Reified Bayesian Modelling

Michael Goldstein and Jonathan Rougier
Dept. Mathematical Sciences, University of Durham, UK

How to relate models and physical systems?

- Basic ingredients:

How to relate models and physical systems?

- Basic ingredients:
  x^*:system properties

How to relate models and physical systems?

- Basic ingredients:

  x^*:system properties

  y: system behaviour (influenced by x^*)

How to relate models and physical systems?

- Basic ingredients:
  x^* :system properties
  y: system behaviour (influenced by x^*)
  z: partial observation of y (with error)

How to relate models and physical systems?

- Basic ingredients:
  x^*:system properties
  y: system behaviour (influenced by x^*)
  z: partial observation of y (with error)
- Ideally, we would like to construct a deterministic model f, embodying the **laws of nature**, which satisfies

$$y = f(x^*).$$

How to relate models and physical systems?

- Basic ingredients:
  x^*:system properties
  y: system behaviour (influenced by x^*)
  z: partial observation of y (with error)
- Ideally, we would like to construct a deterministic model f, embodying the **laws of nature**, which satisfies

$$y = f(x^*).$$

- In practice, however, the our actual model f is inadequate:
  - f simplifies the physics;
  - f approximates the solution of the physical equations

How to relate models and physical systems?

- Basic ingredients:
  x^*:system properties
  y: system behaviour (influenced by x^*)
  z: partial observation of y (with error)
- Ideally, we would like to construct a deterministic model f, embodying the **laws of nature**, which satisfies

$$y = f(x^*).$$

- In practice, however, the our actual model f is inadequate:
  - f simplifies the physics;
  - f approximates the solution of the physical equations
- The fundamental question: What does the imperfect f tell us about the system values (x^*, y)?

How to relate models and physical systems?

- Basic ingredients:
  x^*:system properties
  y: system behaviour (influenced by x^*)
  z: partial observation of y (with error)
- Ideally, we would like to construct a deterministic model f, embodying the **laws of nature**, which satisfies

$$y = f(x^*).$$

- In practice, however, the our actual model f is inadequate:
  - f simplifies the physics;
  - f approximates the solution of the physical equations
- The fundamental question: What does the imperfect f tell us about the system values (x^*, y)?
  [and what do several imperfect simulators for the system jointly tell us about (x^*, y)?]

Model inadequacy

- The best approach so far requires us to introduce a *model discrepancy*, i.e.

$$y = f(x^*) + \epsilon \qquad \epsilon \perp\!\!\!\perp (f, x^*)$$

see, e.g., Kennedy and O'Hagan (JRSSB 2001)
Craig, Goldstein, Rougier and Seheult (JASA 2001).

Model inadequacy

- The best approach so far requires us to introduce a *model discrepancy*, i.e.

$$y = f(x^*) + \epsilon \qquad \epsilon \perp\!\!\!\perp (f, x^*)$$

  see, e.g., Kennedy and O'Hagan (JRSSB 2001)
  Craig, Goldstein, Rougier and Seheult (JASA 2001).
- Clearly, adding model disrepancy term is a great improvement over ignoring model discrepancy.
  BUT observe that this implies that all of the information about y in the model f is contained in the result of the single evaluation $f(x^*)$.

Model inadequacy

- The best approach so far requires us to introduce a *model discrepancy*, i.e.

$$y = f(x^*) + \epsilon \qquad \epsilon \perp\!\!\!\perp (f, x^*)$$

  see, e.g., Kennedy and O'Hagan (JRSSB 2001)
  Craig, Goldstein, Rougier and Seheult (JASA 2001).
- Clearly, adding model disrepancy term is a great improvement over ignoring model discrepancy.
  BUT observe that this implies that all of the information about y in the model f is contained in the result of the single evaluation $f(x^*)$.
- Is this reasonable? Sometimes yes, but very often no.
  Intuitively, this assumption seems to be arbitrary and overly restrictive.
  A simple thought experiment can demonstrate the nature of the problem with this formulation.

