The Lyα forest flux correlation function: a perturbation theory perspective

Shi-Fan Chen[a] **Zvonimir Vlah**[b,c] **Martin White**[a]

[a]Department of Physics, University of California, Berkeley, CA 94720
[b]Kavli Institute for Cosmology, University of Cambridge, Cambridge CB3 0HA, UK.
[c]Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge CB3 0WA, UK.

E-mail: shifan_chen@berkeley.edu, zv217@cam.ac.uk, mwhite@berkeley.edu

Abstract. The Lyα forest provides one of the best means of mapping large-scale structure at high redshift, including our tightest constraint on the distance-redshift relation before cosmic noon. We describe how the large-scale correlations in the Lyα forest can be understood as an expansion in cumulants of the optical depth field, which itself can be related to the density field by a bias expansion. This provides a direct connection between the observable and the statistics of the matter fluctuations which can be computed in a systematic manner. We discuss the way in which complex, small-scale physics enters the predictions, the origin of the much-discussed velocity bias and the 'renormalization' of the large-scale bias coefficients. Our calculations are within the context of perturbation theory, but we also make contact with earlier work using the peak-background split. Using the structure of the equations of motion we demonstrate, to all orders in perturbation theory, that the large-scale flux power spectrum becomes the linear spectrum times the square of a quadratic in the cosine of the angle to the line of sight. Unlike the case of galaxies, both the isotropic and anisotropic pieces receive contributions from small-scale physics.

Contents

1   Introduction

The Lyα forest refers to the structure imprinted in the spectra of high redshift galaxies and quasars by absorption of photons by neutral hydrogen along the line of sight [1–3]. Most of the signal comes from warm gas near mean density that is in photoionization equilibrium with an almost uniform ultra-violet background and as such probes density fluctuations on scales ranging from the Jeans scale ($\mathcal{O}(100\,\mathrm{kpc})$) to the size of the survey (up to Gpc). It is currently our best probe of large-scale structure from sub-Mpc to hundreds of Mpc scales at high redshift, where galaxy redshift surveys are sparse and cover limited area and intensity mapping surveys have yet to report detections.

There are two regimes in which the Lyα forest provides critical constraints on our cosmological models. The first is at small scales, where the sensitivity of the forest to sub-Mpc scales allows constraints on the spectral index, massive neutrinos, warm and fuzzy dark matter and other

candidates that suppress small-scale power [4–14]. The second, which is the focus of this work, is on large scales where the Lyα forest can be thought of as a biased tracer of the density field, much like galaxies or QSOs [15, 16]. On the very largest scales[1] it has been argued that the flux, averaged over a large region, should depend linearly on the matter overdensity and peculiar velocity gradient. In this limit the flux correlation function, or power spectrum, can be computed analytically. The question then arises on what scales such an approximation is valid, how do non-linearities enter, to what extent do the numerous complex physical processes affecting the intergalactic medium on small scales modify the large-scale behavior and how does one systematically extend a linear theory calculation?

Traditionally models of the Lyα forest power spectrum addressing the above questions have involved fits to, or forms inspired by, numerical simulations (see e.g. refs. [26–28] for recent examples and the reviews cited above for a more complete overview). In large part this is because the mapping between the underlying density fluctuations and the observed flux is highly non-linear, mixing small- and large-scale effects, which complicates an analytic or perturbative treatment (though see refs. [29–35] for some notable analytic treatments). However the 3D power spectrum of the Lyα forest is observed to be very close to a linearly biased tracer of the linear power spectrum on large scales [16], which suggests that the mixing of small- and large-scale physics does not alter the flux power spectrum beyond recognition. In this paper we develop a systematic means of computing the flux correlation function beyond linear theory. We show how the effects of this non-linear mapping can be encapsulated into a set of nuisance terms, and isolate the parts of the Lyα forest power spectrum and correlation function that come from large scales and are expected to be amenable to analytic or perturbative treatment and those which depend sensitively on small-scale physics (i.e. "astrophysics"). We discuss how these effects arise in the 1-loop flux power spectrum and demonstrate the approach to linear theory with a specific μ dependence on large scales, including the manner in which higher-order effects enter. We demonstrate how small-scale physics impacts the bias parameters and how redshift-space distortions induce anisotropy in the clustering. Our approach utilizes the fact that the non-linear mapping relating density to flux is of a known form (an exponential of a biased tracer of the density field), and thus should prove to be more general than the particular application we highlight here.

Earlier work [31, 32] also studied the consequences of dynamical nonlinearities under exponential maps, specifically within the peak-background split. These approaches needed to treat redshift-space distortions approximately. Our calculations expand upon these by considering a more general set of nonlinearities within an effective-theory framework, but we make connection with these previous results where appropriate (especially in the context of real space). We also discuss the structure of the theory beyond the 1-loop level, and the allowed angular dependence in the large-scale limit. Recently, refs. [34, 36] proposed a second-order anisotropic bias basis in the context of Lyα and a third-order one for generic biased tracers with line-of-sight selection effects. We show how the combination of dynamics and the exponential map generates the terms in those expansions.

---

[1]Throughout we shall neglect large-scale effects arising from general-relativistic corrections [17, 18] or fluctuations in the ultraviolet background field [19–25], focusing instead on dynamical non-linearities and the way in which small- and large-scale physics couple into the observable flux.

The outline of the paper is as follows. Section 2 shows how the transmitted flux correlation function, ξ_F, is related to cumulants of the optical depth fluctuation, $\delta\tau$. Since $\delta\tau$ can be modeled in a manner quite similar to galaxy or quasar density fields in perturbation theory, this establishes the key connection between our observable and objects which are under good theoretical control. We validate the connection between ξ_F and the cumulants in Section 3 using some highly idealized simulations that illustrate our key points. How the cumulants of $\delta\tau$ behave, how non-linear scales enter and the impact of redshift-space distortions is the topic of Section 4. In this section we present several approaches so as to bring out the key physics and make connection with earlier work. We conclude in Section 5. Details of our calculations and explicit formulae are given in a series of appendices.

2 Formalism

2.1 Optical depth and density

The relationship between the redshift-space optical depth, τ, and the dark matter density, δ, and velocity (divergence) θ is set by the physics of the Lyα forest [1–3]. Specifically the low-density gas is in photoionization equilibrium with an almost uniform ultraviolet background field, and thus traces a power of the dark matter density field on large scales, where pressure forces are unimportant. As shown in Fig. 7 of ref. [37], we are interested primarily in gas near the cosmic mean density. Instead of attempting to model the complex physics determining the optical depth field, even approximately, we shall take the approach that τ can be expanded as a series in the (smoothed) density and velocity fields containing all terms allowed by the symmetries [38]. The coefficients of the terms in such an expansion, the bias coefficients, are to be treated as parameters of the theory that must be fit to observations or numerical simulations that resolve the small-scale physics of the forest.

Using an Eulerian biasing prescription, we have

$$\tau(\mathbf{x}) = \tau_0 \left[1 + b_1 \delta(\mathbf{x}) + \frac{1}{2} b_2 \left(\delta^2(\mathbf{x}) - \langle \delta^2 \rangle \right) + \cdots \right] + \epsilon(\mathbf{x}) \tag{2.1}$$

in real space. Here ϵ is the "stochastic" contribution to τ that is uncorrelated with large scales and has zero mean. Note that in the limit of an isothermal gas where $\tau \propto (1+\delta)^\gamma$ with $\gamma = 2$ this bias expansion would be exact, while for $\gamma \neq 2$ we have $b_2 = \gamma(\gamma - 1)$ and there would be higher order terms. The $\cdots$ includes terms depending upon the shear field and other invariants that can be formed at higher order [38]. Throughout this work we will work primarily in terms of fluctuations in the optical depth, defined as $\tau = \tau_0(1 + \delta\tau)$, where $\tau_0 = \langle \tau \rangle$ is the mean optical depth. Since optical depth is conserved in mapping from real to redshift space, assuming large-scale velocities are gravitationally dominated and invoking the equivalence principle $\delta\tau(\mathbf{s}) = b_1 \delta + \theta$ to lowest order, with θ the velocity divergence. We shall consider how this form is modified by higher-order corrections when going to flux in the following.

In the above discussion we have implicitly worked within the approximation where baryons and cold dark matter are treated as a single fluid. In fact, the two species can be subject to both different small-scale forces due to galactic physics (feedback, outflows, star formation etc.) [39] as well as subtle differences in their densities and velocities post-recombination. These differences

enter at linear order and beyond in the statistics of biased tracers [40–42] and can be especially significant at the high redshifts probed by Lyα measurements [43, 44]. A systematic treatment of these terms in the context of the Lyα — and similar considerations for the inclusion of massive neutrinos — is beyond the scope of this work, though we refer readers to ref. [34] who consider a subset of these operators directly proportional to the baryon-CDM relative velocity.

2.2 Flux correlators and the exponential map

We are interested in correlators of the continuum normalized flux field[2], or transmission fraction, which is defined as the ratio of the observed flux to the continuum flux when there is no intervening absorption. We write the flux, F, in terms of an optical depth

$$F(\mathbf{x}) = e^{-\tau(\mathbf{x})} \tag{2.2}$$

and introduce fluctuations about the mean in the usual manner

$$F(\mathbf{x}) = F_0 \left(1 + \delta F(\mathbf{x})\right) \tag{2.3}$$

where the fluctuations are defined to have zero mean so that $\langle F(\mathbf{x}) \rangle = F_0$. The mean flux, F_0, can be related to the cumulants of the τ field evaluated at a single point, specifically

$$F_0 = \langle F(\mathbf{x}) \rangle = \left\langle e^{-\tau(\mathbf{x})} \right\rangle = \exp\left[\sum_{n=1}^{\infty} \frac{(-1)^n}{n!} \langle (\tau(\mathbf{x}))^n \rangle_c \right] \tag{2.4}$$

and of course F_0 is independent of the chosen point, $\mathbf{x}$. Note that the relation between F_0 and τ_0 depends upon the full PDF of F, and involves non-perturbative physics. It could be measured directly, predicted from numerical simulations or left as a degree of freedom in the model. If we write $F_0 = e^{-\tau_0}(1 + \delta F_0)$ then by Jensen's inequality we have $\delta F_0 \geq 0$.

Next let us consider the 2-point function of the flux field, which is given by

$$\langle F(\mathbf{x})F(\mathbf{x}') \rangle = F_0^2 \left(1 + \langle \delta F(\mathbf{x})\delta F(\mathbf{x}') \rangle \right) = \left\langle e^{-\tau(\mathbf{x}) - \tau(\mathbf{x}')} \right\rangle. \tag{2.5}$$

Using the previous relation of F_0 to the point correlators of τ we can write (see Appendix A)

$$\ln\left(1 + \langle \delta F \delta F' \rangle \right) = \sum_{n=2}^{\infty} \tau_0^n \left[\frac{1 + (-1)^n}{2(n/2)!} \xi^{(n/2,n/2)}(\mathbf{r}) + 2\frac{(-1)^n}{n!} \sum_{m=1}^{\lfloor \frac{n-1}{2} \rfloor} \binom{n}{m} \xi^{(m,n-m)}(\mathbf{r}) \right] \tag{2.6}$$

which expresses the flux correlation function in terms of correlators of the optical depth fluctuations, $\xi_\tau^{(ij)}(\mathbf{r}) = \langle \delta\tau^i \delta\tau^j \rangle_c$. These cumulants encode the manner in which the fluctuations of the optical depth on large scales (that can be modeled e.g. using conventional perturbation theory) affect large-scale flux correlations. Note that at a fixed value of F_0 the value of τ_0 is sensitive to small-scale physics. As in the bias expansion, one should treat τ_0 and F_0 as numbers that cannot be determined from large-scale physics alone – they must be fit from data or simulations

---

[2]Throughout we shall neglect observational issues such as continuum fitting, high column density systems, metal lines, etc. The reader is referred to ref. [16] for a discussion of these issues and references to the observational literature.

– but the functional dependence of $\langle \delta F \delta F' \rangle$ is not free. The large-scale physics (e.g. the baryon acoustic peak) resides in the $\xi_\tau^{(ij)}$. We give some examples of cumulants in the following.

