



Stockholms
universitet

Taming the Stochastic Ocean

A Numerical Study of the Stochastic
Primitive Equations with Varying Bathymetry

Sebastijan Babic

Kandidatuppsats 2026:14
Matematisk statistik
Maj 2026

www.math.su.se

Matematisk statistik
Matematiska institutionen
Stockholms universitet
106 91 Stockholm

Taming the Stochastic Ocean: A Numerical Study of the Stochastic Primitive Equations with Varying Bathymetry

Sebastijan Babić*

May 2026

Abstract

Ocean simulations cannot resolve every small-scale process (wind, turbulence, fine currents), so these are represented by a *random* forcing added at every time step, modelled here as a Q-Wiener process. This creates a numerical puzzle: making the time step smaller, which normally improves accuracy, can instead cause the surface pressure to fluctuate wildly from one step to the next. Since currents are driven by the pressure gradient, this instability feeds straight into them.

Theory shows the instability does not appear if the forcing is *solenoidal*, meaning the horizontal gradient of its vertical integral through the water column vanishes. On a flat sea floor, a *splitting scheme* controls the pressure gradients even for non-solenoidal forcing. This thesis tests by experiment what happens when the forcing is non-solenoidal and the sea floor is no longer flat. The instability is measured by the squared horizontal variation of the surface pressure gradients, summed over time steps and averaged over Monte Carlo samples. Across the four cases (solenoidal vs non-solenoidal, flat vs bumpy), splitting shrinks the measure by a factor of about 30 in the non-solenoidal cases and has no effect in the solenoidal ones. The measure does not diverge as Δt is refined to 10^{-3} ; it remains bounded across the tested range. A factor-of-ten gap between solenoidal-flat and solenoidal-bumpy is traced to extra terms that appear only on a sloped sea floor; rigorous bounds remain future work.

*Postal address: Mathematical Statistics, Stockholm University, SE-106 91, Sweden.
E-mail: sebastijan.sthlm@gmail.com. Supervisor: Josefin Ahlkrona.

Sammanfattning

Havssimuleringar kan inte upplösa alla småskaliga processer (vind, turbulens, små strömmar), så dessa representeras av en *slumpmässig* pådrivning vid varje tidssteg, här modellerad som en Q-Wienerprocess. Detta skapar en numerisk paradox: mindre tidssteg, som vanligtvis förbättrar noggrannheten, kan istället få yttrycket att fluktuera kraftigt från ett steg till nästa. Eftersom strömmarna drivs av tryckgradienten, sprids denna instabilitet direkt in i dem. Teori visar att instabiliteten inte uppträder om pådrivningen är *solenoidal*, det vill säga om den horisontella gradienten av dess vertikala integral genom vattenpelaren försvinner. På en plan havsbotten kontrollerar ett *splittingschema* tryckgradienterna även för icke-solenoidal pådrivning. Denna uppsats undersöker genom experiment vad som händer när pådrivningen är icke-solenoidal och havsbotten inte längre är plan. Instabiliteten mäts med den kvadrerade horisontella variationen av yttryckets gradienter, summerad över tidssteg och medelvärdesbildad över Monte Carlo-stickprov. Över de fyra fallen (solenoidal kontra icke-solenoidal, plan kontra kuperad) minskar splittingschemat måttet med en faktor på cirka 30 i de icke-solenoidala fallen och har ingen effekt i de solenoidala. Måttet divergerar inte när Δt förfinas till 10^{-3} ; det förblir begränsat över det testade intervallet. Ett gap på cirka en faktor tio mellan solenoidal-plan och solenoidal-kuperad spåras till extra termer som endast uppstår på en sluttande havsbotten; rigorösa gränser lämnas som framtida arbete.

Postal address: Mathematical Statistics, Stockholm University, SE-106 91, Sweden.

E-mail: sebastijan.sthlm@gmail.com. **Supervisor:** Josefin Ahlkrona

Acknowledgements

I would like to begin by acknowledging a few people who have helped me throughout my studies and throughout this thesis.

First, I want to thank my supervisor Josefin Ahlkrona, who has helped me throughout this entire thesis as though we were collaborating on a shared research project. Without her help this thesis would be, frankly, impossible.

I also want to thank Claudine von Hallern, who, in her consulting role during this thesis, provided insights into the stochastic aspects of the theory, especially the computational parts regarding the parallelisation.

A big thank you also goes to Zara Fuhrmann who helped me with proofreading and improving the overall structure of the thesis through our continuous discussions.

And last but certainly not least, thank you to my dad Sinisa Babic. Throughout my entire studies and in life you have always been there for me, supporting me regardless.

AI tools (Claude Opus 4.6 and 4.7, NotebookLM) were used to optimize MATLAB parallelisation, generate documentation, and assist in overcoming drafting blocks. A thorough final read-through was also performed by Claude 4.7.

Contents

Notation	1
Introduction	2
1.1 Motivation	2
1.2 Research background	3
1.3 Problem formulation	4
1.4 Aim and research questions	4
1.5 Scope and delimitations	5
1.6 Contributions	5
1.7 Outline of the thesis	7
Background and Mathematical Model	8
2.1 Ocean modelling and the primitive equations	8
2.2 Stochastic forcing	8
2.3 Geometric setting	12
2.4 Stochastic primitive equations	13
2.5 Boundary conditions	14
Numerical Methods	15
3.1 Key terminology	15
3.2 The finite element method	15
3.3 Discretization in time	17
3.4 Monte Carlo estimation	17
3.5 Why splitting methods are introduced	17
3.6 The splitting scheme	18
3.7 Coordinate transformation used to achieve varying bathymetry	18
3.8 Domain geometry and aspect ratio	19
Theoretical Results	20
4.1 Flat-bottom reference case	20
4.2 Stability bound in the flat case	21
4.3 Varying bathymetry	21
4.4 Remarks on bathymetry and boundary terms	23
Numerical Experiments	24
5.1 Experimental setup	24
5.2 Results	25
5.3 Robustness of the observed effects	31
Discussion and Conclusion	36
6.1 Final remarks	36
Appendix	39
A.1 Robustness to domain aspect ratio	39
A.2 Implementation parameters	39
A.3 Computational implementation details	41
Bibliography	42

Notation

This page collects the recurring symbols used throughout the thesis. Concepts that need a sentence rather than a symbol (test function, weak form, stiff problem, and similar) are defined in Section 3. Numerical parameters with concrete values used in the experiments (h_{\max} , τ_T , K_{sim} , etc.) are listed in Table 3.

Symbol	Meaning
Domain and geometry	
\mathcal{D}	Three-dimensional fluid domain, $\omega \times (-D(\vec{x}), 0)$.
ω	Horizontal cross-section, $\omega \subset \mathbb{R}^2$.
$\Gamma_s, \Gamma_b, \Gamma_l$	Surface, bottom, and lateral boundaries of \mathcal{D} .
$\vec{x} = (x, y)$	Horizontal position.
z	Vertical coordinate; surface at $z = 0$, seabed at $z = -D(\vec{x})$.
H	Reference (constant) vertical depth.
L	Horizontal extent of ω .
$D(\vec{x})$	Local water depth (bathymetry); $D \equiv H$ on a flat bottom.
Fields and physical parameters	
\mathbf{u}_H	Horizontal velocity vector, $\mathbf{u}_H = (u_x, u_y)$.
p_s	Surface pressure (lives on Γ_s , extended constantly in z).
p	Full pressure inside \mathcal{D} .
\mathbf{f}	Deterministic forcing term.
ν	Viscosity (taken constant and isotropic).
∇_H, Δ_H	Horizontal gradient and horizontal Laplacian.
$\bar{\mathbf{u}}_H$	Vertical average of \mathbf{u}_H over the local column $[-D(\vec{x}), 0]$.
Stochastic forcing	
Ω	Sample space of the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ (distinct from the spatial domain \mathcal{D}).
$W(t)$	Standard Wiener (Brownian) process.
$W^{Q(t)}$	Q-Wiener process; truncated to $N \times N$ modes in this thesis.
Q	Covariance operator of the noise; analogue of a covariance matrix.
$\beta_{k(t)}$	i.i.d. standard scalar Brownian motions.
e_{ij}	Spatial basis functions of the noise (sine modes on $\omega \times (-H, 0)$).
$\sigma_{ij}, \lambda_{ij}$	Noise amplitude and covariance eigenvalue, with $\lambda_{ij} = \sigma_{ij}^2$.
$B(\mathbf{u}_H) dW$	Stochastic forcing term in the momentum equation. Represents multiplicative noise in the general formulation; in this thesis only additive noise is used.
Splitting scheme	
ξ^n	Splitting (gradient) parameter at time step n ; vanishes on a flat column when the forcing is solenoidal.
η^n	Divergence-free remainder of the splitting.
r_s^{n+1}	Shifted pressure, $r_s^{n+1} = p_s^{n+1} - \Delta t^{-1} \xi^n$.
Discretisation	
Δt	Time step.
$\Delta_{n+1} W$	Brownian increment over one step, $W(t_{n+1}) - W(t_n) \sim \mathcal{N}(0, \Delta t)$.
N	Number of time steps; also the truncation index for noise modes (context-dependent).
φ_j, ψ_j	Finite-element basis functions on the volume mesh and surface mesh respectively.
$\mathbb{E}[\cdot]$	Expectation over the random forcing.
Function spaces	
$L^2(\mathcal{D})$	Functions square-integrable on the spatial domain \mathcal{D} .
$H^1(\mathcal{D})$	Sobolev space on \mathcal{D} : values and first derivatives both in $L^2(\mathcal{D})$.
$H_b^1(\mathcal{D})$	Subspace of $H^1(\mathcal{D})$ satisfying the boundary conditions of the problem.
(f, g)	L^2 inner product on the current domain, $\int fg$. With a subscript $(f, g)_\Gamma$ the integration is over the boundary piece Γ instead.

Introduction

1.1 Motivation

A central well-posedness question for stochastic partial differential equations is whether their numerical approximations remain controlled as the time step is refined. Adding random forcing to a PDE typically reduces the *regularity* of its solutions compared with the deterministic case (loosely, the solutions are less smooth, having fewer well-behaved derivatives, because each random kick injects roughness that the diffusion cannot fully smooth out), and this reduced regularity can interact badly with standard time-stepping schemes. In incompressible fluid problems, the pressure is the variable most exposed to this interaction. Irregular discrete pressure behaviour at small time steps causes large numerical errors in both pressure and velocity, an issue that refining the spatial mesh does not eliminate, see (Feng et al., 2021) and (Carelli & Prohl, 2012). Whether such pressure control carries over to the *stochastic primitive equations* is the subject of this thesis.

The primitive equations themselves are a standard model for large-scale oceanic and atmospheric flows (Lions et al., 1992) and are widely used in geophysical fluid dynamics, appearing for instance in operational ocean models such as NEMO (Madec et al., 2024) and MITgcm (MITgcm Group, 2018). When stochastic forcing is included, the model is better suited for representing unresolved fluctuations and uncertainty in the forcing, and provides a more faithful representation of real-world ocean dynamics, where sub-grid-scale processes and uncertain boundary data introduce inherent randomness (Tucciarone et al., 2025).

In the simplified form used throughout this thesis (linearised momentum equation, hydrostatic and Boussinesq simplifications), the stochastic primitive equations read

$$d\mathbf{u}_H = [\nu\Delta\mathbf{u}_H - \nabla_H p_s + \mathbf{f}] dt + B(\mathbf{u}_H) dW(t) \quad \text{in } \mathcal{D},$$

where $\mathbf{u}_H = (u_x, u_y)$ is the horizontal velocity, p_s the surface pressure, \mathbf{f} the deterministic forcing, ν the viscosity, and $W(t)$ a Q-Wiener process modelling the unresolved stochastic forcing. Writing

$$\overline{\mathbf{u}_H}(\vec{x}) := \frac{1}{D(\vec{x})} \int_{-D(\vec{x})}^0 \mathbf{u}_H(\vec{x}, z) dz$$

for the vertical average of the horizontal velocity over the local water column, the system is closed by the rigid-lid incompressibility constraint $\nabla_H \cdot \overline{\mathbf{u}_H} = 0$ on the surface. The operator B describes how the noise acts on the velocity field; this thesis uses additive noise throughout, so B does not in fact depend on \mathbf{u}_H in our experiments, but the general $B(\mathbf{u}_H)$ notation is retained for compatibility with the multiplicative-noise literature. A precise statement with the full geometric setting and assumptions is given in Section 2.4.

The horizontal gradient of the surface pressure, $\nabla_H p_s$, measures how rapidly the surface pressure varies along the ocean surface. Splitting schemes (schemes that decompose the stochastic forcing into a part that can be absorbed into the pressure variable and a residual that cannot) have been developed for the stochastic Stokes equations by (Feng et al., 2021) and extended to the stochastic primitive equations by (Hallern & Ahlkrone, 2026). The latter is what we study here, and it is designed to control precisely $\nabla_H p_s$: the finite-element¹ error analysis there bounds the discrete velocity error in terms of the time-accumulated squared

¹FEM: solve a PDE on a computer by approximating the solution as a linear function on each triangle of a mesh; the unknowns are the corner values.

norm of $\nabla_H p_s$, so any Δt -dependent growth in that norm propagates directly into the velocity-error bound.

The behavior of the numerical surface pressure is therefore studied through the quantity

$$E \left(\sum_{n=1}^N \Delta t \|\nabla_H p_s^n\|^2 \right), \quad (1.3)$$

which measures the expected value of the accumulated size of the discrete horizontal surface-pressure gradient over time on the ocean surface. Here p_s^n denotes the discrete surface pressure at time n , obtained from the full discrete pressure at the top vertical level p^n .

The argument of the expectation in Equation (1.3) is itself a random variable. It aggregates squared-gradient contributions from every time step of a single realisation of the discrete stochastic system, and changes from one realisation to the next. The stability question is therefore whether its expected value stays bounded as $\Delta t \rightarrow 0$.

1.2 Research background

For the stochastic Stokes equations, previous work shows that whether the discrete pressure remains controlled as the time step shrinks depends on the structure of the stochastic forcing (Carelli & Prohl, 2012; Feng et al., 2021). In the classical Stokes setting, a forcing f is called *solenoidal* when its divergence vanishes, $\nabla \cdot f = 0$, and *non-solenoidal* otherwise. For the primitive equations the relevant condition takes a slightly different form, since only the value of the pressure on the ocean surface enters the horizontal momentum equation: it is the depth-integrated horizontal forcing that must be horizontally divergence-free, and it is this primitive-equation notion of solenoidality that the splitting analysis below relies on. The relevant quantity is

$$E \left(\sum_{n=1}^N \Delta t \|\nabla p^n\|^2 \right), \quad (1.4)$$

and (Feng et al., 2021) show that, in the stochastic Stokes setting, this quantity can blow up as $\Delta t \rightarrow 0$ on a sufficiently fine mesh whenever the noise is non-solenoidal, an instability that ordinary implicit time-stepping² does not cure. (Carelli & Prohl, 2012) confirm the same divide for the stochastic Navier-Stokes equations: their Lemma 3.2 separates the solenoidal and non-solenoidal cases explicitly when establishing convergence rates, with the bounded-pressure conclusion available only in the solenoidal one. The main takeaway is that controlling the full discrete pressure p^n cannot be achieved by implicit time-stepping alone; the structure of the stochastic forcing is what decides whether the pressure stays bounded.

For the stochastic primitive equations, a splitting-based analysis in the flat-bottom case has been developed in (Hallern & Ahlkrona, 2026). The idea of the splitting is to rewrite the equations so that, instead of solving for the full pressure p^n , one solves a reformulated system in which the stochastic forcing on the right-hand side is solenoidal by construction. The original problem is thereby moved into the favourable regime of the Stokes analysis, and a useful orthogonality property specific to the flat-bottom geometry then allows the resulting surface pressure to be controlled. Concretely, (Hallern & Ahlkrona, 2026) establishes the bound

$$\mathbb{E} \left[\sum_{n=1}^N \Delta t \|\nabla_H p_s^n\|^2 \right] \leq \begin{cases} C, & \text{if the forcing is solenoidal} \\ C\Delta t + C\Delta t^{-1}, & \text{otherwise} \end{cases} \quad (1.5)$$

where C denotes any constant independent of Δt (this form, without a square root in the exponent, matches the corresponding bound for the stochastic Stokes equations in (Feng et al.,

²*Implicit time-stepping*: at each step the new value is defined by an equation that involves itself, so each step is found by solving for the new state rather than computing it directly from the current one. The pay-off is stability under large time steps.

2021) (eq. 1.5); the squared norm on the left-hand side is what doubles the exponent compared with a bound stated on $\|\nabla_H p_s\|$ rather than its square). The theoretical prediction is therefore qualitatively different in the two cases: under solenoidal forcing the bound is uniform in Δt , whereas under non-solenoidal forcing the upper bound grows as $\Delta t \rightarrow 0$ like Δt^{-1} . The non-solenoidal bound is still a substantial improvement over the non-solenoidal Stokes case, whose bound diverges faster. As we shall see in Section 5, the empirical Δt -dependence is much milder than the theoretical worst case in both regimes and similar across them (only the prefactor differs), which suggests that neither bound is tight at the tested Δt levels and motivates further analytical work.

While the theoretical pressure bound is established in the flat-bottom case by (Hallern & Ahlkrona, 2026), numerical experiments verifying the behavior of Equation (1.3) for the stochastic primitive equations have, to the author’s knowledge, not previously been carried out.

1.3 Problem formulation

The main quantity studied in this thesis is Equation (1.3), where p_s^n denotes the discrete surface pressure at time level n and Δt is the time step. The central question is whether this quantity remains bounded independently of the time step, or whether it grows as the discretization is refined. In the flat-bottom theory, boundedness can be obtained when the stochastic forcing satisfies a compatibility condition that may be described as being solenoidal in a primitive-equation sense. This suggests a natural contrast between an idealized setting, where pressure control is theoretically available, and more realistic settings, where the forcing may be non-solenoidal and the bathymetry may vary. The problem studied here is therefore twofold: first, to understand the mechanism behind the flat-bottom pressure estimate, and second, to investigate what happens when the assumptions behind that estimate are weakened.

