Unveiling Our Early Universe Using Cross-Correlation Analyses of Multiple Cosmological Probes

Hannah Fronenberg

Department of Physics
McGill University
Montréal, Québec, Canada

February 15, 2025

A thesis presented for the degree of Doctor of Philosophy of Physics

Abstract

Line intensity mapping (LIM) is a novel observational technique that allows us to probe cosmological and astrophysical information over the majority of cosmic history. Despite its promise, these observations are challenged by a variety of bright foreground contaminants as well as by instrument systematics. Luckily, cross-correlations are a powerful tool for mitigating these effects and extracting cosmological information. In this thesis, we demonstrate how combining LIM observations with cross-correlation analyses can maximize the yield of these surveys, showcasing three salient examples. First, we show how combining multiple LIM observations, particularly that of neutral hydrogen and ionized carbon, allows us to extract important information about the early epoch of reionization. Next, to resolve the problem of bright foreground contaminants, we show how intra-dataset cross-correlations can be used to perform foreground removal with percent-level residuals in a model-independent way. Finally, LIM observations can be combined with cosmic microwave background (CMB) lensing observations to make direct measurements of the high-redshift matter density field. These new LIM-CMB observables enable model-independent test of cosmology beyond ΛCDM and can help constrain various properties of the high-redshift matter power spectrum, most notably, making unbiased measurements of the baryon acoustic oscillation (BAO) scale across a large portion of cosmic history.

Abrégé

La cartographie de l'intensité des raies (LIM) est une nouvelle technique d'observation qui nous permet de sonder des informations cosmologiques et astrophysiques sur la majeure partie de l'histoire cosmique. Malgré ses promesses, ces observations sont entravées par une variété de contaminants d'avant-plan brillants ainsi que par la systématique des instruments. Heureusement, les corrélations croisées sont un outil puissant pour atténuer ces effets et extraire des informations cosmologiques. Dans cette thèse, nous démontrons comment la combinaison d'observations LIM avec des analyses de corrélations croisées peut maximiser le rendement de ces enquêtes, en présentant trois exemples marquants. Tout d'abord, nous montrons comment la combinaison de plusieurs observations LIM, en particulier celle de l'hydrogène neutre et du carbone ionisé, nous permet d'extraire des informations importantes sur l'époque de la réionisation. Ensuite, pour résoudre le problème des contaminants d'avant-plan lumineux, nous montrons comment les corrélations croisées d'un ensemble de données avec lui-même peuvent être utilisées pour éliminer l'avant-plan avec des résidus au niveau du pourcentage d'une manière indépendante du modèle. Enfin, les observations LIM peuvent être combinées avec les observations du fond diffus cosmologique (CMB) pour effectuer des mesures directes du champ de densité de matière à un décalage horaire élevé. Ces nouvelles observations LIM-CMB permettent de tester la cosmologie indépendamment du modèle au-delà du modèle ΛCDM et peuvent aider à contraindre diverses propriétés du spectre de puissance de la matière à un décalage élevé, notamment en effectuant des mesures non biaisées de l'échelle de l'oscillation acoustique du baryon sur une grande partie de l'histoire cosmique.

Acknowledgements

These past several years have been the most personally fulfilling of my life thus far. Thinking back to when I started my graduate career roughly five and a half years ago, I could not have imagined they would turn out the way they did. It has been such a privilege to be a part of the Trottier Space Institue and the McGill Physics Department and to be surrounded by the many people who shaped me as a scientist.

First and foremost, this thesis would not exist without my advisor, Adrian Liu. Adrian embodies all of the qualities of a great advisor and educator. He is a great explainer of things, he encourages us to step out of our comfort zones and take on new challenges, he provides unwavering support when things go wrong and when things go right, and in many instances, he believes in us more than we may believe in ourselves. Adrian is an advisor who I believe is equaled by few and surpassed by none. I feel lucky for the opportunity to have been a member of his group.

I would also like to thank my close collaborators Anthony Pullen and Abhishek Maniyar whose work spawned Chapters 4 and 5. Anthony welcomed me into his group at New York University over the summer of 2022, an experience for which I am forever grateful. He made me feel as though I was really just another one of his grad students and I am grateful for the time he devoted to our work and to giving me career advice this past year. Thank you to Abhishek who spent countless hours on Zoom working through bugs and computations for Chapters 4 and 5 and for all of your life advice over the years.

I would also like to thank my supervisory committee members, Matt Dobbs and Vicky Kaspi.

Matt’s expertise in all things instrumentation and data analysis has pushed me as a theorist to think about the “real” problems and inspired much of the analysis in Chapter 3. Vicky’s probing questions during our meetings have always reminded me that we cannot do good cosmology without doing good astrophysics.

I will deeply miss our cosmology community at McGill. From pencil and paper theory to instrumentation, I truly feel as though I have had the opportunity to learn about every aspect of the crazy science we all attempt to do. I have felt so supported by so many but would like to especially thank Robert Brandenberger who has been a cherished collaborator, mentor, and Swiss tour guide, Katelin Schutz who is an inspiration to me and who provides so much support during stressful times (usually in the form of cupcakes), Jon Sievers who in a single sentence can convey a more profound (albeit convoluted) thought than I could in hundreds, and Cynthia Chiang who has inspired me to dream big and who actually got me to use a screwdriver at least once during this degree!

My graduate career has also been shaped by my time working as the TSI Outreach Coordinator, a job that was made so rich thanks to our devoted attendees and volunteers. I always looked forward to my time spent talking about science and space with the greater Montreal community. My interactions with enthusiastic science lovers, both young and old, always filled me with a sense of wonder about our Universe, a feeling that can easily be lost over the course of a graduate degree. This work would not have been possible without Carolina Cruz-Vinaccia who has taught me so much about what it means to be a great educator, Kim Metera who always saved the day, Thomas Brunner, our Outreach Chair, who supported all my crazy ideas and gave up part of his lab space to build countless demos, and to my fellow grad coordinators Alice Curtin and David Gallacher who worked tirelessly to support our programs and made it look easy.

My experience as a graduate student was shaped overwhelmingly by the McGill Cosmic Dawn Group. It has been an honour to work alongside so many talented scientists. I am eternally grateful to my academic siblings Michael Pagano, Lisa McBride, Bobby Pascua, Matteo Blamart, Kai-Feng Chen (beloved step-sibling), Rebecca Ceppas de Castro, Sophia Rubens, Kim Morel, and Josh Goodeve, and to our genius postdocs (although their modesty would never allow them to describe themselves as such) Jordan Mirocha, Saurabh Singh, Adélie Gorce, Ronniy Joseph, Arnab Chakraborty, Debanjan Sarkar, Mike Wilensky, and Libby Berkhout. I feel so lucky to have this group of people in my corner, who always provide deep new insights to any problem. I have also been inspired by our undergraduate researchers who have impressed me with the science they accomplish while also taking a full course load, and who have taught me more than they probably realize. To our group, thank you. My academic older brother, Michael Pagano, said it perfectly, these people aforementioned have "made the term "academic family" feel less academic and more like family".

I am forever grateful for my time spent in New York City which was a period of immense scientific and personal growth. To my New York family, Janna Dutaud and Nora Von Moltke Simms, you made me feel at home from the day I moved in. Without you there would be no Chapters 4 and 5. To Jaclyn Lavy and Karl Valentini, thank you for being a little piece of Montreal in the Big City.

To my office mates, I have immense gratitude for your support, day in and day out. To Rutherford room 208 season 1: Félix Proulx-Giraldeau, Francois Bourassa, Laurent Jutras-Dubé, Victoria Mochulska, thank you for bringing a cosmologist into your biophysics group, I learned more about the immune system and embryo development than I ever thought I would and I am grateful for it. To Rutherford room 208 season 2: Matteo Blamart, Rebecca Ceppas de Castro, Sophia Rubens, Kim Morel, Jen Glover, and Josh Goodeve, I would not have survived this last year of my degree without you all.

I would also like express my gratitude to Prof. Tim Pollack. Tim was the first "real" scientist I ever knew and is someone who truly embodies the scientific spirit. From a young age, he encouraged me to learn about the natural world and is somehow able to discuss any topic in great detail, something to which I always aspired. Although I did not follow in his footsteps in becoming a pharmacologist and clinician (despite my best attempt), Tim always supported and encouraged my pursuit for academic research. I would not be a scientist today without Tim.

And finally, to my parents, thank you for everything.

Statement of Contribution

This thesis is based on four articles:

- ***Chapter 2***, **Fronenberg, H.**, Liu, A. (2024) *Forecasts and Statistical Insights for Line Intensity Mapping Cross-Correlations: A Case Study with 21cm × [CII]*, ApJ **975** 222
- ***Chapter 3***, **Fronenberg, H.**, Liu, A. (2025) *A Data-Driven Technique to Mitigate the Foregrounds of Line Intensity Maps*, in prep.
- ***Chapter 4***, **Fronenberg, H.**, Maniyar, A.S., Pullen, A.R., Liu, A. (2024) *Constraining Cosmology With the CMB × LIM-Nulling Convergence*, Phys. Rev. D **109** 123518
- ***Chapter 5***, **Fronenberg, H.**, Maniyar, A.S., Liu, A., Pullen, A.R. (2024) *New Probe of the High-z BAO scale: BAO tomography With CMB × LIM-Nulling Convergence*, Phys. Rev. Lett. **132** 241001

These articles represent a significant component of my research. For all of these chapters, I ran all simulations, developed the necessary computational frameworks for analysis, and performed the analyses. I developed the methodologies used in the research, in collaboration with and with input from my coauthors. I generated all figures presented in the articles. I wrote the first full draft of each paper, with valuable feedback and edits provided by my coauthors. An exception to this is Appendix B of Chapter 2, which was written entirely by my coauthor Adrian Liu. The additional text of this thesis, including Chapter 1 (Introduction), Chapter 6 (Conclusion), and the intermezzos between chapters, was written entirely by me.

While these articles were the result of a collaborative effort, the core analyses, writing, and presentation of results were primarily my own. I am deeply grateful to my coauthors for their insightful contributions, discussions, and feedback, which helped shape this body of work.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Cosmology's Precision Era

In recent decades, cosmology has entered the era of precision measurement. Observations of the cosmic microwave background (CMB), large-scale galaxy surveys, measurements of primordial chemical abundances, gravitational lensing, and most recently, gravitational wave interferometry, have given us a clearer picture of the Universe in which we live (Planck Collaboration, 2016; Peebles, 2014; Abbott et al., 2016; Cyburt et al., 2016; Bundy et al., 2015). From this abundance of data, we have learned that we live in a spatially flat universe dominated by a cosmological constant leading to today's accelerated expansion (Perlmutter and Project, 1999). Our matter content is dominated by dark matter, and we have constrained the age of the observable universe to 13.772 ± 0.040 Gyr (Planck Collaboration, 2016). These simple properties that describe our complex Universe constitute the Λ Cold Dark Matter (ΛCDM) model, our current concordance model of cosmology. This simple six parameter model has been able to provide an account for the existence of the CMB and its statistics, for the large-scale structure of the distribution of galaxies, for the primordial chemical abundances, and for our current accelerated expansion. Despite cosmology's observational and theoretical successes, the vast majority of the volume of our Universe remains unexplored, which has left us with unanswered questions about the origins of all the structures we see today. How did a universe, which at early times was filled only with individual atoms, come to form enormous galaxies, clusters, and filaments? What were the initial conditions which seeded this structure in the first place? What is our Universe's expansion history and what is the dark energy that drives it?

Despite much uncertainty about the details of the very early universe, the CMB has provided us with strong evidence that quantum fluctuations seeded matter perturbations (see Section 1.3 for an overview of the CMB). Under gravity, these matter overdensities grew to form the large astrophysical objects we see today. Through the process of gravitational structure growth, the Universe's first stars ignited from the collapse of hydrogen around 100 million years after the Big Bang. In term of cosmological redshift (z), which is defined as the amount by which the wavelength of light has been stretched due to cosmic expansion, the first stars are expected to have been born at $z \sim 20$. The end of this cosmic dawn would mark the beginning of reionization, when ultraviolet (UV) photons emanating from the first galaxies ionized the surrounding neutral hydrogen (HI) in the intergalactic medium (IGM) left over from the time when the first atoms formed called recombination. These early epochs: the cosmic Dark Ages (a period after recombination but before galaxy formation), Cosmic Dawn (when the first luminous objects ignited), and the Epoch of Reionization (EoR, when the first galaxies ionized the IGM) remain poorly constrained and understood. Without them, we fail to piece together a complete and uninterrupted history of our Universe.

Over the last two decades, there has been sustained effort to make observations of these early post-recombination epochs. Most notably, the James Webb Space Telescope (JWST), since its launch in 2021, has provided us with some of the most detailed observations of high-redshift galaxies. Additionally, observations of the Gunn-Peterson trough, an extended hydrogen absorption feature in the spectra of high redshift quasars, have revealed that reionization likely persisted until $z \sim 5.5$ (Becker et al., 2015). Even with these strides forward, the sources observed with JWST only comprise the brightest population of galaxies in rest-frame UV, and the small population of observed quasars raises questions about the homogeneity of the ending of reionization. In order to progress towards an understanding of the early period of our cosmic timeline, it is imperative that we develop novel observables specifically designed to probe the statistical properties of the early Universe.

Line intensity mapping (LIM) offers a path forward, allowing us to explore vast volumes of the uncharted Universe thus bridging the gaps in our knowledge. This nascent technique relies on observing the integrated intensity of a single atomic or molecular transition line emanating from galaxies and the IGM. When it comes to mapping large-scale structure, LIM is a transformative tool which, unlike galaxy surveys, does not necessitate resolving individual sources. This means that high-sensitivity, high-redshift LIM measurements can be achieved with relatively modest instruments, which are nonetheless still capable of capturing the spatially fluctuating line emission from the underlying structure, including contributions from unresolved galaxies. What is more, by observing a single spectral line with known rest frequency, LIM yields precise line-of-sight information since the observed frequency is directly related to the redshift at which the line was emitted.

LIM is still in its early days and most efforts are focused on producing statistical measurements while striving to reach the goal of producing the first-ever maps. These early measurements have not only highlighted LIM's scientific potential but also revealed the real-world challenges it faces. In particular, virtually all lines being targeted suffer from bright foreground contamination that drowns out the LIM signal. Due to both the nature of the signal itself as well as instrument's response and systematics, existing foreground removal methods, for instance those developed for the CMB, have not seen much success in teasing out high-redshift line emission. We expand on LIM and its challenges in Sections 1.2 and 1.5 respectively.

Another avenue to the Dark Ages, Cosmic Dawn, and the EoR is through the CMB. The CMB photons which were emitted roughly 400,000 years after the Big Bang ($z \sim 1100$), travel through the vast majority of the Universe's history on their journey to us. Along the way, these photons experience a host of scattering and gravitational interactions with intervening matter, picking up late-time or *secondary* signals in addition to their *primary* signal from the surface of last scattering.

A particularly powerful CMB secondary is that of gravitational lensing whereby CMB photons are lensed by the gravitational potentials they encounter. Since lensing is sensitive to both baryonic and dark matter, encoded in the CMB photons is information about the total underlying mass distribution. This makes lensing an exquisite probe of the matter density field as it provides a representative census of all matter in the Universe without using a biased tracer. Unfortunately, the observed lensing signal is a integrated projection of the entire line-of-sight (LOS) onto a single plane, making it impossible to disentangle contributions from different epochs and study their evolution.

These probes, LIM and CMB lensing, hold immense potential, which can be fully realized by leveraging synergies. Figure 1.1 showcases just how much of the cosmic timeline is probed by the various techniques. Zooming in, Figure 1.2 provides an illustration of the EoR, showing how line emission can trace both the neutral IGM (with the 21cm line of netural hydrogen) and galactic populations (with, for example, the Lyman-α (Lyα) line of hydrogen, various carbon monoxide (CO) lines, or with ionized carbon (CII) lines), while CMB lensing probes the foundation of it all, the matter density field. Taken together, these approaches allow us to construct a comprehensive view of much of the Universe's history.

This thesis aims to demonstrate how we can enhance the scientific return of LIM by leveraging synergies, that is, by combining LIM datasets with one another or with well-established CMB observations. These approaches not only maximize the information we extract from LIM but also help mitigate key observational challenges. In the following sections, we take a deep dive into the theoretical and observational background, laying the foundation for the methods developed in this work.

1.2 Line Intensity Mapping

In this section we provide a theoretical overview of LIM and as well as an update on the current status of LIM measurements. Generally, the field can be subdivided into two main branches: 21 cm

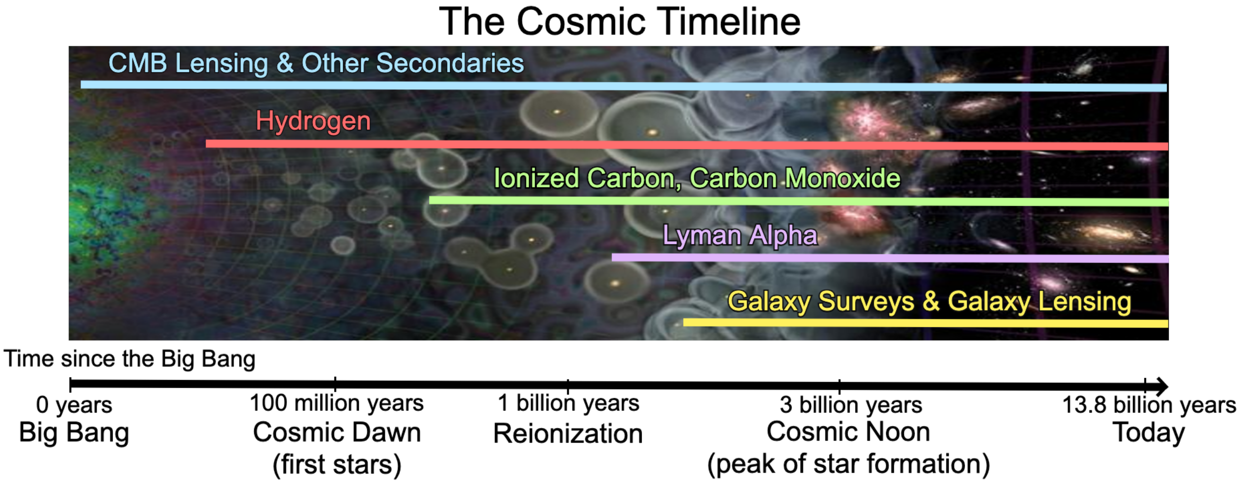

Figure 1.1: The cosmic timeline with key epochs listed below. Superimposed coloured lines denote what portions of our history can be probed by different observables.

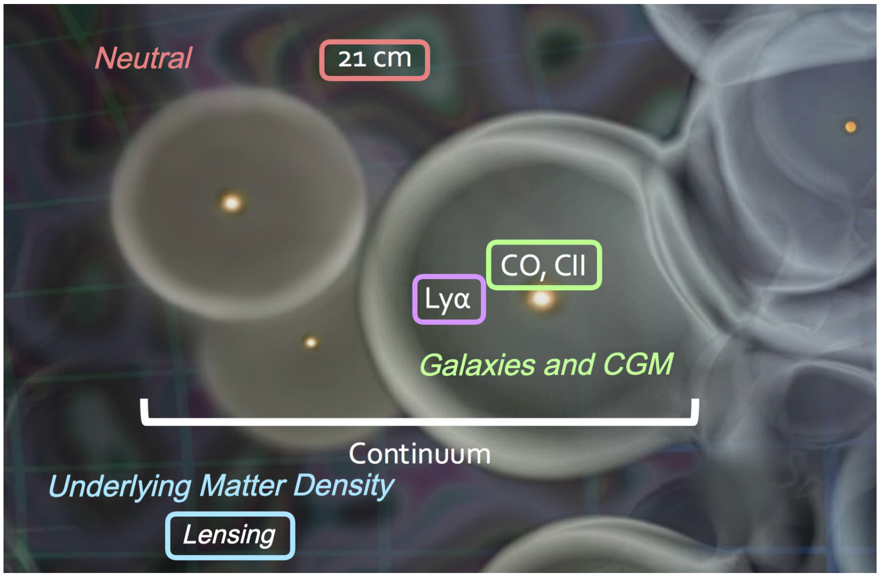

Figure 1.2: Illustration of signals from lensing (blue) and tracers of the neutral IGM (red), and of galaxies and the circumgalactic medium (CGM) (green and purple).

intensity mapping (referring to the mapping of the hyperfine transition line of neutral hydrogen) and the intensity mapping of all the other lines that trace galactic populations. These other lines include the Lyα line, the Hα line, a host of rotational transition lines of carbon monoxide (CO), and the ionized carbon line ([CII]) just to name a few. Each line has its strengths and weaknesses in terms of detectability and astrophysical and cosmological utility.

1.2.1 21 cm Intensity Mapping

As shown in Figure 1.1, 21 cm intensity mapping is a powerful probe of the Universe across an immense cosmic timeline, from deep into the Dark Ages ($z \sim 200$) all the way to the present day. This technique relies on the hyperfine transition line of neutral hydrogen to trace the large-scale distribution of hydrogen throughout cosmic history. One of its most remarkable advantages is the ability to probe the neutral IGM, providing a rare window into both the Dark Ages, an era before the formation of galaxies, and Cosmic Dawn and the EoR, allowing us to study early galaxy formation and reionization as it unfolds in real time.

Hydrogen emission does not occur in isolation but rather on the cosmological stage. Hydrogen atoms are backlit by the CMB leading to the observed 21 cm signal appearing in either emission or absorption. Whether the signal is detected as an emission or absorption feature (or remains altogether undetectable) depends on the temperature contrast between the CMB temperature, T_γ, and the hydrogen spin temperature, T_s, which is defined as

$$\frac{n_1}{n_0} = \frac{g_1}{g_0} \exp\left(\frac{h\nu_{21}}{kT_s}\right) \tag{1.1}$$

where n_1 is the number density of hydrogen atoms in the excited hyperfine state, n_0 is the number density of hydrogen atoms in the ground hyperfine state, $g_1 = 1$ and $g_0 = 3$ are the statistical weights, h is Planck's constant, k is Boltzmann's constant, and $\nu_{21} = 1420.406$ MHz is the rest frequency of the 21 cm line. The spin temperature quantifies the relative number of hydrogen atoms in the excited and ground hyperfine states. The key observable is the differential brightness temperature, which measures the temperature contrast between the spin temperature and the CMB temperature, and is given by

$$T_b(\mathbf{r}, z) = 27\mathrm{mK}\left(\frac{1+z}{10}\right)^{1/2}\frac{H(z)/(1+z)}{dv_{||}/dr_{||}}\left(1-\frac{T_\gamma(z)}{T_s(\mathbf{r}, z)}\right)x_{\mathrm{HI}}(\mathbf{r}, z)(1+\delta_m(\mathbf{r}, z)). \tag{1.2}$$

where $\mathbf{r}$ is the position vector perpendicular to the line of sight, $H(z)$ is the Hubble parameter, $v_{||}$ is the LOS velocity, x_{HI} is the neutral hydrogen fraction, and δ_m is the matter overdensity field [1](Furlanetto et al., 2006).

The spatially fluctuating 21 cm signal, a simulation of which is shown in the top panel of Figure 1.3, exhibits complex morphology, shaped by both cosmological and astrophysical processes. These fluctuations can be statistically characterized by the power spectrum, $P(k)$, defined as

$$\langle\tilde{T}(\mathbf{k})\tilde{T}^*(\mathbf{k}')\rangle = (2\pi)^3\delta_D^{(3)}(\mathbf{k}-\mathbf{k}')P(k) \tag{1.3}$$

where $\tilde{T}$ is the three-dimensional Fourier transform of T, $\mathbf{k}$ is the comoving wavenumber dual to $\mathbf{r}$, and $\delta_D^{(3)}$ is the three-dimensional Dirac delta function. It is often convenient to also define the dimensionless[2] power spectrum

$$\Delta^2(k) = \frac{k^3}{2\pi^2}P(k). \tag{1.4}$$

The spectrum quantifies the variance of the field T as a function of comoving length scale. Measuring the power spectrum has been the primary focus of current-generation 21 cm experiments across a wide range of redshifts, from low ($0 < z < 5$) to high ($5 < z < 20$) to ultra-high ($20 < z < 200$). We now provide a brief overview of the key cosmic events that drive

---

[1] This is in fact the baryon density field although due to the large scales involved, many codes model the baryon density field as tracing the matter density field

[2] The dimensionless power spectrum is only truly dimensionless if the original field is dimensionless which is not the case for intensity mapping. In this context, the dimensionless power spectrum often has temperature squared (or temperature squared equivalent) units.

the variations in brightness temperature seen in Figure 1.3 while also commenting on the current observational landscape.

By the time of recombination, the Universe expanded and cooled sufficiently for the first atoms to form, the overwhelming majority of which were hydrogen atoms. Still, there remained a small fraction of free electrons which Compton scatter off of CMB photons. These electrons would then collide with hydrogen atoms resulting in the tight coupling between the spin temperature and the CMB temperature. Due to this coupling, which persisted until $z \sim 200$, no hydrogen signal is observed (Pritchard and Loeb, 2012). As the Universe continues to expand and the density of matter decreases, CMB photons cannot find electrons sufficiently frequently, and thus the CMB decouples from the spin temperature at $z \sim 200$. Now, the baryon temperature is able to cool much faster than the photons. Due to collisional coupling between the kinetic gas temperature and the spin temperature, a hydrogen absorption signal is visible from $30 < z < 80$. This can be seen on the right side of Figure 1.3 in both the lightcone (where the negative brightness temperature is shown in deep red) as well as in the middle panel plotting the sky-averaged brightness temperature, $\bar{T}_b$, also known as the global signal. During these cosmic Dark Ages, no astrophysical structures have yet formed and the hydrogen traces the dark matter distribution, meaning a measurement of the 21 cm power spectrum is a good proxy for the matter power spectrum, $P_m(k) \propto \langle \tilde{\delta}_m(\mathbf{k}) \tilde{\delta}^*_m(\mathbf{k}) \rangle$, where $\tilde{\delta}_m$ denotes the Fourier transform of the matter density field.

From a theoretical standpoint, measurements of this epoch are highly desirable since more Fourier modes are expected to be in the linear regime meaning each Fourier mode evolves independently, greatly simplifying the modeling (Tegmark and Zaldarriaga, 2009; Scott and Rees, 1990). This clean probe of the matter field would allow for the study of primordial non-Gaussianity, of the small scale properties of dark matter, and would improve constraints on the neutrino mass and the amplitude of primordial fluctuations (Muñoz et al., 2015; Mao et al., 2008; Loeb and Zaldarriaga, 2004).

Despite being a period full of promise, at the ultra-low radio frequencies needed to access these redshifts, measurements quickly become arduous. The ionosphere becomes increasingly opaque

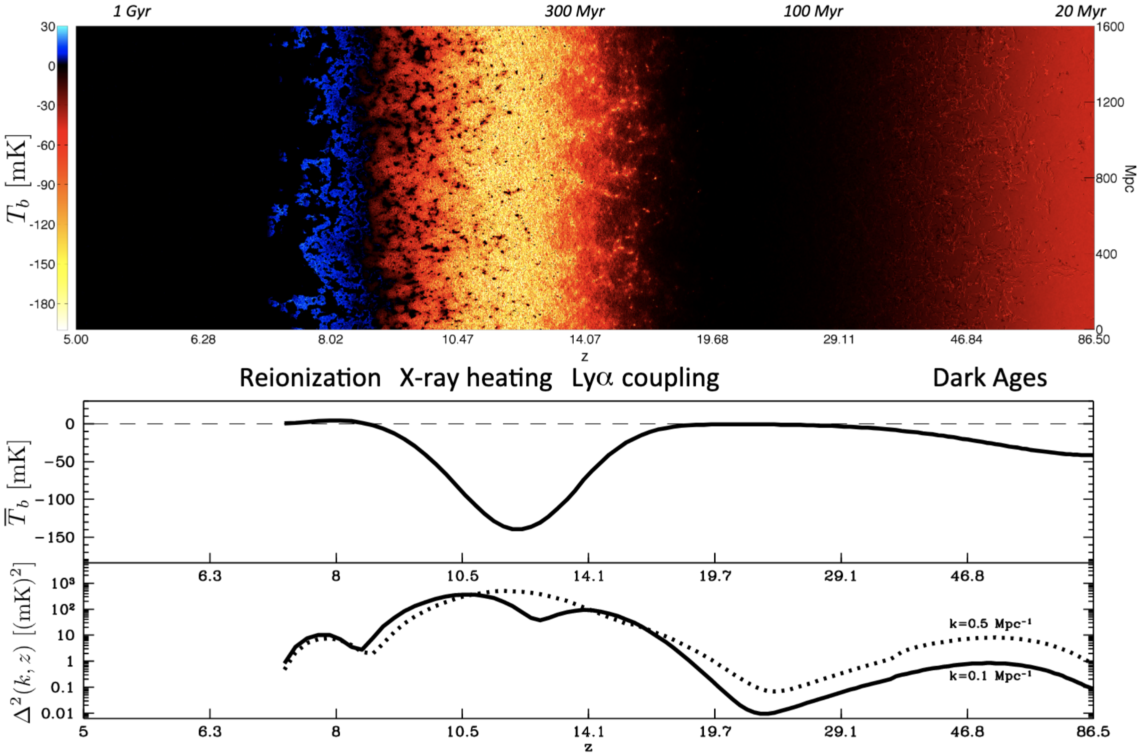

Figure 1.3: A simulated 21 cm signal spanning $7 < z < 90$. The top panel presents a 2D slice of the 3D lightcone of the 21 cm brightness temperature field, where the x-axis tracks cosmic evolution across different redshifts, and the corresponding age of the Universe is shown on the top horizontal axis. The middle panel displays the global 21 cm signal as a function of redshift, computed by averaging the brightness temperature over each slice of the lightcone at a given redshift. The bottom panel shows the dimensionless power spectrum as a function of redshift, with the modes $k = 0.5$ Mpc^{-1}(dotted) and $k = 0.1$ Mpc^{-1} (solid) highlighted. This figure is reproduced from Liu and Shaw (2020).

to frequencies below ∼ 30 MHz and can cause distortions at frequencies above but still close to this limit. While there are efforts to make measurements down to ∼ 1 MHz from the ground with instruments like Array of Long Baseline Antennas for Taking Radio Observations from the Sub-antarctic/Seventy-ninth parallel (ALBATROS, Chiang et al. 2020), many are looking to the sky for new opportunities. Space-based experiments, either in Earth orbit, in lunar orbit, on the far side of the moon aim to overcome these limitations (Burns et al., 2019).

By $z \sim 20$ it is expected that the first luminous objects would have formed, ushering in Cosmic Dawn. By the latter part of the Dark Ages, the Universe had expanded and diluted to the point where collisional coupling between electrons and hydrogen was ineffective, causing the spin temperature to decouple from the gas temperature. At this stage, the only remaining mechanism is the direct absorption or emission of 21 cm by CMB photons, causing the 21 cm signal to go away and therefore the brightness temperature to go to zero, as seen in Figure 1.3. Subsequently, the first stars produced Lyα photons traveled through the IGM and interacted with hydrogen atoms, once again coupling the gas temperature to the spin temperature. Since the gas has continued to cool, the spin temperature also experiences rapid cooling leading to the characteristic Cosmic Dawn absorption trough seen at $z \sim 12$ in the middle panel of Figure 1.3. As star formation becomes more efficient, higher energy photons, namely X-rays, propagate into the IGM, heating the neutral hydrogen. This phase of X-ray heating increases the gas temperature and consequently brings the brightness temperature back into the emission regime. While the precise order and timing of heating events is still unknown, upper limits on the spatially fluctuating 21 cm signal from the Hydrogen Epoch of Reionization Array (HERA, DeBoer et al. 2017) have placed the first constraint on the X-ray luminosity per star formation rate and have found that early galaxies are much more efficient at producing X-rays than their late-universe counterparts (Abdurashidova et al., 2022).

Following X-ray heating, the Universe undergoes reionization and transitions from mostly neutral to mostly ionized. Ultraviolet photons emanating from the first generation of galaxies ionize the surrounding neutral hydrogen producing rich bubble structure with complex morphology. Studying the 21 cm power spectrum during this epoch provides constraints on the timing of reionization and the properties of the ionizing sources (Furlanetto et al., 2006; Loeb and Zaldarriaga, 2004; Madau et al., 1998; Furlanetto et al., 2004; Barkana and Loeb, 2005; Furlanetto et al., 2008). While this epoch has yet to be detected with 21 cm, a host of experiments have placed ever tighter upper limits on the power spectrum, as summarized in Figure 1.4.

In the post-reionization Universe, although the IGM no longer contains a significant density of neutral hydrogen, hydrogen that remained self-shielded in galaxies can still be targeted and used as a tracer of large-scale structure. On this front, the field of hydrogen intensity mapping has seen tremendous experimental progress, yielding multiple detection of the 21 cm line cross-correlation with large-scale structure surveys (Masui et al., 2013; Teng et al., 2013; Rhee et al., 2016; Anderson et al., 2018; Chowdhury et al., 2021; Wolz et al., 2022; CHIME Collaboration et al., 2022; Amiri et al., 2024) as well as a tentative auto-spectrum measurement (Paul et al., 2023). These low-redshift LIM measurements play a vital role, as they can be compared with well-established methods, such as galaxy surveys, to validate LIM as a viable and reliable technique.

1.2.2 Intensity Mapping of Other Lines

In many ways, the intensity mapping of other lines (and of low-z 21 cm emission) can be grouped into a single category. While the physical mechanisms driving their emission vary, these lines are generally used as tracers of galactic populations, providing valuable insights into the distribution and evolution of cosmic structures and are also touted as sensitive astrophysical probes.

Once again, the power spectrum serves as the primary statistical tool for LIM, though alternative approaches exist (see Breysse et al. (2017) on one-point statistics). In the regime where line emission originates from galaxies smaller than the angular resolution of one's instrument and the line width is narrower than a single frequency channel, a simple halo model (Cooray and Sheth, 2002) can be applied, allowing the line power to be expressed as

$$P(k,z) = \bar{T}^2(z)\bar{b}^2(z)P_m(k,z) + P_{\text{shot}}(k,z) \tag{1.5}$$

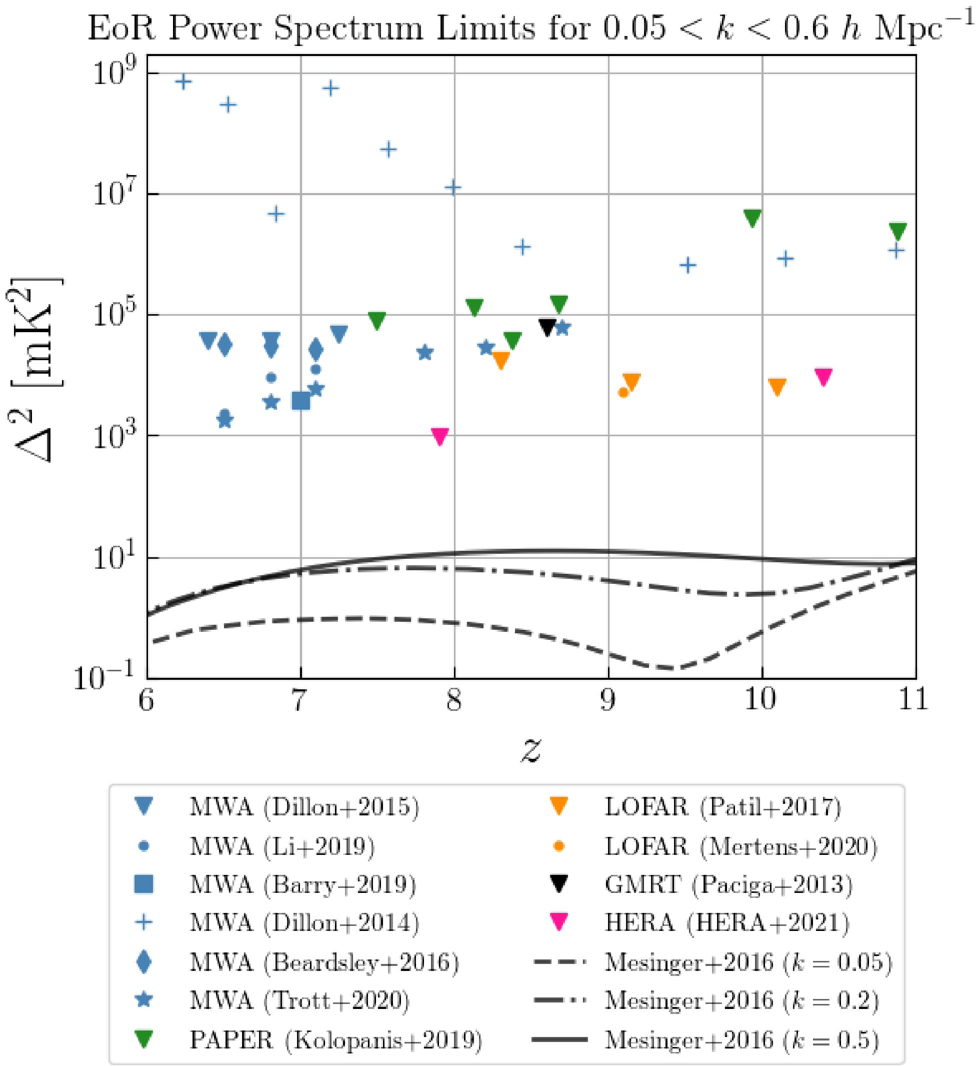

Figure 1.4: Upper limits and theoretical predictions for the 21 cm power spectrum during the EoR. Plot generated from `eor_limits`.

where $\bar{T}$ is the mean brightness temperature of the line, $\bar{b}$ is the luminosity weighted galaxy bias, $P_m(k,z)$ is the linear matter power spectrum of the underlying (and mostly dark) matter distribution, and $P_{\text{shot}}(k,z)$ is the galaxy shot noise power.

In order to build Eq. 1.5 we need only a few simple ingredients. First is the matter power spectrum, $P_m(k,z)$, which, as briefly mentioned in the previous section, describes the variance in density fluctuations as a function of comoving length scale. Assuming some fiducial cosmology, the matter power spectrum can be easily obtained from a Boltzmann code such as CAMB (Lewis and Challinor, 2011). Next, assuming some luminosity function for the source galaxies, $\Phi(L,z)$, we can write down the mean intensity of the line

$$\bar{I}(z) = \frac{c}{4\pi\nu_{\text{emit}}H(z)}\int L\Phi(L,z)dL \tag{1.6}$$

where ν_{emit} is the rest frequency of the line and L is the line luminosity. The intensity can be easily converted to the mean brightness temperature which, in the Rayleigh-Jeans limit, yields

$$\bar{T}_b(z) = \frac{c^3(1+z)^2}{8\pi\nu_{\text{emit}}^3 k_B H(z)}\int L(M)\frac{dn}{dM}dM. \tag{1.7}$$

Here we have taken the extra leap in assuming that the luminosity of a galaxy is a function of its mass, replacing $\Phi(L,z)$ with the mass function dn/dM.

The shot noise power represents the contribution to the power spectrum in the absence of clustering and is proportional to the mean square of galaxy luminosities. In brightness temperature units it is given by

$$P_{\text{shot}}(k,z) = \left(\frac{c^3(1+z)^2}{8\pi\nu_{\text{emit}}^3 k_B H(z)}\right)^2\int L^2\Phi(L,z)dL. \tag{1.8}$$

The final piece is the luminosity averaged galaxy bias which represents the observed clustering of galaxies relative to the clustering of the underlying dark matter

$$\bar{b} = \frac{\int Lb(L)\Phi(L)dL}{\int L\Phi(L)dL}. \tag{1.9}$$

Combining the matter power spectrum with Eqs. 1.7-1.9, we obtain the total line power.

Measurement of this line power has been the goal of many LIM experiments across the electromagnetic spectrum. Experimental progress in intensity mapping has largely been driven by efforts to detect CO emission, which scales directly with the star formation rate. Since CO is one of the primary coolants that enables star formation, CO is an excellent tracer of the molecular gas from which stars form. From an astrophysical perspective, detecting the CO intensity mapping signal would provide crucial insights into the star formation process across cosmic time. Given that star formation peaks around $z \sim 2-3$ (an epoch known as cosmic noon) many CO mapping experiments specifically target this redshift range. At these redshifts, CO emission redshifts into millimeter wavelengths, allowing technology originally developed for CMB observations to be adapted for intensity mapping, accelerating progress in the field. Some studies have even leveraged CMB datasets themselves to search for CO, which is typically treated as a contaminant in CMB analyses (Pullen et al., 2013, 2018; Switzer, 2016; Roy et al., 2024). Currently, the CO shot noise power has been detected by the CO Power Spectrum Survey (COPSS, Keating et al. 2015, 2016) at $z = 2.3 - 3.3$, while the CO Mapping Array Pathfinder (COMAP) has placed stringent upper limits on the CO clustering power spectrum, further constraining models of molecular gas evolution around cosmic noon (Stutzer et al., 2024; Chung et al., 2024; Ihle et al., 2022).

In addition to CO, another line of interest that is observed at millimeter-wavelengths is the 158 μm [CII] line. As the dominant cooling mechanism of interstellar gas, [CII] serves as a powerful probe of both ionized and neutral media in star-forming galaxies. At high redshifts, [CII] is the brightest of the carbon lines, roughly six times brighter than the next strongest CO line at $z = 7$ (Bernal and Kovetz, 2022). Several upcoming instruments including the Experiment for Cryogenic

Large-Aperture Intensity Mapping (EXCLAIM, Oxholm et al. (2020), the Tomographic Intensity Mapping Experiment (TIME, Crites et al. 2014), the Terahertz Intensity Mapper (TIM, Marrone et al. 2022), CONCERTO (Concerto Collaboration et al., 2020), and Prime-Cam on the Fred Young Sub-millimeter Telescope (Aravena et al., 2019), are set to map [CII] in the coming years. Another major target for intensity mapping is the Lyα line, which is being studied from cosmic noon out to reionization. At these redshifts, the rest-frame UV Lyα line redshifts into the near-infrared, making it the primary intensity mapping target for the deep-field survey by the Spectro-Photometer for the History of the Universe, Epoch of Reionization, and Ices Explorer (SPHEREx, Crill et al. 2020). Scheduled for launch in early 2025, SPHEREx will also have sensitivity to the Hα and ionized oxygen (OII) lines, providing valuable insights into galaxy emission spectra (Fonseca et al., 2017).

1.3 The Cosmic Microwave Background

While the true beginning of our Universe remains unknown, reasonable speculations such as the inflationary paradigm exist and what is on extremely firm observational footing is knowledge of our cosmic history back to $t \sim 10^{-2}$ seconds. By this time, the Universe began to cool adiabatically. Within fractions of a second, the Universe had cooled sufficiently for protons and neutrons to form, soon followed by Big Bang nucleosynthesis (BBN), which produced light atomic nuclei: ^{1}H, ^{3}He, ^{4}He, ^{7}Li. From these first few minutes until roughly 400,000 years later, the Universe remained hot and ionized, consisting of a photon-baryon fluid and an underlying distribution of dark matter seeded by the initial quantum fluctuations present at the time of the Big Bang. In this primordial plasma, photons underwent Thomson and Compton scattering off charged particles, and the resulting radiation pressure prevented the ionized plasma from clustering under gravity. This pressure led to the formation of relativistic sound waves known as baryon acoustic oscillations (BAO). As the Universe expanded, it cooled, eventually reaching a temperature where protons and electrons could combine to form neutral hydrogen in a process known as recombination. With the formation of neutral atoms, free electrons became scarce, dramatically increasing the mean free path of photons. This allowed photons to escape, effectively decoupling from the baryonic matter. As a result, radiation pressure ceased to support the photon-baryon fluid, causing the acoustic oscillations within it to freeze into the baryon distribution. Additionally, since the photons last scattered off the baryons at this moment, these patterns imprinted in the surface of last scattering became observable. These photons have since traveled largely unimpeded, apart from secondary effects that will be discussed later, and today these same photons are observed as the cosmic microwave background.

The energy distribution of CMB photons follows a blackbody spectrum[3]. At the time of recombination, the peak of this spectrum was in the infrared, corresponding to an effective temperature of approximately 3000 K. Due to cosmic expansion, this spectrum has been redshifted, and today, CMB photons are observed to peak at microwave wavelengths with a blackbody temperature of 2.726 K +/- 0.010 K with these error bars at 95% confidence level (as shown in Figure 1.5, Mather et al. 1994; Fixsen et al. 1996; Fixsen 2009).

At early times, the Universe was remarkably homogeneous and isotropic, a fact corroborated by a crude measurement of the CMB. However, a more careful examination reveals slight temperature fluctuations at the level of 1 part in 10^5. These anisotropies arise from primordial quantum fluctuations which ultimately shape large-scale structure formation. Observational efforts have focused on making precise measurements of the statistical properties of the CMB. Figure 1.6 illustrates the steady progress in mapping these anisotropies, from their initial detection by the Cosmic Background Explorer (COBE, Janssen and Gulkis 1992) to the higher-resolution observations of the Wilkinson Microwave Anisotropy Probe (WMAP, Bennett et al. 2003) and the *Planck* satellite (Planck Collaboration et al., 2014).

The temperature fluctuations in the CMB are expected to follow a Gaussian distribution, meaning they can be statistically characterized by just two quantities: the mean and the variance.[4] The mean

---

[3] A vast body of literature explores spectral distortions in the CMB arising from both early-time (pre-recombination) and late-time (post-recombination) processes. For a review, see Chluba and Sunyaev (2011) and references therein.

[4] However, certain early Universe models predict primordial non-Gaussianity and searching for its signatures remains a key objective of current-generation CMB experiments Bartolo et al. (2010).

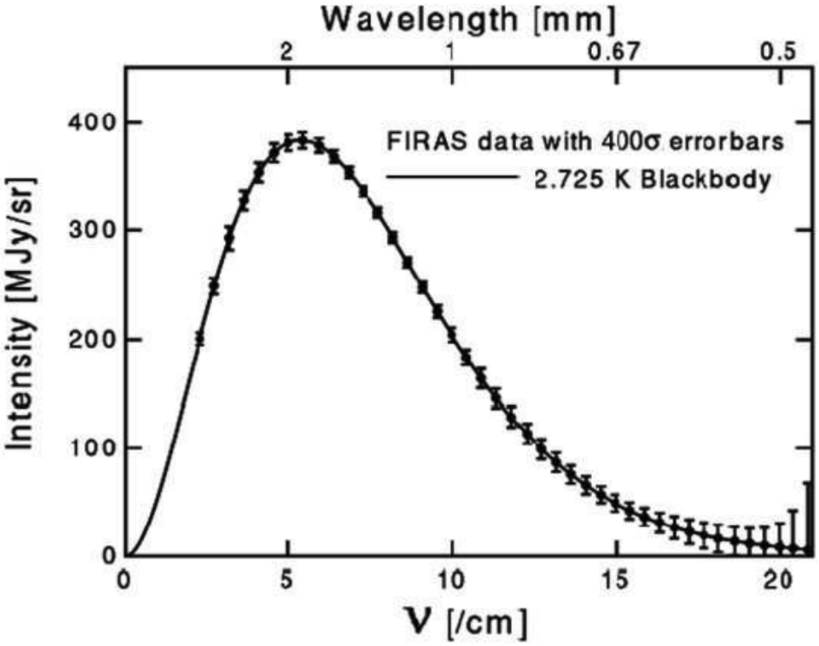

Figure 1.5: The CMB spectrum measured by the Far InfraRed Absolute Spectrophotometer (FIRAS) on the Cosmic Background Explorer (COBE) satellite. Data points with 400σ error bars are shown with black markers. The theory curve for a 2.725 Planck distribution is shown in black. Figure from NASA/FIRAS Science Team.

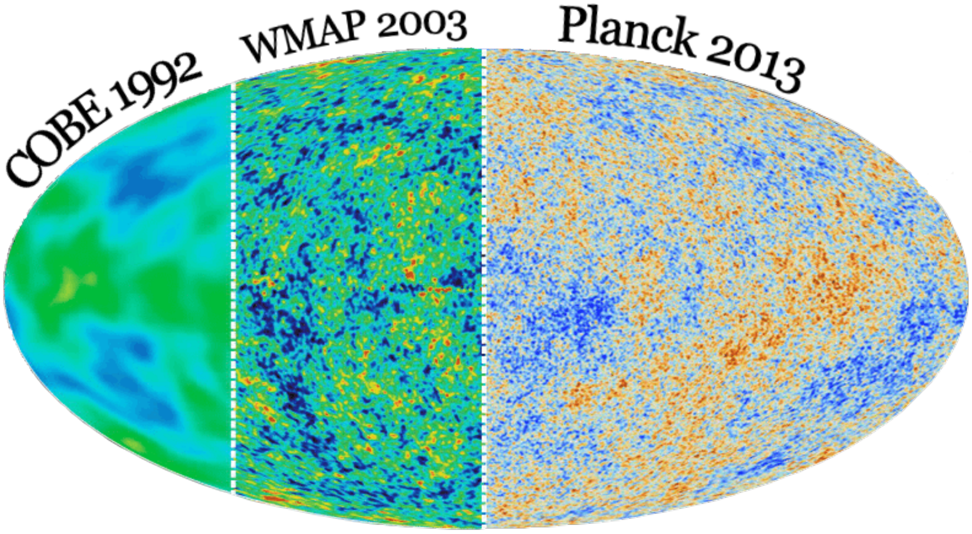

Figure 1.6: CMB temperature map as measured by COBE, WMAP, and *Planck*. As instrumentation improved, smaller angular scales were resolved. Figure adapted from Vazquez Gonzalez et al. (2020).

temperature of the CMB, as we have seen, is deduced from its blackbody spectrum while the variance is encapsulated in the angular power spectrum. The all-sky temperature map, $T(\hat{\mathbf{n}})$, can be decomposed into spherical harmonic modes, given by

$$T(\hat{\mathbf{n}}) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}) \tag{1.10}$$

where $Y_{\ell m}(\hat{\mathbf{n}})$ is the spherical harmonic function of degree ℓ and order m and the expansion coefficients, $a_{\ell m}$, are given by

$$a_{\ell m} = \int_{4\pi} T(\hat{\mathbf{n}}) Y^*_{\ell m}(\hat{\mathbf{n}}) d\Omega. \tag{1.11}$$

The angular power spectrum, C_ℓ, is then defined as

$$\langle a_{\ell m} a^*_{\ell' m'} \rangle = \delta_{\ell\ell'} \delta_{mm'} C_\ell \tag{1.12}$$

where δ is the Dirac delta function. In practice, the estimated power spectrum as a function of multipole, $\ell \sim 180^\circ/\theta$, typically denoted with a hat, is given by the square of the spherical harmonic amplitudes

$$\hat{C}_\ell = \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell} |a_{\ell m}|^2. \tag{1.13}$$

Power at low ℓ corresponds to large-scale fluctuations, while power at high ℓ traces smaller-scale variations. The resulting CMB temperature power spectrum, shown in Figure 1.7, exhibits distinct oscillatory features: baryon acoustic oscillations (BAO). These wiggles in harmonic space correspond to a characteristic scale in configuration space, the BAO scale, which, as the reader may recall, arose from the frozen sound waves sharply imprinted on the CMB at recombination. The angular BAO scale, θ_*, can be compared to the comoving BAO scale, r_d, calculated from linear theory [5] thus acting as a standard ruler. The measurement of BAOs as standard rulers have become a cornerstone of cosmology, performed across multiple redshifts to probe the geometry and expansion history of the Universe (Eisenstein et al., 2005). We will return to a more pointed discussion of BAO measurements in Chapter 5.

Perhaps most remarkable of all is the striking agreement between the red data points and the green theory curve, a testament to the predictive power of our cosmological models and the precision of CMB measurements. Strikingly, the primary signal (again, that is the signal from the surface of last scattering) of the CMB temperature power spectrum can be accurately parameterized using just six independent parameters. These include, A_s, the amplitude of primordial fluctuations; n_s the slope of the primordial power spectrum; Ω_b the relative abundance of baryons; Ω_c the relative abundance of dark matter; τ, the optical depth to reionization; and θ_* the acoustic scale angle. Changing the value of any one of these parameters will affect the amplitude, location, and relative heights of the acoustic peaks.

While alternative parameterizations exist and are used in different contexts—for example, re-

---

[5] In reality, there are non-linear contributions. See Padmanabhan et al. (2009) for an example of how these are treated.

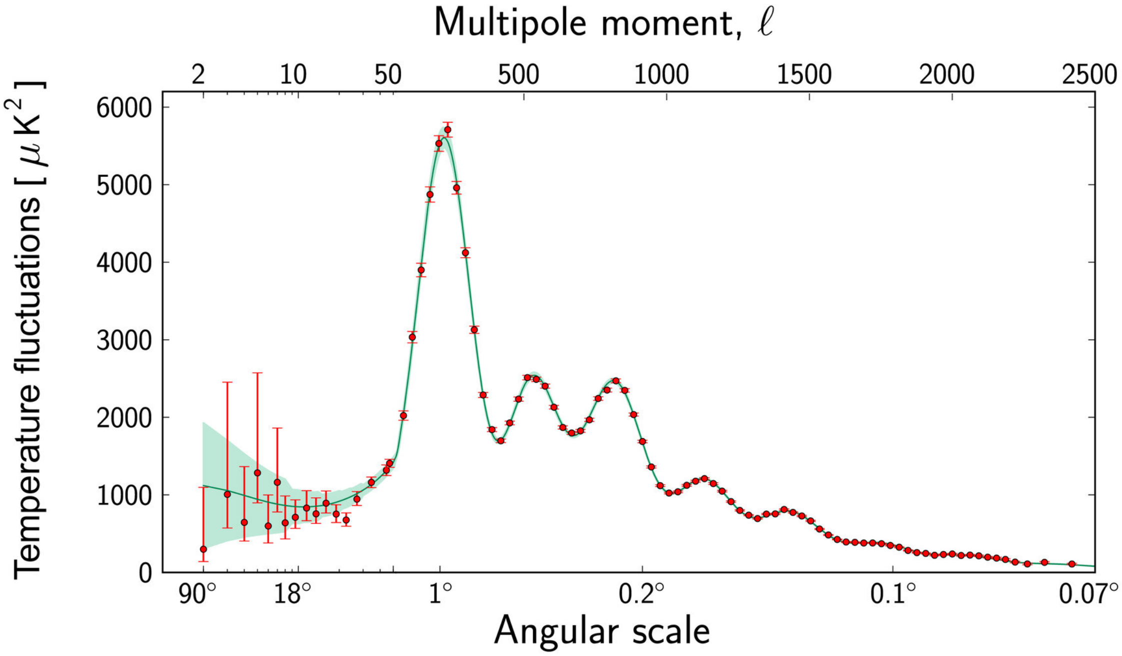

Figure 1.7: The CMB temperature angular power spectrum with *Planck* data shown in red. Fiducial ΛCDM theory curve is plotted in green showing clear agreement with the data. Figure from ESA and the Planck Collaboration.

expressing the parameters in terms of derived quantities such as σ_8 or the Hubble constant, the fundamental point remains that only six independent parameters are needed to describe the primary CMB anisotropies. However, as CMB measurements continue to improve in precision, additional effects that were once negligible are becoming increasingly important. In particular, constraints on the sum of neutrino masses are reaching the point where it is becoming essential to include this quantity in cosmological fits, marking a shift in the standard cosmological parameterization.

While the discussion thus far has focused on temperature anisotropies, it is important to briefly mention that CMB photons are also polarized. This polarization is categorized into E-modes (the curl-free component), and B-modes (the divergence-free component). Considerable effort has

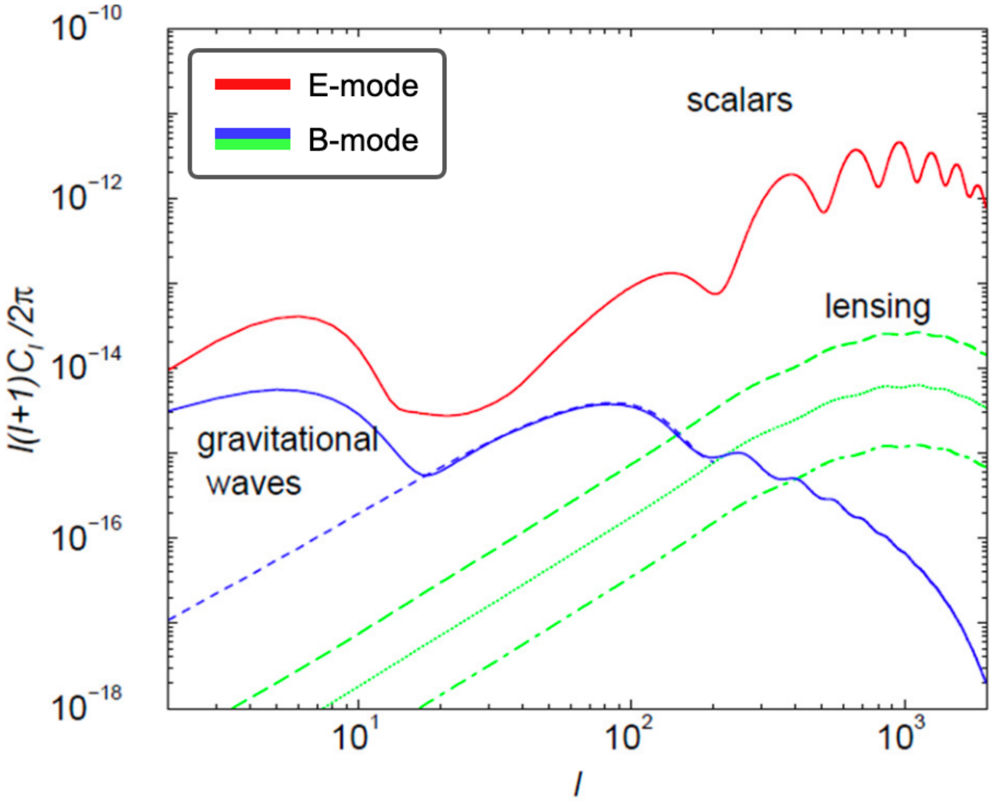

Figure 1.8: Theoretical predictions for the E-mode spectrum (red), the primordial B-mode spectrum (blue), and the lensing B-mode spectrum (green). Figure adapted from Lewis and Challinor (2006).

been dedicated to measuring B-modes in particular, as a primordial gravitational wave background can source B-mode polarization, with the precise tilt of its spectrum being indicative of the early universe paradigm, whether inflationary or otherwise (Kamionkowski and Kovetz, 2016). This measurement is challenging, as gravitational lensing by the late-time Universe also produces B-modes, especially on small angular scales, as illustrated in Figure 1.8. In the following section, we expand on the lensing signals imparted on the CMB and how they can be exploited for cosmology.

1.4 Weak Lensing by Large Scale Structure

While on large scales the Universe is statistically homogeneous, on small scales the Universe consists of dense objects, galaxies and clusters, separated by vast cosmic voids. The bending of light by a locally inhomogeneous gravitation field is known as gravitational lensing. First predicted by Henry Cavendish in 1784 but subsequently given a sound mathematical foundation by Einstein in 1919, lensing occurs when light from a distant source object is bent by the gravitation potentials of intervening massive objects. This effect was first detected in 1919 by Arthur Eddington and Frank Watson Dyson when the Sun was obstructed during a solar eclipse and the light of distant stars could be observed to be distorted by the Sun's gravity. Decades later in 1979, the first gravitationally lensed quasar, known as the Twin Quasar, was observed. This quasar was lensed by a massive elliptical galaxy, producing a double image. This observation was the first to identify both a source and a lens of cosmological origin, exemplifying how lensing can help to reveal both the properties distant objects and help us study the intervening matter distribution of our Universe.

Broadly speaking, lensing comes in two flavours: strong and weak lensing. Strong lensing, as the name would suggest, occurs when the gravitational potential is strong enough to produce visible distortions to a single source object. For point-like sources, multiple images are produced while extended sources may appear as arcs or rings (see Figure 1.9). Weak lensing, on the other hand, can only be detected *statistically*, that is, it is impossible to observe the distortion of a single weakly lensed background source. Weak lensing is a powerful cosmological probe as it probes the statistical properties of the intervening mass density without the need to characterize the composition or dynamics of each object in the population of lenses. Of all the weak lensing source planes, the CMB is perhaps the most powerful. CMB photons, having been emitted at just 400,000 years after the Big Bang, travel through the vast majority of our cosmic volume, picking up lensing signals from every epoch of our Universe's evolution. Therefore, encoded in each CMB photon is information about the entire Universe's history.


Multiple Images

Einstein Ring

Figure 1.9: In the left panel is an artist's conception of the gravitational lensing of a distant quasar by an intervening galaxy. The quasar is observed as multiple images (Image credit NASA, ESA, and D. Player (STScI)). In the right panel is an image taken by Hubble's Wide Field Camera 3 of the blue horseshoe galaxy at $z = 2.4$ lensed by a Luminous Red Galaxy in the foreground. The blue horseshoe galaxy is distorted by the lens into a ring (Image credit NASA, ESA).

Weak gravitational lensing of the CMB arises when photons from the surface of last scattering get deflected by the gravitational potentials they encounter on their journey to us. The deflection angle, α, is proportional to the gradient of the lensing potential, ϕ, the total potential of the projected mass distribution along the line of sight. Using CMB temperature and polarization maps to reconstruct ϕ gives us direct observation of the total matter distribution of the universe, both baryonic and dark, without the use of a biased tracer (Seljak and Zaldarriaga, 1999). Measuring the power spectrum of the lensing potential, either in auto- or in cross-correlation with large scale structure surveys, enables us to probe the growth of matter fluctuations, place limits on primordial non-Gaussianity, constrain the sum of the neutrino masses, and even test theories of modified gravity (Lewis and Challinor, 2007; Schmittfull and Seljak, 2018; Allison et al., 2015).

As a result of the interaction between the CMB photons and later time cosmic structure, several unique effects are induced on the CMB anisotropies. Unlike rings and multiple images, which are characteristic features of strong lensing, perhaps the most notable feature of weak lensing is that is it can only be detected statistically through the measurement of either the lensing shear or the lensing convergence (Dodelson, 2017). Shear, denoted by γ occurs when the background image is stretched from its original shape as a result of interacting with the intervening gravitational potential, as illustrated in Figure 1.10. The amount of shear is related to the potential by

$$\gamma_+ = \frac{1}{2c^2}\left(\frac{\partial^2\phi}{\partial\theta_x^2} - \frac{\partial^2\phi}{\partial\theta_y^2}\right) \tag{1.14}$$

$$\gamma_\times = \frac{1}{c^2}\frac{\partial^2\phi}{\partial\theta_x\partial\theta_y} \tag{1.15}$$

where + and × denote the two shear polarizations and θ_x and θ_y denote angular coordinates on the sky. Precision cosmic shear measurements have been made not with the CMB, but instead with galaxy lensing. By measuring the ellipticity of a large population of galaxies, it is possible to make inferences about the intervening mass distribution. This of course comes with challenges. The source galaxies are not intrinsically round but rather elliptical disks, and even if they were round face-on, they may still appear elliptical unlensed depending on their projection on the sky. The uncertainty related to not knowing how much of the observed ellipticity is intrinsic and how much is a result of lensing is referred to as shape noise. To eliminate this source of uncertainty, it may be noted that the ellipticities of galaxies are random (that is with respect to one another) and therefore it is possible to beat down the shape noise by averaging together the random intrinsic ellipticities, thus teasing out their common ellipticity due to shear. However, galaxy shapes and galaxy spins are correlated with the underlying dark matter distribution giving rise to intrinsic alignment where the ellipticities of the whole populations are aligned.

Convergence, also depicted in Figure 1.10, is a measure of magnification. Denoted by κ, the convergence is simply the integrated matter density along the line of sight weighted by the lensing kernel, $W(\chi, \chi_s)$, and is given by

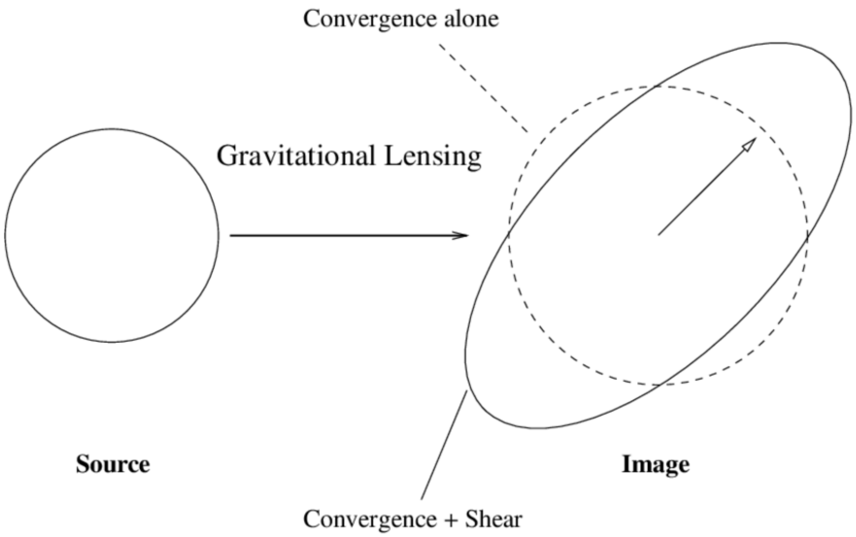

Figure 1.10: Here the source undergoes weak gravitational lensing which both magnifies the source (convergence) and extends the source tangentially to the line of sight (shear). Figure from Umetsu (2010).

$$\kappa(\hat{\mathbf{n}}) = \int_0^{\chi_s} W(\chi, \chi_s)\delta_m(\chi(z)\hat{\mathbf{n}}, z)d\chi \tag{1.16}$$

where χ is the comoving distance, χ_s is the comoving distance to the source, and δ_m is the matter overdensity field. The lensing kernel characterizes which portion of the line of sight matter density contributes most to the magnification effect. Like shear, the convergence is also related to the projected gravitational potential, $\kappa \propto \nabla^2\phi$, although the beauty in convergence is that it is related precisely to the LOS projection of underlying matter density by the Poisson equation, $\nabla^2\phi = 4\pi G \int \rho(\mathbf{r}, z)dz$.

By studying the statistical artifacts on T, E, and B maps of the CMB due to weak lensing, it is possible to infer κ and in turn, the total intervening matter distribution of our Universe (Lewis and Challinor, 2007). Devoid of lensing, if one were to take different patches of the CMB and compute their power spectra, all the spectra should agree within cosmic variance. Lensing, however, distorts these spectra such that they are slightly different on different patches of the sky. Again, it is not individual weak lensing features that can be detected in the CMB, but rather it is the *the statistics* of the CMB that are distorted by weak lensing. The deflections induced on CMB photons as they travel to us are small, on the order of arcminutes ($\ell \sim 1000$) . However, the structures responsible for the deflection are large, on the order of degrees ($\ell \sim 50$). Therefore, somewhat counter intuitively, to study the large scale structure of the Universe one actually has to study the small scale anisotropies of the CMB.

There are two lensing effects that are expected to occur: peak smearing and the induction of non-Gaussianities. To first illustrate the effect of peak smearing, we can consider an individual distortion due to some potential ϕ on a small patch of the CMB. This distortion will shift the power spectrum to lower ℓ if κ is positive (magnifies), or shifts the peaks to higher ℓ if κ is negative (contracts). This should be intuitive since magnifying, for example, makes features appear larger and therefore peak at lower ℓ. This fractional change in the angular power spectrum due to κ is given by

$$\frac{\Delta C_\ell}{C_\ell} = \kappa \frac{d\ln(\ell^2 C_\ell)}{d\ln \ell}. \tag{1.17}$$

Averaged over the sky, however, this shifting effect cancels out since there are just as many positive convergence regions as there are negative. Computing first and second order corrections to the CMB temperature power spectrum due to lensing shows that the resulting observed effect is a smoothing of the peaks and troughs. The amplitude of the smoothing is given by

$$\Delta C_\ell^{\rm smooth} = (\bar{C}_\ell - C_\ell) \int \frac{d^2\ell'}{(2\pi)^2} (\vec{\ell} \cdot \vec{\ell}')^2 \frac{C_{\ell'}^{\phi\phi}}{c^4} \tag{1.18}$$

and depends on the intervening gravitational potential, statistically characterized by its spectrum, $C_{\ell'}^{\phi\phi}$. Here, $\bar{C}_\ell$ is the spectrum smoothed over a region $\ell - \Delta\ell \rightarrow \ell + \Delta\ell$ where $\Delta\ell \simeq 50$ since that is the scale at which the lensing spectrum peaks, and C_ℓ is the unsmoothed spectrum.

Importantly, both the temperature and polarization anisotropies experience this effect. In fact, the smoothing of E-modes is more distinguishable than that of temperature since the E-mode spectrum exhibits more prominent peaks (Lewis and Challinor, 2007)[6]. In addition to smoothing, lensing produces unique polarization effects on CMB photons. E-modes, the curl-less component to polarization, give rise to divergence-less B-modes after lensing. In fact, most of the B-mode power is generated by lensing and therefore lives at high ℓ (~ 1000). The presence of lensing induced B-modes in the B-modes power spectrum is therefore the primary threat to the measurement of low-amplitude primordial B-modes. While primordial B-modes exhibit large scale correlations at low ℓ, even a small lensing contribution can wash out their signal. This is only made worse by other contaminants like polarization from dust and other systematics. Understanding the lensing B-mode signal therefore remains an important part of constraining early universe scenarios (BICEP/Keck Collaboration et al., 2022).

