Optimal
Uncertainty Quantification

Houman Owhadi

Technology, in common with many other activities, tends toward avoidance of risks by investors. Uncertainty is ruled out if possible. People generally prefer the predictable. Few recognize how destructive this can be, how it imposes severe limits on variability and thus makes whole populations fatally vulnerable to the shocking ways our universe can throw the dice.

Frank Herbert (Heretics of Dune)

Caltech Nov 2013

Optimal Uncertainty Quantification. H. Owhadi, Clint Scovel, T. Sullivan, M. McKerns and M. Ortiz. **SIAM Review** Vol. 55, No. 2 : pp. 271-345, 2013

CALTECH PSAAP Predictive Science Academic Alliance Program

The UQ challenge in the certification context

You want to certify that

$$\mathbb{P}[G(X) \geq a] \leq \epsilon$$

Problem

- You don't know G.

and

- You don't know $\mathbb{P}$

The UQ challenge in the certification context (safety of a new model of airplane)

You want to certify that

$$\mathbb{P}\Big[\text{Crash per hour of flight}\Big] \leq 10^{-9}$$

Problem

- You don't know all possible causes of a crash
- You don't know $\mathbb{P}$

The UQ challenge in the certification context (Performance of a weapon system)

Figure 2-1. Generic Design Elements of a Modern Nuclear Weapon

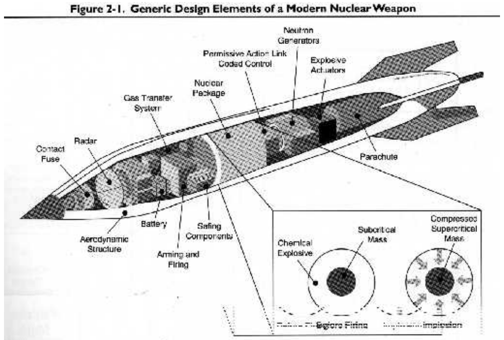

You want to certify that

$$\mathbb{P}\left[\text{failure}\right] \le \text{treshold}$$

Problem

- You cannot test it.
- You don't know all possible causes of a failure
- You don't know $\mathbb{P}$

BUT

- You can simulate 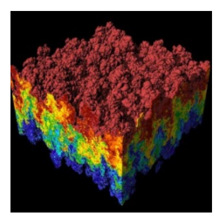
- You have 20 samples from the old system

The UQ challenge in the prediction context (Climate modeling)

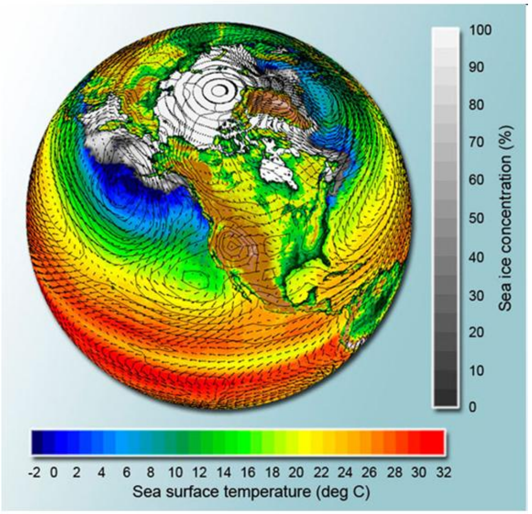

You want to find a 95% interval of confidence on average global temperatures in 50 years

Problem

- Incomplete information on underlying processes
- Limited computation capability
- You don't know $\mathbb{P}$

The UQ challenge in the prediction context (Deepwater Horizon Disaster)

You want to find a 95% interval of confidence on the spill rate

Problem

- You don't know $\mathbb{P}$
- No one really knows how to measure deep water spills of this type.

You want to certify that

$$\mathbb{P}[G(X) \geq a] \leq \epsilon$$

Problem

and

- You don't know G.
- You don't know $\mathbb{P}$

Best thing to do **Compute**

optimal bounds $\mathbb{P}[G(X) \geq a]$
given available information.

Best and Worst case scenarios

Seismic Safety Assessment of a Truss Structure

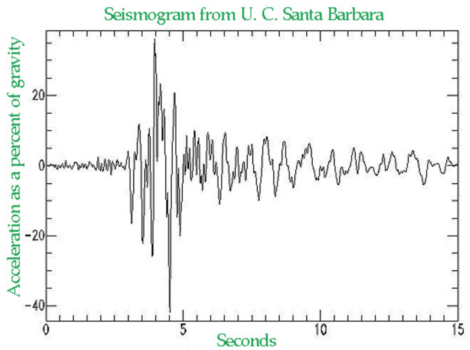

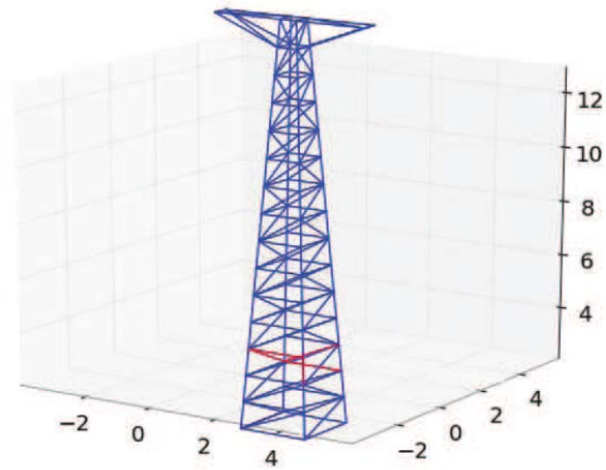

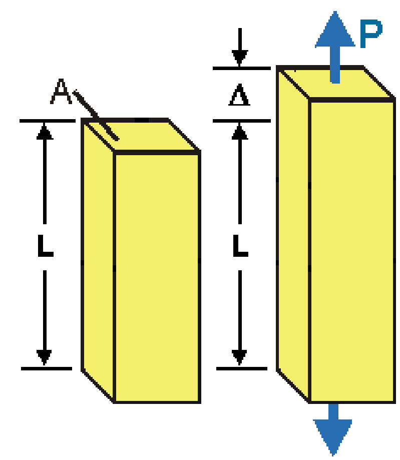

$$a(t) \longrightarrow F(a)$$

Ground Acceleration

F

min(Yield Strain - Axial Strain)