Surface pressure turns out to be harder to control than the full pressure once the bathymetry varies: the cancellation mechanism that makes the flat-bottom proof work depends on the sea floor being level, and this assumption fails as soon as the depth changes with horizontal position.

1.4 Aim and research questions

The aim of this thesis is to study the behavior of the numerical surface pressure in a stochastic primitive-equation model, with particular focus on the quantity Equation (1.3). On the theoretical side, the goal is to re-derive the pressure estimate Equation (1.5) in a self-contained way under the primitive-equation analogue of the solenoidal condition, and to examine how its analytical structure changes when the bathymetry is no longer flat.

On the numerical side, the goal is to investigate whether the pressure remains controlled when the assumptions behind Equation (1.5) are relaxed, in particular for non-solenoidal forcing, more complicated noise terms, and varying bathymetry. For non-solenoidal forcing the bound Equation (1.5) is not expected to hold *without* the splitting scheme, since it is precisely the splitting that rewrites the equations so the effective stochastic forcing becomes solenoidal; the non-solenoidal experiments therefore test whether the splitting succeeds in restoring the bound.

The thesis is guided by the following research questions:

1. Under which assumptions on the stochastic forcing can one obtain the $C\Delta t + C\Delta t^{-1}$ pressure bound on the accumulated surface-pressure-gradient norm in the flat-bottom case, see Equation (1.5)?
2. How is the analytical structure of the pressure estimate affected by non-flat bathymetry?
3. Across the four configurations (solenoidal versus non-solenoidal forcing, flat versus varying bathymetry), does the accumulated surface-pressure-gradient norm Equation (1.3) remain bounded as the time step becomes small, or does it show signs of blow-up?

1.5 Scope and delimitations

The present study is restricted in several ways, both by design and by the constraints of a bachelor’s thesis.

On the modelling side, the stochastic primitive equations considered here are a simplified version of the full system. The non-linear convection term seen in the full non-linear primitive equations is omitted, yielding a linearised momentum equation. Also, temperature and salinity dynamics are not included. By removing the non-linear dynamics which naturally produce fluid turbulence and chaotic secondary instabilities, we were able to isolate the pressure instability. If convection had been included, any resulting numerical blow-ups could have been ambiguously attributed to standard non-linear advective instabilities rather than the fractional splitting scheme failing geometrically. Similarly, the viscosity is taken to be isotropic and constant, whereas in realistic ocean models it is often anisotropic and may depend on the local flow. Due to a lack of theory regarding the case of non-isotropic viscosity we will stick to this assumption.

The stochastic forcing implemented in these experiments relies on a relatively simple, purely mathematical construction: a truncated Q-Wiener process built from $N \times N$ two-dimensional sine modes, with prescribed amplitude coefficients that damp high-frequency contributions. In reality, actual oceanic stochastic forcing arises from thermodynamic changes, unpredictable wind stresses, see (Binz et al., 2024) for a more careful handling of specifically wind, turbulence, and much more. Many of these are chaotic and inherently contain both solenoidal and non-solenoidal elements. Because realistic geophysical noise naturally violates the idealized zero-mean assumption at all times, the splitting scheme’s pressure reconstruction would in practice operate in the non-solenoidal regime, where the experiments of this thesis indicate it provides a constant reduction of the surface-pressure-gradient norm but does not eliminate it. Worth noting that we may be ignoring some possibly important physical properties that turbulence brings to this model.

On the analytical side, the investigation of the varying-bathymetry case is carried out at a formal level. The main result is the identification of the additional boundary and volume terms that arise when the depth depends on the horizontal variable. A complete rigorous proof of an analogue of the flat-bottom pressure estimate in the variable-bathymetry case was not achieved within the scope of this thesis, in part because the orthogonality property that underpins the flat-bottom argument does not carry over directly, and the correction terms that appear have not been controlled in a closed form.

On the numerical side, the domain is taken with aspect ratio $H/L = 1/5$, anisotropic enough to give the bathymetric warping a genuine effect on the mesh while keeping resolution requirements manageable, and similar to ratios reported in (Chacón Rebollo et al., 2012). The bathymetry configurations tested are limited to a mild warp (low-amplitude sinusoidal bumps), since a mild warp is already enough to break the flat-bottom orthogonality that underpins the analytical estimate without obscuring the underlying mechanism with steep-slope effects. Furthermore, the number of Monte Carlo samples (independent simulation runs whose sample average estimates the expected values) used to approximate expectations is moderate, inspired directly by (Feng et al., 2021), and no comparison with real oceanographic data is attempted as that is not the focus of this thesis.

Despite these limitations, the study provides a clear numerical picture of how the surface-pressure stability is affected by the type of stochastic forcing and the presence of varying bathymetry, and identifies the analytical mechanism responsible for the loss of the flat-bottom orthogonality property.

1.6 Contributions

Prior work. The splitting scheme of (Hallern & Ahlkrona, 2026) delivers the bound Equation (1.5) on the accumulated surface-pressure-gradient norm in the flat-bottom case, resting on a flat-bottom orthogonality identity (Lemma 3.2 of (Hallern & Ahlkrona, 2026)). The earlier

stochastic-Stokes analyses of (Feng et al., 2021) and (Carelli & Prohl, 2012) set the stage by isolating non-solenoidal noise as the source of pressure blow-up at small Δt .

This thesis. The novel contributions are:

1. Implementation of varying bathymetry via a vertical coordinate transformation $z \mapsto zD(\vec{x})/H$ that warps a single triangulation to fit either a flat or a bumpy floor, so the four experimental configurations can be run on the same mesh without remeshing.
2. Parallelisation of the Monte Carlo loop using MATLAB's `parfor`, with the Stokes-solver LU factorisation precomputed once per Δt level and shared across all samples; this turned an otherwise prohibitive $K_{\text{sim}} \times M$ batch into a tractable run on commodity hardware.
3. Numerical verification of the surface-pressure-gradient stability across all four configurations (solenoidal versus non-solenoidal forcing crossed with flat versus varying bathymetry), confirming that the splitting scheme remains effective empirically beyond the regime in which it is currently proven.
4. Identification that a naive transfer of the truncated Q-Wiener basis from the reference column $[-H, 0]$ to the physical column $[-D(\vec{x}), 0]$ silently breaks the zero-mean property the splitting relies on, and that building the basis on the local depth $D(\vec{x})$ restores it pointwise in \vec{x} .
5. Analytic identification of the boundary integral proportional to $\nabla_H D$ and of the Leibniz-rule volume corrections that survive in the varying-bathymetry analogue of the flat-bottom orthogonality identity.
6. A robustness check showing that the surface-pressure stability is qualitatively insensitive to the high-frequency damping rate of the truncated Q-Wiener noise: rerunning the splitting experiment across three different decay exponents on the noise eigenvalues leaves the boundedness behaviour as $\Delta t \rightarrow 0$ intact in each case (Section 5.3.4).

A closed-form bound for the varying-bathymetry case is left open.

1.7 Outline of the thesis

The remainder of this thesis is structured as follows:

[Section 2](#) provides the theoretical foundation: ocean modelling, the derivation of the primitive equations, and the truncated Q-Wiener process used to model unresolved stochastic fluctuations. It then specifies the mathematical framework, including the geometric setting, boundary conditions, and the formulation of the stochastic primitive equations, highlighting the distinction between flat and varying bathymetry.

[Section 3](#) details the numerical discretisation strategy: the finite-element method, Euler-Maruyama time stepping, Monte Carlo estimation of expected values, the splitting scheme, and the coordinate transformation used to realise varying bathymetry.

[Section 4](#) presents the analytical investigation. It reviews the orthogonality property in the flat-bottom case and demonstrates how this property is lost under varying bathymetry, due to additional boundary terms and the failure of the volume-integral cancellation.

[Section 5](#) presents the numerical experiments. The experimental setup and the Monte Carlo algorithm are summarised first. The main results then compare the stability of the accumulated surface pressure gradient norm across the four combinations of bathymetry and forcing, with spatial-field snapshots, fitted Δt -scaling slopes, and variance behaviour across Monte Carlo samples. A separate robustness section follows: a mesh-refinement diagnostic distinguishing numerical from formulation error, a Monte Carlo reliability check on sample-size adequacy, a paired Wilcoxon signed-rank test on the statistical significance of the splitting effect, and a robustness check across three different decay rates for the truncated Q-Wiener noise.

Finally, [Section 6](#) summarises the numerical and analytical conclusions, maps the findings to the research questions, and outlines directions for future work. Supplementary material is collected in [Appendix A](#): a velocity-norm sanity check, a robustness check on the domain aspect ratio ([Appendix A.1](#)), and the implementation parameters used in the experiments ([Appendix A.2](#)).

Background and Mathematical Model

2.1 Ocean modelling and the primitive equations

Ocean modelling is a central area of fluid dynamics and climate science. Its development has been driven both by advances in mathematical theory and by improvements in computational technology.

On the computational side, increasingly efficient numerical algorithms, better software environments, and more powerful hardware have made it possible to simulate ocean flows on finer meshes, over longer time intervals, and with more realistic physical effects included. Large-scale ocean circulation models now play an important role in climate science, weather prediction, and the study of transport phenomena in the ocean, much of which is discussed in (Griffies, 2004) and (Randall et al., 2007).

The (deterministic) *primitive equations*, in the form used by (Hallern & Ahlkrona, 2026) (and the form solved in the implementation), take the form

$$\begin{aligned} \partial_t \mathbf{u}_H &= \nu \Delta \mathbf{u}_H - \nabla_H p_s + \mathbf{f} & \text{in } \mathcal{D}, \\ \nabla_H \cdot \overline{\mathbf{u}_H} &= 0 & \text{on } \Gamma_s, \end{aligned} \quad (2.1)$$

where \mathcal{D} is the three-dimensional fluid domain (a horizontal cross-section over a water column of local depth $D(\vec{x})$) and Γ_s is its top surface; $\mathbf{u}_H = (u_x, u_y)$ is the horizontal velocity, p_s the surface pressure (defined on Γ_s and extended constantly in the vertical), \mathbf{f} a deterministic forcing, ν the viscosity, and

$$\overline{\mathbf{u}_H}(\vec{x}) := \frac{1}{D(\vec{x})} \int_{-D(\vec{x})}^0 \mathbf{u}_H(\vec{x}, z) \, dz$$

is the vertical average of \mathbf{u}_H over the water column $[-D(\vec{x}), 0]$.

This form is obtained from the full Navier-Stokes equations by introducing physically motivated simplifying assumptions, most notably the *hydrostatic approximation* and the *Boussinesq approximation*; see (Madec et al., 2024) Chapter 1.1 for the full set of assumptions.

- The hydrostatic approximation: at large horizontal scales the ocean's vertical motion is small, so the vertical pressure balance reduces to $\partial_z p = -\rho g$ (with ρ the fluid density and g the gravitational constant). Integrating vertically gives $p(\vec{x}, z) = p_s(\vec{x}) + \rho g z$, so only the surface pressure appears in the horizontal momentum equation.
- The Boussinesq approximation: density variations are neglected except where they contribute to buoyancy, so the fluid is treated as effectively incompressible. Combined with the rigid-lid condition at the surface and no-penetration at the seabed, vertical integration of the incompressibility constraint eliminates the vertical velocity w and yields the divergence-free constraint $\nabla_H \cdot \overline{\mathbf{u}_H} = 0$ on the surface.

In addition, the non-linear convection term $(\mathbf{u} \cdot \nabla) \mathbf{u}_H$ is dropped from the momentum equation, since it is not relevant to the pressure-regularity question studied here.

2.2 Stochastic forcing

In deterministic fluid models, external effects are usually represented by a prescribed forcing term. However, such forcing is often not known exactly and may vary in an irregular manner. This motivates the introduction of stochastic forcing, where the uncertainty is modeled through a random term in the momentum equation.

This stochastic forcing may be interpreted as representing small-scale effects that are not resolved explicitly by the deterministic model, for example rapidly varying wind stress such as in (Binz et al., 2024) or other uncertain external influences such as turbulence.

Stochastic forcing may be introduced into a deterministic model in many ways, we will however use the way of *Brownian motion* in the form of a *truncated Q -Wiener process*. This will lead us to the concept of stochastic differential equations, and by extension to partial differential equations since we then acquire a way of representing the space.

2.2.1 Brownian motion and Wiener processes

We define *Brownian motion*, also known as a *Wiener process*, in one dimension as follows.

Definition 2.1: A standard one-dimensional Wiener process $W(t)$ is a stochastic process, satisfying:

1. $W(0) = 0$,
 2. $W(t) - W(s) \sim \mathcal{N}(0, t - s)$ for $0 \leq s < t$,
 3. increments on disjoint intervals are independent,
 4. sample paths $t \mapsto W(t)$ are continuous almost surely.
- (2.3)

By a stochastic process it is meant that $W(t)$ is a parametrized collection of random variables $\{W(t)\}_{t \in T}$ defined on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ and assuming values in \mathbb{R} .

A key consequence is that W is nowhere differentiable with probability one, which means stochastic differential equations cannot be interpreted in the classical sense and require Itô calculus³.

Figure 1 shows five independent realisations on $t \in [0, 1]$, each constructed by cumulating i.i.d. $\mathcal{N}(0, \Delta t)$ increments. All paths start at $W(0) = 0$ and have mean zero at every t , but individual trajectories diverge from one another and exhibit the rough, jagged structure that nowhere-differentiability imposes.




Figure 1: Five independent sample paths of a standard one-dimensional Wiener process $W(t)$ on $[0, 1]$. All paths share $W(0) = 0$ and $\mathbb{E}[W(t)] = 0$; the visible spread reflects the variance $\text{Var}(W(t)) = t$, and the jaggedness reflects the almost-sure non-differentiability of Equation (2.3).

³*Itô calculus*: a framework for integrating against a Brownian path W . Because W is nowhere differentiable, the ordinary rules of calculus break down; the Itô calculus defines what $\int g(s) dW(s)$ means and provides a modified chain rule (the Itô formula) for working with stochastic integrals.

2.2.2 Truncated Q-Wiener processes

The forcing in the primitive equations varies in both space and time, so we need a random object that varies across both. Brownian motion $W(t)$ alone produces only a single random function of time. The *Q-Wiener process* is the natural spatial generalisation, built as a sum of fixed spatial patterns each driven by its own independent Brownian motion. The full rigorous construction is in (Lord et al., 2014) Chapter 2.3; the form below is enough for what follows.

Definition 2.2: Let $\{e_k\}_{k \geq 1}$ be an orthonormal basis on the spatial domain, $\beta_k(t)$ independent standard scalar Wiener processes, and $\sigma_k \geq 0$ amplitude coefficients with

$$\sum_{k=1}^{\infty} \sigma_k^2 < \infty. \quad (2.4)$$

The associated Q-Wiener process is

$$W^Q(t) = \sum_{k=1}^{\infty} \sigma_k e_k \beta_k(t). \quad (2.5)$$

The infinite sum cannot be simulated as written, so in practice we work with its N -mode truncation $W_N(t) = \sum_{k=1}^N \sigma_k e_k \beta_k(t)$, keeping only the first N spatial patterns. The discarded modes contribute zero variance to W_N ; everything else is unchanged. Indexing the kept modes by a pair (i, j) rather than a single index k , we use

$$W(t) = \sum_{i=1}^N \sum_{j=1}^N \sigma_{ij} e_{ij}(x, z) \beta_{ij}(t), \quad (2.6)$$

A similar form can be seen in (Carelli & Prohl, 2012). Here H denotes the reference vertical depth and L the horizontal extent of the domain (both made precise in the Geometric Setting of Section 2.3 below). The three ingredients of Definition 2.2 take the following form in this context:

- **Amplitude coefficients,**

$$\sigma_{ij} = \frac{1}{\pi^2(i^2/H^2 + j^2/L^2)}, \quad (2.7)$$

control how much energy each mode (i, j) contributes; the decay in i, j damps high-frequency modes and concentrates the noise at large scales. The squared coefficients $\lambda_{ij} = \sigma_{ij}^2$ are the eigenvalues of the covariance operator Q .⁴

- $e_{ij}(x, z) = \sin(\pi j x/L) \psi_i(z)$ are the **spatial basis functions**, an orthonormal set of sine modes on the (x, z) -domain. They specify the spatial patterns the noise is allowed to excite. The vertical factor $\psi_i(z)$ is chosen differently in the solenoidal and non-solenoidal cases (see Equation (2.10)).
- $\beta_{ij}(t)$ are **independent standard scalar Wiener processes**, one per mode. They are the source of randomness: each mode evolves in time as its own independent Brownian motion, and it is through the β_{ij} that different sample paths of $W(t)$ differ.

Figure 2 shows four independent realisations from the same Q-Wiener process. The shared large-scale structure reflects the eigenvalue decay introduced in Equation (2.7), which suppresses high-frequency modes and concentrates energy at low wavenumbers. The variability between realisations illustrates the random nature of the forcing: same covariance operator Q , different sample paths.

⁴*Covariance operator Q* : the infinite-dimensional analogue of a covariance matrix; its eigenvalues say how much each spatial mode contributes to the noise.




Figure 2: Four independent realisations of the Q -Wiener stochastic forcing field $dW(\vec{x}, z)$ used to drive Equation (2.16). Red regions indicate forcing in the positive x -direction, blue in the negative.

2.2.3 Solenoidal compatibility of the stochastic forcing

We use additive noise throughout: the operator B does not depend on \mathbf{u} . The notation $B(\mathbf{u}) dW(t)$ is retained in the general formulation, where the more general multiplicative case is treated in the SPDE literature under Lipschitz and linear-growth assumptions on B .

A separate, and central, requirement on the noise is that it be compatible with the divergence structure of the primitive equations. The pressure estimate in the flat-bottom analysis (and by extension in the varying-bathymetry analysis) relies on a condition that may be described as the forcing being solenoidal in a primitive-equation sense, by which we mean that the noise coefficient B satisfies

$$\nabla_H \cdot \int_{-D(\vec{x})}^0 B(\mathbf{u}_H) dW(t) dz = 0,$$

i.e. the depth-integrated horizontal forcing is divergence-free in the horizontal sense. This is the forcing-side counterpart to the velocity constraint Equation (2.17) imposed on the dynamics.

