UQ and Inverting Imperfect Models

Tony O'Hagan

Outline

- Uncertainty Quantification
  - Applied maths versus statistics UQ
- The objectives of inversion
  - Physical and tuning parameters
  - Imperfect simulators – two examples
- Ignoring model discrepancy
  - Results of inverting the examples
  - Implications for learning about physical parameters
- The simple machine and model discrepancy
  - Simple model discrepancy
  - Nonidentifiability
- Conclusions

Uncertainty Quantification

UQ

- Uncertainty Quantification (UQ) is a relatively new area of study in applied mathematics and engineering
  - Becoming a major focus of research (and funding)
  - Applications wherever complex simulation models are used
- It is concerned with uncertainties in the predictions of models
  - Mechanistic, science-based models (simulators)
    - Unlike empirical, statistical models, these do not intrinsically account for any of the uncertainties in their predictions
  - Often based on differential equations
- The idea of quantifying (some of) the uncertainties is quite a new concept in these fields
  - But uncertainty quantification is what statisticians have always done!

The simulator as a function

- Using computer language, a simulator takes a number of inputs and produces a number of outputs

- We can represent any output y as a function

$$y = \eta(x)$$

  of a vector x of inputs

Where is the uncertainty?

- How might the simulator output $y = \eta(x)$ differ from the true real-world value z that the simulator is supposed to predict?
  - Error in inputs x
    - Initial values
    - Forcing inputs
    - Model parameters
  - Error in model structure or solution
    - Wrong, inaccurate or incomplete science
    - Bugs, solution errors
- The problem is made particularly challenging by the fact that simulators are often computationally expensive
  - From minutes to months for a single evaluation of $\eta(.)$

Quantifying uncertainty

- The ideal is to provide a probability distribution $p(z)$ for the true real-world value
  - The centre of the distribution is a best estimate
  - Its spread shows how much uncertainty about z is induced by uncertainties on the previous slide

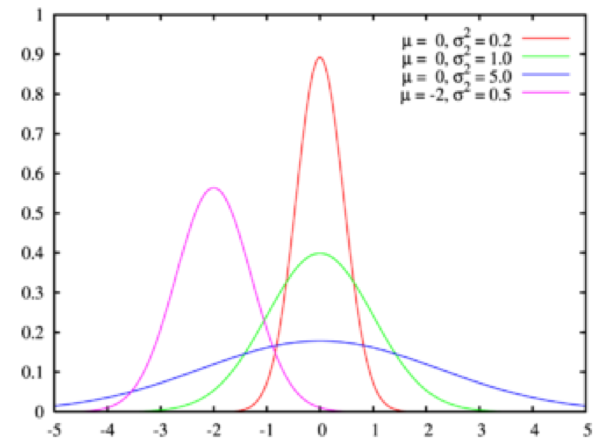

- How do we get this?
  - Input uncertainty: characterise $p(x)$, propagate through to $p(y)$
  - Structural uncertainty: characterise $p(z - y)$

Managing uncertainty

- To understand the implications of different uncertainty sources
  - Probabilistic, variance-based sensitivity analysis
  - Helps with targeting and prioritising research
- To reduce uncertainty, get more information!
- Informal – more/better science
  - Tighten $p(x)$ through improved understanding
  - Tighten $p(z - y)$ through improved modelling or programming
- Formal – using real-world data
  - Calibration – learn about model parameters
  - Data assimilation – learn about the state variables
  - Learn about structural error $z - y$
  - Validation

Applied maths versus Statistics UQ

- Applied Maths
  - Tools for propagating input uncertainty
    - Polynomial chaos expansions, stochastic collocation etc.
  - No study of how to quantify input uncertainty
  - No accounting for structural uncertainty (model inadequacy)
  - No accounting for uncertainty due to approximation of the simulator
    - E.g. through truncating expansions
- Statistics
  - Total UQ
    - All sources of uncertainty are studied and can be addressed
    - And we have our own smart tools (emulators) for propagation
  - MUCM project (http://mucm.ac.uk)

The Objectives of Inversion

Inversion as nonlinear regression

- We have a simulator $\eta(x,\theta)$ and observations

$$z_i = \eta(x_i,\theta) + \varepsilon_i$$

- In statistical language this is a nonlinear regression model
  - The inversion problem is one of inference about θ
- I'll be assuming the Bayesian paradigm
  - Requires a prior distribution for θ
    - Often assumed to be non-informative
  - Produces a posterior distribution
- Very common approach, but has a major flaw
  - The observations are of the real physical system $\zeta(.)$
  - And the simulator is invariably imperfect: $\eta(.,\theta) \neq \zeta(.)\ \forall\theta$

Model discrepancy

- We should write

$$z_i = \zeta(x_i) + \varepsilon_i = \eta(x_i,\theta) + \delta(x_i) + \varepsilon_i$$

