Received: 13 August 2024 | Revised: 31 January 2025 | Accepted: 18 March 2025

DOI: 10.1002/qj.4983

Quarterly Journal of the Royal Meteorological Society RMetS

RESEARCH ARTICLE

A stochastic differential equation framework for gravity wave parametrisation with testing in an idealised setting

K. Xie | **M. Ewetola** | **J. G. Esler**

Department of Mathematics, University College London, London, UK

Correspondence
K. Xie, Department of Mathematics, University College London, London, UK.
Email: kelvin.xie.23@ucl.ac.uk

Funding information
Natural Environment Research Council, Grant/Award Number: NE/S00985X/1

Abstract

Parametrisations of unresolved gravity waves used in general circulation models can be made more computationally efficient by introducing a stochastic component to the forcing. An additional advantage of introducing stochasticity is that intermittency associated with the scheme could be tuned to resemble the intermittency of observed gravity wave sources, and could therefore act to improve the physical fidelity of the scheme. Here, it is argued that using stochastic differential equations (SDEs) to drive the stochastic component provides a natural general framework to develop such schemes. The Holton–Lindzen–Plumb model of the quasi-biennial oscillation (QBO) is used to demonstrate the flexibility of the approach. The QBO generated in a (computationally expensive) deterministic broadband multiwave simulation is accurately reproduced using a number of quite different (cheap) stochastic schemes. The method of stochastic averaging is used to prove a matching result that shows that a wide class of such schemes, driven by different SDEs, can each reproduce the deterministic QBO provided that the characteristic time-scale τ of the SDEs is sufficiently short. However, each scheme has different intermittency properties: as τ is increased, their QBOs are shown to diverge, despite the time-averaged source spectrum in each case remaining unchanged. The SDE framework therefore provides great flexibility to tune a stochastic parametrisation to match observed intermittencies, meaning that future parametrisations can be developed that can account for non-steady gravity wave forcing in a physically consistent manner.

KEYWORDS

gravity waves, parametrisation, stochastic differential equations

1 | INTRODUCTION

In the Earth's stratosphere, momentum and temperature fluxes due to gravity waves are known to make significant contributions to the budget of the mean circulation. An important part of the budget is due to gravity waves with horizontal spatial scales comparable to or less than the resolution of state-of-the-art general circulation models. Therefore, in order to reproduce fundamental features of stratospheric dynamics such as the quasi-biennial oscillation (QBO; e.g., Baldwin *et al.*, 2001), the sources, propagation, and dissipation of these subgrid waves must be parametrised in models. At present (Plougonven *et al.*, 2020), the source terms used in these parametrisations are not well constrained by observations and are usually "tuned" to some extent in models to produce the

desired results. In future, as improved observations begin to better constrain the source terms, there is a clear need to improve understanding of how effects associated with the statistical properties of the wave sources should best be accounted for in parametrisations. For example, it is already understood that account needs to be taken of intermittency in wave sources, leading to additional tuning parameters in model parametrisations, such as the "efficiency parameter" used in the scheme of Alexander and Dunkerton (1999). The efficiency parameter of Alexander and Dunkerton (1999) is a constant correction factor, usually taken to be approximately 0.1, by which the calculated (non-intermittent) forcing is multiplied in order to capture the effect of intermittency.

An interesting innovation in the development of parametrisation schemes is the introduction of stochasticity (Eckermann, 2011; Lott *et al.*, 2012; Piani *et al.*, 2004) in the treatment of the source terms. The idea is that instead of treating a broadband wave source by discretising the wave number–frequency $(\boldsymbol{k}, \omega)$ spectrum and modelling the propagation and dissipation of each discretised wave at every model time step—a so-called multi-wave scheme (e.g., Alexander & Dunkerton, 1999; Lindzen, 1981; Richter *et al.*, 2010; Song & Chun, 2008)—the spectrum is sampled randomly. Since the stochastic scheme needs to compute fluxes for a relatively small number of waves at each model step compared with the multi-wave scheme, large computational savings are possible. A further potential advantage, highlighted by Lott and co-workers (Lott *et al.*, 2012; Lott & Guez, 2013), is that the intermittent behaviour of the stochastic scheme could somehow be tuned to match the intermittency of the observed source, negating the need for a tuning parameter such as that used by Alexander and Dunkerton (1999). How this is best achieved is an open question that is a key motivation for this work.

The central thesis of the present work is that stochastic differential equations (SDEs) provide a natural and flexible framework for the future development of stochastic parametrisation schemes. SDEs offer the following advantages:

1. SDEs generate random processes in continuous time, meaning that grid and time-step independent equations can be used to define the scheme, which is an evidently desirable property for reproducible modelling. Note that the scheme by Eckermann (2011), which is based upon re-sampling the forcing frequency at every time-step, is necessarily time-step dependent. Numerical noise associated with the time-step dependence of the scheme of Eckermann (2011) led subsequent researchers such as Lott and co-workers (Lott *et al.*, 2012; Lott & Guez, 2013) to introduce more waves into the scheme, with their impact being smoothed over multiple time-steps using a first-order autoregressive (AR(1)) process. In the following we will show that the SDE approach naturally leads to a more stable scheme, without the use of additional waves.
2. The freedom to choose the drift and noise functions in an SDE means that, potentially, SDEs can be found that generate time series with the desired intermittent behaviour. This flexibility means that SDEs have great potential to match the intermittency of observed sources.
3. The analysis of SDEs is a mature field; in particular, asymptotic results exist that can be exploited to give mathematical insight into the best approaches to parametrisation. For example, the method of homogenisation (e.g., Pavliotis & Stuart, 2008) can be used to evaluate the effects of fast stochastic processes on slow dynamics, and has been used recently in the QBO setting by Ewetola and Esler (2024).

The aim herein is to develop the SDE framework for parametrisation in as simple a model as possible, namely the Holton–Lindzen–Plumb (HLP) model of the QBO Lindzen and Holton (1968); Holton and Lindzen (1972); Plumb (1977); also, see the review of Renaud and Venaille (2020). Choosing this simplified framework allows concrete results to be established and verified. Despite using this simplification, however, it will be made clear that there is no obvious obstacle to extending the same approach to more comprehensive models. The HLP model is first adapted from its standard formulation, in which the QBO is forced by two monochromatic waves generated at the model lower boundary, to a broadband multi-wave version (BBM-HLP hereafter) in which the forcing is due to a more realistic continuous spectrum of waves in both wave number and frequency space. A recent systematic study (Léard *et al.*, 2020) has confirmed that broadening the source spectrum in this manner does not change the essential nature of the QBO, except to make the periodic signal more robust, increasing the range of values of the Reynolds number that lead to stable QBO-like oscillations (as opposed to frequency locking or chaotic behaviour). In addition, broader source spectrums resulted in typically longer QBO periods and increased amplitudes. Broadening the source spectrum has also been found (Garfinkel *et al.*, 2022) to lead to longer periods and larger QBO amplitudes in a more comprehensive model that uses the Alexander and Dunkerton (1999) scheme.

In the following, it will first be demonstrated that the QBOs generated in BBM-HLP can be reproduced to high accuracy using a stochastic scheme with a small number of waves (two here), controlled by carefully chosen SDEs. Our main mathematical results, which exploit

the principle of stochastic averaging (e.g., Pavliotis & Stuart, 2008, chap. 10), show that the QBO in the stochastic model converges accurately to that in BBM-HLP when the time-scale τ_* of the SDEs in question becomes much shorter than the QBO time-scale T_*. Further, for a given broadband forcing spectrum, it is discovered that there is considerable flexibility in the choice of SDEs that can be used. In fact, the SDEs are constrained only by a matching condition that relates their associated invariant density (i.e., their probability density function in the long-time limit) to the broadband forcing spectrum to be emulated.

The flexibility in the choice of SDEs presents a natural framework to seek stochastic parametrisations that can go beyond just reproducing the QBO generated by the deterministic BBM-HLP model, in order to simulate QBOs due to more realistic intermittent broadband wave forcing. The importance of intermittency effects on the QBO is clear from the numerical calculations of Couston *et al.* (2018) in a two-dimensional stratified Boussinesq system. These calculations compare the QBO that emerges from an intermittent wave source due to a "tropospheric" lower layer of active convection with that from a non-intermittent wave forcing with the same (time-averaged) energy spectrum in a "stratosphere-only" model forced at the bottom boundary. The results show a dramatic difference in the emergent QBO between the simulations—the convectively forced model (their M_1) has a QBO with significantly smaller amplitude and around half the period of that in the steadily forced model (their M_2). The M_1 QBO also exhibits considerable variability not captured by M_2.

A recent step towards understanding the impact of intermittent forcing on the QBO has been made by Ewetola and Esler (2024). There, intermittent forcing is introduced to the standard (two-wave) HLP model, by allowing the amplitude of each of the two forcing waves to evolve stochastically in time according to an SDE. The main mathematical result is that the impact of this amplitude intermittency on the QBO can be captured by a single intermittency parameter λ, which depends only on the details of the SDE controlling the wave amplitudes. Extending this result to broadband forcing, in the general setting, is a complex and lengthy problem that we do not attempt to solve here. To make a first step in this direction, our aim here, using numerical simulations with stochastic schemes based on different families of SDEs, is to demonstrate that the intermittency effects associated with each scheme, which are introduced by varying the SDE time-scale τ_*, can have quite different impacts on the QBO period and amplitude. These proof-of-concept simulations demonstrate that solving the optimisation problem of finding the SDE-based parametrisation that best reproduces the QBO generated by a specified intermittent broadband wave source could be a highly promising avenue for the future development of parametrisations.

The plan of this work is as follows. First, the HLP model and its broadband multi-wave extension (BBM-HLP) are described in Section 2. Section 3 then explains the theory of how QBOs generated in BBM-HLP can be reproduced using a stochastic parametrisation based on an adapted version of the two-wave HLP coupled to an SDE (broadband stochastic HLP, BB-sHLP hereafter). The main mathematical result, which defines the family of SDEs that converge to the correct QBO when τ_* is small, is presented and the implications are discussed. In Section 4, proof-of-concept numerical simulations are described that show the convergence of three different families of BB-sHLP simulations to the BBM-HLP QBO, and the contrasting intermittency effects within each family are highlighted. Finally, in Section 5, the outlook for future work is discussed and conclusions are drawn.

