

THE DIVIDE-AND-CONQUER SEQUENTIAL MONTE CARLO ALGORITHM: THEORETICAL PROPERTIES AND LIMIT THEOREMS

BY JUAN KUNTZ¹, FRANCESCA R. CRUCINIO¹ AND ADAM M. JOHANSEN¹

¹*Department of Statistics, University of Warwick, Coventry, U.K,
juan.kuntz-nussio@warwick.ac.uk; francesca.crucinio@warwick.ac.uk; a.m.johansen@warwick.ac.uk.*

We provide a comprehensive characterisation of the theoretical properties of the divide-and-conquer sequential Monte Carlo (DaC-SMC) algorithm. We firmly establish it as a well-founded method by showing that it possesses the same basic properties as conventional sequential Monte Carlo (SMC) algorithms do. In particular, we derive pertinent laws of large numbers, L^p inequalities, and central limit theorems; and we characterize the bias in the normalized estimates produced by the algorithm and argue the absence thereof in the unnormalized ones. We further consider its practical implementation and several interesting variants; obtain expressions for its globally and locally optimal intermediate targets, auxiliary measures, and proposal kernels; and show that, in comparable conditions, DaC-SMC proves more statistically efficient than its direct SMC analogue. We close the paper with a discussion of our results, open questions, and future research directions.

CONTENTS

| | | |
|-----|--|----|
| 1 | Introduction | 2 |
| 1.1 | Relation to the literature and contributions | 3 |
| 1.2 | Paper structure | 4 |
| 2 | Background: Product-form estimators and sequential Monte Carlo | 4 |
| 2.1 | Product-form estimators | 4 |
| 2.2 | Sequential Monte Carlo | 6 |
| 3 | The divide-and-conquer sequential Monte Carlo algorithm | 8 |
| 3.1 | Problem definition | 8 |
| 3.2 | The algorithm | 10 |
| 3.3 | Efficient computation | 13 |
| 3.4 | Variants | 14 |
| 4 | Theoretical characterization | 16 |
| 4.1 | Methods of proof | 18 |
| 4.2 | Optimal intermediate targets, auxiliary measures, and proposal kernels | 19 |
| 4.3 | Comparison with ASMC | 21 |
| 5 | Discussion | 22 |
| A | A further example | 24 |
| A.1 | The DaC-SMC approach | 25 |
| A.2 | An alternative approach | 26 |
| B | Important notation for the proofs | 26 |
| C | Proofs of Theorems 1, 2, and 5 | 27 |
| D | Proof of Theorem 3 | 32 |
| E | Proof of Theorem 4 | 33 |
| F | Proof of Theorem 6 | 37 |

MSC2020 subject classifications: Primary 65C05; secondary 60F05, 60F15, 62F15, 68W15.

Keywords and phrases: strong law of large numbers, central limit theorem, interacting particle systems, product-form estimators, distributed computing, Bayesian inference.

| | | |
|---|----------------------------------|----|
| G | Proof of the Theorem 7 | 40 |
| H | Proof of Theorem 8 | 46 |
| I | Proof of Theorem 9 | 48 |
| | Funding | 50 |
| | References | 50 |

1. Introduction. Distributed inference algorithms that scale efficiently to large models and data sets are of widespread interest at the present time. Divide-and-conquer methods, which recursively decompose an inference problem into smaller problems whose solutions can be combined in a principled manner (e.g. see [50]), provide a promising approach in this setting. In computational statistics many methods have been developed which allow for distributed inference either by performing inference for a subset of the data at each processor and ultimately combining their results [17, 52, 55, 62, 66] or by carrying out part of the computation associated with each step of an algorithm at a separate processor and minimizing the necessary communication [34, 58, 60, 64]. Another strategy which is beginning to find applications ranging from general distributed Bayesian inference [8] through inference in broad model classes such as general state space hidden Markov models [13, 24, 25] to inference in specific domains such as phylogenetics [35] is the *divide-and-conquer sequential Monte Carlo (DaC-SMC)* algorithm introduced in [47]. The DaC-SMC algorithm recursively breaks down large inferential problems into smaller, more manageable ones and merges their solutions in a systematic manner. It is particularly amenable to distributed or parallelized implementations (e.g. [13] or [47, Section 5.3]).

The goal of this paper to provide a comprehensive theoretical foundation for this algorithm. The key conceptual innovation that enables us to do so is re-interpreting the DaC-SMC algorithm as a careful combination of sequential Monte Carlo (SMC) and a little-known type of ‘product-form’ Monte Carlo estimators (which, once identified, we studied separately in a preliminary paper [43]).

Product-form estimators encompass a class of generalized U-statistics [45, 41] that exploit conditional independences present in the measure targeted, or the proposal distribution employed, to achieve lower variances than standard Monte Carlo estimators [43]. The idea underpinning them is very simple: if $(X_1^1, \dots, X_1^N), \dots, (X_K^1, \dots, X_K^N)$ are K independent sequences of samples respectively drawn from probability distributions μ_1, \dots, μ_K , then every ‘permutation’ of these samples has the same product law:

$$(X_1^{n_1}, \dots, X_K^{n_K}) \sim \mu := \mu_1 \times \dots \times \mu_K \quad \forall (n_1, \dots, n_K) \in \{1, \dots, N\}^K.$$

In other words, using only N samples from each marginal μ_1, \dots, μ_K , we can construct N^K samples from their product μ . Product-form estimators approximate μ , and averages thereof, using the empirical distribution of these N^K samples (as opposed to the default choice involving the empirical distribution of the N ‘unpermuted’ samples $(X_1^n, \dots, X_K^n)_{n=1}^N$). This increase in sample number leads to a decrease in estimator variance, and we say that product-form estimators are more *statistically efficient* than their conventional counterparts. The empirical distribution μ_{\times}^N of the N^K samples can be expressed as the product of the ‘marginal’ empirical distributions $\mu_1^N := N^{-1} \sum_{n=1}^N \delta_{X_1^n}, \dots, \mu_K^N := N^{-1} \sum_{n=1}^N \delta_{X_K^n}$ (hence, the term ‘product-form estimators’). For this reason, product-form estimators embody the fork-and-join model [12]: they ‘fork’ the problem of approximating μ into that of approximating its marginals, solve these smaller problems by computing μ_1^N, \dots, μ_K^N separately, and then ‘join’ their solutions by taking their product, so obtaining a solution (μ_{\times}^N) for the original problem.

Product-form estimators find their greatest use not in isolation, but when embedded within more sophisticated Monte Carlo methodology to tackle the aspects of the problem exhibiting product structure or conditional independences (e.g. [63, 1, 61, 43]). Here, we study

such embeddings within SMC algorithms, a class of Monte Carlo methods that approximate sequences of distributions defined on state spaces of increasing dimension by iteratively mutating and selecting weighted populations of samples or *particles* (see, for example, [11] for a recent book-length treatment). These algorithms are very widely employed in the context of hidden Markov models [27, 38, 29] and, following [9, 23] and related works, have been shown to be a useful alternative to Markov chain Monte Carlo for approximating general distributions.

The DaC-SMC algorithm itself estimates targets on medium-to-high dimensional spaces that are products of collections of lower dimensional spaces. In order to accomplish this, an auxiliary tree is defined. Each node of this tree has some subset of the variables within the original model assigned to it and each subtree of the tree has associated with it an intermediate model over the variables assigned to the nodes in the subtree. The algorithm progressively evolves clouds of weighted particles up the tree (from leaves to root) so that the corresponding weighted empirical distributions target the intermediate distributions. At each node, the clouds corresponding to the node’s children are merged using a product-form estimator. Hence, the algorithm exemplifies the divide-and-conquer paradigm: carrying computations in the partial order dictated by the indexing tree, DaC-SMC recursively constructs product-form estimators for each of the lower-dimensional intermediate targets and, ultimately, for the original target itself.

While some preliminary work was done in [47], the algorithm’s theoretical properties remain underexplored. The main aim of this paper is to fill this gap and firmly establish DaC-SMC as a well-founded algorithm: we show that it possesses the same fundamental theoretical properties as conventional SMC (Section 1.1). We also have several other, more minor, associated aims: (Section 3.3) exploring efficient implementations of the algorithm; (Section 3.4) identifying several interesting variants thereof in addition to those already pointed out in [47, Section 4]; (Section 4.2) deriving the optimal and ‘locally optimal’ intermediate targets, auxiliary measures, and proposal kernels and giving guidance on their choice in practice; and (Section 4.3) formally arguing that DaC-SMC estimators are more statistically efficient than their direct SMC analogues whenever both algorithms are placed on an equal footing (this does not, however, necessarily mean that DaC-SMC is more computationally efficient than SMC as the former’s cost is generally greater than the latter’s; see Sections 3.3 and 4.3 for more on this).

1.1. *Relation to the literature and contributions.* The most relevant publication is [47] where the DaC-SMC algorithm was introduced, various variants were considered, several numerical examples were given, and it was shown that some forms of DaC-SMC estimators for the normalized targets are consistent and those for the normalizing constants are unbiased (cf. Section 3.1 for definitions the ‘unnormalized targets’, ‘normalized targets’, and ‘normalizing constants’). Both in the original paper [47] and also in later work, e.g. [24, 13], the algorithm was found to have good empirical properties. Applications of the product-form estimators it builds on can be found peppered throughout the Monte Carlo literature [63, 47, 1, 61, 43], almost always unnamed and specialized to particular contexts. Their theoretical properties, at least in any generality, were seemingly only studied recently in [43] which established their consistency, unbiasedness, asymptotically normality, and statistical efficiency, examined their computational efficiency, and considered several variants and extensions. Here, we extend these theoretical properties to DaC-SMC in much the same way that the properties of basic importance sampling were extended to standard SMC over the course of a series of papers (e.g. [19, 20, 53, 14, 10, 42, 15]).

The contributions of this paper are several:

- *DaC-SMC’s theoretical characterization.* We obtain a comprehensive description of the algorithm’s theoretical properties (Section 4). In particular, we show that DaC-SMC estimators for both the unnormalized and normalized targets retain the key properties of their SMC analogues: they are consistent (Theorems 1, 2), asymptotically normal (Theorem 6), and satisfy L^p inequalities (Theorem 5). In the case of the unnormalized targets, we also show that the estimators are unbiased (Theorem 3), while in that of the normalized ones, we show that the estimators’ bias decays linearly with the number of particles (Theorem 4). Our proofs here combine well-known arguments used to analyse SMC algorithms with novel techniques required to deal with the product-form aspects of the estimators. To be more specific about the former, for the laws of large numbers and L^p inequalities we adapt results from [14, 51, 21], for the asymptotic normality we borrow from [21, 22], for the unbiasedness of the unnormalized target estimators we emulate [2], and for the bias bound of the normalized target estimator we follow [56]. As for the latter, the new techniques involve carefully exploiting conditional independences and lead to results (e.g. Theorem 7) that are, to the best of our knowledge, unprecedented in the SMC literature. See Section 4.1 for further details.
- *Optimal intermediate targets, auxiliary measures, and proposal kernels.* Using the asymptotic variance expressions obtained in Theorem 6, we derive (Section 4.2) the intermediate targets, auxiliary measures, and proposal kernels (cf. Sections 3.1–3.2 for the definitions of these objects) that lead to zero *asymptotic* variance DaC-SMC estimators for the normalizing constants. A bit unexpectedly, and in contrast with the standard case, we find that these choices do not necessarily lead to these estimators possessing zero *finite-sample* variance. We also generalize the well-known characterization of the ‘locally optimal’ auxiliary measures and proposal kernels of standard SMC to the DaC-SMC setting (Theorem 8).
- *Statistical efficiency of DaC-SMC.* Also using the asymptotic variances expressions in Theorem 6, we show that DaC-SMC estimators are, at least asymptotically, more statistically efficient than their SMC counterparts (Section 4.3): the asymptotic variances of the former are bounded above by those of the latter.

1.2. *Paper structure.* Section 2 reviews the two building blocks underpinning the DaC-SMC algorithm: product-form estimators (Section 2.1) and sequential Monte Carlo (Section 2.2). Section 3 presents the DaC-SMC algorithm: Section 3.1 describes the target measures the algorithm approximates, Section 3.2 the algorithm itself, Section 3.3 efficient implementations thereof, and Section 3.4 several interesting variations (Appendix A contains an example illustrating concepts discussed in Sections 3.3, 3.4). Section 4 details the algorithm’s theoretical properties: Section 4.1 describes the methods of proof we use to argue the properties (to ease the reading, we defer our proofs to Appendices B–I), Section 4.2 studies the choice of intermediate targets, auxiliary measures, and proposal kernels, and Section 4.3 compares DaC-SMC with standard SMC theoretically and computationally. We close the paper with a discussion in Section 5.

2. Background: Product-form estimators and sequential Monte Carlo. The DaC-SMC algorithm combines ideas from two distinct Monte Carlo methods: product-form estimators and sequential Monte Carlo algorithms. Here, we provide a brief overview of each.

2.1. *Product-form estimators.* Product-form estimators [43] are Monte Carlo methods that exploit any conditional independences present in the measure they target, or the proposal distribution they use, to achieve higher statistical efficiency than normal. In particular, consider the task of estimating

$$\mu(\varphi) := \int \varphi(x)\mu(dx)$$

for a given probability distribution μ on a measurable space (E, \mathcal{E}) and a regular enough function φ . Suppose that μ is product-form: it is the product K distributions μ_1, \dots, μ_K on respective measurable spaces $(E_1, \mathcal{E}_1), \dots, (E_K, \mathcal{E}_K)$ satisfying $(E, \mathcal{E}) = \left(\prod_{k=1}^K E_k, \prod_{k=1}^K \mathcal{E}_k \right)$, where $\prod_{k=1}^K \mathcal{E}_k$ denotes the corresponding product σ -algebra. To accomplish this task, suppose that we are given N samples $(\mathbf{X}^n)_{n=1}^N = (X_1^n, \dots, X_K^n)_{n=1}^N$ independently drawn from μ . Instead of approximating $\mu(\varphi)$ with the standard Monte Carlo estimator,

$$\mu^N(\varphi) := \frac{1}{N} \sum_{n=1}^N \varphi(\mathbf{X}^n),$$

consider using the *product-form estimator*,

$$\mu_{\times}^N(\varphi) := \frac{1}{NK} \sum_{\mathbf{n} \in [N]^K} \varphi(\mathbf{X}^{\mathbf{n}}),$$

obtained by averaging over all of the N^K element-wise permutations of samples $(\mathbf{X}^n)_{n=1}^N$, $\mathbf{X}^{\mathbf{n}} := (X_1^{n_1}, \dots, X_K^{n_K}) \quad \forall \mathbf{n} = (n_1, \dots, n_K) \in [N]^K$, where $[N] := \{1, \dots, N\}$.

It is a strongly consistent, unbiased, and asymptotically normal estimator for $\mu(\varphi)$ ([43, Theorem 1]) that achieves a smaller variance than $\mu^N(\varphi)$ ([43, Corollary 1]),

$$\text{Var}(\mu_{\times}^N(\varphi)) \leq \text{Var}(\mu^N(\varphi)) \quad \forall N > 0,$$

for all square μ -integrable test functions φ . (For that matter, $\mu_{\times}^N(\varphi)$ yields the best unbiased estimates of $\mu(\varphi)$ that can be achieved using only the knowledge that μ is product-form and N i.i.d. samples drawn from μ , cf. [43, Theorem 2].) In a nutshell, the standard estimator uses the (non-product-form) empirical distribution of the unpermuted samples,

$$\mu^N := \frac{1}{N} \sum_{n=1}^N \delta_{\mathbf{X}^n}$$

as an approximation to μ , while the product-form estimator uses the (product-form) empirical distribution obtained by taking the product of the marginal empirical distributions,

$$\mu_{\times}^N := \left(\frac{1}{N} \sum_{n_1=1}^N \delta_{X_1^{n_1}} \right) \times \dots \times \left(\frac{1}{N} \sum_{n_K=1}^N \delta_{X_K^{n_K}} \right) = \frac{1}{NK} \sum_{\mathbf{n} \in [N]^K} \delta_{\mathbf{X}^{\mathbf{n}}};$$

and the latter provides a better approximation to μ for the reasons outlined in the introduction (see [43, Section 2.2] for further details).

The above applies similarly to the case where μ is an importance sampling (IS) proposal rather than a target. In particular, if our target γ is a change of measure from μ with Radon-Nikodym derivative $w := d\gamma/d\mu$, then

$$\gamma_{\times}^N(\varphi) := \mu_{\times}^N(w\varphi) = \frac{1}{NK} \sum_{\mathbf{n} \in [N]^K} w(\mathbf{X}^{\mathbf{n}}) \varphi(\mathbf{X}^{\mathbf{n}}),$$

is a consistent, unbiased, and asymptotically normal estimator for $\gamma(\varphi) = \mu(w\varphi)$ whose variance is bounded above by that of the standard IS estimator $\gamma^N(\varphi) := \mu^N(w\varphi)$; cf. [43, Section 3.1]. Because this is the case for all sufficiently integrable φ , we regard

$$(1) \quad \gamma_{\times}^N(d\mathbf{x}) := w(\mathbf{x}) \mu_{\times}^N(d\mathbf{x}) = \frac{1}{NK} \sum_{\mathbf{n} \in [N]^K} w(\mathbf{X}^{\mathbf{n}}) \delta_{\mathbf{X}^{\mathbf{n}}}(d\mathbf{x})$$

as an approximation to γ . Somewhat abusively, we also refer to this type of approximation as a ‘product-form estimator’ (e.g. γ_{\times}^N is a product-form estimator for γ).

Algorithm 1 Sequential importance resampling

- 1: *Propose*: for $n \leq N$, draw $\mathbf{X}_0^{n,N}$ independently from K_0 .
 - 2: *Correct*: compute $\rho_0^N := N^{-1} \sum_{n=1}^N w_0(\mathbf{X}_0^n) \delta_{\mathbf{X}_0^n}$, where $w_0 := d\rho_0/dK_0$, and $\mu_0^N := \rho_0^N / \rho_0^N(\mathbf{E}_0)$.
 - 3: **for** $t = 1, \dots, T$ **do**
 - 4: *Resample*: for $n \leq N$, draw $\mathbf{X}_{t-}^{n,N}$ independently from μ_{t-1}^N .
 - 5: *Mutate*: for $n \leq N$, draw $X_t^{n,N}$ independently from $K_t(\mathbf{X}_{t-}^{n,N}, dx_t)$ and set $\mathbf{X}_t^{n,N} := (X_t^{n,N}, \mathbf{X}_{t-}^{n,N})$.
 - 6: *Correct*: compute ρ_t^N in (2) and $\mu_t^N := \rho_t^N / \rho_t^N(\mathbf{E}_t)$.
 - 7: **end for**
-

2.2. *Sequential Monte Carlo.* Sequential Monte Carlo (SMC) algorithms constitute a class of Monte Carlo methods that, by progressively evolving clouds of weighted particles, approximate sequences ρ_0, \dots, ρ_T of intractable unnormalized measures defined on measurable spaces $(\mathbf{E}_0, \mathcal{E}_0), \dots, (\mathbf{E}_T, \mathcal{E}_T)$ of increasing dimension (in particular, $\mathbf{E}_t = E_0 \times \dots \times E_t$ for some low-dimensional E_0, \dots, E_T and $\mathcal{E}_t = \mathcal{E}_0 \times \dots \times \mathcal{E}_t$ is the appropriate product σ -algebra). They also yield weighted particle approximations of the corresponding normalized distributions μ_0, \dots, μ_T and normalization constants Z_0, \dots, Z_T .

The basic SMC approach, often referred to as sequential importance resampling [27, p. 15], is detailed in Algorithm 1 (we use slightly unconventional notation here to keep things consistent with Section 3). It is a simple iterative procedure: given a weighted-particle approximation ρ_{t-1}^N to ρ_{t-1} , normalize ρ_{t-1}^N to obtain an approximation μ_{t-1}^N of μ_{t-1} , *resample* μ_{t-1}^N to obtain unweighted particles $\mathbf{X}_{t-}^{1,N}, \dots, \mathbf{X}_{t-}^{N,N}$ (whose empirical measure also approximates μ_{t-1}), and extend the path of each resampled particle $\mathbf{X}_{t-}^{n,N}$ by drawing $X_t^{n,N}$ from a proposal kernel $K_t(\mathbf{X}_{t-}^{n,N}, dx_t)$ and setting $\mathbf{X}_t^{n,N} := (X_t^{n,N}, \mathbf{X}_{t-}^{n,N})$. Then, re-weight the particles using the Radon-Nykodim derivative $w_t := d\rho_t/d(\rho_{t-1} \otimes K_t)$ of $\rho_t(dx_t)$ w.r.t. $(\rho_{t-1} \otimes K_t)(dx_t) := \rho_{t-1}(d\mathbf{x}_{t-1})K_t(\mathbf{x}_{t-1}, dx_t)$ to obtain an approximation of ρ_t :

$$(2) \quad \rho_t^N := \frac{Z_{t-1}^N}{N} \sum_{n=1}^N w_t(\mathbf{X}_t^{n,N}) \delta_{\mathbf{X}_t^{n,N}},$$

where $Z_{t-1}^N := \rho_{t-1}^N(\mathbf{E}_{t-1})$ denotes the estimate of the normalizing constant $Z_{t-1} = \rho_{t-1}(\mathbf{E}_{t-1})$ afforded by ρ_{t-1}^N . To initialize the algorithm, we approximate ρ_0 using importance sampling with a proposal distribution K_0 .

Resampling is the process of stochastically replicating some particles within a weighted collection in such a way that the expected number of offspring of each particle is proportional to its weight. By doing so, we obtain an unweighted sample suitable for approximating the same target distribution as the original weighted sample. The simplest way to achieve this, known as multinomial resampling because the vector of offspring numbers follows a multinomial distribution, is to independently draw N times from the weighed empirical measure of the original sample. To not complicate this paper's exposition and proofs, and to keep the focus on the differences between standard SMC and DaC-SMC, we restrict ourselves throughout to the case of multinomial resampling. However, using more sophisticated resampling schemes within the DaC-SMC setting is sensible, incurs no further complications than for SMC, and can in principle be understood using similar theoretical techniques to those employed in the conventional setting (see, in particular, [31]).

SMC algorithms were originally introduced in the filtering literature (see [27, 29] for reviews in this context) and have been successfully employed to sample from general distributions [9, 23]. The filtering problem has an intrinsic sequential structure and the sequence ρ_0, \dots, ρ_T arises naturally. For general distributions μ lacking this sort of structure, SMC

Algorithm 2 Auxiliary sequential importance resampling

-
- 1: *Propose*: for $n \leq N$, draw $\mathbf{X}_0^{n,N}$ independently from K_0 .
 - 2: *Compute*: $\gamma_0^N := N^{-1} \sum_{n=1}^N \delta_{\mathbf{X}_0^{n,N}}$.
 - 3: **for** $t = 1, \dots, T$ **do**
 - 4: *Correct*: compute $\gamma_{t-}^N(d\mathbf{x}_{t-1}) := w_{t-}(\mathbf{x}_{t-1})\gamma_{t-1}^N(d\mathbf{x}_{t-1})$ and $\pi_{t-}^N := \gamma_{t-}^N/\gamma_{t-}^N(\mathbf{E}_{t-1})$.
 - 5: *Resample*: for $n \leq N$, draw $\mathbf{X}_{t-}^{n,N}$ independently from π_{t-}^N .
 - 6: *Mutate*: for $n \leq N$, draw $X_t^{n,N}$ independently from $K_t(\mathbf{X}_{t-}^{n,N}, d\mathbf{x}_t)$ and set $\mathbf{X}_t^{n,N} := (X_t^{n,N}, \mathbf{X}_{t-}^{n,N})$.
 - 7: *Compute*: $\gamma_t^N := N^{-1} \mathcal{Z}_t^N \sum_{n=1}^N \delta_{\mathbf{X}_t^{n,N}}$ where $\mathcal{Z}_t^N := \gamma_{t-}^N(\mathbf{E}_{t-1})$.
 - 8: **end for**
- Note:** At each step t , one can use (4) to compute estimates of ρ_t , Z_t , and μ_t .
-

can instead be applied by building an ‘artificial’ sequence of distributions such that μ_T (or one of its marginals) coincides with μ . This can be achieved by incrementally morphing an easy-to-sample-from distribution into μ , often by tempering [54] or sequentially introducing observations to transition from a prior to a posterior [9]. At first glance, Algorithm 1 may not appear applicable in such settings because the artificial distributions are all defined on the same space. However, it is possible to construct auxiliary sequences of distributions on spaces of increasing dimension such that the t^{th} distribution of interest coincides with the t^{th} ‘coordinate marginal’ of the corresponding auxiliary distribution, hence allowing *exactly* this algorithm to be applied to these problems [23]. The same kind of constructions can be used in the DaC-SMC framework (see [47, Section 4.2] for a brief discussion) and we focus on analysing the fundamental algorithm (which, combined with appropriately extended target distributions, applies directly to these seemingly more complex settings).

We mention one variant of the SMC algorithm which is of particular relevance to what follows: the auxiliary SMC (ASMC) scheme originally proposed for filtering [59]. Although it was initially presented as an alternative proposal method involving an auxiliary variable construction, this scheme can be understood as the standard algorithm with an additional re-weighting prior to the resampling step [7]—the motivation behind the additional re-weighting being to incorporate as much information about the next distribution in the sequence as possible *prior* to resampling. Simply put, it is an SMC algorithm that targets a slightly different sequence of distributions and employs an extra importance sampling correction to obtain estimates with respect to the distributions of interest [36].

In particular, we introduce *auxiliary measures* $\gamma_{1-}, \dots, \gamma_{T-}$ on $\mathbf{E}_0, \dots, \mathbf{E}_{T-1}$, and their normalized counterparts $\pi_{1-}, \dots, \pi_{T-}$, and run Algorithm 1 with these measures as the targets. That is, with the weights used for resampling at time t set to $w_{t-} := d\gamma_{t-}/d\gamma_{t-1}$, where $\gamma_0 := K_0$ and $\gamma_t := \gamma_{t-} \otimes K_t$ for $t = 1, \dots, T$. This gives rise to the approximation

$$(3) \quad \gamma_{t-}^N := \frac{\mathcal{Z}_{t-1}^N}{N} \sum_{n=1}^N w_{t-}(\mathbf{X}_{t-1}^{n,N}) \delta_{\mathbf{X}_{t-1}^{n,N}}$$

of the auxiliary measure γ_{t-} , where $\mathcal{Z}_0^N := 1$ and $\mathcal{Z}_t^N := \gamma_{t-}^N(\mathbf{E}_{t-1})$ for $t = 1, \dots, T$. In the ASMC literature, it is common to find the correction step at the beginning of iteration t instead of the end of iteration $t-1$ (cf. Algorithm 2). Such cyclical permutations of the resample, mutate, and correct steps have no impact beyond labelling; we adopt this ordering because it eases the presentation of the DaC-SMC algorithm in Section 3.

The *extended auxiliary measure* γ_t is then approximated by

$$\gamma_t^N := \frac{\mathcal{Z}_t^N}{N} \sum_{n=1}^N \delta_{\mathbf{X}_t^{n,N}},$$

where, in the case of multinomial resampling, $\mathbf{X}_t^{1,N}, \dots, \mathbf{X}_t^{N,N}$ denote samples drawn independently from the normalization of $\gamma_{t-}^N \otimes K_t$. (We note in passing that γ_{t-1} and γ_{t-} can be interpreted as the ‘predictive’ and ‘updated’ time marginals of a Feynman-Kac flow in the sense of [21].) The above is then corrected using importance sampling and the weight function $w_t := d\rho_t/d\gamma_t$ to obtain approximations of the objects of interest ρ_t , Z_t , and μ_t :

$$(4) \quad \rho_t^N(d\mathbf{x}_t) := w_t(\mathbf{x}_t)\gamma_t^N(d\mathbf{x}_t), \quad Z_t^N := \rho_t^N(\mathbf{E}_t), \quad \mu_t^N(d\mathbf{x}_t) := \frac{1}{Z_t^N}\rho_t^N(d\mathbf{x}_t),$$

with γ_0^N denoting the empirical distribution of the samples drawn from $\gamma_0 = K_0$. This construction gives rise to the auxiliary SMC algorithm described in Algorithm 2.

3. The divide-and-conquer sequential Monte Carlo algorithm. In this section, we define our problem setting (Section 3.1), give the DaC-SMC algorithm (Section 3.2), consider its implementation (Section 3.3), and discuss various variants thereof (Section 3.4).

3.1. Problem definition. We tackle a generalization of the problem in Section 2.2: obtaining weighted-particle approximations to a collection $(\rho_u)_{u \in \mathbb{T}}$ of intractable unnormalized (but finite) measures indexed by the nodes of a finite rooted tree \mathbb{T} . In doing so, we also obtain tractable approximations to the corresponding collections of normalizing constants $(Z_u)_{u \in \mathbb{T}}$ and normalized measures $(\mu_u)_{u \in \mathbb{T}}$. Typically, our main interest will lie in approximating the target ρ_τ indexed by the tree’s root τ and all other targets have been introduced to facilitate this process. Hence, we refer to ρ_τ and μ_τ as the ‘final targets’ and to all other ρ_u and μ_u as the ‘intermediate targets’.

Although this may seem a somewhat specialized setting at first glance, a great many problems (including, for example, essentially any in which one wishes to approximate a probability distribution over \mathbb{R}^d) may be viewed within this framework. In some cases, the tree structure is intrinsic to the model at hand and there are natural candidates for the intermediate targets (e.g. [35, 57, 30] or Example 1 below). In others, the tree decomposition and intermediate targets are entirely artificial constructs introduced for the purpose of providing a computationally tractable approach to approximating the final targets (e.g. see [47, Section 3.4] or [24]) and DaC-SMC serves as a generalization of the approaches in [23, 9].

In all of these cases, the intermediate targets transition incrementally from a collection of measures on low-dimensional spaces indexed by \mathbb{T} ’s leaves to the final target. Hence, the measures are defined on spaces whose dimension grows as we progress up the tree. In particular, the underlying measurable space $(\mathbf{E}_u, \mathcal{E}_u)$ for ρ_u is defined by taking the corresponding partial product of a collection $(E_v, \mathcal{E}_v)_{v \in \mathbb{T}}$ of low-dimensional spaces indexed by \mathbb{T} ’s nodes:

$$\mathbf{E}_u = \prod_{v \in \mathbb{T}_u} E_v, \quad \mathcal{E}_u = \prod_{v \in \mathbb{T}_u} \mathcal{E}_v, \quad \forall u \in \mathbb{T},$$

where \mathbb{T}_u denotes the sub-tree of \mathbb{T} rooted at u (obtained by removing all nodes from \mathbb{T} except for u and its descendants).

Similarly to Algorithm 2, we introduce a collection $(\gamma_{u-})_{u \in \mathbb{T}^\circ}$ of *auxiliary measures* indexed by the nodes of our tree that are not leaves. In particular, for any non-leaf u with set of children \mathcal{C}_u , γ_{u-} is a finite measure on the corresponding intermediate product space:

$$E_{\mathcal{C}_u} := \prod_{v \in \mathbb{T}_u^\circ} E_v, \quad \mathcal{E}_{\mathcal{C}_u} := \prod_{v \in \mathbb{T}_u^\circ} \mathcal{E}_v,$$

where $\mathbb{T}_u^\circ := \mathbb{T}_u \setminus \{u\}$. These auxiliary distributions fulfil essentially the same role as the auxiliary distributions in the ASMC algorithm: introducing as much information as possible

about the measure ρ_u associated with the current node u *before* resampling. As we will see in the next section, this mechanism is particularly important in the divide and conquer context because it allows us to correlate the components of the resampled particles associated with u 's children.

REMARK 1 (A connection with SMC). If the tree is a line (i.e. each node has a single child, except for the first which has none), then our setting here reduces to that of standard ASMC (Algorithm 2). If we choose auxiliary distributions $\gamma_{t_-} := \rho_{t-1}$, then it further reduces to a standard SMC (Algorithm 1).

