Prediction under uncertainty: from models for marine-terminating glaciers to Bayesian computation

by

Andrew D. Davis
B.S., Clarkson University (2010)
S.M., Massachusetts Institute of Technology (2012)

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Computational Science and Engineering
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June 2018

Signature redacted

Author .....
Department of Aeronautics and Astronautics
May 15, 2018

Signature redacted

Certified by .....
Youssef Marzouk
Associate Professor, Massachusetts Institute of Technology
Thesis Supervisor

Certified by ..... Signature redacted
Patrick Heimbach
Associate Professor, The University of Texas at Austin
Thesis Supervisor

Certified by ..... Signature redacted
Karen Willcox
Professor, Massachusetts Institute of Technology
Thesis Committee Member

Accepted by ..... Signature redacted
Hamsa Balakrishnan
Chairman, Graduate Program Committee

Accepted by ..... Signature redacted
Nicolas Hadjiconstantinou
Co-Director, Ph.D. Program in Computational Science and Engineering

Prediction under uncertainty: from models for marine-terminating glaciers to Bayesian computation

by

Andrew D. Davis

Submitted to the Department of Aeronautics and Astronautics
on May 15, 2018, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Computational Science and Engineering

Abstract

The polar ice sheets have enormous potential impact on future global mean sea level rise. Recent observations suggest they are losing mass to the ocean at an accelerated rate. Skillful prediction of the ice sheets' future mass loss remains difficult, however; observations of key variables are insufficient and physical processes are poorly understood. Even when a relatively accurate dynamical model is available, computational limitations make it difficult to characterize uncertainties associated with the model's predictions. To address this prediction challenge, this thesis presents complementary developments in glaciology and in Bayesian computation. In particular, (i) we develop new models of marine-terminating glaciers whose dynamics are controlled by an extended set of physical processes and geometric constraints; and (ii) we develop new sampling algorithms to efficiently characterize selected marginals of a high-dimensional probability distribution describing uncertain parameters. The latter algorithms have broader utility in Bayesian modeling and inference with computationally intensive models.

We begin by studying laterally confined ice streams that terminate in the ocean, where they may form floating ice shelves. Such marine-terminating outlet glaciers are the main conduits by which Greenland and Antarctica drain their ice mass into the ocean. Ice shelves play an important role in buttressing the grounded inland ice. The seaward ice flow is typically accompanied by acceleration and thinning. Increased thinning eventually leads to flotation of the ice supported by buoyant forces from the ocean. The transition region from grounded to floating ice is referred to as the grounding line (or zone), and the mass transport across the grounding line as the output flux. Previous work by Weertman (1974) and Schoof (2007) considers laterally unconfined ice streams, showing that their output flux is a monotonically increasing function of the bedrock rock depth at the grounding line. This scenario leads to the marine ice sheet instability (MISI): retreating into deeper water increases the output flux, and retreat accelerates. Therefore, stable steady states cannot exist on downward sloping beds. We extend this analysis to laterally confined glaciers and investigate when side-wall drag is sufficient to stabilize glaciers on downward sloping

beds. Additionally, we include a parameterization of sub-shelf melt. We find that, whereas lateral drag can stabilize glaciers that would otherwise be subject to the MISI, sub-shelf melt can destabilize them.

Our ultimate goal is to predict future ice sheet volume and to quantify its uncertainty. We do so in the Bayesian statistical setting, conditioning our prediction on available observations. Yet characterizing a posterior distribution—using, for example, Markov chain Monte Carlo (MCMC)—involves repeated evaluations of an ice stream model, which are prohibitively expensive. Furthermore, the model parameters that need to be inferred are high dimensional, even though we are primarily interested in a *low dimensional* quantity: the future ice volume. We address this computational challenge by developing new structure-exploiting Monte Carlo methods that combine marginalization with surrogate modeling. Given a high-dimensional (posterior) distribution on the model parameters, whose density evaluations are computationally intensive, we construct an MCMC chain that directly targets a particular low-dimensional marginal of interest. In general, the marginal density is not available analytically. Instead, we can compute unbiased noisy estimates of this density. Our MCMC algorithm incrementally constructs a local regression approximation of the target marginal density using these estimates. Continual refinement of the approximation, as MCMC sampling proceeds, leads to an asymptotically exact characterization of the desired marginal distribution. Analysis of the bias-variance tradeoff guides an ideal refinement strategy that balances the decay rates of different components of the error. Our approach exploits regularity in the marginal density to significantly reduce computational expense relative to both full-dimensional and pseudo-marginal MCMC.

Thesis Supervisor: Youssef Marzouk
Title: Associate Professor, Massachusetts Institute of Technology

Thesis Supervisor: Patrick Heimbach
Title: Associate Professor, The University of Texas at Austin

Acknowledgments

I would like to thank my advisors Professors Youssef Marzouk and Patrick Heimbach. Their advice, support, and insight have made this thesis possible. I would also like to thank my committee member Professor Karen Willcox and external readers Professors Aaron Smith and Christian Schoof. Finally, I would like to thank other co-authors and collaborators Dr. Matt Parno, Dr. Patrick Conrad, and Professor Natesh Pillai for their contributions and support. This thesis would not have been possible without their help!

I also benefitted from discussions with my colleagues in the Aerospace Computational Design Laboratory and in the Center for Computational Engineering at MIT. I would like to thank all of the students, faculty, and staff that made my experience at MIT productive, educational, and positive.

Finally, I would like to thank the organizations and grants that funded my work. The National Science Foundation and the U.S. Department of Energy funded portions of this thesis, specificially under the grants: NSF OCE #1737759 and DoE/SciDAC SC0008060. I was also were supported by the Air Force Office of Scientific Research Computational Mathematics program and from the National Science Foundation Software Infrastructure for Sustained Innovation (SI2:SSI) program, under grant number ACI-1550487.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Glaciers and ice sheets, most notably those in West Antarctica, have begun rapidly retreating [Joughin et al., 2014a, Rignot et al., 2011]. Mass lost by land-based glaciers is transported into the ocean, where it has huge potential impact on global mean sea level [Gardner et al., 2013, Meier et al., 2007, Pfeffer et al., 2008]. Yet, limited observations (data are often sparse in space and/or time), uncertain ice dynamics, and computationally intensive models make characterizing uncertainty about predicted glacier behavior difficult. This thesis will introduce new models relating glacier retreat and volume loss rates to glacier geometry, and will develop new Bayesian computational methodologies to characterize uncertainty in model predictions.

Large-scale industrial activity leads to dramatic, potentially catastrophic, changes in climate [Allen et al., 2014]. Despite consensus that the Earth's climate is currently undergoing unprecedented changes [Alley et al., 2003], predictive models that specifically determine what these changes will be and their societal impact remain ellusive. Climate and environmental science is plagued with uncertainty: measurement error, errors in mathematical models, and missing knowledge of underlying processes. These largely unquantified uncertainties translate to an unpredictable future. In spite of uncertainty and unpredictability, governments, policymakers, and communities must make crucial decisions about their role in mitigating and adapting to the risks associated with climate and environmental changes. Uncertainty must be quantified but physical models alone are inadequate. Coupling mathematical and physical models

with observational data and rigorous statistical methods enables scientists to identify and reduce these sources of uncertainty.

Decision making (e.g., developing policy) requires quantified uncertainty along with physical predictions. Models predicting increasing ocean temperature, leading to thermal expansion and sea level rise [Church et al., 1991, Intergovernmental Panel on Climate Change, 2014, Meehl et al., 2005], have prompted coastal regions to rely on flood defense to prevent shoreline retreat and subsidence, while governments plan for coastal flooding, decline of salt-marshes, and saltwater intrusion into groundwater supply [Nicholls and Cazenave, 2010]. However, many models do not account for glacial and ice sheet retreat [Meehl et al., 2005, Nicholls and Cazenave, 2010]. Scientists must, at minimum, indicate if their predictions are a worst-case, best-case, or most likely scenarios. An unlikely worst-case scenario may warrant investing in coastal flood prevention or changing aquifer management if the risk is high enough [Allen et al., 2014].

The Intergovernmental Panel on Climate Change's (IPCC) qualitative uncertainty description predicts with "medium confidence" that the Antarctic ice sheet significantly contributed to mean sea level during the last interglacial period (116 000 yr to 129 000 yr ago), suggesting models not accounting for Antarctic retreat are conservative [Intergovernmental Panel on Climate Change, 2014]. Recent research suggests West Antarctica has begun collapsing [Joughin et al., 2014b, Rignot et al., 2014] and the Amundsen Sea sector could contribute 0.2 millimeters per year to mean sea level [Thomas et al., 2004]. Quantifying the probability that this and other scenarios occur determines how much more mean sea level rise is expected due to West Antarctic contribution.

Currently, ice sheet models do not directly relate ice sheet volume to input parameters. Many models predict remarkably fast West Antarctic retreat by evolving ice thickness and velocity fields [Hughes, 1975, Joughin et al., 2014b, Pollard and DeConto, 2009, Rignot and Thomas, 2002, Williams Jr. and Ferrigno, 2013]. Although observations confirm the ice sheet has recently been retreating [Fretwell et al., 2013, Li et al., 2013, Rignot et al., 2014, Schutz et al., 2005]. Large-scale models, simulating

the entire ice sheet, attempt to recreate Antarctic behavior on a continental scale but cannot relate volume change to key parameters, such as local basal topography around the glacier's terminus and a calving law (a law that prescribes how ice breaks off the terminus). Additionally computational challenges often force modelers to neglect relevant physics (e.g., hydrological process or ice-ocean interaction) or make limiting steady-state assumptions. Smaller-scale models are used to model specific physical processes—such as, relating output flux to the geometry of the bedrock [Schoof, 2007b, Schoof, 2007a]—but such models are not sufficient to predict the ice sheet volume.

Glaciological inverse problems often reconstruct basal properties, friction or topography, which have arbitrarily high, potentially infinite, dimension [Isaac et al., 2015, MacAyeal, 1993, Martin et al., 2012, Petra et al., 2014, Petra et al., 2012]. Theoretically, uncertainty over the reconstructions should be propagated onto the quantity of interest (QoI) through predictive models. Simulating this, however, is often computationally infeasible.

Bayesian methods provide the theoretical framework and computational tools to make probabilistic predictions using physical models and data [Jaynes, 1957, Kennedy and O'Hagan, 2001]. Joint probability distributions over high dimensional inferred parameters (e.g., basal topography) and unknown or uncertain quantities of interest (Lois) (e.g., ice volume) characterize the uncertainty in both the inferred parameters and the QoI. Marginalizing over the inferred parameters defines a distribution over the scalar (or, more generally, low dimensional) QoI. The end goal of Bayesian prediction is to characterize these marginal distributions. However, marginal densities rarely have an analytic form and must be approximated stochastically (e.g., using Monte Carlo methods). Current sampling methods, including approximate Bayesian computation [Beaumont et al., 2002, Marin et al., 2012] and pseudo-marginal Markov chain Monte Carlo [Andrieu and Roberts, 2009, Beaumont, 2010], are infeasible if the target density is computationally expensive.

1.1 Background

Our research lies at the intersection of three disciplines: (i) environmental observation, data collection, and data analysis, (ii) ice sheet modeling and viscous fluid mechanics and (iii) Bayesian statistics. Figure 1-1 provides some context as to how these fields intersect. In our application, physical models predict ice sheet retreat/advance in a changing climate and the Bayesian paradigm provides a framework to quantify associated uncertainties. Ice sheet models and the their supporting theories explain observations and, ideally, predict future behavior. Bayesian tools use physical models, uncertainty models, and data to define probability distributions simulating errors and uncertainties. Often probability densities associated with these distributions are computationally expensive to evaluate, high dimensional, and complicated (e.g., non-Gaussian), making them difficult to characterize and/or sample.

Previous work combines observations of ice surface velocity and a continental scale ice sheet model to infer sub-glacial properties, such as basal friction [Isaac et al., 2015, Martin et al., 2012, Petra et al., 2014, Petra et al., 2012]. These methods reconstruct the observed velocities using a computationally intensive Stokes' model with nonlinear viscosity and nonlinear forcing. Abstractly, their model takes a basal friction field, which parameterizes hydrological and thermal process at the ice-bedrock interface, and computes the surface velocities by solving a partial differential equation. Using a similar problem formulation as [MacAyeal, 1993], the most likely input parameter is found by minimizing the model-data mismatch. Uncertainties are quantified by examining the cost function's curvature at the optimal point—information that is obtained by computing the Hessian of the cost function. A Gaussian distribution using the optimal parameter value as the mean and the Hessian as the inverse covariance matrix *approximates* a Bayesian posterior over the input parameters. The quality of this approximation is unknown; it depends on the nonlinear relationship between the basal frictional field and the observations. Never the less, this method provides useful information about how well constrained the inferred field is by data and can be used to identify a subspace of parameter space that is well constrained by observations.

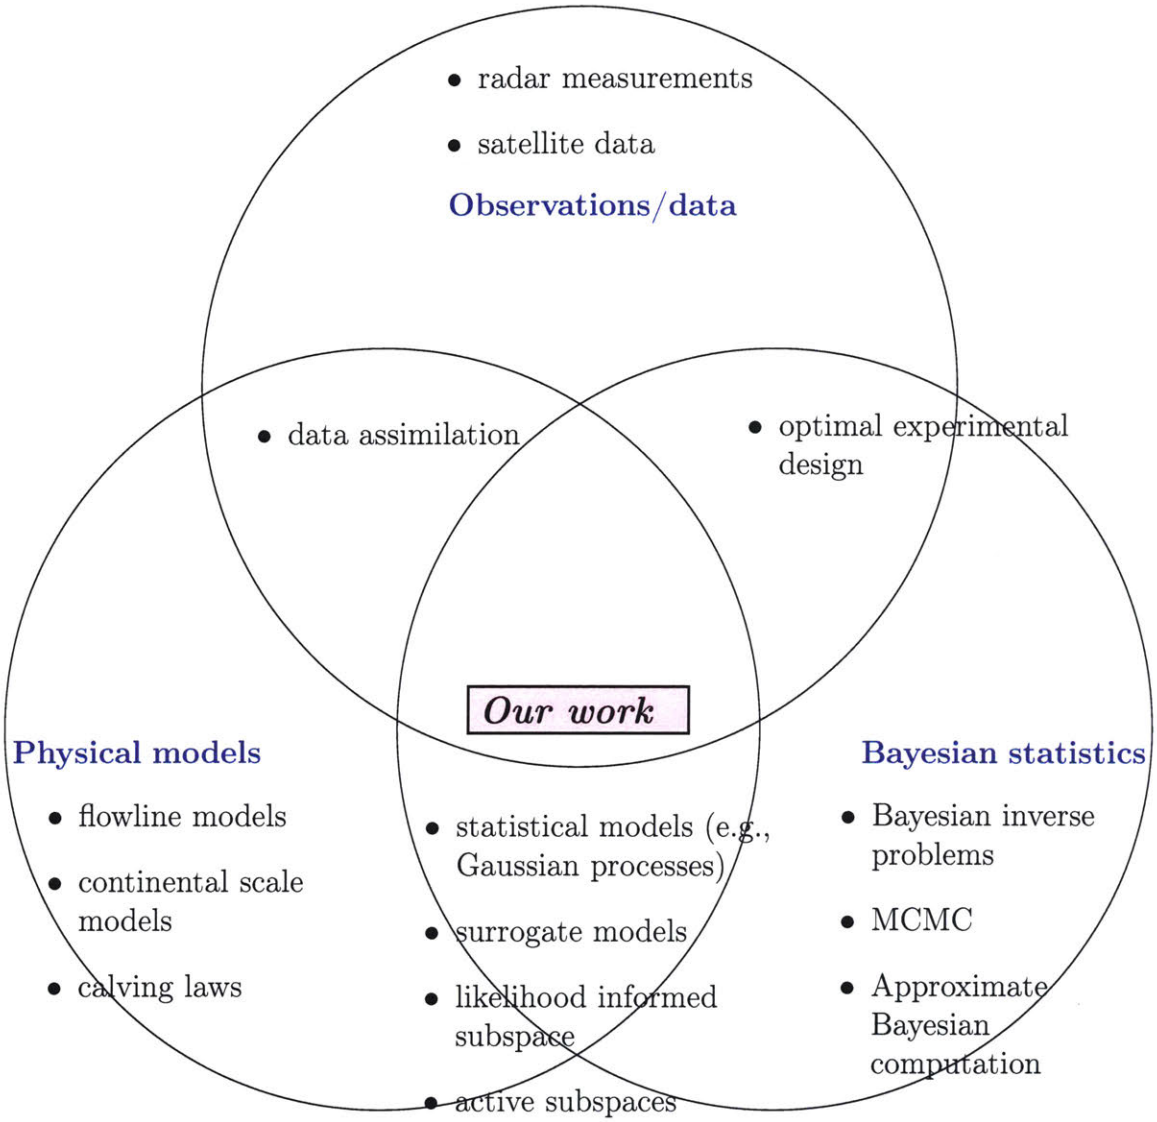

Figure 1-1: A schematic showing how our research intersects with data collection, ice sheet modeling, and Bayesian statistics. Citations: observations/data [Fretwell et al., 2013, Li et al., 2013, Schutz et al., 2005], flowline models [Dupont and Alley, 2005, Hindmarsh, 2012, Jamieson et al., 2012, Pegler et al., 2013, Pegler, 2016], continental scale models [Isaac et al., 2015, MacAyeal, 1993, Martin et al., 2012, Petra et al., 2014, Petra et al., 2012], calving laws [Bassis and Jacobs, 2013, Bassis and Walker, 2011, Ma et al., 2017, Pollard et al., 2015, Nick et al., 2010, Van der Veen, 1998a, Van der Veen, 1998b, Weertman, 1973, Weertman, 1980], Bayesian statistics [Brooks et al., 2011, Kennedy and O'Hagan, 2001, Owen, 2013], statistical models [Bernardo et al., 2008, Santner et al., 2003, Sacks et al., 1989], surrogate models [Bliznyuk et al., 2012, Joseph, 2012, Smolyak, 1963, Wiener, 1938], data informed subspace [Constantine et al., 2014, Cui et al., 2014, Zahm et al., 2018], and experimental design [Chaloner and Verdinelli, 1995, Huan and Marzouk, 2013]

Often, these methods make steady-state assumptions, which are both physically inaccurate and make time-evolution impossible. We can reduce the computational expense of ice sheet models by targeting specific regions within the ice sheet. Ice streams—river-like, fast-flowing ice embedded in a larger, slow-flowing ice sheet—are individually modeled by flowline models [Dupont and Alley, 2005, Hindmarsh, 2012, Jamieson et al., 2012, Pegler et al., 2013, Pegler, 2016, Schoof et al., 2017]. The transition from the slow-flowing ice sheet to the fast-flowing ice stream creates a shear margin [Suckale et al., 2014]. The ice stream is modeled as the width-averaged thickness and velocity field flowing over deformable till with deformable side margins (although our work will assume that the side margin is geographically constrained). Flowline models can be used to make predictions about future behavior; combining multiple flowline models can approximate the behavior of an *entire* ice sheet.

Computational limitations reduce the utility of such methods in the context of inference-for-prediction. Often, inferring high-dimensional fields, such as basal friction, is only an intermediate step in quantifying uncertainty over a low-dimensional QoI. Typically, data inform a small subspace of parameter space and we can restrict our inference problem to this reduced space [Constantine et al., 2014, Cui et al., 2014, Zahm et al., 2018]. Additional dimension reduction can be achieved by restricting the inference problem to dimensions that are *both* informed by data and inform the QoI [Lieberman et al., 2010, Spantini et al., 2017]. However, these methods rely on knowing such subspaces—either assuming linearity and computing a singular value decomposition or learning the data-informed subspace from posterior model evaluations.

Our approach will define a joint posterior distribution over the inferred parameter and the QoI. This defines a high dimensional joint space that we reduce by targeting the marginal distribution over the QoI. We use sampling methods to characterize the joint distribution. However, these methods necessitate repeatedly calling the physical model (often $\gg 10^5$ times), which is computationally prohibitive. Sampling methods such as approximate Bayesian computation [Beaumont et al., 2002, Marin et al., 2012] and delayed acceptance MCMC [Christen and Fox, 2005, Sherlock et al., 2017]

attempt to mitigate this computational expense, but the number of model evaluations still scales at least linearly with the number of generated samples. Even in settings where model evaluations are cheap, targeting the marginal distribution directly is difficult. Psuedo-marginal MCMC exactly characterizes marginal distributions [Andrieu and Roberts, 2009, Beaumont, 2010], but this method can be even more computationally intensive, due to the cost of estimating marginal densities.

1.2 Outline

This thesis will advance the state of the art in ice sheet modeling and Bayesian computation, addressing some of the challenges associated with quantifying the uncertainty about predictions of future ice sheet behavior. Chapter 2 describes a flowline model that includes lateral stress due to side wall drag, subshelf melt, and multiple calving laws. We numerical solve this model and analyze steady-state stability. Chapter explores the relationship between ice stream geometry and output flux by analyzing the previously described flowline model. Chapter 4 derives a Bayesian computational method that characterizes distributions with computationally intensive probability density functions and that can directly target selected marginals. Chapter 5 extends our algorithm to a parallel implementation, improving efficiency as a function of run-time, and provides practical examples. These advances yield physical models and statistical tools that, when combined with data, will help quantify and reduce uncertainty over predictions of future ice sheet behavior.

Chapter 2

Flowline models for laterally confined marine ice streams: numerical solutions and stability analysis

We are interested in whether or not a glacier is subject to the *marine ice instability*. If the output flux increases as the glacier retreats on a retrograde bed, then retreat accelerates and the glacier rapidly collapses. A number of mechanisms have been proposed that could stabilize the marine ice instability: back-stress from the ice shelf reduces the flux at the grounding line, isostatic rebound that reduces the bedrock slope, thermomechanical coupling, and lateral drag due to geometric confinement. We are interested in relating the glacier's stability to geometric properties: the width, the size of the ice shelf, and basal drag.

Calving—the process by which ice breaks off at the terminus—likely plays the most important role in determining the glacier's stability. Unfortunately, calving is also a poorly understood process and there is no universally accepted calving model. Investigating calving physics is worthwhile but beyond the scope of this thesis. Instead, we consider two general calving models: (i) glaciers do not calve—mass is only lost via melt along the ice shelf, and (ii) glaciers instantaneously calving at a critical thickness. Furthermore, we split the calving-limited case into two more sub-cases: (ii.a) glaciers calve at the critical thickness when their weight is supported by the

ocean and (ii.b) the critical calving thickness is determined by meltwater penetration into ice crevasses.

We consider *flowline models*: models that resolve the along-flow coordinate. Such models are widely used and variations parameterize different physical processes. In this chapter we devise a flowline model that includes (i) parameterized lateral drag and (ii) sub shelf melt. We combine this model with three calving laws (i) the melt-limited extreme, (ii.a) calving-at-flotation, and (ii.b) a parameterization of the degree to which meltwater has penetrated the glacier. The particular model considered here is a combination of widely used parameterizations of physical processes.

Our contribution is to devise numerical methods that compute steady-state and transient results. Specifically, we compute steady-state solutions with a shooting method and, for calving-limited glaciers, use discontinuous Galerkin methods to compute transient solutions. Previous work solves similar flowline models with finite volume methods. We propose discontinuous Galerkin methods so that we can take advantage of adaptive mesh refinement associated with finite element methods. This improves computational efficiency, allowing us to investigate glacier stability using high-frequency forcing over long timescales.

In section 2.3, we describe the flowline model and calving laws; specializing existing models to include the physical process we believe are crucial in determining glacier stability (e.g., lateral drag and subshelf melt). In section 2.4 we describe a shooting methods and discontinuous Galerkin method to numerically compute steady state and transient solutions. In section 2.5, we show numerical results that predict glacier stability.

2.1 Motivation

This chapter analyzes stress and mass balance equations that are integrated over the glacier's transverse direction and assume no vertical velocity. These so-called 'flowline models' allow laterally integrated velocity and thickness fields to vary along a transect of the along-flow direction. In this chapter, we first describe the flowline

model, which must also include an assumed calving law. Since there is no universally accepted calving law we consider two generic scenarios: a (i) melt- and (ii) calving-limited glacier. In the melt-limited case, we assume the glacier never calves; instead, ice is only lost via melt underneath the ice shelf. In the calving-limited case, we consider to two models for a calving ice front: (ii.a) calving-at-flotation and (ii.b) calving due to meltwater penetration into ice crevasses. We nondimensionalize the flowline model to reduce the number of free parameters. In chapter 3, the relatively small magnitude of the nondimensional parameter that multiplies the viscous stress will motivate reducing the complexity of this model by neglecting viscous stress in the interior of the glacier.

We numerically compute solutions to the flowline model using a discontinuous Galerkin method. We also integrate in time using an operator splitting method: alternating between updating the velocity and thickness fields and the terminus location. This allows use to (numerically) simulate transient behavior of the glacier, which we will use to assess the glacier's stability in various climate scenarios. Specifically, we are interested in how the glacier responds to changes in time dependent parameters such as surface mass balance, subshelf melt, or the depth that meltwater has penetrated the ice.

2.2 Background

We relate the rate that marine ice streams—fast-flowing systems that transport mass from an ice sheet's interior to its periphery [Bennett, 2003, Truffer and Echelmeyer, 2003]—gain/lose mass to their geometric and climatic parameters. In a basin around an ice stream, accumulated ice deforms under its own weight and creep flow slowly moves mass toward river-like ice streams. Anisotropy in basal properties, such as geographic features or changes in basal friction, cause ice streams to flow significantly faster than the surrounding region. Ice streams are lubricated flows: a gravitational driving stress is balanced by basal and lateral drag. Mass could be lost via surface runoff, but we assume the region/climate is relatively cold and that most of the

mass is lost when ice calves—breaks off—at the terminus or via subglacial melt along the ice shelf. In the marine setting, the ice stream begins to float when it becomes sufficiently thin, which divides it into grounded and floating portions. We refer to the 'extensional stress' as a normal stress on the vertical surface that is not due to cryostatic pressure. Rapid thinning at the terminus or as ice is transported into the floating ice shelf causes the extensional stress to play a leading order role and the ice stream becomes a viscous jet. The rate that mass is transported from the ice stream's interior—the grounded portion—either directly into the ocean or into a floating ice shelf is determined by the local geometry, climatic parameters, and material properties.

The rate that mass leaves the grounded portion, which we refer to as the grounding line flux, is independent of upstream dynamics. Grounding line flux for laterally unconfined marine ice sheets is an increasing function of the bedrock depth at the grounding line [Schoof, 2007b, Schoof, 2007a, Thomas and Bentley, 1978, Weertman, 1974]. Retreat into deeper water increases the grounding line flux, which accelerates retreat. This is the so-called *marine ice instability*: glaciers cannot have stable steady states on a sufficiently steep downward sloping bed [Fowler, 2011, Schoof, 2012]. Given this instability, ice streams only have stable steady states when the bedrock slopes upward toward the interior and many existing ice streams are, therefore, unstable. A number of mechanisms have been proposed that could alter the flux-to-bedrock-depth relationship: hoop stresses in the ice shelf that are exerted on the grounding line [Pegler and Worster, 2012, Pegler, 2016], lithosphere unloading can cause isostatic rebound that changes bedrock depth [Gomez et al., 2010], thermomechanical coupling [Robel et al., 2014, Robel et al., 2016], and lateral drag due to geometric confinement [Dupont and Alley, 2005, Jamieson et al., 2012]. Here, we focus on lateral confinement.

We consider a laterally confined, marine-terminating outlet glacier in a fixed channel with parameterized drag along its side walls. The drag exerted on the glacier reduces the extensional stress at the grounding line and along the ice shelf, where viscous stresses play a leading-order role. The ice flux at the grounding is a function of local bedrock depth and extensional stress [Schoof, 2007b, Schoof, 2007a]. Therefore,

long ice shelves that reduce the extensional stress can also reduce the grounding line flux. Consequently, the ice shelf length plays an important role on the grounding line flux and the glacier's overall stability. The ice shelf length is determined—potentially indirectly—by the choice of calving law. Given a calving law, we analyze the flowline model for a marine terminating glacier using direct numerical simulation and matched asymptotic expansion. This analysis relates grounding line flux to the local bedrock depth, ice shelf geometry, basal and lateral drag, and calving parameters. Furthermore, we incorporate ice-ocean interactions beneath the ice shelf by parameterizing subshelf melt and analyze how this affects grounding line flux.

Buttressing due to lateral confinement can lead to stable steady states when the bedrock slopes down toward the interior. [Goldberg et al., 2009, Gudmundsson et al., 2012] consider ice streams through uniform width channels with a fixed terminus location. In this case, retreating into deeper water lengthens the ice shelf, also increasing the backstress from lateral drag along the shelf, which leads to stable steady states if the backstress sufficiently reduces the extensional stress. This process, where the ice shelf offsets the potential flux increase from the deepened grounding line location, is called *ice shelf buttressing.*

We determine if similar stabilization is possible due to ice sheet buttressing if the terminus location evolves according to a prescribed calving law [Schoof et al., 2017]. In this case, the terminus location is free to evolve and grounding line retreat does not ensure that the ice shelf's length will increase. Unfortunately, there is no universally accepted calving law and we, therefore, investegate this problem for a variety of calving parameterizations. Specifically, we are interested in two widely used calving laws: (i) calving-at-flotation and (ii) a calving law that depends on meltwater penetration into surface and basal crevasses. The model lacks the sophistication to resolve individual crevasses and, therefore, we rely on parameterizations [Dupont and Alley, 2005, Hindmarsh, 2012, Jamieson et al., 2012, Nick et al., 2010, Pegler, 2016, Pegler et al., 2013, Robel et al., 2014, Robel et al., 2016]. We will demonstrate, at least in the limit where extensional stress is increasingly negligible in the ice sheet's interior, that the grounding line flux can be computed cheaply as a function of local

geometry.

Additionally, we investigate how sub-shelf melt, an additional out-flux along the ice shelf, affects the grounding line flux and the transient behavior of the ice stream. The assumed calving law—which was originally proposed in [Nick et al., 2010]—parameterizes hydrological processes at the terminus. Calving events are triggered when water penetrates through the ice. We investigate how including mass loss due to subshelf melt affects the behavior of the flux at the grounding line.

We consider three calving laws: (i) calving-at-flotation, (ii) an extreme melt-limited case when mass is *only* lost via sub-shelf melt, and (iii) the meltwater calving law assumed by [Schoof et al., 2017] including sub-shelf melt. We find that the calving-at-flotation case is simply the melt-limited extreme with an infinite melt rate and the meltwater calving law changes the flux-to-bedrock-depth relationship, potentially stabilizing the marine ice instability. We show that the ice stream has multiple steady states and there exists a stable region, regardless of bedrock slope, where the ice stream evolves to a unique stable steady state.

2.3 Flowline model

We consider a flowline model for a channelized outlet glacier, where mass, accumulated over the entire domain, is advected seaward. Similar flowline models are found in [Dupont and Alley, 2005, Hindmarsh, 2012, Jamieson et al., 2012, Nick et al., 2010, Pegler et al., 2013, Pegler, 2016, Schoof et al., 2017] and figure 2-1a shows the physical domain. Importantly, our model includes parameterized lateral drag and allows mass loss via melt along the ice shelf base. Lateral drag potentially stabilizes the marine ice instability—either slowing or preventing retreat. Let x be the along-flow coordinate and t time, while $u(x,t)$ and $h(x,t)$ are unknown width-averaged surface velocity and ice thickness, respectively. We compute u and h by solving the stress

and mass conservation equations

$$\underbrace{2\left(A^{-\frac{1}{n}}h\left|u_x\right|^{\frac{1}{n}-1}u_x\right)_x}_{\text{Extensional stress}} - \underbrace{C_l w^{-\frac{1}{n}-1}h\left|u\right|^{\frac{1}{n}-1}u}_{\text{Lateral stress}} - \underbrace{\theta C_b\left|u\right|^{m-1}u}_{\text{Basal stress}}$$

$$= \underbrace{\rho_i g\left(1-(1-\theta)\frac{\rho_i}{\rho_w}\right)h\left(h_x+\theta b_x\right)}_{\text{Driving stress}} \tag{2.1a}$$

$$wh_t + (wuh)_x = w\left(a-(1-\theta)\mu\right) \tag{2.1b}$$

for $0 < x < x_c(t)$, where $b(x)$ and $w(x)$ are the (temporally constant) bed elevation and outlet channel width, respectively, $a(t)$ and $\mu(t)$ are the (spatially constant) surface mass balance and sub-shelf melt, respectively, and subscripts indicate partial derivatives. The dimensional parameters (see table 2.1) in (2.1a) and (2.1b) are the flow law coefficient A, the lateral friction coefficient C_l, the basal friction coefficient C_b, the ice density ρ_i, the ocean density ρ_w, and the gravitational acceleration g. We assume, for simplicity, that the flow law parameter A does not depend on temperature or moisture and that the frictional coefficients C_l and C_b are also independent of temperature or other degrees of freedom. The nondimensional parameters n and m are the flow law and friction law exponents, respectively.

The parameterized lateral stress term (second term of 2.1a) models side wall drag. We assume that the side wall exerts a symmetric stress on both grounded and floating portions. Other models consider more complicated parameterization that resolve the transverse direction (e.g., [Wearing et al., 2015]). The stress balance equation (2.16d) becomes more accurate in either the limit of very wide or very narrow channels. In very wide channels, the lateral stress is negligible and the model reduces to the flowline model considered in [Schoof, 2007b, Schoof, 2007a]. In longer, narrow channels the extensional stress is negligible and there is little flow transverse to the long axis.

Let $x_g(t)$ denote the end of the grounded portion, where the ice either begins to float or the glacier terminates and let $x_c(t)$ denote the terminus (either floating or grounded), where ice calves into the ocean or the glacier has thinned until the

g

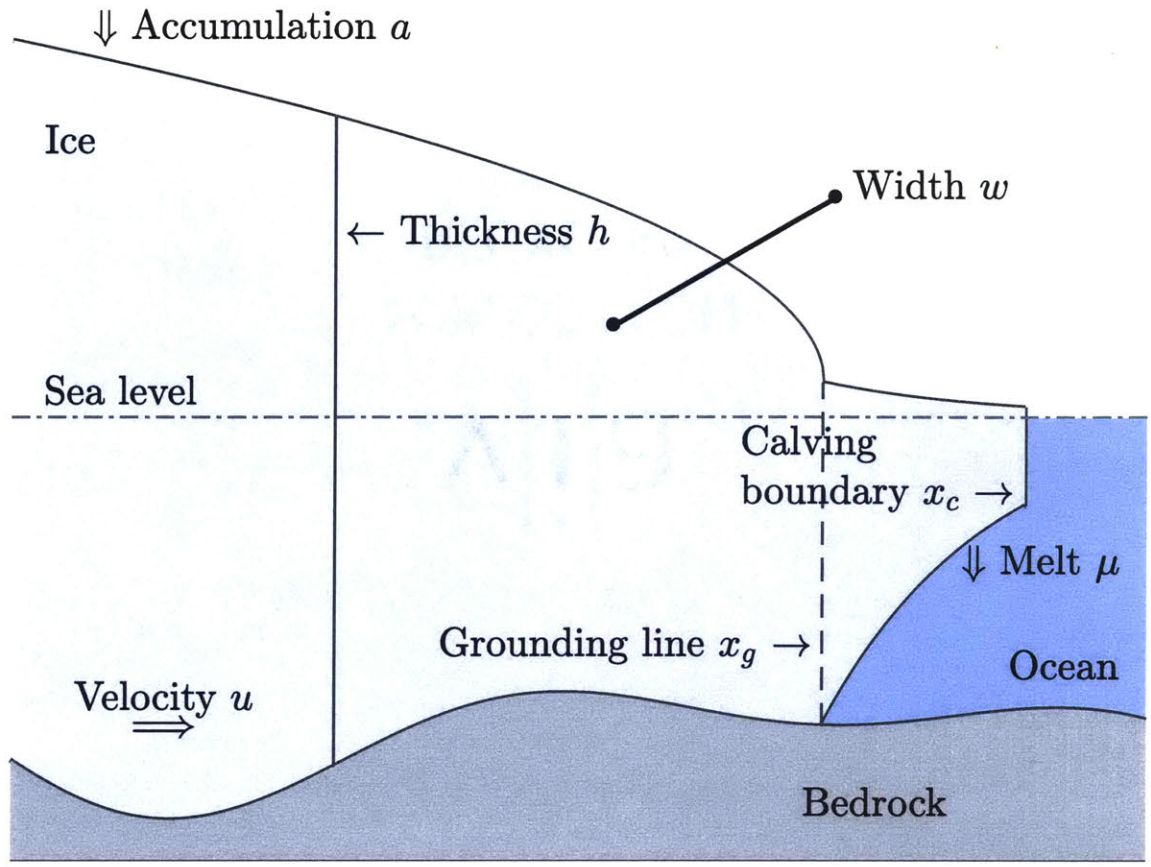

(a) The physical domain

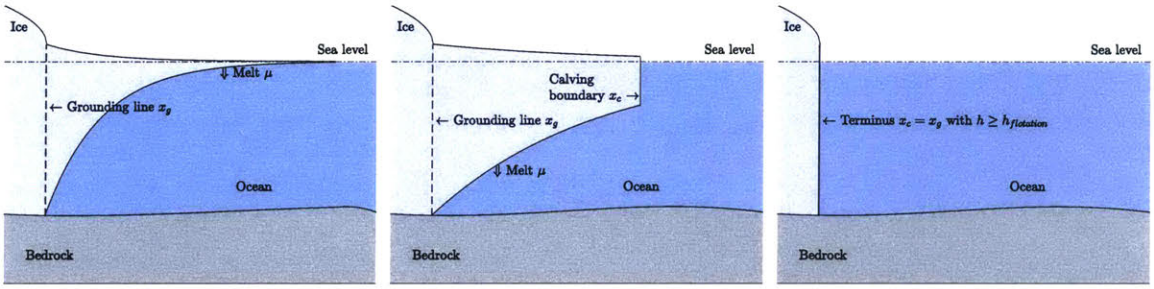

(b) Melt-limited terminus (c) Floating terminus (d) Grounded terminus

Figure 2-1: The physical domain and illustrations of different calving laws.

flux $q = uhw = 0$. When a floating ice shelf exists, the grounded portion ends at the grounding line, where the thickness is equal to the critical flotation thickness $h(x_g(t), t) = -(\rho_w/\rho_i)b(x_g(t))$. In our notation, the thickness may not be equal to the flotation thickness at $x_g(t)$ if there is no floating ice shelf. Often, ice begins to float at the grounding line $x_g(t)$ and terminates at a floating calving boundary $x_c(t)$ (figure 2-1c). However, only allowing mass loss via sub-shelf melt implies there is no calving

	Symbol	Meaning	Value & units
Independent variables	x	Along flow coordinate	m
	t	Time	s
Unknown variables	$u(x,t)$	Ice velocity	$\mathrm{m\,s^{-1}}$
	$h(x,t)$	Ice thickness	m
Known variables	$w(x)$	Width	m
	$b(x)$	Bedrock depth	m
	$a(t)$	Surface mass balance	$\mathrm{m\,s^{-1}}$
	$\mu(t)$	Sub-shelf melt	$\mathrm{m\,s^{-1}}$
	$d_w(t)$	Crevasse water depth	m
	$\theta(x,t)$	Grounded/floating indicator	—
Constants	m	Friction law coefficient	$\frac{1}{3}$
	n	Glen's flow law exponent	3
	A	Glen's flow law coefficient	$2.8 \times 10^{-23}\,\mathrm{Pa^{-3}\,s^{-1}}$
	C_l	Lateral friction coefficient	$8.9 \times 10^{7}\,\mathrm{Pa\,s^{1/3}}$
	C_b	Basal friction coefficient	$2.6 \times 10^{6}\,\mathrm{Pa\,m^{-1/3}\,s^{1/3}}$
	ρ_i	Ice density	$918\,\mathrm{kg/m^3}$
	ρ_w	Ocean density	$1020\,\mathrm{kg/m^3}$
	g	Gravitational acceleration	$9.8\,\mathrm{m/s^2}$

Table 2.1: Dimensional parameters, variables, and constants. Note that the numbers given here are 'reasonable examples.' The exact value with vary for individual glaciers and often there is not an accepted constant value.

boundary (figure 2-1b) and sufficient calving causes the glacier to terminate before reaching its critical flotation thickness (figure 2-1d). After the ice begins to float, the basal frictional stress disappears, the driving stress changes, and mass loss due to sub-shelf melt begins. These changes are modeled by the discontinuous parameter

$$\theta(x) = \begin{cases} 1 & \text{if } \rho_i h(x) \geq -\rho_w b(x) \text{ (grounded ice)} \\ 0 & \text{else (floating ice).} \end{cases} \tag{2.1c}$$

The boundary conditions are determined by the calving law, which is a modeling choice made before solving the equations.

2.3.1 Boundary conditions

We impose boundary conditions at the ice divide $x = 0$ and the terminus $x = x_c(t)$. At the ice divide $x = 0$, require $u = 0$; there is no flow into the ice stream. The boundary conditions at the terminus $x = x_c(t)$ are determined by the calving law. We consider two calving scenarios: a (i) calving- and (ii) melt-limited terminus.

Calving-limited terminus

Our aim is to relate glacier dynamics to calving. Calving is poorly understood and several models have been developed based on either tensile [Nick et al., 2010, Van der Veen, 1998a, Van der Veen, 1998b, Weertman, 1973, Weertman, 1980] or shear [Bassis and Jacobs, 2013, Bassis and Walker, 2011, Ma et al., 2017, Pollard et al., 2015] failure. There is no universally accepted calving law; it is a modeling decision. Different calving laws result in qualitatively different outlet glacier dynamics. Modeling calving dynamics and improving our understanding of the underlying physics is worthwhile but beyond the scope of this study. We relate glacier dynamics to existing calving laws, developing a generic way of assessing a glacier's susceptibility to the marine ice instability.

In the calving-limited case, buoyant forces from the ocean support the ice at the terminus and the terminal thickness is determined by a parameterized calving law. Specifically, the extensional stress at the calving terminus is balanced by hydrostatic stress

$$2A^{-\frac{1}{n}} h \left|u_x\right|^{\frac{1}{n}-1} u_x = \tfrac{1}{2}\left(\rho_i g\left(1-(1-\theta)\tfrac{\rho_i}{\rho_w}\right)h^2 - \theta\rho_w g b^2\right) \quad \text{at} \quad x = x_c(t). \tag{2.2}$$

The calving law prescribes a critical calving thickness $h_c(t)$ that is a fraction of the local bedrock depth. If the terminal thickness is less than $h_c(t)$, the glacier instantaneously calves such that $h(x_c(t), t) = h_c(t)$ and if $h(x_c(t), t) > h_c(t)$, the glacier advances and thins until $h(x_c(t), t) = h_c(t)$. We assume that the critical calving

thickness is a fraction $\phi(t)$ of bedrock depth

$$h_c(t) = -\phi(t) b_c(t) \quad \text{at} \quad x = x_c(t), \tag{2.3a}$$

where $b_c(t) = b(x_c(t))$ is the bedrock depth at the terminus. We consider two choices for ϕ:

$$\phi = \frac{\rho_w}{\rho_i} \quad \text{and} \quad \phi(b_c; d_w) = \begin{cases} -2\frac{\rho_w}{\rho_i} d_w b_c^{-1} & \text{if } -d_w b_c^{-1} < \frac{1}{2} \text{ (floating)} \\ \eta \pm \sqrt{\eta^2 - \frac{\rho_w}{\rho_i}} & \text{else (grounded),} \end{cases} \tag{2.3b}$$

where $\eta(b_c; d_w) = 1 - \left(\frac{\rho_w}{\rho_i} - 1\right) d_w b_c^{-1}$ and $d_w(t)$ is the depth that water has penetrated into ice crevasses. These two definitions (figure 2-2) of $\phi(t)$ correspond to calving-at-flotation and calving when the combined surface and basal crevasse depth equals the ice thickness, respectively [Nick et al., 2010, Schoof et al., 2017]. The terminal thickness $h_c(t) = h(x_c(t), t)$ depends only on the crevasse water depth d_w if the terminus is floating (i.e., $h_c \leq -\frac{\rho_w}{\rho_i} b_c$).

Calving when surface and basal crevasse penetration equals ice thickness is a result of tensile failure. The second choice defined in (2.3b)—which is adapted from [Nick et al., 2010][1]—models this calving scenario. The extensional stress $\Sigma = 2A^{-\frac{1}{n}} |u_x|^{\frac{1}{n}-1} u_x$ is a depth independent normal stress acting on a vertical plane and the ice overburden pressure is a depth dependent compressive stress $\rho_i g(z - h)$ ($z \in [0, h]$ is the vertical coordinate such that $z = 0$ is the base). Therefore, the net normal stress on a vertical plane is $\Sigma + \rho_i g(z - h)$. Surface crevasses penetrate until the net normal stress is larger than the water pressure within the cracks; at this critical depth $d_s = h - z^*$, the normal stress is sufficient to close the cracks. Given that water has penetrated to a depth d_w, the surface water pressure in the cracks is $p_s = \rho_w g d_w$.

[1]Here, we highlight important details from [Nick et al., 2010] to rewrite the calving model using our notation. Specifically, we need $h_c = -\phi(b_c; d_w) b_c$. However, we refer to [Nick et al., 2010] for a more detailed discussion/derivation of this calving law.

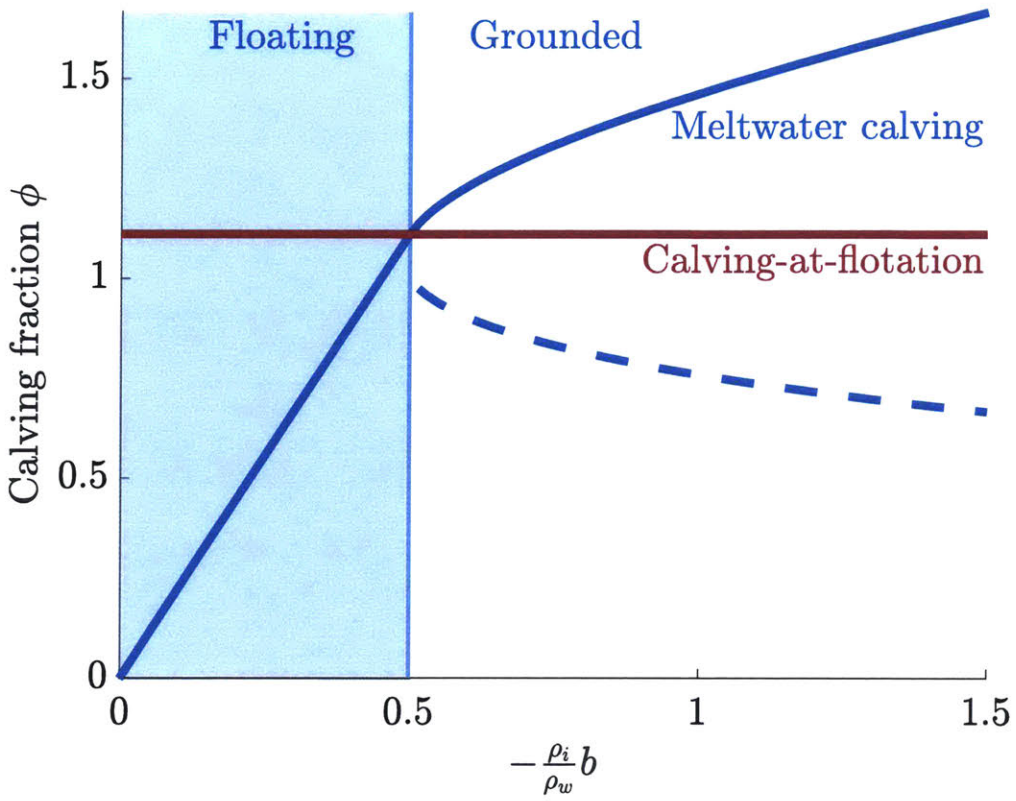

Figure 2-2: The fraction ϕ that defines the calving thickness, as defined by (2.3b). Glaciers either calve at flotation (red lines) or the calving thickness depends on the meltwater penetration d_w (blue lines). In the meltwater penetration case, the glacier either has a floating (blue shaded region) or grounded (unshaded) terminus. The fraction is less than the calving-at-flotation case in the unstable branch (dotted blue line) and, therefore, the glacier is actually floating. However, the dotted blue line corresponds to a terminal condition that violates (2.5).

Surface cracks penetrate to depth d_s, corresponding to $p_s = \rho_i g d_s - \Sigma$:

$$d_s = \frac{\Sigma}{\rho_i g} + \frac{\rho_w}{\rho_i} d_w. \tag{2.4a}$$

Similarly, the depth that crevasses penetrate from the base d_b is the critical point where the basal water pressure $p_b = \rho_w g(h_b - d_b)$ equals the normal stress, where h_b is depth of the ice base

$$h_b = \begin{cases} \frac{\rho_i}{\rho_w} h & \text{if } h < -\frac{\rho_w}{\rho_i} b \text{ (floating)} \\ -b & \text{if } h \geq -\frac{\rho_w}{\rho_i} b \text{ (grounded).} \end{cases} \tag{2.4b}$$

The basal crevasse penetration is, therefore, defined by $p_b = \rho_i g(h - d_b) - \Sigma$:

$$d_b = \frac{\Sigma}{\rho_w - \rho_i} - (\rho_i g h - \rho_w g h_b). \tag{2.4c}$$

At the ice front, hydrostatic balance with the ocean defines the extensional stress Σ via (2.3a). The terminal thickness h_c is equal to the total depth that crevasses have penetrated $h_c = d_s(\Sigma(h_c), d_w) + d_b(\Sigma(h_c), d_w)$; solving for h_c in terms of d_w defines second choice of ϕ in (2.3b). The fraction ϕ and the terminal thickness h_c are plotted in figures 2-2 and 2-3.

The meltwater penetration calving law requires surfaces crevasses to connect with basal crevasses. The model assumes that surface crevasse depth is determined by the depth meltwater has penetrated into the glacier. This calving law predicts no calving when $d_w = 0$—surface crevasses contain no meltwater. More realistically, extensional stress becomes concentrated around crack tips leading to a positive feedback that could cause crack propagation in water-free surface crevasses [Van der Veen, 1998b, Van der Veen, 1998a, Weertman, 1973, Weertman, 1980]. Rather than resolving individual crevasses, we follow the assumption of [Nick et al., 2010] that the extensional stress depends linearly on depth. Additionally, torque near the calving front may allow taller icebergs to calve purely by extensional stress [Hanson and Hooke, 2000, Ma et al., 2017]. Recent work incorporates additional mechanisms such as the role of viscous deformation at crack tips and torque near the calving front into the calving law [Jiménez et al., 2017, Krug et al., 2014, Mobasher et al., 2016, Yu et al., 2017]. However, parameterizing these models in a form analogous to (2.3a) is difficult without resolving individual crevasses. We, therefore, focus on calving that occurs because of meltwater driven surface crevasse propagation and determine the qualitative effect this tensile failure mechanism has on glacier dynamics.

We also note a final pitfall: even assuming that the meltwater calving law from

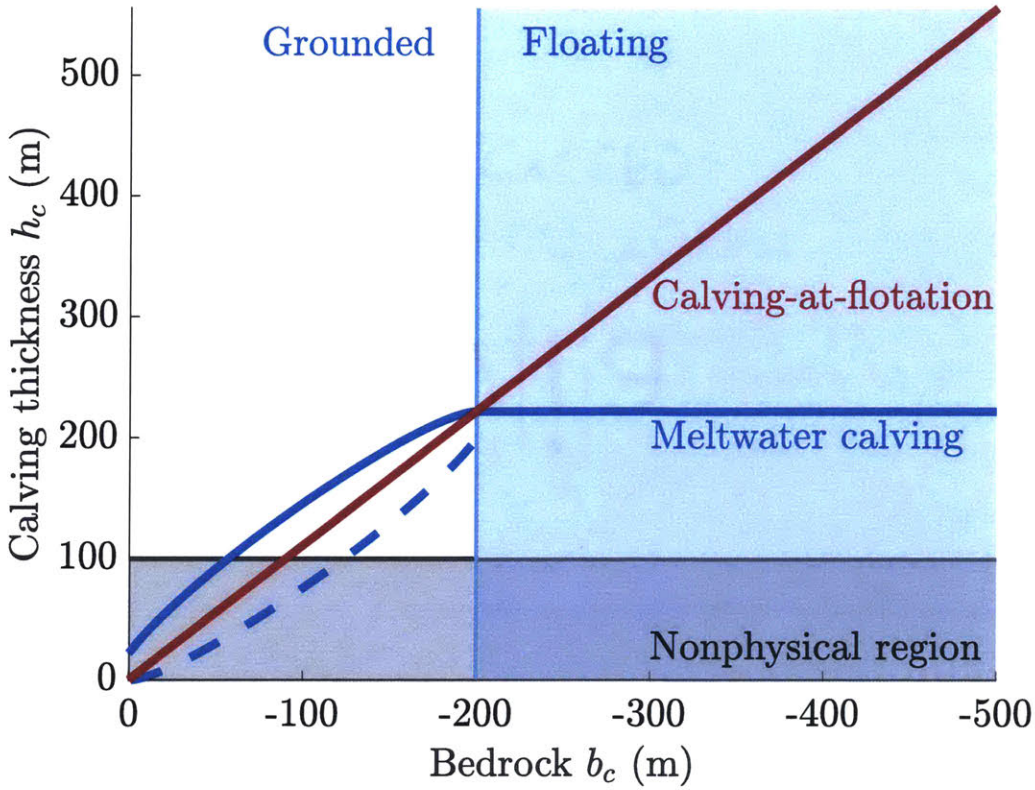

Figure 2-3: The calving thickness h_c as a function of bedrock depth given $d_w = 100$m and $\rho_i/\rho_w = 0.9$. Once the glacier begins to float, the calving thickness is independent of the bedrock depth. Glaciers either calve at flotation (red lines) or the calving thickness depends on the meltwater penetration d_w (sold blue lines). In the meltwater penetration case, the glacier either has a floating (blue shaded region) or grounded (unshaded) terminus. Dotted blue lines corresponds to a terminal condition that violates (2.5) and the grey shaded region corresponds to the nonphysical scenario where $h_c < d_w$.

[Nick et al., 2010], the parameter d_w is rarely known but the model is most sensitive to this parameter. Hydrological models coupled with observations could infer/predict this parameter, but not without considerable effort. For now, we proceed as if d_w is given.

Given a fixed d_w, the meltwater parameterized calving law becomes nonphysical for sufficiently shallow beds. The terminal thickness cannot be smaller than the depth water has penetrated, however, figure 2-3 shows for small enough b_c the calving law predicts $h_c < d_w$. We assume this is a physical breakdown of a mathematical model;

we require that $d_w \leq h_c$.

This model does not capture calving-driven retreat: if the calving front is thicker than the critical thickness $h_c = -\phi b_c$, calving ceases. The calving cliff instability—described by [Pollard et al., 2015]—is triggered if a glacier calves when it terminates *above* the critical thickness h_c. The resulting terminal thickness is still larger than h_c and, therefore, another calving event is triggered, leading to collapse. The glacier is not subject to the calving cliff instability if the leading portion of the glacier *below* the critical thickness h_c is instantaneously removed. In the meltwater penetration calving law, a stably calving glacier satisfies

$$\frac{d(d_s + d_b)}{dh_c} < 1. \tag{2.5}$$

Only the larger fraction ϕ in figure 2-2 satisfies this requirement and figure 2-3 shows that the unstable branch violates the assumption that this portion is grounded. The calving cliff instability is a result of shear failure, which predicts calving when the terminal thickness is above a critical thickness [Bassis and Jacobs, 2013, Bassis and Walker, 2011, Ma et al., 2017, Pollard et al., 2015]. Glaciers that terminate above the critical shear failure thickness cannot persist and, therefore, we only consider the stably calving glaciers.

Terminus speed To describe the evolution of the terminus location, we need an expression for how quickly the terminus is advancing/retreating. In the calving-at-flotation and meltwater calving cases, the calving law prescribes a terminal thickness h_c. We assume calving happens on much faster time scales than the glacier's evolution and, therefore, if $h(x_c) < h_c$, then the glacier immediately calves such that $h(x_c) = h_c$. Assuming the terminus thickness satisfies $h(x_c(t), t) = -b_c(t)\phi(t)$ (noting that $b_c(t) = b(x_c(t))$), and differentiating the terminal thickness condition with respect to time implicitly defines the terminal speed $\dot{x}_c$:

$$\frac{d}{dt}\left\{h(x_c(t), t) = -b_c(t)\phi(t)\right\} \tag{2.6}$$

which implies

$$h_x|_{x=x_c(t)}\,\dot{x}_c + h_t|_{x=x_c(t)} = -\,b_x|_{x=x_c(t)}\,\dot{x}_c\phi(t) - b_c(t)\phi_t. \tag{2.7}$$

We manipulate this to express the grounding line speed $\dot{x}_c$ explicitly as a function of the terminus geometry and the input/output flux at x_c. Letting $\zeta(t) = -d_w(t)b_c^{-1}(t)$ and either

$$\phi = r^{-1} \quad \text{or} \quad \phi(\zeta) = \begin{cases} 2r^{-1}\zeta & \text{if } \zeta < \frac{1}{2} \text{ (floating)} \\ \eta + \sqrt{\eta^2 - r^{-1}} & \text{if } \zeta \geq \frac{1}{2} \text{ (grounded)} \end{cases} \tag{2.8}$$

where $\eta = 1 + (r^{-1} - 1)\zeta$. The derivative of ϕ with respect to time is $\phi_t = \phi_\zeta \zeta_t$, noting that either

$$\phi_\zeta = 0 \quad \text{or} \quad \phi_\zeta = \begin{cases} 2r^{-1} & \text{if } \zeta < \frac{1}{2} \text{ (floating)} \\ r^{-1} - 1 + \frac{1+(r^{-1}-1)\zeta}{\sqrt{(1+(r^{-1}-1)\zeta)^2 - r^{-1}}} & \text{if } \zeta \geq \frac{1}{2} \text{ (grounded)} \end{cases} \tag{2.9}$$

and

$$\zeta_t = d_w b_c^{-2}\, b_x|_{x=x_c(t)}\,\dot{x}_c - \dot{d}_w b_c^{-1}. \tag{2.10}$$

Therefore, either

$$\phi_t = 0 \quad \text{or} \quad \phi_t = \phi_\zeta \left(d_w b_c^{-2}\, b_x|_{x=x_c}\,\dot{x}_c - \dot{d}_w b_c^{-1} \right). \tag{2.11}$$

The evolution of the calving boundary is defined by

$$\dot{x}_c = \left(w^{-1}\,(uwh)_x|_{x=x_c} - a \right) \left(h_x|_{x=x_c} + b_x|_{x=x_c} r^{-1} \right)^{-1} \tag{2.12a}$$

or

$$\dot{x}_c = \left(w^{-1}\,(uwh)_x|_{x=x_c} + (1-\theta)\mu - a + \dot{d}_w\phi_\zeta\right)\left(h_x|_{x=x_c} + \left(\phi + \phi_\zeta d_w b_c^{-1}\right) b_x|_{x=x_c}\right)^{-1} \tag{2.12b}$$

for the calving-at-flotation and meltwater calving terminal conditions, respectively.

The terminus speed (2.12) assumes that the terminus thickness h equals the critical thickness prescribed by the calving law $h_c = -b_c\phi$. In general,

$$\dot{x}_c = \begin{cases} u(x_c(t),t) & \text{if } h > h_c \\ s(t) & \text{if } h = h_c \\ -\infty & \text{if } h < h_c, \end{cases} \tag{2.13}$$

where $s(t)$ is the terminal speed derived in (2.12). This determines the final boundary condition in the calving-limited case.

Melt-limited terminus

In the extreme case of a melt-terminating glacier (figure 2-1b), a finite glacier length results when the cumulative mass loss due to subshelf melt exceeds the cumulative surface mass balance. The terminus is the location where the ice flux (hence, the ice thickness) vanishes, which also requires the extensional stress to vanish. Since (2.1a) is singular when $h = 0$, we impose

$$uh \to 0 \quad \text{and} \quad h|u_x|^{\frac{1}{n}-1}u_x \to 0 \quad \text{as} \quad x \to x_c(t). \tag{2.14}$$

Warm ocean currents beneath the ice shelf may cause significant subshelf melt. If d_w is small and the melt rate μ is large, the glacier may have a small calving front. However, different calving laws lead to qualitatively different upstream dynamics. Analyzing the melt limited extreme provides insight about how subshelf melt affects glacier dynamics.

2.3.2 Nondimensionalization

Considering the dimensionless form of the flowline model (2.1) reduces the number of free parameters and suggests approximations based on the relatively small size of some dimensionless quantities. Internal stresses—the extensional stress in (2.1a)—are negligible in the ice sheet's interior [Schoof, 2007b, Schoof, 2007a].

We first assign relative scales to each dimensional quantity. Assume we know length, accumulation, and width scales

$$x = [x]x^*, \quad a = [a]a^*, \quad \mu = [a]\mu^*, \quad \text{and} \quad w = [w]w^*, \tag{2.15a}$$

such that the asterisks indicate nondimensional quantities. We choose scales for the velocity, time, and thickness

$$u = [u]u^*, \quad t = [t]t^*, \quad \text{and} \quad h = [h]h^* \tag{2.15b}$$

such that, assuming $b = [h]b^*$ and $d_w = [h]\lambda$,

$$C_l[w]^{-\frac{1}{n}-1}[u]^{\frac{1}{n}} = \rho_i g[h][x]^{-1}, \quad [u][h] = [a][x], \quad \text{and} \quad [u] = [x][t]^{-1}, \tag{2.15c}$$

which balances the lateral stress with the driving stress in (2.1a) and the advection with both the mass balance and time derivative terms in (2.1b). The solution to this system implies

$$[u] = [w]\left(\frac{[a]g\rho_i}{C_l}\right)^{\frac{n}{n+1}}, \quad [t] = [x][u]^{-1}, \quad \text{and} \quad [h] = [a][t]. \tag{2.15d}$$

Typical magnitudes of the scaling parameters, given constant parameter values from table 2.1, are given in table 2.2.

We derive the nondimensional form of (2.1) by scaling (2.1a) by the lateral stress'

	Symbol	Meaning	Value & units
Prescribed scaling coefficients	$[x]$	Length scale	250×10^3 m
	$[w]$	Width scale	5×10^3 m
	$[a]$	Mass balance scale	2.5 m/a
Computed scaling coefficients	$[u]$	Velocity scale	≈ 750 m/a
	$[t]$	Time scale	≈ 333 a
	$[h]$	Thickness scale	≈ 833 m
Nondimensional parameters	ε	Extensional stress coefficient	$\approx 10^{-3}$
	γ	Lateral stress coefficient	≈ 3
	λ	Filling level	—
	r	Ice-to-ocean density ratio	≈ 0.9

Table 2.2: Dimensional scaling constants and nondimensional parameters. Again, we note that the numbers given here are examples given 'reasonable' physical scenarios. Our purpose to the approximate magnitude of each quantity, not to claim that these values are universally accepted truths.

dimensional coefficient. We define nondimensional parameters

$$\varepsilon = \tfrac{1}{2} A^{-\frac{1}{n}} C_l^{-1} \left(\frac{[w]}{[x]}\right)^{\frac{1}{n}+1}, \quad \gamma = \left(\frac{[w]^{\frac{1}{n}+1}[u]^{m-\frac{1}{n}}}{[h]}\right) \frac{C_b}{C_l}, \quad \lambda = \frac{d_w}{[h]}, \quad \text{and} \quad r = \frac{\rho_i}{\rho_w}, \tag{2.16a}$$

which represent the a scaled width-to-length aspect ratio, a scaled ratio of the basal-to-lateral frictional coefficients, a rescaled meltwater penetration parameter (we refer to λ as the 'filling level') and the ice-to-ocean density ratio, respectively. Typical values are given in table 2.2. If there is no slip at the lateral boundary, (2.16) models the width-averaged thickness and velocity by setting

$$C_l = (n+2)^{\frac{1}{n}} 2^{\frac{n+1}{n}} A^{-\frac{1}{n}}. \tag{2.16b}$$

However, lateral slip justifies choosing smaller values [Raymond, 1996]. Substituting (2.16b) into (2.16a) implies that

$$\varepsilon = (n+2)^{-\frac{1}{n}} 2^{\frac{-2n-1}{n}} \left(\frac{[w]}{[x]}\right)^{\frac{1}{n}+1}. \tag{2.16c}$$

Therefore, ε measures the lateral aspect ratio $[w][x]^{-1}$. A narrow channel ensures

that ε is small and that creep flow due to the extensional stress is negligible. In the interior, the flow is well approximated by a model where gravitational stress is balanced by friction along the bed and side walls. Dropping the asterisks immediately, the nondimensional flowline model is

$$4\varepsilon \left(h \left| u_x \right|^{\frac{1}{n}-1} u_x \right)_x - w^{-\frac{1}{n}-1} h \left| u \right|^{\frac{1}{n}-1} u - \theta\gamma \left| u \right|^{m-1} u = \left(1 - (1-\theta)\, r\right) h \left(h_x + \theta b_x\right) \quad (2.16d)$$

$$w h_t + (wuh)_x = w \left(a - (1-\theta)\, \mu\right). \quad (2.16e)$$

with boundary conditions

$$4\varepsilon h \left| u_x \right|^{\frac{1}{n}-1} u_x = \tfrac{(1-r)h^2}{2} + \tfrac{\theta\left(r^2h^2 - b^2\right)}{2r} \quad \text{and} \quad \dot{x}_c = \begin{cases} u(x_c(t), t) & \text{if } h > h_c \\ s(t) & \text{if } h = h_c \\ -\infty & \text{if } h < h_c, \end{cases} \quad (2.16f)$$

at $x = x_c(t)$ with $s(t)$ defined as the right hand side of (2.12) (appropriately nondimensionalizing) and

$$\phi = r^{-1} \quad \text{or} \quad \phi(b_c; \lambda) = \begin{cases} -2r^{-1}\lambda b_c^{-1} & \text{if } -\lambda b_c^{-1} < \frac{1}{2} \text{ (floating)} \\ \eta + \sqrt{\eta^2 - r^{-1}} & \text{else (grounded),} \end{cases} \quad (2.16g)$$

where $\eta(b_c; \lambda) = 1 - (r^{-1} - 1)\, \lambda b_c^{-1}$, in the calving-limited case, or

$$uh \to 0 \quad \text{and} \quad h|u_x|^{\frac{1}{n}-1} u_x \to 0 \quad \text{as} \quad x \to x_c(t), \quad (2.16h)$$

in the melt-limited case. Henceforth, we will analyze the nondimensional flowline

model defined by (2.16).

2.4 Numerical methods

We implement numeric solvers to approximate solutions to the full flowline model (2.16). We use a shooting method to compute steady state solution and finite element methods evolve the system in time.

2.4.1 Steady state solutions of the flowline model

We use a shooting method to solve the steady state version of the model defined by the stress and mass conservation equations along a flowline (2.16). For numerical convenience, let L be the maximum value of x—the end of the numerical domain, perhaps at the continental shelf. We begin by transforming the flowline model into a first order ordinary differential equation (ODE), noting that, in steady state, the flowline model is a second order ODE. Define dependent variables

$$q = uhw, \quad \sigma = |u_x|^{\frac{1}{n}-1} u_x, \quad \text{and} \quad \chi = \tfrac{x}{L}, \tag{2.17}$$

noting that $|\sigma|^{n-1}\sigma = u_x$, and independent variable η

$$x_\eta = q, \quad \text{implying} \quad \tfrac{\partial}{\partial x} = L^{-1}\tfrac{\partial}{\partial \chi} \quad \text{and} \quad \tfrac{\partial}{\partial x} = q^{-1}\tfrac{\partial}{\partial \eta}, \tag{2.18}$$

therefore, $\chi_\eta = L^{-1}q$. From the (steady state) mass balance equation $q_x = w(a - (1-\theta)\mu)$, which, after transforming into η coordinates, implies $q_\eta = qw(a-(1-\theta)\mu)$. Applying the product rule to $(uwh)_x = w(a-(1-\theta)\mu)$ derives

$$h_\eta = wh(a-(1-\theta)\mu) - L^{-1}qw^{-1}hw_\chi - h^2 w\,|\sigma|^{n-1}\,\sigma. \tag{2.19}$$

Substituting (2.17) into the stress balance equation

$$4\varepsilon\left(h\left|u_x\right|^{\frac{1}{n}-1}u_x\right)_x - w^{-\frac{1}{n}-1}h\left|u\right|^{\frac{1}{n}-1}u - \gamma\theta\left|u\right|^{m-1}u = (1-(1-\theta)r)\,h\,(h_x+\theta b_x) \tag{2.20}$$

implies σ_η is a function of h_η, which is defined in (2.19),

$$\begin{aligned}\sigma_\eta = &-\sigma h^{-1}h_\eta + (4\varepsilon)^{-1}w^{-\frac{2}{n}-1}h^{-\frac{1}{n}}\left|q\right|^{\frac{1}{n}+1} \\ &+ (4\varepsilon)^{-1}\gamma\theta w^{-m}h^{-m-1}\left|q\right|^{m+1} + (4\varepsilon)^{-1}(1-(1-\theta)r)\left(h_\eta + \theta L^{-1}qb_\chi\right). \end{aligned} \tag{2.21}$$

The first order ODE is

$$\begin{aligned}
h_\eta &= wh(a-(1-\theta)\mu) - L^{-1}qw^{-1}hw_\chi - h^2w\left|\sigma\right|^{n-1}\sigma && \text{(2.22a)}\\
\sigma_\eta &= -w(a-(1-\theta)\mu)\sigma + L^{-1}qw^{-1}w_\chi\sigma + hw\left|\sigma\right|^{n+1} + (4\varepsilon)^{-1}w^{-\frac{2}{n}-1}h^{-\frac{1}{n}}\left|q\right|^{\frac{1}{n}+1} \\
&\quad + (4\varepsilon)^{-1}\gamma\theta w^{-m}h^{-m-1}\left|q\right|^{m+1} + (4\varepsilon)^{-1}(1-(1-\theta)r)\,(wh(a-(1-\theta)\mu) \\
&\quad -L^{-1}qw^{-1}hw_\chi - h^2w\left|\sigma\right|^{n-1}\sigma) + L^{-1}(4\varepsilon)^{-1}\theta qb_\chi && \text{(2.22b)}\\
q_\eta &= qw(a-(1-\theta)\mu) && \text{(2.22c)}\\
\chi_\eta &= L^{-1}q && \text{(2.22d)}
\end{aligned}$$

At the ice divide $\chi = 0$ and we assume $q = 0$. The ice divide, therefore, corresponds to the fixed point

$$(h,\sigma,q,\chi) = \left(h_0,\ \left(\frac{a}{h_0}\right)^{\frac{1}{n}},\ 0,\ 0\right) \tag{2.23}$$

that occurs when $\eta \to -\infty$. We use a shooting method to determine h_0 such that the unique orbit emerging from the fixed point at $x = 0$ satisfies the boundary conditions at the terminus somewhere along its trajectory. The boundary conditions are defined by either the calving- or melt- limited terminus.

2.4.2 Calving conditions

The calving-limited boundary conditions require that

$$4\varepsilon h\sigma = \frac{(1-r)h^2}{2} + \frac{\theta(r^2h^2-b^2)}{2r} \quad \text{and} \quad h = -b\phi(-\lambda b^{-1}) \tag{2.24}$$

at some finite η. We need to find the unique orbit emerging from (2.23) that satisfies these conditions. Let $(h',\, \sigma',\, q',\, \chi')$ be the point on the orbit emerging from the fixed point that satisfies the hydrostatic balance condition (first expression in (2.24)), which we compute given h_0. We find the root of the function

$$G_{\text{calving}}(h_0) = h' + b(\chi')\phi(-\lambda b^{-1}(\chi')). \tag{2.25}$$

The orbit corresponding to h_0^* such that $F(h_0^*) = 0$ satisfies the steady state stress and mass conservation equations and the calving-limited boundary conditions.

We see from plots of (2.25)—figures 2-4 and 2-5—that steady states are *not* unique. The roots in figures 2-4 and 2-5 are candidate steady state thickness h_0, which defines the fixed point (2.23). There is a unique orbit emerging from this fixed point and evolving according (2.22) that defines the glacier's profile.

2.4.3 Melting conditions

The melt-limited boundary conditions require that

$$q \to 0 \quad \text{and} \quad h\sigma \to 0 \quad \text{as} \quad \eta \to \infty, \tag{2.26}$$

which is another fixed point of the dynamical system (2.22). This is not obvious because the system becomes singular as $h \to 0$. We, therefore, switch independent variables using $u = \frac{q}{wh}$. Setting $\theta = 0$ since the glacier must be floating in a neighborhood around the terminal location and assuming $q > 0$, (2.22) is transformed

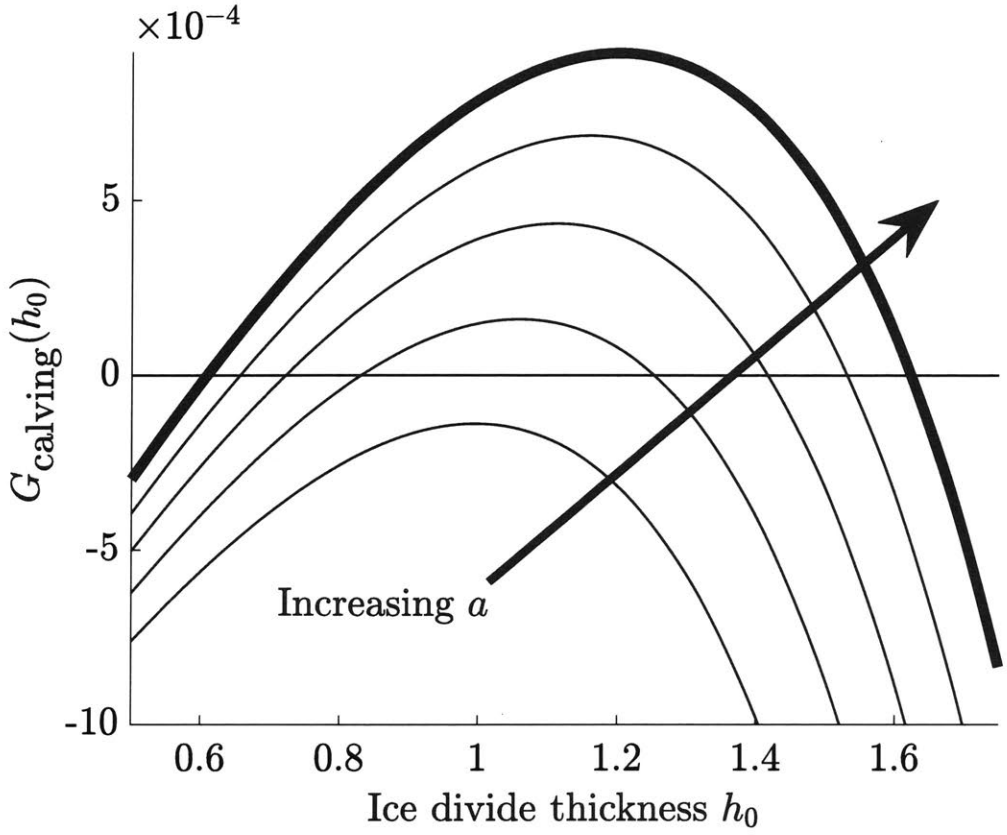

Figure 2-4: The steady state root function (2.25) given nondimensional parameters $\varepsilon = 10^{-3}$, $\gamma = 3$, and $r = 0.9$, constant width $w = 10$, bed rock elevation $b(x) = 0.08x^2 - 0.35$, and varied values of $a \in [0.6, 1]$. The bold lines correspond to $a = 0.8$. Roots of this function define steady state solutions to the flowline model (2.16) in the calving-at-flotation case. Increasing the surface mass balance a increases the net flux into the glacier, which must be balanced by increasing the outgoing flux at the terminus or decreasing the glacier length.

into

$$u_\eta = L^{-1}quw^{-1}w_\chi + q\,|\sigma|^{n-1}\,\sigma \tag{2.27a}$$

$$\begin{aligned}\sigma_\eta = {} & -w(a-\mu)\sigma + L^{-1}qw^{-1}w_\chi\sigma + q\,|\sigma|^{n+1}\,u^{-1} + (4\varepsilon)^{-1}\,w^{-\frac{1}{n}-1}u^{\frac{1}{n}}q \\ & + (4\varepsilon)^{-1}\,(1-r)\,\big(qu^{-1}(a-\mu) \\ & - L^{-1}q^2w^{-2}w_\chi u^{-1} - q^2u^{-2}w^{-1}\,|\sigma|^{n-1}\,\sigma\big)\end{aligned} \tag{2.27b}$$

$$q_\eta = qw(a-\mu) \tag{2.27c}$$

$$\chi_\eta = L^{-1}q. \tag{2.27d}$$

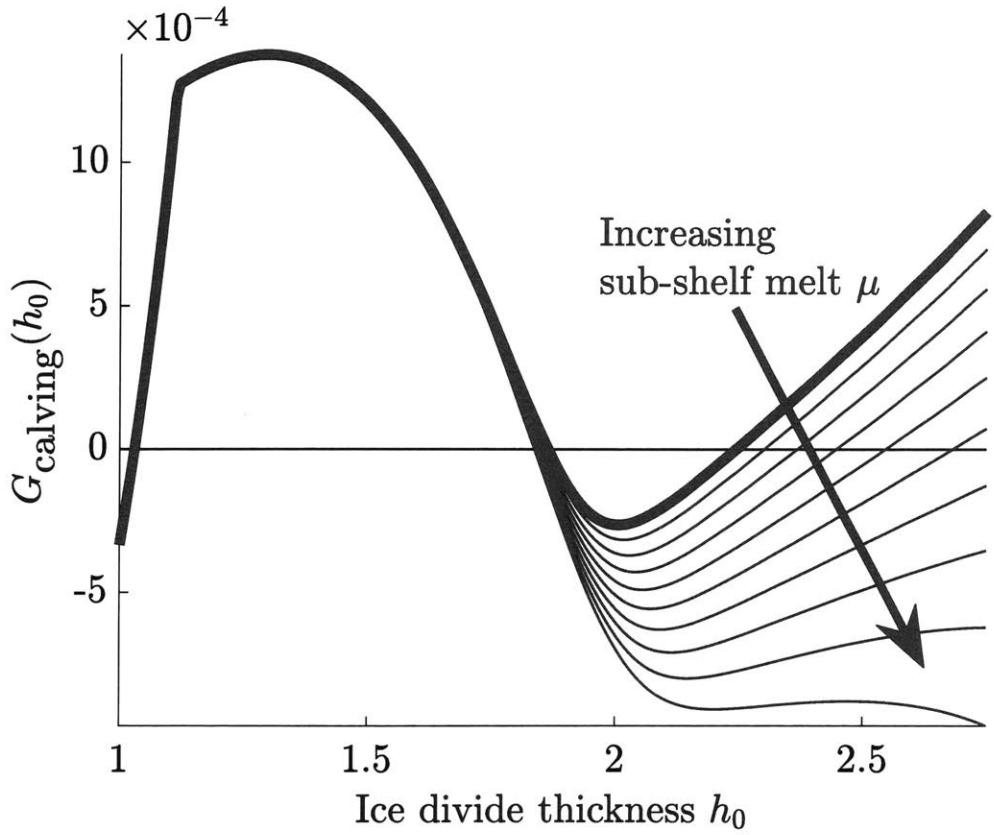

Figure 2-5: The steady state root function (2.25) given nondimensional parameters $\varepsilon = 10^{-3}$, $\gamma = 3$, and $r = 0.9$, constant width $w = 10$, and bed rock elevation $b(x) = 0.08x^2 - 0.35$. We have set $a = 1$, $\lambda = 0.11$, and varied $\mu \in [0, 2]$. The bold line corresponds to $\mu = 0$.

