Warsaw, 08 October 2023

dr hab. Maciej Bilicki, prof. CFT PAN
Center for Theoretical Physics
Polish Academy of Sciences
al. Lotników 32/46
02-668 Warsaw, Poland

Review of the doctoral thesis
by Chandra Shekhar Saraf
entitled
"Studying tomographic cross-correlations between CMB gravitational lensing potential and galaxy surveys"

The doctoral thesis of Mr. Chandra Shekhar Saraf deals with observational and theoretical aspects of large-scale clustering of galaxies and their cross-correlation with the gravitational lensing potential of the cosmic microwave background (CMB). The thesis has the form of a book manuscript, written in English, and is composed of 6 chapters: 1) short introduction; 2) theory overview; 3) cross-correlation study between galaxy distribution in the Herschel Extragalactic Legacy Project (HELP) and Planck CMB lensing; 4) theoretical analysis of photometric-redshift induced corrections to the measurements; and 5) another observational analysis of correlations between galaxy distribution (this time from DESI Legacy Surveys) and Planck CMB lensing; this is followed by 6) a brief summary; bibliography; and appendix. The original results are presented in Chapters 3, 4, 5, and of these, the cross-correlation study between HELP and Planck had been described in a peer-reviewed publication co-authored by the candidate and his PhD supervisors [Saraf, Bielewicz & Chodorowski, MNRAS, 515, 1993 (2022)].

The dissertation by Mr. Saraf concerns an important probe of cosmology, namely using CMB lensing to trace the large-scale matter distribution and in a broader perspective, to learn more about properties of the universe and of gravity. The CMB lensing signal arises due to CMB photons being gravitationally lensed by the large-scale structure all the way between the last scattering surface and us. By cross-correlating this signal with galaxy surveys in "tomographic" bins in redshift, we can learn about the properties and evolution of matter distribution on large, cosmological scales. In particular, this could give access to the so-called growth rate of structure, independently of other measurements. This growth rate is a sensitive probe of gravity on large scales and can help us test general relativity in this regime and the validity of the cosmological model.

In his thesis, Mr. Saraf addresses some of the key aspects of such cross-correlations between galaxies and CMB lensing. The first one is the so-called galaxy bias "*b*" – an important nuisance parameter arising because galaxies are biased tracers of the underlying dark matter distribution. Theory gives us prescriptions for statistical properties of the latter (e.g. matter power spectrum), and measurements made with galaxies need to be properly related to such predictions. Constraining galaxy bias, for instance as a function of redshift, is of great relevance here.

The second issue addressed in the dissertation is the fact that in such cross-correlations, the galaxies usually do not have exact, spectroscopic redshifts. As this type of measurements require wide-angle sky coverage and considerable depth, at present the most appropriate here are imaging-based datasets, where redshifts are estimated from multi-band photometry and come with considerable scatter and various biases. Such photometric redshifts (photo-zs) have to be re-calibrated so that the true redshift distribution of the employed galaxies is derived, as the latter is the key ingredient in theoretical predictions for the measured signal.

Last but not least, one can measure the amplitude of galaxy-CMB lensing cross-correlation, "*A*". In the standard model, this amplitude is by definition equal to 1; however, several studies, including those performed by Mr. Saraf and described in the thesis, lead to $A \neq 1$ and occasionally this departure is statistically significant. As such a value of *A* cannot be accommodated in the standard model, this could hint at some unaccounted for systematics or even the failure of the cosmological model.

Mr. Saraf presents a number of original contributions to resolving these issues, both from the theoretical and observational perspective. The thesis nicely combines applications of both observed and simulated galaxy catalogs and CMB lensing maps, as well as theoretical developments in modeling the above-mentioned systematics related to photometric redshifts.

Let me start the more detailed review of the dissertation by a comment on what is given in the **Abstract** and later in the thesis, namely that the galaxy bias and lensing amplitude are considered "cosmological parameters". I would dispute this nomenclature; galaxy bias is a phenomenological way to quantify the galaxy–dark matter connection and is certainly not a quantity that describes a cosmological model such as LCDM, especially that its value depends on the particular parametrization. Similarly, the amplitude A is a simple rescaling between measurements and theoretical predictions, and its departure from unity points at some issues with the former or the latter (or both), but again, it cannot be considered a "cosmological parameter". I would only consider the amplitude of fluctuations at 8 Mpc/h scales, σ8, discussed and measured in the thesis, to be an actually *cosmological* parameter.

