Predicting Snow Velocity in Large Chute Flows Under Different Environmental Conditions

Jonathan Rougier*
Department of Mathematics
University of Bristol, UK

Martin Kern
Swiss Federal Institute for Snow and Avalanche Research
Davos, Switzerland

April 26, 2008

Abstract

Experimental observations, model evaluations, and expert judgements are combined to make predictions of snow velocity in large chute experiments. Different experimental treatments, namely the environmental conditions snow density and snow-surface temperature, affect all aspects of this inference. We show how treatments can be incorporated into our judgements regarding the uncertain parameters of the physical model, the discrepancy between the physical model and reality, and the observation error. We adopt a Bayes linear approach to avoid the necessity of fully-probabilistic belief specifications, and demonstrate visual tools for statistical validation. Our results represent an important first step in improving the specification of uncertainty in model-based avalanche hazard mapping.

KEYWORDS: Bayes linear, computer experiment, Herschel-Bulkley rheology, model discrepancy, multiple treatments

*Department of Mathematics, University Walk, Bristol BS8 1TW, U.K. Email j.c.rougier@bristol.ac.uk.

1 Introduction

It is very important to understand the behaviour of avalanches: how, for example, the speed of an avalanche depends upon the inclination of the slope, on the snow density, and on the snow-surface temperature. One source of data is from avalanches themselves but, for obvious reasons, these opportunistic observations are hard to come by. Experiments provide a second source of data. In this paper we focus on chute experiments carried out at the Swiss Federal Institute for Snow and Avalanche Research (SLF).

These experiments take place under different environmental conditions, partly by design and partly through circumstance. We would like to extrapolate from the experiments we have, to learn about snow velocities across a range of plausible environmental conditions. For this purpose we introduce a mathematical model which expresses the velocity profile as a function of these conditions. The general idea is to use the mathematical model to construct a joint probability distribution over the experiments we have and those environmental conditions we would like to predict, and then the observations are assimilated into the prediction by probabilistic conditioning.

The statistical field of Computer Experiments is concerned with combining model evaluations and observations. A particular challenge in this field is to account for the fact that some of the model parameters are imperfectly known, and that the model itself is imperfect (Kennedy and O'Hagan, 2001; Craig et al., 2001; Goldstein and Rougier, 2004, 2008). This challenge becomes more acute when the model-outputs and the system behaviour are multivariate. In our application, for example, the model output is functional, and the system is observed at a discrete set of abscissae (heights). But our application also introduces a further complication: variations in the environmental conditions. Thinking of these as 'treatments', the methodological contribution of this paper is to show how treatments can be included in a computer experiment, taking account of the fact that the treatment can affect all aspects of the statistical model that links the model parameters, the model evaluations, the system, and the system observations. A second statistical contribution is to demonstrate a detailed elicitation, including validation, for a complex physical process.

The outline of the paper is as follows. Section 2 describes the background to the experiments, and the mathematical model we adopt. Sec-

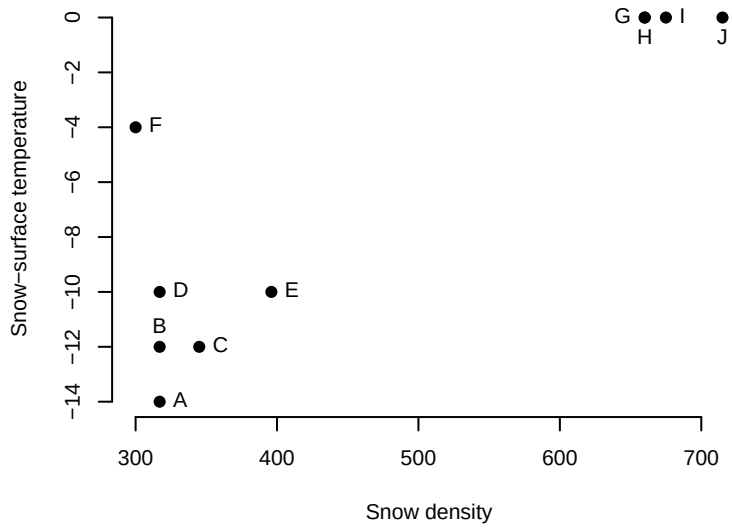

Figure 1: Environmental conditions in the ten experiments (labelled $A, \ldots, J$), indexed by snow density (kg/m^3) and snow-surface temperature (°C).

tion 3 describes the statistical framework that we use to combine model evaluations and observations, over a range of different environmental conditions. Section 4 describes our choice of inferential treatment, the Bayes linear approach, and the simplifications that follow from it. Sections 5 and 6 describe our statistical modelling choices, and the results of our analysis, including diagnostic assessment. Section 7 concludes.

2 The experiments and the physical model

In this section we present a summary of the chute and the experiments, and describe the notational modifications we have made to the physical model. A more detailed account can be found in Kern et al. (2004), and the references below.

The chute and the experiments. Our experiments were performed on the SLF snow chute at Weissfluhjoch, which is 34 m long and 2.5 m wide. In each experiment, 8.4 m^3 of snow was released from a hopper at the top of the chute. The snow then accelerated along a 10 m section, at the end of which it was approximately in steady state. Downslope velocities of the

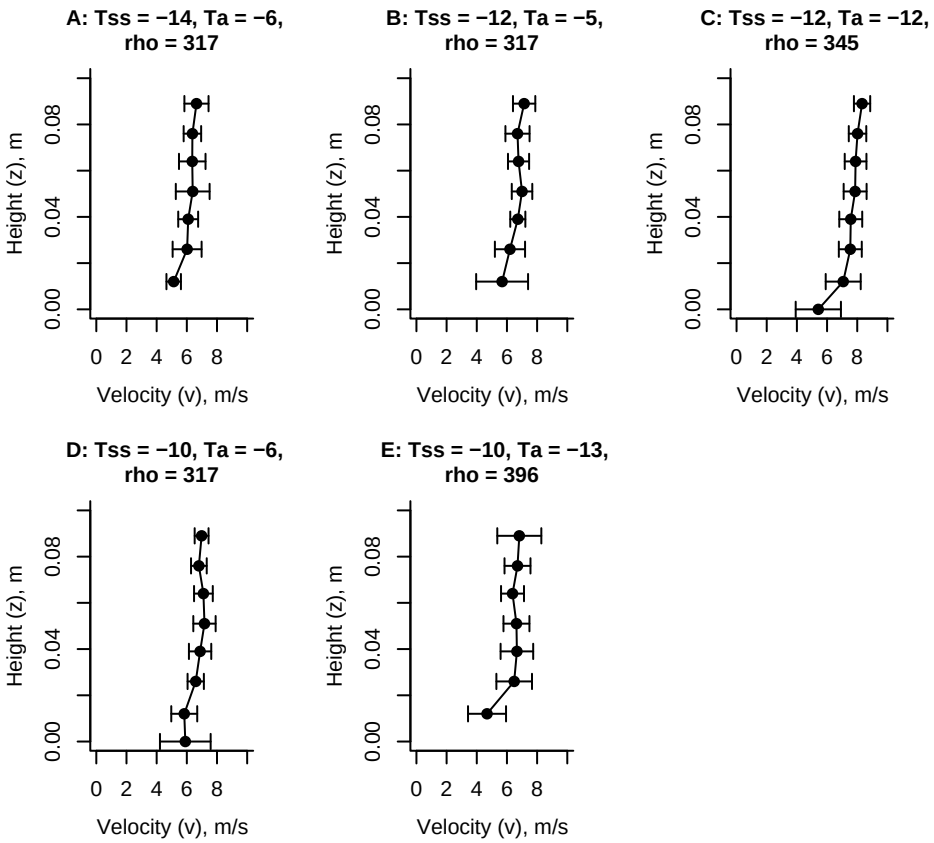

Figure 2: Those experiments with low snow-surface temperatures. Tss: snow-surface temperature (°C); Ta: atmospheric temperature (°C); rho: snow density (kg/m^3). The dots indicate the measurements and the error bars ±2 standard deviations; the measurements have been interpolated to make the velocity profile easier to see.