We note that the structure of Eq. (2.6) is dictated by the form of the relationship between F and τ, rather than the physics of the Lyα forest or cosmological dynamics. We thus expect it to be of much more general utility than simply modeling the Lyα forest. While it is that modeling that we investigate further in the rest of the paper, let us end this section by mentioning one (slight) generalization of the general formalism towards cross-correlation of the Lyα forest with other biased[3] tracers (e.g. quasars or dampled Lyα systems), i.e. $\langle (1 + \delta_Q) e^{-\tau} \rangle = F_0 (1 + \xi_\times)$. By promoting the cumulant expansion in Eq. (2.6) to include two random variables (in this case $\delta\tau$ and δ_Q) it is possible to show that

$$\langle \delta^n e^{-\tau} \rangle = F_0 \sum_m \frac{(-1)^m}{m!} \, \kappa_{m,n} \quad , \quad F_0 = \exp\left[\sum_m \frac{(-1)^m}{m!} \, \kappa_{m,0} \right] \quad . \tag{2.7}$$

where $\kappa_{m,n}(\mathbf{r})$ are joint cumulants of the two relative overdensities (see Appendix A.2). If we use this formula with the normal bias expansions for δ_Q and $\delta\tau$ the form of $\xi_\times$ quickly follows.

3 Numerical experiments

3.1 Simulations

To test the convergence of the moment and cumulant expansions we use a suite of N-body simulations. A detailed modeling of the Lyα forest would require high-resolution hydrodynamic simulations, and we are not in a position to run a large volume of such simulations in order to study the large-scale behavior of the flux correlations. Instead we have chosen a toy model which has some (though definitely not all) of the properties of the Lyα problem. In particular, our focus is on the exponential mapping involved in going from $\delta\tau$ to F, and the behavior of the cumulant expansion. We shall artificially increase the pressure smoothing present in the Lyα forest so that the moments and correlators of the mock flux field will be well behaved and we can run very large volumes with low computational cost.

To this end we employ the suite of simulations described in ref. [45] to construct mock optical depth and flux fields. These are 10 FastPM boxes with sidelength 1.536 h^{-1} Gpc run with 2048^3 particles. To produce mock optical depths we take the matter density field at $z = 2$ smoothed with a Gaussian of $R = 5$ or $10\,h^{-1}$Mpc and transformed them via the FGPA approximation [2, 3]

$$\tau(\mathbf{x}) = A \left[1 + \delta_m(\mathbf{x}) \right]^\gamma , \tag{3.1}$$

with $\gamma = 1.5$ and A chosen such that $F_0 = 0.7$. For these simulations $\tau_0 \approx 0.37$ so that $e^{-\tau_0} \simeq 0.69$ and $\delta F_0 \simeq 0.01$. Note that our smoothing scale is roughly two orders of magnitude larger than the actual Jeans scale in the Lyα forest. We want to use such a large smoothing because the correlators in our expansions become increasingly UV-sensitive and numerically ill-behaved with increased powers of τ and we wanted to ensure numerical convergence within our limited computational resources. In addition, these simulations have significantly larger volume

---

[3]The auto-correlation of galaxies or QSOs is already routinely modeled with perturbative methods so such measurements could be naturally included in our formalism.

(though lower resolution) than typical Lyα simulations so that we can better understand the large-scale behavior of the Lyα signal and series convergence. While these simulations contain qualitatively similar nonlinearities and nonlinear maps, it was not our intent that they be faithful imitations of Lyα physics.

3.2 Results

Figures 1 and 2 compare the moment and cumulant expansions to the direct measurement of the pseudo-flux correlation function, ξ_{FF}. In each case the left hand panel shows the expansion in moments:

$$F_0^2(1 + \xi_{FF}) = \left\langle e^{-\tau_1 - \tau_2} \right\rangle = \sum_{n=0}^{\infty} \frac{(-1)^n}{n!} \left\langle (\tau_1 + \tau_2)^n \right\rangle \tag{3.2}$$

truncated at $n = 2$, 3, 4 and 5. In order to minimize the difference due to the "1" piece on the left-hand side, we have subtracted off the contributions from the disconnected contributions summing up to the mean flux squared, i.e.

$$F_0^2 \xi_{FF} = \sum_{n=0}^{\infty} \frac{(-1)^n}{n!} \left(\left\langle (\tau_1 + \tau_2)^n \right\rangle - \sum_{m=0}^{n} \binom{n}{m} \left\langle \tau^m \right\rangle \left\langle \tau^{n-m} \right\rangle \right), \tag{3.3}$$

where we have used $F_0^2 = \sum_{n=0}^{\infty} (-1)^n \left\langle \tau^n \right\rangle / n!$.

The right panel shows the same correlators resummed via the cumulant expansion. Specifically, we have that

$$1 + \xi_{FF} = \exp \left\{ \sum_{n=0}^{\infty} \frac{(-1)^n}{n!} \left(\left\langle (\tau_1 + \tau_2)^n \right\rangle_c - 2 \left\langle \tau^n \right\rangle_c \right) \right\} \tag{3.4}$$

where the constant F_0^2 piece are divided out using $\ln F_0 = \sum_{n=0}^{\infty} (-1)^n \left\langle \tau^n \right\rangle_c / n!$. This is equivalent to Equation 2.6. Note that, comparing to the moment expansion above, there is no explicit dependence on the mean flux F_0^2 in the cumulant expression.

The low order moments provide quite poor approximations to the measured flux correlations. By contrast the cumulant expansion, Eq. (3.4), shown in the right hand panel converges very rapidly, at least for our artificially large smoothing scales. The difference between the two expansions, even at $n = 2$, might be somewhat surprising given that the two-point function of the optical depth $\langle \delta\tau_1 \delta\tau_2 \rangle$ is much smaller than unity; indeed the difference between the blue curves in the right and left panels comes mostly from the normalization factors at second order

$$\xi_{FF}^{n=2} = \frac{\tau_0^2}{F_0^2} \left\langle \delta\tau_1 \delta\tau_2 \right\rangle \text{ (moment) and } \xi_{FF}^{n=2} \approx \tau_0^2 \left\langle \delta\tau_1 \delta\tau_2 \right\rangle \text{ (cumulant)}, \tag{3.5}$$

where in the second relation we have used $|\xi_{11}| \ll 1$. While these two expressions are equal order-by-order in τ, the convergence of F_0 is slow.

At this point it is worth pointing out a particularly interesting feature of the convergence of the cumulant expansion, most easily seen in the right panel of Figure 2. From long experience with Taylor series, or intuition built from cosmological perturbation theory for the matter or galaxy power spectrum, we are used to series converging more or less rapidly as a function of scale. For example linear perturbation theory suffices to model the matter power spectrum at very low k, then 1-loop improves this to intermediate k and so on. We see a different pattern
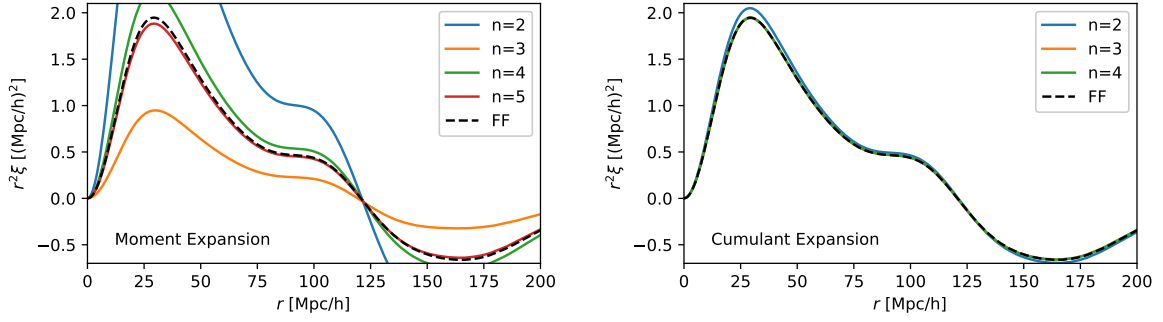
Figure 1. Series convergence for the flux two-point function in our idealized simulations with $R = 10\,h^{-1}\mathrm{Mpc}$ using the moment (left) and cumulant (right) expansions. The cumulant expansion (Eq. 3.4) converges much more rapidly than the moment expansion (Eq. 3.3).
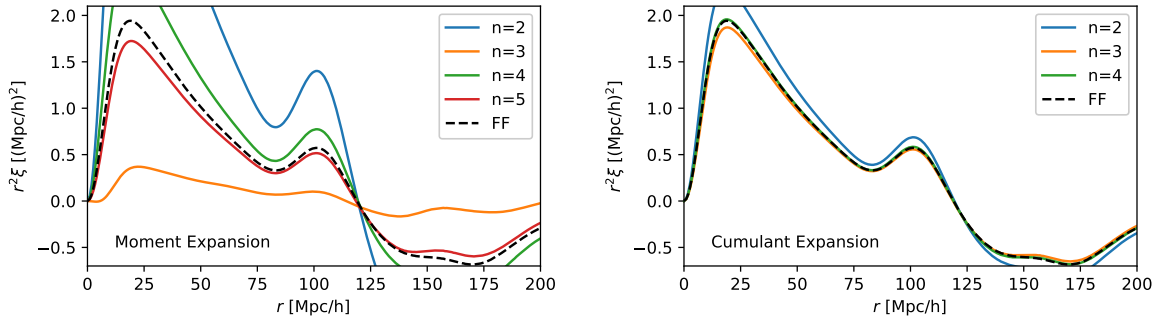

Figure 2. As for Fig. 1 but with $R = 5\,h^{-1}\mathrm{Mpc}$. The cumulant expansion converges much more rapidly than the moment expansion, though not as rapidly as for $10\,h^{-1}\mathrm{Mpc}$ smoothing (Fig. 1). However note that the nature of the convergence is different than might have been expected, since large scales don't converge more quickly than small scales (see text for further discussion).

in Figure 2: the $n = 2$ expression doesn't match ξ_{FF} even above $150\,h^{-1}\mathrm{Mpc}$, including $n = 3$ improves this across all scales but there is still disagreement at both large and small scales (e.g. $25\,h^{-1}\mathrm{Mpc}$ and $160\,h^{-1}\mathrm{Mpc}$), then $n = 4$ improves the agreement even further and so on. This raises an interesting question about why a linear theory form matches Lyα forest data at large scales, that we will discuss fruther in the next section.

Our numerical experiments validate our cumulant expansion and give some insight into the behavior of the individual terms. If we were to continue to reduce the pressure smoothing scale to better approximate our Universe we expect to see the higher cumulants come to play an increasingly important role (e.g. by generating an order-unity velocity bias effect). To understand the implications of this for measurements of the flux correlation function we need to understand the general structure of these cumulants, which we turn to in the next section.

4 Structure of the theory

Key to understanding the behavior of the flux correlation function is that of the correlators, $\xi_\tau^{(mn)}$. In this section we first show that for Gaussian τ the flux correlation function becomes a series in $\xi_\tau^{(11)}$. For the more general, and relevant, case where τ is non-Gaussian the behavior is more complex and more interesting. In particular we find that when either non-linear dynamics or non-linear bias causes τ to be non-Gaussian — even when the linear δ is Gaussian — the higher order correlators ($\xi_\tau^{(mn)}$ for m or n greater than 1) can "renormalize" the lowest order correlator $\xi_\tau^{(11)}$ at large scales. This behavior will be key to understanding the large-scale behavior of the flux correlation function and power spectrum. We present this calculation first in a toy model, the peak-background split, and then motivate the behavior with perturbation theory.

4.1 Gaussian limit

Our Eq. (2.6) becomes particularly simple in the limit that $\delta\tau$ is a Gaussian field, since in that case only $\xi_\tau^{(11)}$ is non-zero. For Gaussian $\delta\tau$ therefore

$$1 + \xi_{FF} = \exp\left[\tau_0^2 \xi_\tau^{(11)}\right] \quad \Rightarrow \quad \xi_{FF} = \sum_{n=1}^{\infty} \frac{\tau_0^{2n}}{n!} \left(\xi_\tau^{(11)}\right)^n \tag{4.1}$$

a result that was previously derived by ref. [46] and can be proved directly using Price's theorem [47]. If $\xi_\tau^{(11)}$ drops rapidly at large scales we see that the flux correlation function is simply proportional to the linear theory correlation function on large scales (see below) while at smaller scales higher powers of $\xi_\tau^{(11)}$ introduce a scale dependent bias. While this simple example neatly illustrates how the exponential map (going from τ to F) generates non-linearity, this model is not very relevant to the Lyα forest so we turn our attention to more complex cases.

