

Collegio Carlo Alberto



Unbiased approximations of products of expectations

Anthony Lee

Simone Tiberi

Giacomo Zanella

No. 541

December 2017

Carlo Alberto Notebooks

www.carloalberto.org/research/working-papers

Unbiased approximations of products of expectations

Anthony Lee*, Simone Tiberi†, and Giacomo Zanella‡

September 4, 2017

Abstract

We consider the problem of approximating the product of n expectations with respect to a common probability distribution μ . Such products routinely arise in statistics as values of the likelihood in latent variable models. Motivated by pseudo-marginal Markov chain Monte Carlo schemes, we focus on unbiased estimators of such products. The standard approach is to sample N particles from μ and assign each particle to one of the expectations. This is wasteful and typically requires the number of particles to grow quadratically with the number of expectations. We propose an alternative estimator that approximates each expectation using most of the particles while preserving unbiasedness. We carefully study its properties, showing that in latent variable contexts the proposed estimator needs only $\mathcal{O}(n)$ particles to match the performance of the standard approach with $\mathcal{O}(n^2)$ particles. We demonstrate the procedure on two latent variable examples from approximate Bayesian computation and single-cell gene expression analysis, observing computational gains of the order of the number of expectations, i.e. data points, n .

Keywords: Latent variable models, Markov chain Monte Carlo, pseudo-marginal, approximate Bayesian computation

1 Introduction

Let X be a random variable with probability measure, or distribution, μ on a measurable space (E, \mathcal{E}) , and $L^1(\mu)$ the class of integrable, real-valued functions, i.e. $L^1(\mu) = \{f : \int_E |f(x)| \mu(dx) < \infty\}$. For a sequence of non-negative “potential” functions $G_1, \dots, G_n \in L^1(\mu)$, we consider approximations of products of n expectations

$$\gamma := \prod_{p=1}^n \mathbb{E}[G_p(X)] = \prod_{p=1}^n \mu(G_p), \quad (1)$$

*School of Mathematics, University of Bristol and The Alan Turing Institute

†Institute of Molecular Life Sciences and SIB Swiss Institute of Bioinformatics, University of Zurich

‡Department of Decision Sciences, BIDS and IGIER, Bocconi University

where we denote $\mu(f) := \int_E f(x)\mu(dx)$ for $f \in L^1(\mu)$. These arise, e.g., as values of the likelihood function in latent variable models. We concentrate on unbiased approximations of γ ; these can be used, e.g., within pseudo-marginal Markov chain methods for approximating posterior expectations (Beaumont 2003, Andrieu & Roberts 2009). To motivate this general problem, and because our main result in the sequel relates to latent variable models, we provide the following generic example of such a model.

Example 1. Let g be a Markov transition density and Y_1, \dots, Y_n be i.i.d. \mathcal{Y} -valued random variables distributed according to the probability density function ν where

$$\nu(y) := \mathbb{E}[g(X, y)] = \int_E g(x, y)\mu(dx), \quad y \in \mathcal{Y}.$$

That is, the Y_p are independent and distributed according to $g(X_p, \cdot)$ where $X_p \sim \mu$. For observations y_1, \dots, y_n , respectively, of Y_1, \dots, Y_n , we can write

$$\prod_{p=1}^n \nu(y_p) = \prod_{p=1}^n \mathbb{E}[g(X, y_p)] = \prod_{p=1}^n \mathbb{E}[G_p(X)] = \gamma,$$

where the potential functions are defined via $G_p(x) := g(x, y_p)$, for $p \in \{1, \dots, n\}$.

Remark 1. To see that γ can be viewed as a value of the likelihood function, let $\theta \in \Theta$ be a statistical parameter and $\{(\mu_\theta, g_\theta) : \theta \in \Theta\}$ parameterized families of distributions and Markov transition densities. The likelihood function L is then $L(\theta) := \prod_{p=1}^n \nu_\theta(y_p)$ where

$$\nu_\theta(y) := \mathbb{E}[g_\theta(X, y)] = \int_E g_\theta(x, y)\mu_\theta(dx), \quad y \in \mathcal{Y},$$

and clearly $L(\theta)$ is of the form (1) for any $\theta \in \Theta$.

The focus of this paper is approximations of γ using N independent and μ -distributed random variables $\zeta := (\zeta_1, \dots, \zeta_N)$, which we will refer to throughout as particles. A straightforward approach to constructing an unbiased approximation of γ is to approximate each expectation $\mathbb{E}[G_p(X)] = \mu(G_p)$ independently using M particles, where $N = Mn$. To be precise, we define

$$\gamma_{\text{simple}}^N := \prod_{p=1}^n \frac{1}{M} \sum_{i=1}^M G_p(\zeta_{(p-1)M+i}). \quad (2)$$

We will often refer to the second moment condition

$$\max_{p \in \{1, \dots, n\}} \mu(G_p^2) < \infty, \quad (3)$$

and in order to simplify the presentation we define the normalized sequence of potential functions $\bar{G}_1, \dots, \bar{G}_n$ via $\bar{G}_p := G_p/\mu(G_p)$. The following lack-of-bias, consistency, second moment and variance properties are easily established. We denote convergence in probability by \xrightarrow{P} .

Proposition 1. We have $\mathbb{E} \left[\gamma_{\text{simple}}^N \right] = \gamma$, $\gamma_{\text{simple}}^N \xrightarrow{P} \gamma$ as $N \rightarrow \infty$ and

$$\mathbb{E} \left[\left(\gamma_{\text{simple}}^N / \gamma \right)^2 \right] = \prod_{p=1}^n \left\{ 1 + \left[\mu(\bar{G}_p^2) - 1 \right] / M \right\}, \quad (4)$$

so $\text{var}(\gamma_{\text{simple}}^N)$ is finite and converges to 0 as $M \rightarrow \infty$ if and only if (3) holds.

The approximation γ_{simple}^N is straightforward to compute and analyze since it is a product of averages of independent random variables. However, each particle is only used to approximate one of the expectations in the product, and in situations where these particles are expensive to obtain this may be wasteful. An alternative approach is to use

$$\gamma_{\text{biased}}^N := \prod_{p=1}^n \frac{1}{N} \sum_{i=1}^N G_p(\zeta_i), \quad (5)$$

which is consistent and not wasteful, but also not unbiased in general.

Proposition 2. We have $\gamma_{\text{biased}}^N \xrightarrow{P} \gamma$ as $N \rightarrow \infty$ but $\mathbb{E} \left[\gamma_{\text{biased}}^N \right] \neq \gamma$ in general.

We propose in the sequel an approximation $\gamma_{\text{recycle}}^N$ that is unbiased like γ_{simple}^N but which is closer to γ_{biased}^N in that it uses most of the particles to approximate each expectation in the product while remaining computationally tractable. The approximation $\gamma_{\text{recycle}}^N$ can be viewed as an unbiased approximation of γ_{perm}^N , the rescaled permanent of a particular rectangular matrix of random variables, which is never worse in terms of variance than γ_{simple}^N but is very computationally costly to compute in general. The approximation $\gamma_{\text{recycle}}^N$ is an extension of the importance sampling approximation of the permanent of a square matrix proposed by [Kuznetsov \(1996\)](#) to the case of rectangular matrices. While it is possible for $\gamma_{\text{recycle}}^N$ to have a higher variance than γ_{simple}^N , we show that in many statistical scenarios it requires far fewer particles to obtain a given variance, e.g. in the latent variable setting described above. In particular, under weak assumptions, one needs to take $N = \mathcal{O}(n)$ to control the relative variance of $\gamma_{\text{recycle}}^N$ but one requires $N = \mathcal{O}(n^2)$ to control the relative variance of γ_{simple}^N .

The paper is structured as follows. In [Section 2](#) we draw the connection with matrix permanents, define the proposed estimator $\gamma_{\text{recycle}}^N$, and analyze basic properties such as unbiasedness and consistency. In [Section 3](#) we compare the behavior of γ_{simple}^N and $\gamma_{\text{recycle}}^N$ as n increases under various assumptions on the potential functions, while in [Section 4](#) we consider latent variable models. Finally, [Section 5](#) provides simulation studies showing the effect of using $\gamma_{\text{recycle}}^N$ rather than γ_{simple}^N in pseudo-marginal Markov chain Monte Carlo methods for estimating the parameters of a g-and-k model, commonly used as a test application for approximate Bayesian computation methodology, and of a Poisson-Beta model for single-cell gene expression. In both cases we observe computational speedups of the order of the number of data points n . [Section 6](#) provides a discussion and potential future works. All proofs are housed in the appendix or the supplementary materials.

2 The associated permanent and its approximation

For integers $i \leq j$ we denote $\llbracket i, j \rrbracket = \{i, \dots, j\}$. We adopt the convention that $\prod_{j \in \emptyset} G_j = 1$, and will occasionally use the notation $x_{p:q} = (x_p, \dots, x_q)$ for $p, q \in \mathbb{N}$ with $p \leq q$. An alternative approximation of γ on the basis of the particles $\zeta = (\zeta_1, \dots, \zeta_N)$ is obtained by first rewriting γ in (1) as, with X_1, \dots, X_p independent μ -distributed random variables,

$$\gamma = \prod_{p=1}^n \mathbb{E}[G_p(X)] = \mathbb{E} \left[\prod_{p=1}^n G_p(X_p) \right].$$

Indeed, γ_{biased}^N is a V-statistic of order n for γ , and the corresponding U-statistic for γ is

$$\gamma_{\text{perm}}^N := \sum_{k \in P(N, n)} |P(N, n)|^{-1} \prod_{p=1}^n G_p(\zeta_{k_p}), \quad (6)$$

where $P(N, n) = \{k \in \llbracket 1, N \rrbracket^n : k_i = k_j \iff i = j\}$ is the set of n -permutations of N , whose cardinality is $|P(N, n)| = N!/(N-n)!$. We observe that γ_{perm}^N is exactly $|P(N, n)|^{-1}$ times the permanent of the rectangular matrix A (see, e.g., [Ryser 1963](#), p. 25) with entries $A_{ij} = G_i(\zeta_j)$ since then

$$\text{perm}(A) = \sum_{k \in P(N, n)} \prod_{p=1}^n A_{p, k_p} = \sum_{k \in P(N, n)} \prod_{p=1}^n G_p(\zeta_{k_p}).$$

The approximation γ_{perm}^N is unbiased and consistent since it is a U-statistic and moreover it is less variable than γ_{simple}^N in terms of the convex order (see, e.g., [Shaked & Shanthikumar 2007](#), Section 3.A), defined by $X \preceq_{\text{cx}} Y$ if $\mathbb{E}[\phi(X)] \leq \mathbb{E}[\phi(Y)]$ for all convex functions $\phi : \mathbb{R} \rightarrow \mathbb{R}$ such that the expectations are well-defined. Since $x \mapsto x$ and $x \mapsto -x$ are convex functions, convex-ordered random variables necessarily have the same expectation, and since $x \mapsto x^2$ is convex, $X \preceq_{\text{cx}} Y$ implies $\text{var}(X) \leq \text{var}(Y)$. Convex-ordered families of random variables also allow one to order the asymptotic variances of associated pseudo-marginal Markov chains ([Andrieu & Vihola 2016](#), Theorem 10). In order to express the second moments of $\gamma_{\text{perm}}^N/\gamma$ and $\gamma_{\text{recycle}}^N/\gamma$, we define the function $\psi_N : \llbracket 1, N \rrbracket^n \rightarrow \mathbb{R}_+$ by

$$\psi_N(r) := \left\{ \prod_{p=1}^n \mu(\bar{G}_p \prod_{j:r_j=p} \bar{G}_j) \right\} \left\{ \prod_{i=n+1}^N \mu(\prod_{j:r_j=i} \bar{G}_j) \right\}, \quad r \in \llbracket 1, N \rrbracket^n, \quad (7)$$

where fg denotes pointwise product so that $\mu(fg) = \int f(x)g(x)\mu(dx)$. We now state basic properties of γ_{perm}^N , which can be compared with Proposition 1.