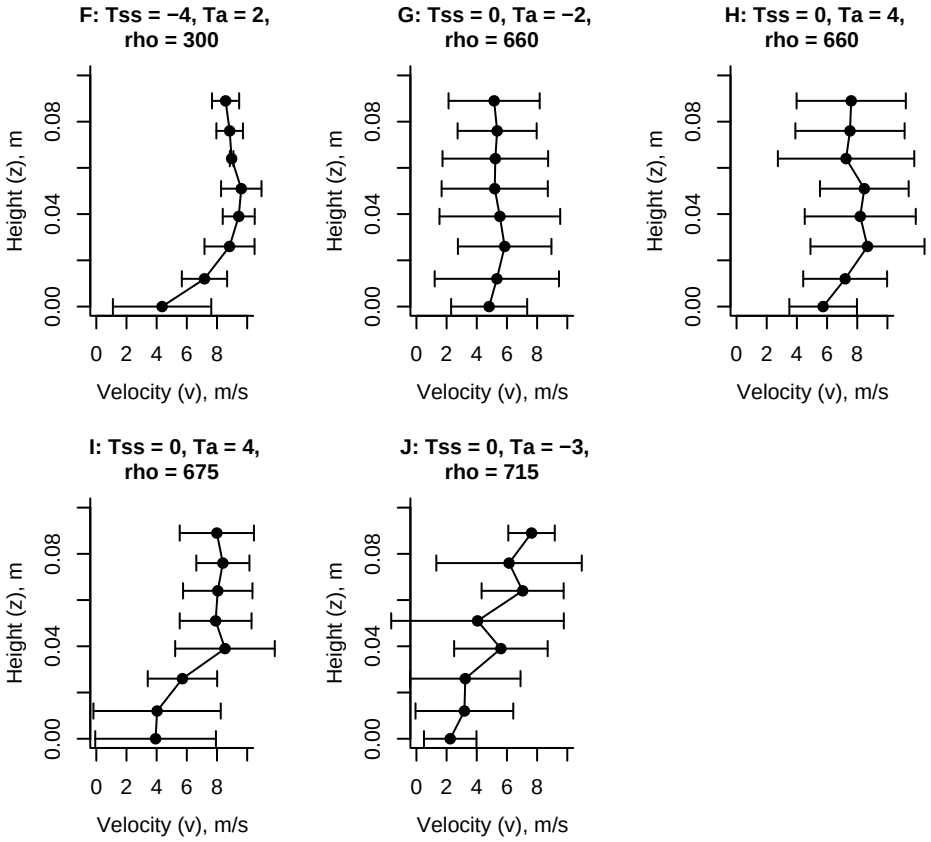

Figure 3: Those experiments with high snow-surface temperatures. See the caption to Figure 2 for details.

flow were measured by an optical velocity sensor array placed in a half wedge at the centreline of the chute. The velocity sensors were placed at 8 equidistant surface-normal elevations in the range between 0 and 0.089 m. From the time-series of the velocity sensors, time-averaged velocity profiles were extracted, along with a measure of uncertainty.

For our analysis we use ten experiments which were performed under different environmental conditions, indexed by the snow density ρ, the ambient air temperature T_{a} and the snow-surface temperature T_{ss}. For the snow and air temperatures, we use measurements recorded by an automatic weather station at the nearby Weisfluhjoch experimental site. The environmental conditions for the ten experiments are summarised in Figure 1, and the observations are given in Figures 2 and 3.

For a detailed description of the measurement setup, the working principle of the optical velocity sensors, and a discussion of the systematic measurement errors, see Tiefenbacher and Kern (2004), Kern et al. (2004), and McElwaine and Tiefenbacher (2005).

The Herschel-Bulkley model. Observations on flowing avalanches (Dent et al., 1997; Sovilla et al., 2008) and from chute experiments suggest that the flow of snow is characterised by a relatively thick layer where the shear rates are low or even zero. This so-called plug-layer travels on a comparatively thin shear-layer. Basal sliding may occur in this shear-layer, depending on the environmental conditions. The simplest rheology which is able to reproduce such a combination of solid- and fluid-like behaviour is the Bingham rheology (Bingham, 1922; Oldroyd, 1947). However, Kern et al. (2004) show that qualitatively better fits can be obtained using a generalised Bingham rheology, the so-called Herschel-Bulkley model (see, e.g., Barnes et al., 1989).

The equations of the Herschel-Bulkley model, adapted for our purposes, are presented in Figure 4, with a simple schematic of the main features in Figure 5. It is described in detail in Kern et al. (2004); here we outline the adaptation we have made relative to this description. Our objective is to identify a set of model parameters about which we are uncertain, and for which we can specify a marginal distribution. The uncertain parameters in the original model are: v_0, denoting basal slip velocity (m); τ_c, the stress at zero shear (Pa); α, a unitless stress coefficient; and K, which enters into the

stress equation for positive shear:

$$\text{stress} = \tau_c + K\left(\frac{\mathrm{d}v}{\mathrm{d}z}\right)^{\alpha} \qquad \frac{\mathrm{d}v}{\mathrm{d}z} > 0, \tag{1}$$

where $v(z)$ is the velocity profile at height z. K is a complicated quantity which has units that depend on α. This makes it hard to elicit K: it would be simpler if we could decouple K from α, in such a way that we would be comfortable treating these two quantities as probabilistically independent. Therefore we reparameterise the stress relationship as

$$\text{stress} = \tau_c\left[1 + \left(t_c\,\frac{\mathrm{d}v}{\mathrm{d}z}\right)^{\alpha}\right] \qquad \frac{\mathrm{d}v}{\mathrm{d}z} > 0, \tag{2}$$

where t_c has units of time (seconds). According to (2), $t_c = (K/\tau_c)^{1/\alpha}$, and t_c replaces K as our fourth uncertain model parameter.

3 Outline of the statistical inference

3.1 General features

The experimental conditions at the point where the experiment is conducted are referred to as the *treatment*. We use 'treatment' to denote both the formal variables in the experiment, and also the concomicant variables that can be measured prior to the experiment. Partly this is for simplicity, but it also reflects the lack of a clear-cut distinction between the two. Some treatment variables are under the control of the experimenter, such as the chute angle. Others are determined by the environment, such as the temperature and the temperature history. Others fall somewhere in-between these two, such as the snow density. What identifies the treatment is that it is known at the point where the experiment is conducted, and it must be specified in order for the outcome of any experiment to be predicted. We denote the treatment as $\boldsymbol{m}$; in our case

$$\boldsymbol{m} = (\rho, T_{\mathrm{ss}})\,,$$

where ρ is the snow density and T_{ss} is snow-surface temperature; we also have information atmospheric temperature but we have not used it in this analysis. We might also consider θ, the inclination, to be a treatment vari-

The Herschel-Bulkley Model

The velocity profile has the form:

$$v(z) = \begin{cases} v_h + (v_0 - v_h)\left(1 - \dfrac{z}{h}\right)^{\frac{1+\alpha}{\alpha}} & 0 \le z < h \\ v_h & z \ge h, \end{cases} \tag{HB-a}$$

where

$$v_h = v_0 + \frac{h}{t_c}\frac{\alpha}{1+\alpha}\left(\frac{h}{H-h}\right)^{1/\alpha} \tag{HB-b}$$

and h solves

$$\tau_c = (H-h)g\rho\sin\theta \tag{HB-c}$$

subject to $h < H$.

Notation

z	Height ordinate (m)	$v(z)$	Velocity (m/s)
h	Height to plug-layer (m)	v_h	plug-layer velocity
	Treated as known		
θ	Inclination (32°)	g	Acceleration ($9.8\,\mathrm{m/s^2}$)
H	Height to top of flow (0.4 m)		
	Environmental variables		
ρ	Snow density ($\mathrm{kg/m^3}$)	T_{ss}	Snow-surface temperature (°C)
	Uncertain model parameters		
v_0	Basal velocity (m/s)	α	Stress coefficient
τ_c	Stress (Pa)	t_c	Time constant (s)

(the variable T_{ss} does not appear explicitly in the model).

Figure 4: Velocity profile of a steady 2D Herschel-Bulkley flow of snow. Our alterations to the standard notation are described in Section 2.

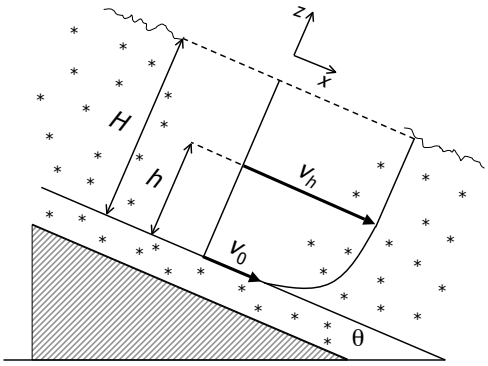

Figure 5: A simple schematic of a typical velocity profile of steady Herschel-Bulkley flow.

able, but it does not vary in our experiments.

Our prediction for the velocity profile under any treatment takes the form of a function relating velocity to height. In this case height is termed an *index variable*, because it indexes the model output. Denoting height as z, our object is to predict $v(\boldsymbol{m}, z)$, the velocity of the snow at height z, in a large chute under treatment $\boldsymbol{m}$. We have three sources of information. First, we have the outcome of experiments on large chutes, for a variety of treatments. These experimental results include measurement error, and are distinguished from the true results by writing $v^{\text{obs}}(\boldsymbol{m}, z)$. We have ten experiments, and for simplicity we will suppose that we record velocities at the same eight heights for each experiment—in fact not all the heights were recorded for every experiment: we make this simplification purely to avoid an extra layer of subscripts. Therefore our observations comprise the matrix

$$V^{\text{obs}} = \begin{pmatrix} v^{\text{obs}}(\boldsymbol{m}_1, z_1) & \dots & v^{\text{obs}}(\boldsymbol{m}_1, z_8) \\ \vdots & \ddots & \vdots \\ v^{\text{obs}}(\boldsymbol{m}_{10}, z_1) & \dots & v^{\text{obs}}(\boldsymbol{m}_{10}, z_8) \end{pmatrix} \tag{3}$$

a matrix in which the rows correspond to the treatments and the columns to the heights of the index variable.

