Uncertainty analysis for computer models.

Michael Goldstein
Durham University*

*Thanks to EPSRC, NERC, MRc, Leverhulme, for funding. Thanks to Ian Vernon for pictures. Energy systems work with Antony Lawson, Amy Wilson, Chris Dent.

Complex physical models

Most large and complex physical systems are studied by mathematical models, implemented as high dimensional computer simulators. Some examples are:

Complex physical models

Most large and complex physical systems are studied by mathematical models, implemented as high dimensional computer simulators. Some examples are:

Oil reservoirs An oil reservoir simulator is used to manage assets associated with the reservoir. The aim is commercial: to develop efficient production schedules, determine whether and where to sink new wells, and so forth.

Complex physical models

Most large and complex physical systems are studied by mathematical models, implemented as high dimensional computer simulators. Some examples are:

Oil reservoirs An oil reservoir simulator is used to manage assets associated with the reservoir. The aim is commercial: to develop efficient production schedules, determine whether and where to sink new wells, and so forth.

Natural Hazards Floods, volcanoes, tsunamis and so forth, are all studied by large computer simulators.

Complex physical models

Most large and complex physical systems are studied by mathematical models, implemented as high dimensional computer simulators. Some examples are:

Oil reservoirs An oil reservoir simulator is used to manage assets associated with the reservoir. The aim is commercial: to develop efficient production schedules, determine whether and where to sink new wells, and so forth.

Natural Hazards Floods, volcanoes, tsunamis and so forth, are all studied by large computer simulators.

Climate change Large scale climate simulators are constructed to assess likely effects of human intervention upon future climate behaviour.

Complex physical models

Most large and complex physical systems are studied by mathematical models, implemented as high dimensional computer simulators. Some examples are:

Oil reservoirs An oil reservoir simulator is used to manage assets associated with the reservoir. The aim is commercial: to develop efficient production schedules, determine whether and where to sink new wells, and so forth.

Natural Hazards Floods, volcanoes, tsunamis and so forth, are all studied by large computer simulators.

Climate change Large scale climate simulators are constructed to assess likely effects of human intervention upon future climate behaviour.

Galaxy formation The study of the development of the Universe is carried out by using a Galaxy formation simulator. The aim is scientific - to gain information about the physical processes underlying the Universe.

Complex physical models

Most large and complex physical systems are studied by mathematical models, implemented as high dimensional computer simulators. Some examples are:

Oil reservoirs An oil reservoir simulator is used to manage assets associated with the reservoir. The aim is commercial: to develop efficient production schedules, determine whether and where to sink new wells, and so forth.

Natural Hazards Floods, volcanoes, tsunamis and so forth, are all studied by large computer simulators.

Climate change Large scale climate simulators are constructed to assess likely effects of human intervention upon future climate behaviour.

Galaxy formation The study of the development of the Universe is carried out by using a Galaxy formation simulator. The aim is scientific - to gain information about the physical processes underlying the Universe.

Energy planning Simulators of future energy demand and provision are key components of planning for energy investment.

Complex physical models

Most large and complex physical systems are studied by mathematical models, implemented as high dimensional computer simulators. Some examples are:

Oil reservoirs An oil reservoir simulator is used to manage assets associated with the reservoir. The aim is commercial: to develop efficient production schedules, determine whether and where to sink new wells, and so forth.

Natural Hazards Floods, volcanoes, tsunamis and so forth, are all studied by large computer simulators.

Climate change Large scale climate simulators are constructed to assess likely effects of human intervention upon future climate behaviour.

Galaxy formation The study of the development of the Universe is carried out by using a Galaxy formation simulator. The aim is scientific - to gain information about the physical processes underlying the Universe.

Energy planning Simulators of future energy demand and provision are key components of planning for energy investment.

The science in each of these applications is completely different. However, the underlying methodology for handling uncertainty is the same.

General structure of problem

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

General structure of problem

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

Each simulator can be conceived as a function $f(x)$, where
x: input vector, representing unknown properties of the physical system;
$f(x)$: output vector representing system behaviour.

General structure of problem

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

Each simulator can be conceived as a function $f(x)$, where
x: input vector, representing unknown properties of the physical system;
$f(x)$: output vector representing system behaviour.

Interest in general qualitative insights plus some of the following.

General structure of problem

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

Each simulator can be conceived as a function $f(x)$, where
x: input vector, representing unknown properties of the physical system;
$f(x)$: output vector representing system behaviour.

Interest in general qualitative insights plus some of the following.

the “appropriate” (in some sense) choice, x^*, for the system properties x,

General structure of problem

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

Each simulator can be conceived as a function $f(x)$, where
x: input vector, representing unknown properties of the physical system;
$f(x)$: output vector representing system behaviour.

Interest in general qualitative insights plus some of the following.

the “appropriate” (in some sense) choice, x^*, for the system properties x,

how informative $f(x^*)$ is for actual system behaviour, y.

General structure of problem

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

Each simulator can be conceived as a function $f(x)$, where
x: input vector, representing unknown properties of the physical system;
$f(x)$: output vector representing system behaviour.

Interest in general qualitative insights plus some of the following.

the "appropriate" (in some sense) choice, x^*, for the system properties x,

how informative $f(x^*)$ is for actual system behaviour, y.

the use that we can make of historical observations z, observed with error on a subset y_h of y, corresponding to a portion $f_h(x)$ of the model, both to test the model and to constrain the input space.

General structure of problem

Different physical models vary in many aspects, but the formal structures for analysing the physical system through computer simulators are very similar (which is why there is a common underlying methodology).

Each simulator can be conceived as a function $f(x)$, where
x: input vector, representing unknown properties of the physical system;
$f(x)$: output vector representing system behaviour.

Interest in general qualitative insights plus some of the following.

the "appropriate" (in some sense) choice, x^*, for the system properties x,

how informative $f(x^*)$ is for actual system behaviour, y.

the use that we can make of historical observations z, observed with error on a subset y_h of y, corresponding to a portion $f_h(x)$ of the model, both to test the model and to constrain the input space.

the optimal assignment of any decision inputs, d, in the model.

Simple 1D Exponential Growth Example

- Say we are interested in the concentration of a chemical which evolves in time. We will model this concentration as $f(x,t)$ where x is a rate parameter and t is time.

Simple 1D Exponential Growth Example

- Say we are interested in the concentration of a chemical which evolves in time. We will model this concentration as $f(x,t)$ where x is a rate parameter and t is time.
- We think $f(x,t)$ satisfies the differential equation or model:

$$\frac{df(x,t)}{dt} = x f(x,t) \quad \Longrightarrow \quad f(x,t) = f_0 \exp(xt)$$

Simple 1D Exponential Growth Example

- Say we are interested in the concentration of a chemical which evolves in time. We will model this concentration as $f(x,t)$ where x is a rate parameter and t is time.
- We think $f(x,t)$ satisfies the differential equation or model:

$$\frac{df(x,t)}{dt} = xf(x,t) \qquad \Longrightarrow \qquad f(x,t) = f_0 \exp(xt)$$

- We will temporarily assume the initial conditions are $f_0 = f(x, t = 0) = 1$.

Simple 1D Exponential Growth Example

- Say we are interested in the concentration of a chemical which evolves in time. We will model this concentration as $f(x,t)$ where x is a rate parameter and t is time.

- We think $f(x,t)$ satisfies the differential equation or model:

$$\frac{df(x,t)}{dt} = x f(x,t) \qquad \Longrightarrow \qquad f(x,t) = f_0 \exp(xt)$$

- We will temporarily assume the initial conditions are $f_0 = f(x, t = 0) = 1$.

- Model features an input parameter x which we want to learn about.

Simple 1D Exponential Growth Example

- Say we are interested in the concentration of a chemical which evolves in time. We will model this concentration as $f(x,t)$ where x is a rate parameter and t is time.

- We think $f(x,t)$ satisfies the differential equation or model:

$$\frac{df(x,t)}{dt} = xf(x,t) \qquad \Longrightarrow \qquad f(x,t) = f_0 \exp(xt)$$

- We will temporarily assume the initial conditions are $f_0 = f(x, t=0) = 1$.

- Model features an input parameter x which we want to learn about.

- Note that normally we would not have the analytic solution for $f(x,t)$.

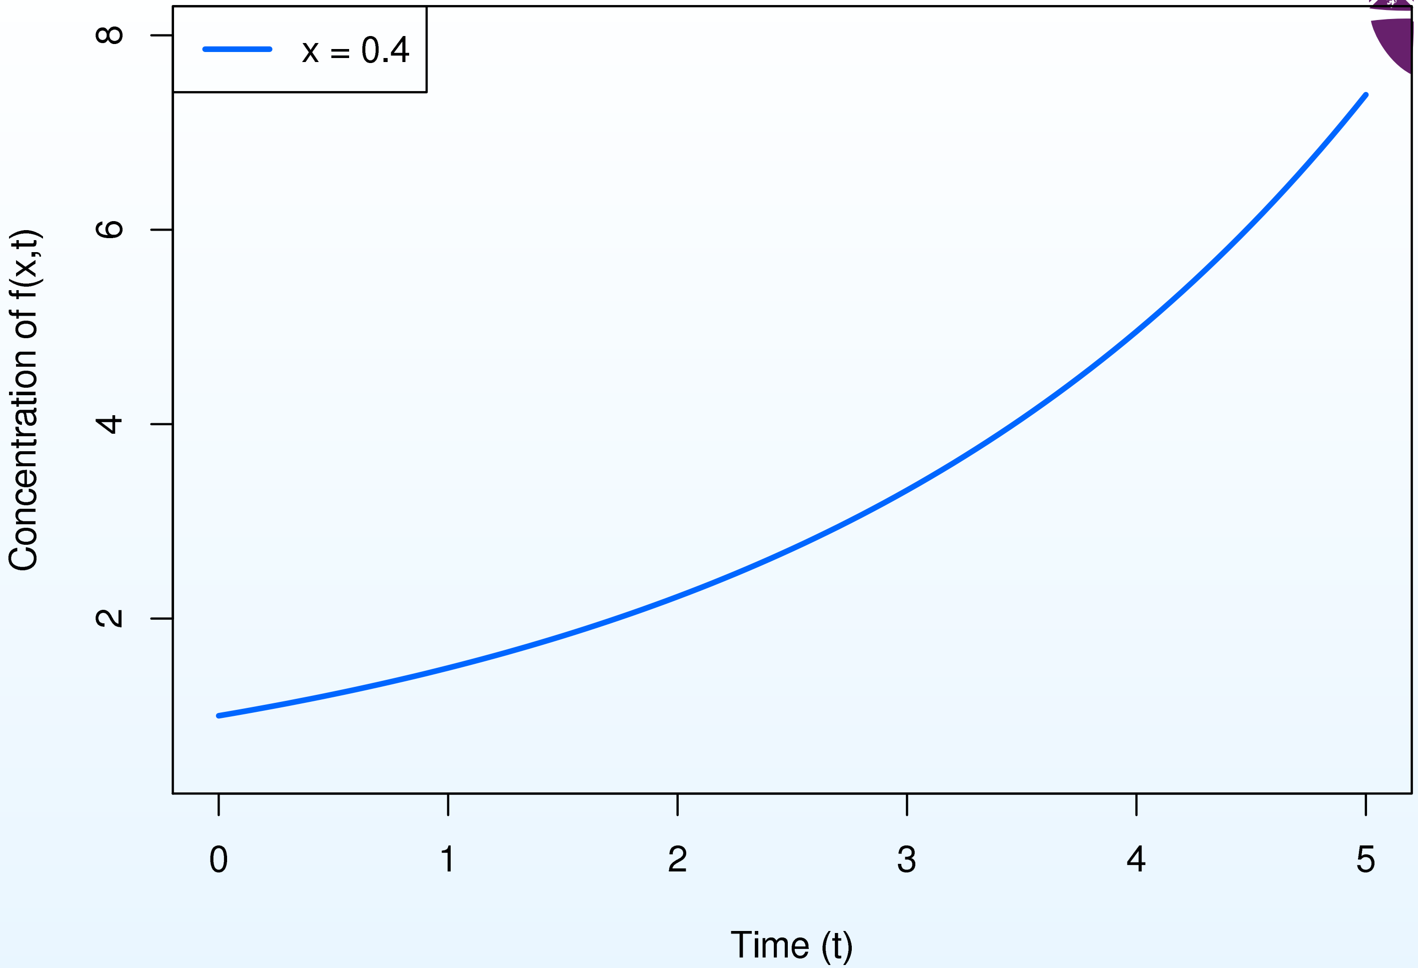

- One “model run” with the input parameter $x = 0.4$

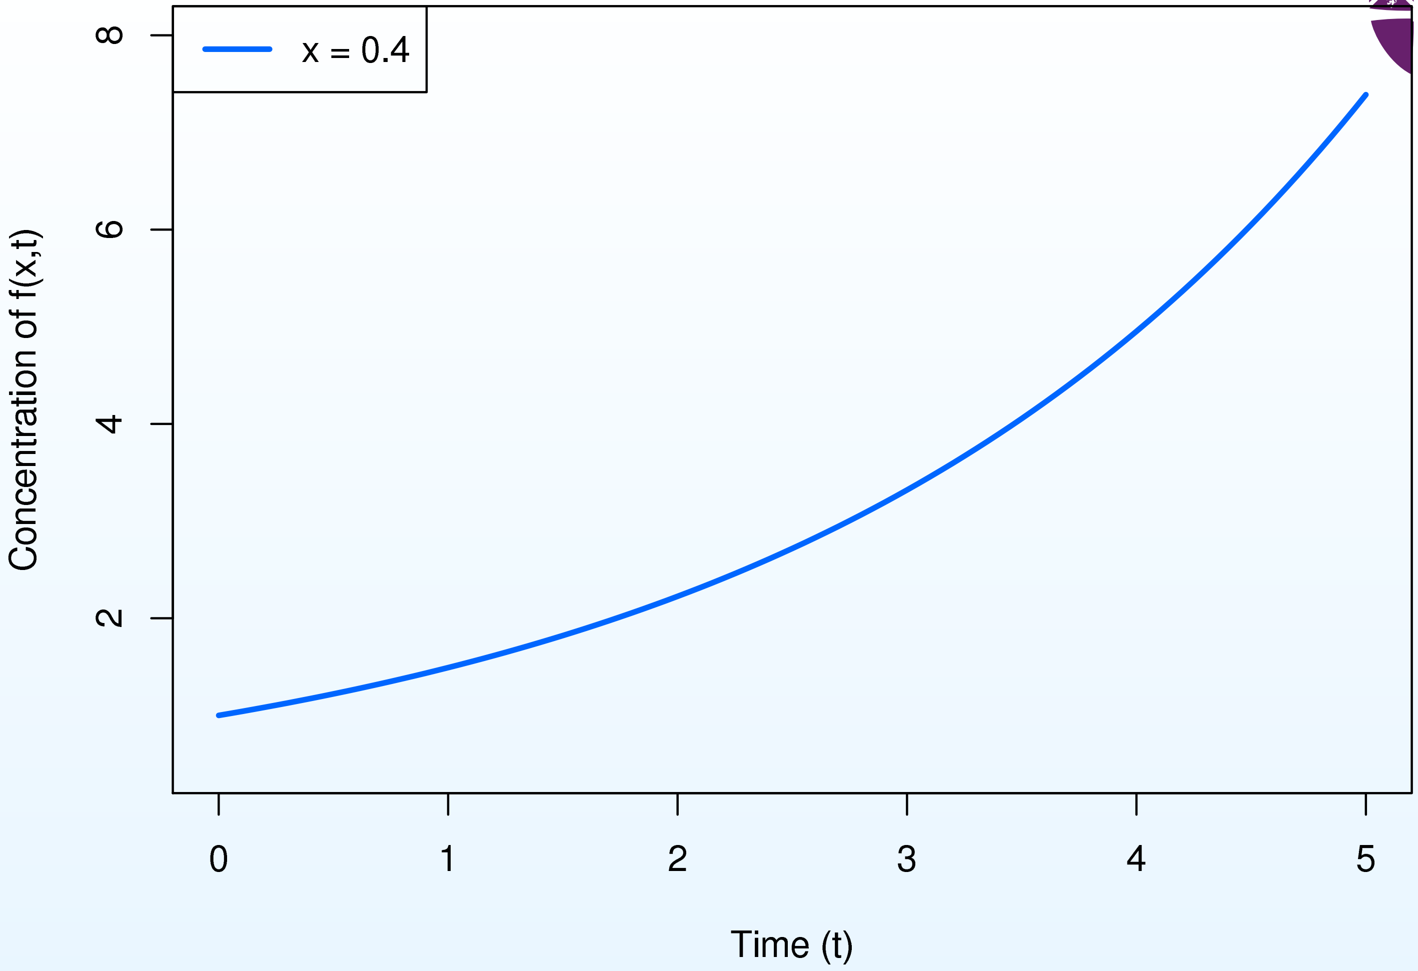

- One “model run” with the input parameter $x = 0.4$
- If we did not know the analytic solution for $f(x,t)$ this would be generated by numerically solving the differential equation.

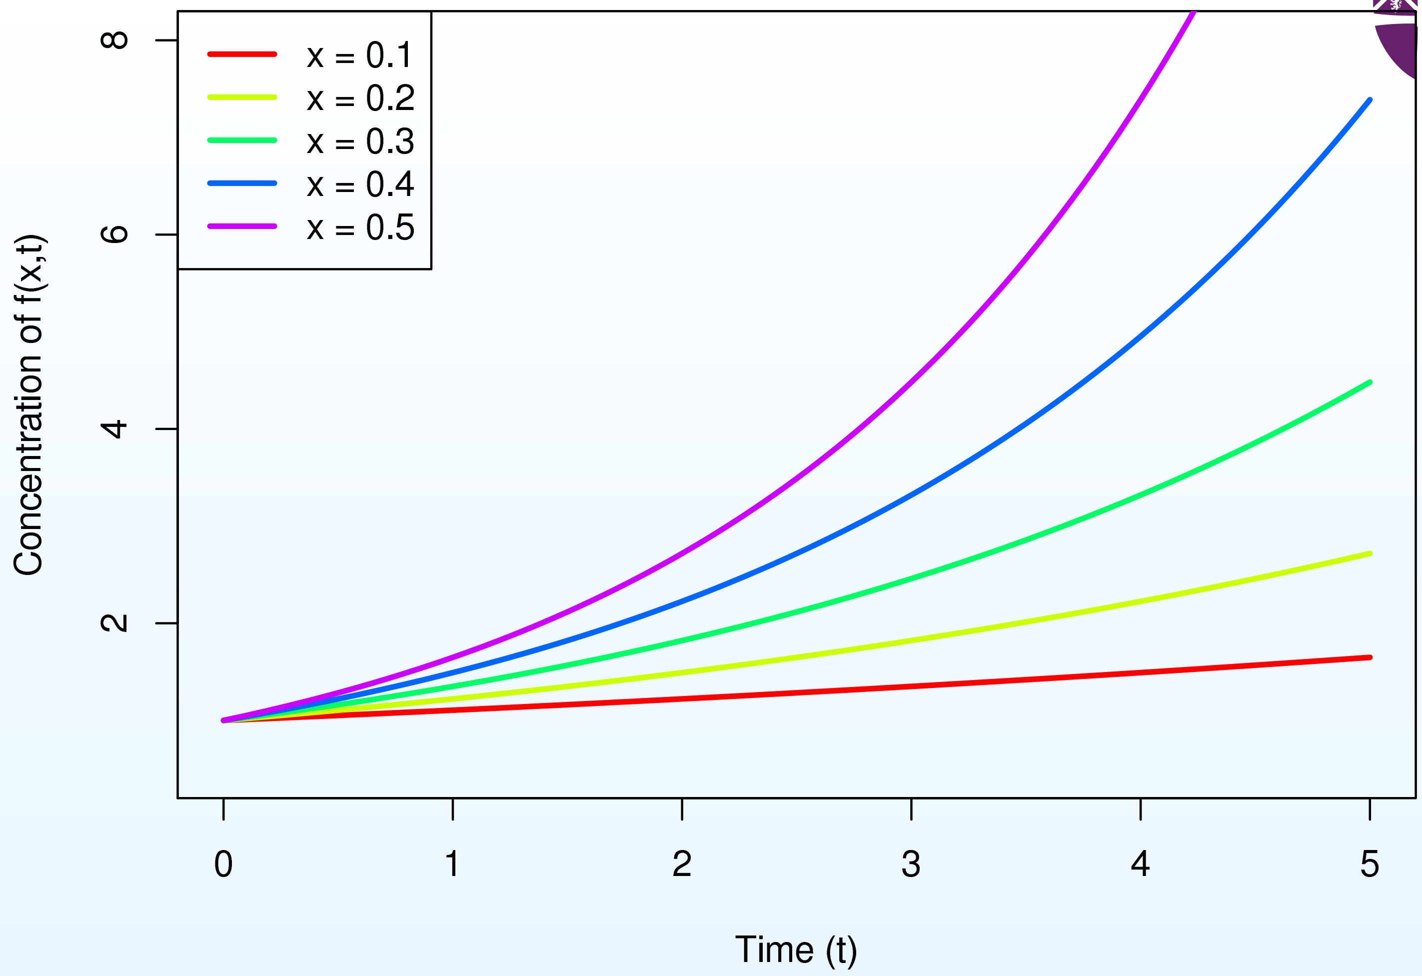

- Five model runs with the input parameter varying from $x = 0.1$ to $x = 0.5$

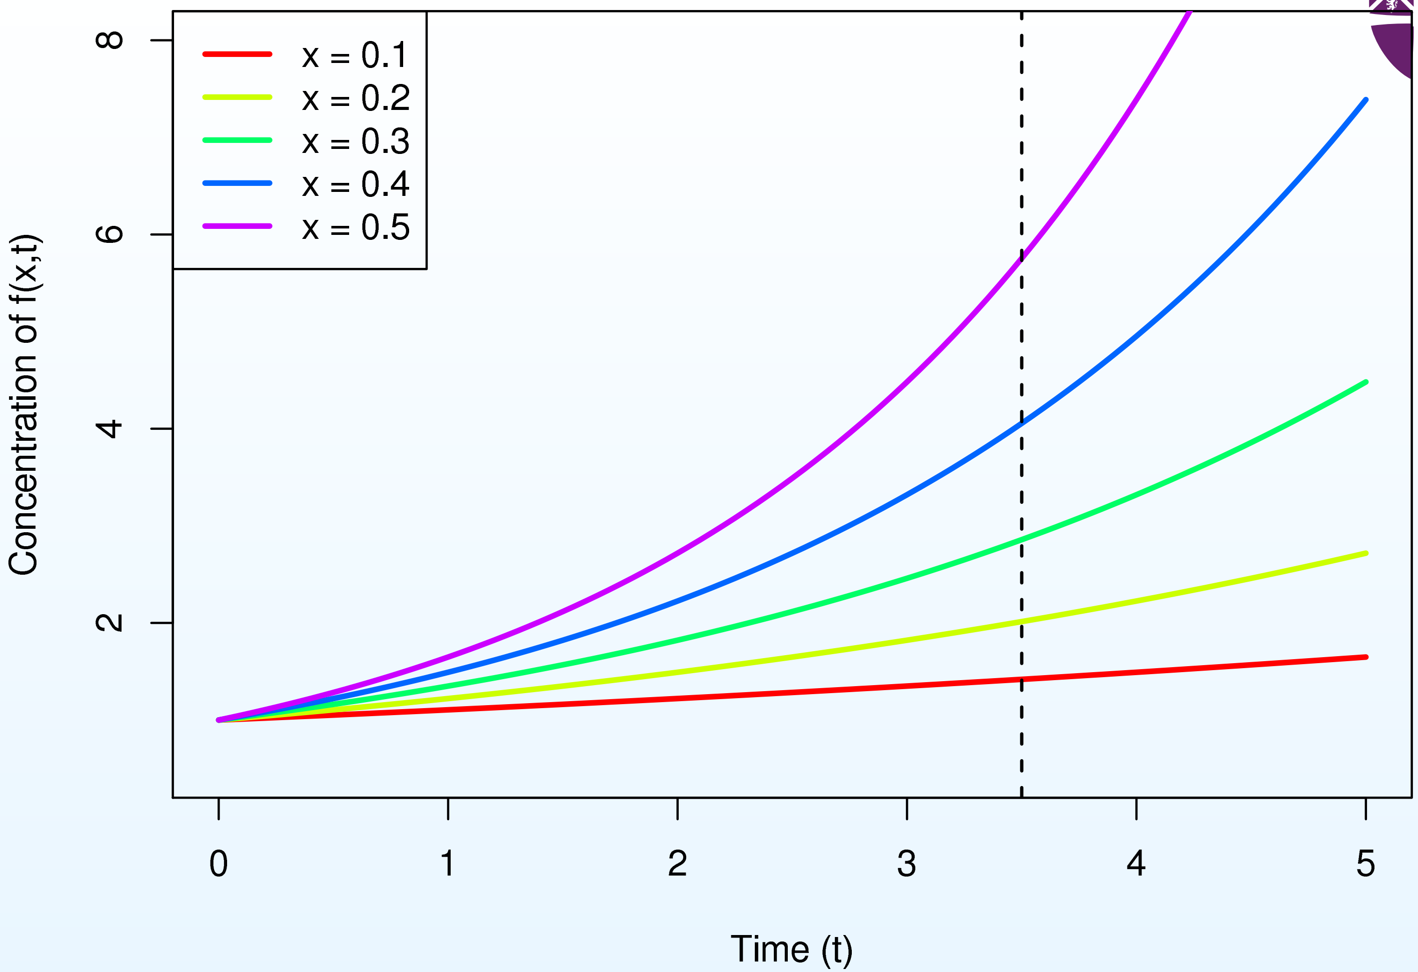

- Five model runs with the input parameter varying from $x = 0.1$ to $x = 0.5$
- We are going to measure $f(x, t)$ at $t = 3.5$

