simonsobs/pixell

et al., 2022; Uddin et al., 2023; Scolnic et al., 2023). The distances to these galaxies are usually estimated with standardizable candles such as Type Ia supernovae (SNIa), which are extremely bright catastrophic events. We do not know the brightness of these supernovae a priori, there is observational evidence that the shape of their light curves is correlated with their peak luminosity (Fausnaugh et al., 2021; Scolnic et al., 2022; Uddin et al., 2023). The luminosities of these supernovae may be calibrated if they are observed in host galaxies for which we have a distance estimate established from other method. These nearby distance measures generally also require some sort of calibration, so the multistep approach to estimating the distance to the far-away galaxies is usually referred to as the "distance ladder". While the particular objects and techniques may vary from analysis to analysis, the general theme remains of calibrating distance measures to increasingly distant objects.

Distance ladder measurements of the Hubble constant have come a long way since its humble beginning: Edwin Hubble originally estimated $H_0 = 500$ km s$^{-1}$ Mpc$^{-1}$ (Hubble, 1929). The **S**upernova $\mathbf{H_0}$ for the **E**quation of **S**tate (SH0ES) team presents perhaps the most prominent contemporary study, which employs an empirical period–luminosity (PL) relation for Cepheid variable stars just as Hubble did (Riess et al., 2022, 2023). Cepheids are remarkably bright variable stars with a stable pulsation period that can be tied to their luminosity through well understood physics (Madore & Freedman, 1991). The PL relation for Cepheid variable stars is calibrated with geometric distance measures to 'nearby' Cepheids in either the Milky Way itself, or local galaxies like the Large Magellanic Cloud (LMC), M31, or NGC4258 (Riess et al., 2022). Once a PL relation is established, it can be used to determine distances to intermediate distance galaxies with both observable Cepheids and SNIa. Supernova observations are typically recorded in catalogues such as the Pantheon(+) sample (Scolnic et al., 2022) or the Carnegie Supernova Project (Uddin et al., 2023). The intermediate distances are once again used to calibrate the absolute brightness of SNIa. Once calibrated, observations of SNIa light curves in galaxies in the Hubble flow can be used to determine the distance to these galaxies. The combination of these distance measures with spectroscopic redshifts allows for the estimation of the Hubble constant $H_0$. The full distance ladder constructed by the SH0ES team is shown in Figure 1.6 (Riess et al., 2022, Figure 12).

Riess et al. (2022) use their empirical determination of the Cepheid PL relation to calibrate the Pantheon+ sample of SNIa (Scolnic et al., 2022) and estimate $H_0 = 73.04 \pm 1.04$ km s$^{-1}$ Mpc$^{-1}$, corresponding to just over $\sim 1\%$ precision. They also present their "error budget" which breaks down their uncertainty estimate into contributions from each component of the analysis, i.e. each step in the distance ladder. The largest share of the error budget belongs to the calibration of supernovae brightness (shown in Figure 24 of Riess et al. 2022). The population of observed supernovae is already limited by randomness, and the subset of supernovae which can be used for calibration is small in correspondence with the local volume containing galaxies with observable objects in the lower steps of the distance ladder. In other words, the sample size of sources in the center panel of Figure 1.6 needs to be larger to better empirically calibrate the SNIa luminosity. The supernova calibration is expected to improve over time. Sample sizes will grow significantly with wide searches being performed by surveys such as the upcoming Vera Rubin Observatory (Bianco et al., 2022), the Zwicky Transient Facility (Bellm et al., 2019), and Transiting Exoplanet Survey Satellite (TESS; Ricker et al. 2014; Holwerda et al. 2021). Individual observations of supernovae will also improve with the advent of the James Webb Space Telescope (JWST; Gardner et al. 2006; Riess et al. 2023).

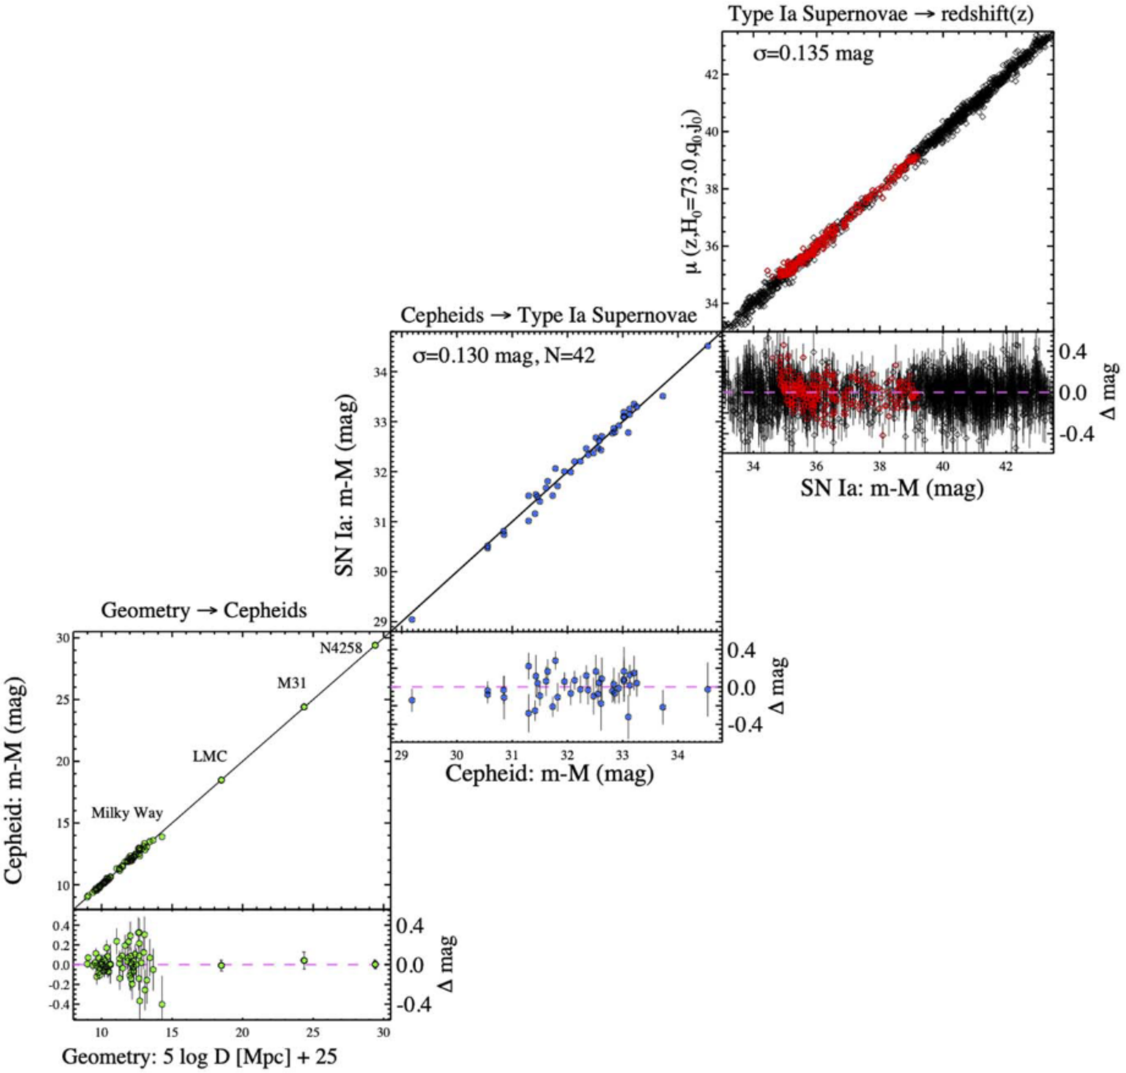

Figure 1.6: The distance ladder constructed by Riess et al. (2022, Figure 12, reproduced by permission of the AAS). The first rung of the ladder is on the bottom left, and contains the sample of nearby Cepheids used to calibrate the PL relation. The LMC contains 70 observable Cepheids that are used in the calibration. The second rung, in the center, contains the sample of nearby galaxies with observable Cepheid variables and at an observed SNIa. The final rung, in the top right, contains the sample of SNIa observed in the Hubble flow used to estimate $H_0$.

### Absolute distance calibration and the Large Magellanic Cloud

Another important component of the $H_0$ error budget is the absolute scale calibration of distances to the closest objects in the distance ladder (shown in Figure 24 of Riess et al. 2022), since any discrepancy at the earliest steps propagates further down to each subsequent step in the distance ladder. The absolute scale calibration relies on geometric measures of distance; for Milky Way Cepheids, the parallax distance is used. A calibrator source (e.g., Cepheid) at the lowest rung in the distance ladder has their luminosity calibrated with respect to their geometric distance measure $r_{\text{cal}}$: $L_{\text{cal}} \propto r_{\text{cal}}^2$. The same types of calibrator sources are then used to calibrate the luminosity of SNIa in 'nearby' galaxies: $L_{\text{SN}} \propto L_{\text{cal}}$. The SNIa luminosity is then used to infer the distance to 'distant' galaxies $r_{\text{gxy}}$ in the Hubble flow: $L_{\text{SN}} \propto r_{\text{gxy}}^2$. The Hubble constant is then inversely proportional to the distances at the farthest part of the distance ladder, and we see that the geometric distance determination of the calibrator sources has qualitatively the same influence: $H_0 \approx v_{\text{gxy}}/r_{\text{gxy}} \propto r_{\text{cal}}^{-1}$. Local extragalactic populations of Cepheids are particularly useful for the first step of the distance calibration. The Riess et al. (2022) study in particular uses samples of Cepheids in the LMC, M31, and NGC4258. If a geometric distance to one of these galaxies can be obtained, then one is able to calibrate the PL relation for a larger sample ($\mathcal{O}(10)$) of Cepheids with a single distance measurement to the host galaxy.

The Large Magellanic Cloud (LMC) is a dwarf galaxy in the outer halo of the Milky Way, with its first documented observation in history recorded by al-Sufi (964). Since then, it has been the subject of study as one of our closest examples of a galaxy separate from our own Milky Way (Alves, 2004; Gaia Collaboration et al., 2018a, 2021a). Its proximity also has profound effects on the dynamics of the Milky Way and its components (Erkal et al., 2019; Petersen & Peñarrubia, 2020; Garavito-Camargo et al., 2021; Correa Magnus & Vasiliev, 2022). It is particularly relevant to the calibration of the distance ladder because it contains a sample of 70 Cepheid variables that can be included in the geometric calibration of the PL relation (Riess et al., 2019, 2022). An incredibly precise distance to the LMC is reported to be $R_{\text{LMC}} = 49.59 \pm 0.09$ (statistical) $\pm\, 0.54$ (systematic) kpc by Pietrzyński et al. (2019), using eclipsing binary distance measures. Recent observations of sources in the LMC by the *Gaia* mission open up an opportunity to independently determine a LMC distance with *Gaia* astrometry, which is the focus of Chapter 3 in this thesis. We provide a summary of the *Gaia* mission in the following section. An accurate distance to the LMC allows one to utilize a large sample of distance calibrators as it is one of the nearest extragalactic sources, so it is considered a vital anchor in the cosmic distance ladder.

## 1.3.2 *Gaia*: Ultra-Precision Astrometry

The *Gaia* satellite is tasked with performing the world's largest simultaneous astrometric, photometric, and spectroscopic survey (Gaia Collaboration et al., 2016). In the first three of its five planned data releases, the *Gaia* mission has significantly advanced our understanding of Galactic dynamics, as well as our understanding of the formation and evolution of our own Milky Way. Some of the most influential discoveries in the *Gaia* era come from projecting the high-precision astrometric data into kinematic phase spaces (Antoja et al., 2018; Helmi et al., 2018). The *Gaia* data have revealed resonances in certain phase spaces that can be traced to our Galaxy's self interactions with its internal structure (Binney & Schönrich, 2018; Bennett & Bovy, 2019; Khoperskov et al., 2019; Hunt et al., 2022; Frankel et al., 2023; Tremaine et al., 2023), or to merger events with dwarf galaxies in the distant past (Lane et al., 2023; Helmi et al., 2018). *Gaia* data have also been used to precisely measure the distance to the Galactic Center as well as the pattern speed of the Galactic bar (Leung et al., 2023), as well as characterize the orbits of dwarf galaxies (Fritz et al., 2018). The *Gaia* data can also be combined with high-resolution spectroscopic surveys to uncover the distribution of chemical structure and stellar populations in the Milky Way (Patil et al., 2023).

### Data Release 2 and the Parallax Zero Point

With its second data release (DR2) in April 2018, the survey accrued astrometric and photometric measurements for over one billion sources down to a magnitude of $G \lesssim 21$ (Gaia Collaboration et al., 2018b). The *Gaia* satellite measures absolute parallaxes by comparing positions of stars in two fields of view (FOVs) widely separated by the "basic" angle $\Gamma_\mathrm{f} = 106.5°$ along the plane of its scanning motion. For a source passing through the center of one of the FOVs, its position is described by $\phi = \pm\Gamma_\mathrm{f}/2$ with respect to the axis bisecting the *formal* basic angle in the same plane. An analytical solution shows that perturbations to such a source's observed parallax ($\varpi$) is degenerate with perturbations to the true basic angle $\Gamma$ (Equation 15 in Butkevich et al. 2017), repeated here

$$\delta\Gamma = \zeta \sin(\phi)\delta\varpi = \zeta \sin(\Gamma_\mathrm{f}/2)\delta\varpi, \tag{1.34}$$

where $\zeta$ is a function of the satellite's orientation with respect to the Solar System barycenter. In other words, oscillations in the instrument's basic angle are degenerate with an absolute perturbation in the observed parallax. This effect is largely corrected by the on-board basic angle monitor, but the scanning motions of the spacecraft leave residual contributions to the parallax zero point that are difficult to model.

A similar effect is also expected to arise from the use of distinct calibration units for sources of different apparent magnitudes (Riello et al., 2018). In addition, *Gaia*'s astrometric solution relies on the position of each source's centroid on the CCD, which is affected by the brightness of the source. There were two methods of determining the positions of centroids: both dim and bright sources ($G \gtrsim 13$ and $G < 13$ respectively) have the positions of their centroids measured as they cross a fiducial line on the CCD, and bright sources $G > 13$ also have their orthogonal positions determined with the entire column traced out by the centroids as they drift across the detector (Lindegren et al., 2018). This extra dimension of position measurement could leave different systematics within the astrometric measurements. In addition, the astrometric calibration used in DR2 makes use of an effective wavenumber determined using mean integrated blue ($G_{BP}$) and red ($G_{RP}$) photometric magnitudes. The unique observed colour of each source is expected to contribute to fluctuations to the astrometric solution, equivalent to a parallax zero point dependence on observed colour. Finally, the parallax zero point has also been shown to vary on large scales across the sky by mapping the observed parallaxes of quasars in the DR2 sample (Arenou et al., 2018).

In general, *Gaia* parallaxes have been reported to be too small. The parallax zero point offset in DR2 has been reported by the *Gaia* collaboration to be $\varpi_0 = -29 \pm 1$ $\mu$as, measured using the parallaxes of quasars in the DR2 sample (Lindegren et al., 2018). This is further supported when they compare DR2 parallaxes of various globular clusters to those found in literature ($\varpi_0 \approx -25$ $\mu$as), but a slightly different result was found using dwarf spheroidal galaxies ($\varpi_0 \approx -49$ $\mu$as) (Khan et al.,

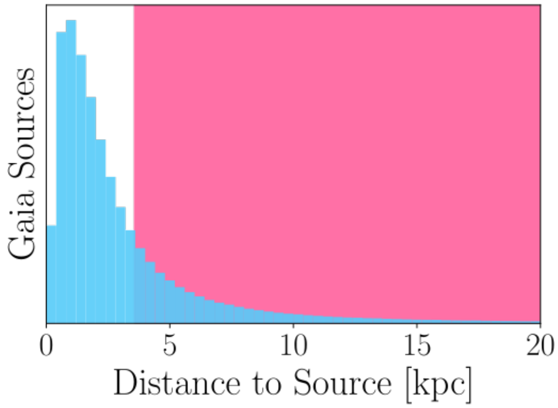

Figure 1.7: Distribution of inverse parallaxes in *Gaia* DR2. The shaded pink region indicates sources for which a 50 $\mu$as systematic in the parallax would correspond to at least a 20% error. There is overlap of this region with approximately 25% the *Gaia* data.

2019). Riess et al. (2018) found $\varpi_0 = -46 \pm 13$ $\mu$as by combining *Gaia* parallaxes with Hubble Space Telescope photometry of Milky Way Cepheids. Stassun & Torres (2018) report $\varpi_0 = -82 \pm 33$ $\mu$as, and Graczyk et al. (2019) determine a zero point of $\varpi_0 = -31 \pm 11$ $\mu$as from analyses comparing existing measurements of parallax for eclipsing binary stars with *Gaia* parallaxes. Zinn et al. (2019) report a modelled $\varpi_0 = -52.8 \pm 2.4$ (random) $\pm 8.6$ (systematic) $-(150.7 \pm 22.7)(\nu_{\rm eff} - 1.5) - (4.21 \pm 0.77)(G - 12.2)\mu$as with significant dependence on the effective wavenumber and observed magnitude using asteroseismology of red giant branch stars. Similarly, Sahlholdt & Silva Aguirre (2018) report a parallax zero point of $\varpi_0 = -35 \pm 16$ $\mu$as using the asteroseismology of dwarf stars, and Hall et al. (2019) use the asteroseismology of red clump stars to determine a mean parallax zero point of $\varpi_0 = -41 \pm 10$ $\mu$as, with individual estimates of $\varpi_0 = -38 \pm 13$ $\mu$as in $K_s$ and $\varpi_0 = -42 \pm 13$ $\mu$as in $G$. Using deep learning of spectro-photometric distances, Leung & Bovy (2019a) determine a modeled constant zero point of $\varpi_0 = -52.3 \pm 2.0$ $\mu$as and they present quadratic parameterizations of the zero point's dependences on $G$, observed colour, and effective temperature.

If one considers a constant parallax zero point of $\varpi_0 \sim 50$ $\mu$as in *Gaia* DR2 data, then Figure 1.7 indicates that approximately 25% of the entire DR2 sample would be impacted with a 20% systematic. A solid understanding of the *Gaia* parallax zero point is necessary to realize the full potential of its astrometric dataset, and we present a robust estimate of the parallax zero point in DR2 ($\varpi_0 = -48 \pm 1$ $\mu$as) along with characterization of its dependence on observed magnitude, colour, and source position in Chapter 2 of this dissertation.

### Data Release 3

The *Gaia* collaboration published its third data release in two stages. The Early Data Release 3 (EDR3) was put out in late 2020 (Gaia Collaboration et al., 2021b), and it contains improved photometric (Riello et al., 2021) and astrometric (Lindegren et al., 2021a) data for over 1.8 billion sources down to a magnitude limit of $G < 21$. It notably features a more robust analysis of the parallax zero point in the wake of a harsh requirement for accurate astrometry in the DR2 era. The *Gaia* collaboration now provides the `gaiadr3-zeropoint`[7] to query the parallax zero point, and

---

[7] https://pypi.org/project/gaiadr3-zeropoint/

takes into account the various dependencies on each source's features (Lindegren et al., 2021b).

The full Data Release 3 (DR3) was published in mid-2022 (Gaia Collaboration et al., 2023). It contains low-resolution $XP$ spectra within its photometric bands, as well as the high-resolution RVS spectra used to measure the radial velocities. The low-resolution $XP$ spectra come from within their $BP \in [330, 680]$ nm, as well as $RP \in [640, 1050]$ nm photometry bands. The $XP$ spectral data are available in a continuous representation of 55 coefficients per band which are readily transformed into the conventional sampled spectra with provided basis functions (Carrasco et al., 2021; Gaia Collaboration et al., 2023). $XP$ sampled spectra occupy 343 bins from 336 nm to 1020 nm. There are 220 million objects with $XP$ spectra available out of 1.8 billion total sources in the DR3 catalogue. The *Gaia* $XP$ spectra are very low-resolution, and the RVS spectra cover only a small wavelength range, so the information contained within an individual object's *Gaia* spectra may not be much. However, the sheer size of the *Gaia* catalogue makes it a promising resource if one is able to statistically combine many objects with similar features.

Chapter 3 presents a calibration of bright giant stars in the LMC using the $XP$ spectra in *Gaia* DR3 combined with precise astrometry of nearby, similar giant stars. The calibration is then used to estimate the distance to the LMC, which has profound implications for the local estimate of $H_0$. An independent estimate of the LMC distance either lower than or in agreement with the Pietrzyński et al. (2019) value would exacerbate or verify tensions between local and early-Universe measurements $H_0$; whereas, an independent estimate of the LMC distance higher than the Pietrzyński et al. (2019) value would alleviate the tension. We further discuss the $H_0$ tension in the following sections.

#### Future Data Releases

While the *Gaia* mission is past its original end date, it has been extended well past the time of writing. The *Gaia* collaboration has announced at least two more data releases past 2025, and 2030[8]. These future releases are expected to contain even more sources, a fainter magnitude limit, as well as all around improved astrometry, photometry, and spectra simply due to the longer observation period.

### 1.3.3 *Planck*: The Inverse Distance Ladder

Measures of cosmic distance can also work forwards from the edge of the observable Universe, in contrast to the distance ladder. The observed CMB power spectra are sensitive to a wide range of cosmological parameters, and the Hubble constant is one of them. The BAO sound horizon is a feature in the matter power spectrum set by the photon-baryon interactions described in §1.2.2. The sound horizon at recombination, described by Equation (1.17), is related to the angular size of the peak of the CMB $TT$ power spectrum through the angular diameter distance. The angular diameter distance to recombination can be parameterized with respect to the redshift of recombination if the spatial curvature of the Universe is known (Hu & White, 1997). The ΛCDM model imposes a flat Universe with zero curvature, so one may estimate $H_0$ using the location of the peaks of the CMB power spectra under this restriction (Bennett et al., 2013; Planck Collaboration et al., 2020a). The *Planck* collaboration reports $H_0 = 67.37 \pm 0.54\,\mathrm{km\ s^{-1}\ Mpc^{-1}}$ under the ΛCDM model (Planck Collaboration et al., 2020a), which is in $5\sigma$ disagreement with the SH0ES result (Riess et al., 2022).

---

[8] `https://www.cosmos.esa.int/web/gaia/release`

The main methodology to add flexibility to the CMB estimate of $H_0$ is to relax the restrictions on spatial curvature. This is done by removing the zero curvature restriction, and supplementing the CMB likelihood with an alternate measure of angular diameter distance to break the degeneracy between $H_0$ and the spatial curvature (Aubourg et al., 2015). This information can be included in the form of measuring the size of the BAO sound horizon through correlations in the distributions of galaxies in clustering data (see Bassett & Hlozek 2010 for a review). The BAO scale, in this case, is used as a 'standard ruler' in the 'inverse distance ladder' without the need to restrict the analysis to a particular cosmology with fixed curvature (Aubourg et al., 2015; Lemos et al., 2019; Camarena & Marra, 2020). These analyses generally agree with the CMB + ΛCDM measurements, or they estimate a value of $H_0$ that is 1-2$\sigma$ away from either the CMB or SH0ES value.

As a final note, the future of CMB measurements of the Hubble constant can be improved after applying delensing of the CMB maps (Hotinli et al., 2022; Ange & Meyers, 2023). We discussed in §1.2.4 the effect of peak smearing in the CMB temperature power through the effects of gravitational lensing. Since these peaks are used to estimate the angular diameter of the BAO sound horizon at recombination, the effect of the lensing is to add uncertainty to the BAO sound horizon. The process of delensing a CMB map requires a reconstruction of the lensing deflection field $\boldsymbol{\alpha}$, and an inverse lensing operation can be applied to recover an estimate of the unlensed CMB fields (Reinecke et al., 2023).

### 1.3.4 Alternative Distance Measures and Probes of the Hubble Constant

The early-Universe estimates of $H_0$ are generally restricted to include to CMB measurement of the BAO scale, since it is one of the probes from the earliest moments of the Universe's history that we can use. One other way to include information from the early Universe is to place limits on the baryon abundance at early times through Big Bang Nucleosynthesis (BBN) abundances (Abbott et al., 2018; Schöneberg et al., 2019). This information, combined with BAO measurements allows for an estimate of $H_0$ without CMB information, but is consistent with those estimates.

There are several other methods for estimating the distances to galaxies in the Hubble flow (and by extension estimating the Hubble constant). Several studies have replaced the Cepheids in the early steps of the distance ladder calibration with Tip of the Red Giant Branch (TRGB) stars, which are a standard candle that are expected to be less affected by dust extinction due to their red colour Freedman et al. (2019); Scolnic et al. (2023); Uddin et al. (2023). A calibration of TRGBs originally concluded $H_0 = 69.8 \pm 0.8\ (\text{stat.}) \pm 1.7\ (\text{sys.})\,\text{km s}^{-1}\ \text{Mpc}^{-1}$, and seemed to provide relief to the tension between early- and late-Universe estimates of $H_0$ (Freedman et al., 2019; Freedman, 2021). However, subsequent studies of the TRGB calibration (combined with other standard candles like Cepheids in Uddin et al. 2023, or as the sole calibrator for SNIa in Scolnic et al. 2023) have placed estimates of $H_0$ back in agreement with other late-Universe estimates.

Another late-Universe estimate of the Hubble constant uses time delay cosmography. This bypasses the requirement to construct and calibrate a distance ladder by taking advantage of the angular diameter distance dependence on the Hubble parameter $H(z)$. The basic setup requires observations of a quasar that has gravitationally lensed by a foreground galaxy to produce multiple images. The photons that are observed as multiple images take different paths around the host galaxy, and as such arrive at the observer at different times. The time delay observed between the images can then be observed through the natural variability of the quasar, and the time delay can be related to the angular diameter distance to the lensing galaxy, which is sensitive to the Hubble constant. A more detailed description of time delay cosmography is reviewed by Birrer et al. (2022). A recent estimate of $H_0 = 73.3^{+1.7}_{-1.8}$ km s$^{-1}$ Mpc$^{-1}$ with time delay cosmography was reported by the H0LiCOW collaboration (Wong et al., 2020), which is in good agreement with other late-Universe measurements. This was later corroborated by Millon et al. (2020); Birrer et al. (2020), although the same team showed that the time delay cosmography estimate of $H_0$ can share a strong dependence on the assumed mass profile of the lensing galaxy (Gilman et al., 2020; Yıldırım et al., 2023). This is known as the mass sheet degeneracy, and subsequent analysis using separate observations of the stellar kinematics in a lensing galaxy to resolve the mass sheet degeneracy have produced results consistent with the original time delay cosmography estimate (Shajib et al., 2023).

Among a host of possible distance ladder calibrations, methods like time delay cosmography seem appealing if one wishes to avoid compounding measurement uncertainties through the combination of observations. There is potential to use the observations of gravitational wave (sometimes shortened to GW) transients as 'standard sirens'. The absolute luminosity of catastrophic merger event signals is associated with the amplitude, or 'strain' of the gravitational wave signal under the assumption of a model for gravity, typically standard general relativity (Holz et al., 2018). If the particular merger event happens to produce an electromagnetic counterpart, then the event may be localized to a host galaxy whose spectroscopic redshift may be measured. The combination of distance and redshift information can then be used to infer the Hubble constant, as was performed by Abbott et al. (2017). Their estimate is not very precise, but it proves that gravitational wave merger events can be used as standard sirens if the host galaxy can be localized. Continued observations of gravitational wave merger events will inevitably produce a large enough sample of localized events (Chen et al., 2018).

In a similar vein, the use of fast radio bursts (FRBs) to estimate distance has also been proposed (Walters et al., 2018; Wu et al., 2020). FRBs are extremely short, broadband radio signals that are likely to be extragalactic in origin (see Cordes & Chatterjee 2019 for a detailed review). An effect of radio signals travelling through the intergalactic medium (IGM) is that they experience a frequency-dependent time delay. The time delay is characterized by the dispersion measure (DM), which is a measure of the total integrated column of free electrons in the IGM along the line of sight. FRB signals from host galaxies farther away are expected to experience longer time delays between frequencies as a result of a larger integrated column of free electrons. The measured DM of an FRB may then be used as a very rough measure of the distance to its source. One must localize an FRB signal to its host galaxy in order to measure its spectroscopic redshift, as with the GW standard sirens. Very Long Baseline Interferometry (VLBI) techniques combining observations from existing experiments like the Canadian Hydrogen Intensity Mapping Experiment (CHIME; CHIME Collaboration et al. 2022) make the localization of FRBs possible (Cassanelli et al., 2022; Michilli et al., 2023). The Hubble constant can be inferred with only information provided by FRBs, as shown in James et al. (2022); Hagstotz et al. (2022), and the constraints are expected to improve as more FRBs are localized.

## 1.4 Contemporary Challenges in Cosmology

### 1.4.1 Dark Matter and the Growth of Structure

The *Planck* results (Planck Collaboration et al., 2020a) indicate that approximately 20% of the Universe's matter is baryonic matter (i.e., matter that interacts directly with electromagnetism). The rest (approx. 80%) of the matter is known to interact gravitationally, but not electromagnetically, hence the name dark matter. Dark matter has yet to be observed directly, but its abundance relative to baryonic matter means that it dominates the growth and dynamics of large scale structures. The existence and abundance of dark matter is well documented and supported by a wide variety of evidence across cosmological scales ranging from the CMB (Bennett et al., 2013; Planck Collaboration et al., 2020a), to the Lyman-$\alpha$ forest (Croft et al., 1998; McDonald et al., 2000, 2005; Zaroubi et al., 2006; Palanque-Delabrouille et al., 2013; Iršič et al., 2017b; Villasenor et al., 2023; Rogers et al., 2023), to the distribution and clustering of galaxies (Gil-Marín et al., 2015; Ross et al., 2020; Alam et al., 2021), to the dynamics of galaxies and substructures within galaxies (Posti & Helmi, 2019; Zavala & Frenk, 2019; Gilman et al., 2020; Garavito-Camargo et al., 2021). There have also been extensive studies cross-matching the data between these experiments (Baxter et al., 2018; Chabanier et al., 2019; Sun et al., 2022; Karim et al., 2022; Chang et al., 2022; Givans et al., 2022; Madhavacheril et al., 2023; Qu et al., 2023).

The true nature of dark matter is yet to be fully understood. The ΛCDM model imposes that all of the dark matter is 'cold dark matter' (CDM), which means that the velocity dispersion of whatever particles or objects makes up the dark matter is relatively low. In contrast, 'hot dark matter' (HDM) or 'warm dark matter' (WDM) models have been proposed consider the possibility that the dark matter components are moving at relativistic speeds, or somewhere in between (Primack & Gross, 2001; Boyarsky et al., 2009; Viel et al., 2013; Lovell et al., 2014; Drewes et al., 2017; Iršič et al., 2017c; Murgia et al., 2018; Palanque-Delabrouille et al., 2020; Garzilli et al., 2021). Models of 'fuzzy dark matter' (fDM) have also been proposed, which invoke ultralight particles (such as ultralight axions) to suppress the formation and growth of structure at scales below a few kpc (Hui et al., 2017; Iršič et al., 2017a; Armengaud et al., 2017; Kobayashi et al., 2017; Ferreira, 2021; Dentler et al., 2022; Laguë et al., 2022; Vogt et al., 2023). We refer the reader to Arbey & Mahmoudi (2021) for a deeper review of dark matter in the context of cosmology.

A majority of the observational evidence indicates that at least most of the dark matter is consistent with behaving like CDM at the largest scales of the Universe, but we have begun to observe inconsistencies between the observed structure growth and clustering at the smaller scales and predictions with CDM (Weinberg et al., 2015). The core-cusp problem describes a discrepancy between the innermost shape of predicted mass profiles of massive dark matter halos with CDM (constant-density cores) and the estimated mass profiles from dynamical observations (an increased, sharper, cusp-like density towards the center). A deep discussion of the core-cusp problem is provided by Ludlow et al. (2013). CDM models also appear to predict an overabundance of dark matter substructures compared to the amount of dwarf galaxies we do observe in groups of galaxies (Stoehr et al., 2002; Kravtsov, 2010). This is commonly referred to as the missing satellites problem, although recent observations suggest that the 'missing satellites' around the Milky Way maybe within statistical tolerances (Kim et al., 2018; Nashimoto et al., 2022).

As briefly covered at the end of §1.2.4, some WDM or fDM models present an opportunity to explain the suppression of structures small-scales and simultaneously provide an explanation for what makes up the dark matter (Hui et al., 2017; Drewes et al., 2017; Ferreira, 2021; Vogt et al., 2023). Other explanations include the effects of neutrino mass and related quantities (Kaplinghat et al., 2003; Lesgourgues & Pastor, 2012; Green & Meyers, 2021; Abazajian et al., 2022), dark matter interactions (Tulin & Yu, 2018; Gluscevic et al., 2019; Buen-Abad et al., 2022), and baryonic feedback (Chisari et al., 2019). One of the major motivations for the development and application of a new small-scale CMB lensing estimator Chapters 5-6 is to better provide observations of the small-scale clustering statistics in the late Universe, which in turn provides more robust constraints for models of dark matter and the growth of structure.

### 1.4.2 Discordance between early- and late-Universe observations

The Hubble tension describes apparent disagreement between early- and late- Universe estimates of $H_0$, and it is a hot topic in cosmology. Figure 1.8 is presented in Riess et al. (2019, Figure 4), and it compares $H_0$ measurements from various experiments. The early-Universe measurements are based on the CMB power spectra, and either the ΛCDM model or the inverse distance ladder. The late-Universe measurements are based on the distance estimates to galaxies in the Hubble flow, and they are calibrated with various astrophysical sources. Refer back to §1.3 for an overview of the measurements themselves. The point of tension in Figure 1.8 is that the CMB-based estimates generally agree with each other, and the local estimates are generally consistent with each other. However, there is presently a 4-5$\sigma$ disagreement between CMB-based estimates and local estimates.

The Hubble tension has sparked extensive and careful studies of the treatment and comparison of statistics and systematics in $H_0$ estimates (Chen et al., 2003; Marshall et al., 2006; Chen & Ratra, 2011; Verde et al., 2013; Hee et al., 2016; Charnock et al., 2017), or reconciling the differences with the extension of cosmological models (Di Valentino et al., 2015; Bernal et al., 2016; Di Valentino et al., 2016; Huang & Wang, 2016). An exciting prospect for bringing early- and late-Universe measurements into agreement is the consideration of new physics or extensions to the standard ΛCDM model, some of which are shown in Figure 1.8 along with the direction in which they would shift early-Universe estimates. Thorough reviews compiling the wide variety of theoretical models proposed as solutions to the Hubble tension can be found here (Di Valentino et al., 2021; Hu & Wang, 2023). One compelling direction that is being considered is the possibility of an early dark energy component that contributes to the early Hubble parameter much before recombination ($z \gtrsim 3000$). Such a model serves to reduce the BAO sound horizon at recombination ($r_s$ in Equation (1.17) goes down if $H(z)$ at early times goes up), while leaving late-time BAO predictions unaffected. This shifts the CMB+BAO estimate of $H_0$ higher, and can provide relief to the tension between early- and late-time measurements (Hojjati et al., 2013; Karwal & Kamionkowski, 2016; Poulin et al., 2019; Smith et al., 2021; Kamionkowski & Riess, 2022).

Another source of discordance between early- and late-Universe cosmological measurements is the $\sigma_8$ or $S_8$ tension. The amplitude of matter fluctuations at late-times on a smoothed distance scale of $8h^{-1}$ Mpc is parameterized by $\sigma_8$. Here $h = H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1})$ is a rescaled Hubble constant. There is no profound reason for defining a quantity at specifically the distance scale $8h^{-1}$ Mpc. The importance of $\sigma_8$ is rather to act as a quantitative anchor for the overall clustering of matter that is independent of cosmological modelling. Similarly, the parameter $S_8 \equiv \sigma_8\sqrt{\Omega_m/0.3}$ is a rescaled version of $\sigma_8$ that is parameterized to mitigate its degeneracy with the overall matter

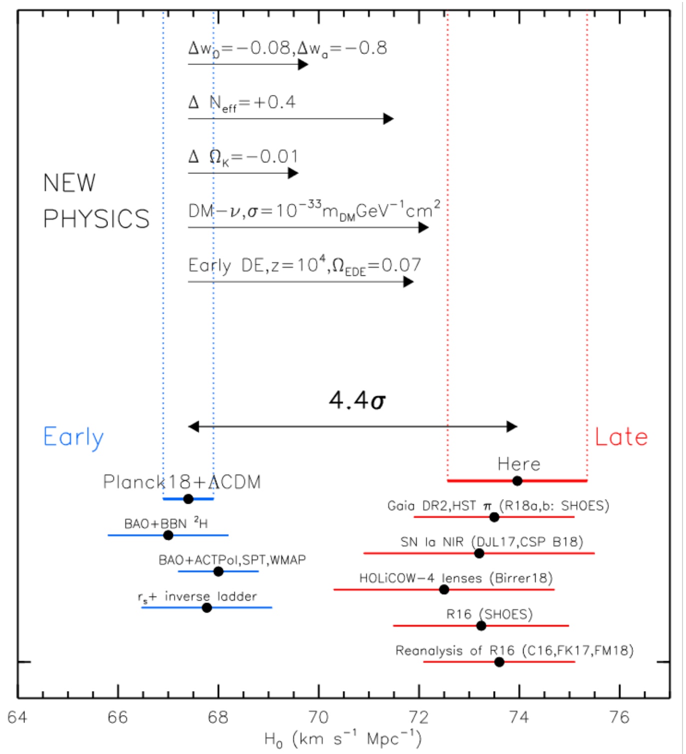

Figure 1.8: A comparison of early- and late-Universe Hubble constant measurements presented by Riess et al. (2019, Figure 4, reproduced by permission of the AAS). The early- and late-Universe measurements are in 4.4$\sigma$ disagreement, and there are several proposed extensions to the standard cosmological model to relieve the tension.

density. The $\sigma_8$ or $S_8$ tension describes the disagreement between the values of $\sigma_8$ or $S_8$ measured by matter-clustering experiments at late-times when compared to expected values from ΛCDM following parameters from CMB measurements. Overall, the measured clustering of late-time matter fluctuations is lower than expected (Gómez-Valent & Solà Peracaula, 2018; Nunes & Vagnozzi, 2021; Poulin et al., 2023).

Chapters 2-3 of this dissertation strive to improve geometric distance estimates with *Gaia* data. Further, independent cosmic distance measures, especially those at the lowest step of the distance ladder, can have a profound impact on the local determination of the Hubble constant as discussed in §1.3. The implications for the Hubble tension are interesting regardless of whether *Gaia* data can be used to resolve the tension, or if its application verifies the discrepancy.

### 1.4.3 Beyond Standard Cosmological Models

The ΛCDM model has, for the most part, had resounding success in describing the evolution of the Universe from the earliest times to the present day. We have seen in the previous sections that the model has some difficulties tying together measurements between observations from the early- and late-Universe. This has motivated a variety of extensions to the so-called standard model to explain the discrepancies, which we have already covered.

An important and popular extension to the standard model is the addition of the inflationary epoch described in §1.2.2 (see Tsujikawa 2003; Baumann & McAllister 2015 for reviews; Baumann 2022 also introduces inflation in Chapter 2). The inflationary epoch is called upon to explain the flatness problem, which lack of spatial curvature ($\Omega_k < 10^{-5}$) in the Universe (Planck Collaboration et al., 2020a). It is also used to explain why the CMB is the same temperature down to a few parts in $10^{-4}$ K. There is a need to explain why the CMB appears to have been causally connected in all directions, which is the horizon problem. Theories from particle physics also predict the existence of exotic species such as magnetic monopoles which we do not observe in nature. We discussed earlier that the initial scalar perturbations in the early Universe could be generated by an epoch of inflation. One may expect such a drastic expansion of space to disperse the Universe's energy away, and inflationary theories typically solve this with a large injection of energy from the physical field or phenomenon driving the expansion. This is often referred to as reheating. Cosmic inflation is typically proposed very shortly after the Big Bang, and facilitates the extremely rapid expansion of the Universe, usually by a factor of at least $e^{\sim 60}$ (Tsujikawa, 2003; Baumann, 2022). The inflationary epoch is invoked to intensely stretch out the (initially causally connected, or close enough together light to travel from one point to the other in the time since the Big Bang $\eta$) fabric of spacetime to smooth out the curvature of space. At the same time, this dilutes the exotic species that are predicted to exist to less than one expected particle per observable volume of the Universe. Finally, as we explained in §1.2.2 inflation be employed to bring quantum fluctuations of curvature perturbations to cosmic scales, which are the initial seeds of density perturbations.

One of the 'smoking gun' pieces of evidence predicted by inflationary theories is the generation of tensor perturbations, or primordial gravitational waves. These are predicted to be observable as the source of a primary signal in the CMB $BB$ power spectrum, see Figure 1.2. Up until now, the ratio between the amplitudes of the tensor and scalar perturbations have been constrained to be $r \lesssim 0.06$ (Planck Collaboration et al., 2020a). The primary challenge with detecting the primary $BB$ signal is its low amplitude, which we can see in Figure 1.2 is several orders of magnitude lower than the

$EE$ signal. To compound to this, the generation of secondary $BB$ power from gravitational lensing dominates over the primary signal to the point that it is no longer visible in Figure 1.2 (Seljak & Hirata, 2004). CMB delensing techniques can serve to remove the secondary $BB$ signal, as discussed in §1.2.4 (Kesden et al., 2002; Knox & Song, 2002; Hotinli et al., 2022). Another additional challenge is the contamination of the CMB polarization signal from dusty foregrounds, which is known to be a nuisance (BICEP2 Collaboration et al., 2014; Flauger et al., 2014; Bonaldi et al., 2014; Mortonson & Seljak, 2014).

We described the LiteBIRD mission at the end of §1.2.5, which aims to observe the CMB polarization signal at a wide variety of frequencies to characterize the foreground contamination (LiteBIRD Collaboration et al., 2023). Their target is to constrain values of $r \gtrsim 10^{-3}$, which requires a robust understanding of its instrument systematics. Chapter 4 of this dissertation describes a study on the impact of electrical crosstalk, an instrument-specific systematic effect, on LiteBIRD observables. We performed realistic simulations of the CMB observation process described in §1.2.5 through the `TOAST` software package that we modified to include the effects of detector crosstalk to investigate the potential impact of crosstalk on the LiteBIRD target sensitivity on $r$.

## 1.5 Modern Statistical and Machine Learning Techniques

### 1.5.1 Bayesian Inference

Precise constraints of physical models now demand robust statistical analysis of cosmological observations. Standard techniques of modern cosmological and astrophysical methodology revolve around Bayes theorem (Dodelson, 2003; Hob, 2009; Joyce, 2021). Within this paradigm, physical models are described with *likelihoods* $p(\mathbf{d}|\boldsymbol{\theta})$ which describe the probability that a set of data $\mathbf{d}$ are observed when given a particular set of parameters $\boldsymbol{\theta}$ which govern the model. Bayes theorem defines the *posterior*, which is the probability distribution of all parameters $\theta$ conditioned on the set of observed data $\mathbf{d}$:

$$p(\boldsymbol{\theta}|\mathbf{d}) = \frac{p(\mathbf{d}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\mathbf{d})}. \tag{1.35}$$

The posterior contains information about the physical model within the likelihood, combined with *prior* knowledge about each individual parameter through $p(\boldsymbol{\theta})$. The *model evidence* $p(\mathbf{d})$ is the marginalized likelihood of the data over all possible values of the parameter space, and it is usually ignored unless a study is comparing the efficacy of different models.

A general data-set $\mathbf{d} = \{d_i\}$ may contain $N$ observables, where $i \in 0, 1, \ldots, N$. The full likelihood function (if the data are independent of each other) can then be a product of distribution functions which describe the probability of each observation $d_i$ being observed under a given set of parameters $\boldsymbol{\theta}$:

$$p(\mathbf{d}|\boldsymbol{\theta}) = \prod_i^N p(d_i|\boldsymbol{\theta}). \tag{1.36}$$

There is a great deal of flexibility when it comes to building physical models that can be described by a likelihood function. One of the simplest (and widely used) likelihood functions is the normal, or Gaussian, distribution. It is parameterized in one-dimension as follows:

$$\mathcal{N}(x|\mu_x, \sigma_x^2) = \frac{1}{\sqrt{2\pi\sigma_x^2}} e^{\frac{(x-\mu_x)^2}{\sigma_x^2}} . \tag{1.37}$$

A quantity $x$ is considered “normally distributed” if it follows a statistical distribution described by Equation (1.37). This set of distributions is characterized by a central value, or mean, $\mu_x$ which at which the probability peaks, and there are symmetrical, exponential tails on either side whose width, or variance, is set by $\sigma_x^2$. Similarly, the multivariate normal distribution is parameterized with a covariance matrix $\mathbf{\Sigma}$ which describes the correlations between observables:

$$\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \mathbf{\Sigma}) = \frac{1}{\sqrt{(2\pi)^N \det(\mathbf{\Sigma})}} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \mathbf{\Sigma}^{-1} (\mathbf{x}-\boldsymbol{\mu})} . \tag{1.38}$$

Normal distributions can be generalized with higher-order statistics, which can introduce skewness (asymmetric tails) and kurtosis (heavier/lighter tails).

Prior distributions for parameters may be informative or uninformative. One may choose to use prior knowledge about a parameter if it was measured to be a particular value in a previous study. In this scenario, an informative prior distribution may be chosen in a form similar to Equation (1.37), where the previous study’s measurement is used to set the mean and variance of the prior. Uninformative priors are usually chosen to be uniform, or flat, distributions, which means that all values of the parameter are equally likely. One may also choose to use a log-uniform distribution if it is necessary to provide similar weight to a particularly large range of parameter space spanning multiple decades.

The power of constructing probabilistic models in the Bayesian framework is the ability to model the dependence of parameters on each other, or on new *hyperparameters* that are not directly observed. Including this extra complexity transforms the model into a hierarchical model, and it is a powerful tool for describing complex physical systems. One note of caution is that the inclusion of hyperparameters leads to an increase in the dimensionality of the parameter space, which can make it difficult to compute or explore the full posterior distribution. One must also be wary of including too many parameters, in which case the results may become overfitted.

### 1.5.2 Sampling Methods

It is possible to analytically or numerically compute the posterior distribution of simple models for small data sets using Equation (1.35). Modern astrophysical and cosmological models are described by many parameters as well as data-sets are commonly highly dimensional, which both contribute to posterior distributions which are intensely complex. While it is still possible to numerically compute these highly dimensional posterior distributions, it is computationally impractical to explore large volumes of parameter space through brute force. In this case, it is often more practical to ‘sample’ the density of posterior distributions with clever algorithms.

A common, and simple, way to sample posterior distributions is the Markov chain Monte Carlo (MCMC) method. Popular Python implementations of MCMC algorithms include `emcee`[9] (Foreman-Mackey et al., 2013), and `PyMC`[10] (Salvatier et al., 2015). Cosmology specific packages include

---

[9] https://github.com/dfm/emcee

[10] https://github.com/pymc-devs/pymc

`CosmoMC`[11] (Lewis & Bridle, 2002), and `Cobaya`[12] (Torrado & Lewis, 2019, 2021). We make particular use of `emcee` for a simple cosmological parameter estimator in Chapter 6 due to its simplicity and ability to construct log-probabilities with blackbox functions. MCMC algorithms are generally simple to implement, but they are not always efficient at sampling complex parameter spaces, especially for models with high dimensionality.

An alternative to more traditional MCMC algorithms is a sub-class called Hamiltonian Monte Carlo (HMC) algorithms. These make clever use of Hamiltonian dynamics to more efficiently sample the parameter space, especially for models with highly correlated parameters (Neal, 2012). The samples produced by an HMC algorithm also generally experience lower levels of auto-correlation than those produced by MCMC algorithms. A popular software implementation of HMC algorithms is `Stan`[13] with Python implementation `PyStan`[14] (Carpenter et al., 2017), which we make use of in Chapters 2-3.

### 1.5.3 Deep Learning and Emulators

Datasets across astronomy and cosmology are increasingly growing in size and complexity. Deep learning tools have seen widespread use among other machine learning techniques in the era of data-driven astronomy. The application of Neural Networks (NN) have seen success in the classification of astronomical transients from light curves (Hložek et al., 2023), the prediction of stellar parameters from observed spectra (Leung & Bovy, 2019b), and more (Smith & Geach 2023 presents an extensive overview of deep learning in astronomy; see Alzubaidi et al. 2021 for deep learning in general). We make use of stellar parameters predicted by `astroNN`[15] (Leung & Bovy, 2019b) in Chapter 2. These NNs were designed to mimic the structure of animal brains, which are made up of a large number of neurons that are inter-connected in non-trivial ways. In astronomy, they are often used to map a set of $i$ input features $\mathbf{x} = \{x_i\}$ to a set of $j$ output features $\mathbf{y} = \{y_j\}$ through a series of hidden layers. These hidden layers are composed of 'connections' which combine inputs from the previous layer with a non-linear function to produce outputs for the next layer. A NN can be trained to learn the output features with a set of labelled training data $\{\mathbf{x}_t, \mathbf{y}_t\}$ for which the output features are known. During the training process, the NN is optimized to minimize the difference between the predicted output features $\mathbf{y}_p$ and the known output features $\mathbf{y}_t$ by adjusting the weights for each connection's contribution to their next layer. The nature of the training process allows the NN to learn complex relationships between the input and output features, which can be used to predict the output features for new input features. The caveat is that the NN is only able to make predictions within the range of input features that it was trained on, similar to an interpolation. The `TensorFlow`[16] (Abadi et al., 2015) and `PyTorch`[17] (Paszke et al., 2019) packages provide frameworks for constructing and training NNs in Python.

An application of modern machine learning techniques in cosmology is the emulation of cosmological power spectra. The computation of power spectra from Boltzmann codes was briefly discussed in §1.2, and they are able to produce outputs from cosmological parameters on the order of $\mathcal{O}(1\,\mathrm{s})$. This

---

[11] https://github.com/cmbant/CosmoMC
[12] https://github.com/CobayaSampler/cobaya
[13] https://github.com/stan-dev/stan
[14] https://github.com/stan-dev/pystan
[15] https://github.com/henrysky/astroNN
[16] https://github.com/tensorflow/tensorflow
[17] https://github.com/pytorch/pytorch

kind of speed is unattainable for power spectra predicted from cosmological simulations which hope to include non-linear effects in the growth of structure, so attention has been turned to emulation to predict power spectra from cosmological parameters (Lawrence et al., 2010; Rogers et al., 2019; Pedersen et al., 2021; Jamieson et al., 2023). Applications of Boltzmann codes for the computation of power spectra can also become a bottleneck for parameter estimation pipelines using the techniques described in §1.5.2. The sampling process may require new power spectra to be generated at every step in the parameter space, and replacing a Boltzmann code with an emulator trained on its output may yield significant speedups (Spurio Mancini et al., 2022). We train two power spectrum emulators within the `CosmoPower`[18] (Spurio Mancini et al., 2022) framework in Chapter 6 to speed up cosmological parameter estimation.

## 1.6 Thesis Overview

In this chapter, we reviewed the state of the art in cosmological observations, the theoretical models that are used to describe them, as well as some statistical and computational methods that are used in the field. We put a particular focus on observations of the cosmic microwave background, as well as measures of cosmic distance, in order to set the stage for the problems in contemporary studies for which we aim to contribute new understanding for through the contents of this thesis.

In Chapters 2-3, we present work on using data from the *Gaia* satellite to improve knowledge of the earliest steps of the local distance ladder, with the ultimate goal of improving the local determination of the Hubble constant $H_0$. We begin in Chapter 2 with a study of the systematic parallax zero point that is persistent in *Gaia* DR2 astrometry with a probabilistic model that simultaneously calibrates a luminosity model for red clump stars and infers the parallax zero point. We infer the *Gaia* DR2 parallax zero point to be $\varpi_0 = -48 \pm 1\,\mu$as, which is the most precise determination of the systematic for DR2. We also show extensions to our probabilistic model that allow us to characterize the variations of the *Gaia* parallax zero point as a function of a source's observed magnitude, observed colour, and position on the sky. Chapter 3 continues the study of *Gaia* distances by searching for bright, giant stars in the Large Magellanic Cloud that have spectra similar to nearby, bright giants with high quality parallax measurements in *Gaia* DR3. We infer the distance to the LMC to be $R_{\rm LMC} = 48.90^{+0.52}_{-0.48}$ kpc using a model that assumes the matched stars share the same luminosity and calibrates the absolute brightness using the nearby parallaxes. Our estimate is in agreement with the accepted value reported by Pietrzyński et al. (2019) at similar precision, and it represents the possibility for using *Gaia* data to independently validate the absolute distance calibration at the lowest step(s) of the local distance ladder.

We move on to studies of the cosmic microwave background in Chapter 4, where we tailor existing simulations of CMB satellite observations to include the effects of detector crosstalk. These simulations allow us study the impact of realistic crosstalk on the LiteBIRD mission's observables, and we find that the impact of *unmitigated* crosstalk will affect CMB polarization power spectra by at most a few percent. The results of this chapter bode well for the LiteBIRD mission's prospects of meeting its target requirements for constraining the tensor-to-scalar ratio to $r \gtrsim 10^{-3}$.

Chapter 5 shifts our focus one final time to the study of the small-scale ($\ell \gg 3000$) gravitational lensing of the CMB. We present the development of the Small Correlated Against Large Estimator,

---

[18] https://github.com/alessiospuriomancini/cosmopower

which we show can outperform traditional quadratic estimators in terms of the signal-to-noise of recovered CMB lensing potential power spectrum at $L \sim 7000$ at the noise levels of upcoming, and future experiments. Chapter 6 continues with the application of SCALE for parameter estimation in simulated maps of the lensed CMB at noise levels similar to the proposed CMB-S4 experiment. We show that the inclusion of SCALE can improve constraints cosmological parameters sensitive to small-scale lensing, with the neutrino mass $m_\nu$ as an example. We show that SCALE can provide enough information to provide a $2.4\sigma$ detection of the minimum neutrino mass $m_\nu = 0.06\,\mathrm{eV}$, which goes undetected if the only observables are CMB power spectra. We conclude in Chapter 7 with a summary of the main results of this thesis in the context of the challenges we set out to address.

# Chapter 2

# *Gaia* DR2 parallax zero point: Hierarchical modelling of the red clump

*Originally published as: “The* Gaia *DR2 parallax zero point: Hierarchical modeling of red clump stars” Chan, Victor. C., & Bovy, Jo 2020, MNRAS, 493, 4367, doi: 10.1093/mnras/staa571. The body of the published manuscript has been adapted for this dissertation with some minor clarifications, and a majority of the introduction has been moved to Chapter 1 (see §1.3.2).*

## Abstract

The systematic offset of *Gaia* parallaxes has been widely reported with *Gaia*’s second data release, and it is expected to persist in future *Gaia* data. In order to use *Gaia* parallaxes to infer distances to high precision, we develop a hierarchical probabilistic model to determine the *Gaia* parallax zero point offset along with the calibration of an empirical model for luminosity of red clump stars by combining astrometric and photometric measurements. Using a cross-matched sample of red clump stars from the Apache Point Observatory Galactic Evolution Experiment (APOGEE) and *Gaia* Data Release 2 (DR2), we report the parallax zero point offset in DR2 to be $\varpi_0 = -48 \pm 1$ $\mu$as. We infer the red clump absolute magnitude to be $M_K = -1.622 \pm 0.004$ in $K_s$, $M_G = 0.517 \pm 0.004$ in $G$, $M_J - 1.019 \pm 0.004$ in $J$, and $M_H - 1.516 \pm 0.004$ in $H$. The intrinsic scatter of the red clump is $\sim 0.09$ mag in $J$, $H$ and $K_s$, or $\sim 0.17$ mag in $G$. We tailor our models to accommodate more complex analyses such as investigating the variations of the parallax zero point with each source’s observed magnitude, observed colour, and sky position. In particular, we find fluctuations of the zero point across the sky to be of order or less than a few $10\,s$ of $\mu$as.

## 2.1 Introduction

SYSTEMATICS in the parallaxes reported in the *Gaia* mission's Data Release 2 (DR2) have been intensely studied to make full use of the astrometric data available. In this chapter, we describe hierarchical Bayesian models inspired by Sesar et al. (2017) and Hawkins et al. (2017) to simultaneously estimate the *Gaia* parallax zero point and constrain an empirical relation for the luminosity of red clump stars. We outline the red clump sample and the data used in these analyses in §2.2. The hierarchical model and several add-ons for additional detailed analysis are described in §2.3. Results of each model/analysis are presented in §2.4, with inferred parameters collected in Table 2.1. We discuss the internal consistencies between each model as well as compare our measurements with other reports in §2.5. Finally, we conclude with a summary in §2.6.

## 2.2 Data

### 2.2.1 Red-clump sample

As a part of SDSS-III/IV (Eisenstein et al., 2011; Blanton et al., 2017), the Apache Point Observatory Galactic Evolution Experiment (APOGEE; Majewski et al. 2017) is a spectroscopic survey in the near infrared. Observing in the infrared is a major advantage for its data, as this allows for measurements that are less affected by dust extinction when compared to optical measurements. The survey data has been pre-processed and is publicly available online as part of the SDSS Data Release 14 (DR14; Holtzman et al. 2015; García Pérez et al. 2016; Abolfathi et al. 2018). This data set includes detailed measurements of each source's chemical abundances (with $S/N > 100$) as well as the stellar parameters $T_{\rm eff}$ and $\log g$ due to its spectroscopic resolution of $R \approx 22\ 500$ (Wilson et al., 2019). We make use of the Two Micron All Sky Survey (2MASS) $J$, $H$, and $K_s$ photometry (Skrutskie et al., 2006). The 2MASS photometry has been previously corrected for reddening with the Rayleigh-Jeans Colour Excess method (Equation 1 in Majewski et al. 2011), yielding measured extinction values $A_J$, $A_H$, and $A_K$ for each object.

We also make use of inferred values for $T_{\rm eff}$, $\log g$, as well as the chemical abundance data, specifically [Fe/H], [O/Fe], [Mg/Fe], [Si/Fe], [S/Fe], and [Ca/Fe]. The abundances for $Z = \{$O, Mg, Si, S, Ca$\}$ are used to construct a value of [$\alpha$/Fe] by calculating

$$[\alpha/\mathrm{Fe}] = \frac{\sum_{z\in Z} w_z [z/\mathrm{Fe}]}{\sum_{z\in Z} w_z}, \tag{2.1}$$

where $w = 1$ or $w = 0$ depending on whether the measurement of [Z/Fe] exists respectively. The abundance measurements upon which [$\alpha$/Fe] is based are not taken from SDSS DR14, but they are instead products of an artificial neural network `astroNN` trained on APOGEE spectra from DR14 (Leung & Bovy, 2019b). Typical sources in `astroNN` have inferences of chemical abundances to $\approx 0.03$ dex, $T_{\rm eff}$ to $\approx 30$ K, and $\log g$ to $\approx 0.05$ dex.

We use only a sub-sample of the APOGEE data set that has been classified as red-clump stars by Bovy et al. (2014). The red-clump stars were identified with strict cuts first in effective-temperature–surface-gravity–metallicity space followed by further cuts in colour–surface-gravity–metallicity space. This was done to minimize contamination from other red giant branch stars within the population.

The Bovy et al. (2014) cuts leave a small amount of contamination ($\lesssim 10\%$) by non red-clump stars, which could be identified with better measurements of $\log g$. In order to further purify the red-clump sample, we apply cuts to sources in the sample for which $(\log g)_{\mathrm{APOGEE}} - (\log g)_{\texttt{astroNN}} < -0.2$. As this quality cut serves to remove suspected contamination sources in the sample, it should improve our calibrations of the red clump luminosity. We do not expect there to be any bias induced on statistical inferences. We provide a brief discussion on the importance of an accurate red clump luminosity calibration in §2.5.1.

### 2.2.2 *Gaia* DR2 data

With the advent of DR2, the *Gaia* mission currently has measurements of sky position and photometry in its $G$-band (330–1050 nm) for nearly 1.7 billion sources (Gaia Collaboration et al., 2018b). Over 1.3 billion of the sources in DR2 also have measurements of parallax ($\varpi$) and proper motion. The current reported limiting magnitude in DR2 is $G \approx 21$, and the data set is reported to be complete over $3 \lesssim G \lesssim 17$.

The *Gaia* mission only reports the 5-parameter astrometric data (position, parallax, and proper motion) for sources that satisfy three requirements: (1) they must be brighter than the limiting magnitude ($G < 21$), (2) each source must have been observed on at least six occasions that are separated by at least four days, and (3) the astrometric parameters must be measured to within a magnitude-dependent uncertainty. The full details of the astrometric solution are discussed by Lindegren et al. (2018).

We cross-match the *Gaia* DR2 sample with the red-clump sample to obtain the $G$ magnitudes, $G_{BP}$ (330–680 nm; also referred to as $BP$) and $G_{RP}$ (630–1050 nm; also referred to as $RP$) magnitudes, as well as parallaxes and their uncertainties. The *Gaia* data set includes measurements of negative parallaxes, which can still contain useful information when combined with their uncertainties. An exception must be made for unphysical measurements of negative parallaxes that are too confident because these must be outliers, so we make further cuts on the quality of the parallax measurements by removing any sources with a measured $\varpi/\sigma_{\varpi} < -3$, or $\varpi = 0$ from our sample.

Additionally, the *Gaia* collaboration has advised that the $G$ magnitudes be corrected according to the following procedure[1]:

$$G_{\mathrm{corr}} = \begin{cases} -0.0473 + 1.164G - 0.0468G^2 + 0.0035G^3, \ 2 < G \leq 6 \\ G - 0.0032(G - 6), \ 6 < G \leq 16 \\ G - 0.032, \ G > 16. \end{cases} \tag{2.2}$$

This correction is applied only when the $G$ magnitudes are used in the context of using $G$ to determine the red clump luminosity through the distance modulus (further elaborated upon in §2.3). Our final data set contains 27,934 red clump stars with values for $\mathbf{d} = \{\varpi, \sigma_{\varpi}, G, G_{BP}, G_{RP}, J, H, K_s, A_J, A_H, A_{K_s}, T_{\mathrm{eff}}, [\mathrm{Fe/H}], [\alpha/\mathrm{Fe}]\}$.

---

[1] `https://www.cosmos.esa.int/web/gaia/dr2-known-issues`

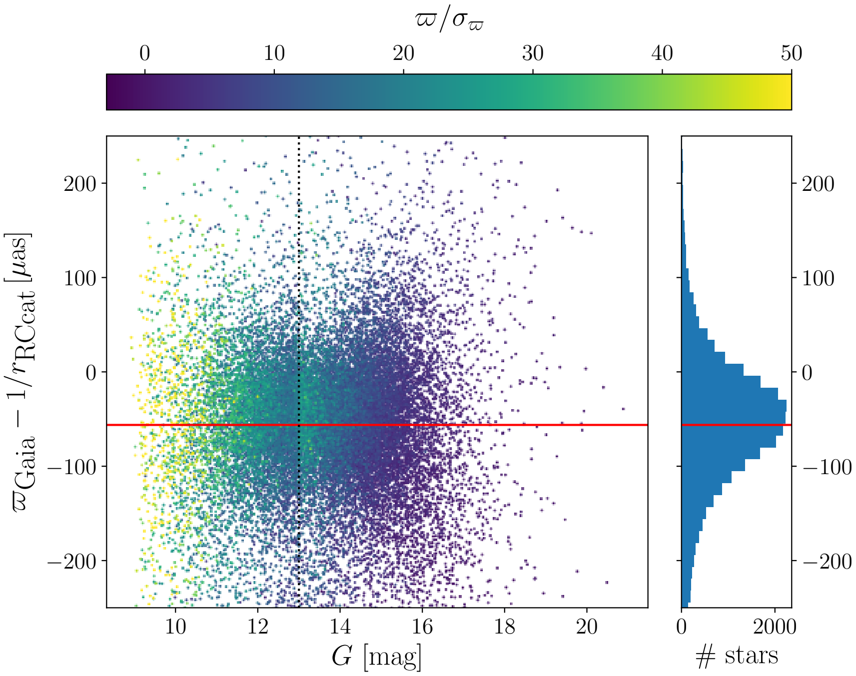

Figure 2.1: The *Gaia* parallax offset assuming the distances to each star in the red clump catalogue are correct. The red line indicates the median of the sample at $-56.2$ $\mu$as. The dotted black line is located at $G = 13$ mag, where the astrometric solution pipeline is reported to change.

### 2.2.3 *Gaia* Extinction Model

As mentioned previously, the 2MASS photometry from the red clump catalog ($J$, $H$, and $K_s$ bands) have been corrected for extinction using the Rayleigh-Jeans Colour Excess method, so each star has associated $A_J$, $A_H$, and $A_K$ extinction values that have been implicitly included in their respective apparent magnitudes as $m_0 = m - A_m$. The *Gaia* $G$ band has *not* been extinction corrected, so we use the following procedure to estimate the $G$ band extinction coefficient.

For each star in the red clump sample, we compute synthetic $G$ and $K_s$ band photometry using the `pystellibs`[2], `pyphot`[3], and `pyextinction`[4] tools, and we determine the $A_G/A_K$ extinction ratios with the following steps:

1. We generate initial synthetic stellar spectra (denoted $F_0(\lambda)$) with the Castelli & Kurucz (2003)/Kurucz (2005) stellar model library in `pystellibs` for each source using their respective $T_{\rm eff}$, $\log g$, and $Z$.

2. The extinction-free $K_s$ band magnitude (denoted $K_0$) is computed with `pyphot` for each stellar spectrum $F_0(\lambda)$ using the 2MASS $K_s$ passband (Cohen et al., 2003).

3. Assuming a value of $A_V$, we use the Fitzpatrick (1999) extinction law $A(\lambda)/A_V$ computed using the `pyextinction` tool, redden each stellar spectrum with $A(\lambda) = A_V \times A(\lambda)/A_V$, yielding

---

[2] https://github.com/mfouesneau/pystellibs
[3] https://github.com/mfouesneau/pyphot
[4] https://github.com/mfouesneau/pyextinction

$F_r(\lambda) = F_0(\lambda) \times 10^{-0.4A(\lambda)}$. We then require knowledge of $A_V$ for each star in order to apply the correct extinction to each stellar spectrum.

4. The extinguished $K_s$ band magnitude (denoted $K$) is computed with `pyphot` for each reddened spectrum $F_r(\lambda)$ using the 2MASS $K_s$ passband.
5. By repeated application of the previous two steps, we solve for each $A_V$ value that matches the given $A_K$ value obtained from the RJCE method. We then redden each stellar spectrum again using this value of $A_V$, yielding $F_R(\lambda)$.
6. The extinction-free $G_0$ magnitude is computed with `pyphot` for each un-reddened stellar spectrum $F_0(\lambda)$ using the Maíz Apellániz & Weiler (2018) revised $G$ passband. Similarly, we compute each extinguished $G$ from the reddened spectrum $F_R(\lambda)$.
7. The extinction ratio for each star is then produced as $A_G/A_K = (G - G_0)/A_K$.

The $G$ band extinction correction is then simply given by $(A_G/A_K)A_K$, the extinction ratio determined using the synthetic photometry multiplied by the $A_K$ value obtained from the RJCE method. We find that our sample of red clump extinction ratios are well approximated by

$$\begin{aligned} A_G/A_K = 7.09 - 2.16A_K + 0.68A_K^2 + 0.18A_K^3 - 0.12A_K^4 \\ + 0.49/1000\mathrm{K}(T_{\mathrm{eff}} - 4835\mathrm{K}). \end{aligned} \tag{2.3}$$

This parameterization is shown along with the sample and its residuals in Figure 2.2.