That simple thought experiment

- *Imagine:* An improved version of f, say f' which has a better solver resolution. We would also want

$$y = f'(x^*) + \epsilon' \qquad \epsilon' \perp\!\!\!\perp (f, f', x^*)$$

from which we can deduce that

$$f'(x^*) - f(x^*) \perp\!\!\!\perp (f, x^*).$$

That simple thought experiment

- *Imagine:* An improved version of f, say f' which has a better solver resolution. We would also want

$$y = f'(x^*) + \epsilon' \qquad \epsilon' \perp\!\!\!\perp (f, f', x^*)$$

from which we can deduce that

$$f'(x^*) - f(x^*) \perp\!\!\!\perp (f, x^*).$$

- *But:* Knowledge of f is surely relevant to knowledge of f', so (f, x^*) is usually informative for $f'(x^*) - f(x^*)$.

That simple thought experiment

- *Imagine:* An improved version of f, say f' which has a better solver resolution. We would also want

$$y = f'(x^*) + \epsilon' \qquad \epsilon' \perp\!\!\!\perp (f, f', x^*)$$

from which we can deduce that

$$f'(x^*) - f(x^*) \perp\!\!\!\perp (f, x^*).$$

- *But:* Knowledge of f is surely relevant to knowledge of f', so (f, x^*) is usually informative for $f'(x^*) - f(x^*)$.
- For example, if x is a scalar, $f(x) \approx bx$ and we consider that $f'(x) \approx b'x$ where b is informative for b', then

$$f'(x^*) - f(x^*) \approx (b' - b)x^*$$

Even if $(b' - b) \perp\!\!\!\perp b$, we can't judge $f'(x^*) - f(x^*)$ independent of x^*.

Reified modelling

- The simple statistical approach that asserts $y = f(x^*) + \epsilon$ cannot usually have $\epsilon \perp\!\!\!\perp (f, x^*)$. Instead, model discrepancy ϵ will often have a very complicated structural relation with f and x^*.

Reified modelling

- The simple statistical approach that asserts $y = f(x^*) + \epsilon$ cannot usually have $\epsilon \perp\!\!\!\perp (f, x^*)$. Instead, model discrepancy ϵ will often have a very complicated structural relation with f and x^*.
- Imagine constructing a *reified model* f^* which has both (i) sufficiently good physics, and (ii) sufficiently accurate solutions that we have no additional insights as to the nature of any further improvements to f^*.
  Therefore, model discrepancy on top of our reified model will be unstructured.
  (DEFN.) [**REIFY**: to consider an abstract concept to be real.]

Reifield modelling

- The simple statistical approach that asserts $y = f(x^*) + \epsilon$ cannot usually have $\epsilon \perp\!\!\!\perp (f, x^*)$. Instead, model discrepancy ϵ will often have a very complicated structural relation with f and x^*.
- Imagine constructing a *reified model* f^* which has both (i) sufficiently good physics, and (ii) sufficiently accurate solutions that we have no additional insights as to the nature of any further improvements to f^*.
  Therefore, model discrepancy on top of our reified model will be unstructured.
  (DEFN.) [**REIFY**: to consider an abstract concept to be real.]
- We apply the discrepancy model to f^* and to f^* alone, i.e.

$$y = f^*(x^*, w^*) + \epsilon^* \qquad \epsilon^* \perp\!\!\!\perp (f, f^*, x^*, w^*)$$

where w are any extra model parameters that might be introduced.

The Reifying Principle.

Our model f is informative for y

because

f is informative for the reified model f^*

The Reifying Principle.

Our model f is informative for y

because

f is informative for the reified model f^*

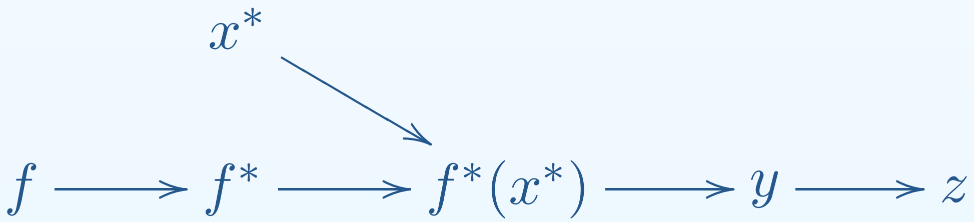

The Reifying Principle.