While peak smearing is an important lensing effect that is routinely taken into account when fitting for CMB power spectra, due to its degeneracy with the primary CMB signal, it is difficult to isolate this signature for lensing studies. Luckily, lensing also affects the higher order statistics of the CMB by producing correlations between Fourier modes. It is possible to perform lensing reconstruction from CMB maps with the use of quadratic estimators. It is assumed that the unlensed CMB is statistically isotropic and therefore the Fourier modes of this field are uncorrelated. Inhomogeneity is then induced by convergence which subsequently leads to correlations between Fourier modes and one can use this fact to learn about the inhomogeneity, that is, the lensing field [7]. Most famously, the Hu and Okamoto quadratic estimator and the Okamoto and Hu estimator have been used to preform lensing reconstruction (Hu and Okamoto, 2002; Okamoto and Hu, 2003). While these estimators have been the main workhorse for field reconstruction, it has been shown that they will be sub-optimal for upcoming CMB experiments (Hirata and Seljak, 2003a). Other techniques like gradient-inversion methods and likelihood based methods have been explored in addition to other quadratic estimators which yield a lower reconstruction noise (Hirata and Seljak,

---

[6] The same argument holds for the TE spectrum.

[7] The lensing field can be referred to by the convergence field, κ, the deflection field α, or the potential field ϕ.

2003a; Carron and Lewis, 2017; Hirata and Seljak, 2003b; Hadzhiyska et al., 2019; Millea et al., 2020; Maniyar et al., 2021).

To date, a number of lensing detections have been made, the first of which was by the Atacama Cosmology Telescope (ACT) in 2011 using the Hu and Okamoto estimator on the temperature maps (Das et al., 2011). Subsequent measurements have been made of the lensing signal in temperature as well as polarization maps by ACT, the South Pole Telescope (SPT), *Planck*, Background Imaging of Cosmic Extragalactic Polarization (BICEP), and the POLARization of the Background Radiation experiment (POLARBEAR) (van Engelen et al., 2012; Hanson et al., 2013; Sherwin et al., 2017; Planck Collaboration et al., 2016, 2020; Omori et al., 2017; Story et al., 2015; BICEP2 Collaboration et al., 2016; Wu et al., 2019; Ade et al., 2014; Adachi et al., 2020). Excitingly, current and next generation wide-field CMB experiments like SPT-3G, SPT-3G+, AdvACT, the Simons Observatory (SO) and CMB-Stage 4, will provide high signal-to-noise lensing measurement with unprecedented angular resolution (Ward and Advanced ACT Collaboration, 2017; Benson et al., 2014; Anderson et al., 2022; Ade et al., 2019; Abazajian et al., 2022).

While CMB lensing measurements have seen significant progress and are poised to become even more sensitive, they remain fundamentally limited. A key challenge is that the observed lensing signal represents an integrated projection along the entire line of sight, preventing one to separate contributions from different epochs thus hindering the study of their evolution. To mitigate this, many efforts have combined CMB lensing with external tracers such as galaxy positions or galaxy lensing (Hand et al., 2015; Liu and Hill, 2015; Singh et al., 2016; Chang et al., 2023; Kalaja et al., 2024; Shaikh et al., 2024). However, these approaches primarily isolate low-redshift lensing signals, and in the case of galaxy positions, they compromise the inherently unbiased nature of the lensing convergence. In Chapters 4 and 5 of this thesis, we introduce a novel method that leverages LIM lensing to construct a high-redshift CMB lensing observable that is tomographic. This observable is an unbiased probe of the matter density field and we show how it can be used to constrain the geometry of our universe through the measurement of lensing BAOs, and also to constrain cosmology beyond ΛCDM.

1.5 Observational Challenges and the Analyst's Solution

Thus far, we outlined the compelling potential of LIM in offering new observational windows into previously unexplored epochs across cosmic time with sensitivity to both cosmology and astrophysics. However, the relative scarcity of detections is no accident; LIM is inherently challenging. The distant line emission being targeted is often faint, while bright foreground contaminants dominate the observed signal. Moreover, each spectral line, depending on the observed wavelength, suffers from its own unique sources of foreground contamination which is only further exacerbated by the interaction between the instrument response, additional systematics, and on-sky contaminants.

In the case of 21 cm intensity mapping, the 21 cm line is heavily contaminated by bright foregrounds, including Galactic synchrotron emission, free-free emission, and both resolved and unresolved radio point sources. Together, these contaminants can dominate over the signal by 4–5 orders of magnitude. This is particularly severe for high-redshift 21 cm observations since the foregrounds, which can be approximated by a spatially-dependent power law distribution, increase in intensity at lower frequencies (i.e. higher redshifts). This foreground model is given by

$$T(\mathbf{r}, \nu) = T_0 \left(\frac{\nu}{\nu_0}\right)^{-\alpha(\mathbf{r})} . \tag{1.19}$$

Here, ν_0 is the reference frequency, T_0 is the mean sky temperature at ν_0, and α is the spectral index[8](Liu and Tegmark, 2012; Wang et al., 2006). This broadband frequency structure, characteristic of the foregrounds, results in a large correlation length along the frequency axis. This, in turn, leads to severe contamination of large-scale LOS modes, which are rich with cosmological information. To analyze these effects, the power spectrum is often decomposed into Fourier wavenumbers perpendicular ($k_\perp$) and parallel ($k_\parallel$) to the LOS, forming what is known as the cylindrically averaged power spectrum. As shown in Figure 1.11, this power spectrum reveals

---

[8] It should be noted that approximating the foregrounds by such a simple model generally breaks down for observations with deep integrations that are probing a very high (10^5) dynamic range

a prominent stripe of excess power at low $k_{\parallel}$, a direct result of bright foreground contamination. Additionally, the frequency-dependent response of the synthesized beam (also known as the point spread function) of radio interferometers introduces a feature known as the foreground wedge (also seen in Figure 1.11), which leaks power from large LOS modes into otherwise uncontaminated small-scale modes. This is a direct result of the fact that, at different frequencies, a given baseline (i.e. the correlation of a pair of antennas separated by some distance) measures different Fourier modes. This frequency evolution of the measured mode is baseline dependent, leading to the wedge-like structure (see Liu et al. (2014a) and Liu et al. (2014b) for more detail). While numerous efforts have been made to filter out foreground contamination, the complex interplay between the instrument response, foregrounds, and additional systematics has led many to adopt the more severe, yet conservative, approach of foreground excision. Unfortunately, this also removes the brightest signal modes which carry significant constraining power and leaves only a small window in the $k_{\perp}$-$k_{\parallel}$ plane to look for the EoR signal (while it may look like a large area of parameter space due to being plotted on a logarithmic scale, this is in fact a very small window). In Chapter 3, we will return to the problem of foreground removal and propose a new path forward.

Higher frequency LIM observations also suffer from broadband contaminants but are more severely impacted by line interloper contamination. Line interlopers are lower-redshift, lower-energy transition lines that redshift into the same observed band as the high-redshift line of interest, resulting in line confusion. For example, [CII] emission from the EoR is highly contaminated by a ladder of CO rotational transitions from $0 < z \lesssim 3$. While broadband contaminants remain a challenge, single-aperture instruments, commonly used at these shorter wavelengths, exhibit more slowly evolving spectral behavior compared to radio interferometers, reducing distortions in smooth spectral features and making broadband foreground removal more feasible. To provide concrete examples, the cosmic infrared background (CIB) is a broadband contaminant to [CII] intensity mapping and in the infrared, zodiacal light, that is, diffuse sunlight scattered by interplanetary dust, is a primary contaminant to Lyα intensity mapping. Fortunately, the CIB is expected to be

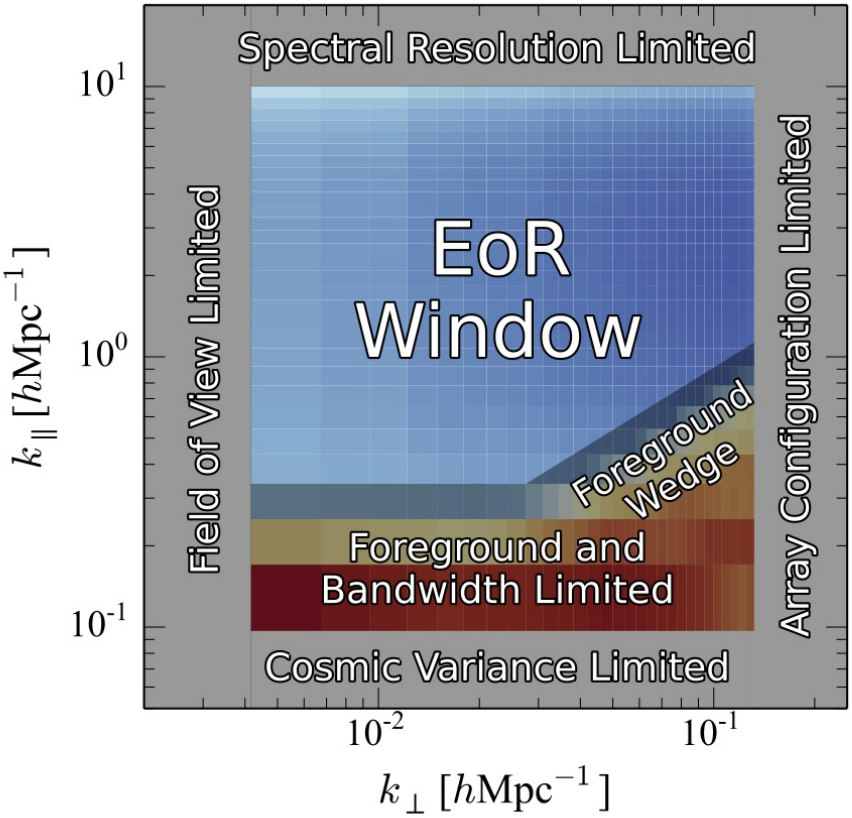

Figure 1.11: A schematic of the cylindrical power spectrum. The foreground contaminated region, including “the wedge”, can be seen in red and the “EoR window” in light blue. Figure from Liu et al. (2014c).

removable through traditional decomposition-based methods (Yue et al., 2015) and extensive efforts have been made to mitigate zodiacal light contamination in upcoming SPHEREx observations. These methods leverage multiple exposures to isolate the common cosmological signal while filtering out the zodiacal component which exhibits day-to-day variability (Williams et al., 2024). Despite these advancements, foreground removal remains a primary concern across all wavelengths. Achieving precision science with LIM requires equally precise foreground removal, and it is not yet immediately clear whether existing techniques can reach the necessary level of accuracy. Moreover, real data often present unexpected challenges, making it essential to develop a diverse portfolio of approaches to ensure robustness across different observational conditions.

So how do we resolve this problem? How do we peel back the foregrounds to reveal the cosmic signal beneath? How do we detect something faint amidst bright, unwanted contaminants? One of the most robust solutions is cross-correlation, a well-established and widely used technique in cosmology. Cross-correlations work by comparing independent datasets that trace the same underlying structure, allowing one to extract the common signal while mitigating uncorrelated sources of contamination. Analogous to the power spectrum defined in Eq. 1.3, the cross-spectrum quantifies the statistical relationship between two fields T_i and T_j as a function of comoving length scale and is given by

$$\langle \tilde{T}_i(\mathbf{k})\tilde{T}_j(\mathbf{k}')\rangle = (2\pi)^3\delta_D^{(3)}(\mathbf{k}-\mathbf{k}')P_{ij}(k). \tag{1.20}$$

Importantly, since any non-signal contributions (i.e. noise and foregrounds) to T_i and T_j are uncorrelated with one another, the cross-spectrum is *unbiased* in the infinite ensemble average limit (a detailed discussion of this can be found in Chapter 2). The variance of the cross-spectrum, as described in Furlanetto and Lidz (2007), is given by

$$\delta P_{ij}^2(k) = \frac{1}{2\sqrt{N_\mathbf{k}}}\left[P_{ij}^2(k) + \delta P_i(k)\delta P_j(k)\right] \tag{1.21}$$

where $\delta P_i(k)$ and $\delta P_j(k)$ is the error on the individual auto-spectra. The factor of $1/\sqrt{N_\mathbf{k}}$ is the reduction of sample variance from binning vectors $\mathbf{k}$ into a k-bin. The factor of $1/2$ comes from only sampling the upper half-plane since the power spectrum is the square of the Fourier transform of a real-valued field. This discussion has hopefully demonstrated that cross-correlations are a powerful, unbiased probe and showcased how bringing in a secondary probe, especially one that is detected at high signal-to-noise (i.e. has small δP) can yield a detection of a faint probe.

Historically, cross-correlations have been instrumental in securing the first detections of faint cosmological signals that were otherwise buried within complex datasets. For example, CMB lensing (Smith et al., 2007), the kinematic Sunyaev-Zel'dovich effect[9] (Hand et al., 2012), the

---

[9] The kinematic Sunyaev-Zel'dovich effect is a CMB secondary and arises from the scattering of CMB photons off

integrated Sachs-Wolfe effect[10] (Boughn and Crittenden, 2004), and, as we have seen, the low-redshift 21 cm signal (Pen et al., 2009), were all first detected in cross-correlation with large-scale structure (LSS) surveys. With high-redshift LIM emerging as a new frontier, cross-correlation analyses are poised to play a similarly crucial role in securing the first detections at high redshift. However, at high-z, traditional cross-correlation techniques face new challenges: high-redshift galaxy surveys are sparse and cover relatively small fields of view, making it difficult to robustly correlate LIM maps with LSS datasets. As a result, relying solely on these conventional approaches will likely not be sufficient to secure a detection (La Plante et al., 2023). This necessitates the development of innovative cross-correlation strategies tailored to the high-redshift regime.

1.6 Roadmap

This thesis explores new methods for maximizing the scientific return of LIM measurements and their cross-correlations, with emphasis on their role in constraining cosmology. The work presented here is divided into two major objectives. The first objective, covered in Chapters 2 and 3, is to establish robust frameworks for analyzing LIM cross-correlations and demonstrate their effectiveness in both mitigating systematics and probing fundamental cosmological and astrophysical processes. The second objective, addressed in Chapters 4 and 5, is to introduce novel applications of LIM lensing that enable access to high-redshift CMB lensing information, providing new avenues for constraining cosmology. Each chapter is summarized as follows:

- **Chapter 2**, "*Forecasts and Statistical Insights for Line Intensity Mapping Cross-Correlations: A Case Study with 21cm* $\times$ *[CII]*" focuses on the theoretical and practical aspects of cross-correlating LIM surveys. With multiple spectral lines probing overlapping cosmic volumes, cross-correlations offer a promising avenue for studying

---

of high-energy electrons with bulk velocities.

[10] This is a gravitational CMB secondary which causes the redshifting and blueshifting of CMB photons by the gravitational potential wells they encounter in an expanding universe.

astrophysical processes while mitigating observational systematics. However, many of the desirable statistical properties of these correlations hold only under idealized conditions. In this chapter, we develop an end-to-end pipeline for analyzing LIM cross-correlations, enabling instrumental effects, residual foregrounds, and analysis choices to be propagated through Monte Carlo simulations. This framework rigorously quantifies error properties, including covariances, window functions, and full probability distributions for power spectrum estimates, analysis steps that, for simplicity, are typically not carried out in existing literature. This then allows one to test, rather than assume, various simplifying assumptions that have yet to be critically examined. As case studies, we apply this methodology to forecast the sensitivity of upcoming 21 cm-[CII] cross-correlation measurements and provide recommendations for survey design.

- **Chapter 3**, “*A Data-Driven Technique To Mitigate The Foregrounds Of Line Intensity Maps*”, addresses one of the most significant challenges in LIM: foreground contamination. Despite extensive efforts, many existing mitigation techniques have struggled to effectively remove the dominant broadband contaminants, particularly those with strong frequency coherence. In this chapter, we introduce a novel, fully data-driven foreground removal method that leverages intra-dataset frequency-frequency cross-correlations to empirically estimate the foreground power spectrum. We test this approach using simulations of 21 cm observations under various experimental configurations, including single-dish and interferometric setups, while incorporating realistic systematic contaminants. Under idealized conditions, we demonstrate that this method can suppress foregrounds by 4–5 orders of magnitude, achieving percent-level signal residuals in a single analysis step. Additionally, we explore the limitations of this estimator, identifying conditions under which it remains robust and scenarios where it is expected to break down. This work establishes a new framework for foreground mitigation that is widely applicable across LIM surveys, ensuring that future analyses can achieve the precision necessary for high-redshift detections.

- **Chapter 4**, "*Constraining Cosmology With the CMB × LIM-Nulling Convergence*" explores the scientific potential of nulling estimators that combine CMB lensing and LIM lensing to selectively remove low-redshift contributions from CMB lensing maps. Standard CMB lensing maps provide a direct probe of the matter distribution without relying on biased tracers, but they are dominated by low-redshift contributions, making it difficult to isolate high-redshift structure. This chapter builds on the formalism of Maniyar et al. (2022), demonstrating that LIM-nulling estimators can provide competitive constraints on ΛCDM parameters and neutrino mass. Furthermore, we highlight LIM-nulling as a powerful tool for model-independent tests of physics beyond ΛCDM and for studying the high-redshift matter power spectrum.

- **Chapter 5**, "*A New Probe of the High-z BAO scale: BAO tomography With CMB × LIM-Nulling Convergence*" extends the application of LIM-nulling techniques to precision cosmology. Standard rulers like the BAO scale are critical for constraining the geometry and expansion history of the Universe, yet most existing BAO measurements operate at low redshifts and rely on biased tracers of the matter field. In this chapter, we propose a novel approach to measuring BAO features at high redshifts ($z \sim 5$) by applying nulling estimators to LIM-CMB lensing cross-correlations. We forecast the sensitivity of next-generation instruments to such a measurement, finding that upcoming surveys could constrain the BAO scale to 7.2% precision, while a more futuristic scenario could reach 4% precision. This constitutes a fundamentally new BAO measurement during an epoch largely unexplored by traditional methods.

Together, these studies advance the use of LIM for precision cosmology, demonstrating how cross-correlations and nulling techniques can unlock new high-redshift observables and overcome key observational challenges. Finally, we summarize our conclusions and and provide additional outlook in Chapter 6.

Bibliography

Planck Collaboration, **594**, A13 (2016).

P. J. E. Peebles, Journal of Physics: Conference Series **484**, 012001 (2014), ISSN 1742-6596, URL http://dx.doi.org/10.1088/1742-6596/484/1/012001.

B. Abbott et al., Physical Review Letters **116** (2016), ISSN 1079-7114, URL http://dx.doi.org/10.1103/PhysRevLett.116.061102.

R. H. Cyburt et al., Reviews of Modern Physics **88** (2016), ISSN 1539-0756, URL http://dx.doi.org/10.1103/RevModPhys.88.015004.

K. Bundy et al., **798**, 7 (2015), 1412.1482.

S. Perlmutter and T. S. C. Project, **517**, 565 (1999), astro-ph/9812133.

G. D. Becker, J. S. Bolton, P. Madau, M. Pettini, E. V. Ryan-Weber, and B. P. Venemans, **447**, 3402 (2015), 1407.4850.

S. R. Furlanetto, S. Peng Oh, and F. H. Briggs, Physics Reports **433**, 181–301 (2006), ISSN 0370-1573, URL http://dx.doi.org/10.1016/j.physrep.2006.08.002.

A. Liu and J. R. Shaw, **132**, 062001 (2020), 1907.08211.

J. R. Pritchard and A. Loeb, Reports on Progress in Physics **75**, 086901 (2012), ISSN 1361-6633, URL http://dx.doi.org/10.1088/0034-4885/75/8/086901.

M. Tegmark and M. Zaldarriaga, Phys. Rev. D **79**, 083530 (2009), URL https://link.aps.org/doi/10.1103/PhysRevD.79.083530.

D. Scott and M. J. Rees, **247**, 510 (1990).

J. B. Muñoz, Y. Ali-Haïmoud, and M. Kamionkowski, Phys. Rev. D **92**, 083508 (2015), URL https://link.aps.org/doi/10.1103/PhysRevD.92.083508.

Y. Mao, M. Tegmark, M. McQuinn, M. Zaldarriaga, and O. Zahn, Phys. Rev. D **78**, 023529 (2008), URL https://link.aps.org/doi/10.1103/PhysRevD.78.023529.

A. Loeb and M. Zaldarriaga, Phys. Rev. Lett. **92**, 211301 (2004), URL https://link.aps.org/doi/10.1103/PhysRevLett.92.211301.

H. C. Chiang, T. Dyson, E. Egan, S. Eyono, N. Ghazi, J. Hickish, J. M. Jáuregui-Garcia, V. Manukha, T. Ménard, T. Moso, et al., Journal of Astronomical Instrumentation **9**, 2050019-564 (2020), 2008.12208.

J. O. Burns, G. Hallinan, J. Lux, L. Teitelbaum, J. Kocz, R. MacDowall, R. Bradley, D. Rapetti, W. Wu, S. Furlanetto, et al., arXiv e-prints arXiv:1911.08649 (2019), 1911.08649.

D. R. DeBoer et al., **129**, 045001 (2017), 1606.07473.

Z. Abdurashidova, J. E. Aguirre, P. Alexander, Z. S. Ali, Y. Balfour, R. Barkana, A. P. Beardsley, G. Bernardi, T. S. Billings, J. D. Bowman, et al., **924**, 51 (2022), 2108.07282.

P. Madau, L. Pozzetti, and M. Dickinson, **498**, 106 (1998), astro-ph/9708220.

S. R. Furlanetto, M. Zaldarriaga, and L. Hernquist, **613**, 16 (2004), astro-ph/0404112.

R. Barkana and A. Loeb, **624**, L65 (2005), astro-ph/0409572.

S. R. Furlanetto, Z. Haiman, and S. P. Oh, **686**, 25 (2008), 0803.3454.

K. W. Masui, E. R. Switzer, N. Banavar, K. Bandura, C. Blake, L. M. Calin, T. C. Chang, X. Chen, Y. C. Li, Y. W. Liao, et al., **763**, L20 (2013), 1208.0331.

S. H. Teng, S. Veilleux, and A. J. Baker, **765**, 95 (2013), 1301.5642.

J. Rhee, P. Lah, J. N. Chengalur, F. H. Briggs, and M. Colless, **460**, 2675 (2016), 1605.02006.

C. J. Anderson, N. J. Luciw, Y. C. Li, C. Y. Kuo, J. Yadav, K. W. Masui, T. C. Chang, X. Chen, N. Oppermann, Y. W. Liao, et al., **476**, 3382 (2018), 1710.00424.

A. Chowdhury, N. Kanekar, B. Das, K. S. Dwarakanath, and S. Sethi, **913**, L24 (2021), 2105.06773.

L. Wolz, A. Pourtsidou, K. W. Masui, T.-C. Chang, J. E. Bautista, E.-M. Müller, S. Avila, D. Bacon, W. J. Percival, S. Cunnington, et al., **510**, 3495 (2022), 2102.04946.

CHIME Collaboration, M. Amiri, K. Bandura, T. Chen, M. Deng, M. Dobbs, M. Fandino, S. Foreman, M. Halpern, A. S. Hill, et al., arXiv e-prints arXiv:2202.01242 (2022), 2202.01242.

M. Amiri, K. Bandura, A. Chakraborty, M. Dobbs, M. Fandino, S. Foreman, H. Gan, M. Halpern, A. S. Hill, G. Hinshaw, et al., **963**, 23 (2024), 2309.04404.

S. Paul, M. G. Santos, Z. Chen, and L. Wolz, *A first detection of neutral hydrogen intensity mapping on mpc scales at* $z \approx 0.32$ *and* $z \approx 0.44$ (2023), 2301.11943.

P. C. Breysse, E. D. Kovetz, P. S. Behroozi, L. Dai, and M. Kamionkowski, Monthly Notices of the Royal Astronomical Society p. stx203 (2017), ISSN 1365-2966, URL http://dx.doi.org/10.1093/mnras/stx203.

A. Cooray and R. Sheth, **372**, 1 (2002), astro-ph/0206508.

A. Lewis and A. Challinor, *CAMB: Code for Anisotropies in the Microwave Background*, Astrophysics Source Code Library, record ascl:1102.026 (2011).

A. R. Pullen, T.-C. Chang, O. Doré, and A. Lidz, **768**, 15 (2013), 1211.1397.

A. R. Pullen, P. Serra, T.-C. Chang, O. Doré, and S. Ho, **478**, 1911 (2018), 1707.06172.

E. Switzer, *Constraining star formation through redshifted CO and CII emission in archival CMB data*, NASA Proposal id.16-ADAP16-33 (2016).

A. Roy, N. Battaglia, and A. R. Pullen, arXiv e-prints arXiv:2406.07861 (2024), `2406.07861`.

G. K. Keating, G. C. Bower, D. P. Marrone, D. R. DeBoer, C. Heiles, T.-C. Chang, J. E. Carlstrom, C. H. Greer, D. Hawkins, J. W. Lamb, et al., **814**, 140 (2015), `1510.06744`.

G. K. Keating, D. P. Marrone, G. C. Bower, E. Leitch, J. E. Carlstrom, and D. R. DeBoer, **830**, 34 (2016), `1605.03971`.

N. O. Stutzer, J. G. S. Lunde, P. C. Breysse, D. T. Chung, K. A. Cleary, D. A. Dunne, H. K. Eriksen, H. T. Ihle, H. Padmanabhan, D. Tolgay, et al., arXiv e-prints arXiv:2406.07511 (2024), `2406.07511`.

D. T. Chung, P. C. Breysse, K. A. Cleary, D. A. Dunne, J. G. S. Lunde, H. Padmanabhan, N. O. Stutzer, D. Tolgay, J. R. Bond, S. E. Church, et al., arXiv e-prints arXiv:2406.07512 (2024), `2406.07512`.

H. T. Ihle, J. Borowska, K. A. Cleary, H. K. Eriksen, M. K. Foss, S. E. Harper, J. Kim, J. G. S. Lunde, L. Philip, M. Rasmussen, et al., The Astrophysical Journal **933**, 185 (2022), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/ac63c5`.

J. L. Bernal and E. D. Kovetz, arXiv (2022).

T. Oxholm, P. A. Ade, C. J. Anderson, E. M. Barrentine, N. G. Bellis, A. D. Bolatto, P. C. Breysse, B. T. Bulcha, J. A. Connors, G. Cataldo, et al., in *American Astronomical Society Meeting Abstracts #236* (2020), vol. 236 of *American Astronomical Society Meeting Abstracts*, p. 244.02.

A. T. Crites, J. J. Bock, C. M. Bradford, T. C. Chang, A. R. Cooray, L. Duband, Y. Gong, S. Hailey-Dunsheath, J. Hunacek, P. M. Koch, et al., in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, edited by W. S. Holland and J. Zmuidzinas

(2014), vol. 9153 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 91531W.

D. P. Marrone, J. E. Aguirre, J. S. Bracks, C. M. Bradford, B. S. Brendal, B. Bumble, A. J. Corso, M. J. Devlin, N. Emerson, J. P. Filippini, et al., in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*, edited by J. Zmuidzinas and J.-R. Gao (2022), vol. 12190 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 1219008.

Concerto Collaboration, P. Ade, M. Aravena, E. Barria, A. Beelen, A. Benoit, M. Béthermin, J. Bounmy, O. Bourrion, G. Bres, et al., **642**, A60 (2020).

M. Aravena, , et al., *The ccat-prime submillimeter observatory* (2019), `1909.02587`.

B. P. Crill, M. Werner, R. Akeson, M. Ashby, L. Bleem, J. J. Bock, S. Bryan, J. Burnham, J. Byunh, T.-C. Chang, et al., in *Space Telescopes and Instrumentation 2020: Optical, Infrared, and Millimeter Wave*, edited by M. Lystrup and M. D. Perrin (2020), vol. 11443 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 114430I, `2404.11017`.

J. Fonseca, M. B. Silva, M. G. Santos, and A. Cooray, **464**, 1948 (2017), `1607.05288`.

J. Chluba and R. A. Sunyaev, Monthly Notices of the Royal Astronomical Society **419**, 1294–1314 (2011), ISSN 0035-8711, URL `http://dx.doi.org/10.1111/j.1365-2966.2011.19786.x`.

J. C. Mather, E. S. Cheng, D. A. Cottingham, R. E. Eplee, Jr., D. J. Fixsen, T. Hewagama, R. B. Isaacman, K. A. Jensen, S. S. Meyer, P. D. Noerdlinger, et al., **420**, 439 (1994).

D. J. Fixsen, E. S. Cheng, J. M. Gales, J. C. Mather, R. A. Shafer, and E. L. Wright, **473**, 576 (1996), `astro-ph/9605054`.

D. J. Fixsen, The Astrophysical Journal **707**, 916–920 (2009), ISSN 1538-4357, URL http://dx.doi.org/10.1088/0004-637X/707/2/916.

M. A. Janssen and S. Gulkis, in *The Infrared and Submillimetre Sky after COBE*, edited by M. Signore and C. Dupraz (1992), vol. 359 of *NATO Advanced Study Institute (ASI) Series C*, p. 391.

C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, D. N. Spergel, G. S. Tucker, et al., **148**, 1 (2003), astro-ph/0302207.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. I. R. Alves, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, H. Aussel, et al., **571**, A1 (2014), 1303.5062.

J. A. Vazquez Gonzalez, L. E. Padilla, and T. Matos, Revista Mexicana de Física E **17**, 73–91 (2020), ISSN 1870-3542, URL http://dx.doi.org/10.31349/RevMexFisE.17.73.

N. Bartolo, S. Matarrese, and A. Riotto, Advances in Astronomy **2010** (2010), ISSN 1687-7977, URL http://dx.doi.org/10.1155/2010/157079.

N. Padmanabhan, M. White, and J. D. Cohn, Physical Review D **79** (2009), ISSN 1550-2368, URL http://dx.doi.org/10.1103/PhysRevD.79.063523.

D. J. Eisenstein, I. Zehavi, D. W. Hogg, R. Scoccimarro, M. R. Blanton, R. C. Nichol, R. Scranton, H. Seo, M. Tegmark, Z. Zheng, et al., The Astrophysical Journal **633**, 560–574 (2005), ISSN 1538-4357, URL http://dx.doi.org/10.1086/466512.

M. Kamionkowski and E. D. Kovetz, Annual Review of Astronomy and Astrophysics **54**, 227–269 (2016), ISSN 1545-4282, URL http://dx.doi.org/10.1146/annurev-astro-081915-023433.

A. Lewis and A. Challinor, **429**, 1 (2006), astro-ph/0601594.

U. Seljak and M. Zaldarriaga, **82**, 2636 (1999), astro-ph/9810092.

A. Lewis and A. Challinor, Phys. Rev. D **76**, 083005 (2007), URL https://link.aps.org/doi/10.1103/PhysRevD.76.083005.

M. Schmittfull and U. Seljak, **97**, 123540 (2018), 1710.09465.

R. Allison, P. Caucal, E. Calabrese, J. Dunkley, and T. Louis, **92**, 123535 (2015), 1509.07471.

S. Dodelson, *Gravitational Lensing* (Cambridge University Press, 2017).

K. Umetsu, arXiv e-prints arXiv:1002.3952 (2010), 1002.3952.

BICEP/Keck Collaboration, :, P. A. R. Ade, Z. Ahmed, M. Amiri, D. Barkats, R. Basu Thakur, D. Beck, C. Bischoff, J. J. Bock, et al., arXiv e-prints arXiv:2203.16556 (2022), 2203.16556.

W. Hu and T. Okamoto, **574**, 566 (2002), astro-ph/0111606.

T. Okamoto and W. Hu, **67**, 083002 (2003), astro-ph/0301031.

C. M. Hirata and U. Seljak, **68**, 083002 (2003a), astro-ph/0306354.

J. Carron and A. Lewis, **96**, 063510 (2017), 1704.08230.

C. M. Hirata and U. Seljak, **67**, 043001 (2003b), astro-ph/0209489.

B. Hadzhiyska, B. D. Sherwin, M. Madhavacheril, and S. Ferraro, **100**, 023547 (2019), 1905.04217.

M. Millea, E. Anderes, and B. D. Wandelt, **102**, 123542 (2020), 2002.00965.

A. S. Maniyar, Y. Ali-Haïmoud, J. Carron, A. Lewis, and M. S. Madhavacheril, Phys. Rev. D **103**, 083524 (2021), URL https://link.aps.org/doi/10.1103/PhysRevD.103.083524.

S. Das, B. D. Sherwin, P. Aguirre, J. W. Appel, J. R. Bond, C. S. Carvalho, M. J. Devlin, J. Dunkley, R. Dünner, T. Essinger-Hileman, et al., **107**, 021301 (2011), 1103.2124.

A. van Engelen, R. Keisler, O. Zahn, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, T. M. Crawford, et al., **756**, 142 (2012), 1202.0546.

D. Hanson, S. Hoover, A. Crites, P. A. R. Ade, K. A. Aird, J. E. Austermann, J. A. Beall, A. N. Bender, B. A. Benson, L. E. Bleem, et al., **111**, 141301 (2013), 1307.5830.

B. D. Sherwin, A. van Engelen, N. Sehgal, M. Madhavacheril, G. E. Addison, S. Aiola, R. Allison, N. Battaglia, D. T. Becker, J. A. Beall, et al., Phys. Rev. D **95**, 123529 (2017), URL `https://link.aps.org/doi/10.1103/PhysRevD.95.123529`.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, et al., **594**, A15 (2016), 1502.01591.

Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, et al., **641**, A8 (2020), 1807.06210.

Y. Omori, R. Chown, G. Simard, K. T. Story, K. Aylor, E. J. Baxter, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, et al., **849**, 124 (2017), 1705.00743.

K. T. Story, D. Hanson, P. A. R. Ade, K. A. Aird, J. E. Austermann, J. A. Beall, A. N. Bender, B. A. Benson, L. E. Bleem, J. E. Carlstrom, et al., **810**, 50 (2015), 1412.4760.

BICEP2 Collaboration, Keck Array Collaboration, P. A. R. Ade, Z. Ahmed, R. W. Aikin, K. D. Alexander, D. Barkats, S. J. Benton, C. A. Bischoff, J. J. Bock, et al., **833**, 228 (2016), 1606.01968.

W. L. K. Wu, L. M. Mocanu, P. A. R. Ade, A. J. Anderson, J. E. Austermann, J. S. Avva, J. A. Beall, A. N. Bender, B. A. Benson, F. Bianchini, et al., **884**, 70 (2019), 1905.05777.

P. A. R. Ade, Y. Akiba, A. E. Anthony, K. Arnold, M. Atlas, D. Barron, D. Boettger, J. Borrill, S. Chapman, Y. Chinone, et al., **113**, 021301 (2014), 1312.6646.

S. Adachi, M. A. O. Aguilar Faúndez, Y. Akiba, A. Ali, K. Arnold, C. Baccigalupi, D. Barron, D. Beck, F. Bianchini, J. Borrill, et al. (POLARBEAR Collaboration), Phys. Rev. Lett. **124**, 131301 (2020), URL `https://link.aps.org/doi/10.1103/PhysRevLett.124.131301`.

J. Ward and N. S. T. R. F. Advanced ACT Collaboration, in *American Astronomical Society Meeting Abstracts #230* (2017), vol. 230 of *American Astronomical Society Meeting Abstracts*, p. 117.01.

B. A. Benson, P. A. R. Ade, Z. Ahmed, S. W. Allen, K. Arnold, J. E. Austermann, A. N. Bender, L. E. Bleem, J. E. Carlstrom, C. L. Chang, et al., in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, edited by W. S. Holland and J. Zmuidzinas (2014), vol. 9153 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 91531P, 1407.2973.

A. J. Anderson, P. Barry, A. N. Bender, B. A. Benson, L. E. Bleem, J. E. Carlstrom, T. W. Cecil, C. L. Chang, T. M. Crawford, K. R. Dibert, et al., arXiv e-prints arXiv:2208.08559 (2022), 2208.08559.

P. Ade, J. Aguirre, Z. Ahmed, S. Aiola, A. Ali, D. Alonso, M. A. Alvarez, K. Arnold, P. Ashton, J. Austermann, et al., **2019**, 056 (2019), 1808.07445.

K. Abazajian, G. E. Addison, P. Adshead, Z. Ahmed, D. Akerib, A. Ali, S. W. Allen, D. Alonso, M. Alvarez, M. A. Amin, et al., **926**, 54 (2022), 2008.12619.

N. Hand, A. Leauthaud, S. Das, B. D. Sherwin, G. E. Addison, J. R. Bond, E. Calabrese, A. Charbonnier, M. J. Devlin, J. Dunkley, et al., **91**, 062001 (2015), 1311.6200.

J. Liu and J. C. Hill, **92**, 063517 (2015), 1504.05598.

S. Singh, R. Mandelbaum, and J. R. Brownstein, Monthly Notices of the Royal Astronomical Society **464**, 2120–2138 (2016), ISSN 1365-2966, URL http://dx.doi.org/10.1093/mnras/stw2482.

C. Chang, Y. Omori, E. Baxter, C. Doux, A. Choi, S. Pandey, A. Alarcon, O. Alves, A. Amon, F. Andrade-Oliveira, et al., Physical Review D **107** (2023), ISSN 2470-0029, URL http://dx.doi.org/10.1103/PhysRevD.107.023530.

A. Kalaja, I. Harrison, and W. R. Coulton, arXiv e-prints arXiv:2412.14713 (2024), 2412.14713.

S. Shaikh, I. Harrison, A. van Engelen, G. A. Marques, T. M. C. Abbott, M. Aguena, O. Alves, A. Amon, R. An, D. Bacon, et al., **528**, 2112 (2024), 2309.04412.

A. Liu and M. Tegmark, Monthly Notices of the Royal Astronomical Society **419**, 3491–3504 (2012), ISSN 1365-2966.

X. Wang, M. Tegmark, M. G. Santos, and L. Knox, **650**, 529 (2006), astro-ph/0501081.

A. Liu, A. R. Parsons, and C. M. Trott, Physical Review D **90** (2014a), ISSN 1550-2368, URL http://dx.doi.org/10.1103/PhysRevD.90.023018.

A. Liu, A. R. Parsons, and C. M. Trott, Physical Review D **90** (2014b), ISSN 1550-2368, URL http://dx.doi.org/10.1103/PhysRevD.90.023019.

A. Liu, A. R. Parsons, and C. M. Trott, Physical Review D **90** (2014c), ISSN 1550-2368, URL http://dx.doi.org/10.1103/PhysRevD.90.023018.

B. Yue, A. Ferrara, A. Pallottini, S. Gallerani, and L. Vallini, Monthly Notices of the Royal Astronomical Society **450**, 3829 (2015), URL https://doi.org/10.1093%2Fmnras%2Fstv933.

A. Williams, H. Hui, and J. Bock, in *American Astronomical Society Meeting Abstracts* (2024), vol. 243 of *American Astronomical Society Meeting Abstracts*, p. 258.15.

S. R. Furlanetto and A. Lidz, The Astrophysical Journal **660**, 1030–1038 (2007), ISSN 1538-4357, URL `http://dx.doi.org/10.1086/513009`.

K. M. Smith, O. Zahn, and O. Doré, Physical Review D **76** (2007), ISSN 1550-2368, URL `http://dx.doi.org/10.1103/PhysRevD.76.043510`.

N. Hand, G. E. Addison, E. Aubourg, N. Battaglia, E. S. Battistelli, D. Bizyaev, J. R. Bond, H. Brewington, J. Brinkmann, B. R. Brown, et al., Physical Review Letters **109** (2012), ISSN 1079-7114, URL `http://dx.doi.org/10.1103/PhysRevLett.109.041101`.

S. Boughn and R. Crittenden, Nature **427**, 45–47 (2004), ISSN 1476-4687, URL `http://dx.doi.org/10.1038/nature02139`.

U.-L. Pen, L. Staveley-Smith, J. B. Peterson, and T.-C. Chang, Monthly Notices of the Royal Astronomical Society: Letters **394**, L6–L10 (2009), ISSN 1745-3925, URL `http://dx.doi.org/10.1111/j.1745-3933.2008.00581.x`.

P. La Plante, J. Mirocha, A. Gorce, A. Lidz, and A. Parsons, The Astrophysical Journal **944**, 59 (2023), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/acaeb0`.

A. S. Maniyar, E. Schaan, and A. R. Pullen, **105**, 083509 (2022), `2106.09005`.

Intermezzo 1

At the end of Chapter 1, we explored how cross-correlations help extract faint signals from otherwise contaminated datasets and how LIM will likely require us to rethink both our analysis methods and the probes we use. In this next chapter, we investigate the potential of LIM-LIM cross-correlations, where two line intensity mapping observations are cross-correlated instead of LIM with another LSS probe. This approach has significant advantages, as LIM, in principle, provides a more representative view of LSS than traditional galaxy surveys. However, implementing LIM-LIM cross-correlations in reality presents several challenges.

First, unlike cross-correlations with external tracers, LIM-LIM analyses involve two probes that are both subdominant to the foreground contamination. While the resulting power spectrum remains unbiased, the variance on these measurements can be significantly high—an issue that, until now, has yet to be fully quantified. Second, differences in instrumental specifications, such as frequency and angular resolution, complicate the analysis by requiring the careful selection of commonly measured Fourier modes. Third, although LIM in principle can probe vast cosmological volumes, many initial surveys, particularly those at millimeter and sub-millimeter wavelengths, are conducted over relatively small fields of view to enable deep integrations. Finally, since most current-generation LIM instruments were primarily designed for auto-spectrum measurements, many surveys lack large overlapping sky areas, further exacerbating the challenge of performing high-sensitivity cross-correlations.

These difficulties raise crucial questions: If we attempt LIM-LIM cross-correlations with existing instrumentation, what can we realistically achieve? And what would an ideal experimental setup look like to fully realize the precision science that cross-correlations promise?

In Chapter 2, we address these questions and systematically explore the four key challenges outlined above. To do so, we perform the first end-to-end simulation and analysis of LIM-LIM cross-correlations, constructing a robust framework to assess the impact of foregrounds, instrument response, instrumental noise, and other observational effects on the measured cross-spectrum. This pipeline allows us to propagate these effects through Monte Carlo simulations, providing a rigorous characterization of error properties, including covariances, window functions, and statistical uncertainties. With this comprehensive approach, we forecast the sensitivity of both near-term and futuristic 21 cm × [CII] EoR surveys, offering insight into the feasibility of early high-redshift detections. Additionally, our framework enables a critical evaluation of commonly held forecasting assumptions and analysis choices, ensuring that future cross-correlation studies are built on a solid methodological foundation. Through this work, we establish a path forward for LIM-LIM analyses, demonstrating their viability as a powerful tool for high-redshift cosmology.

This chapter is based on:

- **Fronenberg, H.**, Liu, A. (2024) *Forecasts and Statistical Insights for Line Intensity Mapping Cross-Correlations: A Case Study with 21cm × [CII]*, ApJ **975** 222

Chapter 2

Forecasts and Statistical Insights for Line Intensity Mapping Cross-Correlations: A Case Study with 21cm × [CII]

Hannah Fronenberg[1,2], **Adrian Liu**[1,2]

[1]Department of Physics, McGill University, Montréal, QC, Canada

[2]Trottier Space Institute, Montréal, QC, Canada

Abstract

Intensity mapping—the large-scale mapping of selected spectral lines without resolving individual sources—is quickly emerging as an efficient way to conduct large cosmological surveys. Multiple surveys covering a variety of lines (such as the hydrogen 21 cm hyperfine line, CO rotational lines, and [CII] fine structure lines, among others) are either observing or will soon be online, promising a panchromatic view of our Universe over a broad redshift range. With multiple lines potentially

covering the same volume, cross-correlations have become an attractive prospect, both for probing the underlying astrophysics and for mitigating observational systematics. For example, cross correlating 21 cm and [CII] intensity maps during reionization could reveal the characteristic scale of ionized bubbles around the first galaxies, while simultaneously providing a convenient way to reduce independent foreground contaminants between the two surveys. However, many of the desirable properties of cross-correlations in principle emerge only under ideal conditions, such as infinite ensemble averages. In this paper, we construct an end-to-end pipeline for analyzing intensity mapping cross-correlations, enabling instrumental effects, foreground residuals, and analysis choices to be propagated through Monte Carlo simulations to a set of rigorous error properties, including error covariances, window functions, and full probability distributions for power spectrum estimates. We use this framework to critically examine the applicability of simplifying assumptions such as the independence and Gaussianity of power spectrum errors. As worked examples, we forecast the sensitivity of near-term and futuristic 21 cm-[CII] cross-correlation measurements, providing recommendations for survey design.

2.1 Introduction

During the period of cosmic dawn and and the epoch of reionization (EoR), our Universe's first stars ignited and the first galaxies formed. Emanating from this first generation of galaxies were UV photons that ionized the surrounding neutral hydrogen (HI) over some several hundred million years. The precise timing, duration, morphology, and sources of reionization remains unknown, representing a crucial missing piece in our understanding of our cosmic history.

Line intensity mapping (LIM) has the potential to revolutionize our status quo of considerably incomplete information. LIM traces the brightness of specific spectral lines over large three-dimensional volumes, mapping out radial fluctuations using the redshift of a line and transverse fluctuations via angular sky plane information. A judicious choice of spectral lines (balancing practical feasibility with sensitivity to physical processes of interest) opens up heretofore unexplored epochs and scales to direct observation. Particularly promising for the study of the EoR is the 21 cm hyperfine transition line of neutral hydrogen, since it is precisely neutral hydrogen that is being ionized and the relevant redshift range is in principle observable (Furlanetto et al., 2006; Morales and Wyithe, 2010; Pritchard and Loeb, 2012; Loeb and Furlanetto, 2013; Mesinger, 2019; Liu and Shaw, 2020).

A number of 21 cm observations have been made in the range $0 \lesssim z \lesssim 20$ (Liu and Shaw, 2020; Trott and Pober, 2019). Global signal experiments have targeted the sky-averaged 21 cm monopole as a function of redshift. At the high redshifts of cosmic dawn (i.e., during the formation of first stars and galaxies but prior to the systematic reionization of the intergalactic medium), a global signal detection has been claimed by the Experiment to Detect the Global EoR Signature (EDGES) collaboration (Bowman et al., 2018), although a search for this claimed signal with the Shaped Antenna measurement of the background RAdio Spectrum 3 (SARAS-3) experiment has not yielded a detection (Singh et al., 2022). At reionization redshifts, global signal experiments have attempted to detect the disappearance of neutral hydrogen as reionization proceeds, and have successfully placed lower limits on the duration of reionzation (Bowman and Rogers, 2010; Monsalve et al., 2017; Singh et al., 2017). Even without a positive detection, such measurements have the potential to set powerful constraints by combining with observations of the kinetic Sunyaev-Zel'dovich in the Cosmic Microwave Background, which tend to place upper limits on the duration (Bégin et al., 2022).

Complementing global signal measurements are experiments targeting spatial fluctuations in the 21 cm brightness temperature. Current-generation instruments typically aim to make a measurement of the power spectrum of these fluctuations, and stringent upper limits have been placed by the Hydrogen Epoch of Reionization Array (HERA; Abdurashidova et al. 2022a; HERA Collaboration et al. 2023), the Giant Meter Wave Radio Telescope (GMRT; Paciga et al. 2013), the Low Frequency Array (LOFAR; van Haarlem et al. 2013; Patil et al. 2017; Gehlot et al. 2019; Mertens et al. 2020), the Donald C. Backer Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al. 2010; Cheng et al. 2018; Kolopanis et al. 2019), the Owens Valley Long

Wavelength Array (LWA; Eastwood et al. 2019; Garsden et al. 2021), the Murchison Widefield Array (MWA; Tingay et al. 2013; Ewall-Wice et al. 2016; Beardsley et al. 2016; Barry et al. 2019; Trott et al. 2020; Li et al. 2019; Yoshiura et al. 2021; Trott et al. 2022; Rahimi et al. 2021; Wilensky et al. 2023a), and the New Extension in Nançay Upgrading LOFAR (NenuFAR; Munshi et al. 2024). These upper limits have begun to constrain theoretical models (HERA Collaboration et al., 2023; Ghara et al., 2020; Greig et al., 2021; Abdurashidova et al., 2022b), but despite this progress, there remains no positive detection of any spatially fluctuating signal during the EoR.

The lack of an EoR detection despite such concerted effort is no coincidence; it is a challenging feat. The highly redshifted 21 cm line falls in low-frequency radio bands, where there are severe foreground contaminants from more local sources of emission (Santos et al., 2005; Wang et al., 2006; Jelić et al., 2008; Chapman and Jelić, 2019). Examples of such foregrounds include Galactic synchrotron emission, bright radio point sources, Bremsstahlung emission, ionospheric effects, and radio frequency interference (RFI), which dominate over the cosmological signal by up to five orders of magnitude (Liu and Tegmark, 2012). The effects of these foregrounds must therefore be mitigated, and a vast majority of proposed methods rely on the relative spectral smoothness of foregrounds compared to the more spectrally jagged cosmological signal (see Liu and Shaw 2020 for an overview of methods). While this is a sound idea in theory, in practice instrument systematics can imprint extra spectral structure on the intrinsically spectrally smooth foregrounds, rendering them more signal-like (Lanman et al., 2020). Therefore, a successful detection of the 21 cm signal will necessitate well-calibrated instruments and careful foreground mitigation. These requirements are difficult to execute both to high precision and without being subject to modeling biases. As such, some have looked to use complementary probes to bolster a 21 cm detection, while others have called for far more realistic foreground and instrument simulations in order to better understand and analyse real cosmological data.

One alternative that has seen considerable success in the low-redshift ($z < 6$) 21 cm cosmology community is to detect the 21 cm signal via cross-correlations. On the journey to detecting the low-redshift 21cm power spectrum, a number of 21 cm cross-correlations (Masui et al., 2013; Wolz et al., 2022; Anderson et al., 2018; Chowdhury et al., 2021; Cunnington et al., 2023; Amiri et al., 2023; Amiri et al., 2024) validated the existence of post-reionization 21 cm emission in galaxies while providing an easier path towards mitigating systematics such as foregrounds. After about a decade of detecting the low-z 21 cm signal cross-correlation with various other probes, the MeerKAT telescope has recently claimed the first detection of the 21 cm auto-spectrum (Paul et al., 2023).

In the context of high-z EoR 21 cm cosmology, cross-correlations have similarly been touted as promising avenue for systematics mitigation, while also being scientifically interesting in furthering our understanding of the interplay between ionizing sources and the intergalactic medium (IGM; Bernal and Kovetz 2022). Possible partners for such cross-correlations are the large suite of emission lines that are being targeted for LIM over a wide variety of redshifts (Gong et al., 2017; Yang et al., 2021; Karkare et al., 2022a). These include CO rotational lines (Lidz et al., 2011; Carilli, 2011; Keating et al., 2015, 2016, 2020; Karkare et al., 2022b; Cleary et al., 2022; Ihle et al., 2022; Lunde et al., 2024; Stutzer et al., 2024; Chung et al., 2024; Roy et al., 2024), the Lyα line (Silva et al., 2013; Pullen et al., 2014; Croft et al., 2018; Mas-Ribas and Chang, 2020; Kakuma et al., 2021; Renard et al., 2021; Lin et al., 2022; Kikuchihara et al., 2022), the 158 μm fine structure line of ionized carbon [CII] (Gong et al., 2012; Crites et al., 2014; Uzgil et al., 2014; Yue et al., 2015; Karoumpis et al., 2022; Padmanabhan, 2019; Chung et al., 2020; Pullen et al., 2018; Yang et al., 2019; Concerto Collaboration et al., 2020; Ade et al., 2020; Vieira et al., 2020; Cataldo et al., 2020; Essinger-Hileman et al., 2022; Béthermin et al., 2022; CCAT-Prime Collaboration et al., 2023; Van Cuyck et al., 2023), and [OIII] (Padmanabhan et al., 2022), among other proposals such as [OII], HD, Hα, and Hβ (Gong et al., 2017; Doré et al., 2014; Breysse et al., 2022). Each of these prospects will not only be a powerful probe of astrophysics (e.g., Sun et al. 2021; Mirocha et al. 2022; Parsons et al. 2022; Bernal and Kovetz 2022; Sun et al. 2023; Horlaville et al. 2024) and cosmology (e.g., Karkare et al. 2022a; Fonseca et al. 2017; Bernal et al. 2019a,b; Schaan and White 2021a; Moradinezhad Dizgah et al. 2019, 2022a,b) on their own via single-line auto power spectrum measurements, but with the large number of lines available for cross-correlation, multi-tracer techniques will be powerfully constraining (Pullen et al., 2013; Chang et al., 2015; Sun et al., 2019; Schaan and White, 2021b; Anderson et al., 2022; Maniyar et al., 2022; Sato-Polito et al., 2023; Mas-Ribas et al., 2023; Roy et al., 2023; Lujan Niemeyer et al., 2023; Fronenberg et al., 2024a,b).

A particularly interesting prospect is the cross-correlation (and other creative combinations, e.g., Beane and Lidz 2018; Beane et al. 2019; McBride and Liu 2023) of the 21 cm line and the [CII] 158 μm fine structure line (Padmanabhan, 2023). The 21 cm-[CII] cross-spectrum $P_{21\text{cm}\times[\text{CII}]}(k)$ encodes crucial information as a function of comoving wavenumber k, being negative (anti-correlated) on large scales and positive (correlated) on small scales. A measurement of the cross-correlation would yield constraints on the size of ionized regions, the ionization history, and also help constrain astrophysical parameters such as the minimum mass of ionizing haloes, M_{turn} and the number of ionizing photons produced by those haloes (Gong et al., 2012; Dumitru et al., 2019). Most recently, Moriwaki et al. (2024) have explored the prospects of studying the redshift evolution of these negatively and positively correlated regions, showing that they provide insight on the heating and ionization histories of the IGM.

While the science case for 21 cm–[CII] cross-correlations is strong, it is essential to carefully quantify the statistical properties of what we measure. Consider for example the question of whether foregrounds can be suppressed to a reasonable level simply through cross-correlations. If we consider two lines $A \equiv s_a + n_a + f_a$ and $B \equiv s_b + n_b + f_b$ each comprised of their respective cosmological signal, s, foregrounds, f, and noise, n, then their cross-spectrum can be expressed as

$$\begin{aligned}\langle \tilde{A}\tilde{B} \rangle &= \langle (\tilde{s}_a + \tilde{n}_a + \tilde{f}_a)(\tilde{s}_b + \tilde{n}_b + \tilde{f}_b)^* \rangle \\ &= \langle \tilde{s}_a \tilde{s}_b^* \rangle + \langle \tilde{s}_a \tilde{n}_b^* \rangle + \langle \tilde{s}_a \tilde{f}_b^* \rangle + \langle \tilde{s}_b \tilde{n}_a^* \rangle + \langle \tilde{s}_b \tilde{f}_a^* \rangle \\ &\quad + \langle \tilde{n}_a \tilde{f}_b^* \rangle + \langle \tilde{n}_b \tilde{f}_a^* \rangle + \langle \tilde{n}_a \tilde{n}_b^* \rangle + \langle \tilde{f}_a \tilde{f}_b^* \rangle \end{aligned} \tag{2.1}$$

where the tildes denote Fourier transforms. Since lines A and B suffer from different sources of foreground contamination and instrument noise, these contributions will be, on average, uncorrelated with one another and with the cosmological signal. Therefore, Equation (2.1) becomes

$$\langle \tilde{A}\tilde{B}^* \rangle = \langle \tilde{s}_a \tilde{s}_b^* \rangle. \tag{2.2}$$

However, Equation (2.2) only follows in the case of infinite ensemble averaging. In reality, one can never measure infinitely many samples of each Fourier mode on the sky and it is therefore expected that any cross-spectrum estimation will suffer from residual systematic effects. Equivalently, while the power spectrum is statistically unbiased in the ensemble averaged limit, the *variance* $\mathrm{Var}(\mathrm{P}_{21\times[\mathrm{CII}]})$ of 21 cm-[CII] cross spectrum will contain terms proportional to each probe's systematics and foregrounds:

$$\begin{aligned}
\mathrm{Var}\left(\mathrm{P}_{21\times[\mathrm{CII}]}\right) &= 2\left(P^S_{21\times[\mathrm{CII}]}\right)^2 + P^S_{21}P^S_{[\mathrm{CII}]} + P^N_{21}P^N_{[\mathrm{CII}]} \\
&\quad + P^F_{21}P^F_{[\mathrm{CII}]} + P^S_{21}P^N_{[\mathrm{CII}]} + P^N_{21}P^S_{[\mathrm{CII}]} \\
&\quad + P^S_{21}P^F_{[\mathrm{CII}]} + P^F_{21}P^S_{[\mathrm{CII}]} + P^F_{21}P^N_{[\mathrm{CII}]} \\
&\quad + P^N_{21}P^F_{[\mathrm{CII}]},
\end{aligned} \tag{2.3}$$

where each term is a product of power spectra, with subscripts indicating the relevant lines and superscripts indicating the power spectrum of cosmological signal (S), foregrounds (F), and noise (N). For example, the cross-spectrum signal $P^S_{21\times[\mathrm{CII}]}$ is equal to $\langle \tilde{s}_a \tilde{s}_b^* \rangle$. The residual systematics can be thought of as either an extra bias that is randomly drawn from these variance terms, or as increased uncertainties on one's final constraints that can (and should) be included in an error budget. Importantly, the form seen here involves *products* of power spectra, which means that non-trivial interactions exist between the different ingredients of one's measurements, making it crucial to include all of them in our explorations.

In this paper we present an end-to-end pipeline for studying LIM cross-correlations. This pipeline takes as input the power spectra of the lines of interest as well as their correlation coefficient (as a function of scale) in order to produce mock cosmological fields. For concreteness, we focus on 21 cm × [CII] cross-correlations. We add a host of simulated foreground contaminants for each line, include the effect of instrumental responses (e.g., their point spread functions; PSFs), model their thermal noise, and finally estimate auto- and cross-spectra. Drawing an ensemble of Monte Carlo realizations (over thermal noise, cosmological fields, and foregrounds) then allows a rigorous quantification of error statistics. In other words, although Equation (2.3) motivates and guides our thinking, it is *not* used to generate our results, which contain higher-order statistical information beyond the variance of our observables. We use this technology to investigate the effects of various systematic contaminants and the interplay between them. Along the way, we critically examine various commonly used simplifications in forecasting methodologies, such as (but not limited to) the assumption of Gaussian errors, or the independence of foregrounds in different parts of the sky. As such, our paper is highly complementary to (and builds on) Roy and Battaglia (2024), which conducted a series of detailed forecasts for LIM cross-correlations. Roy and Battaglia (2024) focused on signal-to-noise and bias statistics, while we also consider non-Gaussian error distributions and ancillary error properties such as power spectrum error covariances and window functions (defined in Section 2.6.2). Additionally, whereas we consider the 21 cm line and therefore must quantify the attendant complications of interferometry, Roy and Battaglia (2024) consider a broader array of millimeter/sub-millimeter lines. Together, our frameworks can aid in the design of future experiments for cross-correlation, optimizing their sensitivity in the face of systematics. Indeed, as a worked example in this paper we perform a forecast for hypothetical current- and next-generation 21 cm-[CII] cross-correlations.

The rest of this paper is organized as follows. In Section 2.2 we provide an overview of our pipeline and in Sections 2.3 to 2.6 we describe in detail each component of this pipeline and the models involved. Readers that are interested in the final forecasts rather than detailed methodology may wish to skip to Sections 2.7 and 2.8. There, we show the results of mock measurements of 21 cm-[CII] cross-spectra (and their accompanying statistical properties) under realistic observing conditions for current-generation and futuristic scenarios. We comment on the challenges and opportunities for pursuing such measurements. Along the way, we investigate the extent to which commonly employed approximations are appropriate, and we summarize the results of these investigations for practitioners of forecasting in Section 2.9. Our conclusions are presented in Section 2.10.

2.2 Motivation and Pipeline Overview

In this work we explore the potential of using LIM-LIM cross-correlations to probe the epoch of reionization. As a case study, we explore cross-correlations between 21 cm and [CII] LIM observations, although our simulation framework is general and can be easily adapted to explore most LIM-LIM pairs.

The ionized carbon fine structure line, [CII], is of notable interest because it is spatially anti-correlated with the 21 cm line. Carbon is first produced in our universe by Population III (Pop III) stars, the first generation stars. Carbon has an ionization energy of 11.26 eV which is below that of hydrogen, meaning that neutral carbon can be more easily ionized. Once ionized, either in the interstellar medium (ISM) or in the IGM, [CII] can undergo the spin-orbit coupling transition ${}^2P_{3/2} \rightarrow {}^2P_{1/2}$, emitting a photon of wavelength 157.7 μm. This transition can occur by three different mechanisms: collisional emission, spontaneous emission, and stimulated emission (Suginohara et al., 1999; Basu et al., 2004). Within galaxies, the main mechanism for emission is collision with ionized gas in the ISM. In the more diffuse IGM, radiative processes are primarily responsible for [CII] line emission. The excitations here are due to spontaneous emission and stimulated emission from collisions with CMB photons. While emission of the CII line occurs both in the ISM and the IGM, the spin temperature of CII line in the IGM is indistinguishable from the CMB photon temperature during the EoR. In the ISM, the spin temperature of [CII] is much greater than the CMB photon temperature during the EoR and the [CII] brightness temperature signal can be observed (Gong et al., 2012). Since the 21 cm brightness temperature emanates from the IGM and the [CII] brightness temperature from the ISM, these two signals are spatially anti-correlated on large scales during the EoR.

Just as with the 21 cm line, [CII] also suffers from foreground contaminants. Luckily, its spectrally smooth far infrared (FIR) continuum foregrounds are not as detrimental to an observation as diffuse synchrotron emission is to 21 cm. FIR continuum foregrounds are believed to be removable with negligible residuals via spectral decomposition (Yue et al., 2015). Still, the [CII] line suffers from bright low-z line interlopers. For example, carbon monoxide (CO) molecules residing at low redshifts can undergo spontaneous rotational transitions, emitting photons that redshift into the same observed frequency band as the photon emitted from the [CII] transition at high redshift. Blind and guided masking techniques have been explored to remove voxels of data thought to be dominated by interlopers (Yue et al., 2015; Visbal et al., 2011; Gong et al., 2014; Silva et al., 2015; Breysse et al., 2015; Sun et al., 2018; Silva et al., 2021). Masking, however, brings about its own host of challenges since it complicates the survey geometry and alters the bias and amplitude of shot noise due to the down-sampling of the intrinsic line-luminosity function (Bernal and Kovetz, 2022). Other interloper mitigation strategies include line identification (Moriwaki et al., 2024; Kogut et al., 2015), cleaning using external tracers of large-scale structure (Bernal and Baleato Lizancos, 2024), analysis of redshift space distortions (Lidz and Taylor, 2016; Liu et al., 2016; Cheng et al., 2020), spectral deconfusion, and of course the subject of this paper, cross-correlations (Gong et al., 2012; Roy and Battaglia, 2024; Gong et al., 2014; Visbal and Loeb, 2010).

While the prospects of measuring a cross-correlation signal are exciting, current-generation experiments are not optimized for joint analyses. In Figure 2.1, the 21cm-[CII] cross spectrum is plotted as a function of $k_\perp$ (spatial wavenumber perpendicular to the line of sight) and $k_\parallel$ (wavenumber parallel to the line of sight) for the fiducial reionization scenario described in Gong et al. (2012). As is evident, on large scales the fields are anti-correlated, resulting in negative power shown in blue. On small scales the fields become positively correlated, resulting in positive power shown in red. To help guide the eye, a yellow line has been plotted at the cross-over scale where the power spectrum transitions from negative to positive at $k \sim 1.5\,h\text{Mpc}^{-1}$. Of course, there remains considerable uncertainty as to how reionization occurred, and as a contrasting case, we also show the cross-over scale of $k \sim 0.03\ h\mathrm{Mpc}^{-1}$ for the high-mass reionization scenario described in Dumitru et al. (2019).

In addition to these theory markers, we plot the Fourier modes accessible by current generation 21 cm and [CII] instruments, HERA (DeBoer et al., 2017; Berkhout et al., 2024) and CCAT-prime (CCAT-Prime Collaboration et al., 2023) respectively. These instruments will serve as our fiducial instruments for the forecasts done in this paper. The modes that are jointly accessible by the two instruments, assuming a 2×2 square-degree overlap in sky coverage, are shown in the cross-hatched region. These are the only modes for which a cross-spectrum can be estimated with upcoming data and, as is evident, the cross-over scale of the Gong et al. (2012) fiducial model is far outside the detectable zone. Therefore, it is entirely possible that a near-future measurement may not detect the cross-over scale, although even a null result may rule out certain scenarios (e.g., the Dumitru et al. 2019 high-mass scenario).

Of course, an overlap in Fourier scales is merely a necessary, not a sufficient condition for a successful cross-correlation measurement. Our simulation pipeline is designed to be an end-to-end effort for a full evaluation of the 21 cm-[CII] cross-correlation science yield in the presence of systematics and highly limited Fourier coverage. This pipeline was assembled using our custom software package which we call LIMSTAT[1]. LIMSTAT is a statistical framework for the simulation and analysis of line intensity maps. A flowchart of this pipeline can be found in Figure 2.2. We start by generating statistically correlated 3D cubes of the 21 cm and [CII] brightness temperature fields using our built-in quick correlation model. This produces Gaussian random fields that contain the right fluctuation properties without modeling the physics. These input maps can of course be replaced with boxes from any suite of cosmological simulations. On the 21 cm side, Galactic synchrotron emission, bremsstrahlung emission, an unresolved point source background, and bright radio point sources are added to the cosmological signal map. We then simulate an interferometric observation of the final map, given that the vast majority of 21 cm cosmology instruments, like HERA, are interferometers. On the [CII] side, CO line interlopers are simulated and it is assumed

---

[1] `https://github.com/McGill-Cosmic-Dawn-Group/LIMstat`

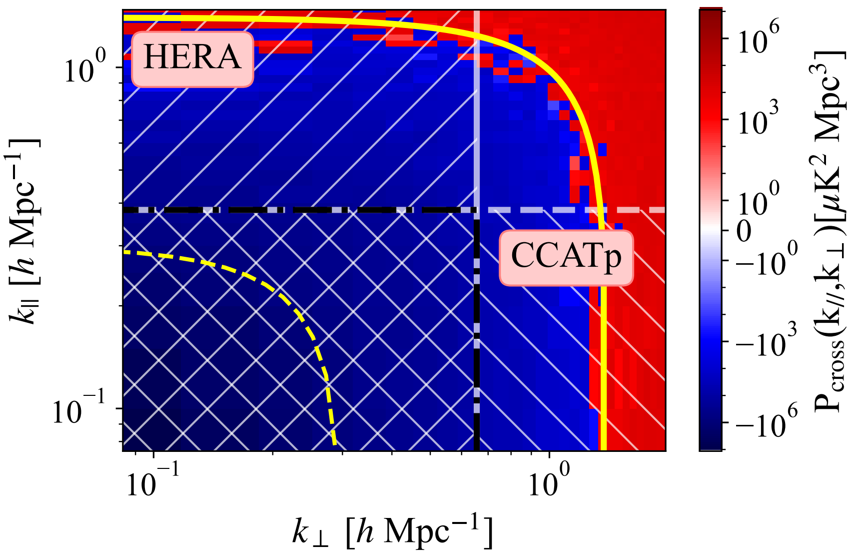

Figure 2.1: Cylindrical 21cm-[CII] cross-spectrum with theory and instrument regions overplotted. The yellow lines correspond to the theoretically predicted cross-over scale for a fiducial reionization scenario (solid yellow line at $k \sim 1.5\ h$Mpc^{-1}) and for high-mass reionization (dashed yellow line at $k \sim 0.03\ h$Mpc^{-1}). The modes accessible by HERA and CCAT-prime are boxed off with the jointly accessible modes visible in the cross-hatched region. With only limited Fourier space overlap, cross-correlation measurements may suffer from sensitivity limitations in addition to systematics.

that continuum emission has been removed. Observation of this interloper-contaminated map by a single dish instrument is then simulated. Finally, these observations are combined to produce cross-spectra.

With the exception of the modelling of Galactic synchrotron emission, which is done using the PYGSM package (de Oliveira-Costa et al., 2008; Zheng et al., 2016), each module of our pipeline was custom built from the ground up for the express purpose of producing realistic simulations of cross-correlations. That being said, the LIMSTAT software package is modular and the user has the complete freedom to use any component individually depending on the task at hand. In addition, while the surveys used here are HERA and CCAT-prime, our framework is a general one that can

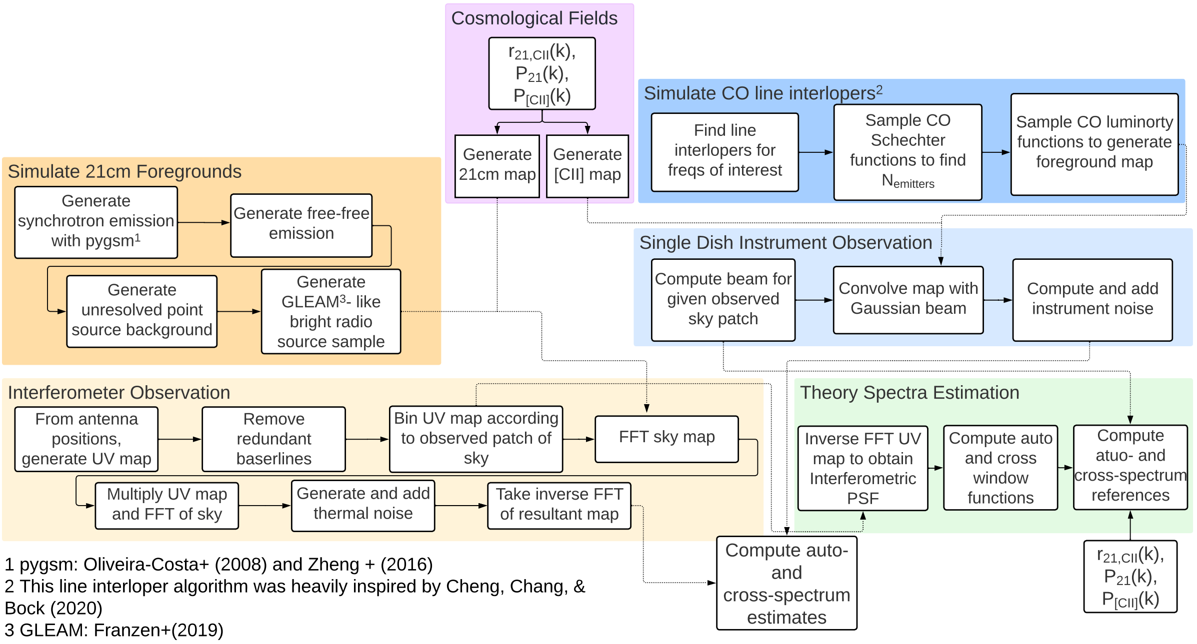

1 pygsm: Oliveira-Costa+ (2008) and Zheng + (2016)
2 This line interloper algorithm was heavily inspired by Cheng, Chang, & Bock (2020)
3 GLEAM: Franzen+(2019)

Figure 2.2: Flowchart depicting our end-to-end cross-correlation pipeline. Processes run from top to bottom and left to right. This pipeline produces correlated cosmological fields, simulates 21 cm foreground contaminants, simulates CO line interlopers (i.e. [CII] foregrounds), models both single dish instruments and interferometers, and finally estimates auto- and cross-spectra.

be used to model any single dish or interferometric instrument; all of the input parameters can be changed to the design specifications of the instrument one is interested in modelling. In the following three sections, each component of the pipeline is discussed in detail.

