Systems Biology 1

Deterministic Modelling

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

Timetable

- Friday: 11-2
- Monday: 10-12, 1-2
- Both days: mixture of ‘lecture’ and ‘lab’ work.

Overview

- Friday
  - Brief introduction to Systems Biology and modelling.
  - Brief introduction to deterministic modelling with *Ordinary Differential Equations*
- Monday
  - Brief introduction to exact stochastic simulation with the *Gillespie algorithm*

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

Practical work

- We will use a solver written in Java
- I'm assuming that everyone is happy with compiling and running Java code from the command line?
  - Speak now if not…
- We do not have much time: I strongly advise you to complete the lab exercises (and build other models) outside the official teaching time.

Lecture notes

- On Moodle.
- They are not exhaustive.

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

What is Systems Biology?

University of Glasgow

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

What is modelling? What is a model?

The purpose of building a **model** of a **biological system** is to formulate in either a mathematical or diagrammatical manner, a **representation** of our **understanding** of the system. This process is useful in itself (it pulls together knowledge of the system into a single coherent place), but can also be used to test our understanding via making **predictions** of system behaviour that can be **verified** in the laboratory.

Q: can we gain knowledge from a model?

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

Types of model in SB

University of Glasgow

Networks

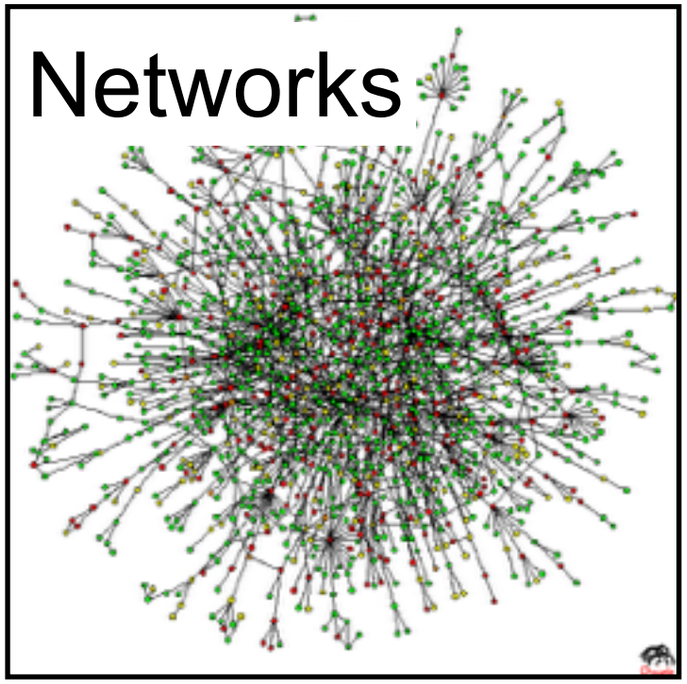

Data analysis

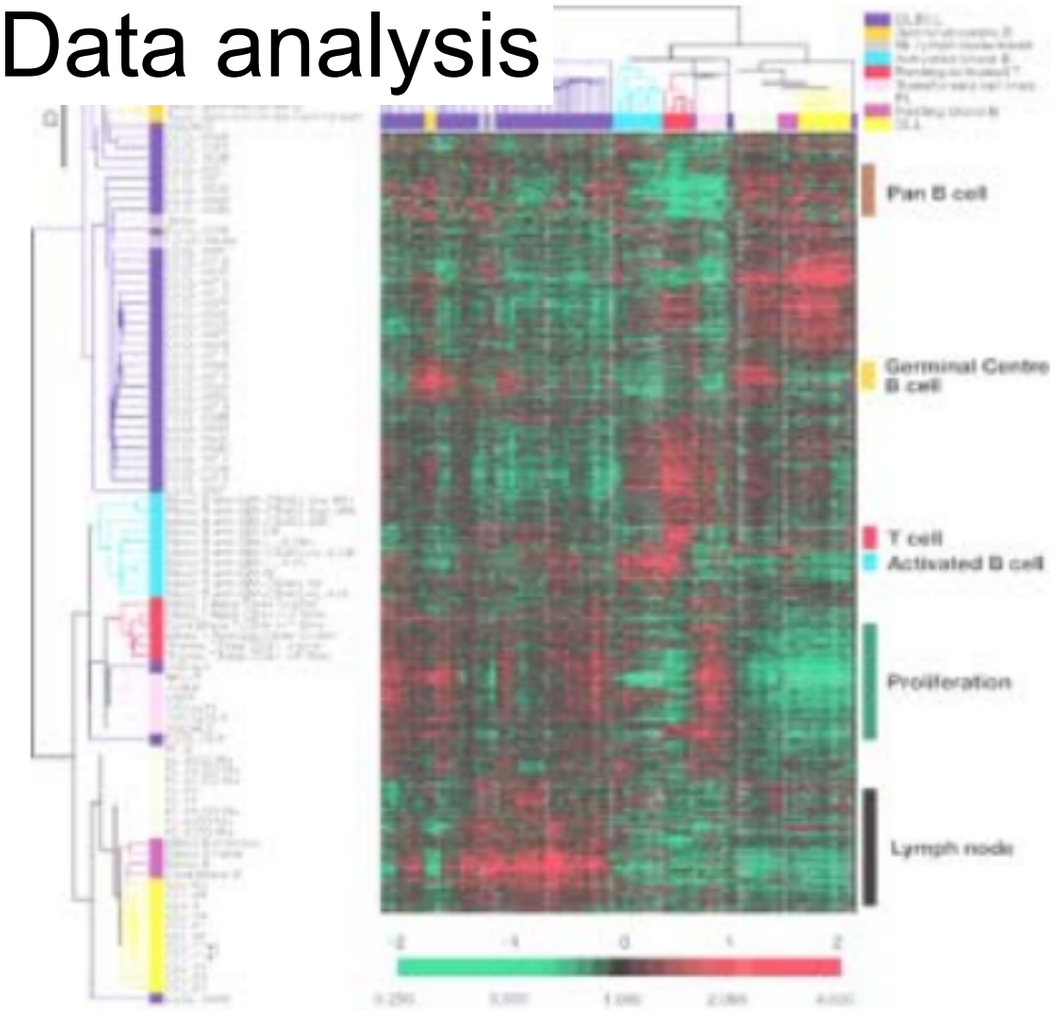

Dynamic/temporal

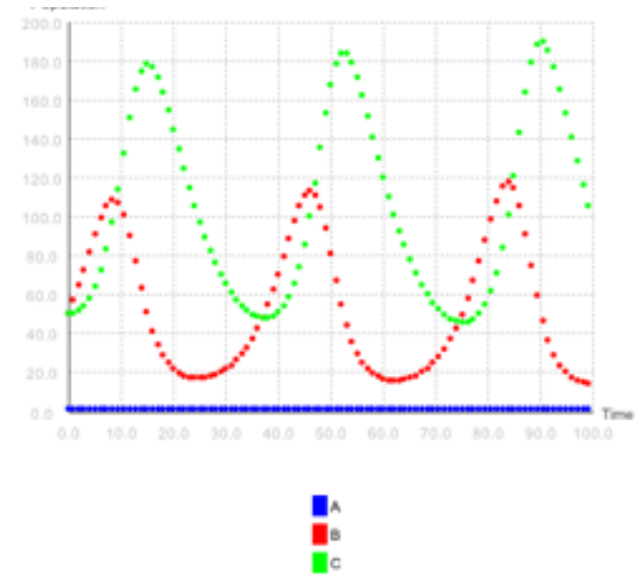

Next semester

This module

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

What is modelling? What is a model?

- For us, model = mathematical model
  - In particular, how things change over time.

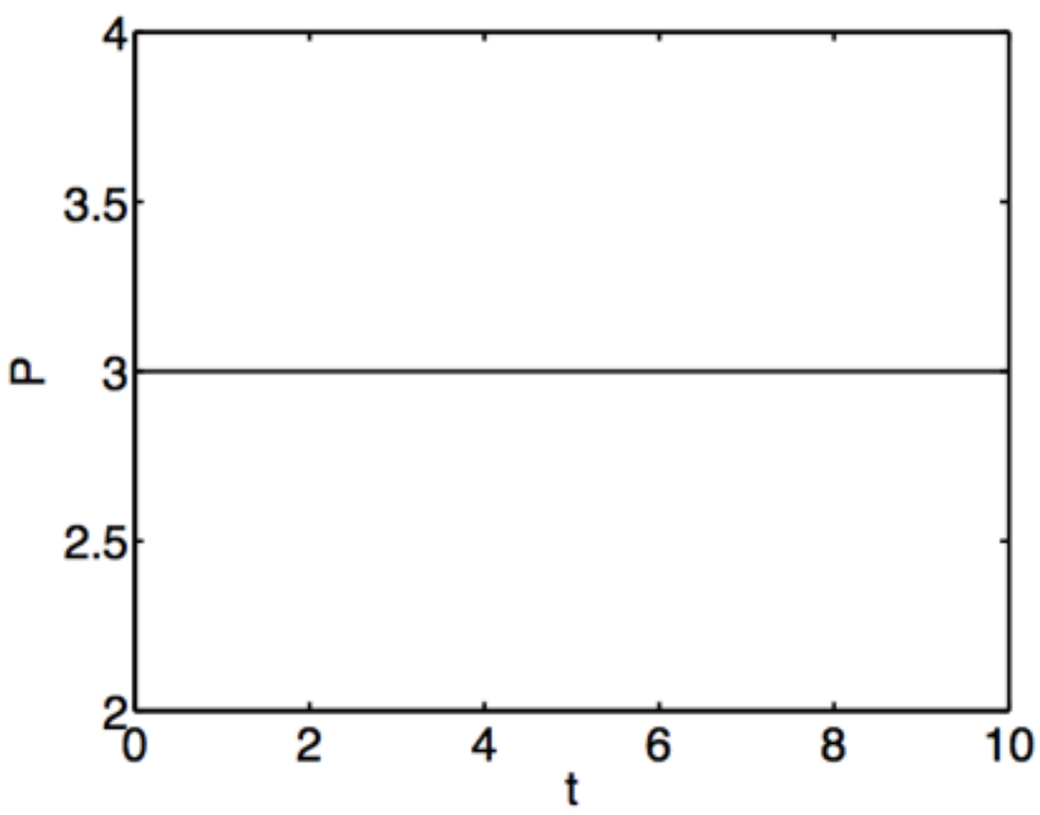

(a) $P = 3$

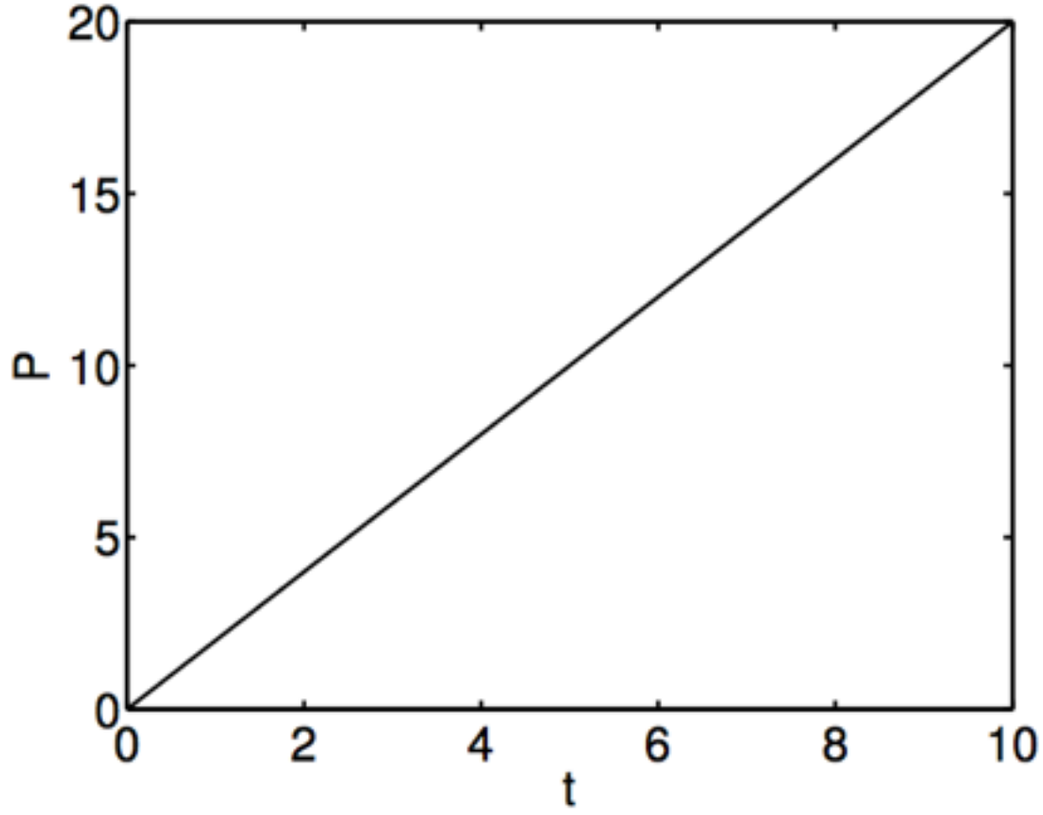

(b) $P = 2t$

[Two mathematical models describing how P behaves over time (t)]

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

What kind of equations?