Our second source of information is evaluations of a physical model, namely the Herschel-Bulkley (HB) model described in Figure 4. This model can be seen as a function mapping $(\boldsymbol{m}, z)$ into a scalar output. But there will be a further set of inputs, namely those model parameters about which we

are uncertain. This uncertainty has two sources. First, the model may contain empirical relationships that stand in for physics we do not understand, or which we choose not to represent fully. These empirical relationships may have uncertain coefficients. Second, deficiencies in the model compromise the interpretation of the model's parameters, even those with well-defined physical meanings. For example, the value of shear stress in actual snow in the chute may not be the best value to use for shear stress in the model. Therefore, although we are guided by the physical interpretation of the parameters, we do not necessarily want to fix them at their physical values. These uncertain parameters are denoted $\boldsymbol{x}$,

$$\boldsymbol{x} = (v_0, \tau_c, \alpha, t_c)\,.$$

The model output is then denoted $g(\boldsymbol{x}, \boldsymbol{m}, z)$.

Our third source of information is our judgements about the physical model, the actual behaviour of snow in the chute, the observations, and the relationships between them. Specifying these judgements occurs in two stages: first we construct a joint statistical model over all uncertain and observed quantities, describing our conditional independence choices. Then we quantify the marginal and conditional distributions that occur in this statistical model. Our choices will be informed by the physics of snow in large chutes, by the feasibility of the elicitation, and by the tractability of the resulting statistical inference. In this respect constructing a statistical model is no different from the process of constructing a physical model.

3.2 Statistical framework

For clarity, we start by describing a statistical framework for a single treatment and a single height, dropping $\boldsymbol{m}$ and z from the notation, and writing $g(\boldsymbol{x})$ for the physical model, v for the actual value, and v^{obs} for the observation. The standard approach to constructing a joint statistical framework is to assert the existence of a 'best' value of the model parameters, denoted $\boldsymbol{x}^*$ (Goldstein and Rougier, 2006; Rougier, 2007). Then we link the evaluations of the physical model and the actual system behaviour through the model evaluated at this 'best' input:

$$\boldsymbol{x}^* \xrightarrow{g} v \longrightarrow v^{\text{obs}} \tag{4}$$

where the only simplification we have made in the joint structure is to choose $v^{\text{obs}} \perp\!\!\!\perp \boldsymbol{x}^* \mid v$, a completely standard and uncontroversial choice. The superscript on the edge from $\boldsymbol{x}^*$ to v indicates the 'location' of the HB model in the inference.

The first edge represents our statistical model of the *discrepancy* between the model and reality: a typical form of this conditional distribution might be

$$\pi(v \mid x^*) = \varphi\big(v; g(\boldsymbol{x}^*), \sigma^2(\boldsymbol{x}^*)\big) \tag{5}$$

where φ is the Gaussian density function, and we specify the variance $\sigma^2(\boldsymbol{x})$ as an explicit function of the model parameters. A common simplification in (5) is to make $\sigma^2(\boldsymbol{x})$ invariant to $\boldsymbol{x}$, i.e. set $\sigma^2(\boldsymbol{x}) = \sigma^2$, a scalar. A regrettable further simplification is to set $\sigma^2 = 0$, which asserts that the model has no structural error, and that it is only uncertainty about $\boldsymbol{x}^*$ that prevents us from performing a perfect evaluation. This, unfortunately, has been the dominant practice in much applied science, where practitioners have either not been aware that it is possible to incorporate structural error into their analysis, or have been reluctant to quantify it. This reluctance to quantify can also be seen in the choice of marginal distribution $\pi(\boldsymbol{x}^*)$, which is often taken to be rectangular with specified limits. This class of distribution is supposed, mistakenly, to be the 'neutral' choice.

The second edge represents our statistical model of the measurement processes: a typical form for this distribution might be

$$\pi(v^{\text{obs}} \mid v) = \varphi\big(v^{\text{obs}}; v, \tau^2\big) \tag{6}$$

where for simplicity we treat measurement error as invariant to flow speed, and specify the observation error in terms of a standard deviation τ.

Multiple heights. Generalising from one value of the index variable height to a collection of heights presents no conceptual problems. Write (4) as

$$\boldsymbol{x}^* \xrightarrow{\ g\ } \boldsymbol{v} \longrightarrow \boldsymbol{v}^{\text{obs}} \tag{7}$$

where $\boldsymbol{v} = \big(v(z_1), \ldots v(z_8)\big)$ and $\boldsymbol{v}^{\text{obs}} = \big(v^{\text{obs}}(z_1), \ldots, v^{\text{obs}}(z_8)\big)$. The scalar variance function $\sigma^2(\boldsymbol{x})$ is replaced by a more general relationship such as

$$\text{Cov}\big(v(z_j), v(z_{j'}) \mid \boldsymbol{x}^*\big) = \sigma(\boldsymbol{x}^*, z_j) \times \sigma(\boldsymbol{x}^*, z_{j'}) \times \kappa_z(z_j, z_{j'}), \tag{8}$$

where $\kappa_z(\cdot)$ is a correlation function, and for simplicity we treat the correlation structure as invariant to $\boldsymbol{x}^*$. We can use $\sigma(\boldsymbol{x}, z)$ to describe our judgements of how the physical model's performance varies according to both $\boldsymbol{x}$ and z. For example, the HB velocity profile comprises two parts, one for heights below the shear layer height h, and a much simpler (vertical) part for heights equal to or above h, where h is a known function of $\boldsymbol{x}$ (see Figure 4). If we judge the simpler model more likely to be in error, then our $\sigma(\boldsymbol{x}, z)$ will be a function of both $\boldsymbol{x}$ and z.

In a similar way, the scalar measurement error variance τ^2 is replaced by the matrix T. We need to include common sources of variation in the off-diagonal elements of T. The dominant source of these is measurements made by the same instrument, which might have a bias. In our large chute, though, there is a different instrument at each height, so this is not an issue. However, the instruments are identical, and so there might be a common source of variation from a fault that is particular to instruments of this general type.

3.3 Including treatments

The purpose of including treatments in our framework and in our experiments is to allow us to predict the behaviour of snow in a chute under environmental conditions we have not observed. For this to be possible, we must believe that there is some relationship between $\boldsymbol{v}(\boldsymbol{m})$ and $\boldsymbol{v}(\boldsymbol{m}')$ where $\boldsymbol{m}'$ is not dissimilar to $\boldsymbol{m}$, so that an experiment under treatment $\boldsymbol{m}$ is informative about what happens under $\boldsymbol{m}'$.

Explicit single treatment. We start by including the treatment explicitly in our statistical model for a single treatment:

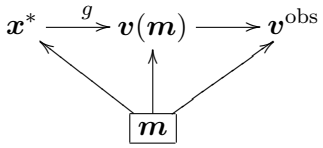

(9)

where the box around $\boldsymbol{m}$ indicates that it is specified, not uncertain. This indicates that our judgements about $\boldsymbol{x}^*$, about our physical model's discrepancy, and about the observation error might all depend on the treatment.

When we introduce multiple treatments, though, we are faced with a

problem. We would like to abstract the best value of the model parameters from the treatment, so that this best value becomes a repository of information about the physical model that can be informed by a range of experiments at different treatments, and can be used to predict the behaviour of snow in the chute at new treatments. Standing in the way of this is the edge from $\boldsymbol{m}$ to $\boldsymbol{x}^*$. The solution is to propose a higher level of treatment-invariant model parameters. Formally, we suppose that we can specify an uncertain vector $\boldsymbol{w}^*$ with marginal density $\pi(\boldsymbol{w}^*)$, and a vector-valued deterministic function $\boldsymbol{x} = \boldsymbol{k}(\boldsymbol{w}, \boldsymbol{m})$, such that $\boldsymbol{k}(\boldsymbol{w}^*, \boldsymbol{m})$ has the distribution $\pi(\boldsymbol{x}^*; \boldsymbol{m})$, for all $\boldsymbol{m}$. This gives

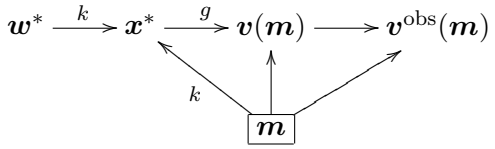

(10)

Therefore learning about the model parameters is learning about $\boldsymbol{w}^*$, from which $\boldsymbol{x}^*$ is then inferred for particular $\boldsymbol{m}$. We will illustrate judgements of this form in section 5.1, for the two model parameters v_0 and τ_c, which will depend on ρ and T_{ss}, respectively.

Multiple treatments. In order to handle multiple treatments we group the actual values for snow in the chute together over the set of treatments we have observed, plus the treatments for which we want to make predictions. Denote these latter treatments as $\boldsymbol{m}_a, \boldsymbol{m}_b, \ldots$. We write the collection as

$$V = \left(\begin{array}{ccc} v(\boldsymbol{m}_1, z_1) & \ldots & v(\boldsymbol{m}_1, z_8) \\ \vdots & \ddots & \vdots \\ v(\boldsymbol{m}_{10}, z_1) & \ldots & v(\boldsymbol{m}_{10}, z_8) \\ \hline v(\boldsymbol{m}_a, z_1) & \ldots & v(\boldsymbol{m}_a, z_8) \\ v(\boldsymbol{m}_b, z_1) & \ldots & v(\boldsymbol{m}_b, z_8) \\ \vdots & \ddots & \vdots \end{array}\right). \tag{11}$$

In this format, the statistical model over multiple treatments can be written

$$\boldsymbol{w}^* \xrightarrow{g,k} V \longrightarrow V^{\mathrm{obs}} \tag{12}$$

where the treatments are now internalised in V and V^{obs}. Looking back to (4), we seem to have come full-circle, with the important difference that we have had to redefine the model-parameters to ensure that they can be expressed independently of the treatments.