2.3 Generating Cosmological Fields

In this section, we present how the cosmological signals are modeled in our pipeline. We implement a quick signal simulator that produces maps of two cosmological signals that exhibit accurate cross-correlations. In addition, the individual output maps, on their own, exhibit their respective

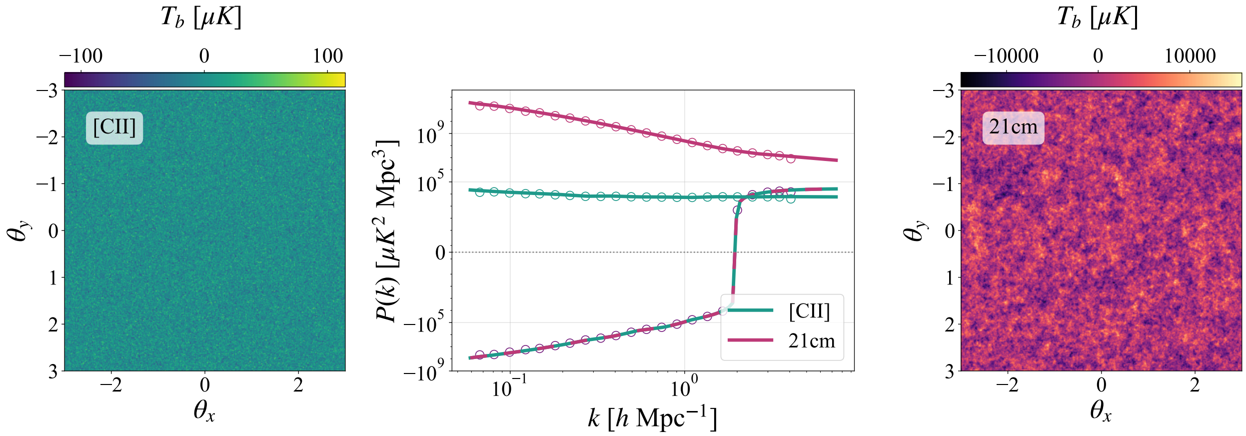

Figure 2.3: Simulated [CII] field (left) and 21 cm field (right) over 6 × 6 degrees with 1' resolution. In the middle panel, we plot the input theory auto-spectra for 21 cm (solid pink) and [CII] (solid turquoise) and also the input theory cross-spectrum (solid alternating pink and turquoise). The spectra estimated from the simulation cubes are plotted with hollow circle markers for 21 cm (pink), [CII] (turquoise), and their cross (purple). The correspondence between the spectra from simulated cubes and the input theory spectra validates our approach for generating cosmological realizations with the correct statistical properties.

auto-spectrum statistics. This is purely a statistical model; there is no underlying physics in these simulations beyond what is encoded in the theoretical auto- and cross-spectra of these lines being simulated. This implies, for instance, our simulated fields will not contain any of the non-Gaussianities that are expected in reality. However, we again emphasize that this step of the pipeline can be replaced by more realistic cosmological simulations should parameter estimation be of interest.

Here we generalize the decorrelation parameter formalism from Pagano and Liu (2020) in order to simulate fields that exhibit the correct scale-dependent correlations. The methodology is as follows. First, we produce the 21 cm field as a Gaussian random field with mean zero and power spectrum $P_{21}(k)$. The Fourier space 21 cm field $\tilde{T}_{21}(\mathbf{k})$ is then used to generate the Fourier space

[CII] field via the relation[2]

$$\tilde{T}_{[\mathrm{CII}]}(\mathbf{k}) = f(k)\tilde{T}_{21}(\mathbf{k})e^{-i\phi(\mathbf{k})}, \tag{2.4}$$

where

$$f(k) \equiv \sqrt{\frac{P_{[\mathrm{CII}]}(k)}{P_{21}(k)}}, \tag{2.5}$$

and $\phi(\mathbf{k})$ is a phase that is drawn randomly from a Gaussian distribution with zero mean and standard deviation $\sigma(k)$ (to be specified later). The $f(k)$ factor ensures that the [CII] field will have the [CII] auto-spectrum, $P_{[\mathrm{CII}]}(k)$. However, the randomly drawn phase means that when $\tilde{T}_{[\mathrm{CII}]}(k)$ is transformed back to configuration space, its bright and dim regions will have been shifted from their original positions, effecting a (partial) decorrelation of the [CII] field away from the 21 cm field.

Positive and negative correlations are assigned by applying an overall sign flip to Equation (2.4), but the degree of correlation or anticorrelation is governed by the parameter $\sigma(k)$. If $\sigma(k) = 0$, every draw of a random ϕ in that k bin is 0. The [CII] field then remains perfectly correlated with the 21 cm field in that k bin. On the other hand, if $\sigma(k)$ is large, for each Fourier pixel in that k bin, ϕ values are essentially uniformly distributed between 0 and 2π. This completely randomizes the [CII] field with respect to the 21cm field in that k bin, resulting in uncorrelated fields. In Appendix 2.10 we show that in order for two fields labelled by a and b to have a correlation coefficient

$$r_{ab}(k) \equiv \frac{P_{ab}(k)}{\sqrt{P_a(k)P_b(k)}}, \tag{2.6}$$

one should pick

$$\sigma(k) = \sqrt{\ln(|r(k)|^{-2})}. \tag{2.7}$$

In Figure 2.3, we show an example set of simulated 21 cm and [CII] fields with $r(k)$ tuned to match that of Gong et al. (2012). This will be the fiducial model that we use throughout the paper. The theory auto-spectra and theory cross-spectrum from Gong et al. (2012) are plotted with solid

---

[2] We follow the standard Fourier convention used in cosmology, as explicitly written out in Appendix 2.10.

curves while the power spectra estimated from the simulation cubes are shown with hollow round markers. As one can see, all of the simulated cubes possess the expected statistical properties.

In summary, our simulation procedure enables the fast generation of paired Gaussian fields that are guaranteed to reproduce chosen auto- and cross-power spectra. The advantage of a fast simulator is that it allows the quick generation of an ensemble of simulations, allowing for a quantification of sample/cosmic variance errors.

2.4 Generating Foreground Contaminants

In this section, we provide an overview of the foreground models used in this pipeline. These are added to the cosmological signals generated in the previous section.

2.4.1 21 cm Foreground Model

Here we lay out how each of the 21 cm foreground contaminants are modeled in the simulation pipeline. In Figure 2.4, we show the simulated field as well as each of the 21 cm foreground components. Even though HERA and CCATp are expected to overlap over only a small $\sim 4\,\mathrm{deg}^2$ patch of sky (the Chandra Deep Field South), we simulate a larger $\sim 36\,\mathrm{deg}^2$ patch. The motivation for doing so is that many radio interferometers designed for 21 cm have large fields of view (see Section 2.5.1) and point spread functions (PSFs[3]) that have side lobe structures far away from the central peak. Although these side lobes are generally orders of magnitude lower in amplitude than the central peak, a sufficiently bright point source caught in a side lobe can still leak in from the edges of a field to the center (Pober et al., 2016; Xu et al., 2022; Burba et al., 2023). In order to take this effect into account, we simulate a larger field but perform our cross-correlation analysis on a small analysis field.

---

[3] Also known as synthesized beams in the radio interferometry parlance.

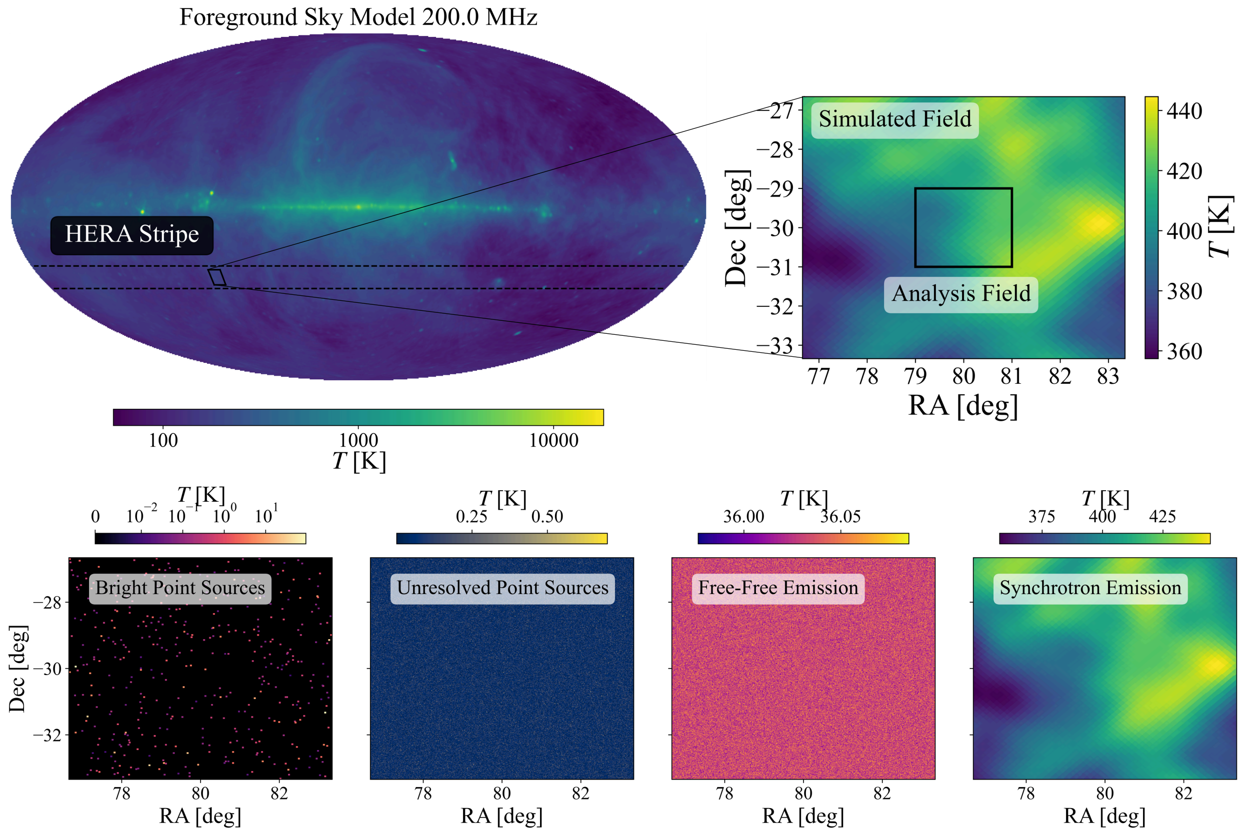

Figure 2.4: Top: Example 21 cm foregrounds from our simulation pipeline at 200 MHz. The full-sky map shows the Galactic synchrotron contribution with the HERA observing stripe superimposed. Each observation is simulated over a $6 \times 6\,\mathrm{deg}^2$ simulation patch to capture bright foreground sources caught in low-level sidelobes of an instrument's PSF, in contrast to the smaller $2 \times 2\,\mathrm{deg}^2$ analysis patch. Bottom: Example realizations from different constituent components of 21 cm foregrounds described in Section 2.4.1.

Galactic Synchrotron Emission

Galactic synchrotron radiation is the brightest of the diffuse 21 cm foreground contaminants. In this work, galactic synchrotron radiation is simulated using the PYGSM package (de Oliveira-Costa et al., 2008; Zheng et al., 2016). Since there do not exist all-sky maps of our galaxy at all frequencies from observation, this package interpolates over gaps in coverage using a set of principal components that are trained on 29 sky maps between 10 MHz and 5 THz. For the forecasts done here we use the PYGSM2008, but PYGSM2016 is also compatible with this pipeline.

Beyond an example map of the synchrotron sky, it is necessary to have a prescription for the statistical distribution from which the map is drawn. This is because one of the goals of this paper is to quantify the increased uncertainty on a power spectrum measurement that is due to foregrounds that survive cross-correlation. Unfortunately, there does not exist a first-principles method to capture the strongly correlated non-Gaussian nature of synchrotron foregrounds. Our approach is therefore to construct an empirical distribution by randomly selecting from different patches of the sky, each of which is the same size as our original simulated field, as illustrated in Figure 2.5. In fact, this is arguably the more realistic approach, since the strategy of using cross-correlations to mitigate foregrounds is one that implicitly relies on different parts of the sky being different random draws that have different phases in Fourier space. In other words, the relevant probability distribution is the probability distribution of foregrounds in *our* Galaxy, not the probability distribution of a generic synchrotron process in our Universe (Tan et al., 2021).

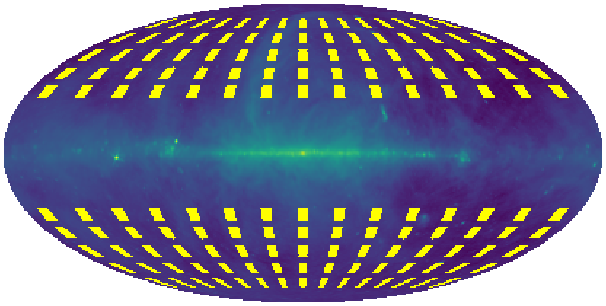

Figure 2.5: Patches of the sky (yellow) from which Galactic synchrotron foregrounds are randomly drawn to form an ensemble of foreground realizations.

Free-Free Emission

The next contaminant to be modeled is Bremsstrahlung or free-free emission from ionized gas mostly originating from the local ISM. Here we present a simplified model of free-free emission and while this does not capture the morphology, it encapsulates the statistical and spectral properties of this bright diffuse radio foreground. For the free-free emission model, each pixel is assigned a brightness temperature T according to the power law spectrum

$$T(\nu) = A_{\rm ff} \left(\frac{\nu}{\nu_*}\right)^{-\alpha_{\rm ff}}, \tag{2.8}$$

where here, $A_{\rm ff} = 33.5$ K and $\nu_* = 150$ MHz. A spectral index $\alpha_{\rm ff}$ is independently drawn for each pixel from a Gaussian distribution, with mean $\overline{\alpha}_{\rm ff} = 2.15$ and standard deviation $\Delta\alpha_{\rm ff} = 0.01$ (Wang et al., 2006).

Radio Point Sources

Next we consider extragalactic point sources, of which there are two main populations: bright radio point sources and unresolved point sources. It should be noted that a third sub-category also exists, bright extended sources, that is, bright and nearby extragalactic sources that have a spatial extent greater than a single pixel. For 21 cm observations in the southern hemisphere, an example of such a source is Fornax A. This sub-category of bright sources is not included in our 21 cm foreground models but a treatment of extended sources can be found in Line et al. (2020).

For both populations of point sources, we model a synthetic point source catalog with the same statistics as the Galactic and Extragalactic All-sky MWA (GLEAM; Hurley-Walker et al. 2017) survey based on Franzen et al. (2016, 2019). We draw the number of sources per pixel n from a source count distribution given by

$$\log_{10}\left(S^{2.5}\right)\frac{dn}{dS} = \sum_{i=0}^{5} a_i \left(\log_{10} S\right)^i, \tag{2.9}$$

where the parameters a_i={3.52, 0.307, -0.388, -0.0404, 0.0351, 0.006} are the best-fit parameters for the GLEAM source count distribution and S is the source flux. This source count distribution is valid for frequencies between 72 MHz and 231 MHz (Franzen et al., 2019). The source spectra were obtained using a similar expression to Equation 2.8,

$$S(\nu) \propto \left(\frac{\nu}{\nu_*}\right)^{-\alpha_{\rm ps}} \tag{2.10}$$

where the spectral index $\alpha_{\rm ps} = 0.8$ and the reference frequency $\nu_* = 150$ MHz.

In order to simulate both unresolved and bright extragalactic sources, we simulate two separate point sources maps with different flux limits. For the unresolved point source map, we draw sources with fluxes less than 100 mJy since most peeling techniques can only remove sources whose flux is greater than ~ 10 to 100 mJy. Of course, the exact flux cutoff depends on the resolution and sensitivity of the instrument since, for example, an instrument with lower resolution will smear more sources into the unresolved background. As in Liu and Tegmark (2012), the more conservative bound of 100 mJy is used here. For the bright point source map, we draw sources with fluxes in the range $100\,\mathrm{mJy} < S \lesssim 300\,\mathrm{Jy}$.

2.4.2 [CII] Foreground Model

Line interlopers are the most problematic foreground contaminant for [CII] observations and this is therefore where we focus our simulation efforts. Carbon monoxide (CO) at low redshifts undergoes spontaneous rotational transitions, emitting a photon that redshifts into the same frequency band as the high-redshift [CII] line. In particular, CO (6-5), CO (5-4), CO (4-3), CO (3-2), CO (2-1), are line interlopers for [CII] observations during the EoR, where the pairs of numbers denote quantum numbers, J, indicating the change in the total angular momentum state of the molecule. To simulate all of these CO lines, we follow a similar prescription to Cheng et al. (2020). The main difference with the method employed here is that we approach the modelling from an observational point of view, which has the added benefit of decreasing the computational cost. Instead of building a large and dense lightcone, we instead populate the spectral channels of the instrument with the lines that will have redshifted into that frequency channel. Currently, each channel's interloper map is computed independently; therefore, any line-of-sight interloper correlations are not simulated. In addition, our interlopers do not exhibit any clustering.

In this foreground model, we populate each pixel with individual CO sources and then draw a luminosity for each source from the Schechter luminosity function of each CO line. The Schechter luminosity function Φ for CO luminosities L_{CO} is given by

$$\Phi(L_{\mathrm{CO}})dL_{\mathrm{CO}} = \phi_* \left(\frac{L_{\mathrm{CO}}}{L_*}\right)^{\alpha} e^{-(L_{\mathrm{CO}}/L_*)} d(L_{\mathrm{CO}}/L_*), \tag{2.11}$$

and is characterised by the Schechter parameters ϕ_*, L_*, and α. The quantity ϕ_* is the normalization density, L_* is a characteristic luminosity, and α is the power-law slope at low luminosity. The Schechter parameters for the various CO lines as well as the [CII] line used here can be found in Popping et al. (2016).

We begin by determining which lines redshift into the observed frequency channel and then proceed to interpolate the Schechter parameters for those lines in order to find all of the parameters at z_{emit}, the redshift at which the line was emitted. Next we proceed to use the Schechter luminosity function to populate the luminosity bins of each line with sources. Having obtained the quantity of interest, $\Phi(L_{\mathrm{CO}})dL_{\mathrm{CO}}$ given by Equation 2.11, which is the number of galaxies with luminosity $L_{\mathrm{CO}} + dL_{\mathrm{CO}}$ per unit volume, we must adjust the number density for the physical size of the voxel that has been simulated. We do this by simply multiplying by the voxel volume

$$n(L_{\mathrm{CO}})dL_{\mathrm{CO}} \equiv \Phi(L_{\mathrm{CO}})dL_{\mathrm{CO}}V_{\mathrm{vox}} \tag{2.12}$$

where V_{vox} is

$$V_{\mathrm{vox}} = dV_{\mathrm{vox}}\Omega_{\mathrm{pix}}, \tag{2.13}$$

with Ω_{pix} being the solid angle subtended by a single voxel in our survey and dV_{vox} being the comoving volume per unit solid angle. The quantity $n(L_{\rm CO})dL_{\rm CO}$ is the average number of sources per voxel in each luminosity bin. To ensure there is source count variation from voxel to voxel, we draw the number of sources per luminosity from a Poisson distribution with mean $n(L_{\rm CO})dL_{\rm CO}$. The final luminosity is assigned by multiplying the number of sources in each voxel by the luminosity of that bin and then summing over all luminosity bins. Lastly, once the luminosity maps for each line are obtained, they are combined into a single foreground map that has contributions from all line interlopers in a particular channel. In Figure 2.6, the characteristic intensity,[4] I_*, of the various CO lines as well as the [CII] line is shown as a function of redshift and observed frequency. This serves to highlight just how foreground dominated high-z [CII] measurements are.

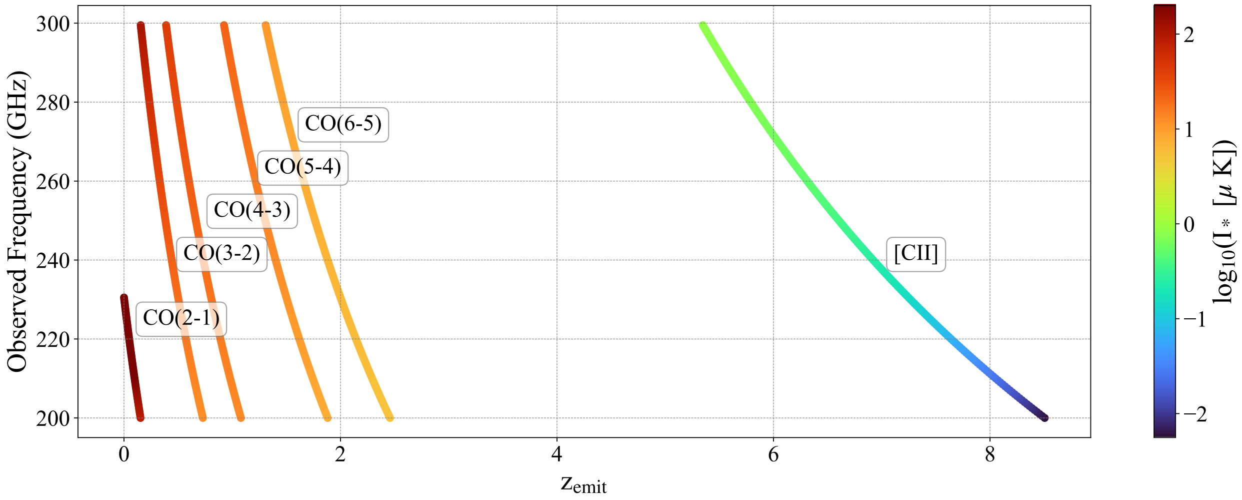

Figure 2.6: Brightness of [CII] and relevant CO emission lines, as functions of emission redshift $z_{\rm emit}$ and observed frequency. At all observed frequencies considered in this paper, CO interlopers will be a strong contaminant to [CII] observations and at the highest redshifts ($z > 6$) the CO(2-1) transition is a particularly strong nuisance.

One potential concern for cross-correlation-based foreground mitigation strategies is the possibility of correlated foregrounds. For example, it is possible that some of the [CII] line interlopers may in fact be emanating from the same galaxies at $0 \lesssim z \lesssim 3$ that are responsible for

---

[4] This is obtained through a simple conversion of the characteristic luminosity, L_*.

21 cm point source foregrounds. While this is certainly not impossible, it is not likely. There do exist a handful of hydrogen and deuterium fine structure lines whose rest frequencies are such that if they came from the same $0 \lesssim z \lesssim 3$ galaxies as the line interlopers, they would redshift into the observed 21 cm frequency range (Fabjan et al., 1971; Kramida, 2010). That being said, the transition rate of these lines is unknown and the transitions have only been observed under strict laboratory conditions. Still, one may consider broadband emission that could produce these correlation, though we leave such a scenario for future work. For the simulations done here, we generate realizations of radio point source foreground contaminants and CO line interlopers independently.

2.5 Simulating Instrument Response

With foregrounds added to our simulated cosmological fields, we proceed to simulating the instruments that observe these foreground-contaminated skies. For concreteness, the conclusions of this paper will be demonstrated using simulations of HERA and CCAT-prime, but we stress that our pipeline is generally applicable to any combination of interferometers and/or single dish experiments.

2.5.1 Radio Interferometers

In order to simulate both the response and the noise of the HERA array, we compute a set of noisy visibilities. Interferometers measure a primary beam- and fringe-weighted integral of the sky intensity known as the visibility, V_{ij}, defined as

$$V_{ij}(\nu) = \int A_{ij}(\hat{\mathbf{s}}, \nu) I(\hat{\mathbf{s}}, \nu) \exp\left(-2\pi i \frac{\mathbf{b}_{ij} \cdot \hat{\mathbf{s}}}{\lambda_{\text{obs}}}\right) d\Omega \tag{2.14}$$

where $d\Omega$ is the differential solid angle, and $\mathbf{b}_{ij}$ is the baseline vector that characterizes the separation and orientation of the ith and jth receiving elements (such as dishes). The unit vector $\hat{\mathbf{s}}$

points to the direction of the incoming radiation on the sky. The observing wavelength is denoted by $\lambda_{\rm obs}$, $I(\hat{\mathbf{s}}, \nu)$ is the specific intensity, and $A_{ij}(\hat{\mathbf{s}}, \nu)$ is the geometric mean between the primary beams of the ith and jth elements. Equation (3.15) can be compared to a two-dimensional Fourier transform $\tilde{I}$ of the specific intensity, namely

$$\tilde{I}(\mathbf{u}, \nu) = \int I(\boldsymbol{\theta}, \nu) \exp\left(-i2\pi \mathbf{u} \cdot \boldsymbol{\theta}\right) d^2\theta, \tag{2.15}$$

where we have invoked the flat-sky approximation to describe positions on the sky in terms of Cartesian angular coordinates $\boldsymbol{\theta} \equiv (\theta_x, \theta_y)$ and have defined a Fourier dual $\mathbf{u} \equiv (u, v)$ to this. One sees that in the limit that the sky is flat and the primary beam is reasonably uniform, each baseline of an interferometer measures a single Fourier mode in the plane of the sky with wavenumbers $u = b_x/\lambda_{\rm obs}$ and $v = b_y/\lambda_{\rm obs}$ on the uv plane, where b_x and b_y are the x and y components of $\mathbf{b}_{ij}$, respectively. These modes can then be easily mapped to the comoving wavevector perpendicular to the line of sight, $\mathbf{k}_\perp$, by

$$\mathbf{k}_\perp = \frac{2\pi \mathbf{b}_{ij}}{\lambda_{\rm obs} D_c(z_{\rm emit})} \tag{2.16}$$

where $D_c(z_{\rm emit})$ is the comoving distance to the source emission. Since our analysis field is much smaller than the area covered by HERA's primary beam (see Table 2.1), we omit it in our simulations. In other words, we assume uniform sensitivity over our 4 deg^2 analysis field.

To simulate the action of an interferometer, we start by denoting the sky at a particular frequency as $\mathbf{m}$ which stores all of the true pixel intensities on the sky. The measured visibilities then form a map $\mathbf{v}$ in the uv plane given by

$$\mathbf{v} = \mathbf{D}\mathcal{F}\mathbf{m} + \mathbf{n} \tag{2.17}$$

where $\mathcal{F}$ denotes the 2D Fourier transform in the angular directions, $\mathbf{D}$ is a binary mask of the uv coverage (recording which uv modes are measured or missed based on the baselines present in an interferometer), and $\mathbf{n}$ is a noise realization in uv space. To generate a noise realization, we populate each pixel in uv space with a complex Gaussian random number with zero mean and standard deviation given by

$$\sigma_{\mathrm{rms}}(u,v) = \frac{T_{\mathrm{sys}}}{\sqrt{N_{\mathrm{red}} t_{\mathrm{obs}} \Delta\nu}} \Omega_{\mathrm{surv}}, \tag{2.18}$$

where it is understood that the real and imaginary components are drawn independently and half of the variance is in each component. Here, T_{sys} is the system temperature, t_{obs} is the total integration time, $\Delta\nu$ is the bandwidth, and Ω_{surv} is the survey area. The quantity N_{red} is the number of baselines that fall into a particular uv cell, and accounts for the fact that HERA is a highly redundant array (see the top panel of Figure 2.7 and Dillon and Parsons 2016) with multiple identical copies of the same baseline. Therefore the noise of any given measurement in the uv-plane is reduced by a factor of $1/\sqrt{N_{\mathrm{red}}}$. The result of this is that the visibilities of short baselines, which are measured many times over, will generally be less noisy than the visibilities of long baselines that are less redundant.

The instrument specifications used to generate HERA-like noise are listed in Table 2.1. It should be noted that in order to match existing HERA noise estimates in the literature, the observing time was reduced by a factor of six. This is to account for the fact that Equation (2.18) assumes that all data collected over the duration t_{obs} is coherently averaged. In reality, HERA is a non-pointing drift-scan telescope where target fields can only be observed for short durations per sidereal day, and thus the coherent averaging down of noise can only happen in short bursts. Observing different fields will also reduce noise, but only via an incoherent average (at the power spectrum stage). Accounting for this is complicated, given that it crucially depends on simulating rotation synthesis effects and partial redundancy between baselines (Zhang et al., 2018), accounting for subtleties associated with having a drift-scan telescope (Liu and Shaw, 2020). While this can be done, it would lead to a slight inconsistency with our approach of using different portions of the Galaxy to represent the ensemble properties of possible foregrounds: as Figure 2.5 illustrates, these patches are distributed across the entire sky (excepting the Galactic plane), and in principle would require multiple arrays deployed at different latitudes for all of them to be observable. For simplicity, we thus forgo a detailed rotation synthesis by scaling the integration time to account for the mix of coherent and incoherent integration. We calibrate this scaling so that our code reproduces the 21 cm power spectrum sensitivities published in Pober et al. (2014) when matching their instrumental configurations.

As another approximation, we use a simple fit for T_{sys} that is given in Table 2.1. Since T_{sys} is partly determined by the receiver temperature ($\sim 100\,\mathrm{K}$) and partly by the sky temperature, it in principle changes when observing the different parts of the sky illustrated in Figure 2.5. We neglect this variation when computing the noise contribution to our observations, finding that the power-law fit $60\lambda_{\mathrm{obs}}^{2.55}$ to be an excellent fit to the mean temperature of our foreground maps. To be clear, this approximation is used only in the determination of noise amplitude in Equation (2.18); the foreground signals that are present in simulated observations are of course specific to each patch of the sky.

Finally, with our simulated visibilities $\mathbf{v}$, we perform an inverse Fourier transform to obtain our observed HERA map. In radio astronomy parlance, this would constitute a dirty map of the observations, i.e., one where the point spread function (PSF) is not deconvolved from the map. In the bottom panel of Figure 2.7, we show HERA's PSF (obtained by Fourier transforming the uv plane's binary mask $\mathbf{D}$). One clearly sees hexagonal features due to the geometry of the array layout. We choose not to remove the effects of the PSF since the blurring effect of the PSF is not a formally invertible operation. Instead, we account for the PSF downstream in our data analysis (Section 2.6) via the normalization of our power spectra.

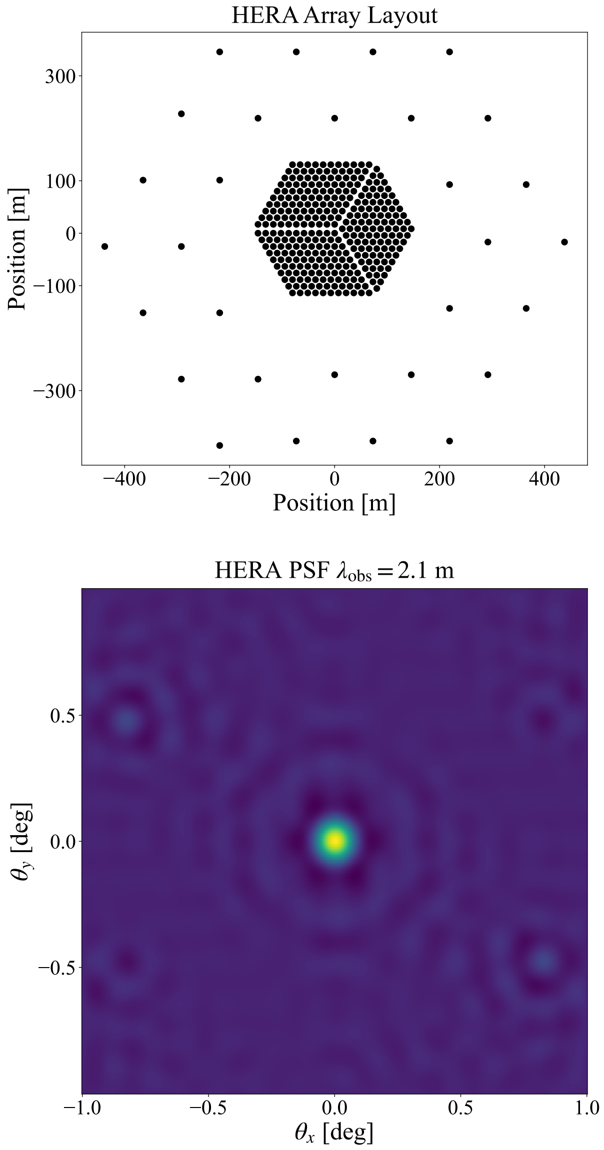

Figure 2.7: Top: HERA array layout assumed for the forecasts in this paper. Bottom: the resulting point spread function.

Paramters	HERA
System Temperature, $T_{\rm sys}$	$100 + 60\lambda_{\rm obs}^{2.55}$ K
Beam FWHM at 150 MHz, $\theta_{\rm FWHM}$	8.7 deg
Element Diameter	14 m
Shortest Baseline	14.6 m
Longest Baseline (core)	292 m
Longest Baseline (outrigger)	876 m
EoR Frequency Range, $\nu_{\rm obs}$	100-200 MHz
Channel Width, $\delta\nu$	97.8 kHz
Integration time on field, $t_{\rm obs}$	170** hours

Table 2.1: Parameters for the HERA array assumed in this paper. These parameters are based on DeBoer et al. (2017) and Pober et al. (2014). **This observing time was further reduced by a factor of six to match existing HERA sensitivity estimates (see Section 2.5.1)

2.5.2 Single Dish Instruments

The other instrument that we simulate in this paper is CCATp, which is a single dish telescope. This instrument is modeled by convolving the sky with a 2D Gaussian beam with standard deviation equal to the diffraction limited angular resolution, that is, $\theta_{\rm CCAT} \sim \lambda_{\rm obs}/D$ where D is the diameter of the dish. After the field has been convolved, a noise realization is added. This noise map consists simply of white noise drawn from a Gaussian with mean zero and standard deviation $\sigma_{\rm rms}$, which is given by

$$\sigma_{\rm rms} = \frac{\sigma}{\sqrt{t_{\rm pix}}} \tag{2.19}$$

where $t_{\rm pix}$ is the total integration time per pixel, and

$$\sigma = \frac{\sigma_{\rm pix}}{\sqrt{\Omega_{\rm sim}/\Omega_{\rm pix,CCAT}}} \tag{2.20}$$

where $\sigma_{\rm pix}$ is the quoted per pixel sensitivity of CCATp, $\Omega_{\rm sim}$ is the size of a pixel in the simulated cube, and $\Omega_{\rm pix,CCAT} = \pi\theta_{\rm FWHM}^2/4\ln 2$ is the nominal pixel size of CCATp observations. Note that

Parameters	CCATp
System temperature, σ_{pix} MJy sr^{-1} s$^{1/2}$	0.86*
Beam FWHM, θ_{FWHM} (arcmin)	0.75*
Dish Diameter (m)	6
Frequency Range, ν_{obs} (GHz)	210-300
Channel Width, $\delta\nu$ (GHz)	2.5
Number of Detectors, N_{det}	20
EoR Survey Area, Ω_{surv} (deg^2)	4

Table 2.2: Parameters for [CII] survey CCATp assumed in this paper. These parameters are based on Breysse and Alexandroff (2019), Chung et al. (2020). It should be noted that the parameters with the superscript ($*$) are frequency-dependent quantities and the values in the table were computed at 237 GHz.

despite the confusing notation, σ is not a standard deviation. This can be seen by examining its dimensions in Table 2.2. There, we summarize the instrument specifications for CCATp that are used in this paper, which are based on parameters given in Chung et al. (2020); Breysse and Alexandroff (2019).

2.6 Simulating a Data Analysis Pipeline

After simulating our cosmological signals (Section 2.3), adding foregrounds (Section 2.4), and putting the resulting sky through simulated instruments (Section 2.5), a crucial last component of an end-to-end analysis is to include a mock data analysis. Nuanced data analysis choices affect one's final uncertainties on a cross-power spectrum and in this section we make our choices explicit.

2.6.1 Power Spectrum Estimation

After Section 2.5, we have our instrument-processed three-dimensional surveys. These are the inputs to our cross power spectrum estimator $\hat{P}_{ab}$. As a first step in our power spectrum estimation, we multiply both surveys by a Blackman-Harris taper function B (Harris, 1978) in the radial direction, as is often done in 21 cm analyses (Abdurashidova et al., 2022a; HERA Collaboration et al., 2023; Abdurashidova et al., 2022b; Thyagarajan et al., 2013, 2016). This is done to prevent harsh non-periodic discontinuities at the edges of one's data cubes, which occur because the 21 cm foregrounds are considerably brighter on the low-frequency end of one's survey than the high-frequency end. Leaving such discontinuities in the data causes ringing upon a Fourier transform, which undesirably scatters foreground contaminants from the spectrally smooth modes at low $k_{\parallel}$ to spectrally unsmooth modes at high $k_{\parallel}$ that should otherwise be reasonably contaminant-free. Denoting our post-tapering survey data to be $T^{\text{tap}}(\mathbf{r})$, we then employ an estimator that is very similar to one that would be used for perfect theoretical data, namely one where the surveys are Fourier transformed, cross-multiplied, and normalized to form

$$\hat{P}_{ab}(\mathbf{k}) \equiv \frac{\tilde{T}_a^{\text{tap}}(\mathbf{k})\tilde{T}_b^{\text{tap}}(\mathbf{k})^*}{VN(\mathbf{k})}, \tag{2.21}$$

where V is the survey volume and $N(\mathbf{k})$ is a normalization factor derived in Appendix 2.10. This is then cylindrically binned by folding $\pm k_z$ into a single $k_{\parallel}$ coordinate and averaging over rings of constant $k_{\perp} \equiv (k_x^2 + k_y^2)^{1/2}$ to form $\hat{P}_{ab}(k_{\perp}, k_{\parallel})$. In some (but not all) of the scenarios that we examine in Sections 2.7 and 2.8, this is the stage at which an extra foreground mitigation step is included. This step is the excision of the 21 cm foreground wedge (Liu et al., 2016; Parsons et al., 2012; Morales et al., 2012; Trott et al., 2012; Pober et al., 2013; Dillon et al., 2014; Liu et al., 2014a; Chapman et al., 2016), which consists of all Fourier modes satisfying

$$k_{\parallel} \lesssim k_{\perp} \frac{H_0 D_c E(z)}{c(1+z)} \tag{2.22}$$

where $E(z) \equiv \sqrt{\Omega_\Lambda + (1+z)^3\Omega_m}$, D_c is comoving line-of-sight distance to the midpoint of the survey volume, H_0 is the Hubble parameter, Ω_m is the normalized matter density, and Ω_Λ is the normalized dark energy density. Following (possible) foreground wedge excision, we spherically bin our power spectrum in rings of constant $k \equiv (k_\perp^2 + k_\parallel^2)^{1/2}$ to obtain our final estimator of the isotropic power spectrum $\hat{P}_{ab}(k)$.

Our power spectrum estimation methodology was chosen for its simplicity, in order to best showcase trends that are intrinsic to the nature of 21 cm-LIM cross-correlations. Our forecasts will therefore be on the conservative side in terms of sensitivity, given that we have neither used optimal frequentist techniques such as an optimal quadratic estimator (Dillon et al., 2014; Tegmark, 1997a; Bond et al., 1998; Liu and Tegmark, 2011; Shaw et al., 2015; Liu et al., 2014b), nor have we implemented a state-of-the-art Bayesian analysis pipeline (Burba et al., 2023; Zhang et al., 2016; Sims et al., 2016; Sims and Pober, 2019; Sims et al., 2019; Burba et al., 2024; Cheng et al., 2024).

2.6.2 Window Functions

Having gone through the imperfections of our simulated instruments and the various choices of our power spectrum analysis pipeline, our estimated power spectrum $\hat{P}_{ab}(k)$ will almost certainly differ in a systematic way (beyond noise fluctuations) from the true power spectrum $P_{ab}(k)$. Our pipeline therefore includes a theory module that computes the joint window function $W_{ab}(k, k')$ characterizing the joint cross power response of the two instruments. Specifically, the window function[5] relates the ensemble average of the estimated power spectrum $\hat{P}_{ab}(k)$ to the true power spectrum P_{ab} via

$$\langle \hat{P}_{ab}(k) \rangle = \int dk' W_{ab}(k, k') P_{ab}(k'). \tag{2.23}$$

---

[5] It is unfortunate that the term *window function* is an overloaded one in the literature, with multiple meanings and definitions. In digital signal processing and radio astronomy papers (e.g., Lanman et al. 2020) it is often used to denote what we define in this paper to be our *tapering function* B. In certain portions of the intensity mapping literature (e.g., Padmanabhan et al. 2022; Bernal et al. 2019a; Li et al. 2016) it quantifies the instrument's multiplicative response to the true power spectrum, accounting for depressions in power due to a telescope's finite resolution (whether spectral or angular). In our nomenclature, this would correspond to our normalization factor $N(\mathbf{k})$. The normalization factor will only coincide with our more general definition of a window function, defined in Equation (2.23), when $W_{ab}(k, k')$ can be approximated as diagonal operator [i.e., proportional to $\delta^D(k - k')$] and the window functions are assumed to be unnormalized (see Appendix 2.10 for details).

Thus, this window function serves as a transfer function that acts on a theoretical input power spectrum to give an output that can be directly compared to the output of our simulation and analysis pipelines. Note that because of the normalization factor included when we estimated our power spectra using Equation (2.21), integrating the window functions over k' from $-\infty$ to $+\infty$ at fixed k will yield unity by construction. As such, our power spectrum estimates are weighted *averages* (rather than just unnormalized weighted *sums*) over the true power. In Appendix 2.10 we derive the window functions that are computed by our pipeline and used in the interpretation of our results in Sections 2.7 and 2.8.

2.7 Forecasting HERA × CCAT-prime

Using the simulation pipeline presented in the last few sections, we forecast the yield of upcoming cross-correlation measurements from HERA and CCATp. We simulate three redshift bins centered on $z = \{6, 7, 8\}$ with a width of $\Delta z \sim 0.5$ corresponding to a bandwidth of $\sim 10\,\mathrm{MHz}$ for 21 cm observations and a bandwidth of $13.6\,\mathrm{GHz}$ for [CII] observations. In order to capture the non-Gaussian covariance of these upcoming measurements, we perform a set of Monte Carlo simulations over 140 patches of sky as shown in Figure 2.5 to build up a distribution for each measurement. We have three goals here: the first is of course the forecast itself; the second is to use the forecast as a worked example of our pipeline; the final goal is to use our full set of non-Gaussian end-to-end Monte Carlo simulations to extract useful lessons for LIM cross-correlation forecasting, although we defer a description of those results to Section 2.9.

2.7.1 Avoiding the Foreground Wedge

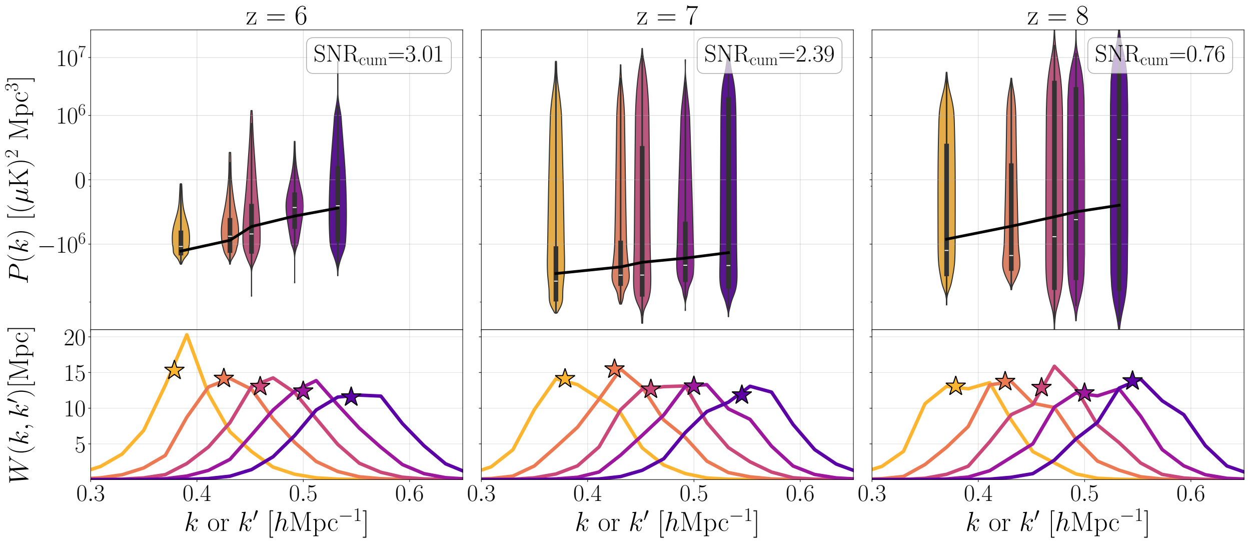

Figure 2.8: Forecasted power spectrum measurement at $z = 6$ (left), $z = 7$ (middle), and $z = 8$ (right) for the HERA × CCAT-prime scenario described in Section 2.7.1, where foreground-contaminated wedge modes are cut. Top: distributions of power spectrum estimates from our Monte Carlo ensemble, with violins showing the shapes of the distributions, the boxes giving 25th and 75th percentiles, the white bar giving the 50th percentile, and the whiskers showing the full extent of the distributions. Black lines indicate the true inputted power spectra. Note that the vertical axis is a hybrid scale that is linear for $|P(k)| \leq 10^6\,\mu$K Mpc^3 and logarithmic otherwise. The cumulative signal-to-noise ratio, computed using Equation (2.24), is also indicated. The $z = 6$ sensitivity is limited by thermal noise, whereas the $z = 7$ and $z = 8$ measurements are more severely limited by a CO (2-1) interloper. Bottom: $W(k, k')$ as a function of k' for each power spectrum estimate in matching colors. Stars indicate the k bin center used in power spectrum estimation. The non-trivial shapes of the $P(k)$ distributions and the window functions showcase the importance of forecasting via end-to-end simulations. The limited Fourier overlap between HERA and CCAT-prime results in a limited cumulative SNR for these measurements.

Given that we are forecasting measurements from HERA, we first examine a scenario that uses HERA's current foreground avoidance strategy in which modes that lie within the foreground wedge given by Equation (2.22) are discarded when forming a power spectrum (HERA Collaboration et al., 2023; Abdurashidova et al., 2022b). With our particular instrumental, observational, and data analysis assumptions here, this amounts to computing the cross-power spectrum for $k_{\parallel} > 0.29h\text{Mpc}^{-1}$ and $k_{\perp} < 0.29h\text{Mpc}^{-1}$ in 5 k-bins. Apart from the wedge cut, we did not perform any other foreground mitigation on either the 21 cm or [CII] observations. This scenario is a rather aggressive one that succeeds in avoiding a complete domination of the cross-power spectra by foregrounds, and will serve as a contrasting case when we adjudicate whether a cross-correlation on its own is capable of correlating away foregrounds, as argued in previous work.

In Figure 2.8 we show the cross-power measurements in each of the redshift bins. Our Monte Carlo ensemble enables us to understand the full distribution of each power measurement, which we portray using violin plots. The box and whiskers of each point represent the 25th and 75th percentiles and the extrema of the distributions respectively, and the short white bar within each box is the median value. The solid black line shows the expected truth obtained by evaluating Equation (2.23). Also shown are the corresponding window functions for each measured power spectrum point. Since these quantify the wavenumber range that is actually probed by each point, we center each violin at the 50th percentile of $W(k, k')$. Each window function curve corresponds to $W(k, k')$ for a fixed value of k (given by the star symbol) plotted as a function of k'.

Several generic features are immediately apparent from Figure 2.8. Broadly speaking, we see from the window functions that when evaluating our power spectrum estimator $\hat{P}_{ab}$ at some target k value (star points), one is indeed roughly probing true power at the correct wavenumber. However, the window functions are rather broad, which is unsurprising for experiments with non-trivial spectral responses (Liu et al., 2014a). Thus, it would be incorrect to assume that the power is being entirely sourced from the target k. In addition, it is not uncommon for there to be some skew to $W(k, k')$ as function of k', where the stars coincide with neither the peak nor the median. This highlights the importance of computing window functions. Another crucial feature is the presence of visibly non-Gaussian distributions in the violins (although we caution the reader that part of the visual distortion is due to the hybrid vertical scale that transitions from linear to logarithmic at

$P(k) = \pm 10^6\,\mu\mathrm{K}^2\,\mathrm{Mpc}^3$).

The most salient feature of the actual forecasting in Figure 2.8 is the high variance of the simulated measurements, especially for $z = 7$ and $z = 8$. This can be (crudely) captured by computing a total signal-to-noise ratio (SNR) that is cumulative accross the entire spectrum and is given by

$$\mathrm{SNR}_{\mathrm{cum}} = \sqrt{\langle \hat{\mathbf{P}} \rangle^t \mathbf{\Sigma}^{-1} \langle \hat{\mathbf{P}} \rangle}, \tag{2.24}$$

where $\hat{\mathbf{P}}$ is a vector (in this case with five elements) storing the power spectrum measurements for a particular Monte Carlo realization and $\mathbf{\Sigma} \equiv \langle \hat{\mathbf{P}} \hat{\mathbf{P}}^t \rangle - \langle \hat{\mathbf{P}} \rangle \langle \hat{\mathbf{P}} \rangle^t$ is the covariance between power spectrum bins, computed by ensemble averaging over our realizations. The $\mathrm{SNR}_{\mathrm{cum}}$ values are 3.01, 2.38, and 0.75 for the $z = 6$, 7, and 8 bins, respectively. In our present scenario, therefore, we have (at best) a marginal detection in the $z = 6$ and $z = 7$ bins.

To understand our lack of a highly significant detection, we can use the flexibility of our simulations to decompose our error bars into their constituent sources of uncertainty. Although the error distributions in Figure 2.8 naturally represent the total uncertainty arising from foreground residuals, cosmic variance, and instrumental noise, we can hold each contribution fixed in order to discern which source of variance poses the biggest threat to our measurements. This allows one to prioritize their mitigation strategies. We find that across all redshift bins, cosmic variance is negligible and that (as expected) the foregrounds and thermal noise are the biggest problem. At $z = 6$, the cross-correlation measurement is noise-limited. While HERA and CCATp are each sensitive instruments on their own, the lack of Fourier overlap (plus the wedge cut) leaves very few $(k_\perp, k_\|)$ bins over which to average down the noise (recall Figure 2.1). For this reason, the noise variance is comparably high across the $z = 7$ and $z = 8$ bins; however, in those two higher redshift bins, foregrounds also become a significant source of uncertainty. Although the wedge cut does an excellent job at removing 21 cm Galactic synchrotron foregrounds across all redshift bins, a bright CO (2-1) interloper line that is absent at $z \sim 6$ appears for the [CII] surveys in our $z \sim 7$ and $z \sim 8$ bins (see Figure 2.6), causing the detection significance to deteriorate.

2.7.2 Working in the Wedge

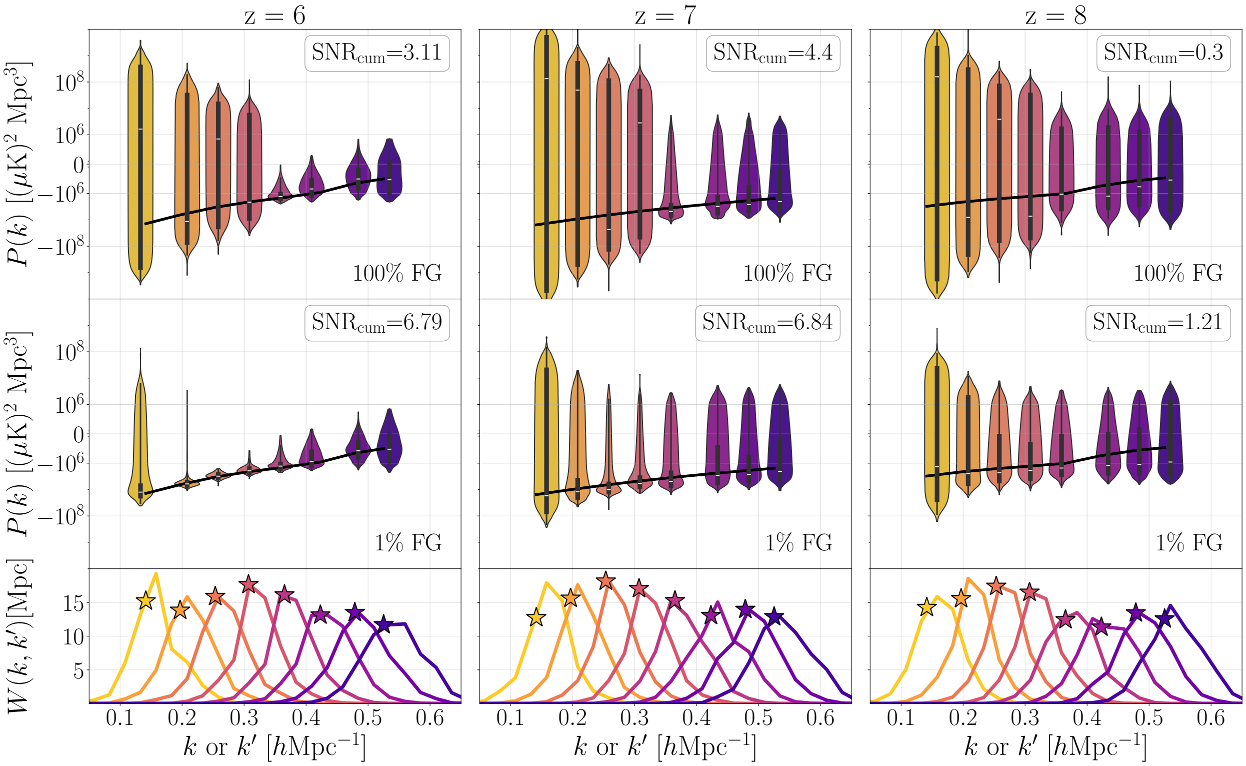

Figure 2.9: Top and bottom: same as Figure 2.8, but for the scenario described in Section 2.7.2 where 21 cm foreground wedge modes are included. Middle: same as the top, but assuming that 21 cm foregrounds have been cleaned to 1% of their original amplitude in the original maps. One sees that cross-correlations alone cannot suppress foregrounds to an acceptable level, but can be powerfully combined with other mitigation techniques to yield statistically significant detections. However, many bins that are important for a detailed characterization of the cross power spectrum remain noise limited.

Within the confines of the scenario in Section 2.7.1 (limited Fourier overlap plus a wedge cut), one ultimately needs to increase the number of independent measurements in order to increase the signal-to-noise ratio. One method for doing so is to forgo the wedge cut and to push to lower k (Chapman et al., 2016; Pober, 2015). This is an approach with trade-offs. Including lower k

modes has the effect of reducing noise variance, but the residual foreground variance increases, and one is relying more heavily on cross-correlations to reduce what starts out as higher foreground contamination in the data.

In Figure 2.9 we show the results of including modes in the wedge. In the top panel, where no foreground mitigation has been performed, the large scale modes are foreground variance dominated at all redshifts. At $z = 6$ and $z = 7$, there is only a slight increase in the cumulative SNR compared to the wedge cut case of the previous section, which highlights the fact that the increased foreground variance outweighs most of the decreased noise variance. In the $z = 8$ bin, the foreground contamination is so extreme that the significance of the detection deteriorates despite the increased number of modes. Relying purely on cross-correlations is simply not enough to suppress foregrounds to a viable level.

In the middle panel of Figure 2.9 we present the results of simulating some 21 cm foreground cleaning. To do so we run the same set of Monte Carlo simulations but reduce the 21 cm foreground maps to 1% of their original brightness. This increases the cumulative SNR significantly at $z = 6$ and $z = 7$ to 6.79 and 6.84 respectively due to the decreased foreground variance at low k values. At $z = 8$ the cumulative SNR barely breaches unity, which is still not a significant detection.

That a high-significance detection can be made at $z = 6$ and $z = 7$ having only performed 21 cm foreground mitigation is encouraging. In the 21 cm auto-correlation literature, much more extreme foreground mitigation is necessary for an auto-spectrum measurement. As a general rule of thumb, the foregrounds are 10^4 to 10^5 times brighter than the cosmological signal, necessitating the same $\sim 10^4$ to 10^5-level foreground suppression for an auto-spectrum measurement (Liu and Shaw, 2020). Here, we have shown that with 1% foreground residuals in our 21 cm maps, one is able to achieve detections with the assistance of a cross-correlation. In addition, we have performed tests where the CO interlopers were also reduced to the 1% level and found that it had a negligible effect on the cumulative SNR. This reinforces our intuition that, among the various terms in Equation (2.3), the most problematic are the ones that involve 21 cm foregrounds.

2.8 A designer's guide for the future

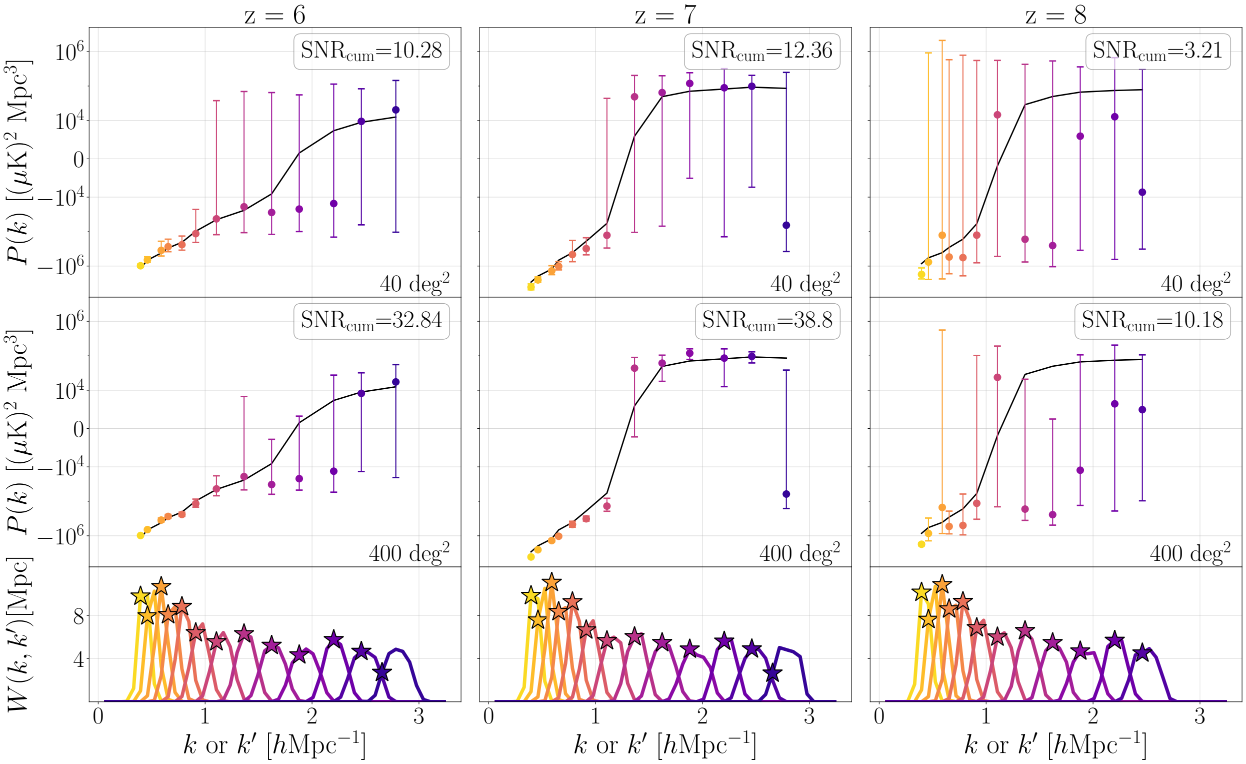

Figure 2.10: Top and bottom: same as Figure 2.8, but for the futuristic scenario described in Section 2.8 with a wedge cut and where the survey area is 40 deg^2 and the spectral resolution of the [CII] is increased. Middle: same as the top, but survey area is increased to 400 deg^2. With the increased overlap both in Fourier space (due to increased spectral resolution) and in configuration space (due to increased survey area), the detection significance rises markedly compared to current-generation surveys.

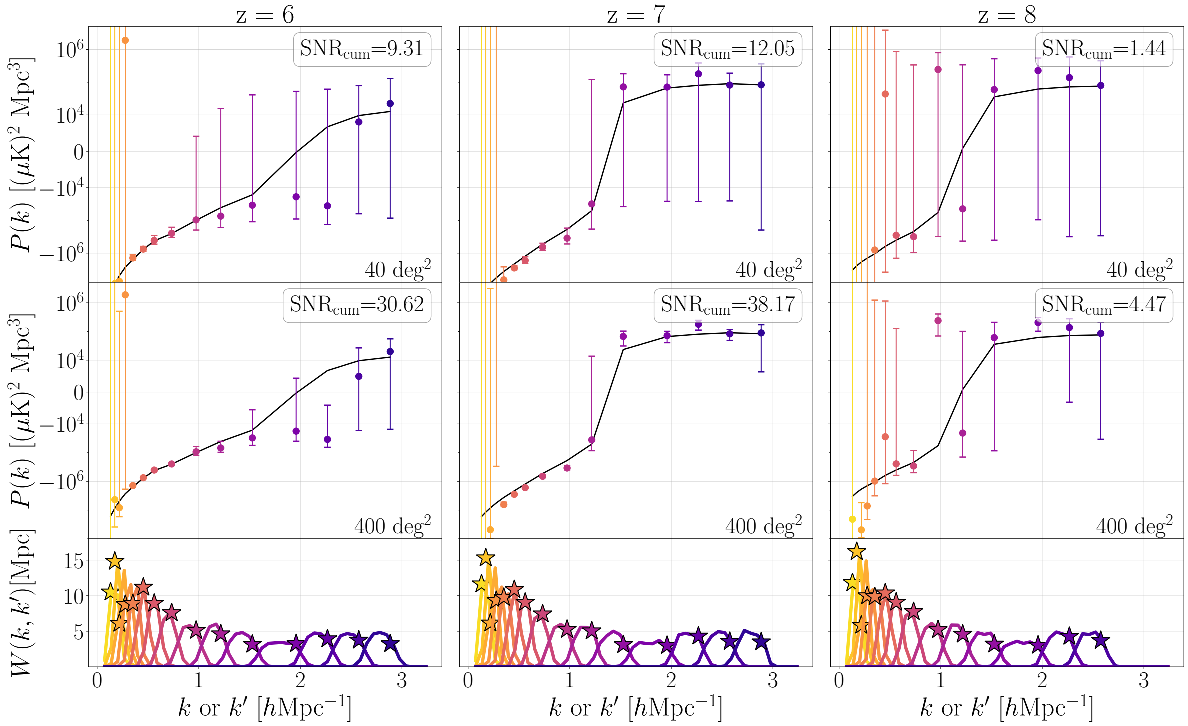

Figure 2.11: Top and bottom: same as Figure 2.9, but for the futuristic scenario described in Section 2.8 with increased survey area to 40 deg^2 and a [CII] instrument with greater spectral resolution, but no 21 cm foreground mitigation nor avoidance of the wedge. Middle: same as the top, but survey area is increased to 400 deg^2.

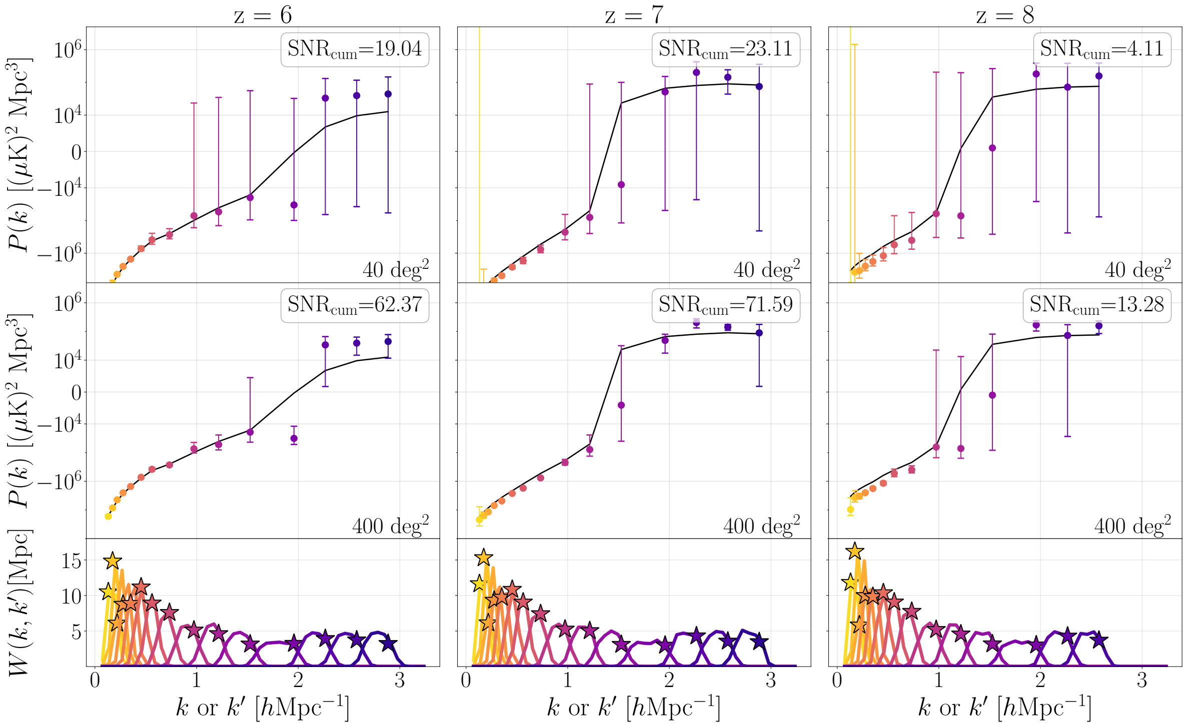

Figure 2.12: Same as Figure 2.11, but assuming that foregrounds in the 21 cm maps have been reduced to 1% of their original level. Compared to Figures 2.10 and 2.11, the extra foreground cleaning allows individual bins of the cross power spectrum to be detected to high significance, permitting a detailed characterization of the spectrum and inference of parameters such as the crossover scale from negative to positive correlation.

Figure 2.9 suggests that pushing to lower k has some utility when combined with some 21 cm foreground cleaning. A future-looking complement that we now examine is to go to wider survey areas and to increase the number of independent Fourier modes by going to higher k. The latter has several advantages. First, from the perspective of increasing cumulative SNR, going to higher k simply adds more modes that can be used to beat down the variance. Second, accessing higher $k_{\|}$ provides access to new independent modes that are less foreground-contaminated. Finally, going to a broader range of k increases the possibility of catching the (as-yet unknown) crossover scale when the cross-correlation power spectrum transitions from being negative to positive. Probing these scales would give us access to information about the morphology of ionization field and allow us to constrain a wider range of reionization scenarios.

In this section we forecast a futuristic 21 cm-[CII] cross-correlation measurement where HERA and CCATp have increased Fourier overlap in the $k_\parallel$ direction up to $k_\parallel \approx 3\,h\mathrm{Mpc}^{-1}$. This requires increasing the CCATp spectral resolution to $\Delta\nu = 0.5\,\mathrm{GHz}$ ($R \approx 500$). Apart of this modification, the specifications for each instrument remain the ones listed in Tables 2.1 and 2.2. For our more futuristic scenarios, we simulate measurements over $40\,\mathrm{deg}^2$ and $400\,\mathrm{deg}^2$. Rather than simulating a contiguous survey area, we continue to simulate the patches shown in Figure 2.5 but average together the power spectrum results from multiple patches to accumulate the larger survey areas. Although this represents an incoherent average of power spectra (rather than a coherent increase in contiguous survey area prior to forming power spectra), our calculations should serve as a reasonable proxy for the expected sensitivity since we are not seeking to access lower $k_\perp$ modes that are angularly coherent over large parts of the sky. To gather statistics for increased survey areas, we perform bootstrap sampling over the different foreground patches. As in the previous section, we simulate two analyses: one with a wedge cut and one without.

In Figure 2.10 the cross spectra with the wedge cut are shown. To avoid visual clutter, we forgo the violin plots and instead show conventional error bars, with the expectation that the increased averaging implicit in our new scenarios will serve to somewhat Gaussianize the distributions (see Section 2.9.2 for a critical examination of this). One sees that over $40\,\mathrm{deg}^2$, precise measurements can be made of the large scale modes at $z = 6$ and $z = 7$. While in principle the increased frequency resolution of the [CII] survey ought to allow one to measure the crossover scale, in practice the high-k bins remain noise limited. A brief examination of Figure 2.1 reveals the reason for this. Since our futuristic scenario does not entail adding coverage at higher $k_\perp$ (which would be extremely difficult for 21 cm experiments), the new high k bins are formed from relatively small sections of contours of constant $k \equiv \sqrt{k_\perp^2 + k_\parallel^2}$ that are almost horizontal on the $k_\perp$-$k_\parallel$ plane. Thus, relatively few independent modes are added, limiting one's ability to average down the noise. For surveys conducted over $400\,\mathrm{deg}^2$ many of the trends remain, although measurements across the crossover scale are now possible in the $z = 7$ bin.