Theorem 1. *The following hold:*

1. $\mathbb{E}[\gamma_{\text{perm}}^N] = \gamma$ and $\gamma_{\text{perm}}^N \preceq_{\text{cx}} \gamma_{\text{simple}}^N$.
2. $\gamma_{\text{perm}}^N \xrightarrow{P} \gamma$ as $N \rightarrow \infty$.

Algorithm 1 Computing $\gamma_{\text{recycle}}^N$

1. Sample $\zeta_1, \dots, \zeta_N \sim \mu$ independently, and set $Z_0 \leftarrow 1$.
 2. For $p = 1, \dots, n$:
 - (a) Set $Z_p \leftarrow Z_{p-1} \sum_{j=1}^N G_p(\zeta_j) \mathbb{I}(j \notin \{K_1, \dots, K_{p-1}\}) / (N - p + 1)$.
 - (b) Sample $K_p \mid (K_1, \dots, K_{p-1})$ according to (9).
 3. Set $\gamma_{\text{recycle}}^N \leftarrow Z_n$.
-

3. The second moment of $\gamma_{\text{perm}}^N / \gamma$ is, with $K \sim \text{Uniform}(P(N, n))$,

$$\mathbb{E}[(\gamma_{\text{perm}}^N / \gamma)^2] = \mathbb{E}\left[\prod_{p=1}^n \bar{G}_p(\zeta_p) \bar{G}_p(\zeta_{K_p})\right] = \mathbb{E}[\psi_N(K)] = \mathbb{E}\left[\prod_{p=1}^n \mu(\bar{G}_p \prod_{j:K_j=p} \bar{G}_j)\right].$$

4. $\text{var}(\gamma_{\text{perm}}^N)$ is finite and $\text{var}(\gamma_{\text{perm}}^N) \rightarrow 0$ as $N \rightarrow \infty$ if and only if (3) holds.

This suggests that γ_{perm}^N is a superior approximation of γ in comparison to γ_{simple}^N . However, computing γ_{perm}^N is equivalent to computing the permanent of a rectangular matrix, which has no known polynomial-time algorithm. In fact, computing the permanent of a square matrix is #P-hard (Valiant 1979). Using an extension of the importance sampling estimator of the permanent of a square matrix due to Kuznetsov (1996), we define the following unbiased approximation of γ_{perm}^N and hence γ ,

$$\gamma_{\text{recycle}}^N := \prod_{p=1}^n \frac{1}{N - p + 1} \sum_{j=1}^N G_p(\zeta_j) \mathbb{I}(j \notin \{K_1, \dots, K_{p-1}\}), \quad (8)$$

where $K := (K_1, \dots, K_n)$ is a $[[1, N]]^n$ -valued random variable whose distribution given ζ is defined by the sequence of conditional probabilities

$$\mathbb{P}[K_p = i \mid \zeta, K_1, \dots, K_{p-1}] \propto G_p(\zeta_i) \mathbb{I}(i \notin \{K_1, \dots, K_{p-1}\}). \quad (9)$$

In (9) we take $G_p(\zeta_i) / \sum_{j=1}^N G_p(\zeta_j) \mathbb{I}(j \notin \{K_1, \dots, K_{p-1}\})$ to be 1 when the denominator is equal to 0. That is, when the denominator is 0, then $K_p \mid (\zeta, K_1, \dots, K_{p-1}) \sim \text{Uniform}([1, N] \setminus \{K_1, \dots, K_{p-1}\})$. The choice of the conditional distribution of K_p when $\sum_{j=1}^N G_p(\zeta_j) \mathbb{I}(j \notin \{K_1, \dots, K_{p-1}\}) = 0$ is in some sense arbitrary, as in any case $\gamma_{\text{recycle}}^N = 0$ whenever this happens. We now state basic properties of $\gamma_{\text{recycle}}^N$, which can be compared with Theorem 1.

Theorem 2. *The following hold:*

1. $\mathbb{E}[\gamma_{\text{recycle}}^N \mid \zeta] = \gamma_{\text{perm}}^N$, $\mathbb{E}[\gamma_{\text{recycle}}^N] = \gamma$ and $\gamma_{\text{perm}}^N \preceq_{\text{cx}} \gamma_{\text{recycle}}^N$.
2. $\gamma_{\text{recycle}}^N \xrightarrow{P} \gamma$ as $N \rightarrow \infty$.

3. Let $S = (S_1, \dots, S_n)$ be a vector of independent random variables with $S_p \sim \text{Uniform}(\llbracket p, N \rrbracket)$ for $p \in \llbracket 1, n \rrbracket$. The second moment of $\gamma_{\text{recycle}}^N/\gamma$ is

$$\mathbb{E}[(\gamma_{\text{recycle}}^N/\gamma)^2] = \mathbb{E}\left[\prod_{p=1}^n \bar{G}_p(\zeta_p)\bar{G}_p(\zeta_{S_p})\right] = \mathbb{E}[\psi_N(S)].$$

4. $\text{var}(\gamma_{\text{recycle}}^N)$ is finite and $\text{var}(\gamma_{\text{recycle}}^N) \rightarrow 0$ as $N \rightarrow \infty$ if and only if

$$\max_{p \in \llbracket 1, n \rrbracket, B \subseteq \llbracket 1, p \rrbracket} \mu(G_p \prod_{j \in B} G_j) < \infty. \quad (10)$$

Corollary 1. If $\max_{p \in \llbracket 1, n \rrbracket} \mu(G_p^{n+1}) < \infty$ then $\text{var}(\gamma_{\text{recycle}}^N) \rightarrow 0$ as $N \rightarrow \infty$.

Remark 2. We observe that $\mathbb{E}[(\gamma_{\text{perm}}^N/\gamma)^2] = \mathbb{E}[\psi_N(K)] \leq \mathbb{E}[\psi_N(S)] = \mathbb{E}[(\gamma_{\text{recycle}}^N/\gamma)^2]$, where K and S are defined in the statements of Theorems 1 and 2, respectively.

Remark 3. While (3) is sufficient for γ_{perm}^N and γ_{simple}^N to have finite variance converging to 0 as $N \rightarrow \infty$, this is not sufficient in general for $\gamma_{\text{recycle}}^N$, which requires (10) instead.

The estimator γ_{simple}^N uses only N/n out of N particles to estimate each expectation in the product; in contrast $\gamma_{\text{recycle}}^N$ uses $N - p$ particles for the p th expectation $\mu(G_p)$. In this sense, the latter recycles most of the particles for each term, and we therefore refer to $\gamma_{\text{recycle}}^N$ as the recycled estimator in the sequel. While Remark 3 implies that it is not possible for $\text{var}(\gamma_{\text{recycle}}^N) \leq \text{var}(\gamma_{\text{simple}}^N)$ in general, we show in the coming section that $\text{var}(\gamma_{\text{recycle}}^N)$ can be orders of magnitude smaller than $\text{var}(\gamma_{\text{simple}}^N)$ in many statistical settings.

3 Scaling of the number of particles with n

We investigate the variance of $\gamma_{\text{recycle}}^N$ in comparison to γ_{simple}^N in the large n regime. In particular, we show that only $N = \mathcal{O}(n)$ particles are required to control the relative variance of $\gamma_{\text{recycle}}^N$ in some scenarios in which $N = \mathcal{O}(n^2)$ particles are required to control the relative variance of γ_{simple}^N . We also show that this cannot always be true, in some situations $N = \mathcal{O}(n^2)$ is a lower bound on the number of particles required to control the relative variance of γ_{perm}^N , and therefore $\gamma_{\text{recycle}}^N$. To simplify the presentation, we define $c_p := \mu(\bar{G}_p^2) - 1$, for $p \in \llbracket 1, n \rrbracket$. We will occasionally make reference to the following assumption when considering the large n regime

$$1 < \inf_{p \geq 1} c_p \leq \sup_{p \geq 1} c_p < \infty. \quad (11)$$

We begin by observing that from Proposition 1, if (11) holds and $M = \lceil \alpha n^\beta \rceil$ then the second moment of $\gamma_{\text{simple}}^N/\gamma$ is bounded above as $n \rightarrow \infty$ if and only if $\alpha > 0$ and $\beta \geq 1$. Since $N = Mn$, this implies that to stabilize the relative variance of γ_{simple}^N in the large n regime one must take $N = \mathcal{O}(n^2)$.

The second moment of $\gamma_{\text{recycle}}^N/\gamma$ is more complex to analyze because it involves interactions between different potential functions. The following results consider three particular situations. The first is a favorable scenario, in which it follows that if $N = \alpha n$ for $\alpha > 1$ and (11) holds with $c = \sup_{p \geq 1} c_p$, then $\mathbb{E}[(\gamma_{\text{recycle}}^N/\gamma)^2] \leq \exp(c/[\alpha - 1])$.

Proposition 3. *Assume $G_1(X), \dots, G_n(X)$ are mutually independent when $X \sim \mu$. Then*

$$\mathbb{E} \left[(\gamma_{\text{recycle}}^N/\gamma)^2 \right] = \prod_{p=1}^n [1 + c_p/(N - p + 1)]. \quad (12)$$

The second scenario is also favorable: if the potential functions are negatively correlated in a specific sense, $\gamma_{\text{recycle}}^N = \gamma_{\text{perm}}^N$ and again it is sufficient to take $N = \mathcal{O}(n)$ to control the second moment of $\gamma_{\text{recycle}}^N/\gamma$.

Proposition 4. *Assume that for any distinct $p, q \in \llbracket 1, n \rrbracket$, $G_p(X)G_q(X) = 0$ almost surely. Then $\gamma_{\text{recycle}}^N = \gamma_{\text{perm}}^N$ almost surely, and $\mathbb{E} \left[(\gamma_{\text{recycle}}^N/\gamma)^2 \right] \leq \prod_{p=1}^n [1 + c_p/(N - p + 1)]$.*

The third scenario is not favorable, corresponding to the case where the potential functions are identical. At least a quadratic in n number of particles is required to control $\mathbb{E}[(\gamma_{\text{recycle}}^N/\gamma)^2]$ in this setting. Loosely speaking, positive correlations between $G_p(X)$ and $G_q(X)$ tend to increase the second moment of $\gamma_{\text{recycle}}^N/\gamma$, while correlations have no effect on the second moment of $\gamma_{\text{simple}}^N/\gamma$. Nevertheless, we show in Proposition 6 that when the moments of $\bar{G}_1(X)$ increase no more quickly than those associated to rescaled Bernoulli random variables, $\gamma_{\text{recycle}}^N$ has a smaller variance than γ_{simple}^N . Hence the recycled estimator may be useful, even if not orders of magnitude better, in some applications involving the same potential functions, e.g. the Poisson estimator of Beskos et al. (2006), based on general methods by Bhanot & Kennedy (1985) and Wagner (1988).

Proposition 5. *Assume $G_1 = \dots = G_n$. Then $\mathbb{E} \left[(\gamma_{\text{recycle}}^N/\gamma)^2 \right] \geq \mathbb{E} \left[(\gamma_{\text{perm}}^N/\gamma)^2 \right] \geq (1 + c_1)^{n^2/N}$.*

Proposition 6. *Let $G_1 = \dots = G_n$ and assume that $\mu(\bar{G}_1^\ell) \leq \mu(\bar{G}_1^2)^{\ell-1}$ for $\ell \in \llbracket 3, n \rrbracket$. Then $\mathbb{E} \left[(\gamma_{\text{recycle}}^N/\gamma)^2 \right] \leq \mathbb{E} \left[(\gamma_{\text{simple}}^N/\gamma)^2 \right]$.*

Our final general result is motivated by approximate Bayesian computation applications, in which it is often the case that the potential functions are indicator functions. In this case, it is also true that $\gamma_{\text{recycle}}^N$ has a smaller variance than γ_{simple}^N .