With this modification, the statistical modelling of the two conditional distributions can be extended quite naturally from the single-treatment. For the discrepancy we might take the collection $\{v(\boldsymbol{m}_i, z_j) \mid \boldsymbol{w}^*\}$ to be jointly Gaussian with mean $g\big(\boldsymbol{k}(\boldsymbol{w}^*, \boldsymbol{m}_i), \boldsymbol{m}_i, z_j\big)$ and variance function

$$\begin{aligned}\operatorname{Cov}\big(v(\boldsymbol{m}_i, z_j), v(\boldsymbol{m}_{i'}, z_{j'}) \mid \boldsymbol{w}^*\big) = \\ \sigma\big(\boldsymbol{k}(\boldsymbol{w}^*, \boldsymbol{m}_i), \boldsymbol{m}_i, z_j\big) \times \sigma\big(\boldsymbol{k}(\boldsymbol{w}^*, \boldsymbol{m}_{i'}), \boldsymbol{m}_{i'}, z_{j'}\big) \times \\ \kappa_m(\boldsymbol{m}_i, \boldsymbol{m}_{i'}) \times \kappa_z(z_j, z_{j'})\end{aligned} \tag{13}$$

where $\kappa_m(\cdot)$ and $\kappa_z(\cdot)$ are correlation functions, and for simplicity we take the joint correlation structure as invariant to $\boldsymbol{x}^*$ and separable in the treatment variables and the index variable.

For the observation error, we might take the collection $\{v^{\text{obs}}(\boldsymbol{m}_i, z_j) \mid V\}$ to be jointly Gaussian, and simplify by treating $\boldsymbol{v}(\boldsymbol{m}_i)$ as sufficient for $\boldsymbol{v}^{\text{obs}}(\boldsymbol{m}_i)$ for each treatment, with mean $\boldsymbol{v}(\boldsymbol{m}_i)$, and variance function

$$\operatorname{Cov}\big(v^{\text{obs}}(\boldsymbol{m}_i, z_j), v^{\text{obs}}(\boldsymbol{m}_{i'}, z_{j'}) \mid V\big) = \begin{cases} T_{jj'} & i = i' \\ 0 & \text{otherwise.} \end{cases} \tag{14}$$

We summarise the statistical requirements described in this section, under the simplifications we have made. Note that the simplifications are not critical: we have made them to clarify the presentation, and also because they seem reasonable in our application and, perhaps, more widely. More importantly, the Gaussian form of the conditional distributions is not critical either. In fact, in the following sections we will dispense with it, and proceed using a Bayes Linear approach.

1. We specify the determinstic function $\boldsymbol{k}(\cdot)$ and the marginal distribution $\boldsymbol{w}^*$ so that the inferred marginal distribution of $\boldsymbol{x}^* = \boldsymbol{k}(\boldsymbol{w}^*, \boldsymbol{m})$ is consistent with our judgements about $\boldsymbol{x}^*$ in treatment $\boldsymbol{m}$.

2. We quantify our judgement about the discrepancy in our physical model in terms of the standard deviation function $\sigma(\boldsymbol{x}, \boldsymbol{m}, z)$, and the two correlation functions $\kappa_m(\boldsymbol{m}, \boldsymbol{m}')$ and $\kappa_z(z, z')$; see eq. (13).

3. We quantify our judgements about the measurement errors in our observations in terms of the variance matrix T; see eq. (14).

4 Bayes linear inference

Our intention is to predict velocities under treatments we have not observed. To make this more concrete, in this paper we will predict the velocity for a range of densities and snow-surface temperatures. If we want to predict on, say, a 21×21 grid in these two variables then V will have $10 + 21^2$ rows and eight columns. i.e. comprise about 3600 components. Therefore this is quite a large inference. At the same time, though, we are aware that our data, although the best of their kind, are noisy, and our physical model is rather simple. Therefore our judgements will play a large part in our predictions. This prioritises diagnostic information. It also makes us cautious about putting more structure in our judgements than we would willingly specify. For these reasons we strongly favour a Bayes linear analysis. The Bayes linear approach is outlined in Goldstein (1999) and described in detail in Goldstein and Wooff (2007); it has proved very powerful in large computer experiments (Craig et al., 1997, 2001; Goldstein and Rougier, 2004, 2006, 2008). It also underpins standard techniques such as Dynamic Linear Models (West and Harrison, 1997).

In the Bayes linear approach, expectation is taken as primitive and judgements are specified in terms of the mean and variance of a collection of quantities. Therefore we are not required to make higher-order specifications, in contrast to the fully probabilistic approach. The resulting updating equations have a simple form that allows rapid computation, including of diagnostic information. Our predictions take the form of a mean vector and a variance matrix over the product of our specified treatments $\boldsymbol{m}_a, \boldsymbol{m}_b, \ldots$ and the abscissae of the velocity profile.

To implement the Bayes linear approach for prediction we must specify a mean and variance over the collection $\{V, V^{\text{obs}}\}$. If we want to do model calibration we must also include $\boldsymbol{w}^*$ in that collection. Calibration is more complicated in a Bayes linear framework, because of the strong non-linearities that can exist between $\boldsymbol{w}^*$ and V, induced mainly by non-linearities in the physical model $g(\cdot)$. The Bayes linear approach to calibration is described in Goldstein and Rougier (2006). In this paper we will focus on prediction,

using the approach described in Craig et al. (2001).

We can write our statistical model in the general form

$$V \equiv G^* + D^* \tag{15a}$$

$$V^{\text{obs}} \equiv V + E \tag{15b}$$

where G^* is the collection of model evaluations with typical component $g(\boldsymbol{x}^*, \boldsymbol{m}_i, z_j)$, D^* is the collection of discrepancies with typical component

$$d(\boldsymbol{x}^*, \boldsymbol{m}_i, z_j) = v(\boldsymbol{m}_i, z_j) - g(\boldsymbol{x}^*, \boldsymbol{m}_i, z_j), \tag{15c}$$

and E is the collection of measurement errors with typical component

$$e(\boldsymbol{m}_i, z_j) = v^{\text{obs}}(\boldsymbol{m}_i, z_j) - v(\boldsymbol{m}_i, z_j). \tag{15d}$$

We induce a mean and variance on the collection $\{V, V^{\text{obs}}\}$ in terms of our choices for $\{G^*, D^*, E\}$. According to our choices in section 3.1, $E \perp\!\!\!\perp \{G^*, D^*\}$, has mean zero, and variance given by (14). D^* and G^*, however, covary, because they share a common source of uncertainty, namely $\boldsymbol{x}^*$. Including this covariance is a challenge for the Bayes linear approach. We judge that our purpose is better served by removing the dependence of D^* on $\boldsymbol{x}^*$. In fact, this is the standard approach in computer experiments where, as far as we are aware, no analysis has yet included the influence of the model parameters on the model's discrepancy, not even in a fully probabilistic approach; see, e.g., the standard set-up in Kennedy and O'Hagan (2001) and the discussion in Rougier (2007). In our statistical model this means removing the effect of $\boldsymbol{x}$ from the standard deviation function $\sigma(\boldsymbol{x}, \boldsymbol{m}, z)$, used in (13); we will show in section 5.3 how the lack of $\boldsymbol{x}$ can be partially mitigated through the creative use of $\boldsymbol{m}$. Now we have $E \perp\!\!\!\perp G^* \perp\!\!\!\perp D$, where we have dropped the '*' on D.

With this simplification we can write the joint mean and variance as

$$E\begin{pmatrix} V \\ V^{\text{obs}} \end{pmatrix} = \begin{pmatrix} \boldsymbol{\mu}^* \\ H\boldsymbol{\mu}^* \end{pmatrix} \tag{16a}$$

$$\text{Var}\begin{pmatrix} V \\ V^{\text{obs}} \end{pmatrix} = \begin{pmatrix} \Xi & \Xi H^T \\ H\Xi & H\Xi H^T + \text{Var}(E) \end{pmatrix} \tag{16b}$$

where H is the incidence matrix, which picks out the observations, $\Xi = \mathrm{Var}(V) = \Sigma^* + \mathrm{Var}(D)$, and $\boldsymbol{\mu}^*$ and Σ^* are the mean and variance of G^*. This still leaves us to determine the mean and variance of G^*. For this purpose we find it helpful to specify a probability distribution for $\boldsymbol{w}^*$, and then to infer $\boldsymbol{\mu}^*$ and Σ^* using samples drawn from G^*. This is a cheap calculation in our application because the physical model is quick to evaluate; where the physical model is expensive we would use an emulator (see, e.g., Craig et al., 2001; Kennedy and O'Hagan, 2001; O'Hagan, 2006; Higdon et al., 2007; Rougier, 2008). This approach is not a simple case of estimating $\boldsymbol{\mu}^*$ and Σ^*, though. Our initial choice for the distribution $\pi(\boldsymbol{w}^*)$ will be informed by our judgements, but we are likely to tune this choice in the light of the implied mean and variance for G^* and, possibly, to modify the resulting mean and variance in the light of other judgements. Thus the distribution $\pi(\boldsymbol{w}^*)$ is a contrivance that helps us to specify $\boldsymbol{\mu}^*$ and Σ^*, rather than a core part of our inference.