Our model f is informative for y

because

f is informative for the reified model f^*

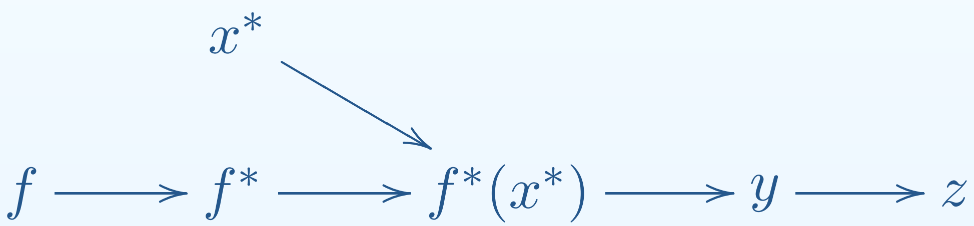

Comment The reifying principle is a sensible pragmatic compromise which retains the essential tractability in linking computer evaluations and system values with system behaviour, while removing logical problems in simple treatments of discrepancy, and providing guidance for discrepancy modelling.

The Reifying Principle.

Our model f is informative for y

because

f is informative for the reified model f^*

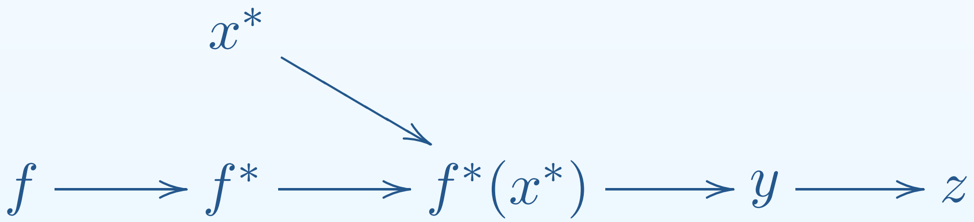

Comment The reifying principle is a sensible pragmatic compromise which retains the essential tractability in linking computer evaluations and system values with system behaviour, while removing logical problems in simple treatments of discrepancy, and providing guidance for discrepancy modelling. If we have several simulators $f_1, f_2, \ldots, f_r$, then the reifying principle suggests that we combine their information by treating each simulator as informative for the single reified form f^*.

Linking f and f^* using emulators

- An *emulator* is a probabilistic belief specification for a deterministic function. Our emulator for component i of f might be

$$f_i(x) = \sum_j \beta_{ij}\, g_{ij}(x) + u_i(x)$$

where $B = \{\beta_{ij}\}$ are unknown scalars, g_{ij} are known deterministic functions of x, and $u(x)$ is a weakly stationary stochastic process (maybe gaussian).

Linking f and f^* using emulators

- An *emulator* is a probabilistic belief specification for a deterministic function. Our emulator for component i of f might be

$$f_i(x) = \sum_j \beta_{ij}\, g_{ij}(x) + u_i(x)$$

  where $B = \{\beta_{ij}\}$ are unknown scalars, g_{ij} are known deterministic functions of x, and $u(x)$ is a weakly stationary stochastic process (maybe gaussian).
- Our simplest emulator for f^* would then be

$$f_i^*(x, w) = \sum_j \beta_{ij}^*\, g_{ij}(x) + u_i^*(x) + u_i^*(x, w)$$

  where we might model $B^* = \{\beta_{ij}^*\}$ as $\beta_{ij}^* = c_{ij}\, \beta_{ij} + \nu_{ij}$ for known c_{ij} and uncertain ν_{ij}, and correlate $u(x)$ and $u^*(x)$, but leave $u^*(x, w)$ uncorrelated.

Structural reification

Usually, we can think more carefully about the reasons for our model's inadequacy.

Structural reification

Usually, we can think more carefully about the reasons for our model's inadequacy.

- Often we can imagine specific generalisations for f with extra model components and parameters giving an better model $f'(x, v)$.

Structural reification

Usually, we can think more carefully about the reasons for our model's inadequacy.

