

Yet another look at Harris' ergodic theorem for Markov chains

Martin Hairer and Jonathan C. Mattingly

The aim of this note is to present an elementary proof of a variation of Harris' ergodic theorem of Markov chains. This theorem, dating back to the fifties [Har56] essentially states that a Markov chain is uniquely ergodic if it admits a “small” set (in a technical sense to be made precise below) which is visited infinitely often. This gives an extension of the ideas of Doeblin to the unbounded state space setting. Often this is established by finding a Lyapunov function with “small” level sets [Has80, MT93]. If the Lyapunov function is strong enough, one has a spectral gap in a weighted supremum norm [MT92, MT93]. In particular, its transition probabilities converge exponentially fast towards the unique invariant measure, and the constant in front of the exponential rate is controlled by the Lyapunov function [MT93].

Traditional proofs of this result rely on the decomposition of the Markov chain into excursions away from the small set and a careful analysis of the exponential tail of the length of these excursions [Num84, Cha89, MT92, MT93]. There have been other variations which have made use of Poisson equations or worked at getting explicit constants [KM05, DMR04, DMLM03]. The present proof is very direct, and relies instead on introducing a family of equivalent weighted norms indexed by a parameter β and to make an appropriate choice of this parameter that allows to combine in a very elementary way the two ingredients (existence of a Lyapunov function and irreducibility) that are crucial in obtaining a spectral gap. Use of a weighted total-variation norm has been important since [MT92].

The original motivation of this proof was the authors' work on spectral gaps in Wasserstein metrics. The proof presented in this note is a version of our reasoning in the total variation setting which we used to guide the calculations in [HM08]. While we initially produced it for this purpose, we hope that it will be of interest in its own right.

1. Setting and main result

Throughout this note, we fix a measurable space \mathbf{X} and a Markov transition kernel $\mathcal{P}(x, \cdot)$ on \mathbf{X} . We will use the notation \mathcal{P} for the operators defined as usual on both the set of

bounded measurable functions and the set of measures of finite mass by

$$(\mathcal{P}\varphi)(x) = \int_{\mathbf{X}} \varphi(y) \mathcal{P}(x, dy), \quad (\mathcal{P}\mu)(A) = \int_{\mathbf{X}} \mathcal{P}(x, A) \mu(dx).$$

Hence we are using \mathcal{P} both to denote the action on functions and its dual action on measure. Note that \mathcal{P} extends trivially to measurable functions $\varphi: \mathbf{X} \rightarrow [0, +\infty]$. We first assume that \mathcal{P} satisfies the following geometric drift condition:

Assumption 1. There exists a function $V: \mathbf{X} \rightarrow [0, \infty)$ and constants $K \geq 0$ and $\gamma \in (0, 1)$ such that

$$(\mathcal{P}V)(x) \leq \gamma V(x) + K, \quad (1)$$

for all $x \in \mathbf{X}$.

Remark 1.1. One could allow V to also take the value $+\infty$. However, since we do not assume any particular structure on \mathbf{X} , this case can immediately be reduced to the present case by replacing \mathbf{X} by $\{x: V(x) < \infty\}$.

Assumption 1 ensures that the dynamics enters the “center” of the state space regularly with tight control on the length of the excursions from the center. We now assume that a sufficiently large level set of V is sufficiently “nice” in the sense that we have a uniform “minorization” condition reminiscent of Doeblin’s condition, but localized to the interior of the level set.

Assumption 2. There exists a constant $\alpha \in (0, 1)$ and a probability measure ν so that

$$\inf_{x \in \mathcal{C}} \mathcal{P}(x, \cdot) \geq \alpha \nu(\cdot),$$

with $\mathcal{C} = \{x \in \mathbf{X}: V(x) \leq R\}$ for some $R > 2K/(1 - \gamma)$ where K and γ are the constants from Assumption 1.

In order to state the version Harris’ theorem under consideration, we introduce the following weighted supremum norm:

$$\|\varphi\| = \sup_x \frac{|\varphi(x)|}{1 + V(x)}. \quad (2)$$

With this notation at hand, one has:

Theorem 1.2. *If Assumptions 1 and 2 hold, then \mathcal{P} admits a unique invariant measure μ_* . Furthermore, there exist constants $C > 0$ and $\gamma \in (0, 1)$ such that the bound*

$$\|\mathcal{P}^n \varphi - \mu_*(\varphi)\| \leq C \gamma^n \|\varphi - \mu_*(\varphi)\|$$

holds for every measurable function $\varphi: \mathbf{X} \rightarrow \mathbf{R}$ such that $\|\varphi\| < \infty$.