2 | THE BBM-HLP MODEL

The HLP model is, arguably, the simplest system to capture the fundamental physics of the QBO. Its parent system is a stratified Boussinesq fluid, with constant buoyancy frequency N and kinematic (eddy) viscosity ν, forced at the bottom boundary by leftward and rightward-propagating waves of amplitude a_*, wave number k_*, and angular frequencies $\pm\omega_*$ (asterisks will be used throughout to denote dimensional quantities). The gravity waves excited by the oscillating lower boundary propagate upwards and are eventually dissipated in the fluid interior by thermal damping at rate α. Using the Wentzel–Kramers–Brillouin method to approximate the momentum flux convergence $-(\overline{u'w'})_z$ due to these waves results in the following non-dimensional integro-differential equation that describes the time evolution of the zonal mean zonal wind $U(z,t)$:

$$\partial_t U = -\sum_{\pm} \partial_z \left\{ \pm \exp\left[-\int_0^z \frac{1}{(U(z',t) \mp 1)^2}\, dz' \right] \right\} + \frac{1}{\mathrm{Re}}\, \partial_z^2 U. \tag{1}$$

Note that the two terms in the sum correspond to the momentum flux convergence due to the rightward and leftward-propagating waves. The complete derivation and discussion of the physical basis of the associated approximations leading to Equation (1) is reviewed in Renaud and Venaille (2020). Following Renaud and Venaille (2020), the dimensional velocity and height and time scales are respectively $U_* = \omega_*/k_*$, $Z_* = \omega_*^2/\alpha k_* N$, and $T_* = 2\omega_*^2/\alpha N^2 k_*^2 a_*^2$. The height scale Z_*, known as the attenuation length, is the scale over which waves

of frequency ω_* and wave number k_* are damped. The time-scale T_*, known as the streaming time-scale, measures the time-scale over which the wave momentum flux interacts with the mean flow. Vallis (2017) discusses approximate dimensional values for these characteristic scales: the choices $U_* \approx 25\,\mathrm{m}\cdot\mathrm{s}^{-1}$, $Z_* \approx 6\,\mathrm{km}$, and $T_* \approx 160$ days result in the HLP QBO having period, amplitude, and vertical scale close to that observed in the stratosphere. A detailed study of the role of the only free parameter, the Reynolds number $\mathrm{Re} = \omega_*^2 a_*^2/(2\alpha\nu)$, in the QBO dynamics of HLP is given in Renaud *et al.* (2019). To summarise, increasing Re past a critical value has the effect of destabilising the $U = 0$ solution to the HLP equation, and the system enters a regime of regular QBO-like oscillations. As Re is increased further, there is a break down of periodicity, with a transition to quasi-periodic oscillations, then to frequency locking and finally onto chaos.

To extend HLP to include a broad spectrum of waves, the (non-dimensional) lower boundary condition applied at $z = h(x,t)$ in the underlying Boussinesq system must be modified. In the derivation of HLP in Equation (1), $h(x,t)$ consists of a leftward and rightward-propagating wave

$$h(x,t) = \mathcal{R}\left\{\mathrm{e}^{\mathrm{i}(x-t/\varepsilon)} + \mathrm{e}^{\mathrm{i}(x+t/\varepsilon)}\right\}, \tag{2}$$

where $\mathcal{R}$ denotes the real part and $\varepsilon = (\omega_* T_*)^{-1} \ll 1$. The small parameter ε, which is central to the derivation of HLP, allows for a natural separation of time-scales into a fast wave time-scale described by the fast time variable $t' = t/\varepsilon$ and the slow QBO time-scale described by the slow time variable t. Using the dimensional values given by Vallis (2017) gives $\varepsilon \approx 0.014$. In a more general setting, the evolving lower boundary height $h(x,t',t)$ can be interpreted as a function of both t' and t, in the spirit of the method of multiple scales (e.g., Holmes, 2013). The associated spectral power density can be defined to be

$$A(k,\omega,t) = \lim_{L,T\to\infty} \frac{4}{LT}\,|\hat{h}_{LT}(k,\omega,t)|^2, \tag{3}$$

where

$$\hat{h}_{LT}(k,\omega,t) = \frac{1}{2\pi}\int_0^T\int_0^L h(x,t',t)\mathrm{e}^{-\mathrm{i}(kx-\omega t')}\,dx\,dt'. \tag{4}$$

The separation of time-scales means that the $T\to\infty$ limit in Equation (3) is formally long compared with the short t' time-scale, but is short compared with the long t time-scale, which allows for the treatment of a fully intermittent (i.e., time-evolving) spectrum of waves on the QBO time-scale. In this work, however, we will be primarily concerned with non-intermittent spectra, and henceforth consider only time-independent spectra with spectral densities $A(k,\omega)$.

Extending the derivation of Renaud and Venaille (2020) to this broadband setting is straightforward, and the HLP equation, Equation (1), is modified to become the BBM-HLP:

$$\begin{aligned}\partial_t U = -\partial_z\Bigg\{&\int_{-\infty}^{\infty}\int_0^{\infty}\omega A(k,\omega)\\ &\times\exp\left[-\int_0^z \frac{k}{(kU(z',t)-\omega)^2}\,dz'\right]dk\,d\omega\Bigg\}\\ &+\frac{1}{\mathrm{Re}}\,\partial_z^2 U,\end{aligned} \tag{5}$$

where Equation (5) is non-dimensionalised using the same characteristic scales as earlier herein, based on the forcing spectrum having a characteristic frequency ω_* and wave number k_*. The HLP case is recovered from Equation (5) when h is given by Equation (2), because in this case the spectral density is (in terms of the Dirac delta) given by $A(k,\omega) = \delta(k-1)(\delta(\omega-1)+\delta(\omega+1))$. It is to be emphasised that Equation (5) requires the evaluation of a triple integral, making explicit the point made in Section 1 that deterministic broadband multi-wave parametrisations of gravity waves are in general computationally expensive.

3 | THE BB-SHLP MODEL

Next we will introduce a broad class of stochastic HLP models each of which is capable of reproducing the QBO behaviour of BBM-HLP, Equation (5), in the limit in which the stochastic process is fast. The stochastic models have the advantages of being simpler to implement and of being significantly cheaper computationally than BBM-HLP is. The idea is to return to forcing the QBO at $z = h$ with two propagating waves, but to allow the amplitudes, wave numbers, and frequencies $\boldsymbol{Q}_t^{\pm} = (A_t^{\pm}, K_t^{\pm}, \text{ and } \Omega_t^{\pm})^T$ of the two waves to evolve according to a joint stochastic process defined on the domain $\mathbb{R}\times\mathbb{R}_+\times\mathbb{R}$. It should be stated clearly at this point that the choice of two waves is just a convenience for studying systems that are close to the original HLP set-up. Similar formulations are possible with any number of evolving stochastic waves, with the optimal choice determined by the complexity of the forcing spectrum $A(k,\omega)$ to be captured.

In the stochastic model, the boundary condition, Equation (2), is replaced by

$$h(x,t) = \mathcal{R}\left\{\sum_{\pm} A_t^{\pm}\exp\left(\mathrm{i}K_t^{\pm}x - \frac{\mathrm{i}}{\varepsilon}\int^t \Omega_s^{\pm}\,ds\right)\right\}. \tag{6}$$

Provided that the (dimensional) time-scale of the stochastic processes satisfies $\tau_* \gg \omega_*^{-1}$, the derivation of the

resulting equation remains essentially unaltered from that described in Renaud and Venaille (2020) and results in the BB-sHLP model:

$$\partial_t U = -\sum_{\pm} \partial_z \left\{ \Omega_t^{\pm} |A_t^{\pm}|^2 \times \exp\left[-\int_0^z \frac{K_t^{\pm}}{(K_t^{\pm} U(z',t) - \Omega_t^{\pm})^2}\, dz' \right] \right\} + \frac{1}{\mathrm{Re}}\, \partial_z^2 U. \tag{7}$$

Notice that the deterministic HLP, Equation (1), is recovered when $(A_t^{\pm}, K_t^{\pm}, \Omega_t^{\pm}) = (1, 1, \pm 1)$.

The main result here is to highlight a close relationship between BB-sHLP, Equation (7), and BBM-HLP, Equation (5). This result is formally valid in the limit in which the time-scale of the stochastic process τ_* is much shorter than the QBO streaming time-scale ($\tau_* \ll T_*$), and the stochastic processes $\boldsymbol{Q}_t^{\pm}$ have stationary probability densities $p_s^{\pm}(a, k, \omega)$, known as invariant densities. In this limit, assuming a standard ergodicity property for $\boldsymbol{Q}_t^{\pm}$ (e.g., Pavliotis & Stuart, 2008, chap. 6.4 discussion), the principle of stochastic averaging (Pavliotis & Stuart, 2008, chap. 10) applies, because the stochastic process $\boldsymbol{Q}_t^{\pm}$, will effectively sample its entire state space before $U(z, t)$ evolves significantly. Applying stochastic averaging to Equation (7), following the methodology of Pavliotis and Stuart (2008), results in

$$\partial_t U = -\sum_{\pm} \partial_z \left\{ \int_{-\infty}^{\infty} \int_0^{\infty} \int_{-\infty}^{\infty} \omega a^2 \times \exp\left[-\int_0^z \frac{k}{(kU(z',t) - \omega)^2}\, dz' \right] \times p_s^{\pm}(a, k, \omega)\, da\, dk\, d\omega \right\} + \frac{1}{\mathrm{Re}}\, \partial_z^2 U. \tag{8}$$

Equation (8) is obtained from Equation (7) because stochastic averaging involves simply replacing the stochastic variables in Equation (7) with their time-averaged values. For example, for any function $f(A_t^+, K_t^+, \Omega_t^+)$ of the ergodic process $\boldsymbol{Q}_t^+$, its time average is given by

$$\langle f \rangle = \int_{-\infty}^{\infty} \int_0^{\infty} \int_{-\infty}^{\infty} f(a, k, \omega) p^+(a, k, \omega)\, da\, dk\, d\omega. \tag{9}$$

This time-average formula holds because $p_s^+(a, k, \omega)\, da\, dk\, d\omega$ is by definition the proportion of time that the process $\boldsymbol{Q}_t^+$ spends in the interval $[a, a + da] \times [k, k + dk] \times [\omega, \omega + d\omega]$. The result, Equation (8), follows from the application of Equations (9) to (7).