In Figure 2.11 the spectra without the wedge cut and without any foreground mitigation are shown. Due to the highly contaminated foreground wedge modes, there is a decrease in the SNR in comparison to when the wedge cut was performed.[6] That being said, many modes can still be measured to high significance especially over the larger 400 deg^2 survey area. If one, once again, performs 21 cm foreground mitigation with 1% residuals, the situation drastically improves. In Figure 2.12 we present the spectra when foreground mitigation is performed. In this case, precision measurements of large scale modes can be made across all redshifts for a 40 deg^2 survey. Over 400 deg^2, precision measurements of both large-scale modes and small-scale modes can be made at all redshifts. By designing a CCAT-like instrument with increased spectral resolution, optimizing survey strategies to ensure extensive sky coverage overlap, and coupling these improvements with a modest foreground mitigation strategy, a high-significance detection of the 21 cm–[CII] cross spectrum is likely achievable.

2.9 Lessons Learned

In conducting the forecasts of Sections 2.7 and 2.8, our Monte Carlo simulations allow considerable flexibility for exploring forecasting methodology. In this section, we summarize some of the "lessons learned" along the way, with an eye towards understanding which commonly

---

[6] At first glance, this may seem to be a counterintuitive result: one might expect that including extra Fourier modes would only increase the cumulative SNR, even if those modes have an extremely large variance to them. In some ways, this is an artifact of our simple power spectrum estimator and the way we compute SNR_{cum}. Our cumulative SNR is computed in a spherically averaged space, that is, for the power spectrum estimator $\hat{P}_{ab}$ as a function of k rather than $k_\perp$ and $k_\parallel$. Many of our k bins straddle the foreground wedge, with some modes coming from within the wedge and some modes coming from outside the wedge. Now, recall that our simple estimator $\hat{P}_{ab}(k)$ simply averages $\hat{P}_{ab}(k_\perp, k_\parallel)$ in rings of constant k *with uniform weights all modes*. This means that without the excision of the wedge, modes with extremely high foreground variance can otherwise contaminate what would be a k bin with reasonable SNR. A more advanced treatment would be to perform power spectrum estimation using an optimal quadratic estimator. Such an estimator would downweight the data by its total inverse covariance matrix (including the foreground covariance) *before* binning (Liu et al., 2014a), thereby organically preventing a contaminated $k_\perp$-$k_\parallel$ foreground mode from polluting a k bin (Liu and Tegmark, 2011). In fact, the Monte Carlo approach that we espouse in this paper can provide precisely the needed covariance information to do this, but for simplicity we leave this investigation to future work.

used forecasting approximations are justifiable and which are not.

2.9.1 A $1/\sqrt{N}$ Scaling Works for Foregrounds

Implicit in the use of cross-correlations for foreground mitigation is the assumption that different patches of the sky have independent foregrounds. In effect, one is using angular averages as an approximation for the ensemble averaging in Equation (2.1), and the hope is that by increasing sky coverage, Equation (2.2) becomes a better and better approximation. Crucially, this relies on different parts of the sky having independent foregrounds.

In principle, foregrounds are most certainly not independent from one part of the sky to another on large scales. This is clear from a visual inspection of Figure 2.4, or from the existence of large-scale power in previous foreground observations (Gold et al., 2009, 2011; Planck Collaboration et al., 2016). However, in practice it may be a sufficiently good approximation for foreground *fluctuations*—as quantified by the foreground contribution in a particular $\mathbf{k}$ mode—to be treated as independent when sourced from different parts of a survey.

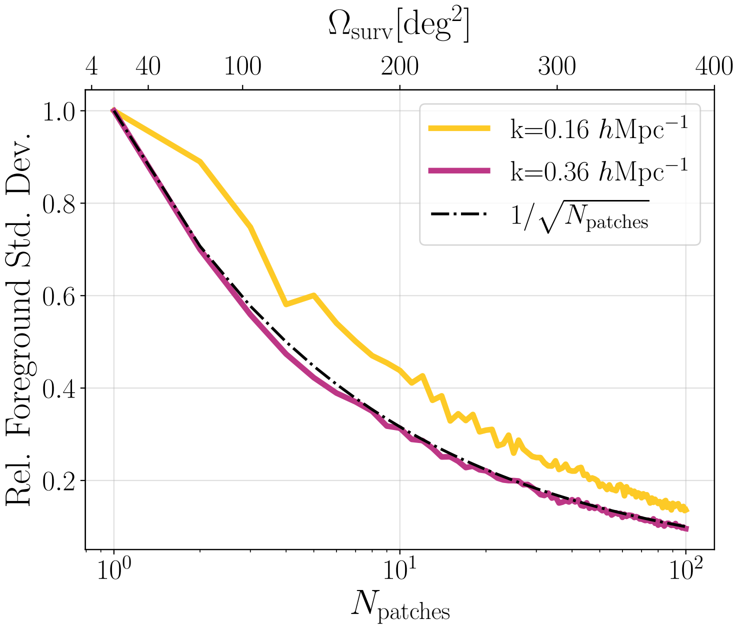

Figure 2.13: Relative foreground standard deviation as a function of number of simulation patches averaged together (equivalently survey area). In yellow the foreground standard deviation for the $k = 0.16\ h\text{Mpc}^{-1}$ is plotted, and in pink the same quantity for $k = 0.36\ h\text{Mpc}^{-1}$ is plotted. As a reference, the expected $1/\sqrt{N_{\rm patches}}$ relation for $N_{\rm patches}$ independent sample is shown in dashed black. One sees that the foreground residual power averages down with increasing sky area.

With our Monte Carlo samples including different draws from different parts of the sky (recall Figure 2.5), our pipeline is equipped to test (rather than assume) the independence of different foreground patches. We examine the averaging down of foregrounds in two different power spectrum bins: one centered on $k = 0.16\,h\text{Mpc}^{-1}$ (intended to be at low enough k to be representative of a mode where strong foreground contamination resides) and one centered on $k = 0.36\,h\text{Mpc}^{-1}$ (at higher k where there is weaker but still non-negligible foregrounds from leakage effects). For each mode, we vary the number of $6 \times 6\,\text{deg}^2$ patches that are used to estimate the power spectrum. The patches are combined in an incoherent manner, such that power spectra are computed for each individual patch and then averaged together.[7] This yields a single power spectrum estimate with

---

[7] In principle, this will result in a lower sensitivity than a setup where a contiguous volume is analyzed coherently.

suppressed (but still existent) foreground residuals that have not entirely cross-correlated away. To compute the increase in foreground variance that this implies, we perform a bootstrap sampling over our random selection of patches in Figure 2.5.

In Figure 2.13 we show the results of such an experiment, with the variance of our bootstrapped samples normalized to the variance for a single $6 \times 6\,\mathrm{deg}^2$ patch. The dark purple line (showing the behaviour of the $k = 0.36\,h\mathrm{Mpc}^{-1}$ bin) clearly follows a $1/\sqrt{N}$ form, implying that it is reasonable to assume that the foregrounds average down independently. The lighter yellow line (for $k = 0.16\,h\mathrm{Mpc}^{-1}$) deviates slightly from $1/\sqrt{N}$ but still averages down quite well as more patches are included. One may thus conclude that although the gold standard for forecasting would be to simulate precisely the planned sky coverage for one's survey, for rough back-of-the-envelope estimates it is likely acceptable to treat different foreground patches as independent in a cross-correlation experiment. This validates, for example, the approach used in Roy and Battaglia (2024) of simulating smaller volumes and then appropriately scaling the resulting SNRs to match proposed surveys, even if foregrounds are involved.

2.9.2 Some Gaussianization Of Errors Distribution Occurs But With Residual Non-Trivial Behaviour In The Tails

A common assumption in many forecasting efforts is that error bars can be treated as Gaussian (although see Wilensky et al. 2023b for a critical examination in the context of 21 cm auto power spectrum measurements). Again, our pipeline is equipped to test this assumption, since we include the non-Gaussian effects of foregrounds and furthermore propagate its interactions with other sources of uncertainty in our measurements. One should expect that with sufficient binning of power spectra that the error distributions would eventually Gaussianize via the Central Limit Theorem (CLT). However, it is crucial to understand whether the probability distributions of foregrounds

---

As an extreme example of this, one's sensitivity to angular modes larger than the extent of an individual patch is formally zero when the patches are analyzed separately but non-zero when analyzed together. In practice, our forecasts do not deal with sufficiently small $k_\perp$ values for this to be a concern.

(in various Fourier modes) are sufficiently well-behaved to satisfy the assumptions of the CLT. This is not *a priori* true for three reasons. First, the probability distributions of foregrounds may have heavy tails to be well-behaved enough for the CLT. Second, the precise mix of foreground, noise, and cosmological power within a spherical Fourier k shell will depend on the direction of $\mathbf{k}$ (indeed, this is the central fact leveraged by wedge-excision foreground mitigation schemes). Third, as we will see in Section 2.3, power spectrum estimates in different parts of Fourier space tend to be correlated with each other. These last two caveats mean that the act of binning from $P(\mathbf{k})$ to $P(k)$ is tantamount to the summation of *non-independent* and *non-identically* distributed random variables—a stark contrast to the independent and identically distributed draws assumed by the CLT.

Although there exist some generalizations of the CLT that can be used for an analytic treatment (Wilensky et al., 2023b; Billingsley, 2012), here it is simpler to test for Gaussianization (or lack thereof) empirically using our simulations. To do so, we use the data from our futuristic case with no wedge cut (Figure 2.11), but with the original survey area[8] of $2 \times 2\,\mathrm{deg}^2$. We use our suite of realizations to study the distribution of the cross spectrum power in individual voxels in three-dimensional Fourier space (i.e., the power spectrum estimates prior to any binning). These Fourier voxels are selected such that they would be binned into the same bin in a spherical binning. The distributions of individual voxel values are then compared to the distribution of the final spherically averaged power spectrum. The hope is that while the data in individual Fourier voxels may not be Gaussian distributed, the distribution of the binned value Gaussianizes.

In Figure 2.14 we plot the various distributions.[9] Shown in different shades of grey are the distributions of cross power at Fourier voxels with different $\mathbf{k}$ but roughly identical k that will eventually be averaged into the same $k \approx 0.74h\,\mathrm{Mpc}^{-1}$ bin. We choose to plot the distributions of

---

[8] Our motivation for using a smaller survey area was to allow for a larger number of samples without the added layer of complication that comes with bootstrapping, which recall from Section 2.8 is how we perform our forecasts for large survey areas.

[9] For visualization purposes, in both Figure 2.14 and 2.15 we use a kernel density estimator (with bandwidth parameter set by Scott's rule; Scott 2015) rather than directly plotting histograms.

three voxels that are representative of the type of modes in the bin: a mode with a low $k_{\|}$ component that is consistently foreground dominated across all realizations, a mode with a high $k_{\|}$ component that is consistently much cleaner across all realizations, and finally an intermediate mode. These distributions have extremely different widths to them, since the foregrounds are so much brighter than instrumental noise. Thus, to highlight the shapes of the distributions (rather than their widths), we scale each distribution by its maximum power value $P_{\max}$. Plotted in red is the distribution of the final binned power. The distributions of the individual Fourier voxels are clearly non-Gaussian, exhibiting a multitude of peaks and valleys. The distribution of the resultant bin shown in red is better behaved and is well fit by a Gaussian distribution (dashed black curve) near the peak; however, it fails to Gaussianize in the tails. We thus conclude that a Gaussian approximation is a reasonable treatment out to several standard deviations, but the behaviour beyond that should be treated with care.

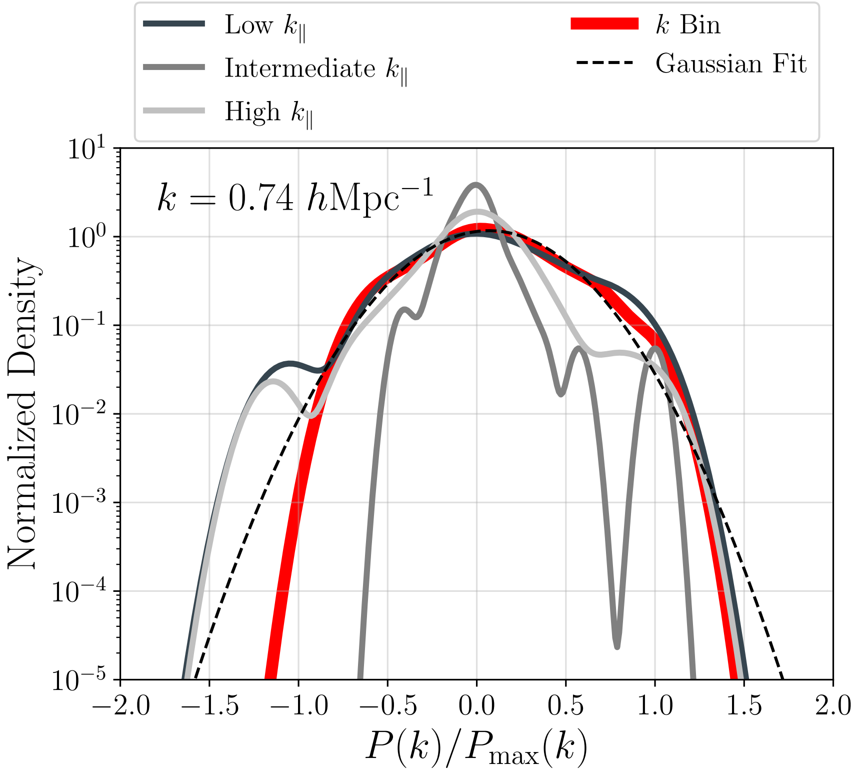

Figure 2.14: The normalized probability density of a selection of power spectrum voxels (various shades of grey) and the resultant power spectrum bin (red). As a reference, the Gaussian fit to the red curve is plotted in dashed black. Binning is seen to Gaussianize power spectrum distributions to some extent, but there are residual discrepancies in the tails.

2.9.3 Cross-Terms are Important, but Separate Simulations are Acceptable

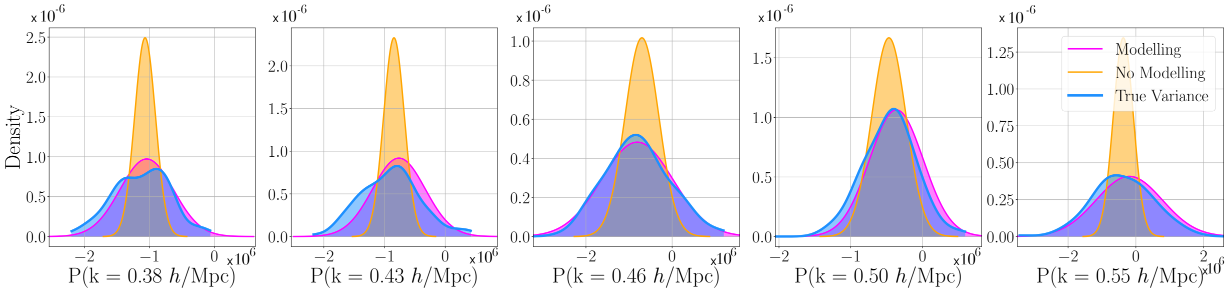

Figure 2.15: From left to right, the probability distribution of the measured cross-power for increasing wavenumber k. In blue are the distributions obtained from full MC simulations, in pink are Gaussianized approximations in which each model component is included but not co-varied, and in orange each source of variance is simply computed independently. As more approximations are made, the distributions deviate further from the blue distributions.

In this paper, we have stressed the importance of simultaneously varying multiple components of a measurement (i.e., the cosmological signal, foregrounds, and noise) in order to capture the full probability distributions of final power spectrum band powers. The non-Gaussianity of the foregrounds (or in principle the cosmological signal) means that the analogous expression to Equation (2.3) for higher-order moments includes terms with non-trivial products between sky components. That the components are statistically independent allows, say, a six-point function with four copies of the foregrounds and two copies of the signal to be reduced a four-point function of foregrounds multiplied by the cosmological power spectrum. That said, taking advantage of such simplifications still require evaluating higher order moments of each component, which means that in practice it may be easier to simply simulate all components together.

For an approximate treatment, however, one can use the fact (established above) that the non-Gaussianities in the final errors are small in many cases of interest. This suggests that simulating the components separately and then applying Equation (2.3) may be sufficient. Figure 2.15 shows that this is indeed the case for non-wedge modes. In blue we plot the full probability distribution of each mode for the measured HERA and CCATp cross-spectrum from the Monte Carlo simulations used to make Figure 2.8. These distributions are non-Gaussian, especially at low k values where foregrounds are important. In pink, we plot the error distributions obtained by individually varying the foreground, noise, and cosmic variance components in our simulation. While varying a given component, the others are still modeled and fixed to a single realization. We then compute the standard deviation of each such distribution (effectively making a Gaussian approximation) and add these individual error components in quadrature to obtain the total standard deviation. This standard deviation is inserted into a Gaussian form for the probability distribution. While these Gaussianized distributions obviously do not capture the full shape of the distributions, they are a relatively faithful description of its width and therefore the variance. Finally, we again vary each component separately but in doing so, do not include a realization of any other components. In other words the cosmic variance, noise variance, and foreground variance are computed separately and then added in quadrature. This is equivalent to computing only the first four terms of Equation 2.3. The resultant Gaussian distribution is plotted in orange. In this case, since no cross terms have been modeled, the variance is heavily underestimated. Neither the width nor the shape of the distribution captures the variance of the data.

In the first few columns of Table 2.3, we summarize the results of this exploration for current-generation surveys. Approximate methods tend to have the effect of overestimating SNR. Across all redshifts, not modeling the full variance ("No Modeling" column) leads to a overestimation of the cumulative SNR. When no joint modelling is performed, the cumulative SNR is overestimated by over 5.41σ at $z = 6$, by 0.40σ at $z = 7$, and by 0.8σ at $z = 8$. When some modeling of the other simulation components is included ("Modeling" column), the cumulative SNR is a much more faithful representation of the true variance, although one loses the full shape of the distribution. The results from simultaneously varying all components and computing the variance directly form the full non-Gaussian distributions in Figure 2.15 are shown in the column labeled "Var". One sees that the SNR decreases further. In the following section, we will see yet another effect when we explore the importance of not only considering the non-Gaussianities but also error covariances between bins.

z	No Modeling	Modeling	Var.	**Cov.**
6	8.85	4.36	3.44	**3.01**
7	2.09	1.78	1.69	**2.38**
8	0.75	1.70	1.43	**0.81**

Table 2.3: Cumulative SNR for each redshift bin computed using various approximations. The column labelled "Cov." makes use of the full covariance matrix and Equation 2.24 to compute the SNR. The neighbouring column labelled "Var." only takes into account the diagonal elements of the covariance matrix (i.e. the variance). The remaining two columns implement further approximations as described in Section 2.9.3.

2.9.4 Error Correlations Between Bins are Non-Trivial

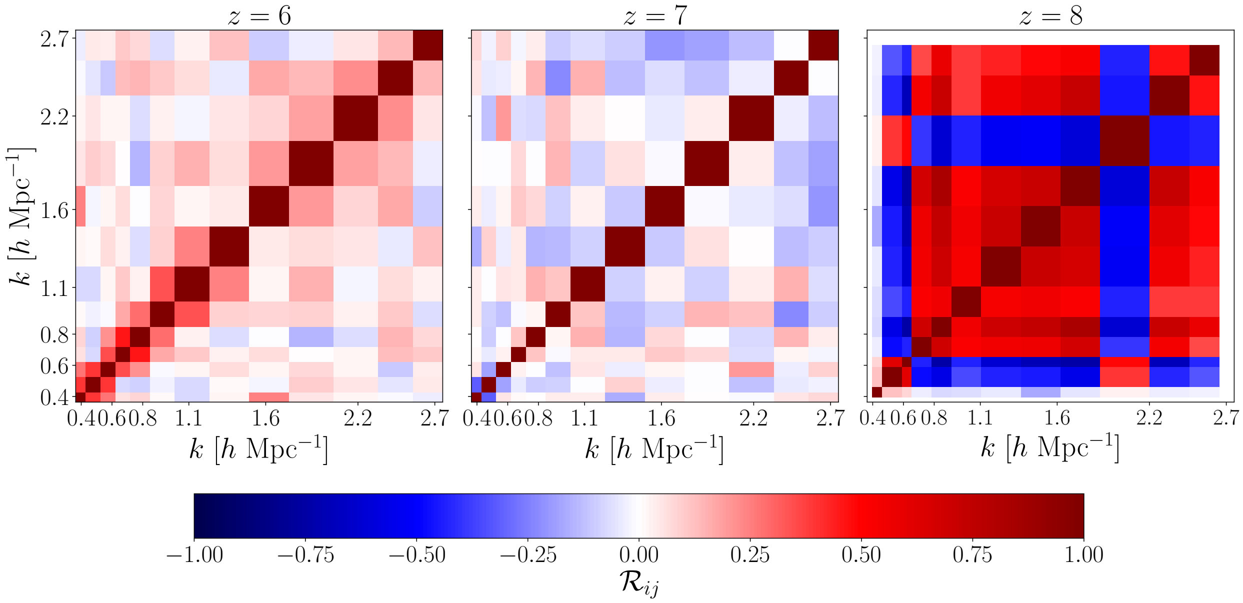

Figure 2.16: Dimensionless correlation matrices $\mathcal{R}$, defined in Equation (2.25), for the futuristic HERA × CCAT-prime forecasts shown in Figure 2.10. Substantial error correlations exist between all k bins, necessitating the incorporation of full covariance information in evaluating signal-to-noise or in propagating power spectrum estimates to downstream parameter inferences.

A frequently used rule-of-thumb in performing power spectrum sensitivity forecasts is to assume that the errors in different Fourier bins are independent as long as their wavenumber spacing Δk is greater than $\Delta k \sim 2\pi/L$, where L is a characteristic size of a survey (Tegmark et al., 1998). With our suite of simulations, we can put this assumption to the test.

In order to highlight the interdependence (or lack thereof) of different Fourier bins, we compute the dimensionless correlation matrix $\mathcal{R}$, given by

$$\mathcal{R}_{ij} \equiv \frac{\mathbf{\Sigma}_{ij}}{\sqrt{\mathbf{\Sigma}_{ii}\mathbf{\Sigma}_{jj}}}, \tag{2.25}$$

rather than the covariance matrix $\mathbf{\Sigma}$. The results are shown in Figure 2.16, using the forecasts of

Section 2.8 with a wedge cut as an example. One sees that substantial, non-negligible correlations exist between different bins. This is not particularly surprising, since the assumption of independent bins can be violated by complicated instrumental responses (Liu et al., 2016), or analysis steps such as apodization and the downweighting or projecting out of systematics (such as foregrounds; Liu et al. 2014a). Bernal and Baleato Lizancos (2024) have also shown that residual interloper contaminants will naturally result in error correlations. Continuum foregrounds are also likely to induce correlation between bins (particularly in the regime where $k \approx k_{\parallel}$), since a key source of variance comes from the leakage of residual foreground power from low k (where the foregrounds intrinsically reside) to high k. With all these complicated and interconnected effects, it is no surprise that the covariance structure seen in Figure 2.16 is non-trivial, which speaks to the necessity of multi-component end-to-end simulations. For example, these error correlations are important to incorporate especially in low-SNR regimes. In the column labeled "Cov." of Table 2.3, we provide the SNRs computed using the full error covariances for the scenario described in Section 2.7.1 (i.e. HERA × CCATp with a wedge cut). At all redshifts, these values differ from the previously discussed ones that assumed no error covariance (column labeled "Var.").

What we have seen here is that modeling assumptions do matter in determining the cumulative SNR. Although we expect this to be true even for advanced futuristic experiments, our illustrative example here was computed for current-generation experiments because the metric of cumulative SNR is most appropriate when one is chasing an initial detection. Once one is *characterizing* a high-significance measurement, the details of precisely which parts of a spectrum are measured begin to matter, and a cumulative SNR becomes too blunt a metric. For example, while it is clear that a high-precision 21 cm × [CII] measurement at *precisely* the crossover scale between the negative and positive correlation regimes would be extremely useful, Equation (2.24) would ascribe precisely zero value to such a measurement! Said differently, SNR studies are necessary—but not sufficient—for understanding instrument performance, and in the following section we explore the effects of error covariances on a simple toy example of parameter estimation.

2.9.5 Window Functions Should Be Computed and Propagated Downstream for Parameter Estimation

Given the extra effort required to compute window functions as well as full error covariances, one may wonder how crucial this is when it comes to parameter estimation. In order to demonstrate the importance of modelling these quantities, we present a toy model whose parameters we seek to constrain and showcase how those constraints change in the presence of various error correlations and window functions.

Suppose one seeks to constrain the crossover scale k_0 at which the cross spectrum between [CII] and 21 cm transitions from negative to positive, as discussed in Section 2.2 and illustrated in Figure 2.3. As a simplified model, we imagine taking the two power spectrum points $\mathbf{P}_\pm \equiv (P_-, P_+)$ on either side and fitting a straight line model $P^{\rm mod}(k)$ between them to find the crossover, such that

$$P^{\rm mod}(k) \approx m(k - k_0), \tag{2.26}$$

where m is the slope of the line. Finding the crossover scale is then tantamount to computing the posterior distribution for k_0 given the measured power spectrum values, the error covariance between them, and their associated window functions. This posterior distribution $p(k_0|\mathbf{P}_\pm, \mathbf{\Sigma}, W)$ is given by

$$p(k_0|\mathbf{P}_\pm, \mathbf{\Sigma}, W) = \int dm\, p(k_0, m|\mathbf{P}_\pm, \mathbf{\Sigma}, W), \tag{2.27}$$

where W contains the window functions for the two bandpowers, and $p(k_0, m|\mathbf{P}_\pm, \mathbf{\Sigma}, W)$ is the joint posterior for the slope and the crossover scale. Bayes' theorem (Bayes, 1763) enables this to be written as

$$p(k_0, m|\mathbf{P}_\pm, \mathbf{\Sigma}, W) \propto p(\mathbf{P}_\pm|k_0, m, \mathbf{\Sigma}, W)\, p(k_0, m|\mathbf{\Sigma}, W), \tag{2.28}$$

where $p(k_0, m|\mathbf{\Sigma}, W)$ is the prior and $p(\mathbf{P}_\pm|k_0, m, \mathbf{\Sigma}, W)$ is the likelihood. Given our previous demonstrations that the binned power spectra have Gaussianized to some degree (at least away from the low-level tails of the distribution), the logarithm of the likelihood can be written as[10]

$$\ln p(\mathbf{P}_\pm | k_0, m, \mathbf{\Sigma}, W) \propto -\frac{1}{2}(\mathbf{P}_\pm - \mathbf{P}_\pm^{\mathrm{mod}})^t \mathbf{\Sigma}^{-1} (\mathbf{P}_\pm - \mathbf{P}_\pm^{\mathrm{mod}}), \tag{2.29}$$

where $\mathbf{P}_\pm^{\mathrm{mod}} \equiv (P_-^{\mathrm{mod}}, P_+^{\mathrm{mod}})$ is Equation (2.26), i.e., our model, evaluated at the k_- and k_+, the wavenumbers of the bandpowers just shy of and just beyond the crossover scale, respectively. These are given by

$$\begin{aligned} P_\pm^{\mathrm{mod}} &= \int dk\, W(k_\pm, k) P^{\mathrm{mod}}(k) \\ &= m\left[\int dk\, k\, W(k_\pm, k) - k_0 \int dk\, W(k_\pm, k)\right] \\ &\equiv m(k_\pm^{\mathrm{eff}} - k_0), \end{aligned} \tag{2.30}$$

where in the penultimate line the second integral is by construction unity for properly normalized window functions, and we have defined the first integral (the centre of mass of the window function) to be $k_\pm^{\mathrm{eff}}$. Inserting this into Equation (2.29) and subsequently into Equation (2.28) then provides the full posterior for k_0 and m.

Alternatively, since Equation (2.29) is a quadratic in both m and k_0, the marginalization over m can be performed analytically to give

$$p(k_0 | \mathbf{P}_\pm, \mathbf{\Sigma}, W) \propto \sqrt{\frac{2\pi}{a}} \exp\left(\frac{b^2}{8a}\right), \tag{2.31}$$

where

$$a \equiv (\mathbf{\Delta k}^{\mathrm{eff}})^t \mathbf{\Sigma}^{-1} \mathbf{\Delta k}^{\mathrm{eff}}, \quad b \equiv -2\mathbf{P}_\pm^t \mathbf{\Sigma}^{-1} \mathbf{\Delta k}^{\mathrm{eff}}, \tag{2.32}$$

with $\mathbf{\Delta k}^{\mathrm{eff}} \equiv (k_-^{\mathrm{eff}} - k_0, k_+^{\mathrm{eff}} - k_0)$.

---

[10] In Equation (2.28), we have omitted the normalization term of the Gaussian likelihood. In principle, cosmological information is present in this term, since $\mathbf{\Sigma}$ is the total covariance—including contributions from the cosmological signal (i.e., cosmic variance). In practice, for most experiments the constraining power from this piece of the likelihood is negligible (Tegmark, 1997b), so we neglect it.

Using Equation 2.29, we compute the marginalized posterior distributions for parameters m and k_0 and explore various assumptions for window functions and error covariances. With the window functions we explore two different computations. In the first one, we take the correct $z = 6$ window functions from Figure 2.12 that correspond to the two data points straddling the $k_0 \approx 2\,h\mathrm{Mpc}^{-1}$ crossover scale. In the second computation, we erroneously assume that the window functions are Dirac delta functions centered on the starred points in Figure 2.12. This is in effect what is assumed if one does not bother to compute window functions for one's power spectrum estimator. For the error covariance, we examine three possibilities. We consider a case where P_- and P_+ are positively correlated by using the true covariance between the points (which happens to have positive off-diagonal terms in this case). Our second possibility is an uncorrelated case where the off-diagonal terms are artificially set to zero. The final possibility is the negatively correlated version where the off-diagonal terms are given the opposite sign compared to the true covariance. The final posteriors, marginalized separately for m and k_0, are shown in Figure 2.17 for all the aforementioned scenarios. Also shown (in vertical dashed lines) are the true values of the parameters inferred by acting on the true theory cross power spectra with the window functions.

Several trends are clear. Immediately, one sees the importance of including window functions in one's analyses. Without taking them into account, significant biases are present in the inference of both m and k_0. The error covariances strongly affect the widths of the posteriors. Positive covariances give rise to narrow posteriors on the slope m because coherent (positively correlated) perturbations move P_- and P_+ up and down in concert, leaving the slope relatively unaffected. The reverse is true for inferences of k_0, where the two points straddling either side of the horizontal axis must be perturbed in opposite directions to keep the intercept relatively constant. Putting this together, one sees that if non-negligible covariances exist between P_- and P_+, their mismodeling can lead to parameter inferences that can be either overconfident or underconfident. For example, in our current scenario the covariance is positive, and thus the omission of off-diagonal terms in $\mathbf{\Sigma}$ would result in an overconfident error bar on k_0.

In this section we have computed a few quantitative examples to demonstrate a qualitative point—that it is crucial to quantify one's window functions and error covariances. However, it is important not to overgeneralize our conclusions, since they can easily change depending on the subtle details of one's scenario. For example, suppose the crossover scale occurred at $k \sim 1.5\, h\mathrm{Mpc}^{-1}$ at $z = 7$, and one was modeling the experiment corresponding to Figure 2.16. In that case, P_- and P_+ would be negatively correlated, leading to the opposite conclusions regarding the question of overconfidence versus underconfidence in one's inferences.

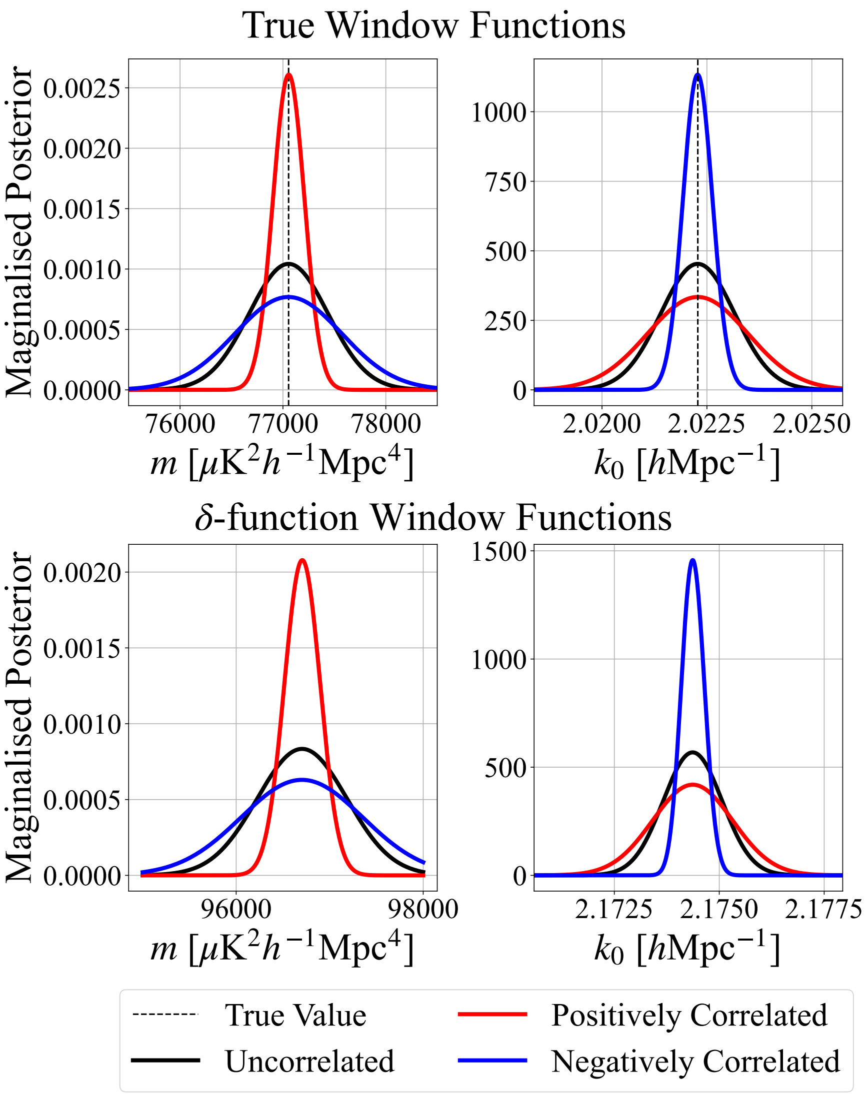

Figure 2.17: Mock constraints on the slope m and crossover scale k_0, expressed as marginalized posteriors where all variables except for the parameter in question have been marginalized over. Top: Constraints using the true window functions of the power spectrum estimate, as computed in Figure 2.12. Bottom: Biased inferences made assuming that the window functions are Dirac delta functions. Different solid lines show the effect of different error covariances while the true parameters values are shown with vertical dotted lines. One sees that neglecting window functions in analyses can lead to biased inferences, whereas the mismodeling of error covariances can lead to parameter constraints that are overly pessimistic or overly optimistic.

2.10 Conclusion

Line-intensity mapping is rapidly becoming a promising technique for conducting efficient yet sensitive surveys of large-scale structure, with the scientific reach of potential surveys spanning both cosmological and astrophysical applications. A variety of different emission lines are being targeted by present and upcoming instruments, opening the door to a large number of possible cross-correlations. Such cross-correlations would likely represent not only a valuable scientific product, but also a means to suppress systematics such as foregrounds. However, forecasting the value of LIM cross-correlations requires the detailed characterization of hypothetical measurements and their statistical properties, quantitatively taking into account instrumental effects and the nuances of one's data analysis prescriptions.

In this paper, we present an end-to-end pipeline for simulating LIM cross-correlations. In using this pipeline to carry out a detailed forecast for upcoming 21 cm-[CII] cross-correlation measurements, we explored the impact of various analysis choices and forecasting methodologies. This resulted in the following lessons learned:

1. When incoherently averaging together power spectrum estimates from N_{patches} different parts of the sky, the foreground residual errors average down with a $1/\sqrt{N_{\text{patches}}}$ scaling (even though one might assume *a priori* that foregrounds are not entirely independent in different parts of the sky).

2. Although foreground residual errors in the power spectrum are not expected to be Gaussian distributed, some level of Gaussianization occurs in the binning step of one's power spectrum estimation. However, there typically remains some residual non-trivial behaviour in the tails of the distributions.

3. In general, cross-terms in power spectrum variances mean that it is advisable to simultaneously vary noise, the cosmological signal, and foregrounds in one's simulations. However, in certain regimes (detailed in Section 2.9.3) this is reasonably similar to error

estimates obtained by performing a simulation where each component is varied while others are held fixed, and then combining the resulting variations in quadrature.

4. Error covariances between different k bins of one's power spectrum estimate exist, and have a non-negligible impact when computing SNR.

5. Window functions and error covariances should be computed and must be propagated downstream to parameter estimation. Doing so avoids parameter inference biases and the possibility of final constraints that are either underconfident or overconfident.

We believe these lessons to be broadly applicable to a variety of forecasts using LIMs or other cosmological probes. To that end, the publicly available pipeline, LIMSTAT, outlined in this work can easily be adapted to not only explore cross-correlations between interferometric and single dish observations, but any combination of such measurements. For instance, one might instead be interested in understanding the cross-correlation between Lyα emitting galaxies detected by the Spectro-Photometer for the History of the Universe, Epoch of Reionization and Ices Explorer (SPHEREx, Doré et al. 2014) and [CII] emission from CCATp.

Cognizant of the aforementioned subtleties, we forecast the possibility of an upcoming cross-correlation measurement with HERA and CCATp, which are set to observe a small overlapping region of sky subtending 4 deg^2. Given this small survey area and the limited Fourier coverage of these two instruments, the commonly employed 21 cm foreground mitigation strategy of excising the footprint of the so-called "foreground wedge" is undesirable. While it is possible with such a strategy to make a marginal detection at $z = 6$ and $z = 7$, these measurements remain noise-limited. Opting to include the wedge modes in the hopes of decreasing the variance without any foreground mitigation strategy does not provide much relief. However, suppose one is able to address the 21 cm foreground contaminants using other means and reduce them by a factor of 100 at the map level. Although ambitious, this requirement is still considerably easier than what is required for a 21 cm auto-spectrum measurement. If such a level of mitigation can be achieved, we find that the cross-correlations can be a huge benefit in teasing out a 6.79σ detection at $z = 6$

and a 6.84σ detection at $z = 7$. Going to higher redshifts remains a challenge since both the noise and the foreground spectra brighten significantly with decreasing frequency. Nonetheless, these forecasts indicate that this first generation of 21 cm-[CII] cross-correlations may help us rule our certain reionization scenarios which produce ionization bubbles on very large scales. That being said, much of the parameter space of possible cross-over scales is not probed by the poor Fourier coverage of these instruments.

Given these limitations, we look to the future and provide a forecast for a set of hypothetical new instruments. We go to larger survey areas and also achieve better Fourier coverage by improving the spectral resolution of CCATp. This allows one to make measurements of higher statistical significance, as well as to cover a greater variety of possible theoretical models (which for example often make different predictions for the crossover scale at which 21 cm × [CII] correlations transition from negative to positive). Over increased survey areas of 40 deg^2 and 400 deg^2, we once again find that the most successful strategy is to include wedge modes and mitigate the 21 cm foregrounds to roughly the percent level. Over the larger survey area of 400 deg^2 precision measurements can be made across the cross-over scale, thus potentially providing constraints on both the timing and morphology of reionization.

While this pipeline has allowed us to perform a relatively realistic forecast, there remain a myriad of additional systematics that were not considered. Currently, 21 cm observations are challenged by a variety of systematics including, but not limited to, radio frequency interference (Wilensky et al., 2019), ionospheric effects (Jelić et al., 2010; Moore et al., 2013; Martinot et al., 2018), calibration errors (Orosz et al., 2019; Byrne et al., 2019; Barry et al., 2016), and mutual coupling (Rath et al., 2024). While only a small number of sub-mm CO and [CII] mapping instruments have come online, we have already been able to learn a great deal about the systematic contaminants that may be detrimental to these types of surveys. Side lobe pickup from the ground (Foss et al., 2022), atmospheric variability associated with Lissajous scanning strategies (Ihle et al., 2022), and mechanical degradation (Lunde et al., 2024) have all shown to decrease data quality. It is still unclear how these systematics will affect a cross-correlation measurement.

To conclude, this work underscores the significant potential of LIM-LIM cross-correlations in constraining the EoR in the near future. Using our simulation and analysis pipeline, we have elucidated the importance of end-to-end modeling and the need for the rigorous quantification of statistical properties. Our forecasts for upcoming HERA and CCATp cross-correlation measurements demonstrate both the challenges and opportunities ahead. Moving forward, it is crucial to pursue unified efforts among various experiments and coordinate common survey areas. By fostering collaborative initiatives, we maximize scientific returns and inch closer to achieving our shared vision of understanding the enigmatic early stages of structure formation.

Acknowledgements

HF would like to thank Adélie Gorce, Robert Pascua, Rebecca Ceppas de Castro, and Kai-Feng Chen for being the first users of this pipeline and acknowledges their work in co-authoring this piece of software, greatly expanding its functionality and user-friendliness. The authors would also like to thank Simon Foreman, Jessica Avva Zebrowski, Kirit Karkare, Cathryn Trott, Anirban Roy, Kana Moriwaki, Miguel Morales, Bryna Hazelton, Steven Murray, Anthony Pullen, Patrick Breysse, Jordan Mirocha, Saurabh Singh, and Michel Adamič for many insightful discussions.

HF is supported by the Fonds de recherche du Québec Nature et Technologies (FRQNT) Doctoral Research Scholarship award number 315907. AL and HF acknowledge support from the Trottier Space Institute, the New Frontiers in Research Fund Exploration grant program, an FRQNT New University Researchers Grant, the Canadian Institute for Advanced Research (CIFAR) Azrieli Global Scholars program, a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant and a Discovery Launch Supplement, the Sloan Research Fellowship, and the William Dawson Scholarship at McGill. This research was enabled in part by support provided by Compute Canada (www.computecanada.ca).

Appendix A: The Statistical Modeling Of Correlated Fields

In Pagano and Liu (2020), a simple statistical model for decorrelating fields was presented in the context of varying the degree of correlation between the ionization and density fields during reionization. This model generated fields that produce the same degree of correlation on all length scales. Here we extend this model to generate fields with scale-dependent correlations.

Suppose we are interested in generating two fields $T_a(\mathbf{r})$ and $T_b(\mathbf{r})$ such that they are consistent with having power spectra $P_a(k)$ and $P_b(k)$, respectively, i.e.,

$$\langle \tilde{T}_a(\mathbf{k})\tilde{T}_a(\mathbf{k}')^*\rangle = (2\pi)^3\delta^D(\mathbf{k}-\mathbf{k}')P_a(k) \quad \text{and} \quad \langle \tilde{T}_b(\mathbf{k})\tilde{T}_b(\mathbf{k}')^*\rangle = (2\pi)^3\delta^D(\mathbf{k}-\mathbf{k}')P_b(k), \tag{2.33}$$

where δ^D denotes a Dirac delta function and we adopt the standard cosmological Fourier convention, such that

$$\tilde{T}(\mathbf{k}) \equiv \int d^3r\, e^{-i\mathbf{k}\cdot\mathbf{r}}T(\mathbf{r}) \quad \text{and} \quad T(\mathbf{r}) = \int \frac{d^3k}{(2\pi)^3}\, e^{i\mathbf{k}\cdot\mathbf{r}}\tilde{T}(\mathbf{k}). \tag{2.34}$$

To ensure that the our two fields obey not only the right autocorrelation statistics but also have the right cross power spectrum, we generate the fields one at a time. For $T_a(\mathbf{r})$ we simply follow the usual procedure and draw random realizations in Fourier space with variance consistent with Equation (2.33). The second field is then created by computing

$$\tilde{T}_b(\mathbf{k}) = \tilde{T}_a(\mathbf{k})f(k)e^{-i\phi}. \tag{2.35}$$

The factor $f(k)$ adjusts the power of field b such that it has the auto-spectrum $P_b(k)$. Computing the two-point function of both sides of Equation (2.35) reveals that

$$P_b(k) = P_a(k)f^2(k) \quad \Rightarrow \quad f(k) = \sqrt{\frac{P_b(k)}{P_a(k)}}, \tag{2.36}$$

giving us the $f(k)$ required to satisfy our constraints on the auto spectra.

The factor $e^{-i\phi}$ adjusts the relative phases of the fields and ensures that the cross spectrum of fields a and b can be tuned to a desired $P_{ab}(k)$. In Pagano and Liu (2020), the values of ϕ were drawn from a Gaussian distribution with mean zero and standard deviation σ for all comoving wavenumber, k.

Here we seek to solve for σ as a function of k by relating it to known k-dependent quantities, the auto- and cross-spectrum. Forming the cross power spectrum between our two fields gives

$$P_{ab}(k) \propto \langle \tilde{T}_a(k)\tilde{T}_b^*(k)\rangle = f(k)\langle e^{-i\phi}\rangle\langle \tilde{T}_a(k)\tilde{T}_a^*(k)\rangle \propto f(k)\langle e^{-i\phi}\rangle P_a(k) = \langle e^{-i\phi}\rangle\sqrt{P_a(k)P_b(k)} \tag{2.37}$$

If ϕ is drawn from a Gaussian distribution, its expectation value can be readily evaluated to give

$$\langle e^{-i\phi}\rangle = \int P(\phi)e^{-i\phi} = \frac{1}{\sqrt{2\pi\sigma^2}}\int d\phi\, e^{-\phi^2/2\sigma^2}e^{-i\phi} = \exp\left[\frac{\sigma^2(k)}{2}\right], \tag{2.38}$$

where in the last expression we have explicitly included the potential k dependence of σ. It is precisely this k dependence that enables one to engineer a scale-dependence on the degree of correlation between our two fields. Comparing Equations (2.37) and (2.38), one sees that getting the right cross-correlation power spectrum $P_{ab}(k)$ can be achieved by setting

$$\frac{\sigma^2}{2} \equiv \ln|r(k)| + \frac{\pi}{2}\left(1 - \mathrm{sgn}[\mathrm{r(k)}]\right), \tag{2.39}$$

where sgn is the sign function. This is equivalent to Equation (2.7) given in Section 2.3, where one simply sets $\sigma(k) = \sqrt{\ln(|r(k)|^{-2})}$ and multiplies the right hand side of Equation (2.35) by $\mathrm{sgn}[r(k)]$.

In summary, drawing random phases for each location in $\mathbf{k}$ according to Equation (2.39) and then evaluating Equation (2.39) gives two fields with the desired auto- and cross-correlation spectra.

Appendix B: Power Spectrum Window Functions

In this appendix we derive expressions that relate the true cross-correlation power spectrum $P_{ab}(k)$ to the measured power spectrum $\hat{P}_{ab}(k)$ that is distorted by one's instrument response. Our goal is to find a window function $W_{ab}(k, k')$ such that

$$\langle \hat{P}_{ab}(k) \rangle = \int dk' W_{ab}(k, k') P_{ab}(k'). \tag{2.40}$$

To do so, we first write the observed sky map of an instrument $T^{\mathrm{obs}}(\mathbf{r})$ as

$$T^{\mathrm{obs}}(\mathbf{r}) = \int d^2u \exp\left(i2\pi \frac{\mathbf{u} \cdot \mathbf{r}_\perp}{r_z} \right) \tilde{G}(\mathbf{u}, \nu_z) \int d^2\phi \, \exp(-i2\pi \mathbf{u} \cdot \boldsymbol{\phi}) A(\boldsymbol{\phi}, \nu_z) T(\boldsymbol{\phi}, \nu_z), \tag{2.41}$$

where we are working in the flat sky limit given the (relatively) narrow fields considered in this paper. This allows us to identify a near-unambiguous line of sight direction that we can align with the z direction of our coordinate system, leaving the x and y directions to span the transverse space perpendicular to the line of sight, such that $\mathbf{r}_\perp \equiv (r_x, r_y)$. Again utilizing the flat-sky limit, we can write $\mathbf{r} \equiv (r_x, r_y, r_z) = (\mathbf{r}_\perp, r_z) = (r_z \boldsymbol{\phi}, r_z)$. Here, r_z is the radial comoving distance and $\boldsymbol{\phi} \equiv (\phi_x, \phi_y)$ refers to small angles corresponding to r_x and r_y. We denote the frequency of observation as ν_z, with the subscript "z" to remind ourselves of the one-to-one correspondence between frequency and radial distance in intensity mapping. In words, Equation (2.41) says that the sky is multiplied by some instrumental configuration space response $A(\boldsymbol{\phi}, \nu_z)$, then Fourier transformed in the transverse direction before being sampled by an instrumental Fourier space response $\tilde{G}$ (which will often be frequency dependent) and Fourier transformed back to configuration space.

Although we deal with two reasonably different types of instruments in this paper (interferometers and single dish telescopes), Equation (2.41) is general enough to accommodate both. For example, for a radio interferometer $A(\boldsymbol{\phi}, \nu_z)$ would represent the primary beam and $\tilde{G}(\mathbf{u}, \nu_z)$ would encode the sampling of the uv-plane by the baselines of the interferometer (hence the suggestive use of $\mathbf{u}$ as the Fourier dual to sky angle $\boldsymbol{\phi}$). For a single dish telescope, one might set $A(\boldsymbol{\phi}, \nu_z)$ to unity, whereas $\tilde{G}(\mathbf{u}, \nu_z)$ would be the Fourier transform of the telescope's point spread function.

To estimate the power spectrum, the data analyst can (optionally) multiply by some tapering/apodization function $B(\mathbf{r})$, defining $T^{\text{tap}}(\mathbf{r}) \equiv B(\mathbf{r})T^{\text{obs}}(\mathbf{r})$. This is then Fourier transformed to obtain

$$\tilde{T}^{\text{tap}}(\mathbf{k}) \equiv \int d^3r\, e^{-i\mathbf{k}\cdot\mathbf{r}} T^{\text{tap}}(\mathbf{r}) = \int \frac{d^3k'}{(2\pi)^3} \tilde{T}(\mathbf{k}') F(\mathbf{k}, \mathbf{k}'), \tag{2.42}$$

where

$$\begin{aligned} F(\mathbf{k}, \mathbf{k}') &\equiv \int d^3r\, e^{-i\mathbf{k}\cdot\mathbf{r}} e^{ik'_z r_z} B(\mathbf{r}) \int d^2u\, e^{i2\pi\mathbf{u}\cdot\mathbf{r}_\perp/r_z} \tilde{G}(\mathbf{u}, \nu_z) \int d^2\theta A(\boldsymbol{\theta}, \nu_z) e^{-i(2\pi\mathbf{u}-\mathbf{k}'_\perp r_z)\cdot\boldsymbol{\theta}} \\ &= \int dr_z \int d^2\mathbf{u}\, e^{-i(k_z - k'_z)r_z} \tilde{B}\left(\mathbf{k}_\perp - \frac{2\pi\mathbf{u}}{r_z}, r_z\right) \tilde{G}(\mathbf{u}, \nu_z) \tilde{A}\left(\mathbf{u} - \frac{\mathbf{k}'_\perp r_z}{2\pi}, \nu_z\right), \end{aligned} \tag{2.43}$$

with $\mathbf{k}_\perp \equiv (k_x, k_y)$ as the Fourier dual to $\mathbf{r}_\perp$. We note that we are mixing two different Fourier conventions here: for quantities that mostly pertain to an instrumental response, we use the Fourier convention where a factor of 2π appears in the exponent and the forward and backward transforms differ only by a sign in the exponent; for cosmological fields and $B(\mathbf{r})$ we adopt the standard cosmological convention stated in Appendix 2.10.

Forming an estimator $\hat{P}_{ab}(\mathbf{k})$ for the cross power spectrum of two surveys involves cross-correlating two copies of $\tilde{T}_{\text{obs}}(\mathbf{k})$, one from each instrument, to give

$$\begin{aligned} \hat{P}_{ab}(\mathbf{k}) &\equiv \frac{\langle \tilde{T}_a^{\text{tap}}(\mathbf{k}) \tilde{T}_b^{\text{tap}}(\mathbf{k})^* \rangle}{VN(\mathbf{k})} \\ &= \frac{1}{VN(\mathbf{k})} \int \frac{d^3k'}{(2\pi)^3} P_{ab}(k') F_a(\mathbf{k}, \mathbf{k}') F_b(\mathbf{k}, \mathbf{k}')^* \\ &= \int d^3k'\, P_{ab}(k') W_{ab}^{\text{3D}}(\mathbf{k}, \mathbf{k}') \end{aligned} \tag{2.44}$$

where $N(\mathbf{k})$ is a normalization factor (to be derived later), and V is the overlap volume of our two surveys, which are labelled using subscripts a and b for each of the two surveys. In the second equality we used the definition of the power spectrum, namely $\langle \tilde{T}_a(\mathbf{k})\tilde{T}_b(\mathbf{k})^* \rangle = (2\pi)^3 \delta^D(\mathbf{k} - \mathbf{k}')P(k)$, where δ^D signifies a Dirac delta function. In the final equality, we defined a window function $W^{\rm 3D}_{ab}$ that is analogous to the one defined in Equation (2.40), except in three-dimensional Fourier space. This version of the window function captures the way in which the estimated power spectrum $\hat{P}_{ab}(\mathbf{k})$ can be anisotropic in $\mathbf{k}$ thanks to our survey instruments, even though we have assumed that the true power spectrum $P_{ab}(k)$ is statistical isotropic and depends only on k.

Deriving explicit expressions for $W^{\rm 3D}_{ab}$ is important for two reasons. First, it allows us to verify that our power spectra are correctly normalized by enforcing that $\int d^3k'\, W^{\rm 3D}_{ab}(\mathbf{k}, \mathbf{k}') = 1$. In practice, this can be done by reverse engineering, where one imposes the normalization condition in order to derive the normalization $N(\mathbf{k})$, which is given by

$$N(\mathbf{k}) \equiv \int d^3k'\, W^{\rm 3D}_{ab}(\mathbf{k}, \mathbf{k}') = \frac{1}{V} \int \frac{d^3k'}{(2\pi)^3}\, F_a(\mathbf{k}, \mathbf{k}') F_b(\mathbf{k}, \mathbf{k}')^*. \tag{2.45}$$

The second reason for computing $W^{\rm 3D}_{ab}(\mathbf{k}, \mathbf{k}')$ is to eventually bin it to yield $W_{ab}(k, k')$, which determines the width of horizontal error bars in our final power spectrum estimates.

Consider first the issue of normalization. To simplify the expressions, we will assume that $B(\mathbf{r})$ is a function of frequency only. This is not an approximation, but simply a reflection of how we use $B(\mathbf{r})$ in this paper: as a spectral apodization function to avoid edge effects from bright foregrounds abruptly dropping to zero outside the survey volume. With this assumption, the spatial part of $\tilde{B}$ is proportional to a Dirac delta function, and one obtains

$$F(\mathbf{k}, \mathbf{k}') \approx \int dr_z\, e^{-i(k_z - k_z')r_z} r_z^2 B(r_z) \tilde{G}\left(\frac{\mathbf{k}_\perp r_z}{2\pi}, \nu_z\right) \tilde{A}\left(\frac{\mathbf{k}_\perp r_z}{2\pi} - \frac{\mathbf{k}'_\perp r_z}{2\pi}, \nu_z\right), \tag{2.46}$$

which eventually yields

$$N(\mathbf{k}) = \frac{1}{V} \int dr_z\, \tilde{G}_a \left(\frac{\mathbf{k}_\perp r_z}{2\pi}, \nu_z \right) \tilde{G}_b \left(\frac{\mathbf{k}_\perp r_z}{2\pi}, \nu_z \right)^* B_a(r_z) B_b(r_z) r_z^2 \int d^2\theta\, A_a(\boldsymbol{\theta}, \nu_z) A_b(\boldsymbol{\theta}, \nu_z). \tag{2.47}$$

For transverse modes on larger lengthscales than the resolution of one's instrument, we have $\tilde{G} \longrightarrow 1$, and this normalization will be close to order unity because the integrals reduces approximately to the survey volume (cancelling out the $1/V$ term), with just a slight modification for the spatial response $A(\boldsymbol{\theta}, \nu_z)$ and the frequency taper $B(r_z)$.

With a correctly normalized power spectrum estimator, the only interesting information in a window function is its shape. This is given by

$$\begin{aligned} W_{ab}^{3\mathrm{D}}(\mathbf{k}, \mathbf{k}') \quad = \quad & \int \frac{dr_z dr'_z}{V(2\pi)^3} e^{-i(k_z - k'_z) r_z} e^{i(k_z - k'_z) r'_z} r_z^2 r_z'^2 B_a(r_z) B_b(r'_z) \tilde{G}_a \left(\frac{\mathbf{k}_\perp r_z}{2\pi}, \nu_z \right) \tilde{G}_b \left(\frac{\mathbf{k}_\perp r'_z}{2\pi}, \nu'_z \right)^* \\ & \times \tilde{A}_a \left(\frac{\mathbf{k}_\perp r_z}{2\pi} - \frac{\mathbf{k}'_\perp r_z}{2\pi}, \nu_z \right) \tilde{A}_b \left(\frac{\mathbf{k}_\perp r'_z}{2\pi} - \frac{\mathbf{k}'_\perp r'_z}{2\pi}, \nu'_z \right)^* . \end{aligned} \tag{2.48}$$

While it is exact (up to the flat-sky approximation), this expression is unfortunately computational infeasible to use in practice. Instead, we must compute the binned version of the window function analytically. This can be done by making a few simplifying assumptions. First, we will take $A(\boldsymbol{\phi}, \nu_z)$ to be a frequency-independent quantity for our interferometer. In general, this will not be the case for a real observation. However, recall from Section 2.4.1 that our analysis involves carving out just a small central portion of our interferometer's large field-of-view, making a frequency-independent treatment more justifiable. Second, whenever r_z appears because it is used as a conversion factor between angles and transverse distances, we replace it with D_c, which we define as r_z evaluated at the midpoint of our survey volume. With both of these approximations we are essentially treating slowly varying functions of frequency as being perfectly constant. In contrast, we certainly cannot do the same for $T(\boldsymbol{\phi}, \nu_z)$ or $\tilde{G}(\mathbf{u}, \nu_z)$. (The latter's frequency dependence, for instance, is responsible for the phenomenology of the foreground wedge in interferometeric measurements; Liu et al. 2014a). Finally, we will assume that the response of our instruments are azimuthally symmetric (precluding the possibility of elliptical point spread functions, for example). This enables us to take $\tilde{G}(\mathbf{u}, \nu_z)$ to $\tilde{G}(|\mathbf{u}|, \nu_z)$.

To implement our approximation of a narrow field cut out from a wide field, we imagine that both $A_a(\boldsymbol{\theta})$ and $A_b(\boldsymbol{\theta})$ take Gaussian forms with standard deviations θ_a and θ_b, respectively. The last two terms of Equation (2.48) then become

$$\tilde{A}_a(\mathbf{q})\tilde{A}_b(\mathbf{q})^* \propto 2\pi^2(\theta_a^2 + \theta_b^2)\exp\left[2\pi^2(\theta_a^2 + \theta_b^2)|\mathbf{q}|^2\right] \rightarrow \delta^D(\mathbf{q}), \tag{2.49}$$

where $\mathbf{q} \equiv (\mathbf{k}_\perp - \mathbf{k}'_\perp)D_c/2\pi$ and in the last step we took the limit $\theta_a, \theta_b \rightarrow \infty$ to encode the flat response in the field of interest. (Essentially, we are taking the very flat top portion of a Gaussian). Inserting this into Equation (2.48) and imposing our azimuthally symmetric approximation on $\tilde{G}$ then gives

$$W_{ab}^{\text{3D}}(\mathbf{k}, \mathbf{k}') \propto \delta^D(\mathbf{k}_\perp - \mathbf{k}'_\perp)W_{ab}^{\|}(k_\perp, k_z - k'_z), \tag{2.50}$$

where

$$W_{ab}^{\|}(k_\perp, k_z) \equiv F_a^{\|}(k_\perp, k_z)F_b^{\|}(k_\perp, k_z)^*, \tag{2.51}$$

with $k_\perp \equiv |\mathbf{k}_\perp|$, and

$$F^{\|}(k_\perp, k_z) \equiv \int dr_z e^{-ik_z r_z}B(r_z)\tilde{G}\left(\frac{k_\perp D_c}{2\pi}, \nu_z\right), \tag{2.52}$$

with a and b subscripts adorning each term as appropriate. At this point, we have arrived at a computationally tractable form $W_{ab}^{\text{3D}}(\mathbf{k}, \mathbf{k}')$, our approximations having resulted in symmetries that reduce the complexity of the expressions (such as the fact that only the *difference* between k_z and k'_z matters).

To move towards a final expression for the binned window function $W_{ab}(k, k')$, we first bin azimuthally in the $k_x k_y$ plane down to $k_\perp$. Defining φ_k as the azimuthal angle within a cylindrical coordinate system defined by $k_\perp$, φ_k, and k_z (and similarly for primed coordinates), we can define

$$W_{ab}^{\text{2D}}(k_\perp, k_z; k'_\perp, k'_z) \propto \int d\varphi_k d\varphi'_k W_{ab}^{\text{3D}}(\mathbf{k}, \mathbf{k}'), \tag{2.53}$$

and using the fact that $\delta^D(\mathbf{k}_\perp - \mathbf{k}'_\perp) = \delta^D(k_\perp - k'_\perp)\delta^D(\varphi_k - \varphi'_k)/k_\perp$, we have

$$W^{\rm 2D}_{ab}(k_\perp, k_z; k'_\perp, k'_z) \propto \delta^D(k_\perp - k'_\perp) W^{\parallel}_{ab}(k_\perp, k_z - k'_z). \tag{2.54}$$

From here, there are two further binning operations that one can perform. The first is relatively straightforward, which is to fold $\pm k_z$ and $\pm k'_z$ into always-positive $k_\parallel$ and $k'_\parallel$ values. The second is to bin in circular $k = \sqrt{k_\perp^2 + k_\parallel^2}$ rings. This step is best done numerically—unlike the azimuthal binning that considerably simplified our expressions by taking advantage of azimuthal symmetries (or approximate symmetries) in our survey instruments, spherical binning has no such symmetry to leverage, since intensity mapping surveys probe line-of-sight fluctuations in a different way than they do transverse fluctuations.

Bibliography

S. R. Furlanetto, S. P. Oh, and F. H. Briggs, **433**, 181 (2006), astro-ph/0608032.

M. F. Morales and J. S. B. Wyithe, **48**, 127 (2010), 0910.3010.

J. R. Pritchard and A. Loeb, Reports on Progress in Physics **75**, 086901 (2012), 1109.6012.

A. Loeb and S. R. Furlanetto, *The First Galaxies in the Universe* (Princeton University Press, 2013).

A. Mesinger, ed., *The Cosmic 21-cm Revolution*, 2514-3433 (IOP Publishing, 2019), ISBN 978-0-7503-2236-2, URL `https://dx.doi.org/10.1088/2514-3433/ab4a73`.

A. Liu and J. R. Shaw, **132**, 062001 (2020), 1907.08211.

C. M. Trott and J. Pober, in *The Cosmic 21-cm Revolution* (IOP Publishing, 2019), 2514-3433, pp. 8–1 to 8–26, ISBN 978-0-7503-2236-2, URL `https://dx.doi.org/10.1088/2514-3433/ab4a73ch8`.

J. D. Bowman, A. E. E. Rogers, R. A. Monsalve, T. J. Mozdzen, and N. Mahesh, **555**, 67 (2018), 1810.05912.

S. Singh, N. T. Jishnu, R. Subrahmanyan, N. Udaya Shankar, B. S. Girish, A. Raghunathan, R. Somashekar, K. S. Srivani, and M. Sathyanarayana Rao, Nature Astronomy **6**, 607 (2022), 2112.06778.

J. D. Bowman and A. E. E. Rogers, **468**, 796 (2010), 1209.1117.

R. A. Monsalve, A. E. E. Rogers, J. D. Bowman, and T. J. Mozdzen, **847**, 64 (2017), 1708.05817.

S. Singh, R. Subrahmanyan, N. Udaya Shankar, M. Sathyanarayana Rao, A. Fialkov, A. Cohen, R. Barkana, B. S. Girish, A. Raghunathan, R. Somashekar, et al., **845**, L12 (2017), 1703.06647.

J.-M. Bégin, A. Liu, and A. Gorce, **105**, 083503 (2022), 2112.06933.

Z. Abdurashidova, J. E. Aguirre, P. Alexander, Z. S. Ali, Y. Balfour, A. P. Beardsley, G. Bernardi, T. S. Billings, J. D. Bowman, R. F. Bradley, et al., **925**, 221 (2022a), 2108.02263.

HERA Collaboration, Z. Abdurashidova, T. Adams, J. E. Aguirre, P. Alexander, Z. S. Ali, R. Baartman, Y. Balfour, R. Barkana, A. P. Beardsley, et al., **945**, 124 (2023), 2210.04912.

G. Paciga, J. G. Albert, K. Bandura, T.-C. Chang, Y. Gupta, C. Hirata, J. Odegova, U.-L. Pen, J. B. Peterson, J. Roy, et al., Monthly Notices of the Royal Astronomical Society **433**, 639–647 (2013), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stt753`.

M. P. van Haarlem, M. W. Wise, A. W. Gunst, G. Heald, J. P. McKean, J. W. T. Hessels, A. G. de Bruyn, R. Nijboer, J. Swinbank, R. Fallows, et al., Astronomy & Astrophysics **556**, A2 (2013), ISSN 1432-0746, URL `http://dx.doi.org/10.1051/0004-6361/201220873`.

A. H. Patil, S. Yatawatta, L. V. E. Koopmans, A. G. d. Bruyn, M. A. Brentjens, S. Zaroubi, K. M. B. Asad, M. Hatef, V. Jelić, M. Mevius, et al., The Astrophysical Journal **838**, 65 (2017), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/aa63e7`.

B. K. Gehlot, F. G. Mertens, L. V. E. Koopmans, M. A. Brentjens, S. Zaroubi, B. Ciardi, A. Ghosh, M. Hatef, I. T. Iliev, V. Jelić, et al., Monthly Notices of the Royal Astronomical Society **488**, 4271–4287 (2019), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stz1937`.

F. G. Mertens, M. Mevius, L. V. E. Koopmans, A. R. Offringa, G. Mellema, S. Zaroubi, M. A. Brentjens, H. Gan, B. K. Gehlot, V. N. Pandey, et al., **493**, 1662 (2020), `2002.07196`.

A. R. Parsons, D. C. Backer, G. S. Foster, M. C. H. Wright, R. F. Bradley, N. E. Gugliucci, C. R. Parashare, E. E. Benoit, J. E. Aguirre, D. C. Jacobs, et al., The Astronomical Journal **139**, 1468–1480 (2010), ISSN 1538-3881, URL `http://dx.doi.org/10.1088/0004-6256/139/4/1468`.

C. Cheng, A. R. Parsons, M. Kolopanis, D. C. Jacobs, A. Liu, S. A. Kohn, J. E. Aguirre, J. C. Pober, Z. S. Ali, G. Bernardi, et al., **868**, 26 (2018), `1810.05175`.

M. Kolopanis, D. C. Jacobs, C. Cheng, A. R. Parsons, S. A. Kohn, J. C. Pober, J. E. Aguirre, Z. S. Ali, G. Bernardi, R. F. Bradley, et al., **883**, 133 (2019), `1909.02085`.

M. W. Eastwood, M. M. Anderson, R. M. Monroe, G. Hallinan, M. Catha, J. Dowell, H. Garsden, L. J. Greenhill, B. C. Hicks, J. Kocz, et al., The Astronomical Journal **158**, 84 (2019), ISSN 1538-3881, URL `http://dx.doi.org/10.3847/1538-3881/ab2629`.

H. Garsden, L. Greenhill, G. Bernardi, A. Fialkov, D. C. Price, D. Mitchell, J. Dowell, M. Spinelli, and F. K. Schinzel, Monthly Notices of the Royal Astronomical Society **506**, 5802–5817 (2021), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stab1671`.

S. J. Tingay, R. Goeke, J. D. Bowman, D. Emrich, S. M. Ord, D. A. Mitchell, M. F. Morales, T. Booler, B. Crosse, R. B. Wayth, et al., **30**, e007 (2013), `1206.6945`.

A. Ewall-Wice, J. S. Dillon, J. N. Hewitt, A. Loeb, A. Mesinger, A. R. Neben, A. R. Offringa, M. Tegmark, N. Barry, A. P. Beardsley, et al., Monthly Notices of the Royal Astronomical

Society **460**, 4320–4347 (2016), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stw1022`.

A. P. Beardsley, B. J. Hazelton, I. S. Sullivan, P. Carroll, N. Barry, M. Rahimi, B. Pindor, C. M. Trott, J. Line, D. C. Jacobs, et al., **833**, 102 (2016), `1608.06281`.

N. Barry, M. Wilensky, C. M. Trott, B. Pindor, A. P. Beardsley, B. J. Hazelton, I. S. Sullivan, M. F. Morales, J. C. Pober, J. Line, et al., The Astrophysical Journal **884**, 1 (2019), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/ab40a8`.

C. M. Trott, C. H. Jordan, S. Midgley, N. Barry, B. Greig, B. Pindor, J. H. Cook, G. Sleap, S. J. Tingay, D. Ung, et al., **493**, 4711 (2020), `2002.02575`.

W. Li, J. C. Pober, N. Barry, B. J. Hazelton, M. F. Morales, C. M. Trott, A. Lanman, M. Wilensky, I. Sullivan, A. P. Beardsley, et al., The Astrophysical Journal **887**, 141 (2019), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/ab55e4`.

S. Yoshiura, B. Pindor, J. L. B. Line, N. Barry, C. M. Trott, A. Beardsley, J. Bowman, R. Byrne, A. Chokshi, B. J. Hazelton, et al., Monthly Notices of the Royal Astronomical Society **505**, 4775–4790 (2021), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stab1560`.

