Lecture 3

Evidence for Policy in the face of Uncertainty?

Roman Frigg, LSE

Plan

A Primer on Models and Chaos

Part I – UKCP09

Part II – Model Error and Prediction

Plan

A Primer on Models and Chaos

Part I – UKCP09

→ The Myopia of Imperfect Climate Models: The Case of UKCP09, forthcoming in Philosophy of Science (December 2013), with David A. Stainforth and Leonard A. Smith.

Part II – Model Error and Prediction

→ Laplace's Demon and the Adventures of His Apprentices, forthcoming in Philosophy of Science (January 2014), with Seamus Bradley, Hailiang Du and Leonard A. Smith.

Available at www.romanfrigg.org

Part I: In Collaboration With

Dave Stainforth

CATS, LSE

Lenny Smith

CATS, LSE

Dept Stats, LSE

Part II: In Collaboration With

Hailiang Du

CATS, LSE

Dept Stats, LSE

Seamus Bradley

Dept Phil, LSE

Lenny Smith

CATS, LSE

Dept Stats, LSE

With Illustrations By

Fiorella Lavado

Independent Artist

www.fiorellalavado.com

A Primer on Models and Chaos

Dynamical system (X, ϕ_t, μ)

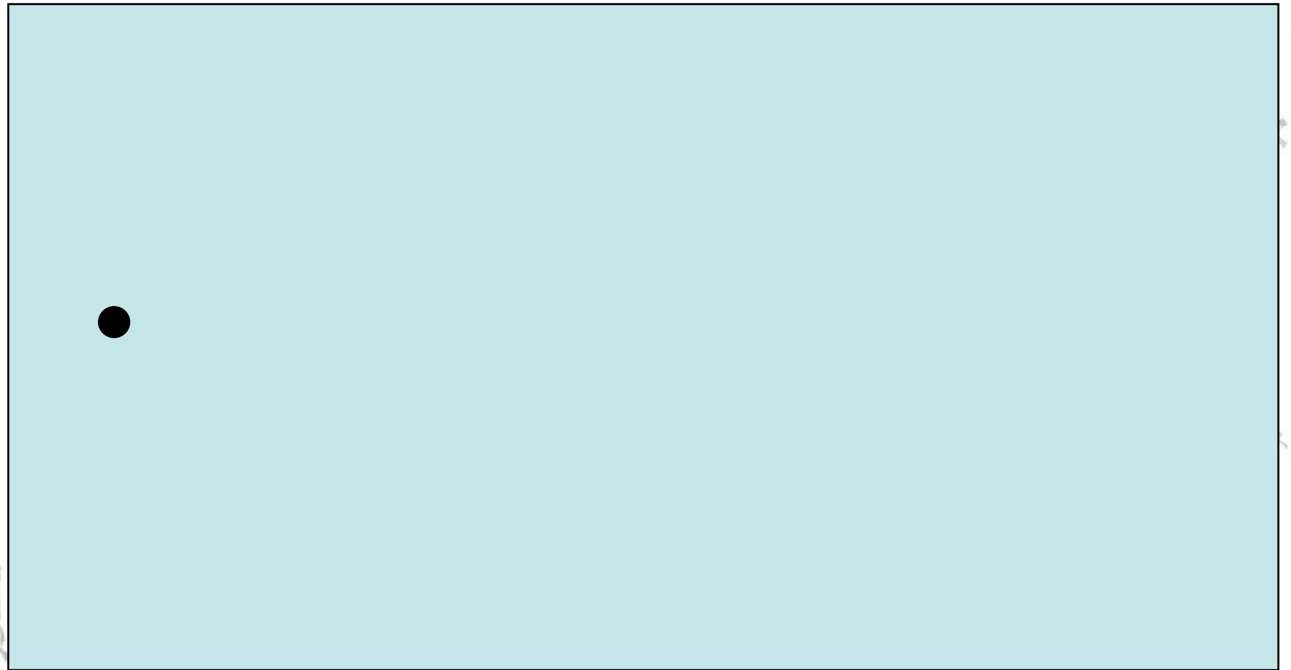

A Primer on Models and Chaos

Dynamical system (X, ϕ_t, μ)

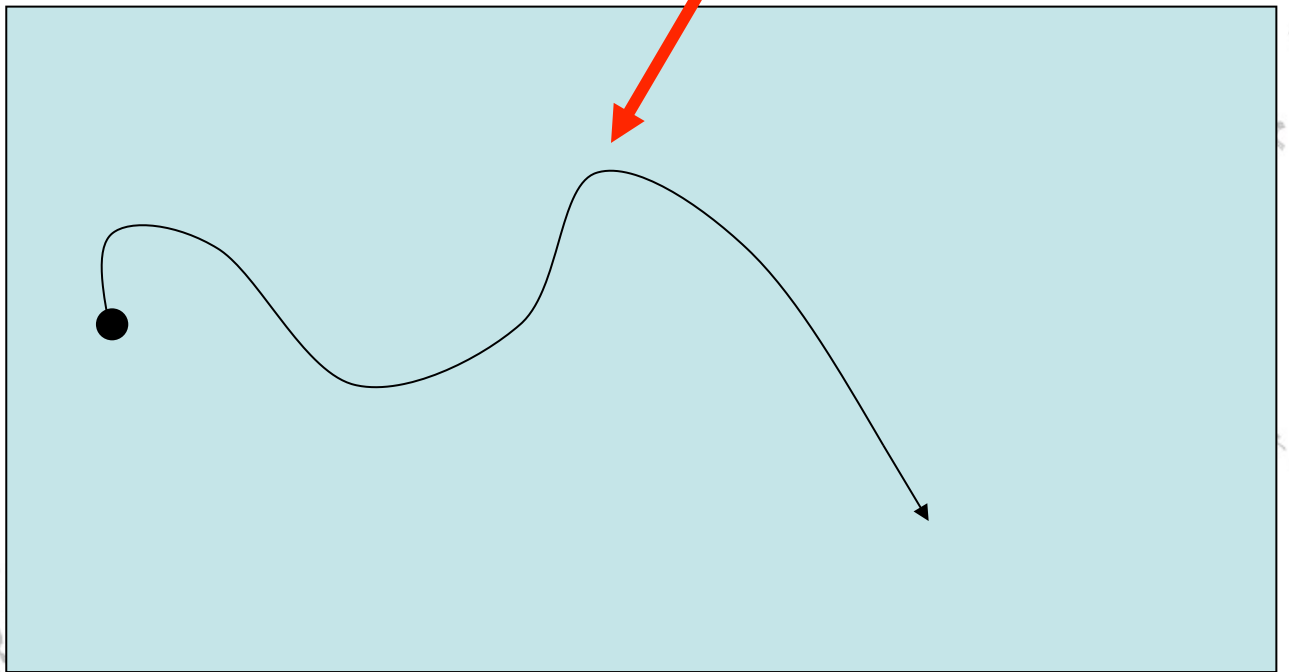

A Primer on Models and Chaos

Dynamical system (X, ϕ_t, μ)

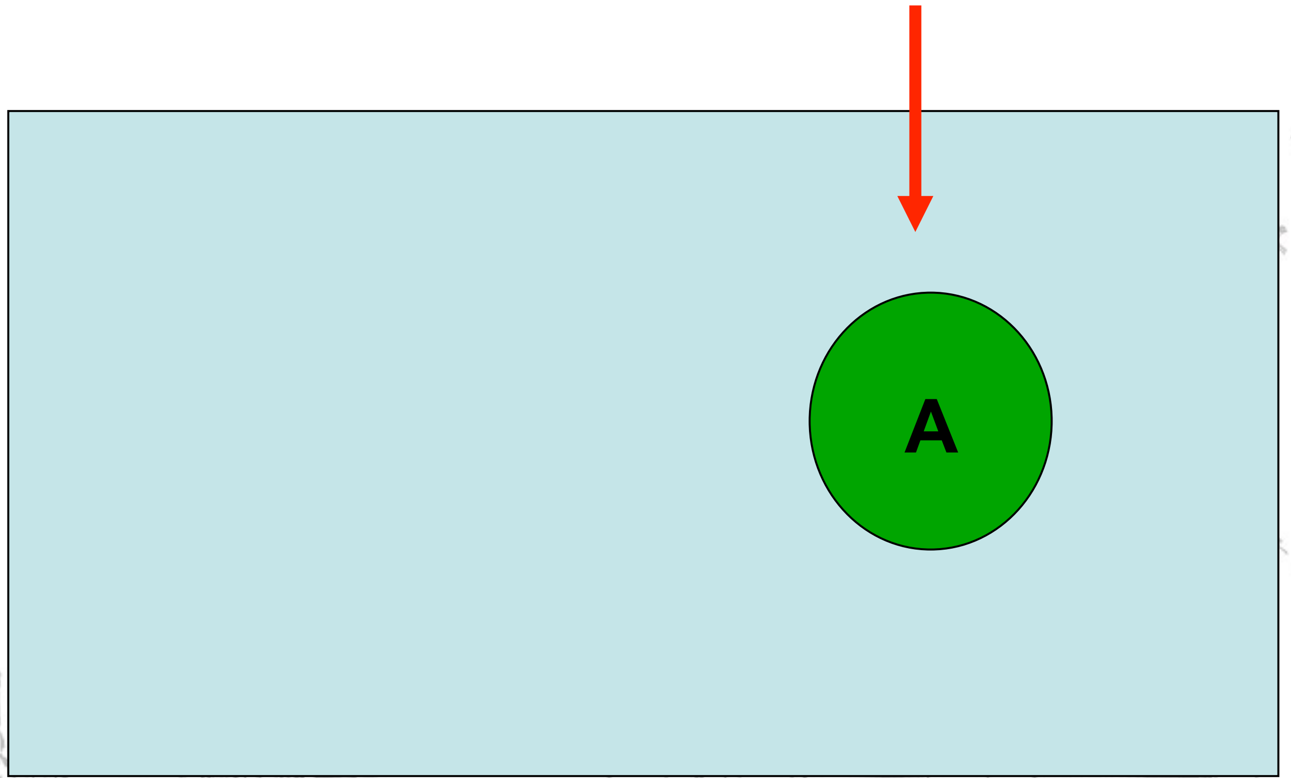

Simple example: stone falling from tower

Simple example: stone falling from tower

Simple example: stone falling from tower

height x

Simple example: stone falling from tower

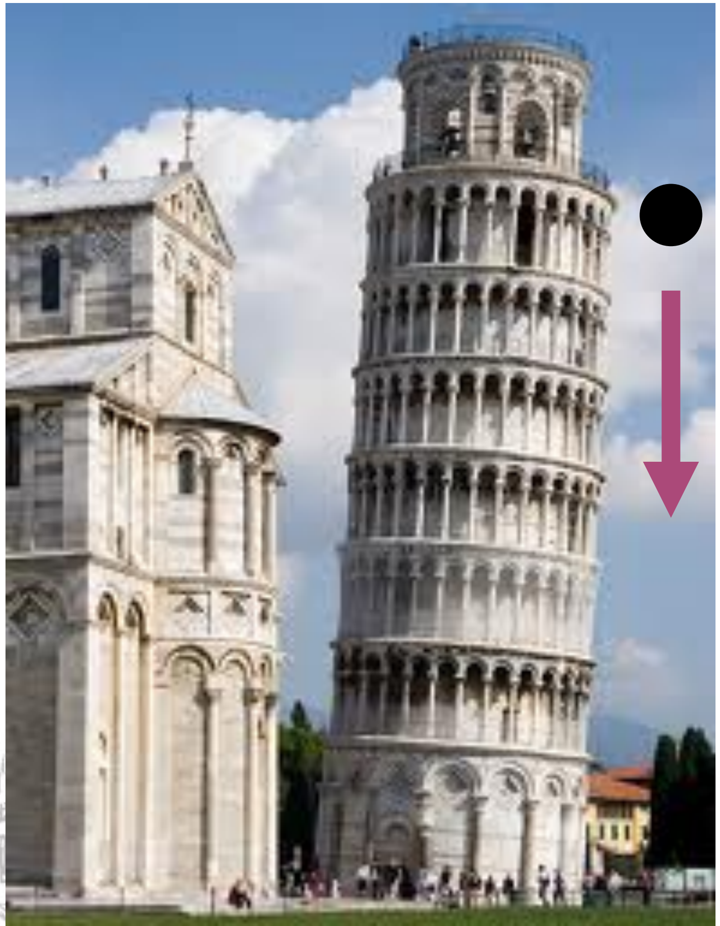

height x

speed / momentum p

Simple example: stone falling from tower

$$(\underline{X},\phi_t,\mu)$$

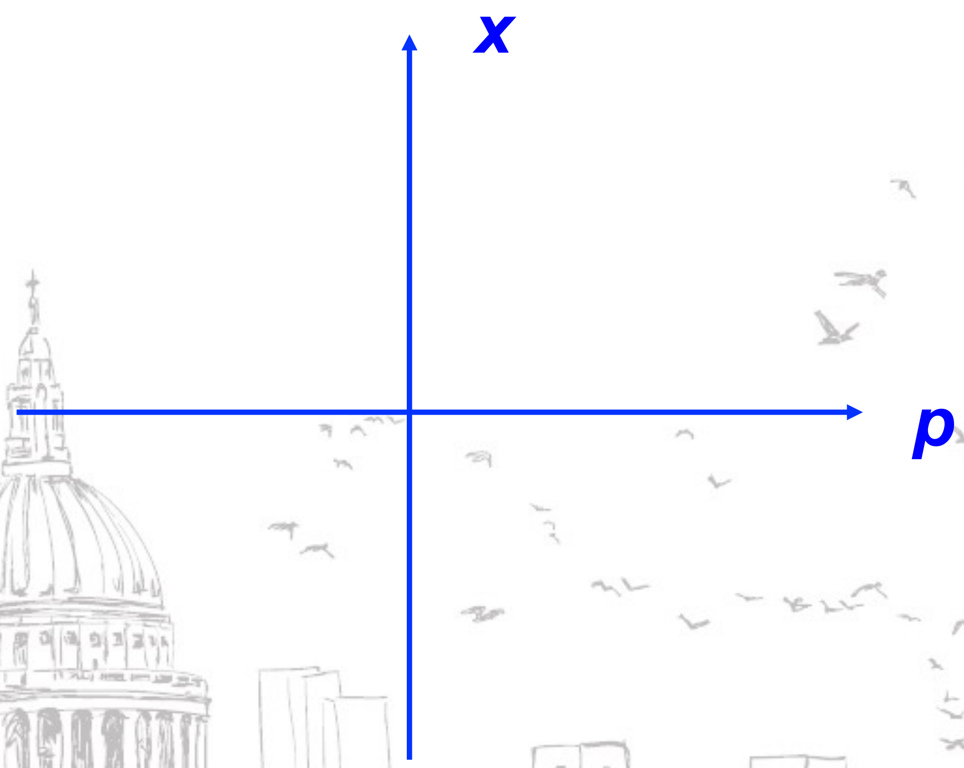

Simple example: stone falling from tower

$$(\underline{X}, \underline{\phi_t}, \mu)$$

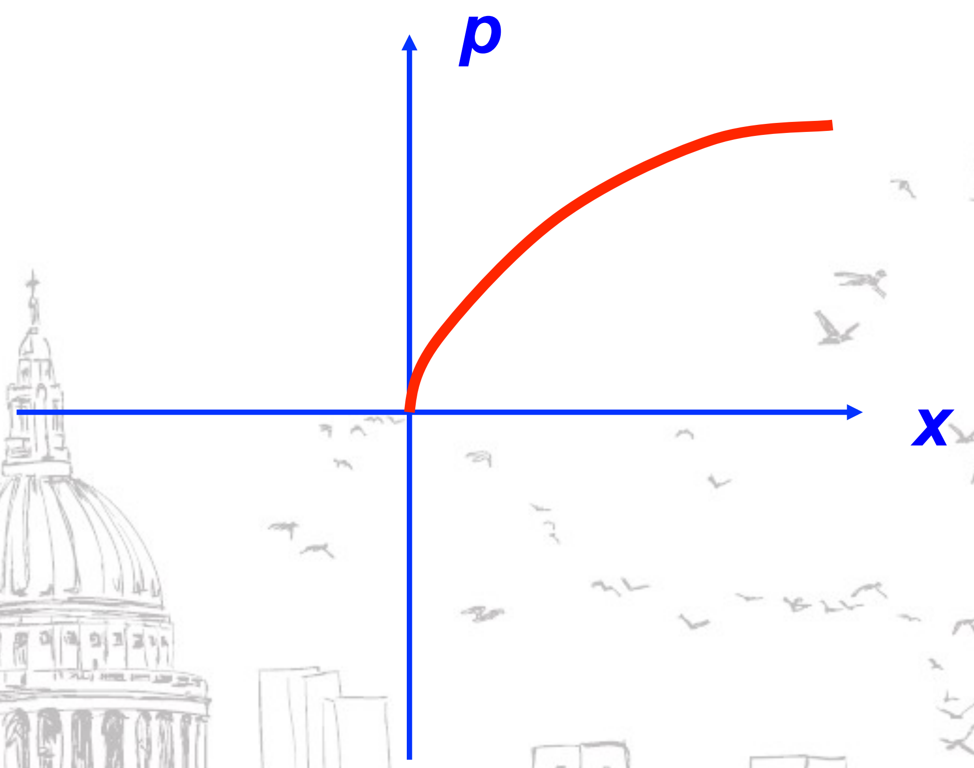

Simple example: stone falling from tower

$$(X, \phi_t, \mu)$$

“normal”
surface

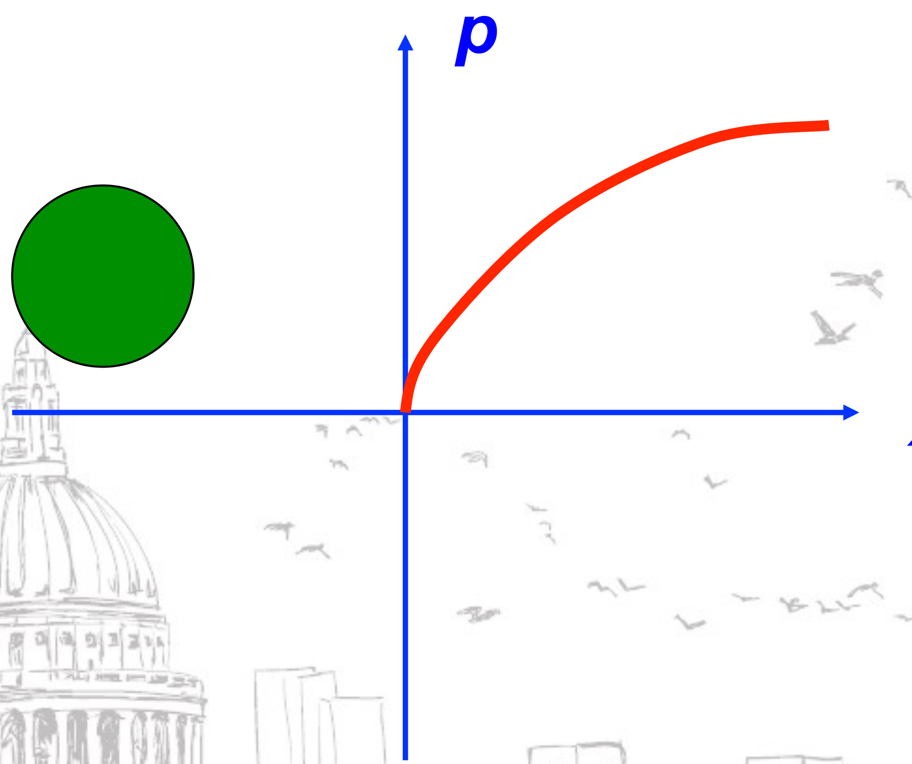

Difficult example: global climate model

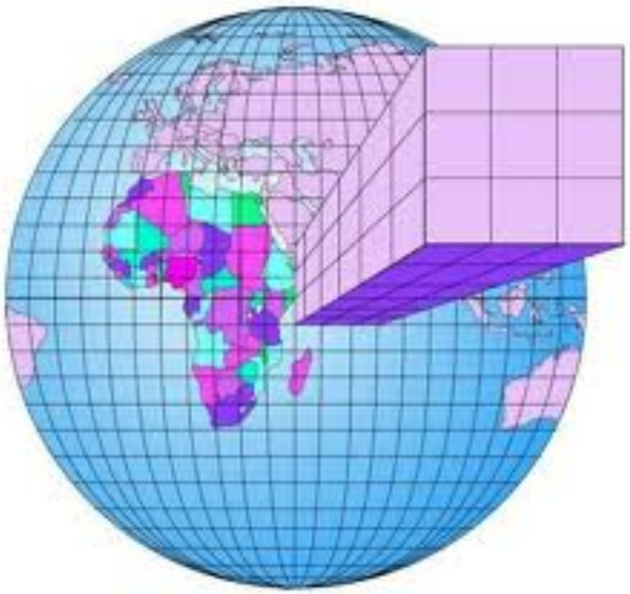

Difficult example: global climate model

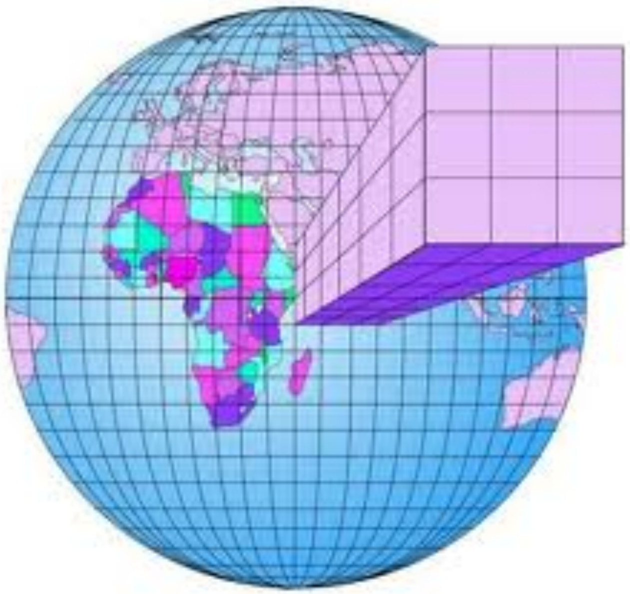

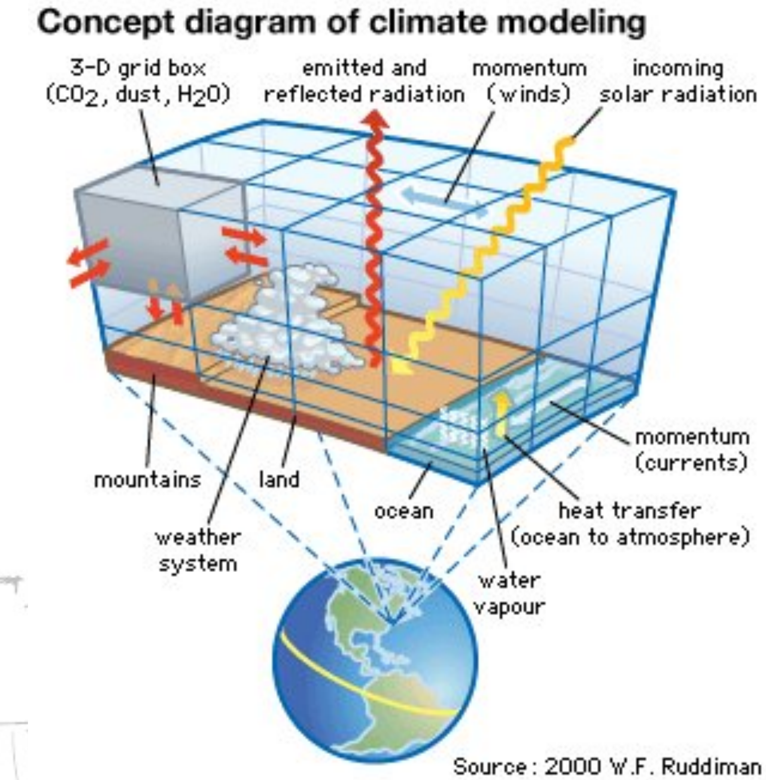

Difficult example: global climate model

Literally 10,000s of climate variables for the entire world

$$(\underline{X}, \phi_t, \mu)$$

Difficult example: global climate model

$$(\underline{X}, \underline{\phi}_t, \mu)$$

Literally 10,000s of climate variables for the entire world

The evolution of these variables over time

Difficult example: global climate model

$$(X, \phi_t, \mu)$$

Literally 10,000s of climate variables for the entire world

The evolution of these variables over time

The so-called invariant measure of the dynamics

Deterministic Chaos

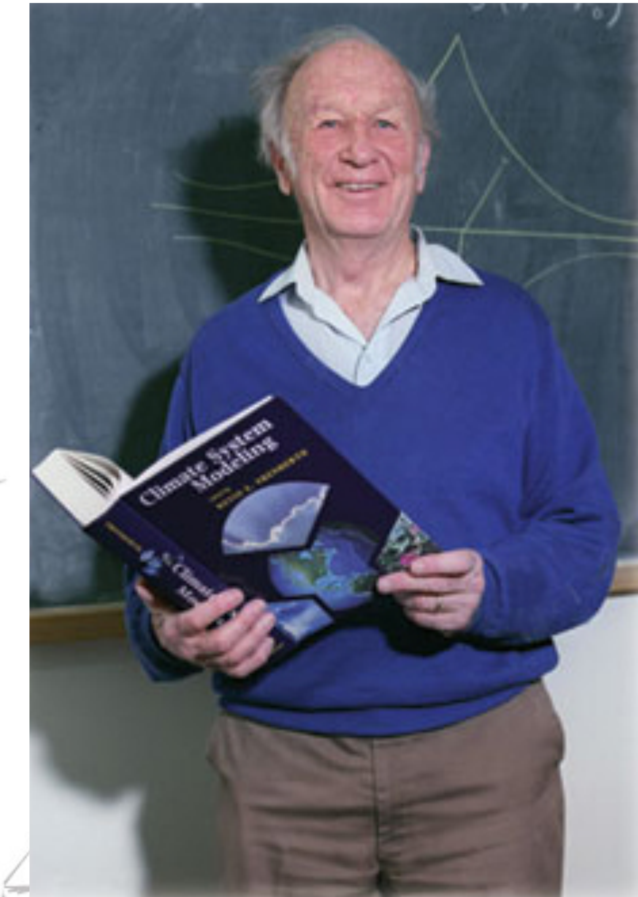

In non-linear equations we encounter so-called “sensitive dependence on initial conditions”

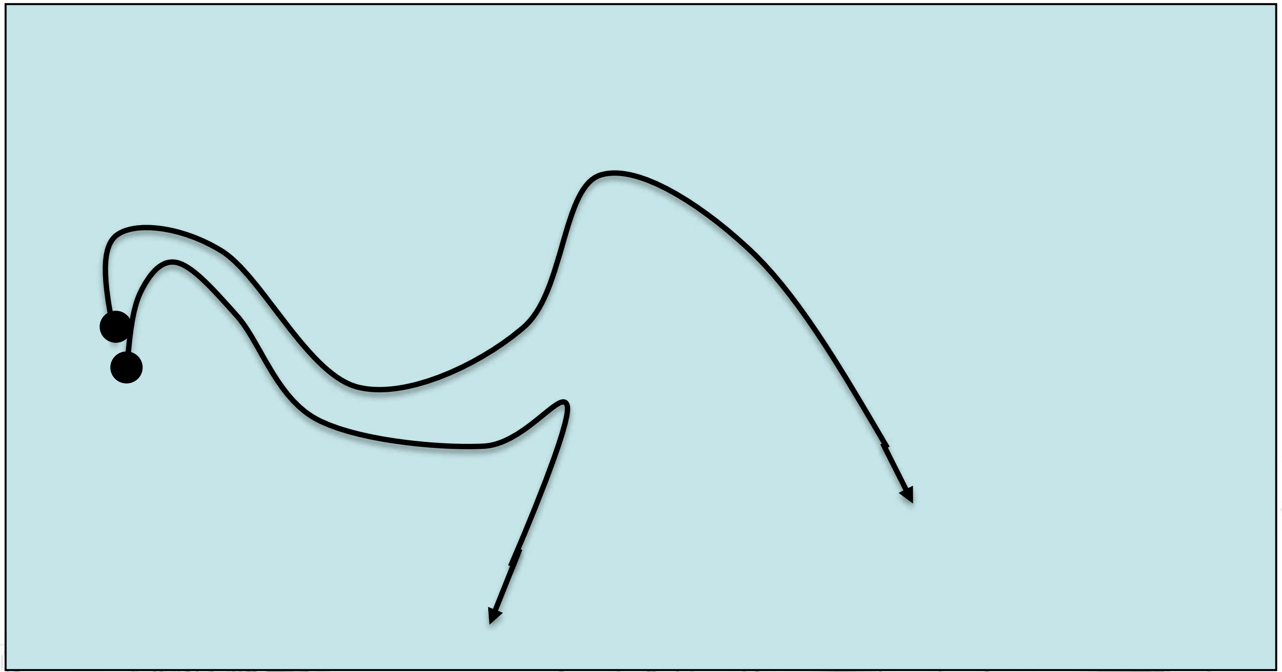

And notice: most target systems have non-linear equations.

- Simple 2 dimensional pendulum with a magnet
- …
- Our solar system
- The world’s climate system.