The scheme for adjusting our mean and variance for V on the basis of observed value for V^{obs} is described in Goldstein and Wooff (2007), ch. 3. The updating equations will be familiar because they are also the conditioning relations of a multivariate Gaussian distribution. However, as already explained, we have adopted the Bayes linear approach partly because of our reluctance to provide fully-probabilistic descriptions of our uncertainty, and therefore we do not judge (V, V^{obs}) to be Gaussian, and nor will we be using the extra structure that this would imply in our predictions. For example, our predictions are explicitly in terms of means, variances, and covariances: we cannot report quantiles without further restrictions.

5 Statistical modelling and results

5.1 Model parameters

In this subsection we specify $\boldsymbol{w}^*$, the function $\boldsymbol{k}(\cdot)$, and the marginal distribution $\pi(\boldsymbol{w}^*)$, as described in section 3.3. Recall from section 3.1 that $\boldsymbol{x} = (v_0, \tau_c, \alpha, t_c)$ and $\boldsymbol{m} = (\rho, T_{\mathrm{ss}})$. The precise manner in which we arrive at quantified distributions is described in section 5.4. Here we outline our reasoning and our general statistical framework.

First we consider the dependence of v_0^* on ρ, which appears to be a feature of our experimental data. For snow in the coexistence regime of snow and

water at 0°C, snow density and snow water content are typically positively related. High water content causes additional basal friction by lubrication-adhesion effects on the ground, especially if the flow is water-dominated, as in slush flows (Jaedicke et al., 2008). Compact, dry snow exhibits slightly less basal friction on dry ground. On the other hand, for low snow densities, a lower ratio between gravitational forces and basal friction may contribute to slower basal slip velocities: low snow densities are frequently associated with a finer ("felty") structure which may result in a higher effective basal friction coefficient. We judge that these two effects will probably combine in a concave relationship between ρ and v_0^* over our range of densities.

Second, we consider the dependence of τ_c on $T_{\rm ss}$. Foehn (1998) and Schweizer (1998) observed that τ_c depends both on the snow temperature and on the applied shear rate. For our Herschel-Bulkley flow, τ_c is the threshold stress for failure of the plug at $z = h$ where the shear rate $\mathrm{d}v/\mathrm{d}z$ vanishes (note that, under our experimental conditions, this transition from shear flow to plug flow might not be perfectly smooth). That is, we can use the experimental results for τ_c under low shear deformation rates, of the order $\mathrm{d}v/\mathrm{d}z \sim 10^{-3}s^{-1}$. Under these conditions, τ_c is decreasing in temperature.

The Appendix outlines a simple approach for quantifying both of these dependencies in terms of some basic elicitations. In summary, we restrict the form of the dependence to a quadratic with uncertain coefficients, and then we constrain the mean and variance of the coefficients. In the case of the basal slip velocity v_0^*, our constraints are (1) $E(v_0^*; \rho = 250) = 4.6$ and $E(v_0^*; \rho = 800) = 2.5$; (2) extremum at $\rho = 400$; (3) $\Pr(\text{concave}) = 0.95$. In the case of the shear τ_c^*, our constraints are (1) $E(\tau_c^*; T_{\rm ss} = -15) = 900$ and $E(\tau_c^*; T_{\rm ss} = -4) = 700$; (2) extremum at $T_{\rm ss} = 1$; (3) $\Pr(\text{concave}) = 0.05$. In each case we have one free statistical parameter with which to tune our choices. The results are shown in Figure 6, as realisations of the random functions. In fact, these realisations were our primary tool in setting the constraints to achieve a distribution of random functions that reflected our judgements.

Finally, we consider α^* and t_c^*, both of which are strictly positive quantities. In both cases we use Gamma distributions. We fix the means at our prior best guesses, 2 and $0.03\,\mathrm{s}^{-1}$, respectively. We then specify the shape and scale parameters to achieve a reasonable description of our uncertainty.

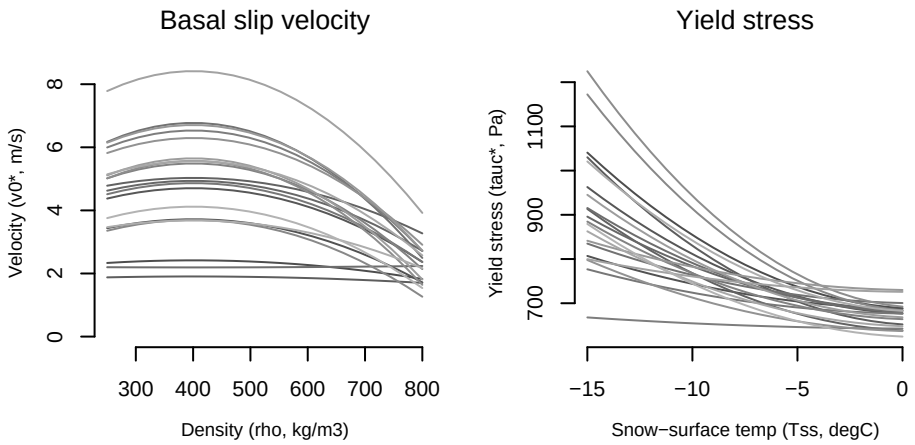

Figure 6: Realisations of the stochastic relationship between treatment and the 'best' value of the model parameters.

For α we choose shape 16 and scale 2/16 (giving a standard deviation of 1/2), and for t_c shape 9 and scale 0.03/9 (std dev. $0.01\,\mathrm{s}^{-1}$).

5.2 Observation error variance matrix

We treat the observation errors as uncorrelated, and use pooled estimates for the variances, based on whether $T_{\mathrm{ss}} \leq -2°\mathrm{C}$: there are simple physical reasons for thinking that the accuracy of the measurements will depend on temperature. At a snow temperature of 0°C, snow is in the phase transition regime. That is, both snow and water are present in the flow. The measurement principle of the velocity sensors is based on a correlation analysis of signals obtained from infra-red (IR) reflectivity sensors. The quality of the signal is a function of the water content of the snow: the ambient water blurs the signal peaks related to passing snow particles. Moreover, the transition spectrum of fluid water has a gap in the IR band used for the reflectivity measurements which causes the water film between snow particles and sensor to dampen the signals. For a detailed discussion of this technical problem, see Guenther (2006). Strictly speaking, fluid water is present only for $T_{\mathrm{ss}} = 0°\mathrm{C}$. But since T_{ss} was not measured in the snow chute but in a nearby snowpit, we use the slightly lower value of −2°C. The estimated standard deviations are 0.536 m/s for the low temperatures, and 1.670 m/s for the high temperatures (see Figures 2 and 3).

5.3 Discrepancy variance matrix

Standard deviation function. We would like to express the standard deviation of $d(\boldsymbol{x}^*, \boldsymbol{m}_i, z_j)$, denoted $\sigma(\boldsymbol{x}^*, \boldsymbol{m}_i, z_j)$ in general, as a two-level function of h^* and z, as explained in section 3.2. From (HB-c),

$$h = H - \frac{\tau_c}{g\rho \sin\theta} \tag{17}$$

and so h^* depends on both $\tau_c^* \in \{\boldsymbol{x}^*\}$ and $\rho \in \{\boldsymbol{m}\}$. However, as explained in section 4, for tractability we choose to exclude $\boldsymbol{x}^*$ from $\sigma(\cdot)$. But since h is linear in τ_c and the expectation of τ_c^* depends on $\boldsymbol{m}$, we can express $E(h^*; \boldsymbol{m})$ as a function of $\boldsymbol{m}$. Thus we replace h^* with its expectation, to give

$$\sigma(\boldsymbol{m}, z) = \begin{cases} \sigma_\ell & z < E(h^*; \boldsymbol{m}) \\ \sigma_u & z \geq E(h^*; \boldsymbol{m}) \,. \end{cases} \tag{18}$$

We choose the values $\sigma_\ell = 1.25\,\mathrm{m/s}$ and $\sigma_u = 1.75\,\mathrm{m/s}$.

Correlation functions. We choose to treat $\kappa_m(\cdot)$ as separable in ρ and T_{ss}, so that

$$\kappa_m(\boldsymbol{m}, \boldsymbol{m}') = \kappa_\rho(\rho, \rho') \times \kappa_t(T_{\mathrm{ss}}, {T_{\mathrm{ss}}}') \tag{19}$$

We use the Matérn correlation function to specify the three correlation functions: each one is parameterised in terms of a range r and a smoothness ν in the general form

$$\text{Matérn}(d; r, \nu) = \frac{1}{2^{\nu-1}\Gamma(\nu)} (d/r)^\nu B(d/r, \nu) \tag{20}$$

where d is Euclidean distance, and $B(\cdot)$ is the modified Bessel function. The smoothness ν is the less important parameter: we use $3/2$ in each case, reflecting our judgement that none of the relationships is particularly smooth. For r we use one quarter of the range for T_{ss}, one sixteenth of the range for ρ, and one half of the range for z, reflecting our judgement that the discrepancy is more systematic in z than in the two environmental variables. Originally we set the r parameter for ρ to be one quarter of the range (like T_{ss}), but we revised this value downward in the light of diagnostic information (see section 5.4). Quantifying these correlation lengths is not easy, but in our case the variance in V is dominated by variance in G^*, and

so these choices are not critical.