2.2.4 Stochastic forcing term

To obtain a concrete noise model, we discretise the abstract Q -Wiener process introduced in Definition 2.2. Recall that a Q -Wiener process has the form $W^{Q(t)} = \sum_k \sigma_k e_k \beta_k(t)$, where e_k are orthonormal basis functions, β_k are independent scalar Wiener processes, and σ_k are the amplitude coefficients. In our setting, the spatial domain is two-dimensional (one horizontal and one vertical direction), so we index the eigenfunctions by a pair (i, j) and choose a product basis of sine functions,

$$e_{ij}(x, z) = \sin\left(\frac{\pi j x}{L}\right) \psi_i(z),$$

where the horizontal part $\sin(\pi j x/L)$ is a standard Fourier mode on the domain of length L , and the vertical part $\psi_i(z)$ is defined as follows:

$$\psi_i(z) = \begin{cases} \sin\left(\frac{\pi i z}{H}\right) & \text{if non-solenoidal} \\ \sin\left(\frac{2\pi i z}{H}\right) & \text{if solenoidal.} \end{cases} \quad (2.10)$$

With the amplitude coefficients Equation (2.7) introduced earlier, higher-frequency modes contribute less to the noise. The corresponding covariance eigenvalues are $\lambda_{ij} = \sigma_{ij}^2$, and the summability condition Equation (2.4) is satisfied since the sum decays as $(i^2 + j^2)^{-2}$ for large i, j , ensuring finite total variance as required by Definition 2.2.

The amplitude coefficients in Equation (2.7) are our own choice, designed to damp high-frequency modes. The precise decay rate is not inherited from a particular reference but a modelling choice we made for this thesis, and a robustness check in Section 5.3.4 shows that the qualitative pressure-stability behaviour is insensitive to it.

2.2.4.1 Solenoidal versus non-solenoidal forcing

The flat-bottom pressure estimate Equation (1.5) requires the stochastic forcing to be *solenoidal in the primitive-equation sense* (the compatibility condition from Section 2.2.3). The splitting scheme (developed in Section 3.6) rewrites the equations so that the effective forcing satisfies this condition even when the noise itself does not; for present purposes it suffices to know that the bound holds whenever the noise already satisfies the condition, and the vertical basis ψ_i in Equation (2.10) is what determines whether it does.

The condition is on the *stochastic forcing* $B(\mathbf{u}) dW$, not on the velocity \mathbf{u} itself. The velocity obeys its own divergence-free constraint Equation (2.17), which is a separate matter; setting $\bar{\mathbf{u}} = 0$ is one way to satisfy it, but not the only way. A simple sufficient way to enforce solenoidality of the **forcing** is to choose the vertical noise basis so that each mode ψ_i has vanishing vertical mean, $\int_{-H}^0 \psi_i(z) dz = 0$. On a flat column this makes $\overline{B(\mathbf{u}) dW} = 0$ and the compatibility condition holds. The construction below is designed to deliver this property.

The two basis choices in Equation (2.10) are where this distinction enters. In the non-solenoidal case,

$$\psi_i(z) = \sin\left(\frac{\pi iz}{H}\right),$$

which has nonzero vertical mean for odd i ,

$$\int_{-H}^0 \sin\left(\frac{\pi iz}{H}\right) dz = \frac{H}{\pi i}(1 - \cos(\pi i)) \neq 0 \quad \text{for odd } i,$$

so the compatibility condition fails. In the solenoidal case,

$$\psi_i(z) = \sin\left(\frac{2\pi iz}{H}\right),$$

which satisfies $\int_{-H}^0 \psi_i(z) dz = 0$ for all $i \geq 1$, the condition holds. This zero-mean property is the analytic content of “solenoidal in the primitive-equation sense”, and it is the condition under which the flat-bottom pressure estimate (Equation (1.5)) can be obtained. A vanishing vertical mean is therefore a **sufficient**, not necessary, route to compatibility.

Two routes are available to enforce this zero-mean condition under varying bathymetry: one can either keep the basis on the reference height H , $\sin(2\pi iz/H)$, in which case the vertical integration domain becomes $[-D(x), 0]$ and the analytic zero-mean property is destroyed whenever $D(x) \neq H$; or rebuild the basis on the local depth, $\sin(2\pi iz/D(x))$, so that $\int_{-D(x)}^0 \sin(2\pi iz/D(x)) dz = 0$ holds pointwise in x . The implementation takes the second route throughout.

2.3 Geometric setting

We work on a three-dimensional fluid domain representing a vertical water column above a horizontal region. The geometry is chosen so as to reflect the oceanographic interpretation of the primitive equations while still remaining simple enough for analysis and computation.

Let $\omega \subset \mathbb{R}^2$ denote the horizontal domain. Its points are written as $\vec{x} = (x, y)$. The vertical coordinate is denoted by z , where the upper boundary corresponds to the ocean surface and the lower boundary corresponds to the seabed. Hence the full three-dimensional domain is written as

$$\mathcal{D} := \{(\vec{x}, z) \in \mathbb{R}^3 : \vec{x} \in \omega, -D(\vec{x}) < z < 0\}. \quad (2.14)$$

The upper boundary $z = 0$ is referred to as the rigid lid or surface, while the lower boundary $z = -D(\vec{x})$ describes the bathymetry. We denote by $H > 0$ the reference vertical depth (the constant depth against which the bathymetry is compared throughout) and by L the horizontal extent of ω .




Figure 3: The geometric setting as defined by Equation (2.14)

2.3.1 Flat bathymetry versus varying bathymetry

In the flat-bottom case one assumes that

$$D(\vec{x}) = H,$$

where $H > 0$ is constant. Then the bottom boundary is horizontal and several geometric boundary terms simplify. The flat-bottom assumption is used at several places in the existing derivation since it leads to proper cancellation of terms.

In the varying-bathymetry case, however, $D = D(\vec{x})$ is allowed to depend on the horizontal position. Then $\nabla_H D \neq 0$ in general, and the outward normal vector at the bottom acquires a non-zero horizontal component. **This geometric change is precisely what generates the additional boundary contributions** that the Analysis chapter investigates in detail. Extending the method to variable bathymetry is one of the main open directions: the splitting scheme was originally developed and proven stable under the flat-bottom assumption, and its behavior under spatially varying $D(\vec{x})$ has not yet been analyzed.

2.4 Stochastic primitive equations

Combining the deterministic momentum balance of Section 2.3 with the Q-Wiener forcing of Section 2.2.4 turns the model into a stochastic PDE in Itô form, with the velocity \mathbf{u}_H now a stochastic process and the increment dW replacing a classical right-hand side. In the simplified setting considered in this thesis, the stochastic primitive equations take the form

$$d\mathbf{u}_H = [\nu \Delta \mathbf{u}_H - \nabla_H p_s + \mathbf{f}] dt + B(\mathbf{u}_H) dW(t) \in \mathcal{D}, \quad (2.16)$$

together with

$$\nabla_H \cdot \overline{\mathbf{u}_H} = 0 \quad \text{on } \Gamma_s. \quad (2.17)$$

Here, $\overline{\mathbf{u}_H}$ denotes the vertical average of \mathbf{u}_H , defined by

$$\overline{\mathbf{u}_H}(\vec{x}) := \frac{1}{D(\vec{x})} \int_{-D(\vec{x})}^0 \mathbf{u}_H(\vec{x}, z) dz. \quad (2.18)$$

p_s denotes the surface pressure and \mathbf{f} is the deterministic forcing introduced in Equation (2.1). The stochastic part is given by $W(t)$, a Q-Wiener process as defined in Equation (2.5), and $B(\mathbf{u}_H)$, which describes how the noise acts on the velocity field, see Section 2.2.3 and Section 2.2.4.

Existence and uniqueness of solutions to Equation (2.16) are taken as given. The analytical foundations for stochastic primitive equations of this type are treated in (Hallern & Ahlkrona, 2026) and the SPDE literature cited there. This thesis focuses on the behaviour of the discrete approximation rather than on the continuous well-posedness theory.

It is the latter that proves problematic in our model since it provides an uncontrolled element. The non-solenoidal branch of Equation (2.10) turns out, in our experiments, to exhibit a mild Δt -dependent growth in the accumulated pressure-gradient norm (much slower than the theoretical Δt^{-1} worst case of Equation (1.5)), while the solenoidal branch, provided the basis is built on the local depth $D(\vec{x})$, keeps that norm bounded as the flat-bottom theory predicts.

The horizontal variables describe the large-scale motion along the surface of the ocean, and the horizontal differential operators are denoted by ∇_H and Δ_H representing $\begin{pmatrix} \partial_x \\ \partial_y \end{pmatrix}$ and $\partial_{xx} + \partial_{yy}$ respectively.

2.5 Boundary conditions

To complete the model, we specify boundary conditions on the three pieces of the domain boundary. At the ocean surface

$$\Gamma_s := \{(\vec{x}, z) \in \mathbb{R}^3 : z = 0, \vec{x} \in \omega\}$$

we impose a rigid-lid condition $w = 0$ on Γ_s : the surface is taken to be fixed in time so that surface waves are neglected, which is standard for large-scale circulation since wave amplitudes are far too small to affect the dominant dynamics (Griffies, 2004). The horizontal velocity at the surface is left unconstrained in the analysis; the experiments use a *homogeneous Dirichlet* boundary condition there (the value of the field is prescribed and set to zero on the boundary). On the lateral boundary

$$\Gamma_l := \{(\vec{x}, z) \in \mathbb{R}^3 : \vec{x} \in \partial\omega, -D(\vec{x}) \leq z \leq 0\}$$

periodic conditions are imposed in the horizontal directions, $\mathbf{u}_H(x + L, z) = \mathbf{u}_H(x, z)$ and similarly for w and p , so the lateral boundary becomes inert in all subsequent integrations by parts. On the bottom boundary

$$\Gamma_b := \{(\vec{x}, z) \in \mathbb{R}^3 : z = -D(\vec{x}), \vec{x} \in \omega\}$$

we impose *no-slip*, $\mathbf{u} = 0$ on Γ_b , which sets the full velocity to zero at the seabed (so the fluid does not slide along the bottom) and simultaneously enforces the physical impermeability requirement that no fluid crosses the seabed. The geometric obstruction from varying bathymetry analysed below comes from Γ_b rather than from Γ_l or Γ_s .

Numerical Methods

The stochastic primitive equations considered in this thesis are too complicated to solve analytically in closed form, and one must therefore rely on numerical time-stepping and spatial-discretization methods. We discretise in space using the finite-element method (FEM, Section 3.2) and in time using the Euler-Maruyama scheme (Section 3.3).

The FEM discretisation itself is not central to this thesis: the analytical and empirical findings would carry through under any reasonable spatial discretisation. It does, however, motivate the central quantity tracked throughout the thesis. The FEM velocity-error analysis of (Hallern & Ahlkrona, 2026) bounds the discrete velocity error by a sum of standard spatial-discretisation terms together with a contribution proportional to the time-accumulated squared norm of the discrete pressure gradient, $\mathbb{E}[\sum_n \Delta t \|\nabla_{HP} p_s^n\|^2]$. Controlling this proxy therefore controls the velocity error itself, which is why it serves as the stability diagnostic studied throughout Section 5.

The splitting method we are using in our analysis comes directly from (Hallern & Ahlkrona, 2026) which itself extends the splitting method for Stokes equations of (Feng et al., 2021).

3.1 Key terminology

Before proceeding, we briefly recall some terms that appear throughout this chapter.

- **Test function:** A function chosen from a suitable function space, used to reformulate equations in *weak* (integrated) form. Rather than requiring an equation to hold at every point, we require it to hold when multiplied by any test function and integrated over the domain.
- **Testing an equation:** Multiplying both sides by a test function and integrating over the domain.
- **Order of convergence:** A method has order p if its error behaves like $(\Delta t)^p$. Halving the step size reduces the error by a factor 2^p .
- **Strong vs weak convergence:** *Strong convergence* measures the pathwise error, $\mathbb{E}[\|u(t_n) - u^n\|] = O(\Delta t^p)$, that is, how close individual trajectories are to the exact solution. *Weak convergence* measures the error in expected values, $|\mathbb{E}[\varphi(u(t_n))] - \mathbb{E}[\varphi(u^n)]| = O(\Delta t^p)$ for sufficiently smooth test functionals φ , that is, how close the statistics of the discrete solution are to those of the exact solution.
- **Explicit method:** A scheme that computes the next state purely from the current state, without solving an equation. The simplest example is the explicit Euler method.
- **Stiff problem:** A problem where some terms vary on much faster time scales than others, forcing explicit methods to take very small time steps to remain stable. Implicit methods (which solve for the next state instead) handle stiffness without this restriction.
- **Energy estimate:** An inequality bounding a quadratic norm of the solution (its “energy”) by the data and initial values. It is derived by testing the equation against a carefully chosen function, typically the solution itself or one of its gradients, and using integration by parts so that the remaining terms can be controlled.

3.2 The finite element method

As noted at the start of this chapter, the spatial discretisation in this thesis uses the finite-element method (FEM) and the time stepping uses Euler-Maruyama; the FEM details collected in this section are standard and not central to the thesis’s main concepts, but are recorded for

completeness and because they fix the function spaces that the velocity-error bound in Section 3 is stated in.

The finite-element method approximates the velocity and pressure as piecewise-linear functions on a triangular mesh. The catch is that a piecewise-linear function has no second derivative anywhere inside a triangle, so a PDE term like $\nu \Delta \mathbf{u}_H$ cannot be evaluated pointwise. The standard fix is to rewrite the equation in an *integrated* form that only requires *first* derivatives, which the piecewise-linear approximation does supply as constant slopes on each triangle.

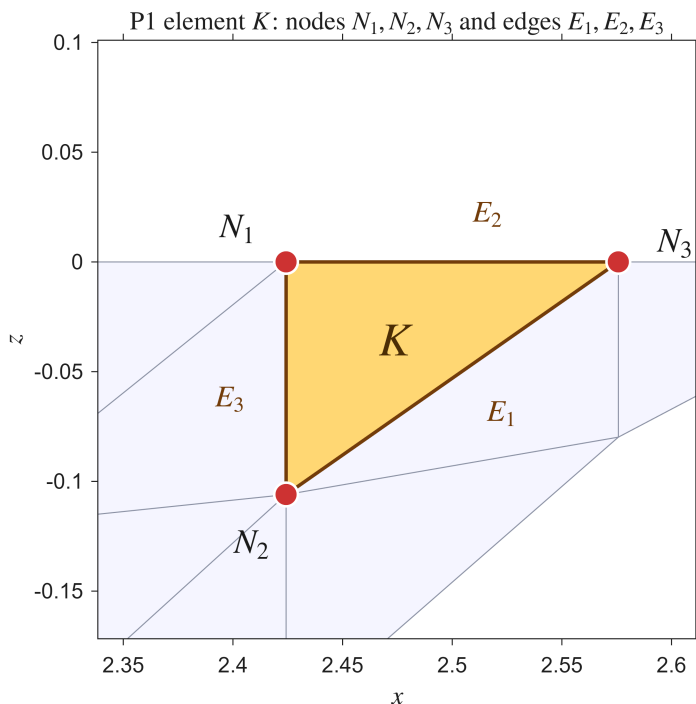
The recipe is:

1. Multiply both sides of the PDE by a smooth *test function* v .
2. Integrate over the domain \mathcal{D} .
3. Use integration by parts to move one derivative from \mathbf{u}_H onto v .

The resulting integrated identity is the *weak formulation*; the full derivation is standard and we refer to (John, 2016). We use piecewise-linear elements for both velocity and pressure (the cheapest equal-order choice, P^1 - P^1); since this pair fails the inf-sup stability criterion⁵ and the pressure can develop spurious oscillations, we add the standard *Brezzi-Pitkäranta stabilisation* with parameter $\tau_T = 0.1h^2$, where h is the maximum element size.

3.2.1 Triangular mesh and finite element spaces

The domain \mathcal{D} is partitioned into non-overlapping triangles of maximum size h_{\max} . On each triangle, the velocity and pressure are piecewise linear (P^1), so each field is determined by its values at the triangle vertices (the nodes). Under the coordinate transformation of Section 3.7 the originally uniform mesh is warped to follow the bathymetry, producing anisotropic elements wherever the depth varies, see Figure 4.



The mesh is generated with maximum element size h_{\max} controlling the spatial resolution, and warped to follow the bathymetry under the coordinate transformation of Section 3.7. Substituting this piecewise-linear ansatz into the weak form reduces the continuous PDE to a finite-dimensional system of stochastic ODEs in the nodal values, which the Euler-Maruyama scheme of the next section advances one Δt -step at a time.

Figure 4: A visualization of the triangular mesh of the finite element method, after the vertical coordinate transformation $z \mapsto zD(\bar{x})/H$ we implemented so that a single triangulation fits either a flat or a bumpy floor without remeshing.

⁵*Inf-sup (LBB) condition:* a compatibility requirement between the velocity and pressure spaces that mixed FEM pairs must satisfy to produce a stable pressure; equal-order P^1 - P^1 fails it generically, which is what motivates the stabilisation term below.

3.3 Discretization in time

With the spatial discretisation of Section 3.2 reducing the PDE to a finite-dimensional system of stochastic ODEs, what remains is to advance through time. The simulation evaluates the equations at the discrete times $t_n = n\Delta t$ for $n = 0, 1, \dots, N$ with $N\Delta t = T$ the final simulation time.

For a generic ordinary differential equation $\dot{x} = f(x)$, the (explicit) Euler scheme reads

$$x^{n+1} = x^n + \Delta t f(x^n), \quad (3.1)$$

advancing one Δt -step at a time from the current state. Augmenting this with a Brownian increment $\Delta_{n+1}W := W(t_{n+1}) - W(t_n) \sim \mathcal{N}(0, \Delta t)$ to incorporate stochastic forcing yields the *Euler-Maruyama* scheme

$$x^{n+1} = x^n + \Delta t f(x^n) + \Delta_{n+1}W. \quad (3.2)$$

For the discretised stochastic primitive equations we use the *implicit* variant of Equation (3.2): the deterministic drift is evaluated implicitly at time $n + 1$, while the noise increment $\Delta_{n+1}W$ is treated explicitly using current-time data, see (Lord et al., 2014).