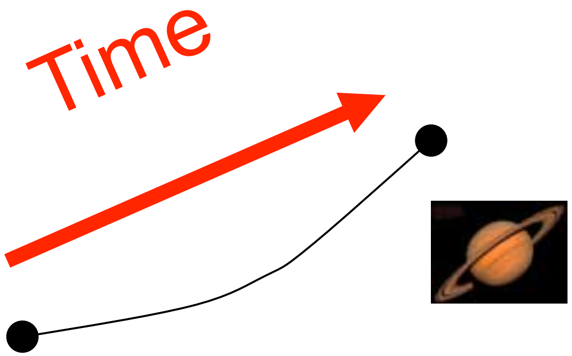
Time

Like-for-like:

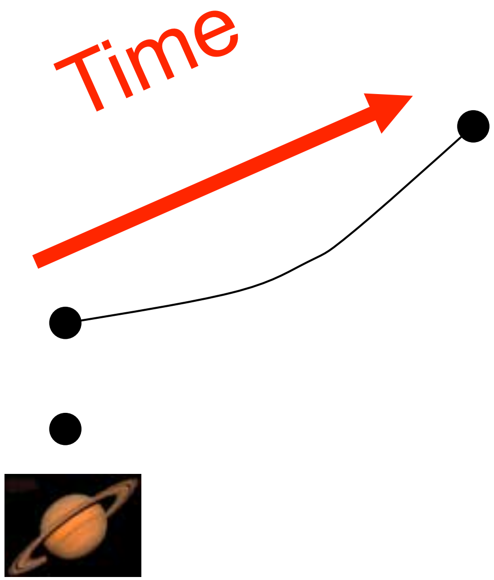

Expectation:

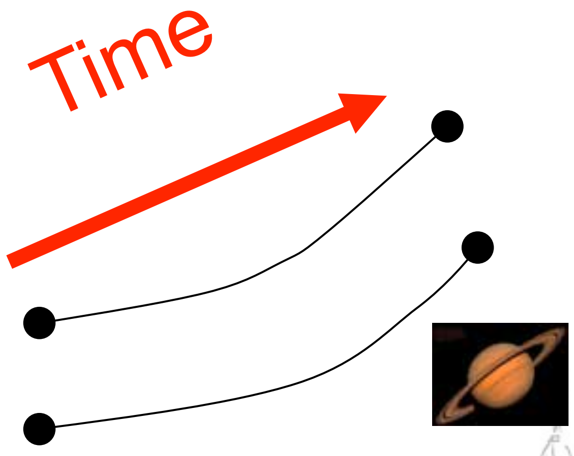

Expectation:

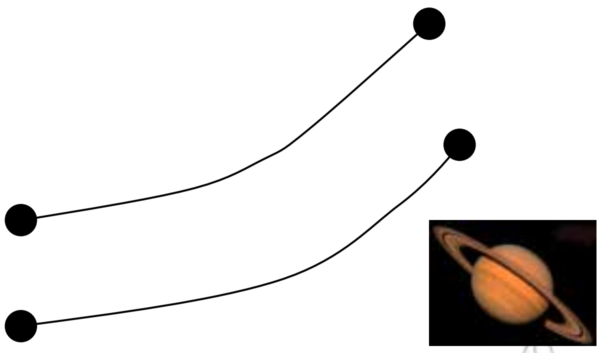

Chaos:

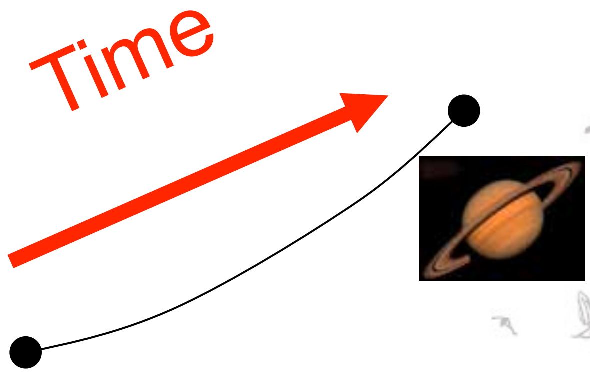

Expectation:

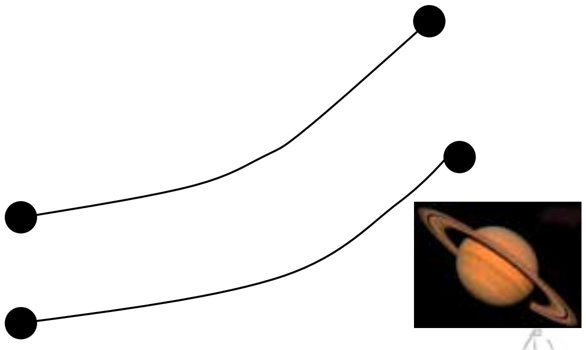

Chaos:

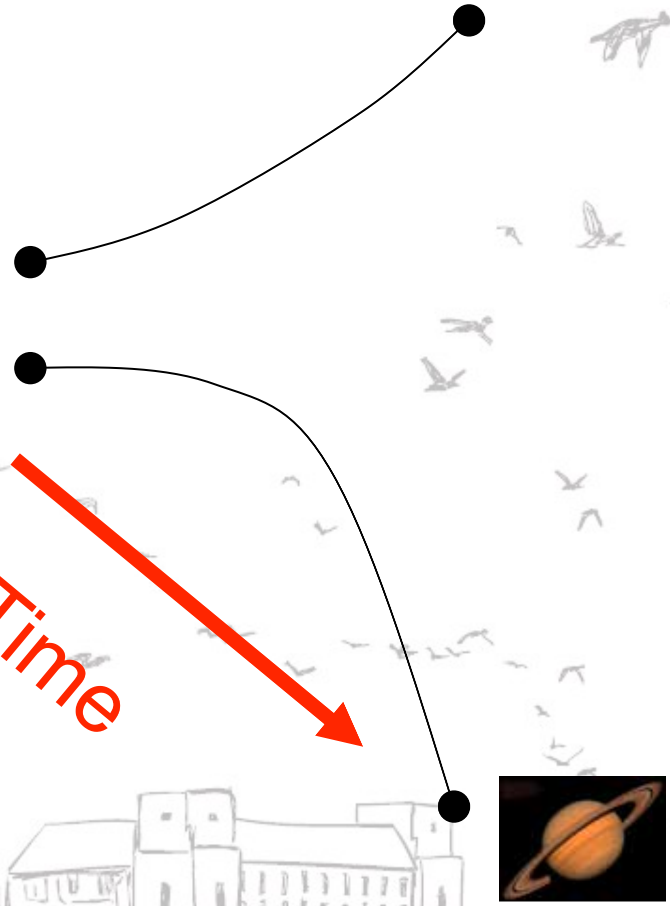

Notice: as commonly presented, chaos is an effect having to do with initial condition uncertainty.

We will see in Part II that there is more to it.

Section I: UKCP09

Starting point: Climate change is real!

But what does climate change mean at a ***local*** level?

Global: global mean temperature, average sea level rise, maybe melting of arctic ice.

Local: pretty much what happens in front of you door!

So if you work at LSE you want to know …

... is this what we will face?

… or this?

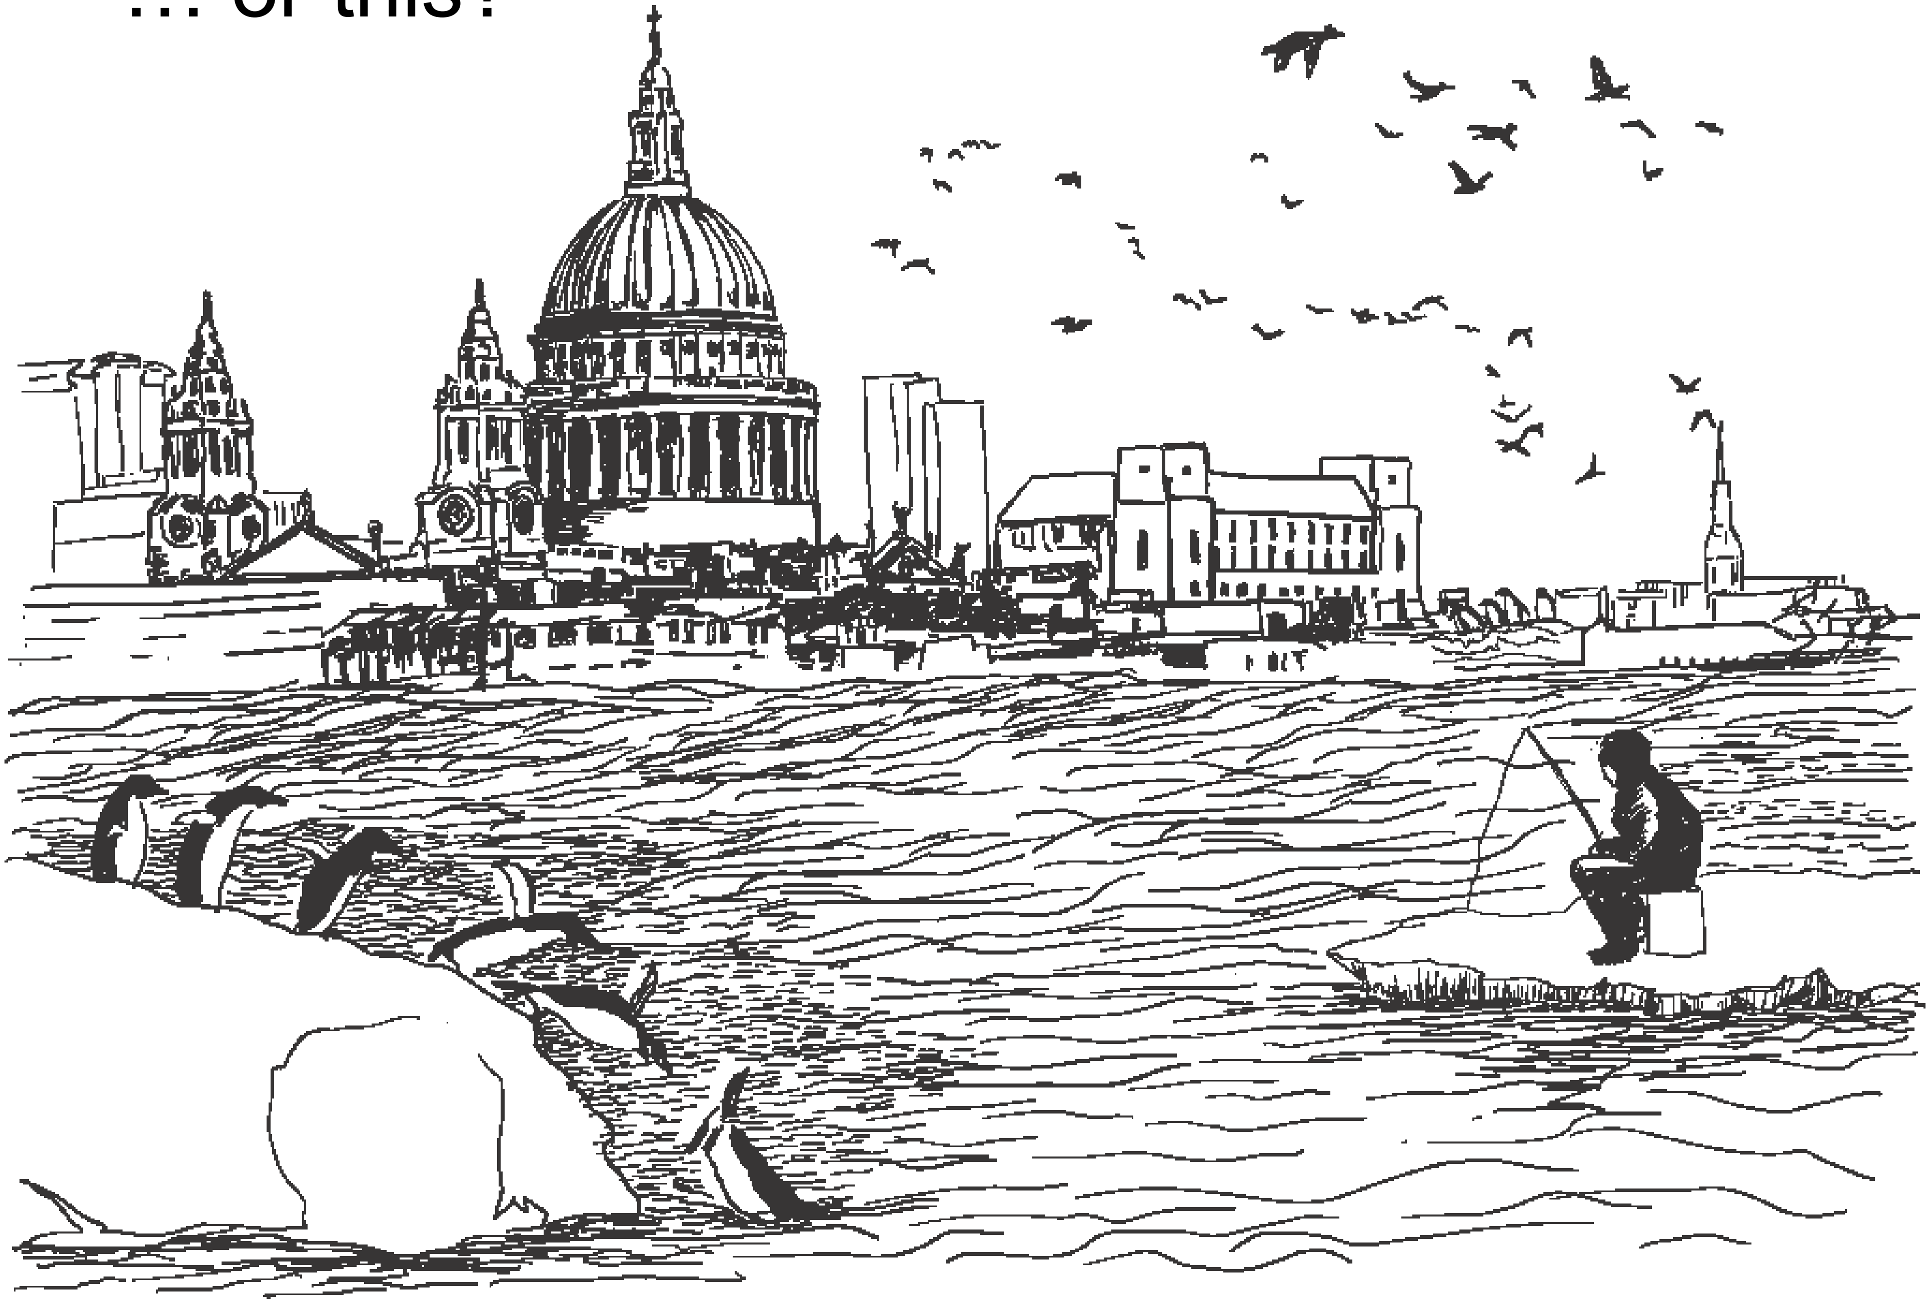

We want to know how the local climate changes because policy is made at the local level.

Make provisions:

- Adaption: flood walls, water provision, etc.
- Mitigation: implement changes and ideally stop bad things from happening.

Scientific impulse:

Build a model of the world’s climate taking all relevant factors into account and use it to predict what the future will be like.

Result: General Circulation Models (GCM’s)

These are run on the largest available super-computers in large scale multi-million projects such as UKCP, an initiative of the UK government to predict ***local effects*** of climate change in the UK up 2080.

Preview of the argument:

Running these models is costly both in terms of man/woman hours and £££.

→ **Central Question (1st formulation)**:

Can these models deliver as advertised?

(a) Are the outcomes of such large scale models trustworthy and reliable?
(b) Can they form the basis of responsible policy making?

***Philosophical impulse*:**

These questions are best answered by reflecting on the methods used in generating results.

This is the project for this talk.

→ Identify methods used

→ Reflect on their trustworthiness

→ Come to a verdict

Climate models: The details don't matter …

… what you need to bear in mind is

that ϕ_t has the following properties:

1. Strong simplifications are made to construct the model. So we are faced with ***model error***.
2. As a matter of fact the dynamics is ***nonlinear***.

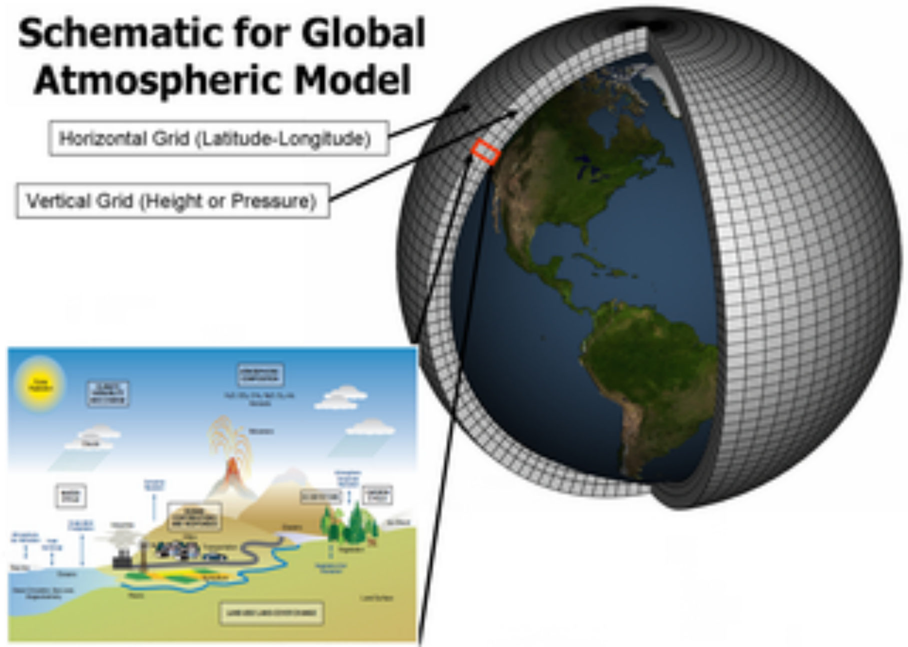

These are the two crucial features.

The initial question can now be reformulated as

Central Question (2nd formulation):

(a) Are the outcomes of ***nonlinear models with structural model error*** trustworthy and reliable?

(b) Can ***nonlinear models with structural model error*** form the basis of responsible policy making?

Overview: UKCP09

The *United Kingdom Climate Impacts Program's* UKCP09 project aims to answer questions about the local impact of global climate change by making high resolution forecasts of the local climate out to 2100.

The declared aim and purpose of UKCP09 is to provide decision-relevant forecasts, on which industry and policy makers can base their future plans.

The launch document says:

The projections have been designed as input to the difficult choices that planners and other decision-makers will need to make, in sectors such as transport, healthcare, water-resources and coastal defences, to ensure that UK is adapting well to the changes in climate that have already begun and are likely to grow in future.’ (Jenkins et al 2009, 9)

Probabilistic predictions are given on a 25km grid for finely defined events such as

- changes in the temperature of the warmest day of a summer
- precipitation of the wettest day of the winter

It is predicted, for instance, that under a medium emission scenario the probability for a 20-30% reduction in summer mean precipitation in central London in 2080 is 0.5

These probabilities are calculated using HadCM3.

The details are involved → See the Frigg&Stainforth&Smith paper in the folder.

Plan for here: emphasis some crucial assumptions.

Fact: HadCM3 involves strong idealising assumptions

→ It has structural model error.

→ Discussion last week.

UKCP09 acknowledges the presence of model error and suggests a way of dealing with it.

The message is that the uncertainties due to SME can be estimated and taken into account in projections.

First introduce a so-called discrepancy term:

$$\underbrace{c}_{\text{World}} = \underbrace{\varphi_{t*}^{S}(x_0;\alpha*)}_{\text{Model}} + \underbrace{d}_{\text{Difference}}$$

First introduce a so-called discrepancy term:

$$c = \varphi_{t*}^{S}(x_0;\alpha*) + d$$

Where:

c = true climate (in the world)

φ = model dynamics

x_0 = initial condition

α* = parameters that best simulate the target

d = discrepancy term

The discrepancy term tells us

‘what the model output would be if all the inadequacies in the climate model were removed, without prior knowledge of the observed outcome’ (Sexton *et al* 2012, 2515).

The discrepancy term tells us

‘what the model output would be if all the inadequacies in the climate model were removed, without prior knowledge of the observed outcome’ (Sexton *et al* 2012, 2515).

And then follows the crucial assumption:

‘Our key assumption is that sampling the effects of structural differences between the model chosen for the PPE and alternative models provides a reasonable proxy for the effects of structural errors in the chosen model relative to the real world.’ (*ibid*., 2516)

That is: UKCP09 compares the output of the model to 12 other models. The claim then is that measuring the average distance of HadSM3 to a set of different models yields a similar result as measuring its distance to the real world – hence, *d* can be determined by measuring by how much HadSM3 diverges from those other models.

Furthermore: the error in *d* is assumed to be Gaussian.

How good are these assumptions?

1. Is comparison with a model ensemble a good proxy for comparison with the world?

2. Is the error Gaussian?

Basic line of argument:

'Indeed, the multimodel ensemble mean has been shown to be a more skilful representation of the present-day climate than any individual member' (ibid.)

'the structural errors in different models can be taken to be independent' (ibid.)

Back to the discussion about ensembles:

- Is more skilful' close to being 'skilful'?
- Independence?

Recall the model ensemble on global mean temperature.

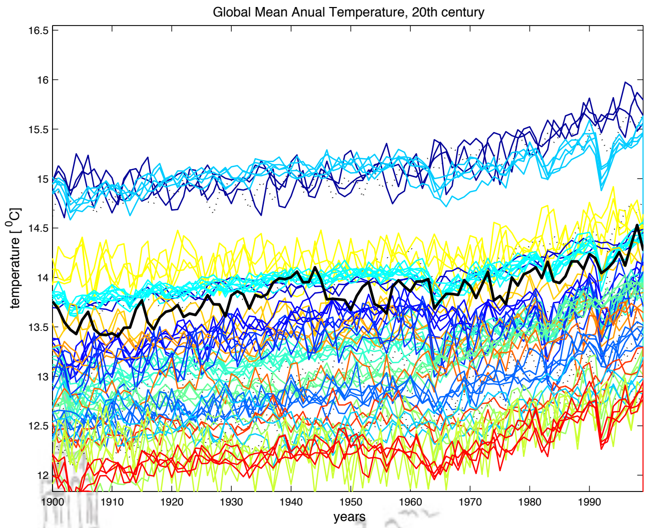

Models show warming but average temperatures vary tremendously. The magnitude of the error in the global mean in a hindcast casts significant doubt on the viability of the informativeness assumption on a 25 km forecast to the end of this century.

But even if we could reproduce the past incredibly well, would that be a reason to believe that we can predict the future well?

Main idea behind UKCP: getting close to the truth is good enough.

But is close enough good enough (even if one is willing to set all other worries aside)?

→ **Part II**

Part II:
Prediction and Model Error

Recall: **Central Question:**

(a) Are the outcomes of ***nonlinear models with structural model error*** trustworthy and reliable?

(b) Can ***nonlinear models with structural model error*** form the basis of responsible policy making?

Take-Home Message - Part 1

If **chaotic** models have even the slightest SME, their capacity to make meaningful forecasts is seriously compromised.

This has dramatic consequences for our ability to make the kind of forecasts about the future that policy makers would like to have.

Attention: *not* the same old story.

So far non-linearity (chaos) has been studied in connection with uncertainty about ***initial conditions***.

We ask what happens if we are uncertain about the correct ***model structure.***

These are completely different problems!

Butterfly effect:

Error in initial conditions

Butterfly effect:

Error in initial conditions

Hawkmoth Effect:

Error in the model structure (equations)

(Erica Thompson)

Take-Home Message – Part 2

We can mitigate against the butterfly effect by making probabilistic predictions rather than point predictions.

This route is foreclosed in the case of the hawkmoth effect: nothing can mitigate against that effect!

So structural model error and not uncertainty in the initial conditions is what truly limits predictive power.

Or: butterflies are pretty; hawkmoths are ugly.

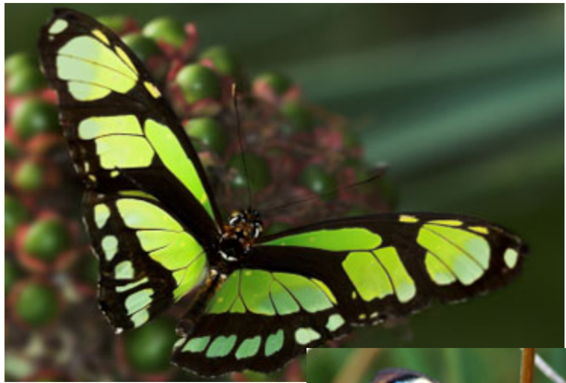




Locating the Issues

Dynamical system (X, ϕ_t, μ)

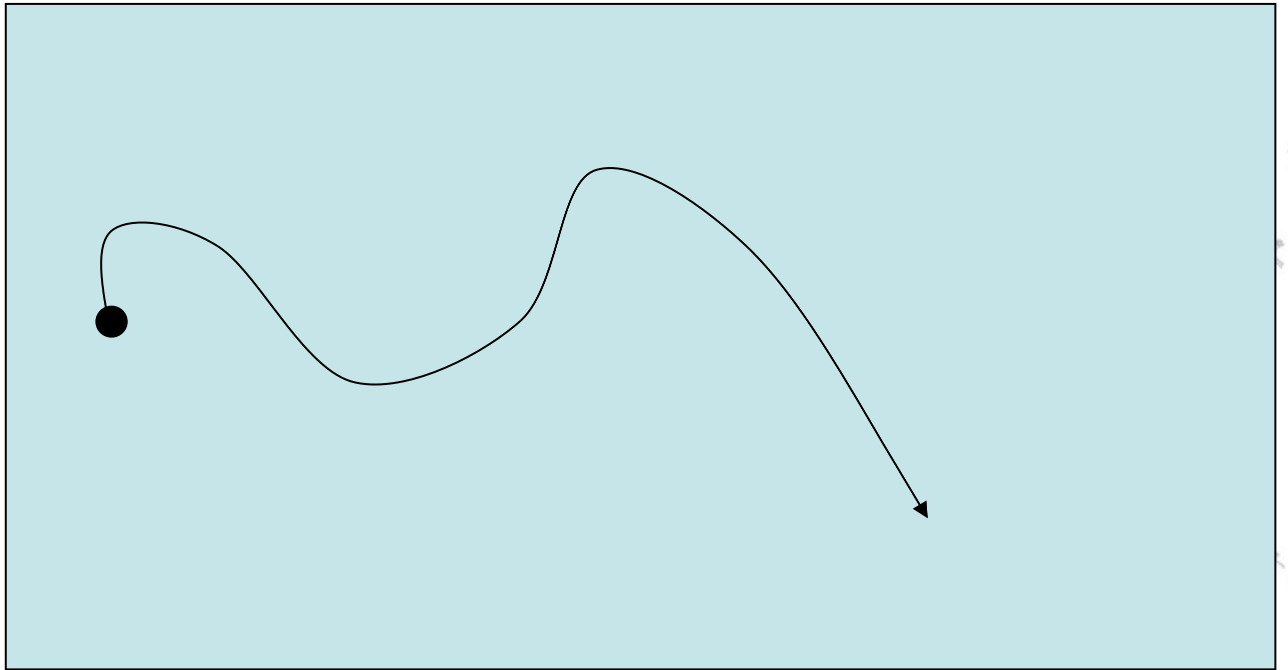

Locating the Issues

Dynamical system (X, ϕ_t, μ)

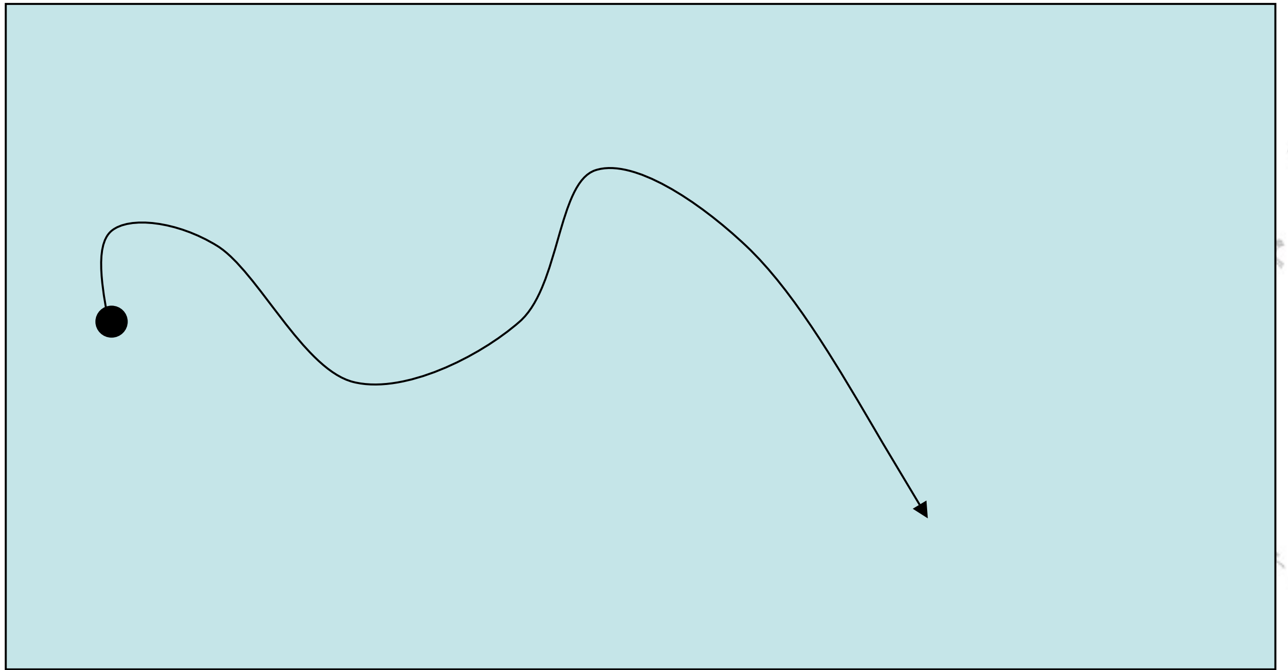

Locating the Issues

Dynamical system (X, ϕ_t, μ)

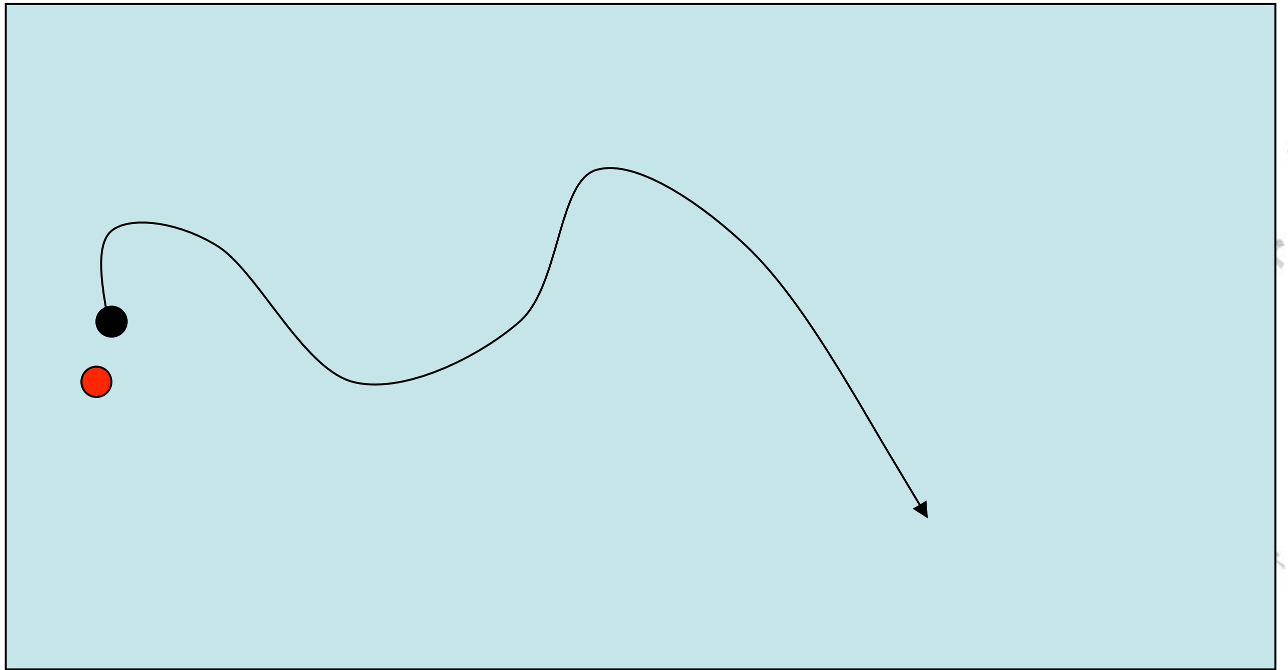

Initial Condition Error (ICE)

