

University of Wollongong

Research Online

National Institute for Applied Statistics
Research Australia Working Paper Series

Faculty of Engineering and Information
Sciences

2014

Re-thinking soil carbon modelling: a stochastic approach to quantify uncertainties

D Clifford

CSIRO Computational Informatics

D Pagendam

CSIRO Computational Informatics

J Baldock

CSIRO Land and Water

N Cressie

University of Wollongong

R Farquharson

CSIRO Land and Water

See next page for additional authors

Follow this and additional works at: <https://ro.uow.edu.au/niasrawp>

Recommended Citation

Clifford, D; Pagendam, D; Baldock, J; Cressie, N; Farquharson, R; Farrell, M; Macdonald, L; and Murray, L, Re-thinking soil carbon modelling: a stochastic approach to quantify uncertainties, National Institute for Applied Statistics Research Australia, University of Wollongong, Working Paper 04-14, 2014, 39.
<https://ro.uow.edu.au/niasrawp/7>

Research Online is the open access institutional repository for the University of Wollongong. For further information contact the UOW Library: research-pubs@uow.edu.au

Re-thinking soil carbon modelling: a stochastic approach to quantify uncertainties

Abstract

The benefits of sequestering carbon are many, including improved crop productivity, reductions in greenhouse gases, and financial gains through the sale of carbon credits. Achieving better understanding of the sequestration process has motivated many deterministic models of soil carbon dynamics, but none of these models addresses uncertainty in a comprehensive manner. Uncertainty arises in many ways - around the model inputs, parameters, and dynamics, and subsequently model predictions. In this paper, these uncertainties are addressed in concert by incorporating a physical-statistical model for carbon dynamics within a Bayesian hierarchical modelling framework. This comprehensive approach to accounting for uncertainty in soil carbon modelling has not been attempted previously. This paper demonstrates proof-of-concept based on a one-pool model and identifies requirements for extension to multi-pool carbon modelling. Our model is based on the soil carbon dynamics in Tarlee, South Australia. We specify the model conditionally through its parameters, soil carbon input and decay processes, and observations of those processes. We use a particle marginal Metropolis-Hastings approach specified using the LibBi modelling language. We highlight how samples from the posterior distribution can be used to summarise our knowledge about model parameters, to estimate the probabilities of sequestering carbon, and to forecast changes in carbon stocks under crop rotations not represented explicitly in the original field trials.

Authors

D Clifford, D Pagendam, J Baldock, N Cressie, R Farquharson, M Farrell, L Macdonald, and L Murray

NIASRA

NATIONAL INSTITUTE FOR APPLIED
STATISTICS RESEARCH AUSTRALIA



***National Institute for Applied Statistics Research
Australia***

The University of Wollongong

Working Paper

04-14

**Re-thinking soil carbon modelling:
A stochastic approach to quantify uncertainties**

D. Clifford, D. Pagendam, J. Baldock, N. Cressie, R. Farquharson, M. Farrell,
L. Macdonald, and L. Murray

*Copyright © 2013 by the National Institute for Applied Statistics Research Australia, UOW.
Work in progress, no part of this paper may be reproduced without permission from the Institute.*

National Institute for Applied Statistics Research Australia, University of Wollongong,
Wollongong NSW 2522. Phone +61 2 4221 5435, Fax +61 2 4221 4845. Email:
anica@uow.edu.au

Re-thinking soil carbon modelling: A stochastic approach to quantify uncertainties

D. Clifford^{a*} D. Pagendam^a, J. Baldock^b, N. Cressie^c, R. Farquharson^b, M. Farrell^b, L. Macdonald^b, and L. Murray^a

Summary:

The benefits of sequestering carbon are many, including improved crop productivity, reductions in greenhouse gases, and financial gains through the sale of carbon credits. Achieving better understanding of the sequestration process has motivated many deterministic models of soil carbon dynamics, but none of these models addresses uncertainty in a comprehensive manner. Uncertainty arises in many ways - around the model inputs, parameters, and dynamics, and subsequently model predictions. In this paper, these uncertainties are addressed in concert by incorporating a physical-statistical model for carbon dynamics within a Bayesian hierarchical modelling framework. This comprehensive approach to accounting for uncertainty in soil carbon modelling has not been attempted previously. This paper demonstrates proof-of-concept based on a one-pool model and identifies requirements for extension to multi-pool carbon modelling. Our model is based on the soil carbon dynamics in Tarlee, South Australia. We specify the model conditionally through its parameters, soil carbon input and decay processes, and observations of those processes. We use a particle marginal Metropolis-Hastings approach specified using the LibBi modelling language. We highlight how samples from the posterior distribution can be used to summarise our knowledge about model parameters, to estimate the probabilities of sequestering carbon, and to forecast changes in carbon stocks under crop rotations not represented explicitly in the original field trials.

Keywords: physical-statistical model; greenhouse gas mitigation; decomposition; exponential decay

^aCSIRO Computational Informatics, GPO Box 2583, Brisbane, QLD 4001, Australia

^bCSIRO Land and Water, Waite Campus, Waite Road - Gate 4, Glen Osmond, SA 5064, Australia

^cNational Institute for Applied Statistics Research Australia (NIASRA), School of Mathematics and Applied Statistics, University of Wollongong, Wollongong, NSW 2522, Australia

* Correspondence to: Commonwealth Scientific and Industrial Research Organisation (CSIRO) Computational Informatics, GPO Box 2583, Brisbane, QLD 4001, Australia. E-mail: David.Clifford@csiro.au

1. INTRODUCTION

Organic carbon stored in soils represents a significant global sink that accounts for 1260-2400 petagrams C depending on the depth over which the stocks were integrated (Eswaran *et al.*, 1995; Jobbágy and Jackson, 2000; Houghton, 2003; Todd-Brown *et al.*, 2013), where 1 petagram = 10^{15} g. Using estimates of 1500 Pg C and 720 Pg C respectively, for organic carbon in soils and atmospheric CO₂ (Powlson, 2005) and an atmospheric CO₂ concentration of 390 ppm (Mauna Loa Observatory, Dec 2010, <http://www.esrl.noaa.gov/gmd/ccgg/trends/global.html>), losses or gains of 1% of carbon stocks in soils would alter atmospheric concentrations of CO₂ by up to 2% or nearly 8 ppm, although actual variations may be lower due to buffering effects of vegetation and oceanic processes. Given reports that the introduction of agricultural production in Australia has resulted in losses of 20-75% of the organic carbon that was present in undisturbed native soil (Lal, 2004; Luo *et al.*, 2010), the potential exists to store increases in atmospheric CO₂ in soil and reduce the net effect of anthropogenic greenhouse gas emissions on atmospheric and global processes.

At the farm scale, the positive effect that soil organic matter (i.e., organic carbon, nitrogen, phosphorus, and sulphur) has on a variety of soil properties and on processes that are key to ensuring productivity, sustainability, and resilience of the soil resource, provides additional incentive to landowners to enhance storage of organic carbon in soil (Baldock and Skjemstad, 1999; Hoyle *et al.*, 2011). Furthermore, the development of carbon crediting systems (e.g., the Australian Government's legislated Carbon Farming Initiative (CFI), <http://www.climatechange.gov.au/reducing-carbon/carbon-farming-initiative>) encourages storage and retention of additional organic carbon in soil through the provision of financial incentives for landowners who add and store carbon for one hundred years.

Changes in soil organic carbon (SOC) stocks occur as a result of alterations to the balance between the magnitudes of carbon input and loss. In the absence of soil erosion or

deposition, inputs are generated when atmospheric carbon is captured by plants through photosynthesis and is deposited on or in soil as plant material, leaf material, roots or root exudates (Jones *et al.*, 2009). Losses occur through decomposition of plant material and subsequent mineralisation of SOC to CO₂. Erosion and deposition may contribute significantly to any changes in soil carbon stocks and may confound attempts to quantify the impact of management practices (Sanderman and Chappell, 2013). Given strong effects on SOC magnitudes due to soil, climate conditions, and diverse management practices, measurements over more than ten years may be required to quantify trends in SOC stocks due to applied management practices (Sanderman and Baldock, 2010).

In response to this requirement of a lengthy sampling period, computer-simulation models such as Century (Parton *et al.*, 1988) and RothC (Jenkinson *et al.*, 1987) were constructed and have been used to estimate changes in soil carbon stocks induced by variations in applied agricultural management practices (Liu *et al.*, 2009, 2011; Robertson and Nash, 2013; Wang *et al.*, 2013) or variations in climatic conditions (Smith *et al.*, 2009; Wan *et al.*, 2011; Gottschalk *et al.*, 2012; Senapati *et al.*, 2013). Sierra *et al.* (2012) describe a modelling framework that provides easy access to a variety of multi-pool deterministic models including Century and RothC. These models handle differences in observed carbon dynamics through the conceptualisation of multiple pools that differ in terms of turnover time and chemical composition. Recent work has linked conceptual pools to measurable pools for the calibration of RothC based on data from Australia (Skjemstad *et al.*, 2004) and Switzerland (Zimmermann *et al.*, 2007). In the former case, Skjemstad *et al.* (2004) calibrated the RothC model using measurable SOC fractions across the continuous-wheat and wheat-fallow rotations of the Tarlee dataset. This dataset is the result of a twenty-year field trial in the mid-north of South Australia, designed to assess agronomic productivity for which observations of SOC were also recorded. The calibration referred to above, which also involved a similar-sized dataset from Brigalow Research Station in Queensland, led to

the development of the Full Carbon Accounting Model (FullCAM) and database system (Richards and Evans, 2004), which were created to provide retrospective estimates of annual changes in Australian national soil carbon stocks that satisfy the United Nations Framework Convention on Climate Change reporting requirements (<http://unfccc.int>).

Despite the widespread use of FullCAM, RothC, and Century in estimating temporal changes in soil carbon stocks for carbon-accounting purposes at a range of spatial scales, there has been little work to account for uncertainty in the data nor to quantify uncertainty in the model outcomes. FullCAM's estimates of changes in stocks can be linked to a point in the landscape or across a series of paddocks with defined areas. To provide meaningful outputs for carbon-accounting purposes, sources of uncertainty associated with input data and the values assigned to model parameters must be identified, quantified, and propagated through modelling scenarios, allowing one to establish a "level of confidence" in the model outputs (Refsgaard *et al.*, 2007). Attaching a certainty measure to a model's output would help landowners understand the risks associated with putting their sequestered carbon into a trading scheme.

Some effort has been made to quantify certainty in model outputs, for example, either through a sensitivity analysis where models are run for various sets of parameter values, or under simulated climate scenarios (e.g., Paul *et al.*, 2003; Juston *et al.*, 2010; Stamati *et al.*, 2013). A statistical modelling approach also makes it possible to quantify certainties in modelled carbon stocks (Koo *et al.*, 2007; Post *et al.*, 2008; Jones *et al.*, 2007). Various forms of data assimilation/inverse modelling have been undertaken to provide information on crop inputs (Meermans *et al.*, 2013), sensitivity of decomposition to temperature (Sakurai *et al.*, 2012), pool structure (Schädel *et al.*, 2013), initialisation of SOC pools (Scharnagl *et al.*, 2010; Yeluripati *et al.*, 2009), and carbon transfer coefficients (Xu *et al.*, 2006). Others demonstrate the difficulties in constraining soil carbon parameters (Richardson *et al.*, 2010; Rahn *et al.*, 2011; Wang and Chen, 2013).

Here we highlight the Bayesian hierarchical modelling (BHM) framework as a viable approach for achieving this goal. One advantage of BHMs is the ability to combine coherently the underlying soil carbon process and its exponential decay over time, with observations of soil carbon and our understanding of its associated measurement errors.