  - where $\delta(.)$ is model discrepancy
  - and is an unknown function
- Inference about θ is now clearly more complex
  - No longer just a nonlinear regression problem
  - Some literature on correlated errors
- How important is it?
  - That depends on the objectives of the inversion
  - And in particular on the nature of θ

Inversion and the nature of parameters

- Parameters may be physical or just for tuning
  - We adjust tuning parameters so the model fits reality better
    - We are not really interested in their ‘true’ values
  - Physical parameters are different
    - We are often really interested in true physical values
- What are we inverting for?
  - To learn about physical parameter values
    - Model discrepancy is hugely important and needs care and thought
  - To predict reality – within context and range of observations
    - Interpolation: model discrepancy is important but easily addressed
  - To predict reality – outside context/range of observations
    - Extrapolation: discrepancy hugely important, needs care and thought

Example 1: A simple machine (SM)

- A machine produces an amount of work y which depends on the amount of effort x put into it
  - Ideally, $y = \beta x$
    - Parameter β is the rate at which effort can be converted to work
    - It's a physical parameter
  - True value of β is $\beta^* = 0.65$
- Graph shows observed data
  - Points lie below $y = 0.65x$
    - For large enough x
  - Because of losses due to friction etc.

$$\zeta(x) = 0.65x(1 + 0.05x)^{-1}$$

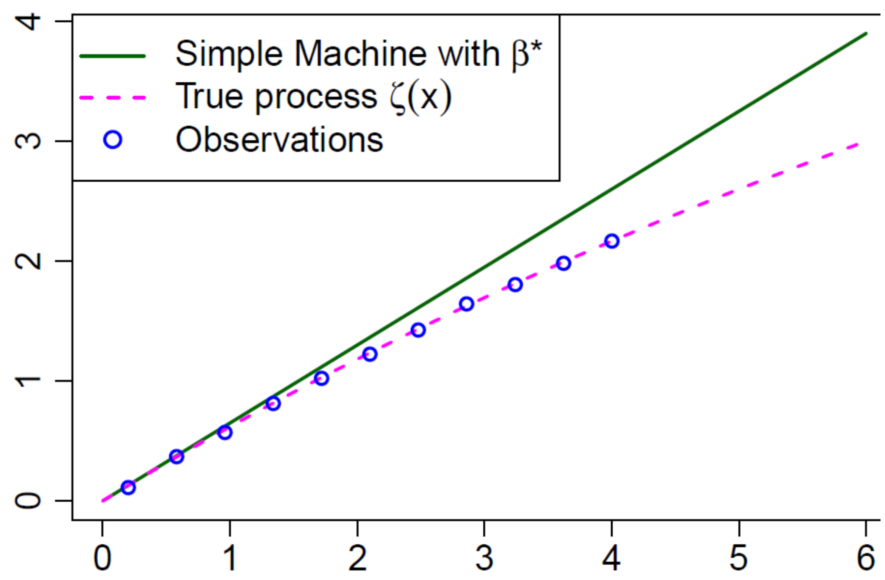

Example 2: Hot and cold (HC)

- An object is placed in a hot medium
  - Initially it heats up but then cools as the medium cools
- Simulator

$$\eta(t,\theta) = \theta_1 t \exp(-\theta_2 t)$$

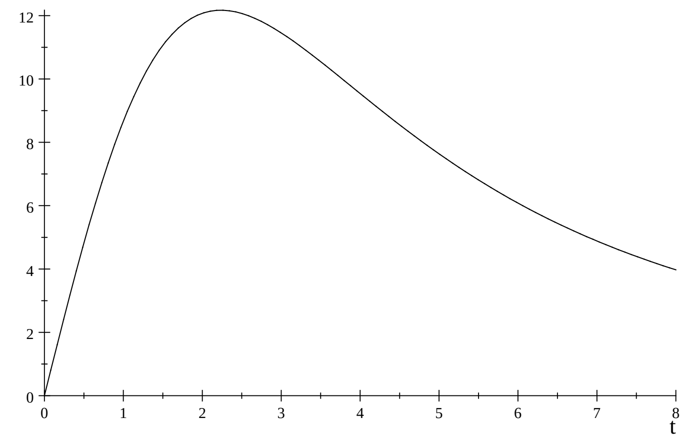

- Reality

$$\zeta(t) = 10t(1 + t^2/10)^{-1.5}$$

- Parameters are physical
  - θ_1 is initial heating rate, a property of the object
  - θ_2 controls the cooling, a property of the medium and setup
- Interested in parameters but also in
  - Maximum temperature ζ_{max} and time t_{max} when max is reached