Locating the Issues

Dynamical system (X, ϕ_t, μ)

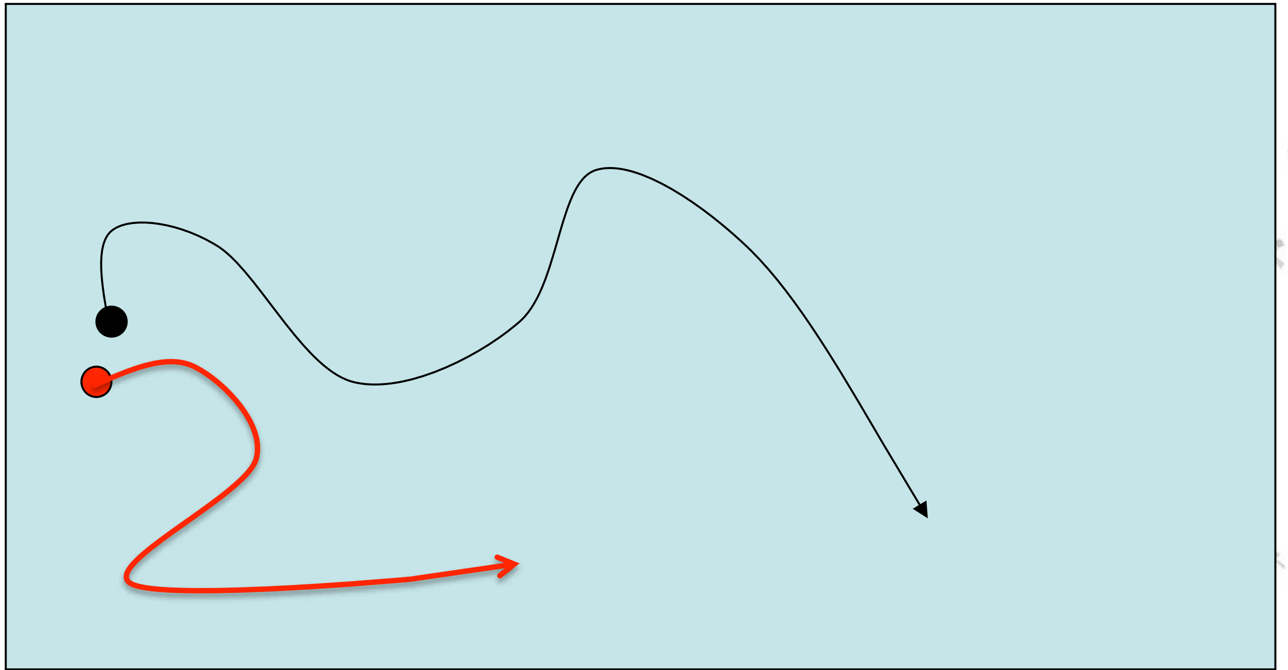

Initial Condition Error (ICE)

Locating the Issues

Initial condition error

↓

Butterfly Effect

Locating the Issues

SME: the time evolution of the model differ from the time evolution of the system under study:

$$\phi_t^S = \phi_t^M + \delta_t$$

True dynamics Model "Difference"

Locating the Issues

Dynamical system (X, ϕ_t, μ)

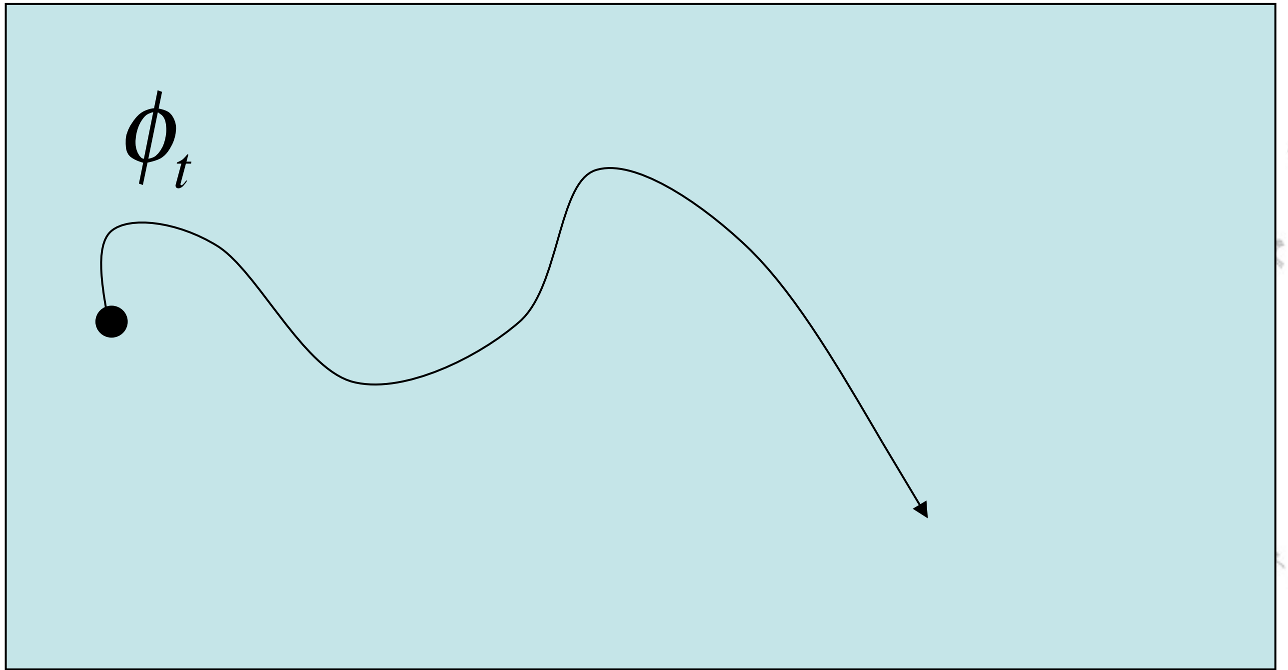

Locating the Issues

Dynamical system (X, ϕ_t, μ)

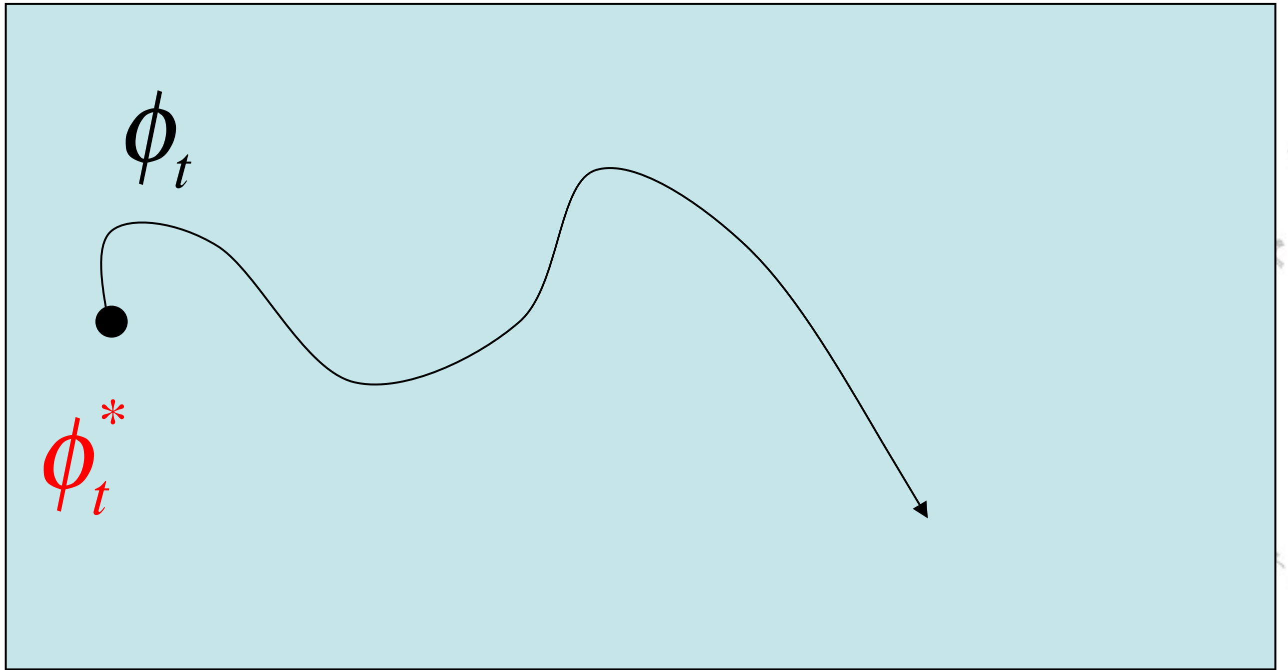

Locating the Issues

Dynamical system (X, ϕ_t, μ)

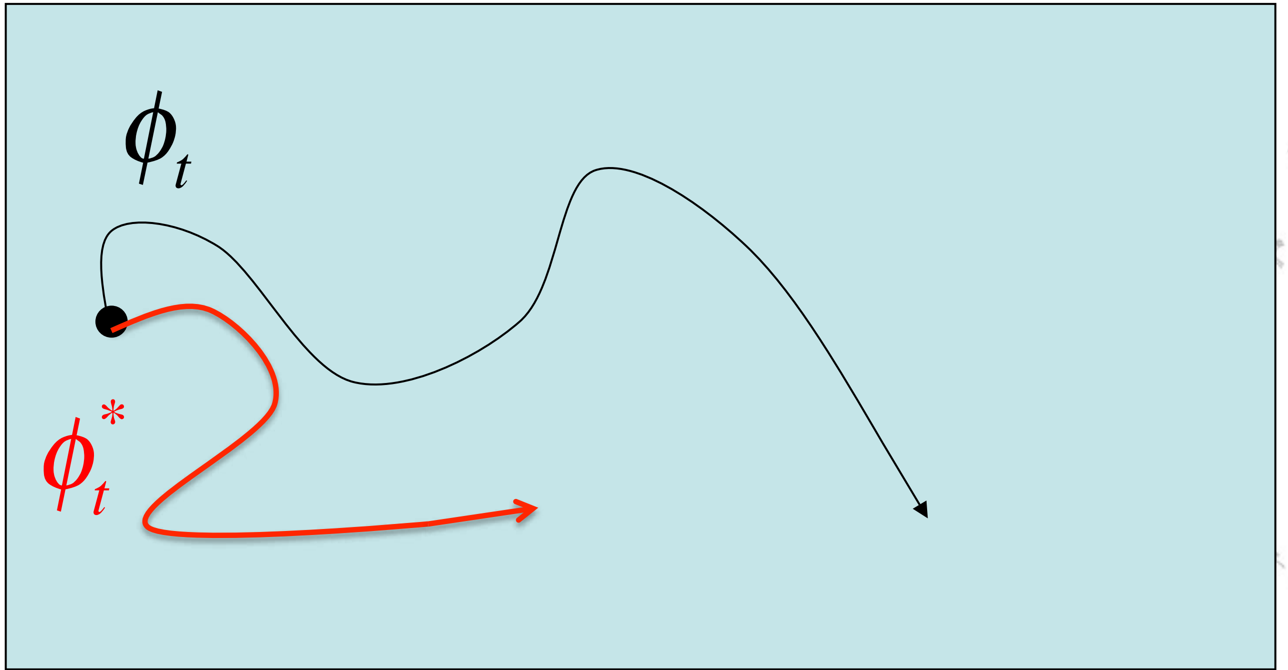

Locating the Issues

Structural Model Error

↓

Hawkmoth Effect

ICE versus SME

Meet Laplace’s Demon

1. Unlimited computational power
2. Unlimited dynamical knowledge
3. Unlimited observational power

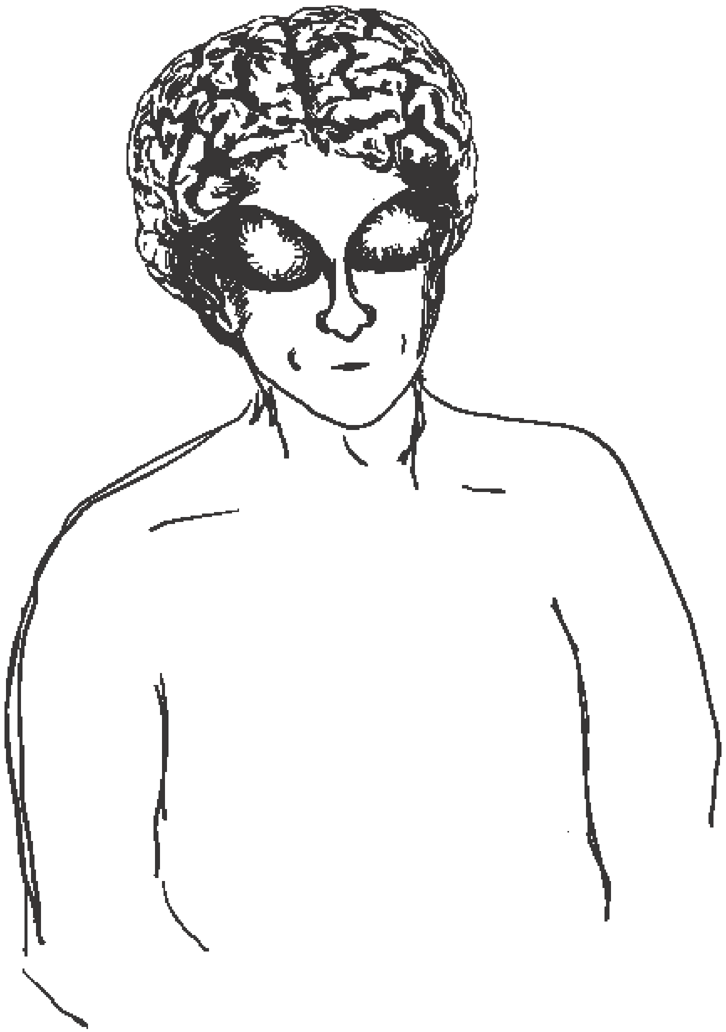

The Demon knows everything: ‘nothing would be uncertain and the future, as the past, would be present to [his] eyes’ (Laplace 1814).

So the Demon’s model of the world’s climate would be trustworthy because it provides the full truth.

But what happens if we are less capable than the Demon?

Meet the Senior Apprentice

1. Unlimited computational power
2. Unlimited dynamical knowledge
3. No unlimited observational power

Meet the Senior Apprentice

1. Unlimited computational power
2. Unlimited dynamical knowledge
3. No unlimited observational power

In other words, the Senior Apprentice only has noisy observations.

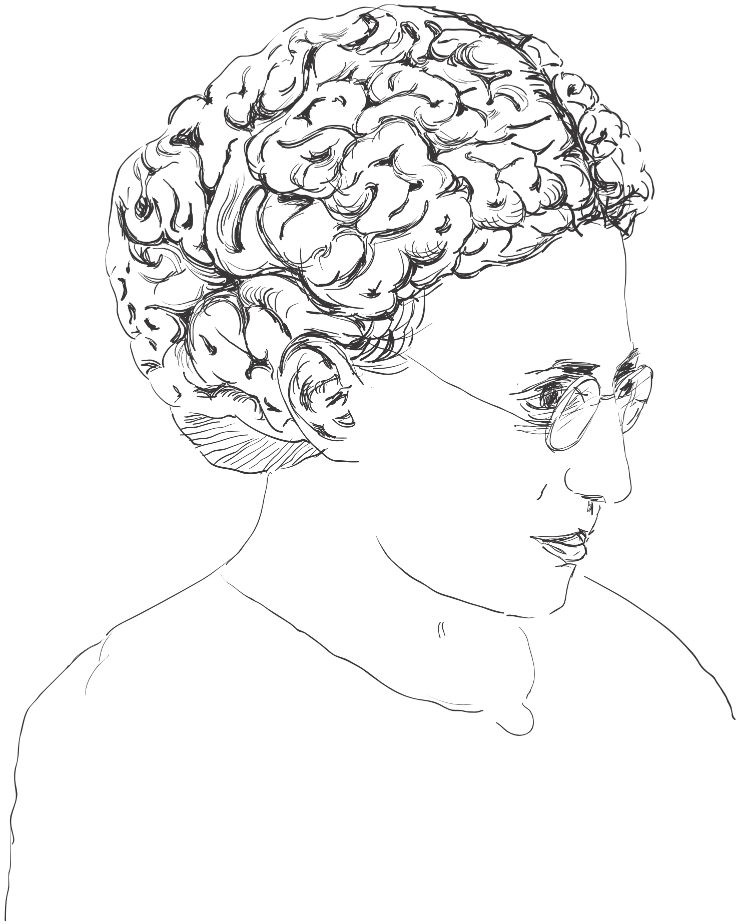

How could the limitation of not having unlimited observational power be overcome?

Reply:

Initial Condition Ensemble

That is, she puts a probability distribution over an approximate initial condition.

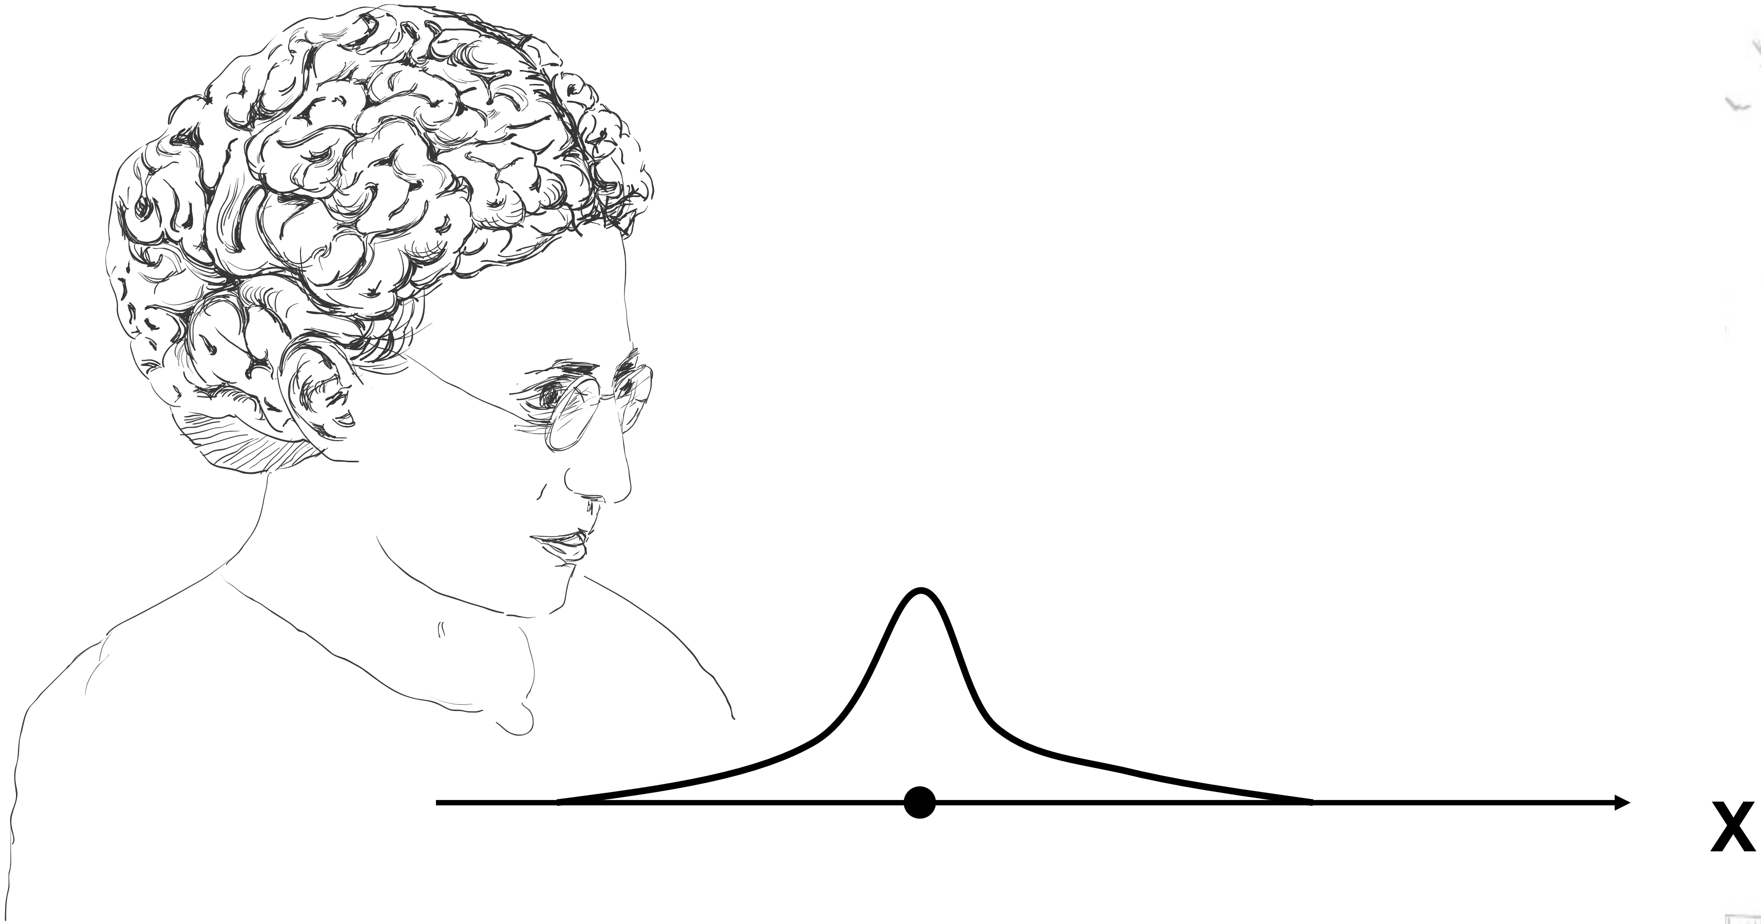

Prediction? Time Evolution?

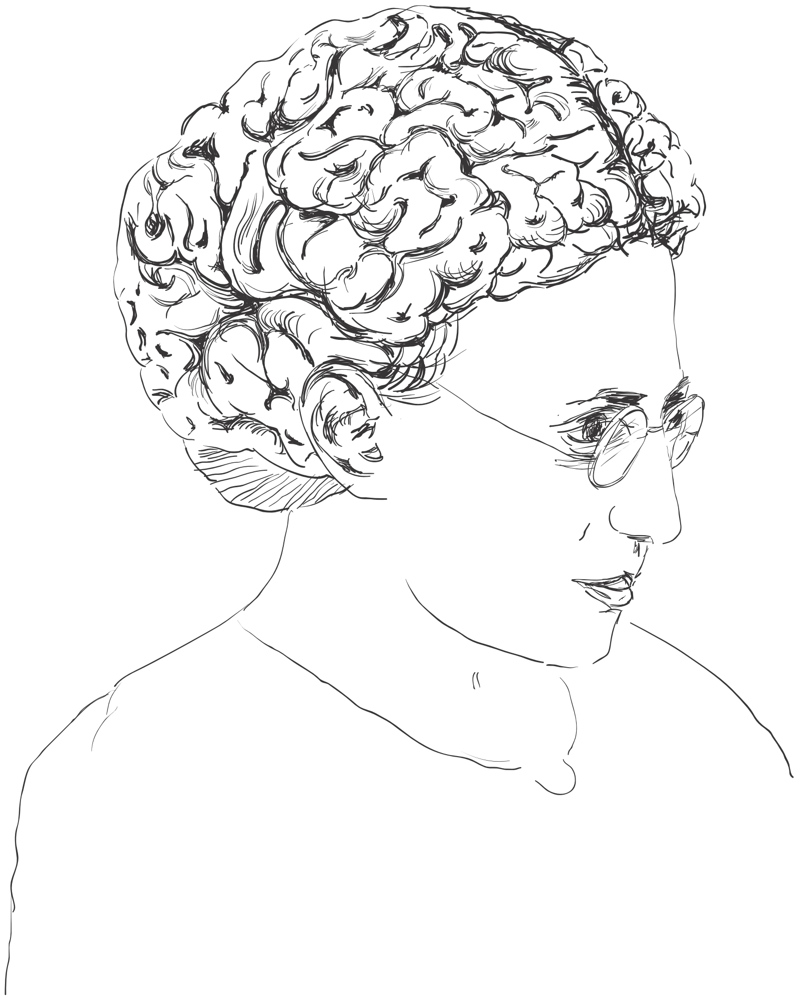

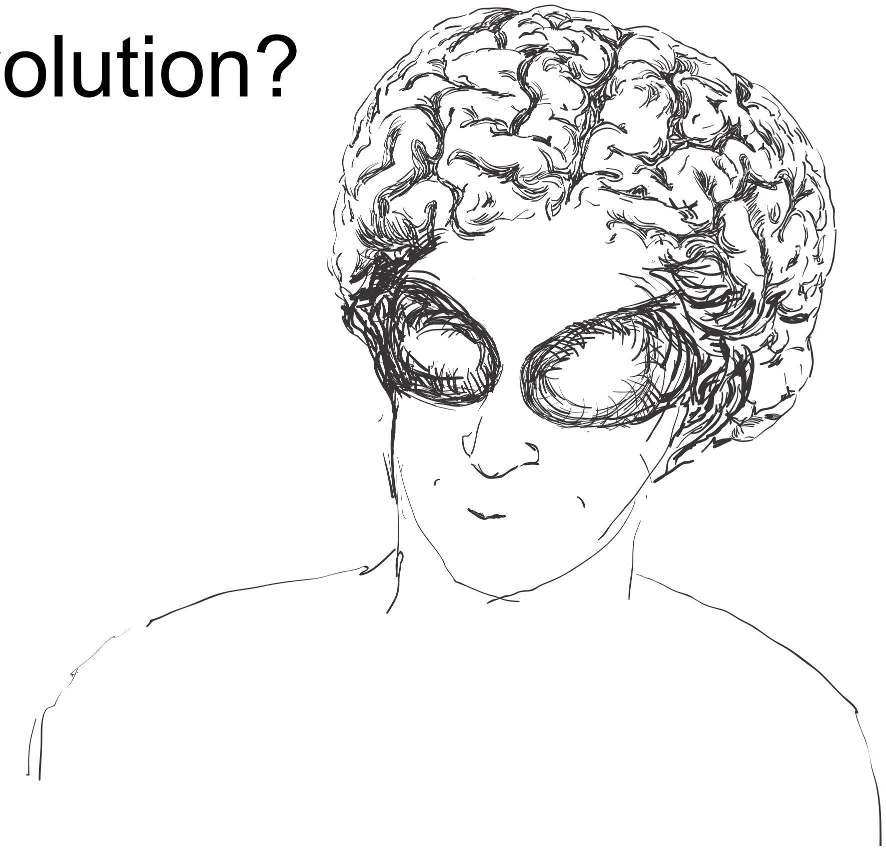

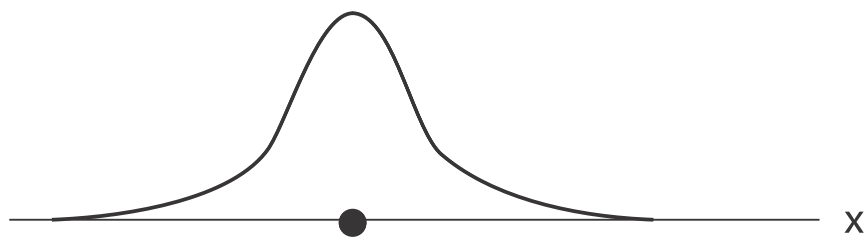

Generate probabilistic predictions by moving the initial probability distribution forward in time:

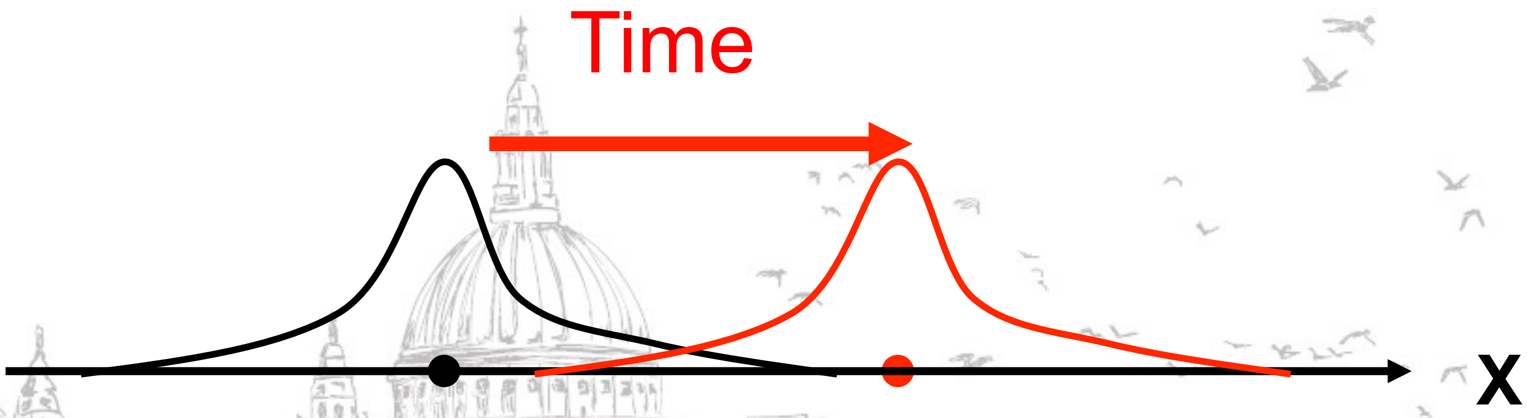

Implications for prediction? They figure that in non-linear systems we expect the probability distribution to disperse.