## 2.3 Joint Luminosity and *Gaia* parallax zero point offset calibration methodology

### 2.3.1 Hierarchical modeling of the red clump in *Gaia* DR2

We construct a probabilistic model based on previous analyses of *Gaia* observations of standard candles. In particular, our model is inspired by Sesar et al. (2017), in which RR Lyrae were used to simultaneously validate *Gaia* DR1 parallaxes and fit a luminosity function for the stars. With parallax calibration parameters $\theta_\varpi$, red clump luminosity function parameters $\theta_{\mathrm{RC}}$, and distance prior parameters $\theta_r$, the posterior probability of all the model parameters $\theta = (\theta_\varpi, \theta_{\mathrm{RC}}, \theta_r)$ is then proportional to the likelihood of the data $\mathbf{d}$, and the prior probability of the model parameters.

$$p(\theta|\mathbf{d}) \propto p(\mathbf{d}|\theta)p(\theta). \tag{2.4}$$

The likelihood of the entire red clump data set can be split into the parallax and magnitude likelihoods for each $i$-th star independently:

$$p(\mathbf{d}|\theta) = \prod_i p(\varpi_i|\theta_\varpi)p(m_i|\theta_{\mathrm{RC}}) \tag{2.5}$$

We model the parallax measurements of each red clump star in *Gaia* DR2 as being drawn from

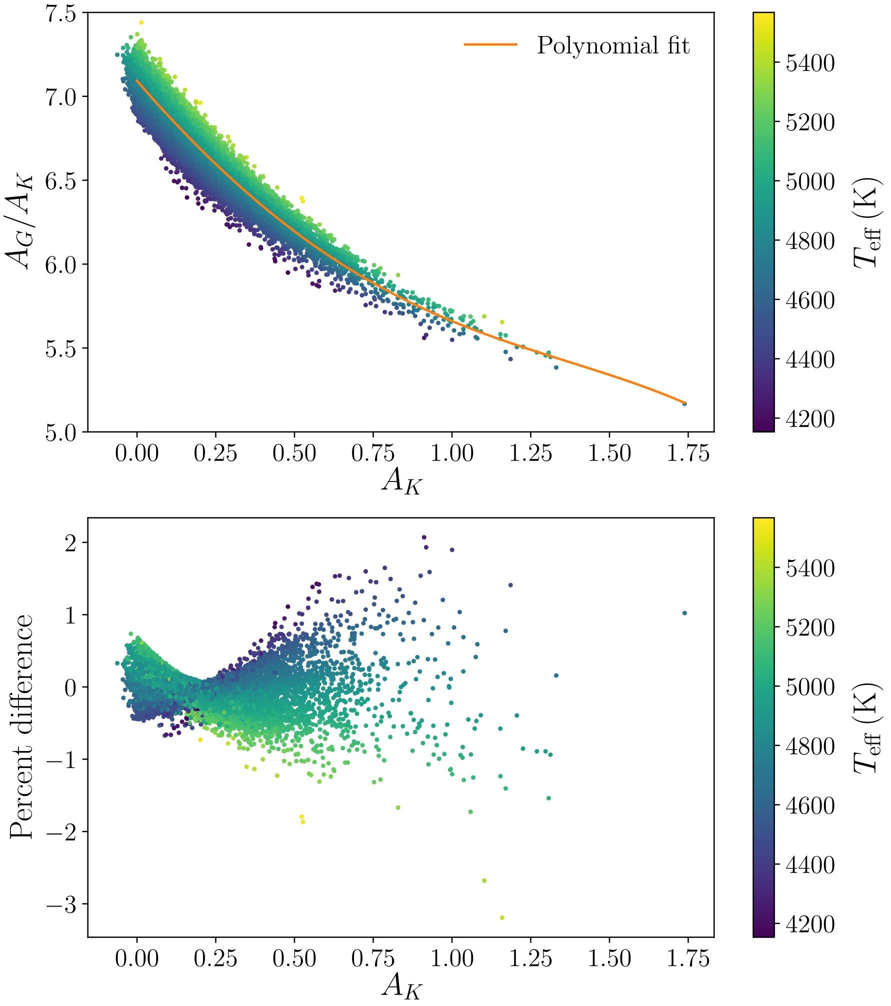

Figure 2.2: The stellar extinction ratios $A_G/A_K$ for the red clump sample. *Top:* The sample of extinction ratios using the steps from §2.2.3, along with the polynomial fit described by Equation (2.3). *Bottom:* The residuals from the polynomial fit. Equation (2.3) describes the sample of modeled extinction ratios to within a few percent.

a normal distribution. In other words, the likelihood of the DR2 parallax measurements is

$$p(\varpi_i|\theta_\varpi) \sim \mathcal{N}\big(\varpi_i\big|\varpi_i', \varsigma_{\varpi_i}^2\big), \text{ where} \tag{2.6}$$

$$\mathcal{N}\big(x\big|\mu_x, \sigma_x^2\big) = \frac{1}{\sqrt{2\pi\sigma_x^2}} e^{-\frac{(x-\mu_x)^2}{2\sigma_x^2}}, \tag{2.7}$$

$$\varpi_i' = 1/r_i + \varpi_0, \text{ and} \tag{2.8}$$

$$\varsigma_{\varpi_i}^2 = (f_\varpi \sigma_{\varpi_i})^2 + \sigma_{\varpi,+}^2. \tag{2.9}$$

The distribution for each star is centered around its "true observed parallax" $\varpi_i'$ set by Equation (2.8), which is the inverse of its true heliocentric distance $r_i$ summed with the systematic parallax zero point offset $\varpi_0$. This is the parallax that *Gaia* would observe for a source taking into account the parallax zero point in the limit of no other measurement uncertainties. We later consider the significance of a non-constant parallax zero point $\varpi_0(T_{\text{eff}}, G, \alpha, \delta)$ that is dependent on colour, magnitude, and sky position. The uncertainty for each parallax measurement is assumed to be Gaussian, with $\sigma_{\varpi_i}^2$ being the reported parallax uncertainty from *Gaia* DR2, but we allow for adjustments to the reported uncertainties to account for mis-estimated uncertainties. Equation (2.9) includes two error correction parameters $f_\varpi$, and $\sigma_{\varpi,+}^2$, which were used to inflate the Tycho-Gaia Astrometric Solution (TGAS) parallax uncertainties reported in *Gaia* DR1 (Lindegren et al., 2016). The values used in *Gaia* DR1 were $f_\varpi = 1.4$, and $\sigma_{\varpi,+} = 0.2$ mas. The reported values for *Gaia* DR2 are $f_\varpi = 1.08$, $\sigma_{\varpi,+}(G < 13) = 0.021$ mas, and $\sigma_{\varpi,+}(G > 13) = 0.043$ mas[5]. We include these parameters $\theta_\varpi = \{\varpi_0, f_\varpi, \sigma_{\varpi,+}\}$ in our model as a validation of the reported parallax uncertainties and corrections in *Gaia* DR2. It is also of note that the astrometric solutions for bright ($G < 13$) and dim ($G \geq 13$) sources are different (Arenou et al., 2018); thus, we also consider a few models with separate *Gaia* systematic parameters $\theta_\varpi = \{\varpi_0, f_\varpi, \sigma_{\varpi,+}\}$ for bright/dim sources.

The red clump requires a model for the luminosity as a function of intrinsic stellar properties such as metallicity, and colour. The absolute magnitude $M$ of the $i$-th red clump star can be described as

$$M_i = M_{\text{ref}} + \alpha\big([J_0 - K_0]_i - [J_0 - K_0]_{\text{ref}}\big) + \beta\big([\text{Fe/H}]_i - [\text{Fe/H}]_{\text{ref}}\big), \tag{2.10}$$

where $(J_0 - K_0)$ is the extinction corrected colour. $M_{\text{ref}}$ is a reference value parameter which represents the absolute magnitude of a typical red clump star. Both terms denoted with a subscript "ref" are *fixed* representative values of the red clump population. In every one of our models, we use the median of each corresponding property in the sample: $(J_0 - K_0)_{\text{ref}} = 0.60$ mag, and $[\text{Fe/H}]_{\text{ref}} = -0.12$ dex.

We then choose to model the absolute magnitude of each red clump star as being drawn from a Student's $t$-distribution centered around $M_{\text{ref}}$, or

$$p\Big(M_i\Big|[\text{Fe/H}]_i, (J_0 - K_0)_i\Big) \sim \mathcal{S}\Big(M_i\Big|M_{\text{ref}}, \sigma_M^2, \nu\Big), \text{ where} \tag{2.11}$$

$$\mathcal{S}(t_i|\nu) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\,\Gamma\left(\frac{\nu}{2}\right)}\left(1 + \frac{t_i^2}{\nu}\right)^{-\frac{\nu+1}{2}}, \text{ and} \tag{2.12}$$

$$t_i = \frac{M_i - M_{\text{ref}}}{\sigma_M}.$$

---

[5] https://www.cosmos.esa.int/documents/29201/1770596/Lindegren_GaiaDR2_Astrometry_extended.pdf

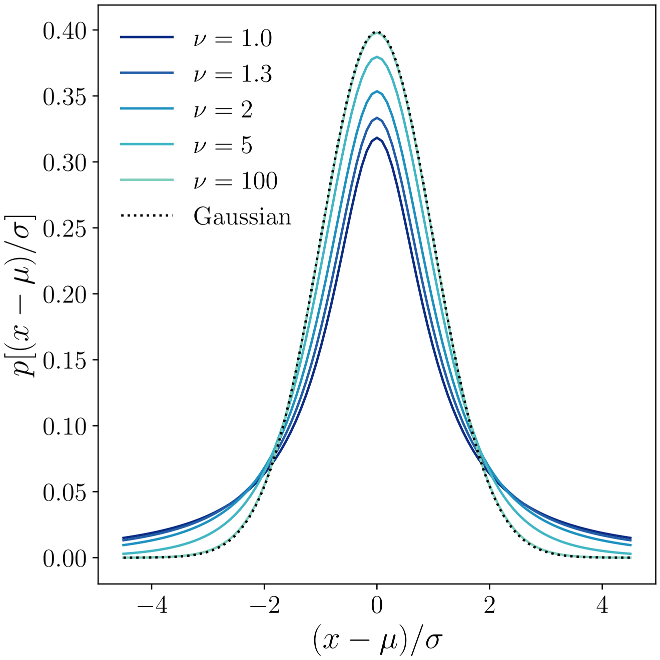

Figure 2.3: The Student's t-distribution in comparison with a Gaussian distribution for different values of $\nu$. The dashed line shows a standard normal distribution. For large values of $\nu$ the distribution is approximately Gaussian (the $\nu = 100$ distribution essentially overlaps the Gaussian curve), but for small values of $\nu$ the distribution has much heavier tails than a Gaussian. We use the Student's t-distribution as a flexible model for the luminosity function of the red clump.

Here, $\nu$ is a parameter that controls the shape of the distribution. The role of $\nu$ in the Student's t-distribution is illustrated in Figure 2.3. A lower value of $\nu$ introduces a positive excess kurtosis to the distribution; in other words, the peak at the central value $M_{\rm ref}$ gets narrower, and the tails become elevated. This allows the luminosity calibration to be less susceptible to outliers which still contaminate the red clump sample after quality cuts. The Student's $t$-distribution converges to a normal distribution as $\nu \to \infty$. The uncertainties associated with $m_0$, [Fe/H], $(J_0 - K_0)$, and the luminosity model chosen in Equation (2.10) are captured in $\sigma_M^2$ as a model parameter, which then sets the width of the distribution.

To tie the luminosity and parallax calibration models together, we unite the observed magnitude of each star with its observed parallax through the distance to the star, $r_i$. This was already done for the parallaxes through Equation (2.8), and can be done for the magnitudes through the distance modulus,

$$\mu_i = m_{0,i} - M_i = 5\log_{10} r_i - 5, \tag{2.13}$$

where $m_0$ is the extinction corrected apparent magnitude. When using the $G$ magnitudes in particular, we use the corrected $G$ magnitudes using Equation (2.2), and apply the extinction correction from our model. We then model the likelihood of the observed distance modulus as a Student's t-distribution centered around the true distance modulus, or

$$p\Big(\mu_i(m_i)\Big|\mu_i(r_i), \sigma_M^2\Big) = \mathcal{S}\Big(\mu_i(m_i)\Big|\mu_i(r_i), \sigma_M^2, \nu\Big). \tag{2.14}$$

The approximation from Equation (2.11) to Equation (2.14) is valid only if measurement uncertainties for $m_0$, $(J_0 - K_0)$ and [Fe/H] are sufficiently small compared to $\sigma_M^2$. We provide a discussion for this later on.

The true distance to each star is unknown, so we include each true distance $r_i$ as a model parameter under an exponentially decreasing volume density prior with scale distance $L$ as described by Bailer-Jones (2015):

$$p(r_i|L) = \frac{r_i^2}{2L^3}\exp(-r_i/L). \tag{2.15}$$

The likelihood for the entire data set $\mathbf{d}$ is then

$$p(\mathbf{d}|\theta) = \prod_i \mathcal{N}(\varpi_i|\varpi_i', \varsigma_{\varpi_i}^2)\mathcal{S}\Big(\mu_i(m_i)\Big|\mu_i(r_i), \sigma_M^2, \nu\Big)p(r_i|L). \tag{2.16}$$

A schematic of our probabilistic model illustrating the dependencies of each parameter and observed quantity is shown in Figure 2.4. Each of the single circled parameters require prior probabilities. The prior for $r$ is described by Equation (2.15). We assign the following broad priors for the remaining parameters:

- Zero point parallax:
  Uniform prior between $-100 < \varpi_0/\mu\text{as} < 100$
- *Gaia* parallax error scaling:
  Uniform prior between $0.2 < f_\varpi < 2$
- *Gaia* parallax error offset:
  Log-uniform prior with $0.1 < \sigma_{\varpi,+}/\mu\text{as}$ and no upper bound

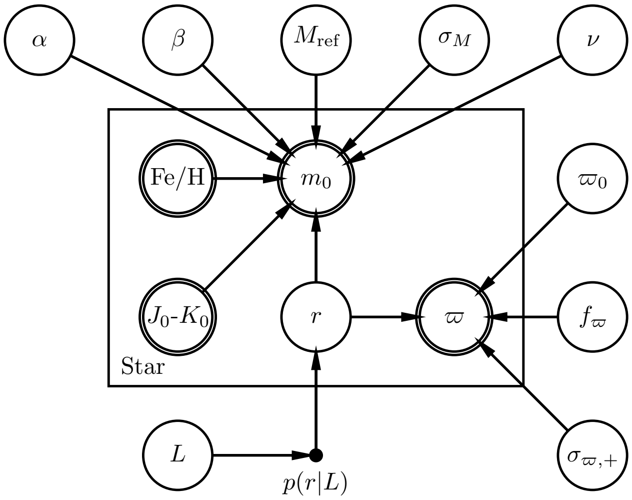

Figure 2.4: A probabilistic graphical model illustrating the base luminosity and parallax calibration for red clump stars. Double circled nodes indicate observed parameters (or likelihoods), and single circled nodes indicate fit parameters (which require priors). Nodes inside the rectangle are different for each star. The direction of each arrow indicates the conditional dependence of each parameter. For example, the arrow pointing from $r$ to $m_0$ indicates $p(m_0|r)$.

- Distance prior scale length for Equation (2.15):
  Uniform prior between $300 < L/\mathrm{pc} < 3000$
- Colour $(J_0 - K_0)$ slope for red clump luminosity:
  Uniform prior between $-5 < \alpha < 5$
- Metallicity [Fe/H] slope for red clump luminosity:
  Uniform prior between $-5 < \beta/(\mathrm{mag\ dex}^{-1}) < 5$
- Reference absolute magnitude for red clump stars in the $K_s$ band:
  Uniform prior between $-2 < M_{K,\mathrm{ref}}/\mathrm{mag} < -1$
- Reference absolute magnitude for red clump stars in the $J$ band:
  Uniform prior between $-2.5 < M_{J,\mathrm{ref}}/\mathrm{mag} < -1.5$
- Reference absolute magnitude for red clump stars in the $H$ band:
  Uniform prior between $-2 < M_{H,\mathrm{ref}}/\mathrm{mag} < -1$
- Reference absolute magnitude for red clump stars in the *Gaia* $G$ band:
  Uniform prior between $0 < M_{G,\mathrm{ref}}/\mathrm{mag} < 1$
- Spread in red clump luminosities in each photometric band:
  Log-uniform prior between $0.01 < \sigma_{M_m}/\mathrm{mag} < 0.8$
- Student's $t$-distribution degrees of freedom parameter for each photometric band:
  Log-uniform prior between $0 < \nu < 1000$

### 2.3.2 Multivariate Photometry Models

The *Gaia* parallax validation model described so far uses only one photometric measurement for each star. In general, $k$ apparent magnitudes may be included in a single analysis through the multivariate $t$-distribution

$$p\Big(\vec{\mu}_i(\vec{m}_i)\Big|\vec{\mu}_i(r_i), \mathbf{\Sigma}_m\Big) = \mathcal{S}\Big(\vec{\mu}_i(\vec{m}_i)\Big|\vec{\mu}_i(r_i), \mathbf{\Sigma}_M, \nu\Big), \text{ where} \tag{2.17}$$

$$\mathcal{S}\Big(\vec{x}\Big|\vec{\mu}, \mathbf{\Sigma}, \nu\Big) = \frac{\Gamma[(\nu+k)/2]}{\Gamma(\nu/2)\sqrt{\nu^k \pi^k |\mathbf{\Sigma}|}}\left[1 + \frac{1}{\nu}(\vec{x}-\vec{\mu})^T \mathbf{\Sigma}^{-1}(\vec{x}-\vec{\mu})\right]. \tag{2.18}$$

Each component of $\vec{\mu}_i(\vec{m}_i)$ can be a distance modulus computed with one of the photometric bands using the left-hand side of Equation (2.13), and each component of $\vec{\mu}_i(r_i)$ is the true distance modulus computed with the distance using the right-hand side of Equation (2.13). As with the single-variable Student's $t$, $\mathbf{\Sigma}$ is related to the width of the distribution like $\sigma_M$ while also taking into account covariances across each photometric band.

We implement the multivariate $t$-distribution to create a model which accepts both $K_s$ and $G$ band photometry simultaneously. This model should act as a check for the *Gaia* $G$-band extinction model discussed in §2.2.3. For any given star, photometric measurements in different bands are expected to be highly correlated, meaning that we do not expect to infer the model parameters to higher precision with the inclusion of multiple photometric information. Rather, this multivariate model serves as a method to validate the extinction estimates for $A_G$ described in §2.2.3, which are expected to be less accurate than the given $A_K$ values. We specifically choose the $K_s$ information to match with $G$ because it is expected to be the least extinguished as the reddest photometric band available.

We also consider a different model which is more intuitively comparable to the single photometry models. While the luminosity of the red clump in each band is expected to be highly correlated, the luminosity of each star is not expected to be as strongly correlated with the intrinsic colour. We can therefore model photometry in different bands (say, $G$ and $K_s$) by modeling the absolute magnitude distribution of one band (say, $K_s$) with one Student's t-distribution and the colour distribution (say, $G - K_s$) with another Student's t-distribution, with the latter given by

$$\begin{aligned} p(G_{0,i} - K_{0,i}|M_{G,i} - &M_{K,i}, \sigma_{GK}^2) \\ &= \mathcal{S}(G_{0,i} - K_{0,i}|M_{G,i} - M_{K,i}, \sigma_{GK}^2), \end{aligned} \tag{2.19}$$

where $M_{K,i}$ and $M_{G,i}$ are still modeled with separate versions of Equation (2.10). This version of the multiple photometry model is implemented by simply including Equation (2.19) in the likelihood set by Equation (2.16).

### 2.3.3 Variation of the *Gaia* Zero Point Parallax

We consider various ways in which the *Gaia* parallax zero point may depend on other quantities. To account for the differences in astrometric solutions for sources with $G < 13$ and $G \geq 13$, we first implement a single photometry model with separate parallax-related parameters $\theta_\varpi = \{\varpi_0, f_\varpi, \sigma_{\varpi,+}\}$ for each of the two cases. This and every other model discussed in this section uses only the $K_s$

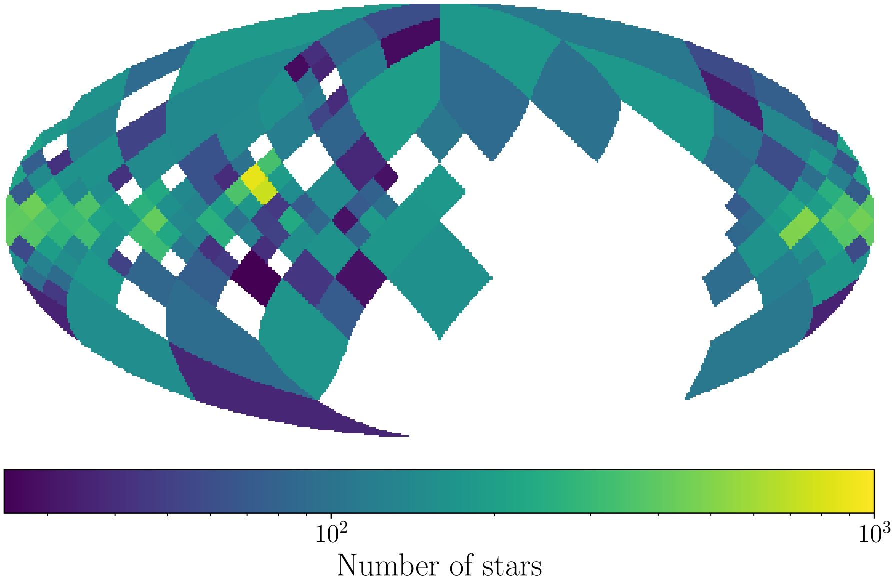

Figure 2.5: The `HEALPIX` projection of the sky distribution of the red clump sample divided into varying resolutions through conditions described in §2.3.3. We use this to determine the parallax zero point offset's variation on the sky.

photometric information to calibrate the red clump luminosity model.

As mentioned previously, the *Gaia* parallax zero point should exhibit continuous and multivariate dependences on properties such as the observed magnitude $G$, the observed colour $G_{BP} - G_{RP}$, and the position of the source on the sky. Note that we investigate dependencies of the parallax zero point with respect to the *uncorrected* (using neither Equation (2.2) nor the extinction corrections) $G$ magnitudes because we want to model how the DR2 pipeline leaves systematic residuals in the astrometry. The photometric/colour dependencies are easily included in the model by introducing a functional form for $\varpi_0$:

$$\varpi_0 = z(G, G_{BP} - G_{RP}) = z_0 + z_G(G) + z_c(G_{BP} - G_{RP}). \tag{2.20}$$

In particular, we investigate simple quadratic parameterizations for the dependences on both $G$, and $G_{BP} - G_{RP}$. Due to the change in the astrometric processing at $G = 13$, we also consider models with independent zero point functional parameterizations for sources $G < 13$ and $G \geq 13$.

We further consider models in which the *Gaia* parallax parameters $\theta_\varpi = \{\varpi_0, f_\varpi, \sigma_{\varpi,+}\}$ are not required to follow any specific functional parameterization, but rather we model them as separate constants in binned $G$ or $G_{BP} - G_{RP}$ space. In particular, we consider 17 bins of width 0.5 mag from 9.5 – 18 mag in $G$ space, and we consider 7 bins of width 0.5 mag from 1 – 4.5 mag in $G_{BP} - G_{RP}$ space. Stars outside of these ranges are discarded for each respective analysis. For each of these binned models, we also consider separate $L$ parameters for each bin to account for the distance dependence of $G$, or any possible 3D position clustering of red clump stars of similar observed colour.

The dependence on sky position can be probed by projecting the red clump sample onto `HEALPIX` maps (Górski et al., 2005). The `HEALPIX` framework divides the surface of a sphere into 12 diamond-shaped patches of equal solid angle at lowest order (called `NSIDE`, and beginning at 1). Higher resolution patches on the sphere are obtained by further dividing each patch into 4. Each division of map patches into 4 increases the `NSIDE` of the map by a factor of 2. The zero point dependence on sky position can then be modelled through unique values of $\varpi_0$, and $L$ for each `HEALPIX` patch on a sky map. In particular, we consider maps of order $\texttt{NSIDE} = \{1, 2, 4, 8\}$, corresponding to patches of approximately $\{3438, 860, 215, 54\}$ square degrees respectively. The equal-area property of each patch in a `HEALPIX` projection is useful for regularly sampling the variation of $\varpi_0$ across the sky, but the effectiveness of this method is restricted by the need for a sufficiently large subsample of red clump stars within each patch. Much of the red clump sample is located within the Galactic disk, so the most precise measurements of $\varpi_0$ on the sky will come from the Galactic disk.

In order to retrieve similar quality inferences of the parallax zero point across the sky, we have developed a method for analyzing a single sky map with varying `HEALPIX` resolutions dependent on the number of stars within each patch. `HEALPIX` patches are recursively split into higher resolution sub-patches if the parent patch contains greater than 200 stars. Each daughter sub-patch is further broken up if it still contains greater than 200 stars and is of order $\texttt{NSIDE} < 8$. If the sub-patch contains between $25 - 200$ stars or has reached $\texttt{NSIDE} = 8$, then it is kept at that resolution. If the sub-patch contains fewer than 25 stars, then it is discarded. The resulting star density map across our sample can be seen in Galactic coordinates in Figure 2.5. The highest resolution patches allow for detailed analysis of the zero point's variation along the Galactic plane, while the lower resolution patches should group enough stars away from the Galactic plane together to provide reasonable inferences of the parallax zero point. All parameters remain global with the exception of $\varpi_0$ and $L$, which are specific to each patch. Again, we also consider a model with independent sky variations of $\varpi_0$ for sources $G < 13$ and $G \geq 13$.

### 2.3.4 Red Clump Luminosity Calibration: Dependence on [$\alpha$/Fe]

To determine the red-clump luminosity function for sub-populations within the red clump sample, we separate the data into bins in [$\alpha$/Fe] (calculated with Equation (2.1) using APOGEE abundance data) space. The sample is split into 7 evenly sized bins from [$\alpha$/Fe] $= -0.05$ dex to 0.3 dex that are 0.05 dex wide. The chemical distribution of the entire sample can be seen in Figure 2.6 along with the boundaries of the [$\alpha$/Fe] bins described. The low-alpha ([$\alpha$/Fe] $\lesssim 0.15$ dex) and high-alpha ([$\alpha$/Fe] $\gtrsim 0.15$ dex) sequences can clearly be seen in abundance space.

To test possible differences in the luminosity calibration between red clump stars of varying chemical compositions, we alter the model described by Figure 2.4 to include independent red clump luminosity parameters $\theta_{\rm RC} = \{\alpha, \beta, M_{\rm ref}, \sigma_M, \nu\}$ for each [$\alpha$/Fe] bin while the remaining parameters are still relevant to the global sample. Keep in mind that we are still taking into account the red clump luminosity dependence on $(J_0 - K_0)$ and [Fe/H] separately for each [$\alpha$/Fe] bin through Equation (2.10). We expect small changes in the red clump luminosity calibration for each sub-population, which can be seen as variations of $\theta_{\rm RC}$ parameters in [$\alpha$/Fe] space.

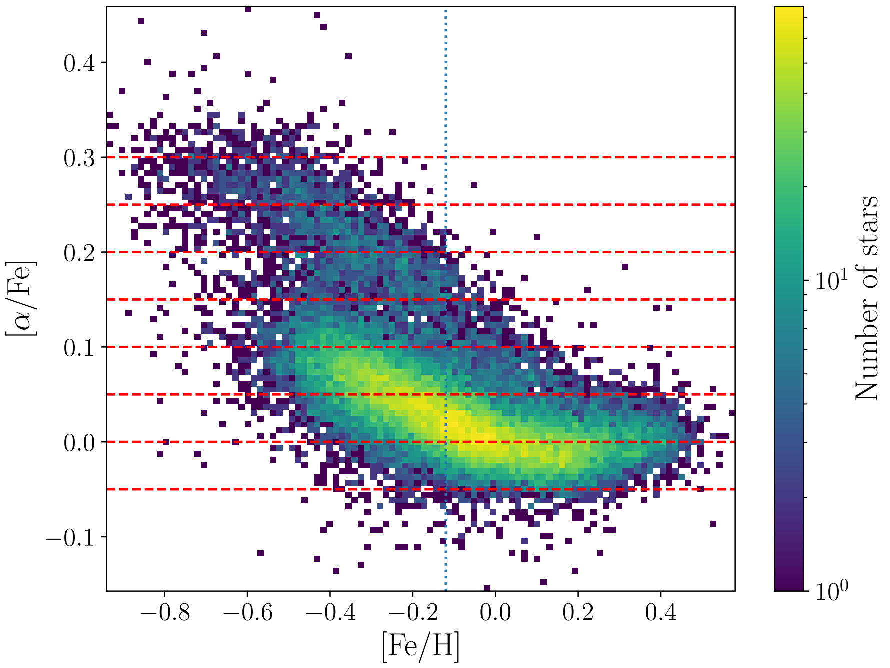

Figure 2.6: The distribution of red clump sample stars in chemical abundance space. The dotted blue line indicates the median metallicity [Fe/H] of the full sample. The boundaries of the [$\alpha$/Fe] bins described in §2.3.4 are shown as dashed red lines. These bins are used in the analysis of variations of the red clump luminosity calibration across sub-populations.

### 2.3.5 Tests of red clump stellar evolution models

Bovy et al. (2014) determined distances to the red clump in our sample using colour and metallicity trends determined from PARSEC stellar models (Bressan et al., 2012), applying a constant calibration offset obtained from a Hipparcos red clump sample. We test the stellar model used in Bovy et al. (2014) as follows: We adjust Equation (2.13) to

$$\mu_i = m_{0,i} - M_i', \tag{2.21}$$

where

$$M_i' = M_i\big([J_0 - K_0]_i, [\mathrm{Fe/H}]_i\big) + m_{0,i} - \mu_i^*, \tag{2.22}$$

and $M_i^* = m_{0,i} - \mu_i^*$ is the absolute magnitude of a red clump star predicted with the stellar model. The new observed distance modulus which will replace Equation (2.13) for this validation model is then

$$\mu_i = \mu_i^* - M_i. \tag{2.23}$$

Here, $M_i$ still represents a red clump luminosity calibrated with the probabilistic model, but it is not tied to any photometric band. Instead, it is related to the accuracy of the stellar models used to determine $\mu_i^*$ in Bovy et al. (2014). We expect $M_i \to 0$, with each parameter within Equation (2.10) now describing possible residuals of the stellar model; i.e, the parameters $M_{\mathrm{ref}}$, $\alpha$, and $\beta$ are expected to be 0 if the stellar models used in the catalogue describe the red clump at least as well as the empirical models in this paper.

### 2.3.6 Implementation of Models

A full probabilistic graphical model in Figure 2.7 illustrates each extension to the base model in Figure 2.4. The sheer number of parameters that need to be sampled for each model (especially the distances to $\sim 28000$ stars) hinders the use of traditional Markov chain Monte Carlo (MCMC) techniques. We choose instead to use `Stan`, a Hamiltonian Monte Carlo (HMC) software (Carpenter et al., 2017). Some advantages of using HMC instead of MCMC include a generally more efficient sampling of parameter space, especially in the case of highly correlated parameters. HMC also provides more accurate and precise sampling of parameter space for both models with complex distributions and models with many parameters. The principles and benefits of HMC are further described in great detail by Neal (2012).

Each of the models described in this paper was implemented in the `Python 3` wrapper for `Stan`, also known as `PyStan`[6]. The posterior distribution space for each model was sampled with 4 chains for 2500 steps, and the first 1000 steps of each chain were discarded as warm-up steps. In total, 6000 samples were collected for the posterior distributions of each model we report. The parameters were intialized randomly according to their respective priors with the exception of the distances. The true distance parameter for each star was initialized using the absolute inverse of its parallax reported in *Gaia* DR2.

## 2.4 Results

### 2.4.1 Basic results

The posterior distributions for all models are summarized in Table 2.1. We report the median and $\pm 1\sigma$ ranges for each model parameter. It is worthwhile to note that we also obtain posterior distributions for the distances to each individual star in our sample. Individual distances are typically constrained to $\sim 10\%$, which may seem insignificant, but the combined data set allows for the parallax zero point to be inferred to $\sim 1\%$ across every model.

We find that zero points inferred across each base model using $K_s$, $J$, and $H$ photometry are consistent ($\varpi_0 = -48 \pm 1$ $\mu$as), with an inconsistent result from the zero point inferred with the base model using $G$ photometry ($\varpi_0 = -38.30^{+0.88}_{-0.80}$ $\mu$as). This inconsistency is alleviated with the joint photometry model using the multivariate $t$-distribution to include $K_s$ and $G$ data simultaneously. The resulting zero point inferred with joint photometry is $\varpi_0 = -48.94^{+0.93}_{-0.96}$ $\mu$as. We also infer consistent *Gaia* parallax uncertainty correction parameters ($f_\varpi$ and $\sigma_{\varpi,+}$) across all base model and the joint photometry analyses. A representative posterior distribution is shown in Figure 2.8 along with the correlations between each model parameter in the base model using $K_s$ photometry.

We infer the absolute magnitude of the red clump to be $M_{\text{ref}} = -1.622 \pm 0.004$ in $K_s$, $M_{\text{ref}} = 0.447 \pm 0.005$ in $G$, $M_{\text{ref}} = -1.019 \pm 0.004$ in $J$, and $M_{\text{ref}} = -1.516 \pm 0.004$ in $H$. Red clump luminosity calibrations infer significant colour ($J_0 - K_0$) and metallicity ([Fe/H]) dependences in every photometric band considered. These dependences are implied again with the validation of reported APOGEE distances discussed later.

The Student's $t$-distribution appears to be able to capture outliers in the sample by widening the distribution about $\mu_i(m_i)$; the degrees-of-freedom parameter is inferred to be approximately $\nu \approx 1.3$

---

[6] Stan Development Team. 2018. PyStan: the Python interface to Stan, Version 2.17.1.0. `http://mc-stan.org`

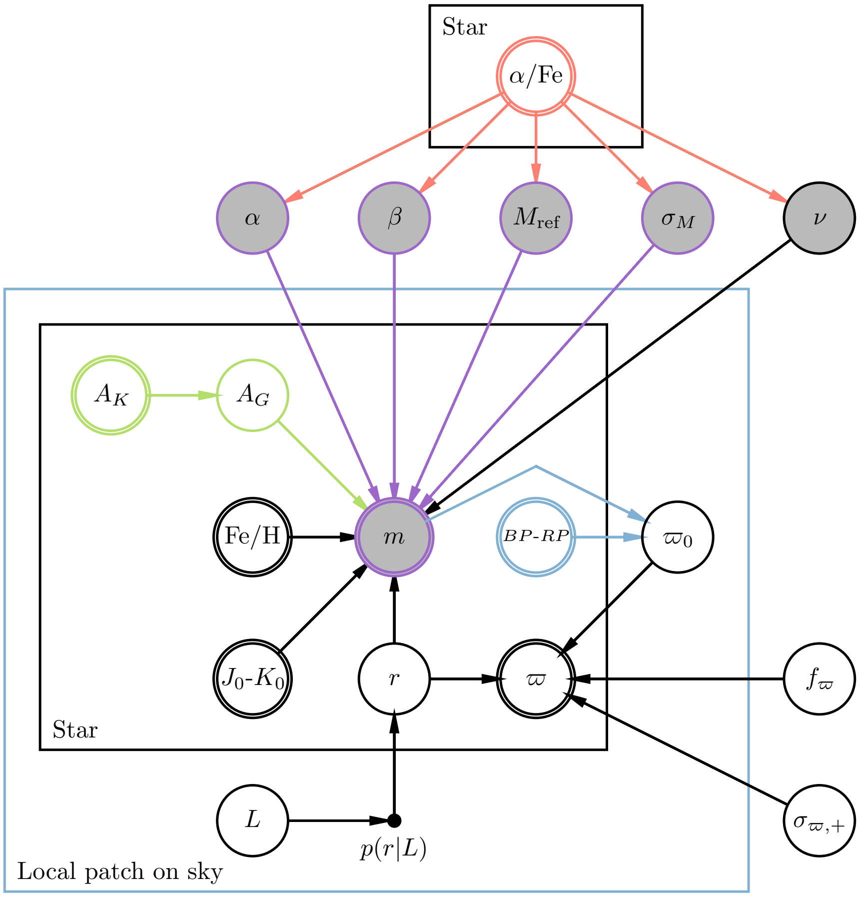

Figure 2.7: A probabilistic graphical model illustrating the extensions to the base luminosity and parallax calibration of red clump stars Figure 2.4. The colours illustrate an independent add-on to the existing base model shown in Figure 2.4. *Blue*: Adding functional dependencies to the parallax zero point $\varpi_0$, or modeling the spatial variations of $\varpi_0$ across the sky with `HEALPIX` patches. *Green*: Adding an extinction model for the *Gaia G* band photometry. *Red*: Adding a [$\alpha$/Fe] dependence of the red clump luminosity calibration set by Equation (2.10). *Purple*: Changing the red clump luminosity calibration to a validation of the distance moduli reported in the red clump catalogue (Bovy et al., 2014). *Gray fill*: Including multiple photometric bands simultaneously in the analysis. A unique copy of each of these parameters is added for each corresponding photometric band.

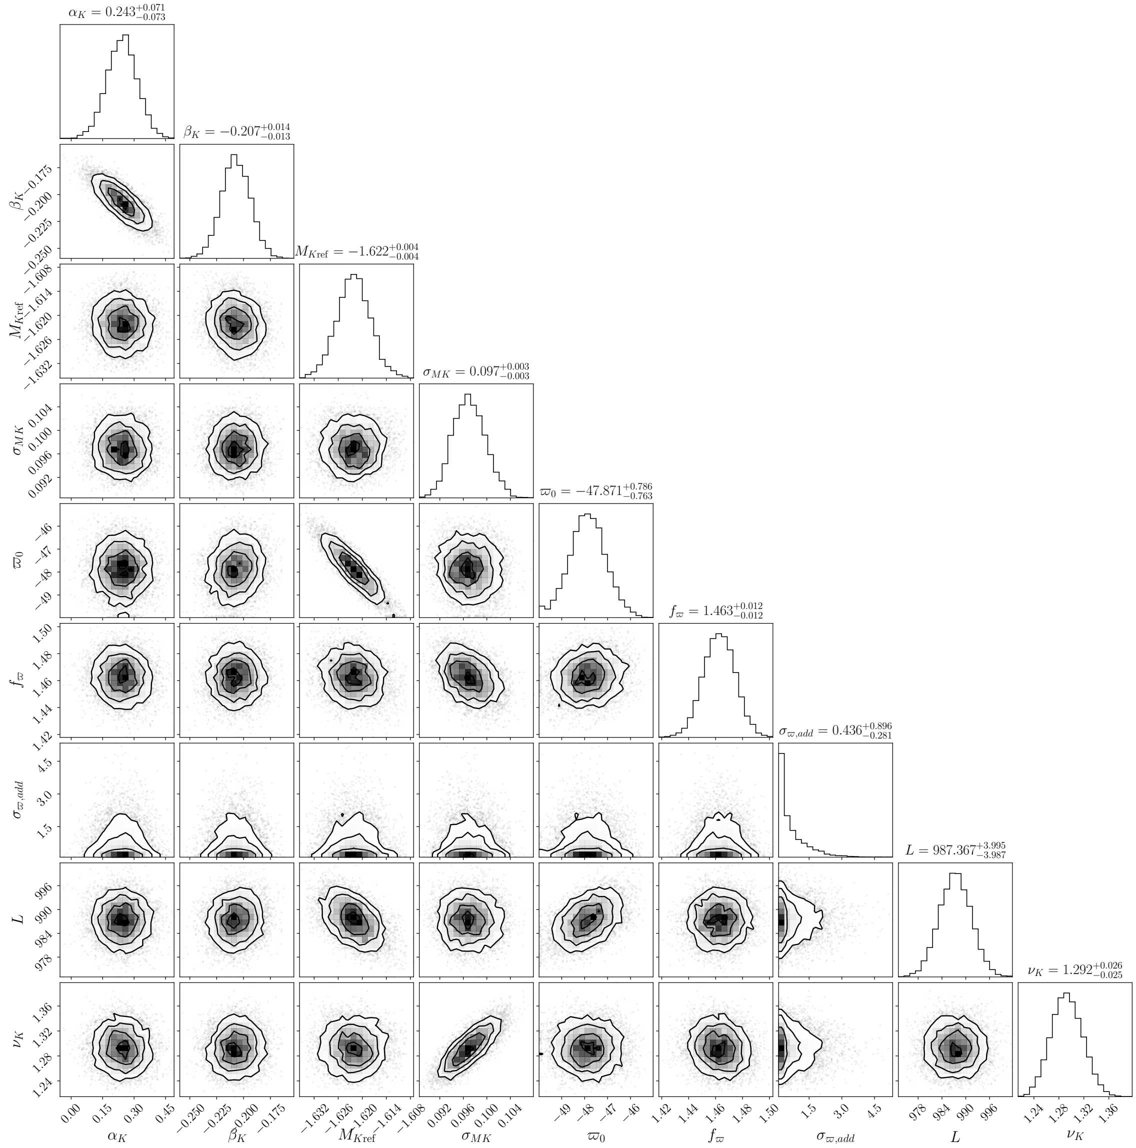

Figure 2.8: The posterior distributions of global parameters with respect to the $K_s$ data for constant $\varpi_0$. All parameters are precisely constrained by the data, and $\varpi_0$ is strongly correlated with the absolute magnitude $M_{K,\mathrm{ref}}$ and is weakly correlated with the distance prior parameter $L$. The parallax zero point has close to vanishing correlations with the remaining parameters.

Table 2.1: Inferred model parameters for all models discussed in this paper along with $\pm 1\sigma$ uncertainties. Adjacent rows from the same model/analysis are indicated with shared background highlighting.

| Photometric band | $\varpi_0$ | $f_\varpi$ | $\sigma_{\varpi,+}$ | $M_{\rm ref}$ | $\sigma_M$ | $\alpha$ | $\beta$ | $\nu$ | FWHM/2.355 | $L$ |
|---|---|---|---|---|---|---|---|---|---|---|
| and model | [$\mu$as] | – | [$\mu$as] | [mag] | [mag] | – | [mag/dex] | – | [mag] | [pc] |
| $K_s$ (Base) | $-47.87^{+0.79}_{-0.76}$ | $1.46^{+0.01}_{-0.01}$ | $0.44^{+0.90}_{-0.28}$ | $-1.622^{+0.004}_{-0.004}$ | $0.097^{+0.003}_{-0.003}$ | $0.24^{+0.07}_{-0.07}$ | $-0.21^{+0.01}_{-0.01}$ | $1.29^{+0.03}_{-0.03}$ | $0.088^{+0.003}_{-0.003}$ | $987^{+4}_{-4}$ |
| $G$ (Base) | $-39.91^{+0.81}_{-0.79}$ | $1.49^{+0.01}_{-0.01}$ | $0.56^{+1.25}_{-0.39}$ | $0.435^{+0.004}_{-0.005}$ | $0.123^{+0.003}_{-0.003}$ | $2.93^{+0.08}_{-0.08}$ | $-0.11^{+0.02}_{-0.02}$ | $1.48^{+0.03}_{-0.03}$ | $0.116^{+0.004}_{-0.004}$ | $1024^{+4}_{-4}$ |
| $J$ (Base) | $-47.91^{+0.79}_{-0.76}$ | $1.46^{+0.01}_{-0.01}$ | $0.41^{+0.87}_{-0.25}$ | $-1.019^{+0.004}_{-0.004}$ | $0.097^{+0.003}_{-0.003}$ | $1.25^{+0.08}_{-0.07}$ | $-0.21^{+0.01}_{-0.01}$ | $1.29^{+0.03}_{-0.03}$ | $0.088^{+0.003}_{-0.003}$ | $987^{+4}_{-4}$ |
| $H$ (Base) | $-49.69^{+0.82}_{-0.81}$ | $1.46^{+0.01}_{-0.01}$ | $0.41^{+0.85}_{-0.25}$ | $-1.516^{+0.004}_{-0.004}$ | $0.098^{+0.003}_{-0.003}$ | $0.50^{+0.07}_{-0.08}$ | $-0.19^{+0.01}_{-0.01}$ | $1.30^{+0.03}_{-0.03}$ | $0.089^{+0.004}_{-0.004}$ | $981^{+4}_{-4}$ |
| $K_s$ (Base; Fixed $f_\varpi$, $\sigma_{\varpi,+}$) | $-50.25^{+0.86}_{-0.85}$ | 1.08 | $21, G<13$<br>$43, G>13$ | $-1.613^{+0.004}_{-0.004}$ | $0.093^{+0.002}_{-0.002}$ | $0.28^{+0.08}_{-0.07}$ | $-0.22^{+0.01}_{-0.01}$ | $1.27^{+0.03}_{-0.02}$ | $0.084^{+0.003}_{-0.003}$ | $983^{+4}_{-4}$ |
| $K_s$ (Joint $K_s$ & $G$; Multi-$t$) | $-48.76^{+0.94}_{-0.98}$ | $1.45^{+0.01}_{-0.01}$ | $0.41^{+0.75}_{-0.25}$ | $-1.628^{+0.006}_{-0.005}$ | ${}^a0.041^{+0.001}_{-0.001}$ | $-0.01^{+0.09}_{-0.09}$ | $-0.19^{+0.02}_{-0.02}$ | $2.69^{+0.05}_{-0.04}$ | ${}^a0.061^{+0.001}_{-0.001}$ | $984^{+4}_{-4}$ |
| $G$ (Joint $K_s$ & $G$; Multi-$t$) | $-48.76^{+0.94}_{-0.98}$ | $1.45^{+0.01}_{-0.01}$ | $0.41^{+0.75}_{-0.25}$ | $0.484^{+0.006}_{-0.005}$ | ${}^a0.049^{+0.001}_{-0.001}$ | $2.75^{+0.09}_{-0.09}$ | $-0.11^{+0.02}_{-0.02}$ | $2.69^{+0.05}_{-0.04}$ | ${}^a0.066^{+0.001}_{-0.001}$ | $984^{+4}_{-4}$ |
| $K_s$ (Joint $K_s$ & $G$; $G-K_s$) | $-47.94^{+0.87}_{-0.81}$ | $1.46^{+0.01}_{-0.01}$ | $0.42^{+0.88}_{-0.27}$ | $-1.622^{+0.004}_{-0.004}$ | $0.219^{+0.003}_{-0.003}$ | $0.25^{+0.07}_{-0.07}$ | $-0.21^{+0.01}_{-0.01}$ | $1.29^{+0.03}_{-0.03}$ | $0.066^{+0.061}_{-0.061}$ | $987^{+4}_{-4}$ |
| $G$ (Joint $K_s$ & $G$; $G-K_s$) | $-47.94^{+0.87}_{-0.81}$ | $1.46^{+0.01}_{-0.01}$ | $0.42^{+0.88}_{-0.27}$ | $0.494^{+0.004}_{-0.004}$ | $0.156^{+0.001}_{-0.001}$ | $3.04^{+0.08}_{-0.08}$ | $-0.12^{+0.01}_{-0.01}$ | $2.54^{+0.03}_{-0.03}$ | $0.162^{+0.002}_{-0.002}$ | $987^{+4}_{-4}$ |
| $K_s$ (Joint $K_s$ & $J$; $J-K_s$) | $-47.89^{+0.82}_{-0.83}$ | $1.46^{+0.01}_{-0.01}$ | $0.42^{+0.89}_{-0.27}$ | $-1.622^{+0.004}_{-0.004}$ | $0.097^{+0.003}_{-0.003}$ | $0.25^{+0.07}_{-0.08}$ | $-0.21^{+0.01}_{-0.01}$ | $1.29^{+0.03}_{-0.03}$ | $0.088^{+0.003}_{-0.003}$ | $987^{+4}_{-4}$ |
| $J$ (Joint $K_s$ & $J$; $J-K_s$) | $-47.89^{+0.82}_{-0.83}$ | $1.46^{+0.01}_{-0.01}$ | $0.42^{+0.89}_{-0.27}$ | $-1.020^{+0.004}_{-0.004}$ | $0.098^{+0.003}_{-0.003}$ | $1.25^{+0.07}_{-0.08}$ | $-0.21^{+0.01}_{-0.01}$ | $1.51^{+0.03}_{-0.03}$ | $0.093^{+0.001}_{-0.001}$ | $987^{+4}_{-4}$ |
| $K_s$ ($G<13$) | $-35.74^{+1.55}_{-1.53}$ | $1.23^{+0.03}_{-0.03}$ | $0.51^{+1.16}_{-0.34}$ | $-1.665^{+0.006}_{-0.006}$ | $0.103^{+0.003}_{-0.003}$ | $0.34^{+0.07}_{-0.08}$ | $-0.20^{+0.01}_{-0.01}$ | $1.34^{+0.03}_{-0.03}$ | $0.095^{+0.004}_{-0.004}$ | $581^{+4}_{-4}$ |
| $K_s$ ($G\geq 13$) | $-42.36^{+0.84}_{-0.86}$ | $1.52^{+0.01}_{-0.01}$ | $0.49^{+1.12}_{-0.32}$ | $-1.665^{+0.006}_{-0.006}$ | $0.103^{+0.003}_{-0.003}$ | $0.34^{+0.07}_{-0.08}$ | $-0.20^{+0.01}_{-0.01}$ | $1.34^{+0.03}_{-0.03}$ | $0.095^{+0.004}_{-0.004}$ | $1262^{+7}_{-6}$ |
| $K_s$ ($G<13$; Fixed $f_\varpi$, $\sigma_{\varpi,+}$) | $-35.39^{+1.57}_{-1.54}$ | 1.08 | 21 | $-1.666^{+0.006}_{-0.005}$ | $0.094^{+0.003}_{-0.003}$ | $0.32^{+0.08}_{-0.08}$ | $-0.20^{+0.01}_{-0.01}$ | $1.28^{+0.03}_{-0.03}$ | $0.086^{+0.003}_{-0.003}$ | $581^{+4}_{-4}$ |
| $K_s$ ($G\geq 13$; Fixed $f_\varpi$, $\sigma_{\varpi,+}$) | $-44.75^{+0.87}_{-0.86}$ | 1.08 | 43 | $-1.666^{+0.006}_{-0.005}$ | $0.094^{+0.003}_{-0.003}$ | $0.32^{+0.08}_{-0.08}$ | $-0.20^{+0.01}_{-0.01}$ | $1.28^{+0.03}_{-0.03}$ | $0.086^{+0.003}_{-0.003}$ | $1261^{+6}_{-6}$ |
| $K_s$ ($G<13$; Fixed $M_{\rm ref}$) | $-46.68^{+0.63}_{-0.63}$ | $1.20^{+0.03}_{-0.02}$ | $0.52^{+1.13}_{-0.35}$ | $-1.622$ | $0.104^{+0.003}_{-0.003}$ | $0.33^{+0.07}_{-0.07}$ | $-0.21^{+0.01}_{-0.01}$ | $1.34^{+0.03}_{-0.03}$ | $0.096^{+0.004}_{-0.004}$ | $569^{+3}_{-3}$ |
| $K_s$ ($G>13$; Fixed $M_{\rm ref}$) | $-48.17^{+0.46}_{-0.44}$ | $1.51^{+0.01}_{-0.01}$ | $0.46^{+1.02}_{-0.30}$ | $-1.622$ | $0.104^{+0.003}_{-0.003}$ | $0.33^{+0.07}_{-0.07}$ | $-0.21^{+0.01}_{-0.01}$ | $1.34^{+0.03}_{-0.03}$ | $0.096^{+0.004}_{-0.004}$ | $1236^{+5}_{-5}$ |
| $K_s$ ($G$ dep.; All $G$; Fixed $M_{\rm ref}$) | ${}^b-60.91$ | $1.42^{+0.01}_{-0.01}$ | $0.41^{+0.80}_{-0.25}$ | $-1.622$ | $0.104^{+0.003}_{-0.003}$ | $0.15^{+0.07}_{-0.07}$ | $-0.21^{+0.01}_{-0.01}$ | $1.31^{+0.03}_{-0.02}$ | $0.095^{+0.004}_{-0.004}$ | $986^{+3}_{-3}$ |
| $K_s$ ($G$ dep.; $G<13$; Free $M_{\rm ref}$) | ${}^b72.51$ | $1.31^{+0.03}_{-0.02}$ | $0.52^{+1.21}_{-0.35}$ | $-1.943^{+0.013}_{-0.013}$ | $0.111^{+0.003}_{-0.003}$ | $0.04^{+0.08}_{-0.08}$ | $-0.12^{+0.02}_{-0.02}$ | $1.38^{+0.03}_{-0.03}$ | $0.103^{+0.004}_{-0.004}$ | $660^{+5}_{-6}$ |

Continued on next page

Table 2.1 – Continued from previous page

| Photometric band | $\varpi_0$ | $f_\varpi$ | $\sigma_{\varpi,+}$ | $M_{\mathrm{ref}}$ | $\sigma_M$ | $\alpha$ | $\beta$ | $\nu$ | FWHM/2.355 | $L$ |
|---|---|---|---|---|---|---|---|---|---|---|
| and model | [$\mu$as] | – | [$\mu$as] | [mag] | [mag] | – | [mag/dex] | – | [mag] | [pc] |
| $K_s$ ($G$ dep.; $G \geq 13$; Free $M_{\mathrm{ref}}$) | [b] $-24.31$ | $1.53^{+0.01}_{-0.01}$ | $0.56^{+1.54}_{-0.38}$ | $-1.943^{+0.013}_{-0.013}$ | $0.111^{+0.003}_{-0.003}$ | $0.04^{+0.08}_{-0.08}$ | $-0.12^{+0.02}_{-0.02}$ | $1.38^{+0.03}_{-0.03}$ | $0.103^{+0.004}_{-0.004}$ | $1433^{+11}_{-11}$ |
| $K_s$ ($G$ dep.; $G < 13$; Fixed $M_{\mathrm{ref}}$) | [b] $-42.93$ | $1.20^{+0.03}_{-0.03}$ | $0.50^{+1.14}_{-0.33}$ | $-1.622$ | $0.106^{+0.003}_{-0.003}$ | $0.30^{+0.07}_{-0.07}$ | $-0.21^{+0.01}_{-0.01}$ | $1.34^{+0.03}_{-0.03}$ | $0.098^{+0.004}_{-0.004}$ | $569^{+3}_{-3}$ |
| $K_s$ ($G$ dep.; $G \geq 13$; Fixed $M_{\mathrm{ref}}$) | [b] $-60.90$ | $1.49^{+0.01}_{-0.01}$ | $0.46^{+1.05}_{-0.29}$ | $-1.622$ | $0.106^{+0.003}_{-0.003}$ | $0.30^{+0.07}_{-0.07}$ | $-0.21^{+0.01}_{-0.01}$ | $1.34^{+0.03}_{-0.03}$ | $0.098^{+0.004}_{-0.004}$ | $1236^{+5}_{-6}$ |
| $K_s$ ($G$ dep.; $G < 13$; $M_{\mathrm{ref}}$ prior) | [b] $-29.20$ | $1.20^{+0.02}_{-0.03}$ | $0.52^{+1.17}_{-0.34}$ | $-1.658^{+0.003}_{-0.003}$ | $0.106^{+0.003}_{-0.003}$ | $0.30^{+0.08}_{-0.08}$ | $-0.21^{+0.01}_{-0.01}$ | $1.34^{+0.03}_{-0.03}$ | $0.098^{+0.004}_{-0.004}$ | $579^{+4}_{-4}$ |
| $K_s$ ($G$ dep.; $G \geq 13$; $M_{\mathrm{ref}}$ prior) | [b] $-56.58$ | $1.50^{+0.01}_{-0.01}$ | $0.48^{+1.04}_{-0.32}$ | $-1.658^{+0.003}_{-0.003}$ | $0.106^{+0.003}_{-0.003}$ | $0.30^{+0.08}_{-0.08}$ | $-0.21^{+0.01}_{-0.01}$ | $1.34^{+0.03}_{-0.03}$ | $0.098^{+0.004}_{-0.004}$ | $1256^{+6}_{-6}$ |
| $K_s$ ($G$ dep.; Binned; Fixed $M_{\mathrm{ref}}$) | [c] $-47.60^{+0.37}_{-0.37}$ | [c] $1.39^{+0.01}_{-0.01}$ | [c] $1.56^{+0.59}_{-0.59}$ | $-1.622$ | $0.108^{+0.003}_{-0.003}$ | $0.31^{+0.07}_{-0.08}$ | $-0.21^{+0.01}_{-0.01}$ | $1.35^{+0.03}_{-0.03}$ | $0.100^{+0.004}_{-0.004}$ | [c] $603^{+2}_{-2}$ |
| $K_s$ ($BP-RP$ dep.; All $G$) | [b] $-100.33$ | $1.43^{+0.01}_{-0.01}$ | $0.44^{+0.97}_{-0.28}$ | $-1.651^{+0.005}_{-0.004}$ | $0.101^{+0.003}_{-0.003}$ | $0.13^{+0.07}_{-0.07}$ | $-0.13^{+0.01}_{-0.01}$ | $1.34^{+0.03}_{-0.03}$ | $0.093^{+0.004}_{-0.004}$ | $998^{+4}_{-4}$ |
| $K_s$ ($BP-RP$ dep.; $G < 13$) | [b] $-204.21$ | $1.26^{+0.03}_{-0.03}$ | $0.48^{+1.11}_{-0.32}$ | $-1.658^{+0.006}_{-0.006}$ | $0.105^{+0.003}_{-0.003}$ | $0.16^{+0.07}_{-0.08}$ | $-0.08^{+0.02}_{-0.02}$ | $1.41^{+0.03}_{-0.03}$ | $0.098^{+0.004}_{-0.004}$ | $574^{+4}_{-4}$ |
| $K_s$ ($BP-RP$ dep.; $G \geq 13$) | [b] $-106.98$ | $1.50^{+0.01}_{-0.01}$ | $0.51^{+1.16}_{-0.34}$ | $-1.657^{+0.006}_{-0.006}$ | $0.105^{+0.003}_{-0.003}$ | $0.16^{+0.07}_{-0.08}$ | $-0.08^{+0.02}_{-0.02}$ | $1.41^{+0.03}_{-0.03}$ | $0.098^{+0.004}_{-0.004}$ | $1256^{+7}_{-6}$ |
| $K_s$ ($BP-RP$ dep.; Binned) | [c] $-49.40^{+0.58}_{-0.58}$ | [c] $1.31^{+0.01}_{-0.01}$ | [c] $1.52^{+0.99}_{-0.99}$ | $-1.634^{+0.004}_{-0.004}$ | $0.106^{+0.003}_{-0.003}$ | $0.26^{+0.07}_{-0.07}$ | $-0.19^{+0.01}_{-0.01}$ | $1.37^{+0.03}_{-0.03}$ | $0.098^{+0.004}_{-0.004}$ | [c] $953^{+4}_{-4}$ |
| $K_s$ (Sky dep.; All $G$) | [c] $-43.06^{+0.37}_{-0.37}$ | $1.44^{+0.01}_{-0.01}$ | $0.00^{+0.00}_{-0.00}$ | $-1.650^{+0.004}_{-0.005}$ | $0.095^{+0.003}_{-0.003}$ | $-0.08^{+0.07}_{-0.08}$ | $-0.14^{+0.01}_{-0.01}$ | $1.29^{+0.03}_{-0.03}$ | $0.087^{+0.003}_{-0.003}$ | [c] $891^{+3}_{-3}$ |
| $K_s$ (Sky dep.; $G < 13$) | [c] $-34.92^{+0.69}_{-0.69}$ | $1.23^{+0.03}_{-0.03}$ | $0.00^{+0.00}_{-0.00}$ | $-1.668^{+0.006}_{-0.005}$ | $0.101^{+0.003}_{-0.003}$ | $-0.17^{+0.07}_{-0.08}$ | $-0.09^{+0.01}_{-0.01}$ | $1.37^{+0.03}_{-0.03}$ | $0.093^{+0.004}_{-0.004}$ | [c] $571^{+3}_{-3}$ |
| $K_s$ (Sky dep.; $G \geq 13$) | [c] $-42.04^{+0.46}_{-0.46}$ | $1.51^{+0.01}_{-0.01}$ | $0.00^{+0.00}_{-0.00}$ | $-1.668^{+0.006}_{-0.005}$ | $0.101^{+0.003}_{-0.003}$ | $-0.17^{+0.07}_{-0.08}$ | $-0.09^{+0.01}_{-0.01}$ | $1.37^{+0.03}_{-0.03}$ | $0.093^{+0.004}_{-0.004}$ | [c] $1229^{+6}_{-6}$ |
| $K_s$ ([$\alpha$/Fe] bins) | $-48.60^{+0.77}_{-0.80}$ | $1.46^{+0.01}_{-0.01}$ | $0.49^{+1.04}_{-0.32}$ | [c] $-1.629^{+0.003}_{-0.003}$ | [c] $0.086^{+0.003}_{-0.003}$ | [c] $-0.18^{+0.08}_{-0.08}$ | [c] $0.01^{+0.02}_{-0.02}$ | [c] $1.29^{+0.03}_{-0.03}$ | $0.078^{+0.003}_{-0.003}$ | $985^{+4}_{-4}$ |
| $\mu$ (Stellar model validation) | $-47.21^{+0.80}_{-0.85}$ | $1.47^{+0.01}_{-0.01}$ | $0.44^{+0.89}_{-0.28}$ | [d] $-0.036^{+0.004}_{-0.004}$ | [d] $0.105^{+0.003}_{-0.003}$ | [d] $1.48^{+0.07}_{-0.07}$ | [d] $-0.39^{+0.01}_{-0.01}$ | [d] $1.33^{+0.03}_{-0.03}$ | $0.096^{+0.004}_{-0.004}$ | $990^{+4}_{-4}$ |

[a] Diagonal components of $\mathbf{\Sigma}$ reported here. $G$ and $K_s$ share a covariance $\sigma_{GK} = 0.041^{+0.001}_{-0.001}$. FWHM are computed in one dimension along the zero point of the other band.

[b] Reported values are integrated means of the modeled/inferred functional forms of each dependence, which are further described in §2.4.2.

[c] Reported values are computed inverse variance weighted means of each all inferred binned distributions. See Figure 2.10 and Figure 2.11 for the binned distributions.

[d] Parameters reported here are not the same red clump luminosity calibration parameters, and a detailed interpretation is discussed in §2.4.3.

across all models. Consequently, the scaling parameter $\sigma_M$ does not fully describe the dispersion of stars about $M_{\mathrm{ref}}$. The Student's $t$-distribution has a formal variance of $\infty$ for $1 < \nu < 2$, but we wish to describe the width of the Student's $t$-distribution with a combination of $\sigma_M$ and $\nu$. We are only interested in the dispersion of stars that have not been considered outliers (which exist in the heightened tails of the distribution), so we report dispersions about the red clump luminosity calibration given by the full width at half maximum (FHWM) of each inferred distribution. Table 2.1 includes the FWHM$/2\sqrt{2\ln(2)}$ for each red clump luminosity distribution, which is equivalent to $1\sigma$ for a Gaussian distribution.

Note that the spread in the red clump luminosity distribution describes the spread in distance moduli as in Equation (2.14). This includes uncertainties from photometric ($\sigma_m \lesssim 0.025$ mag) and metallicity ($\sigma_{[\mathrm{Fe/H}]} \lesssim 0.01$ dex) measurements; however, these uncertainties are much smaller than the spread in the distributions (FWHM$/2.355 \sim 0.1$ mag). Taking this into account, the dispersion in the red clump distributions should be dominated by a combination of intrinsic dispersion and modeling uncertainties.

The inconsistencies in inferred *Gaia* parallax zero points between $G$ and 2MASS photometry appears to be resolved with the implementation of the multiple photometry models. The inferred absolute magnitude of the red clump in $G$ changes from $M_G = 0.447 \pm 0.004$ to $M_G = 0.508 \pm 0.005$ in the multivariate $t$ model, and to $M_G = 0.517 \pm 0.004$ in the model with separate single-variable $t$-distributions for $K_s$ and $G - K_s$. As a result, the inferred *Gaia* parallax zero point becomes consistent with the previous estimates of $\varpi_0 \approx -49\mu$as. This is likely an indication that our model for the $G$ band extinction is not accurate enough for an independent analysis of the *Gaia* parallax zero point using only $G$ photometry. Only by supplementing the model with $K_s$ photometry are we able to achieve consistent results.

In the case of considering 2 constant zero points for $G < 13$ and $G \geq 13$, we infer $\varpi_0 = -35.74 \pm 1.55$ $\mu$as for sources $G < 13$ and $\varpi_0 = -42.36 \pm 0.86$ $\mu$as for sources $G \geq 13$. We also find once again that the *Gaia* parallax uncertainty parameters are discrepant from those reported, although we see a larger $f_\varpi$ inferred for dim sources compared to bright sources. This is similar to the reported values of $\sigma_{\varpi,+}$. In an identical analysis with the exception of fixing $f_\varpi$ and $\sigma_{\varpi,+}$ to the reported values, we find very similar posterior distributions, implying that our probabilistic model prefers inflating *Gaia* parallax errors with $f_\varpi$ rather than $\sigma_{\varpi,+}$ to similar effect. We support this by running the base model analysis using $K_s$ photometry, but adding a gamma distribution prior on $\sigma_{\varpi,+}$:

$$p(\sigma_{\varpi,+}|u,v) = \frac{v^u {\sigma_{\varpi,+}}^{u-1} e^{-v\sigma_{\varpi,+}}}{\Gamma(u)}, \tag{2.24}$$

where $u > 0$ and $1/(10\ \mu\mathrm{as}) < v < 1/(0.1\ \mu\mathrm{as})$ are hyperparameters with the described priors. In particular, we find $v$, the 'rate parameter', to hug the upper bound of the prior $1/(0.2^{+0.5}_{0.1}\ \mu\mathrm{as})$. This in turn causes $\sigma_{\varpi,+}$ to again be consistent with 0 ($\sigma_{\varpi,+} = 1.2^{+1.2}_{-0.8}\ \mu$as). All remaining parameters in this model remain consistent with the original model.

Finally, we repeat the above analysis with free $f_\varpi$ and $\sigma_{\varpi,+}$, but we fix $M_{\mathrm{ref}} = -1.622$. Justification for this analysis is presented in the following section. We infer more consistent values of $\varpi_0 = -46.68 \pm 0.63$ $\mu$as for $G < 13$ and $\varpi_0 = -48.17 \pm 0.46$ $\mu$as for $G \geq 13$.

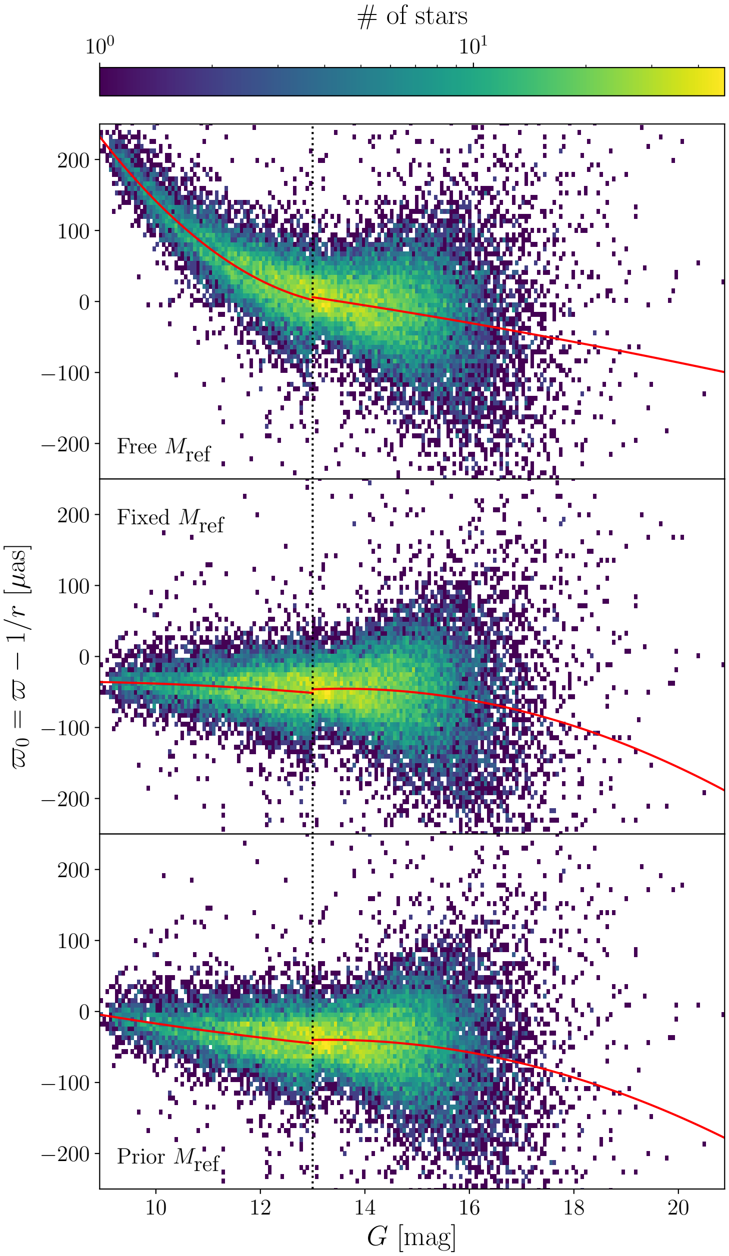

Figure 2.9: Inferred parallax zero point variations $G$. Individual estimates of the zero point for each star are shown by solving Equation 2.8 for $\varpi_0$, using the mean of the posterior samples of the distances to each star as well as its measured parallax. Red lines show the resulting quadratic parameterization of the zero point variations in $G < 13$ and $G \geq 13$. The model with a liberal $M_{K,\mathrm{ref}}$ prior appears to infer less inaccurate distances and parallax zero point offset than the other two models with more conservative priors on $M_{K,\mathrm{ref}}$. See Figure 2.1 for a comparison using distances from the APOGEE red clump catalogue Bovy et al. (2014).