We want to certify that

$$\mathbb{P}\left[F(a) \leq 0\right] \leq \epsilon$$

Historical Data Method

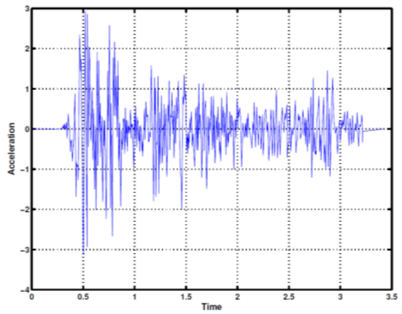

1940 Elcentro

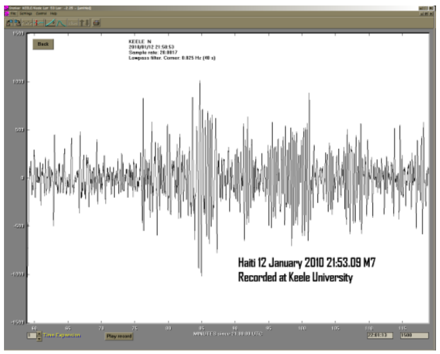

2010 Haiti

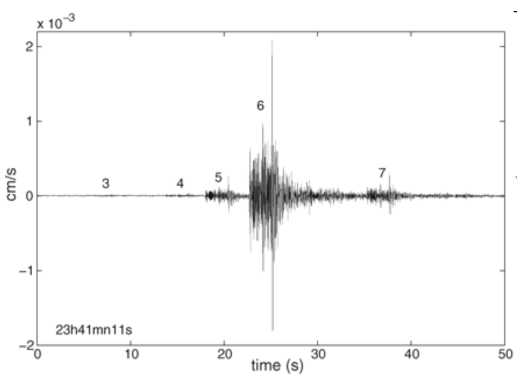

1999 Izmit

Power Spectrum

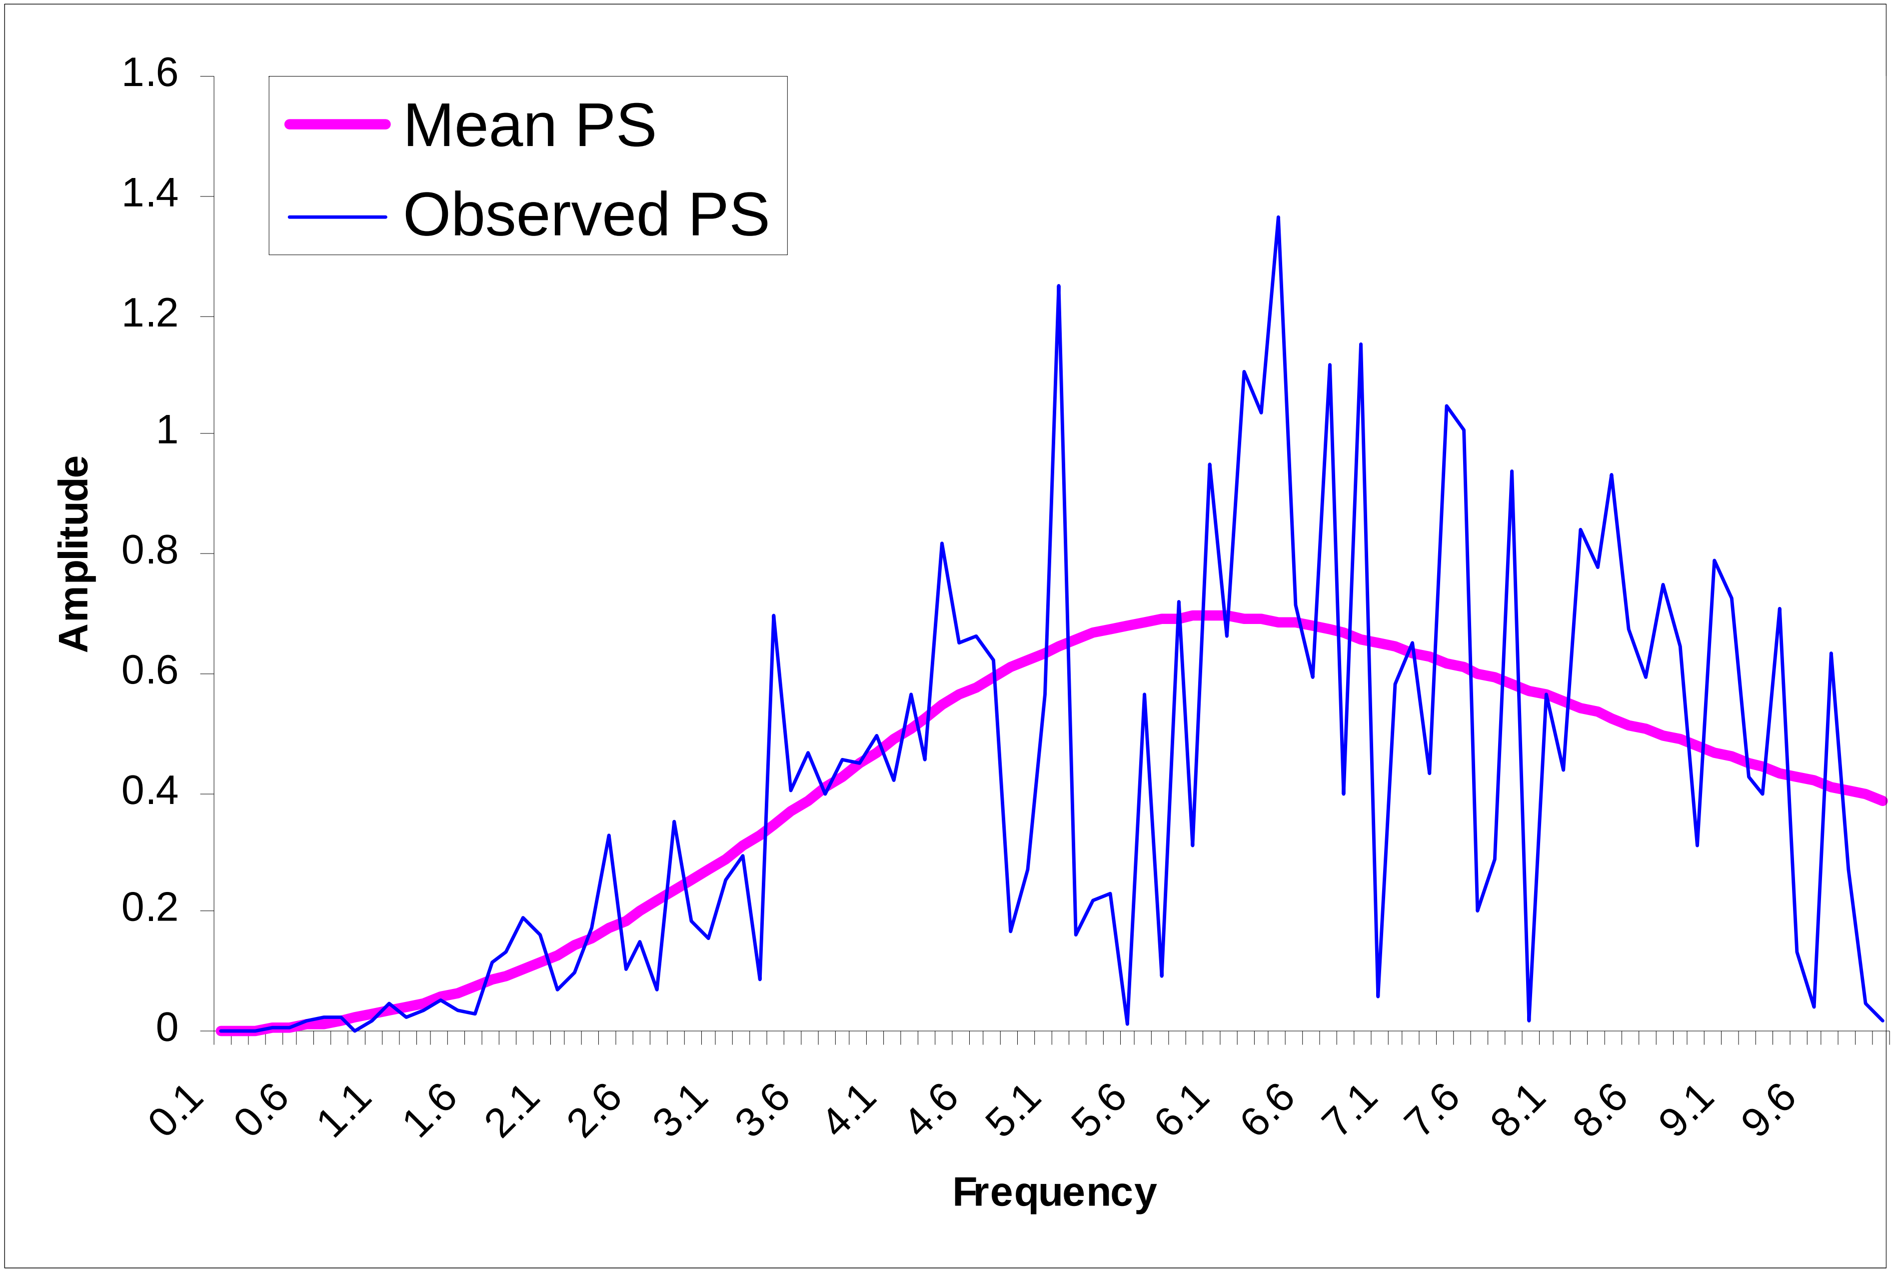

Power Spectrum

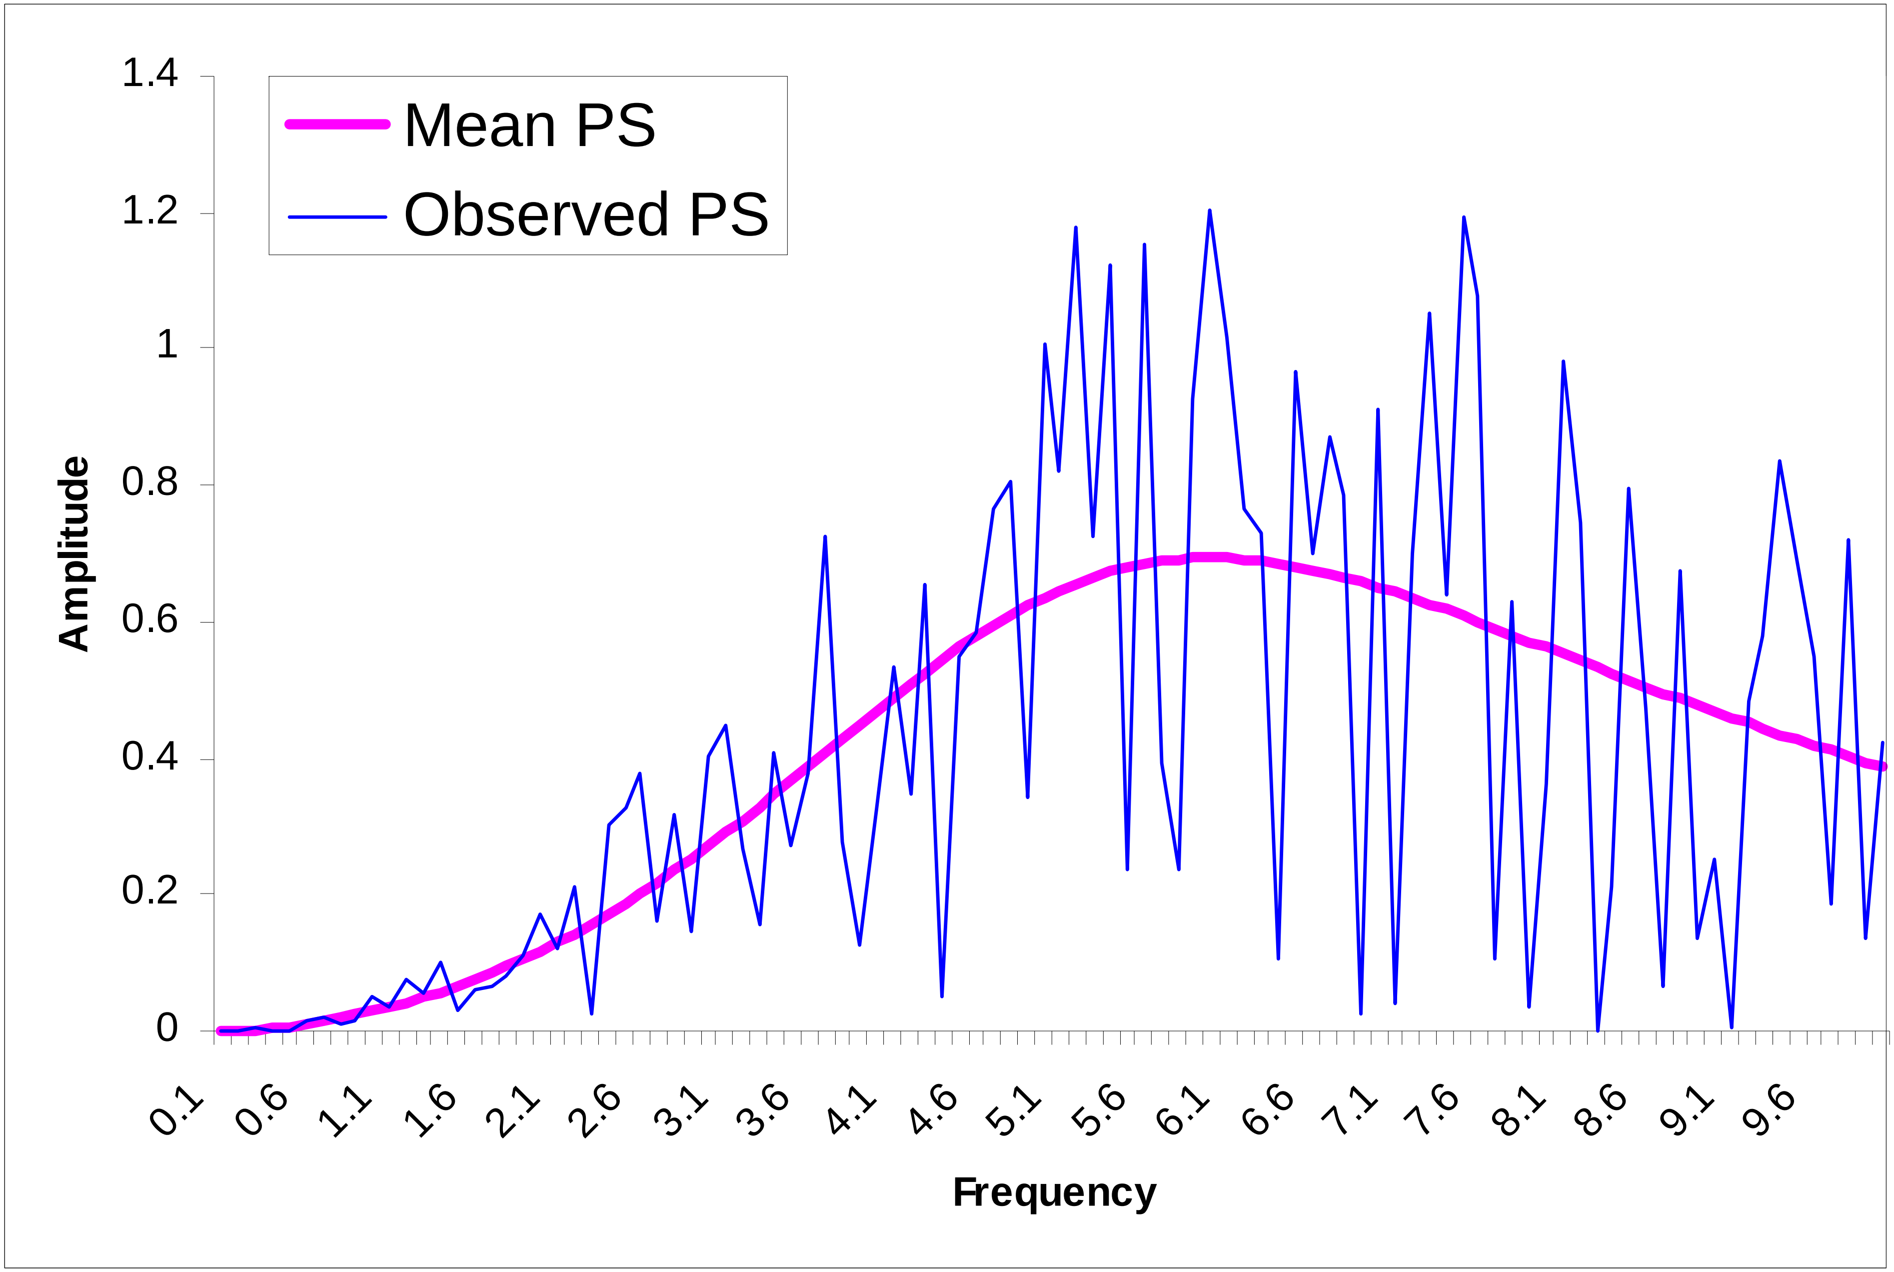