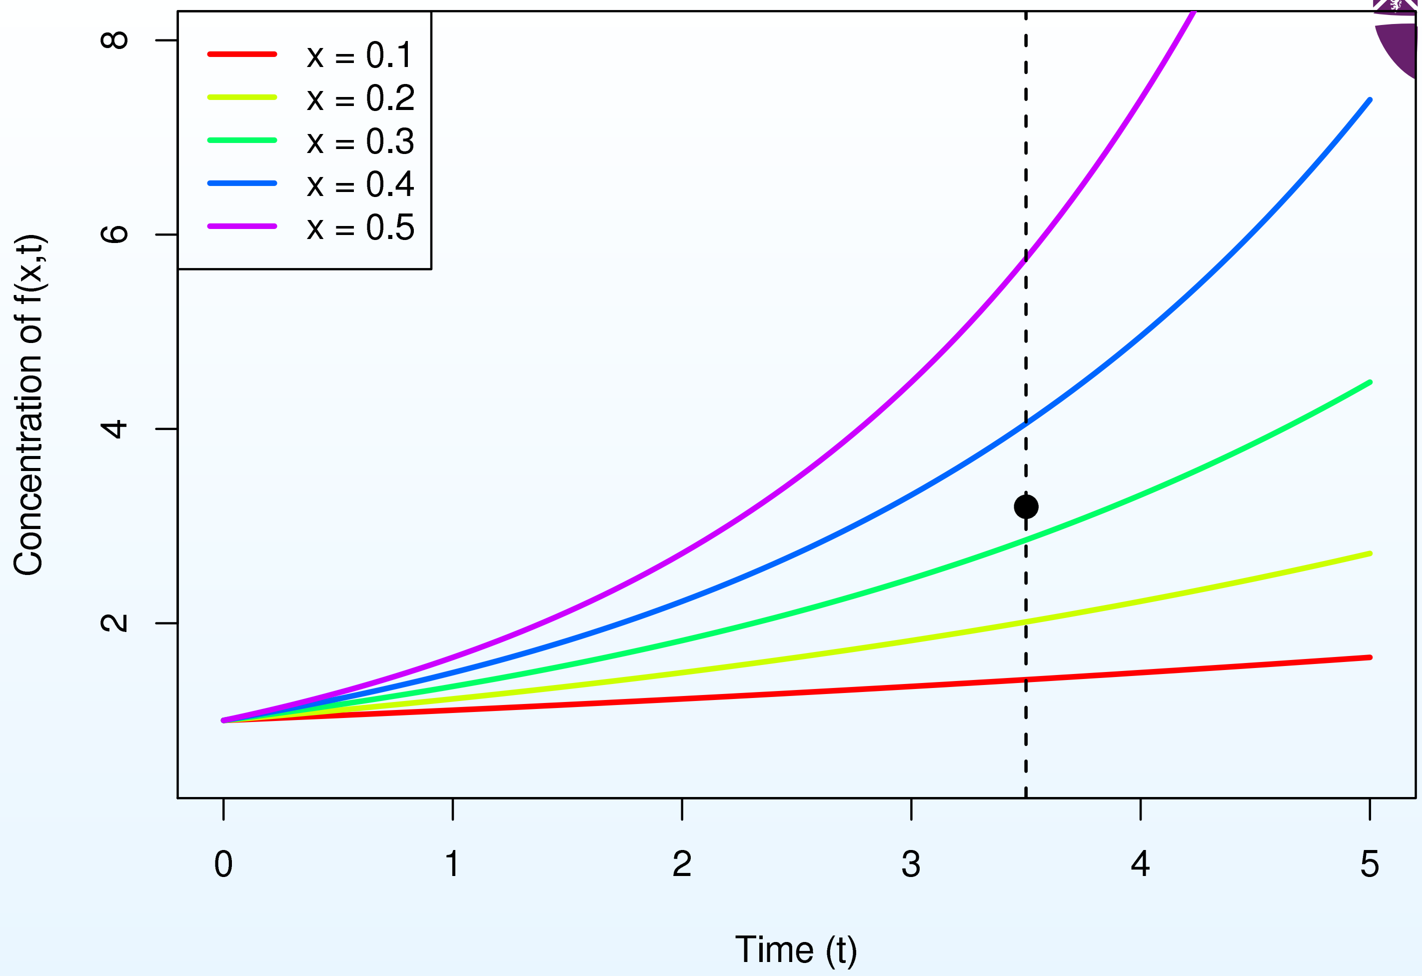

- Five model runs with the input parameter varying from $x = 0.1$ to $x = 0.5$
- We are going to measure $f(x, t)$ at $t = 3.5$

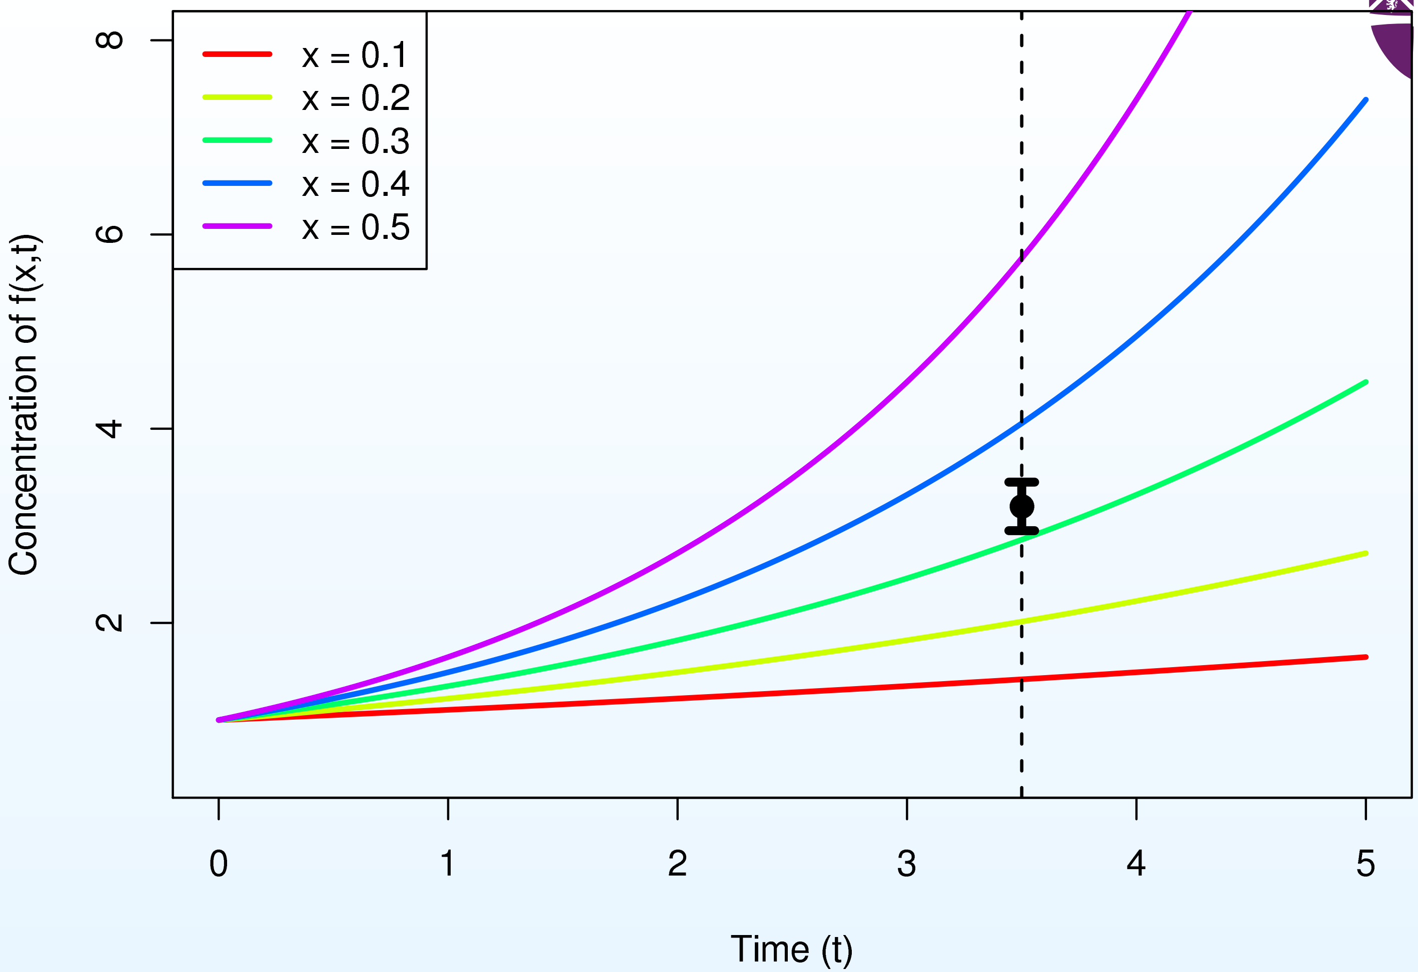

- Five model runs with the input parameter varying from $x = 0.1$ to $x = 0.5$
- We are going to measure $f(x, t)$ at $t = 3.5$
- The measurement is not a point but comes with measurement error.

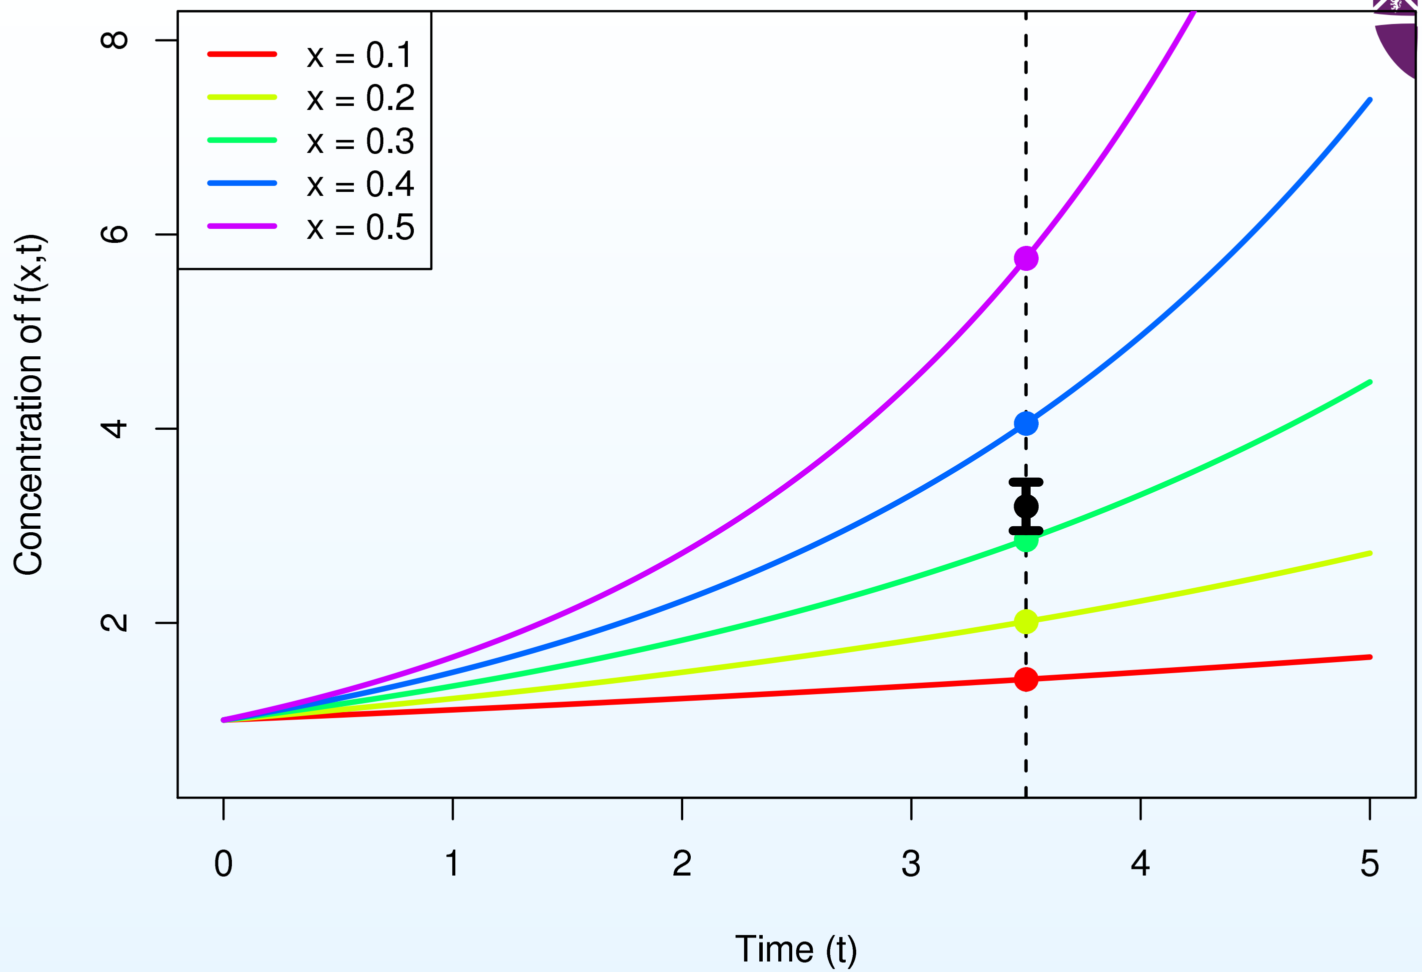

- Question: which values of x ensure the output $f(x, t = 3.5)$ is consistent with the observations?

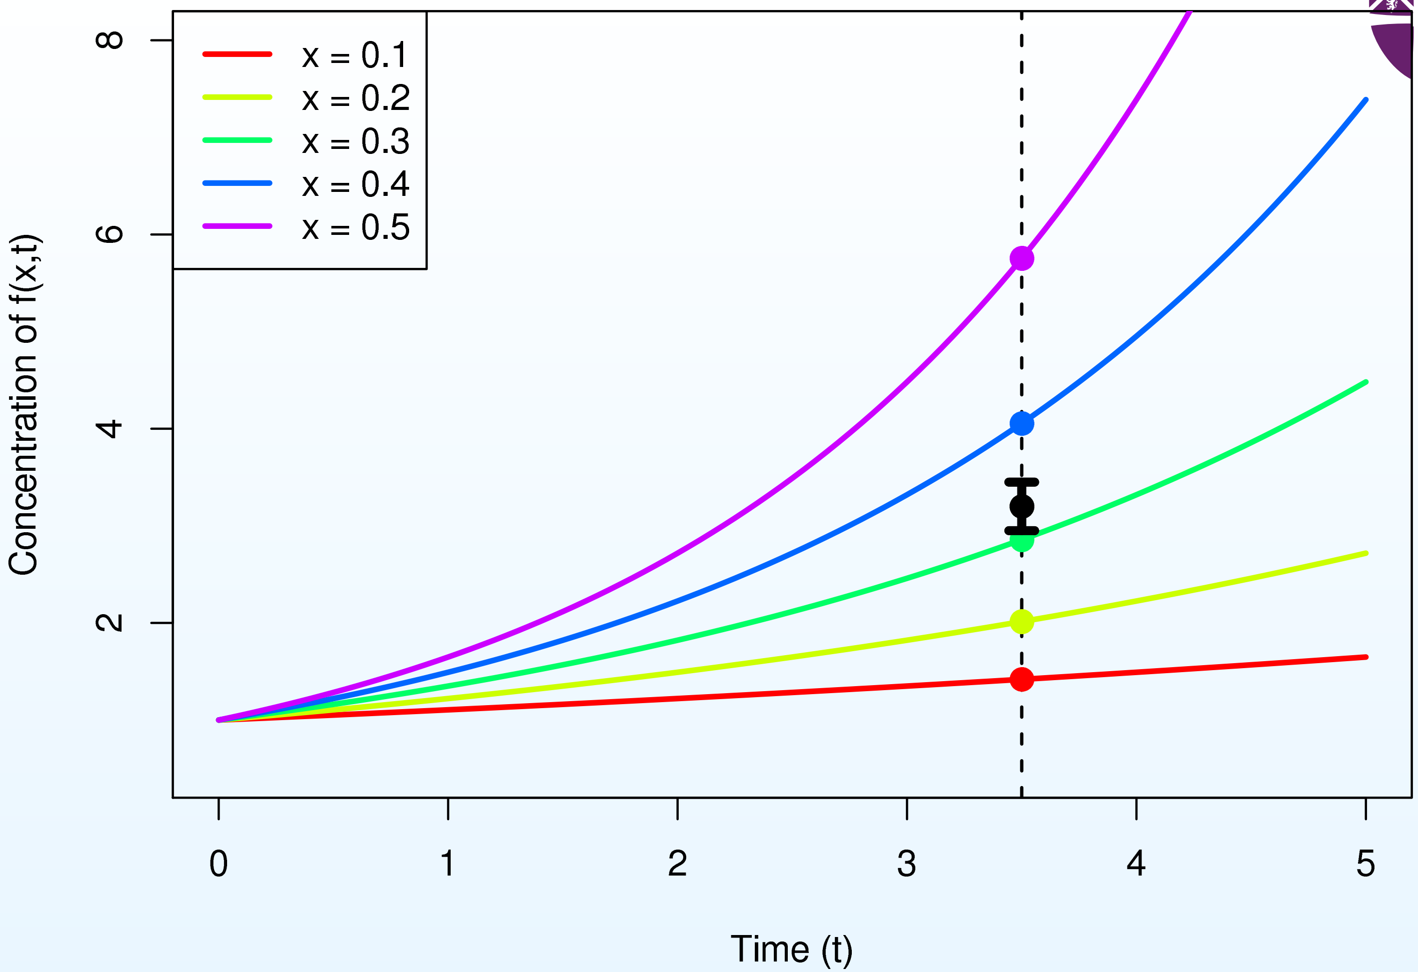

- Question: which values of x ensure the output $f(x, t = 3.5)$ is consistent with the observations?
- It would seem that x has to be at least between 0.3 and 0.4.

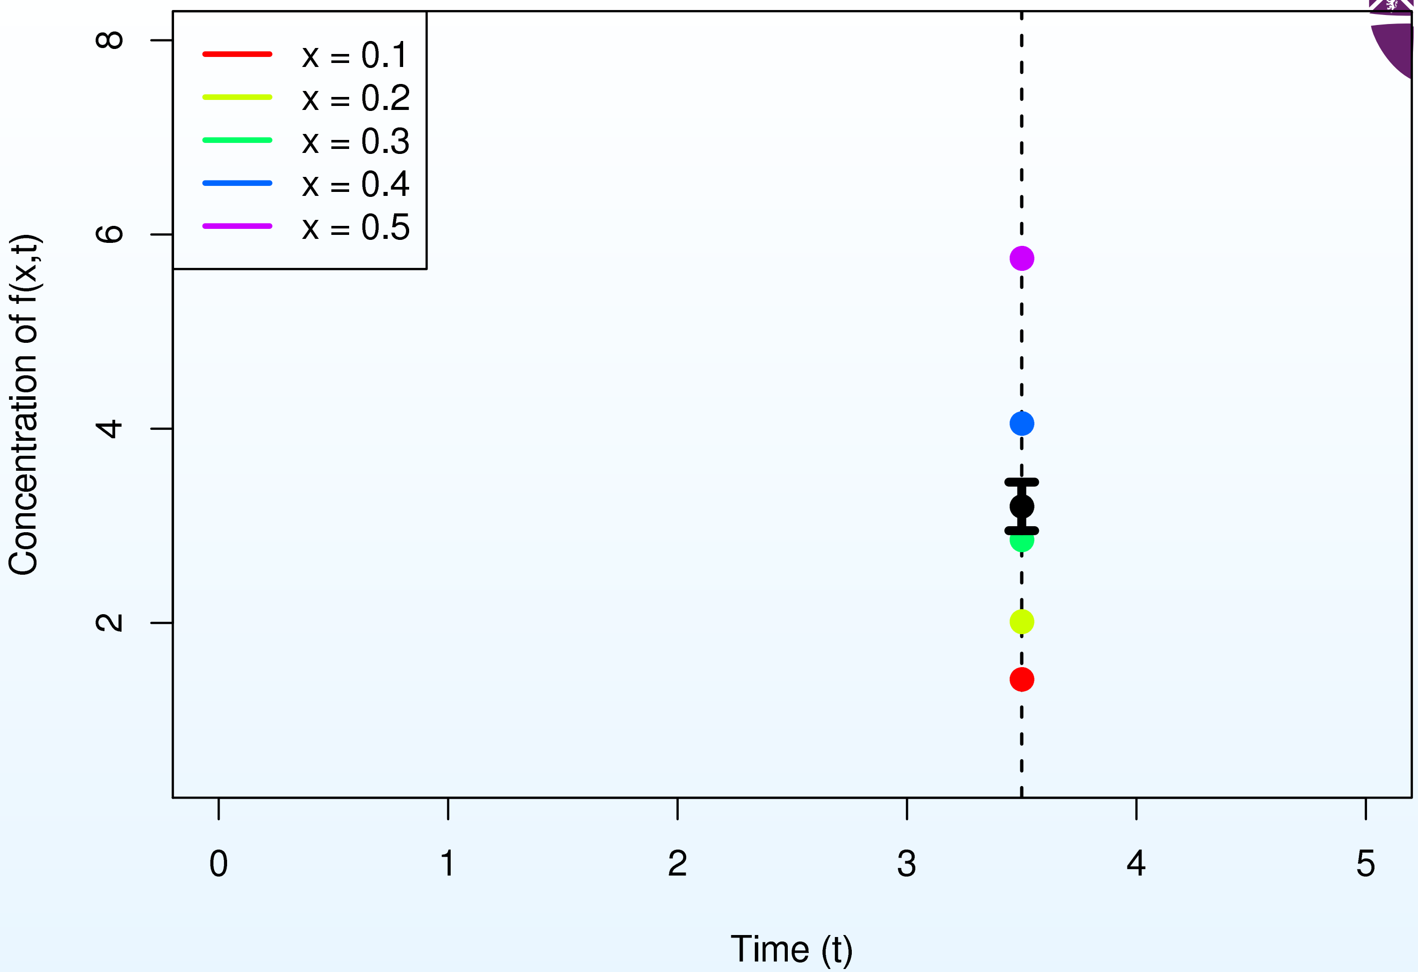

- To answer this, we can now discard other values of $f(x,t)$ and think of $f(x,t=3.5)$ as a function of x only.