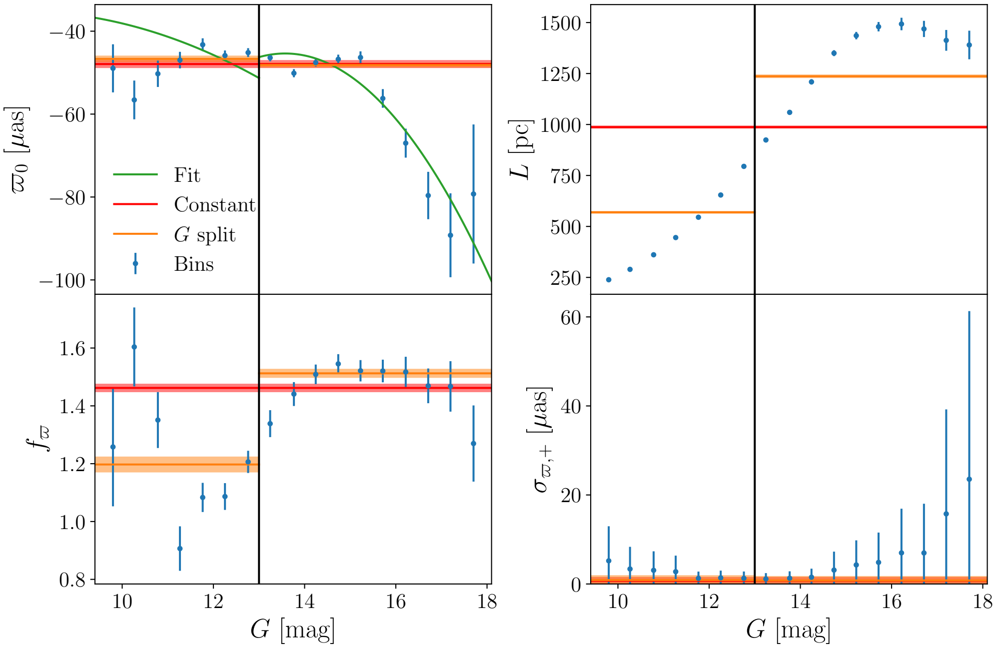

Figure 2.10: The variations in parallax model parameters with $G$. All results shown are from $K_s$-based models in which $M_{\mathrm{ref}} = -1.622$ is fixed. *Blue:* Inferred model parameters across bins in $G$. *Green:* Inferred quadratic dependences for $\varpi_0$ as a function of $G$. The separate parameterizations for $G < 13$ and $G \geq 13$ are shown.*Red:* Inferred distributions of constant parameters for all $G$. *Orange:* Inferred parameters from a model in which parameters are considered constant in $G$, with these four parameters being considered separately for $G < 13$ and $G \geq 13$. The parallax zero point offset is approximately constant down to $G \lesssim 15$, where the majority of our sample lives.

### 2.4.2 Variation of the *Gaia* Zero Point

We begin by reporting the results from inferring a quadratic form of the zero point dependence on $G$ separately for sources $G < 13$ and $G \geq 13$ to account for differences in astrometric solution. This was done using the $K_s$ photometry to calibrate the red clump luminosity. We find

$$\begin{aligned}\varpi_0/\mu\text{as} =& 9.29^{+0.88}_{-0.91}(G-12)^2 - 37.45^{+1.59}_{-1.57}(G-12)\\ &+ 30.47^{+2.92}_{-3.00} \qquad (G < 13),\end{aligned} \tag{2.25}$$

$$\begin{aligned}\varpi_0/\mu\text{as} =& -0.27^{+0.51}_{-0.49}(G-14.5)^2 - 12.04^{+0.61}_{-0.59}(G-14.5)\\ &- 11.28^{+1.46}_{-1.49} \qquad (G \geq 13).\end{aligned} \tag{2.26}$$

This is illustrated in Figure 2.9 in the panel labeled “Free $M_{\text{ref}}$”. The inferred $G < 13$ fit appears to behave very differently from the constant parallax zero points inferred in the prevoius models, allowing for relatively large positive values of the zero point for sources of low $G$. This appears to be an effect of using photometric information to infer the red clump luminosity calibration, the distances to each star, as well as the zero point dependence on $G$ all at the same time. The fact that this model infers $M_{\text{ref}} = -1.943 \pm 0.013$ in $K_s$ (See Table 2.1) in contrast to the previous model further suggests that this method of modeling the parallax zero point offset variation with $G$ may be inaccurate to some degree.

To alleviate this possible degeneracy of photometric information, we implement an identical model with the exception of applying a strict prior by fixing $M_{\text{ref}} = -1.622$. We find

$$\begin{aligned}\varpi_0/\mu\text{as} =& -0.59^{+0.83}_{-0.80}(G-12)^2 - 4.96^{+1.01}_{-1.03}(G-12)\\ &- 45.66^{+0.88}_{-0.86} \qquad (G < 13),\end{aligned} \tag{2.27}$$

$$\begin{aligned}\varpi_0/\mu\text{as} =& -2.68^{+0.49}_{-0.49}(G-14.5)^2 - 4.90^{+0.55}_{-0.54}(G-14.5)\\ &- 47.62^{+0.59}_{-0.61} \qquad (G \geq 13),\end{aligned} \tag{2.28}$$

which appears to agree with previous estimates much better. This is illustrated in Figure 2.9 in the panel labeled “Fixed $M_{\text{ref}}$”.

To demonstrate a model in which the prior on $M_{\text{ref}}$ is slightly more relaxed, we repeat the above analyses. This time, $M_{\text{ref}}$ is given a prior which reflects the inferred value from the base model. In other words, we implement:

$$p(M_{\text{ref}}) = \mathcal{N}(-1.622, (0.004)^2), \tag{2.29}$$

and we find

$$\begin{aligned}\varpi_0/\mu\text{as} =& 0.55^{+0.81}_{-0.84}(G-12)^2 - 8.82^{+1.09}_{-1.08}(G-12)\\ &- 36.54^{+1.24}_{-1.26} \qquad (G < 13),\end{aligned} \tag{2.30}$$

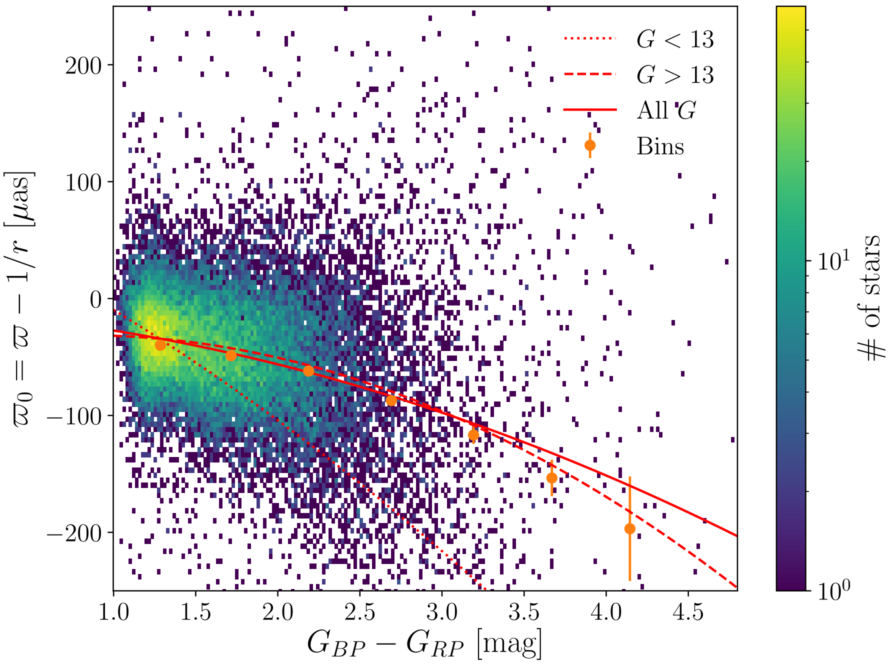

Figure 2.11: Variations of the *Gaia* parallax zero point with observed colour. *Red:* The quadratic parameterizations of the variation. *Orange:* The parallax zero point across colour bins for all $G$. A density of individual estimates from each star using Equation (2.8) with measured parallax and distance posteriors from the "All $G$" quadratic parameterization model is shown in the background. The $G < 13$ quadratic parameterization appears to not be a good model of the parallax zero point variation with colour; whereas, the remaining models agree with one another.

$$\begin{aligned}\varpi_0/\mu\text{as} = &-2.38^{+0.50}_{-0.49}(G-14.5)^2 - 5.77^{+0.54}_{-0.55}(G-14.5)\\ &-43.32^{+0.74}_{-0.76} \qquad (G \geq 13).\end{aligned} \tag{2.31}$$

We infer $M_{\text{ref}} = -1.658 \pm 0.003$ from this model, and the resulting fit can be seen in Figure 2.9 in the panel labeled "Prior $M_{\text{ref}}$". The inferred parameterization appears similar to the Fixed $M_{\text{ref}}$ model, with a slightly stronger slope in $G < 13$. Both models appear to favour a linear solution to the $G$ dependence of the zero point, with quadratic coefficients consistent with 0.

We also report results from a model considering a single quadratic parameterization of the zero point dependence on $G$ for all $G$. For this model, we fix $M_{\text{ref}} = -1.622$. We find

$$\begin{aligned}\varpi_0/\mu\text{as} = &-0.52^{+0.17}_{-0.16}(G-16)^2 - 7.65^{+0.89}_{-0.87}(G-16)\\ &-62.50^{+1.18}_{-1.10}.\end{aligned} \tag{2.32}$$

Finally, we investigate a model in which we do not enforce a specific functional parameterization of the parallax zero point's dependence on $G$. The results are shown in Figure 2.10, in which the parallax parameters $\theta_\varpi = \{\varpi_0, f_\varpi, \sigma_{\varpi,+}\}$ and $L$ are modeled as separate constants along bins in $G$ space. We fix $M_{\text{ref}} = -1.622$ in this model. The variations of $\varpi_0$ with $G$ appear to be well modeled by the quadratic parametrization for $G \geq 13$, while it appears to be better parameterized as a constant for $G < 13$.

The *Gaia* parallax zero point also appears to exhibit variations with respect to observed colour. The following models allow $M_{\mathrm{ref}}$ to be free, as the observed colour and magnitude of the sources should not share the same degeneracies as in the $G$ dependent models. In the quadratic parameterization for all sources, we find

$$\begin{aligned}\varpi_0/\mu\mathrm{as} = &-6.21^{+1.68}_{-1.66}([G_{BP}-G_{RP}]-1)^2\\ &-22.71^{+3.00}_{-2.96}([G_{BP}-G_{RP}]-1)-27.35^{+1.49}_{-1.43}.\end{aligned} \tag{2.33}$$

In considering separate quadratic parameterizations for sources $G < 13$ and $G \geq 13$, we find

$$\begin{aligned}\varpi_0/\mu\mathrm{as} = &-8.74^{+13.18}_{-13.42}([G_{BP}-G_{RP}]-1)^2 \qquad (G<13)\\ &-85.67^{+13.25}_{-13.05}([G_{BP}-G_{RP}]-1)-9.72^{+3.18}_{-3.06},\end{aligned} \tag{2.34}$$

$$\begin{aligned}\varpi_0/\mu\mathrm{as} = &-13.58^{+1.88}_{-1.86}([G_{BP}-G_{RP}]-1)^2 \qquad (G\geq 13)\\ &-5.19^{+3.35}_{-3.52}([G_{BP}-G_{RP}]-1)-31.82^{+1.59}_{-1.58}.\end{aligned} \tag{2.35}$$

All three parameterizations are shown in Figure 2.11, where the inferred functional forms are compared to individual estimates of $\varpi_0$ for each star using its measured parallax and the mean of its inferred distance posterior in the original $K_s$ photometry model in Equation (2.8). The separate parameterization for $G < 13$ appears not to be a good fit to the data; whereas, the parameterization for all $G$ appears to behave better. This is further supported with the analysis of the variations with observed colour in bins, the results from which are also shown in Figure 2.11. We also find significant variation of the parallax zero point across the sky. Sky maps of zero points with hierarchical `HEALPIX` resolutions are shown in Figure 2.12. The maps appear to show variations with respect to Galactic latitude, which we attribute to the correlation of the characteristic scale length associated with the distance prior (Equation (2.15)). The APOGEE data set reaches much deeper along the Galactic plane, meaning inferred values of $L$ are larger close to the Galactic plane. The inferred $\varpi_0$ across each patch is then slightly correlated with $L$.

It is highly unlikely that the parallax zero point offset would have a pattern that follows Galactic latitude, as the *Gaia* satellite does not a priori know about the Galactic plane. The observed variation with sky position is therefore most likely due to intrinsic differences between the red clump stars at different Galactic latitudes that are not fully captured by our model. But the observed zero point variations across the sky still serve as an upper limit on the true variation of the zero point across the sky, because a large variation would be picked up by our model. In an analysis of the variations of $\varpi_0$ across the sky for all $G$, we find the median parallax zero point (with $\pm 1\sigma$ dispersion) across patches to be $\tilde{\varpi}_0 = -42.18^{+13.10}_{-14.96}$ $\mu$as. Similarly we find the spreads in patches for each split in $G$ to be $\tilde{\varpi}_0(G < 13) = -39.92^{+17.82}_{-24.96}$ $\mu$as and $\tilde{\varpi}_0(G \geq 0) = -41.03^{+13.11}_{-14.80}$ $\mu$as. We do not expect the *Gaia* parallax zero point to fluctuate by more than a few 10s of $\mu$as for any given observed magnitude and colour, as such fluctuations would be the ones captured by the inference rather than those induced by the degeneracies with $L$.

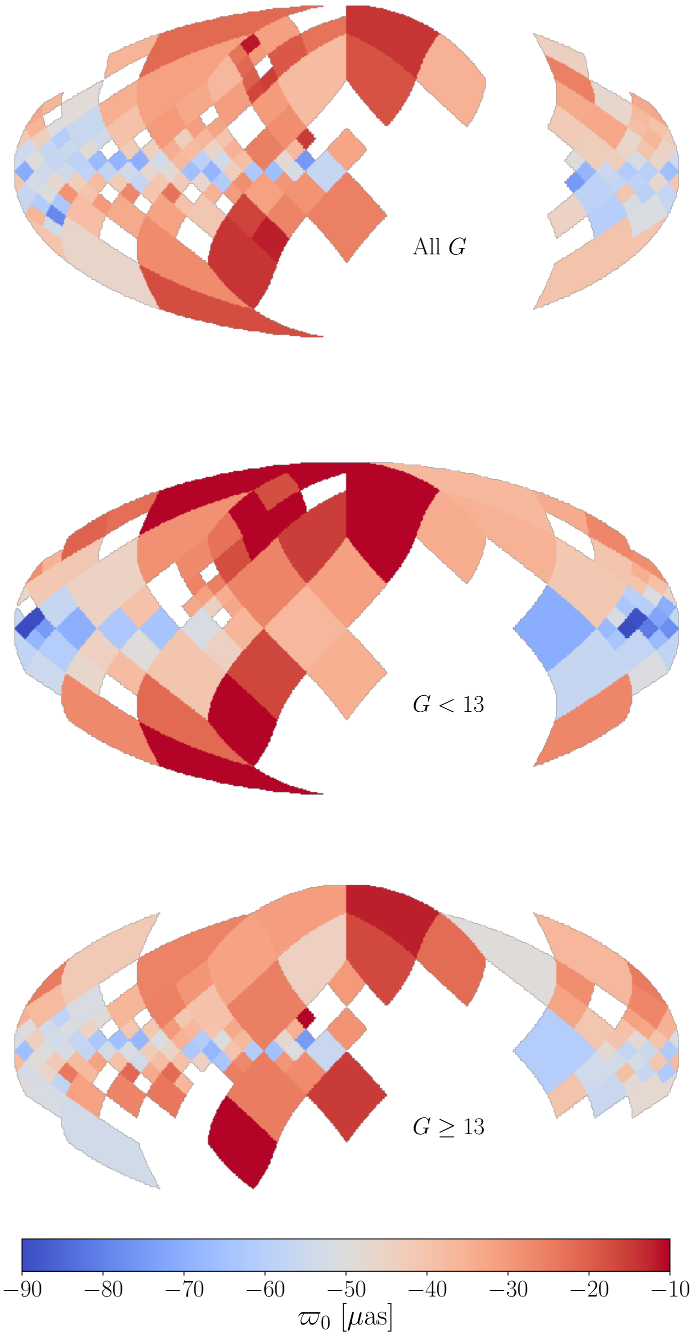

Figure 2.12: Variations of the *Gaia* parallax zero point across the sky, shown here in Galactic coordinates. Patches with $\varpi_0/\sigma_{\varpi_0} < 3$ have been masked. *Top:* Inferred zero points for all $G$. *Middle:* Inferred zero points for $G < 13$. *Bottom:* Inferred zero points for $G \geq 13$.

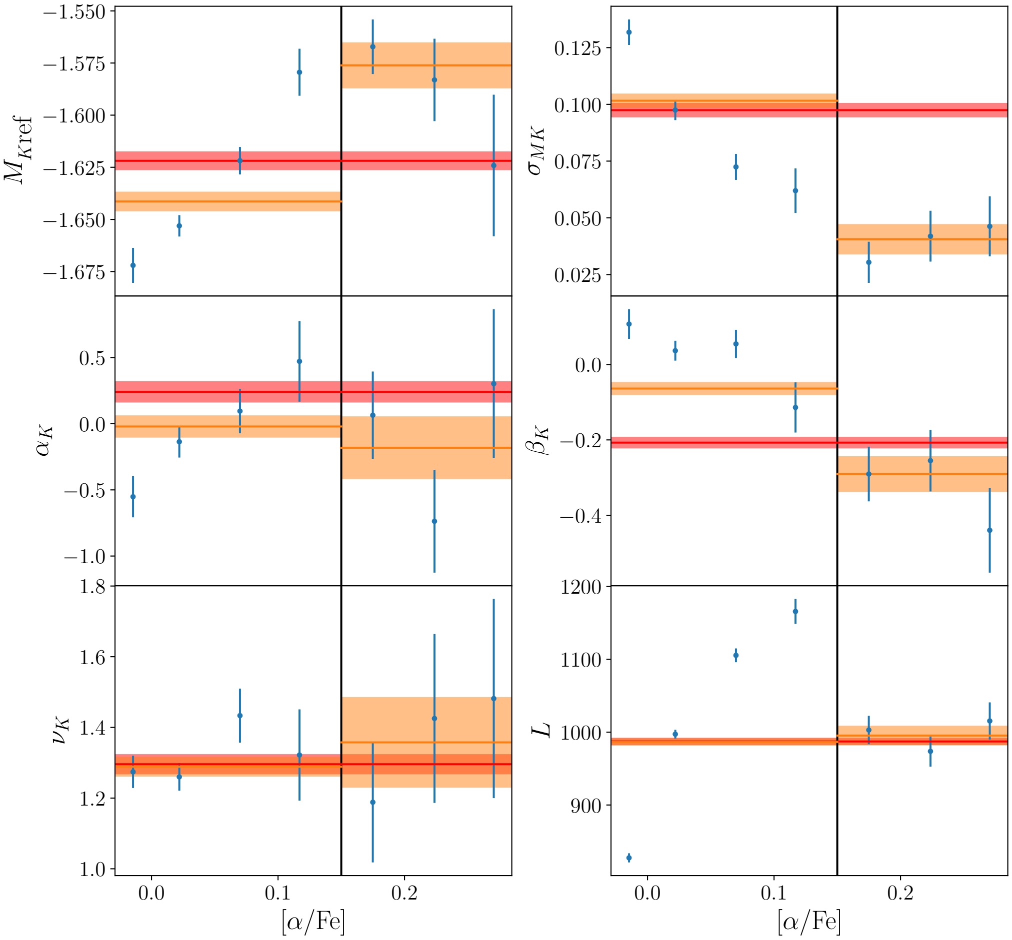

Figure 2.13: The dependence of red clump luminosity calibration parameters with [$\alpha$/Fe]. *Blue:* Median and $\pm1\sigma$ error bars for posteriors in each [$\alpha$/Fe] bin. *Red:* Median and $\pm1\sigma$ ranges for posterior distributions from the base $K_s$ model for all $G$. *Orange:* Median and $\pm1\sigma$ ranges for posterior distributions for simply splitting between the low- and high-[$\alpha$/Fe] sub-populations. *Black:* Split at [$\alpha$/Fe] = 0.15 dex between the low- and high-[$\alpha$/Fe] sub-populations.

### 2.4.3 Modeling the Red Clump

The analysis of the red clump luminosity calibration in $K_s$ as a function of [$\alpha$/Fe] are shown in Figure 2.13. The absolute magnitude exhibits clear trends for the low-alpha ([$\alpha$/Fe] $< 0.15$ dex) and high-alpha ([$\alpha$/Fe] $> 0.15$ dex) populations. In addition, the inverse variance averages of the calibration scatter in each population are FWHM/2.355 $= 0.103 \pm 0.004$ mag for [$\alpha$/Fe] $< 0.15$ dex, and FWHM/2.355 $= 0.040 \pm 0.008$ mag for [$\alpha$/Fe] $\geq 0.15$. While the luminosity dependence on [Fe/H], $\beta_K$, appears to become stronger with [$\alpha$/Fe], the evolution of the $(J_0 - K_0)$ slope $\alpha_K$ does not seem to be as well constrained. This is understandable, as $\alpha_K$ has also been less constrained than $\beta_K$ for every other analysis. Nevertheless, this is further evidence for different behaviours between at least two sub-populations of red clump stars, and suggests that more detailed modeling of these populations may be necessary for precise use of red clump stars in the future. The small luminosity scatter for high-[$\alpha$/Fe] stars means that highly precise red-clump distances can be obtained for them.

Finally, we discuss the results of the verification of stellar models used in the APOGEE red clump catalogue discussed in §2.3.5. We find that the inferred residual $M_{\mathrm{ref}} = -0.036 \pm 0.004$ with a dispersion of FWHM/2.355 $= 0.096 \pm 0.004$ shows the stellar evolution models describe the overall luminosity of the red clump quite well. However, we infer significant non-zero values of $\alpha = 1.48 \pm 0.07$ and $\beta = -0.39 \pm 0.01$ mag/dex, implying that the stellar evolution models leave significant residual luminosity dependencies on stellar temperature and metallicity [Fe/H].

## 2.5 Discussion

### 2.5.1 Consistency of zero point calibrations

Throughout most of our analyses, the *Gaia* parallax zero point is usually inferred to be within the range of $\varpi_0 \approx -47$ to $-49$ $\mu$as when modeled as a constant. There are two main exceptions. First, the model using $G$ photometry only infers $\varpi_0 \approx -38$ $\mu$as. We attribute this to an incomplete understanding of interstellar extinction in $G$, and we resolved the inconsistency by supplementing the model with $K_s$ photometry. The $K_s$ measurements include robust measurements of extinction, and shift the inference with $G$ photometry to match with the other models. The second inconsistency comes from inferring two constant zero points for sources dimmer and brighter than $G = 13$. As similar inconsistencies were seen with the modeling of the zero point variations with $G$, this seemed to be caused by a degeneracy associated with modeling both the zero point dependence and the red clump luminosity calibration simultaneously with photometric information. By either fixing the red clump absolute magnitude to a previously inferred value including an informed prior based on a previous posterior, these models give inferred zero points consistent with the all other models.

Finally, we discuss the possibility of attributing the variations in $\varpi_0$ across both $G$ and sky position to a correlation with the given extinction corrections (namely $A_K$ for most of our models). Running the base model using $K_s$ photometry with an additional sample cut for $0 < A_K < 0.1$ results in an inferred $\varpi_0 = -40 \pm 1$ $\mu$as with a corresponding $M_{K,\mathrm{ref}} \approx -1.611$. Conversely, the sample containing $A_K \geq 0.1$ returns $\varpi_0 = -36 \pm 2$ $\mu$as with a corresponding $M_{K,\mathrm{ref}} = -1.720$. This seems to contradict what the inferred parallax zero points with respect to $G$ may suggest, but they in fact support each other. Within our model, a good estimate of the parallax zero point requires: 1) A good calibration of the red clump luminosity using nearby stars, and 2) A large enough sample of stars far away such that $\varpi_0$ is significant relative to the star's parallax. We can see the effects of removing either of these pillars in our analyses. The zero point inference split between bright/faint stars uses all stars to calibrate the luminosity; therefore, it retains a similar $M_{K,\mathrm{ref}}$ as the original model. The parallax zero points inferred for $G < 13$ are generally more inconsistent because these brighter sources tend to be closer to us; whereas, the fainter sources tend to have parallaxes that are more significanly affected by $\varpi_0$. Similarly, the low extinction sample above retains a consistent luminosity calibration, but fails to estimate a consistent parallax zero point because the stars are mostly closer to us. The high extinction sample is naively expected to provide a good $\varpi_0$, but it fails because it is unable to infer $M_{K,\mathrm{ref}}$ well because it has few bright, nearby stars. In conclusion, in order to use our models to obtain an estimate of the *Gaia* parallax zero point to both high accuracy *and* precision, we require a large sample of stars that contains populations both close by (to get a good handle on the red clump luminosity) as well as far away (more significantly affected by $\varpi_0$).

### 2.5.2 Comparison to other zero point work

With the resolution of tensions between the models presented in this paper, we choose to report a constant *Gaia* DR2 parallax zero point of $\varpi_0 = -47.97 \pm 0.79$ $\mu$as as inferred by our base model with $K_s$ photometry. This measurement can be seen in Figure 2.14 in comparison to other values reported in the literature discussed previously. Note that Khan et al. (2019) also finds a systematic offset in *Gaia* DR2 parallaxes in the range of $-45$ to $-55$ $\mu$as. Our inferred parallax zero point is the most precise to date, and is in good agreement with most other measurements. Note that the red clump occupies a specific space in both observed colour ($1 \lesssim G_{BP} - G_{RP} \lesssim 3.5$) and magnitude ($9.5 \lesssim G \lesssim 18$), so the inferred zero point reported in this paper becomes less accurate outside of this space. This can explain why the *Gaia* zero point reported by Lindegren et al. (2018) is quite discrepant with our own, as quasars typically occupy a dimmer and bluer part of the *Gaia* observation space.

Variations of the parallax zero point appear to be best described using models discussed in this paper with Equation (2.27) and Equation (2.28) for $G$ dependent fluctuations as well as Equation (2.33) for observed colour dependent fluctuations. This is in contrast to the functional parameterizations of the dependences reported by Leung & Bovy (2019a). The data sets in both studies occupy similar regions of observed colour and magnitude, yet we infer variations that are much larger in amplitude. In addition, both quadratic dependences analyzed in this paper reflect parabolas with negative curvature; whereas, the parameterizations reported in Leung & Bovy (2019a) exhibit positive curvature. We are uncertain as to why this may be the case.

Finally the variations of $\varpi_0$ with sky position are inferred here to fluctuate between $-80$ $\mu$as to $-10$ $\mu$as, with a typical scatter across the sky of approximately 15 $\mu$as about $-40$ $\mu$as. We conclude that the *Gaia* parallax zero point must not vary by more than a few 10s of $\mu$as across the sky. Our range of systematic fluctuations is similar to the $\sim 100\mu$as peak-to-peak fluctuations reported by Arenou et al. (2018). It is also worth noting that Khan et al. (2019) also find peak-to-peak variations of the same order for the parallax zero point across the sky.

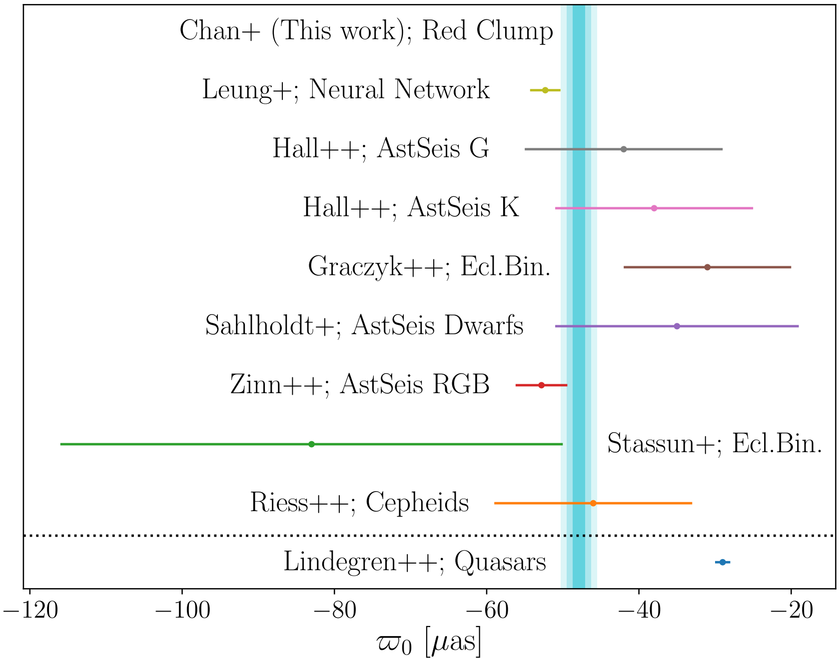

Figure 2.14: A comparison of reported *Gaia* parallax zero point values associated with Data Release 2. These measurements include those which make use of asteroseismology, neural networks applied to other red giants in APOGEE, eclipsing binaries, and Cepheid variables. The Zinn et al. (2019) result is shown at $\nu_{\mathrm{eff}} = 1.5$ and $G = 12.2$. The turquoise vertical stripes indicate the 1, 2, and $3\sigma$ regions of the *Gaia* DR2 parallax zero point that we report. Our measurement of the global *Gaia* zero point offset is consistent with all previous determinations using stars, and it is the most precise so far.

### 2.5.3 Comparison to previous calibrations of the RC

We infer the absolute magnitude of the red clump in the 2MASS as well as *Gaia* photometry bands. The inferred absolute magnitudes in $K_s$ and $G$ appear to be consistent with those reported by Hall et al. (2019). We remark that, while the inferred absolute magnitude in $K_s$ are in agreement with that reported in Hawkins et al. (2017), the absolute magnitudes we infer in $J$, $H$, and $G$ are not compatible. Of course, this only applies to the estimate of the peak of the red clump luminosity distribution or the absolute magnitude of a typical red clump star. The scatters in absolute magnitude inferred in both this paper and in Hawkins et al. (2017) are large enough to contain both estimates. It is possible that incomplete modeling of the red clump sub-populations as described in §2.3.4 could contribute to such differences in red clump calibrations.

## 2.6 Conclusions

The advent of the *Gaia* mission has provided the astronomical community with an excellent opportunity to use ultra precision astrometry; however, a lack of understanding of the parallax zero point has prevented us from unlocking the full potential of *Gaia* data. A precise measurement of the *Gaia* parallax zero point requires a large data set, as well as a robust understanding of the methods involved. We have presented several hierarchical probabilistic models using red clump stars to infer the parallax zero point, while simultaneously calibrating an empirical parameterization for the red clump luminosity. In doing so, we infer the *Gaia* DR2 parallax zero point to be $\varpi_0 = -48 \pm 1$ $\mu$as using $K_s$ photometry. Models using other 2MASS or *Gaia* photometry allow for consistent estimates of the zero point. The use of a Student's $t$-distribution appears to describe the distribution of luminosities in our red clump sample quite well. We also report the absolute magnitude of the red clump to be $M_{\rm ref} = -1.622 \pm 0.004$ in $K_s$, $M_{\rm ref} = 0.517 \pm 0.004$ in $G$, $M_{\rm ref} = -1.019 \pm 0.004$ in $J$, and $M_{\rm ref} = -1.516 \pm 0.004$ in $H$. We find the intrinsic spread of the red clump to be $\sim 0.09$ in $J$, $H$, and $K_s$. The scatter in $G$ is $\sim 0.17$, and this larger value can be attributed to an incomplete understanding of interstellar extinction. Each probabilistic model also infers the distance to every star used as input, yielding typical distance estimates of $\sim 10\%$.

Additions to the base probabilistic model allow for more detailed investigation into either into the variations of the *Gaia* parallax zero point or the red clump luminosity calibration. We find that the variations in the zero point are most significant for dim sources $G \gtrsim 16$, while the zero point offset is constant at brighter magnitudes. Note that less than 7% of our total sample lies within $G \gtrsim 16$, so the inferred parallax systematics may not be as accurate as in brighter bins. The dependence of the zero point on observed colour is can also be parameterized with a quadratic form. Fluctuations of the zero point across the sky are difficult to infer, but we limit them to be less than a few 10s of $\mu$as. We also find significant variations of the red-clump luminosity model with $[\alpha/\mathrm{Fe}]$ and in particular with different trends across low- and high-$\alpha$/Fe sub-populations.

The sheer size of the red clump sample has allowed us to estimate the *Gaia* DR2 parallax zero point to approximately 1.6%. This is the highest precision estimate of the parallax zero point to date, and presents the community with a fantastic outlook for using *Gaia* for high-precision distance estimates. Although it is expected to be smaller in amplitude and better understood, the parallax zero point will still be present in future *Gaia* data releases. We also expect the quality of *Gaia* data to be better in the future, implying that these probabilistic methods will eventually allow us to infer distances to stars to high enough accuracy and precision to be impactful in areas of galactic dynamics, constructing local distance ladders, and much more.

# Chapter 3

# Modelling the LMC distance with *Gaia* DR3

### Abstract

A precise distance to the Large Magellanic Cloud (LMC) is a key component of the absolute scale of the cosmic distance ladder. We present an independent estimate of the LMC distance at $R_{\mathrm{LMC}} = 48.90^{+0.52}_{-0.48}$ kpc using only data from *Gaia*'s Data Release 3. We select a sample of LMC stars in *Gaia* DR3, and compare their $XP$ spectra to bright giants nearby with high quality parallaxes. Matching stars between samples with similar spectra allows a probabilistic model to infer the LMC distance on the assumption that they share the same intrinsic brightness. The model works best when a quantitative estimate for the probability that 'matched' stars are truly equally bright is also provided. Our inferred LMC distance is in agreement with the value reported by Pietrzyński et al. (2019).

## 3.1 Introduction

The Large Magellanic Cloud (LMC) serves as an important distance calibrator used in the 'absolute scale' of cosmic distance ladders, and it is the second most significant contributor to the error budget in the local estimate of the Hubble constant $H_0$ (Riess et al., 2022). Pietrzyński et al. (2019) report the highest precision estimate at $R_{\mathrm{LMC}} = 49.59 \pm 0.09\,(\mathrm{statistical}) \pm 0.54\,(\mathrm{systematic})\,\mathrm{kpc}$. In this chapter, we independently estimate the distance to the LMC using data from the *Gaia* mission's third data release (DR3). The *Gaia* mission provides high precision astrometry for its extensive catalogue, and this work is a study of whether *Gaia* data can provide an alternative calibration to the absolute scale. We select a sample of stars within the LMC through an astrometric and kinematic selection procedure in §3.2, and a second sample of nearby, bright stars with high quality parallaxes is also selected therein. We present a simple method for comparing the sampled spectral data between stars in *Gaia* DR3 to find 'twin' stars in §3.3. Spectroscopic twins have previously been shown to be valid for applications in the cosmic distance ladder (Jofré et al., 2015), and they have also been used to determine the distance to the Pleiades (Mädler et al., 2016). Our probabilistic model for the LMC distance is presented, along with the results they produce in §3.4. We conclude with a discussion of the results in §3.5.

## 3.2 Data

The *Gaia* mission published their Data Release 3 (DR3) in 2022 (Gaia Collaboration et al., 2023). This includes low resolution $XP$ spectra from within their $BP \in [330, 680]$ nm, as well as $RP \in [640, 1050]$ nm photometry bands. The $XP$ spectral data are available in a continuous representation of 55 coefficients per band which are readily transformed into the conventional sampled spectra with provided basis functions (Carrasco et al., 2021; Gaia Collaboration et al., 2023). $XP$ sampled spectra occupy 343 bins from 336 nm to 1020 nm, and will be used for the sources selected in this chapter.

This section describes the selection of two samples from the DR3 catalogue: one sample of stars in the LMC, and another sample of stars with high quality parallax measurements in the solar neighbourhood (also referred to as Milky Way or MW stars). $XP$ sampled spectra as well as apparent magnitudes $G$ are used for all stars in both samples. From the MW sample, parallax measurements $\varpi$ as well as their associated errors $\sigma_\varpi$ are also used in the distance model for the LMC (described later in §3.4). Note that the *Gaia* reported parallaxes are given in units of mas.

### 3.2.1 LMC sample selection

Out of the 220 million objects with $XP$ spectra, we select a subsample of the DR3 table of stars in the general area of the LMC using the following query to the *Gaia* archive[1]:

```
SELECT * FROM gaiadr3.gaia_source as g
WHERE 1=CONTAINS(POINT('ICRS',g.ra,g.dec),
CIRCLE('ICRS',81.28,-69.78,20))
AND g.parallax IS NOT NULL
AND g.has_xp_sampled='True'
AND g.phot_g_mean_mag < 20.5
```

---

[1] https://gea.esac.esa.int/archive/

The query is adapted from Gaia Collaboration et al. (2021a), and it selects all objects that satisfy the following conditions:

- Observed in a circle of radius 20° centered around the LMC position $(\alpha, \delta) = (81.28°, -69.78°)$
- Has reported $XP$ spectra and parallax, and
- $G < 20.5$

The total solid angle of the circle in this query is intentionally large to include the outer regions of the LMC, as well as the LMC-SMC bridge (Gaia Collaboration et al., 2021a). The query returns 682 209 rows, and the table can be accessed at this link[2]. We make further cuts to this sample following the steps in Section 2 of Gaia Collaboration et al. (2021a). In summary, the sample is reprojected into coordinates centered on the dynamical motions of HI gas in the LMC, and then cuts are made based on the relative parallaxes, proper motions and their errors such that the final sample is astrometrically and kinematically consistent with all objects being in one collective group with some bulk motion. The steps are as follows:

1. Transform $(\alpha, \delta)$, $(\mu_{\alpha*}, \mu_\delta)$ into an orthographic projection, and remove all sources with $\sqrt{x^2 + y^2} > \sin(5\,\mathrm{deg})$. The dynamical center of $\mathrm{H_I}$ gas in the LMC $(\alpha_{C,\mathrm{LMC}}, \delta_{C,\mathrm{LMC}}) = (78.77°, -69.01°)$ is used here (Gaia Collaboration et al., 2021a). Note this is slightly different from the center used in the original sample query above, as this center used in the definition of the projection in Gaia Collaboration et al. (2018a).

$$x = \cos\delta \sin(\alpha - \alpha_C), \tag{3.1}$$

$$y = \sin\delta \cos\delta_C - \cos\delta \sin\delta_C \cos(\alpha - \alpha_C), \tag{3.2}$$

$$\mu_x = \mu_{\alpha*} \cos(\alpha - \alpha_C) - \mu_\delta \sin\delta \sin(\alpha - \alpha_C), \tag{3.3}$$

$$\mu_y = \mu_{\alpha*} \sin\delta_C \sin(\alpha - \alpha_C) + \mu_\delta(\cos\delta \cos\delta_C + \sin\delta \sin\delta_C \cos(\alpha - \alpha_C)). \tag{3.4}$$

2. Remove foreground stars with parallax signal-to-noise $\varpi/\sigma_\varpi > 5$, since stars in LMC should have poorly measured parallax.

3. For the next step, remove stars with magnitude $G > 19$ to select only for relatively brighter LMC stars that should have higher quality astrometry.

4. Compute median proper motion values for the LMC $(\mu_{x,\mathrm{med}}, \mu_{y,\mathrm{med}})$ from $(\mu_x, \mu_y)$, and compute their covariance $\Sigma_{\mu_x,\mu_y}$.

5. With the sample obtained at the end of Step 3, choose only stars such that $\mu' \Sigma^{-1} \mu' < 9.21$, where $\mu' = (\mu_x - \mu_{x,\mathrm{med}}, \mu_y - \mu_{y,\mathrm{med}})$. This corresponds to retaining stars that are within a $3\sigma$ confidence region of the median proper motion of the stars that are likely to belong to the LMC.

6. Compute new proper motions $(\hat{\mu}_{\alpha*}, \hat{\mu}_\delta)$ **for the full sample** which are conditional on the median sample parallax $\varpi_{\mathrm{med}}$ being the true parallax of each star. In other words, we shift the observed proper motions of each star in the sample by under the assumption that they all

---

[2] `https://gea.esac.esa.int/tap-server/tap/async/16647294362180/results/result`

truly at the LMC distance. This can be done with errors $\sigma$ and correlations $\rho$ provided in the *Gaia* data:

$$\hat{\mu}_{\alpha*} = \mu_{\alpha*} - (\varpi - \varpi_{\mathrm{med}})\rho_{\mu_{\alpha*}\varpi}\frac{\sigma_{\mu_{\alpha*}}}{\sigma_\varpi}, \tag{3.5}$$

$$\hat{\mu}_{\delta} = \mu_{\delta} - (\varpi - \varpi_{\mathrm{med}})\rho_{\mu_{\delta}\varpi}\frac{\sigma_{\mu_{\delta}}}{\sigma_\varpi}. \tag{3.6}$$

7. Repeat 1-4 with $(\hat{\mu}_{\alpha*}, \hat{\mu}_\delta)$, this time with the magnitude limit $G < 20.5$.
8. Make a final cut such that only stars which satisfy $\hat{\mu}'\Sigma^{-1}\hat{\mu}' < 9.21$ are retained.

After applying these steps, we are left with 22 004 sources out of the original 682 209 from the query. The proper motions of the final sample are shown in Figure 3.1. The selection appears to match the results in Gaia Collaboration et al. (2021a), and the rotation of the LMC is apparent in the astrometric data alone.

As a consistency check, we attempt the same kinematic selection procedure on a sample of sources in a null field. This is a field chosen such that its center (in Galactic coordinates $(l, b)$) is exactly 180° away in $l$ from the LMC center, with all other parameters exactly the same. The ADQL query for the null field is as follows:

```
SELECT * FROM gaiadr3.gaia_source as g
WHERE 1=CONTAINS(POINT('ICRS',g.ra,g.dec),
CIRCLE('ICRS',58.77,8.93,20))
AND g.parallax IS NOT NULL
AND g.has_xp_sampled='True'
AND g.phot_g_mean_mag < 20.5
```

The query returns 358 976 sources, and can be accessed here[3]. Applying the same astrometric and kinematic cuts as the LMC sample reduces this sample to 248 sources. The size of the LMC sample after cuts is approximately 100× larger than the size of the null sample after cuts, meaning that we ought to be quite confident that the stars in the LMC sample are really in the LMC.

The final LMC sample from *Gaia* DR3 contains 22 004 stars, and a colour-magnitude diagram of the sample is shown in Figure 3.2 along with a distribution of the reported parallaxes. It is apparent that the parallax measurements of the sample are not consistent with the LMC distance of $R_{\mathrm{LMC}} = 49.59\,\mathrm{kpc}$ measured by Pietrzyński et al. (2019) even after applying the parallax zero point offset of $\varpi_0 = -39$ $\mu$as reported in Groenewegen (2021), so we are motivated to supplement the data in the *Gaia* LMC subsample to estimate a more accurate distance to the LMC.

### 3.2.2 High quality parallax sample selection

We select a sample of stars with 'high quality' parallax measurements from the *Gaia* DR3 catalogue, later referred to as the Milky Way (MW) sample. The requirement is a measured parallax signal-to-noise $\varpi/\sigma_\varpi > 3$. The idea is to search for twin stars with similar spectra, and therefore brightness, so we limit the search according to the absolute magnitude limit of the LMC stars. The magnitude cutoff for $XP$ spectra is $G < 15$, which, at the distance of the LMC, corresponds to an absolute magnitude of

---

[3] https://gea.esac.esa.int/tap-server/tap/async/16672412001770/results/result

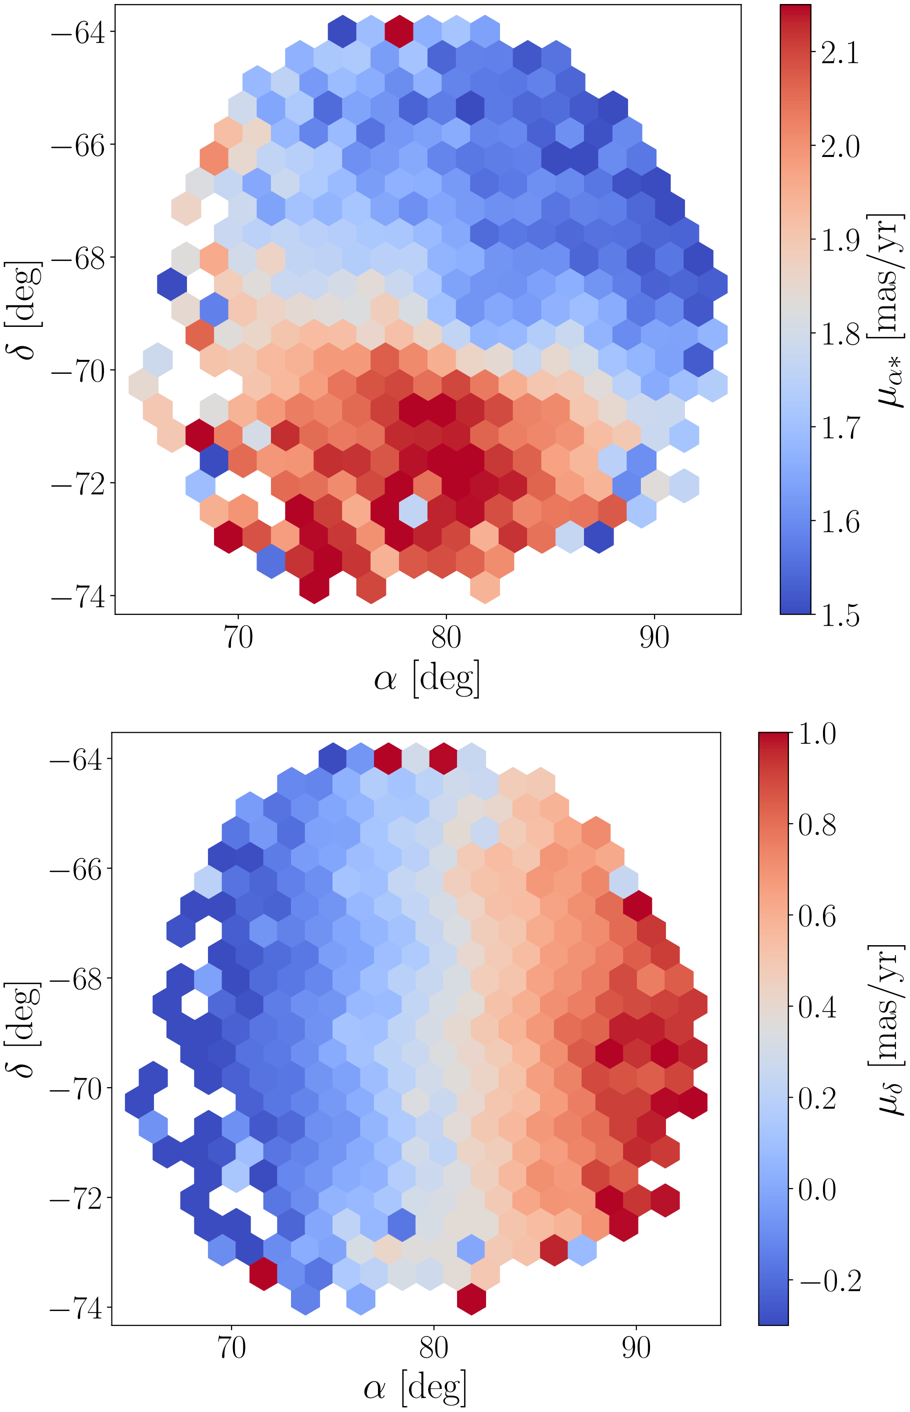

Figure 3.1: Proper motion distribution of the LMC sample. The LMC kinematic selection produces a clean sample whose astrometry directly reveals the bulk motion of the LMC.

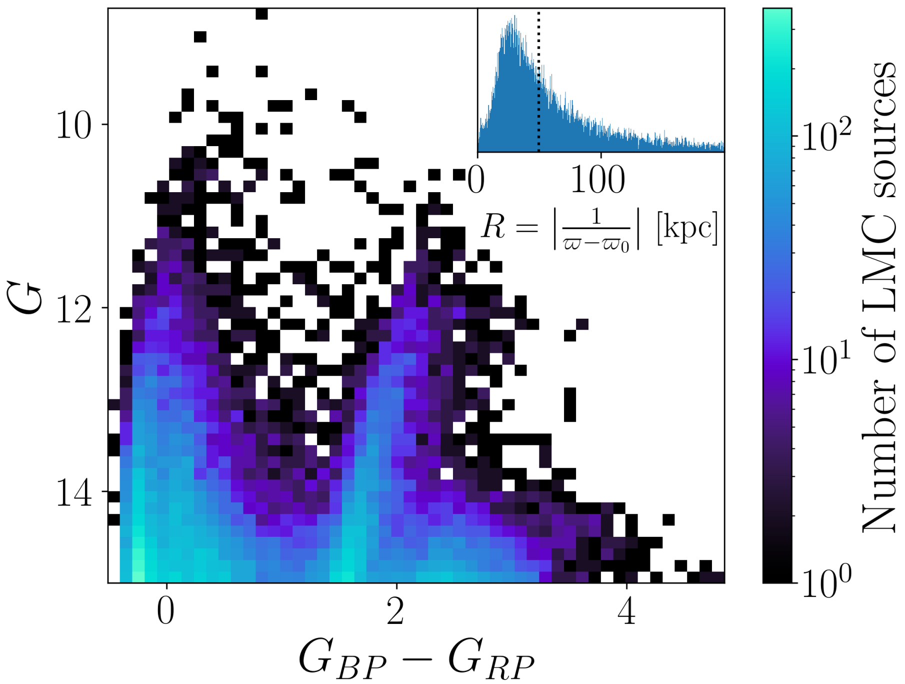

Figure 3.2: The colour-magnitude diagram of the LMC sample. The inset on the upper-right shows the distance distribution of the sample assuming that each measured parallax is truly the inverse distance to its source after zero point correction (Lindegren et al., 2021b). The distribution of LMC parallax distances peaks at approximately $R \sim 30\,\mathrm{kpc}$, and the median parallax distance is $R \sim 53\,\mathrm{kpc}$. We can see that the simply relying on the parallaxes in the LMC sample is not sufficient to estimate the LMC distance to the same accuracy and precision as the value $R_{\mathrm{LMC}} = 49.59\,\mathrm{kpc}$ reported by Pietrzyński et al. (2019) is indicated by the dotted black line.

$$
\begin{aligned}
M_G &= G - 5\log_{10}(R_{\mathrm{LMC}}) + 5 && (3.7)\\
&\lesssim 20 - 5\log_{10}(50\,\mathrm{kpc}), && (3.8)
\end{aligned}
$$

where we have substituted in the apparent magnitude limit of the LMC sample $G < 15$, as well as an approximate distance to the LMC of $R_{\mathrm{LMC}} = 50\,\mathrm{kpc}$. The equivalent apparent magnitude limit of an equivalent star in the Milky Way is then

$$
\begin{aligned}
G_{\mathrm{MW}} &= 5\log_{10}(r_{\mathrm{MW}}) - 5 + M_G && (3.9)\\
&\lesssim 5\Big(\log_{10}(1/\varpi) - \log_{10}(50\,\mathrm{kpc})\Big) + 15 && (3.10)\\
&\lesssim 5\log_{10}\left(\frac{1}{(50\,\mathrm{kpc})\varpi}\right) + 15. && (3.11)
\end{aligned}
$$

Several rough approximations for both the LMC distance and the nearly star distances were made here, and they serve only to quickly choose a subsample of the *Gaia* DR3 dataset that could be treated more carefully later on in an analysis. The only additional requirement for the subsample are the availability of $XP$ spectra, and the final query for the high quality parallax subsample is as follows:

```
SELECT * FROM gaiadr3.gaia_source as g
WHERE g.parallax IS NOT NULL
AND g.parallax / g.parallax_error > 3
AND g.has_xp_sampled='True'
AND g.phot_g_mean_mag < 5 * log10(1/(50*abs(g.parallax))) + 15
```

This returns 5 632 sources, which can be accessed here[4]. The sample size is relatively small compared to the LMC sample due to the restrictions that we've made. The search was limited to stars that are very bright (bright enough to have observable spectra from the LMC), and the parallax signal-to-noise requirement limits the physical search radius around the Sun to a much smaller volume than the LMC (sufficiently bright stars are not considered if they are too distant to have precise parallaxes). We further refine the MW sample by making an extra on the inverse parallax, retaining only stars for which $r = 1/(\varpi - \varpi_0) < 5\,\mathrm{kpc}$. The volume-limited MW sample contains 2 318 stars.

### 3.2.3 Correcting for dust extinction

The reported $G$ magnitudes have not been corrected for extinction (Fitzpatrick et al., 2019; Danielski et al., 2018; Gaia Collaboration et al., 2021b). All $G$ magnitudes used in this chapter have been approximately corrected for extinction $G_{\mathrm{corrected}} = G_{\mathrm{reported}} - A_G$. The extinction in $G$ was approximated with reddening values $E(B-V)$ scaled by a factor of 2.45 (computed from Equation 4 using Table 1 both from Wang & Chen (2019)). For the LMC, this is generally

---

[4] https://gea.esac.esa.int/tap-server/tap/async/16702663198720/results/result

$$A_{G,\mathrm{LMC}} = 2.45 \times E(B-V)_{\mathrm{LMC}}, \tag{3.12}$$

where $E(B-V)_{\mathrm{LMC}} = 0.11$ as reported by Imara & Blitz (2007) (see Table 2 therein). The MW extinction is simliarly computed as

$$A_{G,\mathrm{MW}} = 2.45 \times E(B-V)_{\mathrm{MW}}, \tag{3.13}$$

where each $E(B-V)_{\mathrm{MW}}$ were sampled from the `Combined19` dust map in the `mwdust`[5] package (Bovy et al., 2016) with the `filter='E(B-V)'` option. Note that leaving `filter=None` would otherwise add an extra conversion factor of 0.884 defined by Schlafly & Finkbeiner (2011). The $E(B-V)_{\mathrm{MW}}$ values were sampled at the reported *Gaia* coordinates $(\alpha, \delta)$ along with the inverse parallax distance $r = 1/(\varpi - \varpi_0)$ of each star in the MW sample. As a final cut on the MW sample, we remove any stars with $E(B-V)_{\mathrm{MW}} > 1$ to avoid any stars with extreme extinction. The final MW sample contains 2 139 stars.

## 3.3 Comparing *Gaia* XP spectra

The *Gaia XP* sampled spectra are available in terms of the total observed flux $F_i$ per pixel $i$ in the spectrum. An example spectrum from each of the LMC and MW samples are shown in Figure 3.3. The observed flux of each star is dependent on the distance, so they are not readily comparable to each other. We normalize each star's $XP$ flux $F_i$ as well as their reported uncertainties $\sigma_{F_i}$ by the median value of its flux $\overline{F}$, shown for the same two stars in the middle panel of Figure 3.3:

$$f_i = \frac{F_i}{\overline{F}}, \tag{3.14}$$

$$\sigma_{f_i} = \frac{\sigma_{F_i}}{\overline{F}}. \tag{3.15}$$

We can then compare two normalized spectra by computing the $\chi^2$ per pixel between them as follows:

$$\chi_i^2 = \frac{(f_{i,\mathrm{LMC}} - f_{i,\mathrm{MW}})^2}{\sigma^2_{f_{i,\mathrm{LMC}}} + \sigma^2_{f_{i,\mathrm{MW}}}}. \tag{3.16}$$

A simple method for choosing matches with the full set of $\chi_i^2$ values between the two samples is to combine all 343 pixels into the reduced $\chi^2$ value. If we are to trust the reported uncertainties of the *Gaia XP* spectra, then we expect that $\chi^2_{\mathrm{red}} \approx 1$ for two stars with truly identical spectra. Another way to consider the spectral comparisons, is that each pixel of the $\chi_i^2$ comparison for identical spectra should follow a standard $\chi^2$ distribution. The Kolmogorov–Smirnov test (KS test) can be used in this case to make quantitative statements about the similarity between any two spectra. The KS test is a non-parametric statistic that compares the cumulative distribution functions (CDFs) of two samples. In this case, we would like to compare the set of computed $\chi_i^2$ from all pixels of any two spectra to a standard $\chi^2$ distribution. The KS test returns a $p$-value that quantifies the probability that every pixel of the computed $\chi_i^2$ are drawn from a $\chi^2$ distribution with one degree of freedom. In

---

[5] https://github.com/jobovy/mwdust

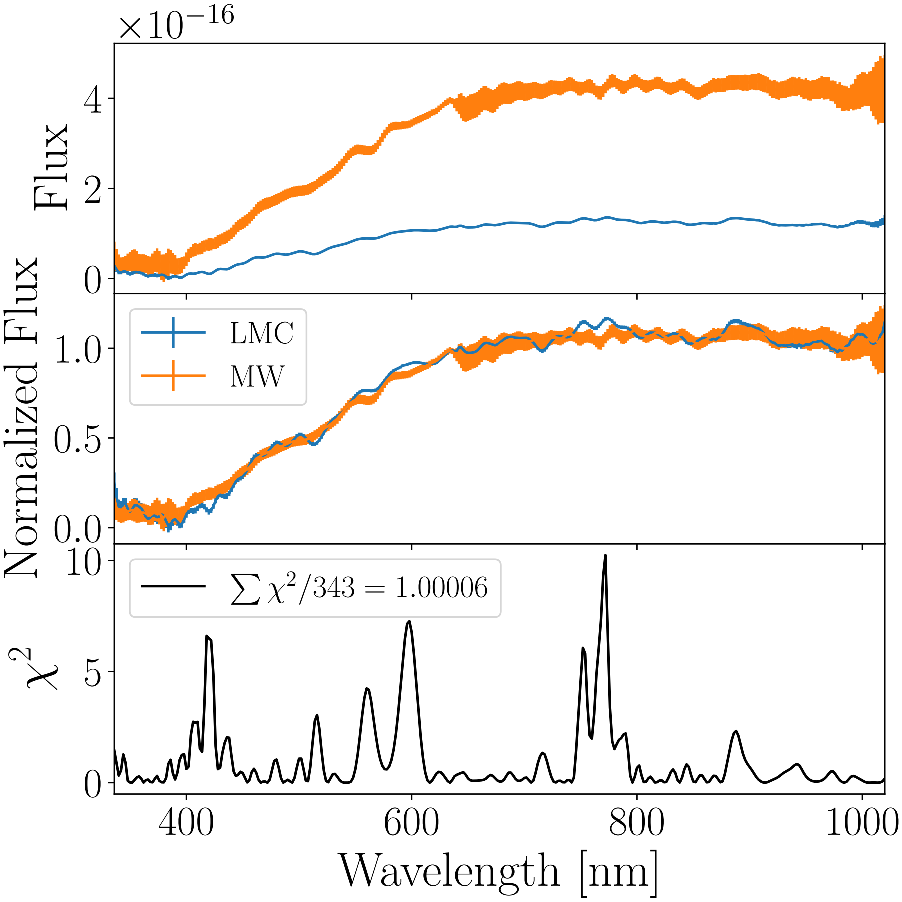

Figure 3.3: *Top*: $XP$ sampled spectrum of a star in the LMC sample in *blue*, and a star in the MW sample in *orange*. *Middle*: Both spectra normalized by their respective median flux. *Bottom*: The $\chi^2$ comparison per pixel between the two spectra. The reduced $\chi^2$ between these two particular spectra is very close to 1, indicating that these stars are strong candidates to be stellar twins.

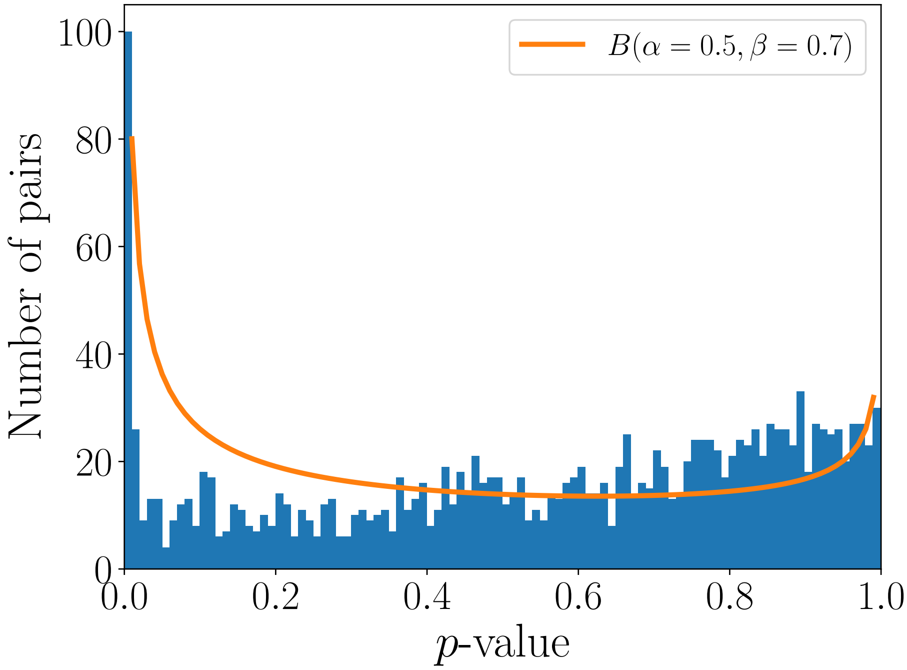

Figure 3.4: Distribution of the highest $p$-values between LMC and MW spectra for all $N_{\text{match}} = 1\,550$ matched stars with $0.99 < \chi^2_{\text{red}} < 1.01$. Also shown is the probability density function of a Beta distribution with parameters $\alpha = 0.5$ and $\beta = 0.7$. This does not exactly match the distribution of $p$-values, but it captures the qualitative description that *most* of the matches are not likely to be true twins, while *some* are.

other words, the $p$-value can be interpreted as the probability that the two stars have observationally similar spectra.

After comparing each star in the LMC sample ($N_{\text{LMC}} = 22\,004$) to every star in the MW sample ($N_{\text{MW}} = 2\,139$), there are a total of $N_{\text{match}} = 1\,550$ stars in the LMC sample with at least one spectral comparison to the MW sample such that $0.99 < \chi^2_{\text{red}} < 1.01$. This is the set of 'matched' stars that we use in the following section. In the case of LMC stars with more than one spectral comparison satisfying $0.99 < \chi^2_{\text{red}} < 1.01$, the MW star with the highest $p$-value is chosen as the 'matched' star.

## 3.4 LMC distance model

If we are truly able to find twin stars between the LMC and MW samples, then the distance to the LMC can be modelled by calibrating the absolute magnitudes of each pair of matched stars, and comparing with the apparent magnitude $G$ of each LMC star. For each pair of matched stars, the absolute magnitude $M_G$ is

$$M_G = G - 5\log_{10}(r) + 5, \tag{3.17}$$

where $r$ is the distance to the star in parsecs. Alternatively, one may choose to rearrange Equation (3.17) to solve for the distance $r$ to the star:

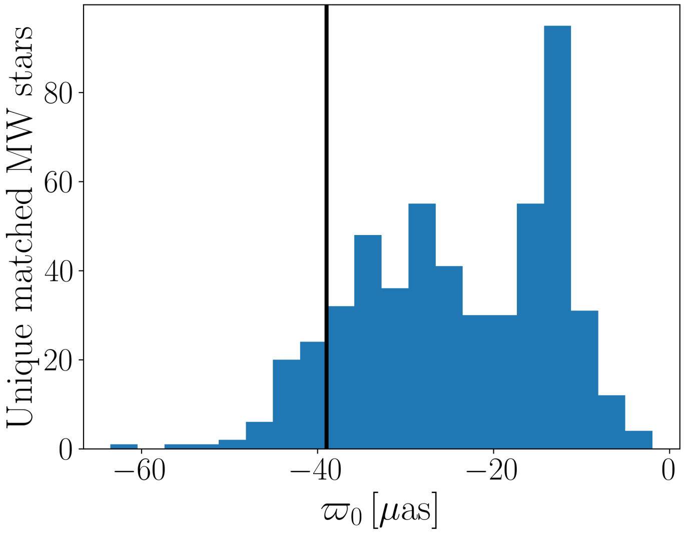

Figure 3.5: Computed parallax zero points for the set of MW stars that are spectrally matched to at least one LMC star. A vertical black line indicates the approximate value $\varpi_0 = -39\,\mu$as reported by Groenewegen (2021). The constant parallax zero point they report is not representative of the distribution of computed parallax zero points we compute for the nearby sample of matched MW stars.

$$r = 10^{\frac{G - M_G + 5}{5}}. \tag{3.18}$$

If a star in the LMC sample is matched to a star in the MW sample, then we may choose to substitute Equation (3.17) for the MW star into Equation (3.18) for a representation of the distance to the LMC star:

$$r_{\mathrm{LMC}} = 10^{\frac{G_{\mathrm{LMC}} - G_{\mathrm{MW}} + 5 \log_{10} r_{\mathrm{MW}}}{5}}. \tag{3.19}$$

The distance to the MW star may be modelled with respect its observed parallax $\varpi$. Each *Gaia* parallax can be modelled being drawn from a Gaussian distribution centered on the true parallax $\varpi'$ with reported uncertainty $\sigma_\varpi$:

$$p(\varpi|\varpi', \sigma_\varpi) \sim \mathcal{N}(\varpi', \sigma_\varpi), \tag{3.20}$$

where the true parallax of the MW star is dependent on its distance $r_{\mathrm{MW}}$ and its parallax zero point $\varpi_0$. We perform two inferences with each of our models. The first inferences use parallax zero points $\varpi_0$ computed from the `gaiadr3_zeropoint`[6] package provided by the *Gaia* collaboration (Lindegren et al., 2021b). The second inferences assume a constant $\varpi_0 = -39$ $\mu$as as reported in Groenewegen (2021) as a test of the sensitivity of our results to the parallax zero point. A visual comparison between the computed and constant zero point values is shown in Figure 3.5. The true parallax of the MW star can be expressed as

$$\varpi' = \frac{1}{r_{\mathrm{MW}}} + \varpi_0. \tag{3.21}$$

---

[6] `https://pypi.org/project/gaiadr3-zeropoint/`

The true distance to each MW star is unknown, so we may choose to model it as a free parameter, with the exponentially decreasing volume density prior with scale distance $L$ (Bailer-Jones, 2015):

$$p(r|L) = \frac{r^2}{2L^3} \exp(-r/L). \tag{3.22}$$

Returning to Equation (3.19), the distance to each LMC star can be modelled as being drawn from a Gaussian distribution centered on the true distance $R_{\mathrm{LMC}}$ with some uncertainty $\sigma_R$:

$$p(r_{\mathrm{LMC}}|R_{\mathrm{LMC}}, \sigma_R) \sim \mathcal{N}(R_{\mathrm{LMC}}, \sigma_R). \tag{3.23}$$

The estimated distance to each LMC star is dependent on the apparent magnitudes of both stars, and the distance to the MW star. The uncertainty in *Gaia* magnitudes at $G = 17$ is $\sigma_G = 1\,\mathrm{mmag}$ (Gaia Collaboration et al., 2021b). As seen in Figure 3.2, the LMC stars all have $G < 15$, so we choose $\sigma_{G,\mathrm{LMC}} = 1\,\mathrm{mmag}$ as a conservative upper limit on the uncertainty in $G_{\mathrm{LMC}}$. The stars in the MW sample have magnitudes $G_{\mathrm{MW}} < 13$, which correspond to reported uncertainty $\sigma_{G,\mathrm{MW}} = 0.3\,\mathrm{mmag}$ (Gaia Collaboration et al., 2021b). The distance to the MW star is modelled as a free parameter, so its uncertainty is not immediately clear. In our model, the distance to the MW star is dependent on its parallax, so we may choose to model the uncertainty in the distance to the MW star as being dependent on the uncertainty in its parallax (refer to Equation (3.21)). This will make up the bulk of the contribution to $\sigma_R$. Propagating uncertainties through Equation (3.19) yields the following expression for the uncertainty in the distance to each LMC star:

$$\sigma_R = \log(10) \times 10^{\frac{G_{\mathrm{LMC}} - G_{\mathrm{MW}} + 5\log_{10} r'}{5}} \times \sqrt{\frac{\sigma^2_{G,\mathrm{LMC}} + \sigma^2_{G,\mathrm{MW}}}{25} + \frac{r'^2 \sigma^2_{\varpi}}{(\log 10)^2}}, \tag{3.24}$$

where $r' = 1/(\varpi - \varpi_0)$ is computed once using the MW parallax in each match. For all spectral pairs with $0.99 < \chi^2_{\mathrm{red}} < 1.01$, the median value of $\sigma_R$ computed with Equation (3.24) is approximately 3 kpc, and the distribution of $R$s and $\sigma_R$s computed with Equation (3.19)-(3.24) is shown in Figure 3.6.