Chapter 1 of the thesis is the Introduction. As indicated earlier, it is very short (9 pages) and in addition, I find it rather superficial. While the context of the standard cosmological model and its various probes is given, including CMB, its weak lensing and cross-correlation with the large scale structure, we are not offered even a single equation to support the discussion. There is one figure (Fig.1.1) – schematic overview of the analysis – however further illustrations would be desirable. Moreover, several statements made in the Introduction are not accurate or not properly supported. Finally, the bibliography could be occasionally extended or updated. Below I provide more details:

- In *Redshift and the Expanding Universe*: it was not Hubble to notice a systematic shift towards longer wavelengths (of distant galaxies); this fact had been know at least since the spectroscopic measurements of Vesto Slipher from 1910s, which Hubble actually used in his analysis, together with those from Milton Humason.

- In *Large-Scale Structure of the Universe*: the thesis suggests Peebles, 1980; Coil, 2013; and Bernardeau et al., 2002 as "current updates on the topic". While these are certainly seminal and classic works, I believe "current updates" should be looked for elsewhere, taking into account the observational and theoretical developments of the recent decade.

- In *Dark Matter and Gravitational Lensing*: the references Natarajan et al. 2017 and Tyson et al. 1998 in the context of gravitational lensing seem rather random; instead, such reviews as Bartelmann & Schneider 2001 (mentioned only in Sec.2.2) would be much more relevant.

- Finally, the Section 1.1 would be in my opinion much more complete and useful for the rest of the thesis, if the following were at least briefly described: theoretical basis of the cosmological model (Friedman equations, FLRW metric); what are the key cosmological parameters that galaxy clustering and CMB lensing give access to; matter perturbations, power spectrum and σ8 parameter. At present, the key theoretical entity for the topic of this thesis – matter power spectrum – is only implicitly given later (first time in Eq. 2.8), while the σ8 parameter, discussed in several places and even measured in Chapter 5, is never defined anywhere in the entire thesis!

Section 1.2:

- Paragraph dealing with the CMB power spectrum would be more complete if it was accompanied by an illustration of this power spectrum. Without it, it is hard to visualize what peaks and their ratios are meant.

- "the Universe is flat or very close to flat, implying that it contains precisely the critical density required for it to remain nearly flat over cosmic time" – as the thesis does not provide much background on the ingredients of the LambdaCDM cosmological model, and previously only matter content was discussed in the CMB context, this could suggest that the observations point to a flat universe filled with matter only. I believe some at least brief discussion of dark energy, and how its necessity is deduced from CMB itself, would be in place here.

Section 1.3:

- "Weak lensing broadens the acoustic peak, thus, changing the statistics of the size distribution and diffusing the size of the under-dense or over-dense regions." I would be grateful for clarification what is meant in the final part of this sentence.

Section 1.4:

- "The cross-correlation approach employed over redshift slices is a powerful tool to study the dynamical evolution of dark energy from the onset of cosmic acceleration and to test the validity of the cosmological models as a function of redshift." – it would be good to have references supporting this statement. This would also benefit from some more quantitative discussion *how* we can use such cross-correlation for this type of studies. What in such measurements allows us to study the acceleration and test the validity of the cosmological models?

- "Cross-correlation measurements are affected by the amplitude and growth of the matter power spectrum and how modifications are made to General Relativity." – this would benefit from relevant equations showing how these modifications would affect such measurements. Also, what is in particular meant by "modifications to GR"?

- "Any deviations from the underlying theory of gravity, General Relativity, will reflect directly on CMB lensing and hence, on cross-correlations." – how (quantitatively or at least qualitatively) would such deviations affect CMB lensing?

Section 1.5:

- "Even though GR has been successfully tested on many occasions, most of its bounds come from the observations made within our Solar System." – this would benefit from references. Also, I think we have equally important constraints on GR from astrophysical phenomena, such as binary pulsars, gravitational wave detections, or black hole observations (SGR A*, M87)

- I do not find Fig.1.1 very helpful as it is not clarified in the thesis how in detail its left-hand side part translates to what is in the most right-hand side (modified theory of gravity, dark energy models).