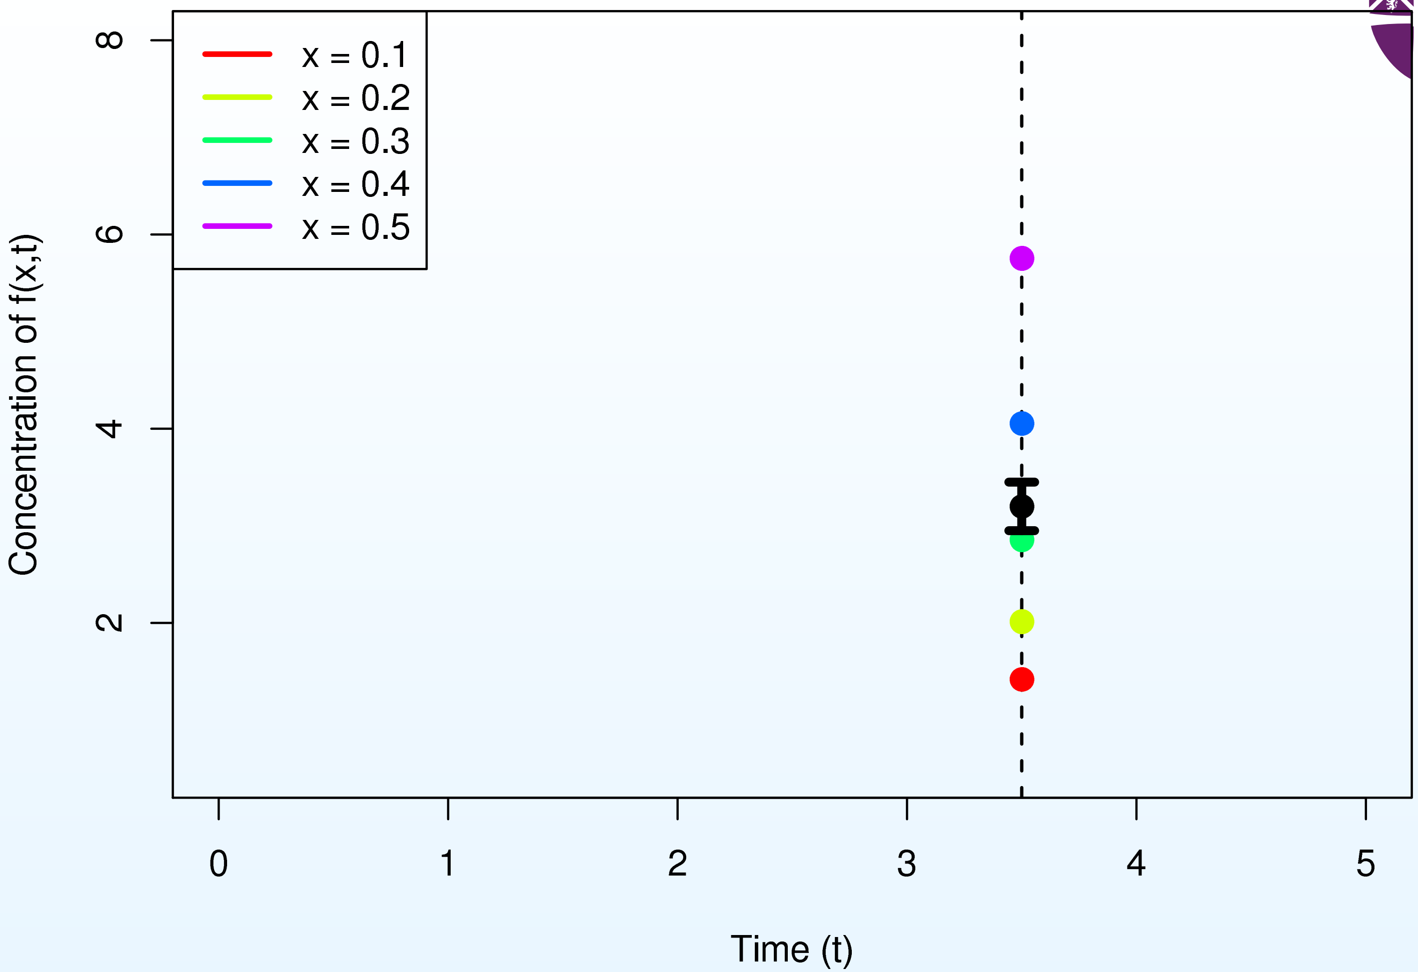

- To answer this, we can now discard other values of $f(x,t)$ and think of $f(x,t=3.5)$ as a function of x only.
- That is take $f(x) \equiv f(x,t=3.5)$

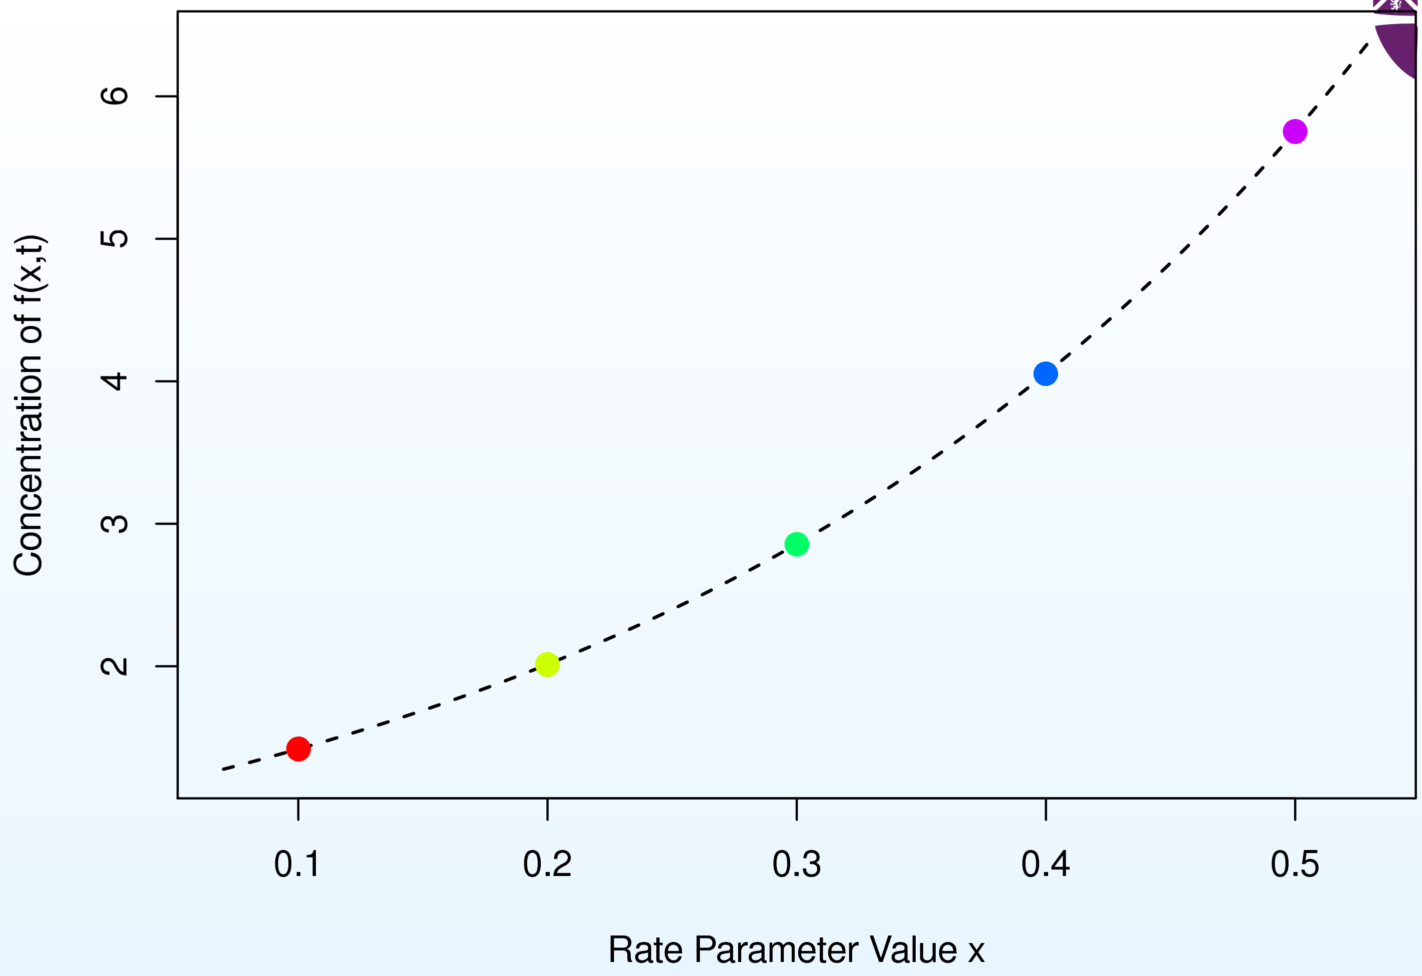

- We can now plot the concentration $f(x)$ as a function of the input parameter x.

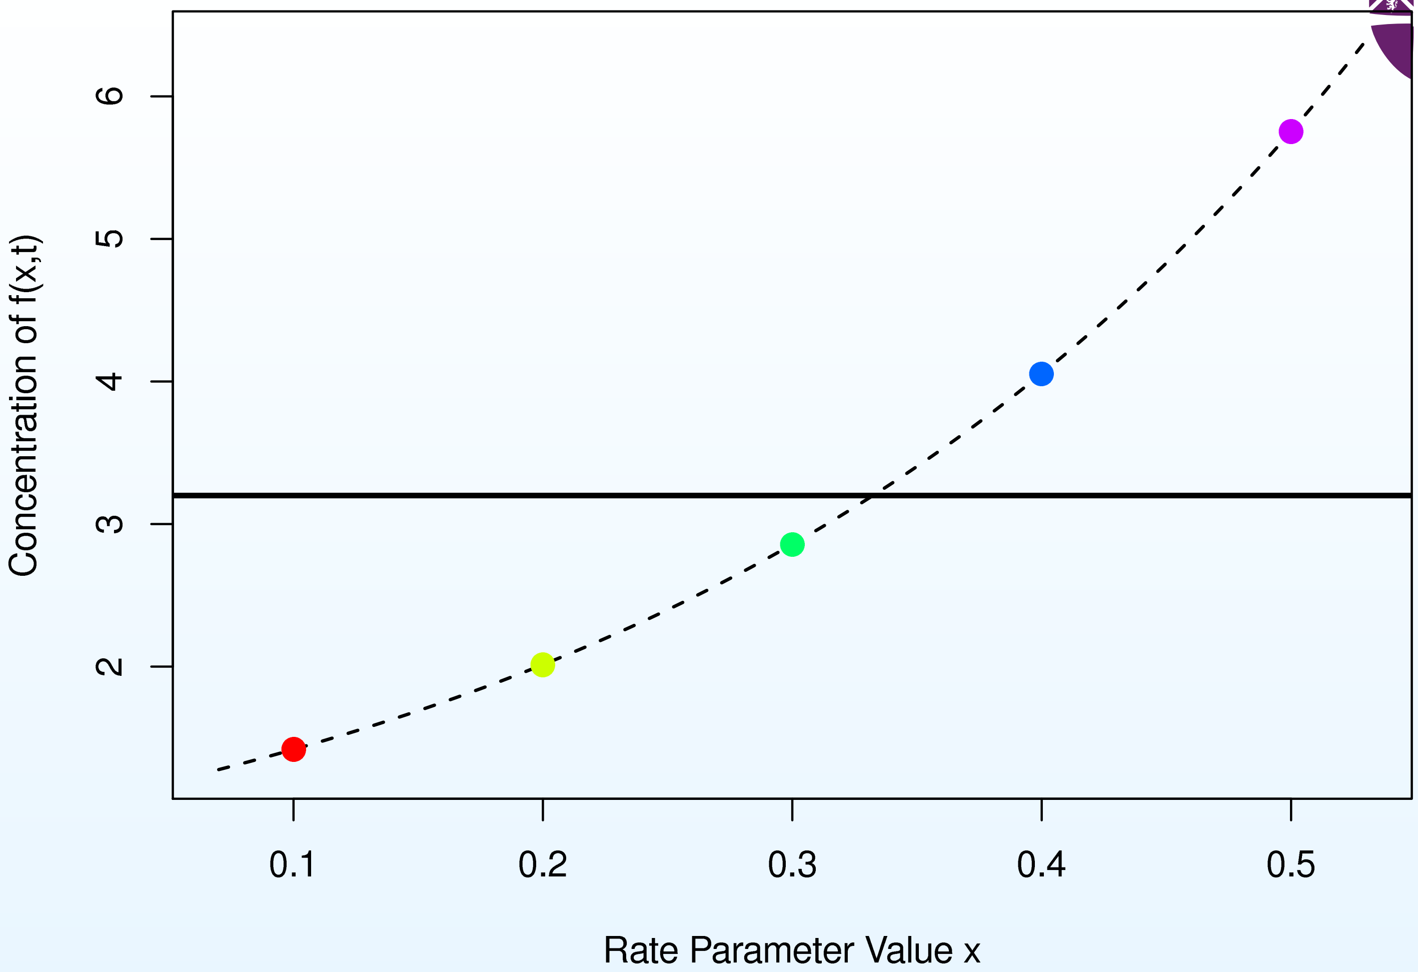

- We can now plot the concentration $f(x)$ as a function of the input parameter x.
- Black horizontal line: the observed measurement of f

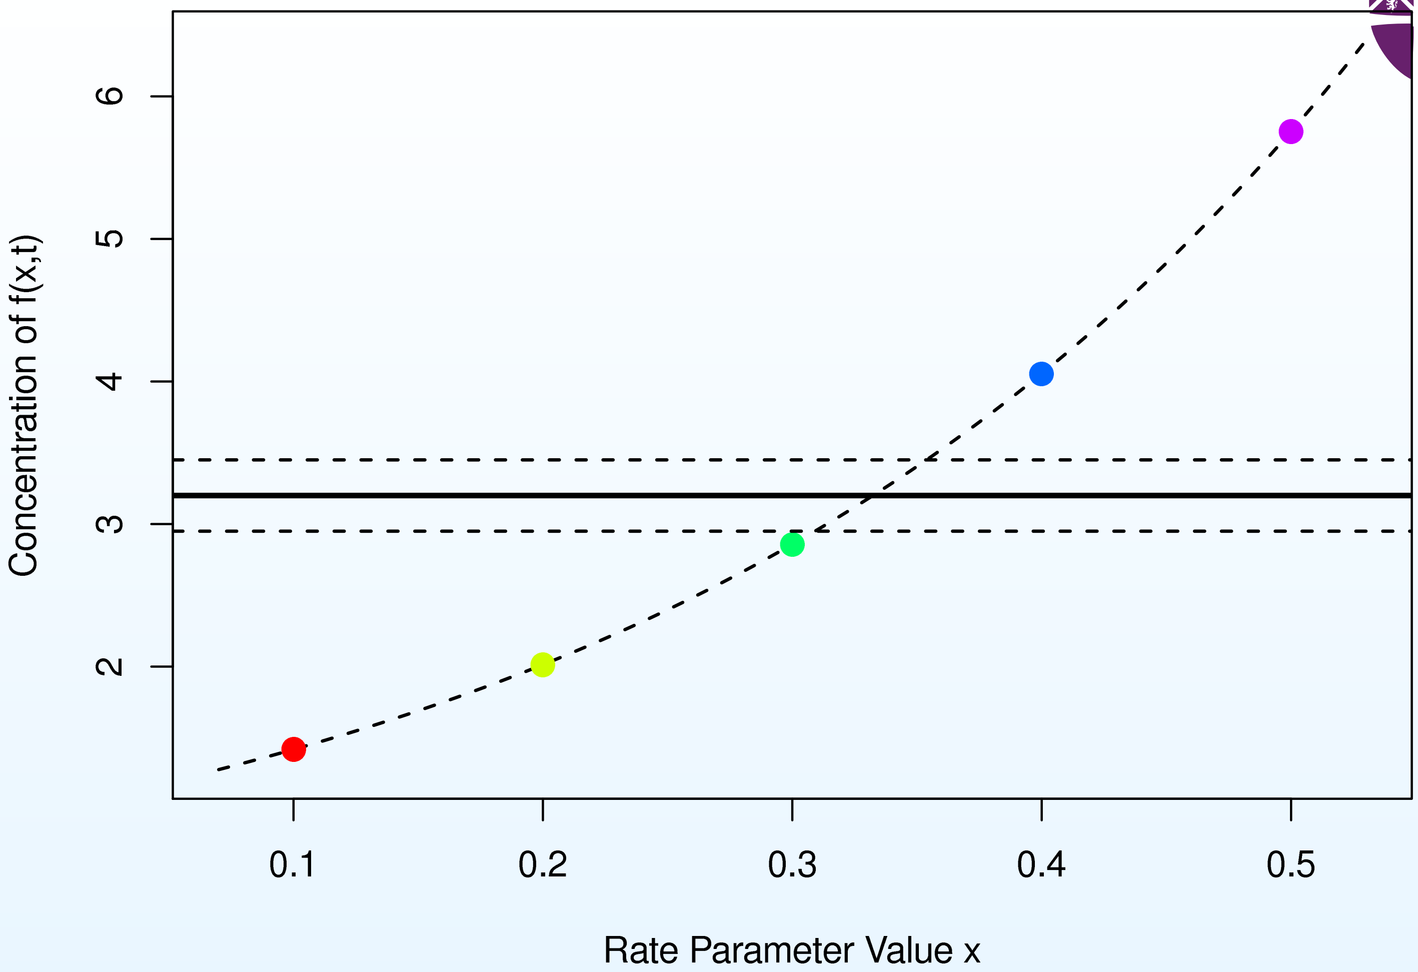

- We can now plot the concentration $f(x)$ as a function of the input parameter x.
- Black horizontal line: the observed measurement of f
- Dashed horizontal lines: the measurement errors

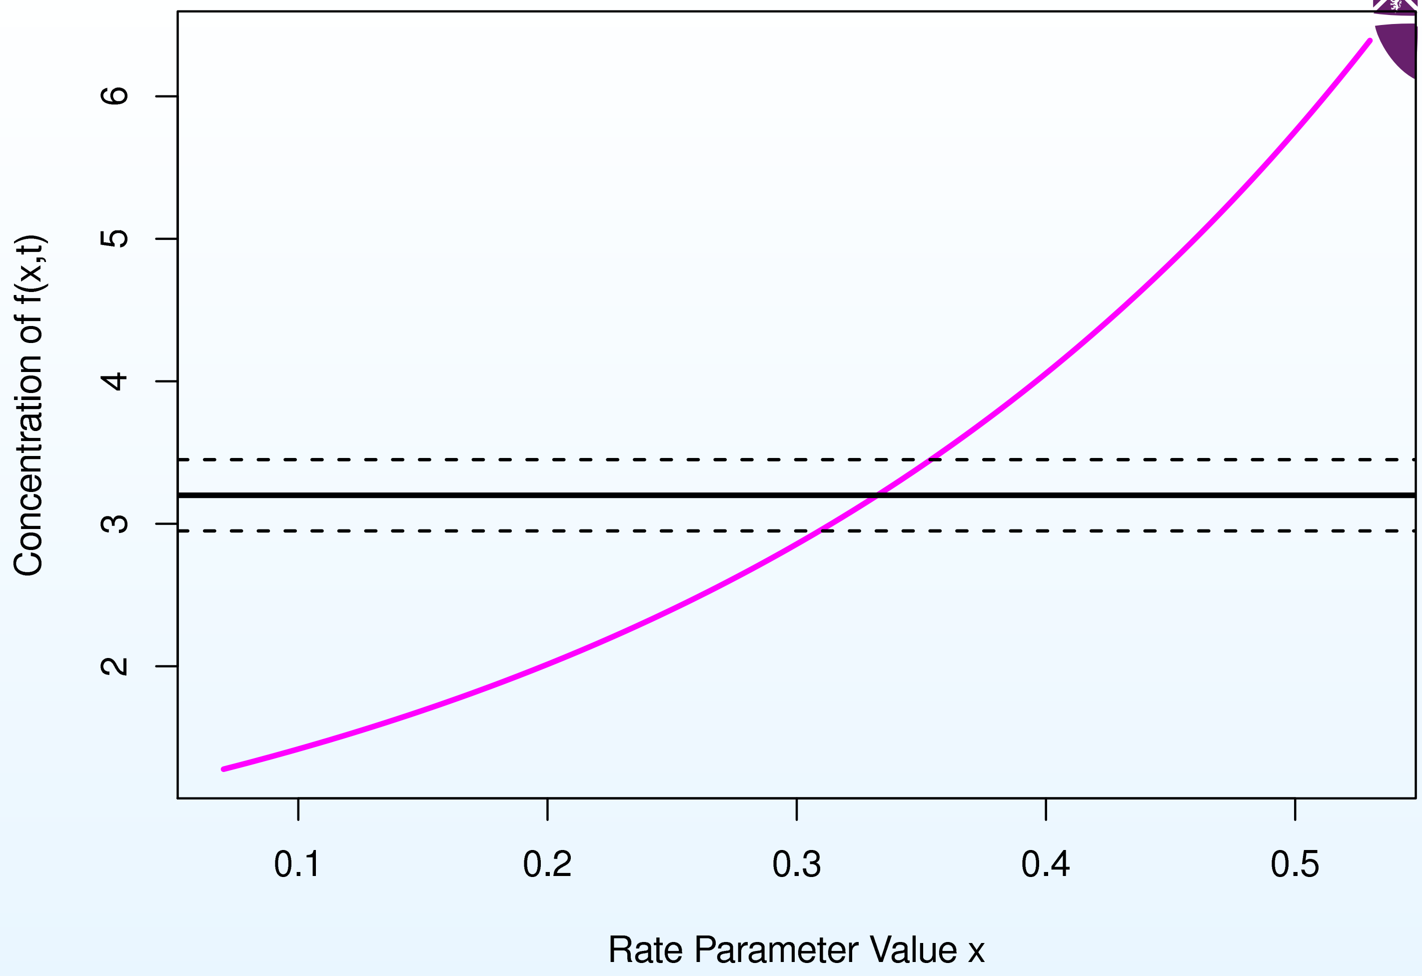

- If we know the analytical expression for $f(x) = \exp(3.5x)$, then we can identify the values of x of interest.

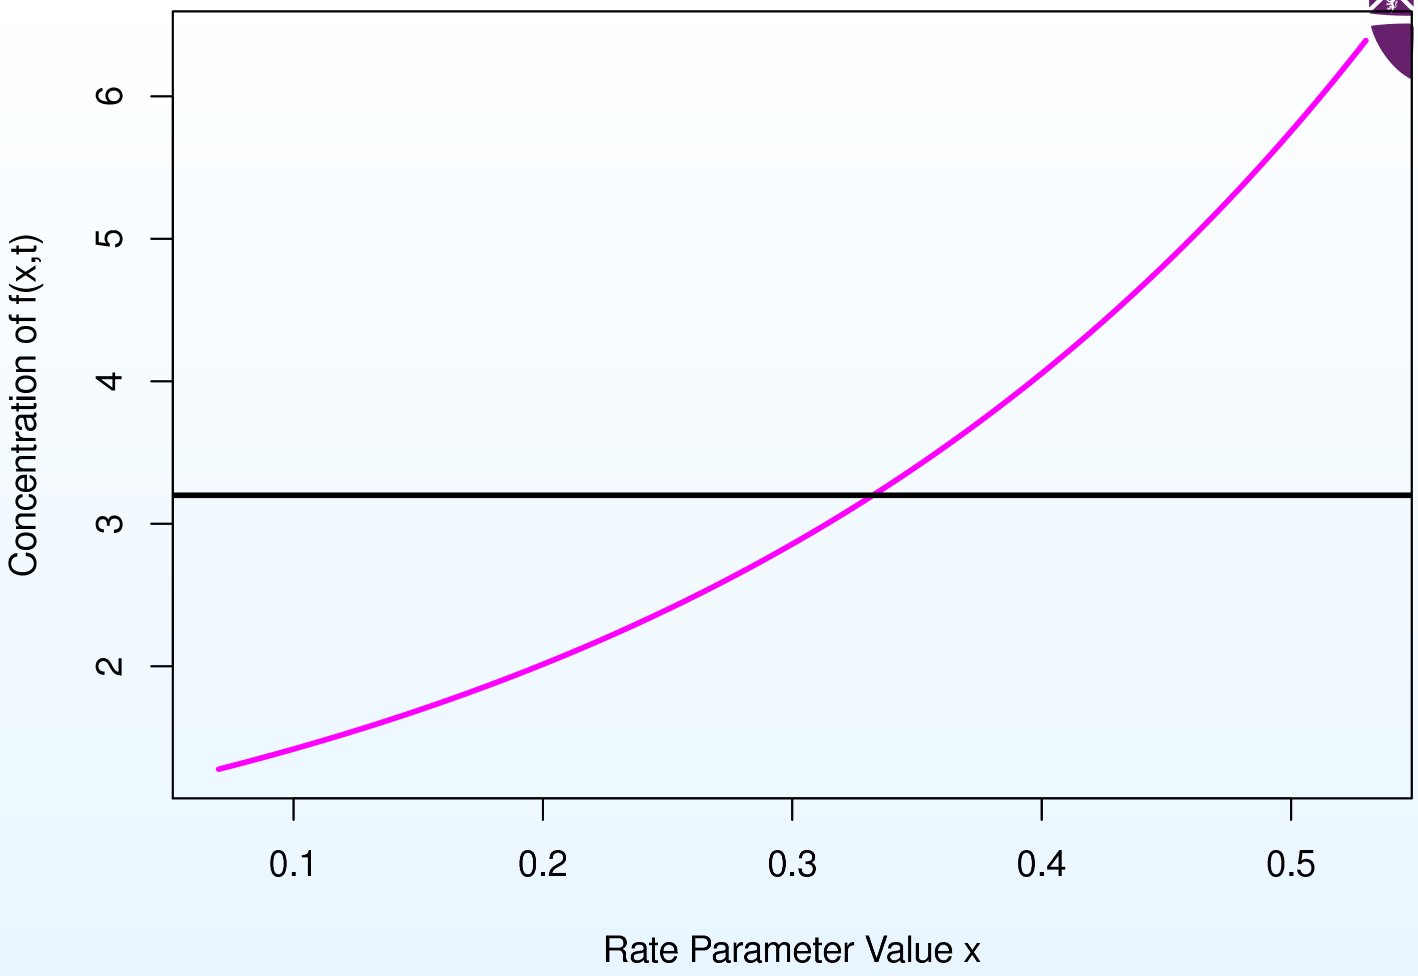

- If we know the analytical expression for $f(x) = \exp(3.5x)$, then we can identify the values of x of interest.
- Ignoring the measurement error would lead to a single value for x but this is incorrect: we have to include the errors.

