

UC Berkeley

UC Berkeley Previously Published Works

Title

Accurate predictions from small boxes: variance suppression via the Zel'dovich approximation

Permalink

<https://escholarship.org/uc/item/2pr6p7g8>

Journal

Journal of Cosmology and Astroparticle Physics, 2022(09)

ISSN

1475-7516

Authors

Kokron, Nickolas

Chen, Shi-Fan

White, Martin

et al.

Publication Date

2022-09-01

DOI

10.1088/1475-7516/2022/09/059

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Accurate predictions from small boxes: variance suppression via the Zel'dovich approximation

Nickolas Kokron^{a,b} Shi-Fan Chen^c Martin White^{c,d} Joseph DeRose^d Mark Maus^c

^aKavli Institute for Particle Astrophysics and Cosmology and Department of Physics, Stanford University, Stanford, CA, USA

^bKavli Institute for Particle Astrophysics and Cosmology, SLAC National Accelerator Laboratory, Menlo Park, CA, USA

^cDepartment of Physics, University of California, Berkeley, CA, USA

^dPhysics Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA

E-mail: kokron@stanford.edu

Abstract. Simulations have become an indispensable tool for accurate modelling of observables measured in galaxy surveys, but can be expensive if very large dynamic range in scale is required. We describe how to combine Lagrangian perturbation theory models with N-body simulations to reduce the effects of finite computational volume in the prediction of ensemble average properties in the simulations within the context of control variates. In particular we use the fact that Zel'dovich displacements, computed during initial condition generation for any simulation, correlate strongly with the final density field. Since all the correlators of biased tracers can be computed with arbitrary precision for these displacements, pairing the Zel'dovich ‘simulation’ with the N-body realization allows hundredfold reductions in sample variance for power spectrum or correlation function estimation. Zel'dovich control variates can accurately extend matter or tracer field emulators to larger scales than previously possible, as well as improving measurements of statistics in simulations which are inherently limited to small volumes, such as hydrodynamical simulations of galaxy formation and reionization.

Contents

1	Introduction	1
2	Control variates and variance reduction	4
2.1	Standard control variates	4
2.2	The Zel'dovich approximation as a control variate	5
3	Lagrangian Perturbation Theory Two Ways	7
3.1	Analytic predictions in the Zel'dovich approximation	9
3.2	Grid-based predictions in the Zel'dovich approximation	10
4	Control variates for LPT basis spectra	11
5	Results and discussion	14
5.1	Variance reduction for the matter power spectrum	15
5.2	Reducing the variance of biased tracer samples	16
5.3	Variance reduction for hybrid EFT spectra	18
5.4	Comparison with previous results	21
6	Conclusions	22
A	Basis spectra in the Zel'dovich approximation	24
B	Matching analytic and grid-based LPT	26
C	Challenges in going beyond Zel'dovich	29

1 Introduction

Computer simulations of the formation of cosmic structures, from those that include dark matter only to those solving for full radiative hydrodynamics, have become essential tools in understanding the evolution of the Universe across different cosmic eras. Simulations shed light on complicated non-linear phenomena that evade analytic descriptions. These include: the non-linear dynamics of gravitational collapse, leading to the formation of bound halos; the interplay between gas and light, leading to the formation of stars and eventually galaxies within these halos; the radiative processes underlying reionization and countless more.

In numerical cosmology there exists an eternal tug-of-war between accurately resolving small-scale physics and running simulations at large enough volumes that one has a statistically robust result. For example, suites of simulations run at numerous cosmologies, used for the construction of emulators, must not only balance the dynamic range of their simulations but also include the expense of maintaining accuracy while spanning a large space of cosmological parameters, where a new simulation has to be run for each point in this parameter space. For emulators of large-scale structure designed to accurately predict summary statistics measured in galaxy surveys, the challenges associated with limited volumes become even more severe. Next-generation galaxy surveys will probe unprecedented cosmic volumes, requiring

highly accurate (and precise) predictions at large scales, imposing additional requirements on simulation-based inference tools.

To surmount these challenges, a plethora of techniques have been introduced that try to ameliorate either resolution or volume requirements of simulations. On the resolution front, statistical learning techniques have been leveraged to produce algorithms that generate so-called *super-resolution* simulations which “fill in” information on scales smaller than what the original simulation was capable of resolving in an inexpensive way [1–5]. However, super-resolution techniques are still in their infancy and must be understood more before being used to confront data from these galaxy surveys. On the other hand, the field of statistics has a rich literature on the subject of variance reduction, and these techniques have begun to be imported in a cosmological context in order to relax the requirements on either the volumes of simulations or sheer quantity which must be run. Perhaps the most popular variance reduction tool in cosmology is that of Latin Hypercube Sampling, which allows for efficient sampling of high-dimensional parameter spaces such as the 7(8)-dimensional $w(\nu)$ CDM space over which emulator suites are constructed [6–9]. Two other techniques which have seen widespread adoption are “paired phase” and “fixed amplitude” simulations. Paired-phase, or, “pairing”, involves simulating two Universes whose initial conditions are exactly the same up to a minus sign [10]. The mean of statistics computed in each simulation then has its variance significantly reduced relative to the expectation of Gaussian initial conditions. “Fixing”, on the other hand, involves initializing simulations where the amplitude of density fluctuations follows a Dirac delta distribution as opposed to the standard Rayleigh distribution [11]. Fixing also significantly reduces the large-scale variance of an N -body simulation. The combination of these two techniques, “paired–fixed” simulations, has become the object of significant study in recent years [12–15].

Another technique which has recently seen use in reducing sample variance of simulations in cosmology is the method of control variates [16–18]. Control variates are particularly powerful when correlated, inexpensive surrogates of the statistics one wishes to measure can be produced. The method is well-understood from a theoretical point of view, and the potential variance reduction that can be achieved through its optimal application is proportional to the degree of correlation between the surrogate adopted and the costly statistic whose variance we wish to reduce. The success of control variates is predicated on a thorough understanding of the statistics of the surrogate adopted, including a well-characterized mean, variance and its co-variance with the desired statistic. So far, the surrogate of choice adopted in cosmological applications has been an approximate simulation which, while significantly less expensive than a full N -body simulation to produce, still incurs a substantial computational cost due to the need of simulating hundreds of approximate mocks in order to estimate the mean of the surrogate. Thus, the current bottleneck of applying control variates to cosmological simulations lies in the requirement of simulating large numbers of approximate simulations in order to characterize well its statistical properties.

At sufficiently large scales, analytic descriptions of large-scale structure statistics are highly accurate and arbitrarily precise, and are thus powerful tools to study the large-scale regime of large-volume surveys. While traditionally treated as two disparate ways of studying structure formation, analytic and simulation-based descriptions are inherently linked. For example, every simulation of cosmic structure requires initial conditions which are generated from the aforementioned analytic descriptions, specifically using Lagrangian Perturbation Theory (LPT). LPT is not only a potent framework to describe the statistics of the densities and velocities of biased tracers [19–28], but it has recently also been used in combination with

N -body simulations to produce *hybrid* models of structure formation [29]. Hybrid models (also called hybrid effective field theory, or HEFT) use the tracer–matter connection as specified by LPT with displacements that are accurate at small scales from N -body simulations. Their combination leads to a powerful field-level description of the building blocks of structure formation. Recent applications of Hybrid EFT include: constructing emulators of clustering and lensing which are accurate to $k \simeq 0.6 h\text{Mpc}^{-1}$ [30–32], modelling higher order statistics beyond the power spectrum [33], and characterizing the tracer–matter connection of simulated samples of galaxies [34, 35].

In this work we propose a novel way to leverage the intimate connection between Lagrangian perturbation theory and N -body simulations in order to improve the precision and accuracy of simulation-based predictions significantly. Specifically, we utilize the principle of control variates in order to create realizations of surrogate Universes with the same large-scale noise as those measured in N -body simulations. We use first order LPT (also known as the Zel’dovich Approximation [36]) to analytically predict the means of summary statistics. While the use of control variates in cosmology is not new [16–18], previous work has relied on running ensembles of simulations in order to employ this technique. Our methods, in contrast, are computationally inexpensive. They rely only on data outputs that are a standard part of producing initial conditions for cosmological simulations, and an additional post-processing step that is identical in the full simulation and in the surrogate simulation.

This paper is structured as follows: in § 2 we review the control variates technique, the variance reduction tool we have adopted in this publication. We discuss how this variance reduction is driven by the cross-correlation coefficient between the expensive simulation we wish to improve and the surrogate version we run. We also discuss how even the simplest rendition of LPT, the Zel’dovich approximation (ZA), produces Universes which are highly correlated with the results of full N -body simulations at low redshifts. In § 3 we give a brief overview of LPT and how it is used to predict the summary statistics of biased tracers. We focus on how to predict observables within LPT using both analytic calculations as well as grid-based realizations of the same expressions. In § 4 we then proceed to re-formulate the control variates problem within the context of improving measurements of the basis spectra that make up the two-point statistics of biased tracers in LPT. We present the results of our implementation of control variates for paired ZA realizations in § 5. We apply our methodology to three classes of tests, of increasing complexity. We use $N = 100$ high-resolution N -body simulations, as well as paired ZA realizations, to assess the statistical performance of our technique. We begin by applying control variates to the matter power spectrum, extend to the statistics of a sample of galaxies populated by a halo occupation distribution (HOD) procedure, and conclude by looking at all ten basis spectra that span second order Lagrangian bias models. Additionally, we quantify the reduction in variance from this technique from our ensemble of simulations and consistently find strong reduction in uncertainty for all of the forms of power spectra assessed in this work. The improvements range from a factor of $10\times$ reduction to nearly $1000\times$ depending on the specific basis spectrum in question. We also discuss potential applications of these techniques beyond just improving predictions from cosmological emulation boxes. In § 6 we summarize our results and identify promising future directions.

Our appendices discuss technical aspects of this technique in order to ensure a simulation will benefit from its use. In Appendix A we derive the necessary expressions to model the power spectrum of biased tracers in the Zel’dovich approximation, which is crucial in ensuring the success of our techniques. In Appendix B we discuss the requirements imposed

on simulations in order to ensure their paired ZA realizations will accurately match analytic predictions. In Appendix C we discuss the potential of extending our techniques beyond the Zel’dovich approximation, as well as challenges that must be circumvented before this can be achieved in practice.

2 Control variates and variance reduction

Control variates are a statistical technique employed to reduce the variance of quantities estimated with limited samples of data. They’re applicable when one can create correlated approximate realizations with well-characterized means [37]. In the following sections we give a brief introduction to the theory of control variates and discuss their current applications within a cosmological context, as well as limitations to the technique as it is currently formulated. We will propose that the Zel’dovich approximation can be used as a control variate, and explore its correlation with the matter density field. In cosmology, control variates have recently been applied in the CARPool technique [16–18] as well as reducing the variance in statistics measured from the AbacusSummit [9] suite of simulations [38].

2.1 Standard control variates

Say we are interested in precisely estimating the mean of a simulated observable, \hat{x} . Suppose also that we have a related quantity, the *control variate* \hat{c} , that is significantly cheaper to produce than \hat{x} but is correlated with it. Then, we may define a quantity

$$\hat{y} \equiv \hat{x} - \beta(\hat{c} - \mu_c) \quad , \quad (2.1)$$

with $\mu_c = \langle \hat{c} \rangle$. Taking the expectation value of Eq. 2.1 shows that $\langle \hat{y} \rangle$ is an unbiased estimator of $\langle \hat{x} \rangle$ for any β . The covariance between \hat{c} and \hat{x} can allow $\text{Var}(\hat{y}) < \text{Var}(\hat{x})$ and in fact $\text{Var}(\hat{y})$ can be minimized by taking

$$\beta^* = \frac{\text{Cov}[\hat{x}, \hat{c}]}{\text{Var}[\hat{c}]} \quad . \quad (2.2)$$

This value of β^* leads to a variance reduction

$$\frac{\text{Var}[\hat{y}]}{\text{Var}[\hat{x}]} = 1 - \frac{\text{Cov}^2[\hat{x}, \hat{c}]}{\text{Var}[\hat{x}]\text{Var}[\hat{c}]} = 1 - \rho_{xc}^2 \quad , \quad (2.3)$$

where we’ve defined the cross-correlation coefficient $\rho_{xc} = \text{Cov}[\hat{x}, \hat{c}]/(\sigma_x\sigma_c)$. Therefore, using a highly correlated surrogate that is inexpensive to produce can lead to significant improvements in estimation of such quantities without having to produce many realizations of \hat{x} , which could be computationally expensive. Within the context of cosmology, \hat{x} can be the power spectrum of an N -body simulation [16], or a quantity such as the covariance matrix of that power spectrum [17]. Our interest will be in the case when \hat{x} is a given element of the basis spectrum $P_{ij}(k)$, which we define shortly in § 3. The basis spectra $P_{ij}(k)$ are the building blocks of the statistics of biased tracers within LPT. This scenario includes as special cases the matter power spectrum and the spectrum of a biased tracer (e.g. galaxies).

One of the key limitations in applying control variates to problems in computational cosmology so far has been the computational cost of estimating the mean μ_c of the variate. If improperly estimated, the estimator \hat{y} could become biased and the variance reduction becomes significantly hindered by the variance associated with estimating μ_c . Even when using fast approximate N -body solvers such as FastPM [39] or COLA [40, 41], the computational




Figure 1. Comparison between the distribution of particles from a full N -body simulation and a matched realization using the Zel'dovich approximation, at $z = 0.5$. The panels showing the Zel'dovich and nonlinear density fields δ^{ZA} and δ^N , respectively, range from $[-1, 4\delta_c]$ with $\delta_c = 1.686$ the critical threshold density for spherical collapse. The right-most panel shows the difference between the two fields, with the color bar for the rightmost plot encompassing the range $\delta^{ZA}(\mathbf{x}) - \delta^N(\mathbf{x}) \in [-1, 1]$. The maps shown are roughly $390 \times 390 h^{-1}\text{Mpc}$, projected across roughly $20h^{-1}\text{Mpc}$.

cost incurred from having to run many realizations can quickly limit possible gains from employing the technique. Another limitation comes in estimating the optimal form of β^* , especially in the multivariate problem. This involves computing many realizations of \hat{x} and \hat{c} in order to estimate the full $\text{Cov}[\hat{x}, \hat{c}]$ matrix across multiple modes. When \hat{x} is an observable measured from a cosmological N -body simulation, the number of realizations for a *single* cosmology required to estimate a numerically stable β^* can easily exceed the total number of simulations typically produced for a whole emulation suite.