For the reasons discussed in Section 4.2, we would like to choose the intermediate targets and auxiliary measures to be as close as possible to the corresponding marginals of the final target ρ_τ . Of course, in general, this is more easily said than done and, in practice, ρ_u and γ_{u_-} will often be obtained by ignoring [47, Section 3.4], rather than ‘integrating-out’, the terms in ρ_τ featuring variables associated with u 's ancestors (and u itself in the case of γ_{u_-}). To illustrate some of these concepts, consider the following simplified version of the example in [47, Section 5.2]:

EXAMPLE 1 (Simplified mathematics test dataset). Suppose we have the following data: the number M_{sy} of students who took a certain exam in school s and year y and the number m_{sy} which passed the exam, for all years y that the exam was taken in school s and all schools s in a given city (we denote the set of the former by \mathcal{Y}_s and that of the latter by \mathcal{S}). A standard hierarchical model [30] for inference on such a dataset is as follows: m_{sy} is binomially distributed with success probability $\alpha(\theta_{sy})$ (i.e. m_{sy} has law $\mathcal{B}(\cdot; M_{sy}, \alpha(\theta_{sy}))$), where θ_{sy} denotes a latent parameter and α denotes the standard logistic function. The parameters are related as follows:

$$\theta_{sy} = \theta_s + \varepsilon_{sy} \quad \forall y \in \mathcal{Y}_s, \quad s \in \mathcal{S}, \quad \theta_s = \theta_\tau + \varepsilon_s \quad \forall s \in \mathcal{S},$$

where $(\theta_s)_{s \in \mathcal{S}}$ and θ_τ denote further latent parameters and $((\varepsilon_{sy})_{y \in \mathcal{Y}_s}, \varepsilon_s)_{s \in \mathcal{S}}$ a collection of independent noises with $\varepsilon_{sy} \sim \mathcal{N}(0, \sigma_s^2)$ and $\varepsilon_s \sim \mathcal{N}(0, \sigma_\tau^2)$ for some unknown variances $(\sigma_s^2)_{s \in \mathcal{S}}$ and σ_τ^2 . For each parameter variable, we choose as a prior the pullback $g(d\theta) = \alpha(\theta)(1 - \alpha(\theta))d\theta$ of the uniform distribution on $(0, 1)$ via α ; while for each variance variable, we choose a unit-mean exponential distribution $f(d\sigma^2; 1)$. The full (unnormalized) posterior is then given by

$$(5) \quad f(d\sigma_\tau^2; 1)g(d\theta_\tau) \prod_{s \in \mathcal{S}} \left[\mathcal{N}(\theta_s - \theta_\tau; 0, \sigma_\tau^2) f(d\sigma_s^2; 1) g(d\theta_s) \right. \\ \left. \times \prod_{y \in \mathcal{Y}_s} \mathcal{N}(\theta_{sy} - \theta_s; 0, \sigma_s^2) \mathcal{B}(m_{sy}; M_{sy}, \alpha(\theta_{sy})) g(d\theta_{sy}) \right].$$

This distribution's dependence structure is encoded by a tree \mathbb{T} whose paths from root to leaves have the following form: root, school, year. In other words, \mathbb{T} 's root node τ has a child for each school ($\mathcal{C}_\tau = \mathcal{S}$), each of these has a child for each year the exam took place in that school ($\mathcal{C}_s = \mathcal{Y}_s$ for all s in \mathcal{S}), and the latter are all leaves (note that we treat elements in different \mathcal{Y}_s s as distinct nodes, even if they represent the same year). With each leaf node, we associate the corresponding latent variable (so that $E_u := \mathbb{R}$ if u is a leaf), and, with each non-leaf node, we associate the corresponding parameter variable and unknown variance (so that $E_u := \mathbb{R} \times [0, \infty)$ if u is not a leaf).

We set the final target ρ_τ to be our measure of interest (5). Analytical expressions for its marginals are unavailable, so we instead pick the intermediate targets as follows: we obtain ρ_u by removing from (5) all terms featuring variables not associated with u or its descendants:

$$(6) \quad \rho_s(d\theta_s, d\sigma_s^2, (d\theta_{sy})_{y \in \mathcal{Y}_s}) := f(d\sigma_s^2; 1)g(d\theta_s) \\ \times \prod_{y \in \mathcal{Y}_s} \mathcal{N}(\theta_{sy} - \theta_s; 0, \sigma_s^2) \mathcal{B}(m_{sy}; M_{sy}, \alpha(\theta_{sy}))g(d\theta_{sy}), \\ \rho_{sy}(d\theta_{sy}) := \mathcal{B}(m_{sy}; M_{sy}, \alpha(\theta_{sy}))g(d\theta_{sy}) \quad \forall y \in \mathcal{Y}_s,$$

for all s in \mathcal{S} . It is also not possible to analytically integrate over $(\theta_\tau, \sigma_\tau^2)$ in (5) or (θ_s, σ_s^2) in (6), so we instead pick the auxiliary measures by further doing away with the terms in (5,6) involving the variables associated with the indexing node:

$$\gamma_{\tau_-} := \prod_{s \in \mathcal{S}} \rho_s, \quad \gamma_{s_-} := \prod_{y \in \mathcal{Y}_s} \rho_{sy} \quad \forall s \in \mathcal{S}.$$

We finish this section with some notation that we use throughout the rest of the paper:

- (Trees) By a ‘tree’ \mathbb{T} we mean a rooted tree: a connected directed acyclic graph in which no node has more than one parent and possessing a single root node, τ , that has no parent. We denote the number of children a node u has as c_u , refer to them as $u1, u2, \dots, uc_u$, and denote the set $\{u1, \dots, uc_u\}$ of them by \mathcal{C}_u . For any subset A of the tree \mathbb{T} , we use $|A|$ to denote its cardinality, $A^\partial := \{v \in A : c_v = 0\}$ to denote the set of all leaves in A , and $A^\partial := A \setminus A^\partial$ that of all other nodes in A . In particular, \mathbb{T}^∂ denotes the set of \mathbb{T} ’s leaves and \mathbb{T}^{∂^c} that of all non-leaf nodes.
- (Products) We assign the subscript ‘ A ’ to a symbol, where A is any subset of \mathbb{T} , if the object represented by the symbol depends on the nodes in A . If the symbol is bold, the object also depends on the descendants of the nodes in A . With few exceptions, this dependence takes the form of a product. For example,

$$E_A := \prod_{u \in A} E_u, \quad \mathbf{E}_A := E_{\cup_{u \in A} \mathbb{T}_u}, \quad \mathcal{E}_A := \prod_{u \in A} \mathcal{E}_u, \quad \mathcal{E}_A := \mathcal{E}_{\cup_{u \in A} \mathbb{T}_u}, \\ \rho_A := \prod_{u \in A} \rho_u, \quad \mu_A := \prod_{u \in A} \mu_u, \quad Z_A := \prod_{u \in A} Z_u.$$

for any non-empty subset A of \mathbb{T} . The exceptions are x_A and \mathbf{x}_A which denote elements in E_A and \mathbf{E}_A rather than any sort of product over A .

3.2. The algorithm. To introduce the divide-and-conquer SMC algorithm, we require one last ingredient: the generalization of the proposal kernels used in SMC/ASMC (cf. Section 2.2) to extend the path of the resampled particles. In our case, we have, for each non-leaf node u , one Markov kernel $K_u : \mathbf{E}_{\mathcal{C}_u} \times \mathcal{E}_u \rightarrow [0, 1]$ mapping from the product $\mathbf{E}_{\mathcal{C}_u} = \prod_{v \in \mathcal{C}_u} \mathbf{E}_v$ of the intermediate spaces \mathbf{E}_v indexed by u ’s children to the low-dimensional space E_u indexed by u itself. For the reasons discussed in Section 4.2, the *proposal kernel* K_u should be chosen as close as possible to the conditional distribution over E_u under μ_u given values for its remaining coordinates.

Algorithm 3 describes DaC-SMC¹ recursively: the first thing it does when called on any given non-leaf node is to call itself on each of that node’s children. Hence, the algorithm

¹Note that, in the original paper [47], Algorithm 3 was referred to as ‘DaC-SMC with mixture resampling’ while the term ‘DaC-SMC’ was reserved for cases where the auxiliary measures are product-form and the algorithm’s running time can be sped up through careful implementation (cf. Section 3.3). For all, at least theoretical, purposes the latter is a special case of the former and we simply refer to both as the ‘DaC-SMC Algorithm’.

Algorithm 3 $\text{dac_smc}(u)$ for u in \mathbb{T} .

- 1: **if** u is a leaf (i.e. $u \in \mathbb{T}^\partial$) **then**
- 2: *Propose*: for $n \leq N$, draw $\mathbf{X}_u^{n,N}$ independently from K_u .
- 3: *Return*: $\gamma_u^N := N^{-1} \sum_{n=1}^N \delta_{\mathbf{X}_u^{n,N}}$.
- 4: **else**
- 5: **for** v in \mathcal{C}_u **do**
- 6: *Recurse*: set $\gamma_v^N := \text{dac_smc}(v)$.
- 7: **end for**
- 8: *Correct*: compute $\gamma_{u_-}^N$ in (8) and $\pi_{u_-}^N := \gamma_{u_-}^N / \gamma_{u_-}^N(\mathbf{E}_{\mathcal{C}_u})$.
- 9: *Resample*: for $n \leq N$, draw $\mathbf{X}_{u_-}^{n,N}$ independently from $\pi_{u_-}^N$.
- 10: *Mutate*: for $n \leq N$, draw $X_u^{n,N}$ independently from $K_u(\mathbf{X}_{u_-}^{n,N}, dx_u)$ and set $\mathbf{X}_u^{n,N} := (X_u^{n,N}, \mathbf{X}_{u_-}^{n,N})$.
- 11: *Return*: $\gamma_u^N := N^{-1} \mathcal{Z}_u^N \sum_{n=1}^N \delta_{\mathbf{X}_u^{n,N}}$ where $\mathcal{Z}_u^N := \gamma_{u_-}^N(\mathbf{E}_{\mathcal{C}_u})$.
- 12: **end if**

Note: One can use (11) to compute estimates of ρ_u , Z_u , and μ_u .

proper begins ‘at the leaves’ by drawing, for each leaf u , N samples $\mathbf{X}_u^{1,N}, \dots, \mathbf{X}_u^{N,N}$ from a proposal distribution K_u on (E_u, \mathcal{E}_u) and setting $\gamma_u^N := N^{-1} \sum_{n=1}^N \delta_{\mathbf{X}_u^{n,N}}$ to be the corresponding unweighted sample approximation to K_u . It then works its way from leaves to root, at each node combining and augmenting the approximations indexed by its children.

More precisely, at any non-leaf node u , and for each child v of u , the algorithm first retrieves an unweighted particle approximation $\gamma_v^N = \mathcal{Z}_v^N N^{-1} \sum_{n=1}^N \delta_{\mathbf{X}_v^{n,N}}$ (of mass \mathcal{Z}_v^N) to the *extended auxiliary measure* γ_v obtained by taking the outer product of the auxiliary measure γ_{v_-} and the proposal kernel K_v :

$$(7) \quad \gamma_v(dx_v) := (\gamma_{v_-} \otimes K_v)(dx_v) = \gamma_{v_-}(d\mathbf{x}_{\mathcal{C}_v}) K_v(\mathbf{x}_{\mathcal{C}_v}, dx_v) \quad \forall v \in \mathbb{T}^\partial.$$

(If v is a leaf then the algorithm just retrieves the sample approximation to $\gamma_v := K_v$ mentioned in the preceding paragraph.) It then takes the product $\gamma_{\mathcal{C}_u}^N := \prod_{v \in \mathcal{C}_u} \gamma_v^N$ of these and re-weights it using $w_{u_-} := d\gamma_{u_-} / d\gamma_{\mathcal{C}_u}$ to obtain a product-form estimator (cf. (1)) for the auxiliary distribution γ_{u_-} associated with node u :

$$(8) \quad \gamma_{u_-}^N(d\mathbf{x}_{\mathcal{C}_u}) := w_{u_-}(\mathbf{x}_{\mathcal{C}_u}) \gamma_{\mathcal{C}_u}^N(d\mathbf{x}_{\mathcal{C}_u}) = \frac{\mathcal{Z}_{\mathcal{C}_u}^N}{N^{c_u}} \sum_{\mathbf{n} \in [N]^{c_u}} w_{u_-}(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}) \delta_{\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}}(d\mathbf{x}_{\mathcal{C}_u})$$

where $\mathcal{Z}_{\mathcal{C}_u}^N := \prod_{v \in \mathcal{C}_u} \mathcal{Z}_v^N$ and, for each $\mathbf{n} = (n_1, \dots, n_{c_u})$ in $[N]^{c_u}$, $\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}$ denotes the ‘per-mutated particle’ $(\mathbf{X}_{u_1}^{n_1}, \dots, \mathbf{X}_{u_{c_u}}^{n_{c_u}})$. Just as in the standard ASMC case (Algorithm 2), the algorithm proceeds by normalizing $\gamma_{u_-}^N$ to obtain an approximation $\pi_{u_-}^N$ of the *normalized auxiliary measure* π_{u_-} ,

$$(9) \quad \pi_{u_-}^N := \frac{\gamma_{u_-}^N}{\mathcal{Z}_{\mathcal{C}_u}^N} \approx \frac{\gamma_{u_-}}{\mathcal{Z}_u} =: \pi_{u_-} \quad \text{where } \mathcal{Z}_u^N := \gamma_{u_-}^N(\mathbf{E}_{\mathcal{C}_u}), \quad \mathcal{Z}_u := \gamma_{u_-}(\mathbf{E}_{\mathcal{C}_u}),$$

and it draws N particles $\mathbf{X}_{u_-}^{1,N}, \dots, \mathbf{X}_{u_-}^{N,N}$ from $\gamma_{u_-}^N$ using (multinomial) resampling. The empirical distribution $N^{-1} \sum_{n=1}^N \delta_{\mathbf{X}_{u_-}^{n,N}}$ of these particles approximates π_{u_-} . Hence, by extending the path of each resampled particle from $\mathbf{E}_{\mathcal{C}_u}$ to \mathbf{E}_u using the kernel K_u ,

$$\mathbf{X}_u^{n,N} := (X_u^{n,N}, \mathbf{X}_{u_-}^{n,N}) \quad \text{where } X_u^{n,N} \sim K_u(\mathbf{X}_{u_-}^{n,N}, dx_u),$$

and defining

$$(10) \quad \gamma_u^N(d\mathbf{x}_u) := \frac{\mathcal{Z}_u^N}{N} \sum_{n=1}^N \delta_{\mathbf{X}_u^{n,N}} \quad \text{and} \quad \pi_u^N(d\mathbf{x}_u) := \frac{1}{N} \sum_{n=1}^N \delta_{\mathbf{X}_u^{n,N}},$$

Algorithm 3 returns a finite-sample approximation of the extended auxiliary measure $\gamma_u := \gamma_{u_-} \otimes K_u$ indexed by u itself.

Also similarly to the standard ASMC case (Section 2.2), we obtain approximations to the targets ρ_u and μ_u by running `dac_smc(u)` in Algorithm 3 and applying a simple importance sampling correction:

$$(11) \quad \rho_u^N(d\mathbf{x}_u) = w_u(\mathbf{x}_u)\gamma_u^N(d\mathbf{x}_u), \quad Z_u^N = \rho_u^N(\mathbf{E}_u), \quad \mu_u^N(d\mathbf{x}_u) = \frac{\rho_u^N(d\mathbf{x}_u)}{Z_u^N},$$

where $w_u := d\rho_u/d\gamma_u$ denotes the appropriate weight function. As we will see in Section 4, these approximations have appealing theoretical properties. The computation time of Algorithm 3 can be lowered by parallelizing the for loop in lines 5–7. Of course, for the approach to work, we must choose the auxiliary measures $(\gamma_{u_-})_{u \in \mathbb{T}^\vartheta}$ and proposal kernels $(K_u)_{u \in \mathbb{T}}$ such that the relevant Radon-Nikodym derivatives exist. To avoid the technical machinery necessary to allow for the possibility of all particles simultaneously being assigned zero weight (cf. [21, Section 7.4]), we further assume that these derivatives are positive everywhere:

ASSUMPTION 1. For all u in \mathbb{T} and v in \mathbb{T}^ϑ , ρ_u is absolutely continuous w.r.t. γ_u , γ_{v_-} is absolutely continuous w.r.t. γ_{C_v} , and the Radon-Nikodym derivatives $w_u := d\rho_u/d\gamma_u$ and $w_{v_-} := d\gamma_{v_-}/d\gamma_{C_v}$ are positive everywhere.

To finish this section, we illustrate how one may choose the proposal kernels and obtain the weight functions by revisiting the example of the previous section:

EXAMPLE 2 (Simplified mathematics test dataset, kernels and weights). To apply DaC-SMC to Example 1, we need to pick proposal kernels. In the case of the leaves, the choice is clear: for all s, y, n , we generate $\theta_{sy}^{n,N}$ by drawing a sample from a Beta distribution with parameters $m_{sy} + 1$ and $M_{sy} - m_{sy} + 1$ and mapping the sample through the logit function; so that

$$K_{sy}(d\theta_{sy}) = (M_{sy} + 1) \binom{M_{sy}}{m_{sy}} \mathcal{B}(m_{sy}; M_{sy}, \alpha(\theta_{sy})) g(d\theta_{sy})$$

and the weight functions are constant (for all s in \mathcal{S} , $w_{s_-} = \prod_{y \in \mathcal{Y}_s} M_{sy}^{-1}$ and, for all y in \mathcal{Y}_s , $w_{sy} = M_{sy}^{-1}$). In the case of the other nodes, we would ideally set the kernel to be the appropriate conditional distribution (cf. Section 4.2),

$$H_s((\theta_{sy})_{y \in \mathcal{Y}_s}, d\theta_s, d\sigma_s^2) \propto f(d\sigma_s^2; 1) g(d\theta_s) \prod_{y \in \mathcal{Y}_s} \mathcal{N}(\theta_{sy} - \theta_s; 0, \sigma_s^2) \quad \forall s \in \mathcal{S};$$

$$H_\tau((\theta_s, \sigma_s^2)_{s \in \mathcal{S}}, d\theta_\tau, d\sigma_\tau^2) \propto f(d\sigma_\tau^2; 1) g(d\theta_\tau) \prod_{s \in \mathcal{S}} \mathcal{N}(\theta_s - \theta_\tau; 0, \sigma_\tau^2).$$

These, however, are intractable. So, in an effort to approximate them, we note that

$$\prod_{s \in \mathcal{S}} \mathcal{N}(\theta_s - \theta_\tau; 0, \sigma_\tau^2) = \frac{e^{\frac{1}{2\sigma_\tau^2} [\sum_{s \in \mathcal{S}} \theta_s (\theta_s - \frac{1}{|\mathcal{S}|} \sum_{s' \in \mathcal{S}} \theta_{s'})]}}{\sqrt{|\mathcal{S}|} \sqrt{2\pi\sigma_\tau^2}^{|\mathcal{S}|-1}} \mathcal{N}\left(d\theta_\tau; \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \theta_s, \frac{\sigma_\tau^2}{|\mathcal{S}|}\right)$$

and similarly for $\prod_{y \in \mathcal{Y}_s} \mathcal{N}(\theta_{sy} - \theta_s; 0, \sigma_s^2)$ and all s in \mathcal{S} ; and we instead set

$$K_s((\theta_{sy})_{y \in \mathcal{Y}_s}, d\theta_s, d\sigma_s^2) := f(d\sigma_s^2; 1) \mathcal{N}\left(d\theta_s; \frac{1}{|\mathcal{Y}_s|} \sum_{y \in \mathcal{Y}_s} \theta_{sy}, \frac{\sigma_s^2}{|\mathcal{Y}_s|}\right),$$

$$K_{\tau}((\theta_s, \sigma_s^2)_{s \in \mathcal{S}}, d\theta_{\tau}, d\sigma_{\tau}^2) := f(d\sigma_{\tau}^2; 1) \mathcal{N} \left(d\theta_{\tau}; \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \theta_s, \frac{\sigma_{\tau}^2}{|\mathcal{S}|} \right);$$

in which case $w_{\tau_-} = \prod_{s \in \mathcal{S}} w_s$,

$$w_s(\theta_s, \sigma_s^2, (\theta_{sy})_{y \in \mathcal{Y}_s}) = \frac{g(\theta_s) e^{\frac{1}{2\sigma_s^2} [\sum_{y \in \mathcal{Y}_s} \theta_{sy} (\theta_{sy} - \frac{1}{|\mathcal{Y}_s|} \sum_{y' \in \mathcal{Y}_s} \theta_{sy'})]}}{\sqrt{|\mathcal{Y}_s|} \sqrt{2\pi\sigma_s^2}^{|\mathcal{Y}_s|-1}} \quad \forall s \in \mathcal{S},$$

$$w_{\tau}(\theta_{\tau}, \sigma_{\tau}^2, (\theta_s, \sigma_s^2, (\theta_{sy})_{y \in \mathcal{Y}_s})_{s \in \mathcal{S}}) = \frac{g(\theta_{\tau}) e^{\frac{1}{2\sigma_{\tau}^2} [\sum_{s \in \mathcal{S}} \theta_s (\theta_s - \frac{1}{|\mathcal{S}|} \sum_{s' \in \mathcal{S}} \theta_{s'})]}}{\sqrt{|\mathcal{S}|} \sqrt{2\pi\sigma_{\tau}^2}^{|\mathcal{S}|-1}}.$$

To sample (θ_s, σ_s^2) from $K_s((\theta_{sy})_{y \in \mathcal{Y}_s}, d\theta_s, d\sigma_s^2)$, first sample σ_s^2 from $f(d\sigma_s^2; 1)$ and then sample θ_s from $\mathcal{N}(d\theta_s; |\mathcal{Y}_s|^{-1} \sum_{y \in \mathcal{Y}_s} \theta_{sy}, |\mathcal{Y}_s|^{-1} \sigma_s^2)$. Similarly for K_{τ} .

3.3. Efficient computation. The main drawback of Algorithm 3 is the $\mathcal{O}(N^{c_u})$ computational cost of resampling from the product-form $\pi_{u_-}^N$ in (9) for a general auxiliary measure γ_{u_-} (to draw N samples from $\pi_{u_-}^N$ in $\mathcal{O}(N^{c_u})$ operations one can, for example, use the alias method [65]). This results in a total algorithmic cost of $\mathcal{O}(N^d)$, where d denotes the tree's degree (i.e. the largest number of children that any of the tree's nodes posses). In some cases, d is small enough that the cost is manageable. For instance, when the collection of target distributions $(\rho_u)_{u \in \mathbb{T}_u}$ is an artificial construct introduced for computational purposes and the interest lies in the final target, it is common [25, 24, 13] to choose the indexing tree to be binary, in which case $d = 2$ and the algorithm's cost is $\mathcal{O}(N^2)$.

In cases where each weight function is bounded above by a known constant, one can lower the cost to $\mathcal{O}(N)$ using rejection sampling as proposed in [13, Section 4.2]. Otherwise, one can avoid the $\mathcal{O}(N^{c_u})$ cost by choosing auxiliary measures that factorize. For instance, suppose that, as in Example 1, they factorize fully:

$$\gamma_{u_-} := \prod_{v \in \mathcal{C}_u} \bar{\gamma}_v \quad \forall u \in \mathbb{T}^{\partial},$$

where $(\bar{\gamma}_v)_{v \in \mathbb{T}}$ denotes a collection of measures on the respective spaces $(\mathbf{E}_v, \mathbf{E}_v)_{v \in \mathbb{T}}$. In this case, the weight function w_{u_-} decomposes into the product

$$w_{u_-} = \frac{d\gamma_{u_-}}{d\gamma_{\mathcal{C}_u}} = \prod_{v \in \mathcal{C}_u} \frac{d\bar{\gamma}_v}{d\gamma_v} =: \prod_{v \in \mathcal{C}_u} \bar{w}_v,$$

and the correction step (line 8 in Algorithm 3) breaks down into the following: compute

$$\gamma_{u_-}^N := \prod_{v \in \mathcal{C}_u} \bar{\gamma}_v^N, \quad \pi_{u_-}^N := \prod_{v \in \mathcal{C}_u} \bar{\pi}_v^N,$$

where, for all children v of u ,

$$(12) \quad \bar{\gamma}_v^N(d\mathbf{x}_v) := \bar{w}_v(\mathbf{x}_v) \gamma_v^N(d\mathbf{x}_v) = \frac{Z_v^N}{N} \sum_{n=1}^N \bar{w}_v(\mathbf{x}_v) \delta_{\mathbf{X}_v^{n,N}}(d\mathbf{x}_v).$$

Resampling from $\pi_{u_-}^N := \bar{\gamma}_v^N / \bar{\gamma}_v^N(\mathbf{E}_v)$ can then be achieved in $\mathcal{O}(N)$ operations by independently drawing N samples from $\bar{\pi}_{u_1}^N, \dots, \bar{\pi}_{u_{c_u}}^N$, e.g. by applying the alias method [65] to each approximation separately, and concatenating them. In other words, we can replace lines

Algorithm 4 $\mathcal{O}(N)$ -cost replacement of lines 5–9 in Algorithm 3 for fully-factorized γ_{u_-} .

- 1: **for** v in \mathcal{C}_v **do**
 - 2: *Recurse*: set $\gamma_v^N := \text{dac_smc}(v)$.
 - 3: *Correct*: compute $\bar{\gamma}_v^N$ in (12) and $\bar{\pi}_v^N := \bar{\gamma}_v^N / \gamma_v^N(\mathbf{E}_v)$.
 - 4: *Resample*: for $n \leq N$, $\bar{\mathbf{X}}_v^{n,N}$ independently from $\bar{\pi}_v^N$.
 - 5: **end for**
 - 6: *Concatenate*: for $n \leq N$, set $\mathbf{X}_{u_-}^{n,N} := (\bar{\mathbf{X}}_v^{n,N})_{v \in \mathcal{C}_u}$.
-

5–9 in Algorithm 3 with those in Algorithm 4 (and obtain the algorithm referred to as ‘DaC-SMC’ in [47], cf. Footnote 1). The running time can then be lowered further by parallelizing these operations across u ’s children.

These choices result in DaC-SMC employing extended auxiliary measures with partial product structure:

$$\gamma_u(d\mathbf{x}_u) = K_u(\mathbf{x}_{\mathcal{C}_u}, d\mathbf{x}_u) \prod_{v \in \mathcal{C}_u} \bar{\gamma}_v(d\mathbf{x}_v) \quad \forall u \in \mathbb{T}^\partial.$$

(I.e. if $(Y_u, \mathbf{Y}_{u_1}, \dots, \mathbf{Y}_{u_{c_u}}) \sim \pi_u$, then $\mathbf{Y}_{u_1}, \dots, \mathbf{Y}_{u_{c_u}}$ are independent.) If the targets ρ_u and μ_u do not share the above structure, then this $\mathcal{O}(N)$ approach will generally require large $\gamma_u^N \mapsto \rho_u^N$ corrections and, consequently, result in greater estimator variance (which may, or may not, be compensated by the computational gains). As suggested in [47, Section 4.2], one way to mitigate this issue is via tempering [23].

Another way to incorporate non-product structure into the auxiliary measures without sacrificing the $\mathcal{O}(N)$ cost is setting them to be sums of fully-factorized measures:

$$\gamma_{u_-} := \sum_{i=1}^I \prod_{v \in \mathcal{C}_u} \bar{\gamma}_v^i =: \sum_{i=1}^I \gamma_{u_-}^i \quad \forall u \in \mathbb{T}^\partial,$$

where $(\bar{\gamma}_v^1, \dots, \bar{\gamma}_v^I)_{v \in \mathbb{T}}$ denote collections of measures on the respective spaces $(\mathbf{E}_v, \mathcal{E}_v)_{v \in \mathbb{T}}$. In this case, the correction step becomes: return

$$\gamma_{u_-}^N := \sum_{i=1}^I \gamma_{u_-}^{i,N}, \quad \pi_{u_-}^N := \sum_{i=1}^I \frac{\gamma_{u_-}^{i,N}(\mathbf{E}_{\mathcal{C}_u})}{\gamma_{u_-}^N(\mathbf{E}_{\mathcal{C}_u})} \pi_{u_-}^{i,N} =: \sum_{i=1}^I \omega_{u_-}^{i,N} \pi_{u_-}^{i,N},$$

where $\gamma_{u_-}^{i,N} := \prod_{v \in \mathcal{C}_u} \bar{\gamma}_v^{i,N}$ and $\pi_{u_-}^{i,N} := \prod_{v \in \mathcal{C}_u} \bar{\pi}_v^{i,N}$, with $\bar{\gamma}_v^{i,N}, \bar{\pi}_v^{i,N}$ defined analogously to $\bar{\gamma}_v^N, \bar{\pi}_v^N$ in (12). We can then draw N samples from $\pi_{u_-}^N$ in $\mathcal{O}(N)$ operations by, for instance, drawing (m_1, \dots, m_I) from a multinomial distribution with weights $(\omega_{u_-}^{i,N})_{i=1}^I$ and, for each i , generating m_i samples from $\pi_{u_-}^{i,N}$ as described above for the $I = 1$ case.

Similar considerations apply if the auxiliary measures are chosen to be sums of partially factorized functions rather than fully factorized ones. The main difference is that resampling from $\pi_{u_-}^N$ has a ‘conditional aspect’ to it and costs $\mathcal{O}(N^a)$ where the exponent $1 < a < c_u$ depends on the amount of factorization, see Appendix A.1 for an example.

3.4. Variants. To simplify the algorithm’s presentation and analysis, we focus throughout on the particular version given in Algorithm 3. There are, however, many variants that one might wish to consider in different settings and whose analysis may be tackled using straightforward modifications of the arguments given in this paper.

First off, we have the low-cost variant referred to as ‘lightweight mixture resampling’ in [47, Section 4.1] (see ‘multiple matching’ in [46, Section 2.2] for similar ideas) where the

product-form estimator $\gamma_{u-}^N(dx_{C_u})$ in (8) is replaced with its ‘incomplete’ version,

$$(13) \quad \frac{\mathcal{Z}_{C_u}^N}{|\mathcal{M}|} \sum_{n \in \mathcal{M}} w_{u-}(\mathbf{X}_{C_u}^{n,N}) \delta_{\mathbf{X}_{C_u}^{n,N}}(dx_{C_u}),$$

with $\mathcal{M} \subseteq [N]^{c_u}$ indexing a user-chosen subset of permuted particles $\mathbf{X}_{C_u}^{n,N}$. In particular, by picking \mathcal{M} carefully, it might be possible to substantially reduce the cost without sacrificing too much of the variance reduction (cf. [40] and references therein for similar feats in the U-statistics literature). While one may opt for more sophisticated approaches (e.g. ones along the lines of those in [40]), we can offer two simple rules of thumb for choosing \mathcal{M} that are motivated by the considerations in [43, Section 4]: (a) set \mathcal{M} ’s size to be such that the cost of resampling from (13)’s normalization is comparable to that of generating the unpermuted particles, $(\mathbf{X}_{u_1}^{n,N}, \dots, \mathbf{X}_{u_{c_u}}^{n,N})_{n=1}^N$; and (b) minimize the number of over-lapping components in the elements of \mathcal{M} . In terms of theoretical analysis for this variant, one can combine the methods in Appendices C–G with the techniques used to study incomplete U-statistics [45, Chapter 4.3] and their generalizations.

We also need not resample at every step and may instead choose to do so adaptively (e.g. resample only when the effective sample size drops underneath a predetermined threshold [39]). To prevent the number of particles from blowing up, any product-form approximation π_{u-}^N to a normalized auxiliary distribution π_{u-} that is not resampled must be replaced with an N sample one, e.g.

$$(14) \quad \frac{\sum_{n=1}^N w_{u-}(\mathbf{X}_{u_1}^{n,N}, \dots, \mathbf{X}_{u_{c_u}}^{n,N}) \delta_{(\mathbf{X}_{u_1}^{n,N}, \dots, \mathbf{X}_{u_{c_u}}^{n,N})}}{\sum_{n=1}^N w_{u-}(\mathbf{X}_{u_1}^{n,N}, \dots, \mathbf{X}_{u_{c_u}}^{n,N})}$$

as in standard ASMC. We should point out here that, for DaC-SMC, and in contrast with SMC/ASMC, there are benefits to resampling beyond just mitigating weight degeneracy. In particular, by permuting and resampling the particles, we explore areas of the target space that would be missed were we to only use the N unpermuted particles [43, Section 3.1].

There is also flexibility in the resampling scheme employed: the theoretical analysis in Section 4 applies to Algorithm 3 in which multinomial resampling is explicitly encoded. But one could extend our analysis along the lines of [31] to cover a much broader class of resampling schemes, including substantially lower-variance ones. An interesting question of practical importance is how best to efficiently implement low-variance resampling schemes when the auxiliary measures decompose into sums of partially-factorized measures (e.g. consider the example in Appendix A.1).

One may also wish to implement minor variations to Algorithm 3 dictated by the particular target of interest. For instance, in some cases, it might be much easier to specify the auxiliary measure γ_{u-} over the variables taking values in \mathbf{E}_u than over those taking values in \mathbf{E}_{C_u} . In this case, mutation would need to be carried out before correction and resampling (line 10 in Algorithm 3 before lines 8–9) leading to a slightly different algorithm with essentially the same properties. This approach allows one to specify an analogue of the extended auxiliary distribution independently of the proposal kernel and directly exploit the information it encodes in the resampling step. The cost of doing so is having to perform mutation N^{c_u} times, rather than N times, although this does not increase the overall complexity of the basic algorithm (at least for unfactorized auxiliary measures).