C. M. Trott, R. Mondal, G. Mellema, S. G. Murray, B. Greig, J. L. B. Line, N. Barry, and M. F. Morales, **666**, A106 (2022), `2208.06082`.

M. Rahimi, B. Pindor, J. L. B. Line, N. Barry, C. M. Trott, R. L. Webster, C. H. Jordan, M. Wilensky, S. Yoshiura, A. Beardsley, et al., Monthly Notices of the Royal Astronomical Society **508**, 5954–5971 (2021), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stab2918`.

M. J. Wilensky, M. F. Morales, B. J. Hazelton, P. L. Star, N. Barry, R. Byrne, C. H. Jordan, D. C. Jacobs, J. C. Pober, and C. M. Trott, **957**, 78 (2023a), `2310.03851`.

S. Munshi, F. G. Mertens, L. V. E. Koopmans, A. R. Offringa, B. Semelin, D. Aubert, R. Barkana, A. Bracco, S. A. Brackenhoff, B. Cecconi, et al., **681**, A62 (2024), 2311.05364.

R. Ghara, S. K. Giri, G. Mellema, B. Ciardi, S. Zaroubi, I. T. Iliev, L. V. E. Koopmans, E. Chapman, S. Gazagnes, B. K. Gehlot, et al., **493**, 4728 (2020), 2002.07195.

B. Greig, A. Mesinger, L. V. E. Koopmans, B. Ciardi, G. Mellema, S. Zaroubi, S. K. Giri, R. Ghara, A. Ghosh, I. T. Iliev, et al., **501**, 1 (2021), 2006.03203.

Z. Abdurashidova, J. E. Aguirre, P. Alexander, Z. S. Ali, Y. Balfour, R. Barkana, A. P. Beardsley, G. Bernardi, T. S. Billings, J. D. Bowman, et al., **924**, 51 (2022b), 2108.07282.

M. G. Santos, A. Cooray, and L. Knox, **625**, 575 (2005), astro-ph/0408515.

X. Wang, M. Tegmark, M. G. Santos, and L. Knox, **650**, 529 (2006), astro-ph/0501081.

V. Jelić, S. Zaroubi, P. Labropoulos, R. M. Thomas, G. Bernardi, M. A. Brentjens, A. G. de Bruyn, B. Ciardi, G. Harker, L. V. E. Koopmans, et al., **389**, 1319 (2008), 0804.1130.

E. Chapman and V. Jelić, in *The Cosmic 21-cm Revolution* (IOP Publishing, 2019), 2514-3433, pp. 6–1 to 6–29, ISBN 978-0-7503-2236-2, URL https://dx.doi.org/10.1088/2514-3433/ab4a73ch6.

A. Liu and M. Tegmark, Monthly Notices of the Royal Astronomical Society **419**, 3491–3504 (2012), ISSN 1365-2966.

A. E. Lanman, J. C. Pober, N. S. Kern, E. de Lera Acedo, D. R. DeBoer, and N. Fagnoni, **494**, 3712 (2020), 1910.10573.

K. W. Masui, E. R. Switzer, N. Banavar, K. Bandura, C. Blake, L. M. Calin, T. C. Chang, X. Chen, Y. C. Li, Y. W. Liao, et al., **763**, L20 (2013), 1208.0331.

L. Wolz, A. Pourtsidou, K. W. Masui, T.-C. Chang, J. E. Bautista, E.-M. Müller, S. Avila, D. Bacon, W. J. Percival, S. Cunnington, et al., **510**, 3495 (2022), 2102.04946.

C. J. Anderson, N. J. Luciw, Y. C. Li, C. Y. Kuo, J. Yadav, K. W. Masui, T. C. Chang, X. Chen, N. Oppermann, Y. W. Liao, et al., **476**, 3382 (2018), 1710.00424.

A. Chowdhury, N. Kanekar, B. Das, K. S. Dwarakanath, and S. Sethi, **913**, L24 (2021), 2105.06773.

S. Cunnington, Y. Li, M. G. Santos, J. Wang, I. P. Carucci, M. O. Irfan, A. Pourtsidou, M. Spinelli, L. Wolz, P. S. Soares, et al., **518**, 6262 (2023), 2206.01579.

M. Amiri, K. Bandura, T. Chen, M. Deng, M. Dobbs, M. Fandino, S. Foreman, M. Halpern, A. S. Hill, G. Hinshaw, et al., The Astrophysical Journal **947**, 16 (2023), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/acb13f`.

M. Amiri, K. Bandura, A. Chakraborty, M. Dobbs, M. Fandino, S. Foreman, H. Gan, M. Halpern, A. S. Hill, G. Hinshaw, et al., **963**, 23 (2024), 2309.04404.

S. Paul, M. G. Santos, Z. Chen, and L. Wolz, *A first detection of neutral hydrogen intensity mapping on mpc scales at* $z \approx 0.32$ *and* $z \approx 0.44$ (2023), 2301.11943.

J. L. Bernal and E. D. Kovetz, The Astronomy and Astrophysics Review **30** (2022), ISSN 1432-0754, URL `http://dx.doi.org/10.1007/s00159-022-00143-0`.

Y. Gong, A. Cooray, M. B. Silva, M. Zemcov, C. Feng, M. G. Santos, O. Dore, and X. Chen, **835**, 273 (2017), 1610.09060.

S. Yang, R. S. Somerville, A. R. Pullen, G. Popping, P. C. Breysse, and A. S. Maniyar, **911**, 132 (2021), 2009.11933.

K. S. Karkare, A. Moradinezhad Dizgah, G. K. Keating, P. Breysse, and D. T. Chung, arXiv e-prints arXiv:2203.07258 (2022a), 2203.07258.

A. Lidz, S. R. Furlanetto, S. P. Oh, J. Aguirre, T.-C. Chang, O. Doré, and J. R. Pritchard, **741**, 70 (2011), `1104.4800`.

C. L. Carilli, **730**, L30 (2011), `1102.0745`.

G. K. Keating, G. C. Bower, D. P. Marrone, D. R. DeBoer, C. Heiles, T.-C. Chang, J. E. Carlstrom, C. H. Greer, D. Hawkins, J. W. Lamb, et al., **814**, 140 (2015), `1510.06744`.

G. K. Keating, D. P. Marrone, G. C. Bower, E. Leitch, J. E. Carlstrom, and D. R. DeBoer, **830**, 34 (2016), `1605.03971`.

G. K. Keating, D. P. Marrone, G. C. Bower, and R. P. Keenan, **901**, 141 (2020), `2008.08087`.

K. S. Karkare, A. J. Anderson, P. S. Barry, B. A. Benson, J. E. Carlstrom, T. Cecil, C. L. Chang, M. A. Dobbs, M. Hollister, G. K. Keating, et al., Journal of Low Temperature Physics **209**, 758 (2022b), `2111.04631`.

K. A. Cleary, J. Borowska, P. C. Breysse, M. Catha, D. T. Chung, S. E. Church, C. Dickinson, H. K. Eriksen, M. K. Foss, J. O. Gundersen, et al., **933**, 182 (2022), `2111.05927`.

H. T. Ihle, J. Borowska, K. A. Cleary, H. K. Eriksen, M. K. Foss, S. E. Harper, J. Kim, J. G. S. Lunde, L. Philip, M. Rasmussen, et al., The Astrophysical Journal **933**, 185 (2022), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/ac63c5`.

J. G. S. Lunde, N. O. Stutzer, P. C. Breysse, D. T. Chung, K. A. Cleary, D. A. Dunne, H. K. Eriksen, S. E. Harper, H. T. Ihle, J. W. Lamb, et al., *Comap pathfinder – season 2 results i. improved data selection and processing* (2024), `2406.07510`, URL `https://arxiv.org/abs/2406.07510`.

N. O. Stutzer, J. G. S. Lunde, P. C. Breysse, D. T. Chung, K. A. Cleary, D. A. Dunne, H. K. Eriksen, H. T. Ihle, H. Padmanabhan, D. Tolgay, et al., arXiv e-prints arXiv:2406.07511 (2024), `2406.07511`.

D. T. Chung, P. C. Breysse, K. A. Cleary, D. A. Dunne, J. G. S. Lunde, H. Padmanabhan, N. O. Stutzer, D. Tolgay, J. R. Bond, S. E. Church, et al., arXiv e-prints arXiv:2406.07512 (2024), `2406.07512`.

A. Roy, N. Battaglia, and A. R. Pullen, arXiv e-prints arXiv:2406.07861 (2024), `2406.07861`.

M. B. Silva, M. G. Santos, Y. Gong, A. Cooray, and J. Bock, **763**, 132 (2013), `1205.1493`.

A. R. Pullen, O. Doré, and J. Bock, **786**, 111 (2014), `1309.2295`.

R. A. C. Croft, J. Miralda-Escudé, Z. Zheng, M. Blomqvist, and M. Pieri, **481**, 1320 (2018), `1806.06050`.

L. Mas-Ribas and T.-C. Chang, **101**, 083032 (2020), `2002.04107`.

R. Kakuma, M. Ouchi, Y. Harikane, Y. Ono, A. K. Inoue, Y. Komiyama, H. Kusakabe, C.-H. Lee, Y. Matsuda, Y. Matsuoka, et al., **916**, 22 (2021), `1906.00173`.

P. Renard, E. Gaztanaga, R. Croft, L. Cabayol, J. Carretero, M. Eriksen, E. Fernandez, J. García-Bellido, R. Miquel, C. Padilla, et al., **501**, 3883 (2021), `2006.07177`.

X. Lin, Z. Zheng, and Z. Cai, **262**, 38 (2022), `2207.10682`.

S. Kikuchihara, Y. Harikane, M. Ouchi, Y. Ono, T. Shibuya, R. Itoh, R. Kakuma, A. K. Inoue, H. Kusakabe, K. Shimasaku, et al., **931**, 97 (2022), `2108.09288`.

Y. Gong, A. Cooray, M. Silva, M. G. Santos, J. Bock, C. M. Bradford, and M. Zemcov, **745**, 49 (2012), `1107.3553`.

A. T. Crites, J. J. Bock, C. M. Bradford, T. C. Chang, A. R. Cooray, L. Duband, Y. Gong, S. Hailey-Dunsheath, J. Hunacek, P. M. Koch, et al., in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, edited by W. S. Holland and J. Zmuidzinas

(2014), vol. 9153 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 91531W.

B. D. Uzgil, J. E. Aguirre, C. M. Bradford, and A. Lidz, **793**, 116 (2014), `1407.4860`.

B. Yue, A. Ferrara, A. Pallottini, S. Gallerani, and L. Vallini, **450**, 3829 (2015), `1504.06530`.

C. Karoumpis, B. Magnelli, E. Romano-Díaz, M. Haslbauer, and F. Bertoldi, **659**, A12 (2022), `2111.12847`.

H. Padmanabhan, **488**, 3014 (2019), `1811.01968`.

D. T. Chung, M. P. Viero, S. E. Church, and R. H. Wechsler, **892**, 51 (2020), `1812.08135`.

A. R. Pullen, P. Serra, T.-C. Chang, O. Doré, and S. Ho, **478**, 1911 (2018), `1707.06172`.

S. Yang, A. R. Pullen, and E. R. Switzer, **489**, L53 (2019), `1904.01180`.

Concerto Collaboration, P. Ade, M. Aravena, E. Barria, A. Beelen, A. Benoit, M. Béthermin, J. Bounmy, O. Bourrion, G. Bres, et al., **642**, A60 (2020).

P. A. R. Ade, C. J. Anderson, E. M. Barrentine, N. G. Bellis, A. D. Bolatto, P. C. Breysse, B. T. Bulcha, G. Cataldo, J. A. Connors, P. W. Cursey, et al., Journal of Low Temperature Physics **199**, 1027 (2020), `1912.07118`.

J. Vieira, J. Aguirre, C. M. Bradford, J. Filippini, C. Groppi, D. Marrone, M. Bethermin, T.-C. Chang, M. Devlin, O. Dore, et al., arXiv e-prints arXiv:2009.14340 (2020), `2009.14340`.

G. Cataldo, P. A. R. Ade, C. J. Anderson, A. Barlis, E. M. Barrentine, N. G. Bellis, A. D. Bolatto, P. C. Breysse, B. T. Bulcha, J. A. Connors, et al., in *Ground-based and Airborne Telescopes VIII*, edited by H. K. Marshall, J. Spyromilio, and T. Usuda (2020), vol. 11445 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 1144524.

T. Essinger-Hileman, P. Ade, C. J. Anderson, A. Barlis, E. M. Barrentine, J. Beeman, N. Bellis, A. D. Bolatto, P. C. Breysse, B. T. Bulcha, et al., in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*, edited by J. Zmuidzinas and J.-R. Gao (2022), vol. 12190 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 1219009.

M. Béthermin, A. Gkogkou, M. Van Cuyck, G. Lagache, A. Beelen, M. Aravena, A. Benoit, J. Bounmy, M. Calvo, A. Catalano, et al., **667**, A156 (2022), `2204.12827`.

CCAT-Prime Collaboration, M. Aravena, J. E. Austermann, K. Basu, N. Battaglia, B. Beringue, F. Bertoldi, F. Bigiel, J. R. Bond, P. C. Breysse, et al., **264**, 7 (2023), `2107.10364`.

M. Van Cuyck, N. Ponthieu, G. Lagache, A. Beelen, M. Béthermin, A. Gkogkou, M. Aravena, A. Benoit, J. Bounmy, M. Calvo, et al., **676**, A62 (2023), `2306.01568`.

H. Padmanabhan, P. Breysse, A. Lidz, and E. R. Switzer, **515**, 5813 (2022), `2105.12148`.

O. Doré, J. Bock, M. Ashby, P. Capak, A. Cooray, R. de Putter, T. Eifler, N. Flagey, Y. Gong, S. Habib, et al., arXiv e-prints arXiv:1412.4872 (2014), `1412.4872`.

P. C. Breysse, S. Foreman, L. C. Keating, J. Meyers, and N. Murray, **105**, 083009 (2022), `2104.06422`.

G. Sun, T. C. Chang, B. D. Uzgil, J. J. Bock, C. M. Bradford, V. Butler, T. Caze-Cortes, Y. T. Cheng, A. Cooray, A. T. Crites, et al., **915**, 33 (2021), `2012.09160`.

J. Mirocha, A. Liu, and P. La Plante, **516**, 4123 (2022), `2205.14168`.

J. Parsons, L. Mas-Ribas, G. Sun, T.-C. Chang, M. O. Gonzalez, and R. H. Mebane, **933**, 141 (2022), `2112.06407`.

J. L. Bernal and E. D. Kovetz, **30**, 5 (2022), `2206.15377`.

G. Sun, L. Mas-Ribas, T.-C. Chang, S. R. Furlanetto, R. H. Mebane, M. O. Gonzalez, J. Parsons, and A. C. Trapp, **950**, 40 (2023), `2206.14186`.

P. Horlaville, D. T. Chung, J. R. Bond, and L. Liang, **531**, 2958 (2024), `2309.15733`.

J. Fonseca, M. B. Silva, M. G. Santos, and A. Cooray, **464**, 1948 (2017), `1607.05288`.

J. L. Bernal, P. C. Breysse, H. Gil-Marín, and E. D. Kovetz, **100**, 123522 (2019a), `1907.10067`.

J. L. Bernal, P. C. Breysse, and E. D. Kovetz, **123**, 251301 (2019b), `1907.10065`.

E. Schaan and M. White, **2021**, 067 (2021a), `2103.01971`.

A. Moradinezhad Dizgah, G. K. Keating, and A. Fialkov, **870**, L4 (2019), `1801.10178`.

A. Moradinezhad Dizgah, G. K. Keating, K. S. Karkare, A. Crites, and S. R. Choudhury, **926**, 137 (2022a), `2110.00014`.

A. Moradinezhad Dizgah, F. Nikakhtar, G. K. Keating, and E. Castorina, **2022**, 026 (2022b), `2111.03717`.

A. R. Pullen, T.-C. Chang, O. Doré, and A. Lidz, **768**, 15 (2013), `1211.1397`.

T. C. Chang, Y. Gong, M. Santos, M. B. Silva, J. Aguirre, O. Doré, and J. Pritchard, in *Advancing Astrophysics with the Square Kilometre Array (AASKA14)* (2015), p. 4, `1501.04654`.

G. Sun, B. S. Hensley, T.-C. Chang, O. Doré, and P. Serra, **887**, 142 (2019), `1907.02999`.

E. Schaan and M. White, **2021**, 068 (2021b), `2103.01964`.

C. J. Anderson, E. R. Switzer, and P. C. Breysse, **514**, 1169 (2022), `2202.00203`.

A. S. Maniyar, E. Schaan, and A. R. Pullen, **105**, 083509 (2022), `2106.09005`.

G. Sato-Polito, N. Kokron, and J. L. Bernal, **526**, 5883 (2023), `2212.08056`.

L. Mas-Ribas, G. Sun, T.-C. Chang, M. O. Gonzalez, and R. H. Mebane, **950**, 39 (2023), 2206.14185.

A. Roy, D. Valentín-Martínez, K. Wang, N. Battaglia, and A. van Engelen, **957**, 87 (2023), 2304.06748.

M. Lujan Niemeyer, J. L. Bernal, and E. Komatsu, **958**, 4 (2023), 2307.08475.

H. Fronenberg, A. S. Maniyar, A. Liu, and A. R. Pullen, **132**, 241001 (2024a), 2309.07215.

H. Fronenberg, A. S. Maniyar, A. R. Pullen, and A. Liu, **109**, 123518 (2024b), 2309.06477.

A. Beane and A. Lidz, **867**, 26 (2018), 1806.02796.

A. Beane, F. Villaescusa-Navarro, and A. Lidz, **874**, 133 (2019), 1811.10609.

L. McBride and A. Liu, arXiv e-prints arXiv:2308.00749 (2023), 2308.00749.

H. Padmanabhan, **523**, 3503 (2023), 2212.08077.

S. Dumitru, G. Kulkarni, G. Lagache, and M. G. Haehnelt, Monthly Notices of the Royal Astronomical Society **485**, 3486–3498 (2019), ISSN 0035-8711.

K. Moriwaki, A. Beane, and A. Lidz, **530**, 3183 (2024), 2404.08266.

A. Roy and N. Battaglia, **969**, 2 (2024), 2312.08471.

M. Suginohara, T. Suginohara, and D. N. Spergel, **512**, 547 (1999), astro-ph/9803236.

K. Basu, C. Hernández-Monteagudo, and R. A. Sunyaev, **416**, 447 (2004), astro-ph/0311620.

B. Yue, A. Ferrara, A. Pallottini, S. Gallerani, and L. Vallini, Monthly Notices of the Royal Astronomical Society **450**, 3829 (2015), URL https://doi.org/10.1093%2Fmnras%2Fstv933.

E. Visbal, H. Trac, and A. Loeb, Journal of Cosmology and Astroparticle Physics **2011**, 010 (2011), URL `https://doi.org/10.1088%2F1475-7516%2F2011%2F08%2F010`.

Y. Gong, M. Silva, A. Cooray, and M. G. Santos, The Astrophysical Journal **785**, 72 (2014), URL `https://doi.org/10.1088%2F0004-637x%2F785%2F1%2F72`.

M. Silva, M. G. Santos, A. Cooray, and Y. Gong, **806**, 209 (2015), `1410.4808`.

P. C. Breysse, E. D. Kovetz, and M. Kamionkowski, **452**, 3408 (2015), `1503.05202`.

G. Sun, L. Moncelsi, M. P. Viero, M. B. Silva, J. Bock, C. M. Bradford, T.-C. Chang, Y.-T. Cheng, A. R. Cooray, A. Crites, et al., The Astrophysical Journal **856**, 107 (2018), URL `https://doi.org/10.3847%2F1538-4357%2Faab3e3`.

M. B. Silva, B. Baumschlager, K. A. Cleary, P. C. Breysse, D. T. Chung, H. T. Ihle, H. Padmanabhan, L. C. Keating, J. Kim, and L. Philip, arXiv e-prints arXiv:2111.05354 (2021), `2111.05354`.

A. Kogut, E. Dwek, and S. H. Moseley, **806**, 234 (2015), `1505.00266`.

J. L. Bernal and A. Baleato Lizancos, arXiv e-prints arXiv:2406.12979 (2024), `2406.12979`.

A. Lidz and J. Taylor, **825**, 143 (2016), `1604.05737`.

A. Liu, Y. Zhang, and A. R. Parsons, **833**, 242 (2016), `1609.04401`.

Y.-T. Cheng, T.-C. Chang, and J. J. Bock, The Astrophysical Journal **901**, 142 (2020), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/abb023`.

E. Visbal and A. Loeb, **2010**, 016 (2010), `1008.3178`.

D. R. DeBoer, A. R. Parsons, J. E. Aguirre, P. Alexander, Z. S. Ali, A. P. Beardsley, G. Bernardi, J. D. Bowman, R. F. Bradley, C. L. Carilli, et al., **129**, 045001 (2017), `1606.07473`.

L. M. Berkhout, D. C. Jacobs, Z. Abdurashidova, T. Adams, J. E. Aguirre, P. Alexander, Z. S. Ali, R. Baartman, Y. Balfour, A. P. Beardsley, et al., **136**, 045002 (2024), 2401.04304.

A. de Oliveira-Costa, M. Tegmark, B. M. Gaensler, J. Jonas, T. L. Landecker, and P. Reich, Monthly Notices of the Royal Astronomical Society **388**, 247–260 (2008), ISSN 1365-2966, URL `http://dx.doi.org/10.1111/j.1365-2966.2008.13376.x`.

H. Zheng, M. Tegmark, J. S. Dillon, D. A. Kim, A. Liu, A. R. Neben, J. Jonas, P. Reich, and W. Reich, Monthly Notices of the Royal Astronomical Society **464**, 3486–3497 (2016), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stw2525`.

M. Pagano and A. Liu, **498**, 373 (2020), 2005.11326.

J. C. Pober, B. J. Hazelton, A. P. Beardsley, N. A. Barry, Z. E. Martinot, I. S. Sullivan, M. F. Morales, M. E. Bell, G. Bernardi, N. D. R. Bhat, et al., **819**, 8 (2016), 1601.06177.

Z. Xu, J. N. Hewitt, K.-F. Chen, H. Kim, J. S. Dillon, N. S. Kern, M. F. Morales, B. J. Hazelton, R. Byrne, N. Fagnoni, et al., **938**, 128 (2022), 2204.06021.

J. Burba, P. H. Sims, and J. C. Pober, **520**, 4443 (2023), 2302.04058.

J. Tan, A. Liu, N. S. Kern, Z. Abdurashidova, J. E. Aguirre, P. Alexander, Z. S. Ali, Y. Balfour, A. P. Beardsley, G. Bernardi, et al., **255**, 26 (2021), 2103.09941.

J. L. B. Line, D. A. Mitchell, B. Pindor, J. L. Riding, B. McKinley, R. L. Webster, C. M. Trott, N. Hurley-Walker, and A. R. Offringa, **37**, e027 (2020), 2005.09316.

N. Hurley-Walker, J. R. Callingham, P. J. Hancock, T. M. O. Franzen, L. Hindson, A. D. Kapińska, J. Morgan, A. R. Offringa, R. B. Wayth, C. Wu, et al., **464**, 1146 (2017), 1610.08318.

T. M. O. Franzen, C. A. Jackson, A. R. Offringa, R. D. Ekers, R. B. Wayth, G. Bernardi, J. D. Bowman, F. Briggs, R. J. Cappallo, A. A. Deshpande, et al., Monthly Notices of the Royal

Astronomical Society **459**, 3314–3325 (2016), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stw823`.

T. M. O. Franzen, T. Vernstrom, C. A. Jackson, N. Hurley-Walker, R. D. Ekers, G. Heald, N. Seymour, and S. V. White, Publications of the Astronomical Society of Australia **36** (2019), ISSN 1448-6083, URL `http://dx.doi.org/10.1017/pasa.2018.52`.

G. Popping, E. van Kampen, R. Decarli, M. Spaans, R. S. Somerville, and S. C. Trager, **461**, 93 (2016), `1602.02761`.

C. W. Fabjan, F. M. Pipkin, and M. Silverman, Phys. Rev. Lett. **26**, 347 (1971), URL `https://link.aps.org/doi/10.1103/PhysRevLett.26.347`.

A. Kramida, Atomic Data and Nuclear Data Tables **96**, 586 (2010), ISSN 0092-640X, URL `https://www.sciencedirect.com/science/article/pii/S0092640X10000458`.

J. S. Dillon and A. R. Parsons, **826**, 181 (2016), `1602.06259`.

Y. G. Zhang, A. Liu, and A. R. Parsons, **852**, 110 (2018), `1707.07668`.

J. C. Pober, A. Liu, J. S. Dillon, J. E. Aguirre, J. D. Bowman, R. F. Bradley, C. L. Carilli, D. R. DeBoer, J. N. Hewitt, D. C. Jacobs, et al., The Astrophysical Journal **782**, 66 (2014), ISSN 1538-4357, URL `http://dx.doi.org/10.1088/0004-637X/782/2/66`.

D. R. DeBoer, A. R. Parsons, J. E. Aguirre, P. Alexander, Z. S. Ali, A. P. Beardsley, G. Bernardi, J. D. Bowman, R. F. Bradley, C. L. Carilli, et al., Publications of the Astronomical Society of the Pacific **129**, 045001 (2017), ISSN 1538-3873, URL `http://dx.doi.org/10.1088/1538-3873/129/974/045001`.

P. C. Breysse and R. M. Alexandroff, Monthly Notices of the Royal Astronomical Society **490**, 260–273 (2019), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stz2534`.

F. J. Harris, Proceedings of the IEEE **66**, 51 (1978).

N. Thyagarajan, N. Udaya Shankar, R. Subrahmanyan, W. Arcus, G. Bernardi, J. D. Bowman, F. Briggs, J. D. Bunton, R. J. Cappallo, B. E. Corey, et al., **776**, 6 (2013), `1308.0565`.

N. Thyagarajan, A. R. Parsons, D. R. DeBoer, J. D. Bowman, A. M. Ewall-Wice, A. R. Neben, and N. Patra, **825**, 9 (2016), `1603.08958`.

A. R. Parsons, J. C. Pober, J. E. Aguirre, C. L. Carilli, D. C. Jacobs, and D. F. Moore, **756**, 165 (2012), `1204.4749`.

M. F. Morales, B. Hazelton, I. Sullivan, and A. Beardsley, **752**, 137 (2012), `1202.3830`.

C. M. Trott, R. B. Wayth, and S. J. Tingay, **757**, 101 (2012), `1208.0646`.

J. C. Pober, A. R. Parsons, J. E. Aguirre, Z. Ali, R. F. Bradley, C. L. Carilli, D. DeBoer, M. Dexter, N. E. Gugliucci, D. C. Jacobs, et al., **768**, L36 (2013), `1301.7099`.

J. S. Dillon, A. Liu, C. L. Williams, J. N. Hewitt, M. Tegmark, E. H. Morgan, A. M. Levine, M. F. Morales, S. J. Tingay, G. Bernardi, et al., **89**, 023002 (2014), `1304.4229`.

A. Liu, A. R. Parsons, and C. M. Trott, **90**, 023018 (2014a), `1404.2596`.

E. Chapman, S. Zaroubi, F. B. Abdalla, F. Dulwich, V. Jelić, and B. Mort, **458**, 2928 (2016).

M. Tegmark, **480**, L87 (1997a), `astro-ph/9611130`.

J. R. Bond, A. H. Jaffe, and L. Knox, **57**, 2117 (1998), `astro-ph/9708203`.

A. Liu and M. Tegmark, **83**, 103006 (2011), `1103.0281`.

J. R. Shaw, K. Sigurdson, M. Sitwell, A. Stebbins, and U.-L. Pen, **91**, 083514 (2015), `1401.2095`.

A. Liu, A. R. Parsons, and C. M. Trott, **90**, 023019 (2014b), `1404.4372`.

L. Zhang, E. F. Bunn, A. Karakci, A. Korotkov, P. M. Sutter, P. T. Timbie, G. S. Tucker, and B. D. Wandelt, **222**, 3 (2016), 1505.04146.

P. H. Sims, L. Lentati, P. Alexander, and C. L. Carilli, **462**, 3069 (2016), 1607.07628.

P. H. Sims and J. C. Pober, **488**, 2904 (2019), 1907.02608.

P. H. Sims, L. Lentati, J. C. Pober, C. Carilli, M. P. Hobson, P. Alexander, and P. M. Sutter, **484**, 4152 (2019).

J. Burba, P. Bull, M. J. Wilensky, F. Kennedy, H. Garsden, and K. A. Glasscock, arXiv e-prints arXiv:2403.13767 (2024), 2403.13767.

Y.-T. Cheng, K. Wang, B. D. Wandelt, T.-C. Chang, and O. Dore, arXiv e-prints arXiv:2403.19740 (2024), 2403.19740.

T. Y. Li, R. H. Wechsler, K. Devaraj, and S. E. Church, **817**, 169 (2016), 1503.08833.

J. C. Pober, **447**, 1705 (2015), 1411.2050.

B. Gold, C. L. Bennett, R. S. Hill, G. Hinshaw, N. Odegard, L. Page, D. N. Spergel, J. L. Weiland, J. Dunkley, M. Halpern, et al., **180**, 265 (2009), 0803.0715.

B. Gold, N. Odegard, J. L. Weiland, R. S. Hill, A. Kogut, C. L. Bennett, G. Hinshaw, X. Chen, J. Dunkley, M. Halpern, et al., **192**, 15 (2011), 1001.4555.

Planck Collaboration, R. Adam, P. A. R. Ade, N. Aghanim, M. I. R. Alves, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, et al., **594**, A10 (2016), 1502.01588.

M. J. Wilensky, J. Brown, and B. J. Hazelton, **521**, 5191 (2023b), 2211.13576.

P. Billingsley, *Wiley series in probability and statistics* (2012).

D. W. Scott, *Multivariate density estimation: theory, practice, and visualization* (John Wiley & Sons, 2015).

M. Tegmark, A. J. S. Hamilton, M. A. Strauss, M. S. Vogeley, and A. S. Szalay, **499**, 555 (1998), astro-ph/9708020.

T. Bayes, Philosophical transactions of the Royal Society of London **53**, 370 (1763).

M. Tegmark, **55**, 5895 (1997b), astro-ph/9611174.

M. J. Wilensky, M. F. Morales, B. J. Hazelton, N. Barry, R. Byrne, and S. Roy, Publications of the Astronomical Society of the Pacific **131**, 114507 (2019), ISSN 1538-3873, URL http://dx.doi.org/10.1088/1538-3873/ab3cad.

V. Jelić, S. Zaroubi, P. Labropoulos, G. Bernardi, A. G. De Bruyn, and L. V. E. Koopmans, Monthly Notices of the Royal Astronomical Society **409**, 1647–1659 (2010), ISSN 0035-8711, URL http://dx.doi.org/10.1111/j.1365-2966.2010.17407.x.

D. F. Moore, J. E. Aguirre, A. R. Parsons, D. C. Jacobs, and J. C. Pober, **769**, 154 (2013), 1302.0876.

Z. E. Martinot, J. E. Aguirre, S. A. Kohn, and I. Q. Washington, The Astrophysical Journal **869**, 79 (2018), ISSN 1538-4357, URL http://dx.doi.org/10.3847/1538-4357/aaeac6.

N. Orosz, J. S. Dillon, A. Ewall-Wice, A. R. Parsons, and N. Thyagarajan, Monthly Notices of the Royal Astronomical Society **487**, 537–549 (2019), ISSN 1365-2966, URL http://dx.doi.org/10.1093/mnras/stz1287.

R. Byrne, M. F. Morales, B. Hazelton, W. Li, N. Barry, A. P. Beardsley, R. Joseph, J. Pober, I. Sullivan, and C. Trott, The Astrophysical Journal **875**, 70 (2019), ISSN 1538-4357, URL http://dx.doi.org/10.3847/1538-4357/ab107d.

N. Barry, B. Hazelton, I. Sullivan, M. F. Morales, and J. C. Pober, Monthly Notices of the Royal Astronomical Society **461**, 3135–3144 (2016), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stw1380`.

E. Rath, R. Pascua, A. T. Josaitis, A. Ewall-Wice, N. Fagnoni, E. de Lera Acedo, Z. E. Martinot, Z. Abdurashidova, T. Adams, J. E. Aguirre, et al., *Investigating mutual coupling in the hydrogen epoch of reionization array and mitigating its effects on the 21-cm power spectrum* (2024), `2406.08549`, URL `https://arxiv.org/abs/2406.08549`.

M. K. Foss, H. T. Ihle, J. Borowska, K. A. Cleary, H. K. Eriksen, S. E. Harper, J. Kim, J. W. Lamb, J. G. S. Lunde, L. Philip, et al., The Astrophysical Journal **933**, 184 (2022), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/ac63ca`.

Intermezzo 2

In the previous chapter, we explored whether cross-correlations can effectively mitigate the foreground contamination of LIM measurements. We learned that while cross-correlations help suppress uncorrelated contaminants and reduce the barrier to a detection, foreground removal, particularly that of 21 cm, is necessary. Whether performing a cross-correlation or an auto-correlation measurement, foreground removal is a non-negotiable data processing step. As briefly mentioned in Chapter 1, existing mitigation methods have yet to yield a high-z LIM detection, highlighting the need for creative new approaches. Motivated by the success of cross-correlation techniques in general, in this chapter, we introduce a novel foreground removal method designed to eliminate broadband contaminants, such as synchrotron emission and the CIB from LIM observations. This method leverages a set of intra-dataset frequency-frequency cross-correlations to empirically estimate the foreground power, offering a fully data-driven approach. A key advantage of this technique is its robustness to beam and systematic effects which are common challenges for traditional foreground removal methods. This presents an exciting avenue toward achieving the necessary foreground suppression to unlock the cross-correlation measurements described in the previous chapter.

Chapter 3

A Data-Driven Technique to Mitigate the Foregrounds of Line Intensity Maps

Hannah Fronenberg[1,2], Adrian Liu[1,2]

[1]Department of Physics, McGill University, Montréal, QC, Canada

[2]Trottier Space Institute, Montréal, QC, Canada

Abstract

Line intensity mapping (LIM) is an emerging observational technique for mapping the large-scale structure of the Universe across a broad redshift range. Multiple surveys targeting various spectral lines are now underway or soon to come online, aiming to make some of the first early statistical measurements. Despite this progress, foreground contamination remains a significant challenge for LIM across virtually all wavelengths, and many existing mitigation techniques have struggled to yield successful detections. In particular, astrophysical foregrounds with broad frequency structure pose a major obstacle. In this work, we present a novel foreground mitigation scheme designed to remove broadband contaminants in LIM observations while accounting for instrumental response and systematic effects. Using a fully data-driven approach, we demonstrate that a trio of intra-dataset frequency-frequency cross-correlations enables an empirical estimate of the foreground power spectrum, which can then be subtracted from the data. As a case study, we simulate 21 cm observations under various experimental configurations, such as single-dish instruments and radio interferometers, while incorporating systematic contaminants. Under idealized conditions, we achieve percent-level signal residuals, demonstrating that nearly 4–5 orders of magnitude of foreground contamination can be removed in a single analysis step. Finally, we discuss the limitations of this method and identify scenarios where the estimator is expected to limited.

3.1 Introduction

Line-intensity mapping (LIM) is a promising new technique for studying a variety of science cases over a large portion of cosmic history. By targeting transition lines with known rest frequencies, one can map their total integrated emission, tracing cosmological structures such as galactic populations or the intergalactic medium (IGM). Many LIM surveys targeting redshifts $z \simeq 0 - 30$ focus on key spectral lines, including the 158 μm line of ionized carbon ([CII]), multiple rotational transitions of carbon monoxide (CO), and several hydrogen emission lines such as the Lyman-α (Lyα) and Hα lines, as well as the 21 cm hyperfine transition of neutral hydrogen (HI). Mapping galactic tracers allows for the study of the star formation history, of the geometry and expansion history of the Universe through the measurement of baryon acoustic oscillations (BAOs, Bernal et al. 2019; Fronenberg et al. 2024), and of the neutral gas fraction. At higher redshifts, the 21 cm line also offers a unique probe of the IGM during the epoch of reionization (EoR) and cosmic dawn, and can be used as a biased tracer of the matter density field well into the cosmic dark ages. The many lines available for study makes LIM a powerful probe of both cosmology and astrophysics throughout many key epochs that are currently poorly constrained.

In recent years, the field of LIM has experienced tremendous experimental progress, yielding both preliminary detections and stringent upper limits. Post EoR, there have been multiple measurements of the 21 cm line of HI in cross-correlation with large-scale structure surveys (Masui et al., 2013; Teng et al., 2013; Rhee et al., 2016; Anderson et al., 2018; Chowdhury et al., 2021; Wolz et al., 2022; Amiri et al., 2023; Amiri et al., 2024) as well as a tentative auto-spectrum measurement (Paul et al., 2023). At sub-mm wavelengths, CO has also been measured in cross-correlation (Pullen et al., 2013, 2018; Roy et al., 2024) while a number of experiments have placed upper limits on the CO luminosity around cosmic noon ($z \sim 2-3$) (Stutzer et al., 2024; Chung et al., 2024; Keating et al., 2016; Ihle et al., 2022). In the optical and infrared (IR) band, Lyα intensity mapping has been achieved through cross-correlation with galaxy surveys. Moving to higher redshifts, a host of upper limits have been placed on the 21 cm power spectrum during cosmic dawn and the EoR (Paciga et al., 2013; Gehlot et al., 2019; Patil et al., 2017; Mertens et al., 2020; Kolopanis et al., 2019; Cheng et al., 2018; Garsden et al., 2021; Ewall-Wice et al., 2016; Beardsley et al., 2016; Barry et al., 2019; Li et al., 2019; Trott et al., 2020; Abdurashidova et al., 2022; Rahimi et al., 2021). The tightest of these upper limits by the Hydrogen Epoch of Reionization Array (HERA, DeBoer et al. 2017; Berkhout et al. 2024) has allowed for the first constraints to be placed on the spin temperature of the IGM at $z \sim 8$ and subsequently on the efficiency of X-ray heating by early galaxies (Abdurashidova et al., 2022; HERA Collaboration et al., 2023; Lazare et al., 2024).

Despite this progress, LIM suffers from host of observational challenges, chief among them, foreground contamination. While many lines, particularly those observed at wavelengths ranging from the sub-mm to the optical, suffer from line interloper contamination (i.e. other lines that redshift into the same observed band as the target line), LIM observations across all frequencies suffer from continuum foregrounds originating from both the Milky Way and from extragalactic sources. For instance, 21 cm observations face some of the most severe foreground contamination by galactic synchrotron emission which is only worsened by the chromatic beam of radio interferometers and by non-trivial instrument systematics. These galactic radio foregrounds can dominate over the cosmological signal by as much as 4-5 orders of magnitude. Similarly, sub-mm observations suffer from cosmic infrared background (CIB) contamination, and IR observations of early times suffer from contamination from zodiacal light (i.e. sunlight that has been scattered by dust grains in the solar system).

In order to combat this, a multitude of foreground mitigation strategies have been proposed. Many exploit the spectral smoothness of the astrophysical contaminant, filtering out the smooth component of the data to reveal the fluctuating cosmological signal (Petrovic and Oh, 2011). In simulation, decomposition based methods such as principal component analysis (PCA) and singular value decomposition (SVD) have been shown to be successful at separating the signal from the foregrounds (Yue et al., 2015; Alonso et al., 2014). The main feature of these schemes is to decompose the data into an eigenbasis whose eigenvalues map to the explained variance of the field. The largest variance components (assumed to be associated with foregrounds) are removed. While these methods are widely used, instruments that have complicated beam patterns, for instance with frequency-dependent sidelobe structure, and systematic contaminants can introduce spectral structure into the data which has been shown to degrade the effectiveness of these mitigation strategies (Rath et al., 2024). Additionally, in all of these approaches where portions of the data are filtered or excised, there is always a risk of accidentally removing signal along with the contaminants. This has prompted extensive efforts to rigorously quantify signal loss and propagate its effects throughout the analysis (Pascua et al., 2024). While much progress has been made on developing novel beam and systematics aware removal strategies (Wang et al., 2024, 2022; Ding et al., 2024), foreground contamination, and its coupling to other observational effects, remains the overarching roadblock to precision LIM measurements.

In this paper, we propose a novel data-driven technique for removing the broadband contaminants of line intensity maps in the presence of complex instrument effects. This method estimates a foreground power spectrum from the data directly. The estimated foreground power can then be subtracted off the total estimated power to reveal the signal power. This method relies on only two key assumptions. First, that the narrow line emission being targeted is uncorrelated between widely spaced channels[1], while the broadband emitting foregrounds are highly correlated across the observing band. This ensures that the cross-correlation between two distant channels will contain no signal power. More explicitly, if we consider two slices of data from two widely separated frequency channels ν_i and ν_j, which contain both a signal (s) and a foreground (f) component, then their cross angular power spectrum will be proportional to

$$C_l(\nu_i, \nu_j) \propto \langle a^f_{lm}(\nu_i) a^{*f}_{lm}(\nu_j)\rangle \tag{3.1}$$

where a_{lm} is the spherical harmonic coefficients of the spherical harmonic decomposition of the data, and l is the angular multipole. Since, in the ensemble average limit, the signal and foreground components will be uncorrelated, and the signal components will be uncorrelated, the only correlations that remain are due to broadband foreground emission, as shown on the right hand side of Eq. 3.1 by the superscript "f".

Next, by adapting the auto spectrum estimator proposed by Beane et al. (2019) (hereon B19) it is possible to isolate the foreground only auto-power using a trio of such cross-correlations between well-separated frequency channels: i, j, and k

$$C^f_l(\nu_i, \nu_i) = \frac{C^f_l(\nu_i, \nu_j) C^f_l(\nu_i, \nu_k)}{C^f_l(\nu_j, \nu_k)}. \tag{3.2}$$

This allows for the estimate of the auto spectrum at frequency i using three cross-powers. This step assumes that the foregrounds are well described by this simple model

$$T^f(\hat{\mathbf{n}}, \nu) = \beta(\nu)\delta_m(\hat{\mathbf{n}}) \tag{3.3}$$

whereby a frequency dependent bias factor, $\beta(\nu)$, modulates the brightness of the emission being painted over some common underlying structure, $\delta_m(\hat{\mathbf{n}})$. For many sources of broadband contaminants, such as various sources of galactic emission and the CIB, we believe this

---

[1] In the following section we make precise exactly how widely these channels must be spaced.

assumption to be reasonable and even if this assumptions is not strictly obeyed we expect this mitigation scheme to work decently well in practice. The questions, of course, is how well and that is precisely the subject of this chapter.

In the what follows, we detail how this adapted B19 estimator (Eq. 3.2) can be used to generate foreground power spectrum estimates that can then be subtracted from the total measured power to reveal the underlying signal. In Section 3.2, we establish this formalism and present some of its key features specifically with regard to how it works in the presence of both a chromatic beam and systematics. With this formalism in hand, we showcase this estimator in action by producing mock 21 cm observations and clean them of their foreground contaminant, paying close attention to how this foreground estimator performs under realistic instrument response and in the presence of complex systematics. In Section 3.3 we outline the models used to simulate the signal, the foregrounds, and the instrument response and systematics. In Section 3.4 we present our results and finally provide concluding remarks in Section 3.5.

3.2 Formalism

3.2.1 Basic Foreground Removal Formalism

The foreground removal scheme we present here relies on only a few key assumptions, one of which is a linear foreground model. While in using this estimator, foregrounds are never explicitly modeled, the estimator we employ to achieve foreground removal relies on the ability to describe the foregrounds by Eq. 3.3. The frequency dependent factor

$$\beta(\nu) = \langle I_\nu \rangle b_\nu \tag{3.4}$$

where $\langle I_\nu \rangle$ is the mean intensity of the field, and b_ν is the "bias", modulates the brightness temperature of the foreground field as a function of frequency ν. The spatial variation of the field is dictated by the underlying mass distribution responsible for the emission, $\delta_m(\hat{\mathbf{n}})$. This should not be taken to mean the cosmological matter density field but rather the mass distribution of the contaminant. For instance, δ_m can represent the distribution of matter in Milky Way galaxy if that is the contaminant in question.

Next, this method seeks to exploit the differences in frequency to frequency spatial correlations between cosmological narrow line emission and broadband foreground emission. Line emission emanating from high redshift galaxies and the IGM has short frequency correlation lengths; for example, during the EoR, 21 cm emission is expected to have a correlation length of $\Delta\nu^s_{\mathrm{corr}} \sim 10\mathrm{MHz}$ (Santos et al., 2005). Conversely, foreground emission with broad spectral structure, for example galactic synchrotron emission, is highly correlated over much of the electromagnetic spectrum.

To use the concrete example of 21 cm observations of the EoR contaminated with galactic synchrotron radiation to illustrate this point, consider an observation of the differential brightness temperature, δT_b, that is decomposed into spherical harmonic modes

$$\delta T_b(\hat{\mathbf{n}}, \nu) = \sum_{l,m} a_{lm}(\nu) Y_{lm}(\hat{\mathbf{n}}). \tag{3.5}$$

The coefficient a_{lm} is comprised of both a signal and a foreground component, $a_{lm}(\nu) = a^s_{lm}(\nu) + a^f_{lm}(\nu)$, denoted by the superscript as in Section 3.1. The multi-frequency angular power spectrum (MAPS), $C_l(\nu, \nu')$, where $|\nu - \nu'| < \Delta\nu^s_{\mathrm{corr}}$, is therefore

$$\begin{aligned} C_l(\nu, \nu') &= \langle a_{lm}(\nu) a^*_{lm}(\nu') \rangle && \text{(3.6a)} \\ &= \langle a^s_{lm}(\nu) a^{s*}_{lm}(\nu') \rangle + \langle a^s_{lm}(\nu) a^{f*}_{lm}(\nu') \rangle && \text{(3.6b)} \\ &+ \langle a^f_{lm}(\nu) a^{s*}_{lm}(\nu') \rangle + \langle a^f_{lm}(\nu) a^{f*}_{lm}(\nu') \rangle \\ &= C^s_l(\nu, \nu') + 2C^{sf}_l(\nu, \nu') + C^f_l(\nu, \nu'). && \text{(3.6c)} \end{aligned}$$

In the infinite ensemble average, denoted by "⟨⟩", the signal and the foregrounds are uncorrelated, and the total angular power is comprised of the signal power plus the foreground power (i.e. the first and last terms of Eq. 3.6c). In practice, however, we never take true ensemble averages, but instead, average over the finite number of modes that are measured on the sky. This means that terms that average down to zero in the infinite ensemble average limit, do not average down in reality and leave behind spurious residual correlations. These residual signal-foreground cross-terms in Eq. 3.6b can be overwhelmingly bright compared to the signal. For example, if the 21 cm signal is $\sim$ 10 mK and the foregrounds are $\sim$ 100 K, then the cross-terms can be as much as two orders of magnitude brighter than the signal. Therefore, in practice, it is never adequate to only model and subtract the foreground auto-spectrum alone, all foreground cross-terms must also be treated. Importantly, these cross-terms arise from the coincidental correlation of the signal and the foregrounds, unique to the particular realization of the signal and of the foregrounds in our Universe. These terms, then, can never be accurately simulated even if the statistical model from which they are being simulated is absolutely correct.

Luckily, since the foreground spectrum obtained using Eq. 3.2 is estimated from the data directly, it yields both foreground auto- and cross-spectra. Following the derivation in Appendix A, all foreground components of the observed spectrum can be obtained using the adapted B19 estimator, that is, where a trio of frequency channels pairwise separated by more than $\Delta\nu^s_{\mathrm{corr}}$ are correlated. This separation requirement ensures that no signal correlations enter the foreground estimate. The cleaning channels ν_j and ν_k are correlated with a set of target frequencies ν, ν'. A single channel, ν can be targeted, but ideally ν indexes a small sub-band whose spectral width is less than the frequency correlation length of the signal. With all of this in mind, the foreground estimator becomes

$$\begin{aligned}\hat{C}^f_l(\nu,\nu') \quad = \quad & \langle a^f_{lm}(\nu)a^{f*}_{lm}(\nu')\rangle + \langle a^s_{lm}(\nu)a^{f*}_{lm}(\nu')\rangle + \langle a^f_{lm}(\nu)a^{s*}_{lm}(\nu')\rangle + \\ & \frac{\beta^f}{\beta^f_j}\langle a^s_{lm}(\nu_j)a^{f*}_{lm}(\nu)\rangle + \frac{\beta^f}{\beta^f_k}\langle a^s_{lm}(\nu_k)a^{f*}_{lm}(\nu)\rangle\end{aligned} \tag{3.7}$$

which we can rewrite as,

$$\hat{C}_l^f(\nu,\nu') = C_l^f(\nu,\nu') + 2C_l^{sf}(\nu,\nu') + \frac{\beta^f}{\beta_j^f}C_l^{sf}(\nu_j,\nu) + \frac{\beta^f}{\beta_k^f}C^{sf}(\nu_k,\nu). \tag{3.8}$$

The first two terms of the estimated foreground power are precisely the two foreground contaminated terms in Eq. 3.6c that need to be removed. The last two terms of the estimate are undesired and bias the estimate but can be suppressed by exploiting the fact that the estimated foreground power $\hat{C}_l^f(\nu,\nu')$ can be computed many times over using different pairs of cleaning channels ν_j,ν_k. By averaging over many cleaning channels across the observed band, these two terms can be suppressed by a factor of $1/\sqrt{N}$ where N is the number of channel pairs used to compute Eq. 3.8. Finally, the recovered signal power is simply the total power (Eq. 3.6c) minus the estimated foreground power (Eq. 3.8) which yields

$$\hat{C}_l^s(\nu,\nu') = C_l^s(\nu,\nu') - \sum_{j,k}\left[\frac{\beta^f}{\beta_j^f}C_l^{sf}(\nu_j,\nu) + \frac{\beta^f}{\beta_k^f}C^{sf}(\nu_k,\nu)\right] \tag{3.9}$$

where here we explicitly write the sum over cleaning channels ν_j,ν_k.

At this stage, this multi-frequency angular power spectrum (MAPS), $\hat{C}_l^s(\nu,\nu')$, can be used to compute the more standard power spectra. Following Datta et al. (2007), the MAPS is then used to construct the frequency-difference angular power spectrum $C_l(\Delta\nu) = C_l(\nu-\nu')$ which is related to the standard cylindrical power spectrum $P(k_\perp,k_\parallel)$ as a function of comoving wave number perpendicular, $k_\perp$, and parallel, $k_\parallel$, to the line-of-sight (LOS) by

$$P(k_\perp,k_\parallel) = r^2r'\int_{-\infty}^{\infty}d(\Delta\nu)e^{-ik_\parallel r'\Delta\nu}C_l(\Delta\nu) \tag{3.10}$$

where r is the comoving distance to the central frequency, ν_c, of the target range, $r' = dr/d\nu$ evaluated at ν_c (Mondal et al., 2018). Here we have also employed the flat sky approximation where the angular multipole, l, is proportional to the perpendicular wave number $k_\perp$. Further averaging in bins of constant $k = \sqrt{k_{\parallel}^2 + k_{\perp}^2}$ yields the spherically averaged power spectrum, $P(k)$.

It should be noted that the standard power spectrum statistic is yielded from the MAPS only when the signal (here signal means whatever has been measured, in this case, the cosmological signal and the foregrounds) is statistically stationary in the LOS direction. Foregrounds generally do not satisfy these properties. Nonetheless, even with a frequency evolving sky, Eq. 3.10 can always be used to compute a power spectrum-like statistic, albeit with clear artifacts that are present due to the non-stationarity. We will see such features in Section 3.4.

3.2.2 Foreground Removal In The Presence Of A Chromatic Beam and Systematics

Perhaps the most desirable feature of this estimator is its ability not simply to obtain a foreground estimate, but to obtain an estimate of the foreground power *as seen by the instrument*. Consider now, the sky observed with an instrument that has a frequency-dependent point spread function (PSF, also known as the synthesized beam). The observed brightness temperature is therefore

$$\delta T_b(\hat{\mathbf{n}}, \nu) = \sum_{l,m} B_{lm}(\nu) a_{lm}(\nu) Y_{lm}(\hat{\mathbf{n}}) \tag{3.11}$$

where $B_{lm}(\nu)$ is the beam. Since this beam factor is not a random variable, it sails out of the ensemble average and Eq. 3.2 becomes

$$\hat{C}_l^{f,\text{obs}}(\nu, \nu') = \frac{B_l(\nu) B_l(\nu_j) B_l(\nu') B_l(\nu_k)}{B_l(\nu_j) B_l(\nu_k)} \frac{C_l^f(\nu, \nu_j) C_l^f(\nu', \nu_k)}{C_l^f(\nu_j, \nu_k)} \tag{3.12a}$$

$$= B_l(\nu) B_l(\nu') \hat{C}_l^f(\nu, \nu'), \tag{3.12b}$$

a fully empirical estimate of the foreground power.

It then directly follows that if there are any systematics, S, that are linear in frequency, that is,

$$\delta T_b(\hat{\mathbf{n}}, \nu) = \sum_{l,m} S_{lm}(\nu) B_{lm}(\nu) a_{lm}(\nu) Y_{lm}(\hat{\mathbf{n}}) \tag{3.13}$$

then this estimator also captures the foregrounds coupled to these systematics.

In this section, we outlined how one can easily obtain a fully empirical estimate of the foreground power by exploiting frequency-frequency cross-correlations. In order to critically evaluate the performance and subsequent utility of this foreground removal scheme, we will showcase this estimator in action by simulating mock observations of high-z LIM observations under various experimental setups. In the following section we outline the sky and instrument models we consider.

3.3 Simulation Setup

In this section, we outline the various components of the simulated observations. To demonstrate foreground removal using the B19 estimator, we consider observations of the 21 cm line during the Epoch of Reionization (EoR) with a range of instrument types. Although we focus on the 21 cm signal, our simulations include instruments commonly used in other LIM measurements, such as single-dish telescopes. As a result, we expect the explorations in the following section to be broadly applicable to the wider LIM community.

3.3.1 Sky Modeling

To generate the 21 cm brightness temperature signal, we use the publicly available code 21cmFAST (Mesinger et al., 2011). This is a semi-numerical code that employs both excursion set formalism and perturbation theory to produce simulations of various cosmological fields, including of course the 21 cm field, that agree with more physically motivated numerical simulations on large scales. Using the fiducial input parameters corresponding to a reionization scenario that starts at $z \sim 12$, and ends at $z \sim 5.5$, we simulate a large $200 \times 200 \times 244$ box with a 10 deg field of view, 3 arcminute pixel size, and 244 frequency channels with channel widths of 0.3 MHz. This box spans a redshift range of $z \sim 6 - 10$ corresponding to a total bandwidth of 73.8 MHz.

Next, we simulate galactic synchrotron emission as the broadband contaminant since it is the brightest of the diffuse contaminants for low-frequency 21 cm observations. We consider two foreground models: one idealized and one more realistic. For the idealized foreground model, each pixel is assigned a brightness temperature T according to the power law spectrum

$$T(\nu) = A_{\rm sync} \left(\frac{\nu}{\nu_*} \right)^{-\alpha(\mathbf{r})}, \tag{3.14}$$

where here, $A_{\rm sync} = 335.4$ K and $\nu_* = 150$ MHz, and α is the spectral index that varies from pixel to pixel (Wang et al., 2006; Liu and Tegmark, 2012). The large-scale spatial fluctuations in α are generated by drawing a Gaussian realization from a power spectrum $P_\alpha(k) \sim \mathcal{N}(0, 0.01)$. The final field $\alpha(\mathbf{r})$ is then obtained by adjusting the mean value of α field to 2.8 (Wang et al., 2006). While this model does not strictly satisfy the requirement of separability, since the frequency and spatial components are coupled by the spectral index, it strikes a good balance between idealization and realism, making it well-suited for demonstrating how the estimator works in the following section.

For the more realistic model, galactic synchrotron radiation is simulated using the PYGSM package (de Oliveira-Costa et al., 2008; Zheng et al., 2016). Since there do not exist all-sky maps of our galaxy at all frequencies from observation, this package interpolates over gaps in coverage using a set of principal components that are trained on 29 sky maps between 10 MHz and 5 THz.

3.3.2 Modeling Point Spread Functions

In order to study the effects of various types of instrumentation that make up the LIM experimental landscape, here, we simulate both single dish telescopes as well as interferometric arrays using simulation code developed for Chapter 2 (now the publicly available code LIMSTAT[2]). To begin, the

---

[2] https://github.com/McGill-Cosmic-Dawn-Group/LIMstat

single dish instrument is modeled by convolving the simulated sky with a 2D Gaussian beam with standard deviation equal to the diffraction limited angular resolution, that is, $\theta \sim \lambda_{\rm obs}/D$ where D is the diameter of the dish. Since the signal we have simulated is that of 21 cm during the EoR, we consider our single dish instrument to be similar to the Five-Hundred-Meter Aperture Spherical Telescope (FAST, Nan et al. 2011) and take $D = 500$ m. In addition to this frequency-dependent beam, we also simulate a frequency-independent beam to use as a benchmark. This is done by simply using the 2D Gaussian corresponding to the central frequency channel at all frequencies.

Next, to simulate the response an array of dishes, we compute a set of visibilities. Interferometers measure a primary beam- and fringe-weighted integral of the sky intensity known as the visibility, V_{ij}, defined as

$$V_{ij}(\nu) = \int A_{ij}(\hat{\mathbf{s}}, \nu) I(\hat{\mathbf{s}}, \nu) \exp\left(-2\pi i \frac{\mathbf{b}_{ij} \cdot \hat{\mathbf{s}}}{\lambda_{\rm obs}}\right) d\Omega \tag{3.15}$$

where $d\Omega$ is the differential solid angle, and $\mathbf{b}_{ij}$ is the baseline vector that characterizes the separation and orientation of the ith and jth receiving elements (such as dishes). The unit vector $\hat{\mathbf{s}}$ points to the direction of the incoming radiation on the sky. The observing wavelength is denoted by $\lambda_{\rm obs}$, $I(\hat{\mathbf{s}}, \nu)$ is the specific intensity, and $A_{ij}(\hat{\mathbf{s}}, \nu)$ is the geometric mean between the primary beams of the ith and jth elements. Equation (3.15) can be compared to a two-dimensional Fourier transform $\tilde{I}$ of the specific intensity, namely

$$\tilde{I}(\mathbf{u}, \nu) = \int I(\boldsymbol{\theta}, \nu) \exp\left(-i 2\pi \mathbf{u} \cdot \boldsymbol{\theta}\right) d^2\theta, \tag{3.16}$$

where we have invoked the flat-sky approximation to describe positions on the sky in terms of Cartesian angular coordinates $\boldsymbol{\theta} \equiv (\theta_x, \theta_y)$ and have defined a Fourier dual $\mathbf{u} \equiv (u, v)$ to this. One sees that in the limit that the sky is flat and the primary beam is reasonably uniform, each baseline of an interferometer measures a single Fourier mode in the plane of the sky with wavenumbers $u = b_x/\lambda_{\rm obs}$ and $v = b_y/\lambda_{\rm obs}$ on the uv plane, where b_x and b_y are the x and y components of $\mathbf{b}_{ij}$, respectively. These modes can then be easily mapped to the comoving wavevector perpendicular

Parameters	HERA-like
Beam FWHM at 150 MHz, θ_{FWHM}	8.7 deg
Element Diameter	14 m
Shortest Baseline	14.6 m
Longest Baseline (core)	292 m
Longest Baseline (outrigger)	876 m
Channel Width, $\delta\nu$	0.3 MHz

Table 3.1: Parameters for the HERA-like array assumed in this paper. These parameters are based on DeBoer et al. (2017) and Pober et al. (2014).

to the line of sight, $\mathbf{k}_\perp$, by

$$\mathbf{k}_\perp = \frac{2\pi \mathbf{b}_{ij}}{\lambda_{\text{obs}} D_c(z_{\text{emit}})} \tag{3.17}$$

where $D_c(z_{\text{emit}})$ is the comoving distance to the source emission. We also assume uniform sensitivity over the field of view.

To simulate the action of an interferometer, we start by denoting the sky at a particular frequency as $\mathbf{m}$ which stores all of the true pixel intensities on the sky. The measured visibilities then form a map $\mathbf{v}$ in the uv plane given by

$$\mathbf{v} = \mathbf{D}\mathcal{F}\mathbf{m} \tag{3.18}$$

where $\mathcal{F}$ denotes the 2D Fourier transform in the angular directions and $\mathbf{D}$ is a binary mask of the uv coverage (recording which uv modes are measured or missed based on the baselines present in an interferometer). Finally, with our simulated visibilities $\mathbf{v}$, we perform an inverse Fourier transform to obtain our observed map (i.e. the dirty map where the PSF is convolved with the sky). The instrument specifications are listed in Table 3.1.

3.3.3 Modeling Systematics

In order to evaluate the robustness of the proposed foreground mitigation scheme to systematics, we simulate cable reflections, a systematic that has been measured in various radio arrays (e.g. Barry et al. 2019; Kern et al. 2019; Kern et al. 2020; Aguirre et al. 2022; Murphy et al. 2023) and which is expected to remain a concern in future instrumentation (O'Hara et al., 2024). Ideally, signal collected by the antenna should travel through a connecting cable toward the amplifiers, digitizers, and correlators. However, due to an impedance mismatch between these components, a portion of the signal at the end of the connecting cable can reflect back toward the antenna, and then once again travel back down the cable. This reflected signal is then added to the data stream. This results in copies of the sky signal being present in the data at some time delay, τ. The length of this delay is simply twice the cable length divided by the speed of light in the cable. Typically the reflected signal has significantly lower amplitude than the original signal, but due to the overwhelming brightness of the foregrounds, this reflected sky signal can still swamp the original EoR signal. Importantly, cable reflections introduce additional chromatic structure across the observed band meaning this is precisely the type of systematic that the B19 foreground estimator should be able to handle and that existing methods relying on spectral smoothness are ill-equipped to deal with. Here, the sky signal with cable reflections is modeled by

$$T^{\text{obs}}(\hat{\mathbf{n}}, \nu) = T(\hat{\mathbf{n}}, \nu)\left(1 + \epsilon \cos(2\pi\nu\tau)\right) \tag{3.19}$$

where $\tau = 1000$ ns, consistent with what has been seen in instruments like the Murchison Widefield Array (MWA, Barry et al. 2019). To truly test whether these reflected signals can be captured in the foreground estimate, we take $\epsilon = 1$ and do not suppress the amplitude of the reflected signal. It should be noted that this is simultaneously a conservative test and one that needs refinement. It's conservative because the amplitude of the cable reflections have been increased significantly from what is seen in real data. On the other hand, multiple reflections with different delays are typically present which add complexity to this systematic.

3.4 Results

Using the simulations presented in the last section, we perform foreground removal using the B19 formalism outlined in Section 3.2, namely, by computing the estimated foreground power using Eq. 3.8 and subtracting it off the total measured power. All the results in this section have a target frequency sub-band of width 10 MHz centered at $z = 7.5$. The main metric used here to quantify the performance is the residual which is defined as the fractional difference between the true signal power and the recovered signal power and is given by

$$\%\mathrm{Residual} = 100 \times \frac{\mathrm{P}^{\mathrm{True}}_{\mathrm{signal}}(\mathrm{k}) - \mathrm{P}^{\mathrm{recov.}}_{\mathrm{signal}}(\mathrm{k})}{\mathrm{P}^{\mathrm{True}}_{\mathrm{signal}}(\mathrm{k})}. \tag{3.20}$$

We present these results in order of increasing realism, moving from idealized cases to more realistic scenarios.

3.4.1 Idealized Foregrounds

To begin with the simplest sky model, we consider a sky that has the 21cmFAST cosmological signal contaminated with galactic foregrounds modeled using Eq. 3.14. To place a benchmark on performance, we first compute the recovered signal power without any instrument response. These results are shown in Figure 3.1. The black line shows the absolute value of the total measured power. The reason for plotting the absolute value here is mostly aesthetic since the total power does have certain modes with negative power. This negative power is a result of the non-stationary assumption not being satisfied by the foreground contaminants. The red line is the true signal power and on large scales, the foreground power dominates over the signal by 4 orders of magnitude. The recovered signal shown in blue is in good agreement with the true signal power and the residuals shown in the bottom panel confirm that the signal is recovered within a few percent on all scales. This confirms that the foreground removal scheme we brought forth is, in principle, able to remove virtually all the foreground power from these mock data.

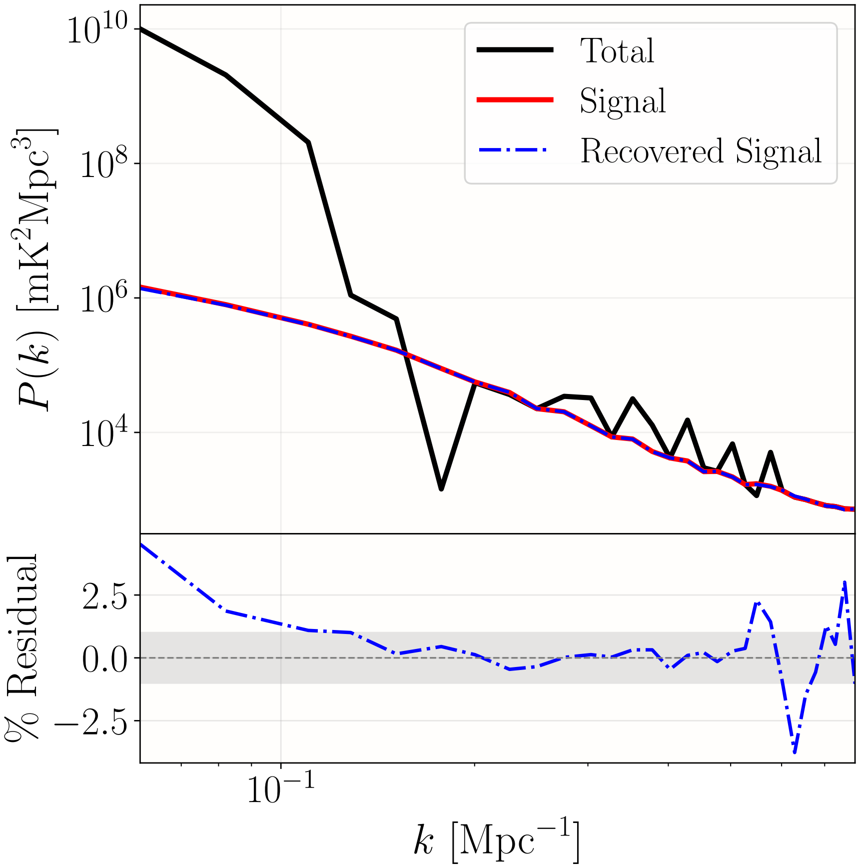

Figure 3.1: Top panel: Absolute value of the power spectrum for the total (signal + foregrounds) (black), the true underlying cosmological signal (red), and the recovered signal after foreground removal (blue dashed). Bottom panel: Fractional residual (in percent) between the true signal and the recovered signal. The grey band indicates the 1% threshold, serving as a benchmark for residual contamination. The foregrounds were modeled using the power law in Eq. 1.19.

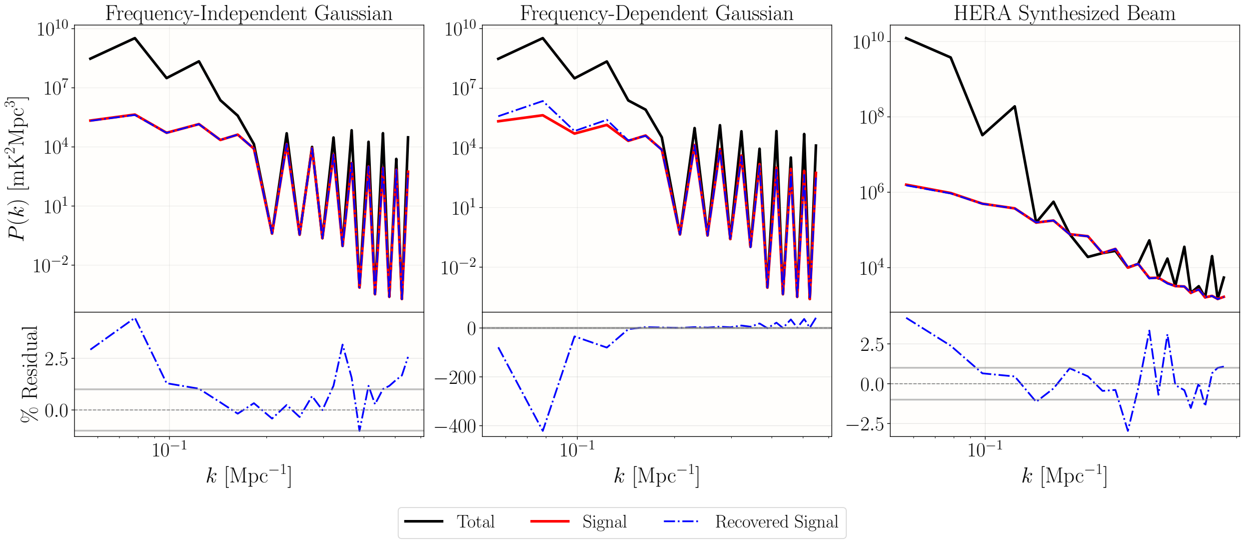

Figure 3.2: Same as Figure 3.1, but shown for different beam models: frequency-independent Gaussian beam (left), frequency-dependent beam (middle), and HERA-like array (right). The foreground model used in these cases follows the power-law prescription given in Eq. 1.19. In all cases, the foregrounds are almost entirely suppressed.