2.2 The Zel'dovich approximation as a control variate

It is highly desirable to define a control variate which simultaneously correlates significantly with the N -body statistic, but whose analytic properties are known exactly. As alluded to in the introduction, cosmological N -body simulations are often initialized using LPT. The first order solution in LPT is also known as the Zel'dovich approximation [36, 42]. While we will describe LPT and the Zel'dovich approximation in further detail in § 3, the approximation qualitatively states that fluid elements simply move in straight lines through the Universe with a displacement that is proportional to the linear growth factor. The direction of these displacements is seeded by the initial density fluctuations, and is found by solving for the linearized continuity equation in Fourier space. The properties of density fields in LPT are known to arbitrary precision, as they are analytic in nature. It is instructive, then, to consider whether the same theory used to initialize an N -body simulation can be extended to lower redshifts and used as a control variate. In Fig. 1 we show $20 h^{-1}\text{Mpc}$ projections of the Zel'dovich and N -body density fields, which share the same initial conditions, evolved to $z = 0.5$. As one can see, the Zel'dovich approximation captures the structure of the cosmic web in a striking fashion. The initial conditions and non-linear distributions used come from the Quijote [43] ‘high resolution’ (HR) suite of dark matter-only N -body simulations, which have a volume of $V = 1(h^{-1}\text{Gpc})^3$ and $N = 1024^3$ particles. The fundamental grid size of these simulations is given by $L_{\text{cell}} \approx 1h^{-1}\text{Mpc}$.

If the correlation between the Zel'dovich density and the non-linear density extends itself




Figure 2. Comparison of the $z = 0.5$ matter-matter power spectrum between linear theory, simulated ZA, analytic ZA and the full non-linear prediction from an N -body simulation. *Top left:* Matter power spectra, P_{mm} , for the linear density field that seeded the initial conditions, the Zel’dovich-displaced density field, and the fully non-linear density field at $z = 0.5$. *Bottom left:* Cross-correlation coefficients between the non-linear density field and the linear (in blue) and Zel’dovich (in orange) densities, respectively. We see that despite having a similar power spectrum to the linear field, the Zel’dovich realization is much more highly correlated with the non-linear density down to small scales. The cross-correlation coefficient is 90% at $k = 0.64 h\text{Mpc}^{-1}$ and 75% at $k = 1.0 h\text{Mpc}^{-1}$, while the linear density rapidly decorrelates. *Top Right:* Matter power spectrum for the same Zel’dovich and non-linear densities, but showing in blue the result of the analytic Zel’dovich spectrum. *Bottom Right:* Ratio between the two schemes used to make ZA predictions, as well as the ratio between the N -body and Zel’dovich realizations with the same initial conditions. The shaded region denotes $2\text{-}\sigma$ error bars assuming a Gaussian disconnected covariance for the power spectrum. We see that the analytic and realization-based Zel’dovich calculations agree to high accuracy down to scales of $k_{\text{max}} \approx 1 h\text{Mpc}^{-1}$.

beyond the visual correlation of Fig. 1, a promising picture of its use as a control variate arises. The Zel’dovich approximation is a surrogate for structure formation that is typically available for all simulations and whose analytic results are known exactly. The basis spectra $P_{ij}(k)$ that are the building blocks for clustering and lensing statistics can be computed with no approximations within Zel’dovich. At the same time, Zel’dovich “realizations” of full N -body simulations are extremely inexpensive to generate (usually using codes that are already part of the N -body pipeline), and form a field-level description for *any* summary statistic that one wishes to measure. Thus, if one can ensure that the two different approaches (analytic and grid-based) to calculating statistics in the Zel’dovich approximation are in numerical agreement, then the control variates technique can be employed with negligible computational overhead to greatly improve the precision of measurements of non-linear basis spectra.

A more quantitative comparison elucidates the structure of this correlation. Specifically, we compute the cross-correlation coefficient between the density fields

$$\rho_{xc} = \frac{P^{ZA,N}(k)}{\sqrt{P^{N,N}(k)P^{ZA,ZA}(k)}}, \quad (2.4)$$

as well as the power spectra of these fields, and show them in the bottom left panel of Fig. 2. We observe a strong correlation coefficient between the Zel’dovich and nonlinear densities, which is still at 75% for $k = 1 h\text{Mpc}^{-1}$. The linear density distribution, by comparison, reaches 75% correlation with the nonlinear density at $k \approx 0.15 h\text{Mpc}^{-1}$.

The fact that the Zel’dovich approximation produces distributions of matter that correlate highly with the non-linear matter distribution is not new. Indeed, this has been explored previously in many publications [44–48]. That the cross-correlation with the Zel’dovich approximation is high while the linear theory prediction is low, despite having similar two-point statistics, is due to the nature of the comparison made. In a Universe with only linear displacements, fluid elements will move across distances of roughly Σ_Ψ , where $\Sigma_\Psi^2 = \int dk P(k)/6\pi^2$. Comparing the fields before and after these displacements leads to decorrelation on scales smaller than Σ_Ψ , as dark matter particles have been displaced by this distance from their initial conditions. The chief impact of gravitational non-linearities is to introduce accelerations and slight deviations in the motions of fluid elements on scales smaller than this dispersion, and thus the Zel’dovich approximation describes most of this displacement. Curiously, it has also been noted that higher-order Lagrangian perturbation theories correlate *more poorly* with the density field at lower redshifts, despite describing the matter power spectrum more accurately.

In this section we reviewed the method of control variates and its applications so far in cosmology. We pointed out the expense of estimating $\langle \hat{c} \rangle$ and β when \hat{c} is a fast N -body simulation as one of the main challenges of deploying this technique today. We proceeded to suggest that the Zel’dovich approximation, whose ingredients are generated when initializing any N -body simulation, could be a powerful control variate due to its high correlation with the non-linear density field.

We now proceed with a discussion of computing power spectra in LPT to § 3. We leave Figures 1 and 2 as tantalizing figures that point to the power of the Zel’dovich approximation as a control variate. We will formulate in § 4 the control variate problem for the basis spectra of biased tracers, which will include as special cases both the power spectrum of matter density fluctuations and the statistics of *any* biased tracer measured from simulations.

3 Lagrangian Perturbation Theory Two Ways

In the Lagrangian picture of structure formation, the density contrast field is calculated from the movement of Lagrangian fluid elements across cosmic time. That is, particles located at a position \mathbf{q} at initial conditions are advected by a displacement Ψ to their final position

$$\mathbf{x}(a) = \mathbf{q} + \Psi(\mathbf{q}, a). \quad (3.1)$$

where $a = (1+z)^{-1}$ is the scale factor. If the initial distribution of densities is approximately uniform, $\rho(\mathbf{q}) \approx \bar{\rho}$, then at late times the mapping from Eulerian to Lagrangian coordinates gives an evolved density distribution from the continuity equation

$$1 + \delta(\mathbf{x}, a) = \int d^3q \delta^D(\mathbf{x} - \mathbf{q} - \Psi(\mathbf{q}, a)), \quad (3.2)$$

where δ^D is the Dirac delta function. In the presence of biased tracers, Eqn. 3.2 is extended by including a functional $F[\delta(\mathbf{q})]$ that specifies the tracer–matter connection at early times. The components of this functional are similarly advected, and the late-time tracer density $\delta_t(\mathbf{x}, a)$ is given by

$$1 + \delta_t(\mathbf{x}, a) = \int d^3q F[\delta(\mathbf{q})] \delta^D(\mathbf{x} - \mathbf{q} - \Psi(\mathbf{q}, a)). \quad (3.3)$$

The functional $F[\delta(\mathbf{q})]$ is normally expanded to second order as [20, 24, 49]

$$F[\delta(\mathbf{q})] \approx 1 + b_\delta \delta(\mathbf{q}) + b_{\delta^2} (\delta^2(\mathbf{q}) - \langle \delta^2 \rangle) + b_s (s^2(\mathbf{q}) - \langle s^2 \rangle) + b_{\nabla^2} \nabla^2 \delta(\mathbf{q}) + \dots \quad (3.4)$$

where $s^2(\mathbf{q}) = s_{ij}(\mathbf{q}) s^{ij}(\mathbf{q})$ is the tidal field strength, and $s_{ij}(\mathbf{q})$ is the tidal tensor defined as

$$s_{ij}(\mathbf{q}) = \left(\frac{\partial_i \partial_j}{\partial^2} - \frac{\delta_{ij}}{3} \right) \delta(\mathbf{q}). \quad (3.5)$$

Advecting each component of Eqn. 3.4 individually leads to the description of the late-time tracer field in terms of a set of advected operators

$$\delta_t(\mathbf{x}) = \delta_m(\mathbf{x}) + b_\delta \mathcal{O}_\delta(\mathbf{x}) + b_{\delta^2} \mathcal{O}_{\delta^2}(\mathbf{x}) + b_s \mathcal{O}_{s^2}(\mathbf{x}) + b_{\nabla^2} \mathcal{O}_{\nabla^2 \delta}(\mathbf{x}) + \epsilon(\mathbf{x}) \quad (3.6)$$

$$= \sum_{i \in \{m, \delta, \delta^2, \dots\}} b_i \mathcal{O}_i + \epsilon(\mathbf{x}), \quad (3.7)$$

where \mathcal{O}_i are the advected operators and $\epsilon(\mathbf{x})$ is a stochastic field which quantifies both inherent randomness in the process of tracer formation as well as the impact of neglected higher-order operators. Note that in this notation $\mathcal{O}_m(\mathbf{x}) = \delta_m(\mathbf{x})$ and $b_m \equiv 1$. We further follow the convention in the LPT literature and use as aliases $b_1 \equiv b_\delta$ and $b_2 \equiv b_{\delta^2}$. Tracer power spectra will receive contributions from correlations between the advected operators that compose the functional F . This decomposition will have the form

$$P^{tt}(k) = \sum_{i, j \in \{m, \delta, \delta^2, \dots\}} b_i b_j P_{ij}(k), \quad (3.8)$$

$$P^{tm}(k) = \sum_{j \in \{m, \delta, \delta^2, \dots\}} b_j P_{mj}(k), \quad (3.9)$$

where P_{ij} is the cross-spectrum $\langle \mathcal{O}_i \mathcal{O}_j \rangle$ and P_{mj} are cross-correlations between the matter density field and the bias operators¹. Since the matter field is obtained by weighting each particle with weight 1, the matter power spectrum (P_{mm}) is often written as P_{11} in the language of Lagrangian bias. For this work, we'll neglect the $\nabla^2 \delta(\mathbf{q})$ operator in our analysis and only work with $\{b_1, b_2, b_s\}$ as bias parameters. The reasons are two-fold; we expect that basis spectra of the type $P_{\nabla^2 \delta X}$ will scale as $-k^2 P_{mX}$ analytically [30], and numerically realizing the $\nabla^2 \delta(\mathbf{q})$ fields is challenging due to the extreme sensitivity of the field to small scales (high k). More careful numerical studies of the $\nabla^2 \delta(\mathbf{q})$ field in the context of hybrid EFT models are of great interest and will be explored in future work.

In Lagrangian Perturbation Theory, the displacements Ψ whose correlations serve as input for predictions of the spectra P_{ij} are computed order-to-order in perturbation theory. To third order, this expansion is commonly written as [51]

$$\Psi^{3\text{LPT}}(\mathbf{q}, a) \approx D(a) \Psi^{(1)}(\mathbf{q}) + D^{(2)}(a) \Psi^{(2)}(\mathbf{q}) + D^{(3)}(a) \Psi^{(3)}(\mathbf{q}), \quad (3.10)$$

¹In the main text we neglect the effect of massive neutrinos, which affect the clustering of matter and galaxies in distinct ways — while neutrinos contribute to the matter field, to a very good approximation galaxies trace the baryon-cold dark matter field δ_{cb} [50], which clusters on small scales (unlike the neutrinos, δ_ν). In order to accurately capture this effect we need to distinguish between two “m” fields: $\delta_{cb}(\mathbf{x})$, the overdensity of cb particles obtained by advecting them via N-body displacements, and $\delta_m(\mathbf{x})$, the nonlinear matter field obtained by mass weighting the cb and ν particles. All of the bias operators should be built from the initial $\delta_{cb}(\mathbf{q})$. Then we simply swap in these two versions of the “m” term in the above equations, e.g. for the cross spectrum $P^{tm}(k) = \sum_j b_j P_{mj}(k)$ where $j \in \{cb, \delta_{cb}, \delta_{cb}^2, \dots\}$.

where $D^{(n)}(a)$ is the n -th order solution to the growth factor in perturbation theory [52]. At any given order (n), $\Psi^{(n)}$ can be computed from suitable convolutions of n powers of the first-order solution to displacements, also known as the Zel'dovich approximation. Today, efficient codes exist that implement the analytic equations for one-loop (combined fourth-order) power spectra in a numerically efficient manner, such as `velocileptors`² [25], and also codes that compute higher-order initial conditions for N -body simulations by numerically evaluating the displacements in Eqn. 3.10 such as `monofonIC` [51] and `LEFTfield` [48]³. For the rest of this text we will focus on only the Zel'dovich approximation, and defer a discussion of implementing our methodology for higher order Lagrangian Perturbation Theory in Appendix C.

3.1 Analytic predictions in the Zel'dovich approximation

In the first order solution to Lagrangian Perturbation Theory, the Zel'dovich approximation, fluid elements in the Universe propagate in straight lines, with a direction set by the potential sourced by the initial matter distribution. These displacements are obtained by solving the linearized continuity equation, and read

$$\Psi^{\text{ZA}}(\mathbf{q}, a) = D(a) \int \frac{d^3k}{(2\pi)^3} e^{i\mathbf{k}\cdot\mathbf{q}} \frac{i\mathbf{k}}{k^2} \delta(\mathbf{k}). \quad (3.11)$$

The late-time tracer density, in Fourier space, is then given by the Fourier transform of Eqn. 3.2

$$\delta_t(\mathbf{k}) = \int d^3q e^{i\mathbf{k}\cdot\mathbf{q}} \left[F[\delta(\mathbf{q})] e^{i\mathbf{k}\cdot\Psi(\mathbf{q})} - 1 \right] \quad (3.12)$$

$$\equiv \sum_{i \in \{m, \delta, \delta^2, \dots\}} b_i \mathcal{O}_i(\mathbf{k}), \quad (3.13)$$

where we've defined the *advected* operators $\mathcal{O}_i(\mathbf{k})$

$$\mathcal{O}_i(\mathbf{k}) \equiv \int d^3q e^{i\mathbf{k}\cdot(\mathbf{q} + \Psi(\mathbf{q}))} F_i(\mathbf{q}), \quad (3.14)$$

and $F_i(\mathbf{q})$ is a Lagrangian element of the functional that composes Eqn. 3.4. The right-most term of Eqn. 3.12 can be neglected, as it only contributes when $\mathbf{k} = 0$. The tracer-tracer power spectrum may then be written as [20, 21, 53]

$$P^{tt}(k) = \sum_{i,j} b_i b_j \langle \mathcal{O}_i(\mathbf{k}) \mathcal{O}_j^*(\mathbf{k}) \rangle \quad (3.15)$$

$$= \sum_{i,j} b_i b_j \underbrace{\int d^3q e^{i\mathbf{k}\cdot\mathbf{q}} \langle F_i(\mathbf{q}_1) F_j(\mathbf{q}_2) e^{i\mathbf{k}\cdot\Delta(\mathbf{q})} \rangle}_{P_{ij}(k)}, \quad (3.16)$$

where we have defined $\Delta(\mathbf{q}) = \Psi(\mathbf{q}_1) - \Psi(\mathbf{q}_2)$ and $\mathbf{q} \equiv \mathbf{q}_1 - \mathbf{q}_2$ ⁴. The expectation value above is computed by using the cumulant expansion as well as defining appropriate source

²Available at <https://github.com/sfschen/velocileptors>.