A posterior distribution for the distance to the LMC can be expressed by combining the parallax likelihood (Equation (3.20)), the LMC likelihood (Equation (3.23)), and the MW prior (Equation (3.22)):

$$p(R_{\mathrm{LMC}}|G_{\mathrm{LMC}}, G_{\mathrm{MW}}, \varpi) \propto \mathcal{N}(\varpi|\varpi', \sigma_\varpi)\mathcal{N}(r_{\mathrm{LMC}}|R_{\mathrm{LMC}}, \sigma_R)p(r_{\mathrm{MW}}|L). \tag{3.25}$$

A probabilistic graphical model depicting Equation (3.25) along with the dependencies between its observables and parameters is shown in Figure 3.7.

As was in Chapter 2, the distance calibration means that there is a free parameter $r_{\mathrm{MW}}$ for every pair of stars in the data in addition to the global parameters. This motivates the use of `Stan` once again to sample our parameter space. The only restriction on the parameters ($r_{\mathrm{MW}}$, $L$, and $R_{\mathrm{LMC}}$) beyond the previously described priors are that they be strictly positive. The results after running four chains for 2000 samples each, and discarding the first half of each chain are shown in Figure 3.8.

The LMC distance inferred by the base model with computed parallax zero points provided by the *Gaia* collaboration is $R_{\mathrm{LMC}} = 48.84 \pm 0.10\,\mathrm{kpc}$, which is in agreement with the value $R_{\mathrm{LMC}} = 49.59 \pm 0.09\,(\mathrm{statistical}) \pm 0.54\,(\mathrm{systematic})\,\mathrm{kpc}$ reported by Pietrzyński et al. (2019). Meanwhile, the value inferred with the base model and a constant parallax zero point is $R_{\mathrm{LMC}} = 47.07 \pm 0.10\,\mathrm{kpc}$,

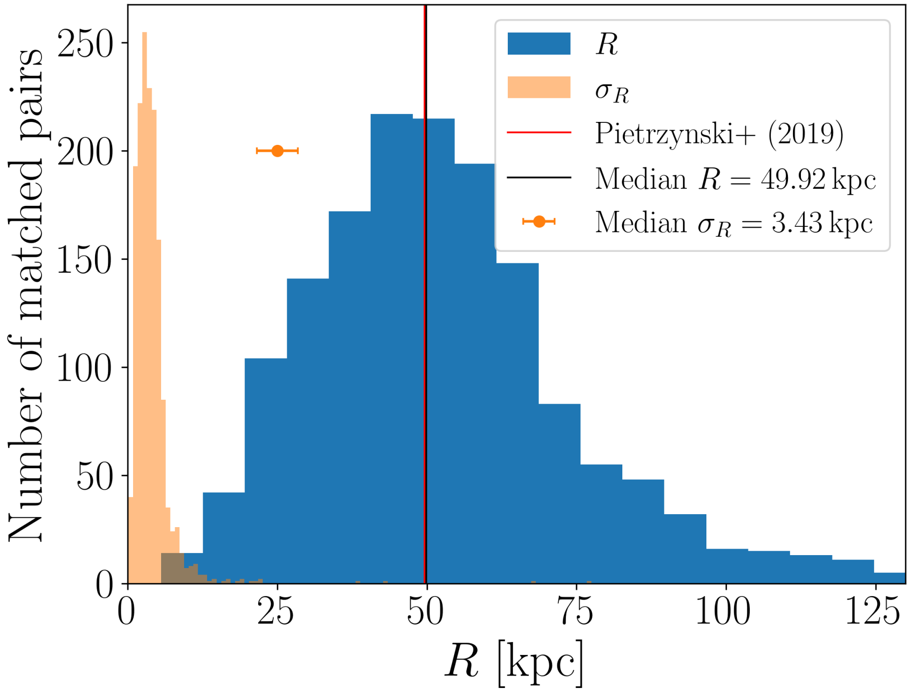

Figure 3.6: Distribution of naive LMC distances $R$ computed with Equation (3.19) assuming the inverse parallax distance of the matched MW star is correct. The distribution of $\sigma_R$ computed with Equation (3.24) is also shown. An example error bar indicates the approximate observational uncertainty predicted for a typical pair of matched stars using Equation (3.24), which is the median value of the orange histogram. The median $R$ is already in agreement with value reported by Pietrzyński et al. (2019).

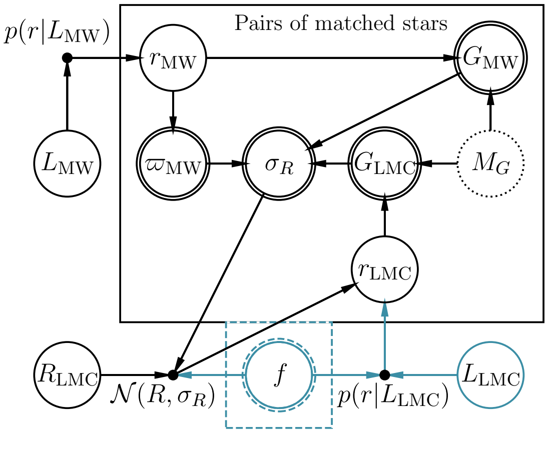

Figure 3.7: Probabilistic graphic model of the LMC distance model described in this chapter. Observed quantities (data) are depicted with double circles. Model parameters are depicted with single circles. Certain priors and likelihoods are depicted with solid points. The mixture model extension is depicted in blue, with the mixture parameter $f$ being either a global parameter or a unique parameter for each pair of matched stars. In the latter model, it is also possible for $f$ to be an observable in the model rather than a free parameter. The absolute magnitude of each matched star $M_G$ is depicted with a dotted circle to indicate its influence while not being an explicit parameter in the model.

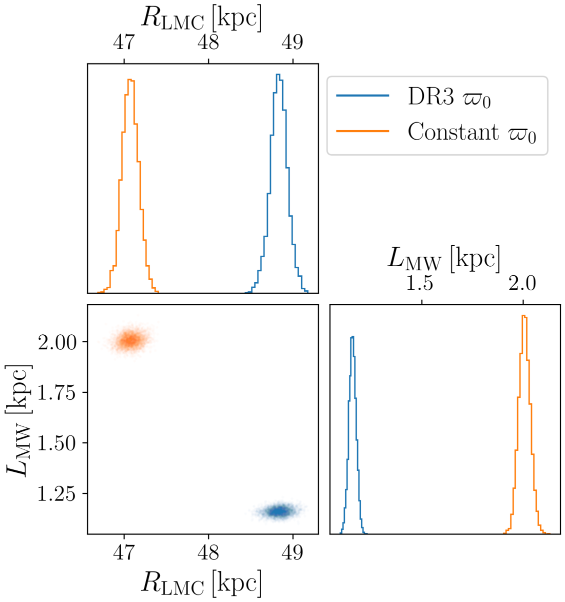

Figure 3.8: Posterior samples for the base LMC distance model. The inferred parameters exhibit a strong dependence on the parallax zero point.

which is in significant disagreement ($\sim 4.5\sigma$) with the literature value. It is, however, slightly more accurate than simply relying on the reported parallaxes of the LMC sample (refer to Figure 3.2). Regardless of the parallax zero point, both distances are inferred to much higher precision than the uncertainties reported by Pietrzyński et al. (2019). This may be attributed in part to our determination of the LMC distance uncertainty in Equation (3.24), which is purely a statement about the uncertainty of the photometry and astrometry used in the model. The base model operates under the rigid assumption that the pairs of matched stars are truly twin stars that can be calibrated to have identical intrinsic brightnesses. It may at first seem reasonable that 1 550 of the stars in the full LMC sample of 22 004 stars, or approximately 7% of the LMC sample, were matched. Conversely, there were only 2 139 stars in the MW to compare to in the first place, and it seems far-fetched to think that approximately 72% of nearby bright stars have spectrally identical stars in the LMC.

The model can be more flexible to the outliers present in the data if it can allow for the possibility that the matched stars are not truly identical. One way to do this is to introduce an additional uncertainty term $\tilde{\sigma}_R$ to Equation (3.24) as a free parameter in the model, which would grow if the model prefers a solution in which there are more discrepancies between 'matched' stars. Another way to add flexibility to the model so that it is robust to outliers in the data is to replace the LMC distance likelihood (Equation (3.23)) with a mixture model:

$$p(r_{\rm LMC}|R_{\rm LMC},\sigma_R) \sim f\mathcal{N}(r_{\rm LMC}|R_{\rm LMC},\sigma_R) + (1-f)p(r_{\rm LMC}|L_{\rm LMC}), \tag{3.26}$$

where $p(r_{\rm LMC}|L_{\rm LMC})$ follows the same distance prior described in Equation (3.22). This second distribution serves as an outlier distribution of LMC 'distances' from poor matches, rather than a separate estimate of $R_{\rm LMC}$. The mixture parameter $f$ may be a global parameter. If $f$ is a global parameter and $f \to 1$, then the model is confident that most stellar matches are between true stellar

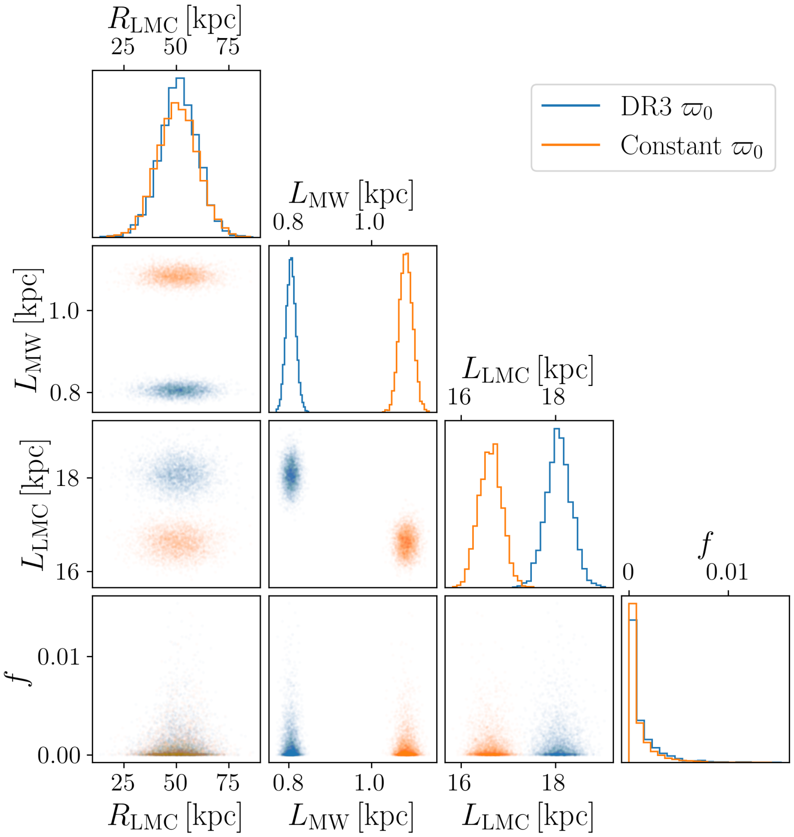

Figure 3.9: Posterior samples for the LMC distance model with a free global mixture parameter $f$. All distance parameters are presented in units of kpc, and vertical lines indicate the median and 68% confidence intervals. The LMC distance is unconstrained with a free global mixture parameter, and its sampled posterior distribution simply follows its prior.

twins in the LMC and Mw sample. Conversely, if $f \to 0$, then the model predicts that most of the 'matched' pairs of stars are not true twins. As a global parameter, this would be a statement about the set of 'matched' stars as a whole. The mixture parameter must be bounded between 0 and 1, so it is modelled as a free parameter following the Beta distribution:

$$p(f) \sim B(\alpha, \beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} f^{\alpha-1}(1-f)^{\beta-1}. \tag{3.27}$$

There is no parameterization of the Beta distribution that fully matches the distribution of $p$-values that we have obtained, but we see in Figure 3.4 that $\alpha = 0.5$ and $\beta = 0.7$ follows a similar shape. Most of the 'matches' should have low confidence, and we would hope that at least a few of them are true matches. It is not necessary for the prior distribution to exactly match the distribution of $p$-values, since the posterior distribution ought to be dominated by the likelihoods of the data. Applying the matches to a mixture model with a global mixture fraction $f$ as a free parameter gives no constraints for $R_{\rm LMC}$. This is best seen when adding a prior distribution such that $R_{\rm LMC} \sim \mathcal{N}(50\,{\rm kpc}, 10\,{\rm kpc})$, which yields the results shown in Figure 3.9.

Allowing for a global mixture fraction as a free parameter causes the LMC distance model to prefer a scenario where none of the 'matched' LMC stars are truly comparable to their MW

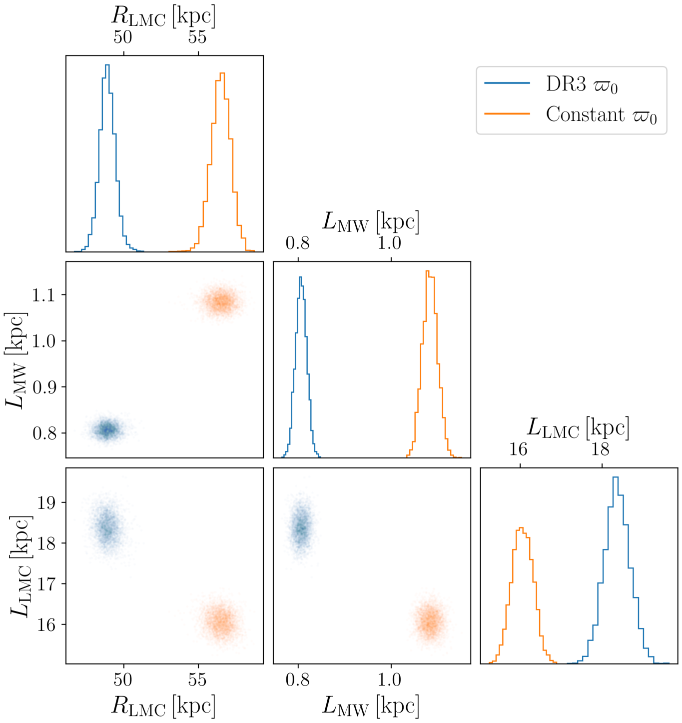

Figure 3.10: Posterior samples for the LMC distance model with a free mixture parameter $f$ for each 'matched' pair of LMC and MW stars. All distance parameters are presented in units of kpc, and vertical lines indicate the median and 68% confidence intervals. The LMC distance is constrained with individual mixture parameters, with a larger scatter in the sampled posterior as desired.

counterpart ($f \to 0$). This leaves $R_{\rm LMC}$ mostly unconstrained, as can be seen by its marginalized sampled posterior essentially following the prior distribution that we've set. One may choose instead to use mixture fractions $f'$ that are unique to each pair of matched stars, which would be a statement about the true similarity between each individual pair of stars. In this case, we have the opportunity to use the mixture parameter $f'$ either as a free parameter or as an observable. Figure 3.10 shows the model with individual mixture fractions as free parameters, which no longer requires an additional prior to constrain $R_{\rm LMC}$. The inferred LMC distance with official *Gaia* parallax zero points is similar to base model, but with a wider statistical scatter as desired: $R_{\rm LMC} = 48.90^{+0.52}_{-0.48}$ kpc. In contrast, the inference with constant parallax zero point returns $R_{\rm LMC} = 56.52^{+0.62}_{-0.66}$, which also has a wider uncertainty but is now in much larger disagreement ($\sim 8\sigma$) with the Pietrzyński et al. (2019) measurement.

Rather than allowing the mixture fraction in Equation (3.26) to be a free parameter, it is possible to set them as observables. We have previously computed $p$-values in §3.3 which can be interpreted as the probability that the $\chi^2_i$ comparison at each pixel was drawn from the $\chi^2$ distribution. The $p$-values are therefore a statement about the similarity between the two spectra, and using them as the mixture fraction $f$ in Equation (3.26) yields the sampled posterior distribution in Figure 3.11. This model does *not* an additional prior distribution on $R_{\rm LMC}$.

The inferred distance to the LMC with $p$-values is $R_{\rm LMC} = 54.55^{+0.64}_{-0.78}$ kpc with computed parallax zero points, and $R_{\rm LMC} = 52.73^{+0.92}_{-0.97}$ kpc with the constant $\varpi_0$. These are in $5.2\sigma$ disagreement and $\sim 2.8\sigma$ tension with the Pietrzyński et al. (2019) value, respectively. A potential source of error is

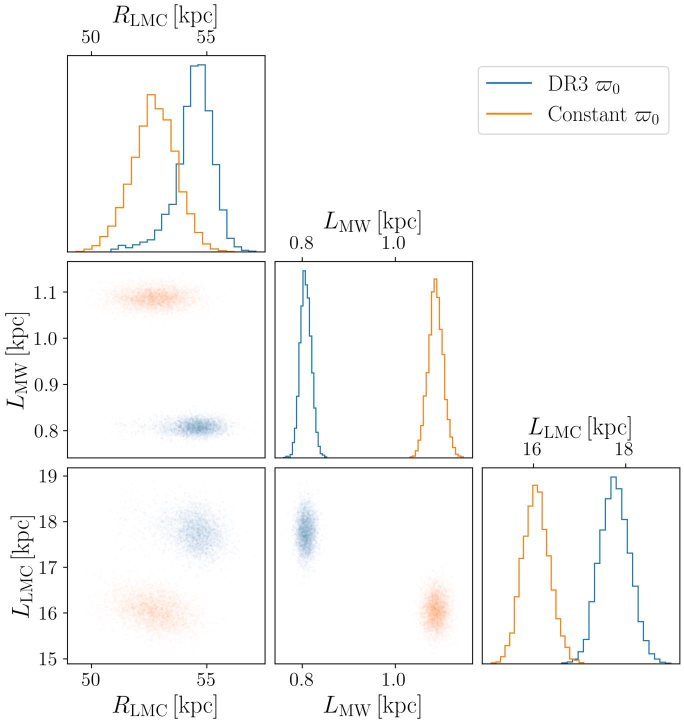

Figure 3.11: Posterior samples for the LMC distance model using each 'matched' pair's $p$-value as the mixture fraction. All distance parameters are presented in units of kpc, and vertical lines indicate the median and 68% confidence intervals. There is a large shift in the estimated LMC distance when using the $p$-values as the mixture fraction.

Table 3.1: Summary of inferred parameters from the LMC distance model. Median values are listed with 68% confidence intervals.

| Zero point | Model | $R_{\rm LMC}$ [kpc] | $L_{\rm MW}$ [kpc] | $L_{\rm LMC}$ [kpc] |
|---|---|---|---|---|
| Lindegren et al. (2021b) | Base | $48.84^{+0.10}_{-0.10}$ | $1.16^{+0.02}_{-0.02}$ | – |
| Lindegren et al. (2021b) | Global $f$ | $^{a}51.20^{+8.69}_{-8.72}$ | $0.81^{+0.01}_{-0.01}$ | $18.07^{+0.27}_{-0.26}$ |
| Lindegren et al. (2021b) | Indiv. $f$ | $48.90^{+0.52}_{-0.48}$ | $0.81^{+0.01}_{-0.01}$ | $18.37^{+0.32}_{-0.31}$ |
| Lindegren et al. (2021b) | $p$-values | $54.55^{+0.64}_{-0.78}$ | $0.81^{+0.01}_{-0.01}$ | $17.77^{+0.33}_{-0.32}$ |
| Groenewegen (2021) | Base | $47.07^{+0.10}_{-0.10}$ | $2.01^{+0.03}_{-0.03}$ | – |
| Groenewegen (2021) | Global $f$ | $^{a}50.53^{+9.29}_{-9.52}$ | $1.08^{+0.02}_{-0.02}$ | $16.62^{+0.24}_{-0.25}$ |
| Groenewegen (2021) | Indiv. $f$ | $56.52^{+0.62}_{-0.66}$ | $1.08^{+0.02}_{-0.02}$ | $16.06^{+0.26}_{-0.25}$ |
| Groenewegen (2021) | $p$-values | $52.73^{+0.64}_{-0.78}$ | $1.09^{+0.02}_{-0.02}$ | $16.06^{+0.30}_{-0.29}$ |

[a] Unconstrained; follows prior distribution.

the fact that the $\chi_i^2$ computed between pixels are not completely independent of each other. There is an avenue for improvement if the KS test can be altered to either account for the covariances between nearby pixels in spectra, or if there is another distribution that can describe the statistics of the spectral comparisons. Alternatively, one could work with the *Gaia* collaboration's continuous representation of the $XP$ spectra, which consist of 110 coefficients for each spectrum with respect to a set of basis functions that can recreate the observed spectrum. The coefficients are reported with uncertainties, so it is possible to compute a set of independent $\chi_i^2$ between two spectra by comparing their coefficients. We leave this study for future work, and report our result using the individual mixture fractions as free parameters in Equation (3.26) which gives us $R_{\rm LMC} = 48.90^{+0.52}_{-0.48}$ kpc.

## 3.5 Discussion and Conclusions

In this chapter, we:

- Selected a (LMC) sample of 22 004 stars in the Large Magellanic Cloud from the *Gaia* DR3 catalogue with reported $XP$ spectra.
- Selected a (MW) sample of 2 139 stars in the solar neighbourhood from the *Gaia* DR3 catalogue with reported $XP$ spectra, high quality parallax measurements, and low levels of expected dust extinction.
- Selected a set of 1 550 pairs of stars from the LMC and MW samples with similar spectra. These were used as candidate 'twin' stars to calibrate the absolute magnitudes of the LMC stars with MW parallaxes.
- Developed a probabilistic model to estimate the distance to the LMC by calibrating the absolute magnitudes of the LMC stars with MW parallaxes.
- Inferred the distance to the LMC to be $R_{\rm LMC} = 48.90^{+0.52}_{-0.48}$ kpc, which is in agreement with the Pietrzyński et al. (2019) value of $R_{\rm LMC} = 49.59 \pm 0.09$ (statistical) $\pm 0.54$ (systematic) kpc.
- Demonstrated that a careful treatment of the *Gaia* parallax zero point is necessary for a robust inference of the LMC distance (see Table 3.1).

This chapter provides an optimistic outlook for future *Gaia* data releases. The model constructed in this chapter uses only *Gaia* photometry and parallaxes, along with some extinction corrections for nearby stars with dust maps accumulated by Bovy et al. (2016). The set of candidate twin stars between the LMC and the solar neighbourhood relies solely on low-resolution $XP$ spectra. Future *Gaia* data will feature fainter magnitude limits, higher resolution spectra, and more precise parallaxes. This will likely increase the sizes of both the MW and LMC samples, as well as improve the spectral comparisons detailed in §3.3. There is also opportunity to use the *Gaia* collaboration's continuous representation of $XP$ spectra for improved comparisons with more accurate $p$-values. A mostly independent measurement of the distance to the LMC using *Gaia* data would be a significant contribution to the absolute scale of the cosmic distance ladder, either further solidifying the Hubble tension or providing some relief.

# Chapter 4

# Crosstalk simulations for cosmic microwave background observations

*Adapted from a report by the Canadian LiteBIRD team: Dobbs, Matt, Hložek, Renée, Smecher, Graeme, Chan, Victor C., Montgomery, Joshua, Nerval, Simran, & Scott, Douglas. The report was a funded Science Maturation Study of the LiteBIRD mission submitted to the Canadian Space Agency in 2019. The contents of this chapter contain original work by Chan, Victor C.*

## Abstract

Upcoming observations of the cosmic microwave background (CMB) face increasingly stringent requirements in order to probe extremely faint signals. The push to higher sensitivity has renewed the community's drive to understand instrument systematics. Crosstalk is a phenomenon in which detectors can become correlated and observe signals intended for other detectors, and it can occur from overlapping detector beams (optical crosstalk) or electrical leakage (electrical crosstalk). We characterize here the effects of electrical crosstalk by through simulations of realistic CMB observations. The simulations work by imitating the scanning motions of an experiment, and we modify them to mix detector timestreams through a predetermined crosstalk matrix. The mixed timestreams are then combined into an 'observed' CMB map, and we conclude that the effects of unmitigated crosstalk on CMB power spectra are at the level of a few percent at large angular scales for a LiteBIRD-like experiment. We also find that the readout ordering of the detectors (which controls which pairs of detectors crosstalk) has a non-trivial effect on the observed CMB power spectra.

## 4.1 Introduction

OBSERVATIONS of the cosmic microwave background are becoming increasingly sensitive, and it is more important than ever to understand the systematic effects present in our data. An example of such a systematic is detector crosstalk. Crosstalk is a phenomenon where detectors can become correlated and 'see' the same signals as other detectors, either through overlapping beams (*optical crosstalk*) or through voltage leaks during electrical readout (*electrical crosstalk*). Crosstalk has been well known and studied in the design and analysis of previous experiments, and it was considered to be a low level systematic effect as recently as SPTpol (Henning et al., 2018). Upcoming experiments are facing more stringent requirements in order to probe extremely faint signals in the CMB, such as the tensor-to-scalar ratio $r$ (LiteBIRD Collaboration et al., 2023), and the community is motivated to better understand the effects of crosstalk as a potential obstacle when upgrading existing experiments such as SPT-3G (Montgomery et al., 2022), or building new experiments such as LiteBIRD. This chapter describes work that we performed to characterize realistic effects of electrical detector crosstalk on observations of CMB polarization for the planned LiteBIRD mission (LiteBIRD Collaboration et al., 2023). This work involved a thorough understanding of the entire CMB observation pipeline in order to simulate the process. The contents of this chapter were included in report by the Canadian LiteBIRD team (PI: Dobbs) as part of a funded Science Maturation Study of the LiteBIRD mission for the Canadian Space Agency[1]. We characterize the general effects of crosstalk on CMB observables with a simple toy model in §4.2. We then describe our implementation of crosstalk into a suite of CMB observation simulations in §4.3, and present the effects of crosstalk propagated into CMB angular power spectra in §4.4. We conclude with a discussion of our results in §4.5.

## 4.2 Effects of crosstalk on detector timestreams

The effects of detector crosstalk on observed timestreams is similar whether it is optical or electrical. For a given detector $i$ in a set of $N_{\text{det}}$ detectors, the total observed timestream that it observes $d_i(t)$ at a given time $t$ can be modelled as a combination of the sky signal $s_i(t)$ that it observes, detector white noise $n_i(t)$, and crosstalk modulated by a matrix of crosstalk amplitudes $W_{ij}$:

$$d_i(t) = \sum_{j=1}^{N_{\text{det}}} W_{ij}\big(s_j(t) + n_j(t)\big). \tag{4.1}$$

The crosstalk matrix $W_{ij}$ is dependent on the detector properties, and a comprehensive study on the simulations and measurement of detector-to-detector electrical crosstalk amplitudes can be found in Montgomery et al. (2022). For the purposes of this work, we assume that the crosstalk matrix is known. We also assume that the diagonal of the crosstalk matrix is strictly one ($W_{ii} = 1$); i.e., an individual detector may not crosstalk with itself. In principle, full knowledge of the crosstalk matrix allows one to invert Equation (4.1) to solve for detector timestreams such that $d_i = s_i + n_i$; however, this operation is only as precise as our understanding of the crosstalk matrix itself. Uncertainties in the crosstalk matrix propagate through into the recovered timestreams, and leave residual crosstalk. In this chapter, we simulate the effects of detector crosstalk under the

---

[1] Private report

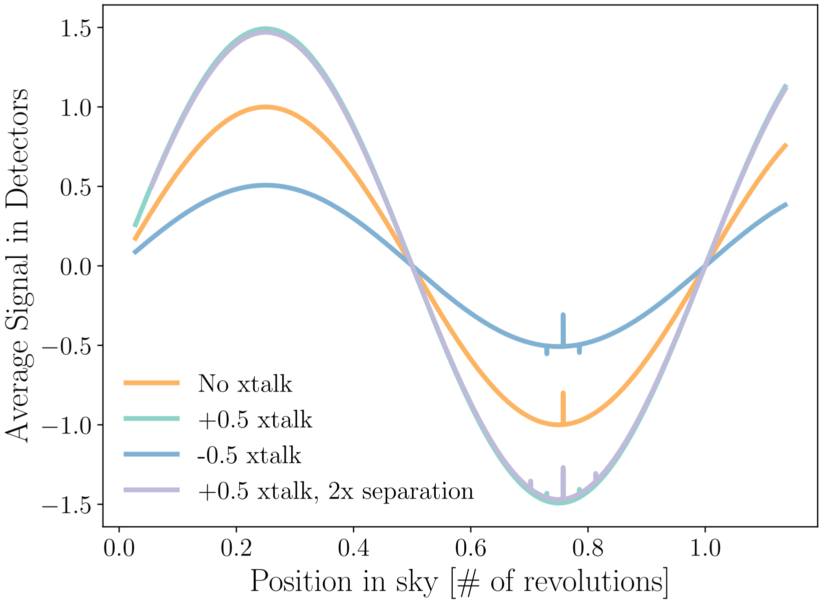

Figure 4.1: Average pointing-corrected TODs for a two detector instrument sharing several crosstalk configurations. A sinusoidal feature represents large angular scale features, while a point source near the sinusoidal minimum represents small angular scale anisotropies. Large-scale features are generally enhanced or suppressed by detector crosstalk, while small-scale features generate phantom signals.

assumption of perfect knowledge of the crosstalk matrix *without* applying the inverse operation, and study the upper limits on the impact of crosstalk on CMB observables.

We begin with a cartoon example of two detectors which are scanning a smooth ‘sky signal’ modelled by a sinusoid. If these detectors are pointed $10^\circ$ away from each other along the scanning direction, then their Time Ordered Data (TODs) will be offset by $t_{\text{off}} = 10^\circ/\omega$, where $\omega$ is the scan speed of the instrument. Detector orientations are typically dependent on the configuration of the instrument’s focal plane, and pointing corrections to the TODs are usually done in the mapmaking stage of a CMB data processing pipeline (Keihänen et al., 2005). Figure 4.1 shows averaged pointing-corrected TODs for our two detectors in several overexaggerated crosstalk configurations.

We see that the true sky signal is captured by the pointing-corrected TOD in a configuration with no crosstalk. Figure 4.1 demonstrates that features in the sky signal with angular sizes greater than the separation of the detectors are amplified (in the case of positive crosstalk) or suppressed (in the case of negative crosstalk). In contrast, features with angular sizes smaller than the detector separation generate phantom signals reminiscent of multiple images generated by gravitational lensing. The amplitude of the phantom signals of any small scale features is proportional to the magnitude of the crosstalk, and the locations of the phantom images are proportional to the angular separation of the detectors.

The crosstalk configurations presented in Figure 4.1 are quite exaggerated, but they illustrate the need for TODs to be renormalized if we wish to recover a higher fidelity sky signal. Figure 4.2 demonstrates two possible choices in renormalization. If one chooses to renormalize TODs by calibrating to the amplitude of a well-understood point source, then extra features generated by crosstalk at both large and small angular scales are somewhat less pronounced. One may instead choose to renormalize TODs by scaling by the overall RMS, or variance, of the crosstalked TODs. Both choices are agnostic to the underlying crosstalk configuration, but the latter choice is more effective at recovering the true large angular scale features in the sky signal. The drawback of this choice is that it may greatly amplify the phantom signals generated by crosstalk from small angular scale features in the sky signal. This can translate to excess power in a CMB power spectrum generated at multipoles smaller than the detector separation.

In summary, our toy model has provided us with the following insight:

1. CMB anisotropies of *larger* angular scales than the detector separation are amplified or suppressed depending on the sign of the crosstalk between each detector.
2. CMB anisotropies of *smaller* angular scales than the detector separation generate phantom signals that are proportional to the magnitude of the crosstalk and the angular separation of the detectors. Consequently, crosstalk can generate small scale power similar to gravitational lensing.
3. Renormalizing TODs can be done such that large scale features are prioritized over small scale features, but the reverse may not be as simple.

## 4.3 Simulations of detector crosstalk in CMB observations

A more detailed analysis of the impact of crosstalk on CMB observables requires a more realistic simulation of an experiment. In this section, we describe a procedure for simulating CMB observations using the `Time Ordered Astrophysics Scalable Tools`, or `TOAST`,[2] package (Kisner et al., 2023). There was no implementation for simulating crosstalk at the time this work was conducted, so we introduced modifications to the package which ultimately became the basis for the current implementation of crosstalk in `TOAST`. The inputs for a given `TOAST` simulation are a CMB map, and a set of parameters describing the experiment configuration. We generated a lensed input CMB map for the simulations in this chapter with the `Python Sky Model`'s[3] (or `PySM`) `c1` model (Thorne et al., 2017) which uses power spectra generated with the `Code for Anisotropies in the Microwave Background`[4], or `CAMB` (Lewis & Challinor, 2011). The experiment parameters are listed in Table 4.1, and they reflect the proposed LiteBIRD satellite mission (LiteBIRD Collaboration et al., 2023) parameters at the time this work was conducted[5]. Additional focal plane parameters such as detector pointing and polarization sensitivity orientations were also chosen from the proposed LiteBIRD configuration at the time.

The original `TOAST` pipeline is illustrated in Figure 4.3, and it follows these steps:

1. A data object containing all the experiment (in this case a satellite) parameters is created. Refer to Table 4.1 for the parameters used in this chapter.

---

[2] https://github.com/hpc4cmb/toast

[3] https://github.com/bthorne93/PySM_public

[4] https://camb.info/

[5] From the LiteBIRD Concept Design (internal document)

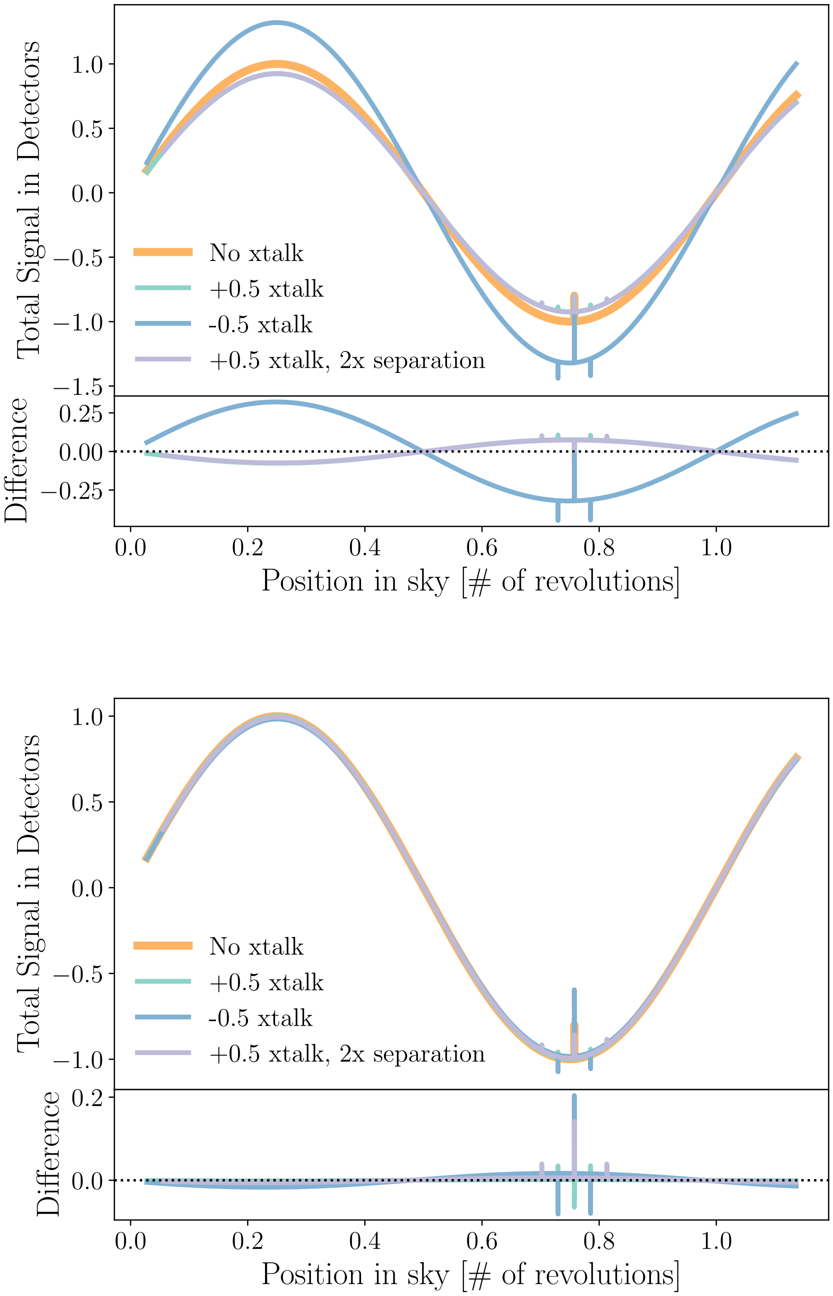

Figure 4.2: Average pointing-corrected TODs for a two detector instrument sharing several crosstalk configurations. TODs have been renormalized by calibrating a point source to an expected amplitude (*top*) and scaling by the amplitude of the TODs' total variance (*bottom*). Scaling with respect to the total variance of the TODs is more effective at recovering the true large-scale features in the sky signal, but it may greatly amplify the phantom signals generated by crosstalk from small-scale features in the sky signal.

| | |
|---|---|
| $N_{\text{det}}$ | 36 |
| Duration of observation | 90 days |
| Observing frequency | 140 GHz |
| Detector sample rate | 22 Hz |
| Satellite spin period | 10 min |
| Initial spin angle | 30° |
| Satellite precession period | 10 min |
| Initial precession angle | 65° |
| Half wave plate rotation frequency | 88 rpm |
| Detector beam size | 23.7 arcmin |
| $f_{\text{knee}}$ | 50 mHz |
| $\alpha$ | 1 |
| $\text{NET}_{\text{array}}$ | $2.1\ \mu\text{K}/\sqrt{\text{s}}$ |
| $f_{\text{min}}$ | 0 Hz |
| Output map NSIDE | 1024 |

Table 4.1: Parameters used to simulate LiteBIRD-like satellite observations in `TOAST`.

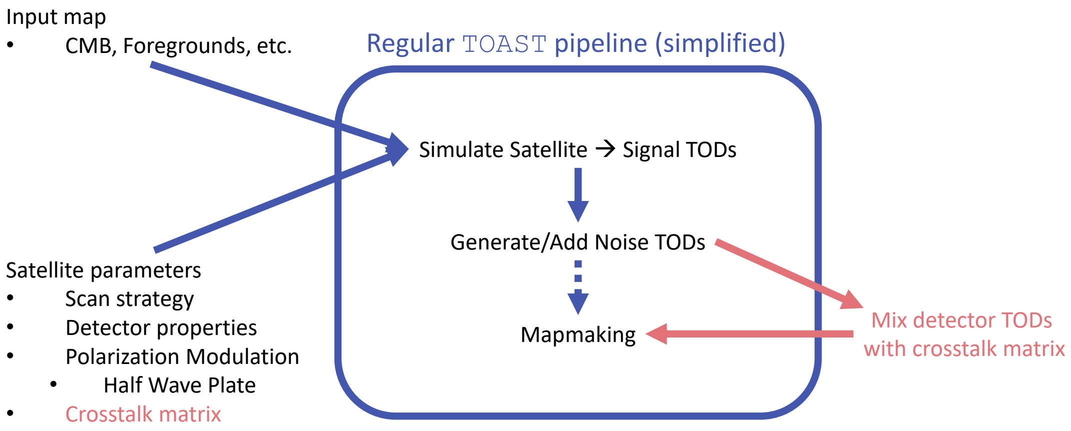

Figure 4.3: Flowchart of the `TOAST` pipeline and its inputs/parameters. Modifications to include crosstalk in the simulations are highlighted in red.

2. Sky signal TODs are generated by simulating the satellite's scanning pattern over the input CMB map. Detector orientations on the focal plane are used to determine which sky pixels are observed by each detector at a given time.
3. Individual detector noise TODs are generated with the noise parameters in the experiment data object. These are then added onto the sky signal TODs for each detector.
4. The full set of detector TODs is transformed into map space using the `MADAM` mapmaker (Keihänen et al., 2005). The result is a set of CMB temperature and polarization maps as they would have been observed by the satellite.

The full flow of the `TOAST` simulation procedure is, of course, more complicated, but these are the most relevant steps for the crosstalk implementation. The implementation of crosstalk to `TOAST` as indicated in Figure 4.3, occurs just before Step 4. Each detector TOD in the simulation contributes some crosstalk to every other detector in the simulation according to Equation (4.1). Two example crosstalk matrices are shown in Figure 4.4, generated by Montgomery et al. (2022). They show that realistic electrical crosstalk levels between detectors are strongest for those closest to each other in frequency scheduling, and additional contributions quickly drop off for detectors farther apart in frequency space. The sign and amplitudes of these crosstalk weights are dependent on the design of the detector readout electronics, as well as choices in the spacing between detectors in the frequency multiplexing. A logarithmic frequency spacing is currently used at the South Pole Telescope (Montgomery et al., 2022), and a linear frequency spacing was recommended for the LiteBIRD mission at the time this work was conducted due to the overall lower crosstalk levels. As a final note, all simulations in this chapter were generated with the same random seed to ensure reproducibility, as well as genuine comparisons between crosstalk configurations.

## 4.4 Impact of detector crosstalk on CMB power spectra

The impact of detector crosstalk on CMB observables may not be readily apparent at the map level; however, the effects are clear in comparisons of CMB power spectra. Figure 4.5 shows $TT$ power spectra for simulations following the procedure outlined in §4.3 using the parameters listed in Table 4.1, along with the application of the first $36 \times 36$ entries of the crosstalk matrix with *linear* frequency spacing in Figure 4.4.

The fractional comparisons in the bottom panel of Figure 4.5 validate our conclusions from §4.2. The introduction of crosstalk generally produces a constant offset in the observed CMB power at large angular scales ($\ell \lesssim 500$), and the magnitude of the offset is dependent on the crosstalk configuration. This change is at the level of less than a few percent, and as discussed in §4.2, they can be mitigated with a proper TOD renormalization strategy (which has *not* been applied here). There is also an $\ell$-dependent increase in the observed CMB temperature power at small angular scales ($\ell \gtrsim 500$) that is reminiscent of the effects of weak gravitational lensing at small scales. As a reminder, these simulations did not carry out the step in which the crosstalk matrix is typically inverted to recover the true detector TODs. The effects of crosstalk in Figure 4.5 serve as an upper limit on the worst-case scenario for crosstalk systematics on realistic CMB observations.

The crosstalk matrices in Figure 4.4 indicate that detectors experience the strongest crosstalk from neighbouring detectors in the frequency schedule. It is then worth considering which detectors

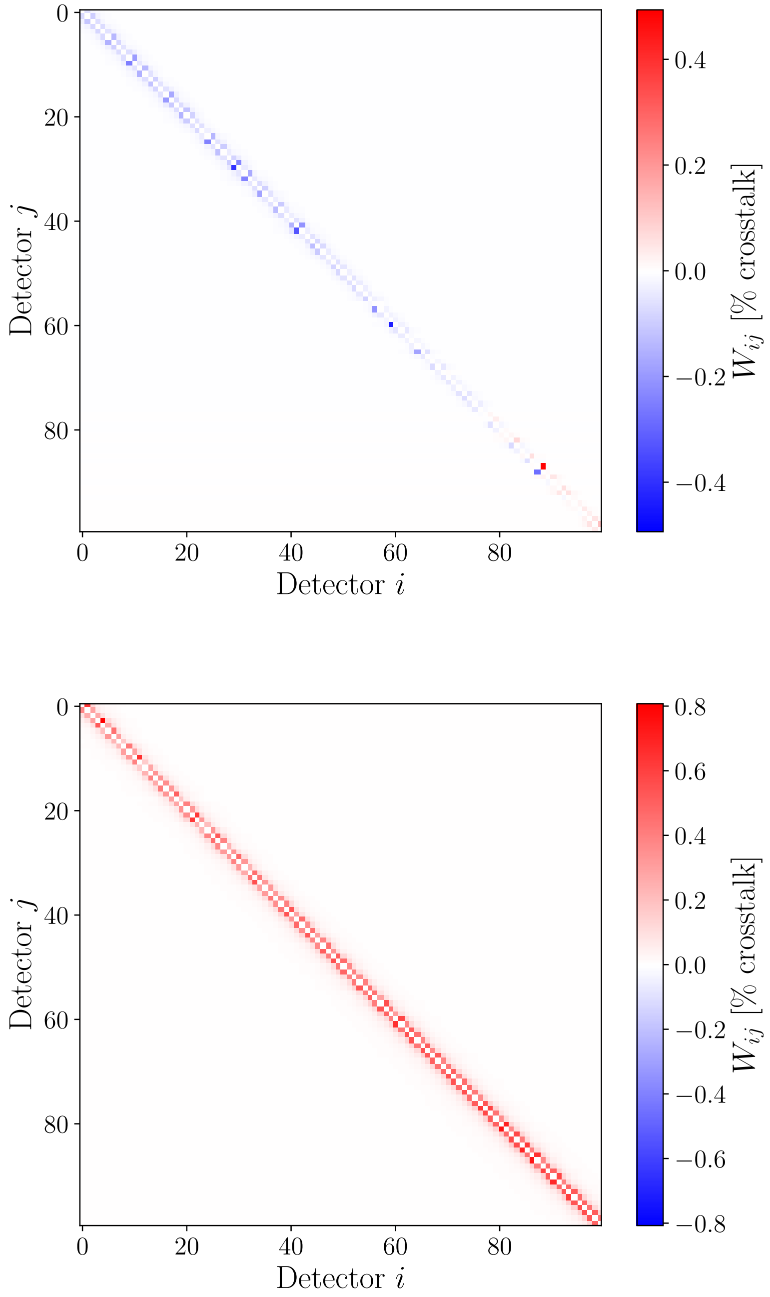

Figure 4.4: Two examples of realistic electrical crosstalk matrices with **linear** spacing in the frequency scheduling (*top*) and **logarithmic** spacing in the frequency scheduling (*bottom*). These representative crosstalk configurations were simulated by Montgomery et al. (2022).

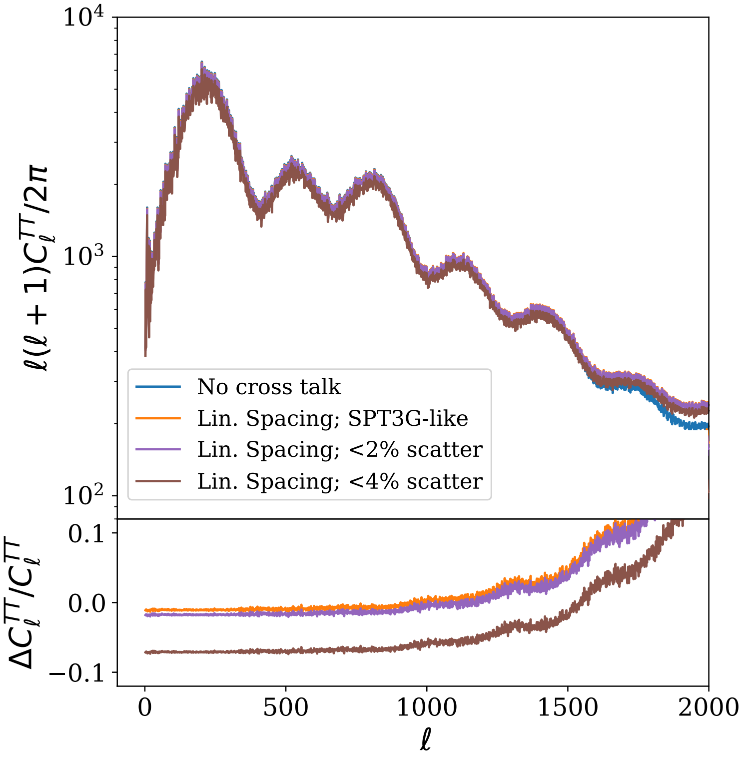

Figure 4.5: Power spectra of CMB simulations with crosstalk compared to a realization without crosstalk. The fractional difference compared to version without crosstalk is shown in the bottom panel, where a constant offset at multiples $\ell \lesssim 500$ can be held to less than a few percent.

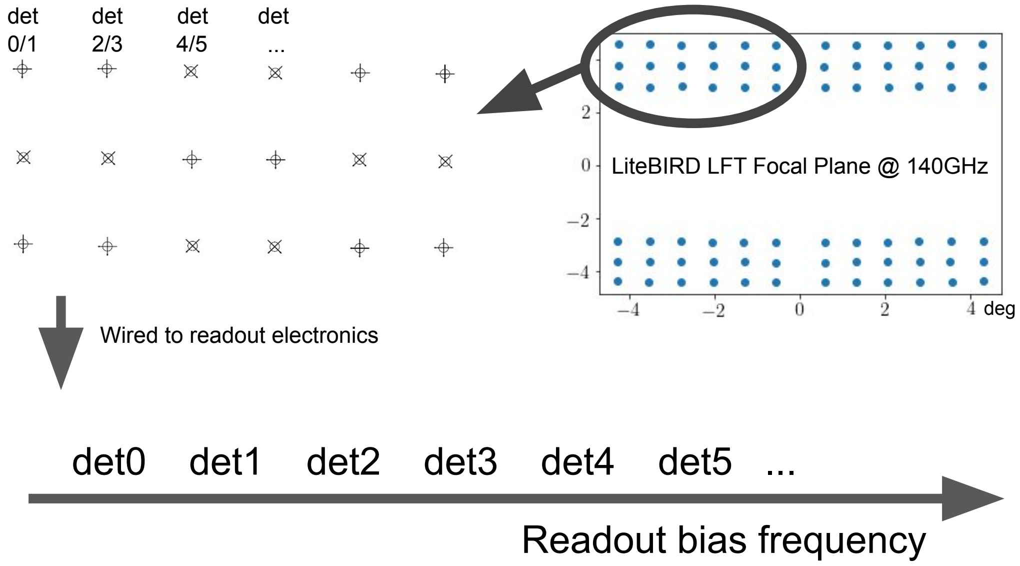

Figure 4.6: Illustration of a proposed LiteBIRD focal plane configuration. An example of a simple choice in detector readout ordering is also shown.

are allowed to crosstalk to each other the most. In other words, one has the freedom to decide which detectors crosstalk to each other when designing an instrument. A straightforward configuration can be to simply order detectors as they are structured physically on the focal plane, as illustrated in an example with a proposed LiteBIRD focal plane in Figure 4.6. In this scenario, a given detector experiences the strongest crosstalk with a detector observing the perpendicular polarization at the same part of the sky, as well as another detector pointed slightly away. The simulated CMB power spectra shown in Figure 4.5 follow this readout configuration.

There are many choices for mapping the detectors from the physical focal plane layout onto the readout frequency schedule. A few simpler choices are shown in Figure 4.7 to illustrate the impact of the readout order. A simpler crosstalk matrix is used here with consistent crosstalk amplitudes such that the strongest crosstalk is $-0.1\%$ for the nearest neighbouring detectors in the frequency schedule, and exponentially decaying amplitudes for detectors farther away in the frequency schedule.

The fractional changes in the observed CMB power spectra are shown for each readout configuration in Figure 4.7. In every scenario, the $TT$ power experiences a similar decrement of low-$\ell$ power due to negative crosstalk seen previously. The effects of different crosstalk readout ordering on the $EE$ and $BB$ power are less clear, but both polarization power spectra appear to experience the same changes. This is likely due to the fact that detectors are crosstalking linear polarization information, which is then converted to $EE$ or $BB$ power at a later stage. It appears that the choice of readout ordering can be a bit nuanced when optimizing for lower effects on observed polarization power spectra, and the random ordering of detectors in Figure 4.7 has the best performance.

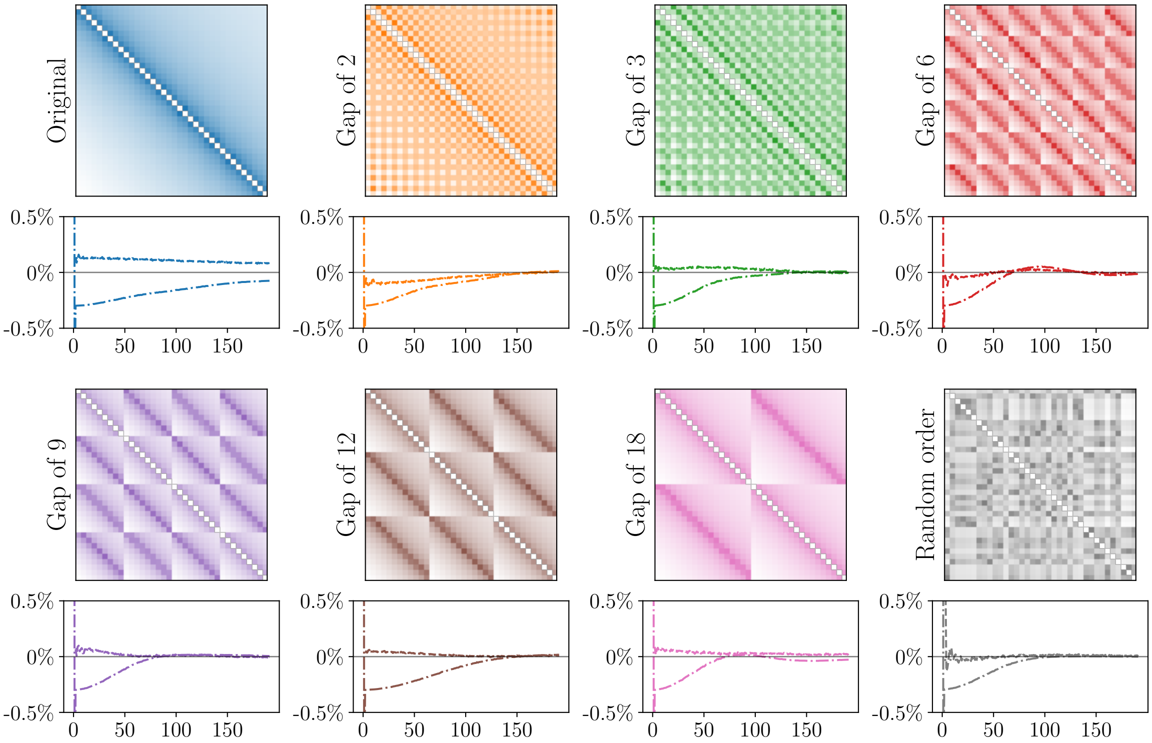

Figure 4.7: A comparison of various crosstalk configurations based on readout ordering and their fractional effects on simulated CMB power spectra. The most significant crosstalk amplitude between pairs of detectors is $-0.1\%$, and the matrix entries are shown on a logarithmic scale for better visibility. Fractional change in observed power spectra compared to a crosstalk free simulation are shown below each matrix for multipoles $2 < \ell < 200$. $EE$ spectra are shown as dashed lines, $BB$ spectra are shown as dotted lines, and $TT$ spectra are shown as dot-dashed lines. The $EE$ and $BB$ spectra overlap in every scenario. Changing the readout ordering of detectors can have a non-trivial impact on the observed CMB power spectra.

## 4.5 Discussion and Conclusions

This chapter is summarized by the following key points:

1. Detector crosstalk generally amplify or suppress CMB power consistently at large angular scales, and generate power at small angular scales. The effects at large angular scales can be mitigated by renormalizing detector TODs by the total RMS or variance of the expected sky signal.

2. We created an implementation of crosstalk in `TOAST` which has since been used as the basis for a module of the package.

3. Full simulations of the CMB observation procedure with realistic crosstalk matrices show that the effects of *unprocessed* crosstalk on CMB power spectra will be, at worst, on the level of a few percent at large angular scales for a LiteBIRD-like experiment.

4. The ordering of detectors in the frequency schedule can have a non-trivial impact on CMB observables, and a random ordering of detectors appears to be a good choice for minimizing the effects of crosstalk on observed polarization $EE$ and $BB$ power spectra. This investigation in particular was continued by graduate student Kolen Cheung at the University of California, Berkeley, and his findings are summarized in an unpublished (at the time of writing) presentation[6].

---

[6] `https://ickc.github.io/LiteBIRD-R9X4K4CNHPSE2CWBH3AFAXV64BO5WACYPSBJ6CDOGOSW/crosstalk/intermediate-summary-2/`

# Chapter 5

# The Small Correlated Against Large Estimator for CMB lensing

*Originally submitted and accepted for publication in Phys. Rev. D. Pre-print available as: "The Small-Correlated-Against-Large-Estimator (SCALE) for the Lensing of the Cosmic Microwave Background" Chan, Victor C., Hložek, Renée, Meyers, Joel, & van Engelen, Alexander 2023, arXiv e-prints, arXiv:2302.13350, doi: 10.48550/arXiv.2302.13350. The text from the accepted manuscript has been adapted for this dissertation with some minor clarifications, and a majority of the introduction has been moved to Chapter 1 (see §1.2.4). The original appendix of the accepted manuscript is presented in Appendix B.*

## Abstract

Weak gravitational lensing of the cosmic microwave background (CMB) carries imprints of the physics operating at redshifts much lower than that of recombination and serves as an important probe of cosmological structure formation, dark matter physics, and the mass of neutrinos. Reconstruction of the CMB lensing deflection field through use of quadratic estimators has proven successful with existing data but is known to be suboptimal on small angular scales ($\ell > 3000$) for experiments with low noise levels. Future experiments will provide better observations in this regime, but these techniques will remain statistically limited by their approximations. We show that correlations between fluctuations of the large-scale temperature gradient power of the CMB sourced by $\ell < 2000$, and fluctuations to the local small-scale temperature power reveal a lensing signal which is prominent in even the real-space pixel statistics across a CMB temperature map. We present the development of the Small Correlated Against Large Estimator (SCALE), a novel estimator for the CMB lensing spectrum which offers promising complementary analysis alongside other reconstruction techniques in this regime. The SCALE method computes correlations between both the large/small-scale temperature gradient power in harmonic space, and it is able to quantitatively recover unbiased statistics of the CMB lensing field without the need for map-level reconstruction. SCALE can outperform quadratic estimator signal-to-noise by a factor of up to 1.5 in current and upcoming experiments for CMB lensing power spectra $C^{\phi\phi}_{6000<L<8000}$.

## 5.1 Introduction

GRAVITATIONAL lensing of the cosmic microwave background (CMB) are the dominant signal in the observed CMB at small angular scales. Our goal in this chapter is to devise a simple estimator that can leverage the low-noise and high-resolution maps expected from future CMB surveys to measure the small-scale lensing power spectrum. We present the Small Correlated Against Large Estimator (SCALE), a new method of obtaining the small-scale lensing power constructed from the cross-correlation between maps of the local large-scale and small-scale temperature power. It was previously shown in Ref. Zaldarriaga (2000) how cross-correlating the large scale temperature gradient with the small-scale temperature power can be used to estimate the lensing power. We discuss some key differences between SCALE and the method used in Ref. Zaldarriaga (2000) in §5.4. This estimator is complementary to reconstruction techniques aimed at estimating a map of the lensing potential. It is similar in spirit to the maximum likelihood, maximum a posteriori, Gradient Inversion, and Bayesian techniques (Hirata & Seljak, 2003b; Carron & Lewis, 2017; Hadzhiyska et al., 2019; Millea et al., 2020) in that it aims to make optimal use of lensing information at small-scales of a CMB temperature map. In contrast to the QE method, SCALE is designed to work on small angular scales, which leverages the ongoing improvements to detectors and telescopes in the coming decade (Abazajian et al., 2016; Ade et al., 2019; Hanany et al., 2019; Abazajian et al., 2019; Sehgal et al., 2019). SCALE specifically aims to avoid the extra variance incurred by QE techniques due to cosmic variance of the large-scale CMB temperature gradient, while also circumventing the highly correlated nature of QE errors at small angular scales. In contrast to the gradient inversion method described above, the SCALE pipeline consists of high-pass and low-pass filtered maps which are squared and then cross-correlated to estimate the lensing power spectrum directly, rather than a map of the lensing potential.

We start with a brief review of CMB lensing in §5.2 and develop a simple test in real space to illustrate the principles of our proposed method in §5.3. We further develop this method, and present the SCALE procedure in §5.4. After introducing our data simulations in §5.5, we present our results in §5.6 and conclude in 5.7.

## 5.2 Review of CMB lensing

In the absence of foregrounds and noise, the observed CMB temperature field $\tilde{T}$ at a particular line of sight $\hat{\mathbf{n}}$ is the unlensed temperature $T$ at a lensing deflection angle $\mathbf{d}(\hat{\mathbf{n}})$ away from the line of sight. The lensing deflection angle $\mathbf{d} = \boldsymbol{\nabla}\phi$ is the gradient of the lensing potential $\phi$ when working within the Born approximation. We denote a gradient across the sky (i.e. along the plane perpendicular to the line of sight) with $\boldsymbol{\nabla}$. The lensed temperature is

$$\tilde{T}(\hat{\mathbf{n}}) = T(\hat{\mathbf{n}} + \boldsymbol{\nabla}\phi(\hat{\mathbf{n}})) = T(\hat{\mathbf{n}}) + \boldsymbol{\nabla}\phi(\hat{\mathbf{n}}) \cdot \boldsymbol{\nabla}T(\hat{\mathbf{n}}) + \ldots\,. \tag{5.1}$$

The lensing potential is directly related to the lensing convergence $\kappa = -\nabla^2\phi/2$, with power spectra related by $C_L^{\kappa\kappa} = (L(L+1))^2 C_L^{\phi\phi}/4$. In our conventions, the Fourier transform of the temperature gradient is

$$\boldsymbol{\nabla}T(\hat{\mathbf{n}}) = i\int \frac{\mathrm{d}^2\boldsymbol{\ell}}{2\pi}\boldsymbol{\ell}T(\boldsymbol{\ell})e^{i\boldsymbol{\ell}\cdot\hat{\mathbf{n}}}\,. \tag{5.2}$$

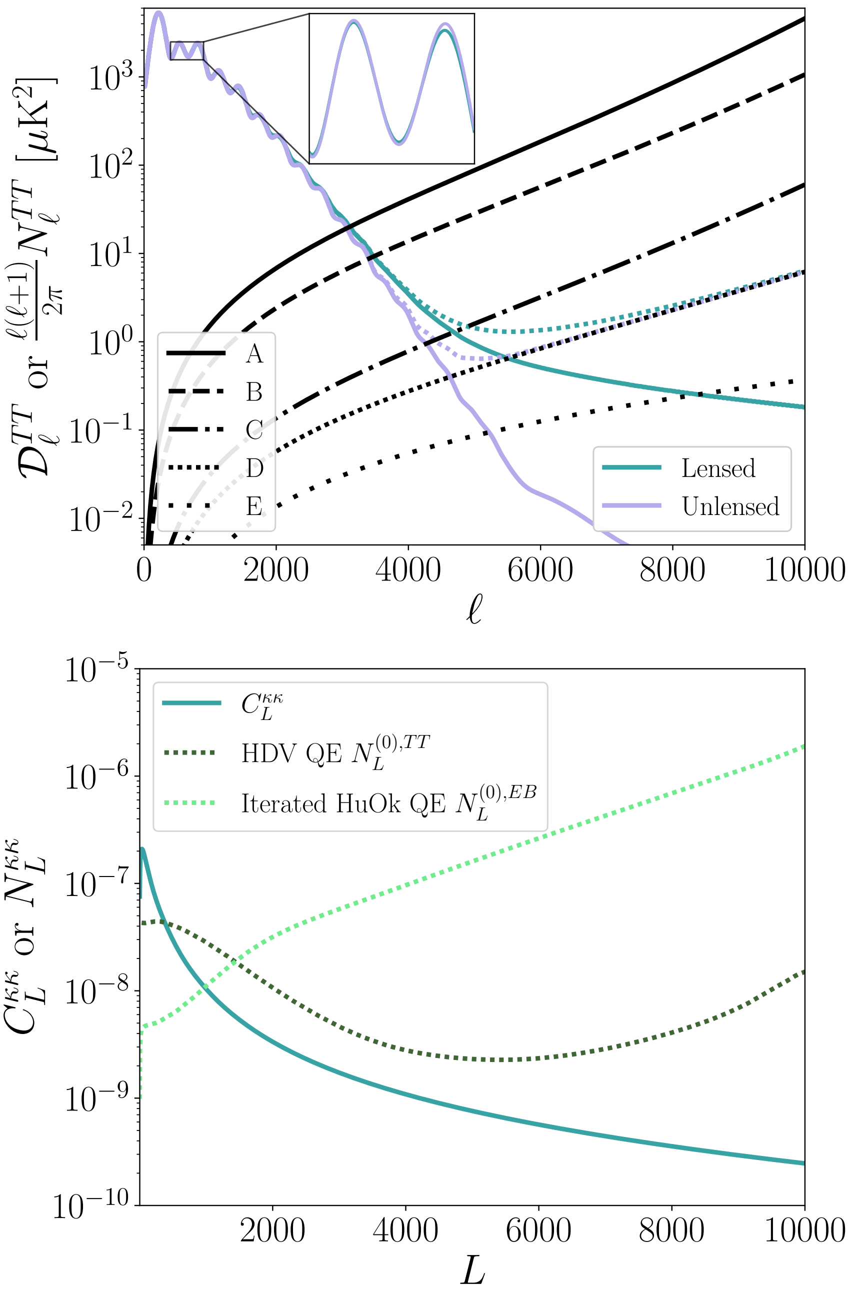

Figure 5.1: *Top*: The input power spectrum for simulated raw CMB maps is shown in solid purple. The resulting lensed CMB power spectrum after applying a lensing field corresponding to the spectrum below is shown in solid teal. Noise spectra corresponding to different experiments are shown in black, and the lensed/unlensed spectra combined with configuration-D (defined below, in Table 5.2) noise are shown in their respective colours and dense dotted lines. *Bottom*: The input lensing convergence power spectrum for simulated lensing potential fields. Also shown is the optimal reconstruction noise $N_L^{(0)}$ for the Hu, DeDeo & Vale $TT$ (HDV, Hu et al. (2007)) and iterative Hu & Okamoto $EB$ quadratic estimators (HuOk, Hu & Okamoto (2002); Okamoto & Hu (2003)) computed with noise from configuration-D.

Note that the majority of the CMB temperature gradient comes from modes with $\ell \lesssim 2000$ (Hu et al., 2007). Taking Equation (5.1) into Fourier space, we apply Equation (5.2) in combination with the convolution theorem to get

$$\begin{aligned}\tilde{T}(\boldsymbol{\ell}) &= \int \mathrm{d}\hat{\mathbf{n}}\tilde{T}(\hat{\mathbf{n}})e^{-i\boldsymbol{\ell}\cdot\hat{\mathbf{n}}} \\ &= T(\boldsymbol{\ell}) - \int \frac{\mathrm{d}^2\boldsymbol{\ell}'}{2\pi}\boldsymbol{\ell}'\cdot(\boldsymbol{\ell}-\boldsymbol{\ell}')\phi(\boldsymbol{\ell}-\boldsymbol{\ell}')T(\boldsymbol{\ell}') + \mathcal{O}(\phi^2)\end{aligned} \tag{5.3}$$

In Equation (5.3), we see that at first order in $\phi$ the lensed temperature field $\tilde{T}$ is a convolution between the lensing potential field $\phi$ and the original unlensed temperature field $T$. Taking the two-point auto-correlation of the temperature field (e.g., steps (4.7) - (4.11) in Lewis & Challinor (2006)) yields the lensed power spectrum:

$$\begin{aligned}\tilde{C}_\ell^{TT} \approx& \left(1 - \int \frac{\mathrm{d}^2\boldsymbol{\ell}'}{(2\pi)^2}C_{\ell'}^{\phi\phi}(\boldsymbol{\ell}\cdot\boldsymbol{\ell}')^2\right)C_\ell^{TT} \\ &+ \int \frac{\mathrm{d}^2\boldsymbol{\ell}'}{(2\pi)^2}[\boldsymbol{\ell}'\cdot(\boldsymbol{\ell}-\boldsymbol{\ell}')]^2 C_{\ell'}^{TT}C_{|\boldsymbol{\ell}-\boldsymbol{\ell}'|}^{\phi\phi}\,.\end{aligned} \tag{5.4}$$

Note that the first term is a cross-term between the zeroth and second order terms of Equation (5.3), and the second term is a product of the first order term with itself. Similar to Equation (5.3), it expresses that the lensed CMB temperature power contains a convolution between the lensing potential power and the original CMB temperature power. The effects of weak gravitational lensing on the CMB do not add or remove from the total CMB temperature variance $\int \mathrm{d}\ell\ell C_\ell^{TT}/2\pi$ across the sky. Instead, lensing serves to redistribute power $C_\ell^{TT}$ between angular modes $\ell$ in a way that "smooths out" the peaks and troughs in the observed power spectrum (as can be seen in Figure 5.1); the power redistributed to scales $\ell \gtrsim 4000$ dominates the signal compared to the unlensed temperature modes which are suppressed by diffusion damping. Traditional estimators of the lensing potential take advantage of the correlations between angular modes that have been introduced, and they work to reconstruct the lensing potential field through measurement of these off-diagonal couplings. We can make approximations to simplify Equation (5.4) in the small-scale limit $\ell \gg 2000$. The CMB temperature gradient variance, which we denote $\langle|\boldsymbol{\nabla}T_L|^2\rangle$, is made up of an integral over the larger scale modes $C_{\ell\lesssim 2000}^{TT}$ of the original CMB temperature field, given by $\int \mathrm{d}\ell\ell^2 C_\ell^{TT}/2\pi$. This background temperature gradient is approximately constant at small scales, which can be enforced with $\ell' \ll \ell$ in Equation (5.4). One can apply these approximations to arrive at a simplified representation of the lensed CMB temperature power on small scales (e.g., §4.1.3 of Lewis & Challinor (2006)):

$$\begin{aligned}\tilde{C}_{\ell\gg 2000}^{TT} &\approx \ell^2 C_\ell^{\phi\phi}\int \frac{\mathrm{d}\ell'}{\ell'}\frac{\ell'^4 C_{\ell'}^{TT}}{4\pi} + C_{\ell,r}^{TT} \\ &= \frac{1}{2}\langle|\boldsymbol{\nabla}T_L|^2\rangle\ell^2 C_\ell^{\phi\phi} + C_{\ell,r}^{TT}\,.\end{aligned} \tag{5.5}$$

Here, we define $C_{\ell,r}^{TT}$ which represents all remaining contributions to the observed temperature power which are not expected to strongly correlate with the large-scale gradient of the CMB temperature field. This includes the first term of Equation (5.4) which contains a small amount of the unlensed CMB temperature power suppressed by diffusion damping crossed with a second-order lensing contribution. We may also include contributions from instrument noise, foregrounds, and other secondaries in $C^{TT}_{\ell,r}$.