Along similar lines, for targets with particular dependence structures, one may use generalizations of product-form estimators (e.g. [43, Section 3.2 and Appendix F]) that directly account for these structures; see Appendix A.2 for an example. The idea is that, by doing so, one can ease the burden placed on the correction steps which, in Algorithm 3, are the sole responsables for introducing dependencies among children into the approximations.

Lastly, one need not generate the same number of particles at each node and may instead opt to allocate more computational power to the nodes that prove most problematic (e.g. the u s whose associated space E_u have greatest dimension), or even take a more sophisticated approach of the sort in [44]. One would just have to replace γ_u^N in (8) with

$$\left(\prod_{v \in \mathcal{C}} \frac{Z_v^{N_v}}{N_v} \right) \sum_{n_1=1}^{N_{u_1}} \cdots \sum_{n_{uc_u}=1}^{N_{uc_u}} w_{u_-}(\mathbf{X}_{u_1}^{n_1, N_{u_1}}, \dots, \mathbf{X}_{uc_u}^{n_{uc_u}, N_{uc_u}}) \delta_{(\mathbf{X}_{u_1}^{n_1, N_{u_1}}, \dots, \mathbf{X}_{uc_u}^{n_{uc_u}, N_{uc_u}})}(d\mathbf{x}_{\mathcal{C}_u})$$

where, for each child v of u , N_v denotes the amount of particles resampled at node v .

4. Theoretical characterization. We now turn to the main results of the paper showing that Algorithm 3 is well-founded. We give an overview of the proofs for these results in Section 4.1, postponing the full details until Appendices C–G. To simplify the exposition, we only state results for the particle approximations to the targets; however, analogous statements for the approximations to the auxiliary measures and their extensions can be extracted from Appendices C–F. Also to keep the exposition simple, we focus throughout on the case of bounded test functions, writing $\mathcal{B}_b(S)$ for the space of bounded measurable real-valued test functions on a measurable space (S, \mathcal{S}) and $\|\varphi\| := \sup_{x \in S} |\varphi(x)| < \infty$ for the supremum norm on $\mathcal{B}_b(S)$, and we assume that the weight functions are bounded:

ASSUMPTION 2. For all u in \mathbb{T}^∂ and v in \mathbb{T} , $w_{u_-} = d\gamma_{u_-}/d\gamma_{\mathcal{C}_u}$ and $w_v = d\rho_v/d\gamma_v$ are bounded: $\|w_{u_-}\| < \infty$ and $\|w_v\| < \infty$.

To begin with, Algorithm 3 produces strongly consistent estimators for the targets:

THEOREM 1 (Strong laws of large numbers). *If Assumptions 1–2 are satisfied, u belongs to \mathbb{T} , and φ belongs to $\mathcal{B}_b(E_u)$, then*

$$\lim_{N \rightarrow \infty} \rho_u^N(\varphi) = \rho_u(\varphi), \quad \lim_{N \rightarrow \infty} \mu_u^N(\varphi) = \mu(\varphi), \quad \lim_{N \rightarrow \infty} Z_u^N = Z_u, \quad \text{almost surely.}$$

PROOF. See Appendix C. □

If the underlying spaces possess nice enough topological properties, the above almost sure pointwise convergence can be strengthened to almost sure weak convergence:

THEOREM 2 (Almost sure weak convergence). *If, in addition to Assumptions 1–2, the spaces $(E_u)_{u \in \mathbb{T}}$ are Polish and $(\mathcal{E}_u)_{u \in \mathbb{T}}$ are the corresponding Borel sigma algebras, then*

$$\rho_u^N \rightharpoonup \rho_u, \quad \mu_u^N \rightharpoonup \mu_u, \quad \text{almost surely,}$$

for each u in \mathbb{T} , where \rightharpoonup denotes weak convergence as $N \rightarrow \infty$.

PROOF. See Appendix C. □

The estimators for the unnormalized targets are unbiased:

THEOREM 3 (Unbiasedness of $(\rho_u^N)_{u \in \mathbb{T}}$). *If Assumptions 1–2 hold, then*

$$\mathbb{E}[\rho_u^N(\varphi)] = \rho_u(\varphi), \quad \mathbb{E}[Z_u^N] = Z_u, \quad \forall N > 0, \quad \varphi \in \mathcal{B}_b(E_u), \quad u \in \mathbb{T}.$$

PROOF. See Appendix D. □

The non-linearity in the correction step introduces a bias in the estimators for the normalized targets. However, the bias decays linearly with the number of particles N :

THEOREM 4 (Bias estimates for $(\mu_u^N)_{u \in \mathbb{T}}$). *If Assumptions 1–2 hold, the weight functions are bounded below (i.e., for all u in \mathbb{T}^∂ and v in \mathbb{T} , $w_{u_-} \geq \beta_{u_-}$ and $w_v \geq \beta_v$ for some constants $\beta_{u_-}, \beta_v > 0$), and u belongs to \mathbb{T} , then there exists a constant $C_u < \infty$ such that*

$$|\mathbb{E} [\mu_u^N(\varphi)] - \mu_u(\varphi)| \leq \frac{C_u \|\varphi\|}{N} \quad \forall N > 0, \varphi \in \mathcal{B}_b(\mathbf{E}_u).$$

PROOF. See Appendix E. □

All estimators converge at a rate proportional the square root of the number of particles:

THEOREM 5 (L^p inequalities). *If Assumptions 1–2 hold, then, for each $p \geq 1$ and u in \mathbb{T} , then there exist constants $C_u^\rho, C_u^\mu < \infty$ such that*

$$\mathbb{E} \left[|\rho_u^N(\varphi) - \rho(\varphi)|^p \right]^{\frac{1}{p}} \leq \frac{C_u^\rho \|\varphi\|}{N^{1/2}}, \quad \mathbb{E} \left[|\mu_u^N(\varphi) - \mu_u(\varphi)|^p \right]^{\frac{1}{p}} \leq \frac{C_u^\mu \|\varphi\|}{N^{1/2}},$$

for all $N > 0$ and φ in $\mathcal{B}_b(\mathbf{E}_u)$. In particular, $\mathbb{E} [|Z_u^N - Z_u|^p]^{1/p} \leq C_u^\rho / N^{1/2}$ for all $N > 0$.

PROOF. See Appendix C. □

To further characterize the \sqrt{N} rate of convergence of ρ_u^N and μ_u^N , we obtain a central limit theorem for each (given in Theorem 6 below). To state these, we first need to introduce some notation: given two measurable spaces (A, \mathcal{A}) and (B, \mathcal{B}) , a kernel $M : A \times \mathcal{B} \rightarrow [0, \infty)$ from the first to the second, and a test function $\varphi : B \rightarrow \mathbb{R}$, we use $M\varphi$ to denote the real-valued function on A defined by

$$(M\varphi)(a) := \int M(a, db) \varphi(b) \quad \forall a \in A,$$

assuming that the above integrals are well-defined.

As we will see in Theorem 6, the asymptotic variances of our estimators decompose into sums of $|\mathbb{T}_u|$ terms, each one accounting for the variance introduced by the computations carried out in a different descendant v of u . The term corresponding to a particular descendant v involves a kernel $\Gamma_{v,u} : \mathbf{E}_v \times \mathcal{E}_u \rightarrow [0, \infty)$ mapping γ_v to γ_u :

$$(15) \quad \gamma_v(\Gamma_{v,u}\varphi) = \gamma_u(\varphi) \quad \forall \varphi \in \mathcal{B}_b(\mathbf{E}_u).$$

The kernel is defined recursively: $\Gamma_{u,u}(\mathbf{x}_u, d\mathbf{y}_u) = \delta_{\mathbf{x}_u}(d\mathbf{y}_u)$,

$$(16) \quad \Gamma_{v,u}(\mathbf{x}_v, d\mathbf{y}_u) := \delta_{\mathbf{x}_v}(d\mathbf{y}_v) \gamma_{\mathcal{C}_u}^v(d\mathbf{y}_{\mathcal{C}_u}^v) w_{u_-}(\mathbf{y}_{\mathcal{C}_u}) K_u(\mathbf{x}_{\mathcal{C}_u}, d\mathbf{y}_u) \quad \forall v \in \mathcal{C}_u$$

where $\gamma_{\mathcal{C}_u}^v := \prod_{r \in \mathcal{C}_u^v} \gamma_r$ with $\mathcal{C}_u^v := \mathcal{C}_u \setminus \{v\}$. For all other descendants $v \neq u$, we set

$$(17) \quad \Gamma_{v,u} = \Gamma_{v,r_1} \Gamma_{r_1,r_2} \cdots \Gamma_{r_l,u},$$

where v, r_1, \dots, r_l, u denotes the branch of \mathbb{T}_u connecting v and u . Because $\Gamma_{v,u}$ trivially satisfies (15) if v is a child of u , it follows from (17) that (15) also holds for all other descendants v of u . Note that the collection $(\Gamma_{u,v})_{u \in \mathbb{T}, v \in \mathbb{T}_u}$ of all these kernels amounts to a generalization of the semigroup describing the propagation of local sampling errors in standard SMC that is central to much of its analysis (e.g. see [21]).

THEOREM 6 (Central limit theorems). *If Assumptions 1–2 hold, then, as $N \rightarrow \infty$,*

$$N^{1/2} (\rho_u^N(\varphi) - \rho_u(\varphi)) \Rightarrow \mathcal{N}(0, \sigma_{\rho_u}^2(\varphi)), \quad N^{1/2} (\mu_u^N(\varphi) - \mu_u(\varphi)) \Rightarrow \mathcal{N}(0, \sigma_{\mu_u}^2(\varphi)),$$

for any given u in \mathbb{T} and φ in $\mathcal{B}_b(\mathbf{E}_u)$, where \Rightarrow denotes convergence in distribution,

$$\begin{aligned} \sigma_{\rho_u}^2(\varphi) &:= \sum_{v \in \mathbb{T}_u} \pi_v([\mathcal{Z}_v \Gamma_{v,u}[w_u \varphi] - \rho_u(\varphi)]^2), \\ \sigma_{\mu_u}^2(\varphi) &:= \sum_{v \in \mathbb{T}_u} \pi_v([\mathcal{Z}_v \Gamma_{v,u}[w_u Z_u^{-1}[\varphi - \mu_u(\varphi)]]]^2). \end{aligned}$$

In particular, $N^{1/2} (Z_u^N - Z_u) \Rightarrow \mathcal{N}(0, \sigma_{Z_u}^2)$ as $N \rightarrow \infty$ with

$$(18) \quad \sigma_{Z_u}^2 := Z_u^2 \sum_{v \in \mathbb{T}_u} \pi_v \left(\left[\frac{d\mu_u^v}{d\pi_v} - 1 \right]^2 \right),$$

where μ_u^v denotes the \mathbf{E}_v -marginal of μ_u (i.e. $\mu_u^v(A) := \mu_u(A \times E_{\mathbb{T}_u \setminus \mathbb{T}_v})$ for all A in \mathcal{E}_v).

PROOF. See Appendix F. □

4.1. Methods of proof. The approximations to the targets are obtained from those to the extended auxiliary measures via a single importance sampling step, cf. (11). Hence, just as most results for ASMC are easily extracted from those for SMC [36], the fundamental objects we need to analyze here are the approximations to the extended auxiliary measures.

To this end, note that the main difference between DaC-SMC (Algorithm 3) and standard ASMC (Algorithm 2) is the correction step: in the former case, we employ the product-form estimator (8) while, in the latter, we instead use the usual estimator (3). Consequently, our proofs for the above results combine well-known methods previously used to establish analogous results for standard SMC and ASMC with novel techniques that control the errors introduced by the product-form estimators embedded within DaC-SMC.

At a conceptual level, we expand the approximation error into a sum of products of ‘local errors’, each accounting for the approximations made at a different node, and we control them separately (in particular, we show that, just as in the standard case, these errors are $\mathcal{O}(N^{-1/2})$). However, in contrast with standard SMC and ASMC whose local errors only get ‘propagated forward in time’, those of DaC-SMC get multiplied together as we move up the tree. Consequently, we end up with a far greater number of terms in our expansion than normal. Herein, lies the main novelty of our analysis: dispatching these extra products-of-local-errors terms. Because the products are taken over nodes on separate branches, the errors in the products are independent of each other. Hence, the product of k of these $\mathcal{O}(N^{-1/2})$ errors should intuitively be $\mathcal{O}(N^{-k/2})$.

More specifically, we begin by proving the L^p inequalities (Theorem 5). To do so, we follow the iterative approach taken in [14, 51]: we call on the Marcinkiewicz-Zygmund inequality ([21, Lemma 7.3.3]) to obtain L^p inequalities for the empirical distributions of the particles indexed by the tree’s leaves and we show, step-by-step, that the algorithm preserves these inequalities. The resampling and mutation steps then follow from arguments similar to those in [14, 51]. For the correction step, we need to do some extra work and show that the product $\gamma_{\mathcal{C}_u}^N = \prod_{v \in \mathcal{C}_u} \gamma_v^N$ also maintains the inequalities. That is, we need to prove that the L^p norm of the product of the errors is $\mathcal{O}(N^{-1/2})$; something we do roughly in Lemma 3 by using the boundedness of the test and weight functions to uniformly bound all but one of the errors in the product, and exploiting the $\mathcal{O}(N^{-1/2})$ size of the remaining error. To complete the argument for the correction step, we then resume with the approach of [14, 51] and show

that the re-weighting in (8) preserves the inequalities. Next, we extract the laws of large numbers (Theorem 1) from the L^p inequalities using the usual approach involving Chebyshev's inequality and the Borel-Cantelli lemma (e.g. [6, p. 17]). The weak convergence (Theorem 2) then follows using standard techniques described in detail for SMC in [61, Section S1].

To prove the unbiasedness in Theorem 3, we generalize the arguments used in [11, Section 16.4.1] to establish the analogous result for standard SMC. The proof also works recursively: we show that $\gamma_{u_-}^N$ is an unbiased estimator for γ_{u_-} if $\gamma_{v_-}^N$ is an unbiased estimator for γ_{v_-} for each non-leaf child v of u . The unbiasedness of ρ_u^N then follows with a simple use of the tower property.

To obtain the bias bounds in Theorem 4, we follow the approach of [56]. This argument also works its way through the algorithm step by step, this time showing that each step increases the bias by at most $\mathcal{O}(N^{-1})$. Just as with the L^p inequalities, our innovation here (Lemma 8) deals with the product $\gamma_{\mathcal{C}_u}^N = \prod_{v \in \mathcal{C}_u} \gamma_v^N$. As in that case, we obtain the necessary bias bound for the product by uniformly bounding the bias of all but one of the approximations in the product and using a previously derived $\mathcal{O}(N^{-1})$ bound for the remaining approximation.

Our proof for the central limit theorems (Theorem 6) follows the conceptual approach outlined above and extends the arguments in [21, Chapter 9]. Similarly to in that chapter, we obtain an expression for the approximation error in terms of local errors (and their 'propagations up the tree'), the difference being that the expression is not just a sum of the local errors, but of their products too. An argument similar to [21, Corollary 9.3.1] then shows that the local errors are asymptotically normal and independent. Hence, if we are able to demonstrate that their products are $o(N^{-1/2})$, then the CLTs follow as in [21, p. 301] by applying Slutsky's theorem and the continuous mapping theorem, and exploiting the fact that a linear combination of independent normal random variables is normal. To this end, we prove the following product version of the L^2 inequality in Theorem 5:

THEOREM 7 (Product L^2 inequality). *If, in addition to Assumptions 1–2, u and v lie in separate branches (i.e. $u \notin \mathbb{T}_v$ and $v \notin \mathbb{T}_u$), then there exists a constant $C_{u,v} < \infty$ such that*

$$\mathbb{E} [(\gamma_u^N - \gamma_u) \times (\gamma_v^N - \gamma_v)(\varphi)^2]^{\frac{1}{2}} \leq \frac{C_{u,v} \|\varphi\|}{N} \quad \forall N > 0, \quad \varphi \in \mathcal{B}_b(\mathbf{E}_u \times \mathbf{E}_v).$$

PROOF. See Appendix G. □

To the best of our knowledge, this theorem and its proof are unprecedented in the SMC literature. To argue it, we use an inductive approach reminiscent of those we take for the L^p inequalities, unbiasedness, and bias bounds: one by one, we show that each step of algorithm preserves the above inequality. For the base case dealing with empirical distributions of the particles indexed by the tree's leaves, we call on [43, Lemma 1].

4.2. Optimal intermediate targets, auxiliary measures, and proposal kernels. While the optimal choice of intermediate targets, auxiliary measures, and proposal kernels generally depends on the particular average we are interested in estimating, there is one choice that leads to zero variance estimates of the final target's normalizing constant Z_τ . In particular, suppose that the final target ρ_τ and the underlying space $(\mathbf{E}_\tau, \mathcal{E}_\tau)$ are nice enough (e.g. $(\mathbf{E}_\tau, \mathcal{E}_\tau)$ is Borel [37, Theorem 8.5], see also [28]) that, for each u in $\mathbb{T}^\mathcal{J}$, there exists a regular conditional probability distribution mapping the $\mathbf{E}_{\mathcal{C}_u}$ -marginal μ_τ^{u-} (i.e. $\mu_\tau^{u-}(A) := \mu_\tau(A \times \mathbf{E}_{\mathbb{T} \setminus \mathbb{T}_u^u})$ for all A in $\mathcal{E}_{\mathcal{C}_u}$) to its \mathbf{E}_u -marginal μ_τ^u . That is, a Markov kernel $M_u : \mathbf{E}_{\mathcal{C}_u} \times \mathcal{E}_u \rightarrow [0, 1]$ satisfying $\mu_\tau^{u-} \otimes M_u = \mu_\tau^u$. In this case, setting

$$\mu_u := \mu_\tau^u, \quad K_u := M_u \quad \forall u \in \mathbb{T}, \quad \pi_{u_-} := \mu_\tau^{u-} \quad \forall u \in \mathbb{T}^\mathcal{J},$$

we have that

$$\pi_v = \pi_{v_-} \otimes M_v = \mu_{\tau}^{v_-} \otimes M_v = \mu_{\tau}^v = \mu_u^v \quad \forall v \in \mathbb{T}_u, \quad u \in \mathbb{T}.$$

It follows from (18) that our estimator Z_{τ}^N for Z_{τ} achieves zero (asymptotic) variance (and similarly for Z_u^N for all other nodes u). The above choices generalize those for which ASMC yields zero-variance estimates of the marginal likelihood in the context of filtering, cf. [33, Proposition 2]. Note also that, just as in the ASMC case, the normalizing constants in the intermediate targets auxiliary measures are immaterial and we are free to choose them as we find convenient: something unsurprising given that they do not influence the resampling operation, and ultimately cancel in the computation of the normalizing constant estimates. However, perhaps a bit unexpectedly, and in contrast with the ASMC case, these choices do not necessarily lead to zero *finite-sample* variance estimates for the normalizing constants. In particular, given the above,

$$\begin{aligned} w_u &= \frac{d\rho_u}{d\gamma_u} = \frac{Z_u d\mu_u}{Z_u d\pi_u} = \frac{Z_u}{Z_u}, & w_{u_-} &= \frac{d\gamma_{u_-}}{d\gamma_{C_u}} = \frac{Z_u d\pi_{u_-}}{Z_{C_u} d\pi_{C_u}} = \frac{Z_u}{Z_{C_u}} \frac{d\mu_{\tau}^{u_-}}{d\prod_{v \in C_u} \mu_{\tau}^v}, \\ \Rightarrow Z_u^N &= \gamma_u^N(w_u) = \frac{Z_u}{Z_u} Z_u^N = \frac{Z_u}{Z_u} \gamma_{C_u}^N(w_{u_-}) = \frac{Z_u}{Z_{C_u}} \gamma_{C_u}^N \left(\frac{d\mu_{\tau}^{u_-}}{d\prod_{v \in C_u} \mu_{\tau}^v} \right), \end{aligned}$$

and the rightmost term will not be constant unless $\mu_{\tau}^{u_-} = \prod_{v \in C_u} \mu_{\tau}^v$. On the other hand, were this to be the case for all nodes u , we could iterate the above down the tree to find that $Z_u^N = Z_u$ for all u .

In short, the above choices lead to zero-finite-sample-variance estimates for the normalizing constants if and only if \mathbb{T} precisely matches the target's dependence structure. Some insight as to why this might be the case can be gleaned by comparing the asymptotic and finite variance expressions given in [43, Theorem 1] for product-form estimators. In essence, these choices ensure that the ‘one-dimensional’ marginals of the sampling distribution π_u at node u coincide with those of μ_u and negate the corresponding $\mathcal{O}(N^{-1/2})$ terms in the finite-sample variance expansions (i.e. those also featuring in the asymptotic variance expansions). However, the sampling distribution factorizes over u 's children and, unless the target does so too, the higher-dimensional marginals are mismatched and the corresponding $o(N^{-1/2})$ terms in the expansion persevere.

Just as is the case in standard ASMC [33], the above choices can rarely be implemented in practice: with few exceptions, μ_{τ} 's marginals cannot be computed explicitly and the conditional distributions are unknown (much less can be sampled from). A more common situation in practice is that where the intermediate targets are fixed and we seek to choose only the proposal kernels and/or auxiliary measures. In standard ASMC (cf. [11, Chapter 10] and references therein), this is typically done by deriving the ‘locally optimal’ proposal kernel and auxiliary measure for each step of the algorithm (i.e. those that minimize the finite-sample variance of the normalizing constant estimate obtained at that step) and, if these prove intractable, approximating them as best as one can. To extend this approach to DaC-SMC, we provide the following generalization of the characterization for the locally optimal kernels and measures previously obtained for SMC and ASMC (e.g. see [11, Theorems 10.1–10.2]).

THEOREM 8. *Fix any u in \mathbb{T}^{∂} and $N > 0$, and let \mathcal{S} denote the set of all (γ_{u_-}, K_u) pairs satisfying Assumption 1 and $\mathcal{S}(K_u)$ and $\mathcal{S}(\gamma_{u_-})$ be its slices:*

$$\begin{aligned} \mathcal{S} &:= \{(\gamma_{u_-}, K_u) : \gamma_{u_-} \sim \gamma_{C_u}, \quad \rho_u \sim \gamma_{u_-} \otimes K_u\}, \\ \mathcal{S}(K_u) &:= \{\gamma : (\gamma, K_u) \in \mathcal{S}\}, \quad \mathcal{S}(\gamma_{u_-}) := \{K : (\gamma_{u_-}, K) \in \mathcal{S}\}, \quad \forall (\gamma_{u_-}, K_u) \in \mathcal{S}. \end{aligned}$$

If $f(\gamma_{u_-}, K_u)$ denotes $\text{Var}(Z_u^N)$ written explicitly as a function of (γ_{u_-}, K_u) , then, for all (γ_{u_-}, K_u) in \mathcal{S} and $C > 0$,

$$(19) \quad C\sqrt{K_u\omega_u(K_u)^2\rho_{\mathcal{C}_u}} \in \underset{\gamma \in \mathcal{S}(K_u)}{\text{arginf}} f(\gamma, K_u),$$

$$(20) \quad M_u \in \underset{K \in \mathcal{S}(\gamma_{u_-})}{\text{arginf}} f(\gamma_{u_-}, K),$$

$$(21) \quad (C\mu_u^{u_-}, M_u) \in \underset{(\gamma, K) \in \mathcal{S}}{\text{arginf}} f(\gamma, K),$$

where $\omega_u(K_u) := d\rho_u/d\rho_{\mathcal{C}_u} \otimes K_u$ and M_u denotes the regular conditional probability distribution mapping $\mu_u^{u_-}$ to μ_u (i.e. $\mu_u^{u_-} \otimes M_u = \mu_u$), assuming that it exists.

PROOF. See Appendix H. □

4.3. *Comparison with ASMC.* Theorems 1–6 demonstrate that DaC-SMC is a well-founded algorithm: it leads to estimators for the unnormalized and normalized targets with the same basic properties as those possessed by SMC and ASMC estimators [21]. It is only natural to now ask whether it performs better than its ASMC analogue (Algorithm 2).

To answer this question, suppose we are interested in approximating the final target ρ_τ and its normalization μ_τ . We first need to figure out how to approximate ρ_τ and μ_τ using ASMC. This requires somehow lumping the nodes in \mathbb{T} to obtain a line \mathbb{L} (Remark 1). While there are many ways that it could be done, we focus on perhaps the most obvious: merge together all nodes in each level of the tree to obtain a line $\mathcal{L}_0 \rightarrow \mathcal{L}_1 \rightarrow \dots \rightarrow \mathcal{L}_T = \{\tau\}$ whose length equals \mathbb{T} 's depth. Taking the corresponding products of the auxiliary measures,

$$\gamma_{t_-} := \prod_{u \in \mathcal{L}_t^{\vartheta}} \gamma_{u_-}, \quad \pi_{t_-} := \prod_{u \in \mathcal{L}_t^{\vartheta}} \pi_{u_-}, \quad \forall t = 1, \dots, T,$$

we obtain sequences $(\gamma_{t_-})_{t=1}^T$ and $(\pi_{t_-})_{t=1}^T$ that can be used in a standard ASMC set-up. In particular, replacing $(\gamma_{u_-})_{u \in \mathbb{T}^{\vartheta}}$, $(\pi_{u_-})_{u \in \mathbb{T}^{\vartheta}}$, and $(K_u)_{u \in \mathbb{T}}$ with $(\gamma_{t_-})_{t=1}^T$, $(\pi_{t_-})_{t=1}^T$, and $(K_t)_{t=0}^T$, for some proposal kernels $K_t: \mathbf{E}_{t_-} \times \mathcal{E}_t \rightarrow [0, 1]$ mapping from $\mathbf{E}_{t_-} := \prod_{u \in \mathcal{L}_t^{\vartheta}} \mathbf{E}_{\mathcal{C}_u}$ to $\mathcal{E}_t := \mathcal{E}_{\mathcal{L}_t}$, Algorithm 3 reduces to Algorithm 2. To keep both algorithms on as equal a footing as possible, we set the ASMC proposal kernels to be the corresponding products of the DaC-SMC ones:

$$(22) \quad K_t(\mathbf{x}_{t_-}, dx_t) := \left(\prod_{u \in \mathcal{L}_t^{\vartheta}} K_u(dx_u) \right) \left(\prod_{u \in \mathcal{L}_t^{\vartheta}} K_u(\mathbf{x}_{\mathcal{C}_u}, dx_u) \right) \quad \forall t = 0, \dots, T.$$

Theoretically, Theorems 1–6 tell us that both ASMC and DaC-SMC produce consistent estimators of ρ_τ and μ_τ , unbiased in the case of ρ_τ and with an $\mathcal{O}(N^{-1})$ bias in that of μ_τ , whose errors are $\mathcal{O}(N^{-1/2})$ and asymptotically normal. Comparing the asymptotic variance expressions in Theorem 6 for both algorithms, we find that the DaC-SMC estimators are more statistically efficient than the ASMC ones:

THEOREM 9. *If Assumptions 1–2 are satisfied, φ belongs to $\mathcal{B}_b(\mathbf{E}_\tau)$, and $\sigma_{\rho_\tau, \text{smc}}^2(\varphi)$ and $\sigma_{\mu_\tau, \text{smc}}^2(\varphi)$ denote the respective asymptotic variances of the ASMC estimators for $\rho_\tau(\varphi)$ and $\mu_\tau(\varphi)$, then*

$$\sigma_{\rho_\tau}^2(\varphi) \leq \sigma_{\rho_\tau, \text{smc}}^2(\varphi), \quad \sigma_{\mu_\tau}^2(\varphi) \leq \sigma_{\mu_\tau, \text{smc}}^2(\varphi).$$

In particular, the asymptotic variance of the DaC-SMC normalizing constant estimate is bounded above by that of the SMC estimate.

PROOF. See Appendix I. □

Of course, the above theorem, and this entire comparison, applies only if we are using factorized proposal kernels in (22) for ASMC, a needless restriction for this algorithm and one that will often be detrimental to its performance. Indeed, unless the intermediate targets exhibit the same kind of conditional independence encoded in the factorized kernels, standard results for ASMC (e.g. [11, Theorems 10.1,10.2] or [33, Proposition 2]) tell us that these kernels will not be optimal in any usual sense.

However, for such factorized proposals, the gains in statistical efficiency can sometimes be drastic (e.g. exponential in the tree’s degree), see Examples 2 and 4 and Section 3.4 in [43]. In exchange, DaC-SMC generally has a higher computational cost. More specifically, the mutation steps of DaC-SMC require drawing N samples from K_u for each u in \mathbb{T} , while ASMC requires drawing N samples from K_t for each $t = 0, \dots, T$. However, K_t ’s definition implies that drawing a sample from it is equivalent to drawing one from each K_u with u in \mathcal{L}_t and concatenating them. Hence, the cost incurred by the mutation steps of both algorithms is the same. The cost of the correction and resampling steps, however, is not the same. In the case of ASMC, we need to sample N particles from π_{t-}^N for each t in $0, \dots, T$ which requires only N evaluations of $w_{t-} = \prod_{u \in \mathcal{L}_t^\vartheta} w_{u-}$ per t . For DaC-SMC, we must instead sample N particles from π_{u-}^N for each u in \mathbb{T}^ϑ . In the absence of any sort of special structure in the auxiliary measures, this involves N^{c_u} evaluations of w_{u-} per u in \mathbb{T}^ϑ . Hence, for general auxiliary measures, DaC-SMC’s cost is $\mathcal{O}(N^d)$, where d denotes \mathbb{T} ’s degree, while ASMC’s is just $\mathcal{O}(N)$.

However, this extra cost materializes only in nodes with more than one child and concentrates in those with d children. If these nodes only feature high up in the tree, then the high cost incurred by the generation of the particles indexed by the nodes’ children can outweigh the extra overhead incurred by the resampling and correction steps. Similarly if the proposal kernels for those nodes are particularly expensive to sample from. It is worth noting that in many settings of practical interest (e.g. [13, 24, 25]), $d = 2$ and this cost can be borne. Otherwise, as described in Section 3.3, DaC-SMCs cost may be brought down by choosing auxiliary measures that are partially or fully factorized (or sums of such measures)—and when these differ too much from the optimal auxiliary measures the impact of this discrepancy upon the variance can be mitigated via tempering [23]. Lastly, in cases where d is not small and the preferred auxiliary measures do not factorize sufficiently, one can instead turn to the ‘incomplete’ estimators (13) similar to those successfully employed in the U-statistic literature to tackle analogous issues (cf. [40] and references therein).

5. Discussion. In this paper, we study DaC-SMC algorithm introduced in [47] and show that it is theoretically well-founded: it possesses the same basic properties that standard SMC algorithms do (Section 4). To achieve this, we combine well-known methods previously used to study standard SMC algorithms with novel techniques that control the errors introduced by the product-form estimators (8) embedded within DaC-SMC (see Section 4.1 for an overview and Appendices C–G for the details). Our analysis here can be sharpened and refined in various ways. The positivity requirements in Assumptions 1 can be circumvented by emulating the use of stopping times in [21]. The boundedness ones in Assumptions 2 can be avoided by introducing families of appropriately integrable functions at each node, e.g. similarly as in [10, 26]. While the additional requirement in Theorem 4 that the weight functions are lower bounded is common in the SMC literature (e.g. [22, Proposition 9.5.6] and [56]), we anticipate that it could be relaxed by considering a more restricted class of test functions (than that of all bounded and measurable ones) and dealing with the possible ‘extinction’ of the particle cloud as done in [21, Theorem 7.4.3]. Lastly, we do not believe that the particular

lumping construction in Section 4.3 is crucial for the variance bounds in Theorem 9 to hold, only that the ASMC auxiliary measures and proposals are obtained by taking appropriate products of the DaC-SMC ones. In fact, in Appendix H we actually show that lumping the children of a node never reduces the variance of DaC-SMC estimators. However, to argue the bounds in full generality one would likely have to further show that lumping ‘cousins’ does not lower the variance, and we opted to focus on that particular lumping construct so not overly complicate Section 4.3 and Theorem 8’s proof.