While this result is well-known, the proofs found in the literature are often quite involved and rely on careful estimates of the return times to small sets, combined with a clever application of Kendall’s lemma. See for example [MT93, Section 15].

The aim of this note is to provide a very short and elementary proof of Theorem 1.2 based on a simple trick. Instead of working directly with (2), we define a whole family of

weighted supremum norms depending on a scale parameter $\beta > 0$ that are all equivalent to the original norm (2):

$$\|\varphi\|_\beta = \sup_x \frac{|\varphi(x)|}{1 + \beta V(x)} .$$

We also define the associated dual metric ρ_β on probability measures given by

$$\rho_\beta(\mu_1, \mu_2) = \sup_{\varphi: \|\varphi\|_\beta \leq 1} \int_{\mathbf{X}} \varphi(x) (\mu_1 - \mu_2)(dx) . \quad (3)$$

It is well-known that ρ_β is nothing but a weighted total variation distance:

$$\rho_\beta(\mu_1, \mu_2) = \int_{\mathbf{X}} (1 + \beta V(x)) |\mu_1 - \mu_2|(dx) .$$

With these notations, our main result is:

Theorem 1.3. *If Assumptions 1 and 2 hold, then there exists $\bar{\alpha} \in (0, 1)$ and $\beta > 0$ so that*

$$\rho_\beta(\mathcal{P}\mu_1, \mathcal{P}\mu_2) \leq \bar{\alpha} \rho_\beta(\mu_1, \mu_2)$$

for any probability measure μ_1 and μ_2 on \mathbf{X} . In particular, for any $\alpha_0 \in (0, \alpha)$ and $\gamma_0 \in (\gamma + 2K/R, 1)$ one can choose $\beta = \alpha_0/K$ and $\bar{\alpha} = (1 - (\alpha - \alpha_0)) \vee (2 + R\beta\gamma_0)/(2 + R\beta)$.

Remark 1.4. The interest of this result lies in the fact that it is possible to tune β in such a way that \mathcal{P} is a strict contraction for the distance ρ_β . In general, this does *not* imply that \mathcal{P} is a contraction for ρ_1 , say, even though the equivalence of the norms $\|\cdot\|_\beta$ does of course imply that there exists $n > 0$ such that \mathcal{P}^n is such a contraction.

2. Alternative formulation of metric ρ_β

We now introduce an alternative definition of the weighted total variation norm ρ_β . We begin by defining a metric d_β between points in \mathbf{X} by

$$d_\beta(x, y) = \begin{cases} 0 & x = y \\ 2 + \beta V(x) + \beta V(y) & x \neq y \end{cases}$$

Though slightly odd looking, the reader can readily verify that since $V \geq 0$, d_β indeed satisfies the axioms of a metric. This metric in turn induces a Lipschitz seminorm on measurable functions and a metric on probability measures defined respectively by

$$\|\varphi\|_\beta = \sup_{x \neq y} \frac{|\varphi(x) - \varphi(y)|}{d_\beta(x, y)} ,$$

$$d_\beta(\mu_1, \mu_2) = \sup_{\varphi: \|\varphi\|_\beta \leq 1} \int \varphi(x) (\mu_1 - \mu_2)(dx) .$$

It turns out that these norms are almost identical to the ones from the previous section. More precisely, one has:

Lemma 2.1. *One has the identity $\|\varphi\|_\beta = \inf_{c \in \mathbf{R}} \|\varphi + c\|_\beta$. In particular, $d_\beta = \rho_\beta$.*

Proof. It is obvious that $\|\varphi\|_\beta \leq \|\varphi\|_\beta$ and therefore $\|\varphi\|_\beta \leq \inf_{c \in \mathbf{R}} \|\varphi + c\|_\beta$, so it remains to show the reverse inequality.