Power Spectrum

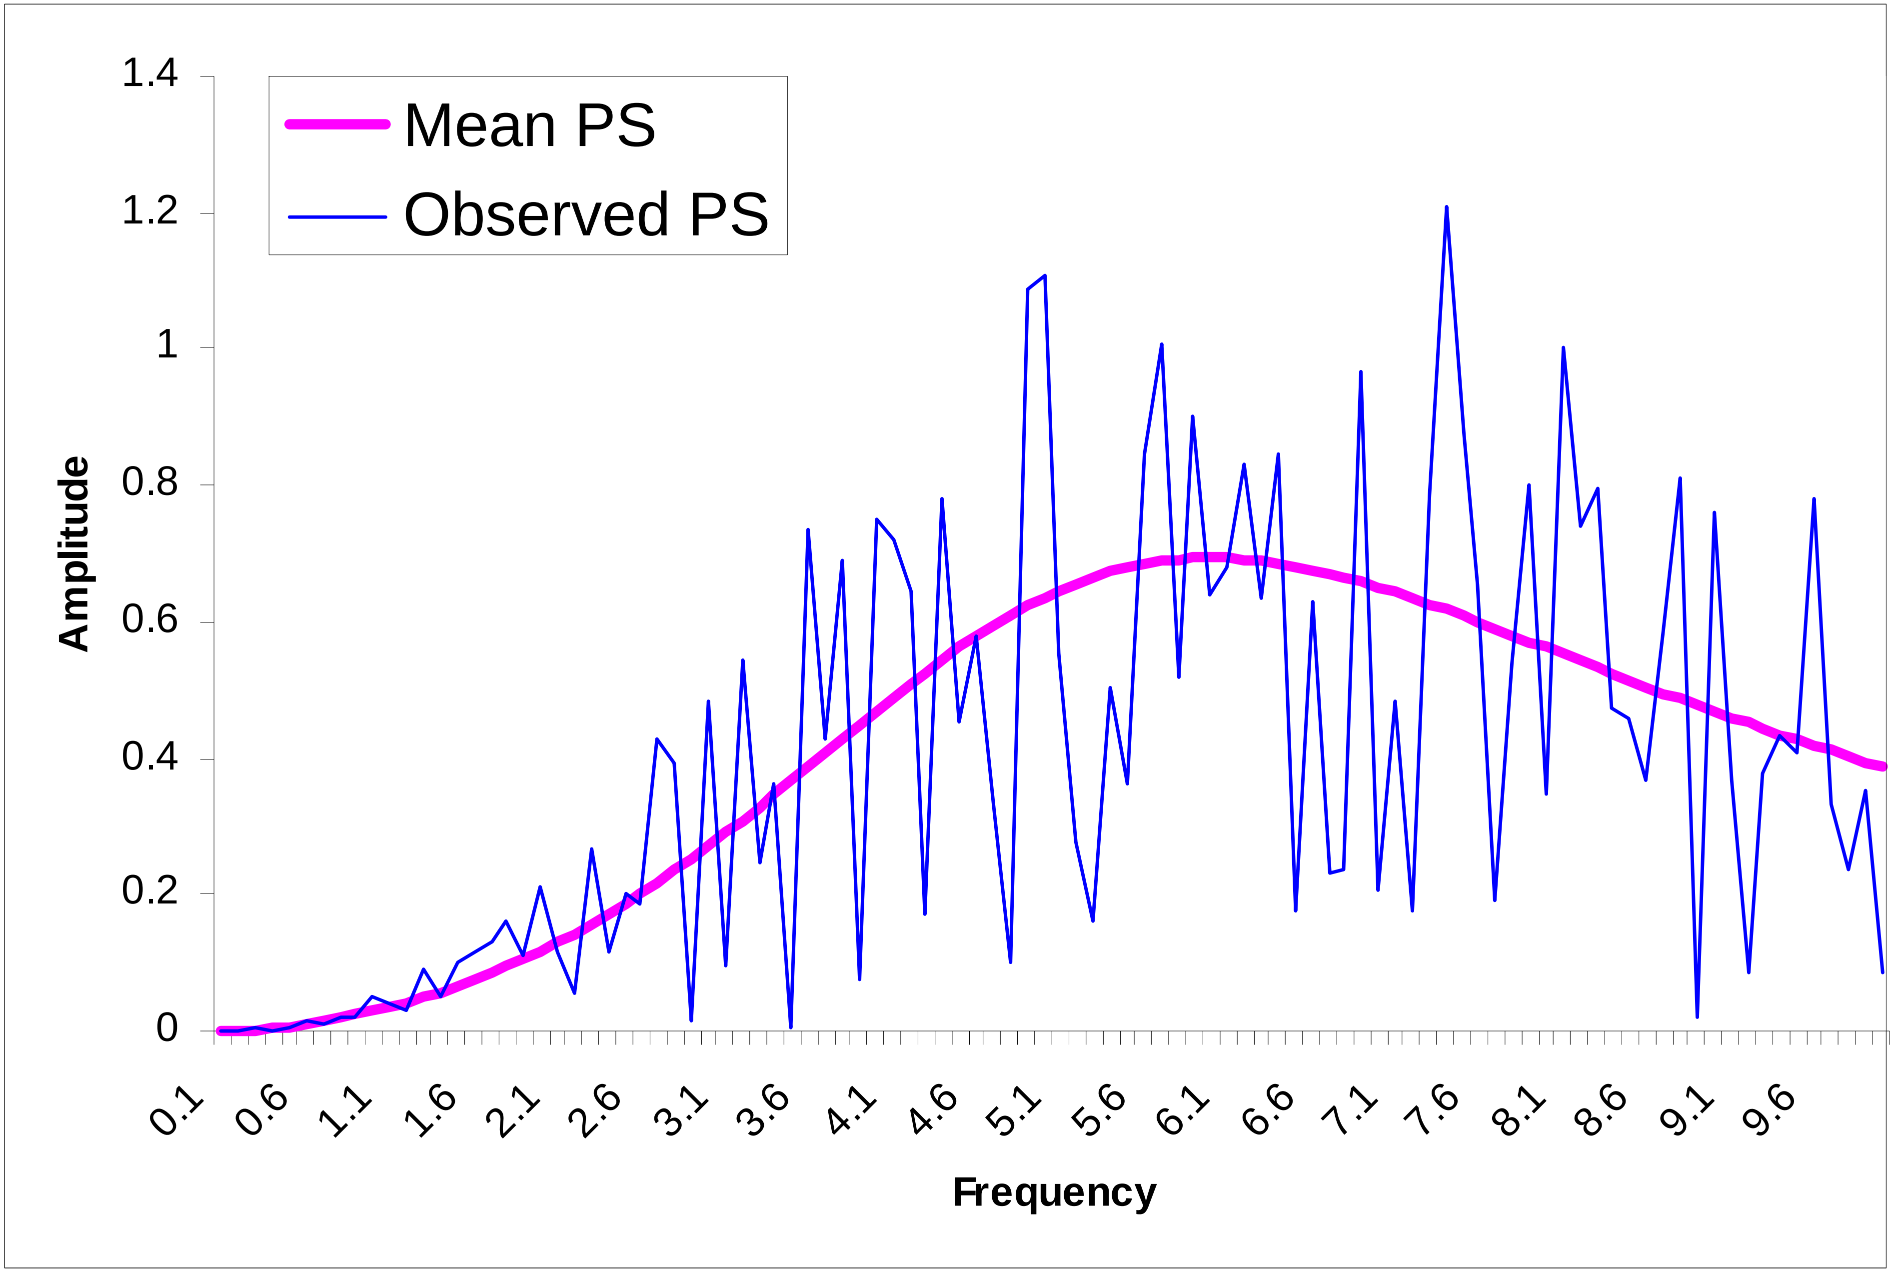

Filtered White Noise Model

White noise

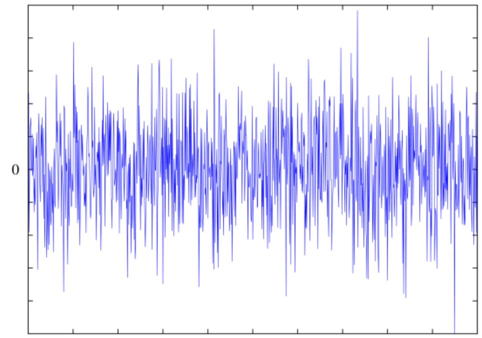

Filter

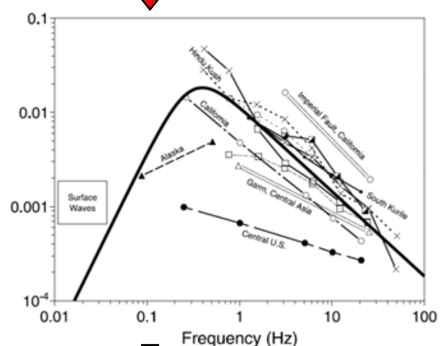

Ground acceleration

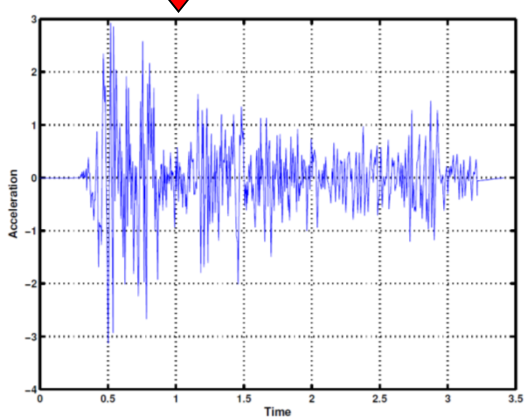

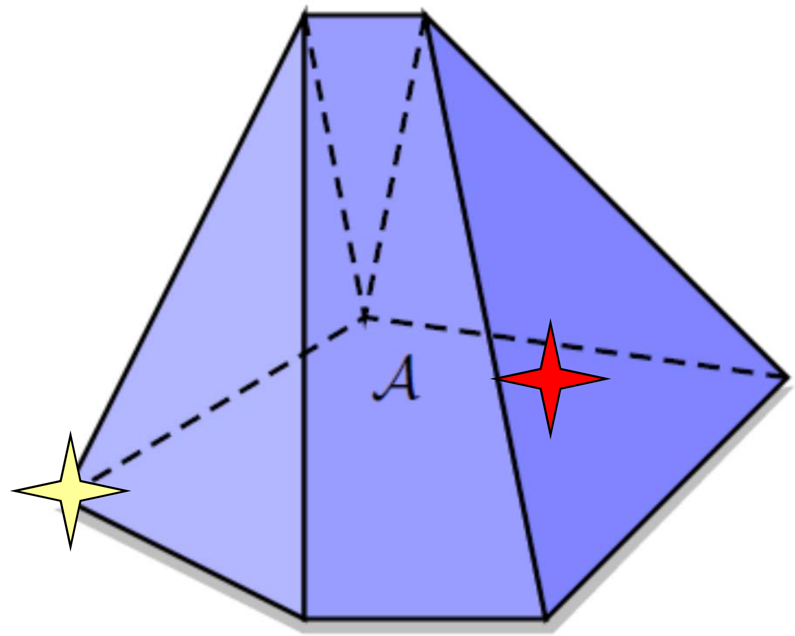

Vulnerability Curves (vs earthquake magnitude)