- P = 2t tells us the value of P at any time, t.
- Our understanding of biological phenomena is not like this.
- e.g. reactions tell us how things **change**
- **things** = species like proteins and other interesting molecules

$$\mathrm{A} + \mathrm{B} \xrightarrow{k_1} \mathrm{A} + \mathrm{B} + \mathrm{B}$$

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

What kind of equations?

- How things change = gradients

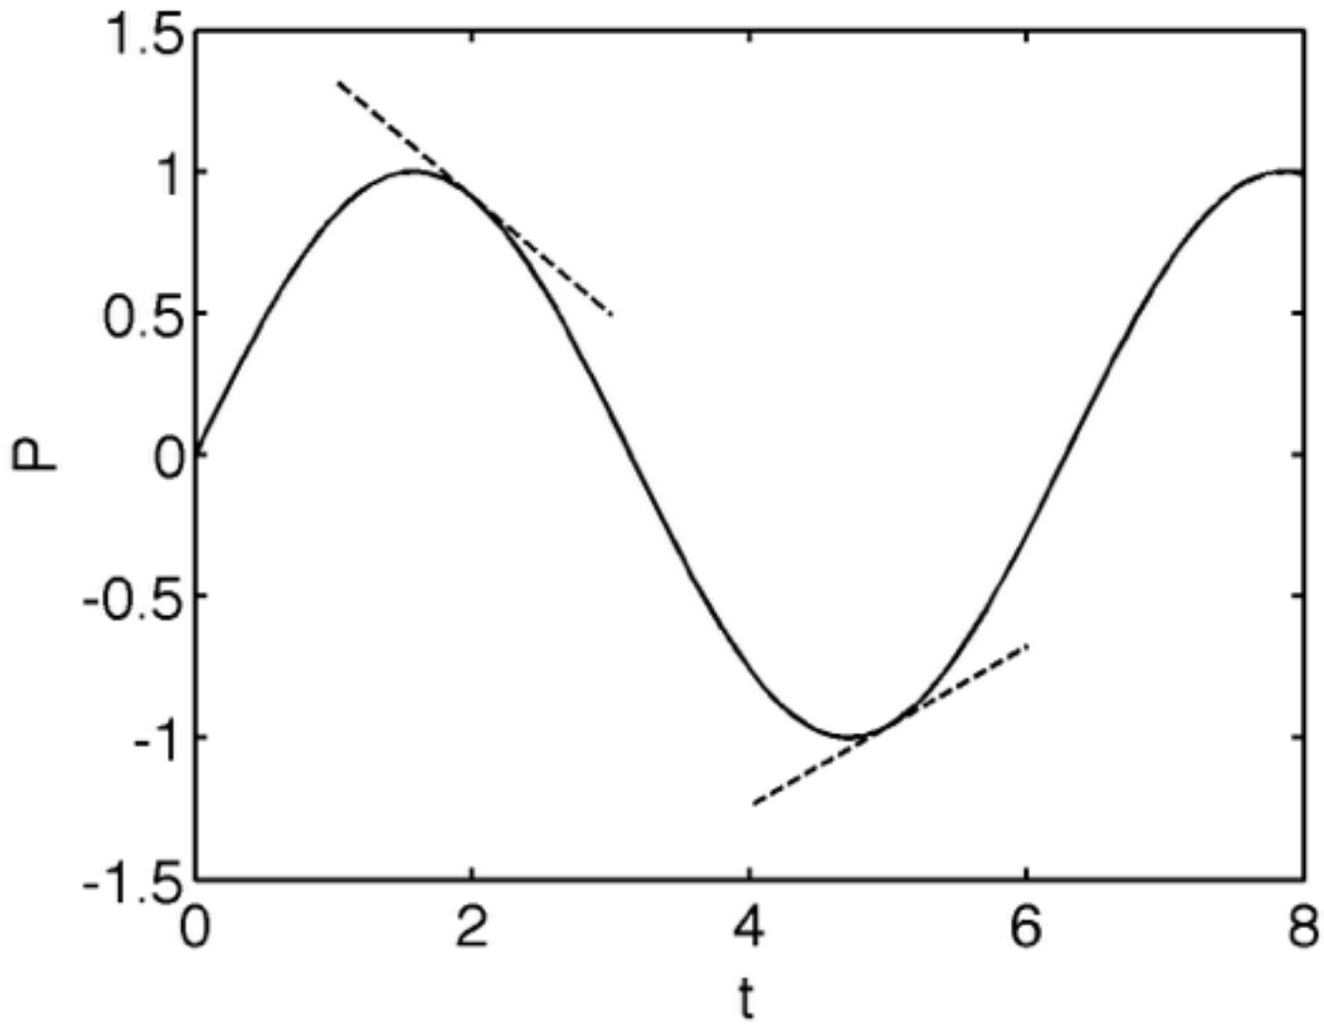

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

Ordinary Differential Equations

- Ordinary Differential Equations define gradients instead of values.

$$\frac{dP}{dt}$$

The rate of change of P with respect to t

Shorthand

$$\dot{P}$$

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

Ordinary Differential Equations

University of Glasgow

Typically:

$$\frac{dP}{dt} = \dot{P} = f(P, t, A, Z, \ldots)$$

[the rate of change of P w.r.t t is a function of P, t, and other things]

$$\frac{dP_t}{dt} = \dot{P}_t = f(P_t, t, A_t, Z_t, \ldots)$$

[the rate of change of P (at time t) w.r.t t is a function of P (at time t), t, and other things (at time t)]

$$\mathrm{A} + \mathrm{B} \xrightarrow{k_1} \mathrm{A} + \mathrm{B} + \mathrm{B}$$

the amount that B increases at some time t depends on how much A is in the system at time t…

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

Roadmap

- Given a set of reactions
- we can define a set of differential equations
- which we can solve (on a computer)
- to tell us how much of each thing there is at any time t.
- There will be one equation for each species (e.g. protein) of interest.

What **is** P?

- Our variables (e.g. P,A,B etc) are **concentrations** of things we want to model.
- Mathematically we treat them as continuous values - each one can take **any** value.
- Q: implications of this? Assumptions?
- [they should also be +ve]

Example: protein decay

- P = concentration of protein.

$$\dot{P}_t = -\lambda P_t$$

Q: where does that come from?

Q: what else do we need to know?

Q: what does changing lambda do?

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

Protein decay

This particular differential equation can be solved **analytically,** the result, for an **initial concentration** of 100 can be seen below:

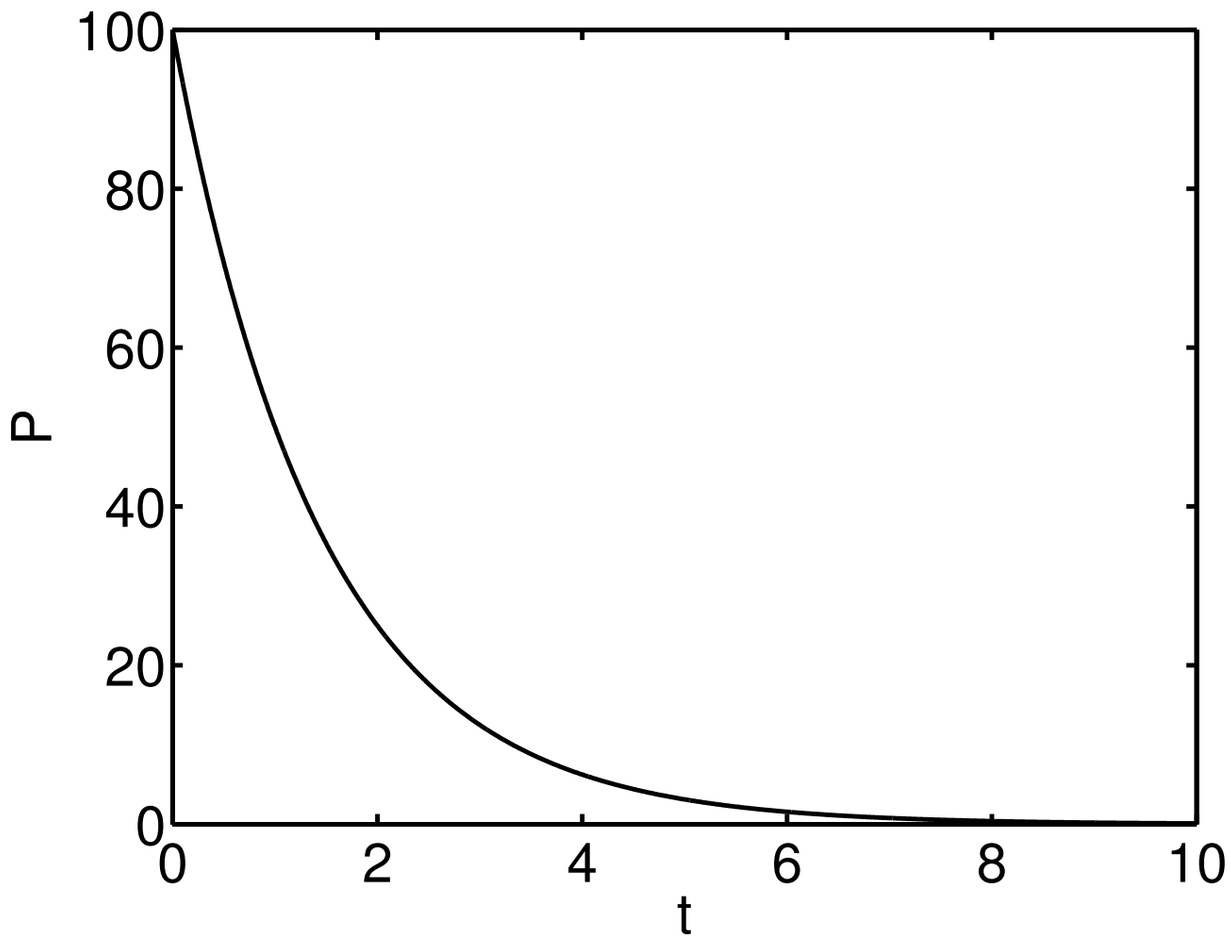

$$P_t = P_0 \exp(-\lambda t).$$

No useful models can be solved analytically. So we have to resort to numerically solving on a computer.

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

ODE Solvers

- Many packages exist.
- Basic mathematical idea is quite simple.
- I've provided a [very] simple solver written in Java.
  - Do not use for anything 'real'!!