The Q-Wiener expansion is truncated to $N \times N$ modes, enough to satisfy the Nyquist criterion⁶ on the chosen mesh (see (Trefethen, 2000)). For the convergence tests that vary Δt across multiple levels, the fine-grid Brownian increments are drawn once and aggregated to obtain the coarser-grid increments, so all Δt levels share the same underlying Brownian path and any difference between curves reflects the time-step refinement rather than sampling variability.

3.4 Monte Carlo estimation

The central quantity $\mathbb{E} \left[\sum_n \Delta t \|\nabla_H p_s^n\|^2 \right]$ is an expected value over the random Q-Wiener forcing. There is no closed-form expression for it: even computing one realisation requires running the FEM-discretised SDE solver from $t = 0$ to $t = T$. We therefore approximate the expectation by a sample average. Draw K_{sim} independent Brownian paths, run a full Euler-Maruyama simulation for each, and average the resulting per-sample values $X^{(k)} = \sum_n \Delta t \|\nabla_H p_s^{n,(k)}\|^2$:

$$\hat{X}_{K_{\text{sim}}} = \frac{1}{K_{\text{sim}}} \sum_{k=1}^{K_{\text{sim}}} X^{(k)} \approx \mathbb{E}[X]. \quad (3.3)$$

By the law of large numbers $\hat{X}_{K_{\text{sim}}} \rightarrow \mathbb{E}[X]$ as $K_{\text{sim}} \rightarrow \infty$, and the central limit theorem gives a 95% confidence interval $\hat{X}_{K_{\text{sim}}} \pm 1.96\sigma_X/\sqrt{K_{\text{sim}}}$ for the true expectation, where σ_X is the per-sample standard deviation. The experiments use $K_{\text{sim}} = 501$ samples per configuration and per Δt level (Table 3), a sample-size choice inspired directly by (Feng et al., 2021); the empirical adequacy of this choice is verified in Section 5.3.2.

To assess whether the splitting and no-split estimators differ by a statistically significant amount, we use the Wilcoxon signed-rank test on their paired differences (the two solvers share the same underlying Brownian path at each sample index, so the estimators are paired by construction). This non-parametric paired test makes no distributional assumption on the differences; see Section 5.3.3 for application.

3.5 Why splitting methods are introduced

A standard implicit scheme handles deterministic stiffness well, but the stochastic forcing introduces a problem implicit methods alone cannot solve. The pressure has *reduced temporal regularity* in the stochastic problem (it is less smooth in time than its deterministic counterpart) (Carelli & Prohl, 2012; Feng et al., 2021). When the standard stability argument is applied to

⁶*Nyquist criterion*: a mesh of resolution h can only meaningfully represent spatial modes of wavelength $\sim 2h$. Modes finer than this alias onto coarser ones, so there is no point in including them in the expansion.

the discrete equation, one term in the resulting bound grows like $C/\Delta t$ as the time step shrinks: a coupling between the noise increment and the pressure gradient that cannot be absorbed into a quantity we already control. Because the velocity-error bound of (Hallern & Ahlkrona, 2026) is proportional to the accumulated norm $\mathbb{E}[\sum \Delta t \|\nabla_H p_s\|^2]$, the blow-up propagates into a velocity error that gets *worse* under refinement. The splitting scheme is introduced precisely to remove this offending coupling.

3.6 The splitting scheme

The idea, introduced for stochastic Stokes in (Feng et al., 2021) and adapted to the stochastic primitive equations in (Hallern & Ahlkrona, 2026), is to perform a hydrostatic Helmholtz decomposition of the stochastic forcing at each time step,

$$B(\mathbf{u}^n)\Delta_{n+1}W = \nabla_H \xi^n + \eta^n, \quad (3.4)$$

into a gradient part $\nabla_H \xi^n$ (the scalar potential ξ^n is obtained from a Helmholtz-type problem on the surface; the construction is in (Hallern & Ahlkrona, 2026)) and a remainder η^n that is solenoidal in the primitive-equation sense:

$$\left(\int_{-D(x)}^0 \eta^n dz, \nabla_H q \right)_{\Gamma_s} = 0 \quad \forall q \in H^1(\Gamma_s). \quad (3.5)$$

The two pieces play very different roles in the energy estimate that produced the $C/\Delta t$ blow-up. The gradient part $\nabla_H \xi^n$ is folded into a *shifted pressure* $r_s^{n+1} = p_s^{n+1} - \Delta t^{-1} \xi^n$, so it never appears as a separate stochastic forcing; the discrete momentum equation is simply rewritten in terms of r_s in place of p_s . The remainder η^n , by Equation (3.5), has zero inner product against any horizontal surface gradient: when the momentum equation is tested against $\nabla_H r_s^{n+1}$, the offending cross-term identified above vanishes by construction.

The time stepper used throughout this thesis is implicit Euler-Maruyama (as introduced in Section 3.3); the splitting scheme operates on top of it, modifying only the forcing applied at each step rather than replacing the time stepper. In the *solenoidal* case the original forcing already satisfies the compatibility condition, so $\xi^n \equiv 0$ and the decomposition is trivial. The shift gives $r_s = p_s$, and the noise drives the equation as usual but no longer enters the elliptic problem that determines the pressure with a divergent contribution. The splitting therefore only does real work in the *non-solenoidal* case, where ξ^n is non-trivial and the shift to r_s is what removes the blow-up.

3.7 Coordinate transformation used to achieve varying bathymetry

A note on notation. The mathematical model in Section 2.3 is three-dimensional, with horizontal coordinates (x, y) and a vertical coordinate z . The numerical experiments are run on a 2D vertical slice, so the second axis of the slice plays the role of the vertical direction. In the remainder of Section 3 and throughout Section 5 we therefore relabel this second axis as y , matching the convention used in the MATLAB implementation and in all figure captions. In what follows x is the single horizontal coordinate of the slice and y is its vertical coordinate, with $-D(x) \leq y \leq 0$.

The varying bathymetry is implemented via a coordinate transformation applied to the reference mesh. Given a reference domain with constant depth H , the vertical coordinates of all mesh nodes are warped according to

$$y \mapsto y \cdot D(x)/H,$$

where $D(x)$ is the local depth at horizontal position x . This maps the flat reference domain onto the physical domain with spatially varying depth. The transformation is applied to both the pressure (P^1) and velocity (P^1) node sets. All standard finite-element matrices (the *stiffness matrix* for the viscous term, the *mass matrix* for the time-derivative term, and the divergence

operator for the continuity constraint) are then assembled on the warped mesh, so that the geometric effect of the bathymetry is incorporated automatically into the discrete system.

For the experiments reported in this thesis, the local depth $D(x)$ is taken to be a sinusoidal warping at 20% of the domain height corresponding to $D(x) = H(0.70 + 0.20 \sin(10\pi x/L) + 0.05 \sin(3\pi x/L))$ with H the reference depth and L the horizontal extent.

3.8 Domain geometry and aspect ratio

In the experiments, the domain dimensions are chosen as $L = 5$ and $H = 1$, giving an aspect ratio of $H/L = 1/5$. Similar aspect ratios appear in (Chacón Rebollo et al., 2012) for finite-element discretisations of the primitive equations. While real ocean basins exhibit aspect ratios on the order of 10^{-3} or smaller, the value $1/5$ is sufficient to break the symmetry between horizontal and vertical directions while keeping the mesh resolution requirements manageable. It ensures that the coordinate transformation used to implement varying bathymetry (see Section 4.3.1) introduces a genuine anisotropy in the mesh, rather than acting on a square domain where horizontal and vertical effects are artificially balanced. A robustness check against the square-domain case is reported in Appendix A.1.

Theoretical Results

This chapter is the analytical core of the thesis. The flat-bottom case treats the setting where an orthogonality identity makes the viscous term drop out, allowing the stability bound from (Hallern & Ahlkrona, 2026) to be reproduced. The varying-bathymetry case repeats the same analysis without the flat-bottom assumption: the orthogonality fails, and explicit correction terms proportional to $\nabla_H D$ appear in both the boundary integral and the volume integral. Establishing an analogous bound under varying bathymetry is beyond the scope of this thesis.

4.1 Flat-bottom reference case

We first recall the flat-bottom situation, which is the setting treated in the reference analysis. In this case the depth is constant, so the lower boundary is horizontal and several geometric terms simplify.

4.1.1 The known orthogonality

In the flat-bottom case, the depth function is constant as defined in Section 2.5. The following lemma is the key cancellation behind the flat-bottom stability proof: it states that the viscous term, integrated against any horizontal gradient, gives zero.

Lemma 4.1: Assume a flat bathymetry, i.e. $D(\vec{x}) = H$ is constant, and periodic boundary conditions in the horizontal directions. Then, for sufficiently regular velocity fields \mathbf{u}_H satisfying the assumptions of the primitive-equation formulation, the identity

$$(\Delta \mathbf{u}_H, \nabla_H r) = 0 \quad (4.1)$$

holds for all $r \in H^1(\omega) := \{r \in L^2(\omega) : \nabla_H r \in L^2(\omega)\}$.

This identity, established as Lemma 3.2 in (Hallern & Ahlkrona, 2026), is the load-bearing ingredient of the flat-bottom stability proof in Section 4.2 below: it is what makes the viscous term drop out when one tests the discrete momentum equation against $\nabla_H r$.

Proof: We begin with an integration by parts over \mathcal{D} . Since $r = r(\vec{x})$ depends only on the horizontal variables, we obtain

$$\begin{aligned} (\Delta \mathbf{u}_H, \nabla_H r) &= \int_{\mathcal{D}} \Delta \mathbf{u}_H \cdot \nabla_H r \, d(\vec{x}, z) \\ &= - \int_{\mathcal{D}} (\nabla_H \cdot \Delta \mathbf{u}_H) r \, d(\vec{x}, z) + \int_{\partial \mathcal{D}} (\Delta \mathbf{u}_H \cdot \mathbf{n}_H) r \, dS. \end{aligned} \quad (4.2)$$

Boundary term. We decompose

$$\partial \mathcal{D} = \Gamma_s \cup \Gamma_l \cup \Gamma_b,$$

where Γ_s is the surface, Γ_l the lateral boundary, and Γ_b the bottom boundary. On $\Gamma_s = \{z = 0\}$ the outward unit normal is $\mathbf{n} = (0, 0, 1)$, so its horizontal component $\mathbf{n}_H = 0$ and the surface integral vanishes. On Γ_l the contributions cancel under the imposed periodic boundary conditions. With constant depth, the bottom $\Gamma_b = \{z = -D_0\}$ is also a horizontal plane, with outward unit normal $\mathbf{n} = (0, 0, -1)$, so $\mathbf{n}_H = 0$ on Γ_b as well. All three boundary pieces vanish.

Volume term. For sufficiently regular fields, the horizontal gradient commutes formally with the Laplacian,

$$\begin{aligned}
\nabla_H \cdot \Delta \mathbf{u}_H &= \nabla_H \cdot (\Delta u_x, \Delta u_y) \\
&= \partial_x (\partial_{xx} u_x + \partial_{yy} u_x + \partial_{zz} u_x) + \partial_y (\partial_{xx} u_y + \partial_{yy} u_y + \partial_{zz} u_y) \\
&= \partial_{xx} (\partial_x u_x + \partial_y u_y) + \partial_{yy} (\partial_x u_x + \partial_y u_y) + \partial_{zz} (\partial_x u_x + \partial_y u_y) \\
&= \Delta (\nabla_H \cdot \mathbf{u}_H).
\end{aligned} \tag{4.4}$$

Using the hydrostatic incompressibility relation

$$\nabla_H \cdot \mathbf{u}_H + \partial_z w = 0,$$

we obtain $\nabla_H \cdot \mathbf{u}_H = -\partial_z w$, so

$$\int_{\mathcal{D}} (\nabla_H \cdot \Delta \mathbf{u}_H) r \, d(\vec{x}, z) = - \int_{\mathcal{D}} \Delta (\partial_z w) r \, d(\vec{x}, z).$$

With constant depth, the vertical integration commutes with horizontal derivatives, and the volume integral reduces to the depth-averaged divergence, which vanishes by the constraint Equation (2.17) together with $w|_{z=0} = w|_{z=-D_0} = 0$.

Both the boundary and volume terms vanish, yielding Equation (4.1). \blacksquare

4.2 Stability bound in the flat case

In the flat-bottom case, under a solenoidal compatibility condition on the stochastic forcing, the accumulated surface-pressure-gradient norm satisfies

$$E \left(\sum_{n=0}^{N-1} \Delta t \|\nabla_H r_s^{n+1}\|^2 \right) \leq C,$$

with C independent of Δt . The variable bounded is the *shifted pressure* $r_s^{n+1} := p_s^{n+1} - \Delta t^{-1} \xi^n$, the one the splitting actually controls. In the solenoidal cases $\xi \equiv 0$ so $r_s = p_s$ and the splitting acts trivially.

The argument is a standard energy estimate: test the discrete momentum equation against $\nabla_H r_s^{n+1}$, use Lemma 4.1 to kill the viscous term, and use the solenoidal property of η^n (Equation (3.5)) to kill the stochastic cross-term. The convective term⁷ is absent in the no-convection setting we work in. Full details are in Theorem 5.4 of (Hallern & Ahlkrona, 2026).

Under varying bathymetry the orthogonality fails (see Section 4.3 below), and this bound no longer follows in closed form.

4.3 Varying bathymetry

We now repeat the calculation underlying Lemma 4.1 without assuming constant depth. As before, r denotes a generic horizontal test function in $H^1(\omega)$; in a stability proof we would specialise to $r = r_s$, but the orthogonality identity itself holds (or fails) for any such r . The flat-bottom proof used $\nabla_H D = 0$ in two places: it forced $\mathbf{n}_H = 0$ on the bottom boundary, and it made the vertical integration in the volume term commute with horizontal derivatives. Dropping those simplifications, the integration by parts Equation (4.2) is unchanged, so

$$(\Delta \mathbf{u}_H, \nabla_H r) = - \int_{\mathcal{D}} (\nabla_H \cdot \Delta \mathbf{u}_H) r \, d(\vec{x}, z) + \int_{\partial \mathcal{D}} (\Delta \mathbf{u}_H \cdot \mathbf{n}_H) r \, dS,$$

⁷*Convective term*: the nonlinear advection $(\mathbf{u} \cdot \nabla) \mathbf{u}_H$ in the full primitive equations, which describes how the fluid carries itself along. This thesis works in a *no-convection setting*, meaning the linearised momentum equation in which this term has been dropped, as introduced in Section 1.

but the geometry of the bottom boundary and the Leibniz step in the volume term now produce surviving contributions, which we identify below.

4.3.1 Geometry of the bottom boundary

We decompose the boundary as

$$\partial\mathcal{D} = \Gamma_s \cup \Gamma_l \cup \Gamma_b,$$

where Γ_s denotes the surface, Γ_l the lateral boundary, and Γ_b the bottom boundary. The bottom is given by $z = -D(\vec{x})$. Equivalently, it is the level surface of the function

$$\Phi(\vec{x}, z) = z + D(\vec{x}),$$

so that $\Gamma_b = \{(\vec{x}, z) : \Phi(\vec{x}, z) = 0\}$. Since the gradient of a level-set function is normal to the surface,

$$\nabla\Phi = \begin{pmatrix} \nabla_H D \\ 1 \end{pmatrix}.$$

Hence the outward-pointing unit normal on Γ_b is

$$\mathbf{n} = \frac{1}{\sqrt{1 + |\nabla_H D|^2}} \begin{pmatrix} -\nabla_H D \\ -1 \end{pmatrix},$$

and its horizontal component is therefore

$$\mathbf{n}_H = -\frac{\nabla_H D}{\sqrt{1 + |\nabla_H D|^2}}.$$

Moreover, the induced surface element on Γ_b is

$$dS = \sqrt{1 + |\nabla_H D|^2} d\vec{x}.$$

4.3.2 The surviving bottom contribution

On Γ_s and Γ_l the boundary integral still vanishes by the same arguments as in the proof of Lemma 4.1 (purely vertical normal on the surface; periodic cancellation on the lateral boundary). Only the bottom contribution changes. Substituting the expressions for \mathbf{n}_H and dS from Section 4.3.1 into the surface integral over Γ_b gives

$$\int_{\Gamma_b} (\Delta\mathbf{u}_H \cdot \mathbf{n}_H) r dS = - \int_{\omega} (\Delta\mathbf{u}_H(\vec{x}, -D(\vec{x})) \cdot \nabla_H D(\vec{x})) r(\vec{x}) d\vec{x}, \quad (4.15)$$

with $\Delta\mathbf{u}_H$ evaluated on the bottom boundary $z = -D(\vec{x})$. With constant depth this contribution vanishes because $\nabla_H D = 0$; with variable depth it is generically nonzero.

4.3.3 The volume term and the loss of cancellation

By the commutator identity Equation (4.4) and the hydrostatic incompressibility relation, the volume term reduces, exactly as in the flat case, to

$$\int_{\mathcal{D}} (\nabla_H \cdot \Delta\mathbf{u}_H) r d(\vec{x}, z) = - \int_{\mathcal{D}} \Delta(\partial_z w) r d(\vec{x}, z).$$

In the flat case, the vertical integration commuted with horizontal derivatives and the right-hand side collapsed to the depth-averaged divergence, which vanishes. In the variable-bathymetry case the lower limit of the vertical integral depends on \vec{x} through $D(\vec{x})$, and applying the Leibniz rule with variable lower boundary generates additional terms involving $\nabla_H D$ and the values of the solution on the sea floor $z = -D(\vec{x})$. Hence the volume term no longer collapses, and bathymetry-dependent corrections remain.

4.3.4 Consequences of non-flat geometry

Combining the surviving bottom contribution Equation (4.15) with the surviving volume term, the varying-bathymetry analogue of Equation (4.1) reads

$$(\Delta \mathbf{u}_H, \nabla_H r) = - \int_{\mathcal{D}} (\nabla_H \cdot \Delta \mathbf{u}_H) r \, d(\vec{x}, z) - \int_{\omega} (\Delta \mathbf{u}_H(\vec{x}, -D(\vec{x})) \cdot \nabla_H D(\vec{x})) r(\vec{x}) \, d\vec{x},$$

The second term is proportional to $\nabla_H D$. The first term carries the same proportionality less obviously: applying the Leibniz rule to the vertical integration over $[-D(\vec{x}), 0]$ generates correction terms involving $\nabla_H D$, so it too fails to vanish whenever $\nabla_H D \neq 0$. The orthogonality Equation (4.1) is therefore lost, and the testing argument of Section 4.2 no longer closes the stability bound in the same way.