The objective of this paper is to create a BHM for SOC, drawing on management practices, crop production, and soil data collected from a research trial located near Tarlee in South Australia. While our model may not have broad applicability from a soil-science perspective, the approach taken is broadly applicable. A simple one-pool model for SOC inside a BHM was developed to demonstrate its potential utility and outcomes. Stochastic errors associated with the one-pool model account for heterogeneity in the decay time, although its generalisation to multiple pools will then fully capture the diverse composition and decomposition rates of SOC. In Section 2, we give a one-pool physical-statistical model for soil carbon dynamics within the BHM framework. In Section 3, this BHM is made specific to the Tarlee dataset, and statistical computational aspects are presented. In Section 4, we use the results of Section 3 to answer important questions that may be asked by a landowner, a scientist, or a policy-maker. Section 5 contains discussion and conclusions, particularly a discussion of what would be required to extend our BHM approach to a multi-pool model. Technical aspects can be found in the Appendix.

2. A PHYSICAL-STATISTICAL APPROACH TO MODELLING SOIL CARBON DYNAMICS

Soil organic carbon is heterogeneous, with various chemistries and protective physical and chemical associations resulting in functionally discrete pools of carbon of varying stabilities. To capture this heterogeneity, the dynamics are typically modelled using a fixed number of discrete pools or compartments that interact over time. Each of these conceptual pools can be considered to represent organic molecules of a particular complexity, size, and resistance

to decomposition. For example, simple carbohydrate molecules are more readily utilised by soil microorganisms than say complex polymers such as lignins, whilst even more resistant carbon molecules (e.g., natural charcoal from fires or added biochar) may belong to an inert pool that degrades very slowly, generally over centennial or millennial timescales (Lehmann *et al.*, 2008). Commonly, soil carbon models make use of simple, first-order reaction kinetics, such that the masses in each of the carbon pools undergo exponential decay over time. The rates of decomposition associated with the pools are modified by changes in moisture and temperature. As organic carbon in a pool decays, it is metabolised and either enters a different pool, or it is mineralised to CO_2 and lost to the atmosphere. In this article we assume a single pool where heterogeneity is captured stochastically; hence as organic carbon decays it goes directly into the atmosphere. This allows us to demonstrate unambiguously the importance of generalising from *deterministic* decay models to *stochastic* decay models.

Alternative to the compartment models, some models have cohorts of carbon where the SOC decomposes as a continuous function of its age (e.g., the Q model by Bosatta and Ågren, 1985), reflecting the changes in quality of each cohort over time. Regardless of whether carbon is pooled based on its decay-rate or based on its age, assumptions are made about the biophysical processes that govern decomposition, and the validity of these assumptions has not been adequately tested because the models contain parameters that cannot be easily estimated (Feng, 2009). Models based on conceptual pools need to be calibrated using data collected on physical fractions. Skjemstad *et al.* (2004) and Zimmermann *et al.* (2007) demonstrate two different approaches for linking physical fractions with conceptual pools, but such pools are not homogeneous and models of the soil organic continuum remain the ultimate research challenge.

In what follows, we show how a one-pool dynamical model of SOC can be incorporated into the BHM framework (Section 1) to account for: (i) uncertainties in model parameters; (ii) randomness inherent to the underlying physical processes (plant growth, carbon inputs,

and carbon decomposition); and (iii) errors in the observations of above-ground and below-ground crop inputs and below-ground SOC. Uncertainties arising from each of these three sources are modelled respectively through the prior $[\Theta]$, the process model $[Y|\Theta = \theta]$, and the data model $[Z|Y = y, \Theta = \theta]$, where the notation $[\cdot]$ denotes the probability density function of the enclosed random variable (e.g., Cressie and Wikle, 2011, Ch. 2), $[\cdot|E]$ denotes the conditional probability density function given the event E , Θ denotes unknown parameters, Y denotes an unobserved state process, and Z denotes observations. Under this hierarchical framework, we can write the joint distribution of all of the quantities in our model as, $[\Theta, Y, Z] = [\Theta] \times [Y|\Theta] \times [Z|Y, \Theta]$. Note that all the uncertainty in the model is captured by the joint distribution, which conveniently can be written as a simple product of the components of the BHM.

Unlike deterministic models, where parameters (Θ) are set and remain fixed, a model analysed within the BHM framework allows prior knowledge about the parameters to be incorporated into the analysis, but the distributions (and hence uncertainty) of these parameters are updated as data are observed. Moreover, the true physical state of the SOC (Y) at different times t can also be updated. That is, we build a physical-statistical model based on the BHM framework, and we make inferences about parameters, soil carbon, and functions of these through the posterior distribution, $[Y, \Theta|Z]$, which is the conditional distribution of Y and Θ given the data (Z). Bayes' Theorem (Bayes and Price, 1763) allows us to write the posterior distribution as $[Y, \Theta|Z] = [Z|Y, \Theta] \times [Y, \Theta]/[Z]$, which, using the law of total probability, can be rewritten as $[Y, \Theta|Z] = [Z|Y, \Theta] \times [Y|\Theta] \times [\Theta]/[Z]$. The denominator, $[Z]$, which depends only on the data, ensures that the posterior distribution integrates to 1; this may be difficult to calculate analytically or numerically, which in turn means that the posterior itself may be difficult to evaluate. Fortunately, it is not necessary to evaluate the posterior as long as one can simulate from it. In the next section, the physical-statistical modelling of soil carbon dynamics, and our method for simulating from

the associated posterior distribution, are made more specific through an application to a dataset from Tarlee, South Australia.

3. BAYESIAN HIERARCHICAL MODELLING AND COMPUTATIONAL STATISTICS: THE TARLEE DATASET

3.1. Background to the Tarlee Dataset

Tarlee is located 80 km north of Adelaide, South Australia. A long-term trial was established in 1977 on a hard-setting red-brown earth soil (1% C, pH 6.8, clay 14%; Schultz, 1995) with sandy loam texture (Alfisols as classified by the USDA soil classification system; Soil Survey Staff, 1975). The trial site has winter-dominated rainfall (average of 355 mm, April-October) with an annual average of 475 mm. The effects of rotation (cereal, grain legume, pasture, fallow), stubble (with and without removal), cultivation (with and without tillage), and nitrogen fertiliser rate (0, 40, 80 kg N ha⁻¹) on crop production and soil properties were monitored over a twenty-year period. The long-term-monitoring dataset has been used in carbon-modelling studies across Australia (Skjemstad *et al.*, 2004; Luo *et al.*, 2011). It appears that from 1979-1997 there is a greater decline in SOC under wheat-fallow rotation (from 38.6 t ha⁻¹ to 33.8 t ha⁻¹) compared to continuous wheat (39.2 t ha⁻¹ to 35.9 t ha⁻¹).

3.2. Notation

We let $Y_C(i, t)$, $Y_G(i, t)$, $Y_W(i, t)$, and $Y_P(i, t)$ denote respectively, the masses of SOC, total grain dry matter, total wheat dry matter, and total pasture dry matter produced in field i at time t (all in t ha⁻¹). We let $Y(i, t)$ serve as shorthand for all processes in field i at time t , that is, $Y(i, t) \equiv \{Y_C(i, t), Y_G(i, t), Y_W(i, t), Y_P(i, t)\}$. In order to indicate all processes in all fields at time t we use $Y(., t) \equiv \{Y(1, t), Y(2, t), Y(3, t)\}$, and to refer to a specific process in all fields we use $Y_C(., t)$, etc. All processes in all fields and all times is denoted by Y , and all

SOC process are denoted by Y_C . Similar index notation is used when referring to the data associated with each process; for example $Z_C(i, t)$ is the SOC data in field i at time t , and Z refers to all data (in all fields and all times).

3.3. Process Model

Let $I_C(i, t)$ denote the mass of carbon inputs (t ha^{-1}) for field i at time t . We model the soil carbon dynamics (i.e., the state process) for field i using the following process model:

$$\log(Y_C(i, t + \Delta t)) = \log(Y_C(i, t)e^{-K\Delta t} + I_C(i, t + \Delta t)) + \eta(i, t); \quad \eta(i, t) \stackrel{i.i.d.}{\sim} N(0, \sigma_\eta^2). \quad (1)$$

Notice that, at its core, the model has exponential decay with rate K (like the models used by RothC and Century), but it also has a stochastic component, represented by $\eta(i, t)$, that captures the uncertainty we have in the way that $Y_C(i, t + \Delta t)$ evolves from $Y_C(i, t)$. A component of that uncertainty captures the heterogeneity that the one-pool model does not account for.

The soil carbon dynamics of each field are modelled independently, with each field's initial conditions at $t = 1978$ specified as part of the prior distribution for Θ . For the purposes of this paper, the rate of decay K is assumed to be the same in each field, and the distribution of the stochastic component $\eta(i, t)$ does not change over time or from one field to another. Spatial models for the underlying SOC process are also possible. However, the information required to parameterise such a model is not available, since the Tarlee study does not contain data to learn about the nature of the spatial relationships of the SOC process.

The carbon input, $I_C(i, t)$, to field i at time t , is dependent on the crops growing that year and how they were harvested. The total wheat dry matter includes the total grain dry matter (i.e., $Y_G(i, t) \leq Y_W(i, t)$), but we separate the two processes here due to the fact that grain yield is usually measured, and the associated wheat dry matter is inferred from the grain yield. The conditionally independent process models for $Y_G(i, t)$, $Y_W(i, t)$, and $Y_P(i, t)$

are:

$$Y_G(i, t + \Delta t) \sim \text{LogN}(\mu_G + \rho_G (Y_G(i, t) - \mu_G), \sigma_G^2); \quad (2)$$

$$Y_W(i, t + \Delta t) \sim \text{LogN}(\log h + \log y_G(i, t + \Delta t), \sigma_W^2); \quad \text{and} \quad (3)$$

$$Y_P(i, t + \Delta t) \sim \text{LogN}(\mu_P + \rho_P (Y_P(i, t) - \mu_P), \sigma_P^2), \quad (4)$$

where $Y_W(i, t + \Delta t)$ is defined conditional on $Y_G(i, t + \Delta t) = y_G(i, t + \Delta t)$ and $\text{LogN}(\mu, \sigma^2)$ denotes the lognormal distribution with parameters μ and σ^2 . The parameters ρ_G and ρ_P are used to model grain and pasture yields as autoregressive processes. The parameter h is known as the *harvest index*, and is used in a model for the expected wheat yield as a function of the grain yield. That is, Y_W has no direct autoregressive relationship, but it does have temporal variability through its dependence on Y_G .

Finally, the input function, $I_C(i, t)$, can now be specified: It depends upon the crop type and includes carbon that comes into the soil both from plant-matter that remains above-ground after harvesting and from plant-matter that is already below-ground (i.e., roots). When wheat is sown for grain, the carbon that enters the soil from the plant-matter that remains above-ground after harvesting is, $c(Y_W(i, t) - Y_G(i, t))$, where c is the carbon content of generic plant material and $(Y_W(i, t) - Y_G(i, t))$ is the above-ground plant-matter biomass. Root-to-shoot ratios are used to infer below-ground plant-matter biomass based on above-ground plant-matter biomass. The amount of carbon below-ground for wheat is then $cr_W Y_W(i, t)$, where r_W is the root-to-shoot ratio for wheat. Five management practices based on three crops were used in the Tarlee study. The carbon input to field i at time t

based on the management practice is:

$$I_C(i, t) = \begin{cases} c(Y_W(i, t) - Y_G(i, t)) + cr_W Y_W(i, t) & \text{Wheat for Grain} \\ cpY_W(i, t) + cr_W Y_W(i, t) & \text{Wheat for Hay} \\ cY_P(i, t) + cr_P Y_P(i, t) & \text{Pasture} \\ cpY_P(i, t) + cr_P Y_P(i, t) & \text{Pasture for Hay} \\ 0 + 0 & \text{Fallow,} \end{cases} \quad (5)$$

where p is the proportion of the crop left above-ground as stubble following harvesting as hay, and r_W and r_P are respectively, the root-to-shoot ratios for wheat crops and for pasture. Each entry in Equation 5 is of the form $A + B$ where A is the above-ground carbon input and B is the below-ground carbon input. Wheat and pasture crops were harvested for hay in 1988 and 1989 only.