Given any φ with $\|\varphi\|_\beta \leq 1$, we set $c = \inf_x (1 + \beta V(x) - \varphi(x))$. Observe that for any x and y , $\varphi(x) \leq |\varphi(y)| + |\varphi(x) - \varphi(y)| \leq |\varphi(y)| + 2 + \beta V(x) + \beta V(y)$. Hence $1 + \beta V(x) - \varphi(x) \geq -1 - \beta V(y) - |\varphi(y)|$. Since there exists at least one point with $V(y) < \infty$ we see that c is bounded from below and hence $|c| < \infty$.

Observe now that

$$\varphi(x) + c \leq \varphi(x) + 1 + \beta V(x) - \varphi(x) = 1 + \beta V(x),$$

and

$$\begin{aligned} \varphi(x) + c &= \inf_y \varphi(x) + 1 + \beta V(y) - \varphi(y) \\ &\geq \inf_y 1 + \beta V(y) - \|\varphi\|_\beta \cdot d_\beta(x, y) \geq -(1 + \beta V(x)), \end{aligned}$$

so that $|\varphi(x) + c| \leq 1 + \beta V(x)$ as required.

It follows that the sets $\{\varphi : \|\varphi\|_\beta \leq 1\}$ and $\{\varphi : \|\varphi\|_\beta \leq 1\}$ only differ by additive constants, so that one has indeed $d_\beta = \rho_\beta$. \square

Remark 2.2. Note that of course $d_\beta = \rho_\beta$ only for probability measures, or at least positive measures of equal mass. Otherwise, d_β is $+\infty$ in general, while ρ_β need not be.

3. Proof of main theorem

Theorem 3.1. *If Assumptions 1 and 2 hold there exists an $\bar{\alpha} \in (0, 1)$ and $\beta > 0$ such that*

$$\|\mathcal{P}\varphi\|_\beta \leq \bar{\alpha}\|\varphi\|_\beta.$$

Actually, setting $\gamma_0 = \gamma + 2K/R < 1$, for any $\alpha_0 \in (0, \alpha)$ one can choose $\beta = \alpha_0/K$ and $\bar{\alpha} = (1 - \alpha + \alpha_0) \vee (2 + R\beta\gamma_0)/(2 + R\beta)$.

Proof. Fix a test function φ with $\|\varphi\|_\beta \leq 1$. By Lemma 2.1, we can assume without loss of generality that one also has $\|\varphi\|_\beta \leq 1$. The claim then follows if we can exhibit $\bar{\alpha} < 1$ so that

$$|\mathcal{P}\varphi(x) - \mathcal{P}\varphi(y)| \leq \bar{\alpha}d_\beta(x, y).$$

If $x = y$, the claim is true. Henceforth we assume $x \neq y$. We begin by assuming that x and y are such that

$$V(x) + V(y) \geq R. \tag{4}$$

Fixing γ_0 as in the statement of the theorem, for any $\beta > 0$ we set $\gamma_1 = (2 + \beta R\gamma_0)/(2 + \beta R)$. Observe that for $\beta \in (0, 1)$ and $R > 0$, one has $\gamma_1 \in (\gamma_0, 1)$. With these choices, we have from (1) and (2) the bound

$$\begin{aligned} |\mathcal{P}\varphi(x) - \mathcal{P}\varphi(y)| &\leq 2 + \beta\mathcal{P}V(x) + \beta\mathcal{P}V(y) \\ &\leq 2 + \beta\gamma V(x) + \beta\gamma V(y) + 2\beta K \\ &\leq 2 + \beta\gamma_0 V(x) + \beta\gamma_0 V(y) \\ &\leq 2\gamma_1 + \beta\gamma_1 V(x) + \beta\gamma_1 V(y) = \gamma_1 d_\beta(x, y). \end{aligned}$$

The third line follows from our choice of γ_0 and the fact that by (4) we know that $2K \leq (\gamma_0 - \gamma)(V(x) + V(y))$. The last line follows from the fact that $2(1 - \gamma_1) = \beta R(\gamma_1 - \gamma_0) \leq \beta(\gamma_1 - \gamma_0)(V(x) + V(y))$ given our choice of γ_1 . We emphasise that up to now β could be any positive number; only the precise value of γ_1 depends on it (and gets “worse” for small values of β). The second part of the proof will determine a choice of $\beta > 0$.