Implications for prediction? They figure that in non-linear systems we expect the probability distribution to disperse.

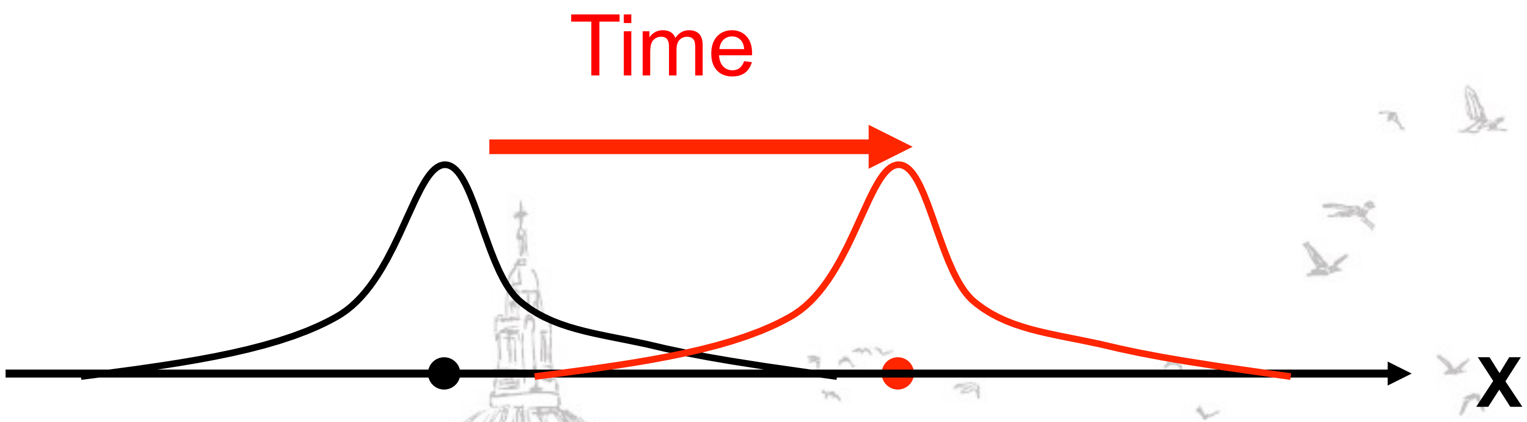

Implications for prediction? They figure that in non-linear systems we expect the probability distribution to disperse.

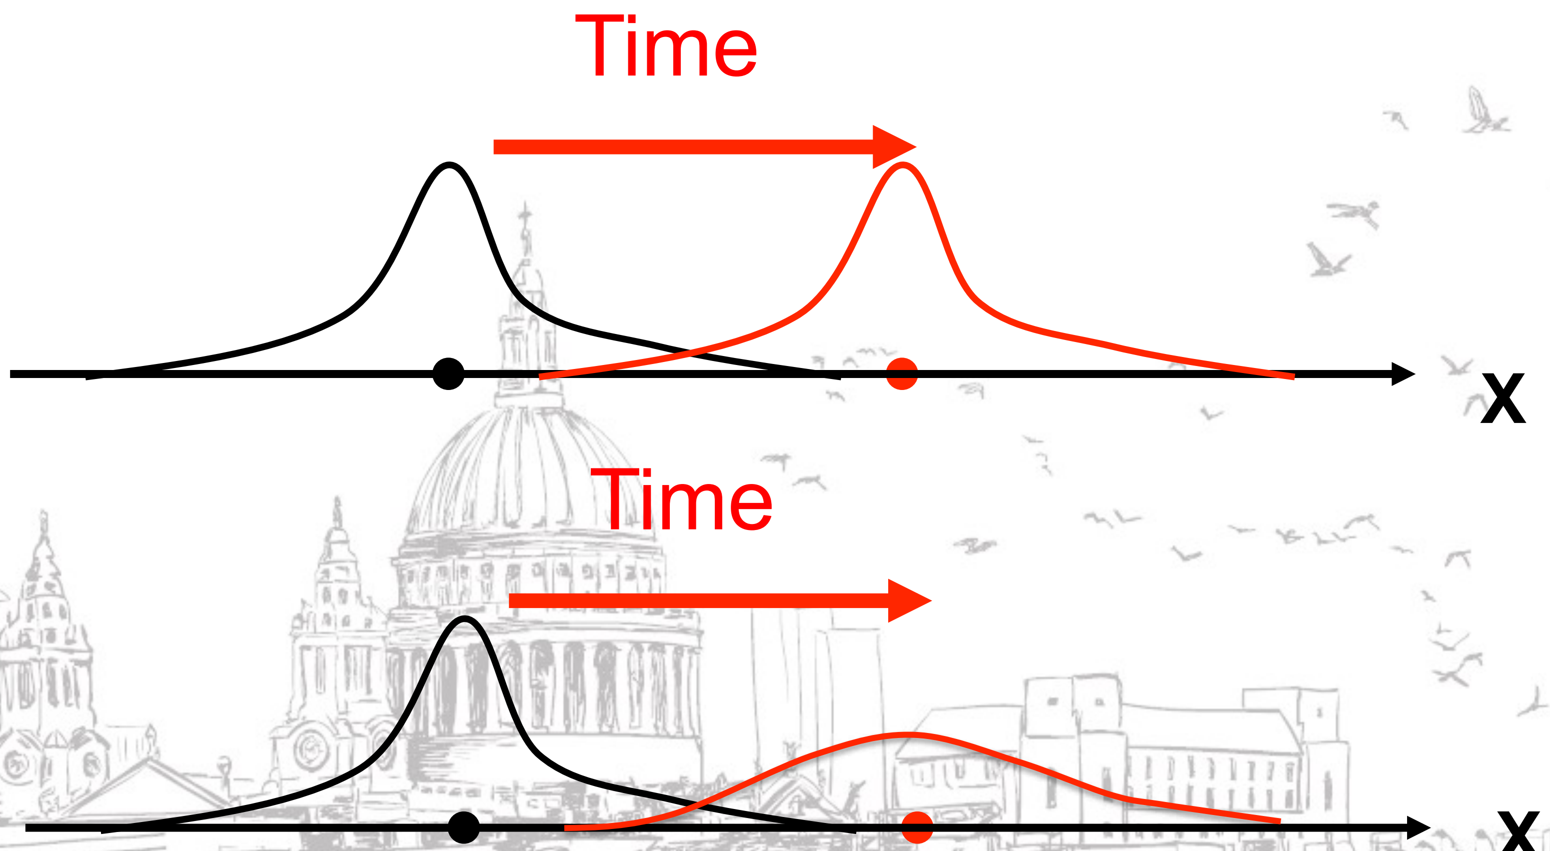

Why dispersion?

Why dispersion?

Distributions become *uninformative* as time passes, but they *do not become misleading*.

The Senior Apprentice realises that this is the limitation that she has to accept.

It is the price to pay for not having unlimited observational power.

Or: butterflies are pretty; hawkmoths are ugly.

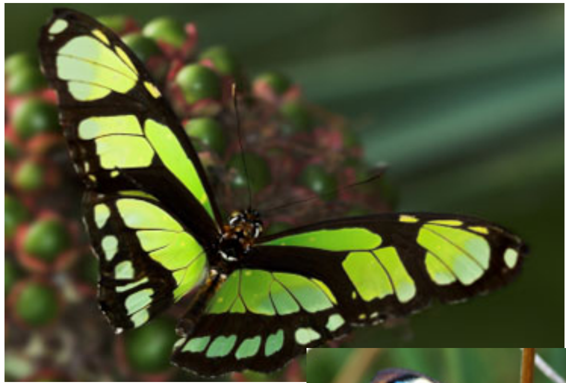




Meet the Freshman Apprentice

1. Unlimited computational power
2. No unlimited dynamical knowledge
3. No unlimited observational power

The Freshman Apprentice now claims he can do everything that the Senior Apprentice can do, his additional limitation notwithstanding

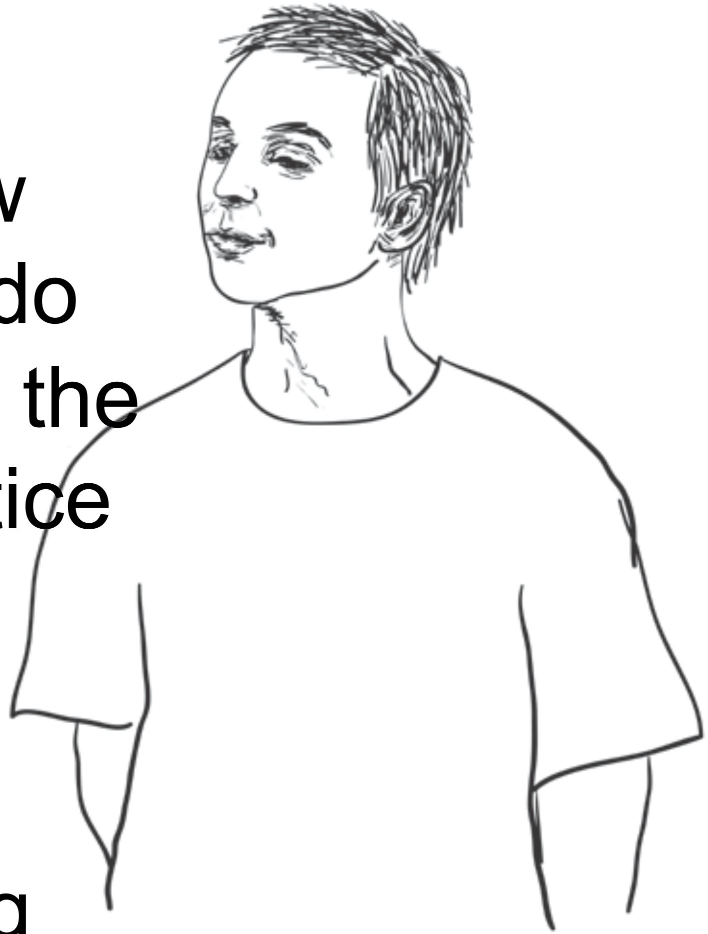

Recall: The Freshman can't formulate the exact dynamics of a system.

Reaction: Distortions and idealisations of all kind are acceptable as long as the resulting model is close enough to the truth.

This is the ***closeness-to-goodness link***.

→ This is a crucial part!

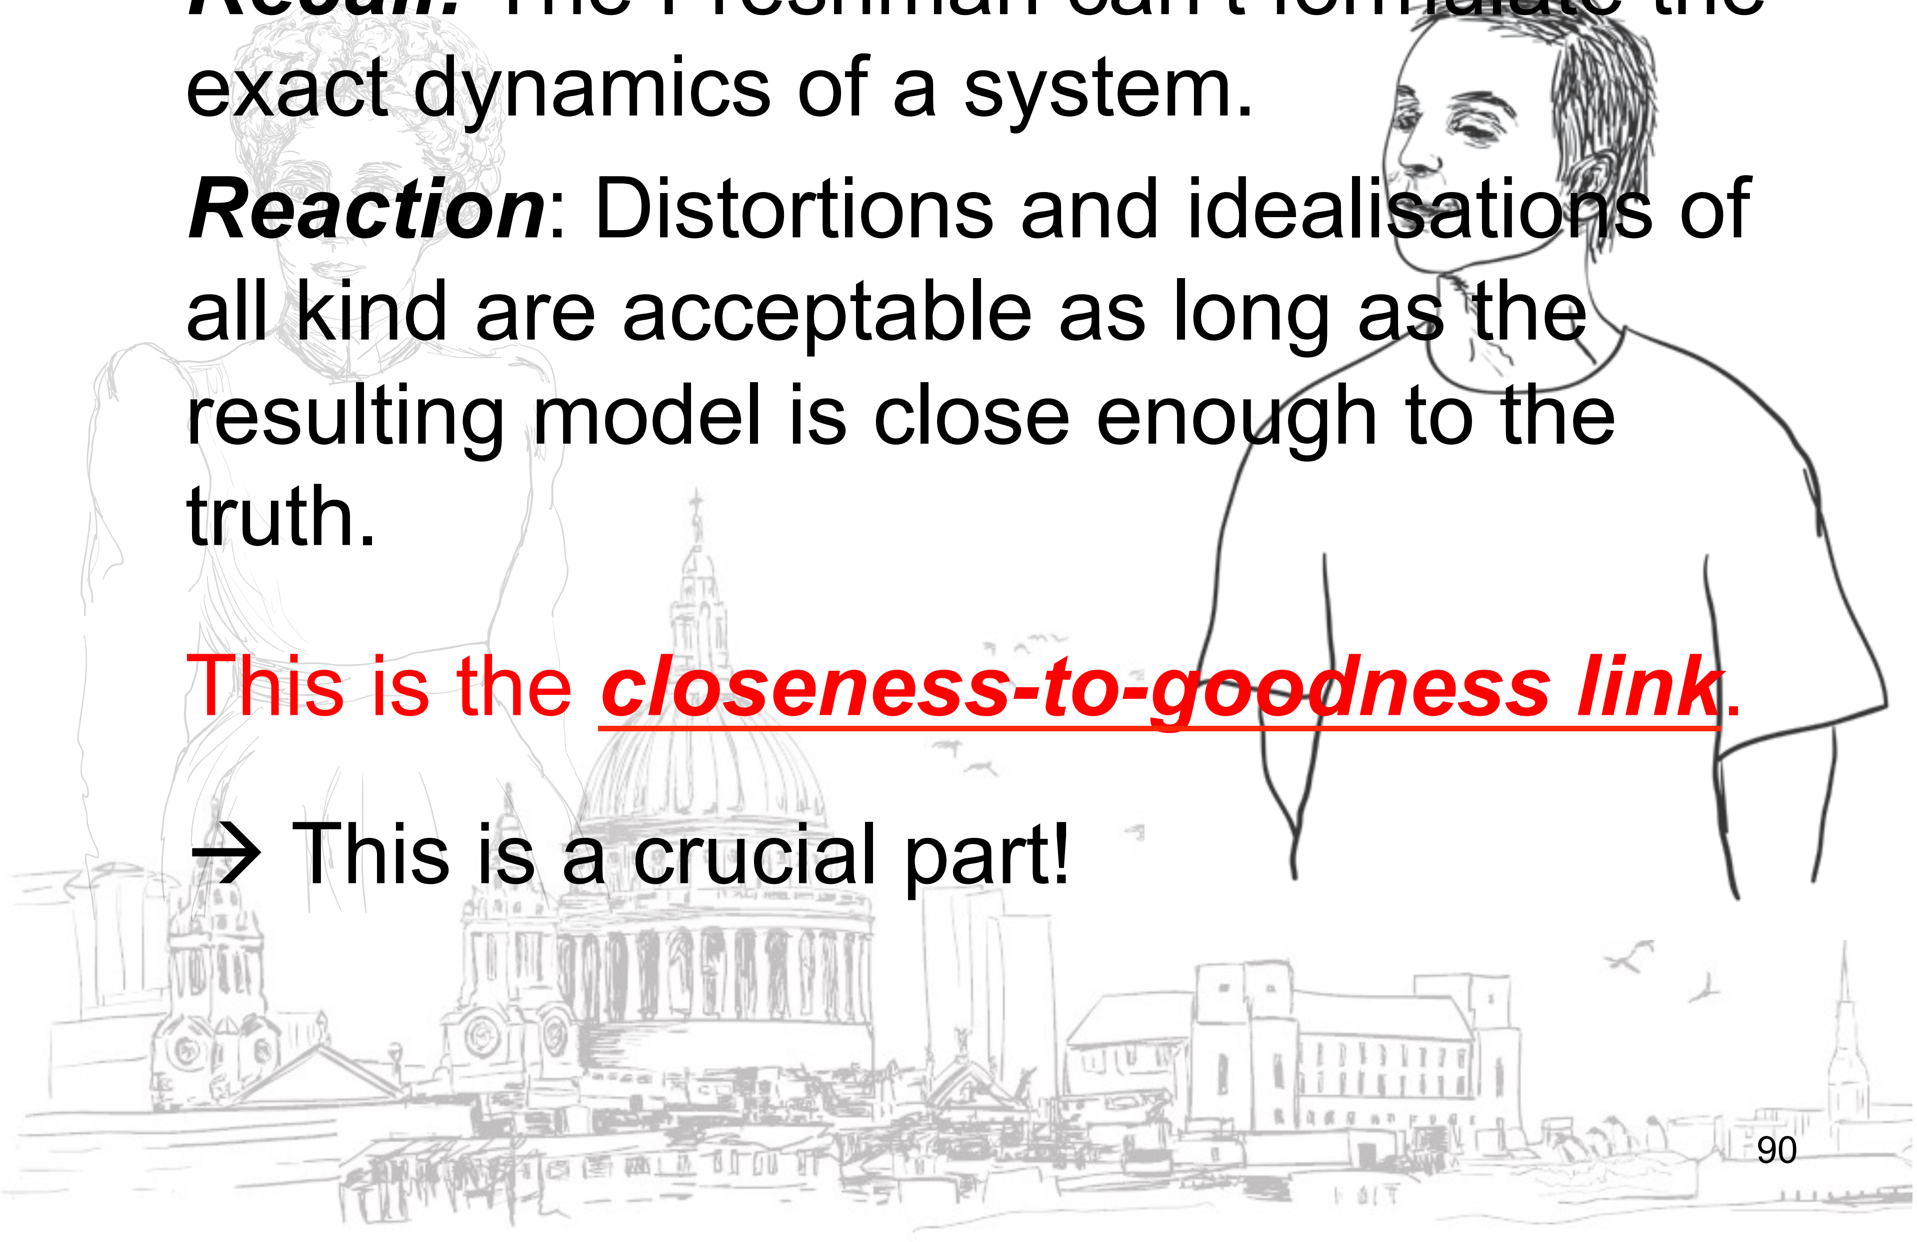

That is, the Freshman claims that *his* probabilistic predications are as good as the *Senior Apprentice*'s because he can rely on the closeness to goodness link.

Notice: Real scientists bear striking similarities to the Apprentice.

Question: is the Apprentice right?

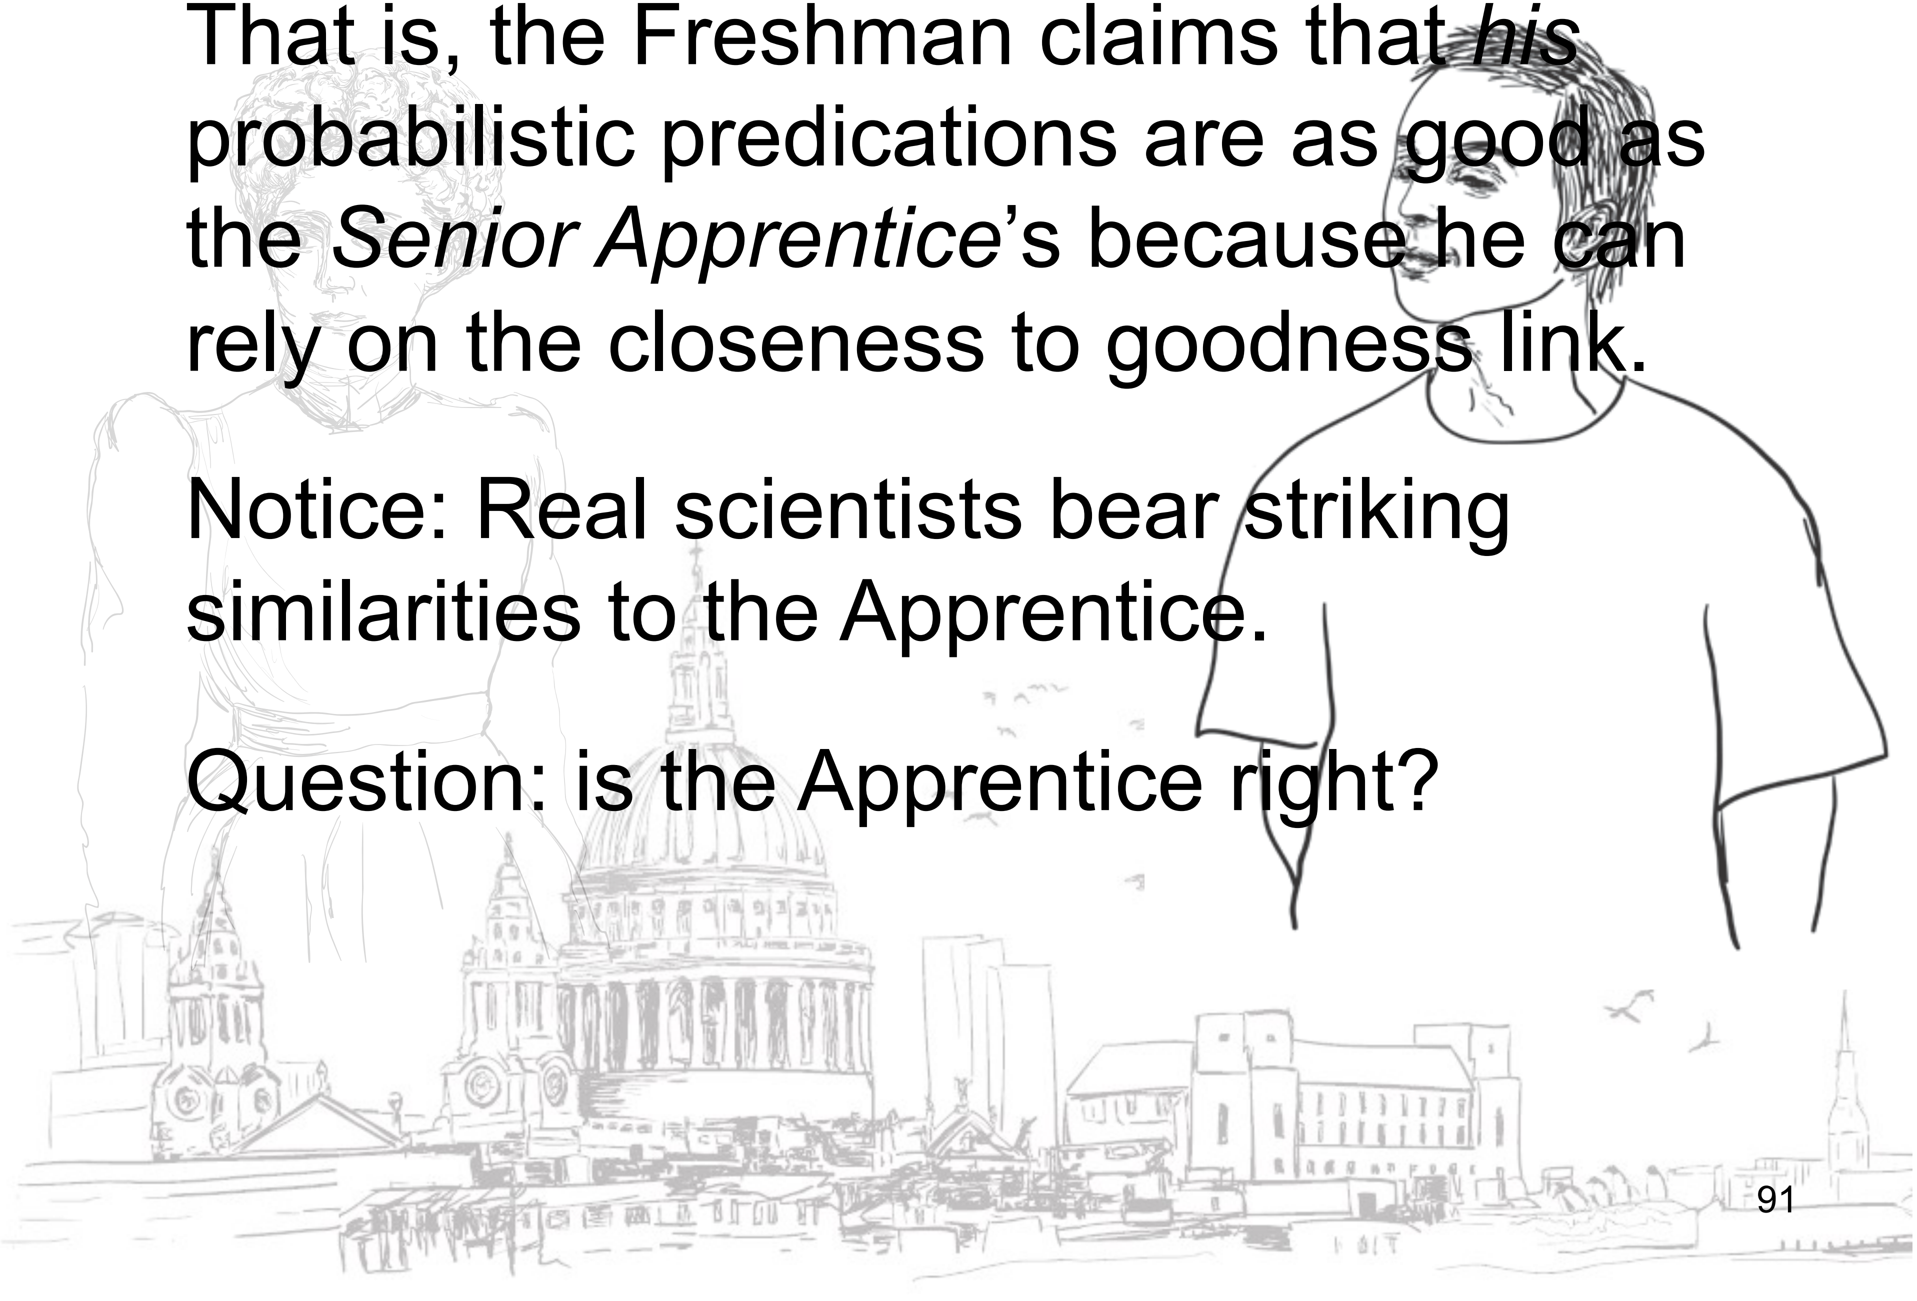


No way!