Moving now to increased levels of realism, we turn to Figure 3.2 to examine the effects of various PSFs. As expected the frequency-independent Gaussian beam has virtually no effect on the performance of the foreground mitigation and the residuals remain in a similar range to what they were without any beam at all. Similarly, the foreground mitigation is successful even for the HERA-like synthesized beam with complex uv sampling. In the case of the frequency-dependent Gaussian, however, the recovered spectrum on large scales, is visibly different from and is still roughly 1-2 orders of magnitude off from the true power. Nonetheless, the vast majority of the foreground power has been suppressed. It should be noted that the structure seen in the signal power as compared to Figure 3.1 is a result of the instrument response. While this mitigation scheme removes the foregrounds in the presence of a beam, it does not undo the effects of the beam itself.

To add one last element of realism, consider a HERA-like observation with a cable reflection with 1000 ns delay. These results are shown in Figure 3.3. When compared to the rightmost plot of Figure 3.2, there is clearly additional structure in both the total and the signal power spectrum when cable reflections are added. These systematics, however, are no challenge for the B19 estimator and for a HERA-like instrument, the residuals always remain within a few percent wether systematics are present or not. This demonstrates the robustness of this foreground removal method under realistic observing conditions. Although the foreground model used here is somewhat idealized, it is still not separable, making the achieved level of foreground removal a non-trivial benchmark to surpass. Having successfully cleared this initial test, we now proceed to an even more realistic model.

3.4.2 PYGSM Foregrounds

In Figure 3.4, we present the results obtained using a more realistic foreground model simulated with PYGSM. Compared to the previous section, foreground removal is notably less effective across all beam models. While some suppression is achieved, significant residuals remain. For the Gaussian beam models, these residuals are most pronounced on large scales, whereas for the

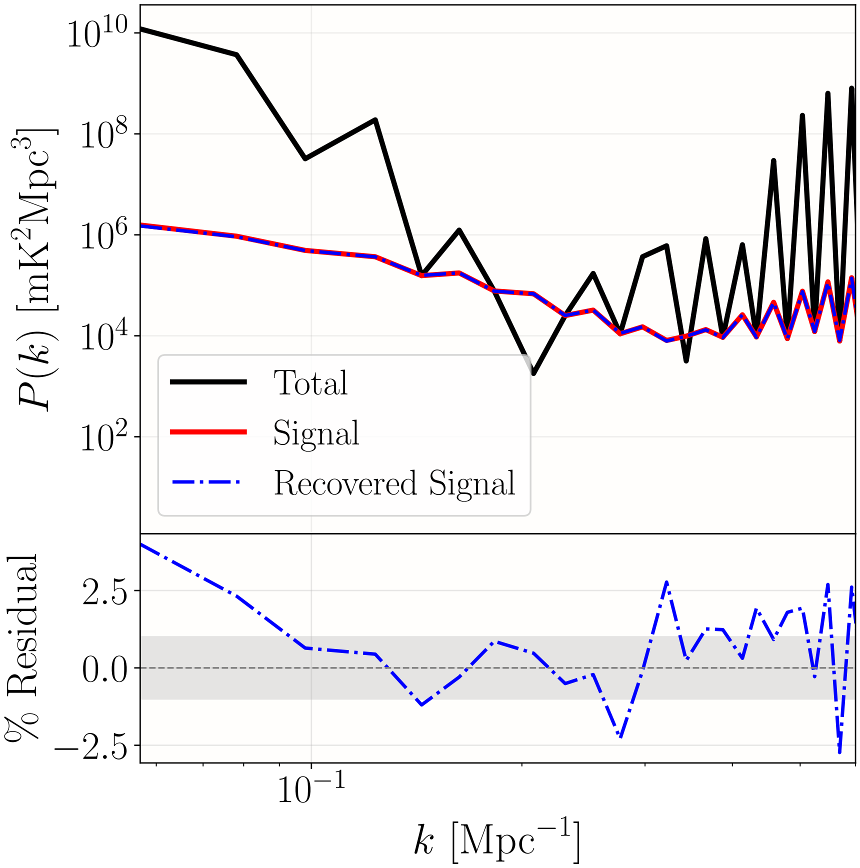

Figure 3.3: Same as Figure 3.1, but with the introduction of cable reflections as a systematic contaminant. This demonstrates the impact of instrumental systematics on the recovered signal and the effectiveness of the foreground removal method in their presence.

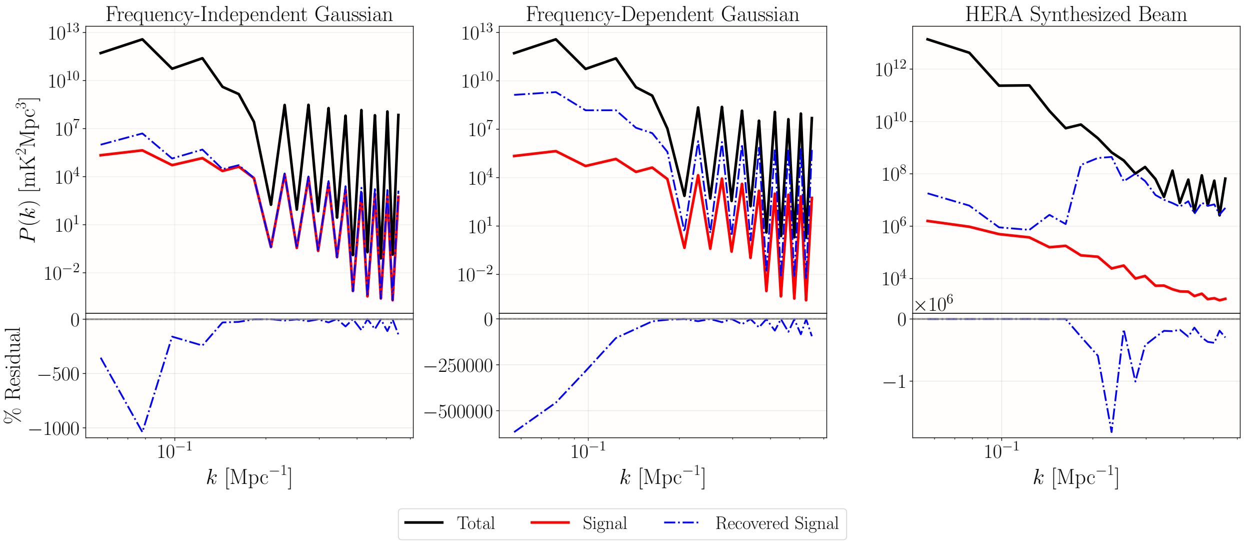

Figure 3.4: Same as Figure 3.2, but using the more realistic PYGSM foreground model instead of the power-law model from Eq. 1.19. The three panels correspond to the frequency-independent Gaussian beam (left), frequency-dependent beam (middle), and HERA-like array (right), illustrating the impact of different instrumental responses on foreground removal.

HERA-like PSF, they are higher on small scales. These effects arise from the complex interplay between the spatial and frequency correlations of the foregrounds and the instrument response. A further challenge is that, as an empirical model, PYGSM inherently includes a broad range of foreground contaminants. While it is primarily composed of Galactic synchrotron emission, it also encapsulates other sources, such as free-free emission and extragalactic contributions. This additional complexity disrupts the linear biasing model required for effective foreground removal. When compounded with instrumental effects, these factors lead to the large biases observed in the residuals.

3.5 Conclusion

In this work, we introduced a novel foreground mitigation scheme designed to handle the broadband contaminants of line intensity mapping (LIM) observations. Building on the B19 estimator presented in Beane et al. (2019), we demonstrated that a fully data-driven approach using intra-dataset frequency-frequency cross-correlations enables the estimation of a foreground-only power spectrum, which can then be subtracted from the total measured power to reveal the underlying cosmological signal.

This method offers several key advantages. (1) It is based on minimal assumptions, requiring only that foregrounds be separable in their frequency and spatial components and that their frequency correlation length be longer than that of the signal. (2) It remains effective in the presence of a chromatic (or achromatic) beam, making it robust across different experimental configurations. (3) By operating directly on the *measured* power spectrum, it removes foregrounds and their coupled systematics simultaneously, providing a systematics-hardened estimator.

To evaluate the performance of this method in practice, we simulated 21 cm observations across different experimental setups, including both single-dish instruments and radio interferometers, while incorporating realistic systematic contaminants, namely cable reflections. Importantly, this method does not attempt to undo the effects of systematics in the recovered signal spectrum; rather, it estimates all foreground components in the presence of these effects, allowing for a direct subtraction from the observed spectrum which inherently contains the same systematic effects. While the resultant recovered signal does contain these instrumental effects, systematics on the signal itself have been shown to be non-detrimental to a detection. For instance, in the case of 21 cm, where the expected signal power is on the order of $P(k) \sim 30$ mK, a 1% systematic contamination from cable reflections would only alter the measured power by ~ 0.1 mK. Thus, this technique effectively removes foregrounds and their systematics in one step, enabling a clearer path to detection.

Under idealized but physically motivated conditions, we achieved percent-level signal residuals, demonstrating that 4–5 orders of magnitude of foreground contamination can be removed in a single analysis step. However, when tested with a more realistic foreground model, while the technique still provided suppression, the interaction between the foregrounds and the instrument response led to imperfect subtraction. This highlights the need for further investigation into the types of contaminants and observational conditions under which this approach is most effective. Nonetheless, this work introduces an entirely new way to study foregrounds. Unlike traditional approaches that rely on explicit foreground modeling or decomposition based techniques, this is a fully data-driven method that learns the foreground components directly from the data itself.

LIM remains an observationally challenging endeavor, requiring a diverse set of tools to enable robust detections across different surveys at different wavelengths. This work contributes yet another powerful tool to the growing analysis toolbox, ensuring that we continue to refine our strategies for extracting cosmological signals from complex datasets. As the field advances, leveraging multiple techniques in tandem will be critical for fully realizing the potential of LIM as a powerful probe of cosmic structure and evolution.

Appendix A: Foreground Power Spectrum Estimation with the B19 Estimator

In this appendix, we will derive Eq. 3.8, showing that the foreground power spectrum estimator does indeed recover foreground cross-terms. To begin we assume a separable foreground model as described by Eq. 3.3 Section 3.1, that is,

$$T^f(\hat{\mathbf{n}}, \nu) = \beta(\nu)\delta_m(\hat{\mathbf{n}}). \tag{3.21}$$

If we consider a sky model with both a signal and a foreground component, $T^{\text{obs}} = T^s + T^f$, we can attempt to recover the foreground power, including foreground-signal cross-terms, using the B19 estimator from Beane et al. (2019) adapted for the MAPS

$$C_l(\nu,\nu') = \frac{C_l(\nu,\nu_j)C_l(\nu',\nu_k)}{C_l(\nu_j,\nu_k)}. \tag{3.22}$$

Decomposing T^{obs} into spherical harmonics, we obtain,

$$C_l(\nu,\nu') = \frac{\langle(a^s_{lm}(\nu)+a^f_{lm}(\nu))(a^{s*}_{lm}(\nu_j)+a^{f*}_{lm}(\nu_j))\rangle\langle(a^s_{lm}(\nu')+a^f_{lm}(\nu'))(a^{s*}_{lm}(\nu_k)+a^{f*}_{lm}(\nu_k))\rangle}{\langle(a^s_{lm}(\nu_j)+a^f_{lm}(\nu_j))(a^{s*}_{lm}(\nu_k)+a^{f*}_{lm}(\nu_k))\rangle}. \tag{3.23}$$

Since the denominator is at most quadratic in the foreground spherical harmonic coefficients, then any terms in the numerator that are quadratic or less will be order unity or less. Therefore, keeping only higher order foreground terms we obtain the estimated foreground power, $\hat{C}^f_l(\nu,\nu')$

$$\hat{C}^f_l(\nu,\nu') = C^f_l(\nu,\nu') + 2C^{sf}_l(\nu,\nu') + \frac{\beta^f}{\beta^f_j}C^{sf}_l(\nu_j,\nu) + \frac{\beta^f}{\beta^f_k}C^{sf}(\nu_k,\nu) \tag{3.24}$$

as required. The first two terms constitute the foreground auto- and cross-terms while the last two terms are a residual bias which can be mitigated by averaging over many cleaning channels ν_j, ν_k.

Bibliography

J. L. Bernal, P. C. Breysse, and E. D. Kovetz, **123**, 251301 (2019), 1907.10065.

H. Fronenberg, A. S. Maniyar, A. Liu, and A. R. Pullen, Phys. Rev. Lett. **132**, 241001 (2024), URL https://link.aps.org/doi/10.1103/PhysRevLett.132.241001.

K. W. Masui, E. R. Switzer, N. Banavar, K. Bandura, C. Blake, L. M. Calin, T. C. Chang, X. Chen, Y. C. Li, Y. W. Liao, et al., **763**, L20 (2013), 1208.0331.

S. H. Teng, S. Veilleux, and A. J. Baker, **765**, 95 (2013), 1301.5642.

J. Rhee, P. Lah, J. N. Chengalur, F. H. Briggs, and M. Colless, **460**, 2675 (2016), 1605.02006.

C. J. Anderson, N. J. Luciw, Y. C. Li, C. Y. Kuo, J. Yadav, K. W. Masui, T. C. Chang, X. Chen, N. Oppermann, Y. W. Liao, et al., **476**, 3382 (2018), 1710.00424.

A. Chowdhury, N. Kanekar, B. Das, K. S. Dwarakanath, and S. Sethi, **913**, L24 (2021), 2105.06773.

L. Wolz, A. Pourtsidou, K. W. Masui, T.-C. Chang, J. E. Bautista, E.-M. Müller, S. Avila, D. Bacon, W. J. Percival, S. Cunnington, et al., **510**, 3495 (2022), 2102.04946.

M. Amiri, K. Bandura, T. Chen, M. Deng, M. Dobbs, M. Fandino, S. Foreman, M. Halpern, A. S. Hill, G. Hinshaw, et al., The Astrophysical Journal **947**, 16 (2023), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/acb13f`.

M. Amiri, K. Bandura, A. Chakraborty, M. Dobbs, M. Fandino, S. Foreman, H. Gan, M. Halpern, A. S. Hill, G. Hinshaw, et al., **963**, 23 (2024), 2309.04404.

S. Paul, M. G. Santos, Z. Chen, and L. Wolz, *A first detection of neutral hydrogen intensity mapping on mpc scales at* $z \approx 0.32$ *and* $z \approx 0.44$ (2023), 2301.11943.

A. R. Pullen, T.-C. Chang, O. Doré, and A. Lidz, **768**, 15 (2013), 1211.1397.

A. R. Pullen, P. Serra, T.-C. Chang, O. Doré, and S. Ho, **478**, 1911 (2018), 1707.06172.

A. Roy, N. Battaglia, and A. R. Pullen, arXiv e-prints arXiv:2406.07861 (2024), 2406.07861.

N. O. Stutzer, J. G. S. Lunde, P. C. Breysse, D. T. Chung, K. A. Cleary, D. A. Dunne, H. K. Eriksen, H. T. Ihle, H. Padmanabhan, D. Tolgay, et al., arXiv e-prints arXiv:2406.07511 (2024), 2406.07511.

D. T. Chung, P. C. Breysse, K. A. Cleary, D. A. Dunne, J. G. S. Lunde, H. Padmanabhan, N. O. Stutzer, D. Tolgay, J. R. Bond, S. E. Church, et al., arXiv e-prints arXiv:2406.07512 (2024), 2406.07512.

G. K. Keating, D. P. Marrone, G. C. Bower, E. Leitch, J. E. Carlstrom, and D. R. DeBoer, **830**, 34 (2016), `1605.03971`.

H. T. Ihle, J. Borowska, K. A. Cleary, H. K. Eriksen, M. K. Foss, S. E. Harper, J. Kim, J. G. S. Lunde, L. Philip, M. Rasmussen, et al., The Astrophysical Journal **933**, 185 (2022), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/ac63c5`.

G. Paciga, J. G. Albert, K. Bandura, T.-C. Chang, Y. Gupta, C. Hirata, J. Odegova, U.-L. Pen, J. B. Peterson, J. Roy, et al., Monthly Notices of the Royal Astronomical Society **433**, 639–647 (2013), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stt753`.

B. K. Gehlot, F. G. Mertens, L. V. E. Koopmans, M. A. Brentjens, S. Zaroubi, B. Ciardi, A. Ghosh, M. Hatef, I. T. Iliev, V. Jelić, et al., Monthly Notices of the Royal Astronomical Society **488**, 4271–4287 (2019), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stz1937`.

A. H. Patil, S. Yatawatta, L. V. E. Koopmans, A. G. d. Bruyn, M. A. Brentjens, S. Zaroubi, K. M. B. Asad, M. Hatef, V. Jelić, M. Mevius, et al., The Astrophysical Journal **838**, 65 (2017), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/aa63e7`.

F. G. Mertens, M. Mevius, L. V. E. Koopmans, A. R. Offringa, G. Mellema, S. Zaroubi, M. A. Brentjens, H. Gan, B. K. Gehlot, V. N. Pandey, et al., **493**, 1662 (2020), `2002.07196`.

M. Kolopanis, D. C. Jacobs, C. Cheng, A. R. Parsons, S. A. Kohn, J. C. Pober, J. E. Aguirre, Z. S. Ali, G. Bernardi, R. F. Bradley, et al., **883**, 133 (2019), `1909.02085`.

C. Cheng, A. R. Parsons, M. Kolopanis, D. C. Jacobs, A. Liu, S. A. Kohn, J. E. Aguirre, J. C. Pober, Z. S. Ali, G. Bernardi, et al., **868**, 26 (2018), `1810.05175`.

H. Garsden, L. Greenhill, G. Bernardi, A. Fialkov, D. C. Price, D. Mitchell, J. Dowell, M. Spinelli, and F. K. Schinzel, Monthly Notices of the Royal Astronomical Society **506**, 5802–5817 (2021), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stab1671`.

A. Ewall-Wice, J. S. Dillon, J. N. Hewitt, A. Loeb, A. Mesinger, A. R. Neben, A. R. Offringa, M. Tegmark, N. Barry, A. P. Beardsley, et al., Monthly Notices of the Royal Astronomical Society **460**, 4320–4347 (2016), ISSN 1365-2966, URL http://dx.doi.org/10.1093/mnras/stw1022.

A. P. Beardsley, B. J. Hazelton, I. S. Sullivan, P. Carroll, N. Barry, M. Rahimi, B. Pindor, C. M. Trott, J. Line, D. C. Jacobs, et al., **833**, 102 (2016), 1608.06281.

N. Barry, M. Wilensky, C. M. Trott, B. Pindor, A. P. Beardsley, B. J. Hazelton, I. S. Sullivan, M. F. Morales, J. C. Pober, J. Line, et al., The Astrophysical Journal **884**, 1 (2019), ISSN 1538-4357, URL http://dx.doi.org/10.3847/1538-4357/ab40a8.

W. Li, J. C. Pober, N. Barry, B. J. Hazelton, M. F. Morales, C. M. Trott, A. Lanman, M. Wilensky, I. Sullivan, A. P. Beardsley, et al., The Astrophysical Journal **887**, 141 (2019), ISSN 1538-4357, URL http://dx.doi.org/10.3847/1538-4357/ab55e4.

C. M. Trott, C. H. Jordan, S. Midgley, N. Barry, B. Greig, B. Pindor, J. H. Cook, G. Sleap, S. J. Tingay, D. Ung, et al., **493**, 4711 (2020), 2002.02575.

Z. Abdurashidova, J. E. Aguirre, P. Alexander, Z. S. Ali, Y. Balfour, R. Barkana, A. P. Beardsley, G. Bernardi, T. S. Billings, J. D. Bowman, et al., **924**, 51 (2022), 2108.07282.

M. Rahimi, B. Pindor, J. L. B. Line, N. Barry, C. M. Trott, R. L. Webster, C. H. Jordan, M. Wilensky, S. Yoshiura, A. Beardsley, et al., Monthly Notices of the Royal Astronomical Society **508**, 5954–5971 (2021), ISSN 1365-2966, URL http://dx.doi.org/10.1093/mnras/stab2918.

D. R. DeBoer, A. R. Parsons, J. E. Aguirre, P. Alexander, Z. S. Ali, A. P. Beardsley, G. Bernardi, J. D. Bowman, R. F. Bradley, C. L. Carilli, et al., **129**, 045001 (2017), 1606.07473.

L. M. Berkhout, D. C. Jacobs, Z. Abdurashidova, T. Adams, J. E. Aguirre, P. Alexander, Z. S. Ali, R. Baartman, Y. Balfour, A. P. Beardsley, et al., **136**, 045002 (2024), 2401.04304.

HERA Collaboration, Z. Abdurashidova, T. Adams, J. E. Aguirre, P. Alexander, Z. S. Ali, R. Baartman, Y. Balfour, R. Barkana, A. P. Beardsley, et al., **945**, 124 (2023), 2210.04912.

H. Lazare, D. Sarkar, and E. D. Kovetz, Phys. Rev. D **109**, 043523 (2024), URL `https://link.aps.org/doi/10.1103/PhysRevD.109.043523`.

N. Petrovic and S. P. Oh, Monthly Notices of the Royal Astronomical Society **413**, 2103–2120 (2011), ISSN 0035-8711, URL `http://dx.doi.org/10.1111/j.1365-2966.2011.18276.x`.

B. Yue, A. Ferrara, A. Pallottini, S. Gallerani, and L. Vallini, Monthly Notices of the Royal Astronomical Society **450**, 3829 (2015), URL `https://doi.org/10.1093%2Fmnras%2Fstv933`.

D. Alonso, P. Bull, P. G. Ferreira, and M. G. Santos, Monthly Notices of the Royal Astronomical Society **447**, 400–416 (2014), ISSN 0035-8711, URL `http://dx.doi.org/10.1093/mnras/stu2474`.

E. Rath, R. Pascua, A. T. Josaitis, A. Ewall-Wice, N. Fagnoni, E. de Lera Acedo, Z. E. Martinot, Z. Abdurashidova, T. Adams, J. E. Aguirre, et al., *Investigating mutual coupling in the hydrogen epoch of reionization array and mitigating its effects on the 21-cm power spectrum* (2024), 2406.08549, URL `https://arxiv.org/abs/2406.08549`.

R. Pascua, Z. E. Martinot, A. Liu, J. E. Aguirre, N. S. Kern, J. S. Dillon, M. J. Wilensky, N. Fagnoni, E. de Lera Acedo, and D. DeBoer, arXiv e-prints arXiv:2410.01872 (2024), 2410.01872.

H. Wang, K. Masui, K. Bandura, A. Chakraborty, M. Dobbs, S. Foreman, L. Gray, M. Halpern, A. Joseph, J. MacEachern, et al., arXiv e-prints arXiv:2408.08949 (2024), 2408.08949.

H. Wang, J. Mena-Parra, T. Chen, and K. Masui, **106**, 043534 (2022), 2203.07184.

J. Ding, X. Wang, U.-L. Pen, and X.-D. Li, **274**, 44 (2024), 2408.06682.

A. Beane, F. Villaescusa-Navarro, and A. Lidz, The Astrophysical Journal **874**, 133 (2019), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/ab0a08`.

M. G. Santos, A. Cooray, and L. Knox, **625**, 575 (2005), `astro-ph/0408515`.

K. K. Datta, T. R. Choudhury, and S. Bharadwaj, **378**, 119 (2007), `astro-ph/0605546`.

R. Mondal, S. Bharadwaj, and K. K. Datta, **474**, 1390 (2018), `1706.09449`.

A. Mesinger, S. Furlanetto, and R. Cen, **411**, 955 (2011), `1003.3878`.

X. Wang, M. Tegmark, M. G. Santos, and L. Knox, **650**, 529 (2006), `astro-ph/0501081`.

A. Liu and M. Tegmark, Monthly Notices of the Royal Astronomical Society **419**, 3491–3504 (2012), ISSN 1365-2966.

A. de Oliveira-Costa, M. Tegmark, B. M. Gaensler, J. Jonas, T. L. Landecker, and P. Reich, Monthly Notices of the Royal Astronomical Society **388**, 247–260 (2008), ISSN 1365-2966, URL `http://dx.doi.org/10.1111/j.1365-2966.2008.13376.x`.

H. Zheng, M. Tegmark, J. S. Dillon, D. A. Kim, A. Liu, A. R. Neben, J. Jonas, P. Reich, and W. Reich, Monthly Notices of the Royal Astronomical Society **464**, 3486–3497 (2016), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stw2525`.

R. Nan, D. Li, C. Jin, Q. Wang, L. Zhu, W. Zhu, H. Zhang, Y. Yue, and L. Qian, International Journal of Modern Physics D **20**, 989 (2011), `1105.3794`.

D. R. DeBoer, A. R. Parsons, J. E. Aguirre, P. Alexander, Z. S. Ali, A. P. Beardsley, G. Bernardi, J. D. Bowman, R. F. Bradley, C. L. Carilli, et al., Publications of the Astronomical Society of the Pacific **129**, 045001 (2017), ISSN 1538-3873, URL `http://dx.doi.org/10.1088/1538-3873/129/974/045001`.

J. C. Pober, A. Liu, J. S. Dillon, J. E. Aguirre, J. D. Bowman, R. F. Bradley, C. L. Carilli, D. R. DeBoer, J. N. Hewitt, D. C. Jacobs, et al., The Astrophysical Journal **782**, 66 (2014), ISSN 1538-4357, URL `http://dx.doi.org/10.1088/0004-637X/782/2/66`.

N. Barry, A. P. Beardsley, R. Byrne, B. Hazelton, M. F. Morales, J. C. Pober, and I. Sullivan, **36**, e026 (2019), `1901.02980`.

N. S. Kern, A. R. Parsons, J. S. Dillon, A. E. Lanman, N. Fagnoni, and E. de Lera Acedo, The Astrophysical Journal **884**, 105 (2019), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/ab3e73`.

N. S. Kern, A. R. Parsons, J. S. Dillon, A. E. Lanman, A. Liu, P. Bull, A. Ewall-Wice, Z. Abdurashidova, J. E. Aguirre, P. Alexander, et al., **888**, 70 (2020), `1909.11733`.

J. E. Aguirre, S. G. Murray, R. Pascua, Z. E. Martinot, J. Burba, J. S. Dillon, D. C. Jacobs, N. S. Kern, P. Kittiwisit, M. Kolopanis, et al., The Astrophysical Journal **924**, 85 (2022), ISSN 1538-4357, URL `http://dx.doi.org/10.3847/1538-4357/ac32cd`.

G. G. Murphy, P. Bull, M. G. Santos, Z. Abdurashidova, T. Adams, J. E. Aguirre, P. Alexander, Z. S. Ali, R. Baartman, Y. Balfour, et al., *Bayesian estimation of cross-coupling and reflection systematics in 21cm array visibility data* (2023), `2312.03697`, URL `https://arxiv.org/abs/2312.03697`.

O. S. D. O'Hara, F. Dulwich, E. de Lera Acedo, J. Dhandha, T. Gessey-Jones, D. Anstey, and A. Fialkov, Monthly Notices of the Royal Astronomical Society **533**, 2876–2892 (2024), ISSN 1365-2966, URL `http://dx.doi.org/10.1093/mnras/stae1952`.

Intermezzo 3

Up until now, our focus has been on addressing the near-term challenges of LIM in order to enable the first high-z detections. In doing so, we have presented real solutions to outstanding problems in the field in Chapters 2 and 3. But of course, the ultimate goal is not just to measure spectra for the sake of measuring spectra; we seek to use these measurements to constrain cosmology. With this in mind, we now turn to a key question: once LIM becomes a mature observational field, akin to the CMB, what groundbreaking science will it enable?

In Chapters 4 and 5, we look toward the future, considering an era when LIM measurements have advanced to the point where we can extract LIM secondaries, analogous to the secondary anisotropies in the CMB. As line emission propagates along the line of sight, it acquires additional signals which can, in principle, be isolated. In this work, we demonstrate how LIM lensing and CMB lensing observables can be combined synergistically to remove the low-z lensing signal from the CMB. This process yields a high-z-only CMB lensing observable, providing an unbiased probe of the high-z matter density field.

We then explore how this new observable can be leveraged to test cosmology beyond ΛCDM. Notably, for the first time, we demonstrate an observable that exhibits lensing BAOs, that is, BAO features in a lensing measurement that are typically washed out due to the line-of-sight projection of lensing observables. We show how this novel standard ruler can be used to constrain cosmology. The observables presented in this work introduce entirely new tools for high-redshift studies, opening exciting opportunities for precision cosmology in the LIM era.

These chapters are based on:

- **Fronenberg, H.**, Maniyar, A.S., Pullen, A.R., Liu, A. (2024) *Constraining Cosmology With the CMB × LIM-Nulling Convergence*, Phys. Rev. D **109** 123518
- **Fronenberg, H.**, Maniyar, A.S., Liu, A., Pullen, A.R. (2024) *New Probe of the High-z BAO scale: BAO tomography With CMB × LIM-Nulling Convergence*, Phys. Rev. Lett. **132** 241001.

Chapter 4

Constraining Cosmology With the CMB × LIM-Nulling Convergence

Hannah Fronenberg[1,2], Abhishek S. Maniyar[3,4,5], Anthony R. Pullen[6,7], Adrian Liu[1,2]

[1]Department of Physics, McGill University, Montréal, QC, Canada

[2]Trottier Space Institute, Montréal, QC, Canada

[3]Department of Physics, Stanford University, Stanford, CA, USA

[4] SLAC National Accelerator Laboratory, Menlo Park, CA, USA

[5]Kavli Institute for Particle Astrophysics and Cosmology, Stanford, CA, USA

[6]Department of Physics, New York University, New York, NY, USA

[7]Center for Computational Astrophysics, Flatiron Institute, New York, NY, USA

Abstract

Lensing reconstruction maps from the cosmic microwave background (CMB) provide direct observations of the matter distribution of the universe without the use of a biased tracer. Such maps, however, constitute projected observables along the line of sight that are dominated by their

low-redshift contributions. To cleanly access high-redshift information, Maniyar et al. (2022) showed that a linear combination of lensing maps from both CMB and line intensity mapping (LIM) observations can exactly null the low-redshift contribution to CMB lensing convergence. In this paper we explore the scientific returns of this nulling technique. We show that LIM-nulling estimators can place constraints on standard ΛCDM plus neutrino mass parameters that are competitive with traditional CMB lensing. Additionally, we demonstrate that as a clean probe of the high-redshift universe, LIM-nulling can be used for model-independent tests of cosmology beyond ΛCDM and as a probe of the high-redshift matter power spectrum.

4.1 Introduction

In recent decades, there has been a sustained effort to make precision measurements of the large scale universe over a vast portion of its history. Line intensity mapping (LIM) is an emergent technique for studying large scale structure. Here, one observes the integrated intensity of a single spectral line emanating from galaxies and the intergalactic medium (IGM). By virtue of observing lines with known rest frequencies, line intensity mapping allows one to obtain precise redshift information. Mapping line emission over a large bandwidth can therefore yield unprecedentedly large maps of the universe in three dimensions, allowing us to observe cosmic evolution in action. A number of lines are being targeted by current and upcoming experiments including Lyman-α, H-α, the 21 cm line of neutral hydrogen (HI), the 3727 Å, 3729 Å lines of singly ionized oxygen ([OII]), the forbidden 88.4 μm and 51.8 μm transitions of doubly ionized oxygen ([OIII]), a host of rotational line transitions of carbon monoxide (CO), and the forbidden 158 μm line of ionized carbon ([CII]). Each line traces a biased matter density field as well as regions of the IGM and of the galaxy related to their specific emission or absorption mechanisms. This makes line intensity mapping a powerful probe of both cosmology and astrophysics.

Along with LIMs, gravitational lensing is a promising probe of the matter density field. Weak gravitational lensing of the cosmic microwave background (CMB) arises when CMB photons from the surface of last scattering get deflected by the gravitational potentials that they encounter on their journey to the observer. Using CMB temperature and polarization maps to reconstruct the lensing potential, ϕ, gives us direct observation of the total matter distribution of the universe, both baryonic and dark, without the use of a biased tracer (Seljak and Zaldarriaga, 1999). Measuring the power spectrum of the lensing potential, either in auto- or in cross-correlation with large scale structure surveys, has the ability to probe the growth of matter fluctuations, place limits on primordial non-Gaussianity, constrain the sum of the neutrino masses, and even test theories of modified gravity (Lewis and Challinor, 2007; Schmittfull and Seljak, 2018; Allison et al., 2015). This information, however, is collapsed onto a single observable, the convergence, and the high-redshift contribution to the convergence is dwarfed by that of the low-redshift universe ($z \lesssim 2$). Since the CMB lensing convergence contains information about how matter is distributed along the entire line of sight, it has the potential to help us trace out the matter distribution of the early universe.

There are several proposed techniques to disentangle the redshift integrated lensing signal and to extract information from particular redshift intervals. For instance, cross-correlating the CMB convergence field with another tracer, such as a galaxy survey or line intensity map, allows one to pick out common matter density correlations at their common redshift. This method, however, has its drawbacks. The redshifts available for study are limited to those of the non-lensing probe and by virtue of cross-correlating with a biased tracer, the resulting cross-correlation is likewise biased, losing out on the unbiased nature of the lensing convergence. One can make progress on the latter by considering not a correlation with the tracer itself but rather, for example, with the LIM lensing convergence. Just like the CMB, LIMs also experience weak lensing by large scale structure as the photons pass through the cosmos on their way to our instruments. These lines, however, are only lensed by a portion of large scale structure that lenses the CMB, namely the low redshift universe. Cross-correlating LIM lensing and CMB lensing allows one to study the common low redshift matter density field that lenses both the LIM and the CMB. While the resulting correlation is unbiased it again is limited to the redshifts between the observer and the source plane of the LIM.

In order to access the high redshift information, Maniyar et al. (2022) proposes using the lensing information of two LIMs, to not just suppress, but exactly null out the low redshift contribution to the CMB convergence. This "nulling" method has been explored in the context of galaxy lensing (Huterer and White, 2005; Bernardeau et al., 2014; Barthelemy et al., 2020) as well as CMB lensing (McCarthy et al., 2021; Qu et al., 2023; Zhang et al., 2023). For instance, McCarthy et al. (2021) show that one can "null" out the imprint of uncertain baryonic effects from CMB lensing maps using cosmic shear surveys at $z < 1$. Similarly, Qu et al. (2023) showed that one can also use cosmic shear surveys to subtract off the imprint of uncertain dark energy physics from CMB lensing maps. Zhang et al. (2023) explores the potential of subtracting the imprint of gravitational nonlinearity at low redshift to help measure primordial bispectra. While never implemented with real data, these nulling techniques could be an important new tool for studying the high redshift universe.

What is more, Maniyar et al. (2022) show that the CMB × LIM-nulling convergence spectrum, $\langle \hat{\kappa}\hat{\kappa}_{\rm null}\rangle$, does not contain so-called line interloper bias when the LIM convergence maps are estimated with "LIM-pair" estimators of Section 4.2.2. Line interlopers are one of the chief systematic contaminants in LIMs, and consist of low-redshift spectral lines that redshift into the same observed frequency channel as the high redshift target line. In addition to estimators such as the LIM-pair estimators that mitigate interloper bias by construction, other strategies such as line identification, analysis of redshift space distortions, spectral deconfusion, and cross-correlations have all been shown to help reduce line interloper contamination (Silva et al., 2021; Sun et al., 2018; Visbal et al., 2011; Breysse et al., 2015; Gong et al., 2014; Liu et al., 2016). Provided that some combination of these strategies is able to bring line interlopers (and other potential systematics) under control, the CMB × LIM-nulling convergence spectrum has the potential to reveal exclusive information about the early universe. Exactly what information is revealed is the subject of this paper.

In this work, we explore the parameter space of LIM-nulling measurements. In section 4.2 we derive, for the first time, the CMB × LIM-nulling variance as well as discuss the potential use for this probe in constraining cosmology. In a companion paper, Fronenberg et al. (2024), we forecast using the CMB × LIM-nulling convergence spectrum to detect baryon acoustic oscillations in the early universe, which may serve as a standard ruler over vast portion of cosmic history. In Sections 4.3 and 4.4, calculate the signal-to-noise ratio of this cross-spectrum statistic as a function of various observing parameters. After exploring this vast parameter space, we converge on three possible observing scenarios some of which would allow one to constrain cosmology at high redshift in Section 4.4. Using these scenarios, we present several forecasts. Section 4.5 consists of a series of forecasts on ΛCDM+M_ν cosmology. First is a Fisher forecast to test the sensitivity of the CMB × LIM-nulling to the concordance model of cosmology and to compare it with traditional CMB lensing forecasts. In this section, we also explore how this LIM-nulling estimator behaves, in comparison to the regular CMB lensing convergence, in universes with time-evolving cosmologies. In addition, we forecast the sensitivity of this probe to the matter power spectrum at high-z in Section 4.6. Unless otherwise explicitly stated, our fiducial cosmology is that of *Planck 2015*.

4.2 LIM Lensing and LIM Nulling

In this section, we outline the key lensing estimators and observables used throughout the text which follow Maniyar et al. (2022). In the first subsection, we provide a brief overview of weak lensing by large scale structure in the context of the CMB and extend the discussion to include LIMs. In Section 4.2.2, we quickly review LIM lensing estimator and LIM-pair estimators. Finally, in section 4.2.3, we build upon the existing LIM-nulling estimator formalism to provide a derivation of the CMB × LIM-nulling variance.

4.2.1 Weak Lensing By Large-Scale Structure

The CMB acts as a source image which is lensed by the intervening matter density field. The deflection angle, α, is proportional to the gradient of the lensing potential, ϕ, which is the total gravitational potential of the projected mass distribution along the line of sight. This gradient of the potential is related to the convergence, $\kappa = -1/2\nabla\phi$, which is the line-of-sight-integrated matter density field, given by

$$\kappa(\hat{\mathbf{n}}) = \int_0^{z_\mathrm{s}} W(z', z_\mathrm{s})\delta_m(\chi(z')\hat{\mathbf{n}}, z')\frac{c\,dz'}{H(z')} \tag{4.1}$$

where z_s is the redshift of the source, $W(z, z_s)$ is the lensing kernel, $H(z)$ is the Hubble parameter, c is the speed of light, χ denotes the comoving distance, and $\delta_m(\mathbf{r}, z)$ is the matter density field at position $\mathbf{r}$ and redshift z. The lensing kernel for a source at a single comoving slice is given by

$$W(z, z_s) = \frac{3}{2}\left(\frac{H_0}{c}\right)^2 \frac{\Omega_{m,0}}{a}\chi(z)\left(1 - \frac{\chi(z)}{\chi(z_s)}\right) \tag{4.2}$$

where H_0 is the Hubble constant, $\Omega_{m,0}$ is the matter fraction today, a is the scale factor, and $\chi(z_s)$ is the comoving distance to the source.

The deflections induced on CMB photons are small, on the order of arcminutes. However, the structures responsible for the deflection are large, on the order of degrees. Therefore, somewhat counter-intuitively, to study the large scale structure of the universe one actually has to study the small scale anisotropies of the CMB. Lensing induces correlations between the otherwise uncorrelated CMB spherical harmonic coefficients, a_{lm}. With the use of quadratic estimators like those derived in Refs. Hu and Okamoto (2002) and Maniyar et al. (2021), an estimate of κ can be obtained which we denote by $\hat{\kappa}$.

One can then compute the angular power spectrum of the convergence which is given by

$$C_L^{\hat{\kappa}\hat{\kappa}} = \int_0^{z_s} \frac{W(z', z_\mathrm{s})^2}{\chi(z')^2} P_\mathrm{m}\left(k = \frac{L + 1/2}{\chi(z')}, z'\right)\frac{c\,dz'}{H(z')} \tag{4.3}$$

where P_m is the matter power spectrum. In this expression, we assume the Limber approximation (Lemos et al., 2017).

Obtaining lensing measurements are challenging yet have seen tremendous progress in recent years. To date, a number of lensing detections have been made, the first of which was by the Wilkinson Microwave Anisotropy Probe (WMAP) in 2007 using the Hu and Okamoto estimator on temperature maps and cross-correlating the resultant κ map with radio galaxy counts (Smith et al., 2007). Subsequent measurements have been made of the lensing signal in temperature as well as polarization maps by ACT (Das et al., 2011; Sherwin et al., 2017; Shaikh et al., 2023), the South Pole Telescope (SPT, van Engelen et al. 2012; Hanson et al. 2013; Omori et al. 2017; Story et al. 2015; Wu et al. 2019; Pan et al. 2023), *Planck* (Planck Collaboration et al., 2016, 2020), Background Imaging of Cosmic Extragalactic Polarization (BICEP, BICEP2 Collaboration et al. 2016), and the POLARization of the Background Radiation experiment (POLARBEAR, Ade et al. 2014; Adachi et al. 2020). Excitingly, current and next generation wide-field CMB experiments like SPT-3G, SPT-3G+, AdvACT, the Simons Observatory (SO) and CMB-Stage 4 (CMBS4), will provide high signal-to-noise lensing measurement with unprecedented angular resolution (Ward and Advanced ACT Collaboration, 2017; Benson et al., 2014; Anderson et al., 2022; Ade et al., 2019; Abazajian et al., 2022). These upcoming detections will enable further analyses such as the LIM-nulling measurement we propose in Section 4.2.3 and forecast in Section 4.5.

4.2.2 LIM Lensing Estimators

Just like the CMB, lower redshift LIMs also incur correlations between Fourier coefficients as a result of lensing. In the same spirit as CMB lensing reconstruction, the LIM lensing convergence can be estimated with LIM lensing estimators which are extensions of those developed for the CMB. LIMs however, suffer from significant foreground bias, be it from diffuse extended sources, or from line interlopers. This has been shown to cause significant foreground bias to the LIM lenisng convergence Maniyar et al. (2022). Luckily, Maniyar et al. (2022) showed that by using a LIM-pair estimator, one can perform LIM lensing reconstruction free of interloper bias (to first order). This LIM-pair estimator makes use of the fact that two LIMs from the same redshift slice and the same patch of the sky will contain the same correlations due to lensing. However, since each line is observed at a different frequency, they will suffer from different sources of foreground contamination which will be uncorrelated. The LIM-pair lensing estimator, using LIMs X and Y, is given by

$$\hat{\kappa}_{XY}(\boldsymbol{L}) = \int \frac{d^2 l_1}{(2\pi)^2} \frac{d^2 l_2}{(2\pi)^2} \, F_{XY}(\boldsymbol{l}_1, \boldsymbol{l}_2) \, X_{l_1} Y_{l_2} \,, \tag{4.4}$$

where $\boldsymbol{l}$ are the two dimensional Fourier wavenumbers for the LIM in the flat sky approximation, and $\boldsymbol{L} = \boldsymbol{l}_1 + \boldsymbol{l}_2$ are the wavenumbers for the lensing potential. The quantities $X(\boldsymbol{l})$ and $Y(\boldsymbol{l})$ are the observed LIM fields in Fourier space. The function F_{XY} is uniquely determined to ensure that $\hat{\kappa}_{XY}$ is unbiased to first order and to ensure that $\hat{\kappa}_{XY}$ is the minimum variance estimate of κ_{XY}. This solution to F_{XY} is given by

$$F_{XY}(\boldsymbol{l}_1, \boldsymbol{l}_2) = \lambda_{XY}(L) \frac{C_{l_1}^{YY} C_{l_2}^{XX} f_{XY}(\boldsymbol{l}_1, \boldsymbol{l}_2) - C_{l_1}^{XY} C_{l_2}^{XY} f_{XY}(\boldsymbol{l}_2, \boldsymbol{l}_1)}{C_{l_1}^{XX} C_{l_2}^{YY} C_{l_1}^{YY} C_{l_2}^{XX} - \left(C_{l_1}^{XY} C_{l_2}^{XY}\right)^2} \tag{4.5}$$

where C_l^{XX} and C_l^{YY} are the total auto-spectra for LIMs X and Y including noise, while C_l^{XY} is their cross-spectrum. The Lagrange multiplier $\lambda_{XY}(L)$ is given by

$$\lambda_{XY}(L) \equiv \left[\int_{\substack{\boldsymbol{l}_1+\boldsymbol{l}_2 \\ =\boldsymbol{L}}} f_{XY}(\boldsymbol{l}_1, \boldsymbol{l}_2) \frac{C_{l_1}^{YY} C_{l_2}^{XX} f_{XY}(\boldsymbol{l}_1, \boldsymbol{l}_2) - C_{l_1}^{XY} C_{l_2}^{XY} f_{XY}(\boldsymbol{l}_2, \boldsymbol{l}_1)}{C_{l_1}^{XX} C_{l_2}^{YY} C_{l_1}^{YY} C_{l_2}^{XX} - \left(C_{l_1}^{XY} C_{l_2}^{XY}\right)^2} \right]^{-1} . \tag{4.6}$$

where for brevity, we introduce the notation

$$\int_{\substack{\boldsymbol{l}_1+\boldsymbol{l}_2 \\ =\boldsymbol{L}}} \ldots \equiv \iint \frac{d^2 l_1 d^2 l_2}{(2\pi)^2} \delta(\boldsymbol{l}_1 + \boldsymbol{l}_2 - \boldsymbol{L}) \ldots . \tag{4.7}$$

The factor $f_{XY}(\boldsymbol{l}, \boldsymbol{l}')$ is the coupling coefficient

$$f_{XY}(\boldsymbol{l}, \boldsymbol{l}') = -\frac{2}{L^2} \left[\widetilde{C}_{l_1}^{XY} (\boldsymbol{L} \cdot \boldsymbol{l}_1) + \widetilde{C}_{l_2}^{XY} (\boldsymbol{L} \cdot \boldsymbol{l}_2) \right] \tag{4.8}$$

where $\widetilde{C}_l^{XY}$ is the unlensed cross-power spectrum. The reconstruction noise, $N_{XY}(L)$, of the LIM-pair estimator is given by

$$N_{XY}(L) = \int_{\substack{\boldsymbol{l}_1+\boldsymbol{l}_2 \\ =\boldsymbol{L}}} F_{XY}(\boldsymbol{l}_1,\boldsymbol{l}_2)\Big(F_{XY}(\boldsymbol{l}_1,\boldsymbol{l}_2)C_{l_1}^{XX}C_{l_2}^{YY} + F_{XY}(\boldsymbol{l}_2,\boldsymbol{l}_1)C_{l_1}^{XY}C_{l_2}^{XY}\Big). \tag{4.9}$$

For the interested reader, a detailed derivation of this estimator can be found in Appendix B of Maniyar et al. (2022).

4.2.3 CMB × LIM-nulling

In the last two subsections, we outlined how one could make use of both the CMB and of LIMs to extract information about the intervening matter density field. In the case of the CMB the resulting field is the matter density field over cosmic history since the surface of last scattering, projected onto a single plane. It is important to note that LIMs are lensed by the same low-z gravitational potentials that lens the CMB and therefore these probes share common low-redshift induced correlations. This can be exploited in order to make use of the LIM lensing information to "clean" the CMB convergence of its low redshift contribution.

From Eq. (4.1), it should be clear that it is possible to construct some kernel that vanishes over the low redshift interval $[0, z_{\rm null}]$. Since W is quadratic in χ, a linear combination of three such kernels suffices to find a non-trivial null solution for the coefficients of this polynomial. As shown in Maniyar et al. (2022), using two convergence maps each estimated from two LIMs at redshifts z_1 and z_2 ($z_1 < z_2$), and one CMB convergence map sourced at the surface of last scattering, $z_{\rm CMB}$, the LIM-nulling kernel is given by

$$W_{\rm null} = W(z, z_{\rm CMB}) + \alpha W(z, z_2) - (1+\alpha)W(z, z_1) \tag{4.10}$$

where

$$\alpha \equiv \frac{1/\chi(z_{\rm CMB}) - 1/\chi(z_1)}{1/\chi(z_1) - 1/\chi(z_2)}. \tag{4.11}$$

In Fig. 4.1, the LIM lensing kernels, the CMB lensing kernel, and the LIM-nulling kernel are plotted. The LIM-nulling kernel is exactly null between $0 < z < 4.5$, meaning that when integrating over the whole redshift range in Eq. (4.1), $\kappa_{\text{null}} = \kappa_{\text{CMB}} + \alpha\kappa_{z_2} - (1+\alpha)\kappa_{z_1}$ provides a map of the line of sight integrated matter density field between $4.5 < z < 1100$, providing a pristine view of early times.

LIM-nulling can be thought of as a type of foreground cleaning where the LIM-lensing information is used to clean the low redshift contribution to the CMB lensing. The data product that results from LIM-nulling does not itself contain any LIM information, whether from the original map or from its lensing reconstruction. Of course, LIMs cannot partake in tracing out the matter density field at a time before the line emission was emitted. It is composed of CMB lensing information from $z > z_{\text{null}}$.

This high-redshift information can be captured statistically by computing the CMB × LIM-nulling convergence spectrum,

$$C_L^{\hat{\kappa}\hat{\kappa}_{\text{null}}} = \int_0^{z_s} \frac{W(z', z_{\text{s}})W_{\text{null}}(z', z_1, z_2, z_{\text{CMB}})}{\chi(z')^2} P_{\text{m}}\left(k = \frac{L+1/2}{\chi(z')}, z'\right) \frac{c\,dz'}{H(z')}. \tag{4.12}$$

The motivation for always computing the cross-correlation of the LIM-nulling convergence with the CMB convergence and not simply in auto-correlation, is that cross-spectrum is free of all interloper bias, as shown in Maniyar et al. (2022). In addition, the variance of the cross-spectrum contains fewer cross terms than the variance of the LIM-nulling auto-spectrum. This cross-spectrum, however, effectively contains the same cosmological information as the LIM-nulling auto-spectrum when Δz is small. One therefore gains in signal-to-noise.

In Fig. 4.2 both the CMB convergence spectrum and the CMB × LIM-nulling convergence spectrum are plotted in black and red solid lines respectively. Upon first glance it is immediately evident that the CMB convergence has an order of magnitude more power than the CMB × LIM-nulling convergence spectrum. This is expected since a significant portion of the power has been nulled in the CMB × LIM-nulling spectrum. Perhaps more subtle is the re-emergence of acoustic peaks in the nulling spectrum. In order to help elucidate this feature, we also show the CMB × LIM-nulling spectrum using the no-wiggle Eisentein and Hu fitting function in place of the typical matter power spectrum. This fitting function is essentially the matter power spectrum without the baryon acoustic oscillations (BAO). One can see the solid red line oscillating about the no-wiggle spectrum which is plotted in dashed black. We have refrained from forecasting the potential to constrain the BAO scale in this work as we believe it merits a dedicated discussion. In our companion paper, Fronenberg et al. (2024), we perform an Alcock-Paczynski test on mock CMB × LIM-nulling data sets in order to forecast whether it is possible to measure BAO features with such a probe.

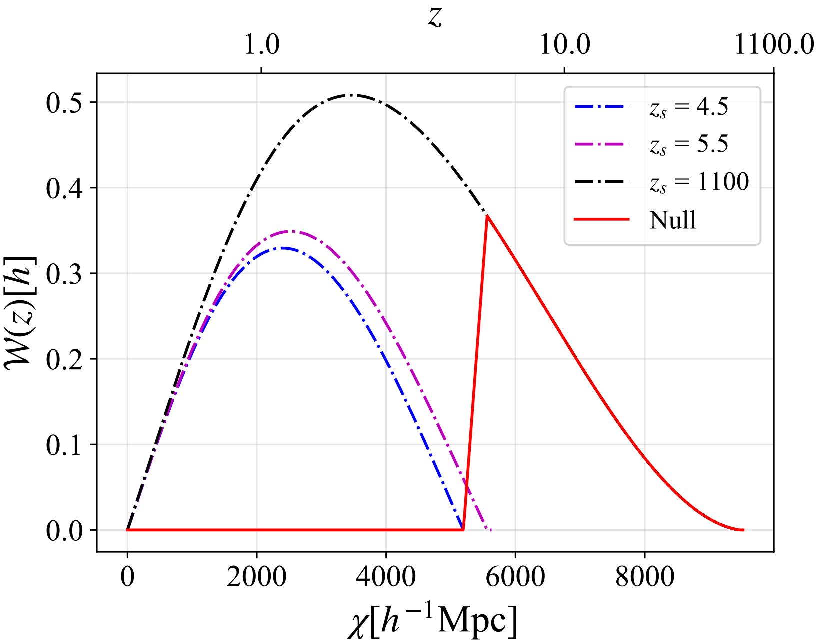

Figure 4.1: The rescaled nulling kernel, defined as $\mathcal{W}(z, z_s) = W(z, z_s)\frac{c}{H(z)}$ for three sources. The CMB kernel is shown in black and the two LIM kernels at redshifts 4.5 and 5.5 are shown in blue and magenta, respectively. Finally, the LIM-nullling kernel is shown here in red and is null from $0 < z < 4.5$ which corresponds to a complete insensitivity to the matter density field over that redshift range.

In the next two sections, we explore how the choice of various observing parameters affects the signal-to-noise ratio (SNR) of this probe. In order to do so, we require the variance of the CMB × LIM-Nulling cross-spectrum. It is given by

$$\begin{aligned}\mathrm{var}\left(C_L^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_{\rm null}}\right) = \frac{1}{f_{\rm sky}(2L+1)}\Bigg\{&2\Big[(C_L^{\hat{\kappa}_{\rm CMB}}+N_L^{\hat{\kappa}_{\rm CMB}})^2 + 2\alpha(C_L^{\hat{\kappa}_{\rm CMB}}+N_L^{\hat{\kappa}_{\rm CMB}})C_L^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_2} \\ &- 2(1+\alpha)(C_L^{\hat{\kappa}_{\rm CMB}}+N_L^{\hat{\kappa}_{\rm CMB}})C_L^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_1}\Big] + \alpha^2\Big[(C_L^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_2})^2 + (C_L^{\hat{\kappa}_{\rm CMB}}+N_L^{\hat{\kappa}_{\rm CMB}})(C_L^{\hat{\kappa}_2}+N_L^{\hat{\kappa}_2})\Big] \\ &- 2\alpha(1+\alpha)\Big[C_L^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_1}C_L^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_2} + (C_L^{\hat{\kappa}_{\rm CMB}}+N_L^{\hat{\kappa}_{\rm CMB}})C_L^{\hat{\kappa}_1\hat{\kappa}_2}\Big] \\ &+ (1+\alpha)^2\Big[(C_L^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_1})^2 + (C_L^{\hat{\kappa}_{\rm CMB}}+N_L^{\hat{\kappa}_{\rm CMB}})(C_L^{\hat{\kappa}_1}+N_L^{\hat{\kappa}_1})\Big]\Bigg\}. \qquad (4.13)\end{aligned}$$

where we use “var” to denote the variance. The quantity $N_L^{\hat{\kappa}_i}$ denote the lensing reconstruction noise corresponding to the estimated convergence $\hat{\kappa}_i$. To be explicit, the CMB reconstruction noise is denoted by $N_L^{\hat{\kappa}_{CMB}}$ and $N_L^{\hat{\kappa}_1}$ is shorthand for the LIM-pair lensing reconstruction noise, $N_{XY}(L)$, using lines X and Y at z_1. Finally, $f_{\rm sky}$ is the fraction of the sky area observed. A complete derivation can be found in Appendix 4.7.

4.3 Cosmic Variance Limited SNR

Before considering the effects of instrument noise on the CMB × LIM-nulling SNR, we first wish to illustrate the importance of the choice of line redshifts and of the redshift separation between LIMs used for nulling. To do this, we work in harmonic space and compute the LIM-pair reconstruction noise at different redshifts as well as the CMB lensing reconstruction noise, both of which enter into Eq. (4.13). For various cases, we compute both the SNR per mode, as well as the cumulative SNR assuming uncorrelated errors which is given by

$$\text{Cumulative SNR} = \left[\sum_L \left(\frac{C_L}{\sqrt{\mathrm{var}(C_L)}}\right)^2\right]^{1/2}. \qquad (4.14)$$

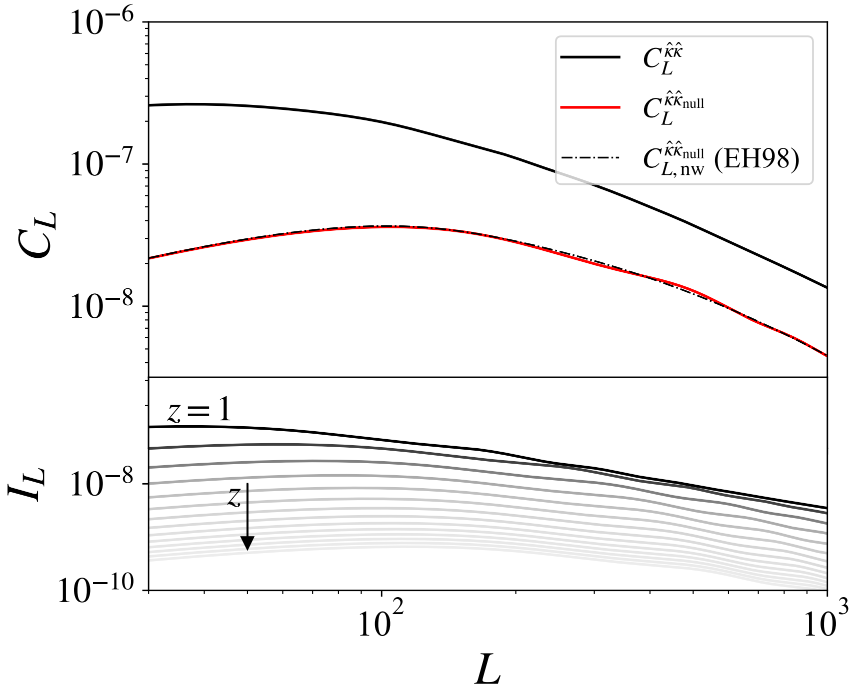

Figure 4.2: Top: the CMB convergence spectrum, $C_L^{\hat{\kappa}\hat{\kappa}}$ in black and the CMB × LIM-nulling convergence spectrum, $C_L^{\hat{\kappa}\hat{\kappa}_{\rm null}}$, in red. The dot-dashed black curve shows the CMB × LIM-nulling convergence spectrum computed with the no-wiggle Eisenstein & Hu fitting function. Bottom: the integrand of Eq. (5.3) evaluated at increasing redshifts from top to bottom starting at z = 1. Since the BAO scale is a fixed comoving scale, its angular projection changes as a function of z. The BAO features evolve gradually to lower L as z decreases which, when integrated over redshift, result in the washing out of BAO wiggles in $C_L^{\hat{\kappa}\hat{\kappa}}$. It is for this reason that the CMB convergence spectrum $C_L^{\hat{\kappa}\hat{\kappa}}$ is smooth with no discernible BAO features in the top panel. In contrast, the LIM-nulled convergence $C_L^{\hat{\kappa}\hat{\kappa}_{\rm null}}$ sees the reemergence of acoustic peaks (especially apparent when viewed against the reference no-wiggle nulled spectrum). These acoustic features are the result of the much slower angular evolution of BAO wiggles at early times.

We compute C_L using the matter power spectrum from the publicly available code `CAMB`.[1] The line auto-spectra and cross-spectra which enter into the LIM-pair lensing reconstruction noise through Eqs. (4.5) and (4.6), are obtained from the publicly available code `HaloGen`.[2] As the name suggests, `HaloGen` uses a halo model formalism based on conditional luminosity functions (Schaan and White, 2021a,b). For more information on how we model these lines here, we refer the interested reader to Appendix A of Maniyar et al. (2022). We take the spherically averaged power spectra of each line and convert them into the corresponding angular spectra. Working in the thin shell approximation, this conversion is written as,

$$C_l = \mathcal{V}^{-1} P(k = l/\chi(z), z) \tag{4.15}$$

where $\mathcal{V} = \chi^2(z)\Delta\chi$ is the comoving volume per steradian of a shell centered at z. We set the width of the shell to the comoving distance corresponding to a redshift width of 0.05.

---

[1] https://github.com/cmbant/CAMB

[2] https://github.com/EmmanuelSchaan/HaloGen/tree/LIM

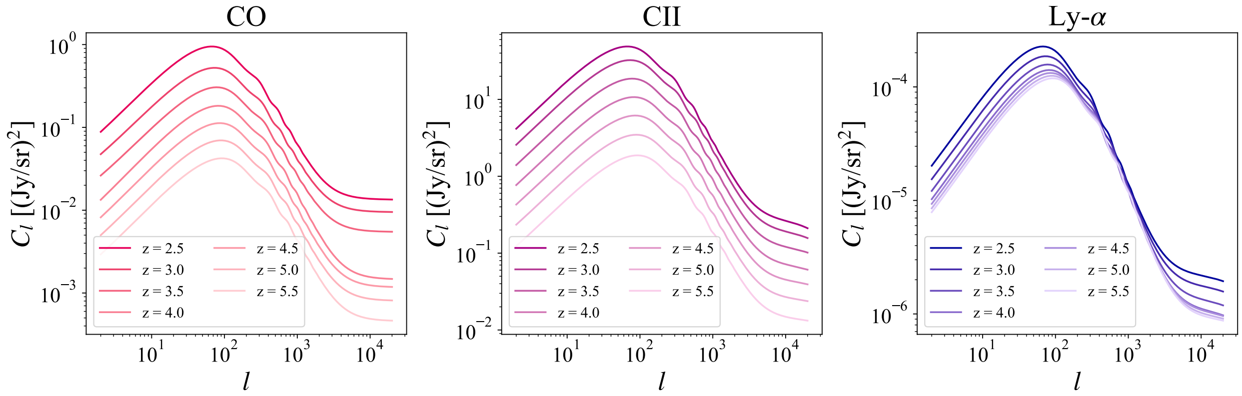

Figure 4.3: Angular power spectra of CO 4-3 (left, red), [CII] (middle, purple), and Ly-α (right, blue) as functions of angular multipole l. We plot these spectra at redshifts ranging from $z = 2.5$ to 5.5. Darker line colour denotes low-z while paler line colour denoted high-z. To produce these spectra, spherically averaged line power spectra are obtained from `HaloGen` and then converted to angular power spectra using Eq. (4.15). Although the achromaticity of lensing (to first order) means that the amplitudes of the line power spectra cancel out in the noiseless cosmic variance-limited regime considered in Section 4.3, this changes in the realistic noisy scenarios that we examine in Section 4.4.

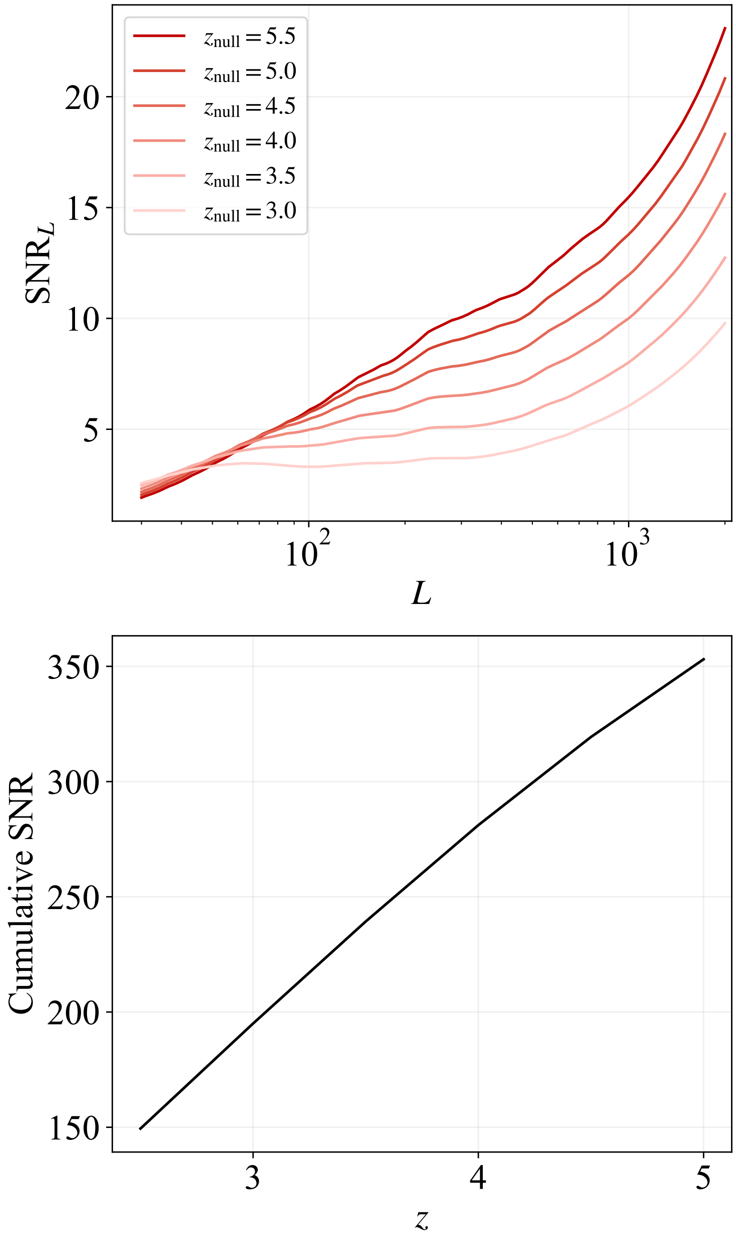

Figure 4.4: Top: signal-to-noise ratio of the nulling estimator per lensing multipole L. Each line denotes a different nulling redshift going from high redshift in dark red to low redshift in pale red. Bottom: The cumulative SNR as a function of nulling redshift. In all cases, the line separation is fixed at $\Delta z = 0.5$. Even aside from the science applications at high redshifts, one sees that there is preference for high-redshift nulling in order to maximize sensitivity.

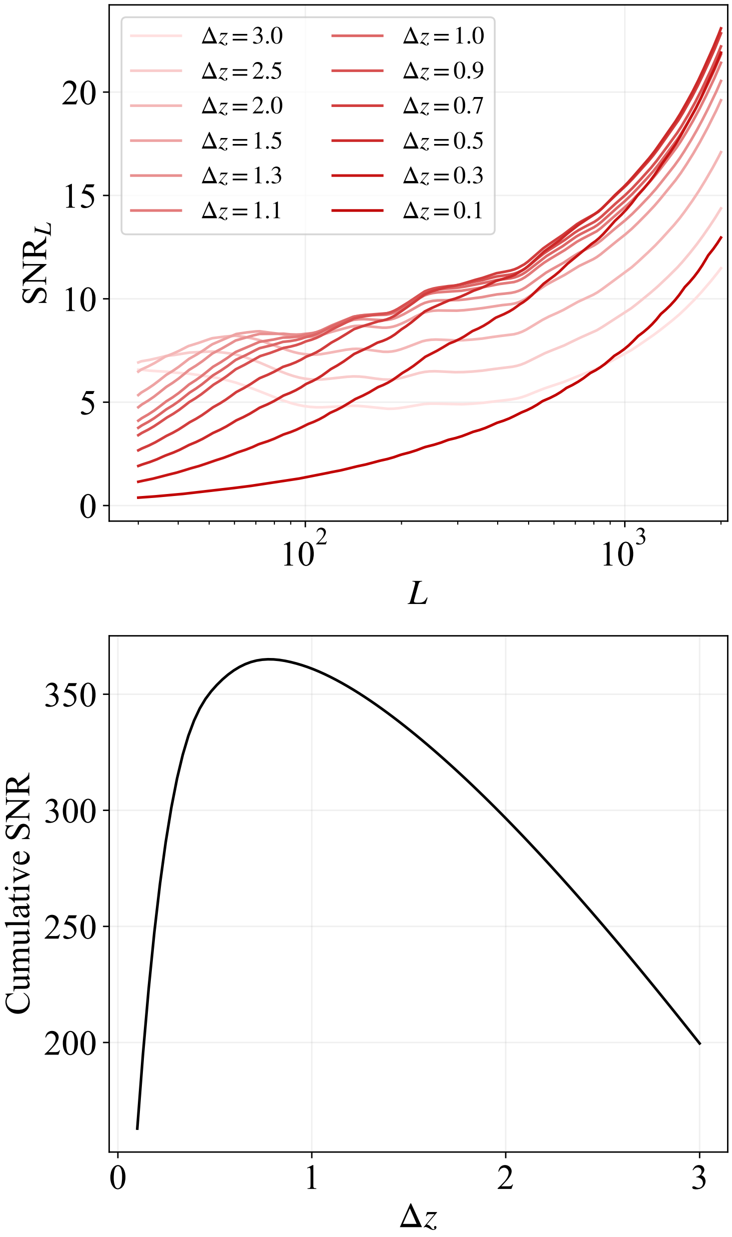

Figure 4.5: Same as Fig. 4.4, but optimizing for the redshift separation Δz between LIMs rather than the nulling redshift. In all cases, the higher redshift line is fixed at $z = 5.5$. In the noiseless cosmic variance-limited case the optimal SNR is achieved at $\Delta z = 0.7$, but this changes when instrumental noise is introduced in Section 4.4.

We consider CO, [CII], and Ly-α as possible lines with which nulling could be performed. Their angular power spectra as a function of redshift are plotted in Fig. 4.3. While the amplitude of the power spectra vary greatly between lines, devoid of noise and systematics, any pair of lines will yield the same SNR in the cosmic-variance-limited case; in this idealized case, the LIM-pair lensing reconstruction noise, $N_{XY}(L)$, is the same for any line pair, therefore the LIM-lensing SNR is fixed. Somewhat more intuitively, without the effects of noise and systematics, no line traces the line-of-sight gravitational potentials it encounters better or worse than any other since gravitational lensing is achromatic. This propagates through to nulling. Referring the reader back to Eq. (4.12), nowhere do the individual line spectra enter into $C_L^{\hat{\kappa}\hat{\kappa}_{\rm null}}$.

The quantities of paramount importance, then, which do enter into both the CMB × LIM-nulling convergence as well as its variance, are the line redshifts. The redshifts, or comoving distance to the source plane, of the lines enters into the nulling kernel in two ways: through the LIM lensing kernels themselves and through the parameter α. Referring back to Fig. 4.1 and 4.2, it might seem, by construction, that nulling at lower z would increase the amplitude $C_L^{\hat{\kappa}\hat{\kappa}_{\rm null}}$, bringing it closer to $C_L^{\hat{\kappa}\hat{\kappa}}$, thus resulting in higher SNR. However, the line redshifts also enter into the nulling variance through the cross spectrum terms $C_L^{\hat{\kappa}_i\hat{\kappa}_j}$. This nulling variance decreases as a redshift increases. We find that these effects combined, the increase in both the nulling spectrum amplitude and the LIM-nulling variance as z decreases, lead to a preference for high-z nulling. Plotted in Fig. 4.4 is the CMB × LIM-nulling SNR as a function of z and L, where the redshift separation of the lines is fixed at $\Delta z = 0.5$. As redshift increases the cumulative SNRs likewise increase.