Our analysis sheds light on the algorithm, its limitations, and practical use. First, one can improve the algorithm’s performance by approximating optimal intermediate targets, auxiliary measures, and proposal kernels identified in Section 4.2 (this provides a formal explanation for the success of the ad hoc strategies taken in [13, Section 4.1] and [24, Sections 3.7.4, 3.10.5]). Next, whenever the auxiliary measures and proposal kernels used for DaC-SMC are the corresponding products of those used for standard ASMC, the former proves more statistically efficient than the latter (Section 4.3): using the same number of particles, DaC-SMC estimators achieve lower variances than their standard counterparts. In exchange, DaC-SMC generally carries a higher computational cost. How much higher depends on the amount of product structure present in the auxiliary measures $(\gamma_{u_-})_{u \in \mathbb{T}^\theta}$ employed (Section 3.3). If they are fully-factorized (or sums of such measures), then the DaC-SMC’s cost is $\mathcal{O}(N)$ where N denotes the number of particles (the same as ASMC). If they do not factorize at all (nor break down into sums of partially-factorized functions), then DaC-SMC’s cost is $\mathcal{O}(N^d)$ where d denotes the tree’s degree. For anything between these two extremes, the cost is $\mathcal{O}(N^k)$ where $1 < k < d$ depends on amount of factorization (the greater it is, the lower k is).

At each node u , the algorithm approximates the auxiliary measure γ_{u_-} with the product-form $\gamma_{\mathcal{C}_u}^N$ and then corrects this approximation using the weight function w_{u_-} (8). The corrected approximation, $\gamma_{u_-}^N$, is then extended using a mutation step, and the resulting extended approximation, γ_u^N , is corrected using the weight function w_u (11) to produce an approximation, ρ_u^N , of the intermediate target ρ_u . This means that any non-product structure in ρ_u must be introduced into our approximations via these two re-weighting steps (aside from structure linking the components indexed by u ’s children with u itself, which can be introduced using the mutation kernels). Consequently, if the intermediate target ρ_u possesses pronounced non-product structure, we end up with competing interests: (a) we would like to choose the auxiliary measure γ_{u_-} as close to factorized as possible so that the algorithm’s cost is low (Section 3.3) and the first re-weighting does not necessitate large corrections and lead to high-variance approximations of γ_{u_-} ; and (b) we would like to choose the auxiliary measure γ_{u_-} close ρ_u . In fact, the analysis in Section 4.2 suggests that, purely in terms of variance and disregarding any considerations of the algorithm’s cost, γ_{u_-} should be chosen as close as possible to the corresponding marginal of ρ_u . How to balance these two competing considerations is an interesting question beyond the scope of this paper and likely best dealt with on a case-by-case basis.

Alternatively, it might also be possible to introduce non-product structure without blowing up the computational cost or the variance in the correction steps using tempering [23], ‘incomplete’ versions of product-form estimators (13), or generalizations of these estimators that account for more complicated dependence structures (see Appendix A.2 for an example). We find each of these variants, and the question of which one should be used under what circumstances, worthy of further investigation; and we believe this can be done using extensions of the techniques we employed in Appendices C–G. Similarly for the other, perhaps more familiar, refinements that will also likely improve the algorithm’s performance in practice (e.g. adaptive resampling, more efficient resampling schemes, and varying the computational power spent at each node; see Section 3.4).

More broadly, we believe that the techniques and intermediate results developed in this paper might find applications beyond DaC-SMC and its variants. Here, we have in mind recently developed Monte Carlo algorithms for inference in models defined as a merger of several submodels, and motivated by either a genuine interest in the submodel merger (as in [32, 49]) or computational reasons such as facilitating distributed implementations and privacy concerns (e.g. [17, 60, 18]).

Lastly, we anticipate that a variety of interesting applications of the DaC-SMC algorithm will involve its merger with other well-known Monte Carlo methodology and our results pave the way for such combinations. For instance, its merger with pseudo-marginal MCMC [3] to tackle targets with intractable densities and that with particle-marginal Metropolis-Hastings [2] to allow for Metropolis-Hastings proposals with intractable densities; two approaches justified by the unbiasedness in Theorem 3. It can also be used within general particle MCMC [2], including particle Gibbs (e.g. pertinent for models such as Markov random fields with collections of correlated parameter variables lacking time-series structure), the formal justification of which requires the characterization of the joint law of all random variables generated during the running of the Algorithm 3. This characterization can be found in [16, Chapter 8] which generalizes the argument given in [47, Appendix A.1] for balanced binary trees.

In summary, DaC-SMC is a theoretically-sound algorithm and a promising extension of SMC. Its development, however, still lies in its infancy and there remain many open questions regarding its use. For instance, ‘how exactly should we carry out the resampling and which nodes should we allocate more computational power to?’, ‘how do we introduce non-product structure into the approximations while balancing the algorithm’s cost and the variance of the correction steps?’, and ‘with which other Monte algorithms would it prove fruitful to merge DaC-SMC and for what type of targets should which merger be used for?’. We look forward to their resolution.

Appendices

Appendix A contains an example illustrating several concepts discussed in Sections 3.3–3.4. The other appendices contain the proofs of Theorems 1–9 given in the main text. In particular, Appendix B contains notation used throughout Appendices C–I, Appendix C contains the proofs for the DaC-SMC estimators’ consistency and L^p inequalities (Theorems 1, 2, and 5), Appendix D that for the unbiasedness of the unnormalized target estimators (Theorem 3), Appendix E that for the bias bound of the normalized target ones (Theorem 4), Appendix F that for the asymptotic normality (Theorem 6), Appendix G that for the product L^2 inequality (Theorem 7), Appendix H the derivation of the locally optimal proposals and auxiliary measures (Theorem 8), and Appendix I the argument showing that the asymptotic variances of DaC-SMC estimators are bounded above by those of their SMC counterparts (Theorem 9).

APPENDIX A: A FURTHER EXAMPLE

Consider the following time-varying version of the toy hierarchical model in [43, Section 3.1]:

$$Y_{t,l} \sim \mathcal{N}(X_{t,l}, 1), \quad X_{t,l} \sim \mathcal{N}(0, \theta_t), \quad \forall l \in [L], \quad \theta_t \sim f(d\theta; \theta_{t-1}), \quad \forall t \in [T], \quad \theta_0 \sim f(d\theta; 1),$$

where $T, L > 0$ denote integers, $Y_{0:T,1:L} := ((Y_{t,l})_{l=1}^L)_{t=0}^T$ some observed variables, $X_{0:T,1:L} := ((X_{t,l})_{l=1}^L)_{t=0}^T$ some latent variables, $\theta_{0:T} = (\theta_t)_{t=0}^T$ some unknown parameters, and

$f(d\theta; \eta)$ the exponential distribution with mean η . Suppose we observe data $y_{0:T,1:L}$ and wish to draw inferences from the smoothing distribution proportional to

$$\begin{aligned} \rho(d\theta_{0:T}, dx_{0:T,1:L}) := & f(d\theta_0; 1) \left(\prod_{t=1}^T f(d\theta_t; \theta_{t-1}) \right) \\ & \times \left(\prod_{t=0}^T \prod_{l=1}^L \mathcal{N}(y_{t,k}; x_{t,l}, 1) \mathcal{N}(dx_{t,l}; 0, \theta_t) \right). \end{aligned}$$

We consider two approaches with which we could tackle this problem:

A.1. The DaC-SMC approach. One way to obtain tractable approximations of the posterior ρ is to apply DaC-SMC with the following tree. Each parameter variable θ_t has a node t assigned to it (with $E_t = [0, \infty)$), each latent variable $X_{t,l}$ has a node (t, l) assigned to it (with $E_{t,l} = \mathbb{R}$), and the tree has an extra ‘dummy’ root node $T+1$ (which we introduce purely for notational convenience). The nodes assigned to θ_0 and $X_{1:T,1:L}$ are the tree’s leaves and, for each t in $[T+1]$, the children of t are $t-1$ and $(t-1, 1), \dots, (t-1, L)$. For the auxiliary measure γ_{t-} associated with the node t assigned to θ_t , we simply remove all terms in ρ that involve variables θ_s and $x_{s,1}, \dots, x_{s,L}$ with $s \geq t$:

$$\begin{aligned} \gamma_{(t+1)-}(d\theta_{0:t}, dx_{0:t,1:L}) := & f(d\theta_0; 1) \left(\prod_{t'=1}^t f(d\theta_{t'}; \theta_{t'-1}) \right) \\ & \times \left(\prod_{t'=0}^t \prod_{l=1}^L \mathcal{N}(y_{t',l}; x_{t',l}, 1) \mathcal{N}(dx_{t',l}; 0, \theta_{t'}) \right) \quad \forall t = 0, \dots, T. \end{aligned}$$

(Note that $\gamma_{(T+1)-}$ coincides with the unnormalized smoothing distribution ρ , so $\pi_{(T+1)-}^N$ approximates the smoothing distribution.) For the proposal kernels we pick

$$(23) \quad K_0(d\theta_0) := f(d\theta_0; 1), \quad K_t(\theta_{t-1}, d\theta_t) := f(d\theta_t; \theta_{t-1}) \quad \forall t = 1, \dots, T,$$

$$(24) \quad K_{t,k}(dx_{t,l}) := \mathcal{N}(dx_{t,l}; y_{t,l}, 1) \quad \forall t = 0, \dots, T, \quad l = 1, \dots, L;$$

in which case

$$w_{(t+1)-}(\theta_t, x_{t,1:L}) = \prod_{l=1}^L \mathcal{N}(x_{t,l}; 0, \theta_{t-1}) \quad \forall t = 0, \dots, T.$$

Hence, with $g_{t,l}(\theta) := \sum_{n_l=1}^N \mathcal{N}(X_{t,l}^{n_l}; 0, \theta)$ and $g_t(\theta) := \prod_{l=1}^L g_{t,l}(\theta)$,

$$\begin{aligned} (25) \quad \frac{N^{L+1}}{\mathcal{Z}_{\mathcal{C}_t}^N} \gamma_{(t+1)-}^N &= \sum_{n=1}^N \left(\prod_{l=1}^L \sum_{n_l=1}^N \mathcal{N}(X_{t,l}^{n_l}; 0, \theta_t^n) \delta_{X_{t,l}^{n_l}} \right) \delta_{(\theta_{1:t}^n, X_{1:t-1,1:L}^n)} \\ &= \sum_{n=1}^N \left(\prod_{l=1}^L \sum_{n_l=1}^N \frac{\mathcal{N}(X_{t,l}^{n_l}; 0, \theta_t^n)}{g_{t,l}(\theta_t^n)} \delta_{X_{t,l}^{n_l}} \right) g_t(\theta_t^n) \delta_{(\theta_{1:t}^n, X_{1:t-1,1:L}^n)}. \end{aligned}$$

(Note that we are omitting the N superscripts from θ_t^n and $X_{t,l}^n$ to simplify the notation.)

Setting $\omega_{t,l}^{n,n_l} := \frac{\mathcal{N}(X_{t,l}^{n_l}; 0, \theta_t^n)}{g_{t,l}(\theta_t^n)}$ and $\omega_t^n := \frac{g_t(\theta_t^n)}{\sum_{m=1}^N g_t(\theta_t^m)}$, and normalizing, we find that

$$\pi_{(t+1)-}^N = \sum_{n=1}^N \left(\prod_{l=1}^L \sum_{n_l=1}^N \omega_{t,l}^{n,n_l} \delta_{X_{t,l}^{n_l}} \right) \omega_t^n \delta_{(\theta_{1:t}^n, X_{1:t-1,1:L}^n)}.$$

Hence, we are able to draw N samples from $\pi_{(t+1)-}^N$ in $\mathcal{O}(N^2)$ operations by: (a) evaluating $(\omega_t^n)_{n=1}^N$; (b) drawing N indices $(m_n)_{n=1}^N$ from $(\omega_t^n \delta_n)_{n=1}^N$ using the alias method; (c) for each n and l , evaluating $(\omega_t^{m_n, k})_{k=1}^N$; (d) for each n and l , drawing an index $k_{n,l}$ from $(\omega_t^{m_n, k} \delta_k)_{k=1}^N$; and (e) setting

$$\mathbf{X}_{(t+1)-}^n := (\theta_{1:t}^{m_n}, X_{1:t-1,1:L}^{m_n}, X_{t,1}^{k_{n,1}}, \dots, X_{t,L}^{k_{n,L}}) \quad \forall n \in [N].$$

This process can be trivially sped up by storing and re-using the weights $(\omega_t^{m_n, k})_{k=1}^N$ for indices m_n drawn more than once. If there are m_n s drawn many times, one may be better off computing probability and alias tables for the corresponding weights $(\omega_t^{m_n, k})_{k=1}^N$ and using these to generate the corresponding $k_{n,l}$ s.

A.2. An alternative approach. This model has two types of direct dependencies linking its variables: θ_{t+1} is drawn from an exponential distribution with mean θ_t and $X_{t,l}$ is drawn from a zero-mean Gaussian distribution with variance θ_t . The first of these two is encoded into the kernels in (23,24): we propose a value for θ_{t+1}^n by sampling $f(d\theta_{t+1}; \theta_t^n)$. The second type, however, is not: $X_{t,l}^n$ are drawn independently of θ_t^n by sampling from $\mathcal{N}(y_{t,l}, 1)$.

To account for the second type of dependency, we could instead try drawing

$$\theta_t^n \sim f(d\theta_t; \theta_{t-1}^n) \text{ (or } f(d\theta_0; 1) \text{ if } t=0), \quad X_{t,1:L}^{n,m} \sim \prod_{l=1}^L \mathcal{N}\left(\frac{\theta_t^n}{1+\theta_t^n} y_{t,l}, \frac{\theta_t^n}{1+\theta_t^n}\right),$$

for each for all n and m in $[N]$; resulting in a total cost of $\mathcal{O}(N^2)$. We would then have to replace the product-form approximation to $\gamma_{(t+1)-}$ in (25) with the following ‘partially product-form’ one [43, Section 3.2]:

$$\begin{aligned} \gamma_{(t+1)-}^N &:= \frac{\mathcal{Z}_t^N}{N^{L+1}} \sum_{n=1}^N \left[\prod_{l=1}^L \sum_{n_l=1}^N \mathcal{N}(y_{t,l}; 0, 1 + \theta_t^n) \delta_{X_{t,l}^{n, n_l}} \right] \delta_{(\theta_{1:t}^n, X_{1:t-1,1:L}^n)} \\ &= \frac{\mathcal{Z}_t^N}{N^{L+1}} \sum_{n=1}^N \left[\prod_{l=1}^L \frac{1}{N} \sum_{n_l=1}^N \delta_{X_{t,l}^{n, n_l}} \right] g_t(\theta_t^n) \delta_{(\theta_{1:t}^n, X_{1:t-1,1:L}^n)}, \end{aligned}$$

where $g_t(\theta) := \prod_{l=1}^L \mathcal{N}(y_{t,l}; 0, 1 + \theta)$ and \mathcal{Z}_t^N denotes the mass of the previously computed γ_{t-}^N (with $\mathcal{Z}_0^N := 1$). Drawing N samples from the approximation’s normalization,

$$\pi_{(t+1)-}^N := \frac{\gamma_{(t+1)-}^N}{\mathcal{Z}_{(t+1)-}^N} = \sum_{n=1}^N \left[\prod_{l=1}^L \frac{1}{N} \sum_{n_l=1}^N \delta_{X_{t,l}^{n, n_l}} \right] \omega_t^n \delta_{(\theta_{1:t}^n, X_{1:t-1,1:L}^n)}$$

where $\omega_t^n := g_t(\theta_t^n) / \sum_{m=1}^N g_t(\theta_t^m)$, can then be done in $\mathcal{O}(N)$ operations by: (a) evaluating $(\omega_t^n)_{n=1}^N$; (b) drawing N indices $(m_n)_{n=1}^N$ from $(\omega_t^n \delta_n)_{n=1}^N$ using the alias method; (c) for each n and l , drawing an index $k_{n,l}$ from $(N^{-1} \delta_k)_{k=1}^N$; and (d) setting

$$\mathbf{X}_{(t+1)-}^n := (\theta_{1:t}^{m_n}, X_{1:t-1,1:L}^{m_n}, X_{t-1,1}^{m_n, k_{n,1}}, \dots, X_{t-1,L}^{m_n, k_{n,L}}) \quad \forall n \in [N].$$

APPENDIX B: IMPORTANT NOTATION FOR THE PROOFS

Throughout Appendices C–I, we employ the notation introduced in Sections 3.1–3.2 and the following:

- Unless specified otherwise, $(\Omega, \mathcal{F}, \mathbb{P})$ denotes the underlying probability space on which all random variables introduced in Algorithm 3 (including those indexed by different N s) are jointly defined, and \mathbb{E} denotes expectations with respect to \mathbb{P} .

- C denotes a generic positive constant dependent on u , v , and p but not on φ and N that may change from one line to the next. We use this notation whenever we are interested only in the existence of such a constant and its precise value is unimportant. To not overly repeat ourselves, throughout the appendices, we often introduce C without mention (e.g. Lemma 1 below) and leave the corresponding ‘there exists some $C < \infty$ such that’ or ‘for some $C < \infty$ ’ statements implicit.
- Given any two subsets $A \subseteq B \subseteq \mathbb{T}$, a measure ν on (E_A, \mathcal{E}_A) , and measurable function ψ on (E_B, \mathcal{E}_B) , we use $\nu(\psi)$ to denote the measurable function on $(E_{B \setminus A}, \mathcal{E}_{B \setminus A})$ obtained by integrating the arguments of ψ indexed by A with respect to ν :

$$\nu(\psi)(x_{B \setminus A}) := \int \psi(x_A, x_{B \setminus A}) \nu(dx_A) \quad \forall x_{B \setminus A} \in E_{B \setminus A},$$

under the assumption that the integral is well-defined for all $x_{B \setminus A}$ in $E_{B \setminus A}$.

- For any u in \mathbb{T}^∂ and $N > 0$, $\epsilon_u^N := N^{-1} \sum_{n=1}^N \delta_{\mathbf{X}_{u_-}^{n,N}}$ denotes the empirical distribution of the resampled particles (not to be confused with π_u^N in (10) which is empirical distribution of these particles after mutation) and $\pi_{\mathcal{C}_u}^N := \prod_{v \in \mathcal{C}_u} \pi_v^N$ denotes the product of the particle approximations to the extended normalized auxiliary measures indexed by u 's children. Additionally, $\mathcal{F}_{\mathcal{C}_u}^N$ denotes the sigma algebra generated by the mutated particles $(\mathbf{X}_{\mathcal{C}_u}^{n,N})_{n=1}^N := (\mathbf{X}_{u_1}^{n,N}, \dots, \mathbf{X}_{u_{c_u}}^{n,N})_{n=1}^N$ indexed by u 's children, and that $\mathcal{F}_{u_-}^N$ that generated by the resampled particles $(\mathbf{X}_{u_-}^{n,N})_{n=1}^N$ indexed by u .

APPENDIX C: PROOFS OF THEOREMS 1, 2, AND 5

Throughout this appendix, we use the notation described in Appendix B. The aim of the appendix is to prove L^p inequalities in Theorem 5. The laws of large numbers in Theorem 1 follow from the inequalities and the ensuing consequence of the Borel-Cantelli lemma:

LEMMA 1 (e.g. [6, p. 17]). *Let $(Z_N)_{N=1}^\infty$ be a sequence of real-valued random variables. If there exists a $p > 2$ such that*

$$\mathbb{E}[|Z_N|^p] \leq \frac{C}{N^{p/2}} \quad \forall N > 0,$$

then $Z_N \rightarrow 0$ with probability one as $N \rightarrow \infty$.

The weak convergence in Theorem 2 follows directly:

PROOF OF THEOREM 2. Because the product of (finitely-many) Borel σ -algebras on Polish spaces coincides with the Borel σ -algebra on the product space, as the former are separable cf. [37, Lemma 1.2], and the latter is also Polish, this follows immediately from Theorem 1 and the equivalence established in [5, Theorem 2.2]. See [61, Section S1, Theorem 1] for details in an SMC context. \square

To establish the L^p inequalities in Theorem 5, we first obtain such results for $(\pi_u^N)_{u \in \mathbb{T}}$:

THEOREM 10 (L^p inequalities for $(\pi_u^N)_{u \in \mathbb{T}}$). *If Assumptions 1–2 hold, then*

$$\mathbb{E} \left[\left| \pi_u^N(\varphi) - \pi_u(\varphi) \right|^p \right]^{\frac{1}{p}} \leq \frac{C \|\varphi\|}{N^{1/2}}, \quad \forall N > 0, \quad \varphi \in \mathcal{B}_b(\mathbf{E}_u), \quad p \geq 1, \quad u \in \mathbb{T}.$$

To prove the above we follow the approach for standard SMC in [14] adapted to cover the case of general p (as in, for example, [51]). In particular, we derive an L^p inequality for the empirical distributions of the particles indexed by the tree's leaves and show that each step of the algorithm preserves this inequality. Let's start: the Marcinkiewicz-Zygmund-type inequality in [21, Lemma 7.3.3] shows that

$$\mathbb{E} \left[\left| \pi_u^N(\varphi) - \pi_u(\varphi) \right|^m \right]^{\frac{1}{m}} \leq \frac{C \|\varphi\|}{N^{1/2}} \quad \forall N > 0, \varphi \in \mathcal{B}_b(\mathbf{E}_u),$$

for all u in \mathbb{T}^∂ and positive integers m . We extend the above to all $p \geq 1$ using Jensen's inequality: if m is the smallest integer no smaller than p , then

$$\begin{aligned} \mathbb{E} \left[\left| \pi_u^N(\varphi) - \pi_u(\varphi) \right|^p \right]^{\frac{1}{p}} &= \mathbb{E} \left[\left(\left| \pi_u^N(\varphi) - \pi_u(\varphi) \right|^m \right)^{\frac{p}{m}} \right]^{\frac{1}{p}} \leq \mathbb{E} \left[\left| \pi_u^N(\varphi) - \pi_u(\varphi) \right|^m \right]^{\frac{1}{m}} \\ (26) \quad &\leq \frac{C \|\varphi\|}{N^{1/2}} \quad \forall N > 0, \varphi \in \mathcal{B}_b(\mathbf{E}_u). \end{aligned}$$

Taking the product of π_v^N over all children v of a node u preserves (26):

LEMMA 2 (Product step). *If, in addition to Assumptions 1–2, (26) is satisfied for each child v (i.e. it holds with v replacing u therein) of a node u in \mathbb{T}^∂ and some $p \geq 1$, then*

$$(27) \quad \mathbb{E} \left[\left| \pi_{\mathcal{C}_u}^N(\varphi) - \pi_{\mathcal{C}_u}(\varphi) \right|^p \right]^{\frac{1}{p}} \leq \frac{C \|\varphi\|}{N^{1/2}} \quad \forall N > 0, \varphi \in \mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u}).$$

The key to proving the above is the following lemma giving an L^p inequality for products of independent random probability measures satisfying their own individual L^p inequalities:

LEMMA 3. *Suppose that $(\eta_1^N)_{N=1}^\infty$ and $(\eta_2^N)_{N=1}^\infty$ are independent sequences of random probability measures defined on some common probability triplet $(\Omega, \mathcal{F}, \mathbb{P})$, respectively taking values in some measurable spaces (S_1, \mathcal{S}_1) and (S_2, \mathcal{S}_2) , and satisfying*

$$(28) \quad \mathbb{E} \left[\left| \eta_k^N(\varphi) - \eta_k(\varphi) \right|^p \right]^{\frac{1}{p}} \leq \frac{C \|\varphi\|}{N^{1/2}} \quad \forall N > 0, \varphi \in \mathcal{B}_b(S_k), k \in \{1, 2\},$$

for some $p \geq 1$, with limits η_1 and η_2 that are also probability measures. Then,

$$(29) \quad \mathbb{E} \left[\left| (\eta_1^N \times \eta_2^N)(\varphi) - (\eta_1 \times \eta_2)(\varphi) \right|^p \right]^{\frac{1}{p}} \leq \frac{C \|\varphi\|}{N^{1/2}} \quad \forall N > 0, \varphi \in \mathcal{B}_b(S_1 \times S_2).$$

PROOF. Fix any $N > 0$ and φ in $\mathcal{B}_b(S_1 \times S_2)$ and note that

$$(30) \quad \begin{aligned} \eta_1^N \times \eta_2^N(\varphi) - \eta_1 \times \eta_2(\varphi) &= (\eta_1^N - \eta_1) \times (\eta_2^N - \eta_2)(\varphi) \\ &\quad + (\eta_1^N - \eta_1) \times \eta_2(\varphi) + \eta_1 \times (\eta_2^N - \eta_2)(\varphi). \end{aligned}$$

Hence, Minkowski's inequality implies that

$$(31) \quad \begin{aligned} \mathbb{E} \left[\left| \eta_1^N \times \eta_2^N(\varphi) - \eta_1 \times \eta_2(\varphi) \right|^p \right]^{\frac{1}{p}} &\leq \mathbb{E} \left[\left| (\eta_1^N - \eta_1) \times (\eta_2^N - \eta_2)(\varphi) \right|^p \right]^{\frac{1}{p}} \\ &\quad + \mathbb{E} \left[\left| \eta_1^N(\eta_2(\varphi)) - \eta_1(\eta_2(\varphi)) \right|^p \right]^{\frac{1}{p}} \\ &\quad + \mathbb{E} \left[\left| \eta_2^N(\eta_1(\varphi)) - \eta_2(\eta_1(\varphi)) \right|^p \right]^{\frac{1}{p}}. \end{aligned}$$

Because η_1 and η_2 are probability measures, $\eta_1(\varphi)$ and $\eta_2(\varphi)$ are bounded above by $\|\varphi\|$. For this reason, the L^p inequalities in (28) imply that

$$(32) \quad \mathbb{E} \left[\left| \eta_1^N(\eta_2(\varphi)) - \eta_1(\eta_2(\varphi)) \right|^p \right]^{\frac{1}{p}} \leq \frac{C \|\eta_2(\varphi)\|}{N^{1/2}} \leq \frac{C \|\varphi\|}{N^{1/2}},$$

$$(33) \quad \mathbb{E} \left[\left| \eta_2^N(\eta_1(\varphi)) - \eta_2(\eta_1(\varphi)) \right|^p \right]^{\frac{1}{p}} \leq \frac{C \|\eta_1(\varphi)\|}{N^{1/2}} \leq \frac{C \|\varphi\|}{N^{1/2}}.$$

To control the remaining term in (31), let \mathcal{F}_2 denote the σ -algebra generated by the η_2^N s. For each ω in Ω , $x_1 \mapsto \int \varphi(x_1, x_2) \eta_2^N(\omega, dx_2) - \eta_2(\varphi)(x_1)$ is a bounded function on S_1 . Hence, (28) and the independence of $(\eta_1^N)_{N=1}^\infty$ and $(\eta_2^N)_{N=1}^\infty$ imply that

$$(34) \quad \begin{aligned} & \mathbb{E} \left[\left| (\eta_1^N - \eta_1) \times (\eta_2^N - \eta_2)(\varphi) \right|^p \right]^{\frac{1}{p}} \\ &= \mathbb{E} \left[\left| \eta_1^N(\eta_2^N(\varphi) - \eta_2(\varphi)) - \eta_1(\eta_2^N(\varphi) - \eta_2(\varphi)) \right|^p \right]^{\frac{1}{p}} \\ &= \mathbb{E} \left[\mathbb{E} \left[\left| \eta_1^N(\eta_2^N(\varphi) - \eta_2(\varphi)) - \eta_1(\eta_2^N(\varphi) - \eta_2(\varphi)) \right|^p \mid \mathcal{F}_2 \right] \right]^{\frac{1}{p}} \\ &\leq \frac{C \mathbb{E} \left[\left| \eta_2^N(\varphi) - \eta_2(\varphi) \right|^p \right]^{\frac{1}{p}}}{N^{1/2}} \leq \frac{C \mathbb{E} \left[\left(\|\eta_2^N(\varphi)\| + \|\eta_2(\varphi)\| \right)^p \right]^{\frac{1}{p}}}{N^{1/2}} \\ &\leq \frac{C \mathbb{E} \left[\left(2 \max \{ \|\eta_2^N(\varphi)\|, \|\eta_2(\varphi)\| \} \right)^p \right]^{\frac{1}{p}}}{N^{1/2}} \leq \frac{2C \|\varphi\|}{N^{1/2}}. \end{aligned}$$

Putting (31–34) together, we obtain the L^p inequality (29) for the product. \square

PROOF OF LEMMA 2. Because $(\pi_{u_1}^N)_{N=1}^\infty, \dots, (\pi_{u_{c_u}}^N)_{N=1}^\infty$ are independent sequences of probability measures by construction, so are $(\pi_{u[k]}^N)_{N=1}^\infty$ and $(\pi_{u(k+1)}^N)_{N=1}^\infty$ for all $k < c_u$, where $u[k] := \{u_1, \dots, u_k\}$ denotes the set containing the first k children of u and $\pi_{u[k]}^N := \prod_{v \in u[k]} \pi_v^N$ the corresponding product of π_v^N s. Hence, starting from the premise and repeatedly applying Lemma 3, we obtain the L^p inequality for $\pi_{u[c_u]}^N = \pi_{c_u}^N$. \square

Emulating the approach of [14, Lemma 4] and [51, Lemma 1], we find that the correction step also respects the inequality:

LEMMA 4 (Correction step). *If, in addition to Assumptions 1–2, (27) is satisfied for some u in \mathbb{T}^θ and $p \geq 1$, then*

$$(35) \quad \mathbb{E} \left[\left| \pi_{u_-}^N(\varphi) - \pi_{u_-}(\varphi) \right|^p \right]^{\frac{1}{p}} \leq \frac{C \|\varphi\|}{N^{1/2}} \quad \forall N > 0, \quad \varphi \in \mathcal{B}_b(\mathbf{E}_{c_u}).$$

PROOF. Fix any $p \geq 1$, $N > 0$, and φ in $\mathcal{B}_b(\mathbf{E}_{c_u})$. Recall the definitions in Section 3.2:

$$\gamma_{u_-}^N(d\mathbf{x}_{c_u}) = w_{u_-}(\mathbf{x}_{c_u}) \gamma_{c_u}^N(d\mathbf{x}_{c_u}), \quad \gamma_{c_u}^N(d\mathbf{x}_{c_u}) = \mathcal{Z}_{c_u}^N \pi_{c_u}^N(d\mathbf{x}_{c_u}).$$

It follows that

$$\pi_{u_-}^N(\varphi) = \frac{\gamma_{u_-}^N(\varphi)}{\gamma_{u_-}^N(\mathbf{E}_{c_u})} = \frac{\gamma_{c_u}^N(w_{u_-} \varphi)}{\gamma_{c_u}^N(w_{u_-})} = \frac{\pi_{c_u}^N(w_{u_-} \varphi)}{\pi_{c_u}^N(w_{u_-})}.$$

Because $w_{u_-} = d\gamma_{u_-}/d\gamma_{c_u}$, we similarly have that

$$\pi_{u_-}(\varphi) = \frac{\gamma_{u_-}(\varphi)}{\gamma_{u_-}(\mathbf{E}_{c_u})} = \frac{\gamma_{c_u}(w_{u_-} \varphi)}{\gamma_{c_u}(w_{u_-})} = \frac{\pi_{c_u}(w_{u_-} \varphi)}{\pi_{c_u}(w_{u_-})}.$$

Hence,

$$(36) \quad \left| \pi_{u_-}^N(\varphi) - \pi_{u_-}(\varphi) \right| \leq \left| \pi_{u_-}^N(\varphi) - \frac{\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)}{\pi_{\mathcal{C}_u}(w_{u_-})} \right| + \left| \frac{\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)}{\pi_{\mathcal{C}_u}(w_{u_-})} - \frac{\pi_{\mathcal{C}_u}(w_{u_-}\varphi)}{\pi_{\mathcal{C}_u}(w_{u_-})} \right|.$$

To control the first term on the right-hand side, we use

$$(37) \quad \left| \pi_{u_-}^N(\varphi) - \frac{\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)}{\pi_{\mathcal{C}_u}(w_{u_-})} \right| \leq \frac{|\pi_{u_-}^N(\varphi)| |\pi_{\mathcal{C}_u}(w_{u_-}) - \pi_{\mathcal{C}_u}^N(w_{u_-})|}{\pi_{\mathcal{C}_u}(w_{u_-})} \\ \leq \frac{\|\varphi\| |\pi_{\mathcal{C}_u}(w_{u_-}) - \pi_{\mathcal{C}_u}^N(w_{u_-})|}{\pi_{\mathcal{C}_u}(w_{u_-})}.$$