- In the paragraph discussing “cross-correlation analyses between CMB weak lensing and galaxy samples”: “These low-redshift cross-correlation probes consistently measure a lower value for S8 as compared to the high-redshift CMB-only measurements from Planck satellite (Planck Collaboration et al., 2020a), resulting in the so-called S8 tension.” I would like to point out that so-called S8 tension originally arose from measurements of *galaxy* lensing (cosmic shear) and not that of CMB. In addition, some papers cited in this paragraph either do not provide S8 measurements (e.g. Peacock & Bilicki 2018) or do not employ CMB lensing (e.g. Pandey et al. 2022).

Chapter 2 shortly overviews the theory and methodology employed in the thesis. It provides the key equations and framework for parameter inference performed in the thesis (via maximum likelihood estimation). Occasionally, some more in-depth discussion and clarifications could improve the presentation, though.

- Equation 2.8 gives the theoretical reference for cross-correlations. Where would the possible departures from GR, mentioned several times in the Introduction, show up in here?

- Which power spectrum is generated from CAMB for the analysis? Linear, non-linear? Does it matter for the analysis? This is relevant also for subsequent Chapters.

- In my opinion Eq.2.8 could be followed by explicit equations for Cgg and Cgκ, which would show us how these depend on the bias and lensing amplitude. This is discussed later in Sec.2.4 in the context of parameter estimation, but the equations are not provided.

- In Section 2.3, it would be good to see some more background and discussion of other approaches to measuring the angular power spectrum:

* First, I consider it an important omission that the seminal works on the topic by Peebles in 1970s are not mentioned.

* Second, many more teams developed methods to estimate the full-sky power spectrum from partial-sky data and references to those are lacking. Among them there is actively developed and maintained NaMASTER code used for instance by the LSST collaboration. Why has the candidate decided to implement his own angular power spectrum estimator, instead of using something “off-the-shelf”, but well-tested, such as NaMASTER?

* Last but not least, what are the assumptions and limitations of the approach summarized in Fig.2.1? In particular, how small sky coverage could be still used to estimate the full-sky power spectrum without considerable biases? Is in particular the sky coverage of the HELP samples, used in the subsequent chapter, sufficient?

- Section 2.4 would in my opinion benefit most from improvements in this chapter:

* It’s surprising to see no mention of the fact that the discussed methodology is Bayesian. Also, an equation relating the probabilities, likelihood, priors etc. would greatly support the subsequent discussion.

* The likelihood function of what is maximized? (first sentence of the section)

* “The amplitude parameter can then be used to test the validity of the cosmological model.” – it would be good to have it explicitly shown how this could be done, with references.

* “The galaxy linear bias, on the other hand, can be used to put constrains on models of structure formation in the Universe.” – again, this would benefit from references.

* The data vector is composed of Cgg and Cgκ. Is then the CMB lensing auto-correlation Cκκ not used, and if not – why? See also below.

Chapter 3, based on Saraf et al. 2022 (MNRAS), presents a cross-correlation between Planck CMB lensing and galaxy samples from the HELP project, joined with auto-correlation of the latter. This is used to estimate the galaxy bias of HELP galaxies and the amplitude of cross-correlation. The latter is in some cases found to be much below the fiducial value of unity, and several systematics and analysis choices are inspected to verify if one (or more) of them is not responsible for shifting A from the expected value of 1. In most cases the estimation of A is immune to these; one exception is however when different CMB lensing maps are used. In this latter case the inferred value of A can vary considerably and in particular become consistent with 1.

My general comments on this chapter are that we are not given motivation and conclusions. Namely, why was that study done in the first place? Why this particular galaxy dataset and not some other? And what are the conclusions from this work? Is there an issue with the lensing amplitude being inconsistent with unity, or not? Especially taking into account Sec.3.7.6.

For some more details:

- Among the possible systematics tested are effects of photometric redshifts. Two aspects were verified: shifting the mean of the redshift distribution, and injecting catastrophic redshift errors. Regarding the latter, let me note that often they take the form of *systematic* rather than random shifts, as done currently in the thesis. For instance, some spectral features might be misidentified and galaxies assigned systematically over- or underestimated photo-zs. In addition, as addressed later in the thesis, the assumption of Gaussian photo-z errors may be inaccurate due to considerable non-Gaussian “wings” in the error distribution. Testing also for these effects could be potentially useful when looking for possible influence on the derived amplitude A.