Comparing Equations (8) and (5), it follows that, if SDEs governing $\boldsymbol{Q}_t^{\pm}$ can be found so that their invariant densities $p_s^{\pm}$ satisfy the "matching condition"

$$\sum_{\pm} \int_{-\infty}^{\infty} a^2 p_s^{\pm}(a, k, \omega)\, da = A(k, \omega), \tag{10}$$

then BB-sHLP, Equation (7), should reproduce the behaviour of BBM-HLP, Equation (5), in the limit $\tau_* \ll T_*$. The matching condition Equation (10) has an alternative interpretation as a conditional expectation:

$$\sum_{\pm} \mathbb{E}((A_t^{\pm})^2 \mid (K_t^{\pm}, \Omega_t^{\pm}) = (k, \omega)) = A(k, \omega). \tag{11}$$

It turns out that the matching condition is satisfied by a broad class of SDEs. To simplify the following discussion, however, a restricted class of SDEs will be considered in which the evolution of $A_t^{\pm}$ is slaved to the wave number and frequency $\boldsymbol{\Omega}_t^{\pm} = (K_t^{\pm}, \Omega_t^{\pm})^T$ according to

$$A_t^{\pm} = G^{\pm}(\boldsymbol{\Omega}_t^{\pm}), \tag{12}$$

where the "amplitude modulation" functions $G^{\pm}(\cdot)$ must be chosen with the matching condition in mind. Further, the wave number and frequency $\boldsymbol{\Omega}_t^{\pm}$ are governed by a type of SDE known as an overdamped Langevin (ODL) equation:

$$d\boldsymbol{\Omega}_t^{\pm} = -\frac{\nabla_{\boldsymbol{\omega}} V^{\pm}(\boldsymbol{\Omega}_t^{\pm})}{\tau}\, dt + \left(\frac{2}{\tau}\right)^{1/2} d\boldsymbol{B}_t^{\pm}. \tag{13}$$

Here, the functions $V^{\pm}(\boldsymbol{\omega})$, known as potential functions, are prescribed functions of $\boldsymbol{\omega} = (k, \omega)^T$, $\nabla_{\boldsymbol{\omega}} \equiv (\partial_k, \partial_\omega)^T$ and $\tau = \tau_*/T_*$ is the non-dimensional time-scale of the process (required here to satisfy $\varepsilon \ll \tau \ll 1$), and $\boldsymbol{B}_t^{\pm}$ are independent Brownian processes in two dimensions. The ODL equation, Equation (13), is convenient to work with here because, provided that the potential functions satisfy the conditions required for normalisation (essentially that $V^{\pm}(\boldsymbol{\omega}) \to \infty$ as $|\boldsymbol{\omega}| \to \infty$ sufficiently quickly), its invariant density is proportional to $\exp(-V^{\pm}(\boldsymbol{\omega}))$ (e.g., Gardiner, 2009).

The restricted class of SDE defined by Equations (12) and (13) represents a considerable simplification compared with the general case, but nevertheless offers considerable flexibility in satisfying the matching condition. Using the ODL result, the invariant density for $\boldsymbol{Q}_t^{\pm}$ is given by

$$p_s^{\pm}(a, \boldsymbol{\omega}) = p_0^{\pm} \delta(a - G^{\pm}(\boldsymbol{\omega})) \exp(-V^{\pm}(\boldsymbol{\omega})).$$

Here, $\delta(\cdot)$ is the Dirac delta and

$$p_0^{\pm} = \left(\int_{\mathbb{R}_+ \times \mathbb{R}} \exp(-V^{\pm}(\boldsymbol{\omega}))\, d\boldsymbol{\omega} \right)^{-1}$$

TABLE 1 Choices for the amplitude modulation and potential functions for three schemes which satisfy the matching condition in Equation (15).

Scheme	Amplitude modulation function $G^{\pm}(\boldsymbol{\omega})$	Potential function $V^{\pm}(\boldsymbol{\omega})$	Normalisation constant $(p_0^{\pm})^{-1}$
ODL	$(p_0^{\pm})^{-1/2}$	$-\log(A^{\pm}(\boldsymbol{\omega}))$	$\int_{\mathbb{R}_+\times\mathbb{R}} A^{\pm}(\boldsymbol{\omega})\, d\boldsymbol{\omega}$
H	$(p_0^{\pm})^{-1/2} A^{\pm}(\boldsymbol{\omega})^{1/4}$	$-\frac{1}{2}\log(A^{\pm}(\boldsymbol{\omega}))$	$\int_{\mathbb{R}_+\times\mathbb{R}} A^{\pm}(\boldsymbol{\omega})^{1/2}\, d\boldsymbol{\omega}$
ELGM	$(p_0^{\pm})^{-1/2} A^{\pm}(\boldsymbol{\omega})^{1/2}$	$\begin{cases} 0 & \boldsymbol{\omega} \in D^{\pm} \\ \infty & \boldsymbol{\omega} \notin D^{\pm} \end{cases}$	$\lvert D^{\pm}\rvert$

Note: ODL: overdamped Langevin; H: hybrid; ELGM: Eckermann–Lott–Guez–Maury. Note that the final scheme (ELGM) is defined only when the broadband spectra $A^{\pm}(\boldsymbol{\omega})$ are non-zero only within finite simply connected subsets $D^{\pm} \subset \mathbb{R}_+ \times \mathbb{R}$ (i.e., $A^{\pm}(\boldsymbol{\omega})$ has compact support). In this context $|D^{\pm}|$ denotes the area of these subsets.

are normalising constants. Inserting the expression for $p_s^{\pm}$ into the matching condition Equation (10), and then integrating over a, results in

$$\sum_{\pm} G^{\pm}(\boldsymbol{\omega})^2 p_0^{\pm} \exp(-V^{\pm}(\boldsymbol{\omega})) = A(\boldsymbol{\omega}). \tag{14}$$

In the HLP context it is natural to assume that the broadband forcing spectrum can be decomposed into leftward and rightward components according to $A(\boldsymbol{\omega}) = A^{+}(\boldsymbol{\omega}) + A^{-}(\boldsymbol{\omega})$, where $A^{+}(\boldsymbol{\omega})$ is the spectrum associated with rightward-propagating waves and $A^{-}(\boldsymbol{\omega})$ is associated with leftward-propagating ones. This further simplifies the matching condition to

$$G^{\pm}(\boldsymbol{\omega})^2 p_0^{\pm} \exp(-V^{\pm}(\boldsymbol{\omega})) = A^{\pm}(\boldsymbol{\omega}). \tag{15}$$

Since there is considerable freedom to choose the functions $G^{\pm}$ and $V^{\pm}$, it should be clear Equation (15) defines a broad family of schemes each of which has the property that, in the limit $\tau \to 0$, it should reproduce the deterministic QBO generated by the multi-wave model BBM-HLP. Table 1 details three possible schemes, which we have named ODL, hybrid (H), and Eckermann–Lott–Guez–Maury (ELGM). The ELGM scheme is so named because the choice of the "square-well" potential function means that $\boldsymbol{\Omega}_t^{\pm}$ evolves according to a regular Brownian process within its (respective) domain $D^{\pm}$, with a reflecting boundary condition applying at the boundary of $D^{\pm}$. Our ELGM scheme is therefore quite close in concept to the random update procedure for the wave phase speed used in the scheme of Eckermann (2011), with the key difference being that the wave frequency evolves continuously in time in our ELGM scheme, as opposed to being updated every time step in the scheme of Eckermann (2011). It should be noted that the abrupt time-step dependence associated with the random updates in Eckermann's scheme is associated with numerical problems, which led to subsequent works (Lott *et al.*, 2012; Lott & Guez, 2013) adapting the scheme by retaining a larger number of stochastic waves, with their impact being smoothed over multiple time steps using an AR(1) process. In the following it will be demonstrated that the SDE framework leads to more stable schemes without the need for additional waves.

The three schemes in Table 1 will be the focus of the next section.

4 | PROOF-OF-CONCEPT SIMULATIONS

In this section, numerical simulations of the deterministic multi-wave model BBM-HLP, Equation (5), and of the stochastic model of BB-sHLP, Equation (7), driven by the three stochastic schemes detailed in Table 1, are investigated with the primary aim of demonstrating that:

1. For each of the three BB-sHLP stochastic schemes, the BB-sHLP QBO converges accurately to the deterministic BBM-HLP QBO in the "fast" ($\tau \to 0$) limit in which the SDE effectively samples its state space on a time-scale much shorter than the QBO evolution time-scale.
2. As τ is increased, each of the three BB-sHLP schemes exhibits quite different intermittency properties, as can be seen clearly by observing changes in QBO amplitude and period.

The details of the integrations follow.

4.1 | Set-up of the numerical simulations

For the BBM-HLP simulations, the following choice for the spectral density function $A(k,\omega) = A^{+}(k,\omega) + A^{-}(k,\omega)$ is made:

$$A^{\pm}(k,\omega) = \delta(k-1)\frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(\omega \mp 1)^2}{2\sigma^2}\right]. \tag{16}$$

The choice to make the spectrum narrowband in wave number (i.e., fixing $k = 1$) and broadband in frequency

is made with computational efficiency in mind, because in the general case the numerical solution of BBM-HLP requires a discretisation of the spectrum in both wave number and frequency space, which makes the model expensive to integrate. The parameter σ controls the width of the broadband frequency spectrum; notice that the original HLP model, Equation (1), is recovered in the limit $\sigma \to 0$.