Now consider the case of x and y such that $V(x) + V(y) \leq R$ and hence $x, y \in \mathcal{C}$. For such x and y we define the Markov transition $\tilde{\mathcal{P}}$ by $\tilde{\mathcal{P}}(x, \cdot) = \frac{1}{1-\alpha}\mathcal{P}(x, \cdot) - \frac{\alpha}{1-\alpha}\nu(\cdot)$. Now we have $\mathcal{P}\varphi(x) = (1 - \alpha)\tilde{\mathcal{P}}\varphi(x) + \alpha \int \varphi d\nu$ and $\mathcal{P}\varphi(y) = (1 - \alpha)\tilde{\mathcal{P}}\varphi(y) + \alpha \int \varphi d\nu$. Subtracting the second of these expressions from the first and using that since V is a non-negative function $\tilde{\mathcal{P}}V(x) \leq \frac{1}{1-\alpha}\mathcal{P}V(x)$ produces

$$\begin{aligned} |\mathcal{P}\varphi(x) - \mathcal{P}\varphi(y)| &= (1 - \alpha)|\tilde{\mathcal{P}}\varphi(x) - \tilde{\mathcal{P}}\varphi(y)| \\ &\leq (1 - \alpha)2 + (1 - \alpha)\beta(\tilde{\mathcal{P}}V(x) + \tilde{\mathcal{P}}V(y)) \\ &\leq (1 - \alpha)2 + \beta(\mathcal{P}V(x) + \mathcal{P}V(y)) \\ &\leq (1 - \alpha)2 + \gamma\beta V(x) + \gamma\beta V(y) + 2\beta K. \end{aligned}$$

Hence fixing $\beta = \alpha_0/K$ for any $\alpha_0 \in (0, \alpha)$ and setting $\gamma_2 = (1 - (\alpha - \alpha_0)) \vee \gamma \in (0, 1)$ produces

$$\begin{aligned} |\mathcal{P}\varphi(x) - \mathcal{P}\varphi(y)| &\leq 2(1 - (\alpha - \alpha_0)) + \gamma\beta V(x) + \gamma\beta V(y) \\ &\leq \gamma_2 d_\beta(x, y). \end{aligned}$$

Setting $\bar{\alpha} = \gamma_1 \vee \gamma_2$ and recalling that $\gamma_1 \geq \gamma$ concludes the proof. \square

Theorem 1.3 now follows as a corollary since $d_\beta = \rho_\beta$ and d_β is the norm dual to $\|\cdot\|_\beta$. In order to conclude that Theorem 1.2 holds, it only remains to show that our assumptions imply that an invariant measure μ_\star actually exists and that the integral of V with respect to μ_\star is finite.

3.1. Existence of an invariant measure

We have already shown that Assumptions 1 and 2 allow to prove that for some $\beta > 0$, \mathcal{P} is a strict contraction in the weighted total variation metric ρ_β defined by (3). We now show that the same assumptions are also sufficient to ensure the existence of an invariant measure:

Theorem 3.2. *If Assumptions 1 and 2 hold then there exists a probability measure μ_∞ on \mathbf{X} such that $\int V d\mu_\infty < \infty$ and which is invariant in that $\mathcal{P}\mu_\infty = \mu_\infty$.*

Proof. Fixing any $x \in \mathbf{X}$, for $n \in \mathbf{N}$ define $\mu_n = \mathcal{P}^n \delta_x$. By Theorem 1.3, we know that for some $\bar{\alpha} \in (0, 1)$ and some $\beta > 0$,

$$\rho_\beta(\mu_{n+1}, \mu_n) \leq \bar{\alpha}^n \rho_\beta(\mu_1, \delta_x).$$

Hence, μ_n is a Cauchy sequence. Since ρ_β is complete for the space of probability measures integrating V (because the total variation distance is complete for the space of measures with finite mass) there exists a probability measure μ_∞ so that $\rho_\beta(\mu_n, \mu_\infty) \rightarrow 0$ as $n \rightarrow \infty$. Since this implies that $\mu_n \rightarrow \mu_\infty$ in total variation and \mathcal{P} is always a contraction

in the total variation distance, it follows that $\mathcal{P}\mu_\infty = \lim \mathcal{P}\mu_n = \lim \mu_{n+1} = \mu_\infty$ as required. \square

3.2. A slightly different set of assumptions

Many results in the theory of Harris chains results are proved under a slightly different set of assumptions. The Lyapunov function condition in Assumption 1 is replaced with the following:

Assumption 3. There exists a function $V : \mathbf{X} \rightarrow [1, \infty)$ and constants $b \geq 0$, $\tilde{\gamma} \in (0, 1)$ and a subset $S \subset \mathbf{X}$ such that

$$(\mathcal{P}V)(x) \leq \tilde{\gamma}V(x) + b\mathbf{1}_S(x), \quad (5)$$

for all $x \in \mathbf{X}$.