³Available at <https://bitbucket.org/ohahn/monofonic> and <https://gitlab.mpcdf.mpg.de/leftfield/release/leftfield>, respectively.

⁴The integral for the basis spectra $P_{ij}(k)$ depend only on \mathbf{q} due to translation invariance.

currents for the Lagrangian fields, $F_i(\mathbf{q})$. When working strictly within the Zel’dovich approximation, $\Delta(\mathbf{q})$ is a Gaussian random variable and only the second connected moment remains, significantly simplifying the calculation of the basis spectra $P_{ij}(k)$.

In Appendix. A we provide a detailed derivation of the structure of each basis spectrum, P_{ij} , within the Zel’dovich approximation. We include several new terms that exist beyond the standard “one-loop” (quadratic in the two-point correlation function) order. We implement them in a new code, ZeNBu⁵. We use ZeNBu for the analytic calculations in the rest of this paper unless otherwise specified.

3.2 Grid-based predictions in the Zel’dovich approximation

An alternative way of computing the basis spectra P_{ij} comes from employing *grid-based* Lagrangian Perturbation Theory. Given a fixed realization of the initial Lagrangian density $\delta(\mathbf{q})$, one computes the displacements $\Psi(\mathbf{q}, a)$ by numerically evaluating the Fourier transform in Eqn. 3.11. Each element of the initial grid is then advected using these displacements, and they are used to reconstruct the final density field as in Eqn. 3.2. Grid-based LPT has been well-studied in the context of setting up initial conditions for cosmological N -body simulations, with codes for ZA, 2LPT [54] and 3LPT [51] being widely available. Recent works have also investigated n -LPT at arbitrary order [48], including to assess the fundamental convergence properties of LPT [55]. Other applications of grid-based LPT include studying the fundamental limits of perturbative models [56] and using LPT as forward models to reconstruct the initial conditions of data taken from cosmic surveys [57–59]. Relative to analytic approaches, grid-based schemes allow for “field-level” assessments of perturbation theories. Computing any summary statistic in grid-based PT is achieved by analyzing the output as a mock dataset. Grid-based schemes have also recently seen interest within the context of Eulerian perturbation theories [60–62]. However, Eulerian theory does not capture large-scale displacements which are naturally included in Lagrangian schemes.

In the right panel of Fig. 2 we show a comparison between the matter power spectrum measured from the Zel’dovich density field, the fully nonlinear field and the matter power spectrum as predicted by ZeNBu. Both ZA predictions have the same Gaussian smoothing applied to their linear power spectra, at a scale of $k_{\text{cut}} = \pi N_{\text{mesh}}/L_{\text{box}} \approx 3.2 \text{ hMpc}^{-1}$, given by Eqn. B.2. The predictions agree exquisitely to small scales of $k = 1 \text{ hMpc}^{-1}$.

Emulators based on Lagrangian Perturbation Theory (hybrid EFT) have recently been introduced as powerful models for describing the basis spectra of Eqn. 3.8, where the non-linear displacements from N -body simulations are used, as opposed to perturbative displacements [29–32]. While these emulators are promising tools, there are still limitations which have prevented their wide adoption. For example, most LPT-based emulators eventually revert to analytic predictions at large scales. This happens when sample variance in the suite of N -body simulations used to construct the emulator becomes too large, preventing accurate predictions from being made. Many algorithms for performing N -body simulations exhibit slight discrepancies in how they evolve large-scale growth [8, 63–66] relative to linear theory. A small-mismatch between LPT predictions and measurements in simulations can introduce unphysical features in emulated spectra⁶. These features, while local in Fourier space, significantly affect configuration space-based emulators using these same suites. Another challenge

⁵“Ze(ldovich calculations for) N-B(oddy Em)u(lators)”, available at <https://github.com/sfschen/ZenBu>

⁶For example, refs. [30, 31] had to take care to connect LPT predictions to HEFT spectra without introducing discontinuities in the emulator predictions.

is the trade-off between covering a cosmological parameter space adequately while also having sufficiently large volumes that emulator accuracy is below the requirements for the next generation of large-scale surveys. Indeed, attempts to use `anzu`⁷, the emulator of ref. [30] based on the `Aemulus` suite [8], in an analysis of data [28] ran into the issue that the bounds of the emulator were too restrictive relative to the posteriors obtained from the analysis.

As such, reducing sample variance at large scales, especially in the transition regime between LPT and full non-linear displacements, for basis spectra measured from N -body based emulators could have the potential to enable larger parameter space coverage and more stable emulators than have been currently constructed. The striking correlation between the matter density fields constructed in the Zel’dovich approximation and the final result of the N -body simulation suggest that cheap Zel’dovich realizations can enable substantial reduction in sample variance of emulated quantities. This was pointed out, without control variates, in ref. [46] for the case of the matter power spectrum, but this principle holds for all basis spectra, as well as any other statistic that can be computed within the Zel’dovich approximation. We will illustrate how this is the case in the remainder of the text. In Fig. 3 we show a similar plot to the lower right-hand plot of Fig. 2 but for all component spectra in ZA. We find near sub-percent agreement for almost all basis spectra to scales of $k = 1 h\text{Mpc}^{-1}$. We note that the discrepancy for cubic basis spectra at large scales is due to these quantities being noisily measured despite using $N=100$ `Quijote` HR simulations. The small scale discrepancies for $\langle s^2 \delta^2 \rangle$ are due to issues in the smallest scales probed by our simulations⁸. These small discrepancies are not an issue, as

1. Contributions from higher order basis spectra are sub-leading relative to the full predictions of galaxy clustering and lensing.
2. Higher order basis spectra fall off rapidly at high k , and thus their difference will induce a very small bias in the variance-reduced estimate of that basis spectra.
3. The regression coefficient $\hat{\beta}(k)$ will be small at high k , when the correlation between Zel’dovich and N -body is small.

The sub-percent agreement for most basis spectra to $k = 1 h\text{Mpc}^{-1}$ shown in Fig. 3 shows there are no potential biases in using analytic calculations as the mean when applying ZA as a control variate for Lagrangian basis spectra. The use of analytic predictions for the mean of the control variate significantly reduces computational expenses in applying the technique. The other main potential challenge, as discussed, is computing the regression matrix β . We will discuss this in the following section.

4 Control variates for LPT basis spectra

Returning to the question of variance reduction, we now formulate the problem of control variates within the notation of basis spectra in LPT. The control variate problem as defined in Eqn. 2.1 may be written as⁹

⁷Available at <https://github.com/kokron/anzu>

⁸We discuss in Appendix B what is the dynamic range in scales that an N -body simulation should contain in order to accurately reproduce the analytic calculations of ZA.

⁹In principle one should consider the full multivariate control variate problem for the data vector $\{P_{ij}(k_1) \cdots P_{ij}(k_n)\}$, however for power spectrum estimation the diagonal approximation works reasonably well [16].




Figure 3. Ratio between the grid-based ZA basis spectra as measured from 100 Quijote HR boxes and spectra computed analytically in the Zel’dovich approximation. Each line corresponds to the cross-spectrum of the advected particles weighted by fields (i, j) in the initial conditions. The measurements are binned with $\Delta k = 4$ modes, corresponding to approximately $\Delta k = 0.025 h\text{Mpc}^{-1}$. The shaded bands are the standard deviations measured from the Quijote HR boxes divided by $\sqrt{100}$.

$$\hat{P}_{ij}(k) = \hat{P}_{ij}^N(k) - \beta_{ij}^*(k) \left(\hat{P}_{ij}^{\text{ZA}}(k) - P_{ij}^{\text{ZA}}(k) \right), \quad (4.1)$$

where \hat{P}_{ij}^N denotes the measured basis spectrum from full N -body simulations, \hat{P}_{ij}^{ZA} is the basis spectrum from ZA given the same initial conditions, and P_{ij}^{ZA} is the average ZA prediction, which can be computed either analytically or from several ZA realizations averaged together. The β^* term in this case is slightly more complicated. Assuming only contributions diagonal in Fourier wavenumber k contribute we have

$$\text{Cov}[\hat{P}_{ij}^N, \hat{P}_{ij}^{\text{ZA}}](k) \propto \left\langle \hat{\mathcal{O}}_i^N(\mathbf{k}) \hat{\mathcal{O}}_j^N(-\mathbf{k}) \hat{\mathcal{O}}_i^{\text{ZA}}(\mathbf{k}) \hat{\mathcal{O}}_j^{\text{ZA}}(-\mathbf{k}) \right\rangle - P_{ij}^N(k) P_{ij}^{\text{ZA}}(k), \quad (4.2)$$

$$\text{Var}[\hat{P}_{ij}^{\text{ZA}}](k) \propto \left\langle \hat{\mathcal{O}}_i^{\text{ZA}}(\mathbf{k}) \hat{\mathcal{O}}_j^{\text{ZA}}(-\mathbf{k}) \hat{\mathcal{O}}_i^{\text{ZA}}(\mathbf{k}) \hat{\mathcal{O}}_j^{\text{ZA}}(-\mathbf{k}) \right\rangle - (P_{ij}^{\text{ZA}}(k))^2. \quad (4.3)$$

The operator $\hat{\mathcal{O}}_i(\mathbf{k})$ corresponds to the advected, late-time Lagrangian bias field whose cross-correlations form the basis spectra, as defined in Eqn. 3.14. When the ratio defining β^* is taken, we can neglect the proportionality constants related to the volume of the simulation box¹⁰. Since any choice of β leads to an unbiased estimator, we adopt a compromise between the optimal β^* and a surrogate that is easier to compute while still providing substantial variance reduction. To do this, we will employ approximations in computing the covariances of Eqns. 4.2 and 4.3. Notably, we opt to select only the disconnected contributions from the above four-point correlations, which leads to a more amenable form

$$\hat{\beta}_{ij}(k) \approx \frac{\hat{P}_{ii}^{\text{NZ}} \hat{P}_{jj}^{\text{NZ}} + \hat{P}_{ij}^{\text{NZ}} \hat{P}_{ij}^{\text{ZN}}}{(\hat{P}_{ij}^{\text{ZZ}})^2 + \hat{P}_{ii}^{\text{ZZ}} \hat{P}_{jj}^{\text{ZZ}}}, \quad (4.4)$$

where P_{ij}^{ZN} is defined as the spectrum $\langle \mathcal{O}_i^{\text{ZA}} \mathcal{O}_j^N \rangle$.

¹⁰There are also contributions which depend on the connected tri-spectrum between N body and ZA fields which don’t depend on the volume, but they are neglected when taking only disconnected terms.




Figure 4. Plotting β as measured empirically from `Quijote` HR, compared to the analytic approximation of Eqn. 4.4 as well as using the tanh filter function of Eqn. 4.7 where k_0 and Δ_k have been calibrated to β_{11} . Despite calibrating to β_{11} we find the filtered analytic β are in good agreement with the empirical results for all other basis spectra.

The quantity $\hat{\beta}$ is simple to compute from the simulation outputs used to generate the basis spectra of Lagrangian bias emulators. The main over-head is associated with producing the Zel’dovich-advected component fields. Additionally, since the estimator in Eqn. 4.4 is a ratio of power spectra measured from boxes which share the same initial phases, we expect $\hat{\beta}_{ij}$ to be a relatively noiseless quantity, despite depending on only a single ZA realization. We have verified that, while in principle the quantity $\hat{\beta}_{ij}(k)$ could vary from simulation to simulation the variance cancellation from the ratios of simulated quantities results in a measured $\hat{\beta}_{ij}$ with negligible variance. Additionally we note that, formally, using a $\hat{\beta}_{ij}(k)$ measured in the same simulation as the one whose variance is being reduced could affect the unbiased properties of the estimator in Eqn. 4.1. To avoid this issue, we use $\hat{\beta}_{ij}$ measured from Box 0 of `Quijote` HR as the regression matrix for all other boxes.

In order to test the approximation of $\hat{\beta}$ in Eqn. 4.4, we also compute the numerical β^* from the `Quijote` HR boxes. We compute the covariance between ZA and N -body simulations explicitly and compute the multivariate estimator for β^* described in [16]. In the context of basis spectra estimation, this β^* reads

$$\beta_{ij}^*(k, k') = \sum_{k''} \text{Cov}[\hat{P}_{ij}^N(k), \hat{P}_{ij}^{ZA}(k'')] \text{Cov}[P_{ij}^{ZA}(k''), P_{ij}^{ZA}(k')]^{-1}. \quad (4.5)$$

The inverse covariance is computed using the Moore-Penrose pseudoinverse due to the small number of realizations. While the entire matrix β^* is poorly estimated with only $N = 100$ simulations, we find that the diagonal component is relatively well measured. Thus, we proceed to compare our analytic approximation with the empirical result. This comparison is shown between the orange and blue lines of Fig. 4 for the ten basis spectra which span second order Lagrangian bias. The analytic approximation of Eqn. 4.4 holds well until scales of $k \approx 0.3 h\text{Mpc}^{-1}$, after which the empirically measured β^* damp to zero. This damping is not observed in the analytic approximation. The disagreement between the analytic form and our measurement can be understood by re-casting β^* as

$$\beta^* = \rho_{N,ZA} \frac{\sigma_N}{\sigma_{ZA}}. \quad (4.6)$$

The dominant contribution to the second term will scale as the “transfer function” between linear and non-linear power, and typically grows as a function of scale in Λ CDM. On the other hand, the disconnected approximation for $\rho_{N,ZA}$ we employ does not capture the full de-correlation between N -body dynamics and the Zel’dovich approximation. However, for all basis spectra we also find that this damping is very well approximated by a tanh function of the form

$$F(k; k_0, \Delta_k) = \frac{1}{2} \left[1 - \tanh \left(\frac{k - k_0}{\Delta_k} \right) \right]. \quad (4.7)$$

Fitting values of k_0 and Δ_k to β_{11} , measured for the dark matter densities, we find fiducial values $k_0 = 0.618 h\text{Mpc}^{-1}$ and $\Delta_k = 0.167 h\text{Mpc}^{-1}$. These values give a very good description of high- k damping for all other β functions and we use them throughout the rest of this work. Thus, we adopt as an inexpensive approximation to the full empirical β the product between Eqn. 4.4 and the filter function $F(k; k_0, \Delta_k)$. Choosing different values of k_0 and Δ_k also allows for control over the trade-off between any bias between the predicted and measured variate, and cancellation of sample variance. This damping of $\hat{\beta}$ also means any small-scale biases in our estimation of the mean Zel’dovich of the control variate can be safely neglected. In Fig. 4 we show in green the resulting damped β for the ten basis spectra of Lagrangian bias. Despite calibrating the filter F to the regression matrix of the matter power spectrum, its combination with the analytic approximation of Eqn. 4.4 provides a strong agreement for the diagonals of regression matrices of all 10 basis spectra.

We have thus formulated the control variate problem for basis spectra in Lagrangian bias theory (which includes the matter power spectrum and clustering/lensing of biased tracers as special cases) under the assumption the Zel’dovich approximation is a suitable control variate. We have derived an approximation for the regression matrix β_{ij} which reproduces empirical results obtained from a suite of high resolution simulations. We therefore have all of the ingredients required to estimate $\hat{P}_{ij}(k)$ and quantify the variance reduction obtained from Zel’dovich control variates.