- Often we can imagine specific generalisations for f with extra model components and parameters giving an better model $f'(x, v)$.
- Suppose $f'(x, v_0) = f(x)$. We might emulate f' 'on top' of f, i.e. as

$$f_i'(x, v) = f_i(x) + \sum_k \gamma_{ik} \, g_{ik}(x, v) + u_i(x, v)$$

where $g_{ik}(x, v_0) = u_i(x, v_0) = 0$.

Structural reification

Usually, we can think more carefully about the reasons for our model's inadequacy.

- Often we can imagine specific generalisations for f with extra model components and parameters giving an better model $f'(x, v)$.
- Suppose $f'(x, v_0) = f(x)$. We might emulate f' 'on top' of f, i.e. as

$$f_i'(x, v) = f_i(x) + \sum_k \gamma_{ik}\, g_{ik}(x, v) + u_i(x, v)$$

  where $g_{ik}(x, v_0) = u_i(x, v_0) = 0$.
- The reified emulator for $f_i^*(x, v, w)$ would then be

$$\sum_j \beta_{ij}^*\, g_{ij}(x) + \sum_k \gamma_{ik}^*\, g_{ik}(x, v) + u_i^*(x) + u_i^*(x, v) + u_i^*(x, v, w)$$

  where we now model the relationship between the coefficients in the three emulators.

The Zickfeld et al (2004) model of the Atlantic

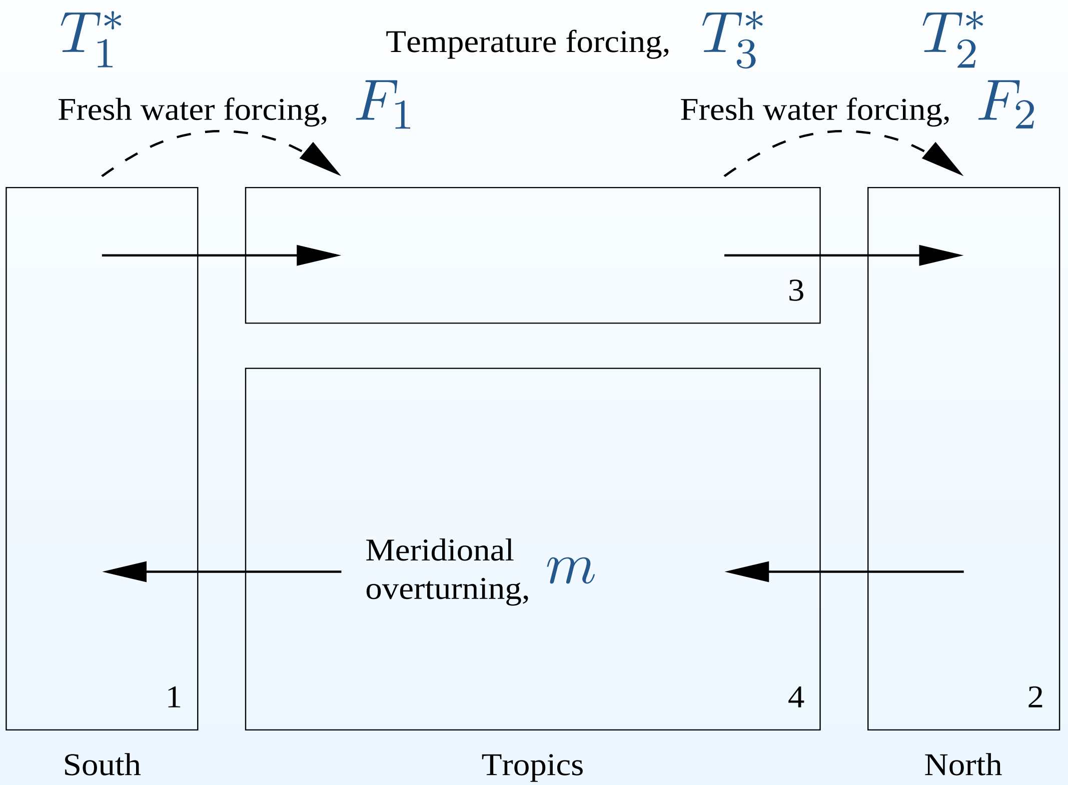

- Model parameters: $T_1^*, T_2^*, T_3^*, \Gamma, k$ (last two not shown above).
- Model outputs: Steady state (SS) temperature for compartments 1, 2 and 3, SS salinity differences between compartments 1 & 2 and between 2 & 3. SS. overturning m, and critical freshwater forcing F_1^{crit}.