Within each field, the dynamics of all four processes are modelled jointly as follows.

$$\begin{aligned} [Y(i, t + \Delta t) | Y(i, t), \Theta] &= [Y_C(i, t + \Delta t) | Y_W(i, t + \Delta t), Y_G(i, t + \Delta t), Y_P(i, t + \Delta t), Y(i, t), \Theta] \\ &\quad \times [Y_W(i, t + \Delta t), Y_G(i, t + \Delta t), Y_P(i, t + \Delta t) | Y(i, t), \Theta] \\ &= [Y_C(i, t + \Delta t) | I_C(i, t + \Delta t), Y_C(i, t), \Theta] \\ &\quad \times [Y_G(i, t + \Delta t) | Y_G(i, t), \Theta] \\ &\quad \times [Y_W(i, t + \Delta t) | Y_G(i, t + \Delta t), \Theta] \\ &\quad \times [Y_P(i, t + \Delta t) | Y_P(i, t), \Theta] \end{aligned}$$

These four components are defined, respectively, by Equations 1, 2, 3, and 4. We model the dynamics of all processes independently for each field, that is,

$$[Y(., t + \Delta t) | Y(., t), \Theta] = \prod_{i=1}^3 [Y(i, t + \Delta t) | Y(i, t), \Theta]. \quad (6)$$

We make the initial conditions of the four processes a part of $[\Theta]$, and so products of terms such as those represented in Equation 6 give $[Y|\Theta]$, our overall process model.

3.4. Data Model

The data model captures the uncertainty in the observations, given the state of the underlying processes and parameters. The components of the data model are specified to be conditionally independent, namely,

$$\begin{aligned} [Z|Y, \Theta] &= [Z_C|Y, \Theta] \times [Z_G|Y, \Theta] \times [Z_W|Y, \Theta] \times [Z_P|Y, \Theta] \\ &= [Z_C|Y_C, \Theta] \times [Z_G|Y_G, \Theta] \times [Z_W|Y_W, \Theta] \times [Z_P|Y_P, \Theta]. \end{aligned}$$

For each process, we invoke conditional independence across time and across fields; for example, in the case of the SOC process,

$$[Z_C|Y_C, \Theta] = \prod_{i=1}^3 \prod_t [Z_C(i, t)|Y_C(i, t), \Theta].$$

Finally, we use multiplicative error structures for observations of each process in field i at time t . We write the components of the data model as follows:

$$\text{given } Y_C(i, t) = y_C(i, t), \quad Z_C(i, t) \sim \text{LogN}(\log y_C(i, t), \sigma_{\epsilon_C}^2); \quad (7)$$

$$\text{given } Y_G(i, t) = y_G(i, t), \quad Z_G(i, t) \sim \text{LogN}(\log y_G(i, t), \sigma_{\epsilon_G}^2); \quad (8)$$

$$\text{given } Y_W(i, t) = y_W(i, t), \quad Z_W(i, t) \sim \text{LogN}(\log y_W(i, t), \sigma_{\epsilon_W}^2); \text{ and} \quad (9)$$

$$\text{given } Y_P(i, t) = y_P(i, t), \quad Z_P(i, t) \sim \text{LogN}(\log y_P(i, t), \sigma_{\epsilon_P}^2), . \quad (10)$$

3.5. Parameter Model

The data and process models defined above involve 28 independent parameters, which includes the initial conditions of seven processes at $t = 1978$. Table 1 describes the

priors for each of these parameters and our parameter model, $[\Theta]$, is the product of the probability density functions listed therein. These priors and hyperpriors could be considered informative, weakly informative, or uninformative; and they were set according to the current state of scientific knowledge about the 28 parameters. The state of scientific knowledge was codified into priors by soliciting typical values and associated confidence intervals from those of us with expertise in soil science. Our elicitation process was based around 95% confidence intervals, other elicitation processes such as Sheffield Elicitation Framework (Oakley and O’Hagan, 2010) work in a more structured manner with intervals based on quartiles or tertiles; such approaches, for example, have been applied to models of the terrestrial carbon cycle (O’Hagan, 2012).

The four variance parameter values from Equations 7 - 10 were set based on analysis of independent data sets. In the case of $\sigma_{\epsilon_C}^2$, observations of soil carbon from four paddocks (areas 20 to 66 hectares) in the Tarlee region, collected as part of Australia’s National Soil Carbon Research Programme (Sanderman *et al.*, 2011), were used to evaluate the spatial variability of SOC. Data from each paddock included 10 observations within a 25m by 25m grid as well as 10 observations spread uniformly across the paddock. Robust estimates of the variogram were found based on these samples (Cressie and Hawkins, 1980). The nugget variance, which is a combination of measurement error variance and microscale variation, could not be disentangled into its components and so the estimate of nugget variance was used for $\sigma_{\epsilon_C}^2$. For the remaining variance parameters, independent field trial data from the region were analysed using mixed effects models for block and treatment effects. The estimates of measurement error variance from these models were used for $\sigma_{\epsilon_G}^2$, $\sigma_{\epsilon_W}^2$, and $\sigma_{\epsilon_P}^2$.

[Table 1 about here.]

3.6. Obtaining the Posterior Distribution

Our interest lies in drawing samples from the posterior distribution $[Y_C, \Theta | Z = z]$. We achieved this by drawing samples from $[Y, \Theta | Z = z]$ and retaining the components pertaining to the SOC process and its parameters, that is, Y_C and Θ . The approach we used to draw these samples is known as particle marginal Metropolis-Hastings (PMMH), one of a number of particle Markov chain Monte Carlo (PMCMC) methods (e.g., Andrieu *et al.*, 2010). We provide a detailed overview of our particle-sampling approach in Appendix A.

In brief, the posterior distribution $[Y, \Theta | Z]$ can be factored into two components, namely, $[Y, \Theta | Z] = [Y | \Theta, Z] \times [\Theta | Z]$. The idea of PMMH is to use a Metropolis-Hastings algorithm (Metropolis *et al.*, 1953; Hastings, 1970) to draw samples of the parameters from $[\Theta | Z]$ using a particle filter (Doucet *et al.*, 2001) to estimate the marginal likelihood $[Z | \Theta]$ required to achieve this; then that particle filter is reused to draw a sample of the state process from $[Y | \Theta, Z]$. The simplest particle filter, the bootstrap particle filter (Gordon *et al.*, 1993), provides marginal likelihood and state estimates that work sufficiently well for the model and dataset, so that it has been unnecessary to explore more elaborate schemes. The Metropolis-Hastings algorithm requires the selection of an appropriate proposal distribution for generating each candidate parameter; these proposal distributions are given in Table 2. As is usual practice, we tuned the proposal by using a few pilot runs until a reasonable acceptance rate was achieved. The rule-of-thumb acceptance rate for as many parameters as we have here is 23% (Gelman *et al.*, 1996). However, because PMMH uses a likelihood estimator, it is appropriate to aim lower than this. The final run has an acceptance rate of 16%, which we find gives good mixing. We drew 500,000 such samples, of which we discarded the initial 50,000 as burn-in.

[Table 2 about here.]

An advantage of the PMMH method, using the bootstrap particle filter, is that it requires only that the process model can be simulated. That is, it requires that samples can be

drawn from the conditional density, $p(Y(., t) | Y(., t - \Delta t) = y(., t - \Delta t), \Theta = \theta)$, but it does not require that the conditional density function be evaluated. This allows for quite complex process models to be handled. In our application, soil carbon models were not developed with the evaluation of the conditional density function in mind, and indeed requiring so may impose unjustifiable constraints on the process model. Hence we were able to compute based on the posterior distribution without compromising our model of SOC dynamics.

The model is specified using the LibBi modelling language (Murray, 2013, www.LibBi.org), and LibBi is also used to sample from the posterior distribution using PMMH. The results of our implementation of PMMH to the BHM that incorporates our physical-statistical model and the Tarlee dataset are given in the next section.

4. INFERENCE ON THE HIDDEN PROCESS AND ITS FUNCTIONALS

The ultimate goal of this project is to quantify the uncertainty associated with statements about the amount of carbon that lies within the soil in the landscape, about how much of that SOC was added through the management practices of the landowner, and about the parameters driving the sequestration of carbon. We have chosen the one-pool SOC model to illustrate inference on the (hidden) carbon sequestration and decay processes and their associated parameters. In this section, we show how the BHM approach allows us to predict the amount of SOC present under various management practices as well as obtain associated credible intervals.

The carbon credits associated with Australia's Carbon Farming Initiative (CFI) are tied to changes in SOC over time, so both farmers and government are interested in the mass of SOC added over say a twenty-year period, depending on management practices. The uncertainty of our estimate can be quantified in many ways, through a 95% credible interval for example, or through the estimated probability of functionals of interest. We might estimate a change of

4 t ha⁻¹ during a twenty-year period, but that number is not certain. A landowner's decision about committing to Australia's CFI is a personal one, depending on that individual's risk tolerance. Consequently, quantifying the risk through the probability that at least k t ha⁻¹ is sequestered, for various k , would be a highly informative tool.

All of these inference tools can be developed under a BHM framework. Recall from Section 3.6 that we are sampling from the posterior distribution using PMMH. Consequently, the posterior distribution, $[Y, \Theta|Z = z]$, is represented by N samples $\{(y^n, \theta^n) : n = 1, \dots, N\}$ drawn from $[Y, \Theta|Z = z]$, with superscripts used for indexing each sample. The samples can be used to estimate the posterior expectation of any function $f(Y, \Theta)$:

$$\mathbb{E}(f(Y, \Theta)|Z = z) \approx \hat{\mathbb{E}}(f(Y, \Theta)|Z = z) \equiv \frac{1}{N} \sum_{n=1}^N f(y^n, \theta^n). \quad (11)$$

The only limit to the accuracy of such estimates is the size of N . The error is made negligible with N sufficiently large.

4.1. Inference on SOC Dynamics

Consider the change in SOC to field i between 1978 and a subsequent year t :

$$f_1(Y, \Theta) \equiv Y_C(i, t) - Y_C(i, 1978).$$

Making inference on f_1 depends on the posterior distribution, $[Y_C(i, t) - Y_C(i, 1978)|Z = z]$.

For example, a common (Bayes) estimate of f_1 is,

$$\hat{f}_1 \equiv \mathbb{E}(Y_C(i, t) - Y_C(i, 1978)|Z = z),$$

and a measure of uncertainty associated with this estimate is the posterior variance, $\text{var}(Y_C(i, t) - Y_C(i, 1978)|Z = z)$. A 95% credible interval for f_1 can be found by computing the 2.5th percentile and the 97.5th percentile of the posterior distribution and defining them

to be respectively, the lower and upper limits of the interval. Then the probability that the true value f_1 is contained in the interval is 0.95.

Of course, f_1 could be negative, so suppose we wish to make inference on

$$f_2(Y, \Theta) \equiv I(Y_C(i, t) - Y_C(i, 1978) > k); k \geq 0,$$

where I is the indicator function that takes the value 1 when its argument is true (i.e., the carbon content has increased by more than k units) and 0 when its argument is false. This can be interpreted as the risk of sequestering k t ha⁻¹ or more (e.g., $k = 4$). The Bayes estimate is

$$\hat{f}_2(Y, \Theta) = \Pr(Y_C(i, t) - Y_C(i, 1978) > k | Z = z).$$

The competing risk of sequestering no carbon at all is based on

$$f_3(Y, \Theta) \equiv I((Y_C(i, t) - Y_C(i, 1978) < 0),$$

resulting in

$$\hat{f}_3(t) \equiv \Pr(Y_C(i, t) - Y_C(i, 1978) < 0 | Z = z).$$

Clearly, any event of interest can be estimated by the posterior probability of that event.