The next choice to consider is the redshift separation of the lines. It would seem, looking at Fig. 4.1, that an ideal nulling scneario would have the two lines be as close as possible in redshift in order to get a sharp cutoff. In Fig. 4.5, the SNR as a function of redshift separation, Δz, is plotted. It is immediately obvious that both very small and very large separations of the lines are sub-optimal. The cumulative SNR peaks at $\Delta z = 0.7$ when nulling performed with the higher redshift LIM at z = 5.5.

It is important to note that $\Delta z = 0.7$ is not the optimal solution in general, since what must be optimized in the nulling equation is the interplay between the comoving distance separation of the three probes being used. While throughout this work we consider nulling performed on the CMB lensing convergence by LIMs, any three convergence maps may be used for nulling. For instance galaxy lensing convergence maps from $z \sim 2$, can be used to null a LIM convergence map at $z = 10$. The redshift optimization must be performed on a case-by-case basis since what is at work here is the interplay between the comoving distances of the three probes used in the nulling estimator. What is more, the addition of interlopers, diffuse foreground contaminants, and instrument systematics, which have non-trivial frequency evolution, complicate the problem. Care must be taken to optimize for the observing scenario at hand.

In this section, we provide the reader with intuition about how the LIM-nulling estimator depends on LIM parameters. We show that in the idealized case where noise and systematics are not included, the SNR depends only on the choice of line redshift and the redshift separation of the lines. Illustrating this with a concrete example, we presented the optimal redshift solutions for CMB nulling taking place at $z = 4.8$ in the cosmic variance limited case. We found that the LIM redshift which maximize the cumulative SNR are $z_1 = 4.8$ and $z_2 = 5.5$. In the presence of the frequency dependent noise of LIM experiments, this breaks down. In the following section, we explore how the CMB × LIM-nulling SNR scales with the sensitivity and the area of mm, sub-mm and IR surveys. We choose three fiducial observing scenarios to use in the subsequent forecasts. These include instrument noise and line interlopers, and we therefore re-optimize the LIM redshifts given these contaminants.

4.4 Survey Area and Sensitivity

In this section we explore the dependence of survey specifications on the CMB × LIM-nulling SNR. As previously mentioned, when devoid of instrument noise and systematics, all lines yield the same lensing reconstruction noise and therefore the same nulling SNR. Once such effects are added, this is no longer the case since the map-level SNR enters into $N_L^{\hat{\kappa}}$. In Section 4.4.1, we present the models used for computing the noise power of CO, [CII], and Ly-α-type experiments. In Section 4.4.2 we study how our nulling statistic depends on the intensity mapping survey sensitivities, while in Section 4.4.3, we explore how the SNR depends on their survey areas. Subsequently, we optimize the survey area for a given survey sensitivity to maximize the nulling SNR.

While the final forecasts in Sections 4.5 and 4.6 make use of the LIM-pair estimator, this exploratory section assumes each LIM lensing convergence map is estimated using a single line at a time. This is done to isolate the effects of a particular instrument in order to explore how the CMB × LIM-nulling SNR varies as a function of survey area and sensitivity for a single instrument. For example, we use [CII] observations from two channels of a given [CII] survey to perform nulling.

4.4.1 Instrument Noise Power

CO Experiments

In the case of CO experiments, we typically write the noise power as

$$P_N^{\mathrm{CO}} = \sigma_{\mathrm{vox}}^2 V_{\mathrm{vox}}^{\mathrm{CO}}, \tag{4.16}$$

where σ_{vox} is the noise in a single voxel and is given by

$$\sigma_{\mathrm{vox}} = \frac{T_{\mathrm{sys}}}{\sqrt{N_{\mathrm{det}} t_{\mathrm{pix}}^{\mathrm{CO}} \delta\nu}}. \tag{4.17}$$

Here, T_{sys} is the system temperature, N_{det} is the number of detector feeds, and $t_{\mathrm{pix}}^{\mathrm{CO}}$ is the observing time of a single pixel (Li et al., 2016). The time per pixel is related to the total observing time t_{obs} of the survey via $t_{\mathrm{pix}}^{\mathrm{CO}} = t_{\mathrm{obs}}(\Omega_{\mathrm{pix}}^{\mathrm{CO}}/\Omega_{\mathrm{surv}})$ where Ω_{surv} is the total survey area. A "pixel" here is defined to cover a solid angle $\Omega_{\mathrm{pix}}^{\mathrm{CO}} = \sigma_{\mathrm{beam}}^2$, where σ_{beam}^2 is the variance of the instrument's Gaussian beam. We compute σ_{beam} using the relation $\sigma_{\mathrm{beam}} = \theta_{\mathrm{FWHM}}/\sqrt{8\ln 2}$ where θ_{FWHM} is the beam full width half maximum. The comoving volume $V_{\mathrm{vox}}^{\mathrm{CO}}$ of a single voxel is the volume subtended by a pixel with angular size $\Omega_{\rm pix}^{\rm CO}$ and frequency resolution $\delta\nu$. Since the number of detectors, $N_{\rm det}$, constitutes a measure of the instantaneous sensitivity of the instrument, we parameterize the total survey sensitivity in terms of spectrometer-hours, $N_{\rm det}t_{\rm obs}$.

In the following subsection, we demonstrate how the SNR of the CMB × LIM-nulling convergence spectrum scales with survey area (or equivalently, $f_{\rm sky}$, the fractional sky coverage) and sensitivity, $N_{\rm det}t_{\rm obs}$. To that end, we compute the CO noise power using Eq. (4.16) and include it in our computation of the LIM lensing reconstruction noise, the single line analog to the LIM-par lensing estimator in Eq. (4.4). When varying the sensitivity and survey area, we anchor the remaining instrument specifications to the CO Mapping Array Project phase 2 (COMAP2) (Ihle et al., 2019). The specifications of this instrument as well as its current generation phase 1 counterpart (COMAP1) can be found in Table 4.1. COMAP aims to detect spectral lines from various rotational line transitions of CO, including the CO(1-0) during the peak of star formation around $z \sim 3$ when the CO luminosity function peaks, and will have some sensitivity of other CO transition lines out to $z \sim 8$ (Cleary et al., 2022; Breysse et al., 2022; Chung et al., 2022).

Parameter	COMAP1	COMAP2
$T_{\rm sys}$ (K)	40	40
$N_{\rm det}$	19	95
$\theta_{\rm FWHM}$ (arcsec)	4	4
$\Delta\nu$ (MHz)	15.6	15.6
$t_{\rm obs}$ (h)	6000	9000
$\Omega_{\rm surv}$ (deg^2)	2.5	2.5

Table 4.1: Instrument parameters for COMAP1 and COMAP2 experiments. These values are taken from Ihle et al. (2019).

[CII] Experiments

In the case of [CII] surveys, we follow Chung et al. (2020) and write

$$P_N^{\mathrm{CII}} = \frac{\sigma_{\mathrm{pix}}^2}{t_{\mathrm{pix}}^{\mathrm{CII}}} V_{\mathrm{vox}}^{\mathrm{CII}}. \tag{4.18}$$

Here, $t_{\mathrm{pix}}^{\mathrm{CII}}$ and $V_{\mathrm{vox}}^{\mathrm{CII}}$ are defined (by convention) slightly differently than the corresponding CO quantities. A "pixel" is defined to cover a solid angle $\Omega_{\mathrm{pix}}^{\mathrm{CII}} = 2\pi\sigma_{\mathrm{beam}}^2$. This new definition of the pixel size is then used, along with the frequency channel width, to compute $V_{\mathrm{vox}}^{\mathrm{CII}}$. Here,

$$t_{\mathrm{pix}}^{\mathrm{CII}} = \frac{N_{\mathrm{det}} t_{\mathrm{obs}}}{\Omega_{\mathrm{surv}}/\Omega_{\mathrm{pix}}^{\mathrm{CII}}}. \tag{4.19}$$

Analogously to the previous section on CO, when varying the sensitivity and survey area we fix the rest of the instrument parameters to those of the the CarbON CII line in post-rEionisation and ReionisaTiOn epoch (CONCERTO), a current generation high-z [CII] mapping experiment (Concerto Collaboration et al., 2020). CONCERTO will detect the [CII] line from $6 \lesssim z \lesssim 11$ over 1.4 deg^2 on the sky. The instrument specification used for simulating CONCERTO noise are summarized in Table 4.2. We also compute the SNR for a handful of other [CII] mapping experiments whose specifications are also summarized in Table 4.2.

Ly-α Experiments

For the noise power associated with infrared (IR) intensity mapping experiments targeting the high-z Ly-α line, we write the noise power spectrum $P_N^{\mathrm{Ly}\text{-}\alpha}$ as

$$P_N^{\mathrm{Ly}\text{-}\alpha} = \sigma_{\mathrm{vox}}^2 V_{\mathrm{vox}}^{\mathrm{Ly}\text{-}\alpha} \tag{4.20}$$

Parameter	FYST	CONCERTO	TIME	StageII
$\sigma_{\rm pix}$ (MJy/sr s$^{1/2}$)	0.86	11.0	11.0	0.21
$N_{\rm det}$	20	3000	32	16000
$\theta_{\rm FWHM}$ (arcsec)	46.0	22.5	22.5	30.1
$\delta\nu$ (GHz)	2.5	1.5	1.9	0.4
$t_{\rm obs}$ (hr)	4000	1200	1000	2000
$\Omega_{\rm surv}$ (deg^2)	16	1.4	0.6	100

Table 4.2: Instrument parameters for [CII] mapping experiments, including the Fred Young Submillimeter Telescope (FYST), CONCERTO, Tomographic Ionized Carbon Intensity Mapping Experiment (TIME), and a next-generation "Stage II" concept. The Stage II parameters are based on Silva et al. (2015), while the rest are from Chung et al. (2020).

where $V_{\rm vox}^{\text{Ly-}\alpha}$ is the single-voxel volume (defined in an identical way to analogous quantities as the [CII] and CO cases) and $\sigma_{\rm vox}$ is defined as

$$\sigma_{\rm vox} = s\sqrt{\frac{4\pi f_{\rm sky}}{t_{\rm obs}\Omega_{\rm pix}}}. \tag{4.21}$$

Here the instantaneous pixel sensitivity is given by s. Since the IR experiments we consider scan the sky one pixel at a time, the number of detectors is simply unity and we thus parameterize our total survey sensitivity by the total observing time, $t_{\rm obs}$, instead of spectrometer-hours. We set the other instrument specifications in Eqs. (4.20) and (4.21) to that of the Cosmic Dawn Intensity Mapper (CDIM) (Cooray et al., 2019). CDIM is a next generation optical and IR instrument aimed at detecting high redshift galaxies and quasars as well as spectral lines during cosmic dawn and reionization. We also compute the SNR for Spectro-Photometer for the History of the Universe, Epoch of Reionization and Ices Explorer (SPHEREx), an upcoming intensity mapping mission with Ly-α mapping capabilities (Doré et al., 2015). The specification for both CDIM and SPHEREx are summarized in Table 4.3.

Parameter	SPHEREx	CDIM
s (Jy/sr $s^{1/2}$)	231.46	38.58
θ_{FWHM} (arcsec)	6	2
t_{obs} (h)	2000	2000
Ω_{surv} (deg^2)	100	100
$R \equiv \lambda/\Delta\lambda$	150	300

Table 4.3: Instrument parameters for Ly-α experiments. Instrument specification for SPHEREx and CDIM are obtained from Doré et al. (2015) and Cooray et al. (2019) respectively. We compute the sensitivity, s, for a nominal 2000 hour survey given the quoted noise powers for both instruments. For SPHEREx, the spectral resolving power ($R \equiv \lambda/\Delta\lambda$, where λ is the observing wavelength and $\Delta\lambda$ is the wavelength resolution) is computed at $\lambda \simeq 4.5\,\mu m$. CDIM achieves R $\geq$ 300 over its whole bandwidth.

4.4.2 Dependence on Sensitivity

We compute the CMB × LIM-nulling convergence spectrum and its variance when nulling is performed with LIMs from $z = 5.5$ and $z = 3.5$. The LIM lensing reconstruction noise is computed using the single line analog of Eq. (4.9). We take $l_{\mathrm{min,LIM}} = 30$ and $l_{\mathrm{max,LIM}} = 5000$ where l_{min} is driven by the area of the survey and l_{max} is driven by the angular resolution of the instrument.For the CMB lensing reconstruction noise, we assume that of SO and use $N_L^{\hat{\kappa}}$ from the SO noise calculator.[3] It may be noted that some of the instruments we consider in this section do not probe modes as large as $l = 30$. For the purposes of singling out the effect of increasing sensitivity, we fix $l_{\mathrm{min,LIM}} = 30$ when varying the sensitivity. When computing the LIM-nulling SNR of particular experiments in following sections, we adjust $l_{\mathrm{min,LIM}}$ accordingly.

In Fig. 4.6, we plot the SNR per L mode as a function of L and spectrometer-hours $N_{\mathrm{det}}t_{\mathrm{obs}}$. Summing over all L, the *cumulative* SNR as a function of the number of spectrometer hours is plotted in Fig. 4.7. As expected, cumulative SNR increases as $N_{\mathrm{det}}t_{\mathrm{obs}}$ increases until it saturates to a plateau. For COMAP2-type instruments, this plateaus at $N_{\mathrm{det}}t_{\mathrm{obs}} \sim 10^4$ h, while for CONCERTO-

---

[3] `https://github.com/simonsobs/so_noise_models/tree/master/LAT_lensing_noise`

and CDIM-type experiments this occurs at $N_{\rm det}t_{\rm obs} \sim 10^6$ h and $N_{\rm det}t_{\rm obs} \sim 10^2$ h respectively .

Using the nominal survey specifications listed in Table 4.1, COMAP1 yields a cumulative SNR of 1.11 while COMAP2 has a cumulative SNR of 1.30. As for how current and upcoming [CII] experiments fare, Stage II has a cumulative SNR of 8.33 while all other experiments yield a cumulative SNR < 1. Nulling with SPHEREx results in a noise dominated measurement; however, CDIM achieves a cumulative SNR of 7.05. For CDIM, COMAP2, and CONCERTO, the nominal survey configurations are denoted with the diamond, circle, and square in Fig. 4.7, respectively. It is clear that perhaps with the exception of CONCERTO, these configurations are in the regime where the cumulative SNR has plateaued as a function of spectrometer-hours. Therefore, in order to increase the SNR further, other experimental parameters must be altered.

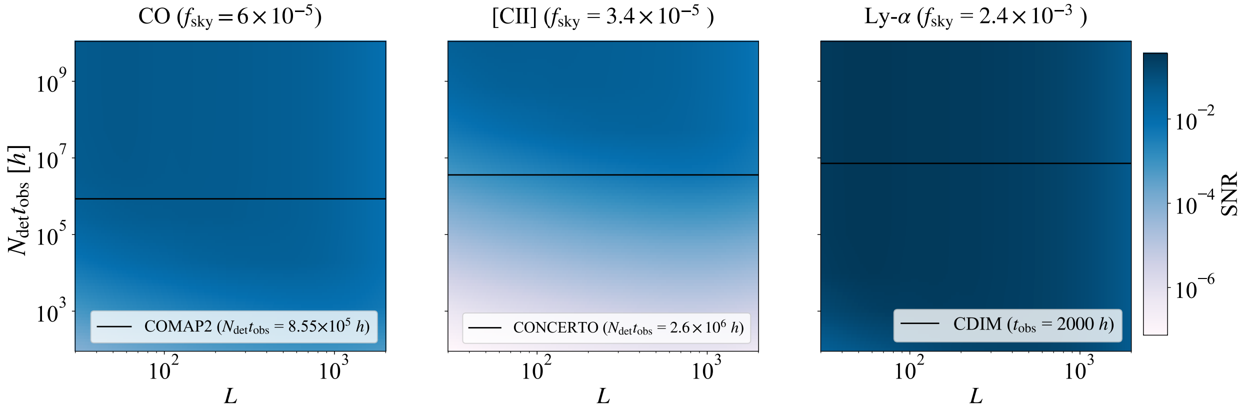

Figure 4.6: SNR as a function of lensing multipole, L, and sensitivity in terms of spectrometer-hours, $N_{\rm det}t_{\rm obs}$, for CO-type experiments (left), [CII]-type experiments (middle), and Ly-α-type experiments (right). For Ly-α, since $N_{\rm det} = 1$, the sensitivity is simply parameterized by $t_{\rm obs}$. The black horizontal line in each panel denotes the sensitivity of each line's nominal survey.

4.4.3 Dependence on Survey Area

Next we consider how the CMB × LIM-nulling SNR varies as a function of the survey area, parametrized by the fractional sky area $f_{\rm sky} \equiv \Omega_{\rm surv}/4\pi$. For any cosmological measurement

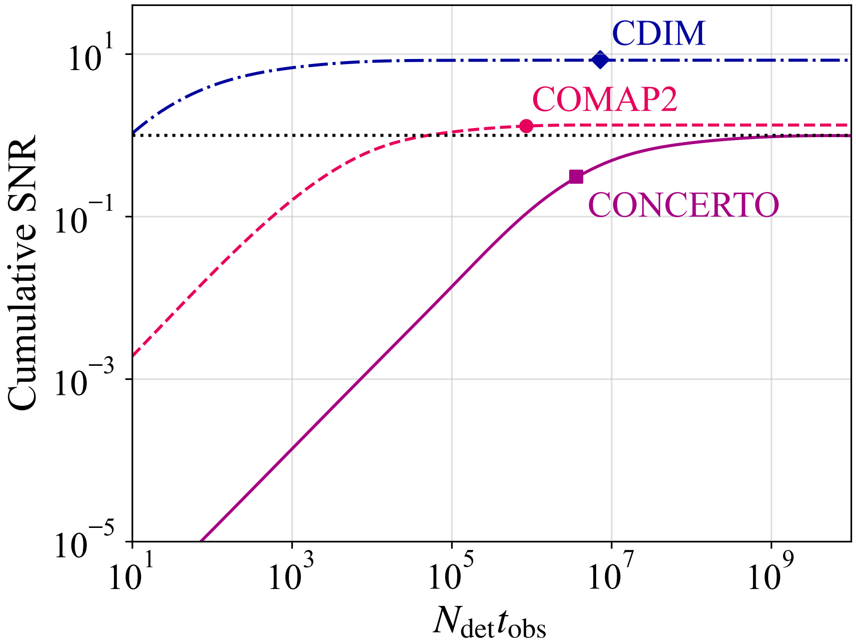

Figure 4.7: Cumulative SNR as a function of $N_{\rm det}t_{\rm obs}$ for Ly-α-type experiments (blue dot-dashed), CO-type experiments (pink dashed), and [CII]-type experiments (solid purple). The dotted black line denotes SNR = 1. It should be noted that for Ly-α experiments, $N_{\rm det} = 1$ and so the horizontal axis can simply be interpreted as $t_{\rm obs}$. The cumulative SNR for the nominal CDIM (blue diamond), COMAP2 (pink circle), and CONCERTO (purple square) surveys are also shown here. With the (mild) exception of CONCERTO, all three experiments roughly sit on the plateau where further increases in the number of spectrometer-hours do not increase SNR.

there is always a trade-off when increasing $f_{\rm sky}$. Keeping the sensitivity and observing time fixed, increasing the sky coverage of a survey means that less time is spent integrating on each pixel. This results in a shallower survey. This can be seen by examining Eqs. (4.16) and (4.17), which show that the noise power is proportional to $f_{\rm sky}$. However, increasing the sky coverage also increases the number of Fourier modes sampled and results in decreased sample variance per mode. This is reflected in Eq. (4.13) where the nulling variance sees a factor of $1/f_{\rm sky}$ out front. It is the optimization of these two effects that determines the optimal survey coverage for an instrument of a given sensitivity (Tegmark, 1997).

In Fig. 4.8 we plot the SNR as a function of L and $f_{\rm fsky}$ while in Fig. 4.9 , we simply plot the cumulative SNR as a function of $f_{\rm sky}$. We compute L up to 2000. This somewhat arbitrary cutoff was informed by our reconstruction noise curves which increase at high L. In addition, the small angular scales at the map level used in the reconstruction contain non-linear baryonic effects which are difficult to model accurately. For all three types of instruments, the cumulative SNR curves follow the same shape. The slope increases until it reaches a maximum, and subsequently decreases. This maximum survey area balances the two competing effects: the survey depth and the sample variance per mode. If the survey area is too small for a fixed survey duration, not enough modes are sampled and the SNR decreases. The region of the cumulative SNR curves to the left of the maximum value constitute this regime. If the survey area is too big for a fixed survey duration, many modes are sampled, but the survey is shallow, thus decreasing the cumulative SNR. This constitutes the region to the right of the maximum value.

CONCERTO is an excellent example of a survey whose area yields the maximum CMB × LIM-nulling SNR given its sensitivity. This is not by chance. The 1.4 deg^2 survey area for CONCERTO optimizes the SNR of its line power spectrum and we find that as an approximate rule of thumb, optimizing the line power spectrum SNR with respect to $f_{\rm sky}$ also maximizes the nulling SNR.

In contrast, we see from Fig. 4.9 that for COMAP2 and CDIM, the survey areas are not optimal. At present, the sky coverage of these surveys place them in the sample variance limited regime where the instantaneous sensitivity is relatively high, but simply not enough modes are sampled to truly maximize the capabilities of the instrument.

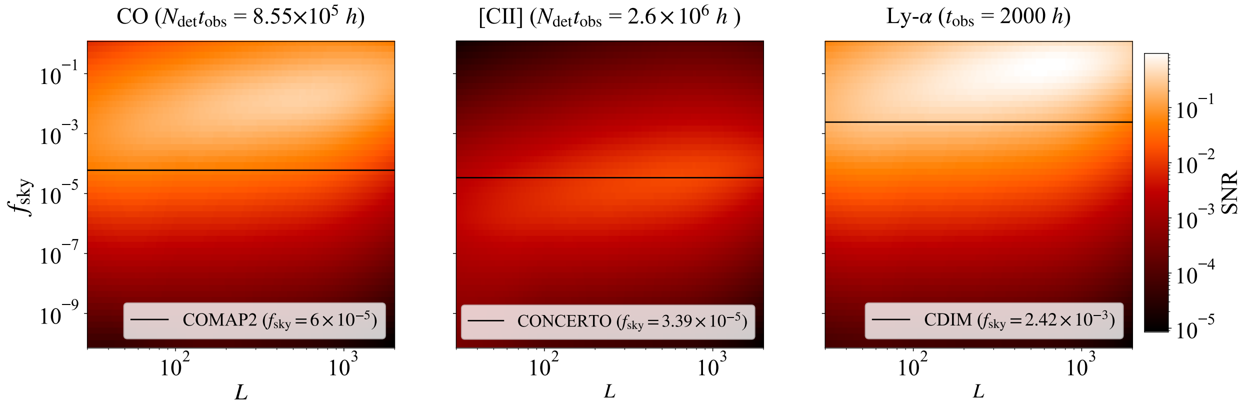

Figure 4.8: SNR as a function of lensing multipole, L, and survey area, $f_{\rm sky}$, for CO-type experiments (left), [CII]-type experiments (middle), and Ly-α-type experiments (right). The black horizontal line in each panel denotes the survey area of each line's nominal survey.

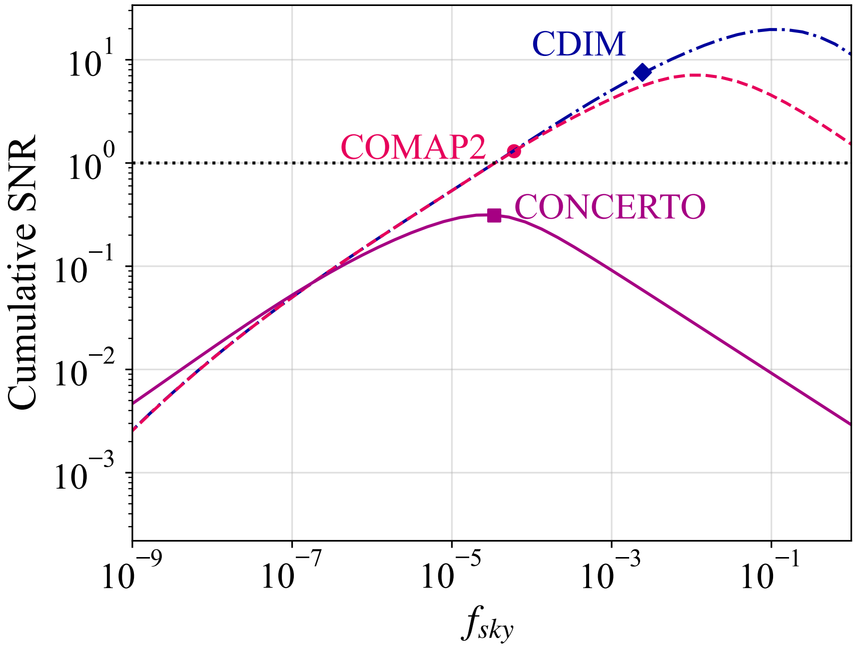

Figure 4.9: Cumulative SNR as a function of f_{sky} for Ly-α-type experiments (blue dot-dashed), CO-type experiments (pink dashed), and [CII]-type experiments (solid purple). The dotted black line denotes SNR = 1. The cumulative SNR for the nominal CDIM (blue diamond), COMAP2 (pink circle), and CONCERTO (purple square) surveys are also shown here. While CONCERTO is reasonably optimized to balance instrumental noise and cosmic variance, CDIM and COMAP2 would benefit from additional sky coverage as far as a nulling measurement is concerned.

4.4.4 Forecast Scenarios

Given the exploration into the observational parameter space of the last few sections, we now present the three observing scenarios that are used in our cosmological parameter estimation forecasts. Here we return to computing the full CMB × LIM-nulling convergence using the LIM-*pair* estimator, where the convergence maps at each redshift are constructed with two different lines via Eq. (4.4). We denote the scenarios as Current Generation, Next Generation, and Futuristic. We choose to perform LIM-nulling with LIM observations from [CII] and Ly-α mapping experiments given that these lines yielded higher SNRs in the near term. These scenarios are summarized in Table 4.4.

For our Current Generation scenario, we take [CII] to be observed by CONCERTO and Ly-α to be observed by SPHEREx at $z = 3.5$ and $z = 5.5$. In Section 4.3, we optimized the line separation in the noiseless case to be $\Delta z = 0.7$. In the presence of frequency-dependent noise, that is no longer the optimal line separation. Given the noise power of CONCERTO and SPHEREx, the line redshifts which maximize the SNR are $z = 5.5$ and $z = 3.5$. Of course, aside from pure sensitivity concerns, one must also account for systematics such as interloper lines. At high redshifts, Ly-α is contaminated by low-redshift H-α emission while the high-redshift [CII] line is contaminated by low-redshift CO emission. Using `HaloGen`, we generate H-α and CO spectra at the appropriate redshifts and include these when computing the LIM-pair lensing reconstruction noise. We assume the observations have undergone foreground removal, leaving behind a 10% interloper residual power for each line. To simulate this, we simply multiply the generated interloper spectra by an overall factor of 0.1.

Although SPHEREx has two large 100 deg^2 deep fields at the poles, CONCERTO has a much smaller field of view, at 1.4 deg^2. Therefore, $\kappa_{\rm null}$ can only be computed over the small overlapping field of 1.4 deg^2. As a result, we compute LIM-lensing reconstruction noise with the largest angular scale $l_{\rm min,LIM} = 153$ (~1 deg). We take the finest angular scale to be $l_{\rm max,LIM} = 5000$ (~ 0.04 deg), which is much coarser than the angular resolution of both instruments. For the CMB instrument we choose SO where lensing reconstruction is performed with minimum and maximum spherical harmonic l of $l_{\rm min,CMB} = 30$ and $l_{\rm max,CMB} = 5000$, respectively. The total SNR in this scenario is 0.1, suggesting that nulling estimation will likely be a future endeavour.

For the Next Generation scenario, we take [CII] to be observed by a Stage II instrument at $z = 5.5$ and Ly-α to be observed by CDIM at $z = 4.5$. Again, these redshifts are obtained by optimizing the nulling SNR for a given noise level. We assume the observations contain 5% interloper residual power for each line. Here, both instruments are expected to survey a 100 deg^2 field and we assume that they overlap entirely. We compute the LIM-lensing reconstruction noise with $l_{\rm min,LIM} = 30$ (~ 6 deg) and $l_{\rm max,LIM} = 10000$ (~ 0.02 deg), again, using scales coarser than to angular resolution of both instruments. For the CMB instrument we simulate the noise power of CMBS4 where lensing reconstruction is performed with $l_{\mathrm{min,CMB}} = 30$ and $l_{\mathrm{max,CMB}} = 5000$. The CMBS4 lensing reconstruction noise used here can be found in Sailer et al. (2021). The total SNR for the Next Generation scenario is 9.5, representing a firm detection that will be an important proof-of-concept for the nulling technique. However, it will perhaps still not quite be the high-precision measurement that unlocks high-precision science.

Finally, we construct a Futuristic scenario that guarantees a high SNR measurement. We consider [CII] observed by a Stage II-like instrument at $z = 5.5$ and Ly-α to be observed by a CDIM-like instrument at $z = 4.5$, over a quarter of the sky. By this we mean that the [CII] maps are the same depth as those expected from Stage II but over a larger area of the sky. Likewise the CDIM-like instrument is one that produces maps at the same depth as the nominal CDIM survey but again over a larger portion of the sky. This can be achieved by using an instrument with the same instantaneous sensitivity and increasing the total observing time of the survey until the instantaneous integration time reaches that of the nominal surveys. This can also be achieved by increasing the scanning rate of the instrument such as to observe a larger portion of the sky in the same total observing time, but correspondingly increasing the instantaneous sensitivity in order to obtain the same depth. Of course, a combination of these two strategies will also suffice. We assume the observations contain 1% interloper residual power for each line. Like the last scenario, we compute the LIM-lensing reconstruction noise with $l_{\mathrm{min,LIM}} = 30$ and $l_{\mathrm{max,LIM}} = 10000$. For the CMB instrument we again choose the CMBS4 where lensing reconstruction is performed with $l_{\mathrm{min,CMB}} = 30$ and $l_{\mathrm{max,CMB}} = 5000$. The Futuristic scenario has a total SNR of 110. As we will demonstrate in Section 4.5, in this regime one is able to obtain competitive parameter constraints.

In Table 4.4, we also quote the total SNR for a cosmic variance limited case to showcase the upper bound of what is achievable in an idealized case. Here we include no instrument noise nor interloper contaminants. Like the Futuristic scenario, we assume $l_{\mathrm{min,LIM}} = 30$ and $l_{\mathrm{max,LIM}} = 10000$ and $l_{\mathrm{min,CMB}} = 30$ and $l_{\mathrm{max,CMB}} = 5000$, but we set $f_{\mathrm{sky}} = 1$. In this case, the CMB × LIM-nulling cumulative SNR is 408.

Scenario	Ly-α	[CII]	CMB	Survey Area	$l_{\mathrm{min,LIM}}$	$l_{\mathrm{max,LIM}}$	Interloper Residual	$\sum$SNR
Current Generation	SPHEREx	CONCERTO	SO	1.4 deg^2	153	5,000	10%	0.1
Next Generation	CDIM	Stage II	CMBS4	100 deg^2	30	10,000	5%	9.5
Futuristic	CDIM	Stage II	CMBS4	$f_{\mathrm{sky}} = 0.25$	30	20,000	1%	110
CV Limited	–	–	–	$f_{\mathrm{sky}} = 1$	30	20,000	–	408

Table 4.4: Observational scenarios used for our forecasts, in addition to the cosmic variance limited case as a high-SNR reference. While current instruments will not make nulling detections, their successors will be capable of not just detections but also high-SNR characterizations that will be scientifically interesting.

4.5 Fisher Forecast: ΛCDM + M_ν Cosmology

In this section we present a Fisher forecast on potential parameter constraints from CMB × LIM-nulling when $0 < z \lesssim 5$ has been nulled. We take our parameter set to be the standard ΛCDM parameters plus the sum of the neutrino masses M_ν. We begin by providing a brief overview of the Fisher formalism to establish notation, and follow by presenting the parameter covariances for our fiducial cosmology. Despite $C_l^{\hat{\kappa}\hat{\kappa}_{\rm null}}$ having worse constraints than regular CMB lensing measurements by construction, we discuss how $C_l^{\hat{\kappa}\hat{\kappa}_{\rm null}}$ uniquely probes high-redshift physics. We show how such a probe could serve as a model-independent test of non-standard time evolution, illustrating such a test by constructing an *ad hoc* cosmology with deviations on the scale of the current Hubble tension and the σ_8 tension (Valentino et al., 2021; Poulin et al., 2023; Douspis et al., 2019).

4.5.1 General Fisher Formalism

The Fisher information matrix captures how much information an observable, O, measured to some precision, carries about a set of model parameters that are grouped into a vector $\boldsymbol{\theta}$. The elements of the Fisher matrix are given by

$$F_{ij} = \sum_l \frac{1}{\sigma_l^2} \frac{\partial O_l}{\partial \theta_i} \frac{\partial O_l}{\partial \theta_j} \tag{4.22}$$

where i and j index the parameters in the model, l indexes each measured mode of the observable, and σ_l is the error on the measurement of that mode. The Fisher matrix is the inverse of the covariance matrix, $\mathbf{F}^{-1} = \mathbf{C}$, for the set of model parameters. In our case, we perform two independent forecasts, one with $O = \{C_l^{\hat{\kappa}\hat{\kappa}_{\rm null}}\}$ (the forecast of interest) and one with $O = \{C_l^{\hat{\kappa}\hat{\kappa}}\}$ (which serves as a reference). We compute these spectra using Eqs. (4.3) and (4.12) and take their numerical derivatives using finite differences. The Fisher matrix encodes Gaussian parameter uncertainties; by the Cramer-Rao bound, this provides an optimistic approximation of the true posterior distributions (Coe, 2009). In reality, the true uncertainties may be larger than those modeled due to other systematics and may be non-Gaussian.

In addition to the parameter covariance matrix, one can compute the bias on each parameter in the presence of a systematic. Given some observable C_l^{obs} that contains the the signal of interest, C_l, and also some systematic contaminant, C_l^{cont}, the total observed quantity is given by

$$C_l^{\mathrm{obs}} = C_l + C_l^{\mathrm{cont}}. \tag{4.23}$$

Following Huterer and Takada (2005) the ith component of the parameter bias vector $\mathbf{b}$ is given by

$$b_i = \langle \hat{\theta}_i \rangle - \langle \theta_i^{\mathrm{true}} \rangle = \sum_j (\mathbf{F}^{-1})_{ij} \mathbf{B}_j \tag{4.24}$$

where $\hat{\boldsymbol{\theta}}$ contains the best fit parameter values, $\boldsymbol{\theta}^{\mathbf{true}}$ is vector containing the true underlying values, and $\mathbf{B}_j$ is

$$\mathbf{B}_j = \sum_l \frac{C_l^{\mathrm{cont}}}{\sigma_l^2} \frac{\partial C_l}{\partial \theta_j}. \tag{4.25}$$

In the following section, we present the results of a Fisher forecast for the Next Generation and Futuristic nulling scenarios and compare it to constraints from regular CMB lensing measurements. Following this, in Sections 4.5.3 and 4.5.4 we compute the parameter bias vector in a slightly unusual application of the formalism: we consider the case where an incorrect cosmological model results in a “theory systematic” that perturbs the inferred parameter values.

4.5.2 Concordance Cosmology

Here we present the results of the Fisher forecast with respect to a concordance model of cosmology whose parameters and fiducial values are summarized in Table 4.5. In addition to the independent model parameters used to define the fiducial cosmology, we also compute four derived quantities: Ω_{m}, σ_8, $S_8 \equiv \sigma_8(\Omega_{\mathrm{m}}/0.3)^{0.5}$, and $S_8^{\mathrm{CMBL}} \equiv \sigma_8(\Omega_{\mathrm{m}}/0.3)^{0.25}$. The quantity σ_8 is the root-mean-squared variance of density perturbations on 8 h^{-1} Mpc scales, and is given by

$$\sigma_8^2 = \int_0^\infty \frac{k^2 dk}{2\pi^2} P_m^{\rm lin}(k) \left[\frac{3 j_1(kR)}{kR}\right]^2 \tag{4.26}$$

where $P_m^{\rm lin}(k)$ is the matter power spectrum at $z = 0$ assuming *linear* theory, $R \equiv 8\,h^{-1}$ Mpc, and j_1 is the first-order spherical Bessel function of the first kind.

In Fig. 4.10, the posterior distributions are shown. In orange, the constraints from regular CMB lensing measurements forecast for SO are plotted. The green contours show the constraints from the Next Generation nulling scenario and while the Futuristic nulling scenario is in blue. *Planck* 2015 priors excluding *Planck* lensing have been applied to all cases (Planck Collaboration, 2016). The black dashed lines mark the fiducial values of the model.

When comparing the CMB lensing contours to the nulling ones, it is immediately obvious that the LIM-nulling probes contain less information than that of regular CMB lensing. This is entirely expected, since nulling by construction removes the low redshift information from the CMB convergence. This low-redshift information does of course have constraining power. Even with this being the case, both LIM-nulling measurements do add non-negligible information to the *Planck* prior and still provide comparable constraints to CMB lensing. For both CMB and CMB × LIM-nulling cases, A_s and τ are not well constrained; they trace the prior. The parameters to

Parameter	Definition	Fiducial Value
H_0	Hubble constant [km/s/Mpc]	67.5
$\Omega_b h^2$	Fractional baryon density	0.022
$\Omega_c h^2$	Fractional dark matter density	0.120
M_ν	Sum of the neutrino masses [eV/c^2]	0.06
A_s	Primordial fluctuation amplitude	2×10^{-9}
n_s	Spectral index	0.965
τ	CMB Optical depth	0.06

Table 4.5: Model parameters and their fiducial values.

which they are most sensitive are $\Omega_c h^2$, H_0, and M_ν, in that order. As for the derived quantities (Ω_m, σ_8, S_8, and $S_8^{\rm CMBL}$), one does see some hints that select combinations of Ω_m and σ_8 are better constrained by our nulling estimator—as is usually the case for lensing. However, we caution that the Fisher formalism is not equipped to fully capture the shapes of degenerate joint posteriors between parameters. In order to make more definitive claims about these parameters, a full sampling of the posterior (e.g., via Markov Chain Monte Carlo techniques) would be more appropriate, and so we omit these parameters from the analyses in subsequent sections.

It may seem on the face of it that LIM-nulling is not worth the effort, given that it requires one to make high significance detections of not just one, but three cosmological probes—just to obtain constraints on ΛCDM parameters that are less competitive (albeit comparable) to those of CMB lensing! Yet, it is important to appreciate that CMB lensing and LIM-nulling lensing are not measuring the same thing. The power of LIM-nulling is that it is a clean probe of the high-redshift universe exclusively. To illustrate this, we forecast how the constraints of these probes differ in a universe in which we have an unexpected time evolution of parameters.

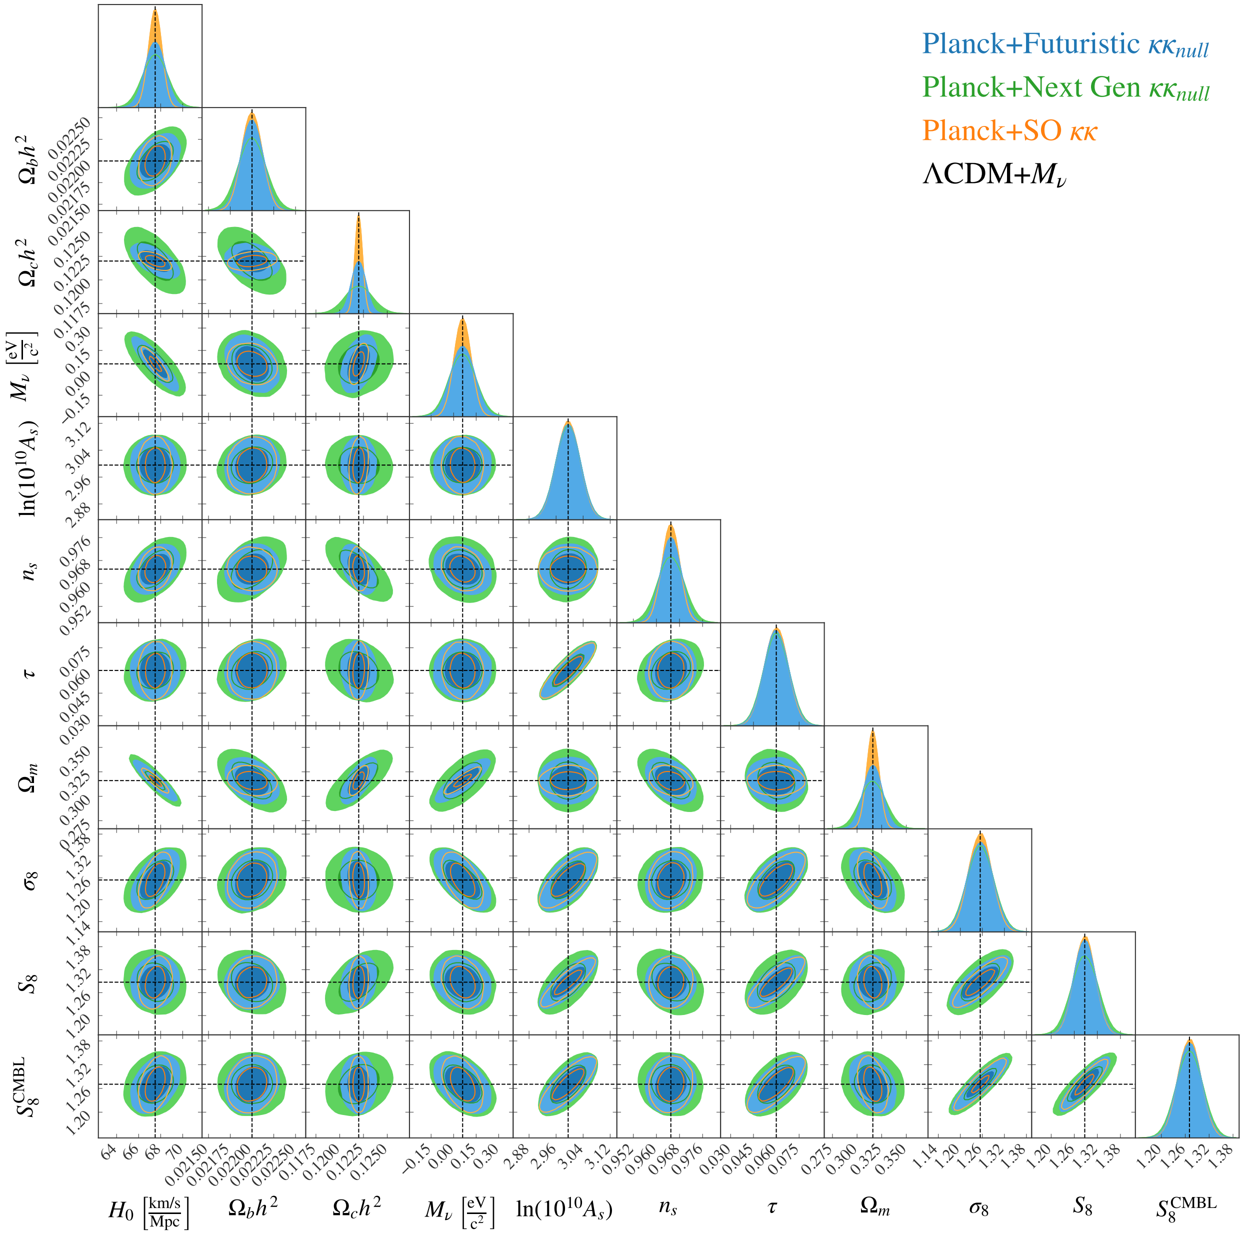

Figure 4.10: Forecasted posterior distributions for the concordance cosmology scenario of Section 4.5.2 from SO lensing (orange), LIM-nulling in the Futuristic scenario (blue), and LIM-nulling in the Next Generation scenario (green). The dark inner region of the contours indicate the 68% confidence region while the light coloured outer contours denote the 95% confidence region. The dashes black lines denote the fiducial parameter values. All contours share the same Planck prior on ΛCDM + M_ν.

4.5.3 Early- and Late-Time Parameter Consistency: The Hubble Parameter

Here we construct mock lensing data in a universe that obeys a cosmology which deviates from standard ΛCDM. Our toy model is inspired by the Hubble tension. When computing the lensing convergence spectrum (i.e. Eq. 5.3), for the integration steps where $z > z_{\rm null}$ the value of H_0 that enters into $H(z) = H_0\sqrt{\Omega_m(1+z)^3 + \Omega_\Lambda}$ and into the initialization of `CAMB` to obtain the matter power spectrum is 67.7 km/s/Mpc. This means that for $z > z_{\rm null}$ the matter power spectrum evolves as usual in a ΛCDM cosmology and *if* evolved all the way to $z = 0$ *would* reach the value 67.7 km/s/Mpc. When evaluating integration steps where $z < z_{\rm null}$ the same procedure is followed, but the value of H_0 that enters into $H(z)$ and into the initialization of `CAMB` to obtain the matter power spectrum is 72 km/s/Mpc. This model can also be described as a scenario where the growth factor, $D(z)$, and therefore the amplitude of $P(k)$ undergoes a sudden change at $z = z_{\rm null}$.

The two values of H_0 that were considered, 67.7 km/s/Mpc and 72 km/s/Mpc, constitute the discrepancy between low- and high-redshift measurements of the Hubble constant (Valentino et al., 2021). This Hubble tension remains one of the outstanding problems of the last decade. Some argue that a time evolving cosmology may be to blame (Benetti et al., 2019, 2021). If this were the case, LIM-nulling may be able to help elucidate this mystery as it contains no information about the late-time evolution of the matter density field, allowing a clean measurement of what happens at high redshifts. To be clear, we do not argue that the model proposed here is a genuine solution to the Hubble tension, nor that the LIM-nulling technique is destined to detect the tension. We are simply using this commonly known open cosmological problem as inspiration for how one might look for parameter consistency between early- and late-time measurements.

We fit CMB lensing measurements, which are sensitive to both values of H_0, and CMB × LIM-nulling measurements, which are only sensitive to the high redshift value of H_0, to the same fiducial model in which the Hubble parameter has not undergone an abrupt shift and therefore has one H_0 parameter. After all, one's initial null hypothesis in a real data analysis pipeline will likely not assume any abrupt shifts in cosmological parameters. Performing parameter fits using this incorrect cosmological model, we expect the CMB lensing measurements to be biased. We quantify the parameter biases using Eq. (4.24) and Eq. (4.25), treating the systematic contaminant to be the *difference* between the true lensing contribution from $0 < z < z_{\rm null}$ and the incorrect contribution assuming a single value of H_0.

In Fig. 4.11, we show the constraints on CMB lensing from CMBS4 and on LIM-nulling measurements from our Futuristic scenario. By construction, the constraints from the $C_l^{\hat{\kappa}\hat{\kappa}_{\rm null}}$ fit contain no bias with respect to the fiducial model parameters. This is not the case for the constraints from CMB lensing. It is immediately clear that there is a discrepancy between the constraints from CMB lensing and those from LIM-nulling. Such measurements would constitute a tension that provides evidence that the behaviour underlying the CMB lensing data is better described by another model. And indeed, the CMB lensing data contains information about the universe before and after the abrupt change to $H(z)$, while the model does not.

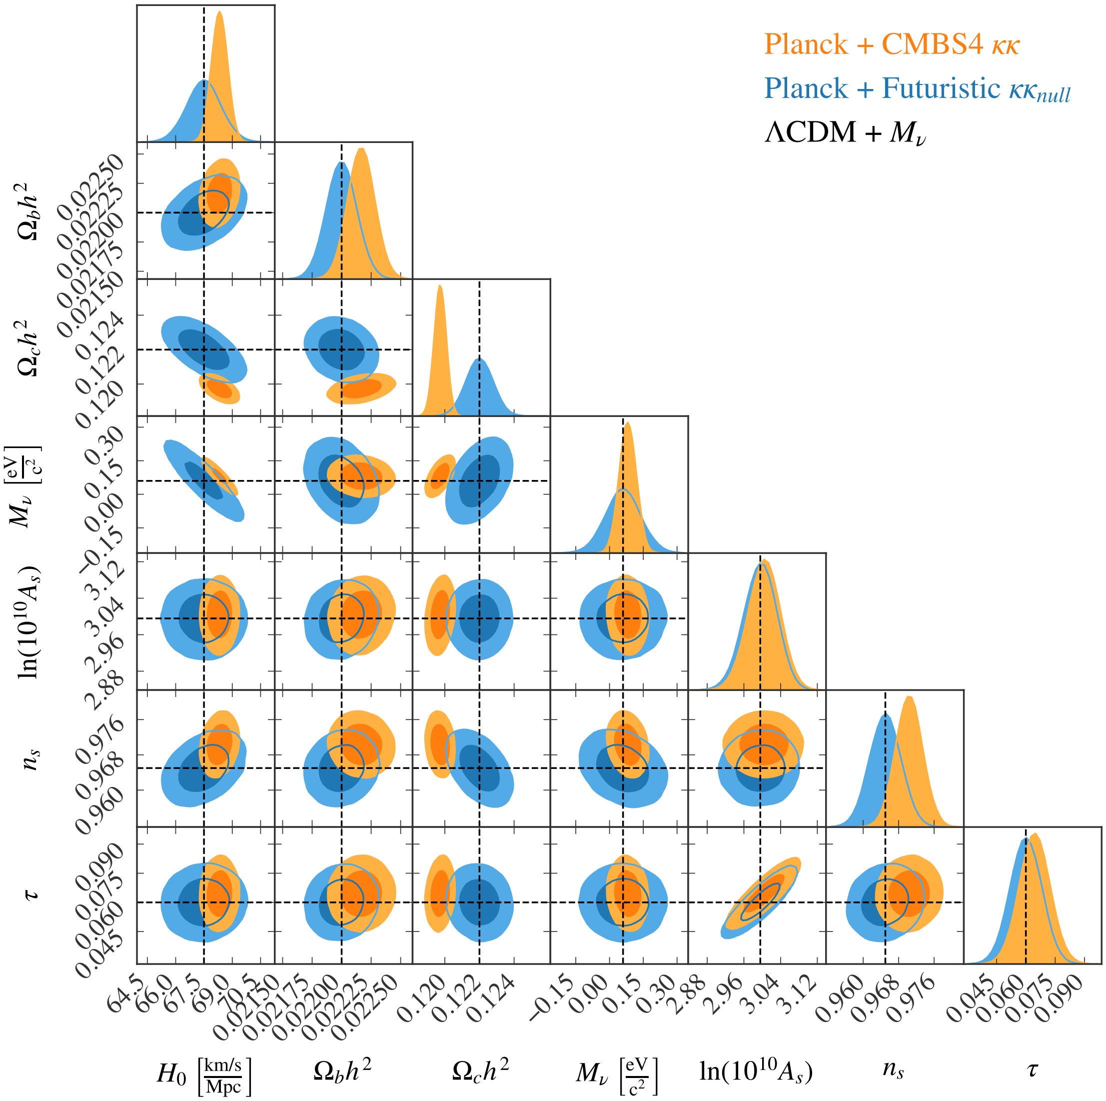

Figure 4.11: Forecasted posterior constraints from the abruptly evolving $H(z)$ cosmology of Section 4.5.3. Constraints are from CMBS4 lensing (orange) and LIM-nulling in the Futuristic scenario (blue). The dark inner region of the contours indicate the 68% confidence region while the light coloured outer contours denote the 95% confidence region. The dashes black lines denote the fiducial parameter values. These contours showcase a tension between the CMB lensing measurements and the CMB × LIM-nulling measurements. Both contours share the same Planck prior on ΛCDM + M_ν. Since LIM nulling is sensitive only to the high-redshift universe, a comparison between the contours allows for model-independent tests of unexpected differences between high and low redshifts.

The parameter for which the tension is the strongest is $\Omega_c h^2$ at just over 2σ. Recall that this parameter sees the most improvement when lensing data is added to the *Planck* prior. It is also worth commenting on the role the prior plays in this analysis. If one had some credence that the universe did not follow ΛCDM, one may choose to relax the *Planck* prior as has been done for a number of weak lensing analyses. We found that the prior covariance can be inflated by an overall factor of 10 and still yield a $\sim 2\sigma$ tension in the inferred value of $\Omega_c h^2$ while inflating the prior by an overall factor of 20 results in a $\sim 1\sigma$ tension of $\Omega_c h^2$. When the prior is relaxed any further the tension is no longer significant. We acknowledge that in practice, more rigorous analysis would be required to claim a tension between measurements (Handley and Lemos, 2019; Park and Rozo, 2020; Marshall et al., 2006; Lemos et al., 2021).

An advantage of this type of measurement is that making a higher significance detection of the bias between CMB lensing and LIM-nulling can result from either improving the CMB lensing measurement *or* the LIM-nulling measurement *or* both. Finally, it is important to note that nowhere in the analysis did we need to make any assumption about exactly what new cosmological model *did* in fact fit both data sets to detect a deviation from the standard cosmology. This constitutes a model independent test of cosmology beyond ΛCDM.

The reader may have noticed that in order to null the low-redshift contribution to CMB lensing, one needs to compute the nulling coefficients, α, given by Eq. 4.11, which depend on the comoving distance, $\chi(z)$, and therefore the cosmology. This may seem to negate the claim of model independence if an assumption about the fiducial cosmology is needed to null in the first place. In practice, nulling can be performed through minimizing the amplitude of the LIM-nulling convergence spectrum which preserves the model independence. What nulling does is makes use of LIM lensing information by removing it from CMB lensing maps. In doing so, the CMB × LIM-nulling convergence spectrum will have a smaller amplitude than the total CMB convergence spectrum as seen in Figure 4.2. However, there is a limit to how low that spectrum can go; one can only remove as much information as is present in the LIM lensing maps. Therefore, an equivalent nulling estimator is one that minimizes the amplitude of the CMB × LIM-nulling convergence spectrum. We have explicitly verified that this minimization scheme is equivalent to the one presented in Section 4.2.3. In practice, one can solve for the cosmology which yields the values of α that actually minimize the nulling convergence spectrum. Not only does this scheme ensure true model independence, but solving for the nulling coefficients, α, is itself a test of cosmology and can also serve as a test for residual systematics.

4.5.4 Early- and Late-Time Parameter Consistency: Matter Fluctuation Amplitude

Here we present a similar example to that of the previous section. We construct a scenario in which there is an abrupt change to the growth factor at $z_{\rm null}$ as a result of varying the value of A_s for integration steps before and after $z = z_{\rm null}$. If one were to compute σ_8 with and without taking into account the changing value of A_s one would obtain different values of σ_8 which corresponds roughly to the current σ_8 tension. Once again if the value of A_s used to initialize the matter power spectrum for integration steps at high-z is evolved to today (i.e. the standard ΛCDM scenario), one would infer $\sigma_8 = 0.79$. If one were to take into account the abrupt evolution of the matter power spectrum at $z = z_{\rm null}$ (i.e. our exotic cosmology) the inferred value of $\sigma_8 = 0.72$ (Mohanty et al., 2018). Similar to the previous section, LIM-nulling would not know about the shift since it is only sensitive to the high-redshift value of A_s while CMB lensing is sensitive to both values as it probes the matter density field before and after this sudden evolution. We therefore use the same parameter bias formalism to perform our forecasts.

In Fig. 4.12, we show the constraints on CMB lensing from CMBS4 and on LIM-nulling measurements from our Futuristic scenario. The constraints from $C_l^{\hat{\kappa}\hat{\kappa}_{\rm null}}$ are not biased with respect to the fiducial model parameters and, while not as extreme as the results from the previous section, there remains a slight bias to the CMB lensing constraints. There is a $\sim 1\sigma$ tension between CMB lensing and LIM-nulling measurements for $\Omega_c h^2$, perhaps providing a slight hint that the CMB lensing data and the LIM-nulling data are not adequately described by the same model. If one were to relax the prior by a factor of a few ($\lesssim 5$) the slight tension is preserved but, since the tension is not as large as it was in the last section, the tension is no longer significant if relaxed any further.

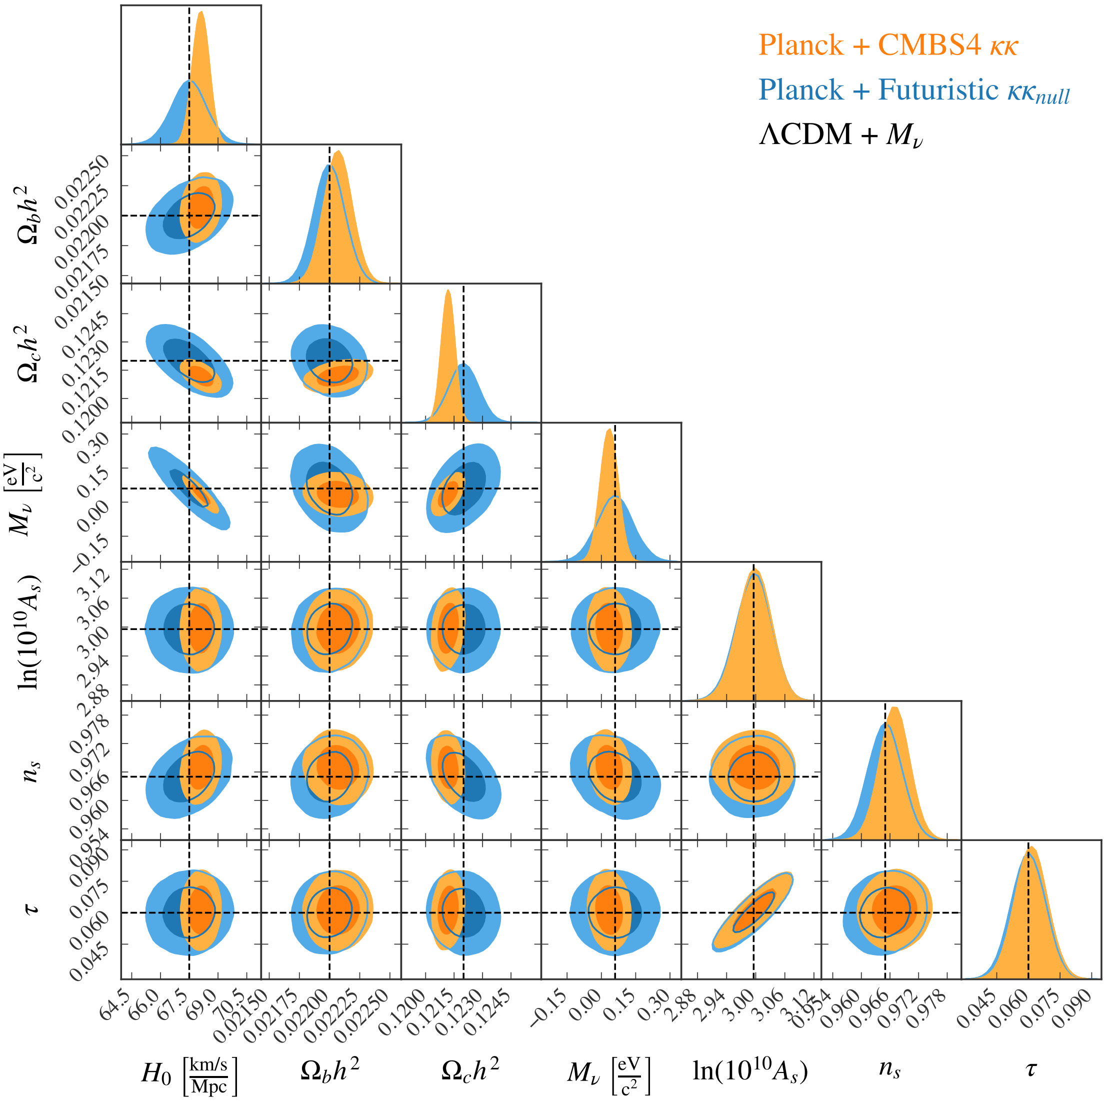

Figure 4.12: Same as Fig. 4.11, except for the A_s-evolving cosmology of Section 4.5.4. While a slight tension is still evident, the results are less statistically significant than for the abruptly evolving $H(z)$ cosmology.

4.6 Sensitivity To the Matter Power Spectrum

Ultimately, the CMB × LIM-nulling convergence is a kernel weighted map of the high redshift matter density field. While lensing measurements from galaxy shear and convergence and the CMB have provided the first unbiased measurements of the matter density field out to $z \sim 3$, the majority of our knowledge of the matter density field comes from the use of biased luminous tracers. LIM-nulling has the potential to reveal the unbiased high-redshift matter density field over large cosmological volumes. This leads to the additional advantage that the high-redshift matter perturbations are more linear, making the power spectrum easier to model. We forecast the sensitivity of $C_L^{\hat{\kappa}\hat{\kappa}_{\text{null}}}$ to the matter power spectrum at various length-scales and redshifts. Using the same Fisher formalism from the previous section, we define a set of parameters which are the amplitude of the matter power spectrum in various (k, z) bins. In other words, we take $\boldsymbol{\theta} = [P_m(k_1, z_1), P_m(k_2, z_1), ..., P_m(k_{\text{max}}, z_{\text{max}})]$, where k_{max} and z_{max} are the maximum k and z values, respectively. We define 2000 (k, z) bins, ranging from $0 < z < 1100$ and 10^{-2} h Mpc^{-1} $< k < 1$ h Mpc^{-1}.

Once again, for the sake of comparison we forecast the sensitivity of both $C_L^{\hat{\kappa}\hat{\kappa}}$ and $C_L^{\hat{\kappa}\hat{\kappa}_{\text{null}}}$ to the parameter vector $\boldsymbol{\theta}$. While CMB lensing probes length-scales across the whole redshift range, the LIM-nulling spectrum is insensitive to $z < z_{\text{null}}$ leaving a large portion of its Fisher information matrix null. Unable to invert such a matrices, we perform principal component analysis (PCA) on the Fisher matrices. We diagonalize each of the Fisher matrices and plot the six principal eigenmodes (i.e. the eigenvectors which have the largest eigenvalues).

In Figs. 4.13 and 4.14, we present the principal components for CMB lensing observed with CMBS4 and for LIM-Nulling in the Futuristic scenario, respectively. As expected, CMB lensing is sensitive to low redshift modes as its lensing kernel peaks at $z \sim 2$. The principal components are all sensitive to roughly the same range in L, as one can see from the lines of constant L superimposed on the figures. Roughly speaking, each successive principal component probes finer and finer oscillatory modes as a function of L. For this CMB lensing case, the first ~500 eigenmodes can be detected with SNR > 1.

Because the principal components follow contours of constant L, they contain information about many modes of wavenumber k from various redshifts. This mixing of matter power spectrum modes into a single L is the reason why the CMB convergence spectrum is smooth despite it being made up of the matter power spectrum which does contain acoustic peaks. This mode-mixing effect can be seen in the bottom panel of Fig. 4.2.

In the LIM-nulling case, things are appreciably different. As expected, the LIM-nulling eigenmodes peak at $z \sim z_{\rm null}$ since this is precisely where the nulling kernel peaks. The principal components therefore contain information about the high redshift modes of the matter power spectrum. Like in the CMB lensing case, the LIM-nulling eigenmodes trace lines of constant L; however: as redshift increases the relationship between k and L tends to one-to-one. This is driven by the fact that the relationship between angular scales and transverse comoving distances evolves slowly at high-z. Therefore, measuring a single L in the LIM-nulling spectrum contains information about a much narrower range of power spectrum modes, which allow for one to more cleanly trace matter fluctuations. Again, this is why BAO peaks emerge in the LIM-nulling convergence spectrum as seen in Fig. 4.2. It should also be noted the LIM-nulling eigenmodes are not simply the CMB lensing ones but with the low-z portion removed. Due to the differences in the shapes of their SNR curves, these spectra probe slightly different L. However, in a conceptually similar situation to the CMB lensing case, in LIM-nulling each successive eigenmode roughly corresponds to probing finer and finer features in the matter power spectrum. For the nulling case, the first ~60 eigenmodes can be detected with SNR > 1. As expected this is a lower cumulative SNR detection than with CMB lensing alone, but again, the value of the nulling estimator is its clean sensitivity to the high-z matter power spectrum. It is encouraging that one is able to attain a high-significance detection of principal modes of the unbiased matter power spectrum at $z \gtrsim 5$.

The choice to constrain power spectrum modes in the range $10^{-2}\ h\ \mathrm{Mpc}^{-1} < k < 1\ h\ \mathrm{Mpc}^{-1}$ was no accident. Indeed this is the range of k where the BAO live and where the power spectrum peaks. While we do not discuss the feasibility of measuring BAO at high redshift with LIM-nulling here, we direct the reader to Fronenberg et al. 2024[4]. In that work we find that it is indeed possible to constrain the BAO scale with a level of confidence comparable to other high-z probes. In addition, tomographic LIM-nulling can be used as a tool to constrain other features in the matter power spectrum (not just BAO), to place upper limits on the amplitude of the matter power spectrum as a function of redshift, and to constrain the matter transfer function.

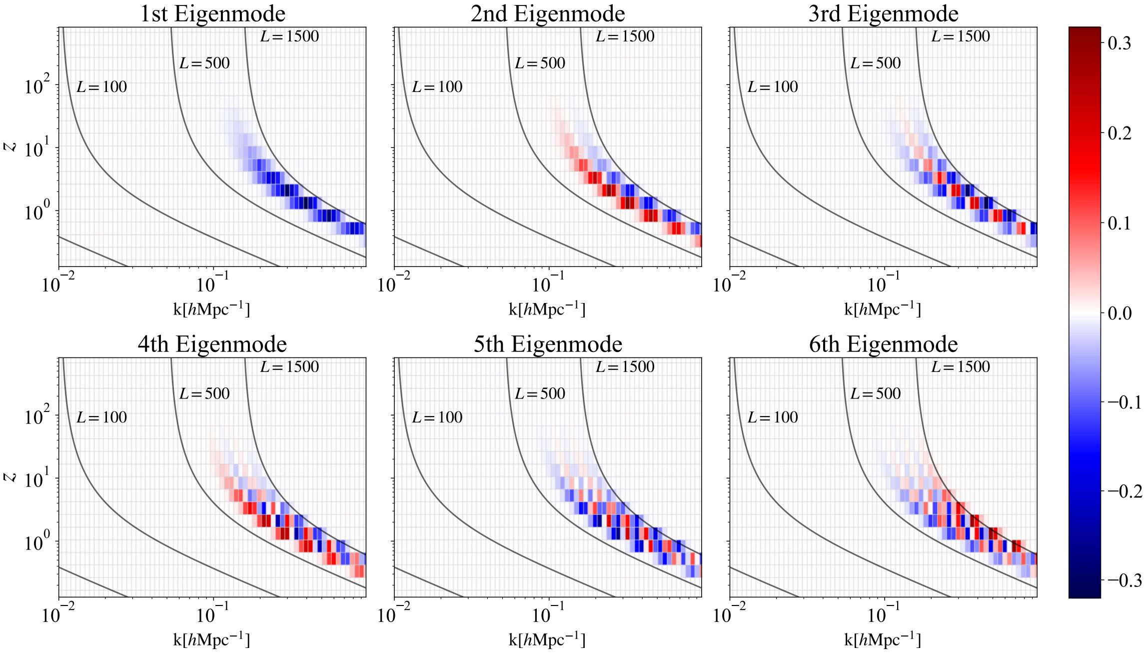

Figure 4.13: First six principal eigenmodes of the matter power spectrum $P_m(k,z)$ for regular CMB lensing measurements from CMBS4. Lines of constant L are shown in black to guide the eye, and reveal that one is essentially sensitive to an approximately fixed range of angular scales. As one goes to higher eigenmodes, one probes finer and finer features as a function of k. Since CMB lensing is sensitive to the integrated matter density from $z = 0$ to the surface of last scattering, there is broad support as a function of redshift.

---

[4] This is Chapter 5.

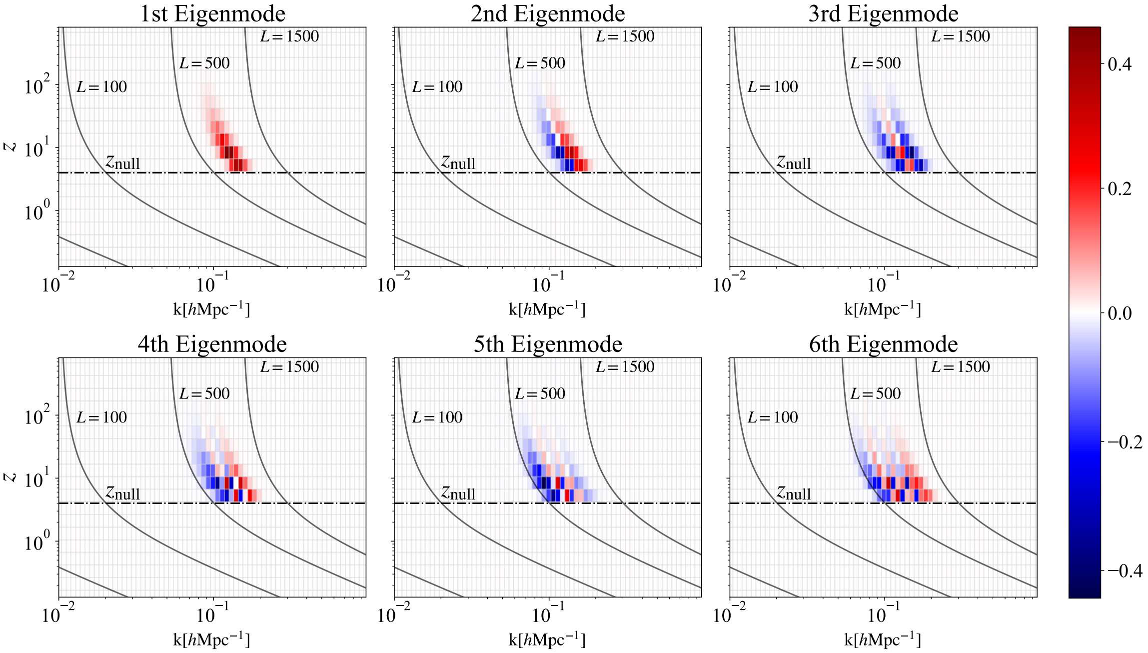

Figure 4.14: Same as Fig. 4.13, but for the CMB × LIM-nulling lensing measurements in the Futuristic scenario. By construction there is no sensitivity to redshifts below z_{null}. The range of angular scales is also shifted slightly because of different SNR characteristics of nulling versus regular lensing measurements. The slower evolution of the mapping between L and k at high redshifts hints at a potential BAO measurement with nulling estimators.