We require that $q \to 0$ and $u \to u_c > 0$ as $\eta \to \infty$, which corresponds to the fixed point $(u_c, 0, 0, \chi_c)$. The Jacobian of (2.27) evaluated at $(u_c, 0, 0, \chi_c)$ is

$$\mathbf{J} = \begin{bmatrix} 0 & 0 & \frac{u_c w_\chi}{Lw} & 0 \\ 0 & -w(a-\mu) & (4\varepsilon)^{-1}\left(w^{-\frac{1}{n}-1}u_c^{\frac{1}{n}} + (1-r)\left(\frac{a-\mu}{u_c}\right)\right) & 0 \\ 0 & 0 & w(a-\mu) & 0 \\ 0 & 0 & L^{-1} & 0 \end{bmatrix} \tag{2.28}$$

The Eigenvalues of this matrix are $\{0, 0, w(a-\mu), -w(a-\mu)\}$, implying that there is a stable, unstable, and two center manifolds. The Eigenvectors corresponding to

the two center manifolds are

$$\begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \tag{2.29}$$

which have no dynamics—varying u and χ corresponds to other fixed points. Since $w(a-\mu) < 0$, the Eigenvectors corresponding to the stable and unstable manifolds are

$$\begin{bmatrix} \frac{u_c w_\chi}{w} \\ L\left(8\varepsilon\right)^{-1}\left(w^{-\frac{1}{n}-1}u_c^{\frac{1}{n}} + (1-r)\frac{a-\mu}{u_c}\right) \\ Lw(a-\mu) \\ 1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \tag{2.30}$$

respectively. The unique orbit-of-interest corresponds to the stable manifold entering this fixed point. Using this Jacobian within a linear expansion gives the local form

$$\begin{aligned} u &\sim u_c + \frac{u_c w_\chi}{Lw}\Delta q && \text{(2.31a)} \\ \sigma &\sim -w(a-\mu)\Delta\sigma + (4\varepsilon)^{-1}\left(w^{-\frac{1}{n}-1}u_c^{\frac{1}{n}} + (1-r)\left(\frac{a-\mu}{u_c}\right)\right)\Delta q && \text{(2.31b)} \\ q &\sim w(a-\mu)\Delta q && \text{(2.31c)} \\ \chi &\sim \chi_c + L^{-1}\Delta q, && \text{(2.31d)} \end{aligned}$$

which means that $\Delta q = x - x_c$. Setting $\Delta\sigma = 0$ to avoid perturbing in the direction

of the unstable manifold, the local form becomes

$$u \sim u_c + \frac{u_c w_x}{w}(x - x_c) \tag{2.32a}$$

$$\sigma \sim (4\varepsilon)^{-1}\left(w^{-\frac{1}{n}-1}u_c^{\frac{1}{n}} + (1-r)\left(\frac{a-\mu}{u_c}\right)\right)(x - x_c) \tag{2.32b}$$

$$q \sim w(a-\mu)(x - x_c) \tag{2.32c}$$

We need to find the orbit that connects the fixed points $(h_0,\ (ah_0^{-1})^{\frac{1}{n}},\ 0,\ 0)$ as $\eta \to -\infty$ and $(u_c,\ 0,\ 0,\ \chi_c)$ as $\eta \to \infty$, where u and h are related via $q = uhw$. Given a guess for the ice divide thickness h_0, we define the critical point $(h_g,\ \sigma_g,\ q_g,\ \chi_g)$ by integrating (2.22) forward with respect to η from $(h_0,\ (ah_0^{-1})^{\frac{1}{n}},\ 0,\ 0)$ until $h = -r^{-1}b(x)$, which corresponds to the grounding line location. The terminus location χ_c is determined by mass balance—mass can only be gained/lost via surface mass balance or subshelf melt and, therefore, the terminal location satisfies

$$\int_0^{\chi_c} aw(\chi)\,d\chi = \int_{\chi_g}^{\chi_c} \mu w(\chi)\,d\chi. \tag{2.33}$$

Given a guess for the terminal velocity u_c, we define a second critical point $(h_g',\ \sigma_g',\ q_g',\ \chi_g')$ by integrating (2.27) backward with respect to η from $(u_c,\ 0,\ 0,\ \chi_c)$ until $\sigma = \sigma_g$. Continuity requires that $(h_g' = \frac{q_g'}{w(\chi_g')u_g'},\ \sigma_g',\ q_g',\ \chi_g') = (h_g,\ \sigma_g,\ q_g,\ \chi_g)$. We have already guaranteed that $\sigma_g' = \sigma_g$. We have two free parameters (h_0 and u_c) that are constrained by requiring $h_g' = h_g$ and $\chi_g' = \chi_g$—the residual functions $F(h_0, u_c) = \chi_g - \chi_g'$ and $G(h_0, u_c) = h_g - h_g'$ are shown in figures 2-6 and 2-7. Since $\chi_g' = \chi_g$ and we have chosen χ_c to satisfy (2.33), then $q_g = q_g'$.

2.4.4 Transient solutions of the flowline model

We compute the evolution of the velocity and thickness fields using a Galerkin method and an implicit time discretization. The location of the calving boundary, the right boundary, is time-dependent; the spatial domain is $[0,\ x_c(t)]$. We, therefore, define (ν, τ)-coordinates:

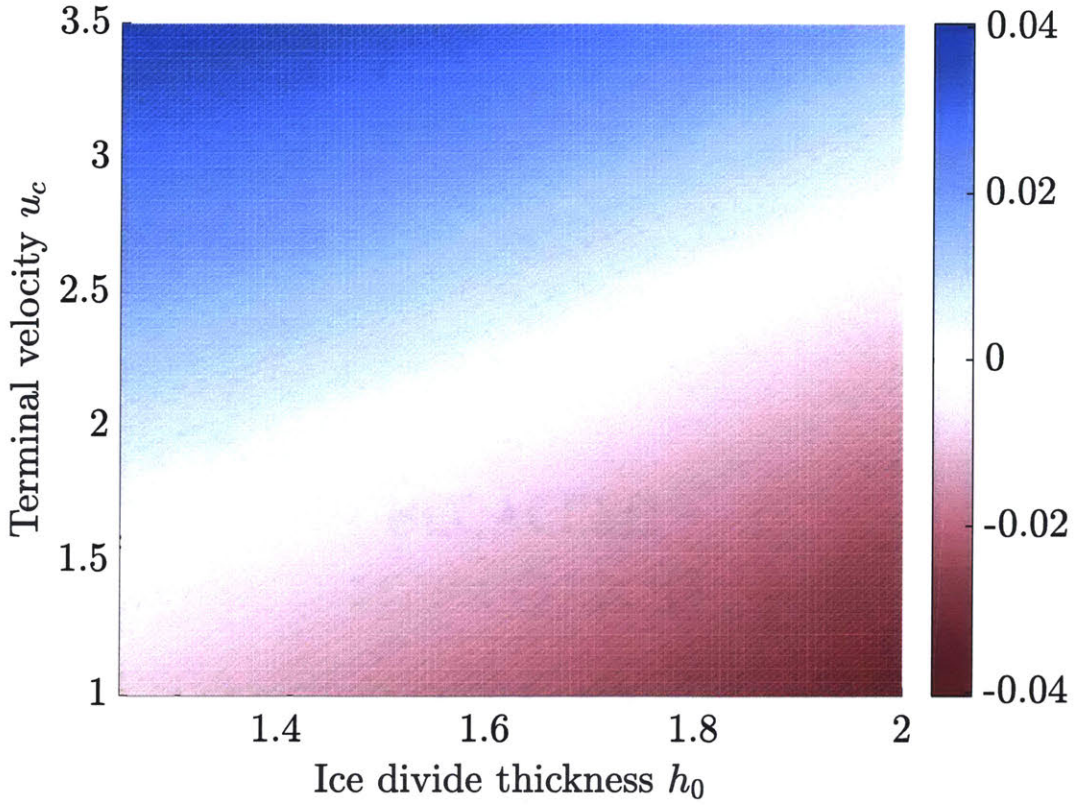

Figure 2-6: A plot of $F(h_0, u_c) = \chi_g - \chi_g'$, which is the difference is the predicted grounding line location χ_g computed by integrating from each end of the glacier, the ice divide and the terminus. In steady state, integrating from either end of the glacier should predict the same grounding line location (e.g., $F(h_0, u_c) = 0$). In this example, we have chosen $b(x) = 0.08x^2 - 0.35$, $w = 10$, $a = 0.8$, and $\mu = 8$. The unique steady state orbit corresponds to the point where both functions are zero.

$$\nu(x,t) = \tfrac{x}{x_c(t)} \quad \text{and} \quad \tau(t) = t, \tag{2.34}$$

which implies

$$\tfrac{\partial}{\partial x} = x_c^{-1}\tfrac{\partial}{\partial \nu} \quad \text{and} \quad \tfrac{\partial}{\partial t} = \tfrac{\partial}{\partial \tau} - \tfrac{\nu \dot{x}_c}{x_c}\tfrac{\partial}{\partial \nu}, \tag{2.35}$$

so the domain (grounded and floating) is [0, 1]. We derive an expression that describes how the terminus location evolves and use the stress and mass balance equations to evolve the velocity and thickness fields.

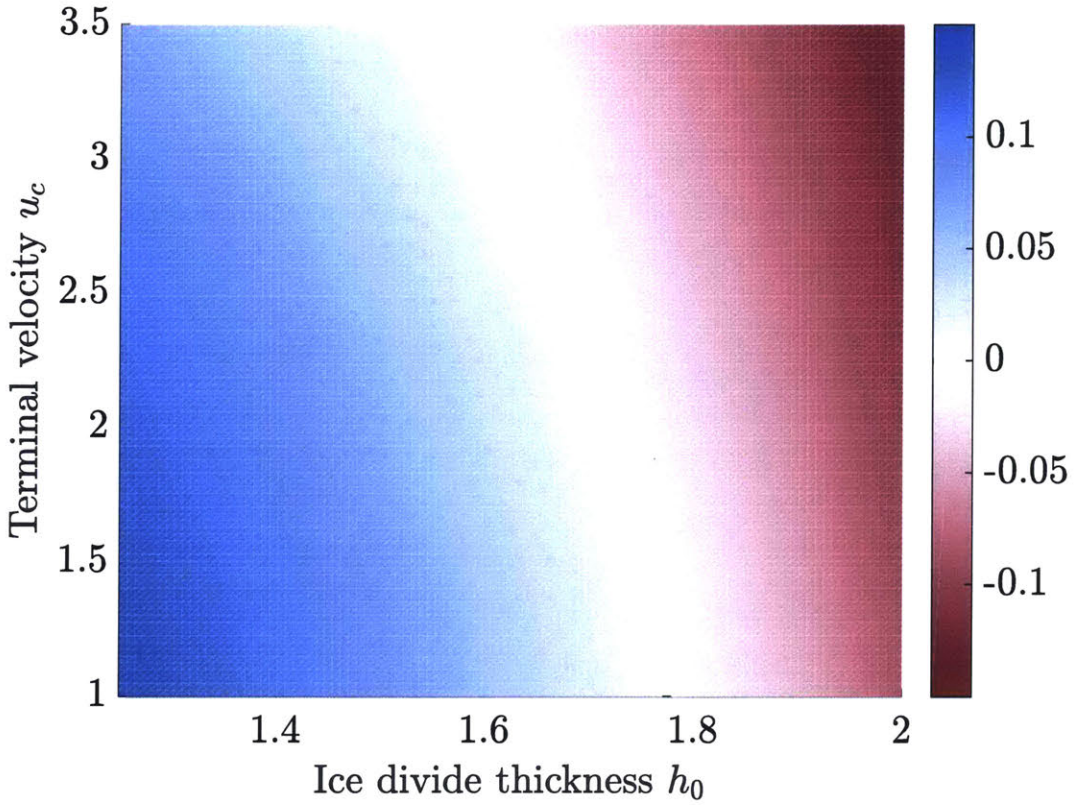

Figure 2-7: A plot of $G(h_0, u_c) = h_g - h'_g$, which quantifies the size of the discontinuity in thickness at the grounding line afterintegrating from each end of the glacier, the ice divide and the terminus. At steady state locations, the thickness should be continuous (e.g., $G(h_0, u_c) = 0$). In this example, we have chosen $b(x) = 0.08x^2 - 0.35$, $w = 10$, $a = 0.8$, and $\mu = 8$. The unique steady state orbit corresponds to the point where both functions are zero.

Terminus speed

We need to regularize (2.13) so that the function governing how the terminus evolves in time is continuous. Recall, the terminus speed $\dot{x}_c = \frac{dx_c}{dt}$ is determined by the calving law. Let $h_c(t) = -\phi b(x_c(t)$ denote the critical thickness imposed by the calving law. The terminus cannot evolve faster than the terminal velocity; therefore, we assume that if the terminal thickness is above this critical thickness that the terminus advances at the same rate as the terminal velocity. Conversely, if the terminal thickness is below the critical thickness h_c, the ice instantaneously calves. Finally, we assume that the glacier terminus at steady state if $h(x_c(t), t) = h_c(t)$. This leads to

the first step of our relaxation:

$$\dot{x}_c = \begin{cases} u(x_c(t), t) & \text{if } h > h_c \\ 0 & \text{if } h = h_c \\ -\infty & \text{if } h < h_c. \end{cases} \tag{2.36}$$

We have removed the term that depends explictly on the terminus speed $s(t)$ from (2.12). We note that $s(t)$ is derived from mass balance and $s(t) = 0$ in steady state. Any parameter change that forces the glacier out of steady state will cause a change in the terminal thickness and lead to terminus advance/retreat.

We new relax this expression for terminal evolution so that the right hand side is continuous. Let $\eta = h(x(t), t) - h_c(t)$ by the difference between the terminal thickness and the critical terminal thickness. Set

$$\dot{x}_c = \alpha(\eta) u \tag{2.37a}$$

such that $\alpha(\eta) \to 1$ as $\eta \to \infty$, $\alpha(\eta) \to -\infty$ as $\eta \to -\infty$, and $\alpha(\eta) = 0$ if $\eta = 0$. We prescribe

$$\alpha(\eta) = 1 - \exp(-\alpha_1 \eta), \tag{2.37b}$$

where α_1 is a prescribed algorithmic constant (see figure 2-8).

Stress and mass balance equations: weak form

We numerically compute solutions to the stress and mass balance equations using a symmetric interior penalty (SIP) discontinuous Galerkin method [Arnold et al., 2002, Cockburn et al., 2000]. In (ν, τ)-coordinates, the stress and mass balance equations

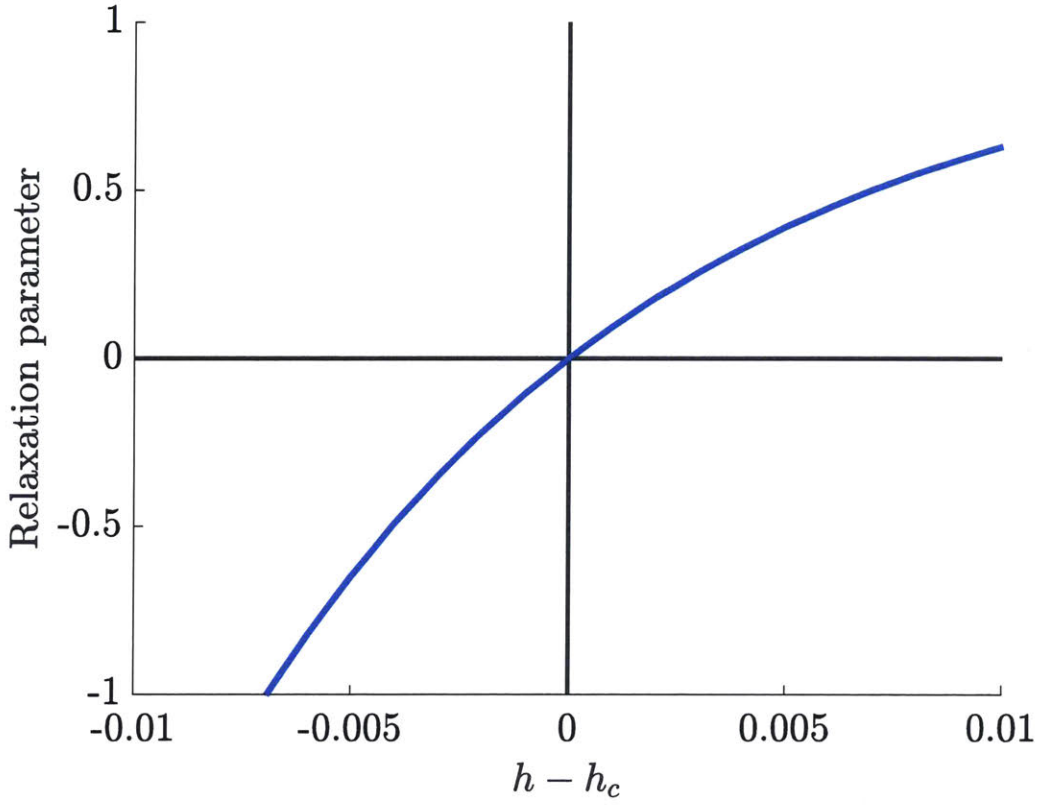

Figure 2-8: The relaxation parameter (2.37b) as a function of the mismatch between the terminal thickness and the critical thickness prescribed by the calving law $h(x(t), t) - h_c(t)$ with $\alpha_1 = 100$.

are

$$4\varepsilon x_c^{-\frac{1}{n}-1}\left(h\left|u_\nu\right|^{\frac{1}{n}-1}u_\nu\right)_\nu - w^{-\frac{1}{n}-1}h\left|u\right|^{\frac{1}{n}-1}u - \gamma\theta\left|u\right|^{m-1}u$$
$$= \left(1-(1-\theta)r\right)x_c^{-1}h\left(h_\nu + \theta b_\nu\right) \quad (2.38a)$$

$$wh_\tau - \frac{\nu\dot{x}_c}{x_c}wh_\nu + x_c^{-1}(uwh)_\nu = w(a-(1-\theta)\mu) \quad (2.38b)$$

with boundary conditions (assuming a calving-limited terminus, given a form of ϕ)

$$u = 0 \quad \text{at} \quad \nu = 0 \quad (2.38c)$$

and

$$h_c = -b_c\phi \quad \text{and} \quad 4\varepsilon x_c^{-\frac{1}{n}} h \left|u_\nu\right|^{\frac{1}{n}-1} u_\nu = \tfrac{(1-r)h^2}{2} + \tfrac{\theta(r^2h^2-b^2)}{2r} \quad \text{at} \quad \nu = 1, \tag{2.38d}$$

where $b_c(t) = b(x_c(t))$. Defining the viscous parameter

$$\zeta = 4\varepsilon x_c^{-\frac{1}{n}-1} \tag{2.39}$$

and the diffusive and advective numerical fluxes

$$\begin{aligned}
F(u, u_\nu, h) &= \zeta h \left|u_\nu\right|^{\frac{1}{n}-1} u_\nu, && \text{(2.40a)}\\
f(u, h) &= (1-(1-\theta)r)x_c^{-1} h(h_\nu + \theta b_\nu) && \text{(2.40b)}\\
&\quad + w^{-\frac{1}{n}-1} h \left|u\right|^{\frac{1}{n}-1} u + \theta\gamma \left|u\right|^{m-1} u,\\
G(u, h) &= x_c^{-1}\left(uwh - \dot{x}_c \nu w h\right), \text{ and} && \text{(2.40c)}\\
g(h) &= w(a-(1-\theta)\mu) - \frac{\dot{x}_c}{x_c}\left(\nu w\right)_\nu h && \text{(2.40d)}
\end{aligned}$$

allow us to rewrite the strong form equations as

$$\begin{aligned}
\left(F(u, u_\nu, h)\right)_\nu &= f(u, h) && \text{(2.41a)}\\
w h_\tau + \left(G(u, h)\right)_\nu &= g(h) && \text{(2.41b)}
\end{aligned}$$

with boundary conditions

$$u = 0 \quad \text{at} \quad \nu = 0 \tag{2.41c}$$

and

$$h = -b\phi \quad \text{and} \quad F(u, u_\nu, h) = \tfrac{(1-r)h^2}{2x_c} + \tfrac{\theta(r^2h^2-b^2)}{2rx_c} \quad \text{at} \quad \nu = 1. \tag{2.41d}$$

We represent unknown functions u and h as discontinuous piecewise polynomials over a triangulation $\Omega_{\Delta\nu}$ of [0, 1]. The triangulation is a set of N_e elements indexed by e_i such that

$$\bigcup_{i=1}^{N_e} \Omega_{e_i} = [0,\ 1] \quad \text{and} \quad \Omega_{e_i} \cap \Omega_{e_j} = \emptyset \tag{2.42}$$

for $i \neq j$. The space of discontinuous piecewise polynomials over the triangulation is

$$S^p_{\Delta\nu}(\Omega_{\Delta\nu}) = \{v|_{\Omega_{e_i}} \in P^p(\Omega_{e_i}), \text{ for all } \Omega_{e_i} \in \Omega_{\Delta\nu}\}, \tag{2.43}$$

where $P^p(\Omega_{e_i})$ is the space of polynomials with degree at most p over element Ω_{e_i}. For shape functions $\varphi^i_j \in P^p(\Omega_{e_i})$ the velocity and thickness are approximated by

$$u(\nu) = \sum_{j=1}^{p+1} u^i_j \varphi^i_j(\nu) \quad \text{and} \quad h(\nu) = \sum_{j=1}^{p+1} h^i_j \varphi^i_j(\nu) \quad \text{for all} \quad \nu \in \Omega_{e_i} \tag{2.44}$$

and, similarly, we choose test functions

$$v(\nu) = \sum_{j=1}^{p+1} v^i_j \varphi^i_j(\nu) \quad \text{and} \quad z^i(\nu) = \sum_{j=1}^{p+1} z^i_j \varphi^i_j(\nu) \quad \text{for all} \quad \nu \in \Omega_{e_i}, \tag{2.45}$$

such that $v(0) = 0$ and $z(1) = 0$. Define the jump operator, weighted jump operator, and the average operators at the element boundaries:

$$[\![v]\!] = v^+ - v^-, \quad [\![v]\!]_{\mathbf{n}} = v^+n^+ - v^-n^-, \quad \text{and} \quad \{v\} = \tfrac{1}{2}(v^+ + v^-), \tag{2.46}$$

respectively—where $\mathbf{n}$ is the outward pointing normal vector. Here, the '$-$' or '$+$' su-

perscript indicates the function's value on element to the left or right of the boundary. Also define the inner products for this triangulation

$$\langle v, u \rangle_{\partial\Omega_{\Delta\nu}} = \sum_{i=1}^{N_e} vu|_{\partial\Omega_{e_i}} \quad \text{and} \quad \langle v, u \rangle_{\Omega_{\Delta\nu}} = \sum_{i=1}^{N_e} \int_{\Omega_{e_i}} vu \, d\nu. \tag{2.47}$$

We derive the weak form by integrating against the test functions v and z

$$\frac{1}{2x_c}\Big((1-(1-\theta)r)h^2 - \theta r^{-1}b^2) \Big)v|_{\nu=1} + \frac{\zeta}{\alpha} \langle [\![v]\!], [\![u]\!] \rangle_{\partial\Omega_{\Delta\nu}} + \langle \{F(u, u_\nu, h)\}, [\![v]\!]_{\mathbf{n}} \rangle_{\partial\Omega}$$
$$- \langle v_\nu, F(u, u_\nu, h) \rangle_{\Omega_{\Delta\nu}} - \langle v, f(u, h) \rangle_{\Omega_{\Delta\nu}} = 0 \tag{2.48a}$$

$$\Big\langle \hat{G}(u^+, h^+, n^+) - \hat{G}(u^-, h^-, n^-), [\![z]\!] \Big\rangle_{\partial\Omega_{\Delta\nu}}$$
$$+ \langle z, wh_\tau \rangle_{\Omega_{\Delta\nu}} - \langle z_\nu, G(u, h) \rangle_{\Omega_{\Delta\nu}} - \langle z, g(h) \rangle_{\Omega_{\Delta\nu}} = 0, \tag{2.48b}$$

and

$$\hat{G}(u, h, n) = \frac{1}{2} \left(G(u, h) \cdot \mathbf{n} + |G(u, h)| \cdot \mathbf{n} \right) \tag{2.49}$$

is the upwind numerical flux. The symmetric interior penalty method adds the term multiplied by $\zeta\alpha^{-1}$ to penalize large discontinuities in the discontinuous piecewise polynomial approximation. Typically, we choose $\alpha = 4 \times 10^{-3}$. We find solutions to (3.5) using a discontinuous Galerkin implementation in the Fenics software package [Alnæs et al., 2015, Logg, 2007, Logg and Wells, 2010, Ølgaard et al., 2008].

Time integration

We integrate with respect to time using an operator splitting method. Each timestep must compute the instantaneous velocities, update the thickness field, and update the terminus location. Since the terminus speed (2.12) depends on solutions to (2.48) and

vice versa, we employ an operator splitting method within each timestep to update u, h, and x_c. Algorithm 1 describes how we evolve the velocity, thickness, and terminus location in time.

Algorithm 1 Numerical time integration with operator splitting of the flowline model (2.16).

1: Set the initial velocity $u^{(0)}$, thickness $h^{(0)}$, and terminus location $x_c^{(0)}$ and the number of timesteps T such that the final time is $T\Delta t$, where Δt is the time step size.
2: **for** $t \leftarrow 1$ to T **do**
3: Compute the terminus speed $\dot{x}_c$ via (2.37) using $u^{(t)}$ and set $x_c' = x_c^{(t)} + \Delta t \dot{x}_c$
4:
5: Compute the thickness and velocity fields at the next timestep by solving (2.48), discretizing the time derivative such that $h_\tau \approx \Delta t^{-1}\left(h^{(t+1)} - h^{(t)}\right)$ and replacing h with $h^{(t+1)}$, u with $u^{(t+1)}$, and x_c with x_c'.
6:
7: Re-compute the terminus speed $\dot{x}_c$ via (2.37) using $u^{(t+1)}$ and set $x_c^{(t+1)} = x_c^{(t)} + \Delta t \dot{x}_c$
8: **end for**

2.5 Solutions to the flowline model

We investigate steady state and transient behavior of glaciers modeled using (2.16) given two idealized geometries: (i) upward and (ii) downward sloping beds—

$$b(x) = -0.25x^6 - 0.15x^2 - 0.2 \quad \text{and} \quad b(x) = 0.08x^2 - 0.25, \tag{2.50}$$

respectively. We use the convention that 'upward' and 'downward' sloping indicates that the bed slopes up/downward toward the ice divide, respectively.

2.5.1 Steady state solution

We compute steady state solutions using both the shooting method from section 2.4 and by solving the weak form with a discontinuous Galkerin method. In the melt-limited case, we only use the shooting method.

Melt-limited extreme

In the melt-limited scenario we compute steady state solution using the shooting method from section 2.4. The unique steady states for the idealized upward and downward sloping beds are shown in figures 2-9 and 2-10.

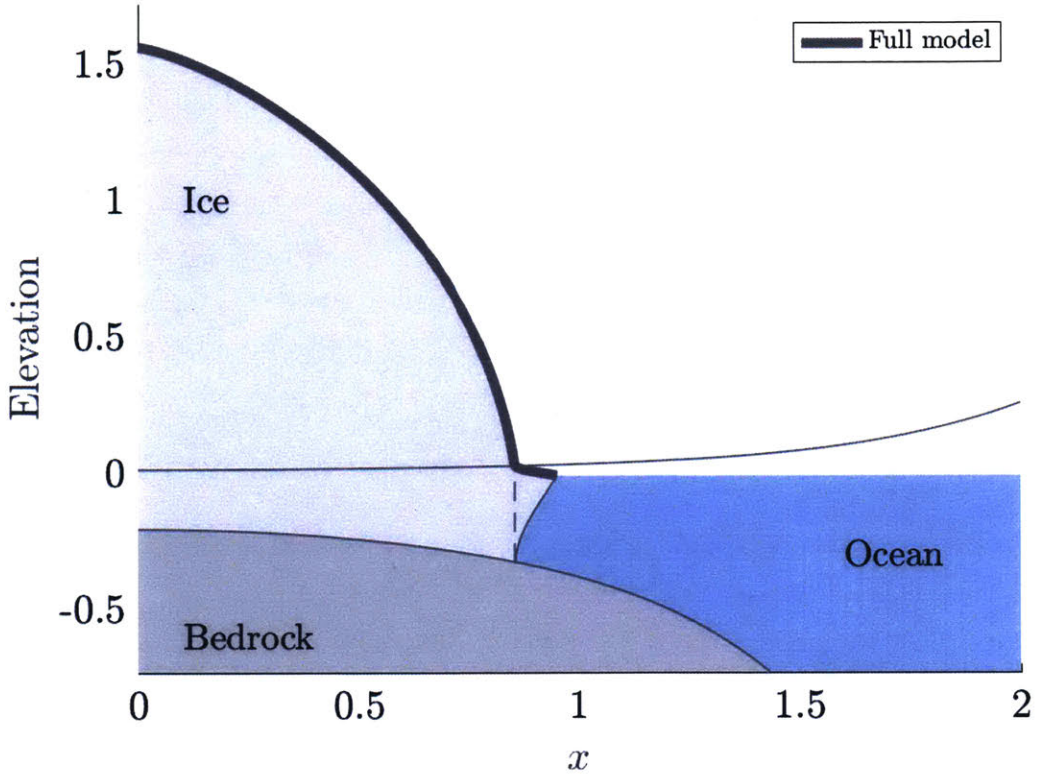

Figure 2-9: Steady state glacier profile for an upward ($b(x) = -0.25x^6 - 0.15x^2 - 0.2$) sloping bed given that mass is only lost via subshelf melt. The grey line is the steady state solution to the full model using the shooting method described in section 2.4. The thin grey line is the flotation thickness $h_g(x) = -r^{-1}b(x)$. We have set $a = 0.8$ and $\mu = 8$.

Calving-at-flotation

In the calving-at-flotation case, we can again compute steady state solutions using the shooting method from section 2.4. However, we can also compute the steady state by solve the weak form of the flowline model with a discontinuous Galerkin method.

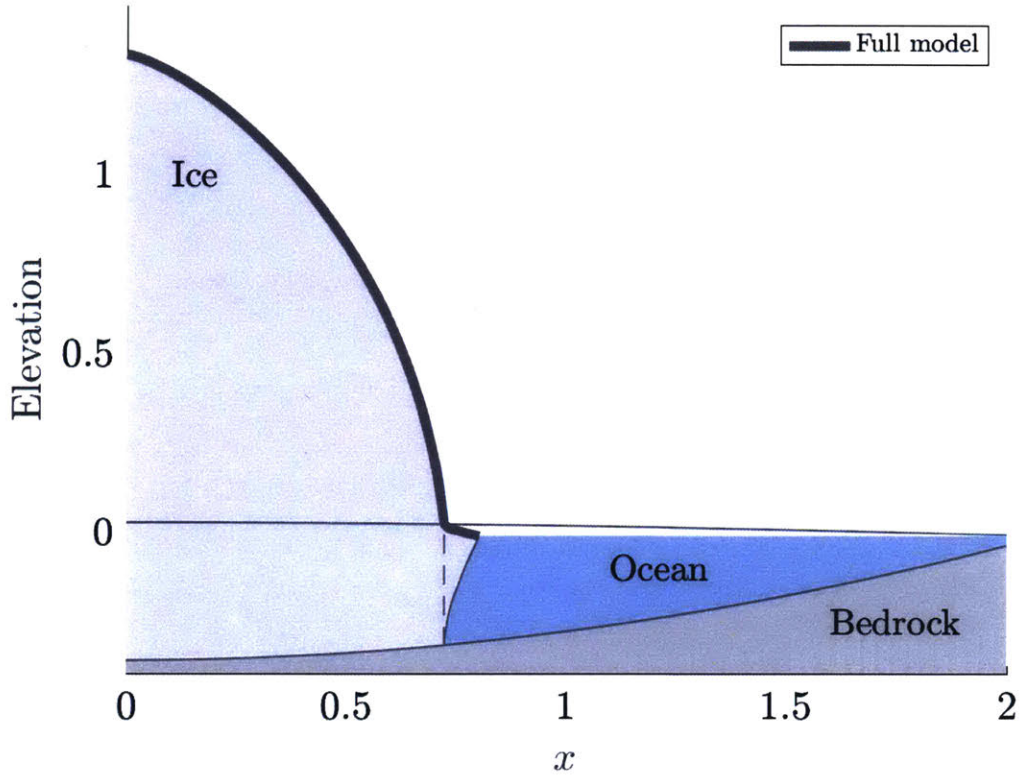

Figure 2-10: Steady state glacier profiles for a downward ($b(x) = 0.08x^2 - 0.25$) sloping bed given that mass is only lost via subshelf melt. The grey line is the steady state solution to the full model using the shooting method described in section 2.4. The thin grey line is the flotation thickness $h_g(x) = -r^{-1}b(x)$. We have set $a = 0.8$ and $\mu = 8$.

Both solutions are shown in figures 2-11 and 2-12. Both methods compute the same solution.

Meltwater calving

In the meltwater calving case we can again compute steady state solutions using both methods: (i) the shooting method from section 2.4 and (ii) solving the weak from with a discontinuous Galerkin method. The results are shown in figures 2-13 and 2-14. The key difference in this case versus the melt-limited extreme and calving-at-flotation is that are multiple steady state for *both* upward and downward sloping beds.

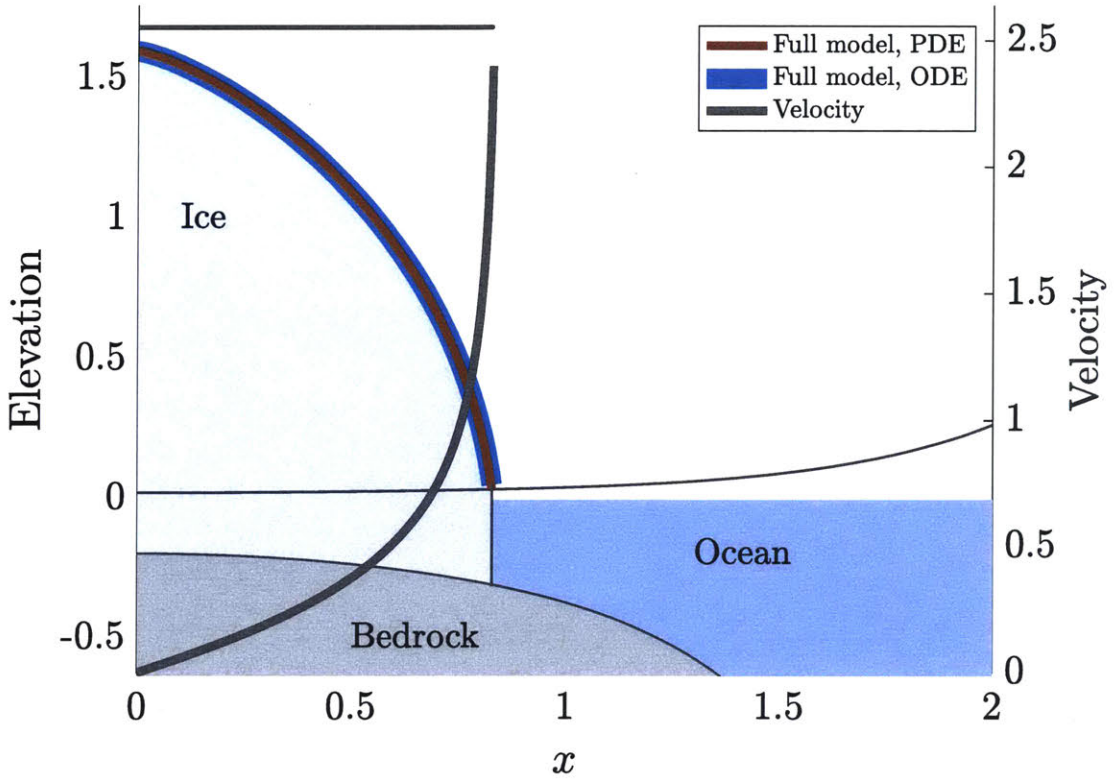

Figure 2-11: Steady state glacier profile for an upward ($b(x) = -0.25x^6 - 0.15x^2 - 0.2$) sloping bed given that the glacier calves at its critical flotation thickness. Solutions are computed using the shooting method from section 2.4 (blue line) and a discontinuous Galkerin method (brown and grey lines). The thin grey line is the flotation thickness $h_g(x) = -r^{-1}b(x)$. Grey dots along the top of each figure are vertexes of the triangulation after adaptive mesh refinement on the full model. We have set $a = 1$ and $w = 10$.

2.5.2 Adaptive mesh refinement

The numerical mesh must have enough elements to resolve the boundary layer, suggesting that the computed solutions may be sensitive to mesh refinement. Define the goal functional

$$M(u, h) = \int_{\nu=0}^{\nu=1} h + \eta u \, d\nu, \tag{2.51}$$

(recall that ν is the transformed coordinate $\nu = x x_c^{-1}$) typically, choosing $\eta = 1$, and an *a posteriori* error estimate

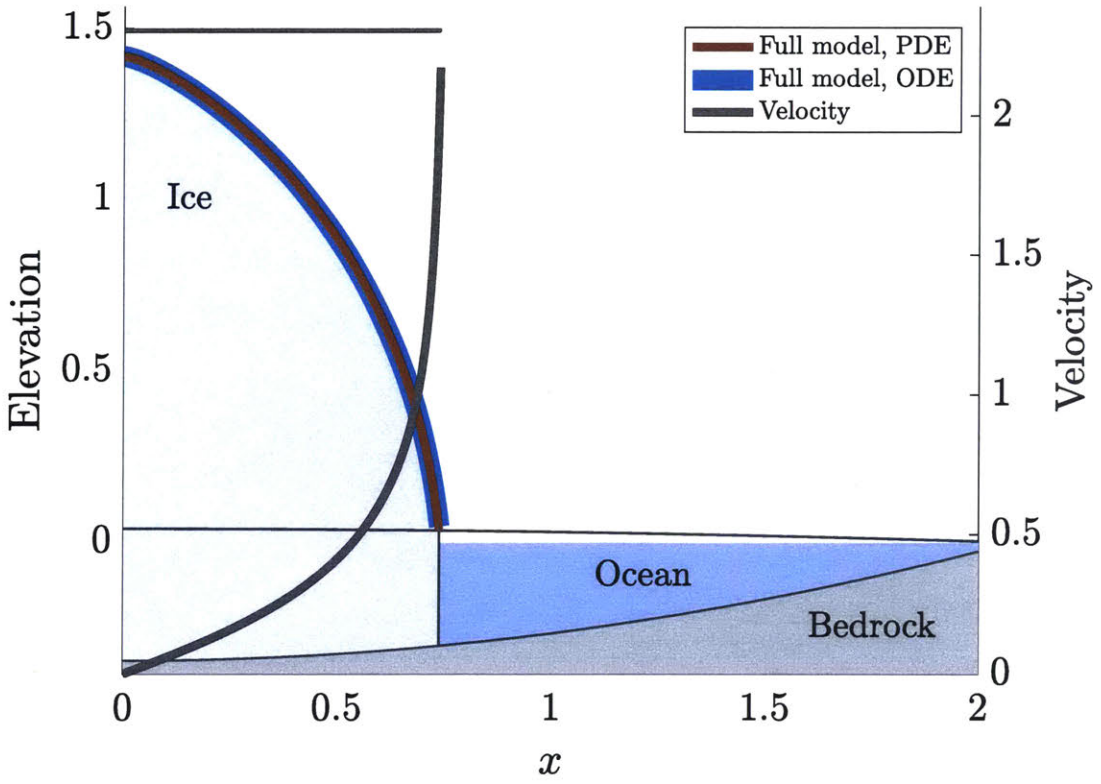

(a)

Figure 2-12: Steady state glacier profile for an upward ($b(x) = 0.08x^2 - 0.25$) sloping bed given that the glacier calves at its critical flotation thickness. Solutions are computed using the shooting method from section 2.4 (blue line) and a discontinuous Galkerin method (brown and grey lines). The thin grey line is the flotation thickness $h_g(x) = -r^{-1}b(x)$. Grey dots along the top of each figure are vertexes of the triangulation after adaptive mesh refinement on the full model. We have set $a = 1$ and $w = 10$.

$$e(\Omega_{\Delta\nu}) = |M(u,h) - M(u(\Omega_{\Delta\nu}), h(\Omega_{\Delta\nu}))|, \tag{2.52}$$

where u and h are the true solution and $u(\Omega_{\Delta\nu})$ are $h(\Omega_{\Delta\nu})$ are computed on the triangulation $\Omega_{\Delta\nu}$. We compute a bound on this error estimate using a dual-weighted residual error estimator [Rognes and Logg, 2013]. We expect the SIP Galerkin method using linear elements to exhibit second order convergence with respect to the number of degrees of freedom under uniform refinement—increasing the number of uniformly sized elements by an order of magnitude should result in a two order of magni-

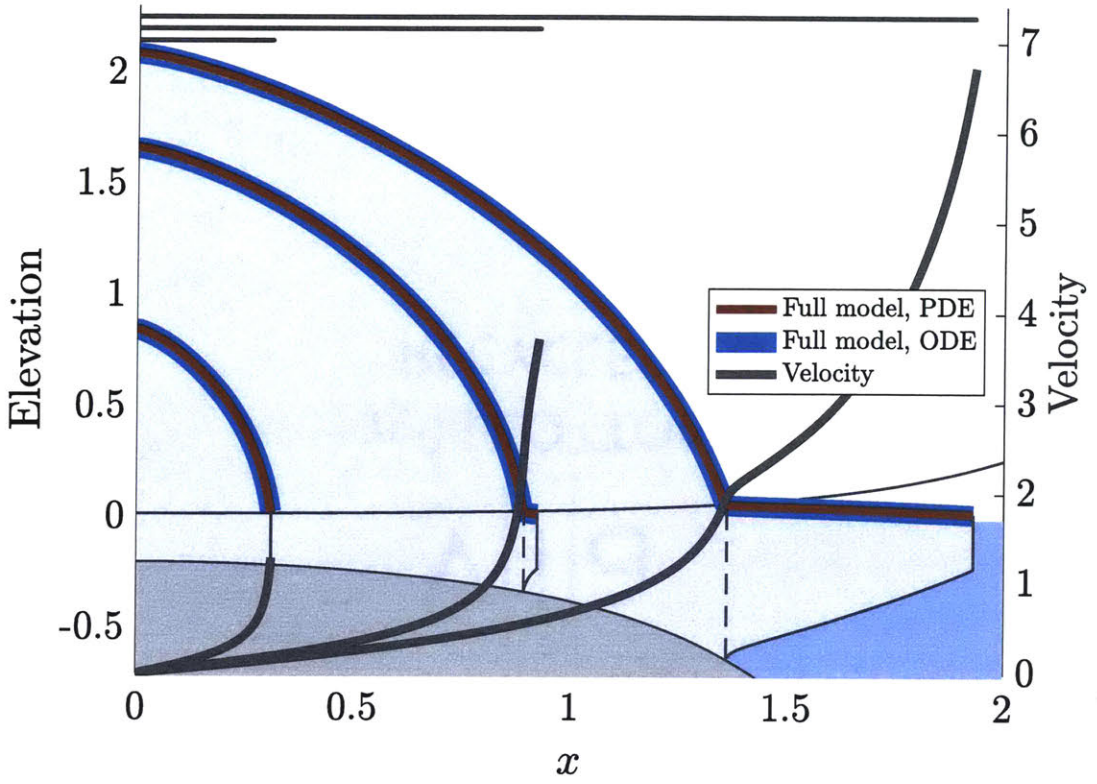

Figure 2-13: Steady state glacier profile for an upward ($b(x) = -0.25x^6 - 0.15x^2 - 0.2$) sloping bed given that the glacier calves at the critical thickness defined by the meltwater calving law. Solutions are computed using the shooting method from section 2.4 (blue line) and a discontinuous Galkerin method (brown and grey lines). The thin grey line is the flotation thickness $h_g(x) = -r^{-1}b(x)$. Grey dots along the top of each figure are vertexes of the triangulation after adaptive mesh refinement on the full model. We have set $a = 1$, $\mu = 0.5$, $\lambda = 0.11$, and $w = 10$.

tude decrease of the error e. The calving-at-flotation case exhibits this behavior—demonstrated by the solid red line in figure 2-15. However, the meltwater calving law exhibits slower convergence—the solid red line in figure 2-16 because of the discontinuity at the grounding line.

2.5.3 Transient solutions

We initialize the model at steady state solutions and force transient behavior by varying surface mass balance a, subshelf melt μ, or the filling level λ.

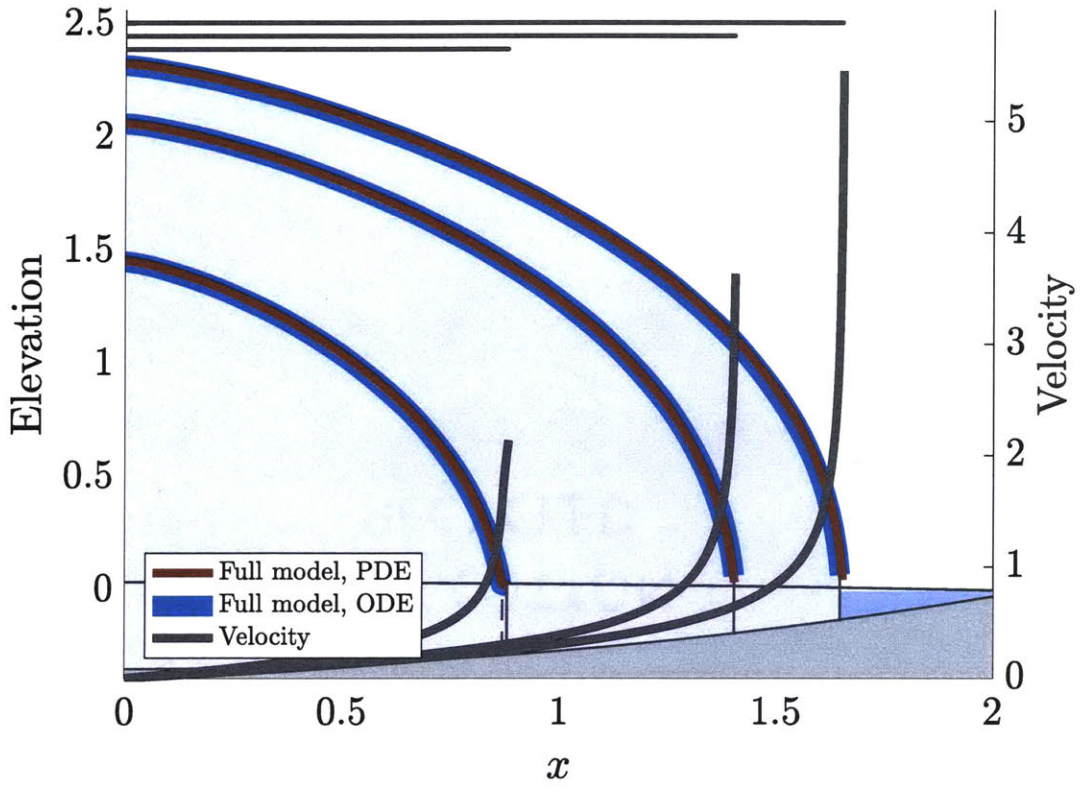

Figure 2-14: Steady state glacier profile for an upward ($b(x) = 0.08x^2 - 0.25$) sloping bed given that the glacier calves at the critical thickness defined by the meltwater calving law. Solutions are computed using the shooting method from section 2.4 (blue line) and a discontinuous Galkerin method (brown and grey lines). The thin grey line is the flotation thickness $h_g(x) = -r^{-1}b(x)$. Grey dots along the top of each figure are vertexes of the triangulation after adaptive mesh refinement on the full model. We have set $a = 0.6$, $\mu = 0.5$, $\lambda = 011$, and $w = 10$.

Glacier stability

We analyze glacier stability by varying the surface mass balance a and comparing four scenarios: (i.a) calving-at-flotation on an upward sloping bed, (i.b) calving-at-flotation on a downward sloping bed, (ii.a) the meltwater calving law on an upward sloping bed, and (ii.b) the meltwater calving law on a downward sloping bed. Let the surface mass balance vary according to

$$a(t) = a_0 - a_1 \sin(2f\pi t), \tag{2.53}$$

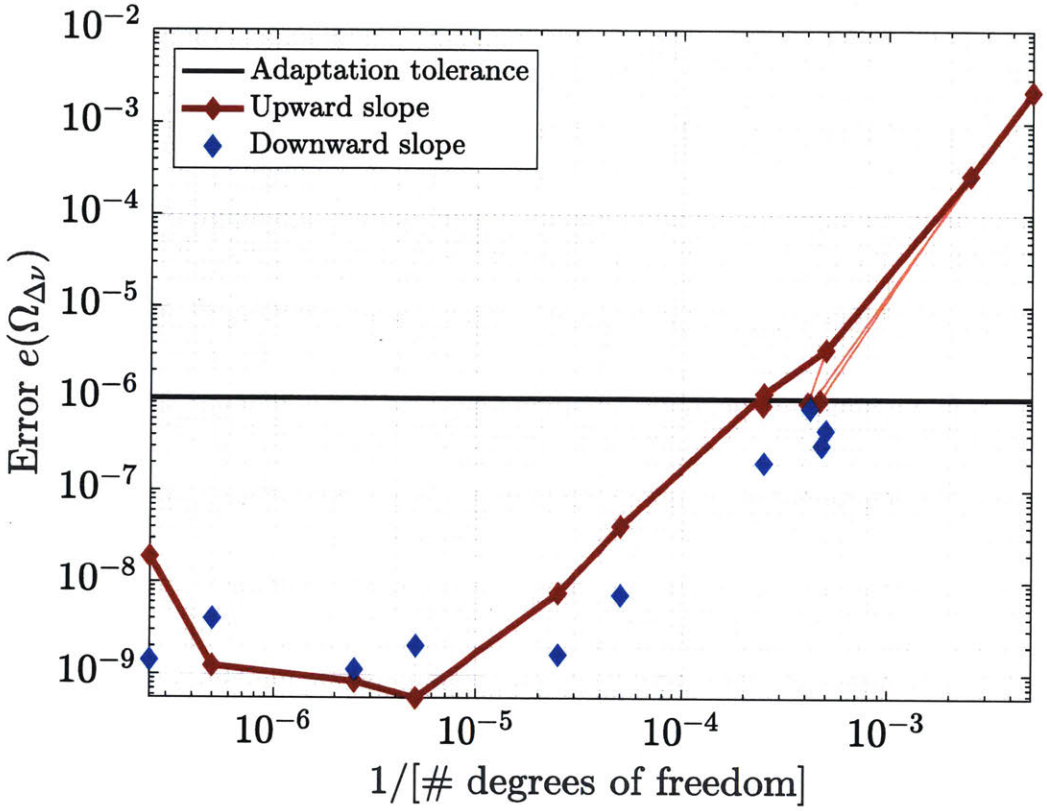

Figure 2-15: We plot the rate that the error estimate (2.52) decays under mesh refinement for the calving-at-flotation calving laws. Each steady state solution corresponds to a glacier profile in figure 2-11 and 2-12. The solid red lines corresponds to uniform mesh refinement (no adaptive refinement) with 5×10^1, 10^2, 5×10^2, 10^3, 5×10^3, 10^4, 5×10^4, 10^5, 5×10^5, or 10^6 elements. The number of degrees of freedom (horizontal axis) is four times the number of elements since we are using discontinuous piecewise linear approximations for u and h. We set an error threshold of 10^{-6} and, beginning from a uniform mesh, let adaptive mesh refinement optimize the mesh.

where f is the frequency, chosen to be relatively slow compared to the evolution of the grounding line, a_0 is the initial surface mass balance, and a_1 is the magnitude of the variation. All simulations begin at steady state setting $a = a_0$. We show numerically that the calving-at-flotation case can only have stable steady states on upward sloping beds. However, the meltwater calving scenario has a stable steady state regardless of bed slope.

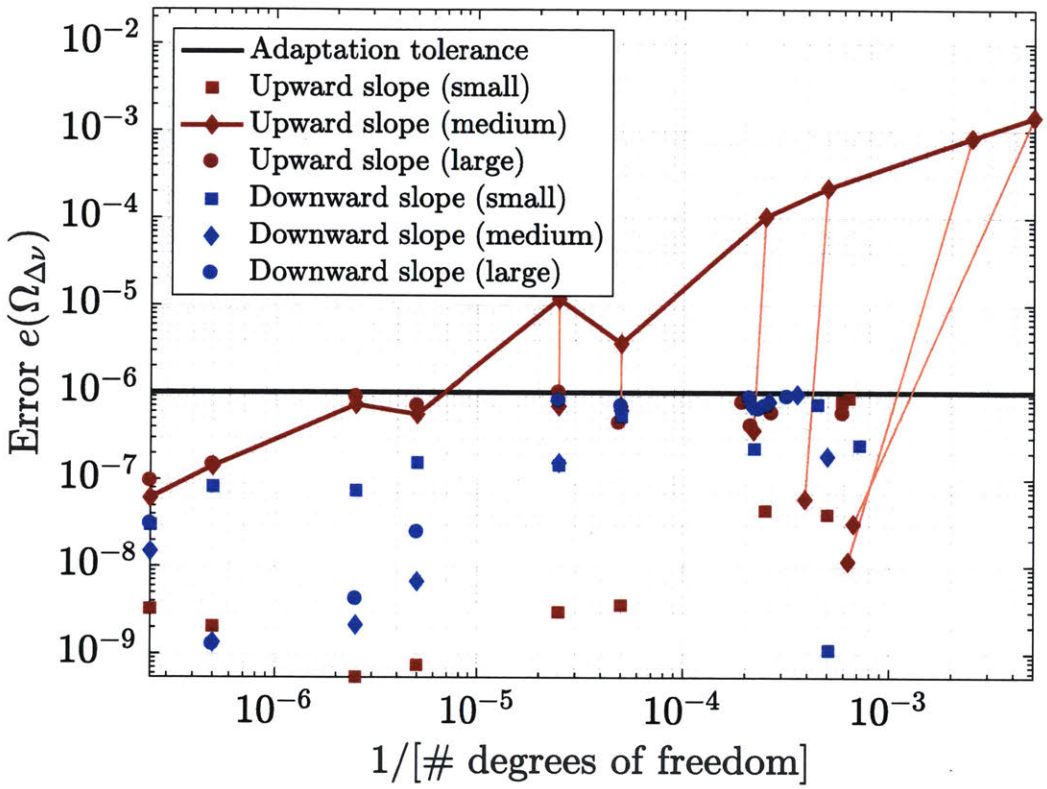

Figure 2-16: We plot the rate that the error estimate (2.52) decays under mesh refinement for the (b) meltwater calving laws. Each steady state solution corresponds to a glacier profile in figures 2-13 and 2-14. The solid red lines corresponds to uniform mesh refinement (no adaptive refinement) with 5×10^1, 10^2, 5×10^2, 10^3, 5×10^3, 10^4, 5×10^4, 10^5, 5×10^5, or 10^6 elements. The number of degrees of freedom (horizontal axis) is four times the number of elements since we are using discontinuous piecewise linear approximations for u and h. We set an error threshold of 10^{-6} and, beginning from a uniform mesh, let adaptive mesh refinement optimize the mesh.

Calving-at-flotation In the calving-at-flotation case, [Schoof, 2007b, Schoof, 2007a] showed that there is a monotonic relationship between bedrock depth at the terminus and the grounding line flux[2]. Therefore, the stability of the steady-state to depend on the bedrock slope. Decreasing surface mass balance a decreases the net in-flux, requiring the *steady state* terminus location to move to shallower water—decreasing the out-flux. Decreasing a over a steady-state glacier means the net in-flux is less than the out-flux. Therefore, glaciers on downward sloping beds retreat *away* from

[2]We obtain a similar result for the melt limited extreme in chapter 3.

the new steady state and glaciers on upward sloping beds retreat *toward* the new steady state. This stability behavior is evident in figures 2-17 and 2-18, which show how the terminus location responds to fluctuations in a. This so-called marine ice instability can be stabilized when the grounding line flux is no longer a monotonic function of bedrock depth.

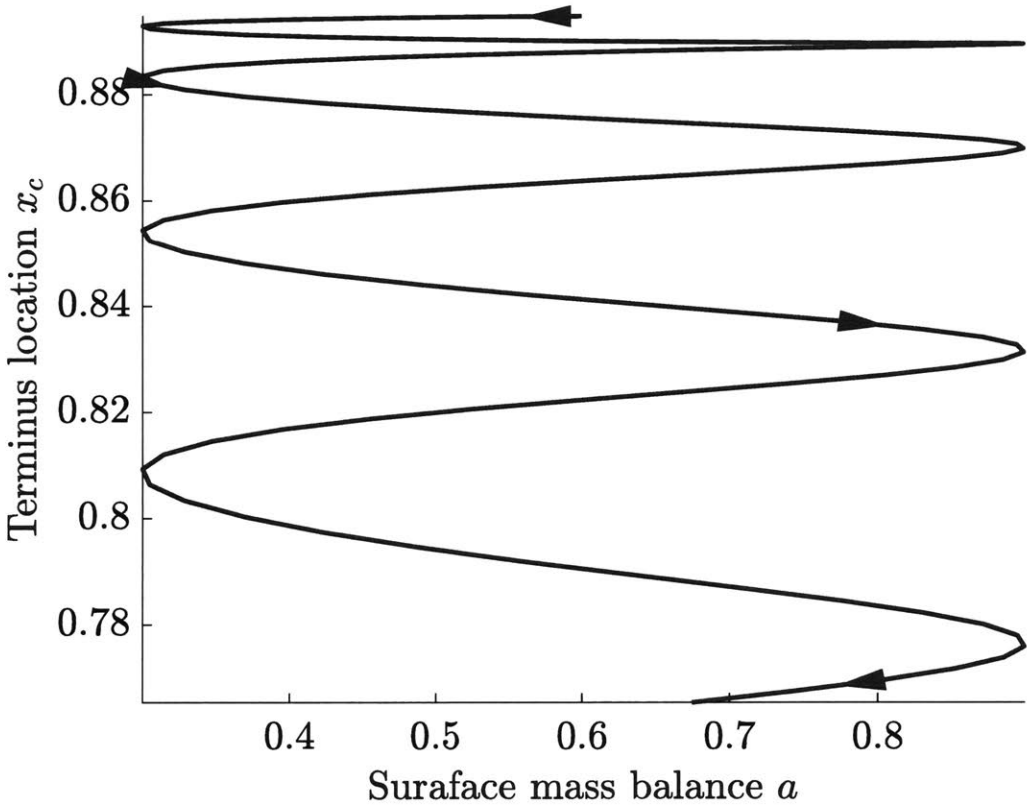

Figure 2-17: Calving-at-flotation on a downward sloping bed: a hysteresis plot of the terminus location against the surface mass balance $a(t) = 1 - 0.4\sin(8\pi t)$ for $t \in [0, 1]$ given a glacier that calves at flotation. We assume constant width $w = 10$ and that this glacier sits on the downward sloping bed defined by (2.50), which is unstable.

Meltwater calving law Given the meltwater calving law, the flux at the grounding line is no longer a monotonic function of bedrock depth, which potentially leads to stable glaciers regardless of the bed's slope. For now, we make this claim based on numerical observations, however, in chapter 3 we will formally show that the flux-to-bedrock relationship is nonmonotonic. The flowline model with the meltwater calving

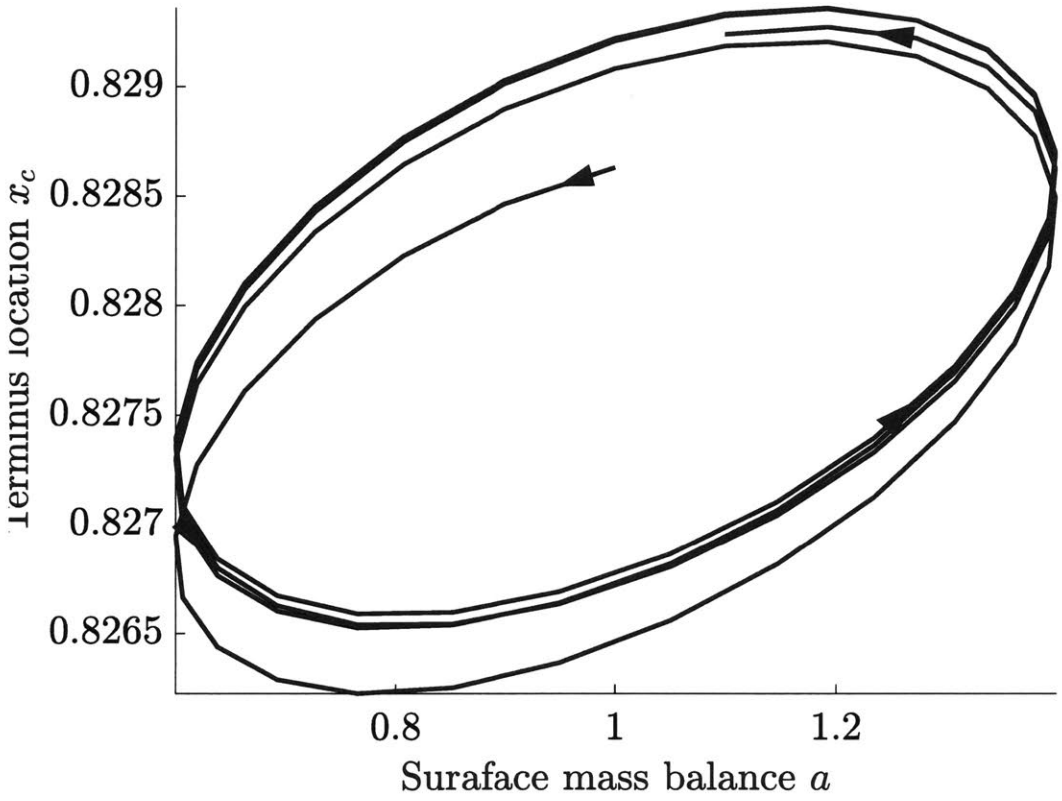

Figure 2-18: Calving-at-flotation on an upward sloping bed: a hysteresis plot of the terminus location against the surface mass balance $a(t) = 1 - 0.4\sin(8\pi t)$ for $t \in [0, 1]$ given a glacier that calves at flotation. We assume constant width $w = 10$ and that this glacier sits on the upward sloping bed defined by (2.50), which is stable.

law potentially has multiple steady states; we initialize the simulations at the stable state, which corresponds to the middle profile in figures 2-13 and 2-14. In both cases, there are two unstable steady states and one stable one. We focus on the stable steady state.

In figures 2-19 and 2-20, we show the stability of glaciers on downward and upward sloping beds given the meltwater calving law. We assume there is no subshelf melt ($\mu = 0$) and drive transient behavior by varying the surface mass balance. Both upward and downward sloping beds have stable steady states; neither glacier retreats or advances unboundedly. Therefore, the meltwater calving law can have a stable steady state *regardless of bed slope.* We also note that in the downward sloping case the stable steady state does *not* have an ice shelf (figure 2-14) and, therefore, ice shelf

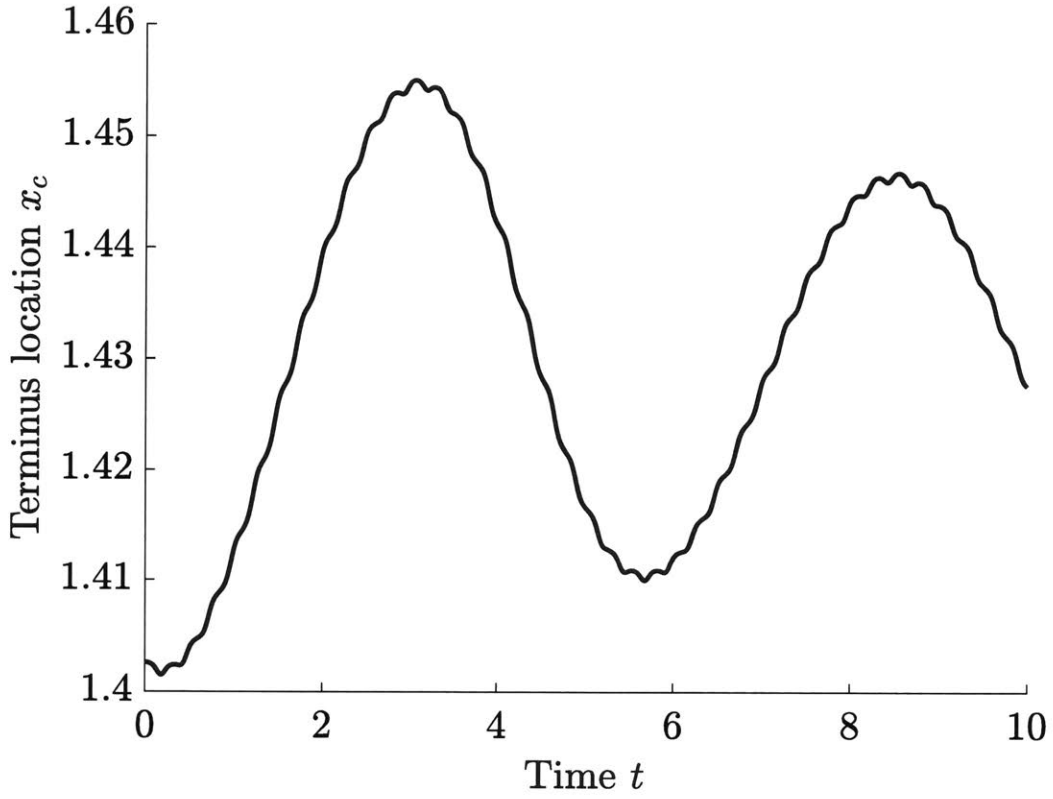

Figure 2-19: The meltwater calving law on a downward sloping bed: a plot of the terminus location evolution given surface mass balance $a(t) = 0.6 - 0.3\sin(8\pi t)$ for $t \in [0, 10]$ and the meltwater calving law. We assume constant width $w = 10$ and that this glacier sits on the downward sloping bed defined by (2.50). We also set $\mu = 0$ and $\lambda = 0.11$.

buttressing is not the stabilizing this system. In chapter 3, we show that in the case of a grounded calving cliff lateral drag can still reduce the extensional stress at the grounding line flux and, thus, stabilize the glacier. Intuitively, a glacier that calves before reaching its critical flotation thickness has a large ice cliff (the height above sea level at the terminus), which is not supported by the ocean. Lateral drag acting the ice cliff reduces the extensional stress and, therefore, also reduces the grounding line flux.

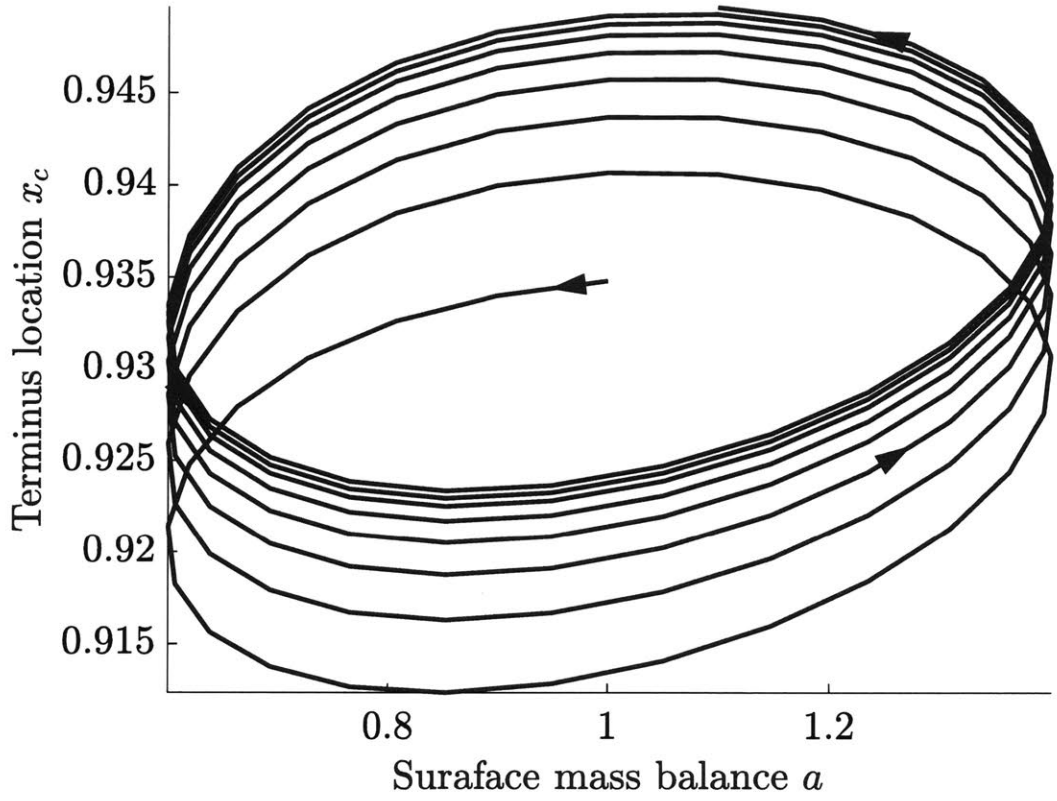

Figure 2-20: The meltwater calving law on an upward sloping bed: a hysteresis plot of the terminus location against the surface mass balance $a(t) = 1 - 0.4\sin(8\pi t)$ for $t \in [0, 2]$ given a glacier the meltwater calving law. We assume constant width $w = 10$ and that this glacier sits on the upward sloping bed defined by (2.50), which is stable. We also set $\mu = 0$ and $\lambda = 0.11$.

Effect of subshelf melt μ and the filling level λ

Increasing the subshelf melt μ decreases ice shelf buttressing and, therefore, causes retreat. Figure 2-21 shows the result of evolving (2.16) forcing transient behavior with

$$\mu(t) = 0.75 - 0.75\cos(8\pi t). \tag{2.54}$$

We initialize the glacier at steady state with no subshelf melt; the system evolves toward the average steady state location.

We similarly compute the glacier's response to changes in the filling level λ. Figure

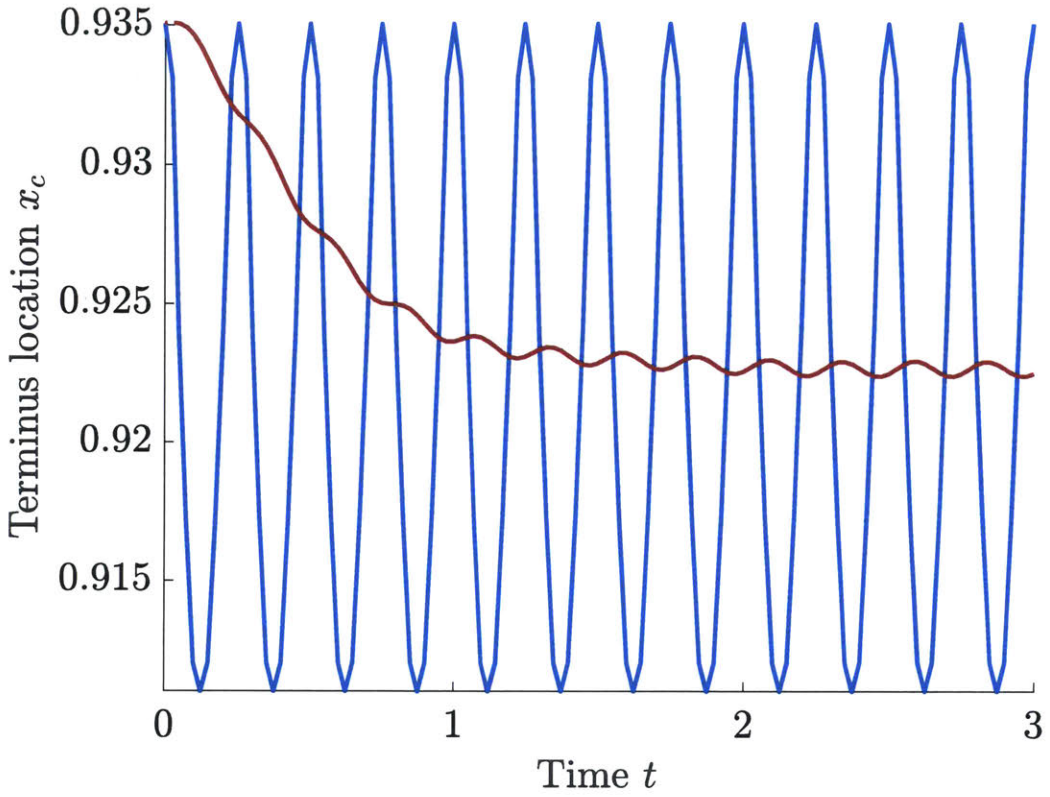

Figure 2-21: The meltwater calving law on an upward sloping bed: a plot of the terminus location evolution given subshelf melt $\mu(t) = 0.75 - 0.75\cos(8\pi t)$ for $t \in [0, 3]$ and the meltwater calving law. The blue line tracks the terminus location of the stable steady state and the red line tracks the terminus location as a result of evolving the flowline model (2.16). We assume constant width $w = 10$ and that this glacier sits on the upward sloping bed defined by (2.50). We also set $a = 1$ and $\lambda = 0.11$.