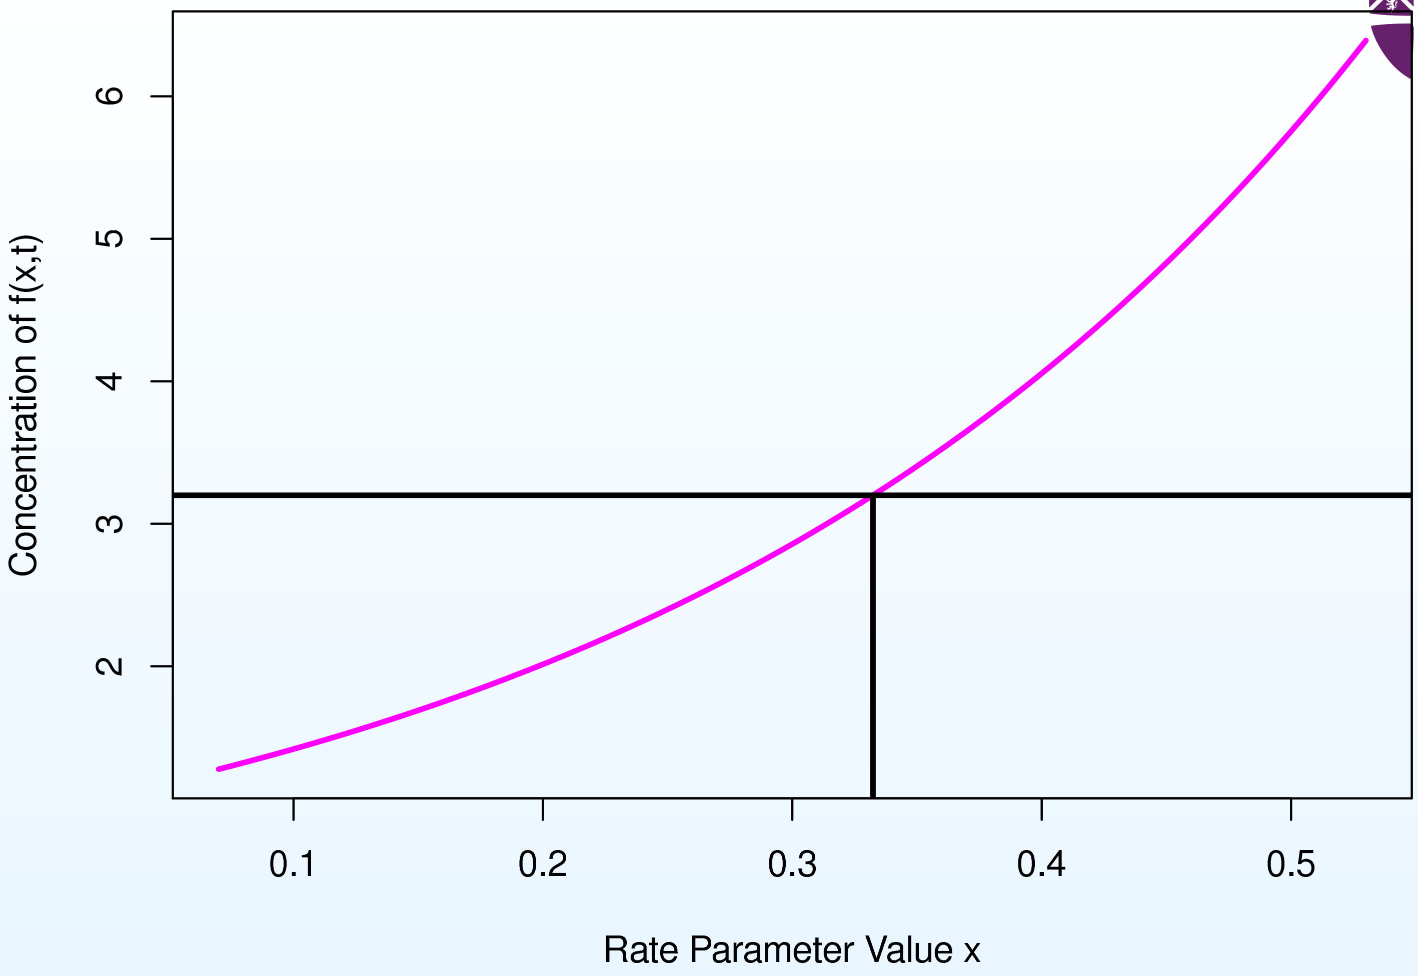

- If we know the analytical expression for $f(x) = \exp(3.5x)$, then we can identify the values of x of interest.
- Ignoring the measurement error would lead to a single value for x but this is incorrect: we have to include the errors.

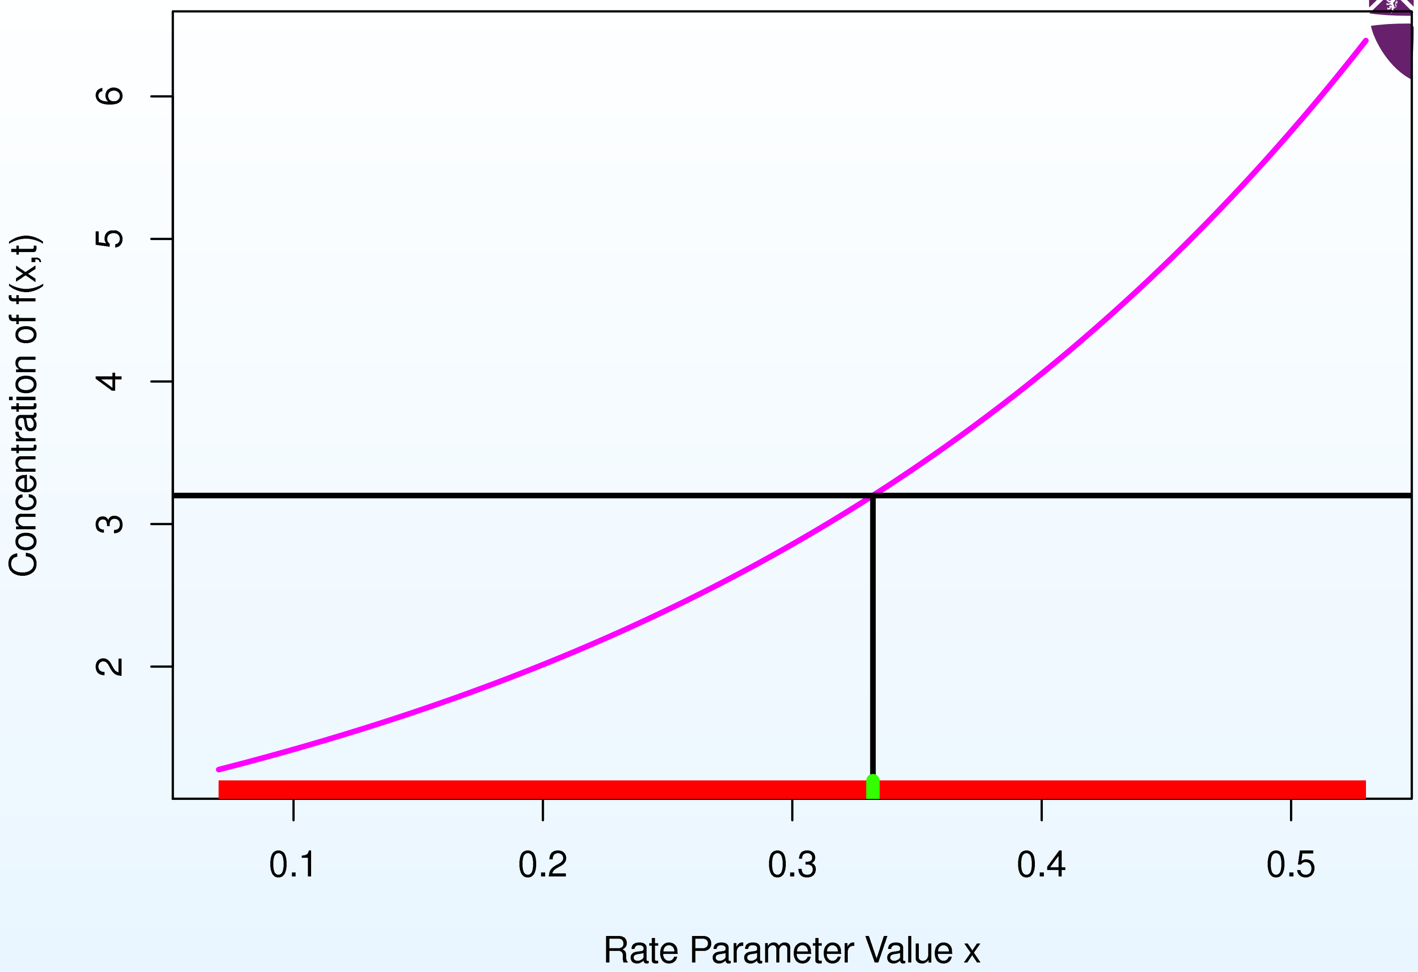

- If we know the analytical expression for $f(x) = \exp(3.5x)$, then we can identify the values of x of interest.
- Ignoring the measurement error would lead to a single value for x but this is incorrect: we have to include the errors.

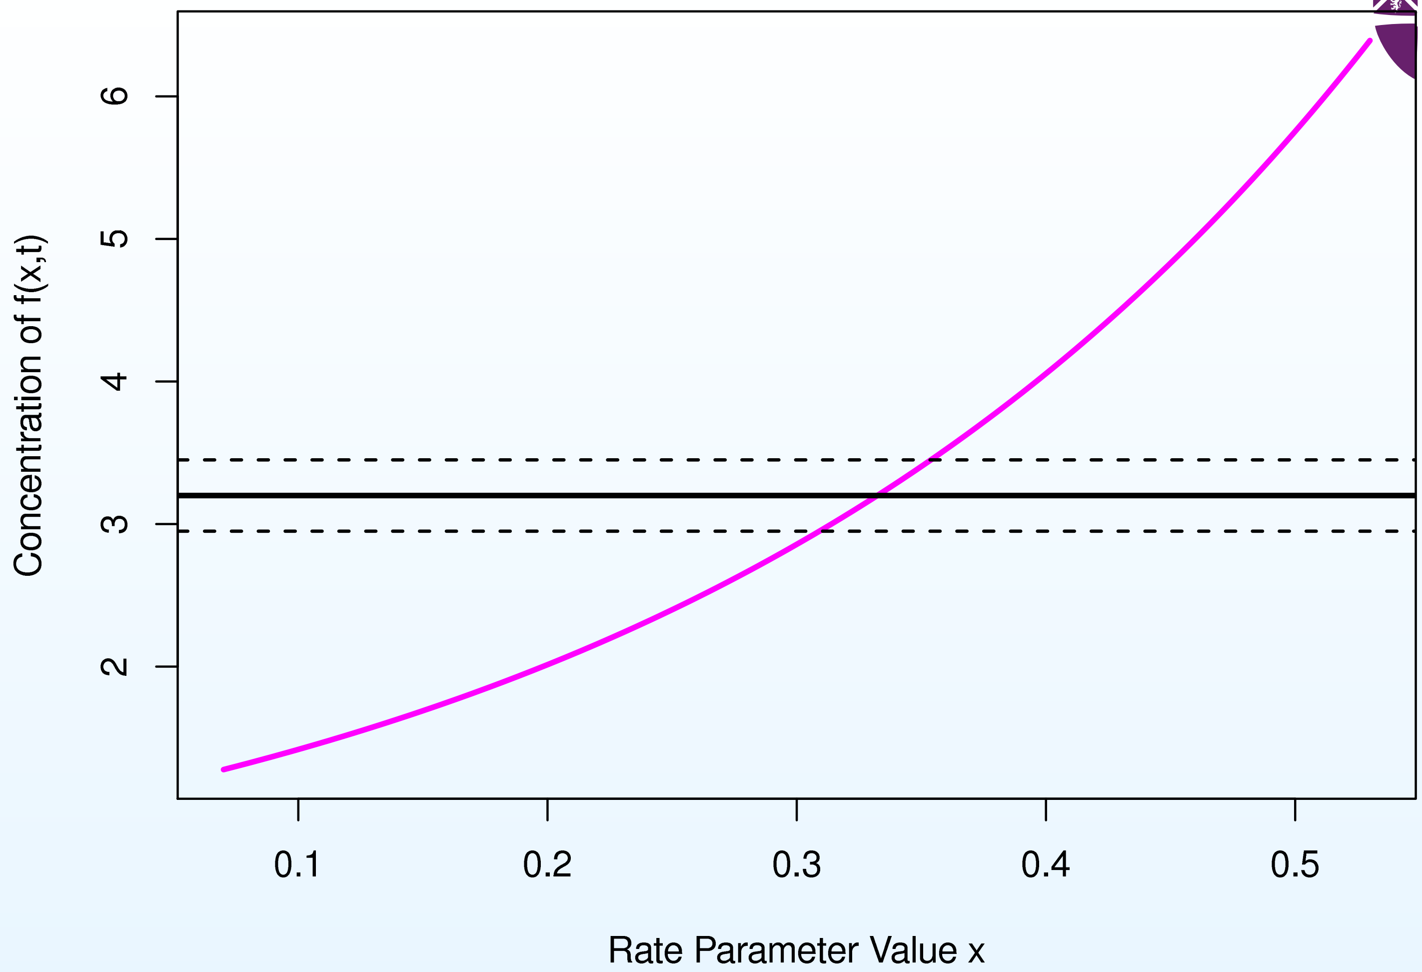

- Uncertainty in the measurement of $f(x, t = 3.5)$ leads to uncertainty in the inferred values of x.

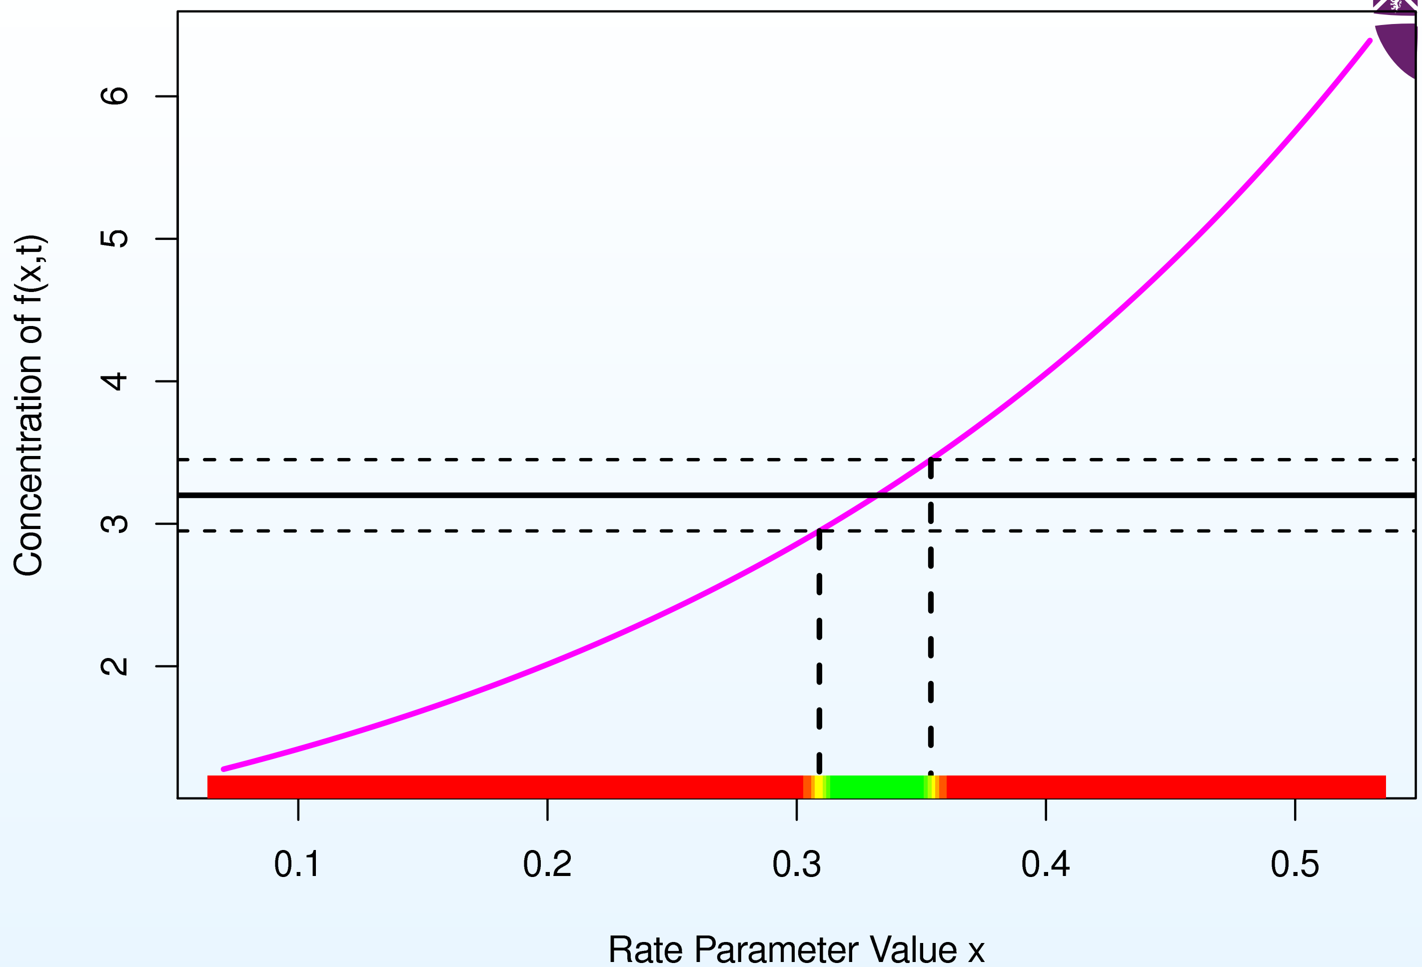

- Uncertainty in the measurement of $f(x, t = 3.5)$ leads to uncertainty in the inferred values of x.
- Hence we see a range (green/yellow) of possible values of x consistent with the measurements, with all the implausible values of x in red.

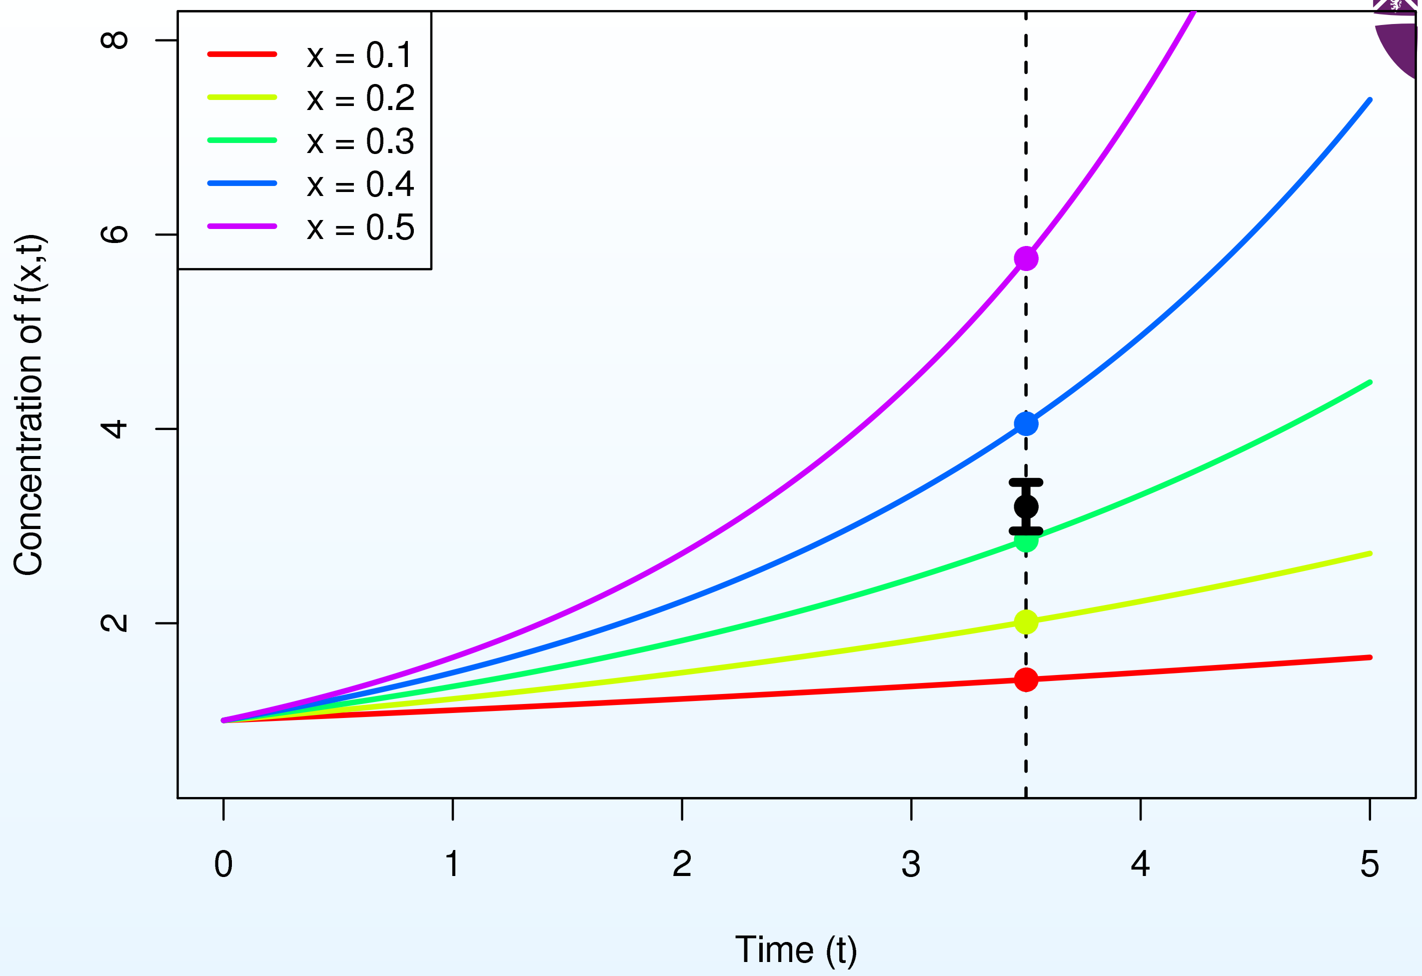

- Constraints on x from observations impose constraints on $f(x,t)$ in the future.

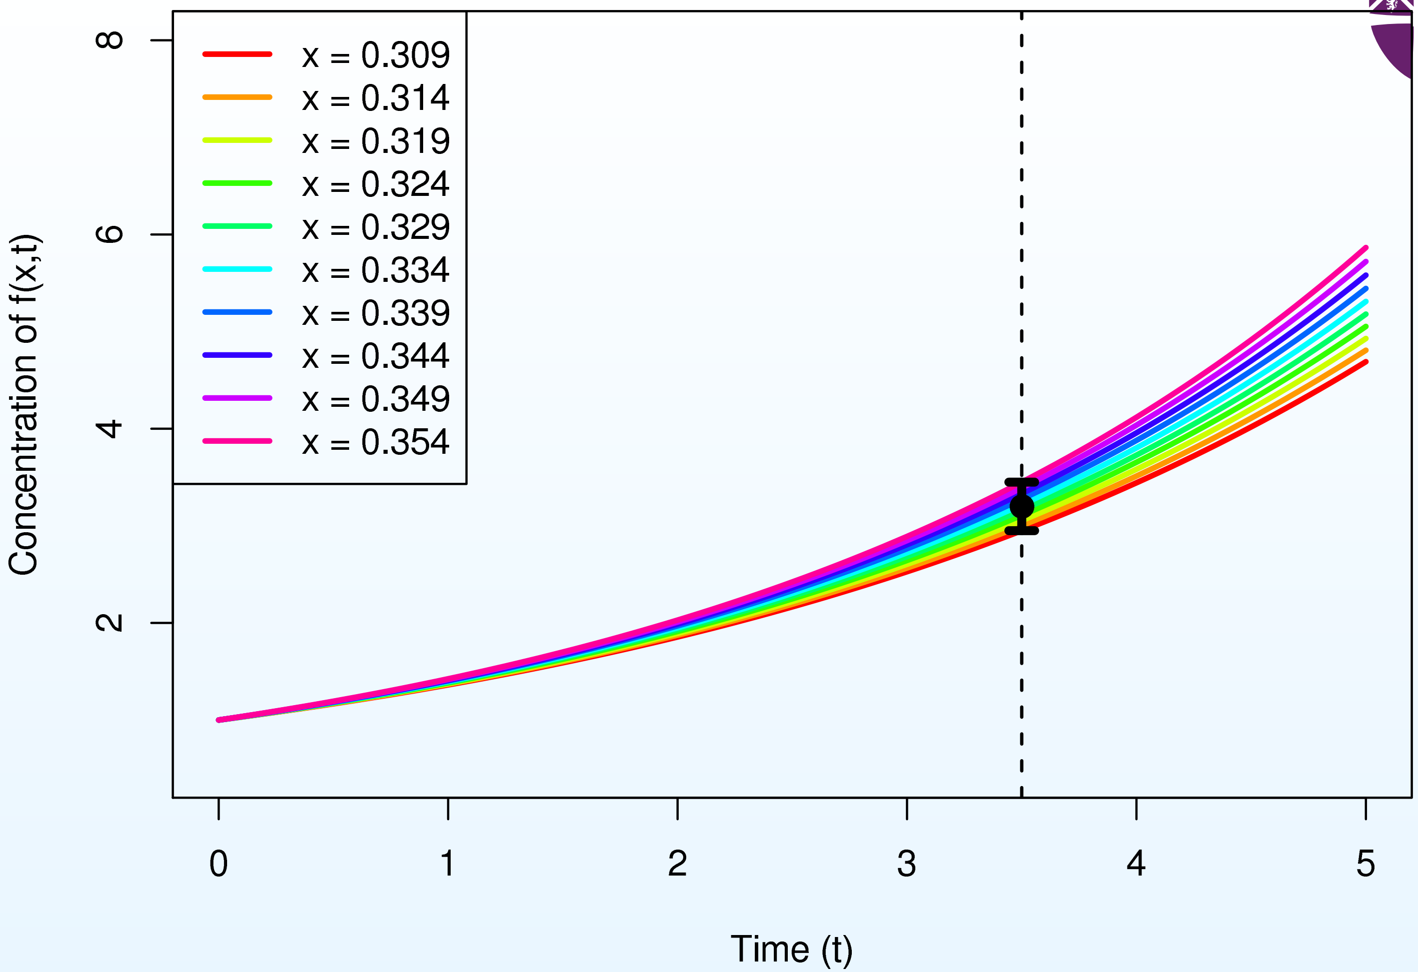

- We choose values of x consistent with the measurement of $f(x,t)$ at $t = 3.5$, and perform corresponding runs of the model.