A straightforward method to estimate the small-scale lensing power is to simply divide the observed excess small-scale temperature power by the average unlensed temperature gradient power on large scales. That is, we can rework Equation (5.5) and estimate the small scale lensing power spectrum as

$$C^{\phi\phi}_{\ell} \approx \frac{\tilde{C}^{TT}_{\ell} - C^{TT}_{\ell,r}}{\ell^2 \left(\frac{1}{2}\left\langle |\boldsymbol{\nabla} T|^2 \right\rangle\right)}\,, \tag{5.6}$$

for $\ell \gg 2000$. The motivation for our SCALE technique is that we can do better than Equation (5.6), even without reconstructing a map of the lensing field. In any given patch of sky, the large-scale temperature gradient power around the line-of-sight $\hat{\mathbf{n}}$ will deviate from the sky average due to random fluctuations. As a consequence, the *local* small-scale temperature power that results from lensing will also deviate from the sky average. By correlating the spatial variations in the locally measured large-scale temperature gradient power with the spatial variations in the small-scale temperature power, we can construct an improved estimate of the small-scale lensing power. Furthermore, variations in the observed small-scale temperature power that are due to sources other than lensing (such as non-stationary noise or astrophysical foregrounds), are not expected to correlate with variations in the large-scale temperature gradient power since these effects result from survey choices or local physics unrelated to the long wavelength fluctuations responsible for the large-scale temperature gradients.

In summary, we propose a new lensing estimator with a similar form to Equation (5.6). The key difference is allowing the *local* small-scale lensed temperature power to fluctuate according to the steepness of the background temperature gradient in the same part of the sky:

$$\tilde{C}^{TT,\mathrm{local}}_{\ell\gg 2000}(\hat{\mathbf{n}}) \approx \frac{1}{2}\left|\boldsymbol{\nabla} T_L(\hat{\mathbf{n}})\right|^2 \ell^2 C^{\phi\phi}_{\ell} + C^{TT}_{\ell,r}\,. \tag{5.7}$$

Equation (5.6) is recovered by taking the sky average of this version. We consider combinations of CMB temperature maps because temperature-based lensing reconstruction out-performs polarization-based estimators on small angular scales, due to the fact that polarization maps become dominated by noise at these angular scales (see Figure 5.1). The lensing estimator we propose shares some similarities with techniques used to measure the kinetic Sunyaev-Zel'dovich (kSZ) effect through variations in small-scale temperature power (Smith & Ferraro, 2017).

In the following section, we present a simple proof-of-concept that takes advantage of the **local map-space correlations between the small-scale temperature power and the square of the observed temperature gradient amplitude** in order to tease out the statistics of the underlying lensing potential field.

## 5.3 Introductory concepts in Real Space

Before describing the harmonic-space implementation of SCALE in this work, we start with an illustrative real-space description of the concept to provide the intuition and motivation for the techniques we develop in the following section. We introduce the notation $\overline{X}(\hat{\mathbf{n}})$ along with shorthand $\overline{X}$ to indicate the *local* average of a quantity $X$ near a line-of-sight $\hat{\mathbf{n}}$.

The total observed small-scale CMB power along a given line of sight $\overline{T_S^2}(\hat{\mathbf{n}})$ is given by the 2-dimensional (2D) angular integral over Equation (5.7). We also specify the large-scale temperature gradient amplitude $\overline{|\boldsymbol{\nabla} T_L(\hat{\mathbf{n}})|}$ to still be approximately constant near each single line of sight while allowing for the small fluctuations between different lines of sight as described in the previous section. Since we do not expect the lensing potential field to strongly correlate with the remaining contributions to the small-scale temperature power, we generally expect

$$\overline{T_S^2}(\hat{\mathbf{n}}) \simeq a_1 \overline{|\boldsymbol{\nabla} T_L(\hat{\mathbf{n}})|}^2 + a_0, \tag{5.8}$$

where we have the following contributions:

1. A term containing the lensing contribution which scales with the amplitude of the large-scale CMB temperature gradient $\overline{|\boldsymbol{\nabla} T_L(\hat{\mathbf{n}})|}^2$, and
2. a term containing all remaining contributions which *do not* scale with the CMB temperature gradient.

Equation (5.8) motivates a simple, map-space approach to gather local statistics for $\overline{T_S^2}(\hat{\mathbf{n}})$ and $\overline{|\boldsymbol{\nabla} T_L(\hat{\mathbf{n}})|}^2$ and to take advantage of their correlations to bring out the lensing signal. The expectation is that $a_1 \to 0$ in a CMB temperature map *without* lensing, and $a_1$ should increase with a stronger lensing signal (i.e., $a_1 \propto C_L^{\phi\phi}$, c.f., 5.7.). Contributions to the temperature map such as noise and foregrounds that do not come from lensing should directly contribute to $a_0$, but not $a_1$ because they are not expected to correlate with the large-scale CMB temperature gradient. We should therefore be able to infer the small-scale lensing power from the measured value of $a_1$.

### 5.3.1 Map reduction to local patches

Of the two observable quantities in Equation (5.8), we begin with measuring background temperature gradient $\overline{\boldsymbol{\nabla} T_L}$ as well as the small-scale temperature $T_S$ from a single input CMB temperature map. Local statistics of $\overline{|\boldsymbol{\nabla} T_L|}^2$ and $\overline{T_S^2}$ then need to be gathered in small cutouts of the observed field.

To ensure that we are only including the smooth component of the temperature gradient, we filter the maps in Fourier space. We compute maps of the observed CMB temperature gradient using

$$\boldsymbol{\nabla} T_L(\boldsymbol{\ell}) = i\boldsymbol{\ell} T(\boldsymbol{\ell}). \tag{5.9}$$

We apply a low-pass top hat filter before returning the map to pixel space: $\boldsymbol{\nabla} T_L(\ell > \ell_{\nabla T}) = 0$.

It is not immediately obvious what scale $\ell_{\nabla T}$ should be used in the low-pass filter. Enough modes should be included such that the resulting $\boldsymbol{\nabla} T_L$ maps contain enough information about the background temperature gradient, and $\ell_{\nabla T}$ should be small enough such that there are no direct correlations from similar shared modes between $\boldsymbol{\nabla} T_L$ and the small scale temperature $T(\ell \gg 2000)$. The large-scale CMB temperature gradient $\boldsymbol{\nabla} T_L|$ that we are looking for is mostly constituted by modes $\ell \lesssim 2000$, so we consider a low-pass cutoff at $\ell_{\nabla T} = 3000$. It is important to keep in mind that pixels near the boundary of the $\boldsymbol{\nabla} T_L$ map may be unusable if the original $T$ map does not have the appropriate repeating boundary conditions.

In order to pick out only the small scale lensing power contributions which we wish to correlate with the background temperature gradient, the original CMB temperature map also needs to be

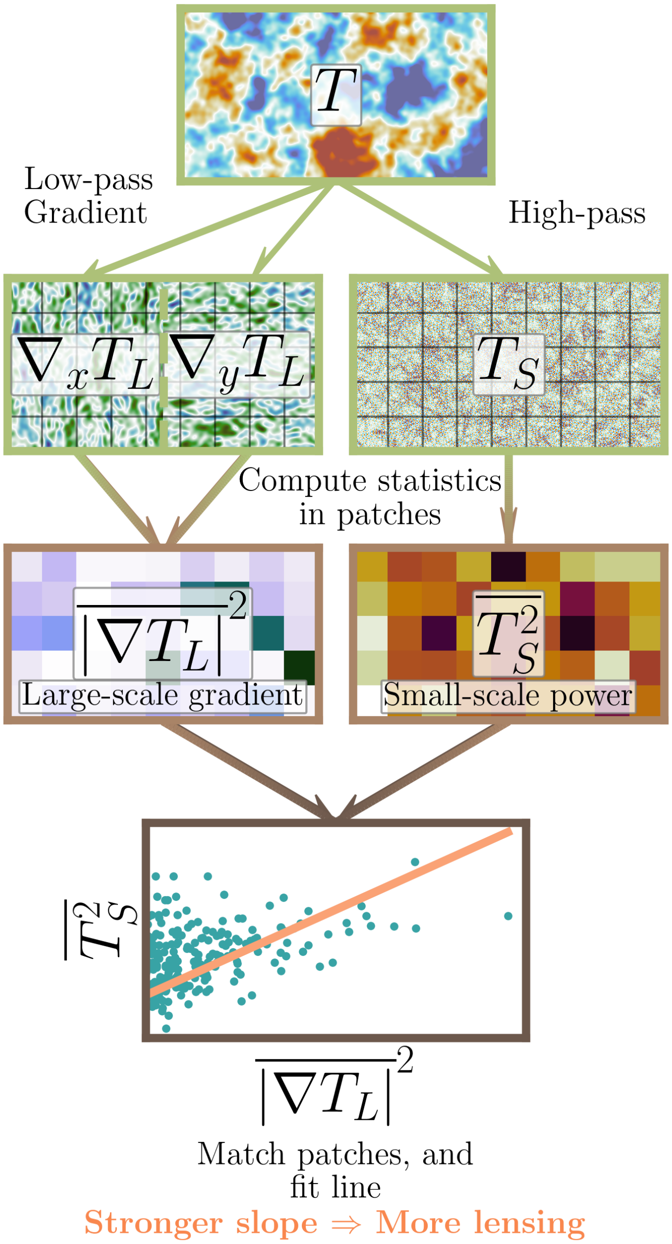

Figure 5.2: Schematic of a procedure to reduce a CMB temperature map to local, real-space statistics which can then be correlated to infer lensing effects. Maps in the original resolution are denoted in *green* borders, and the *light brown* borders indicate a degraded resolution after computing relevant statistics within local patches of width $40'$. Note that the patches shown here are for visual presentation, and they are larger than those chosen later on.

filtered, this time with a high-pass window ($\ell_{\min} < \ell < \ell_{\max}$) for the scales relevant to the analysis. A variety of choices for $\ell_{\min}$ and $\ell_{\max}$ can be made as long as scales are in a regime where lensing is dominant ($\ell \gg 2000$), and we require $\ell_{\min} > \ell_{\nabla T}$ to ensure that any correlations between the large-scale gradient and small-scale power fields is strictly from lensing. The choice in $\ell_{\min}$ and $\ell_{\max}$ affect the expected $a_1$ and $a_0$ in Equation (5.8) through the angular integral for the total CMB temperature power.

We then split both filtered maps into patches in order to estimate the quantities $\overline{|\boldsymbol{\nabla} T_L|}^2$ and $\overline{T_S^2}$ along different lines of sight $\hat{\mathbf{n}}$. The required patch size is also not obvious, but needs to be small such that the assumption of $\overline{|\boldsymbol{\nabla} T_L|}^2$ being constant within a patch is reasonable. In considering the characteristic angular scale $\theta \sim 2\pi/\ell$ for $\ell = 2000$, the patches should be $\lesssim 10'$ wide. The patches also need to be large enough such that there are sufficient pixels of the map within each patch to make good estimates of $\overline{|\boldsymbol{\nabla} T_L|}^2$ and $\overline{T_S^2}$. This is also dependent on the resolution of the original map, which will be discussed further in § 5.5. For the example shown here, we choose a patch size of $10' \times 10'$, which is $20 \times 20$ pixels in our simulated maps with resolution $0.5'$.

The quantity $\overline{|\boldsymbol{\nabla} T_L(\hat{\mathbf{n}})|}^2$ can now be computed for every patch across the map. For each of the two perpendicular directions on the map $\hat{x}$ and $\hat{y}$, we first compute the average gradient across each patch $\overline{\boldsymbol{\nabla} T_L}(\hat{\mathbf{n}}) = \overline{\nabla_x T_L}\hat{x} + \overline{\nabla_y T_L}\hat{y}$. We can then readily compute $\overline{|\boldsymbol{\nabla} T_L(\hat{\mathbf{n}})|}^2 = (\overline{\nabla_x T_L})^2 + (\overline{\nabla_y T_L})^2$ for each patch.

### 5.3.2 Lensing from patch statistics

The remaining quantity to compute within each patch is $\overline{T_S^2}(\hat{\mathbf{n}})$. We compute the auto-variance of the high-pass filtered $T_S$ map within each patch. The $\overline{|\boldsymbol{\nabla} T_L|}^2$ and $\overline{T_S^2}$ from each patch, in principle, provides a very noisy and approximate estimate of the overall small-scale lensing power present in the map, following a distribution about Equation (5.8). By combining the statistics of many noisy patches across the map, there is opportunity to more rigorously quantify the slope in Equation (5.8), and relate it to the 2D-angular integral of Equation (5.5) in the appropriate space of chosen $\ell_{\nabla T}$, $\ell_{\min}$, and $\ell_{\max}$.

The outputs of this real-space procedure applied to both a lensed and unlensed realization of the CMB temperature are shown in Figure 5.3 to illustrate the relationship between them. A positive correlation between $\overline{T_S^2}$ and $\overline{|\boldsymbol{\nabla} T_L|}^2$ can be clearly seen in the ensemble of lensed patches, while no significant correlation is seen in the sample of unlensed patches. Two example patches are also highlighted in Figure 5.3, and the corresponding cutouts of the lensed CMB temperature filtered to the relevant scales of the small-scale temperature and the large-scale gradient are shown. The cutouts show that, even upon visual inspection, typical areas on the lensed CMB with a steep background temperature gradient usually have a higher small scale temperature power than typical areas with a relatively weak background temperature gradient.

We briefly consider quantifying the slope of the real-space lensing correlations in Appendix A. There are several challenges that must be overcome in order to use this real-space method as a reliable estimator of CMB lensing. First, it is important to note that the strictly positive nature of the auto-variance forces the observed distribution of $\overline{T_S^2}(\hat{\mathbf{n}})$ across the map to be positively skewed. It is also non-trivial to determine a choice in patch size that optimally includes as many pixels per patch while keeping the large-scale gradient and lensing statistics consistent within each patch. In fact, we show in §5.4 and Figure 5.5 that there is no single patch size which can be chosen to

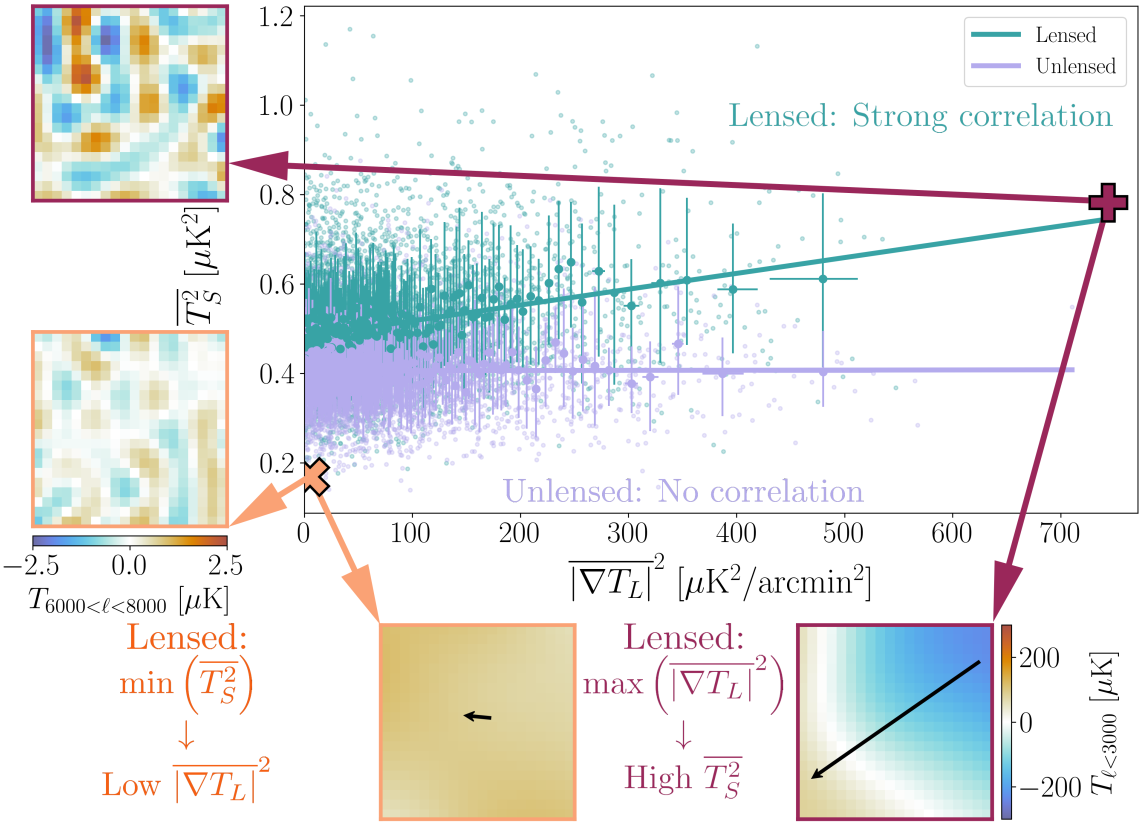

Figure 5.3: *Main*: Local small scale ($6000 < \ell < 8000$) temperature variance $\overline{T_S^2}$ vs. average large scale ($\ell < 3000$) temperature gradient amplitude squared $\overline{|\boldsymbol{\nabla} T_L|}^2$ for lensed (*teal*) and unlensed (*purple*) realizations are shown here as faint, small points. A low-variance/small-gradient patch from the lensed realization (*orange* x), and a high-variance/large-gradient patch from the lensed realization (*maroon* +) are highlighted. The larger points are centered on the medians within bins of $\overline{|\boldsymbol{\nabla} T_L|}^2$ containing an equal number of patches, with error bars corresponding to 68% quantiles. The lines of best fit through the binned points are also shown. *Left*: $10' \times 10'$ cutouts of the lensed CMB temperature map filtered for small scales ($6000 < \ell < 8000$) corresponding to the highlighted patches. *Bottom*: $10' \times 10'$ cutouts of the lensed CMB temperature map filtered for large scales ($\ell < 3000$) corresponding to the highlighted patches. The average gradient direction and relative amplitudes across each patch are shown with the overlaid arrows. Gravitational lensing adds a strong correlation between the local large-scale temperature gradient and the small-scale temperature power that is clear from just map-space statistics.

effectively capture all the correlations between the large and small scale temperature fluctuations. These factors, combined with the fact that the temperatures observed in neighbouring real-space pixels across each patch are highly correlated, suggest that the expected distribution of observed $\overline{T_S^2}(\hat{\mathbf{n}})$ about Equation (5.8) is non-trivial. One option to make this distribution better behaved is to compute the covariance of two observations of the same CMB temperature field. By splitting up time-ordered CMB observations into two or more maps of the same area of sky, one can take advantage of the fact that the maps contain the same CMB realization (which should contain the same lensing information and correlations) and different noise realizations (which should not co-vary across maps).

One more challenge with quantifying this method is the loss of information coming from the local large-scale gradient direction when computing $\overline{|\boldsymbol{\nabla} T_L(\hat{\mathbf{n}})|}^2$. One may choose to construct individual filters for each real-space patch and its observed gradient direction in order to focus on the expected lensing signal(s). Directional filtering is motivated by the fact that lensing induces small-scale temperature gradient fluctuations proportional to the local large-scale gradient amplitude and direction. One such example is choosing a filter $f_\ell = \cos\alpha = \hat{\nabla} \cdot \hat{\ell}$ in addition to the high-pass filter for the small-scale temperature. In practice, this means that the small-scale temperature patches must each be filtered separately and uniquely based on the observed gradient direction in each patch $\hat{\nabla} T_L(\hat{\mathbf{n}})$. This approach once again faces the previous challenge of the large-scale temperature gradient fluctuations not being fully represented within a single patch size. We find that applying such a filter also introduces edge-effects along the borders of each patch, which alters information from an already limited set of pixels within each patch. For these reasons, we choose to present Figure 5.3 without additional gradient information in the filtering of the small-scale temperature field.

While we limit our current presentation of this real-space procedure to a qualitative analysis, it provides significant intuition and motivation for the development of SCALE. Figure 5.3 demonstrates that the small-scale CMB temperature fluctuations are intricately tied to the statistics of the underlying lensing field as well as the large-scale temperature fluctuations of the CMB itself. In other words, information about the lensing field naturally comes out when correlating small-scale CMB temperature fluctuations to large-scale CMB temperature fluctuations. While the observed CMB temperature field is expected to be contaminated by foregrounds and noise, we do not expect such contributions to be strongly correlated between small and large scales. These properties of a lack of noise correlation and the direct lensing correlation between the large-scale gradient and small-scale temperature are central to SCALE. This estimator overcomes the weaknesses of the real-space method, and we will show that it provides a quantitative estimate of the underlying small-scale lensing statistics in the following section.

## 5.4 The SCALE Method

The small-scale lensing signal is reflected in CMB temperature maps through local correlations between the small-scale temperature power fluctuations and the large-scale temperature power fluctuations across the sky. For a given map of the observed CMB temperature field $T(\boldsymbol{\ell})$, we begin by constructing the fields containing the relevant information at small, and large scales. Here, we will sketch the procedure of forming the optimal direct estimate of the lens-induced correlation between small-scale gradient power and large-scale gradient power, in a way that minimizes the variance of

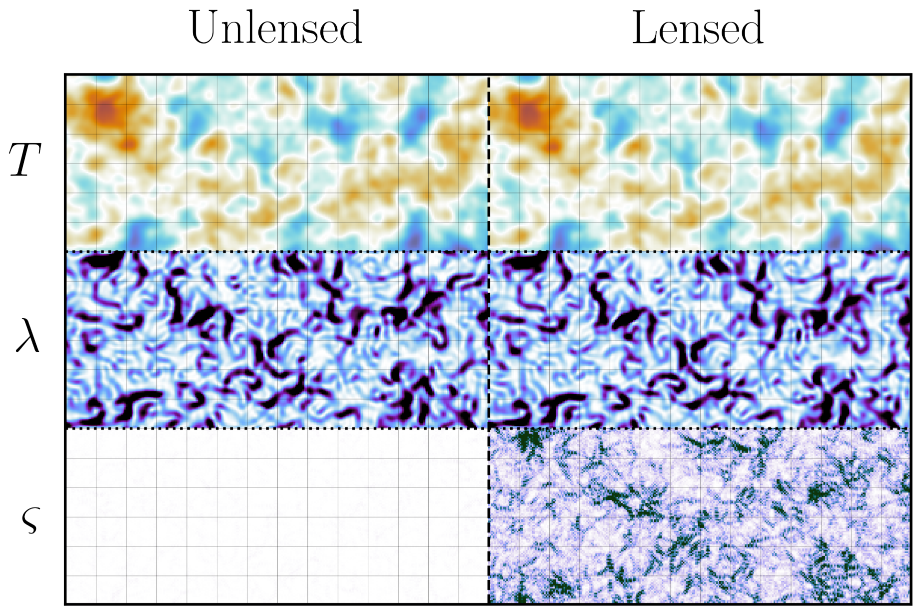

Figure 5.4: A comparison of a CMB temperature realization before and after lensing in the absence of noise and foregrounds. The same area of sky is shown for all panels, including a visualization of the $\lambda$ and $\varsigma$ maps derived from each version. Panels on the same row are shown with the same colormap and limits. A grid with $20'$ spacing is overlaid, which illustrates patches twice the width of the chosen patches for the real-space proof of concept. The $\lambda$ field is largely untouched by the effects of lensing, and lensing generates features in the $\varsigma$ field that are visibly correlated with the $\lambda$ field.

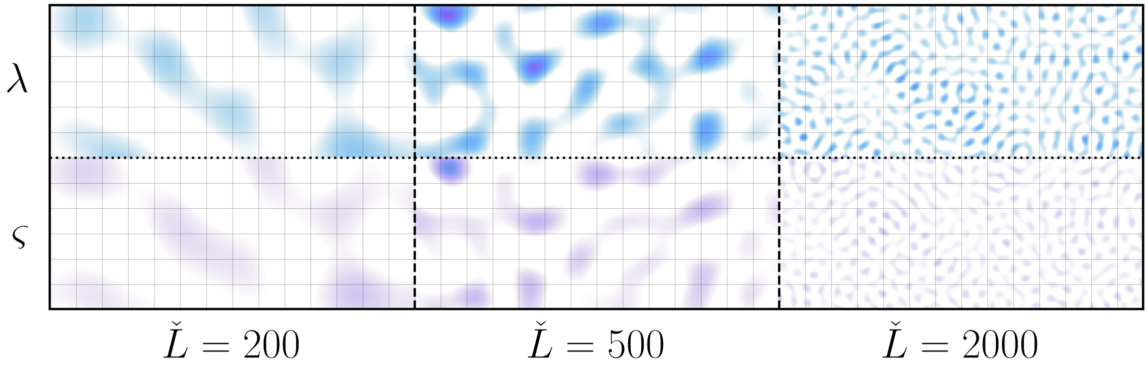

Figure 5.5: A visualization of several $\check{L}$ bands of width $\Delta\check{L} = 300$ centered at the $\check{L}$ labelled highlights the correlations between $\lambda$ and $\varsigma$ induced by lensing. The same lensed CMB realization from Figure 5.4 is shown, with the same colormaps and limits for $\lambda$ and $\varsigma$. SCALE quantifies the correlations between the top and bottom panels in its estimates of the underlying lensing statistics. A grid with $20'$ spacing is overlaid, which illustrates patches twice the width of the chosen patches for the real-space proof of concept. The correlations between $\lambda$ and $\varsigma$ exist across many angular scales.

the result. The full details of the derivation are presented in Appendix B.

We begin by constructing *large-scale* temperature gradient fields for two perpendicular directions on the map, $\boldsymbol{\nabla} T_L(\boldsymbol{\ell})$, by applying the top hat filter $W_\lambda(\boldsymbol{\ell})$ combined with a Wiener filter to the original temperature field,

$$W_\lambda(\boldsymbol{\ell}) = \begin{cases} 1, & \ell_{2,\mathrm{min}} \leq |\boldsymbol{\ell}| < \ell_{2,\mathrm{max}} \\ 0, & \mathrm{else}\,, \end{cases} \tag{5.10}$$

$$\boldsymbol{\nabla} T_L(\boldsymbol{\ell}) = \frac{i\,\boldsymbol{\ell}\,W_\lambda(\boldsymbol{\ell})\,C_\ell^{TT}\,T(\boldsymbol{\ell})}{C_\ell^{TT,\mathrm{obs}}}\,. \tag{5.11}$$

Note that a fiducial temperature power spectrum $C_\ell^{TT}$ is required for the Wiener filter in this step. It is not imperative that the assumed model exactly matches the underlying cosmology, as the results are not sensitive to this choice. The observed CMB temperature power spectrum, $C_\ell^{TT,\mathrm{obs}}$, of the map is also required for our filters. We construct a field containing the large scale temperature power fluctuations after returning each gradient component to real space, squaring each component, and then adding them together:

$$\lambda(\hat{\mathbf{n}}) = \big(\nabla_x T_L(\hat{\mathbf{n}})\big)^2 + \big(\nabla_y T_L(\hat{\mathbf{n}})\big)^2\,. \tag{5.12}$$

Similarly, we construct *small-scale* temperature gradient fields in two perpendicular directions on the map, $\boldsymbol{\nabla} T_S(\boldsymbol{\ell})$, by applying a top hat filter $W_\varsigma(\boldsymbol{\ell})$ combined with an inverse-variance filter to the observed temperature field,

$$W_\varsigma(\boldsymbol{\ell}) = \begin{cases} 1, & \ell_{1,\mathrm{min}} \leq |\boldsymbol{\ell}| < \ell_{1,\mathrm{max}} \\ 0, & \mathrm{else}\,, \end{cases} \tag{5.13}$$

$$\boldsymbol{\nabla} T_S(\boldsymbol{\ell}) = \frac{i\,\boldsymbol{\ell}\,W_\varsigma(\boldsymbol{\ell})\,T(\boldsymbol{\ell})}{C_\ell^{TT,\mathrm{obs}}}\,. \tag{5.14}$$

We construct a field containing the small-scale temperature power fluctuations after returning each gradient component to real space, squaring each component, and then adding them together:

$$\varsigma(\hat{\mathbf{n}}) = \big(\nabla_x T_S(\hat{\mathbf{n}})\big)^2 + \big(\nabla_y T_S(\hat{\mathbf{n}})\big)^2\,. \tag{5.15}$$

The field $\varsigma$ strictly contains the small-scale temperature power fluctuations at the scales allowed by the filter $W_\varsigma(\boldsymbol{\ell})$. We expect this field to correlate with the large-scale temperature power fluctuations captured by the field $\lambda$. Such correlations are only expected as a result of lensing on the original CMB temperature field because the filters are chosen to have disjoint support in $\ell$. As a result the cross-spectrum between each field, $C_{\check{L}}^{\varsigma\lambda}$, is a four-point function that estimates the power of the lensing potential $C_L^{\phi\phi}$. Each mode $\check{L}$ of the cross-spectrum represents a particular scale over which the fields $\lambda$ and $\varsigma$ correlate. This is illustrated in Figures 5.4 and 5.5. The $\lambda$ fields look visually similar between unlensed and lensed realizations of the CMB, but in the absence of noise,

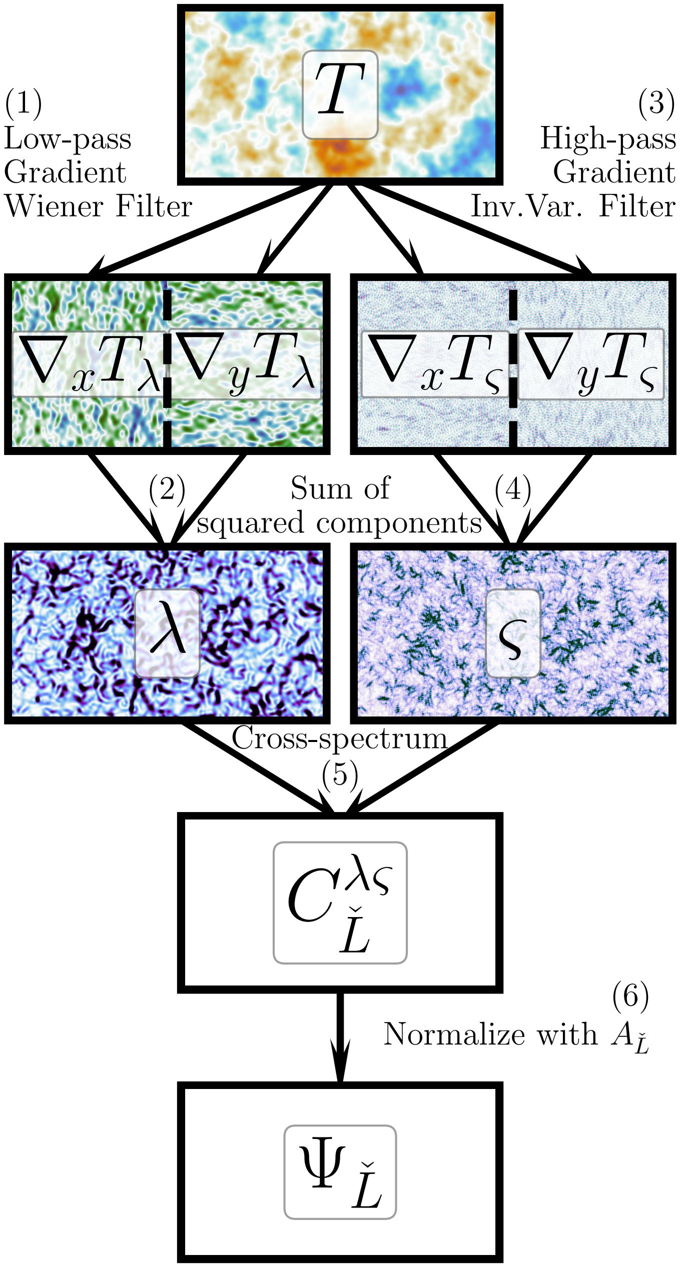

Figure 5.6: Schematic of the steps taken in SCALE pipeline.

the $\varsigma$ fields show much stronger fluctuations in the lensed realization (unlensed small scale power is suppressed by diffusion damping). Fluctuations in the lensed $\varsigma$ field visibly correlate with the $\lambda$ field. Further filtering the $\lambda$ and $\varsigma$ fields illustrates the $\check{L}$ modes probed by the cross-spectrum, and the lensing-induced correlation between $\lambda$ and $\varsigma$ becomes striking.

Figure 5.5 also visualizes how the real-space method from Section 5.3 was combining the information from many modes of $\check{L} \gtrsim 500$ within each patch, and was including information from modes $\check{L} \lesssim 500$ when performing the fit with many patches.

The introduction of noise (and foregrounds) adds power to both $\lambda$ and $\varsigma$ fields, and it can become the dominant source of power in the $\varsigma$ field. Noise contributions to $\lambda$ and $\varsigma$ are not expected to correlate with each other, meaning the cross-spectrum $C_{\check{L}}^{\varsigma\lambda}$ is expected to be largely insensitive to noise (though noise will contribute to its variance).

For the SCALE cross-spectrum to be an unbiased estimate of the lensing power, it needs to be normalized to take into account the filtering that was applied, as well as the expected action of lensing on the fields:

$$\Psi_{\check{L}} = A_{\check{L}} C_{\check{L}}^{\varsigma\lambda} \,. \tag{5.16}$$

The normalization $A_{\check{L}}$ is computed as a double integral of both the observed temperature power spectrum $C_\ell^{TT,\mathrm{obs}}$, and the fiducial temperature power spectrum $C_\ell^{TT}$ used in the Wiener filter above:

$$\begin{aligned} A_{\check{L}} = \Bigg[ 2 \int & \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} W_\varsigma(\boldsymbol{\ell}_1) W_\varsigma(\check{\boldsymbol{L}} - \boldsymbol{\ell}_1) \\ & \times \left(\boldsymbol{\ell}_1 \cdot (\boldsymbol{\ell}_1 - \check{\boldsymbol{L}})\right) \frac{1}{C_{\ell_1}^{TT,\mathrm{obs}}} \frac{1}{C_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_1|}^{TT,\mathrm{obs}}} \\ & \times \int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2} W_\lambda(\boldsymbol{\ell}_2) W_\lambda(\check{\boldsymbol{L}} - \boldsymbol{\ell}_2) \left(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \boldsymbol{\ell}_1)\right) \\ & \quad \times \left((\check{\boldsymbol{L}} - \boldsymbol{\ell}_2) \cdot (\boldsymbol{\ell}_1 - \boldsymbol{\ell}_2)\right) \left(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \check{\boldsymbol{L}})\right) \\ & \quad \times \frac{\left(C_{\ell_2}^{TT}\right)^2}{C_{\ell_2}^{TT,\mathrm{obs}}} \frac{\left(C_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|}^{TT}\right)^2}{C_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|}^{TT,\mathrm{obs}}} \Bigg]^{-1} , \end{aligned} \tag{5.17}$$

The bounds of each integral correspond to the scales allowed by the small-scale window function $W_\varsigma$, and the large-scale window function $W_\lambda$. See Appendix B for the steps leading to the definition of $A_{\check{L}}$ in Equation (B.24). The expected value of $\langle \Psi_{\check{L}} \rangle$ can be similarly computed with the lensing power $C_\ell^{\phi\phi}$:

$$\begin{aligned}\langle \Psi_{\check{L}} \rangle = 2A_{\check{L}} & \int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} W_\varsigma(\boldsymbol{\ell}_1) W_\varsigma(\boldsymbol{\check{L}} - \boldsymbol{\ell}_1) \\ & \times \left(\boldsymbol{\ell}_1 \cdot (\boldsymbol{\ell}_1 - \boldsymbol{\check{L}})\right) \frac{1}{C_{\ell_1}^{TT,\mathrm{obs}}} \frac{1}{C_{|\boldsymbol{\check{L}}-\boldsymbol{\ell}_1|}^{TT,\mathrm{obs}}} \\ & \times \int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2} W_\lambda(\boldsymbol{\ell}_2) W_\lambda(\boldsymbol{\check{L}} - \boldsymbol{\ell}_2) \left(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \boldsymbol{\ell}_1)\right) \\ & \quad \times \left((\boldsymbol{\check{L}} - \boldsymbol{\ell}_2) \cdot (\boldsymbol{\ell}_1 - \boldsymbol{\ell}_2)\right) \left(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \boldsymbol{\check{L}})\right) \\ & \quad \times \frac{\left(C_{\ell_2}^{TT}\right)^2}{C_{\ell_2}^{TT,\mathrm{obs}}} \frac{\left(C_{|\boldsymbol{\check{L}}-\boldsymbol{\ell}_2|}^{TT}\right)^2}{C_{|\boldsymbol{\check{L}}-\boldsymbol{\ell}_2|}^{TT,\mathrm{obs}}} C_{|\boldsymbol{\ell}_1-\boldsymbol{\ell}_2|}^{\phi\phi} \,. \end{aligned} \tag{5.18}$$

Note that unlike for, e.g., the QE estimator, with the SCALE estimator we do not directly recover the signal of immediate interest (namely, $C_\ell^{\phi\phi}$ in this case). This means that non-trivial physical changes to $C_\ell^{\phi\phi}$, such as from extensions to the cosmological model, would appear in the $\Psi_{\check{L}}$ statistic in the SCALE estimator only indirectly, via this integral relation. Nevertheless, we will show below that the expected $\Psi_{\check{L}}$ is readily computed for any cosmological model, and shows excellent agreement with simulated reconstructions.

The expected noise variance of $\Psi_{\check{L}}$, i.e., the variance in the absence of any lensing, is $N_{\check{L}} \approx 4A_{\check{L}}$, and the expected minimum uncertainty on an estimated $\hat{\Psi}_{\check{L}}$ in the limit of low covariance between $\check{L}$ modes is

$$\Delta\hat{\Psi}_{\check{L}} = \sqrt{\frac{\Psi_{\check{L}}^2 + 4A_{\check{L}}}{f_{\mathrm{sky}} \Delta\check{L}(2\check{L}+1)}}. \tag{5.19}$$

The details of $A_{\check{L}}$, $\langle \Psi_{\check{L}} \rangle$ and its expected variances are derived in Appendix B leading up to Equation (B.25). We also demonstrate that the inverse-variance and Wiener filters are the optimal filters to minimize the noise variance $N_{\check{L}}$ of the SCALE estimator. Note that in this procedure we correlate the small-scale *gradient* power with the large-scale gradient power; whereas, in Figure 5.3 we demonstrate that the small-scale temperature power is strongly correlated to the large-scale gradient through lensing. The use of small-scale gradient power is motivated by the discussion at the end of §5.3, where we argue that lensing imparts a directional perturbation to the small-scale temperature field which is correlated with the large-scale temperature gradient in both amplitude and direction. We further find that if we follow the optimization steps with the small-scale and large-scale temperatures (rather than their gradients) in Appendix B, we inevitably conclude that the optimal filters would include factors of $\ell$ for each field which imply the two gradient powers are being correlated.

The general flow of the SCALE pipeline is illustrated in Figure 5.6, and summarized here beginning with a temperature map $T(\hat{\mathbf{n}})$:

1. Transform $T(\hat{\mathbf{n}})$ into harmonic space $T(\boldsymbol{\ell})$, apply the operations described by Equation (5.11), and return $\boldsymbol{\nabla} T_L$ to map space.

2. Compute $\lambda(\hat{\mathbf{n}})$ using Equation (5.12).

3. Apply the operations described by Equation (5.14) to $T(\boldsymbol{\ell})$, and return $\boldsymbol{\nabla} T_S$ to map space.

4. Compute $\varsigma(\hat{\mathbf{n}})$ using Equation (5.15).
5. Compute the cross-spectrum $C_{\check{L}}^{\lambda\varsigma}$.
6. Apply the normalization $\Psi_{\check{L}} = A_{\check{L}} C_{\check{L}}^{\lambda\varsigma}$.

The end result $\Psi_{\check{L}}$ is a set of separate estimates of the lensing power spectrum $C_\ell^{\phi\phi}$ weighted by the normalization $A_{\check{L}}$ along a range of scales set by $W_\varsigma(\boldsymbol{\ell})$ and $W_\lambda(\boldsymbol{\ell})$. We note that the nature of our effectively four-point correlator is reminiscent of the trispectrum calculations for $N_L^{(1)}$ (Kesden et al., 2003) which are typically discarded in CMB lensing power spectrum analyses. This provides a hint that the non-Gaussian signatures of CMB lensing are being considered as part of the SCALE signal. This is further shown by Ref. Zaldarriaga (2000), wherein they derive the lensing signal present in the three- and four-point functions of the small scale CMB temperature field. The cross-spectrum between large- and small-scales is particularly shown to extract information from the connected four-point function corresponding to two large-$\ell$ and two low-$\ell$ modes (see Eq. (37)-(42) and Figure 9a therein). A lensing estimator featuring the cross-spectrum between the large-scale temperature gradient and the small-scale temperature power is also presented in Ref. Zaldarriaga (2000), and some key differences in SCALE are the usage of small-scale *gradient* power, the inclusion of the optimal filters, as well as the normalization $A_{\check{L}}$ to correct for bias. We show below that our method allows us to extract lensing information from the cross-spectrum out to higher $\check{L}$. Finally, it is shown in Ref. Zaldarriaga (2000) that the cross-spectrum between large- and small-scale temperature can be treated as approximately Gaussian, meaning that Eq. (5.19) should accurately predict SCALE uncertainties. We also note that the SCALE procedure draws parallels with the estimator constructed in Smith & Ferraro (2017), wherein the locally measured small-scale ($\ell \gtrsim 3000$) temperature power varies across the sky due to the patchiness of the kinetic Sunyaev-Zel'dovich effect. The main difference is that here we correlate the fluctuations in power between large and small scales, whereas the kSZ estimator of Smith & Ferraro (2017) studies the angular power spectrum of the locally measured small-scale temperature power.

## 5.5 Simulated observables

We test SCALE on simulated CMB maps to determine the robustness of the method. The input power spectra for the simulated maps (shown in Figure 5.1) in all of our analyses were generated with `CAMB`[1] (Lewis et al., 2000; Howlett et al., 2012) and the parameters listed in Table 5.1. We choose parameters to approximately match results from the *Planck* results (Planck Collaboration et al., 2020b), as well as accuracy factors suggested by McCarthy et al. (2022). We generate all simulated raw CMB maps using the `rand_map` method from `pixell`[2] at a resolution of 0.5'. Simulated maps are generally $10^\circ \times 10^\circ$ and centered at the equator. These smaller maps are well within the flat-sky approximation and can be quickly simulated in large quantities. The `rand_map` method imposes repeating boundary conditions in each realization, so the filtering steps do not generate any edge effects. When gathering power spectrum and/or cross-spectrum statistics, we choose bins of $\Delta\ell$, $\Delta L$ and $\Delta\check{L}$ that are integer multiples of the fundamental mode $\ell_{\rm fun} = 36$ for our maps. This is to ensure that we gather values at bin widths which are commensurate with both the grids in which

---

[1] https://camb.info/

[2] https://github.com/simonsobs/pixell

| Parameter | Value |
|---|---|
| `H0` | 67.5 km/s |
| `ombh2` | 0.022 |
| `omch2` | 0.122 |
| `tau` | 0.06 |
| `As` | 2e-9 |
| `ns` | 0.965 |
| `r` | 0 |
| `lmax` | 20000 |
| `lens_potential_accuracy` | 8 |

Table 5.1: The set of non-default arguments given to `CAMB` when simulating power spectra chosen to approximately match results from Planck Collaboration et al. (2020b). Lensing accuracy parameters were chosen as suggested by McCarthy et al. (2022).

| Config. | $w$ [$\mu$K-arcmin] | $b$ [arcmin] | Analogous Exp. |
|---|---|---|---|
| A | 10.5 | 1.3 | ACT (Aiola et al., 2020) |
| B | 6.3 | 1.4 | SO (Ade et al., 2019) |
| C | 1.5 | 1.4 | CMB-S4 (Abazajian et al., 2016) |
| D | 1.0 | 1.0 | Comparison with (Hadzhiyska et al., 2019) |
| E | 0.5 | 0.25 | CMB-HD (Sehgal et al., 2019) |

Table 5.2: The set of simulated noise configurations chosen to be representative of existing or upcoming experiments from ACT (Configuration A, Aiola et al. (2020)) to CMB-HD (Configuration E, Sehgal et al. (2019))

the realizations themselves were generated and in which the correlation statistics are evaluated. We apply lensing to the raw CMB maps using `pixell`'s `lensing` package, and a lens potential field corresponding to the $C_L^{\kappa\kappa}$ spectrum shown in Figure 5.1.

For every CMB map, we generate noise realizations with experiment-relevant values listed in Table 5.2. Configuration A represents a Stage III-like survey like ACT and SPT, while Configuration B is illustrative of the Simons Observatory (SO; Ade et al. 2019). Configuration C gives noise and beam corresponding to CMB Stage IV-like properties (Abazajian et al., 2016), and Configuration D shows a slightly more futuristic experiment corresponding with some tests made for the gradient inversion estimator in Hadzhiyska et al. (2019). Configuration E represents a low-noise, high-resolution experiment like the proposed CMB-HD (Sehgal et al., 2019). We also briefly consider tests in the noise-free limit.

We consider a range of current to future experiments, and we present a particular focus of results for Configuration D. We generally choose a window function for $W_\varsigma(\ell)$ to include modes $\ell_1 \in [6000, 8000]$ (unless otherwise shown) for a balance between being in a regime where the lensing signal is expected to be high, and noise is not too dominant (refer to Figure 5.1). This reasoning is demonstrated against several choices of $\ell_1$ windows in §5.6, but we expect the SCALE methodology to be effective as long as $\ell_1$ satisfies our small-scale approximations (i.e., $\ell_1 \gg 2000$). We also generally choose $W_\lambda(\ell)$ to include modes $\ell_2 \in [0, 3000]$ to ensure that $\lambda$ maps include most of the information about the large-scale temperature gradient power.

We later compare the results of our SCALE lensing method to those of the Hu, DeDeo & Vale

(HDV, Hu et al. (2007)) quadratic estimator with the $TT$ fields. We choose to compare with a quadratic estimator since it provides well-understood and established benchmark. We choose the HDV quadratic estimator in particular due to its behaviour in the small-scale regime. Small angular scale ($\ell \gtrsim 2000$) contributions to the gradient power are removed in the HDV implementation to avoid a bias introduced by higher order cross-terms between the temperature gradient and the lensing convergence in this regime; this is less of a concern with the original Hu & Okamoto estimators applied at larger angular scales (refer to Figure 5.1 and Hu et al. (2007)). The HDV quadratic estimator and its principles have also been applied in studies of cluster lensing including some using *Planck* (Raghunathan et al., 2018), SPT data (Baxter et al., 2018), ACT data (Madhavacheril et al., 2020b), as well as forecasts of lensing results with the proposed CMB-HD experiment (Han & Sehgal, 2022). We perform reconstructions of the lensing convergence field $\hat{\kappa}$ with the HDV method for a subset of simulated lensed CMB maps, along with computations of the optimal noise $N_L^{(0)}$ and realization dependent noise $\hat{N}_L^{(0)}$ using the `symlens`[3] package. In particular, we choose `xmask` with $\ell_{\min} = 2$ and $\ell_{\max} = 3000$, `ymask` with $\ell_{\min} = 2$ and $\ell_{\max} = 10000$, and `kmask` with $L_{\min} = 100$ and $L_{\max} = 10000$. The first two masks are applied to each version of the temperature field of the $TT$ quadratic estimator, and the final mask is applied to the reconstructed convergence field. We apply our method to 100 000 realizations for each suite of tests to obtain stable statistics of the SCALE output. In particular, we found that we need at least 100 000 simulations to reach a converged inverse covariance matrix that we use later to compute signal-to-noise. We apply the HDV quadratic estimator on a subset containing 10 000 of the full set of realizations when making comparisons, choosing a smaller sample size because it is computationally more intensive to run quadratic estimators. We also found that the inverse covariance matrix for the HDV output converges with a sample size of 10 000. Each set of simulations applying the SCALE procedure shares the following general flow:

1. Generate primordial CMB temperature power spectrum $C_\ell^{TT}$, lensing power spectrum $C_L^{\phi\phi}$, and lensed CMB temperature power spectrum $\tilde{C}_\ell^{TT}$ with `CAMB`.
2. Generate $N_\ell^{TT}$ according to one of the experiment configurations in Table 5.2.
3. Compute $A_{\check{L}}$ and $\langle \Psi_{\check{L}} \rangle$ with the above power spectra. $A_{\check{L}}$ are different for lensed/unlensed maps, and $\langle \Psi_{\check{L}} \rangle = 0$ for unlensed maps.
4. For each of 100 000 sims:
   (a) Generate realization of CMB temperature $T$ with $C_\ell^{TT}$, and lensing field $\phi$ with $C_L^{\phi\phi}$.
   (b) Apply lensing to the CMB temperature field to get the lensed temperature field $\tilde{T}$. (Not done in the null test.)
   (c) Generate noise field $N$ with $N_\ell^{TT}$, and add to $\tilde{T}$. (Add to $T$ in the null test with no lensing.)
   (d) Follow the steps in Figure 5.6 to estimate $\hat{\Psi}_{\check{L}}$ for this given realization.

   The next three steps are unnecessary for SCALE, but are performed for 10 000 iterations if we wish to compare SCALE with the HDV quadratic estimator.

---

[3] https://github.com/simonsobs/symlens

| Symbol | Description |
|---|---|
| $T$ | CMB temperature field |
| $\tilde{T}$ | Lensed CMB temperature field |
| $T_L$ | Large-scale temperature field |
| $T_S$ | Small-scale temperature field |
| $\lambda$ | Large-scale temperature gradient power field |
| $\varsigma$ | Small-scale temperature gradient power field |
| $\phi$ | CMB lensing potential field |
| $\kappa$ | CMB lensing convergence field |
| $\ell$ | CMB multipole |
| $L$ | Lensing field multipole |
| $\check{L}$ | SCALE cross spectrum multipole |
| $\ell_1$ | Small-scale filter multipole |
| $\ell_2$ | Large-scale filter multipole |
| $\hat{\mathbf{n}}$ | Line-of-sight direction |
| $C_\ell^{XY}$ | Cross (or auto) spectrum of fields $X$ and $Y$ |
| $N_\ell^{XX}$ | Noise spectrum of $XX$ |
| $\Psi_{\check{L}}$ | Normalized SCALE cross spectrum |
| $\Delta\Psi_{\check{L}}$ | Minimum expected SCALE uncertainty |
| $A_{\check{L}}$ | SCALE normalization |
| $\mathbf{X}$ | Vector quantity $X$ |
| $\boldsymbol{\nabla} X$ | Gradient of field $X$ |
| $\overline{X}$ | Average of quantity $X$ (possibly around $\hat{\mathbf{n}}$) |
| $\Delta X$ | Size/width of bin for quantity $X$ |
| $\hat{X}$ | Estimated quantity or reconstructed field $X$ |
| $\langle X\rangle$ | Expected value of quantity $X$ |

Table 5.3: Summary of notation relevant to the SCALE method for quick reference. Symbols appearing first take precedence in the case of apparent conflict.

(e) Reconstruct the lensing convergence field $\hat{\kappa}$ with the HDV quadratic estimator described above.

(f) Estimate the lensing power spectrum $\hat{C}_L^{\kappa\kappa}$ with the reconstructed $\hat{\kappa}$ field.

(g) Compute the realization-dependent reconstruction noise $\hat{N}_L^{(0)}$.

We provide a summary of all our SCALE-relevant notation in Table 5.3 for quick reference. The summary statistics for SCALE $\hat{\Psi}_{\check{L}}$ and the HDV quadratic estimator $\hat{C}_L^{\kappa\kappa} - \hat{N}_L^{(0)}$ over 100 000 simulations are collected and presented in the following section. We make our code publicly available[4] along with example scripts and an example tutorial notebook.

## 5.6 Results

The SCALE estimator is exceptional at detecting the presence of small-scale lensing in CMB temperature maps. We illustrate this in Figure 5.7, which shows the summary statistics of recovered

---

[4] `https://github.com/victorcchan/cmbpix`

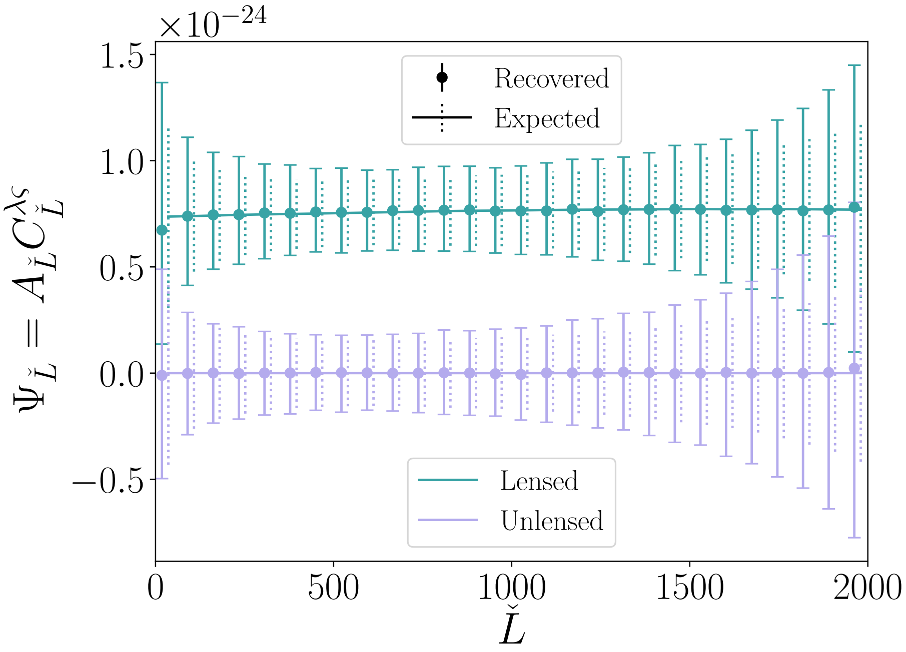

Figure 5.7: Comparison of expected and recovered $\Psi_{\check{L}}$ band-powers with $\Delta\check{L} = 72$ from 100 000 simulations of 100 sq.deg. temperature maps in noise configuration D. Estimates of recovered $\hat{\Psi}_{\check{L}}$ are the median and 68% scatter of the band-power at each bin. The procedure we have outlined in §5.4 accurately recovers the expected lensing signal predicted by Equation (5.18), and the recovered 'signal' from unlensed realizations are consistent with zero.

$\hat{\Psi}_{\check{L}}$ from 100 000 simulations of 10° × 10° lensed and unlensed CMB temperature maps in the experiment D noise configuration. The combination of all bins in Figure 5.7 corresponds to a detection of lensing in 100 sq.deg. maps over the null test with an approximate SNR of 9. We also show the expected theoretical $\langle \Psi_{\check{L}} \rangle$ computed with Equation (5.18). The vertical extent of the error bars represents the 68% width of the distribution (centered at the median within the $\check{L}$ bin) of estimated $\hat{\Psi}_{\check{L}}$ band-powers, and they describe the statistical scatter of estimated band-powers for a given CMB realization of similar total area. The scatter of these band powers is comparable to, but slightly in excess of, the minimal expectation $\Delta\Psi_{\check{L}}$ given by Equation (5.19). Even with relatively small maps of the CMB temperature, the SCALE estimator is able to make a clear distinction of whether or not lensing is present in maps generated with the parameters described in §5.5 and Figure 5.1.

In principle, the SCALE method can be applied to any $\ell_1$ regime so long as the small-scale lensing approximations are appropriate. Figure 5.8 shows comparisons between expected $\langle \Psi_{\check{L}} \rangle$ computed with Equation (5.18) and Equation (5.19) and recovered $\hat{\Psi}_{\check{L}}$ band-power statistics from simulations for different shifts in the small-scale $\ell_1$ window while keeping $\Delta\ell_1 = 2000$. This roughly corresponds to shifting (in the same direction) which $C_\ell^{TT}$, $N_\ell^{TT}$ and $C_L^{\kappa\kappa}$ modes contribute to $\Psi_{\check{L}}$. The expected and recovered band-powers agree to the same extent as the results from Figure 5.7. The recovered band-powers begin to exhibit a positive bias as the $\ell_1$ window is shifted towards lower $\ell$, which can be explained by a departure from the small-scale lensing approximations made in §5.2. In particular, the Taylor series expansion in Equation (5.1) becomes inaccurate at scales $\ell \sim 2000$ because of the similarity of scales with the average deflection angle (Lewis & Challinor, 2006). Higher accuracy

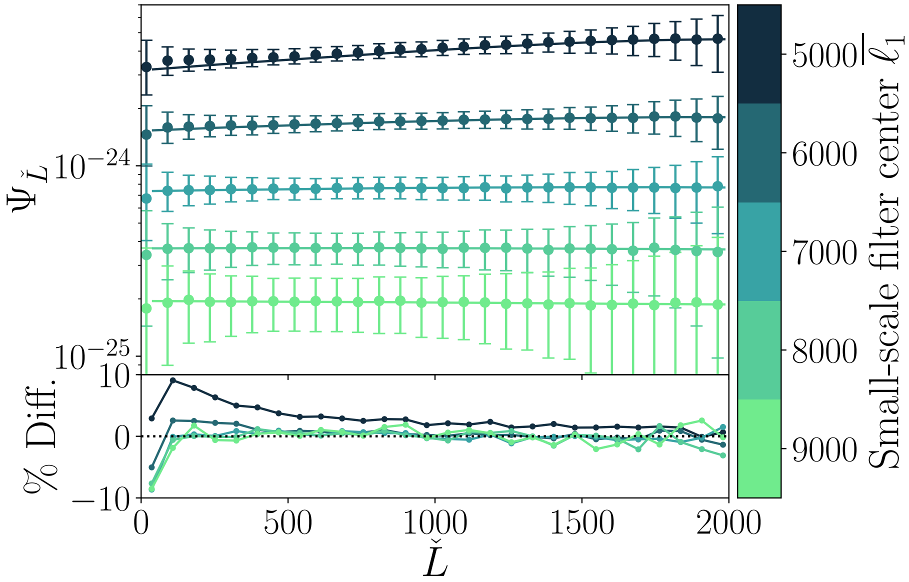

Figure 5.8: *Top*: Comparisons of expected $\Psi_{\check{L}}$ and spread from theory (solid lines with dotted error-bars) and median recovered $\Psi_{\check{L}}$ and 68% spread from simulation (points with capped error-bars) when shifting the centre of the $\ell_1$ window which defines the small-scale filter while keeping the width of the filter constant at $\Delta\ell_1 = 2000$. Error-bars shown are reduced by a factor of 2 for improved visual comparison. *Bottom*: The bias of the recovered $\Psi_{\check{L}}$ when compared to the expected $\Psi_{\check{L}}$ shown as a percentage of the total signal. At fixed window width, increasing the central $\ell_1$ reduces the strength of the recovered signal.

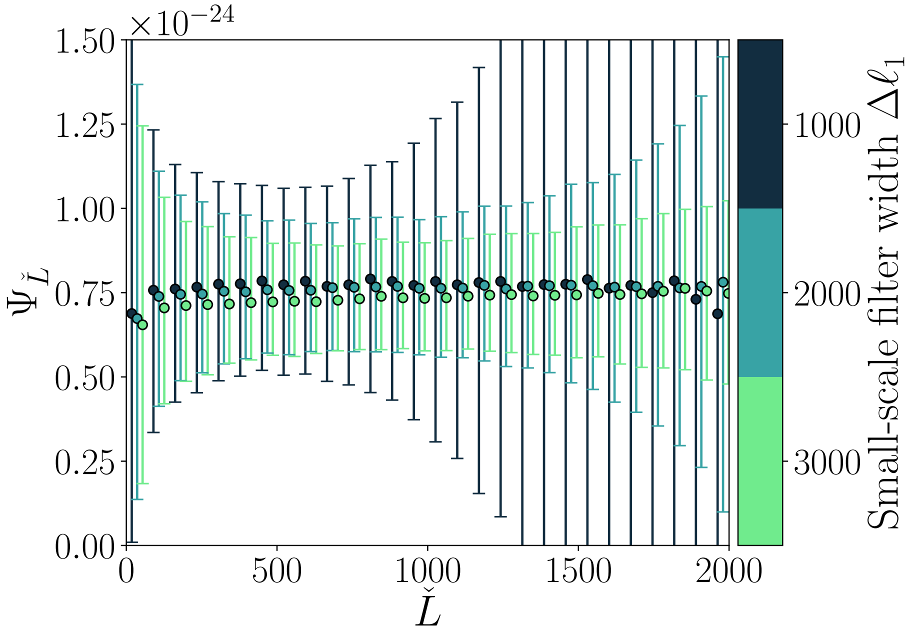

Figure 5.9: As a Complement to Figure 5.8, the comparisons of expected $\Psi_{\check{L}}$ and spread from theory (solid lines with dotted error-bars) and median recovered $\Psi_{\check{L}}$ and 68% spread from simulation (points with capped error-bars) when altering the size of the small-scale filter window $\Delta\ell_1$ used to compute the signal, this time keeping the centre of the $\ell_1$ window at $\overline{\ell_1} = 7000$. For a fixed central $\ell_1$, reducing the width of the $\ell-$window merely increases the error bar on the recovered signal.

in this regime would require consideration of higher-order terms of the expansion. The overall amplitude of each $\Psi_{\check{L}}$ curve decreases as the $\ell_1$ window shifts to higher $\ell$, which reflects the shape of the $C_L^{\kappa\kappa}$ lensing power spectrum in Figure 5.1. The statistical spread of $\Psi_{\check{L}}$ band-powers grows as the $\ell_1$ window is shifted to higher $\ell$ because of increased contributions from experiment noise $N_\ell^{TT}$ at high $\ell$ (see Figure 5.1). The default $6000 < \ell_1 < 8000$ window presented in Figure 5.7 offers a balance between satisfying the small-scale lensing approximations while not appearing to be statistically dominated by experiment noise.

Mathematically, Equation (5.17) and Equation (5.18) are constructed such that the SCALE output is a normalized estimate of the average lensing power within and slightly around the small-scale $\ell_1$ window. The $\ell_1$ window width determines how many $C_L^{\kappa\kappa}$ modes contribute to what we consider signal, but the amplitude of $\Psi_{\check{L}}$ centered at the same $\overline{\ell_1}$ should not change significantly after normalization with $A_{\check{L}}$. Similarly, we add more contributions from $N_\ell^{TT}$ modes, but they are not expected to strongly correlate between large/small-scales. The overall effect of widening the $\ell_1$ window is to use the increased presence of lensing signal to reduce statistical scatter of recovered $\Psi_{\check{L}}$ band-powers. Figure 5.9 illustrates this quite well by comparing $\Psi_{\check{L}}$ after narrowing or widening $\Delta\ell_1$ while keeping the window centered at $\overline{\ell_1} = 7000$. The overall amplitude of the normalized $\Psi_{\check{L}}$ curves appears mostly unchanged, and a wider window does indeed result in tighter distributions of recovered $\Psi_{\check{L}}$ band-powers. The changes we do see in the amplitude of $\Psi_{\check{L}}$ are set by the shape of the underlying lensing power spectrum (i.e., the slope of $C_L^{\kappa\kappa}$ is slightly steeper on one side of the bin center when compared to the other).

While the goal of conventional methods is to reconstruct the underlying lensing field, the SCALE method's output $\Psi_{\check{L}}$ is an indirect estimate of the statistics of the lensing field. Figure 5.5 illustrates the space of $\check{L}$ modes in which the small-scale temperature power fluctuations correlate with those on the large-scales. These $\check{L}$ modes are *not equivalent* to the space of $L$ modes describing the lensing field itself. As shown in Equation (5.18), each band-power of $\Psi_{\check{L}}$ contains information from a wide range of lensing field statistic modes $C_L^{\kappa\kappa}$. Figure 5.10 shows the covariance between estimated $\Delta\check{L} = 72$ band-powers using SCALE on 100 000 simulations with noise configuration D, as well as the covariance between estimated $\Delta L = 100$ band-powers of noise-subtracted lensing power spectra $\hat{C}_L^{\phi\phi} - \hat{N}_L^{(0),\phi\phi}$ from 10 000 reconstructions using the HDV quadratic estimator on a subset of simulations.

Each set of recovered $\hat{\Psi}_{\check{L}}$ band-powers appears to contain correlations induced by lensing that have weak correlations between low $\check{L}$ modes, and those at higher $\check{L}$ modes appear mostly independent from other modes. This is in contrast to the estimated band-powers of the quadratic estimator, which have relatively strong correlations between all bands.

Each covariance matrix $\mathbf{C}$ allows us to compute a signal-to-noise ratio

$$\mathrm{SNR} = \sqrt{a^{\mathrm{T}}\mathbf{C}^{-1}a} \tag{5.20}$$

using either $a = \langle\Psi_{\check{L}}\rangle$ for SCALE or $a = C_L^{\phi\phi}$ for the quadratic estimator. Using covariance matrices for SCALE and HDV QE outputs of 100 000 and 10 000 simulations respectively with the noise configurations in Table 5.2, we compute the expected signal-to-noise ratios and compare them in Figure 5.11 as well as Table 5.4. We scale the covariance matrices used to compute each of these signal-to-noise values by a factor of map-area ratios from 100 sq.deg. to 20 000 sq.deg.