Clearly Assumption 3 implies Assumption 1 with $K = b$. The question is whether Assumption 2 holds with that choice of K and with \mathcal{C} defined as in Assumption 2. If it does then our main theorem holds. However, Assumption 3 is most naturally paired with the following modified version of Assumption 2.

Assumption 4. There exists a constant $\tilde{\alpha} \in (0, 1]$ and a probability measure $\tilde{\nu}$ so that the lower bound

$$\inf_{x \in S} \mathcal{P}(x, \cdot) \geq \tilde{\alpha}\tilde{\nu}(\cdot)$$

holds. Here, the set S is the same as in Assumption 3.

It is relatively clear that Assumptions 1 and 2 together imply Assumptions 3 and 4. In particular, if one picks a $\tilde{\gamma} \in (\gamma, 1)$ sufficiently close to one, then $R \geq K/(\tilde{\gamma} - \gamma)$ and setting $S = \{x : V(x) \leq K\}$ we see that the desired implication holds.

Remark 3.3. In general, one cannot hope for Assumptions 4 and 3 to imply Assumptions 1 and 2 and hence the existence of a spectral gap without any further assumptions. A trivial example is given by $\mathbf{X} = \{0, 1\}$ with the (deterministic) transition probabilities $\mathcal{P}(x, \cdot) = \delta_{1-x}$. This Markov operator has spectrum $\{-1, 1\}$ and has therefore no spectral gap. On the other hand, Assumptions 4 and 3 are satisfied with $\tilde{\alpha} = 1$, $\tilde{\gamma} = 1/2$, and $b = 3/2$ if one makes for example the choice $S = \{0\}$, $\tilde{\nu} = \delta_1$, and $V(x) = 1 + x$.

In spite of the preceding remark, we are now going to show that Assumptions 4 and 3 are essentially equivalent to Assumptions 1 and 2 from the previous section. More precisely, for $N > 0$, define the ‘‘averaged’’ Markov operator

$$\mathcal{Q} = \frac{1}{N+1} \sum_{k=0}^N \mathcal{P}^k.$$

Then we have:

Theorem 3.4. *If \mathcal{P} satisfies Assumptions 4 and 3, then there exists a choice of N such that \mathcal{Q} satisfies Assumptions 1 and 2.*

Proof. Fix some arbitrary R with $R > 2b/(1 - \tilde{\gamma})$. Our aim is to show that we can find $N > 0$, a probability measure ν and a constant $\alpha > 0$ such that $\mathcal{Q}(x, \cdot) \geq \alpha\nu(\cdot)$ for every x with $V(x) \leq R$.

Iterating (5), we find that one has the bound

$$1 \leq \mathcal{P}^{n+1}V \leq \tilde{\gamma}^{n+1}V + b \sum_{k=0}^n \tilde{\gamma}^k \mathcal{P}^{n-k} \mathbf{1}_S, \quad (6)$$

so that, on the set $\mathcal{S}_n = \{x : V(x) \leq \tilde{\gamma}^{-n-1}/2\}$, one has the lower bound

$$\inf_{x \in \mathcal{S}_n} \sum_{k=0}^n \tilde{\gamma}^k \mathcal{P}^{n-k}(x, S) \geq \frac{1}{2b}. \quad (7)$$

In particular this implies that for every $x \in \mathbf{X}$, there exists n such that $\mathcal{P}^n(x, S) > 0$. Combining this with our two assumptions shows that $\int V(x) \tilde{\nu}(dx) = C < \infty$ so that, integrating (6) with respect to $\tilde{\nu}$, we obtain

$$1 \leq C\tilde{\gamma}^{n+1} + b \sum_{k=0}^n \tilde{\gamma}^k (\mathcal{P}^{n-k} \tilde{\nu})(S).$$