4.7 Conclusion

In this paper we have shown how the CMB × LIM-nulling convergence cross-spectrum (proposed in Ref. Maniyar et al. (2022)) can be used to constrain cosmology. This probe combines convergence maps of the CMB and of LIMs to exactly null out the low-redshift contribution to the CMB lensing convergence, leaving behind a direct and unbiased probe of the matter density field at high-redshift. This probe may serve to complement other high-redshift probes such as LIMs themselves, or high-redshift galaxy surveys, both of which are biased tracers of the matter density field.

Building upon the work in Ref. Maniyar et al. (2022), we have computed the variance of the CMB × LIM-nulling convergence cross-spectrum. This enabled the SNR parameter-space studies of Sections 4.3 and 4.4, which provided useful intuition for the expected sensitivity of nulling estimator measurements. Optimizing a set of experiments with some rough rules of thumb, we found that next-generation experiments may be able to make a detection.

Moving onto a set of futuristic and aspirational—but still potentially realizable—experiments, we explored the potential of nulling estimators to place constrains on standard ΛCDM parameters plus M_ν. These experiments can place constraints that are comparable to those from regular CMB lensing. However, parameters derived from nulling estimators will always be slightly worse by construction, since the low-redshift information captured by regular CMB lensing does has sensitivity to cosmological parameters.

The true benefit of nulling estimators, then, is not in raw statistical sensitivity. Instead, it is in one's ability to probe high redshifts cleanly. We showed that nulling estimators can be compared to traditional CMB lensing constraints, and that such comparisons can serve as model-independent tests of cosmology beyond ΛCDM. While we do not claim that these tests can solve outstanding problems in cosmology such as the Hubble tension and the σ_8 tension, we use these examples to illustrate how nulling estimators can probe cosmology at early times.

Additionally, we explicitly show that the CMB × LIM-nulling convergence probes high-redshift modes of the matter power spectrum, which can in turn be used to place limits on the matter power amplitude, the matter transfer function, and measure important features of the matter power spectrum such as BAOs. In our companion paper, Fronenberg et al. (2024), we forecast a BAO measurement with the CMB × LIM-nulling convergence and find it to be encouraging. Moreover, in all of our forecasts, we assume that one uses just one pair of LIM frequency channels to perform nulling. In future work, one can imagine taking advantage of the large bandwidth of LIM experiments, leveraging all frequency channels to perform LIM-nulling tomography. This would enable the direct study of the growth of structure.

While the prospects of using LIM lensing are promising, they do not come without serious challenges. Due to the complex astrophysical processes associated with the line emission and absorption mechanisms, LIMs are highly non-Gaussian and therefore using existing quadratic lensing estimators designed for the Gaussian CMB are sub-optimal Vafaei et al. (2010). It has been shown that attempting to use such estimators results in biases in the LIM lensing convergence Foreman et al. (2018); Schaan et al. (2018). These biases can be mitigated through various means, for instance, using filters that Gaussianize the field, or through bias-hardening, a method which makes use of our knowledge about the non-Gaussianity of LIMs in the estimation Vafaei et al. (2010); Foreman et al. (2018).

In addition, LIMs suffer from foreground contamination beyond the line interloper foregrounds included in this paper. For example, continuum foregrounds (such as the cosmic infrared background or galactic synchrotron emission) are expected to be present (although at a level that is generally considered to be less of a concern than line interlopers). One possible strategy for removing continuum foregrounds is to use their spectrally smooth nature to be removed. For example, Yue et al. (2015) show that the spectrally smooth far infrared (FIR) continuum foregrounds of [CII] can be removed with a negligible residual via spectral decomposition Yue et al. (2015). However, a fuller treatment ought to explicitly model these residuals in the context of nulling estimators.

In summary, the early epochs of the universe remain a treasure trove of cosmological information. LIM-nulling, in principle, has the potential to provide us with a clean window into the epoch of reioniziation, cosmic dawn, and even into the cosmic dark ages, allowing for unbiased measurements of the matter density field before the time of galaxy formation. While this endeavor presents a serious challenge, we have shown that the result offers a unique way to constrain cosmology beyond what is offered by existing probes, providing yet another pathway to the ultimate goal of understanding our Universe on all scales and at all redshifts.

Acknowledgements

The authors would like to thank Eiichiro Komatsu, Simon Foreman, Adélie Gorce and Giulio Fabbian for helpful discussion and comments. HF is supported by the Fonds de recherche du

Québec Nature et Technologies (FRQNT) Doctoral Research Scholarship award number 315907 and acknowledges support from the Mitacs Globalink Research Award for this work. AL and HF acknowledge support from the Trottier Space Institute, the New Frontiers in Research Fund Exploration grant program, the Canadian Institute for Advanced Research (CIFAR) Azrieli Global Scholars program, a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant and a Discovery Launch Supplement, the Sloan Research Fellowship, and the William Dawson Scholarship at McGill. ARP was supported by NASA under award numbers 80NSSC18K1014, NNH17ZDA001N, and 80NSSC22K0666, and by the NSF under award number 2108411. ARP was also supported by the Simons Foundation. This research was enabled in part by support provided by Compute Canada (www.computecanada.ca).

Appendix A: CMB × LIM-Nulling Variance

In Maniyar et al. (2022), it has been shown that the cross-correlation, $C_L^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_{\rm null}}$, is an unbiased estimator. Here, we compute the variance of this estimator in the presence of uncorrelated Gaussian noise at the convergence map level. We also assume, as in Maniyar et al. (2022), that the estimated convergence maps, $\hat{\kappa}_{\rm CMB}$ and $\hat{\kappa}_{\rm null}$, are Gaussian.

We begin with the standard formula for the variance, i.e.,

$$\mathrm{var}(\mathrm{C}_{\mathrm{L}}^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_{\rm null}}) = \langle(\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_{\rm null})^2\rangle - \langle\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_{\rm null}\rangle^2. \tag{4.27}$$

The first term can be written as

$$\begin{aligned}\langle(\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_{\rm null})^2\rangle &= \langle\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_{\rm null}\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_{\rm null}\rangle \\ &= \langle\hat{\kappa}_{\rm CMB}(\hat{\kappa}^*_{\rm CMB}+\alpha\hat{\kappa}^*_2-(1+\alpha)\hat{\kappa}^*_1)\hat{\kappa}_{\rm CMB}(\hat{\kappa}^*_{\rm CMB}+\alpha\hat{\kappa}^*_2-(1+\alpha)\hat{\kappa}^*_1)\rangle \\ &= \langle\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_{\rm CMB}\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_{\rm CMB}\rangle+2\alpha\langle\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_{\rm CMB}\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_2\rangle-2(1+\alpha)\langle\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_{\rm CMB}\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_1\rangle \\ &+\alpha^2\langle\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_2\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_2\rangle-2\alpha(1+\alpha)\langle\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_2\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_1\rangle+(1+\alpha)^2\langle\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_1\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_1\rangle,\end{aligned} \tag{4.28}$$

where here $\hat{\kappa}_1$ and $\hat{\kappa}_2$ are the Fourier transform of the estimated convergence maps from LIMs at comoving distances χ_1 and χ_2 where $\chi_1 < \chi_2$ and the asterisks denote the complex conjugate. Each term in this expression can be evaluated using the fourth-order moment relation

$$\langle x_1x_2x_3x_4\rangle = \langle x_1x_3\rangle\langle x_2x_4\rangle + \langle x_1x_2\rangle\langle x_3x_4\rangle + \langle x_1x_4\rangle\langle x_2x_3\rangle \tag{4.29}$$

where x_1, x_2, x_3, x_4 are Gaussian random variables with mean zero. Keeping in mind that the ith convergence map contains a cosmological signal, s_i, and is contaminated with uncorrelated Gaussian random noise with mean zero, n_i, we have $\kappa_i = s_i + n_i$. Eq. (4.28) then simplifies to

$$\begin{aligned}\langle(\hat{\kappa}_{\rm CMB}\hat{\kappa}^*_{\rm null})^2\rangle &= 3[C_L^{\hat{\kappa}_{\rm CMB}}+N_L^{\hat{\kappa}_{\rm CMB}}]^2+3[2\alpha(C_L^{\hat{\kappa}_{\rm CMB}}+N_L^{\hat{\kappa}_{\rm CMB}})C_L^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_2}] \\ &- 3[2(1+\alpha)(C_L^{\hat{\kappa}_{\rm CMB}}+N_L^{\hat{\kappa}_{\rm CMB}})C_L^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_1}]+\alpha^2[2(C_L^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_2})^2+(C_L^{\hat{\kappa}_{\rm CMB}}+N_L^{\hat{\kappa}_{\rm CMB}})(C_L^{\hat{\kappa}_2}+N_L^{\hat{\kappa}_2})] \\ &- 2\alpha(1+\alpha)[2C^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_1}C^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_2}+(C_L^{\hat{\kappa}_{\rm CMB}}+N_L^{\hat{\kappa}_{\rm CMB}})C^{\hat{\kappa}_1\hat{\kappa}_2}] \\ &+ (1+\alpha)^2[2(C_L^{\hat{\kappa}_{\rm CMB}\hat{\kappa}_1})^2+(C_L^{\hat{\kappa}_{\rm CMB}}+N_L^{\hat{\kappa}_{\rm CMB}})(C_L^{\hat{\kappa}_1}+N_L^{\hat{\kappa}_1})]\end{aligned} \tag{4.30}$$

where $C_L^{\hat{\kappa}_i} \equiv \langle s_is_i^*\rangle$, $C_L^{\hat{\kappa}_i\hat{\kappa}_j} \equiv \langle s_is_j^*\rangle$, and $N_L^{\hat{\kappa}_i} \equiv \langle n_in_i^*\rangle$. Similarly, using Eq. (4.29) the second term of Eq. (4.27) is

$$\begin{aligned}\langle\hat{\kappa}_{\rm CMB}\hat{\kappa}^{*}_{\rm null}\rangle^2 &= (C_L^{\hat{\kappa}{\rm CMB}}+N_L^{\hat{\kappa}{\rm CMB}})^2+2\alpha(C_L^{\hat{\kappa}{\rm CMB}}+N_L^{\hat{\kappa}{\rm CMB}})C_L^{\hat{\kappa}{\rm CMB}\hat{\kappa}_2}\\ &- 2(1+\alpha)(C_L^{\hat{\kappa}{\rm CMB}}+N_L^{\hat{\kappa}{\rm CMB}})C_L^{\hat{\kappa}{\rm CMB}\hat{\kappa}_1}+\alpha^2(C_L^{\hat{\kappa}{\rm CMB}\hat{\kappa}_2})^2\\ &-2\alpha(1+\alpha)C^{\hat{\kappa}{\rm CMB}\hat{\kappa}_1}C^{\hat{\kappa}{\rm CMB}\hat{\kappa}_2}+(1+\alpha)^2(C_L^{\hat{\kappa}{\rm CMB}\hat{\kappa}_1})^2. \end{aligned} \tag{4.31}$$

Plugging Eq. (4.30) and Eq. (4.31) back into (4.27), and accounting for cosmic variance we obtain

$$\begin{aligned}{\rm var}({\rm C}_{\rm L}^{\hat{\kappa}{\rm CMB}\hat{\kappa}_{\rm null}}) &= \left(\frac{1}{f_{\rm sky}(2L+1)}\right)2[(C_L^{\hat{\kappa}{\rm CMB}}+N_L^{\hat{\kappa}{\rm CMB}})^2+2\alpha(C_L^{\hat{\kappa}{\rm CMB}}+N_L^{\hat{\kappa}{\rm CMB}})C_L^{\hat{\kappa}{\rm CMB}\hat{\kappa}_2}\\ &- 2(1+\alpha)(C_L^{\hat{\kappa}{\rm CMB}}+N_L^{\hat{\kappa}{\rm CMB}})C_L^{\hat{\kappa}{\rm CMB}\hat{\kappa}_1}]+\alpha^2[(C_L^{\hat{\kappa}{\rm CMB}\hat{\kappa}_2})^2+(C_L^{\hat{\kappa}{\rm CMB}}+N_L^{\hat{\kappa}{\rm CMB}})(C_L^{\hat{\kappa}_2}+N_L^{\hat{\kappa}_2})]\\ &- 2\alpha(1+\alpha)[C^{\hat{\kappa}{\rm CMB}\hat{\kappa}_1}C^{\hat{\kappa}{\rm CMB}\hat{\kappa}_2}+(C_L^{\hat{\kappa}{\rm CMB}}+N_L^{\hat{\kappa}{\rm CMB}})C^{\hat{\kappa}_1\hat{\kappa}_2}]\\ &+ (1+\alpha)^2[(C_L^{\hat{\kappa}{\rm CMB}\hat{\kappa}_1})^2+(C_L^{\hat{\kappa}{\rm CMB}}+N_L^{\hat{\kappa}{\rm CMB}})(C_L^{\hat{\kappa}_1}+N_L^{\hat{\kappa}_1})], \end{aligned} \tag{4.32}$$

which is precisely Eq. (4.13).

Bibliography

A. S. Maniyar, E. Schaan, and A. R. Pullen, **105**, 083509 (2022), 2106.09005.

U. Seljak and M. Zaldarriaga, **82**, 2636 (1999), astro-ph/9810092.

A. Lewis and A. Challinor, Phys. Rev. D **76**, 083005 (2007), URL `https://link.aps.org/doi/10.1103/PhysRevD.76.083005`.

M. Schmittfull and U. Seljak, **97**, 123540 (2018), 1710.09465.

R. Allison, P. Caucal, E. Calabrese, J. Dunkley, and T. Louis, **92**, 123535 (2015), 1509.07471.

D. Huterer and M. White, **72**, 043002 (2005), astro-ph/0501451.

F. Bernardeau, T. Nishimichi, and A. Taruya, **445**, 1526 (2014), `1312.0430`.

A. Barthelemy, S. Codis, C. Uhlemann, F. Bernardeau, and R. Gavazzi, **492**, 3420 (2020), `1909.02615`.

F. McCarthy, S. Foreman, and A. van Engelen, **103**, 103538 (2021), `2011.06582`.

F. J. Qu, B. D. Sherwin, O. Darwish, T. Namikawa, and M. S. Madhavacheril, **107**, 123540 (2023), `2208.04253`.

Z. J. Zhang, Y. Omori, and C. Chang, **524**, 6392 (2023), `2211.09617`.

M. B. Silva, B. Baumschlager, K. A. Cleary, P. C. Breysse, D. T. Chung, H. T. Ihle, H. Padmanabhan, L. C. Keating, J. Kim, and L. Philip, *Synergies between the comap co line intensity mapping mission and a ly-alpha galaxy survey: How to probe the early universe with voxel based analysis of observational data* (2021), URL `https://arxiv.org/abs/2111.05354`.

G. Sun, L. Moncelsi, M. P. Viero, M. B. Silva, J. Bock, C. M. Bradford, T.-C. Chang, Y.-T. Cheng, A. R. Cooray, A. Crites, et al., The Astrophysical Journal **856**, 107 (2018), URL `https://doi.org/10.3847%2F1538-4357%2Faab3e3`.

E. Visbal, H. Trac, and A. Loeb, Journal of Cosmology and Astroparticle Physics **2011**, 010 (2011), URL `https://doi.org/10.1088%2F1475-7516%2F2011%2F08%2F010`.

P. C. Breysse, E. D. Kovetz, and M. Kamionkowski, *Masking line foregrounds in intensity mapping surveys* (2015), URL `https://arxiv.org/abs/1503.05202`.

Y. Gong, M. Silva, A. Cooray, and M. G. Santos, The Astrophysical Journal **785**, 72 (2014), URL `https://doi.org/10.1088%2F0004-637x%2F785%2F1%2F72`.

A. Liu, Y. Zhang, and A. R. Parsons, **833**, 242 (2016), `1609.04401`.

H. Fronenberg, A. S. Maniyar, A. Liu, and A. R. Pullen, Physical Review Letters **132** (2024), ISSN 1079-7114, URL `http://dx.doi.org/10.1103/PhysRevLett.132.241001`.

W. Hu and T. Okamoto, **574**, 566 (2002), `astro-ph/0111606`.

A. S. Maniyar, Y. Ali-Haïmoud, J. Carron, A. Lewis, and M. S. Madhavacheril, Phys. Rev. D **103**, 083524 (2021), URL `https://link.aps.org/doi/10.1103/PhysRevD.103.083524`.

P. Lemos, A. Challinor, and G. Efstathiou, Journal of Cosmology and Astroparticle Physics **2017**, 014–014 (2017), ISSN 1475-7516, URL `http://dx.doi.org/10.1088/1475-7516/2017/05/014`.

K. M. Smith, O. Zahn, and O. Doré, **76**, 043510 (2007), `0705.3980`.

S. Das, B. D. Sherwin, P. Aguirre, J. W. Appel, J. R. Bond, C. S. Carvalho, M. J. Devlin, J. Dunkley, R. Dünner, T. Essinger-Hileman, et al., **107**, 021301 (2011), `1103.2124`.

B. D. Sherwin, A. van Engelen, N. Sehgal, M. Madhavacheril, G. E. Addison, S. Aiola, R. Allison, N. Battaglia, D. T. Becker, J. A. Beall, et al., Phys. Rev. D **95**, 123529 (2017), URL `https://link.aps.org/doi/10.1103/PhysRevD.95.123529`.

S. Shaikh, I. Harrison, A. van Engelen, G. A. Marques, T. M. C. Abbott, M. Aguena, O. Alves, A. Amon, R. An, D. Bacon, et al., arXiv e-prints arXiv:2309.04412 (2023), `2309.04412`.

A. van Engelen, R. Keisler, O. Zahn, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, T. M. Crawford, et al., **756**, 142 (2012), `1202.0546`.

D. Hanson, S. Hoover, A. Crites, P. A. R. Ade, K. A. Aird, J. E. Austermann, J. A. Beall, A. N. Bender, B. A. Benson, L. E. Bleem, et al., **111**, 141301 (2013), `1307.5830`.

Y. Omori, R. Chown, G. Simard, K. T. Story, K. Aylor, E. J. Baxter, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, et al., **849**, 124 (2017), `1705.00743`.

K. T. Story, D. Hanson, P. A. R. Ade, K. A. Aird, J. E. Austermann, J. A. Beall, A. N. Bender, B. A. Benson, L. E. Bleem, J. E. Carlstrom, et al., **810**, 50 (2015), 1412.4760.

W. L. K. Wu, L. M. Mocanu, P. A. R. Ade, A. J. Anderson, J. E. Austermann, J. S. Avva, J. A. Beall, A. N. Bender, B. A. Benson, F. Bianchini, et al., **884**, 70 (2019), 1905.05777.

Z. Pan et al. (SPT) (2023), 2308.11608.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, et al., **594**, A15 (2016), 1502.01591.

Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, et al., **641**, A8 (2020), 1807.06210.

BICEP2 Collaboration, Keck Array Collaboration, P. A. R. Ade, Z. Ahmed, R. W. Aikin, K. D. Alexander, D. Barkats, S. J. Benton, C. A. Bischoff, J. J. Bock, et al., **833**, 228 (2016), 1606.01968.

P. A. R. Ade, Y. Akiba, A. E. Anthony, K. Arnold, M. Atlas, D. Barron, D. Boettger, J. Borrill, S. Chapman, Y. Chinone, et al., **113**, 021301 (2014), 1312.6646.

S. Adachi, M. A. O. Aguilar Faúndez, Y. Akiba, A. Ali, K. Arnold, C. Baccigalupi, D. Barron, D. Beck, F. Bianchini, J. Borrill, et al. (POLARBEAR Collaboration), Phys. Rev. Lett. **124**, 131301 (2020), URL `https://link.aps.org/doi/10.1103/PhysRevLett.124.131301`.

J. Ward and N. S. T. R. F. Advanced ACT Collaboration, in *American Astronomical Society Meeting Abstracts #230* (SPIE, 2017), vol. 230 of *American Astronomical Society Meeting Abstracts*, p. 117.01.

B. A. Benson, P. A. R. Ade, Z. Ahmed, S. W. Allen, K. Arnold, J. E. Austermann, A. N. Bender, L. E. Bleem, J. E. Carlstrom, C. L. Chang, et al., in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, edited by W. S. Holland and

J. Zmuidzinas (SPIE, 2014), vol. 9153 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 91531P, `1407.2973`.

A. J. Anderson, P. Barry, A. N. Bender, B. A. Benson, L. E. Bleem, J. E. Carlstrom, T. W. Cecil, C. L. Chang, T. M. Crawford, K. R. Dibert, et al., arXiv e-prints arXiv:2208.08559 (2022), `2208.08559`.

P. Ade, J. Aguirre, Z. Ahmed, S. Aiola, A. Ali, D. Alonso, M. A. Alvarez, K. Arnold, P. Ashton, J. Austermann, et al., **2019**, 056 (2019), `1808.07445`.

K. Abazajian, G. E. Addison, P. Adshead, Z. Ahmed, D. Akerib, A. Ali, S. W. Allen, D. Alonso, M. Alvarez, M. A. Amin, et al., **926**, 54 (2022), `2008.12619`.

E. Schaan and M. White, Journal of Cosmology and Astroparticle Physics **2021**, 068 (2021a), URL `https://doi.org/10.1088%2F1475-7516%2F2021%2F05%2F068`.

E. Schaan and M. White, Journal of Cosmology and Astroparticle Physics **2021**, 067 (2021b), URL `https://doi.org/10.1088%2F1475-7516%2F2021%2F05%2F067`.

T. Y. Li, R. H. Wechsler, K. Devaraj, and S. E. Church, The Astrophysical Journal **817**, 169 (2016), URL `https://doi.org/10.3847%2F0004-637x%2F817%2F2%2F169`.

H. T. Ihle, D. Chung, G. Stein, M. Alvarez, J. R. Bond, P. C. Breysse, K. A. Cleary, H. K. Eriksen, M. K. Foss, J. O. Gundersen, et al., **871**, 75 (2019), `1808.07487`.

K. A. Cleary, J. Borowska, P. C. Breysse, M. Catha, D. T. Chung, S. E. Church, C. Dickinson, H. K. Eriksen, M. K. Foss, J. O. Gundersen, et al., The Astrophysical Journal **933**, 182 (2022), URL `https://doi.org/10.3847%2F1538-4357%2Fac63cc`.

P. C. Breysse, D. T. Chung, K. A. Cleary, H. T. Ihle, H. Padmanabhan, M. B. Silva, J. R. Bond, J. Borowska, M. Catha, S. E. Church, et al., The Astrophysical Journal **933**, 188 (2022), URL `https://doi.org/10.3847%2F1538-4357%2Fac63c9`.

D. T. Chung, P. C. Breysse, K. A. Cleary, H. T. Ihle, H. Padmanabhan, M. B. Silva, J. R. Bond, J. Borowska, M. Catha, S. E. Church, et al., The Astrophysical Journal **933**, 186 (2022), URL `https://doi.org/10.3847%2F1538-4357%2Fac63c7`.

D. T. Chung, M. P. Viero, S. E. Church, and R. H. Wechsler, **892**, 51 (2020), `1812.08135`.

Concerto Collaboration, P. Ade, M. Aravena, E. Barria, A. Beelen, A. Benoit, M. Béthermin, J. Bounmy, O. Bourrion, G. Bres, et al., **642**, A60 (2020).

M. Silva, M. G. Santos, A. Cooray, and Y. Gong, **806**, 209 (2015), `1410.4808`.

A. Cooray, T.-C. Chang, S. Unwin, M. Zemcov, A. Coffey, P. Morrissey, N. Raouf, S. Lipscy, M. Shannon, G. Wu, et al., *Cdim: Cosmic dawn intensity mapper final report* (2019), `1903.03144`.

O. Doré, J. Bock, M. Ashby, P. Capak, A. Cooray, R. de Putter, T. Eifler, N. Flagey, Y. Gong, S. Habib, et al., *Cosmology with the spherex all-sky spectral survey* (2015), `1412.4872`.

M. Tegmark, Physical Review D **56**, 4514–4529 (1997), ISSN 1550-7998.

N. Sailer, E. Castorina, S. Ferraro, and M. White, **2021**, 049 (2021), `2106.09713`.

E. D. Valentino, O. Mena, S. Pan, L. Visinelli, W. Yang, A. Melchiorri, D. F. Mota, A. G. Riess, and J. Silk, Classical and Quantum Gravity **38**, 153001 (2021), URL `https://doi.org/10.1088%2F1361-6382%2Fac086d`.

V. Poulin, J. L. Bernal, E. D. Kovetz, and M. Kamionkowski, **107**, 123538 (2023), `2209.06217`.

M. Douspis, L. Salvati, and N. Aghanim, *On the tension between large scale structures and cosmic microwave background* (2019), `1901.05289`.

D. Coe, *Fisher matrices and confidence ellipses: A quick-start guide and software* (2009), `0906.4123`, URL `https://arxiv.org/abs/0906.4123`.

D. Huterer and M. Takada, Astroparticle Physics **23**, 369 (2005), astro-ph/0412142.

Planck Collaboration, **594**, A13 (2016).

M. Benetti, W. Miranda, H. Borges, C. Pigozzo, S. Carneiro, and J. Alcaniz, Journal of Cosmology and Astroparticle Physics **2019**, 023 (2019), URL https://doi.org/10.1088%2F1475-7516%2F2019%2F12%2F023.

M. Benetti, H. Borges, C. Pigozzo, S. Carneiro, and J. Alcaniz, Journal of Cosmology and Astroparticle Physics **2021**, 014 (2021), URL https://doi.org/10.1088%2F1475-7516%2F2021%2F08%2F014.

W. Handley and P. Lemos, Phys. Rev. D **100**, 023512 (2019), URL https://link.aps.org/doi/10.1103/PhysRevD.100.023512.

Y. Park and E. Rozo, Monthly Notices of the Royal Astronomical Society **499**, 4638–4645 (2020), ISSN 1365-2966, URL http://dx.doi.org/10.1093/mnras/staa2647.

P. Marshall, N. Rajguru, and A. Slosar, Physical Review D **73** (2006), ISSN 1550-2368, URL http://dx.doi.org/10.1103/PhysRevD.73.067302.

P. Lemos, M. Raveri, A. Campos, Y. Park, C. Chang, N. Weaverdyck, D. Huterer, A. R. Liddle, J. Blazek, R. Cawthon, et al., Monthly Notices of the Royal Astronomical Society **505**, 6179–6194 (2021), ISSN 1365-2966, URL http://dx.doi.org/10.1093/mnras/stab1670.

S. Mohanty, S. Anand, P. Chaubal, A. Mazumdar, and P. Parashari, J. Astrophys. Astron. **39**, 46 (2018).

S. Vafaei, T. Lu, L. van Waerbeke, E. Semboloni, C. Heymans, and U.-L. Pen, Astroparticle Physics **32**, 340 (2010), URL https://doi.org/10.1016%2Fj.astropartphys.2009.10.003.

S. Foreman, P. D. Meerburg, A. van Engelen, and J. Meyers, **2018**, 046 (2018), 1803.04975.

E. Schaan, S. Ferraro, and D. N. Spergel, Phys. Rev. D **97**, 123539 (2018), URL `https://link.aps.org/doi/10.1103/PhysRevD.97.123539`.

B. Yue, A. Ferrara, A. Pallottini, S. Gallerani, and L. Vallini, Monthly Notices of the Royal Astronomical Society **450**, 3829 (2015), URL `https://doi.org/10.1093%2Fmnras%2Fstv933`.

Chapter 5

A New Probe of the High-z BAO scale: BAO tomography With CMB × LIM-Nulling Convergence

Hannah Fronenberg[1,2], Abhishek S. Maniyar[3,4,5], Adrian Liu[1,2], Anthony R. Pullen[6,7]

[1]Department of Physics, McGill University, Montréal, QC, Canada

[2]Trottier Space Institute, Montréal, QC, Canada

[3]Department of Physics, Stanford University, Stanford, CA, USA

[4] SLAC National Accelerator Laboratory, Menlo Park, CA, USA

[5]Kavli Institute for Particle Astrophysics and Cosmology, Stanford, CA, USA

[6]Department of Physics, New York University, New York, NY, USA

[7]Center for Computational Astrophysics, Flatiron Institute, New York, NY, USA

Abstract

Standard rulers such as the baryon acoustic oscillation (BAO) scale serve as workhorses for precision tests of cosmology, enabling distance measurements that probe the geometry and expansion history of our Universe. Aside from BAO measurements from the cosmic microwave background (CMB), most standard ruler techniques operate at relatively low redshifts and depend on biased tracers of the matter density field. In a companion paper, we explored the scientific reach of nulling estimators, where CMB lensing convergence maps are cross-correlated with linear combinations of similar maps from line intensity mapping (LIM) to precisely null out the low-redshift contributions to CMB lensing. We showed that nulling estimators can be used to constrain the high redshift matter power spectrum and showed that this spectrum exhibits discernible BAO features. Here we propose using these features as a standard ruler at high redshifts that does not rely on biased tracers. Forecasting such a measurement at $z \sim 5$, we find that next-generation instruments will be able to constrain the BAO scale to 7.2% precision, while our futuristic observing scenario can constrain the BAO scale to 4% precision. This constitutes a fundamentally new kind of BAO measurement during early epochs in our cosmic history.

5.1 Introduction

At the time of recombination ($z \sim 1100$), the first atoms formed and photons streamed freely through the universe. Today, we observe those photons as the cosmic microwave background (CMB), a map which provides unparalleled insight into the early moments at the surface of last scattering. These photons, however, have not traveled unimpeded. Weak gravitational lensing of the CMB arises when photons from the surface of last scattering are deflected by the gravitational potentials they encounter on their way to us, resulting in distortions to the statistics of the CMB. With the use of lensing estimators, one can reconstruct the gravitational potential of the projected mass distribution along the line of sight (LOS) (Hu and Okamoto, 2002; Okamoto and Hu, 2003).

Reconstructing the lensing potential, or equivalently the convergence, yields direct constraints on the total matter distribution of the universe, both baryonic and dark, without the use of a biased tracer. As such, the CMB convergence has the ability to probe the growth of matter fluctuations, place limits on primordial non-Gaussianity, constrain the sum of the neutrino masses, and even test theories of modified gravity (Seljak and Zaldarriaga, 1999; Lewis and Challinor, 2007; Schmittfull and Seljak, 2018; Allison et al., 2015).

This information, however, is projected onto a single observable and the high-redshift contribution to the convergence is dwarfed by that of the low-redshift universe ($z \lesssim 2$). This places limitations on the inferences that can be made about the matter density field. Most notably, the fact that the convergence is a LOS integrated quantity results in the washing out of baryon acoustic oscillations (BAOs). During the radiation–dominated era, dark matter began to cluster while the photon–baryon fluid continued to oscillate, producing BAOs which were then left imprinted on the CMB at recombination. Through large scale structure formation, the BAO scale remains embossed in the distribution of galaxies, and provides a standard distance measure across cosmic time.

Luckily, line intensity maps (LIMs) also experience weak lensing by large scale structure. These lines, however, are only lensed by a portion of large scale structure that lenses the CMB, namely the low redshift universe. In Maniyar et al. (2022), it has been shown that by using the lensing information of two LIMs, one could not just suppress, but exactly null out the low redshift contribution to the CMB convergence. This nulling method has been explored in the context of galaxy lensing (Huterer and White, 2005; Bernardeau et al., 2014; Barthelemy et al., 2020) and has been shown to be able to remove the imprint of various effects from CMB lensing maps for which the physics is uncertain (McCarthy et al., 2021; Qu et al., 2023; Zhang et al., 2023). While never implemented with real data, these techniques can be an important new tool for studying the high redshift universe. What is more, Maniyar et al. (2022) show that the CMB × LIM-nulling convergence spectrum, $\langle \hat{\kappa} \hat{\kappa}_{\rm null} \rangle$, is free of LIM interloper bias. Most recently, we showed in Fronenberg et al. (2024) that the CMB × LIM-nulling convergence can be compared to traditional CMB lensing constraints, and that such comparisons can serve as model-independent tests of cosmology beyond ΛCDM. Additionally, we explicitly showed that the CMB × LIM-nulling convergence probes high-redshift modes of the matter power spectrum, which can in turn be used to place limits on the matter power amplitude, the matter transfer function, and measure important features of the matter power spectrum. These results suggest that performing LIM-nulling of CMB lensing observations has the potential to probe a vast amount of high-redshift information.

In this *Letter*, we propose a new method to detect BAOs using the CMB × LIM-nulling convergence spectrum which can be used as a standard ruler over a large cosmological window ($z > 2$). This method, in principle, allows one to directly probe the matter density field during cosmic noon, the epoch of reionization (EoR), cosmic dawn, and even during the cosmic dark ages. The procedure to measure BAOs in this cross-convergence spectrum is similar to that of using the matter power spectrum or the galaxy correlation function; however, this probe has unique benefits. In Fig. 5.1, the window in which CMB × LIM-nulling can be used for studying the high-z universe is shaded in grey. LIM surveys typically have fine spectral resolution and are conducted over a large bandwidth. This therefore allows for direct, large scale studies of the matter density field over a vast and relatively unexplored cosmic window, filling a void between BAO measurements from $z < 3$ and high-z CMB measurements (Beutler et al., 2011; Ross et al., 2015; Alam et al., 2017; Kazin et al., 2014; Abbott et al., 2018a; Bautista et al., 2017; Ata et al., 2017; du Mas des Bourboux et al., 2017; Abbott et al., 2018b).

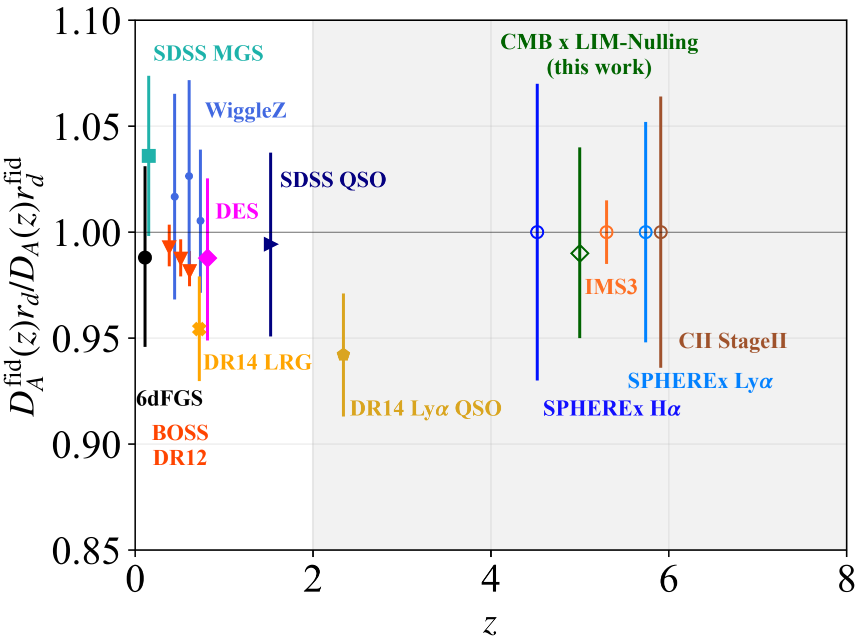

Figure 5.1: Current and projected BAO measurements as a function of redshift. Current measurements from galaxy, quasar, and Lyman-α surveys are shown at $z < 3$ using solid markers (Beutler et al., 2011; Ross et al., 2015; Alam et al., 2017; Kazin et al., 2014; Abbott et al., 2018a; Bautista et al., 2017; Ata et al., 2017; du Mas des Bourboux et al., 2017; Abbott et al., 2018b). Forecast constraints are shown with a hollow marker. The dark green diamond shows the projected measurement from CMB × LIM-nulling while the remaining constraints are from high-z LIM surveys (Bernal et al., 2019). The grey shaded region denotes the redshifts where LIM-nulling can be performed for high-z studies ($2 < z \lesssim 14$).

5.2 CMB convergence and LIM-nulling

The CMB acts as a source image which is lensed by the intervening matter density field. The deflection of CMB photons produces correlations between the otherwise uncorrelated CMB spherical harmonic coefficients, $a_{\ell m}$. Just like the CMB, lower redshift LIMs also incur such correlation as a result of lensing and, most importantly, LIMs and the CMB share common low redshift induced correlations. With the use of quadratic estimators (or more advanced techniques Hirata and Seljak (2003a,b)), the lensing convergence can be estimated. This estimated quantity is related to the total matter density integrated along the line of sight. The convergence κ is given by

$$\kappa(\hat{n}) = \int_0^{z_\mathrm{s}} W(z', z_\mathrm{s}) \delta_m(\chi(z')\hat{n}, z') \frac{c\, dz'}{H(z')} \tag{5.1}$$

where z_s is the redshift of the source, $W(z, z_s)$ is the lensing kernel, $H(z)$ is the Hubble parameter, c is the speed of light, χ denotes the comoving distance, and $\delta_m(\mathbf{r}, z)$ is the matter overdensity field. The lensing kernel for a source at a single comoving slice is given by

$$W(z, z_s) = \frac{3}{2}\left(\frac{H_0}{c}\right)^2 \frac{\Omega_{m,0}}{a} \chi(z) \left(1 - \frac{\chi(z)}{\chi(z_s)}\right), \tag{5.2}$$

where H_0 is the Hubble constant, $\Omega_{m,0}$ is the matter fraction today, a is the scale factor, and c is the speed of light, and $\chi(z_s)$ is the comoving distance to the source. One can then compute the angular power spectrum $C_L^{\kappa_i\kappa_j}$ of any two convergence maps κ_i and κ_j, given by

$$C_L^{\kappa_i\kappa_j} = \int_0^{z_s} \frac{W_i(z', z_\mathrm{s}) W_j(z', z_\mathrm{s})}{\chi(z')^2} P_\mathrm{m}\left(k = \frac{L + 1/2}{\chi(z')}, z'\right) \frac{c\, dz'}{H(z')}, \tag{5.3}$$

where P_m is the matter power spectrum in the Limber approximation, and i and j index the maps. These maps might, for example, be constructed from the linear combination of multiple convergence maps from different probes.

From Eq. (5.1), it should be clear that it is possible to construct some $W_j(z, z_\mathrm{s})$ that vanishes over the low redshift interval $[0, z_\mathrm{null}]$. Since W is quadratic in χ, a linear combination of three such kernels suffices to find a non-trivial null solution for the coefficients of this polynomial. As shown in Maniyar et al. (2022), using convergence maps estimated from two LIMs and from the CMB at redshifts $z_1 < z_2 < z_\mathrm{CMB}$ respectively, the LIM-nulling kernel is given by

$$W_\mathrm{null} = W(z, z_\mathrm{CMB}) + \alpha W(z, z_2) - (1 + \alpha) W(z, z_1), \tag{5.4}$$

where $\alpha \equiv [1/\chi(z_{\text{CMB}}) - 1/\chi(z_1)]/[1/\chi(z_1) - 1/\chi(z_2)]$.

Fig. 5.2 shows the relevant estimated convergence spectra. The CMB convergence spectrum, $C_L^{\hat{\kappa}\hat{\kappa}}$, is smooth with no discernible BAO features which is due to the angular evolution of the BAO wiggles which is depicted in the bottom panel. Examining the integrand of Eq. 5.3 evaluated at several redshifts, one can see that the BAO wiggles evolve to lower L as z decreases which, when integrated over, results in the washing out of BAO wiggles in $C_L^{\hat{\kappa}\hat{\kappa}}$. While $C_L^{\hat{\kappa}\hat{\kappa}}$ can provide information about the scale of structures which dominate the deflection of CMB photons, it cannot act as a standard ruler.

In the CMB × LIM-nulled convergence spectrum, $C_L^{\hat{\kappa}\hat{\kappa}_{\text{null}}}$, acoustic peaks emerge which is apparent in the top panel of Fig. 5.2 where the fractional difference between the wiggle and no-wiggle spectra are plotted. This BAO feature is the result of the much slower angular evolution of the BAO scale at early times. Again, referring now to the bottom panel of Fig. 5.2, the pale high-z curves share largely overlapping acoustic peaks compared to the darker low-z curves whose peaks and troughs mix.

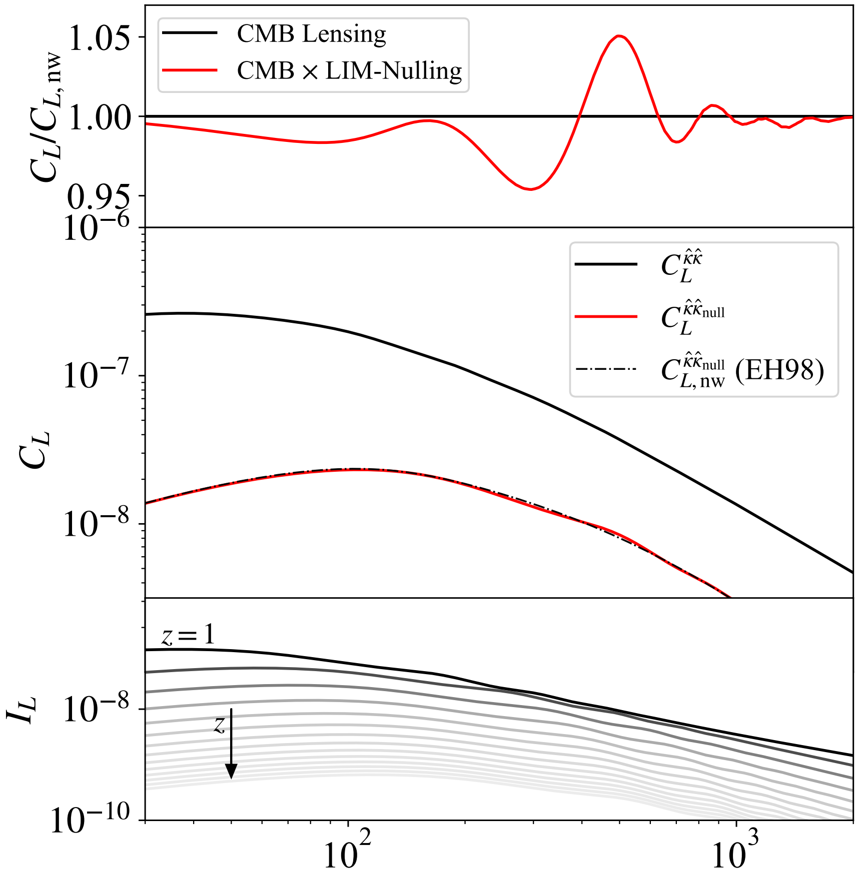

Figure 5.2: Top: Fractional difference between convergence spectra computed with the linear power spectrum and with the no-wiggle Eisenstein & Hu fitting function for regular CMB lensing (black) and CMB × LIM-nulling (red). Middle: The CMB convergence spectrum $C_L^{\hat\kappa\hat\kappa}$ (solid black) and the CMB × LIM-nulling convergence spectrum $C_L^{\hat\kappa\hat\kappa\mathrm{null}}$ (solid red). The dot-dashed black line corresponds to the CMB × LIM-nulling convergence spectrum computed with the no-wiggle Eisenstein & Hu fitting function. Bottom: The plot shows the integrand of Eq. (5.3) I_L at increasing redshifts from top (z = 1) to bottom.

What we argue is that the acoustic peak in the CMB × LIM-nulling spectrum is, to good approximation, a proxy for measuring the BAO scale at $z = z_{\mathrm{null}}$. Since the nulling kernel is sharply peaked near z_{null} and the matter transfer function is monotonically increasing, the matter density field is weighted most heavily near this redshift. Therefore, the location of the BAO wiggles in the CMB × LIM-nulling spectrum can be used to measure the BAO scale at $z \sim z_{\rm null}$. To test this hypothesis, we perform an Alcock-Paczynski (AP) test on a mock data set in order to place constraints on the BAO scale which we describe in the following section.

5.3 BAO Model and AP Test

To model the CMB × LIM-nulling convergence spectrum, one can write down the spectrum using Eq. (5.3) with i = CMB and j =LIM-nulling, and parameterise the matter power spectrum using $P_{\rm model}(k, z) = P_{\rm nw}(k, z) + AP_{\rm BAO}(k' = \omega k, z)$ where the BAO wiggles are independently parameterised by A and ω. These parameters control the amplitude and location of the wiggles respectively. The BAO spectrum is given by $P_{\rm BAO} = P_{\rm lin} - P_{\rm nw}$ where $P_{\rm lin}$ and $P_{\rm nw}$ are the Eisentsein & Hu linear and no-wiggle fitting functions computed using the publicly available code `nbodykit` (Eisenstein and Hu, 1998; Hand et al., 2018). Using this power spectrum model, we fit our two BAO parameters, A and ω. The parameter A is the amplitude of the BAO wiggles in the matter power spectrum and ω stretches the position of the wiggles as a function of wave-number k. The parameter ω is of particular interest since a change in the configuration space BAO scale is captured by our dilatation parameter ω.

We perform an AP test on mock CMB × LIM-nulling convergence spectra in order to make use of the BAO wiggle as a standard ruler. Such a test constitutes altering the location of features at some wave number k to k/α. Typically the AP parameter α is decomposed into LOS, $k_{||}/\alpha_{||}$, and perpendicular modes, $k_\perp/\alpha_\perp$, and these are related to the following physical parameters via the relations

$$\alpha_\perp = \frac{D_A^{\rm fid}(z) r_d}{D_A(z) r_d^{\rm fid}} \tag{5.5}$$

$$\alpha_{||} = \frac{H^{\rm fid}(z) r_d^{\rm fid}}{H(z) r_d} \tag{5.6}$$

where H is the Hubble parameter, D_A is the angular diameter distance, r_d is the acoustic scale, and the superscript "fid" denotes the value of the parameter in the fiducial cosmology.

When nulling is performed with a single pair of LIMs, however, we are only sensitive to the perpendicular AP parameter. This is because $C_L^{\hat{\kappa}\hat{\kappa}_{\rm null}}$ is a LOS integrated quantity. Thus, instead of fitting for both AP parameters, we focus on $\alpha_\perp$ since this is precisely our parameter ω. While in this work we focus on a single LIM-nulling pair, it is in principle possible to perform LIM-nulling tomography at several different redshifts, using pairs of frequency channels from two LIM experiments, to access the LOS information and $\alpha_{||}$. Given the high spectral resolution of current and upcoming LIM surveys, which resolve the BAO scale in the LOS direction, one can perform nulling at several redshifts, subtract the nulled convergence maps, and build a CMB convergence cube instead of just a single map. We leave the full 3D convergence estimation and analysis for future work.

5.4 Mock Data and MCMC Set-up

We simulate nulling performed with Lyman-α (Ly-α) and ionized carbon ([CII]) observations from $z = 3$ and $z = 5$ respectively. Line interlopers (where other spectral lines may redshift into one's observational bands) can be a serious concern for LIMs. Here we include H-α and CO (J = 4–3) as line interlopers for Ly-α and [CII] respectively. We simulate the auto- and cross- spectra of these lines using the publicly available code `Halogen`[1] which uses a halo model formalism based on conditional luminosity functions (Schaan and White, 2021a,b).

We define Next Generation and Futuristic observing scenarios. For the former we simulate noise from Cosmic Dawn Intensity Mapper (CDIM) for Ly-α observations, the Stage II instrument for [CII] observations, and the Simons Observatory (SO) for CMB observations, over a 100 deg^2

---

[1] https://github.com/EmmanuelSchaan/HaloGen/tree/LIM

field (the nominal survey area for these LIM experiments) (Fronenberg et al., 2024; Cooray et al., 2019; Silva et al., 2015; Ade et al., 2019). We assume interloper residuals to be at the 5% level and we perform LIM lensing reconstruction using $\ell_{\rm LIM,max} = 10,000$. For the Futurisic scenario, we model noise from CDIM, Stage II and from CMB Stage-4 (CMBS4), over one quarter of the sky ($f_{\rm sky} = 0.25$) Abazajian et al. (2022). In this scenario we assume 1% interloper residuals and perform LIM lensing reconstruction with $\ell_{\rm LIM,max} = 20,000$. For the SO noise, we use $N_\ell^{\hat{\kappa}{\rm CMB}}$ from the SO noise calculator[2] and for CMBS4 the lensing noise is obtained from the CMBS4 Wiki (Sailer et al., 2021). Throughout, our fiducial cosmology is that of *Planck* 2015 (Planck Collaboration, 2016). Additional details can be found in Section IV of Fronenberg et al. (2024).

Using Eq. (5.3), our power spectrum model, and the line and instrument models discussed above, we generate mock $C_L^{\hat{\kappa}\hat{\kappa}_{\rm null}}$ data and draw Gaussian random noise consistent with the CMB × LIM-nulling reconstruction variance, $\mathrm{var}_{\rm L}^{\hat{\kappa}\hat{\kappa}_{\rm null}}$, which we derive in Appendix A of Fronenberg et al. (2024). In both cases, the data samples multipoles $L_{\rm min} = 30$ to $L_{\rm max} = 1500$. We define a Gaussian likelihood $\mathcal{L}$

$$\ln \mathcal{L}(\lambda_i) = -\frac{1}{2}\sum_L \frac{(C_L^{\rm data} - C_L^{\rm model}(\lambda_i))^2}{\mathrm{var}_{\rm L}^{\rm data}} \tag{5.7}$$

where $C_L^{\rm data}$ is the data set used to constrain the model $C_L^{\rm model}(\lambda_i)$ with parameters $\lambda_i = \{A, \alpha_\perp\}$ in our case. We sample the likelihood using the Python package `emcee` (Foreman-Mackey et al., 2013) and impose a prior on A that it be non-negative, $A \geq 0$. Given that the BAO scale has been measured to the percent level by galaxy and quasar surveys and to the sub-percent level by *Planck*, we place a Gaussian prior on $\alpha_\perp$ with 10% error which is consistent with current observations yet conservative.

---

[2] `https://github.com/simonsobs/so_noise_models/tree/master/LAT_lensing_noise`

5.5 Results

We present the forecast sensitivities on the model parameters A and $\alpha_\perp$ in Table 5.1 and we show the one- and two-dimensional posterior distributions for these parameters in Fig. 5.3. In both the Next Generation and Futuristic observing scenarios a BAO detection is possible albeit with relatively low signal-to-noise. The parameter A which characterises the amplitude and therefore the existence of BAO features in the spectrum is weakly constrained in both scenarios, but nonetheless rules out a featureless spectrum at over 3.6σ in the optimistic case 1.7σ in the moderate case. For the AP parameter $\alpha_\perp$, the situation is more encouraging. In the Next Generation scenario, the BAO scale can be measured with to a precision of 7.2%. In the Futuristic scenario, things are even more promising. It is possible with future generation experiments to measure the BAO scale with 4% precision solely using CMB lensing information. These constitute BAO scale measurements at $z \sim 5$. While there is still a small amount of decorrelation of BAO wiggles due to the evolution of the comoving distance at $z > 5$, our tests (which involved artificially aligning the BAO features in the bottom panel of Fig. 5.2) indicate that this is a negligible effect. This test also indicates that any smearing to the effective redshift of this measurement due to the extended kernel is also negligible.

	A		$\alpha_\perp$	
	best fit	$\sigma_{\rm rel}$%	best fit	$\sigma_{\rm rel}$%
Next Gen.	$5.12^{+3.05}_{-3.09}$	59.6	$0.97^{+0.08}_{-0.07}$	7.2
Futuristic	$1.67^{+0.48}_{-0.46}$	27.5	$0.99^{+0.05}_{-0.04}$	4.0

Table 5.1: Posterior on the BAO model parameters for Next Generation and Futuristic cases with 68% credibility error bars and corresponding relative percent errors.

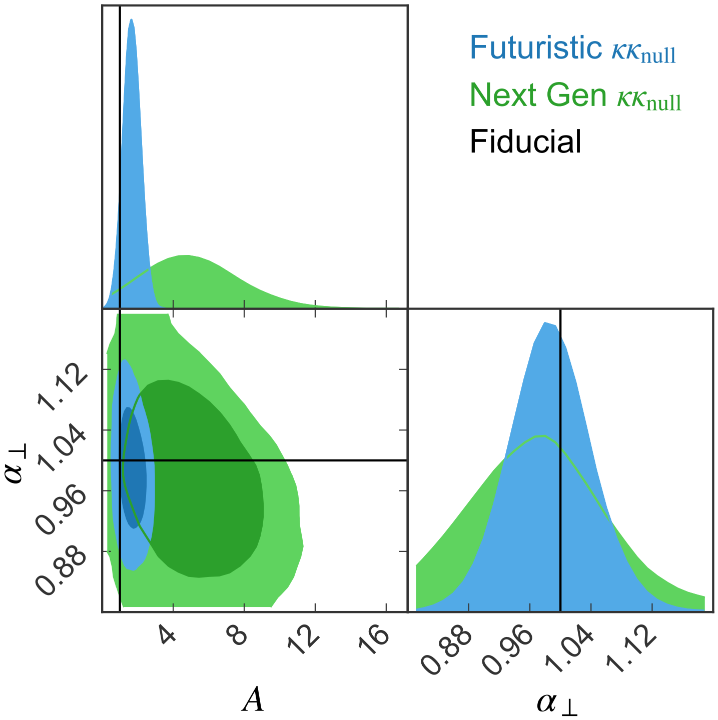

Figure 5.3: Projected 68% (darker colour) and 95% (lighter colour) credibility contours for the BAO model parameters A and $\alpha_\perp$ for Next Generation (green) and Futuristic (blue) scenarios. The black line denotes the fiducial values.

To put these results in context, we comment on the current status of BAO measurements from high-redshift LIM surveys. While there are currently no BAO detections at $3 < z < 1100$, there have been a number of forecasts around $z \sim 5$ to which we can compare our results. Karkare and Bird (2018) explore the possibility of a BAO detection with CO and [CII] line emission and show that while current generation experiments will be unable to achieve a detection, next-generation and more futuristic experiments could achieve 5% and ~ 1% level detection respectively, contingent on the intensity of the line emission. Bernal et al. (2019) show that using SPHEREx H-α measurements at $z = 4.52$ it is possible to measure $\alpha_\perp$ to a precision of 7.0%, the same level as our Next Generation scenario. Using SPHEREx Ly-α they show it is possible to reach a level of precision of 5.2% at $z = 5.74$. They too include a conception of a futuristic CO mapping instrument which, similar to Karkare and Bird (2018), is able to achieve a 1.5% level detection at $z = 5.3$. While we do not achieve quite this level of precision in our Futuristic scenario, it should be noted that lensing measurements are not subject to the same astrophysical modelling biases to which LIM BAO measurements are subjected. Since lensing is an unbiased tracer, there is no disentangling the matter power spectrum component from the astrophysical components of the measured spectrum as is the case with LIM surveys.

Considering now BAO forecasts from HI experiments, SKA will not be capable of making BAO detection from the angular direction at frequencies below 800 MHz ($z > 0.78$) (Bacon et al., 2020). This is due to the angular smoothing of BAO features by the SKA beam. SKA will, however, be able to make a BAO detection in the radial direction out to $z \sim 3$. The detection of velocity-induced acoustic oscillations (VAOs) in the HI power spectrum has also been shown to be a promising standard ruler at cosmic dawn; however, VAOs are damped and undetectable by $z \sim 12$ though some work has suggested their signature may persist to lower z (Muñoz, 2019a,b; Sarkar and Kovetz, 2023; Cain et al., 2020). Measuring the BAO scale during and soon after reionization is indeed a challenging feat.

While the constraints we present here are already competitive with existing forecasts, they can still be improved upon. Our forecast here uses just a single pair of LIM frequency channels. By making use of the fine spectral resolution of LIM experiments and combining this spectral information while nulling, it is possible to yield a higher significance detection. The precise gain in SNR one would obtain from combining spectral information we leave for future work. In addition, the tried and true methods for increasing one's constraints apply here. Increasing the survey areas, performing lensing reconstruction with higher $\ell_{\max}$, and improving foreground removal techniques for line interlopers can all lower the CMB × LIM-nulling variance. Of course, another avenue would be to place more stringent priors on the fit, although care must be taken to ensure that the posterior is not prior-dominated.

5.6 Conclusion

We have shown that it is possible to measure the BAO scale over a wide redshift range from a CMB lensing observable. The BAO features that emerge in the CMB × LIM-nulling convergence spectrum serve as a proxy for the BAO features at $z \sim z_{\rm null}$. Using next generation instruments, we have shown, by way of an AP test, that the detection of this feature can act as a standard ruler, and can constrain the perpendicular AP parameter, $\alpha_\perp$, to the 4% level in a Futuristic observing scenario and to 7.2% in the Next Generation scenario. These constrains are competitive with existing high-z forecasts of BAO measurements from LIM surveys. This technique may be used to tease out information about the matter density field over a large period of cosmic history from difficult to reach redshifts purely using information from lensed CMB photons.

Acknowledgements

The authors are delighted to thank Marta Spinelli and Alexandre Refregier for useful discussions and would especially like to thank José Luis Bernal, Gabriela Sato-Polito, Julian Muñoz, Simon Foreman, and Eiichiro Komatsu for their insightful comments on an earlier draft. HF is supported by the Fonds de recherche du Québec Nature et Technologies (FRQNT) Doctoral Research Scholarship award number 315907 and acknowledges support from the Mitacs Globalink Research Award for this work. AL and HF acknowledge support from the Trottier Space Institute, the New Frontiers in Research Fund Exploration grant program, the Canadian Institute for Advanced Research (CIFAR) Azrieli Global Scholars program, a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant and a Discovery Launch Supplement, the Sloan Research Fellowship, and the William Dawson Scholarship at McGill. ARP was supported by NASA under award numbers 80NSSC18K1014, NNH17ZDA001N, and 80NSSC22K0666, and by the NSF under award number 2108411. ARP was also supported by the Simons Foundation. This research was enabled in part by support provided by Compute Canada (www.computecanada.ca).

Bibliography

W. Hu and T. Okamoto, **574**, 566 (2002), astro-ph/0111606.

T. Okamoto and W. Hu, **67**, 083002 (2003), astro-ph/0301031.

U. Seljak and M. Zaldarriaga, **82**, 2636 (1999), astro-ph/9810092.

A. Lewis and A. Challinor, Phys. Rev. D **76**, 083005 (2007), URL https://link.aps.org/doi/10.1103/PhysRevD.76.083005.

M. Schmittfull and U. Seljak, **97**, 123540 (2018), 1710.09465.

R. Allison, P. Caucal, E. Calabrese, J. Dunkley, and T. Louis, **92**, 123535 (2015), 1509.07471.

A. S. Maniyar, E. Schaan, and A. R. Pullen, **105**, 083509 (2022), 2106.09005.

D. Huterer and M. White, **72**, 043002 (2005), astro-ph/0501451.

F. Bernardeau, T. Nishimichi, and A. Taruya, **445**, 1526 (2014), 1312.0430.

A. Barthelemy, S. Codis, C. Uhlemann, F. Bernardeau, and R. Gavazzi, **492**, 3420 (2020), 1909.02615.

F. McCarthy, S. Foreman, and A. van Engelen, **103**, 103538 (2021), 2011.06582.

F. J. Qu, B. D. Sherwin, O. Darwish, T. Namikawa, and M. S. Madhavacheril, **107**, 123540 (2023), 2208.04253.

Z. J. Zhang, Y. Omori, and C. Chang, **524**, 6392 (2023), 2211.09617.

H. Fronenberg, A. S. Maniyar, A. R. Pullen, and A. Liu, Physical Review D **109** (2024), ISSN 2470-0029, URL http://dx.doi.org/10.1103/PhysRevD.109.123518.

F. Beutler, C. Blake, M. Colless, D. H. Jones, L. Staveley-Smith, L. Campbell, Q. Parker, W. Saunders, and F. Watson, Monthly Notices of the Royal Astronomical Society **416**, 3017 (2011), URL `https://doi.org/10.1111%2Fj.1365-2966.2011.19250.x`.

A. J. Ross, L. Samushia, C. Howlett, W. J. Percival, A. Burden, and M. Manera, Monthly Notices of the Royal Astronomical Society **449**, 835 (2015), URL `https://doi.org/10.1093%2Fmnras%2Fstv154`.

S. Alam, M. Ata, S. Bailey, F. Beutler, D. Bizyaev, J. A. Blazek, A. S. Bolton, J. R. Brownstein, A. Burden, C.-H. Chuang, et al., Monthly Notices of the Royal Astronomical Society **470**, 2617 (2017), URL `https://doi.org/10.1093%2Fmnras%2Fstx721`.

E. A. Kazin, J. Koda, C. Blake, N. Padmanabhan, S. Brough, M. Colless, C. Contreras, W. Couch, S. Croom, D. J. Croton, et al., Monthly Notices of the Royal Astronomical Society **441**, 3524 (2014), URL `https://doi.org/10.1093%2Fmnras%2Fstu778`.

T. M. C. Abbott, F. B. Abdalla, A. Alarcon, S. Allam, F. Andrade-Oliveira, J. Annis, S. Avila, M. Banerji, N. Banik, K. Bechtol, et al., Monthly Notices of the Royal Astronomical Society **483**, 4866 (2018a), URL `https://doi.org/10.1093%2Fmnras%2Fsty3351`.

J. E. Bautista, N. G. Busca, J. Guy, J. Rich, M. Blomqvist, H. du Mas des Bourboux, M. M. Pieri, A. Font-Ribera, S. Bailey, T. Delubac, et al., **603**, A12 (2017), `1702.00176`.

M. Ata, F. Baumgarten, J. Bautista, F. Beutler, D. Bizyaev, M. R. Blanton, J. A. Blazek, A. S. Bolton, J. Brinkmann, J. R. Brownstein, et al., Monthly Notices of the Royal Astronomical Society **473**, 4773 (2017), URL `https://doi.org/10.1093%2Fmnras%2Fstx2630`.

H. du Mas des Bourboux, J.-M. Le Goff, M. Blomqvist, N. G. Busca, J. Guy, J. Rich, C. Yèche, J. E. Bautista, É. Burtin, K. S. Dawson, et al., **608**, A130 (2017), `1708.02225`.

T. M. C. Abbott, F. B. Abdalla, J. Annis, K. Bechtol, J. Blazek, B. A. Benson, R. A. Bernstein, G. M. Bernstein, E. Bertin, D. Brooks, et al., Monthly Notices of the Royal Astronomical Society **480**, 3879 (2018b), URL `https://doi.org/10.1093%2Fmnras%2Fsty1939`.

J. L. Bernal, P. C. Breysse, and E. D. Kovetz, Phys. Rev. Lett. **123**, 251301 (2019), URL `https://link.aps.org/doi/10.1103/PhysRevLett.123.251301`.

C. M. Hirata and U. Seljak, **68**, 083002 (2003a), `astro-ph/0306354`.

C. M. Hirata and U. Seljak, **67**, 043001 (2003b), `astro-ph/0209489`.

D. J. Eisenstein and W. Hu, **496**, 605 (1998), `astro-ph/9709112`.

N. Hand, Y. Feng, F. Beutler, Y. Li, C. Modi, U. Seljak, and Z. Slepian, The Astronomical Journal **156**, 160 (2018), URL `https://doi.org/10.3847%2F1538-3881%2Faadae0`.

E. Schaan and M. White, Journal of Cosmology and Astroparticle Physics **2021**, 068 (2021a), URL `https://doi.org/10.1088%2F1475-7516%2F2021%2F05%2F068`.

E. Schaan and M. White, Journal of Cosmology and Astroparticle Physics **2021**, 067 (2021b), URL `https://doi.org/10.1088%2F1475-7516%2F2021%2F05%2F067`.

A. Cooray, T.-C. Chang, S. Unwin, M. Zemcov, A. Coffey, P. Morrissey, N. Raouf, S. Lipscy, M. Shannon, G. Wu, et al., *Cdim: Cosmic dawn intensity mapper final report* (2019), URL `https://arxiv.org/abs/1903.03144`.

M. Silva, M. G. Santos, A. Cooray, and Y. Gong, **806**, 209 (2015), `1410.4808`.

P. Ade, J. Aguirre, Z. Ahmed, S. Aiola, A. Ali, D. Alonso, M. A. Alvarez, K. Arnold, P. Ashton, J. Austermann, et al., **2019**, 056 (2019), `1808.07445`.

K. Abazajian, G. E. Addison, P. Adshead, Z. Ahmed, D. Akerib, A. Ali, S. W. Allen, D. Alonso, M. Alvarez, M. A. Amin, et al., **926**, 54 (2022), `2008.12619`.

N. Sailer, E. Castorina, S. Ferraro, and M. White, **2021**, 049 (2021), 2106.09713.

Planck Collaboration, **594**, A13 (2016).

D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, **125**, 306 (2013), 1202.3665.

K. S. Karkare and S. Bird, **98**, 043529 (2018), 1806.09625.

D. J. Bacon, R. A. Battye, P. Bull, S. Camera, P. G. Ferreira, I. Harrison, D. Parkinson, A. Pourtsidou, M. G. Santos, L. Wolz, et al., Publications of the Astronomical Society of Australia **37**, e007 (2020).

J. B. Muñoz, **100**, 063538 (2019a), 1904.07881.

J. B. Muñoz, **123**, 131301 (2019b), 1904.07868.

D. Sarkar and E. D. Kovetz, Phys. Rev. D **107**, 023524 (2023), URL `https://link.aps.org/doi/10.1103/PhysRevD.107.023524`.

C. Cain, A. D'Aloisio, V. Iršič, M. McQuinn, and H. Trac, **898**, 168 (2020), 2004.10209.

Chapter 6

Conclusion

In this thesis, we have developed novel techniques for unveiling the early Universe by maximizing the scientific return of LIM measurements through the powerful use of cross-correlations and synergies between multiple cosmological probes. As we discussed in Chapter 1, LIM holds immense potential as an unprecedented probe of large-scale structure across the majority of cosmic history, and yet, numerous observational and theoretical challenges stand between its promise and its full realization. These challenges include foreground contamination, instrumental systematics, and the careful accounting of these effects in statistical analysis and error budgeting, all of which must be carefully addressed to enable high-precision cosmology.

In Chapter 2, we developed a rigorous end-to-end simulation and analysis framework for LIM cross-correlations. This enabled, for the first time, the robust statistical treatment of the joint effects of foregrounds, instrumental response and instrument noise from both LIMs. This framework allowed us to critically examine key assumptions in LIM power spectrum analyses, quantify error properties, and forecast the feasibility of upcoming 21 cm-[CII] cross-correlation measurements. By moving beyond the simplifying assumptions often made in the literature, this work also provides a more realistic and practical foundation for designing future LIM surveys.

Foreground contamination is one of the most significant barriers to LIM detections, and in

Chapter 3, we tackled it head-on by introducing a fully data-driven foreground removal method that leverages intra-dataset frequency-frequency cross-correlations to empirically estimate and subtract the foreground power spectrum. Using simulations of 21 cm observations across different experimental setups, we demonstrated that this technique can suppress foregrounds by 4–5 orders of magnitude, achieving percent-level signal residuals in somewhat idealized conditions. We also explored its limitations, identifying observational scenarios where the method remains robust and those where further improvements are necessary. This approach represents an important step toward precision foreground mitigation in LIM analyses.

While the first half of this thesis focused on near-term challenges in LIM analysis, the second half looked ahead to the scientific potential of LIM once it becomes a mature observational field. In Chapter 4, we began our exploration of LIM-nulling estimators which were first developed in Maniyar et al. (2022). This estimator exploits LIM lensing signals to remove the dominant low-redshift lensing contributions to CMB lensing, enabling high-redshift studies of the matter density field. We demonstrated that this CMB × LIM nulling estimator can place competitive constraints on ΛCDM parameters plus the neutrino mass, while also serving as a model-independent probe of physics beyond ΛCDM.

Building on this, Chapter 5 revealed how LIM-nulling techniques unearth high-redshift BAO signals that are otherwise buried in CMB lensing observables, presenting a fundamentally new standard ruler for cosmology. We showed that by applying the CMB × LIM-nulling estimator, it is possible to detect lensing BAOs, features that are typically washed out due to the line-of-sight projection of standard CMB lensing observables. We forecast that next-generation surveys could measure the BAO scale at $z \sim 5$ to 7.2% precision, with more futuristic scenarios improving constraints to 4% precision. This constitutes an altogether new class of BAO measurements that can be used to probe epochs that are largely unexplored by traditional techniques, paving fresh pathway to high-redshift cosmology.

By peeling back the foregrounds (with techniques from Chapter 2 and 3), we uncover the line luminosity, and by peeling back the line luminosity (with techniques from Chapter 4 and 5), we reveal the very foundation of our Universe: the matter density field, imprinted with signatures of the primordial quantum fluctuations that were seeded at the time of the Big Bang. This work demonstrated that synergies between probes create observables that are greater than the sum of their parts, enabling us to study phenomena that may not be accessible through any single probe alone. Synergies are not merely a means to an end but a powerful tool in their own right, unlocking new windows into the physics that governs the evolution of our Universe.

Bibliography

A. S. Maniyar, E. Schaan, and A. R. Pullen, **105**, 083509 (2022), 2106.09005.