Deriving a closed form bound on the residual contributions involving $\nabla_H D$ is beyond the scope of this thesis. As such, we do not establish whether these surviving terms admit a uniform bound on $\mathbb{E} \left[\sum \Delta t \|\nabla_H p_s\|^2 \right]$ as $\Delta t \rightarrow 0$. The same kind of boundary contributions also arise in the proof of Theorem 5.4 in (Hallern & Ahlkrona, 2026) under the flat-bottom assumption; with varying bathymetry both arguments pick up additional $\nabla_H D$ terms that would need to be controlled. The role of the numerical experiments in Section 5 is to understand the magnitude of the resulting gap empirically, not to substitute for that analytical step.

4.4 Remarks on bathymetry and boundary terms

The effect of bathymetry enters the analysis through the geometry of the bottom boundary. For flat bathymetry, certain boundary contributions vanish automatically, whereas for varying bathymetry additional terms appear. This mechanism was derived explicitly in Section 4.3.

A second remark concerns a discrepancy between the analysis presented above and the implementation used in Section 5. While the lateral and bottom boundary conditions in Section 2.5 are stated in their physically motivated form (periodic in the horizontal directions, no-penetration with free-slip at the seabed), the code instead imposes homogeneous Dirichlet data on both the lateral and bottom boundaries, leaving only the surface unconstrained in keeping with the rigid-lid assumption.

The analytical role of the lateral boundary is purely to make the lateral integrals from integration by parts vanish: periodicity achieves this through cancellation between matching boundaries, while homogeneous Dirichlet achieves the same through vanishing test functions, so all identities used in Section 4.3 carry over unchanged. The geometric obstruction associated with $\nabla_H D$ originates from the bottom boundary integrand $\Delta \mathbf{u}_H \cdot \nabla_H D$, which depends on $\Delta \mathbf{u}_H$ rather than on \mathbf{u}_H itself, and therefore survives even when \mathbf{u}_H is forced to vanish at the seabed. Moreover, the stochastic basis $\sin(\pi j x / L), \psi_i(z)$ already vanishes at $x = 0$ and $x = L$, so the noise itself is compatible with both choices. Dirichlet data is preferred in the implementation because the *manufactured-solution* test case used to verify the deterministic solver (a verification technique in which one picks an exact solution analytically and inserts the corresponding forcing into the equations to check that the solver recovers it) prescribes nonzero boundary values that are incompatible with periodicity.

Numerical Experiments

5.1 Experimental setup

5.1.1 Parameter choices

The simulation parameters used to generate Figure 9 are listed in Table 3 and the implementation-level parameters (viscosity, FEM order, manufactured-solution test case, boundary-condition type) in Appendix A.2. The headline choices are $K_{\text{sim}} = 501$ Monte Carlo samples, $N = 15$ stochastic modes per spatial direction, mesh size $h_{\text{max}} = 0.2$, domain $H \times L = 1 \times 5$, final time $T = 5$, and ten time-step levels $\Delta t \in \{10^{-3}, \dots, 10^0\}$. The combination produces a total of order 10^8 Stokes solves; statistical uncertainty is dominated by K_{sim} , as quantified in Section 5.3.2.

5.1.2 Experiments

For orientation, the experimental work is organised into six experiments. The first two are the main numerical studies, the third provides qualitative spatial context, the fourth and fifth probe sampling and statistical reliability, and the last is a robustness check reported in the appendix.

1. **Main Δt -sweep** (Section 5.2.3). All four configurations (solenoidal/non-solenoidal \times flat/varying) under both the splitting and the no-split solver, evaluated at ten time-step sizes $\Delta t \in \{10^{-3}, \dots, 10^0\}$ with the parameters of Table 3. Results: Figure 9, Figure 10.
2. **Mesh refinement test** (Section 5.3.1). All four configurations at seven mesh sizes $h_{\text{max}} \in \{0.5, 0.35, 0.25, 0.2, 0.15, 0.1, 0.05\}$, fixed $\Delta t = 0.01$, with a summary re-run at $\Delta t = 10^{-3}$. Results: Figure 12, Table 1.
3. **Spatial-field snapshots**. Single-realisation snapshots of velocity, split surface pressure p_s , and reconstructed pressure $p_{\text{orig},s}$ at $t = 5$ and $\Delta t = 10^{-3}$, for each of the four configurations under both solver modes. Results: Figure 5, Figure 6, Figure 7, Figure 8.
4. **Statistical reliability of the Monte Carlo estimator** (Section 5.3.2, Section 5.3.3). Confidence-interval and rate diagnostics on the running estimator, and a paired Wilcoxon signed-rank test for the significance of the splitting effect. Both pieces operate on the samples produced by Experiment 1, no new simulations. Results: Figure 13, Figure 14.
5. **Monte-Carlo variance and velocity-norm diagnostics** (Section 5.2.5). ± 1 standard-deviation bands across $K_{\text{sim}} = 501$ samples and the ensemble-mean velocity norm $\mathbb{E}[\|\mathbf{u}\|_{L^2}]$, on the same parameter set as Experiment 1. Results: Figure 11, Figure 15.
6. **Robustness to the noise decay rate** (Section 5.3.4). The accumulated pressure-gradient norm recomputed under the splitting scheme for three decay exponents $\{-1/2, -1, -2\}$ on the truncated Q-Wiener eigenvalues, otherwise as in Experiment 1. Results: Table 2.

5.1.3 Algorithm

The Monte Carlo procedure used in Experiment 1 is summarised in Algorithm 1. Two design choices are worth highlighting because they directly motivate the statistical analysis in Section 5.3.2 and Section 5.3.3:

1. **Hierarchical Brownian path**. Each sample k generates a single fine-grid noise sequence which is then aggregated to all coarser time-step grids by `buildCoarseRVs`. Samples at different Δt levels within a realisation are therefore coupled by construction, whereas samples across k remain independent. The dependence across Δt is what makes inter- Δt comparisons (slopes, mesh-refinement curves) free of sampling noise.
2. **Same path for both solvers**. At each sample k the no-split and split solvers process the same underlying Brownian path; only the solver step differs. This pairing isolates the solver effect from the path effect and is what licenses the paired Wilcoxon analysis in Section 5.3.3.

Implementation-level details (parallelisation strategy, RNG seeding, and the rigour of the seeding approach) are deferred to [Appendix A.3](#).

Algorithm 1: Monte Carlo estimation of the measured quantity

```

1: ▷ Precompute once per configuration: mesh, FEM operators, noise basis, eigenvalues,
   and the Stokes-matrix factorisation at each time-step level  $\Delta t_\ell$ .
2:
3: for each sample  $k = 1, \dots, K_{\text{sim}}$  do
4:   ▷ One fine-grid Brownian path, shared across all  $\Delta t$  levels:
5:    $\zeta^{\text{fine}} \leftarrow$  draw Gaussian increments on the finest grid
6:    $\zeta_\ell \leftarrow$  sum  $\zeta^{\text{fine}}$  into coarser increments for each  $\Delta t_\ell$ 
7:
8:   for each mode  $\in \{\text{no-split, split}\}$  (same path, different solver) do
9:     for each time-step level  $\Delta t_\ell$  do
10:       $\mathbf{u}, p_s, A_k \leftarrow 0$  (reset state and accumulator)
11:      for each step  $n = 1, \dots, T/\Delta t_\ell$  do
12:        load  $F^n \leftarrow$  assemble from the increment  $\zeta_\ell$  at step  $n$ 
13:        if mode = split then
14:           $\xi^n \leftarrow$  solve a small 1D surface problem for the split correction
15:           $F^n \leftarrow F^n -$  contribution of  $\xi^n$ 
16:        end
17:         $(\mathbf{u}^{n+1}, p_s^{n+1}) \leftarrow$  one Stokes step using the precomputed factorisation
18:         $A_k \leftarrow A_k + \Delta t_\ell \|\nabla_H p_s^{n+1}\|^2$ 
19:      end
20:    end
21:  end
22: end
23:
24: ▷ After the loop, average over samples:
25:  $\hat{X}(\text{mode}, \Delta t_\ell) \leftarrow (1/K_{\text{sim}}) \sum_{k=1}^{K_{\text{sim}}} A_k(\text{mode}, \Delta t_\ell)$ 

```

5.2 Results

5.2.1 Effects of varying bathymetry

In order to get a feel of how the transformation affects the velocity and pressure in our model we may plot a spatial field showcasing the change in velocity over the horizontal plane. As a reference, we may use the flat domain.

[Figure 5](#) and [Figure 6](#) show single-realisation snapshots of the velocity field, split pressure p_s , and original pressure $p_{\text{orig},s}$ at $\Delta t = 0.001$, for solenoidal forcing on a flat and mildly varying domain respectively. Both figures use the corrected basis $\sin(2\pi jy/D(x))$ evaluated on the local depth, so that the forcing is genuinely zero-mean over the physical column.

Two pressure quantities are reported because the splitting scheme actually solves for the shifted pressure (the variable r_s from [Section 3.6](#), plotted in the figures as p_s for consistency with the no-split solver’s variable), while the original pressure $p_{\text{orig},s} = p_s + \Delta t^{-1}\xi$ is recovered by inverting the shift. In the no-split solver no shift is applied, so $\xi \equiv 0$ and the two coincide bit-for-bit; this is why panel (c) duplicates panel (b) in the no-split row of every figure. In

the splitting scheme they generally differ whenever ξ is nonzero: the stability bound proved in Section 4.2 controls p_s (which is what the analysis was designed for), but $p_{\text{orig},s}$ is the pressure the original (un-split) momentum equation would produce. Showing both is what makes the splitting's trade-off visible, since the scheme can shrink one of the two while inflating the other.

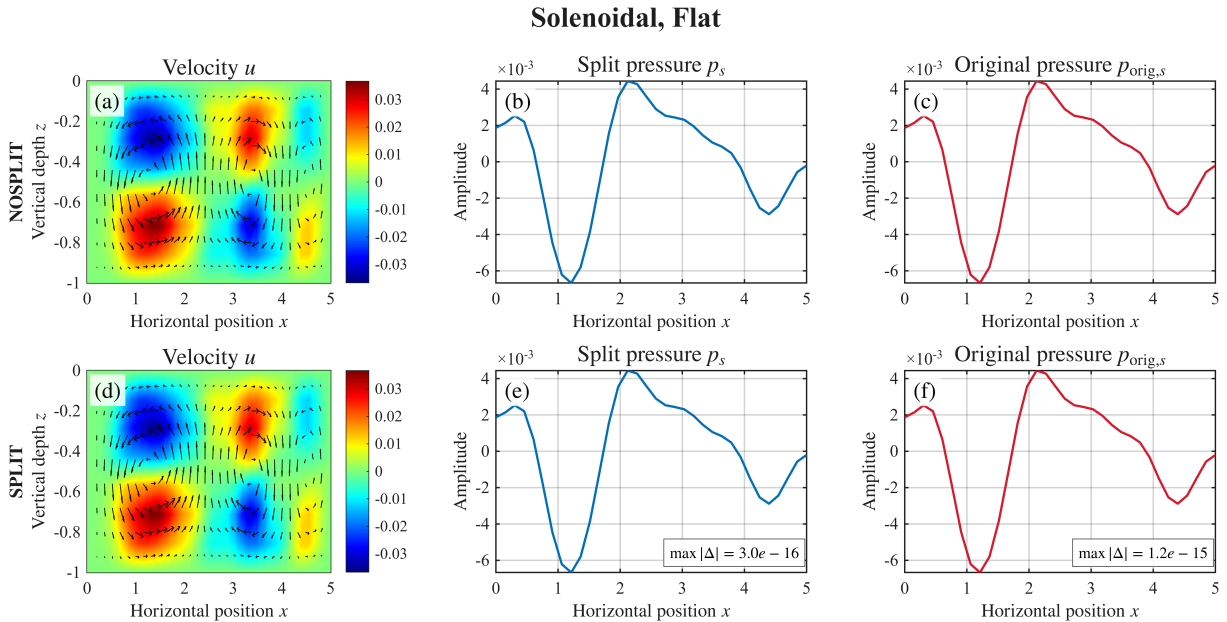


Figure 5: Spatial flow fields at $t = 5.0$ (solenoidal forcing, flat domain). **Cols:** velocity magnitude u , split pressure p_s , original pressure $p_{\text{orig},s}$. **Rows:** no-split solver (top), splitting scheme (bottom). Both rows driven by the same Q -Wiener realisation.

On the flat domain the no-split and splitting rows are visually indistinguishable. The original pressure is smooth and small (amplitude $\sim 5 \times 10^{-3}$), consistent with the theoretical bound Equation (1.5), and the splitting correction ξ is zero to floating-point precision as can be seen in the bottom right corners. Implementing the varying bathymetry we see the following:

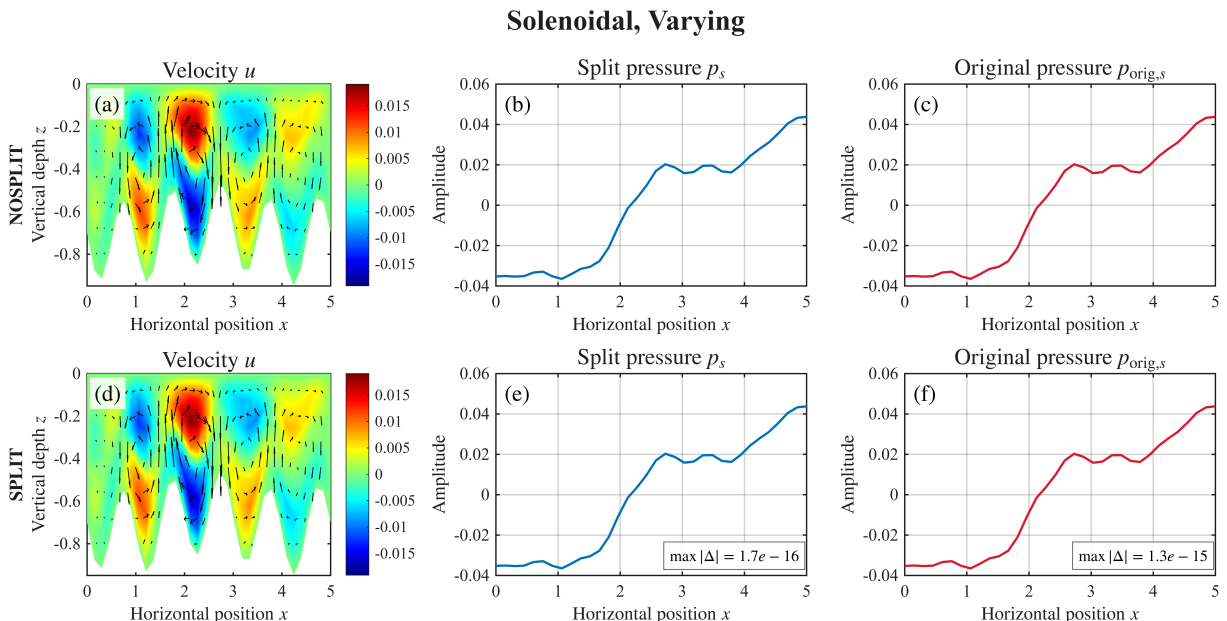


Figure 6: Spatial flow fields at $t = 5.0$ (solenoidal forcing, varying bathymetry). Layout as in Figure 5. With the corrected basis the splitting reduces to the no-split solver and the two rows coincide to floating-point precision.

Once the basis is reformulated on the local depth $D(x)$, the forcing is again genuinely solenoidal on the physical column and the splitting correction ξ vanishes. The two rows of Figure 6 therefore agree to floating-point precision. The bathymetric warping is visible in the velocity

panel as a notched lower boundary, but the surface pressure profile is smooth and bounded (amplitude $\sim 4 \times 10^{-2}$), about an order of magnitude larger than the flat case but still several orders below the non-solenoidal configurations shown next.

5.2.2 Effects of solenoidal versus non-solenoidal

We have now seen the general effect that varying bathymetry has on the variables of interest. As mentioned earlier, we incorporate the solenoidal property in order to attempt to eliminate some of these extra terms that arise from not having flat bathymetry. Let us now look at the effects of this solenoidal property versus the lack of it.

Unlike the accumulated norm $\mathbb{E}[\sum \Delta t \|\nabla_H p_s\|^2]$ used elsewhere as the stability diagnostic, the spatial figures in this section report the pressure p_s itself rather than its horizontal gradient. The reason is purely a presentational one: at the single-realisation, single-snapshot level the pressure profile $p_s(x)$ is what is directly plotted in the panels, and its amplitude is easier to read off and compare across configurations than that of the gradient. The accumulated-gradient norm remains the quantity of interest for stability, for which see Section 5.2.3, but for the qualitative comparison made here the pressure itself is the more transparent measure.

We report two amplitude measures of the surface pressure along x : the **peak** amplitude $\max_x |p_s(x)|$ and the **root-mean-square** amplitude

$$\text{RMS}(p_s) = \sqrt{\frac{1}{N_s} \sum_{i=1}^{N_s} p_s(x_i)^2},$$

where the sum runs over the N_s surface nodes. Peak captures the worst-case excursion and is sensitive to localised spikes; RMS captures the typical magnitude across the whole surface and is insensitive to single outliers. Reporting both lets the reader distinguish a profile that is globally inflated (peak and RMS scale together) from one whose peak is driven by a localised feature (peak much larger than RMS).