For the three fields in the Tarlee trial (Field 1, Field 2, and Field 3) our estimates of the amount of carbon added between 1978 and 1997 are $\hat{f}_1 = -11.4$ t ha⁻¹, -10.8 t ha⁻¹, and 3.25 t ha⁻¹ respectively, with negative values indicating that the first two fields are expected to lose carbon over a twenty-year period. The 95% credible intervals for the amount of carbon added are $(-24.1, 2.30)$, $(-23.6, 2.75)$, and $(-12.5, 18.6)$, respectively.

In terms of the functionals f_2 and f_3 , we find that Fields 1 and 2 have only a 2% chance

each of sequestering more than 4 t ha^{-1} and a 95% chance of sequestering no carbon or losing carbon. Field 3 (which is under a wheat-pasture rotation) has a 46% chance of sequestering more than 4 t ha^{-1} and a 33% chance of sequestering no carbon or losing carbon.

Other visual summaries of the posterior distribution may be of interest in highlighting what can be learned in a BHM framework. In Figure 1, we plot 200 samples drawn from the posterior distribution of the SOC process for each field. Crop-yield data collection ceased at Tarlee in 1996 but we can run the model forward under the same management practices for each field (forecasting). Figure 1 also highlights the 2.5th, 25th, 50th, 75th, and 97.5th percentiles for the SOC process of each field based on all 450,000 draws from the posterior distribution.

[Figure 1 about here.]

We can also make inference about the carbon inputs to the soil each year, which are directly informed by the measurements of crop production from 1979 to 1996; see Figure 2.

[Figure 2 about here.]

Information about the initial SOC level, the final SOC level, and the change in SOC level can be gleaned from the posterior distributions of, respectively, $Y_C(i, 1978)$, $Y_C(i, 1997)$, and $(Y_C(i, 1997) - Y_C(i, 1978))$. We highlight histograms of the change in SOC level for each field in the top row of Figure 3. The bottom row of Figure 3 highlights the cumulative distribution function for the change in SOC for each field. The dotted lines indicate the chance of sequestering zero or fewer tonnes of carbon per hectare for the three fields.

[Figure 3 about here.]

Finally, for each model parameter, we can compare the prior distribution with a histogram of samples drawn from the posterior distribution to highlight what we have learned about the parameters; see Figure 4, which includes eighteen such plots for the key model parameters.

[Figure 4 about here.]

Based on Figure 4, it is clear that we learn quite a lot about some parameters, namely the initial SOC conditions in the three fields, the harvest index, the mean grain yield, and various variance parameters. It is also clear that we learn little new about other parameters, something that is evident when the posterior and prior are very similar. Examples of such parameters include c , p , and the root-to-shoot ratios. That there is little to learn about these parameters should come as no surprise when one considers the form of the data. Nevertheless, these parameters remain important and useful components of the model; because they are well known to the agricultural research community, we are able to specify priors for them quite easily.

Another aspect of Figure 4 is worth pointing out. The posterior for K indicates that the rate of decomposition of SOC is slower than what we specified in our prior. The prior mean for K is 0.067 year^{-1} , which should be compared to the posterior mean for K equal to 0.061 year^{-1} .

The posterior distributions noted in Figure 4 are based on a field trial carried out on red-brown earth soil in South Australia. The information captured in these posterior distributions could be used when modelling soil carbon dynamics within similar soils under similar climatic conditions.

4.2. Forecasting and Counterfactuals

One of the benefits of the BHM framework is the ability to make inference with quantified uncertainty about future scenarios, that is, to forecast SOC stocks in the absence of observations and to do so under scenarios that were not part of the initial data collection (counterfactuals). Figures 1 and 2 highlight forecasted SOC dynamics without changes to the management practices. To implement these forecasts, we ran the posterior model forward from 1998 to 2018, a period for which no observations are available. Hence there is large

uncertainty around the SOC inputs and, as a result, large uncertainty on SOC levels for each field. Figures 5 highlights forecasted SOC dynamics for Fields 2 and 3 under changed management practices, pasture-pasture for Field 2 and wheat-wheat-pasture for Field 3.

[Figure 5 about here.]

Together, the panels of Figures 1, 2, and 5 tell a coherent story about SOC dynamics. The most obvious feature across the three figures, and one that is to be expected, is increased uncertainty in SOC levels the further into the future we forecast. Under continuous wheat (Figure 1A) we forecast the gradual sequestration of carbon, which rises from a median value of 30.8 t ha^{-1} in 1997 to about 35.7 t ha^{-1} by 2018. We also see that the rate at which SOC is building up within the soil is slowing over time and it is not hard to imagine the SOC levels eventually levelling out at perhaps 41 t ha^{-1} .

Contrast this with what happens under a wheat-fallow rotation in Field 2 (Figure 1B). Here we see continued loss of carbon from the soil. The loss of carbon levels off slightly over time, reflecting the exponential decay that is central to the model. We also see less variability in the SOC dynamics in Figure 1B, because in fallow years there is no randomness in the inputs (the inputs are identically zero).

In Figure 1C, the wheat-pasture rotation has resulted in high SOC levels rising from 44.1 t ha^{-1} to around 45 t ha^{-1} between 1997 and 2018. Individual dynamics are quite variable but an examination of the median forecast shows that during years when wheat is sown, the amount of carbon added to the soil fails to offset the carbon lost to decay, and we see a slight dip in median SOC levels. In subsequent years, the pasture crop adds more than the amount that decays and we see a slight rise in median SOC levels. Over time the median SOC level rises, showing net sequestration of carbon.

The median carbon sequestered by wheat is 1.9 t ha^{-1} , and the median carbon sequestered by pasture is 2.9 t ha^{-1} (Figure 2). Seeing the behavior in Figure 1, we might ask what would happen to SOC levels under alternative management practices such as continuous

pasture or wheat-wheat-pasture. These practices are explored in Figure 5 for Fields 2 and 3, respectively.

Under continuous pasture (Figure 5A) we are clearly adding more SOC than we are losing to decay every year (hence the increase in SOC values). In the case of wheat-wheat-pasture (Figure 5B), like with wheat-pasture in Figure 1C, the median SOC levels are not changing very much from one year to the next. Field 3 was under a wheat-pasture rotation up until 1997 and, as such, it is starting out with the highest SOC levels at that point in time with levels at approximately 44.1 t ha^{-1} . Figure 5B shows subtle oscillations in the median SOC levels from 1998 onwards, where two years of decline in median SOC are followed by one year of growth, corresponding to two years of wheat followed by one year of pasture. As we are above the putative equilibrium hinted at in Figure 1A, the wheat crops are failing to add more carbon than what is decaying and so, in terms of median-level dynamics, we forecast two years of overall loss followed by one year of significant gain under pasture. A closer examination of the median levels shows overall slow growth in median SOC levels under this wheat-wheat-pasture strategy, though this is not immediately obvious when looking at Figure 5B.

5. DISCUSSION AND CONCLUSIONS

In this paper, we incorporate a one-pool model of carbon dynamics into a physical-statistical model that is part of a BHM framework and that allows us to think critically and conditionally about the data, the process, and the parameters, and it is a natural way to engage with researchers in a science-team setting. The model structure created in Section 3 has physical decay at its core (Equation 1) yet most of the modelling effort is placed on the input processes, reflecting the fact that most of the data is informative about the inputs (Equations 2–5). Our success with this proof-of-concept model so far is due in no small part

to the modelling flexibility provided by the LibBi modelling language and the computational efficiency of sampling from the posterior distributions once that model has been properly set up. As is evident in Section 4, the ability to sample from the posterior distribution of the process and parameters, conditional on the data, enables one to answer any question related to the process and parameters in the BHM.

Quite complex models can be tackled in this manner, and our next research goal is to continue development of the model to account for multiple pools of carbon that are decaying and interacting within the soil on a finer time scale. Currently we capture this heterogeneity as a component of the model error, σ_{η}^2 . The fact that multi-pool models based on conceptual pools are calibrated using data from measurable pools indicates some uncertainty within the modelling community. Ultimately we aim to represent the soil organic continuum directly, including interactions between SOC and other soil constituents (Lehmann *et al.*, 2007). To achieve this in future work, the availability of data from which we can learn about different aspects of the models and continue with model development is critical. While the Tarlee dataset does have information about individual carbon pools, it is limited in that spatial effects cannot be modelled, and the temporal resolution is coarse relative to a scientifically more meaningful monthly time-step used by RothC, Century, and FullCAM.

The availability of appropriate data is key to future development of these models, but the reverse of this statement is also true. Models such as this can highlight gaps in our knowledge of SOC dynamics, and field trials can be designed to directly address those gaps. One such gap will centre around what effect given may have on SOC stocks. Climate information such as rainfall and temperature are used to modify the rate of decay K , and historical weather information can be used to extract information about how the rate of decay relates to changes in temperature and rainfall. But climate-change scenarios go beyond the annual variation one sees in historical weather information, in which case the right kinds of controlled experiments within laboratory settings will be required to more accurately model SOC dynamics.

Re-thinking soil carbon modelling presents several challenges. For soil scientists, these challenges include thinking conditionally and probabilistically about components of the models and model structure, and moving from deterministic models with fixed parameters to thinking about how to represent knowledge about model parameters in terms of prior probability distributions. For statisticians, the challenge becomes one of translation, converting conditional-probability models written in terms of mathematics into a model specification where posterior-probability sampling can be implemented in a computationally efficient manner. These are challenges we face in all forms of physical-statistical modelling, and overcoming them paves the way to insights and scientific inferences based on all available data and knowledge.

Acknowledgements

The authors thank Murray Unkovich for data on field measurement for pasture and wheat, CSIRO's Sustainable Agriculture Flagship for funding this research, CSIRO's Flagship Collaboration Fund for supporting Prof Cressie's involvement, and CSIRO's Computational and Simulation Sciences Transformational Capability Platform for funding Lawrence Murray's involvement with this research.

REFERENCES

- Andrieu C, Doucet A, Holenstein R, 2010. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **72**(3): 269–342.
- Baldock J, Skjemstad J, 1999. Soil organic carbon/soil organic matter. In Peverill K, Sparrow L, Reuter D (eds.), *Soil Analysis: An Interpretation Manual*, CSIRO Publishing, Collingswood, Vic, Australia, 159–170.
- Bayes T, Price R, 1763. An essay towards solving a problem in the doctrine of chances. by the Late Rev. Mr. Bayes, F.R.S. Communicated by Mr. Price, in a letter to John Canton, A.M. F.R.S. *Philosophical Transactions* **53**: 370–418.

-
- Bosatta E, Ågren GI, 1985. Theoretical analysis of decomposition of heterogeneous substrates. *Soil Biology and Biochemistry* **17**(5): 601–610.
- Cressie N, Hawkins DM, 1980. Robust estimation of the variogram: I. *Journal of the International Association for Mathematical Geology* **12**(2): 115–125.
- Cressie N, Wikle CK, 2011. *Statistics for Spatio-Temporal Data*. John Wiley and Sons., Hoboken, NJ, 624 pp.
- Del Moral P, 2004. *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. Springer, New York, 574 pp.
- Doucet A, de Freitas N, Gordon N (eds.), 2001. *Sequential Monte Carlo Methods in Practice*. Springer, New York.
- Eswaran H, Van den Berg E, Reich P, Kimble J, 1995. Global soil C resources. In Lal R, Kimble J, Levine E, Stewart B (eds.), *Soils and Global Change*, Lewis Publishers, Boca Raton, FL, USA, 27–43.
- Feng Y, 2009. Fundamental considerations of soil organic carbon dynamics: A new theoretical framework. *Soil Science* **174**(9): 467–481.
- Gelman A, Roberts GO, Gilks WR, 1996. Efficient Metropolis jumping rules. In Bernardo JM, Berger JO, Dawid AP, Smith AFM (eds.), *Bayesian Statistics 5, Proceedings of the Fifth Valencia International Meeting, June 5-9, 1994*, Oxford University Press, 599–607.
- Gordon N, Salmond D, Smith AFM, 1993. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *Radar and Signal Processing, IEE Proceedings F* **140**(2): 107–113.
- Gottschalk P, Smith J, Wattenbach M, Bellarby J, Stehfest E, Arnell N, Osborn TJ, Jones C, Smith P, 2012. How will organic carbon stocks in mineral soils evolve under future climate? Global projections using RothC for a range of climate change scenarios. *Biogeosciences* **9**(8): 3151–3171.
- Hastings WK, 1970. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57**(1): 97–109.
- Houghton RA, 2003. The contemporary carbon cycle. *Treatise on Geochemistry* **8**: 473–513.
- Hoyle F, Baldock J, Murphy D, 2011. Soil organic carbon - role in rainfed farming systems: with particular reference to Australian conditions. In Tow P, Cooper I, Partridge I, Birch C (eds.), *Rainfed Farming Systems*, Springer, New York, 339–361.
- Jenkinson D, Hart P, Rayner J, Parry L, 1987. Modelling the turnover of organic matter in long-term experiments at Rothamsted. *Intecol Bulletin* **15**: 1–8.