2-22 shows the result of setting

$$\lambda(t) = 0.12 + 0.01\sin(8\pi t). \tag{2.55}$$

Again, the stable steady state responds as expected.

The glacier's response to changing input parameters also depends on the frequency on which the parameter varies. Here, the characteristic timescale of a glacier is hundreds of years; one unit of time in these simulations corresponds to $\mathcal{O}(10^2)$ years. The surface mass balance, subshelf melt, and filling level can all vary on much faster time scales (e.g., seasonally). Figure 2-23 shows the long-time scale evolution of a

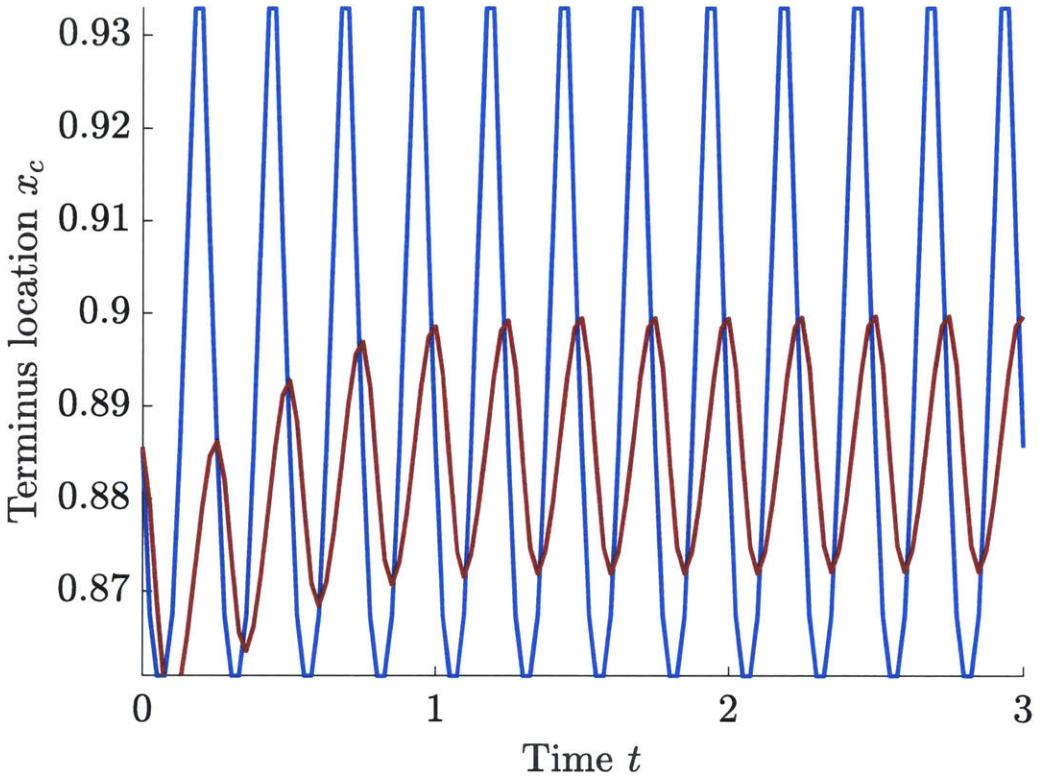

Figure 2-22: The meltwater calving law on an upward sloping bed: a plot of the terminus location evolution given filling level $\lambda(t) = 0.12 - 0.01\sin(8\pi t)$ for $t \in [0, 3]$ and the meltwater calving law. The blue line tracks the terminus location of the stable steady state and the red line tracks the terminus location as a result of evolving the flowline model (2.16). We assume constant width $w = 10$ and that this glacier sits on the upward sloping bed defined by (2.50). We also set $a = 1$ and $\mu = 0$.

glacier whose transient behavior is being forced by seasonal variations in the filling level λ. Figure 2-24 shows the same simulation zoomed-in on the last couple cycles. The glacier evolves toward a stable periodic cycle.

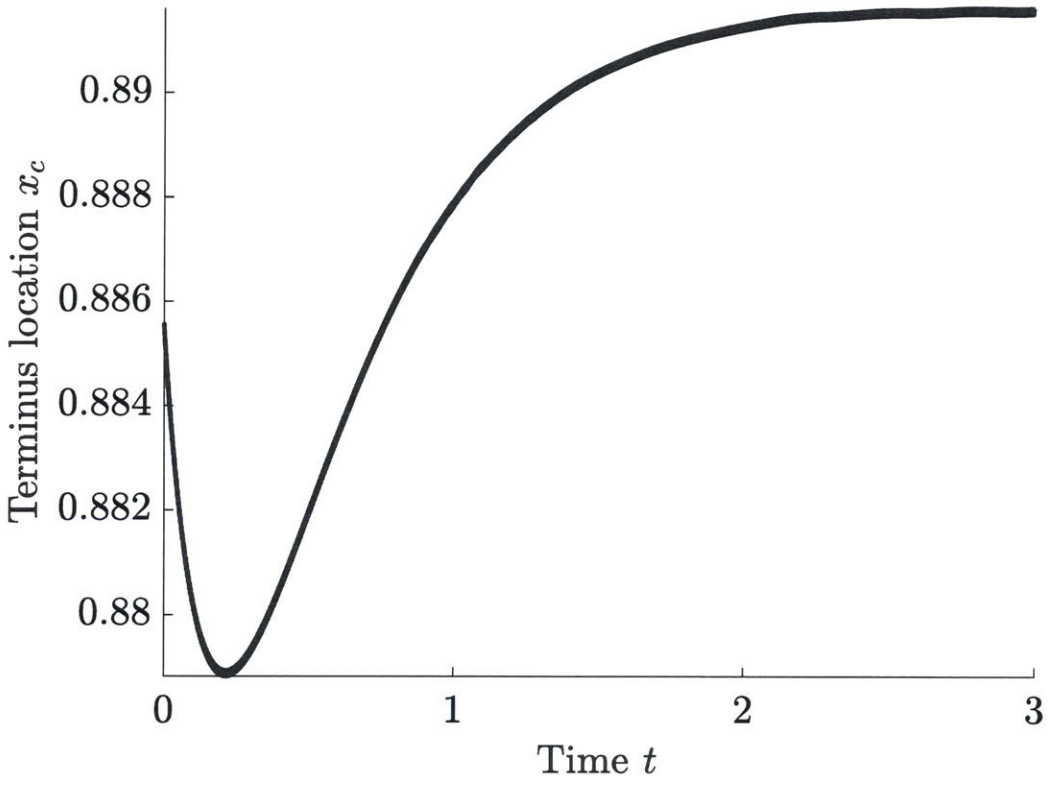

Figure 2-23: The meltwater calving law on an upward sloping bed: a plot of the terminus location evolution given filling level $\lambda(t) = 0.12 - 0.01 \sin(1 \times 10^3 \pi t)$ for $t \in [0, 3]$ and the meltwater calving law. The blue line tracks the terminus location of the stable steady state and the red line tracks the terminus location as a result of evolving the flowline model (2.16). We assume constant width $w = 10$ and that this glacier sits on the upward sloping bed defined by (2.50). We also set $a = 1$ and $\mu = 0$.

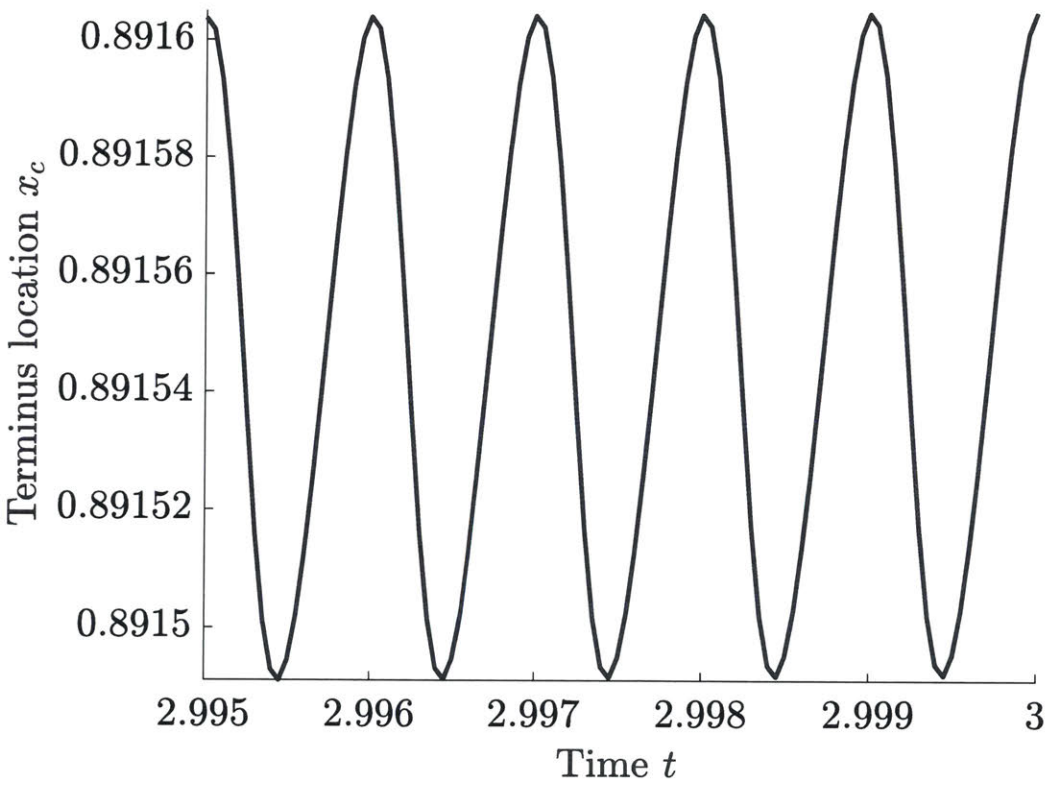

Figure 2-24: Zoomed-in on the stable portion of the simulation shown in figure 2-23.

Chapter 3

The effect of ice shelf buttressing on grounding line flux via matched asymptotic analysis of flowline models

Flowline models are computationally burdensome and provide minimal insight as to which physical processes are responsible for observed behavior. Here, we use analytical methods to relate output flux to glacier geometry; allowing us to compute the output flux *without* solving the full flowline model. We show[1] that the flowline model can be separated into an 'interior' and 'boundary layer' problem. In the interior of the glacier, extensional (viscous) stresses are negligible. This region, where the driving stress is primarily balanced by basal and lateral drag, is termed the 'interior' problem. At the grounding line (or terminus if there is no floating ice shelf) rapid acceleration and thinning makes the extensional stress non-negligible in a small region around the grounding line (or terminus) and along the ice shelf. This 'boundary layer' problem relates the grounding line flux (the rate mass is transported into the ocean) to local bedrock geometry and ice shelf buttressing.

In chapter 2, we showed that a glacier's susceptibility to the marine ice instability depends on the assumed calving law. Here, we explain *why* we observe this behavior. In the melt-limited and calving-at-flotation cases, the output flux monotonically de-

[1]Portions of this chapter have been published in [Schoof et al., 2017].

pends on bedrock depth: retreating into a deeper bed increases the flux and triggers the marine ice instability. However, assuming the meltwater calving law changes this fundamental relationship. When the output flux does not depend monotonically on bedrock depth, retreating into deeper beds may not necessarily trigger the instability. Analyzing limiting cases (e.g., large lateral or basal frictions) provides some insight as to why we observe this non-monotonic relationship.

Previous results determine a similar flux-to-bedrock-depth relationship, however, they neglect lateral drag and assume the glacier calves at flotation. We generalize previous results that relate output flux to glacier geometry to include: (i) side-wall drag, (ii) subshelf melt, and (iii) generalized calving laws. Section 3.1 splits the model into the 'interior' problem (section 3.1.1) and the 'boundary layer' problem (section 3.1.2). Section 3.2 shows solutions to the boundary layer problem and discusses their physical implications. Section 3.3 combines the interior and boundary layer problems and compares this result to the full model (see chapter 2).

3.1 Model approximations

Extensional stresses play an insignificant role in the force balance in the interior of the glacier. Mathematically, we assume that $\varepsilon \ll 1$, which is justified by a scaling argument that suggests $\varepsilon \approx 10^{-3}$—similar to [Schoof, 2007b, Schoof, 2007a, Schoof et al., 2017]. At the grounding line x_g, the ice rapidly thins, accelerates, and begins to float [Schoof, 2007b, Schoof, 2007a, Schoof et al., 2017]. There is, therefore, a small boundary layer around x_g where extensional stress becomes non-negligible. We extend previous boundary layer analysis to consider a generic calving law $h(x_c(t), t) = -\phi(t)b(x_c(t))$ and to include sub-shelf melt. We analyze the extreme case when melt accounts for all of the mass loss (no calving) and determining how sub-shelf melt affects the flux over the grounding line when calving occurs.

Assuming $\varepsilon \ll 1$ leads to a reduced model that evolves the thickness without needing to explicitly compute the velocities u but is also a singular perturbation at the terminus. There is, therefore, a small boundary layer around x_g where extensional

stress becomes non-negligible.

3.1.1 Reduced model

The reduced model is an approximation for the grounded portion of (2.16) whose accuracy depends on the relative importance of the extensional stress. Neglecting the first term in (2.16d) and setting $\theta = 1$ derives an expression for the flux in the grounded portion

$$q = wuh = -\frac{w^{n+2}h^{n+1}}{\left(h + \gamma w^{\frac{1}{n}+1}\right)^n} \left|h_x + b_x\right|^{n-1} (h_x + b_x) \,. \tag{3.1a}$$

We no longer need to explicitly compute the velocity profile and, since the ice is grounded, there are no discontinuities in physics. The flux can be substituted into (6.4b) to derive a diffusive model for the grounded ice thickness

$$wh_t + q_x = wa \tag{3.1b}$$

with boundary conditions

$$q = 0 \quad \text{at} \quad x = 0 \quad \text{and} \quad h = -\max(r^{-1}, \phi)b \quad \text{at} \quad x = x_g(t). \tag{3.1c}$$

The terminal condition is either the flotation thickness—if there is an ice shelf—or the critical terminal thickness prescribed by the meltwater calving law, whichever is larger.

Weak form and numerical methods

We numerically compute solutions to the reduced model using a continuous Galerkin method. We redefine (ν, τ)-coordinates[2] such that the domain of the reduced model is [0, 1]: $\nu = x_g^{-1}x$ and $\tau = t$. In (ν, τ)-coordinates, the reduced model is

[2]Note that these coordinates are different than those used to rescale the full model onto [0, 1]. Here, we normalize by the grounding line location rather than the terminal location.

$$wh_\tau - \frac{\dot{x}_g}{x_g}\nu h_\nu + x_g^{-1} q_\nu = aw \tag{3.2}$$

with boundary conditions

$$q = 0 \quad \text{at} \quad \nu = 0 \quad \text{and} \quad h = -\phi_r b \quad \text{at} \quad \nu = 1, \tag{3.3}$$

where the flux is

$$q = wuh = -\frac{w^{n+2} h^{n+1}}{\left(x_g \left(h + \gamma w^{\frac{1}{n}+1}\right)\right)^n} \left|h_\nu + b_\nu\right|^{n-1} \left(h_\nu + b_\nu\right). \tag{3.4}$$

and $\phi_r = \max\left(\phi(-\lambda b^{-1}), r^{-1}\right)$. Integrating again the test function $z \in H^1([0,1])$ such that $z(1) = 0$ the weak form becomes

$$\int_0^1 wh_\tau z - \frac{\dot{x}_g}{x_g}\nu h_\nu z - x_g^{-1} q z_\nu - awz \, d\nu = 0. \tag{3.5}$$

We find solutions to (3.5) using a continuous Galerkin implementation in the Fenics software package [Alnæs et al., 2015, Logg, 2007, Logg and Wells, 2010, Ølgaard et al., 2008].

Neglecting the extensional stress is a singular perturbation at the terminus: in the calving-limited case we cannot satisfy hydrostatic balance at x_c and in the melt-limited case we cannot enforce that the extensional stress disappears as $x \to x_c$. Since we are no longer simulating the floating portion, the domain of the reduced model does not even reach the relevant boundary location. Given a grounding line location x_g the conditions in (3.1c) are all required to solve the reduced model. We can compute grounded ice thickness profiles with (3.1), but we need a third condition to evolve the grounding line location. This boundary conditions comes out of a boundary layer around the grounding line that allows us to satisfy the stress condition. Generalizing boundary layer from [Schoof, 2007b, Schoof, 2007a], we can compute the flux at x_g as a function of geometric, climate, and material parameters, allowing us to compute x_g independent of (3.1).

3.1.2 Boundary layer analysis

We construct boundary layer theory for the flow of an outlet glacier near its grounding line. The solution to the boundary layer problem is the flux out of the grounded portion of the glacier. Specifically, we neglect the extensional stress in the glacier's interior and compute the flux into the grounding line via matched asymptotic analysis. We assume that lateral drag and sub-shelf melt play a leading-order role in order for the floating portion to have a significant effect in buttressing the grounding line.

Rescaling inside the boundary layer

We begin by rescaling the problem in a boundary layer around the grounding line x_g. Define rescaled variables

$$X = \varepsilon^{\beta}(x - x_g), \quad H = \varepsilon^{\iota} h, \quad U = \varepsilon^{\alpha} U, \quad \text{and} \quad T = t \tag{3.6}$$

and let

$$X_c = \varepsilon^{\beta}(x_c - x_g) \tag{3.7}$$

be the finite ice shelf length. Rescaling H implies that ice thickness near the grounding line is small compared with the interior of the glacier. If there is an ice shelf, the glacier must reach its flotation thickness at the grounding line and we assume the glacier is at least near flotation if there is no floating ice shelf. This implies that around the grounding line the bed elevation must be similar in size to the rescaled thickness—

$$B = \varepsilon^{\iota} b, \tag{3.8}$$

where we assume $B \sim \mathcal{O}(1)$. We also assume the width is constant over the short length scale associated with the 'inner' boundary layer problem

$$W = w(x_g). \tag{3.9}$$

We substitute these rescaled variables into the force balance equation (note that $\frac{\partial}{\partial x} = \frac{\partial X}{\partial x}\frac{\partial}{\partial X} = \varepsilon^{\beta}\frac{\partial}{\partial X}$)

$$4\varepsilon^{1-\iota+\beta+\frac{\beta-\alpha}{n}}\left(H\left|U_X\right|^{\frac{1}{n}-1}U_X\right)_X - W^{-\frac{1}{n}-1}\varepsilon^{-\iota-\frac{\alpha}{n}}H\left|U\right|^{\frac{1}{n}-1}U - \iota\theta\varepsilon^{-m\alpha}\left|U\right|^{m-1}U = \left(1-(1-\theta)r\right)\varepsilon^{-2\iota}H\left(\varepsilon^{\beta}H_X + \theta B_x\right), \quad (3.10)$$

where we have assumed the bedrock varies on length scales associated with the outer problem, and the mass balance equation

$$\varepsilon^{-\iota}H_T - \varepsilon^{-\iota-\beta}H_X\varepsilon^{\beta-\alpha-\iota}(UH)_X = \left(a-(1-\theta)\mu\right). \quad (3.11)$$

We assume that extensional stress, lateral stress and driving stress all balance and we assume $UH \sim \mathcal{O}(1)$ to derive

$$1+\beta+\frac{\beta}{n} = 0, \quad -\frac{\alpha}{n} = -\iota+\beta, \quad \text{and} \quad \iota = -\alpha, \quad (3.12a)$$

whose solution is

$$\beta = -\frac{n}{n+1}, \quad \iota = -\frac{n^2}{(n+1)^2}, \quad \text{and} \quad \alpha = \frac{n^2}{(n+1)^2}. \quad (3.12b)$$

The rescaled equations to leading order in ε are

$$4\left(H\left|U_X\right|^{\frac{1}{n}-1}U_X\right)_X - W^{-\frac{1}{n}-1}H\left|U\right|^{\frac{1}{n}-1}U - \theta\Gamma\left|U\right|^{m-1}U = \left(1-(1-\theta)r\right)HH_X \quad (3.13a)$$

$$(UH)_X = -(1-\theta)M, \quad (3.13b)$$

where

$$\Gamma = \varepsilon^{\frac{n-mn^2-n^2}{(n+1)^2}}\gamma \quad \text{and} \quad M = \varepsilon^{\frac{n}{n+1}}\mu. \tag{3.14}$$

We have retained two factors of ε: the coefficients Γ and M (later we will also retain a factor of ε in the calving-limited terminal condition). We retain these because they represent physical processes that we expect to be potentially relevant near the grounding line. Formally, we have assumed $\Gamma, M \sim \mathcal{O}(1)$. Also defining $H_g = -r^{-1}B(0)$, the grounded versus floating indicator becomes

$$\theta = \begin{cases} 1 & \text{if } H \geq H_g \\ 0 & \text{otherwise.} \end{cases} \tag{3.15}$$

The rescaled model (3.13) is a second order system of ordinary differential equations with dependent variables U and H and independent variable X.

We derive terminal conditions for the 'inner' problem by applying the rescaling to the boundary conditions. In the calving-limited case,

$$H = rH_g\phi(\Lambda H_g^{-1}) \quad \text{and} \quad 4H\,|U_X|^{\frac{1}{n}-1}U_X = \tfrac{1-(1-\theta)r}{2}H^2 - \tfrac{\theta r}{2}H_g^2 \quad \text{at} \quad X = X_c \tag{3.16}$$

(defining $\Lambda = r\lambda\varepsilon^{\iota}$ — another assumed $\mathcal{O}(1)$ parameter that includes ε) or, in the melt-limited case,

$$UH \to 0 \quad \text{and} \quad H\,|U_X|^{\frac{1}{n}-1}U_X \to 0 \quad \text{as} \quad X \to X_c. \tag{3.17}$$

Having rescaled inside a boundary layer around x_g, we must define matching conditions such that the 'outer' and 'inner' problems match as the ice approaches or

leaves the boundary layer, respectively. These conditions take the form

$$Q = \lim_{X\to-\infty} UH = \lim_{x\to x_g^-} -\frac{w^{n+1}h^{n+1}}{(h+\gamma w^{\frac{1}{n}+1})^n}|h_x+b_x|^{n-1}(h_x+b_x) \tag{3.18a}$$

$$W^{-\frac{1}{n}-1}Q|U|^{\frac{1}{n}-1} \sim -\frac{Q}{U}\left(\left(\frac{Q}{U}\right)_X + D\right) \tag{3.18b}$$

$$\lim_{X\to-\infty} U = 0. \tag{3.18c}$$

These require the flux into the boundary layer to match the flux leaving the boundary layer in the upstream direction, the extensional stress to disappear leaving the boundary layer in the upstream direction, and the rescaled velocity to disappear, respectively. The flux at the boundary of the domain of the outer problem Q is to be determined as the solution of the boundary layer problem.

Further rescaling

The boundary layer analysis should provide a recipe for Q in terms of local geometric properties—for example, the flotation thickness at the grounding line H_g—and known climate parameters—such as melt rate M or friction coefficient Γ. We simplify the form of this relationship by applying a further rescaling:

$$U = Q^{\frac{n}{n+1}}W^{\frac{1}{n+1}}\mathcal{U}, \quad H = Q^{\frac{1}{n+1}}W^{-\frac{1}{n+1}}\mathcal{H}, \quad \text{and} \quad X = W\mathcal{X} \tag{3.19}$$

and defining

$$\mathcal{M} = \tfrac{MW}{Q}, \quad \mathcal{C} = Q^{\frac{nm-2}{n+1}}W^{\frac{m+n+3}{n+1}}\Gamma, \quad \mathcal{H}_g = Q^{-\frac{1}{n+1}}W^{\frac{1}{n+1}}H_g, \quad \text{and} \quad \mathcal{L} = \Lambda Q^{-\frac{1}{n+1}}W^{\frac{1}{n+1}}. \tag{3.20}$$

Substituting the secondary rescaling into (3.13) derives

$$4\left(\mathcal{H}\left|\mathcal{U}_{\mathcal{X}}\right|^{\frac{1}{n}-1}\mathcal{U}_{\mathcal{X}}\right)_{\mathcal{X}}-\mathcal{H}\left|\mathcal{U}\right|^{\frac{1}{n}-1}\mathcal{U}-\theta\mathcal{C}\left|\mathcal{U}\right|^{m-1}\mathcal{U}$$
$$=(1-(1-\theta)r)\,\mathcal{H}\mathcal{H}_{\mathcal{X}} \quad (3.21a)$$

$$(\mathcal{U}\mathcal{H})_{\mathcal{X}}=-(1-\theta)\mathcal{M} \quad (3.21b)$$

($\theta = 1$ if $\mathcal{H} \geq \mathcal{H}_g$ and $\theta = 0$ otherwise). We need $\mathcal{U}\mathcal{H}$ to be constant over the grounded portion of the boundary layer. Here, we have kept $\mathcal{M}$, assuming that subshelf melt plays an important role in the determining the ice shelf geometry. The terminal conditions become (letting $\mathcal{X}_c = WX_c$)

$$\mathcal{H}=r\mathcal{H}_g\phi(\mathcal{L}\mathcal{H}_g^{-1}) \quad \text{and} \quad 4\mathcal{H}\left|\mathcal{U}_{\mathcal{X}}\right|^{\frac{1}{n}-1}\mathcal{U}_{\mathcal{X}}=\tfrac{1-(1-\theta)r}{2}\mathcal{H}^2-\tfrac{\theta r}{2}\mathcal{H}_g^2 \quad \text{at} \quad \mathcal{X}=\mathcal{X}_c \quad (3.22)$$

or

$$\mathcal{U}\mathcal{H}\to 0 \quad \text{and} \quad \mathcal{H}\left|\mathcal{U}_{\mathcal{X}}\right|^{\frac{1}{n}-1}\mathcal{U}_{\mathcal{X}}\to 0 \quad \text{as} \quad \mathcal{X}\to\mathcal{X}_c. \quad (3.23)$$

The matching conditions (3.18) become

$$\lim_{\mathcal{X}\to-\infty}\mathcal{U}\mathcal{H}\sim 1, \quad \left|\mathcal{U}\right|^{\frac{1}{n}-1}\mathcal{U}\sim\mathcal{U}^{-2}\mathcal{U}_{\mathcal{X}}-\mathcal{D}, \quad \text{and} \quad \lim_{\mathcal{X}\to-\infty}\mathcal{U}=0. \quad (3.24)$$

The transformed boundary layer equations in (3.21) constitute a second order system, which we will transform into a first order autonomous system. Define another transformation

$$\mathcal{Q}=\mathcal{U}\mathcal{H}, \quad \xi=\mathcal{U}^{\frac{(n+1)^2}{n^2}}, \quad \text{and} \quad \Psi=\mathcal{Q}^{-1}\mathcal{U}^{\frac{-(2n+1)}{n^2}}\left|\mathcal{U}_{\mathcal{X}}\right|^{\frac{1}{n}-1}\mathcal{U}_{\mathcal{X}}. \quad (3.25)$$

The matching conditions (3.24) imply that $\mathcal{Q} \to 1$ and $\mathcal{U} \to 0$ as $\mathcal{X} \to -\infty$; therefore, $\xi \to 0$. For convenience, we note that

$$\mathcal{U} = \xi^{\frac{n^2}{(n+1)^2}} \tag{3.26a}$$

$$\mathcal{H} = \mathcal{Q}\xi^{-\frac{n^2}{(n+1)^2}} \tag{3.26b}$$

$$\mathcal{U}_{\mathcal{X}} = \mathcal{Q}^n \xi^{\frac{2n^2+n}{(n+1)^2}} \Psi^n \tag{3.26c}$$

$$\mathcal{H}\left|\mathcal{U}_{\mathcal{X}}\right|^{\frac{1}{n}-1}\mathcal{U}_{\mathcal{X}} = \mathcal{Q}^2 \xi^{\frac{-n^2+2n+1}{(n+1)^2}} \Psi. \tag{3.26d}$$

The derivatives of $\mathcal{Q}$ and ξ with respect to the independent variable $\mathcal{X}$ are

$$\mathcal{Q}_{\mathcal{X}} = (\mathcal{U}\mathcal{H})_{\mathcal{X}} = -(1-\theta)\mathcal{M} \tag{3.27a}$$

$$\xi_{\mathcal{X}} = \frac{(n+1)^2}{n^2}\mathcal{U}^{\frac{(n+1)^2}{n^2}-1}\mathcal{U}_{\mathcal{X}} = \frac{(n+1)^2}{n^2}\mathcal{Q}^n \xi^{\frac{2n+1}{n+1}} \Psi^n. \tag{3.27b}$$

Substituting the transformed variables into (3.21a) derives the derivative of Ψ with respect to $\mathcal{X}$:

$$4\left(\mathcal{Q}^2 \xi^{\frac{-n^2+2n+1}{(n+1)^2}} \Psi\right)_X - \mathcal{Q}\left|\xi\right|^{\frac{-n^2+n}{(n+1)^2}} - \theta\mathcal{C}\left|\xi\right|^{\frac{mn^2-n^2}{(n+1)^2}} \xi^{\frac{n^2}{(n+1)^2}}$$
$$= (1-(1-\theta)r)\,\mathcal{Q}\xi^{-\frac{n^2}{(n+1)^2}}\left(\mathcal{Q}\xi^{-\frac{n^2}{(n+1)^2}}\right)_{\mathcal{X}} \tag{3.28}$$

Requiring $\xi > 0$ (the velocity must be positive), defining $\delta = 1-(1-\theta)r$, and solving for $\Psi_{\mathcal{X}}$ derives

$$\Psi_X = \frac{n^2-2n-1}{n^2}\mathcal{Q}^n \xi^{\frac{n}{n+1}} \Psi^{n+1} - 2\mathcal{Q}^{-1}\Psi\mathcal{Q}_{\mathcal{X}}$$
$$+ \frac{1}{4}\mathcal{Q}^{-1}\xi^{-\frac{1}{n+1}} + \frac{1}{4}\theta\mathcal{C}\mathcal{Q}^{-2}\xi^{\frac{mn^2+n^2-2n-1}{(n+1)^2}} + \frac{1}{4}\delta\mathcal{Q}^{-1}\xi^{-1}\mathcal{Q}_{\mathcal{X}} - \frac{1}{4}\delta\mathcal{Q}^n\xi^{-\frac{1}{n+1}}\Psi^n \tag{3.29}$$

Factor out $\alpha = \mathcal{Q}^{-1}\xi^{-\frac{1}{n+1}}\Psi^n$ and the first order system becomes

$$\mathcal{Q}_{\mathcal{X}} = -\alpha(1-\theta)\mathcal{M}\mathcal{Q}\xi^{\frac{1}{n+1}}\Psi^{-n} \tag{3.30a}$$

$$\xi_{\mathcal{X}} = \alpha\frac{(n+1)^2}{n^2}\mathcal{Q}^{n+1}\xi^2 \tag{3.30b}$$

$$\begin{aligned}\Psi_{\mathcal{X}} = \alpha\Bigg[&\frac{n^2-2n-1}{n^2}\mathcal{Q}^{n+1}\xi\Psi + 2(1-\theta)\mathcal{M}\xi^{\frac{1}{n+1}}\Psi^{-n+1} \\ &+\frac{1}{4}\Psi^{-n} + \frac{1}{4}\theta\mathcal{C}\mathcal{Q}^{-1}\xi^{\frac{mn^2+n^2-n}{(n+1)^2}}\Psi^{-n} \\ &-\frac{1}{4}\delta(1-\theta)\mathcal{M}\xi^{-\frac{n}{n+1}}\Psi^{-n} - \frac{1}{4}\delta\mathcal{Q}^{n+1}\Bigg]\end{aligned} \tag{3.30c}$$

where the indicator $\theta = 1$ if $\mathcal{Q}\xi^{-\frac{n^2}{(n+1)^2}} \geq \mathcal{H}_g$ and $\theta = 0$ otherwise. We seek orbits of (3.30) that satisfy the matching conditions

$$(\mathcal{Q},\xi,\Psi) \to (1,0,1) \quad \text{as} \quad \mathcal{X} \to -\infty \tag{3.31a}$$

and either the melt-limited terminal conditions

$$\mathcal{Q} \to 0 \quad \text{and} \quad \mathcal{Q}^2\xi^{\frac{-n^2+2n+1}{(n+1)^2}}\Psi \to 0 \quad \text{as} \quad \mathcal{X} \to \mathcal{X}_c \tag{3.31b}$$

or the calving-limited terminal conditions

$$\mathcal{Q}\xi^{-\frac{n^2}{(n+1)^2}} = r\mathcal{H}_g\phi(\mathcal{L}\mathcal{H}_g^{-1}) \quad \text{and} \quad \xi\Psi = \tfrac{1-(1-\theta)r}{8} - \tfrac{\theta r}{8}\mathcal{H}_g^2\mathcal{Q}^{-2}\xi^{\frac{2n^2}{(n+1)^2}} \quad \text{at} \quad \mathcal{X} = \mathcal{X}_c. \tag{3.31c}$$

A functional relationship that defines the flux at the grounding line exists if and only if an orbit of (3.30) satisfies the matching and terminal conditions. This functional relationship expresses the flux at x_g (or, equivalently, the flux at $X = 0$) as a function of geometric, climatic, and material parameters—

$$Q_g = Q(\underbrace{H_g, W,}_{\substack{\text{geometric}\\ \text{parameters}}} \underbrace{M, \Lambda,}_{\substack{\text{climate}\\ \text{parameters}}} \underbrace{\Gamma, n, m, r}_{\substack{\text{material}\\ \text{parameters}}}). \tag{3.32}$$

The relationship is implicitly defined; we find Q_g that satisfies

$$F(\mathcal{H}_g, \mathcal{M}, \mathcal{C}, \mathcal{L}, m, n, r) = 0, \tag{3.33}$$

for an F whose form depends on the calving law—note that $\mathcal{H}_g$, $\mathcal{M}$, $\mathcal{C}$, and $\mathcal{L}$ depend on Q. We show that finding an orbit of (3.30) that satisfies the matching and terminal conditions is equivalent to finding the roots of F. This task is simplified by absorbing α into to the independent variable by defining $\omega = \int_0^{\mathcal{X}} \alpha(\mathcal{X}')d\mathcal{X}'$, which makes (3.30) a nonsingular dynamical system with independent variable ω. We must find an orbit that satisfies the matching conditions (3.31a) and either the calving-limited terminal conditions (3.31c) or the melt-limited terminal conditions (3.31b).

In both cases, the orbit-of-interest must first satisfy the matching conditions (3.31a). There must be an orbit emerging from the fixed point $(\mathcal{Q}, \xi, \Psi) = (1, 0, 1)$ with $\theta = 1$. The glacier remains grounded around this fixed point and, therefore, we can restrict the orbit to the $\mathcal{Q} = 1$ plane with $\theta = 1$. This simplifies (3.30) to

$$\xi_\omega = \frac{(n+1)^2}{n^2}\xi^2 \tag{3.34a}$$

$$\Psi_\omega = \frac{n^2 - 2n - 1}{n^2}\xi\Psi + \frac{1}{4}\Psi^{-n} + \frac{1}{4}\mathcal{C}\xi^{\frac{mn^2+n^2-n}{(n+1)^2}}\Psi^{-n} - \frac{1}{4}, \tag{3.34b}$$

whose Jacobian has one negative and one zero Eigenvalue at the fixed point $(\xi, \Psi) = (0, 1)$. Therefore, there is a stable manifold and an unstable centre manifold around the fixed point, which ensures the centre manifold is unique [Sijbrand, 1985]. The orbit-of-interest is associated with the centre manifold and, therefore, is also unique. Any orbit satisfying (3.31) must emerge from $(\mathcal{Q}, \xi, \Psi) = (1, 0, 1)$ and also satisfy either (3.31c) or (3.31b). We seek functions of the form (3.33) such that finding root

of F corresponds to finding the orbit-of-interest.

Melt-limited terminus Orbits that satisfy the melt-limited condition (3.31b) terminate at another fixed point of (3.30). This fixed point corresponds to the limit $\mathcal{Q} \to 0$ and is not unique. It satisfies

$$\xi\Psi = -\frac{1}{8\mathcal{M}}\xi^{\frac{n}{n+1}} + \frac{1-r}{8} \tag{3.35}$$

with $\mathcal{Q} = 0$ and $\theta = 0$. This defines a one parameter family of fixed points that corresponds to a centre manifold with no dynamics. The remaining invariant manifolds suggest there that, in addition to the centre manifold, there is a stable and an unstable manifold. We use a shooting method to find ξ_t such that the unique stable orbit entering from the terminal point

$$(\mathcal{Q}_t, \xi_t, \Psi_t) = \left(0, \xi_t, -\frac{1}{8\mathcal{M}}\xi_t^{-\frac{1}{n+1}} + \frac{1-r}{8}\xi_t^{-1}\right), \tag{3.36a}$$

integrating (3.30) backward with respect to ω, intersects the $\mathcal{Q} = 1$ plane when $\xi_f = \mathcal{H}_g^{-\frac{(n+1)^2}{n^2}}$. This defines a critical point

$$(\mathcal{Q}_f, \xi_f, \Psi_f) = \left(1, \mathcal{H}_g^{-\frac{(n+1)^2}{n^2}}, \Psi_f\right). \tag{3.36b}$$

This portion of the orbit corresponds to the region of the boundary layer where the ice is floating. We find a second critical point by integrating (3.34) through the grounded portion of the boundary layer from $(\xi, \Psi) = (0, 1)$ to the flotation thickness—

$$(\mathcal{Q}_g, \xi_g, \Psi_g) = \left(1, \mathcal{H}_g^{-\frac{(n+1)^2}{n^2}}, \Psi_g\right). \tag{3.36c}$$

Both portions of the orbit begin at a fixed point that satisfies either the terminal or matching conditions. Continuity requires that they meet on the $\mathcal{Q} = 1$ plane at the critical flotation thickness. We, therefore, define

$$F(\mathcal{H}_g, \mathcal{M}, \mathcal{C}, \mathcal{L}, m, n, r) = \Psi_f - \Psi_g, \tag{3.37}$$

which has the desired form (3.33). When $F = 0$, the two orbits correspond to the unique orbit that satisfies both the terminal and matching conditions.

Calving-limited terminus There are two calving-limited cases: the terminus can either be grounded or floating. If the terminus is grounded because $\mathcal{L}\mathcal{H}_g^{-1} \geq \frac{1}{2}$, then we integrate (3.34) from $(\xi, \Psi) = (0, 1)$ until the first condition of (3.31c) is met, which defines the critical point

$$(\mathcal{Q}', \xi', \Psi') = \left(1, \left(r\mathcal{H}_g\phi(\mathcal{L}\mathcal{H}_g^{-1})\right)^{-\frac{(n+1)^2}{n^2}}, \Psi'\right). \tag{3.38a}$$

In the calving-at-flotation case or if the terminus is floating because $\mathcal{L}\mathcal{H}_g^{-1} < \frac{1}{2}$, we integrate (3.34) from $(\xi, \Psi) = (0, 1)$ to the flotation thickness, which defines the grounding line point

$$(\mathcal{Q}_g, \xi_g, \Psi_g) = \left(1, \mathcal{H}_g^{-\frac{(n+1)^2}{n^2}}, \Psi_g\right). \tag{3.38b}$$

In the calving-at-flotation case, we define the critical point $(\mathcal{Q}', \xi', \Psi') = (\mathcal{Q}_g, \xi_g, \Psi_g)$. If the terminus is floating, the system leaves the $\mathcal{Q} = 1$ plane and we integrate (3.30) from $(\mathcal{Q}_g, \xi_g, \Psi_g)$ to the critical point

$$(\mathcal{Q}', \xi', \Psi') = \left(\mathcal{Q}', \left(r\mathcal{H}_g\phi(\mathcal{L}\mathcal{H}_g^{-1})\ \mathcal{Q}'^{-1}\right)^{-\frac{(n+1)^2}{n^2}}, \Psi'\right). \tag{3.38c}$$

This computes an orbit emerging from $(\mathcal{Q}, \xi, \Psi) = (1, 0, 1)$ that includes $(\mathcal{Q}', \xi', \Psi')$, satisfies the matching conditions (3.31) and the first condition of (3.31c). We define

$$F(\mathcal{H}_g, \mathcal{M}, \mathcal{C}, \mathcal{L}, m, n, r) = \xi'\Psi' - \frac{1 - (1-\theta)r}{8} + \frac{\theta r}{8}\mathcal{H}_g^2\ \mathcal{Q}'^{-2}\ \xi'^{\frac{2n^2}{(n+1)^2}}, \tag{3.39}$$

which, again, has the desired form (3.33). The roots of this function correspond to an orbit that satisfies the matching and terminal conditions.

Finally, we find the roots of F to compute the flux Q entering the boundary layer. We evaluate F given Q by first computing the scaling parameters $\mathcal{H}_g$, $\mathcal{M}$, $\mathcal{C}$, and $\mathcal{L}$

then evaluating either (3.37) or (3.39). We approximate derivatives of F using finite difference

$$\frac{dF}{dQ} \approx \frac{F(Q + \Delta Q) - F(Q)}{\Delta Q} \tag{3.40}$$

and use Newton's method to find $F(Q^*) = 0$. Numerically, we use PETSc's [Balay et al., 2017a, Balay et al., 2017b, Balay et al., 1997] implementation of Newton's method combined with a backtracking line search. This defines the flux into the boundary layer implicitly as a function of geometric, climate, and material properties, as desired.

3.1.3 Boundary layer analysis: limiting cases

In some cases, we can derive an explicit formula for the flux at the grounding line. We formally derive the formulas here, although, we note that later we will discuss these results in a more intuitive way. This section may be skipped by readers who are not interested in the formal derivation of expressions for grounding line flux in limiting cases of the boundary layer analysis. For calving-at-flotation, calving fronts near flotation, and short ice shelves, we assume that $1 - r \ll 1$, which is justified since ice and ocean water have similar densities—$r \approx 0.9$. For notational convenience, define $\delta = 1 - r$.

Calving near flotation

We first consider the case where the terminal thickness $H_c \approx H_g$. From the definition of $\mathcal{H}_g$ in (3.19), we see that if $\mathcal{H}_g \sim \mathcal{O}(1)$ then taking $1 - r \to 0$ has little effect on the relationship between Q, H_g, and W. This is of course not true since the flotation thickness H_g depends on r. Therefore, we let $\bar{\mathcal{H}}_g = \delta^{\frac{n^2}{(n+1)^2}} \mathcal{H}_g$. We rescale again

$$\bar{\mathcal{X}} = \delta^{\frac{n}{n+1}} \mathcal{X}, \quad \bar{\mathcal{U}} = \delta^{-\frac{n^2}{(n+1)^2}} \mathcal{U}, \quad \text{and} \quad \bar{\mathcal{H}} = \delta^{\frac{n^2}{(n+1)^2}} \mathcal{H}, \tag{3.41}$$

which transforms (3.21) into (noting that $\frac{\partial}{\partial \bar{\mathcal{X}}} = \delta^{\frac{n}{n+1}} \frac{\partial}{\partial \mathcal{X}}$ and the grounded terminus

implies $\theta = 1$)

$$4\delta\left(\bar{\mathcal{H}}\left|\bar{\mathcal{U}}_{\bar{\mathcal{X}}}\right|^{\frac{1}{n}-1}\bar{\mathcal{U}}_{\bar{\mathcal{X}}}\right)_{\bar{\mathcal{X}}} - \bar{\mathcal{H}}\left|\bar{\mathcal{U}}\right|^{\frac{1}{n}-1}\bar{\mathcal{U}} - \bar{\mathcal{C}}\left|\bar{\mathcal{U}}\right|^{m-1}\bar{\mathcal{U}} = \bar{\mathcal{H}}\bar{\mathcal{H}}_{\bar{\mathcal{X}}} \tag{3.42a}$$

$$(\bar{\mathcal{U}}\bar{\mathcal{H}})_{\bar{\mathcal{X}}} = 0. \tag{3.42b}$$

We assume $\bar{\mathcal{C}} = \delta^{\frac{n^2m+n^2-n}{(n+1)^2}}\mathcal{C}$ is $\mathcal{O}(1)$—we expect the basal drag to play a leading-order role. Suppose the calving thickness is a small perturbation from the floating thickness

$$\delta^{\frac{n^2}{(n+1)^2}}\mathcal{H}_c = \bar{\mathcal{H}}_g + \delta\bar{\mathcal{H}}_g'. \tag{3.43}$$

At leading order in δ, the boundary conditions become

$$\bar{\mathcal{H}} = \bar{\mathcal{H}}_g \quad \text{and} \quad 4|\bar{\mathcal{U}}_{\bar{\mathcal{X}}}|^{\frac{1-n}{n}}\bar{\mathcal{U}}_{\bar{\mathcal{X}}} = \bar{\mathcal{T}} \tag{3.44}$$

at the grounding line $\bar{\mathcal{X}} = 0$, where $\bar{\mathcal{T}}$ is the extensional stress at the grounding line. Linearizing ϕ around the flotation thickness implies that

$$\bar{\mathcal{T}} = \frac{1}{2}\bar{\mathcal{H}}_g - \bar{\mathcal{H}}_g'. \tag{3.45}$$

The matching conditions (3.24) are the same as $\bar{\mathcal{X}} \to \infty$:

$$\lim_{\bar{\mathcal{X}}\to-\infty}\bar{\mathcal{U}}\bar{\mathcal{H}} \sim 1, \quad |\bar{\mathcal{U}}|^{\frac{1}{n}-1}\bar{\mathcal{U}} \sim \bar{\mathcal{U}}^{-2}\bar{\mathcal{U}}_{\bar{\mathcal{X}}}, \quad \text{and} \quad \lim_{\bar{\mathcal{X}}\to-\infty}\bar{\mathcal{U}} = 0. \tag{3.46}$$

Neglecting the $\mathcal{O}(\delta)$ term in (3.42a) and choosing $m = n^{-1}$ implies that

$$\left|\bar{\mathcal{U}}\right|^{\frac{n+1}{n}} = \frac{\bar{\mathcal{H}}^2}{\bar{\mathcal{H}}+\bar{\mathcal{C}}}\bar{\mathcal{U}}_{\bar{\mathcal{X}}} \tag{3.47}$$

Applying the boundary conditions (3.44) implies

$$\bar{\mathcal{U}}(0)^{\frac{n+1}{n}} = \frac{\bar{\mathcal{H}}_g^{n+2}}{8^n(\bar{\mathcal{H}}_g+\bar{\mathcal{C}})} \tag{3.48}$$

and, therefore,

$$1 = \bar{\mathcal{U}}(0)\bar{\mathcal{H}}(0) = \frac{\bar{\mathcal{H}}_g^{\frac{3n+1}{n+1}} \bar{\mathcal{T}}^{\frac{n^2}{n+1}}}{4^{\frac{n^2}{n+1}} \left(\bar{\mathcal{C}} + \bar{\mathcal{H}}_g\right)^{\frac{n}{n+1}}}. \tag{3.49}$$

We now consider two scenarios: (i) a grounded and (ii) and floating calving front.

Grounded calving front Substituting the definitions of $\bar{\mathcal{H}}_g$, $\bar{\mathcal{H}}_c$, $\bar{\mathcal{C}}$, and $\bar{\mathcal{T}}$, we can express the flux as

$$Q \approx \left(\frac{1-r}{8}\right)^{\frac{n^2}{n+1}} \frac{H_g^{\frac{3n+1}{n+1}} \left(H_g - 2(1-r)^{-1} H_g'\right)^{\frac{n^2}{n+1}}}{\left(W^{-\frac{n+1}{n}} H_g + \Gamma\right)^{\frac{n}{n+1}}}. \tag{3.50}$$

Setting $H_g' = 0$ recovers the calving-at-flotation case:

$$Q = \left(\frac{1-r}{8}\right)^{\frac{n^2}{n+1}} \left(W^{-\frac{n+1}{n}} H_g + \Gamma\right)^{-\frac{n}{n+1}} H_g^{\frac{n^2+3n+1}{n+1}}. \tag{3.51}$$

This is a variant of the grounding flux derived by [Schoof, 2007b, Schoof, 2007a]; we rederive it here to compare this result to our results that include subshelf melt.

Floating calving front Now suppose that the terminal thickness remains near flotation $H_c \approx H_c$ such that there is a small ice shelf. Define

$$\breve{\mathcal{X}} = \delta^{-1}\bar{\mathcal{X}}, \quad \breve{\mathcal{U}} = \bar{\mathcal{U}}, \quad \text{and} \quad \breve{H} = \bar{H} \tag{3.52}$$

and define $\breve{\mathcal{T}}$ such that $\breve{\mathcal{U}}_{\breve{\mathcal{X}}} = 4^{-n}\delta|\breve{\mathcal{T}}|^{n-1}\breve{\mathcal{T}}$.

The rescaled equations (3.21) for the ice tongue become ($\theta = 0$ since the ice is floating at the terminus)

$$\left(\breve{\mathcal{H}}\breve{\mathcal{T}}\right)_{\breve{\mathcal{X}}} - \breve{\mathcal{H}}|\breve{\mathcal{U}}|^{1/n-1}\breve{\mathcal{U}} - \breve{\mathcal{H}}\breve{\mathcal{H}}_{\breve{\mathcal{X}}} = 0 \tag{3.53}$$

$$\left(\breve{\mathcal{H}}\breve{\mathcal{U}}\right)_{\breve{\mathcal{X}}} = 0 \tag{3.54}$$

with boundary conditions (at leading order of δ)

$$\breve{T} = \frac{1}{2}\breve{\mathcal{H}} \tag{3.55a}$$

$$\breve{\mathcal{H}} = \bar{\mathcal{H}}_g \tag{3.55b}$$

at the calving front $\breve{\mathcal{X}} = \breve{\mathcal{X}}_c$ and

$$\breve{\mathcal{U}} = \bar{\mathcal{U}}, \quad \breve{T} = \bar{T}, \quad \text{and} \quad \breve{\mathcal{H}} = \bar{\mathcal{H}}_g + \delta\bar{\mathcal{H}}'_g \tag{3.55c}$$

at the grounding line $\breve{\mathcal{X}} = 0$.

We need to determine the extensional stress $\breve{T}$ to substitute into (3.49) in order to solve for the flux. We, therefore, expand to zeroth order for $\breve{T}$ and $\breve{\mathcal{U}}$ and first order for $\breve{\mathcal{H}}$:

$$\breve{T} = \breve{T}^{(0)} + \mathcal{O}(\delta), \quad \breve{\mathcal{U}} = \breve{\mathcal{U}}^{(0)} + \mathcal{O}(\delta), \quad \text{and} \quad \breve{\mathcal{H}} = \breve{\mathcal{H}}^{(0)} + \delta\breve{\mathcal{H}}^{(1)} + \mathcal{O}(\delta^2). \tag{3.56}$$

Substituting the expansions into (3.54), the boundary conditions (3.55) require

$$\breve{\mathcal{U}}^{(0)} = \lim_{\breve{\mathcal{X}}\to 0^-} \breve{\mathcal{U}} = \text{constant} \quad \text{and} \quad \breve{\mathcal{H}}^{(0)} = \breve{\mathcal{H}}_g = \text{constant} \tag{3.57}$$

(i.e., the thickness and velocity are constant at leading order). Furthermore, at leading order

$$\breve{\mathcal{H}}^{(0)}\breve{T}^{(0)}_{\breve{\mathcal{X}}} - \breve{\mathcal{H}}^{(0)}(\breve{\mathcal{U}}^{(0)})^{\frac{1}{n}} = 0 \tag{3.58}$$

and at first order

$$\breve{\mathcal{U}}^{(0)}\breve{\mathcal{H}}^{(1)}_{\breve{\mathcal{X}}} = -\breve{\mathcal{H}}^{(0)}|\breve{T}^{(0)}|^{n-1}\breve{T}^{(0)} \tag{3.59}$$

with boundary conditions

$$\breve{\mathcal{H}}^{(1)}(0) = \breve{\mathcal{H}}'_g, \quad \breve{\mathcal{H}}^{(1)}(\breve{\mathcal{X}}_c) = 0, \quad \text{and} \quad \breve{\mathcal{T}}^{(0)}(\breve{\mathcal{X}}_c) = \tfrac{1}{8}\breve{\mathcal{H}}_g \tag{3.60}$$

We solve the leading order equation (3.58) for

$$\breve{\mathcal{T}} = \frac{1}{2}\breve{\mathcal{H}}_g - (\breve{\mathcal{U}})^{\frac{1}{n}}\left(\breve{\mathcal{X}}_c - \breve{\mathcal{X}}\right), \tag{3.61}$$

which yields the extensional stress at the grounding line $\breve{\mathcal{X}} = 0$

$$\breve{\mathcal{T}}(0) = \frac{1}{2}\breve{\mathcal{H}}_g - (\breve{\mathcal{U}})^{\frac{1}{n}}\breve{\mathcal{X}}_c. \tag{3.62}$$

Given this expression for $\breve{\mathcal{T}}$, we solve the first order equation (3.59) for the flux

$$\breve{\mathcal{U}}^{(0)}\breve{\mathcal{H}}^{(1)}(\breve{\mathcal{X}}) = \breve{\mathcal{U}}^{(0)}\breve{\mathcal{H}}'_g - \frac{4\breve{\mathcal{H}}_g}{(n+1)(\breve{\mathcal{U}}^{(0)})^{\frac{1}{n}}}\left[\frac{1}{8}\breve{\mathcal{H}}_g - \frac{1}{4}(\breve{\mathcal{U}}^{(0)})^{\frac{1}{n}}\left(\breve{\mathcal{X}}_c - \breve{\mathcal{X}}\right)\right]^{n+1} \tag{3.63}$$

The boundary conditions $\breve{\mathcal{H}}^{(1)}(\breve{\mathcal{X}}_c) = 0$ and $\breve{\mathcal{H}}^{(1)}(0) = \breve{\mathcal{H}}'_g$ in imply that

$$\left(\left(\frac{\breve{\mathcal{H}}_g}{8}\right)^{n+1} - \frac{(n+1)(\breve{\mathcal{U}}^{(0)})^{\frac{n+1}{n}}\breve{\mathcal{H}}'_g}{4\breve{\mathcal{H}}_g}\right)^{\frac{1}{n+1}} = \frac{1}{8}\breve{\mathcal{H}}_g - \frac{1}{4}(\breve{\mathcal{U}}^{(0)})^{\frac{1}{n}}\breve{\mathcal{X}}_c = \frac{1}{4}\breve{\mathcal{T}}(0). \tag{3.64}$$

Given that $\bar{\mathcal{T}} = \breve{\mathcal{T}}$, we see from (3.49)

$$1 = \frac{\breve{\mathcal{H}}_g^{\frac{3n+1}{n}}}{\bar{C} + \breve{\mathcal{H}}_g}\left(\left(\frac{\breve{\mathcal{H}}_g}{8}\right)^{n+1} - \frac{(n+1)(\breve{\mathcal{U}}^{(0)})^{\frac{n+1}{n}}\breve{\mathcal{H}}'_g}{4\breve{\mathcal{H}}_g}\right)^{\frac{n}{n+1}} \tag{3.65}$$

Substituting the definitions of $\breve{\mathcal{H}}_g$ and $\breve{\mathcal{H}}'_g$, we obtain an implicit expression for the grounding line flux

$$Q \approx \frac{H_g^{\frac{3n+1}{n+1}}}{\left(\Gamma + W^{-\frac{n+1}{n}}H_g\right)^{\frac{n}{n+1}}}\left(\left(\frac{(1-r)H_g}{8}\right)^{n+1} - \frac{(n+1)Q^{\frac{n+1}{n}}H'_g}{W^{\frac{n+1}{n}}H_g^{\frac{2n+1}{n}}}\right)^{\frac{n^2}{(n+1)^2}}. \tag{3.66}$$

Short ice shelves: large melt rate

Large melt rates, specifically, melt rates such that $\mathcal{M} \sim \mathcal{O}(\delta^{-\frac{1}{n+1}})$, result in short ice shelves. We derive an implicit formula for the grounding line flux Q. First, define

$$\tilde{\mathcal{M}} = \delta^{\frac{1}{n+1}}\mathcal{M}, \quad \tilde{\mathcal{L}} = \delta^{\frac{n^2}{(n+1)^2}}\mathcal{L}, \quad \text{and} \quad \tilde{\mathcal{H}}_g = \delta^{\frac{n^2}{(n+1)^2}}\mathcal{H}_g. \tag{3.67}$$

Letting $\mathcal{S} = |\mathcal{U}_\mathcal{X}|^{\frac{1}{n}-1}\mathcal{U}_\mathcal{X}$ be the extensional stress, we define the rescaling

$$\tilde{\mathcal{X}} = \delta^{-\frac{1}{n+1}}\mathcal{X}, \quad \tilde{\mathcal{U}} = \delta^{-\frac{n^2}{(n+1)^2}}\mathcal{U}, \quad \tilde{\mathcal{H}} = \delta^{\frac{n^2}{(n+1)^2}}\mathcal{H}, \quad \text{and} \quad \tilde{\mathcal{S}} = \delta^{-\frac{2n+1}{(n+1)^2}}\mathcal{S}, \tag{3.68}$$

noting that this is different than the calving-at-flotation case. In the ice shelf, $\theta = 0$ and (3.21) becomes (noting that $\frac{\partial}{\partial \mathcal{X}} = \delta^{-\frac{1}{n+1}}\frac{\partial}{\partial \tilde{\mathcal{X}}}$)

$$4\left(\tilde{\mathcal{H}}\tilde{\mathcal{S}}\right)_{\tilde{\mathcal{X}}} - \tilde{\mathcal{H}}\left|\tilde{\mathcal{U}}\right|^{\frac{1}{n}-1}\tilde{\mathcal{U}} = \tilde{\mathcal{H}}\tilde{\mathcal{H}}_{\tilde{\mathcal{X}}} \tag{3.69a}$$

$$(\tilde{\mathcal{U}}\tilde{\mathcal{H}})_{\tilde{\mathcal{X}}} = -\tilde{\mathcal{M}} \tag{3.69b}$$

$$\tilde{\mathcal{U}}_{\tilde{\mathcal{X}}} = \delta\left|\tilde{\mathcal{S}}\right|^{n-1}\tilde{\mathcal{S}} \tag{3.69c}$$

with boundary conditions

$$\tilde{\mathcal{H}} = \tilde{\mathcal{H}}_c = r\tilde{\mathcal{H}}_g\phi(\tilde{\mathcal{L}}\tilde{\mathcal{H}}_g^{-1}) \quad \text{and} \quad 8\tilde{\mathcal{H}}\tilde{\mathcal{S}} = \tilde{\mathcal{H}}^2 \quad \text{at} \quad \tilde{\mathcal{X}} = \tilde{\mathcal{X}}_c. \tag{3.70}$$

We neglected terms multiplied by δ, which implies $\tilde{\mathcal{U}}_{\tilde{\mathcal{X}}} = 0$ and $\tilde{\mathcal{U}} > 0$ is constant. Therefore, $\tilde{\mathcal{U}}\tilde{\mathcal{H}}_{\tilde{\mathcal{X}}} = -\tilde{\mathcal{M}}$ and

$$\tilde{\mathcal{H}} = \tilde{\mathcal{H}}_g - \tilde{\mathcal{M}}\tilde{\mathcal{U}}^{-1}\tilde{\mathcal{X}} \quad \text{and} \quad \mathcal{X}_c = \tilde{\mathcal{U}}\tilde{\mathcal{M}}^{-1}\left(\tilde{\mathcal{H}}_g - \tilde{\mathcal{H}}_c\right). \tag{3.71}$$

Integrating (3.69a) implies

$$4\tilde{\mathcal{H}}(\tilde{\mathcal{X}})\tilde{\mathcal{S}}(\tilde{\mathcal{X}}) = \left(\tilde{\mathcal{U}}^{\frac{1}{n}} - \tilde{\mathcal{M}}\tilde{\mathcal{U}}^{-1}\right)\left(\tilde{\mathcal{H}}_g\tilde{\mathcal{X}} - \frac{1}{2}\tilde{\mathcal{M}}\tilde{\mathcal{U}}^{-1}\tilde{\mathcal{X}}^2\right) + \frac{1}{2}\tilde{\mathcal{H}}_g^2 - \frac{1}{2}\tilde{\mathcal{U}}^{\frac{1}{n}+1}\tilde{\mathcal{M}}^{-1}\left(\tilde{\mathcal{H}}_g^2 - \tilde{\mathcal{H}}_c^2\right) \tag{3.72}$$

and, therefore, at $\tilde{\mathcal{X}} = 0$

$$\tilde{\mathcal{S}}(0) = \frac{1}{8}\tilde{\mathcal{H}}_g - \frac{1}{8}\tilde{\mathcal{U}}^{\frac{1}{n}+1}\tilde{\mathcal{M}}^{-1}\left(\tilde{\mathcal{H}}_g - \frac{\tilde{\mathcal{H}}_c^2}{\tilde{\mathcal{H}}_g}\right) \tag{3.73}$$

Following the steps from the calving-at-flotation case—reusing the rescaling (3.41)—we can derive an expression for the extensional stress at $\tilde{\mathcal{X}} = 0$. Effectively, we substitute (3.68) into (3.47) at $\tilde{\mathcal{X}} = 0$

$$\tilde{\mathcal{U}}^{\frac{n+1}{n}} = \frac{\tilde{\mathcal{H}}_g^2}{\tilde{\mathcal{H}}_g + \tilde{\mathcal{C}}}\tilde{\mathcal{S}}(0)^n. \tag{3.74}$$

We substitute this into the matching conditions that require $\tilde{\mathcal{U}}\tilde{\mathcal{H}}_g = 1$

$$1 = \frac{\tilde{\mathcal{H}}_g^{\frac{3n+1}{n+1}}}{\left(\tilde{\mathcal{H}}_g + \tilde{\mathcal{C}}\right)^{\frac{n}{n+1}}}\tilde{\mathcal{S}}(0)^{\frac{n^2}{n+1}} = \frac{\tilde{\mathcal{H}}_g^{\frac{3n+1}{n+1}}}{\left(\tilde{\mathcal{H}}_g + \tilde{\mathcal{C}}\right)^{\frac{n}{n+1}}}\left(\frac{1}{8}\tilde{\mathcal{H}}_g - \frac{1}{8}\tilde{\mathcal{H}}_g^{-\frac{2n+1}{n}}\tilde{\mathcal{M}}^{-1}\left(\tilde{\mathcal{H}}_g^2 - \tilde{\mathcal{H}}_c^2\right)\right)^{\frac{n^2}{n+1}}, \tag{3.75}$$

where we have also substituted $\tilde{\mathcal{U}} = \tilde{\mathcal{H}}_g^{-1}$. Assuming $m = n^{-1}$, we derive an implicit expression for the flux Q

$$Q = \frac{H_g^{\frac{3n+1}{n+1}}}{\left(W^{-\frac{n+1}{n}}H_g + \Gamma\right)^{\frac{n}{n+1}}}\left(\frac{1-r}{8}H_g - \frac{Q^{\frac{n+1}{n}}}{8W^{\frac{n+1}{n}}H_g^{\frac{2n+1}{n}}M}\left(H_g^2 - (rH_g\phi(\Lambda H_g^{-1}))^2\right)\right)^{\frac{n^2}{n+1}}. \tag{3.76}$$

Setting $\phi = 0$ recovers the melt-limited extreme. We also note that as $M \to \infty$, this expression becomes equivalent to (3.51). Calving-at-flotation is the extreme case of infinite subshelf melt.

Large basal friction

A large calving cliff protruding above sea level causes a sharp increase in the extensional stress, which leads to an anomalous flux-to-bedrock-depth relationship. Importantly, this behavior happens because the calving cliff protrudes above the flotation thickness and not because of the lateral drag. We show this by investigating the limit of large basal friction. The relationship derived in (3.49) is still valid, but the denominator is dominated by $\bar{C}$. However, (3.49) no longer holds if we have a floating calving front, which is why the anomalous flux-to-bedrock depth relationship does not apply if there is an ice shelf.

Formally, we need another rescaling. We see from (3.45) that $\breve{T} \sim \breve{H}$ and, therefore, (3.49) with large $\bar{C}$ implies that $\breve{U} \sim \breve{H}^{\frac{2n}{n+1}} \breve{T}^{\frac{n^2}{n+1}} \bar{C}^{-\frac{n}{n+1}}$. The constraint that $\breve{U}\breve{H}$ implies a rescaling

$$\breve{H} \sim \bar{C}^{\frac{n}{n^2+3n+1}} \breve{H}, \quad \breve{U} \sim \bar{C}^{-\frac{n}{n^2+3n+1}} \breve{U}, \quad \text{and} \quad \breve{T} \sim \bar{C}^{\frac{n}{n^2+3n+1}} \breve{T}. \tag{3.77}$$

and, in the floating ice tongue,

$$\breve{x} = \bar{C}^{(n+1)/(n^2+3n+1)} \breve{x}. \tag{3.78}$$

The ice tongue model (3.54) becomes

$$\begin{aligned} (\breve{H}\breve{T})_{\breve{x}} - \breve{H}|\breve{U}|^{\frac{1-n}{n}}\breve{U} - \breve{H}\breve{H}_{\breve{x}} &= 0 && \text{(3.79a)} \\ \left(\breve{H}\breve{U}\right)_{\breve{x}} &= 0 && \text{(3.79b)} \\ \breve{U}_{\breve{x}} &= 4^{-n}\bar{C}^{\frac{(n+1)^2}{n^2+3n+1}} \delta |\breve{T}|^{n-1}\breve{T}, && \text{(3.79c)} \end{aligned}$$

Unlike (3.54), if $\bar{C} \sim \delta^{-\frac{n^2+3n+1}{(n+1)^2}}$, the stretching rate $\breve{U}_{\breve{x}}$ and the surface slope $\breve{H}_{\breve{x}}$ are not negligible. Additionally, we can no longer neglect the driving stress in the force balance of the ice shelf: when $\bar{C} \gg \delta^{-\frac{n^2+3n+1}{(n+1)^2}}$, the driving stress dominates lateral shear stress.

Recall, $\bar{C} = \delta^{\frac{n^2}{(n+1)^2}} C$, and $\breve{H} = \delta^{\frac{n^2}{n+1}} H \bar{C}^{\frac{n}{n^2+3n+1}} \sim \mathcal{O}(1)$. As we require $H \sim L$, it

follows that $\mathcal{L} \sim \delta^{-\frac{n^2}{n+1}} \bar{\mathcal{C}}^{-\frac{n}{n^2+3n+1}}$, and we have from the definitions of $\mathcal{C}$ and $\mathcal{L}$ with $m = 1/n$ that

$$\frac{\Gamma W^{\frac{n+1}{n}}}{\Lambda} = \frac{\mathcal{C}}{\mathcal{L}} \sim \bar{\mathcal{C}}^{\frac{(n+1)^2}{n^2+3n+1}} \sim \delta. \tag{3.80}$$

Therefore, the parameter regime in which a floating shelf no longer simply buttresses the grounding line is $\Gamma \sim \delta^{-\frac{2n^2+3n+1}{(n+1)^2}} \Lambda W^{-\frac{n+1}{n}}$.

Long ice shelves

In the case of long ice shelves, we are concerned with the limit of a deep bed at the grounding line coupled with relatively small subshelf melt and critical calving thickness—mathematically, $\mathcal{H}_g \gg 1$ while $\mathcal{M}$ and $\mathcal{L}$ are relatively small. We assume the $\mathcal{L}$ is small enough that the terminal thickness $\mathcal{H}_c \ll 1$ and that the system is well approximated by the melt-limited case. Rescaling

$$\hat{\mathcal{H}} = \mathcal{H}_g^{-1}\mathcal{H}, \quad \hat{\mathcal{U}} = \mathcal{H}_g\mathcal{U}, \quad \text{and} \quad \hat{\mathcal{X}} = \mathcal{H}_g^{-\frac{n+1}{n}}\mathcal{X}, \tag{3.81}$$

implies (again, noting that $\frac{\partial}{\partial \mathcal{X}} = \mathcal{H}_g^{-\frac{n+1}{n}} \frac{\partial}{\partial \hat{\mathcal{X}}}$ and assuming that $m = n^{-1}$)

$$4\mathcal{H}_g^{-\frac{(n+1)^2}{n^2}} \left(\hat{\mathcal{H}} \left| \hat{\mathcal{U}}_{\hat{\mathcal{X}}} \right|^{\frac{1}{n}-1} \hat{\mathcal{U}}_{\hat{\mathcal{X}}} \right)_{\hat{\mathcal{X}}} - \hat{\mathcal{H}} \left| \hat{\mathcal{U}} \right|^{\frac{1}{n}-1} \hat{\mathcal{U}} - \theta \mathcal{C} \mathcal{H}_g^{\frac{1-mn-n}{n}} \left| \hat{\mathcal{U}} \right|^{m-1} \hat{\mathcal{U}}$$
$$= (1 - (1-\theta)r)\, \hat{\mathcal{H}} \hat{\mathcal{H}}_{\hat{\mathcal{X}}} \tag{3.82a}$$

$$(\hat{\mathcal{U}}\hat{\mathcal{H}})_{\hat{\mathcal{X}}} = -(1-\theta)\mathcal{H}_g^{\frac{n+1}{n}} \mathcal{M}, \tag{3.82b}$$

where $\theta = 1$ if $\hat{\mathcal{H}} \geq 1$ and $\theta = 0$ otherwise. The matching conditions become

$$\lim_{\hat{\mathcal{X}} \to -\infty} \hat{\mathcal{U}}\hat{\mathcal{H}} \sim 1, \quad |\hat{\mathcal{U}}|^{\frac{1}{n}-1}\hat{\mathcal{U}} \sim \hat{\mathcal{U}}^{-2}\hat{\mathcal{U}}_{\hat{\mathcal{X}}}, \quad \text{and} \quad \lim_{\hat{\mathcal{X}} \to -\infty} \hat{\mathcal{U}} = 0. \tag{3.83}$$

At leading order, $\hat{\mathcal{H}}_g^{-\frac{(n+1)^2}{n^2}} \ll 1$ and we neglect the first term of (3.82a). Therefore, over the floating portion ($\theta = 0$) the flux has the form

$$\hat{\mathcal{U}}\hat{\mathcal{H}} = -(1-r)^n \hat{\mathcal{H}} \left|\hat{\mathcal{H}}_{\hat{\mathcal{X}}}\right|^{n-1} \hat{\mathcal{H}}_{\hat{\mathcal{X}}} \tag{3.84}$$

Integrating the rescaled mass balance equation (3.82b) over the grounded portion ($\hat{\mathcal{X}} < 0$ to $\hat{\mathcal{X}} = 0$ and $\theta = 1$) implies that

$$\hat{\mathcal{H}}(0)\hat{\mathcal{U}}(0) = 1. \tag{3.85}$$

Assuming continuity in $\hat{\mathcal{H}}$ at the grounding line $\hat{\mathcal{X}} = 0$, implies that $\hat{\mathcal{H}}(0) = 1$ since this is the floating condition. Integrating (3.82b) over the floating portion implies that

$$\hat{\mathcal{U}}\hat{\mathcal{H}} = 1 - \mathcal{M}\hat{\mathcal{H}}_g^{\frac{n+1}{n}}\hat{\mathcal{X}} = -(1-r)^n \hat{\mathcal{H}} \left|\hat{\mathcal{H}}_{\hat{\mathcal{X}}}\right|^{n-1} \hat{\mathcal{H}}_{\hat{\mathcal{X}}}. \tag{3.86}$$

Therefore,

$$\hat{\mathcal{H}}^{\frac{n+1}{n}} = 1 + (1-r)^{-1}\mathcal{M}^{-1}\mathcal{H}_g^{-\frac{n+1}{n}} \left(\left(1 - \mathcal{M}\mathcal{H}_g^{\frac{n+1}{n}}\mathcal{X}\right)^{\frac{n+1}{n}} - 1\right). \tag{3.87}$$

We need to find an expression for the flux that also satisfies the terminal boundary conditions. In order to enforce the boundary conditions, we see that there is another boundary layer at the shelf terminus where the extensional stress again plays a leading-order role. We have decomposed the original boundary layer problem into an outer and inner region.

We rescale inside the second boundary layer so that the extensional stress is non-negligible. The rescaling is

$$\breve{\mathcal{H}} = \mathcal{H}_g^{-\frac{n+1}{n}}\hat{H}, \quad \breve{U} = \hat{U}, \quad \text{and} \quad \breve{X} = \mathcal{H}_g^{\frac{n+1}{n}}\left(\hat{\mathcal{X}} - \hat{\mathcal{X}}_c\right), \tag{3.88}$$

where $\hat{\mathcal{X}}_c$ is the terminus location. The rescaled stress and mass balance equations

become (noting $\frac{\partial}{\partial \breve{\mathcal{X}}} = \mathcal{H}_g^{\frac{n+1}{n}} \frac{\partial}{\partial \breve{\mathcal{X}}}$)

$$4\left(\breve{\mathcal{H}}\left|\breve{\mathcal{U}}_{\breve{\mathcal{X}}}\right|^{\frac{1}{n}-1}\breve{\mathcal{U}}_{\breve{\mathcal{X}}}\right)_{\breve{\mathcal{X}}} - \breve{\mathcal{H}}\left|\breve{\mathcal{U}}\right|^{\frac{1}{n}-1}\breve{\mathcal{U}} = (1-r)\breve{\mathcal{H}}\breve{\mathcal{H}}_{\breve{\mathcal{X}}} \tag{3.89a}$$

$$(\breve{\mathcal{U}}\breve{\mathcal{H}})_{\breve{\mathcal{X}}} = -\mathcal{H}_g^{\frac{n+1}{n}}\mathcal{M}, \tag{3.89b}$$

$$\tag{3.89c}$$

with boundary conditions

$$\breve{\mathcal{H}}\left|\breve{\mathcal{U}}_{\breve{\mathcal{X}}}\right|^{\frac{1}{n}-1}\breve{\mathcal{U}}_{\breve{\mathcal{X}}} \to 0 \quad \text{and} \quad \breve{\mathcal{U}}\breve{\mathcal{H}} \to 0 \quad \text{as} \quad \breve{\mathcal{X}} \to 0. \tag{3.90}$$

Matching this inner problem with the outer problem (3.87) requires the flux and thickness to vanish approaching the terminus—

$$\hat{\mathcal{U}}\hat{\mathcal{H}} \to 0, \quad \hat{\mathcal{H}} \to 0, \quad \text{and} \quad \breve{\mathcal{H}}\left|\breve{\mathcal{U}}\right|^{\frac{1}{n}-1}\breve{\mathcal{U}} \sim -(1-r)\breve{\mathcal{H}}\breve{\mathcal{H}}_{\breve{\mathcal{X}}} \tag{3.91}$$

as $\hat{\mathcal{X}} \to \hat{\mathcal{X}}_c$ and $\breve{\mathcal{X}} \to -\infty$. Requiring the flux (3.93) to vanish at $\hat{\mathcal{X}}_c$, implies that $\hat{\mathcal{X}}_c = \mathcal{M}^{-1}\hat{\mathcal{H}}_g^{-\frac{n+1}{n}}$ and requiring $\hat{\mathcal{H}}(\hat{\mathcal{X}}_c) = 0$ using (3.87), implies that

$$1 = (1-r)^{-1}\mathcal{M}^{-1}\mathcal{H}_g^{-\frac{n+1}{n}}. \tag{3.92}$$

We can solve for Q—

$$Q = ((1-r)M)^{\frac{n}{n+1}}WH_g. \tag{3.93}$$

3.2 Solutions to the boundary layer model

The boundary layer connects to the 'outer' problem (3.1) via asymptotic matching [Holmes, 2012]. The flux into the boundary layer from the outer problem Q is the solution of the boundary layer problem. Matching conditions (3.18) require that the outer problem's flux near the boundary layer is equal to the flux in the upstream

end of the boundary layer, the extensional stress gradients becomes negligible at the up stream end of the boundary layer, and velocities in the boundary layer are large compared to the glacier's interior.