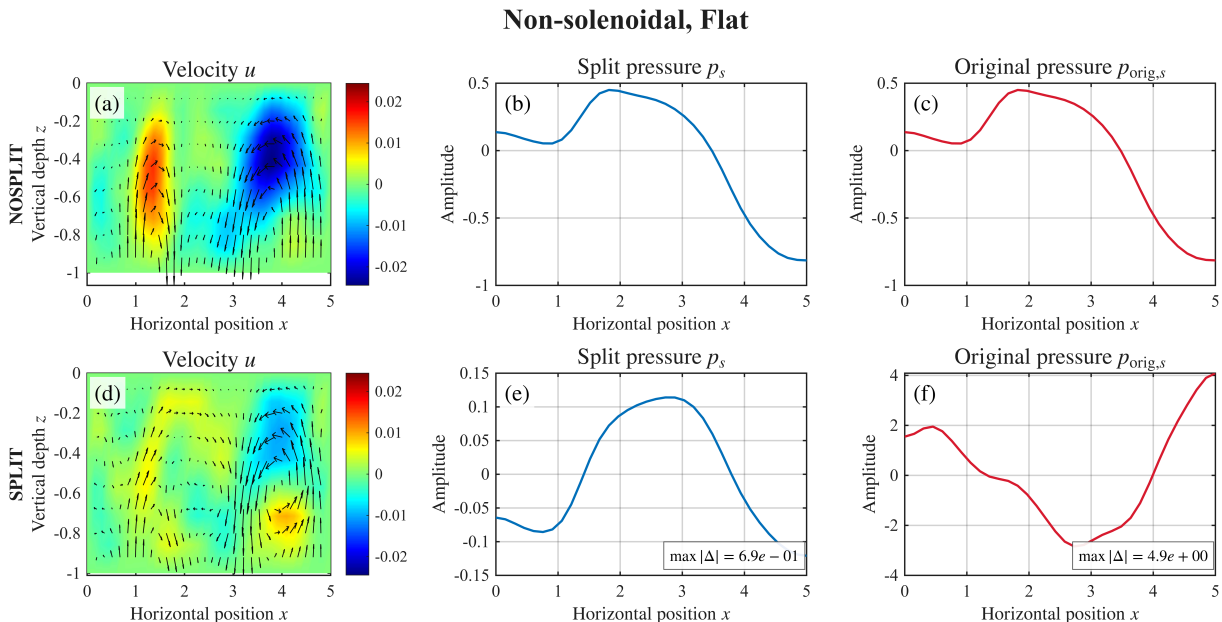


Figure 7: Spatial flow fields at $t = 5.0$ (non-solenoidal forcing, flat domain). Layout as in Figure 5.

The two rows of Figure 7 now differ markedly. The splitting scheme reduces the split pressure p_s from a peak amplitude $\sim \pm 0.8$ (RMS ≈ 0.4) in the no-split row down to $\sim \pm 0.1$, but the **original** pressure $p_{\text{orig},s}$ that the splitting reconstructs grows from the same no-split profile up to a peak range of roughly $[-2.8, 4.1]$ (RMS ≈ 2.0) under the split, a roughly five-fold inflation in both peak and RMS. The velocity field is also visibly attenuated under the split (note the order-of-magnitude smaller colour scale on the bottom-left panel). This is the signature of a

non-zero ξ correction: the split controls the p_s degree of freedom it was designed for, but pays for it in the reconstructed pressure.

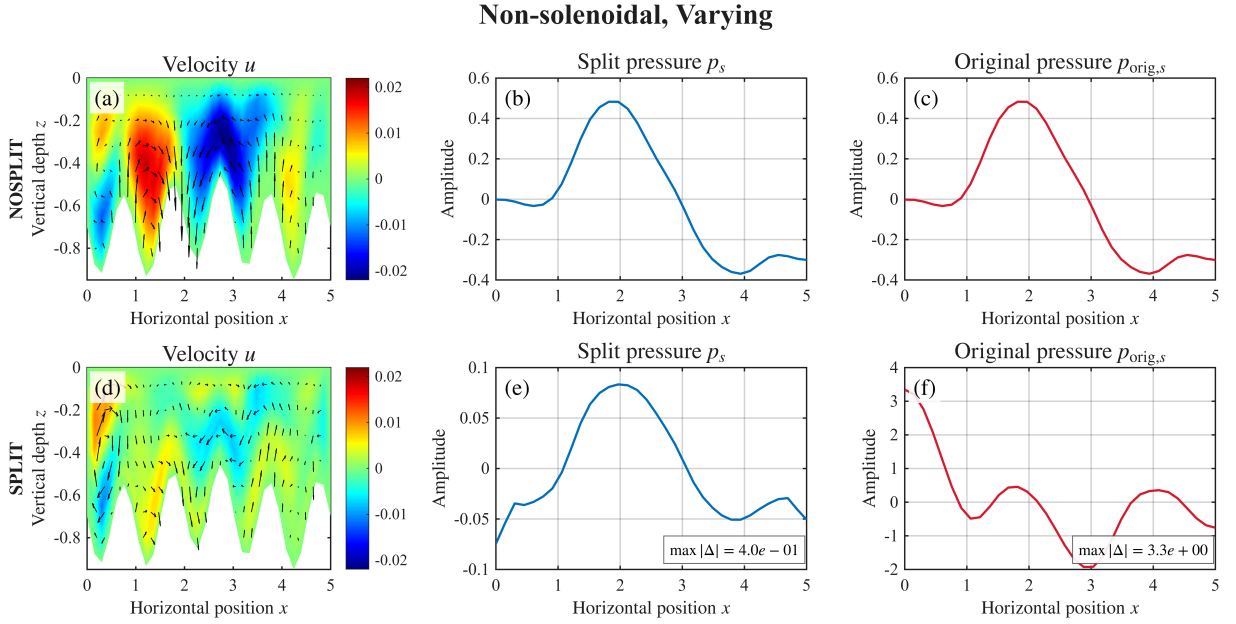


Figure 8: Spatial flow fields at $t = 5.0$ (non-solenoidal forcing, varying bathymetry). Layout as in Figure 5.

Adding varying bathymetry to the non-solenoidal case amplifies the same pattern. The no-split solver produces a bounded multi-humped profile of amplitude $\sim \pm 0.5$, whereas the reconstructed pressure $p_{\text{orig},s}$ under the splitting scheme acquires a structurally different profile: a peak of $\approx +3.4$ at the left wall, a trough of ≈ -1.9 near $x \approx 2.9$, and a partial recovery to ≈ -0.8 at the right wall. The p_s panel shrinks (as in the flat case) but the reconstruction $p_{\text{orig},s} = p_s + \xi$ inherits the large ξ contribution. The qualitative picture this paints, already visible before any convergence plots, is that the splitting scheme is not merely producing a noisier estimate of the same field; it is producing a different field altogether.

5.2.3 Accumulated pressure-gradient norm

A complete overview of our findings using parameters defined in Table 3 can be visualized in the following figure:

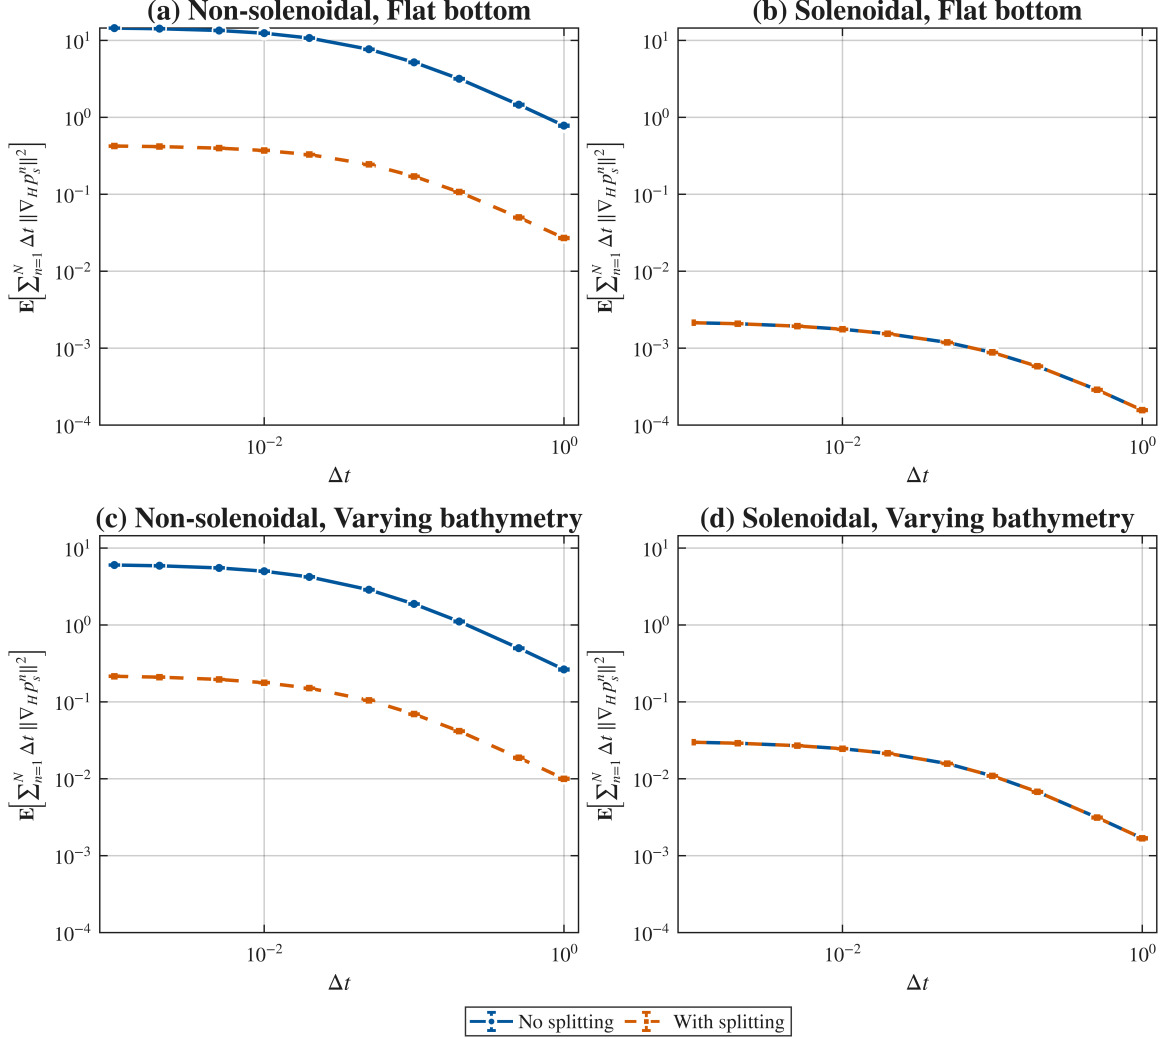


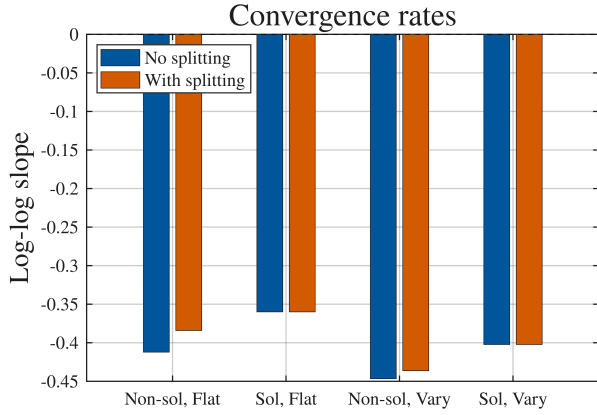
Figure 9: Accumulated surface pressure gradient norm $\mathbb{E}\left[\sum \Delta t \|\nabla_H p_s\|^2\right]$ for all four configurations: non-solenoidal flat, non-solenoidal varying, solenoidal flat, solenoidal varying. Error bars show one Monte-Carlo standard error of the mean, $\text{SE}(\Delta t) = \text{std}_{\text{samples}} / \sqrt{K_{\text{sim}}}$, and lie within the marker size at this scale.

Figure 9 tells us that all four configurations separate cleanly into two “regimes” by amplitude and not by rate. The two solenoidal configurations sit several orders of magnitude below the two non-solenoidal configurations. None of the four configurations exhibit $\Delta t \rightarrow 0$ divergence.

Figure 9 compares the splitting scheme against the no-split solver on a flat domain for both forcing types. In the non-solenoidal case, the splitting scheme produces a curve about thirty times smaller in amplitude than the no-split solver. Both curves remain bounded and slowly increase as $\Delta t \rightarrow 0$; fitted log-log slopes are -0.41 (no split) and -0.38 (split), so the splitting reduces the norm uniformly in Δt rather than altering its scaling. In the solenoidal case both curves coincide to floating-point precision as can be seen later in the following section. This confirms that the splitting correction ξ vanishes when the forcing is genuinely zero-mean on the integration interval.

5.2.4 Δt -scaling and fitted slopes

The amplitude hierarchy seen in Figure 9 can be quantified more precisely by fitting log-log slopes for each configuration. The velocity field also remains bounded across all configurations on the tested range, as seen in Figure 15.



Config	No-split	Split
Non-sol Flat	-0.412	-0.384
Sol Flat	-0.360	-0.360
Non-sol Vary	-0.447	-0.436
Sol Vary	-0.402	-0.402

Figure 10: Fitted log-log slopes of the accumulated pressure gradient norm versus Δt , across all four configurations and both solvers. A more negative slope indicates a stronger Δt -dependence on the tested range.

Figure 10 quantifies the Δt -scaling of the accumulated pressure gradient norm. All eight fitted slopes (four configurations \times two splitting modes) are negative and fall in the narrow range $[-0.45, -0.36]$, far milder than the theoretical Δt^{-1} worst case for the non-solenoidal bound of Equation (1.5). The configurations therefore do not separate by **rate**; they separate by **amplitude**. The splitting scheme leaves the slope essentially unchanged in every configuration, it does not cure the Δt -scaling, only shifts the intercept.

5.2.5 Variance behavior

The expected-value diagnostic $\mathbb{E}[\sum \Delta t \|\nabla_{Hp_s}\|^2]$ summarises the average stability behavior across realisations, but the spread across Monte Carlo samples carries additional stochastic information that the mean alone does not reveal. Figure 11 shows the same accumulated norm together with ± 1 standard-deviation bands across $K_{\text{sim}} = 501$ samples, separately on the flat and varying-bathymetry domains.

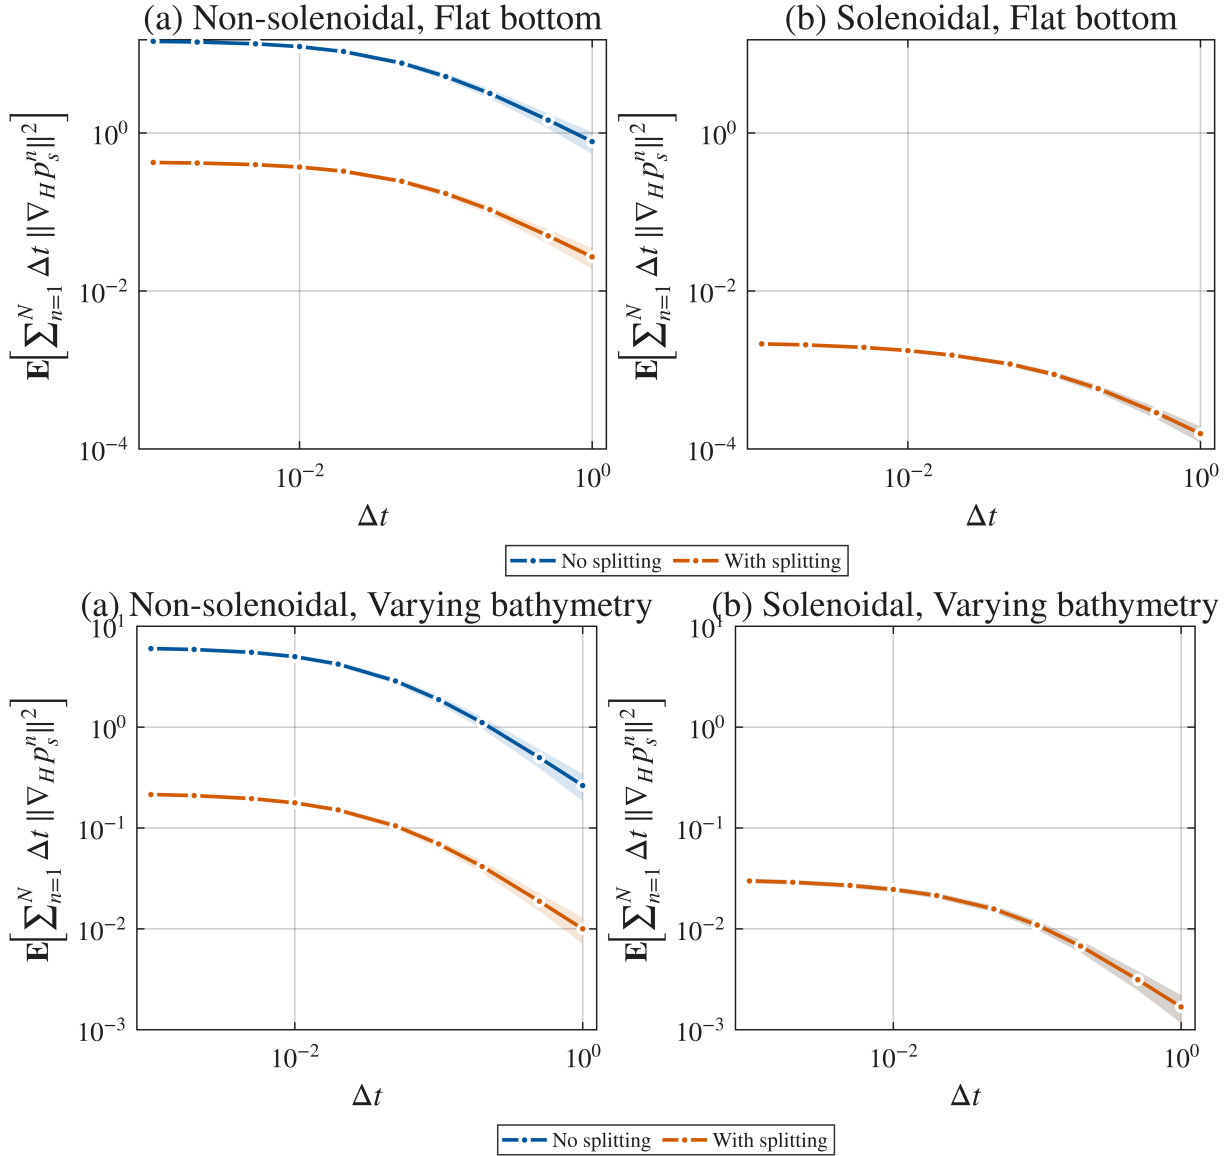


Figure 11: Accumulated pressure gradient norm with ± 1 standard deviation bands ($K_{\text{sim}} = 501$). Top: flat domain. Bottom: varying bathymetry.