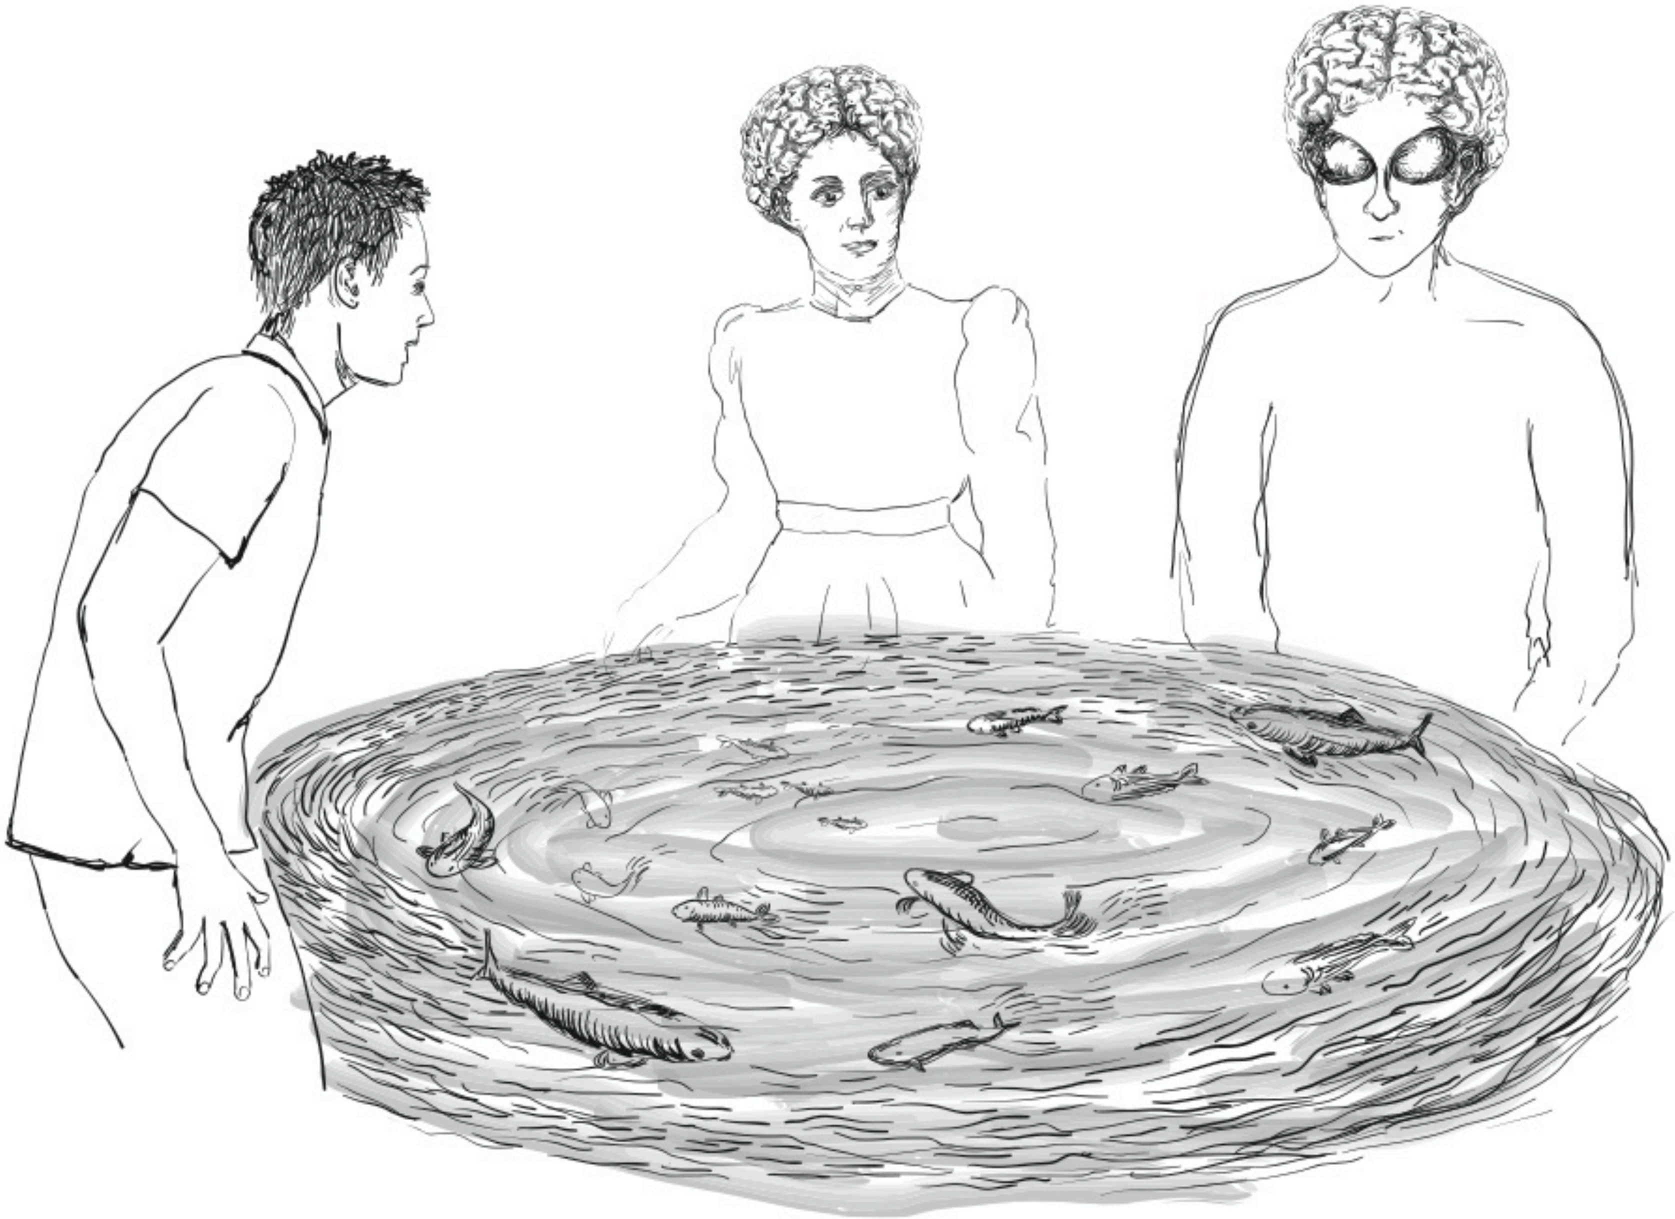

Population density:

$$\rho = \frac{\#\,\text{fish} / m^3}{\#\,\text{max fish} / m^3}$$

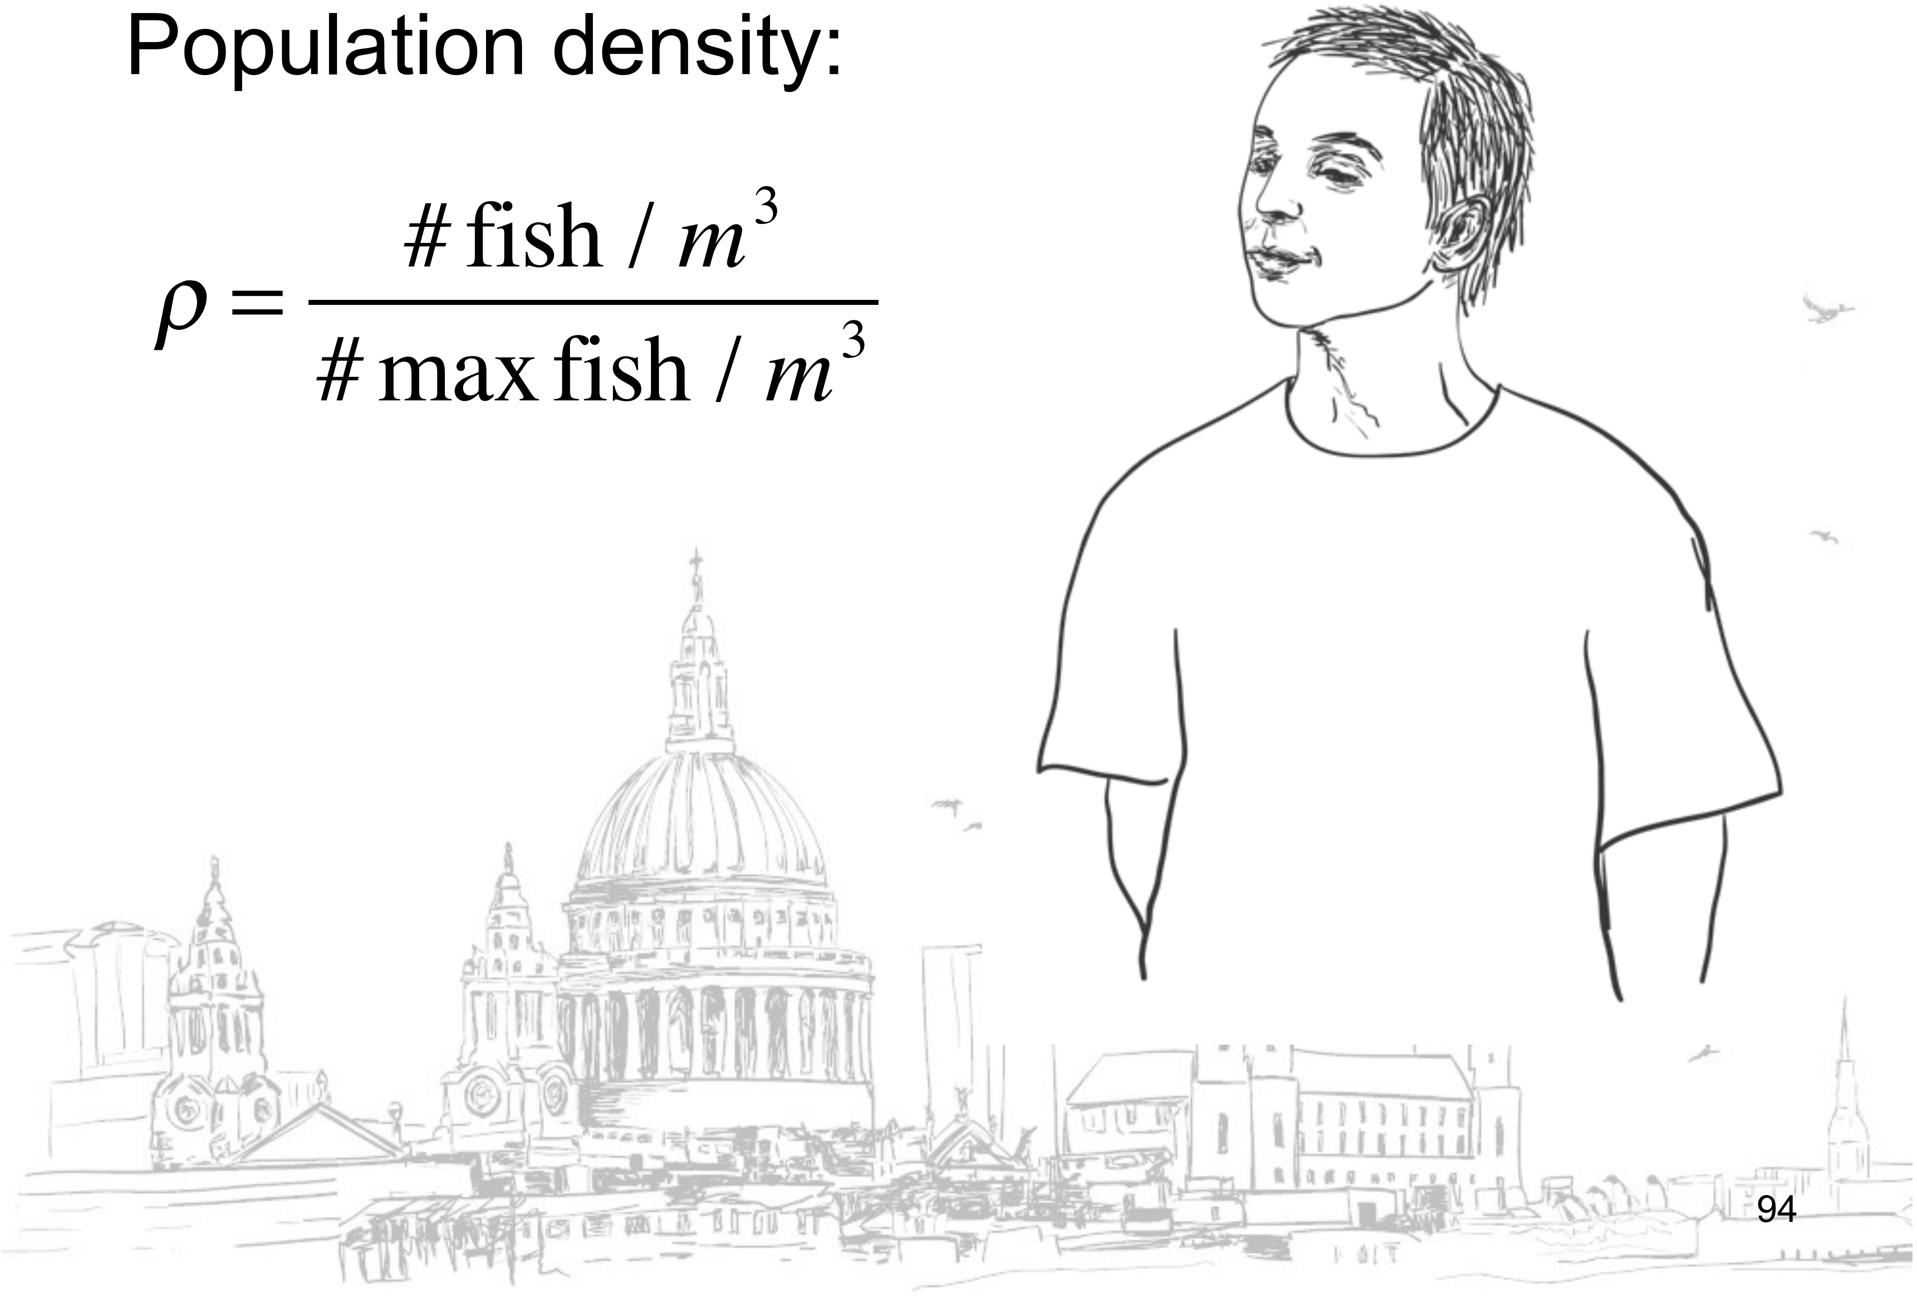

Population density:

$$\rho = \frac{\#\,\text{fish} \,/\, m^3}{\#\,\text{max fish} \,/\, m^3}$$

Hence: $\rho \in [0,1]$

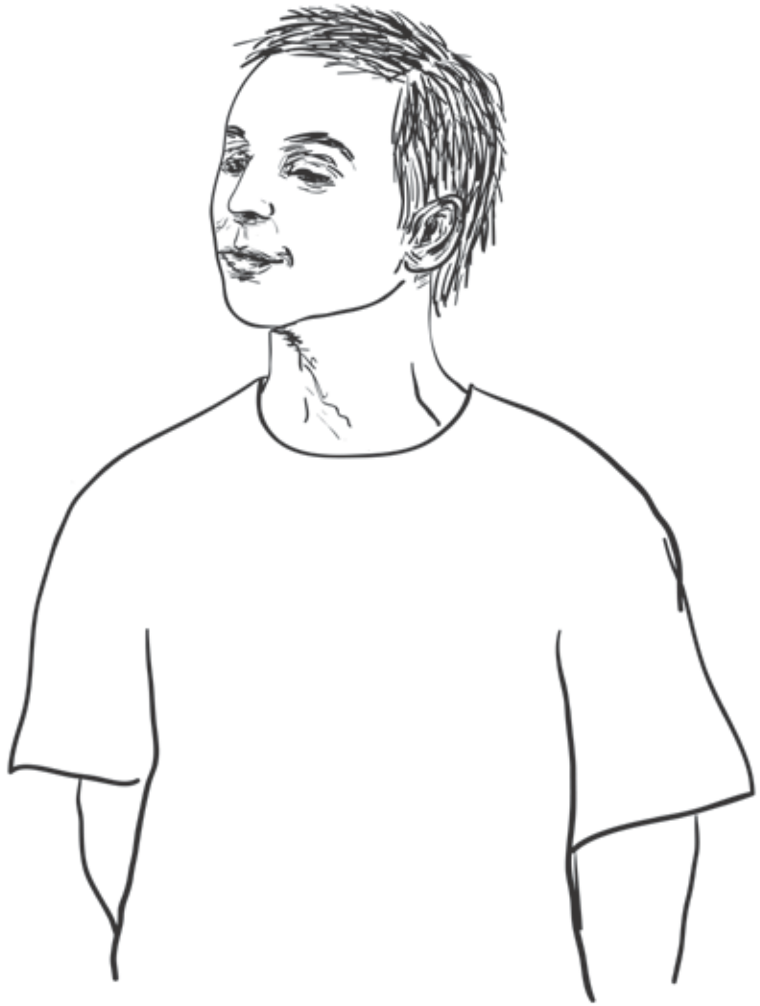

Population density:

$$\rho = \frac{\#\,\text{fish} / m^3}{\#\,\text{max fish} / m^3}$$

Hence: $\rho \in [0,1]$

Model:

$$\rho_{t+1} = 4\rho_t(1-\rho_t)$$

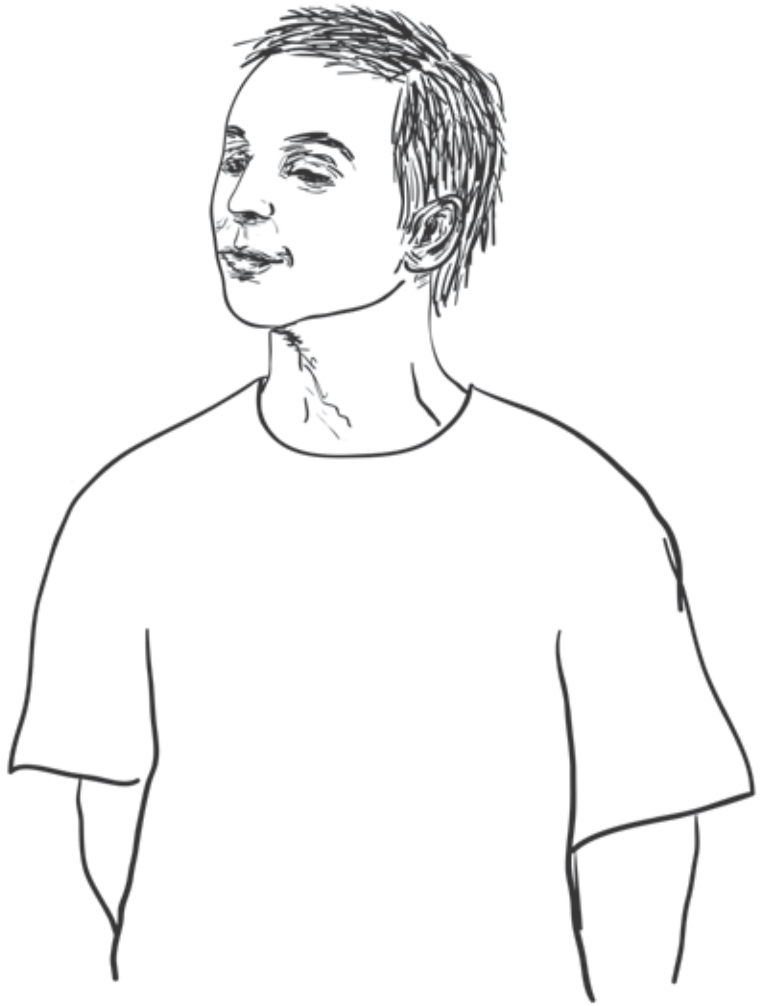

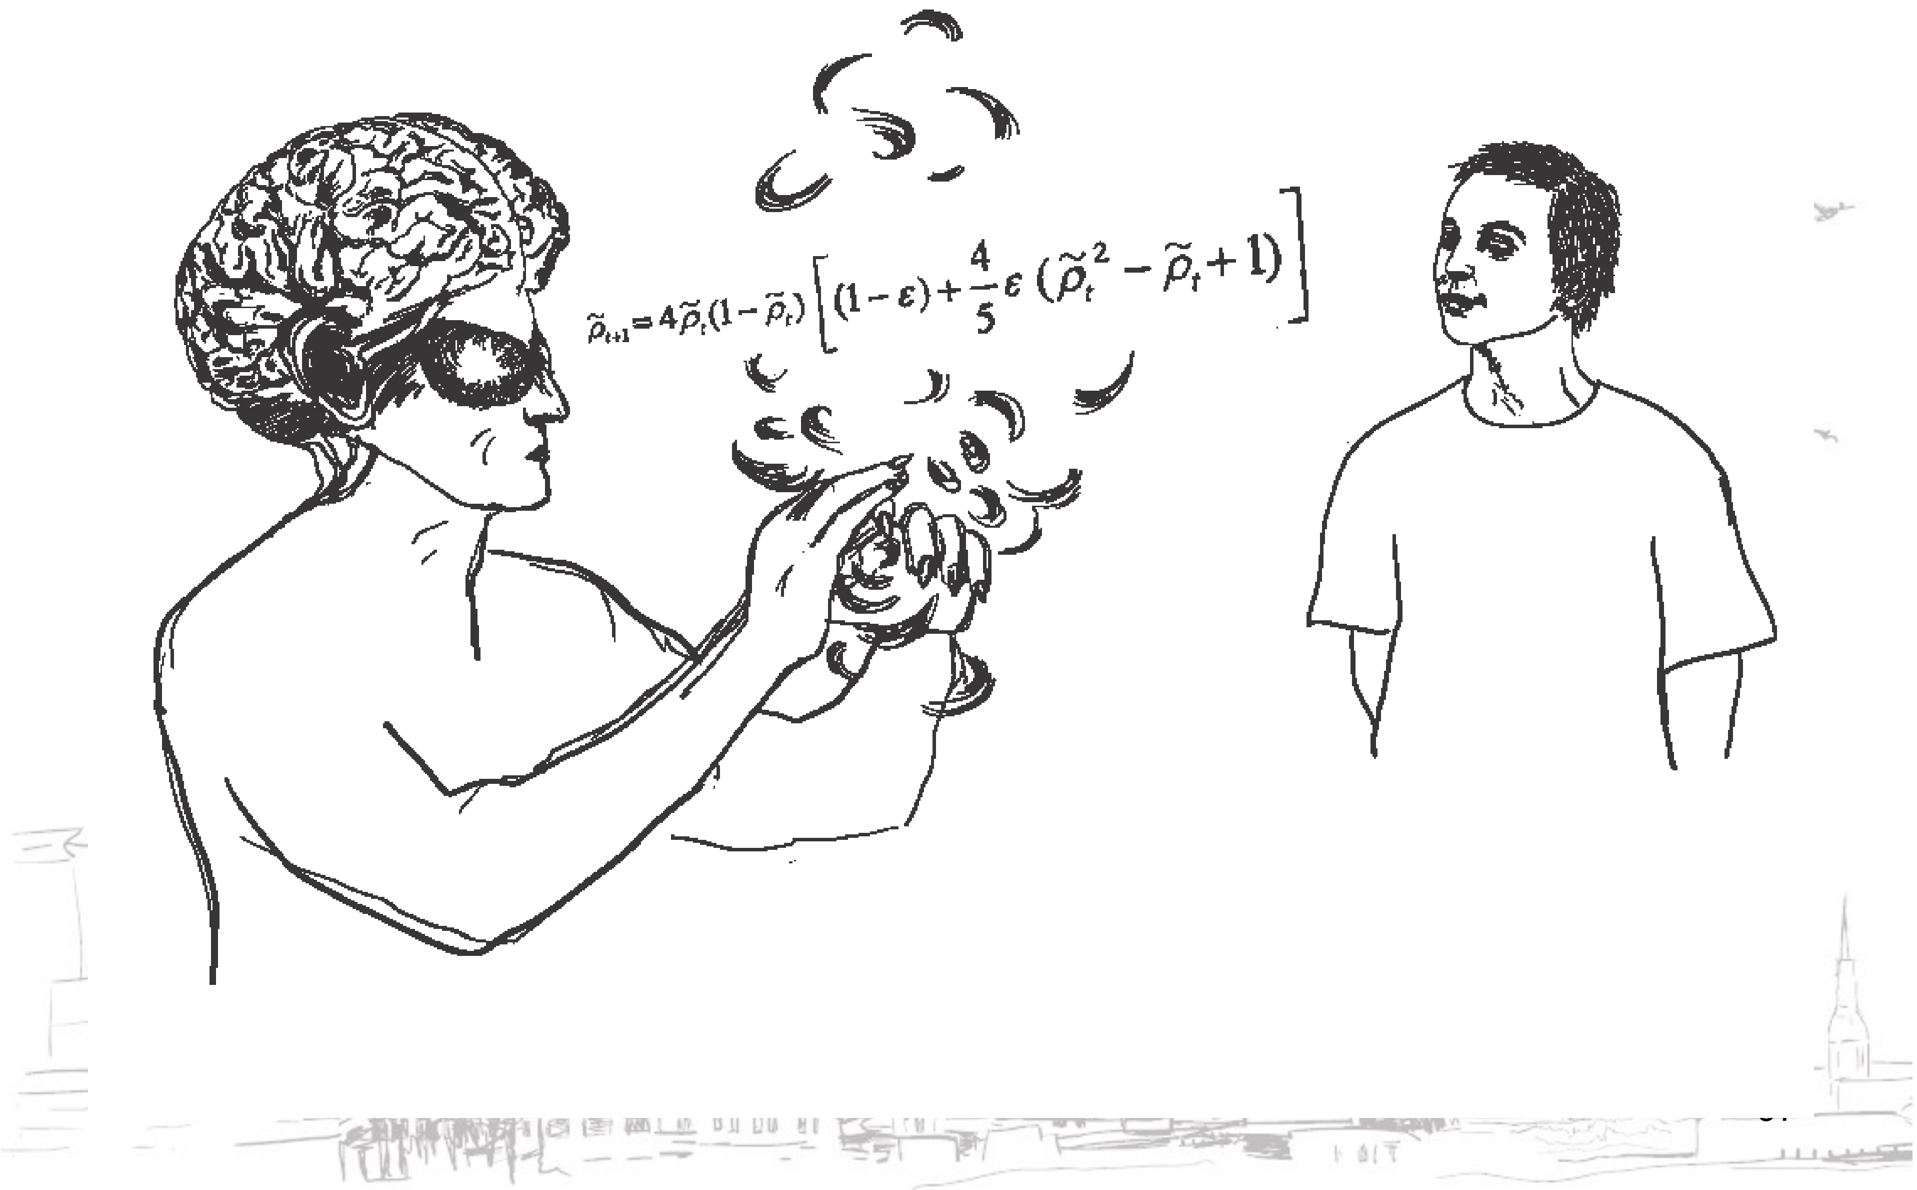
$\tilde{\rho}_{t+1} = 4\tilde{\rho}_t(1-\tilde{\rho}_t)\left[(1-\varepsilon)+\frac{4}{5}\varepsilon\,(\tilde{\rho}_t^{\,2}-\tilde{\rho}_t+1)\right]$

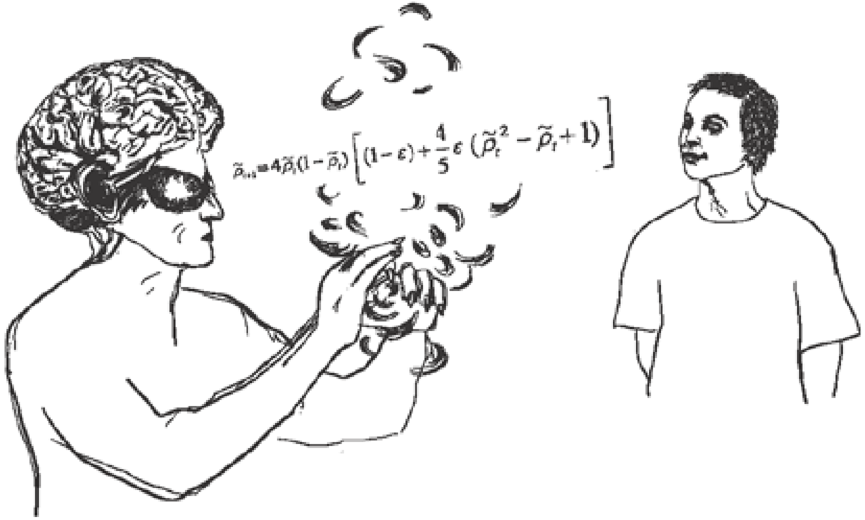

$$\widetilde{\rho}_{t+1} = 4\widetilde{\rho}_t(1-\widetilde{\rho}_t)\left[(1-\varepsilon)+\frac{4}{5}\varepsilon(\widetilde{\rho}_t^{\,2}-\widetilde{\rho}_t+1)\right]$$

where $\varepsilon = 0.1$

The Apprentice remains defiant:

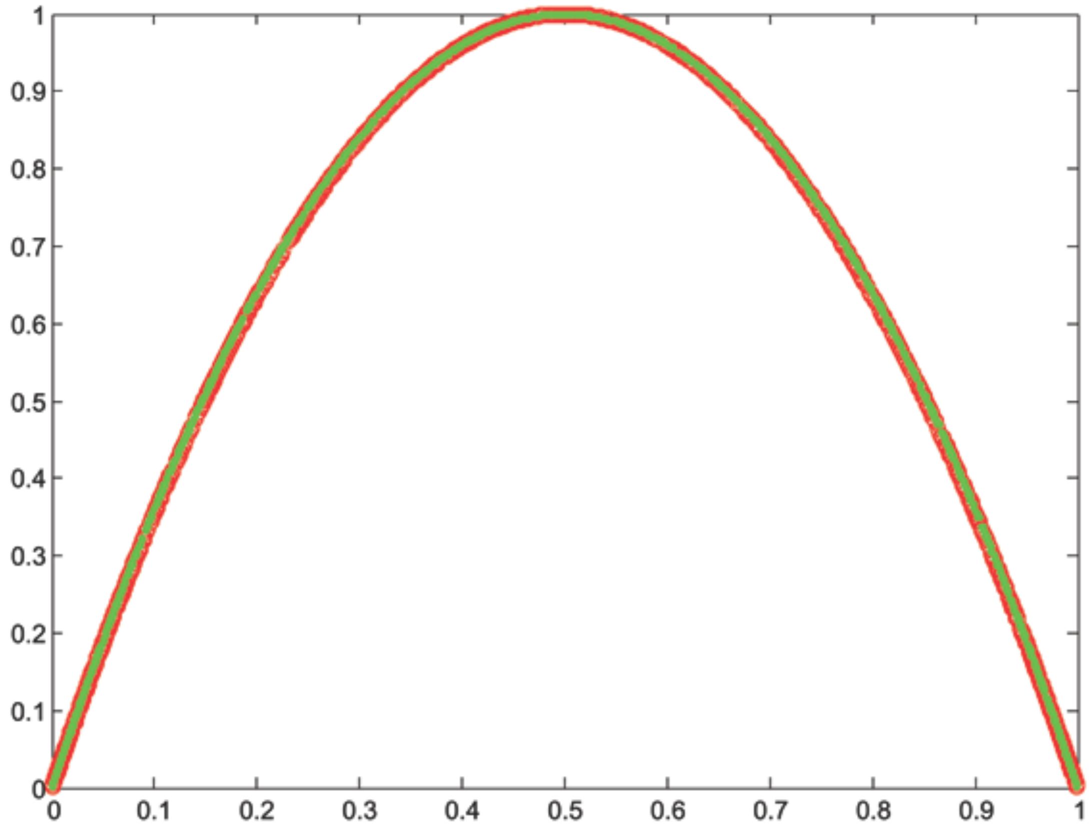

Green – Apprentice and Red - Demon

Mathematically:

$$\rho_{t+1} = 4\rho_t(1-\rho_t)$$ + small perturbation

= $$\tilde{\rho}_{t+1} = 4\tilde{\rho}_t(1-\tilde{\rho}_t)\left[(1-\varepsilon)+\frac{4}{5}\varepsilon(\tilde{\rho}_t^{\,2}-\tilde{\rho}_t+1)\right]$$

Mathematically:

$$\rho_{t+1} = 4\rho_t(1-\rho_t)$$ + small perturbation

$$= \quad \tilde{\rho}_{t+1} = 4\tilde{\rho}_t(1-\tilde{\rho}_t)\left[(1-\varepsilon) + \frac{4}{5}\varepsilon(\tilde{\rho}_t{}^2 - \tilde{\rho}_t + 1)\right]$$

One step error: 0.001

Mathematically:

$$\rho_{t+1} = 4\rho_t(1-\rho_t) \text{ + small perturbation}$$

$$= \quad \tilde{\rho}_{t+1} = 4\tilde{\rho}_t(1-\tilde{\rho}_t)\left[(1-\varepsilon)+\frac{4}{5}\varepsilon(\tilde{\rho}_t^{\,2}-\tilde{\rho}_t+1)\right]$$

One step error: 0.001

Closeness-to-goodness link: this is close enough and predictions are reliable.