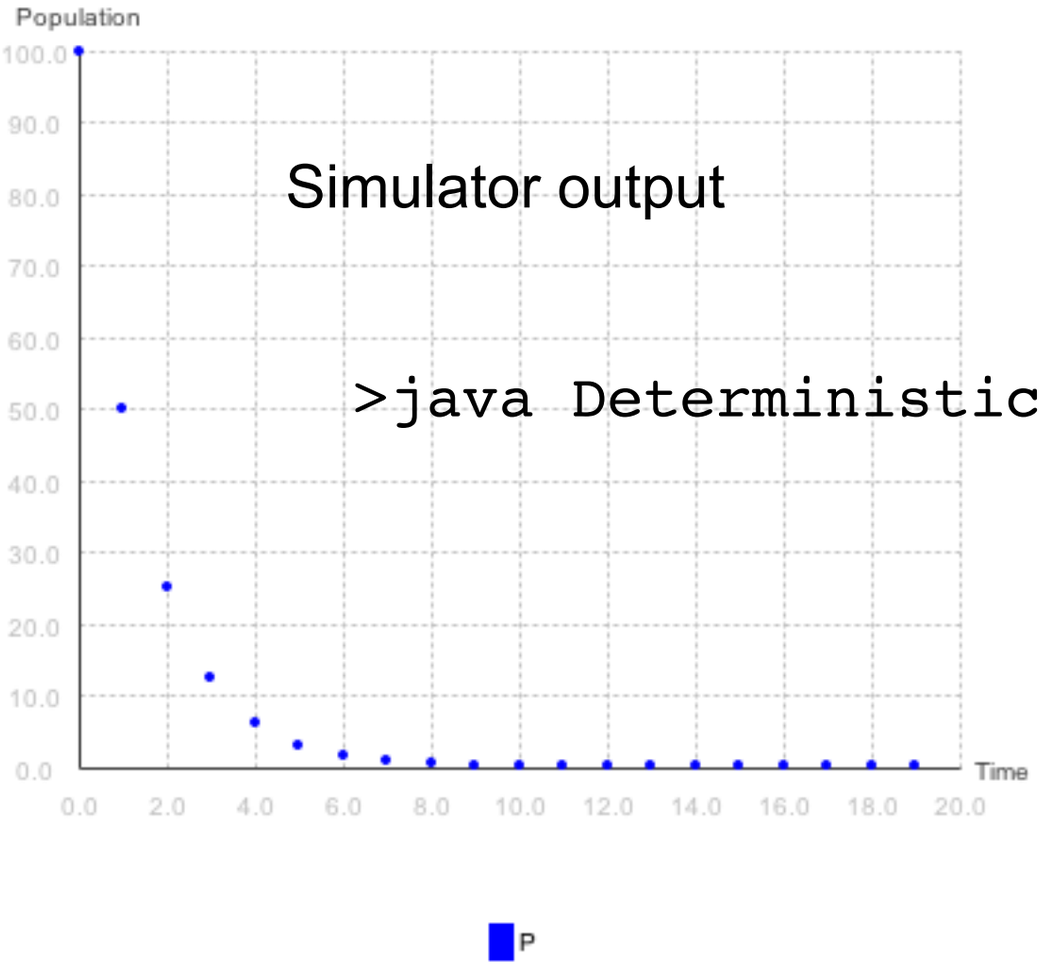

Simulator output

```
>java DeterministicSolver decay 20 1e-3 1 out.txt
```

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

More complex example

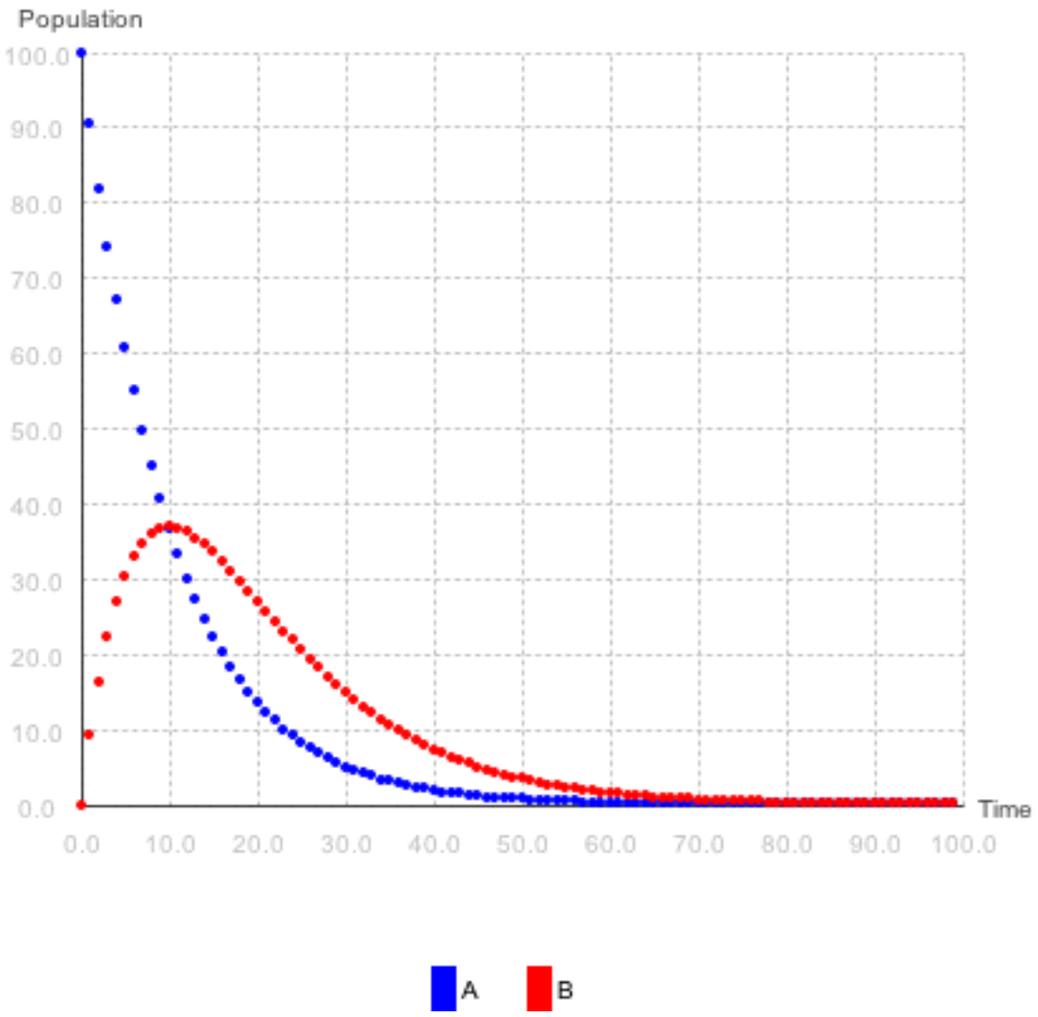

$$\mathrm{A} \xrightarrow{k_1} \mathrm{B}$$

$$\mathrm{B} \xrightarrow{k_2} 0$$

Q: can you predict what will happen if you change k1 or k2?

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

From reactions to equations

$$\mathrm{A} + \mathrm{B} \xrightarrow{k_1} \mathrm{A} + \mathrm{B} + \mathrm{B}$$

$$\mathrm{A} \xrightarrow{k_2} 0$$

Q: how do the reactions change the species?

Q: how fast do the reactions happen?

Kinetic models define how fast reactions happen.
e.g. mass action, Michaelis-Menten

A model inside a model!

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

Mass Action

$$\mathrm{A} + \mathrm{B} \xrightarrow{k_1} \mathrm{A} + \mathrm{B} + \mathrm{B}$$

$$\mathrm{A} \xrightarrow{k_2} 0$$

Reaction 1 leaves A unchanged. Reaction 2 **decreases** A.

Reaction 1 **increases** B. Reaction 2 leaves B unchanged.

$$\begin{aligned} \dot{A}_t &= -k_2 A_t \\ \dot{B}_t &= k_1 A_t B_t \end{aligned}$$

Mass Action: rate of reaction is proportional to the product of the 'masses' of the reactants.

Q: what are the main assumptions under-pinning mass action?

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

Mass Action Example

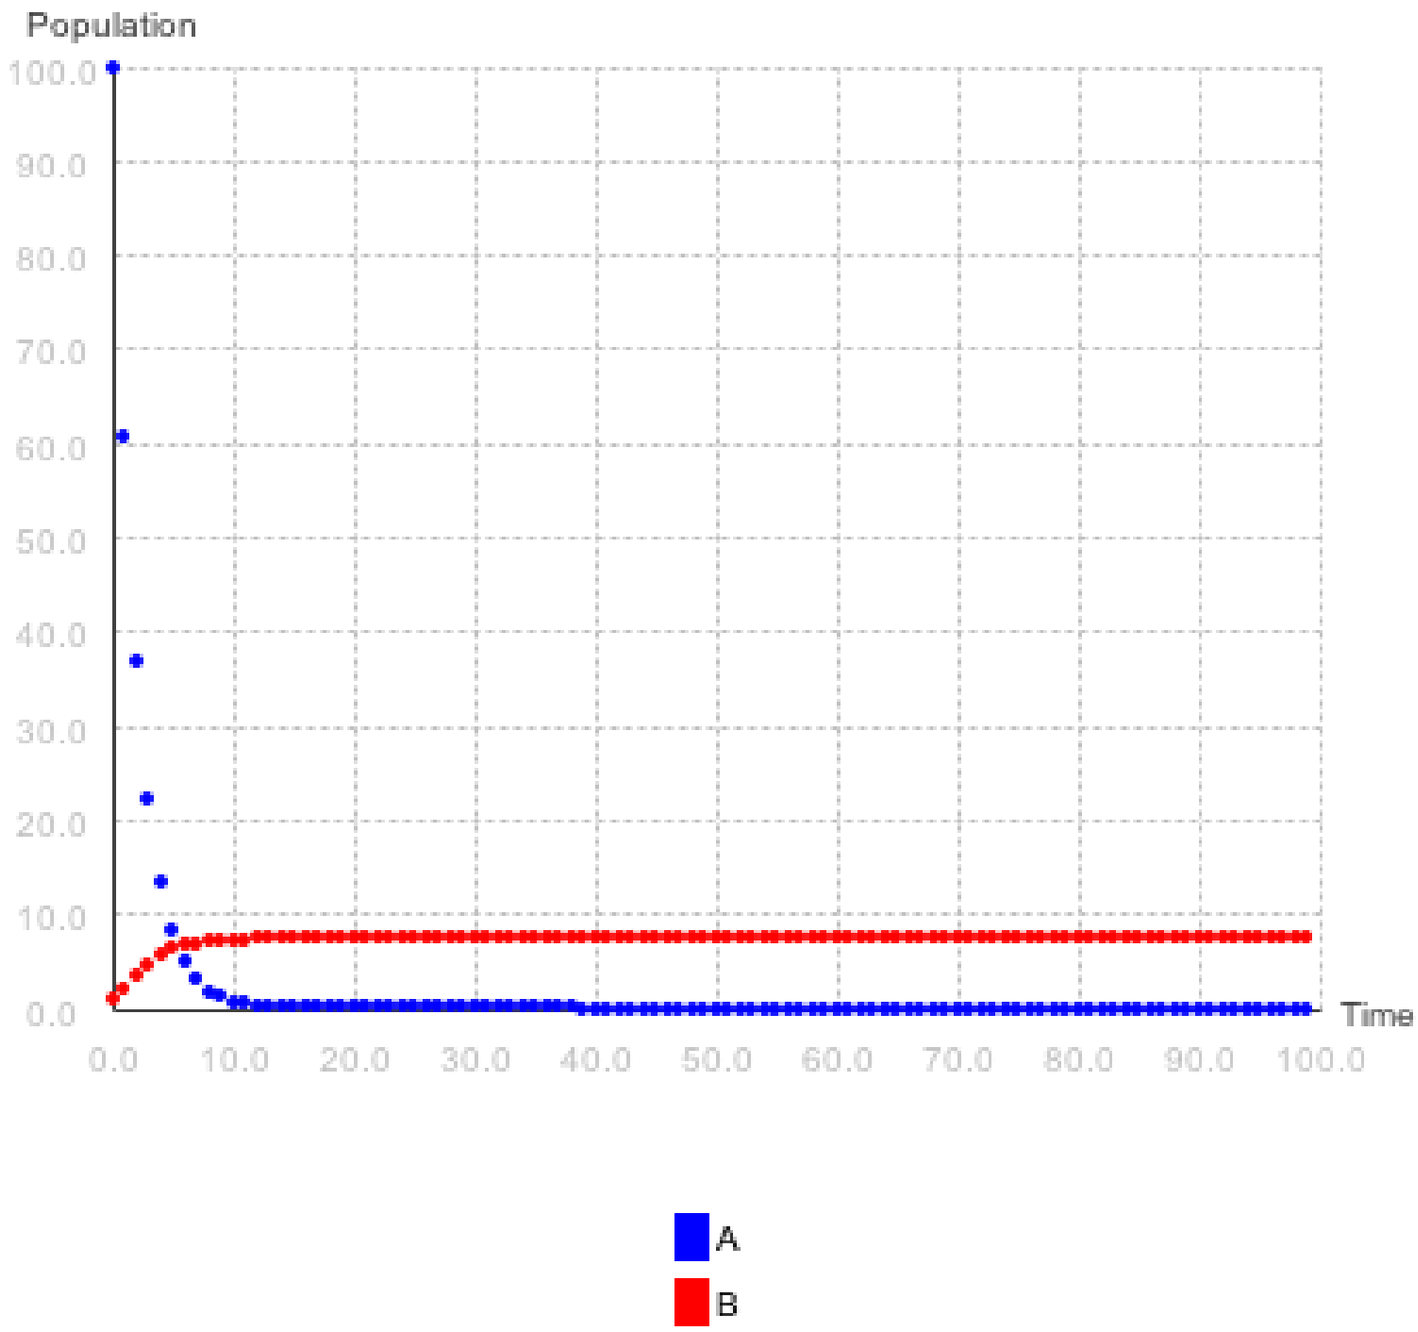

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

Class exercise

- What are the mass action ODEs for the following system of reactions:

$$\mathrm{A} + \mathrm{B} \xrightarrow{k_1} \mathrm{A} + \mathrm{B} + \mathrm{B}$$

$$\mathrm{B} + \mathrm{C} \xrightarrow{k_2} \mathrm{C} + \mathrm{C}$$

$$\mathrm{C} \xrightarrow{k_3} 0$$

Tip: one equation per species, not per reaction…

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

Run it

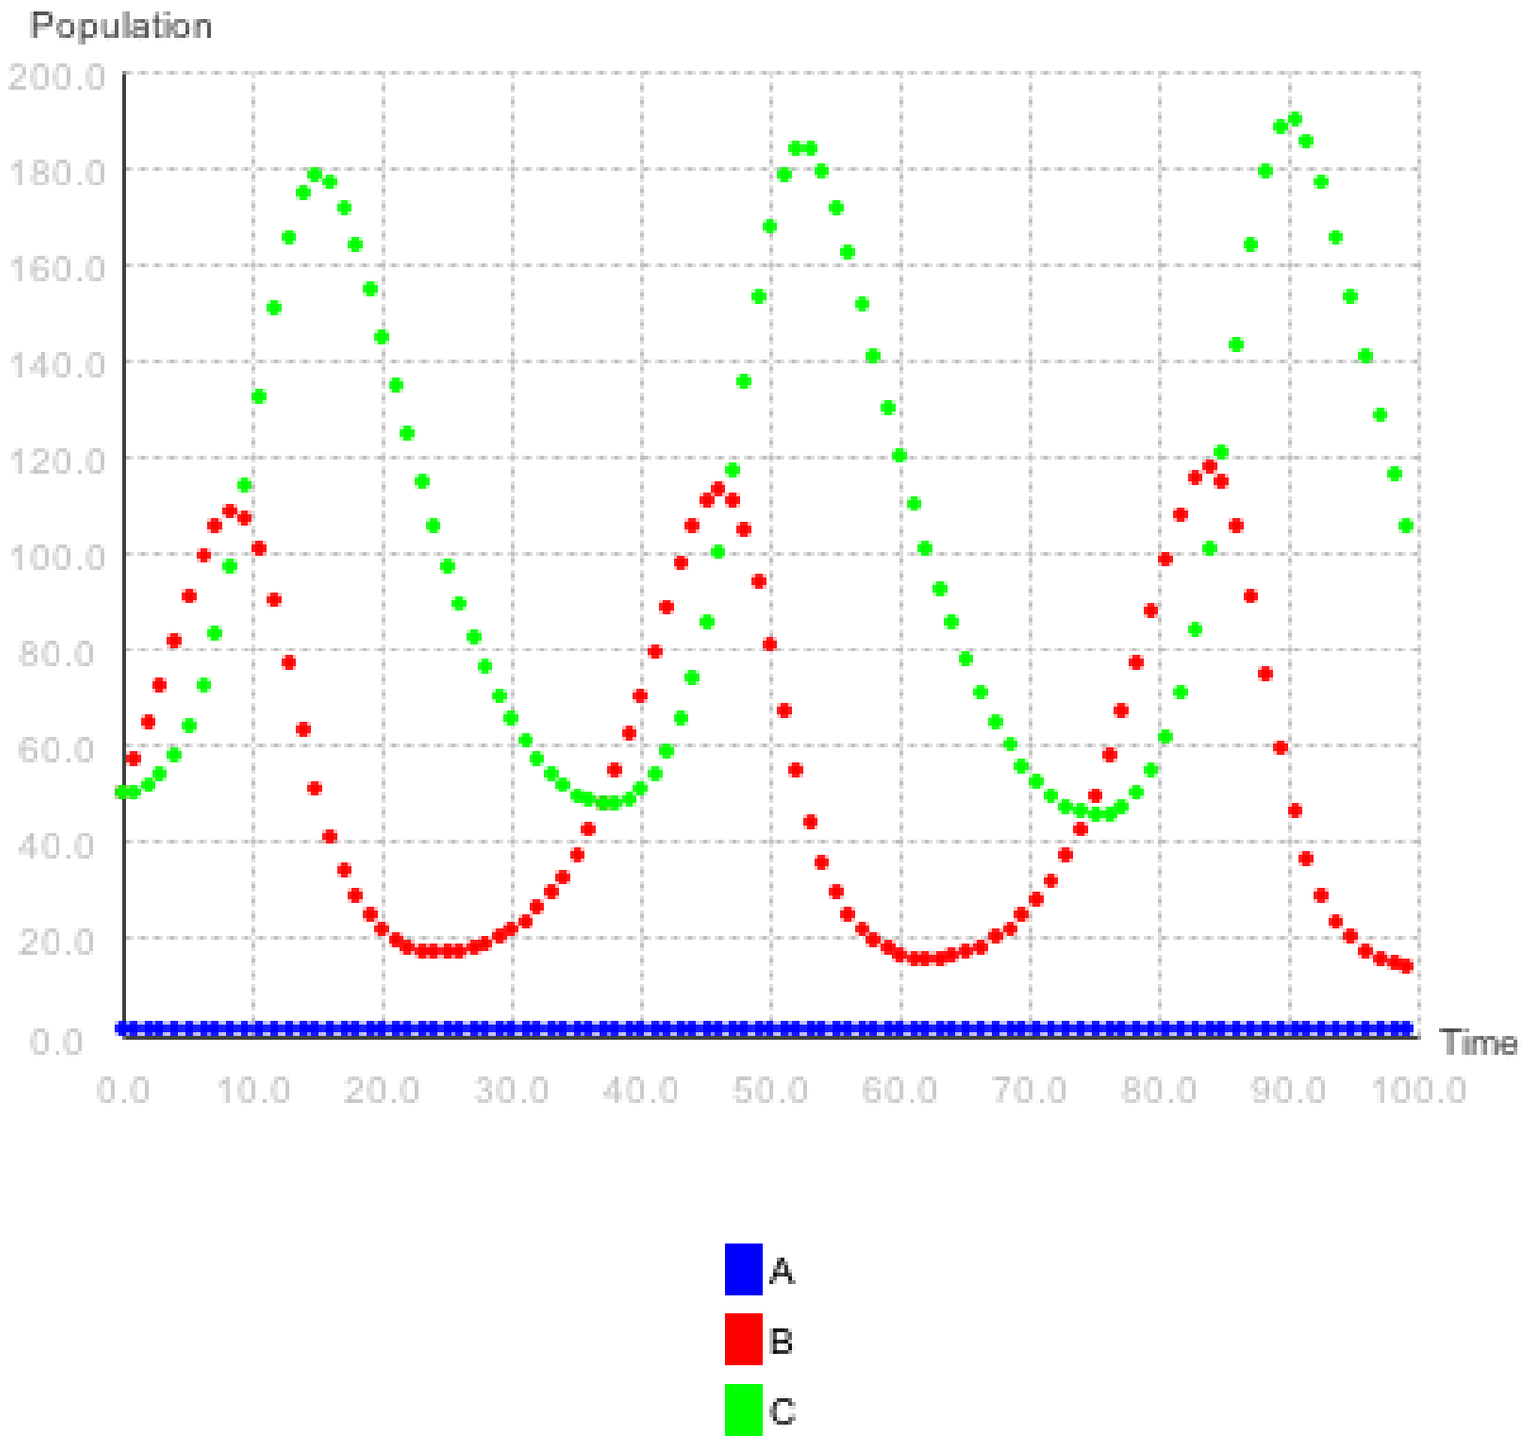

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

Other kinetic models

$$\mathrm{A} \longrightarrow \mathrm{A} + \mathrm{B}$$

Mass Action

$$\dot{B}_t = kA_t$$

Michaelis-Menten

$$\dot{B}_t = \frac{k_1 A_t}{k_2 + A_t}$$

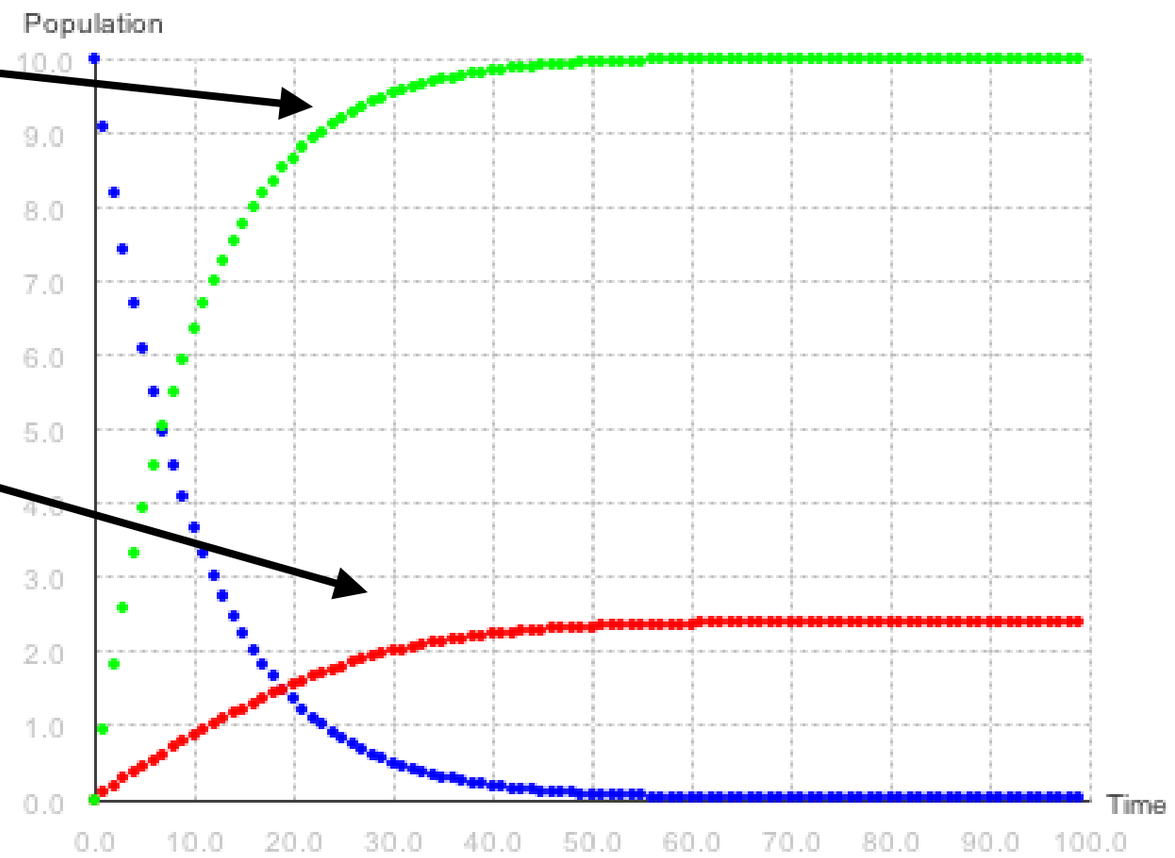

Try other parameter values in the lab

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

Other kinetic models II

- Michaelis-Menten:
- Q: can you work out what real phenomena the parameters correspond to?

Summary

- We have (briefly) covered some of the ideas behind deterministic modelling with ordinary differential equations.
- In the notes, you'll see some exercises to try in the lab (now).
- After the lab, we will look over these exercises and possible move onto discussing **data** and **parameter/model inference**

Data

Note: what follows is another whole lecture…we will not cover it all!

simon.rogers@glasgow.ac.uk www.dcs.gla.ac.uk/~srogers @sdrogers

To simulate, we need a model

- All simulations start with:
  - A model
  - Some parameters
  - Some initial conditions:

$$
\begin{aligned}
A + B &\rightarrow 2B + A \\
B + C &\rightarrow 2C \\
C &\rightarrow \emptyset
\end{aligned}
$$

$$
A|_{t=0} = 1 \qquad B|_{t=0} = 50 \qquad C|_{t=0} = 50
$$

$$
k_1 = 0.25 \qquad k_2 = 0.0025 \qquad k_3 = 0.125
$$

- What if we start with:
  - A model and no parameters?
  - No model?