4.2 Peak-background split

Let us consider a simple toy model that illustrates the beyond-Gaussian and non-linear behavior. We will employ the peak-background split (PBS), first at the level of the density field, δ, and finally at the level of the optical depth, $\delta\tau$. To start we will neglect redshift-space distortions, deferring them until we discuss perturbation theory.

First we approximate the density field as a sum of long- and short-wavelength modes: $\delta = \delta_\ell + \delta_s + \nu_2 \delta_\ell \delta_s$ with $\nu_2 = 34/21$ the angle average of the Eulerian perturbation theory kernel, F_2 [48]. The ν_2 term approximates the effects of non-linear mode coupling in this model. We consider the case where δ_s is uncorrelated on the separation scale, $\mathbf{r}$, of interest and δ_ℓ is a Gaussian random variable (i.e. a "linear" mode). For an optical depth $\tau[\delta]$ that is an arbitrary local function of the nonlinear density we can write to first order in the long-wavelength density

$$\tau[\delta] = \tau[\delta_s + (1 + \nu_2 \delta_s)\delta_\ell] = \tau[\delta_s] + \tau'[\delta_s](1 + \nu_2 \delta_s)\delta_\ell + \mathcal{O}(\delta_\ell^2). \tag{4.2}$$

suggesting that the (unnormalized) linear bias is $b_\tau = \langle \tau' + \nu_2 \delta \ \tau' \rangle_s$ where primes stand for derivatives with respect to the argument and the average is carried out in the absence of long modes. This recovers the results of ref. [31], and indeed carries over without modification to the linear flux bias in this limit as well given that it is just given by the composite map $F[\delta] =$
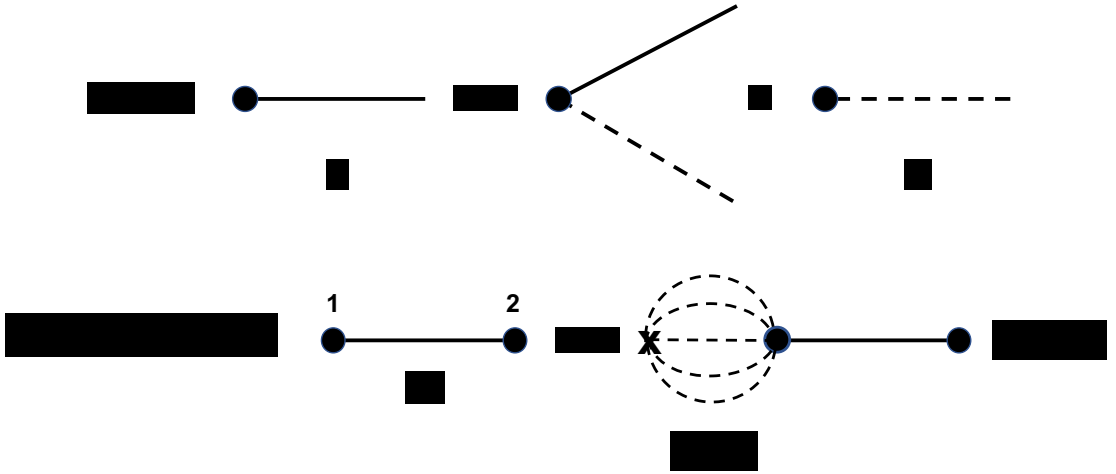
Figure 3. (Top) Contributions to the nonlinear density in the peak-background split. Solid points indicate points in configuration space, while the solid and dashed lines refer to long and short modes which are coupled with strength ν_2 in the middle diagram. (Bottom) Diagrams indicating the terms that contribute to the flux correlation function in the PBS approximation at order n in the optical depth. While the long modes are Gaussian, the short modes can couple to itself at a point to arbitrary order, though not at other points. We have omitted a similar diagram proportional to ν_2^2 that has a corresponding short-mode "bubble" on point on the right.

$\exp(-\tau[\delta])$. This ease of translation does not carry over to redshift space however, as we will see below.

Let us now treat the specific case where $\tau[\delta]$ is given by a bias expansion. This will allow us to anticipate qualitatively what happens within the full perturbation theory argument. Recalling the δ_s are uncorrelated at separation $\mathbf{r}$, the lowest order cumulant in real-space is

$$\xi_\tau^{(11)} = \langle \delta\tau_1 \delta\tau_2 \rangle = b_1^2 \langle \delta_{\ell1}\delta_{\ell2} \rangle + \frac{b_1 b_2}{2} \left\langle \delta_1 \, \delta_2^2 \right\rangle + \cdots \tag{4.3}$$

where $\delta\tau_j = \delta\tau(\mathbf{r}_j)$ and similarly for δ. While the higher order correlators of the purely Gaussian piece, $b_1\delta_\ell$, vanish, the higher order bias and mode coupling terms in the higher cumulants generate contributions like $\langle \delta_{\ell1}\delta_{\ell2} \rangle$. Consider the $b_1 b_2$ term above. Since δ_2^2 contains $2\delta_{s2}\,\nu_2\delta_{\ell2}\delta_{s2}$ we have

$$\left\langle \delta_1 \, \delta_2^2 \right\rangle \supset 2\nu_2 \left\langle \delta_{\ell1}\delta_{s2}\delta_{\ell2}\delta_{s2} \right\rangle = 2\nu_2 \left\langle \delta_s^2 \right\rangle \; \xi_L \tag{4.4}$$

with $\xi_L \equiv \langle \delta_{\ell1}\delta_{\ell2} \rangle$ the long-mode density correlation function. This contribution depends upon small-scale physics through $\langle \delta_s^2 \rangle$ and "renormalizes" the lowest order expression $b_1^2\xi_L$ [49]. We will see a direct analog of this term in the next section. Another such term is

$$\xi_\tau^{(12)} = \left\langle \delta\tau_1 \, \delta\tau_2^2 \right\rangle \supset 2b_1^3\nu_2 \left\langle (\delta_{\ell1} + \delta_{s1})\,(\delta_{\ell2} + \delta_{s2})\delta_{\ell2}\delta_{s2} \right\rangle \supset 2b_1^3\nu_2 \left\langle \delta_s^2 \right\rangle \; \xi_L \tag{4.5}$$

which also depends upon small-scale physics and "renormalizes" the lowest order term. Again, there will be a direct analog in the next section. In fact there are an infinite number of such terms that arise in the cumulant expansion.

We can see the manner in which the higher order terms combine in the two-point function calculation by simplifying our peak-background split model even further. Let us focus on the case where the optical depth is some constant multiple of the nonlinear density, i.e. $\tau = \tau_0(1 + \delta(\mathbf{x}))$, with the density expressed via the PBS, as this form illustrates our point quite clearly. In this limit the correlation function is simply

$$\langle F(\mathbf{x}_1)F(\mathbf{x}_2)\rangle = e^{-2\tau_0}\sum_{n=0}^{\infty}\frac{(-\tau_0)^n}{n!}\left(\sum_{i=1,2}\left[\delta_\ell + (1 + \nu_2\delta_\ell)\delta_s\right](\mathbf{x}_i)\right)^n. \tag{4.6}$$

Our assumption that the long and short modes are uncorrelated, and that the former is Gaussian, then leads to three kinds of contributions to the correlation function, shown diagramatically in Figure 3. First, the long modes can be contracted between the two points in the second cumulant. Then, the contributions proportional to ν_2 can couple either with the pure long-mode or itself. Finally, the pure short-mode contributions must be contracted at a single point and, along with the point contractions of the long modes, amount to the normalization, F_0^2. From these considerations we can extract terms directly proportional to the linear theory correlation function in the cumulants at order n:

$$2n(n-1)\nu_2\left\langle \delta_s^{n-1}\right\rangle \xi_L + n(n-1)\nu_2^2\sum_{m=0}^{n-2}\left\langle \delta_s^{m+1}\right\rangle\left\langle \delta_s^{n-2-m+1}\right\rangle \xi_L$$

where the combinatorial factors denote the number of possible contractions of long and short modes. In particular, we will see in the next section that the $n = 3$ term exactly corresponds to the UV-sensitive piece in perturbation theory for real space, with ν_2 standing in for the matter density kernel F_2 and the bubble with $\left\langle \delta_s^2\right\rangle$ standing in for the large momentum modes in the integral. Summing up all the contributions and normalizing by the total flux yields

$$F_0^2\,\xi_{FF} = \left(\left\langle F'[\delta]\right\rangle_s + \nu_2\left\langle \delta\,F'[\delta]\right\rangle_s\right)^2\xi_L + \mathcal{O}(\delta_\ell^4). \tag{4.7}$$

This agrees with the field-level derivation in Eq. 4.2, as well as similar results in refs. [27, 31, 32].

Let us close by noting the limitations of the above calculations. Firstly, the averages over the short modes (e.g. $\left\langle \delta^2\right\rangle_s$) depend on small scale physics like the suppression of density power below the Jeans scale due to gas pressure and additional highly nonlinear astrophysical processes. Moreover, the PBS calculations above assume that the nonlinearities involved can be modeled as a single density coupling between long and short modes, given by a coefficient ν_2 fixed by second-order perturbation theory when there is in fact no restriction on higher-order couplings or those involving velocities, which are themselves subject to different small scale nonlinearities (e.g. turbulence). Despite these complications, however, the fact that on large scales the theory prediction boils down to bias coefficients multiplying the linear correlation function is significant, showing that the small scales can be modeled via a handful of effective-theory parameters even without a full modelling of small scale astrophysics.

4.3 One-loop perturbation theory

In the previous section we looked at how the cumulants behaved in a simplified model. In this section we will show that similar behavior occurs at 1-loop order in perturbation theory, where

the proper inclusion of redshift-space distortions also gives rise to the phenomenon of velocity bias. In redshift space we assume that the optical depth fluctuation is determined by shifting the real space fluid element[4]

$$1 + \delta\tau_s(\mathbf{r}) = \int d^3\tilde{r} \ (1 + \delta\tau(\tilde{\mathbf{r}})) \, \delta^D \left(\mathbf{r} - \tilde{\mathbf{r}} - \mathbf{u}_\parallel \hat{n} \right), \tag{4.8}$$

where $\mathbf{u}_\parallel$ is the tracer velocity component projected along the line of sight. This is equivalent to the galaxy overdensity redshift space mapping and we are thus able to use the same Eulerian perturbative kernels Z_n (see e.g. [48]) to describe the perturbative orders of $\delta\tau_s$ field.

Compared to the standard 2-point function calculation for biased tracers in redshift space, the Lyα calculation at 1-loop involves only one additional cumulant. This is because at 1-loop order cumulants with more than three powers of the optical depth $\delta\tau$ vanish since, for example, the fourth moment $\langle \delta\tau\delta\tau\delta\tau\delta\tau \rangle \sim \langle \delta\tau\delta\tau \rangle \langle \delta\tau\delta\tau \rangle$ is given at leading order by disconnected contributions. Starting from the Eq.(2.6) which constitutes the cumulant expansion for the flux correlators, then, we have

$$\ln\left(1 + \left\langle \delta F \delta F' \right\rangle_s \right) = \tau_0^2 \xi_s^{(11)}(\mathbf{s}) - \tau_0^3 \xi_s^{(12)}(\mathbf{s}) + \mathcal{O}(P_L^3). \tag{4.9}$$

Exponentiating, expanding and keeping only the one-loop results we obtain the equivalent expression for the flux power spectrum

$$P_{FF}(\mathbf{k}) = \tau_0^2 P_s^{(11)}(\mathbf{k}) - \tau_0^3 P_s^{(12)}(\mathbf{k}) + \frac{1}{2}\tau_0^4 \int_{\mathbf{p}} P_s^{(11)}(\mathbf{p})P_s^{(11)}(\mathbf{k}-\mathbf{p}) + \cdots, \tag{4.10}$$

where we have introduced the cumulant spectra

$$P_s^{(11)}(\mathbf{k}) = \int d^3s \ \xi_s^{(11)}(\mathbf{s})e^{i\mathbf{k}\cdot\mathbf{s}} \quad \text{and} \quad P_s^{(12)}(\mathbf{k}) = \int d^3s \ \xi_s^{(12)}(\mathbf{s})e^{i\mathbf{k}\cdot\mathbf{s}}, \tag{4.11}$$

and the last term is obtained as the Fourier transform of the second term in the expansion of the first cumulant. In order for this term to be consistently evaluated in one-loop perturbation theory only linear level (Kaiser) contributions to $P_s^{(11)}$ are required, besides the appropriate counterterms that are also discussed in Appendix B.