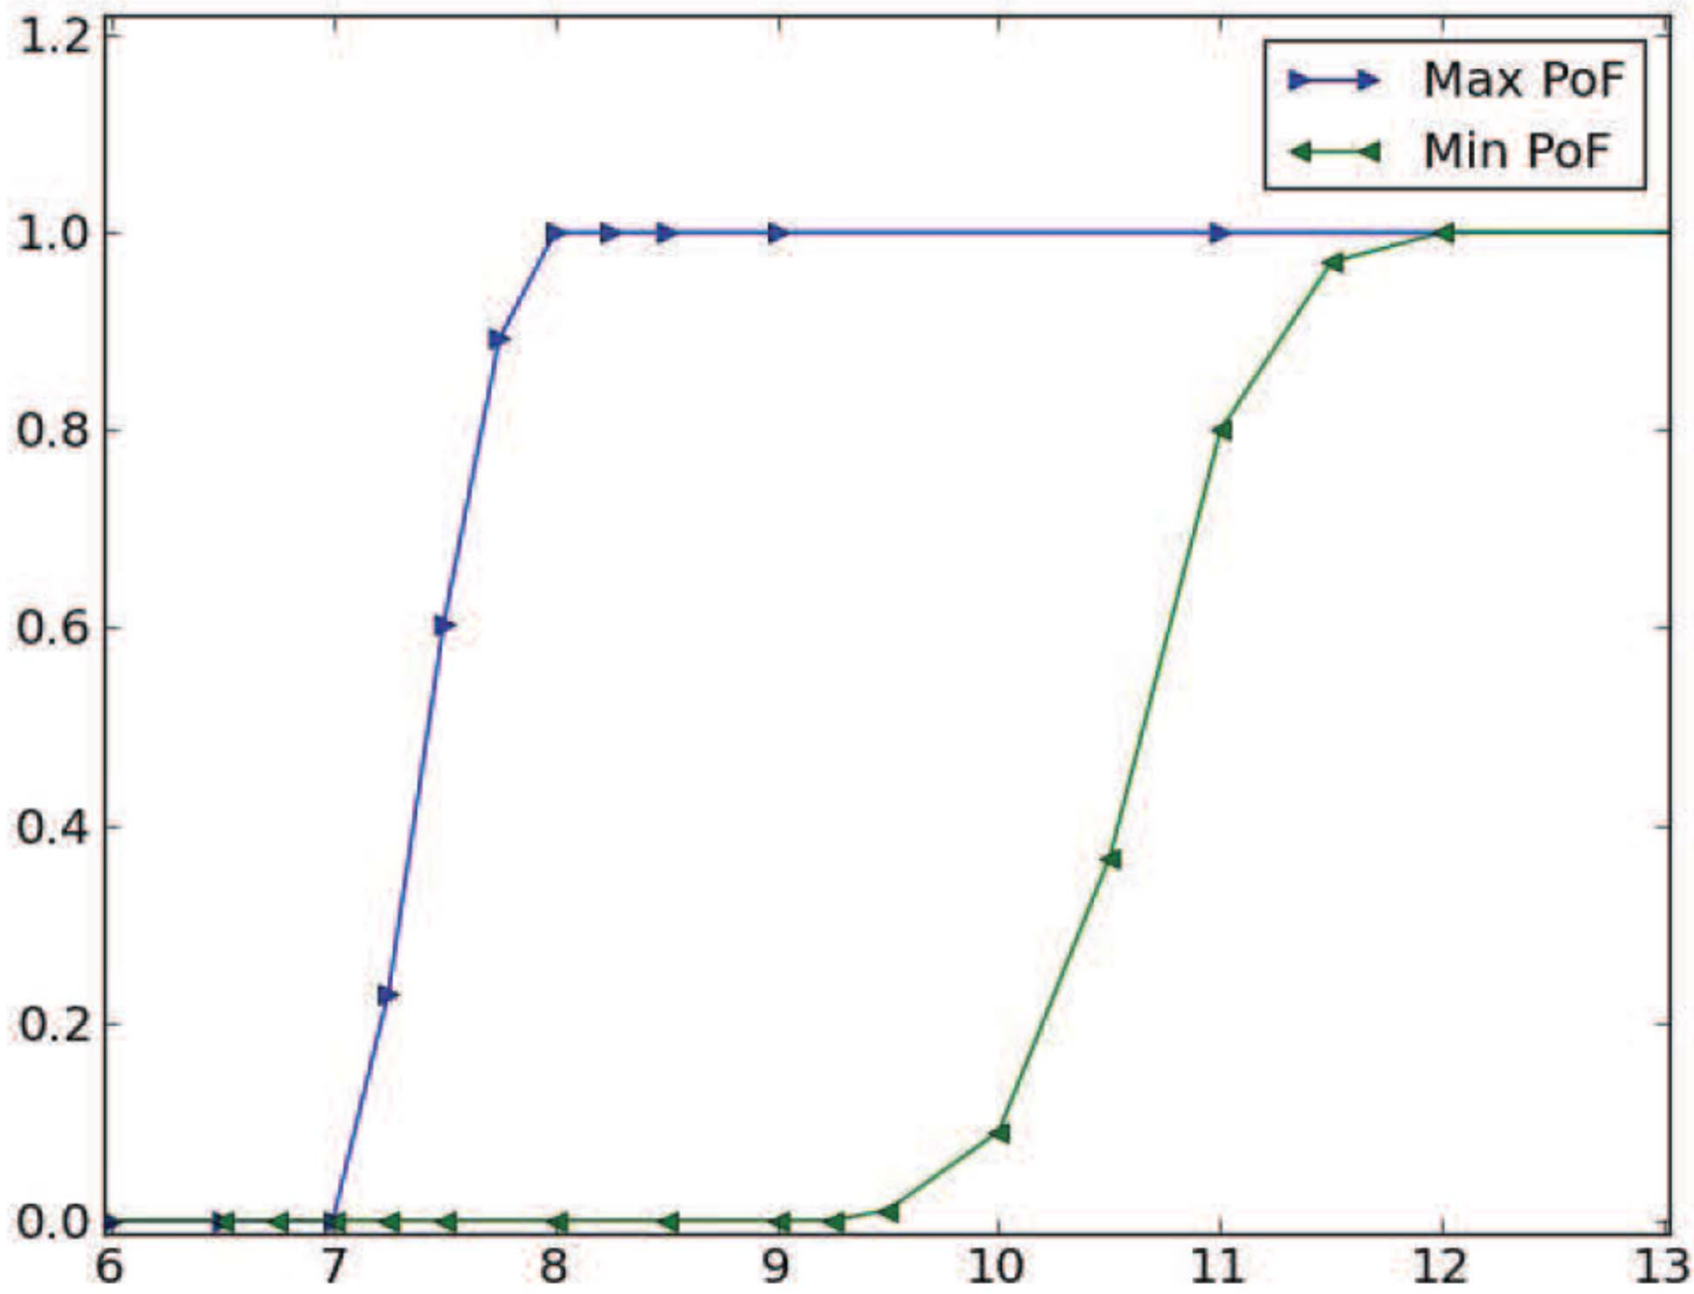

A simple example

What is the least upper bound on $\mathbb{P}[X \geq a]$

If all you know is $\mathbb{E}[X] \leq m$

and $\mathbb{P}[0 \leq X \leq 1] = 1$?

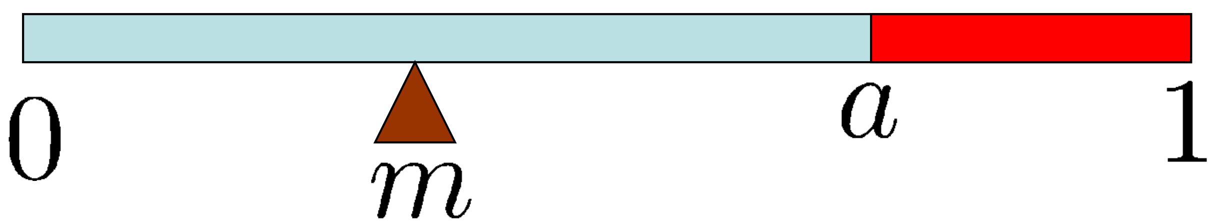

Answer

$$\sup_{\mu \in \mathcal{A}} \mu[X \geq a]$$

$$\mathcal{A} = \{\mu \in \mathcal{M}([0,1]) \,|\, \mathbb{E}_{\mu}[X] \leq m\}$$

You are given one pound of play-doh.
How much mass can you put above <u>a</u> while keeping the seesaw balanced around <u>m</u>?

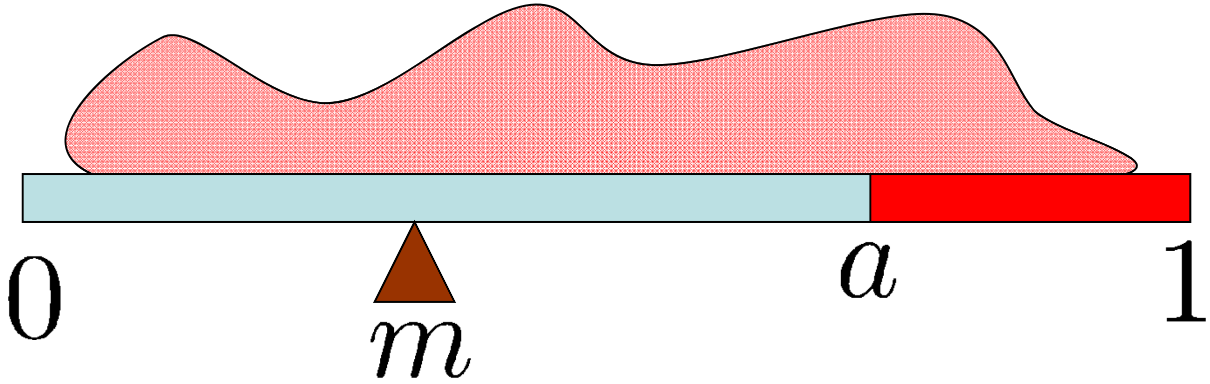

You have to use the whole pound.
The play-doh can be spread arbitrarily over the seesaw.

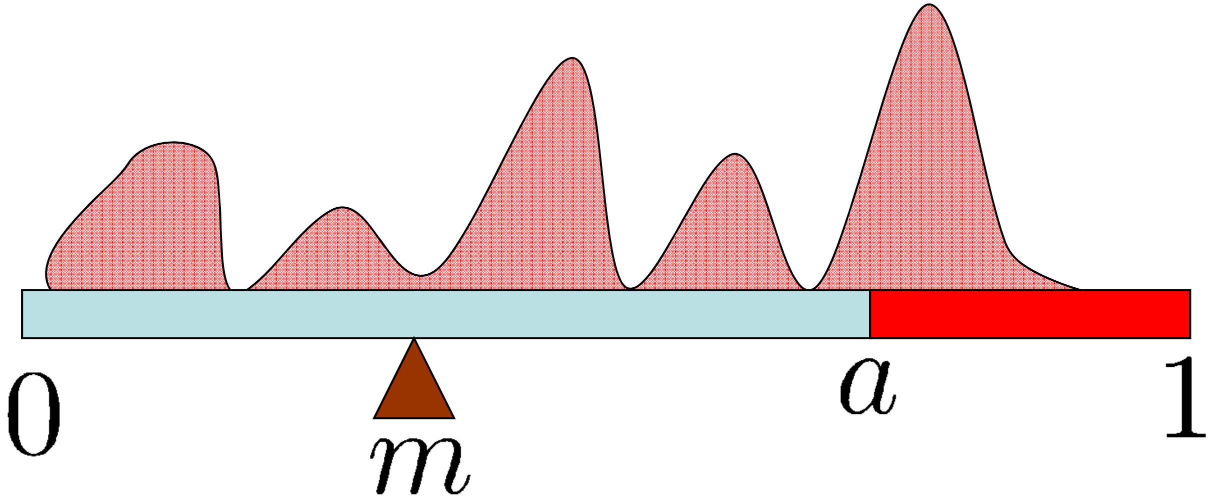

You are given one pound of play-doh.
How much mass can you put above a while keeping the seesaw balanced around m?

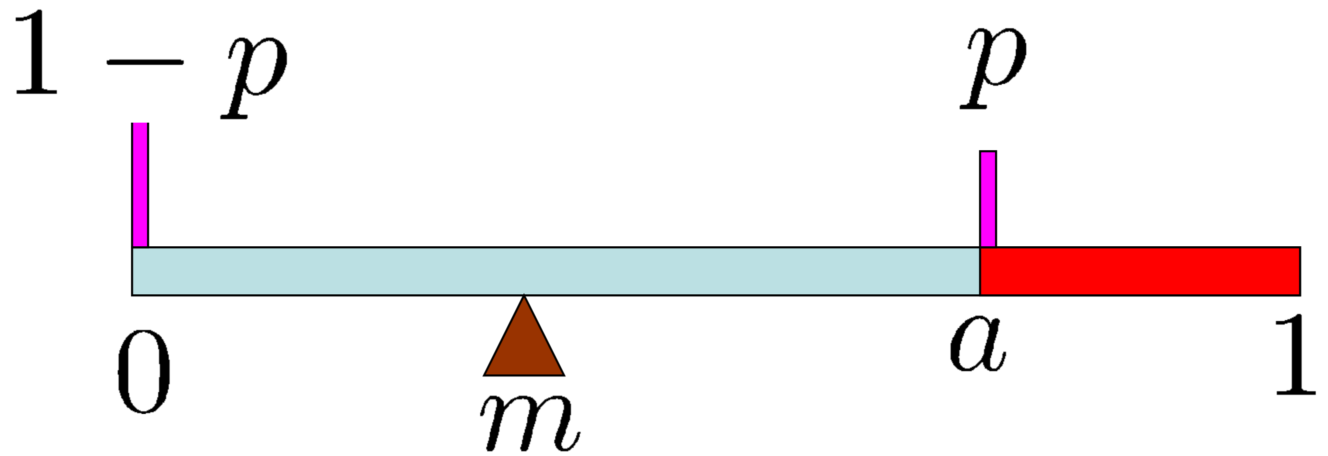

Answer

$$\begin{cases} \max p \\ \text{subject to } a\,p \leq m \end{cases}$$

Markov's inequality

$$\sup_{\mu \in \mathcal{A}} \mu\big[X \geq a\big] = \frac{m}{a}$$

$$\mathcal{A} = \{\mu \in \mathcal{M}([0,1]) \,|\, \mathbb{E}_{\mu}[X] \leq m\}$$