Proposition 7. *Let $A_1, \dots, A_n \in \mathcal{E}$ satisfy $\mu(A_p) > 0$ for $p \in \llbracket 1, n \rrbracket$. Let $G_p := \mathbb{I}_{A_p}$ for $p \in \llbracket 1, n \rrbracket$. Then $\mathbb{E} \left[(\gamma_{\text{recycle}}^N/\gamma)^2 \right] \leq \mathbb{E} \left[(\gamma_{\text{simple}}^N/\gamma)^2 \right]$.*

Remark 4. An alternative approximation of γ can be obtained by sampling a number of permutations $\sigma_1, \dots, \sigma_q$ of $\llbracket 1, n \rrbracket$ and calculating γ_{simple}^N using each permutation. That is, if we define

$$\gamma_{\text{simple}}^N(\sigma) = \prod_{p=1}^n \frac{1}{M} \sum_{i=1}^M G_p(\zeta_{\sigma((p-1)M+i)}),$$

then $q^{-1} \sum_{i=1}^q \gamma_{\text{simple}}^N(\sigma_i)$ is also an approximation of γ_{perm}^N and hence γ . This strategy does not scale well with n , however. For example, if (11) holds and $M = \lceil \alpha n^\beta \rceil$ with $\alpha > 0$ and $\beta \in (0, 1)$ we require q to grow exponentially with n to stabilize the relative variance of $q^{-1} \sum_{i=1}^q \gamma_{\text{simple}}^N(\sigma_i)$. The crucial observation to establish this is that for non-negative random variables W_1, \dots, W_q with identical means and variances, we have $\mathbb{E}[(q^{-1} \sum_{i=1}^q W_i)^2] \geq q^{-1} \mathbb{E}[W_1^2]$, and the argument follows from Proposition 1.

4 Latent variable models

The assumption of mutual independence in Proposition 3 is very strong in statistical settings. However, we show now that in latent variable models the expected second moment of $\gamma_{\text{recycle}}^N/\gamma$ is very similar to (12), where the expectation \mathbb{E} is w.r.t. the law of the observations $Y_1, \dots, Y_n \stackrel{i.i.d.}{\sim} \nu$. For the remainder of this section, we denote by \bar{G}_1 the random function $x \mapsto g(x, Y_1)/\nu(Y_1)$ for $Y_1 \sim \nu$. We begin by verifying (10) for latent variable models under a finite expected second moment condition for $\bar{G}_1(X)$ when $X \sim \mu$. This condition has appeared in the literature in a variety of places, see e.g. Breiman & Friedman (1985), Buja (1990), Schervish & Carlin (1992), Liu et al. (1995) and Khare & Hobert (2011).

Proposition 8. *In the setting of Example 1, assume that $\mathbb{E}[\mu(\bar{G}_1^2)] < \infty$. If $Y_1, \dots, Y_n \stackrel{i.i.d.}{\sim} \nu_0$ where ν_0 is absolutely continuous with respect to ν , then (10) holds almost surely.*

The following Theorem is our main result in terms of applicability to statistical scenarios. It suggests that when considering the expected second moment of $\gamma_{\text{recycle}}^N/\gamma$, it is as if the random variables $G_1(X), \dots, G_n(X)$ are “mutually independent on average”, and allows easy comparison with the corresponding expected second moment of $\gamma_{\text{simple}}^N/\gamma$.

Theorem 3. *In the setting of Example 1, and letting \mathbb{E} denoting expectation w.r.t. Y_1, \dots, Y_n ,*

$$\mathbb{E} \left[\mathbb{E} \left[(\gamma_{\text{recycle}}^N/\gamma)^2 \right] \right] = \prod_{p=1}^n [1 + C/(N - p + 1)], \quad C = \mathbb{E}[\mu(\bar{G}_1^2)] - 1.$$

Remark 5. In the setting of Example 1, it is straightforward to obtain from Proposition 1 that $\mathbb{E}[\mathbb{E}[(\gamma_{\text{simple}}^N/\gamma)^2]] = (1 + C/M)^n$, where C is as in Theorem 3. Hence, one requires $N = \lceil \alpha n \rceil$ for $\alpha > 1$ to control the expected relative variance of $\gamma_{\text{recycle}}^N$ but one requires $M = \mathcal{O}(n)$ and hence $N = \mathcal{O}(n^2)$ to control the expected relative variance of γ_{simple}^N when $1 < C < \infty$. In addition, it is clear that $\mathbb{E}[\mathbb{E}[(\gamma_{\text{recycle}}^N/\gamma)^2]] < \mathbb{E}[\mathbb{E}[(\gamma_{\text{simple}}^N/\gamma)^2]]$ for any N that is an integer multiple of $n > 1$.

Remark 6. The condition $\mathbb{E}[\mu(\bar{G}_1^2)] < \infty$ is not very strong, but is not always satisfied. For example, if μ is Uniform(0, 1) and $g(x, \cdot)$ is Uniform(0, x) for each $x \in (0, 1)$ then simple calculations show that $\mathbb{E}[\mu(\bar{G}_1^2)] = \infty$.

5 Applications

We consider Bayesian inference in two latent variable model applications, employing $\gamma_{\text{recycle}}^N$ or γ_{simple}^N to approximate $L(\theta)$ in a pseudo-marginal version of a random-walk Metropolis Markov chain. General guidelines for tuning the value of N in such chains have been proposed by [Doucet et al. \(2015\)](#) and [Sherlock et al. \(2015\)](#), who suggest that one should choose N such that the relative variance of the estimator is roughly 2. While the relative variance typically varies with θ , if the posterior distribution for θ is reasonably concentrated near the true parameter θ_0 , in practice one can often choose N so that the estimator has a relative variance of 2 at some point close to θ_0 . In both applications below, following [Roberts & Rosenthal \(2001\)](#), we tune the proposal for the random-walk Metropolis algorithm using a shorter run of the Markov chain. Specifically, we choose the proposal density

$$q(\theta, \theta') = \mathcal{N}(\theta'; \theta, d^{-1/2} 2.38 \hat{\Sigma}),$$

where $\hat{\Sigma}$ is the estimated covariance matrix of the posterior distribution and $\theta \in \mathbb{R}^d$. All computations were performed in the R programming language with C++ code accessed via the ‘Rcpp’ package ([Eddelbuettel & François 2011](#)). Effective sample sizes were computed using the ‘mcmcse’ package ([Flegal et al. 2017](#)).

Using γ_{perm}^N instead of γ_{simple}^N to approximate each $L(\theta)$ in a pseudo-marginal Markov chain can only decrease the asymptotic variance of ergodic averages of functions φ with $\text{var}_{\pi}(\varphi) < \infty$. This is a consequence of [Andrieu & Vihola \(2016, Theorem 10\)](#) and [Theorem 1](#). Using $\gamma_{\text{recycle}}^N$ does not have the same guarantee in general, but [Theorem 3](#) suggests that if the estimators perform similarly for a set of θ with large posterior mass, then this should result in greatly improved performance over γ_{simple}^N for large n .

5.1 Approximate Bayesian computation: g-and-k model

Approximate Bayesian computation (ABC) is a branch of simulation-based inference used when the likelihood function cannot be evaluated pointwise but one can simulate from the model for any value of the statistical parameter. While there are a number of variants, in general the methodology involves comparing a summary statistic associated with the observed data with summary statistics associated with pseudo-data simulated using different parameter values (see [Marin et al. 2012](#), for a recent review). When the data are modelled as n observations of i.i.d. random variables with distribution μ , it is commonplace to summarize the data using some fixed-dimensional summary statistic independent of n , for computational rather than statistical reasons. This summarization, or dimension reduction, can in principle involve little loss of information about the parameters — in exponential families sufficient statistics of fixed dimension exist and could be computed or approximated — but in practice this is not always easy to achieve. An alternative approach that we adopt here is to eschew dimension reduction altogether and treat the model as a standard latent

variable model, essentially using the noisy ABC methodology of [Fearnhead & Prangle \(2012\)](#). This may be viewed as an alternative to the construction of summaries using the Wasserstein distance recently proposed by [Bernton et al. \(2017\)](#). One possible outcome of this is that less data may be required to achieve a given degree of posterior concentration; a theoretical treatment of this is beyond the scope of this paper.

The g -and- k distribution has been used as an example application for ABC methods since [Allingham et al. \(2009\)](#). The distribution is parameterized by $\theta = (A, B, g, k)$ and a sample X from this distribution can be expressed as

$$X = A + B \left\{ 1 + c \frac{1 - \exp(-gZ)}{1 + \exp(-gZ)} \right\} \{1 + Z^2\}^k Z,$$

where $Z \sim \mathcal{N}(0, 1)$ is a standard normal random variable and we fix the value $c = 4/5$. We consider here observations y_1, \dots, y_n of n independent random variables $Y_i = X_i + \epsilon U_i$ where $X_i \stackrel{i.i.d.}{\sim} g\text{-and-}k(\theta_0)$ and $U_i \stackrel{i.i.d.}{\sim} \text{Uniform}(-\epsilon, \epsilon)$ with $\theta_0 = (3, 1, 2, 0.5)$ and $\epsilon = 1/5$. The values of c and θ_0 follow [Allingham et al. \(2009\)](#). We let μ_θ denote the distribution of $X \sim g\text{-and-}k(\theta)$, and define $G_p = \mathbb{I}_{(y_p - \epsilon, y_p + \epsilon)}$, so that $\gamma(\theta) = \prod_{p=1}^n \mu_\theta(G_p)$ is equivalent to the likelihood $L(\theta)$ associated with θ . We take $n = 100$ and, following [Allingham et al. \(2009\)](#), we put independent $\text{Uniform}(0, 10)$ priors on each component of θ . In order to have a relative variance of $\gamma_{\text{recycle}}^N(\theta_0)$ of roughly 2, it was sufficient to take $N = 100n = 10^4$ whereas for $\gamma_{\text{simple}}^N(\theta_0)$ we required $N = 100n^2 = 10^6$. Using both estimators resulted in very similar Markov chains, but the computational cost of using the simple estimator was over 30 times greater; it took 25.6 hours to simulate a simple chain of length 10^6 and 8.4 hours to simulate a recycled chain of length 10^7 . It would have taken over 10 days to simulate a simple chain of length 10^7 . Figure 1 shows posterior density estimates associated with the recycled chain; effective sample sizes for each component were above 80,000. In this example, simulating from μ is approximately 1200 times more expensive than evaluating a potential function. Finally, we observe that the posterior distribution for θ places most of its mass near θ_0 despite using $n = 100$; in contrast [Allingham et al. \(2009\)](#) used $n = 10^5$ and Figure 1 shows more concentration overall and better identification of the g parameter than their Figure 3. This suggests that this type of latent variable approach may be preferable to dimension-reducing summaries in some i.i.d. ABC models.