The (11) cumulant is simply the redshift-space two-point function for biased tracers, which is well known. Let us focus on the cumulant $\xi_s^{(12)}$. There are two contributions to this cumulant at 1-loop:

$$\left\langle \delta\tau_1\delta\tau_2^2 \right\rangle = 2\left\langle \delta\tau_1^{(1)}\delta\tau_2^{(1)}\delta\tau_2^{(2)} \right\rangle + \left\langle \delta\tau_1^{(2)}\delta\tau_2^{(1)}\delta\tau_2^{(1)} \right\rangle.$$

These are

$$P_s^{(12)} \supset 2 \times 2Z_1(\mathbf{k})P(k) \int \frac{d^3\mathbf{p}}{(2\pi)^3} \ Z_1(\mathbf{p})Z_2(\mathbf{k},\mathbf{p})P(p). \tag{4.12}$$

and

$$P_s^{(12)} \supset 2 \int \frac{d^3p}{(2\pi)^3} \ Z_2(p,k-p)Z_1(-p)Z_1(p-k)P(p)P(k-p). \tag{4.13}$$

---

[4]See also ref. [32] for some further discussion on the validity of this mapping for Lyα forests and the effects of thermal broadening.

The Z_n are the order n perturbation theory kernels in redshift space ([48] and Appendix B) with $Z_1(\mathbf{k}_1) = b_1 + f\mu_1^2$ and

$$Z_2(\mathbf{k}_1, \mathbf{k}_2) = b_1 F_2(\mathbf{k}_1, \mathbf{k}_2) + f\mu_k^2 G_2(\mathbf{k}_1, \mathbf{k}_2) + \cdots$$
$$+ \frac{fk\mu_k}{2}\left[\frac{\mu_1}{k_1}(b_1 + f\mu_2^2) + \frac{\mu_2}{k_2}(b_1 + f\mu_1^2)\right] \tag{4.14}$$

where $\mu_k = \hat{k} \cdot \hat{n}$, $\mathbf{k} = \mathbf{k}_1 + \mathbf{k}_2$ and F_2 and G_2 are the standard density and velocity kernels in Eulerian perturbation theory [48]. (In the definitions above we have omitted the contributions from higher-order bias such as b_2 and b_s for the sake of brevity. These do not qualitatively affect our conclusions and the full calculation with all relevant terms included are detailed in Appendix B.)

The integral in Equation 4.12 has a form where the linear power spectrum is multiplied by a scale-dependent term that is sensitive to contributions from small-scale modes $p \gg k$. In the large-scale ($k \to 0$) limit we can write

$$Z_1 Z_2 \supset \left(b_1 + f(\hat{p} \cdot \hat{n})^2\right)\left[\left(\frac{5}{7} + \frac{2}{7}\,[\hat{p} \cdot \hat{k}]^2\right)b_1 + \left(\frac{3}{7} + \frac{4}{7}[\hat{p} \cdot \hat{k}]^2\right)f\mu_p^2 + \frac{f\mu_p^2}{2}\left(b_1 + f(\hat{k} \cdot \hat{n})^2\right)\right].$$

The dependence on the long-wavelength wavevector $\hat{k}$ can be pulled out of the above integrand to give (see Appendix B)

$$P_s^{(12)}(\mathbf{k}) \to 2(b_1 + f\mu_k^2)(A_0 + A_2\mu_k^2)\,\sigma^2 P_L(k) \tag{4.15}$$

where $\sigma^2 = \int_{\mathbf{p}} P_L(p)$ is the rms density contrast and $A_{0,2}$ are functions of the bias parameters and growth rate given in the Appendix B. The above form strongly suggests a modification to the Kaiser form of redshift-space distortions

$$\left[(b_1 + \Delta b_1) + (f + \Delta b_\mu)\mu_k^2\right]^2 P_L(k) = (b_1 + f\mu_k^2)^2 P_L(k) + 2(b_1 + f\mu_k^2)(\Delta b_1 + \Delta b_\mu \mu_k^2)P_L(k) + \dots$$

However, it is important to note that the small-scale contributions $\Delta b_{1,\mu}$ need to have their UV sensitivities renormalized by appropriate counterterms, so that their value cannot be determined by the 1-loop calculation above.

Before moving on to the most general case let us comment briefly on some features of the above calculation. Most importantly, the inclusion of redshift-space distortions produces a novel feature, velocity bias, wherein the linear-theory prediction for biased tracers is modified when applied to the Lyα flux. Previous studies [31, 32] of this effect have largely focused on the effect of dynamical nonlinearities in the density; indeed, in the absence of RSD (i.e. $f = 0$) our derivation recovers the PBS result in Equation 4.7 via the angular average $\nu_2 = 2\left\langle 5/7 + 2(\hat{p} \cdot \hat{k})^2/7\right\rangle$. A new result in this paper is that in redshift space contributions coupling the orientation $\hat{k}$ of the long mode and the short modes contribute at equal order owing to both the nonlinear velocity kernel (G_2) and displacements due to the real-to-redshift-space mapping. Finally, a salient feature of our derivation is that the $\hat{k}$ dependence of the integrand can be written as

$$\int d\Omega_{\mathbf{p}}\, Z_1(\mathbf{p})Z_2(\mathbf{p}, \mathbf{k}) \to a^{(0)} + a_{ij}^{(2)}\hat{k}_i\hat{k}_j$$

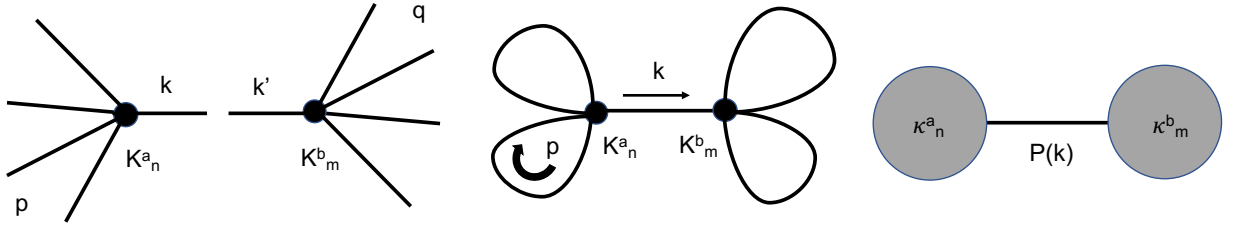**Figure 4**. Diagramatic structure of nonlinearities contributing linear-theory-like terms. These consist of pairs of n^{th} order vertices with one long-wavelength leg $(\mathbf{k}, \mathbf{k}')$, with the other small-scale modes $(\mathbf{p}, \mathbf{q})$ integrated out at each point, leading to a factorizable structure where nonlinear bubbles dress linear theory at each point.

where the tensor $a^{(2)}$ is expressible in terms of δ_{ij} and $\hat{n}_i$ due to the axial symmetry of redshift-space distortions. Thus while the coefficients themselves depend on complicated small-scale physics the parametric contribution is only quadratic in μ^2.

Lastly let us look at the stochastic contributions and Equation 4.13. As we have mentioned, the structure of the $P_s^{(11)}$ cumulant is equivalent to the galaxy case, see e.g. ref. [50] for expressions in the same bias convention including a detailed discussion of stochastic contributions. Equation 4.13 generates the same functional form for stochastic contributions in the second cumulant in the large-scale limit (Appendix B)

$$P_s^{(12)} \supset \text{``const}_0\text{''} + \text{``const}_2^{(0)}\text{''}\frac{k^2}{k_\star^2} + \text{``const}_2^{(2)}\text{''}\mu^2\frac{k^2}{k_\star^2} + \dots, \tag{4.16}$$

which consequently renormalizes the contributions given in $P_s^{(11)}$. However, note that the higher cumulants introduce an explicit dependence on the logarithmic growth rate f to the counterterms "const_0" and "const_2". Importantly, the scale independent, white noise contribution remains isotropic.

4.4 General structure of the large-scale limit

While we have shown it explicitly for the 1-loop spectrum, the large-scale angular dependence (the square of a quadratic in μ) holds beyond 1-loop. This can be established by looking at how the mode-coupling kernels behave in the long-wavelength limit. The contribution going as $P(k)$ as $k \to 0$ can be isolated by considering diagrams linking one power of the linear field at each point with all other momenta contracted with themselves to give loops at each of the two vertices (see Fig. 4). Defining

$$\kappa_n^{(a)}(\mathbf{k}) = \int_{p_i} K_n^{(a)}(k, p_1, ..., p_{n-1}) \left\langle \delta_0(p_1)...\delta_0(p_{n-1}) \, (2\pi)^3 \delta^D(\sum p_i) \right\rangle, \tag{4.17}$$

as the integral over the $(n-1)/2$ short-wavelength loops at one vertex that comes from $\delta\tau^a$ at n^{th} order in the linear density (the last mode being the $\delta(\mathbf{k})$ to be contracted with the other vertex) we see as $k \to 0$ the diagram in Fig. 4 becomes

$$\sum_{a,b}\sum_{n,m}\kappa_n^{(a)}\kappa_m^{(b)}P_L(k) = \left(\sum_a\sum_n\kappa_n^{(a)}\right)^2 P_L(k) \to b^2(\mu)P_L(k) \quad (k \to 0) \tag{4.18}$$

– 13 –

where $b(\mu)$ is the large-scale asymptote of the sum of $\kappa_n^{(a)}$'s which will tend to have non-constant k dependence outside of that limit. That we can factor these diagrams simply into the square of an angular-dependent bias factor $b(\mu)$ is a result of the independence of the loops on each side of Fig. 4. This also directly implies that cross correlations with the Lyα forest can be parametrized in the usual way, i.e.

$$P_s^{g,Ly\alpha}(\mathbf{k}) = (b_1^g + f\mu^2)\, b^{Ly\alpha}(\mu)P_L(k)$$

at linear order. This raises the possibility of improving constraints on the growth rate through "sample variance cancellation" [51] with biased tracers like quasars that obey the Kaiser form even if Lyα does not[5]. Even though Lyα does not have a simple dependence on the growth rate, $f(z)$, as we will now show the functional dependence of $b(\mu)$ is itself limited to two free parameters which can be leveraged in the modeling. Futhermore, we note that at the 1-loop level $\kappa_n^{(a)}$ can also be used to capture all $P^{(13)}$ type contributions by expanding $\kappa_n^{(a)} \supset \alpha_0 k^2 + \alpha_2 k^2 \mu^2 + \cdots$ beyond zeroth order in the wavevector.

To complete the picture we need to show that $b(\mu)$ is at most quadratic in μ. As a first step we note that, as was the case in the 1-loop case (Eq. 4.12), the n^{th} order kernel contributing to $\delta\tau^a$, $K_n^{(a)}$, can be written as a product of redshift-space kernels Z_n, only one of which contains the long mode $\mathbf{k}$. The problem can then be reduced to the angular structure of the redshift-space kernels alone, which indeed must truncate at $\hat{k}_i\hat{k}_j$ as we show in Appendix C as a consequence of the structure of dynamical nonlinearities in structure formation. This implies that $b(\mu)$ must be quadratic in μ to all orders in perturbation theory.

Let us close with some comments on how this argument may be extended to operators beyond linear bias. As we have shown above, modifications to the Kaiser formula arising from short modes contracted at a point lead to a factorable but anisotropic bias term. We can extend this argument to the case of multiple long modes

$$\kappa_n^{(a)}(\mathbf{k}_1, ..., \mathbf{k}_m) = \int_{p_i} K_n^{(a)}(k_1, ...k_m, p_1, ..., p_{n-m}) \left\langle \delta_0(p_1)...\delta_0(p_{n-m})\, (2\pi)^3 \delta^D(\sum p_i) \right\rangle. \quad (4.19)$$

Such contributions will need to be renormalized with their own (anisotropic) counterterms which will again be limited in their forms by the structure of dynamical nonlinearities in the equations of motion (i.e. Z_n). Recently, refs. [34, 36] proposed a second-order anisotropic bias basis in the context of Lyα and a third-order one for generic biased tracers with line-of-sight selection effects. In Appendix C, we briefly sketch the kinds of terms that are generated from the procedure outlined above, showing that they appear to be well-parametrized by the bases proposed in those works. We intend to return to a more systematic study of these terms and further constraints on their forms due to fundamental symmetries [53] in a future work.