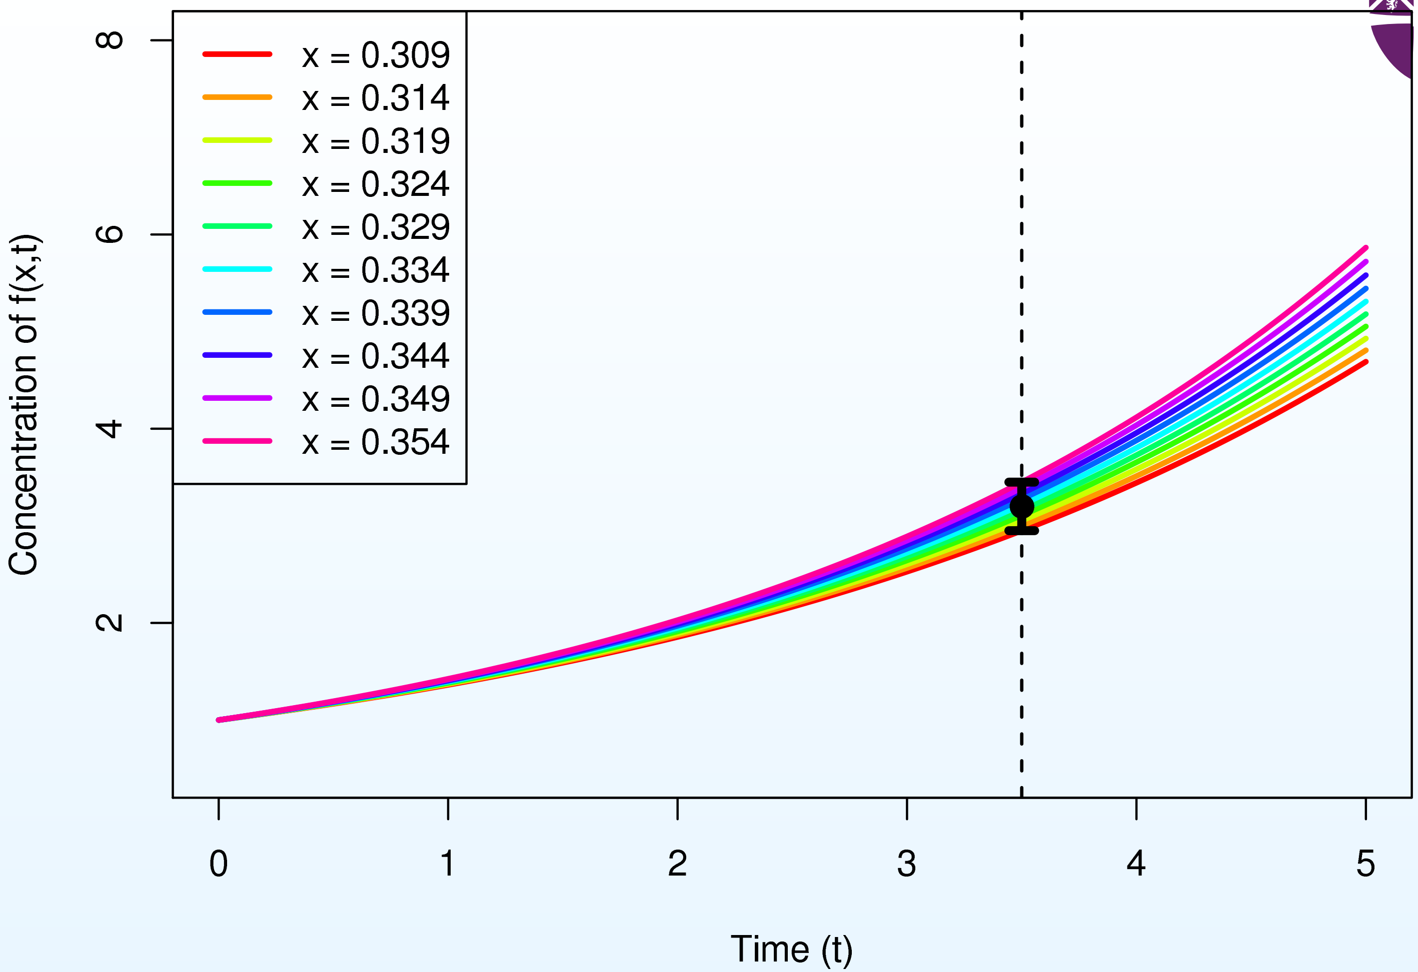

- This shows us the range of possible predictions of future system behaviour consistent with the model structure and the past measurements.

Function emulation

Uncertainty analysis, for high dimensional problems, is particularly challenging if $f(x)$ is expensive, in time and computational resources, to evaluate for any choice of x.
[For example, large energy models.]

Function emulation

Uncertainty analysis, for high dimensional problems, is particularly challenging if $f(x)$ is expensive, in time and computational resources, to evaluate for any choice of x.
[For example, large energy models.]

In such cases, f must be treated as uncertain for all input choices except the small subset for which an actual evaluation has been made.

Function emulation

Uncertainty analysis, for high dimensional problems, is particularly challenging if $f(x)$ is expensive, in time and computational resources, to evaluate for any choice of x.
[For example, large energy models.]

In such cases, f must be treated as uncertain for all input choices except the small subset for which an actual evaluation has been made.

Therefore, we must construct a description of the uncertainty about the value of $f(x)$ for each x.

Function emulation

Uncertainty analysis, for high dimensional problems, is particularly challenging if $f(x)$ is expensive, in time and computational resources, to evaluate for any choice of x.
[For example, large energy models.]

In such cases, f must be treated as uncertain for all input choices except the small subset for which an actual evaluation has been made.

Therefore, we must construct a description of the uncertainty about the value of $f(x)$ for each x.

Such a representation is often termed an emulator of the simulator.

Function emulation

Uncertainty analysis, for high dimensional problems, is particularly challenging if $f(x)$ is expensive, in time and computational resources, to evaluate for any choice of x.
[For example, large energy models.]

In such cases, f must be treated as uncertain for all input choices except the small subset for which an actual evaluation has been made.

Therefore, we must construct a description of the uncertainty about the value of $f(x)$ for each x.

Such a representation is often termed an emulator of the simulator.

The emulator both contains
(i) an approximation to the simulator and

Function emulation

Uncertainty analysis, for high dimensional problems, is particularly challenging if $f(x)$ is expensive, in time and computational resources, to evaluate for any choice of x.
[For example, large energy models.]

In such cases, f must be treated as uncertain for all input choices except the small subset for which an actual evaluation has been made.

Therefore, we must construct a description of the uncertainty about the value of $f(x)$ for each x.

Such a representation is often termed an emulator of the simulator.

The emulator both contains
(i) an approximation to the simulator and

(ii) an assessment of the likely magnitude of the error of the approximation.

Function emulation

Uncertainty analysis, for high dimensional problems, is particularly challenging if $f(x)$ is expensive, in time and computational resources, to evaluate for any choice of x.
[For example, large energy models.]

In such cases, f must be treated as uncertain for all input choices except the small subset for which an actual evaluation has been made.

Therefore, we must construct a description of the uncertainty about the value of $f(x)$ for each x.

Such a representation is often termed an emulator of the simulator.

The emulator both contains
(i) an approximation to the simulator and

(ii) an assessment of the likely magnitude of the error of the approximation.

Unlike the original simulator, the emulator is fast to evaluate for any choice of inputs.

Form of the emulator

We may represent beliefs about component f_i of f, using an emulator:

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

Form of the emulator

We may represent beliefs about component f_i of f, using an emulator:

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

where $B = \{\beta_{ij}\}$ are unknown scalars, g_{ij} are known deterministic functions of x, (for example, polynomials)

Form of the emulator

We may represent beliefs about component f_i of f, using an emulator:

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

where $B = \{\beta_{ij}\}$ are unknown scalars, g_{ij} are known deterministic functions of x, (for example, polynomials)
$Bg(x)$ expresses global variation in f.

Form of the emulator

We may represent beliefs about component f_i of f, using an emulator:

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

where $B = \{\beta_{ij}\}$ are unknown scalars, g_{ij} are known deterministic functions of x, (for example, polynomials)
$Bg(x)$ expresses global variation in f.

$u_i(x)$ is a weakly second order stationary stochastic process, with (for example) correlation function

$$\mathrm{Corr}(u_i(x), u_i(x')) = \exp(-(\tfrac{\|x-x'\|}{\theta_i})^2)$$

Form of the emulator

We may represent beliefs about component f_i of f, using an emulator:

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

where $B = \{\beta_{ij}\}$ are unknown scalars, g_{ij} are known deterministic functions of x, (for example, polynomials)

$Bg(x)$ expresses global variation in f.

$u_i(x)$ is a weakly second order stationary stochastic process, with (for example) correlation function

$$\mathrm{Corr}(u_i(x), u_i(x')) = \exp(-(\tfrac{\|x-x'\|}{\theta_i})^2)$$

$u(x)$ expresses local variation in f

Emulating the Model: Simple 1D Example

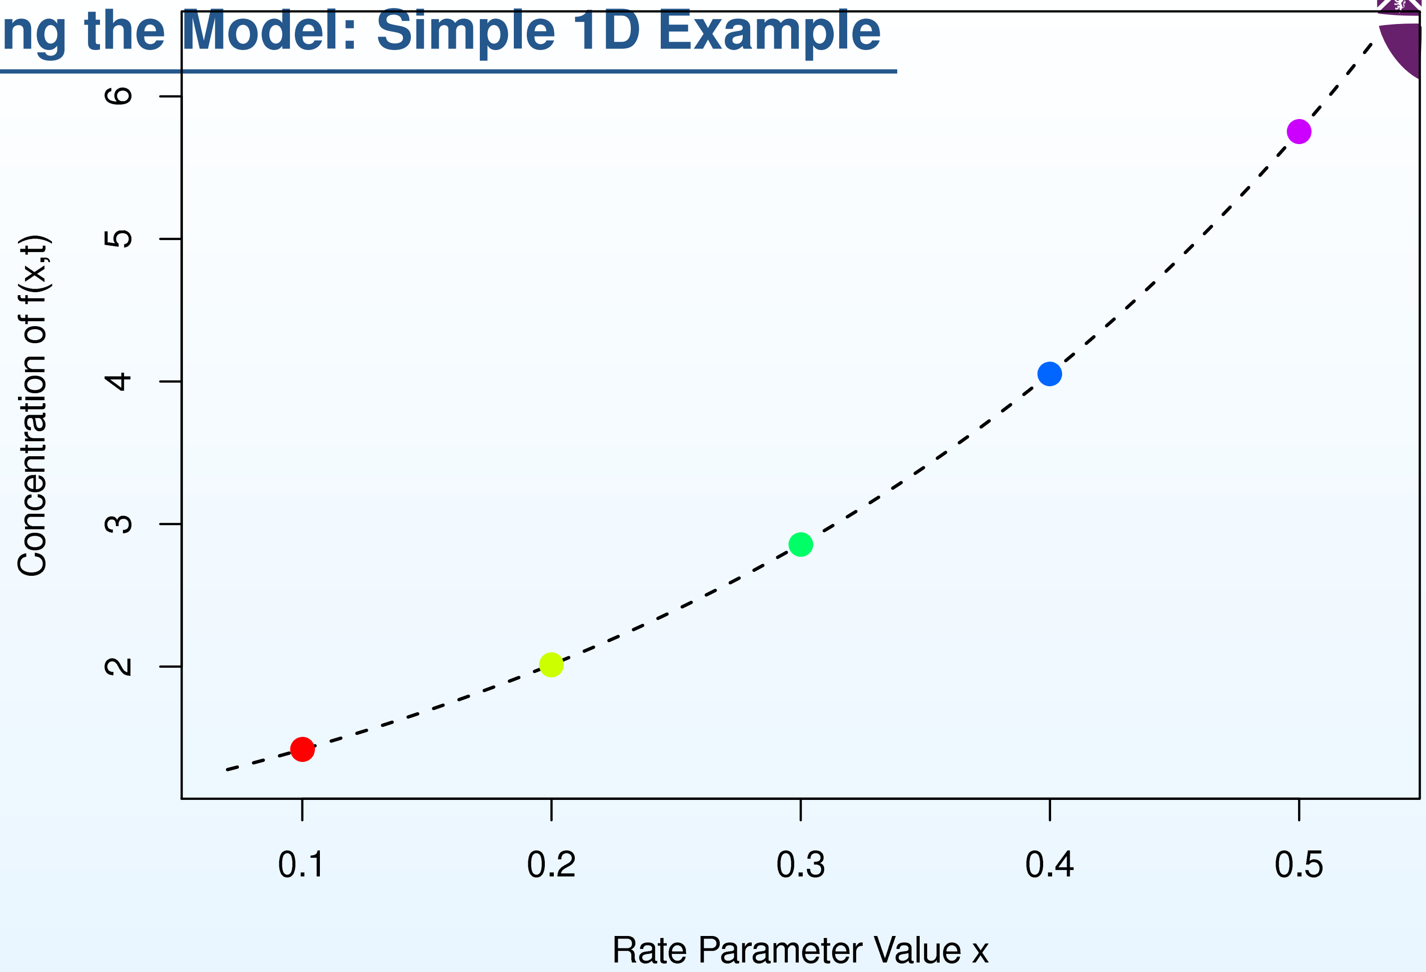

- Consider the graph of $f(x)$: in general we do not have the analytic solution of $f(x)$, here given by the dashed line.

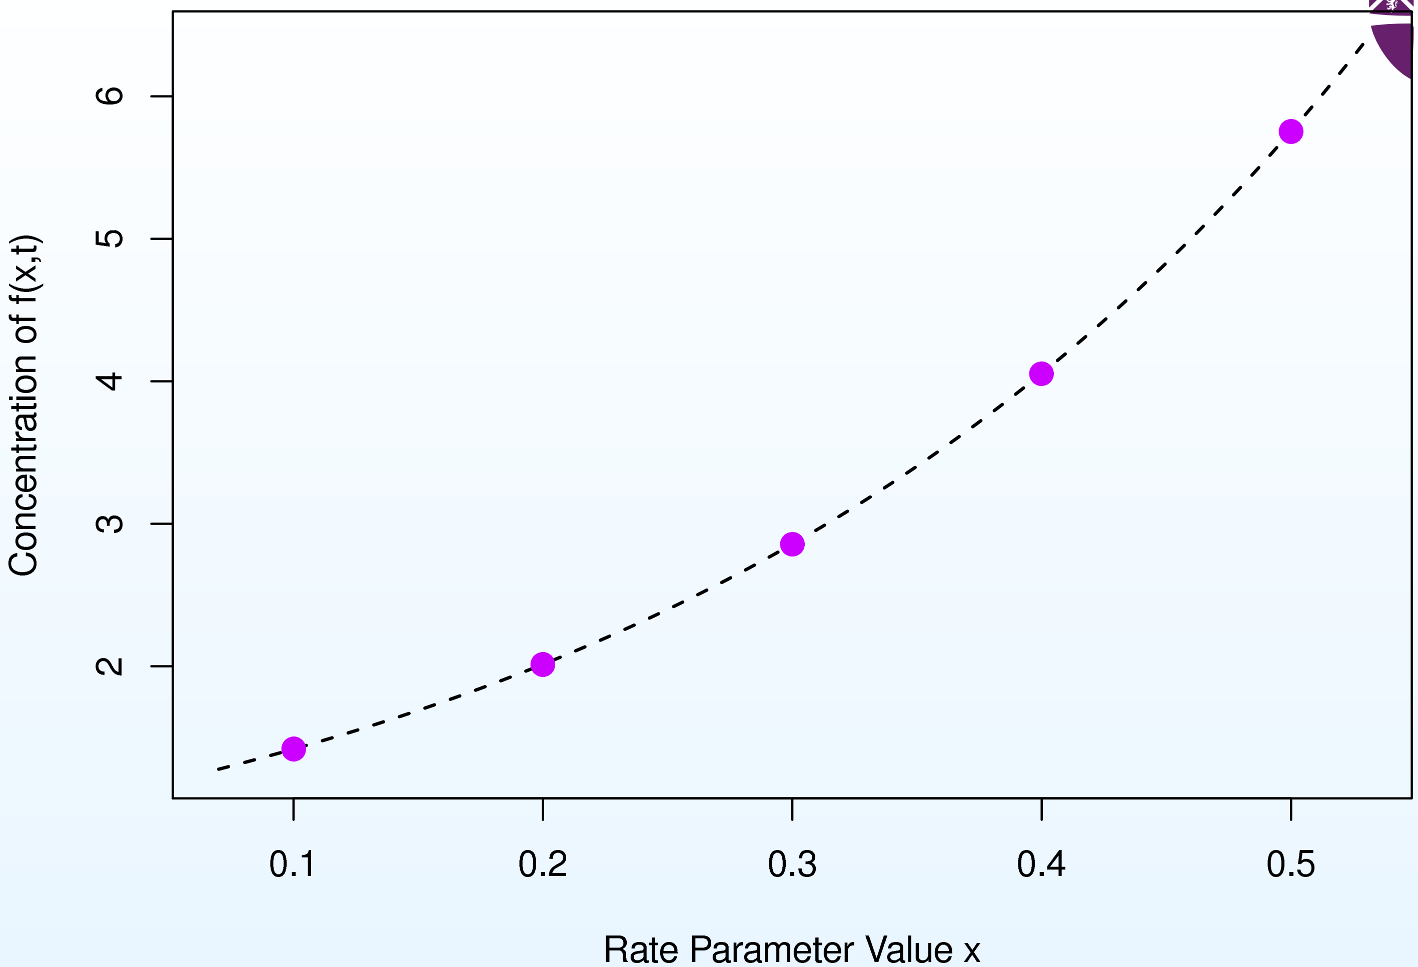

- Consider the graph of $f(x)$: in general we do not have the analytic solution of $f(x)$, here given by the dashed line.

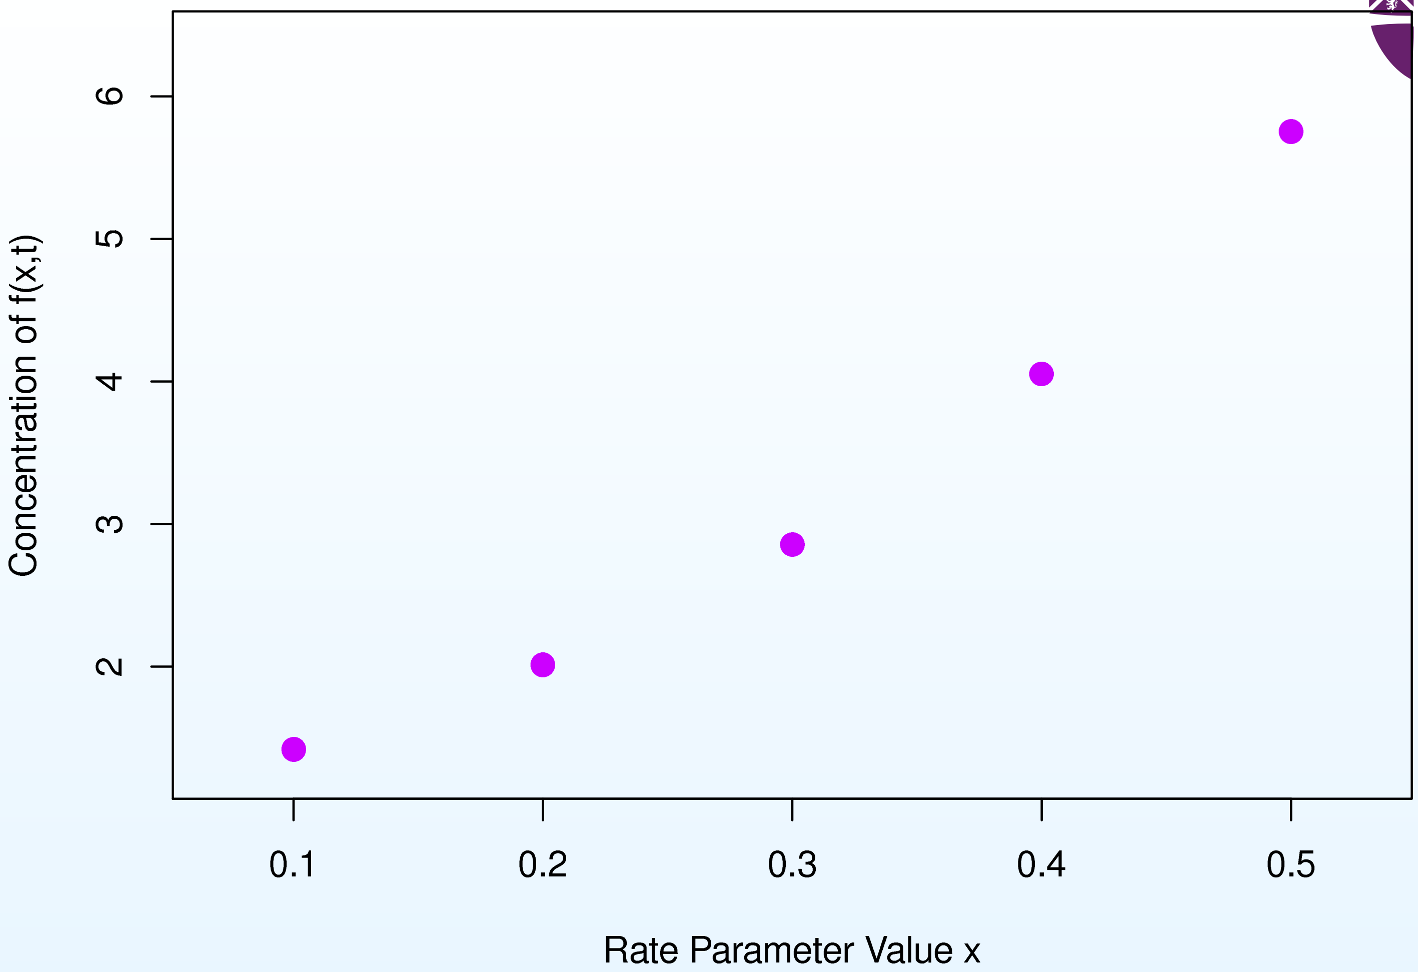

- Consider the graph of $f(x)$: in general we do not have the analytic solution of $f(x)$, here given by the dashed line.

- Instead we only have a finite number of runs of the model, in this case five.

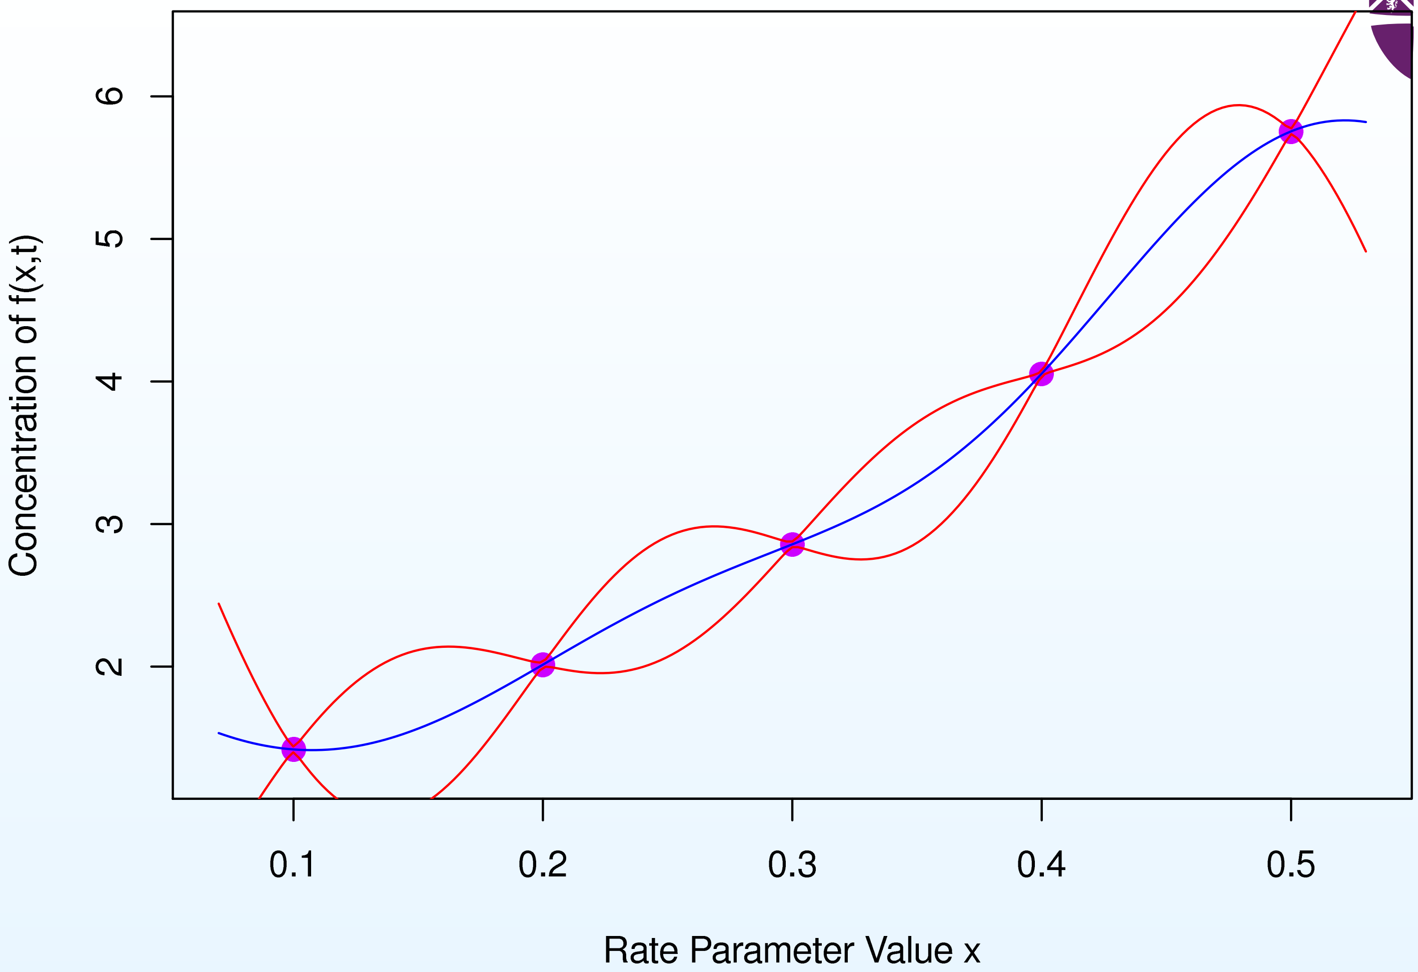

- The emulator can be used to represent our beliefs about the behaviour of the model at untested values of x, and is fast to evaluate.