5.2 Poisson-Beta model for gene expression

[Peccoud & Ycart \(1995\)](#) proposed a continuous-time birth-and-death process in a random environment to model single-cell gene expression levels; this model enjoys strong experimental support ([Delmans & Hemberg 2016](#)). Letting $V = (V_t)_{t \geq 0}$ denote the \mathbb{Z}_+ -valued process counting the amount of transcribed mRNA and $W = (W_t)_{t \geq 0}$ denote the $\{0, 1\}$ -valued process indicating whether the gene is inactive or active, the Markov process $X = (X_t)_{t \geq 0} = (V, W)$ is described by,




Figure 1: Posterior density estimates for θ for the g-and-k model.

with $X_t = (v, w)$,

$$\mathbb{P}[X_{t+dt} = (v', w')] = \begin{cases} v dt + o(dt) & (v', w') = (v - 1, w), \\ [k_{\text{on}}(1 - w) + k_{\text{off}}w] dt + o(dt) & (v', w') = (v, 1 - w), \\ \lambda dt + o(dt) & (v', w', w) = (v + 1, 1, 1), \\ 0 & \text{otherwise.} \end{cases}$$

The statistical parameters k_{on} , k_{off} and λ are respectively the rates at which: the gene switches from inactive to active, the gene switches from active to inactive, and, mRNA is transcribed when the gene is active. The rate of mRNA degradation is assumed here to be 1. [Peccoud & Ycart \(1995\)](#) derive the probability generating function of the stationary distribution of this process, from which one obtains the probability mass function of the stationary marginal distribution of V

$$\mu(v) = \frac{\lambda^v \int_0^1 t^{k_{\text{on}}+v-1} (1-t)^{k_{\text{off}}-1} e^{-\lambda t} dt}{v! \text{Beta}(k_{\text{on}}, k_{\text{off}})}, \quad v \in \mathbb{Z}_+.$$

As observed by, e.g., [Kim & Marioni \(2013\)](#), straightforward calculations provide that this is equivalent to the probability mass function of $X \sim \text{Poisson-Beta}(\lambda, k_{\text{on}}, k_{\text{off}})$, defined hierarchically by $X | S \sim \text{Poisson}(\lambda S)$ where $S \sim \text{Beta}(k_{\text{on}}, k_{\text{off}})$. This model was also mentioned in [Wills et al. \(2013\)](#), who described the Poisson-Beta model directly.

In an experiment, one might observe in the stationary regime mRNA counts with noise for n independent cells, which can therefore be modelled as n independent random variables Y_1, \dots, Y_n with distribution $Y_i = X_i + \sigma Z_i$ where $X_i \stackrel{\text{ind}}{\sim} \text{Poisson-Beta}(\lambda, k_{\text{on}}, k_{\text{off}})$ and Z_1, \dots, Z_n are independent standard normal random variables. Hence, this can be viewed as a latent variable model with $\theta = (\lambda, k_{\text{on}}, k_{\text{off}})$, $\mu_\theta = \text{Poisson-Beta}(\theta)$ and $G_p(x) = \mathcal{N}(y_p; x, \sigma^2)$, and likelihood function values $L(\theta)$ are exactly of the form described in Remark 1. We simulated data y_1, \dots, y_n with $n = 1000$, $\theta_0 = (500, 2, 8)$ and $\sigma = 5$, and proceeded to conduct Bayesian inference via pseudo-marginal MCMC with independent exponential priors on λ , k_{on} and k_{off} with means of 1000, 10 and 10, respectively. In order to have a relative variance of $\gamma_{\text{recycle}}^N(\theta_0)$ of roughly 2, it was sufficient to




Figure 2: Posterior density estimates for θ for the Poisson-Beta model.

take $N = 20n = 2 \times 10^4$ whereas for $\gamma_{\text{simple}}^N(\theta_0)$ we required $N = 20n^2 = 2 \times 10^7$. Using both estimators resulted in very similar Markov chains, but the computational cost of using the simple estimator was approximately 600 times greater; it took 12.25 hours to simulate a simple chain of length 10^4 and 6.15 hours to simulate a recycled chain of length 3×10^6 . It would have taken over 153 days to simulate a simple chain of length 3×10^6 . Figure 2 shows posterior density estimates associated with the recycled chain; effective sample sizes for each component were above 23,000. In this example, simulating from μ is approximately 2700 times more expensive than evaluating a potential function.

6 Discussion

We have proposed an unbiased estimator of a product of expectations that involves using, or recycling, most of the random variables simulated. This results in considerable decreases in the computational time required to approximate such a product accurately when the number of terms in the product, n , is large and the computational cost of simulating random variables is significantly larger than that of evaluating functions involved in the terms. In latent variable models, we have shown that the number of samples N required for a given relative variance is proportional to n , while for a simple estimator one requires N to be proportional to n^2 . We have demonstrated that the use of the recycled estimator proposed here successfully reduces computational time for Bayesian inference using pseudo-marginal Markov chain Monte Carlo from days or months to hours in some situations. It would be interesting to see if the methodology could be combined with the correlated particle filter methodology of [Deligiannidis et al. \(2015\)](#) to bring further improvements.

Relating the results on numbers of samples required to common notions of asymptotic time complexity requires some care. For a given relative variance in the setting of [Theorem 3](#), one can choose α such that the following approximately holds. The number of samples required for the recycled

estimator is αn and the number of function evaluations is slightly less than αn^2 while for the simple estimator we require αn^2 samples and αn^2 function evaluations. The computational time for the recycled estimator can be expressed as $\alpha n(c_s + nc_g + nc_r)$ where c_s is the cost of sampling from μ , c_g the cost of evaluating a potential function, and c_r is the problem-independent time per particle associated with Step 2(b) in Algorithm 1. For the simple estimator, the computational time is $\alpha n^2(c_s + c_g)$ and so the recycled estimator is $(c_g + c_s)/(c_g + c_r)$ times faster than the simple estimator as $n \rightarrow \infty$, so that the improvement depends almost entirely on the relative differences between c_s , c_g and c_r . In both our applications, c_s is over a thousand times larger than c_g .

There are alternative unbiased approximations of the permanent of a rectangular matrix that could be used in place of the approach due to Kuznetsov (1996). In particular, it is straightforward to extend the algorithm of Kou & McCullagh (2009) to the rectangular case. However, we have found the corresponding approximation to be orders of magnitude worse than Kuznetsov's for the rectangular matrices used here. This is due to the fact that Kou & McCullagh's algorithm is specifically designed to overcome deficiencies of Kuznetsov's algorithm for square matrices by emphasizing the importance of large values in relation to others in the same column. In the rectangular case, Kou & McCullagh's algorithm overcompensates in this regard as this consideration is less important. There are also much more computationally expensive approximations of the permanent, such as Wang & Jasra (2016), which may be useful in situations where simulations are very expensive in comparison to function evaluations.

It would be of interest to obtain accurate, general lower bounds for the second moment of γ_{perm}^N to complement the upper bounds for $\gamma_{\text{recycle}}^N$, particularly in the setting of Example 1 to complement Theorem 3. We have been able to show that in the setting of Proposition 3, there we have $\mathbb{E}[(\gamma_{\text{perm}}^N/\gamma)^2] \geq \prod_{p=1}^n [1 + c_p/N]$ but the argument did not extend naturally to the setting of Theorem 3. Finally, it is straightforward to define $\gamma_{\text{recycle}}^N$ alternatively by choosing a permutation σ of $\llbracket 1, n \rrbracket$ according to any distribution and re-ordering the G_1, \dots, G_n as $G_{\sigma(1)}, \dots, G_{\sigma(n)}$. The corresponding condition to (10), if the distribution for σ places mass on every possible permutation of $\llbracket 1, n \rrbracket$, is then $\max_{p \in \llbracket 1, n \rrbracket, B \subseteq \llbracket 1, n \rrbracket} \mu(G_p \prod_{j \in B} G_j) < \infty$.

It is straightforward to define a recycled estimator of a product of n expectations, each with respect to a different distribution. Letting μ_1, \dots, μ_n denote the distributions, one can define a common dominating probability distribution $\tilde{\mu}$ and take, for each $p \in \llbracket 1, n \rrbracket$, $\tilde{G}_p = G_p \cdot d\mu_p/d\tilde{\mu}$ so that $\tilde{\mu}(\tilde{G}_p) = \mu_p(G_p)$. That is, one can re-express the product of expectations as a product of expectations all with respect to $\tilde{\mu}$. The results of Sections 2 and 3 then apply, and the recycled estimator could be very useful when $\tilde{\mu}(\tilde{G}_p^2)/\tilde{\mu}(\tilde{G}_p)^2$ is not too large for any $p \in \llbracket 1, n \rrbracket$. One can also compare variances of the simple estimator with the recycled estimator through Proposition 1 and Theorem 2, even though in this case the simple estimator would use blocks of independent random variables from μ_1, \dots, μ_n whereas the recycled estimator would use independent $\tilde{\mu}$ -distributed random variables.

Acknowledgments

Some of this research was undertaken while all three authors were at the University of Warwick. The authors acknowledge helpful discussions and comments from Christophe Andrieu, Christopher Jennison and Matti Vihola. AL supported by The Alan Turing Institute under the EPSRC grant EP/N510129/1. GZ supported in part by an EPSRC Doctoral Prize fellowship and by the European Research Council (ERC) through StG “N-BNP” 306406.

Appendices

A Proof of Theorem 1

The following lemma provides a sufficient condition for two random variables to be convex-ordered, which is useful for our analysis.

Lemma 1. *Let X and Y be random variables such that $\mathbb{E}[X]$ and $\mathbb{E}[Y]$ are well-defined. Then $X \preceq_{\text{cx}} Y$ if there exists a probability space with random variables X' and Y' equal in distribution respectively to X and Y , and an additional random variable Z' such that $\mathbb{E}[Y' | Z'] = X'$ almost surely.*

Proof. For convex ϕ , Jensen’s inequality provides

$$\mathbb{E}[\phi(Y')] = \mathbb{E}[\mathbb{E}[\phi(Y') | Z']] \geq \mathbb{E}[\phi(\mathbb{E}[Y' | X'])] = \mathbb{E}[\phi(X')],$$

so $\mathbb{E}[\phi(Y)] \geq \mathbb{E}[\phi(X)]$. □

Remark 7. Lemma 1 is related to the deeper and well-known Strassen Representation Theorem (Strassen 1965, Theorem 8) which states, with the notation of Lemma 1, that $X \preceq_{\text{cx}} Y$ if and only if $\mathbb{E}[Y' | X'] = X'$ almost surely. In fact, $\mathbb{E}[Y' | Z'] = X'$ implies $\mathbb{E}[Y' | X', Z'] = X'$, which is equivalent to a conditional convex order between X and Y , and which implies $X \preceq_{\text{cx}} Y$ (cf. Leskelä & Vihola 2017, Proposition 2.1(ii)).