Because, $\pi_{\mathcal{C}_u}(w_{u_-}) = \mathcal{Z}_{\mathcal{C}_u}^{-1} \gamma_{\mathcal{C}_u}(w_{u_-}) = \mathcal{Z}_{\mathcal{C}_u}^{-1} \gamma_{u_-}(\mathbf{E}_{\mathcal{C}_u}) = \mathcal{Z}_{\mathcal{C}_u}^{-1} \mathcal{Z}_u$, the desired L^p inequality (35) follows from (27,36,37) and Minkowski's inequality:

$$\mathbb{E} \left[\left| \pi_{u_-}^N(\varphi) - \pi_{u_-}(\varphi) \right|^p \right]^{\frac{1}{p}} \\ \leq \frac{\|\varphi\| \mathbb{E} \left[\left| \pi_{\mathcal{C}_u}(w_{u_-}) - \pi_{\mathcal{C}_u}^N(w_{u_-}) \right|^p \right]^{\frac{1}{p}} + \mathbb{E} \left[\left| \pi_{\mathcal{C}_u}(w_{u_-}\varphi) - \pi_{\mathcal{C}_u}^N(w_{u_-}\varphi) \right|^p \right]^{\frac{1}{p}}}{\pi_{\mathcal{C}_u}(w_{u_-})} \\ \leq \frac{C \|\varphi\| \|w_{u_-}\| + C \|w_{u_-}\varphi\|}{N^{1/2} \pi_{\mathcal{C}_u}(w_{u_-})} \leq \left(\frac{2C \mathcal{Z}_{\mathcal{C}_u} \|w_{u_-}\|}{\mathcal{Z}_u} \right) \frac{\|\varphi\|}{N^{1/2}}.$$

□

To show that the resampling step also preserves L^p inequalities, we tweak [14, Lemma 5] (recall that $\epsilon_u^N := N^{-1} \sum_{n=1}^N \delta_{\mathbf{X}_{u_-}^{n,N}}$ denotes the resampled particles' empirical distribution):

LEMMA 5 (Resampling step). *If, in addition to Assumptions 1–2, the inequality (35) is satisfied for some u in \mathbb{T}^{ϑ} and $p \geq 1$, then*

$$(38) \quad \mathbb{E} \left[\left| \epsilon_u^N(\varphi) - \pi_{u_-}(\varphi) \right|^p \right]^{\frac{1}{p}} \leq \frac{C \|\varphi\|}{N^{1/2}} \quad \forall N > 0, \quad \varphi \in \mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u}).$$

Key to proving the above is the following conditional Marcinkiewicz-Zygmund inequality:

LEMMA 6 (e.g. [67, Theorem 3.3]). *Let Y_1, \dots, Y_N denote bounded real-valued random variables on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ that are independent when conditioned on some sigma algebra $\mathcal{G} \subseteq \mathcal{F}$ and satisfy $\mathbb{E}[Y_n | \mathcal{G}] = 0$ almost surely for all n in $[N]$. For any given $p \geq 1$, there exists a constant C independent of Y_1, \dots, Y_N and N such that*

$$\mathbb{E} \left[\left| \sum_{n=1}^N Y_n \right|^p \middle| \mathcal{G} \right] \leq C \mathbb{E} \left[\left(\sum_{n=1}^N Y_n^2 \right)^{\frac{p}{2}} \middle| \mathcal{G} \right] \quad \text{almost surely.}$$

Lemma 5's proof is now straightforward:

PROOF OF LEMMA 5. Fix any $p \geq 1$, $N > 0$, and φ in $\mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u})$. Conditioning on $\mathcal{F}_{\mathcal{C}_u}^N$ (the sigma algebra generated by the mutated particles indexed by u 's children), we have that

$$\mathbb{E} [\varphi(\mathbf{X}_{u_-}^{n,N}) | \mathcal{F}_{\mathcal{C}_u}^N] = \pi_{u_-}^N(\varphi) \quad \text{almost surely, for all } n \leq N.$$

Hence, applying Lemma 6 with $\mathcal{G} := \mathcal{F}_{\mathcal{C}_u}^N$ and $Y_n := N^{-1}[\varphi(\mathbf{X}_{u_-}^{n,N}) - \pi_{u_-}^N(\varphi)]$, the tower rule, and the bound $Y_n^2 \leq 4N^{-2} \|\varphi\|^2$, we obtain

$$(39) \quad \mathbb{E} \left[|\epsilon_u^N(\varphi) - \pi_{u_-}^N(\varphi)|^p \right]^{\frac{1}{p}} \leq \frac{C \|\varphi\|}{N^{1/2}}.$$

Inequality (38) then follows from the above, Minkowski's inequality, and (35). \square

For the mutation step, we adapt [14, Lemma 3]:

LEMMA 7 (Mutation step). *If, in addition to Assumptions 1–2, (38) is satisfied for some u in $\mathbb{T}^{\mathcal{J}}$ and $p \geq 1$, then (26) is also satisfied for the same u and p .*

PROOF. Fix any $p \geq 1$, $N > 0$, and φ in $\mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u})$. Minkowski's inequality implies that

$$(40) \quad \mathbb{E} \left[|\pi_u^N(\varphi) - \pi_u(\varphi)|^p \right]^{\frac{1}{p}} \leq \mathbb{E} \left[|\pi_u^N(\varphi) - \nu_u^N(\varphi)|^p \right]^{\frac{1}{p}} + \mathbb{E} \left[|\nu_u^N(\varphi) - \pi_u(\varphi)|^p \right]^{\frac{1}{p}},$$

where $\nu_u^N := \epsilon_u^N \otimes K_u$. Because $\pi_u = \pi_{u_-} \otimes K_u$ and $\|K_u \varphi\| \leq \|\varphi\|$ as K_u is a Markov kernel, (38) implies that

$$(41) \quad \mathbb{E} \left[|\nu_u^N(\varphi) - \pi_u(\varphi)|^p \right]^{\frac{1}{p}} \leq \frac{C \|K_u \varphi\|}{N^{1/2}} \leq \frac{C \|\varphi\|}{N^{1/2}}.$$

To control the other term in (40), we condition on $\mathcal{F}_{u_-}^N$ (the sigma algebra generated by the resampled particles indexed by u) and obtain

$$\mathbb{E} [\varphi(\mathbf{X}_u^{n,N}) | \mathcal{F}_{u_-}^N] = (K_u \varphi)(\mathbf{X}_{u_-}^{n,N}) \quad \forall n \leq N.$$

Hence, applying Lemma 6 similarly as for (39), only this time with $\mathcal{G} := \mathcal{F}_{u_-}^N$ and $Y_n := N^{-1}[\varphi(\mathbf{X}_u^{n,N}) - (K_u \varphi)(\mathbf{X}_{u_-}^{n,N})]$, we find that

$$(42) \quad \mathbb{E} \left[|\pi_u^N(\varphi) - \nu_u^N(\varphi)|^p \right]^{\frac{1}{p}} \leq \frac{C \|\varphi\|}{N^{1/2}}.$$

Combining (40–42) completes the proof. \square

PROOF OF THEOREM 10. The theorem follows by repeatedly applying Lemmas 2, 4, 5, and 7, starting from (26). \square

PROOF OF THEOREM 5. Suppose that we are able to argue that

$$(43) \quad \mathbb{E} \left[|\mathcal{Z}_u^N - \mathcal{Z}_u|^p \right]^{\frac{1}{p}} \leq \frac{C}{N^{1/2}} \quad \forall N > 0,$$

for all u in \mathbb{T} and $p \geq 1$. Because, for all φ in $\mathcal{B}_b(\mathbf{E}_u)$,

$$\begin{aligned} |\gamma_u^N(\varphi) - \gamma_u(\varphi)| &= |\mathcal{Z}_u^N \pi_u^N(\varphi) - \mathcal{Z}_u \pi_u(\varphi)| \\ &\leq |\pi_u^N(\varphi)| |\mathcal{Z}_u^N - \mathcal{Z}_u| + \mathcal{Z}_u |\pi_u^N(\varphi) - \pi_u(\varphi)| \\ &\leq \|\varphi\| |\mathcal{Z}_u^N - \mathcal{Z}_u| + \mathcal{Z}_u |\pi_u^N(\varphi) - \pi_u(\varphi)| \end{aligned}$$

we would then obtain an L^p inequality for γ_u^N from that for π_u^N (Theorem 10). Given that

$$\rho_u^N(d\mathbf{x}_u) = w_u(\mathbf{x}_u) \gamma_u^N(d\mathbf{x}_u), \quad \mathcal{Z}_u^N = \gamma_u^N(w_u), \quad \mu_u^N(d\mathbf{x}_u) = \frac{w_u(\mathbf{x}_u) \gamma_u^N(d\mathbf{x}_u)}{\mathcal{Z}_u^N},$$

and similarly for γ_u , ρ_u , Z_u , and μ_u , the inequalities in Theorem 5 would then follow using arguments of the type in Lemma 4's proof.

Fix any $p \geq 1$ and $N > 0$. In the case of a leaf u , (43) is trivially satisfied because $Z_u^N = Z_u = 1$ by definition. Suppose, instead, that u is not a leaf and that the L^p inequality holds for each of its children (i.e. (43) holds with v replacing u therein, for each v in \mathcal{C}_u). Because $Z_{u_1}^N, \dots, Z_{u_{C_u}}^N$ are independent by definition, and using the following multinomial expansion

$$Z_{\mathcal{C}_u}^N = \prod_{v \in \mathcal{C}_u} Z_v^N = \prod_{v \in \mathcal{C}_u} [(Z_v^N - Z_v) + Z_v] = \sum_{A \subseteq \mathcal{C}_u} \left(\prod_{v \in A} (Z_v^N - Z_v) \right) Z_{\mathcal{C}_u}^A$$

where $Z_{\mathcal{C}_u}^A := Z_{\mathcal{C}_u \setminus A}$, Minkowski's inequality implies that

$$\begin{aligned} \mathbb{E} \left[|Z_{\mathcal{C}_u}^N - Z_{\mathcal{C}_u}|^p \right]^{\frac{1}{p}} &\leq \sum_{\emptyset \neq A \subseteq \mathcal{C}_u} \mathbb{E} \left[\left| \prod_{v \in A} (Z_v^N - Z_v) \right|^p \right]^{\frac{1}{p}} Z_{\mathcal{C}_u}^A \\ &= \sum_{\emptyset \neq A \subseteq \mathcal{C}_u} Z_{\mathcal{C}_u}^A \prod_{v \in A} \mathbb{E} \left[|Z_v^N - Z_v|^p \right]^{\frac{1}{p}} \leq \frac{\sum_{\emptyset \neq A \subseteq \mathcal{C}_u} Z_{\mathcal{C}_u}^A}{N^{1/2}}. \end{aligned}$$

Given the above and

$$\begin{aligned} |Z_u^N - Z_u| &= |Z_{\mathcal{C}_u}^N \pi_{\mathcal{C}_u}^N(w_{u_-}) - Z_{\mathcal{C}_u} \pi_{\mathcal{C}_u}(w_{u_-})| \\ &\leq |\pi_{\mathcal{C}_u}^N(w_{u_-})| |Z_{\mathcal{C}_u}^N - Z_{\mathcal{C}_u}| + Z_{\mathcal{C}_u} |\pi_{\mathcal{C}_u}^N(w_{u_-}) - \pi_{\mathcal{C}_u}(w_{u_-})| \\ &\leq \|w_{u_-}\| |Z_{\mathcal{C}_u}^N - Z_{\mathcal{C}_u}| + Z_{\mathcal{C}_u} |\pi_{\mathcal{C}_u}^N(w_{u_-}) - \pi_{\mathcal{C}_u}(w_{u_-})|, \end{aligned}$$

(43) follows from the L^p inequality for $\pi_{\mathcal{C}_u}^N$ in (27) (obtained in the proof of Theorem 10) and Minkowski's inequality. Hence, starting from the nodes whose children are leaves and working our way inductively up the tree, we obtain (43) for all u in \mathbb{T} . \square

APPENDIX D: PROOF OF THEOREM 3

Throughout this appendix, we use the notation described in Appendix B. Because $\rho_u^N(\varphi) = \gamma_u^N(w_u \varphi)$ and w_u is bounded (Assumption 2), we need only show that γ_u^N is unbiased. That is, for all u in \mathbb{T} ,

$$(44) \quad \mathbb{E} [\gamma_u^N(\varphi)] = \gamma_u(\varphi) \quad \forall N > 0, \varphi \in \mathcal{B}_b(\mathbf{E}_u).$$

If u is a leaf, the above holds trivially. For any other u , note that $\mathbf{X}_u^{n,N}$, by definition, has law $\pi_{u_-}^N \otimes K_u$ when conditioned on $\mathcal{F}_{\mathcal{C}_u}^N$. Because Z_u^N is $\mathcal{F}_{\mathcal{C}_u}^N$ -measurable, it follows that

$$\begin{aligned} \mathbb{E} [\gamma_u^N(\varphi)] &= \frac{1}{N} \sum_{n=1}^N \mathbb{E} [Z_u^N \varphi(\mathbf{X}_u^{n,N})] = \frac{1}{N} \sum_{n=1}^N \mathbb{E} [\mathbb{E} [Z_u^N \varphi(\mathbf{X}_u^{n,N}) | \mathcal{F}_{\mathcal{C}_u}^N]] \\ &= \frac{1}{N} \sum_{n=1}^N \mathbb{E} [Z_u^N \mathbb{E} [\varphi(\mathbf{X}_u^{n,N}) | \mathcal{F}_{\mathcal{C}_u}^N]] = \frac{1}{N} \sum_{n=1}^N \mathbb{E} [Z_u^N \pi_{u_-}^N(K_u \varphi)] = \mathbb{E} [\gamma_{u_-}^N(K_u \varphi)] \end{aligned}$$

for all $N > 0$ and φ and $\mathcal{B}_b(\mathbf{E}_u)$. For this reason, and because $\gamma_u = \gamma_{u_-} \otimes K_u$, we need only show that $\gamma_{u_-}^N(\varphi)$ is unbiased to argue (44). We achieve this by extending the arguments used in [11, Section 16.4.1] to establish the analogous result for SMC:

THEOREM 11. *For any u in \mathbb{T}^\emptyset ,*

$$(45) \quad \mathbb{E} [\gamma_{u_-}^N(\varphi)] = \gamma_{u_-}(\varphi) \quad \forall N > 0, \varphi \in \mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u}).$$

PROOF. We argue the result inductively, starting from the nodes whose children are leaves and recursively moving our way up the tree. For the base case (u 's children are leaves), note that the independence of $(\mathbf{X}_{u1}^{n,N})_{n=1}^N, \dots, (\mathbf{X}_{uc_u}^{n,N})_{n=1}^N$ and the fact that $\mathbf{X}_v^{1,N}, \dots, \mathbf{X}_v^{N,N} \sim K_v$ for each v in \mathcal{C}_u imply that

$$\begin{aligned} \mathbb{E} [\gamma_{u-}^N(\varphi)] &= \frac{1}{N^{c_u}} \sum_{\mathbf{n} \in [N]^{c_u}} \mathbb{E} \left[w_{u-}(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}) \varphi(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}) \right] = \frac{1}{N^{c_u}} \sum_{\mathbf{n} \in [N]^{c_u}} \left(\prod_{v \in \mathcal{C}_u} K_v \right) (w_{u-} \varphi) \\ &= \left(\prod_{v \in \mathcal{C}_u} K_v \right) (w_{u-} \varphi) = \gamma_{\mathcal{C}_u}(w_{u-} \varphi) = \gamma_{u-}(\varphi) \quad \forall N > 0, \varphi \in \mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u}). \end{aligned}$$

For the inductive step, suppose instead that (45) holds for each of u 's non-leaf children:

$$(46) \quad \mathbb{E} [\gamma_{v-}^N(\varphi)] = \gamma_{v-}(\varphi), \quad \forall N > 0, \varphi \in \mathcal{B}_b(\mathbf{E}_{\mathcal{C}_v}), \quad v \in \mathcal{C}_u^\partial.$$

We need to show that (45) also holds for u itself. To do so, note that $(\mathbf{X}_{u1}^{n,N})_{n=1}^N, \dots, (\mathbf{X}_{uc_u}^{n,N})_{n=1}^N$ are independent and that, for each v in \mathcal{C}_u^∂ , $\mathbf{X}_v^{1,N}, \dots, \mathbf{X}_v^{N,N}$ have law $\pi_{v-}^N \otimes K_v$ conditioned on $\mathcal{F}_{\mathcal{C}_v}^N$. Hence,

$$\begin{aligned} \mathbb{E} [\gamma_{u-}^N(\varphi)] &= \frac{1}{N^{c_u}} \sum_{\mathbf{n} \in [N]^{c_u}} \mathbb{E} \left[\mathcal{Z}_{\mathcal{C}_u}^N w_{u-}(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}) \varphi(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}) \right] \\ &= \frac{1}{N^{c_u}} \sum_{\mathbf{n} \in [N]^{c_u}} \mathbb{E} \left[\mathbb{E} \left[\left(\prod_{v \in \mathcal{C}_u^\partial} \mathcal{Z}_v^N \right) w_{u-}(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}) \varphi(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}) \middle| \bigvee_{v \in \mathcal{C}_u^\partial} \mathcal{F}_{\mathcal{C}_v}^N \right] \right] \\ &= \frac{1}{N^{c_u}} \sum_{\mathbf{n} \in [N]^{c_u}} \mathbb{E} \left[\prod_{v \in \mathcal{C}_u^\partial} \mathcal{Z}_v^N \mathbb{E} \left[w_{u-}(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}) \varphi(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}) \middle| \bigvee_{v \in \mathcal{C}_u^\partial} \mathcal{F}_{\mathcal{C}_v}^N \right] \right] \\ &= \frac{1}{N^{c_u}} \sum_{\mathbf{n} \in [N]^{c_u}} \mathbb{E} \left[\prod_{v \in \mathcal{C}_u^\partial} \mathcal{Z}_v^N \left(\prod_{v \in \mathcal{C}_u^\partial} \pi_{v-}^N \otimes K_v \right) \left(\left(\prod_{v \in \mathcal{C}_u} K_v \right) (w_{u-} \varphi) \right) \right] \\ &= \mathbb{E} \left[\left(\prod_{v \in \mathcal{C}_u^\partial} \gamma_{v-}^N \right) (K_{\mathcal{C}_u} [w_{u-} \varphi]) \right] = \mathbb{E} \left[\left(\prod_{v \in \mathcal{C}_u^\partial} \gamma_{v-} \right) (K_{\mathcal{C}_u} [w_{u-} \varphi]) \right] \\ &= \gamma_{\mathcal{C}_u}(w_{u-} \varphi) = \gamma_{u-}(\varphi), \end{aligned}$$

where the third-to-last equality follows from applying (46) and the tower rule, and $K_{\mathcal{C}_u}$ denotes the kernel defined by

$$(47) \quad K_{\mathcal{C}_u}(\mathbf{x}_{\mathcal{C}_u-}, d\mathbf{x}_{\mathcal{C}_u}) := \left(\prod_{v \in \mathcal{C}_u^\partial} K_v(dx_v) \right) \left(\prod_{v \in \mathcal{C}_u} K_v(\mathbf{x}_{\mathcal{C}_v}, dx_v) \right) \quad \forall \mathbf{x}_{\mathcal{C}_u-} \in \prod_{v \in \mathcal{C}_u^\partial} \mathbf{E}_{\mathcal{C}_v}.$$

□

APPENDIX E: PROOF OF THEOREM 4

Throughout this appendix, we use the notation described in Appendix B. By definition,

$$\mu_u^N(\varphi) = \frac{\rho_u^N(\varphi)}{\rho_u^N(\mathbf{E}_u)} = \frac{\gamma_u^N(w_u \varphi)}{\gamma_u^N(w_u)} = \frac{\pi_u^N(w_u \varphi)}{\pi_u^N(w_u)}, \quad \mu_u(\varphi) = \frac{\rho_u(\varphi)}{\rho_u(\mathbf{E}_u)} = \frac{\gamma_u(w_u \varphi)}{\gamma_u(w_u)} = \frac{\pi_u(w_u \varphi)}{\pi_u(w_u)},$$

for all φ in $\mathcal{B}_b(\mathbf{E}_u)$ and u in \mathbb{T} . For this reason, and because w_u is bounded above and below by theorem's premise, the bias bounds for $(\mu_u^N)_{u \in \mathbb{T}}$ in Theorem 4 follow using standard arguments (see [48, p. 35] or (55–60) below) from those for $(\pi_u^N)_{u \in \mathbb{T}}$ in Theorem 12 below.

THEOREM 12 (Bias estimates for $(\pi_u^N)_{u \in \mathbb{T}}$). *If Assumptions 1–2 hold, the weight functions are bounded below (i.e. for every u in \mathbb{T}^∂ , $w_{u_-} \geq \beta_u$ for some constant $\beta_u > 0$), then for all u in \mathbb{T} ,*

$$(48) \quad |\mathbb{E} [\pi_u^N(\varphi)] - \pi_u(\varphi)| \leq \frac{C \|\varphi\|}{N} \quad \forall N > 0, \quad \varphi \in \mathcal{B}_b(\mathbf{E}_u).$$

To argue the above, we follow an approach similar to the one we took in Appendix C for the L^p inequalities: we show that the bias bound holds for the leaves, prove that each step of algorithm preserves the bound, and inductively work our way up the tree. To this end, we require Lemma 8, our main innovation in this appendix, showing that the product operation respects bias bounds.

LEMMA 8. *Suppose that $(\eta_1^N)_{N=1}^\infty$ and $(\eta_2^N)_{N=1}^\infty$ are independent sequences of random probability measures defined on some common probability triplet $(\Omega, \mathcal{F}, \mathbb{P})$, respectively taking values in some measurable spaces (S_1, \mathcal{S}_1) and (S_2, \mathcal{S}_2) , and satisfying bias estimates:*

$$(49) \quad |\mathbb{E} [\eta_k^N(\varphi_k)] - \eta_k(\varphi_k)| \leq \frac{C \|\varphi_k\|}{N} \quad \forall N > 0, \quad \varphi_k \in \mathcal{B}_b(S_k), \quad k \in \{1, 2\},$$

with limits η_1 and η_2 that are also probability measures. The sequence of products $(\eta_1^N \times \eta_2^N)_{N=1}^\infty$ also satisfies a bias estimate:

$$(50) \quad |\mathbb{E} [(\eta_1^N \times \eta_2^N)(\varphi)] - (\eta_1 \times \eta_2)(\varphi)| \leq \frac{C \|\varphi\|}{N} \quad \forall N > 0, \quad \varphi \in \mathcal{B}_b(S_1 \times S_2).$$

PROOF OF LEMMA 8. Fix any φ in $\mathcal{B}_b(S_1 \times S_2)$ and $N > 0$, and recall the decomposition in (30). Because η_1, η_2 are probability measures, $\|\eta_1(\varphi)\|$ and $\|\eta_2(\varphi)\|$ are both bounded above by $\|\varphi\|$ and we can apply the bias estimates in (49) to control the expected value of the rightmost two terms in (30):

$$(51) \quad |\mathbb{E} [\eta_1^N(\eta_2(\varphi))] - \eta_1(\eta_2(\varphi))| \leq \frac{C \|\eta_2(\varphi)\|}{N} \leq \frac{C \|\varphi\|}{N},$$

$$(52) \quad |\mathbb{E} [\eta_2^N(\eta_1(\varphi))] - \eta_2(\eta_1(\varphi))| \leq \frac{C \|\eta_1(\varphi)\|}{N} \leq \frac{C \|\varphi\|}{N}.$$

To control the expected value of the remaining term in the right-hand side of (30), let \mathcal{F}_2 denote the σ -algebra generated by the η_2^N s. Because, for each ω in Ω , $x_1 \mapsto \int \varphi(x_1, x_2) \eta_2^N(\omega, dx_2) - \eta_2(\varphi)(x_1)$ is a bounded function on S_1 , (49) and the independence of $(\eta_1^N)_{N=1}^\infty$ and $(\eta_2^N)_{N=1}^\infty$ imply that

$$\begin{aligned} & |\mathbb{E} [(\eta_1^N - \eta_1) \times (\eta_2^N - \eta_2)(\varphi)]| \\ &= |\mathbb{E} [\eta_1^N(\eta_2^N(\varphi) - \eta_2(\varphi)) - \eta_1(\eta_2^N(\varphi) - \eta_2(\varphi))]| \\ &= |\mathbb{E} [\mathbb{E} [\eta_1^N(\eta_2^N(\varphi) - \eta_2(\varphi)) - \eta_1(\eta_2^N(\varphi) - \eta_2(\varphi)) | \mathcal{F}_2]]| \\ &\leq \mathbb{E} [|\mathbb{E} [\eta_1^N(\eta_2^N(\varphi) - \eta_2(\varphi)) - \eta_1(\eta_2^N(\varphi) - \eta_2(\varphi)) | \mathcal{F}_2]|] \\ &\leq \frac{C \mathbb{E} [|\eta_2^N(\varphi) - \eta_2(\varphi)|]}{N} \leq \frac{2C \|\varphi\|}{N}. \end{aligned}$$

Putting the above and (51,52) together, we obtain the bias estimate for the product in (50). \square

Armed with Lemma 8, the remainder of Theorem 12's proof follows the approach of [56]:

PROOF OF THEOREM 12. In the case of a leaf u , the bound (48) holds trivially:

$$(53) \quad \mathbb{E} [\pi_u^N(\varphi)] = \frac{1}{N} \sum_{n=1}^N \mathbb{E} [\varphi(\mathbf{X}_u^{n,N})] = K_u(\varphi) = \pi_u(\varphi) \quad \forall \varphi \in \mathcal{B}_b(\mathbf{E}_u),$$

because $\mathcal{Z}_u^N = 1$ by definition and $\mathbf{X}_u^{1,N}, \dots, \mathbf{X}_u^{N,N}$ are drawn directly from K_u .

For all other nodes, we argue the result inductively, starting from those whose children are leaves and recursively moving our way up the tree. For the base case (u 's children are leaves), note that $(\pi_{u_1}^N)_{N=1}^\infty, \dots, (\pi_{u_{c_u}}^N)_{N=1}^\infty$ are independent sequences of probability measures by construction; consequently, for any $k < c_u$, $(\pi_{u[k]}^N)_{N=1}^\infty$ and $(\pi_{u(k+1)}^N)_{N=1}^\infty$ are also independent sequences (where $u[k] := \{u_1, \dots, u_k\}$ and $\pi_{u[k]}^N := \prod_{v \in u[k]} \pi_v^N$). For this reason, and because the bound (48) holds for all of u 's children, repeated applications of Lemma 8 yield

$$(54) \quad |\mathbb{E} [\pi_{\mathcal{C}_u}^N(\varphi)] - \pi_{\mathcal{C}_u}(\varphi)| \leq \frac{C \|\varphi\|}{N} \quad \forall N > 0, \quad \varphi \in \mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u}).$$

To proceed, fix any φ in $\mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u})$ and $N > 0$. As shown at the start of Lemma 4's proof,

$$\pi_{u_-}(\varphi) = \frac{\pi_{\mathcal{C}_u}(w_{u_-}\varphi)}{\pi_{\mathcal{C}_u}(w_{u_-})}, \quad \pi_{u_-}^N(\varphi) = \frac{\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)}{\pi_{\mathcal{C}_u}^N(w_{u_-})}.$$

Hence, applying the triangle inequality, we find that

$$(55) \quad \left| \mathbb{E} [\pi_{u_-}^N(\varphi)] - \pi_{u_-}(\varphi) \right| \leq \left| \mathbb{E} \left[\frac{\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)}{\pi_{\mathcal{C}_u}^N(w_{u_-})} \right] - \frac{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)]}{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]} \right| + \left| \frac{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)]}{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]} - \frac{\pi_{\mathcal{C}_u}(w_{u_-}\varphi)}{\pi_{\mathcal{C}_u}(w_{u_-})} \right|.$$

To control the first term, we apply a bivariate Taylor expansion with Lagrange remainder of $(y_1, y_2) \mapsto y_1/y_2$ at $(\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)], \mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})])$ and obtain

$$(56) \quad \frac{\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)}{\pi_{\mathcal{C}_u}^N(w_{u_-})} = \frac{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)]}{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]} + \frac{\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi) - \mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)]}{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]} - \frac{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)]}{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]^2} (\pi_{\mathcal{C}_u}^N(w_{u_-}) - \mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]) + R^N,$$

where

$$R^N := \frac{\theta_1}{\theta_2^3} (\pi_{\mathcal{C}_u}^N(w_{u_-}) - \mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})])^2 - \frac{1}{\theta_2^2} (\pi_{\mathcal{C}_u}^N(w_{u_-}) - \mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]) (\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi) - \mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)]),$$

for some random point (θ_1, θ_2) on the line segment joining $(\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi), \pi_{\mathcal{C}_u}^N(w_{u_-}))$ and $(\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)], \mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})])$. Taking expectations of (56), we find that

$$(57) \quad \left| \mathbb{E} \left[\frac{\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)}{\pi_{\mathcal{C}_u}^N(w_{u_-})} \right] - \frac{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)]}{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]} \right| \leq \mathbb{E} \left[\left| \frac{\theta_1}{\theta_2^3} (\pi_{\mathcal{C}_u}^N(w_{u_-}) - \mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})])^2 \right| \right] + \mathbb{E} \left[\left| \frac{1}{\theta_2^2} (\pi_{\mathcal{C}_u}^N(w_{u_-}) - \mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]) (\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi) - \mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)]) \right| \right].$$

As we will show in Lemma 9 below, for all non-negative integers satisfying $l + k \geq 1$,

$$(58) \quad \mathbb{E} \left[\left| \pi_{\mathcal{C}_u}^M(\psi_1) - \mathbb{E} [\pi_{\mathcal{C}_u}^M(\psi_1)] \right|^l \left| \pi_{\mathcal{C}_u}^M(\psi_2) - \mathbb{E} [\pi_{\mathcal{C}_u}^M(\psi_2)] \right|^k \right] \leq \frac{C \|\psi_1\|^l \|\psi_2\|^k}{M^{(l+k)/2}}$$

for all $M > 0$ and ψ_1, ψ_2 in $\mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u})$. Given our assumption that w_{u_-} is bounded below by some constant $\beta > 0$, we have that

$$\theta_2 \geq \min\{\pi_{\mathcal{C}_u}^N(w_{u_-}), \mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]\} \geq \min\{\beta \pi_{\mathcal{C}_u}^N(\mathbf{E}_{\mathcal{C}_u}), \mathbb{E} [\beta \pi_{\mathcal{C}_u}^N(\mathbf{E}_{\mathcal{C}_u})]\} = \beta.$$

Similarly, $\theta_1 \leq \|w_{u_-}\| \|\varphi\|$ and it follows from (57,58) that

$$(59) \quad \left| \mathbb{E} \left[\frac{\pi_{\mathcal{C}_u}^N(w_{u_-} \varphi)}{\pi_{\mathcal{C}_u}^N(w_{u_-})} \right] - \frac{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-} \varphi)]}{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]} \right| \leq \left(\frac{\|w_{u_-}\|^3 \|\varphi\|}{\beta^3} + \frac{\|w_{u_-}\|^2 \|\varphi\|}{\beta^2} \right) \frac{C}{N}.$$

To deal with the second term in (55), we apply the triangle inequality:

$$\begin{aligned} & \left| \frac{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-} \varphi)]}{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]} - \frac{\pi_{\mathcal{C}_u}(w_{u_-} \varphi)}{\pi_{\mathcal{C}_u}(w_{u_-})} \right| \\ & \leq \frac{|\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-} \varphi)] - \pi_{\mathcal{C}_u}(w_{u_-} \varphi)|}{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]} + \frac{|\pi_{\mathcal{C}_u}(w_{u_-} \varphi)| |\pi_{\mathcal{C}_u}(w_{u_-}) - \mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]|}{\pi_{\mathcal{C}_u}(w_{u_-}) \mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]} \\ & \leq \frac{|\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-} \varphi)] - \pi_{\mathcal{C}_u}(w_{u_-} \varphi)|}{\beta} + \frac{\|\varphi\|}{\beta} |\pi_{\mathcal{C}_u}(w_{u_-}) - \mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]|, \end{aligned}$$

where, in the final inequality, we used $|\pi_{\mathcal{C}_u}(w_{u_-} \varphi)/\pi_{\mathcal{C}_u}(w_{u_-})| = |\pi_{u_-}(\varphi)| \leq \|\varphi\|$. Applying the bias estimate for $\pi_{\mathcal{C}_u}^N$ in (54), we then find that

$$\left| \frac{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-} \varphi)]}{\mathbb{E} [\pi_{\mathcal{C}_u}^N(w_{u_-})]} - \frac{\pi_{\mathcal{C}_u}(w_{u_-} \varphi)}{\pi_{\mathcal{C}_u}(w_{u_-})} \right| \leq \frac{2 \|w_{u_-}\| \|\varphi\| C}{\beta N}.$$

Combining the above with (55,59), we obtain a bias bound for $\pi_{u_-}^N$:

$$(60) \quad |\mathbb{E} [\pi_{u_-}^N(\varphi)] - \pi_{u_-}(\varphi)| \leq \frac{C \|\varphi\|}{N} \quad \forall N > 0, \quad \varphi \in \mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u}).$$

Lastly, given that $\mathbf{X}_u^{1,N}, \dots, \mathbf{X}_u^{N,N}$ have law $\pi_{u_-}^N \otimes K_u$ when conditioned on $\mathcal{F}_{\mathcal{C}_u}^N$,

$$\mathbb{E} [\pi_{u_-}^N(\varphi)] = \frac{1}{N} \sum_{n=1}^N \mathbb{E} [\varphi(\mathbf{X}_u^{n,N})] = \frac{1}{N} \sum_{n=1}^N \mathbb{E} [\mathbb{E} [\varphi(\mathbf{X}_u^{n,N}) | \mathcal{F}_{\mathcal{C}_u}^N]] = \mathbb{E} [\pi_{u_-}^N(K_u \varphi)]$$

for all $N > 0$ and φ in $\mathcal{B}_b(\mathbf{E}_u)$. Because $\pi_u = \pi_{u_-} \otimes K_u$, the desired bias estimate (48) for π_u^N follows from that for $\pi_{u_-}^N$ in (60) for all non-leaf nodes u whose children are leaves.