5   Conclusions

Modern surveys capable of measuring the spectra of hundreds of thousands or millions of distant objects with modest signal to noise per Å can tightly constrain the flux decrement power spectrum

---

[5]Indeed, soon after this paper was submitted [52] reported promising results on high-redshift growth rate measurements via Lyα-quasar cross correlations using this technique.

on cosmological scales [54, 55]. In turn this provides tight constraints on the distance-redshift relation [16] and provides a new method for measuring the clustering of high redshift objects [56].

The transmitted flux in the Lyα forest probes the underlying density field in a non-linear manner, but on large scales the flux power spectrum is proportional to that predicted in linear perturbation theory [16]. In this paper we have investigated how large- and small-scale physics affect the measured power spectrum by expanding the flux correlation function in cumulants of the optical depth field (Eq. 2.6). Since the optical depth can be related to the density field by a bias expansion (Eq. 2.1), this provides a direct connection between the large-scale, cosmological physics of interest and the measurement. Our formalism gives a systematic means of computing corrections to the flux power spectrum arising from non-linearities, redshift-space distortions and scale dependent bias. It provides an easy way to see the emergence of "velocity bias" in the measured Lyα forest power spectrum in redshift space and the manner in which small-scale physics (i.e. "astrophysics") affects the prediction.

We demonstrated in some idealized numerical experiments (§3) that the convergence of the cumulant expansion does not behave as one might expect from experience with the matter power spectrum in cosmological perturbation theory. Higher cumulants affect the flux correlation function at both large and small scales, rather than being confined to increasingly small scales.

To understand the way in which non-linearities, redshift-space distortions and scale dependent bias affect the flux correlation function, we have presented a perturbative analysis of the flux two-point function due to nonlinearities in the exponential mapping within the peak-background split (§4.2), 1-loop perturbation theory (§4.3), and general considerations of the structure of the equations of motion (§4.4). While the clustering signal for galaxies in redshift space has an "unbiased" contribution due to cosmic velocities given by the Kaiser form $(b_1 + f\mu^2)^2 P_L(k)$, it is well known [15, 16, 57, 58] that the Lyα forest breaks this simple form. Within the framework of effective perturbation theory this is due to the presence of contact terms from higher cumulants which modify the large-scale limit of the power spectrum away from the Kaiser form but preserves the "square of a quadratic in μ" behavior. While previous works have shown this within the context of density nonlinearities and the peak background split, we show using a 1-loop calculation that nonlinear contributions to the velocity, as well as displacements due to redshift-space distortions, contribute at equal order. Extending from this example we show that the quadratic-in-μ form is preserved to all orders in perturbation theory based on the structure of the equations of motion alone and outline how higher-order bias terms can be similarly generated, making connection with the anisotropic bases of refs. [34, 36] and commenting on the possibility for cross correlations and sample variance cancellation. It is worth noting that this analysis of velocity bias should qualitatively extend to less tractable (non-exponential) maps of redshift-space density fields, as has been for example observed in the clustering of cosmic voids [59].

Our calculations demonstrate that for Lyα, unlike for galaxies, both the isotropic and anisotropic contributions to the clustering depend on small-scale astrophysics. In particular, the coefficients of both the μ^0 and μ^2 terms in the linear bias receive contributions from small-scale physics at arbitrary order. While of little direct relevance for measuring distances using baryon acoustic oscillations, this fact has immediate implications for extracting physics from the broad-band shape of the Lyα forest flux power spectrum. If we identify the piece of the spectrum that can be predicted perturbatively with "cosmological information" and the small-scale physics as

"astrophysical information" we see that the Lyα forest contains an irreducibly complex mixture of the two. At the practical level this leads to a larger number of coefficients that need to be modeled or fit to measurements, which necessarily reduces the amount of information that can be obtained. However unlike for the case of galaxies, the Lyα forest can be simulated quite well which may provide a means of using informative priors on the coefficients to improve constraints. An interesting avenue of further investigation may be to measure these coefficients in simulations via the recently developed anisotropic separate universe simulation technique [60, 61], as was already done approximately in [57]. Such methods enable the measurment of responses of e.g. Lyα clustering to large-scale tidal fields with nontrivial angular dependence such as those considered in this work (encoded in $\kappa_n^{(a)}(\mathbf{k})$ of Eqn. 4.18). The feasibility of accurate hydrodynamical modeling of the Lyα forest on small scales and the fact that the impact of these physics on the broadband shape of the correlation function on large scales can be isolated into single coefficients measurable from such simulations suggests that accurate full-shape analyses of the Lyα forest are possible even if simulations and perturbation theory are individually restricted to small and large scales.

Finally, let us briefly comment on the implications of our work for the Lyα forest BAO feature. It is well known that non-linear evolution causes a broadening of the BAO peak in the correlation function of galaxies [62–68] and, similar measurements being one of the premier science goals for the Lyα forest, it is interesting to ask whether the BAO peak in the flux correlation function is similarly broadened. Our formalism provides a means to answer this question, since we can treat $\delta\tau$ as a biased tracer of the density field for which the role of long-wavelength displacements responsible for the damping are well understood and then relate cumulants of $\delta\tau$ to the flux correlation function using Eq. (2.6). What our calculations make clear, however, is that the broadening of the BAO peak in the Lyα forest will be different at each (higher) order in optical depth than it is for tracers such as galaxies or QSOs. Each of the cumulants in Eq. (2.6) responds to the long wavelengths differently, and thus the broadening has the potential to be more complex than a single exponential damping factor sometimes used in the literature. We defer further investigation of this to future work.

Acknowledgments

A Exponential mapping and the 2-point function

A.1 Flux autocorrelation

The two point function of the flux field is given by

$$\left\langle F(\mathbf{x})F(\mathbf{x'})\right\rangle_F = F_0^2 \left\langle 1 + \left\langle \delta F(\mathbf{x})\delta F(\mathbf{x'})\right\rangle_{\delta F}\right\rangle = \left\langle e^{-\tau(\mathbf{x})-\tau(\mathbf{x'})}\right\rangle_\tau \tag{A.1}$$

where the subscripts on $\langle\cdots\rangle$ indicates which distribution is used in the average. We will make use of the cumulant expansion to simplify this expression. Defining

$$Z_\tau(J) = \ln\left\langle \exp\left[\int d\mathbf{x}\, J(\mathbf{x})\tau(\mathbf{x})\right]\right\rangle_\tau \quad \text{and} \quad \left\langle \tau_{\mathbf{x}}^m \tau_{\mathbf{x'}}^{n-m}\right\rangle_{\tau,c} \equiv \frac{\partial^n}{\partial J_{\mathbf{x}}^m \partial J_{\mathbf{x'}}^{n-m}} Z_\tau\bigg|_{J=0} \tag{A.2}$$

we see the cumulant is the coefficient of $J_{\mathbf{x}}^m J_{\mathbf{x'}}^{n-m}/m!/(n-m)!$ in the expansion of Z_τ. Rewriting $1/m!/(n-m)!$ as $\binom{n}{m}/n!$ and using

$$F_0^2 = \exp\left(2\sum_{n=1}^\infty \frac{(-1)^n}{n!}\left\langle \tau^n\right\rangle_c\right) \tag{A.3}$$

we have

$$1 + \left\langle \delta F \delta F'\right\rangle_{\delta F} = \exp\left[\sum_{n=2}^\infty \frac{(-1)^n}{n!}\mathcal{C}_n\right] \quad \text{with} \quad \mathcal{C}_n = \sum_{m=1}^{n-1}\binom{n}{m}\left\langle \tau^m \tau'^{n-m}\right\rangle_{\tau,c} \tag{A.4}$$

where we have used the shorthand $\tau' = \tau(\mathbf{x'})$.

We want to express these correlators in terms of fluctuations of optical depth $\delta\tau$. Replacing $\tau(\mathbf{x}) = \tau_0(1 + \delta\tau(\mathbf{x}))$ in the definition of Z_τ it follows that

$$\left\langle \tau^m \tau'^{n-m}\right\rangle_{\tau,c} = \frac{d^m}{dJ^m}\frac{d^{n-m}}{dJ'^{n-m}}Z_\tau(J)\bigg|_{J=0} = \frac{d^m}{dJ^m}\frac{d^{n-m}}{dJ'^{n-m}}\left[\tau_0\int J + Z_{\delta\tau}(\tau_0 J)\right]\bigg|_{J=0}. \tag{A.5}$$

Since $n \geq 2$ the $\tau_0\int J$ term does not contribute and we have $\left\langle \tau^m \tau'^{n-m}\right\rangle_{\tau,c} = \tau_0^n \left\langle \delta\tau^m \delta\tau'^{n-m}\right\rangle_{\delta\tau,c}$. Introducing

$$\xi^{(ij)}(\mathbf{r}) = \left\langle \delta\tau^i \delta\tau'^j\right\rangle_{\delta\tau,c} \quad , \quad \mathbf{r} = \mathbf{x} - \mathbf{x'}, \tag{A.6}$$

we have

$$\mathcal{C}_n(\mathbf{r}) = \tau_0^n \sum_{m=1}^{n-1}\binom{n}{m}\xi^{(m,n-m)}(\mathbf{r}), \tag{A.7}$$

and thus

$$\ln\left(1 + \left\langle \delta F \delta F'\right\rangle\right) = \sum_{n=2}^\infty \frac{(-1)^n}{n!}\mathcal{C}_n(\mathbf{r}) \tag{A.8}$$

$$= \sum_{n=2}^\infty \tau_0^n\left[\frac{1 + (-1)^n}{2(n/2)!}\xi^{(n/2,n/2)}(\mathbf{r}) + 2\frac{(-1)^n}{n!}\sum_{m=1}^{\lfloor\frac{n-1}{2}\rfloor}\binom{n}{m}\xi^{(m,n-m)}(\mathbf{r})\right]. \tag{A.9}$$

Note that the cumulants, $\xi^{(n,m)}$, contain the large–scale physics but the overall expression is complex which explains why the flux correlation function can look like linear theory on large scales but have complex, scale-dependent bias factors on smaller scales.

A.2 Flux-quasar cross correlations

Let us briefly comment on the modeling of cross-correlations of the flux field with another tracer of large-scale structure, e.g. quasars or damped Lyα systems. As we did for the flux auto-spectrum we shall neglect important observational issues such as continuum fitting, high column density systems, the proximity effect, etc. and focus on the intrinsic signal. Thus we seek to model e.g.

$$\left\langle (1 + \delta_Q) e^{-\tau} \right\rangle = F_0 + \left\langle \delta_Q e^{-\tau} \right\rangle = F_0 \left(1 + \xi_\times \right) \tag{A.10}$$

In perturbation theory we would write δ_Q as an expansion in powers of the matter overdensity, δ. We can proceed by considering the generating function $Z_\times(J, K) = \ln \left\langle \exp\left[J\tau + K\delta\right] \right\rangle$ where $\tau = \tau(\mathbf{x})$ and $\delta = \delta(\mathbf{x}')$. As usual we define the cumulants[6] as

$$\kappa_{m,n} = \left. \frac{\partial^{m+n}}{\partial J^m \partial K^n} Z_\times \right|_{J=K=0} \tag{A.11}$$

so that $\kappa_{m,n}$ is the coefficient of $J^m K^n / m!/n!$ in the expansion of $Z_\times$ and

$$\left\langle e^{K\delta + J\tau} \right\rangle_{\delta,\tau} = \exp\left[\sum_{m,n} \frac{J^m K^n}{m!n!} \kappa_{m,n} \right] . \tag{A.12}$$

Taking $J = -1$ and considering K-derivatives of the above (then setting $K = 0$) it is straightforward to show that

$$\left\langle \delta^n e^{-\tau} \right\rangle = F_0 \sum_m \frac{(-1)^m}{m!} \kappa_{m,n} \quad , \quad F_0 = \exp\left[\sum_m \frac{(-1)^m}{m!} \kappa_{m,0} \right] . \tag{A.13}$$

If we use this formula with the normal bias expansions for δ_Q and $\delta\tau$ the form of $\xi_\times$ quickly follows.

B   Eulerian perturbation theory calculation

In this appendix we give details of the redshift-space τ cumulants at one-loop order in (standard) Eulerian perturbation theory. We include this calculation because the formalism may be familiar to many readers, and to make connection with similar calculations within this framework.