The results from the quadratic estimator directly retrieve information from modes of the lensing

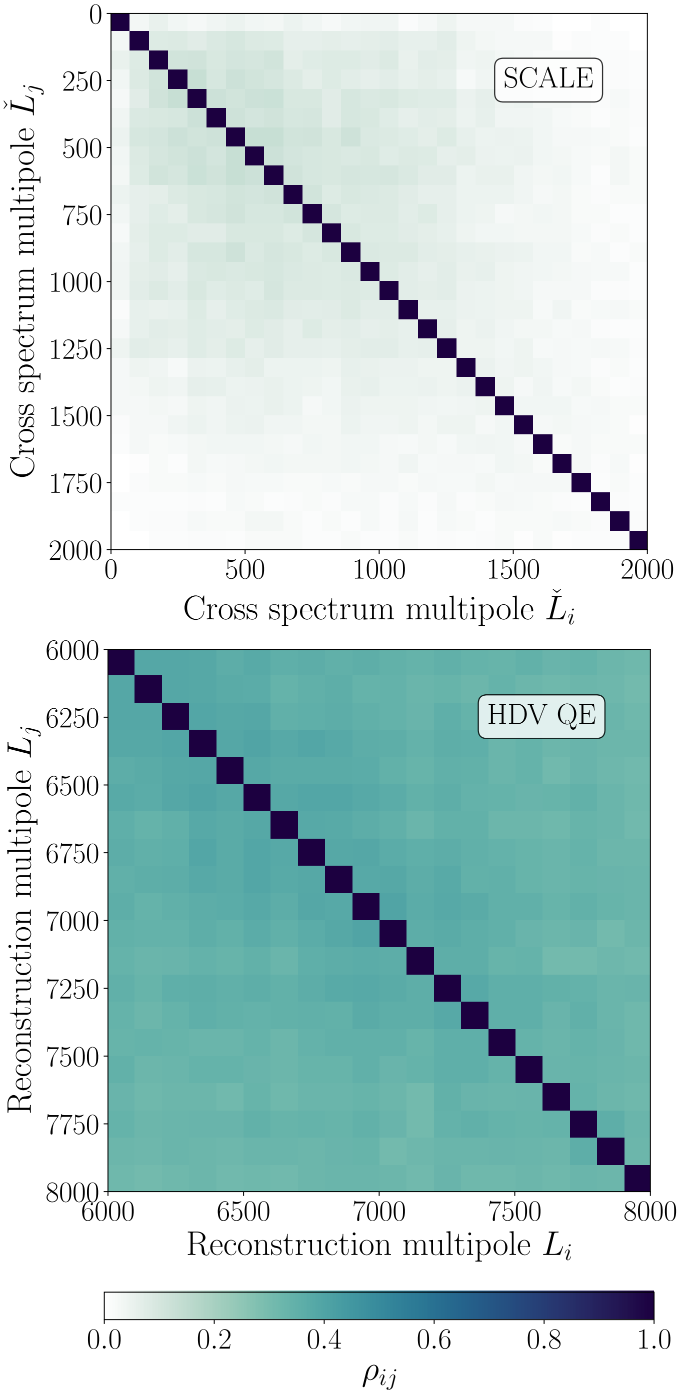

Figure 5.10: *Top:* Correlation matrix of the SCALE estimator output $\Psi_{\check{L}}$ over 100 000 simulations in noise configuration D. The typical variance at a $\check{L}$ bin of width $\Delta\check{L} = 72$ is of order $\sim (10^{-25})^2$. *Bottom:* Correlation matrix of the HDV quadratic estimator reconstructed noise subtracted lensing potential power spectra $\hat{C}_L^{\phi\phi} - \hat{N}_L^{(0),\phi\phi}$ for a subset containing 10 000 of the above simulations. The typical variance at a $L$ bin of width $\Delta L = 100$ is of order $\sim (10^{-25})^2$.

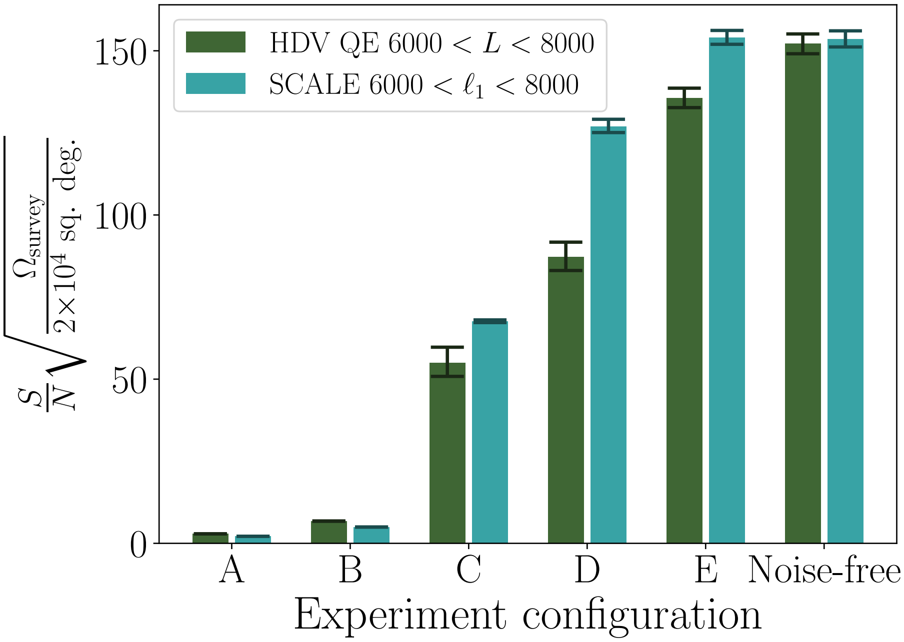

Figure 5.11: Comparison of the values of the signal-to-noise ratio (SNR, Equation (5.20)) across noise configurations between the SCALE estimator for a small-scale window $W_\varsigma(6000 < \ell < 8000)$ and the HDV quadratic estimator applied to $6000 < L < 8000$. SCALE bars indicate the median and 68% range of the bootstrap distribution for SNRs computed 1000 times with a set of 100 000 simulations. HDV bars indicate the median and 68% range of the bootstrap distribution for SNRs computed 1000 times with a set of 10 000 simulations. Each realization has a map area of 100 sq.deg., and SNR values are scaled up to 20 000 sq.deg.

| Config. | HDV QE $S/N$ | SCALE $S/N$ |
|---|---|---|
| A | 2.9 | 2.1 |
| B | 6.8 | 5.0 |
| C | $55.1^{+4.7}_{-4.2}$ | $67.6^{+0.4}_{-0.4}$ |
| D | $87.3^{+4.5}_{-4.2}$ | $127.0^{+2.1}_{-1.9}$ |
| E | $135.7^{+3.0}_{-3.0}$ | $154.0^{+2.2}_{-2.0}$ |
| Noise-free | $152.2^{+2.9}_{-3.1}$ | $153.6^{+2.4}_{-2.5}$ |

Table 5.4: Computed signal-to-noise ratio (SNR, Equation (5.20)) across noise configurations between the SCALE estimator for a small-scale window $W_\varsigma(6000 < \ell < 8000)$ and the HDV quadratic estimator applied to $6000 < L < 8000$. Values are the median and 68% range of the bootstrap distribution for SNRs computed 1000 times with a set of 100 000 simulations for SCALE and 10 000 simulations for the HDV QE. No uncertainty is shown if the 68% range of the bootstrap distribution is smaller than the significant figures provided. Each realization has a map area of 100 sq.deg., and SNR values are scaled up to 20 000 sq.deg.

field between $6000 < L < 8000$ along with some covariance from modes outside of this band; however, we note that we have neglected to consider the effect of subtracting the higher order $N_L^{(1)}$ bias in this simple calculation as is typically done in QE analyses. The same modes of the lensing field are the main contribution to the SCALE results with the $6000 < \ell_1 < 8000$ window due to our choice of $W_\varsigma(\ell)$, but the set of $\Psi_{\check{L}}$s includes contributions from modes of the lensing field outside of this band as shown in Equation (5.18).

The reverse is also true: lensing modes within $6000 < L < 8000$ would also contribute to a lesser extent in other implementations of SCALE with a different choice in $\ell_1$ range. These modes would make up a small part of the SCALE signal if we choose instead to filter for $W_\varsigma(8000 < \ell < 10000)$ rather than the filter choice we make in this paper for $W_\varsigma(6000 < \ell < 8000)$. These contributions to the SCALE estimator from modes outside the filtered band make a direct comparison between both SCALE and other approaches difficult, as it is non-trivial to restrict SCALE to include information only from a certain subset of $L$ modes from the lensing field.

Finally, we consider in a simple example the SCALE method's ability to discriminate between cosmological models that predict changes in the shape/amplitude of the matter power spectrum $P(k)$, and by extension, the lensing power $C_L^{\phi\phi}$ or $C_L^{\kappa\kappa}$ by adjusting the total sum of neutrino masses $\sum m_\nu$. A higher neutrino mass produces effects similar to warm or fuzzy dark matter, suppressing structure formation at small scales, but one key feature is that the lensing power $C_L^{\kappa\kappa}$ is suppressed similarly at high $L$. Figure 5.12 compares a couple of models with neutrino mass heavier than our fiducial model. We see that the fractional changes in $C_L^{\kappa\kappa}$ do not contain much shape information, but in principle, different choices small-scale windows can elucidate potential shape information. We place *approximate* fractional error bands of $\Delta L = 2000$ for SCALE and the HDV QE on the assumption that the SNR values for noise configuration D can be taken at face value. In other words, each fractional error band shown in Figure 5.12 is calculated as 1/SNR, centered on the fiducial curve.

While Figure 5.12 is not meant to be a forecast of either SCALE or QE performance, it provides some insight into the distinguishing power of each method. We note that the SCALE estimator exists in a separate space of $\check{L}$ modes that are each an estimate of a weighted sum of $C_L^{\phi\phi}$ (and by extension $C_L^{\kappa\kappa}$) as prescribed by Equation (5.18). We leave full parameter constraints and a more thorough comparison between methods for future work.

## 5.7 Discussion & Conclusion

In this paper, we:

- showed that fluctuations in the local small-scale ($\ell \gg 3000$) CMB temperature power are intricately tied to the variations in the local large-scale temperature gradient through correlations induced by lensing.

- confirmed that this correlation is readily detectable in the pixel-space statistics of a lensed CMB temperature map, and there is no discernable correlation in a CMB temperature map without lensing (see Figure 5.3).

- visualized correlations between a cross spectrum of large/small-scale CMB temperature power in Figure 5.5.

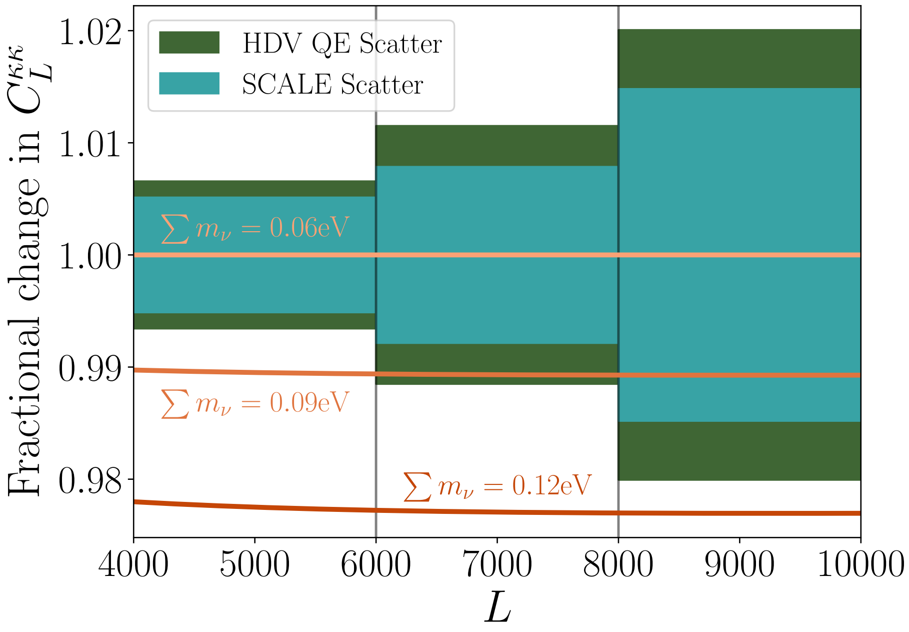

Figure 5.12: Comparison of SCALE vs HDV QE fractional error bands alongside fractional changes in the lensing convergence power $C_L^{\kappa\kappa}$ with neutrino mass. Note that $\sum m_\nu = 0.06$ eV is the fiducial model used in our previous analyses. Different SCALE windows can be used to distinguish between cosmological models, and potentially modifications to the shape of $C_L^{\kappa\kappa}$.

- created the Small Correlated Against Large Estimator (SCALE) which efficiently applies various filters to pick out the relevant small/large-scale in a CMB temperature map, and computes their cross spectrum (Figure 5.6).
- demonstrated that the SCALE method effectively recovers the expected statistics of underlying lensing fields, which matches well with insignificant bias against analytic forms (Figure 5.7).
- tested the properties of the SCALE estimator against different choices of filtering scales (Figures 5.8-5.9).
- determined that the SCALE method can outperform (by a factor of up to 1.5 in signal-to-noise) estimates of the CMB lensing power spectrum $C_L^{\phi\phi}$ through reconstruction with quadratic estimators in noise configurations similar to future experiments.

We find that significant lensing signals at modes $L > 3000$ can be recovered by exploiting the dependence of the CMB temperature power at similar scales $\ell > 3000$ on fluctuations of the CMB temperature gradient (Equation (5.1), illustrated by Figure 5.5). A key advantage of this method is the expectation that any noise, foregrounds, and other CMB secondaries present in observations of the CMB temperature are not expected to correlate with the CMB gradient fluctuations themselves. Foregrounds in particular are known to be a nuisance at all angular scales (e.g., van Engelen et al. (2014); Osborne et al. (2014)), but there have been significant advancements in foreground cleaning techniques in tandem with multi-frequency observations from space experiments such as *Planck*. We expect foregrounds to be well controlled at angular scales relevant to the construction of the large-scale $\lambda$ maps (Madhavacheril & Hill, 2018; Beck et al., 2020; Abylkairov et al., 2021; Darwish et al., 2021a). We expect the effects of foregrounds and foreground cleaning to change the noise structure in the small-scale $\varsigma$ maps, but these contributions should not correlate with $\lambda$ if the foregrounds have been cleaned properly at large angular scales. We do expect the need to account for masking in the normalization $A_{\check{L}}$ and expected value $\langle \Psi_{\check{L}} \rangle$. This would look similar to an additional factor in the filtering, but it would be a convolution because the mask is applied in real space. We may expect a mask to induce spurious correlations between $\lambda$ and $\varsigma$ at $\check{L}$ relevant to the scale of mask apodization ($\check{L} \sim 200$ for an apodization scale of $1^\circ$), but we do not expect this to be the case for *all* $\check{L}$ because such a mask would have support in harmonic space only at the largest scales. At relevant small scales, telescope systematic effects such as the differential and boresight pointing become more important and need to be modelled correctly to remain below a level of $1\sigma$ (Mirmelstein et al., 2021).

We first demonstrated the correlation between large/small-scales in the CMB temperature field with real-space statistics to detect the presence of lensing in a small temperature map covering 100 square degrees. The SCALE procedure builds on the intuitions of the real-space method from §5.3, and it can successfully quantify the correlations induced by lensing in line with expectations.

A simple comparison of signal-to-noise for configurations similar to present and future experiments reveals that the SCALE method demonstrates a marked improvement over the effectiveness of traditional quadratic estimators at low noise levels in the small-scale regime. We do not expect that SCALE will serve as a replacement for existing lensing reconstruction techniques. Map-level reconstruction is useful for delensing and cross-correlation studies, and there are existing reconstruction techniques that are optimal across a wide range of angular scales. SCALE provides the most benefit in the small-scale and low noise lensing regime, so the most precise lensing measurements are likely to come from a combination of different techniques applied to different scales. This could be achieved, for example, by utilizing an estimate of the lensing map from a quadratic estimator, using the estimated lensing map to delens the CMB temperature, and then applying SCALE to estimate the power spectrum of the low noise, small-scale lensing modes that remain in the delensed map. We wish to highlight the simplicity of the SCALE pipeline's steps, which allows it to be quickly applied to any given CMB temperature map. This is in contrast to the maximum likelihood and maximum a posteriori methods which have been shown to be optimal, but they are computationally expensive to perform. These methods, in addition to the Bayesian and Gradient Inversion methods, reconstruct the underlying lensing field $\phi$, which can then be cross-analyzed with other observations such as galaxy clusters. SCALE does not reconstruct a map of the lensing field, but it is a simple and fast method of effectively recovering the statistics of the underlying lensing field to levels of accuracy and precision beyond what is capable with QE techniques.

The SCALE method presents an optimistic outlook for the future of CMB lensing science, providing a fresh opportunity to make high-quality estimates of lensing statistics using a relatively straightforward procedure in a regime that has historically been limited in CMB-only techniques due to limits in techniques and observational noise. The small-scale regime is particularly exciting because the lensing statistics here are sensitive to a wide-range of dark matter and gravitational clustering phenomena.

# Chapter 6

# Applications of SCALE in cosmological parameter estimation

## Abstract

The Small Correlated Against Large Estimator (SCALE) for small-scale lensing of the cosmic microwave background (CMB) is a novel method for measuring the amplitude of CMB lensing power without the need for reconstruction of the lensing field. In our previous study, we showed that the SCALE method can outperform existing reconstruction methods to detect the presence of lensing at small scales ($\ell \gg 3000$). We continue here with a procedure to include information from SCALE in cosmological parameter estimation. We construct a set of emulators to quickly map cosmological parameters to desired CMB observables such as $TT$ power spectra, and SCALE cross spectra. The emulators can provide predictions that are precise to $\sim 0.01\%$ and $\sim$ 2-5% (per mode), respectively. We also outline a method to apply SCALE to full-sky maps of the CMB temperature field, and construct a likelihood for the application of SCALE in parameter estimation. The $TT$ power spectrum as the lone CMB observable, in addition to a *Planck* prior on $\tau$ and BAO information on the total matter density $\Omega_m$, is insufficient for constraining parameters that are sensitive to the small-scale lensing amplitude such as the neutrino mass $m_\nu$. We show that adding SCALE to such an analysis can improve constraints enough to measure the minimum neutrino mass to at least $2.4\sigma$ significance in the scenario of minimal mass, and higher significance for higher mass. A cosmic variance limit measurement of $\tau$ improves the $TT$ + SCALE detection of $m_\nu$ to $4\sigma$.

## 6.1 Introduction

Detections of small-scale lensing in the cosmic microwave background lean in favour towards the Small Correlated Against Large Estimator (SCALE) rather than standard quadratic estimators (QE) in terms of the expected signal-to-noise in upcoming experiments, as we showed in the previous chapter. In this chapter, we build on the SCALE method established in Chapter 5, and study its constraining power when applied to realistic cosmological parameter estimation. Sampling cosmological parameters quickly requires fast predictions of theoretical spectra to compare to within the likelihood, so we construct emulators and present their performance in §6.2. We then present our suite of simulations for the lensed cosmic microwave background in §6.3. We construct a simple likelihood for applying CMB temperature power spectra in parameter estimation, and then extend the model to include SCALE cross spectra in §6.4, and discuss some results in §6.5. Finally, we discuss the results and conclude in §6.6.

## 6.2 Emulation of CMB and SCALE spectra

We begin this chapter with the development of a set of emulators to quickly predict theoretical CMB lensed $TT$ power spectra as well as analytical SCALE products $A_{\check{L}}$ and $\Psi_{\check{L}}$. The posterior sampling process typically requires upwards of $\mathcal{O}(10^3$ to $10^4)$ steps per chain in order to reasonably explore a hyperspace of several cosmological parameters, and that necessitates a quick mapping between the parameters and observables. We find in this section that the CMB angular power spectra can be computed reasonably fast enough with existing software, but SCALE observables require a significant speedup over conventional numerical integration methods. We show that neural network (NN) emulators can predict SCALE observables at the speed required for quick posterior sampling without a significant penalty in terms of accuracy. The emulators can also be trained to predict lensed CMB $TT$ power spectra more quickly than Boltzmann codes, while retaining high accuracy. All computation speeds reported in this chapter are timed with a Ryzen 9 5900X CPU with 12 physical cores and 24 logical cores.

### 6.2.1 Timing of calculations without emulators

Given a set of cosmological parameters, there are now Boltzmann codes that quickly and accurately compute primary CMB power spectra $C_\ell$, lensing power spectra $C_L^{\phi\phi}$, and lensed CMB power spectra $\tilde{C}_\ell$ (`CAMB`[1], Lewis & Challinor 2011; `CLASS`[2], Blas et al. 2011). We opt to use `CAMB` in this work to keep consistency with Chapter 5, but the applications should be comparable to outputs from `CLASS`. Consider computing example power spectra in `CAMB` with `lens_potential_accuracy=8` as the only non-default parameter to ensure lensing accuracy at high-$\ell$ (McCarthy et al., 2022). `CAMB` is able to compute and return the power spectra in $\mathcal{O}(1\,\mathrm{s})$, with some mild dependence on the requested `lmax`. This is fast enough to be used in parameter estimation, but it is possible to speed up the process further by using an emulator. This is especially true if one is interested in running modified versions such as `axionCAMB`[3] which take extra computational steps to include non-standard physics which could affect the small-scale lensing power spectrum (Grin et al., 2022).

---

[1] https://camb.info/

[2] https://github.com/lesgourg/class_public

[3] https://github.com/dgrin1/axionCAMB

A stronger motivation for the construction of emulators for theoretical spectra comes from the application of SCALE in a likelihood. The analytic forms for SCALE products were presented in Chapter 5, and are repeated here:

$$
\begin{aligned}
A_{\check{L}} = \Bigg[ 2 \int & \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} W_\varsigma(\boldsymbol{\ell}_1) W_\varsigma(\check{\boldsymbol{L}} - \boldsymbol{\ell}_1) \left(\boldsymbol{\ell}_1 \cdot (\boldsymbol{\ell}_1 - \check{\boldsymbol{L}})\right) \frac{1}{C^{TT,\mathrm{obs}}_{\ell_1}} \frac{1}{C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}} - \boldsymbol{\ell}_1|}} \\
& \times \int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2} W_\lambda(\boldsymbol{\ell}_2) W_\lambda(\check{\boldsymbol{L}} - \boldsymbol{\ell}_2) \left(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \boldsymbol{\ell}_1)\right) \left((\check{\boldsymbol{L}} - \boldsymbol{\ell}_2) \cdot (\boldsymbol{\ell}_1 - \boldsymbol{\ell}_2)\right) \left(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \check{\boldsymbol{L}})\right) \\
& \times \frac{\left(C^{TT}_{\ell_2}\right)^2}{C^{TT,\mathrm{obs}}_{\ell_2}} \frac{\left(C^{TT}_{|\check{\boldsymbol{L}} - \boldsymbol{\ell}_2|}\right)^2}{C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}} - \boldsymbol{\ell}_2|}} \Bigg]^{-1},
\end{aligned}
\tag{6.1}
$$

$$
\begin{aligned}
\langle \Psi_{\check{L}} \rangle = 2A_{\check{L}} \int & \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} W_\varsigma(\boldsymbol{\ell}_1) W_\varsigma(\check{\boldsymbol{L}} - \boldsymbol{\ell}_1) \left(\boldsymbol{\ell}_1 \cdot (\boldsymbol{\ell}_1 - \check{\boldsymbol{L}})\right) \frac{1}{C^{TT,\mathrm{obs}}_{\ell_1}} \frac{1}{C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}} - \boldsymbol{\ell}_1|}} \\
& \times \int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2} W_\lambda(\boldsymbol{\ell}_2) W_\lambda(\check{\boldsymbol{L}} - \boldsymbol{\ell}_2) \left(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \boldsymbol{\ell}_1)\right) \left((\check{\boldsymbol{L}} - \boldsymbol{\ell}_2) \cdot (\boldsymbol{\ell}_1 - \boldsymbol{\ell}_2)\right) \left(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \check{\boldsymbol{L}})\right) \\
& \times \frac{\left(C^{TT}_{\ell_2}\right)^2}{C^{TT,\mathrm{obs}}_{\ell_2}} \frac{\left(C^{TT}_{|\check{\boldsymbol{L}} - \boldsymbol{\ell}_2|}\right)^2}{C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}} - \boldsymbol{\ell}_2|}} C^{\phi\phi}_{|\boldsymbol{\ell}_1 - \boldsymbol{\ell}_2|},
\end{aligned}
\tag{6.2}
$$

where $A_{\check{L}}$ is the normalization for a cross spectrum $C^{\lambda\varsigma}_{\check{L}}$ between *large-scale* gradient power $\lambda$ and *small-scale* gradient power $\varsigma$ such that $\Psi_{\check{L}} = A_{\check{L}} C^{\lambda\varsigma}_{\check{L}}$. This integral is constructed in a 4-dimensional Fourier space, and an implementation where the integrals are numerically computed with the mid-point rule is provided in our publicly available package `cmbpix`[4]. Consider an example computation of analytic SCALE $A_{\check{L}}$ and $\Psi_{\check{L}}$ for all modes $2 < \check{L} < 2000$ such that the width of the small-scale window is $\ell_{1,\mathrm{max}} - \ell_{1,\mathrm{min}} = 2000$, the large scale window is $0 < \ell_2 < 3000$, and evaluated with a Riemann sum on 2-dimensional grids of $\Delta\ell_1 = 75$ and $\Delta\ell_2 = 100$. This takes $\mathcal{O}(10^4\,\mathrm{s})$ or $\mathcal{O}(10\,\mathrm{min})$ to compute. The computational accuracy and speed of Equation (6.1)-(6.2) is dependent on the resolution of the grid(s) on which it is evaluated. Regardless, numerically integrating the analytic SCALE products with the mid-point rule is slow enough that one would desire considerable speedups. A first approach to speeding up the evaluation of Equation (6.1)-(6.2) is to consider Monte Carlo (MC) integration. We find that a Monte Carlo integration offers a good balance of speed and accuracy, and is implemented in `cmbpix`. One drawback with Monte Carlo integration is its non-deterministic nature, and there is some level of inherent inaccuracy dependent on the number of samples with which the integral is evaluated. Figure 6.1 shows that the accuracy of the MC integration is approximately $\sim 1\%$ scatter when evaluated with $N_{\mathrm{samples}} \sim \mathcal{O}(10^5)$ samples. This is the accuracy for the evaluation at one $\check{L}$ mode, and when we bin SCALE data into bands with width $\Delta\check{L} = 71$, we expect the scatter to be reduced by a factor of $\sqrt{\Delta\check{L}} \approx 8.4$. Evaluating $A_{\check{L}}$ and $\Psi_{\check{L}}$ in the same configuration as above with $N_{\mathrm{samples}} = 2 \times 10^5$ in the MC integration takes $\mathcal{O}(10\,\mathrm{s})$. This is a significant speed-up, but it is still too slow to be used in parameter estimation.

---

[4] https://github.com/victorcchan/cmbpix

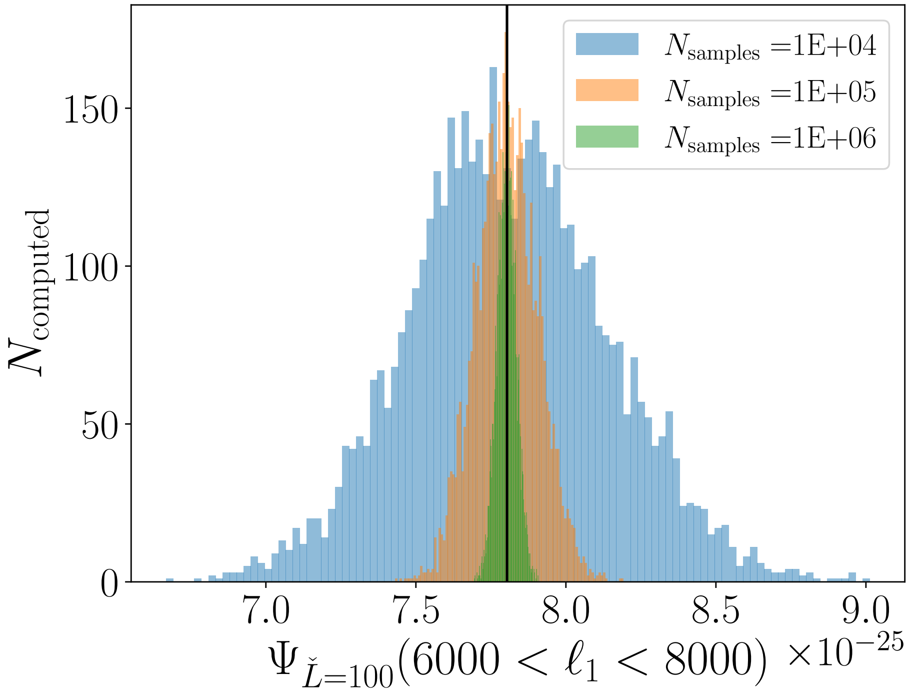

Figure 6.1: A comparison of the accuracy of the Monte Carlo integration of Equation (6.2) at $\check{L} = 100$, $6000 < \ell_1 < 8000$, and $0 < \ell_2 < 3000$ for different choices in sample size. Evaluating with $N_{\text{samples}} \sim \mathcal{O}(10^5)$ yields approximately $\sim 1\%$ scatter, which is the sample size we use to compute the training set for our emulators. A vertical line indicates the expected value computed with the mid-point rule. We choose to use $N_{\text{samples}} = 20000$ to compute the SCALE cross spectra in the training set for our emulator.

Table 6.1: The set of cosmological parameters and their prior ranges used to generate training data for the NN emulators. The ΛCDM parameters are chosen to be centered on the Planck 2018 best-fit cosmology (Planck Collaboration et al., 2020a) with $\pm 4.5\sigma$ on either side. Note that the emulators are trained on $\Omega_\mathrm{b}$, $\Omega_\mathrm{c}$, and $\ln(10^{10} A_s)$.

| Parameter | Prior range |
|---|---|
| $h$ | $[0.6793, 0.6979]$ |
| $\Omega_\mathrm{b} h^2$ | $[0.021695, 0.023045]$ |
| $\Omega_\mathrm{c} h^2$ | $[0.1146, 0.1254]$ |
| $\tau$ | $[0.02155, 0.08725]$ |
| $A_\mathrm{s}$ | $[1.965, 2.235] \times 10^{-9}$ |
| $n_\mathrm{s}$ | $[0.946, 0.9838]$ |
| $m_\nu$ | $[0, 0.18]\,\mathrm{eV}$ |
| $\Delta\ell_1$ | $^a[800, 3200]$ |
| $\overline{\ell_1}$ | $^a[6200, 9800]$ |

$^a$Used only for SCALE emulators.

### 6.2.2 Construction of emulators

We construct a set of emulators for the theoretical CMB lensed $TT$ power spectra $\tilde{C}_\ell^{TT}$, and SCALE spectra $\Psi_{\check{L}}$. The emulators are created in the `COMSOPOWER`[5] framework (Spurio Mancini et al., 2022), which is a Python package for the construction of emulators for cosmological observables. It is based on `TensorFlow`[6] (Abadi et al., 2015), and it provides a structure for training and using neural network (NN) emulators.

#### Training data

The training data span a set of $N_\mathrm{train} = 8192$ cosmological parameter samples from a Latin hypercube for a uniform prior in the ranges set by Table 6.1.

We compute the lensed CMB $TT$ power spectrum out to `lmax=20000` with `CAMB` for all $N_\mathrm{train} = 8192$ sets of parameters in the training range, remembering to set `lens_potential_accuracy=8` for high-$\ell$ accuracy (McCarthy et al., 2022). We also compute the SCALE observable for each set of parameters with the Monte Carlo integration in `cmbpix`. In practice, the only SCALE observable that needs to be predicted as a function of cosmological parameters is the un-normalized cross spectrum $C_{\check{L}}^{\lambda\varsigma}$ which can be computed with Equation (6.2) excluding the $A_{\check{L}}$ factor. This is because the construction of the SCALE data vector $\Psi_{\check{L}}$ depends on a set of optimal filters that are applied directly to a CMB temperature field (further discussed in §6.3). The expected value expressed in Equation (6.2) and the normalization $A_{\check{L}}$ of Equation (6.1) assume a choice of fiducial cosmological parameters in the filtering scheme. We assume a set of fiducial parameters listed in Table 6.2 in the filtering when computing the set of training cross spectra $C_{\check{L}}^{\lambda\varsigma}$. In summary, the training set of SCALE observables $C_{\check{L}}^{\lambda\varsigma}$ is computed with Equation (6.2) without the $A_{\check{L}}$ factor, where the unlensed $TT$ power $C_\ell^{TT}$ and total $TT$ power $C_\ell^{TT,\mathrm{obs}} = \tilde{C}_\ell^{TT} + N_\ell^{TT}$ are dependent on only the fiducial model described in Table 6.2. The only dependence of the SCALE observables on the cosmological parameters is through the lensing power spectrum $C_L^{\phi\phi}$. We additionally add a dependence of

---

[5] https://github.com/alessiospuriomancini/cosmopower

[6] https://www.tensorflow.org/

Table 6.2: The set of cosmological parameters chosen to be the fiducial model.

| Parameter | Value |
|---|---|
| $h$ | 0.675 |
| $\Omega_\mathrm{b}h^2$ | 0.022 |
| $\Omega_\mathrm{c}h^2$ | 0.122 |
| $\tau$ | 0.06 |
| $A_\mathrm{s}$ | $2.1 \times 10^{-9}$ |
| $n_\mathrm{s}$ | 0.965 |
| $m_\nu$ | 0.06 eV |
| $w$ | [a]1 $\mu$K-arcmin |
| $\sigma_b$ | [a]1 arcmin |

[a] Experiment noise chosen to match the analysis of Configuration D in Chapter 5.

Table 6.3: The set of neural network specifications for the emulators.

| Parameter | Argument |
|---|---|
| `hidden_layers` | $[512, 512, 512, 512]$ |
| `validation_split` | 0.125 |
| `learning_rates` | $[10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}]$ |
| `batch_sizes` | [1024, 1024, 1024, 1024, 1024] |
| `gradient_accumulation_steps` | [1, 1, 1, 1, 1] |
| `patience_values` | [100, 100, 100, 100, 100] |
| `max_epochs` | [1000, 1000, 1000, 1000, 1000] |

SCALE observables on the small-scale filter width $\Delta\ell_1$ and center $\overline{\ell_1}$, which is equivalent to altering the limits of the $\ell_1$ integral in Equation (6.2). This allows for some flexibility when applying the likelihood for SCALE data vectors constructed from the same map(s), but with different ranges of small-scale filtering.

### Emulator structure

Each of the emulators we construct contain 4 hidden layers, each with 512 nodes. The input layer contains the 7 cosmological parameters in Table 6.1 for the $\tilde{C}_\ell^{TT}$ emulator, and a set of 3000 nodes for $\tilde{C}_\ell^{TT}$ between $2 < \ell < 3002$ as outputs. Similarly, the SCALE emulator takes all 9 parameters of Table 6.1, and outputs to 2000 nodes for $C_{\check{L}}^{\lambda\varsigma}$ between $2 < \check{L} < 2002$. We note here that the input parameters for the emulators are converted such that they are trained on $\Omega_\mathrm{b}$, $\Omega_\mathrm{c}$, and $\ln(10^{10}A_s)$. Both emulators are trained with a cooling schedule such that the learning rate steps through $[10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}]$. The emulators are trained with a batch size of 1024, and the training skipped ahead to the next learning rate (or ended at the final stage) when the validation loss does not improve for 100 epochs. The emulators are trained with 87.5% of the full training set $N_\mathrm{train} = 8192$, with the remaining 1024 reserved for validation. The training is performed on a single NVIDIA GeForce RTX 3080 GPU with 10 GB of memory. The $TT$ emulator takes approximately 4 min to train, and the SCALE emulator takes approximately 1 min 30 s to train. Each emulator only needs to be trained once, and the trained model can be saved and loaded for future use.

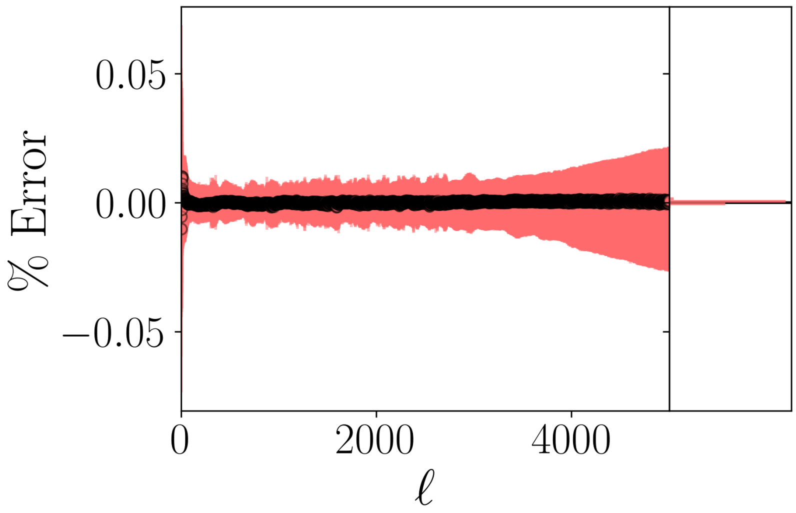

Figure 6.2: The validation of the $\tilde{C}_\ell^{TT}$ emulator. The % error represented here is computed as the difference between the emulator prediction and the `CAMB` output divided by the `CAMB` output. The black markers indicate the median error across realizations of the predicted spectra for the validation set at each $\ell$ mode, and the extent of the error bars indicates the 68-percentile scatter centered on the median. There is a scatter of $\sim 0.1\%$ per mode for the $\tilde{C}_\ell^{TT}$ emulator, which goes down with binning. The right panel depicts the distribution of median errors (black markers) for each $\ell$ mode. It is centered at $2 \times 10^{-4}\%$ with a width of $6 \times 10^{-4}\%$, indicating a lack of significant bias from the emulator.

### Emulator performance

We find that both emulators perform much quicker than their original counterparts, providing predictions of their respective observables in $\mathcal{O}(10^{-3}\,\mathrm{s})$. Figure 6.2 shows that the lensed CMB $TT$ emulator is extremely accurate, with a prediction error of $\sim 0.1\%$ per $\ell$ mode. Similarly, Figure 6.3 shows that the SCALE emulator predicts the cross spectra $C_{\check{L}}^{\lambda\varsigma}$ with an error of $\sim$ 2-5% per $\check{L}$ mode. We attribute the lower precision of the SCALE emulator to the inherent scatter of the input cross spectra computed with Monte Carlo integration. The emulators are expected to perform well when predicting functions that are smooth, as is the case for the $TT$ power spectra (Spurio Mancini et al., 2022). The SCALE cross spectra are expected to be smooth, but the natural scatter associated with the MC integration adds some level of fuzziness relative to the number of samples used in the computation (refer to Figure 6.1). The scatter from the SCALE emulator is at approximately the same level as the precision of the MC integration itself ($\sim 1\%$), and we show later that binning the cross spectrum into band-powers reduces the scatter. Finally, both emulators exhibit an insignificant bias in their predictions, as the prediction errors effectively scatter around zero. A summary of the speedup provided by the emulators is provided in Table 6.4.

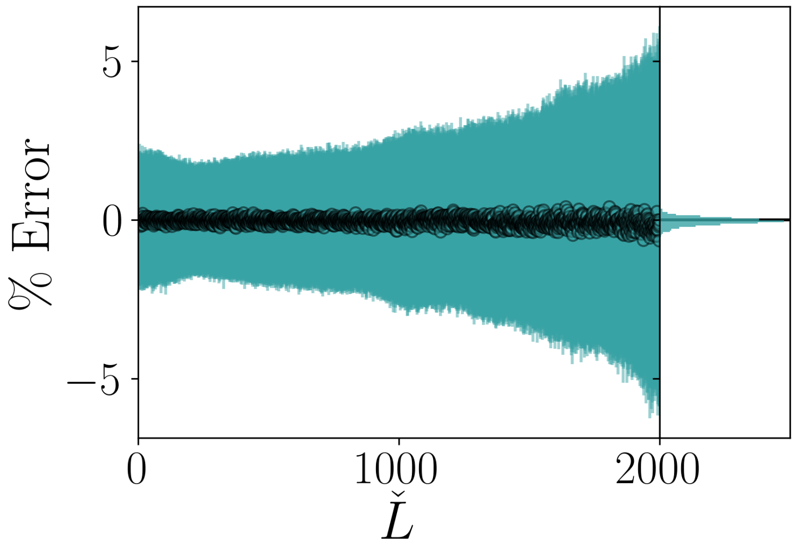

Figure 6.3: The validation of the SCALE $C_{\check{L}}^{\lambda\varsigma}$ emulator. The error represented here is computed as the difference between the emulator prediction and the `cmbpix` output divided by the `cmbpix` output. The black markers indicate the median error across realizations of the predicted spectra for the validation set at each $\check{L}$ mode, and the extent of the error bars indicates the 68-percentile scatter centered on the median. There is an expected scatter of $\sim$ 2-5% per mode for the $C_{\check{L}}^{\lambda\varsigma}$ emulator, which goes down with binning. The right panel depicts the distribution of median errors (black markers) for each $\check{L}$ mode. It is centered at $-0.05\%$ with a width of 0.13%, indicating a lack of significant bias from the emulator.

Table 6.4: A comparison of computation speed for theoretical CMB lensed power spectra and SCALE cross spectra. The emulators provide significant speedup for both observables.

| Computation | Time to evaluate $[\mathcal{O}(\mathrm{s})]$ |
|---|---|
| $\tilde{C}^{TT}_{2<\ell<5\mathrm{k}}$ (`CAMB`) | 1 |
| $\tilde{C}^{TT}_{2<\ell<5\mathrm{k}}$ (Emulator) | $10^{-3}$ |
| $\Psi_{2<\check{L}<2\mathrm{k}}$ (`cmbpix` mid-point) | $10^{4}$ |
| $\Psi_{2<\check{L}<2\mathrm{k}}$ (`cmbpix` Monte Carlo) | 10 |
| $C^{\lambda\varsigma}_{2<\check{L}<2\mathrm{k}}$ (Emulator) | $10^{-3}$ |

## 6.3 Simulations

We compute a suite of full-sky simulations of the lensed CMB temperature field with the `lenspyx`[7] package, which wraps around methods from `DUCC`[8] (Distinctly Useful Code Collection) which allow for efficient and accurate lensing and de-lensing operations with spherical harmonics transforms (Reinecke et al., 2023). A notebook with instructions to simulate the lensed CMB temperature and polarization is provided in the `lenspyx` repository. We choose to simulate maps with `HEALPix` resolution `NSIDE=8192`, as the `lenspyx` accuracy is good out to $\ell \approx 2 \times$ `NSIDE`. The simulations are constructed with the same fiducial cosmological parameters as shown in Table 6.2, and the general procedure (all methods are called from `lenspyx` unless otherwise stated) is as follows:

1. Compute the unlensed CMB $TT$ power spectrum $C_\ell^{TT}$ and lensing potential power spectrum $C_L^{\phi\phi}$ with `CAMB` out to `lmax=20000` with `lens_potential_accuracy=8`.

2. Generate spherical harmonic coefficients $a_{\ell m}$ for both an unlensed temperature $T$ using $C_\ell^{TT}$ and lensing potential $\phi$ field using $C_L^{\phi\phi}$ with `synalm`.

3. Transform the lensing potential field into a spin-1 deflection field $\mathbf{d}$ with `lenspyx`'s `almxfl` method.

4. Compute the lensed temperature field $\tilde{T}$ using the unlensed temperature $T$ and deflection $\mathbf{d}$ coefficients with `alm2lenmap`. This returns a lensed temperature field $\tilde{T}$ in map-space at `NSIDE=8192`.

5. Generate spherical harmonic coefficients $a_{\ell m}$ for a noise temperature field $N$ with `synalm` (noise parameters also in Table 6.2), convert to map-space with `alm2map`, and add to the lensed temperature field $T_{\text{obs}} = \tilde{T} + N$.

6. Convert the observed temperature field $T_{\text{obs}}$ to spherical harmonic coefficients $a_{\ell m}$ using `map2alm`, and compute the total observed $TT$ power spectrum $\tilde{C}_\ell^{TT,\text{obs}}$ with `alm2cl`, which is saved.

7. Apply a low-pass filter such that $0 < \ell_2 < 3000$ along with a Wiener filter to the observed temperature field $T_{\text{obs}}$ (shown below as Equation (6.3)) with `almxfl`. The product is a set of spherical harmonic coefficients, that when converted to map space with a spin-1 inverse transform `alm2map_spin`, produces the two large-scale gradient components $[\nabla_\theta T_L, \nabla_\phi T_L / \sin\theta]$ that make up the $\lambda = (\nabla_\theta T_L)^2 + (\nabla_\phi T_L/\sin\theta)^2$ map of large-scale temperature gradient power required for one half of SCALE. Note that the filter is constructed with the theoretical spectra from Step 1 using `CAMB`, and follow the fiducial cosmology in Table 6.2.

$$W_\lambda(\ell) = \begin{cases} \sqrt{\ell(\ell+1)}\frac{C_\ell^{TT}}{\tilde{C}_\ell^{TT}+N_\ell^{TT}} & , \ \ell < 3000 \\ 0 & , \ \ell \geq 3000 \end{cases} . \tag{6.3}$$

8. Convert the $\lambda$ map into spherical harmonic space with `map2alm`.

---

[7] https://github.com/carronj/lenspyx

[8] https://gitlab.mpcdf.mpg.de/mtr/ducc

Table 6.5: Binning scheme for the CMB $TT$ power spectra $\tilde{C}_\ell^{TT,\mathrm{obs}}$.

| Multipole range | Bin width $\Delta\ell$ | Number of band-powers |
|---|---|---|
| $2 \leq \ell \leq 31$ | 1 | 30 |
| $32 \leq \ell \leq 361$ | 11 | 30 |
| $362 \leq \ell \leq 1471$ | 37 | 30 |
| $1472 \leq \ell \leq 2986$ | 101 | 15 |

9. Apply a high-pass filter such that $\ell_{1,\mathrm{min}} < \ell_1 < \ell_{1,\mathrm{max}}$ along with an inverse variance filter to the observed temperature field $T_{\mathrm{obs}}$ (shown below as Equation (6.4)) with `almxfl`. The product is a set of spherical harmonic coefficients, that when converted to map space with a spin-1 inverse transform `alm2map_spin`, produces the two small-scale gradient components $[\nabla_\theta T_S, \nabla_\phi T_S/\sin\theta]$ that make up the $\varsigma = (\nabla_\theta T_S)^2 + (\nabla_\phi T_S/\sin\theta)^2$ map of small-scale temperature gradient power required for the other half of SCALE. Note that the filter is constructed with the theoretical spectra from Step 1 using `CAMB`, and follow the fiducial cosmology in Table 6.2.

$$W_\varsigma(\ell) = \begin{cases} \sqrt{\ell(\ell+1)}\frac{1}{\tilde{C}_\ell^{TT}+N_\ell^{TT}} & , \ \ell_{1,\mathrm{min}} < \ell_1 < \ell_{1,\mathrm{max}} \\ 0 & , \ \text{else} \end{cases} . \tag{6.4}$$

10. Convert the $\varsigma$ map into spherical harmonic space with `map2alm`.

11. Compute the cross spectrum $C_{\check{L}}^{\lambda\varsigma}$ between $\lambda$ and $\varsigma$ with `alm2cl`, which is saved.

We repeat the above procedure for 400 simulations, and produce a set of 400 observed $TT$ power spectra $\tilde{C}_\ell^{TT,\mathrm{obs}}$ and SCALE cross spectra $C_{\check{L}}^{\lambda\varsigma}$. We show in §6.4 that, in principle, Steps 9–11 may be repeated for the same realization with various choices in the small-scale $\ell_{1,\mathrm{min}} < \ell_1 < \ell_{1,\mathrm{max}}$ filter limits, and included in the same likelihood. For the purposes of this study, we limit ourselves to just one version of SCALE such that $8\,000 < \ell_1 < 10\,000$. We choose this range to ensure that we are comfortably within the small-scale approximations with which SCALE is constructed while also including a healthy amount of small-scale modes. While we save the full spectra for both observables at every mode from $\ell \geq 2$ up to `lmax=20000`, we compute band-powers of each spectrum to reduce the noise of each individual entry in the data vector, as well as to reduce the dimensionality of the data vector. We choose to bin the $TT$ power spectra into 105 band-powers following the prescription outlined in Table 6.5, and the SCALE cross spectra are binned into band-powers of equal width $\Delta\check{L} = 71$ out to $\check{L}_{\mathrm{max}} \approx 2\,000$. The SCALE cross spectra were normalized by $A_{\check{L}}$ computed with Equation (6.1) before binning using the fiducial cosmology, as its purpose is to debias the cross spectrum with respect to the filters which assumed the fiducial model.

The covariance matrices for the 400 binned $TT$ power spectra and SCALE cross spectra are shown in Figure 6.4 and Figure 6.5 respectively. A remarkable feature of Figure 6.4 is an apparent lack of contrast at the high-$\ell$ limit of the covariance. This can likely be explained by the change in binning structure at $\ell \approx 1500$, where the noise is artificially reduced from the bin being approximately three times wider. We find that the SCALE observables, in particular, exhibit very low levels of covariance between band-powers, as we found with the flat-sky simulations in Chapter 5. The highest levels of SCALE variance are towards the edges of the $\check{L}$ range, which we attribute to cosmic variance at low $\check{L}$, and noise at higher $\check{L}$. The covariances between band-powers will be necessary in our

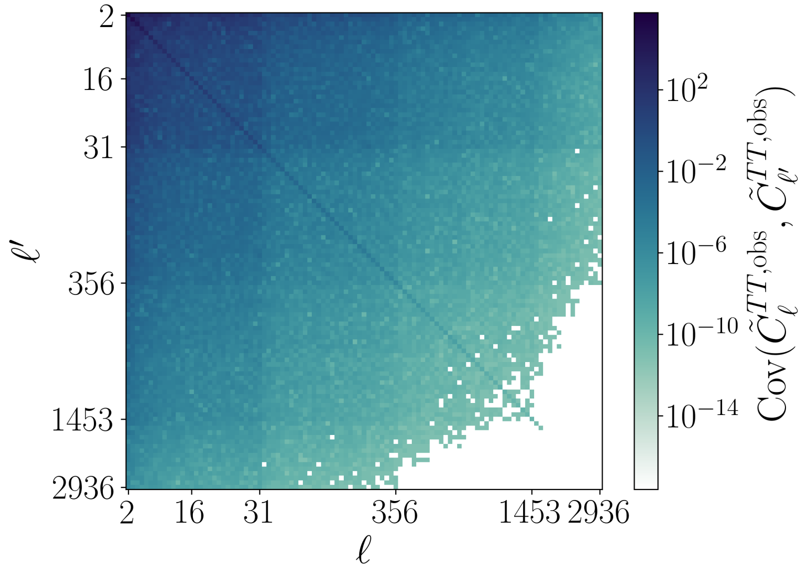

Figure 6.4: Covariance matrix between band-powers of simulated CMB temperature power spectra $\tilde{C}_\ell^{TT,\mathrm{obs}}$. The covariance is computed from 400 simulations, and the band-powers are binned according to the prescription in Table 6.5. Contrast becomes low at high-$\ell$ due to the large bin width. This is the covariance used in our likelihood with $TT$ band-powers.

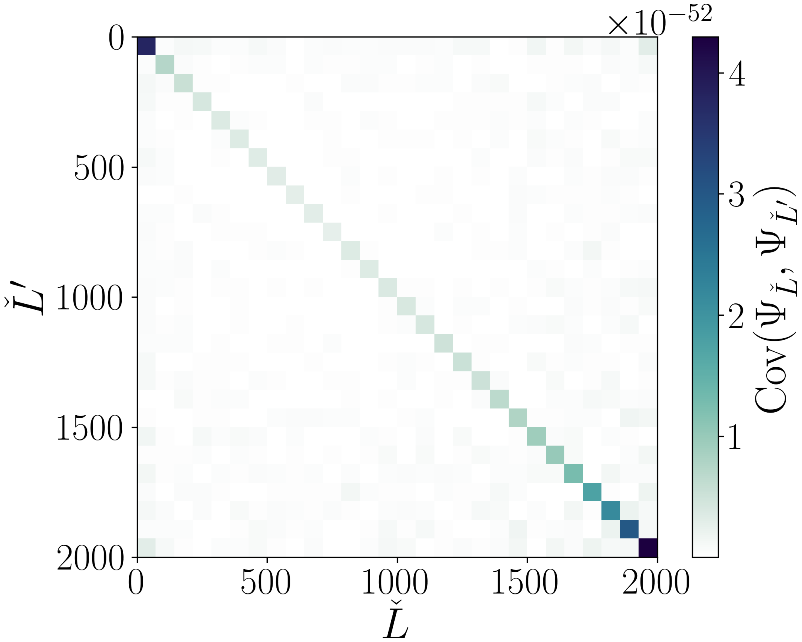

Figure 6.5: Covariance between band-powers of simulated SCALE cross spectra $C_{\check{L}}^{\lambda\varsigma}$. The covariance is computed from 400 simulations, and the band-powers are binned with width $\Delta\check{L} = 71$. cross spectra were normalized by with $A_{\check{L}}$ computed with Equation (6.1) with the fiducial cosmology in Table 6.2.

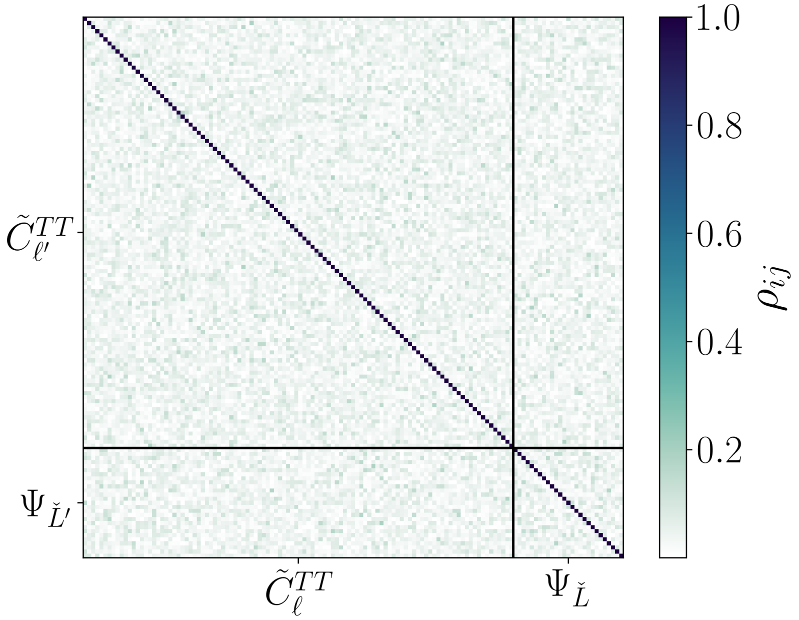

Figure 6.6: Correlation matrix between CMB temperature and SCALE band-powers. The correlation is computed from 400 simulations, and the band-powers are binned according to the prescription in Table 6.5 for $\tilde{C}_{\ell}^{TT,\mathrm{obs}}$ and with width $\Delta\check{L} = 71$ for $\Psi_{\check{L}}$. The covariance matrix is used in our likelihood combining $TT$ band-powers with SCALE band-powers.

likelihood described in §6.4. The correlation matrix (covariance divided by the outer product of the square-root of its diagonal) between the band-powers of all observables is shown in Figure 6.6. We find low levels of correlation generally across all observed band-powers.

Finally, we present a visualization of one realization of the simulated observables in Figure 6.7 in comparison with the theoretical values at the fiducial cosmology. Our choice in binning allows for accurate characterization of approximately the first 5 peaks, as well as the first 4 troughs of the $TT$ power spectra. The normalized SCALE spectrum is relatively featureless, so a simple binning prescription allows for accurate characterization of the overall lensing amplitude within by the range of the small-scale $\ell_1$ filter.

## 6.4 Constructing a Likelihood with SCALE

All the components required for probabilistic sampling of parameters are now in place. We consider a model fit for a ΛCDM cosmology with the addition of one neutrino mass eigenstate parameterized with $m_\nu$. The general suppression of small-scale lensing due to the massive neutrino offers a simple test of SCALE's constraining power. Our vector of parameters is $\theta = \{m_\nu, \Omega_c, \Omega_b, \ln(10^{10}A_s), n_s, h, \tau\}$. There is an implicitly constrained parameter $\Omega_\Lambda$ such that $\Omega_c + \Omega_b + \Omega_\nu + \Omega_\Lambda = 1$. Our data vector $\vec{d} = \{\tilde{C}_{\ell}^{TT,\mathrm{obs}}, \Psi_{\check{L}}\}$ consists of a set of $TT$ and SCALE band-powers, and we have constructed emulators in §6.2 to predict the theoretical expected values of the observables given a set of cosmological parameters: $\vec{t}(\theta)$. We would like to remind the reader that the simulations and emulators return the un-normalized SCALE cross spectrum $C_{\check{L}}^{\lambda\varsigma}$, and that we must first normalize them $\Psi_{\check{L}} = A_{\check{L}} C_{\check{L}}^{\lambda\varsigma}$ using Equation (6.1) with the fiducial cosmology in the filtering **before** binning.

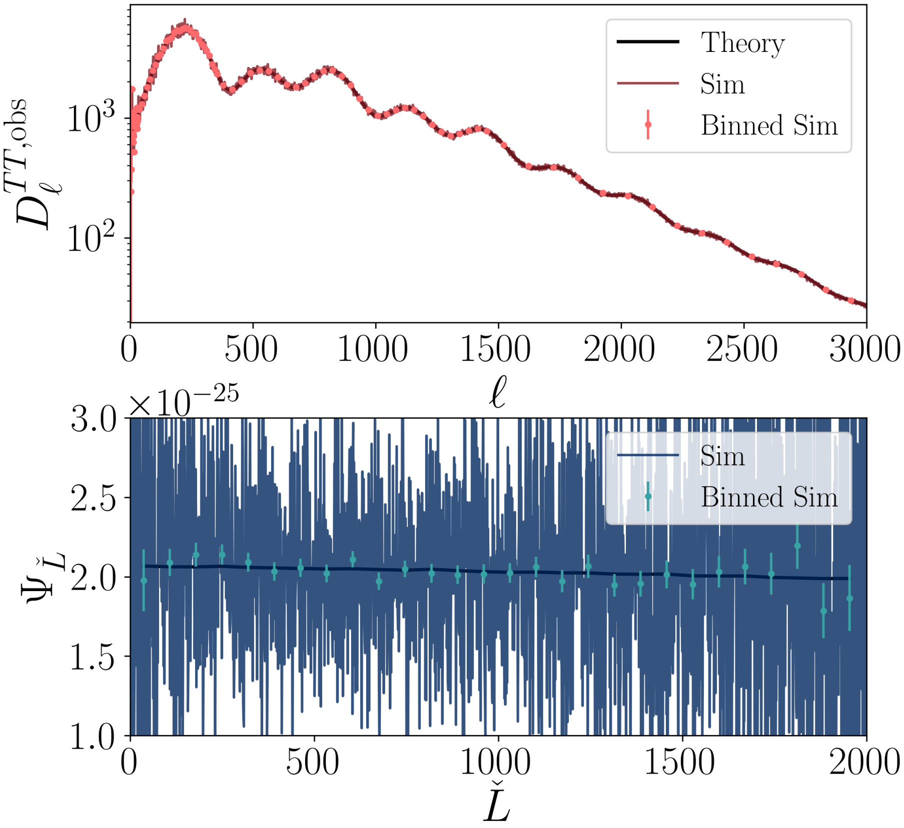

Figure 6.7: *Top*: A comparison between the theoretical $TT$ power spectrum (with added noise $N_\ell^{TT}$) from `CAMB` with an observed power spectrum from one realization. The band-powers are binned with the prescription in Table 6.5 are overlaid with error-bars corresponding to the diagonal of Figure 6.4. *Bottom*: A comparison between the theoretical SCALE observable $\Psi_{\check{L}}$ from `cmbpix` with an observed cross spectrum from one realization. Both are normalized with $A_{\check{L}}$ following the fiducial model in Table 6.2. The band-powers binned with width $\Delta\check{L} = 71$ are overlaid with error-bars corresponding to the diagonal of Figure 6.5. The theory spectrum has also been binned to reduce the scatter from the Monte Carlo integration for visual presentation. Both binning schemes accurately reflect the overall shape of their respective spectra while reducing the noise of each individual entry in the data vector.

Finally, we have empirical estimates of the covariance between all band-powers of the data vector from our set of simulations in §6.3: $\mathbf{C}$. This allows us to construct a multivariate normal log-likelihood $p_l(\vec{d}|\theta) = \ln(p(\vec{d}|\theta))$ for the data vector $\vec{d}$ given a set of parameters $\theta$:

$$p_l(\vec{d}|\theta) \sim \mathcal{N}_l(\vec{d}|\vec{t}(\theta), \hat{\mathbf{C}}^{-1}) \sim -\frac{1}{2}(\vec{d} - \vec{t}(\theta))^T \hat{\mathbf{C}}^{-1}(\vec{d} - \vec{t}(\theta)) \,. \tag{6.5}$$

Note that the covariance $\mathbf{C}$ is constructed with $N_{\text{sims}} = 400$ realizations of the data vector, so the unbiased estimator for the inverse covariance must include the Hartlap factor (Hartlap et al., 2007):

$$\hat{\mathbf{C}}^{-1} = \frac{N_{\text{sims}} - P - 2}{N_{\text{sims}} - 1}\mathbf{C}^{-1} \,, \tag{6.6}$$

where $P$ is the number of band-powers included in the data vector. For an analysis with only $TT$ band-powers $P_{TT} = 105$, and including SCALE increases it to $P_{TT+\text{SCALE}} = 133$. It is recommended that there is a minimum of realizations $N_{\text{sims}} \gtrsim 2P$, which we have satisfied (Hartlap et al., 2007).

We set broad, uniform priors $p_{\text{Uni}}(\theta)$ for every parameter in $\theta$ following Table 6.1 except for $\tau$, for which we impose a Gaussian prior about the fiducial value (Table 6.2):

$$p(\tau) \sim \mathcal{N}(0.06, \sigma_\tau) \,, \tag{6.7}$$

where we choose either $\sigma_\tau = 0.007$ set by the value reported by *Planck* 2018 (Planck Collaboration et al., 2020a), or $\sigma_\tau = 0.002$ set by the cosmic variance limit (LiteBIRD Collaboration et al., 2023). We also include an additional likelihood which include forecasted constraints from Baryon Acoustic Oscillation (BAO) information from the Dark Energy Spectroscopic Instrument (DESI, Font-Ribera et al. (2014)), which mainly constrains the matter density $\Omega_m$. We follow the steps in Appendix V of Allison et al. (2015) to construct a Fisher matrix $\mathbf{F}$ for BAO observables and the covariances between our other parameters, including $m_\nu$, $\Omega_c$, $\Omega_b$, and $\ln(10^{10}A_s)$. The BAO log-likelihood is constructed as follows:

$$p_{l,\text{BAO}}(\theta|\mathbf{F}) \sim -\frac{1}{2}(\theta - \theta_{\text{fid}})^T \mathbf{F}(\theta - \theta_{\text{fid}}) \,, \tag{6.8}$$

where $\theta_{\text{fid}}$ contains the fiducial parameters from Table 6.2. Our final posterior is then expressed by the following:

$$p(\theta|\vec{d}) \sim \mathcal{N}(\vec{d}|\vec{t}(\theta), \hat{\mathbf{C}}^{-1})\mathcal{N}(\tau|0.06, \sigma_\tau) p_{\text{BAO}}(\theta|\mathbf{F}) p_{\text{Uni}}(\theta) \,. \tag{6.9}$$

In principle, the data vector $\vec{d}$, covariance matrix $\mathbf{C}$, and theory vector constructed from emulators $\vec{t}(\theta)$ can contain any combination of $TT$ and SCALE band-powers as long as all three are constructed consistently. We present in §6.5 results from only $TT$ band-powers as well as results from a combination of $TT$ and SCALE band-powers. As mentioned in §6.3, we only consider one version of SCALE with $\ell_{1,\text{min}} = 8000$ and $\ell_{1,\text{max}} = 10000$. We expect that including multiple versions of SCALE with different $\ell_{1,\text{min}}$ and $\ell_{1,\text{max}}$ in a single fit could allow for constraints on models which change the shape of the lensing potential power spectrum $C_L^{\phi\phi}$, which is an exciting direction for future study. Our model including a massive neutrino simply shifts the amplitude of the lensing potential power in the small-scale regime that we consider, so a single SCALE estimator is sufficient for the purposes of this study.

We construct and sample our probabilistic model with the Python implementation of Markov

Chain Monte Carlo (MCMC) techniques in `emcee`[9] (Foreman-Mackey et al., 2013). We choose to use `emcee` rather than more commonly-used software designed specifically for cosmology such as `CosmoMC`[10] (Lewis & Bridle, 2002) or `cobaya`[11] (Torrado & Lewis, 2019, 2021) because it offers a simple way to construct log-probabilities with the added flexibility of allowing for the use of black-box functions in the model. The latter point is essential in order to use the emulators constructed in §6.2 at each step of the chain. We use 14 walkers, or chains, with 10 000 steps each. The first 1 000 steps are discarded as burn-in steps that have yet to converge, and we further thin the chains by a factor of 50 to reduce the auto-correlation between samples. We find that the chains converge (satisfying the Gelman-Rubin ratio requirement $R < 1.1$) after approximately 1 000 steps post-burn-in and before thinning. The results of each model are presented in §6.5.

## 6.5 Results

The results for model fits using the realization of observables presented in Figure 6.7 are summarized in Table 6.6 and Figure 6.8-6.9. We find that the size of the 68% confidence region for the best fit results do not change appreciably if we choose a different realization for the data vector. The center of the best fit can vary slightly between realizations, but the change is generally well within the 68% range. The model presented in this work allows the $TT$ band-powers, combined with the BAO likelihood and $\tau$ prior, to constrain most of the parameters $\theta$ to high precision in the presence of noise levels similar to a futuristic S4 experiment. The exception is $m_\nu$, for which the $TT$ only model with a *Planck* 2018 prior on $\tau$ is not able to detect ($1.7\sigma$). The constraint value we present is a signal-to-noise of the fiducial value $m_\nu = 0.06\,\mathrm{eV}$ divided by half the size of the 68% confidence limit. We also see in Figure 6.9 that the marginalized posterior for $m_\nu$ using only the $TT$ band-powers as observables causes the distribution to hit the edge of the prior at $0\,\mathrm{eV}$. We conclude that the $TT$ band-powers alone cannot provide a significant detection of $m_\nu$ without additional information from lensing.

The addition of SCALE into the data vector affects the parameters most sensitive to lensing (see Figure 6.8): $m_\nu$, $\Omega_c$, and $\Omega_b$. We see that the addition of SCALE alters the degeneracies between these parameters to become more constraining. Perhaps the most salient effect from adding SCALE is the added ability to provide evidence for non-zero $m_\nu$ at $2.4\sigma$. This has yet to be achieved with cosmological evidence. The effect is more prominent if we swap the $\tau$ prior to the cosmic variance limit $\sigma_\tau = 0.002$ (see Table 6.6), which is forecasted to be achievable with upcoming data from the LiteBIRD satellite mission (LiteBIRD Collaboration et al., 2023) or the CLASS ground-based survey (Essinger-Hileman et al., 2014). In this case the detection jumps from $2.2\sigma$ with only $TT$ observables to $4\sigma$ with the inclusion of SCALE. We have thus demonstrated that the addition of small-scale lensing information with SCALE into a cosmological model fit provides extra constraining power on parameters which alter the lensing amplitude at $\ell \gg 3000$.