Choosing n sufficiently large then implies the existence of some $\ell > 0$ such that $(\mathcal{P}^{\ell-1} \tilde{\nu})(S) > 0$. Combining this with Assumption 3 shows that there exists $\hat{\alpha} > 0$ such that $\mathcal{P}^\ell \tilde{\nu} \geq \hat{\alpha} \tilde{\nu}$. Setting now $\nu = \frac{1}{\ell} \sum_{k=0}^{\ell-1} \mathcal{P}^k \tilde{\nu}$, it follows that one has the bound

$$\mathcal{P}\nu = \frac{1}{\ell} \sum_{k=1}^{\ell-1} \mathcal{P}^k \tilde{\nu} + \frac{1}{\ell} \mathcal{P}^\ell \tilde{\nu} \geq \frac{1}{\ell} \sum_{k=1}^{\ell-1} \mathcal{P}^k \tilde{\nu} + \frac{\hat{\alpha}}{\ell} \tilde{\nu} \geq \hat{\alpha} \nu.$$

In particular, this implies that for every $m \geq 1$ there exists a constant α_m such that the lower bound

$$\inf_{x \in S} \sum_{k=m}^{m+\ell} \mathcal{P}^k(x, \cdot) \geq \alpha_m \nu(\cdot) \quad (8)$$

holds. Let now n be sufficiently large such that $\tilde{\gamma}^{-n-1}/2 \geq R$ and set $N = n + 1 + \ell$. Combining (7) and (8) then yields the desired result. \square

Remark 3.5. Keeping track of the constants appearing in the proof of the previous result, we see that one can choose for example any integer N such that

$$N > 1 + \log\left(\frac{2b}{1 - \tilde{\gamma}} \int V(x) \tilde{\nu}(dx)\right) / \log \tilde{\gamma}.$$

Acknowledgments

MH acknowledges the support of an EPSRC Advanced Research Fellowship (grant number EP/D071593/1) and JCM the support of NSF through a PECASE award (DMS-0449910) and the support of the Sloan foundation through a Sloan foundation fellowship.

References

- [Cha89] K. S. CHAN. A note on the geometric ergodicity of a Markov chain. *Adv. in Appl. Probab.* **21**, no. 3, (1989), 702–704.
- [DMLM03] P. DEL MORAL, M. LEDOUX, and L. MICLO. On contraction properties of Markov kernels. *Probability Theory and Related Fields* **126**, no. 3, (2003), 395–420.
- [DMR04] R. DOUC, E. MOULINES, and J. S. ROSENTHAL. Quantitative bounds on convergence of time-inhomogeneous Markov chains. *The Annals of Applied Probability* **14**, no. 4, (2004), 1643–1665.
- [Har56] T. E. HARRIS. The existence of stationary measures for certain Markov processes. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, 1954–1955, vol. II*, 113–124. University of California Press, Berkeley and Los Angeles, 1956.
- [Has80] R. Z. HAS’MINSKII. *Stochastic stability of differential equations*, vol. 7 of *Monographs and Textbooks on Mechanics of Solids and Fluids: Mechanics and Analysis*. Sijthoff & Noordhoff, Alphen aan den Rijn, 1980. Translated from the Russian by D. Louvish.
- [HM08] M. HAIRER and J. MATTINGLY. Spectral gaps in Wasserstein distances and the 2D stochastic Navier-Stokes equations, 2008. To be published in *Ann. Probab.*
- [KM05] I. KONTOYIANNIS and S. P. MEYN. Large deviations asymptotics and the spectral theory of multiplicatively regular markov processes. *Electronic Journal of Probability* **10**, (2005), no. 3, 61123 (electronic).
- [MT92] S. P. MEYN and R. L. TWEEDIE. Stability of Markovian processes. I. Criteria for discrete-time chains. *Adv. in Appl. Probab.* **24**, no. 3, (1992), 542–574.
- [MT93] S. P. MEYN and R. L. TWEEDIE. *Markov chains and stochastic stability*. Communications and Control Engineering Series. Springer-Verlag London Ltd., London, 1993.
- [Num84] E. NUMMELIN. *General irreducible Markov chains and nonnegative operators*, vol. 83 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1984.

Martin Hairer

Mathematics Institute, The University of Warwick, CV4 7AL, UK

e-mail: M.Hairer@Warwick.ac.uk

Jonathan C. Mattingly

Mathematics Department, Center of Theoretical and Mathematical Science, Department of Statistical Science, and Center of Nonlinear and Complex Systems, Duke University, Durham NC, USA

e-mail: jonm@math.duke.edu