One possible generalisation

An extra compartment at the 'southern' end denoting 'Other oceans'.

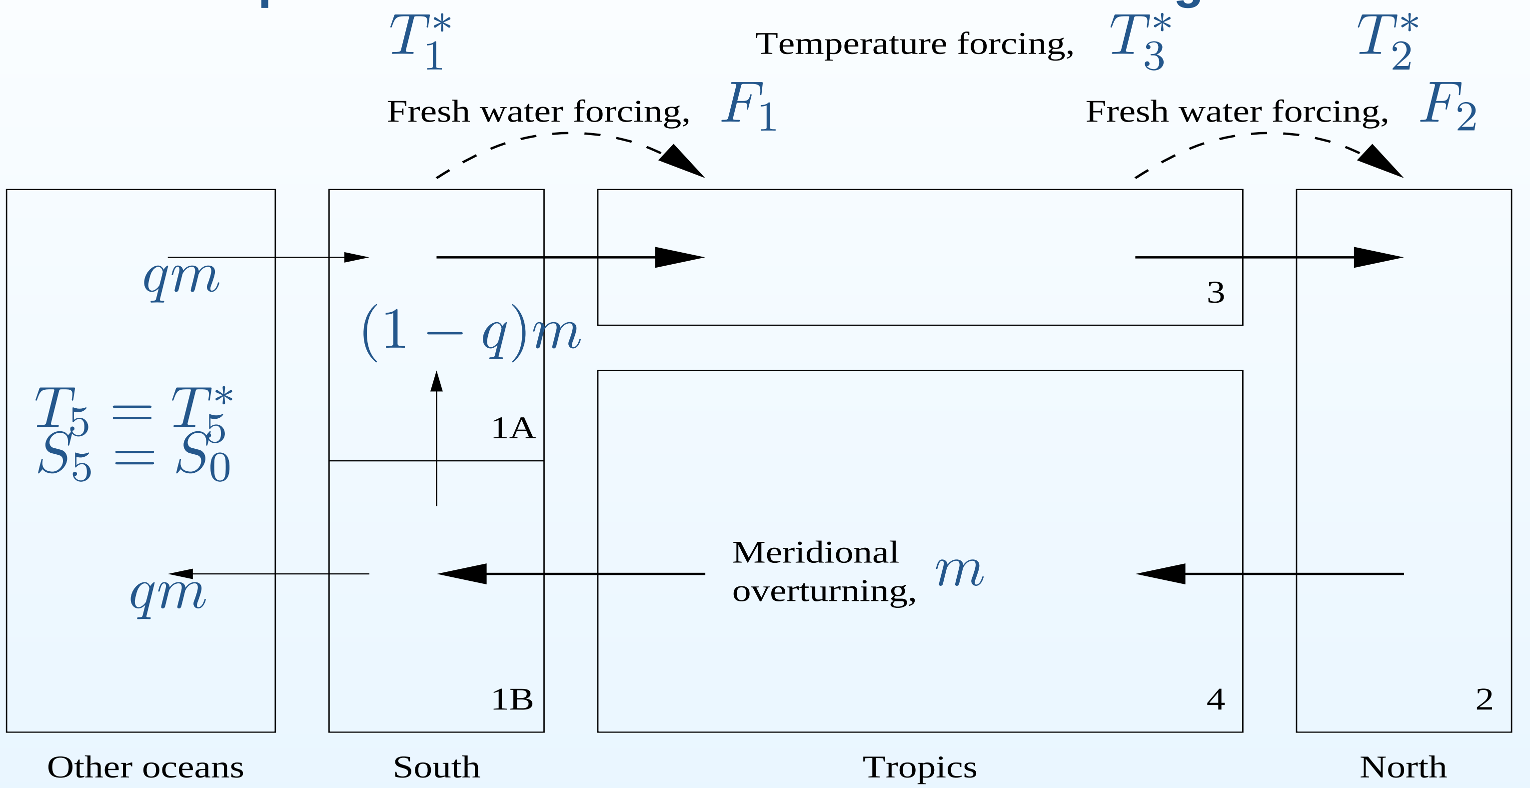

- Two extra model parameters, T_5^* and q, with $q = 0$ returning us to the original model.
- Same model outputs as before.