Correspondingly, BB-sHLP, Equation (5), is simplified by setting $K_t^{\pm} = 1$. The dynamics are then governed by $A_t^{\pm} = G^{\pm}(\Omega_t^{\pm})$ for the amplitude of each wave, and the corresponding frequency $\Omega_t^{\pm}$ is governed by the single variable SDE:

$$d\Omega_t^{\pm} = -\frac{V^{\pm\prime}(\Omega_t^{\pm})}{\tau}\,dt + \left(\frac{2}{\tau}\right)^{1/2} dB_t^{\pm}. \tag{17}$$

For the three schemes the potential and amplitude modulation functions are then given by

$$\begin{aligned}
\text{ODL:}\quad & V^{\pm}(\omega) = \frac{(\omega \mp 1)^2}{2\sigma^2}, \\
& G^{\pm}(\omega)^2 = 1, \\
\text{H:}\quad & V^{\pm}(\omega) = \frac{(\omega \mp 1)^2}{4\sigma^2}, \\
& G^{\pm}(\omega)^2 = \sqrt{2}\exp\left[-\frac{(\omega \mp 1)^2}{4\sigma^2}\right], \\
\text{ELGM:}\quad & V^{\pm}(\omega) = \begin{cases} 0 & |\omega| \in [\omega_-, \omega_+] \\ \infty & \text{otherwise} \end{cases}, \\
& G^{\pm}(\omega)^2 = \frac{(\omega_+ - \omega_-)}{\sqrt{2\pi\sigma^2}}\exp\left[-\frac{(\omega \mp 1)^2}{2\sigma^2}\right].
\end{aligned} \tag{18}$$

Figure 1 illustrates the behaviour of the SDEs governing the three schemes. The left panels show the invariant density (or stationary probability density function) $p_s^+(\omega)$ for the stochastic process Ω_t^+ in each case, together with histograms compiled from numerical realisations of Equation (17). The histograms serve to demonstrate that in each case the stationary probability density function of Ω_t^+ is effectively sampled by the process Ω_t^+ within 50τ time units (corresponding to 1.5×10^5 samples). The middle panels show a single realisation of Ω_t^+ over an interval of length 10τ in order to illustrate the typical behaviours of Equation (17) for the three schemes. Finally, the right panels plot the square of the amplitude modulation function $G^+(\omega)$, which determines the wave amplitude A_t^+ in terms of its frequency Ω_t^+. Notice that, for the ELGM scheme, the process Ω_t^+ is a random walk that is reflected at the boundary points $[\omega_-, \omega_+]$, which means that the tails of the spectrum $A^+(\omega)$ outside of this interval are not captured by the scheme. However, provided that the cut-offs are chosen so most of the spectral power is captured then the approximation is not significant. Here, $[\omega_-, \omega_+] = [0.1, 1.9]$, which is sufficient to capture the Eliassen–Palm (EP) flux at the lower boundary to eight significant figures.

4.2 | Numerical integrations of BBM-HLP

The broadband multi-wave model BBM-HLP, given by Equation (5) and with spectral density $A^{\pm}(k, \omega)$ given by Equation (16), with Reynolds number $\mathrm{Re} = 10$, is integrated numerically using a scheme adapted from that described in the supplementary material of Renaud *et al.* (2019). Time stepping is performed with a standard two-step predictor–corrector scheme. The treatment of the vertical integral is as described in Ewetola and Esler (2024), in which the mean velocity U is replaced by a piecewise linear approximation. The integrals in the right-hand side of Equation (5) can then be evaluated exactly for the piecewise linear flow, greatly reducing the sensitivity of the scheme to the relative location of critical levels and the model grid. The vertical domain is of height 5 (dimensional units Z_*) and is discretised with grid size $\Delta z = 10^{-3}$. A time step $\Delta t = 10^{-2}/\mathrm{Re}$ was found to be adequate across a wide range of model parameter space. The numerical validation of the code for the monochromatic HLP case is described in Ewetola and Esler (2024). For the extension to BBM-HLP, the frequency spectrum integral is discretised into a sum over $N_\omega = 50$ discrete frequencies in the range $0.01 \le \omega \le 1.99$ with $\Delta\omega = 0.0404$, which was found to be sufficient to capture the EP flux associated with the continuous spectrum to high accuracy. The discretisation of the spectrum means that BBM-HLP is close to N_ω times more computationally expensive than the original HLP model.

Figure 2 shows Hovmöller plots of $U(z, t)$ illustrating the typical QBO-like oscillations that occur in BBM-HLP. Each panel shows integrations in which the width σ of the broadband spectrum is varied, from $\sigma = 0$, which corresponds to the original HLP system, Equation (1), and then BBM-HLP calculations with $\sigma = 0.15$ and $\sigma = 0.3$. The effect of spectral broadening in our simulations differs somewhat from that reported recently by Léard *et al.* (2020). Here, increasing σ is seen to increase the QBO period and (slightly) increase the QBO amplitude, whereas Léard *et al.* (2020) found much larger increases in both QBO period and amplitude. There are a number of differences in the experimental set-up that can account for the differences between the two studies. For example, a different dissipation mechanism for the waves is used in Léard *et al.* (2020), which means their governing equations differ from HLP. Arguably even more significantly, in Léard *et al.* (2020) the total wave kinetic energy at the

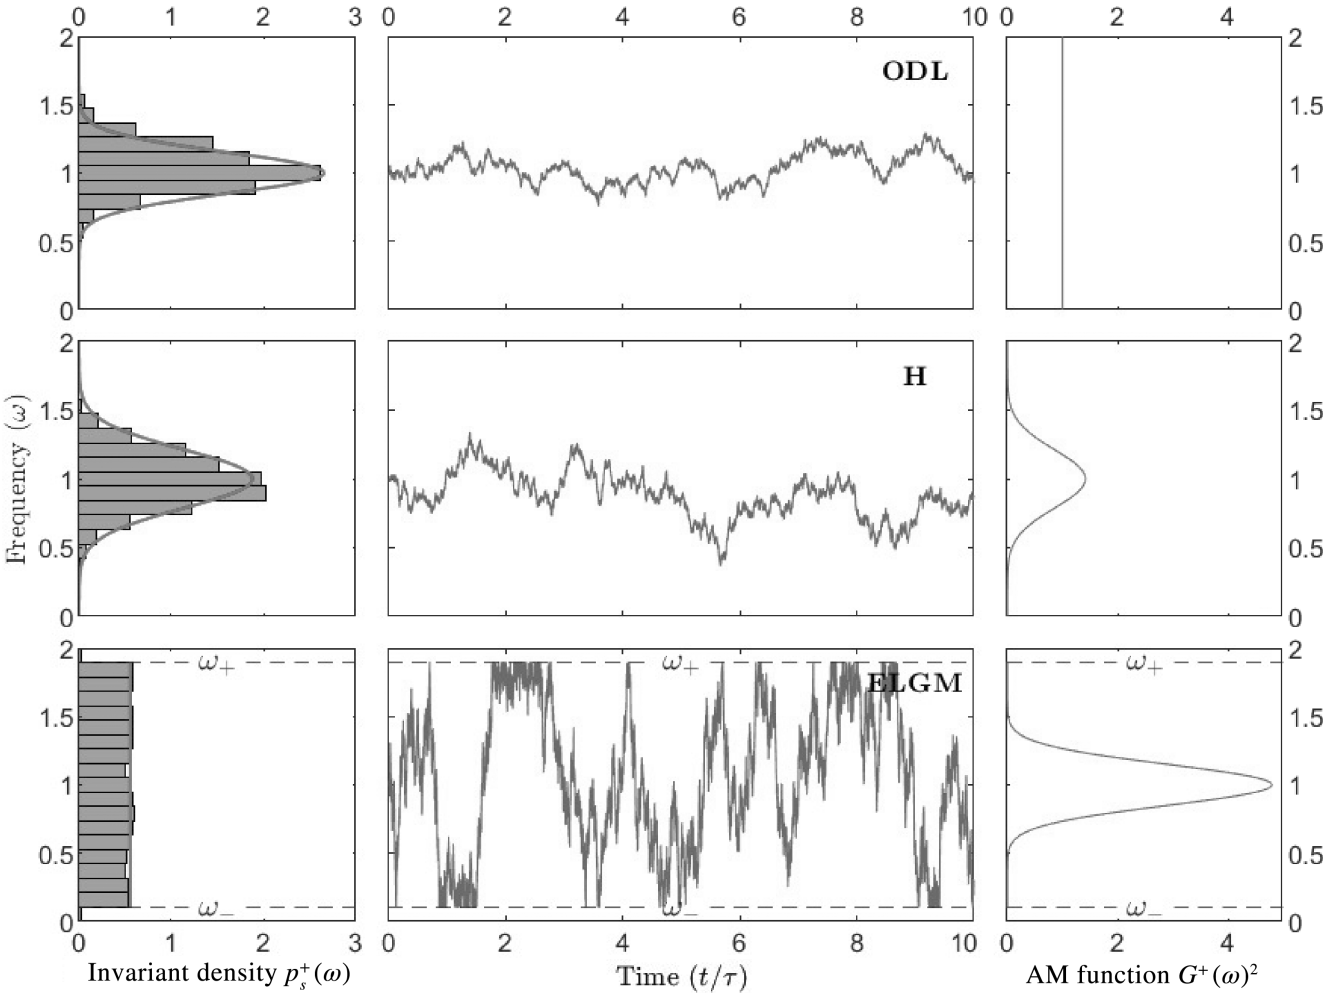

FIGURE 1 Illustrating the behaviour of the stochastic differential equations governing the overdamped Langevin (ODL), hybrid (H), and Eckermann–Lott–Guez–Maury (ELGM) schemes when the spectral density $A^{\pm}(k,\omega)$ is governed by Equation (16) with $\sigma = 0.15$. Left panels: Histograms of Ω_t^+ (compiled from 1.5×10^5 data points sampled at every time step over a time interval 50τ) plotted against the corresponding invariant density $p_s^+(\omega)$ (curves). Middle panels: A single realisation of Ω_t^+ over an interval of length 10 τ. Right panels: The amplitude modulation (AM) function $G^+(\omega)^2$.

lower boundary is held constant as σ is varied, whereas in Figure 2 it is the total spectral density (i.e., the integral of $A(k,\omega)$) held constant. Additionally, a special property of the spectrum in Equation (16) is that the total EP flux associated with the left and right parts of the spectrum, for example

$$\int_{\mathbb{R}_+\times\mathbb{R}} \omega A^+(k,\omega)\, dk\, d\omega, \tag{19}$$

is also conserved as σ is varied. Recently, steps towards understanding what controls the changing stability of the QBO when the spectrum is broadened have been made by Chartrand *et al.* (2024), in a detailed study of a multi-wave HLP system.

Next, the practical question of how these (expensive) BBM-HLP calculations can be emulated by (cheap) BB-sHLP stochastic calculations will be addressed.