At linear order the second cumulant gives the usual super-cluster infall effect which in Fourier space takes $P \to (b_1 + f\mu^2)^2 P$ [69]. Beyond linear theory the second cumulant remains equivalent to the redshift-space correlation function of a biased tracer, so we will not repeat the expression here but refer readers to e.g. ref. [50] and the previous appendix for the expression with the same biasing convention. We note that some of these loop contributions are proportional to the linear $P(k)$ on large scales, modifying the amplitude of this term by factors depending upon small-scale physics. Also, at one-loop order any correlator involving four powers of τ must be disconnected, so the fourth cumulant vanishes. This leaves us with the third cumulant, $\left\langle \delta\tau_1 \delta\tau_2^2 \right\rangle$, which we now compute. Our focus will be on the terms most sensitive to small-scale physics, the so-called UV-sensitive terms.

---

[6]For this section we find it easier to use cumulants of τ rather than $\delta\tau$ as it leads to slightly more compact expressions. The two differ only by powers of τ_0 and a constant for the first cumulant.

There are two contributions to the third cumulant quadratic in the linear power spectrum:

$$\left\langle \delta\tau_1 \delta\tau_2^2 \right\rangle = 2 \left\langle \delta\tau_1^{(1)} \delta\tau_2^{(1)} \delta\tau_2^{(2)} \right\rangle + \left\langle \delta\tau_1^{(2)} \delta\tau_2^{(1)} \delta\tau_2^{(1)} \right\rangle. \tag{B.1}$$

The first is

$$\mathrm{FT} \left\langle \delta\tau_s^{(1)}(\mathbf{r}_1)\, \delta\tau_s^{(1)} \delta\tau_s^{(2)}(\mathbf{r}_2) \right\rangle = 2 Z_1(\mathbf{k}) P(k) \int \frac{d^3\mathbf{p}}{(2\pi)^3}\, Z_1(\mathbf{p}) Z_2(\mathbf{k}, \mathbf{p}) P(p). \tag{B.2}$$

The first and second order redshift space kernels are given by [48]

$$Z_1(\mathbf{k}_1) = b_1 + f\mu_1^2 \tag{B.3}$$

and

$$Z_2(\mathbf{k}_1, \mathbf{k}_2) = b_1 F_2(\mathbf{k}_1, \mathbf{k}_2) + f\mu_k^2 G_2(\mathbf{k}_1, \mathbf{k}_2) + \frac{b_2}{2} + b_s \left(\frac{(\mathbf{k}_1 \cdot \mathbf{k}_2)^2}{k_1^2 k_2^2} - \frac{1}{3} \right)$$
$$+ \frac{fk\mu_k}{2} \left[\frac{\mu_1}{k_1}(b_1 + f\mu_2^2) + \frac{\mu_2}{k_2}(b_1 + f\mu_1^2) \right] \tag{B.4}$$

where $\mu_k = \hat{k} \cdot \hat{n}$, $\mathbf{k} = \mathbf{k}_1 + \mathbf{k}_2$ and F_2 and G_2 are the standard density and velocity kernels in Eulerian perturbation theory [48]:

$$F_2(\mathbf{k}_1, \mathbf{k}_2) = \frac{5}{7} + \frac{1}{2} \frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{k_1 k_2} \left(\frac{k_1}{k_2} + \frac{k_2}{k_1} \right) + \frac{2}{7} \frac{(\mathbf{k}_1 \cdot \mathbf{k}_2)^2}{k_1^2 k_2^2} \tag{B.5}$$

$$G_2(\mathbf{k}_1, \mathbf{k}_2) = \frac{3}{7} + \frac{1}{2} \frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{k_1 k_2} \left(\frac{k_1}{k_2} + \frac{k_2}{k_1} \right) + \frac{4}{7} \frac{(\mathbf{k}_1 \cdot \mathbf{k}_2)^2}{k_1^2 k_2^2} \tag{B.6}$$

To isolate the pieces sensitive to small-scale physics, let us take the $\mathbf{k} \to 0$ limit of the integrand in Equation B.2 (or more correctly consider $k \ll p$). Dropping the terms that are odd in $\hat{p}$ we get

$$Z_2(\mathbf{k}, \mathbf{p}) \to \left[b_1 \left(\frac{5}{7} + \frac{2}{7} [\hat{p} \cdot \hat{k}]^2 \right) + \left(\frac{3}{7} + \frac{4}{7}[\hat{p} \cdot \hat{k}]^2 + \frac{b_1}{2} + \frac{f\mu_k^2}{2} \right) f\mu_p^2 + \frac{b_2}{2} + b_s \left([\hat{p} \cdot \hat{k}]^2 - \frac{1}{3} \right) \right].$$

Using the angular averages

$$\int \frac{d\Omega_p}{4\pi} (\hat{k} \cdot \hat{p})^2 = \frac{1}{3}, \quad \int \frac{d\Omega_p}{4\pi} (\hat{k} \cdot \hat{p})^2 (\hat{n} \cdot \hat{p})^2 = \frac{1 + 2\mu_k^2}{15}, \quad \int \frac{d\Omega_p}{4\pi} (\hat{k} \cdot \hat{p})^2 (\hat{n} \cdot \hat{p})^4 = \frac{1 + 4\mu_k^2}{35} \tag{B.7}$$

Equation (B.2) then becomes

$$2\sigma^2 (b_1 + f\mu_k^2) \left(\left[\frac{5}{7} \left(b_1 + \frac{f}{3} \right) + \frac{2}{7} \left(\frac{b_1}{3} + \frac{1 + 2\mu_k^2}{15} f \right) + \frac{b_2}{2} \right] b_1 \right.$$
$$+ \left[\frac{3}{7} \left(\frac{b_1}{3} + \frac{f}{5} \right) + \frac{4}{7} \left(\frac{1 + 2\mu_k^2}{15} b_1 + \frac{1 + 4\mu_k^2}{35} f \right) + \frac{b_2}{6} + \left(\frac{1 + 2\mu_k^2}{15} - \frac{1}{9} \right) b_s \right] f$$
$$\left. + \frac{f}{2} (b_1 + f\mu_k^2) \left(\frac{b_1}{3} + \frac{f}{5} \right) \right) P(k) \tag{B.8}$$

where

$$\sigma^2 = \int \frac{p^2 \, dp}{2\pi^2} P(p) \quad . \tag{B.9}$$

As $k \to 0$ this is of the form

$$2 \left(A_0 + A_2\mu^2 \right) (b_1 + f\mu^2) \, \sigma^2 \, P(k) \tag{B.10}$$

with

$$A_0 = \frac{17b_1^2}{21} + \frac{46b_1 f}{105} + \frac{b_1^2 f}{6} + \frac{5f^2}{49} + \frac{b_1 f^2}{10} + \frac{b_1 b_2}{2} + \frac{fb_2}{6} - \frac{2fb_s}{45} \tag{B.11}$$

$$A_2 = \frac{4b_1 f}{35} + \frac{16f^2}{245} + \frac{b_1 f^2}{6} + \frac{f^3}{10} + \frac{2fb_s}{15} \quad . \tag{B.12}$$

Note that this has similar angular structure to the linear theory result, but the coefficients of the powers of μ are different and depend upon large momenta in the integral over P, i.e. on small-scale physics.

Similarly, the second contribution is given by

$$\mathrm{FT} \left\langle \tau^{(2)} | \tau^{(1)} \tau^{(1)} \right\rangle = 2 \int \frac{d^3 p}{(2\pi)^3} \, Z_2(p, k-p) Z_1(-p) Z_1(p-k) P(p) P(k-p)$$

$$\to \text{``const}_0\text{''} + \text{``const}_2^{(0)}\text{''} \frac{k^2}{k_\star^2} + \text{``const}_2^{(2)}\text{''} \mu^2 \frac{k^2}{k_\star^2} + \ldots \quad (k \to 0) \tag{B.13}$$

Taking the same limit in this equation simply produces a constant at leading order in $\mathbf{k}$, which leads to a δ-function in configuration space.

C General structure in the large-scale limit

In this appendix we fill in the technical details of the argument in Section 4.4 and consider in general terms the effects of loop contributions beyond those enumerated in the previous appendix. We shall focus on those that renormalize the linear theory predictions while commenting briefly at the end on contributions beyond it.

C.1 Renormalizing linear theory

As in the main text let us write the order n contribution to the $\delta\tau^m$ as

$$\int_{p_1 \ldots p_n} K_n^{(m)}(p_1, \ldots, p_n) \, \delta_0(p_1) \ldots \delta_0(p_n) \, (2\pi)^3 \delta^D(k - \sum p_i). \tag{C.1}$$

The contributions of interest come from those linking only one power of the linear field between two points (Fig. 4)

$$\int_{k, p_i} \int_{k', q_i} K_n^{(a)}(k, p_1, \ldots, p_{n-1}) K_m^{(b)}(k', q_1, \ldots, q_{m-1})$$

$$\left\langle \delta_0(k)\delta_0(p_1)\ldots\delta_0(p_{n-1}) \, (2\pi)^3 \delta^D(\sum p_i) \, \delta_0(k')\delta_0(q_1)\ldots\delta_0(q_{n-1}) \, (2\pi)^3 \delta^D(\sum q_i) \right\rangle.$$

Contracting the k, k' modes yields the contribution

$$(2\pi)^3 \delta^D(k + k') \; \kappa_n^a(\mathbf{k}) \; \kappa_m^b(\mathbf{k}) \; P(k) \tag{C.2}$$

where we have defined

$$\kappa_n^{(a)} = \int_{p_i} K_n^{(a)}(k, p_1, ..., p_{n-1}) \left\langle \delta_0(p_1)...\delta_0(p_{n-1}) \; (2\pi)^3 \delta^D(\textstyle\sum p_i) \right\rangle \tag{C.3}$$

which can depend on μ at leading order in k as evident in our explicit calculation for $a, b = 1, 2$. That these contributions can all be so factored implies that the total contribution to the correlation function logarithm can be summarized by a single renormalized linear bias factor

$$\text{FT}\{\ln(1 + \xi_{FF})\} \supset \sum_{a,b} \kappa^{(a)} \kappa^{(b)} P(k) = \left(\sum_a \kappa^{(a)} \right)^2 P(k) \rightarrow b^2(\mu) P(k) \quad \text{as } k \rightarrow 0 \tag{C.4}$$

where we have defined $\kappa^{(a)} = \sum_n \kappa_n^{(a)}$.

Let us now consider the structure of $b(\mu)$ by inspecting the large scale ($k \rightarrow 0$) limit. To do so we expand the kernels $K_n^{(a)}$ in k to consider the $\mathcal{O}(k^0)$ contributions. Generically this will take the form

$$\lim_{k \to 0} K_n^{(a)}(k, p_1, ..., p_{n-1}) = a^{(0)} + a_{ij}^{(2)} \hat{k}_i \hat{k}_j + a_{ijkl}^{(4)} \hat{k}_i \hat{k}_j \hat{k}_k \hat{k}_l + ... \tag{C.5}$$

The coefficients $a^{(n)}$ come from integrating over the p_i modes and, by the symmetries of the problem, must be factorable into products of the Kronecker δ and the line-of-sight direction $\hat{n}$. They are defined to be order $\mathcal{O}(k^0)$ so are the relevant contributions in the large scale limit. Indeed we can see that the integrals in Equation B.7 fall into this form and truncate at $a_{ij}^{(2)}$. The μ dependence of $b(\mu)$ therefore depends on at what order in $\hat{k}$ this expansion truncates. As a first simplifying step we note that $K_n^{(a)}$ are products of redshift-space densities in configuration space, i.e. convolutions of n^{th} order redshift-space kernels in Fourier space, that is

$$K_n^{(a)}(k, p_1, ..., p_n) = Z_{n_1}(k, p_1, ..., p_{n_1-1}) \; Z_{n_2}(p_{n_1}, ...) \; Z_{n_3}(p_{n_1+n_2}, ...) \; ... \tag{C.6}$$

with conservation of momentum enforced by a delta function when integrating over k, p_i. Since k only appears in one of the Z_{n_i} it is sufficient to look at these independently.

In fact, the redshift-space kernels can themselves be decomposed into constituent density, velocity and bias kernels. Specifically we can write [48]

$$\delta_s(\mathbf{k}) = \sum_{n=0}^{\infty} \int_{\mathbf{k}_i} \left[\delta(\mathbf{k}_1) + f\mu_1^2 \theta(\mathbf{k}_1) \right] \frac{(f\mu k)^{n-1}}{(n-1)!} \frac{\mu_2}{k_2} \theta(\mathbf{k}_2)...\frac{\mu_n}{k_n} \theta(\mathbf{k}_n) (2\pi)^3 \delta^D(\mathbf{k} - \sum_i \mathbf{k}_i). \tag{C.7}$$

From this structure we can deduce that for each power of δ or θ either (a) they are linear modes with momentum $\mathbf{k}_i$, in which case their $\hat{k}_i$ dependence can either come from a μ_i^2 or μ_i/k or (b) they are nonlinear modes with total momentum $\mathbf{k}_i$ dominated by short modes with one long mode $\mathbf{p}$ hidden inside, in which case the angular dependence on the long mode depends solely on the nonlinear kernel of δ or θ since μ_i doesn't depend on $\mathbf{p}$ to lower order.