Reduction of optimization variables

$$\{f:\mathcal{X}\to\mathbb{R},\ \mu\in\mathcal{P}(\mathcal{X})\}$$

$$\left\{f:\mathcal{X}\to\mathbb{R},\ \mu\in\mathcal{P}(\mathcal{X})\ \middle|\ \mu=\sum_{i=1}^{k}\alpha_k\delta_{x_k}\right\}$$

$$\{f:\{1,2,\ldots,n\}\to\mathbb{R},\ \mu\in\mathcal{P}(\{1,2,\ldots,n\})\}$$

$$\{\{1,2,\ldots,q\},\ \mu\in\mathcal{P}(\{1,2,\ldots,n\})\}$$

Caltech Small Particle Hypervelocity Impact Range

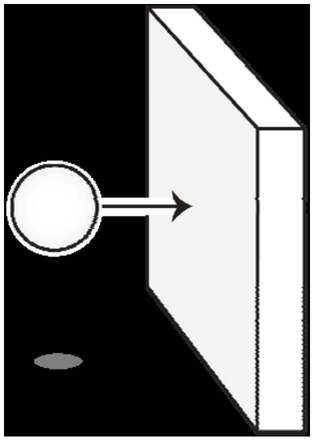

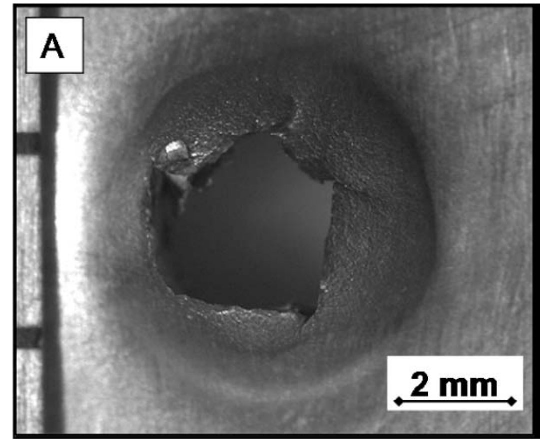

(h, α, v) → $G(h, \alpha, v)$

G

Plate thickness

Plate Obliquity

Projectile velocity

Perforation area

We want to certify that

$$\mathbb{P}[G = 0] \leq \epsilon$$

Marc Adams, Leslie Lamberson, Jonathan Mihaly, Laurence Bodelot, Justin Brown, Addis Kidane, Anna Pandolfi, Guruswami Ravichandran, and Ares Rosakis

Caltech Hypervelocity Impact Surrogate Model

Plate thickness $h \in \mathcal{X}_1 := [1.524, 2.667]\,\mathrm{mm},$

Plate Obliquity $\alpha \in \mathcal{X}_2 := [0, \frac{\pi}{6}],$

Projectile velocity $v \in \mathcal{X}_3 := [2.1, 2.8]\,\mathrm{km}\cdot\mathrm{s}^{-1}.$

Thickness, obliquity, velocity: independent random variables

Mean perforation area: in between 5.5 and 7.5 mm^2

Deterministic surrogate model for the perforation area (in mm^2)

$$G(h, \alpha, v) = K \left(\frac{h}{D_\mathrm{p}} \right)^p (\cos\alpha)^u \left(\tanh\left(\frac{v}{v_\mathrm{bl}} - 1 \right) \right)_+^m,$$

$H_0 = 0.5794\,\mathrm{km}\cdot\mathrm{s}^{-1},\quad s = 1.4004,\quad n = 0.4482,\quad K = 10.3936\,\mathrm{mm}^2,$

$p = 0.4757,\quad u = 1.0275,\quad m = 0.4682.$

$$v_\mathrm{bl} := H_0 \left(\frac{h}{(\cos\alpha)^n} \right)^s$$

The optimization variables can be reduced to the tensorization of 2 Dirac masses on thickness, obliquity and velocity