- As generally discussed in this thesis, the ability to properly calibrate redshift distributions of photometric samples is key for analysis of this type of correlations. Here, this is based on photo-z error estimates from Duncan et al. 2018a,b. First, it would be good to have some details on how these errors were derived. Second, let me emphasize that these errors seem substantial. At the median redshifts of the particular fields, the relative errors are several dozen per cent (e.g. ~50% for NGP). This raises the question how well we can calibrate the true redshift distributions with so uncertain photo-zs.

- Above Fig.3.2, full CMB lensing map is mentioned. But earlier in Chapter 2, the data vector did not include Cκκ. In which part of the analysis is then the full CMB map used if for cross-correlation only the overlap with HELP is needed? Similarly, in Sec.3.3 Cκκ is mentioned. If it is in fact used as part of the data vector, is the discussion of the covariance matrix in Sec.2.4 valid? The same applies to the discussion around Fig.3.8, where again Cκκ is not present.

- Both here and in Chapter 5, the constraints on galaxy bias b have very small errorbars. In addition, as is mentioned, the bias is constrained mostly by the galaxy auto-correlations, and hardly affected by the cross-correlation Cgκ. It would be good to elaborate more on that as such per-cent level errors on the bias certainly do not reflect all sources on the uncertainty. This is evidenced by significant shifts of b0 visible in figs. 3.11, 3.12 , 3.14 & 3.15.

- How much are the results affected by bias parametrization? One possible test it by changing the bias model from the one given by Eq.3.9 to for instance constant effective bias for each HELP patch, or a b(z) model linear in redshift (b0 + b1*z).

- Section 3.7.6 leaves me somewhat puzzled. First, earlier in the Introduction not much discussion was given on how CMB lensing maps are obtained (reconstructed) from input Planck data and that there might be more than one version of such a map from the same survey (here 3 possibilities are discussed). Second, which of these reconstructions is the most robust and why? Last but not least, which map is used significantly affects the derived amplitude A and hence the conclusions. How should one then decide which of the derivations in Table 3.5 is the correct one? Taking into account the scatter between the values of A for 3 different maps and the associated errors, can we still conclude that there is an issue with A being significantly below unity for this analysis?

Chapter 4 present unpublished (as of now) study of an important issue affecting correlation measurements performed with photometric data, namely the “leakage” of objects between redshift bins. This is induced by the fact that when binning is done in photo-zs, which have considerable errors as compared to the truth, galaxies will be scattered between bins. This then affects the amplitude of correlations, especially the galaxy-galaxy ones, as part of the sources used come in fact from a different redshift range and hence do not correlate with those in a given bin (neglecting lensing magnification effects). This was earlier analyzed in e.g. Zhang et al. (2010) and quantified in the form of the so-called “scattering matrix”. Later, a method to compensate for this leakage was proposed in Zhang et al. (2017). Here Mr. Saraf presents an alternative method to compute the scattering matrix, which is faster and more efficient than the approach of Zhang et al. (2017). This allows him to compensate for the leakage correction and, as shown using simulated data, recover the correct auto- and cross-power measurements between mock LSST-like samples and Planck CMB lensing maps.

I think this is a very interesting and valuable approach and as such it gives promise for future studies, especially that its applicability to real data is demonstrated in the subsequent Chapter 5. However, it is not the only possible method in this context and some discussion on that would be valuable. Namely, the signal “leaking out” of the redshift bins can be recovered by cross-correlating various bins, especially the adjacent ones. Such cross-correlations can be modeled theoretically in a similar manner as bin auto-correlations, and joint inference performed. This was for instance done in Balaguera-Antolinez et al. (2018), but more importantly also in Hang et al. (2021), which is the main reference for the study in the subsequent Chapter 5. I find it then a bit surprising that it is not at all mentioned here. A relevant question is also how the two approaches compare and what are their pros and cons? It is suggested in the last sentence of this Chapter that the author’s approach be “strictly used for future tomographic studies”. However, why should one prefer it over the arguably better established analysis of various bin cross-correlations in a joint inference?

Several further comments I have on the contents of this Chapter are as follows:

- In Sec.4.1 some references and specifications could be improved. First, Dey et al. 2019b (the same reference as Dey et al. 2019a) describes the *photometric* DESI Legacy Imaging Surveys, from which target catalogs for the Dark Energy *Spectroscopic* Instrument (DESI) were constructed; the proper reference for DESI itself is for instance arXiv:1611.00036 (rather than Fagrelius 2020 – a conference contribution); this applies also to Sec.4.9 and Chapter 6.