Finding the flux Q into the boundary layer amounts to finding an orbit of the second-order system (3.13) that satisfies the boundary conditions (either (3.16) or (3.17)) and the matching conditions (3.18) along its trajectory. In section 3.1.2, we show that the flux $Q(W, M, H_g, \Gamma, \Lambda)$ is implicitly defined as the root of a function

$$F(Q; W, M, H_g, \Gamma, \Lambda) = 0. \tag{3.94}$$

Given the choice of calving law, calving- or melt-limited, reparameterizing reduces the complexity of this formula and allows us to analyze the flux as a function of fewer parameters. For example, we can analyze how the flux depends on bedrock depth by varying only two parameters. The melt-limited case is unaffected by Λ and, therefore, we reparameterize the boundary layer problem in terms of $Q(WM)^{-1}$, $H_g M^{\frac{-1}{n+1}}$, and $W^{\frac{m+mn+n+1}{n+1}} M^{\frac{nm-2}{n+1}} \Gamma$. Similarly, in the calving-limited case, we reparameterize the boundary layer problem in terms of $Q(W\Lambda^{n+1})^{-1}$, $H_g\Lambda^{-1}$, and $M\Lambda^{-(n+1)}$, assuming that $\Gamma = 0$. The rootfinding problems become

$$\begin{aligned} F\left(Q(WM)^{-1}, H_g M^{-\frac{1}{n+1}} W^{\frac{m+mn+n+1}{n+1}}, M^{\frac{nm-2}{n+1}}\Gamma)\right) &= 0 \qquad (3.95a) \\ F\left(Q(W\Lambda^{n+1})^{-1}, H_g\Lambda^{-1}, M\Lambda^{-(n+1)}\right) &= 0 \qquad (3.95b) \end{aligned}$$

in the melt- and calving- limited cases, respectively.

Collapsing solutions to the boundary layer problem onto a three parameter family, as presented in (3.95), allows us to investigate how the flux Q is related to input parameters. For example, we can plot $Q(WM)^{-1}$ or $Q(W\Lambda^{n+1})^{-1}$ against $H_g M^{-\frac{1}{n+1}} W^{\frac{m+mn+n+1}{n+1}}$ or $H_g\Lambda^{-1}$ given fixed values of $M^{\frac{nm-2}{n+1}}\Gamma$ or $M\Lambda^{-(n+1)}$ to show, roughly speaking, how the flux Q varies as a function of the flotation thickness H_g. The flotation thickness is a proxy for bedrock depth—the two are related via

$H_g = -r^{-1}B(0)$. We explore how the flux varies in the calving- and melt-limited cases in response to changes in flotation thickness H_g and subshelf melt M.

3.2.1 Comparison to the full model

First, we confirm the boundary layer result by comparing solutions to those computed using the full model (2.16) as $\varepsilon \to 0$. Different solutions to the steady state problem are obtained by varying the surface mass balance a. For each solution, we compare the flux across the grounding line to the boundary layer problem's solution using the parameters at that location. The result is shown in figure 3-1. As $\varepsilon \to 0$, the fluxes computed using the full model are indistinguishable from those computed using the boundary layer analysis.

3.2.2 Melt-limited extreme

We begin by restricting our analysis to the melt-limited extreme, where mass is only lost via melt along the ice shelf base. We show that calving-at-flotation is merely a limiting case of the melt limited extreme where $\mu \to \infty$ and that both scenarios are subject to the marine ice instability.

In the calving-at-flotation case and melt-limited extreme, the grounding line flux Q is a monotonically increasing function of bedrock depth (figure 3-2). The calving-at-flotation case is the same as [Schoof, 2007b, Schoof, 2007a]. The dotted brown lines in figure 3-2 (and figure 3-6) shows that, in the calving-at-flotation case, the grounding line flux is well approximated by (3.51). We have extended previous work of [Schoof, 2007b, Schoof, 2007a] to the melt-limited case. Both calving-at-flotation and melt-limited glaciers are subject to the marine ice instability: if the glacier retreats along a bed that slopes downward toward the ice divide the flux increases, accelerating retreat.

Calving-at-flotation is a special case of the melt-limited extreme with an infinite melt rate. Taking $M \to \infty$, we see that (3.76) becomes (3.51). This result was derived formally in section 3.1.3. Figure 3-2 shows that decreasing the melt rate to a finite

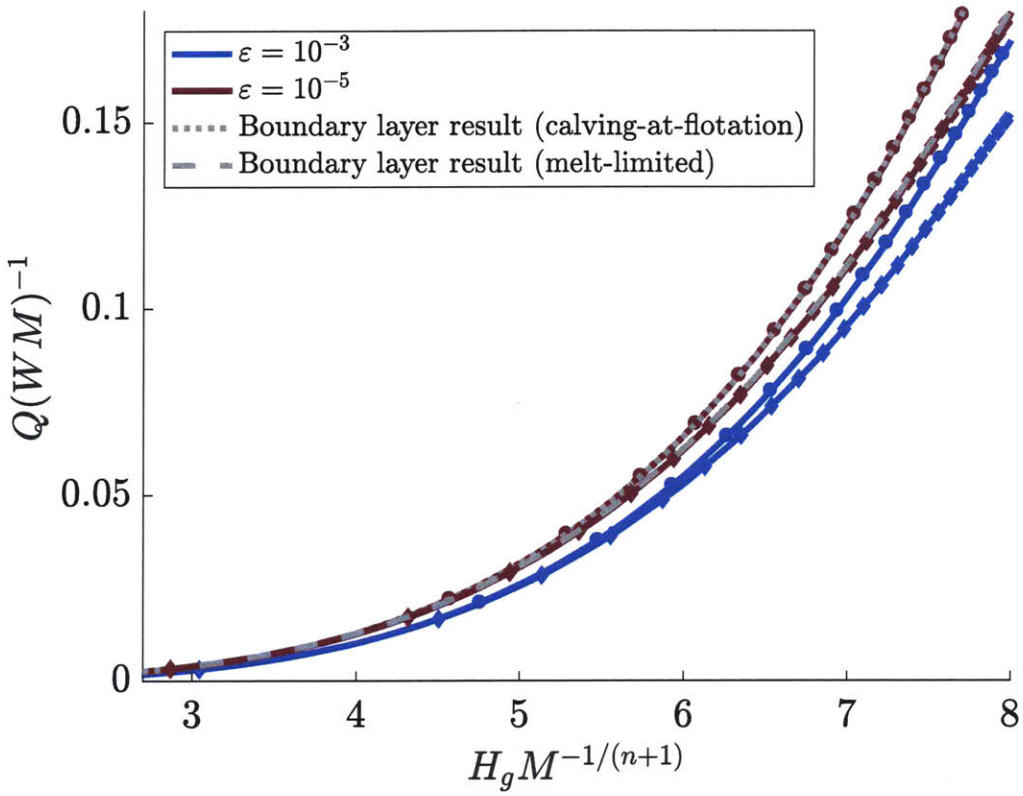

Figure 3-1: Comparison of the solution to the boundary layer problem (3.95) to solutions of the full model (2.16) with $\varepsilon = 10^{-3}$ (blue lines) and $\varepsilon = 10^{-5}$ (red lines). Circles and diamonds show solutions to (2.16) using the calving-at-flotation law and in the melt-limited extreme, respectively. Whereas, dotted and dashed lines show solutions to the boundary layer problem in the calving-at-flotation case and melt-limited extreme. Steady states are computing for the full model, varying the surface mass balance a to force grounding line migration. We set $w = 1$, $n = 3$, $m = n^{-1}$, $r = 0.9$, $\mu = 5\varepsilon^{-\frac{n}{n+1}}$, $\gamma = \frac{1}{4}\varepsilon^{-\frac{n-mn^2-n^2}{(n+1.0)^2}}$, and $b(x) = 3\varepsilon^{\frac{n^2}{(n+1)^2}}(x^2-4)$ and plot $Q(WM)^{-1} = \varepsilon^{-\frac{n}{n+1}}u(x_g)h(x_g)(w(x_g)\mu)^{-1}$ against $H_g M^{\frac{1}{n+1}} = r^{-1}\varepsilon^{-\frac{n}{n+1}}b(x_g)\mu^{-1}$.

value—which grows an ice shelf that buttresses the glacier—decreases the flux out of the grounded portion but the buttressing is not sufficient to stabilize the system; the flux still grows monotonically with bedrock depth.

In the melt-limited extreme, ice shelf buttressing decreases the rate that the grounding line flux grows as a function of bedrock depth. Requiring that $Q = 0$ at the terminus implies that the grounding line flux is $Q(0) = MX_c$, where X_c is the ice shelf length. Given the grounding line flux via (3.93), increasing the bedrock depth

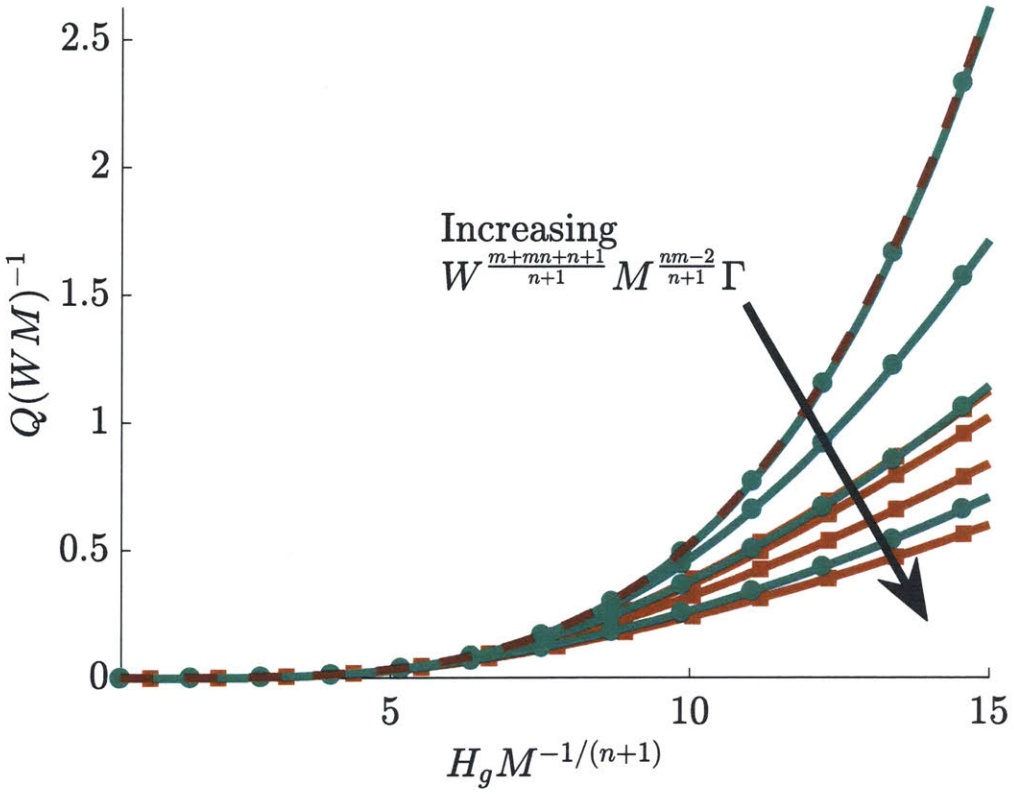

Figure 3-2: Here, we show the relationship between a proxy for grounding line flux $Q(WM)^{-1}$ and a proxy for bedrock depth at the grounding line $H_g M^{-\frac{1}{n+1}}$; H_g is the flotation thickness. We increase the basal friction parameter Γ by setting $W^{\frac{m+mn+n+1}{n+1}} M^{\frac{nm-2}{n+1}} \Gamma \in \{0, 5, 25, 125\}$. The dotted brown line is (3.51) with $\Gamma = 0$.

at the grounding line will also increase the ice shelf length. The degree to which the glacier is buttressed depends on the side-wall drag along the ice shelf; longer shelves induce a larger back-stress that buttresses the glacier. For an *unbuttressed* glacier (3.51) implies $Q \propto H_g^{n+1}$—a result that is consistent with [Schoof, 2007b, Schoof, 2007a]. In our current setting, this corresponds to calving-at-flotation—equivalently $M \to \infty$—and the green lines in figures 3-2 reflect this trend. Decreasing M to a finite value—the orange lines in figure 3-2—requires the glacier to have an ice shelf that buttresses the system and, therefore, decreases the grounding line flux. In the other extreme $M \to 0$, (3.93) suggests that the flux increases linearly with H_g. Figure 3-3 suggests that the limiting behavior as the subshelf melt decreases, approximated by (3.93), is only accurate for very small melt rates. However, the two extremes $M \to \infty$

and $M \to 0$ suggest that the $Q \propto H_g^\alpha$ where $\alpha \in [0, n+1]$ depends on M. Ice shelf buttressing decreases the rate Q increases with H_g but does not change the relationship that deeper beds lead to greater grounding line flux. When the grounding line of a melt-limited glacier retreats into deeper water two processes trigger the marine ice instability: (i) the shorter grounded portion has a smaller net in-flux and (ii) the grounding line flux increases.

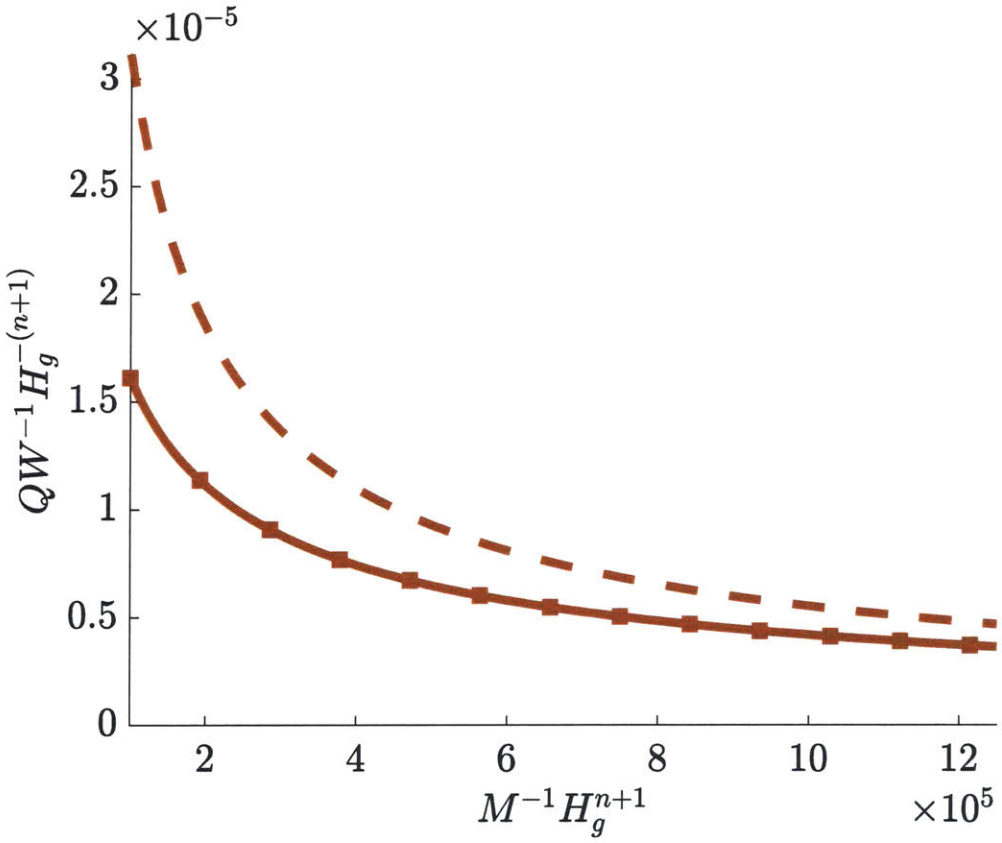

Figure 3-3: Setting $\Gamma = 0$, we show the relationship between a proxy for the grounding line flux and the subshelf melt. The dotted orange line is (3.93), which approximates the grounding line flux when the ice shelf is long—equivalently, when the sub-shelf melt is small ($M \ll 1$).

3.2.3 Calving-limited termini

In the meltwater calving case, the flux is no longer a monotonic function of bedrock depth. Figure 3-4 (green line) shows that the flux can actually *decrease* as a function of

bedrock depth. This contrasts the calving-at-flotation case and melt-limited extreme, where grounding line flux always increases with bedrock depth. This anomalous relationship means the glacier is not subject to the marine ice instability: retreating into deeper water could decrease the flux and lead to re-advance.

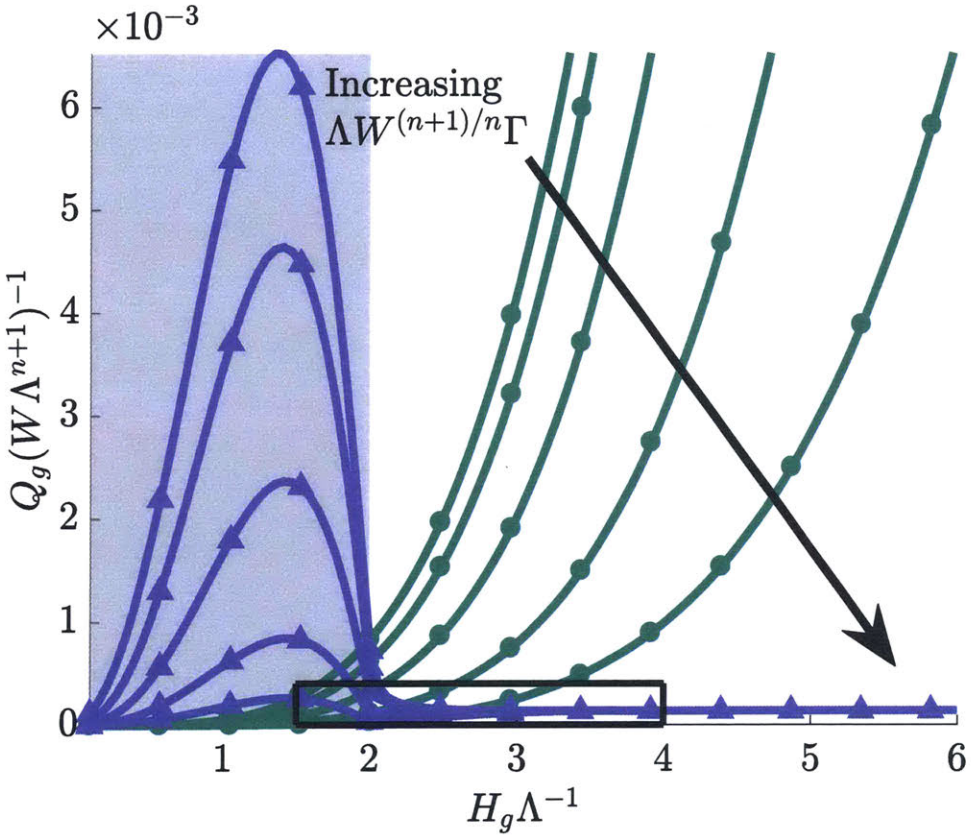

Figure 3-4: A proxy for the grounding line flux Q as a function of the flotation thickness H_g. Here, we set $M\Lambda^{-(n+1)} = 0$ and $W^{\frac{n+1}{n}}\Lambda^{-1}\Gamma \in \{0, 1, 5, 25, 125\}$. The shaded and unshaded regions corresponds to grounded and floating calving fronts, respectively. The green and purple lines are the grounding line flux computed by finding roots of (3.33) in the calving-at-flotation and meltwater calving cases, respectively.

No subshelf melt

We begin by restricting our analysis to the meltwater calving law with no subshelf melt. Figures 3-4 and 3-5 show that Q decreases as a function of H_g near the critical flotation thickness $H_g\Lambda^{-1} = 2$. We consider the grounding and floating cases separately—the flux at a grounded terminus always decreases as a function of

$H_g = -r^{-1}B$ and, in the floating case, the behavior depends on the lateral friction Γ. However, in either case the behavior of the flux depends on the role of extensional stress at the grounding line.

We begin by expressing the flux Q as a function of extensional stress $\Sigma = 4|U_X|^{\frac{1-n}{n}}U_X$, and floating thickness[3] H_g. We assume that Σ is relatively small—specifically, $\Sigma \ll H$. This assumption is justified when $(1-\delta)$ is small. For example, in the calving-at-flotation case $\Sigma_g = \frac{1-\delta}{2}H_g$. From (3.13b), $H_x = -HU_XU^{-1}$ and, neglecting gradients of $H\Sigma$, (3.13a) implies

$$|U|^{\frac{n+1}{n}} \approx \frac{H^2U_X}{\Gamma + W^{-\frac{n+1}{n}}H}. \tag{3.96}$$

Therefore, applying $H = H_g$ at $X = 0$, the flux is

$$Q = \left(\frac{1}{4}\right)^{\frac{n^2}{n+1}} \frac{H_g^{\frac{3n+1}{n+1}}\Sigma_g^{\frac{n^2}{n+1}}}{\left(\Gamma + W^{-\frac{n+1}{n}}H_g\right)^{\frac{n}{n+1}}} \tag{3.97}$$

This result is formally derived in section 3.1.3. This formula is an extension from [Schoof, 2007b, Schoof, 2007a], modifying their result to include lateral stress.

In the calving-at-flotation case, (3.97) provides a formula for the flux in terms of the width, lateral friction, and bedrock depth. Hydrostatic balance implies that $\Sigma_g = \frac{1-\delta}{2}H_g$, therefore

$$Q = \left(\frac{1-r}{8}\right)^{\frac{n^2}{n+1}} \frac{H_g^{\frac{n^2+3n+1}{n+1}}}{\left(\Gamma + W^{-\frac{n+1}{n}}H_g\right)^{\frac{n}{n+1}}}. \tag{3.98}$$

Again, this result is formally derived in section 3.1.3 and is a generalization of a previously mentioned result by [Schoof, 2007b, Schoof, 2007a] that $Q \propto H_g^{n+1}$, which holds if $\Gamma = 0$ (no lateral stress).

Grounded calving fronts We use (3.97) to explain why the flux is a decreasing function of bedrock depth when the glacier has a grounded terminus near flotation.

[3]Recall, the floating thickness $H_g = r^{-1}B$ is an explicit function of bedrock depth; we use the two interchangeably.

In this the case, the terminal thickness H_c is greater than $H_g = -\delta^{-1}B$. However, figure 2-3 shows that H_c increases with H_g. Equation (3.97) suggests that Q increases with H_c and extensional stress. Therefore, the observed decrease in Q requires that Σ_g decreases as H_g increases.

Restricting our analysis to the calving-limited glaciers, figures 3-4 and 3-5 show that Q_g decreases as a function of H_g near the critical flotation thickness $H_g\Lambda^{-1} = 2$. We consider the grounding and floating cases separately—the flux at a grounded terminus always decreases as a function of $H_g = -r^{-1}B$ and, in the floating case, the behavior depends on the lateral friction Γ. However, in either case the behavior of the flux depends on the role of extensional stress at the grounding line.

We perturb the floating thickness near $H_g\Lambda^{-1} = 2$ to understand the behavior of the ice flux around the flotation condition. Recall, the terminal thickness is $H_c = rH_g\phi(\Lambda H_g^{-1})$. Choosing bedrock depth such that $H_{g0} = 2\Lambda$ implies $H_c = H_{g0}$. We perturb the bedrock depth $H_{g0} + \hat{H}_g$ and use a first-order Taylor expansion of ϕ to compute the perturbed terminal thickness

$$H_c = \delta\left(H_{g0} + \hat{H}_g\right)\phi\left(\frac{\Lambda}{H_{g0} + \hat{H}_g}\right) \tag{3.99}$$

$$\approx \delta\left(H_{g0} + \hat{H}_g\right)\left(\phi\left(\frac{\Lambda}{H_{g0}}\right) - \phi'\left(\frac{\Lambda}{H_{g0}}\right)\frac{\Lambda\hat{H}_g}{(H_{g0} + \hat{H}_g)^2} + \mathcal{O}\left(\hat{H}_g^2\right)\right) \tag{3.100}$$

where ϕ' indicates the derivative. Given the form of ϕ and that $\Lambda H_{g0}^{-1} = 0.5$, $\phi(\Lambda H_{g0}^{-1}) = \delta^{-1}$ and $\phi'(\Lambda H_{g0}^{-1}) = 2\delta^{-1}$,

$$H_c = H_{g0} + \mathcal{O}\left(\hat{H}_g^2\right). \tag{3.101}$$

Given the meltwater calving law, a small perturbation in bedrock depth has only second-order effects on the thickness at the calving front.

The extensional stress, however, is sensitive to perturbations in bedrock depth. Let $\hat{H}_g < 0$—the bedrock depth is perturbed such that the terminus remains grounded. The hydrostatic stress condition (3.16) implies that, to first order in $\hat{H}_g$,

$$\Sigma_g \approx \frac{1-\delta}{2} H_{g0} - \hat{H}_g. \tag{3.102}$$

Since $\hat{H}_g < 0$, we observe a first order increase in the extensional stress and, assuming $1 - \delta \ll 1$, small changes in $\hat{H}_g$ can have relatively significant impacts on Σ_g. For example, assuming $H_{g0} \in \mathcal{O}(1)$, perturbing $\hat{H}_g$ by 0.05 causes the stress perturbation to be the same order of magnitude Σ_g at the critical flotation thickness.

The change in extensional stress occurs because the terminal thickness has not changed but the bedrock is shallower. The calving cliff protrudes farther out of the ocean and the depth-averaged normal stress exerted by the ocean is smaller. Therefore, the extensional stress increases. For constant H_g and increasing Σ_g, (3.97) shows that the flux will increase. Figures 3-4 and 3-5 show this behavior. From the critical flotation thickness, the flux increases if the bedrock is perturbed such that the terminus remains constant.

Floating calving fronts Figure 3-5 shows a more complicated picture if the floating thickness is perturbed to have a floating terminus (i.e., $\hat{H}_g > 0$). Whether the flux is an increasing or decreasing function of bedrock depth depends on the basal friction. First, we note that if $\hat{H}_g > 0$ the extensional stress is also constant at first order. Specifically,

$$\Sigma(X_c) = \frac{1-\delta}{2} H_{g0} + \mathcal{O}\left(\hat{H}_g^2\right). \tag{3.103}$$

We use this fact to explain how basal friction affects the grounding line flux in the presence of an ice shelf.

Small basal friction When the basal friction is small, figure 3-5 shows that the grounding line flux is a decreasing function of bedrock depth. Suppose that the basal drag is not large enough to make the lateral drag insignificant over the grounded portion. The ice at the grounding line is slightly thicker than at the calving front, however, the driving stress in the floating ice shelf is small compared to the other forces. Most of the ice shelf is supported by the ocean—the thinning is largely

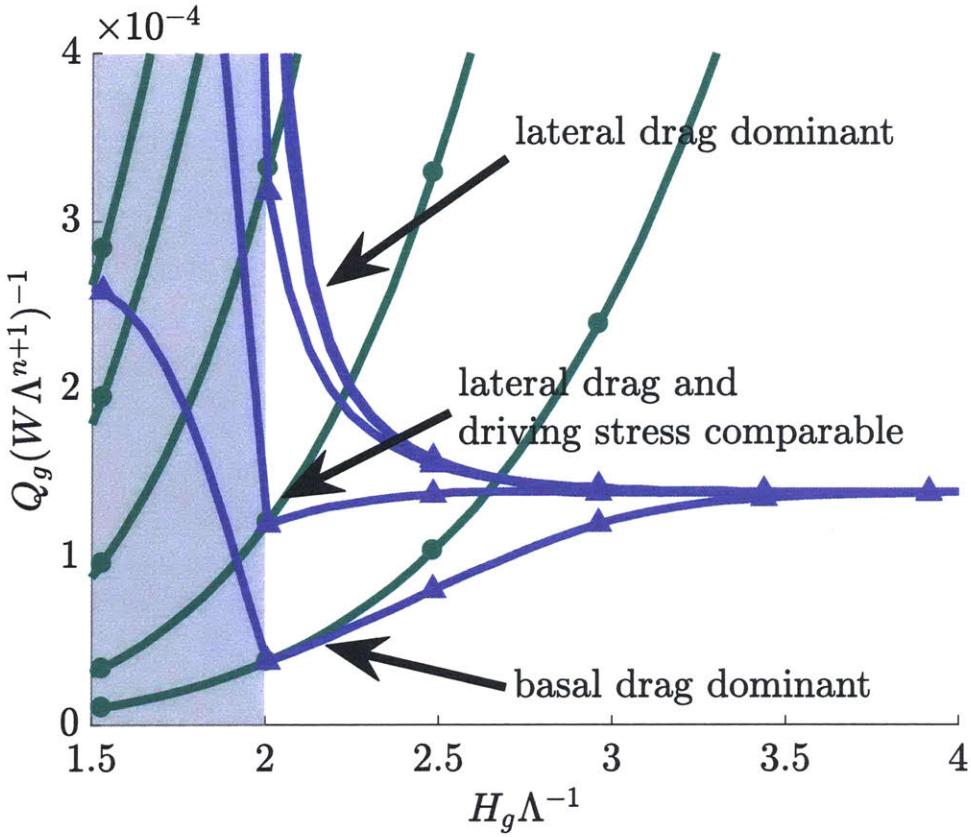

Figure 3-5: Zoomed-in version figure 3-4—we are showing the portion in the grey box. The basal stress' magnitude determines the behavior of flux for small ice shelves.

accounted for by the bottom sloping upward. The surface slope, creates the driving stress due to gravity, is only $(1-\delta)\delta^{-1}$ times the bottom slope. Therefore, the dominate forces in the ice shelf are the depth-integrated extensional stress and the lateral drag; (3.13a) becomes

$$4\left(H\left|U_X\right|^{\frac{1}{n}-1}U_X\right)_X = W^{-\frac{1}{n}-1}H\left|U\right|^{\frac{1}{n}-1}U. \tag{3.104}$$

Integrating this expression assuming $\hat{H}_g$ is small enough that $H \approx H_{g0}$ and $U = QH_{g0}^{-1}$ are constant

$$\Sigma(X) = \Sigma_c - W^{-\frac{n+1}{n}}Q^{\frac{1}{n}}H_{g0}^{\frac{n-1}{n}}(X_c - X) \tag{3.105}$$

$$= \frac{1-\delta}{2}H_{g0} - W^{-\frac{n+1}{n}}Q^{\frac{1}{n}}H_{g0}^{\frac{n-1}{n}}(X_c - X) \tag{3.106}$$

reveals that the ice shelf reduces the extensional stress relative to its value at the calving front. This is the degree to which the system is buttressed.

We need to know the ice shelf length X_c to derive an expression for the extensional stress at the grounding line $\Sigma_g = \Sigma(0)$. A larger flotation thickness requires a longer ice shelf for the glacier to thin to the calving thickness, potentially increasing the amount of buttressing. Assuming $A \ll 1$ and $M = 0$, (3.13b) implies that the rate the ice thins is

$$-H_X = \frac{HU_X}{U} \approx \frac{H_{g0}^2 \Sigma^n}{4^n Q}. \tag{3.107}$$

Therefore, the ice shelf length is computed by solving[4]

$$\hat{H}_g = \int_0^{X_c} -H_X \, dX = \frac{H_{g0}^2}{4^n Q} \int_0^{X_c} \Sigma^n \, dX \tag{3.108}$$

for X_c. Substituting this into (3.106) derives an expression for the extensional stress at the grounding line

$$\Sigma_g = \Sigma(0) = \left(\left(\frac{1-\delta}{2} H_{g0} \right)^{n+1} - \frac{(n+1)4^n Q^{\frac{n+1}{n}}}{W^{\frac{n+1}{n}} H_{g0}^{\frac{n+1}{n}}} \hat{H}_g \right)^{\frac{1}{n+1}} \tag{3.109}$$

Substituting this into (3.97) derives Q implicitly. Furthermore, this shows that a decrease in $\hat{H}_g$ results in a decrease in Σ_g and, therefore, a decrease in the flux Q. The extensional stress decreases because the ice shelf lengths as $\hat{H}_g$ increases and the total amount of lateral drag on the ice shelf increases. Again, we have more formally derived these results in section 3.1.3.

Large basal friction Conversely, when the basal friction is large, figure 3-5 shows that the grounding line flux is an increasing function of bedrock depth. When the basal drag is large, ice velocities near the grounding line become small, which (i) reduces the velocity-dependent lateral drag term and (ii) increases the thickness gradient and, therefore, the driving stress. The driving stress can no longer be neglected

[4]We compute the ice shelf length by substituting (3.106) into (3.108) and solving for X_c.

as in previous case.

Balancing the basal and driving stresses estimates how large Γ must be in order for the driving stress to appear at leading order in the boundary layer. The thickness, velocity, and stress are continuous across the grounding line and, since

$$H_X = -\frac{HU_X}{U} = -\frac{H|\Sigma|^{n-1}\Sigma}{4^n U}, \tag{3.110}$$

so is H_X. However, the surface slope is not continuous. For sufficiently large Γ, the driving stress at the grounding line balances the basal stress $-HH_X \sim \Gamma|U|^{\frac{1}{n}-1}U$. To appear at leading order, the driving stress along the ice shelf should be comparable to lateral drag $-(1-\delta)HH_X \sim W^{-\frac{n+1}{n}}H|U|^{\frac{1}{n}-n}U$. Therefore, (recall, at the grounding line $H \sim 2\Lambda$)

$$\Gamma W^{\frac{n+1}{n}}H^{-1} \sim \Gamma W^{\frac{n+1}{n}}(2\Lambda)^{-1} \sim (1-\delta)^{-1} \tag{3.111}$$

When $\delta = 0.9$, this requires $\Gamma W^{\frac{n+1}{n}}\Lambda^{-1} \sim 20$, which is consistent with 3-5, where the lateral and driving stresses approximately balance when $\Gamma W^{\frac{n+1}{n}}\Lambda^{-1} = 25$.

Finally, consider the limiting case of very large basal friction ($\Gamma W^{\frac{n+1}{n}}\Lambda^{-1} \gg (1-\delta)^{-1}$). Lateral drag become negligible compared to the large basal drag and the force balance simply requires

$$(H\Sigma)_X = (1-\delta)HH_X \tag{3.112}$$

as in [Schoof, 2007b, Schoof, 2007a]. Integrating and applying the hydrostatic boundary condition suggests that $\Sigma_g = \frac{1-\delta}{2}H_g$, which we substitute into (3.97) to estimate the flux. As expected, the result is a scaled of the result from [Schoof, 2007b, Schoof, 2007a], where flux increases monotonically with bedrock depth. However, if the ice shelf is sufficiently long, the lateral stress cannot be neglected, even with large basal friction.

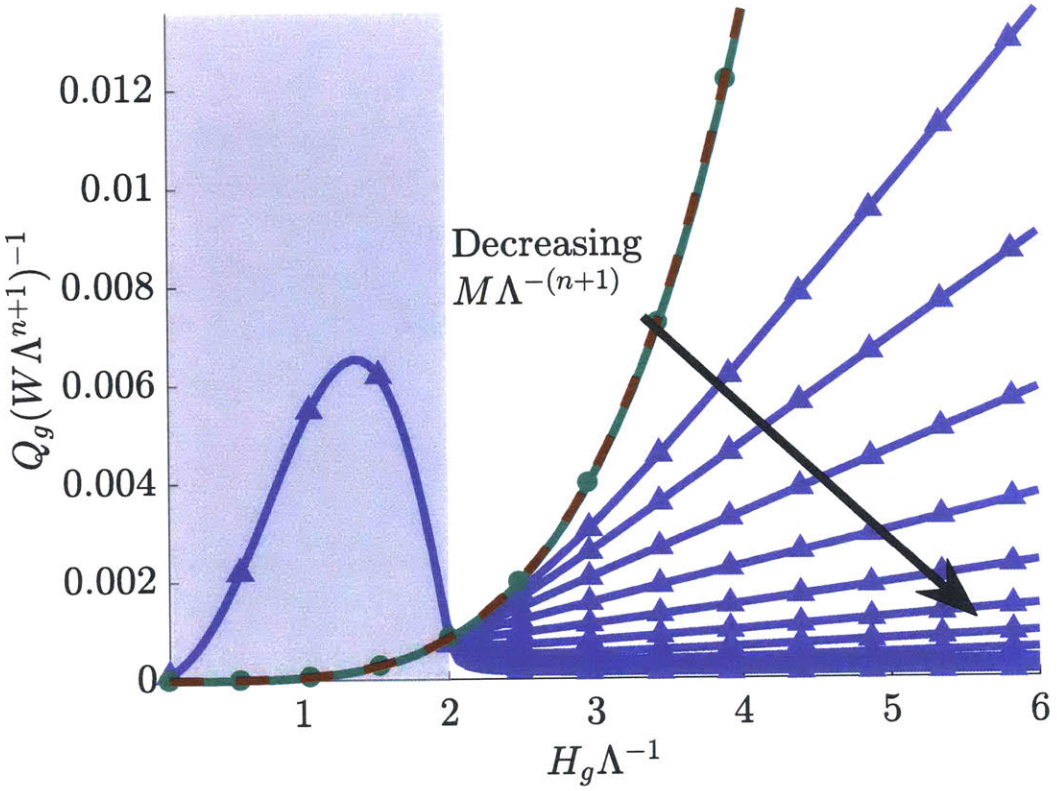

Figure 3-6: A proxy for grounding line flux as a function of a proxy for the flotation thickness. The green line is the solution to the boundary layer problem given calving-at-flotation and the dotted brown line is (3.51). The purple lines are the solution to the boundary layer problem in the meltwater calving case given $M\Lambda^{-(n+1)} = 0$ and $M\Lambda^{-(n+1)} = 2^{1\times 10^{-5}j}$ for $j = 1:10$. The shaded and unshaded regions corresponds to grounded and floating calving fronts, respectively. The dotted brown line is (3.51) with $\Gamma = 0$.

3.2.4 Effects of subshelf melt

Including sub-shelf melt links the grounding line flux to the bedrock depth via a similar mechanism as the melt limited extreme. When M is large enough, most of the flux into the floating ice shelf is lost to sub-shelf melt. We assume $Q_c \ll (M-A)X_c$ and, given (3.113), $Q_g \approx (M-A)X_c$. If the terminal thickness is small enough—$\Lambda \ll 1$—then the system is well approximated by the melt-limited case and (3.93) predicts that Q_g is a linear function of H_g. Figure ?? shows this relationship for $M \gg 1$ and $\Lambda \ll 1$. More generally, increasing H_g requires the ice shelf length to

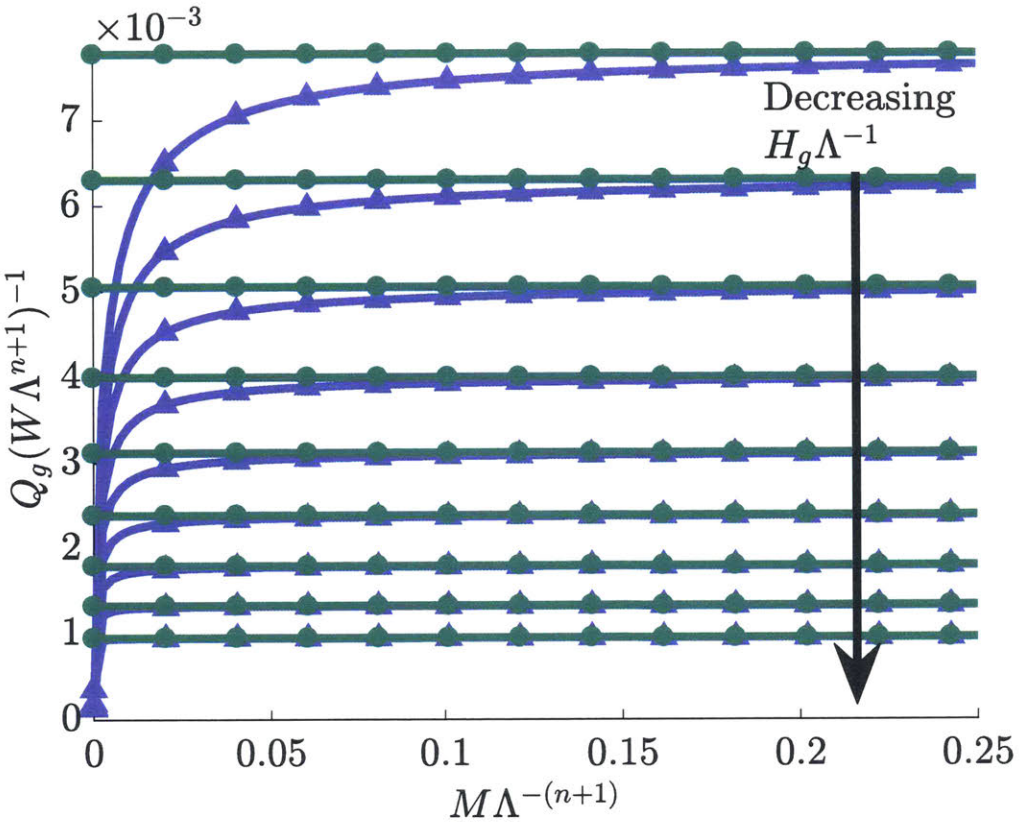

Figure 3-7: The grounding line flux Q_g as a function of sub-shelf melt M. Each subsequent level corresponds to $H_g\Lambda^{-1}$ linearly spaced between 2—which is the smallest value of $H_g\Lambda^{-1}$ that has an attached ice shelf—and 3.5. The green and purple lines are the grounding line flux computed by finding roots of (3.33) in the calving-at-flotation and meltwater calving cases, respectively.

increase, which also increases the mass loss due to sub-shelf melt. The grounding line flux must increase to offset this. Sufficiently small melt rates allow the ice shelf to completely buttress the glacier and Q_g is asymptotically constant. However, for large enough M, the Q_g must increase linearly with bedrock depth to account for the increased out-flux.

The grounding line flux—the volume of ice that enters the floating ice shelf—is removed from the glacier either by sub-shelf melt or a flux at the glacier terminus. By integrating (3.13b) from $X = 0$ to any point X along the shelf we see that

$$Q = Q_g - (M - A)X, \tag{3.113}$$

the flux decreases as ice is lost due to sub-shelf melt and the flux at the terminus is $Q_c = Q_g - (M - A)X_c$. This formalizes the notion that all of the mass lost from the ice shelf must be replaced by a flux at the grounding line. We emphasize that we have *not* assumed the glacier is in steady state. Rather, we have assumed that the boundary layer responds infinitely fast compared to the interior. For example, Q_g responds to a change in sub-shelf melt much faster than the grounded portion responds to the subsequent change in ice shelf buttressing.

When there is no mass loss/gain along the shelf—$(M - A) \ll 1$—the flux is constant along the ice shelf and insensitive to changes in bedrock depth at the grounding line. Previous work by [Hindmarsh, 2012, Pegler, 2016] show that as the bedrock deepens at the grounding line the flux approaches a finite limit that is independent of H_g—specifically,

$$Q_g \to (1 - r)^n C(n) W (2\Lambda)^{n+1}, \tag{3.114}$$

where $C(n)$ is a constant that depends on the flow law exponent. Essentially, for large enough H_g there is a large difference between the grounded thickness H_g and the terminal thickness H_c, which requires a long ice shelf. Along the ice shelf extensional stress gradients are weak except near the terminus—a boundary layer within the boundary layer. Since the flux is constant along the ice shelf, the flux at the terminus determined by the second boundary layer is also the flux at the grounding line. Figures 3-4 and 3-5 show that when $M \ll 1$ the flux is constant with respect to bed rock depth, as predicted by [Hindmarsh, 2012, Pegler, 2016].

As we increase M, the grounding line flux approaches the calving-at-flotation flux. Recall, that the calving-at-flotation case is a special case of the melt-limited calving law with infinite sub-shelf melt. As a function of bedrock depth, the calving-at-flotation case grows according to $Q \propto H_g^{n+1}$ [Schoof, 2007b, Schoof, 2007a]. However, increasing H_g with a finite melt rate also lengthens the ice shelf, which buttresses the glacier and slows the rate at which the flux grows (figure 3-6). If we fix H_g and compute Q as a function of the melt rate M, the calving-at-flotation flux is constant

(green lines in figure 3-7). For finite M, the glacier has a floating ice shelf that buttresses the glacier. Figure 3-7 shows that grounding line flux when $M < \infty$ is less than the calving-at-flotation case but it increases toward the calving-at-flotation limit as M increases.

The meltwater calving case may be able to stabilize the marine ice instability if the sub-shelf melt is sufficiently small. As before, when the flux is an increasing function of bedrock depth the glacier is subject to the marine ice instability. Figure 3-4 suggests the flux may be independent of H_g for sufficiently small melt rates. [Hindmarsh, 2012, Pegler, 2016] confirm that the flux approaches a constant value as $H_g \to \infty$ if there is no sub-shelf melt. However, as M increases the flux once again because an increasing function of H_g, which could potentially destabilize the system due to the marine ice instability.

3.3 Steady state comparison: full and reduced models

We compare steady-state solutions of both the full model (2.16) and the reduced model (3.1) computed using the methods from sections 2.4 and 3.1.1. In steady state, we find x_c to satisfy mass balance. The in-flux is determined by the net surface mass balance and the mass is lost either via subshelf melt or a flux at the terminus. Therefore, we find x_c such that

$$\int_0^{x_c} aw(x)\,dx = \int_{x_g}^{x_c} \mu w(x)\,dx + u(x_c)h(x_c)w(x_c). \tag{3.115}$$

Computing this mass balance depends on u and h—the solution to (2.16). However, solving for u and h requires x_c. We, therefore, employ an iterative method: given a guess for x_c, solve (2.16) for u and h and, then, given u and h find x_c such that (3.115) holds. We iterate until x_c stops changing. We use PETSc's implementation of Newton's method to find roots of (3.115) [Balay et al., 2017b, Balay et al., 2017a, Balay et al., 1997]. Given a terminus location x_c we solve (2.16) using a discontinuous

Galerkin implementation in the Fenics software package [Alnæs et al., 2015, Logg, 2007, Logg and Wells, 2010, Ølgaard et al., 2008].

We similarly compute solutions to the reduced model. We find the grounding line location x_g via mass balance

$$\int_0^{x_g} aw(x)\,dx = u(x_g)h(x_g)w(x_g) \approx w(x_g)q_g(x_g) \tag{3.116}$$

where $q_g(x)$ is the solution the boundary layer problem. Given the grounding line location x_g, we can compute the grounding line flux q_g *without* knowing the velocity u and thickness h fields. In steady state, we find x_g such that (3.116) holds. Given the steady state grounding line location x_g, we solve (3.1) using a continuous Galerkin implementation in the Fenics software package [Alnæs et al., 2015, Logg, 2007, Logg and Wells, 2010, Ølgaard et al., 2008]. The steady state solutions using the different methods result in the same solution (figures 3-8, 3-9, 3-10, 3-11, 3-12, and 3-13).

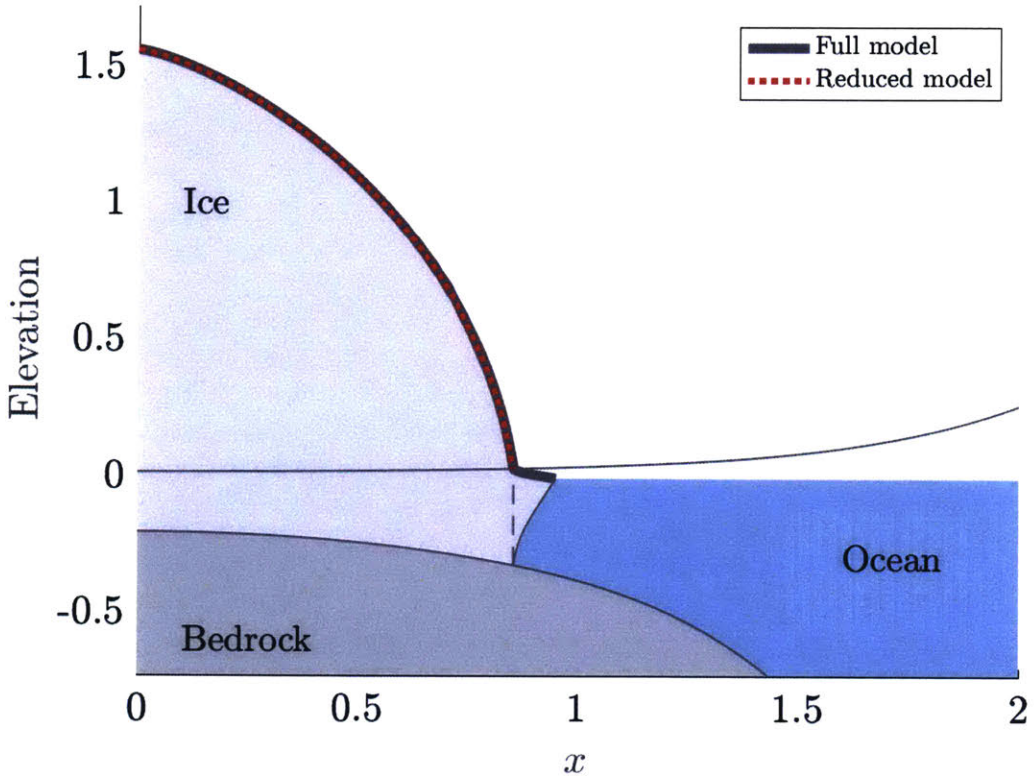

Figure 3-8: Steady state glacier profile for an upward ($b(x) = -0.25x^6 - 0.15x^2 - 0.2$) sloping bed given that mass is only lost via subshelf melt. The grey line is the steady state solution to the full model using the shooting method described in section 2.4 and the dotted red line is the steady state solution to the reduced model (3.116) and (3.1). The thin grey line is the flotation thickness $h_g(x) = -r^{-1}b(x)$. We have set $a = 0.8$ and $\mu = 8$.

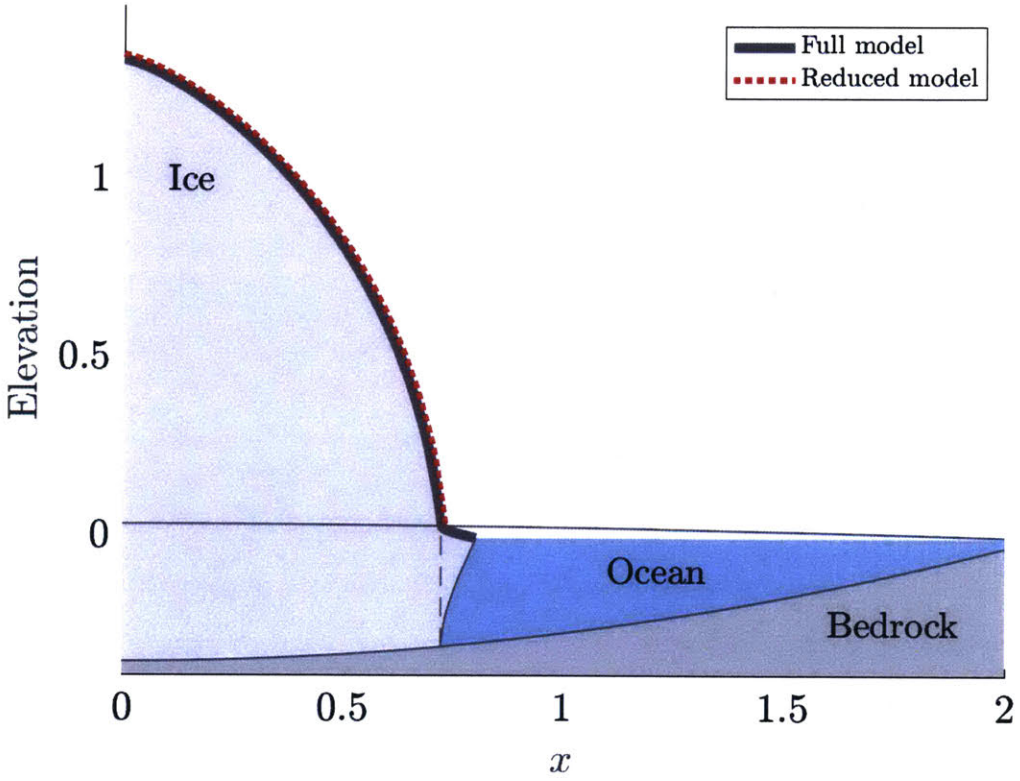

Figure 3-9: Steady state glacier profiles for a downward ($b(x) = 0.08x^2 - 0.25$) sloping bed given that mass is only lost via subshelf melt. The grey line is the steady state solution to the full model using the shooting method described in section 2.4 and the dotted red line is the steady state solution to the reduced model (3.116) and (3.1). The thin grey line is the flotation thickness $h_g(x) = -r^{-1}b(x)$. We have set $a = 0.8$ and $\mu = 8$.

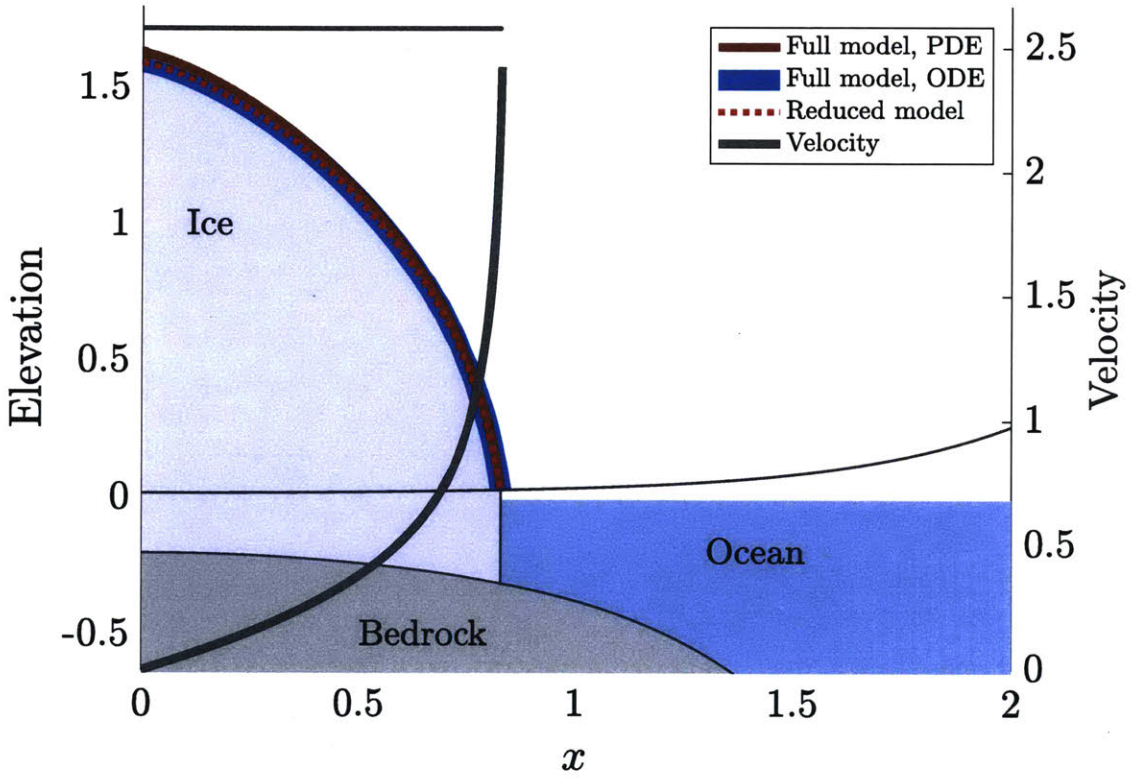

Figure 3-10: Steady state glacier profile for an upward ($b(x) = -0.25x^6 - 0.15x^2 - 0.2$) sloping bed given that the glacier calves at its critical flotation thickness. Solutions are computed using the shooting method from section 2.4 (blue line) and a discontinuous Galkerin method (brown and grey lines). The dotted red line is the steady state solution to the reduced model (3.116) and (3.1). The thin grey line is the flotation thickness $h_g(x) = -r^{-1}b(x)$. Grey dots along the top of each figure are vertexes of the triangulation after adaptive mesh refinement on the full model. We have set $a = 1$ and $w = 10$.

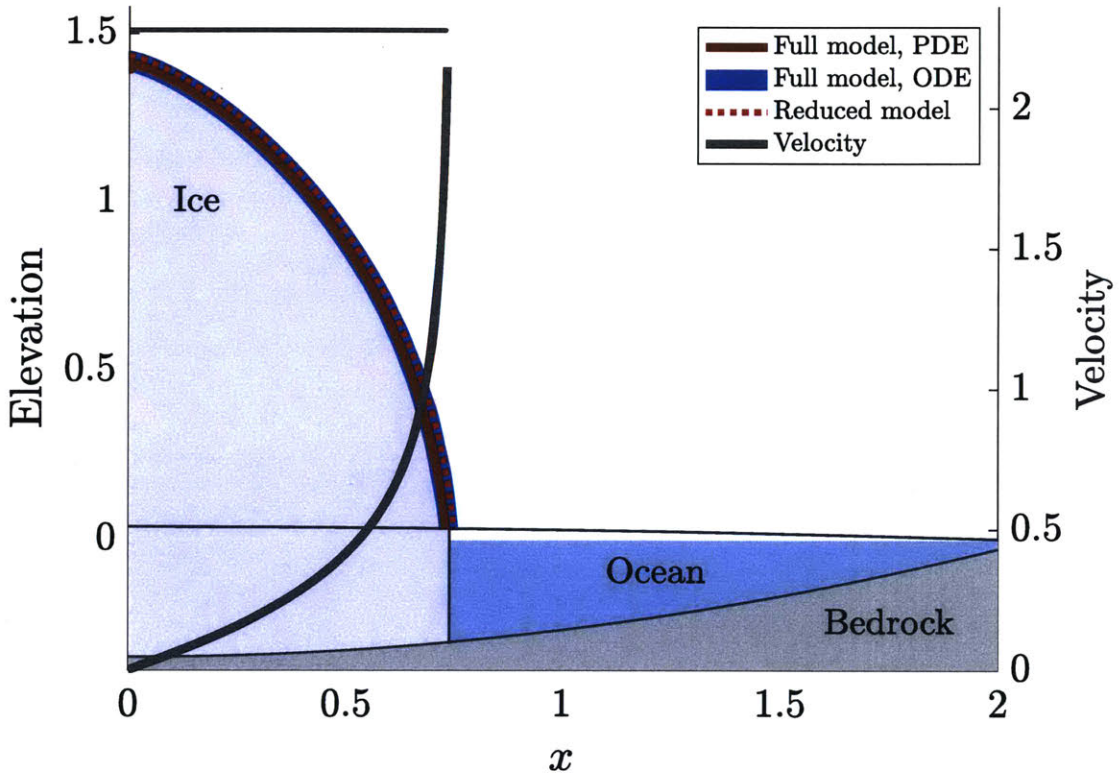

Figure 3-11: Steady state glacier profile for an upward ($b(x) = 0.08x^2 - 0.25$) sloping bed given that the glacier calves at its critical flotation thickness. Solutions are computed using the shooting method from section 2.4 (blue line) and a discontinuous Galkerin method (brown and grey lines). The dotted red line is the steady state solution to the reduced model (3.116) and (3.1). The thin grey line is the flotation thickness $h_g(x) = -r^{-1}b(x)$. Grey dots along the top of each figure are vertexes of the triangulation after adaptive mesh refinement on the full model. We have set $a = 1$ and $w = 10$.

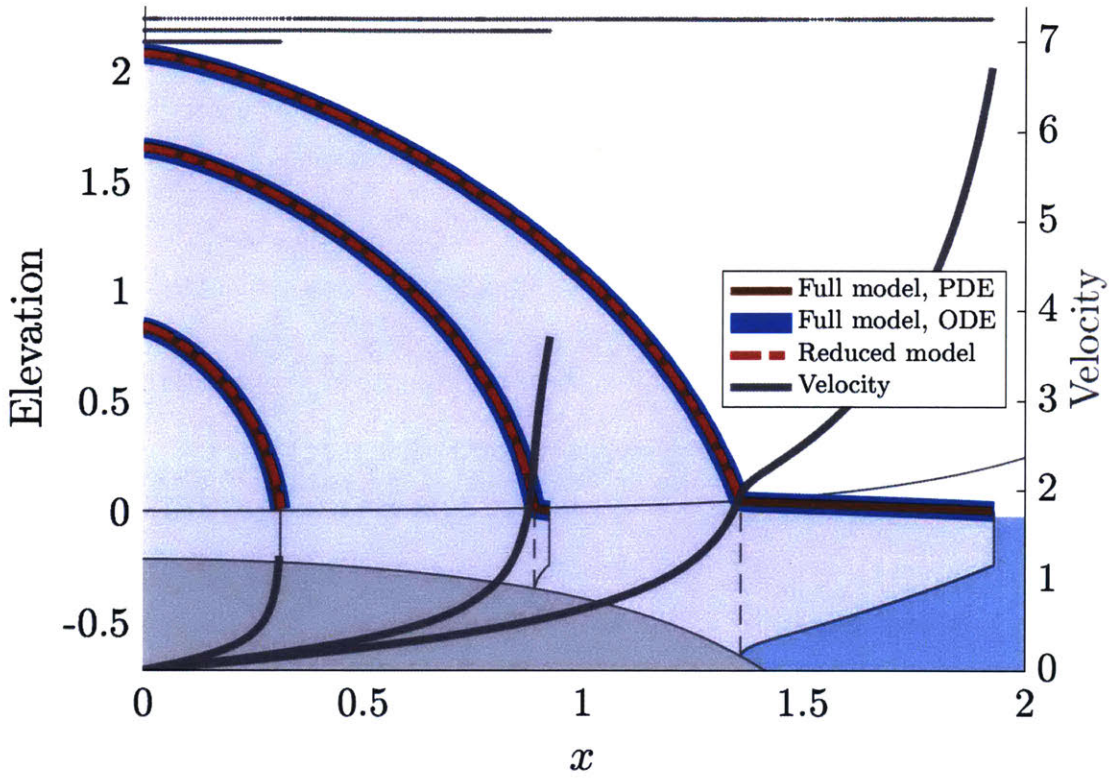

Figure 3-12: Steady state glacier profile for an upward ($b(x) = -0.25x^6 - 0.15x^2 - 0.2$) sloping bed given that the glacier calves at the critical thickness defined by the meltwater calving law. Solutions are computed using the shooting method from section 2.4 (blue line) and a discontinuous Galkerin method (brown and grey lines). The dotted red line is the steady state solution to the reduced model (3.116) and (3.1). The thin grey line is the flotation thickness $h_g(x) = -r^{-1}b(x)$. Grey dots along the top of each figure are vertexes of the triangulation after adaptive mesh refinement on the full model. We have set $a = 1$, $\mu = 0.5$, $\lambda = 0.11$, and $w = 10$.

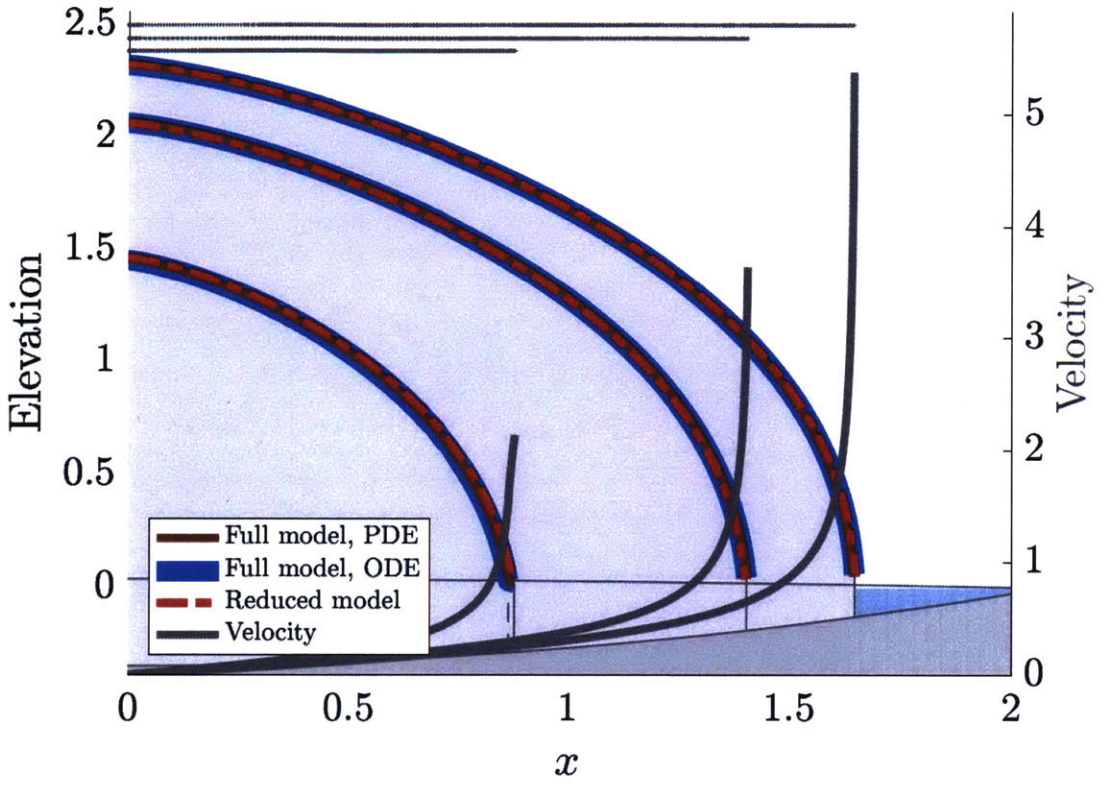

Figure 3-13: Steady state glacier profile for an upward ($b(x) = 0.08x^2 - 0.25$) sloping bed given that the glacier calves at the critical thickness defined by the meltwater calving law. Solutions are computed using the shooting method from section 2.4 (blue line) and a discontinuous Galkerin method (brown and grey lines). The dotted red line is the steady state solution to the reduced model (3.116) and (3.1). The thin grey line is the flotation thickness $h_g(x) = -r^{-1}b(x)$. Grey dots along the top of each figure are vertexes of the triangulation after adaptive mesh refinement on the full model. We have set $a = 0.6$, $\mu = 0.5$, $\lambda = 011$, and $w = 10$.

Chapter 4

Incremental local approximations for computationally intensive sampling

Sampling methods (e.g., Monte Carlo or Markov chain Monte Carlo (MCMC)), are used to estimate expectations with respect to probability distributions. However, these methods typically require many samples and the number of times the probability density function must be evaluated scales linearly with the number of samples. Furthermore, sampling methods become increasingly inefficient with increasing dimension. If the probability density function is computationally expensive and/or the problem is high dimensional, sampling methods are infeasible.

High dimensional problems often have a low-dimensional quantity of interest (QoI). In such cases, we partition parameter space into the QoI and auxiliary variables. Marginalizing the auxiliary variables reduces to dimensionality of the space—potentially making sampling methods feasible. However, deriving an analytical expression for the marginal probability density function is rarely possible. However, unbiased estimates are often available at great computational cost. We now must characterize a low-dimensional distribution with a computationally expensive probability density function that can only estimated stochastically.

We develop a new algorithm that characterizes marginal distributions with computationally intensive probability density functions (Algorithm 4). This algorithm replaces computationally intensive probability density evaluations with a cheaper sur-

rogate model. We show that a bias-variance trade-off implies and ideal refinement rate. Additionally, build the surrogate with noisy density evaluations characterizes selected marginal distributions. Section 4.2 describes existing sampling methods and challenges associated with them. Section 4.3 describes the local polynomial surrogate model that we will use to reduce the computational expense of sampling methods. Finally, section 4.4 describes our algorithm.

4.1 Motivation

Statistical models characterize uncertainty—mathematically describing our state-of-knowledge about unknown parameters. For example, predictions of future mean sea level might be described by a Gaussian distribution with a particular mean and variance. Rather than making a single prediction, we assign each prediction a probability density; in our simple sea level example, the most likely scenario is the predicted mean. However, a large variance indicates other scenarios—sea levels much larger or much smaller than the predicted mean—have non-negligible probability. The probability distribution associated with each parameter models how certain we are about its value.

Probability distributions that describe our state of knowledge are constructed by combining multiple sources of information via Bayes' rule. In general, we consider distributions that combine (i) prior information, (ii) expert knowledge in the form of physical and/or numerical models, and (iii) noisy data. Mathematically, we define random variables[1] X and Y, which take values in $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^n$, respectively; we let X represent a quantity of interest (QoI) and Y represents observational data. Bayes' rule, derived from the law of total probability, defines the posterior distribution:

$$\underbrace{\pi_{X|Y}(x|y)}_{\text{Posterior}} = \frac{\pi_{X,Y}(x,y)}{\pi_Y(y)} = \frac{\pi_{Y|X}(y|x)\pi_X(x)}{\pi_Y(y)} \propto \underbrace{\pi_{Y|X}(y|x)}_{\text{Likelihood}} \underbrace{\pi_X(x)}_{\text{Prior}}. \tag{4.1}$$

[1]Upper case letters denote a random variable and lower case letters denote a specific realization of that random variable.

Here, we assume that all random variables have densities with respect to the Lesbesgue measure. We use $\pi_X(x)$ to denote the probability density of the random variable X, evaluated at x. Henceforth, we will let the arguments of the density dictate its associated distribution (e.g., $\pi(x) = \pi_X(x)$). The posterior distribution $\pi_{X|Y}$ is decomposed into the prior π_X, the likelihood $\pi_{Y|X}$, and the evidence π_Y. The prior represents all knowledge about the parameters before observing the data Y, for example, positivity, smoothness, or other physical constraints.

Expert knowledge and data are incorporated into the likelihood, which quantifies how well the data are explained by given parameters. Let $f : \mathcal{X} \to \mathcal{Y}$ be the forward model—derived from expert knowledge—that maps realizations of X to the observed quantity. Typically, the likelihood is a function of some "distance" between the observations y and the prediction $f(x)$—the better the data y are explained by $f(x)$, the larger the likelihood. This framework constitutes a Bayesian inverse problem; the posterior distribution models uncertainty about the inferred parameter X given the prior π_X, the model f, and the conditional distribution $\pi_{Y|f(x)}$, of the realized data y.

Most applications require expectations with respect to the posterior distribution, which, in general, cannot be computed analytically. The posterior mean, for example, is the expectation of the parameter X under the posterior. We may also be interested in the probability that X is greater than some critical value x

$$\mathbb{P}[X > x|y] = \mathbb{E}_{\pi_{X|Y}}\left[\mathbb{1}(X > x)|y\right], \tag{4.2}$$

where $\mathbb{1}(\cdot)$ is the indicator function. In some cases, if the likelihood belongs to an exponential family and the prior has a specific form, the posterior has known structure. Such prior-posterior pairs are conjugate, allowing i.i.d. posterior sampling or closed-form evaluation of posterior moments [Gelman et al., 2013]. Special distributions, such as uniform or Gaussian, have known sampling algorithms, and one-dimensional distributions can be easily sampled if the inverse cumulative distribution is known. In general, however, computing these integrals analytically is rarely possible. We, there-

fore, resort to numerics: quadrature [Gerstner and Griebel, 1998, Smolyak, 1963] or sampling methods. We focus on general sampling methods that draw correlated and/or weighted samples from a target distribution, which is often a Bayesian posterior.

Our goal is to characterize probability distributions whose densities are computationally expensive and/or when only stochastic estimates of its probability density function are available. For example, we will often be in characterizing *marginal* distributions. Suppose $Z \in \mathcal{Z} \subseteq \mathbb{R}^n$ are auxiliary variables are interdependent with X but are not relevant in our application. The posterior distribution over the target and auxiliary variables—X and Z conditioned on the data Y—is $\pi_{X,Z|Y}$. The marginal of interest is then

$$\pi(x|y) = \int_{\mathcal{Z}} \pi(x, z|y)\, dz. \tag{4.3}$$

Typically, the marginalization cannot be computed analytically; however, unbiased Monte Carlo estimates of the marginal density might be available. Additionally, evaluating the posterior density $\pi(x, z|y)$ is often computationally intensive. We will develop algorithms that target $\pi_{X|Y}$ directly, without needing to fully characterize the distribution over the joint space, while minimizing the number of posterior density evaluations.

4.2 Sampling methods

Sampling methods characterize a target distribution π (often a posterior distribution) by estimating expectations with respect to that distribution. Expectations of $f : \mathcal{X} \to \mathbb{R}$ are approximated with weighted and/or correlated samples such that

$$\mathbb{E}_{\pi_X}[f] = \int f(x)\pi(x)\, dx \approx \sum_{i=1}^{M} w^{(i)} f(x^{(i)}) \tag{4.4}$$

Ideally, the distribution π can be sampled directly, in which case $x^{(i)} \sim \pi$ and $w^{(i)} = M^{-1}$. In this case, the stochastic estimate (4.4) converges almost surely to the true

value as $M \to \infty$ with rate $1/\sqrt{M}$ if $\mathbb{V}\mathrm{ar}_{X\sim\pi_X}(f(X)) = \sigma_\pi^2 < \infty$ [Billingsley, 2009]. This rate is given by the central limit theorem

$$\sqrt{M}\left(\frac{1}{M}\sum_{i=1}^{M} f(x^{(i)}) - \mathbb{E}_{\pi_X}[f]\right) \xrightarrow[M\to\infty]{d} N\left(0, \sigma_\pi^2\right). \tag{4.5}$$

In general, however, $\{x^{(i)}\}$ are either drawn from some other distribution p_X and the error in the estimate is corrected with non-uniform weights or $\{x^{(i)}\}$ are correlated samples from π_X. Common sampling methods are in this broader class are importance sampling [Gelman et al., 2013] and Markov chain Monte Carlo (MCMC) [Brooks et al., 2011]. We will briefly review both.

4.2.1 Importance sampling

Importance sampling computes weighted samples of the target distribution [Owen, 2013]. The expectation (4.4) is approximated by

$$\mathbb{E}_{\pi_X}[f] = \int f(x)\frac{\pi(x)}{q(x)}q(x)\,dx \approx \sum_{i=1}^{M} w^{(i)} f(x^{(i)}); \tag{4.6}$$

the importance samples $x^{(i)}$ are drawn from the biasing distribution q_X and the weights are $w^{(i)} = \pi(x^{(i)})/q(x^{(i)})$. If the target density $\pi(x)$ can only be evaluated up to a multiplicative constant $\pi(x) \propto \bar{\pi}(x)$, the weights are normalized: $w^{(i)} = \bar{w}^{(i)}/\sum_{i=1}^{M}\bar{w}^{(i)}$, where $\bar{w}^{(i)} = \bar{\pi}(x^{(i)})/q(x^{(i)})$. This adjustment is called 'self-normalized importance sampling.' Again, applying the central limit theorem, (4.5) to the importance sampling estimator

$$\sqrt{M}\left(\sum_{i=1}^{M} \bar{w}^{(i)} f(x^{(i)}) - \mathbb{E}_{\pi_X}[f]\right) \xrightarrow[M\to\infty]{d} N\left(0, \mathbb{V}\mathrm{ar}_{X\sim q_X}\left(\frac{f(X)\pi(X)}{q(X)}\right)\right), \tag{4.7}$$

where the variance is assumed finite. In the case of self-normalized importance sampling there is a nonzero bias that vanishes as $M \to \infty$. The importance sampling estimate (4.6), therefore, also converges almost surely at rate $1/\sqrt{M}$. In regions

where $q(x) \ll \pi(x)$ the importance weights are large, but rarely drawn. Conversely, in regions where $q(x) \gg \pi(x)$ the importance weights are small and the sampler frequently draws samples that have little effect on the estimate.

We frequently approximate marginal distributions with importance sampling. Consider the joint distribution $\pi_{X,Z}$ between two random variables X and Z; we assume X is the QoI and Z is an auxiliary variable. We can approximate the marginal density

$$\pi(x) = \int \pi(x,z)\,dz = \int \frac{\pi(x,z)}{q(z)} q(z)\,dz \approx \sum_{i=1}^{M} \frac{\pi(x,z^{(i)})}{q(z^{(i)})}, \tag{4.8}$$

where $z^{(i)} \sim q_Z$. When joint density evaluations are available but analytically computing the marginal is impossible, importance sampling computes stochastic, unbiased estimates of the marginal density.