Conditioning at zero height. We make one further modification to the variance matrix $\mathrm{Var}(D)$. One of the model inputs is v_0, the basal slip velocity. This is also one of the outputs. We identify this input and output by conditioning $d(\boldsymbol{m}, z)$ on $d(\boldsymbol{m}, 0) = 0$ for all $(\boldsymbol{m}, z)$; in practice we condition on $d(\boldsymbol{m}_g, 0) = 0$ for all $\boldsymbol{m}_g$ in a dense regular grid over the treatment space. After this modification, the only source of uncertainty about $v(\boldsymbol{m}, 0)$ is $(v_0^*; \boldsymbol{m})$.

5.4 Diagnostics and re-modelling

Prior predictive mean and standard deviation. It is difficult for us to make judgements about the marginal distribution $\pi(\boldsymbol{w}^*; \boldsymbol{m})$ directly, because the semi-empirical nature of our physical model means that the quantities $\boldsymbol{w}^*$ and $\boldsymbol{x}^*$ are not operationally defined. Therefore, we make these judgements partly indirectly, by examining their implications for the velocity profile at different treatments, which *is* operationally defined. The velocity profile at any particular $\boldsymbol{m}$ will be an uncertain quantity that synthesises our choices for $\boldsymbol{w}^*$, the HB model, and our statistical model for the discrepancy. We treat the HB model as inviolate, so that if the velocity profile does not accord with our judgements, then we need to change the marginal distribution of $\boldsymbol{w}^*$, the mapping function $k(\cdot)$, or the variance function of the discrepancy. Once the velocity profile looks about right (a more formal criterion is not possible here), we know at least that our judgements on the two domains, $\boldsymbol{w}^*$ and $v(\boldsymbol{m}, z)$, are consistent. When we first made this comparison, we found that the velocity profiles looked quite wrong: the uncertainties were far too large. We used this information mainly to alter our choices for the marginal variances of α^* and t_c^*, described in section 5.1.

Figure 7 shows our main prior predictive summary diagnostic: the mean and standard deviation for the velocity profile in a 3×3 layout of snow densities and snow-surface temperatures. Each panel shows both the total uncertainty, and, inside that, the uncertainty attributable to uncertainty in the model parameters: the difference is uncertainty attributable to the discrepancy. The first source of uncertainty dominates, more so at higher densities.

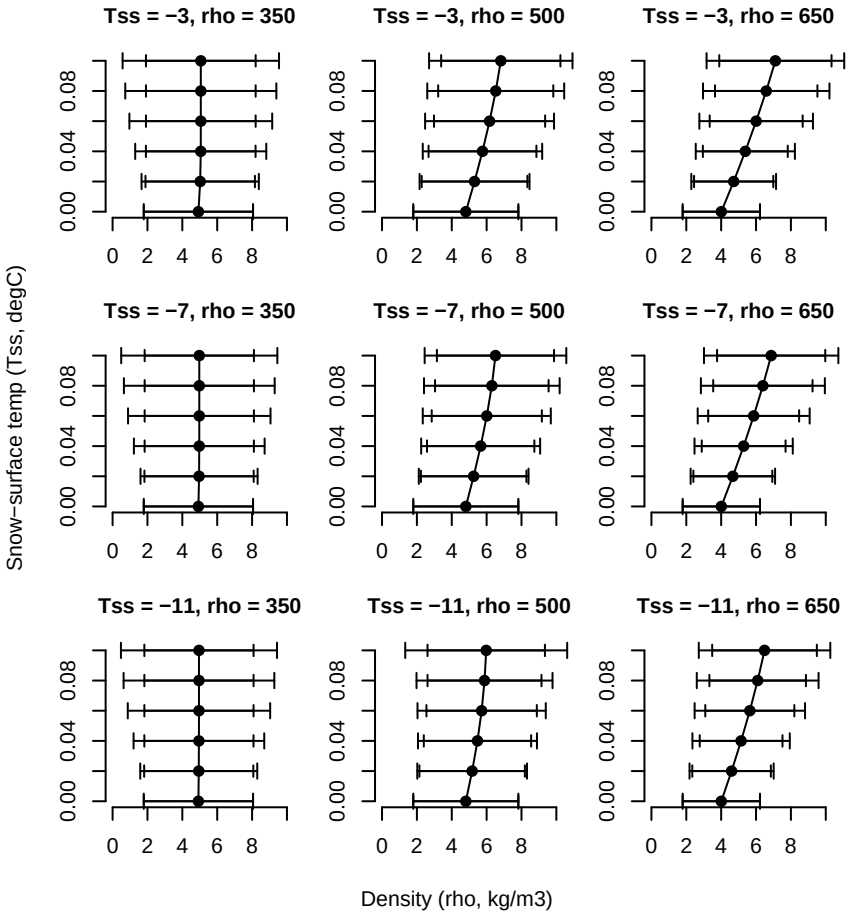

Figure 7: Prior predicted velocities on a regular grid in snow density and snow-surface temperature (*cf.* Figures 2 and 3). The share attributable to uncertainty in the model parameters is shown by the ticks inside each error bar, with the rest being due to the discrepancy between the model and actual snow behaviour.

Joint structure. Figure 7 summarises the marginal structure of our judgements, i.e. taken pointwise at different treatments. Our judgements about the joint structure are much less well-formed: certainly not well-enough formed that we might use them as the basis for further adjustments to our choices for the marginal distribution of $\boldsymbol{w}^*$. However, it is still interesting to see what our choices and the HB model imply for the covariance of velocities over the treatment space.

For simplicity, we restrict attention to a single height, $z = 0.4$m. The velocity $v(\boldsymbol{m}, 0.4)$ is the velocity of the plug-layer, v_h. From (HB-b), this is a linear function of v_0 and a non-linear function of τ_c (through h, see (HB-c)), α, and t_c. Therefore our uncertainty about v_h^* will be affected by both of the environmental variables, since ρ affects h directly and v_0^* indirectly, and T_{ss} affects τ_c^* indirectly. To visualise the joint structure of this uncertainty, we compute the spectral decomposition of the variance matrix of v_h^* over a high-resolution grid in the two environmental variables, and plot the first few eigenvectors. These plots are not shown here: they cannot be effectively reproduced without colour (colour plots are available at `http://www.maths.bris.ac.uk/~mazjcr/snowEigen.pdf`). The first eigenvector accounts for 43% of total variation, and describes uncertainty about the general height of the function, without much differentiation according to the values of the environmental variables. The second eigenvector accounts for 8% of the total variation, and describes the tilt of the function along a fulcrum roughly orthogonal to density. All of the remaining eigenvalues are quite small.

From the point of view of experimental design, we conclude that experiments with extreme values of snow density would be more valuable at reducing uncertainty over the range of environmental variable values than those with extreme values of snow-surface temperature. This accords with our intuition that density plays a larger role that temperature. But note the *caveat* that about 50% of the variation lies outside the first two components, comprising many small contributions. Therefore many experiments will be required to reduce uncertainty substantially everywhere: certainly more than the ten experiments we currently have.

Leave-one-out diagnostic. For our observation-based diagnostic we use leave-one-out plots. For each experiment in turn, we predict the observations using the outcome of the other nine experiments, and compare the prediction

with the actual observations. Each prediction takes the form of a mean vector and a variance matrix, therefore we transform the prediction errors so that they ought to have mean zero and standard deviation one, and be uncorrelated.

The resulting prediction errors for each experiment are shown in Figure 8. This is not the first such plot we made, as we have permitted ourselves a small amount of tuning of this diagnostic. While the arrangement of the points is broadly satisfactory, there is apparent some tendency for the points in any one experiment to be systematically biased away from zero: experiments G and H show this most markedly (although recollect from Figure 1 that these two experiments took place under the same environmental conditions). This tendency was previously stronger, and we traced its source to the correlation length in the ρ contribution to the discrepancy variance. We shortened this correlation length, reducing almost to zero the correlation between experiments of different densities, in the discrepancy. Therefore common effects across treatments arise mainly through uncertainty in the model parameters. Our inferences in the next section, which focus on means and marginal standard deviations, are not very sensitive to the values of the environmental variables' correlation lengths.

6 Results

Our main interest is in the velocity profile across a range of values for the two environmental variables. For simplicity, we restrict attention to the velocity of the plug-layer (i.e. $z = 0.4\,\mathrm{m}$). We present both our prior assessment, and our assessment after adjusting by the observations from the ten experiments. The results are given in Figures 9 and 10. Each Figure shows the mean function, represented in terms of grey-scale and contours, and also an indication of the pointwise standard deviation, in terms of the size of the grey boxes: details are given in the caption to Figure 9.