The system and the data

The system and the data

- Our system y is the Atlantic, but, due to the very aggregate nature of our model it is more natural to use data from a much larger model (CLIMBER-2) that has been carefully tuned to the Atlantic in a separate experiment. The data from CLIMBER-2 comprises SS. temperatures, SS. salinity differences and SS. overturning. We write

$$y = f^*(x^*, v^*) + \epsilon^* \quad \text{and} \quad z = Hy.$$

The system and the data

- Our system y is the Atlantic, but, due to the very aggregate nature of our model it is more natural to use data from a much larger model (CLIMBER-2) that has been carefully tuned to the Atlantic in a separate experiment. The data from CLIMBER-2 comprises SS. temperatures, SS. salinity differences and SS. overturning. We write

$$y = f^*(x^*, v^*) + \epsilon^* \quad \text{and} \quad z = Hy.$$

- For the discrepancy,

$$\mathrm{Var}\epsilon_i^* = \begin{cases} 0.84^2 & i = 1, 2, 3[\text{SS Temps}] \\ 0.075^2 & i = 4, 5[\text{SS Salinity Differences}] \\ 3.30^2 & i = 6[\text{SS Overturning}] \\ 0.044^2 & i = 7[F_1^{\text{crit}}]. \end{cases}$$

Some aspects of our emulation

Some aspects of our emulation

- Our emulator for $F_1^{\text{crit}}(x)$ is ($\times 10^{-3}$)

$$87 - T_2^* + 68(T_1^* - T_2^*) - 5(T_3^* - T_1^*) + 14\,\Gamma + 21\,k + u_7(x)$$

with $\mathrm{Sd}(u_7(x)) = 21$ [based on 30 evaluations of the simulator]

Some aspects of our emulation

- Our emulator for $F_1^{\text{crit}}(x)$ is ($\times 10^{-3}$)

$$87 - T_2^* + 68(T_1^* - T_2^*) - 5(T_3^* - T_1^*) + 14\,\Gamma + 21\,k + u_7(x)$$

  with $\text{Sd}(u_7(x)) = 21$ [based on 30 evaluations of the simulator]

- Assess the difference between f' and f, and between f^* and f' as

$$\begin{aligned} \Delta' &= \text{E}(\text{Var}(f'(x^*, v^*) - f(x^*))|x^*, v^*) \\ \Delta^* &= \text{E}(\text{Var}(f^*(x^*, v^*, w^*) - f'(x^*, v^*))|x^*, v^*, w^*) \end{aligned}$$

Some aspects of our emulation

- Our emulator for $F_1^{\text{crit}}(x)$ is ($\times 10^{-3}$)

$$87 - T_2^* + 68(T_1^* - T_2^*) - 5(T_3^* - T_1^*) + 14\,\Gamma + 21\,k + u_7(x)$$

with $\text{Sd}(u_7(x)) = 21$ [based on 30 evaluations of the simulator]

- Assess the difference between f' and f, and between f^* and f' as

$$\begin{aligned}\Delta' &= \text{E}(\text{Var}(f'(x^*, v^*) - f(x^*))|x^*, v^*)\\ \Delta^* &= \text{E}(\text{Var}(f^*(x^*, v^*, w^*) - f'(x^*, v^*))|x^*, v^*, w^*)\end{aligned}$$

For F_1^{crit} our modelling leads to the following assignments:
$\text{SD}(f(x^*)) = 0.079$
$\sqrt{\Delta'} = 0.033$
$\sqrt{\Delta^*} = 0.066$
$\text{SD}(\epsilon^*) = 0.044$