Meaning of parameters

- What is the relationship between parameters and reality?
  - They don't appear in $\zeta(.)$
  - In the SM example, β is the gradient at the origin
    - Theoretical efficiency only achievable at low inputs
    - This is well defined for reality, $\beta = 0.65$
  - In the HC example, θ_1 is the gradient at the origin
    - Again well defined, $\theta_1 = 10$
  - θ_2 is more difficult because in reality cooling is not exponential
    - We define $\theta_2 = 0.413$ from log gradient at point of inflection
  - ζ_{max} and t_{max} are not really physical
    - From the simulator, $\zeta_{max} = \theta_1 \theta_2^{-1} e^{-1}$, $t_{max} = \theta_2^{-1}$
    - In reality, $\zeta_{max} = 12.172$ and $t_{max} = 2.236$ depend on θ and the setup

Ignoring model discrepancy

SM assuming no discrepancy

- Following the usual approach, inversion is a simple matter of linear regression through the origin

$$z_i = \beta x_i + \varepsilon_i$$

- Here are some results from various sample sizes spread uniformly over 3 ranges of x values

Range	[0.1,2]	[0.1,4]	[2,6]
n=11	0.549 (0.063)	0.562 (0.029)	0.533 (0.023)
n=31	0.656 (0.038)	0.570 (0.017)	0.529 (0.011)
n=91	0.611 (0.021)	0.571 (0.012)	0.528 (0.007)
n infinite	0.605 (0)	0.565 (0)	0.529 (0)

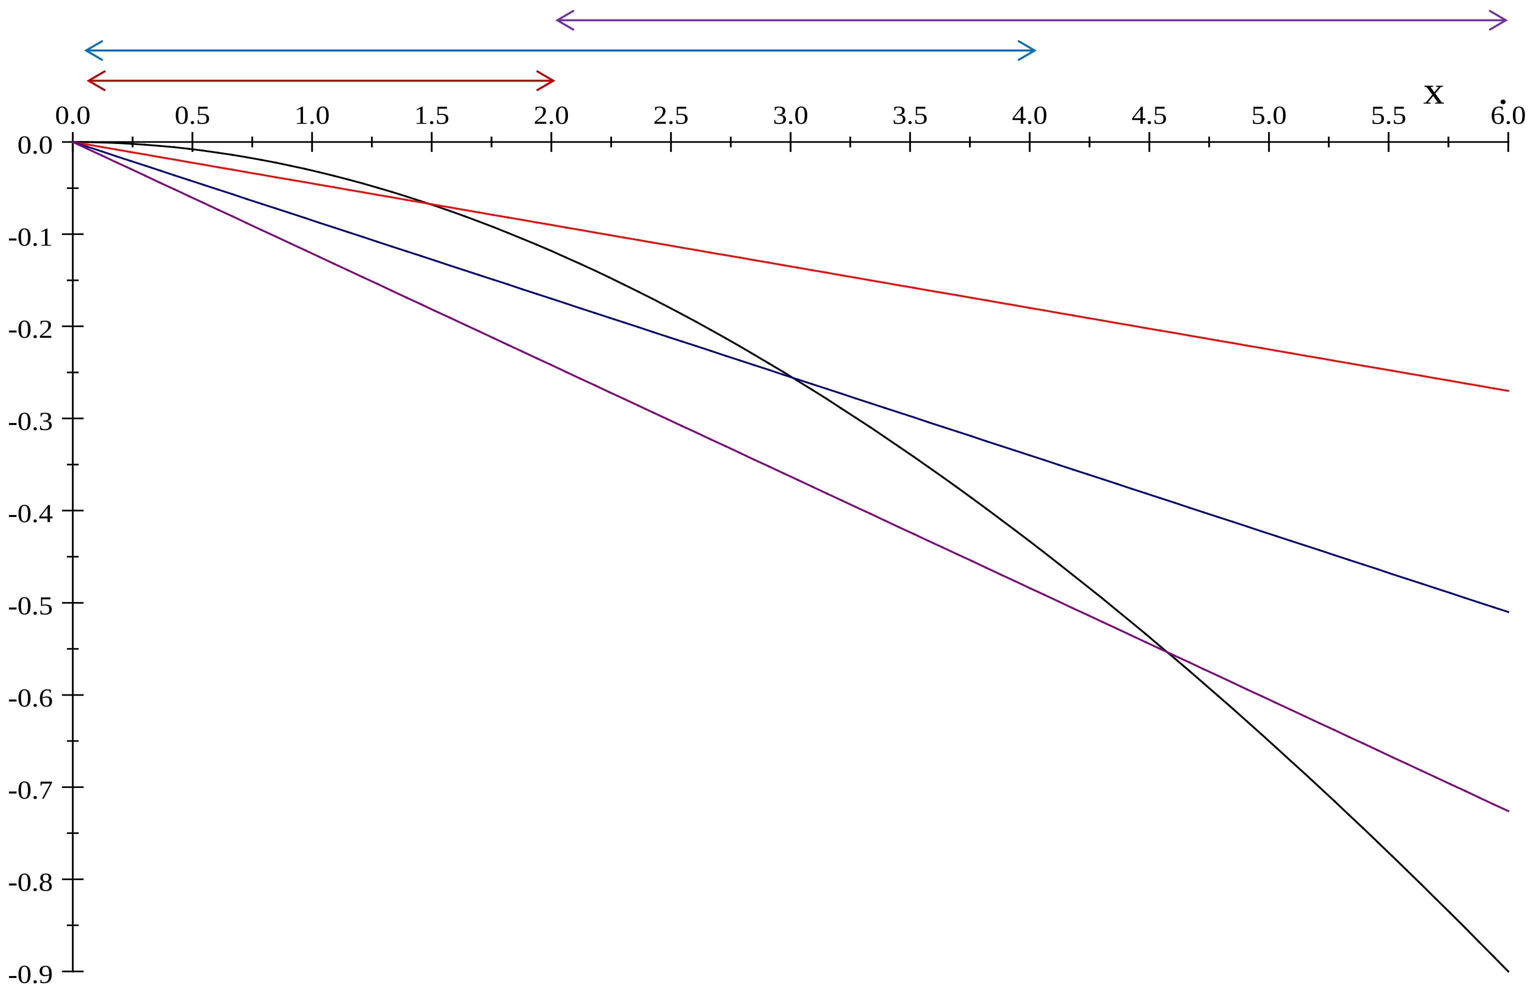
x
0.0
0.5
1.0
1.5
2.0
2.5
3.0
3.5
4.0
4.5
5.0
5.5
6.0
0.0
-0.1
-0.2
-0.3
-0.4
-0.5
-0.6
-0.7
-0.8
-0.9