In Figure 9, the prior mean field shows that snow density is more important than snow-surface temperature, but that there is an interaction between the two, so that temperature is much more influential when density is high. A simple observation, but an important one, is that this mean field could not have been constructed from a representative set of model parameters, even in the absence of a discrepancy term, because the HB model is non-linear

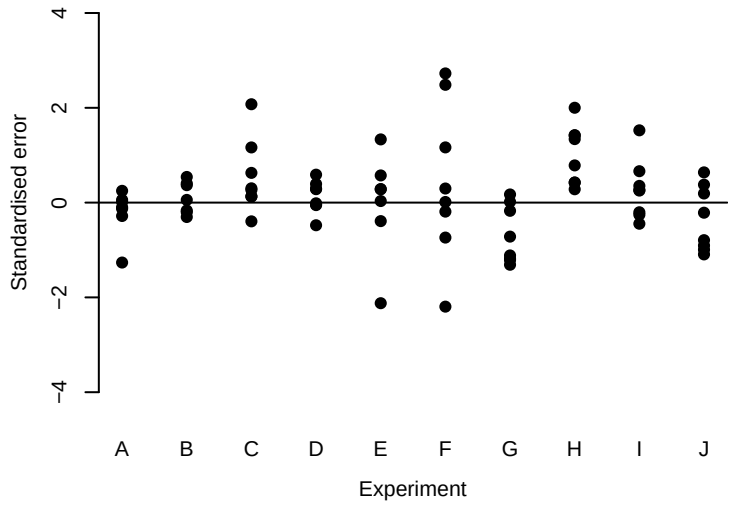

Figure 8: Leave-one-out diagnostic. The observations for each experiment in turn are predicted using observations from the other nine. The prediction errors from each experiment are standardised so that they ought to have mean zero and standard deviation one, and be uncorrelated.

in $\boldsymbol{x}$. Therefore specifying our uncertainty about $\boldsymbol{x}^*$ is an essential part of deriving this mean field. The prior standard deviations range from 2.3 m/s (low density, low temperature) to 3.0 m/s (high density, high temperature).

In Figure 10, the most noticeable effect of adjusting by the observations from our ten experiments is that the mean field has increased by about 1 m/s everywhere. This is consistent with our observation following the spectral analysis of the prior variance, which indicated that the overall level of the function was the dominant source of uncertainty. The 'tilt' of the function does not seem to have changed much. The marginal standard deviations have been reduced to the range 1.8 m/s to 2.2 m/s. This moderate reduction in uncertainty (less than 1 m/s) is also consistent with the spectral analysis.

The spectral analysis also indicates that we should not expect a few additional experiments to have a substantial impact on our uncertainty. The HB model is a simple model of a very complex physical process, and, by the very nature of the process and the experimental set-up, the observations we have are quite noisy. In this situation, it would be unrealistic to expect ten or even twenty experiments to substantially reduce our uncertainty over the whole range of values for the environmental variables. But what is

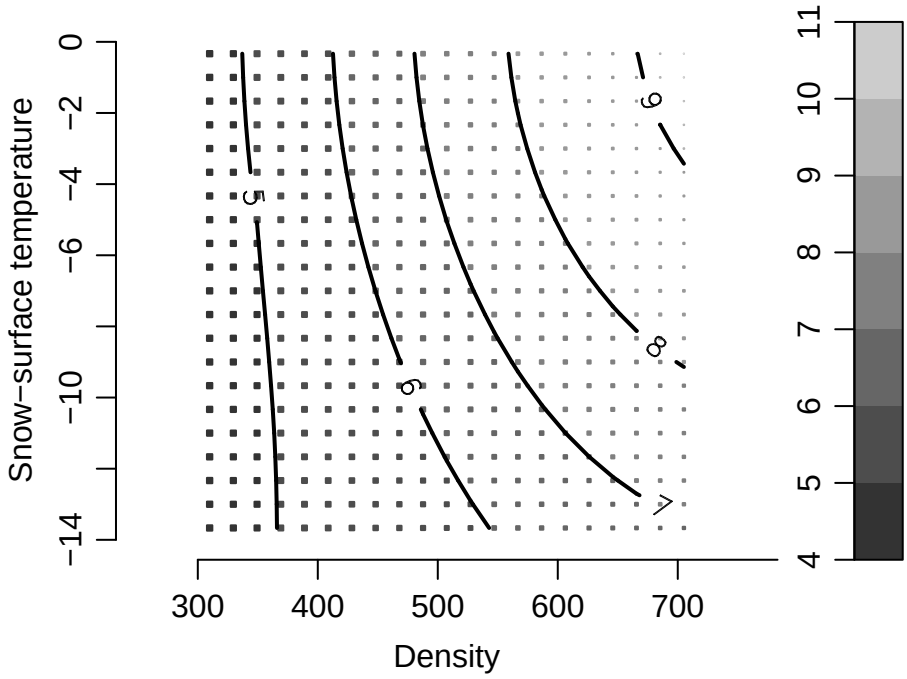

Figure 9: Prior predictive mean and standard deviation of velocity at a height of 0.4 m, by snow density and snow-surface temperature. The grey-scale and contours show the mean velocity. The width of each grey square is inversely proportional to the standard deviation, so that regions with more white indicate more uncertainty; the standard deviations run from 0 m/s (full width) to ≥ 3 m/s (no width, all white).

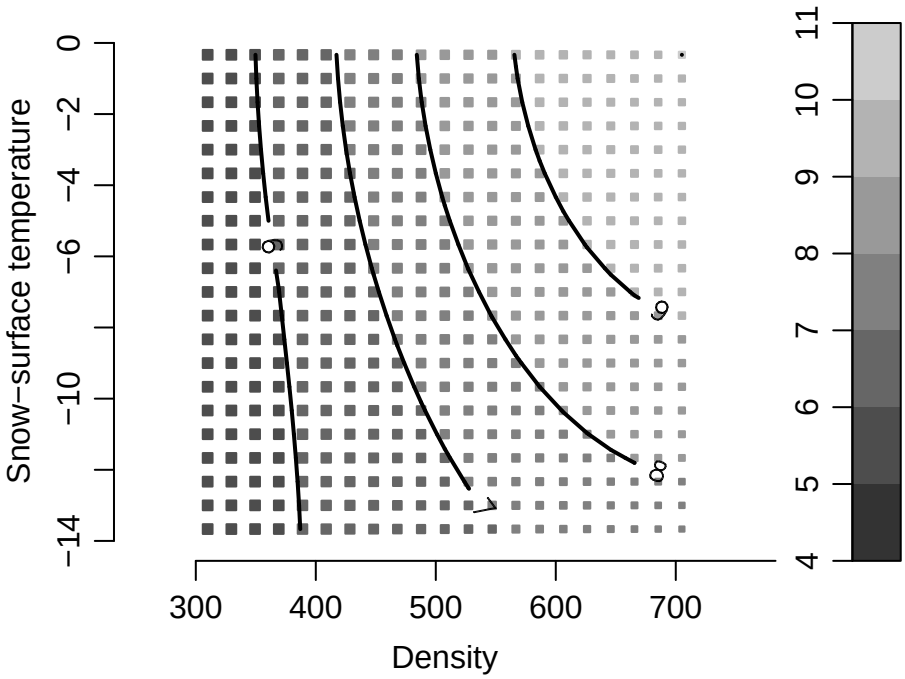

Figure 10: Adjusted predictive mean and standard deviation of velocity at a height of 0.10 m, *cf.* Figure 9.

important is that we have, for the first time, quantified the uncertainty in our predictions for chute velocities that follows from the HB model's uncertain parameters and discrepancy, and the observational errors.

7 Conclusion

In this study we have extended the 'standard' approach to model-based inference for complex physical systems to include multiple experiments under different treatments. This involves introducing an explicit role for the treatment in the assessment of the model's discrepancy with reality, and also in the observation process. In our analysis of ten chute experiments using the Herschel-Bulkley (HB) model the treatment denoted a particular set of environmental conditions, described by snow density and snow-surface temperature. We incorporated the treatment into the variance function of the model discrepancy, where density affected the height of the shear layer, which interacted with the output index-variable, height. We also incorporated the treatment into the observation error, to account for the larger observation errors that arise when the snow is close to melting.

In general the treatment can also affect our judgement about the 'best' value of the model parameters. This requires a transformation of the model parameter space; put simply, we have to construct 'hyper-model-parameters' that are uncertain and treatment-independent. These hyper-parameters are the quantities in which information from experiments under different treatments is combined, to inform predictions over a range of treatment values. For the HB model we required two sets of such hyper-parameters: to account for the effect of density on basal slip velocity, and temperature on stress. In both cases they took the form of uncertain coefficients in quadratic equations. Hyper-parameters are commonly used in hierarchical statistical modelling, but this is, to our knowledge, the first time they have been used in this way in a computer experiment with a physical model.

Throughout our analysis we have made judgements about the HB model and the statistical framework linking model parameters, model evaluations, actual snow behaviour, and observations. These types of judgements are never easy, and can be contentious. But one of the great advantages of a Bayesian analysis is that it forces us to quantify our judgements in a transparent manner, to facilitate discussion, and—in due course—revision.