4.3 | Numerical integrations of BB-sHLP

The BB-sHLP model, Equation (7), can be integrated numerically by making minimal alterations to a standard HLP numerical code. Essentially, the HLP model needs only to be coupled to a pair of SDEs that solve for $\boldsymbol{Q}_t^+$ and $\boldsymbol{Q}_t^-$. In the present set-up, it is sufficient to couple to Equation (17), which solves for Ω_t^+ and Ω_t^-, since $K_t^{\pm} = 1$ and $A_t^{\pm} = G^{\pm}(\Omega_t^{\pm})$ are prescribed. In our experience, it is sufficient to solve Equation (17) using a standard Euler–Maruyama time-stepping scheme. Note that the standard Euler–Maruyama discretisation of the Ornstein–Uhlenbeck (OU) process yields an AR(1) process. Solving the SDE is of very low computational expense compared with evaluating the integral terms in HLP; therefore, in ideal circumstances, BB-sHLP is hardly more expensive than the original HLP model, Equation (1).

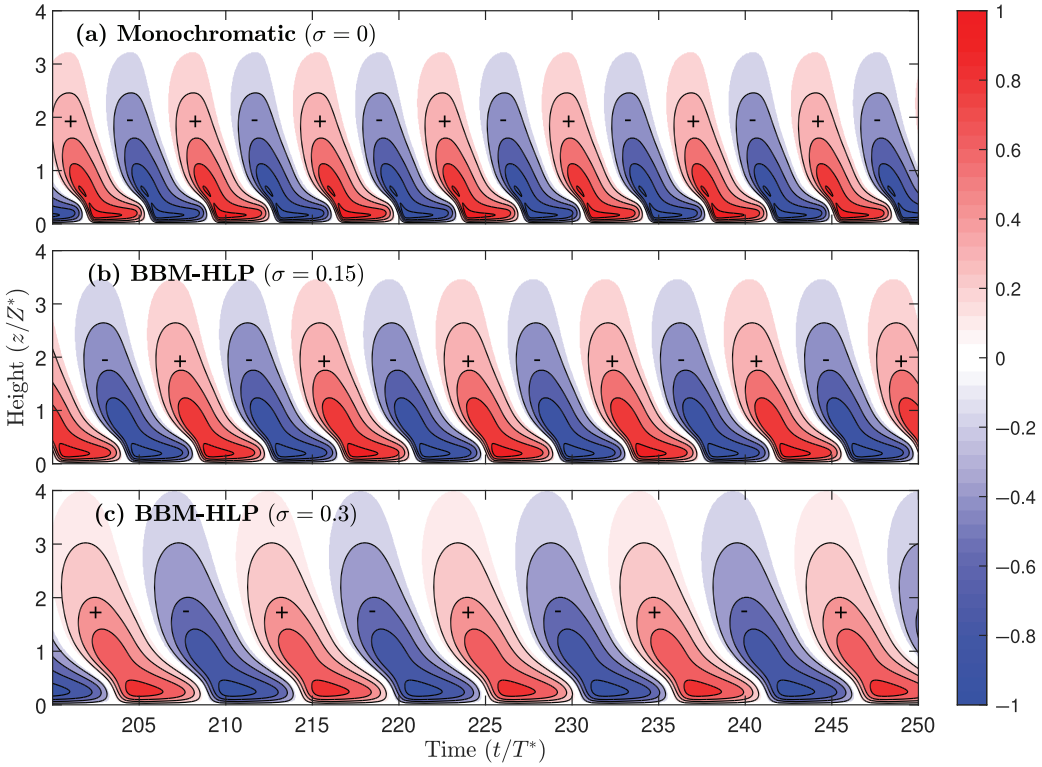

FIGURE 2 Hovmöller plots of $U(z, t)$ in the broadband multi-wave version of the Holton–Lindzen–Plumb (BBM-HLP) model calculations showing the equilibrated quasi-biennial oscillations after a spin-up period. In each case Re = 10 and the width σ of the broadband spectrum is (a) $\sigma = 0$ (the monochromatic HLP case), (b) $\sigma = 0.15$, and (c) $\sigma = 0.3$. [Colour figure can be viewed at wileyonlinelibrary.com]

However, in practice, particularly for lower values of the SDE time-scale τ, it is necessary to reduce the BB-sHLP model time step in order to resolve accurately the time dependence of the integrand appearing in Equation (7). The time step used for the ODL and H schemes is $\Delta t = Min\{\tau/20, 10^{-2}/\mathrm{Re}\}$ with smaller steps required for ELGM (see subsequent discussion).

For Figure 4, the BB-sHLP, Equation (7), is integrated with each of the three stochastic schemes (ODL, H, ELGM), for a range of values of $\tau = \{0.001, 0.002, 0.005, 0.01, 0.02, 0.05, 0.1, 0.2, 0.3\}$. In each case, Equation (7) with Re = 10 is coupled with Equation (17) with the potential function and amplitude modulation functions obtained from Equation (18) with the spectral width $\sigma = 0.15$. In Figure 3, Hovmöller plots of $U(z, t)$ are shown for the shorter SDE time-scale value ($\tau = 0.02$, left panels) and the longest value ($\tau = 0.3$, right panels). Note that, although we are comparing the schemes at fixed τ in Figure 3, the exact value of τ is not directly comparable between the different schemes. Our main results, however, are not qualitatively sensitive to redefining τ for each scheme.

The left panels in Figure 3 should be compared with the corresponding BBM-HLP calculation shown in Figure 2b. Accurate convergence of the BB-sHLP integration to the corresponding BBM-HLP QBO is evident for all three schemes. This convergence occurs because stochastic averaging applies, since the SDE in each case evolves rapidly compared with the QBO time-scale (as $\tau \ll 1$), and each SDE scheme satisfies the matching condition in Equation (15).

The typical effect on the QBO of introducing intermittency (i.e., finite τ) to each of the three schemes is shown in the right panels of Figure 3, for which $\tau = 0.3$ in each case. Based on Figure 3 alone, the most robust difference between the schemes is that they become increasingly noisy from ODL → H → ELGM. Autocorrelation plots of $U(z, t)$ (not shown) reveal that differences in the QBO vertical structure between the three schemes are in fact small. The variations in the QBO structure apparent in Figure 3 are therefore primarily due to the internal variability of each scheme as opposed to differences between them.

A more systematic picture of how the resulting QBO varies as τ is increased in each of the three schemes is presented in Figure 4, where the QBO period $\mathcal{P}$ is plotted against QBO amplitude $\mathcal{A}$. Here, the QBO amplitude is calculated using the following root-mean-square norm:

$$\mathcal{A} = \left(\frac{1}{\Delta T} \int_{T_-}^{T_+} \int_0^{z_t} U(z, t)^2 \, dz \, dt \right)^{1/2}, \qquad (20)$$

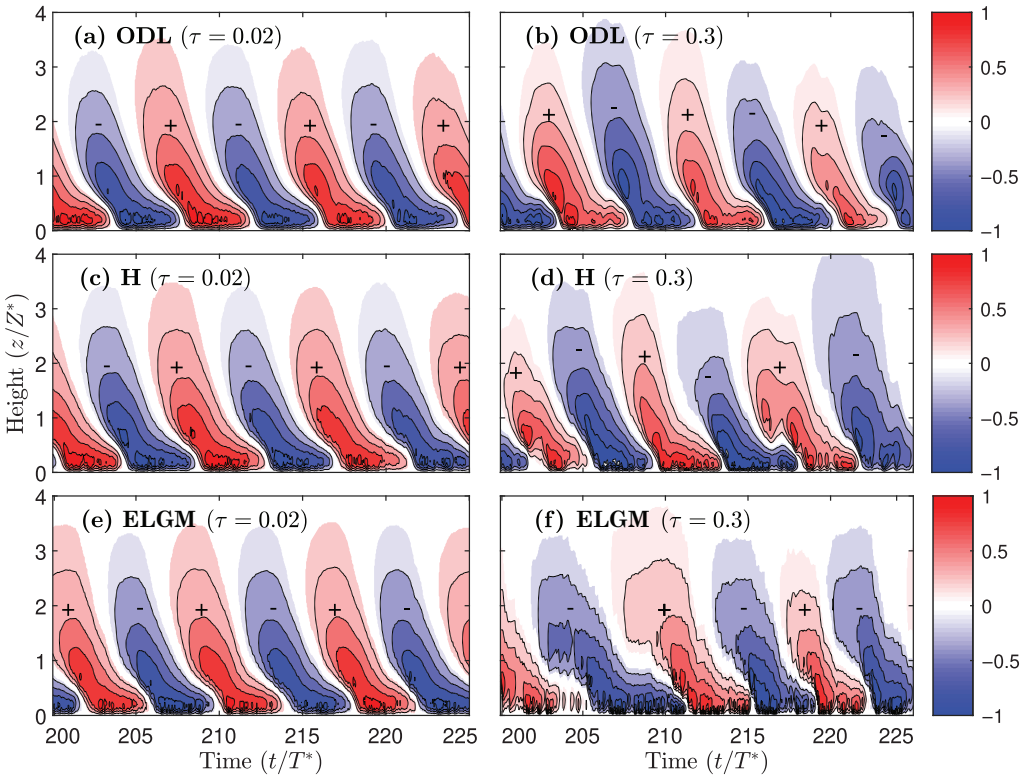

FIGURE 3 Hovmöller plots of $U(z, t)$ from broadband stochastic Holton–Lindzen–Plumb (HLP) model integrations using the overdamped Langevin (ODL), hybrid (H), and Eckermann–Lott—Guez–Maury (ELGM) schemes. Left panels: Calculations with $\tau = 0.02$ and $\sigma = 0.15$ illustrating convergence to the deterministic broadband multi-wave HLP quasi-biennial oscillation (QBO; see Figure 2b). Right panels: Calculations with $\tau = 0.3$ illustrating the QBOs generated by each scheme when significant intermittency is present. [Colour figure can be viewed at wileyonlinelibrary.com]

where $z_t = 5$ is the top of the domain, $(T_-, T_+) = (200, 10, 200)$, and $\Delta T = T_+ - T_- = 10^4$. The QBO period is calculated from the location of the principal maximum of the autocorrelation function of $U(z_*, t)$, where $z_* = 0.2$ is a representative level near the bottom of the domain. Note that both of these definitions are modified from those used in Ewetola and Esler (2024), where the amplitude was taken to be the root-mean-square value of $U(z_m, t)$, where z_m is the level of its maximum value, and the period was obtained from the Fourier transform of $U(z_m, t)$. The new definitions have been adopted because they return more robust results for the most noisy time series; however, the change has minimal impact on any of the results presented here or in the previous work.