Two stochastic-side observations follow. First, on the flat domain the non-solenoidal no-split bands widen substantially as $\Delta t \rightarrow 0$, signalling a stochastic instability that the mean alone does not fully expose: even when the average stays bounded, the variability of individual sample paths is itself growing. The splitting scheme keeps the bands tight in the same regime, so the splitting controls not just the average but also the spread of the pressure response. Second, on varying bathymetry the bands are wider for every configuration, including the solenoidal cases in which $\xi \equiv 0$ and split equals no-split bit-for-bit. The elevated variance therefore cannot be attributed to a basis-level loss of solenoidality; it reflects the same $\nabla_H D$ contributions to the momentum equation identified in Section 4.3. The geometric obstruction shows up at the variance level as well as at the mean level.

5.3 Robustness of the observed effects

5.3.1 Mesh refinement test: numerical versus formulation error

A diagnostic preliminary to the remaining results. We re-run all four configurations at seven mesh sizes $h_{\text{max}} \in \{0.5, 0.35, 0.25, 0.2, 0.15, 0.1, 0.05\}$, fixing $\Delta t = 0.01$ and the parameters of Table 3, in order to settle whether the solenoidal + varying-bathymetry gap of Figure 9 is a numerical artefact (in which case it should vanish under refinement) or a structural feature of

the formulation (in which case it should persist). Figure 12 reports the L^2 norm of the divergence diagnostic $\nabla_H \int_{-D}^0 \mathbf{u}_H dz$ versus h_{\max} in the left panel, and the accumulated pressure-gradient norm $\mathbb{E}[\sum \Delta t \|\nabla_H p_s\|^2]$ in the right panel.

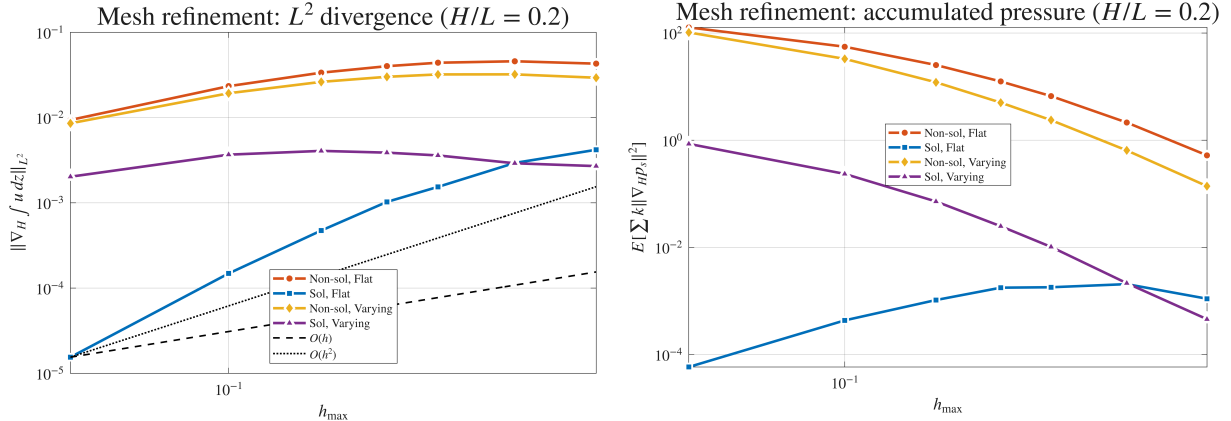


Figure 12: Mesh refinement diagnostic. Left: L^2 norm of $\nabla_H \int_{-D}^0 \mathbf{u}_H dz$ at $t = T$ versus h_{\max} . Right: accumulated pressure-gradient norm $\mathbb{E}[\sum \Delta t \|\nabla_H p_s\|^2]$ versus h_{\max} , all four configurations.

The solenoidal + flat case converges at approximately $O(h^2)$ to zero, confirming that its nonzero divergence is purely a discretisation artefact. The solenoidal + varying-bathymetry case plateaus at $\approx 10^{-3}$, an order of magnitude below the non-solenoidal cases but bounded away from zero. Since the implementation uses the basis $\sin(2\pi jy/D(x))$ built on the local column depth (so the zero-mean property holds pointwise in x), this residual reflects the $\nabla_H D$ corrections of Section 4.3, not a basis-level loss of solenoidality. The right panel of Figure 12 confirms the same picture for $\mathbb{E}[\sum \Delta t \|\nabla_H p_s\|^2]$: the order-of-magnitude gap between solenoidal + flat and solenoidal + varying is preserved at every h_{\max} , with split and no-split agreeing to relative differences of order 10^{-13} in both solenoidal configurations.

Configuration	Mode	$e(h = 0.05)$	\bar{p}
Non-sol, Flat	No split	9.37×10^{-3}	0.53
	Split	1.26×10^{-3}	1.03
Sol, Flat	-	1.55×10^{-5}	2.26
Non-sol, Varying	No split	8.56×10^{-3}	0.41
	Split	2.55×10^{-3}	0.48
Sol, Varying	-	2.02×10^{-3}	-0.04

Table 1: Convergence behavior of the divergence diagnostic. $e(h = 0.05)$ is the L^2 norm of $\nabla_H \int_{-D}^0 \mathbf{u}_H dz$ at the finest mesh; \bar{p} is the mean convergence order across the seven mesh sizes (a value near 2 indicates $O(h^2)$ convergence to zero, near 0 indicates a plateau). The solenoidal cases are split-invariant ($\xi \equiv 0$, so split and no-split coincide bit-for-bit) and listed once.

5.3.2 Monte Carlo reliability

Every point plotted in Figure 9 is the Monte Carlo estimator $\hat{X}_{K_{\text{sim}}}$ from Equation (3.3), computed at $K_{\text{sim}} = 501$ samples. As a finite average, this estimator is itself a random variable: a different draw of 501 Brownian paths would yield a slightly different value. To verify that $K_{\text{sim}} = 501$ is large enough for the inter-configuration differences in Figure 9 to be physically meaningful rather than artefacts of sampling, Figure 13 reports two convergence diagnostics at the finest tested $\Delta t = 10^{-3}$.

Monte Carlo convergence ($K_{\text{sim}} = 501$, $N = 15$, $h = 0.20$)

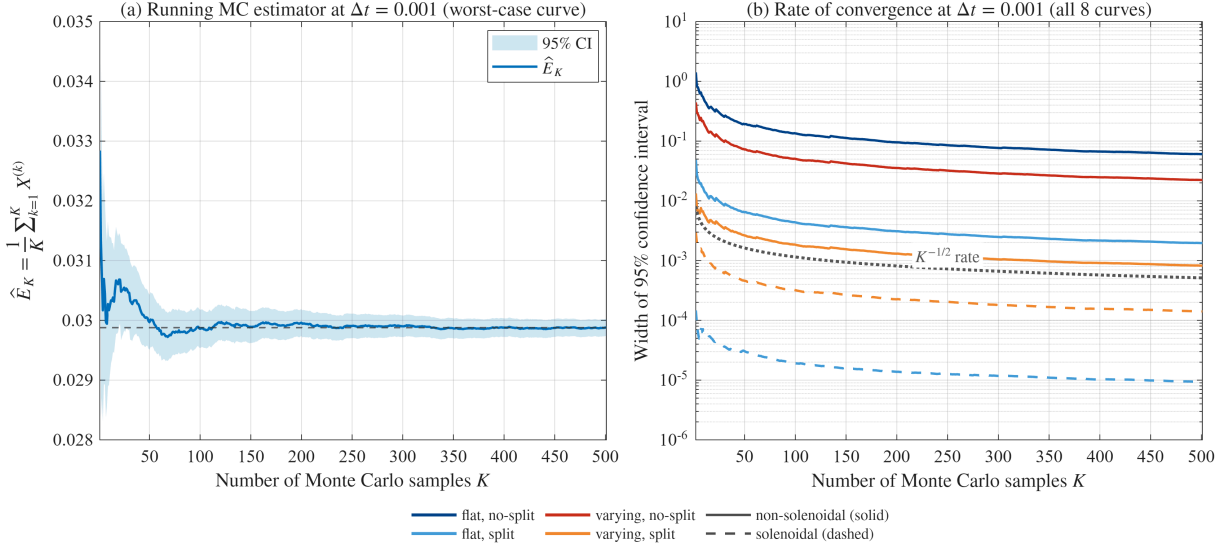


Figure 13: Monte Carlo reliability diagnostic at $\Delta t = 10^{-3}$. (a) Running estimator with 95% confidence band for the worst-case curve. (b) Width of the 95% confidence interval versus K for all eight (config, split) combinations, with the theoretical $K^{-1/2}$ rate as a dotted reference.

Panel (a) plots the running mean as the sample count K grows from 1 to $K_{\text{sim}} = 501$, with the surrounding 95% confidence band $\hat{X}_K \pm 1.96\hat{\sigma}_K/\sqrt{K}$. The running mean stabilises by $K \approx 100$ and the band has closed to a thin sliver by $K_{\text{sim}} = 501$. Only one curve is shown: the configuration with the largest relative confidence-interval width at this Δt . Every other curve converges at least as fast.

Panel (b) plots the width of that confidence band for all eight curves. The dotted reference is the $K^{-1/2}$ rate predicted by standard Monte Carlo theory; every empirical curve runs parallel to it, so the samples behave as expected.

The quantity that matters is the confidence-interval width **relative to the value being estimated**, since the eight curves span several orders of magnitude. Across all eight curves at $\Delta t = 10^{-3}$, the largest 95% confidence interval extends at most $\sim 0.47\%$ around the reported value. Across all eight curves and all ten Δt levels, the largest is $\sim 2.70\%$. The vertical gaps between curves in Figure 9 are factors of 30 to 1000, so a 2.7% envelope cannot bridge them. The visible separations therefore reflect real differences in the underlying expectations $\mathbb{E}[X]$, not which 501 Brownian paths were drawn.

5.3.3 Statistical significance of the splitting effect

Figure 9 shows the splitting scheme reducing the accumulated norm by roughly a factor of 30 in the non-solenoidal configurations and having no visible effect in the solenoidal ones. To verify that the apparent reduction is statistically distinguishable from sampling noise rather than an artefact of comparing two random averages, we apply a paired hypothesis test on the per-sample values.

By construction, the no-split and split simulations at the same sample index k use the same underlying Brownian path (built hierarchically by `buildCoarseRVs` in the suite driver), so the two estimators are **paired by design**: only the choice of solver differs between $X_{\text{no-split}}^{(k)}$ and $X_{\text{split}}^{(k)}$ at fixed k . The pairing isolates the solver effect from the path effect and yields substantially more statistical power than an unpaired comparison.

For each of the 40 (config, Δt) cells we form the paired differences

$$D^{(k)} = X_{\text{no-split}}^{(k)} - X_{\text{split}}^{(k)}$$

and apply the **Wilcoxon signed-rank test**. The Wilcoxon test is the standard non-parametric counterpart to the paired t -test: it replaces the values $D^{(k)}$ by the ranks of $|D^{(k)}|$ together with

their original signs, then asks whether positive and negative ranks are imbalanced. Working with ranks rather than values makes the test statistic’s null distribution depend only on the sample size K_{sim} and not on the shape of the distribution of D , so no Gaussian assumption is required. This is desirable here because, although most of the 80 (config, split, Δt) cells are close to Gaussian (median sample skewness 0.27), a minority of cells show mild right-skew with sample skewness up to 0.8, and the differences $D^{(k)}$ inherit some of that asymmetry. The Wilcoxon test is therefore a defensive option that remains valid across the full grid without relying on a normality assumption. Figure 14 shows the distribution of $D^{(k)}$ at the finest $\Delta t = 10^{-3}$ for each of the four configurations.

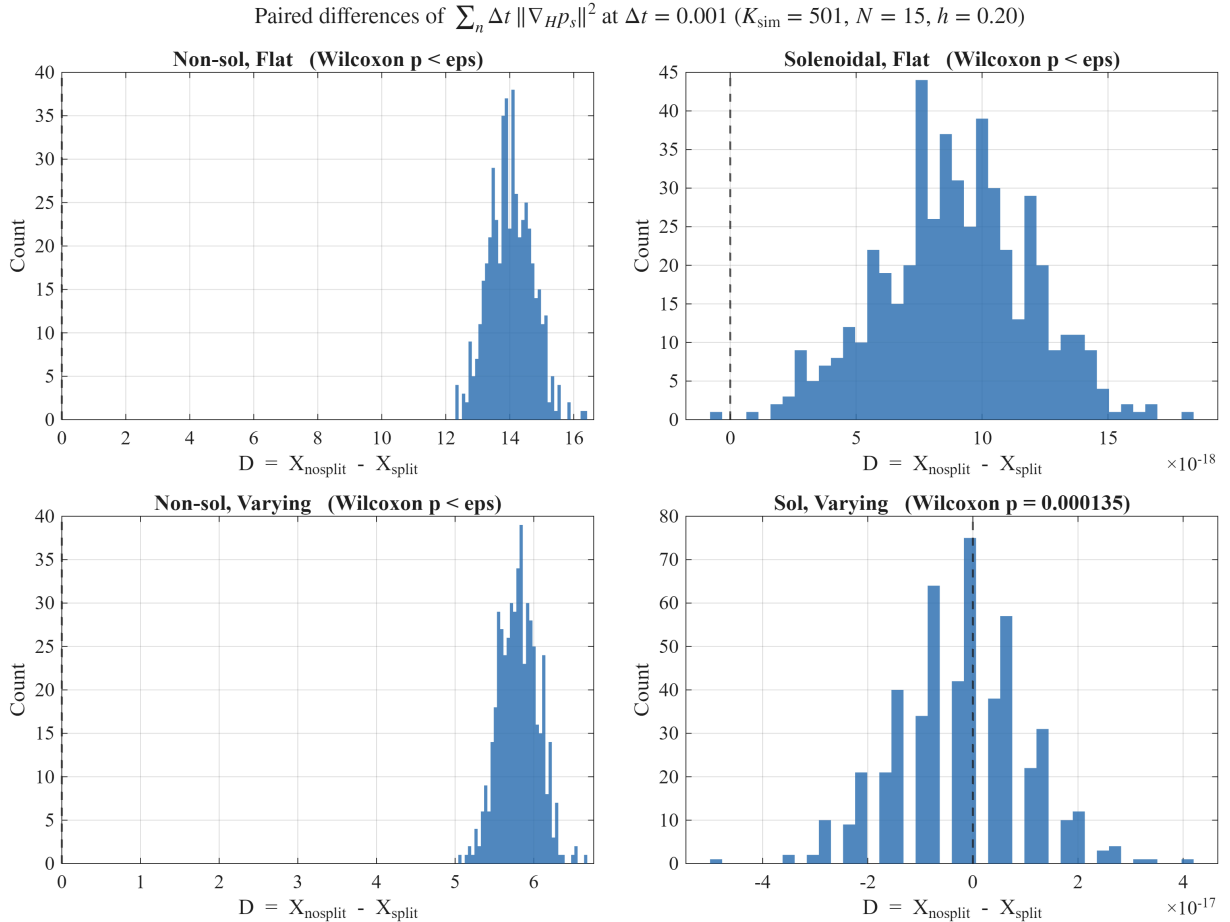


Figure 14: Distribution of paired differences $D^{(k)} = X_{\text{no-split}}^{(k)} - X_{\text{split}}^{(k)}$ at $\Delta t = 10^{-3}$, one panel per configuration. Non-solenoidal panels (left column) show large, consistently positive differences (note the x-axis scale); solenoidal panels (right column) show floating-point-precision differences around zero (note the 10^{-17} and 10^{-18} scales). Wilcoxon signed-rank p-values are reported in each subplot title.

In the two non-solenoidal configurations (top-left and bottom-left panels) the distributions are concentrated far from zero: every one of the 501 paired draws produces a positive $D^{(k)}$, on the order of 5 to 15. The Wilcoxon test rejects H_0 at p below machine epsilon at every one of the ten Δt levels, and the effect-size ratio $|(X)_{\text{no-split}}|/|(X)_{\text{split}}$ stays in the range 26.4 to 34.1. The factor-of-thirty reduction visible in Figure 9 is therefore not a graphical illusion: it is a statistically distinguishable effect at extreme significance.

In the two solenoidal configurations (top-right and bottom-right panels) the differences $D^{(k)}$ are at the level of floating-point precision (order 10^{-17} to 10^{-18}). For Solenoidal Flat every $D^{(k)}$ happens to share the same sign, which causes the Wilcoxon test to formally reject H_0 ; for Sol-Varying the differences scatter around zero and the test gives moderate p-values. In both cases the effect-size ratio is 1.000 to the precision reported. The solenoidal configurations therefore exhibit no physically meaningful splitting effect, consistent with the analytical observation that $\xi \equiv 0$ when the depth-integrated forcing vanishes; the formal rejection in the Sol-Flat panel is

a numerical artefact arising from the test's extreme power at $K_{\text{sim}} = 501$, not a real difference between the schemes.

5.3.4 Robustness to the noise decay rate

The amplitude coefficients σ_{ij} in Equation (2.7) act as a low-pass filter on the Q-Wiener forcing: low-frequency modes (small i, j) get the largest weights and the weights drop off as i, j grow, so the high-wavenumber spikes that would otherwise dominate the realised noise field are suppressed. The faster this drop-off, the smoother the forcing looks in space; the slower it is, the more high-frequency roughness survives, and the spikier each realisation of dW becomes. The drop-off rate is therefore the dial that controls how spatially smooth the noise driving the momentum equation is.

To check that the conclusions of the preceding sections do not hinge on the particular drop-off rate used everywhere else, the splitting experiment was rerun with one rougher and one smoother choice of σ_{ij} , shown as the outer two columns of Table 2; the middle column is the baseline used in every other experiment. All three settings produce a bounded accumulated pressure-gradient norm as $\Delta t \rightarrow 0$, separated only by a vertical shift that tracks the total noise variance $\sum \sigma_{ij}^2$, shown in the bottom row. The splitting mechanism is therefore not tied to one particular noise smoothness.