Indeed, the nonlinear bias and matter kernels also satisfy this structure. Let us begin with the latter as an instructive example. The nonlinear matter density and velocity kernels F_n and G_n obey recursion relations that take the schematic form [48]

$$F_n(\mathbf{q}_1,...,\mathbf{q}_n) \sim G_m(\mathbf{q}_1,...,\mathbf{q}_m)\Big(\alpha(\mathbf{k}_1,\mathbf{k}_2)F_{n-m}(\mathbf{q}_{m+1},...,\mathbf{q}_n) + \beta(\mathbf{k}_1,\mathbf{k}_2)G_{n-m}(\mathbf{q}_{m+1},...,\mathbf{q}_n)\Big)$$

where we have omitted the implied sum over m and numerical factors, and defined [48]

$$\alpha(\mathbf{k}_1,\mathbf{k}_2) = \frac{\mathbf{k}_{12}\cdot\mathbf{k}_1}{k_1^2} \quad , \quad \beta(\mathbf{k}_1,\mathbf{k}_2) = \frac{k_{12}^2(\mathbf{k}_1\cdot\mathbf{k}_2)}{2k_1^2 k_2^2} \quad , \tag{C.8}$$

with $\mathbf{k}_1 = \mathbf{q}_1 + ... + \mathbf{q}_m$, $\mathbf{k}_2 = \mathbf{q}_{m+1} + ... + \mathbf{q}_n$ and $\mathbf{k}_{12} = \mathbf{k}_1 + \mathbf{k}_2$. The velocity kernels, G_n, have the same form. We are interested in the limit where one of the momenta ($\mathbf{p}$) is much smaller than the others, such that it can be ignored when summed with other momenta, $\mathbf{q}_i$. In this limit we see that the leading contribution of $\mathbf{p}$ to the denominators of α and β are at most order $\mathcal{O}(p^2)$ when $m = 1, n-1$, and $\mathcal{O}(p^0)$ otherwise. However in both of these cases they multiply $F_1, G_1 = 1$. Thus in general the terms contributing to kernels at order n have denominators $1/p^2$ coming from α, β or from the F's and G's themselves, which do not produce additional orders when multiplied because they carry different momenta. This implies that each term has at most p^2 in the denominator. This then implies that expansion as in Equation C.5 must terminate at $p_i p_j/p^2 = \hat{p}_i\hat{p}_j$. Recently, [53] showed that all real-space bias operators can be expressed at each order by combining lower-order operators X, Y though

$$XY, \ \left(\partial_i X\right)\left(\frac{\partial_i}{\partial^2}Y\right), \left(\frac{\partial_i\partial_j}{\partial^2}X\right)\left(\frac{\partial_i\partial_j}{\partial^2}Y\right) \tag{C.9}$$

e.g. the quadratic shear bias can be formed by applying the third combination to the linear operator δ. From a similar argument about joining operators that we used in the redshift-space formula above we can conclude that no real-space bias operator can produce a dependence on long modes more complicated than second-order in $\hat{k}$ using just the structure of dynamical nonlinearities in structure formation.

C.2 Two and more soft modes

We can also extend the arguments of the previous subsection to the response of the flux field to more than one long-wavelength mode. To do so it is necessary to consider the n^{th} order kernel above with more than one long-wavelength mode, i.e.

$$\kappa_n^{(a)}(\mathbf{k}_1,\mathbf{k}_2) = \int_{p_i} K_n^{(a)}(k_1, k_2,, p_1, ..., p_{n-2})\left\langle \delta_0(p_1)...\delta_0(p_{n-m})\ (2\pi)^3\delta^D(\sum p_i)\right\rangle. \tag{C.10}$$

We can consider two cases. In one the two soft modes can eventually be reduced to two different redshift-space kernels, i.e. $Z_{n_1}(\mathbf{k}_1, p_1, ...)Z_{n_2}(\mathbf{k}_2, p_2, ...)$ in which case previous arguments imply factorizable forms like

$$\delta(\mathbf{p})\ \frac{q_z q_z}{q^2}\delta(\mathbf{q}), \ \frac{p_z p_z}{p^2}\delta(\mathbf{p})\ \delta(\mathbf{q}), \ \frac{p_z p_z}{p^2}\delta(\mathbf{p})\ \frac{q_z q_z}{q^2}\delta(\mathbf{q}). \tag{C.11}$$

However, these two linear modes can also be coupled by nonlinearities. For this argument it will be useful to write down the general form of anisotropic couplings induced by redshift-space distortions; from Section 3 of ref. [70] we can see that these generally take the form of

$$\alpha^s(\mathbf{k}_1, \mathbf{k}_2) = \frac{k_{1,z}(k_{1,z} + k_{2,z})}{k_1^2} \sim X \ \hat{n}_i \hat{n}_j \frac{\partial_i \partial_j}{\partial^2} Y, \ (\hat{n}_i \partial_i X) \ \hat{n}_j \frac{\partial_j}{\partial^2} Y \quad \text{(C.12)}$$

$$\beta^s(\mathbf{k}_1, \mathbf{k}_2) = \frac{(k_1 + k_2)^2 k_{1,z} k_{2,z}}{2 k_1^2 k_2^2} \sim (\hat{n}_i \partial_i X) \ \hat{n}_j \frac{\partial_j}{\partial^2} Y, \ \left(\frac{\partial_i \partial_{\hat{n}}}{\partial^2} X\right)\left(\frac{\partial_i \partial_{\hat{n}}}{\partial^2} Y\right). \quad \text{(C.13)}$$

There are essentially two cases of this—they can be couple by real-space nonlinearities, with anisotropy introduced "in front" e.g. $\mu_{p+q}^2 G_n(p, q, ...)$ or directly coupled in an anisotropic way i.e. $\hat{p}_i \hat{p}_z \delta(\mathbf{p}) \ \hat{q}_i \hat{q}_z \delta(\mathbf{q})$. At quadratic order this recovers the terms in ref. [34], though we have neglected displacement effects due to operators in the form of ∂_i / ∂^2, whose coefficients can be fixed by equivalence-principle considerations [53]. We intend to return to this topic in a future work.

References

[1] J. A. Peacock, *Cosmological Physics.* Cambridge University Press, 1999.

[2] A. A. Meiksin, *The physics of the intergalactic medium*, *Reviews of Modern Physics* **81** (Oct., 2009) 1405–1469, [arXiv:0711.3358].

[3] M. McQuinn, *The Evolution of the Intergalactic Medium*, *Ann. Rev. Astron.& Astrophys.* **54** (Sept., 2016) 313–362, [arXiv:1512.00086].

[4] R. A. C. Croft, W. Hu, and R. Davé, *Cosmological Limits on the Neutrino Mass from the Lyα Forest*, *PhysRevLett* **83** (Aug., 1999) 1092–1095, [astro-ph/9903335].

[5] M. Viel, J. Weller, and M. G. Haehnelt, *Constraints on the primordial power spectrum from high-resolution Lyman α forest spectra and WMAP*, *MNRAS* **355** (Dec., 2004) L23–L28, [astro-ph/0407294].

[6] E. Armengaud, N. Palanque-Delabrouille, C. Yèche, D. J. E. Marsh, and J. Baur, *Constraining the mass of light bosonic dark matter using SDSS Lyman-α forest*, *MNRAS* **471** (Nov., 2017) 4606–4614, [arXiv:1703.09126].

[7] J. Baur, N. Palanque-Delabrouille, C. Yèche, A. Boyarsky, O. Ruchayskiy, É. Armengaud, and J. Lesgourgues, *Constraints from Ly-α forests on non-thermal dark matter including resonantly-produced sterile neutrinos*, *JCAP* **2017** (Dec., 2017) 013, [arXiv:1706.03118].

[8] R. Murgia, V. Iršič, and M. Viel, *Novel constraints on noncold, nonthermal dark matter from Lyman-α forest data*, *PRD* **98** (Oct., 2018) 083540, [arXiv:1806.08371].

[9] N. Palanque-Delabrouille, C. Yèche, N. Schöneberg, J. Lesgourgues, M. Walther, S. Chabanier, and E. Armengaud, *Hints, neutrino bounds, and WDM constraints from SDSS DR14 Lyman-α and Planck full-survey data*, *JCAP* **2020** (Apr., 2020) 038, [arXiv:1911.09073].

[10] W. Enzi, R. Murgia, O. Newton, S. Vegetti, C. Frenk, M. Viel, M. Cautun, C. D. Fassnacht, M. Auger, G. Despali, J. McKean, L. V. E. Koopmans, and M. Lovell, *Joint constraints on thermal relic dark matter from a selection of astrophysical probes*, *arXiv e-prints* (Oct., 2020) arXiv:2010.13802, [arXiv:2010.13802].

[11] M. Viel, G. D. Becker, J. S. Bolton, and M. G. Haehnelt, *Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman-α forest data*, PRD **88** (Aug., 2013) 043502, [arXiv:1306.2314].

[12] V. Iršič, M. Viel, M. G. Haehnelt, J. S. Bolton, S. Cristiani, G. D. Becker, V. D'Odorico, G. Cupani, T.-S. Kim, T. A. M. Berg, S. López, S. Ellison, L. Christensen, K. D. Denney, and G. Worseck, *New constraints on the free-streaming of warm dark matter from intermediate and small scale Lyman-α forest data*, PRD **96** (July, 2017) 023522, [arXiv:1702.01764].

[13] V. Iršič, M. Viel, M. G. Haehnelt, J. S. Bolton, and G. D. Becker, *First Constraints on Fuzzy Dark Matter from Lyman-α Forest Data and Hydrodynamical Simulations*, PhysRevLett **119** (July, 2017) 031302, [arXiv:1703.04683].

[14] K. K. Rogers and H. V. Peiris, *Strong Bound on Canonical Ultralight Axion Dark Matter from the Lyman-Alpha Forest*, PhysRevLett **126** (Feb., 2021) 071302, [arXiv:2007.12705].

[15] A. Slosar, A. Font-Ribera, M. M. Pieri, J. Rich, J.-M. Le Goff, É. Aubourg, J. Brinkmann, N. Busca, B. Carithers, R. Charlassier, M. Cortês, R. Croft, K. S. Dawson, D. Eisenstein, J.-C. Hamilton, S. Ho, K.-G. Lee, R. Lupton, P. McDonald, B. Medolin, D. Muna, J. Miralda-Escudé, A. D. Myers, R. C. Nichol, N. Palanque-Delabrouille, I. Pâris, P. Petitjean, Y. Piškur, E. Rollinde, N. P. Ross, D. J. Schlegel, D. P. Schneider, E. Sheldon, B. A. Weaver, D. H. Weinberg, C. Yeche, and D. G. York, *The Lyman-α forest in three dimensions: measurements of large scale flux correlations from BOSS 1st-year data*, JCAP **2011** (Sept., 2011) 001, [arXiv:1104.5244].

[16] H. du Mas des Bourboux, J. Rich, A. Font-Ribera, V. de Sainte Agathe, J. Farr, T. Etourneau, J.-M. Le Goff, A. Cuceu, C. Balland, J. E. Bautista, M. Blomqvist, J. Brinkmann, J. R. Brownstein, S. Chabanier, E. Chaussidon, K. Dawson, A. X. González-Morales, J. Guy, B. W. Lyke, A. de la Macorra, E.-M. Mueller, A. D. Myers, C. Nitschelm, A. Muñoz Gutiérrez, N. Palanque-Delabrouille, J. Parker, W. J. Percival, I. Pérez-Ràfols, P. Petitjean, M. M. Pieri, C. Ravoux, G. Rossi, D. P. Schneider, H.-J. Seo, A. Slosar, J. Stermer, M. Vivek, C. Yèche, and S. Youles, *The Completed SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations with Lyα Forests*, ApJ **901** (Oct., 2020) 153, [arXiv:2007.08995].