The solid black point in Figure 4 shows $(\mathcal{A}, \mathcal{P})$ for the $\sigma = 0.15$ BBM-HLP calculation shown in Figure 2b. For context, to show the impact of varying the spectral bandwidth σ in BBM-HLP, the unfilled black circles show the corresponding results for $\sigma = 0.1$, 0.125, and 0.175. The BB-sHLP results for the simulations described earlier, over the full range of τ, are shown by the coloured points. The simulations illustrated in Figure 3 (i.e., those with $\tau =$ 0.02 and 0.3) are given an additional outline for emphasis.

As τ is decreased for each scheme, the QBO amplitude $\mathcal{A}$ and period $\mathcal{P}$ are seen to converge to the BBM-HLP result. At the lowest value of $\tau = 0.001$, the differences in $(\mathcal{A}, \mathcal{P})$ between BBM-HLP and BB-sHLP are $\lesssim 0.2\%$ for all three schemes for $\sigma = 0.15$. It should be noted that the $\tau = 0.001$ integrations have significant associated computational cost compared with HLP, because a shorter time step is required. A time step $\delta t = \min(\tau/20, 10^{-2}/\mathrm{Re})$ is used, with the second value being the standard time step for the deterministic HLP model, giving $\delta t = 10^{-3}$ for our $\mathrm{Re} = 10$ simulations. Consequently, the $\tau = 0.001$ runs require a time step 20 times smaller than for the deterministic case. However, the differences in $(\mathcal{A}, \mathcal{P})$ when $\tau = 0.02$, which is the lowest value of τ for which a shorter time step is not required (for the ODL and H schemes) to properly resolve the SDE, remains less than 1%. Therefore, it has been demonstrated that the BBM-HLP QBO can be effectively reproduced by a BB-sHLP stochastic scheme that is a factor of N_ω (here, 50) cheaper than the broadband multi-wave calculation.

Table 2 shows how the convergence of the schemes perform as the width of the broadband forcing spectrum σ is increased. For broader spectra, the accuracy of the

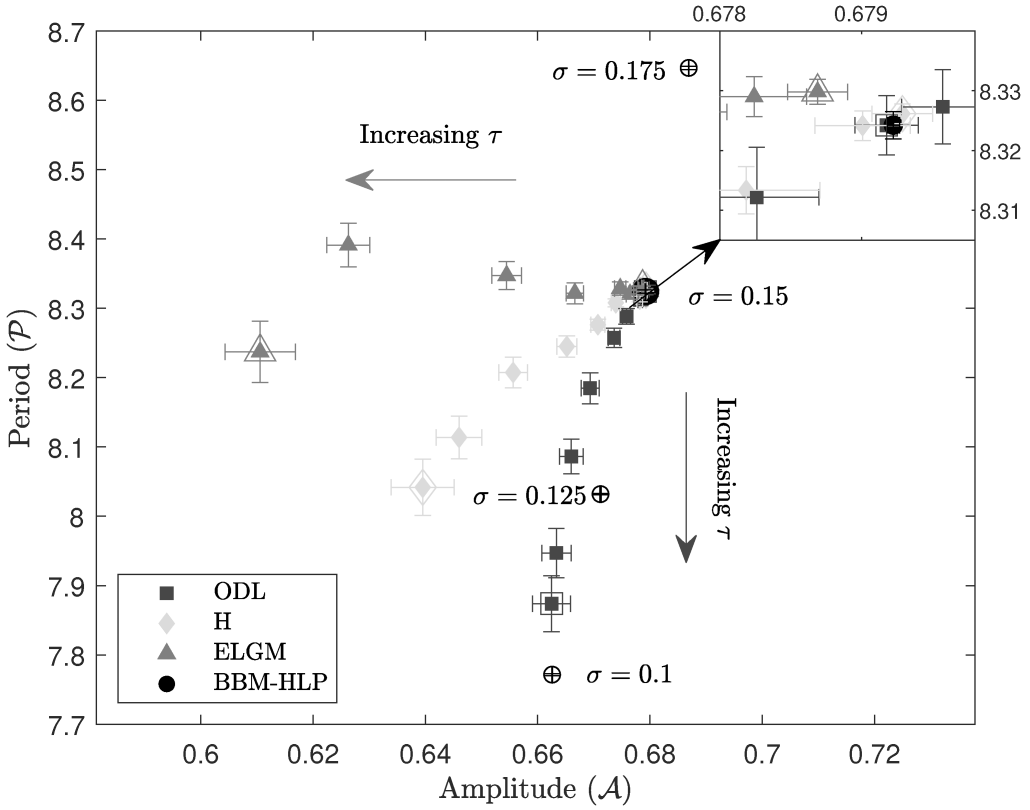

FIGURE 4 Calculated quasi-biennial oscillation (QBO) period $\mathcal{P}$ against QBO amplitude $\mathcal{A}$ for the set of broadband stochastic Holton–Lindzen–Plumb (HLP) model simulations described in the text with $\sigma = 0.15$. Error bars calculated according to the method described by Flyvbjerg and Petersen (1989, section III). The overdamped Langevin (ODL), hybrid (H), and Eckermann–Lott–Guez–Maury (ELGM) schemes are illustrated by triangles, diamonds and squares respectively. The result from the corresponding deterministic broadband multi-wave HLP (BBM-HLP) calculation is shown by the circle. The results for BBM-HLP simulations with different values of σ are shown by unfilled circles. The simulations shown in Figure 3 are labelled with an additional bounding box.

TABLE 2 Table of the percentage error when calculating period and amplitude of the broadband multi-wave Holton–Lindzen–Plumb (HLP) model for several values of σ when using broadband stochastic HLP models with $\tau = 0.001$.

	Error (%)					
	Period			Amplitude		
σ	ODL	H	ELGM	ODL	H	ELGM
0.15	0.0038	0.0194	−0.0633	0.0244	0.0081	0.0962
0.30	0.0495	0.0512	0.0297	0.2902	0.1369	0.0669
0.45	0.9012	1.8435	0.0058	3.9653	2.0359	0.8102

Abbreviations: ODL: overdamped Langevin; H: hybrid; ELGM: Eckermann–Lott–Guez–Maury.

period and amplitude calculated from the ODL and H BB-sHLP model at a fixed value of τ is reduced. This result is expected, as the stochastic wave has a broader spectrum to traverse in order to emulate the results of the deterministic multi-wave model, and it is likely that accuracy could be recovered by increasing the number of stochastic waves. However, even with a single wave the accuracy remains reasonable for a significant range of values of σ.

The ELGM results come with a caveat, because technical issues become important as the spectrum gets broader. In the ELGM scheme, the stochastic processes Ω_t^+ must be restricted to a finite interval $[\omega_-, \omega_+]$. The tails of the spectrum $A^+(\omega)$ outside this interval are cut off and not captured by the scheme. As long as the interval is chosen to be sufficiently broad, the approximation is not significant. For a Gaussian spectrum, the width of interval required to

capture a given percentage of the total flux is linear in the standard deviation σ. For $\sigma = 0.15$, we chose an interval of width 12σ, which captures the total EP flux of the Gaussian spectrum up to eight significant figures. However, for $\sigma > 1/6$, an interval of width 12σ would give $\omega_- < 0$, which means that the wave associated with the rightward spectrum can have a negative frequency ω. As our model fixes $k = 1$, this means that the "rightward wave" can have a leftward phase speed; this causes problems, as the QBO-like oscillations of the HLP model are dependent on two opposing waves. In order to circumvent these issues, for $\sigma > 1/6$ we have set $[\omega_-, \omega_+] = [0, 2]$. This has the side effect of cutting off more of the forcing spectrum compared with an interval of width 12σ, but the majority is still captured. A choice of $[\omega_-, \omega_+] = [0, 2]$ captures 99.9% of the spectrum for $\sigma = 0.3$, and 97.4% of the spectrum for $\sigma = 0.45$. However, as this is not the same percentage for each value of σ, the results for different values of σ for the ELGM model are not directly comparable.

Figure 4 shows that the effect of introducing intermittency (increasing τ) has a different impact on the QBO for each scheme. The effect of increasing τ in the ODL scheme is primarily to shorten the QBO period, whereas the QBO amplitude remains fairly constant. By contrast, in the ELGM scheme the QBO amplitude is reduced as τ increases but the period remains almost unchanged. The behaviour of the H scheme is intermediate between the other two schemes. These results make clear that it is not sufficient simply to introduce randomness into a gravity wave parametrisation in order to improve the modelling of an intermittent source. The nature of the intermittency to be modelled must first be understood, and then the stochastic scheme must be chosen carefully to account for it. Although a full solution to this problem must be addressed in future work, Figure 4 makes clear that even the restricted SDE framework presented here offers the flexibility to capture a wide range of intermittency impacts on idealised QBO amplitude and period.