Multiple importance sampling

Suppose rather than having one biasing distribution q_X, suppose that we have multiple biasing distributions $q_X^{(l)}$ for $l = 1 : L$ [Owen, 2013]. We partition the $M = \sum_{l=1}^{L} m^{(l)}$ importance samples such that $m^{(l)}$ samples are generated samples from each biasing distribution; $X^{(\zeta_{l,i})} \sim q_X^{(l)}$, where $\zeta_{l,i} = \sum_{j=1}^{l-1} m^{(j)} + i$ is an indexing function. Define $w^{(l,i)} = \bar{w}^{(l,i)} / \sum_{j=1}^{m^{(l)}} \bar{w}^{(l,j)}$, where $\bar{w}^{(l,i)} = \bar{\pi}(x^{(\zeta_{l,i})})/q^{(l)}(x^{(\zeta_{l,i})})$; the multiple importance sampling estimate is

$$\mathbb{E}_{\pi}[g] = \sum_{l=1}^{L} \omega^{(l)} \sum_{i=1}^{m^{(l)}} w^{(l,i)} g(x^{(\zeta_{l,i})}), \tag{4.9}$$

where $\omega^{(l)} \geq 0$ is a partition of unity—$\sum_{l=1}^{L} \omega^{(l)} = 1$. For example, $\omega^{(l)} = \frac{m^{(l)}}{M}$ weighs each level by the number of importance samples drawn from the distribution $q_X^{(l)}$.

4.2.2 Markov chain Monte Carlo (MCMC)

MCMC computes correlated samples from the target distribution π by constructing a Markov chain whose stationary distribution is π [Brooks et al., 2011, Metropolis

et al., 1953]. Given the state x_t at step[2] t, a transition kernel determines the next state x_{t+1}. The transition kernel is constructed such that, after infinite steps, states in the chain $\{x_t\}_{t=1}^{\infty}$ are correlated samples from the target distribution π. Equivalently, π is the stationary distribution of the transition kernel K, which is defined such that $K(dx_{t+1}, x_t)$ is the probability of moving to dx_{t+1} from x_t.

MCMC samples are correlated—the current state depends on the previous states. Because MCMC is constructed as a Markov chain, the transition kernel only depends on the current state (although some variations relax this condition). Therefore, the correlation between the current state x_t and previous states x_{t-j} decreases as the lag j increases. Typically, each MCMC step consists of two steps: (i) propose a new point by sampling the proposal distribution $\tilde{x} \sim p_{X|x_t}$ and (ii) move to the proposed point with probability α. For Metropolis-Hastings MCMC (MH-MCMC) [Hastings, 1970, Metropolis et al., 1953], the acceptance probability α is

$$\alpha = \min\left(1, \frac{\pi(\tilde{x})p(x_t|\tilde{x})}{\pi(x_t)p(\tilde{x}|x_t)}\right). \tag{4.10}$$

We typically define a Gaussian proposal $\tilde{x} \sim \mathcal{N}(x_t, \Sigma)$. Figure 4-1 shows a visual representation of a simple MCMC chain. We can also adapt the proposal covariance using previous samples

$$\mu_t = \sum_{\tau=1}^{t} x_\tau \quad \text{and} \quad \Sigma_t = \varepsilon_0 \sum_{\tau=1}^{t} (x_\tau - \mu_t)(x_\tau - \mu_t)^T + \varepsilon_1 I; \tag{4.11}$$

I is the identity and ε_0 and ε_1 are algorithm parameters. Chains generated using an adaptive proposal are not Markov chains: the next state depends on *all* of the previous states. However, this algorithm, adaptive Metropolis MCMC (AM-MCMC), also has stationary distribution π [Haario et al., 2001]. Therefore, samples generated with both MH-MCMC and AM-MCMC are treated as correlated samples from the target distribution.

[2]For notational convenience, we denote importance samples with superscript (i) and MCMC steps with subscript t. This distinction will be useful when we use importance sampling within MCMC.

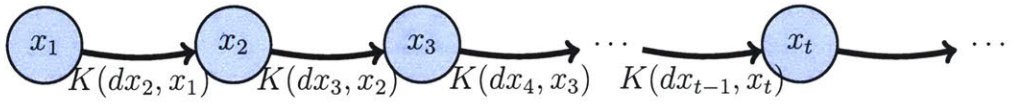

Figure 4-1: A visual representation of a Markov chains with transition kernel K; MCMC constructs a transition kernel such that the stationary distribution π_X is the desired target distribution.

Metropolis adjusted Langevin algorithm (MALA)

The Metropolis adjusted Langevin algorithm (MALA) and its manifold variants (mMALA) leverage gradient information of the target density to propose the next state [Girolami and Calderhead, 2011, Roberts and Rosenthal, 1998]. MALA is based on Langevin diffusion with stationary distribution π_X, which is defined by the stochastic differential equation (SDE)

$$\dot{X} = \frac{1}{2}\nabla \log\left(\pi(X)\right) + dW(t), \tag{4.12}$$

where $dW(t)$ denotes Brownian motion. Discretizing in time with Euler–Maruyama (forward Euler in the SDE context) defines the proposal

$$\tilde{x} = x_t + \frac{\epsilon^2}{2}\nabla \log\left(\pi(x_t)\right) + \epsilon\theta, \tag{4.13}$$

where $\theta \sim \mathcal{N}(0, I)$ and ϵ is the step size of the time integrator. However, integration error introduced by the time discretization means that π_X is no longer guaranteed to be the stationary distribution. We correct this using a Metropolis accept/reject step. We propose the next state by sampling the proposal distribution $p_{X|X} = \mathcal{N}(x_t + \frac{\epsilon^2}{2}\nabla \log\left(\pi(x_t)\right), \epsilon^2 I)$ and accept with probability (4.10).

The simplified manifold Metropolis adjusted Langevin algorithm (mMALA) adapts the proposal to the local geometry of the target distribution. This method is closely related to the preconditioning preformed in the stochastic Newton method [Martin et al., 2012]. This method extends MALA by defining Langevin diffusion with invariant measure π_X on a Riemann manifold with arbitrary metric tensor $M(x)$ [Girolami and Calderhead, 2011]. The proposal is derived by discretizing the SDE defining

Langevin diffusion on the manifold, leading to

$$p_{X|X} = \mathcal{N}\left(x_t + \frac{\epsilon}{2}M^{-1}(x_t)\nabla \log\left(\pi(x_t)\right), \epsilon^2 M^{-1}(x_t)\right). \tag{4.14}$$

We treat the position-dependent symmetric positive definite (SPD) mass matrix $M^{-1}(x)$ as a preconditioner that rescales parameter space to make the distribution locally more isotropic. Again, we correct errors introduced by the time discretization using a Metropolis accept/reject step.

We are free to choose any preconditioner $M(x)$ as long as $M^{-1}(x)$ is SPD. Choosing the identity matrix $M(x) = I$ corresponds to standard MALA. However, [Girolami and Calderhead, 2011] define mMALA by choosing $M(x)$ to reflect a Riemannian metric induced by the posterior distribution:

$$M(X) = -\mathbb{E}_{Y|X}\left[\nabla^2 \log\left(\pi(Y|X)\right)\right] - \nabla^2 \log\left(\pi(X)\right), \tag{4.15}$$

which is the expected Fisher information plus the negative Hessian of the log-prior density. In general, computing the expected Fisher information is nontrivial to compute. However, it becomes feasible if the likelihood is Gaussian conditioned on the forward model result, e.g.,

$$\pi(Y|f(x)) = N(Y; f(x), \Sigma_l) \tag{4.16}$$

for some prescribed covariance matrix Σ_l. If the prior is also Gaussian with covariance Σ_p, then

$$M^{-1}(x) = J^T(x)\Sigma_l^{-1}J(x) + \Sigma_p^{-1}, \tag{4.17}$$

where $J(x) = \nabla f(x)$. Although [Girolami and Calderhead, 2011] observe that this preconditioner can dramatically improve MALA's performance, this algorithm can be difficult to apply because it requires derivative information. Often, we do not have computable derivatives and if we do, they are computationally burdensome.

Pseudo-marginal Markov chain Monte Carlo (PM-MCMC)

Suppose our goal is to characterize the distribution π_X but we cannot directly evaluate its density. Suppose, further, thatwe can evaluate the joint density $\pi(x, z)$. The 'brute force' approach is to generate samples from $\pi_{X,Z}$ using MCMC—for example, MH-MCMC, AM-MCMC, or MALA—and discard the samples of the auxiliary random variable Z. This approach is inefficient since the augmented space $\mathcal{X} \times \mathcal{Z}$ is potentially high dimensional. Often, $\dim(\mathcal{X}) \ll \dim(\mathcal{Z})$ and targeting π_X directly is more feasible.

Pseudo-marginal MCMC (PM-MCMC) directly targets marginal distributions using importance sampling estimates of the target marginal [Andrieu and Roberts, 2009, Beaumont, 2003]. Given that the MCMC chain is at state x_t, PM-MCMC proposes a new state $\tilde{x} \sim p_{\tilde{X}|X}$. The proposal distribution $p_{\tilde{X}|X}$ could be fixed, as in MH-MCMC, or adaptive, as in AM-MCMC. Given $\tilde{x}$, M samples from the biasing distribution $\tilde{z} \sim q_{Z|X}$ are used to compute the unbiased estimate $\hat{\pi}(\tilde{x}) \approx \pi(\tilde{x})$ via (4.8). The proposed point is accepted with probability

$$\alpha = \min\left(1, \frac{\hat{\pi}(\tilde{x})p(x_t|\tilde{x})}{\hat{\pi}(x_t)p(\tilde{x}|x_t)}\right). \tag{4.18}$$

[Andrieu and Roberts, 2009, Beaumont, 2003] observe that the acceptance ratio can be written equivalently as

$$\frac{\hat{\pi}(\tilde{x})p(x_t|\tilde{x})}{\hat{\pi}(x_t)p(\tilde{x}|x_t)} = \frac{\sum_{i=1}^{M}[\pi(\tilde{x}, \tilde{z}^{(i)}) \prod_{j=1;j\neq i}^{M} q(\tilde{z}^{(j)}|\tilde{x})]p(x_t|\tilde{x}) \prod_{j=1}^{M} q(z_t^{(j)}|x_t)}{\sum_{i=1}^{M}[\pi(x_t, z_t^{(i)}) \prod_{j=1;j\neq i}^{M} q(z_t^{(j)}|x_t)]p(\tilde{x}|x_t) \prod_{j=1}^{M} q(\tilde{z}^{(j)}|\tilde{x})}, \tag{4.19}$$

where M is the number of importance samples. We interpret the left side of (4.19) as a Markov chain targeting the marginal distribution π_X, as desired; whereas, we interpret the right side of (4.19) as a Markov chain targeting a distribution over the augmented space $\mathcal{X} \times \mathcal{Z}^{(1)} \times \ldots \times \mathcal{Z}^{(M)}$ [Andrieu and Roberts, 2009]. Figure 4-2 shows a visual representation of a pseudo-marginal chain. Although this augmented space is very high-dimensional, the importance samples $\tilde{y}^{(i)}$ are uncorrelated. At each step, we

can choose to save the marginal estimate $\hat{\pi}(x_t)$ or resample $z_t \sim q_{Z|X}$ and recompute $\hat{\pi}(x_t)$. [Andrieu and Roberts, 2009] show that recomputing the estimate is crucial to the convergence of the algorithm. Therefore, PM-MCMC requires two importance sampling estimates per MCMC step, which is often computationally prohibitive.

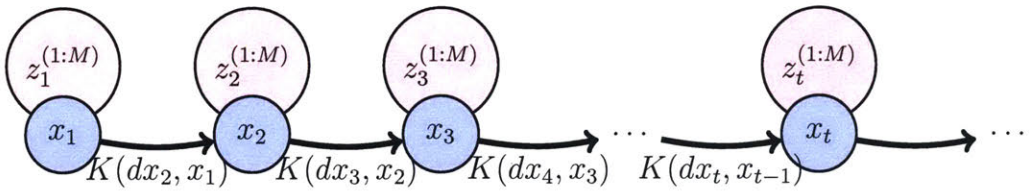

Figure 4-2: A visual representation of a pseudo-marginal MCMC chain with transition kernel K. We propose the next state x_t by sampling from a proposal distribution p_X and the importance samples $z_t^{(1:M)}$ are sampled from a biasing distribution q_Z.

4.2.3 Example of joint and pseudo-marginal MCMC

We demonstrate MCMC (AM-MCMC and PM-MCMC) and importance sampling targeting a two dimensional distribution with the following log-quartic density:

$$\log\left(\pi(X_1, X_2)\right) = -X_1^4 - \frac{(2X_2 - X_1)^2}{2}. \tag{4.20}$$

Typically, evaluating the joint density is the most computationally intensive component of a sampling algorithm. Therefore, we measure the 'cost' of characterizing π_{X_1,X_2} or its marginals by the number of joint density evaluations. For example, generating T MCMC samples using AM-MCMC directly targeting π_{X_1,X_2} requires T joint density evaluations. Note that we will interchangeably use the terms 'joint density evaluations' and 'model evaluations.' Figure 4-3 shows a histogram of the joint density and its marginals computed with AM-MCMC.

We can approximate the marginal densities using importance sampling. The blue dots in figure 4-3 show estimates of $\pi(X_1)$ and $\pi(X_2)$ computed using (4.8). The result is an unbiased estimate of the target density whose variance decreases with the number of importance samples. In this case, each blue dot costs M model evaluations.

PM-MCMC directly characterizes the target marginal distribution. The red line

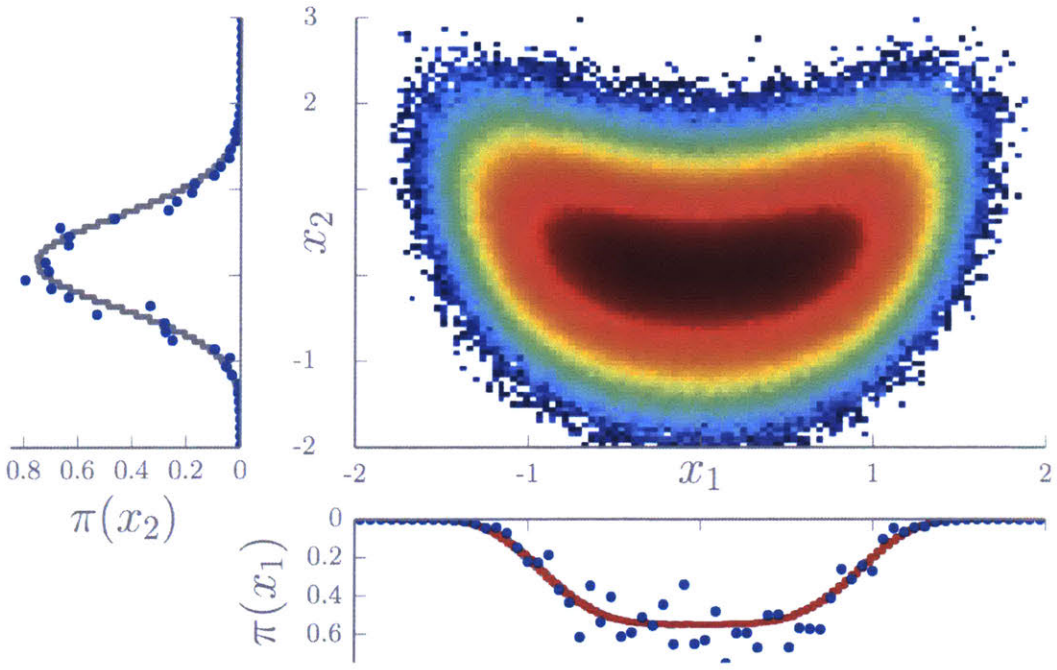

Figure 4-3: The middle plot shows a histogram of $T = 10^7$ samples generated using AM-MCMC; the grey lines are histograms of the same points showing the marginal distributions. The blue dots are importance sampling estimates of the marginal densities computing with $M = 10$ importance samples and a standard normal biasing distribution over the marginalized variable. The red line shows a histogram showing $T = 10^7$ samples generated using PM-MCMC targeting π_{X_1} using $M = 10$ importance samples at each step to estimate the target density. We note that even though the marginal estimates are noisy, the PM-MCMC result and the marginal of AM-MCMC targeting the joint density are indistinguishable.

in figure 4-3 shows a histogram of samples generated with PM-MCMC; the target marginal distribution is estimated with M importance samples. Generating T MCMC samples using PM-MCMC requires MT joint density evaluations. In this example, we only use $M = 10$ importance samples—leading to fairly crude target density estimates. However, this still increases the algorithm's cost by an order of magnitude.

4.3 Surrogate models for sampling algorithms

4.3.1 Background

Computationally efficient sampling methods reduce the required number of target density evaluations [Conrad et al., 2016, Kennedy and O'Hagan, 2001, Sacks et al., 1989]. Rather than repeatedly evaluating $\pi(X)$ within a sampling method, these methods leverage properties of the density (e.g., continuity, smoothness) to build computationally cheaper approximations $\breve{\pi}(X) \propsim \pi(X)$ and replace the target density $\pi(X)$ with the surrogate model. For example, we can approximate the target density with (i) global polynomial approximations [Marzouk and Xiu, 2009, Marzouk et al., 2007], (ii) sparse polynomial chaos expansions [Conrad and Marzouk, 2013, Smolyak, 1963, Wiener, 1938], (iii) radial basis functions [Bliznyuk et al., 2012, Joseph, 2012], or (iv) Gaussian processes [Bernardo et al., 2008, Kennedy and O'Hagan, 2001, Santner et al., 2003, Sacks et al., 1989]. Targeting $\breve{\pi}$ instead of π introduces a bias $b = \mathbb{E}_{\pi}[h] - \mathbb{E}_{\breve{\pi}}[h]$ on the expectation of any function h. We denote the true target density by $\pi(x)$; a Monte Carlo (or importance sampling) estimate of $\pi(x)$ by $\hat{\pi}(x)$; and a surrogate model constructed from evaluations of $\pi(x)$ or $\hat{\pi}(x)$, by $\breve{\pi}(x)$.

An alternative to replacing π with the computationally cheaper surrogate $\breve{\pi}$ is to generate samples targeting π and use the surrogate to correct MCMC estimates [Vihola et al., 2016]. MCMC generates correlated samples from the target distribution π_X, evaluating the true target density $\pi(x)$ every step. However, standard MCMC algorithms do not account for regularity in the underlying target density. [Vihola et al., 2016] save target density evaluations at each step and partition the generated

samples into two subsets. The first subset is used to construct a surrogate model that exploits regularity in the target density; the remaining samples are used to compute MCMC estimates, however, each sample is weighted using importance sampling weights with the surrogate as the biasing distribution. This correction accounts for expected regularity in the underlying target density, improving the accuracy of the estimate. However, the number of model evaluations is still equal to the number of generated samples.

4.3.2 Local approximations

Our goal is to construct an algorithm that leverages regularity in the target density *and* is asymptotically exact. Previous work by [Conrad et al., 2016] show that continually and infinitely refining the surrogate model as the algorithm progresses asymptotically generates samples from the exact target distribution. However, [Conrad et al., 2016] do not derive an optimal refinement rate. Moreover, their approach requires *exact* target density evaluations, rather than only (unbiased) estimates.

We focus on *local polynomial* surrogate models, which approximate $g(x)$ given either exact or noisy evaluations [Conn et al., 2009, Kohler, 2002, Stone, 1977]. Let the evaluated set $\mathcal{S}_n = \{x_1, \ldots, x_n\}$ be the set of points were we have evaluated g (or an unbiased estimator of g). For now, we will take this set as a given and discuss how to choose the points $x_{1:n}$ later. A local approximation, built from nearby function evaluations, allows refinement to focus on regions where the target distribution has greater mass. Polynomial approximations are appealing because they can be built easily, are cheap to evaluate, and have known, analytic derivatives. We begin by building a surrogate model from exact evaluations and then generalize to the noisy evaluation case.

In the Bayesian setting (when the target density is a posterior decomposed via Bayes rule (4.1)), evaluating the prior is often cheap compared to the likelihood. In such cases, building a surrogate model of the likelihood function rather than the posterior density is sufficient. In general, we build a surrogate model for the function $g : \mathcal{X} \to \mathbb{R}$; $g(x)$ may be the target density, likelihood, or any other function such that,

given the surrogate, approximating the target density is trivial and if the surrogate is exact then the target density approximation is exact as well.

4.3.3 Local polynomial approximations with exact function evaluations

We construct a local polynomial approximation using n exact evaluations of g. Define a sequence of functions $\{g_n(x;\mathcal{S}_n)\}_{n\in\mathbb{N}}$ such that in the limit of infinite model evaluations g_n is exactly g. More rigorously, for all x, $\lim_{n\to\infty} g_n(x) = g(x)$; such sequences are asymptotically exact. Define a local average estimator as the weighted sum of model evaluations

$$g_n(x) = \sum_{i=1}^{n} w_i(x) g(x_i). \tag{4.21}$$

We will later assume that the weights $w_i(x)$ depend on how far x is from $x_i \in \mathcal{S}_n$. For example, the simplest local estimator is the k nearest neighbor average, a local averaging estimate with weight function

$$w_i(x;\mathcal{S}_n) = \begin{cases} k^{-1} & \text{if } x_i \in \mathcal{S}_n \text{ is one of the } k \text{ nearest neighbors of } x \\ 0 & \text{else.} \end{cases} \tag{4.22}$$

Kernel-based local approximations

Local averaging estimates are equivalent to weighted regression with constant polynomials when the weight functions are derived from kernel[3] functions $K : \mathcal{X} \times \mathcal{X} \to \mathbb{R}^+$ [Kohler, 2002]. Define the kernel weight function

$$w_i(x;\mathcal{S}_n) = \frac{K(x, x_i)}{\sum_{j=1}^{n} K(x, x_j)}, \tag{4.23}$$

which, via (4.21), induces a corresponding kernel estimate. Choosing

[3]Note that this is a *different* kernel than the MCMC transition kernel.

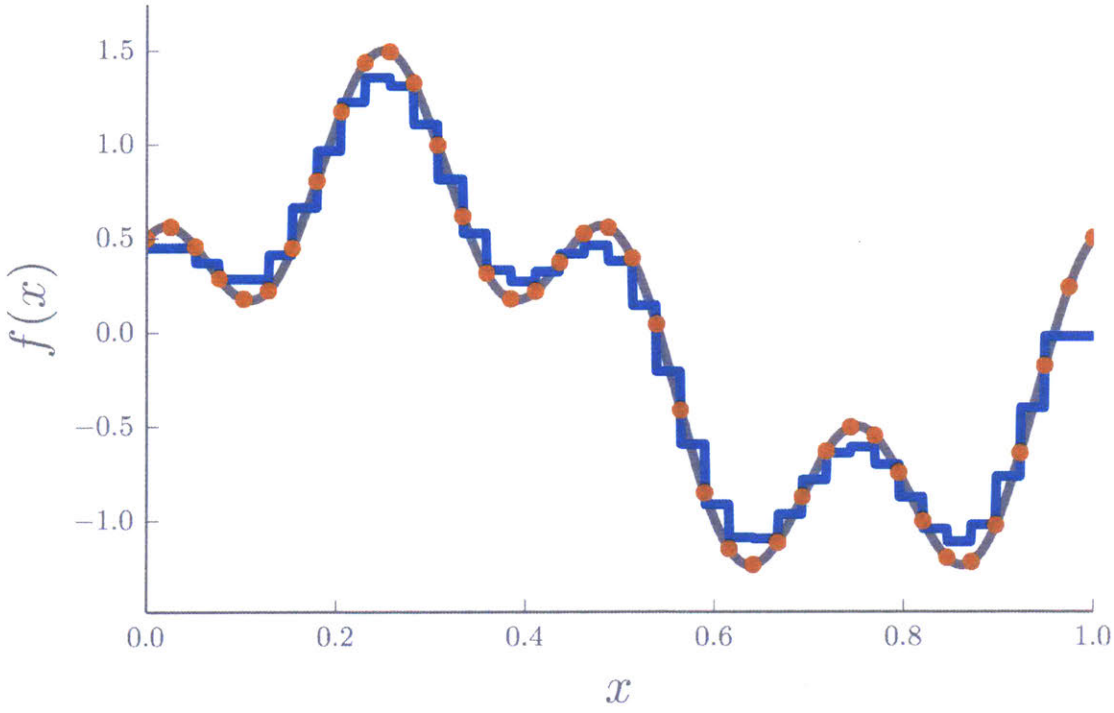

Figure 4-4: An estimate (blue line) of $g(x) = \sin(2\pi x) + 0.5\cos(8\pi x)$ (grey line) using (4.21) with weights induced by the kernel (4.24) with $k = 4$ nearest neighbors. The data are $n = 40$ equally spaced model evaluations (orange dots).

$$K(x, x_i) = \begin{cases} 1 & \text{if } x_i \in \mathcal{S}_n \text{ is one of the } k \text{ nearest neighbors of } x \\ 0 & \text{else} \end{cases} \tag{4.24}$$

is equivalent to the k nearest neighbor weights (4.22). For each $x \in \mathcal{X}$, the kernel estimate g_n corresponding to kernel K is the solution to

$$\sum_{i=1}^{n} (g_n(x) - g(x_i))^2 K(x, x_i) = \min_{a \in \mathbb{R}} \sum_{i=1}^{n} (a - g(x_i))^2 K(x, x_i) \tag{4.25}$$

and g_n is, therefore, a locally constant function [Kohler, 2002]. Figure 4-4 shows an example of this estimate.

Continuous local approximations Despite being locally constant, kernel estimates are continuous if the K is chosen properly. If we assume the kernel has the

form $K(x, x_i) = \tilde{K}(\|x - x_i\|)$, where $\tilde{K} : \mathbb{R}^+ \to \mathbb{R}^+$ is a monotonically decreasing and left-continuous function such that $b\delta_{[0,r^2]}(u) \leq \tilde{K}(u) \leq B\delta_{[0,R^2]}(u)$, where $0 < r \leq R < \infty$, $0 < b \leq B < \infty$, and $\delta_{[a,b]}(u)$ is the indicator function. This assumption ensures that the kernel's support is a radius $0 < \Delta < \infty$ ball centered at x and decreases as x_i moves farther from x. The kernel estimate corresponding to

$$\tilde{K}(u; r) = \begin{cases} 1 & \text{if } u \text{ is in the ball of radius } \hat{u} \text{ centered at } 0 \\ \max\left(\left(1 - \left(\frac{u-\hat{u}}{r}\right)\right)^3, 0\right) & \text{else,} \end{cases} \tag{4.26}$$

is continuous. In practice, we use (4.24), favoring its simplicity over continuity. We require the kernel's support to contain at least one point: $k \geq 1$ for (4.24) and $\hat{u}$ sufficiently large to include at least one point for (4.26). This requirement reflects the locally constant polynomial interpretation and grows when we generalize to higher-order polynomials.

High-order local polynomials

The kernel estimate generalizes by replacing the locally constant approximation with local polynomial approximations. Let $\mathcal{P}_p$ be the space of $j \leq p$ degree polynomials over $\mathcal{X}$. The local surrogate for each $x \in \mathcal{X}$ is $g_n(x)$ such that

$$\sum_{i=1}^{n} (g_n(x) - g(x_i))^2 \tilde{K}(\|x - x_i\|) = \min_{m \in \mathcal{P}_p} \sum_{i=1}^{n} (m(x_i) - g(x_i))^2 \tilde{K}(\|x - x_i\|), \tag{4.27}$$

which is equivalent to weighted regression using p-degree polynomials [Kohler, 2002]. Define basis $\Phi = \{\phi_j\}_{j=1}^{q}$ for $\mathcal{P}_p$,[4] $\boldsymbol{\phi}(x) = [\phi_1(x), \ldots, \phi_q(x)]^T$, and $\boldsymbol{x}_k = \{\tilde{x} \in \mathcal{S}_n : \tilde{K}(\|x - \tilde{x}\|) > 0\}$. The elements of $\boldsymbol{x}_k(x)$ are the k nearest neighbors to x: the k closest points in the evaluated set to x. The optimal polynomial kernel estimate is $g_n(x) = \boldsymbol{\phi}^T(x)\mathbf{a}(x)$ such that

[4]The polynomial space has dimension q; if $\mathcal{X} \subseteq \mathbb{R}$ then $q = p + 1$.

$$\mathbf{a}(x) = \underset{\boldsymbol{\alpha}\in\mathbb{R}^q}{\arg\min} \|\mathbf{V}\boldsymbol{\alpha} - \mathbf{g}\|_{\mathbf{K}}^2, \tag{4.28}$$

where $\|\mathbf{x}\|_{\mathbf{A}}^2 = \mathbf{x}^T\mathbf{A}\mathbf{x}$, $\boldsymbol{K}(x;\boldsymbol{x}_k) = \text{diag}\left[\tilde{K}(\|x - x_1\|), \ldots, \tilde{K}(\|x - x_k\|)\right]$, $\mathbf{V}$ is the Vandermonde matrix

$$\mathbf{V}(\boldsymbol{x}_k) = \begin{bmatrix} — & \phi^T(x_1) & — \\ & \vdots & \\ — & \phi^T(x_k) & — \end{bmatrix}, \quad \text{and} \quad \mathbf{g}(\boldsymbol{x}_k) = \begin{bmatrix} g(x_1) \\ \vdots \\ g(x_k) \end{bmatrix}. \tag{4.29}$$

Assuming $\mathbf{V}$ has full column rank and $k \geq q$, the polynomial kernel estimates exists and is unique [Kohler, 2002]. The solution to (4.28) is $\mathbf{a}(x) = \left(\mathbf{V}^T\mathbf{K}\mathbf{V}\right)^{-1}\mathbf{V}^T\mathbf{g}$. Given $\mathcal{S}_n$, a polynomial kernel estimate is computed by solving the weighted least squares problem (4.28), which is equivalent to (4.27) [Kohler, 2002]. Figure 4-5 shows constant, linear, and quadratic polynomial kernel estimates.

A sequence is $\{g_n\}_{n=k}^{\infty}$ is constructed by adding new points to the evaluated set.[5] This sequence satisfies the required limit $\lim_{n\to\infty} g_n(x) = g(x)$ if the kernel's radius decays $\Delta \to 0$ as $n \to \infty$. Reparameterizing (4.26) so the radius Δ is the minimum radius that contains k points and fixing $k > k^*$, where k^* is the number of points required to interpolate, ensures that both (4.24) and (4.26) induce asymptotically exact estimates as $n \to \infty$ [Conn et al., 2009, Kohler, 2002, Stone, 1977]. We postpone detailed discussion about how this error decays until after we describe how to modify (4.27) to account for noisy function evaluations.

4.3.4 Local polynomial approximations with noisy function evaluations

We now relax the assumption that g can be evaluated exactly. Instead, we assume that we can only evaluate the unbiased estimator $\hat{g} : \mathcal{X} \to \mathbb{R}$ such that

$$\mathbb{E}\left[\hat{g}(x)\right] = g(x) \quad \text{and} \quad \mathbb{V}\text{ar}\left(\hat{g}(x)\right) = \sigma_g^2 < \infty \tag{4.30}$$

[5] New points must be chosen such that $\mathbf{V}$ maintains full column rank.

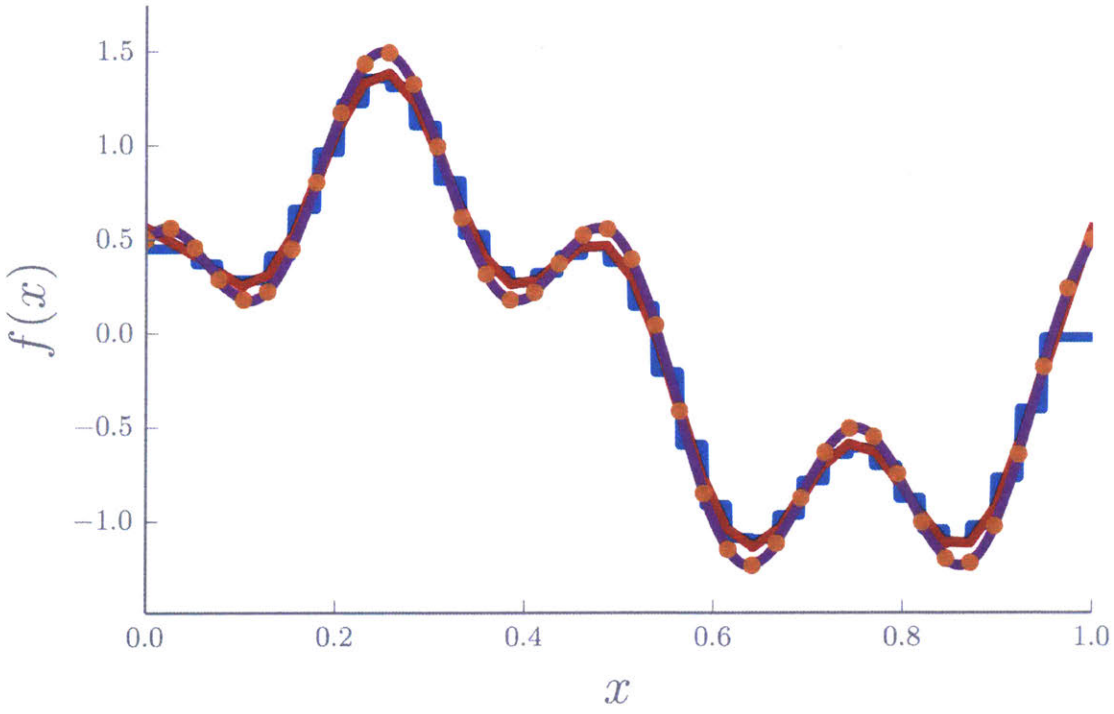

Figure 4-5: Locally constant (blue line), linear (red line), and quadratic (purple line) estimates of $g(x) = \sin(2\pi x) + 0.5\cos(8\pi x)$ (grey line) using (4.27) with weights induced by the kernel (4.24) with $k = 4$ nearest neighbors. Like figure 4-4, the data are $n = 40$ equally spaced model evaluations (orange dots).

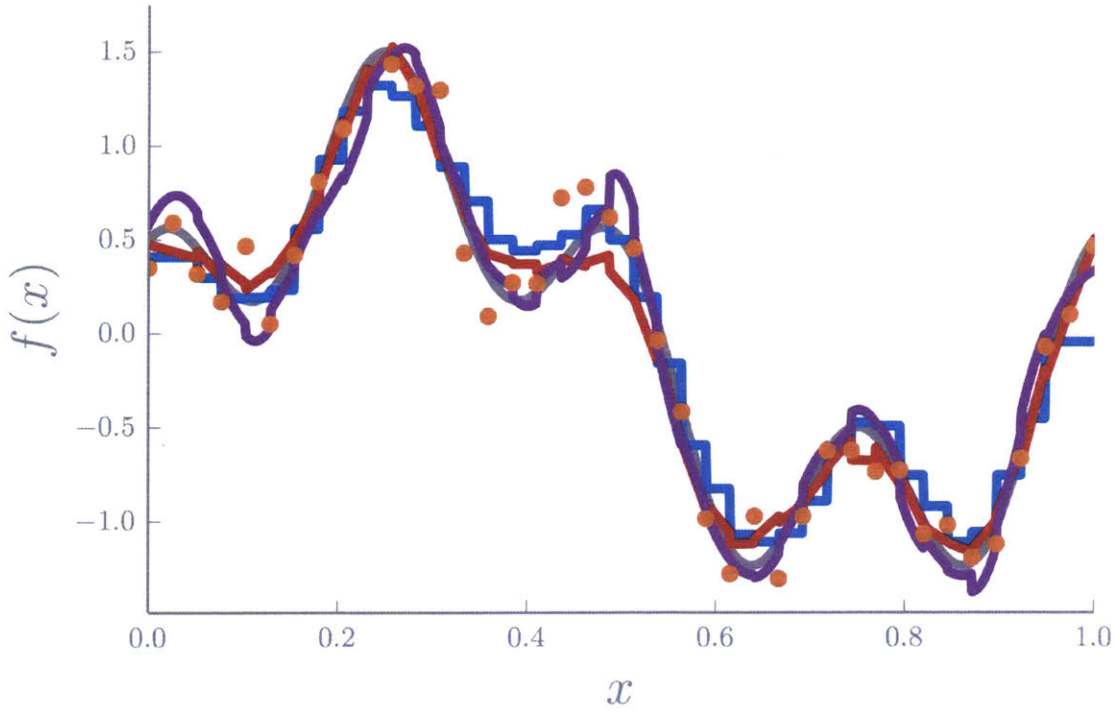

Figure 4-6: Similar estimates to those shown in figure 4-6 but here we have noisy evaluations $\hat{g}(x) = g(x) + \varepsilon$, where $\epsilon \sim \mathcal{N}(0, 2.5 \times 10^{-2})$ and the estimate is computed by solving (4.31).

for all $x \in \mathcal{X}$. We still want to construct a sequence $\{g_n\}_{n \in \mathbb{N}}$ that is asymptotically exact (i.e., for all $x \in \mathcal{X}$, $\lim_{n \to \infty} g_n(x) = g(x)$). Analogous to the estimators defined for a set of exact evaluations in (4.27), define an estimator for a set of noisy evaluations

$$g_n(x) = \underset{m \in \mathcal{P}_p}{\arg\min} \sum_{i=1}^{n} \left(m(x_i) - \hat{g}(x_i)\right)^2 \tilde{K}\left(\|x - x_i\|\right). \tag{4.31}$$

Solving (4.31) is equivalent to the solving regression problem

$$\mathbf{a}(x) = \underset{\boldsymbol{\alpha} \in \mathbb{R}^q}{\arg\min} \|\mathbf{V}\boldsymbol{\alpha} - \hat{\mathbf{g}}\|_{\mathbf{K}}^2, \tag{4.32}$$

where $\hat{\mathbf{g}}\left(\boldsymbol{x}_k\right) = [\hat{g}(x_1), \ldots, \hat{g}(x_k)]^T$, and setting $g_n(x) = \boldsymbol{\phi}^T(x)\mathbf{a}(x)$. The solution to (4.32) is $\mathbf{a}(x) = \left(\mathbf{V}^T\mathbf{K}\mathbf{V}\right)^{-1}\mathbf{V}^T\hat{\mathbf{g}}$.

We build surrogate models by applying $\left(\mathbf{V}^T\boldsymbol{K}\mathbf{V}\right)^{-1}\mathbf{V}^T$ to the vector of exact or noisy function evaluations. However, in the noisy case, guaranteeing the sequence of

estimators built with stochastic model evaluations is asymptotically exact is not as simple as adding points to the evaluated set. The kernel's radius still must decay (i.e., $\Delta \to 0$ as $n \to \infty$), but this is no longer sufficient. Satisfying $\lim_{n\to\infty} g_n(x) = g(x)$ also requires the number of points inside the ball $k \to \infty$ and $\frac{k}{n} \to 0$ as $n \to \infty$ [Conn et al., 2009, Kohler, 2002, Stone, 1977]. More qualitatively, the three convergence requirements for each $x \in \mathcal{X}$ are: (i) the kernel's support shrinks to an arbitrarily small ball around x and (ii) the number of nearest neighbors approaches infinity, but (iii) the rate the number of nearest neighbors grows must be slower than the rate new points are added to the evaluated set.

4.3.5 Error analysis

We derive a local error indicator that depends on the number of nearest neighbors, the size of the ball containing them, and the local behavior of the surrogate. For a fixed evaluated set $\mathcal{S}_n$, the expected difference between the estimators built with exact and noisy function evaluations is zero (substituting the solutions from (4.28) and (4.32))

$$\boldsymbol{\phi}^T \left(\mathbf{V}^T\mathbf{K}\mathbf{V}\right)^{-1} \mathbf{V}^T \left(\mathbb{E}\left[\hat{\mathbf{g}}\right] - \mathbf{g}\right) = 0. \tag{4.33}$$

Previous work [Conn et al., 2009, Kohler, 2002, Stone, 1977] shows that if we refine the surrogate model built with exact evaluations by adding points to the evaluated set then $|g_n(x) - g(x)| \to 0$, which implies that, for all $x \in \mathcal{X}$,

$$\mathbb{E}\left[g_n(x)\right] \to g(x) \quad \text{as} \quad n \to \infty \tag{4.34}$$

in the noisy case. However, we are only constructing a single realization of each sequence, not computing the expectation. To show a single sequence $\{g_n\}_{n\in\mathbb{N}}$ built with noisy model evaluations is asymptotically exact we also require that the variance at each $x \in \mathcal{X}$ decays as the surrogate model is refined. This variance decays as the number of points in the kernel's support increases—

$$\mathbb{V}\mathrm{ar}\left(g_n\right) = \mathbb{V}\mathrm{ar}\left(\boldsymbol{\phi}^T\left(\mathbf{V}^T\mathbf{K}\mathbf{V}\right)^{-1}\mathbf{V}^T\hat{\mathbf{g}}\right) \to 0 \quad \text{as} \quad k \to \infty \tag{4.35}$$

[Conn et al., 2009, Kohler, 2002, Stone, 1977]. Therefore, an additional requirement to ensure asymptotic exactness of our surrogate (in addition to the shrinking ball size) is $k \to \infty$. We derive the other requirements ($n \to \infty$ faster than n) by examining the rates with which (4.34) and (4.35) converge. These rates are also used to construct error indicators that depend on the number of function evaluations and the number of nearest neighbors. The expectation in (4.34) represents error due to the polynomial regression and the variance in (4.35) represents error due to the stochasticity in the function evaluations.

Lagrange polynomials

We first investigate the local behavior of polynomials in the smallest ball containing k points. Define the radius $\Delta_k(x) = \max_{\tilde{x}\in\boldsymbol{x}_k(x)}\left(\|x-\tilde{x}\|\right)$ and nearest neighborhood ball $\mathcal{B}(x;\boldsymbol{x}_k(x)) = \{\tilde{x}\in\mathcal{X} : \|x-\tilde{x}\| \leq \Delta_{k)(x)}\}$. Henceforth, we assume the kernel's support is $\mathcal{B}(x;\boldsymbol{x}_k(x))$, noting that the size of the kernel's support is not uniform in $\mathcal{X}$. We use k Lagrange polynomials

$$\lambda_j(x) = \underset{m\in\mathcal{P}_p}{\arg\min}\sum_{i=1}^{n}\left(m(x_i)-\delta_{ij}\right)^2\tilde{K}\left(\|x-x_i\|^2\right), \tag{4.36}$$

where δ_{ij} is the Dirac delta, to quantify how well behaved the estimator is inside $\mathcal{B}(x;\boldsymbol{x}_k)$. The Lagrange polynomials are defined analogously to (4.27) and (4.31). Equivalently, we can compute the Lagrange coefficients

$$\boldsymbol{\lambda}_j(x) = \underset{\boldsymbol{\alpha}\in\mathbb{R}^q}{\arg\min}\,\|\mathbf{V}\boldsymbol{\alpha}-\mathbf{e}_j\|_{\mathbf{K}}^2, \tag{4.37}$$

where $\mathbf{e}_j$ is the j^{th} Euclidean vector, and set $\lambda_j(x) = \boldsymbol{\phi}^T(x)\boldsymbol{\lambda}_j(x)$. The Lagrange coefficients are $\boldsymbol{\lambda}_j = \left(\mathbf{V}^T\mathbf{K}\mathbf{V}\right)^{-1}\mathbf{V}^T\mathbf{e}_j$; define the Lagrange polynomial vector

$$\boldsymbol{\lambda}(x) = \begin{bmatrix} \lambda_1(x) \\ \vdots \\ \lambda_k(x) \end{bmatrix} = \mathbf{V}\left(\mathbf{V}^T\mathbf{K}\mathbf{V}\right)^{-1}\boldsymbol{\phi}(x). \tag{4.38}$$

We assume the Lagrange polynomials induced by $\boldsymbol{x}_k(x)$ are Λ-poised (as defined in [Conn et al., 2009]); the set $\boldsymbol{x}_k(x)$ is Λ-poised and strongly Λ-poised in $\mathcal{B}(x;\boldsymbol{x}_k(x))$ if and only if

$$\begin{aligned} &\max_{\tilde{x}\in\mathcal{B}(x;\boldsymbol{x}_k)} \|\boldsymbol{\lambda}(\tilde{x})\|_\infty = \Lambda_\infty(x) \leq \Lambda \quad \text{and/or} \\ &\max_{\tilde{x}\in\mathcal{B}(x;\boldsymbol{x}_k)} \|\boldsymbol{\lambda}(\tilde{x})\|_2 = \Lambda_2(x) \leq \frac{q}{\sqrt{k}}\Lambda, \end{aligned} \tag{4.39}$$

for $\Lambda > 0$, respectively. By virtue of norm equivalency, assuming Λ-poisedness implies $\Lambda_2 \leq \sqrt{k}\Lambda_\infty \leq \sqrt{k}\Lambda$ and, similarly, assuming strong Λ-poisedness implies $\Lambda_\infty \leq \Lambda_2 \leq \frac{q}{\sqrt{k}}\Lambda$. If k is constant, Λ-poisedness implies strong Λ-poisedness and vice versa—a constant $\Lambda > 0$ exists such that both poisedness bounds are true. In the interpolation case, when $k = q$, any Λ that satisfies either Λ-poisedness or strong Λ-poisedness implies the other. However, if k increases than strong Λ-poisedness implies Λ-poisedness but Λ-poisedness is not enough to guarantee strong Λ-poisedness.

Λ-poisedness

Λ-poisedness measures how well spread the points are in the ball; figure 4-7 illustrates well and poorly poised sets. For example, consider points i and j in a unit ball and quadratic polynomials $p = 2$. The Lagrange polynomial corresponding to point i minimizes a least squares cost function so that it is close to one at point i and zero at point j. The parabolic Lagrange polynomial will be narrower if the points are close together and wider if they are far apart. The wider parabola will have a smaller poisedness constant than the narrower one. [Conn et al., 2009] explains that strong Λ-poisedness "reflects how well the samples are spread in space to form poised subsets." Guaranteeing strong Λ-poisedness garentuees that as k increases, the points are arranged such that the set of nearest neighbors and its subsets are well spread in

(a) Well poised (b) Poorly poised

Figure 4-7: An illustration of (a) well and (b) poorly poised points in a ball.

the nearest neighbor ball.

Error bounds

Assuming Λ-poisedness is sufficient to bound (4.34). The expectation is bounded by

$$|\mathbb{E}\left[g_n(x)\right] - g(x)| \leq \frac{\nu_p}{(p+1)!}\Lambda_\infty k\Delta^{p+1} = C_E\Lambda k\Delta^{p+1} \tag{4.40}$$

where C_E is a constant and ν_p is an upper bound on the $(p+1)^{th}$ derivative of $g(x)$ [Conn et al., 2009]. If k is constant, the radius Δ decreases as points are added to $\mathcal{S}_n$—assuming they are chosen in such a way that maintains Λ-poisedness. If k grows, the rate that points are added to $\mathcal{S}_n$ must be sufficiently fast that $k\Delta^{p+1}$ decreases.

Strong Λ-poisedness and $k \to \infty$ are both required to bound both (4.34) and (4.35). The expectation is bounded by

$$|\mathbb{E}\left[g_n(x)\right] - g(x)| \leq \frac{\nu_p}{(p+1)!}\Lambda_\infty k\Delta^{p+1} \leq \frac{\nu_p}{(p+1)!}\Lambda_2 k\Delta^{p+1} \leq C_E q\Lambda\sqrt{k}\Delta^{p+1}. \tag{4.41}$$

We refer to error in the expectation of g_n as structural error; the structural error is bounded by the geometry of the evaluated set. If k grows, points need to be added to the evaluated set faster to ensure that $\sqrt{k}\Delta^{p+1}$ decays. Therefore, we require $\frac{k}{n} \to 0$ as $n \to \infty$. The variance is bounded by

$$\mathbb{V}\mathrm{ar}\left(g_n\right) = \mathbb{V}\mathrm{ar}\left(\boldsymbol{\phi}^T\left(\mathbf{V}^T\mathbf{K}\mathbf{V}\right)^{-1}\mathbf{V}^T\hat{\mathbf{g}}\right) = \mathbb{V}\mathrm{ar}\left(\boldsymbol{\lambda}^T(x)\hat{\mathbf{g}}\right) \tag{4.42}$$

$$= \sum_{i=1}^{k}\lambda_i^2(x)\,\mathbb{V}\mathrm{ar}\left(\hat{g}(x_i)\right) = \sigma_g^2\|\lambda(x)\|_2^2 \leq \sigma_g^2\frac{q^2}{k}\Lambda^2. \tag{4.43}$$

Strong Λ-poisedness guarantees that in the noisy case the variance decays as k increases. However, if we can exactly evaluate $g(x)$ (perhaps at considerable computational cost), then Λ-poisedness is sufficient.

4.3.6 Local refinements

Given $g_n(x)$, we can improve the surrogate in a neighborhood around x by preforming a 'local refinement.' A local refinement adds a new point x^* to the evaluated set such that $\boldsymbol{x}_{n+1} = \boldsymbol{x}_n \cup \{x^*\}$ and the local error bound (4.41) decreases. The surrogate model $g_n(x)$ is constructed using the k nearest neighbors to x and the error is bounded by the geometry of the nearest neighbors. Suppose, $g : \mathbb{R}^2 \to \mathbb{R}$ and the $k = 8$ nearest neighbors to $x = [1,\ 0]$ are generated via

$$x_1 = x + \tfrac{1}{2\|s_1\|}s_1 \quad \text{and} \quad x_{i>1} = x + \tfrac{\alpha_i}{2\|s_i\|}s_i \tag{4.44}$$

for $i = 2:8$ and $s_i \sim \mathcal{U}([0,1]^2)$ and $\alpha_i \sim \mathcal{U}([0,1])$; these points are shown in figure 4-8. Randomly chosen points, such as these, tend to form clusters [Rote and Tichy, 1996], which lead to poorly poised sets. The Vandermonde matrix (4.29) is singular if $x_i = x_j$ for any $i \neq j$ and becomes nearly singular if x_i and x_j are sufficiently close to each other. We quantify how well spread the points are in the nearest neighbor ball using the poisedness constant Λ, as defined in (4.39)[6]. The Lagrange polynomial $\lambda_i(x_j) \approx \delta_{ij}$ and if x_i and x_j are nearby then $\lambda_i(x)$ achieves are larger magnitude inside the nearest neighbor ball. Therefore, the new point x^* must be chosen to maintain a well-poised set inside the nearest neighbor ball.

[6](4.39) defines $\Lambda_\infty < \Lambda$ and $\Lambda_2 < \frac{q}{\sqrt{k}}\Lambda$; we use Λ to generically refer to a constant that satisfies these inequalities and Λ_∞ or Λ_2 to specifically refer the poisedness constant associated with a nearest neighborhood ball.

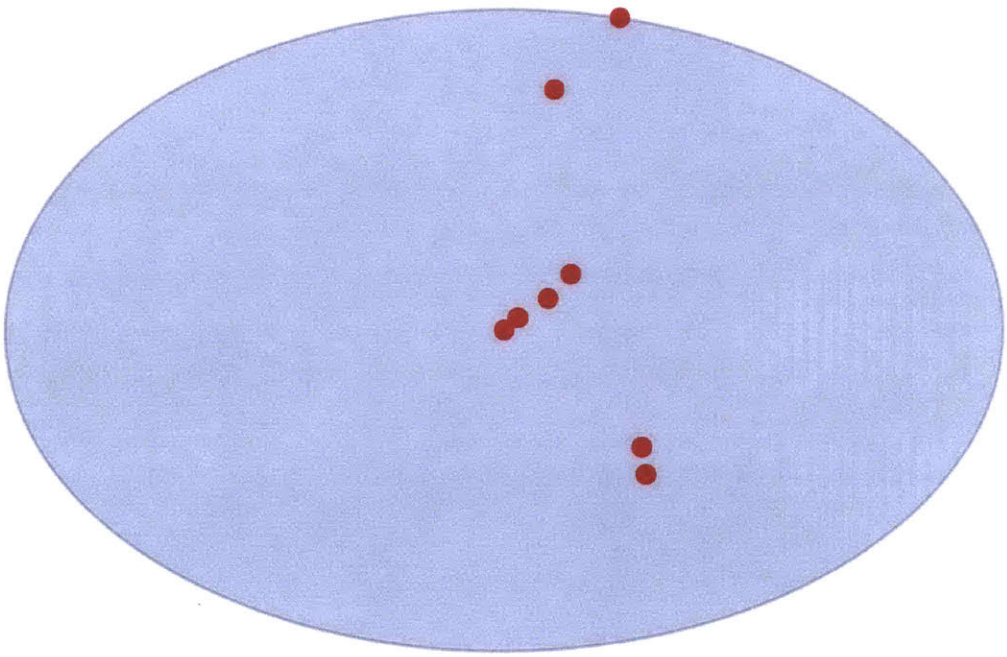

Figure 4-8: We randomly choose $k = 8$ points (red circles) from a ball of radius 0.5 (grey circle) around $x = [1, 0]$ (purple diamond) such that one point lies on the edge of the ball. Using quadratic polynomials, these points form a set with poisedness constant $\Lambda_\infty \approx 35.6$.

Computing the poisedness constant

We first compute the poisedness constant Λ_∞ given the k nearest neighbors $\boldsymbol{x}_k$ to x. We restrict the polynomials to be at most quadratic. Solving

$$\Lambda_\infty(x) = \max_{\tilde{x}\in\mathcal{B}(x;\boldsymbol{x}_k)} \|\boldsymbol{\lambda}(\tilde{x})\|_\infty = \max_{i\in 1:k} \max_{\tilde{x}\in\mathcal{B}(x;\boldsymbol{x}_k)} |\lambda_i(x)| \tag{4.45}$$

requires us to compute

$$\Lambda_\infty^{i,+} = \max_{\tilde{x}\in\mathcal{B}(x;\boldsymbol{x}_k)} \lambda_i(x) \quad \text{and} \quad \Lambda_\infty^{i,-} = \max_{\tilde{x}\in\mathcal{B}(x;\boldsymbol{x}_k)} -\lambda_i(x), \tag{4.46}$$

which amounts to maximizing a quadratic polynomial subject to ball constraints. We employ the algorithm from [Moré and Sorensen, 1983] to compute these quantities. The poisedness constant Λ_∞ is the maximum of value of $\Lambda_\infty^{i,+}$ or $\Lambda_\infty^{i,-}$. Additionally, define

$$x^{i,+} = \arg\max_{\tilde{x}\in\mathcal{B}(x;\boldsymbol{x}_k)} \lambda_i(x) \quad \text{and} \quad x^{i,-} = \arg\max_{\tilde{x}\in\mathcal{B}(x;\boldsymbol{x}_k)} -\lambda_i(x) \tag{4.47}$$

and let x_∞ be $x^{i,+}$ or $x^{i,-}$ corresponding with $\Lambda_\infty^{i,+/-} = \Lambda_\infty$. Figure 4-9 shows the location x_∞ given the randomly chosen points from figure 4-8.

Choosing the refinement location

Replacing the i^{th} point $x_i \in \boldsymbol{x}_k$—the point that corresponds with the Lagrange polynomial $\lambda_i(x)$ that achieves the value Λ_∞—with x_∞ maintains a well-poised set with k elements. Since $\lambda_i(x_j) \approx \delta_{ij}$, the maximal value x_∞ is not near the points remaining points x_j. If the original set $\boldsymbol{x}_k$ is well-poised ($\Lambda \approx 1$), then $x_\infty \approx x_i$ and the updated set remains well-poised. However, if the original set is poorly-poised ($\Lambda \gg 1$), the new point x_∞ will be in a region relatively far from x_i, where the quadratic polynomial $\lambda_i(x)$ is large. The new set $\boldsymbol{x}'_k = \boldsymbol{x}_k \cup \{x_\infty\} - \{x_i\}$ will be better poised than the original.

Adding x_∞ to the evaluated set $\mathcal{S}_n$ also reduces the local error of the surrogate model. We assume that for any x, the k nearest neighbors $\boldsymbol{x}_k \subseteq \mathcal{S}_n$ form a well poised

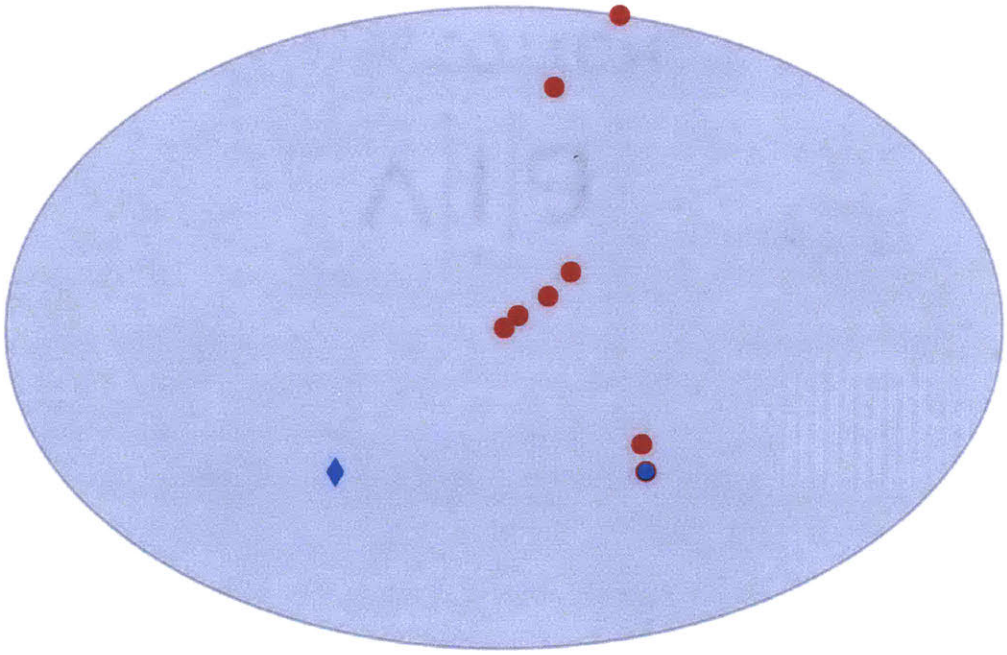

Figure 4-9: Given the $k = 8$ randomly chosen points shown in figure 4-8 (red dots), compute Λ_∞ by finding the maximum value the i Lagrange polynomial—corresponding to i^{th} point (blue dot with red outline). The point where the i^{th} Lagrange polynomial achieves its largest value x_∞ is shown as a blue diamond. Replacing the point corresponding to the Lagrange polynomial with the largest value with the location where that same Lagrange polynomial achieves is maximum value creates a new set of $k = 8$ points whose poisedness constant is $\Lambda_\infty \approx 1.8$. If we replace the farthest point from $x = [1, 0]$ (purple diamond) with x_∞, the new poisedness constant is $\Lambda_\infty \approx 3.2$.

subset. Recall, the structural error bound (4.40) implies that

$$|\mathbb{E}[g_n(x)] - g(x)| \leq C_E \Lambda k \Delta^{p+1}. \tag{4.48}$$

Updating $\boldsymbol{x}_{n+1} = \mathcal{S}_n \cup \{x^*\}$ decreases the local radius Δ, reducing the local error bound, as desired.

4.4 Local approximation sampling methodologies

Recall that our goal is to characterize the distribution π_X via a sampling method (e.g., importance sampling or MCMC). We assume either that either evaluations of the density $\pi(X)$ are computationally expensive or that we want to compute the expectation of a computationally expensive function $f(X)$. Furthermore, we may only be able to compute unbiased, noisy estimates of $\pi(X)$ or $f(X)$. We build a surrogate model for $\pi(X)$ or $f(X)$—leveraging properties of the underlying function, such as continuity—and use the computationally cheaper surrogate to reduce the computational cost of sampling. Below we discuss how to do so in the context of both importance sampling and MCMC.

4.4.1 Local approximation importance sampling (LAIS)

Suppose we want to generated weighted samples using importance sampling to estimate the expectation $\mathbb{E}_{\pi_X}[f]$. Recall, given a biasing distribution q_X, the importance sampling estimate is

$$\mathbb{E}_{\pi_X}[f] \approx \tfrac{1}{M} \sum_{i=1}^{M} f(x^{(i)}) w^{(i)} = \frac{1}{M} \sum_{i=1}^{M} f(x^{(i)}) \pi(x^{(i)}) \frac{1}{q(x^{(i)})}, \quad \text{where} \quad x^{(i)} \sim q_X. \tag{4.49}$$

If $\pi(x^{(i)})$ is computationally expensive, computing importance sampling estimates is prohibitive. For simplicity, assume we are not doing self-normalized importance sampling, but this method generalizes to this more complicated case. We, therefore,

substitute a surrogate model $\breve{\pi}(x) \approx \pi(x)$ to improve computational efficiency. If $f(x)$ is also expensive we could alternatively build $\breve{g}(x) \approx f(x)\pi(x)$ or $\breve{g}(x) \approx f(x)$. For simplicity, however, we will focus on $\breve{\pi}(x) \approx \pi(x)$.

Introducing the surrogate model

We reduce the computational cost by minimizing the number of target density evaluations $\pi(x)$. Importance sampling with no surrogate model costs M density evaluations. We, therefore, use $N < M$ density evaluations to construct $\breve{\pi}(x) \approx \pi(x)$ and then substitute the surrogate into the importance sampling estimate

$$\mathbb{E}_{\breve{\pi}_X}[f] \approx \sum_{i=1}^{M} f(x^{(i)})\breve{\pi}(x^{(i)})\frac{1}{q(x^{(i)})}, \quad \text{where} \quad x^{(i)} \sim q_X. \tag{4.50}$$

We note that, in general, $\breve{\pi}(x)$ is *not* a probability density function. However, in the limit of infinite refinement the surrogate model and the target density are equivalent. Therefore, we treat out analysis for a finite number of samples as a heuristic for the behavior of this method in finite time. If the total variation between π and $\breve{\pi}$ is bounded

$$\frac{1}{2}\int_{\mathcal{X}} |\pi(x) - \breve{\pi}(x)|\, dx \leq e_{\breve{\pi}}, \tag{4.51}$$

then we can, similarly, bound the error in the expectation

$$|\mathbb{E}_{\pi_X}[f] - \mathbb{E}_{\breve{\pi}_X}[f]| \leq \int_{\mathcal{X}} |f(x)(\pi(x) - \breve{\pi}(x))|\, dx \leq C_f e_{\breve{\pi}}, \tag{4.52}$$

where $C_f = \sup_{\mathcal{X}} |f(x)|$ is assumed finite (albeit unknown). Again, we emphasize that $e_{\breve{\pi}}$ may not be the total variation between two distributions because $\breve{\pi}$ is not necessarily a probability distribution. However, the following analysis provides heuristics on how to refine the surrogate. Replacing $\pi(x)$ with a surrogate model $\breve{\pi}(x)$ introduces a bias (4.52) that is bounded by the total variation between π_X and $\breve{\pi}_X$. Assuming π_X and $\breve{\pi}_X$ both have finite variance, the central limit theorem (4.5) guarantees that variance of the importance sampling estimates decay at rate $1/\sqrt{M}$,

$$\mathbb{V}\mathrm{ar}_{X\sim q_X}\left(\frac{1}{M}\sum_{i=1}^{M} f(X^{(i)})\frac{\pi(X^{(i)})}{q(X^{(i)})}\right) \leq \frac{C_q}{\sqrt{M}}, \tag{4.53}$$

where C_q is an unknown constant that depends on the biasing distribution q_X. The mean squared error of the importance sampling estimate—the variance plus the bias squared—is, therefore, bounded

$$\mathbb{E}_{q_X}\left[\left|\frac{1}{M}\sum_{i=1}^{M} f(X^{(i)})\frac{\pi(X^{(i)})}{q(X^{(i)})} - \mathbb{E}_{\pi_X}[f]\right|\right] \leq C_q\frac{1}{M} + C_f^2 e_{\breve{\pi}}^2. \tag{4.54}$$

We emphasize that C_q and C_f are unknown constants that depend on the biasing distribution q and the function whose expectation we are trying to estimate f. Since these are unknown, we cannot ensure that the two sources of error, MCMC variance C_q/M and structural surrogate error $C_f^2 e_{\breve{\pi}}^2$, have the same order of magnitude. However, we can guarantee that they decay at the same rate. Ideally, the bias introduced by the surrogate model does not dominate the total error. However, building a surrogate model such that $C_f^2 e_{\breve{\pi}}^2 \ll C_q M^{-1}$ is unnecessarily expensive since the error will be dominated by the importance sampling variance. This bias-variance trade-off implies that we should ideally construct a surrogate model such that $C_f^2 e_{\breve{\pi}}^2 \sim C_q M^{-1}$.

Building the surrogate model

Now, we describe how to build a local polynomial surrogate model, as described in section 4.3, such that the error at each importance sample satisfies the bias-variance trade-off. Suppose we have generated M samples from the biasing distribution $x^{(i)} \sim q_X$. Let the surrogate $\breve{\pi}(x)$ be a $p \leq 2$ degree local polynomial approximation using a kernel $K(x, x')$ whose support encompasses the k nearest neighbors to x in the evaluated set $\mathcal{S}_n$

$$\breve{\pi}(x) = \underset{m\in\mathcal{P}_p}{\arg\min} \sum_{i=1}^{k} \left(m(x_i) - \pi(x_i)\right)^2 K(x, x_i) \quad \text{for} \quad x_i \in \boldsymbol{x}_k \subseteq \mathcal{S}_n. \tag{4.55}$$

Recall, the local error bound (4.40) implies that $|\breve{\pi}(x) - \pi| \leq C_E \Lambda k \Delta^{p+1}$, assuming

that $\Lambda_\infty < \Lambda$. The total variation is, therefore, bounded by the average ball size

$$\frac{1}{2}\int_{\mathcal{X}} |\pi(x) - \breve{\pi}(x)|\, dx \leq \frac{1}{2} C_E k \Lambda \int_{\mathcal{X}} \Delta^{p+1}\, dx. \tag{4.56}$$

We, therefore, want to choose n points to build the evaluated set $\mathcal{S}_n$ such that subsets of k nearest neighbors are well-posed and the average ball size guarantees the local polynomial surrogate satisfies the bias-variance trade-off.

Refining the surrogate model

We trigger refinement based on a local error indicator. Given $\{x^{(i)}\}_{i=1}^{M}$ samples from the biasing distribution q_X, we compute the local error indicator

$$\epsilon^{(i)} = k\Lambda\Delta^{p+1}. \tag{4.57}$$

If the error indicator is larger than a user-prescribed threshold γ_0, we locally refine the surrogate model (as described in section 4.3.6). We simultaneously sample the biasing distribution and build the surrogate model. The surrogate model is becoming incrementally more accurate as we draw samples from q_X and refine $\breve{\pi}(X) \approx \pi(X)$. Therefore, we wait to compute the importance sampling weights until after drawing the importance samples. Algorithm 2 describes the algorithm in detail.

The error threshold is chosen to balance the surrogate bias squared with the importance sampling variance. Ideally, we choose γ_0 such that

$$\gamma_0^2 \sim (k\Lambda\Delta^{p+1})^2 \sim \frac{C_q}{C_f^2} M^{-1}. \tag{4.59}$$

Again, we rarely know C_q and C_f and, in practice, the fixed error threshold γ_0 must be appropriately tuned. An algorithm that refines the surrogate model when the error indicator is greater than this threshold is described in algorithm 2; we will refer to this as fixed threshold local approximation importance sampling (FT-LAIS).

Choosing $\gamma_0 = \sqrt{C/M}$ for some constant C ensures the error decays at the same rate as importance sampling with exact target density evaluations. For example,

Algorithm 2 Fixed threshold local approximation importance sampling (FT-LAIS): refine the surrogate model when the error indicator is greater than a constant error threshold.

Define the biasing distribution q_X and the error threshold $\gamma_0 > 0$.
Initialize the evaluated set $\mathcal{S}_n = \{x_1, \ldots, x_k\}$, where k is the prescribed number of nearest neighbors and x_j for $1 \leq j \leq k$ are sampled from the biasing distribution.
There are many ways to initialize the evaluated set $\mathcal{S}_n$ such that $n \geq k$. Sampling from the biasing distribution is a simple default method for our purposes.

Sampling from the biasing distribution and build the surrogate model.
for $i \leftarrow 1$ to M **do**
    Draw a sample $x^{(i)} \sim q_X$ from the biasing distribution.
    Compute the error threshold $\epsilon^{(i)} = k\Lambda(x^{(i)}(\Delta(x^{(i)})^{p+1}$.
    if $\epsilon^{(i)} > \gamma_0$ **then**
        Find x^*, as described in section 4.3.6, and set

$$\mathcal{S}_{n+1} = \mathcal{S}_n \cup \{x^*\} \tag{4.58}$$

        by evaluating the target density $\pi(x^*)$.
    end if
end for

Compute importance sampling weights using the surrogate model.
for $i \leftarrow 1$ to M **do**
    Compute the importance sampling weight $w^{(i)} = \breve{\pi}(x^{(i)})/q(x^{(i)})$
end for
Normalize the importance sampling weights $\bar{w} = \sum_{i=1}^{M} w^{(i)}$ and then $w^{(i)} = w^{(i)}/\bar{w}$

Estimate the expectation with respect to π_X.
Return the estimated expectation $\mu = \sum_{i=1}^{M} w^{(i)} f(x^{(i)})$

suppose we want to estimate the covariance of the log-quartic distribution whose density is given by (4.20). We compute a reference 'truth'

$$\mu_0^{(T)} = \sum_{i=1}^{M} w^{(i)} x^{(i)} \quad \text{and} \quad \Sigma_0 = \sum_{i=1}^{M} w^{(i)} \left(x^{(i)} - \mu_0^{(t)}\right) \left(x^{(i)} - \mu_0^{(t)}\right)^T \tag{4.60}$$

using $M = 10^7$ importance samples using a standard normal biasing distribution and exact density evaluations. We construct the importance sampling estimate Σ_T with T samples, either using exact density evaluations or using Algorithm 2, and estimate the squared error relative with the Frobenius norm

$$\epsilon_M = \frac{||\Sigma_M - \Sigma_0||_F^2}{||\Sigma_0||_F^2}. \tag{4.61}$$

Figure 4-10 plots the squared relative error as a function of the number of importance samples using exact density evaluations and using a surrogate model with error threshold $\gamma_0 = CM^{-\frac{1}{4}}$ (purple dots), $\gamma_0 = CM^{-\frac{1}{2}}$ (red dots), or $\gamma_0 = CM^{-1}$ (blue dots). The bias-variance trade-off suggests that the ideal error decay rate corresponds to $\gamma_0 = CM^{-\frac{1}{2}}$. If the decay rate is $\gamma_0 = CM^{-\frac{1}{4}}$, the surrogate is poorly refined and the error is dominated by the surrogate bias. Indeed, the purple dots in figure 4-10 shows that the error is much larger than MCMC with exact evaluations. Conversely, if $\gamma_0 = 1.4 \times 10^3 M^{-1}$, the error threshold decays much faster than the importance sampling variance. The blue dots in figure 4-10 shows that the total error decays at the same rate the MCMC variance *as a function of the number of importance samples.* Figure 4-11 shows that the total number of joint density evaluations only marginally less than the number required to use exact density evaluations. We have unnecessarily refined the surrogate and importance sampling variance is dominating the total error. In this example, we chose $C = 1.4 \times 10^3$.

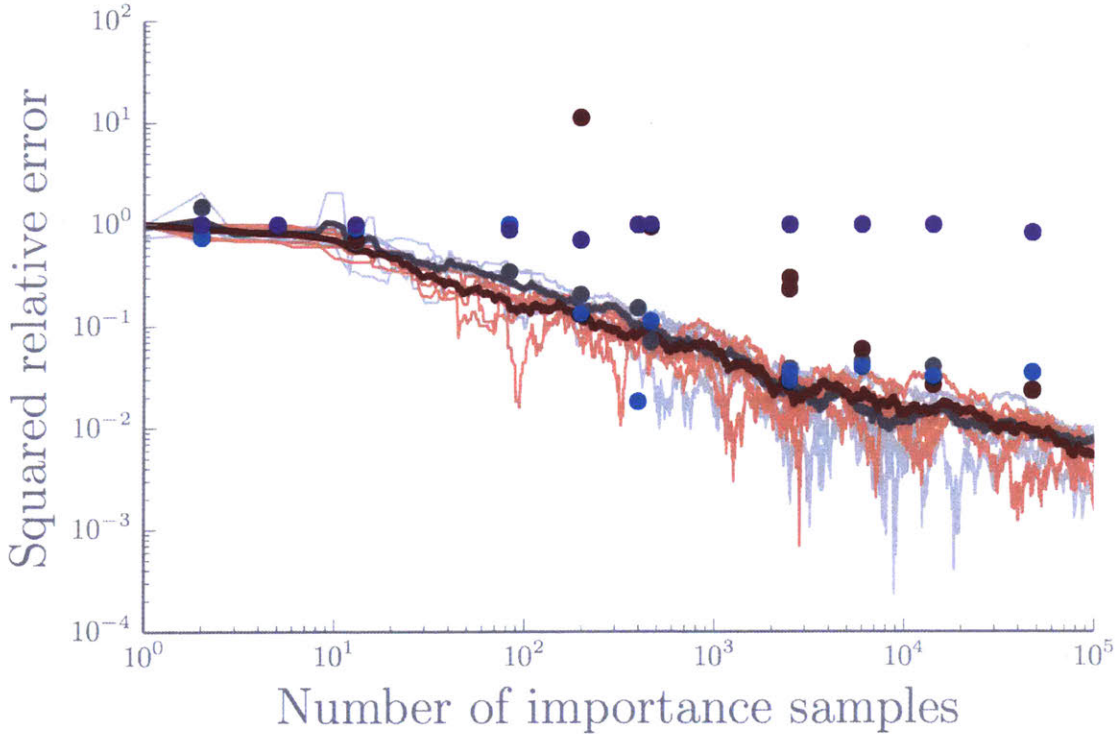

Figure 4-10: The squared relative error (4.61) for estimates of the covariance matrix of the distribution with log-quartic density (4.20) as a function of the number of *importance samples*. Grey dots correspond to importance sampling estimates with exact density evaluations and purple, red, and blue dots are estimates computed using FT-LAIS (algorithm 2) with $\gamma_0 = 1.4 \times 10^3 M^{-\frac{1}{4}}$, $\gamma_0 = 1.4 \times 10^3 M^{-\frac{1}{2}}$, and $\gamma_0 = 1.4 \times 10^3 M^{-1}$, respectively. Light grey lines correspond to importance sampling estimates with exact joint density evaluations and pink lines correspond to importance sampling estimates computed with AT-LAIS (algorithm 3) with $\tau_0 = 1$, $\gamma_0 = 1.4 \times 10^3$, and $\gamma_1 = 1$. Dark grey and red lines are the expected error using exact density evaluations and AT-LAIS, respectively.