5 Results and discussion

In the following section we apply Zel’dovich control variates to progressively more complex-cosmological power spectra. We begin with the matter power spectrum as a warm up, proceed to the case of extending the results for an HOD-like sample of galaxies, and conclude with the fully general case of improving measurements of all 10 basis spectra in second order Lagrangian bias models.

All measurements in this section are carried out across the whole ensemble of $N = 100$ Quijote HR simulations. While small in number relative to the full Quijote suite, we find that we need the higher dynamic range in order to match analytic and grid-based LPT, and prevent biases in the mean control variate estimate. We show why this is the case explicitly in Appendix B. We adopt as a fiducial choice the $z = 0.5$ snapshot but note that there are no impediments to using other snapshots beyond computational costs. The use of $N = 100$ simulations should allow, at least, for quantification of the reduction in the variance of numerical observables. We leave a quantification of their reduction of *co*-variance to future work, as this is a significantly more numerically challenging problem.




Figure 5. *Left panel:* Scale-dependent ratio, between standard deviations of matter power spectra, from paired ZA control-variates and the empirical standard deviations from the $N=100$ Quijote HR boxes. *Right panel:* Visualizing variance reduction in the configuration-space correlation function, $\xi_{mm}(r)$. The blue curve results from Hankel-transforming the Zel’dovich $P_{mm}(k)$ prediction from ZeNBu. The orange and green curves are obtained by similarly Hankel-transforming empirically measured power spectra from a single Quijote HR box. The inset shows the increase in bias and variance for the standard N -body result at very large separations, while the control variate closely tracks the correct large-scale behavior predicted by the Zel’dovich approximation.

5.1 Variance reduction for the matter power spectrum

In Eqn. 2.3 we derived that, in the standard control variates approach, the variance of the estimator is parametrically reduced by the cross-correlation coefficient between the full N -body result and the Zel’dovich surrogate. In Fig. 2 we showed that for the case of the matter density field, this cross-correlation coefficient was substantial at $k = 1 h\text{Mpc}^{-1}$. However, Eqn. 2.3 assumes a form β^* which is different from the final form of $F(k; k_0, \Delta_k) \hat{\beta}_{ij}(k)$ that we used in this work. Thus, we turn to our statistical ensemble of simulations to quantify, in practice, how much variance reduction we can achieve through our technique.

We begin by considering solely the matter density field and its auto-spectrum, $P_{mm}(k)$. In the language of Lagrangian bias this is the $P_{11}(k)$ spectrum, and we will show results for the whole suite of basis spectra shortly. The control variate estimator we’ve derived, explicitly, is given by

$$\hat{P}_{mm}(k) = \hat{P}_{mm}^N(k) - F(k; k_0, \Delta_k) \left(\frac{\hat{P}_{mm}^{*NZ}(k)}{\hat{P}_{mm}^{*ZA}(k)} \right)^2 \left(\hat{P}_{mm}^{ZA}(k) - P_{mm}^{ZA}(k) \right), \quad (5.1)$$

where \hat{P}^{*NZ} and \hat{P}^{*ZA} are spectra measured from a different box in order to prevent biasing the estimator, as discussed in § 4. We measure the matter power spectrum in our Zel’dovich mocks, as well as in the full N -body fields, for all 100 boxes of Quijote HR. We also measure the cross-power spectrum between matter fields in all 100 matched boxes. To quantify the amount of variance reduction (or, equivalently, effective volume increase) we report the ratio of standard deviations

$$\frac{\sigma^{\text{Nbody}}(k)}{\sigma^{\text{CV}}(k)}, \quad (5.2)$$

as a function of scale obtained from implementing ZA control-variates for all basis spectra. These standard deviations are measured empirically from the Quijote HR suite. We show

the result in the right-hand panel of Fig. 5.

In agreement with the intuition of Eqn. 2.3, we see that variance reduction is most substantial at large scales where the Zel’dovich approximation more faithfully captures the dynamics of structure formation. We also find that, due to the damping function $F(k; k_0, \Delta_k)$ at very small scales we observe no reduction in uncertainty from our estimator. At $k \simeq 0.01 h\text{Mpc}^{-1}$, the reduction in standard deviation is of the order of $\sigma^{\text{Nbody}}/\sigma^{\text{CV}} \sim 110$. Since the standard deviation scales as

$$\sigma \propto V^{1/2},$$

this is equivalent to an effective volume increase of a factor of $\sim 10^4$. This reduction is larger at larger scales which are worse affected by sample variance. While the empirical N -body uncertainty grows as $k \rightarrow 0$, we find that for the paired ZA control variate the uncertainty scales as $\sigma^{\text{CV}} \propto P(k)$, after $k \lesssim 0.03 h\text{Mpc}^{-1}$. This can be understood by treating the cross-spectrum between ZA and N -body as the same at large scales as the cross-spectrum between ZA and linear theory. Then, following ref. [10] we can write $\rho_{ZN}^2 = e^{-2(k/k_{\text{NL}})^2}$. The end result is that $\sigma_{\text{CV}}^2 \propto P(k)^2/(V\Delta k k_{\text{NL}}^2)$ at large scales.

To illustrate the degree of sample variance reduction we achieve, we also reconstruct the configuration-space correlation function $\xi_{mm}(r)$ directly from the empirically measured power spectra in boxes. We reconstruct the correlation function by evaluating

$$\hat{\xi}_{mm}(r) = \frac{1}{2\pi^2} \int_0^\infty k^2 dk j_0(kr) \hat{P}_{mm}(k),$$

which is done using `MCFit`¹¹. The empirically measured power spectra are extrapolated beyond $k \in [k_f, k_{\text{Nyq}}]$ ¹² by constructing a linear spline in $(\log k, \log P(k))$, which is then rebinned in log-spaced k bins. We show our reconstructed correlation functions in the left panel of Fig. 5, as well as the prediction from the Zel’dovich approximation. We find the reconstructed $\hat{\xi}$ from ZA control variates to have significantly reduced noise, as well as appropriate damping of the BAO. We also see that at very large scales $r \geq 10^2 h^{-1}\text{Mpc}$ the result is in close agreement with linear theory, despite only extrapolating the power spectrum to lower k . In contrast, the result from attempting this ξ_{mm} reconstruction for the default N -body result leads to significant biases beyond $r > 10^2 h^{-1}\text{Mpc}$, as well as large variances around the trend line. We note that if we extend $\hat{P}(k)$ at $k < 2\pi/L_{\text{box}}$ by including the noiseless $P(k)$ from linear theory, the reconstructed control variate ξ does not change, while the N -body result becomes less biased at large scales but still with equivalently large variance.

The reconstruction of large-scale $\hat{\xi}_{mm}(r)$ in this form, while a simple application, showcases the strength of this technique and the myriad ways in which the variance-reduced spectra could be used. We stress that the only additional steps in going from the green to the orange curves in Fig. 5 are: running an inexpensive Zel’dovich mock at $z = 0.5$ and from this measuring $P_{mm}^{\text{ZA}}(k)$ and the cross-spectrum $P_{mm}^{\text{NZ}}(k)$.

5.2 Reducing the variance of biased tracer samples

Hydrodynamic simulations [67–73] that try to reproduce galaxy formation ab initio, and radiative hydrodynamics simulations that try to understand the structure and evolution of reionization fronts [74] have extreme resolution requirements which significantly limit their

¹¹<https://github.com/eelregit/mcfit>

¹²For `Quijote` HR, which has $L_{\text{box}} = 1000 h^{-1}\text{Mpc}$ and $N_{\text{mesh}} = 1024$ we have $k_f \approx 6.3 \times 10^{-3} h\text{Mpc}^{-1}$ and $k_{\text{Nyq}} \approx 3.2 h\text{Mpc}^{-1}$.




Figure 6. *Left panel:* Measurements of tracer statistics with and without applying Zel’dovich control variates. The points show the power spectrum of galaxies populating halos in a randomly chosen Quijote HR box, with uncertainties estimated from the ensemble. The dark, solid colored lines show measurements from the same realization, but measured using control variates. Dotted black lines correspond to the mean of the Quijote HR ensemble of simulations. The transparent lines correspond to the Zel’dovich predictions, at large scales, using only b_1 . *Right panel:* Ratio of the standard deviations of the spectra shown in the left panel. The dashed lines show the improvements in error reduction from only using matter field as a control variate, while the solid lines show the same technique but where we have also included the b_1 field in building the control variate.

volume. Measuring summary statistics in these simulations with less noise and to larger scales is thus highly desirable. Given the impacts of small-scale baryonic physics do not backreact on large scales, we can use control variates to extend their measurements to larger scales. Extending simulation-based models to larger scales is also highly desirable in light of recent advances in simulation-based inference. While the power spectrum at small scales possesses many modes, which require simulations to accurately describe, summary statistics at large scales provide complementary information and their combination can break parameter degeneracies that arise in complicated models. Purely simulation-based analyses of galaxy statistics such as those in [75–79] achieve impressive constraints on parameters, however are limited to fairly small scales ($r \lesssim 30h^{-1}\text{Mpc}$).

We proceed to apply Zel’dovich control variates to a sample of “galaxies” selected from the Quijote HR boxes to demonstrate how our technique can extend the range of scales of measurements of biased tracer statistics in simulations. Using the public halo catalogs, we create density fields of halos, where each halo is weighted by the expected number of galaxies $\langle N(M) \rangle$ for its mass. The weights are derived from the halo occupation distribution of ref. [80], which was fit to a sample of luminous red galaxies from the Sloan Digital Sky Survey. Since control variates will be most powerful at large scales, this crude approximation to the galaxy–halo connection should be sufficient for the purposes of illustrating applications to tracer samples. While we have used catalogs of mock LRGs, these techniques are applicable to any tracer catalog from a simulation. This includes galaxies in hydrodynamical simulations, or the power spectrum of the electron distribution measured during reionization, for example.

We wish to estimate the power spectra relevant to galaxy clustering and lensing, that is,

$$P_{gg} = \langle \delta_g \delta_g \rangle, \quad P_{gm} = \langle \delta_g \delta_m \rangle \quad (5.3)$$

in our simulations (having discussed P_{mm} already). We first use the matter power spectrum itself as a surrogate as this is extremely cheap and quick to implement. We then define the

control variate estimator for these spectra

$$\hat{P}_{g[g,m]}(k) = P_{g[g,m]} - \hat{\beta}_{mm}^* \left[\hat{P}_{mm}^{\text{ZA}}(k) - P_{mm}^{\text{ZA}}(k) \right], \quad (5.4)$$

where $P_{g[g,m]}$ is a short-hand for either spectrum used. While we expect that the cross-correlation coefficient between the matter power spectrum in Zel’dovich and the statistics of biased tracers to be lower than if we had used a surrogate for the tracer itself, we should still expect some degree of variance reduction. We show the results of this procedure in the right-hand panel of Fig. 6, in the dashed lines. Despite using a crude control variate, we still find substantial reduction in sample variance across all scales $k \leq 0.2 h\text{Mpc}^{-1}$. At $k = 0.2 h\text{Mpc}^{-1}$ the autospectrum has a $2\times$ reduction in its uncertainty, while for P_{gm} we find a $3\times$ reduction at this scale. The reduction is always higher for the galaxy–galaxy lensing spectrum P_{gm} , which can be understood from the absence of shot noise in this spectrum compared to the tracer autospectrum. Shot noise will explicitly decorrelate the matter density field from that of tracers at small scales. At $k \sim 0.01 h\text{Mpc}^{-1}$, we find an average reduction in uncertainty that is of an order of magnitude. This is equivalent to averaging $N = 100$ simulations at this volume.

We also consider variance reduction in biased tracer spectra for a slightly more accurate control variate. Instead of using the matter power spectrum, we use our Zel’dovich component fields to create linearly biased surrogates with power spectra given by

$$P_{gg}^{\text{ZA}}(k) = P_{mm}^{\text{ZA}}(k) + 2b_1 P_{1\delta}^{\text{ZA}}(k) + b_1^2 P_{\delta\delta}^{\text{ZA}}(k), \quad (5.5)$$

$$P_{gm}^{\text{ZA}}(k) = P_{mm}^{\text{ZA}}(k) + b_1 P_{1\delta}^{\text{ZA}}(k). \quad (5.6)$$

We also use $\hat{\beta}$ constructed appropriately from the disconnected approximation of Eqn. 4.4. This step requires the re-estimation of the cross-spectra between the biased tracer and the Zel’dovich fields, and thus incurs some additional computational overhead. However this overhead is quite small, as measurements of power spectra are not a particularly computationally intensive task. For every box we estimate a value of b_1 using the field-level bias estimator of ref. [35] at $k_{\text{max}} = 0.2 h\text{Mpc}^{-1}$, but the variance reduction obtained is insensitive to the exact value of b_1 adopted as long as the analytic and grid-based ZA predictions use consistent values¹³. The results from applying our linearly biased control variate are shown in both panels of Fig. 6. In the left panel, we compare a single realization to the mean of the 100 Quijote HR boxes for this tracer sample and find entirely compatible results. In addition, we see that our largest scale points are wholly in agreement with predictions from the Zel’dovich approximation and our inferred linear bias value. This means that we can extend, as a model, our simulation-based measurement of $P_{g[g,m]}$ to arbitrarily large scales without worrying about sample variance. In addition, in the right panel of Fig. 6 we show, in solid lines, the reduction in sample variance from the linearly biased control variate. The performance is mildly better than the P_{mm} case, but overall comparable. The peak reduction in uncertainty over the P_{mm} case is of the order of 30% (equivalent to $\sim 70\%$ more volume), and occurs at $k = 0.1 h\text{Mpc}^{-1}$.

5.3 Variance reduction for hybrid EFT spectra

Beyond the matter density field, or linearly biased tracers, it is also of great interest to improve measurements of all possible spectra that can contribute to the statistics of biased tracers.

¹³Indeed, the P_{mm} estimator can be thought of as a special case of the linearly biased tracer example where we set $b_1 = 0$. This is an ‘infinitely incorrect’ estimate of b_1 , and yet we still obtain strong variance reduction.