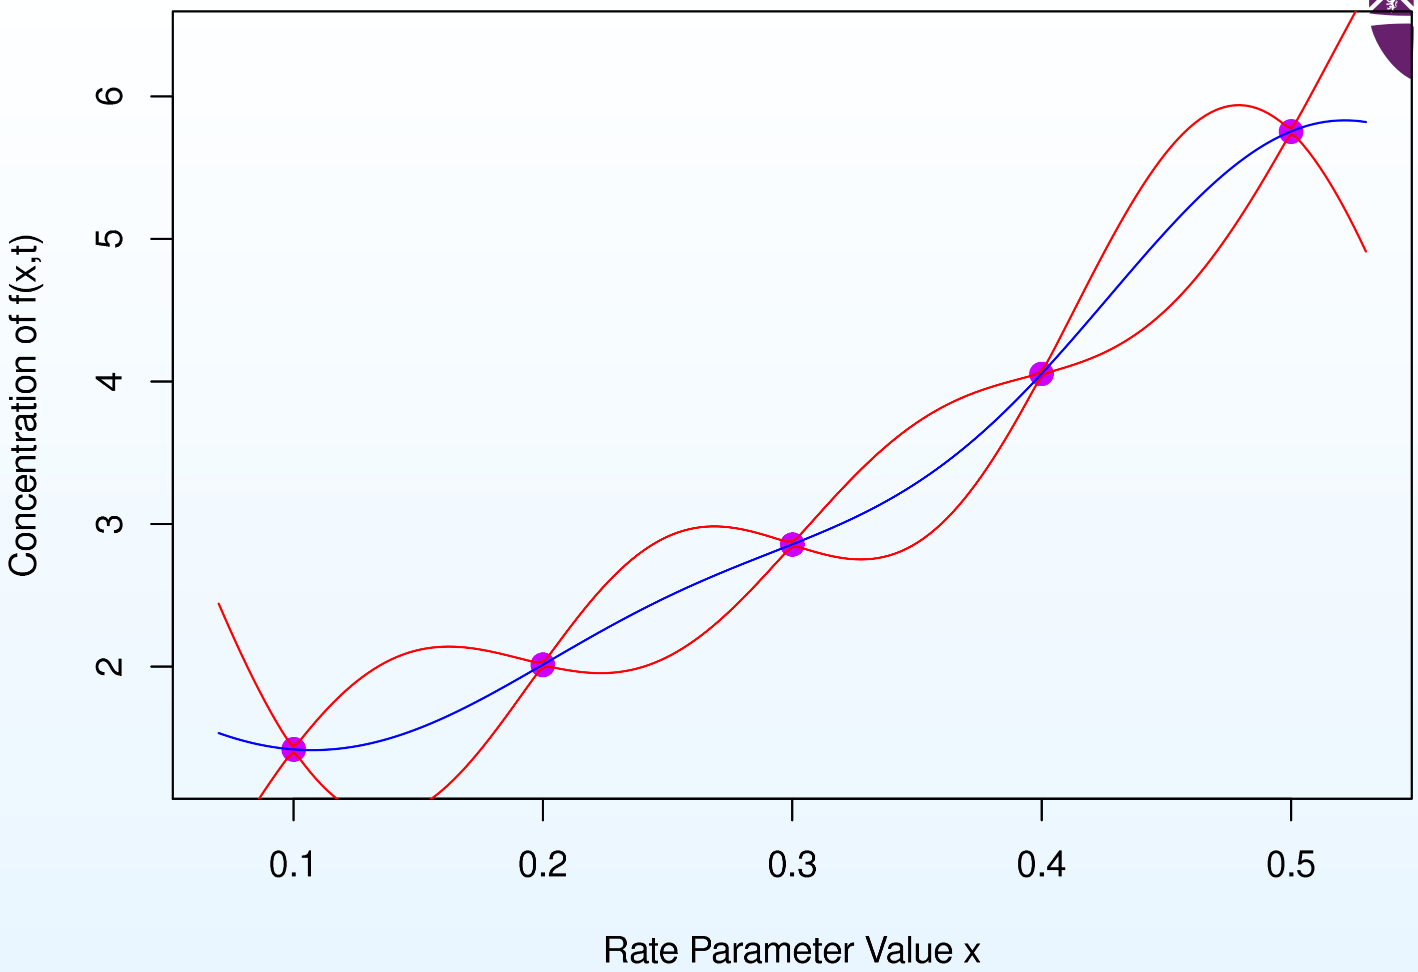

- The emulator can be used to represent our beliefs about the behaviour of the model at untested values of x, and is fast to evaluate.
- It gives both the expected value of $f(x)$ (the blue line) along with a credible interval for $f(x)$ (the red lines) representing the uncertainty about the model's behaviour.

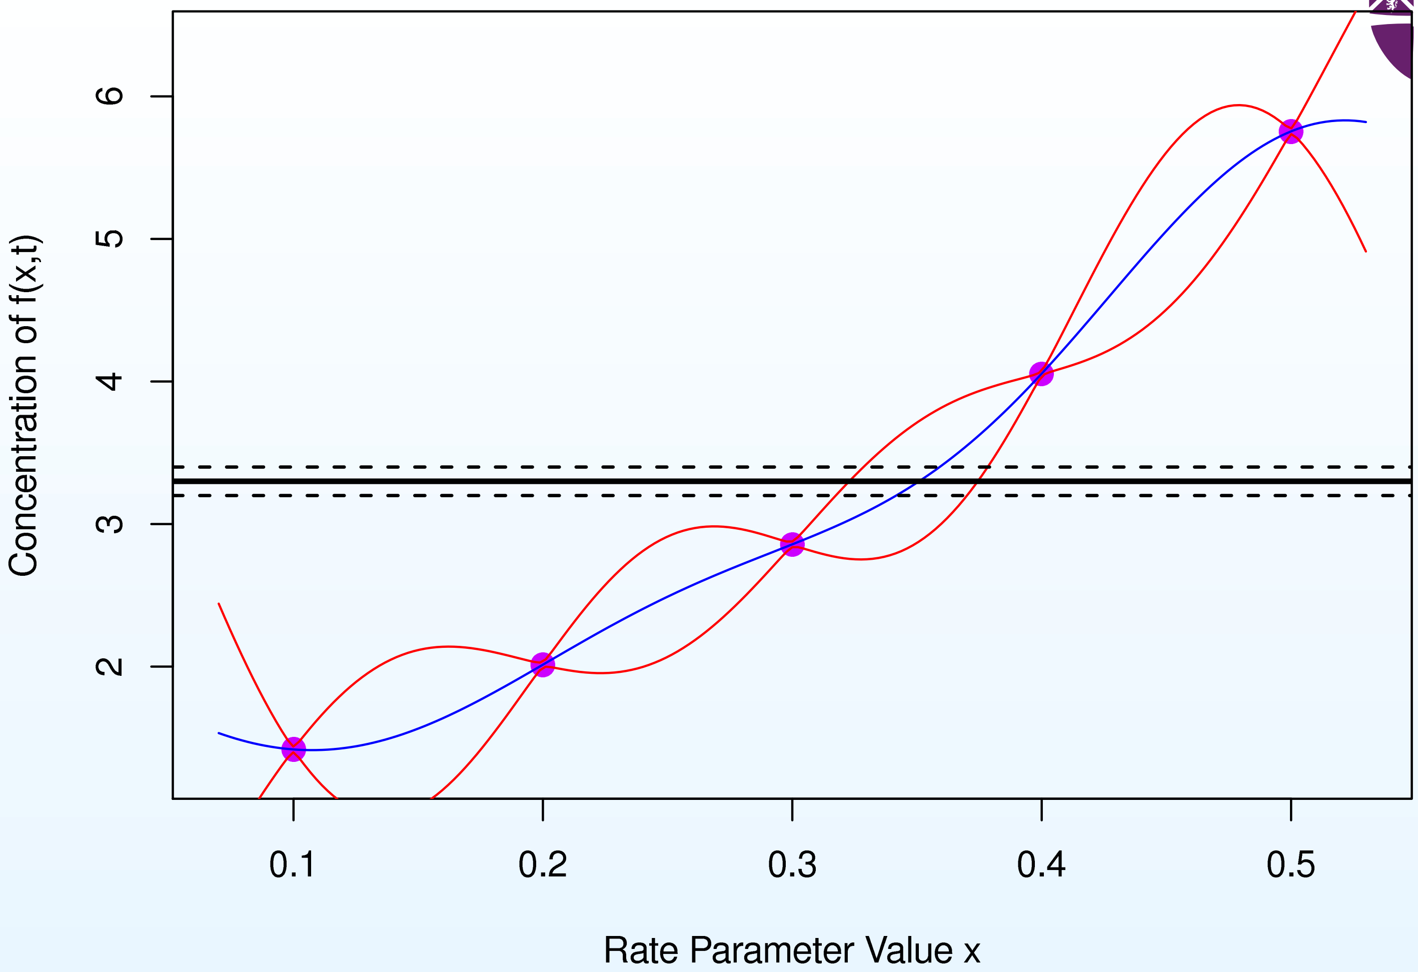

- Comparing the emulator to the observed measurement we again identify the set of x values currently consistent with this data (the observed errors here have been reduced for clarity).

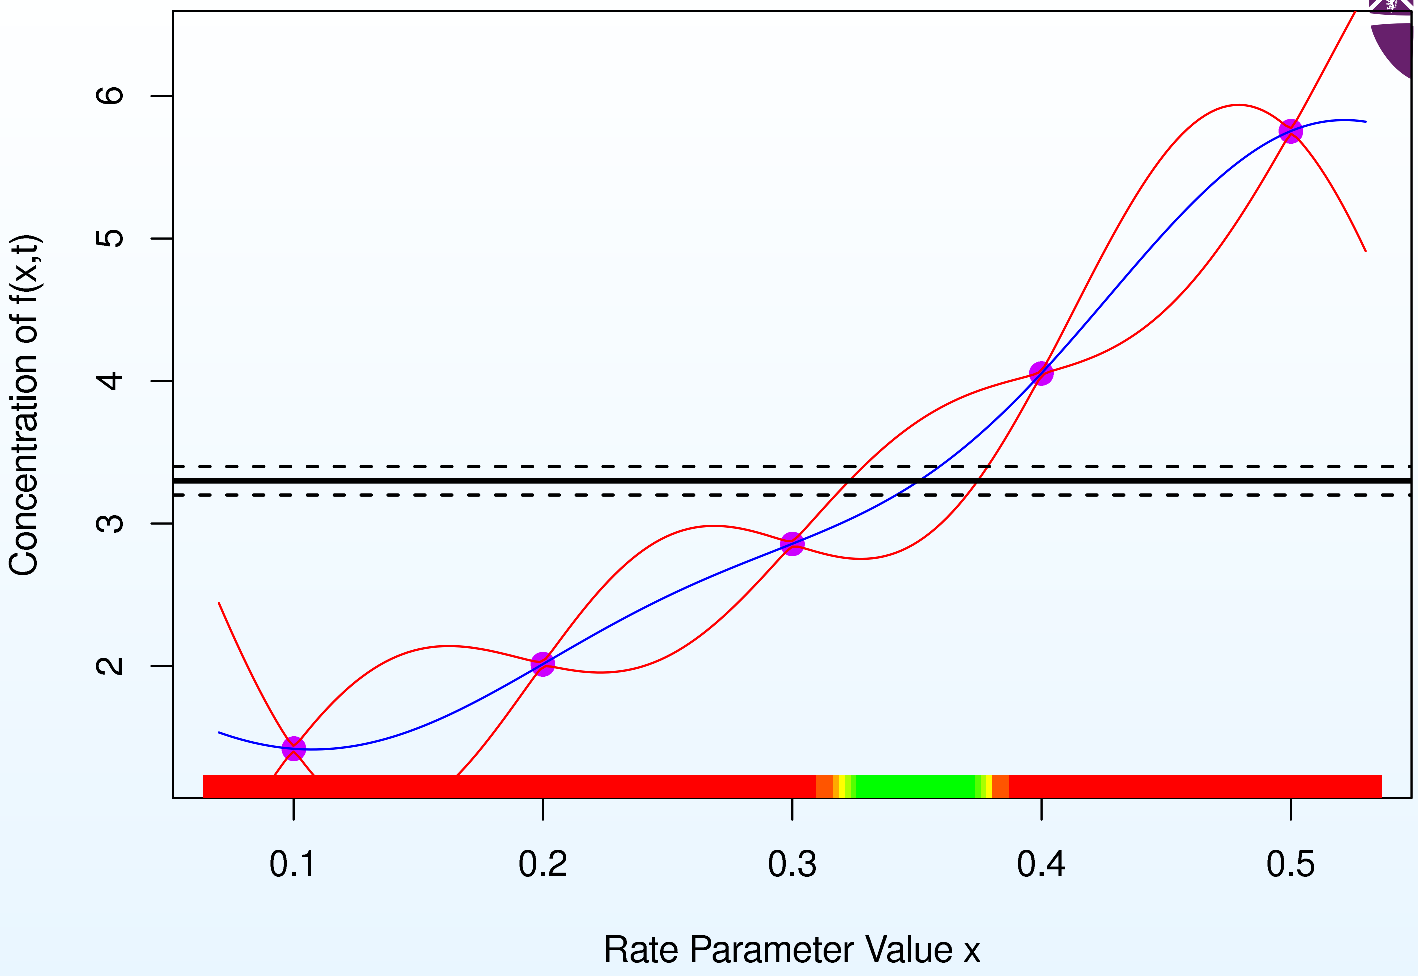

- Comparing the emulator to the observed measurement we again identify the set of x values currently consistent with this data (the observed errors here have been reduced for clarity).
- Note the uncertainty on x now includes uncertainty coming from the emulator.

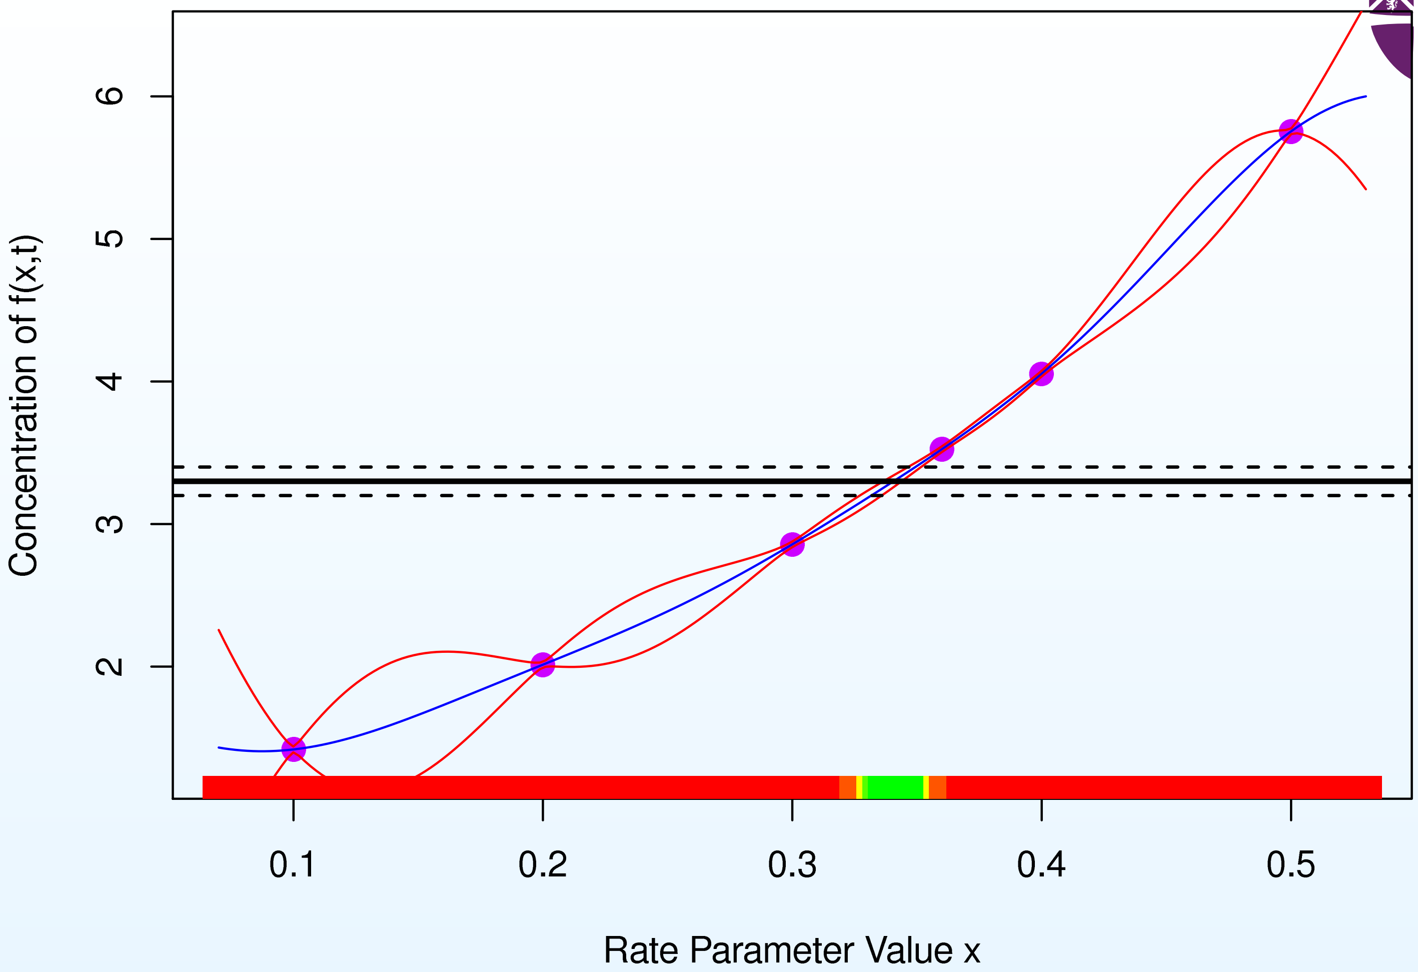

- We perform a 2nd iteration or wave of runs to improve emulator accuracy.

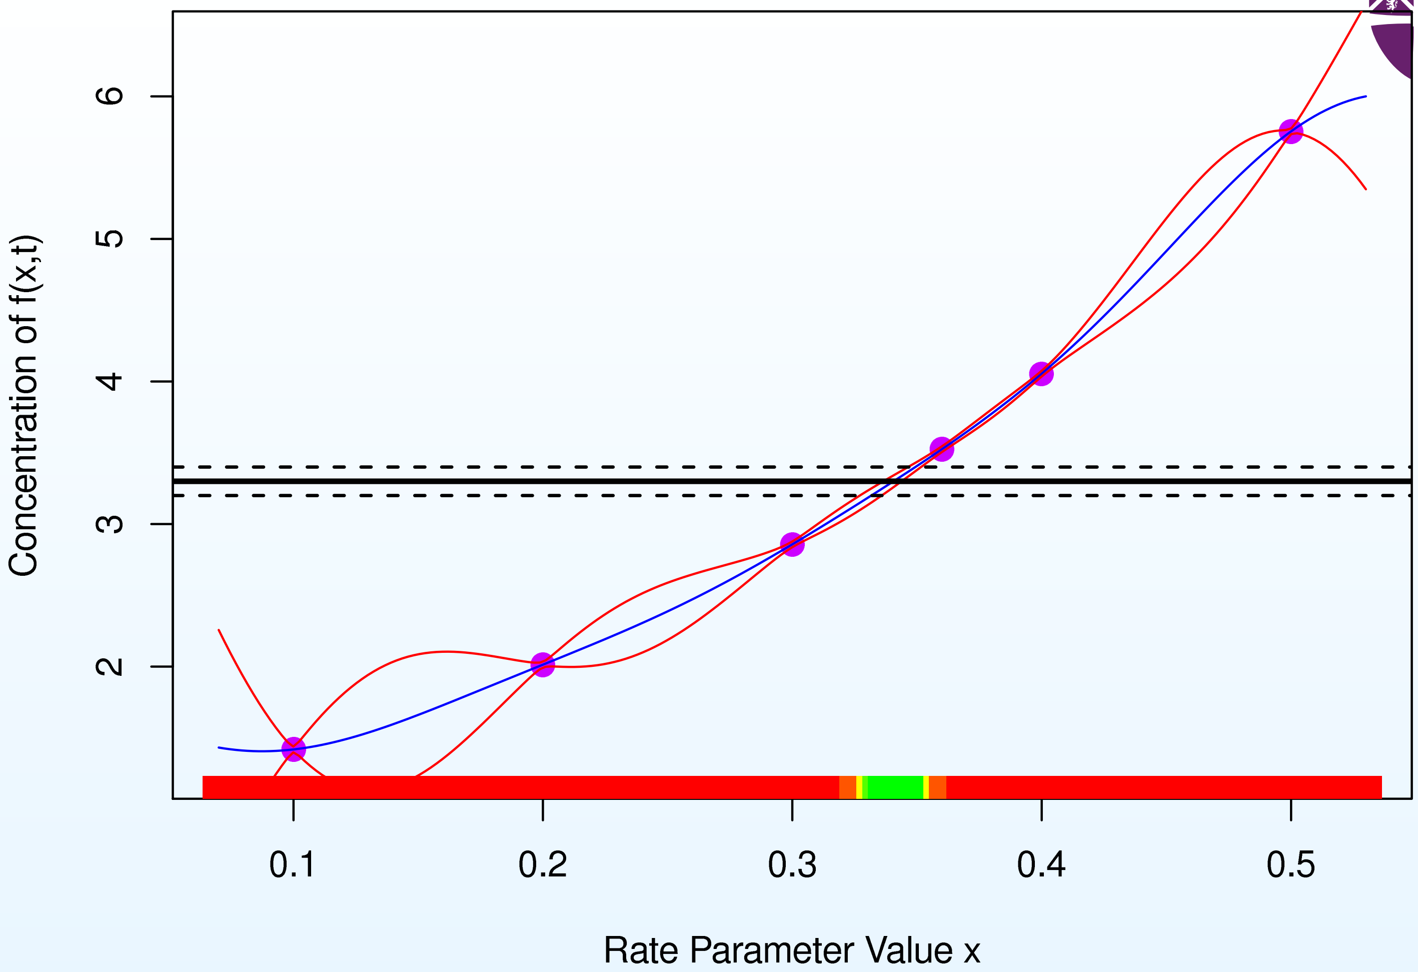

- We perform a 2nd iteration or wave of runs to improve emulator accuracy.
- The runs are located only at non-implausible (green/yellow) points.

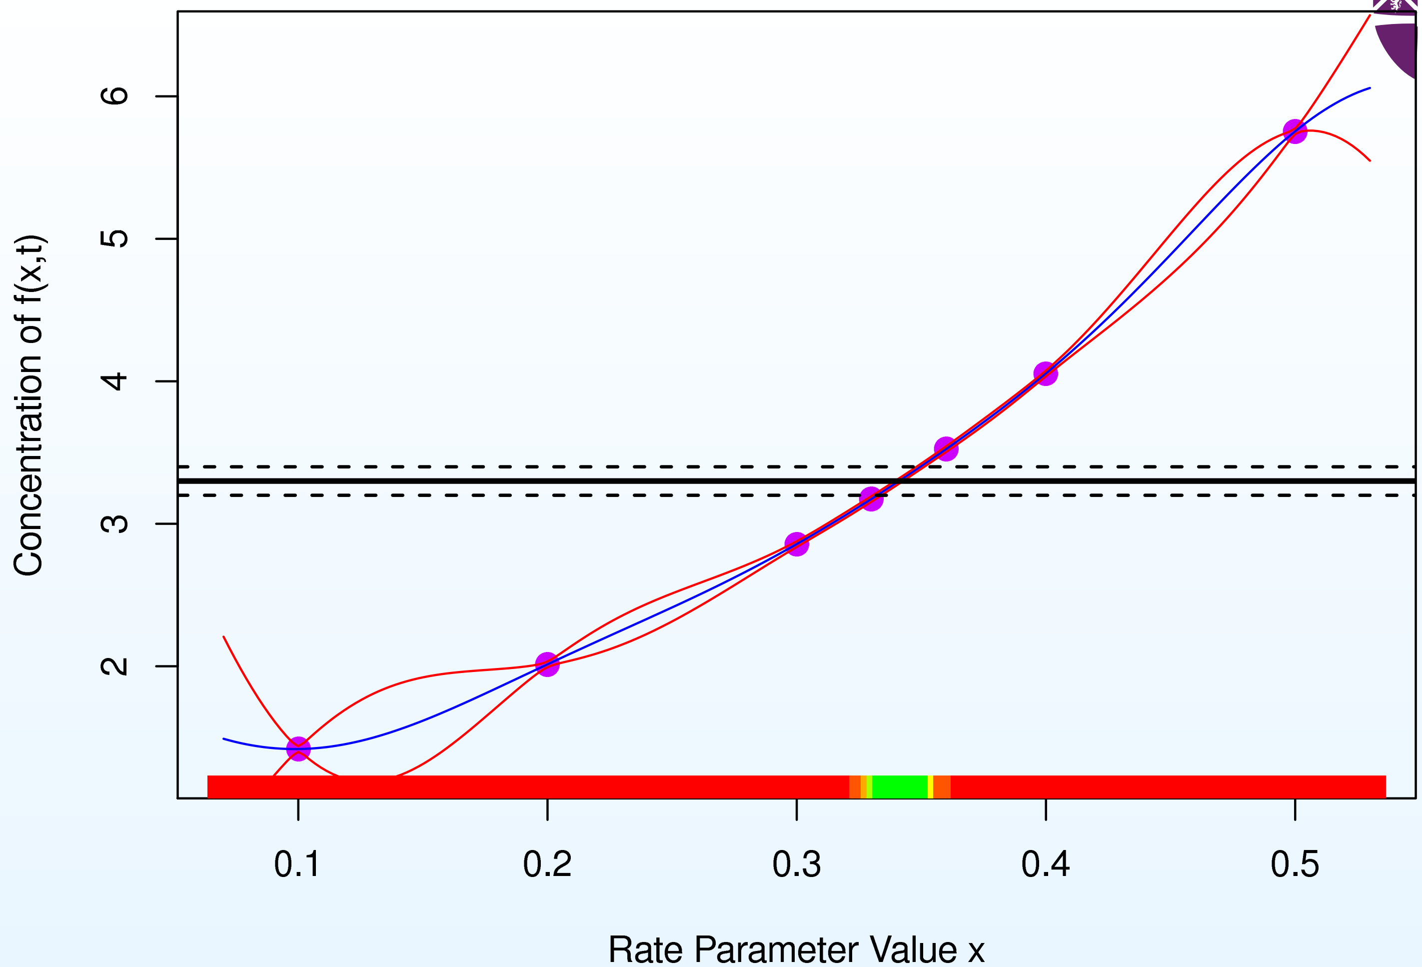

- We perform a 2nd iteration or wave of runs to improve emulator accuracy.
- The runs are located only at non-implausible (green/yellow) points.
- Now the emulator is more accurate than the observation, and we can identify the set of all x values of interest.