- Jobbágy E, Jackson R, 2000. The vertical distribution of soil organic carbon and its relation to climate and vegetation. *Ecological Applications* **10**(2): 423–436.
- Jones D, Nguyen C, Finlay R, 2009. Carbon flow in the rhizosphere: carbon trading at the soil-root interface. *Plant and Soil* **321**(1-2): 5–33.
- Jones J, Koo J, Naab J, Bostick W, Traore S, Graham W, 2007. Integrating stochastic models and in situ sampling for monitoring soil carbon sequestration. *Agricultural Systems* **94**(1): 52–62.
- Juston J, Andrén O, Kätterer T, Jansson PE, 2010. Uncertainty analyses for calibrating a soil carbon balance model to agricultural field trial data in Sweden and Kenya. *Ecological Modelling* **221**(16): 1880–1888.
- Koo J, Bostick WM, Naab JB, Jones JW, Graham WD, Gijsman AJ, 2007. Estimating soil carbon in agricultural systems using ensemble Kalman filter and DSSAT-CENTURY. *Transactions of the ASABE* **50**(5): 1851–1865.
- Lal R, 2004. Soil carbon sequestration impacts on global climate change and food security. *Science* **304**(5677): 1623–1627.
- Lehmann J, Kinyangi J, Solomon D, 2007. Organic matter stabilization in soil microaggregates: implications from spatial heterogeneity of organic carbon contents and carbon forms. *Biogeochemistry* **85**(1): 45–57.
- Lehmann J, Skjemstad J, Sohi S, Carter J, Barson M, Falloon P, Coleman K, Woodbury P, Krull E, 2008. Australian climate-carbon cycle feedback reduced by soil black carbon. *Nature Geoscience* **1**(12): 832–835.
- Liu DL, Chan KY, Conyers MK, 2009. Simulation of soil organic carbon under different tillage and stubble management practices using the Rothamsted carbon model. *Soil and Tillage Research* **104**(1): 65–73.
- Liu DL, Chan KY, Conyers MK, Li G, Poile GJ, 2011. Simulation of soil organic carbon dynamics under different pasture managements using the RothC carbon model. *Geoderma* **165**(1): 69–77.
- Luo Z, Wang E, Sun OJ, 2010. Soil carbon change and its responses to agricultural practices in Australian agro-ecosystems: A review and synthesis. *Geoderma* **155**(3-4): 211–223.
- Luo Z, Wang E, Sun OJ, Smith CJ, Probert ME, 2011. Modeling long-term soil carbon dynamics and sequestration potential in semi-arid agro-ecosystems. *Agricultural and Forest Meteorology* **151**(12): 1529–1544.
- Meermans J, Martin M, Laccarce E, Orton T, De Baets S, Gourrat M, Saby N, Wetterlind J, Bispo A, Quine T, Arrouays D, 2013. Estimation of soil carbon input in France: An inverse modelling approach. *Pedosphere* **23**(4): 422–436.

-
- Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E, 1953. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics* **21**(6): 1087–1092.
- Murray LM, 2013. Bayesian state-space modelling on high-performance hardware using LibBi, in review.
- Oakley J, O’Hagan A, 2010. SHELF: the Sheffield elicitation framework (version 2.0). *School of Mathematics and Statistics, University of Sheffield, Sheffield, United Kingdom*. Available from <http://tonyohagan.co.uk/shelf> (accessed October 2011) .
- O’Hagan A, 2012. Probabilistic uncertainty specification: Overview, elaboration techniques and their application to a mechanistic model of carbon flux. *Environmental Modelling & Software* **36**: 35–48.
- Parton W, Stewart J, Cole C, 1988. Dynamics of C, N, P and S in grassland soils: a model. *Biogeochemistry* **5**(1): 109–131.
- Paul KI, Polglase PJ, Richards GP, 2003. Sensitivity analysis of predicted change in soil carbon following afforestation. *Ecological Modelling* **164**(2–3): 137–152.
- Post J, Hattermann FF, Krysanova V, Suckow F, 2008. Parameter and input data uncertainty estimation for the assessment of long-term soil organic carbon dynamics. *Environmental Modelling & Software* **23**(2): 125–138.
- Powlson D, 2005. Climatology: Will soil amplify climate change? *Nature* **433**(7023): 204–205.
- Rahn KH, Butterbach-Bahl K, Werner C, 2011. Selection of likelihood parameters for complex models determines the effectiveness of Bayesian calibration. *Ecological Informatics* **6**(6): 333–340.
- Refsgaard JC, van der Sluijs JP, Højberg AL, Vanrolleghem PA, 2007. Uncertainty in the environmental modelling process - a framework and guidance. *Environmental Modelling & Software* **22**(11): 1543–1556.
- Richards GP, Evans DM, 2004. Development of a carbon accounting model (FullCAM Vers. 1.0) for the Australian continent. *Australian Forestry* **67**(4): 277–283.
- Richardson A, Williams M, Hollinger D, Moore D, Dail D, Davidson E, Scott N, Evans R, Hughes H, Lee J, Rodrigues C, Savage K, 2010. Estimating parameters of a forest ecosystem C model with measurements of stocks and fluxes as joint constraints. *Oecologia* **164**(1): 25–40.
- Robertson F, Nash D, 2013. Limited potential for soil carbon accumulation using current cropping practices in Victoria, Australia. *Agriculture, Ecosystems & Environment* **165**: 130–140.
- Sakurai G, Jomura M, Yonemura S, Iizumi T, Shirato Y, Yokozawa M, 2012. Inversely estimating temperature sensitivity of soil carbon decomposition by assimilating a turnover model and long-term field data. *Soil Biology and Biochemistry* **46**: 191–199.

-
- Sanderman J, Baldock J, Hawke B, Macdonald L, Puccini A, Szarvas S, 2011. National soil carbon research programme: field and laboratory methodologies. Report EP117207, CSIRO Land and Water, Waite Campus, Urrbrae SA 5064.
- Sanderman J, Baldock JA, 2010. Accounting for soil carbon sequestration in national inventories: a soil scientist's perspective. *Environmental Research Letters* **5**(3): 034003.
- Sanderman J, Chappell A, 2013. Uncertainty in soil carbon accounting due to unrecognized soil erosion. *Global Change Biology* **19**(1): 264–272.
- Schädel C, Luo Y, Evans RD, Fei S, Schaeffer SM, 2013. Separating soil CO₂ efflux into C-pool-specific decay rates via inverse analysis of soil incubation data. *Oecologia* **171**(3): 721–732.
- Scharnagl B, Vrugt JA, Vereecken H, Herbst M, 2010. Information content of incubation experiments for inverse estimation of pools in the Rothamsted carbon model: a Bayesian perspective. *Biogeosciences* **7**(2): 763–776.
- Schultz J, 1995. Crop production in a rotation trial at Tarlee, South Australia. *Australian Journal of Experimental Agriculture* **35**(7): 865–876.
- Senapati N, Smith P, Wilson B, Yeluripati JB, Daniel H, Lockwood P, Ghosh S, 2013. Projections of changes in grassland soil organic carbon under climate change are relatively insensitive to methods of model initialization. *European Journal of Soil Science* **64**(2): 229–238.
- Sierra CA, Müller M, Trumbore SE, 2012. Models of soil organic matter decomposition: the SoilR package, version 1.0. *Geoscientific Model Development* **5**(4): 1045–1060.
- Skjemstad JO, Spouncer LR, Cowie B, Swift RS, 2004. Calibration of the Rothamsted organic carbon turnover model (RothC ver. 26.3), using measurable soil organic carbon pools. *Australian Journal of Soil Research* **42**(1): 79–88.
- Smith W, Grant B, Desjardins R, Qian B, Hutchinson J, Gameda S, 2009. Potential impact of climate change on carbon in agricultural soils in Canada 2000 – 2099. *Climatic Change* **93**(3-4): 319–333.
- Soil Survey Staff, 1975. *Soil Taxonomy: A Basic System of Soil Classification for Making and Interpreting Soil Surveys*. Handbook 436, U.S. Dept. of Agric., 754 pp.
- Stamati FE, Nikolaidis NP, Schnoor JL, 2013. Modeling topsoil carbon sequestration in two contrasting crop production to set-aside conversions with RothC - calibration issues and uncertainty analysis. *Agriculture, Ecosystems & Environment* **165**: 190–200.
- Todd-Brown KEO, Randerson JT, Post WM, Hoffman FM, Tarnocai C, Schuur EAG, Allison SD, 2013.

- Causes of variation in soil carbon simulations from CMIP5 Earth system models and comparison with observations. *Biogeosciences* **10**(3): 1717–1736.
- Wan Y, Lin E, Xiong W, Li Y, Guo L, 2011. Modeling the impact of climate change on soil organic carbon stock in upland soils in the 21st century in China. *Agriculture, Ecosystems & Environment* **141**(1-2): 23–31.
- Wang G, Chen S, 2013. Evaluation of a soil greenhouse gas emission model based on Bayesian inference and MCMC: Parameter identifiability and equifinality. *Ecological Modelling* **253**: 107–116.
- Wang J, Lu C, Xu M, Zhu P, Huang S, Zhang W, Peng C, Chen X, Wu L, 2013. Soil organic carbon sequestration under different fertilizer regimes in north and northeast China: RothC simulation. *Soil Use Manage* **29**(2): 182–190.
- Xu T, White L, Hui D, Luo Y, 2006. Probabilistic inversion of a terrestrial ecosystem model: Analysis of uncertainty in parameter estimation and model prediction. *Global Biogeochem. Cycles* **20**(2): GB2007.
- Yeluripati JB, van Oijen M, Wattenbach M, Neftel A, Ammann A, Parton W, Smith P, 2009. Bayesian calibration as a tool for initialising the carbon pools of dynamic soil models. *Soil Biology and Biochemistry* **41**(12): 2579–2583.
- Zimmermann M, Leifeld J, Schmidt MWI, Smith P, Fuhrer J, 2007. Measured soil organic matter fractions can be related to pools in the RothC model. *European Journal of Soil Science* **58**(3): 658–667.

APPENDIX

A. PARTICLE MARGINAL METROPOLIS-HASTINGS

A brief overview of the particle marginal Metropolis-Hastings (PMMH) method (Andrieu *et al.*, 2010) is provided here, with the aim of motivating its use for this particular model and dataset. In doing so, it is conducive to work with the probability density function corresponding to the posterior distribution, $p(Y, \Theta|Z) \equiv [Y, \Theta|Z]$. This may be factorised as

$$p(Y, \Theta|Z) = p(Y|\Theta, Z)p(\Theta|Z). \quad (12)$$

The idea of PMMH is then to use a Metropolis-Hastings algorithm (Metropolis *et al.*, 1953; Hastings, 1970) to draw samples of the parameters from $[\Theta|Z]$, using a particle filter (Doucet *et al.*, 2001) to estimate the marginal likelihood $[Z|\Theta]$ required to achieve this, and reusing that particle filter to draw a sample of the state process from $[Y|\Theta, Z]$.