## 6.6 Discussion and Conclusions

In this chapter, we:

---

[9] https://github.com/dfm/emcee

[10] https://github.com/cmbant/CosmoMC

[11] https://github.com/CobayaSampler/cobaya

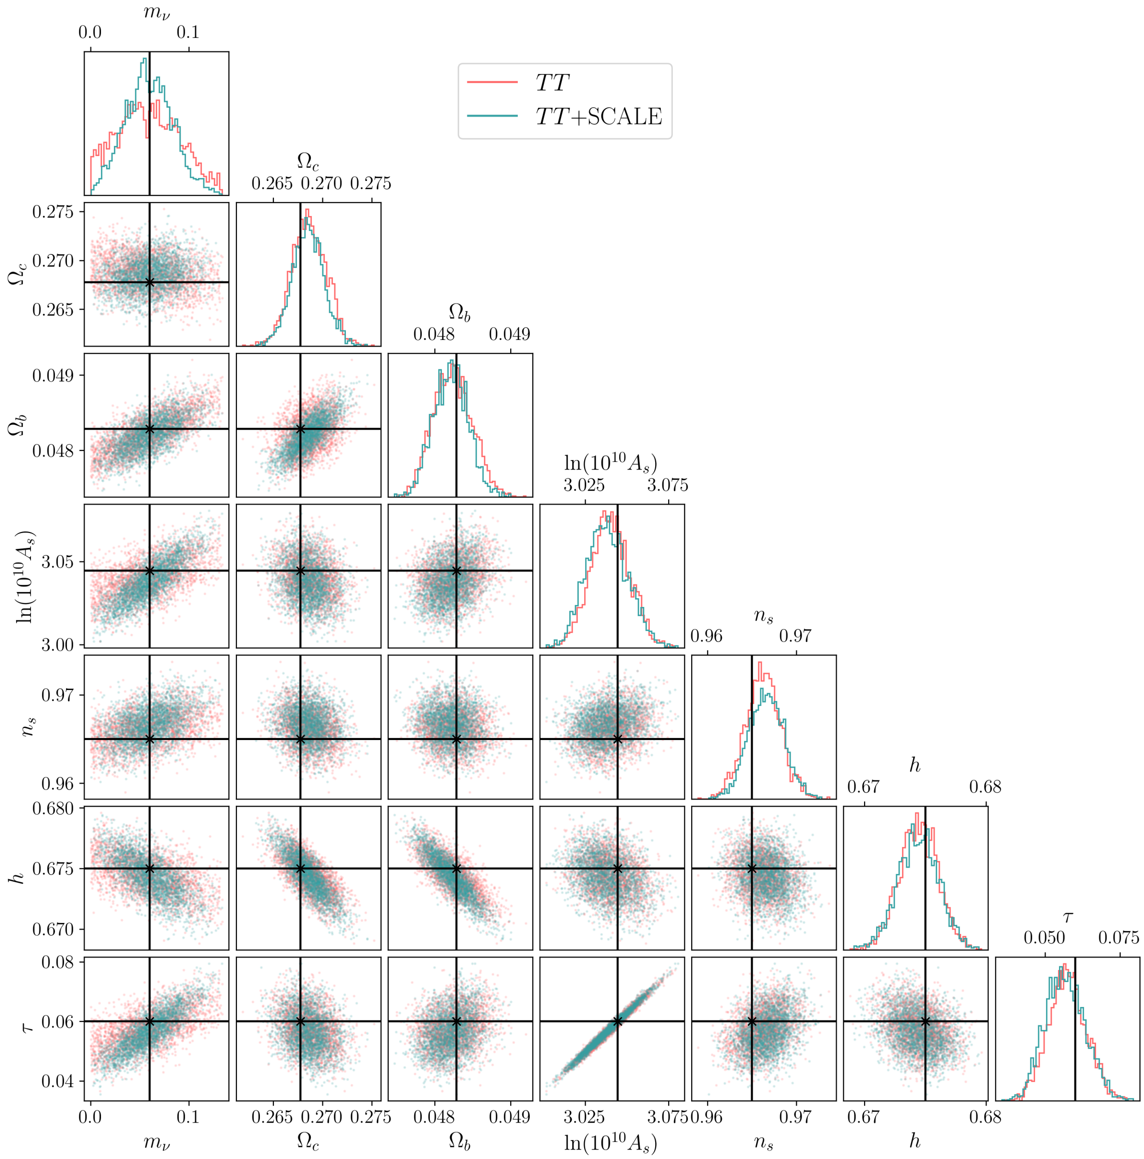

Figure 6.8: The sampled posterior distribution of the probabilistic model described in Equation (6.9) using the $TT$ band powers along with *Planck* 2018 $\tau$ prior $\sigma_\tau = 0.007$ and BAO likelihood is shown in *pink*. The resulting sampled posterior distribution with the addition of SCALE is shown in *green*. The fiducial values are indicated with black lines. The addition of SCALE observables makes the difference in a detection of the minimum neutrino mass $m_\nu$. We also see that the covariances between matter clustering related parameters $m_\nu$, $\Omega_c$ and $\Omega_b$ tighten up with the addition of SCALE.

Table 6.6: Summary of cosmological parameter constraints from the observed `TT` power spectrum and SCALE cross spectrum, along with $\tau$ and BAO priors. Fiducial values from Table 6.2 are also included. Columns marked with $P$ use the *Planck* 2018 $\tau$ prior $\sigma_\tau = 0.007$ (Planck Collaboration et al., 2020a), and columns marked with $CV$ use the cosmic variance limit $\tau$ prior $\sigma_\tau = 0.002$ (LiteBIRD Collaboration et al., 2023). Reported values are the median and 68% confidence interval of the marginalized posterior distribution of each parameter. The median values of the fit can slightly shift around the fiducial values depending on the realization, but in general the shifts are comfortably within the 68% confidence intervals. The widths of the 68% confidence regions do not change between realizations.

| Parameter | Fiducial | $TT^P$ | $TT+\mathrm{SCALE}^P$ | $TT^{CV}$ | $TT+\mathrm{SCALE}^{CV}$ |
|---|---|---|---|---|---|
| $m_\nu$ [eV] | 0.06 | $0.057^{+0.035}_{-0.033}$ | $0.058^{+0.025}_{-0.024}$ | $0.065^{+0.027}_{-0.027}$ | $0.068^{+0.014}_{-0.015}$ |
| $\Omega_c$ | 0.2678 | $0.2686^{+0.0019}_{-0.0017}$ | $0.2685^{+0.0016}_{-0.0016}$ | $0.2687^{+0.0018}_{-0.0018}$ | $0.2683^{+0.0016}_{-0.0015}$ |
| $\Omega_b$ | 0.04829 | $0.04823^{+0.0003}_{-0.0003}$ | $0.04821^{+0.0002}_{-0.0002}$ | $0.04828^{+0.0003}_{-0.0003}$ | $0.04823^{+0.0002}_{-0.0002}$ |
| $\ln(10^{10}A_s)$ | 3.045 | $3.0440^{+0.012}_{-0.011}$ | $3.038^{+0.013}_{-0.012}$ | $3.044^{+0.004}_{-0.004}$ | $3.044^{+0.004}_{-0.004}$ |
| $n_s$ | 0.965 | $0.966^{+0.002}_{-0.002}$ | $0.967^{+0.002}_{-0.002}$ | $0.966^{+0.002}_{-0.002}$ | $0.967^{+0.002}_{-0.002}$ |
| $h$ | 0.675 | $0.674^{+0.002}_{-0.002}$ | $0.674^{+0.002}_{-0.002}$ | $0.674^{+0.002}_{-0.002}$ | $0.674^{+0.002}_{-0.002}$ |
| $\tau$ | 0.06 | $0.057^{+0.006}_{-0.006}$ | $0.056^{+0.007}_{-0.006}$ | $0.060^{+0.002}_{-0.002}$ | $0.060^{+0.002}_{-0.002}$ |

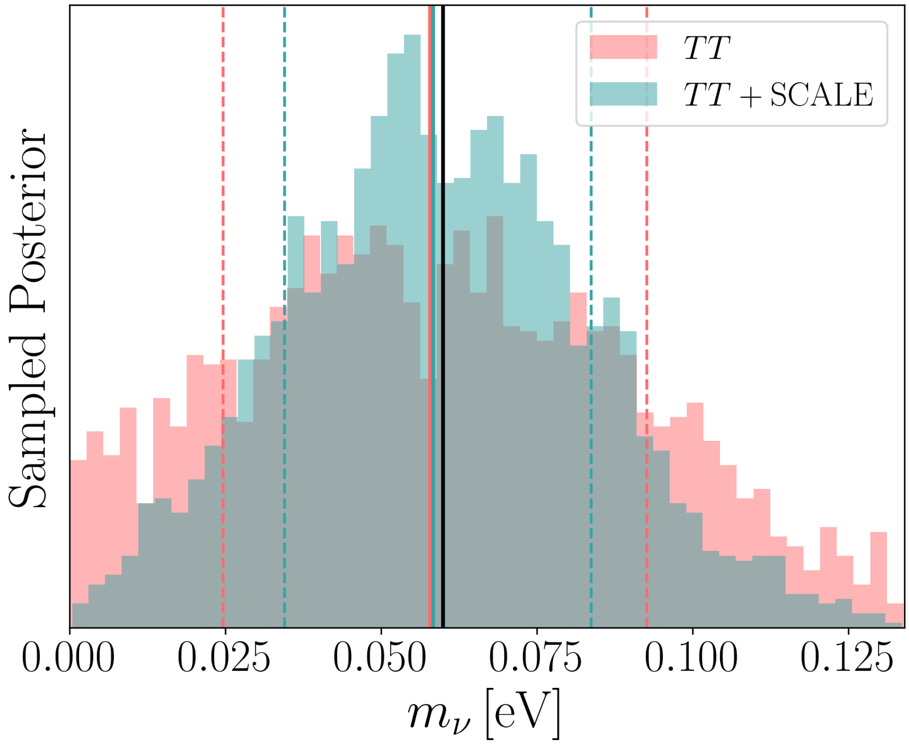

Figure 6.9: The 1-dimensional marginalized posterior for $m_\nu$ with (*green*) and without (*pink*) SCALE (same chains as Figure 6.8). The median of each result is indicated with solid vertical lines of their respective colours, and dashed lines indicate their 68% confidence intervals. A vertical black line indicates the fiducial value at $m_\nu = 0.06\,\mathrm{eV}$. The *Planck* 2018 $\tau$ prior $\sigma_\tau = 0.007$ is used here in addition to the $TT$ band-powers and BAO likelihood. The addition of SCALE not only tightens the distribution around the predicted neutrino mass $m_\nu$, but it also prevents its samples from hitting the 'wall' of the prior at $m_\nu = 0\,\mathrm{eV}$.

- Constructed neural network emulators for the lensed $TT$ power spectrum and the SCALE cross spectrum, which provide quick mapping from cosmological parameters to expected observables at $\sim 0.01\%$ and $\sim$ 2-5% precision per mode respectively (Figure 6.2-6.3 and Table 6.4).

- Presented a procedure to simulate a large sample of high-resolution (`NSIDE=8192`), full-sky simulations of the lensed CMB with the `lenspyx` package. We also present a procedure to compute SCALE observables from these full-sky `HEALPix` representations, and SCALE observables from this procedure match well with the flat-sky results in Chapter 5.

- Developed a likelihood which includes SCALE in parameter estimation using CMB observables for a standard ΛCDM model with the addition of a massive neutrino $m_\nu$.

- Demonstrated that SCALE can directly provide constraining information in the estimation of parameters, such as $m_\nu$, which affect the amplitude of small-scale lensing beyond measurements of the lensed CMB power spectrum.

Some avenues for additional study include:

- Including multiple versions of SCALE with different $\ell_1$ filters in a single fit to constrain models which change the shape of the lensing potential power spectrum $C_L^{\phi\phi}$.

- Comparing the SCALE method with other methods which use small-scale lensing information, such as the $TT$ quadratic estimator (QE; Hu & Okamoto 2002; Hu et al. 2007). An additional comparison with combined CMB + QE + SCALE observables can show whether or not SCALE provides extra lensing information over quadratic estimator techniques.
  - We expect similar constraints on $m_\nu$ can be obtained with the standard QE along with the same $TT$ band-powers, $\tau$ prior, and BAO information. We discussed in Chapter 5 the advantages SCALE over the QE, including its higher signal-to-noise detection of the lensing at similar small-scale regimes.
  - Addition of the QE observables (band-powers of $C_L^{\phi\phi}$) to the likelihood is straightforward, and the emulators constructed in §6.2 can be applied to predict the $C_L^{\phi\phi}$ band-powers similar to $\tilde{C}_\ell^{TT}$.

- Investigating the effects of foreground contamination on the small-scale statistics used in SCALE. Our extension of SCALE to the full-sky in §6.3 can allow for analysis with the WebSky CMB mock simulations (Stein et al., 2020), which include simulations of foreground contamination from realistic dark matter halos generated with the mass-Peak Patch method (Stein et al., 2019).

- Investigating the effects of masking in flat-sky observations in the cross spectrum of SCALE intermediate products $\lambda$ and $\varsigma$. A common mask applied to a CMB temperature mask carries through into the construction of $\lambda$ and $\varsigma$, so a proper treatment of masking is required to ensure that the cross spectrum is not biased in the application of SCALE in ground-based observations.

We have shown that the addition of SCALE into a cosmological model fit can provide significant evidence for non-zero $m_\nu$ with a Stage IV-like experiment, which bodes well for the application of

SCALE in future data. The effect of massive neutrinos on the lensing power spectrum is a nearly scale-independent decrease in amplitude at high-$\ell$, so including one version of SCALE with a single small-scale filter $\ell_1$ is sufficient for a detection. We expect that including multiple versions of SCALE with different $\ell_1$ filters would allow for constraints on models which change the shape of the lensing potential power spectrum $C_L^{\phi\phi}$, which we leave for future study. This opens a wide window of opportunity for SCALE to provide constraints on exotic forms of dark matter or clustering models which are predicted to have non-trivial effects on the shape of the lensing potential power spectrum.

# Chapter 7

# Conclusions

Physical phenomena observed at early- and late-times appear to conflict with ΛCDM when one begins to scrutinize the model's predictions, and this thesis represents a collection of studies to further our understanding of some of these inconsistencies. We cover a range of cosmic scales spanning almost the entire history of the Universe, from the local distance ladder to the cosmic microwave background. We began in Chapter 1 with an introduction into modern cosmological principles. This led into a discussion of the cosmic microwave background, which began with a history of its observations. We then described the physical processes that generate the signals within it, with a particular focus on the signatures left behind by the effects of gravitational lensing. Finally, we outlined a general procedure for observing the cosmic microwave background from detectors all to way to the inference of cosmological parameters from power spectra. We continued with a discussion on cosmic distance measures and their importance in the determination of the Hubble constant. This culminated in our presentation of the key issues in the field which motivate the work in this dissertation, followed by some key statistical and computational techniques that we employed. The key issues we address in this thesis include the Hubble tension, the clustering of matter at small-scales, and the potential detection of primordial gravitational waves in the CMB polarization signal.

In Chapter 2, we presented a hierarchical Bayesian model for simultaneously calibrating a luminosity model for red clump stars and estimating the parallax zero point systematic in *Gaia* Data Release 2. We inferred the *Gaia* DR2 parallax zero point to be $\varpi_0 = -48 \pm 1\,\mu$as, which is the most precise determination for DR2. We also studied extensions to our model to characterize separate luminosity models for subpopulations of red clump stars of different metallicities. We also included several additions to the probabilistic model to determine the parallax zero point's variations with observed magnitude, observed colour, and sky position. This work helped to further understand the systematics present in *Gaia* astrometry, and the subsequent release of *Gaia* Early Data Release 3 included a software to compute unique parallax zero points for each source in the catalogue based on their observed magnitude, observed colour sky position, and calibration coefficient.

Chapter 3 follows in the steps of verifying cosmic measures of distance with *Gaia* data by constructing a model for the distance to the Large Magellanic Cloud. We used the low-resolution $XP$ spectral data in *Gaia* Data Release 3 to select a sample of bright stars in the LMC with similar spectra to stars nearby with high quality *Gaia* parallaxes. We were able to infer the LMC distance to be $R_{\rm LMC} = 48.90^{+0.52}_{-0.48}\,$kpc, which is in agreement with the value reported by Pietrzyński et al. (2019), but we also show that the LMC distance inferred by our models is sensitive to the parallax zero point systematics. This study shows that the high precision and extensive size of the *Gaia* dataset can be exploited to independently solidify uncertainties for geometric distance anchors in the local distance ladder, and future data releases will only improve the precision of *Gaia* distances.

Chapter 4 sees a shift in focus towards observations of the cosmic microwave background. We present a study into the effects of detector crosstalk systematics on forecasted observables for a LiteBIRD-like experiment. We implemented a crosstalk model into the `TOAST` software, which simulates the scanning motions of a CMB satellite experiment to generate TODs and then reproject them into 'observed' CMB maps of temperature and polarization. Our crosstalk model mixes detector TODs in `TOAST` according to a pre-defined crosstalk matrix that represents realistic electrical crosstalk. This study showed that the effects of unmitigated electrical crosstalk can bias observed $EE$ and $BB$ power spectra at large angular scales ($\ell \lesssim 100$) in a LiteBIRD-like experiment by at most a few percent. We also showed that the effects have non-trivial dependence on the ordering of the detectors from the focal plane to the readout electronics. The results of this study will help to inform the design and construction of the electronic readout system for the proposed LiteBIRD mission.

Chapter 5 continues studies of the CMB, but shifts the focus once again to the effects of gravitational lensing at small angular scales ($\ell > 3000$). We presented the development of the Small Correlated Against Large Estimator, which takes advantages of the correlations between the large-scale CMB temperature gradient power, and the small-scale CMB gradient power generated by gravitational lensing. We showed that the SCALE method can recover the lensing statistics in well-controlled maps of simulated CMB temperature data to better precision than traditional quadratic estimator techniques. This study establishes a solid foundation for this novel technique, which will continue to perform existing methods as CMB experiments become increasingly sensitive to the effects of small-scale lensing.

In Chapter 6, we proceeded to apply the SCALE procedure to a realistic parameter estimation pipeline. We constructed and trained emulators that significantly speed up the prediction of CMB $TT$ power spectra, and SCALE cross-spectra from cosmological parameters at a given step in sampling space. We constructed a likelihood for CMB $TT$ power spectra, and combined it with a *Planck* 2018 prior for $\tau$ and a forecasted BAO likelihood for $m_\nu$, $\Omega_c$, $\Omega_b$, and $A_s$ to show that the minimum neutrino mass $m_\nu = 0.06\,$eV is not detected at high significance with only CMB data with CMB-S4-like noise levels. We demonstrate that the inclusion of SCALE in the CMB likelihood adds enough information for a $2.4\sigma$ detection of the minimum neutrino mass in the same conditions, increasing to a $4\sigma$ detection with a cosmic variance limit measurement of $\tau$. This study provides an example of how to include SCALE products in cosmological parameter estimation, and it also shows that SCALE observables are sensitive to parameters that limit their effects to the small-scale regime of the lensing potential power spectrum $C_L^{\phi\phi}$. This opens up the possibility of including multiple versions of SCALE at different multipole ranges to constrain models that affect the shape of $C_L^{\phi\phi}$ at small-scales.

# Appendix A

# Map-based lensing estimator

MAP space statistics of the large-scale CMB temperature gradient and the small-scale CMB temperature power are strongly correlated through the effects of gravitational lensing. We show this in Chapter 5, specifically in §5.3. We present here a prescription to quantify the correlations in real space as well as a discussion on the challenges that it offers. The gradient operator $\boldsymbol{\nabla}_{\hat{\mathbf{n}}}$ contains two components: $\nabla_x$ and $\nabla_y$. The directions $x$ and $y$ generally describe two orthogonal directions on the plane of the sky. These may be the map coordinate directions in a flat-sky representation. For a full-sky map, it is convenient to use the polar angle $\theta$ and the azimuthal angle $\phi$:

$$\nabla_x = \frac{\partial}{\partial\theta}\,, \text{ and} \tag{A.1}$$

$$\nabla_y = \frac{1}{\sin\theta}\frac{\partial}{\partial\phi}\,. \tag{A.2}$$

Beginning with Equation (1.26) in the limit of small scales where $T \approx 0$ due to diffusion damping, the small-scale temperature power that we expect to observe in a given patch of sky is:

$$T^{\mathrm{S}}(\hat{n}) \approx \boldsymbol{\nabla}_{\hat{\mathbf{n}}}\phi \cdot \boldsymbol{\nabla}_{\hat{\mathbf{n}}}T \tag{A.3}$$

$$\begin{aligned}\langle T^2\rangle^{\mathrm{S}} &\approx \langle(\boldsymbol{\nabla}_{\hat{\mathbf{n}}}\phi \cdot \boldsymbol{\nabla}_{\hat{\mathbf{n}}}T)^2\rangle \\ &\approx \langle|\boldsymbol{\nabla}_{\hat{\mathbf{n}}}\phi|^2\rangle\langle|\boldsymbol{\nabla}_{\hat{\mathbf{n}}}T|^2\rangle\langle\cos^2\xi\rangle \\ &\approx \frac{1}{2}\langle|\boldsymbol{\nabla}_{\hat{\mathbf{n}}}\phi|^2\rangle\langle|\boldsymbol{\nabla}_{\hat{\mathbf{n}}}T|^2\rangle\,. \end{aligned} \tag{A.4}$$

In the second step, we expand the dot product using the geometric definition of the dot product, and the separation angle between $\boldsymbol{\nabla}_{\hat{\mathbf{n}}}\phi$ and $\boldsymbol{\nabla}_{\hat{\mathbf{n}}}T$ in the plane of the sky is $\xi$. The third step is taken under the assumption that both the lensing deflection field, and the temperature gradient are isotropic and independent from one another. The expression we end up with is similar to the small-scale lensing temperature power approximated in Lewis & Challinor (2006, §4.1.3, Equation (4.16)-(4.17))

$$C_\ell^{TT} \approx \ell^2 C_\ell^{\phi\phi} R^T\,, \text{ where} \tag{A.5}$$

$$R^T \equiv \frac{1}{2}\langle|\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T|^2\rangle = \frac{1}{4\pi}\int \frac{\mathrm{d}\ell}{\ell}\ell^4 C_\ell^{TT}\,. \tag{A.6}$$

We can see from the above steps that if one collects statistics of the small-scale temperature power $\langle T^2\rangle^{\mathrm{S}}_{\mathrm{patch}}$ as well as the average background temperature gradient $\langle|\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T|^2\rangle$ locally on the sky in small patches (refer back to Figure 5.3), then we roughly expect them to be correlated to each other through Equation (A.4). We can do better by directly computing the two-dimensional integral for the total observed temperature at small-scales between $\ell_{\mathrm{min}}$ and $\ell_{\mathrm{max}}$:

$$\langle T^2\rangle^{\mathrm{S}}_{\mathrm{patch}} = \frac{1}{(2\pi)^2}\int_{\boldsymbol{\ell}_{\mathrm{min}}}^{\boldsymbol{\ell}_{\mathrm{max}}} \mathrm{d}^2\boldsymbol{\ell}\, C_{\ell,\mathrm{S}}^{TT}\,. \tag{A.7}$$

Let us first construct a version of Equation (A.5) while preserving the direction of the gradient rather than simply taking its amplitude. On the flat sky, the 2D Fourier representation ($\mathcal{FT}$) of a scalar field is

$$T(\mathbf{x}) \sim \int \mathrm{d}^2\boldsymbol{\ell}\, T(\boldsymbol{\ell}) e^{-i\boldsymbol{\ell}\cdot\mathbf{x}}\,, \tag{A.8}$$

and the corresponding representation of its gradient is

$$\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T(\mathbf{x}) \sim -i\int \mathrm{d}^2\boldsymbol{\ell}\, T(\mathbf{x})\boldsymbol{\ell} e^{-i\boldsymbol{\ell}\cdot\mathbf{x}}\,. \tag{A.9}$$

We start from Equation A.3, keeping in mind that the CMB temperature gradient gets the lion's share of its contributions from multipoles $\ell \lesssim 2000$ (Lewis & Challinor, 2006; Hu et al., 2007). We also include contributions $T_0^{\mathrm{S}}(\boldsymbol{\ell})$ from the relatively low amplitude contribution from the primary CMB, noise, and foregrounds that do *not* correlate directly with $\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}$. Locally (on a small patch of sky located at position $\mathbf{x}$), its 2D Fourier representation on the flat sky is

$$\tilde{T}^{\mathrm{S}}(\boldsymbol{\ell}) = -i\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}\cdot\boldsymbol{\ell}\phi(\boldsymbol{\ell}) + T_0^{\mathrm{S}}(\boldsymbol{\ell})\,, \tag{A.10}$$

Note that Equation (A.10) without $T_0^{\mathrm{S}}(\boldsymbol{\ell})$ is exactly the estimator constructed by Hadzhiyska et al. (2019). The small-scale CMB power is $\langle TT^*\rangle$. We continue under the assumption that the two terms in Equation (A.10) do not correlate with each other. We can then write the small-scale CMB power as:

$$C_{\ell,\mathrm{S}}^{TT} = (\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}\cdot\boldsymbol{\ell})^2 C_\ell^{\phi\phi} + C_{\ell,0}^{TT}\,. \tag{A.11}$$

Our integral from Equation (A.7) then becomes

$$\begin{aligned}
\langle T^2\rangle^{\mathrm{S}}_{\mathrm{patch}} &= \frac{1}{(2\pi)^2}\int_{\boldsymbol{\ell}_{\mathrm{min}}}^{\boldsymbol{\ell}_{\mathrm{max}}} \mathrm{d}^2\boldsymbol{\ell}\, C_\ell^{TT} &\text{(A.12)}\\
&= \frac{1}{(2\pi)^2}\int_{\boldsymbol{\ell}_{\mathrm{min}}}^{\boldsymbol{\ell}_{\mathrm{max}}} \mathrm{d}^2\boldsymbol{\ell}\Big[(\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}\cdot\boldsymbol{\ell})^2 C_\ell^{\phi\phi} + C_{\ell,0}^{TT}\Big] &\text{(A.13)}\\
&= \frac{1}{(2\pi)^2}\left[|\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}|^2 \int_{\ell_{\mathrm{min}}}^{\ell_{\mathrm{max}}} \mathrm{d}^2\boldsymbol{\ell}(\nabla\hat{T}^{\mathrm{L}}\cdot\hat{\ell})^2\ell^2 C_\ell^{\phi\phi} + \int_{\ell_{\mathrm{min}}}^{\ell_{\mathrm{max}}} \mathrm{d}^2\boldsymbol{\ell}\, C_{\ell,0}^{TT}\right]. &\text{(A.14)}
\end{aligned}$$

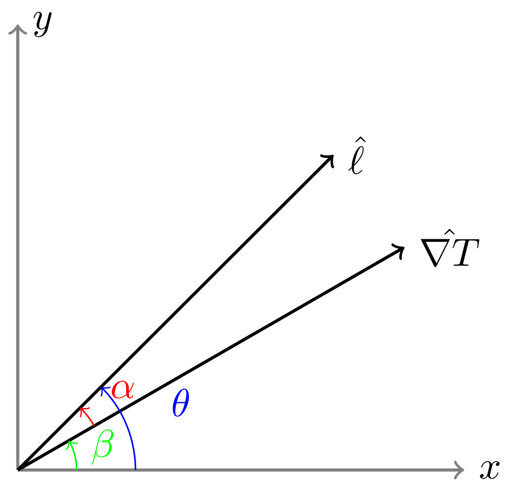

Figure A.1: Illustrating our change of angular coordinates.

Within a specific patch, we chose in §5.3 to make the assumption that $\boldsymbol{\nabla}_{\hat{\mathbf{n}}}T^{\mathrm{L}}$ is constant; therefore, we treat $|\boldsymbol{\nabla}_{\hat{\mathbf{n}}}T^{\mathrm{L}}|$ and $\hat{\nabla T}^{\mathrm{L}}$ as constant here. We now continue by expanding $\mathrm{d}^2\boldsymbol{\ell} = \ell\mathrm{d}\theta\mathrm{d}\ell$. We also replace $(\hat{\nabla T}^{\mathrm{L}} \cdot \hat{\ell})^2 = \cos^2\alpha$, where $\alpha$ is the angle between $\hat{\nabla T}^{\mathrm{L}}$ and $\hat{\ell}$.

$$\langle T^2\rangle^{\mathrm{S}}_{\mathrm{patch}} = \frac{|\boldsymbol{\nabla}_{\hat{\mathbf{n}}}T^{\mathrm{L}}|^2}{(2\pi)^2}\int_{\ell_{\min}}^{\ell_{\max}}\cos^2\alpha\,\ell^3 C_\ell^{\phi\phi}\mathrm{d}\theta\mathrm{d}\ell + \frac{1}{(2\pi)^2}\int_{\ell_{\min}}^{\ell_{\max}}\ell C_{\ell,0}^{TT}\mathrm{d}\theta\mathrm{d}\ell\,. \tag{A.15}$$

Here, we remind ourselves that $\theta$ describes the direction of $\boldsymbol{\ell}$, or more specifically, the angle between $\boldsymbol{\ell}$ and an arbitrary axis (like the $x$ axis). We have the freedom to choose this arbitrary axis, so we can simply change our coordinates to $\theta = \alpha + \beta \to \alpha = \theta - \beta$, where $\beta$ is the angle between $\hat{\nabla T}^{\mathrm{L}}$ and the arbitrary axis $x$. There are many angles and directions here, so Figure A.1 may help with visualizing this. Since $\beta$ is constant within each patch, this change of coordinates is straightforward to apply; i.e. $\mathrm{d}\alpha = \mathrm{d}\theta$. Evaluating the integral is also simple because $C_\ell^{\phi\phi}$ is independent of our choice of coordinate system and any of the angles we are working with.

$$\langle T^2\rangle^{\mathrm{S}}_{\mathrm{patch}} = \frac{|\boldsymbol{\nabla}_{\hat{\mathbf{n}}}T^{\mathrm{L}}|^2}{(2\pi)^2}\int_0^{2\pi}\cos^2\alpha\mathrm{d}\alpha\int_{\ell_{\min}}^{\ell_{\max}}\ell^3 C_\ell^{\phi\phi}\mathrm{d}\ell + \frac{1}{(2\pi)^2}\int_0^{2\pi}\mathrm{d}\theta\int_{\ell_{\min}}^{\ell_{\max}}\ell C_{\ell,0}^{TT}\mathrm{d}\ell \tag{A.16}$$

$$= \frac{|\boldsymbol{\nabla}_{\hat{\mathbf{n}}}T^{\mathrm{L}}|^2}{4\pi}\int_{\ell_{\min}}^{\ell_{\max}}\ell^3 C_\ell^{\phi\phi}\mathrm{d}\ell + \frac{1}{2\pi}\int_{\ell_{\min}}^{\ell_{\max}}\ell C_{\ell,0}^{TT}\mathrm{d}\ell\,. \tag{A.17}$$

This allows us to directly compute an expected value for the slope in Figure 5.3 along with an expression for the residual power which does not correlate with the CMB gradient (i.e. the $y$-intercept).

Consider the application of an additional filter $f_\ell$ to the small scale CMB temperature. We motivate this in Chapter 5 to include additional information from the lensing operation that is neglected when the gradient amplitude is taken in Equation (A.13)-(A.14). In this section, we

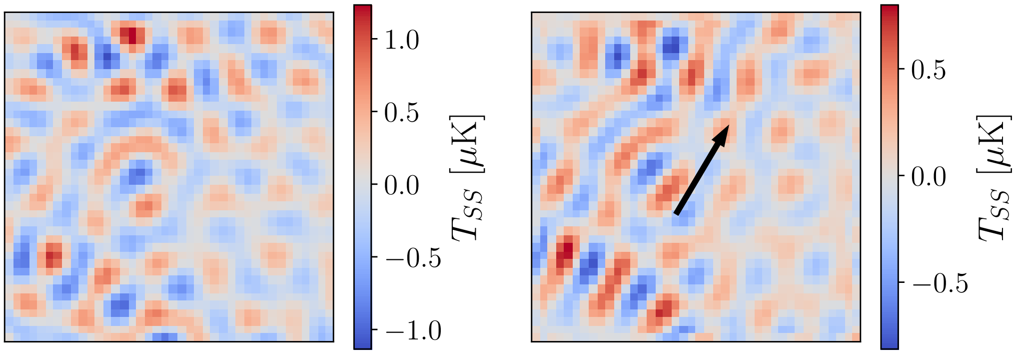

Figure A.2: *Left*: A patch of a CMB temperature field with a top hat filter applied such that only modes $6000 < \ell < 8000$ are included. *Right*: The same patch with an additional filter defined with respect to the average CMB temperature gradient direction within the patch: $f_\ell = (\hat{\nabla T} \cdot \hat{\ell})$. The features in the patch are exaggerated along the direction of the gradient.

compare the above derivation to one with an additional filter $f_\ell = (\hat{\nabla T} \cdot \hat{\ell}) = \cos\alpha$ for each patch. This has the effect of exaggerating the small-scale temperature signal along the direction of the background gradient (see Figure A.2), which are the features we expect to see from lensing. For each patch with the filter $f_\ell = \cos\alpha$:

$$\begin{aligned}
\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}} &= \frac{1}{(2\pi)^2} \int_{\boldsymbol{\ell}_{\mathrm{min}}}^{\boldsymbol{\ell}_{\mathrm{max}}} \mathrm{d}^2\boldsymbol{\ell} f_\ell^2 C_\ell^{TT} && \text{(A.18)} \\
&= \frac{|\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}|^2}{(2\pi)^2} \int_{\boldsymbol{\ell}_{\mathrm{min}}}^{\boldsymbol{\ell}_{\mathrm{max}}} \mathrm{d}^2\boldsymbol{\ell} (\hat{\nabla T^{\mathrm{L}}} \cdot \hat{\ell})^4 \ell^2 C_\ell^{\phi\phi} + \frac{1}{(2\pi)^2} \int_{\boldsymbol{\ell}_{\mathrm{min}}}^{\boldsymbol{\ell}_{\mathrm{max}}} \mathrm{d}^2\boldsymbol{\ell} (\hat{\nabla T^{\mathrm{L}}} \cdot \hat{\ell})^2 C_{\ell,0}^{TT} && \text{(A.19)} \\
&= \frac{|\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}|^2}{(2\pi)^2} \int_0^{2\pi} \cos^4\alpha \mathrm{d}\alpha \int_{\ell_{\mathrm{min}}}^{\ell_{\mathrm{max}}} \ell^3 C_\ell^{\phi\phi} \mathrm{d}\ell + \frac{1}{(2\pi)^2} \int_0^{2\pi} \cos^2\alpha \mathrm{d}\alpha \int_{\ell_{\mathrm{min}}}^{\ell_{\mathrm{max}}} \ell C_{\ell,0}^{TT} \mathrm{d}\ell && \text{(A.20)} \\
&= \frac{3|\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}|^2}{16\pi} \int_{\ell_{\mathrm{min}}}^{\ell_{\mathrm{max}}} \ell^3 C_\ell^{\phi\phi} \mathrm{d}\ell + \frac{1}{4\pi} \int_{\ell_{\mathrm{min}}}^{\ell_{\mathrm{max}}} \ell C_{\ell,0}^{TT} \mathrm{d}\ell\,. && \text{(A.21)}
\end{aligned}$$

We expect to see a factor of 3/4 suppression in the estimated slope in a filtered, lensed CMB map, as well as a factor of 1/2 in the estimated residual power in the $y$-intercept. Similarly, if we instead choose $f_\ell = \cos^2\alpha$, then we expect to find a factor of 5/8 suppression of the estimated slope, and a factor of 3/8 in the estimated residual power in the $y$-intercept:

$$\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}} = \frac{5|\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}|^2}{32\pi} \int_{\ell_{\mathrm{min}}}^{\ell_{\mathrm{max}}} \ell^3 C_\ell^{\phi\phi} \mathrm{d}\ell + \frac{3}{16\pi} \int_{\ell_{\mathrm{min}}}^{\ell_{\mathrm{max}}} \ell C_{\ell,0}^{TT} \mathrm{d}\ell\,. \tag{A.22}$$

For an example simulation of size 10° × 10° with a resolution of pixels with side-length 0.5 arcmin, we find the distribution of patch statistics shown in Figure A.3. For a quick first try at building a model to describe the distribution of patches relative to a line in $\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}}$ vs $|\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}|^2$ space, we may choose to approximate (rather poorly) the errors associated with estimating the $\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}}$ in each patch as Gaussian using the distribution of patches itself. Figure A.3 shows the errors in $\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}}$ taken to be the spread in the distribution of $\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}}$ across 20 bins in $|\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}|^2$ such that each bin

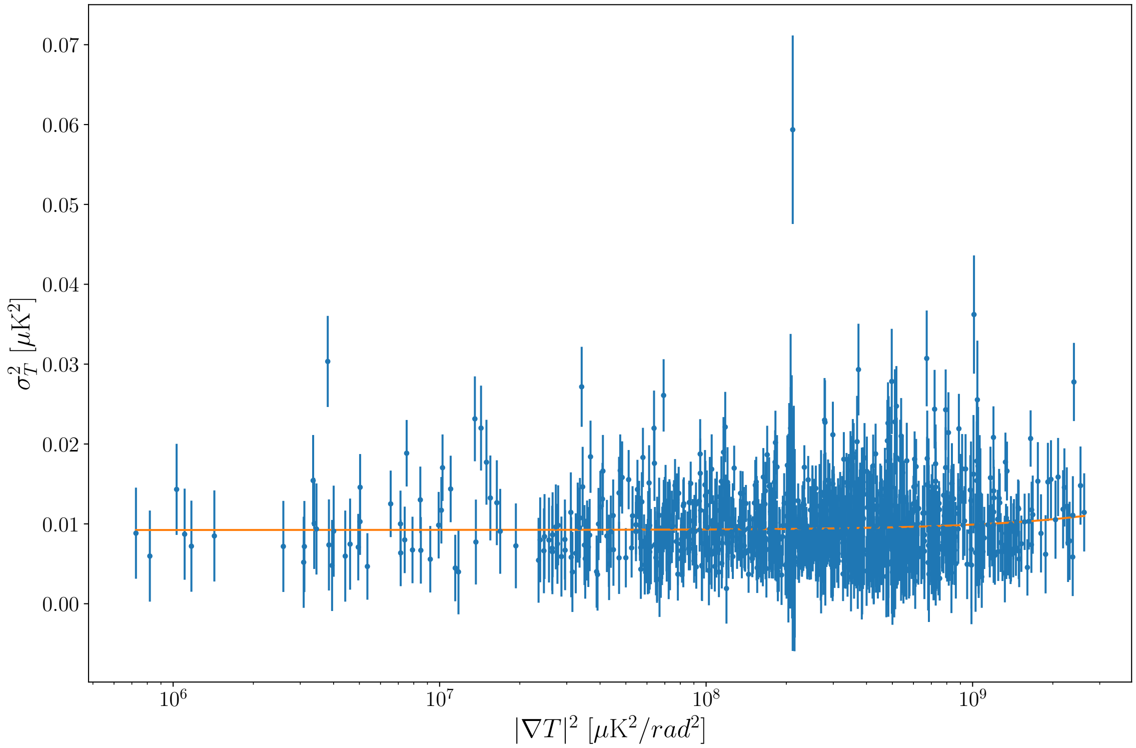

Figure A.3: *Blue*: The distribution of estimated patch statistics for an example simulation filtered for $\ell_{\min} = 6000$ and $\ell_{\max} = 8000$, and using $f_\ell = \cos\alpha$. The errors in $\langle T^2\rangle^{\rm S}_{\rm patch}$ (labelled as $\sigma_T^2$ on the $y$-axis) are taken to be the spread in the distribution of $\langle T^2\rangle^{\rm S}_{\rm patch}$ across 20 bins in $|\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\rm L}|^2$ such that each bin has an equal number of points. *Orange*: The line of best fit given by `scipy`'s `curve_fit` function. Each individual patch provides a poor estimate of the lensing correlations, but the collection of patches across a map allows for a statistically significant detection of lensing through the fit of the slope.

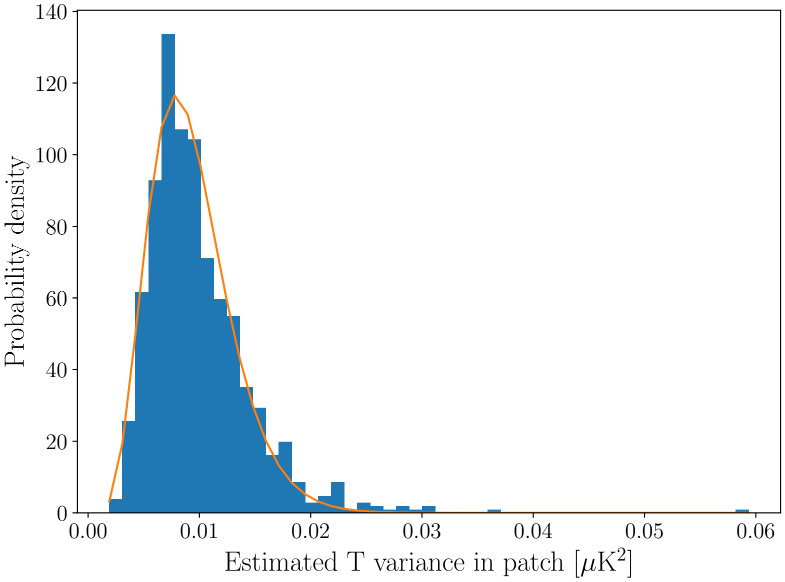

Figure A.4: *Blue:* The distribution of estimated $\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}}$ for an example simulation filtered for $\ell_{\mathrm{min}} = 6000$ and $\ell_{\mathrm{max}} = 8000$, and using the directional filter $f_\ell = \cos\alpha$. Patches were each constructed with $40 \times 40$ pixels of resolution $0.5'$. *Orange:* The best fit probability density function for a Pearson Type III distribution. The best fit parameter relevant to the number of independent samples used to estimate $\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}}$ is approximately $N = 14$.

has an equal number of points. Note that we have included the filter $f_\ell = \cos\alpha$ here.

We may use the distribution of patches along with a Gaussian likelihood to determine the best fit line for Equation (A.21). This likelihood would be:

$$p\big(\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}}, |\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}|^2 \big| m, b\big) \propto \exp\Big(\frac{-\chi^2}{2}\Big), \tag{A.23}$$

where $\chi^2$ is the chi-squared value for the distribution of patches with Gaussian errors described above relative to the line with intercept $b$ and slope $m$ computed with Equation (A.17) at the same $|\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}|^2$, (A.21) or (A.22). In this case, we used the filter $f_\ell = \cos\alpha$, so our best fit line corresponds to Equation (A.21). Application of the Gaussian likelihood to fit for the best fit parameters is equivalent to using `scipy`'s `curve_fit` routine, which we see the result for in Figure A.3. The Gaussian likelihood effortlessly returns a fit, but it does not well represent the distribution of $\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}}$ from the data. Figure A.4 shows the 1D distribution of estimated $\sigma_T^2$ in each patch, which has a significant positive skew.

Because we are using $\sim 40^2$ pixels to estimate the small scale temperature variance within each patch, the estimated sample variance that we actually get should follow a Pearson Type III distribution[1]:

$$p\big(\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}} \big| N, S^2\big) = \frac{\big(\frac{N}{2S^2}\big)^{(N-1)/2}}{\Gamma\big(\frac{N-1}{2}\big)} \big(\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}}\big)^{(N-3)/2} e^{-N\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}}/(2^2)}, \tag{A.24}$$

where $S^2$ is taken to be Equation A.21 evaluated at the $|\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}|^2$ associated with the $\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}}$ of each

---

[1] https://mathworld.wolfram.com/SampleVarianceDistribution.html

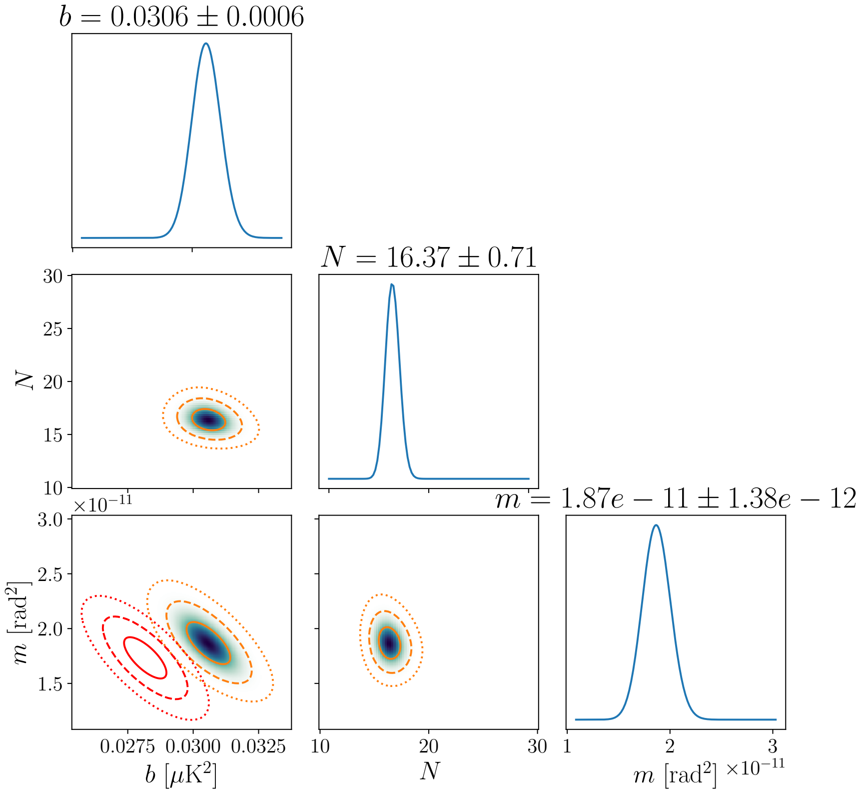

Figure A.5: The 3D-likelihood of the real-space estimator model using the Pearson Type III distribution on a single $10° \times 10°$ simulation of the lensed CMB. Patches were each constructed with $40 \times 40$ pixels of resolution $0.5'$. *Red:* The $1\sigma$, $2\sigma$, and $3\sigma$ contours for the model assuming a Gaussian likelihood. *Orange:* The $1\sigma$, $2\sigma$, and $3\sigma$ contours for this Pearson Type III likelihood. We see a significant effect in shifting the best fit slope and intercept parameters when accounting for the fact that individual pixels within a patch of sky are not fully independent from one another.

patch. Here, $N$ represents the number of independent samples used to estimate the sample variance. Naively, we would simply use $N \approx 40^2$, but it is important to remember that the neighbouring pixels are highly correlated for CMB observations. In practice, the true value of $N$ should be much smaller, as shown in Figure A.4, where the best fit Pearson type III distribution for the sample of estimated $\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}}$ of all patches has $N \approx 14$.

The full likelihood should then be the result of Equation A.24 evaluated using the $\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}}$ and their associated $|\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}|^2$ for all patches. As there are a large number of patches within the sample, this computation is likely to overflow, so the most practical way of computing the likelihood is as a sum in log-likelihood space:

$$\log\Big(p\big(\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}}, |\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\mathrm{L}}|^2 \big| m, b\big)\Big) \propto \sum \log\Big(p\big(\langle T^2 \rangle^{\mathrm{S}}_{\mathrm{patch}} \big| N, S^2\big)\Big). \tag{A.25}$$

We present the results of the fit with the Pearson type III likelihood in Figure A.5. The recovered intercept with the Pearson type III likelihood are approximately $5\sigma$ away from the best fit from the

Table A.1: Comparison of parameter signal-to-noise for different choices of directional filtering in the real-space lensing estimator. Values are averaged over 10 realizations of $10^\circ \times 10^\circ$ CMB observations with patches of size $20' \times 20'$. We find that the filter $f_\ell = \cos\alpha$ performs best at recovering the lensing signal contained within the slope of the $\langle T^2 \rangle^{\rm S}_{\rm patch}$ and $|\boldsymbol{\nabla}_{\hat{\mathbf{n}}} T^{\rm L}|^2$ correlations in real-space.

| Filter $f_\ell$ | Slope SNR | Intercept SNR |
|---|---|---|
| None | 4.8 | 19.7 |
| $\cos\alpha$ | 5.8 | 15.6 |
| $\cos^2\alpha$ | 5.6 | 13.8 |

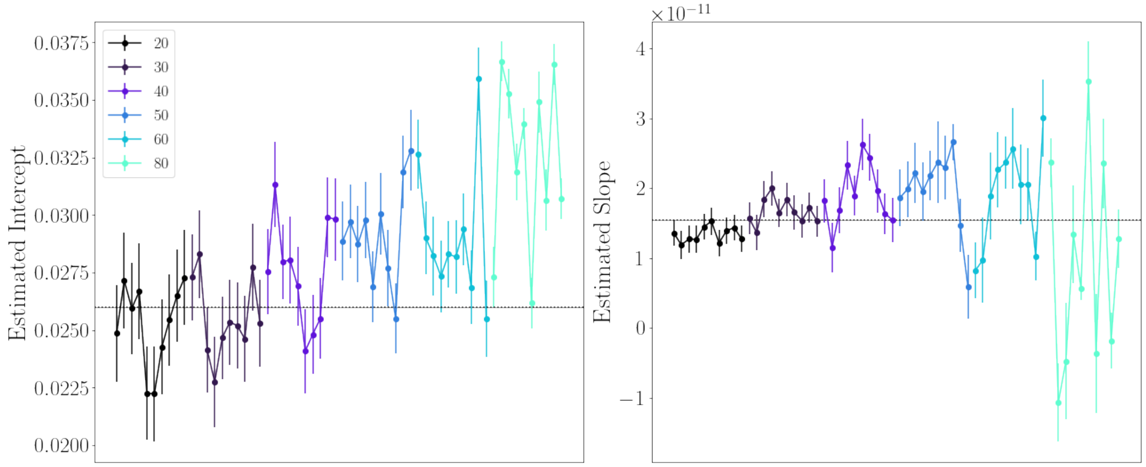

Figure A.6: A comparison of the accuracy of the estimated line parameters for 10 different realizations of simulated CMB observations at $0.5'$ resolution with different choices of patch sizes used to compute input statistics. The filter $f_\ell = \cos\alpha$ is used here, and horizontal lines indicate the slope and intercept predicted by Equation (A.21). Patch sizes indicated are the side length in number of pixels. We see that the patch size must be fine-tuned to mitigate potential biases in the estimated slope which contains the lensing signals.

Gaussian likelihood. There is less of a dramatic shift for the slope, but we see the importance of accounting for the pixel-to-pixel correlations in our model. We summarize the signal-to-noise of estimated slope and intercepts for each choice in directional filtering in Table A.1. We see that the filter $f_\ell = \cos\alpha$ performs best at distinguishing a non-zero lensing signal.

We finally shift our focus to some of the drawbacks of our real-space lensing estimator. These were mainly discussed at the end of §5.3 in Chapter 5. In summary, we argue that there is no single choice in patch size that can achieve both of the following:

- Contain a constant background temperature gradient within the patch, and
- Contain a sufficient number of pixels with which to estimate the small-scale temperature variance.

We show this in Figure A.6, where we compare the statistics of the estimated line parameters against the values predicted by Equation (A.21). We find that one must fine-tune the choice in patch size in order to mitigate potential biases in the estimated parameters from the model fit. Moreover, we find that these choices are specific to the map resolution. These are the obstacles that we overcome in the implementation of the Small Correlated Against Large Estimator in Chapter 5 (§5.4).

# Appendix B

# Derivation of SCALE

*Originally included in an Appendix of the manuscript accepted for publication in Phys. Rev. D. Pre-print available as: "The Small-Correlated-Against-Large-Estimator (SCALE) for the Lensing of the Cosmic Microwave Background" Chan, Victor C., Hložek, Renée, Meyers, Joel, & van Engelen, Alexander 2023, arXiv e-prints, arXiv:2302.13350, doi: 10.48550/arXiv.2302.13350. This section was written by Meyers, Joel and gives context to the analytical side of the methods presented in Chapters 5-6.*

We wish to construct an estimator of the small scale lensing power. We will construct the estimator using the cross correlation between a field constructed from the square of small-scale temperature fluctuations and a field constructed from the square of large-scale temperature fluctuations. In the very small-scale regime, where the effects of lensing dominate the temperature power spectrum, the temperature power is proportional to the product of the large-scale temperature gradient power and the small-scale lensing deflection power. The motivation of the estimator that we construct here is that the small-scale lensing-induced temperature power is non-Gaussian; the locally measured small-scale temperature power is correlated with variations in the large-scale temperature gradient, and the relation between the two is proportional to the small-scale lensing power.

Let us first define a field $\varsigma$ defined by the locally measured small-scale temperature gradient power

$$\varsigma(\check{\boldsymbol{L}}) = \int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{2\pi} g(\boldsymbol{\ell}_1, \check{\boldsymbol{L}}) \left(\boldsymbol{\ell}_1 \cdot (\boldsymbol{\ell}_1 - \check{\boldsymbol{L}})\right) T(\boldsymbol{\ell}_1) T(\check{\boldsymbol{L}} - \boldsymbol{\ell}_1)\,, \tag{B.1}$$

where $g$ is a filter applied to the small-scale temperature fluctuations, to be determined in what follows. We will expand the small-scale temperature fluctuations to first order in the lensing gradient

$$\tilde{T}(\boldsymbol{\ell}) = \int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{2\pi} \left(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \boldsymbol{\ell})\right) T(\boldsymbol{\ell}_2) \phi(\boldsymbol{\ell} - \boldsymbol{\ell}_2)\,, \tag{B.2}$$

where we have dropped the unlensed small-scale temperature, since it is assumed to be negligible compared to the lensing contribution on very small scales. Inserting this into $\varsigma$ gives

$$\begin{aligned}\varsigma(\check{\boldsymbol{L}}) = \int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{2\pi} g(\boldsymbol{\ell}_1, \check{\boldsymbol{L}}) \left(\boldsymbol{\ell}_1 \cdot (\boldsymbol{\ell}_1 - \check{\boldsymbol{L}})\right) \int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{2\pi} \int \frac{\mathrm{d}^2\boldsymbol{\ell}_3}{2\pi} \left(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \boldsymbol{\ell}_1)\right) \left(\boldsymbol{\ell}_3 \cdot (\boldsymbol{\ell}_3 - \check{\boldsymbol{L}} + \boldsymbol{\ell}_1)\right) \\ \times T(\boldsymbol{\ell}_2) T(\boldsymbol{\ell}_3) \phi(\boldsymbol{\ell}_1 - \boldsymbol{\ell}_2) \phi(\check{\boldsymbol{L}} - \boldsymbol{\ell}_1 - \boldsymbol{\ell}_3)\,. \qquad \text{(B.3)}\end{aligned}$$

We will be interested in an estimate of the small-scale lensing power, rather than the realization of the lensing potential, so we will take an average over lensing realizations

$$\begin{aligned}\varsigma(\check{\boldsymbol{L}}) = \int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{2\pi} g(\boldsymbol{\ell}_1, \check{\boldsymbol{L}}) \left(\boldsymbol{\ell}_1 \cdot (\boldsymbol{\ell}_1 - \check{\boldsymbol{L}})\right) \\ \times \int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2} \left(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \boldsymbol{\ell}_1)\right) \left((\check{\boldsymbol{L}} - \boldsymbol{\ell}_2) \cdot (\boldsymbol{\ell}_1 - \boldsymbol{\ell}_2)\right) T(\boldsymbol{\ell}_2) T(\check{\boldsymbol{L}} - \boldsymbol{\ell}_2) C^{\phi\phi}_{|\boldsymbol{\ell}_1 - \boldsymbol{\ell}_2|}\,. \qquad \text{(B.4)}\end{aligned}$$

The local large-scale temperature gradient power can be expressed in harmonic space as

$$\lambda(\check{\boldsymbol{L}}) = \int \frac{\mathrm{d}^2\boldsymbol{\ell}_3}{2\pi} h(\boldsymbol{\ell}_3, \check{\boldsymbol{L}}) \left(\boldsymbol{\ell}_3 \cdot (\boldsymbol{\ell}_3 - \check{\boldsymbol{L}})\right) T(\boldsymbol{\ell}_3) T(\check{\boldsymbol{L}} - \boldsymbol{\ell}_3)\,, \qquad \text{(B.5)}$$

where $h$ is a filter applied to the large-scale temperature fluctuations. On large angular scales, lensing is only a small correction, and so we work to zeroth order in the lensing potential for the temperature fluctuations appearing in $\lambda$.

Our aim is to isolate the small scale lensing power by analyzing correlations between the $\varsigma$ and $\lambda$ fields. The product of $\varsigma$ and $\lambda$ is

$$\begin{aligned}\varsigma(\check{\boldsymbol{L}})\lambda(\check{\boldsymbol{L}}') = \int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{2\pi} g(\boldsymbol{\ell}_1, \check{\boldsymbol{L}}) \left(\boldsymbol{\ell}_1 \cdot (\boldsymbol{\ell}_1 - \check{\boldsymbol{L}})\right) T(\boldsymbol{\ell}_1) T(\check{\boldsymbol{L}} - \boldsymbol{\ell}_1) \\ \times \int \frac{\mathrm{d}^2\boldsymbol{\ell}_3}{2\pi} h(\boldsymbol{\ell}_3, \check{\boldsymbol{L}}') \left(\boldsymbol{\ell}_3 \cdot (\boldsymbol{\ell}_3 - \check{\boldsymbol{L}}')\right) T(\boldsymbol{\ell}_3) T(\check{\boldsymbol{L}}' - \boldsymbol{\ell}_3)\,. \qquad \text{(B.6)}\end{aligned}$$

The cross-spectrum of $\varsigma$ and $\lambda$ is given for $\check{\boldsymbol{L}} \neq 0$ by

$$\begin{aligned}\left\langle \varsigma(\check{\boldsymbol{L}})\lambda(\check{\boldsymbol{L}}') \right\rangle = \int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{2\pi} \int \frac{\mathrm{d}^2\boldsymbol{\ell}_3}{2\pi} g(\boldsymbol{\ell}_1, \check{\boldsymbol{L}}) \left(\boldsymbol{\ell}_1 \cdot (\boldsymbol{\ell}_1 - \check{\boldsymbol{L}})\right) h(\boldsymbol{\ell}_3, \check{\boldsymbol{L}}') \left(\boldsymbol{\ell}_3 \cdot (\boldsymbol{\ell}_3 - \check{\boldsymbol{L}}')\right) \\ \times \left\langle T(\boldsymbol{\ell}_1) T(\check{\boldsymbol{L}} - \boldsymbol{\ell}_1) T(\boldsymbol{\ell}_3) T(\check{\boldsymbol{L}}' - \boldsymbol{\ell}_3) \right\rangle\,. \qquad \text{(B.7)}\end{aligned}$$

Choosing disjoint ranges of multipoles for the small-scale and large-scale temperature fluctuations means that the disconnected part of the temperature four-point function vanishes for $\check{\boldsymbol{L}} \neq 0$. The signal of interest is the one which is first order in the lensing power spectrum, whose dominant contribution comes from the terms of first order in the lensing potential in the gradient expansion of the small-scale temperature fluctuations

$$\begin{aligned}\left\langle \varsigma(\check{\boldsymbol{L}})\lambda(\check{\boldsymbol{L}}')\right\rangle =& \int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{2\pi} g(\boldsymbol{\ell}_1,\check{\boldsymbol{L}})\left(\boldsymbol{\ell}_1\cdot(\boldsymbol{\ell}_1-\check{\boldsymbol{L}})\right)\int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2}\left(\boldsymbol{\ell}_2\cdot(\boldsymbol{\ell}_2-\boldsymbol{\ell}_1)\right)\left((\check{\boldsymbol{L}}-\boldsymbol{\ell}_2)\cdot(\boldsymbol{\ell}_1-\boldsymbol{\ell}_2)\right)C^{\phi\phi}_{|\boldsymbol{\ell}_1-\boldsymbol{\ell}_2|}\\ &\times\int \frac{\mathrm{d}^2\boldsymbol{\ell}_3}{2\pi} h(\boldsymbol{\ell}_3,\check{\boldsymbol{L}})\left(\boldsymbol{\ell}_3\cdot(\boldsymbol{\ell}_3-\check{\boldsymbol{L}})\right)\left\langle T(\boldsymbol{\ell}_2)T(\check{\boldsymbol{L}}-\boldsymbol{\ell}_2)T(\boldsymbol{\ell}_3)T(\check{\boldsymbol{L}}-\boldsymbol{\ell}_3)\right\rangle\delta(\check{\boldsymbol{L}}+\check{\boldsymbol{L}}')\\ =& \int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} g(\boldsymbol{\ell}_1,\check{\boldsymbol{L}})\left(\boldsymbol{\ell}_1\cdot(\boldsymbol{\ell}_1-\check{\boldsymbol{L}})\right)\int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2}\left(\boldsymbol{\ell}_2\cdot(\boldsymbol{\ell}_2-\boldsymbol{\ell}_1)\right)\left((\check{\boldsymbol{L}}-\boldsymbol{\ell}_2)\cdot(\boldsymbol{\ell}_1-\boldsymbol{\ell}_2)\right)C^{\phi\phi}_{|\boldsymbol{\ell}_1-\boldsymbol{\ell}_2|}\\ &\times\left(h(-\boldsymbol{\ell}_2,-\check{\boldsymbol{L}})+h(\boldsymbol{\ell}_2-\check{\boldsymbol{L}},-\check{\boldsymbol{L}})\right)\left(\boldsymbol{\ell}_2\cdot(\boldsymbol{\ell}_2-\check{\boldsymbol{L}})\right)C^{TT}_{\ell_2}C^{TT}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|}\delta(\check{\boldsymbol{L}}+\check{\boldsymbol{L}}')\\ =& C^{\lambda\varsigma}_{\check{L}}\delta(\check{\boldsymbol{L}}+\check{\boldsymbol{L}}')\,. \end{aligned} \tag{B.8}$$

We wish to use this cross-spectrum to obtain an unbiased estimate of the integrated small-scale lensing power. We define a quantity

$$\Psi_{\check{L}} \equiv A_{\check{L}}\left\langle \varsigma(\check{\boldsymbol{L}})\lambda(-\check{\boldsymbol{L}})\right\rangle\,, \tag{B.9}$$

with $A_{\check{L}}$ defined such that $\Psi_{\check{L}}$ is a weighted-average of the small-scale lensing power

$$\begin{aligned}A_{\check{L}} =& \Bigg[\int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} g(\boldsymbol{\ell}_1,\check{\boldsymbol{L}})\left(\boldsymbol{\ell}_1\cdot(\boldsymbol{\ell}_1-\check{\boldsymbol{L}})\right)\int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2}\left(\boldsymbol{\ell}_2\cdot(\boldsymbol{\ell}_2-\boldsymbol{\ell}_1)\right)\left((\check{\boldsymbol{L}}-\boldsymbol{\ell}_2)\cdot(\boldsymbol{\ell}_1-\boldsymbol{\ell}_2)\right)\\ &\times\left(h(-\boldsymbol{\ell}_2,-\check{\boldsymbol{L}})+h(\boldsymbol{\ell}_2-\check{\boldsymbol{L}},-\check{\boldsymbol{L}})\right)\left(\boldsymbol{\ell}_2\cdot(\boldsymbol{\ell}_2-\check{\boldsymbol{L}})\right)C^{TT}_{\ell_2}C^{TT}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|}\Bigg]^{-1}\,. \end{aligned} \tag{B.10}$$

Let us briefly change to a full-sky notation which makes the calculation of the variance more transparent. We wish to estimate the small-scale lensing power from the cross-spectrum of the $\varsigma$ and $\lambda$ fields

$$A_{\check{L}}\left\langle \varsigma_{\check{L}\check{M}}\lambda_{\check{L}'-\check{M}'}\right\rangle \equiv \Psi_{\check{L}}\delta_{\check{L}\check{L}'}\delta_{\check{M}\check{M}'}\,. \tag{B.11}$$

An estimator for $\Psi_{\check{L}}$ can be constructed as

$$\hat{\Psi}_{\check{L}} \equiv A_{\check{L}}\frac{1}{2\check{L}+1}\sum_{\check{M}}\varsigma_{\check{L}\check{M}}\lambda_{\check{L}-\check{M}}\,, \tag{B.12}$$

such that in an isotropic universe

$$\left\langle \hat{\Psi}_{\check{L}}\right\rangle = A_{\check{L}}\frac{1}{2\check{L}+1}\sum_{\check{M}}\left\langle \varsigma_{\check{L}\check{M}}\lambda_{\check{L}-\check{M}}\right\rangle = \Psi_{\check{L}}\,. \tag{B.13}$$

The variance of this estimator can then be computed to be

$$\begin{aligned}\left\langle\left(\hat{\Psi}_{\check{L}}-\Psi_{\check{L}}\right)^2\right\rangle =& A_{\check{L}}^2\frac{1}{(2\check{L}+1)^2}\sum_{\check{M}\check{M}'}\left\langle\varsigma_{\check{L}\check{M}}\lambda_{\check{L}-\check{M}}\varsigma_{\check{L}\check{M}'}\lambda_{\check{L}-\check{M}'}\right\rangle-\Psi_{\check{L}}^2\\ =& A_{\check{L}}^2\frac{1}{(2\check{L}+1)^2}\Bigg[\left(\sum_{\check{M}}\left\langle\varsigma_{\check{L}\check{M}}\lambda_{\check{L}-\check{M}}\right\rangle\right)^2+\sum_{\check{M}}\left(\left\langle\varsigma_{\check{L}\check{M}}\lambda_{\check{L}-\check{M}}\right\rangle\right)^2\\ &+\sum_{\check{M}}\left\langle\varsigma_{\check{L}\check{M}}\varsigma_{\check{L}-\check{M}}\right\rangle\left\langle\lambda_{\check{L}\check{M}}\lambda_{\check{L}-\check{M}}\right\rangle\Bigg]-\Psi_{\check{L}}^2\\ =&\frac{1}{2\check{L}+1}\left[\Psi_{\check{L}}^2+N_{\check{L}}\right]\,,\end{aligned} \tag{B.14}$$

where we have defined

$$N_{\check{L}}\equiv A_{\check{L}}^2\left\langle\varsigma_{\check{L}\check{M}}\varsigma_{\check{L}-\check{M}}\right\rangle\left\langle\lambda_{\check{L}\check{M}}\lambda_{\check{L}-\check{M}}\right\rangle\,. \tag{B.15}$$

Next, we need to choose the filters $g$ and $h$ to minimize the variance of our small-scale lensing estimate. Returning to the flat-sky approximation, the noise variance can be expressed as

$$\begin{aligned}N_{\check{L}} =& A_{\check{L}}^2\int\frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2}g(\boldsymbol{\ell}_1,\check{\boldsymbol{L}})\left(g(-\boldsymbol{\ell}_1,-\check{\boldsymbol{L}})+g(\boldsymbol{\ell}_1-\check{\boldsymbol{L}},-\check{\boldsymbol{L}})\right)\left(\boldsymbol{\ell}_1\cdot(\boldsymbol{\ell}_1-\check{\boldsymbol{L}})\right)^2C_{\ell_1}^{TT,\mathrm{obs}}C_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_1|}^{TT,\mathrm{obs}}\\ &\times\int\frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2}h(\boldsymbol{\ell}_2,\check{\boldsymbol{L}})\left(h(-\boldsymbol{\ell}_2,-\check{\boldsymbol{L}})+h(\boldsymbol{\ell}_2-\check{\boldsymbol{L}},-\check{\boldsymbol{L}})\right)\left(\boldsymbol{\ell}_2\cdot(\boldsymbol{\ell}_2-\check{\boldsymbol{L}})\right)^2C_{\ell_2}^{TT,\mathrm{obs}}C_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|}^{TT,\mathrm{obs}}\,.\end{aligned} \tag{B.16}$$

Differentiating with respect to the choice of $g$ filter we find

$$\begin{aligned}\frac{\partial N_{\check{L}}}{\partial g(\boldsymbol{\ell}',\check{\boldsymbol{L}})} =& \frac{2A_{\check{L}}^2}{(2\pi)^2}\left(g(-\boldsymbol{\ell}',-\check{\boldsymbol{L}})+g(\boldsymbol{\ell}'-\check{\boldsymbol{L}},-\check{\boldsymbol{L}})\right)\left(\boldsymbol{\ell}'\cdot(\boldsymbol{\ell}'-\check{\boldsymbol{L}})\right)^2C_{\ell'}^{TT,\mathrm{obs}}C_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}'|}^{TT,\mathrm{obs}}\\ &\times\int\frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2}h(\boldsymbol{\ell}_2,\check{\boldsymbol{L}})\left(h(-\boldsymbol{\ell}_2,-\check{\boldsymbol{L}})+h(\boldsymbol{\ell}_2-\check{\boldsymbol{L}},-\check{\boldsymbol{L}})\right)\left(\boldsymbol{\ell}_2\cdot(\boldsymbol{\ell}_2-\check{\boldsymbol{L}})\right)^2C_{\ell_2}^{TT,\mathrm{obs}}C_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|}^{TT,\mathrm{obs}}\\ &-\frac{2A_{\check{L}}^3}{(2\pi^2)}\left(\boldsymbol{\ell}'\cdot(\boldsymbol{\ell}'-\check{\boldsymbol{L}})\right)\int\frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2}\left(\boldsymbol{\ell}_2\cdot(\boldsymbol{\ell}_2-\boldsymbol{\ell}')\right)\left((\check{\boldsymbol{L}}-\boldsymbol{\ell}_2)\cdot(\boldsymbol{\ell}'-\boldsymbol{\ell}_2)\right)\\ &\times\left(h(-\boldsymbol{\ell}_2,-\check{\boldsymbol{L}})+h(\boldsymbol{\ell}_2-\check{\boldsymbol{L}},-\check{\boldsymbol{L}})\right)\left(\boldsymbol{\ell}_2\cdot(\boldsymbol{\ell}_2-\check{\boldsymbol{L}})\right)C_{\ell_2}^{TT}C_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|}^{TT}\\ &\times\int\frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2}g(\boldsymbol{\ell}_1,\check{\boldsymbol{L}})\left(g(-\boldsymbol{\ell}_1,-\check{\boldsymbol{L}})+g(\boldsymbol{\ell}_1-\check{\boldsymbol{L}},-\check{\boldsymbol{L}})\right)\left(\boldsymbol{\ell}_1\cdot(\boldsymbol{\ell}_1-\check{\boldsymbol{L}})\right)^2C_{\ell_1}^{TT,\mathrm{obs}}C_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_1|}^{TT,\mathrm{obs}}\\ &\times\int\frac{\mathrm{d}^2\boldsymbol{\ell}_3}{(2\pi)^2}h(\boldsymbol{\ell}_3,\check{\boldsymbol{L}})\left(h(-\boldsymbol{\ell}_3,-\check{\boldsymbol{L}})+h(\boldsymbol{\ell}_3-\check{\boldsymbol{L}},-\check{\boldsymbol{L}})\right)\left(\boldsymbol{\ell}_3\cdot(\boldsymbol{\ell}_3-\check{\boldsymbol{L}})\right)^2C_{\ell_3}^{TT,\mathrm{obs}}C_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_3|}^{TT,\mathrm{obs}}\,.\end{aligned} \tag{B.17}$$