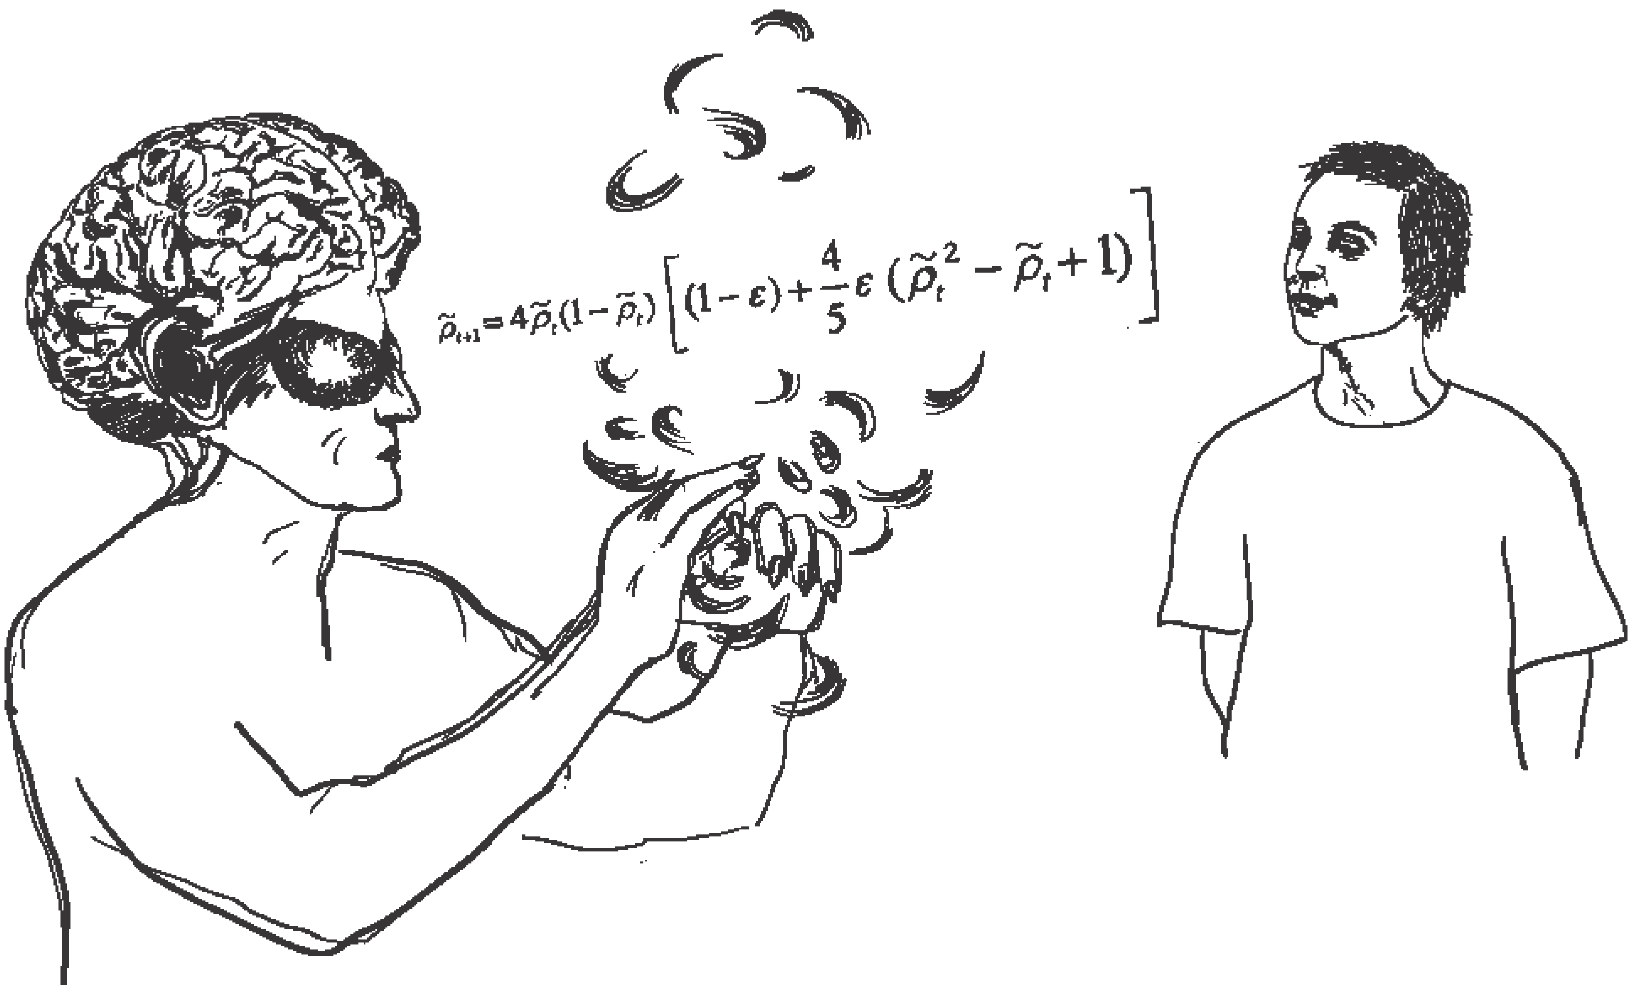
$\tilde{\rho}_{t+1} = 4\tilde{\rho}_t(1-\tilde{\rho}_t)\left[(1-\varepsilon)+\frac{4}{5}\varepsilon\,(\tilde{\rho}_t^{\,2}-\tilde{\rho}_t+1)\right]$


$p_t(x) := \phi_t[p_0(x)]$
ϕ_t
$\rho_{t+1} = 4\rho_t(1-\rho_t)$
$\Delta_U := U_T - U_A$
(X, ϕ_t, μ)

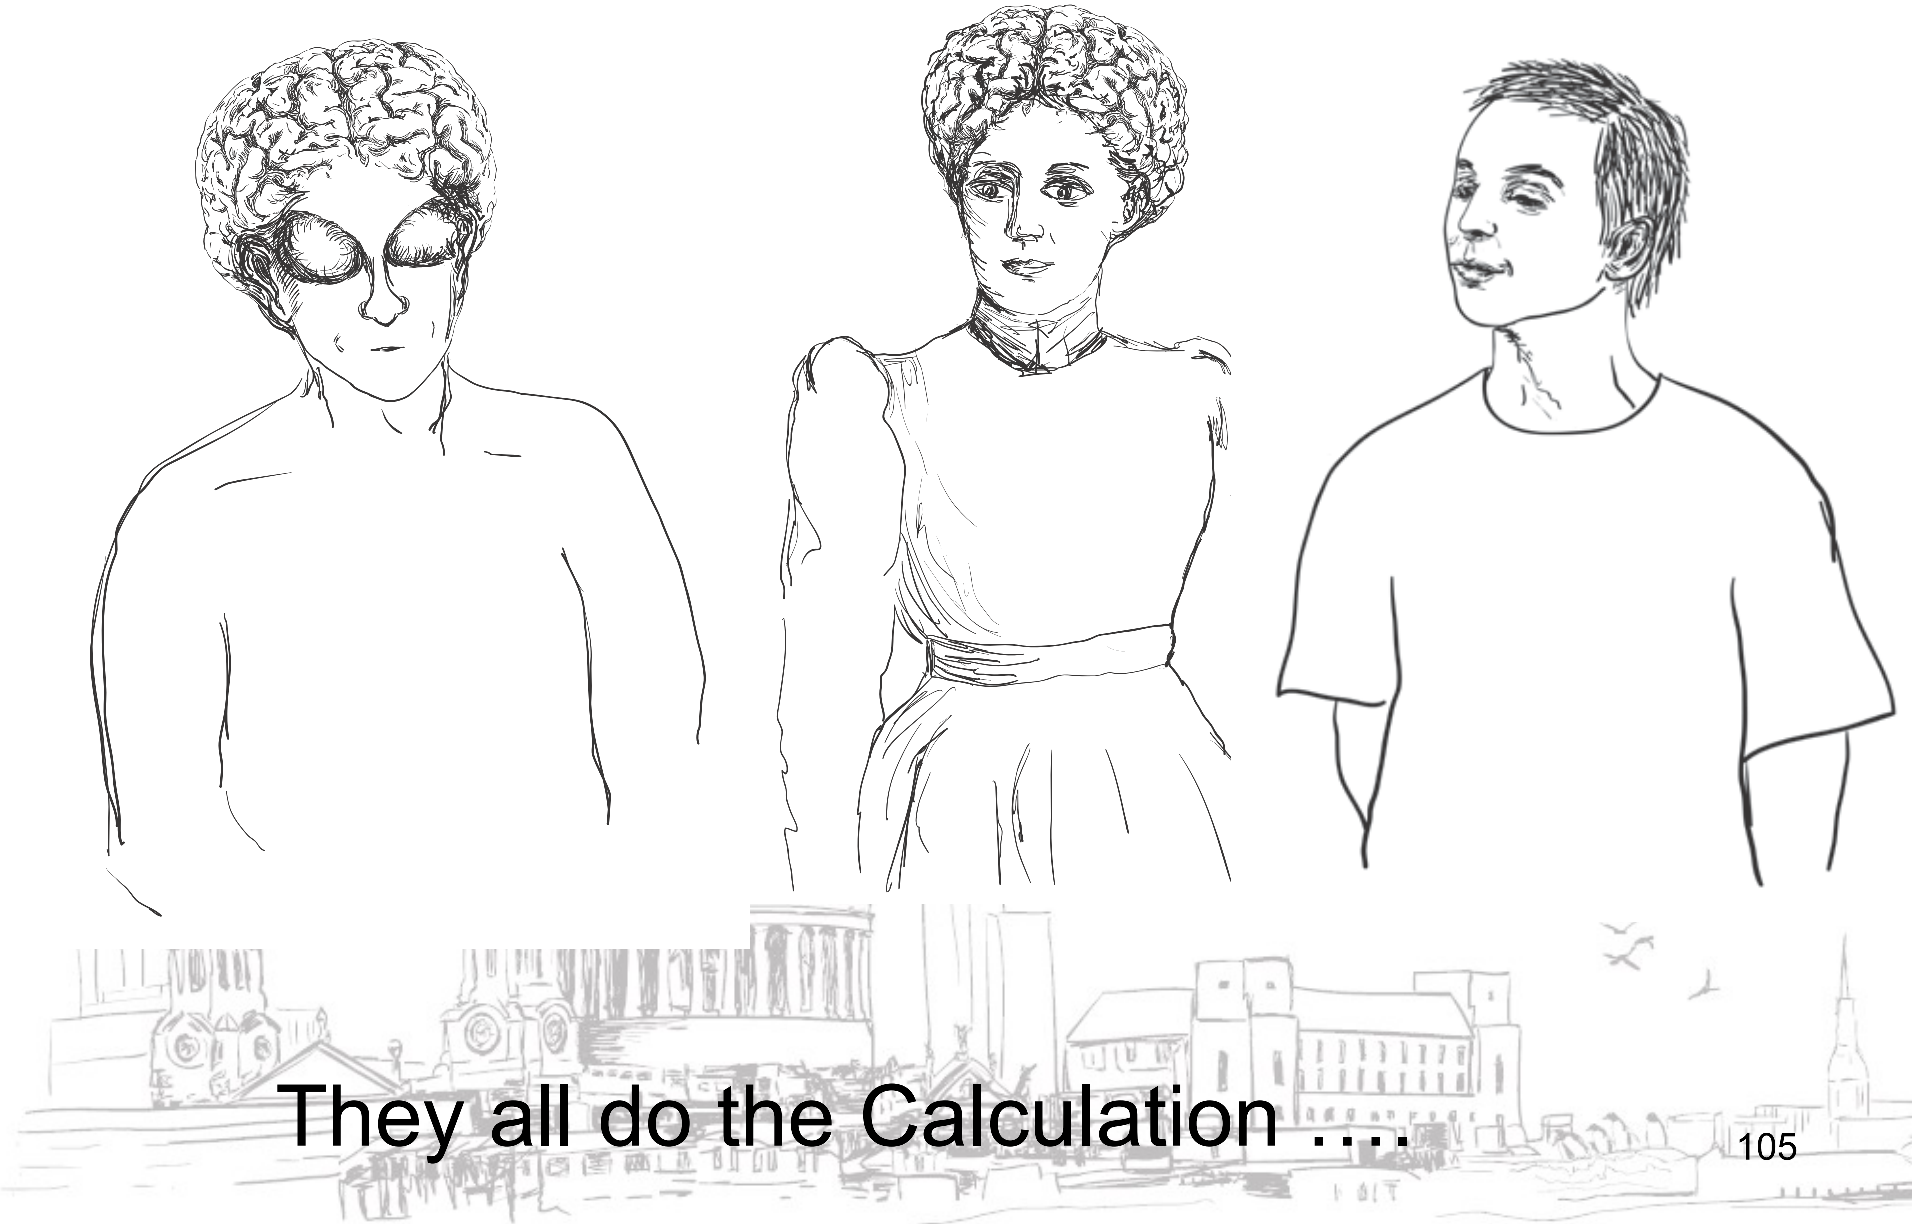

They all do the Calculation ….

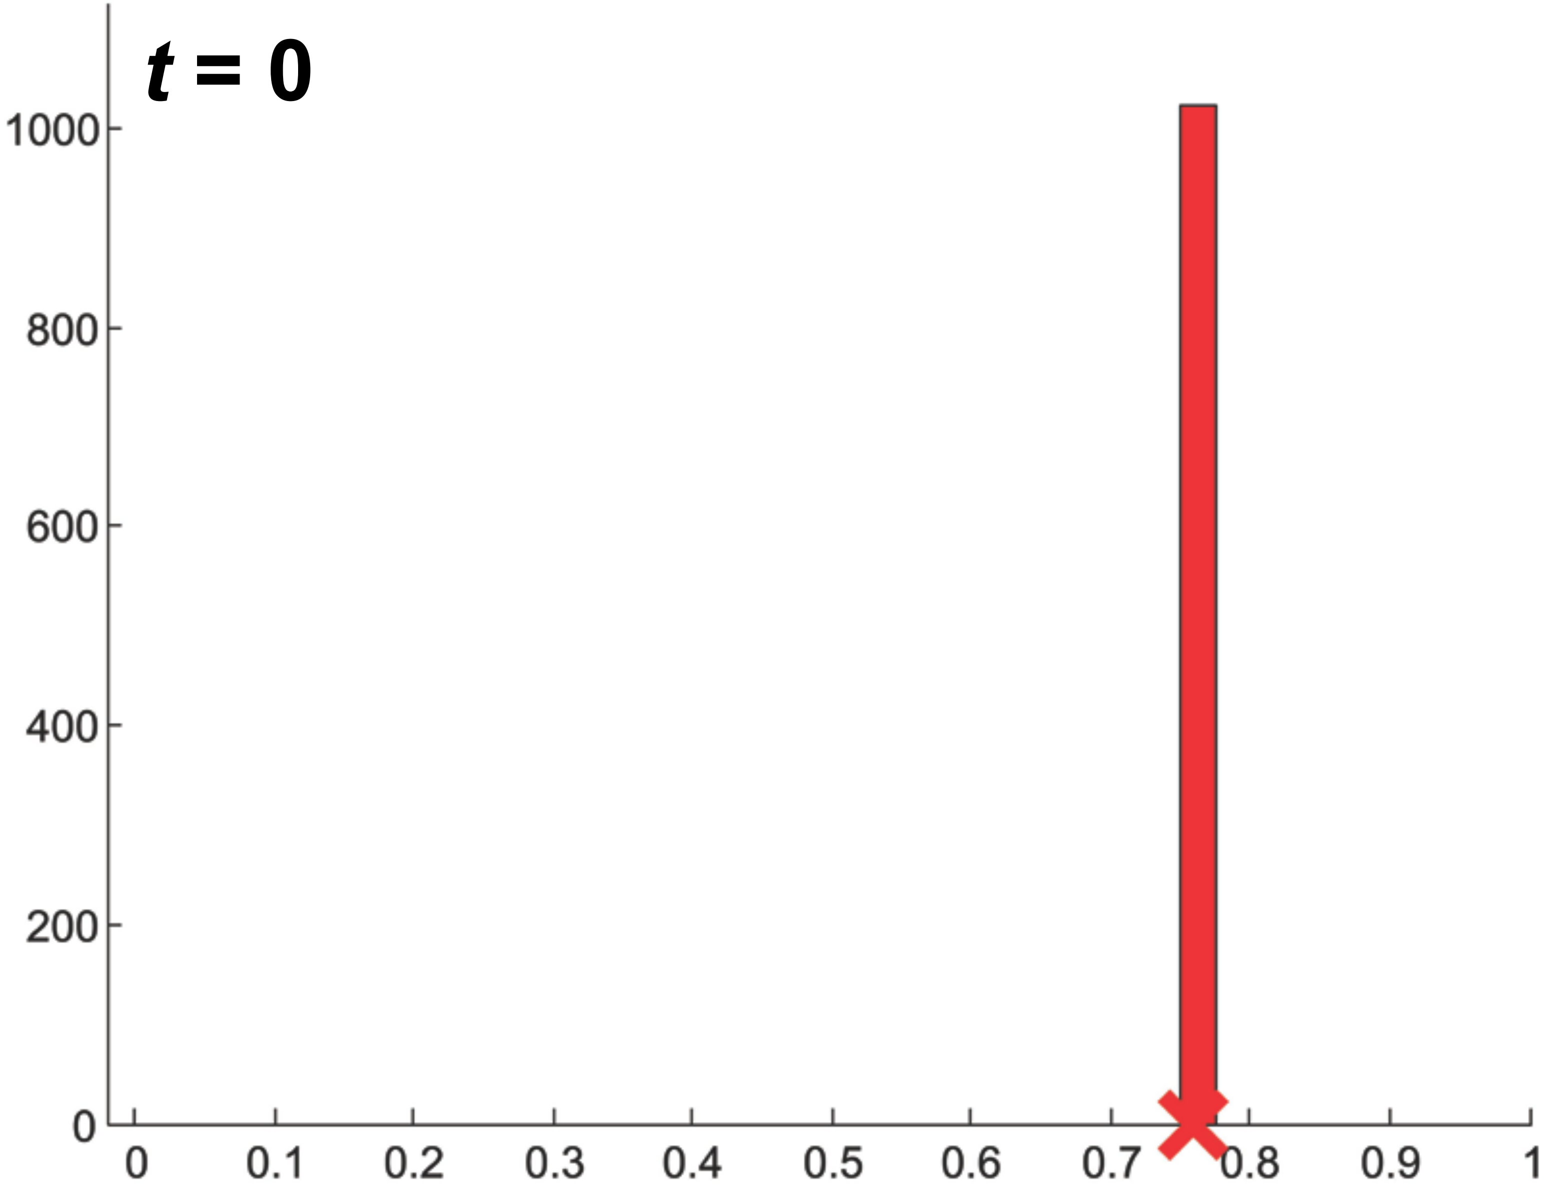
t = 0
1000
800
600
400
200
0
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

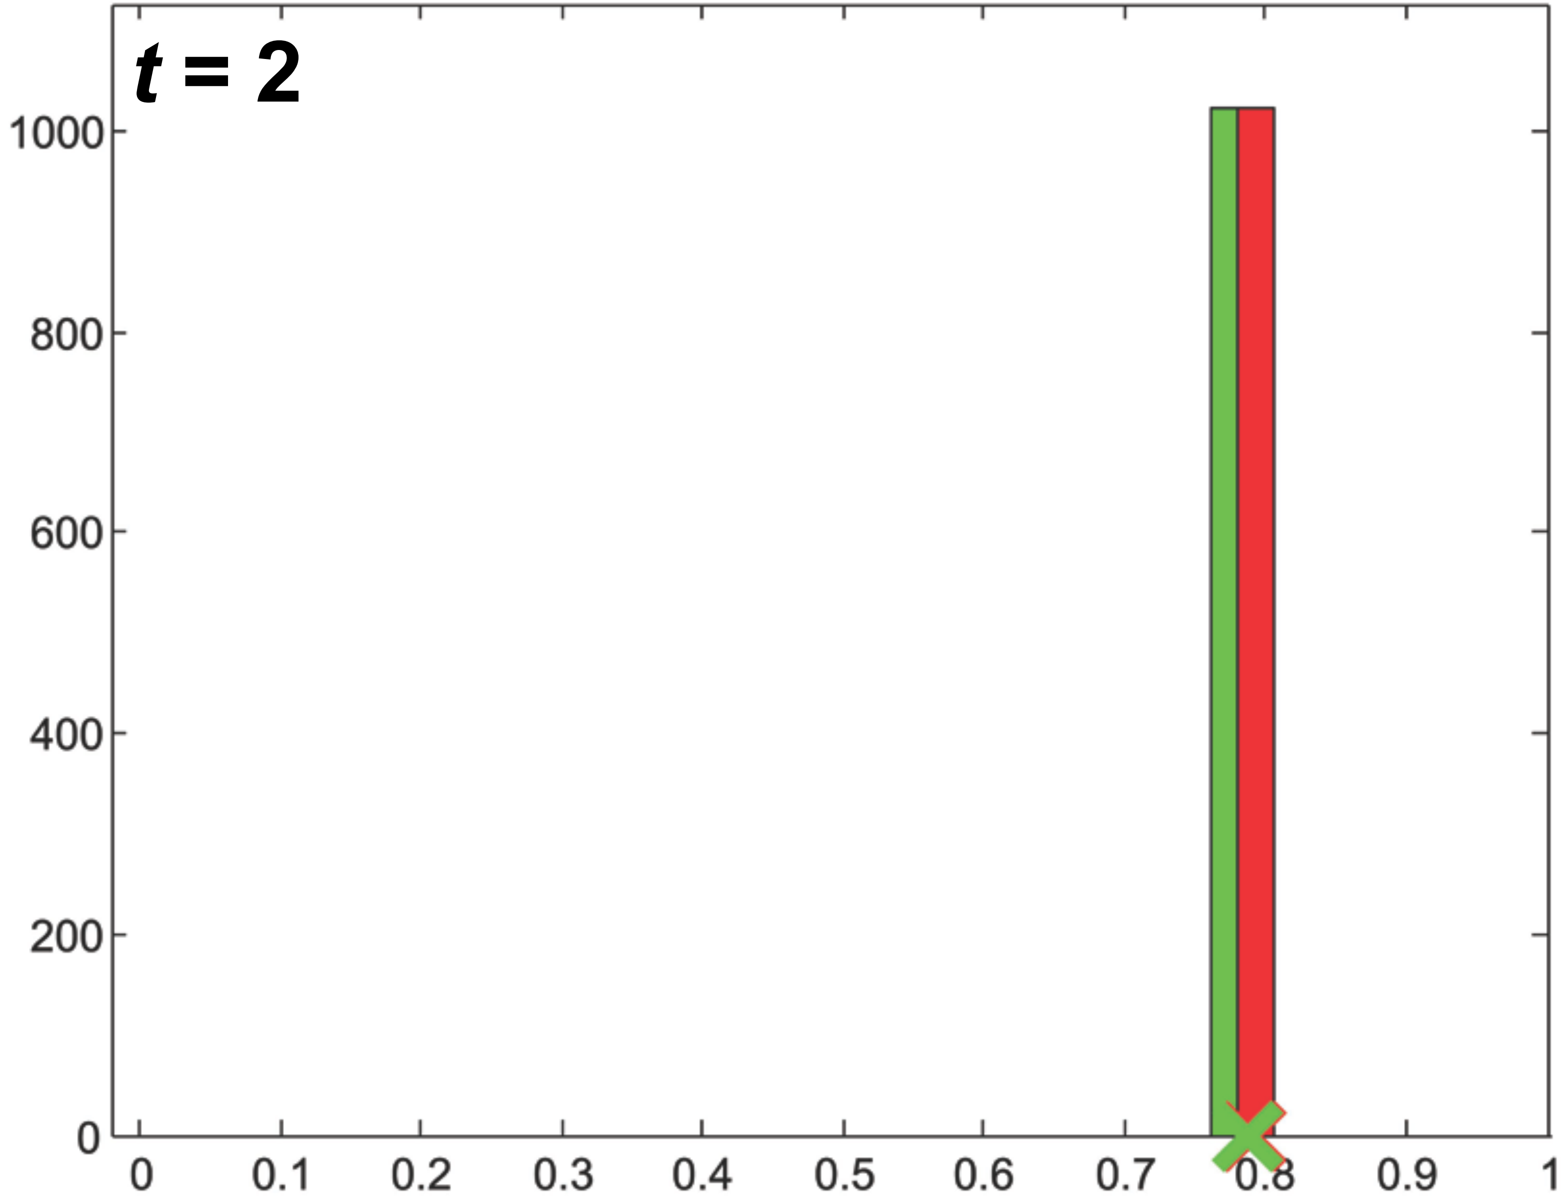
t = 2
1000
800
600
400
200
0
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

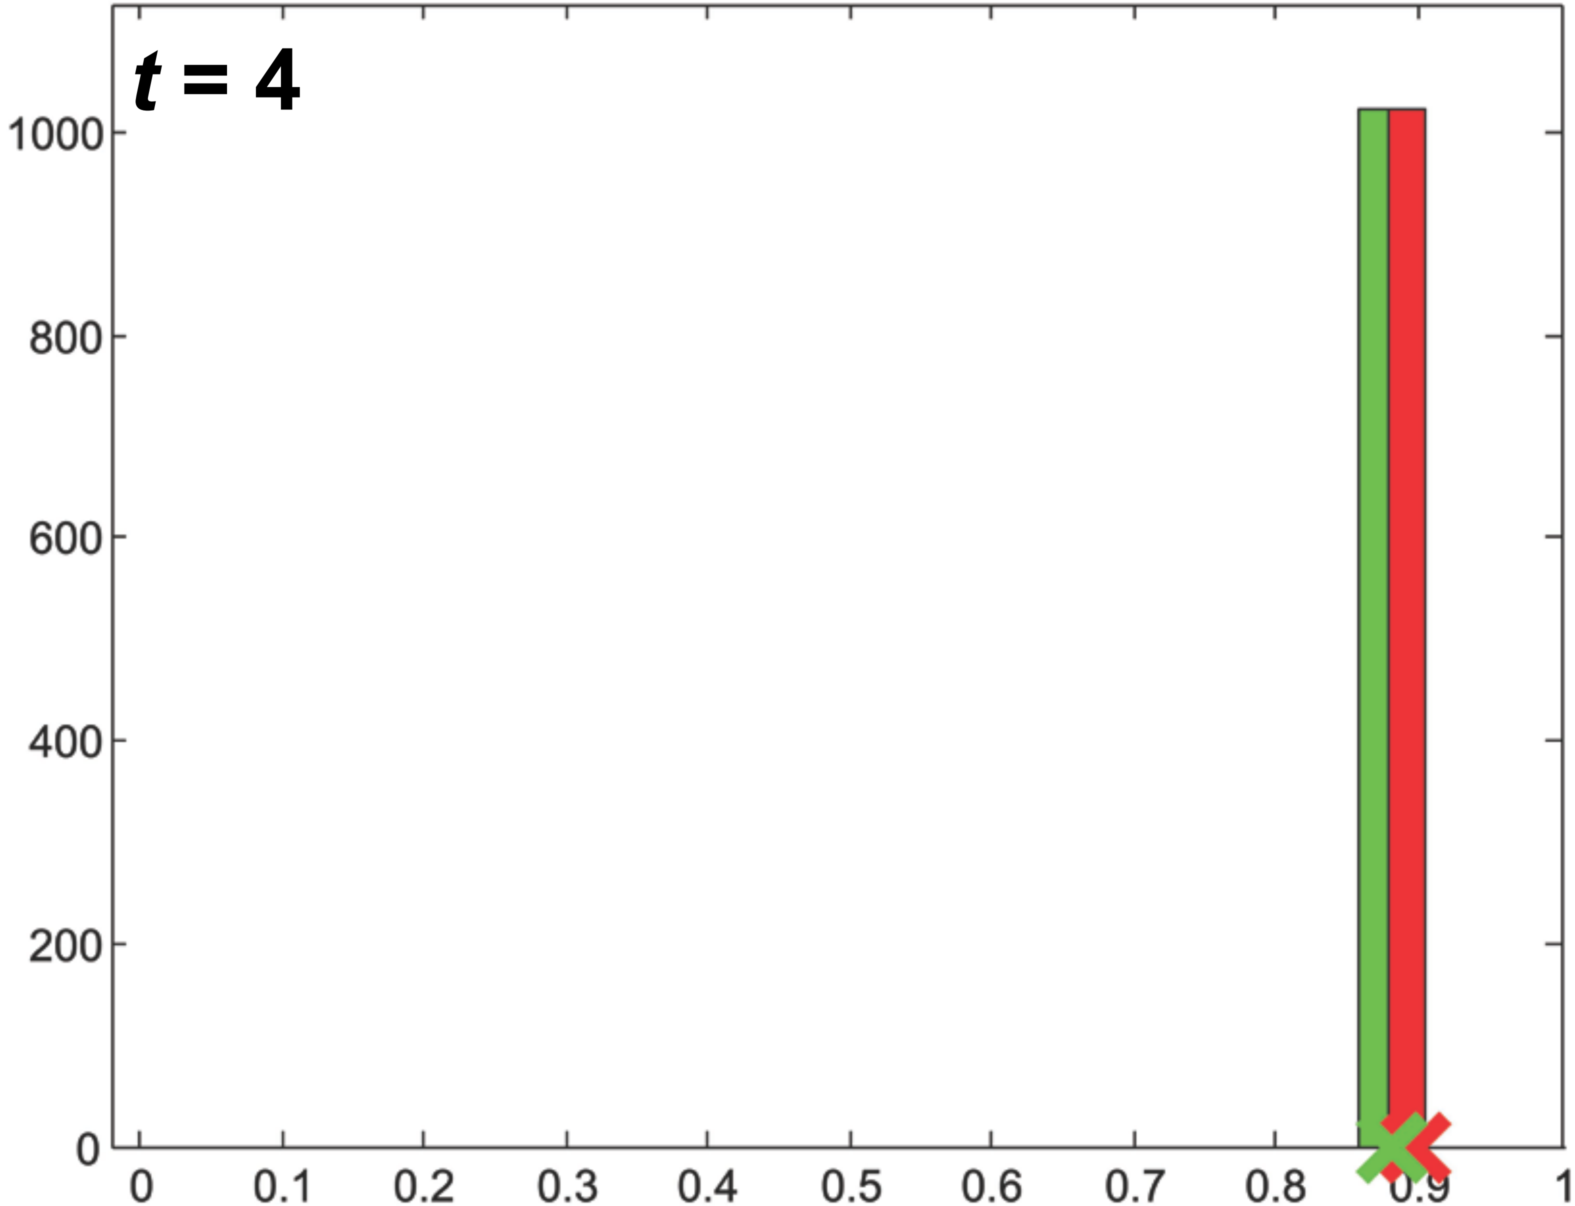

t = 4
1000
800
600
400
200
0
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