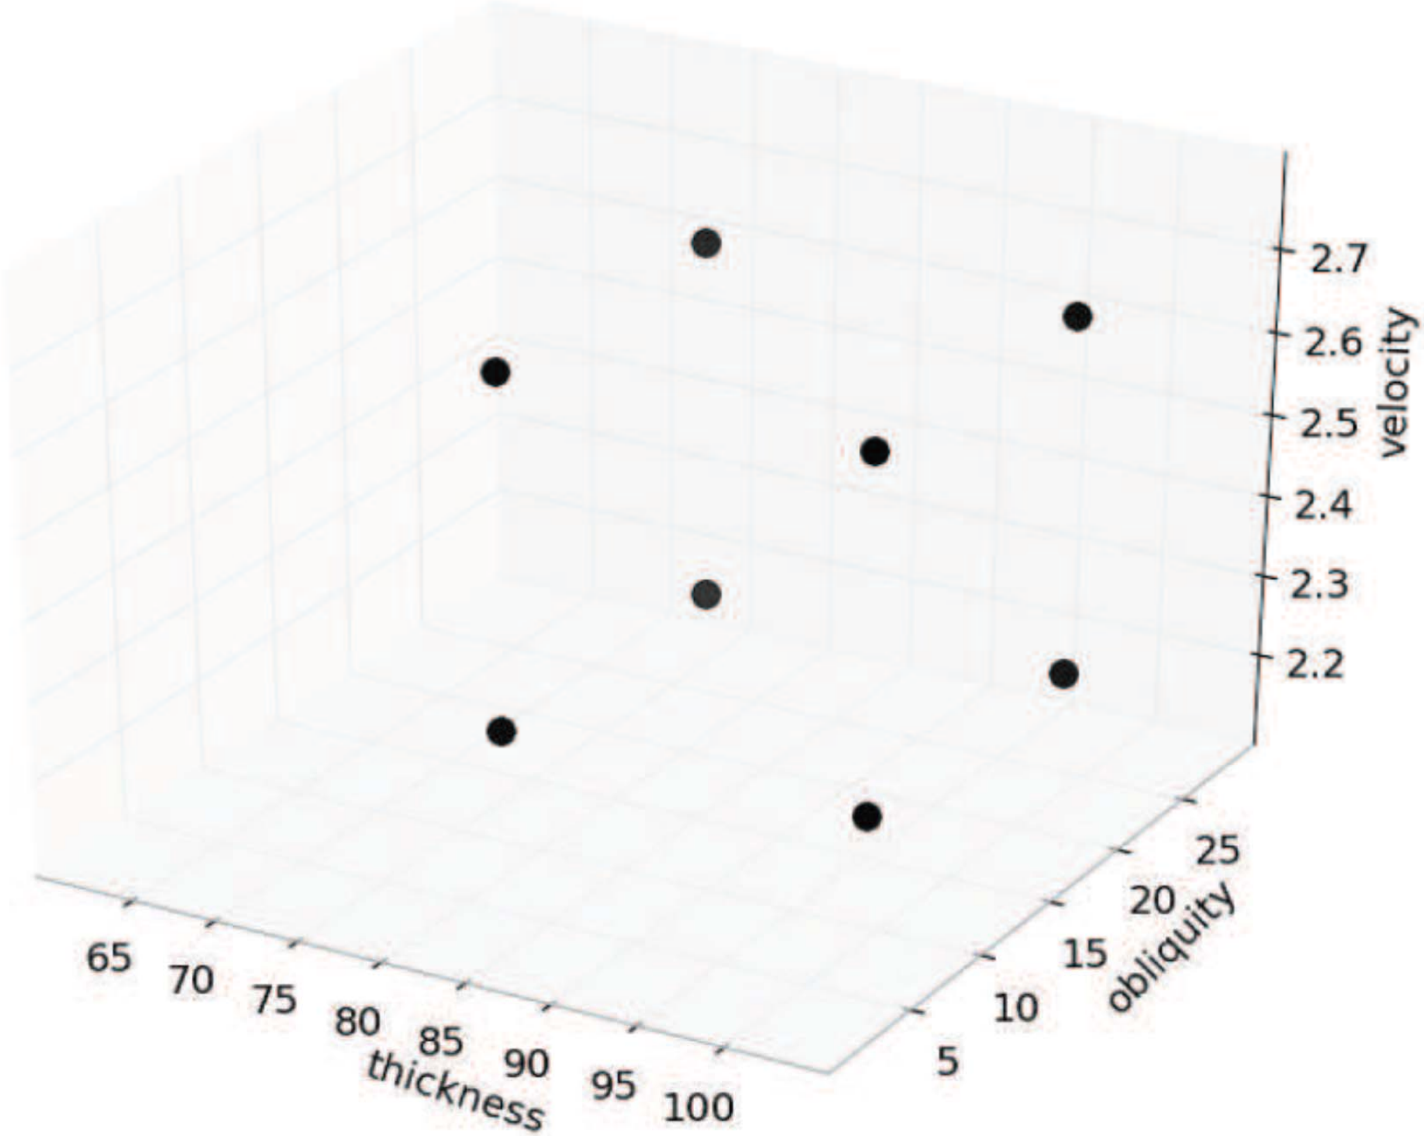

Support Points at iteration 0

Numerical optimization

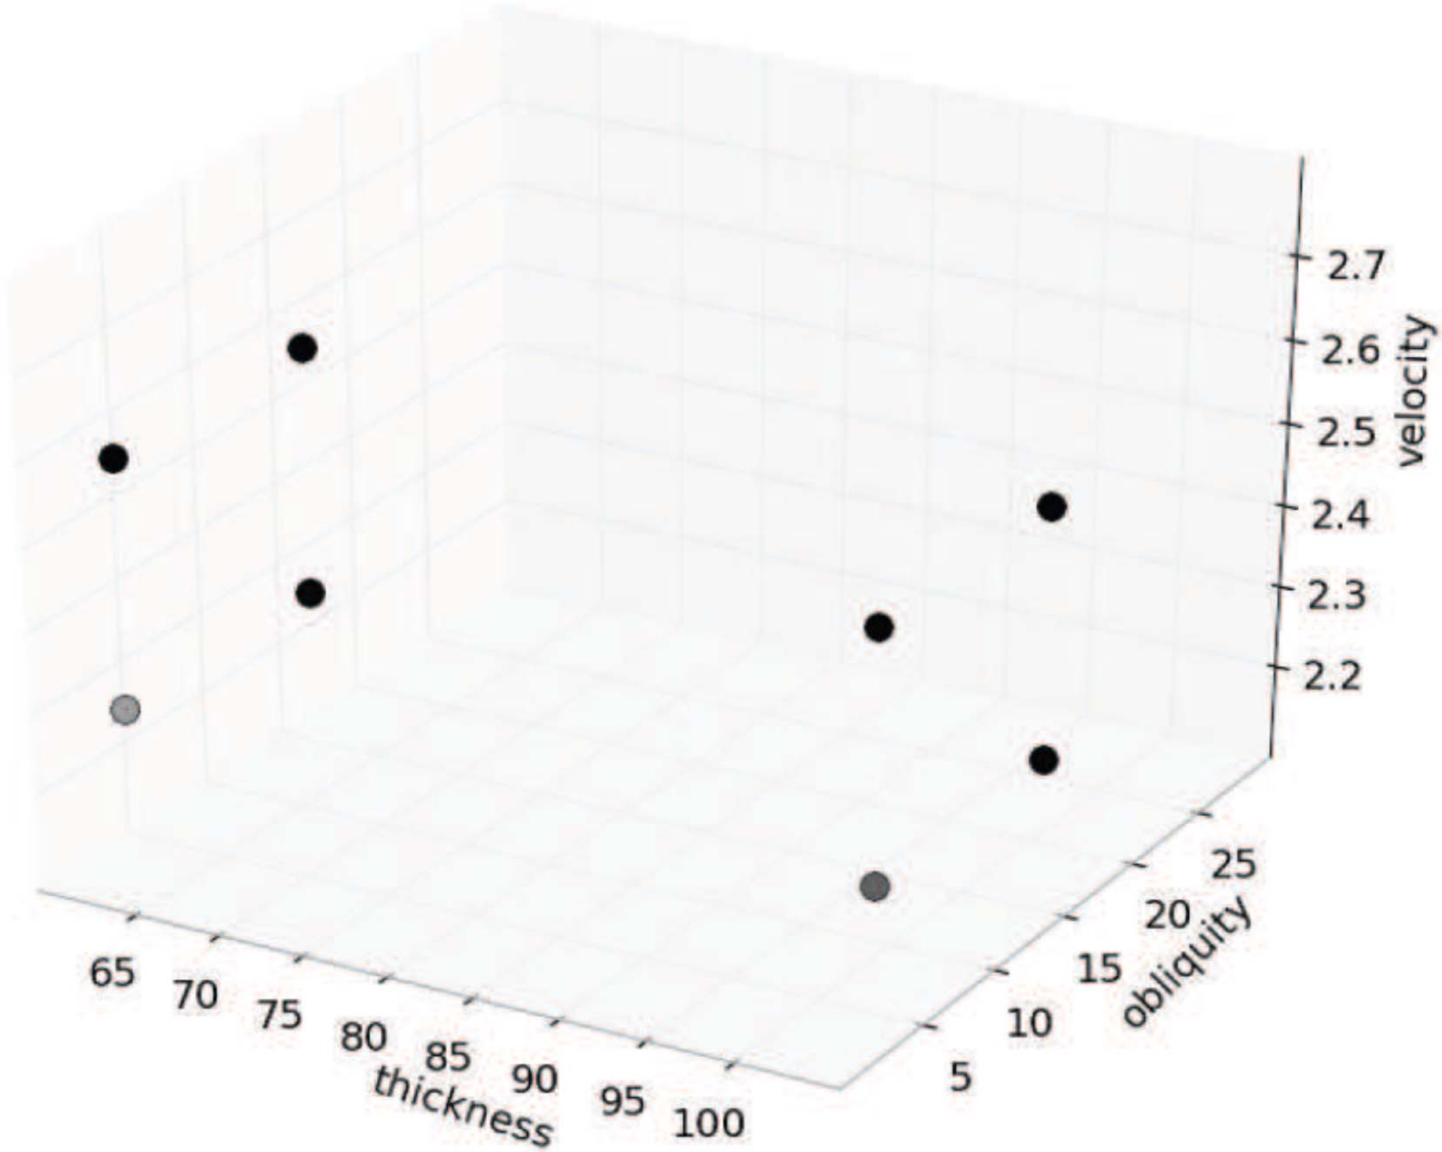

Support Points at iteration 150

Numerical optimization

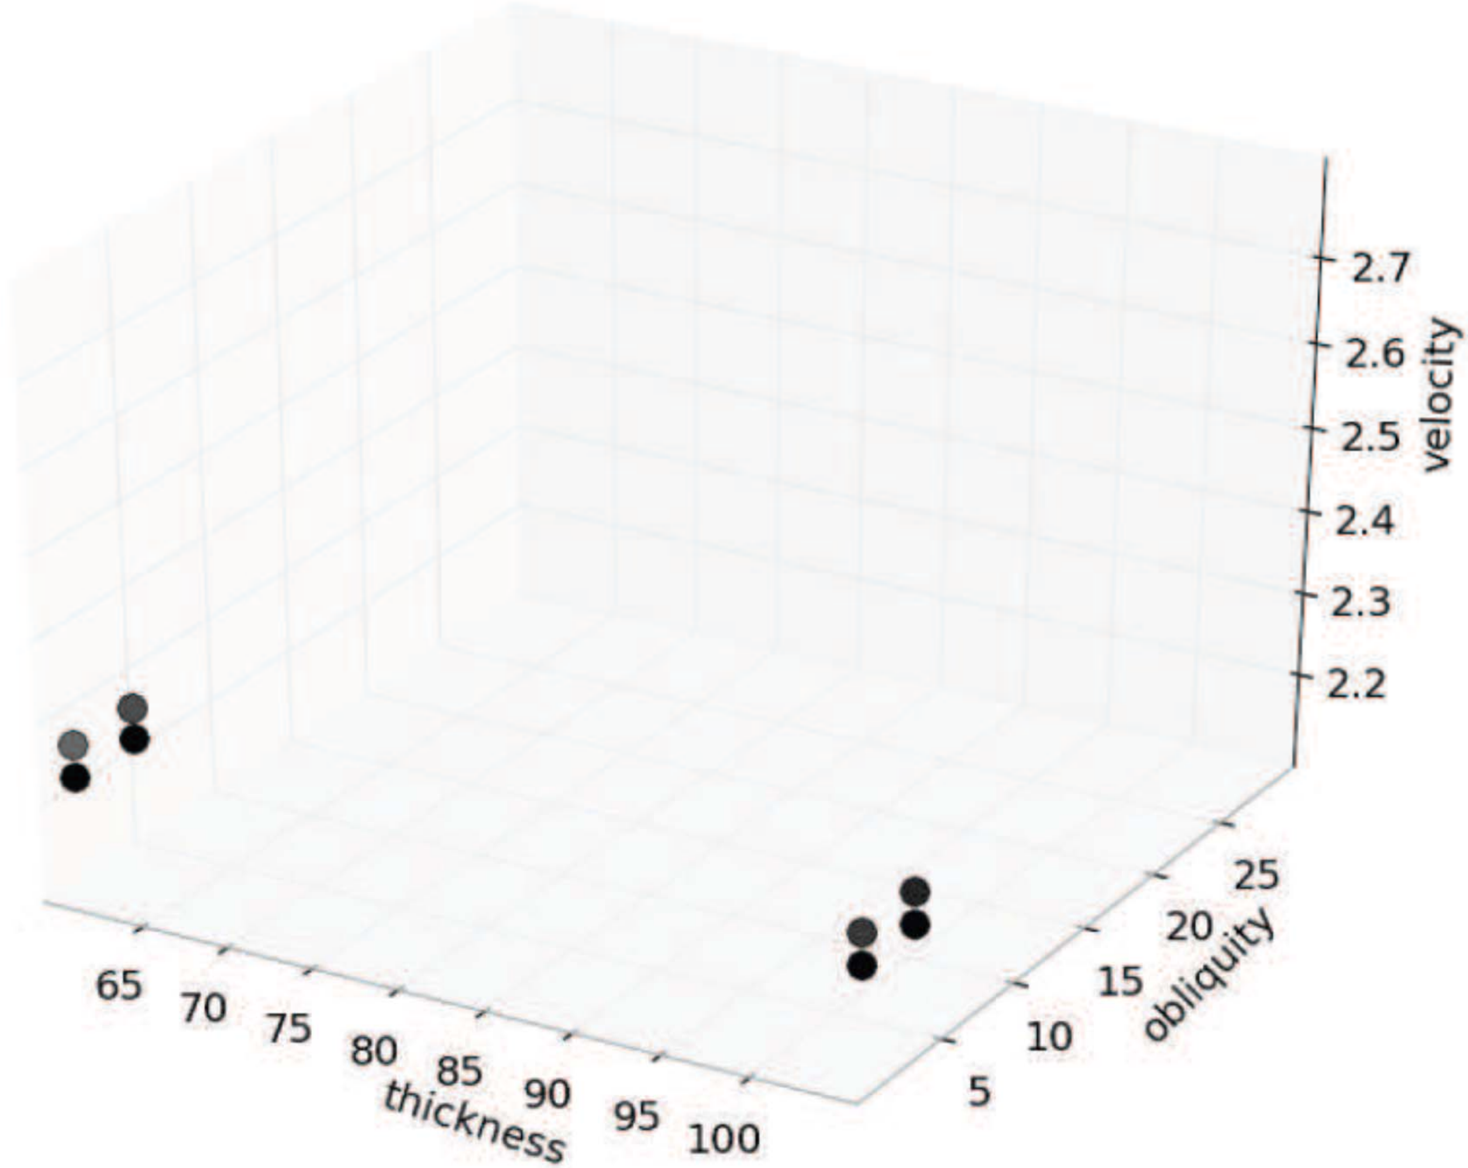

Support Points at iteration 200

Velocity and obliquity marginals each collapse to a single Dirac mass. The plate thickness marginal collapses to have support on the extremes of its range.

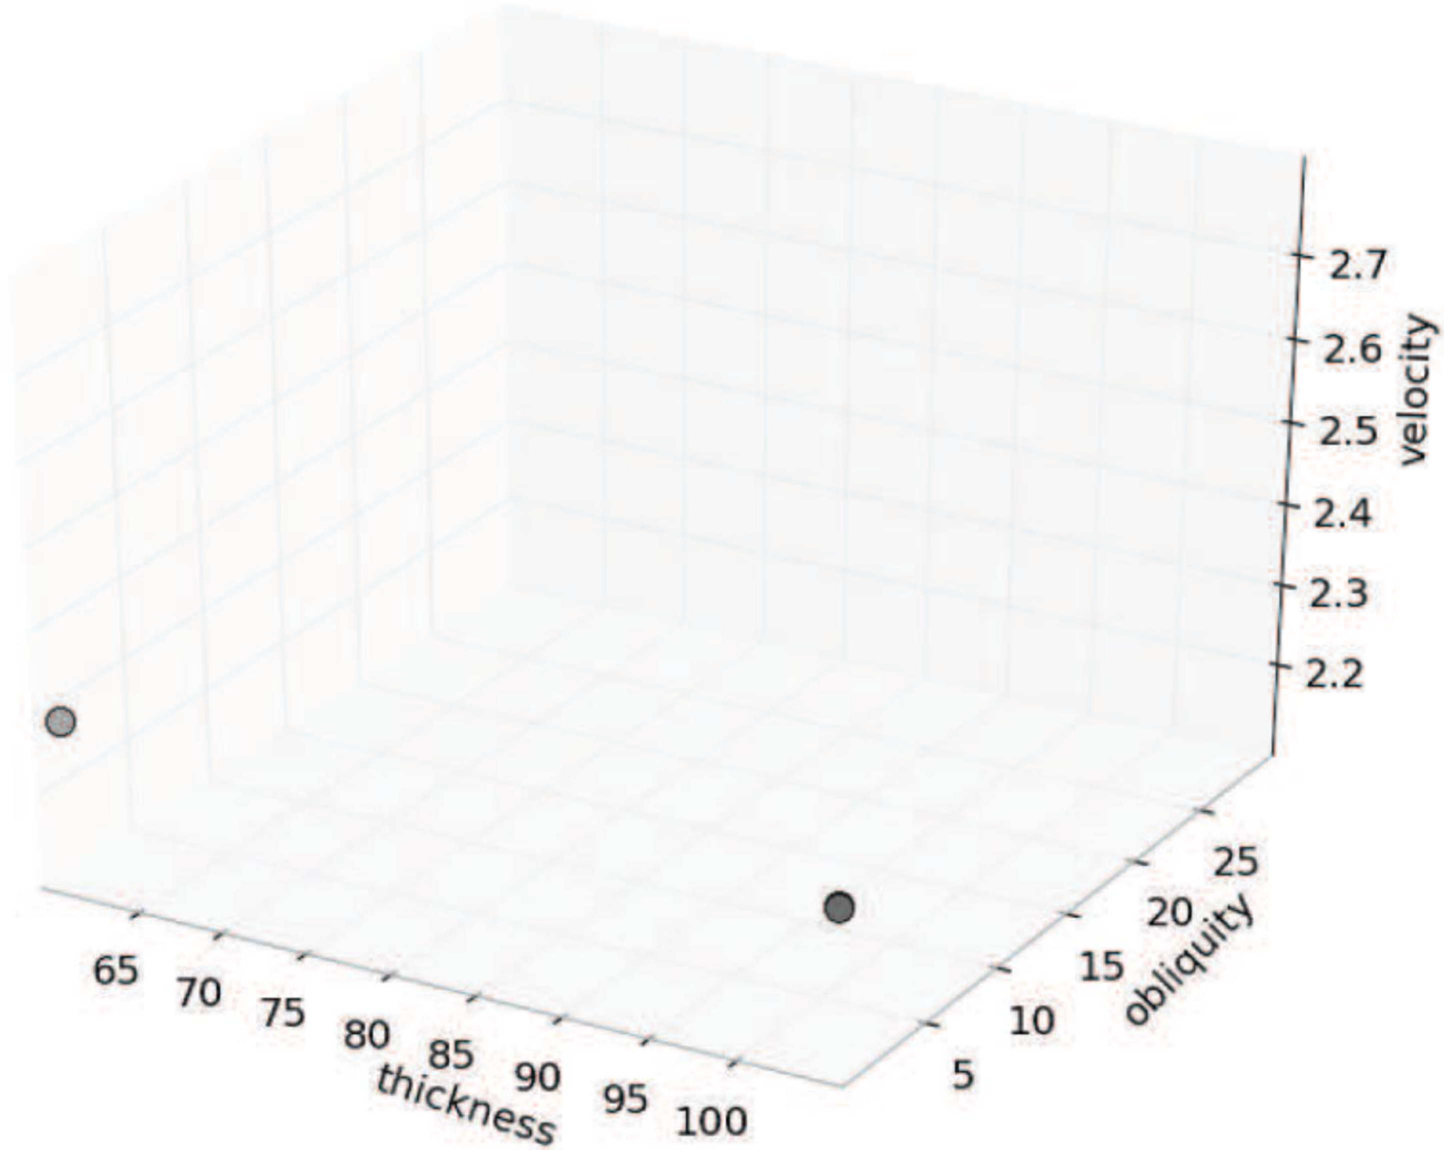

Iteration 1000

Probability non-perforation maximized by distribution supported on minimal, not maximal, impact obliquity. Dirac on velocity at a non extreme value.