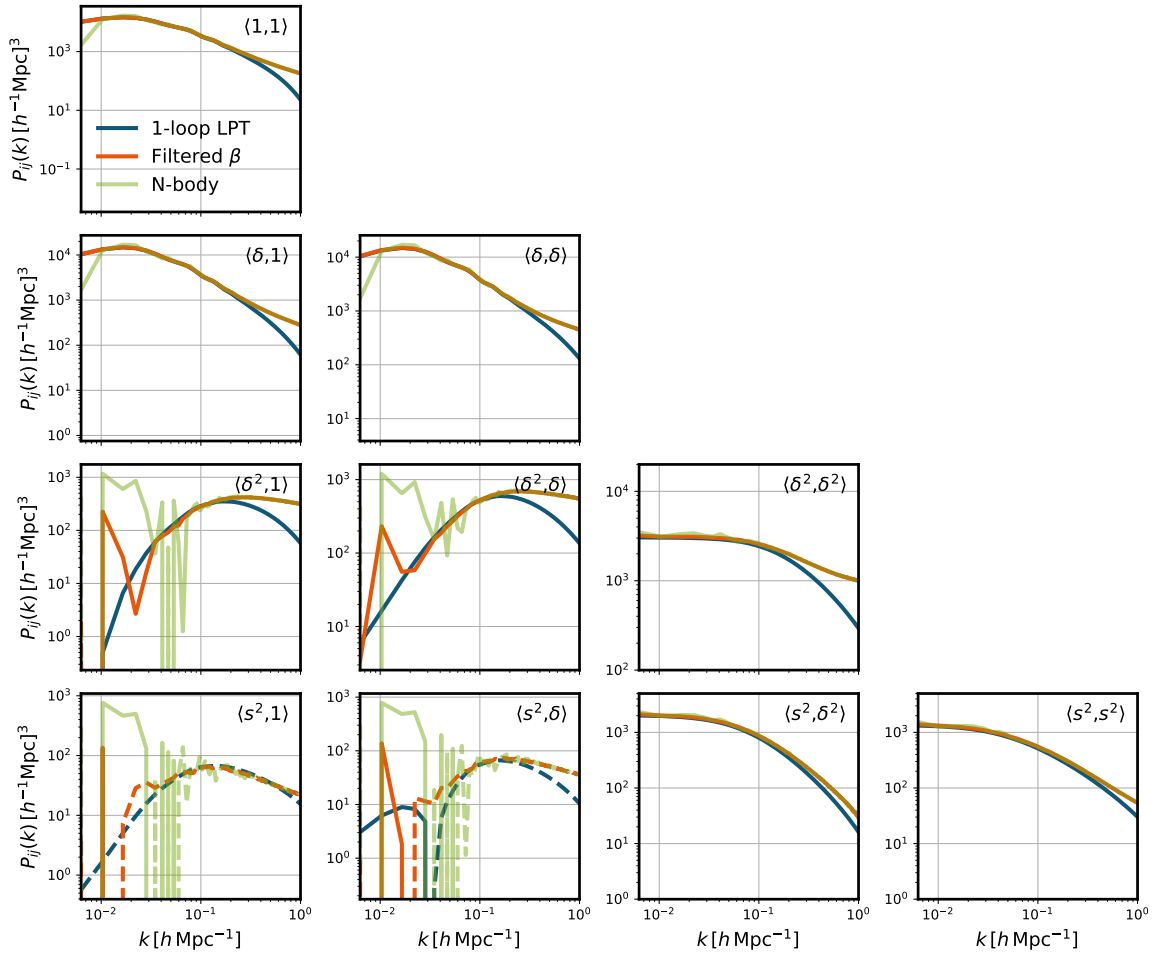


Figure 7. Comparison between the non-linear component spectra and the variance-reduced spectra measured with Zel’dovich control variates. With a single simulation we can extend accurate simulation-based measurements for the basis spectra from $k \simeq 0.1 \text{ hMpc}^{-1}$ down to $k \approx 0.03 \text{ hMpc}^{-1}$, allowing accurate matching to LPT for the extrapolation to $k \rightarrow 0$.

The question of applying control variates to the full set of basis spectra $P_{ij}(k)$ is particularly relevant for the construction of emulators of basis spectra in hybrid EFT, as discussed in § 3.2. In this subsection we report the results of applying ZA control variates to the ten basis spectra of second-order hybrid EFT.

In Fig. 7 we show component spectra from an N -body simulation drawn from this suite, and the result of applying our Zel’dovich control variate scheme, as well as predictions from one-loop LPT computed using `velocileptors`. We observe significant reduction in noise at large scales, especially for the cubic spectra, which enforces significantly better agreement with perturbation theory at large scales than naively obtained from a single realization. Notably, we find that a single paired Zel’dovich realization can extend the range of agreement with LPT from $k \simeq 0.1 \text{ hMpc}^{-1}$ to $k \approx 0.03 \text{ hMpc}^{-1}$. We also note that despite using ZA as the control variate, the final estimator agrees (in the $k \rightarrow 0$ limit) with one-loop Lagrangian Perturbation Theory which is a more accurate model of basis spectra.

A more quantitative assessment of the improvement in accuracy from our approach is




Figure 8. Ratio between basis spectra from an N -body box with the mean of $N = 100$ Quijote HR basis spectra (in blue) and one Zel'dovich control variate estimate (in orange), using the filter of Eqn. 4.7 fit to β_{11}^* . We see significant improvements in the accuracy of the estimation of basis spectra, especially at large scales. We caution that noise in the ratios for *cubic* spectra are due to the fact that even with 100 realizations the mean values of cubic spectra exhibit significant sample variance.

shown in Fig. 8, where we now show residuals of our CV approach compared to the mean of our Quijote HR boxes. We see that for non-cubic spectra the paired curves are in exquisite agreement with the low k values inferred by our large volumes, with no discernible biases. For cubic spectra, this comparison is hampered by significant sample variance despite having $N = 100$ boxes. The means are not well determined, and of comparable noise to a single ZA-paired control variate estimate. Since ZA realizations are extremely inexpensive to generate for a given box (and, indeed, the density field that seeds ZA displacements is a standard output of initial conditions codes), there is a large potential for improvement of the accuracy of simulation-based measurements with little extra effort relative to running the initial simulation.

In Fig. 9 we now show the full degree of scale-dependent variance reduction we obtain from paired ZA control variates for all ten basis spectra that span second-order Lagrangian bias. The $\langle 11 \rangle$ curve was previously shown in Fig. 5. Quantitatively, we find that even for basis spectra of fields that couple small scale modes to large scales in the initial conditions (such as $\delta^2(\mathbf{q})$ and $s^2(\mathbf{q})$) we still find an over ten-fold reduction in error for all basis spectra. For spectra that don't involve convolved fields we find improvements that range from $100\times$ to $500\times$ reductions in power spectra uncertainties from our estimator which augments each simulation with a single ZA mock.

The reduced performance of our variance reduction techniques for fields that involve Fourier-space convolutions can be directly understood as a consequence of the smoothing imposed in initial conditions in order to match grid-based and analytic predictions of basis spectra in the Zel'dovich approximation, as described in Appendix B. The full Quijote HR simulations have been run with un-damped initial conditions, that cut off in power at the Nyquist frequency k_{Nyq} . When producing Lagrangian bias operators which involve Fourier-space convolutions such as

$$\delta^2(\mathbf{k}) = \int \frac{d^3k'}{(2\pi)^3} \delta(\mathbf{k} - \mathbf{k}') \delta(\mathbf{k}'), \quad (5.7)$$

a large scale mode \mathbf{k} will receive contributions from arbitrarily small scales. When the linear density $\delta(\mathbf{k})$ is filtered in order to accurately match grid-based and analytic approaches, the




Figure 9. Ratio of the standard deviations of basis spectra, as measured from $N = 100$ *Quijote* HR simulations, compared to the variance measured from $N = 100$ pairs of simulations and ZA control variates. At large scales, we observe at least a factor of $10\times$ reduction in uncertainty compared to cosmic variance at volumes of $V = 1(h^{-1}\text{Gpc})^3$ for the $\langle s^2, s^2 \rangle$ basis spectrum, up to nearly $1000\times$ for the $\langle \delta, \delta \rangle$ spectrum.

final filtered Lagrangian fields will now be slightly de-correlated even at large scales compared to the fiducial density field used to initialize the N -body simulation. From Eqn. 2.3 we can then see that a lower cross-correlation coefficient will result in less substantial variance reduction than for the case of the standard density field, where the Fourier-space damping is cancelled out. This issue can be circumvented: if one produces ZA realizations from the un-damped density field whose statistics agree with analytic LPT then our technique should result in comparable variance reduction for all basis spectra. We leave a more detailed study of matching analytic and grid-based LPT to future work.

The consequences of applying paired ZA control variates to the problem of simulation-based modelling are manifold. The most immediate consequence is that suites of cosmological N -body simulations used to construct emulators such as *Aemulus* [8], *BACCO* [81] and *AbacusSummit* [9] can drastically improve their large-scale measurements of biased tracer spectra without running additional simulations. These methods also make clear that future suites designed for the purpose of emulating biased tracer statistics should not require prohibitively large simulation volumes to ensure precise measurements of these statistics at large scales. As long as the dynamic range of the simulation is sufficient (and most emulation suites possess better dynamic range than *Quijote* HR) to resolve physical scales that contribute to most basis spectra, paired ZA control variate estimators can allow for accurate emulation of large physical scales with simulations that have significantly reduced volumes. Thus, a larger number of cheaper simulations can be used to develop the suite. This larger number of simulations can then cover a larger cosmological parameter space, or perhaps known extensions to Λ CDM where the second order Lagrangian bias expansion above and Zel’dovich dynamics are still appropriate.

5.4 Comparison with previous results

It is also worth comparing techniques described in this paper with other recent attempts at reducing the variance of basis spectra measured in simulations. In ref. [82], the authors studied the properties of basis spectra as measured in ‘paired-fixed’ simulations. Pairing and fixing involves running, at every cosmology, two sets of simulations with non-Gaussian initial

conditions that exactly fix the amplitude of fluctuations to follow $|\delta_{\mathbf{k}}| = \sqrt{VP(k)}$ and have Fourier phases shifted by a factor of π . The authors of ref. [82] have shown that by using an estimator that averages basis spectra as measured by these two sets of simulations, significant reduction in uncertainties can be obtained for a sub-set of basis spectra.

Our results in Fig. 9 are numerically comparable to those of fig. 3 in ref. [82], however we additionally find substantial reduction in uncertainty for the quartic spectra $\langle \delta^2 \delta^2 \rangle$, $\langle \delta^2 s^2 \rangle$, and $\langle s^2 s^2 \rangle$ unlike ref. [82] who find no improvement from fixing-and-pairing for those spectra. Additionally, the use of paired ZA control variates only requires cheap surrogate realizations for every simulation, with generic initial conditions, as opposed to having to run two sets of simulations with non-Gaussian initial conditions when using paired-and-fixed simulations. However, the tools of ref. [82] are complementary to what we have introduced in this work and future work could investigate jointly using many variance reduction techniques to achieve greater results than what could be accomplished by each individually.

Recently, the DESI collaboration [83] applied control variates in order to expand the volumes of their fiducial suite of simulations, **AbacusSummit**. Their control variate of choice was the approximate N -body solver **FastPM** [39]. Over 500 surrogates were produced throughout their work, generating on the order of 400TB of ancillary data at a computational cost of 24 million NERSC CPU-hours. Their technique results in an improvement of effective volume for the redshift-space clustering of halos on the order of $100\times$. While not an apples-to-apples comparison, our worst-case scenario improvement for biased spectra is comparable to their quoted improvement. In this publication we have not analyzed redshift space correlations or higher order correlations, but we shall present these developments in future work.

The computational expense of implementing ZA control variates in an N -body simulation are substantially more modest than when using an approximate solver such as **FastPM**. For the code that we have developed in this paper, at the **Quijote** HR resolution, producing all relevant Lagrangian component fields, advecting them, and measuring all relevant cross-spectra takes on the order of 50 CPU-hours per snapshot. The largest computational expenses are equally distributed among generating the Lagrangian fields, advecting them to produce late-time realizations, and measuring all relevant cross-spectra to compute β (which also contains all N -body and ZA basis spectra).

6 Conclusions

In this paper we have re-visited the problem of sample variance reduction in N -body simulations through use of the method of control variates. While previous works in this direction have used approximate N -body solvers as surrogates for structure formation, we proposed the use of the Zel'dovich approximation as a surrogate of structure formation that is simultaneously inexpensive to produce and highly correlated with the non-linear density field. We have shown that, for biased tracers in real space, we can compute the mean prediction in Zel'dovich to arbitrary precision. This development sidesteps one of the main limitations of applying control variates previously, where many realizations of the approximate simulation had to be produced in order to reduce the uncertainty on the mean variate. Additionally, we have presented a physically motivated, analytic approximation for the regression matrix β , which is often another source of significant computational expense when applying control variates. We validated this approximation using an ensemble of N -body simulations and showed that a simple empirical damping function applied to the approximation leads to strong agreement with numerical estimates of the regression matrix.

We proceed to quantify any residual biases in our technique, as well as reduction in uncertainty, in three different problems which are applicable to many different classes of N -body simulations. We showed that for the matter power spectrum, we find a reduction in variance that is equivalent to averaging 10^4 simulations at $k = 0.01 h\text{Mpc}^{-1}$ for the *Quijote* volume. We also find strong agreement with the Zel’dovich approximation at large scales, which enables reconstructions of the configuration–space correlation function that are significantly less noisy at large scales.

We then turned to the more complicated case of reducing variance in measurements of tracer statistics. We showed that using the matter power spectrum as the control variate for tracer statistics still leads to substantial variance reduction, comparable to increasing the volume of the simulation by $100\times$ at large scales. We further showed that using a more elaborate but still highly approximate surrogate, a linearly biased tracer, led to additional reductions in uncertainty on the order of 30% at little additional computational cost.

We concluded with the most general case of applying control variates to the ten basis spectra of second-order hybrid EFT. We found improvements in effective volume that were equivalent to $\sim 100 - 10^6\times$, depending on the specific basis spectrum under consideration. This dramatic increase in effective volume has substantial implications in the design and accuracy of emulators of cosmic structure.

There are many future directions to consider stemming from this work. The most immediate, perhaps, is to extend our results to redshift-space. Redshift-space distortions in the Zel’dovich approximation are a simple additive offset in the original displacements. Thus, we should expect that the degree of variance reduction obtained here should be sustained. We shall present this development in a future paper. Another direction in which to extend this work is in pushing to higher order correlations such as the bispectrum of biased tracers. While the 3-point correlation function has previously been computed in Zel’dovich [84], the bispectrum has not. However, even lacking a full description of the bispectrum of biased tracers, the results using matter power spectra as a surrogate show that gains could be made in this regime by computing the matter bispectrum in the Zel’dovich approximation, B_{mmm}^{ZA} . Another potential consideration is extending these techniques to the case of simulations with primordial non-Gaussianity (PNG). Matching phases between simulations with and without PNG has recently been demonstrated to be a powerful tool to isolate the effects of PNG in simulations with substantially less sample variance [85]. As biasing is well-understood in the presence of PNG [86], producing variance-reduced realizations of observables in the presence of PNG is well within the scope of the techniques we have laid out in this paper. We leave these, and other extensions, to future work.

Acknowledgments

We thank Diogo Bragança and Philip Mansfield for helpful conversations. We thank the referee for insightful comments which helped improve the paper. N.K. is supported by the Gerald J. Lieberman Fellowship. J.D. is supported by the Lawrence Berkeley National Laboratory Chamberlain Fellowship. M.W. and S.C. are supported by the DOE and the NSF. We acknowledge the use of the *Quijote* simulations [43] and thank their authors for making these products public. This research has made use of NASA’s Astrophysics Data System and the arXiv preprint server. This research is supported by the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, and by the National Energy Research Scientific Computing Center, a DOE Office

of Science User Facility under the same contract. Some of the computing for this project was performed on the Sherlock cluster at Stanford. We would like to thank Stanford University and the Stanford Research Computing Center for providing computational resources and support that contributed to these research results. Calculations and figures in this work have been made using `nbodykit` [87] and the SciPy Stack [88–90].