In order to understand some of the differences between the three schemes, in particular why ELGM appears to be noisier and more intermittent than ODL and H are at the same value of τ, the time evolution of the lower boundary EP flux $F^+ = \Omega_t^+ |A_t^+|^2$ due to the rightward-propagating wave is shown in Figure 5 for $\tau = 0.3$. It is at once evident that F^+ is relatively constant in time in ODL, is more variable in H, and is extremely variable with large intermittent pulses in ELGM. The requirement of resolving the variable impact on the QBO associated with these intermittent EP flux pulses is the reason the ELGM scheme requires a shorter time step than ODL or H. We have found that for moderate τ the ELGM time step needs to be smaller by at least a factor of 5 compared with ODL, because the intrinsic QBO time-scale T_* is shorter by a factor of 5 during the high EP flux episodes (see scaling of T_* earlier herein). To understand the differences in the EP flux behaviour, consider that the ODL scheme visits frequencies at the tails of the spectrum infrequently but maintains the wave amplitude when it does so, resulting

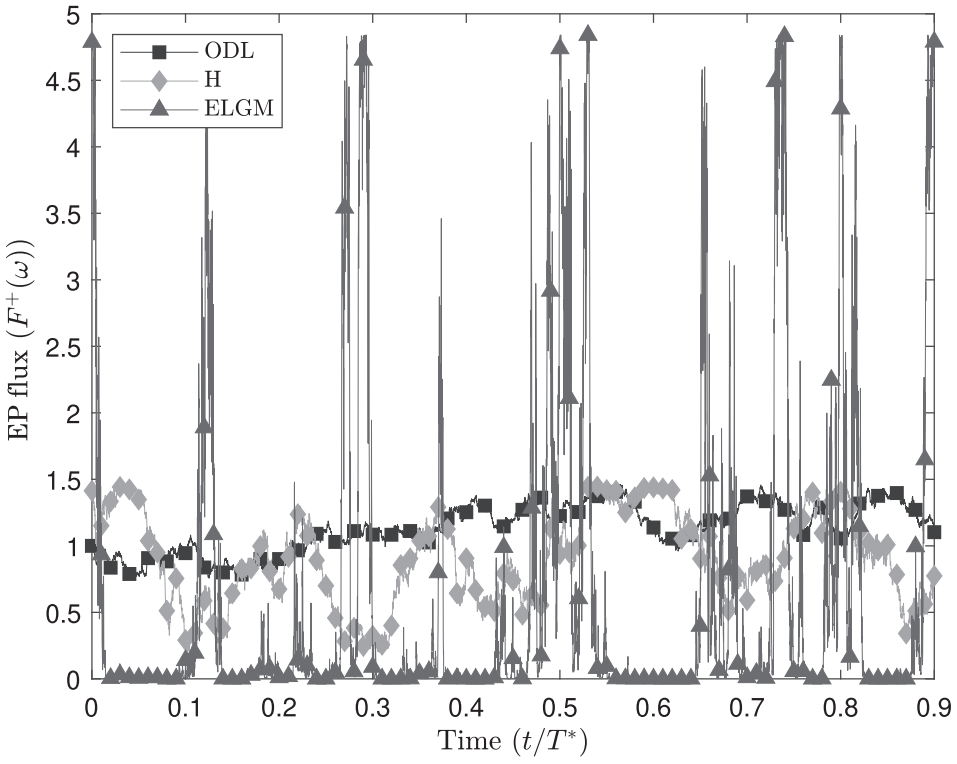

FIGURE 5 Time evolution of the Eliassen–Palm flux at the lower boundary $F^+ = \Omega^+ |A_t^+|^2$ of the rightward-propagating wave in the overdamped Langevin (ODL), hybrid (H), and Eckermann–Lott–Guez–Maur (ELGM) schemes; $\tau = 0.3$.

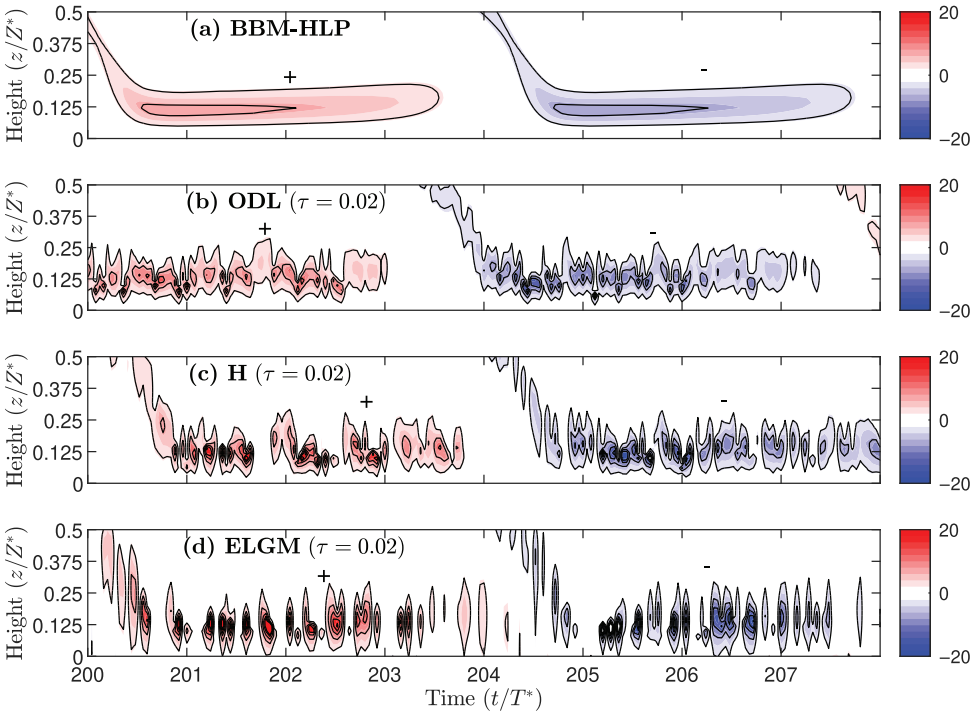

FIGURE 6 Contour plot of the Eliassen–Palm flux in the broadband multi-wave Holton–Lindzen–Plumb (BBM-HLP), overdamped Langevin (ODL), hybrid (H), and Eckermann–Lott–Guez–Maur (ELGM) schemes; $\tau = 0.02$, $\sigma = 0.15$. [Colour figure can be viewed at wileyonlinelibrary.com]

in a relatively constant EP flux. By contrast, the ELGM scheme visits the tails of the spectrum often but attenuates the wave amplitude when it does so, resulting in the long intervals with low EP flux seen in Figure 5.

Figure 6 shows a different perspective on the intermittency of the three schemes by contouring the EP flux as a function of (z, t). This can also be seen in a contour plot of forcing—as given by the right-hand side of Equation (7)—as a function of height and time. The results of this contour plot for the deterministic HLP model are well known. The monochromatic case consists of a sharp spike just below the critical level, whereas the BBM-HLP results in a broader peak. Figure 6 shows that the forcing due to the BB-sHLP schemes is less smooth than that due to the BBM-HLP scheme. In particular, the forcing due to the ELGM scheme is highly localised in time, occurring in several short separated bursts. This corresponds with the behaviour shown in Figure 5, with the ODL, H, and ELGM schemes being ordered by increasing intermittency of their forcing.

5 | CONCLUSIONS

The SDE framework developed here has a number of advantages that have been clarified herein.

1. The SDE framework generates parametrisations that are defined independently of the details of the model discretisation, such as the model time step, meaning that researchers are free to choose a numerical scheme best suited to their needs. Model developers can then draw on the vast literature on the numerical solution of SDEs (e.g., Kloeden & Platen, 1992; Milstein & Tretyakov, 2004).
2. In common with previous stochastic schemes (e.g., Eckermann, 2011; Lott *et al.*, 2012), the SDE framework allows for considerable computational savings. For example, in the HLP setting used here, it has been shown in principle how to emulate the behaviour of a deterministic multi-wave forcing spectrum (BBM-HLP) that is broadband in both frequency and wave number, using a scheme with just two stochastically evolving waves (BB-sHLP). In the broadband frequency setting studied in Section 4, these computational savings are realised in full, as the ODL scheme in BB-sHLP is only marginally more expensive to integrate (provided $\tau \gtrsim 0.02$) than the deterministic two-wave HLP model. By contrast, BBM-HLP is N_ω times more expensive, where $N_\omega = 50$ is the number of waves in the discretised spectrum. For a spectrum that is broadband in both frequency and wave number, much greater savings are realistic. It is worth emphasising that in the general circulation model setting these computational savings could be reinvested in enhancing the realism of the schemes in other respects; for example, to move beyond columnar gravity wave propagation, which is known to be a significant shortcoming of current schemes (Plougonven *et al.*, 2020).
3. The flexibility introduced by the SDE framework allows for the possibility that schemes can be developed in order to match the intermittency of observed gravity

wave sources. Each SDE scheme having the correct spectral density $A(k,\omega)$—that is, satisfying the matching condition in Equation (10)—has its own intermittency properties that become evident as the time-scale τ of the SDE process is increased. The intermittency effects of different schemes can have profoundly different impacts on the QBO period and amplitude, as is seen in Figure 4, showing that the naive approach of just taking any stochastic scheme and "making it intermittent" will not suffice as a robust approach to model observed intermittent effects. A step forward in understanding the effect of intermittency has been made recently by Ewetola and Esler (2024), who showed, in a version of the two-wave HLP in which the wave amplitudes evolve according to an SDE, that the effect of the intermittency can be reduced to a single intermittency parameter depending only on the details of the SDE in question.

4. The flexibility of the SDE framework can also be exploited to enhance the numerical stability of a stochastic gravity wave parametrisation. For example, it was notable in the study of Section 4 that the ODL scheme was considerably more stable and computationally robust (and therefore cheaper) compared with the ELGM scheme, with the probable cause being the much lower variance in EP flux in the ODL scheme. Enhancing the numerical stability properties of such schemes can be a key practical concern of modellers.

Finally, it should be emphasised that although the SDE schemes developed here have been tested in the highly idealised HLP framework, there is no obstacle in principle to adapting the approach to more realistic model settings. Nevertheless, there remain theoretical, practical, and numerical challenges that need to be addressed before new schemes based on the framework described here can be used operationally to their best advantage. Principal among these are the following:

1. A natural extension of the present work is to deal with more general parametrisations of wave forcing. For example, in Ewetola and Esler (2024), the theory is developed for a general forcing functional $F[U(z),z]$ in place of the (monochromatic) HLP wave forcing. Extending the approach of Ewetola and Esler (2024) to the broadband setting of the current work is a priority. Functionals F can be chosen to align with the results of simple wave-breaking parametrisations as a starting point for addressing more realistic parametrisations.
2. The next challenge for future work is to extend the results of Ewetola and Esler (2024), in which a theoretical description of the effects of intermittency in the case of monochromatic forcing in HLP is given, to the more general broadband situation. A theoretical understanding of intermittency will underpin attempts to optimise the choice of the potential function ($V^{\pm}(\omega)$ herein), and time-scale τ, which define the SDE-based parametrisation scheme. The aim is that the optimal choice will best match the QBO generated by the actual intermittent source. Ideally, the theory can be developed so that the optimal choice of $V^{\pm}(\omega)$ and τ follow from the theory as opposed to being determined by parameter tuning.
3. A further challenge is to relax the vertical localisation of the launch spectrum to allow for vertically distributed wave sources; for example, from convective sources; see Lecoanet *et al.* (2015).
4. Finally, careful numerical analysis and testing of any new scheme are required, in order to understand numerical aspects, such as each scheme's time-step restrictions. For example, in the stochastic scheme of Lott *et al.* (2012) it was found necessary to apply a smoother to the time tendency of the wave forcing in order to minimise numerical issues, which were found to be present in previous stochastic schemes (Eckermann, 2011). A more general understanding of this type of numerical issue may be beneficial.