Lemma 2. *Let R be a $[[1, N]]^n$ -valued random variable independent of ζ . Then*

$$\mathbb{E} \left[\prod_{p=1}^n \bar{G}_p(\zeta_p) \bar{G}_p(\zeta_{R_p}) \right] = \mathbb{E}[\psi_N(R)].$$

Proof. From the law of total expectation $\mathbb{E} \left[\prod_{p=1}^n \bar{G}_p(\zeta_p) \bar{G}_p(\zeta_{R_p}) \right] = \mathbb{E} \left[\mathbb{E} \left[\prod_{p=1}^n \bar{G}_p(\zeta_p) \bar{G}_p(\zeta_{R_p}) \mid R \right] \right]$,

and by collecting together terms involving the independent ζ_1, \dots, ζ_N , for any $r \in \llbracket 1, N \rrbracket^n$,

$$\begin{aligned} \mathbb{E} \left[\prod_{p=1}^n \bar{G}_p(\zeta_p) \bar{G}_p(\zeta_{r_p}) \mid R = r \right] &= \mathbb{E} \left[\left\{ \prod_{p=1}^n \bar{G}_p(\zeta_p) \prod_{j:r_j=p} \bar{G}_j(\zeta_p) \right\} \left\{ \prod_{i=n+1}^N \prod_{j:r_j=i} \bar{G}_j(\zeta_i) \right\} \right] \\ &= \left\{ \prod_{p=1}^n \mathbb{E} \left[\bar{G}_p(\zeta_p) \prod_{j:r_j=p} \bar{G}_j(\zeta_p) \right] \right\} \left\{ \prod_{i=n+1}^N \mathbb{E} \left[\prod_{j:r_j=i} \bar{G}_j(\zeta_i) \right] \right\} \\ &= \left\{ \prod_{p=1}^n \mu \left(\bar{G}_p \prod_{j:r_j=p} \bar{G}_j \right) \right\} \left\{ \prod_{i=n+1}^N \mu \left(\prod_{j:r_j=i} \bar{G}_j \right) \right\}. \quad \square \end{aligned}$$

Proof of Theorem 1. The lack-of-bias follows from the fact that that γ_{perm}^N is a U-statistic. Recalling that $N = Mn$, let $C(N, n)$ be the set of partitions of $\llbracket 1, N \rrbracket$ such that each $A \in C(N, n)$ consists of elements of cardinality exactly M denoted A_1, \dots, A_n . Then for any $A \in C(N, n)$, $\gamma_A^N := \prod_{p=1}^n \frac{1}{N} \sum_{i \in A_p} G_p(\zeta_i)$ is equal in distribution to γ_{simple}^N . Moreover, by symmetry we have

$$\gamma_{\text{perm}}^N = \sum_{k \in P(N, n)} \frac{(N-n)!}{(N)!} \prod_{p=1}^n G_p(\zeta_{k_p}) = \sum_{A \in C(N, n)} \frac{1}{|C(N, n)|} \gamma_A^N = \mathbb{E} [\gamma_S^N \mid \zeta],$$

where in the last equality $S \sim \text{Uniform}(C(N, n))$. Hence, Lemma 1 implies $\gamma_{\text{perm}}^N \preceq_{\text{cx}} \gamma_{\text{simple}}^N$. The consistency follows from the Strong Law of Large Numbers for U-statistics (Hoeffding 1961). For the second moment of $\gamma_{\text{perm}}^N/\gamma$, from the definition (6) of γ_{perm}^N , we have

$$\begin{aligned} \mathbb{E} \left[(\gamma_{\text{perm}}^N/\gamma)^2 \right] &= |P(N, n)|^{-2} \mathbb{E} \left[\sum_{k \in P(N, n)} \sum_{r \in P(N, n)} \prod_{p=1}^n \bar{G}_p(\zeta_{k_p}) \bar{G}_p(\zeta_{r_p}) \right] \\ &= |P(N, n)|^{-1} \sum_{k \in P(N, n)} \mathbb{E} \left[\prod_{p=1}^n \bar{G}_p(\zeta_p) \bar{G}_p(\zeta_{k_p}) \right] = \mathbb{E} \left[\prod_{p=1}^n \bar{G}_p(\zeta_p) \bar{G}_p(\zeta_{K_p}) \right], \end{aligned}$$

the second equality following from the exchangeability of ζ . Applying Lemma 2 we obtain

$$\begin{aligned} \mathbb{E} \left[\prod_{p=1}^n \bar{G}_p(\zeta_p) \bar{G}_p(\zeta_{K_p}) \right] &= \mathbb{E} \left[\left\{ \prod_{p=1}^n \mu \left(\bar{G}_p \prod_{j:K_j=p} \bar{G}_j \right) \right\} \left\{ \prod_{i=n+1}^N \mu \left(\prod_{j:K_j=i} \bar{G}_j \right) \right\} \right] \\ &= \mathbb{E} \left[\prod_{p=1}^n \mu \left(\bar{G}_p \prod_{j:K_j=p} \bar{G}_j \right) \right], \end{aligned}$$

the second equality following since $\mu(\bar{G}_p) = 1$ for each $p \in \llbracket 1, n \rrbracket$ and K being $P(N, n)$ -valued implies $\max_{i \in \llbracket n+1, N \rrbracket} |\{j : K_j = i\}| \leq 1$. For the last part, assume first that (3) holds. We observe then that $\max_{k \in P(N, n), p \in \llbracket 1, n \rrbracket} \mu(\bar{G}_p \prod_{j:k_j=p} \bar{G}_j) < \infty$ since for any $k \in P(N, n)$ and $p \in \llbracket 1, n \rrbracket$, $|\{j : k_j = p\}| \leq 1$ and the Cauchy–Schwarz inequality implies that $\mu(\bar{G}_p \bar{G}_q)^2 \leq \mu(\bar{G}_p^2) \mu(\bar{G}_q^2) < \infty$ for any $p, q \in \llbracket 1, n \rrbracket$. It follows that $\text{var}(\gamma_{\text{perm}}^N) < \infty$. We can write

$$\mathbb{E} \left[(\gamma_{\text{perm}}^N/\gamma)^2 \right] = \mathbb{E} [\psi_N(K)] = \mathbb{P}(A) + \mathbb{P}(A^c) \mathbb{E} \left[\psi_N(K) \mid A^c \right],$$

where $A = \{K \in P(N, n) : \min\{K_1, \dots, K_n\} > n\}$. We observe that $\mathbb{P}(A) = \prod_{i=1}^n \frac{N-n-i+1}{N-i+1} \rightarrow 1$ for fixed n as $N \rightarrow \infty$, and for $k \in P(N, n)$ we have

$$\mathbb{E} \left[\psi_N(K) \mid A^c \right] \leq \max_{k \in P(N, n)} \psi_N(k) = \max_{k \in P(N, n)} \prod_{p=1}^n \mu(\bar{G}_p) \prod_{j:k_j=p} \bar{G}_j \leq \prod_{p=1}^n \max_{q \in \llbracket 1, n \rrbracket} \mu(\bar{G}_p \bar{G}_q),$$

where the R.H.S. is finite and independent of N . Hence, $\mathbb{E} [(\gamma_{\text{perm}}^N/\gamma)^2] \rightarrow 1$ and $\text{var}(\gamma_{\text{perm}}^N) \rightarrow 0$, as desired. Finally, assume that $\mu(G_q^2) = \infty$ for some $q \in \llbracket 1, n \rrbracket$ and let $k = (1, \dots, n)$. We have then

$$\mathbb{E} [(\gamma_{\text{perm}}^N/\gamma)^2] = \mathbb{E} [\psi_N(K)] \geq |P(N, n)|^{-1} \psi_N(k),$$

and since $\psi_N(k) = \prod_{p=1}^n \mu(\bar{G}_p^2) = \infty$, we conclude that $\text{var}(\gamma_{\text{perm}}^N) = \infty$. \square

B Proof of Theorem 2

Proof of Theorem 2. This is a consequence of Lemmas 3, 4, 5, and 6 below. \square

Lemma 3. $\mathbb{E} [\gamma_{\text{recycle}}^N \mid \zeta] = \gamma_{\text{perm}}^N$, $\mathbb{E} [\gamma_{\text{recycle}}^N] = \gamma$ and $\gamma_{\text{perm}}^N \preceq_{\text{cx}} \gamma_{\text{recycle}}^N$.

Proof. Combining (8) and (9) we obtain,

$$\mathbb{E} [\gamma_{\text{recycle}}^N \mid \zeta] = |P(N, n)|^{-1} \sum_{k \in P(N, n)} \prod_{p=1}^n G_p(\zeta_{k_p}) = \gamma_{\text{perm}}^N.$$

It follows that $\mathbb{E} [\gamma_{\text{recycle}}^N] = \mathbb{E} [\gamma_{\text{perm}}^N] = \gamma$ and by Lemma 1, $\gamma_{\text{perm}}^N \preceq_{\text{cx}} \gamma_{\text{recycle}}^N$. \square

Lemma 4. $\gamma_{\text{recycle}}^N \xrightarrow{P} \gamma$ as $N \rightarrow \infty$.

Proof. Let $\tilde{Z}_{p,N} := (N-p+1)^{-1} \sum_{j=1}^N G_p(\zeta_j) \mathbb{I}(j \notin \{K_1, \dots, K_{p-1}\})$ for $p \in \llbracket 1, n \rrbracket$, so that $\gamma_{\text{recycle}}^N = \sum_{p=1}^n \tilde{Z}_{p,N}$. We will show $\tilde{Z}_{p,N} \xrightarrow{P} \mu(G_p)$ for every $p \in \llbracket 1, n \rrbracket$ and deduce, by Slutsky's Theorem, that $\gamma_{\text{recycle}}^N \xrightarrow{P} \gamma$. Since $(N-p+1)/N \rightarrow 1$ as $N \rightarrow \infty$, $\tilde{Z}_{p,N} \xrightarrow{P} \mu(G_p)$ is equivalent to $\tilde{Z}_{p,N}(N-p+1)/N \xrightarrow{P} \mu(G_p)$. Since

$$\frac{N-p+1}{N} \tilde{Z}_{p,N} = \frac{\sum_{j=1}^N G_p(\zeta_j)}{N} - \frac{\sum_{i=1}^{p-1} G_p(\zeta_{K_i})}{N},$$

and $N^{-1} \sum_{j=1}^N G_p(\zeta_j) \xrightarrow{P} \mu(G_p)$ as $N \rightarrow \infty$ by the weak Law of Large Numbers, it suffices to show that $N^{-1} \sum_{i=1}^{p-1} G_p(\zeta_{K_i}) \xrightarrow{P} 0$ as $N \rightarrow \infty$. Since $\sum_{i=1}^{p-1} G_p(\zeta_{K_i}) = 0$ for $p = 1$, it suffices to show

the result for $p \geq 2$. For any $\epsilon > 0$,

$$\begin{aligned} \mathbb{P} \left[\frac{\sum_{i=1}^{p-1} G_p(\zeta_{K_i})}{N} \geq \epsilon \right] &\leq \mathbb{P} \left[\frac{(p-1)}{N} \max_{j \in \llbracket 1, N \rrbracket} G_p(\zeta_j) \geq \epsilon \right] \\ &\leq \sum_{j=1}^N \mathbb{P} \left[G_p(\zeta_j) \geq \frac{\epsilon N}{p-1} \right] = \frac{p-1}{\epsilon} \frac{N\epsilon}{p-1} \mathbb{P} \left[G_p(\zeta_1) \geq \frac{\epsilon N}{p-1} \right] \\ &\leq \frac{p-1}{\epsilon} \mathbb{E} \left[G_p(\zeta_1) \mathbb{I} \left\{ G_p(\zeta_1) \geq \frac{\epsilon N}{p-1} \right\} \right]. \end{aligned}$$

Since $\mathbb{E}[G_p(\zeta_1)] = \mu(G_p) < \infty$, the last term converges to 0 as $N \rightarrow \infty$ since it is the tail of a convergent integral, and we conclude. \square

Lemma 5. *The second moment of $\gamma_{\text{recycle}}^N/\gamma$ can be expressed as*

$$\mathbb{E} [(\gamma_{\text{recycle}}^N/\gamma)^2] = \mathbb{E} \left[\prod_{p=1}^n \bar{G}_p(\zeta_p) \bar{G}_p(\zeta_{S_p}) \right] = \mathbb{E} [\psi_N(S)],$$

where $S = (S_1, \dots, S_n)$ is a vector of independent random variables with $S_p \sim \text{Uniform}(\llbracket p, N \rrbracket)$ for $p \in \llbracket 1, n \rrbracket$.