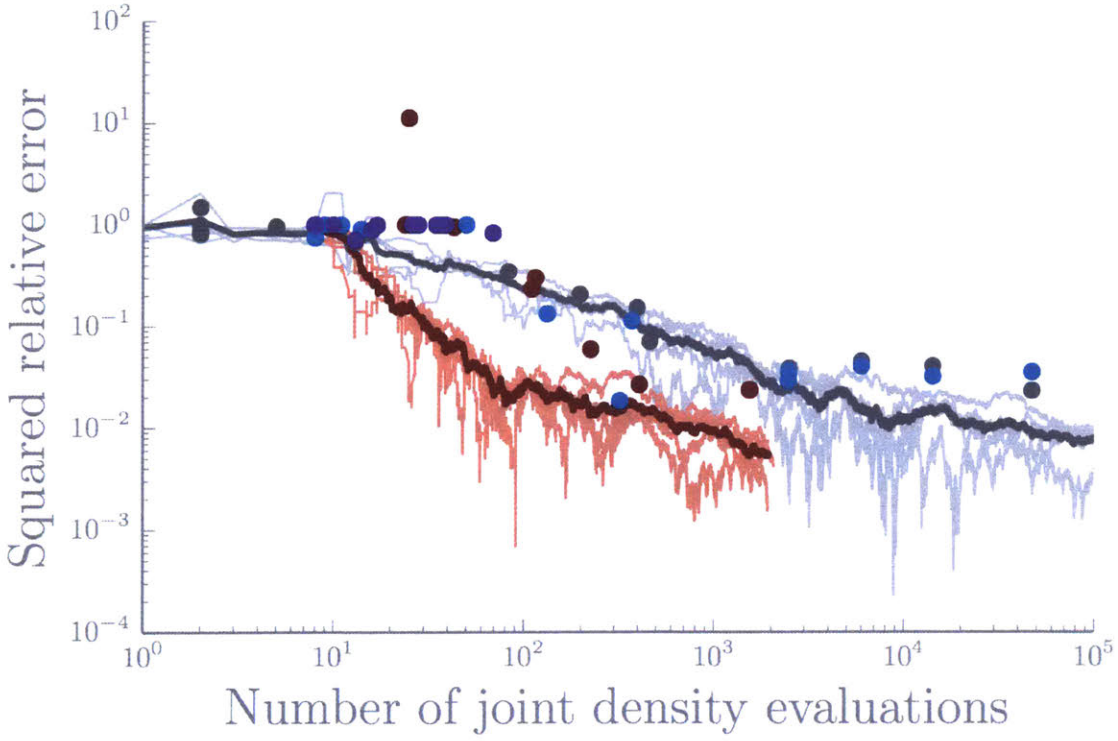

Figure 4-11: The squared relative error (4.61) for estimates of the covariance matrix of the distribution with log-quartic density (4.20) as a function of the number of *density evaluations*. Grey dots correspond to importance sampling estimates with exact density evaluations and purple, red, and blue dots are estimates computed using FT-LAIS (algorithm 2) with $\gamma_0 = 1.4 \times 10^3 M^{-\frac{1}{4}}$, $\gamma_0 = 1.4 \times 10^3 M^{-\frac{1}{2}}$, and $\gamma_0 = 1.4 \times 10^3 M^{-1}$, respectively. Light grey lines correspond to importance sampling estimates with exact joint density evaluations and pink lines correspond to importance sampling estimates computed with AT-LAIS (algorithm 3) with $\tau_0 = 1$, $\gamma_0 = 1.4 \times 10^3$, and $\gamma_1 = 1$. Dark grey and red lines are the expected error using exact density evaluations and AT-LAIS, respectively.

Adaptive error threshold

Now, we generalize Algorithm 2 to allow for the estimate to be progressively improved by adding more importance samples. Suppose we have generated M_1 importance samples and wish to refine our estimate by adding M_2 more samples. Algorithm 2 requires us to rebuild the surrogate model with ideal error threshold $\gamma_0 = \sqrt{C/(M_1 + M_2)}$. Each dot in Figures 4-10 and 4-11 rebuilds the surrogate model. However, if we have already built a surrogate with error threshold $\gamma_0 = \sqrt{C/M_1}$, discarding the evaluated set $\mathcal{S}_{n(M_1)}$ and rebuilding the surrogate, generating a new evaluated set $\mathcal{S}'_{n(M_1+M_2)}$, wastes the computational effort employed to generate $\mathcal{S}_{n(M_1)}$. Instead we can *reuse* the original target density evaluations by imposing a piecewise constant error threshold $\gamma_0 = \sqrt{C/M_1}$ for steps $i \leq M_1$ and $\gamma_1 = \sqrt{C/(M_1 + M_2)}$ for steps $M_1 < i \leq M_2$.

More generally, we divide the samples into L levels and prescribe a decaying error threshold on each level. Let $\gamma_0 > 0$ be the user-prescribed initial error threshold and γ_1 be the error threshold decay rate decays such that

$$\gamma_l = \gamma_0 l^{-\gamma_1}, \tag{4.62}$$

where l is the l^{th} level, is the decaying error threshold. The bias squared error incurred because we are imposing error threshold on the l^{th} level balances the importance sampling variance at sample T_l such that

$$(\gamma_0 l^{-\gamma_1})^2 \sim \breve{C}_2 \frac{1}{T_l}. \tag{4.63}$$

Therefore, we decrease the error threshold after generating

$$T_l = \tau_0 l^{2\gamma_1} \tag{4.64}$$

importance samples. The length of each level—equivalently, the number of samples generated with error threshold $\gamma_0 l^{-\gamma_1}$ is

$$T_{l+1} - T_l = \tau_0 \left((l+1)^{2\gamma_1} - l^{2\gamma_1} \right), \tag{4.65}$$

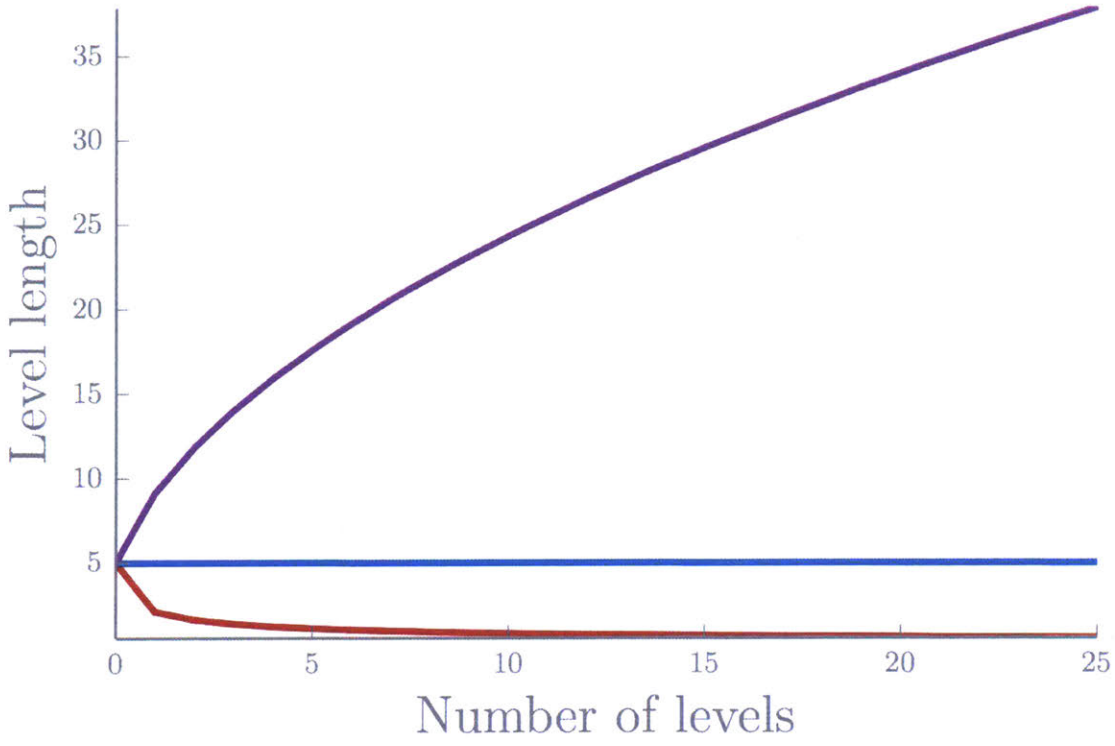

Figure 4-12: The length of each level $T_{l+1} - T_l$, where $T_1 = \tau_0 l^{2\gamma_1}$ with $\gamma_1 = 0.25$ (red line), $\gamma_1 = 0.5$ (blue line), and $\gamma_2 = 0.75$ (purple line). The initial level length is $\tau_0 = T_0 = 5$.

which is shrinking if $\gamma_1 < 0.5$, constant if $\gamma_1 = 0.5$, and increasing if $\gamma_1 > 0.5$ (figure 4-12). Figure 4-13 shows the piecewise constant error threshold imposed by this scheme. If the length of each level is shrinking ($\gamma_1 < 0.5$), then after a sufficient number of steps $T_{l+1} - T_l \ll 1$ and the level number l grows by more than one per step. In practice, we require that the length of each level grows and, therefore, assume $\gamma_1 > 0.5$. Asymptotically, this also allows the surrogate model to refine in the entire domain before decreasing the error threshold. The generalized version of FT-LAIS, which we refer to as adaptive threshold LAIS (AT-LAIS), is described in detail in Algorithm 3.

We no longer need to rebuild the surrogate model if we want to refine our estimate by adding more importance samples. The result of LAIS with a piecewise constant error threshold is shown Figures 4-10 and 4-11. Here we have chosen the initial level length to be $\tau_0 = 1$, $\gamma_0 = 1.4 \times 10^3$, and $\gamma_1 = 1$. This error threshold decay rate

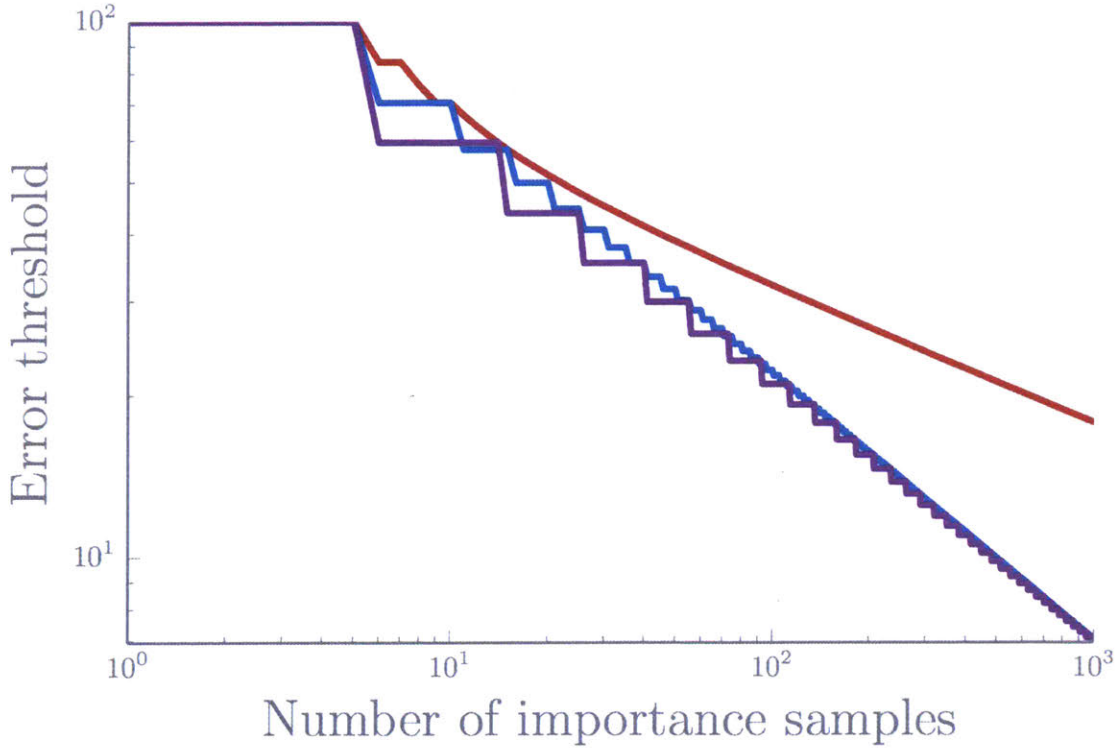

Figure 4-13: The piecewise constant error threshold $\gamma_l = \gamma_0 l^{-\gamma_1}$ with $\gamma_0 = 100$ and $\gamma_1 = 0.25$ (red line), $\gamma_1 = 0.5$ (blue line), and $\gamma_2 = 0.75$ (purple line). The initial level length is $\tau_0 = T_0 = 5$ and incrementally decrease the error threshold after $T_l = \lfloor \tau_0 l^{2\gamma_1} \rfloor$ samples.

balances surrogate bias with importance sampling variance.

Building a surrogate model using local polynomials and replacing target density evaluations with the computationally cheaper surrogate reduces the computational cost of importance sampling. The bias-variance trade-off implies an ideal error threshold and provides a method to incrementally decrease the error threshold as we accumulate more samples. However, importance sampling does *not* evaluate the target density $\pi(X)$ until *after* all of the samples have been collected and the surrogate has been built. This accounts for why the error with fewer importance samples is better in the AT-LAIS cases than when using FT-LAIS in figures 4-10 and 4-11: these are running estimates of the covariance given that we have already obtained $M = \sum_{l=1}^{L} M_l$ samples from the biasing distribution.

4.4.2 Local approximation MCMC (LA-MCMC)

Sequential algorithms, such as MCMC, generate correlated samples; we must evaluate the target density (or a surrogate) every step, *before* completely building the surrogate model. Local approximation MCMC (LA-MCMC) uses a surrogate model $\breve{\pi}(x) \approx \pi(x)$ to reduce the computational expense of MCMC while guaranteeing that the algorithm asymptotically characterizes the exact distribution-of-interest π_X. In contrast to LAIS, LA-MCMC uses the surrogate immediately (rather than computing the weights after sampling) to generated correlated samples from the target density π_X. Although, after a finite number of samples, the algorithm is targeting the approximation $\breve{\pi}_X$, which is, in general, *not* a probability distribution. [Conrad et al., 2016] guarantee that asymptotically (in the limit of infinite surrogate refinements), we characterize the target distribution π_X. However, below we develop heuristics on how to refine the surrogate by treating $\breve{\pi}_X$ as an approximation of π_X.

Each step of MCMC with exact evaluations (e.g., MH-MCMC, AM-MCMC, or MALA) consists of three steps: (i) propose a new state, (ii) compute the acceptance probability, and (iii) accept or reject the proposed state. Computing the acceptance probability costs at least one joint density evaluation; LA-MCMC replaces this expensive function call with a computationally cheaper surrogate model. Previous work

Algorithm 3 Adaptive threshold local approximation importance sampling (ATLAIS) with a piecewise constant error threshold triggering refinement.

Define the biasing distribution q_X, the initial error threshold $\gamma_0 > 0$, error threshold decay rate $\gamma_1 > 0.5$, and the initial level length $\tau_0 \geq 1$.
Initialize the evaluated set $\mathcal{S}_n = \{x_1, \ldots, x_k\}$, where k is the prescribed number of nearest neighbors and x_j for $1 \leq j \leq k$ are sampled from the biasing distribution.
There are many ways to initialize the evaluated set $\mathcal{S}_n$ such that $n \geq k$. Sampling from the biasing distribution is a simple default method for our purposes.

Sampling from the biasing distribution and build the surrogate model.
Set $M_0 = 0$.
for $l \leftarrow 1$ to L **do**
    Compute the length of the next level $M_l = \tau_0 l^{2\gamma_1}$.
    for $i \leftarrow M_{l-1}$ to M_l **do**
        Draw a sample $x^{(i)} \sim q_X$ from the biasing distribution.
        Compute the error threshold $\epsilon^{(i)} = k\Lambda\Delta^{p+1}$.
        if $\epsilon^{(i)} > \gamma_l$ **then**
            Find x^*, as described in section 4.3.6, and set

$$\boldsymbol{x}_{n+1} = \boldsymbol{x}_n \cup \{x^*\} \tag{4.66}$$

            by evaluating the target density $\pi(x^*)$.
        end if
    end for
end for
Set $M = \sum_{l=1}^{L} M_l$.

Compute importance sampling weights using the surrogate model.
for $i \leftarrow 1$ to M **do**
    Compute the importance sampling weight $w^{(i)} = \breve{\pi}(x^{(i)})/q(x^{(i)})$
end for
Normalize the importance sampling weights $\bar{w} = \sum_{i=1}^{M} w^{(i)}$ and then $w^{(i)} = w^{(i)}/\bar{w}$

Estimate the expectation with respect to π_X.
Return the estimated expectation $\mu = \sum_{i=1}^{M} w^{(i)} f(x^{(i)})$

showed that if we continuously and infinitely refine the surrogate—as described in Algorithm 4—then LA-MCMC will asymptotically characterize the true target distribution π_X [Conrad et al., 2016]. We extend this result in two significant ways: (i) we show that a bias-variance trade-off implies an ideal refinement strategy and (ii) we build surrogate models from noisy target density evaluations, which allow us to efficiently characterize marginal distributions.

Algorithm 4 A sketch of local approximation MCMC (LA-MCMC) with generic refinement strategy. If a refinement is triggered at step t, then the estimate $\breve{\pi}(x_{t-1}) \approx \pi(x_t)$ should be updated accordingly (do *not* reuse the initial surrogate estimate).

Define the initial state x_0.
Initialize the evaluated set $\mathcal{S}_n = \{x_1, \ldots, x_k\}$, where k is the prescribed number of nearest neighbors and x_j for $1 \leq j \leq k$ are sampled from a prescribed distribution.

for $t \leftarrow 1$ to ∞ **do**
    Possibly refine the surrogate model $\breve{\pi}(x_t) \approx \pi(x_t)$.
    if Refinement triggered **then**
        Find x^*, as described in section 4.3.6, and set

$$\boldsymbol{x}_{n+1} = \boldsymbol{x}_n \cup \{x^*\} \tag{4.67}$$

    by evaluating the target density $\pi(x^*)$.
    end if

    Propose a new state $\tilde{x} \sim p_{X|X=x_{t-1}}$.

    Compute acceptance probability

$$\alpha = \min\left(1, \frac{\breve{\pi}(\tilde{x})p(x_t|\tilde{x})}{\breve{\pi}(x_t)p(\tilde{x}|x_t)}\right). \tag{4.68}$$

    Accept or reject the proposed state

$$x_{t+1} = \begin{cases} \tilde{x} & \text{with probability } \alpha \\ x_t & \text{else.} \end{cases} \tag{4.69}$$

end for

Refinement strategies

The refinement strategy determines when to refine the surrogate model. We cannot afford to evaluate the true target density $\pi(x)$ at every MCMC step and, ideally, we would like the refinement frequency to decrease as the number of MCMC steps T increases. [Conrad et al., 2016] show that the algorithm asymptotically characterizes the target distribution π_X if the expected number of refinements is infinite in a chain. Next, we describe a series of different criteria that can trigger surrogate refinements: (a) poisedness refinement, (b) random refinement, and (c) the ideal bias-variance trade-off.

(a) Poisedness refinement We assume that the k nearest neighbors, which are used to construct the local polynomial approximation, are well-poised. In practice, we ensure that this by prescribing a constant Λ and triggering a refinement if $\Lambda_\infty > \Lambda$, where Λ_∞ is defined in (4.39). Typically, $\Lambda \gg 1$ and this type of refinement is rarely triggered. Our results below refine if $\Lambda_\infty > \Lambda$ but the total number of refinements triggered in this fashion[7] is typically about 25 for chains of length $T = 10^5$.

(b) Random refinement We need the expected number of refinements to be infinite, but we want the refinement probability to decay. Let the refinement probability be

$$\beta_t = \beta_0 t^{-\beta_1}, \tag{4.70}$$

where $\beta_0 > 0$ is the probability of refinement at step $t = 1$ and $0 < \beta_1 < 1$ is the decay rate. We require $\beta_1 < 1$ to ensure

[7]Not only are few refinements triggered by poisedness refinement, but the few that are typically happen in the first 1×10^3 MCMC steps. We typically sample a small number of points from some prescribed distribution (e.g., the prior in the Bayesian case) to initialize the evaluated set. This leads to poorly poised points, which are quickly corrected by poisedness refinement. Furthermore, poisedness refinement almost never happens if we refine based on the bias-variance trade-off; presumably because the error indicator includes the poisedness constant.

$$\sum_{t=1}^{\infty} \beta_t = \infty \tag{4.71}$$

and, therefore, that the expected number of refinements $N \to \infty$ as the number of MCMC steps $T \to \infty$. Figure 4-14 shows samples from the joint distribution using LA-MCMC with random refinement and Figure 4-15 shows the expected number of refinements for various values of β_1. We have chosen algorithm parameters such that the refinement probability is decaying and the expected number of refinements is infinite; all of these examples will be *asymptotically* exact. However, if the decay rate is sufficiently slow, then after a finite number of steps the number of model evaluations is about the same order of magnitude as the number of MCMC steps (blue line of 4-14). Conversely, if the decay rate is too fast we tend to over-refine the surrogate model (purple line of 4-14). This suggests a bias-variance trade-off may imply an ideal refinement rate.

We note that for all of these examples $\beta_1 < 1$ and, therefore, the expected number of refinements is infinite. Previous work by [Conrad et al., 2016] show that this is guarantees that the algorithm will eventually characterizes the distribution-of-interest. However, they also, show, and we confirm here, that for particular choices of (valid) algorithm parameters, performance can be poor after a *finite* number of MCMC steps.

The refinement rate β_1 significantly affects error in the estimate of $\mathbb{E}_{\pi_X}[f]$. Again, consider the distribution with log-quartic density (4.20). Similar to the LAIS example, we estimate the second moment using the correlated MCMC samples

$$\mu_0^{(T)} = \tfrac{1}{T} \sum_{t=1}^{T} x_t \quad \text{and} \quad \Sigma_0 = \tfrac{1}{T-1} \sum_{t=1}^{T} \left(x_t - \mu_0^{(t)}\right) \left(x_t - \mu_0^{(t)}\right)^T. \tag{4.72}$$

We compute a reference 'truth' using AM-MCMC with exact evaluations and $T = 10^7$ samples and, again, define the estimated squared error via (4.61). Figure 4-16 shows the squared relative error as a function of MCMC steps. The expected error using exact evaluations defines a lower bound on LA-MCMC's expected error; the

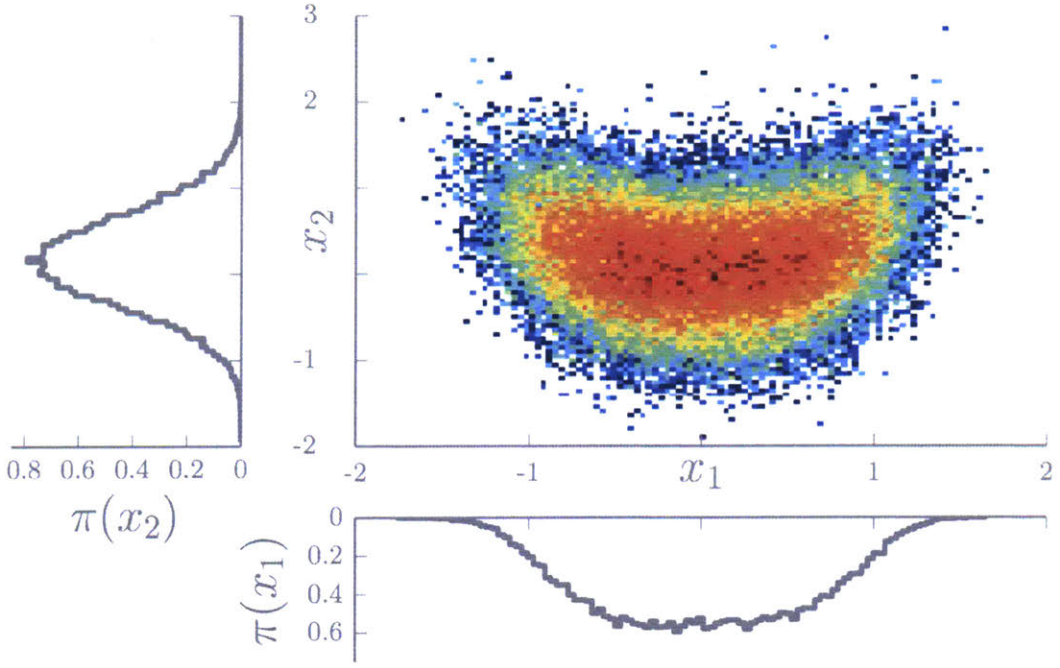

Figure 4-14: $T = 10^5$ samples from the distribution whose associated density is (4.20) generated using LA-MCMC with random refinement ($\beta_0 = 1$ and $\beta_1 = 0.1$) and poisedness refinement ($\Lambda = 25$), building locally quadratic polynomials with $k = 8$ nearest neighbors. These samples characterize the same density as those shown in figure 4-3.

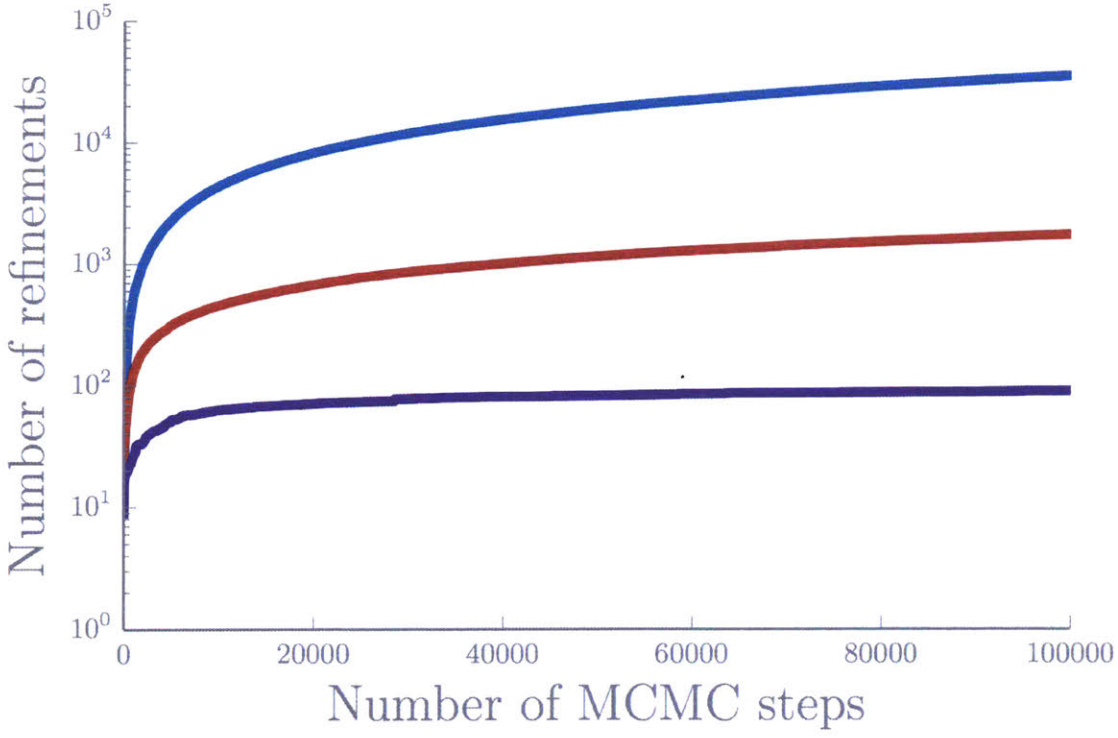

Figure 4-15: The number of refinements—which are equivalent to joint density evaluations in this case—as a function of MCMC steps. Light blue, pink, and light purple lines are single LA-MCMC chains with $\beta_0 = 1$ and $\beta_1 = 0.1$, $\beta_1 = 0.4$, and $\beta_1 = 0.8$, respectively. Darker bold lines are the expected values given these algorithmic parameters.

additional bias from the surrogate model can only increase the total error. However, we can achieve errors of the same order of magnitude by appropriately tuning the refinement rate β_1. If β_1 is too large (purple lines in figure 4-16), the surrogate bias dominates the error, slowing the error decay rate. If surrogate refinement is too infrequent, the chain can occasionally walk into the tails of the distribution. Although, asymptotically randomly triggered refinement will drive the chain back into regions of high probability density, this rarely happens in a reasonable (less than 10^6) number of MCMC steps. However, sufficiently small β_1 (blue and red lines in figure 4-16) achieve similar errors as the exact evaluation case. Figure 4-17 shows the squared relative error as a function of target density evaluations. If β_1 is too small (blue line in figure 4-17), LA-MCMC refines $\breve{\pi}(X)$ too frequently. This reduces the surrogate bias and causes the error to be dominated by the MCMC variance. Increasing β_1 achieves the same sized error using fewer target density evaluations (red line in figure 4-17). This implies a bias-variance trade-off, similar to LAIS.

Random refinement is oblivious to structure in surrogate model. If the MCMC chain is exploring a region of the target density where the surrogate error is large, random refinement cannot identify this feature and trigger a refinement. In figure 4-16 the error of individual chains as a function of MCMC steps can suddenly increase. This is because the chain has wandered into a region where $|\breve{\pi}(x) - \pi(x)| \gg 1$—often this occurs in the tails of the target distribution. Such an error cannot be corrected until a refinement is randomly triggered. Asymptotically this is guaranteed to happen, however, after a finite number of steps without refinement could lead to significant errors. Conversely, random refinement may trigger target density evaluations where $|\breve{\pi}(x) - \pi(x)| \ll 1$, which wastes computational resources improving the surrogate when refinement is not necessary.

(c) Bias-variance trade-off The MCMC estimate has a bias and a variance after a finite number of steps T—even with no approximation. We assume the MCMC bias is unknown and the MCMC variance is bounded by

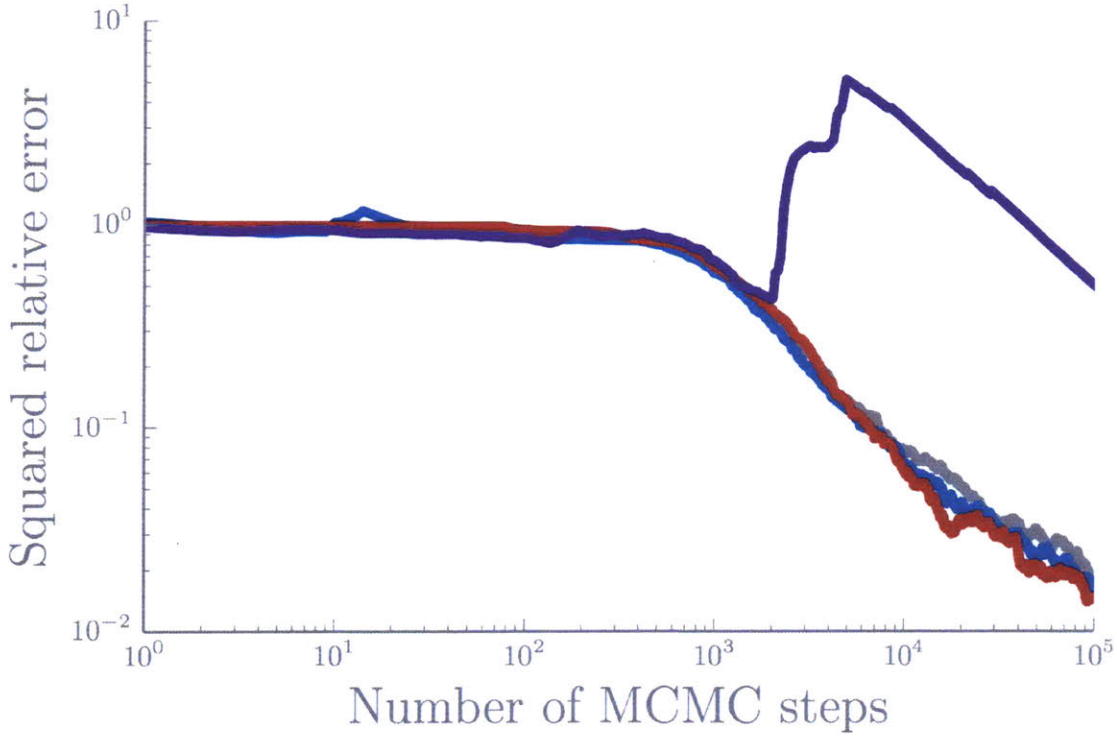

Figure 4-16: The squared relative error (4.61) as a function of MCMC steps. Light grey lines are a single chain of AM-MCMC with exact evaluations, light blue, pink, and light purple lines are single changes of LA-MCMC with $\beta_0 = 1$ and $\beta_1 = 0.1$, $\beta_1 = 0.4$, and $\beta_1 = 0.8$, respectively. In all cases, we are building locally quadratic polynomials with $k = 8$ nearest neighbors. Darker bold lines are the expected values given these algorithmic parameters. The exact evaluation case (grey line) provides a lower bound on the expected error, however, sufficiently small β_1 achieve errors of the same order-of-magnitude. We note that, despite seemingly poor performance for $\beta_1 = 0.8$, all of these are valid choices in the sense that they asymptotically characterize the distribution-of-interest.

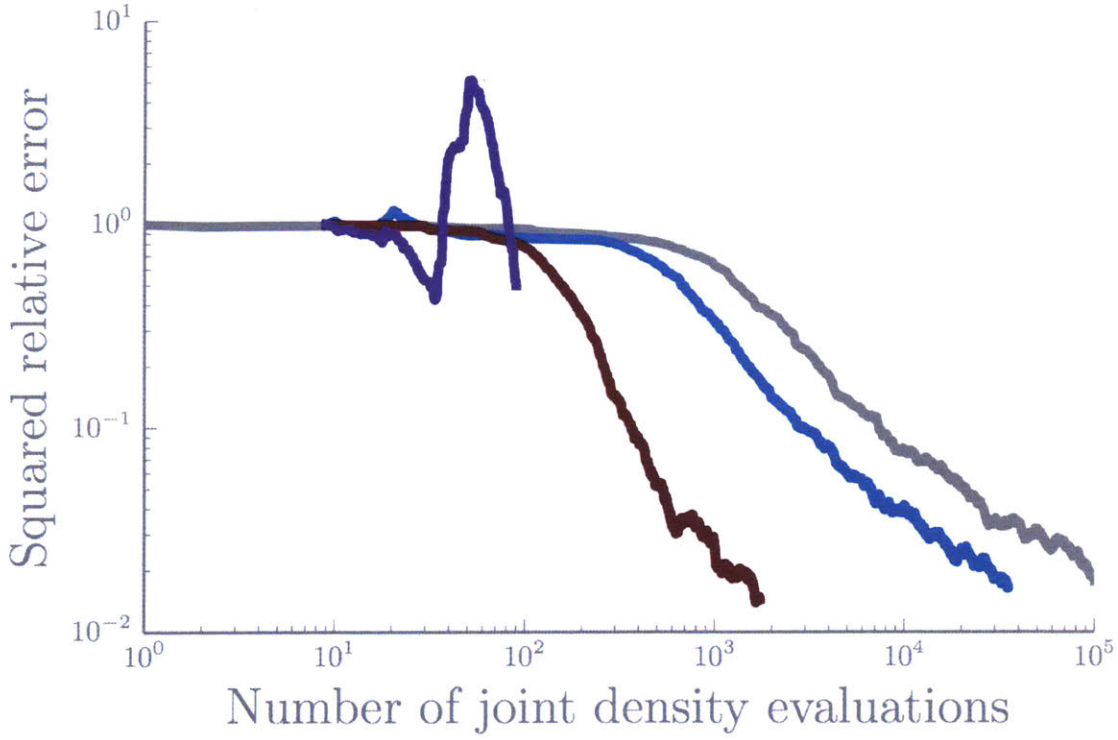

Figure 4-17: The squared relative error (4.61) as a function of target density evaluations. Light grey lines are a single chain of AM-MCMC with exact evaluations, light blue, pink, and light purple lines are single changes of LA-MCMC with $\beta_0 = 1$ and $\beta_1 = 0.1$, $\beta_1 = 0.4$, and $\beta_1 = 0.8$, respectively. Darker bold lines are the expected values given these algorithmic parameters. If β_1 is too small (blue lines), the expected number of true target density evaluations is about the same as MCMC with exact evaluations.

$$\mathbb{V}\mathrm{ar}_{\pi_X}\left(\frac{1}{T}\sum_{t=1}^{T} f(x_t)\right) \leq C_\pi \frac{1}{T}. \tag{4.73}$$

We emphasize that this assumption requires exact target density evaluations—the right hand side of this inequality will likely be infinite using the surrogate model. Recall (4.56), which states that the bias introduced by the surrogate model is bounded by the average ball size. Like LAIS, we want the MCMC variance and the surrogate bias to be the same order of magnitude. If the additional error from the MCMC bias dominates the total error, then introducing the surrogate model does not make the estimate worse. We divide the chain into L levels and prescribe a piecewise constant error threshold

$$\gamma_l = \gamma_0 l^{-\gamma_1} \tag{4.74}$$

for each level l. We decrease the refinement threshold when the MCMC variance balances the surrogate bias; we increment $l \leftarrow l + 1$ at step

$$T_l = \tau_0 l^{2\gamma_1}, \tag{4.75}$$

where τ_0 is length of the first level. Continuously incrementing the refinement threshold as a function of MCMC steps t, rather than defining a piecewise constant threshold, requires

$$\gamma_t = \gamma_0^2 t^{-2\gamma_1} \sim C_\pi t^{-1}. \tag{4.76}$$

This is equivalent to setting $\tau_0 = 1$, $\gamma_0 = \sqrt{C_\pi}/C_f$ (C_f and C_π are the constants defined in (4.53)), and $\gamma_1 = 0.5$ in the piecewise constant framework. However, C_f and C_π are rarely known and, in practice, finding an appropriate choice for γ_0 is difficult. Setting $\tau_0 \geq 1$ and choosing $\gamma_0 \approx k\Lambda\Delta^{p+1}$ at the starting point is heuristically more robust. We also require that $\gamma_1 > 0.5$ so that the length of each level increases and the error threshold decreases at most once per MCMC step (see figures 4-12 and 4-13). Figure 4-18 shows samples from the log-quartic density (4.20) using the bias-variance

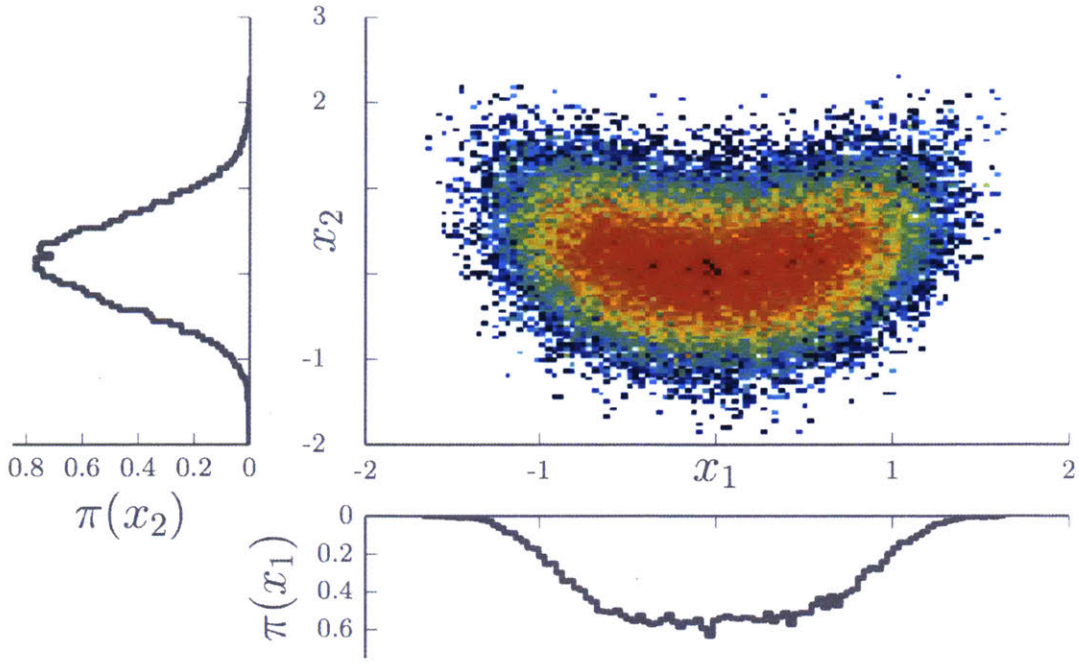

Figure 4-18: $T = 10^5$ samples from the distribution whose associated density is (4.20) generated using LA-MCMC with bias-variance trade-off thresholding (defined by (4.77) and (4.75) with $\tau_0 = 1$, $\gamma_0' = 2$, $k = 8$, $\Lambda = 50$, and $\gamma_1 = 1$). These samples characterize the same density as those shown in figure 4-3.

trade-off thresholding.

Practical choices of γ_0 The ideal value of the constant γ_0 is unknown, but any choice of γ_0 guarantees that the *rate* of surrogate bias squared and MCMC variance decay are the same. In practice, however, a poor choice for γ_0 can have significant effect on LA-MCMC's performance. Therefore, we prescribe the error threshold

$$\gamma_l = k\Lambda\gamma_0' l^{-\gamma_1}, \tag{4.77}$$

where k is the number of nearest neighbors and Λ is the prescribed poisedness threshold. This is equivalent to prescribing $\gamma_0 = k\Lambda\gamma'$; however, we find that choosing k and Λ, then scaling by γ_0' is easier in practice. Intuitively, assuming we maintain Λ-poisedness, this error threshold requires that the ball size raised to the $(p+1)^{th}$

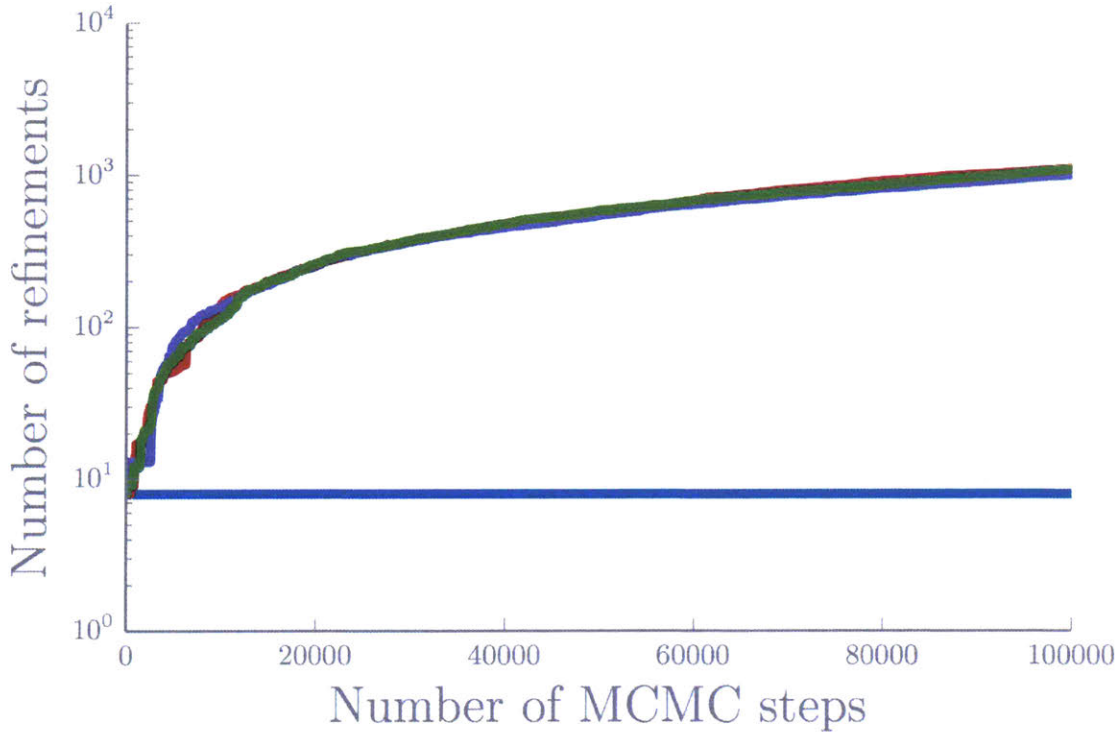

Figure 4-19: The number of refinements—which are equivalent to joint density evaluations in this case—as a function of MCMC steps. Light blue, pink, light purple, and light green lines are single changes of LA-MCMC with $\gamma_1 = 0.25$, $\gamma_1 = 0.5$, $\gamma_1 = 0.75$, and $\gamma_1 = 1$, respectively. We define the piecewise constant error threshold (4.77) with $\gamma_0' = 2$, $k = 8$, and $\Lambda = 50$. We increment the refinement threshold according to (4.75) with $\tau_0 = 1$. Darker bold lines are the expected values given these algorithmic parameters.

power[8] Δ^{p+1} decays according to $\gamma_0' l^{-\gamma_1}$.

Figure 4-19 shows the number of refinements as a function of MCMC steps is less than the number of refinements required by random refinement. Structural refinement—refining based on a bias-variance trade-off—takes into account the quality of the surrogate model. We do not evaluate the expensive target density $\pi(X)$ if the surrogate error indicator is small. We also note that the refinement frequency is *decreasing* as a function of MCMC steps.

For any $\gamma_1 > 0.5$, the bias-variance trade-off ensures that the estimate error will decrease at the same rate as the MCMC variance. Figure 4-20 shows the squared relative error of estimates of the covariance matrix for the joint distribution with quartic

[8]Recall that p is the degree of the local polynomial.

density (4.20), as defined by (4.61). Again, AM-MCMC with exact evaluations defines a lower bound on the expected error—introducing a surrogate can only increase the error. However, the expected error when $\gamma_1 > 0.5$ is of the same order of magnitude as the exact case. Figure 4-21 shows the squared relative error as a function of the number of target density evaluations. Furthermore, the expected number of target density evaluations when $\gamma_1 > 0.5$ is constant. If $\gamma_1 < 0.5$, the error threshold does not decay quickly enough to guarantee the bias-variance trade-off (figure 4-13) and the error is dominated by the surrogate bias. As a function of MCMC steps, the error decays more slowly than the exact case but as a function of model evaluations the error decays more quickly.

4.4.3 LA-MCMC targeting marginal distributions

Suppose our goal is to characterize the marginal distribution of some quantity of interest X. If we can evaluate the joint density of X and an auxiliary random variable Z, then the target density is

$$\pi(x) = \int_{\mathcal{Z}} \pi(x, z)\, dz. \tag{4.78}$$

The sets $\mathcal{X}$ and $\mathcal{Z}$ partition the joint parameter space such that Z are unimportant for our application. Typically, we do not have an analytic expression for $\pi(x)$ but we do have an unbiased estimator $\hat{\pi}(X)$, via, for example an importance sampling approximation. Section 4.3.4 describes how to build asymptotically exact surrogate models using these noisy estimates of the target density. Here, we extend LA-MCMC to target marginal distributions.

Local approximation MCMC targeting marginal distributions uses the same algorithm as the exact case. We use three possible mechanisms to trigger a local refinement: (i) the poisedness constant is too large, (ii) randomly with decaying probability, or (iii) to balance a bias-variance trade-off threshold. These refinement strategies are designed such that the error indicator

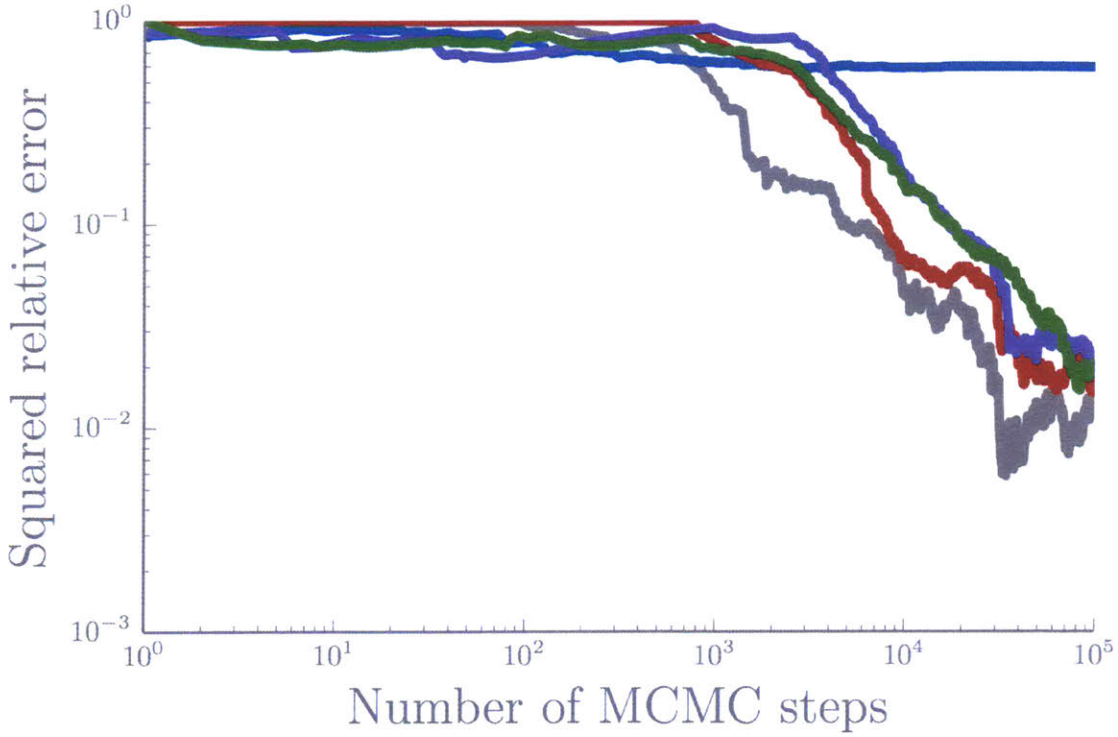

Figure 4-20: The squared relative error (4.61) as a function of MCMC steps. Light grey lines are single chains of AM-MCMC with exact evaluations, light blue, pink, light purple, and light green lines are single chains of LA-MCMC with $\gamma_1 = 0.25$, $\gamma_1 = 0.5$, $\gamma_1 = 0.75$, and $\gamma_1 = 1$, respectively. We define the piecewise constant error threshold (4.77) with $\gamma_0' = 2$, $k = 8$, and $\Lambda = 50$. We increment the refinement threshold according to (4.75) with $\tau_0 = 1$. Darker bold lines are the expected values given these algorithmic parameters. When $\gamma_1 < 0.5$, the error threshold decays slower than the expected MCMC variance decay, therefore, the error is dominated by the surrogate bias.

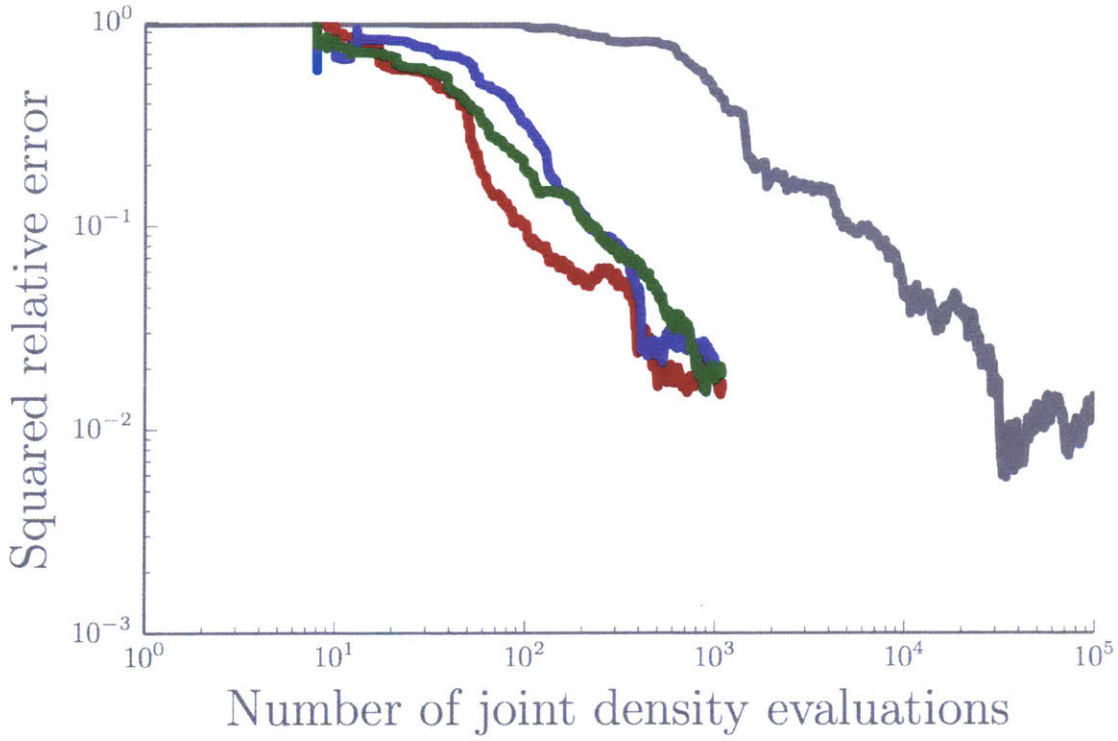

Figure 4-21: The squared relative error (4.61) as a function of target density evaluations. Light grey lines are single chains of AM-MCMC with exact evaluations, light blue, pink, light purple, and light green lines are single chains of LA-MCMC with $\gamma_1 = 0.25$, $\gamma_1 = 0.5$, $\gamma_1 = 0.75$, and $\gamma_1 = 1$, respectively. We define the piecewise constant error threshold (4.77) with $\gamma_0' = 2$, $k = 8$, and $\Lambda = 50$. We increment the refinement threshold according to (4.75) with $\tau_0 = 1$. Darker bold lines are the expected values given these algorithmic parameters. When $\gamma_1 < 0.5$, the surrogate bias dominates the error, therefore, refining the model is more valuable than preforming more MCMC steps.

$$k\Lambda \int_{\mathcal{X}} \Delta^{p+1}\, dx \to 0 \tag{4.79}$$

as the number of refinements $N \to \infty$. We suggest that Algorithm 4, building $\breve{\pi}(x)$ using the unbiased estimator $\hat{\pi}(x)$, asymptotically characterizes the target distribution π_X.

Similar to pseudo-marginal MCMC (PM-MCMC) [Andrieu and Roberts, 2009, Beaumont, 2003], we can interpret LA-MCMC with noisy target density evaluations as a Markov chain targeting an augmented space. We provide a sketch of this argument below, but defer a formal proof to later work. Recall that, in the case of PM-MCMC, we can interpret the chain as targeting an augmented state that includes x_t and the importance samples $y_t^{(1:M)}$ used to compute $\hat{\pi}(x_t)$ (see figure 4-2). We use a similar interpretation, except that we do not update $y_t^{(1:M)}$ every MCMC step. Therefore, we consider a sub-chain that only consists of steps where we have refined the surrogate model (see figure 4-22). This sub-chain also defines a Markov chain on an augmented space.

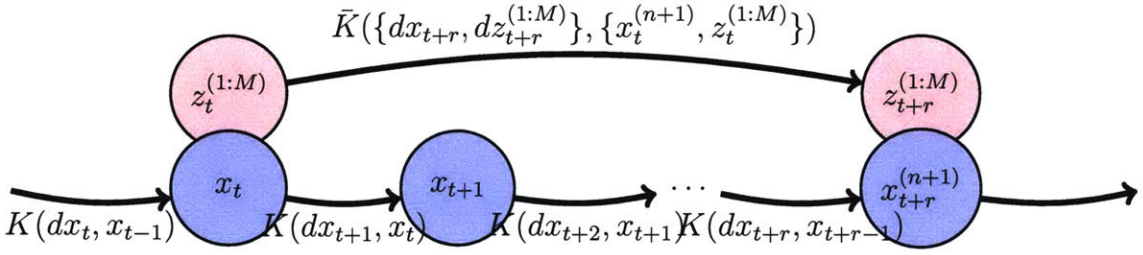

Figure 4-22: A visual representation of a LA-MCMC chain with transition kernel $K(\cdot, \cdot)$ targeting a marginal distribution. We view the sub-chain consisting of steps where the outer chain as refined as a Markov chain targeting an augmented space.

Algorithm 4 constructs a Markov chain $\{x_t\}_{t=1}^{\infty}$ such that the transition kernel from x_i to x_j is

$$K(dx_j, x_i) = \int_{\mathcal{X}} p(x|x_i)(1 - \alpha(x_i, x))\, dx + \alpha(x_i, dx_j)p(dx_j|x_i), \tag{4.80}$$

where $\alpha(x_i, dx_j)$ is the acceptance probability

$$\alpha(x_i, x_j) = \min\left(1, \frac{\breve{\pi}^{(n)}(x_j)p(x_i, x_j)}{\breve{\pi}^{(n)}(x_i)p(x_j x_i)}\right) \tag{4.81}$$

and $p_{X|X}$ is the proposal distribution. Therefore, the probability density of moving from point x_t to x_{t+r} in r steps is

$$f(x_{t+r}|x_t) = \int_{\mathcal{X}} \cdots \int_{\mathcal{X}} \left(\prod_{i=t}^{t+r} K(dx_{i+1}, x_i)\right) dx_{t+2} \ldots dx_{t+r-1}. \tag{4.82}$$

The surrogate model $\breve{\pi}^{(n)}(x)$ depends on n, the number of times we have evaluated the estimator $\hat{\pi}(x)$, and we periodically refine by incrementing $n \leftarrow n+1$ such that $n \to \infty$ as the number of MCMC steps $t \to \infty$. Let $\{x^{(n)}, z^{(n,1)}, \ldots, z^{(n,M)}\}_{n=1}^{\infty}$ be a sub-chain such that $x^{(n)} = x_{t'(n)}$, where $t'(n)$ is the step where we evaluated the estimator $\hat{\pi}(x)$ for the n^{th} time. The sub-chain also includes $\{z^{(n,i)}\}_{i=1}^{M}$, which define the augmented space. Note that $z^{(n,i)}$ is a d_z dimensional quantity and, therefore, the augmented space is $d_x + Md_z$ dimensional. Recall, $q_{Z|X}$ is the biasing distribution such that

$$\hat{\pi}(x^{(n)}) \approx \frac{1}{M}\sum_{i=1}^{M} \frac{\pi(x^{(n)}, z^{(n,i)})}{q(z^{(n,i)}|x^{(n)})}. \tag{4.83}$$

Algorithm 5 defines a Metropolis-Hastings MCMC algorithm that samples this augmented space. We note that the acceptance probability (4.90) in algorithm 5 simplifies to

$$\alpha = \min\left(1, \frac{\breve{\pi}^{(n+1)}(\tilde{x})f(x_n|\tilde{x})}{\breve{\pi}^{(n+1)}(x^{(n)})f(\tilde{x}, x^{(n)})}\right), \tag{4.84}$$

where we are now interpreting f from (4.82) as the proposal density on the sub-chain. Figure 4-22 shows one step of the sub-chain and algorithm 5 outlines the MCMC algorithm for the sub-chain whose stationary distribution is $\pi_{X,Z^{(1)},\ldots,Z^{(M)}}$ with density

$$\pi(x, z^{(1)}, \ldots, z^{(M)}) = \pi(x)\prod_{i=1}^{M} q(z^{(i)}|x). \tag{4.85}$$

At every step n of the sub-chain we refine the surrogate density

$$\pi(x, z^{(1)}, \ldots, z^{(M)}) \approx \breve{\pi}(x, z^{(1)}, \ldots, z^{(M)}) = \breve{\pi}(x) \prod_{i=1}^{M} q(z^{(i)}|x). \tag{4.86}$$

Therefore, the sub-chain is itself a local approximation algorithm targeting an augmented space that includes the QoI and the importance samples. In the case of the sub-chain, we refine every step. The samples generated by LA-MCMC building the surrogate $\breve{\pi}(x)$ from noisy target density evaluations, therefore, asymptotically characterize the target marginal π_X.

We compare the performance of LA-MCMC targeting a marginal distribution to pseudo-marginal MCMC. Figure 4-23 shows the marginal distribution π_{X_1} of the joint density defined by 4.20 characterized with both PM-MCMC and LA-MCMC; the results are indistinguishable. Figure 4-24 shows the number of refinements required for various values of γ_1. We also compare estimates of the variance of this marginal distribution. We compute a reference 'truth' σ_0^2 via (4.60) using pseudo-marginal MCMC with 10^6 MCMC steps and an importance sampling estimate of the target density with a standard normal biasing distribution and $M = 100$ importance samples. We compare the squared error using LA-MCMC with T samples

$$\epsilon_T = \frac{(\sigma_T^2 - \sigma_0^2)^2}{\sigma_0^2}, \tag{4.92}$$

where σ_T^2 is the LA-MCMC estimate.

The error decay as a function of the algorithm parameters γ_0 and γ_1 is very similar LA-MCMC where we build the surrogate from exact evaluations. Figures 4-25 and 4-26 show squared error as a function of the number of MCMC steps and the number of target density evaluations. We see that for $\gamma_1 > 0.5$ the algorithm balances surrogate bias with MCMC variance and the that the algorithm under refines if $\gamma_1 < 0.5$ because the error threshold does not decay fast enough. This only difference between this and the previous result is that the number of joint density evaluations is now multiplied by a factor of $M = 10$, which is the number of importance samples used to estimate the marginal density.

Algorithm 5 Local approximation MCMC (LA-MCMC) targeting an augmented state space consisting of the QoI and the importance samples.

1: Define the initial state $\{x^{(0)}, z^{(0,1)}, \dots z^{(0,M)}\}$.
2: Initialize the evaluated set $\mathcal{S}_n = \{x_1, \dots, x_k\}$, where k is the prescribed number of nearest neighbors and x_j for $1 \leq j \leq k$ are sampled from a prescribed distribution.
3: **for** $n \leftarrow 1$ to ∞ **do**
4: Determine the number of sub-steps $r \sim p_R^{(n)}$ such that $\mathbb{E}_{p_R}[r] \to \infty$ as $n \to \infty$.
5: *Propose $\tilde{x}$ from an 'outer' proposal distribution that amounts to r Metropolis-Hastings sub-steps.*
6: Set $x_1 = x^{(n)}$.
7: **for** $t \leftarrow 1$ to $r - 1$ **do**
8: Propose a new state $\check{x} \sim p_{X|X=x_t}$.
9: Compute acceptance probability

$$\alpha_t = \min\left(1, \frac{\check{\pi}^{(n)}(\check{x})p(x_t|\check{x})}{\check{\pi}^{(n)}(x_t)p(\check{x}|x_t)}\right). \tag{4.87}$$

10: Accept or reject the proposed state

$$x_{t+1} = \begin{cases} \check{x} & \text{with probability } \alpha \\ x_t & \text{else.} \end{cases} \tag{4.88}$$

11: **end for**
12: Set $\tilde{x} = x_r$.
13: Propose $\tilde{z}^{(\cdot,i)} \sim q_{Z|X=x^{(n)}}$ for $i = 1, \dots, M$.
14: Find x^*, as described in section 4.3.6, and set

$$\boldsymbol{x}_{n+1} = \boldsymbol{x}_n \cup \{x^*\} \tag{4.89}$$

by evaluating the target density estimation $\hat{\pi}(x^*)$ using $\{z^{(n,i)}\}_{i=1}^M$ importance samples.
15: Compute acceptance probability

$$\alpha = \min\left(1, \frac{\left(\check{\pi}^{(n+1)}(\tilde{x})\prod_{i=1}^M q(\tilde{z}^{(\cdot,i)}|\tilde{x})\right)\left(f(x_n|\tilde{x})\prod_{i=1}^M q(z^{(n,i)}|x^{(n)})\right)}{\left(\check{\pi}^{(n+1)}(x^{(n)})\prod_{i=1}^M q(z^{(n,i)}|x^{(n)})\right)\left(f(\tilde{x}, x^{(n)})\prod_{i=1}^M q(\tilde{z}^{(\cdot,i)}|\tilde{x})\right)}\right), \tag{4.90}$$

where the outer proposal density $f(x^{(i)}|x^{(j)})$ is defined in (4.82).
16: Accept or reject the proposed state

$$\{x^{(n+1)}, z^{(n+1,1)}, \dots z^{(n+1,M)}\} = \begin{cases} \{\tilde{x}^{(n+1)}, \tilde{z}^{(n+1,1)}, \dots \tilde{z}^{(n+1,M)}\} & \text{with probability } \alpha \\ \{x^{(n)}, z^{(n,1)}, \dots z^{(n,M)}\} & \text{else.} \end{cases} \tag{4.91}$$

17: **end for**

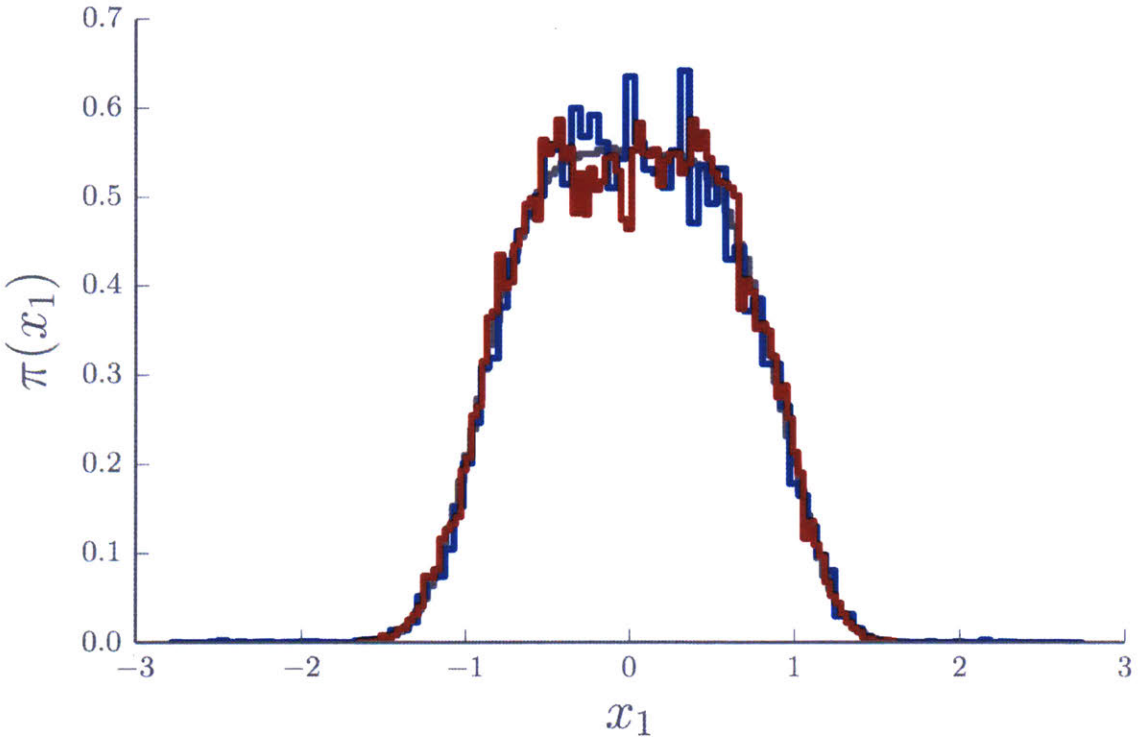

Figure 4-23: 10^5 samples of the marginal distribution π_{X_1} of the distribution π_{X_1,X_2} with quartic density defined in 4.20. We have characterized the distribution with pseudo-marginal MCMC (red line) and LA-MCMC (blue line). In both cases we approximate the target density using a standard normal biasing distribution with $M = 10$ importance samples. In the LA-MCMC case, algorithmic parameters defined by (4.77) and (4.75) with $\tau_0 = 1$, $\gamma_0' = 0.25$, $k = 10$, $\Lambda = 50$, and $\gamma_1 = 1$.

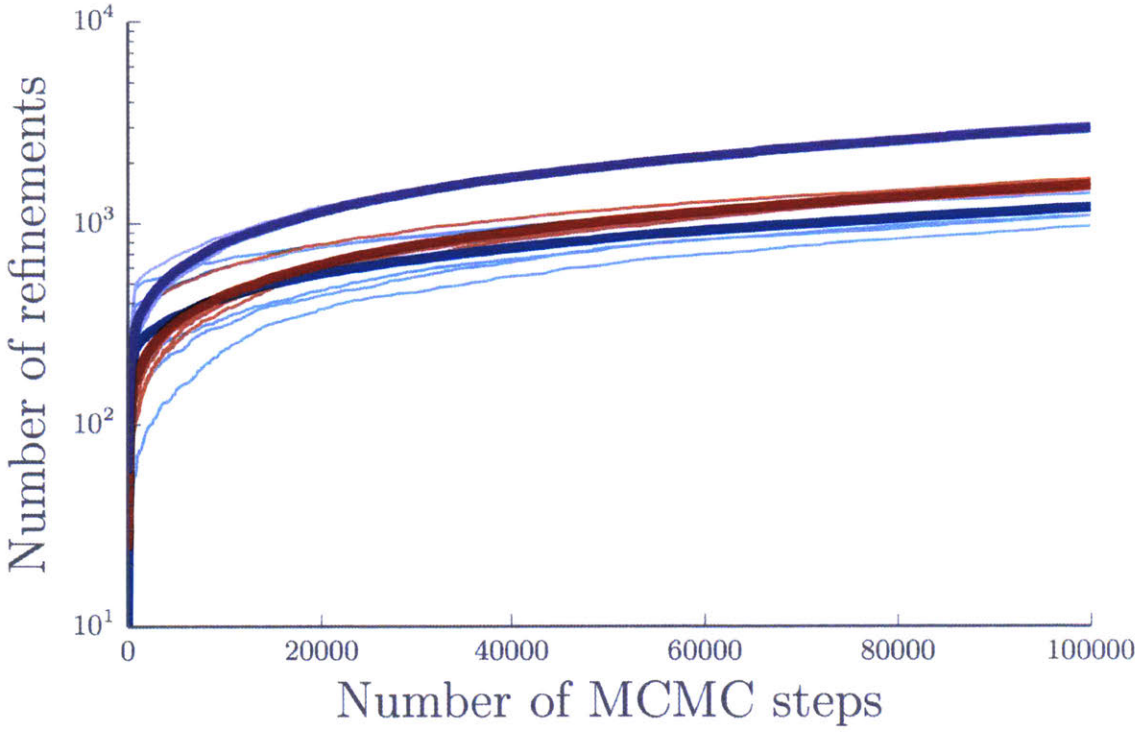

Figure 4-24: The number of refinements as a function of MCMC steps. In this case, each refinement corresponds to computing an importance sampling estimate with $M = 10$ samples, therefore, costing 10 joint density evaluations. We have prescribed algorithmic parameters defined by (4.77) and (4.75) with $\tau_0 = 1$, $\gamma_0' = 0.25$, $k = 10$, and $\Lambda = 50$ and let $\gamma_1 = 0.25$ (blue lines), $\gamma_1 = 0.5$ (red lines), $\gamma_1 = 0.75$ (purple lines), and $\gamma_1 = 1$ (green lines).

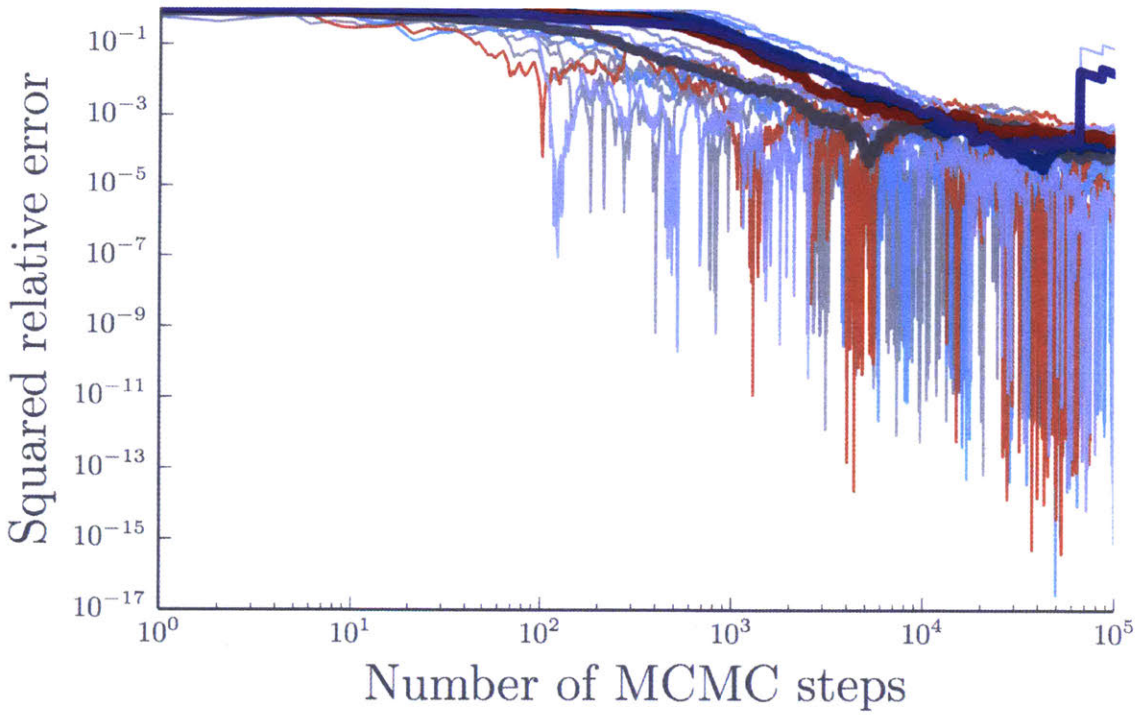

Figure 4-25: The squared relative error (4.61) as a function of MCMC steps. We have prescribed algorithmic parameters defined by (4.77) and (4.75) with $\tau_0 = 1$, $\gamma_0' = 0.25$, $k = 10$, and $\Lambda = 50$ and let $\gamma_1 = 0.25$ (blue lines), $\gamma_1 = 0.5$ (red lines), $\gamma_1 = 0.75$ (purple lines), and $\gamma_1 = 1$ (green lines).

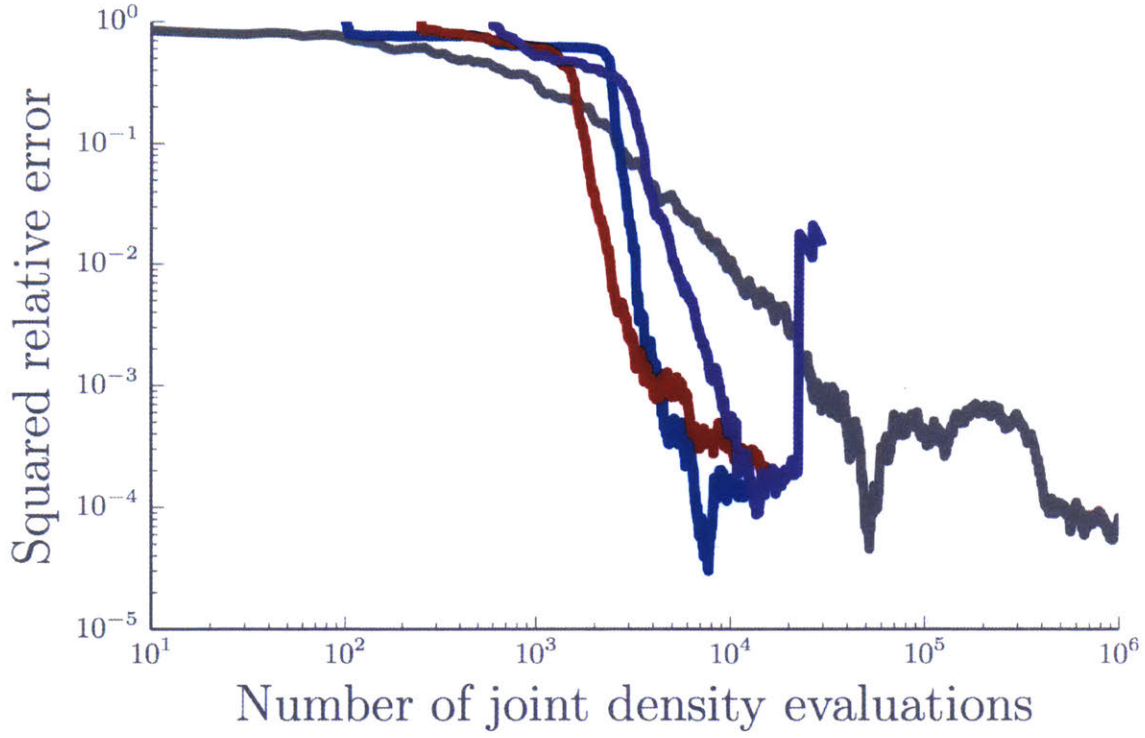

Figure 4-26: The squared relative error (4.61) as a function of target density evaluations. In this case, each refinement corresponds to computing an importance sampling estimate with $M = 10$ samples, therefore, costing 10 joint density evaluations. We have prescribed algorithmic parameters defined by (4.77) and (4.75) with $\tau_0 = 1$, $\gamma_0' = 0.25$, $k = 10$, and $\Lambda = 50$ and let $\gamma_1 = 0.25$ (blue lines), $\gamma_1 = 0.5$ (red lines), $\gamma_1 = 0.75$ (purple lines), and $\gamma_1 = 1$ (green lines).