A model and no parameters

- Given a system, we can hypothesise a model.
  - e.g. I can hypothesise that proteins decay without necessarily knowing how fast this will happen (encoded as a parameter).
- Without parameters we can't simulate.
- But we can measure data.
- It would be useful if we could 'learn' parameters from measured data.

A model and no parameters

- Recall our model of protein decay and the initial condition:

$$\dot{P} = -\lambda P \qquad P_0 = 1$$

- We measure some data in the lab:

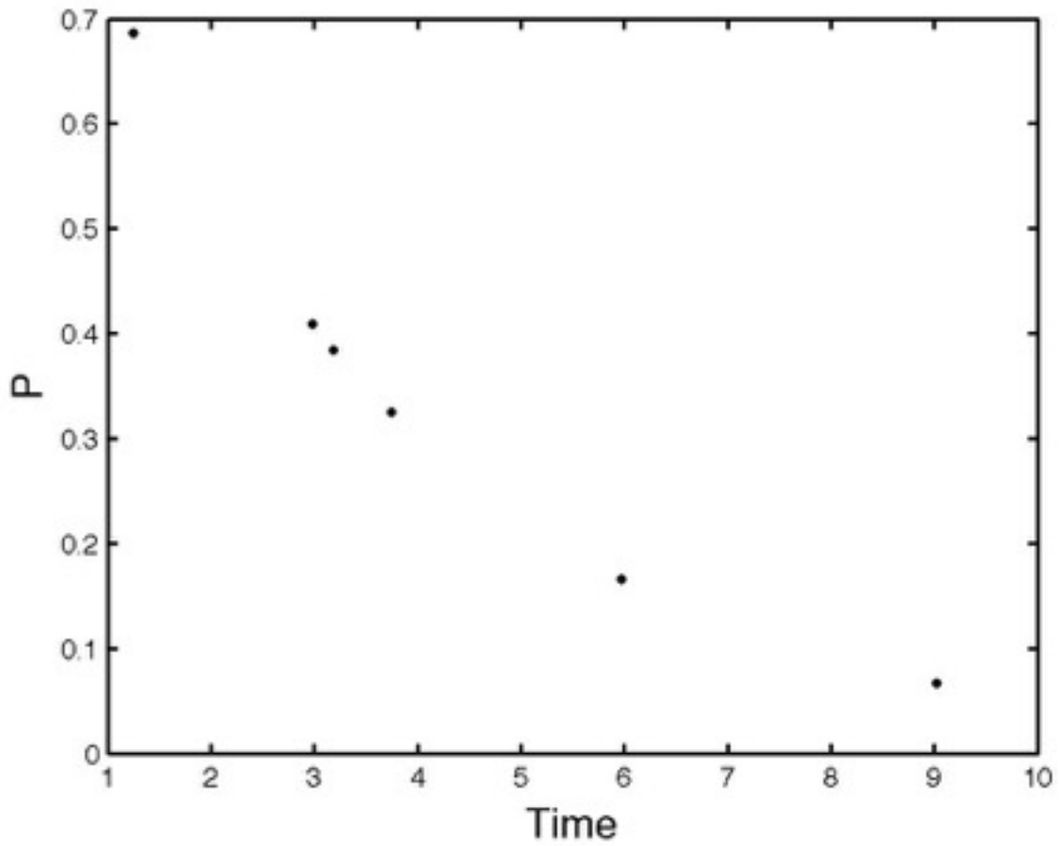

What is λ ?

Points and distributions

$$\text{What is } \lambda \text{ ?}$$

- Two broad strategies:
  - Find the ‘best’ value - *point estimates*

$$\lambda = 0.5$$

  - Find a distribution - *Bayesian inference*

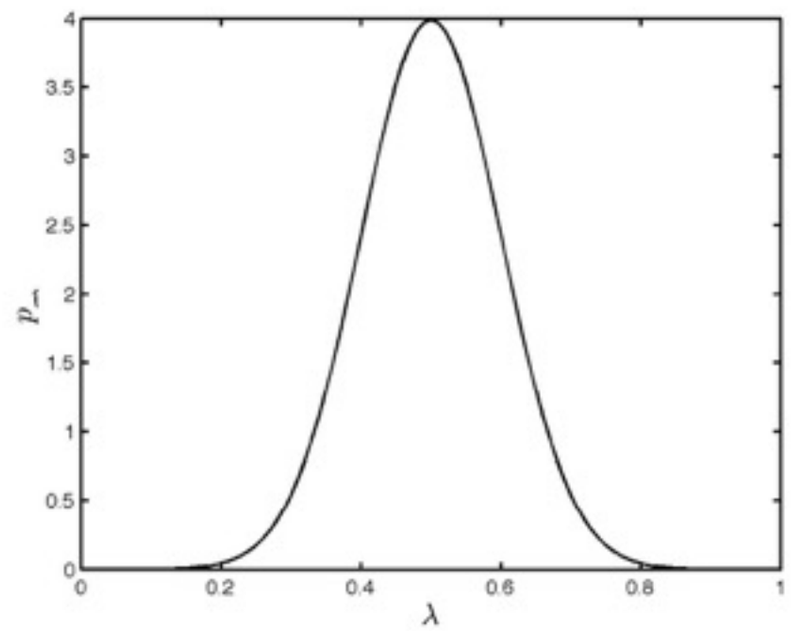

Point estimates

$$\text{What is } \lambda \text{ ?}$$

- Finding the ‘best’ value:
  - What does ‘best’ mean?

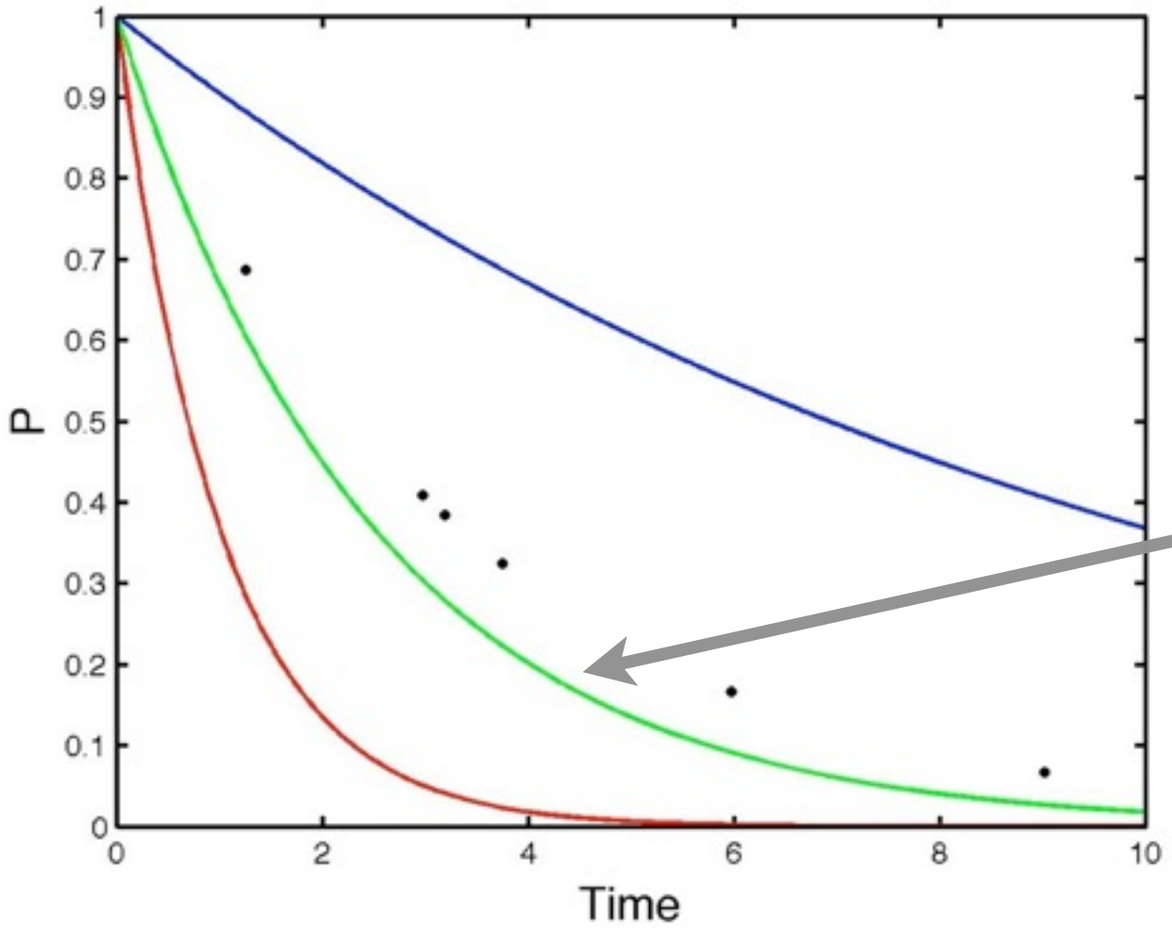

Each value produces a curve.
Which curve is best?

This curve best represents the experimental evidence.

Measures of ‘bestness’

- To automate the process of finding the best curve, we need a formal definition of ‘bestness’.
  - Non-statistical: e.g. *squared error, absolute error* (lower is better)
  - Statistical: e.g. *likelihood* (higher is better)

Squared error

For a particular λ

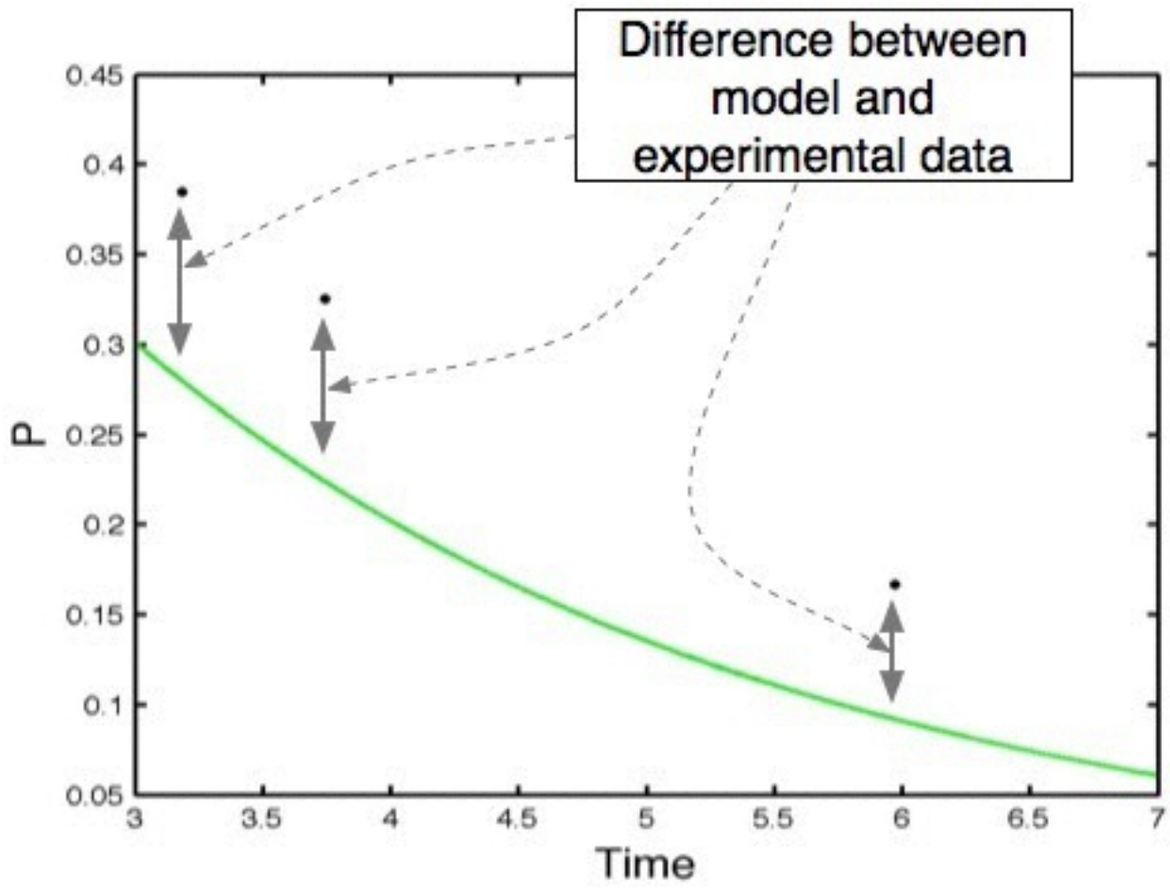

- Compute the *squared* difference between the model and experimental data at the times when the data was measured.
  - Why *squared?*
- Add the values for all data points together. $= L(\lambda)$

Squared error

- Mathematically speaking:

$$\arg\min_{\lambda} \; L(\lambda)$$

- Find the value that minimises L
- How?
  - Guess lots of times and keep the best.