We have used a number of different methods for quantifying our judgements, including a novel but effective way of specifying simple uncertain functions. We have validated these judgements, collectively, in an assessment of prior predictive means and variances, and in a leave-one-out diagnostic analysis. They remain, however, *our* judgements, and we do not expect that they will meet with the approval of all statisticians and snow experts.

From a practical point of view, our results show how we can extend the prediction range of our model to environmental conditions that have not been experimentally observed. In this study we have focused on chute experiments, but our longterm goal is to construct better avalanche hazard maps, which are based on numerical modelling of avalanche velocity and run-out length. At the heart of these models are empirical rheology models similar to the HB model considered here. However, in current practice the discrepancy between these models and actual avalanche behaviour is never formally accounted for, and information on the snow cover properties (e.g. density and temperature) is ignored; both of these create uncertainty in the resulting hazard maps. The approach described here allows us to quantify these uncertainties and should lead to a more time- and site-specific approach to avalanche hazard estimation, by considering the local temporal and spatial properties of the snowcover in the avalanche release zones. Many more treatment variables would have to be taken into account (topography properties, erosional properties, spreading in 3D avalanche flow). But we are convinced that this is a neccessary and feasible task for future development of avalanche hazard management tools.

Acknowledgements

We would like to thank Jürg Schweizer and Martin Schneebeli for guidance on the experimental data for τ_c and their validity for avalanche dynamics modelling, and Ana García Suárez and Cari Kaufman for suggestions about visualising the mean and standard deviation fields, and the eigenvectors.

Appendix: Imposing constraints on random quadratic functions

Suppose we have a random function

$$y = \alpha_0 + \alpha_1 x + \alpha_2 x^2 \tag{A1}$$

where $\boldsymbol{\alpha} = (\alpha_0, \alpha_1, \alpha_2)$ is an uncertain vector subject to the following constraints:

(1) $E(y; x = x_1) = y_1$, and $E(y; x = x_2) = y_2$, where $x_1 \neq x_2$;

(2) y has an extremum at x_e with probability 1, where $x_e \neq (x_1 + x_2)/2$;

(3) y is concave with probability p;

(4) α_0 and α_1 are independent;

(5) $\boldsymbol{\alpha}$ is Multivariate Normal.

These conditions may or may not be consistent. Suppose that they are. Then the mean vector and variance matrix of $\boldsymbol{\alpha}$ are completely determined by one additional value, or free parameter, the standard deviation of α_0.

First, note that (5) implies we need only consider the mean and variance of $\boldsymbol{\alpha}$. Then (2) implies that we must have $\alpha_1 + 2\alpha_2 x_e = 0$, which constrains both the mean and variance of (α_1, α_2). Thus the three components of the mean vector must satisfy three linear constraints, given by (1) and (2):

$$\begin{pmatrix} 1 & x_1 & x_1^2 \\ 1 & x_2 & x_2^2 \\ 0 & 1 & 2x_e \end{pmatrix} \begin{pmatrix} \mu_0 \\ \mu_1 \\ \mu_2 \end{pmatrix} = \begin{pmatrix} y_1 \\ y_2 \\ 0 \end{pmatrix}. \tag{A2}$$

The conditions $x_1 \neq x_2$ and $x_e \neq (x_1 + x_2)/2$ ensure that the square matrix is non-singular.

Consider the case $x_e \neq 0$. For the variance, α_2 is a linear function of α_1, so if we specify the variance matrix of $\{\alpha_0, \alpha_1\}$ then we induce a variance matrix for $\boldsymbol{\alpha}$. By (4), α_0 and α_1 are independent, and so we only need to specify the two standard deviations. One of these is the free parameter $\sigma_0 = \mathrm{Sd}(\alpha_0)$. The other, σ_1, controls the probability that y is concave, or,

equivalently, the probability that α_2 is negative:

$$\Pr(\alpha_2 < 0) = \Pr(-\alpha_1/2x_e < 0) = \begin{cases} \Pr(\alpha_1 < 0) & x_e < 0, \\ \Pr(\alpha_1 > 0) & x_e > 0. \end{cases} \tag{A3}$$

In the first case we have $\sigma_1 = -\mu_1/\Phi^{-1}(p)$, where $\Phi^{-1}(\cdot)$ is the quantile function for the Standard Normal, using (3), and (5) again, and in the second case $\sigma_1 = -\mu_1/\Phi^{-1}(1-p)$. The constraints are consistent if $\sigma_1 > 0$.

In the case where $x_e = 0$, we must have $\mu_1 = \sigma_1 = 0$. In this case $\Pr(\alpha_2 < 0) = \Phi(-\mu_2/\sigma_2)$, and $\sigma_2 = -\mu_2/\Phi^{-1}(p)$. The constraints are consistent if $\sigma_2 > 0$.

References

Barnes, H. A., J. F. Hutton, and K. Walters (1989). *An Introduction to Rheology*. Amsterdam: Elsevier.

Bingham, E. C. (1922). *Fluidity and Plasticity*. New York: McGraw-Hill.

Craig, P., M. Goldstein, J. Rougier, and A. Seheult (2001). Bayesian forecasting for complex systems using computer simulators. *Journal of the American Statistical Association 96*, 717–729.

Craig, P., M. Goldstein, A. Seheult, and J. Smith (1997). Pressure matching for hydrocarbon reservoirs: A case study in the use of Bayes Linear strategies for large computer experiments. In C. Gatsonis, J. Hodges, R. Kass, R. McCulloch, P. Rossi, and N. Singpurwalla (Eds.), *Case Studies in Bayesian Statistics III*, pp. 37–87. New York: Springer-Verlag. With discussion.

Dent, J., K. Burrel, D. Schmidt, M. Louge, E. Adams, and T. Jazbutis (1997). Density, velocity and friction measurements in a dry snow avalanche. *Annals of Glaciology 26*, 247–252.

Foehn, P. (1998). Mechanical and structural properties of weak snow layers measured in situ. *Annals of Glaciology 26*, 1–6.

Goldstein, M. (1999). Bayes linear analysis. In S. Kotz (Ed.), *Encyclopaedia of Statistical Sciences, update vol. 3*, pp. 29–34. London: John Wiley & Sons.

Goldstein, M. and J. Rougier (2004). Probabilistic formulations for transferring inferences from mathematical models to physical systems. *SIAM Journal on Scientific Computing 26*(2), 467–487.

Goldstein, M. and J. Rougier (2006). Bayes linear calibrated prediction for complex systems. *Journal of the American Statistical Association 101*, 1132–1143.

Goldstein, M. and J. Rougier (2008). Reified Bayesian modelling and inference for physical systems. Forthcoming as a discussion paper in the *Journal of Statistical Planning and Inference*, currently available at `http://www.maths.dur.ac.uk/stats/people/jcr/Reify.pdf`.

Goldstein, M. and D. Wooff (2007). *Bayes Linear Statistics: Theory & Methods*. Chichester, England: John Wiley & Sons.

Guenther, G. (2006, March). Optimierung optischer Geschwindigkeitssensoren für slushflows. Technical Report x, Swiss Federal Institute for Snow and Avalanche Research.

Higdon, D., J. Gattiker, B. Williams, and M. Rightley (2007). Computer model calibration using high dimensional output. Technical Report LA-UR-07-1444, Los Alamos National Laboratory. To appear in the *Journal of the American Statistical Association*.

Jaedicke, C., M. Kern, P. Gauer, M. Baillifard, and K. Platzer (2008). Chute experiments on slushflow dynamics. *Cold Regions Science and Technology 51*, 156–156.

Kennedy, M. and A. O'Hagan (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society, Series B 63*, 425–464. With discussion.

Kern, M., F. Tiefenbacher, and J. McElwaine (2004). The rheology of snow in large chute flows. *Cold Regions Science and Technology 39*, 181–192.

McElwaine, J. and F. Tiefenbacher (2005). Calculating internal average velocities from correlation with error analysis. *Surveys in Geophysics 24*, 499–524.

O'Hagan, A. (2006). Bayesian analysis of computer code outputs: A tutorial. *Reliability Engineering and System Safety 91*, 1290–1300.

Oldroyd, J. G. (1947). A rational formulation of the equation of plastic flow for a Bingham solid. *Proceedings of the Cambridge Philosophical Society 43*, 100–105.

Rougier, J. (2007). Probabilistic inference for future climate using an ensemble of climate model evaluations. *Climatic Change 81*, 247–264.

Rougier, J. (2008). Efficient emulators for multivariate deterministic functions. In submission, currently available at `http://www.maths.bris.ac.uk/~mazjcr/OPemulator.pdf`.

Schweizer, J. (1998). Laboratory experiments on the shear failure of snow. *Annals of Glaciology 26*, 97–102.

Sovilla, B., M. Schaer, M. Kern, and P. Bartelt (2008). Impact pressures and flow regimes in dense snow avalanches observed at the Vallée de la Sionne test site. *Journals of Geophysical Research 113*. doi:10.1029/2006JF000688.

Tiefenbacher, F. and M. Kern (2004). Experimental devices to determine snow avalanche basal friction and velocity profiles. *Cold Reg. Sci. Tech. 38*, 17–30.

West, M. and J. Harrison (1997). *Bayesian Forecasting and Dynamic Models* (2nd ed.). New York: Springer-Verlag.