A Basis spectra in the Zel'dovich approximation

In this section we give the power spectra of tracers with bias up to quadratic order within the Zeldovich approximation. Doing so requires including terms beyond 1-loop order which are usually dropped in perturbation-theory calculations to consistently track dynamics and biasing to the same order; here we must include them to properly compute the mean of the basis spectra in simulations where the dynamics are taken to be Zeldovich. Up to 1-loop order the results below are equivalent to those in Eqn. 4.11 from ref. [25] with terms due to beyond-Zeldovich displacements set to zero, while a subset of the higher-order terms for the case of quadratic density bias were calculated in ref. [42, 91].

To evaluate the basis power spectra within the Zeldovich approximation we need to evaluate the functional

$$\mathcal{Z}(\mathbf{q}, \lambda_n, a_n) = \langle e^{\mathcal{M}(\mathbf{q}_1, \mathbf{q}_2, \lambda_n, a_n)} \rangle$$

where the exponent is defined as:

$$\mathcal{M}(\mathbf{q}_1, \mathbf{q}_2) = i\mathbf{k} \cdot \Delta + \sum_{n=1,2} \lambda_n \delta(\mathbf{q}_n) + a_{n,ij} s_{ij}(\mathbf{q}_n), \quad \Delta = \Psi(\mathbf{q}_1) - \Psi(\mathbf{q}_2). \quad (\text{A.1})$$

Since \mathcal{M} is Gaussian we have that \mathcal{Z} is simply given by the exponentiated second cumulant,

$$\begin{aligned} \frac{1}{2} \langle \mathcal{M}^2 \rangle_c = & -\frac{1}{2} k_i k_j A_{ij} + (\lambda_1 + \lambda_2) i k_i U_i + \lambda_1 \lambda_2 \xi_L + \frac{1}{2} (\lambda_1^2 + \lambda_2^2) \sigma_L^2 \\ & + a_{1,ij} a_{2,kl} C_{ijkl} + \frac{1}{2} (a_{1,ij} a_{1,kl} + a_{2,ij} a_{2,kl}) \langle s_{ij} s_{kl} \rangle \\ & + (a_{1,jk} + a_{2,jk}) i k_i B_{ijk} + (\lambda_1 a_{2,ij} + \lambda_2 a_{1,ij}) E_{ij} \end{aligned}$$

where we have defined the functions [21]

$$A_{ij}(\mathbf{q}) = \langle \Delta_i \Delta_j \rangle = X(q) \delta_{ij} + Y(q) \hat{q}_i \hat{q}_j, \quad U_i(\mathbf{q}) = \langle \delta_1 \Delta_i \rangle \quad (\text{A.2})$$

and we also define shear correlators

$$C_{ijkl} = \langle s_{1,ij} s_{2,kl} \rangle, \quad B_{ijk} = \langle \Delta_i s_{1,jk} \rangle, \quad E_{ij} = \langle \delta_1 s_{2,ij} \rangle \quad (\text{A.3})$$

and we have used that $\langle \delta s_{ij} \rangle = 0$. All the correlators are strictly functions of $\mathbf{q} = \mathbf{q}_1 - \mathbf{q}_2$ by translation invariance.

In order to obtain the biased-tracer power spectrum we use the substitutions

$$b_1 \delta_n \rightarrow b_1 \frac{d}{d\lambda_n}, \quad b_2 \delta_n^2 \rightarrow b_2 \frac{d^2}{d\lambda_n^2}, \quad b_s s_n^2 \rightarrow b_s \delta_{ia} \delta_{jb} \frac{d}{da_{n,ij}} \frac{d}{da_{n,ab}}. \quad (\text{A.4})$$

in the bias functional $F(\mathbf{q}_n)$, which becomes an operator \hat{F}_n . The power spectrum is then given by

$$P(k) = \int d^3 \mathbf{q} e^{i\mathbf{k} \cdot \mathbf{q}} \left[\hat{F}_1 \hat{F}_2 \mathcal{Z}(\mathbf{q}, \lambda_n, a_n) \right]_{\lambda_n, a_n=0}. \quad (\text{A.5})$$

This gives the component spectra as

$$P_{ij}(k) = \int d^3\mathbf{q} e^{i\mathbf{k}\cdot\mathbf{q} - \frac{1}{2}k_i k_j A_{ij}} F_{ij}(\mathbf{k}, \mathbf{q}) \quad (\text{A.6})$$

where for each basis spectrum the function $F_{ij}(\mathbf{k}, \mathbf{q})$ is given by

$$\begin{aligned} (1, 1) &: 1 \\ (1, b_1) &: ik_i U_i \\ (b_1, b_1) &: \xi_L - k_i k_j U_i U_j \\ (1, b_2) &: -k_i k_j U_i U_j \\ (b_1, b_2) &: 2ik_i U_i \xi_L - ik_i k_j k_k U_i U_j U_k \\ (b_2, b_2) &: 2\xi_L^2 - 4k_i k_j U_i U_j \xi_L + k_i k_j k_k k_l U_i U_j U_k U_l \\ (1, b_s) &: -k_i k_j B_{iab} B_{jab} \\ (b_1, b_s) &: 2ik_i B_{iab} E_{ab} - ik_i k_j k_k U_i B_{jab} B_{kab} \\ (b_2, b_s) &: 2E_{ab} E_{ab} - 4k_i k_j U_i E_{ab} B_{jab} + k_i k_j k_k k_l U_i U_j B_{kab} B_{lab} \\ (b_s, b_s) &: 2C_{abcd} C_{abcd} - 4k_i k_j B_{iab} B_{jcd} C_{abcd} + k_i k_j k_k k_l B_{iab} B_{jab} B_{kcd} B_{lcd}. \end{aligned} \quad (\text{A.7})$$

Many of the contractions in the above can be written as correlators of the scalar s^2 with densities and displacements. Some of these have been previously defined in other works (see e.g. ref. [26])

$$\begin{aligned} \chi &= \langle \delta_1^2 s_2^2 \rangle_c = 2E_{ab} E_{ab} \quad , \quad \zeta = \langle s_1^2 s_2^2 \rangle_c = 2C_{abcd} C_{abcd} \\ \Upsilon_{ij} &= \langle s^2 \Delta_i \Delta_j \rangle = 2B_{iab} B_{jab} \quad , \quad V_i^{12} = \langle \delta_1 s_2^2 \Delta_i \rangle = 2E_{ab} B_{iab} \end{aligned} \quad (\text{A.8})$$

so we can simplify some of above expressions for basis spectra:

$$\begin{aligned} (1, b_s) &: -\frac{1}{2}k_i k_j \Upsilon_{ij} \\ (b_1, b_s) &: ik_i V_i^{12} - \frac{1}{2}ik_i k_j k_k U_i \Upsilon_{jk} \\ (b_2, b_s) &: \chi - 2k_i k_j U_i V_j^{12} + \frac{1}{2}k_i k_j k_k k_l U_i U_j \Upsilon_{kl} \\ (b_s, b_s) &: \zeta - 4k_i k_j B_{iab} B_{jcd} C_{abcd} + \frac{1}{4}k_i k_j k_k k_l \Upsilon_{ij} \Upsilon_{kl}. \end{aligned} \quad (\text{A.9})$$

The (b_s, b_s) component contains a term that cannot be reduced to previously computed quantities but is rather proportional to $L_{ij} = B_{iab} B_{jcd} C_{abcd}$. This term comes from the expectation value of $(1/2)(ik_i \Delta_i)^2 s_1^2 s_2^2$ which contains

$$\langle \Delta_i \Delta_j s_1^2 s_2^2 \rangle \ni 8 \langle \Delta_i s_{1,ab} \rangle \langle \Delta_j s_{2,cd} \rangle \langle s_{1,ab} s_{2,cd} \rangle$$

leading to the combinatorial factor $(-4) = \frac{1}{2} \times i^2 \times 8$. This contribution can be expressed in terms of its components as

$$L_{ij} = 2K_1 J_3 \delta_{ij} + (2K_1 (J_2 + J_3 + J_4) + K_2 (J_2 + 2J_3 + J_4)) \hat{q}_i \hat{q}_j \quad (\text{A.10})$$

where we the J_n are defined as in Appendix A of ref. [42, 91] and we have defined the quantities

$$K_1 = 2J_3 (J_6 + J_8), \quad K_2 = 2J_3 (J_7 + 2J_8 + J_9) + J_4 (2J_6 + J_7 + 4J_8 + J_9). \quad (\text{A.11})$$

All of the above expressions are implemented in the publicly available code `ZenBu`¹⁴.

¹⁴ZenBu (Idovich calculations for) N-B(ody Em)u(lators); <https://github.com/sfschen/ZenBu>.

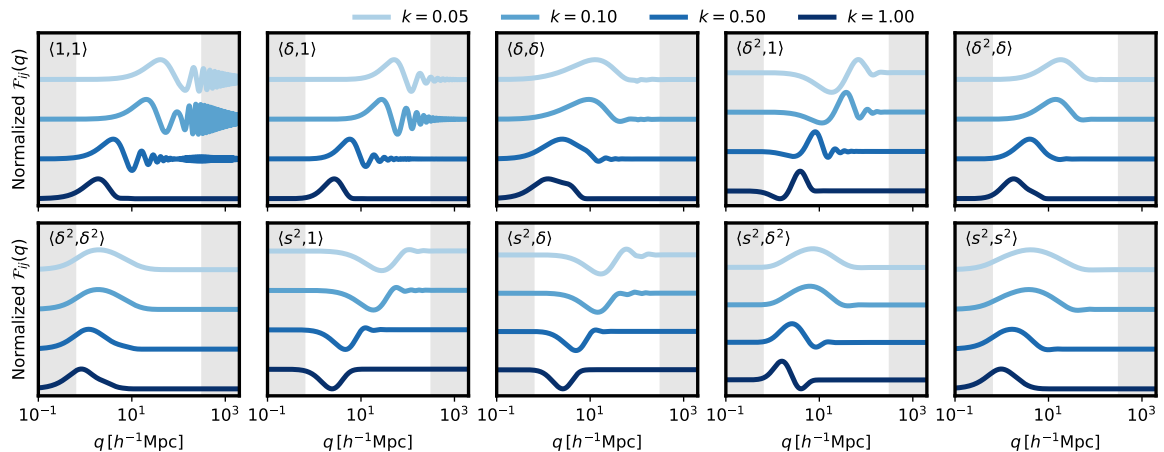


Figure 10. Normalized integrands for Lagrangian basis spectra at scales of $k = [0.05, 0.1, 0.5, 1] h\text{Mpc}^{-1}$. The shaded regions correspond to physical scales that are not well probed by the fiducial Quijote HR simulations. Each integrand is offset (vertically) by an arbitrary value so their features and support can be clearly distinguished.

B Matching analytic and grid-based LPT

The success of the control variates approach we have laid out in this work is predicated on ensuring precise agreement from our two separate predictions for basis spectra in the Zel’dovich approximation. If grid-based calculations do not agree with the analytic result, then the control variate could introduce significant biases, as well as additional uncertainty, in the estimates of the fully non-linear spectra.

In this appendix we explore, in more detail, the structure of the contributions to any given basis spectrum in the analytic code. Any given spectrum $P_{ij}(k)$ is written as in the form of Eqn. A.6. In order to efficiently evaluate the integrals, ZENBU further projects the $F_{ij}(\mathbf{k}, \mathbf{q})$ kernels into a Spherical Fourier-Bessel basis. The angular integrals are then performed analytically, and the radial integrals reduce to Hankel transforms of the form

$$P_{ij}(k) = 4\pi \sum_{\ell=0}^{\infty} \int q^2 dq e^{-k^2(X+Y)/2} \left(\frac{kY}{q}\right)^{\ell} f_{ij}^{(\ell)}(q, k) j_{\ell}(kq), \quad (\text{B.1})$$

where $f_{ij}^{(\ell)}(q)$ is a basis spectrum-dependent kernel. These Hankel transforms can then be rapidly evaluated using the FFTLog algorithm [23, 92, 93]. A summary of the structure of the implementation can be found in Appendix A of ref. [28], with full details available in refs. [24, 25]. In practice, additional regularization is needed to compute this expression stably and this is done by including a cutoff in the linear power spectrum

$$P^{\text{lin}}(k) \rightarrow e^{-(k/k_{\text{cut}})^2} P^{\text{lin}}(k), \quad (\text{B.2})$$

where k_{cut} is a cutoff scale imposed from filtering the linear density modes at initial conditions as

$$\delta(\mathbf{k}) \rightarrow e^{-(k/k_{\text{cut}})^2/2} \delta(\mathbf{k}). \quad (\text{B.3})$$

This same filter is applied to the initial density field used to evaluate Zel’dovich displacements and their subsequent component fields/spectra. We use $k_{\text{cut}} = \pi N_{\text{p}}^{1/3}/L_{\text{box}} \approx 3.2 h\text{Mpc}^{-1}$,

where N_p is the total number of particles in the simulation, corresponding to the smallest Fourier modes probed by the initial conditions of the `Quijote` HR boxes¹⁵. Without this equivalent filtering, we find that we are unable to match grid-based results and analytics, even in the case of $P_{mm}(k)$.

If we re-cast Eq. B.1 in the following form:

$$P_{ij}(k) = \int dq \mathcal{F}_{ij}(q, k), \quad (\text{B.4})$$

where \mathcal{F}_{ij} is given by

$$\mathcal{F}_{ij}(q, k) = \sum_{\ell} e^{-(k^2/2)(X^{\text{lin}}+Y^{\text{lin}})} \left(\frac{kY}{q}\right)^{\ell} f_{ij}^{(\ell)}(q, k) q^2 j_{\ell}(kq), \quad (\text{B.5})$$

then we can assess the scales needed to be probed by our grid-based realization in order to ensure accuracy relative to the analytic calculation. If for a given wavenumber k , $\mathcal{F}_{ij}(q, k)$ has significant support for q that are either larger than the largest resolved mode in the box, q_{max} or smaller than the smallest scale q_{min} , then these contributions will be missed in the grid-based calculation. While at first impression one might assume that $q_{\text{min}} = L_{\text{box}}/N_{\text{mesh}}$, the fundamental grid spacing, and $q_{\text{max}} = L_{\text{box}}/2$, the largest resolvable separation, this is not the case. Power in the N -body simulation is initialized on a Fourier-space grid, which is then Fourier-transformed to configuration space. This imposes two anisotropic window functions which filter small and larger scales. The ‘‘true’’ density field is thus

$$\tilde{\delta}(\mathbf{k}) = W_{\text{UV}}(\mathbf{k})W_{\text{IR}}(\mathbf{k})\delta(\mathbf{k}), \quad (\text{B.6})$$

where the small-scale and large-scale anisotropic window functions are, respectively,

$$W_{\text{UV}}(\mathbf{k}) = \begin{cases} 1, & \text{if } |k_i| \leq \frac{\pi N_{\text{mesh}}}{L_{\text{box}}} \\ 0, & \text{if } |k_i| > \frac{\pi N_{\text{mesh}}}{L_{\text{box}}} \end{cases}, \quad W_{\text{IR}}(\mathbf{k}) = \begin{cases} 1, & \text{if } |k_i| \geq \frac{2\pi}{L_{\text{box}}} \\ 0, & \text{if } |k_i| < \frac{2\pi}{L_{\text{box}}} \end{cases}. \quad (\text{B.7})$$

The sharp Fourier-space windows imposed by the gridded density field will leak power into scales $q < L_{\text{box}}/N_{\text{mesh}}$. In configuration space, these filters are given by sinc functions of the form

$$W_{\text{UV}}(\mathbf{q}) \propto \prod_i \text{sinc}\left(\frac{\pi N_{\text{mesh}} q_i}{L_{\text{box}}}\right). \quad (\text{B.8})$$

If we compare the integral of this window relative to a ‘sharp- q ’ window function, we can define an equivalent q_{min} from assessing when their integrals reach their asymptotic values at $q \rightarrow \infty$. For the sinc function with argument as we’ve defined, this value is reached at roughly $q_{\text{min}} \approx 2/k_{\text{Nyq}}$ ¹⁶. Thus, we adopt $q_{\text{min}} = 2/k_{\text{Nyq}}$ as a heuristic for the smallest configuration-space scales probed by simulations when understanding the extent of integrands $\mathcal{F}_{ij}(q)$. Similarly, we choose $q_{\text{max}} = 2/k_f = L_{\text{box}}/\pi$ as a heuristic scale for the largest configuration-space scale. In Fig. 10, we explicitly visualize the kernels \mathcal{F}_{ij} for $k \in [0.05, 0.1, 0.5, 1]$. We find that across this range of scales, most integrands have support well within the scales probed by the `Quijote` HR boxes. We also note that this same figure explains why we cannot use

¹⁵As `Quijote` is evolved using a tree-based code, nonlinear forces are calculated at scales smaller than k_{cut} .