A.1. Particle Filter

For a given $\Theta = \theta$ and the Tarlee dataset $Z = z$, the bootstrap particle filter (Gordon *et al.*, 1993) proceeds as follows, where N_Y denotes the number of state samples, and superscripts index these samples:

```

1  for each  $n \in \{1, \dots, N_Y\}$ 
2      draw  $y^n(., 0) \sim p(Y(., 0)|\theta)$ 
3      set  $\omega^n(0) \leftarrow 1/N_Y$ 
4  for  $t = 1, \dots, T$ 
5      for each  $n \in \{1, \dots, N_Y\}$ 
6          draw  $a^n(t) \sim \text{MULTINOMIAL}(\omega^1(t-1), \dots, \omega^{N_Y}(t-1))$ 
7          draw  $y^n(., t) \sim p(Y(., t)|y^{a^n(t)}(., t-1), \theta)$ 
8          set  $\omega^n(t) \leftarrow p(z(., t)|y^n(., t), \theta)$ 

```

The MULTINOMIAL on line 6 refers to a multinomial distribution that gives an index $n \in \{1, \dots, N_Y\}$ with the probability of drawing n being proportional to $\omega^n(t-1)$.

The particle filter has two interesting properties that are used below. First, at the conclusion of the algorithm, one can recursively draw a single trajectory y^* from the particle filter as follows:

- 1 draw $b^*(T) \sim \text{MULTINOMIAL}(\omega^1(T), \dots, \omega^{N_Y}(T))$
- 2 set $y^*(\cdot, T) \leftarrow y^{b^*(T)}(\cdot, T)$
- 3 **for** $t = T, \dots, 1$
- 4 set $b^*(t-1) \leftarrow a^{b^*(t)}(t)$
- 5 set $y^*(t-1) \leftarrow y^{b^*(t-1)}(t-1)$

The trajectory y^* , drawn in this way, is distributed according to $p(Y|\theta, z)$ (Andrieu *et al.*, 2010). Second, a simple function of the weights gives an unbiased estimate of the marginal likelihood (Del Moral, 2004):

$$p(z|\theta) \approx \hat{p}(z|\theta) \equiv \prod_{t=1}^T \left(\frac{1}{N_Y} \sum_{n=1}^{N_Y} \omega^n(t) \right). \tag{13}$$

A.2. Metropolis-Hastings Algorithm

Sampling from the second factor in (12) using a Metropolis-Hastings (Metropolis *et al.*, 1953; Hastings, 1970) algorithm proceeds as follows, where N_Θ is the number of parameter samples to draw and $q(\Theta'|\theta)$ is a proposal distribution (details follow):

```

1  draw  $\theta^0 \sim p(\Theta)$ 
2  set  $l^0 \leftarrow p(z|\theta^0)$ 
3  for  $n = 1, \dots, N_{\Theta}$ 
4      draw  $\theta' \sim q(\Theta|\theta^{n-1})$  // propose
5      set  $l' \leftarrow p(z|\theta')$ 
6      draw  $\alpha \sim \text{UNIFORM}([0, 1])$ 
7      if  $\alpha \leq (l'p(\theta')q(\theta^{n-1}|\theta')) / (l^{n-1}p(\theta^{n-1})q(\theta'|\theta^{n-1}))$ 
8          set  $\theta^n \leftarrow \theta'$  // accept proposal
9          set  $l^n \leftarrow l'$ 
10     else
11         set  $\theta^n \leftarrow \theta^{n-1}$  // reject proposal
12         set  $l^n \leftarrow l^{n-1}$ 

```

Where it occurs, the marginal likelihood, $p(z|\theta)$, cannot be readily computed for the soil carbon model. However, by running a particle filter, the unbiased estimate $\hat{p}(z|\theta)$ in (13) can be used as a substitute. When a proposed θ' is accepted, a trajectory $y' = y^*$ can be drawn from the same particle filter that was used to estimate the likelihood of θ' , and (y', θ') is then a complete sample from the posterior distribution. The correctness of the approach is seen by considering the extended space of all auxiliary random variables upon which the particle filter operates, and observing that the posterior distribution is a marginal of this. The proof and further details, are given in Andrieu *et al.* (2010).

FIGURES

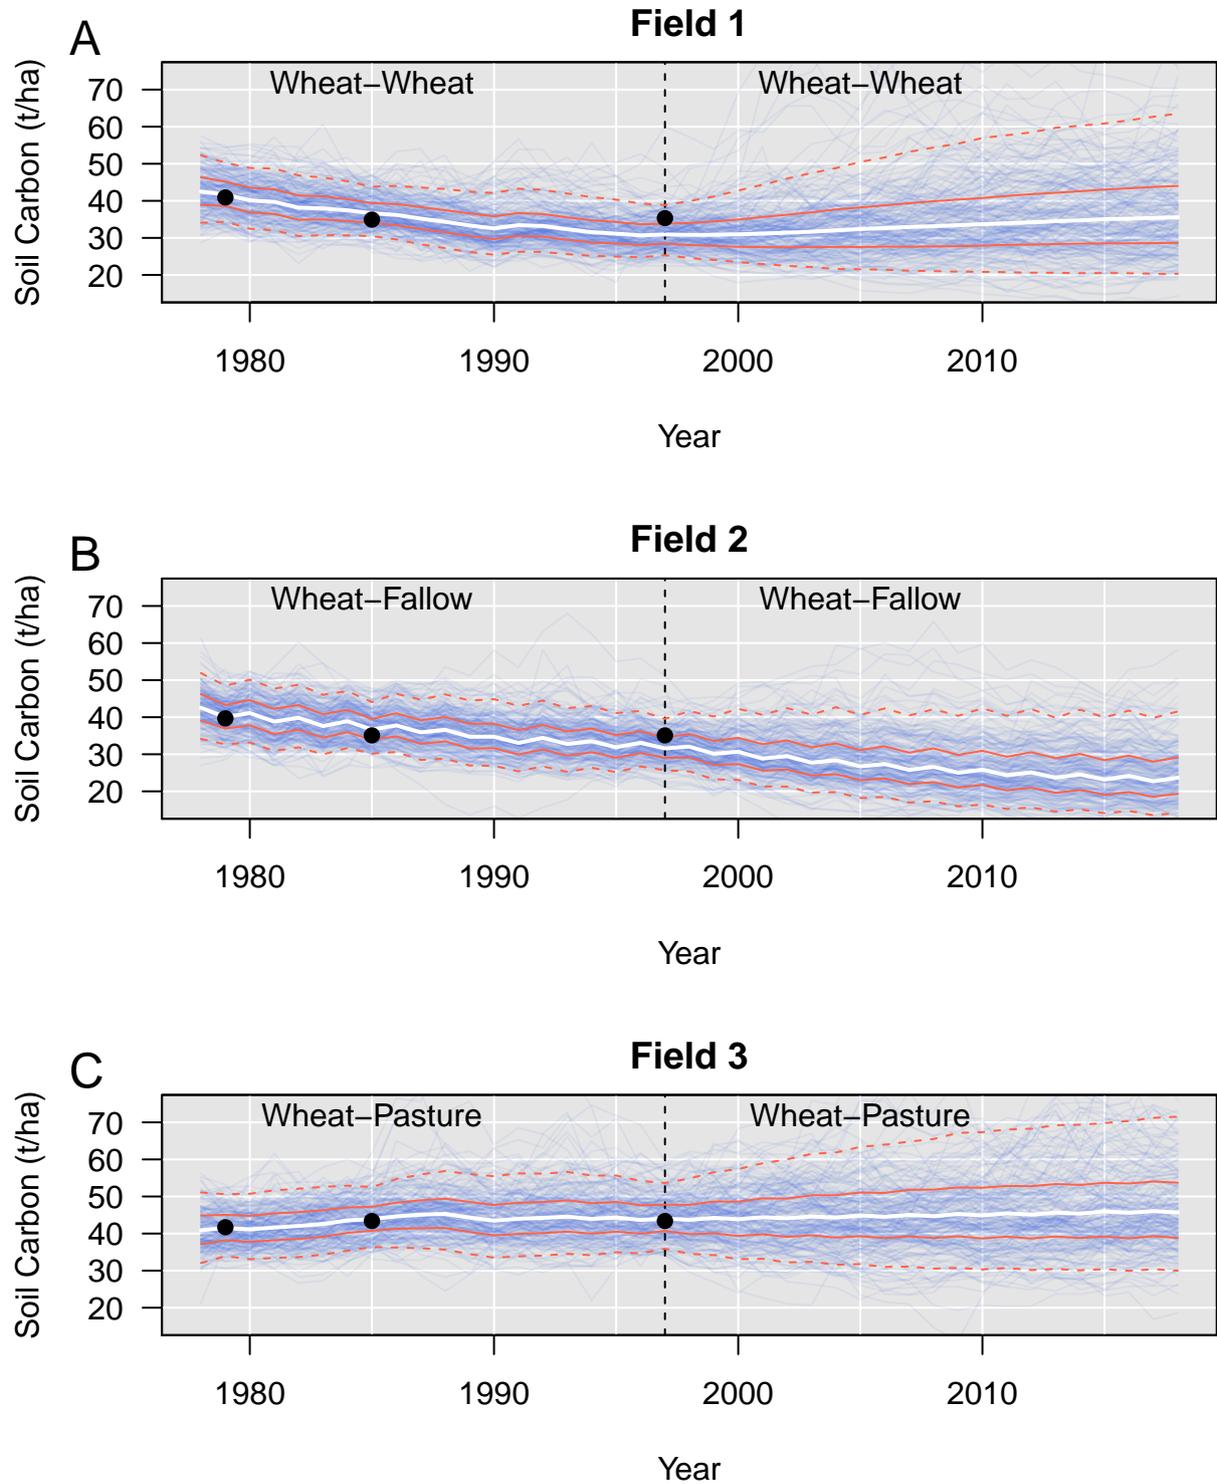


Figure 1. SOC ($t\ ha^{-1}$) dynamics in the three Tarlee fields. The measured SOC values are indicated by black dots. The 2.5th percentiles and the 97.5th percentiles for the SOC process are indicated by the dashed red lines. The 25th percentiles and the 75th percentiles for the SOC process are indicated by the solid red lines. The median carbon mass is indicated by the white line. From 1998 onwards these plots indicate forecasted SOC dynamics in Fields 1, 2, and 3 under the same management strategies.