Proof. We obtain from (8) and (9) that

$$\mathbb{E} \left[(\gamma_{\text{recycle}}^N/\gamma)^2 \mid \zeta \right] = |P(N, n)|^{-2} \sum_{k \in P(N, n)} \prod_{p=1}^n \left[\bar{G}_p(\zeta_{k_p}) \sum_{j \in \llbracket 1, N \rrbracket \setminus \{k_1, \dots, k_{p-1}\}} \bar{G}_p(\zeta_j) \right].$$

From exchangeability of ζ and the law of total expectation, we then obtain

$$\mathbb{E} [(\gamma_{\text{recycle}}^N/\gamma)^2] = |P(N, n)|^{-1} \mathbb{E} \left[\prod_{p=1}^n \bar{G}_p(\zeta_p) \sum_{j=p}^N \bar{G}_p(\zeta_j) \right] = \mathbb{E} \left[\prod_{p=1}^n \bar{G}_p(\zeta_p) \bar{G}_p(\zeta_{S_p}) \right],$$

and conclude by applying Lemma 2. \square

Lemma 6. *$\text{var}(\gamma_{\text{recycle}}^N)$ is finite and $\text{var}(\gamma_{\text{recycle}}^N) \rightarrow 0$ as $N \rightarrow \infty$ if and only if (10) holds.*

Proof. Assume (10) holds. From Lemma 5 and (7),

$$\mathbb{E} [(\gamma_{\text{recycle}}^N/\gamma)^2] = \mathbb{E} [\psi_N(S)] = \mathbb{P}(A) + \mathbb{P}(A^c) \mathbb{E} [\psi_N(S) \mid A^c],$$

where $A = \{S : \min_i S_i > n \text{ and } S_i \neq S_j, \forall i \neq j\}$. We observe that $\mathbb{P}(A) = \prod_{i=1}^n \frac{N-n-i+1}{N-i+1} \rightarrow 1$ for fixed n as $N \rightarrow \infty$. We consider first the term $\prod_{p=1}^n \mu(\bar{G}_p \prod_{j:S_j=p} \bar{G}_j)$ in (7). For each $p \in \llbracket 1, n \rrbracket$, we can write $\mu(\bar{G}_p \prod_{j:S_j=p} \bar{G}_j) = \mu(\bar{G}_p \prod_{j \in B} \bar{G}_j)$ for $B = \{j \in \llbracket 1, n \rrbracket : S_j = p\}$ and $B \subseteq \llbracket 1, p \rrbracket$ from the definition of S . Hence,

$$\prod_{p=1}^n \mu(\bar{G}_p \prod_{j:S_j=p} \bar{G}_j) \leq \prod_{p=1}^n \max_{B \subseteq \llbracket 1, p \rrbracket} \mu(\bar{G}_p \prod_{j \in B} \bar{G}_j) < \infty,$$

by (10). Now, the term $\prod_{i=n+1}^N \mu(\prod_{j:S_j=i} \bar{G}_j)$ is a product of at most n terms different from 1, each of which can be written as $\mu(\prod_{j:S_j=i} \bar{G}_j) = \mu(\prod_{j \in \tilde{B}} \bar{G}_j)$ for some $\tilde{B} \subseteq \llbracket 1, n \rrbracket$ and hence as $\mu(\prod_{j:S_j=i} \bar{G}_j) = \mu(\bar{G}_p \prod_{j \in B} \bar{G}_j)$ for $p = \max(\tilde{B})$ and $B \subseteq \llbracket 1, p \rrbracket$. Therefore,

$$\prod_{i=n+1}^N \mu \left(\prod_{j:S_j=i} \bar{G}_j \right) \leq \left\{ \max_{p \in \llbracket 1, n \rrbracket, B \subseteq \llbracket 1, p \rrbracket} \mu \left(\bar{G}_p \prod_{j \in B} \bar{G}_j \right) \right\}^n < \infty$$

again by (10). It follows that

$$\mathbb{E} \left[\psi_N(S) \mid A^{\mathfrak{G}} \right] \leq \left\{ \max_{p \in \llbracket 1, n \rrbracket, B \subseteq \llbracket 1, p \rrbracket} \mu(\bar{G}_p \prod_{j \in B} \bar{G}_j) \right\}^{2n},$$

where the R.H.S. is finite and independent of N . Hence, $\mathbb{E} \left[(\gamma_{\text{recycle}}^N / \gamma)^2 \right] \rightarrow 1$ and $\text{var}(\gamma_{\text{recycle}}^N) \rightarrow 0$, as desired.

Suppose now that (10) does not hold. Then there exists $q \in \llbracket 1, n \rrbracket$ and $B \in \llbracket 1, q \rrbracket$ such that $\mu(G_q \prod_{j \in B} G_j) = \infty$. We treat separately the case where $q \in B$ and when $q \notin B$. If $q \in B$, then let $s \in \llbracket 1, N \rrbracket^n$ be defined by $s_p = q$ for $p \in B$ and $s_p = p$ for $p \in \llbracket 1, n \rrbracket \setminus B$. It can then be checked that

$$\psi_N(s) = \mu \left(\bar{G}_q \prod_{j \in B} \bar{G}_j \right) \left[\prod_{p \in B} \mu(\bar{G}_p) \right] \left[\prod_{p \in \llbracket 1, n \rrbracket \setminus B} \mu(\bar{G}_p^2) \right],$$

so $\psi_N(s) = \infty$ because $\mu(\bar{G}_q \prod_{j \in B} \bar{G}_j) = \infty$ and the other terms are non-zero. Since $\mathbb{P}(S = s) > 0$, where S is defined in Lemma 5, it follows that $\mathbb{E} \left[(\gamma_{\text{recycle}}^N / \gamma)^2 \right] = \mathbb{E}[\psi_N(S)] = \infty$. If instead $q \notin B$, then let $s \in \llbracket 1, N \rrbracket^n$ be defined by $s_p = q$ for all $p \in B$, $s_q = r$ for some $r \in B \setminus \{q\}$ and $s_p = p$ for $p \in \llbracket 1, n \rrbracket \setminus B$. It can then be checked that

$$\psi_N(s) = \mu \left(\bar{G}_q \prod_{j \in B} \bar{G}_j \right) \mu(\bar{G}_r \bar{G}_q) \left[\prod_{p \in B \setminus \{r\}} \mu(\bar{G}_p) \right] \left[\prod_{p \in \llbracket 1, n \rrbracket \setminus \{r\}} \mu(\bar{G}_p^2) \right],$$

so $\psi_N(s) = \infty$ because $\mu(\bar{G}_q \prod_{j \in B} \bar{G}_j) = \infty$ and the other terms are non-zero. Since $\mathbb{P}(S = s) > 0$, it follows as before that $\mathbb{E} \left[(\gamma_{\text{recycle}}^N / \gamma)^2 \right] = \mathbb{E}[\psi_N(S)] = \infty$. \square

Proof of Corollary 1. By Theorem 2, it suffices to show that $\max_{p \in \llbracket 1, n \rrbracket} \mu(G_p^{n+1}) < \infty$ implies (10). Consider an arbitrary term $\mu(G_p \prod_{j \in B} G_j)$ in (10). The generalized Hölder inequality (see, e.g., Kufner et al. 1977, p. 67) implies that

$$\mu(G_p \prod_{j \in B} G_j) \leq \mu(G_p^{|B|+1})^{1/(|B|+1)} \prod_{j \in S} \mu(G_j^{|B|+1})^{1/(|B|+1)}.$$

Since $|B| + 1 \leq (n + 1)$ we have $\mu(G_j^{|B|+1})^{1/(|B|+1)} \leq \mu(G_j^{n+1})^{1/(n+1)}$ by applying the Hölder

inequality to G_j and the constant random variable equal to 1. Therefore

$$\mu(G_p \prod_{j \in B} G_j) \leq \mu(G_p^{n+1})^{1/(n+1)} \prod_{j \in B} \mu(G_j^{n+1})^{1/(n+1)} \leq \max_{p \in \llbracket 1, n \rrbracket} \mu(G_p^{n+1}),$$

and so $\max_{p \in \llbracket 1, n \rrbracket} \mu(G_p^{n+1}) < \infty$ implies (10). \square

C Proofs of Propositions 3–5

Proof of Proposition 3. From the assumption it follows that $\mu(\bar{G}_p \prod_{j: s_j=p} \bar{G}_j) = \mu(\bar{G}_p^2)^{\mathbb{I}(s_p=p)}$ and $\mu(\prod_{j: s_j=p} \bar{G}_j) = 1$. Therefore by Lemma 5, and with $S = (S_1, \dots, S_n)$ a vector of independent random variables with $S_p \sim \text{Uniform}(\llbracket p, N \rrbracket)$ for $p \in \llbracket 1, n \rrbracket$,

$$\mathbb{E}[(\gamma_{\text{recycle}}^N / \gamma)^2] = \mathbb{E}[\psi_N(S)] = \mathbb{E}\left[\prod_{p=1}^n \mu(\bar{G}_p^2)^{\mathbb{I}(S_p=p)}\right] = \prod_{p=1}^n \mathbb{E}\left[\mu(\bar{G}_p^2)^{\mathbb{I}(S_p=p)}\right].$$

We conclude from $\mathbb{P}(S_p = p) = 1/(N - p + 1)$ and $c_p = \mu(\bar{G}_p^2) - 1$. \square

Proof of Proposition 4. Almost surely, we have $\sum_{i=1}^N G_p(\zeta_i) G_q(\zeta_i) = 0$ so

$$\sum_{j=1}^N G_p(\zeta_j) \mathbb{I}(j \notin \{K_1, \dots, K_{p-1}\}) = \sum_{j=1}^N G_p(\zeta_j),$$

and hence

$$\gamma_{\text{recycle}}^N = \prod_{p=1}^n \frac{1}{N - p + 1} \sum_{j=1}^N G_p(\zeta_j) = \frac{1}{|P(N, n)|} \sum_{k \in P(N, n)} \prod_{p=1}^n G_p(\zeta_{k_p}) = \gamma_{\text{perm}}^N.$$

The inequality holds because $\mu(\bar{G}_p \prod_{j: s_j=p} \bar{G}_j) \leq \mu(\bar{G}_p^2)^{\mathbb{I}(s_p=p)}$ and $\mu(\prod_{j: s_j=p} \bar{G}_j) \leq 1$, following the same reasoning as in the proof of Proposition 3. \square

Proof of Proposition 5. Theorem 2 implies that $\mathbb{E}[(\gamma_{\text{recycle}}^N / \gamma)^2] \geq \mathbb{E}[(\gamma_{\text{perm}}^N / \gamma)^2]$, and Theorem 1 implies that $\mathbb{E}[(\gamma_{\text{perm}}^N / \gamma)^2] = \mathbb{E}[\psi(K)] = \mathbb{E}[(1 + c_1)^Z]$, where $K \sim \text{Uniform}(P(N, n))$ and $Z = \sum_{p=1}^n \mathbb{I}(K_p \leq n)$. Jensen's inequality then provides $\mathbb{E}[(1 + c_1)^Z] \geq (1 + c_1)^{\mathbb{E}[Z]}$, and we conclude since $\mathbb{E}[Z] = \sum_{p=1}^n \mathbb{E}[\mathbb{I}\{K_p \leq n\}] = \sum_{p=1}^n \mathbb{P}[K_p \leq n] = n^2/N$. \square

D Proof of Proposition 8 and Theorem 3

Proof of Proposition 8. First assume that $Y_1, \dots, Y_n \stackrel{i.i.d.}{\sim} \nu$, and let $p \in \llbracket 1, n \rrbracket$ and $B \subseteq \llbracket 1, p \rrbracket$ be arbitrary, and \mathbb{E} denote expectation w.r.t. the law of Y_1, \dots, Y_n . From $\bar{G}_p(x) = g(x, y_p)/\nu(y_p)$ for

$p \in \llbracket 1, n \rrbracket$ we obtain

$$\begin{aligned} \mathbb{E} \left[\mu(\bar{G}_p \prod_{j \in B} \bar{G}_j) \right] &= \int \nu^{\otimes p}(y_{1:p}) \mu(\mathrm{d}x) \frac{g(x, y_p)}{\nu(y_p)} \prod_{j \in B} \frac{g(x, y_j)}{\nu(y_j)} \mathrm{d}y_{1:p} \\ &= \int \mu(\mathrm{d}x) \frac{g(x, y_p)^{1+\mathbb{I}(p \in B)}}{\nu(y_p)^{\mathbb{I}(p \in B)}} \mathrm{d}y_p, \end{aligned}$$

which is equal to 1 if $p \notin B$ and to $\mathbb{E}[\mu(\bar{G}_1^2)]$ if $p \in B$. It follows that $\mathbb{E}[\mu(\bar{G}_p \prod_{j \in B} \bar{G}_j)] < \infty$ and so $\mu(\bar{G}_p \prod_{j \in B} \bar{G}_j)$ is finite almost surely. The extension to $Y_1, \dots, Y_n \stackrel{i.i.d.}{\sim} \nu_0$ is immediate. \square

To prove Theorem 3 we first need the following Lemma.