HC assuming no discrepancy

- These results are from samples of 91 observations over three different ranges
  - Almost every single posterior distribution is concentrated far from the true value

Range	[0.1,1]	[0.2,2]	[0.4,4]	TRUE
θ_1	10.57 (0.11)	11.11 (0.13)	12.77 (0.19)	10
θ_2	0.159 (0.049)	0.237 (0.023)	0.401 (0.033)	0.413
t_{max}	5.00 (0.47)	3.61 (0.13)	2.52 (0.04)	2.24
ζ_{max}	19.42 (1.55)	14.75 (0.38)	11.85 (0.08)	12.17

The problem is completely general

- Inverting (calibrating, tuning, matching) a wrong model gives parameter estimates that are wrong
  - Not equal to their true physical values – biased
  - With more data we become more sure of these wrong values
- The SM and HC are trivial models, but the same conclusions apply to all models
  - All models are wrong
  - In more complex models it is just harder to see what is going wrong
  - Even with the SM, it takes a lot of data to see any curvature in reality
- What can we do about this?

The Simple Machine and Model Discrepancy

SM revisited

- Kennedy and O'Hagan (2001) introduced discrepancy $\delta(.)$
  - Modelled it as a zero-mean Gaussian process
  - They claimed it acknowledges additional uncertainty
  - And mitigates against over-fitting of θ
- So add this model discrepancy term to the linear model of the simple machine

$$z_i = \beta x_i + \delta(x_i) + \varepsilon_i$$

- With $\delta(.)$ modelled as a zero-mean GP
- Posterior distribution of β now behaves quite differently
- Results here from extensive study of SM in
  - Brynjarsdóttir, J. and O'Hagan, A. (2014). Learning about physical parameters: The importance of model discrepancy. *Inverse Problems*, **30**, 114007 (24pp), November 2014.