Emulation methods

We fit the emulators, given a collection of carefully chosen model evaluations, using our favourite statistical tools - generalised least squares, maximum likelihood, Bayes - with a generous helping of expert judgement.

Emulation methods

We fit the emulators, given a collection of carefully chosen model evaluations, using our favourite statistical tools - generalised least squares, maximum likelihood, Bayes - with a generous helping of expert judgement.

If the model is slow to evaluate, we typically create an informed prior assessment based on a fast approximation, then combine with a carefully designed set of runs of the full simulator to construct the emulator.

Emulation methods

We fit the emulators, given a collection of carefully chosen model evaluations, using our favourite statistical tools - generalised least squares, maximum likelihood, Bayes - with a generous helping of expert judgement.

If the model is slow to evaluate, we typically create an informed prior assessment based on a fast approximation, then combine with a carefully designed set of runs of the full simulator to construct the emulator.

We use detailed diagnostics to check emulator validity.

Emulation methods

We fit the emulators, given a collection of carefully chosen model evaluations, using our favourite statistical tools - generalised least squares, maximum likelihood, Bayes - with a generous helping of expert judgement.

If the model is slow to evaluate, we typically create an informed prior assessment based on a fast approximation, then combine with a carefully designed set of runs of the full simulator to construct the emulator.

We use detailed diagnostics to check emulator validity.

We can use the emulator to match the simulator to real system data either by model calibration or by our preferred method of **history matching** i.e. finding the collection of all input choices for which the simulator output agrees with the data to within the assessed uncertainty bounds.

Emulation methods

We fit the emulators, given a collection of carefully chosen model evaluations, using our favourite statistical tools - generalised least squares, maximum likelihood, Bayes - with a generous helping of expert judgement.

If the model is slow to evaluate, we typically create an informed prior assessment based on a fast approximation, then combine with a carefully designed set of runs of the full simulator to construct the emulator.

We use detailed diagnostics to check emulator validity.

We can use the emulator to match the simulator to real system data either by model calibration or by our preferred method of **history matching** i.e. finding the collection of all input choices for which the simulator output agrees with the data to within the assessed uncertainty bounds.

Unlike the model, the emulator is fast to evaluate. Therefore construction of the emulator allows us to go beyond scenario analysis and explore the future behaviour of the model across the whole range of physically meaningful input specifications.

The Bayesian approach

The (subjective) Bayesian approach is based on a very simple collection of ideas.

The Bayesian approach

The (subjective) Bayesian approach is based on a very simple collection of ideas.

You are uncertain about many things in the world.

The Bayesian approach

The (subjective) Bayesian approach is based on a very simple collection of ideas.

You are uncertain about many things in the world.

You can quantify your uncertainties as

The Bayesian approach

The (subjective) Bayesian approach is based on a very simple collection of ideas.

You are uncertain about many things in the world.

You can quantify your uncertainties as
[1] probabilities, for the quantities you are interested in,

The Bayesian approach

The (subjective) Bayesian approach is based on a very simple collection of ideas.

You are uncertain about many things in the world.

You can quantify your uncertainties as
[1] probabilities, for the quantities you are interested in,
[2]conditional probabilities for data that you might collect given the things you are interested in.

The Bayesian approach

The (subjective) Bayesian approach is based on a very simple collection of ideas.

You are uncertain about many things in the world.

You can quantify your uncertainties as
[1] probabilities, for the quantities you are interested in,
[2]conditional probabilities for data that you might collect given the things you are interested in.

When data arrives, Bayes theorem tells you how to move from your prior probabilities to new conditional probabilities for the quantities of interest.

The Bayesian approach

The (subjective) Bayesian approach is based on a very simple collection of ideas.

You are uncertain about many things in the world.

You can quantify your uncertainties as
[1] probabilities, for the quantities you are interested in,
[2]conditional probabilities for data that you might collect given the things you are interested in.

When data arrives, Bayes theorem tells you how to move from your prior probabilities to new conditional probabilities for the quantities of interest.

If you need to make decisions, then you may also specify a utility function, given which your preferred decision is that which maximises expected utility with respect to your conditional probability distribution for these quantities.

The Bayesian approach

The (subjective) Bayesian approach is based on a very simple collection of ideas.

You are uncertain about many things in the world.

You can quantify your uncertainties as
[1] probabilities, for the quantities you are interested in,
[2]conditional probabilities for data that you might collect given the things you are interested in.

When data arrives, Bayes theorem tells you how to move from your prior probabilities to new conditional probabilities for the quantities of interest.

If you need to make decisions, then you may also specify a utility function, given which your preferred decision is that which maximises expected utility with respect to your conditional probability distribution for these quantities.

Bayesian statistics provides a wide range of tools for constructing and validating such uncertainty judgements and decision choices.

Example

Suppose that either
you have a particular (rare) disease (event D)
or you don't have the disease (event $\overline{D}$).

Example

Suppose that either
you have a particular (rare) disease (event D)
or you don't have the disease (event $\overline{D}$).

You take a diagnostic test which gives either
a positive response (event $+$),
or a negative response (event $-$).

Example

Suppose that either
you have a particular (rare) disease (event D)
or you don't have the disease (event $\overline{D}$).

You take a diagnostic test which gives either
a positive response (event $+$),
or a negative response (event $-$).

The test is judged to be 99% reliable, e.g

$$\mathrm{P}(+|D) = 0.99, \; \mathrm{P}(+|\overline{D}) = 0.01.$$

Example

Suppose that either
you have a particular (rare) disease (event D)
or you don't have the disease (event $\overline{D}$).

You take a diagnostic test which gives either
a positive response (event $+$),
or a negative response (event $-$).

The test is judged to be 99% reliable, e.g

$$\mathrm{P}(+|D) = 0.99,\ \mathrm{P}(+|\overline{D}) = 0.01.$$

You take the test and get a positive result. Do you have the disease?

Analysis

Suppose that the proportion, p, of people with the disease is 0.001.

Analysis

Suppose that the proportion, p, of people with the disease is 0.001.

Then we may apply Bayes theorem, giving

$$\mathrm{P}(D|+) = \frac{\mathrm{P}(+|D)\mathrm{P}(D)}{\mathrm{P}(+)}$$

Analysis

Suppose that the proportion, p, of people with the disease is 0.001.

Then we may apply Bayes theorem, giving

$$\begin{aligned} \mathrm{P}(D|+) &= \frac{\mathrm{P}(+|D)\mathrm{P}(D)}{\mathrm{P}(+)} \\ &= \frac{\mathrm{P}(+|D)\mathrm{P}(D)}{\mathrm{P}(+|D)\mathrm{P}(D) + \mathrm{P}(+|\overline{D})\mathrm{P}(\overline{D})} \end{aligned}$$

Analysis

Suppose that the proportion, p, of people with the disease is 0.001.

Then we may apply Bayes theorem, giving

$$\begin{aligned}
\mathrm{P}(D|+) &= \frac{\mathrm{P}(+|D)\mathrm{P}(D)}{\mathrm{P}(+)} \\
&= \frac{\mathrm{P}(+|D)\mathrm{P}(D)}{\mathrm{P}(+|D)\mathrm{P}(D) + \mathrm{P}(+|\overline{D})\mathrm{P}(\overline{D})} = 0.09
\end{aligned}$$

Analysis

Suppose that the proportion, p, of people with the disease is 0.001.

Then we may apply Bayes theorem, giving

$$
\begin{aligned}
\mathrm{P}(D|+) &= \frac{\mathrm{P}(+|D)\mathrm{P}(D)}{\mathrm{P}(+)} \\
&= \frac{\mathrm{P}(+|D)\mathrm{P}(D)}{\mathrm{P}(+|D)\mathrm{P}(D) + \mathrm{P}(+|\overline{D})\mathrm{P}(\overline{D})} = 0.09
\end{aligned}
$$

We see that most likely you do not have the disease, but have scored a false positive.

Analysis

Suppose that the proportion, p, of people with the disease is 0.001.

Then we may apply Bayes theorem, giving

$$
\begin{aligned}
\mathrm{P}(D|+) &= \frac{\mathrm{P}(+|D)\mathrm{P}(D)}{\mathrm{P}(+)} \\
&= \frac{\mathrm{P}(+|D)\mathrm{P}(D)}{\mathrm{P}(+|D)\mathrm{P}(D)+\mathrm{P}(+|\overline{D})\mathrm{P}(\overline{D})} = 0.09
\end{aligned}
$$

We see that most likely you do not have the disease, but have scored a false positive.

Bayes analysis reaches the appropriate conclusion by combining evidence from the data ($\mathrm{P}(+|D)$) with prior judgements ($\mathrm{P}(D)$).

Analysis

Suppose that the proportion, p, of people with the disease is 0.001.

Then we may apply Bayes theorem, giving

$$\begin{aligned} \mathrm{P}(D|+) &= \frac{\mathrm{P}(+|D)\mathrm{P}(D)}{\mathrm{P}(+)} \\ &= \frac{\mathrm{P}(+|D)\mathrm{P}(D)}{\mathrm{P}(+|D)\mathrm{P}(D) + \mathrm{P}(+|\overline{D})\mathrm{P}(\overline{D})} = 0.09 \end{aligned}$$

We see that most likely you do not have the disease, but have scored a false positive.

Bayes analysis reaches the appropriate conclusion by combining evidence from the data ($\mathrm{P}(+|D)$) with prior judgements ($\mathrm{P}(D)$).

The Bayesian approach is well developed, popular and effective.

The Bayes linear approach

The Bayesian approach can be difficult in large problems because of the extreme level of detail which is required in the specification of beliefs.

The Bayes linear approach

The Bayesian approach can be difficult in large problems because of the extreme level of detail which is required in the specification of beliefs.

The Bayes linear approach concerns problems in which we combine prior judgements of uncertainty with observational data, and we use expectation rather than probability as the primitive.

The Bayes linear approach

The Bayesian approach can be difficult in large problems because of the extreme level of detail which is required in the specification of beliefs.

The Bayes linear approach concerns problems in which we combine prior judgements of uncertainty with observational data, and we use expectation rather than probability as the primitive.

This approach is similar in spirit to a full Bayes analysis, but uses a much simpler approach for prior specification and analysis, and so offers a practical methodology for analysing partially specified beliefs for very large problems.

The Bayes linear approach

The Bayesian approach can be difficult in large problems because of the extreme level of detail which is required in the specification of beliefs.

The Bayes linear approach concerns problems in which we combine prior judgements of uncertainty with observational data, and we use expectation rather than probability as the primitive.

This approach is similar in spirit to a full Bayes analysis, but uses a much simpler approach for prior specification and analysis, and so offers a practical methodology for analysing partially specified beliefs for very large problems.

For a detailed treatment, see
Bayes linear Statistics: Theory and Methods, 2007, (Wiley)
Michael Goldstein and David Wooff

Limitations of physical models

A physical model is a description of the way in which system properties (the inputs to the model) affect system behaviour (the output of the model). This description involves two basic types of simplification.

Limitations of physical models

A physical model is a description of the way in which
system properties (the inputs to the model)
affect system behaviour (the output of the model).
This description involves two basic types of simplification.

(i) we approximate the properties of the system (as these properties are too complicated to describe fully and anyway we don't know them)

Limitations of physical models

A physical model is a description of the way in which
system properties (the inputs to the model)
affect system behaviour (the output of the model).
This description involves two basic types of simplification.

(i) we approximate the properties of the system (as these properties are too complicated to describe fully and anyway we don't know them)

(ii) we approximate the rules for finding system behaviour given system properties (because of necessary mathematical simplifications, simplifications for numerical tractability, and because we do not fully understand the physical laws which govern the process).

Limitations of physical models

A physical model is a description of the way in which system properties (the inputs to the model) affect system behaviour (the output of the model). This description involves two basic types of simplification.

(i) we approximate the properties of the system (as these properties are too complicated to describe fully and anyway we don't know them)

(ii) we approximate the rules for finding system behaviour given system properties (because of necessary mathematical simplifications, simplifications for numerical tractability, and because we do not fully understand the physical laws which govern the process).

Neither of these approximations invalidates the modelling process. Problems only arise when we forget these simplifications and confuse the analysis of the model with the corresponding analysis for the physical system itself.

Relating the model and the system

Model evaluations

Actual system

System observations

1. We start with a collection of model evaluations, and some observations on actual system
2. We link the model evaluations to the evaluation of the model at the (unknown) system values x^* for the inputs
3. We link the system evaluation to the actual system by adding model discrepancy
4. We incorporate measurement error into the observations

Relating the model and the system

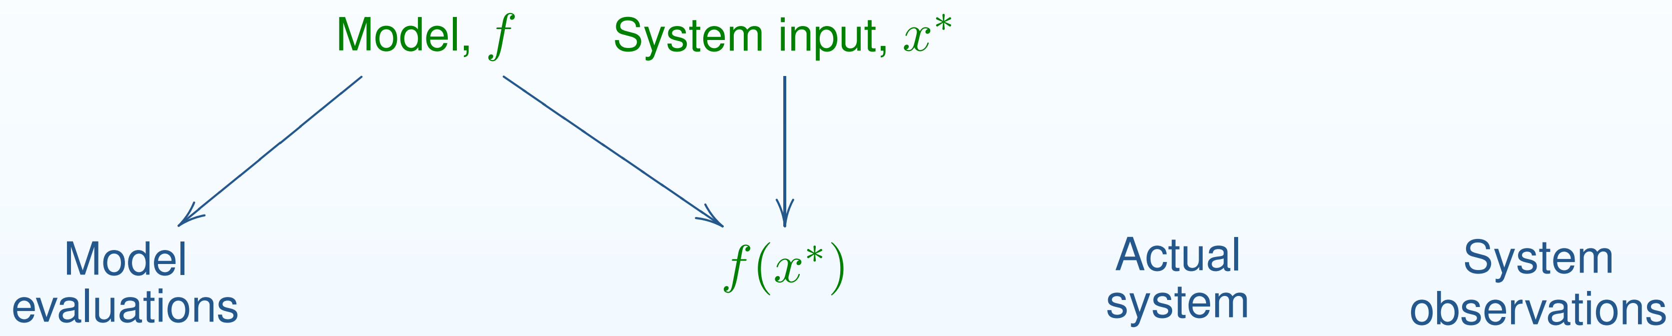

1. We start with a collection of model evaluations, and some observations on actual system
2. We link the model evaluations to the evaluation of the model at the (unknown) system values x^* for the inputs
3. We link the system evaluation to the actual system by adding model discrepancy
4. We incorporate measurement error into the observations

Relating the model and the system

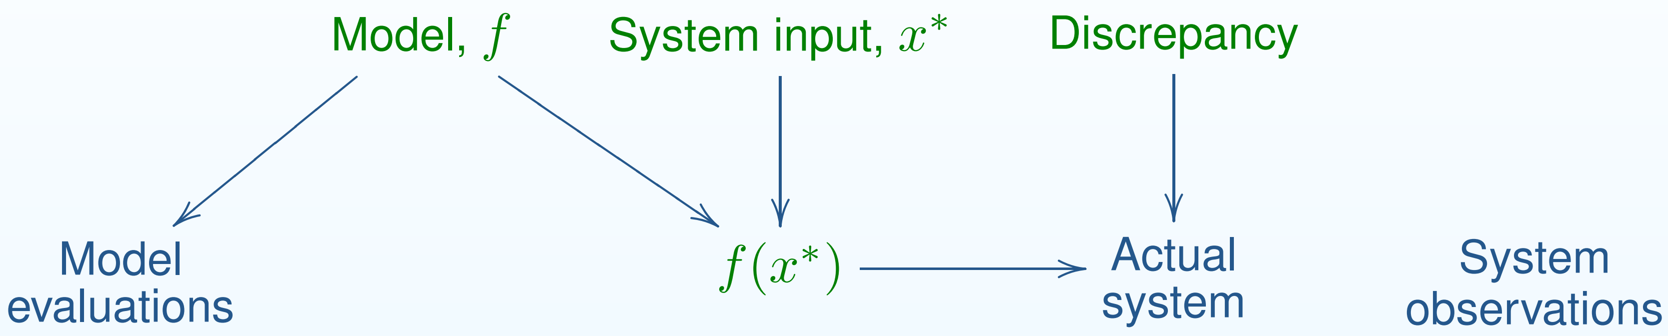

1. We start with a collection of model evaluations, and some observations on actual system
2. We link the model evaluations to the evaluation of the model at the (unknown) system values x^* for the inputs
3. We link the system evaluation to the actual system by adding model discrepancy
4. We incorporate measurement error into the observations

Relating the model and the system

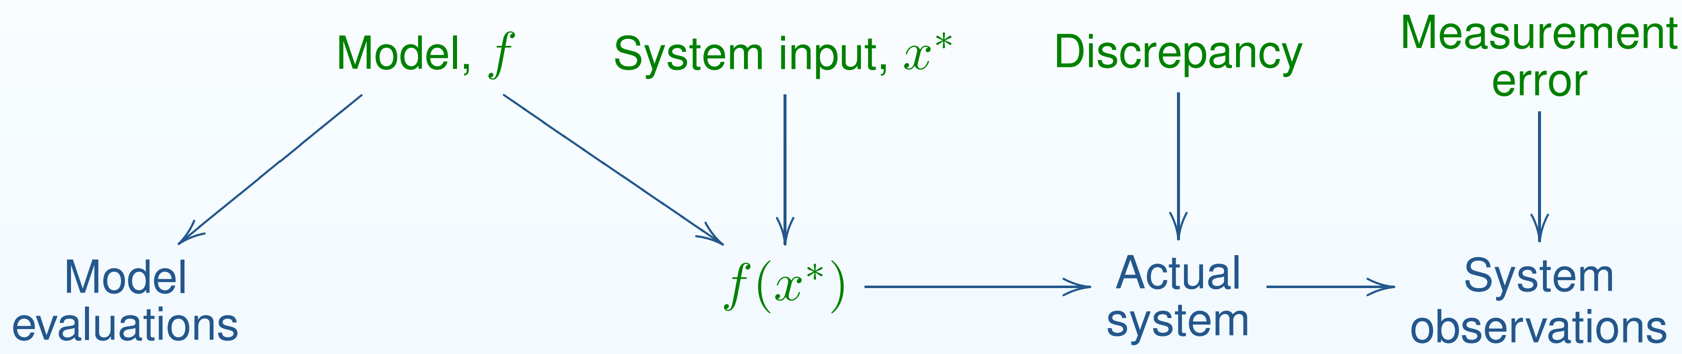

1. We start with a collection of model evaluations, and some observations on actual system
2. We link the model evaluations to the evaluation of the model at the (unknown) system values x^* for the inputs
3. We link the system evaluation to the actual system by adding model discrepancy
4. We incorporate measurement error into the observations

Example: adding model discrepancy

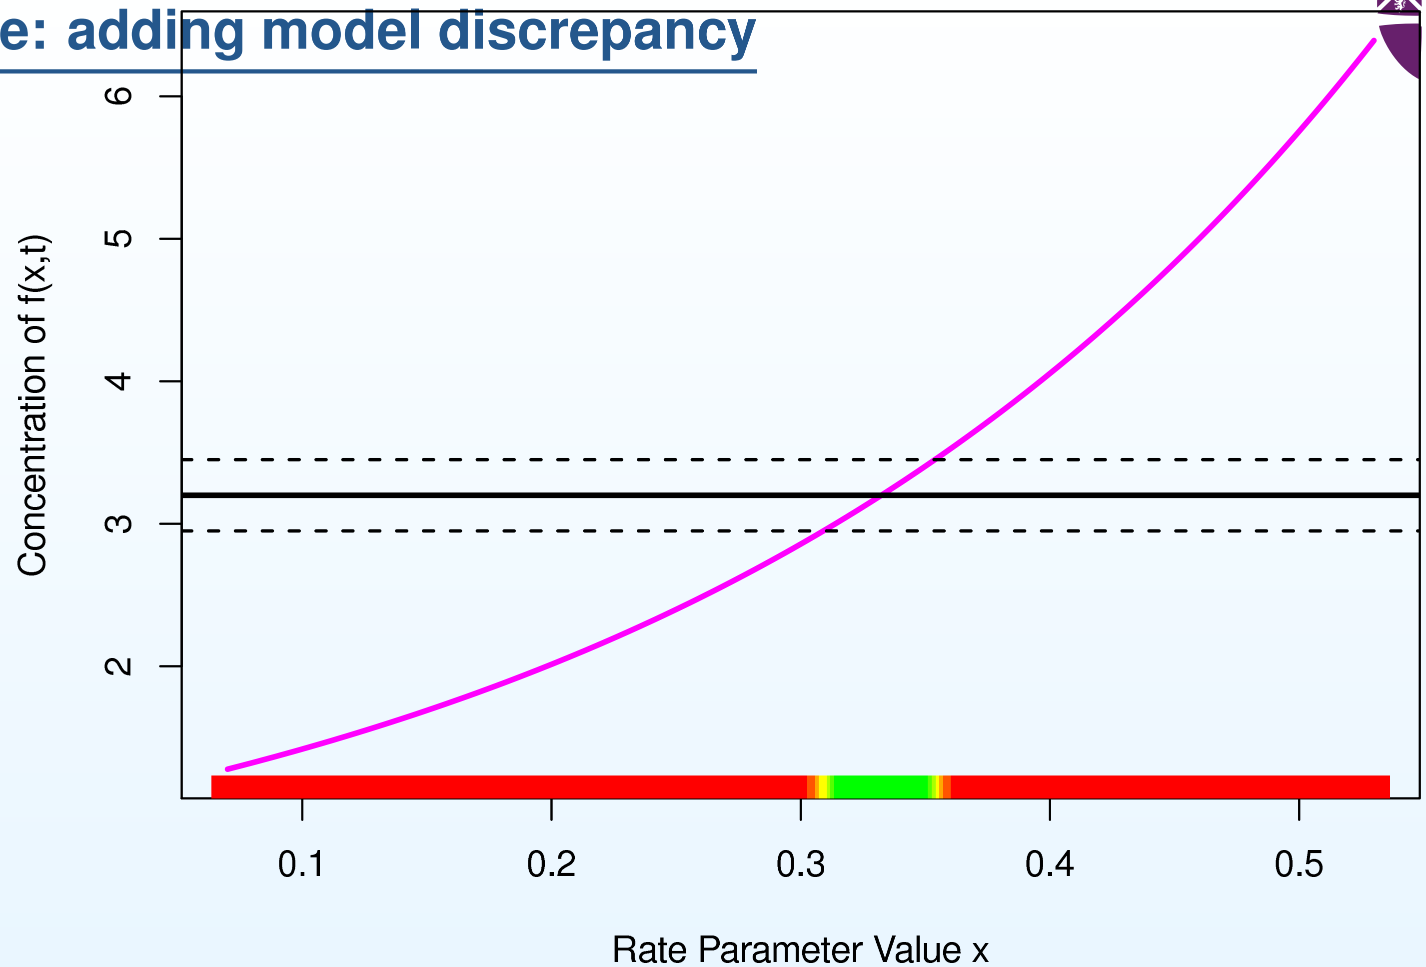

- The notion of model discrepancy is related to how accurate we believe the model to be.