Setting this equal to zero and rearranging, we find

$$\begin{aligned}
0 = {} & \big(g(-\boldsymbol{\ell}', -\check{\boldsymbol{L}}) + g(\boldsymbol{\ell}' - \check{\boldsymbol{L}}, -\check{\boldsymbol{L}})\big) \big(\boldsymbol{\ell}' \cdot (\boldsymbol{\ell}' - \check{\boldsymbol{L}})\big)^2 C^{TT,\mathrm{obs}}_{\ell'} C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}'|} \\
& \times \int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} g(\boldsymbol{\ell}_1, \check{\boldsymbol{L}}) \big(\boldsymbol{\ell}_1 \cdot (\boldsymbol{\ell}_1 - \check{\boldsymbol{L}})\big) \int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2} \left(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \boldsymbol{\ell}_1)\right) \big((\check{\boldsymbol{L}} - \boldsymbol{\ell}_2) \cdot (\boldsymbol{\ell}_1 - \boldsymbol{\ell}_2)\big) \\
& \times \big(h(-\boldsymbol{\ell}_2, -\check{\boldsymbol{L}}) + h(\boldsymbol{\ell}_2 - \check{\boldsymbol{L}}, -\check{\boldsymbol{L}})\big) \big(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \check{\boldsymbol{L}})\big) C^{TT}_{\ell_2} C^{TT}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|} \\
& - \big(\boldsymbol{\ell}' \cdot (\boldsymbol{\ell}' - \check{\boldsymbol{L}})\big) \int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2} \left(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \boldsymbol{\ell}')\right) \big((\check{\boldsymbol{L}} - \boldsymbol{\ell}_2) \cdot (\boldsymbol{\ell}' - \boldsymbol{\ell}_2)\big) \\
& \times \big(h(-\boldsymbol{\ell}_2, -\check{\boldsymbol{L}}) + h(\boldsymbol{\ell}_2 - \check{\boldsymbol{L}}, -\check{\boldsymbol{L}})\big) \big(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \check{\boldsymbol{L}})\big) C^{TT}_{\ell_2} C^{TT}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|} \\
& \times \int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} g(\boldsymbol{\ell}_1, \check{\boldsymbol{L}}) \big(g(-\boldsymbol{\ell}_1, -\check{\boldsymbol{L}}) + g(\boldsymbol{\ell}_1 - \check{\boldsymbol{L}}, -\check{\boldsymbol{L}})\big) \big(\boldsymbol{\ell}_1 \cdot (\boldsymbol{\ell}_1 - \check{\boldsymbol{L}})\big)^2 C^{TT,\mathrm{obs}}_{\ell_1} C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_1|}
\end{aligned} \tag{B.18}$$

A similar procedure for the derivative with respect to the $h$ filter gives

$$\begin{aligned}
0 = {} & \big(h(-\boldsymbol{\ell}', -\check{\boldsymbol{L}}) + h(\boldsymbol{\ell}' - \check{\boldsymbol{L}}, -\check{\boldsymbol{L}})\big) \big(\boldsymbol{\ell}' \cdot (\boldsymbol{\ell}' - \check{\boldsymbol{L}})\big)^2 C^{TT,\mathrm{obs}}_{\ell'} C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}'|} \\
& \times \int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} g(\boldsymbol{\ell}_1, \check{\boldsymbol{L}}) \big(\boldsymbol{\ell}_1 \cdot (\boldsymbol{\ell}_1 - \check{\boldsymbol{L}})\big) \int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2} \left(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \boldsymbol{\ell}_1)\right) \big((\check{\boldsymbol{L}} - \boldsymbol{\ell}_2) \cdot (\boldsymbol{\ell}_1 - \boldsymbol{\ell}_2)\big) \\
& \times \big(h(-\boldsymbol{\ell}_2, -\check{\boldsymbol{L}}) + h(\boldsymbol{\ell}_2 - \check{\boldsymbol{L}}, -\check{\boldsymbol{L}})\big) \big(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \check{\boldsymbol{L}})\big) C^{TT}_{\ell_2} C^{TT}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|} \\
& - \big(\boldsymbol{\ell}' \cdot (\boldsymbol{\ell}' - \check{\boldsymbol{L}})\big) C^{TT}_{\ell'} C^{TT}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}'|} \\
& \times \int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} g(\boldsymbol{\ell}_1, \check{\boldsymbol{L}}) \big(\boldsymbol{\ell}_1 \cdot (\boldsymbol{\ell}_1 - \check{\boldsymbol{L}})\big) \left(\boldsymbol{\ell}' \cdot (\boldsymbol{\ell}' - \boldsymbol{\ell}_1)\right) \big((\check{\boldsymbol{L}} - \boldsymbol{\ell}') \cdot (\boldsymbol{\ell}_1 - \boldsymbol{\ell}')\big) \\
& \times \int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2} h(\boldsymbol{\ell}_2, \check{\boldsymbol{L}}) \big(h(-\boldsymbol{\ell}_2, -\check{\boldsymbol{L}}) + h(\boldsymbol{\ell}_2 - \check{\boldsymbol{L}}, -\check{\boldsymbol{L}})\big) \big(\boldsymbol{\ell}_2 \cdot (\boldsymbol{\ell}_2 - \check{\boldsymbol{L}})\big)^2 C^{TT,\mathrm{obs}}_{\ell_2} C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|}
\end{aligned} \tag{B.19}$$

These equations are difficult to solve in general, but if we restrict attention to cases where $\varsigma$ includes only temperature fluctuations on scales much smaller than fluctuations appearing in $\lambda$ and also much smaller than scales defined by $\check{L}$, then one can see that an approximate solution is provided by

$$g(\boldsymbol{\ell}, \check{\boldsymbol{L}}) = W_\varsigma(\boldsymbol{\ell}) \frac{1}{C^{TT,\mathrm{obs}}_{\ell}} W_\varsigma(\check{\boldsymbol{L}} - \boldsymbol{\ell}) \frac{1}{C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}|}}\,, \tag{B.20}$$

$$h(\boldsymbol{\ell}, \check{\boldsymbol{L}}) = W_\lambda(\boldsymbol{\ell}) \frac{C^{TT}_{\ell}}{C^{TT,\mathrm{obs}}_{\ell}} W_\lambda(\check{\boldsymbol{L}} - \boldsymbol{\ell}) \frac{C^{TT}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}|}}{C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}|}}\,, \tag{B.21}$$

where $W_\varsigma$ and $W_\lambda$ are window functions that restrict the temperature fluctuations to the appropriate scales,

$$W_\varsigma(\boldsymbol{\ell}) = \begin{cases} 1, & \ell_{1,\mathrm{min}} \le |\boldsymbol{\ell}| < \ell_{1,\mathrm{max}} \\ 0, & \text{else}\,, \end{cases} \tag{B.22}$$

$$W_\lambda(\boldsymbol{\ell}) = \begin{cases} 1, & \ell_{2,\mathrm{min}} \le |\boldsymbol{\ell}| < \ell_{2,\mathrm{max}} \\ 0, & \text{else}\,. \end{cases} \tag{B.23}$$

Using this choice of filters gives

$$\begin{aligned}A_{\check{L}} = \Bigg[ & 2\int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} W_\varsigma(\boldsymbol{\ell}_1) W_\varsigma(\check{\boldsymbol{L}}-\boldsymbol{\ell}_1)\left(\boldsymbol{\ell}_1\cdot(\boldsymbol{\ell}_1-\check{\boldsymbol{L}})\right)\frac{1}{C^{TT,\mathrm{obs}}_{\ell_1}}\frac{1}{C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_1|}} \\ & \times\int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2} W_\lambda(\boldsymbol{\ell}_2) W_\lambda(\check{\boldsymbol{L}}-\boldsymbol{\ell}_2)\left(\boldsymbol{\ell}_2\cdot(\boldsymbol{\ell}_2-\boldsymbol{\ell}_1)\right)\left((\check{\boldsymbol{L}}-\boldsymbol{\ell}_2)\cdot(\boldsymbol{\ell}_1-\boldsymbol{\ell}_2)\right)\left(\boldsymbol{\ell}_2\cdot(\boldsymbol{\ell}_2-\check{\boldsymbol{L}})\right) \\ & \times\frac{\left(C^{TT}_{\ell_2}\right)^2}{C^{TT,\mathrm{obs}}_{\ell_2}}\frac{\left(C^{TT}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|}\right)^2}{C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|}}\Bigg]^{-1} ,\end{aligned} \tag{B.24}$$

and the expected value of $\hat{\Psi}_{\check{L}}$ is

$$\begin{aligned}\left\langle\hat{\Psi}_{\check{L}}\right\rangle = & 2A_{\check{L}}\int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} W_\varsigma(\boldsymbol{\ell}_1) W_\varsigma(\check{\boldsymbol{L}}-\boldsymbol{\ell}_1)\left(\boldsymbol{\ell}_1\cdot(\boldsymbol{\ell}_1-\check{\boldsymbol{L}})\right)\frac{1}{C^{TT,\mathrm{obs}}_{\ell_1}}\frac{1}{C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_1|}} \\ & \times\int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2} W_\lambda(\boldsymbol{\ell}_2) W_\lambda(\check{\boldsymbol{L}}-\boldsymbol{\ell}_2)\left(\boldsymbol{\ell}_2\cdot(\boldsymbol{\ell}_2-\boldsymbol{\ell}_1)\right)\left((\check{\boldsymbol{L}}-\boldsymbol{\ell}_2)\cdot(\boldsymbol{\ell}_1-\boldsymbol{\ell}_2)\right)\left(\boldsymbol{\ell}_2\cdot(\boldsymbol{\ell}_2-\check{\boldsymbol{L}})\right) \\ & \times\frac{\left(C^{TT}_{\ell_2}\right)^2}{C^{TT,\mathrm{obs}}_{\ell_2}}\frac{\left(C^{TT}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|}\right)^2}{C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|}}C^{\phi\phi}_{|\boldsymbol{\ell}_1-\boldsymbol{\ell}_2|} ,\end{aligned} \tag{B.25}$$

The noise variance of $\hat{\Psi}_{\check{L}}$ is

$$\begin{aligned}N_{\check{L}} = & 4A^2_{\check{L}}\int \frac{\mathrm{d}^2\boldsymbol{\ell}_1}{(2\pi)^2} W_\varsigma(\boldsymbol{\ell}_1) W_\varsigma(\check{\boldsymbol{L}}-\boldsymbol{\ell}_1)\left(\boldsymbol{\ell}_1\cdot(\boldsymbol{\ell}_1-\check{\boldsymbol{L}})\right)^2\frac{1}{C^{TT,\mathrm{obs}}_{\ell_1}}\frac{1}{C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_1|}} \\ & \times\int \frac{\mathrm{d}^2\boldsymbol{\ell}_2}{(2\pi)^2} W_\lambda(\boldsymbol{\ell}_2) W_\lambda(\check{\boldsymbol{L}}-\boldsymbol{\ell}_2)\left(\boldsymbol{\ell}_2\cdot(\boldsymbol{\ell}_2-\check{\boldsymbol{L}})\right)^2\frac{\left(C^{TT}_{\ell_2}\right)^2}{C^{TT,\mathrm{obs}}_{\ell_2}}\frac{\left(C^{TT}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|}\right)^2}{C^{TT,\mathrm{obs}}_{|\check{\boldsymbol{L}}-\boldsymbol{\ell}_2|}} \\ \simeq & 4A_{\check{L}} ,\end{aligned} \tag{B.26}$$

where in the last line, we used the same approximations as in deriving the $g$ and $h$ filters.

# List of Acronyms

**2MASS** Two Micron All Sky Survey

**ACT** Atacama Cosmology Telescope
**APOGEE** Apache Point Observatory Galactic Evolution Experiment

**BAO** Baryon Acoustic Oscillations
**BOOMERanG** Balloon Observations Of Millimetric Extragalactic Radiation and Geophysics
**BOSS** Baryon Oscillation Spectroscopic Survey

**CAMB** Code for Anisotropies in the Microwave Background
**CCD** Charge-Coupled Device
**CDM** Cold Dark Matter
**CHIME** Canadian Hydrogen Intensity Mapping Experiment
**CIB** Cosmic Infrared Background
**CLASS** Cosmic Linear Anisotropy Solving System, *or*
**CLASS** Cosmology Large Angular Scale Surveyor
**CMB** Cosmic Microwave Background
**CMB-HD** The Cosmic Microwave Background in High Definition
**CMB-S4** Cosmic Microwave Background-Stage IV
**CSA** Canadian Space Agency
**COBE** Cosmic Background Explorer

**DES** Dark Energy Survey
**DESI** Dark Energy Spectroscopic Instrument
**DM** Dark Matter or Dispersion Measure
**DR** Data Release
**DR2** (*Gaia*) Data Release 2
**DR3** (*Gaia*) Data Release 3
**DMR** Differential Microwave Radiometer
**DUCC** Distinctly Useful Code Collection

**eBOSS** Extended Baryon Oscillation Spectroscopic Survey
**EDR3** (*Gaia*) Early Data Release 3

**fDM** fuzzy Dark Matter
**FIRAS** Far-InfraRed Absolute Spectrophotometer
**FLRW** Friedmann-Lemaître-Robertson-Walker
**fMUX** Frequency Multiplexing
**FOV** Field of View
**FRB** Fast Radio Burst
**FWHM** Full Width at Half Maximum

**GI** Gradient Inversion
**GW** Gravitational Wave

**HDM** Hot Dark Matter
**HDV** Hu, DeDeo, and Vale (Quadratic Estimator)
**HuOk** Hu and Okamoto (Quadratic Estimator)
**HEALPix** Hierarchical Equal Area isoLatitude Pixelization
**HFT** High Frequency Telescope
**HMC** Hamiltonian Monte Carlo

**IGM** Intergalactic Medium

**JAXA** Japan Aerospace Exploration Agency
**JWST** James Webb Space Telescope

**kSZ** Kinetic Sunyaev-Zel'dovich

**ΛCDM** Lambda Cold Dark Matter
**LFT** Low Frequency Telescope
**LiteBIRD** Lite satellite for the studies of B-mode polarization and Inflation from cosmic background Radiation Detection
**LMC** Large Magellanic Cloud
**LSS** Large Scale Structure
**LSST** Large Synoptic Survey Telescope (now Vera C. Rubin Observatory)

**MC** Monte Carlo
**MCMC** Markov chain Monte Carlo
**MFT** Medium Frequency Telescope
**MKIDs** Microwave Kinetic Inductance Detectors
**MW** Milky Way

**NET** Noise Equivalent Temperature
**NN** Neural Network

**PL** Period–Luminosity (relation)

**PSD** Power Spectral Density

**QE** Quadratic Estimator

**RC** Red Clump (star)
**RJCE** Rayleigh-Jeans Colour Excess

**SCALE** Small Correlated Against Large Estimator
**SDSS** Sloan Digital Sky Survey
**SH0ES** Supernova $H_0$ for the Equation of State
**SNIa** Type Ia Supernova
**SNR** Signal-to-Noise Ratio
**SO** Simons Observatory
**SPT** South Pole Telescope
**SZ** Sunyaev-Zel'dovich

**TES** Transition Edge Sensor
**TESS** Transiting Exoplanet Survey Satellite
**tSZ** Thermal Sunyaev-Zel'dovich
**TOAST** Time Ordered Astrophysics Scalable Tools
**TOD** Time Ordered Data
**TRGB** Tip of the Red Giant Branch

**VLBI** Very Long Baseline Interferometry
**VRO** Vera C. Rubin Observatory (formerly Large Synoptic Survey Telescope)

**WDM** Warm Dark Matter
**WMAP** Wilkinson Microwave Anisotropy Probe

# Bibliography

2009, Bayesian Methods in Cosmology (Cambridge University Press), doi: 10.1017/CBO9780511802461

Abadi, M., Agarwal, A., Barham, P., et al. 2015, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. https://www.tensorflow.org/

Abazajian, K., et al. 2019. https://arxiv.org/abs/1907.04473

Abazajian, K. N., et al. 2016. https://arxiv.org/abs/1610.02743

—. 2022. https://arxiv.org/abs/2203.07377

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, Nature, 551, 85, doi: 10.1038/nature24471

Abbott, T. M. C., Abdalla, F. B., Annis, J., et al. 2018, MNRAS, 480, 3879, doi: 10.1093/mnras/sty1939

Abolfathi, B., Aguado, D. S., Aguilar, G., et al. 2018, ApJS, 235, 42, doi: 10.3847/1538-4365/aa9e8a

Abylkairov, Y. S., Darwish, O., Hill, J. C., & Sherwin, B. D. 2021, Phys. Rev. D, 103, 103510, doi: 10.1103/PhysRevD.103.103510

Ade, P., et al. 2019, JCAP, 02, 056, doi: 10.1088/1475-7516/2019/02/056

Ade, P. A. R., Ahmed, Z., Amiri, M., et al. 2022, ApJ, 927, 77, doi: 10.3847/1538-4357/ac4886

Aghanim, N., et al. 2020, Astron. Astrophys., 641, A8, doi: 10.1051/0004-6361/201833886

Aiola, S., et al. 2020, JCAP, 12, 047, doi: 10.1088/1475-7516/2020/12/047

al-Sufi, A. a.-R. 964, The Book of Fixed Stars

Alam, S., Aubert, M., Avila, S., et al. 2021, Phys. Rev. D, 103, 083533, doi: 10.1103/PhysRevD.103.083533

Allison, R., Caucal, P., Calabrese, E., Dunkley, J., & Louis, T. 2015, Phys. Rev. D, 92, 123535, doi: 10.1103/PhysRevD.92.123535

Alvarez, M. A., Ferraro, S., Hill, J. C., Hložek, R., & Ikape, M. 2021, Phys. Rev. D, 103, 063518, doi: 10.1103/PhysRevD.103.063518

Alves, D. R. 2004, New A Rev., 48, 659, doi: 10.1016/j.newar.2004.03.001

Alzubaidi, L., Zhang, J., Humaidi, A. J., et al. 2021, Journal of Big Data, 8. https://api.semanticscholar.org/CorpusID:232434552

Ange, J., & Meyers, J. 2023, arXiv e-prints, arXiv:2307.01662, doi: 10.48550/arXiv.2307.01662

Antoja, T., Helmi, A., Romero-Gómez, M., et al. 2018, Nature, 561, 360, doi: 10.1038/s41586-018-0510-7

Arbey, A., & Mahmoudi, F. 2021, Progress in Particle and Nuclear Physics, 119, 103865, doi: 10.1016/j.ppnp.2021.103865

Arenou, F., Luri, X., Babusiaux, C., et al. 2018, A&A, 616, A17, doi: 10.1051/0004-6361/201833234

Armengaud, E., Palanque-Delabrouille, N., Yèche, C., Marsh, D. J. E., & Baur, J. 2017, MNRAS, 471, 4606, doi: 10.1093/mnras/stx1870

Aubourg, É., Bailey, S., Bautista, J. E., et al. 2015, Phys. Rev. D, 92, 123516, doi: 10.1103/PhysRevD.92.123516

Bailer-Jones, C. A. L. 2015, Publications of the Astronomical Society of the Pacific, 127, 994, doi: 10.1086/683116

Baselmans, J. J. A., Facchin, F., Pascual Laguna, A., et al. 2022, A&A, 665, A17, doi: 10.1051/0004-6361/202243840

Bassett, B., & Hlozek, R. 2010, in Dark Energy: Observational and Theoretical Approaches, ed. P. Ruiz-Lapuente, 246, doi: 10.48550/arXiv.0910.5224

Baumann, D. 2022, Cosmology (Cambridge University Press), doi: 10.1017/9781108937092

Baumann, D., & McAllister, L. 2015, Inflation and String Theory, doi: 10.48550/arXiv.1404.2601

Baxter, E. J., Raghunathan, S., Crawford, T. M., et al. 2018, MNRAS, 476, 2674, doi: 10.1093/mnras/sty305

Baxter, E. J., et al. 2022. https://arxiv.org/abs/2203.06795

Beck, D., Errard, J., & Stompor, R. 2020, J. Cosmology Astropart. Phys., 2020, 030, doi: 10.1088/1475-7516/2020/06/030

Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, PASP, 131, 018002, doi: 10.1088/1538-3873/aaecbe

Bennett, C. L., Banday, A. J., Gorski, K. M., et al. 1996, ApJ, 464, L1, doi: 10.1086/310075

Bennett, C. L., Hill, R. S., Hinshaw, G., et al. 2011, ApJS, 192, 17, doi: 10.1088/0067-0049/192/2/17

Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, ApJS, 208, 20, doi: 10.1088/0067-0049/208/2/20

Bennett, M., & Bovy, J. 2019, MNRAS, 482, 1417, doi: 10.1093/mnras/sty2813

Bernal, J. L., Verde, L., & Riess, A. G. 2016, J. Cosmology Astropart. Phys., 2016, 019, doi: 10.1088/1475-7516/2016/10/019

Bianchini, F., Bielewicz, P., Lapi, A., et al. 2015, ApJ, 802, 64, doi: 10.1088/0004-637X/802/1/64

Bianco, F. B., Ivezić, Ž., Jones, R. L., et al. 2022, ApJS, 258, 1, doi: 10.3847/1538-4365/ac3e72

BICEP2 Collaboration, Ade, P. A. R., Aikin, R. W., et al. 2014, Phys. Rev. Lett., 112, 241101, doi: 10.1103/PhysRevLett.112.241101

Binney, J., & Schönrich, R. 2018, MNRAS, 481, 1501, doi: 10.1093/mnras/sty2378

Birrer, S., Millon, M., Sluse, D., et al. 2022, arXiv e-prints, arXiv:2210.10833, doi: 10.48550/arXiv.2210.10833

Birrer, S., Shajib, A. J., Galan, A., et al. 2020, A&A, 643, A165, doi: 10.1051/0004-6361/202038861

Blanton, M. R., Bershady, M. A., Abolfathi, B., et al. 2017, AJ, 154, 28, doi: 10.3847/1538-3881/aa7567

Blas, D., Lesgourgues, J., & Tram, T. 2011, J. Cosmology Astropart. Phys., 2011, 034, doi: 10.1088/1475-7516/2011/07/034

Bleem, L. E., Crawford, T. M., Ansarinejad, B., et al. 2022, ApJS, 258, 36, doi: 10.3847/1538-4365/ac35e9

Bonaldi, A., Ricciardi, S., & Brown, M. L. 2014, MNRAS, 444, 1034, doi: 10.1093/mnras/stu1495

Bovy, J., Rix, H.-W., Green, G. M., Schlafly, E. F., & Finkbeiner, D. P. 2016, ApJ, 818, 130, doi: 10.3847/0004-637X/818/2/130

Bovy, J., Nidever, D. L., Rix, H.-W., et al. 2014, ApJ, 790, 127, doi: 10.1088/0004-637X/790/2/127

Boyarsky, A., Lesgourgues, J., Ruchayskiy, O., & Viel, M. 2009, J. Cosmology Astropart. Phys., 2009, 012, doi: 10.1088/1475-7516/2009/05/012

Bressan, A., Marigo, P., Girardi, L., et al. 2012, MNRAS, 427, 127, doi: 10.1111/j.1365-2966.2012.21948.x

Buen-Abad, M. A., Essig, R., McKeen, D., & Zhong, Y.-M. 2022, Phys. Rept., 961, 1, doi: 10.1016/j.physrep.2022.02.006

Butkevich, A. G., Klioner, S. A., Lindegren, L., Hobbs, D., & van Leeuwen, F. 2017, A&A, 603, A45, doi: 10.1051/0004-6361/201730781

Camarena, D., & Marra, V. 2020, MNRAS, 495, 2630, doi: 10.1093/mnras/staa770

Carpenter, B., Gelman, A., Hoffman, M. D., et al. 2017, Journal of Statistical Software, 76, 1

Carrasco, J. M., Weiler, M., Jordi, C., et al. 2021, A&A, 652, A86, doi: 10.1051/0004-6361/202141249

Carron, J., & Lewis, A. 2017, Phys. Rev. D, 96, 063510, doi: 10.1103/PhysRevD.96.063510

Cassanelli, T., Leung, C., Rahman, M., et al. 2022, AJ, 163, 65, doi: 10.3847/1538-3881/ac3d2f

Castelli, F., & Kurucz, R. L. 2003, in IAU Symposium, Vol. 210, Modelling of Stellar Atmospheres, ed. N. Piskunov, W. W. Weiss, & D. F. Gray, A20. https://arxiv.org/abs/astro-ph/0405087

Chabanier, S., Millea, M., & Palanque-Delabrouille, N. 2019, MNRAS, 489, 2247, doi: 10.1093/mnras/stz2310

Chang, C., Omori, Y., Baxter, E. J., et al. 2022, arXiv e-prints, arXiv:2203.12440, doi: 10.48550/arXiv.2203.12440

Charnock, T., Battye, R. A., & Moss, A. 2017, Phys. Rev. D, 95, 123535, doi: 10.1103/PhysRevD.95.123535

Chen, G., Gott, J. Richard, I., & Ratra, B. 2003, PASP, 115, 1269, doi: 10.1086/379219

Chen, G., & Ratra, B. 2011, PASP, 123, 1127, doi: 10.1086/662131

Chen, H.-Y., Fishbach, M., & Holz, D. E. 2018, Nature, 562, 545, doi: 10.1038/s41586-018-0606-0

CHIME Collaboration, Amiri, M., Bandura, K., et al. 2022, ApJS, 261, 29, doi: 10.3847/1538-4365/ac6fd9

Chisari, N. E., et al. 2019, Open J. Astrophys., 2, 4, doi: 10.21105/astro.1905.06082

Chown, R., Omori, Y., Aylor, K., et al. 2018, ApJS, 239, 10, doi: 10.3847/1538-4365/aae694

Cohen, M., Wheaton, W. A., & Megeath, S. T. 2003, AJ, 126, 1090, doi: 10.1086/376474

Cordes, J. M., & Chatterjee, S. 2019, ARA&A, 57, 417, doi: 10.1146/annurev-astro-091918-104501

Correa Magnus, L., & Vasiliev, E. 2022, MNRAS, 511, 2610, doi: 10.1093/mnras/stab3726

Croft, R. A. C., Weinberg, D. H., Katz, N., & Hernquist, L. 1998, ApJ, 495, 44, doi: 10.1086/305289

Danielski, C., Babusiaux, C., Ruiz-Dern, L., Sartoretti, P., & Arenou, F. 2018, A&A, 614, A19, doi: 10.1051/0004-6361/201732327

Darwish, O., Sherwin, B. D., Sailer, N., Schaan, E., & Ferraro, S. 2021a, arXiv e-prints, arXiv:2111.00462, doi: 10.48550/arXiv.2111.00462

Darwish, O., Madhavacheril, M. S., Sherwin, B. D., et al. 2021b, MNRAS, 500, 2250, doi: 10.1093/mnras/staa3438

de Haan, T., Suzuki, A., Boyd, S. T. P., et al. 2020, Journal of Low Temperature Physics, 199, 754, doi: 10.1007/s10909-020-02403-8

Dentler, M., Marsh, D. J. E., Hložek, R., et al. 2022, MNRAS, 515, 5646, doi: 10.1093/mnras/stac1946

Di Valentino, E., Melchiorri, A., & Silk, J. 2015, Phys. Rev. D, 92, 121302, doi: 10.1103/PhysRevD.92.121302

—. 2016, Physics Letters B, 761, 242, doi: 10.1016/j.physletb.2016.08.043

Di Valentino, E., Mena, O., Pan, S., et al. 2021, Classical and Quantum Gravity, 38, 153001, doi: 10.1088/1361-6382/ac086d

Dicke, R. H., Peebles, P. J. E., Roll, P. G., & Wilkinson, D. T. 1965, ApJ, 142, 414, doi: 10.1086/148306

Diego-Palazuelos, P., Eskilt, J. R., Minami, Y., et al. 2022, Phys. Rev. Lett., 128, 091302, doi: 10.1103/PhysRevLett.128.091302

Dobbs, M., Halpern, M., Irwin, K. D., et al. 2009, in Journal of Physics Conference Series, Vol. 155, Journal of Physics Conference Series, 012004, doi: 10.1088/1742-6596/155/1/012004

Dodelson, S. 2003, Modern Cosmology

Dodelson, S., & Schmidt, F. 2020, Modern Cosmology, doi: 10.1016/C2017-0-01943-2

Dong, F., Yu, Y., Zhang, J., Yang, X., & Zhang, P. 2021, MNRAS, 500, 3838, doi: 10.1093/mnras/staa3194

Drewes, M., et al. 2017, JCAP, 01, 025, doi: 10.1088/1475-7516/2017/01/025

Dunkley, J., Komatsu, E., Nolta, M. R., et al. 2009, ApJS, 180, 306, doi: 10.1088/0067-0049/180/2/306

Dunkley, J., Hlozek, R., Sievers, J., et al. 2011, ApJ, 739, 52, doi: 10.1088/0004-637X/739/1/52

Dvorkin, C., & Smith, K. M. 2009, Phys. Rev. D, 79, 043003, doi: 10.1103/PhysRevD.79.043003

Efstathiou, G., Bond, J. R., & White, S. D. M. 1992, MNRAS, 258, 1P, doi: 10.1093/mnras/258.1.1P

Eimer, J. R., Li, Y., Brewer, M. K., et al. 2023, arXiv e-prints, arXiv:2309.00675, doi: 10.48550/arXiv.2309.00675

Einstein, A. 1916, Annalen der Physik, 354, 769, doi: 10.1002/andp.19163540702

—. 1917, Sitzungsberichte der K&ouml;niglich Preussischen Akademie der Wissenschaften, 142

Eisenstein, D. 2020, Animation of the Formation of the Acoustic Peak. https://lweb.cfa.harvard.edu/~deisenst/acousticpeak/acoustic_anim.html

Eisenstein, D. J., Weinberg, D. H., Agol, E., et al. 2011, AJ, 142, 72, doi: 10.1088/0004-6256/142/3/72

Erkal, D., Belokurov, V., Laporte, C. F. P., et al. 2019, MNRAS, 487, 2685, doi: 10.1093/mnras/stz1371

Errard, J., Ade, P. A. R., Akiba, Y., et al. 2015, ApJ, 809, 63, doi: 10.1088/0004-637X/809/1/63

Eskilt, J. R., & Komatsu, E. 2022, Phys. Rev. D, 106, 063503, doi: 10.1103/PhysRevD.106.063503

Essinger-Hileman, T., Ali, A., Amiri, M., et al. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9153, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII, ed. W. S. Holland & J. Zmuidzinas, 91531I, doi: 10.1117/12.2056701

Fausnaugh, M. M., Vallely, P. J., Kochanek, C. S., et al. 2021, ApJ, 908, 51, doi: 10.3847/1538-4357/abcd42

Ferraro, S., Sherwin, B. D., & Spergel, D. N. 2015, Phys. Rev. D, 91, 083533, doi: 10.1103/PhysRevD.91.083533

Ferreira, E. G. M. 2021, A&A Rev., 29, 7, doi: 10.1007/s00159-021-00135-6

Fitzpatrick, E. L. 1999, PASP, 111, 63, doi: 10.1086/316293

Fitzpatrick, E. L., Massa, D., Gordon, K. D., Bohlin, R., & Clayton, G. C. 2019, ApJ, 886, 108, doi: 10.3847/1538-4357/ab4c3a

Flauger, R., Hill, J. C., & Spergel, D. N. 2014, J. Cosmology Astropart. Phys., 2014, 039, doi: 10.1088/1475-7516/2014/08/039

Font-Ribera, A., McDonald, P., Mostek, N., et al. 2014, J. Cosmology Astropart. Phys., 2014, 023, doi: 10.1088/1475-7516/2014/05/023

Foreman, S., Hotinli, S. C., Madhavacheril, M. S., van Engelen, A., & Kreisch, C. D. 2023, Phys. Rev. D, 107, 083502, doi: 10.1103/PhysRevD.107.083502

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306, doi: 10.1086/670067

Fosalba, P., Gaztañaga, E., & Castander, F. J. 2003, ApJ, 597, L89, doi: 10.1086/379848

Fowler, J. W., Acquaviva, V., Ade, P. A. R., et al. 2010, ApJ, 722, 1148, doi: 10.1088/0004-637X/722/2/1148

Frankel, N., Bovy, J., Tremaine, S., & Hogg, D. W. 2023, MNRAS, 521, 5917, doi: 10.1093/mnras/stad908

Freedman, W. L. 2021, ApJ, 919, 16, doi: 10.3847/1538-4357/ac0e95

Freedman, W. L., Madore, B. F., Hatt, D., et al. 2019, ApJ, 882, 34, doi: 10.3847/1538-4357/ab2f73

Friedmann, A. 1922, Zeitschrift fur Physik, 10, 377, doi: 10.1007/BF01332580

Fritz, T. K., Battaglia, G., Pawlowski, M. S., et al. 2018, A&A, 619, A103, doi: 10.1051/0004-6361/201833343

Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, A&A, 595, A1, doi: 10.1051/0004-6361/201629272

Gaia Collaboration, Helmi, A., van Leeuwen, F., et al. 2018a, A&A, 616, A12, doi: 10.1051/0004-6361/201832698

Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018b, A&A, 616, A1, doi: 10.1051/0004-6361/201833051

Gaia Collaboration, Luri, X., Chemin, L., et al. 2021a, A&A, 649, A7, doi: 10.1051/0004-6361/202039588

Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2021b, A&A, 649, A1, doi: 10.1051/0004-6361/202039657

Gaia Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023, A&A, 674, A1, doi: 10.1051/0004-6361/202243940

Garavito-Camargo, N., Besla, G., Laporte, C. F. P., et al. 2021, ApJ, 919, 109, doi: 10.3847/1538-4357/ac0b44

García Pérez, A. E., Allende Prieto, C., Holtzman, J. A., et al. 2016, AJ, 151, 144, doi: 10.3847/0004-6256/151/6/144

Gardner, J. P., Mather, J. C., Clampin, M., et al. 2006, Space Sci. Rev., 123, 485, doi: 10.1007/s11214-006-8315-7

Garzilli, A., Magalich, A., Ruchayskiy, O., & Boyarsky, A. 2021, MNRAS, 502, 2356, doi: 10.1093/mnras/stab192

Giannantonio, T., Scranton, R., Crittenden, R. G., et al. 2008, Phys. Rev. D, 77, 123520, doi: 10.1103/PhysRevD.77.123520

Gil-Marín, H., Noreña, J., Verde, L., et al. 2015, MNRAS, 451, 539, doi: 10.1093/mnras/stv961

Gilman, D., Birrer, S., & Treu, T. 2020, A&A, 642, A194, doi: 10.1051/0004-6361/202038829

Givans, J. J., Font-Ribera, A., Slosar, A., et al. 2022, J. Cosmology Astropart. Phys., 2022, 070, doi: 10.1088/1475-7516/2022/09/070

Gluscevic, V., et al. 2019, Bull. Am. Astron. Soc., 51, 134. https://arxiv.org/abs/1903.05140

Gómez-Valent, A., & Solà Peracaula, J. 2018, MNRAS, 478, 126, doi: 10.1093/mnras/sty1028

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759, doi: 10.1086/427976

Graczyk, D., Pietrzyński, G., Gieren, W., et al. 2019, ApJ, 872, 85, doi: 10.3847/1538-4357/aafbed

Green, D., & Meyers, J. 2021. https://arxiv.org/abs/2111.01096

Green, D., Meyers, J., & van Engelen, A. 2017, JCAP, 12, 005, doi: 10.1088/1475-7516/2017/12/005

Grin, D., Marsh, D. J. E., & Hlozek, R. 2022, axionCAMB: Modification of the CAMB Boltzmann code, Astrophysics Source Code Library, record ascl:2203.026. http://ascl.net/2203.026

Groenewegen, M. A. T. 2021, A&A, 654, A20, doi: 10.1051/0004-6361/202140862

Hadzhiyska, B., Sherwin, B. D., Madhavacheril, M., & Ferraro, S. 2019, Phys. Rev. D, 100, 023547, doi: 10.1103/PhysRevD.100.023547

Hagstotz, S., Reischke, R., & Lilow, R. 2022, MNRAS, 511, 662, doi: 10.1093/mnras/stac077

Hall, O. J., Davies, G. R., Elsworth, Y. P., et al. 2019, MNRAS, 486, 3569, doi: 10.1093/mnras/stz1092

Han, D., & Sehgal, N. 2022, Phys. Rev. D, 105, 083516, doi: 10.1103/PhysRevD.105.083516

Hanany, S., et al. 2019. https://arxiv.org/abs/1902.10541

Hanson, D., Challinor, A., Efstathiou, G., & Bielewicz, P. 2011, Phys. Rev. D, 83, 043005, doi: 10.1103/PhysRevD.83.043005

Hartlap, J., Simon, P., & Schneider, P. 2007, A&A, 464, 399, doi: 10.1051/0004-6361:20066170

Hawkins, K., Leistedt, B., Bovy, J., & Hogg, D. W. 2017, Monthly Notices of the Royal Astronomical Society, 471, 722, doi: 10.1093/mnras/stx1655

Hee, S., Handley, W. J., Hobson, M. P., & Lasenby, A. N. 2016, MNRAS, 455, 2461, doi: 10.1093/mnras/stv2217

Helmi, A., Babusiaux, C., Koppelman, H. H., et al. 2018, Nature, 563, 85, doi: 10.1038/s41586-018-0625-x

Henning, J. W., et al. 2018, Astrophys. J., 852, 97, doi: 10.3847/1538-4357/aa9ff4

Hirata, C. M., & Seljak, U. 2003a, Phys. Rev. D, 68, 083002, doi: 10.1103/PhysRevD.68.083002

—. 2003b, Phys. Rev. D, 67, 043001, doi: 10.1103/PhysRevD.67.043001

Hložek, R., Malz, A. I., Ponder, K. A., et al. 2023, ApJS, 267, 25, doi: 10.3847/1538-4365/accd6a

Ho, S., Hirata, C., Padmanabhan, N., Seljak, U., & Bahcall, N. 2008, Phys. Rev. D, 78, 043519, doi: 10.1103/PhysRevD.78.043519

Hojjati, A., Linder, E. V., & Samsing, J. 2013, Phys. Rev. Lett., 111, 041301, doi: 10.1103/PhysRevLett.111.041301

Holtzman, J. A., Shetrone, M., Johnson, J. A., et al. 2015, AJ, 150, 148, doi: 10.1088/0004-6256/150/5/148

Holwerda, B. W., Knabel, S., Steele, R. C., et al. 2021, MNRAS, 505, 1316, doi: 10.1093/mnras/stab1370

Holz, D. E., Hughes, S. A., & Schutz, B. F. 2018, Physics Today, 71, 34, doi: 10.1063/PT.3.4090

Horowitz, B., Ferraro, S., & Sherwin, B. D. 2019, Mon. Not. Roy. Astron. Soc., 485, 3919, doi: 10.1093/mnras/stz566

Hotinli, S. C., Meyers, J., Trendafilova, C., Green, D., & van Engelen, A. 2022, JCAP, 04, 020, doi: 10.1088/1475-7516/2022/04/020

Howlett, C., Lewis, A., Hall, A., & Challinor, A. 2012, J. Cosmology Astropart. Phys., 1204, 027, doi: 10.1088/1475-7516/2012/04/027

Hu, J.-P., & Wang, F.-Y. 2023, Universe, 9, 94, doi: 10.3390/universe9020094

Hu, W., DeDeo, S., & Vale, C. 2007, New J. Phys., 9, 441, doi: 10.1088/1367-2630/9/12/441

Hu, W., & Okamoto, T. 2002, Astrophys. J., 574, 566, doi: 10.1086/341110

Hu, W., & White, M. 1997, in Microwave Background Anisotropies, Vol. 16, 333–338, doi: 10.48550/arXiv.astro-ph/9606140

Huang, Q.-G., & Wang, K. 2016, European Physical Journal C, 76, 506, doi: 10.1140/epjc/s10052-016-4352-x

Hubble, E. 1929, Proceedings of the National Academy of Science, 15, 168, doi: 10.1073/pnas.15.3.168

Hui, L., Ostriker, J. P., Tremaine, S., & Witten, E. 2017, Phys. Rev. D, 95, 043541, doi: 10.1103/PhysRevD.95.043541

Hunt, J. A. S., Price-Whelan, A. M., Johnston, K. V., & Darragh-Ford, E. 2022, MNRAS, 516, L7, doi: 10.1093/mnrasl/slac082

Imara, N., & Blitz, L. 2007, ApJ, 662, 969, doi: 10.1086/517911

Iršič, V., Viel, M., Haehnelt, M. G., Bolton, J. S., & Becker, G. D. 2017a, Phys. Rev. Lett., 119, 031302, doi: 10.1103/PhysRevLett.119.031302

Iršič, V., Viel, M., Berg, T. A. M., et al. 2017b, MNRAS, 466, 4332, doi: 10.1093/mnras/stw3372

Iršič, V., Viel, M., Haehnelt, M. G., et al. 2017c, Phys. Rev. D, 96, 023522, doi: 10.1103/PhysRevD.96.023522

James, C. W., Ghosh, E. M., Prochaska, J. X., et al. 2022, MNRAS, 516, 4862, doi: 10.1093/mnras/stac2524

Jamieson, D., Li, Y., de Oliveira, R. A., et al. 2023, ApJ, 952, 145, doi: 10.3847/1538-4357/acdb6c

Jofré, P., Mädler, T., Gilmore, G., et al. 2015, MNRAS, 453, 1428, doi: 10.1093/mnras/stv1724

Joyce, J. 2021, in The Stanford Encyclopedia of Philosophy, Fall 2021 edn., ed. E. N. Zalta (Metaphysics Research Lab, Stanford University)

Kamionkowski, M., Kosowsky, A., & Stebbins, A. 1997, Phys. Rev. D, 55, 7368, doi: 10.1103/PhysRevD.55.7368

Kamionkowski, M., & Riess, A. G. 2022, arXiv e-prints, arXiv:2211.04492, doi: 10.48550/arXiv.2211.04492

Kaplinghat, M., Knox, L., & Song, Y.-S. 2003, Phys. Rev. Lett., 91, 241301, doi: 10.1103/PhysRevLett.91.241301

Karim, T., Singh, S., Rezaie, M., Hadzhiyska, B., & Eisenstein, D. 2022, in APS Meeting Abstracts, Vol. 2022, APS April Meeting Abstracts, H13.006

Karwal, T., & Kamionkowski, M. 2016, Phys. Rev. D, 94, 103523, doi: 10.1103/PhysRevD.94.103523

Keating, B., & Miller, N. 2006, New A Rev., 50, 184, doi: 10.1016/j.newar.2005.11.020

Keihänen, E., Keskitalo, R., Kurki-Suonio, H., Poutanen, T., & Sirviö, A. S. 2010, A&A, 510, A57, doi: 10.1051/0004-6361/200912813

Keihänen, E., Kurki-Suonio, H., & Poutanen, T. 2005, MNRAS, 360, 390, doi: 10.1111/j.1365-2966.2005.09055.x

Kesden, M., Cooray, A., & Kamionkowski, M. 2002, Phys. Rev. Lett., 89, 011304, doi: 10.1103/PhysRevLett.89.011304

Kesden, M., Cooray, A., & Kamionkowski, M. 2003, Phys. Rev. D, 67, 123507, doi: 10.1103/PhysRevD.67.123507

Khan, S., Miglio, A., Mosser, B., et al. 2019, A&A, 628, A35, doi: 10.1051/0004-6361/201935304

Khoperskov, S., Di Matteo, P., Gerhard, O., et al. 2019, A&A, 622, L6, doi: 10.1051/0004-6361/201834707

Kim, S. Y., Peter, A. H. G., & Hargis, J. R. 2018, Phys. Rev. Lett., 121, 211302, doi: 10.1103/PhysRevLett.121.211302

Kisner, T., Keskitalo, R., Zonca, A., et al. 2023, TOAST: Time Ordered Astrophysics Scalable Tools, Astrophysics Source Code Library, record ascl:2307.022. http://ascl.net/2307.022

Knox, L. 1995, Phys. Rev. D, 52, 4307, doi: 10.1103/PhysRevD.52.4307

Knox, L., & Song, Y.-S. 2002, Phys. Rev. Lett., 89, 011303, doi: 10.1103/PhysRevLett.89.011303

Kobayashi, T., Murgia, R., De Simone, A., Iršič, V., & Viel, M. 2017, Phys. Rev. D, 96, 123514, doi: 10.1103/PhysRevD.96.123514

Komatsu, E. 2022, Nature Reviews Physics, 4, 452, doi: 10.1038/s42254-022-00452-4

Kravtsov, A. 2010, Advances in Astronomy, 2010, 281913, doi: 10.1155/2010/281913

Kurki-Suonio, H., Keihänen, E., Keskitalo, R., et al. 2009, A&A, 506, 1511, doi: 10.1051/0004-6361/200912361

Kurucz, R. L. 2005, Memorie della Societa Astronomica Italiana Supplementi, 8, 14

Laguë, A., Bond, J. R., Hložek, R., et al. 2022, J. Cosmology Astropart. Phys., 2022, 049, doi: 10.1088/1475-7516/2022/01/049

Lane, J., Bovy, J., & Mackereth, T. 2023, arXiv e-prints, arXiv:2306.03084, doi: 10.48550/arXiv.2306.03084

Lawrence, E., Heitmann, K., White, M., et al. 2010, CosmicEmu: Cosmic Emulator for the Dark Matter Power Spectrum, Astrophysics Source Code Library, record ascl:1010.030. http://ascl.net/1010.030

Lemaître, G. 1933, Annales de la Soci&eacute;t&eacute; Scientifique de Bruxelles, 53, 51

Lemos, P., Lee, E., Efstathiou, G., & Gratton, S. 2019, MNRAS, 483, 4803, doi: 10.1093/mnras/sty3082

Lesgourgues, J., & Pastor, S. 2012, Adv. High Energy Phys., 2012, 608515, doi: 10.1155/2012/608515

Leung, H. W., & Bovy, J. 2019a, MNRAS, 489, 2079, doi: 10.1093/mnras/stz2245

—. 2019b, MNRAS, 483, 3255, doi: 10.1093/mnras/sty3217

Leung, H. W., Bovy, J., Mackereth, J. T., et al. 2023, MNRAS, 519, 948, doi: 10.1093/mnras/stac3529

Lewis, A., & Bridle, S. 2002, Phys. Rev. D, 66, 103511, doi: 10.1103/PhysRevD.66.103511

Lewis, A., & Challinor, A. 2006, Phys. Rep., 429, 1, doi: 10.1016/j.physrep.2006.03.002

—. 2011, CAMB: Code for Anisotropies in the Microwave Background, Astrophysics Source Code Library, record ascl:1102.026. http://ascl.net/1102.026

Lewis, A., Challinor, A., & Lasenby, A. 2000, ApJ, 538, 473, doi: 10.1086/309179

Li, M., Li, X.-D., Wang, S., & Wang, Y. 2013, Frontiers of Physics, 8, 828, doi: 10.1007/s11467-013-0300-5

Lin, X., Cai, Z., Li, Y., Krolewski, A., & Ferraro, S. 2020, ApJ, 905, 176, doi: 10.3847/1538-4357/abc620

Lindegren, L., Lammers, U., Bastian, U., et al. 2016, A&A, 595, A4, doi: 10.1051/0004-6361/201628714

Lindegren, L., Hernández, J., Bombrun, A., et al. 2018, A&A, 616, A2, doi: 10.1051/0004-6361/201832727

Lindegren, L., Klioner, S. A., Hernández, J., et al. 2021a, A&A, 649, A2, doi: 10.1051/0004-6361/202039709

Lindegren, L., Bastian, U., Biermann, M., et al. 2021b, A&A, 649, A4, doi: 10.1051/0004-6361/202039653

LiteBIRD Collaboration, Allys, E., Arnold, K., et al. 2023, Progress of Theoretical and Experimental Physics, 2023, 042F01, doi: 10.1093/ptep/ptac150

Liu, J., & Hill, J. C. 2015, Phys. Rev. D, 92, 063517, doi: 10.1103/PhysRevD.92.063517

Lokken, M., Hložek, R., van Engelen, A., et al. 2022, ApJ, 933, 134, doi: 10.3847/1538-4357/ac7043

Lovell, M. R., Frenk, C. S., Eke, V. R., et al. 2014, MNRAS, 439, 300, doi: 10.1093/mnras/stt2431

Ludlow, A. D., Navarro, J. F., Boylan-Kolchin, M., et al. 2013, MNRAS, 432, 1103, doi: 10.1093/mnras/stt526

Madhavacheril, M. S., & Hill, J. C. 2018, Phys. Rev. D, 98, 023534, doi: 10.1103/PhysRevD.98.023534

Madhavacheril, M. S., Hill, J. C., Næss, S., et al. 2020a, Phys. Rev. D, 102, 023534, doi: 10.1103/PhysRevD.102.023534

Madhavacheril, M. S., Sifón, C., Battaglia, N., et al. 2020b, ApJ, 903, L13, doi: 10.3847/2041-8213/abbccb

Madhavacheril, M. S., Qu, F. J., Sherwin, B. D., et al. 2023, arXiv e-prints, arXiv:2304.05203, doi: 10.48550/arXiv.2304.05203

Mädler, T., Jofré, P., Gilmore, G., et al. 2016, A&A, 595, A59, doi: 10.1051/0004-6361/201629091

Madore, B. F., & Freedman, W. L. 1991, PASP, 103, 933, doi: 10.1086/132911

Maíz Apellániz, J., & Weiler, M. 2018, A&A, 619, A180, doi: 10.1051/0004-6361/201834051

Majewski, S. R., Zasowski, G., & Nidever, D. L. 2011, ApJ, 739, 25, doi: 10.1088/0004-637X/739/1/25

Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, AJ, 154, 94, doi: 10.3847/1538-3881/aa784d

Marsh, D. J. E. 2016, Phys. Rep., 643, 1, doi: 10.1016/j.physrep.2016.06.005

Marshall, P., Rajguru, N., & Slosar, A. 2006, Phys. Rev. D, 73, 067302, doi: 10.1103/PhysRevD.73.067302

Mather, J. C., Cheng, E. S., Cottingham, D. A., et al. 1994, ApJ, 420, 439, doi: 10.1086/173574

McCarthy, F., Hill, J. C., & Madhavacheril, M. S. 2022, Phys. Rev. D, 105, 023517, doi: 10.1103/PhysRevD.105.023517

McDonald, P., Miralda-Escudé, J., Rauch, M., et al. 2000, ApJ, 543, 1, doi: 10.1086/317079

McDonald, P., Seljak, U., Cen, R., et al. 2005, ApJ, 635, 761, doi: 10.1086/497563

McKellar, A. 1940, PASP, 52, 187, doi: 10.1086/125159

Michilli, D., Bhardwaj, M., Brar, C., et al. 2023, ApJ, 950, 134, doi: 10.3847/1538-4357/accf89

Millea, M., Anderes, E., & Wandelt, B. D. 2020, Phys. Rev. D, 102, 123542, doi: 10.1103/PhysRevD.102.123542

Millea, M., et al. 2021, Astrophys. J., 922, 259, doi: 10.3847/1538-4357/ac02bb

Millon, M., Galan, A., Courbin, F., et al. 2020, A&A, 639, A101, doi: 10.1051/0004-6361/201937351

Mirmelstein, M., Fabbian, G., Lewis, A., & Peloton, J. 2021, Phys. Rev. D, 103, 123540, doi: 10.1103/PhysRevD.103.123540

Montgomery, J., Ade, P. A. R., Ahmed, Z., et al. 2022, Journal of Astronomical Telescopes, Instruments, and Systems, 8, 014001, doi: 10.1117/1.JATIS.8.1.014001

Mortonson, M. J., & Seljak, U. 2014, J. Cosmology Astropart. Phys., 2014, 035, doi: 10.1088/1475-7516/2014/10/035

Murgia, R., Iršič, V., & Viel, M. 2018, Phys. Rev. D, 98, 083540, doi: 10.1103/PhysRevD.98.083540

NASA. 2023, NASA. https://universe.nasa.gov/universe/basics/

Nashimoto, M., Tanaka, M., Chiba, M., et al. 2022, ApJ, 936, 38, doi: 10.3847/1538-4357/ac83a4

Neal, R. M. 2012, arXiv e-prints, arXiv:1206.1901. https://arxiv.org/abs/1206.1901

Netterfield, C. B., Ade, P. A. R., Bock, J. J., et al. 2002, ApJ, 571, 604, doi: 10.1086/340118

Nishizawa, A. J. 2014, Progress of Theoretical and Experimental Physics, 2014, 06B110, doi: 10.1093/ptep/ptu062

Nunes, R. C., & Vagnozzi, S. 2021, MNRAS, 505, 5427, doi: 10.1093/mnras/stab1613

Okamoto, T., & Hu, W. 2003, Phys. Rev. D, 67, 083002, doi: 10.1103/PhysRevD.67.083002

Osborne, S. J., Hanson, D., & Doré, O. 2014, JCAP, 03, 024, doi: 10.1088/1475-7516/2014/03/024

Palanque-Delabrouille, N., Yèche, C., Schöneberg, N., et al. 2020, J. Cosmology Astropart. Phys., 2020, 038, doi: 10.1088/1475-7516/2020/04/038

Palanque-Delabrouille, N., Yèche, C., Borde, A., et al. 2013, A&A, 559, A85, doi: 10.1051/0004-6361/201322130

Paszke, A., Gross, S., Massa, F., et al. 2019, arXiv e-prints, arXiv:1912.01703, doi: 10.48550/arXiv.1912.01703

Patil, A. A., Bovy, J., Jaimungal, S., Frankel, N., & Leung, H. W. 2023, arXiv e-prints, arXiv:2306.09319, doi: 10.48550/arXiv.2306.09319

Pedersen, C., Font-Ribera, A., Rogers, K. K., et al. 2021, J. Cosmology Astropart. Phys., 2021, 033, doi: 10.1088/1475-7516/2021/05/033

Peebles, P. J. E. 1965, ApJ, 142, 1317, doi: 10.1086/148417

Penzias, A. A., & Wilson, R. W. 1965, ApJ, 142, 419, doi: 10.1086/148307

Perlmutter, S., Aldering, G., Goldhaber, G., et al. 1999, ApJ, 517, 565, doi: 10.1086/307221

Petersen, M. S., & Peñarrubia, J. 2020, MNRAS, 494, L11, doi: 10.1093/mnrasl/slaa029

Piccirilli, G., Migliaccio, M., Branchini, E., & Dolfi, A. 2022, arXiv e-prints, arXiv:2208.07774, doi: 10.48550/arXiv.2208.07774

Pietrzyński, G., Graczyk, D., Gallenne, A., et al. 2019, Nature, 567, 200, doi: 10.1038/s41586-019-0999-4

Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014a, A&A, 571, A1, doi: 10.1051/0004-6361/201321529

—. 2014b, A&A, 571, A16, doi: 10.1051/0004-6361/201321591

Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2016a, A&A, 594, A1, doi: 10.1051/0004-6361/201527101

Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016b, A&A, 594, A13, doi: 10.1051/0004-6361/201525830

Planck Collaboration, Aghanim, N., Ashdown, M., et al. 2016c, A&A, 596, A107, doi: 10.1051/0004-6361/201628890

Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020a, A&A, 641, A6, doi: 10.1051/0004-6361/201833910

—. 2020b, A&A, 641, A1, doi: 10.1051/0004-6361/201833880

Posti, L., & Helmi, A. 2019, A&A, 621, A56, doi: 10.1051/0004-6361/201833355

Poulin, V., Bernal, J. L., Kovetz, E. D., & Kamionkowski, M. 2023, Phys. Rev. D, 107, 123538, doi: 10.1103/PhysRevD.107.123538

Poulin, V., Smith, T. L., Karwal, T., & Kamionkowski, M. 2019, Phys. Rev. Lett., 122, 221301, doi: 10.1103/PhysRevLett.122.221301

Primack, J. R., & Gross, M. A. K. 2001, in Current aspects of neutrino physics, ed. D. O. Caldwell, 287–308, doi: 10.48550/arXiv.astro-ph/0007165

Qu, F. J., Sherwin, B. D., Madhavacheril, M. S., et al. 2023, arXiv e-prints, arXiv:2304.05202, doi: 10.48550/arXiv.2304.05202

Raghunathan, S., Bianchini, F., & Reichardt, C. L. 2018, Phys. Rev. D, 98, 043506, doi: 10.1103/PhysRevD.98.043506

Reichardt, C. L. 2016, in Astrophysics and Space Science Library, Vol. 423, Understanding the Epoch of Cosmic Reionization: Challenges and Progress, ed. A. Mesinger, 227, doi: 10.1007/978-3-319-21957-8_8

Reinecke, M., Belkner, S., & Carron, J. 2023, arXiv e-prints, arXiv:2304.10431, doi: 10.48550/arXiv.2304.10431

Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9143, Space Telescopes and Instrumentation 2014: Optical, Infrared, and Millimeter Wave, ed. J. Oschmann, Jacobus M., M. Clampin, G. G. Fazio, & H. A. MacEwen, 914320, doi: 10.1117/12.2063489

Riello, M., De Angeli, F., Evans, D. W., et al. 2018, A&A, 616, A3, doi: 10.1051/0004-6361/201832712

—. 2021, A&A, 649, A3, doi: 10.1051/0004-6361/202039587

Riess, A. G., Casertano, S., Yuan, W., Macri, L. M., & Scolnic, D. 2019, ApJ, 876, 85, doi: 10.3847/1538-4357/ab1422

Riess, A. G., Filippenko, A. V., Challis, P., et al. 1998, AJ, 116, 1009, doi: 10.1086/300499

Riess, A. G., Casertano, S., Yuan, W., et al. 2018, ApJ, 861, 126, doi: 10.3847/1538-4357/aac82e

Riess, A. G., Yuan, W., Macri, L. M., et al. 2022, ApJ, 934, L7, doi: 10.3847/2041-8213/ac5c5b

Riess, A. G., Anand, G. S., Yuan, W., et al. 2023, arXiv e-prints, arXiv:2307.15806, doi: 10.48550/arXiv.2307.15806

Robertson, H. P. 1935, ApJ, 82, 284, doi: 10.1086/143681

Robertson, N. C., Alonso, D., Harnois-Déraps, J., et al. 2021, A&A, 649, A146, doi: 10.1051/0004-6361/202039975

Rogers, K. K., Hložek, R., Laguë, A., et al. 2023, J. Cosmology Astropart. Phys., 2023, 023, doi: 10.1088/1475-7516/2023/06/023

Rogers, K. K., Peiris, H. V., Pontzen, A., et al. 2019, J. Cosmology Astropart. Phys., 2019, 031, doi: 10.1088/1475-7516/2019/02/031

Ross, A. J., Bautista, J., Tojeiro, R., et al. 2020, MNRAS, 498, 2354, doi: 10.1093/mnras/staa2416

Sahlholdt, C. L., & Silva Aguirre, V. 2018, MNRAS, 481, L125, doi: 10.1093/mnrasl/sly173

Salvatier, J., Wiecki, T., & Fonnesbeck, C. 2015, arXiv e-prints, arXiv:1507.08050, doi: 10.48550/arXiv.1507.08050

Schaffer, K. K., Crawford, T. M., Aird, K. A., et al. 2011, ApJ, 743, 90, doi: 10.1088/0004-637X/743/1/90

Schlafly, E. F., & Finkbeiner, D. P. 2011, ApJ, 737, 103, doi: 10.1088/0004-637X/737/2/103

Schmittfull, M., & Seljak, U. 2018, Phys. Rev. D, 97, 123540, doi: 10.1103/PhysRevD.97.123540

Schöneberg, N., Lesgourgues, J., & Hooper, D. C. 2019, J. Cosmology Astropart. Phys., 2019, 029, doi: 10.1088/1475-7516/2019/10/029

Scolnic, D., Brout, D., Carr, A., et al. 2022, ApJ, 938, 113, doi: 10.3847/1538-4357/ac8b7a

Scolnic, D., Riess, A. G., Wu, J., et al. 2023, arXiv e-prints, arXiv:2304.06693, doi: 10.48550/arXiv.2304.06693

Sehgal, N., et al. 2019. https://arxiv.org/abs/1906.10134

Seljak, U., & Hirata, C. M. 2004, Phys. Rev. D, 69, 043005, doi: 10.1103/PhysRevD.69.043005

Seljak, U., & Zaldarriaga, M. 1996, ApJ, 469, 437, doi: 10.1086/177793

Seljak, U., & Zaldarriaga, M. 2000, Astrophys. J., 538, 57, doi: 10.1086/309098

Sesar, B., Fouesneau, M., Price-Whelan, A. M., et al. 2017, ApJ, 838, 107, doi: 10.3847/1538-4357/aa643b

Shajib, A. J., Mozumdar, P., Chen, G. C. F., et al. 2023, A&A, 673, A9, doi: 10.1051/0004-6361/202345878

Sherwin, B. D., Das, S., Hajian, A., et al. 2012, Phys. Rev. D, 86, 083006, doi: 10.1103/PhysRevD.86.083006

Silk, J. 1968, ApJ, 151, 459, doi: 10.1086/149449

Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, AJ, 131, 1163, doi: 10.1086/498708

Smith, K. M., & Ferraro, S. 2017, Phys. Rev. Lett., 119, 021301, doi: 10.1103/PhysRevLett.119.021301

Smith, K. M., Hanson, D., LoVerde, M., Hirata, C. M., & Zahn, O. 2012, JCAP, 06, 014, doi: 10.1088/1475-7516/2012/06/014

Smith, M. J., & Geach, J. E. 2023, Royal Society Open Science, 10, 221454, doi: 10.1098/rsos.221454

Smith, T. L., Poulin, V., Bernal, J. L., et al. 2021, Phys. Rev. D, 103, 123542, doi: 10.1103/PhysRevD.103.123542

Smoot, G. F., Bennett, C. L., Kogut, A., et al. 1991, ApJ, 371, L1, doi: 10.1086/185988

Sobrin, J. A., Anderson, A. J., Bender, A. N., et al. 2022, ApJS, 258, 42, doi: 10.3847/1538-4365/ac374f

SPIDER Collaboration, Ade, P. A. R., Amiri, M., et al. 2021, arXiv e-prints, arXiv:2103.13334, doi: 10.48550/arXiv.2103.13334

Spurio Mancini, A., Piras, D., Alsing, J., Joachimi, B., & Hobson, M. P. 2022, MNRAS, 511, 1771, doi: 10.1093/mnras/stac064

Stassun, K. G., & Torres, G. 2018, ApJ, 862, 61, doi: 10.3847/1538-4357/aacafc

Stein, G., Alvarez, M. A., & Bond, J. R. 2019, MNRAS, 483, 2236, doi: 10.1093/mnras/sty3226

Stein, G., Alvarez, M. A., Bond, J. R., van Engelen, A., & Battaglia, N. 2020, J. Cosmology Astropart. Phys., 2020, 012, doi: 10.1088/1475-7516/2020/10/012

Stoehr, F., White, S. D. M., Tormen, G., & Springel, V. 2002, MNRAS, 335, L84, doi: 10.1046/j.1365-8711.2002.05891.x

Story, K. T., et al. 2015, Astrophys. J., 810, 50, doi: 10.1088/0004-637X/810/1/50

Sun, Z., Yao, J., Dong, F., et al. 2022, MNRAS, 511, 3548, doi: 10.1093/mnras/stac138

Sunyaev, R. A., & Zeldovich, Y. B. 1970, Ap&SS, 7, 20, doi: 10.1007/BF00653472

—. 1972, Comments on Astrophysics and Space Physics, 4, 173

—. 1980, MNRAS, 190, 413, doi: 10.1093/mnras/190.3.413

Suzuki, A., Cothard, N., Lee, A. T., et al. 2020, Journal of Low Temperature Physics, 199, 1158, doi: 10.1007/s10909-019-02325-0

Tegmark, M., de Oliveira-Costa, A., & Hamilton, A. J. 2003, Phys. Rev. D, 68, 123523, doi: 10.1103/PhysRevD.68.123523

Thorne, B., Alonso, D., Naess, S., & Dunkley, J. 2017, PySM: Python Sky Model, Astrophysics Source Code Library, record ascl:1704.007. http://ascl.net/1704.007

Torrado, J., & Lewis, A. 2019, Cobaya: Bayesian analysis in cosmology, Astrophysics Source Code Library, record ascl:1910.019. http://ascl.net/1910.019

—. 2021, J. Cosmology Astropart. Phys., 2021, 057, doi: 10.1088/1475-7516/2021/05/057

Tremaine, S., Frankel, N., & Bovy, J. 2023, MNRAS, 521, 114, doi: 10.1093/mnras/stad577

Tsujikawa, S. 2003, arXiv e-prints, hep, doi: 10.48550/arXiv.hep-ph/0304257

Tulin, S., & Yu, H.-B. 2018, Phys. Rept., 730, 1, doi: 10.1016/j.physrep.2017.11.004

Uddin, S. A., Burns, C. R., Phillips, M. M., et al. 2023, arXiv e-prints, arXiv:2308.01875, doi: 10.48550/arXiv.2308.01875

Ulbricht, G., De Lucia, M., & Baldwin, E. 2021, arXiv e-prints, arXiv:2111.08576, doi: 10.48550/arXiv.2111.08576

van Engelen, A., Bhattacharya, S., Sehgal, N., et al. 2014, Astrophys. J., 786, 13, doi: 10.1088/0004-637X/786/1/13

Verde, L., Protopapas, P., & Jimenez, R. 2013, Physics of the Dark Universe, 2, 166, doi: 10.1016/j.dark.2013.09.002

Viel, M., Becker, G. D., Bolton, J. S., & Haehnelt, M. G. 2013, Phys. Rev. D, 88, 043502, doi: 10.1103/PhysRevD.88.043502

Villasenor, B., Robertson, B., Madau, P., & Schneider, E. 2023, Phys. Rev. D, 108, 023502, doi: 10.1103/PhysRevD.108.023502

Vogt, S. M. L., Marsh, D. J. E., & Laguë, A. 2023, Phys. Rev. D, 107, 063526, doi: 10.1103/PhysRevD.107.063526

Walker, A. G. 1935, The Quarterly Journal of Mathematics, 1, 81, doi: 10.1093/qmath/os-6.1.81

Walters, A., Weltman, A., Gaensler, B. M., Ma, Y.-Z., & Witzemann, A. 2018, ApJ, 856, 65, doi: 10.3847/1538-4357/aaaf6b

Wang, S., & Chen, X. 2019, ApJ, 877, 116, doi: 10.3847/1538-4357/ab1c61

Weinberg, D. H., Bullock, J. S., Governato, F., Kuzio de Naray, R., & Peter, A. H. G. 2015, Proceedings of the National Academy of Science, 112, 12249, doi: 10.1073/pnas.1308716112

Wilson, J. C., Hearty, F. R., Skrutskie, M. F., et al. 2019, PASP, 131, 055001, doi: 10.1088/1538-3873/ab0075

Wong, K. C., Suyu, S. H., Chen, G. C. F., et al. 2020, MNRAS, 498, 1420, doi: 10.1093/mnras/stz3094

Wu, Q., Yu, H., & Wang, F. Y. 2020, ApJ, 895, 33, doi: 10.3847/1538-4357/ab88d2

Yan, Z., van Waerbeke, L., Tröster, T., et al. 2021, A&A, 651, A76, doi: 10.1051/0004-6361/202140568

Yıldırım, A., Suyu, S. H., Chen, G. C. F., & Komatsu, E. 2023, A&A, 675, A21, doi: 10.1051/0004-6361/202142318

Zaldarriaga, M. 2000, Phys. Rev. D, 62, 063510, doi: 10.1103/PhysRevD.62.063510

—. 2001, Phys. Rev. D, 64, 103001, doi: 10.1103/PhysRevD.64.103001

Zaldarriaga, M., & Seljak, U. 1997, Phys. Rev. D, 55, 1830, doi: 10.1103/PhysRevD.55.1830

Zaroubi, S., Viel, M., Nusser, A., Haehnelt, M., & Kim, T. S. 2006, MNRAS, 369, 734, doi: 10.1111/j.1365-2966.2006.10333.x

Zavala, J., & Frenk, C. S. 2019, Galaxies, 7, 81, doi: 10.3390/galaxies7040081

Zel'dovich, Y. B. 1968, Soviet Physics Uspekhi, 11, 381, doi: 10.1070/PU1968v011n03ABEH003927

Zeldovich, Y. B., & Sunyaev, R. A. 1969, Ap&SS, 4, 301, doi: 10.1007/BF00661821

Zinn, J. C., Pinsonneault, M. H., Huber, D., & Stello, D. 2019, ApJ, 878, 136, doi: 10.3847/1538-4357/ab1f66