¹⁶The precise value is closer to $1.926/k_{\text{Nyq}}$ and is related to when the Sine integral $\text{Si}(q_*) = \pi/2$ for the first time.




Figure 11. The resolution dependence of basis spectra at Lagrangian coordinates. *Left:* The ratio of a standard `Quijote` resolution realization, a super-resolution realization and an analytic calculation of the effects of a small scale cutoff on the $\langle s^2 s^2 \rangle$ spectrum compared to grids computed at the fiducial `Quijote` HR resolution of this work. *Right:* The same ratios but for the basis spectrum $\langle s^2 \delta \rangle$. Cutting off the linear theory power spectrum at small scales produces an effect on correlators that is comparable to the effect of reducing grid size.

the fiducial resolution `Quijote` suite: for $k \geq 0.5 h\text{Mpc}^{-1}$ they fail to resolve scales near the peak of the integrands for several of the basis functions. However, we also note that from our heuristic q_{\min} we might be at risk of missing several low- q scales for spectra such as $\langle s^2 s^2 \rangle$ and $\langle \delta^2 s^2 \rangle$.

Another way to assess whether the dynamic range in a grid used for grid-based LPT is sufficient is to consider, instead, Lagrangian correlators. These are the leading contributions to the final advected spectra, but don't mix configuration-space and Fourier-space modes as strongly as in the post-advection case we have considered above. For $\langle \delta^2 s^2 \rangle$ and $\langle s^2 s^2 \rangle$ these correlators are given by Hankel transforms of the functions $\chi(q)$ and $\zeta(q)$ defined in Appendix. A. Discrepancies in Lagrangian correlators will propagate into post-advection measurements, and their origin similarly comes from being unable to resolve all scales that contribute to a given spectrum in a grid-based approach. To illustrate this we consider the *Lagrangian* power spectra

$$\begin{aligned}
 P_{\delta^2 s^2}^{\text{Lag}}(k) &= 4\pi \int q^2 dq j_0(kq) \chi(q), \\
 P_{s^2 s^2}^{\text{Lag}}(k) &= 4\pi \int q^2 dq j_0(kq) \zeta(q).
 \end{aligned}
 \tag{B.9}$$

We measure, in the Lagrangian fields, these basis spectra for three grids that share the same initial seeds and power spectra with varying resolutions. Beyond the fiducial `Quijote` HR resolution, we also consider a grid at the standard `Quijote` resolution of $N_{\text{mesh}} = 512$. To understand whether our `Quijote` HR results are well converged, we also consider correlators measured at a resolution comparable to that of the `Aemulus` suite, using $N_{\text{mesh}} = 1400$. We also compute the spectra in Eqn. B.9 analytically in two different ways. The first is using, as input, the linear theory damped $P(k)$ across all scales. If our grid adequately resolves all relevant physical scales, and the numerical operation of extracting s_{ij} is well-converged, then the comparison between $P_{\delta^2 s^2}^{\text{Lag}}(k)$ computed using Eqn. B.9 and as measured in the Lagrangian fields of our simulations should be excellent at small scales where sample variance

is negligible. We have verified that in our boxes this is the case – residuals between grid-based and analytic Lagrangian correlators agree to within 0.2% at $k = 1 h\text{Mpc}^{-1}$.

Our second calculation uses analytic window functions that are *spherically* sharp in k space, with $k_{\text{UV}} = 512\pi/L_{\text{box}}$ and $k_{\text{IR}} = 2\pi/L_{\text{box}}$. While these window functions are different from how the grid-based calculations are treated, their inclusion in the analytic calculation allows for comparison with the downsampled result at $N_{\text{mesh}} = 512$. We increase k_{UV} by 20% to approximately convert the spherical damping back to the cartesian case employed in grid-based calculations, resulting in $k_{\text{UV}} \approx 2h\text{Mpc}^{-1}$. We show these results in Fig. 11, where we display the ratio of power spectra at various resolutions relative to the fiducial **Quijote** HR realization. The higher resolution realization has less than a percent additional power at $k = 1 h\text{Mpc}^{-1}$ for $\langle s^2 s^2 \rangle$ while the excess is of order a percent for $\langle s^2 \delta^2 \rangle$. This excess for $\langle s^2 \delta^2 \rangle$ can also be seen in Fig. 3, implying for that specific spectrum we would find potentially better convergence with the analytic ZA result at a higher resolution. The comparison with the **Aemulus** resolution boxes shows that we are not missing small-scale modes for the **Quijote** HR-based analysis.

Turning to the standard **Quijote** resolution realizations, we find a large absence of power for both Lagrangian correlators. At $k = 1 h\text{Mpc}^{-1}$ the deficit for $\langle s^2 s^2 \rangle$ is on the order of 10%, and both this amount as well as the shape of the decay is well matched by our theoretical result smoothed by a spherical sharp- k window. For $\langle s^2 \delta^2 \rangle$ the effects are even stronger, we find on the order of 80% missing power at $k = 1 h\text{Mpc}^{-1}$. Again, the observed damping is well described by convolving the input theory spectrum with a spherical sharp- k window function before computing correlators in **ZeNBu**. Fig. 11 shows that the standard resolution **Quijote** simulations do not possess sufficient dynamic range for Zel’dovich control variates to be used effectively to scales of $k = 1 h\text{Mpc}^{-1}$.

Future attempts to use control variates should be cognizant of the analysis carried out in this Appendix. Understanding accuracy requirements at different scales for emulation, as well as what the dynamic range required in a simulation/grid-based LPT in order to achieve satisfactory prediction of the mean of a control variate can be obtained via the arguments laid out here.

C Challenges in going beyond Zel’dovich

In this work we have concerned ourselves solely with using the Zel’dovich approximation as the control variate that correlates with the non-linear density field. It is natural, then, to wonder whether even further reduction in sample variance can be achieved with a surrogate that is more closely correlated to the non-linear density, but is still analytic. The natural candidate for such a surrogate is the density field predicted from displacements carried out at higher order in LPT. Indeed, the state of the field is such that there now exist mature codes that compute basis spectra in one-loop LPT (such as **velocileptors**) and using grid-based schemes, such as **monofonIC** [51].

However, there is an additional subtlety if one wishes to compare 3LPT between an analytic prediction and a grid-based realization at the same order. Most analytic calculation codes for higher order LPT implement so-called *convolution Lagrangian perturbation theory* [21]. CLPT performs a re-summation of certain terms that arise in the expansion of connected moments, $\langle \mathcal{M}^n \rangle_c$, with \mathcal{M} as defined in Eqn. A.1. In the Zel’dovich approximation, all ingredients in the expansion are Gaussian and thus there are no connected moments beyond second order.

However, when Δ contains contributions from higher order displacements the above statement is no longer true. Now, connected moments beyond the second exist and will contribute to the evaluation of \mathcal{Z} . For example, a contribution arises from the correlation of three displacements at second order

$$W_{ijk} = \langle \Delta_i \Delta_j \Delta_k \rangle. \quad (\text{C.1})$$

In CLPT, and other perturbative approaches, higher order terms like W_{ijk} are treated as small and expanded out of the exponential in \mathcal{Z} ¹⁷. On the other hand, grid-based approaches fully evaluate this exponential. These differences lead to discrepancies between grid-based and analytic LPT when trying to extend the use of surrogates to higher order. Attempts to use higher order LPT as a control variate must, then, deal with this issue before being practical. We leave this question to further work, as we find the performance of the Zel'dovich approximation as a control variate to be adequate for the purposes of a first study.

References

- [1] B. Dai, Y. Feng and U. Seljak, *A gradient based method for modeling baryons and matter in halos of fast simulations*, *Journal of Cosmology and Astroparticle Physics* **2018** (2018) 009.
- [2] B. Dai, Y. Feng, U. Seljak and S. Singh, *High mass and halo resolution from fast low resolution simulations*, *Journal of Cosmology and Astroparticle Physics* **2020** (2020) 002.
- [3] B. Dai and U. Seljak, *Learning effective physical laws for generating cosmological hydrodynamics with lagrangian deep learning*, *Proceedings of the National Academy of Sciences* **118** (2021) .
- [4] Y. Li, Y. Ni, R.A.C. Croft, T.D. Matteo, S. Bird and Y. Feng, *AI-assisted superresolution cosmological simulations*, *Proceedings of the National Academy of Sciences* **118** (2021) .
- [5] D. Schaurecker, Y. Li, J. Tinker, S. Ho and A. Refregier, *Super-resolving Dark Matter Halos using Generative Deep Learning*, [2111.06393](#).
- [6] K. Heitmann, D. Bingham, E. Lawrence, S. Bergner, S. Habib, D. Higdon et al., *The Mira-Titan Universe: Precision Predictions for Dark Energy Surveys*, *Astrophys. J.* **820** (2016) 108 [[1508.02654](#)].
- [7] Euclid Collaboration, M. Knabenhans, J. Stadel, D. Potter, J. Dakin, S. Hannestad et al., *Euclid preparation: IX. EuclidEmulator2 - power spectrum emulation with massive neutrinos and self-consistent dark energy perturbations*, *Mon. Not. R. Astron. Soc.* **505** (2021) 2840 [[2010.11288](#)].
- [8] J. DeRose, R.H. Wechsler, J.L. Tinker, M.R. Becker, Y.-Y. Mao, T. McClintock et al., *The Aemulus Project I: Numerical Simulations for Precision Cosmology*, *Astrophys. J.* **875** (2019) 69 [[1804.05865](#)].
- [9] N.A. Maksimova, L.H. Garrison, D.J. Eisenstein, B. Hadzhiyska, S. Bose and T.P. Satterthwaite, *AbacusSummit: a massive set of high-accuracy, high-resolution N-body simulations*, *Mon. Not. Roy. Astron. Soc.* **508** (2021) 4017 [[2110.11398](#)].
- [10] A. Pontzen, A. Slosar, N. Roth and H.V. Peiris, *Inverted initial conditions: Exploring the growth of cosmic structure and voids*, *Physical Review D* **93** (2016) .
- [11] R.E. Angulo and A. Pontzen, *Cosmological N-body simulations with suppressed variance*, *Mon. Not. Roy. Astron. Soc.* **462** (2016) L1 [[1603.05253](#)].

¹⁷See [94] for a discussion specifically about expanding W_{ijk} or not. There are many terms beyond W_{ijk} which are expanded from this exponential in perturbative approaches.

- [12] F. Villaescusa-Navarro, S. Naess, S. Genel, A. Pontzen, B. Wandelt, L. Anderson et al., *Statistical properties of paired fixed fields*, *Astrophys. J.* **867** (2018) 137 [[1806.01871](#)].
- [13] C.-H. Chuang et al., *UNIT project: Universe N-body simulations for the Investigation of Theoretical models from galaxy surveys*, *Mon. Not. Roy. Astron. Soc.* **487** (2019) 48 [[1811.02111](#)].
- [14] L. Anderson, A. Pontzen, A. Font-Ribera, F. Villaescusa-Navarro, K.K. Rogers and S. Genel, *Cosmological Hydrodynamic Simulations with Suppressed Variance in the Ly α Forest Power Spectrum*, *Astrophys. J.* **871** (2019) 144 [[1811.00043](#)].
- [15] A. Klypin, F. Prada and J. Byun, *Suppressing cosmic variance with paired-and-fixed cosmological simulations: average properties and covariances of dark matter clustering statistics*, *Monthly Notices of the Royal Astronomical Society* **496** (2020) 3862.
- [16] N. Chartier, B. Wandelt, Y. Akrami and F. Villaescusa-Navarro, *CARPool: fast, accurate computation of large-scale structure statistics by pairing costly and cheap cosmological simulations*, *arXiv e-prints* (2020) arXiv:2009.08970 [[2009.08970](#)].
- [17] N. Chartier and B.D. Wandelt, *Carpool covariance: Fast, unbiased covariance estimation for large-scale structure observables*, *Monthly Notices of the Royal Astronomical Society* (2021) .
- [18] N. Chartier and B.D. Wandelt, *Bayesian Control Variates for optimal covariance estimation with pairs of simulations and surrogates*, [2204.03070](#).
- [19] T. Matsubara, *Resumming cosmological perturbations via the lagrangian picture: One-loop results in real space and in redshift space*, *Physical Review D* **77** (2008) .
- [20] T. Matsubara, *Nonlinear perturbation theory with halo bias and redshift-space distortions via the Lagrangian picture*, *Phys. Rev. D* **78** (2008) 083519 [[0807.1733](#)].
- [21] J. Carlson, B. Reid and M. White, *Convolution lagrangian perturbation theory for biased tracers*, *Monthly Notices of the Royal Astronomical Society* **429** (2012) 1674–1685.
- [22] R.A. Porto, L. Senatore and M. Zaldarriaga, *The Lagrangian-space Effective Field Theory of large scale structures*, *Journal of Cosmology and Astroparticle Physics* **2014** (2014) 022 [[1311.2168](#)].
- [23] Z. Vlah, M. White and A. Aviles, *A Lagrangian effective field theory*, *Journal of Cosmology and Astroparticle Physics* **2015** (2015) 014 [[1506.05264](#)].
- [24] Z. Vlah, E. Castorina and M. White, *The gaussian streaming model and convolution lagrangian effective field theory*, *Journal of Cosmology and Astroparticle Physics* **2016** (2016) 007–007.
- [25] S.-F. Chen, Z. Vlah and M. White, *Consistent modeling of velocity statistics and redshift-space distortions in one-loop perturbation theory*, *Journal of Cosmology and Astroparticle Physics* **2020** (2020) 062–062.
- [26] S.-F. Chen, Z. Vlah, E. Castorina and M. White, *Redshift-space distortions in lagrangian perturbation theory*, *Journal of Cosmology and Astroparticle Physics* **2021** (2021) 100.
- [27] S.-F. Chen, Z. Vlah and M. White, *A new analysis of galaxy 2-point functions in the BOSS survey, including full-shape information and post-reconstruction BAO*, *Journal of Cosmology and Astroparticle Physics* **2022** (2022) 008 [[2110.05530](#)].
- [28] M. White, R. Zhou, J. DeRose, S. Ferraro, S.-F. Chen, N. Kokron et al., *Cosmological constraints from the tomographic cross-correlation of DESI Luminous Red Galaxies and Planck CMB lensing*, *Journal of Cosmology and Astro-Particle Physics* **2022** (2022) 007 [[2111.09898](#)].
- [29] C. Modi, S.-F. Chen and M. White, *Simulations and symmetries*, *Mon. Not. Roy. Astron. Soc.* **492** (2020) 5754 [[1910.07097](#)].