- Second, DESI itself is not a good example of a galaxy survey that will have “increased depth and better estimation of photometric redshifts” as it is a spectroscopic one. This is also relevant for the last sentence of the summary (Chapter 6) where again DESI is presented as “multi-wavalength photometric”. On the other hand, some discussion could be valuable on how such a new-era spectroscopic survey, that covers large swaths of sky, offers millions of redshifts and could be binned in true redshift slices, may be used for CMB lensing tomographic studies.

- As detailed in Sec.4.2, an effort is made to simulate expected galaxy distribution from the future LSST survey. At the same time, simulated Planck maps are used and no forthcoming CMB survey, that could supersede Planck in this context, is mentioned. Why is that? Will there be no new, better CMB lensing maps available in the LSST era? Taking into account how the various versions of the Planck CMB lensing map affect the amplitude A (Chapters 3 & 5), switching to a new, better map could be worthwhile.

- As illustrated in Fig.4.9, the leakage affects mostly the galaxy auto-spectrum, leading to (sometimes considerably) underestimated measurements. At the same time, the cross-correlations with CMB lensing are hardly affected by the leakage, and if anything, uncorrected spectra are *overestimated*. Some discussion why this happens would be valuable.

- In Fig.4.13, why do the errorbars on σ8 decrease with redshift, and do not seem to reflect the much decreasing number density of mock galaxies at z>1 (Fig.4.1)? Especially in the final redshift bins, at z>2, the relatively small errors on σ8(z) seem unrealistic, taking into account that even LSST will be sparsely sampling the galaxy distribution at this range and that the photo-z errors (scaling here as 1+z) are expected to be considerable there. Are these effects taken into account in deriving the σ8 uncertainties?

- A somewhat puzzling results is presented in Fig.4.14, which shows that there might be situations (here "Set 4") where the leakage leads to *overestimation* of galaxy auto-correlation. Some discussion how this could arise, could be valuable. This applies also to Sec.5.5 and Fig.5.6.

Chapter 5 includes unpublished results on "tomographic" (redshift-binned) auto-correlation of DESI Legacy surveys joined with their cross-correlation with Planck CMB lensing. The galaxy samples and their binning are taken from the earlier study by Hang et al. (2021; H21 hereafter), while the CMB data are as in Chapter 3. The novelty and originality in the study described here, as compared to H21, lies in: 1) a different treatment of photo-z calibration; 2) the application of the leakage correction from Chapter 4; 3) measurement of redshift dependence of σ8. The end result are measurements of galaxy bias, lensing amplitude, and σ8, as a function of redshift between z~0.2 and z~0.7. These results are compared (where possible) with H21 and the importance of leakage correction as well as of exchanging CMB lensing maps is tested. The results are largely consistent with H21, although one puzzling outcome is from bin #2, where both the lensing amplitude and σ8(z) are well below the LCDM-based expectation.

The study presented in this Chapter is a valuable original contribution to the field, nicely demonstrating what possible improvements can be made both to photo-z treatment and the leakage due to photo-z binning. One general comment I would have here is similar as to the previous Chapter – there is no mention of the fact that this leakage could be compensated for by using bin cross-correlations, which was in fact done in H21. In addition, some aspects of the redshift calibration could be possibly discussed in greater detail, especially that it largely relies on the data released in H21. One further concern is that the errors on galaxy bias reported here are several times smaller than in H21, despite the same data and types of measurements being used. This fact is not commented upon.

- In Sec. 5.3 and 5.4, the redshift calibration relies on overlapping spectroscopic subsamples (enumerated in Sec.5.2). As described in H21, this is done by matching the photometric and spectroscopic samples in the color space. I understand that the photo-z error distribution is then also based on this matching, although details are not provided in the thesis. This is important, as a possibility exists that not all the photometric galaxies have well-quantified photo-z errors due to the sparseness of the spectroscopic data in the multi-color space and/or photometric scatter, considerable at the faintest magnitudes. In addition, such kind of mismatch or noise can affect the details of the error distributions shown in Fig.5.2. More discussion on that would be valuable.