Squared error

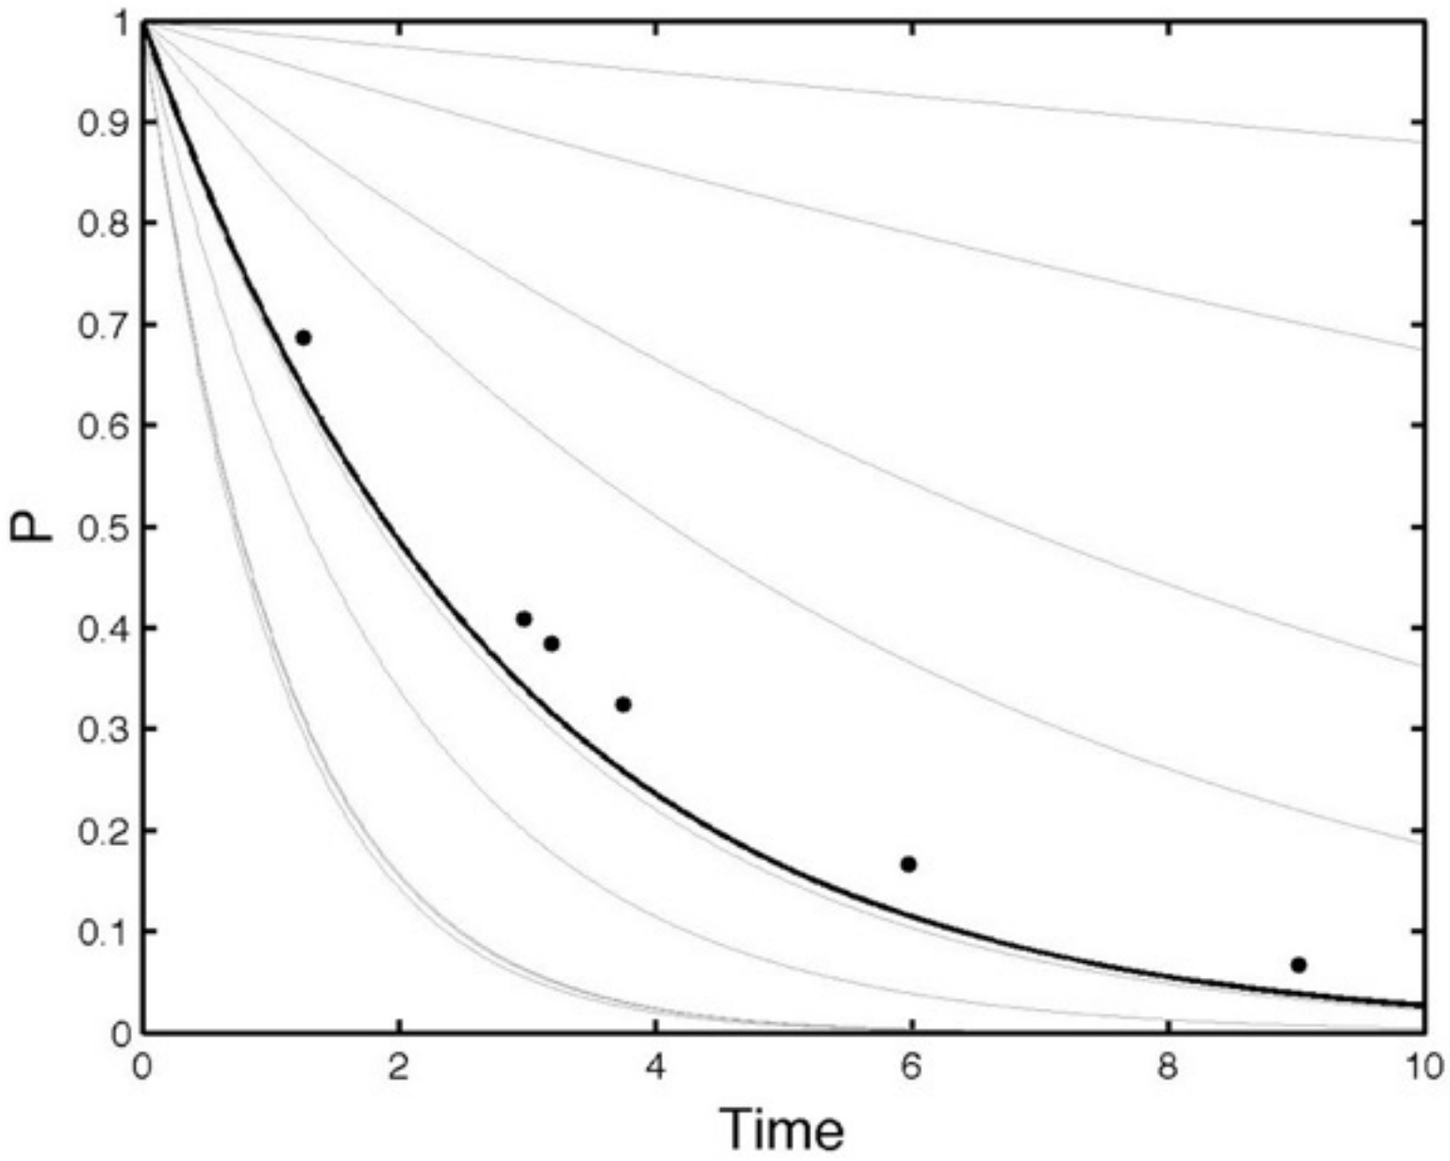

- First 10 guesses (best is highlighted)

Squared error

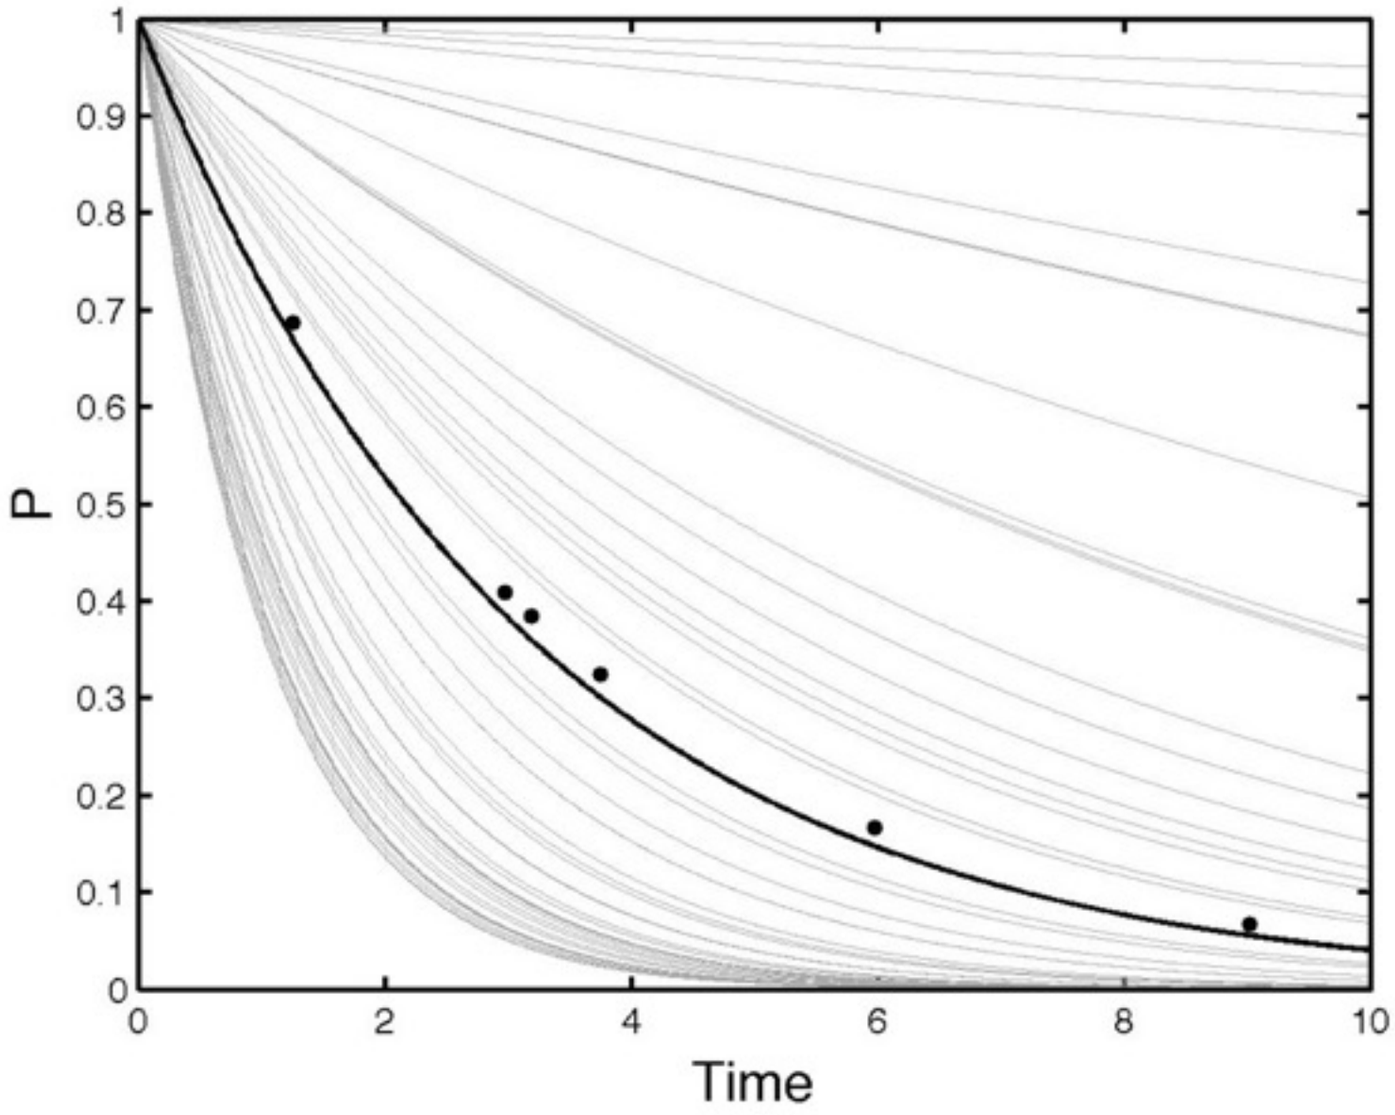

- First 50 guesses (best is highlighted)