- [30] N. Kokron, J. DeRose, S.-F. Chen, M. White and R.H. Wechsler, *The cosmology dependence of galaxy clustering and lensing from a hybrid n-body–perturbation theory model*, *Monthly Notices of the Royal Astronomical Society* **505** (2021) 1422–1440.
- [31] B. Hadzhiyska, C. García-García, D. Alonso, A. Nicola and A. Slosar, *Hefty enhancement of cosmological constraints from the des y1 data using a hybrid effective field theory approach to galaxy bias*, 2021.
- [32] M. Zennaro, R.E. Angulo, M. Pellejero-Ibáñez, J. Stücker, S. Contreras and G. Aricò, *The bacco simulation project: biased tracers in real space*, 2021.
- [33] A. Banerjee, N. Kokron and T. Abel, *Modeling nearest neighbor distributions of biased tracers using hybrid effective field theory*, 2021.
- [34] M. Zennaro, R.E. Angulo, S. Contreras, M. Pellejero-Ibáñez and F. Maion, *Priors on lagrangian bias parameters from galaxy formation modelling*, 2021.
- [35] N. Kokron, J. DeRose, S.-F. Chen, M. White and R.H. Wechsler, *Priors on red galaxy stochasticity from hybrid effective field theory*, [2112.00012](#).
- [36] Y.B. Zel’Dovich, *Gravitational instability: an approximate theory for large density perturbations.*, *Astronomy & Astrophysics* **500** (1970) 13.
- [37] A.B. Owen, *Monte Carlo theory, methods and examples* (2013).
- [38] Z. Ding et al., *The DESI N-body Simulation Project II: Suppressing Sample Variance with Fast Simulations*, [2202.06074](#).
- [39] Y. Feng, M.-Y. Chu, U. Seljak and P. McDonald, *FASTPM: a new scheme for fast simulations of dark matter and haloes*, *MNRAS* **463** (2016) 2273 [[1603.00476](#)].
- [40] S. Tassev, M. Zaldarriaga and D.J. Eisenstein, *Solving large scale structure in ten easy steps with COLA*, *Journal of Cosmology and Astroparticle Physics* **2013** (2013) 036.
- [41] S. Tassev, D.J. Eisenstein, B.D. Wandelt and M. Zaldarriaga, *scola: The n-body cola method extended to the spatial domain*, 2015. 10.48550/ARXIV.1502.07751.
- [42] M. White, *The Zel’dovich approximation*, *Mon. Not. Roy. Astron. Soc.* **439** (2014) 3630 [[1401.5466](#)].
- [43] F. Villaescusa-Navarro, C. Hahn, E. Massara, A. Banerjee, A.M. Delgado, D.K. Ramanah et al., *The quijote simulations*, *The Astrophysical Journal Supplement Series* **250** (2020) 2.
- [44] P. Coles, A.L. Melott and S.F. Shandarin, *Testing approximations for non-linear gravitational clustering*, *Monthly Notices of the Royal Astronomical Society* **260** (1993) 765 [<https://academic.oup.com/mnras/article-pdf/260/4/765/2911119/mnras260-0765.pdf>].
- [45] A.L. Melott, T. Buchert and A.G. Weiss, *Testing higher order Lagrangian perturbation theory against numerical simulations. 2: Hierarchical models*, *Astron. Astrophys.* **294** (1995) 345 [[astro-ph/9404018](#)].
- [46] S. Tassev and M. Zaldarriaga, *The mildly non-linear regime of structure formation*, *Journal of Cosmology and Astroparticle Physics* **2012** (2012) 013–013.
- [47] S. Tassev, *Lagrangian or Eulerian; Real or Fourier? Not All Approaches to Large-Scale Structure Are Created Equal*, *JCAP* **06** (2014) 008 [[1311.4884](#)].
- [48] F. Schmidt, *An n-th order lagrangian forward model for large-scale structure*, *Journal of Cosmology and Astroparticle Physics* **2021** (2021) 033.
- [49] V. Desjacques, D. Jeong and F. Schmidt, *Large-Scale Galaxy Bias*, *Phys. Rept.* **733** (2018) 1 [[1611.09787](#)].

- [50] E. Castorina, C. Carbone, J. Bel, E. Sefusatti and K. Dolag, *DEMNUi: the clustering of large-scale structures in the presence of massive neutrinos*, *Journal of Cosmology and Astroparticle Physics* **2015** (2015) 043 [[1505.07148](#)].
- [51] M. Michaux, O. Hahn, C. Rampf and R.E. Angulo, *Accurate initial conditions for cosmological N-body simulations: Minimizing truncation and discreteness errors*, *Mon. Not. Roy. Astron. Soc.* **500** (2020) 663 [[2008.09588](#)].
- [52] F. Bernardeau, S. Colombi, E. Gaztanaga and R. Scoccimarro, *Large scale structure of the universe and cosmological perturbation theory*, *Phys. Rept.* **367** (2002) 1 [[astro-ph/0112551](#)].
- [53] A.N. Taylor and A.J.S. Hamilton, *Non-linear cosmological power spectra in real and redshift space*, *Monthly Notices of the Royal Astronomical Society* **282** (1996) 767.
- [54] M. Crocce, S. Pueblas and R. Scoccimarro, *Transients from Initial Conditions in Cosmological Simulations*, *Mon. Not. Roy. Astron. Soc.* **373** (2006) 369 [[astro-ph/0606505](#)].
- [55] C. Rampf and O. Hahn, *Shell-crossing in a λ CDM universe*, *Monthly Notices of the Royal Astronomical Society: Letters* **501** (2020) L71.
- [56] T. Baldauf, E. Schaan and M. Zaldarriaga, *On the reach of perturbative methods for dark matter density fields*, *JCAP* **03** (2016) 007 [[1507.02255](#)].
- [57] U. Seljak, G. Aslanyan, Y. Feng and C. Modi, *Towards optimal extraction of cosmological information from nonlinear data*, *Journal of Cosmology and Astroparticle Physics* **2017** (2017) 009.
- [58] C. Modi, M. White, A. Slosar and E. Castorina, *Reconstructing large-scale structure with neutral hydrogen surveys*, *Journal of Cosmology and Astroparticle Physics* **2019** (2019) 023.
- [59] C. Modi, M. White, E. Castorina and A. Slosar, *Mind the gap: the power of combining photometric surveys with intensity mapping*, *Journal of Cosmology and Astroparticle Physics* **2021** (2021) 056.
- [60] A. Taruya, T. Nishimichi and D. Jeong, *Grid-based calculation for perturbation theory of large-scale structure*, *Physical Review D* **98** (2018) .
- [61] A. Taruya, T. Nishimichi and D. Jeong, *Covariance of the matter power spectrum including the survey window function effect: N -body simulations versus fifth-order perturbation theory on grids*, *Phys. Rev. D* **103** (2021) 023501 [[2007.05504](#)].
- [62] A. Taruya, T. Nishimichi and D. Jeong, *Grid-based calculations of redshift-space matter fluctuations from perturbation theory: Uv sensitivity and convergence at the field level*, 2021. 10.48550/ARXIV.2109.06734.
- [63] K. Heitmann, M. White, C. Wagner, S. Habib and D. Higdon, *The Coyote Universe I: Precision Determination of the Nonlinear Matter Power Spectrum*, *Astrophys. J.* **715** (2010) 104 [[0812.1052](#)].
- [64] M. Zennaro, J. Bel, F. Villaescusa-Navarro, C. Carbone, E. Sefusatti and L. Guzzo, *Initial conditions for accurate n-body simulations of massive neutrino cosmologies*, *Monthly Notices of the Royal Astronomical Society* **466** (2016) 3244.
- [65] A. Schneider, R. Teyssier, D. Potter, J. Stadel, J. Onions, D.S. Reed et al., *Matter power spectrum and the challenge of percent accuracy*, *Journal of Cosmology and Astroparticle Physics* **2016** (2016) 047–047.
- [66] L.H. Garrison, D.J. Eisenstein, D. Ferrer, M.V. Metchnik and P.A. Pinto, *Improving Initial Conditions for Cosmological N-Body Simulations*, *Mon. Not. Roy. Astron. Soc.* **461** (2016) 4125 [[1605.02333](#)].
- [67] Y. Dubois, J. Devriendt, A. Slyz and R. Teyssier, *Self-regulated growth of supermassive black holes by a dual jet-heating active galactic nucleus feedback mechanism: methods, tests and*

- implications for cosmological simulations, *Monthly Notices of the Royal Astronomical Society* **420** (2012) 2662.
- [68] A.M.C.L. Brun, I.G. McCarthy, J. Schaye and T.J. Ponman, *Towards a realistic population of simulated galaxy groups and clusters*, *Monthly Notices of the Royal Astronomical Society* **441** (2014) 1270.
- [69] J. Schaye, R.A. Crain, R.G. Bower, M. Furlong, M. Schaller, T. Theuns et al., *The EAGLE project: simulating the evolution and assembly of galaxies and their environments*, *Monthly Notices of the Royal Astronomical Society* **446** (2014) 521.
- [70] N. Khandai, T.D. Matteo, R. Croft, S. Wilkins, Y. Feng, E. Tucker et al., *The MassiveBlack-II simulation: the evolution of haloes and galaxies to $z \sim 0$* , *Monthly Notices of the Royal Astronomical Society* **450** (2015) 1349.
- [71] I.G. McCarthy, J. Schaye, S. Bird and A.M.C. Le Brun, *The BAHAMAS project: calibrated hydrodynamical simulations for large-scale structure cosmology*, *Monthly Notices of the Royal Astronomical Society* **465** (2017) 2936 [1603.02702].
- [72] R. Davé, D. Anglés-Alcázar, D. Narayanan, Q. Li, M.H. Rafieferantsoa and S. Appleby, *simba: Cosmological simulations with black hole growth and feedback*, *Monthly Notices of the Royal Astronomical Society* **486** (2019) 2827.
- [73] D. Nelson, V. Springel, A. Pillepich, V. Rodriguez-Gomez, P. Torrey, S. Genel et al., *The illustrating simulations: Public data release*, 2021.
- [74] H. Trac, R. Cen and P. Mansfield, *SCORCH. i. THE GALAXY–HALO CONNECTION IN THE FIRST BILLION YEARS*, *The Astrophysical Journal* **813** (2015) 54.
- [75] J.U. Lange, F.C. van den Bosch, A.R. Zentner, K. Wang, A.P. Hearin and H. Guo, *Cosmological evidence modelling: a new simulation-based approach to constrain cosmology on non-linear scales*, *Monthly Notices of the Royal Astronomical Society* **490** (2019) 1870.
- [76] S. Yuan, L.H. Garrison, B. Hadzhiyska, S. Bose and D.J. Eisenstein, *AbacusHOD: a highly efficient extended multitracer HOD framework and its application to BOSS and eBOSS data*, *Monthly Notices of the Royal Astronomical Society* **510** (2021) 3301.
- [77] J.U. Lange, A.P. Hearin, A. Leauthaud, F.C. van den Bosch, H. Guo and J. DeRose, *Five per cent measurements of the growth rate from simulation-based modelling of redshift-space clustering in BOSS LOWZ*, *Monthly Notices of the Royal Astronomical Society* **509** (2021) 1779.
- [78] S. Yuan, L.H. Garrison, D.J. Eisenstein and R.H. Wechsler, *Stringent σ_8 constraints from small-scale galaxy clustering using a hybrid MCMC+emulator framework*, [2203.11963](https://arxiv.org/abs/2203.11963).
- [79] Z. Zhai, J.L. Tinker, A. Banerjee, J. DeRose, H. Guo, Y.-Y. Mao et al., *The Aemulus Project V: Cosmological constraint from small-scale clustering of BOSS galaxies*, *arXiv e-prints* (2022) arXiv:2203.08999 [2203.08999].
- [80] Z. Zhai, J.L. Tinker, C. Hahn, H.-J. Seo, M.R. Blanton, R. Tojeiro et al., *The clustering of luminous red galaxies at $z \sim 0.7$ from eboss and boss data*, *The Astrophysical Journal* **848** (2017) 76.
- [81] R.E. Angulo, M. Zennaro, S. Contreras, G. Aricò, M. Pellejero-Ibañez and J. Stücker, *The bacco simulation project: Exploiting the full power of large-scale structure for cosmology*, 2021.
- [82] F. Maion, R.E. Angulo and M. Zennaro, *Statistics of biased tracers in variance-suppressed simulations*, *arXiv e-prints* (2022) arXiv:2204.03868 [2204.03868].
- [83] DESI collaboration, *The DESI Experiment Part I: Science, Targeting, and Survey Design*, *arXiv e-prints* (2016) [1611.00036].

- [84] S. Tassev, *N-point statistics of large-scale structure in the zel'dovich approximation*, *Journal of Cosmology and Astroparticle Physics* **2014** (2014) 012.
- [85] S. Avila and A.G. Adame, *Validating galaxy clustering models with Fixed & Paired and Matched-ICs simulations: application to Primordial Non-Gaussianities*, [2204.11103](#).
- [86] V. Assassi, D. Baumann and F. Schmidt, *Galaxy bias and primordial non-gaussianity*, *Journal of Cosmology and Astroparticle Physics* **2015** (2015) 043.
- [87] N. Hand, Y. Feng, F. Beutler, Y. Li, C. Modi, U. Seljak et al., *nbodykit: An open-source, massively parallel toolkit for large-scale structure*, *The Astronomical Journal* **156** (2018) 160.
- [88] C.R. Harris, K.J. Millman, S.J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau et al., *Array programming with NumPy*, *Nature* **585** (2020) 357–362.
- [89] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau et al., *SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python*, *Nature Methods* **17** (2020) 261.
- [90] J.D. Hunter, *Matplotlib: A 2d graphics environment*, *Computing in Science Engineering* **9** (2007) 90.
- [91] M. White, *Erratum: The Zel'dovich approximation*, *Monthly Notices of the Royal Astronomical Society* **447** (2015) 2169.
- [92] A.J.S. Hamilton, *Uncorrelated modes of the non-linear power spectrum*, *Monthly Notices of the Royal Astronomical Society* **312** (2000) 257.
- [93] M. Schmittfull, Z. Vlah and P. McDonald, *Fast large scale structure perturbation theory using one-dimensional fast fourier transforms*, *Physical Review D* **93** (2016) .
- [94] Z. Vlah, U. Seljak and T. Baldauf, *Lagrangian perturbation theory at one loop order: Successes, failures, and improvements*, *Phys. Rev. D* **91** (2015) 023508 [[1410.1617](#)].