Important observations

Extremizers are singular

They identify key players
i.e. vulnerabilities of the physical system

Extremizers are attractors

Initialization with 3 support points per marginal

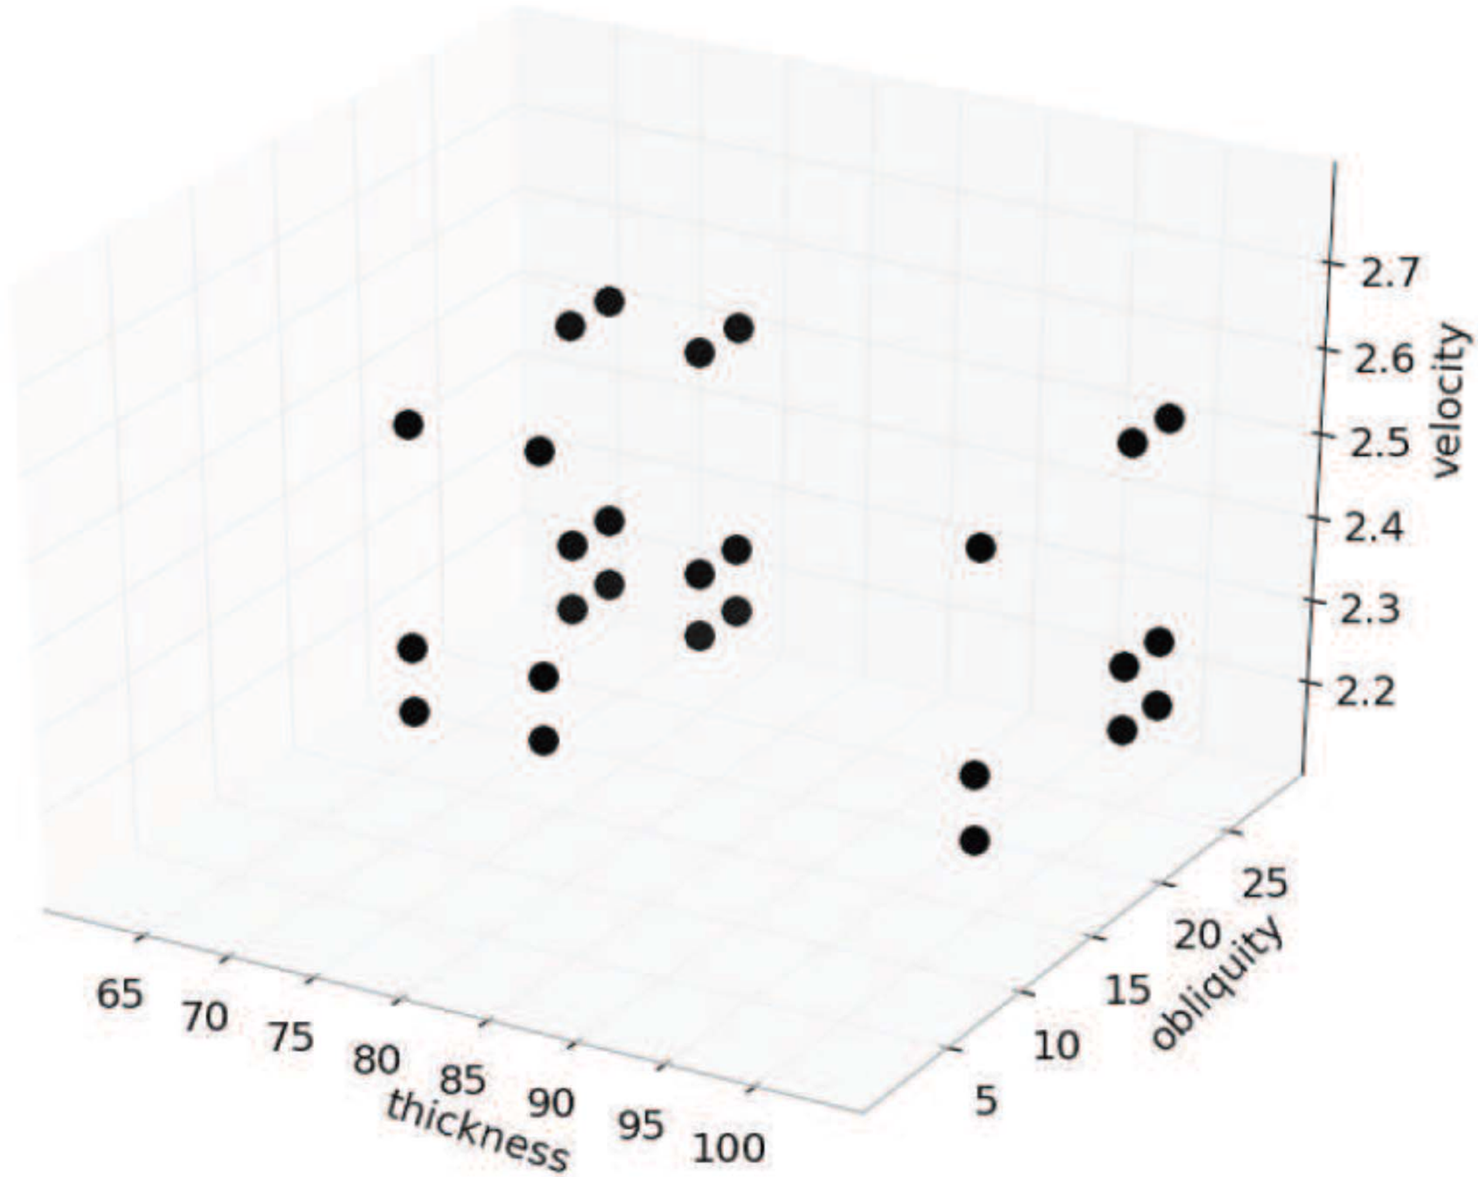

Support Points at iteration 0

Initialization with 3 support points per marginal

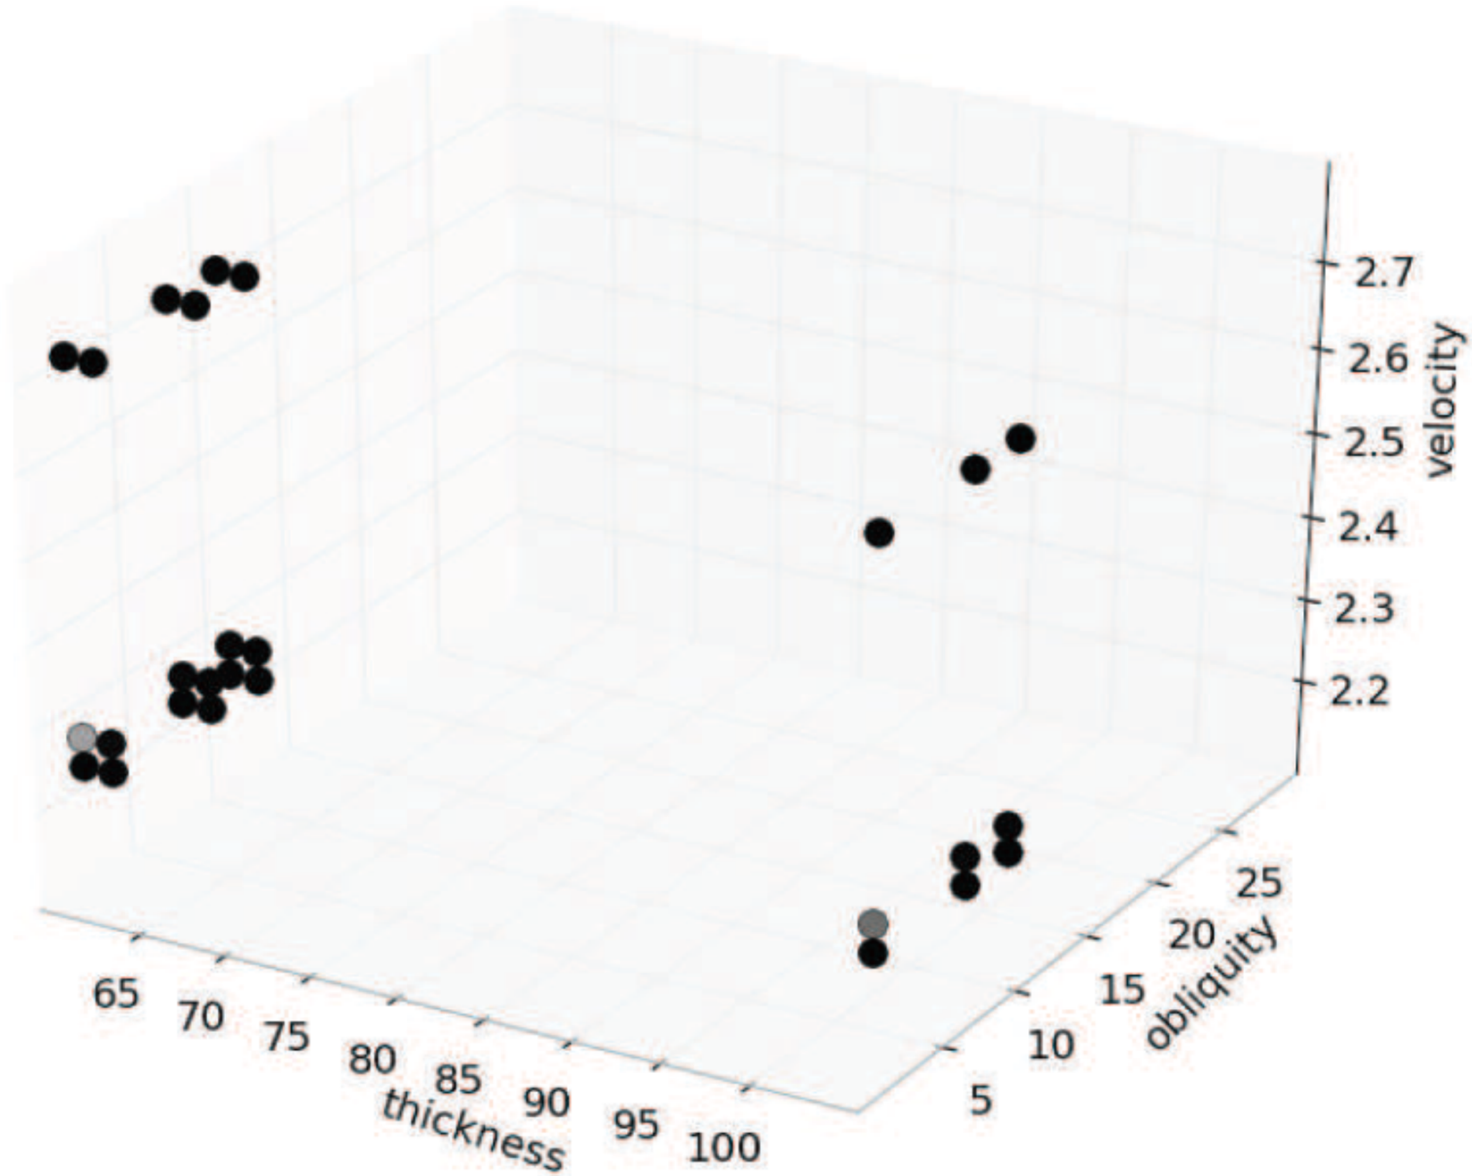