SM – inversion, with discrepancy

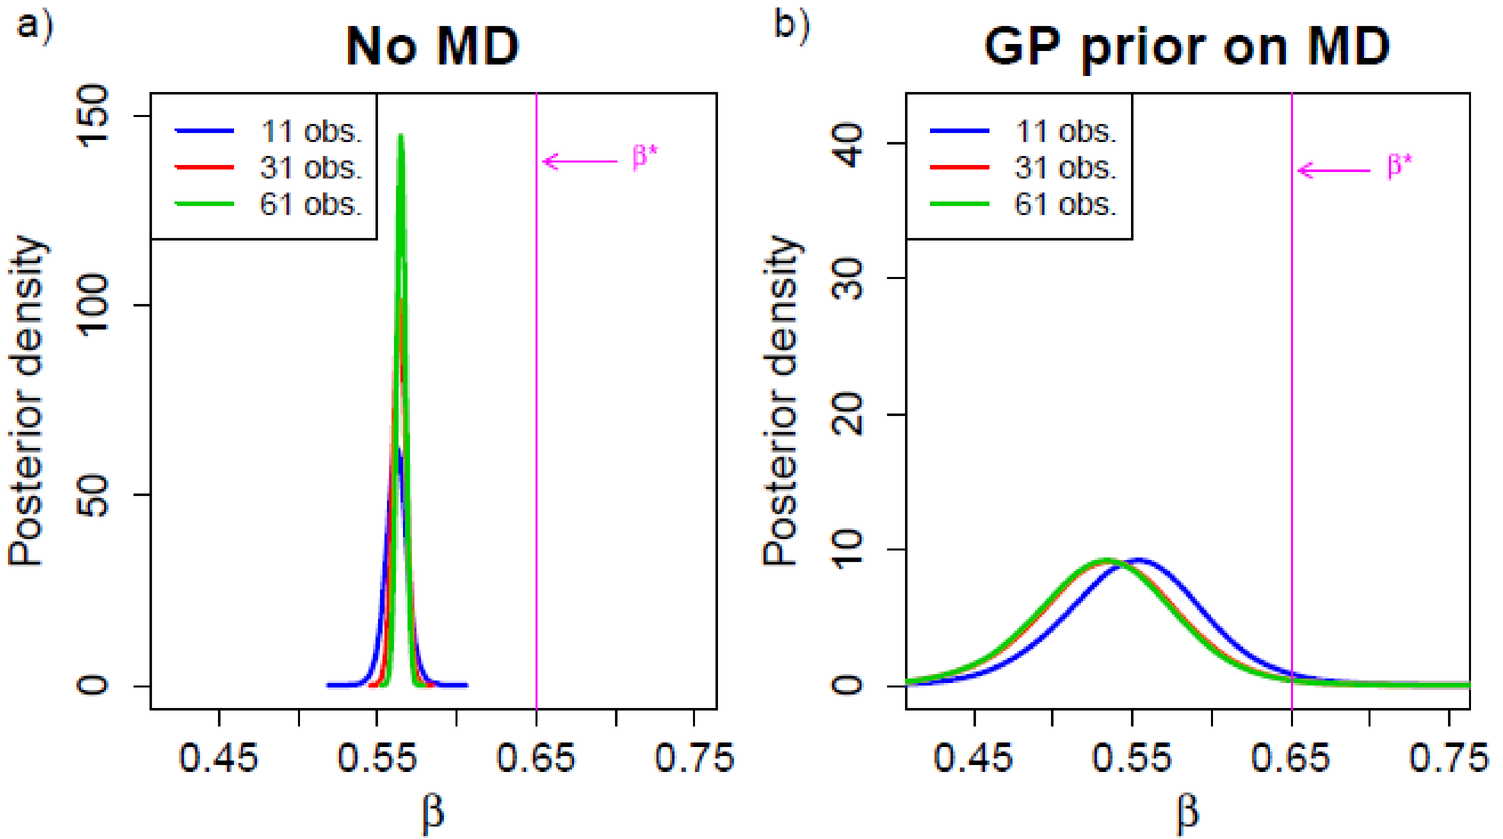

- Posterior distribution much broader and doesn't get worse with more data
  - But still misses the true value

Interpolation

- Main benefit of simple GP model discrepancy is prediction
  - E.g. at $x = 1.5$