Lemma 7. *Let $r \in \llbracket 1, n \rrbracket^\ell$ satisfy $r_1 < \dots < r_\ell$. Then $\mathbb{E}[\mu(\bar{G}_{r_1} \dots \bar{G}_{r_\ell}) \mid Y_{r_1+1}, \dots, Y_n] = \mu(\bar{G}_{r_2} \dots \bar{G}_{r_\ell})$.*

Proof. Without loss of generality, let $r_i = i$ for $i \in \llbracket 1, \ell \rrbracket$. Then,

$$\begin{aligned} \mathbb{E}[\mu(\bar{G}_1 \dots \bar{G}_\ell) \mid Y_2, \dots, Y_n] &= \int_{\mathcal{Y}} \int_E \frac{g(x, y_1)}{\nu(y_1)} \left[\prod_{k=2}^{\ell} \bar{G}_k(x) \right] \mu(\mathrm{d}x) \nu(y_1) \mathrm{d}y_1 \\ &= \int_E \int_{\mathcal{Y}} g(x, y_1) \mathrm{d}y_1 \left[\prod_{k=2}^{\ell} \bar{G}_k(x) \right] \mu(\mathrm{d}x), \end{aligned}$$

and we conclude since $\int_{\mathcal{Y}} g(x, y_1) \mathrm{d}y_1 = 1$. \square

Proof of Theorem 3. We can write $\mathbb{E}[\mathbb{E}[(\gamma_{\text{recycle}}^N / \gamma)^2]] = \mathbb{E}[\mathbb{E}[\psi_N(S, Y)]]$, where

$$\psi_N(s, y) := \left\{ \prod_{p=1}^n \mu \left(\bar{G}_p \prod_{j:s_j=p} \bar{G}_j \right) \right\} \left\{ \prod_{i=n+1}^N \mu \left(\prod_{j:s_j=i} \bar{G}_j \right) \right\}, \quad s \in \llbracket 1, N \rrbracket^n, y \in \mathcal{Y}^n,$$

with $\bar{G}_p(x) = g(x, y_p) / \nu(y_p)$ for $p \in \llbracket 1, n \rrbracket$. Since S and Y are independent, we consider terms of the form $\mathbb{E}[\psi_N(s, Y)]$, and define

$$\psi_{N,q}(s_{q:n}, y_{q:n}) := \left\{ \prod_{p=q}^n \mu \left(\bar{G}_p \prod_{j=q}^p \bar{G}_j^{\mathbb{I}\{s_j=p\}} \right) \right\} \left\{ \prod_{i=n+1}^N \mu \left(\prod_{j=q}^n \bar{G}_j^{\mathbb{I}\{s_j=i\}} \right) \right\},$$

for $q \in \llbracket 1, n \rrbracket$, which satisfies $\psi_{N,1} \equiv \psi_N$. We will show that for $s \in \llbracket 1, N \rrbracket \times \dots \times \llbracket n, N \rrbracket$ and $q \in \llbracket 1, n-1 \rrbracket$,

$$\mathbb{E}[\psi_{N,q}(s_{q:n}, Y_{q:n})] = \mathbb{E}[\mu(\bar{G}_q^2)]^{\mathbb{I}(s_q=q)} \mathbb{E}[\psi_{N,q+1}(s_{q+1:n}, Y_{q+1:n})], \quad (13)$$

by considering the cases $s_q = q$ and $s_q \neq q$. If $s_q = q$ then,

$$\begin{aligned} \mathbb{E}[\psi_{N,q}(s_{q:n}, Y_{q:n})] &= \mathbb{E}[\mathbb{E}[\psi_{N,q}(s_{q:n}, Y_{q:n}) \mid Y_{q+1:n}]] \\ &= \mathbb{E}[\mathbb{E}[\mu(\bar{G}_q^2) \psi_{N,q+1}(s_{q+1:n}, Y_{q+1:n}) \mid Y_{q+1:n}]] \\ &= \mathbb{E}[\mu(\bar{G}_q^2)] \mathbb{E}[\psi_{N,q+1}(s_{q+1:n}, Y_{q+1:n})], \end{aligned}$$

while if $s_q \neq q$ then,

$$\begin{aligned} \mathbb{E}[\psi_{N,q}(s_{q:n}, Y_{q:n})] &= \mathbb{E}[\mathbb{E}[\psi_{N,q}(s_{q:n}, Y_{q:n}) \mid Y_{q+1:n}]] \\ &= \mathbb{E}\left[\mathbb{E}\left[\mu(\bar{G}_q) \left\{ \prod_{p=q+1}^n \mu\left(\bar{G}_p \prod_{j=q}^p \bar{G}_j^{\mathbb{1}\{s_j=p\}}\right)\right\} \left\{ \prod_{i=n+1}^N \mu\left(\prod_{j=q}^n \bar{G}_j^{\mathbb{1}\{s_j=i\}}\right)\right\} \mid Y_{q+1:n}\right]\right] \\ &= \mathbb{E}[\psi_{N,q+1}(s_{q+1:n}, Y_{q+1:n})], \end{aligned}$$

where the last equality follows from $\mu(\bar{G}_q) = 1$ and Lemma 7. Hence, (13) holds for all S with positive probability and $q \in \llbracket 1, n-1 \rrbracket$. From $\mathbb{E}[\psi_{N,n}(s_n, Y_n)] = \mathbb{E}[\mu(\bar{G}_n^2)]^{\mathbb{1}\{s_n=n\}}$ and (13), we conclude that

$$\mathbb{E}[\mathbb{E}[\psi_N(S, Y)]] = \prod_{p=1}^n \mathbb{E}[\mathbb{E}[\mu(\bar{G}_p^2)]^{\mathbb{1}\{S_p=p\}}] = \prod_{p=1}^n [1 + C/(N-p+1)]. \quad \square$$

Supplementary materials

The supplementary materials consist of proofs of Propositions 1, 2, 6 and 7.

E Proofs of Propositions 1 and 2

Proof of Proposition 1. Let $Z_{p,M} := \frac{1}{M} \sum_{i=1}^M G_p(\zeta_{(p-1)M+i})$ for $p \in \llbracket 1, n \rrbracket$. Since $\mathbb{E}[Z_{p,M}] = \mu(G_p)$ for each $p \in \llbracket 1, n \rrbracket$ and $Z_{1,M}, \dots, Z_{n,M}$ are independent random variables we obtain $\mathbb{E}[\gamma_{\text{simple}}^N] = \prod_{p=1}^n \mu(G_p) = \gamma$. As $M \rightarrow \infty$, $Z_{p,M} \xrightarrow{P} \mu(G_p)$ for each $p \in \llbracket 1, n \rrbracket$ by the Weak Law of Large Numbers, and so $\gamma_{\text{simple}}^N \xrightarrow{P} \gamma$ as $M \rightarrow \infty$ by Slutsky's Theorem. To obtain the expression for the second moment (see, e.g., Goodman 1962) we have

$$\mathbb{E}[(\gamma_{\text{simple}}^N)^2] = \prod_{p=1}^n \mathbb{E}[Z_{p,M}^2] = \prod_{p=1}^n [\text{var}(Z_{p,M}) + \mu(G_p)^2],$$

from which we conclude using $\text{var}(Z_{p,M}) = [\mu(G_p^2) - \mu(G_p)^2]/M$ and the definition of $\bar{G}_1, \dots, \bar{G}_n$. \square

Proof of Proposition 2. Let $Z_{p,N} := \frac{1}{N} \sum_{i=1}^N G_p(\zeta_i)$ for $p \in \llbracket 1, n \rrbracket$. The Weak Law of Large Numbers provides that $Z_{p,N} \xrightarrow{P} \mu(G_p)$ for each $p \in \llbracket 1, n \rrbracket$ as $N \rightarrow \infty$ and so $\gamma_{\text{biased}}^N \xrightarrow{P} \gamma$ as $N \rightarrow \infty$ by

Slutsky's Theorem. However, we observe that

$$\mathbb{E}[\gamma_{\text{biased}}^N] = \mathbb{E}\left[\prod_{p=1}^n \frac{1}{N} \sum_{i=1}^N G_p(\zeta_i)\right] = \mathbb{E}\left[\prod_{p=1}^n G_p(\zeta_{K_p})\right],$$

where K is a vector of n independent $\text{Uniform}([1, N])$ random variables and this is not in general equal to γ . \square

F Proofs of Propositions 6 and 7

To prove Propositions 6 and 7 we need some additional lemmas.

Lemma 8 (See, e.g., [Esary et al. 1967](#)). *For any random variable X and non-decreasing real-valued functions g_1 and g_2 , $\mathbb{E}[g_1(X)g_2(X)] \geq \mathbb{E}[g_1(X)]\mathbb{E}[g_2(X)]$.*

Lemma 9. *Let B_1, \dots, B_n be Bernoulli r.v.s with $B_1 \sim \text{Bernoulli}(n/N)$ and*

$$B_p \mid (B_1, \dots, B_{p-1}) \sim \text{Bernoulli}\left(\frac{n - \sum_{j=1}^{p-1} B_j}{N - p + 1}\right). \quad (14)$$

Then, for any m_1, \dots, m_n all greater than or equal to 1, $\mathbb{E}\left[\prod_{p=1}^n m_p^{B_p}\right] \leq \prod_{p=1}^n \mathbb{E}\left[m_p^{B_p}\right]$.