There is a trade-off between increasing the number of nearest neighbors k and increasing the number of importance samples m used to estimate the target density. Recall that the surrogate variance is bounded by (4.43):

$$\mathbb{V}\mathrm{ar}\left(g_n\right) \leq C_q \frac{q^2}{kM} \Lambda^2, \tag{4.93}$$

where C_q is a constant that depends on the biasing distribution and M is the number of importance samples. From this, we see that if $M > k$ we obtain more surrogate variance reduction by increasing the number of nearest neighbors k and vice-versa. Although both PM-MCMC and LA-MCMC targeting a marginal distribution will asymptotically characterize the correct distribution as long as the variance is finite, we find that the noise has practical implication on performance. Heuristically, however, we find that we can still use relatively few importance samples.

Chapter 5

Practical use of local approximation MCMC: parallel implementation and example applications

In chapter 4 we developed an algorithm that characterizes marginal distributions with computationally intensive probability distributions. We showed that, in theory, replacing the computationally expensive target density function with a cheaper surrogate model and continually refining the surrogate during sampling significantly reduces the computational expense. Furthermore, a bias-variance trade-off implies an ideal refinement rate. However, we did not show whether this algorithm works in practice.

First, we extend the result from chapter 4 to a parallel implementation. Naively, running multiple MCMC chains in parallel generated more samples and, therefore, reduces the variance of the estimate. Rather than having each chain build its own surrogate model, we share model evaluations across chains and build *one* global surrogate. We show that this reduces runtime and increases the efficiency of each individual chain.

Finally, We now focus on practical uses of LA-MCMC. We apply local approximation MCMC (LA-MCMC) to two practical examples: (i) inferring the transmissivity of an aquifer given observations of tracer concentration and (ii) inferring static pa-

rameters of a state-space model. Both these examples highlight the usefulness of out algorithm targeting distributions with computationally and potentially noisy probability density functions.

Chapter 4 outlines a novel algorithm for Bayesian computation. In section 5.2, we extend this algorithm to parallel implementations and, in section 5.3 we discuss how to use surrogate models with more general proposals. These two advancements are published in [Conrad et al., 2018]. Finally, in section 5.4, we apply local approximation MCMC to large-scale, practical examples.

5.1 Motivation

Markov chain Monte Carlo (MCMC) is a powerful tool for performing Bayesian inference, but can be computationally prohibitive in many settings, especially when posterior density evaluations involve a computationally expensive step. Glacier models, for instance, often require partial differential equation *forward models*, evaluated using numerical solvers with nontrivial run times. When these solvers must be invoked with each posterior evaluation, direct sampling with MCMC can become intractable.

As we have discussed in Section 4.3, constructing a surrogate of the forward model, likelihood function, or posterior density can reduce this computational burden. We then sample from (or otherwise characterize) the posterior distribution induced by this approximation [Sacks et al., 1989, Kennedy and O'Hagan, 2001, Rasmussen, 2003, Santner et al., 2003, Marzouk et al., 2007, Marzouk and Xiu, 2009, Bliznyuk et al., 2012, Joseph, 2012, Li and Marzouk, 2014, Chen and Schwab, 2015, Cui et al., 2016]. Although such approaches are effective at reducing computational cost, they are difficult to use in practice; they separate the construction of the surrogate from the subsequent inference procedure. [Conrad et al., 2016] demonstrate that surrogate construction and posterior exploration can be *joined*, yielding a framework for incrementally and infinitely refining a surrogate during sampling. This framework allows the surrogate to be tailored to the problem—e.g., made most accurate in regions of high posterior probability—while guaranteeing that the associated Markov chain

asymptotically samples from the *exact* posterior distribution of interest.

In Chapter 4, we extended the ideas of [Conrad et al., 2016] to include a bias-variance trade-off that triggers refinement, and we argued that the same algorithm can be used to target marginals of high-dimensional distributions. Biases introduced into posterior expectation due to the surrogate model biases cannot easily be quantified [Cotter et al., 2010]. However, we have constructed an algorithm that allocates computational effort to ensure that the surrogate bias decays at the same *rate* as the error from posterior sampling. We also demonstrated that this procedure works when only noisy evaluations of the target density (e.g., Monte Carlo estimates) are available. Alternatives such as delayed-acceptance MCMC [Christen and Fox, 2005, Cui et al., 2011] yield asymptotically exact sampling, but still require at least one expensive density evaluation per accepted sample.

Here, we again extend work by [Conrad et al., 2016] and Chapter 4. First, we show that our approximation scheme enables a novel type of MCMC parallelism: concurrent chains can collaboratively develop a *shared* approximation. Effectively exploiting parallel computation in MCMC is often challenging because the core algorithm is inherently sequential, but our strategy directly deploys parallel resources to address the key performance bottleneck: the cost of repeatedly running the forward model.

Second, while our previous work showed how to build a convergent approximation of the target probability density, it did not support the idea of using this approximation to construct a proposal distribution. MCMC performance is highly dependent on the choice of proposal, but sophisticated proposals, such as the Metropolis-adjusted Langevin algorithm (MALA) and its manifold variants [Girolami and Calderhead, 2011], can be expensive to apply because they require gradients (and possibly higher derivatives) of the forward model. This derivative information is often expensive or impossible to compute directly, but is trivial to extract from an approximation. Intuitively, it should then be possible to use our approximation framework to greatly reduce the costs of such proposals. In [Conrad et al., 2018] we do exactly that, extending our previous theoretical results to show that the Monte Carlo estimates obtained by our algorithm converge to the correct value, as long as the convergence

of our approximation to the target distribution yields convergence of the associated approximate Markov transition kernel in a suitably strong norm. As an example, we show how to use simplified manifold MALA within our local approximation scheme, and prove that the resulting stochastic process is convergent in a representative case.

Finally, we construct inference problems that are representative of interesting scientific queries, that involve computationally expensive forward models (such that naïve use of the model in sampling would take days or months), and that have non-trivial posterior structure which must be characterized using MCMC. Our numerical experiments evaluate MCMC efficiency, accuracy, and wallclock time, and benchmark the parallel performance of our algorithms.

We refer to chapter 4 for detailed discussion of local approximation MCMC. Here, we will discussion the parallel extension to the algorithm and demonstrates its effectiveness in practical applications. We also refer to [Conrad et al., 2018] for rigorous proof of the convergence of LA-MCMC using MALA; we will focus on more practical aspects.

5.2 Sharing local approximations for parallel MCMC

The naïve approach to parallelizing MCMC is to run several independent chains in parallel. Although running parallel chains facilitates useful convergence diagnostics [Cowles and Carlin, 1996, Brooks and Roberts, 1998], practical scaling in highly parallel environments is limited because of the serial nature of MCMC and the replication of transient behavior across multiple chains [Rosenthal, 2000].

More sophisticated strategies for parallel MCMC exchange information between the chains. For example, [Craiu et al., 2009] proposes moves to states discovered by other chains. Population MCMC algorithms explore a family of tempered distributions with parallel chains, so that swapping states between the chains can provide long-range moves [Cappe et al., 2004]. These techniques attempt to improve the mixing time of the Markov chain, and when successful, may provide superior performance to the naïve parallelization [Green et al., 2015]. Other constructions, e.g., [Calder-

head, 2014], propose multiple points in parallel and try to make use of all these points in determining subsequent steps of a single chain.

Any of these parallel approaches requires each chain—and, therefore, each processor—to repeatedly evaluate the forward model, which can dominate the overall cost of the algorithm. If multiple copies of algorithm 4 are run in parallel, a natural idea is to allow them to share a common set of evaluations $\mathcal{S}_n$. That is, whenever one chains performs refinement, the result is shared asynchronously with all the chains; each chain receives additional model evaluations "for free." In settings where the limiting computational cost is constructing $\mathcal{S}_t$, parallelizing this process directly impacts the real-world performance of the sampler during the stationary and even the transient phases of the chains. With regard to the latter point, we note that parallelizing $\mathcal{S}_n$ can reduce the number of model evaluations that are triggered by each individual chain during its initial transient phase.

5.2.1 Convergence of the parallel algorithm

Recall that our local approximation MCMC algorithm is detailed in chapter 4; see, in particular, algorithm 4. Below we will show that the sufficient conditions for convergence of a single-chain version of algorithm 4, as described in [Conrad et al., 2016] and chapter 4, are also sufficient conditions for the convergence of the parallel version. The arguments given in [Conrad et al., 2016] are straightforward to extend because we have chosen conditions where enlarging the evaluated set $\mathcal{S}_n$ is always helpful (i.e., ensured Λ-poisedness); thus, the additional refinements contributed by parallel chains cannot hinder convergence. Rather than repeating the entire discussion of convergence from that paper, here we merely extend the simplest and weakest convergence result—for a single chain on a compact state space [Conrad et al., 2016, Theorem 3.4]—to the case of parallel chains. We refer the reader to [Conrad et al., 2016, Theorem 3.3] for related conditions and a treatment of non-compact state spaces that can similarly be extended to the parallel case.

We require some notation before stating the result. Let $\mathcal{L}(X)$ denote the distribution of a random variable X. For fixed $\epsilon > 0$, we say that $\mathcal{S} \subset \mathcal{X}$ is an ϵ-cover of

$\mathcal{X}$ if $\sup_{x\in\mathcal{X}} \min_{s\in\mathcal{S}} \|x-s\|_2 < \epsilon$. We assume that we are using algorithm 4 with *random refinement* and all references to $\mathcal{S}_n$ are replaced by a reference to a single set $\mathcal{S}$, then the sequence $\{x_t\}_{t\geq 0}$ constructed by running the modified algorithm is a Markov chain. We use the $\mathcal{S}$ subscript to denote all approximate objects associated with this Markov chain (e.g., $\pi_{\mathcal{S}}(x) \approx \pi(x)$, where $\pi(x)$ is the true target density). Similarly, ∞ subscript (e.g., $\pi_\infty(x) = \pi(x)$) denotes the values of these objects for the Markov chain with the same proposal kernel as in algorithm 4 and with the *correct* posterior distribution as its target distribution. Our result makes the following assumptions:

Assumption 1 (Sufficient conditions for convergence)

1. *The state space $\mathcal{X}$ is compact.*

2. *The proposal $q(x,\cdot\,|\,\mathbf{f}) = q(x,\cdot)$ does not depend on $\mathbf{f}$, and both the proposal distribution $q(x,\cdot)$ and target distribution $p(\cdot\,|\,\mathbf{d})$ have C^∞ densities that are bounded away from zero uniformly in x.*

3. *The sequence of parameters $\{\beta_t\}_{t\in\mathbb{N}}$ used in Algorithm ?? are of the form $\beta_t \equiv \beta > 0$; any sequence $\{\gamma_t\}_{t\geq 0}$ is allowed.*

4. *The approximation of $\log p(x|\mathbf{d})$ is made via quadratic interpolation on the $N = (d+1)(d+2)/2$ nearest points.*

5. *Rather than choosing new points via the method described in section 4.3.6, we assume that we refine at the current state of the chain.*

6. *We assume Λ-poisedness as defined in* (4.39).

The following result extends Theorem 3.4 of [Conrad et al., 2016] to parallel chains.

Thereom 1 (Convergence with parallel chains) *Let $\{X_t^{(i)}\}_{t\geq 0,\, 1\leq i\leq n}$ be the n stochastic processes obtained in a parallel run of algorithm 4 with n chains, and assume that the algorithm parameters satisfy assumption 1. Then, for all $1 \leq i \leq n$, $\lim_{t\to\infty} \|\mathcal{L}(X_t^{(i)}) - \pi\|_{TV} = 0$.*

Proof: The proof of Lemma B.3 of [Conrad et al., 2016] holds exactly as stated, with the proof as given. The remainder of the proof of Theorem 3.4 from [Conrad et al., 2016] holds for $\{X_t^{(i)}\}_{t\geq 0}$, for each fixed $1 \leq i \leq n$, with the following modifications:

(i) The chain $\{X_t\}_{t\geq 0}$ should be replaced by $\{X_t^{(i)}\}_{t\geq 0}$ wherever it appears; and

(ii) The auxillary process associated with $\{X_t^{(i)}\}_{t\geq 0}$ is

$$\left\{\left(\mathcal{S}_t, X_t^{(1)}, \ldots, X_t^{(i-1)}, X_t^{(i+1)}, \ldots, X_t^{(n)}\right)\right\}_{t\geq 0},$$

rather than $\{\mathcal{S}_t\}_{t\geq 0}$. ■

We emphasize that this proof of the convergence of $\{X_t^{(i)}\}_{t\geq 0}$ is completely indifferent to the points that are added to $\mathcal{S}_t$ by the other chains $\{X_t^{(j)}\}_{t\geq 0}$, $j \neq i$.

5.3 Local approximations and approximating the proposal

Many variations of Markov chain Monte Carlo (MCMC) use the target distribution to propose new states, in addition to computing the acceptance probability. One such algorithm is Metropolis adjusted Langevin algorithm (MALA—see section 4.2.2), which uses gradient information of the target density to propose a the next state. Replacing the, typically expensive, gradient computation with gradients of the surrogate model provide a cheaper way of computing information. Thus far, our work as assumed that the proposal mechanism is *independent* of the surrogate model. We show in [Conrad et al., 2018] that using the surrogate within the proposal asymptotically characterized the desired target distribution; see [Conrad et al., 2018] for details. We will employ local approximation with a MALA proposal and parallel chains to demonstrate some practical uses of LA-MCMC.

5.3.1 Example of local approximation MALA

Recall the target distribution with the following log-quartic density:

$$\log \pi(x_1, x_2) = -x_1^4 - \frac{(2x_2 - x_1^2)^2}{2}, \tag{5.1}$$

also presented in (4.20) and throughout chapter 4 and illustrated in figure 4-3. We simulate from this target distribution in four ways: using (i) adaptive Metropolis (AM) [Haario et al., 2001] and (ii) mMALA, each paired with either (a) evaluations of the exact target density or (b) our local approximation scheme. In other words, the combinations (a+i) and (a+ii) are standard MCMC algorithms with two different proposal schemes, and the combinations (b+i) and (b+ii) pair local approximation MCMC with the same proposal schemes. We call these simulation approaches 'exact+AM,' 'exact+mMALA,' 'LA+AM,' and 'LA+mMALA,' respectively.

For each of the simulation approaches defined above, we run 20 independent chains, each of length 6×10^5 steps. (No parallelism is employed in this example.) We then evaluate an *expected squared relative error* $\bar{\varepsilon}^2$, as a function of the number of target density evaluations, for each approach. The quantity $\bar{\varepsilon}^2$ is defined as follows. Before computing expectations with respect to the target density, we discard the first 10^4 samples of each chain as burn-in. Then we obtain a reference estimate C_0 for the target covariance matrix by pooling post-burn-in samples from all 40 chains that employ exact target density evaluations. Next, for each independent chain (indexed by i) associated with a given simulation approach, we compute a running t-sample estimate of the target covariance $\hat{C}_t^{(i)}$ and define a relative squared error as

$$\varepsilon_t^{2,(i)} = \frac{\|\hat{C}_t^{(i)} - C_0\|_F^2}{\|C_0\|_F^2}, \tag{5.2}$$

where $\|\cdot\|_F$ denotes the Frobenius norm. Then we average over the 20 independent chains to obtain $\bar{\varepsilon}_t^2 = \frac{1}{20}\sum_{i=1}^{20} \varepsilon_t^{2,(i)}$. Figure 5-1 plots $\bar{\varepsilon}_t^2$ versus the number of target *density evaluations*, for each simulation approach. For the exact+mMALA chains, which require direct evaluation of the gradients of $\log \pi$, we count each gradient evaluation as an additional density evaluation. For large-scale models, gradient evaluations (e.g., via an adjoint solve) might be more expensive than density evaluations, so this accounting is a conservative estimate of computational cost.

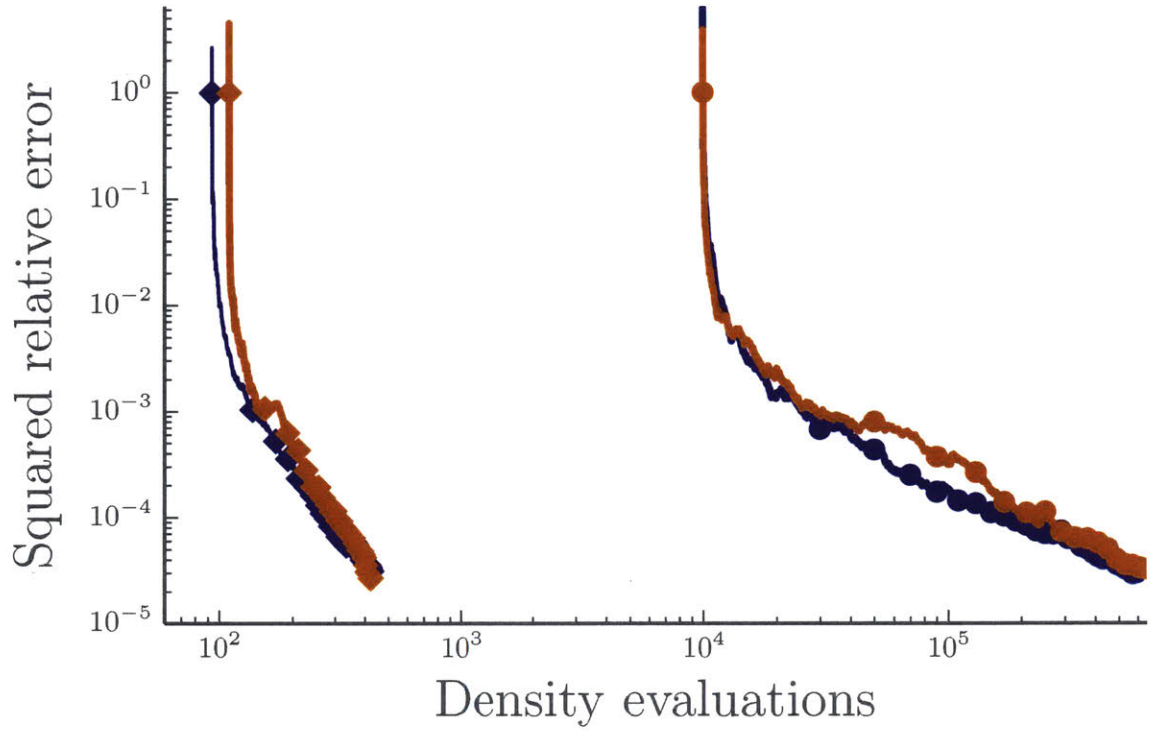

Figure 5-1: Expected squared relative error $\bar{\varepsilon}_t^2$ as a function of the number of target density evaluations. Purple lines correspond to AM chains while gold lines correspond to mMALA. The circles mark chains that employ exact evaluations of the target density, while diamonds mark chains using local approximation. In the exact case, mMALA requires evaluations of the target density and its gradient. We assume gradient evaluations are comparable in cost to density evaluations and, therefore, count them as density evaluations. Errors are obtained by averaging over 20 independent chains from each simulation approach, each of length 6×10^5 steps.

Another useful measure of sample quality is the effective sample size (ESS) of each chain, which we compute from each chain's integrated autocorrelation time [Wolff, 2004]. ESS is a measure of how many "effectively independent" samples have been generated from the target distribution. In Figure 5-2, we plot the ESS for each independently realized chain, using each of the four simulation approaches. In general, the mMALA chains have larger ESS than the AM chains, reflecting their improved mixing for this target distribution. Also, the local approximation chains achieve nearly the same ESS as their exact counterparts, but with nearly three orders of magnitude fewer density evaluations. ESS of course varies from realization to realization; the dark symbols in the middle of the scatter plots illustrate the *average* ESS and cost of

each set of 20 chains. In general, we do not expect that introducing an approximation will improve mixing, and in this example ESS with exact evaluations (exact+AM or exact+mMALA) provides an upper bound on sampling performance. Indeed, Figure 5-2 shows that the ESS is very slightly lower using local approximations; this is apparent in the MALA cases. Nonetheless, the local approximation chains achieve nearly the same ESS as their exact counterparts, but with nearly three orders of magnitude fewer density evaluations. Moreover, the improved mixing of mMALA in the exact case is preserved when using local approximations.

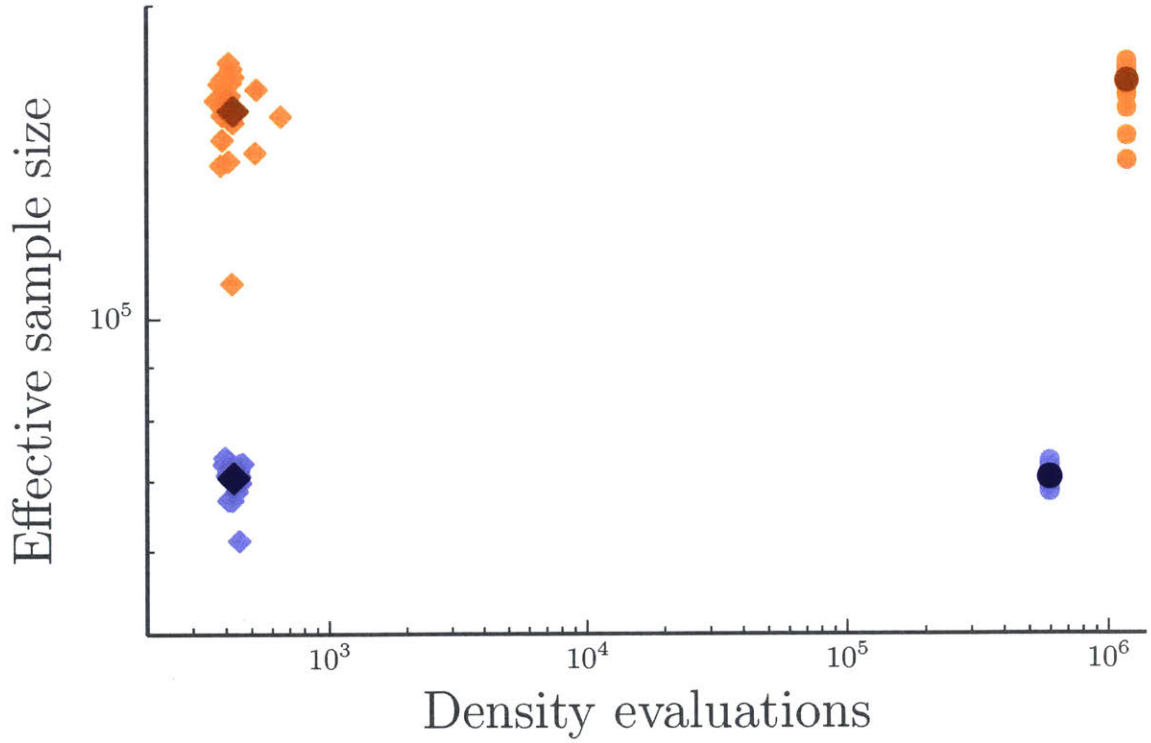

Figure 5-2: Effective sample size for independent MCMC chains, each of length 6×10^5. As in Figure 5-1, purple symbols correspond to AM chains and gold symbols correspond to mMALA chains. Circles indicate chains using exact target density evaluations, while diamonds indicate the use of local approximations. The darker dot in each cluster is its expected value. In the mMALA case with exact evaluations, we count target gradient evaluations as density evaluations.

5.4 Numerical experiments

We present two larger numerical examples to explore the algorithmic ideas developed in the preceding sections and chapter 4. We use more computationally intensive inference problems, with forward models drawn from realistic applications. The first of these is a state space model—specifically, a hidden Markov chain where an observable quantity is related to a unobservable state via known statistical structure. Our goal is to infer static hyperparameters given observations, marginalizing the unobservable state. The second application, a groundwater tracer transport problem, will focus on our parallel MCMC explorations. Though posterior evaluations are quite expensive in this problem, we can still compare results with standard MCMC chains that employ no approximation, and thus verify the accuracy of posterior expectations. Finally, we consider an elliptic PDE. The problem set-up is similar to groundwater problem, however, here we approximate the posterior distribution by characterizing marginals of the posterior in data-informed direction and sampling the remaining parameters from the prior.

5.4.1 Stochastic volatility model

We infer static hyperparameters of a state-space model. Let Z_t and Y_t denote the unobservable state and observations at timestep t, respectively. We model the evolution of Z_t and Y_t with a stochastic volatility model such that

$$Z_1 \sim N(\nu, \tfrac{1}{1-\phi^2}) \quad Z_t|Z_{t-1} \sim N(\mu + \phi(Z_{t-1} - \mu), \sigma^2), \tag{5.3a}$$

with observations

$$Y_t \sim N(0, \exp(Z_t)). \tag{5.3b}$$

The quantity of interest is $x = [x_1, x_2, x_3]$ such that $[\mu, \phi, \sigma] = g(x)$. Specifically,

$$\mu = x_1, \quad \phi = 2\frac{\exp(x_2+3)}{1+\exp(x_2+3)} - 1, \quad \text{and} \quad \sigma^2 = F^{-1}_{\alpha,\beta} \circ \Phi(x_3), \tag{5.4}$$

where

$$f_{\alpha,\beta}(x) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{-\alpha-1} \exp\left(-\frac{\beta}{x}\right) \quad \text{and} \quad F_{\alpha,\beta}(x) = \frac{\Gamma\left(\alpha, \frac{\beta}{x}\right)}{\Gamma(\alpha)} \tag{5.5}$$

are the density and cumulative density functions for an inverse gamma distribution, respectively, and $\Phi(x)$ is the cumulative density function for a standard normal distribution. The composition $F^{-1}_{\alpha,\beta} \circ \Phi(\cdot)$ transforms a sample from a standard normal distribution $x \sim N(\cdot; 0, 1)$ to a sample from the inverse gamma distribution $y \sim f_{\alpha,\beta}$. We set $\alpha = 1$ and $\beta = 0.1$. Figure 5-3 shows how the hyperparameters, states, and observations are related.

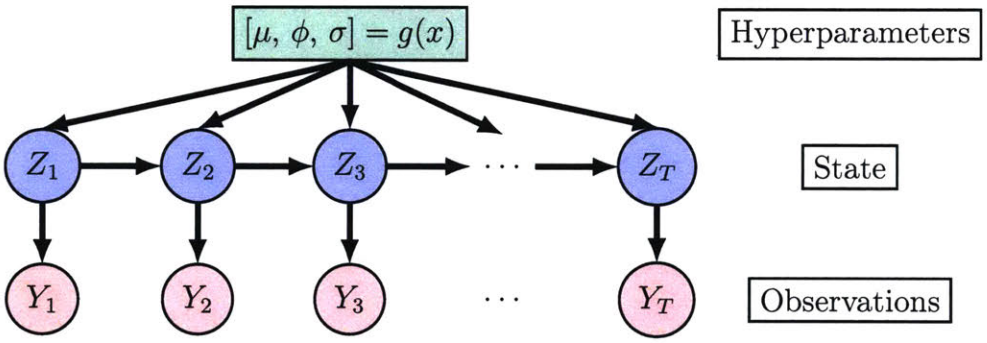

Figure 5-3: The stochastic volatility model described by (5.3).

We infer the static hyperparameters $[x_1, x_2, x_3]$ given observations $Z_{1:T}$. Let

$$\pi(x) = N(x; 0, I) \tag{5.6}$$

be the prior distributions on the hyperparameters. Given hyperparameters, we compute $Y(x)$ by evaluating (5.3). We generate data by prescribing 'true' parameters ν^*, ϕ^*, and σ^* and computing an instance $y = Y(\nu^*, \phi^*, \sigma^*)$; figures 5-4 and 5-5 show an example of the states and observations from one simulation. The joint posterior density is

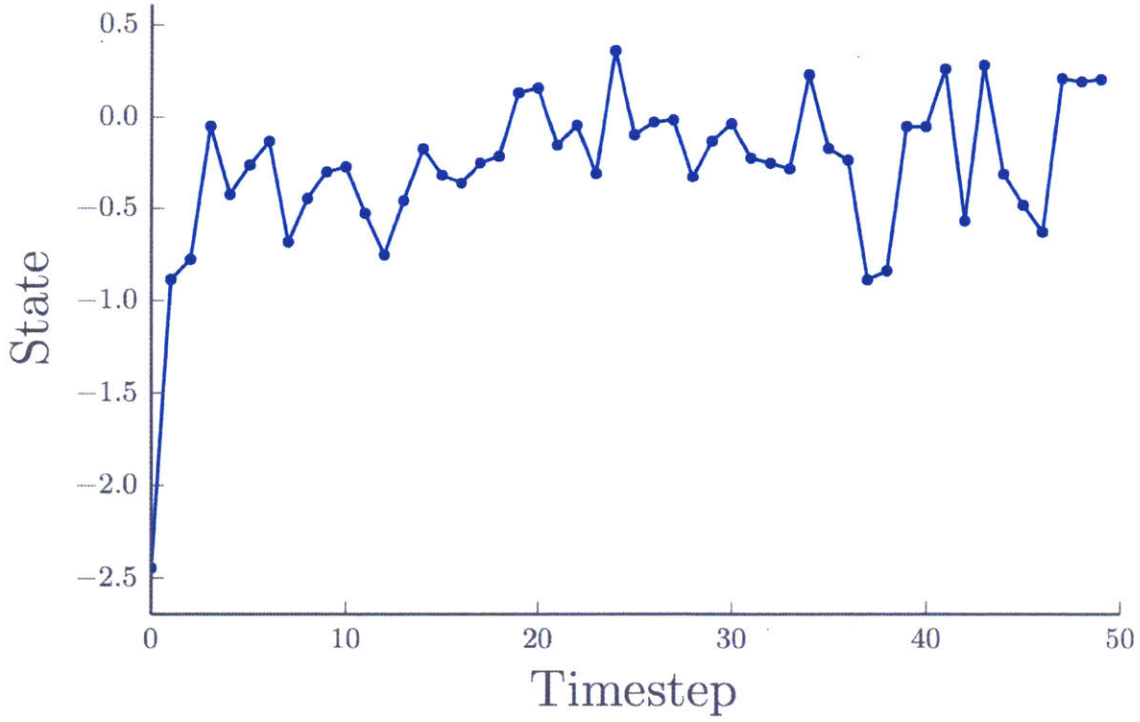

(a)

Figure 5-4: The state of the stochastic volatility model (5.3) conditioned on a single realization of the hyperparameters $x \sim N(0, 1)$.

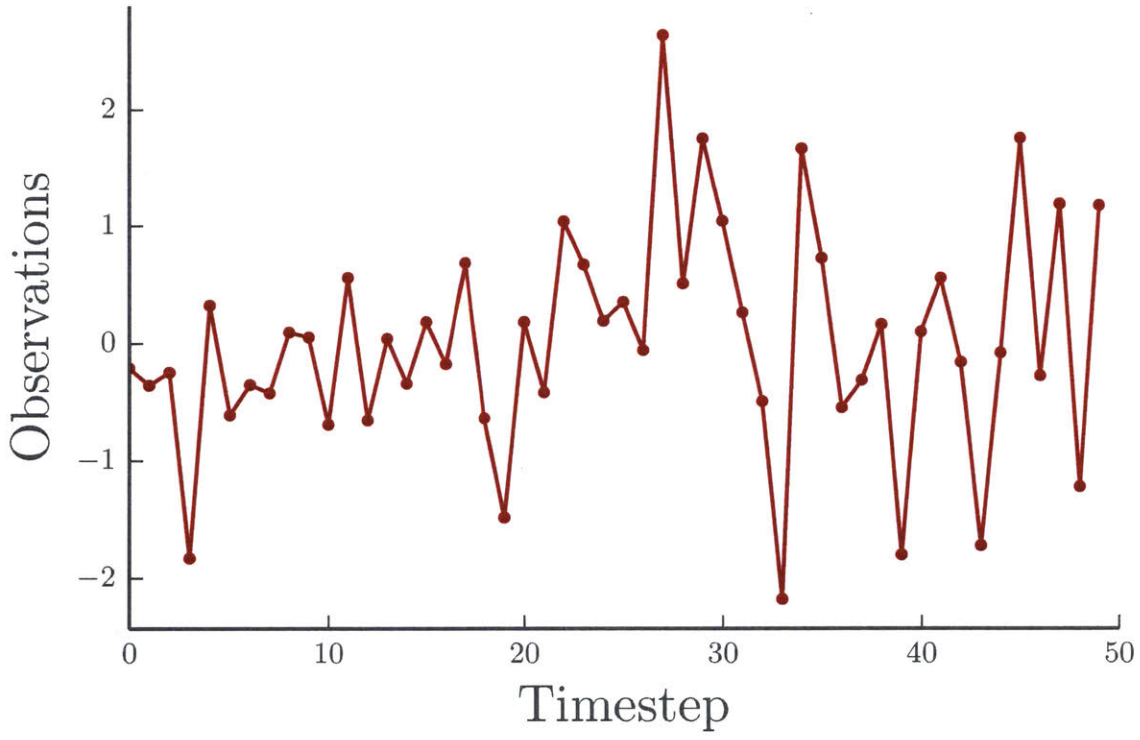

(a)

Figure 5-5: The observations of the stochastic volatility model (5.3) conditioned on a single realization of the hyperparameters $x \sim N(0, 1)$.

$$\pi(x, z|y) \propto \pi(x)\pi(z|x)\pi(y|x, z), \tag{5.7a}$$

where

$$\pi(z|x) = \pi(z_1|x) \prod_{t=2}^{T} \pi(z_t|z_{t-1}, x) \tag{5.7b}$$

and

$$\pi(y|x, z) = \pi(y|z) = \prod_{t=1}^{T} \pi(y_t|z_t). \tag{5.7c}$$

Our goal is to characterize the posterior marginal distribution

$$\pi(x|y) = \int_{\mathcal{Z}} \pi(x, z|y)\, dz \propto \underbrace{\pi(x)}_{\text{prior}} \underbrace{\int_{\mathcal{Z}} \pi(y|z)\pi(z|x)\, dz}_{\text{marginal likelihood}}. \tag{5.8}$$

In general, we cannot directly target (5.7) with MCMC because T is large. However, the quantity of interest x is only three-dimensional. Therefore, we construct unbiased estimates of the marginal likelihood

$$\pi(y|x) = \int_{\mathcal{Z}} \pi(y|z)\pi(z|x)\, dz \approx \hat{\pi}(y|x) = \sum_{i=1}^{M} \frac{\pi(y|z^{(i)})\pi(z^{(i)}|x)}{q(z^{(i)}|x, y)}, \tag{5.9}$$

where $q_{Z|X,Y}$ is a biasing distribution. In general, finding an appropriate biasing distribution is non-trivial. For this example, we choose the conditional Laplace approximation

$$q(z|x, y) = N(z; z^* + \hat{C}_{zx}\hat{C}_{xx}^{-1}(x^* - x), \hat{C}_{zz} - \hat{C}_{zx}\hat{C}_{xx}^{-1}\hat{C}_{xz}), \tag{5.10}$$

where $z^*, x^* = \arg\max_{z\in\mathcal{Z}, x\in\mathcal{X}} \log(\pi(x, z|y))$ and $\hat{C}_{ij} = \frac{\partial}{\partial z_i \partial z_j} \log(\pi(x, z|y))\Big|_{z=z^*, x=x^*}$ and

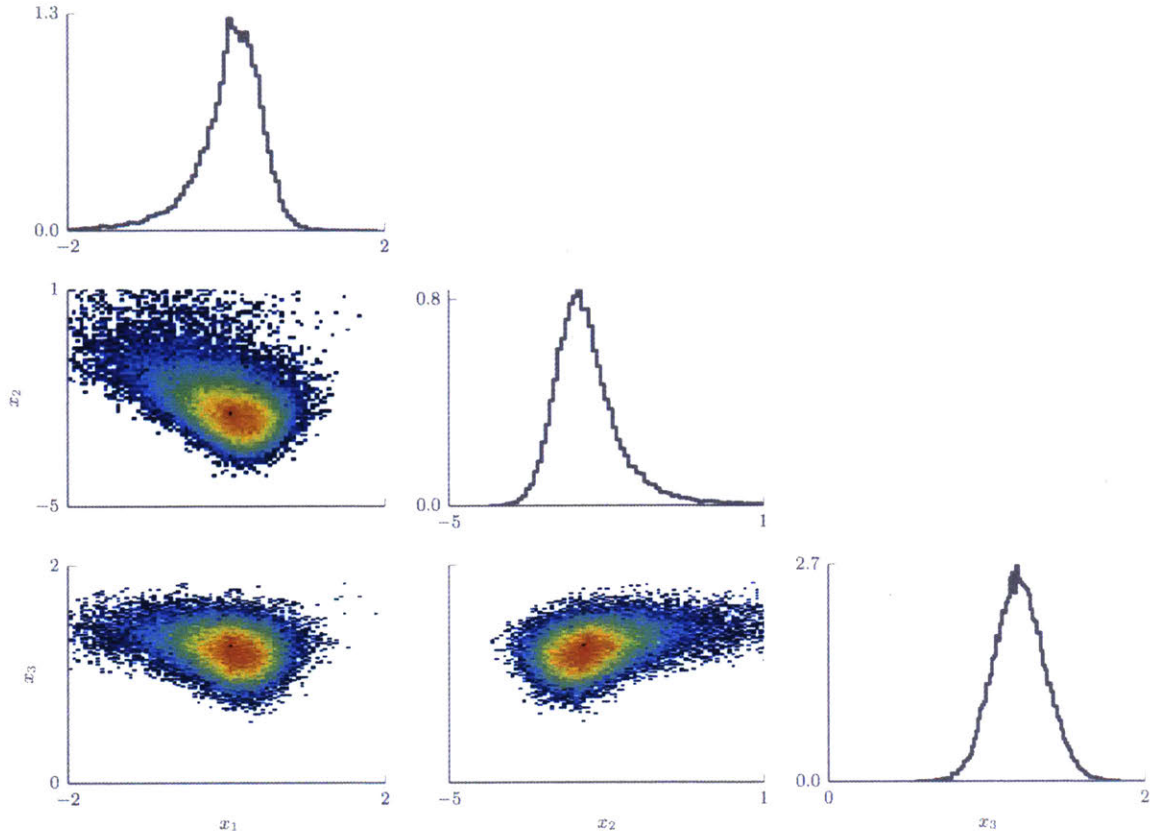

Figure 5-6: The posterior marginal distribution $\pi_{X|Y}$ characterized using 10^5 samples generated using pseudo-marginal MCMC with $M = 15$ importance samples to approximate the target marginal density.

$$\hat{C} = \begin{bmatrix} \hat{C}_{xx} & \hat{C}_{xz} \\ \hat{C}_{zx} & \hat{C}_{zz} \end{bmatrix}. \tag{5.11}$$

We employ both pseudo marginal MCMC and local approximation MCMC. The results are shown in Figures 5-6 and ??.

5.4.2 Tracer transport problem

Predicting the evolution of groundwater contaminant concentrations over time is vital to many monitoring and remediation efforts [Matott, 2012]. A contaminant is typically modeled as a non-reactive tracer that diffuses and is advected by groundwater flow. Here we construct an inverse problem that simulates a monitoring configuration:

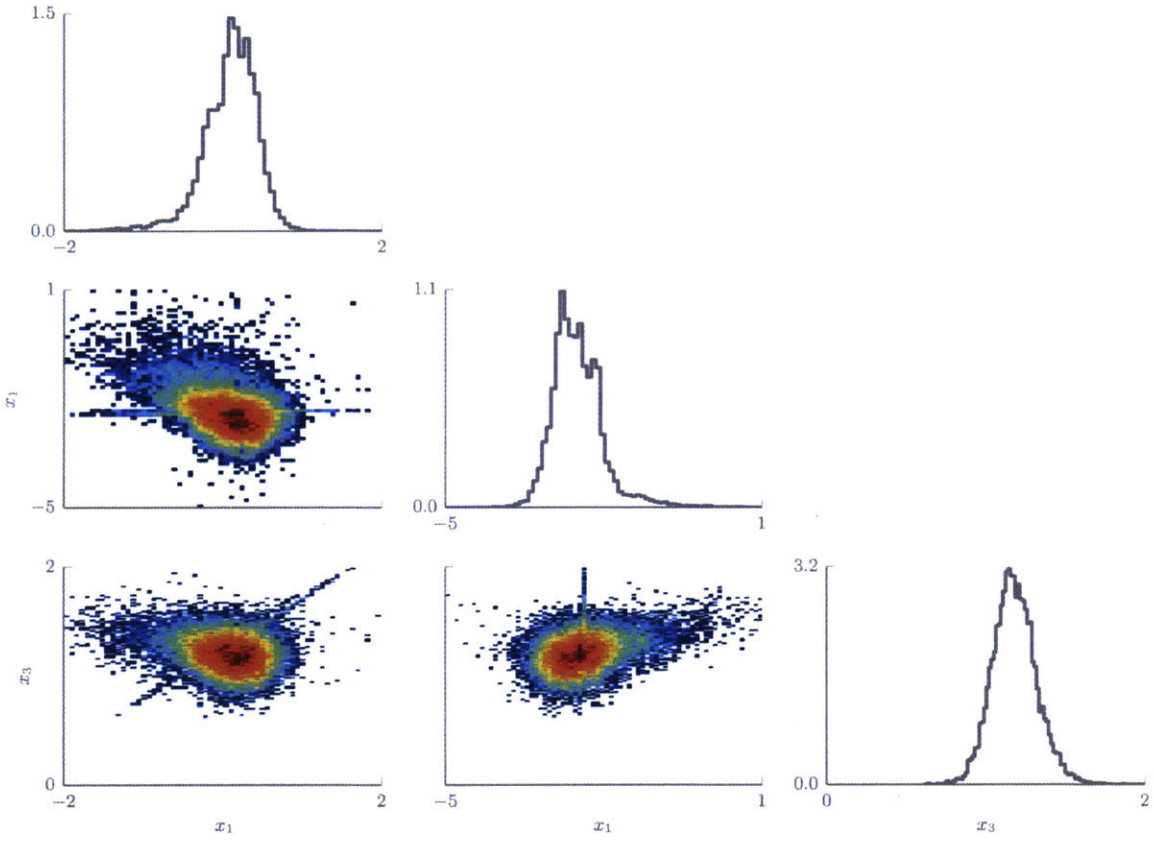

Figure 5-7: The posterior marginal distribution $\pi_{X|Y}$ characterized using 10^5 samples generated using LA-MCMC with $M = 15$ importance samples to approximate the target marginal density. We only required 2.8×10^3 importance sampling estimates to build the surrogate model, which is significantly fewer than the 2×10^5 required by pseudo-marginal MCMC.

Parameter	x_0^j	x_1^j	y_0^j	y_1^j	True value	Prior lower	Prior upper
θ_1	0.6	1	0.15	0.3	-0.75	-1	0
θ_2	0	0.4	0.7	0.1	-0.25	-1	1
θ_3	0.5	1	0	0.5	-0.5	-1	0
θ_4	0.4	1	0.6	1	1	0	2
θ_5	0.2	0.25	0	0.3	-0.25	-1	0
θ_6	0	1	0	1	3	2	5

Table 5.1: True values of the parameters for the tracer problem. The log-conductivity at location (x, y) is $\theta_{\hat{j}(x,y)}$, where $\hat{j}(x, y)$ is the smallest integer j such that $x_0^j \leq x \leq x_1^j$ and $y_0^j \leq y \leq y_1^j$; each parameter value θ_j corresponds to $\log \kappa_j$ in figure 5-8.

the tracer concentration is observed at a small number of wells over a short period of time, and the subsurface conductivity field must be inferred given these data.

The conductivity field is assumed to be piecewise constant in six irregularly-shaped areas, reflecting different subsurface features (e.g., sand, clay, gravel) each with constant but unknown conductivities. We consider a problem domain with two horizontal coordinates $x, y \in [0, 1]^2$. The true log-conductivity is depicted in figure 5-8. The conductivity is parameterized as

$$\kappa(x, y) = \exp \theta_{\hat{j}(x,y)}, \tag{5.12a}$$

where $\hat{j}(x, y) \in \{1, 2, \ldots, 6\}$ is the smallest integer j such that $x_0^j \leq x \leq x_1^j$ and $y_0^j \leq y \leq y_1^j$, where the bounds $(x_0^j, x_1^j, y_0^j, y_1^j)$ are given in table 5.1. The parameters[1] θ_i are endowed with uniform priors; the upper and lower bounds for each prior are also given in table 5.1.

Modeling tracer evolution requires first computing the hydraulic head, which determines the groundwater velocity. Under the Dupuit approximation [Dupuit, 1863], the hydraulic head h obeys the elliptic equation,

$$\nabla \cdot (\kappa h \nabla h) = -f_h, \tag{5.12b}$$

where $\kappa(x, y)$ is the conductivity field and $f_h(x, y)$ is the hydraulic head forcing. In

[1]We have switched notation such that the inferred quantity is denoted by θ, rather than x. This is because, in this example x is a spatial coordinate.

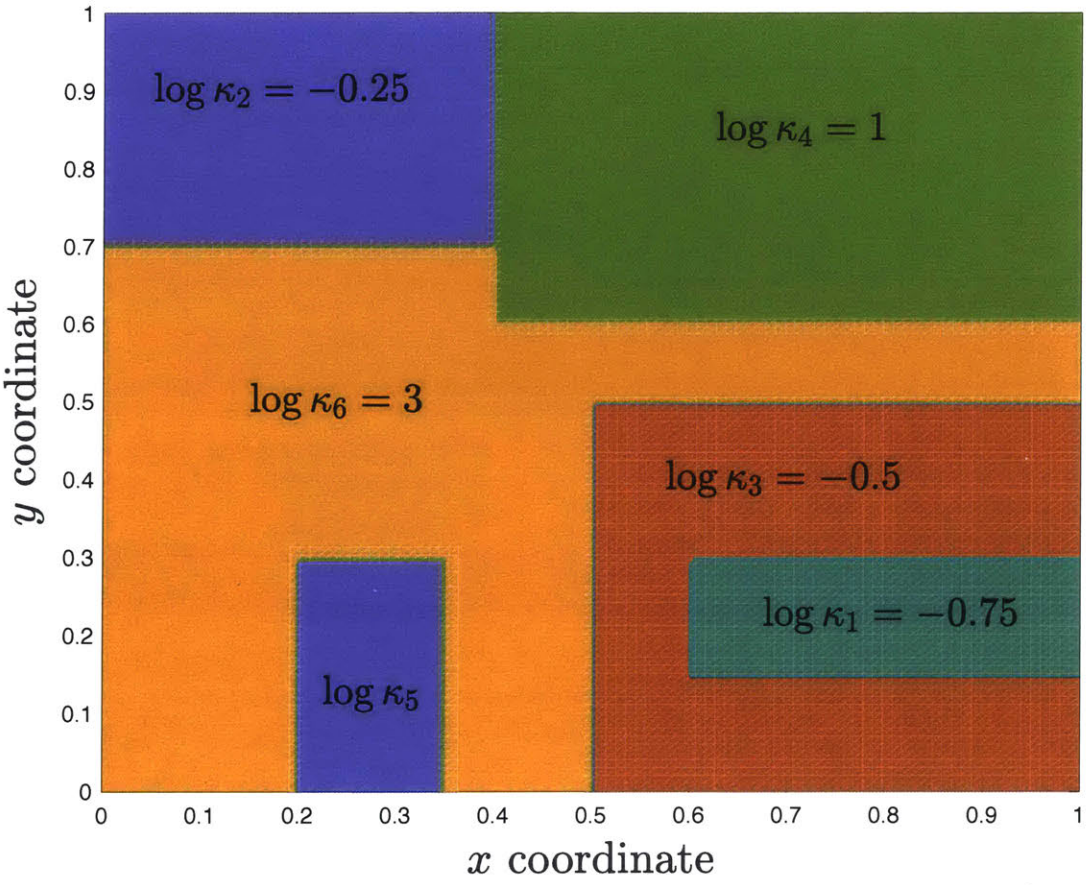

Figure 5-8: The "true" log-conductivity field.

our problem setup, the forcing is created by pumping at four well locations, $(a_i, b_i) \in \{(0.15, 0.15), (0.85, 0.15), (0.85, 0.85), (0.15, 0.85)\}$, such that

$$f_h(x, y) = \sum_{i=1}^{4} p_i \exp\left(\frac{(a_i - x)^2 + (b_i - y)^2}{0.02}\right), \tag{5.12c}$$

where $p_i \in (10, 50, 150, 50)$. The model (5.12b) assumes homogeneous Dirichlet boundary conditions at $y = 0$ and $y = 1$ and homogeneous Neumann conditions at $x = 0$ and $x = 1$. The Darcy velocity is determined by the hydraulic head gradient

$$\begin{bmatrix} u \\ v \end{bmatrix} = -h\kappa\nabla h. \tag{5.12d}$$

The time-dependent tracer concentration $c(x, y, t)$ then evolves given a flow-dependent dispersion tensor, via

$$\frac{\partial c}{\partial t} + \nabla \cdot \left(\left(d_m \mathbf{I} + d_l \begin{bmatrix} u^2 & uv \\ uv & v^2 \end{bmatrix}\right)\nabla c\right) - \begin{bmatrix} u \\ v \end{bmatrix} \cdot \nabla c = -f_t, \tag{5.12e}$$

where $d_m = 2.5 \times 10^{-3}$ and $d_l = 2.5 \times 10^{-3}$ are dispersion coefficients and $f_t(x, y)$ is the tracer forcing. The tracer is forced by injection at each well location. The source term is similar to the one forcing the hydraulic head

$$f_t(x, y) = \sum_{i=1}^{4} r_i \exp\left(\frac{(a_i - x)^2 + (b_i - y)^2}{0.005}\right). \tag{5.12f}$$

where $r_i \in (10, 5, 10, 5)$. The tracer has initial condition $c(x, y, 0) = 0$, and homogeneous Neumann conditions are enforced at all spatial boundaries. Since the hydraulic head forcing, tracer forcing, and dispersion coefficients are known, the forward model simply maps the conductivity to a time-evolving concentration field. Tracer observations are taken at 25 well locations: (x_i, y_j) such that $x_i = 0.1 + \frac{i-1}{5}$ and $y_j = 0.1 + \frac{j-1}{5}$ for $i, j \in \{1, \ldots, 5\}$ at successive times $t \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$.

The forward solver computes the steady state pressure and velocity fields, then simulates the tracer advection/diffusion. Figure 5-9 shows the hydraulic head and ve-

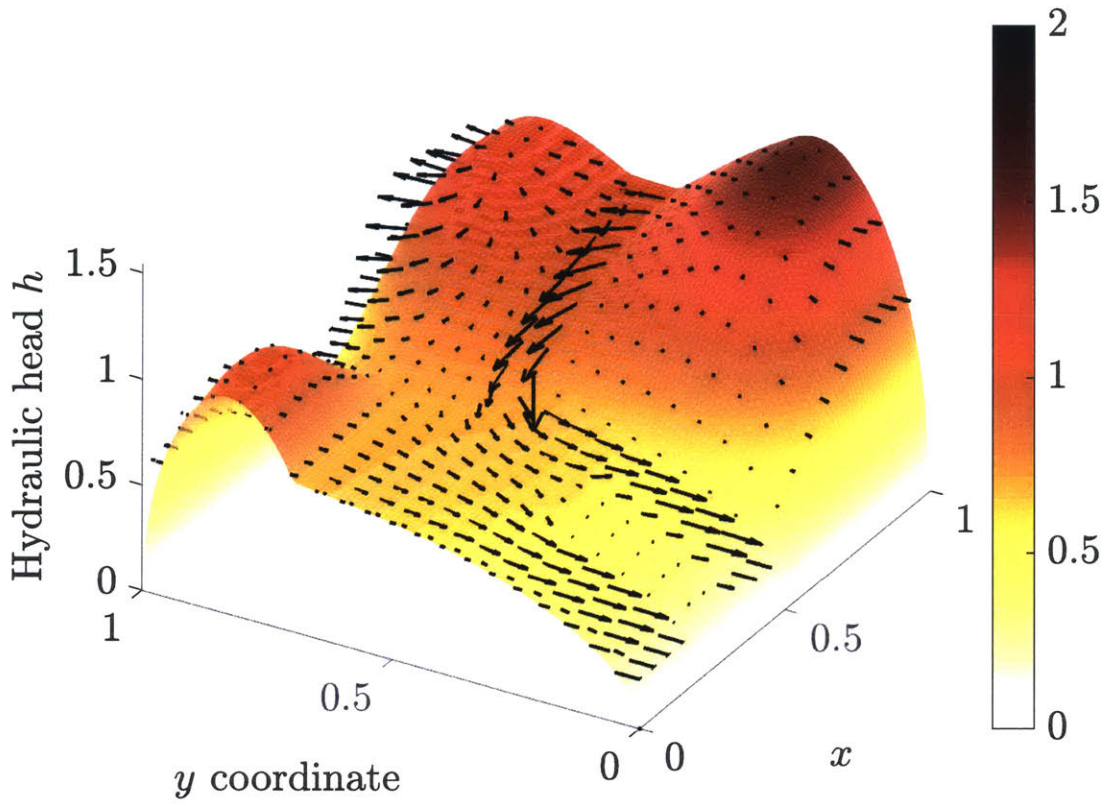

Figure 5-9: Hydraulic head (colormap) $h(x, y)$ computed via (5.12b) and corresponding velocities (5.12d) (arrows), given the conductivity field in Figure 5-8.

locity fields resulting from the true log-conductivity, and figure 5-10 shows the associated tracer concentration field at $t = 0.4$. Overall, the parameter-to-observable map, from the log-conductivities to the time-dependent tracer concentrations, is strongly nonlinear and challenging to approximate. Data for inversion are generated using a standard finite element scheme on a 200×200 mesh. The solver used for inversion (i.e., to evaluate the posterior density at a candidate value of θ) uses a coarser 100×100 mesh. In both cases (generating the data and within the inversion), time integration of the contaminant concentration field uses a Crank-Nicolson scheme. The likelihood assumes additive and independent and identically distributed errors for each observation of tracer concentration, Gaussian with mean zero and variance 1×10^{-2}.

In a serial implementation, each evaluation of the forward model and hence the likelihood requires roughly 13 seconds of computation. Though we will mitigate this cost using local approximations, we also wish to compare our approach with chains

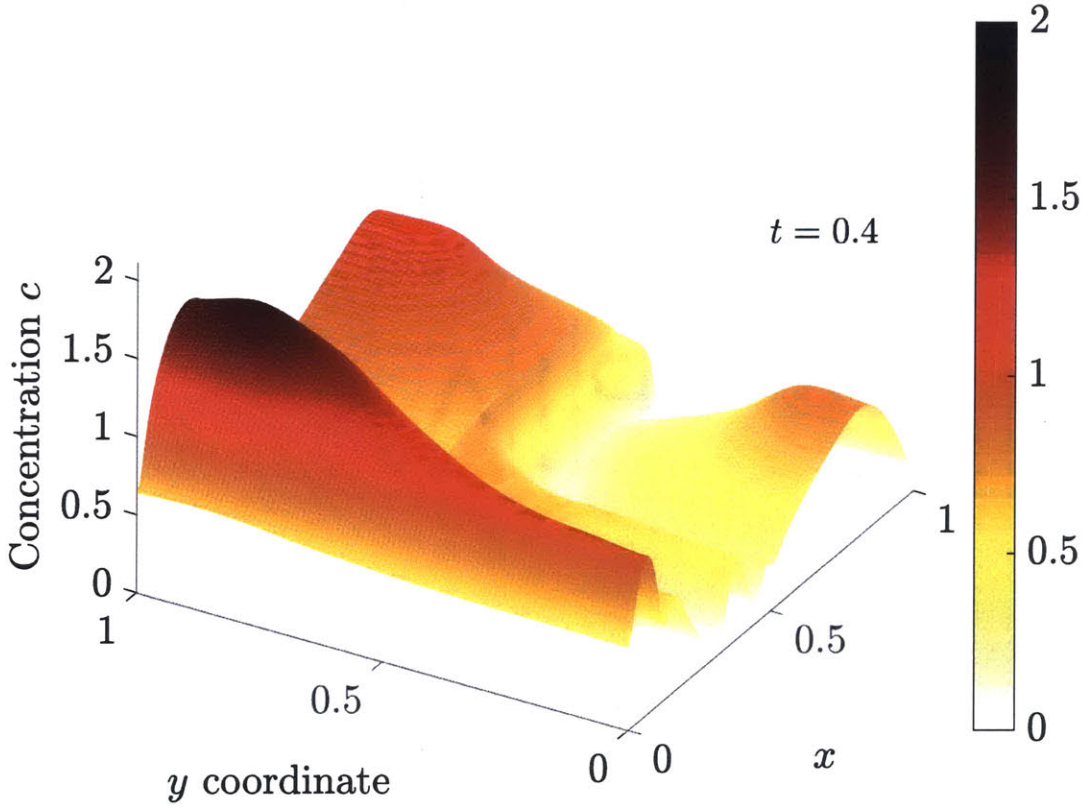

Figure 5-10: The tracer concentration $c(x, y, t = 0.4)$, given the conductivity field in Figure 5-8. The tracer is injected from a well in each corner.

that employ exact evaluations of the forward model. To make such comparisons feasible—and also to reflect computational practice for complex PDE models—we parallelize each forward model evaluation. We use four processors, which reduces the forward model's runtime to roughly 4 seconds of computation. Thus our parallel MCMC scheme actually employs two levels of parallelism: an outer level involving parallel chains, as described in Section 5.2, and an inner level within each forward model evaluation.

The posterior distribution in this problem has no standard analytical form. To establish a baseline for accuracy comparisons, we instead run 31 independent exact+AM chains. Each chain is 10^5 steps long, which requires several days (per chain) of computation. After discarding the first 10^4 samples of each chain as burn-in, the remaining samples are pooled and used to characterize the posterior distribution. Figure 5-11a shows a trace plot of one such exact+AM chain, for all six components of the state. Visually, the transient behavior of the chain appears exhausted well before 10^4 steps, justifying our choice of burn-in. One- and two-dimensional marginals of the posterior distribution, computed using the pooled exact+AM chains, are shown in figure 5-12. The distribution has distinctly non-Gaussian structures, and the regions of high posterior probability seem to concentrate around the "true" parameters given in table 5.1.

While the AM chains appear to mix well for this problem, mMALA proves far less effective. Figure 5-11c shows trace plots of an exact+mMALA chain targeting the same posterior. This calculation is rather laborious (over 415 hours), as direct evaluations of the gradient of the forward model are not available; instead we compute the gradients using finite differences. This simulation is not intended as a practical approach, but rather to assess the performance of mMALA in the absence of local approximations. We find that the chain mixes quite poorly; the ESS after 10^5 MCMC steps is only 80. We do not expect mMALA paired with local approximations to fare any better and, indeed, figure 5-11d shows that mixing is poor for an LA+mMALA chain. Given these results, we focus the rest of this section on AM chains, with a goal of exploring the performance of parallel LA schemes. More broadly, we note that

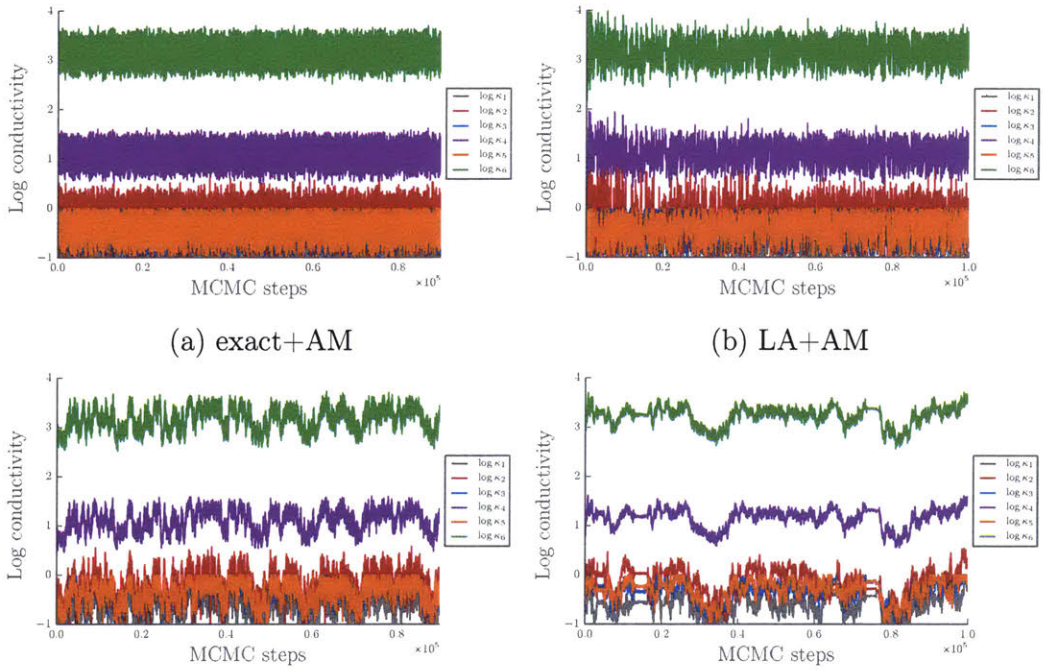

(a) exact+AM

(b) LA+AM

(c) exact+MALA

(d) LA+MALA

Figure 5-11: Tracer transport problem: trace plots for a single MCMC chain (state versus MCMC iteration) either using exact density evaluations or employing a local approximation (LA), paired with either an AM or mMALA proposal.

there is no guarantee that MALA schemes should improve over adaptive Metropolis (or even simple random-walk Metropolis) in low-dimensional problems such as those considered here. The potential for such improvements is problem-dependent and sometimes rather delicate, as was recognized almost immediately when MALA was introduced [Roberts and Rosenthal, 1998].

We first examine the convergence of estimates produced by *single* LA+AM chains. We run 51 independent chains, again discarding the first 10^4 samples of each chain as burn-in. For consistency, we simply choose the same burn-in period for the exact chains and LA chains. If anything, this choice is less favorable to LA—though asymptotically it is immaterial. The mixing of a single LA+AM chain is visualized by the trace plot in Figure 5-11b. Initially, the chain does not mix as quickly as in the exact+AM case, but mixing improves as the approximation is refined, and overall the chain appears to explore the posterior quite efficiently. We also emphasize that the horizontal axis in Figure 5-11b does not reflect computational cost, since the latter is

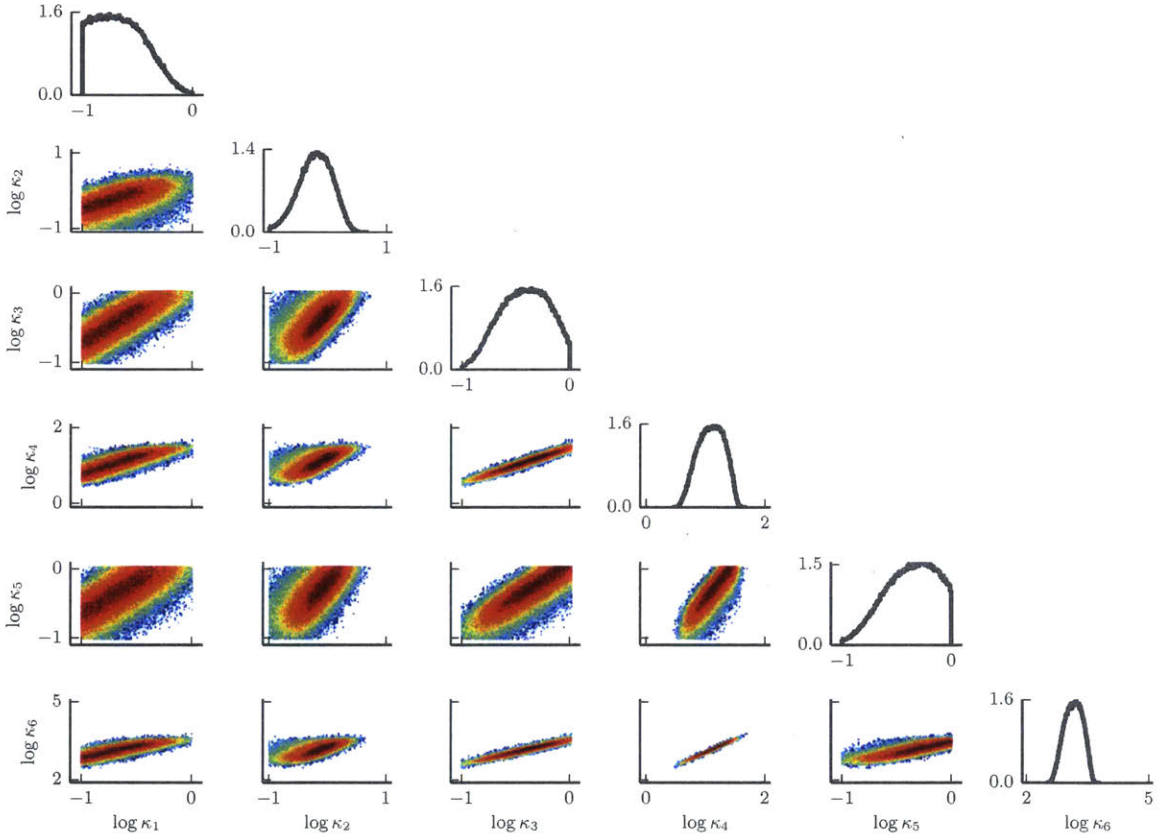

Figure 5-12: One- and two-dimensional posterior marginals of the parameters in the hydrologic tracer transport problem. Bounds on each subplot axis are the upper and lower bounds for the uniform prior on the corresponding parameter (table 5.1).

dominated by target density evaluations rather than MCMC steps.

To assess error versus computational cost, Figure 5-13 shows, for each individual chain, the squared relative error in a running posterior covariance estimate versus wall clock time. The squared relative error $\varepsilon_t^{2,(i)}$ is defined in (5.2), where the reference value C_0 of the posterior covariance is computed by pooling all 2.79×10^6 available exact+AM samples. For comparison, we also plot error versus run time for 31 exact+AM chains. When reporting wall clock times here and below, we include the computational cost of the entire chain, including the cost of portions discarded as burn-in. Error in the LA chains decreases steadily and reaches an accuracy comparable to the exact chains, but with significantly shorter run times. We also notice that decay rate of the expected error (bold red line in figure 5-13) in the LA case seems to accelerate. As noted in the quartic example (where longer chains accentuated this trend), this acceleration is due to the fact that refinements happen less frequently as the chain progresses, while additional MCMC steps continue to reduce the error.

The local approximation sampler becomes even more effective in a parallel chain setting, where concurrent chains are allowed to share posterior density evaluations by building a common $\mathcal{S}_t$. The colored lines in Figure 5-14 show error versus run time for increasing levels of parallelism k, from 1 to 30 chains. To assess the variability of the error, each k-chain simulation is repeated several times; each such realization is shown on the figure. Each individual LA+AM chain (within a group of k) has a fixed length of 10^5 steps and, as before, the first 10^4 samples of each chain are discarded as burn-in. The error plotted on the vertical axis is again the squared relative error in the posterior covariance. Two trends are visible in the colored lines. First, as the number of chains increases, the error decreases. In and of itself, this is not surprising: summing across the chains, we accumulate more MCMC samples and, along the way, seek more model evaluations to refine the local approximations (this will be quantified precisely in subsequent figures). But the colored lines *also* move to the left as the number of parallel chains increases; in other words, both the error *and* the run time are reduced. This trend contrasts with that obtained by simply running exact+AM chains in parallel, an exercise depicted by the gray lines in Figure 5-14. Using this

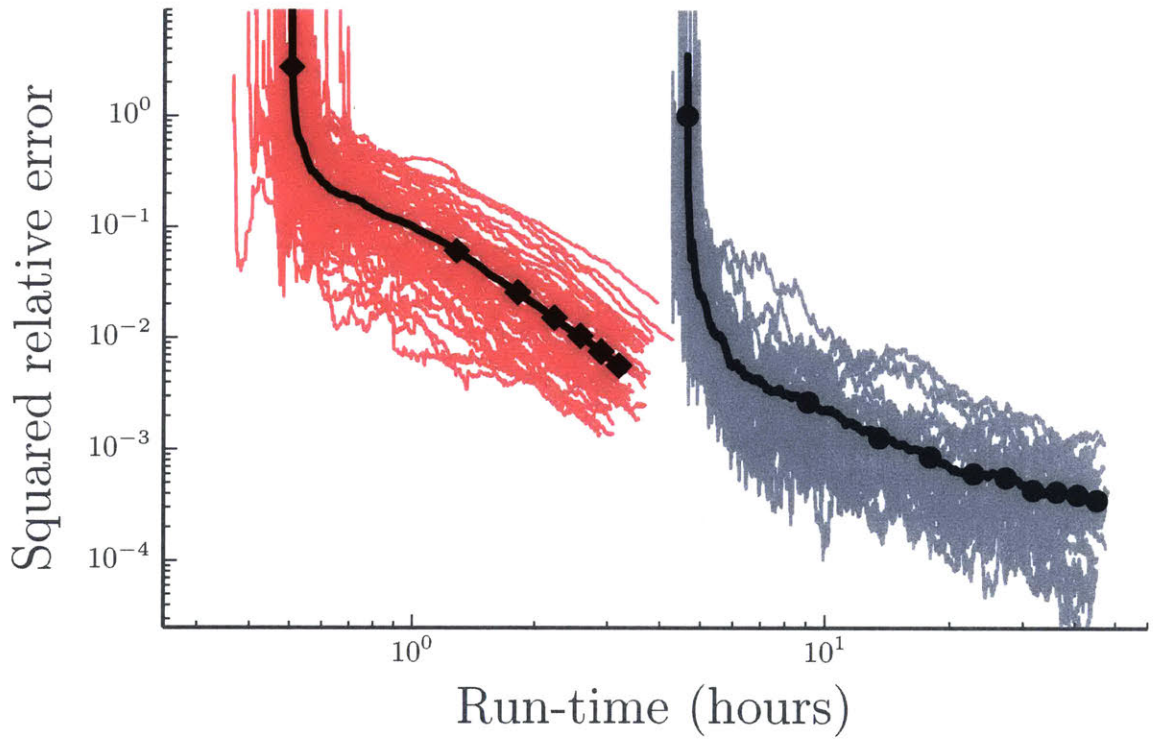

Figure 5-13: Tracer transport problem: relative squared error in the posterior covariance estimates produced by independent single (i.e., not parallel) AM chains, versus run time. The light gray lines correspond to 31 independent exact+AM chains, each of length 10^5. The dark gray line shows the expected error for this exact case. The light red lines correspond to 51 single LA+AM chains, each of length 2×10^5. The dark red line shows the expected error in the approximate case.

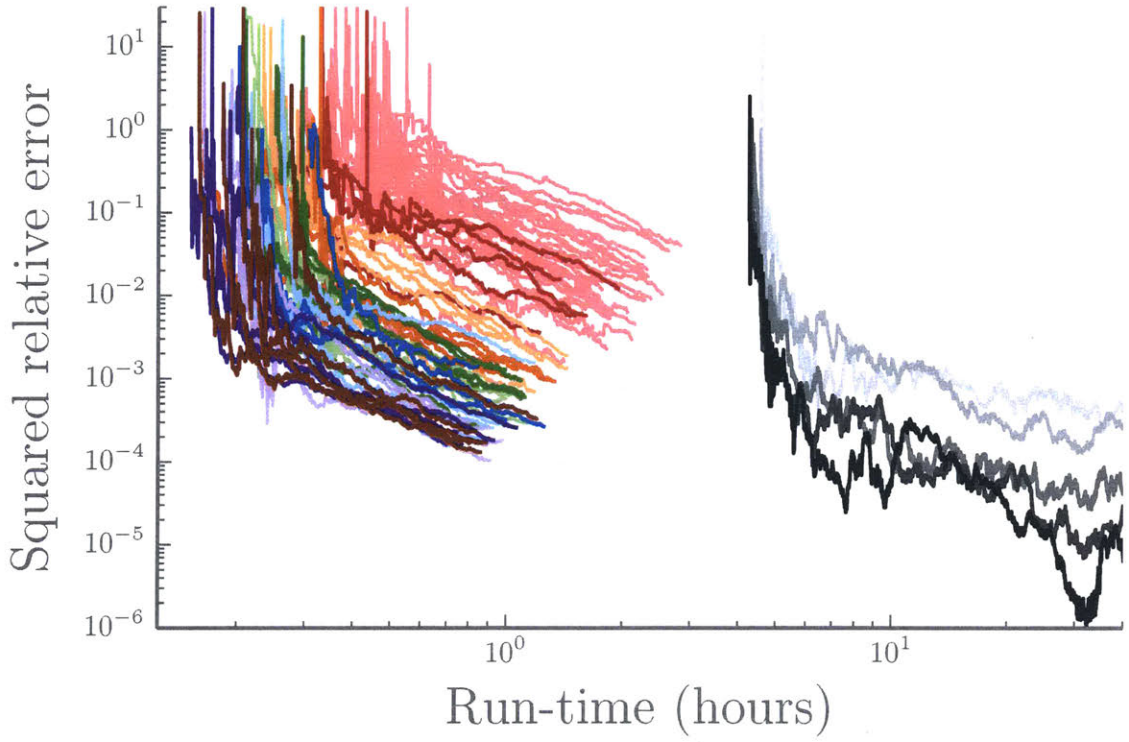

Figure 5-14: Tracer transport problem: relative squared error in the posterior covariance estimates obtained from parallel MCMC chains. The gray lines are computed using exact target density evaluations for $k \in \{1, 2, 4, 8, \text{ or } 16\}$ chains. Darker shades correspond to simulations with more parallel chains. The colored lines are computed using local approximation MCMC. We use $k \in \{1, 2, 4, 6, 8, 10, 13, 16, 20, 25, \text{or } 30\}$ chains corresponding to light red, red, light orange, orange, light green, green, light blue, blue, light purple, purple, and brown, respectively. The error is that of a running covariance estimate obtained by pooling samples from the k concurrent chains. Sharing posterior density evaluations shortens the runtime *and* reduces the error.

naïve parallelization, adding more chains decreases the sampling error but does not affect the run time. Moreover, the run times of LA+AM are one to two orders of magnitude smaller for comparable errors.

We can also characterize the behavior of parallel local approximations by evaluating ESS as a function of computational effort. Figure 5-15 shows ESS as a function of wall clock time. First, as a baseline, consider again running exact+AM chains of length 10^5 in parallel, depicted by gray and black circles. We certainly expect parallel chains to yield a larger ESS once their samples are pooled, and indeed the circles jump upwards as we increase the number of concurrent chains from 1 to 30. Increasing the

number of chains in the exact case does not, however, change the time it takes to simulate each chain; thus the gray and black dots are vertically aligned at the same run times. In the parallel LA+AM cases, depicted by colored diamonds, the story is more interesting. As the number of parallel chains increases, the symbols move upwards *and* to the left, reflecting decreased run times. Several independent realizations of each parallel case are presented, since the simulations are not deterministic. Note that the ESS of a single LA+AM chain (light red) is lower than that of an exact+AM chain of the same length; this is expected, given the mixing comparison at the top of Figure 5-11. Similarly, 30 parallel exact+AM chains have a higher combined ESS than 30 parallel LA+AM chains (the brown diamonds of Figure 5-15). But the latter entail a vastly smaller computational effort. Because of the collaboration among chains, we can compute a larger number of independent samples in less time.

Our second comparison uses a more stringent measure of parallel efficiency: ESS per chain–hour, i.e., the total ESS divided by the number of chains and the wall clock time. This measure removes the intrinsic advantage of having multiple chains. A naïve MCMC parallelization yields no improvement in efficiency according to this metric: the number of independent samples might grow linearly with the number of chains, but this growth is normalized away. Figure 5-16 shows this behavior for exact+AM chains using gray circles. In contrast, the results of parallel local approximation, depicted by colored diamonds, show steady gains in ESS/(chain–hour) with additional parallel chains. This gain is the result of collaboration among the chains in the most computationally expensive element of the inference problem—evaluating the posterior density—by sharing evaluations from which we construct a shared surrogate model. We note that the total number of model evaluations performed during the parallel experiments is still higher than in a single-chain case, but since the additional evaluations are parallelized, the run time is shorter.