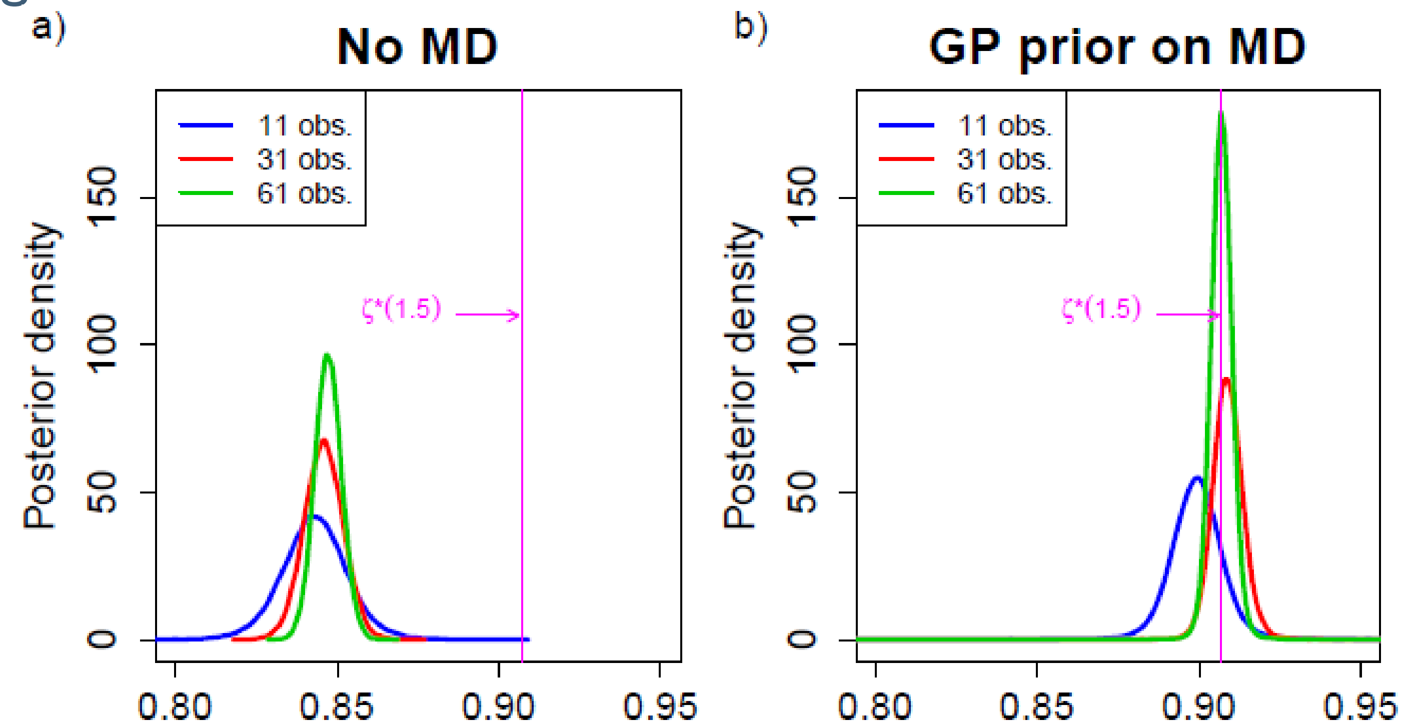

- Prediction within the range of the data is possible
  - And gets better with more data

But when it comes to extrapolation …

- … at $x = 6$

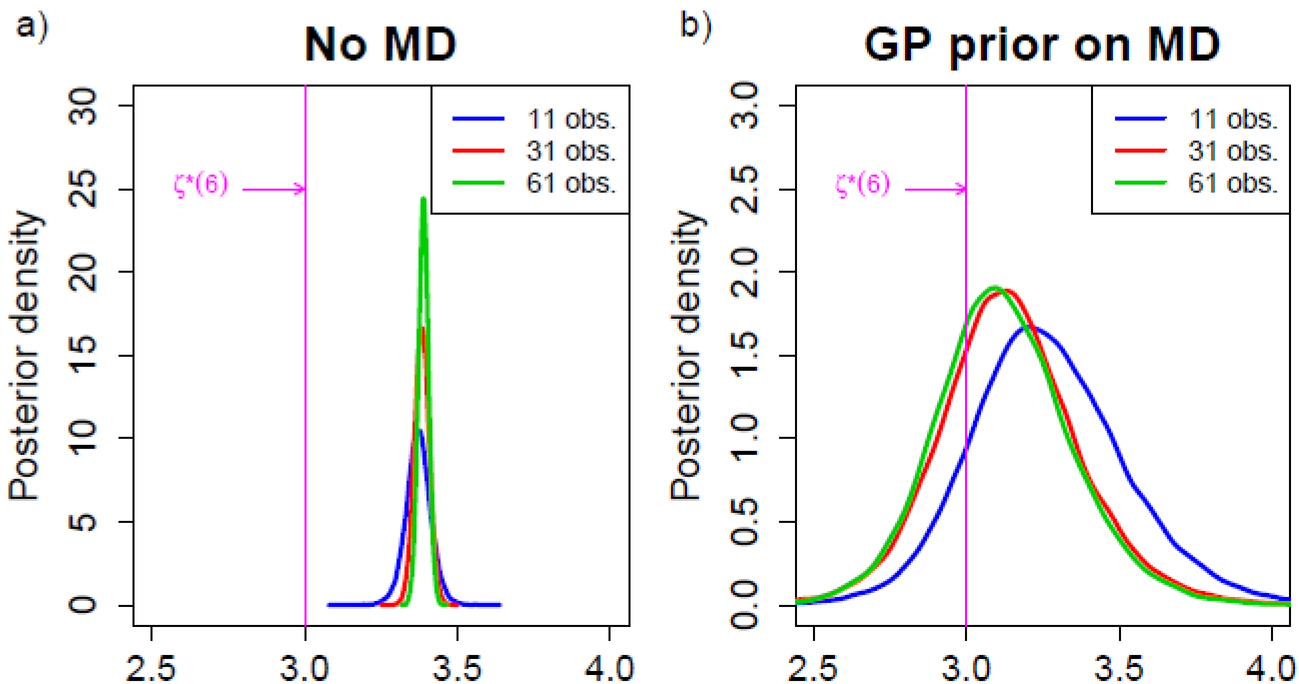

- More data doesn't help because it's all in the range [0, 4]
- Prediction OK here but gets worse for larger x

Extrapolation

- One reason for wish to learn about physical parameters
  - Should be better for extrapolation than just tuning
- Without model discrepancy
  - The parameter estimates will be biased
  - Extrapolation will also be biased
    - Because best fitting parameter values are different in different parts of the control variable space
  - With more data we become more sure of these wrong values
- With GP model discrepancy
  - Extrapolating far from the data does not work
    - No information about model discrepancy
    - Prediction just uses the (calibrated) simulator

We haven't solved the problem

- With simple GP model discrepancy the posterior distribution for θ is typically much wider
  - Increases the chance that we cover the true value
  - But is not very helpful
  - And increasing data does not improve the precision
- Similarly, extrapolation with model discrepancy gives wide prediction intervals
  - And may still not be wide enough
- What's going wrong here?