Overall, the features described herein provide a strong recommendation for exploring the SDE framework as a promising direction for the future development of gravity wave parametrisation schemes.

ACKNOWLEDGMENTS

JGE acknowledges funding by the UK Natural Environment Research Council (grant NE/S00985X/1).

DATA AVAILABILITY STATEMENT

Data sharing is not applicable to this article as no new data were created or analysed in this study.

ORCID

K. Xie https://orcid.org/0009-0008-2784-1236
J. G. Esler https://orcid.org/0000-0002-8577-8422

REFERENCES

Alexander, M.J. & Dunkerton, T.J. (1999) A spectral parameterization of mean-flow forcing due to breaking gravity waves. *Journal of the Atmospheric Sciences*, 56(24), 4167–4182. Available from: https://doi.org/10.1175/1520-0469(1999)056<4167:ASPOMF>2.0.CO;2

Baldwin, M.P., Gray, L.J., Dunkerton, T.J., Hamilton, K., Haynes, P.H., Randel, W.J. et al. (2001) The quasi-biennial oscillation. *Reviews of Geophysics*, 39, 179–229. Available from: https://doi.org/10.1029/1999RG000073

Chartrand, X., Nadeau, L.P. & Venaille, A. (2024) Recovering quasi-biennial oscillations from chaos. *Journal of the Atmospheric Sciences*, 81, 1213–1224. Available from: https://doi.org/10.1175/JAS-D-23-0220.1

Couston, L.A., Lecoanet, D., Favier, B. & Le Bars, M. (2018) Order out of chaos: slowly reversing mean flows emerge from turbulently generated internal waves. *Physical Review Letters*, 120(244), 505. Available from: https://doi.org/10.1103/PhysRevLett.120.244505

Eckermann, S.D. (2011) Explicitly stochastic parameterization of nonorographic gravity wave. *Journal of the Atmospheric Sciences*, 68, 1749–1765.

Ewetola, M. & Esler, J. (2024) The effect of intermittency in wave forcing on the quasi-biennial oscillation. *Journal of Fluid Mechanics*, 988, A16.

Flyvbjerg, H. & Petersen, H. (1989) Error estimates on averages of correlated data. *The Journal of Chemical Physics*, 91, 461–466. Available from: https://doi.org/10.1063/1.457480

Gardiner, C.W. (2009) *Stochastic methods: a handbook for the natural and social sciences*, 4th edition. Berlin: Springer, p. 447.

Garfinkel, C.I., Gerber, E.P., Shamir, O., Rao, J., Jucker, M., White, I. et al. (2022) A QBO cookbook: sensitivity of the quasi-biennial oscillation to resolution, resolved waves, and parameterized gravity waves. *Journal of Advances in Modeling Earth Systems*, 14(3), e2021MS002568. Available from: https://doi.org/10.1029/2021MS002568

Holmes, M.H. (2013) *Multiple scales*. New York NY: Springer, pp. 139–221. Available from: https://doi.org/10.1007/978-1-4614-5477-9_3

Holton, J.R. & Lindzen, R.S. (1972) An updated theory for the quasi-biennial cycle of the tropical stratosphere. *Journal of the Atmospheric Sciences*, 29, 1076–1080. Available from: https://doi.org/10.1175/1520-0469

Kloeden, P.E. & Platen, E. (1992) *Numerical solution of stochastic differential equations*. Berlin Heidelberg: Springer. Available from: https://doi.org/10.1007/978-3-662-12616-5.12

Léard, P., Lecoanet, D. & Le Bars, M. (2020) Multimodal excitation to model the quasibiennial oscillation. *Physical Review Letters*, 125(234), 501. Available from: https://doi.org/10.1103/PhysRevLett.125.234501

Lecoanet, D., Le Bars, M., Burns, K.J., Vasil, G.M., Brown, B.P., Quataert, E. et al. (2015) Numerical simulations of internal wave generation by convection in water. *Physical Review E*, 91(63), 016. Available from: https://doi.org/10.1103/PhysRevE.91.063016

Lindzen, R.S. (1981) Turbulence and stress owing to gravity wave and tidal breakdown. *Journal of Geophysical Research: Oceans*, 86(C10), 9707–9714. Available from: https://doi.org/10.1029/JC086iC10p09707

Lindzen, R.S. & Holton, J.R. (1968) A theory of the quasi-biennial oscillation. *Journal of the Atmospheric Sciences*, 25, 1095–1107. Available from: https://doi.org/10.1175/1520-0469

Lott, F. & Guez, L. (2013) A stochastic parameterization of the gravity waves due to convection and its impact on the equatorial stratosphere. *Journal of Geophysical Research: Atmospheres*, 118, 8897–8909. Available from: https://doi.org/10.1002/jgrd.50705

Lott, F., Guez, L. & Maury, P. (2012) A stochastic parameterization of non-orographic gravity waves: formalism and impact on the equatorial stratosphere. *Geophysical Review Letters*, 39, L06807. Available from: https://doi.org/10.1029/2012GL051001

Milstein, G.N. & Tretyakov, M.V. (2004) *Stochastic numerics for mathematical physics*. New York: Springer.

Pavliotis, G.A. & Stuart, A.M. (2008) *Homogenization for odes and sdes*. New York: Springer, pp. 157–182. Available from: https://doi.org/10.1007/978-0-387-73829-1.11

Piani, C., Norton, W.A. & Stainforth, D.A. (2004) Equatorial stratospheric response to variations in deterministic and stochastic gravity wave parameterizations. *Journal of Geophysical Research: Atmospheres*, 109(D14), D14101. Available from: https://doi.org/10.1029/2004JD004656

Plougonven, R., de la Cámara, A., Hertzog, A. & Lott, F. (2020) How does knowledge of atmospheric gravity waves guide their parameterizations? *Quarterly Journal of the Royal Meteorological Society*, 146(728), 1529–1543. Available from: https://doi.org/10.1002/qj.3732

Plumb, R.A. (1977) The interaction of two internal waves with the mean flow: implications for the theory of the quasi-biennial oscillation. *Journal of the Atmospheric Sciences*, 34, 1847–1858. Available from: https://doi.org/10.1175/1520-0469

Renaud, A. & Venaille, A. (2020) On the Holton–Lindzen–Plumb model for mean flow reversals in stratified fluids. *Quarterly Journal of the Royal Meteorological Society*, 146(732), 2981–2997. Available from: https://doi.org/10.1002/qj.3821

Renaud, A., Nadeau, L. & Venaille, A. (2019) Periodicity disruption of a model quasibiennial oscillation of equatorial winds. *Physical Review Letters*, 122(214), 504. Available from: https://doi.org/10.1103/PhysRevLett.122.214504

Richter, J.H., Sassi, F. & Garcia, R.R. (2010) Toward a physically based gravity wave source parameterization in a general circulation model. *Journal of the Atmospheric Sciences*, 67(1), 136–156. Available from: https://doi.org/10.1175/2009JAS3112.1

Song, I.S. & Chun, H.Y. (2008) A Lagrangian spectral parameterization of gravity wave drag induced by cumulus convection. *Journal of the Atmospheric Sciences*, 65(4), 1204–1224. Available from: https://doi.org/10.1175/2007JAS2369.1

Vallis, G.K. (2017) *Atmospheric and oceanic fluid dynamics: fundamentals and large-scale circulation*, 2nd edition. Cambridge: Cambridge University Press. Available from: https://doi.org/10.1017/9781107588417

How to cite this article: Xie, K., Ewetola, M. & Esler, J.G. (2025) A stochastic differential equation framework for gravity wave parametrisation with testing in an idealised setting. *Quarterly Journal of the Royal Meteorological Society*, e4983. Available from: https://doi.org/10.1002/qj.4983

APPENDIX A. COMPUTATION TIME OF BB-SHLP

Here, further details of the computational costs of the various schemes used to solve BBM-HLP and BB-sHLP are reported. Unsurprisingly, the majority of the computational resources used in solving the HLP equation are

required to evaluate the forcing terms on the right-hand side. Since BBM-HLP requires the simulation of N_ω waves for an accurate representation of the forcing spectrum, it is therefore approximately N_ω times more expensive to simulate. BB-sHLP, on the other hand, only requires the simulation of one pair of waves, the same as the classical, monochromatic HLP model. The only additional cost is the evaluation of the SDE, which can be negligible in comparison with the cost of simulating the HLP equation, depending on whether or not the SDE requires a shorter time step to resolve its variability.

Taking the original monochromatic HLP model as a baseline, Table A.1 confirms that the three stochastic schemes have comparable run time for $\tau = 0.02$, whereas the BBM-HLP model takes $N_\omega = 50$ times as long. In Table A.2, in which the run times for different values of τ are compared, it can be seen that the amount of time required to evaluate a single time step remains constant as τ is varied. Therefore, changing τ only changes the computational cost if it necessitates a reduced time step for the model. Note that there is great potential for improving the performance of the model at low values of τ by choosing a more sophisticated SDE scheme (e.g., Kloeden & Platen, 1992; Milstein & Tretyakov, 2004).

TABLE A.1 Time required to simulate the monochromatic, classical Holton–Lindzen–Plumb (HLP) model; the broadband multi-wave HLP (BBM-HLP) model with $N_\omega = 50$, and the three broadband stochastic HLP schemes (overdamped Langevin [ODL], hybrid [H], Eckermann–Lott–Guez–Maury [ELGM]), with $\tau = 0.02$. Each simulation was run until time $T = 100T_*$ on a single-processor laptop on a vertical grid with 5,000 grid points.

Model	Runtime (s)
HLP	28.56
BBM-HLP	1439
ODL	28.55
H	28.29
ELGM	28.16

TABLE A.2 Time required to simulate the overdamped Langevin model for several values of τ.

	$\tau = 0.005$	$\tau = 0.01$	$\tau = 0.02$	$\tau = 0.05$
Runtime (s)	126.1	66.81	34.06	34.65
Runtime/time step (s)	3.152×10^{-4}	3.341×10^{-4}	3.406×10^{-4}	3.465×10^{-4}