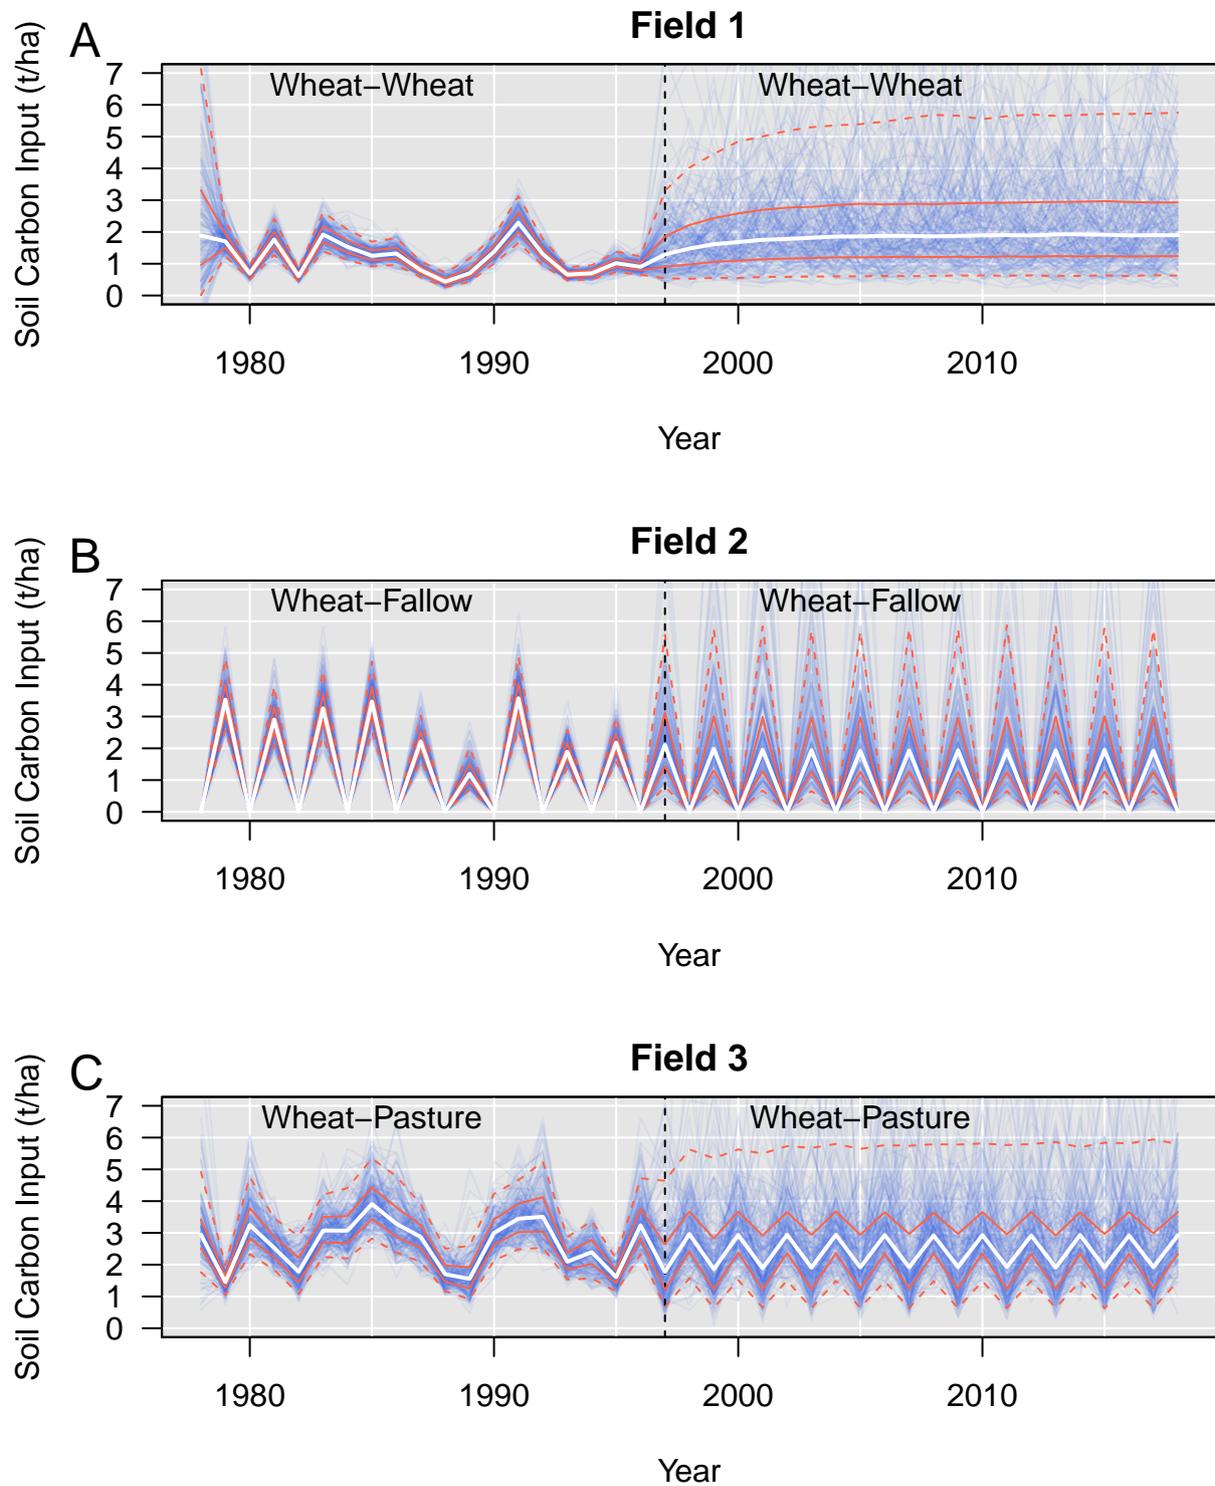


Figure 2. SOC inputs ($t\ ha^{-1}$) in the three Tarlee fields based on the crop sown each year, see Equation 5. The 2.5th percentiles and the 97.5th percentiles for the SOC process are indicated by the dashed red lines. The 25th percentiles and the 75th percentiles for the SOC process are indicated by the solid red lines. The median carbon mass is indicated by the white line. From 1998 onwards these plots indicate forecasted SOC inputs in Fields 1, 2, and 3 under the same management strategies.

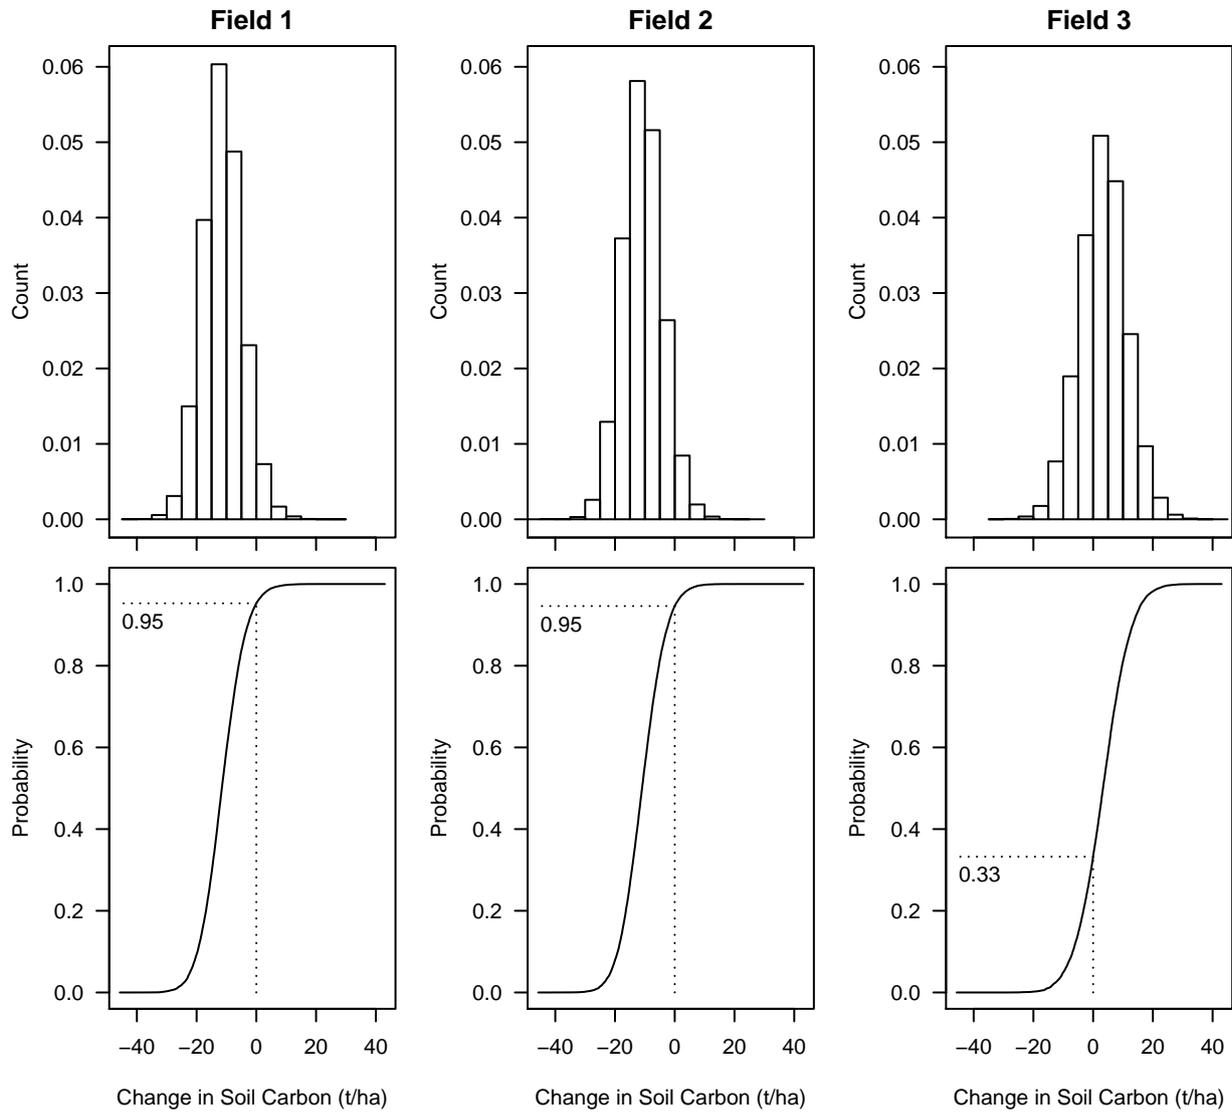


Figure 3. Histograms (top row) and Cumulative Distribution Functions (bottom row) of the change in soil carbon between 1978 and 1997 for each field. The posterior probability of sequestering no carbon under each management scheme (i.e., \hat{f}_3) is indicated in the plots on the bottom row. These probabilities for each field are noted as, respectively, 0.95, 0.95, and 0.33.