Our resulting predictions for F_1^{crit}

Our resulting predictions for F_1^{crit}

1. From 30 evaluations of f, plus the reified statistical modelling, our prior prediction is

$$F_1^{\text{crit}} = 0.085 \pm 0.224 \quad \text{(mean } \pm 2 \text{ std. dev.)}$$

Our resulting predictions for F_1^{crit}

1. From 30 evaluations of f, plus the reified statistical modelling, our prior prediction is

$$F_1^{\text{crit}} = 0.085 \pm 0.224 \quad \text{(mean } \pm 2 \text{ std. dev.)}$$

2. Using the CLIMBER-2 data z in the Bayes linear approach [see, Craig, Goldstein, Rougier and Seheult (JASA 2001)], our adjusted prediction is

$$F_1^{\text{crit}} = 0.119 \pm 0.099$$

[Our conclusions are ‘reasonably insensitive’ to moderate tweaks to our prior values, and our modelling survives ‘reasonable diagnostic testing’.]

Our resulting predictions for F_1^{crit}

1. From 30 evaluations of f, plus the reified statistical modelling, our prior prediction is

$$F_1^{\text{crit}} = 0.085 \pm 0.224 \quad \text{(mean } \pm 2 \text{ std. dev.)}$$

2. Using the CLIMBER-2 data z in the Bayes linear approach [see, Craig, Goldstein, Rougier and Seheult (JASA 2001)], our adjusted prediction is

$$F_1^{\text{crit}} = 0.119 \pm 0.099$$

   [Our conclusions are 'reasonably insensitive' to moderate tweaks to our prior values, and our modelling survives 'reasonable diagnostic testing'.]

3. **Model Design** The effect of making many more evaluations of current simulator f would be to reduce uncertainty about F_1^{crit} by around 2%.
   In comparison, the effect of constructing and evaluating the 5 compartment model f' would be to reduce uncertainty in F_1^{crit} by around 10%.

Conclusions

Reified modelling has the following advantages:

Conclusions

Reified modelling has the following advantages:

1. Provides a logical rationale for linking models to systems

Conclusions

Reified modelling has the following advantages:

1. Provides a logical rationale for linking models to systems
2. Provides a structural modelling approach for quantifying discrepancies, representing whatever physical intuition we consider relevant

Conclusions

Reified modelling has the following advantages:

1. Provides a logical rationale for linking models to systems
2. Provides a structural modelling approach for quantifying discrepancies, representing whatever physical intuition we consider relevant
3. Offers guidance as to which model elaborations are likely to be sufficiently informative for the quantities of interest that they would justify the effort required for careful numerical implementation.

Conclusions

Reified modelling has the following advantages:

1. Provides a logical rationale for linking models to systems
2. Provides a structural modelling approach for quantifying discrepancies, representing whatever physical intuition we consider relevant
3. Offers guidance as to which model elaborations are likely to be sufficiently informative for the quantities of interest that they would justify the effort required for careful numerical implementation.
4. Provides a way to unify analyses from several models of the same system (through joint reification)

Conclusions

Reified modelling has the following advantages:

1. Provides a logical rationale for linking models to systems
2. Provides a structural modelling approach for quantifying discrepancies, representing whatever physical intuition we consider relevant
3. Offers guidance as to which model elaborations are likely to be sufficiently informative for the quantities of interest that they would justify the effort required for careful numerical implementation.
4. Provides a way to unify analyses from several models of the same system (through joint reification)
5. Is the only approach that we know of that satisfies (1)–(4).

Conclusions

Reified modelling has the following advantages:

1. Provides a logical rationale for linking models to systems
2. Provides a structural modelling approach for quantifying discrepancies, representing whatever physical intuition we consider relevant
3. Offers guidance as to which model elaborations are likely to be sufficiently informative for the quantities of interest that they would justify the effort required for careful numerical implementation.
4. Provides a way to unify analyses from several models of the same system (through joint reification)
5. Is the only approach that we know of that satisfies (1)–(4).

References

M. Goldstein and J.C.Rougier. Probabilistic formulations for transferring inferences from mathematical models to physical systems (2005) SIAM journal on scientific computing, 26, 467-487.

M. Goldstein and J.C.Rougier. Reified Bayesian modelling and inference for physical systems (2006) In revision for JSPI.