Example: adding model discrepancy

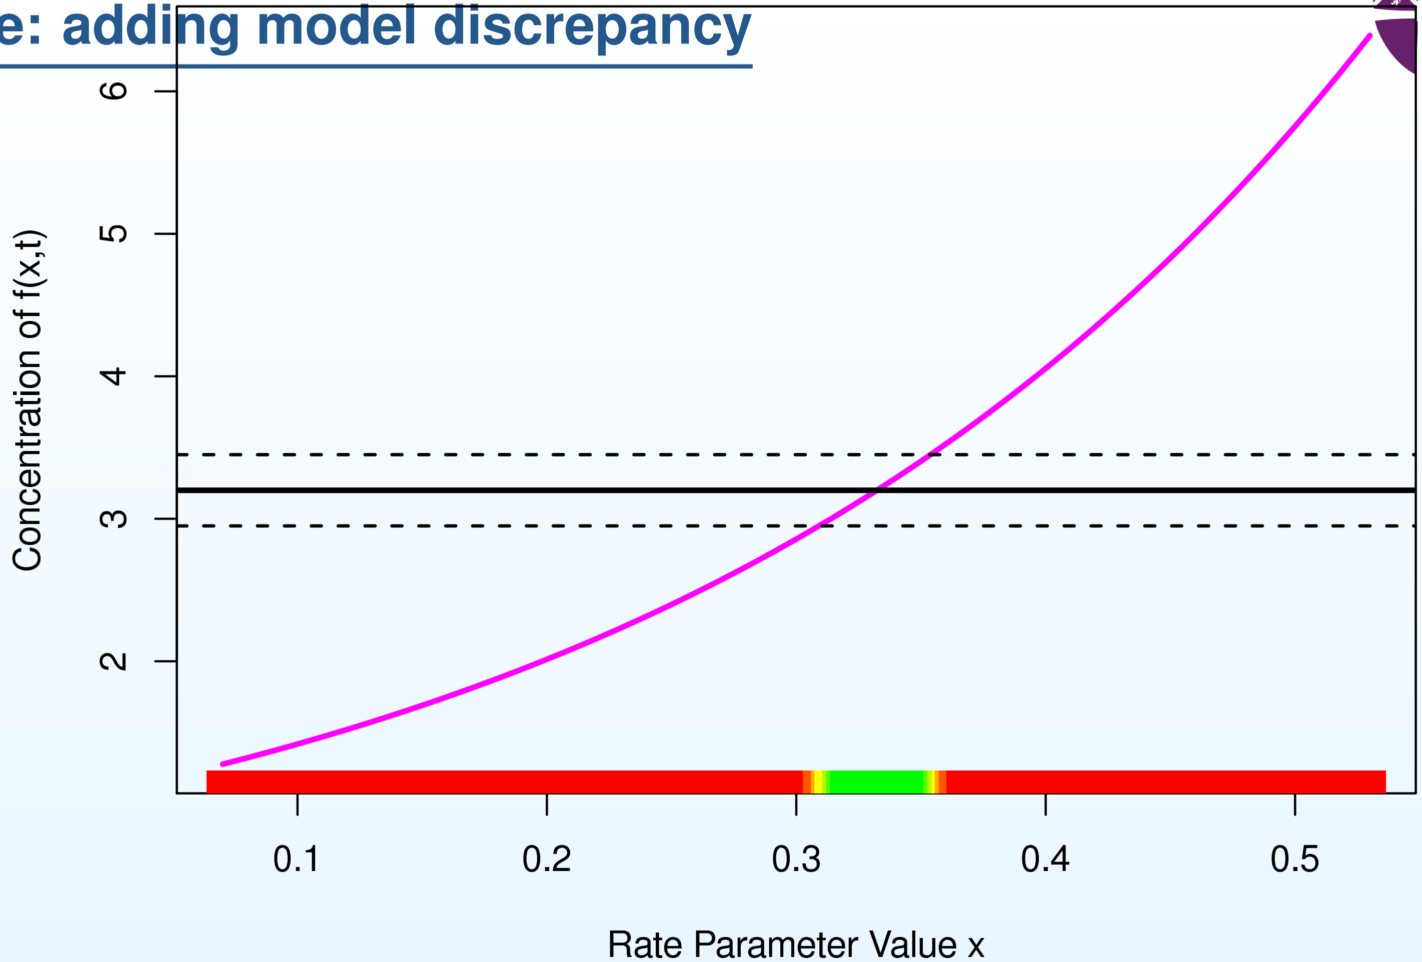

- The notion of model discrepancy is related to how accurate we believe the model to be.
- This uncertainty arises from many issues e.g. is the form of the model (the differential equation) appropriate, is the model a simplified description of a more complex system, is there uncertainty in the initial conditions etc?

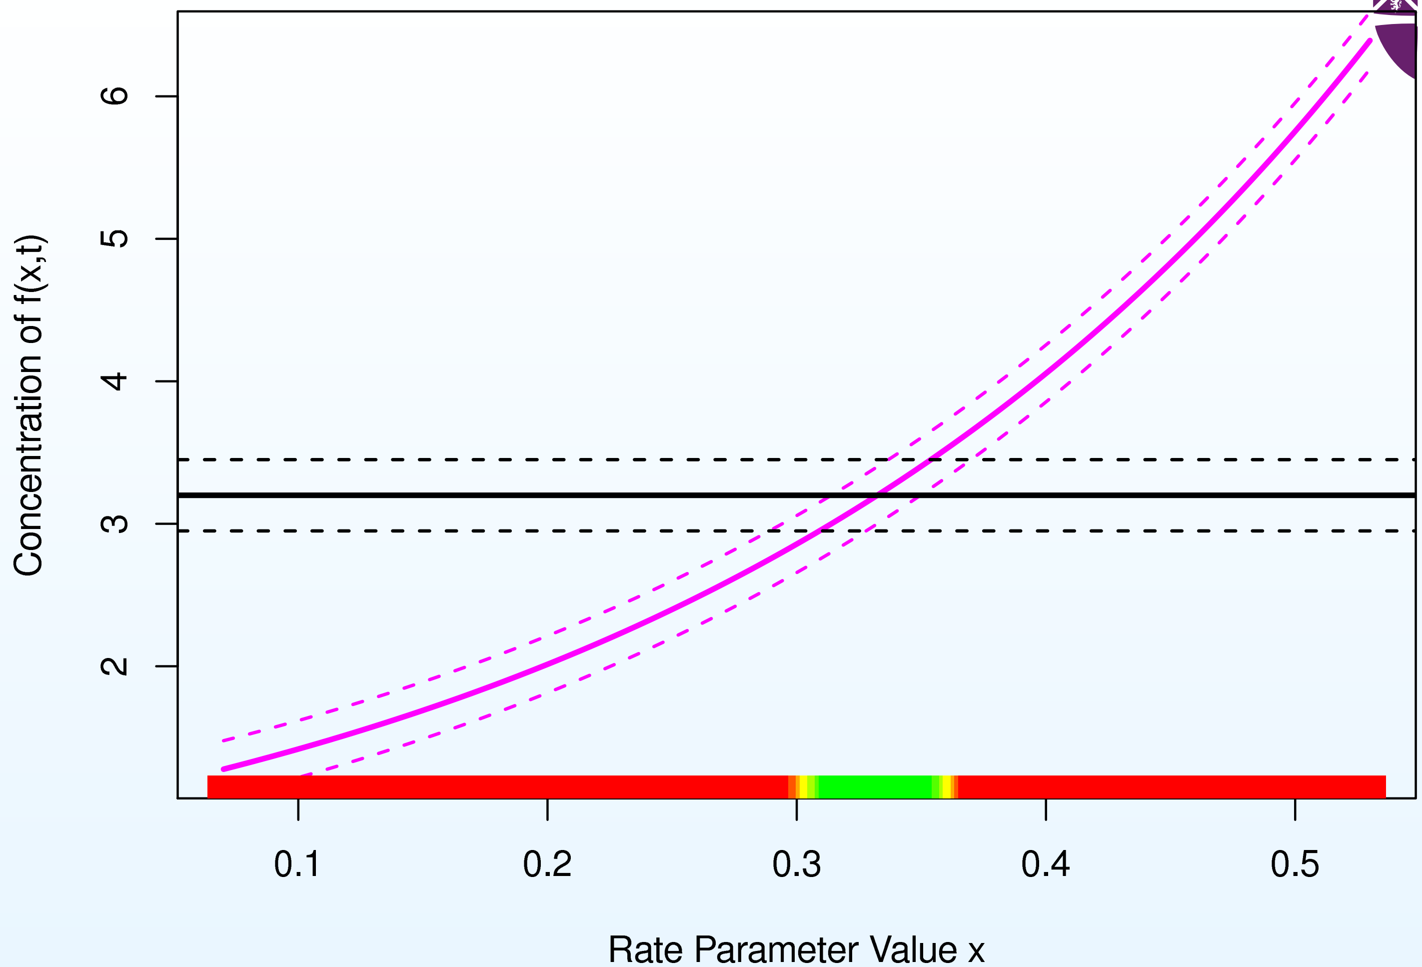

- Model discrepancy is represented as uncertainty around the model output $f(x)$ itself: here the purple dashed lines.

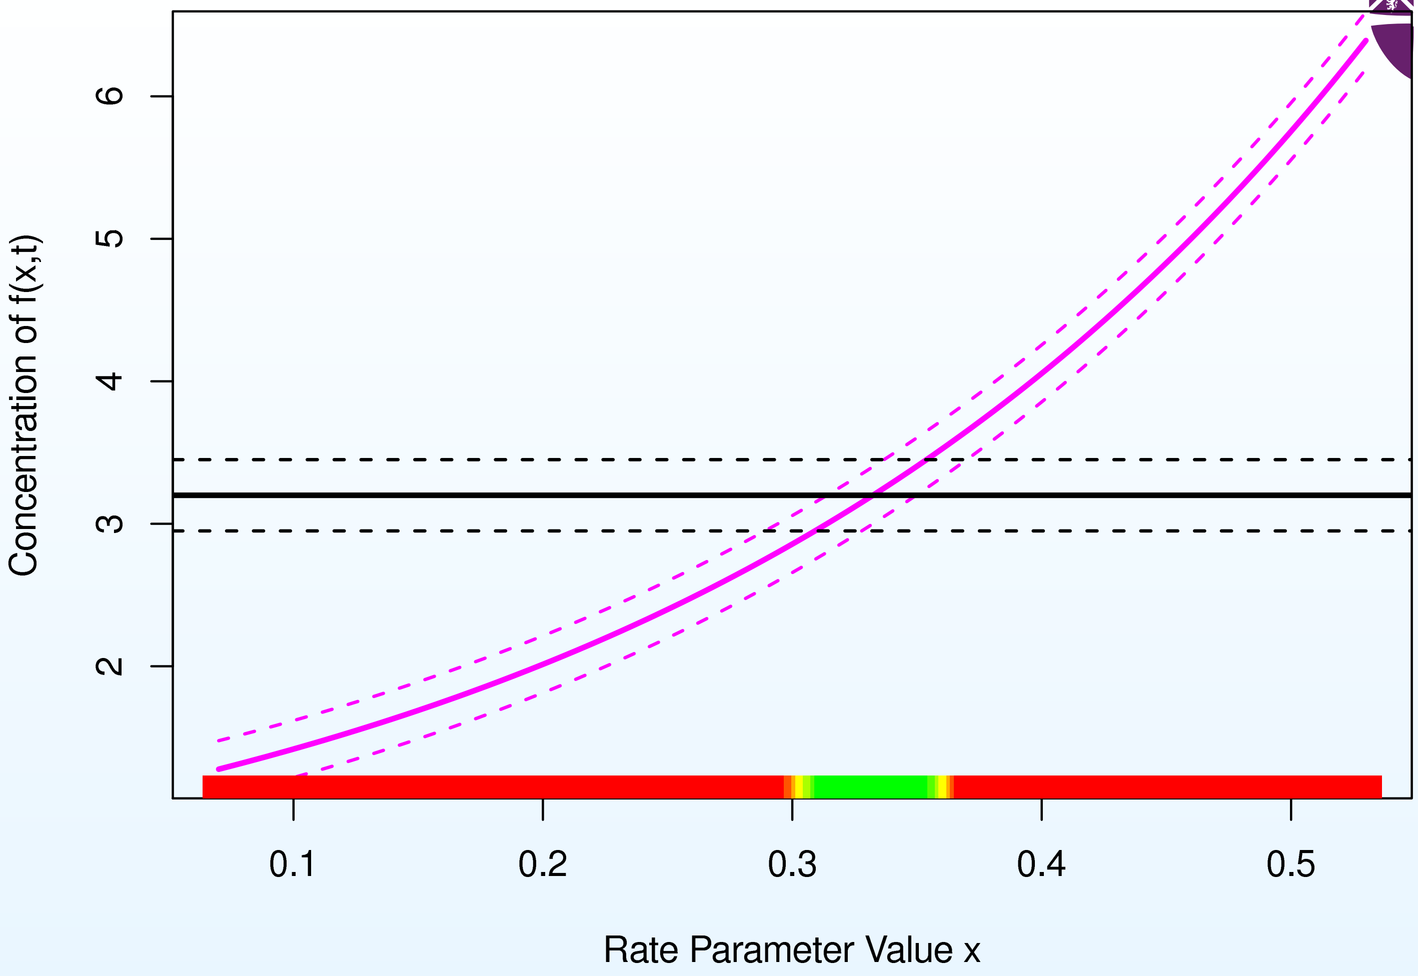

- Model discrepancy is represented as uncertainty around the model output $f(x)$ itself: here the purple dashed lines.

- This results in more uncertainty in x, and hence a larger range of x values.

Internal and external uncertainty

Structural uncertainty assessessment should form a central part of the problem analysis. We may distinguish two types of model discrepancy.

Internal and external uncertainty

Structural uncertainty assessessment should form a central part of the problem analysis. We may distinguish two types of model discrepancy.
(i) **Internal discrepancy** This relates to any aspect of structural discrepancy whose magnitude we may assess by experiments on the computer simulator.

Internal and external uncertainty

Structural uncertainty assessessment should form a central part of the problem analysis. We may distinguish two types of model discrepancy.
(i) **Internal discrepancy** This relates to any aspect of structural discrepancy whose magnitude we may assess by experiments on the computer simulator.
For example, we may vary parameters held fixed in the standard analysis,
we may add random noise to the state vector which the model propagates,
we may allow parameters to vary over time.

Internal and external uncertainty

Structural uncertainty assessessment should form a central part of the problem analysis. We may distinguish two types of model discrepancy.

(i) **Internal discrepancy** This relates to any aspect of structural discrepancy whose magnitude we may assess by experiments on the computer simulator.

For example, we may vary parameters held fixed in the standard analysis,
we may add random noise to the state vector which the model propagates,
we may allow parameters to vary over time.

Internal discrepancy analysis gives a lower bound on the structural uncertainty that we must introduce into our model analyses.

Internal and external uncertainty

Structural uncertainty assessessment should form a central part of the problem analysis. We may distinguish two types of model discrepancy.

(i) **Internal discrepancy** This relates to any aspect of structural discrepancy whose magnitude we may assess by experiments on the computer simulator.
For example, we may vary parameters held fixed in the standard analysis,
we may add random noise to the state vector which the model propagates,
we may allow parameters to vary over time.
Internal discrepancy analysis gives a lower bound on the structural uncertainty that we must introduce into our model analyses.

(ii) **External discrepancy** This relates to inherent limitations of the modelling process embodied in the simulator. There are no experiments on the simulator which may reveal this magnitude. It is determined by a combination of expert judgements and statistical estimation.

Internal and external uncertainty

Structural uncertainty assessessment should form a central part of the problem analysis. We may distinguish two types of model discrepancy.

(i) **Internal discrepancy** This relates to any aspect of structural discrepancy whose magnitude we may assess by experiments on the computer simulator.

For example, we may vary parameters held fixed in the standard analysis,

we may add random noise to the state vector which the model propagates,

we may allow parameters to vary over time.

Internal discrepancy analysis gives a lower bound on the structural uncertainty that we must introduce into our model analyses.

(ii) **External discrepancy** This relates to inherent limitations of the modelling process embodied in the simulator. There are no experiments on the simulator which may reveal this magnitude. It is determined by a combination of expert judgements and statistical estimation.

We assess structural discrepancy for

(i) historical system outputs to improve our model fitting to data and

(ii) for future system outputs to improve our uncertainty statements for future system forecasts.

Sources of Uncertainty

Durham University

Sources of Uncertainty

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)

Sources of Uncertainty

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)
(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),

Sources of Uncertainty

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)

(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),

(iii) functional uncertainty (model evaluations take a long time, so the function is unknown almost everywhere)

Sources of Uncertainty

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)

(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),

(iii) functional uncertainty (model evaluations take a long time, so the function is unknown almost everywhere)

(iv) stochastic uncertainty (either the model is stochastic, or it should be),

Sources of Uncertainty

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)
(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),
(iii) functional uncertainty (model evaluations take a long time, so the function is unknown almost everywhere)
(iv) stochastic uncertainty (either the model is stochastic, or it should be),
(v) solution uncertainty (as the system equations can only be solved to some necessary level of approximation).

Sources of Uncertainty

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)
(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),
(iii) functional uncertainty (model evaluations take a long time, so the function is unknown almost everywhere)
(iv) stochastic uncertainty (either the model is stochastic, or it should be),
(v) solution uncertainty (as the system equations can only be solved to some necessary level of approximation).
(vi) structural uncertainty (the model only approximates the physical system),

Sources of Uncertainty

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)
(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),
(iii) functional uncertainty (model evaluations take a long time, so the function is unknown almost everywhere)
(iv) stochastic uncertainty (either the model is stochastic, or it should be),
(v) solution uncertainty (as the system equations can only be solved to some necessary level of approximation).
(vi) structural uncertainty (the model only approximates the physical system),
(vii) measurement uncertainty (as the model is calibrated against system data all of which is measured with error),

Sources of Uncertainty

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)

(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),

(iii) functional uncertainty (model evaluations take a long time, so the function is unknown almost everywhere)

(iv) stochastic uncertainty (either the model is stochastic, or it should be),

(v) solution uncertainty (as the system equations can only be solved to some necessary level of approximation).

(vi) structural uncertainty (the model only approximates the physical system),

(vii) measurement uncertainty (as the model is calibrated against system data all of which is measured with error),

(viii) multi-model uncertainty (usually we have not one but many models related to the physical system)

Sources of Uncertainty

(i) parametric uncertainty (each model requires a, typically high dimensional, parametric specification)
(ii) condition uncertainty (uncertainty as to boundary conditions, initial conditions, and forcing functions),
(iii) functional uncertainty (model evaluations take a long time, so the function is unknown almost everywhere)
(iv) stochastic uncertainty (either the model is stochastic, or it should be),
(v) solution uncertainty (as the system equations can only be solved to some necessary level of approximation).
(vi) structural uncertainty (the model only approximates the physical system),
(vii) measurement uncertainty (as the model is calibrated against system data all of which is measured with error),
(viii) multi-model uncertainty (usually we have not one but many models related to the physical system)
(ix) decision uncertainty (to use the model to influence real world outcomes, we need to relate things in the world that we can influence to inputs to the simulator and through outputs to actual impacts. These links are uncertain.)

Energy systems integration

Much of the CESI project involves energy systems talking to each other.

Energy systems integration

Much of the CESI project involves energy systems talking to each other.

So, for example, we might have system one, $f_1(x)$, whose output y forms an input to system two , $f_2(y, w)$.

Energy systems integration

Much of the CESI project involves energy systems talking to each other.

So, for example, we might have system one, $f_1(x)$, whose output y forms an input to system two , $f_2(y, w)$.

The way to assess uncertainty for the combined system, over the whole range of input choices, is

Energy systems integration

Much of the CESI project involves energy systems talking to each other.

So, for example, we might have system one, $f_1(x)$, whose output y forms an input to system two , $f_2(y, w)$.

The way to assess uncertainty for the combined system, over the whole range of input choices, is

[1] Build emulators, $\hat{f}_1$ and $\hat{f}_2$ for f_1 and f_2

Energy systems integration

Much of the CESI project involves energy systems talking to each other.

So, for example, we might have system one, $f_1(x)$, whose output y forms an input to system two , $f_2(y, w)$.

The way to assess uncertainty for the combined system, over the whole range of input choices, is

[1] Build emulators, $\hat{f}_1$ and $\hat{f}_2$ for f_1 and f_2

[2] Assess structural discrepancy δ_1 and δ_2 for the two models.

Energy systems integration

Much of the CESI project involves energy systems talking to each other.

So, for example, we might have system one, $f_1(x)$, whose output y forms an input to system two , $f_2(y, w)$.

The way to assess uncertainty for the combined system, over the whole range of input choices, is

[1] Build emulators, $\hat{f}_1$ and $\hat{f}_2$ for f_1 and f_2

[2] Assess structural discrepancy δ_1 and δ_2 for the two models.

[3] Assess uncertainty for the combined system by
(i) making random draws $\hat{y}$ from $\hat{f}_1(x) + \delta_1$

Energy systems integration

Much of the CESI project involves energy systems talking to each other.

So, for example, we might have system one, $f_1(x)$, whose output y forms an input to system two , $f_2(y, w)$.

The way to assess uncertainty for the combined system, over the whole range of input choices, is

[1] Build emulators, $\hat{f}_1$ and $\hat{f}_2$ for f_1 and f_2

[2] Assess structural discrepancy δ_1 and δ_2 for the two models.

[3] Assess uncertainty for the combined system by
(i) making random draws $\hat{y}$ from $\hat{f}_1(x) + \delta_1$
(ii) making random draws $\hat{u}$ from $\hat{f}_2(\hat{y}, w) + \delta_2$

Energy systems integration

Much of the CESI project involves energy systems talking to each other.

So, for example, we might have system one, $f_1(x)$, whose output y forms an input to system two , $f_2(y, w)$.

The way to assess uncertainty for the combined system, over the whole range of input choices, is

[1] Build emulators, $\hat{f}_1$ and $\hat{f}_2$ for f_1 and f_2

[2] Assess structural discrepancy δ_1 and δ_2 for the two models.

[3] Assess uncertainty for the combined system by
(i) making random draws $\hat{y}$ from $\hat{f}_1(x) + \delta_1$
(ii) making random draws $\hat{u}$ from $\hat{f}_2(\hat{y}, w) + \delta_2$

This approach is modular. We can emulate and assess structural discrepancy over each model separately, then combine all of the specifications to carry out the composite uncertainty analysis over any collection of subsystems of interest.

Concluding comments

There is a general Bayesian methodology for performing full uncertainty analyses on collections of complex physical systems (such as integrated energy systems), which are modelled by a collection of computer simulators.

Concluding comments

There is a general Bayesian methodology for performing full uncertainty analyses on collections of complex physical systems (such as integrated energy systems), which are modelled by a collection of computer simulators.

The two key features of this methodology are

Concluding comments

There is a general Bayesian methodology for performing full uncertainty analyses on collections of complex physical systems (such as integrated energy systems), which are modelled by a collection of computer simulators.

The two key features of this methodology are

(i) detailed simulator emulation, to move us beyond scenario analysis

Concluding comments

There is a general Bayesian methodology for performing full uncertainty analyses on collections of complex physical systems (such as integrated energy systems), which are modelled by a collection of computer simulators.

The two key features of this methodology are

(i) detailed simulator emulation, to move us beyond scenario analysis

(ii) careful structural discrepancy modelling, to make reliable uncertainty statements about the real world

Concluding comments

There is a general Bayesian methodology for performing full uncertainty analyses on collections of complex physical systems (such as integrated energy systems), which are modelled by a collection of computer simulators.

The two key features of this methodology are

(i) detailed simulator emulation, to move us beyond scenario analysis

(ii) careful structural discrepancy modelling, to make reliable uncertainty statements about the real world

It is important to design simulators which support these tasks.

Concluding comments

There is a general Bayesian methodology for performing full uncertainty analyses on collections of complex physical systems (such as integrated energy systems), which are modelled by a collection of computer simulators.

The two key features of this methodology are

(i) detailed simulator emulation, to move us beyond scenario analysis

(ii) careful structural discrepancy modelling, to make reliable uncertainty statements about the real world

It is important to design simulators which support these tasks.

For example, if the simulator is slow to evaluate, make it easy to create fast approximations to support emulation

Concluding comments

There is a general Bayesian methodology for performing full uncertainty analyses on collections of complex physical systems (such as integrated energy systems), which are modelled by a collection of computer simulators.

The two key features of this methodology are

(i) detailed simulator emulation, to move us beyond scenario analysis

(ii) careful structural discrepancy modelling, to make reliable uncertainty statements about the real world

It is important to design simulators which support these tasks.

For example, if the simulator is slow to evaluate, make it easy to create fast approximations to support emulation

if the simulator represents a model with many assumptions, make it easy to access and modify the code to allow simulator runs to quantify (internal) structural discrepancy.

References

Goldstein, M., Seheult, A. & Vernon, I. (2013). Assessing Model Adequacy. In Environmental Modelling: Finding Simplicity in Complexity. Mulligan, Mark. & Wainwright, John. Wiley-Blackwell. 435-449.

J. Cumming, M. Goldstein Bayes Linear Uncertainty Analysis for Oil Reservoirs Based on Multiscale Computer Experiments (2009), in the Handbook of Applied Bayesian Analysis,eds A. O'Hagan, M. West, OUP

I. Vernon, M. Goldstein, R. Bower (2010), Galaxy Formation: a Bayesian Uncertainty Analysis (with discussion) , Bayesian Analysis, 5(4): 619–670.

A. Lawson, M.Goldstein, C.Dent (2017) Bayesian Framework for Power Network Planning Under Uncertainty, Sustainable Energy, Grids and Networks, to appear

M. Goldstein and J.C.Rougier (2009). Reified Bayesian modelling and inference for physical systems (with discussion), JSPI, 139, 1221-1239

M. Goldstein Subjective Bayesian analysis: principles and practice (2006) Bayesian Analysis, 1, 403-420 (and 'Rejoinder to discussion': 465-472)