Δt	$\sigma_{ij} \sim (i^2/H^2 + j^2/L^2)^{-1/2}$	$\sigma_{ij} \sim (i^2/H^2 + j^2/L^2)^{-1}$	$\sigma_{ij} \sim (i^2/H^2 + j^2/L^2)^{-2}$
10^{-4}	3.78×10^{-1}	1.38×10^{-2}	1×10^{-4}
5×10^{-4}	3.72×10^{-1}	1.36×10^{-2}	1×10^{-4}
10^{-3}	3.66×10^{-1}	1.34×10^{-2}	1×10^{-4}
5×10^{-3}	3.29×10^{-1}	1.21×10^{-2}	1×10^{-4}
10^{-2}	2.99×10^{-1}	1.10×10^{-2}	1×10^{-4}
5×10^{-2}	1.87×10^{-1}	7.0×10^{-3}	$< 10^{-4}$
10^{-1}	1.31×10^{-1}	4.9×10^{-3}	$< 10^{-4}$
$\sum \sigma_{ij}^2$	1.26	4.11×10^{-2}	2.08×10^{-4}

Table 2: Accumulated pressure gradient norm $\mathbb{E}[\sum_n \Delta t \|\nabla_H p_s^n\|^2]$ under the splitting scheme for three different amplitude scalings σ_{ij} . The middle column corresponds to Equation (2.7). All columns remain bounded as $\Delta t \rightarrow 0$.

Discussion and Conclusion

This thesis studied the numerical surface pressure in the stochastic primitive equations across four configurations (solenoidal versus non-solenoidal forcing crossed with flat versus varying bathymetry), combining analytical derivations with numerical experiments. The splitting scheme of (Hallern & Ahlkrona, 2026) served as the analytical reference for the flat-bottom case and as one of the two solvers compared throughout.

The three research questions of Section 1 are addressed as follows. **(Q1)** The flat-bottom bound Equation (1.5) is reproduced empirically when the forcing is solenoidal in the primitive-equation sense, i.e. when $\nabla_H \cdot \int_{-D(\bar{x})}^0 B(\mathbf{u}_H) dW(t) dz = 0$; for non-solenoidal forcing the splitting scheme rewrites the equations into the solenoidal form, so the bound is recovered. **(Q2)** On varying bathymetry the orthogonality property underpinning the flat-bottom estimate is lost: a boundary integral proportional to $\nabla_H D$ and Leibniz-rule volume corrections appear, and an analogous closed-form bound is left open. **(Q3)** The accumulated norm $\mathbb{E} \left[\sum \Delta t \|\nabla_H p_s\|^2 \right]$ remains bounded across all four configurations on the tested range $\Delta t \in [10^{-3}, 1]$; no Δt^{-1} divergence is observed.

The central findings are four.

Type of randomness dominates. The two solenoidal configurations sit several orders of magnitude below the two non-solenoidal ones across the tested range $\Delta t \in [10^{-3}, 10^0]$, on both bathymetries. None of the four shows a Δt^{-1} -style divergence; the curves in Figure 9 retain a mild Δt dependence but stay bounded as $\Delta t \rightarrow 0$. Since the velocity-error estimate of (Hallern & Ahlkrona, 2026) is proportional to this accumulated norm, the observed boundedness directly bounds the velocity error.

The solenoidal varying-bathymetry gap is geometric, not numerical. The accumulated norm in solenoidal + varying is roughly ten times larger than in solenoidal + flat at $h_{\max} = 0.2$, and the divergence diagnostic $\nabla_H \int_{-D}^0 \mathbf{u}_H dy$ plateaus at $\approx 10^{-3}$ rather than vanishing under mesh refinement. Since ξ vanishes to machine precision in both solenoidal configurations, split and no-split agree bit-for-bit; the gap is a property of the discrete momentum equation on the warped column. The candidate mechanism is the $\nabla_H D$ boundary integral and Leibniz volume corrections of Section 4.3. The gap appears at the **amplitude** level, not in the Δt -scaling (fitted slopes -0.402 and -0.360 , essentially parallel), pointing to a genuine bathymetry-induced contribution rather than a latent Δt^{-1} -style growth.

The splitting scheme remains useful beyond its proven regime. On the non-solenoidal cases it reduces the accumulated norm by roughly a factor of 30 relative to no-split, with comparable improvement on flat and varying bathymetry, even though the orthogonality lemma Lemma 4.1 is lost as soon as $\nabla_H D \neq 0$. The $\nabla_H D$ correction terms identified in Section 4.3 appear to contribute little in absolute size in the non-solenoidal regime: the dominant non-solenoidal cross-term swamps the geometric extras.

Variance follows the same pattern. On the flat domain the non-solenoidal no-split bands of Figure 11 widen substantially as $\Delta t \rightarrow 0$; the splitting controls this widening just as it controls the mean. On varying bathymetry the bands are elevated for every configuration including the solenoidal cases where $\xi \equiv 0$, so the same $\nabla_H D$ contributions that lift the mean also lift the variance.

6.1 Final remarks

One final point about Equation (1.5). Theory predicts qualitatively different behaviour in the two regimes: a uniform-in- Δt bound C in the solenoidal case, versus a worst-case $C\Delta t + C\Delta t^{-1}$ growth (dominated by Δt^{-1} as $\Delta t \rightarrow 0$) in the non-solenoidal case. The empirical curves tell

a more uniform story: all eight fitted log-log slopes fall in the narrow range $[-0.45, -0.36]$, far milder than the Δt^{-1} rate the non-solenoidal bound allows, and only the prefactor C differs between regimes. The solenoidal cases sit several orders of magnitude below the non-solenoidal ones but exhibit qualitatively the **same** mild Δt -dependence, not the uniform-in- Δt behaviour predicted by the solenoidal theory. This discrepancy is not resolved here; plausible explanations include the bounds not being tight in the discrete FEM setting, the asymptote sitting below $\Delta t = 10^{-3}$, or some component of the discretisation introducing a residual Δt -dependence the continuum theory does not capture. Identifying the mechanism is a useful direction for further analytical work.

On bathymetry realism. The 20% sinusoidal profile used here was chosen for diagnostic clarity, not realism. The structural failure $\int_{-D(x)}^0 \sin(2\pi jy/H) dy \neq 0$ whenever $D(x) \neq H$ activates as soon as the local depth deviates from the reference, so these results should be read as a lower bound: realistic bathymetry (shelf breaks, ridges, multi-scale roughness) would produce a **more** severe loss of the zero-mean property, not a milder one. Mesh skewness, boundary-layer separation, and sub-grid topographic effects associated with steep gradients are separate failure modes not addressed here.

Future directions. A rigorous analytical bound for the varying-bathymetry case, controlling the $\nabla_H D$ corrections of Section 4.3, would close the analytical gap. Natural extensions include reinstating the convection term, allowing multiplicative noise $B(\mathbf{u})$, and replicating the experiments at finer mesh resolution and over a wider range of Δt to see whether the predicted Δt^{-1} asymptote eventually appears.

Appendix

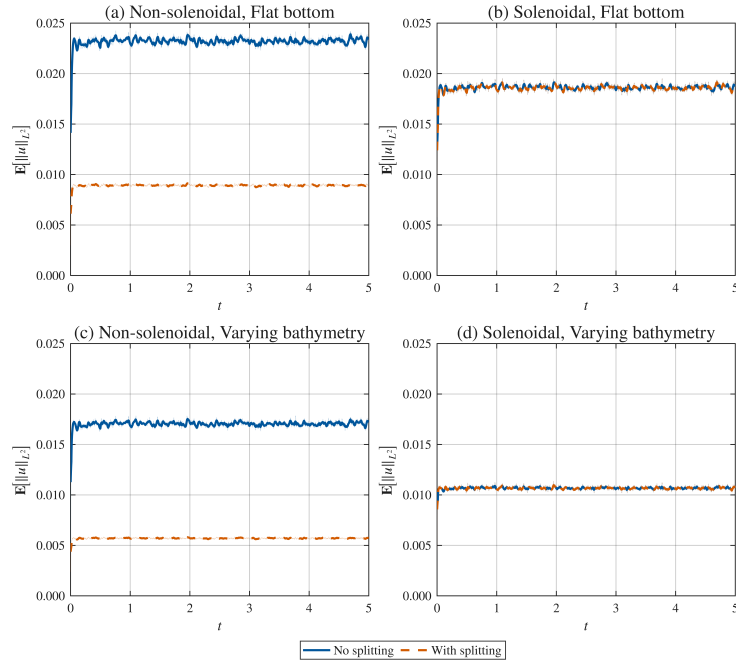


Figure 15: Expected L^2 velocity norm $\mathbb{E}[\|\mathbf{u}\|_{L^2}]$ at $\Delta t = 0.001$. The no-split solver and splitting scheme are shown with a moving-average smoothing applied to the ensemble mean; light-coloured lines show the raw signal.

A.1 Robustness to domain aspect ratio

We verified that the qualitative behavior of all four experimental configurations is preserved under the change from a square domain ($H = L$) to the anisotropic one ($H/L = 1/5$). The convergence rates, the separation between split and non-split schemes, and the relative ordering of the configurations all remain consistent. This confirms that the results reported in Section 5 are not artefacts of a particular domain geometry but reflect the underlying structure of the discretisation.

A.2 Implementation parameters

The simulation parameters used in all numerical experiments are summarised in Table 3.

Parameter	Value	Description
K_{sim}	501	Monte Carlo samples for approximating $\mathbb{E}[\cdot]$. Large enough for the Central Limit Theorem to give a good approximation.
N	15	Stochastic modes per spatial direction; noise expanded in $N \times N$ eigenfunctions. Satisfies the Nyquist criterion ((Trefethen, 2000)).
T	5.0	Final simulation time. Enough to see if divergent or not, any more is too expensive.
h_{max}	0.2	Maximum finite-element diameter. Balance between resolution and cost, see Section 5.3.1.
H	1.0	Vertical depth of domain. Aspect ratio as in (Chacón Rebollo et al., 2012).
L	5.0	Horizontal length of domain.
τ_T	$0.1h_{\text{max}}^2$	Brezzi-Pitkäranta pressure-stabilisation parameter, see (John, 2016).
Δt	10^{-3} to 10^0	10 time-step sizes: 0.001, 0.002, 0.005, 0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1.0.

Table 3: *Simulation parameters used in all numerical experiments.*

For computational efficiency and because anisotropic viscosity is not the focus of this study, we use isotropic viscosity, that is $\text{nu}_1 = \text{mu}_1$, so that horizontal and vertical viscous diffusion are treated identically. We define the finite-element order for velocity and pressure both to be 1, that is $\text{orderv} = 1$ and $\text{orderp} = 1$, giving a lowest-order equal-order pair stabilised by a Brezzi-Pitkäranta term. These orders are used in order to minimize computational costs.

We set the wavelength parameter needed for our testcase mentioned shortly $\text{wl} = 1$. The fluid density $\text{rho} = 1$ following Section 2, and the gravitational acceleration $\text{grav} = 9.81$ (SI units).

The test case is $\text{testcase} = \text{'rebollo2'}$, referring to the manufactured-solutions that is used in order to verify that the deterministic parts of our model Equation (2.16) are solid and that we can then isolate the effects of the stochastic forcing Section 2.2.4 and the bathymetry changes coming from Section 3.7.

The boundary conditions are $\text{BCtype} = \text{'dirichlet'}$, imposing homogeneous Dirichlet data on the lateral and bottom boundaries while leaving the surface unconstrained. As discussed in Section 4.4, this departs from the periodic lateral conditions used in the analysis but does not affect the conclusions drawn there.

A.3 Computational implementation details

The algorithmic structure of the experiment is given in Section 5.1.3 in the main text; this appendix collects the MATLAB-level implementation choices.

Most of the helper functions such as `StochTerm` were given by Josefin Ahlkrona. This codebase was used as the foundation and then expanded on. Personal coding consisted of continuous adjustments to make the given code work with the implemented coordinate transformation, fixing the scaling of the truncated Q-Wiener process, adding the measurements of the surface-pressure-gradient norms used throughout Section 5 (the accumulated norm $\mathbb{E}[\sum \Delta t \|\nabla_{HP_s}\|^2]$ and the divergence diagnostic), and parallelising the Monte-Carlo loop. The collaborative process was done via GitHub with the branch `feature/varying-bathymetry`.

Claude was used to generate documentation, refactoring and the idea of parallelisation. A risk of the parallelisation is how the toolbox handles randomness. A concern was that each `parfor` worker on the CPU would share the same global RNG-state. Even worse that two workers may get the same random number if seeded equally which would lead to correlated samples and therefore a bias in the Monte-Carlo estimation. This was resolved by giving each sample s its own seed s via a Mersenne Twister, see (The MathWorks, Inc., 2024) for more on this. In practice this makes it unlikely to have overlap between samples. This simple fix is

```
1 parfor sample = 1:Ksim
2     rng(sample, 'twister'); % <-- fix
```

matlab

Claude was also used for prettier plots since this is essentially just a miscellaneous task.

Remark on the rigour of the seeding strategy. Seeding each Monte Carlo sample with a distinct integer guarantees that no two `parfor` workers share an initial RNG state, but it does not on its own provide a mathematical guarantee that the resulting sample streams are statistically independent: independence between Mersenne Twister streams initialised from nearby integer seeds is empirical rather than provable. Given the Mersenne Twister period of $2^{19937} - 1$ and that the $K_{\text{sim}} = 501$ samples used here occupy a vanishingly small fraction of that state space, overlap or correlation between streams is exceedingly unlikely in practice, and the Monte Carlo standard error reported in Section 5 is consistent with the assumption of independence. A fully rigorous alternative is to use a generator with proven-independent substreams, such as MRG32k3a (L'Ecuyer, 1999), assigning one substream per sample; in MATLAB this amounts to replacing `rng(s, 'twister')` inside the `parfor` body with a `RandStream('mrg32k3a', 'Seed', 0)` whose `Substream` property is set to s . We adopt the simpler integer-seed approach in the experiments reported here; the distinction is noted for completeness.

Simplified loop structure of the stability suite. The `parfor` loop distributes Monte Carlo samples across CPU workers. Each sample sets its own RNG seed, ensuring independence and reproducibility regardless of execution order. The LU factorization is precomputed once per Δt since the system matrix is time-level-independent.

Bibliography

- Binz, T., Hieber, M., Hussein, A., & Saal, M. (2024). The primitive equations with stochastic wind driven boundary conditions. *Journal De Mathématiques Pures Et Appliquées*, 183, 76–101. <https://doi.org/10.1016/j.matpur.2024.01.001>
- Carelli, E., & Prohl, A. (2012). Rates of Convergence for Discretizations of the Stochastic Incompressible Navier–Stokes Equations. *SIAM Journal on Numerical Analysis*, 50, 2467–2496. <https://doi.org/10.1137/110845008>
- Chacón Rebollo, T., Gómez Mármol, M., & Sánchez Muñoz, I. (2012). Numerical solution of the Primitive Equations of the ocean by the Orthogonal Sub-Scales VMS method. *Applied Numerical Mathematics*, 62(4), 342–359. <https://doi.org/10.1016/j.apnum.2011.06.012>
- Feng, X., Prohl, A., & Vo, L. (2021). Optimally convergent mixed finite element methods for the stochastic Stokes equations. *IMA Journal of Numerical Analysis*, 41(3), 2280–2310. <https://doi.org/10.1093/imanum/drab006>
- Griffies, S. M. (2004). *Fundamentals of Ocean Climate Models*. Princeton University Press. <http://www.jstor.org/stable/j.ctv301gzg>
- Hallern, C. von, & Ahlkrone, J. (2026, March). *A splitting-scheme for the linear stochastic primitive equations*.
- John, V. (2016). *Finite Element Methods for Incompressible Flow Problems* (Issue 51). Springer. <https://doi.org/10.1007/978-3-319-45750-5>
- L'Ecuyer, P. (1999). Good Parameters and Implementations for Combined Multiple Recursive Random Number Generators. *Operations Research*, 47(1), 159–164. <https://doi.org/10.1287/opre.47.1.159>
- Lions, J.-L., Temam, R., & Wang, S. (1992). On the equations of the large-scale ocean. *Nonlinearity*, 5(5), 1007–1053. <https://doi.org/10.1088/0951-7715/5/5/002>
- Lord, G. J., Powell, C. E., & Shardlow, T. (2014). *An Introduction to Computational Stochastic PDEs*. Cambridge University Press.
- Madec, G., Bell, M., Blaker, A., Bricaud, C., Bruciaferri, D., Castrillo, M., Calvert, D., Chanut, J., Clementi, E., Coward, A., Epicoco, I., Éthé, C., Ganderton, J., Huthnance, J., Iovino, D., Lea, D., Lovato, T., Martin, M., Martin, N., ... Wilson, C. (2024, January). *NEMO Ocean Engine*. Zenodo. <https://doi.org/10.5281/zenodo.14515373>
- The MathWorks, Inc. (2024, January). *Control Random Number Streams on Workers*. <https://www.mathworks.com/help/parallel-computing/control-random-number-streams-on-workers.html>
- MITgcm Group. (2018, January). *MITgcm: MIT General Circulation Model*. Zenodo. <https://doi.org/10.5281/zenodo.1409237>
- Randall, D. A., Wood, R. A., Bony, S., Colman, R., Fichet, T., Fyfe, J., Kattsov, V., Pitman, A., Shukla, J., Srinivasan, J., Stouffer, R. J., Sumi, A., Taylor, K. E., AchutaRao, K., Allan, R., Berger, A., Blatter, H., Bonfils, C., Boone, A., ... McAvaney, B. (2007). *Climate Models and Their Evaluation*.
- Trefethen, L. N. (2000). *Spectral Methods in MATLAB*. Society for Industrial, Applied Mathematics. <https://doi.org/10.1137/1.9780898719598>
- Tucciarone, F. L., Li, L., Mémin, E., & Chandramouli, P. (2025). Derivation and Numerical Assessment of a Stochastic Large–Scale Hydrostatic Primitive Equations Model. *Journal of Advances in Modeling Earth Systems*, 17(7), e2024MS004783. <https://doi.org/10.1029/2024MS004783>