Squared error

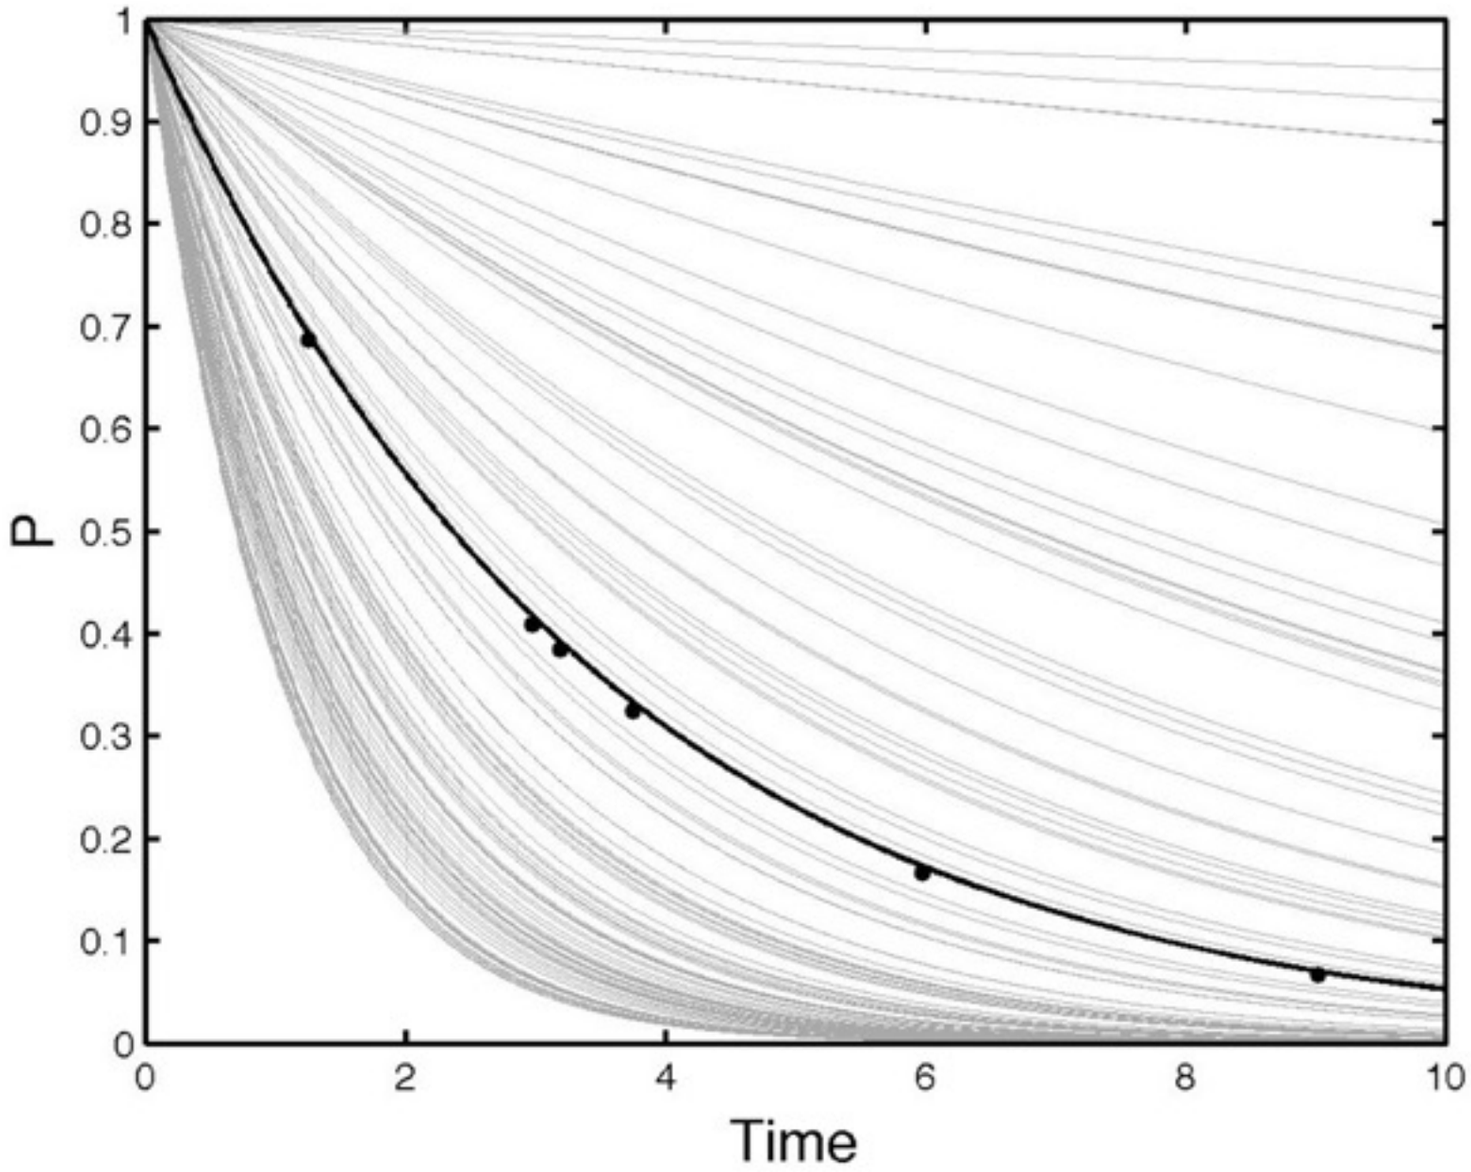

- First 100 guesses (best is highlighted)

$$\lambda = 0.2933$$

Alternatively

- Guess a value, compute L.
- Change it a bit, compute new L.
- If new L is better, keep new value. Otherwise throw it away.

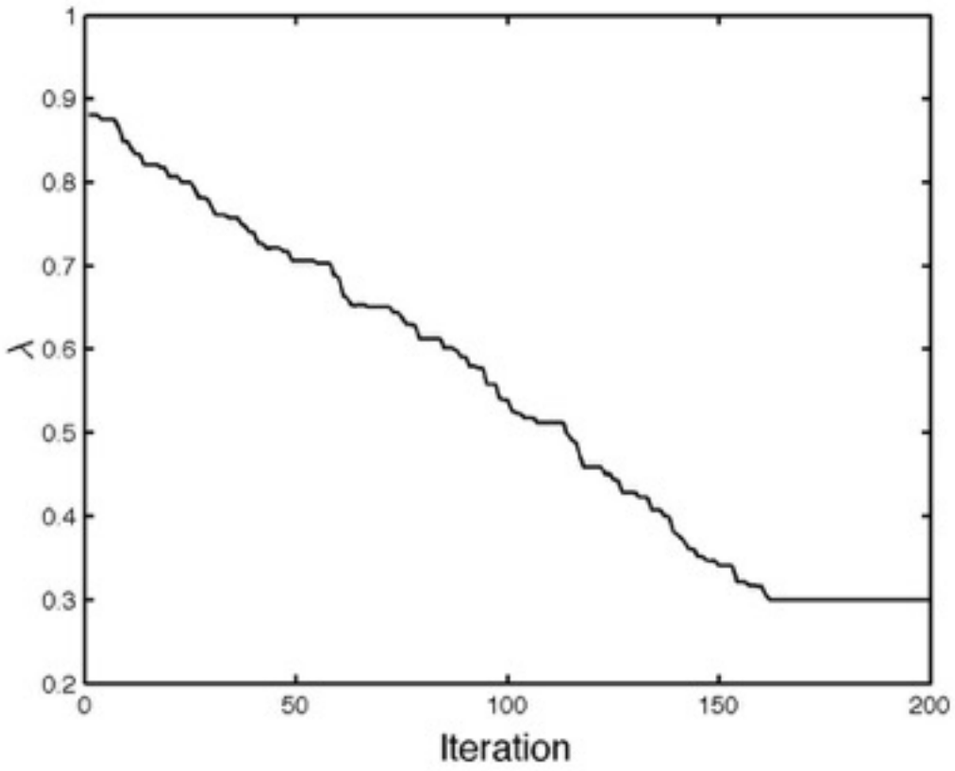

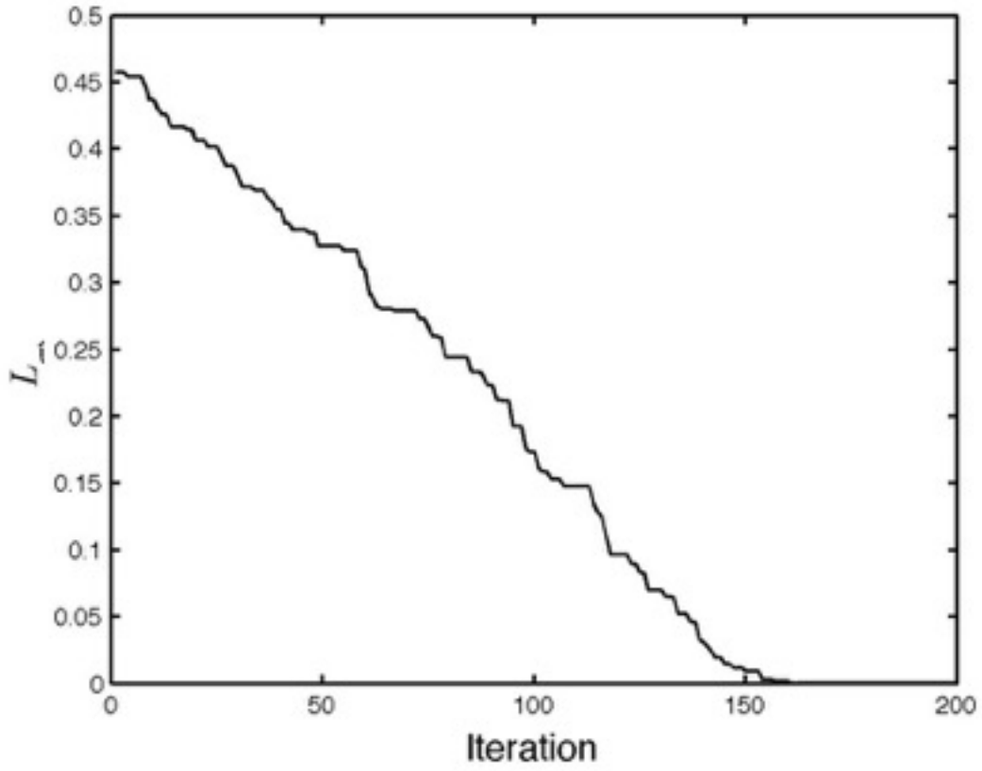

$$\lambda = 0.2997$$

Help is at hand

- Many algorithms exist for optimisation.
- Most are smarter than these two.
  - Different ways of updating the parameters.
  - Necessary when problem gets more *complex*.
  - e.g. Finding best values of many parameters at once.
- Using simbiology in Matlab gives: $\lambda = 0.3$
- This is the ‘correct’ value.

Complexity

- For any optimisation routine:
  - Each evaluation of L requires one model simulation.
  - For models with many parameters, we may require millions of evaluations!
  - *Very* slow.
  - Active research area.

Is there a correct value?