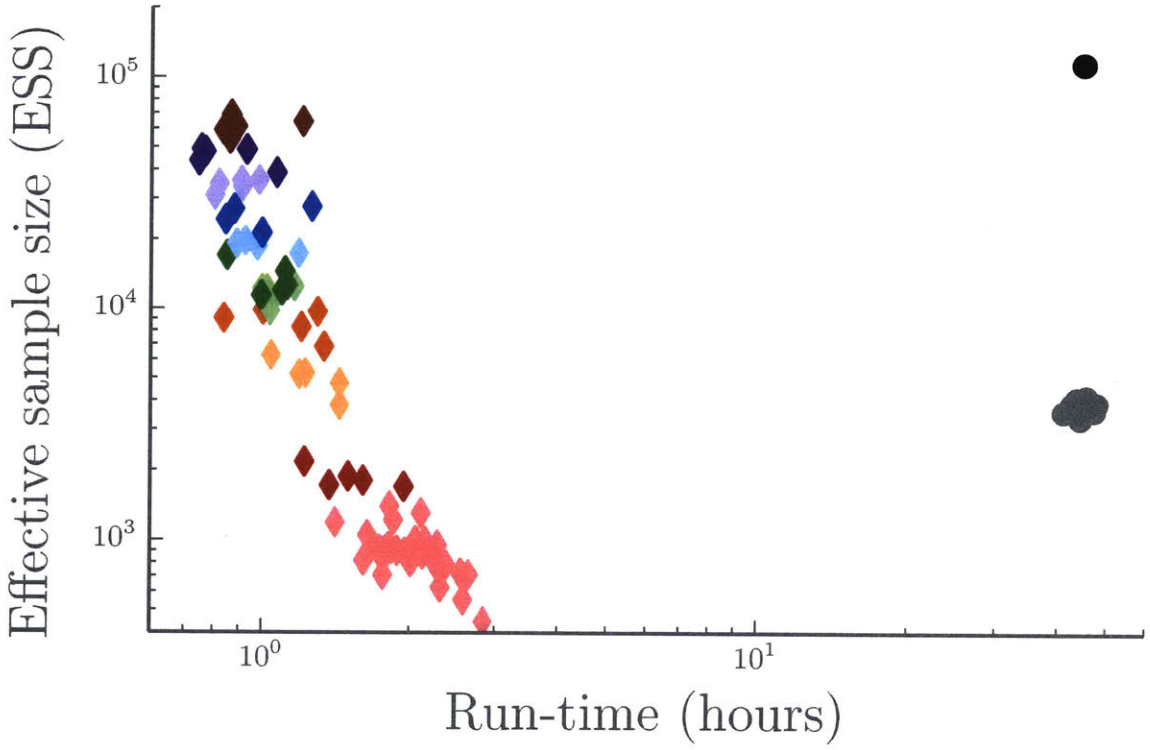

Figure 5-15: Results of the parallel efficiency study on the tracer transport problem, comparing run time to the total ESS across parallel chains. Each symbol represents one (parallel) experiment. Light gray circles and light red diamonds correspond to single chains of length 10^5 (with 10^4 burn-in) using exact evaluations and local approximation, respectively. Each colored diamond represents a different number k of parallel LA+AM chains, $k \in \{1, 2, 4, 6, 8, 10, 13, 16, 20, 25, \text{or } 30\}$, and the colors are as in Figure 5-14. The black circle corresponds to 30 parallel exact+AM chains. Using local approximations, running more chains increases ESS *and* decreases run time.

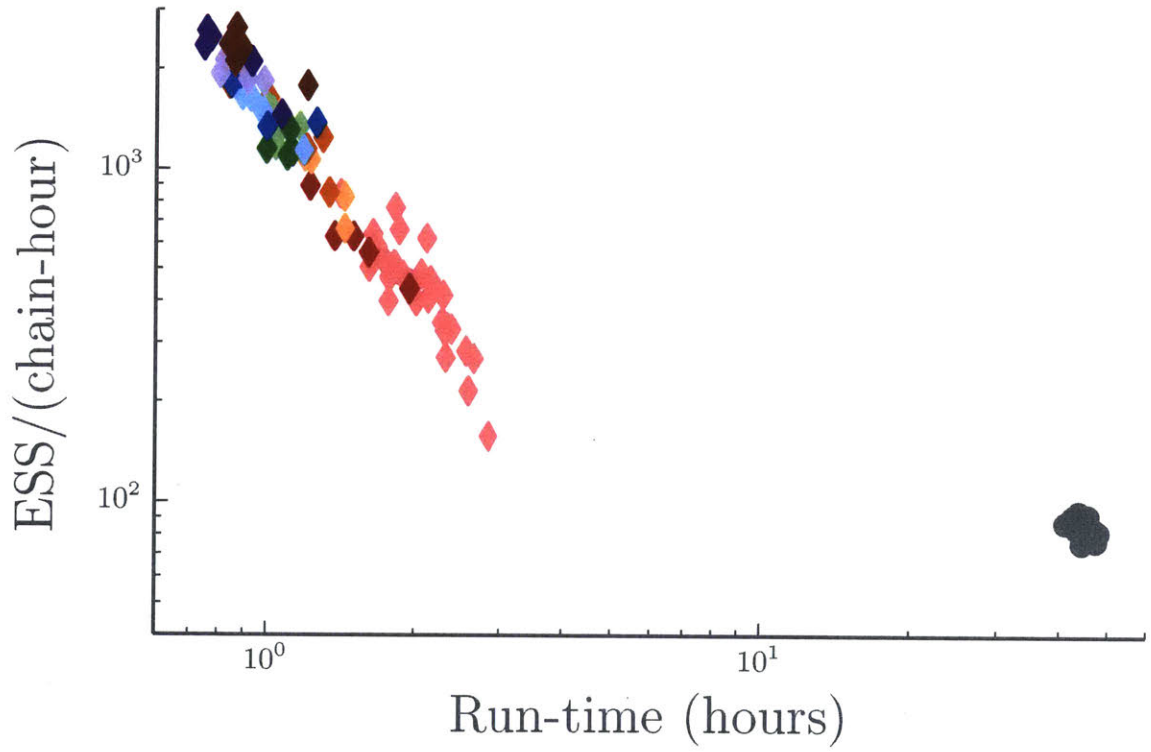

Figure 5-16: Results of the parallel efficiency study on the tracer transport problem, comparing run time to the effective number of samples produced per chain–hour. Each symbol represents one (parallel) experiment. Again, colors from light red to brown correspond to more parallel chains, $k \in \{1, 2, 4, 6, 8, 10, 13, 16, 20, 25, \text{or } 30\}$. Using parallel local approximations, ESS per chain–hour increases with the number of chains.

Chapter 6

Conclusions and future work

6.1 Conclusion

We have focused on two major areas of study: (i) marine ice stream dynamics and (ii) sampling methods targeting distributions with computationally intensive densities. These two, seemingly disparate, areas are tied together within a Bayesian paradigm that combines data and models via statistical tools. In this framework, physical models alone are insufficient; we also model our degree of uncertainty. This allows us to rigorously answer questions like: (i) what is the probability that global mean sea level will rise by a prescribed threshold? (ii) how sensitive is this prediction to changes in input parameters? (iii) what are the primary sources of uncertainty?

Chapters 2 and 3 focused on glacier physics—primarily exploring the glacier stability with respect to bed slope at the terminus. This thesis, along with previous work [Schoof, 2007b, Schoof, 2007a, Thomas and Bentley, 1978, Weertman, 1974], suggests that output flux depends on glacier geometry. Furthermore, the form of this relationship strong impacts whether or not the glacier is subject to the marine ice instability. The marine ice instability is triggered when output flux increases as the glacier retreats into deeper water. This accelerates retreat and leads to rapid deglaciation. Since many glaciers in West Antarctica and Greenland sit on bedrock that slopes downward toward the continent's interior, understanding this relationship is crucial when making predictions about future global mean sea level. We extended previous

result to show that assuming the calving[1] law by [Nick et al., 2010] changes the flux-to-bedrock-depth relationship such that output flux is *not* a monotonic function of bedrock depth. Therefore, a glacier's stability depends on the calving dynamics along with geometric parameters.

This result is, however, not sufficient to quantify uncertainty about our state of knowledge about any given glacier. Calving, for example, is a widely studied but poorly understood process. Predictions are conditioned on calving models that we know do not account for potentially crucial physical processes. Since there is no universally accepted calving law, predictions with any single calving law are inherently uncertain. Furthermore, even if we trust the form of the model, input parameters are often unknown and unobservable. Predictions must account for model and parameter uncertainty. Predicting the future volume of the polar ice sheets and quantifying the associated uncertainty requires a mathematical framework that allows us to combine expert knowledge in the form of physical models with data from multiple sources. Bayesian methods endow unknown and uncertain parameters—model inputs, observed parameters, and quantities of interest (QoI), such as future glacier volume—with probability distributions that quantify our degree of knowledge conditioned on data.

Often, defining a probability distribution that models uncertainty—including model uncertainty, observation noise, and uncertain parameters—is relatively simple. However, computationally intensive physical models and complicated (e.g., non-Gaussian) correlations between QoIs and auxiliary parameters make characterizing these distributions difficult and, thus, limits the practical applicability of these methods. In chapters 4 and 5, we address this by introducing surrogate models that approximate the associated probability density function. Previous work [Conrad et al., 2016] showed that continual refinement of the surrogate model asymptotically characterizes the distribution of interest. We extended this result to include an ideal refinement strategy based on a bias-variance trade-off and to build the surrogate from noisy target density estimates. This allows marginal sampling even if we only have Monte

[1]Calving is the process by which ice break off at the glacier terminus

Carlo estimates of the target density.

The joint development of physical models for marine-terminating glaciers and Bayesian computational methods motivates research in a wide range of topics—from fluid dynamics to sampling methods. In addition, we rely on existing computational tools (e.g., discontinuous Galerkin method) whose improvement would have significant implication on our results. For example, developing faster numerical solvers to approximate predictive flowline models facilities these model's use within Monte Carlo methods. The research presented in this thesis and the future work implied by our results are intermediary, but crucial, steps toward rigorously predicting future polar ice sheet dynamics and quantifying the associated uncertainty.

6.1.1 MIT Uncertainty Quantification (MUQ) software

In addition to being theoretically difficult, using physical models within Bayesian computational methods is practically challenge. Significant expertise in computer science and software engineering is required to numerically solve physical models and implement sampling algorithms. To facilitate future work, we implemented models and statistical tools within a cohesive, called MIT Uncertainty Quantification (MUQ), framework that will facilitate combining these two fields in the future. While a detailed description of the software is beyond the scope of this thesis[2], we believe it is worth noting that this code is available.

6.2 Future work

This thesis motivates a wide array of future work. Each part, ice sheet/stream dynamics and Bayesian computational methods, motivates its own future directions. Additionally, combining these two fields provides a rich set of questions. Here, we will briefly summarize future work for each section and then provide examples of projects that combine the two. We include more detail about the combined work to illustrate how these two fields can be combined.

[2]See `http://muq.mit.edu` for detailed description and documentation.

Generalizing the calving model is the most obvious future direction associated with glacier modeling. We assume a general calving law that prescribes the terminal thickness as a fraction ϕ of the local bedrock depth. We also showed the form of this fraction (e.g., calving-at-flotation versus meltwater calving) strongly impacts whether or not the glacier is susceptible to the marine ice instability. However, as we have emphasized, calving is a poorly understood process and this model is by no means universally accepted. We could generalize this and determine which forms of ϕ lead to stable or unstable systems. Furthermore, calving physics is poorly understood and generating new models—either modifying ϕ or determining new parameterizations altogether—is worthwhile.

The surrogate model used within the sampling methods is *not* a probability distribution after a finite number of refinements. Even more problematic, both in theory and in practice, polynomial approximations have unbounded tails such that

$$\int \breve{\pi}(x)\, dx = \infty, \tag{6.1}$$

where $\breve{\pi}(x)$ is the surrogate model after a finite number of refinements. The unbounded tails explain why, after a finite number of MCMC steps, the chain may be arbitrarily far from areas of nontrivial posterior support. Future work should augment the surrogate model—perhaps combining it with a global surrogate—such that this integral remains finite.

This thesis also motivates, for example, investigating forward uncertainty quantification and sensitivity analysis, inference-for-prediction, and optimal experimental design. All of these projects combine glacier modeling with Bayesian computation. Although rigorously exploring these topics is beyond the scope of this thesis, we outline the problem set up and some initial results that show why these are promising directions for future projects. Ideally, these more detailed outlines show how these two fields can be combined.

6.2.1 Uncertainty quantification and sensitivity analysis

Our primary interest is in how future ice sheet volume[3] responds to changes in climate. Given parameters (basal topography, width, surface mass balance, subshelf melt, and the filling level) that vary in space or time, we can evaluate a forward model (e.g., the flowline model (2.16)) that outputs a single prediction of the future ice volume. If we endow the model inputs with probability distributions—for example, suppose we sample the surface mass balance from a Gaussian process—we can repeatedly sample from these distributions and compute many predictions of future ice volume.

Our first goal is to characterize the predictive distribution given fixed distributions over input parameters. Each prediction of future ice volume—computed by sampling the distribution over input parameters and evaluating the forward model—is a sample from the predictive distribution. Therefore, we can employ sampling methods to estimate properties of the predictive distribution (e.g., the predictive mean, variance, or Sobel indices [Saltelli, 2002, Sobol, 2001]).

We model variability in the QoI as coming from two sources: (i) internal variability and (ii) climate change. The first is modeled by sampling model input parameters from a fixed distribution and the second is modeled by changing the distribution over input parameters. For example, the mean surface mass balance may increase as a function of time in a warming climate. However, we a specific realization of the surface mass balance is sampled from a Gaussian process with linear mean

$$\bar{a}(t) = a_0 + a_1 t. \tag{6.2}$$

On average, $a(t)$ depends linearly on time with slope a_1 but a specific sample may, have periods of increasing or decreasing surface mass balance. Defining a reference climate by setting a_0 we would like to compute the expected future volume as a function of the hyperparameter a_1. In a warming climate, we expect a_1 to be positive; endowing the hyperparameter with a Gamma distribution model variability in climate. Given the hyperparameters (e.g., fixing a_1) models internal variability in a fixed climate.

[3]Since we are interested in future mean global sea level, we are actually only interested in the *grounded* volume. However, for simplicity, we simply say 'volume.'

The goal of this 'forward uncertainty quantification' study is to determine the relationship between the prescribed distribution over hyperparameters (i.e., climate variability) and the predictive distribution, marginalizing internal variability. Monte Carlo methods allow us to characterize the predictive distribution, jointly sampling from distributions over the hyperparameters and distributions over model inputs given the hyperparameters. However, the predictive model is prohibitively computationally intensive. However, building a computationally cheaper surrogate model, as described in chapter 4, makes this study possible.

6.2.2 Inference-for-prediction

We develop a Bayesian inverse modeling framework for predicting future ice sheet volume[4] with associated uncertainty estimates. In this setting, we assume that we have observations (e.g., thickness, surface elevation, basal topography, or surface velocities) that inform uncertain parameters. Recall that Bayesian inverse problems condition uncertain parameters on observations, defining a posterior distribution that updates a prior distribution via the law of total probability. Sampling the posterior (e.g., using Monte Carlo, importance sampling, and/or MCMC) and evaluating a predictive model generates samples of predicted volume. Our goal is to define this distribution in such that we can use local approximation MCMC to directly target the distribution over the predicted volume.

We use the flowline model (2.16)—mass is advected seaward through a laterally confined channel. The thickness and velocity fields are evolved via the force and mass conservation equations. Computational limitations make evolving this model within sampling algorithms intractable. However, we can compute the grounding line flux q_g without computing the thickness/velocity fields by neglecting extensional stress and solve the boundary layer model from section 3.1.2. We then leverage that the quantity of interest is the future grounded volume, not the thickness/velocity fields,

[4]Again, we are actually interested in the *grounded* volume but use 'volume' for simplicity.

which evolves according to

$$\begin{aligned} \dot{V}_g &= \text{[net surface mass balance over the ice stream]} - \text{[grounding line flux]} \\ &= A(x_g) - q_g(x_g), \end{aligned} \tag{6.3a}$$

where $A(x) = \int_0^x aw\,dx'$. Here, $q_g(x)$ denotes the solution to the boundary layer problem. We also need to evolve the grounding line location x_g; we propose a model for the evolution of x_g for the sake of proof-of-concept. However, our model for x_g requires scrutiny. The thickness at the end of the grounded portion is a fraction ϕ of the basal topography: either the flotation thickness or the a grounded calving cliff above flotation. Differentiating this condition with respect to time

$$\frac{d}{dt}\left[h(x_g(t),t) = \varphi(\lambda(t))b(x_g(t))\right] \tag{6.3b}$$

implies

$$h_x|_{x=x_g}\,\dot{x}_g + h_t|_{x=x_g} = \varphi_\lambda \lambda_t b_g + \varphi\, b_x|_{x=x_g}\,\dot{x}_g \tag{6.3c}$$

The flux in the grounded portion is a function of the thickness

$$q = wuh = -\frac{w^{n+2}h^{n+1}}{\left(h + \gamma w^{\frac{1}{n}+1}\right)^n}\left|h_x + b_x\right|^{n-1}\left(h_x + b_x\right). \tag{6.3d}$$

Setting $q(x_g) = q_g(x_g)$ implicitly defines $h_x|_{x=x_g}$. Mass conservation implies that $h_t|_{x=x_g} = -w^{-1}\,q_x|_{x=x_g} + a$. We compute $q_x|_{x=x_g}$ by linearizing the thickness $h \approx h_g + h_x|_{x=x_g}(x - x_g)$ and approximating

$$q_x|_{x=x_g} \approx \frac{q(x_g) - q(x_g - \Delta x)}{\Delta x}. \tag{6.3e}$$

We use (6.3) to evolve the grounded volume V_g and x_g without need to compute the

full thickness and velocity fields.

We also compute a snapshot of the grounded surface velocity $u(x)$ and thickness $h(x)$

$$4\varepsilon \left(h \left| u_x \right|^{\frac{1}{n}-1} u_x \right)_x - w^{-\frac{1}{n}-1} h \left| u \right|^{\frac{1}{n}-1} u - \gamma\theta \left| u \right|^{m-1} u = \left(1 - (1-\theta)r\right) h \left(h_x + \theta b_x \right) \quad (6.4a)$$

$$ws + (wuh)_x = w\left(a - (1-\theta)\mu\right) \quad (6.4b)$$

with $u = 0$ at $x = 0$ and $4\varepsilon h \left| u_x \right|^{\frac{1}{n}-1} u_x = f(x_c)$ and $h = h_c$ at $x = x_c$. The Bayesian inverse problem infers $x_c(0)$, $\{s_k\}$, $\{b_k\}$, and $\{w_k\}$ which reconstructs the change in thickness $s = h_t$, basal topography b, and width w—

$$s(x) = \sum_{k=1}^{K_s} s_k \varphi_k^{(s)}(x), \quad b(x) = \sum_{k=1}^{K_b} b_k \varphi_k^{(b)}(x), \quad \text{and} \quad w(x) = \sum_{k=1}^{K_w} w_k \varphi_k^{(w)}(x), \quad (6.5)$$

given Karhunen-Loève basis functions $\varphi_k^{(s)}(x)$, $\varphi_k^{(b)}(x)$, and $\varphi_k^{(w)}(x)$. We assume that w has large length scales—K_w is small. Most of the inferred parameters are $\{s_k\}$ and $\{b_k\}$. Similarly, we represent the surface mass balance a, subshelf melt μ, and calving law parameter λ

$$a(t) = \sum_{k=1}^{K_a} a_k \varphi_k^{(a)}(t), \quad \mu(x) = \sum_{k=1}^{K_\mu} \mu_k \varphi_k^{(\mu)}(t), \quad \text{and} \quad \lambda(t) = \sum_{k=1}^{K_\lambda} \lambda_k \varphi_k^{(\lambda)}(t). \quad (6.6)$$

We assume that these time-dependent parameters are not informed by the data and, therefore, place no restrictions on the size of K_a, K_μ, or K_λ.

The joint posterior distribution is defined by the conditional predictive, prior, and likelihood distributions. Let $\chi(t) = \{V_g(t), x_g(t)\}$ be the state vector and $z = \{x_c, b_{1:K_b}, w_{1:K_w}, s_{1:K_s}, a_{1:K_a}, \mu_{1:K_\mu}, \lambda_{1:K_\lambda}\}$. Given samples of $\chi(0)$ and z, evolving (6.3)

generates samples of the future grounded volume and grounding line location. Our goal is to characterize the posterior predictive $\pi(\chi(T)|d)$, where d are observations and T is a time-of-interest. The posterior is defined by (letting $\chi = \chi(T)$)

$$\pi^{(T)}(\chi, z|d) \propto \pi^{(T)}(\chi|z)\pi(z)\pi^{(T)}(d|\chi, z) = \underbrace{\pi^{(T)}(\chi|z)}_{\substack{\text{conditional} \\ \text{predictive}}} \underbrace{\pi(z)}_{\text{prior}} \underbrace{\pi(d|z)}_{\text{likelihood}}, \tag{6.7}$$

where we have assumed the data y only depend on z. The density of the posterior predictive distribution is

$$\pi^{(T)}(\chi|y) \propto \int \pi^{(T)}(\chi|z)\pi(z)\pi(y|z)\, dz. \tag{6.8}$$

We assume that we can (cheaply) evaluate the prior density $\pi(z)$. The likelihood density is

$$\pi(y|z) = \mathcal{N}(h(z); y_h, \Sigma_h)\mathcal{N}(u(z); y_u, \Sigma_u), \tag{6.9}$$

where $h(z)$ and $u(z)$ satisfy (6.4) and y_h and y_u are observations of the thickness and velocity, respectively, with corresponding covariance matrices Σ_h and Σ_u. Given z, the initial thickness field h (computed using (6.4)) defines the expected initial grounded volume $\bar{V}_g$ and grounding line location $\bar{x}_g$

$$h(\bar{x}_g) = \phi(\lambda(0))b(\bar{x}_g) \quad \text{and} \quad \bar{V}_g = \textstyle\int_0^{\bar{x}_g} h\, dx; \tag{6.10}$$

let $\bar{\chi} = \{\bar{V}_g, \bar{x}_g\}$. The predictive model (6.3) defines an invertible map—$\chi(t) = \Psi(\chi(0); t, z)$ and $\chi(0) = \Psi^{-1}(\chi(t); t, z)$; the initial conditions are computed by integrating (6.3) backward in time. We simulate model inadequacy—missing physics, truncation error in the KL decompositions, and numerical error—by defining an initial distribution over χ

$$\pi^{(T)}(\chi|z) = \mathcal{N}\left(\Psi^{-1}(\chi; T, z); \bar{\chi}, \Sigma_\chi\right) \det J, \tag{6.11}$$

where J is the Jacobian of the inverse map Ψ^{-1}. We estimate (6.8) using importance

sampling

$$\pi^{(T)}(\chi|y) \propto \hat{\pi}^{(T)}(\chi|y) = \frac{1}{M}\sum_{i=1}^{M} \pi^{(T)}(\chi|z^{(i)}) \frac{\pi(z^{(i)})\pi(y|z^{(i)})}{q(z^{(i)})} \tag{6.12}$$

where $z^{(i)} \sim q_Z$. We use local approximation MCMC to characterize the posterior predictive and, given samples of the initial grounded volume and grounding line location.

6.2.3 Optimal experimental design

Borehole data are essential for calibrating ice sheet models. However, field expeditions for acquiring borehole data are often time-consuming, expensive, and even dangerous. It is thus essential to plan the best sampling locations that maximize the value of data and minimize the costs and risks. We[5] present an uncertainty quantification (UQ) workflow based on rigorous probability framework to achieve these objectives.

First, we employ an optimal experimental design (OED) procedure to compute the borehole locations that yields the highest expected information gain in an attempt to reconstruct past surface temperatures. The problem setup is similar to that of [Dahl-Jensen et al., 1998], who solve a Bayesian inference problem to reconstruct the glacier's past surface temperature given data from a vertical borehole. In their formulation the location of the borehole and, therefore, the data are take as a given. Here, we determine the optimal location to place the borehole to maximize the utility of the data for reconstruction past surface temperature.

We take into account practical considerations of location accessibility (e.g., proximity to research sites, terrain, and ice velocity may affect feasibility of drilling) and robustness (e.g., real-time constraints such as weather may force researchers to drill at sub-optimal locations near those originally planned). To address these issues, we incorporate a penalty reflecting accessibility as well as sensitivity to deviations from the optimal locations. Next, we extract vertical temperature profiles from these boreholes

[5]Joint work with Xun Huan (Assistant Professor at the University of Michigan, Department of Mechanical Engineering)

and formulate a Bayesian inverse problem to reconstruct historical surface temperatures. The procedures employ a model of temperature advection/diffusion. The top boundary condition, corresponding to surface temperatures, is calibrated via efficient Markov chain Monte Carlo (MCMC). The overall procedure can then be iterated to choose new optimal borehole locations for the next expeditions.

We illustrate that UQ offers powerful tools for designing experiments, calibrating models, making predictions, and assessing sensitivity, all performed under an uncertain environment. We note that our frame is broadly applicable; we use historical surface temperature reconstruction as an example. We develop a theoretical framework as well as practical software that leverages state-of-the-art UQ tools within an intuitive workflow. Ultimately, we show how to leverage UQ and Bayesian statistics to combine data and models within the context of environmental and climate modeling.

Problem formulation

Consider the flowline model (2.16); given the bedrock elevation, ice stream width and time-dependent parameters (surface mass balance, subshelf melt, and the filling level), we evolve the ice stream in time. Suppose we take the time-evolution of this model as a given—e.g., we know how the *geometry* of the ice sheet evolves. Embedded within this assumption is, of course, many rich questions that remain unanswered: constructing past and future surface mass balance alone is a nontrivial problem [Dahl-Jensen et al., 1993]. However, for the sake of proof-of-concept, we assume that the evolving geometry is given.

We make a second assumption that ice temperature evolves *independently* from the velocity and the thickness. Temperature would typically enter into flowline models like (2.16) though the flow law parameters; ice viscosity is temperature dependent. The nondimensional parameter

$$\varepsilon = \frac{1}{2} A^{-\frac{1}{n}} C_l^{-1} \left(\frac{[w]}{[x]} \right)^{\frac{1}{n}+1} \tag{6.13}$$

depends on temperature through the flow law coefficient A. We assume that the

temperature does not vary in such a way to make ε large and, therefore, render the extensional stress non-negligible in the ice sheet interior. Again, this potentially limiting, but simplifies the problem set up for this proof-of-concept.

Our 'forward model' evolves the two-dimensional temperature field $T(x, z)$. Given the thickness and velocity fields, we model temperature via an advection-diffusion equation

$$T_t + \kappa(T_{xx} + T_{zz}) + uT_x = f(x, z), \tag{6.14}$$

where κ is a nondimensional parameter and we have assumed that vertical advection is negligible. The forcing term $f(x, z)$ accounts for internal heating; for now, we set $f(x, z) = 0$. The boundary conditions of this model are

$$T(x, h(x, t), t) = T_s(t) \quad \text{and} \quad T_x|_{z=0, x=0, x=x_g} = G(x, z, t) \tag{6.15}$$

where $G(x, t)$ is a heat flux through the glacier boundaries (e.g., geothermal heating) and $T_s(t)$ is the surface temperature. Prescribing[6] a 'true' basal topography, width, surface mass balance, subshelf melt, and filling level, the steady-state profile for the glacier is shown in figure 6-1. We assume constant nondimensional temperature $T = 0.2$; we have nondimensionalized the temperature field given a dimensional parameter $[T]$ such that $T \sim O(1)$. Figures 6-1 and 6-3 show the evolution of this model given boundary conditions in figure 6-2.

We endow the surface temperature $T_s(t)$—which is the target of our underlying inference problem—with a Gaussian process prior using an isotropic, squared-exponential covariance kernel

$$\Sigma(x_1, x_2) = \sigma^2 \exp\left(\frac{(x_1 - x_2)^2}{2L}\right). \tag{6.16}$$

We then parameterize the field with a Karhunen-Loève (KL) expansion [Adler, 1981]:

[6] We prescribe these parameters by sampling from a Gaussian process in either space or time.

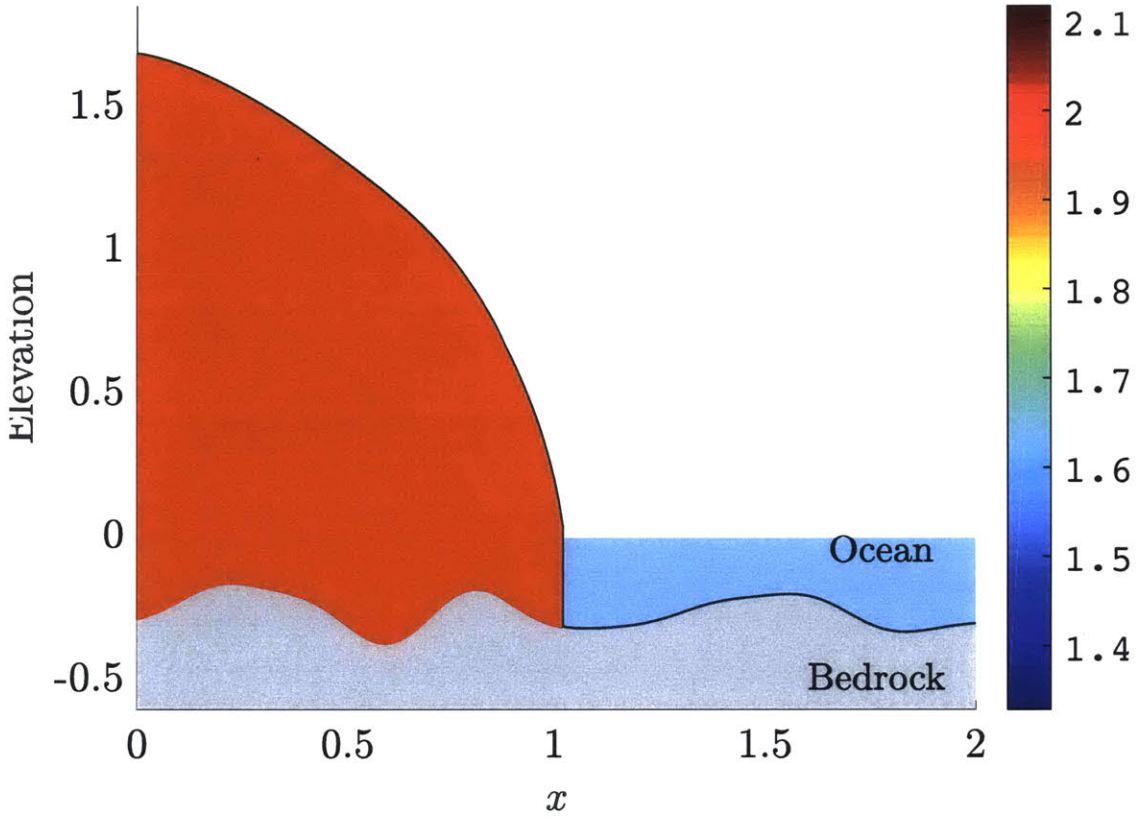

Figure 6-1: The steady state thickness profile for the grounded portion of a glacier given basal topography, width, surface mass balance, subshelf melt, and filling level

$$T_s(t) = \tau_0 + \sum_{i=1}^{\infty} \tau_i \sqrt{\lambda_i} \varphi_i(t), \tag{6.17}$$

where τ_0 is the mean and λ_i and $\varphi_i(x)$ are the Eigenvalues and Eigenfunctions, respectively, of the integral operator on $[0, 1]$ defined by the kernel Σ. The parameters τ_i (for $i > 0$) inherit independent standard normal priors, $\kappa_i \sim \mathcal{N}(0, 1)$. Figure 6-2 shows the 'true' surface temperature field.

Our goal is to find the optimal location to drill a borehole ξ, collecting vertical temperature data at the final time $t = t_f$. Assuming that the we know the geometric evolution of the glacier and that the initial temperature is constant we formulate an optimization problem that determines this location.

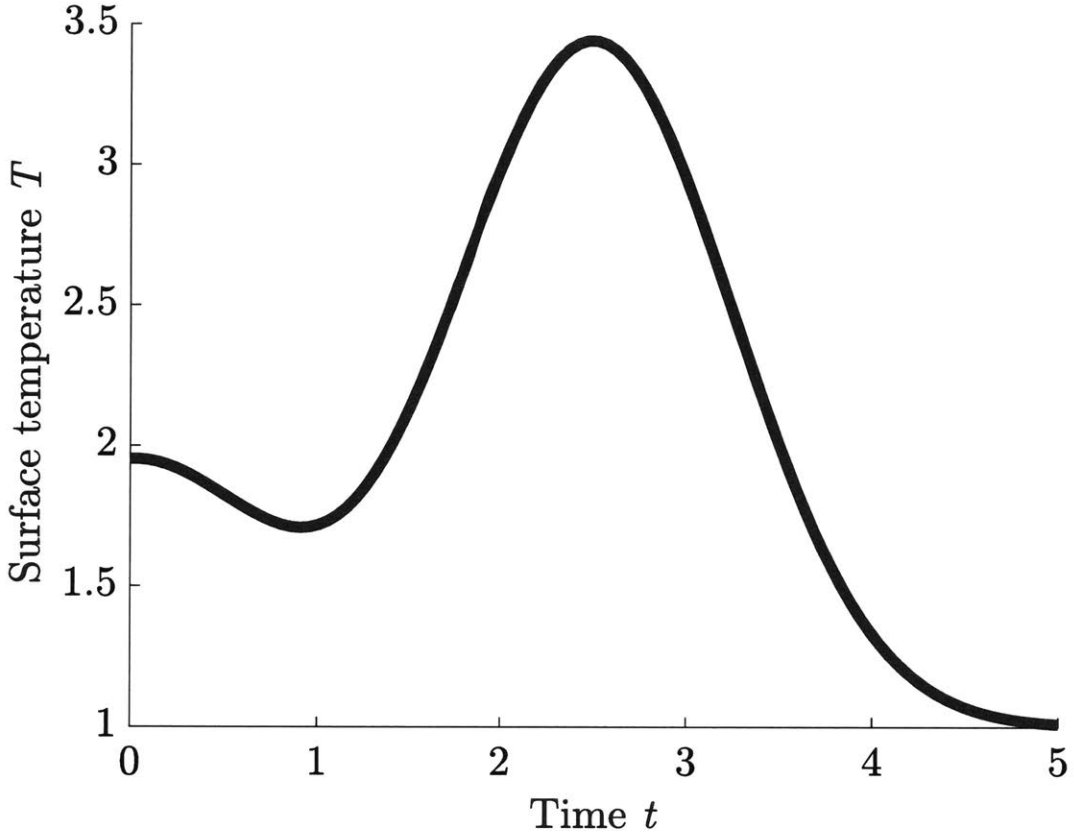

Figure 6-2: The prescribed surface temperature as a function of time, which enters the temperature-dependent model through the boundary conditions (6.15).

Optimal experimental design

We begin by resetting our notation to be consistent with the formulation of local approximation importance sampling (see chapter 4). Let x be the coefficients of the KL expansion for surface temperature (the parameter we want to infer), y the vertical temperature profile (data), and ξ the location to drill a borehole (design variables). We take a Bayesian perspective in characterizing the uncertainty of x, and incorporate the data through Bayes' theorem:

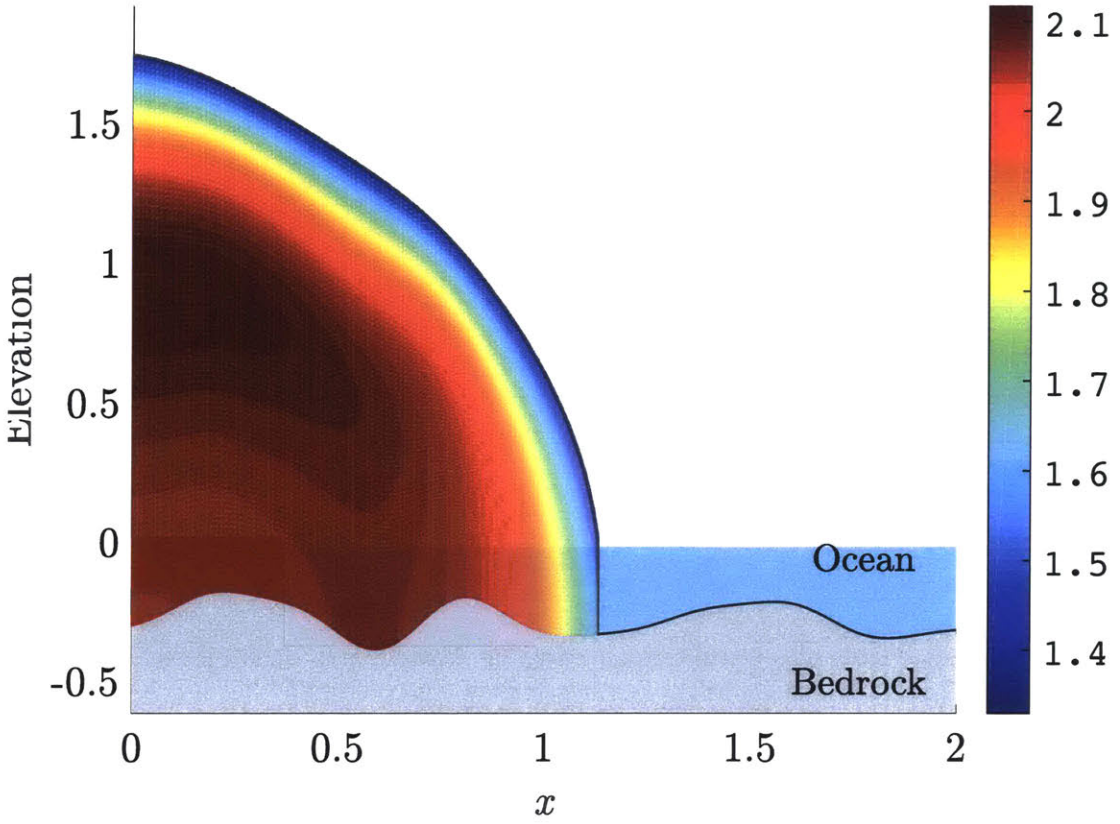

Figure 6-3: The final thickness profile for the grounded portion of a glacier after evolving from the steady state profile in figure 6-1 and prescribing the 'true' boundary condition in figure 6-2.

$$\overbrace{p(x|y,\xi)}^{\text{Posterior}} = \frac{\overbrace{p(y|x,\xi)}^{\text{Likelihood}}\ \overbrace{p(x)}^{\text{Prior}}}{\underbrace{p(y|\xi)}_{\text{Evidence}}}. \tag{6.18}$$

We assume the *prior* is independent of design, i.e., $p(x|\xi) = p(x)$. In other words, the prior should not change based on what experiment we *plan* to do.

Following [Huan and Marzouk, 2013], we establish an *expected utility* based on the Kullback-Leibler (KL) divergence between the posterior and prior, reflecting the average information gained on x from performing the experiment at ξ:

$$U(\xi) = \int_{y|\xi} \int_x \ln\left(\frac{p(x|y,\xi)}{p(x)}\right) p(x|y,\xi)\, dx\, p(y|\xi)\, dy \tag{6.19}$$

$$= \int_{y|\xi} \int_x \left(\ln\left(p(y|x,\xi)\right) - \ln\left(p(y|\xi)\right)\right) p(y|x,\xi) p(x)\, dx\, dy, \tag{6.20}$$

$$= \int_{y|\xi} \int_x \left(\ln\left(p(y|x,\xi)\right) - \ln\left(\mathbb{E}_{p(x)}[p(y|x,\xi)]\right)\right) p(y|x,\xi) p(x)\, dx\, dy, \tag{6.21}$$

$$\leq \int_{y|\xi} \int_x \left(\ln\left(p(y|x,\xi)\right) - \mathbb{E}_{p(x)}[\ln\left(p(y|x,\xi)\right)]\right) p(y|x,\xi) p(x)\, dx\, dy, \tag{6.22}$$

where the inequality follows from Jensen's inequality: $\mathbb{E}[\ln(X)] \leq \ln\left(\mathbb{E}[X]\right)$. The optimal experimental design is then

$$\xi^* = \arg\max_{\xi} U(\xi). \tag{6.23}$$

Preliminary results: evaluation of U with surrogate models

Evaluating the expected utility (6.20) is computationally infeasible. Approximating either (6.20) or (6.22) requires a *doubly* nested Monte Carlo evaluation. Crudely, using $M_1 = M_2 = 10^3$ Monte Carlo samples requires 10^6 model evaluations *per optimization step*. Some computational savings can be achieved by doing these model evaluations in parallel, however, this problem remains intractable without many cores. Introducing a surrogate for the likelihood makes the algorithm computationally cheaper. We reformulate this problem

$$U(\xi) \leq \int_x \underbrace{\int_{y|\xi} \left(\ln\left(p(y|x,\xi)\right) - \mathbb{E}_{p(x)}[\ln\left(p(y|x,\xi)\right)]\right) p(y|x,\xi)\, dy}_{f(\xi,x) \approx \hat{f}(\xi,x)}\ p(x)\, dx \tag{6.24}$$

so that the integrand $f(\xi, x)$is a function of the design variables ξ and the QoI x. We can compute noisy estimates of $f(\xi, x)$ using Monte Carlo; these estimates are still a doubly nested Monte Carlo estimation and, therefore, remain computationally bur-

densome. However, we can apply adaptive threshold local approximation importance sampling (AT-LAIS—algorithm 3 in chapter 4) to build a surrogate model.

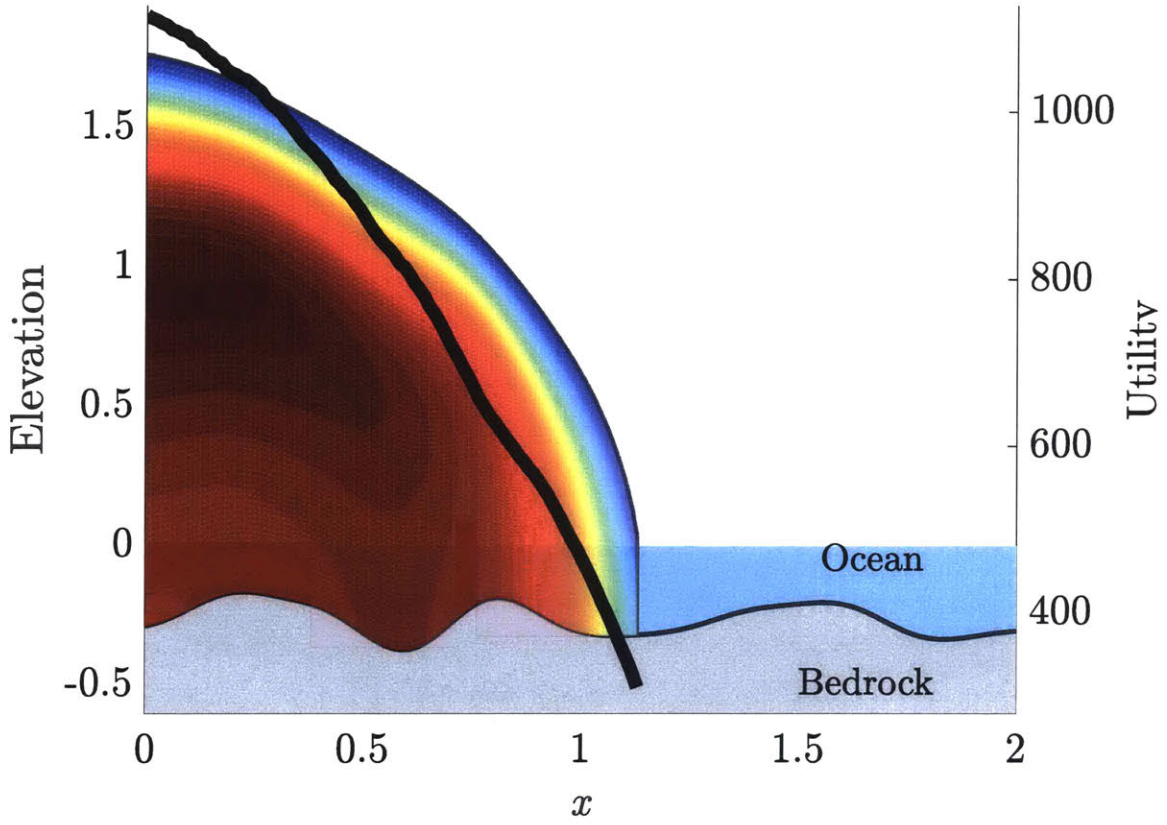

Figure 6-4: The expected utility function computing using AT-LAIS at each point in the final glacier profile.

We apply AT-LAIS to this problem to estimate the expected utility at each point the glacier's domain. Figure 6-4 shows the result. We see, intuitively, that the optimal borehole location to reconstruct the surface temperature boundary condition is at the ice divide. We emphasize that this is proof-of-concept: using a crude (but still computationally intensive) model we are able to rigorously reconstruct an intuitive result. Future work using these tools will include more physically accurate models, improved biasing distributions for the importance sampling estimates, and the ability to drill multiple boreholes.

Bibliography

[Adler, 1981] Adler, R. J. (1981). *The geometry of random fields.* SIAM. 248

[Allen et al., 2014] Allen, M. R., Barros, V. R., Broome, J., Cramer, W., Christ, R., Church, J. A., Clarke, L., Dahe, Q., Dasgupta, P., Dubash, N. K., et al. (2014). IPCC fifth assessment synthesis report-climate change 2014 synthesis report. 33, 34

[Alley et al., 2003] Alley, R. B., Marotzke, J., Nordhaus, W., Overpeck, J., Peteet, D., Pielke, R., Pierrehumbert, R., Rhines, P., Stocker, T., Talley, L., et al. (2003). Abrupt climate change. *Science*, 299(5615):2005–2010. 33

[Alnæs et al., 2015] Alnæs, M. S., Blechta, J., Hake, J., Johansson, A., Kehlet, B., Logg, A., Richardson, C., Ring, J., Rognes, M. E., and Wells, G. N. (2015). The fenics project version 1.5. *Archive of Numerical Software*, 3(100). 74, 96, 134

[Andrieu and Roberts, 2009] Andrieu, C. and Roberts, G. O. (2009). The pseudo-marginal approach for efficient monte carlo computations. *The Annals of Statistics*, pages 697–725. 35, 39, 150, 151, 196

[Arnold et al., 2002] Arnold, D. N., Brezzi, F., Cockburn, B., and Marini, L. D. (2002). Unified analysis of discontinuous galerkin methods for elliptic problems. *SIAM journal on numerical analysis*, 39(5):1749–1779. 70

[Balay et al., 2017a] Balay, S., Abhyankar, S., Adams, M. F., Brown, J., Brune, P., Buschelman, K., Dalcin, L., Eijkhout, V., Gropp, W. D., Kaushik, D., Knepley, M. G., May, D. A., McInnes, L. C., Rupp, K., Sanan, P., Smith, B. F., Zampini, S., Zhang, H., and Zhang, H. (2017a). PETSc users manual. Technical Report ANL-95/11 - Revision 3.8, Argonne National Laboratory. 107, 133

[Balay et al., 2017b] Balay, S., Abhyankar, S., Adams, M. F., Brown, J., Brune, P., Buschelman, K., Dalcin, L., Eijkhout, V., Gropp, W. D., Kaushik, D., Knepley, M. G., May, D. A., McInnes, L. C., Rupp, K., Smith, B. F., Zampini, S., Zhang, H., and Zhang, H. (2017b). PETSc Web page. `http://www.mcs.anl.gov/petsc`. 107, 133

[Balay et al., 1997] Balay, S., Gropp, W. D., McInnes, L. C., and Smith, B. F. (1997). Efficient management of parallelism in object oriented numerical software libraries. In Arge, E., Bruaset, A. M., and Langtangen, H. P., editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhäuser Press. 107, 133

[Bassis and Jacobs, 2013] Bassis, J. and Jacobs, S. (2013). Diverse calving patterns linked to glacier geometry, nat. geosci., 6, 833–836. 11, 37, 50, 55

[Bassis and Walker, 2011] Bassis, J. and Walker, C. (2011). Upper and lower limits on the stability of calving glaciers from the yield strength envelope of ice. *Proc. R. Soc. Lond. A*, 468:913–931. 11, 37, 50, 55

[Beaumont, 2003] Beaumont, M. A. (2003). Estimation of population growth or decline in genetically monitored populations. *Genetics*, 164(3):1139–1160. 150, 196

[Beaumont, 2010] Beaumont, M. A. (2010). Approximate Bayesian computation in evolution and ecology. *Annual review of ecology, evolution, and systematics*, 41:379–406. 35, 39

[Beaumont et al., 2002] Beaumont, M. A., Zhang, W., and Balding, D. J. (2002). Approximate Bayesian computation in population genetics. *Genetics*, 162(4):2025–2035. 35, 38

[Bennett, 2003] Bennett, M. R. (2003). Ice streams as the arteries of an ice sheet: their mechanics, stability and significance. *Earth-Science Reviews*, 61(3-4):309–339. 43

[Bernardo et al., 2008] Bernardo, J., Bayarri, M., Berger, J., Dawid, A., Heckerman, D., Smith, A., West, M., et al. (2008). Gaussian processes to speed up hybrid Monte Carlo for expensive Bayesian integrals. *Bayesian Statistics*, 7:651–659. 11, 37, 153

[Billingsley, 2009] Billingsley, P. (2009). *Convergence of probability measures*, volume 493. John Wiley & Sons. 145

[Bliznyuk et al., 2012] Bliznyuk, N., Ruppert, D., and Shoemaker, C. A. (2012). Local derivative-free approximation of computationally expensive posterior densities. *Journal of Computational and Graphical Statistics*, 21(2):476–495. 11, 37, 153, 206

[Brooks et al., 2011] Brooks, S., Gelman, A., Jones, G. L., and Meng, X.-L. (2011). *Handbook of markov chain monte carlo*. Chapman and Hall/CRC. 11, 37, 145, 147

[Brooks and Roberts, 1998] Brooks, S. P. and Roberts, G. O. (1998). Assessing Convergence of Markov Chain Monte Carlo Algorithms. *Statistics and Computing*, 8(4):319–335. 208

[Calderhead, 2014] Calderhead, B. (2014). A general construction for parallelizing Metropolis-Hastings algorithms. *Proceedings of the National Academy of Sciences*, 111(49):17408–17413. 209

[Cappe et al., 2004] Cappe, O., Guillin, A., Marin, J.-M., and Robert, C. P. (2004). Population Monte Carlo. *Journal of Computational and Graphical Statistics*, 13(4):907–929. 208

[Chaloner and Verdinelli, 1995] Chaloner, K. and Verdinelli, I. (1995). Bayesian experimental design: A review. *Statistical Science*, pages 273–304. 11, 37

[Chen and Schwab, 2015] Chen, P. and Schwab, C. (2015). Sparse-grid, reduced-basis Bayesian inversion. *Computer Methods in Applied Mechanics and Engineering*, 297:84–115. 206

[Christen and Fox, 2005] Christen, J. A. and Fox, C. (2005). Markov chain Monte Carlo Using an Approximation. *Journal of Computational and Graphical Statistics*, 14(4):795–810. 38, 207

[Church et al., 1991] Church, J. A., Godfrey, J. S., Jackett, D. R., and McDougall, T. J. (1991). A model of sea level rise caused by ocean thermal expansion. *Journal of Climate*, 4(4):438–456. 34

[Cockburn et al., 2000] Cockburn, B., Karniadakis, G. E., and Shu, C.-W. (2000). The development of discontinuous galerkin methods. In *Discontinuous Galerkin Methods*, pages 3–50. Springer. 70

[Conn et al., 2009] Conn, A. R., Scheinberg, K., and Vicente, L. N. (2009). *Introduction to derivative-free optimization*, volume 8. Siam. 154, 158, 161, 162, 163, 164

[Conrad et al., 2018] Conrad, P. R., Davis, A. D., Marzouk, Y. M., Pillai, N. S., and Smith, A. (2018). Parallel local approximation mcmc for expensive models. *SIAM/ASA Journal on Uncertainty Quantification*, 6(1):339–373. 206, 207, 208, 211

[Conrad and Marzouk, 2013] Conrad, P. R. and Marzouk, Y. M. (2013). Adaptive smolyak pseudospectral approximations. *SIAM Journal on Scientific Computing*, 35(6):A2643–A2670. 153

[Conrad et al., 2016] Conrad, P. R., Marzouk, Y. M., Pillai, N. S., and Smith, A. (2016). Accelerating asymptotically exact mcmc for computationally intensive models via local approximations. *Journal of the American Statistical Association*, 111(516):1591–1607. 153, 154, 180, 182, 183, 184, 206, 207, 209, 210, 211, 238

[Constantine et al., 2014] Constantine, P. G., Dow, E., and Wang, Q. (2014). Active subspace methods in theory and practice: applications to kriging surfaces. *SIAM Journal on Scientific Computing*, 36(4):A1500–A1524. 11, 37, 38

[Cotter et al., 2010] Cotter, S. L., Dashti, M., and Stuart, A. M. (2010). Approximation of Bayesian Inverse Problems. *SIAM Journal of Numerical Analysis*, 48(1):322–345. 207

[Cowles and Carlin, 1996] Cowles, M. K. and Carlin, B. P. (1996). Markov Chain Monte Carlo Convergence Diagnostics : A Comparative Review. *Journal of the American Statistical Association*, 91(434):883–904. 208

[Craiu et al., 2009] Craiu, R. V., Rosenthal, J., and Yang, C. (2009). Learn From Thy Neighbor : Parallel-Chain and Regional Adaptive MCMC. *Journal of the American Statistical Association*, 104(488):1454–1466. 208

[Cui et al., 2011] Cui, T., Fox, C., and O'Sullivan, M. J. (2011). Bayesian calibration of a large-scale geothermal reservoir model by a new adaptive delayed acceptance Metropolis Hastings algorithm. *Water Resources Research*, 47(10):W10521. 207

[Cui et al., 2014] Cui, T., Martin, J., Marzouk, Y. M., Solonen, A., and Spantini, A. (2014). Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems*, 30(11):114015. 11, 37, 38

[Cui et al., 2016] Cui, T., Marzouk, Y. M., and Willcox, K. (2016). Scalable posterior approximations for large-scale Bayesian inverse problems via likelihood-informed parameter and state reduction. *Journal of Computational Physics*, 315:363–387. 206

[Dahl-Jensen et al., 1993] Dahl-Jensen, D., Johnsen, S., Hammer, C., Clausen, H., and Jouzel, J. (1993). Past accumulation rates derived from observed annual layers in the grip ice core from summit, central greenland. In *Ice in the climate system*, pages 517–532. Springer. 247

[Dahl-Jensen et al., 1998] Dahl-Jensen, D., Mosegaard, K., Gundestrup, N., Clow, G. D., Johnsen, S. J., Hansen, A. W., and Balling, N. (1998). Past temperatures directly from the greenland ice sheet. *Science*, 282(5387):268–271. 246

[Dupont and Alley, 2005] Dupont, T. and Alley, R. (2005). Assessment of the importance of ice-shelf buttressing to ice sheet flows. *Geophys. Res. Letts.*, 32(L04503):doi:10.1029/2004GL022024. 11, 37, 38, 44, 45, 46

[Dupuit, 1863] Dupuit, J. (1863). *Etudes theoriques et pratiques sur le mouvement des eaux dans les canaux decouverts et a travers les terrains permeables.* Dunod. 222

[Fowler, 2011] Fowler, A. (2011). Mathematical geoscience, vol. 36 of interdisciplinary applied mathematics. 44

[Fretwell et al., 2013] Fretwell, P., Pritchard, H. D., Vaughan, D. G., Bamber, J., Barrand, N., Bell, R., Bianchi, C., Bingham, R., Blankenship, D., Casassa, G., et al. (2013). Bedmap2: improved ice bed, surface and thickness datasets for Antarctica. *The Cryosphere*, 7(1):375–393. 11, 34, 37

[Gardner et al., 2013] Gardner, A. S., Moholdt, G., Cogley, J. G., Wouters, B., Arendt, A. A., Wahr, J., Berthier, E., Hock, R., Pfeffer, W. T., Kaser, G., et al. (2013). A reconciled estimate of glacier contributions to sea level rise: 2003 to 2009. *science*, 340(6134):852–857. 33

[Gelman et al., 2013] Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian data analysis.* CRC press. 143, 145

[Gerstner and Griebel, 1998] Gerstner, T. and Griebel, M. (1998). Numerical integration using sparse grids. *Numerical algorithms*, 18(3-4):209. 144

[Girolami and Calderhead, 2011] Girolami, M. and Calderhead, B. (2011). Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214. 148, 149, 207

[Goldberg et al., 2009] Goldberg, D., Holland, D., and Schoof, C. (2009). Grounding line movement and ice shelf buttressing in marine ice sheets. *Journal of Geophysical Research: Earth Surface*, 114(F4). 45

[Gomez et al., 2010] Gomez, N., Mitrovica, J. X., Huybers, P., and Clark, P. U. (2010). Sea level as a stabilizing factor for marine-ice-sheet grounding lines. *Nature Geoscience*, 3(12):850. 44

[Green et al., 2015] Green, P. J., Łatuszyński, K., Pereyra, M., and Robert, C. P. (2015). Bayesian computation: a summary of the current state, and samples backwards and forwards. *Statistics and Computing*, 25(4):835–862. 208

[Gudmundsson et al., 2012] Gudmundsson, G., Krug, J., Durand, G., Favier, L., and Gagliardini, O. (2012). The stability of grounding lines on retrograde slopes. *The Cryosphere*, 6:1497–1505. 45

[Haario et al., 2001] Haario, H., Saksman, E., and Tamminen, J. (2001). An adaptive metropolis algorithm. *Bernoulli*, pages 223–242. 147, 212

[Hanson and Hooke, 2000] Hanson, B. and Hooke, R. L. (2000). Glacier calving: a numerical model of forces in the calving-speed—water-depth relation. *Journal of Glaciology*, 46(153):188–196. 53

[Hastings, 1970] Hastings, W. K. (1970). Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109. 147

[Hindmarsh, 2012] Hindmarsh, R. (2012). An observationally validated theory of viscous flow dynamics at the ice shelf cavling front. *J. Glaciol.*, 58(208):to appear. 11, 37, 38, 45, 46, 132, 133

[Holmes, 2012] Holmes, M. H. (2012). *Introduction to perturbation methods*, volume 20. Springer Science & Business Media. 117

[Huan and Marzouk, 2013] Huan, X. and Marzouk, Y. M. (2013). Simulation-based optimal bayesian experimental design for nonlinear systems. *Journal of Computational Physics*, 232(1):288–317. 11, 37, 251

[Hughes, 1975] Hughes, T. (1975). The West Antarctic ice sheet: Instability, disintegration, and initiation of ice ages. *Reviews of Geophysics*, 13(4):502–526. 34

[Intergovernmental Panel on Climate Change, 2014] Intergovernmental Panel on Climate Change (2014). *Climate Change 2013: The Physical Science Basis: Working Group I Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press. 34

[Isaac et al., 2015] Isaac, T., Stadler, G., and Ghattas, O. (2015). Solution of nonlinear stokes equations discretized by high-order finite elements on nonconforming and anisotropic meshes, with application to ice sheet dynamics. *SIAM Journal on Scientific Computing*, 37(6):B804–B833. 11, 35, 36, 37

[Jamieson et al., 2012] Jamieson, S., Vieli, A., Livingstone, S., O'Cofaigh, C., Stokes, C., Hillenbrand, C.-D., and Dowdeswell, J. (2012). Ice-stream stability on a reverse bed slope. *Nature Geo*, 5(11):799–802. 11, 37, 38, 44, 45, 46

[Jaynes, 1957] Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical review*, 106(4):620. 35

[Jiménez et al., 2017] Jiménez, S., Duddu, R., and Bassis, J. (2017). An updated-lagrangian damage mechanics formulation for modeling the creeping flow and fracture of ice sheets. *Computer Methods in Applied Mechanics and Engineering*, 313:406–432. 53

[Joseph, 2012] Joseph, V. R. (2012). Bayesian computation using design of experiments-based interpolation technique. *Technometrics*, 54(3):209–225. 11, 37, 153, 206

[Joughin et al., 2014a] Joughin, I., Smith, B. E., and Medley, B. (2014a). Marine ice sheet collapse potentially under way for the thwaites glacier basin, west antarctica. *Science*, 344(6185):735–738. 33

[Joughin et al., 2014b] Joughin, I., Smith, B. E., and Medley, B. (2014b). Marine ice sheet collapse potentially under way for the Thwaites glacier basin, West Antarctica. *Science*, 344(6185):735–738. 34

[Kennedy and O'Hagan, 2001] Kennedy, M. C. and O'Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464. 11, 35, 37, 153, 206

[Kohler, 2002] Kohler, M. (2002). Universal consistency of local polynomial kernel regression estimates. *Annals of the Institute of Statistical Mathematics*, 54(4):879–899. 154, 155, 156, 157, 158, 161, 162

[Krug et al., 2014] Krug, J., Weiss, J., Gagliardini, O., and Durand, G. (2014). Combining damage and fracture mechanics to model calving. *The Cryosphere*, 8(6):2101–2117. 53

[Li and Marzouk, 2014] Li, J. and Marzouk, Y. M. (2014). Adaptive construction of surrogates for the Bayesian solution of inverse problems. *SIAM Journal on Scientific Computing*, 36(3):A1163–A1186. 206

[Li et al., 2013] Li, J., Paden, J., Leuschen, C., Rodriguez-Morales, F., Hale, R. D., Arnold, E. J., Crowe, R., Gomez-Garcia, D., and Gogineni, P. (2013). High-altitude radar measurements of ice thickness over the Antarctic and Greenland ice sheets as a part of Operation IceBridge. 11, 34, 37

[Lieberman et al., 2010] Lieberman, C., Willcox, K., and Ghattas, O. (2010). Parameter and state model reduction for large-scale statistical inverse problems. *SIAM Journal on Scientific Computing*, 32(5):2523–2542. 38

[Logg, 2007] Logg, A. (2007). Automating the finite element method. *Archives of Computational Methods in Engineering*, 14(2):93–138. 74, 96, 134

[Logg and Wells, 2010] Logg, A. and Wells, G. N. (2010). Dolfin: Automated finite element computing. *ACM Transactions on Mathematical Software*, 37(2). 74, 96, 134

[Ma et al., 2017] Ma, Y., Tripathy, C. S., and Bassis, J. N. (2017). Bounds on the calving cliff height of marine terminating glaciers. *Geophysical Research Letters*, 44(3):1369–1375. 11, 37, 50, 53, 55

[MacAyeal, 1993] MacAyeal, D. R. (1993). A tutorial on the use of control methods in ice-sheet modeling. *Journal of Glaciology*, 39(131):91–98. 11, 35, 36, 37

[Marin et al., 2012] Marin, J.-M., Pudlo, P., Robert, C. P., and Ryder, R. J. (2012). Approximate bayesian computational methods. *Statistics and Computing*, 22(6):1167–1180. 35, 38

[Martin et al., 2012] Martin, J., Wilcox, L. C., Burstedde, C., and Ghattas, O. (2012). A stochastic newton mcmc method for large-scale statistical inverse problems with application to seismic inversion. *SIAM Journal on Scientific Computing*, 34(3):A1460–A1487. 11, 35, 36, 37, 148

[Marzouk and Xiu, 2009] Marzouk, Y. and Xiu, D. (2009). A stochastic collocation approach to Bayesian inference in inverse problems. *Communications in Computational Physics*. 153, 206

[Marzouk et al., 2007] Marzouk, Y. M., Najm, H. N., and Rahn, L. A. (2007). Stochastic spectral methods for efficient bayesian solution of inverse problems. *Journal of Computational Physics*, 224(2):560–586. 153, 206

[Matott, 2012] Matott, L. S. (2012). Screening-Level Sensitivity Analysis for the Design of Pump-and-Treat Systems. *Ground Water Monitoring and Remediation*, 32(2):66–80. 220

[Meehl et al., 2005] Meehl, G. A., Washington, W. M., Collins, W. D., Arblaster, J. M., Hu, A., Buja, L. E., Strand, W. G., and Teng, H. (2005). How much more global warming and sea level rise? *Science*, 307(5716):1769–1772. 34

[Meier et al., 2007] Meier, M. F., Dyurgerov, M. B., Rick, U. K., O'neel, S., Pfeffer, W. T., Anderson, R. S., Anderson, S. P., and Glazovsky, A. F. (2007). Glaciers dominate eustatic sea-level rise in the 21st century. *Science*, 317(5841):1064–1067. 33

[Metropolis et al., 1953] Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092. 147

[Mobasher et al., 2016] Mobasher, M. E., Duddu, R., Bassis, J. N., and Waisman, H. (2016). Modeling hydraulic fracture of glaciers using continuum damage mechanics. *Journal of Glaciology*, 62(234):794–804. 53

[Moré and Sorensen, 1983] Moré, J. J. and Sorensen, D. C. (1983). Computing a trust region step. *SIAM Journal on Scientific and Statistical Computing*, 4(3):553–572. 167

[Nicholls and Cazenave, 2010] Nicholls, R. J. and Cazenave, A. (2010). Sea-level rise and its impact on coastal zones. *Science*, 328(5985):1517–1520. 34

[Nick et al., 2010] Nick, F., Van der Veen, C., Vieli, A., and Benn, D. (2010). A physically based calving model applied to marine outlet glaciers and implications for the glacier dynamics. *Journal of Glaciology*, 56(199):781–794. 11, 37, 45, 46, 50, 51, 53, 54, 238

[Ølgaard et al., 2008] Ølgaard, K. B., Logg, A., and Wells, G. N. (2008). Automated code generation for discontinuous galerkin methods. *SIAM Journal on Scientific Computing*, 31(2):849–864. 74, 96, 134

[Owen, 2013] Owen, A. B. (2013). Monte carlo theory, methods and examples. *Monte Carlo Theory, Methods and Examples. Art Owen.* 11, 37, 145, 146

[Pegler, 2016] Pegler, S. (2016). The dynamics of confined extensional flows. *J. Fluid Mech.*, 804:24–57. 11, 37, 38, 44, 45, 46, 132, 133

[Pegler et al., 2013] Pegler, S., Kowal, K., Hasenclever, L., and Worster, M. (2013). Lateral controls on grounding-line dynamics. *J. Fluid Mech.*, 722:R1. 11, 37, 38, 45, 46

[Pegler and Worster, 2012] Pegler, S. S. and Worster, M. G. (2012). Dynamics of a viscous layer flowing radially over an inviscid ocean. *Journal of Fluid Mechanics*, 696:152–174. 44

[Petra et al., 2014] Petra, N., Martin, J., Stadler, G., and Ghattas, O. (2014). A computational framework for infinite-dimensional bayesian inverse problems, part ii: stochastic newton mcmc with application to ice sheet flow inverse problems. *SIAM Journal on Scientific Computing*, 36(4):A1525–A1555. 11, 35, 36, 37

[Petra et al., 2012] Petra, N., Zhu, H., Stadler, G., Hughes, T. J., and Ghattas, O. (2012). An inexact gauss–newton method for inversion of basal sliding and rheology parameters in a nonlinear stokes ice sheet model. *Journal of Glaciology*, 58(211):889–903. 11, 35, 36, 37

[Pfeffer et al., 2008] Pfeffer, W. T., Harper, J., and O'Neel, S. (2008). Kinematic constraints on glacier contributions to 21st-century sea-level rise. *Science*, 321(5894):1340–1343. 33

[Pollard and DeConto, 2009] Pollard, D. and DeConto, R. M. (2009). Modelling west antarctic ice sheet growth and collapse through the past five million years. *Nature*, 458(7236):329–332. 34

[Pollard et al., 2015] Pollard, D., DeConto, R. M., and Alley, R. B. (2015). Potential antarctic ice sheet retreat driven by hydrofracturing and ice cliff failure. *Earth and Planetary Science Letters*, 412:112–121. 11, 37, 50, 55

[Rasmussen, 2003] Rasmussen, C. E. (2003). Gaussian Processes to Speed up Hybrid Monte Carlo for Expensive Bayesian Integrals. In *Bayesian Statistics 7*, pages 651–659. Oxford University Press. 206

[Raymond, 1996] Raymond, C. (1996). Shear margins in glaciers and ice sheets. *Journal of Glaciology*, 42(140):90–102. 59

[Rignot et al., 2014] Rignot, E., Mouginot, J., Morlighem, M., Seroussi, H., and Scheuchl, B. (2014). Widespread, rapid grounding line retreat of Pine Island, Thwaites, Smith, and Kohler glaciers, West Antarctica, from 1992 to 2011. *Geophysical Research Letters*. 34

[Rignot and Thomas, 2002] Rignot, E. and Thomas, R. H. (2002). Mass balance of polar ice sheets. *Science*, 297(5586):1502–1506. 34

[Rignot et al., 2011] Rignot, E., Velicogna, I., van den Broeke, M. R., Monaghan, A., and Lenaerts, J. T. (2011). Acceleration of the contribution of the greenland and antarctic ice sheets to sea level rise. *Geophysical Research Letters*, 38(5). 33

[Robel et al., 2014] Robel, A., Schoof, C., and Tziperman, E. (2014). Rapid grounding line migration induced by internal ice stream variability. *Journal of Geophysical Research: Earth Surface*, 119(11):2430–2447. 44, 45

[Robel et al., 2016] Robel, A., Schoof, C., and Tziperman, E. (2016). Persistence and variability of ice-stream grounding lines on retrograde bed slopes. *The Cryosphere*, 10(4):1883. 44, 45

[Roberts and Rosenthal, 1998] Roberts, G. O. and Rosenthal, J. S. (1998). Optimal scaling of discrete approximations to langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(1):255–268. 148, 228

[Rognes and Logg, 2013] Rognes, M. E. and Logg, A. (2013). Automated goal-oriented error control i: Stationary variational problems. *SIAM Journal on Scientific Computing*, 35(3):C173–C193. 79

[Rosenthal, 2000] Rosenthal, J. S. (2000). Parallel computing and Monte Carlo algorithms. *Far east journal of theoretical statistics*, 4(2):207–236. 208

[Rote and Tichy, 1996] Rote, G. and Tichy, R. F. (1996). Quasi-monte-carlo methods and the dispersion of point sequences. *Mathematical and computer modelling*, 23(8-9):9–23. 165

[Sacks et al., 1989] Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P. (1989). Design and analysis of computer experiments. *Statistical science*, pages 409–423. 11, 37, 153, 206

[Saltelli, 2002] Saltelli, A. (2002). Making best use of model evaluations to compute sensitivity indices. *Computer physics communications*, 145(2):280–297. 241

[Santner et al., 2003] Santner, T. J., Williams, B. J., and Notz, W. (2003). *The design and analysis of computer experiments*. Springer. 11, 37, 153, 206

[Schoof, 2007a] Schoof, C. (2007a). Ice sheet grounding line dynamics: Steady states, stability, and hysteresis. *Journal of Geophysical Research: Earth Surface*, 112(F3). 35, 44, 47, 58, 83, 94, 96, 109, 119, 121, 124, 129, 132, 237

[Schoof, 2007b] Schoof, C. (2007b). Marine ice sheet dynamics. Part I. The case of rapid sliding. *J. Fluid Mech.*, 573:27–55. 35, 44, 47, 58, 83, 94, 96, 109, 119, 121, 124, 129, 132, 237

[Schoof, 2012] Schoof, C. (2012). Marine ice sheet stability. *Journal of Fluid Mechanics*, 698:62–72. 44

[Schoof et al., 2017] Schoof, C., Davis, A. D., and Popa, T. V. (2017). Boundary layer models for calving marine outlet glaciers. *The Cryosphere*, 11(5):2283. 38, 45, 46, 51, 93, 94

[Schutz et al., 2005] Schutz, B., Zwally, H., Shuman, C., Hancock, D., and DiMarzio, J. (2005). Overview of the ICESat mission. *Geophysical Research Letters*, 32(21). 11, 34, 37

[Sherlock et al., 2017] Sherlock, C., Golightly, A., and Henderson, D. A. (2017). Adaptive, delayed-acceptance mcmc for targets with expensive likelihoods. *Journal of Computational and Graphical Statistics*, 26(2):434–444. 38

[Sijbrand, 1985] Sijbrand, J. (1985). Properties of center manifolds. *Transactions of the American Mathematical Society*, 289(2):431–469. 104

[Smolyak, 1963] Smolyak, S. (1963). Quadrature and interpolation formulas for tensor products of certain classes of functions. In *Soviet Math. Dokl.*, volume 4, pages 240–243. 11, 37, 144, 153

[Sobol, 2001] Sobol, I. M. (2001). Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates. *Mathematics and computers in simulation*, 55(1-3):271–280. 241

[Spantini et al., 2017] Spantini, A., Cui, T., Willcox, K., Tenorio, L., and Marzouk, Y. (2017). Goal-oriented optimal approximations of bayesian linear inverse problems. *SIAM Journal on Scientific Computing*, 39(5):S167–S196. 38

[Stone, 1977] Stone, C. J. (1977). Consistent nonparametric regression. *The annals of statistics*, pages 595–620. 154, 158, 161, 162

[Suckale et al., 2014] Suckale, J., Platt, J. D., Perol, T., and Rice, J. R. (2014). Deformation-induced melting in the margins of the West Antarctic ice streams. *Journal of Geophysical Research: Earth Surface*. 38

[Thomas et al., 2004] Thomas, R., Rignot, E., Casassa, G., Kanagaratnam, P., Acuña, C., Akins, T., Brecher, H., Frederick, E., Gogineni, P., Krabill, W., et al. (2004). Accelerated sea-level rise from west antarctica. *Science*, 306(5694):255–258. 34

[Thomas and Bentley, 1978] Thomas, R. H. and Bentley, C. R. (1978). A model for holocene retreat of the west antarctic ice sheet. *Quaternary Research*, 10(2):150–170. 44, 237

[Truffer and Echelmeyer, 2003] Truffer, M. and Echelmeyer, K. A. (2003). Of isbrae and ice streams. *Annals of Glaciology*, 36:66–72. 43

[Van der Veen, 1998a] Van der Veen, C. (1998a). Fracture mechanics approach to penetration of bottom crevasses on glaciers. *Cold Regions Science and Technology*, 27(3):213–223. 11, 37, 50, 53

[Van der Veen, 1998b] Van der Veen, C. (1998b). Fracture mechanics approach to penetration of surface crevasses on glaciers. *Cold Regions Science and Technology*, 27(1):31–47. 11, 37, 50, 53

[Vihola et al., 2016] Vihola, M., Helske, J., and Franks, J. (2016). Importance sampling type correction of markov chain monte carlo and exact approximations. *arXiv preprint arXiv:1609.02541*. 153

[Wearing et al., 2015] Wearing, M. G., Hindmarsh, R. C., and Worster, M. G. (2015). Assessment of ice flow dynamics in the zone close to the calving front of antarctic ice shelves. *Journal of Glaciology*, 61(230):1194–1206. 47

[Weertman, 1973] Weertman, J. (1973). Can a water-filled crevasse reach the bottom surface of a glacier. *IASH publ*, 95:139–145. 11, 37, 50, 53

[Weertman, 1974] Weertman, J. (1974). Stability of the junction of an ice sheet and an ice shelf. *Journal of Glaciology*, 13(67):3–11. 44, 237

[Weertman, 1980] Weertman, J. (1980). Bottom crevasses. *Journal of Glaciology*, 25(91):185–188. 11, 37, 50, 53

[Wiener, 1938] Wiener, N. (1938). The homogeneous chaos. *American Journal of Mathematics*, pages 897–936. 11, 37, 153

[Williams Jr. and Ferrigno, 2013] Williams Jr., R. S. and Ferrigno, J. G. (2013). Image atlas of glaciers of the world. State of the Earth's cryosphere at the beginning of the 21st century: glaciers, global snow cover, floating ice, and permafrost and periglacial environments. (USGS professional paper 1386). *Journal of Glaciology*, 59(216):795. 34

[Wolff, 2004] Wolff, U. (2004). Monte Carlo errors with less errors . *Computer Physics Communications*, 156(2):143–153. 213

[Yu et al., 2017] Yu, H., Rignot, E., Morlighem, M., and Seroussi, H. (2017). Iceberg calving of thwaites glacier, west antarctica: full-stokes modeling combined with linear elastic fracture mechanics. *The Cryosphere*, 11(3):1283. 53

[Zahm et al., 2018] Zahm, O., Constantine, P., Prieur, C., and Marzouk, Y. (2018). Gradient-based dimension reduction of multivariate vector-valued functions. *arXiv preprint arXiv:1801.07922*. 11, 37, 38