[17] V. Iršič, E. Di Dio, and M. Viel, *Relativistic effects in Lyman-α forest*, JCAP **2016** (Feb., 2016) 051, [arXiv:1510.03436].

[18] F. Lepori, V. Iršič, E. Di Dio, and M. Viel, *The impact of relativistic effects on the 3D Quasar-Lyman-α cross-correlation*, JCAP **2020** (Apr., 2020) 006, [arXiv:1910.06305].

[19] R. A. C. Croft, *Ionizing Radiation Fluctuations and Large-Scale Structure in the Lyα Forest*, ApJ **610** (Aug., 2004) 642–662, [astro-ph/0310890].

[20] A. Meiksin and M. White, *The effects of ultraviolet background correlations on Lyα forest flux statistics*, MNRAS **350** (May, 2004) 1107–1126, [astro-ph/0307289].

[21] M. McQuinn, L. Hernquist, A. Lidz, and M. Zaldarriaga, *The signatures of large-scale temperature and intensity fluctuations in the Lyman α forest*, MNRAS **415** (July, 2011) 977–992, [arXiv:1010.5250].

[22] S. Gontcho A Gontcho, J. Miralda-Escudé, and N. G. Busca, *On the effect of the ionizing background on the Lyα forest autocorrelation function*, MNRAS **442** (July, 2014) 187–195, [arXiv:1404.7425].

[23] A. Pontzen, *Scale-dependent bias in the baryonic-acoustic-oscillation-scale intergalactic neutral hydrogen*, PRD **89** (Apr., 2014) 083010, [arXiv:1402.0506].

[24] A. Pontzen, S. Bird, H. Peiris, and L. Verde, *Constraints on Ionizing Photon Production from the Large-scale Lyα Forest*, *ApJL* **792** (Sept., 2014) L34, [arXiv:1407.6367].

[25] A. Meiksin and M. McQuinn, *Time-dependent fluctuations in the metagalactic photoionization background*, *MNRAS* **482** (Feb., 2019) 4777–4797, [arXiv:1809.08645].

[26] A. Borde, N. Palanque-Delabrouille, G. Rossi, M. Viel, J. S. Bolton, C. Yèche, J.-M. LeGoff, and J. Rich, *New approach for precise computation of Lyman-α forest power spectrum with hydrodynamical simulations*, *JCAP* **2014** (July, 2014) 005, [arXiv:1401.6472].

[27] A. Arinyo-i-Prats, J. Miralda-Escudé, M. Viel, and R. Cen, *The non-linear power spectrum of the Lyman alpha forest*, *JCAP* **2015** (Dec., 2015) 017, [arXiv:1506.04519].

[28] M. Walther, E. Armengaud, C. Ravoux, N. Palanque-Delabrouille, C. Yèche, and Z. Lukić, *Simulating intergalactic gas for DESI-like small scale Lymanα forest observations*, *JCAP* **2021** (Apr., 2021) 059, [arXiv:2012.04008].

[29] N. Y. Gnedin and L. Hui, *Probing the Universe with the Lyalpha forest - I. Hydrodynamics of the low-density intergalactic medium*, *MNRAS* **296** (May, 1998) 44–55, [astro-ph/9706219].

[30] P. McDonald, J. Miralda-Escudé, M. Rauch, W. L. W. Sargent, T. A. Barlow, R. Cen, and J. P. Ostriker, *The Observed Probability Distribution Function, Power Spectrum, and Correlation Function of the Transmitted Flux in the Lyα Forest*, *ApJ* **543** (Nov., 2000) 1–23, [astro-ph/9911196].

[31] U. Seljak, *Bias, redshift space distortions and primordial nongaussianity of nonlinear transformations: application to Ly-α forest*, *JCAP* **2012** (Mar., 2012) 004, [arXiv:1201.0594].

[32] A. M. Cieplak and A. Slosar, *Towards physics responsible for large-scale Lyman-α forest bias parameters*, *JCAP* **2016** (Mar., 2016) 016, [arXiv:1509.07875].

[33] V. Iršič and M. McQuinn, *Absorber Model: the Halo-like model for the Lyman-α forest*, *JCAP* **2018** (Apr., 2018) 026, [arXiv:1801.02671].

[34] J. J. Givans and C. M. Hirata, *Redshift-space streaming velocity effects on the Lyman-α forest baryon acoustic oscillation scale*, *PRD* **102** (July, 2020) 023515, [arXiv:2002.12296].

[35] M. Garny, T. Konstandin, L. Sagunski, and M. Viel, *Neutrino mass bounds from confronting an effective model with BOSS Lyman-alpha data*, *arXiv e-prints* (Nov., 2020) arXiv:2011.03050, [arXiv:2011.03050].

[36] V. Desjacques, D. Jeong, and F. Schmidt, *The galaxy power spectrum and bispectrum in redshift space*, *JCAP* **2018** (Dec., 2018) 035, [arXiv:1806.04015].

[37] Z. Lukić, C. W. Stark, P. Nugent, M. White, A. A. Meiksin, and A. Almgren, *The Lyman α forest in optically thin hydrodynamical simulations*, *MNRAS* **446** (Feb., 2015) 3697–3724, [arXiv:1406.6361].

[38] V. Desjacques, D. Jeong, and F. Schmidt, *Large-Scale Galaxy Bias*, *ArXiv e-prints* (Nov., 2016) [arXiv:1611.09787].

[39] M. Lewandowski, A. Perko, and L. Senatore, *Analytic prediction of baryonic effects from the EFT of large scale structures*, *JCAP* **2015** (May, 2015) 019, [arXiv:1412.5049].

[40] R. Angulo, M. Fasiello, L. Senatore, and Z. Vlah, *On the statistics of biased tracers in the Effective Field Theory of Large Scale Structures*, *JCAP* **9** (Sept., 2015) 029, [arXiv:1503.08826].

[41] F. Schmidt, *Effect of relative velocity and density perturbations between baryons and dark matter on the clustering of galaxies*, *PRD* **94** (Sept., 2016) 063508, [arXiv:1602.09059].

[42] S.-F. Chen, E. Castorina, and M. White, *Biased tracers of two fluids in the Lagrangian picture*, *JCAP* **2019** (June, 2019) 006, [arXiv:1903.00437].

[43] D. Tseliakhovich and C. Hirata, *Relative velocity of dark matter and baryonic fluids and the formation of the first structures*, *PRD* **82** (Oct., 2010) 083520, [arXiv:1005.2416].

[44] J. Yoo, N. Dalal, and U. Seljak, *Supersonic relative velocity effect on the baryonic acoustic oscillation measurements*, *JCAP* **2011** (July, 2011) 018, [arXiv:1105.3732].

[45] C. Modi, S.-F. Chen, and M. White, *Simulations and symmetries*, *MNRAS* **492** (Mar., 2020) 5754–5763, [arXiv:1910.07097].

[46] A. S. Szalay, *Constraints on the Biasing of Density Fluctuations*, *ApJ* **333** (Oct., 1988) 21.

[47] R. Price, *A useful theorem for nonlinear devices having gaussian inputs*, *IRE Transactions on Information Theory* **4** (1958), no. 2 69–72.

[48] F. Bernardeau, S. Colombi, E. Gaztañaga, and R. Scoccimarro, *Large-scale structure of the Universe and cosmological perturbation theory*, *PhysRep* **367** (Sept., 2002) 1–248, [astro-ph/0112551].

[49] P. McDonald and A. Roy, *Clustering of dark matter tracers: generalizing bias for the coming era of precision LSS*, *JCAP* **8** (Aug., 2009) 020, [arXiv:0902.0991].

[50] S.-F. Chen, Z. Vlah, and M. White, *Consistent Modeling of Velocity Statistics and Redshift-Space Distortions in One-Loop Perturbation Theory*, *JCAP* **07** (2020) 062, [arXiv:2005.00523].

[51] U. Seljak, *Extracting Primordial Non-Gaussianity without Cosmic Variance*, *PhysRevLett* **102** (Jan., 2009) 021302, [arXiv:0807.1770].

[52] A. Cuceu, A. Font-Ribera, B. Joachimi, and S. Nadathur, *Cosmology beyond BAO from the 3D distribution of the Lyman-α forest*, *arXiv e-prints* (Mar., 2021) arXiv:2103.14075, [arXiv:2103.14075].

[53] T. Fujita and Z. Vlah, *Perturbative description of biased tracers using consistency relations of LSS*, *JCAP* **2020** (Oct., 2020) 059, [arXiv:2003.10114].

[54] M. McQuinn and M. White, *On estimating Lyα forest correlations between multiple sightlines*, *MNRAS* **415** (Aug., 2011) 2257–2269, [arXiv:1102.1752].

[55] A. Font-Ribera, P. McDonald, and A. Slosar, *How to estimate the 3D power spectrum of the Lyman-α forest*, *JCAP* **2018** (Jan., 2018) 003, [arXiv:1710.11036].

[56] A. Font-Ribera, J. Miralda-Escudé, E. Arnau, B. Carithers, K.-G. Lee, P. Noterdaeme, I. Pâris, P. Petitjean, J. Rich, E. Rollinde, N. P. Ross, D. P. Schneider, M. White, and D. G. York, *The large-scale cross-correlation of Damped Lyman alpha systems with the Lyman alpha forest: first measurements from BOSS*, *JCAP* **2012** (Nov., 2012) 059, [arXiv:1209.4596].

[57] P. McDonald, *Toward a Measurement of the Cosmological Geometry at z ~2: Predicting Lyα Forest Correlation in Three Dimensions and the Potential of Future Data Sets*, *ApJ* **585** (Mar., 2003) 34–51, [astro-ph/0108064].

[58] A. Slosar, S. Ho, M. White, and T. Louis, *The acoustic peak in the Lyman alpha forest*, *JCAP* **2009** (Oct., 2009) 019, [arXiv:0906.2414].

[59] C.-H. Chuang, F.-S. Kitaura, Y. Liang, A. Font-Ribera, C. Zhao, P. McDonald, and C. Tao, *Linear redshift space distortions for cosmic voids based on galaxies in redshift space*, *PRD* **95** (Mar., 2017) 063528, [arXiv:1605.05352].

[60] S. Masaki, T. Nishimichi, and M. Takada, *Anisotropic separate universe simulations*, *MNRAS* **496** (July, 2020) 483–496, [arXiv:2003.10052].

[61] J. Stücker, A. S. Schmidt, S. D. M. White, F. Schmidt, and O. Hahn, *Measuring the tidal response of structure formation: anisotropic separate universe simulations using TREEPM*, *MNRAS* **503** (May, 2021) 1473–1489, [`arXiv:2003.06427`].

[62] D. J. Eisenstein, H.-J. Seo, and M. White, *On the Robustness of the Acoustic Scale in the Low-Redshift Clustering of Matter*, *ApJ* **664** (Aug., 2007) 660–674, [`astro-ph/0604361`].

[63] M. Crocce and R. Scoccimarro, *Nonlinear evolution of baryon acoustic oscillations*, *PRD* **77** (Jan., 2008) 023533, [`arXiv:0704.2783`].

[64] T. Matsubara, *Resumming cosmological perturbations via the Lagrangian picture: One-loop results in real space and in redshift space*, *PRD* **77** (Mar., 2008) 063530, [`arXiv:0711.2521`].

[65] J. Carlson, B. Reid, and M. White, *Convolution Lagrangian perturbation theory for biased tracers*, *MNRAS* **429** (Feb., 2013) 1674–1685, [`arXiv:1209.0780`].

[66] L. Senatore and M. Zaldarriaga, *The IR-resummed Effective Field Theory of Large Scale Structures*, *JCAP* **2** (Feb., 2015) 13, [`arXiv:1404.5954`].

[67] Z. Vlah, U. Seljak, M. Yat Chu, and Y. Feng, *Perturbation theory, effective field theory, and oscillations in the power spectrum*, *JCAP* **2016** (Mar., 2016) 057, [`arXiv:1509.02120`].

[68] D. Blas, M. Garny, M. M. Ivanov, and S. Sibiryakov, *Time-sliced perturbation theory II: baryon acoustic oscillations and infrared resummation*, *JCAP* **2016** (July, 2016) 028, [`arXiv:1605.02149`].

[69] N. Kaiser, *Clustering in real space and in redshift space*, *MNRAS* **227** (July, 1987) 1–21.

[70] M. M. Ivanov and S. Sibiryakov, *Infrared resummation for biased tracers in redshift space*, *JCAP* **2018** (July, 2018) 053, [`arXiv:1804.05080`].