- Regarding Fig.5.2 and associated fitting. Let me point out that the presentation with logarithmic y-axis exaggerates the situation: the number of galaxies at the "wings" is in fact much smaller than near the peak. Secondly, possible sparse sampling at the "edges" of the color space, mentioned above, could highly influence the exact shapes of these curves at their extremities. This raises the question if, first, detailed multi-Gaussian matching of these error curves does not lead to over-fitting and second, if this matters much for the end results. For instance, the χ-square values in Table 5.3

(as well as visual inspection of Fig.5.8) do not clarify if modified Lorentzian modeling (as in H21) is less reliable than the multi-Gaussian, as suggested in this Chapter.

- What is the aim and usefulness of Sec.5.4.2? This deconvolution approach is here applied only to the full distribution, not binned in photo-z, and it does not seem to be used later in the thesis?

- I find Fig.5.7 and associated results somewhat puzzling. It has been shown both in the previous Chapter and earlier in Fig.5.6, that the galaxy-CMB lensing cross-correlation is largely insensitive to leakage, and the correction affects mostly the Cgg spectrum. At the same time, the lensing amplitude is by definition derived only from Cgκ measurements. How can then such large changes in A arise after the leakage correction is applied (right-hand panel of Fig.5.7)? Is this because of the b-A degeneracy?

- Regarding the degeneracy between bias b and amplitude A, mentioned in Sec.5.6.2. This was not evident earlier in Chapter 3: the contours in the b-A plane were rather round, or slightly elongated in the A direction, and the amplitude was largely insensitive to various systematics tested, while b was very much affected. It would be therefore valuable to show an illustration of the b-A plane in this case, if there is indeed such a degeneracy present.

- In Sec.5.6, why the Cgg measurements and predictions are not shown at all? This is particularly relevant as they are the most affected by the leakage correction, but it’s worth knowing if also photo-z error modeling affects them or not.

- As I indicated above, I believe some discussion is relevant on why the derived errors on galaxy bias in here are several times smaller than in H21 (e.g. Table 5.4), despite the same data, binning and probes (Cgg and Cgκ) being used.

- Sec.5.6.3. Similarly as in Chapter 3, using different CMB lensing maps does affect the derivations of A and the χ-square. However, here this influence on A seems much lower. Some discussion why (what is different) would be valuable.

- Sec.5.6.4. What could be the reasons for so low values of A and σ8 in the second redshift bin? Neither leakage correction nor different CMB lensing maps seem to bring especially the latter up to the fiducial value.

Chapter 6 provides a very short, two-page summary. Even if short, it is occasionally not very much in line with the contents of the thesis.

- “Testing cosmological models” is mentioned, which was never directly addressed in the dissertation.

- Summary of Chapter 2 offers no conclusions.

- Summary of Chapter 4 mentions “biased estimation of cosmological parameters” while the main part of analysis dealt with two “nuisance” parameters: galaxy bias and lensing amplitude.

- “This will lead to apparent tensions on cosmological parameters like σ8.” – while this is indeed a conclusion from the σ8-related analysis in Chapter 4, let me point out that the σ8 measured there, if leakage corrections are not applied, is *overestimated* with respect to the true value. At the same time, as also mentioned a few paragraphs below in the summary, the current observational “S8 tension” lies in the fact that S8/σ8 measured from large-scale structure / lensing studies are typically *lower* than predicted by the Planck-based LCDM model.

- “Different CMB lensing maps” are highlighted as important for the analysis, but in the thesis they were studied only briefly, in additional tests in Chapter 3 and 5. The conclusion “They indicate how big are uncertainties related with possible systematic errors of the CMB lensing maps” is a potential red flag, as these systematic errors and their importance were not very much discussed here. In particular, current conclusions especially from Chapter 3 could be that changing the maps heavily affects the estimation of the amplitude A and does not allow us to say much about its consistency with unity or lack thereof. Is there a hope this will matter less in the future?

- “The tensions on cosmological parameters have been a major motivation towards alternative models of our Universe.” This could be substantiated with literature references. Such beyond-LCDM/GR studies are indeed numerous, but the current dissertation unfortunately does not mention any of them.

Overall, despite all these various issues and drawbacks I identified, I believe the dissertation provides important original and valuable contributions to the field. Mr. Saraf highlights and quantifies the importance of various systematic effects in the study of cross-correlation between galaxy surveys and CMB lensing. Several of such effects are then addressed both from the theoretical and observational perspective. I therefore consider the doctoral thesis of Mr. Chandra Shekhar Saraf to meet the criteria prescribed by the law for a doctoral dissertation. Hence, **I request that this dissertation be admitted to a public defense**.

Maciej Bilicki

(signed electronically)