For the inductive step, take any non-leaf node u , assume that the bound (48) holds for all of its non-leaf children, and repeat the exact same argument given above for the base case. \square

We have one loose end left to tie up:

LEMMA 9. *If Assumptions 1-2 are satisfied and l, k are non-negative integers satisfying $l + k \geq 1$, then (58) holds for all $M > 0$ and ψ_1, ψ_2 in $\mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u})$.*

PROOF. Given the L^p inequality for $\pi_{\mathcal{C}_u}^N$ in (27) (obtained in the proof of Theorem 10), this proof consists of applying Jensen's inequality and the Cauchy-Schwarz inequality as was done in the proof of [56, Lemma 2.2]. \square

APPENDIX F: PROOF OF THEOREM 6

Throughout this appendix, we use the notation described in Appendix B. The aim of the appendix is to establish the CLTs in Theorem 6. The meat of the matter entails deriving a CLT for the estimator γ_u^N of the extended auxiliary measure γ_u :

THEOREM 13 (CLT for the unnormalized flow). *If Assumptions 1–2 hold and u lies in \mathbb{T} ,*

$$N^{1/2} (\gamma_u^N(\varphi) - \gamma_u(\varphi)) \Rightarrow \mathcal{N}(0, \sigma_{\gamma_u}^2(\varphi)) \text{ as } N \rightarrow \infty, \quad \forall \varphi \in \mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u})$$

where $\sigma_{\gamma_u}^2(\varphi) := \sum_{v \in \mathbb{T}_u} \pi_v([\mathcal{Z}_v \Gamma_{v,u} \varphi - \gamma_u(\varphi)]^2)$.

Theorem 6 follows from Theorem 13 using standard arguments:

PROOF OF THEOREM 6. Fix any u in \mathbb{T} and φ in $\mathcal{B}_b(\mathbf{E}_u)$. Because $\rho_u^N(\varphi) = \gamma_u^N(w_u \varphi)$, the CLT for ρ_u^N follows immediately from that for γ_u^N and

$$\sigma_{\gamma_u}^2(w_u \varphi) = \sum_{v \in \mathbb{T}_u} \pi_v([\mathcal{Z}_v \Gamma_{v,u} [w_u \varphi] - \gamma_u(w_u \varphi)]^2) = \sum_{v \in \mathbb{T}_u} \pi_v([\mathcal{Z}_v \Gamma_{v,u} [w_u \varphi] - \rho_u(\varphi)]^2).$$

For $\mu_u^N(\varphi)$ note that, just as in the standard SMC case [21, Section 9.4.2],

$$(61) \quad \mu_u^N(\varphi) - \mu_u(\varphi) = \frac{Z_u}{Z_u^N} \rho_u^N \left(\frac{\varphi - \mu_u(\varphi)}{Z_u} \right).$$

Theorem 1 tells us that Z_u^N converges almost surely to Z_u as N tends to infinity. Moreover,

$$\rho_u \left(\frac{\varphi - \mu_u(\varphi)}{Z_u} \right) = \mu_u(\varphi - \mu_u(\varphi)) = 0.$$

Hence, we obtain the CLT for μ_u^N applying Slutsky's theorem to (61):

$$N^{1/2} (\mu_u^N(\varphi) - \mu_u(\varphi)) \Rightarrow \mathcal{N}(0, \sigma_{\rho_u}^2(Z_u^{-1}[\varphi - \mu_u(\varphi)])) \text{ as } N \rightarrow \infty.$$

To complete the proof, note that

$$\sigma_{Z_u}^2 = \sigma_{\rho_u}^2(1) = \sum_{v \in \mathbb{T}_u} \pi_v([\mathcal{Z}_v \Gamma_{v,u} w_u - Z_u]^2).$$

But, for any v in \mathbb{T}_u and ψ in $\mathcal{B}_b(\mathbf{E}_v)$,

$$\pi_v(\mathcal{Z}_v \Gamma_{v,u} w_u \psi) = \gamma_v(\Gamma_{v,u} w_u \psi) = \gamma_u(w_u \psi) = \rho_u(\psi) = Z_u \mu_u(\psi) = Z_u \mu_u^v(\psi).$$

In other words, $\mathcal{Z}_v \Gamma_{v,u} w_u = Z_u d\mu_u^v/d\pi_v$ and (18) follows. \square

For the proof of Theorem 13, we adapt the approach taken in [21, Chapter 9] to establish the analogous result for standard SMC. In particular, we first show that the *local errors*, $\pi_v^N - \pi_{v_-}^N \otimes K_v$, are $\mathcal{O}(N^{-1/2})$, asymptotically normal, and independent (Lemma 10). Next, we break down the *global error* $\gamma_u^N - \gamma_u$ into a sum of the local errors and their products, both 'propagated up the tree' via $(\Gamma_{v,u})_{v \in \mathbb{T}_u}$ (Lemma 11 and (72)). Using the L^2 product inequality in Lemma 7, we show that the products are $o(N^{-1/2})$ and, hence, vanish in the rescaled limit. An application of the continuous mapping theorem then completes the proof. Let's begin: emulating [22, Theorem 10.6.1]'s proof (or, similarly, those of [21, Theorem 9.31; Corollary 9.3.1]), we obtain the following:

LEMMA 10. Fix any u in \mathbb{T} and a collection $(\varphi_v)_{v \in \mathbb{T}_u}$ with φ_v belonging to $\mathcal{B}_b(\mathbf{E}_v)$ for each v in \mathbb{T}_u . Let

$$(62) \quad V_v^N(\varphi_v) := \begin{cases} N^{1/2}(\pi_v^N(\varphi_v) - K_v(\varphi_v)) & \text{if } v \text{ is a leaf} \\ N^{1/2}(\pi_v^N(\varphi_v) - \pi_{v_-}^N(K_v\varphi_v)) & \text{otherwise} \end{cases} \quad \forall N > 0, \quad v \in \mathbb{T}_u.$$

Under Assumptions 1–2, the collection $(V_v^N(\varphi_v))_{v \in \mathbb{T}_u}$ converges weakly as $N \rightarrow \infty$ to a collection $(V_v(\varphi_v))_{v \in \mathbb{T}_u}$ of independent zero mean random variables with Gaussian laws:

$$(63) \quad V_v(\varphi_v) \sim \mathcal{N}(0, \pi_v([\varphi_v - \pi_v(\varphi_v)]^2)) \quad \forall v \in \mathbb{T}_u.$$

PROOF. Because $(K_v)_{v \in \mathbb{T}_u}$ are Markov kernels, Lemma 1 and the L^p inequality for $\pi_{v_-}^N$ in (35) (obtained in the proof of Theorem 10) imply that

$$\lim_{N \rightarrow \infty} \pi_{v_-}^N(K_v\psi) = \pi_{v_-}(K_v\psi) = \pi_v(\psi) \quad \text{almost surely,} \quad \forall \psi \in \mathcal{B}_b(\mathbf{E}_v), \quad v \in \mathbb{T}_u^\partial.$$

Armed with the above, this proof is entirely analogous to that of [22, Theorem 10.6.1], we only need to tweak the definition of the $U_k^N(\varphi)$ s therein. In order to do this, we define a bijection $b = (b_1, b_2)$ from the one-dimensional index set $\{1, 2, \dots, N|\mathbb{T}_u|\}$ to the two-dimensional index set $\mathbb{T}_u \times \{1, 2, \dots, N\}$ that ‘preserves \mathbb{T}_u ’s structure’ in the sense that:

$$b^{-1}(v', n') \leq b^{-1}(v, n) \quad \forall v' \in \mathbb{T}_v, \quad n', n \leq N, \quad v \in \mathbb{T}_u.$$

That is, $(\mathbf{X}_{b_1(s)}^{b_2(s), N})_{s=1}^r$ contains all particles sub-indexed by descendants of $b_1(r)$ and no particles sub-indexed by ancestors of $b_1(r)$. This mapping imposes an order on the contributions of the local errors, allowing us to introduce the following sequence of martingale increments (defined with respect to an appropriate filtration):

$$U_k^N(\varphi) := \frac{1}{N^{1/2}} \left[\varphi_{b_1(k)}(\mathbf{X}_{b_1(k)}^{b_2(k), N}) - \pi_{b_1(k)}^N(K_{b_1(k)}\varphi_{b_1(k)}) \right];$$

and [22, Theorem 10.6.1]’s proof applies as is. \square

To obtain the aforementioned local error decomposition for $\gamma_u^N - \gamma_u$, it is easiest to first derive an analogous result for $\gamma_{\mathcal{C}_u}^N - \gamma_{\mathcal{C}_u}$. In particular, we express $\gamma_{\mathcal{C}_u}^N - \gamma_{\mathcal{C}_u}$ in terms of the local errors and an $\mathcal{O}(N^{-1})$ remainder term that is a linear combination of their products. When doing so, we find it convenient to introduce the kernel $\Upsilon_{v,u} : \mathbf{E}_{\mathcal{C}_v} \times \mathcal{E}_{\mathcal{C}_u} \rightarrow [0, \infty)$ mapping $\gamma_{\mathcal{C}_v}$ to $\gamma_{\mathcal{C}_u}$ (i.e. $\gamma_{\mathcal{C}_v}(\Upsilon_{v,u}\varphi) = \gamma_{\mathcal{C}_u}(\varphi)$ for all φ in $\mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u})$): $\Upsilon_{u,u}(\mathbf{x}_{\mathcal{C}_u}, d\mathbf{y}_{\mathcal{C}_u}) = \delta_{\mathbf{x}_{\mathcal{C}_u}}(d\mathbf{y}_{\mathcal{C}_u})$ and

$$(64) \quad \Upsilon_{r,u}(\mathbf{x}_{\mathcal{C}_r}, d\mathbf{y}_{\mathcal{C}_u}) = \delta_{\mathbf{x}_{\mathcal{C}_r}}(d\mathbf{y}_{\mathcal{C}_r}) w_{v_-}(\mathbf{y}_{\mathcal{C}_r}) K_v(\mathbf{y}_{\mathcal{C}_r}, d\mathbf{y}_r) \Gamma_{r,v}(\mathbf{y}_r, d\mathbf{y}_v) \gamma_{\mathcal{C}_u}^\vee(d\mathbf{x}_{\mathcal{C}_u}^\vee)$$

for all $r \neq u$ in \mathbb{T}_u^∂ , where v denotes the child of u that is an ancestor of r (with $v = r$ if r itself is a child of u) and $\Gamma_{r,v}$ is as in (15–17). It is straightforward to check that

$$(65) \quad \Upsilon_{r,u}\varphi = \Upsilon_{r,v}\Upsilon_{v,u}\varphi \quad \forall \varphi \in \mathcal{B}_b(\mathbf{E}_u), \quad r \in \mathbb{T}_v^\partial, \quad v \in \mathbb{T}_u^\partial, \quad u \in \mathbb{T}^\partial,$$

$$(66) \quad \Gamma_{r,u}\varphi = \gamma_{\mathcal{C}_v}^\vee(\Upsilon_{v,u}[w_{u_-}K_u\varphi]) \quad \forall \varphi \in \mathcal{B}_b(\mathbf{E}_u), \quad r \in \mathcal{C}_v, \quad v \in \mathbb{T}_u^\partial, \quad u \in \mathbb{T}^\partial,$$

two properties that will be of use in the following proofs.

LEMMA 11. If Assumptions 1–2 hold, u lies in \mathbb{T}^∂ , and φ in $\mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u})$, then

$$(67) \quad \gamma_{\mathcal{C}_u}^N(\varphi) - \gamma_{\mathcal{C}_u}(\varphi) = \frac{1}{N^{1/2}} \sum_{v \in \mathbb{T}_u^\partial} \sum_{r \in \mathcal{C}_v} \mathcal{Z}_r^N V_r^N(\gamma_{\mathcal{C}_v}^\vee(\Upsilon_{v,u}\varphi)) + R_u^N(\varphi),$$

where $(V_v^N(\Upsilon_{v,u}\varphi))_{v \in \mathbb{T}_u^\partial}$ is as in (62) and $R_u^N(\varphi)$ satisfies $N^{1/2}R_u^N(\varphi) \Rightarrow 0$ as $N \rightarrow \infty$.

PROOF. Fix any u in \mathbb{T}^∂ and φ in $\mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u})$. Note that, similarly to (75),

$$(68) \quad \gamma_{\mathcal{C}_u}^N - \gamma_{\mathcal{C}_u} = \prod_{v \in \mathcal{C}_u} [\gamma_v^N - \gamma_v + \gamma_v] - \gamma_{\mathcal{C}_u} = \sum_{\emptyset \neq A \subseteq \mathcal{C}_u} \Delta_A^N \times \gamma_{\mathcal{C}_u}^A,$$

where $\Delta_A^N := \prod_{v \in A} (\gamma_v^N - \gamma_v)$ and $\gamma_{\mathcal{C}_u}^A := \gamma_{\mathcal{C}_u \setminus A}$ for all subsets A of \mathcal{C}_u . The same arguments as those in (76), only with the inequality in (73) replaced by that in Theorem 7, show that

$$\mathbb{E} \left[\Delta_A^N (\gamma_{\mathcal{C}_u}^A(\varphi))^2 \right]^{\frac{1}{2}} \leq N^{-1} C \left\| \gamma_{\mathcal{C}_u}^A(\varphi) \right\| \leq N^{-1} C \mathcal{Z}_{\mathcal{C}_u}^A \|\varphi\|$$

if $|A| > 1$. Markov's inequality, then implies that

$$\mathbb{P} \left(\left\{ \left| \Delta_A^N (\gamma_{\mathcal{C}_u}^A(\varphi)) \right| \geq \varepsilon \right\} \right) \leq \varepsilon^{-2} \mathbb{E} \left[\Delta_A^N (\gamma_{\mathcal{C}_u}^A(\varphi))^2 \right] \leq (\varepsilon N)^{-2} C \mathcal{Z}_{\mathcal{C}_u}^A \|\varphi\| \quad \forall \varepsilon > 0;$$

from which it follows that

$$(69) \quad N^{1/2} \Delta_A^N (\gamma_{\mathcal{C}_u}^A(\varphi)) \Rightarrow 0 \quad \text{as } N \rightarrow \infty \quad \forall A \subseteq \mathcal{C}_u : |A| > 1.$$

We now focus on the case $|A| = 1$ where A is a singleton $\{v\}$. Comparing with (64),

$$\begin{aligned} \gamma_v^N (\gamma_{\mathcal{C}_u}^v(\varphi)) &= \mathcal{Z}_v^N \pi_v^N (\gamma_{\mathcal{C}_u}^v(\varphi)) = N^{-1/2} \mathcal{Z}_v^N V_v^N (\gamma_{\mathcal{C}_u}^v(\varphi)) + \mathcal{Z}_v^N \pi_{v-}^N (K_v \gamma_{\mathcal{C}_u}^v(\varphi)) \\ &= N^{-1/2} \mathcal{Z}_v^N V_v^N (\gamma_{\mathcal{C}_u}^v(\varphi)) + \gamma_{v-}^N (K_v \gamma_{\mathcal{C}_u}^v(\varphi)) \\ &= N^{-1/2} \mathcal{Z}_v^N V_v^N (\gamma_{\mathcal{C}_u}^v(\varphi)) + \gamma_{\mathcal{C}_v}^N (\Upsilon_{v,u} \varphi), \\ \gamma_v (\gamma_{\mathcal{C}_u}^v(\varphi)) &= \gamma_{v-} (K_v \gamma_{\mathcal{C}_u}^v(\varphi)) = \gamma_{\mathcal{C}_v} (\Upsilon_{v,u} \varphi), \end{aligned}$$

for all non-leaf children v of u . For these reasons,

$$(70) \quad (\Delta_v^N \times \gamma_{\mathcal{C}_u}^v)(\varphi) = \frac{\mathcal{Z}_v^N V_v^N (\gamma_{\mathcal{C}_u}^v(\varphi))}{N^{1/2}} \quad \forall v \in \mathcal{C}_u^\partial$$

$$(71) \quad \begin{aligned} (\Delta_v^N \times \gamma_{\mathcal{C}_u}^v)(\varphi) &= \gamma_v^N (\gamma_{\mathcal{C}_u}^v(\varphi)) - \gamma_v (\gamma_{\mathcal{C}_u}^v(\varphi)) \\ &= \mathcal{Z}_v^N V_v^N (\gamma_{\mathcal{C}_u}^v(\varphi)) N^{-1/2} + \gamma_{\mathcal{C}_v}^N (\Upsilon_{v,u} \varphi) - \gamma_{\mathcal{C}_v} (\Upsilon_{v,u} \varphi) \quad \forall v \in \mathcal{C}_u^\partial. \end{aligned}$$

If all of u 's children are leaves, then (67) follows by setting $R_u^N(\varphi) := \sum_{A \subseteq \mathcal{C}_u : |A| > 1} (\Delta_A^N \times \gamma_{\mathcal{C}_u}^A)(\varphi)$ and combining (68–70). For all other u , we argue the claim inductively: suppose that, for each non-leaf child v of u and ψ in $\mathcal{B}_b(\mathbf{E}_{\mathcal{C}_v})$, (67) (with ψ replacing φ) holds for some $R_v^N(\psi)$ satisfying $N^{1/2} R_v^N(\psi) \Rightarrow 0$ as $N \rightarrow \infty$. Given (65), we can re-write (71) as

$$\begin{aligned} N^{1/2} (\Delta_v^N \times \gamma_{\mathcal{C}_u}^v)(\varphi) &= \mathcal{Z}_v^N V_v^N (\gamma_{\mathcal{C}_u}^v(\varphi)) + \sum_{v' \in \mathbb{T}_v^\partial} \sum_{r \in \mathcal{C}_{v'}} \mathcal{Z}_r^N V_r^N (\gamma_{\mathcal{C}_{v'}}^r(\Upsilon_{v',u} \varphi)) + N^{1/2} R_v^N (\Upsilon_{v,u} \varphi) \quad \forall v \in \mathcal{C}_u^\partial. \end{aligned}$$

Because $\mathbb{T}_u^\partial = \{u\} \cup_{v \in \mathcal{C}_u} \mathbb{T}_v^\partial$, we obtain (67) by combining the above with (68–70) and setting

$$R_u^N(\varphi) := \sum_{v \in \mathcal{C}_u^\partial} R_v^N (\Upsilon_{v,u} \varphi) + \sum_{A \subseteq \mathcal{C}_u : |A| > 1} (\Delta_A^N \times \gamma_{\mathcal{C}_u}^A)(\varphi).$$

□

Theorem 13 follows from Lemmas 10–11 and the continuous mapping theorem:

PROOF OF THEOREM 13. If u is a leaf, then the CLT reduces to that for i.i.d. sequences. For any other u , fix φ in $\mathcal{B}_b(\mathbf{E}_u)$ and note that, by definition,

$$(72) \quad \begin{aligned} \gamma_u^N(\varphi) - \gamma_u(\varphi) &= \gamma_u^N(\varphi) - \gamma_{u_-}^N(K_u\varphi) + \gamma_{u_-}^N(K_u\varphi) - \gamma_{u_-}(K_u\varphi) \\ &= N^{-1/2} \mathcal{Z}_u^N V_u(\varphi) + \gamma_{\mathcal{C}_u}^N(w_{u_-} K_u\varphi) - \gamma_{\mathcal{C}_u}(w_{u_-} K_u\varphi). \end{aligned}$$

Using Lemma 11 and (66), we then find that

$$N^{1/2}(\gamma_u^N(\varphi) - \gamma_u(\varphi)) = f((\mathcal{Z}_v^N)_{v \in \mathbb{T}_u}, (V_v^N(\varphi_v))_{v \in \mathbb{T}_u}, N^{1/2} R_u^N(w_{u_-} K_u\varphi))$$

where φ_v denotes $\Gamma_{v,u}\varphi$ and

$$f(\mathbf{z}, \boldsymbol{\nu}, R) := \sum_{v \in \mathbb{T}_u} z_v \nu_v + R \quad \forall \mathbf{z}, \boldsymbol{\nu} \in \mathbb{R}^{|\mathbb{T}_u|}, R \in \mathbb{R}.$$

The L^p inequalities for $(\mathcal{Z}_v^N)_{v \in \mathbb{T}_u}$ together with Lemma 1 imply that $(\mathcal{Z}_v^N)_{v \in \mathbb{T}_u}$ converges almost surely to $(\mathcal{Z}_v)_{v \in \mathbb{T}_u}$ as N tends to infinity. Moreover, Lemmas 10–11 show that $N^{1/2} R_u^N(w_{u_-} K_u\varphi)$ converges weakly to zero and $(V_v^N(\varphi_v))_{v \in \mathbb{T}_u}$ converges to the family $(V_v(\varphi_v))_{v \in \mathbb{T}_u}$ of independent Gaussian random variables in (63). For these reasons, the continuous mapping theorem tells us that

$$N^{1/2}(\gamma_u^N(\varphi) - \gamma_u(\varphi)) \Rightarrow \sum_{v \in \mathbb{T}_u} \mathcal{Z}_v V_v(\varphi_v) = \sum_{v \in \mathbb{T}_u} \mathcal{Z}_v V_v(\Gamma_{v,u}\varphi) \quad \text{as } N \rightarrow \infty.$$

The right-hand side is a sum of zero mean independent Gaussian random variables and, so, also a zero mean Gaussian random with variance

$$\begin{aligned} \text{Var} \left(\sum_{v \in \mathbb{T}_u} \mathcal{Z}_v V_v(\Gamma_{v,u}\varphi) \right) &= \sum_{v \in \mathbb{T}_u} \mathcal{Z}_v^2 \text{Var}(V_v(\Gamma_{v,u}\varphi)) = \sum_{v \in \mathbb{T}_u} \mathcal{Z}_v^2 \pi_v([\Gamma_{v,u}\varphi - \pi_v(\Gamma_{v,u}\varphi)]^2) \\ &= \sum_{v \in \mathbb{T}_u} \pi_v([\mathcal{Z}_v \Gamma_{v,u}\varphi - \gamma_v(\Gamma_{v,u}\varphi)]^2) \\ &= \sum_{v \in \mathbb{T}_u} \pi_v([\mathcal{Z}_v \Gamma_{v,u}\varphi - \gamma_u(\varphi)]^2). \end{aligned}$$

□

APPENDIX G: PROOF OF THE THEOREM 7

Throughout this appendix, we use the notation described in Appendix B. The aim of the appendix is to prove the product L^2 inequality for DaC-SMC in Theorem 7. Similarly to the argument for the L^p inequality in Appendix C, we prove the above by obtaining a product L^2 inequality for the empirical distributions of the particles indexed by the tree's leaves and showing that each step of the algorithm preserves this inequality. To start, note that, if u and v are leaves, then [43, Lemma 1] implies that

$$(73) \quad \mathbb{E} [(\pi_u^N - \pi_u) \times (\pi_v^N - \pi_v)(\varphi)^2]^{\frac{1}{2}} \leq \frac{C \|\varphi\|}{N} \quad \forall N > 0, \varphi \in \mathcal{B}_b(\mathbf{E}_u \times \mathbf{E}_v).$$

Taking the products of $\pi_{u'}^N$ and $\pi_{v'}^N$ over all children u' and v' of nodes u and v to obtain $\pi_{\mathcal{C}_u}^N$ and $\pi_{\mathcal{C}_v}^N$ preserves (73):

LEMMA 12 (Product step). *If, in addition to Assumptions 1–2, (73) is satisfied for each child u' and v' (i.e. it holds with u' and v' replacing u and v therein) of two non-leaf nodes u and v that lying separate branches (i.e. $u \notin \mathbb{T}_v$ and $v \notin \mathbb{T}_u$), then*

$$(74) \quad \mathbb{E} [(\pi_{\mathcal{C}_u}^N - \pi_{\mathcal{C}_u}) \times (\pi_{\mathcal{C}_v}^N - \pi_{\mathcal{C}_v})(\varphi)^2]^{\frac{1}{2}} \leq \frac{C \|\varphi\|}{N} \quad \forall N > 0, \varphi \in \mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u} \times \mathbf{E}_{\mathcal{C}_v}).$$

PROOF. Fix any $N > 0$ and φ in $\mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u} \times \mathbf{E}_{\mathcal{C}_v})$. Using a multinomial expansion as in the proof of Theorem 10, we have that

$$\pi_{\mathcal{C}_u}^N = \prod_{r \in \mathcal{C}_u} [\pi_r^N - \pi_r + \pi_r] = \sum_{A \subseteq \mathcal{C}_u} \Delta_A^N \times \pi_{\mathcal{C}_u}^A, \quad \pi_{\mathcal{C}_v}^N = \cdots = \sum_{A \subseteq \mathcal{C}_v} \Delta_A^N \times \pi_{\mathcal{C}_v}^A,$$

where $\Delta_A^N := \prod_{r \in A} (\pi_r^N - \pi_r)$ and $\pi_A^B := \pi_{A \setminus B}$ for all $B \subseteq A \subseteq \mathbb{T}$. Hence,

$$(75) \quad \pi_{\mathcal{C}_u}^N - \pi_{\mathcal{C}_u} = \sum_{\emptyset \neq A \subseteq \mathcal{C}_u} \Delta_A^N \times \pi_{\mathcal{C}_u}^A, \quad \pi_{\mathcal{C}_v}^N - \pi_{\mathcal{C}_v} = \sum_{\emptyset \neq A \subseteq \mathcal{C}_v} \Delta_A^N \times \pi_{\mathcal{C}_v}^A,$$

and it follows that

$$(\pi_{\mathcal{C}_u}^N - \pi_{\mathcal{C}_u}) \times (\pi_{\mathcal{C}_v}^N - \pi_{\mathcal{C}_v})(\varphi) = \sum_{\emptyset \neq A \subseteq \mathcal{C}_u} \sum_{\emptyset \neq B \subseteq \mathcal{C}_v} \Delta_{A \cup B}^N (\pi_{\mathcal{C}_u \cup \mathcal{C}_v}^{A \cup B}(\varphi)).$$

Given Minkowski's inequality, we need only to show that

$$\mathbb{E} [\Delta_{A \cup B}^N (\pi_{\mathcal{C}_u \cup \mathcal{C}_v}^{A \cup B}(\varphi))^2]^{\frac{1}{2}} \leq \frac{C \|\varphi\|}{N} \quad \forall N > 0, \quad \emptyset \neq A \subseteq \mathcal{C}_u, \quad \emptyset \neq B \subseteq \mathcal{C}_v.$$

To do so, pick any nodes u', v' respectively in any two subsets A, B of $\mathcal{C}_u, \mathcal{C}_v$, and note that

$$\mathbb{E} [\Delta_{A \cup B}^N (\pi_{\mathcal{C}_u \cup \mathcal{C}_v}^{A \cup B}(\varphi))^2]^{\frac{1}{2}} = \mathbb{E} [(\pi_{u'}^N - \pi_{u'}) \times (\pi_{v'}^N - \pi_{v'}) (\psi)^2]^{\frac{1}{2}},$$

where $\psi := \Delta_{A \cup B \setminus \{u', v'\}}^N (\pi_{\mathcal{C}_u \cup \mathcal{C}_v}^{A \cup B}(\varphi))$. Because u and v lie on separate branches, and particles indexed by distinct nodes in \mathcal{C}_u (or \mathcal{C}_v) are independent, $(\mathbf{X}_r^{n, N})_{n=1}^N$ and $(\mathbf{X}_{r'}^{n, N})_{n=1}^N$ are independent collections of particles for all r and r' in $A \cup B$ satisfying $r \neq r'$. Hence, $\pi_{u'}^N$ and $\pi_{v'}^N$ are independent of the sigma algebra \mathcal{G} generated by $((\mathbf{X}_r^{n, N})_{n=1}^N)_{r \in A \cup B \setminus \{u', v'\}}$ and the desired inequality follows from our assumption that (73) holds for u' and v' :

$$(76) \quad \begin{aligned} \mathbb{E} [\Delta_{A \cup B}^N (\pi_{\mathcal{C}_u \cup \mathcal{C}_v}^{A \cup B}(\varphi))^2]^{\frac{1}{2}} &= \mathbb{E} [\mathbb{E} [(\pi_{u'}^N - \pi_{u'}) \times (\pi_{v'}^N - \pi_{v'}) (\psi)^2 | \mathcal{G}]]^{\frac{1}{2}} \leq \frac{C \mathbb{E} [|\psi|^2]^{\frac{1}{2}}}{N} \\ &\leq \frac{C 2^{(|A|+|B|-2)/2} \|\pi_{\mathcal{C}_u \cup \mathcal{C}_v}^{A \cup B}(\varphi)\|}{N} \leq \frac{C 2^{(|A|+|B|-2)/2} \|\varphi\|}{N}. \end{aligned}$$

□

The correction step also preserves these inequalities:

LEMMA 13 (Correction step). *If, in addition to Assumptions 1–2, (74) is satisfied for some non-leaf nodes u and v lying in separate branches (i.e. $u \notin \mathbb{T}_v$ and $v \notin \mathbb{T}_u$), then*

$$(77) \quad \mathbb{E} [(\pi_{u_-}^N - \pi_{u_-}) \times (\pi_{v_-}^N - \pi_{v_-})(\varphi)^2]^{\frac{1}{2}} \leq \frac{C \|\varphi\|}{N} \quad \forall N > 0, \quad \varphi \in \mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u} \times \mathbf{E}_{\mathcal{C}_v}).$$

PROOF. Fix any $N > 0$ and φ in $\mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u} \times \mathbf{E}_{\mathcal{C}_v})$, and let

$$\varrho_r^N(d\mathbf{x}_{\mathcal{C}_r}) := w_{r_-}(\mathbf{x}_{\mathcal{C}_r}) \pi_{\mathcal{C}_r}^N(d\mathbf{x}_{\mathcal{C}_r}), \quad \varrho_r(d\mathbf{x}_{\mathcal{C}_r}) := w_{r_-}(\mathbf{x}_{\mathcal{C}_r}) \pi_{\mathcal{C}_r}(d\mathbf{x}_{\mathcal{C}_r}), \quad \forall r = v, u.$$

Because $\pi_{u_-}^N = \varrho_u^N / \pi_{\mathcal{C}_u}^N(w_{u_-})$ and $\pi_{u_-} = \varrho_u / \pi_{\mathcal{C}_u}(w_{u_-})$,

$$\pi_{u_-}^N - \pi_{u_-} = \pi_{u_-}^N - \frac{\varrho_u}{\pi_{\mathcal{C}_u}(w_{u_-})} = \frac{\varrho_u^N - \varrho_u - (\pi_{\mathcal{C}_u}^N(w_{u_-}) - \pi_{\mathcal{C}_u}(w_{u_-})) \pi_{u_-}^N}{\pi_{\mathcal{C}_u}(w_{u_-})},$$

and similarly for v . Hence,

$$\begin{aligned}
(78) \quad & \pi_{\mathcal{C}_u}(w_{u_-})\pi_{\mathcal{C}_v}(w_{v_-})(\pi_{u_-}^N - \pi_{u_-}) \times (\pi_{v_-}^N - \pi_{v_-})(\varphi) \\
& = (\varrho_u^N - \varrho_u) \times (\varrho_v^N - \varrho_v)(\varphi) \\
& \quad - (\pi_{\mathcal{C}_v}^N(w_{v_-}) - \pi_{\mathcal{C}_v}(w_{v_-}))(\varrho_u^N - \varrho_u)(\pi_{v_-}^N(\varphi)) \\
& \quad - (\pi_{\mathcal{C}_u}^N(w_{u_-}) - \pi_{\mathcal{C}_u}(w_{u_-}))(\varrho_v^N - \varrho_v)(\pi_{u_-}^N(\varphi)) \\
& \quad + (\pi_{\mathcal{C}_u}^N(w_{u_-}) - \pi_{\mathcal{C}_u}(w_{u_-}))(\pi_{\mathcal{C}_v}^N(w_{v_-}) - \pi_{\mathcal{C}_v}(w_{v_-}))(\pi_{u_-}^N \times \pi_{v_-}^N)(\varphi).
\end{aligned}$$