Nonidentifiability

- Formulation with model discrepancy is not identifiable
- For any θ, there is a $\delta(x)$ to match reality perfectly
  - Reality is $\zeta(x) = \eta(x, \theta) + \delta(x)$
  - Given θ and $\zeta(x)$, model discrepancy is $\delta(x) = \zeta(x) - \eta(x, \theta)$
- Suppose we had an unlimited number of observations
  - We would learn reality's true function (x) exactly
    - Within the range of the data
    - Interpolation works
  - But we would still not learn θ
    - It could in principle be anything
  - And we would still not be able to extrapolate reliably

The joint posterior

- Inversion leads to a joint posterior distribution for θ and δ(x)
- But nonidentifiability means there are many equally good fits (θ, δ(x)) to the data
  - Induces strong correlation between θ and δ(x)
  - This may be compounded by the fact that simulators often have large numbers of parameters
    - (Near-)redundancy means that different θ values produce (almost) identical predictions
    - Sometimes called equifinality
- Within this set, the prior distributions for θ and δ(x) count

The importance of prior information

- The nonparametric GP term allows the model to fit and predict reality accurately given enough data
  - Within the range of the data
- But it **doesn't** mean physical parameters are correctly estimated
  - The separation between original model and discrepancy is unidentified
  - Estimates depend on prior information
  - Unless the real model discrepancy is just the kind expected a priori the physical parameter estimates will still be biased
- To learn about θ in the presence of model discrepancy we need better prior information
  - And this is also crucial for extrapolation

Better prior information

- For calibration
  - Prior information about θ and/or δ(x)
  - We wish to calibrate because prior information about θ is not strong enough
  - So prior knowledge of model discrepancy is crucial
    - In the range of the data
- For extrapolation
  - All this plus good prior knowledge of δ(x) *outside* the range of the calibration data
    - That's seriously challenging!
  - In the SM, a model for δ(x) that says it is zero at x = 0, then increasingly negative, should do better

Inference about the physical parameter

- We conditioned the GP
  - $\delta(0) = 0$
  - $\delta'(0) = 0$
  - $\delta'(0.5) < 0$
  - $\delta'(1.5) < 0$