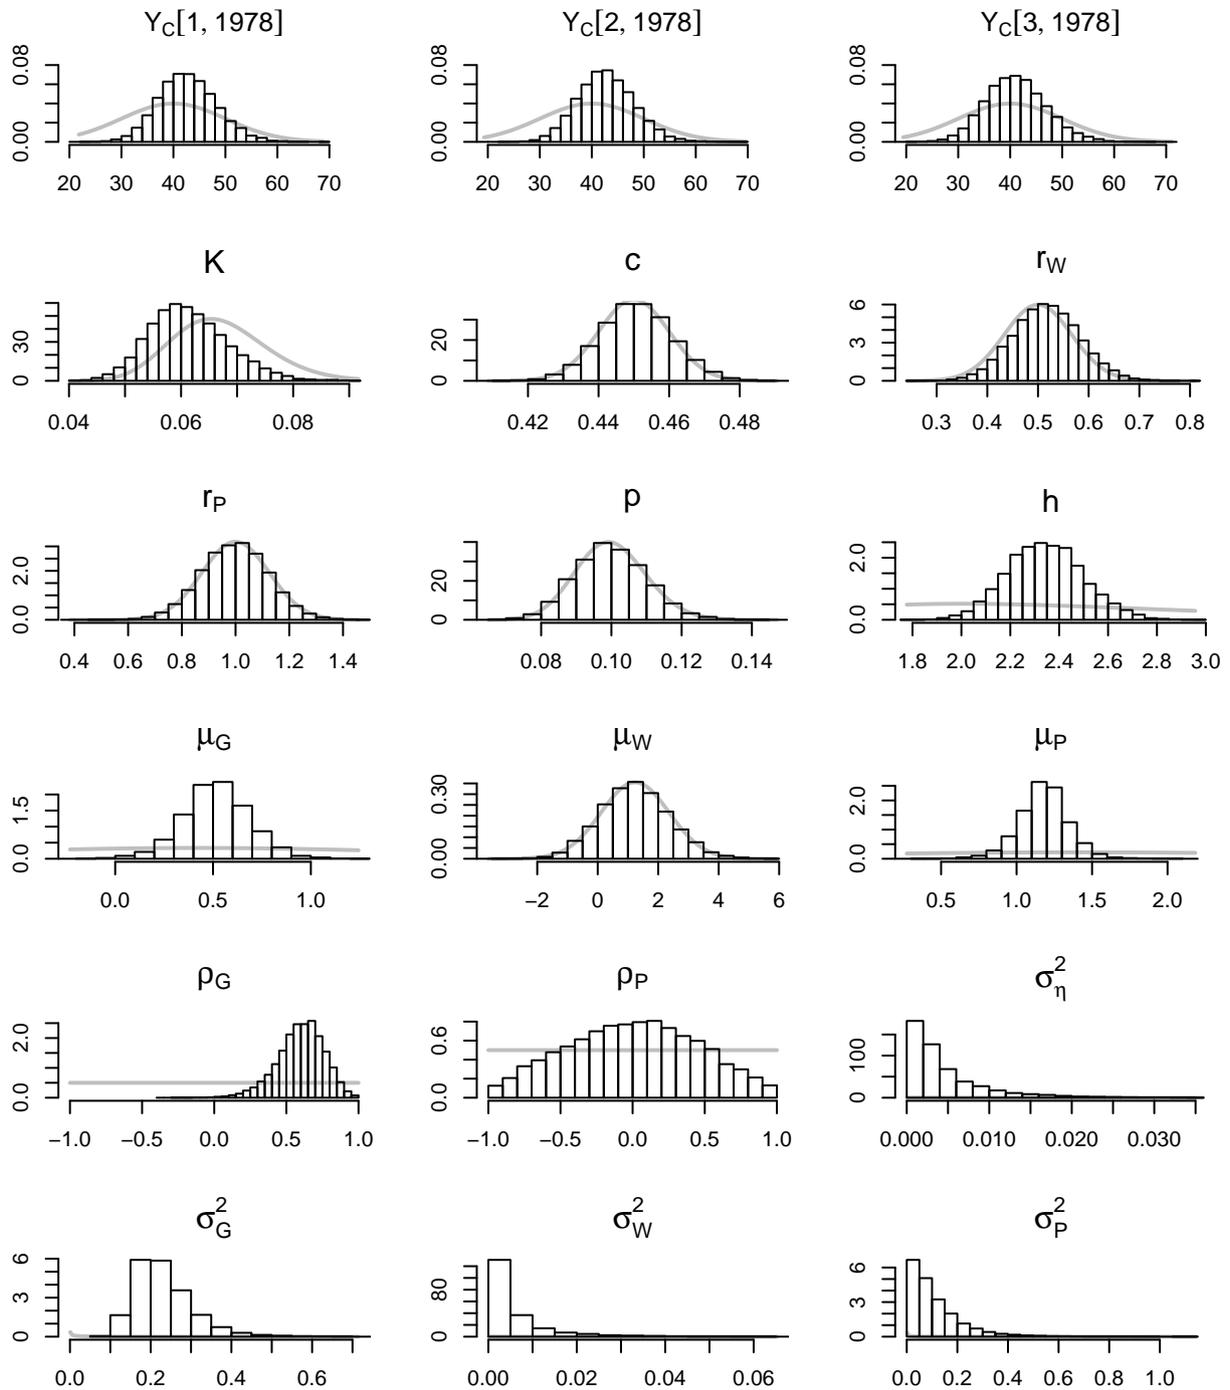


Figure 4. Panels show the prior distribution (grey density) and marginal posterior distribution (histogram) for eighteen model parameters. Uninformative priors for some of the parameters are barely visible when plotted over the histograms due to the fact that much has been learned about those parameters from the data.

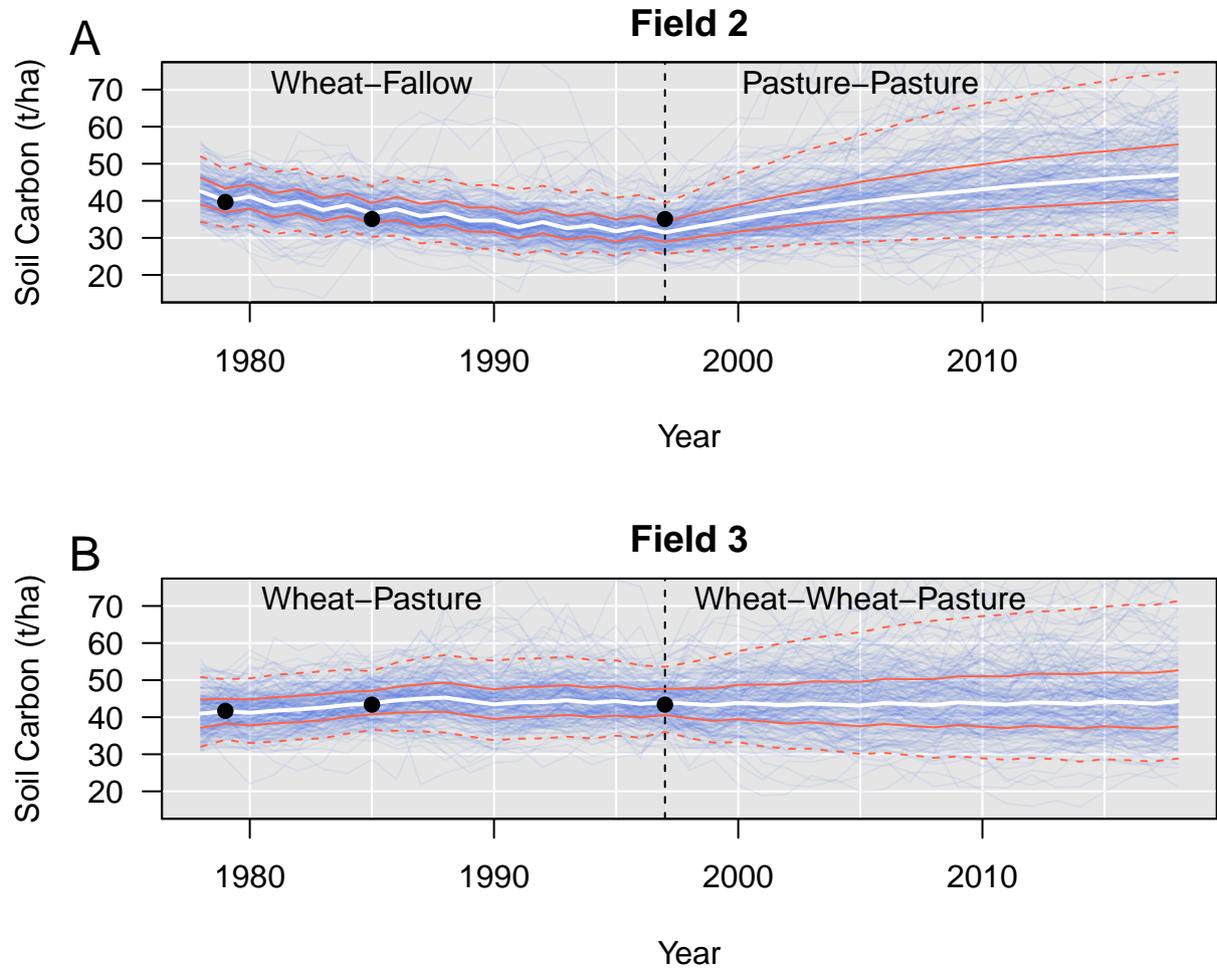


Figure 5. SOC dynamics ($t\ ha^{-1}$) in Fields 2 and 3 based on the crop sown each year with a change in management practice from 1998 onwards. The 2.5th percentiles and the 97.5th percentiles for the SOC process are indicated by the dashed red lines. The 25th percentiles and the 75th percentiles for the SOC process are indicated by the solid red lines. The median carbon mass is indicated by the white line. From 1998 onwards these plots indicate forecasted SOC dynamics under continuous pasture for Field 2 and under a wheat-wheat-pasture management practice in Field 3.

TABLES

Table 1. Treatment of parameters in the model.

Parameter	Prior	Type	Equation
$Y_C(1, 1978)$	Truncated-Normal($40, 10^2, \text{lower}=0$)	Uninformative	1
$Y_C(2, 1978)$	Truncated-Normal($40, 10^2, \text{lower}=0$)	Uninformative	1
$Y_C(3, 1978)$	Truncated-Normal($40, 10^2, \text{lower}=0$)	Uninformative	1
K	Log-Normal($-2.71, (0.127)^2$)	Informative	1
c	Normal($0.45, (0.01)^2$)	Informative	5
r_W	Normal($0.5, (0.067)^2$)	Informative	5
r_P	Normal($1.0, (0.125)^2$)	Informative	5
p	Beta($89.9, 809.1$)	Informative	5
h	Log-Normal($0.825, 0.36^2$)	Weakly Informative	3
μ_G	Normal($0.42, (1.18)^2$)	Weakly Informative	2
μ_W	Normal($1.24, (1.12)^2$)	Weakly Informative	3
μ_P	Normal($1.41, (1.81)^2$)	Weakly Informative	4
ρ_G	Uniform($-1, 1$)	Uninformative	2
ρ_P	Uniform($-1, 1$)	Uninformative	4
σ_η^2	Inv-Gamma($0.001, 0.001$)	Uninformative	1
σ_G^2	Inv-Gamma($0.001, 0.001$)	Uninformative	2
σ_W^2	Inv-Gamma($0.001, 0.001$)	Uninformative	3
σ_P^2	Inv-Gamma($0.001, 0.001$)	Uninformative	4
$\sigma_{\epsilon_C}^2$	0.025	Fixed	7
$\sigma_{\epsilon_G}^2$	0.023	Fixed	8
$\sigma_{\epsilon_W}^2$	0.133	Fixed	9
$\sigma_{\epsilon_P}^2$	0.067	Fixed	10
$Y_G(1, 1978)$	Log-Normal($\mu_G, 4\sigma_G^2$)	Uninformative	2
$Y_G(2, 1978)$	Log-Normal($\mu_G, 4\sigma_G^2$)	Uninformative	2
$Y_G(3, 1978)$	Log-Normal($\mu_G, 4\sigma_G^2$)	Uninformative	2
$Y_P(1, 1978)$	Log-Normal($\mu_P, 4\sigma_P^2$)	Uninformative	4
$Y_P(2, 1978)$	Log-Normal($\mu_P, 4\sigma_P^2$)	Uninformative	4
$Y_P(3, 1978)$	Log-Normal($\mu_P, 4\sigma_P^2$)	Uninformative	4

Table 2. Proposal distributions used in PMMH.

Parameter	Proposal
K	Normal($K, 0.001^2$)
c	Truncated-Normal($c, 0.005^2$, lower = 0, upper = 1)
r_W	Truncated-Normal($r_W, 0.05^2$, lower = 0)
r_P	Truncated-Normal($r_P, 0.05^2$, lower = 0)
p	Truncated-Normal($p, 0.005^2$, lower = 0, upper = 1)
h	Truncated-Normal($h, 0.05^2$, lower = 0)
μ_G	Normal($\mu_G, 0.05^2$)
μ_W	Normal($\mu_W, 0.2^2$)
μ_P	Normal($\mu_P, 0.05^2$)
ρ_G	Truncated-Normal($\rho_G, 0.05$, lower = -1, upper = 1)
ρ_P	Truncated-Normal($\rho_P, 0.1$, lower = -1, upper = 1)
σ_η^2	Truncated-Normal($\sigma_\eta^2, 0.001^2$, lower = 0)
σ_G^2	Truncated-Normal($\sigma_G^2, \sigma_G^2/20^2$, lower = 0)
σ_W^2	Truncated-Normal($\sigma_W^2, 0.001^2$, lower = 0)
σ_P^2	Truncated-Normal($\sigma_P^2, 0.1^2$, lower = 0)