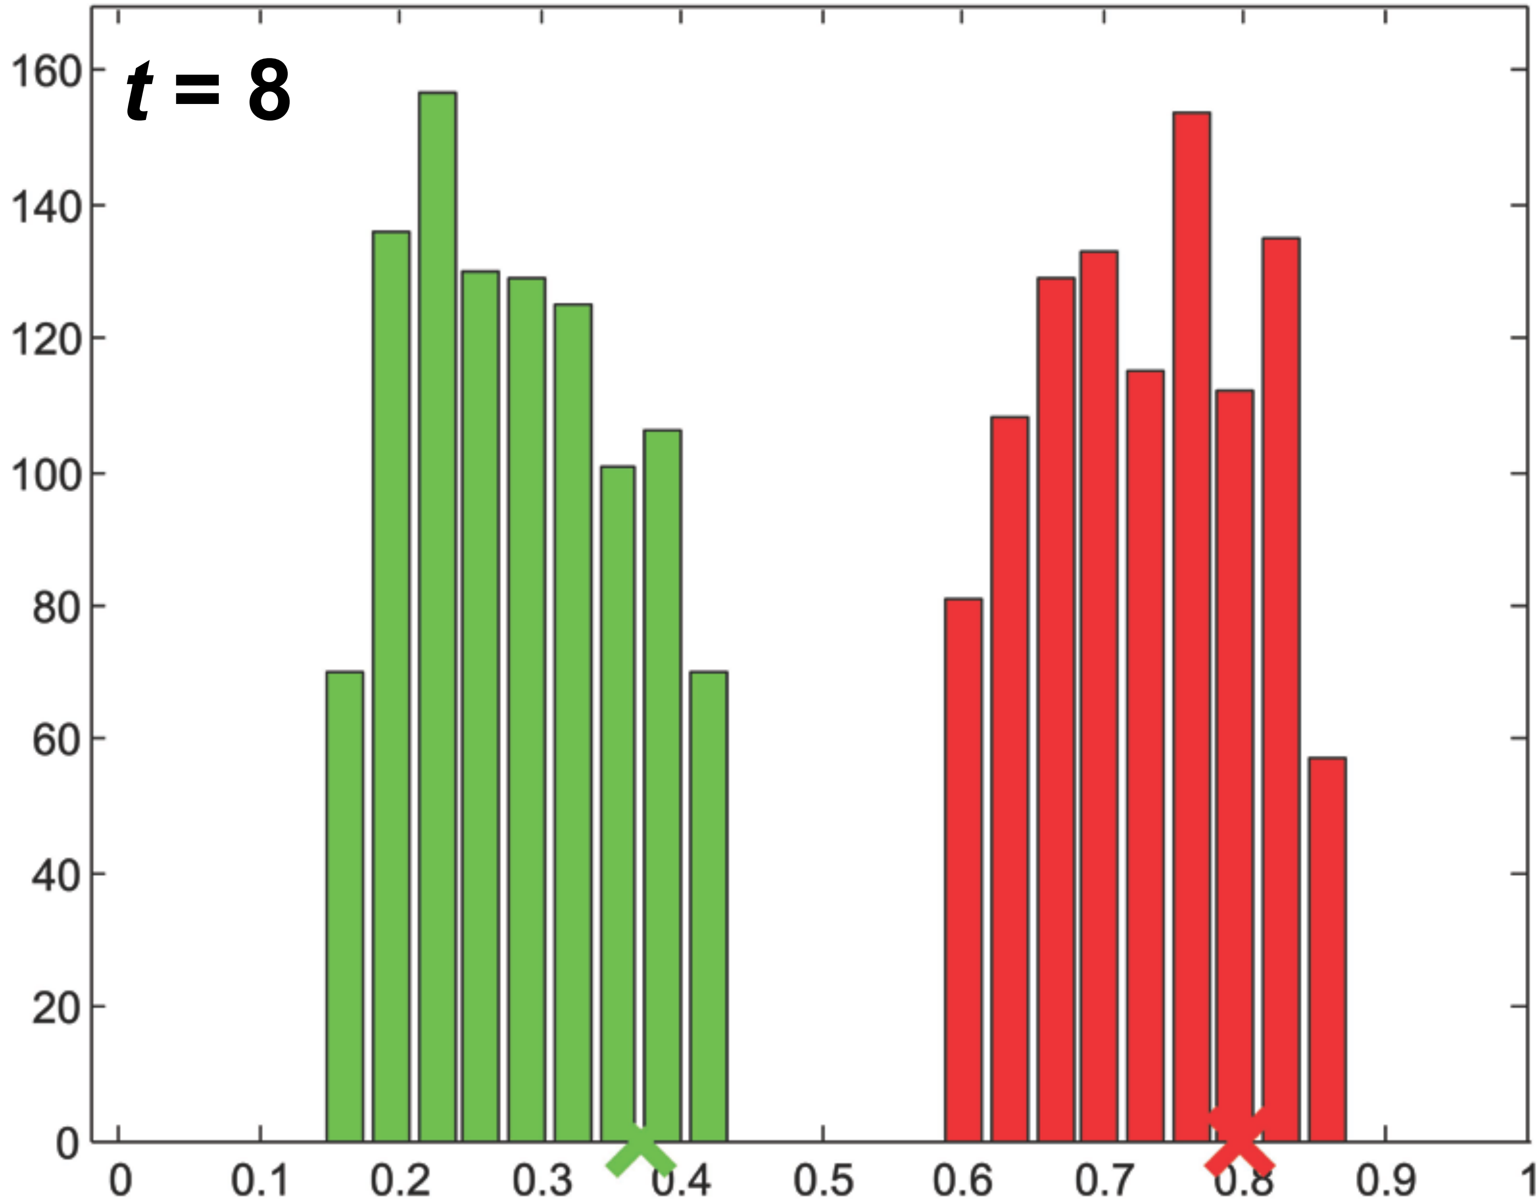
t = 8
0
20
40
60
80
100
120
140
160
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

If you use your model to offer odds on certain events, you get it completely wrong!

- You regard things that never happen as very likely.
- You regard things that happen very often as unlikely.

Question: is this a special case?

Relative Entropy of 2048 initial distributions (t=8)

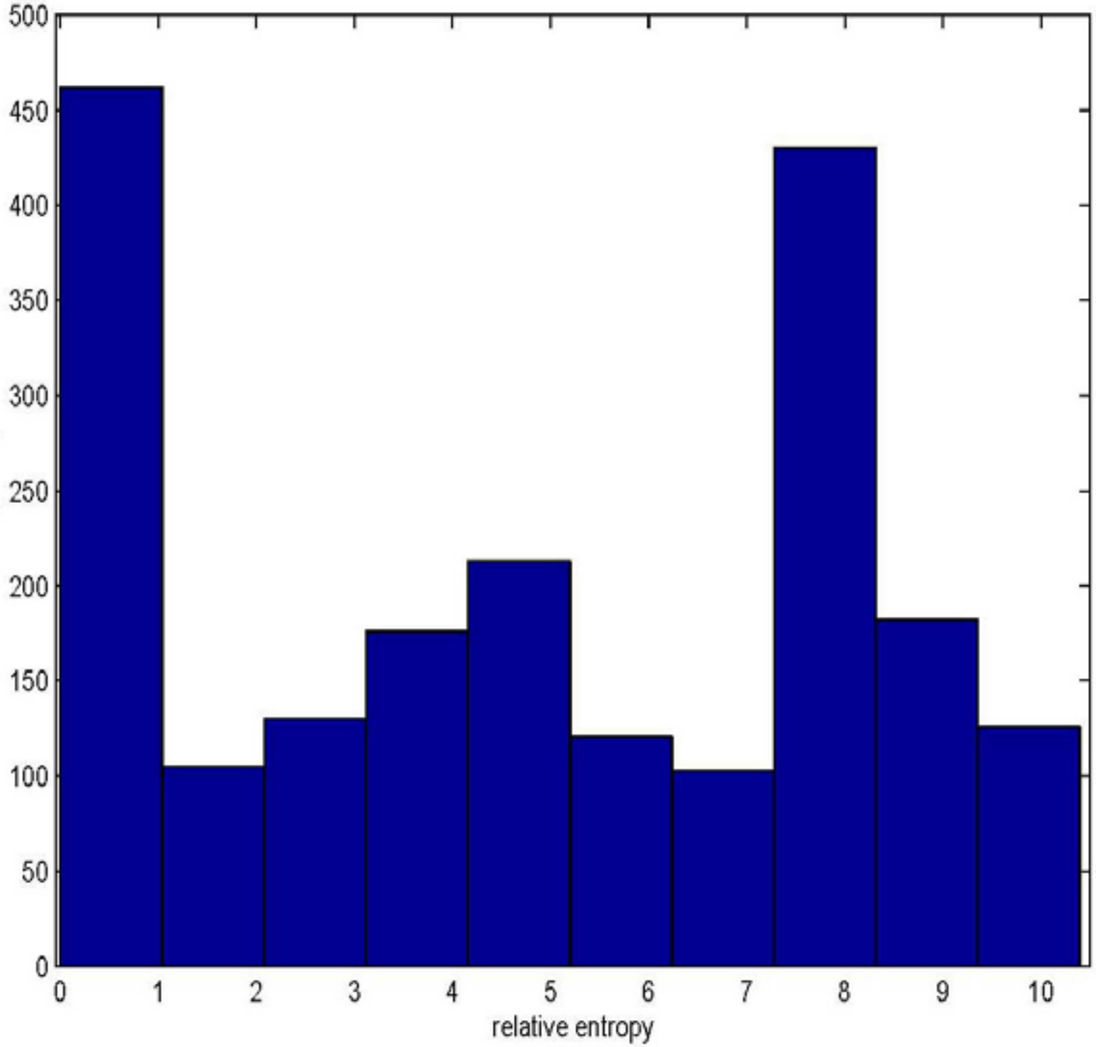

The Pond Casino

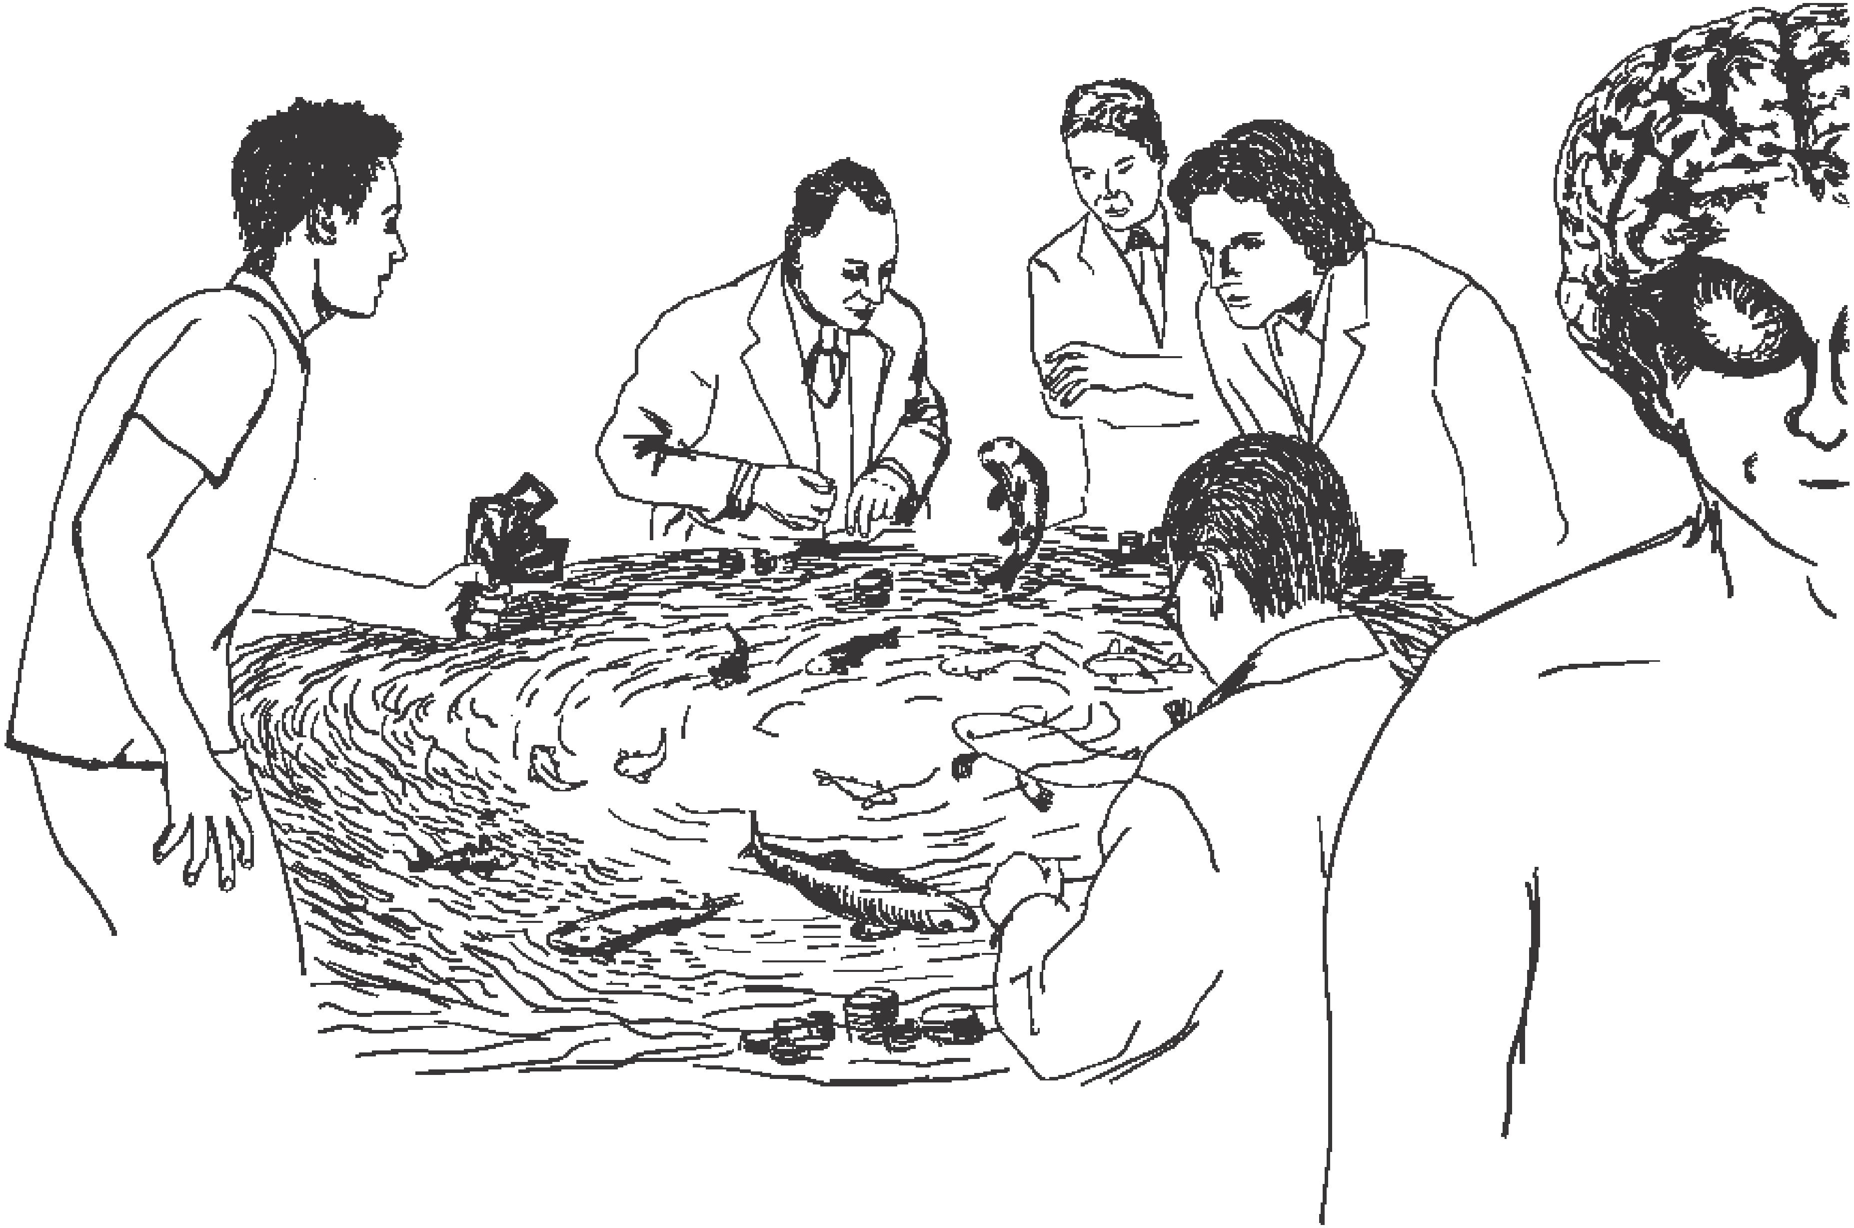

Nine punters with £1000 each.

In every round they bet 10% of their wealth on events with probability in the interval:

1st Punter: [1/2, 1]

2nd Punter: [1/4), 1/2)

…

9th Punter: [0, 1/256)

How are they doing?

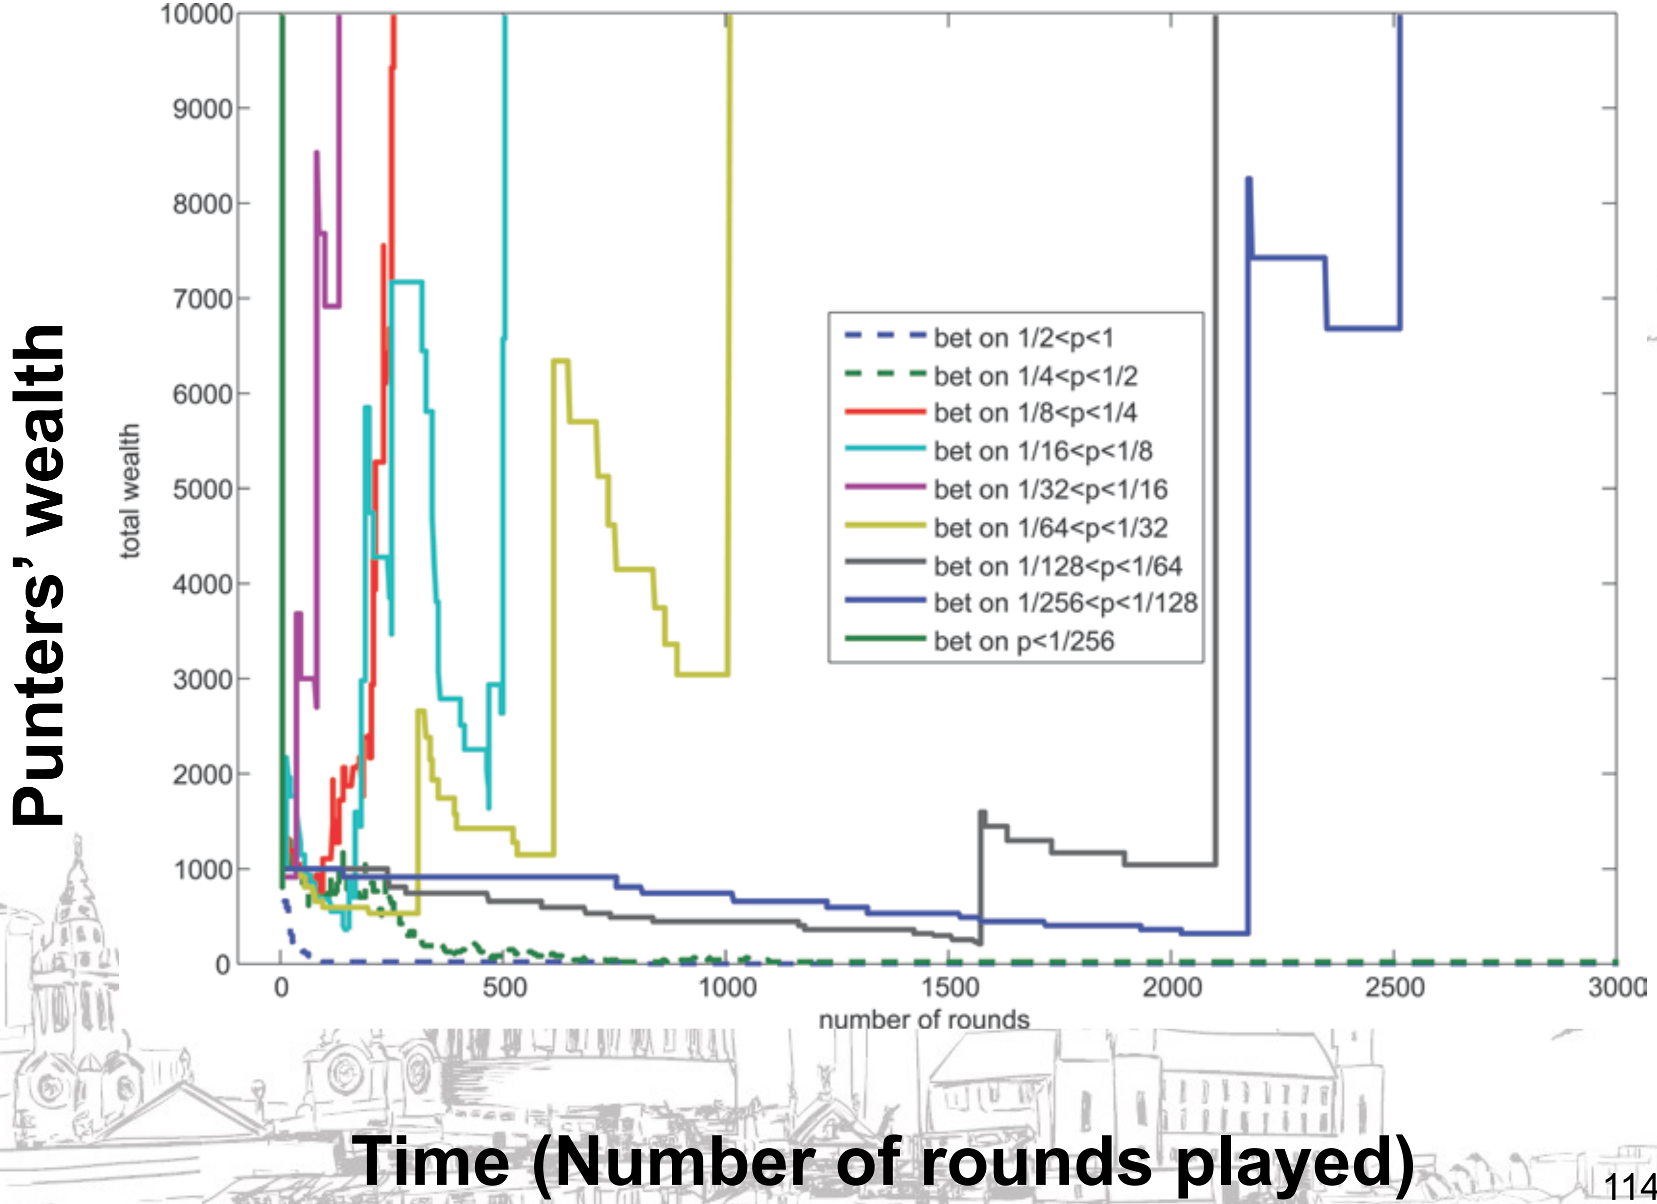

Punters' wealth
total wealth
10000
9000
8000
7000
6000
5000
4000
3000
2000
1000
0
0
500
1000
1500
2000
2500
3000
number of rounds
bet on 1/2<p<1
bet on 1/4<p<1/2
bet on 1/8<p<1/4
bet on 1/16<p<1/8
bet on 1/32<p<1/16
bet on 1/64<p<1/32
bet on 1/128<p<1/64
bet on 1/256<p<1/128
bet on p<1/256
Time (Number of rounds played)

Result:

- 7 out of the 9 punters make enormous gains!
- The casino runs up huge losses.

→ Insurance companies …

But: is this just a bad “bad luck event”?

Time to bust for 2048 casinos:

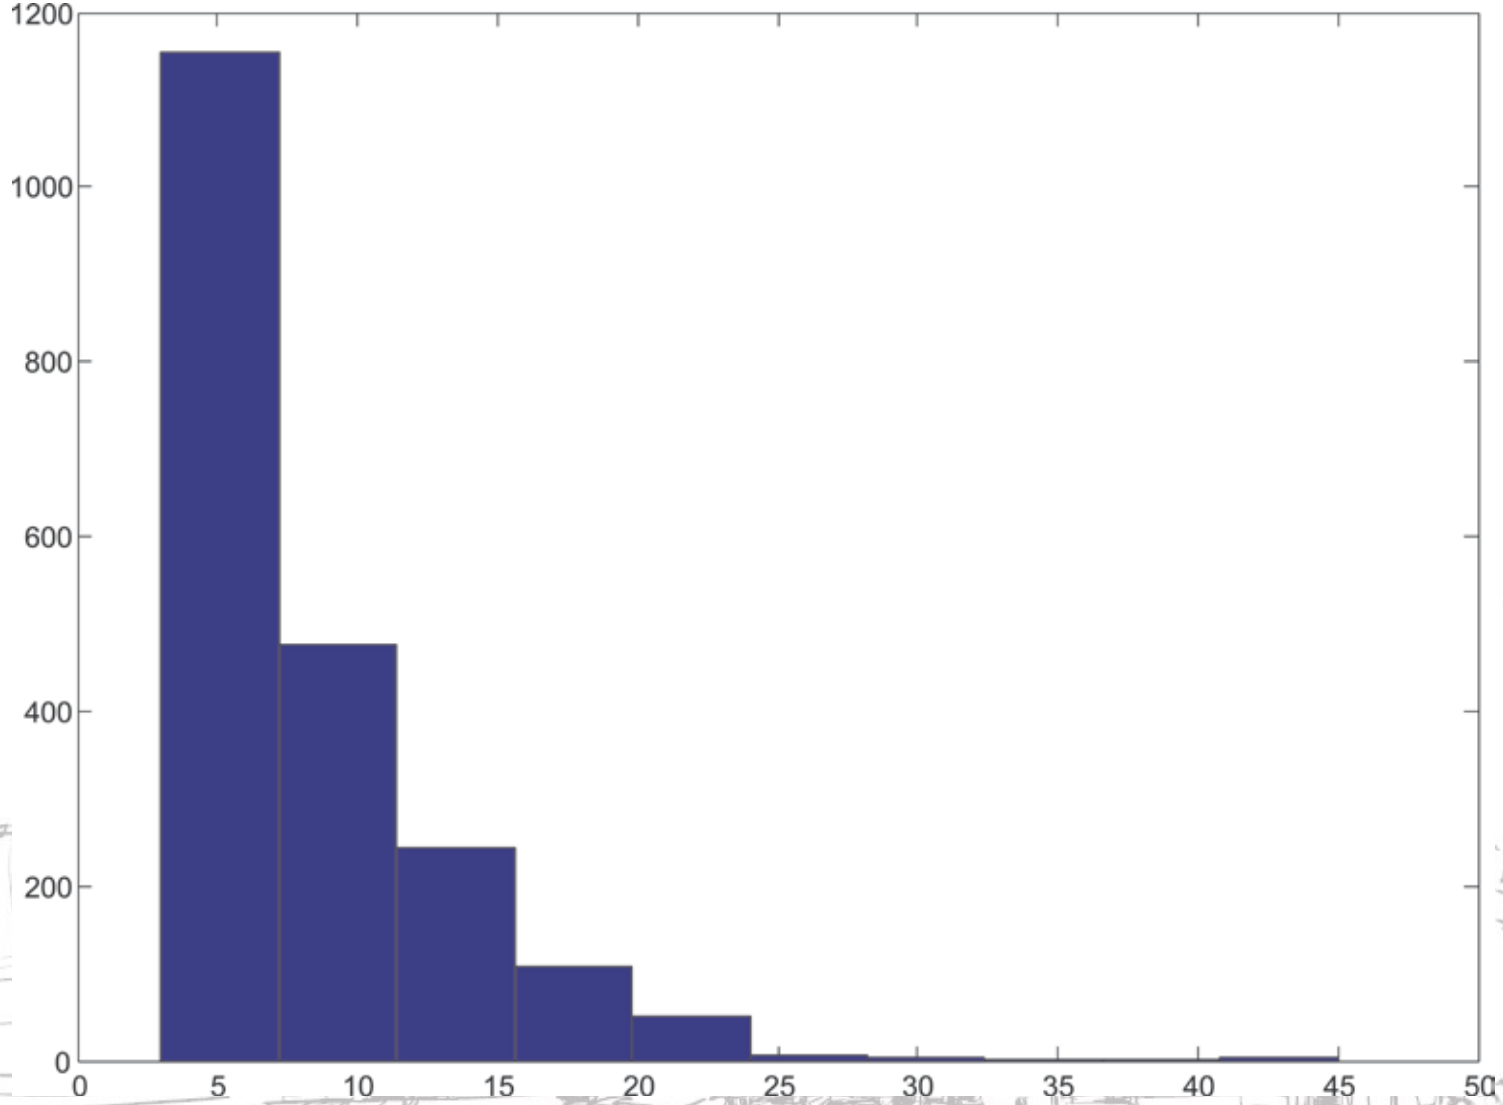

Conclusion:

Even though the model is very close to the truth, it provides ruinous predictions!

Hence: If chaotic models have even the slightest model error, their capacity to make meaningful (and policy relevant!) probabilistic forecasts is lost.

Conclusion:

Even though the model is very close to the truth, it provides ruinous predictions!

Hence: If chaotic models have even the slightest model error, their capacity to make meaningful (and policy relevant!) probabilistic forecasts is lost.

The closeness-to-goodness link is wrong!

The failure of the closeness-to-goodness link gives raise to the ***hawkmoth effect***: the smallest deviation in model structure leads to completely different results, both for deterministic *and* probabilistic forecasts.

Therefore: **an Initial Condition Ensemble and the closeness to goodness link are not an adequate means to deal with <u>structural model error</u>.**

Or: butterflies are pretty; hawkmoths are ugly.

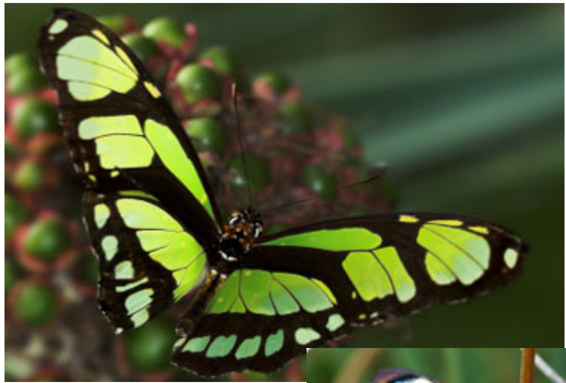

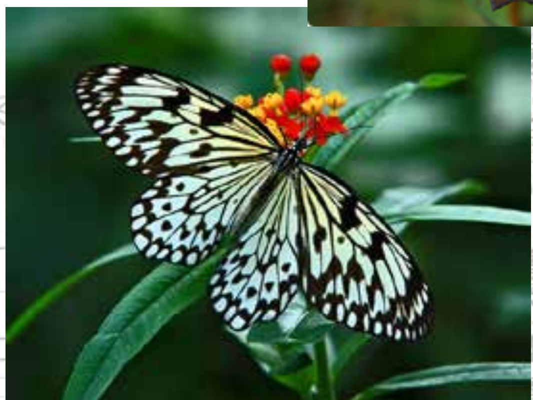

Reinventing the wheel?