Because $\varrho_u^N(\varphi) = \pi_{\mathcal{C}_u}^N(w_{u_-}\varphi)$, $\varrho_u(\varphi) = \pi_{\mathcal{C}_u}(w_{u_-}\varphi)$, and similarly for v , (74) implies that

$$\begin{aligned}
(79) \quad & \mathbb{E} [|(\varrho_u^N - \varrho_u) \times (\varrho_v^N - \varrho_v)|^2]^{\frac{1}{2}} = \mathbb{E} [|(\pi_{\mathcal{C}_u}^N - \pi_{\mathcal{C}_u}) \times (\pi_{\mathcal{C}_v}^N - \pi_{\mathcal{C}_v})(w_{u_-}w_{v_-}\varphi)|^2]^{\frac{1}{2}} \\
& \leq \frac{C \|w_{u_-}w_{v_-}\varphi\|}{N} \leq \frac{C \|w_{u_-}\| \|w_{v_-}\| \|\varphi\|}{N}.
\end{aligned}$$

Next, because u and v lie in different branches, $\mathcal{F}_{\mathcal{C}_u}^N$ and $\mathcal{F}_{\mathcal{C}_v}^N$ are independent. By definition, $\pi_{\mathcal{C}_u}^N(w_{u_-})$ is $\mathcal{F}_{\mathcal{C}_u}^N$ -measurable and $\pi_{\mathcal{C}_v}^N(w_{v_-})$ is $\mathcal{F}_{\mathcal{C}_v}^N$ -measurable. Hence, the L^p inequalities for $\pi_{\mathcal{C}_u}^N$ and $\pi_{\mathcal{C}_v}^N$ in (27) (obtained in the proof of Theorem 10) imply that

$$\begin{aligned}
(80) \quad & \mathbb{E} [(\pi_{\mathcal{C}_v}^N(w_{v_-}) - \pi_{\mathcal{C}_v}(w_{v_-}))^2 (\varrho_u^N - \varrho_u)(\pi_{v_-}^N(\varphi))^2]^{\frac{1}{2}} \\
& = \mathbb{E} [(\pi_{\mathcal{C}_v}^N(w_{v_-}) - \pi_{\mathcal{C}_v}(w_{v_-}))^2 \mathbb{E} [(\varrho_u^N - \varrho_u)(\pi_{v_-}^N(\varphi))^2 | \mathcal{F}_v^N]]^{\frac{1}{2}} \\
& \leq \mathbb{E} \left[(\pi_{\mathcal{C}_v}^N(w_{v_-}) - \pi_{\mathcal{C}_v}(w_{v_-}))^2 \frac{C^2 \|w_{u_-}\pi_{v_-}^N(\varphi)\|^2}{N} \right]^{\frac{1}{2}} \\
& \leq \frac{C \|w_{u_-}\| \|\varphi\|}{N^{1/2}} \mathbb{E} [(\pi_{\mathcal{C}_v}^N(w_{v_-}) - \pi_{\mathcal{C}_v}(w_{v_-}))^2]^{\frac{1}{2}} \\
& \leq \frac{C \|w_{u_-}\| \|w_{v_-}\| \|\varphi\|}{N}.
\end{aligned}$$

Similarly, we find that

$$(81) \quad \mathbb{E} [(\pi_{\mathcal{C}_u}^N(w_{u_-}) - \pi_{\mathcal{C}_u}(w_{u_-}))^2 (\varrho_v^N - \varrho_v)(\pi_{u_-}^N(\varphi))^2]^{\frac{1}{2}} \leq \frac{C \|w_{v_-}\| \|w_{u_-}\| \|\varphi\|}{N},$$

$$\begin{aligned}
(82) \quad & \mathbb{E} [(\pi_{\mathcal{C}_u}^N(w_{u_-}) - \pi_{\mathcal{C}_u}(w_{u_-}))^2 (\pi_{\mathcal{C}_v}^N(w_{v_-}) - \pi_{\mathcal{C}_v}(w_{v_-}))^2 (\pi_{u_-}^N \times \pi_{v_-}^N)(\varphi)^2]^{\frac{1}{2}} \\
& \leq \mathbb{E} [(\pi_{\mathcal{C}_u}^N(w_{u_-}) - \pi_{\mathcal{C}_u}(w_{u_-}))^2]^{\frac{1}{2}} \mathbb{E} [(\pi_{\mathcal{C}_v}^N(w_{v_-}) - \pi_{\mathcal{C}_v}(w_{v_-}))^2]^{\frac{1}{2}} \|\varphi\| \\
& \leq \frac{C \|w_{u_-}\| \|w_{v_-}\| \|\varphi\|}{N}.
\end{aligned}$$

Combining (78–82) and applying Minkowski's inequality then completes the proof. \square

For the resampling step, recall that $\epsilon_u^N := N^{-1} \sum_{n=1}^N \delta_{\mathbf{X}_{u_-}^{n,N}}$ denotes the empirical distribution of the resampled particles. We then have the following:

LEMMA 14 (Resampling step). *If, in addition to Assumptions 1–2, (77) is satisfied for some non-leaf nodes u and v lying in separate branches (i.e. $u \notin \mathbb{T}_v$ and $v \notin \mathbb{T}_u$), then*

$$(83) \quad \mathbb{E} [(\epsilon_u^N - \pi_{u_-}) \times (\epsilon_v^N - \pi_{v_-})(\varphi)^2]^{\frac{1}{2}} \leq \frac{C \|\varphi\|}{N} \quad \forall N > 0, \quad \varphi \in \mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u} \times \mathbf{E}_v).$$

PROOF. Fix any $N > 0$ and φ in $\mathcal{B}_b(\mathbf{E}_{\mathcal{C}_u} \times \mathbf{E}_v)$. Note that

$$S_N := N^2(\epsilon_u^N - \pi_{u_-}) \times (\epsilon_v^N - \pi_{v_-})(\varphi) = \sum_{i=1}^N \sum_{j=1}^N \psi(\mathbf{X}_{u_-}^{i,N}, \mathbf{X}_{v_-}^{j,N}),$$

where $\psi := \varphi - \pi_{v_-}(\varphi) - \pi_{u_-}(\varphi) + (\pi_{u_-} \times \pi_{v_-})(\varphi)$. In other words,

$$S_N = \sum_{i=1}^N \sum_{j=1}^N Y_{i,j}^N \quad \text{where} \quad Y_{i,j}^N := \psi(\mathbf{X}_{u_-}^{i,N}, \mathbf{X}_{v_-}^{j,N}) \quad \forall i, j \leq N.$$

Conditioned on $\mathcal{G} := \mathcal{F}_{\mathcal{C}_u}^N \vee \mathcal{F}_{\mathcal{C}_v}^N$, the resampled particles $\mathbf{X}_{u_-}^{1,N}, \dots, \mathbf{X}_{u_-}^{N,N}, \mathbf{X}_{v_-}^{1,N}, \dots, \mathbf{X}_{v_-}^{N,N}$ are independent with

$$\mathbf{X}_{u_-}^{n,N} \sim \pi_{u_-}^N, \quad \mathbf{X}_{v_-}^{n,N} \sim \pi_{v_-}^N, \quad \forall n \leq N.$$

We proceed by cases:

- ($i_1 \neq i_2$ and $j_1 \neq j_2$). With probability one,

$$\begin{aligned} \mathbb{E}[Y_{i_1, j_1}^N Y_{i_2, j_2}^N | \mathcal{G}] &= \mathbb{E}[\psi(\mathbf{X}_{u_-}^{i_1, N}, \mathbf{X}_{v_-}^{j_1, N}) \psi(\mathbf{X}_{u_-}^{i_2, N}, \mathbf{X}_{v_-}^{j_2, N}) | \mathcal{G}] \\ &= \mathbb{E}[\psi(\mathbf{X}_{u_-}^{i_1, N}, \mathbf{X}_{v_-}^{j_1, N}) | \mathcal{G}] \mathbb{E}[\psi(\mathbf{X}_{u_-}^{i_2, N}, \mathbf{X}_{v_-}^{j_2, N}) | \mathcal{G}] \\ &= \pi_{u_-}^N \times \pi_{v_-}^N (\psi)^2. \end{aligned}$$

But ψ 's definition implies that $\pi_{u_-}(\psi) = \pi_{v_-}(\psi) = 0$ and $\|\psi\| \leq 4\|\varphi\|$. Hence, by (77),

$$(84) \quad \mathbb{E}[Y_{i_1, j_1}^N Y_{i_2, j_2}^N] \leq \frac{C^2 \|\psi\|^2}{N^2} \leq \frac{16C^2 \|\varphi\|^2}{N^2} \quad \text{if } i_1 \neq i_2, \quad j_1 \neq j_2.$$

- ($i_1 \neq i_2$ and $j_1 = j_2$). Let $\mathcal{G}_{j_1} := \mathcal{G} \vee \sigma(\mathbf{X}_{v_-}^{j_1, N})$ so that

$$\mathbb{E}[Y_{i_1, j_1}^N Y_{i_2, j_2}^N | \mathcal{G}_{j_1}] = \pi_{u_-}^N(\psi) (\mathbf{X}_{v_-}^{j_1, N})^2 \quad \text{almost surely.}$$

Because $\pi_{u_-}(\psi) = 0$ by definition, independence and the L^p inequality for $\pi_{u_-}^N$ in (35) (obtained in the proof of Theorem 10) then imply that

$$(85) \quad \mathbb{E}[Y_{i_1, j_1}^N Y_{i_2, j_2}^N] \leq \frac{C^2 \|\psi\|^2}{N} \leq \frac{16C^2 \|\varphi\|^2}{N} \quad \text{if } i_1 \neq i_2, \quad j_1 = j_2.$$

- ($i_1 = i_2$ and $j_1 \neq j_2$). Similarly,

$$(86) \quad \mathbb{E}[Y_{i_1, j_1}^N Y_{i_2, j_2}^N] \leq \frac{C^2 \|\psi\|^2}{N} \leq \frac{16C^2 \|\varphi\|^2}{N} \quad \text{if } i_1 = i_2, \quad j_1 \neq j_2.$$

- ($i_1 = i_2$ and $j_1 = j_2$). Clearly,

$$(87) \quad \mathbb{E}[Y_{i_1, j_1}^N Y_{i_2, j_2}^N] \leq \|\psi\|^2 \leq 16\|\varphi\|^2 \quad \text{if } i_1 = i_2, \quad j_1 = j_2.$$

In the sum

$$\mathbb{E}[S_N^2] = \sum_{i_1=1}^N \sum_{j_1=1}^N \sum_{i_2=1}^N \sum_{j_2=1}^N \mathbb{E}[Y_{i_1, j_1}^N Y_{i_2, j_2}^N],$$

there are $N^2(N-1)^2$ terms of the type in (84), $2N^2(N-1)$ of the type in (85,86), and N^2 of the type in (87). Hence, (84–87) imply that $\mathbb{E}[S_N^2] \leq N^2 C \|\varphi\|^2$ and the result follows. \square

Similarly for the mutation step:

LEMMA 15 (Mutation step). *If, in addition to Assumptions 1–2, u and v lie in separate branches (i.e. $u \notin \mathbb{T}_v$ and $v \notin \mathbb{T}_u$) and either (a) neither u or v is a leaf and (83) holds or (b) u or v is a leaf, then (73) also holds for u and v .*

PROOF. For the sake of brevity, we focus on the case that neither u or v is a leaf and skip the analogous, albeit simpler, case that one of the two is a leaf. Fix any $N > 0$ and φ in $\mathcal{B}_b(\mathbf{E}_u \times \mathbf{E}_v)$, and note that

$$\begin{aligned} (\pi_u^N - \pi_u) \times (\pi_v^N - \pi_v)(\varphi) &= (\pi_u^N - \nu_u^N + \nu_u^N - \pi_u) \times (\pi_v^N - \nu_v^N + \nu_v^N - \pi_v)(\varphi) \\ (88) \quad &= (\pi_u^N - \nu_u^N) \times (\pi_v^N - \nu_v^N)(\varphi) + (\pi_u^N - \nu_u^N) \times (\nu_v^N - \pi_v)(\varphi) \\ &\quad + (\nu_u^N - \pi_u) \times (\pi_v^N - \nu_v^N)(\varphi) + (\nu_u^N - \pi_u) \times (\nu_v^N - \pi_v)(\varphi), \end{aligned}$$

where $\nu_u^N := \epsilon_u^N \otimes K_u$ and $\nu_v^N := \epsilon_v^N \otimes K_v$. Because K_u and K_v are Markov kernels, the inequality (83) bounds the fourth term in (88):

$$\begin{aligned} (89) \quad \mathbb{E} [(\nu_u^N - \pi_u) \times (\nu_v^N - \pi_v)(\varphi)^2]^{\frac{1}{2}} &= \mathbb{E} [(\epsilon_u^N - \pi_{u_-}) \times (\epsilon_v^N - \pi_{v_-})(K_u K_v \varphi)^2]^{\frac{1}{2}} \\ &\leq \frac{C \|K_u K_v \varphi\|}{N} \leq \frac{C \|\varphi\|}{N}. \end{aligned}$$

To control the other three terms in (88), we emulate the approach in Lemma 14's proof. For the first term, we set

$$S_N := N^2 (\pi_u^N - \nu_u^N) \times (\pi_v^N - \nu_v^N)(\varphi) = \sum_{i=1}^N \sum_{j=1}^N Y_{i,j}^N,$$

where $Y_{i,j}^N := \psi(\mathbf{X}_u^{i,N}, \mathbf{X}_v^{j,N})$ with $\psi := \varphi - K_v \varphi - K_u \varphi + K_u K_v \varphi$. Suppose that $i_1 \neq i_2$ and set $\mathcal{G} := \sigma(\mathbf{X}_v^{j_1,N}, \mathbf{X}_v^{j_2,N}) \vee \mathcal{F}_{u_-}^N$. Because u and v lie in different branches, $\mathbf{X}_v^{j_1,N}$ and $\mathbf{X}_v^{j_2,N}$ are independent of $\mathcal{F}_{u_-}^N$. Moreover, conditioned on $\mathcal{F}_{u_-}^N$, $X_u^{i_1,N}$ and $X_u^{i_2,N}$ are independent with respective laws $K_u(\mathbf{X}_{u_-}^{i_1,N}, \cdot)$ and $K_u(\mathbf{X}_{u_-}^{i_2,N}, \cdot)$. For these reasons,

$$\begin{aligned} \mathbb{E} [Y_{i_1,j_1}^N Y_{i_2,j_2}^N | \mathcal{G}] &= \mathbb{E} [\psi(\mathbf{X}_u^{i_1,N}, \mathbf{X}_v^{j_1,N}) \psi(\mathbf{X}_u^{i_2,N}, \mathbf{X}_v^{j_2,N}) | \mathcal{G}] \\ &= \mathbb{E} [\psi(\mathbf{X}_u^{i_1,N}, \mathbf{X}_v^{j_1,N}) | \mathcal{G}] \mathbb{E} [\psi(\mathbf{X}_u^{i_2,N}, \mathbf{X}_v^{j_2,N}) | \mathcal{G}] \\ &= (K_u \psi)(\mathbf{X}_{u_-}^{i_1,N}, \mathbf{X}_v^{j_1,N}) (K_u \psi)(\mathbf{X}_{u_-}^{i_2,N}, \mathbf{X}_v^{j_2,N}) \quad \text{almost surely.} \end{aligned}$$

But ψ 's definition implies that $K_u \psi = 0$ and we have that $\mathbb{E} [Y_{i_1,j_1}^N Y_{i_2,j_2}^N] = 0$ if $i_1 \neq i_2$. The same argument shows that this is also the case if $j_1 \neq j_2$, and it follows that

$$\mathbb{E} [S_N^2] = \sum_{i_1=1}^N \sum_{j_1=1}^N \sum_{i_2=1}^N \sum_{j_2=1}^N \mathbb{E} [Y_{i_1,j_1}^N Y_{i_2,j_2}^N] = \sum_{i=1}^N \sum_{j=1}^N \mathbb{E} [(Y_{i,j}^N)^2] \leq N^2 \|\psi\|^2.$$

Because K_u and K_v are Markov kernels, $\|\psi\| \leq 4\|\varphi\|$ and, so,

$$(90) \quad \mathbb{E} [(\pi_u^N - \nu_u^N) \times (\pi_v^N - \nu_v^N)(\varphi)^2]^{\frac{1}{2}} = \frac{\mathbb{E} [S_N^2]^{\frac{1}{2}}}{N^2} \leq \frac{4\|\varphi\|}{N}.$$

For the remaining two terms in (88), set

$$S_N := N^2 (\pi_u^N - \nu_u^N) \times (\nu_v^N - \pi_v)(\varphi) = \sum_{i=1}^N \sum_{j=1}^N Y_{i,j}^N,$$

where $Y_{i,j}^N := \psi(\mathbf{X}_u^{i,N}, \mathbf{X}_v^{j,N})$ with $\psi := K_v \varphi - \pi_{v_-}(K_v \varphi) - K_u K_v \varphi + K_u \pi_{v_-}(K_v \varphi)$. We proceed by cases:

- ($i_1 \neq i_2$). Because u and v lie in separate branches, conditional on $\mathcal{G} := \mathcal{F}_{u_-}^N \vee \mathcal{F}_{v_-}^N$, $X_u^{i_1, N}$ and $X_u^{i_2, N}$ are independent with respective laws $K_u(X_{u_-}^{i_1, N}, \cdot)$ and $K_u(X_{u_-}^{i_2, N}, \cdot)$. For this reason,

$$\begin{aligned} \mathbb{E} [Y_{i_1, j_1}^N Y_{i_2, j_2}^N | \mathcal{G}] &= \mathbb{E} [\psi(\mathbf{X}_u^{i_1, N}, \mathbf{X}_{v_-}^{j_1, N}) \psi(\mathbf{X}_u^{i_2, N}, \mathbf{X}_{v_-}^{j_2, N}) | \mathcal{G}] \\ &= \mathbb{E} [\psi(\mathbf{X}_u^{i_1, N}, \mathbf{X}_{v_-}^{j_1, N}) | \mathcal{G}] \mathbb{E} [\psi(\mathbf{X}_u^{i_2, N}, \mathbf{X}_{v_-}^{j_2, N}) | \mathcal{G}] \\ &= (K_u \psi)(\mathbf{X}_{u_-}^{i_1, N}, \mathbf{X}_{v_-}^{j_1, N}) (K_u \psi)(\mathbf{X}_{u_-}^{i_2, N}, \mathbf{X}_{v_-}^{j_2, N}) \quad \text{almost surely.} \end{aligned}$$

But ψ 's definition implies that $K_u \psi = 0$ and so

$$(91) \quad \mathbb{E} [Y_{i_1, j_1}^N Y_{i_2, j_2}^N] = 0 \quad \text{if } i_1 \neq i_2.$$

- ($i_1 = i_2$ and $j_1 \neq j_2$). If $\mathcal{G} := \sigma(\mathbf{X}_u^{i_1, N}) \vee \mathcal{F}_{v_-}^N$, then

$$\begin{aligned} \mathbb{E} [Y_{i_1, j_1}^N Y_{i_2, j_2}^N | \mathcal{G}] &= \mathbb{E} [\psi(\mathbf{X}_u^{i_1, N}, \mathbf{X}_{v_-}^{j_1, N}) \psi(\mathbf{X}_u^{i_1, N}, \mathbf{X}_{v_-}^{j_2, N}) | \mathcal{G}] \\ &= \mathbb{E} [\psi(\mathbf{X}_u^{i_1, N}, \mathbf{X}_{v_-}^{j_1, N}) | \mathcal{G}] \mathbb{E} [\psi(\mathbf{X}_u^{i_1, N}, \mathbf{X}_{v_-}^{j_2, N}) | \mathcal{G}] \\ &= \pi_{v_-}^N(\psi)(\mathbf{X}_u^{i_1, N})^2 \quad \text{almost surely.} \end{aligned}$$

But, by definition, $\pi_{v_-}(\psi) = 0$ and the L^p inequality for $\pi_{u_-}^N$ in Theorem 5 shows that

$$(92) \quad \mathbb{E} [Y_{i_1, j_1}^N Y_{i_2, j_2}^N] \leq \frac{16C^2 \|\varphi\|^2}{N} \quad \text{if } i_1 = i_2, \quad j_1 \neq j_2.$$

- ($i_1 = i_2$ and $j_1 = j_2$). Clearly,

$$(93) \quad \mathbb{E} [Y_{i_1, j_1}^N Y_{i_2, j_2}^N] \leq \|\psi\|^2 \leq 16 \|\varphi\|^2 \quad \text{if } i_1 = i_2, \quad j_1 = j_2.$$

In the sum

$$\mathbb{E} [S_N^2] = \sum_{i_1=1}^N \sum_{j_1=1}^N \sum_{i_2=1}^N \sum_{j_2=1}^N \mathbb{E} [Y_{i_1, j_1}^N Y_{i_2, j_2}^N],$$

there are $N^2(N-1)$ terms of the type in (92) and N^2 of that in (93). Hence, (91–93) imply that $\mathbb{E} [S_N^2] \leq N^2 C^2 \|\varphi\|^2$ and we find that

$$\mathbb{E} [(\pi_u^N - \nu_u^N) \times (\nu_v^N - \pi_v)(\varphi)^2]^{\frac{1}{2}} \leq \frac{C \|\varphi\|}{N}.$$

For the same reasons, we also have that

$$\mathbb{E} [(\nu_u^N - \pi_u) \times (\pi_v^N - \nu_v^N)(\varphi)^2]^{\frac{1}{2}} \leq \frac{C \|\varphi\|}{N}.$$

Combining (88–90), the above, and Minkowski's inequality, completes the proof. \square

PROOF OF THEOREM 7. Repeatedly applying Lemmas 13–15 starting from (73), we find that, for any u and v lying in separate branches (i.e. $u \notin \mathbb{T}_v$ and $v \notin \mathbb{T}_u$),

$$(94) \quad \mathbb{E} [(\pi_u^N - \pi_u) \times (\pi_v^N - \pi_v)(\varphi)^2]^{\frac{1}{2}} \leq \frac{C \|\varphi\|}{N} \quad \forall N > 0, \quad \varphi \in \mathcal{B}_b(\mathbf{E}_u \times \mathbf{E}_v).$$

The remainder of this proof is quite similar to that of Lemma 13 and we only sketch it:

$$\begin{aligned} &(\gamma_u^N - \gamma_u) \times (\gamma_v^N - \gamma_v)(\varphi) \\ &= (\gamma_u^N - \mathcal{Z}_u \pi_u^N + \mathcal{Z}_u \pi_u^N - \gamma_u) \times (\gamma_v^N - \mathcal{Z}_v \pi_v^N + \mathcal{Z}_v \pi_v^N - \gamma_v)(\varphi) \end{aligned}$$

$$\begin{aligned}
&= (\mathcal{Z}_u^N - \mathcal{Z}_u)(\mathcal{Z}_v^N - \mathcal{Z}_v)\pi_u^N \times \pi_v^N(\varphi) \\
&\quad - (\mathcal{Z}_u^N - \mathcal{Z}_u)\mathcal{Z}_v(\pi_v^N - \pi_v)(\pi_u^N(\varphi)) \\
&\quad - (\mathcal{Z}_v^N - \mathcal{Z}_v)\mathcal{Z}_u(\pi_u^N - \pi_u)(\pi_v^N(\varphi)) \\
&\quad + \mathcal{Z}_u\mathcal{Z}_v(\pi_u^N - \pi_u) \times (\pi_v^N - \pi_v)(\varphi) \quad \forall N > 0, \quad \varphi \in \mathcal{B}_b(\mathbf{E}_u \times \mathbf{E}_v).
\end{aligned}$$

To control the L^2 norm of the fourth term, use (94). For the first three, use independence, the L^p inequalities for π_u^N and π_v^N in Theorem 10, and those for \mathcal{Z}_u^N and \mathcal{Z}_v^N in (43). Then, apply Minkowski's inequality. \square

APPENDIX H: PROOF OF THEOREM 8

Throughout this appendix, we use the notation described in Appendix B. The aim of the appendix is to prove (19–21) in Theorem 8. Before we do so, we take a moment to verify that all of the terms in the theorem's statement are well-defined. First off, note that $\omega_u(K_u)$ therein is well-defined, and can be chosen to be positive everywhere, because, by \mathcal{S} 's definition and Assumption 1,

$$\rho_u \sim \gamma_{u_-} \otimes K_u \sim \gamma_{\mathcal{C}_u} \otimes K_u \sim \rho_{\mathcal{C}_u} \otimes K_u \quad \forall (\gamma_{u_-}, K_u) \in \mathcal{S}.$$

The above further implies that $C\sqrt{K_u\omega_u^2}\rho_{\mathcal{C}_u} \sim \rho_{\mathcal{C}_u}$. Hence, if (γ_{u_-}, K_u) belongs to \mathcal{S} ,

$$\begin{aligned}
&C\sqrt{K_u\omega_u^2}\rho_{\mathcal{C}_u} \otimes K_u \sim \rho_{\mathcal{C}_u} \otimes K_u \sim \gamma_{\mathcal{C}_u} \otimes K_u \otimes \gamma_{u_-} \otimes K_u \sim \rho_u, \\
&\gamma_{u_-} \otimes K_u \sim \rho_u \sim \mu_u \sim \rho_u^{u_-} \otimes M_u \Rightarrow \gamma_{u_-} \sim \mu_u^{u_-} \Rightarrow \gamma_{u_-} \otimes M_u \sim \mu_u^{u_-} \otimes M_u = \mu_u \sim \rho_u.
\end{aligned}$$

In other words, both $(C\sqrt{K_u\omega_u^2}\rho_{\mathcal{C}_u}, K_u)$ and (γ_{u_-}, M_u) belong to \mathcal{S} whenever (γ_{u_-}, K_u) does and all terms in the statement are well-defined.

To prove (19–21), we need only argue (19–20) as (21) follows immediately:

$$f(\gamma_{u_-}, K_u) \geq f(\gamma_{u_-}, M_u) \geq f(\sqrt{M_u\omega_u(M_u)^2}\rho_{\mathcal{C}_u}, M_u) \quad \forall (\gamma_{u_-}, K_u) \in \mathcal{S},$$

and, given that $\rho_u = Z_u\mu_u = Z_u\mu_u^{u_-} \otimes M_u = \rho_u^{u_-} \otimes M_u$,

$$\begin{aligned}
\omega_u(K_u) &= \frac{d\rho_u}{d\rho_{\mathcal{C}_u} \otimes K_u} = \frac{d\rho_u^{u_-} \otimes M_u}{d\rho_{\mathcal{C}_u} \otimes K_u} = \frac{d\rho_u^{u_-}}{d\rho_{\mathcal{C}_u}} \frac{dM_u}{dK_u} \\
(95) \quad \Rightarrow K_u\omega_u(K_u)^2 &= \left(\frac{d\rho_u^{u_-}}{d\rho_{\mathcal{C}_u}}\right)^2 K_u \left[\frac{dM_u}{dK_u}\right]^2 \Rightarrow \sqrt{K_u\omega_u^2(M_u)}\rho_{\mathcal{C}_u} = \frac{d\rho_u^{u_-}}{d\rho_{\mathcal{C}_u}}\rho_{\mathcal{C}_u} = \rho_u^{u_-}.
\end{aligned}$$

To argue (19–20), we express Z_u^N 's variance as

$$(96) \quad \text{Var}(Z_u^N) = A + \mathbb{E}[Bg(X_{\mathcal{C}_u}^N, \gamma_{u_-}, K_u)]$$

where A is a constant and B is a random variable, neither dependent on (γ_{u_-}, K_u) , $\mathbf{X}_{\mathcal{C}_u}^N := (\mathbf{X}_{\mathcal{C}_u}^{n,N})_{\mathbf{n} \in [N]^{c_u}}$, and g is a real-valued function on $\mathbf{E}_{\mathcal{C}_u}^{N^{c_u}} \times \mathcal{S}$. We will then show that

$$(97) \quad g(\mathbf{x}_{\mathcal{C}_u}^N, \gamma_{u_-}, K_u) \geq g(\mathbf{x}_{\mathcal{C}_u}^N, \sqrt{K_u\omega_u(K_u)^2}\rho_{\mathcal{C}_u}, K_u) \quad \forall \mathbf{x}_{\mathcal{C}_u}^N \in \mathbf{E}_{\mathcal{C}_u}^{N^{c_u}},$$

$$(98) \quad g(\mathbf{x}_{\mathcal{C}_u}^N, \gamma_{u_-}, K_u) \geq g(\mathbf{x}_{\mathcal{C}_u}^N, \gamma_{u_-}, M_u) \quad \forall \mathbf{x}_{\mathcal{C}_u}^N \in \mathbf{E}_{\mathcal{C}_u}^{N^{c_u}},$$

for all (γ_{u_-}, K_u) in \mathcal{S} . Given (96), (19–20) then follow from the above.