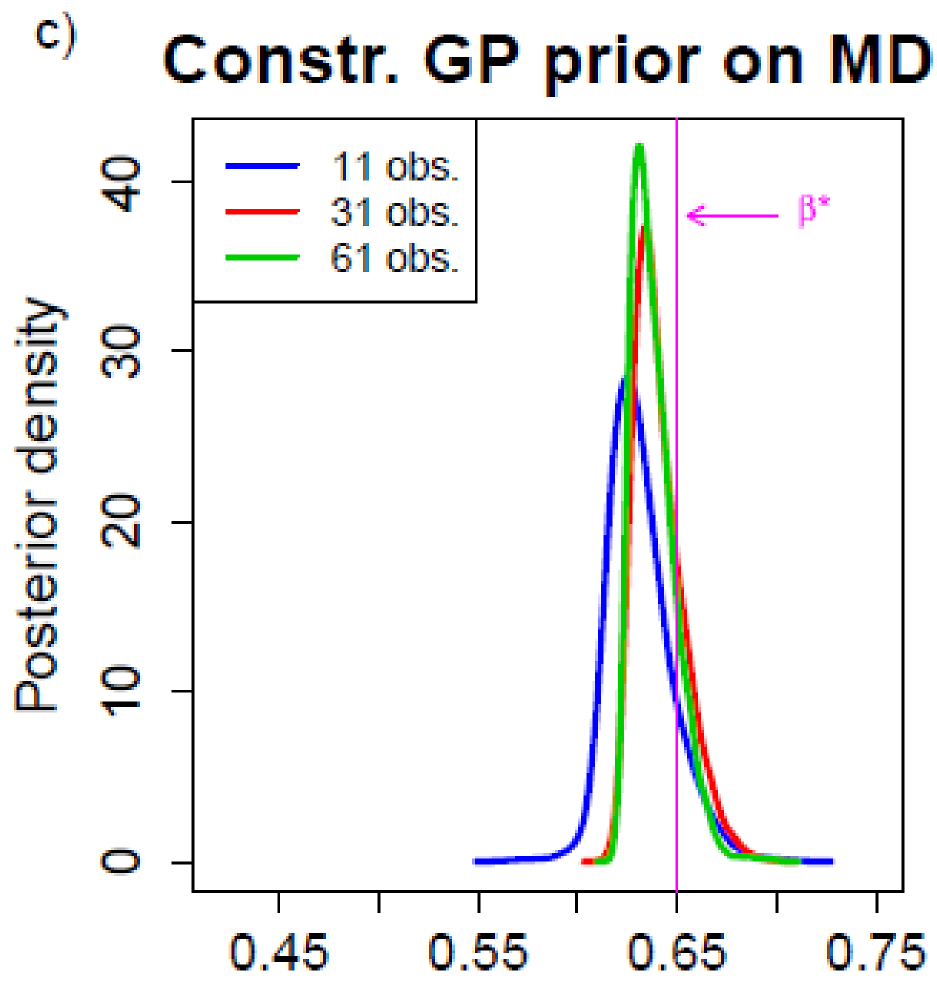

Prediction

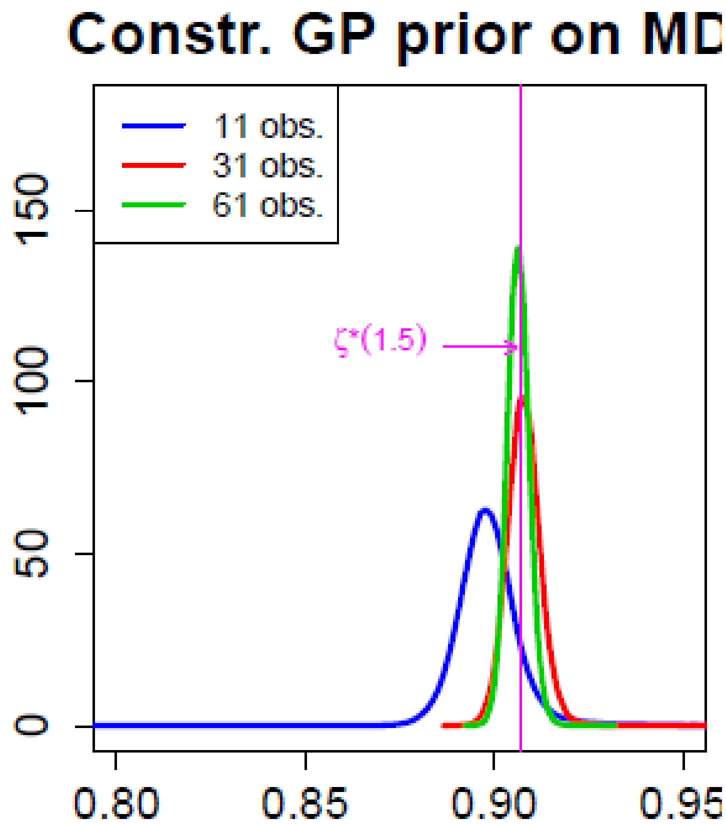

x = 1.5

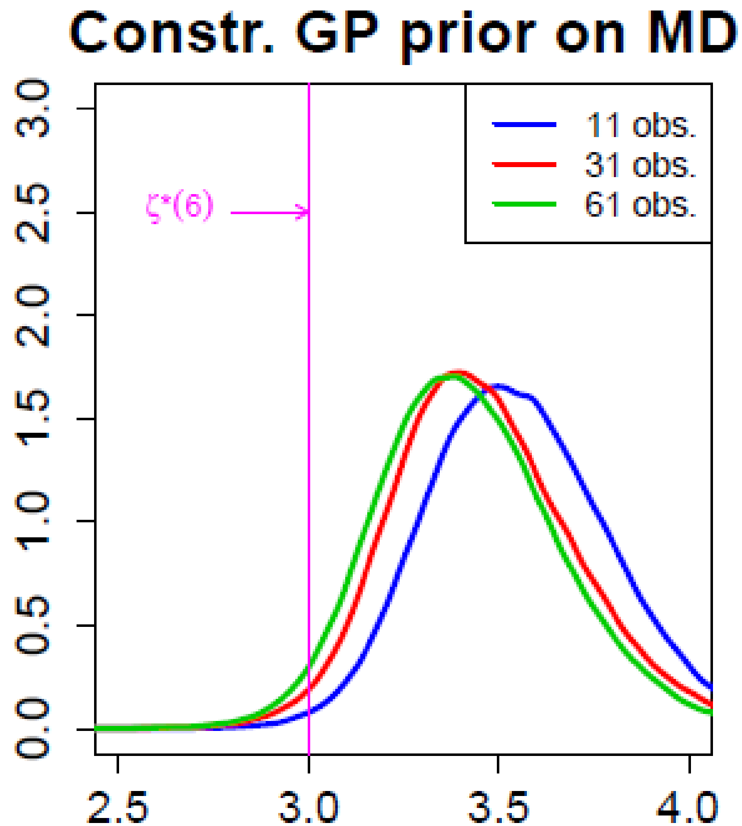

x = 6

Conclusions

Summary

- Without model discrepancy
  - Inference about physical parameters will be wrong
    - And will get worse with more data
  - The same is true of prediction
    - Both interpolation and extrapolation
- With crude GP model discrepancy
  - Interpolation inference is OK
    - And gets better with more data
  - But we still get physical parameters and extrapolation wrong
- The better our prior knowledge about model discrepancy
  - The more chance we have of getting physical parameters right
  - Also extrapolation
    - But then we need even better prior knowledge