Feigenbaum's classical discussion:

$$\rho_{t+1} = \alpha \rho_t (1 - \rho_t)$$

Parameter:

$$\alpha \in [0, 4]$$

Time series for different parameter values:

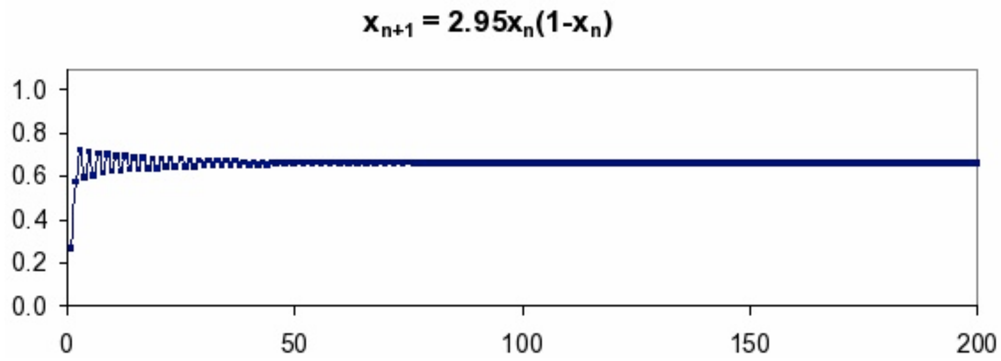

$$\alpha = 2.95$$

Time series for different parameter values:

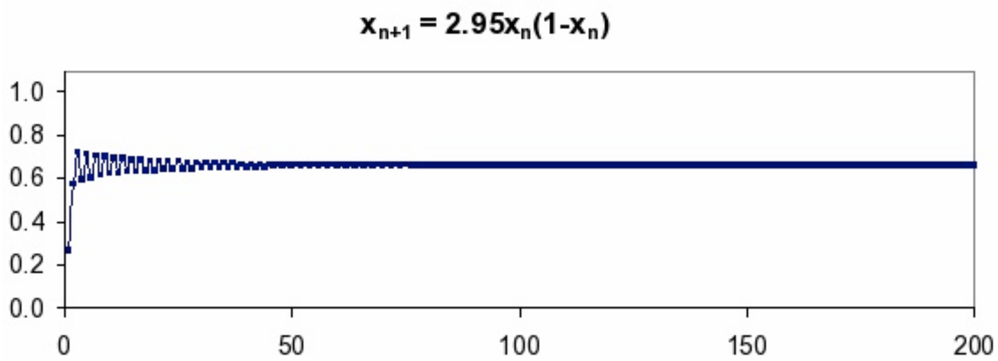

$$\alpha = 2.95$$

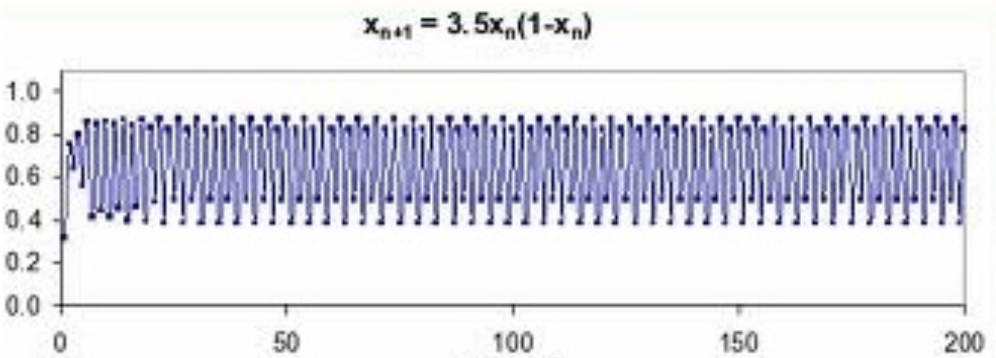

$$\alpha = 3.5$$

Time series for different parameter values:

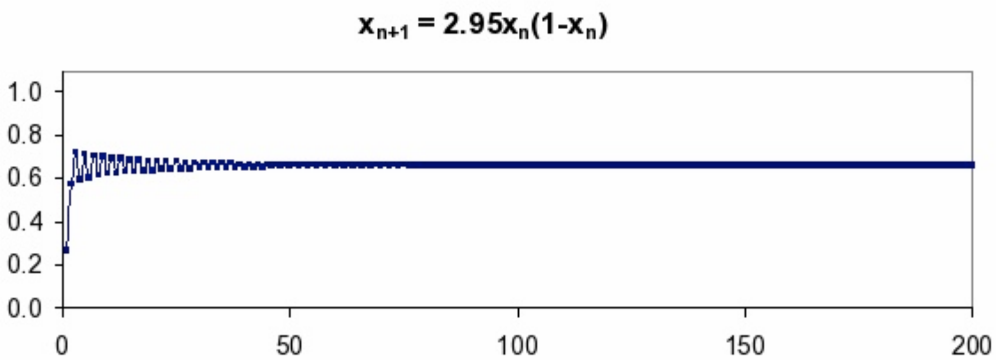

$\alpha = 2.95$

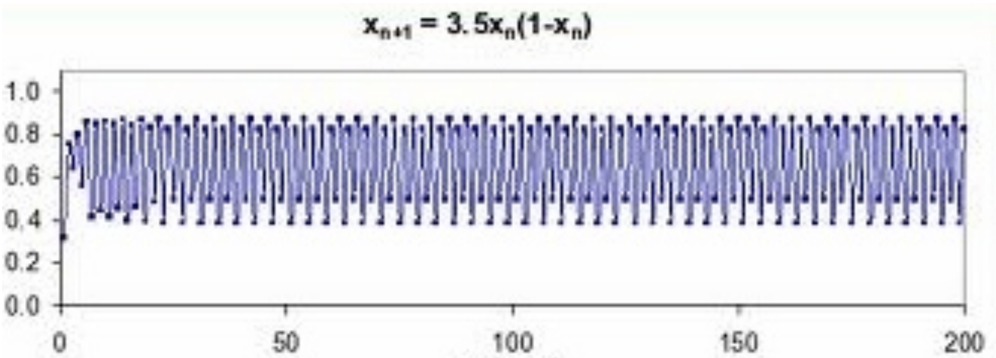

$\alpha = 3.5$

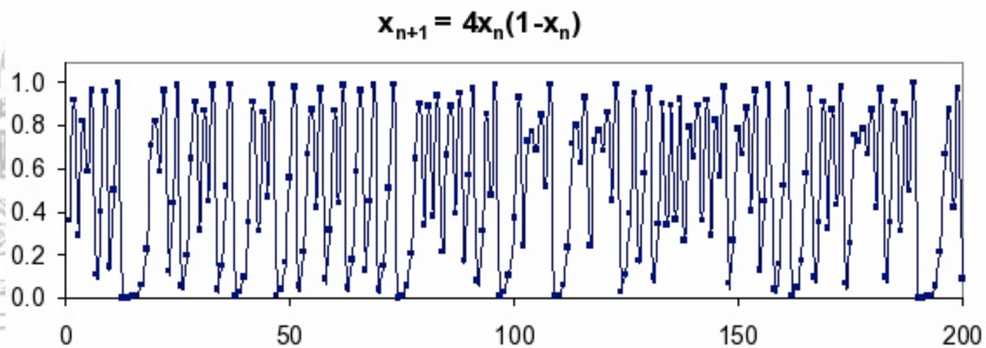

$\alpha = 4$

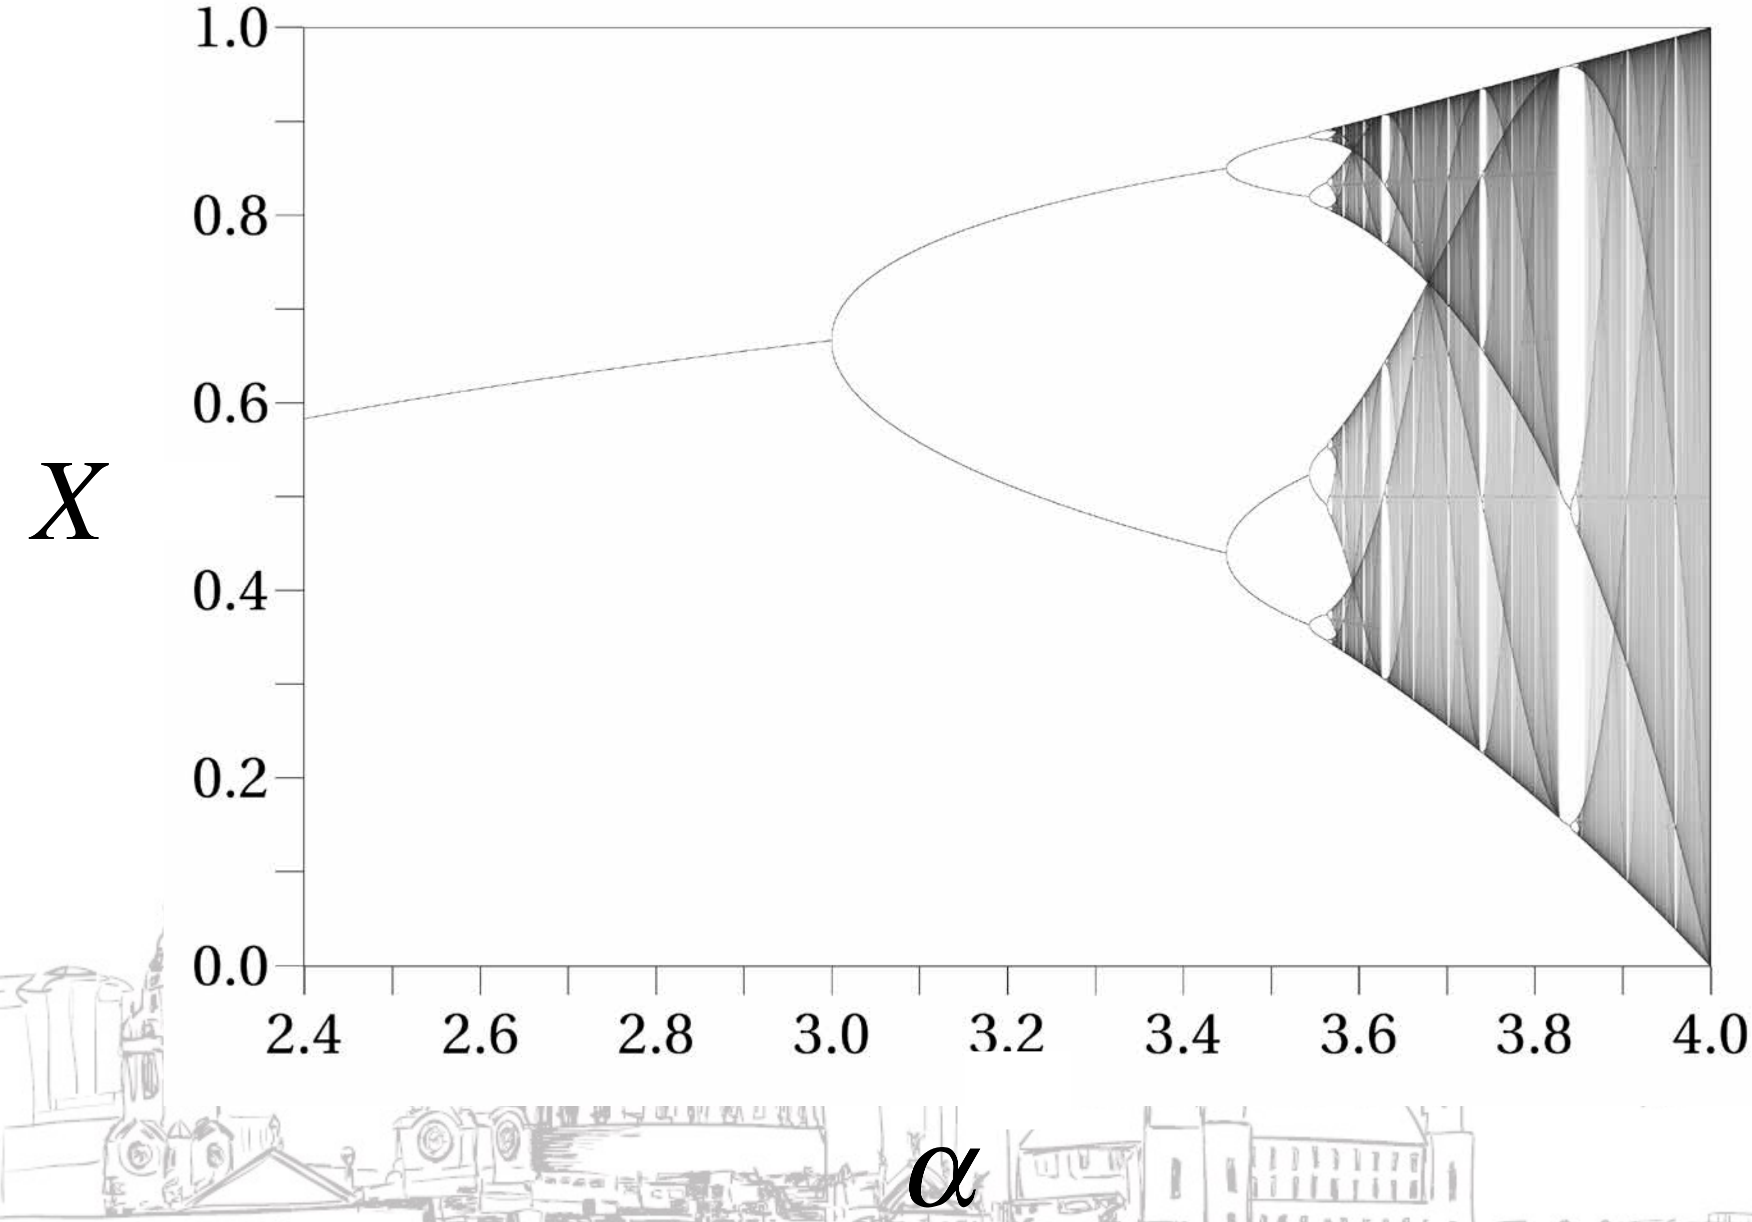
1.0
0.8
0.6
0.4
0.2
0.0
X
2.4
2.6
2.8
3.0
3.2
3.4
3.6
3.8
4.0
α

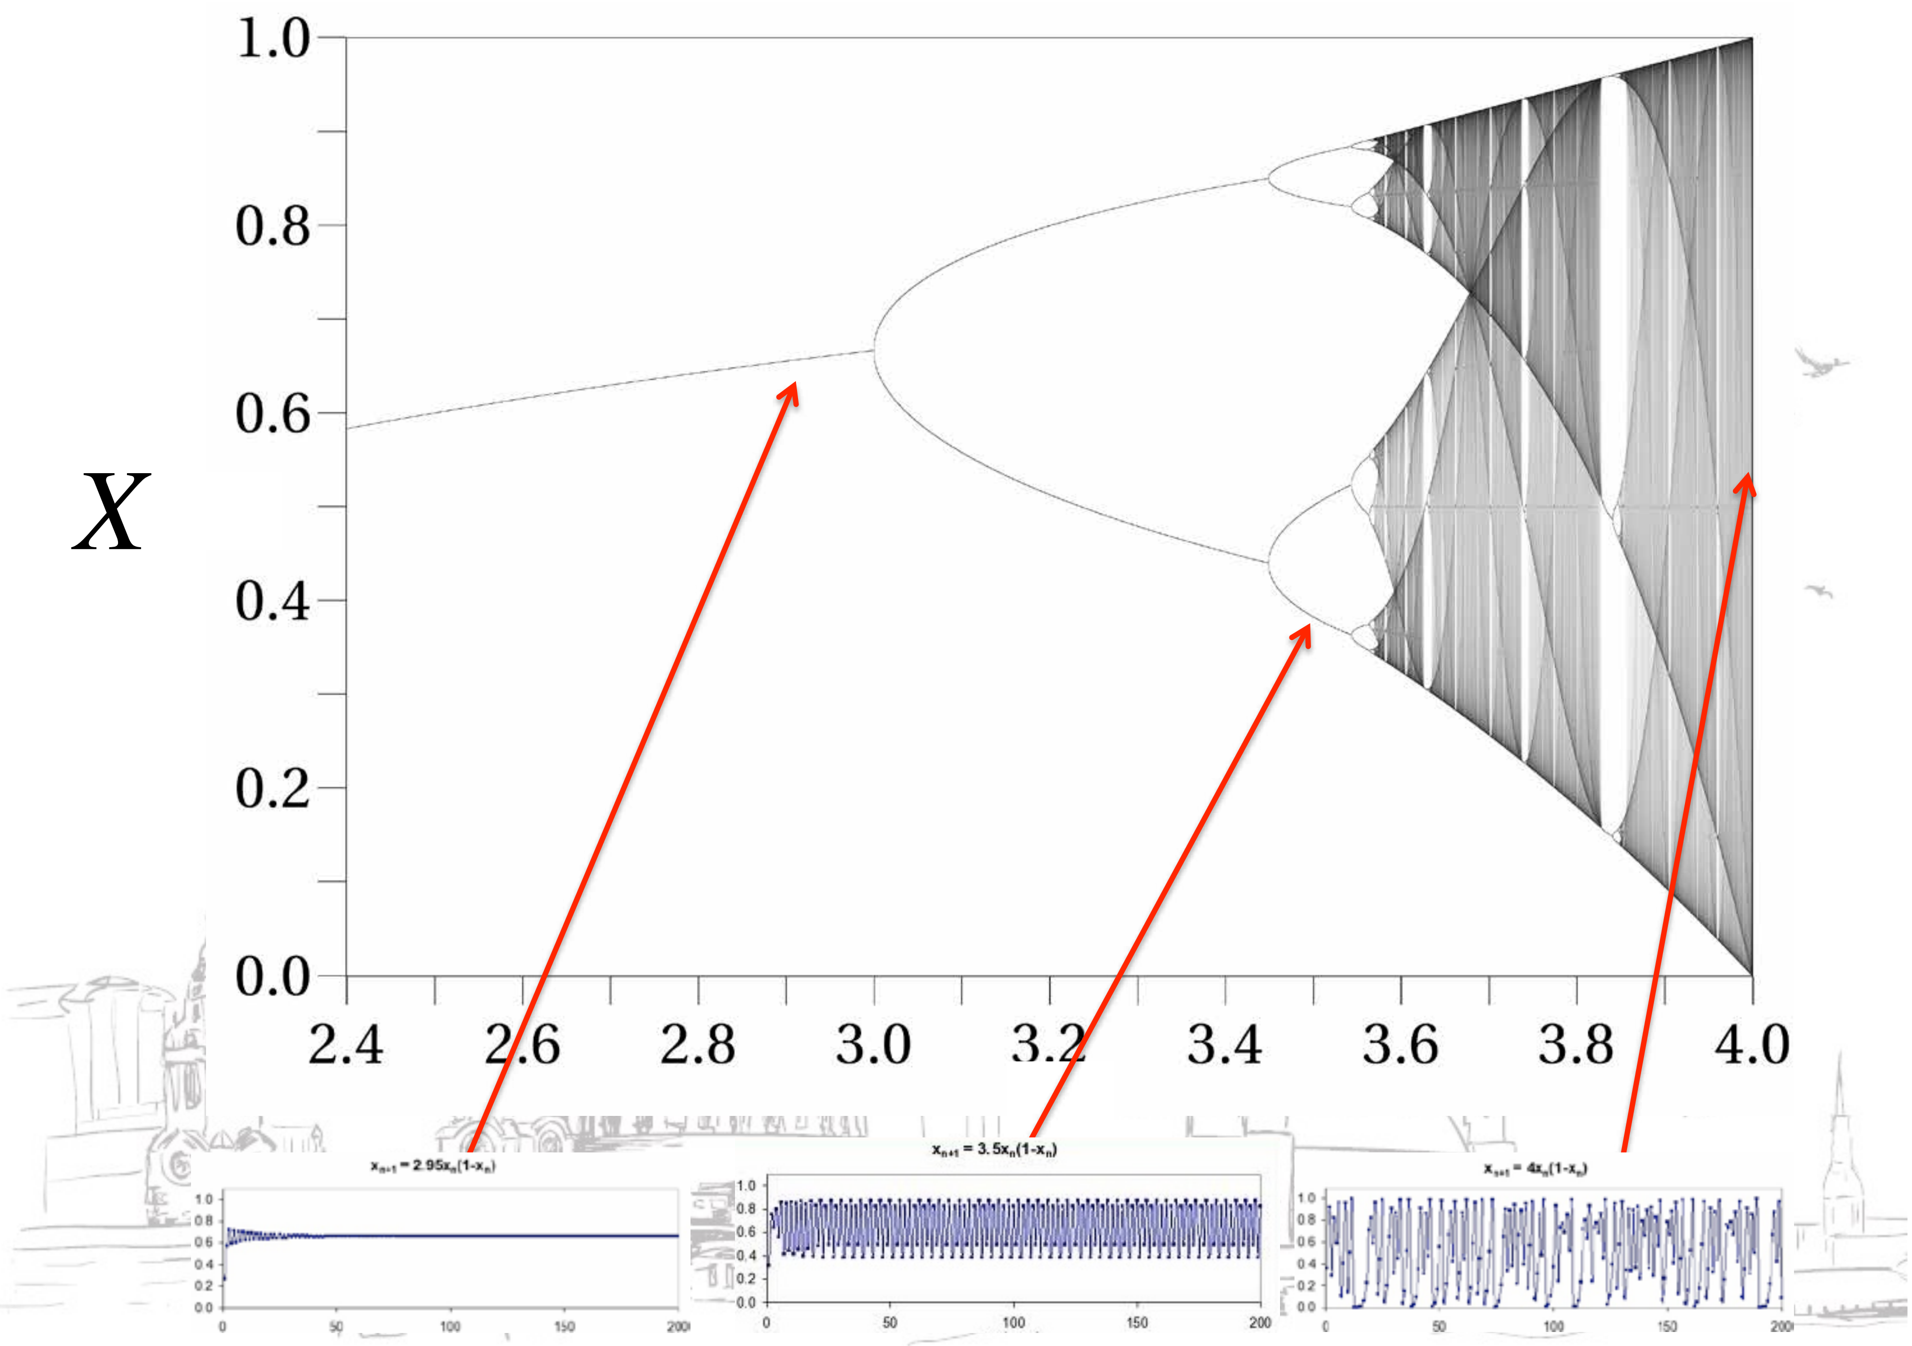

1.0
0.8
0.6
0.4
0.2
0.0
X
2.4
2.6
2.8
3.0
3.2
3.4
3.6
3.8
4.0
$x_{n+1} = 2.95x_n(1-x_n)$
$x_{n+1} = 3.5x_n(1-x_n)$
$x_{n+1} = 4x_n(1-x_n)$

This is a study of **parameter variation**.

It provides information about what happens if we are uncertain about parameter values.

But: it provides **no** information about what happens when we are **uncertain about the model structure**.

What if the true equation is not exactly

$$\rho_{t+1} = \alpha \rho_t (1 - \rho_t) \text{ ?}$$

Overselling an example?

Recall our conclusion: the closeness to goodness link is not an adequate means to deal with <u>structural model error</u>.

Why is this a ***general*** problem and not just a problem of our example?

There is an elaborate mathematical theory of structural stability:

Andronov and Pontrjagin, Peixoto, Palis, Smale, Mañé, Hayashi.

But:

Stability proofs are forthcoming only for two-dimensional flows!

But that is a very special kind of system!

In general the situation is more involved:

Axiom A: the system is uniformly hyperbolic.

Strong transversality condition: stable and unstable manifolds must intersect transversely at every point.

Palis and Smale (1970) conjectured that a system is structurally stable iff it satisfies Axiom A and the strong tranversality condition.

Proof: Mañé (1988) for maps and Hayashi (1997) for flows.


What do Axiom A and the strong transversality condition mean for physical models?

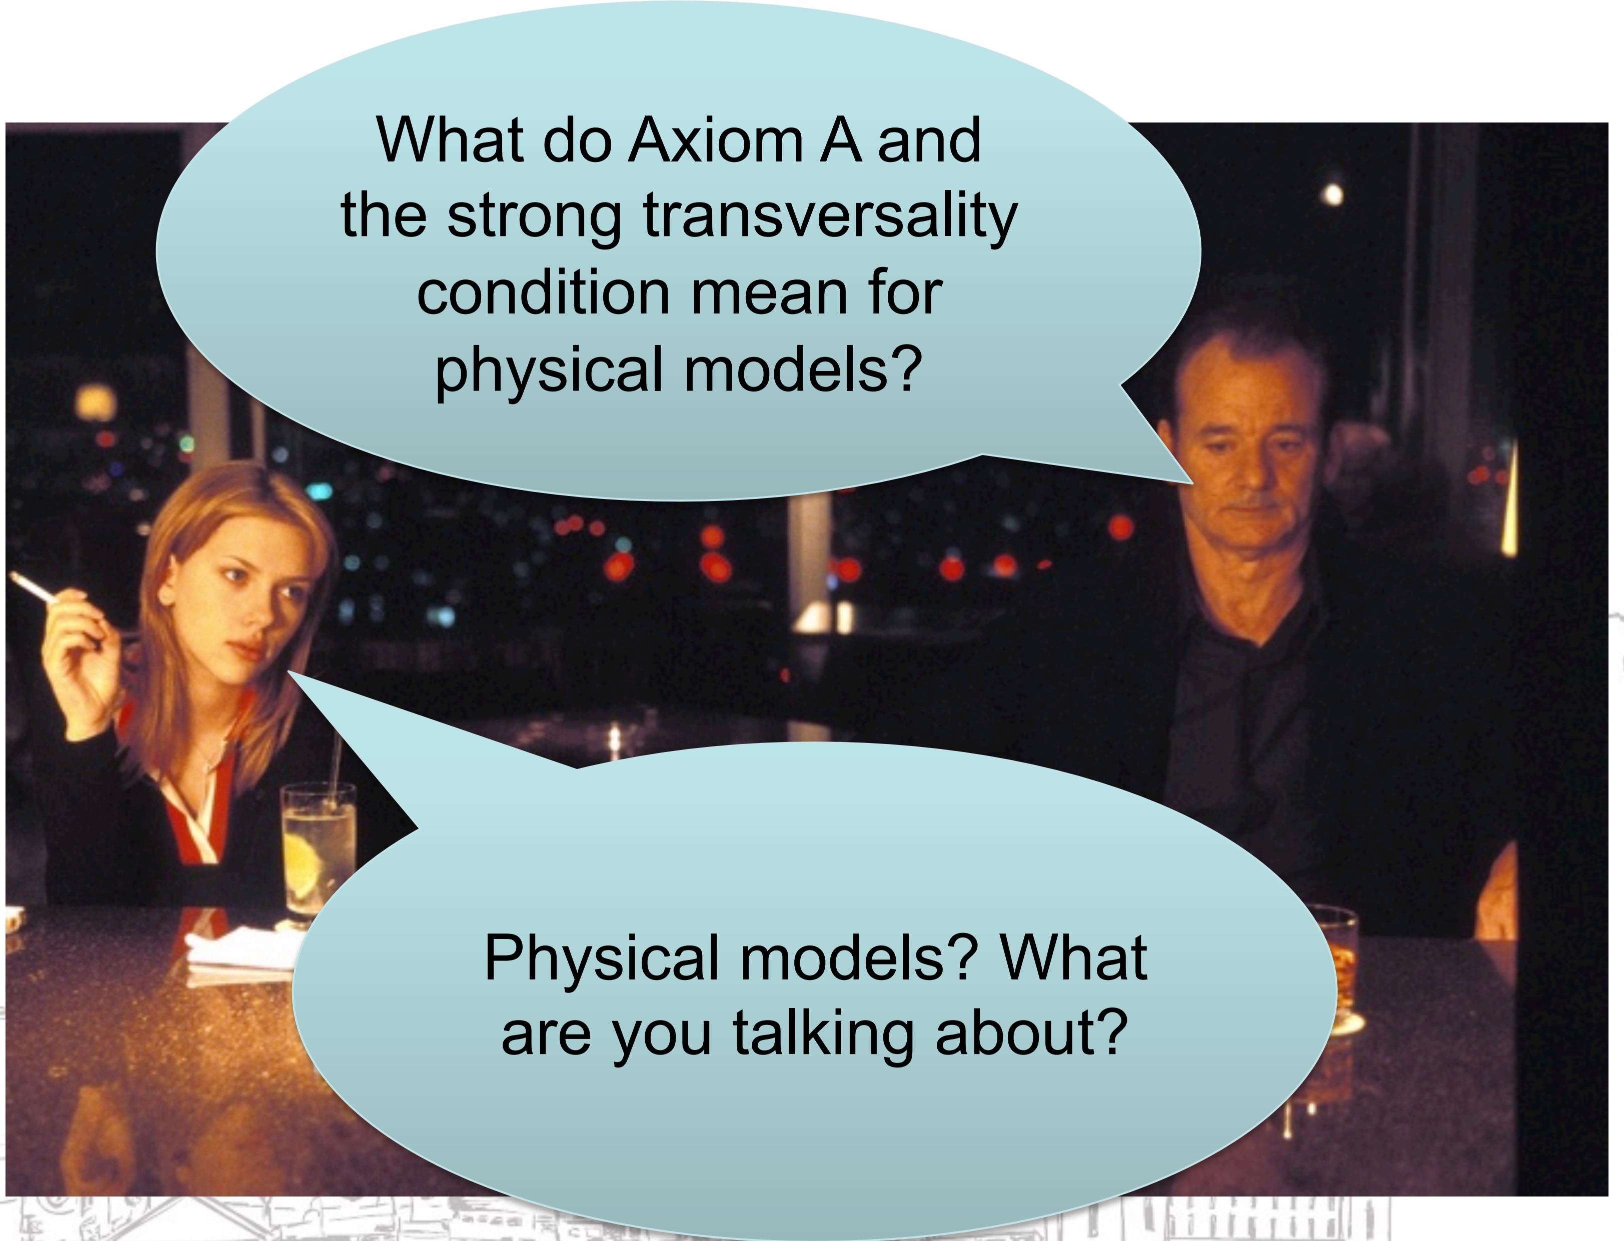
What do Axiom A and the strong transversality condition mean for physical models?
Physical models? What are you talking about?

But:

Smale (1966): structural stability is not generic in the class of diffeomorophisms on a manifold: the set of structurally stable systems is open but not dense.

Smith (2002) and Judd and Smith (2004): if the model's and the system's dynamics are not identical, then 'no state of the model has a trajectory consistent with observations of the system' (2004, 228).

Minimal conclusion: shift of the onus of proof!

Those using non-linear models for predictive purposes owe us an argument that they are structurally stable, not *vice versa*!

Thank you!