Proof. For any $\ell \in [1, n]$, define $Z_\ell = \sum_{p=1}^\ell B_p$. From (14), $(B_1, \dots, B_{\ell-1})$ and B_ℓ are conditionally independent given $Z_{\ell-1}$ and so

$$\mathbb{E}\left[-\prod_{p=1}^\ell m_p^{B_p}\right] = \mathbb{E}\left[\mathbb{E}\left[-m_\ell^{B_\ell} \mid Z_{\ell-1}\right] \mathbb{E}\left[\prod_{p=1}^{\ell-1} m_p^{B_p} \mid Z_{\ell-1}\right]\right].$$

We now show that $g_1(Z_{\ell-1}) := \mathbb{E}\left[-m_\ell^{B_\ell} \mid Z_{\ell-1}\right]$ and $g_2(Z_{\ell-1}) := \mathbb{E}\left[\prod_{p=1}^{\ell-1} m_p^{B_p} \mid Z_{\ell-1}\right]$ are non-decreasing functions of $Z_{\ell-1}$ so that $\mathbb{E}[g_1(Z_{\ell-1})g_2(Z_{\ell-1})] \geq \mathbb{E}[g_1(Z_{\ell-1})]\mathbb{E}[g_2(Z_{\ell-1})]$ by Lemma 8. That g_1 is non-decreasing follows from (14) and $m_p \geq 1$. Interpreting $(B_1, \dots, B_{\ell-1})$ as a draw from a hypergeometric experiment, we can rewrite $g_2(Z_{\ell-1}) = \mathbb{E}\left[\prod_{j=1}^{Z_{\ell-1}} m_{I_j}\right]$ where $I_1, \dots, I_{Z_{\ell-1}}$ are drawn uniformly without replacement from $[1, \ell-1]$. Hence, from $m_p \geq 1$ for all $p \in [1, \ell-1]$ and a simple coupling argument we obtain that g_2 is also non-decreasing. It follows that $\mathbb{E}\left[\prod_{p=1}^\ell m_p^{B_p}\right] \leq \mathbb{E}\left[m_\ell^{B_\ell}\right] \mathbb{E}\left[\prod_{p=1}^{\ell-1} m_p^{B_p}\right]$ and since ℓ is arbitrary we can conclude. \square

Proof of Proposition 6. Let $m = \mu(\bar{G}_1^2)$ and observe that for every s with $s_p \in [p+1, n]$

$$\psi_N(s) \leq \left(\prod_{p=1}^n m^{|\{j:s_j=p\}|}\right) \prod_{p=n+1}^N m^{(|\{j:s_j=p\}|-1) \vee 0} = m^{\sum_{p=1}^n b_p},$$

where $b_p = \mathbb{I}(s_p \leq n \text{ or } s_p \in \{s_1, \dots, s_{p-1}\})$. Hence, $\mathbb{E}\left[(\gamma_{\text{recycle}}^N/\gamma)^2\right] = \mathbb{E}[\psi_N(S)] \leq \mathbb{E}\left[m^{\sum_{p=1}^n B_p}\right]$

where S is as defined in Theorem 2. From Lemma 9, we conclude that

$$\mathbb{E} [(\gamma_{\text{recycle}}^N/\gamma)^2] \leq \mathbb{E} \left[m^{\sum_{p=1}^n B_p} \right] \leq \prod_{p=1}^n \mathbb{E} [m^{B_p}] = \mathbb{E} [(\gamma_{\text{simple}}^N/\gamma)^2]. \quad \square$$

Lemma 10. *Let G_1, \dots, G_n be as in the statement of Proposition 7. Then for any $s \in \llbracket 1, N \rrbracket^n$ such that $s_p \geq p$ for all $p \in \llbracket 1, n \rrbracket$ we have $\psi_N(s) \leq \prod_{p=1}^n \mu(\bar{G}_p^2)^{b_p}$, where $b_p = \mathbb{I}(s_p \in \llbracket 1, n \rrbracket \cup \{s_1, \dots, s_{p-1}\})$.*

Proof. Define $m_p := 1/\mu(A_p)$. It follows that $\bar{G}_p = m_p \mathbb{I}_{A_p}$ and $\mu(\bar{G}_p^2) = m_p$. Moreover, for $i_1, \dots, i_p \in \llbracket 1, n \rrbracket$ we have $\mu(\prod_{j=1}^p \bar{G}_{i_j}) \leq \left(\prod_{j=1}^p m_{i_j} \right) / \max_{j \in \llbracket 1, p \rrbracket} m_{i_j}$, with equality if the sets A_{i_1}, \dots, A_{i_p} are nested, i.e. $A_{i_j} \subseteq A_{i_k}$ or $A_{i_k} \subseteq A_{i_j}$ for distinct $j, k \in \llbracket 1, p \rrbracket$. Since $\psi_N(s)$ is a non-decreasing function of products of expressions of the form $\mu(\prod_{j=1}^p \bar{G}_{i_j})$, we can upper bound $\psi_N(s)$ by assuming henceforth that A_1, \dots, A_n are nested, in which case we observe that $\mu(\prod_{j=1}^p \bar{G}_{i_j}) = \left(\prod_{j=1}^p m_{i_j} \right) / \max_{j \in \llbracket 1, p \rrbracket} m_{i_j} \leq \prod_{j=2}^p m_{i_j}$. Plugging this inequality carefully into $\psi_N(s)$ using the definition of b_1, \dots, b_n gives $\psi_N(s) \leq \prod_{p=1}^n m_p^{b_p}$. \square

Proof of Proposition 7. From Theorem 2 and Lemma 10 we have

$$\mathbb{E} [(\gamma_{\text{recycle}}^N/\gamma)^2] = \mathbb{E} [\psi_N(S)] \leq \mathbb{E} \left[\prod_{p=1}^n \mu(\bar{G}_p^2)^{B_p} \right],$$

where S is as defined in Theorem 2, and $B_p = \mathbb{I}(S_p \in \llbracket 1, n \rrbracket \cup \{S_1, \dots, S_{p-1}\})$ for $p \in \llbracket 1, n \rrbracket$. From Lemma 9,

$$\mathbb{E} \left[\prod_{p=1}^n \mu(\bar{G}_p^2)^{B_p} \right] \leq \prod_{p=1}^n \mathbb{E} [\mu(\bar{G}_p^2)^{B_p}] = \prod_{p=1}^n \mathbb{E} \left[1 + \frac{\mu(\bar{G}_p^2) - 1}{N/n} \right] = \mathbb{E} [(\gamma_{\text{simple}}^N/\gamma)^2]. \quad \square$$

References

- Allingham, D., King, R. A. R. & Mengersen, K. L. (2009), ‘Bayesian estimation of quantile distributions’, *Stat. Comput.* **19**(2), 189–201.
- Andrieu, C. & Roberts, G. O. (2009), ‘The pseudo-marginal approach for efficient Monte Carlo computations’, *Ann. Statist.* **37**(2), 697–725.
- Andrieu, C. & Vihola, M. (2016), ‘Establishing some order amongst exact approximations of MCMCs’, *Ann. Appl. Probab.* **26**(5), 2661–2696.
- Beaumont, M. A. (2003), ‘Estimation of population growth or decline in genetically monitored populations’, *Genetics* **164**(3), 1139–1160.
- Bernton, E., Jacob, P. E., Gerber, M. & Robert, C. P. (2017), ‘Inference in generative models using the Wasserstein distance’, *arXiv preprint arXiv:1701.05146*.

- Beskos, A., Papaspiliopoulos, O., Roberts, G. O. & Fearnhead, P. (2006), ‘Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes (with discussion)’, *J. R. Stat. Soc. B* **68**(3), 333–382.
- Bhanot, G. & Kennedy, A. D. (1985), ‘Bosonic lattice gauge theory with noise’, *Phys. Lett. B* **157**(1), 70–76.
- Breiman, L. & Friedman, J. H. (1985), ‘Estimating optimal transformations for multiple regression and correlation’, *J. Amer. Statist. Assoc.* **80**(391), 580–598.
- Buja, A. (1990), ‘Remarks on functional canonical variates, alternating least squares methods and ACE’, *Ann. Statist.* pp. 1032–1069.
- Deligiannidis, G., Doucet, A. & Pitt, M. K. (2015), ‘The correlated pseudo-marginal method’, *arXiv preprint arXiv:1511.04992*.
- Delmans, M. & Hemberg, M. (2016), ‘Discrete distributional differential expression (D³E) - a tool for gene expression analysis of single-cell RNA-seq data’, *BMC Bioinform.* **17**(1), 110.
- Doucet, A., Pitt, M. K., Deligiannidis, G. & Kohn, R. (2015), ‘Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator’, *Biometrika* **102**(2), 295–313.
- Eddelbuettel, D. & François, R. (2011), ‘Rcpp: Seamless R and C++ integration’, *J. Stat. Softw.* **40**(8), 1–18.
- Esary, J. D., Proschan, F. & Walkup, D. W. (1967), ‘Association of random variables, with applications’, *Ann. Math. Stat.* **38**(5), 1466–1474.
- Fearnhead, P. & Prangle, D. (2012), ‘Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation’, *J. R. Stat. Soc. B* **74**(3), 419–474.
- Flegal, J. M., Hughes, J., Vats, D. & Dai, N. (2017), *mcmcse: Monte Carlo Standard Errors for MCMC*, Riverside, CA, Denver, CO, Coventry, UK, and Minneapolis, MN. R package version 1.3-2.
- Goodman, L. A. (1962), ‘The variance of the product of k random variables’, *J. Amer. Statist. Assoc.* **57**(297), 54–60.
- Hoeffding, W. (1961), The strong law of large numbers for U-statistics, Technical report, North Carolina State University, Dept. of Statistics.
- Khare, K. & Hobert, J. P. (2011), ‘A spectral analytic comparison of trace-class data augmentation algorithms and their sandwich variants’, *Ann. Statist.* **39**(5), 2585–2606.
- Kim, J. K. & Marioni, J. C. (2013), ‘Inferring the kinetics of stochastic gene expression from single-cell rna-sequencing data’, *Genome Biol.* **14**(1), R7.

- Kou, S. C. & McCullagh, P. (2009), ‘Approximating the α -permanent’, *Biometrika* **96**(3), 635–644.
- Kufner, A., John, O. & Fucik, S. (1977), *Function spaces*, Vol. 3, Noordhoff, Leyden.
- Kuznetsov, N. Y. (1996), ‘Computing the permanent by importance sampling method’, *Cybernet. Systems Anal.* **32**(6), 749–755.
- Leskelä, L. & Vihola, M. (2017), ‘Conditional convex orders and measurable martingale couplings’, *Bernoulli*. To appear.
- Liu, J. S., Wong, W. H. & Kong, A. (1995), ‘Covariance structure and convergence rate of the Gibbs sampler with various scans’, *J. R. Stat. Soc. B* pp. 157–169.
- Marin, J.-M., Pudlo, P., Robert, C. P. & Ryder, R. J. (2012), ‘Approximate Bayesian computational methods’, *Stat. Comput.* **22**(6), 1167–1180.
- Peccoud, J. & Ycart, B. (1995), ‘Markovian modeling of gene-product synthesis’, *Theor. Popul. Biol.* **48**(2), 222–234.
- Roberts, G. O. & Rosenthal, J. S. (2001), ‘Optimal scaling for various Metropolis–Hastings algorithms’, *Statist. Sci.* **16**(4), 351–367.
- Ryser, H. J. (1963), *Combinatorial Mathematics*, Vol. 14 of *Carus Mathematical Monographs*, Math. Assoc. America.
- Schervish, M. J. & Carlin, B. P. (1992), ‘On the convergence of successive substitution sampling’, *J. Comput. Graph. Statist.* **1**(2), 111–127.
- Shaked, M. & Shanthikumar, J. G. (2007), *Stochastic Orders*, Springer, New York.
- Sherlock, C., Thiery, A. H., Roberts, G. O. & Rosenthal, J. S. (2015), ‘On the efficiency of pseudo-marginal random walk Metropolis algorithms’, *Ann. Statist.* **43**(1), 238–275.
- Strassen, V. (1965), ‘The existence of probability measures with given marginals’, *Ann. Math. Statist.* pp. 423–439.
- Valiant, L. G. (1979), ‘The complexity of computing the permanent’, *Theoret. Comput. Sci.* **8**(2), 189–201.
- Wagner, W. (1988), ‘Unbiased multi-step estimators for the Monte Carlo evaluation of certain functional integrals’, *J. Comput. Phys.* **79**(2), 336–352.
- Wang, J. & Jasra, A. (2016), ‘Monte Carlo algorithms for computing α -permanents’, *Stat. Comput.* **26**(1-2), 231–248.
- Wills, Q. F., Livak, K. J., Tipping, A. J., Enver, T., Goldson, A. J., Sexton, D. W. & Holmes, C. (2013), ‘Single-cell gene expression analysis reveals genetic associations masked in whole-tissue experiments’, *Nature Biotechnol.* **31**(8), 748–752.