To carry out these steps, we'll need several identities that we collect here:

$$\gamma_{u_-}([K_u w_u]\varphi) = \gamma_{u_-}(K_u[w_u\varphi]) = \gamma_u(w_u\varphi) = \rho_u(\varphi) = \rho_u^{u_-}(\varphi) \quad \forall \varphi \in \mathcal{B}(\mathbf{E}_{\mathcal{C}_u}),$$

where ρ_u^{u-} denotes the $E_{\mathcal{C}_u}$ -marginal of ρ_u , from which it follows that that

$$\begin{aligned}
 K_u w_u &= \frac{d\rho_u^{u-}}{d\gamma_{u-}} \\
 (99) \quad \Rightarrow w_{u-} K_u w_u &= \frac{d\gamma_{u-}}{d\gamma_{\mathcal{C}_u}} \frac{d\rho_u^{u-}}{d\gamma_{u-}} = \frac{d\rho_u^{u-}}{d\gamma_{\mathcal{C}_u}}, \\
 w_u &= \frac{d\rho_u}{d\gamma_u} = \frac{d\rho_{\mathcal{C}_u} \otimes K_u}{d\gamma_u} \frac{d\rho_u}{d\rho_{\mathcal{C}_u} \otimes K_u} = \frac{d\rho_{\mathcal{C}_u} \otimes K_u}{d\gamma_{u-} \otimes K_u} \omega_u(K_u) \\
 &= \frac{d\rho_{\mathcal{C}_u}}{d\gamma_{u-}} \omega_u(K_u) \\
 (100) \quad \Rightarrow w_{u-} K_u w_u^2 &= \frac{d\gamma_{u-}}{d\gamma_{\mathcal{C}_u}} \left(\frac{d\rho_{\mathcal{C}_u}}{d\gamma_{u-}} \right)^2 K_u \omega_u(K_u)^2 = \frac{d\rho_{\mathcal{C}_u}}{d\gamma_{\mathcal{C}_u}} \frac{d\rho_{\mathcal{C}_u}}{d\gamma_{u-}} K_u \omega_u(K_u)^2 \\
 &= \frac{w_{\mathcal{C}_u}^2}{w_{u-}} K_u \omega_u(K_u)^2, \quad \text{where } w_{\mathcal{C}_u} := \frac{d\rho_{\mathcal{C}_u}}{d\gamma_{\mathcal{C}_u}}.
 \end{aligned}$$

Let's start in earnest: by the law of total variance,

$$(101) \quad \text{Var}(Z_u^N) = \text{Var}(\mathbb{E}[Z_u^N | \mathcal{F}_{\mathcal{C}_u}^N]) + \mathbb{E}[\text{Var}(Z_u^N | \mathcal{F}_{\mathcal{C}_u}^N)].$$

Using (99), we find that

$$\begin{aligned}
 \mathbb{E}[Z_u^N | \mathcal{F}_{\mathcal{C}_u}^N] &= \frac{\mathcal{Z}_{\mathcal{C}_u}^N}{N} \sum_{n=1}^N \mathbb{E}[w_u(\mathbf{X}_u^{n,N}) | \mathcal{F}_{\mathcal{C}_u}^N] = \mathcal{Z}_u^N \pi_{u-}^N(K_u w_u) = \gamma_{u-}^N(K_u w_u) \\
 &= \frac{\mathcal{Z}_{\mathcal{C}_u}^N}{N^{c_u}} \sum_{\mathbf{n} \in [N]^{c_u}} w_{u-}(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N})(K_u w_u)(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}) = \frac{\mathcal{Z}_{\mathcal{C}_u}^N}{N^{c_u}} \sum_{\mathbf{n} \in [N]^{c_u}} \frac{d\rho_u^{u-}}{d\gamma_{\mathcal{C}_u}}(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}).
 \end{aligned}$$

It follows that the first term in (101) does not depend on (γ_{u-}, K_u) and we absorb it into A in (96). Similarly,

$$\begin{aligned}
 \text{Var}(Z_u^N | \mathcal{F}_{\mathcal{C}_u}^N) &= \text{Var}(\gamma_{u-}^N(w_u) | \mathcal{F}_{\mathcal{C}_u}^N) = [N^{-1} \mathcal{Z}_u^N]^2 \sum_{n=1}^N \text{Var}(w_u(\mathbf{X}_u^{n,N}) | \mathcal{F}_{\mathcal{C}_u}^N) \\
 &= N^{-1} [\mathcal{Z}_u^N]^2 [\pi_{u-}^N(K_u w_u^2) - \pi_{u-}^N(K_u w_u)^2] \\
 &= N^{-1} \mathcal{Z}_u^N \gamma_{u-}^N(K_u w_u^2) - N^{-1} \gamma_{u-}^N(K_u w_u)^2 \quad \text{a.s.},
 \end{aligned}$$

and, using again (99), we obtain

$$\gamma_{u-}^N(K_u w_u) = \frac{\mathcal{Z}_{\mathcal{C}_u}^N}{N^{c_u}} \sum_{\mathbf{n} \in [N]^{c_u}} w_{u-}(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N})(K_u w_u)(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}) = \frac{\mathcal{Z}_{\mathcal{C}_u}^N}{N^{c_u}} \sum_{\mathbf{n} \in [N]^{c_u}} \frac{d\rho_u^{u-}}{d\gamma_{\mathcal{C}_u}}(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}),$$

another term that does not depend on (γ_{u-}, K_u) and that we absorb into A in (96). Next,

$$\frac{\mathcal{Z}_u^N}{N} \gamma_{u-}^N(K_u w_u^2) = \frac{[\mathcal{Z}_{\mathcal{C}_u}^N]^2}{N^{2c_u+1}} \left(\sum_{\mathbf{m} \in [N]^{c_u}} w_{u-}(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{m},N}) \right) \sum_{\mathbf{n} \in [N]^{c_u}} w_{u-}(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N})(K_u w_u^2)(\mathbf{X}_{\mathcal{C}_u}^{\mathbf{n},N}),$$

and (96) follows with A as above, $B := N^{-1} [\mathcal{Z}_{\mathcal{C}_u}^N]^2$, and

$$g(\mathbf{x}_{\mathcal{C}_u}^N, \gamma_{u-}, K_u) := \frac{1}{N^{2c_u}} \left(\sum_{\mathbf{m} \in [N]^{c_u}} w_{u-}(\mathbf{x}_{\mathcal{C}_u}^{\mathbf{m},N}) \right) \sum_{\mathbf{n} \in [N]^{c_u}} w_{u-}(\mathbf{x}_{\mathcal{C}_u}^{\mathbf{n}})(K_u w_u^2)(\mathbf{x}_{\mathcal{C}_u}^{\mathbf{n}})$$

$$(102) \quad = \frac{1}{N^{2c_u}} \sum_{\mathbf{n} \in [N]^{c_u}} \frac{w_{\mathcal{C}_u}(\mathbf{x}_{\mathcal{C}_u}^{\mathbf{n}})^2}{\bar{w}_{u_-}(\mathbf{x}_{\mathcal{C}_u}^{\mathbf{n}})} (K_u \omega_u(K_u)^2)(\mathbf{x}_{\mathcal{C}_u}^{\mathbf{n},N}),$$

where $\bar{w}_{u_-}(\mathbf{x}_{\mathcal{C}_u}^{\mathbf{n}}) := \frac{w_{u_-}^n(\mathbf{x}_{\mathcal{C}_u}^{\mathbf{n},N})}{\sum_{\mathbf{m} \in [N]^{c_u}} w_{u_-}(\mathbf{x}_{\mathcal{C}_u}^{\mathbf{m}})}$ and the second equality follows from (100).

Because Jensen's inequality implies that

$$K_u \left[\frac{dM_u}{dK_u} \right]^2 \geq \left[K_u \frac{dM_u}{dK_u} \right]^2 = 1 = M_u \left[\frac{dM_u}{dM_u} \right]^2,$$

(95) implies that $K_u \omega_u(K_u)^2$ bounded below by $M_u \omega_u(M_u)^2$, and (98) follows from (102). Lastly, note that the only terms in (102) that depend on γ_{u_-} are $(\bar{w}_{u_-}(\mathbf{x}_{\mathcal{C}_u}^{\mathbf{n}}))_{\mathbf{n} \in [N]^{c_u}}$. Treating these as free variables and using a Lagrange multiplier like in [4, p.153] to minimize (102) subject to the constraint ' $\sum_{\mathbf{n} \in [N]^{c_u}} \bar{w}_{u_-}(\mathbf{x}_{\mathcal{C}_u}^{\mathbf{n}}) = 1$ ', we obtain (98).

APPENDIX I: PROOF OF THEOREM 9

As shown in the proof of Theorem 6 in Appendix F, for all φ in $\mathcal{B}_b(\mathbf{E}_\tau)$,

$$\sigma_{\rho_\tau}^2(\varphi) = \sigma_{\gamma_\tau}^2(w_\tau \varphi), \quad \sigma_{\mu_\tau}^2(\varphi) = \sigma_{\gamma_\tau}^2(w_\tau Z_\tau^{-1}[\varphi - \mu_\tau(\varphi)]),$$

where $\sigma_{\gamma_\tau}^2(\varphi)$ denotes the asymptotic variance of $\gamma_\tau^N(\varphi)$ (cf. Theorem 13). Because, as explained in the main text, Algorithm 3 reduces to Algorithm 2 if we replace \mathbb{T} with \mathbb{L} , the same equations hold if we replace $(\sigma_{\mu_\tau}^2, \sigma_{\rho_\tau}^2, \sigma_{\gamma_\tau}^2)$ with $(\sigma_{\mu_\tau, smc}^2, \sigma_{\rho_\tau, smc}^2, \sigma_{\gamma_\tau, smc}^2)$ where, for any φ in $\mathcal{B}_b(\mathbf{E}_\tau)$, $\sigma_{\gamma_\tau, smc}^2(\varphi)$ denotes the asymptotic variance of $\gamma_\tau^N(\varphi)$ using \mathbb{L} instead of \mathbb{T} . For these reasons, we only need to show that $\sigma_{\gamma_\tau}^2(\varphi) \leq \sigma_{\gamma_\tau, smc}^2(\varphi)$ for all φ in $\mathcal{B}_b(\mathbf{E}_\tau)$.

Because we obtain \mathbb{L} by repeatedly lumping the children of \mathbb{T} 's nodes, it suffices to show that this lumping does not lower the asymptotic variance. Proposition 1 below shows that lumping of the children of a node u only affects the terms in the asymptotic variance sums corresponding to those children, and it follows that

$$(103) \quad \sigma_{\gamma_\tau, L}^2(\varphi) - \sigma_{\gamma_\tau}^2(\varphi) = \pi_{\mathcal{C}_u}([\mathcal{Z}_{\mathcal{C}_u} w_{u_-} K_u \Gamma_{u, \tau} \varphi - \gamma_\tau(\varphi)]^2) - \sum_{v \in \mathcal{C}_u} \pi_v([\mathcal{Z}_v \Gamma_{v, \tau} \varphi - \gamma_\tau(\varphi)]^2),$$

where $\sigma_{\gamma_\tau, L}^2(\varphi)$ denotes the asymptotic variance of the estimator for $\gamma_\tau(\varphi)$ obtained after lumping u 's children together. But,

$$\gamma_\tau(\varphi) = \gamma_u(\Gamma_{u, \tau} \varphi) = \gamma_{u_-}(K_u \Gamma_{u, \tau} \varphi) = \gamma_{\mathcal{C}_u}(w_{u_-} K_u \Gamma_{u, \tau} \varphi) = \mathcal{Z}_{\mathcal{C}_u} \pi_{\mathcal{C}_u}(w_{u_-} K_u \Gamma_{u, \tau} \varphi),$$

$$\mathcal{Z}_v \Gamma_{v, \tau} \varphi = \mathcal{Z}_v \Gamma_{v, u} \Gamma_{u, \tau} \varphi = \mathcal{Z}_v \gamma_{\mathcal{C}_u}^\psi(w_{u_-} K_u \Gamma_{u, \tau} \varphi) = \mathcal{Z}_{\mathcal{C}_u} \pi_{\mathcal{C}_u}^\psi(w_{u_-} K_u \Gamma_{u, \tau} \varphi) \quad \forall v \in \mathcal{C}_u;$$

and it follows that

$$\mathcal{Z}_{\mathcal{C}_u}^{-2}[\sigma_{\gamma_\tau, L}^2(\varphi) - \sigma_{\gamma_\tau}^2(\varphi)] = \pi_{\mathcal{C}_u}([\phi - \pi_{\mathcal{C}_u}(\phi)]^2) - \sum_{v \in \mathcal{C}_u} \pi_v(\pi_{\mathcal{C}_u}^\psi(\phi - \pi_v(\phi))^2),$$

where $\phi := w_{u_-} K_u \Gamma_{u, \tau} \varphi$. However,

$$\pi_{u1}(\pi_{\mathcal{C}_u}^{\psi1}(\phi - \pi_{u1}(\phi))^2) = \pi_{u1}(\pi_{u2}(\psi - \pi_{u1}(\psi))^2),$$

$$\pi_{u2}(\pi_{\mathcal{C}_u}^{\psi2}(\phi - \pi_{u2}(\phi))^2) = \pi_{u2}(\pi_{u1}(\psi - \pi_{u2}(\psi))^2),$$

where $\psi := \pi_{u[3:c_u]}(\phi)$ and $v[3:c_u] := \{u3, \dots, uc_u\}$. Jensen's inequality implies that

$$\pi_{u2}(\pi_{u1}(\psi - \pi_{u2}(\psi))^2) \leq \pi_{u2}(\pi_{u1}([\psi - \pi_{u2}(\psi)]^2)),$$

and it follows that

$$\begin{aligned}
 & \pi_{u1}(\pi_{\mathcal{C}_u}^{y1}(\phi - \pi_{u1}(\phi))^2) + \pi_{u2}(\pi_{\mathcal{C}_u}^{y2}(\phi - \pi_{u2}(\phi))^2) \\
 & \leq \pi_{u1}(\pi_{u2}(\psi - \pi_{u1}(\psi))^2) + \pi_{u2}(\pi_{u1}([\psi - \pi_{u2}(\psi)]^2)) \\
 & = \pi_{u1}([\pi_{u2}(\psi) - \pi_{u[2]}(\psi)]^2) + \pi_{u1}(\pi_{u2}([\psi - \pi_{u2}(\psi)]^2)) \\
 & = \pi_{u1}(\pi_{u2}(\psi)^2) - \pi_{u[2]}(\psi)^2 + \pi_{u1}(\pi_{u2}(\psi^2) - \pi_{u2}(\psi)^2) \\
 & = \pi_{u[2]}(\psi^2) - \pi_{u[2]}(\psi)^2 = \pi_{u[2]}([\psi - \pi_{u[2]}(\psi)]^2).
 \end{aligned}$$

Setting now $\psi := \pi_{u[4:c_u]}(\phi)$, we have that

$$\begin{aligned}
 & \sum_{k=1}^2 \pi_{uk}(\pi_{\mathcal{C}_u}^{yk}(\phi - \pi_{uk}(\phi))^2) \leq \pi_{u[2]}(\pi_{u3}(\psi - \pi_{u[2]}(\psi))^2), \\
 & \pi_{u3}(\pi_{\mathcal{C}_u}^{y3}(\phi - \pi_{u3}(\phi))^2) = \pi_{u3}(\pi_{u[2]}(\psi - \pi_{u3}(\psi))^2).
 \end{aligned}$$

Hence, applying Jensen's inequality again, we obtain

$$\sum_{k=1}^3 \pi_{uk}(\pi_{\mathcal{C}_u}^{yk}(\phi - \pi_{uk}(\phi))^2) \leq \pi_{u[3]}([\psi - \pi_{u[3]}(\psi)]^2).$$

Iterating this argument gives us the desired bound:

$$\begin{aligned}
 \sum_{v \in \mathcal{C}_u} \pi_v(\pi_{\mathcal{C}_u}^{y_v}(\phi - \pi_v(\phi))^2) &= \sum_{k=1}^{c_u} \pi_{uk}(\pi_{\mathcal{C}_u}^{yk}(\phi - \pi_{uk}(\phi))^2) \leq \pi_{u[k]}(\pi_{\emptyset}(\phi - \pi_{u[k]}(\phi))^2) \\
 &= \pi_{\mathcal{C}_u}([\phi - \pi_{\mathcal{C}_u}(\phi)]^2).
 \end{aligned}$$

We have one loose end to tie off: proving (103).

PROPOSITION 1. *If Assumptions 1–2 are satisfied, then (103) holds for all φ in $\mathcal{B}_b(\mathbf{E}_\tau)$.*

PROOF. Fix any φ in $\mathcal{B}_b(\mathbf{E}_\tau)$. By its definition in Theorem 13,

$$(104) \quad \sigma_{\gamma_\tau}^2(\varphi) = \sum_{v \in \mathcal{C}_u} \pi_v([\mathcal{Z}_v \Gamma_{v,\tau} \varphi - \gamma_\tau(\varphi)]^2) + \sum_{v \in \mathbb{T} \setminus \mathcal{C}_u} \pi_v([\mathcal{Z}_v \Gamma_{v,\tau} \varphi - \gamma_\tau(\varphi)]^2)$$

where we have singled out the terms corresponding to u 's children. Because the lumping only affects the terms indexed by u 's children, we have that

$$(105) \quad \sigma_{\gamma_{\tau,L}}^2(\varphi) = \pi_{\mathcal{C}_u}([\mathcal{Z}_{\mathcal{C}_u} \Gamma_{\mathcal{C}_u,\tau}^L \varphi - \gamma_\tau(\varphi)]^2) + \sum_{v \in \mathbb{T} \setminus \mathcal{C}_u} \pi_v([\mathcal{Z}_v \Gamma_{v,\tau}^L \varphi - \gamma_\tau(\varphi)]^2),$$

where $\Gamma_{v,\tau}^L$ denotes the kernel defined analogously to $\Gamma_{v,\tau}$ in (15–17) except that u 's children have been lumped together into a single node which, with an abuse of notation, we denote \mathcal{C}_u (with $\gamma_{\mathcal{C}_u}$ and $K_{\mathcal{C}_u}$ correspondingly set to $\prod_{v \in \mathcal{C}_u} \gamma_v$ and $\prod_{v \in \mathcal{C}_u} K_v$, cf. (47) for the latter). Given (104,105), we need only show that

$$(106) \quad \Gamma_{r,\tau}^L \varphi = \Gamma_{r,\tau} \varphi \quad \forall r \notin \mathcal{C}_u, \quad \Gamma_{\mathcal{C}_u,\tau}^L \varphi = w_{u-} K_u \Gamma_{u,\tau} \varphi.$$

However, the lumping leaves $(\pi_{v-})_{v \notin \mathcal{C}_u}$ and $(K_v)_{v \notin \mathcal{C}_u}$ unchanged, and so, (16) implies that

$$\Gamma_{s,r}^L = \Gamma_{s,r} \quad \forall s \in \mathcal{C}_r, \quad r \in \mathbb{T} \setminus (\{u\} \cup \mathcal{C}_u).$$

Hence, if a is an ancestor of u ($a \in \mathbb{T}$ s.t. $u \in \mathbb{T}_a$), r a grandchild of u ($r \in \cup_{v \in \mathcal{C}_u} \mathcal{C}_v$), and s a descendant of r ($s \in \mathbb{T}_r$), it follows from (17) that

$$(107) \quad \Gamma_{u,r}^L = \Gamma_{u,r}, \quad \Gamma_{a,r}^L = \Gamma_{a,r}.$$

By definition $\Gamma_{\mathcal{C}_u,u}^L = w_{u-} K_u$ and the rightmost equation in (106) follows from that in (107):

$$\Gamma_{\mathcal{C}_u,r}^L \varphi = \Gamma_{\mathcal{C}_u,u}^L \Gamma_{u,r}^L \varphi = w_{u-} K_u \Gamma_{u,r} \varphi.$$

Moreover, for any grandchild r of u with parent v ($v \in \mathcal{C}_u$ and $r \in \mathcal{C}_v$),

$$\begin{aligned} \Gamma_{r,\mathcal{C}_u}^L \varphi &= \gamma_{\cup_{v' \in \mathcal{C}_u} \mathcal{C}_{v'}}^{\mathcal{Y}'}(w_{\mathcal{C}_u-} K_{\mathcal{C}_u} \varphi) = \gamma_{\mathcal{C}_v}^{\mathcal{Y}'} \left(w_{v-} \left(\prod_{v' \in \mathcal{C}_u^{\mathcal{Y}'}} \gamma_{\mathcal{C}_{v'}} \right) \left(\left[\prod_{v' \in \mathcal{C}_u^{\mathcal{Y}'}} w_{v'-} \right] K_{\mathcal{C}_u} \varphi \right) \right) \\ &= \gamma_{\mathcal{C}_v}^{\mathcal{Y}'} \left(w_{v-} K_v \left(\prod_{v' \in \mathcal{C}_u^{\mathcal{Y}'}} \gamma_{v'-} \otimes K_{v'} \right) (\varphi) \right) = \gamma_{\mathcal{C}_v}^{\mathcal{Y}'}(w_{v-} K_v \gamma_{\mathcal{C}_u^{\mathcal{Y}'}}(\varphi)). \end{aligned}$$

Thus,

$$\begin{aligned} \Gamma_{r,u}^L \varphi &= \Gamma_{r,\mathcal{C}_u}^L \Gamma_{\mathcal{C}_u,u}^L \varphi = \gamma_{\mathcal{C}_v}^{\mathcal{Y}'}(w_{v-} K_v \gamma_{\mathcal{C}_u^{\mathcal{Y}'}}(w_{u-} K_u \varphi)) = \gamma_{\mathcal{C}_v}^{\mathcal{Y}'}(w_{v-} K_v \Gamma_{v,u} \varphi) \\ &= \Gamma_{r,v} \Gamma_{v,u} \varphi = \Gamma_{v,u} \varphi, \end{aligned}$$

and the leftmost equation in (106) also follows from (107). \square

Funding. JK and AMJ acknowledge support from the EPSRC (grant # EP/T004134/1) and the Lloyd's Register Foundation Programme on Data-Centric Engineering at the Alan Turing Institute. FRC acknowledges support from the EPSRC and the MRC OXWASP Centre for Doctoral Training (grant # EP/L016710/1). FRC and AMJ acknowledge further support from the EPSRC (grant # EP/R034710/1).

REFERENCES

- [1] AITCHISON, L. (2019). Tensor Monte Carlo: Particle Methods for the GPU era. In *Adv Neural Inf Process Syst* **32** 7148–7157.
- [2] ANDRIEU, C., DOUCET, A. and HOLENSTEIN, R. (2010). Particle Markov chain Monte Carlo methods. *J R Stat Soc B* **72** 269–342. [doi:10.1111/j.1467-9868.2009.00736.x](https://doi.org/10.1111/j.1467-9868.2009.00736.x)
- [3] ANDRIEU, C. and ROBERTS, G. O. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *Ann Stat* **37** 697–725. [doi:10.1214/07-AOS574](https://doi.org/10.1214/07-AOS574)
- [4] ASMUSSEN, S. and GLYNN, W. (2007). *Stochastic Simulation: Algorithms and Analysis*. Springer-Verlag New York. [doi:10.1007/978-0-387-69033-9](https://doi.org/10.1007/978-0-387-69033-9)
- [5] BERTI, P., PRATELLI, L. and RIGO, P. (2006). Almost sure weak convergence of random probability measures. *Stochastics and Stochastics Reports* **78** 91–97. [doi:10.1080/17442500600745359](https://doi.org/10.1080/17442500600745359)
- [6] BOUSTATI, A., AKYILDIZ, O. D., DAMOULAS, T. and JOHANSEN, A. M. (2020). Generalised Bayesian Filtering via Sequential Monte Carlo. *Adv Neural Inf Process Syst* **33** 418–429.
- [7] CARPENTER, J., CLIFFORD, P. and FEARNHEAD, P. (1999). An Improved Particle Filter for Non-linear Problems. *IEE P-Radar Son Nav* **146** 2–7.
- [8] CHAN, R., POLLOCK, M., JOHANSEN, A. M. and ROBERTS, G. O. (2021). Divide-and-Conquer Monte Carlo Fusion. *arXiv preprint arXiv:2110.07265*.
- [9] CHOPIN, N. (2002). A sequential particle filter method for static models. *Biometrika* **89** 539–551. [doi:10.1093/biomet/89.3.539](https://doi.org/10.1093/biomet/89.3.539)
- [10] CHOPIN, N. (2004). Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *Ann Stat* **32** 2385–2411. [doi:10.1214/009053604000000698](https://doi.org/10.1214/009053604000000698)
- [11] CHOPIN, N. and PAPASPILIOPOULOS, O. (2020). *An Introduction to Sequential Monte Carlo*. Springer. [doi:10.1007/978-3-030-47845-2](https://doi.org/10.1007/978-3-030-47845-2)
- [12] CONWAY, M. E. (1963). A multiprocessor system design. In *Proceedings of the November 12-14, 1963, Fall Joint Computer Conference* 139–146. [doi:10.1145/1463822.1463838](https://doi.org/10.1145/1463822.1463838)

- [13] CORNEFLOS, A., CHOPIN, N. and SÄRKKÄ, S. (2022). De-Sequentialized Monte Carlo: a parallel-in-time particle smoother. *arXiv preprint arXiv:2202.02264*.
- [14] CRISAN, D. and DOUCET, A. (2002). A Survey of Convergence Results on Particle Filtering Methods for Practitioners. *IEEE T Signal Proces* **50** 736–746. doi:10.1109/78.984773
- [15] CRISAN, D. and LYONS, T. (1997). Nonlinear filtering and measure-valued processes. *Probab Theory Relat Fields* **109** 217–244. doi:10.1007/s004400050131
- [16] CRUCINIO, F. R. (2021). Some Interacting Particle Methods with Non-Standard Interactions, Ph.D. Thesis, University of Warwick.
- [17] DAI, H., POLLOCK, M. and ROBERTS, G. (2019). Monte Carlo fusion. *J Appl Probab* **56** 174–191. doi:10.1017/jpr.2019.12
- [18] DAI, H., POLLOCK, M. and ROBERTS, G. (2021). Bayesian Fusion: Scalable unification of distributed statistical analyses.
- [19] DEL MORAL, P. (1996). Non-linear filtering: interacting particle resolution. *Markov Process Relat Fields* **2** 555–580.
- [20] DEL MORAL, P. (1998). Measure-valued processes and interacting particle systems. Application to nonlinear filtering problems. *Ann Appl Probab* **8** 438–495. doi:10.1214/aoap/1028903535
- [21] DEL MORAL, P. (2004). *Feynman-Kac Formulae*. Springer. doi:10.1007/978-1-4684-9393-1
- [22] DEL MORAL, P. (2013). *Mean field simulation for Monte Carlo integration*, 1 ed. *Monographs on Statistics and Applied Probability*. CRC Press. doi:10.1201/b14924
- [23] DEL MORAL, P., DOUCET, A. and JASRA, A. (2006). Sequential Monte Carlo samplers. *J R Stat Soc B* **68** 411–436. doi:10.1111/j.1467-9868.2006.00553.x
- [24] DING, D. (2019). Monte Carlo Algorithms for Hypothesis Testing and for Hidden Markov Models, Ph.D. Thesis, Imperial College London.
- [25] DING, D. and GANDY, A. (2018). Tree-based Particle Smoothing Algorithms in a Hidden Markov Model. *arXiv preprint arXiv:1808.08400*.
- [26] DOUC, R. and MOULINES, E. (2008). Limit theorems for weighted samples with applications to sequential Monte Carlo methods. *Ann Stat* **36** 2344–2376. doi:10.1214/07-AOS514
- [27] DOUCET, A. and JOHANSEN, A. M. (2011). A Tutorial on Particle Filtering and Smoothing: Fiteen years later. In *The Oxford Handbook of Nonlinear Filtering* (D. Crisan and B. Rozovsky, eds.) 656–704. Oxford University Press.
- [28] FADEN, A. M. (1985). The Existence of Regular Conditional Probabilities: Necessary and Sufficient Conditions. *Ann Probab* **13** 288–298. doi:10.1214/aop/1176993081
- [29] FEARNHEAD, P. and KÜNSCH, H. R. (2018). Particle Filters and Data Assimilation. *Annu Rev Stat Appl* **5** 421–449. doi:10.1146/annurev-statistics-031017-100232
- [30] GELMAN, A. and HILL, J. (2006). *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press. doi:10.1017/CBO9780511790942
- [31] GERBER, M., CHOPIN, N. and WHITELEY, N. (2019). Negative Association, Ordering and Convergence of Resampling Methods. *Ann Stat* **47** 2236–2260. doi:10.1214/18-AOS1746
- [32] GOUDIE, R. J. B., PRESANIS, A. M., LUNN, D., DE ANGELIS, D. and WERNISCH, L. (2019). Joining and splitting models with Markov melding. *Bayesian Anal* **14** 81–109. doi:10.1214/18-BA1104
- [33] GUARNIERO, P., JOHANSEN, A. M. and LEE, A. (2017). The Iterated Auxiliary Particle Filter. *J Am Stat Assoc* **112** 1636–1647. doi:10.1080/01621459.2016.1222291
- [34] GÜRBÜZBALABAN, M., GAO, X., HU, Y. and ZHU, L. (2021). Decentralized stochastic gradient Langevin dynamics and Hamiltonian Monte Carlo. *J Mach Learn Res* **22** 1–69.
- [35] JEWELL, S. W. (2015). Divide and conquer sequential Monte Carlo for phylogenetics, Master Thesis, University of British Columbia.
- [36] JOHANSEN, A. M. and DOUCET, A. (2008). A Note on the Auxiliary Particle Filter. *Stat Probab Lett* **78** 1498–1504. doi:10.1016/j.spl.2008.01.032
- [37] KALLENBERG, O. (2021). *Foundations of Modern Probability*. Springer, Cham. doi:10.1007/978-3-030-61871-1
- [38] KANTAS, N., DOUCET, A., SINGH, S. S., MACIEJOWSKI, J. and CHOPIN, N. (2015). On Particle Methods for Parameter Estimation in State-Space Models. *Stat Sci* **30** 328–351. doi:10.1214/14-STSS11
- [39] KONG, A., LIU, J. S. and WONG, W. H. (1994). Sequential Imputations and Bayesian Missing Data Problems. *J Am Stat Assoc* **89** 278–288. doi:10.1080/01621459.1994.10476469
- [40] KONG, X. and ZHENG, W. (2021). Design Based Incomplete U-Statistics. *Stat Sinica* **31** 1593–1618. doi:10.5705/ss.202019.0098
- [41] KOROLYUK, V. S. and BOROVSKICH, Y. V. (1994). *Theory of U-statistics*. Springer Science & Business Media.
- [42] KÜNSCH, H. R. (2005). Recursive Monte Carlo filters: Algorithms and theoretical analysis. *Ann Stat* **33** 1983–2021. doi:10.1214/009053605000000426

- [43] KUNTZ, J., CRUCINIO, F. R. and JOHANSEN, A. M. (2022). Product-form estimators: exploiting independence to scale up Monte Carlo. *Statistics and Computing* **32**. doi:10.1007/s11222-021-10069-9
- [44] LEE, A. and WHITELEY, N. (2018). Variance estimation in the particle filter. *Biometrika* **105** 609–625. doi:10.1093/biomet/asy028
- [45] LEE, A. J. (1990). *U-Statistics: Theory and Practice*. CRC Press.
- [46] LIN, M. T., ZHANG, J. L., CHENG, Q. and CHEN, R. (2005). Independent Particle Filters. *J Am Stat Assoc* **100** 1412–1421. doi:10.1198/016214505000000349
- [47] LINDSTEN, F., JOHANSEN, A. M., NAESSETH, C. A., KIRKPATRICK, B., SCHÖN, T. B., ASTON, J. A. D. and BOUCHARD-CÔTÉ, A. (2017). Divide-and-Conquer with sequential Monte Carlo. *J Comput Graph Stat* **26** 445–458. doi:10.1080/10618600.2016.1237363
- [48] LIU, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer-Verlag New York. doi:10.1007/978-0-387-76371-2
- [49] MANDERSON, A. A. and GOUDIE, R. J. B. (2021). Combining chains of Bayesian models with Markov melding. *arXiv preprint arXiv:2111.11566*.
- [50] MATTSON, T. G., SANDERS, B. and MASSINGILL, B. (2004). *Patterns for parallel programming*. Pearson Education.
- [51] MÍGUEZ, J., CRISAN, D. and DJURIĆ, P. M. (2013). On the convergence of two sequential Monte Carlo methods for maximum a posteriori sequence estimation and stochastic global optimization. *Stat Comput* **23** 91–107. doi:10.1007/s11222-011-9294-4
- [52] MINSKER, S., SRIVASTAVA, S., LIN, L. and DUNSON, D. B. (2014). Scalable and robust Bayesian inference via the median posterior. In *31st Int Conf Mach Learn* **32** 1656–1664.
- [53] MORAL, P. D. and GUIONNET, A. (1999). Central Limit Theorem for Nonlinear Filtering and Interacting Particle Systems. *Ann Appl Probab* **9** 275–297. doi:10.1214/aoap/1029962742
- [54] NEAL, R. M. (2001). Annealed importance sampling. *Stat Comput* **11** 125–139. doi:10.1023/A:1008923215028
- [55] NEISWANGER, W., WANG, C. and XING, E. (2014). Asymptotically Exact, Embarrassingly Parallel MCMC. In *13th Conf Uncertain Artif Intell* 623–632.
- [56] OLSSON, J. and RYDÉN, T. (2004). The bootstrap particle filtering bias. *Lund University, Technical Report 929081* **2004:24**.
- [57] PAIGE, B. and WOOD, F. (2016). Inference networks for sequential Monte Carlo in graphical models. In *33rd Int Conf Mach Learn* **48** 3040–3049.
- [58] PARAYIL, A., BAI, H., GEORGE, J. and GURRAM, P. (2020). Decentralized Langevin dynamics for Bayesian learning. *Adv Neural Inf Process Syst* **33**.
- [59] PITT, M. K. and SHEPHARD, N. (1999). Filtering via Simulation: Auxiliary Particle Filters. *J Am Stat Assoc* **94** 590–599.
- [60] RENDELL, L. J., JOHANSEN, A. M., LEE, A. and WHITELEY, N. (2021). Global Consensus Monte Carlo. *J Comput Graph Stat* **30** 249–259. doi:10.1080/10618600.2020.1811105
- [61] SCHMON, S. M., DELIGIANNIDIS, G., DOUCET, A. and PITT, M. K. (2020). Large-sample asymptotics of the pseudo-marginal method. *Biometrika* **108** 37–51. doi:10.1093/biomet/asaa044
- [62] SCOTT, S. L., BLOCKER, A. W., BONASSI, F. V., CHIPMAN, H. A., GEORGE, E. I. and MCCULLOCH, R. E. (2016). Bayes and Big Data: The Consensus Monte Carlo Algorithm. *Int J Manag Sci Eng Manag* **11** 78–88. doi:10.1080/17509653.2016.1142191
- [63] TRAN, M. H., SCHARTH, M., PITT, M. K. and KOHN, R. (2013). Importance sampling squared for Bayesian inference in latent variable models. *arXiv preprint arXiv:1309.3339*.
- [64] VONO, M., DOBIGEON, N. and CHAINAIS, P. (2021). Asymptotically Exact Data Augmentation: Models, Properties, and Algorithms. *J Comput Graph Stat* **30** 335–348. doi:10.1080/10618600.2020.1826954
- [65] WALKER, A. J. (1977). An Efficient Method for Generating Discrete Random Variables with General Distributions. *ACM T Math Software* **3** 253–256. doi:10.1145/355744.355749
- [66] WANG, X. and DUNSON, D. B. (2013). Parallelizing MCMC via Weierstrass Sampler. *arXiv preprint arXiv:1312.4605*.
- [67] YUAN, D. M. and LI, S. J. (2015). Extensions of several classical results for independent and identically distributed random variables to conditional cases. *J Korea Math Soc* **52** 431–445. doi:10.4134/JKMS.2015.52.2.431