- Imagine doing the biological experiment twice

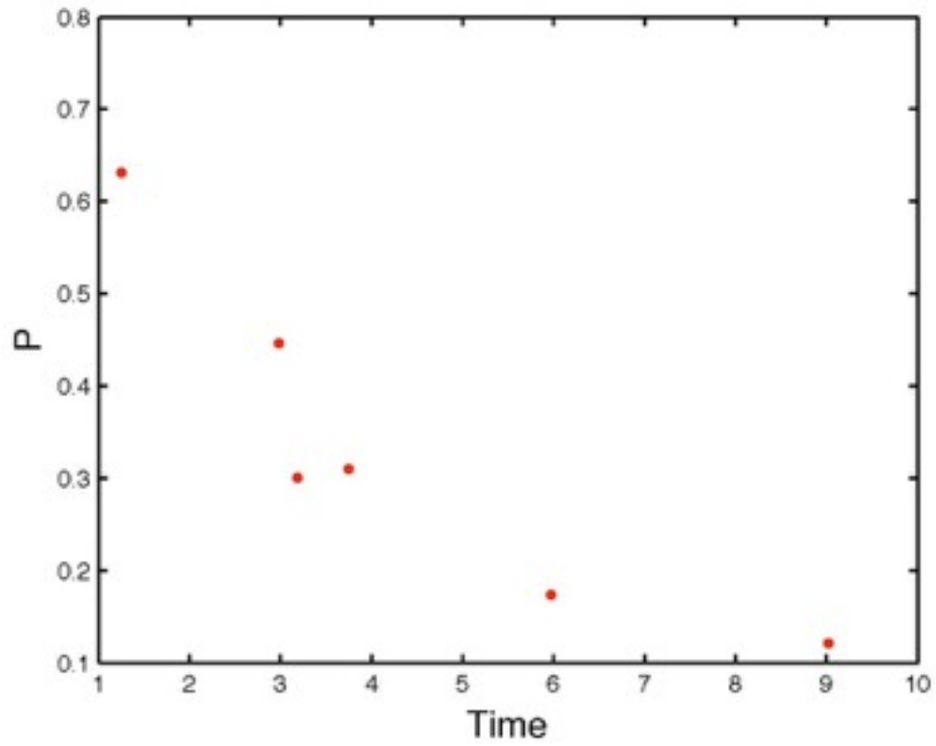

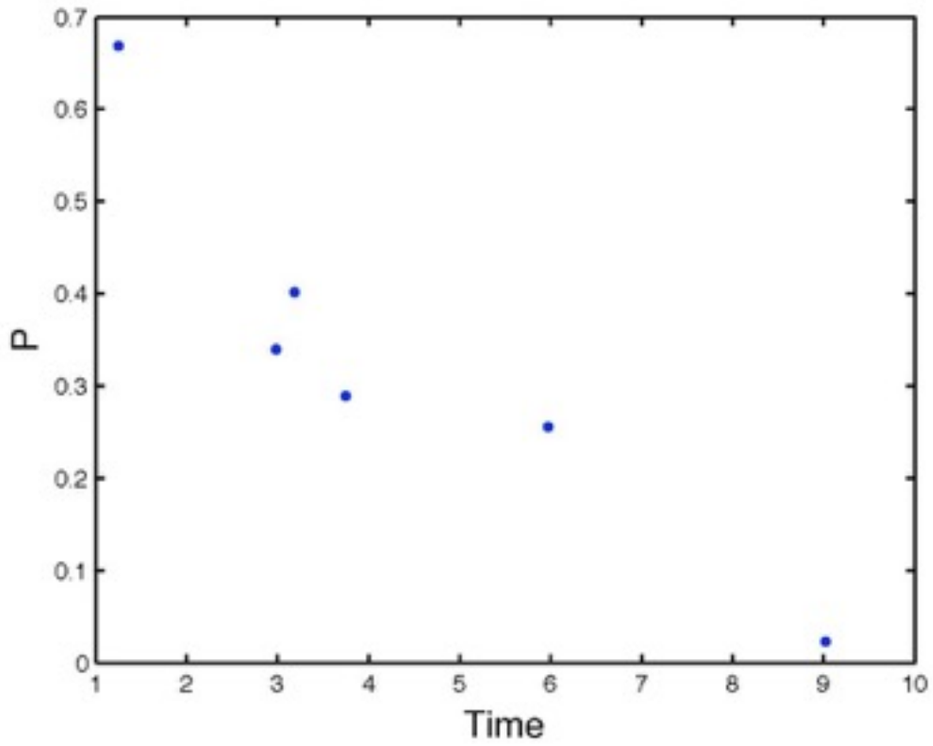

- Not the same - why?
  - Stochasticity, contamination, measurement error

$$\text{What is } \lambda \text{ ?}$$

Distributions, not points

- Bayesian parameter estimation:
  - Some (me included) would argue that point estimates are misleading.
  - They represent a degree of certainty that is unjustified:
    - Repeat the experiment, value changes.
- Alternative: Produce a *distribution* over values.

Do the experiment lots of times...

- Repeat the experiment many times.
- Find the best value for each experiment.
- Histogram the results:

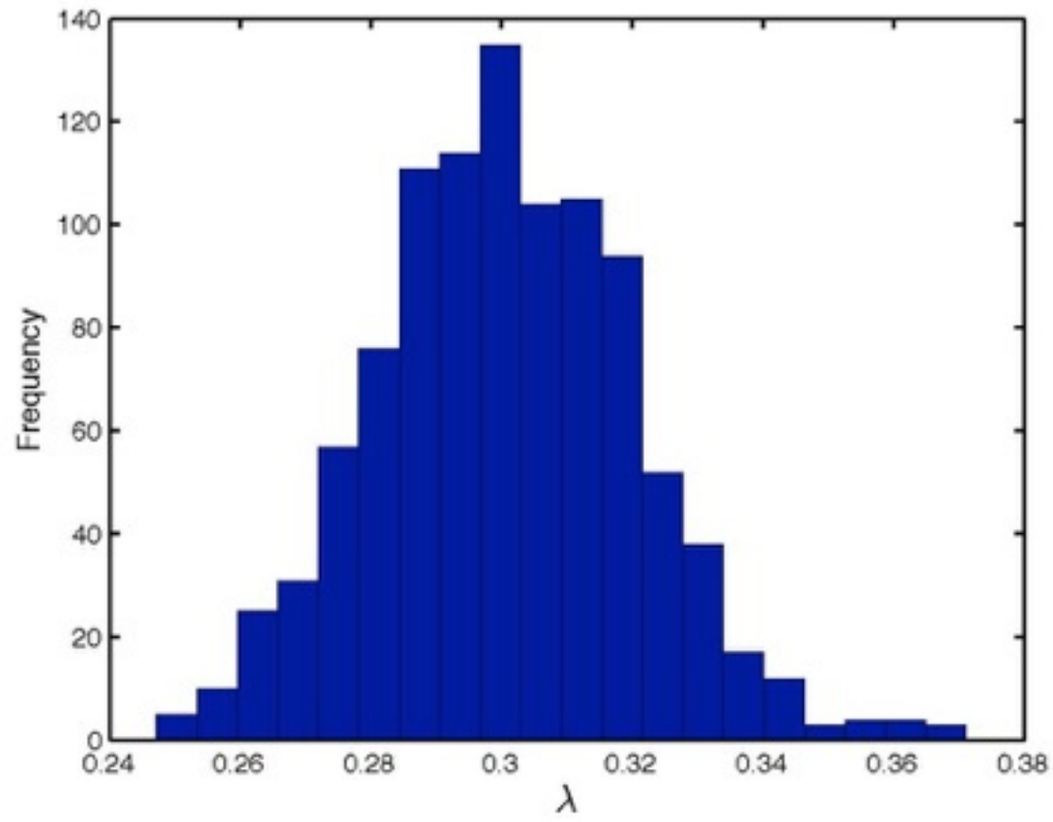

- More informative:
  - Tells us something about the uncertainty caused by things we can't measure

Bayesian Parameter Estimation

- We can get this *without* repeating the experiment.
- Instead, we make *assumptions* about the things we can't measure.
- i.e. all things we can't measure, added together look like random variables of some particular type.
- e.g. deviations from model (noise) are *Gaussian*

Bayesian Parameter Estimation

- Unfortunately (for you) it's hard - not supported by SimBiology.
- Fortunately (for me) it's hard - requires research.
- Other benefits:
  - It's possible to incorporate *prior knowledge*.
  - It's possible to rigorously compare how well different models explain data: *comparing hypotheses* **(science!)**
- Negatives?
  - Even more evaluations of L: even *slower!*
  - Normally requires expert tuning.

But ODE models are bad....

- What about learning parameters in stochastic models?
  - What kind of data would we need?
    - Single cell, molecule counts - very expensive.
    - Microarray would be no good.
  - Discrete data causes problems in standard algorithms (non-smooth).
  - Even *slower* still.
  - Active research area.

No model and no parameters

- Comparing hypotheses:
  - Given some data, can we choose:

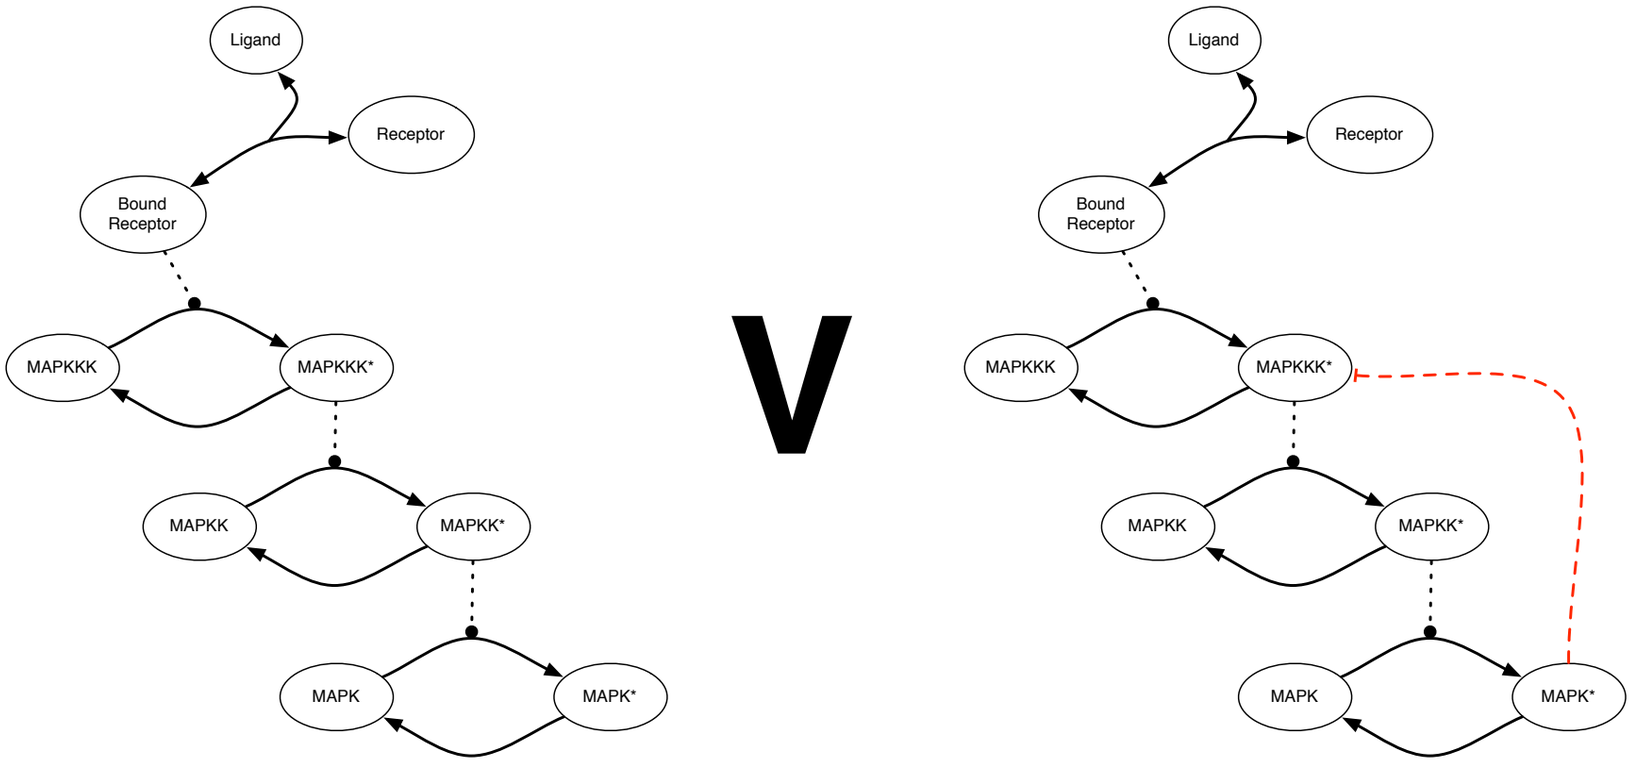

Yes:

BIOINFORMATICS ORIGINAL PAPER

Vol. 24 no. 6 2008, pages 833–839
doi:10.1093/bioinformatics/btm607

Systems biology

Bayesian ranking of biochemical system models

Vladislav Vyshemirsky* and Mark A. Girolami
Department of Computing Science, University of Glasgow, Glasgow, G12 8QQ, UK

No model and no parameters

- (Many) people have also tried to rely more heavily on data:
- Search for statistical (not necessarily biological) relationships between species:

BIOINFORMATICS **ORIGINAL PAPER** *Vol. 21 no. 14 2005, pages 3131–3137*
doi:10.1093/bioinformatics/bti487

Systems biology

A Bayesian regression approach to the inference of regulatory networks from gene expression data

Simon Rogers* and Mark Girolami

Department of Computing Science, Bioinformatics Research Centre, University of Glasgow, Glasgow, UK

Just data

Expand an existing model using data

Genome Biology 2006, 7:R25 **Open Access**

Method

Ranked prediction of p53 targets using hidden variable dynamic modeling

Martino Barenco*†, Daniela Tomescu*, Daniel Brewer*†, Robin Callard*†, Jaroslav Stark†‡ and Michael Hubank*†

Summary

- If we have a model and no parameters, we can find them with data.
- Point estimates:
  - Easy (to understand and do).
  - Too certain.
- Bayesian inference:
  - Hard(er) (to understand and do).
  - Reflects uncertainty.
  - Needs additional assumptions - noise.
- Can use data to compare models.
- Can use data to *build* models.