Support Points at iteration 500

Initialization with 3 support points per marginal

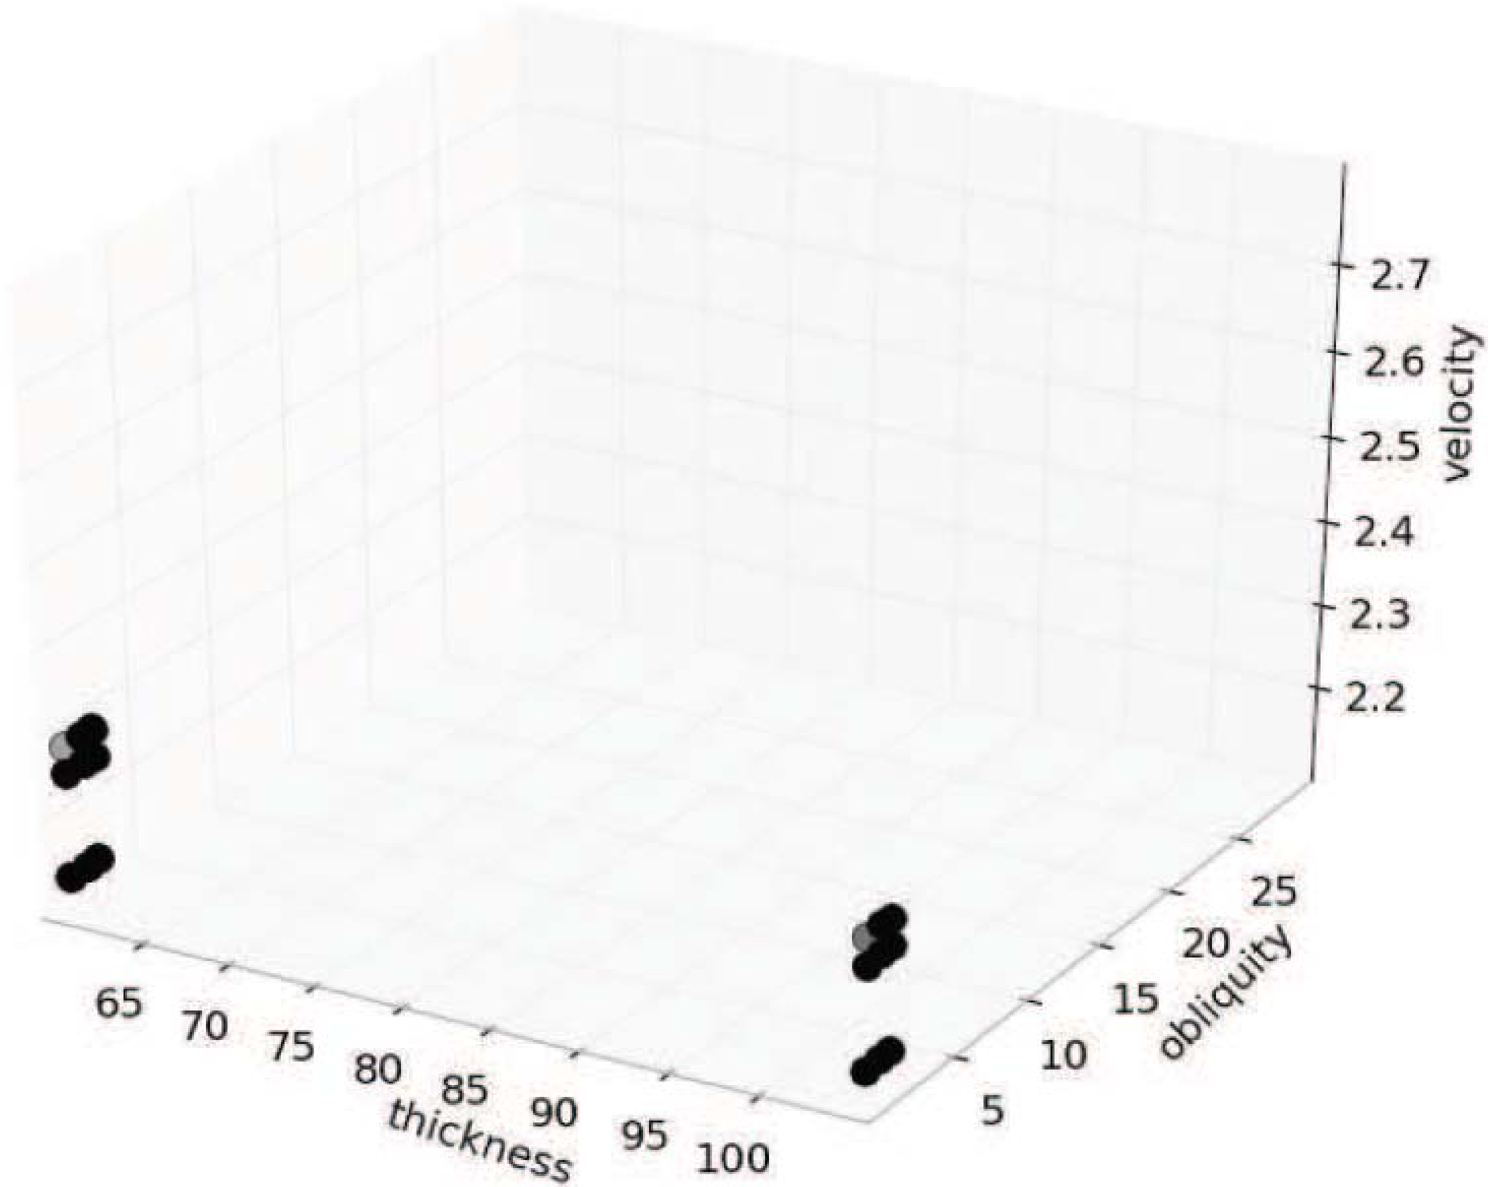

Support Points at iteration 1000

Initialization with 3 support points per marginal

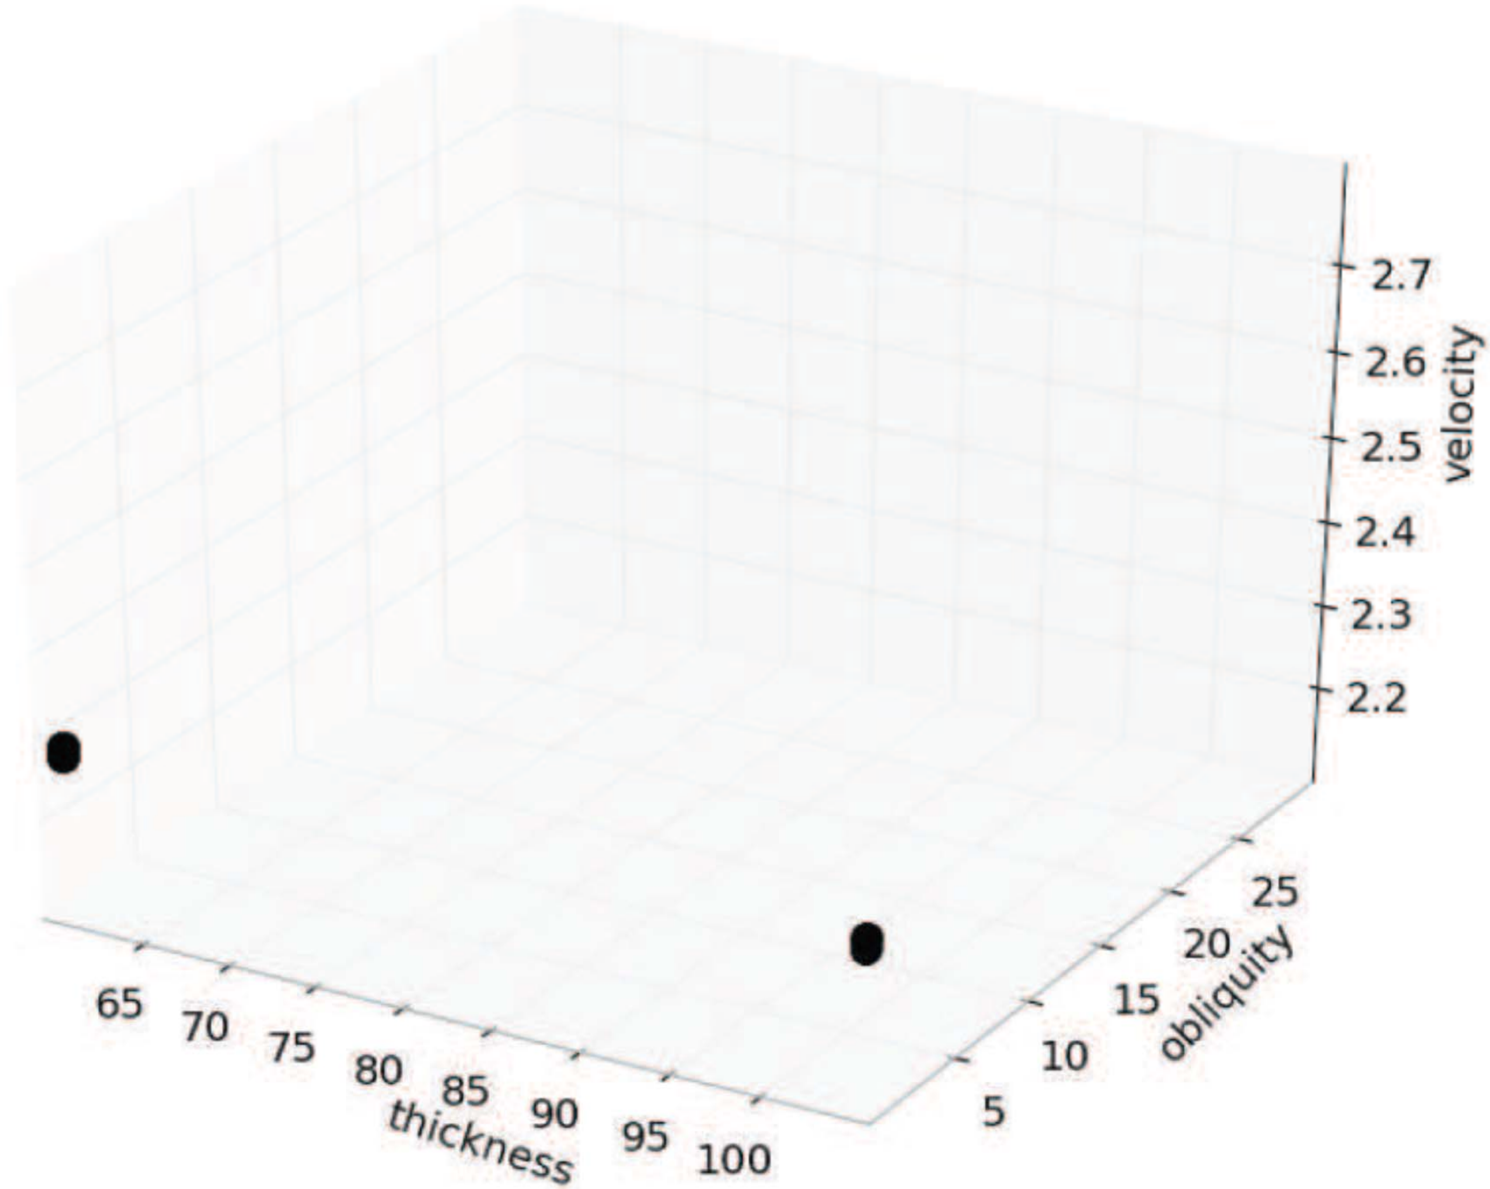

Support Points at iteration 2155

Initialization with 5 support points per marginal

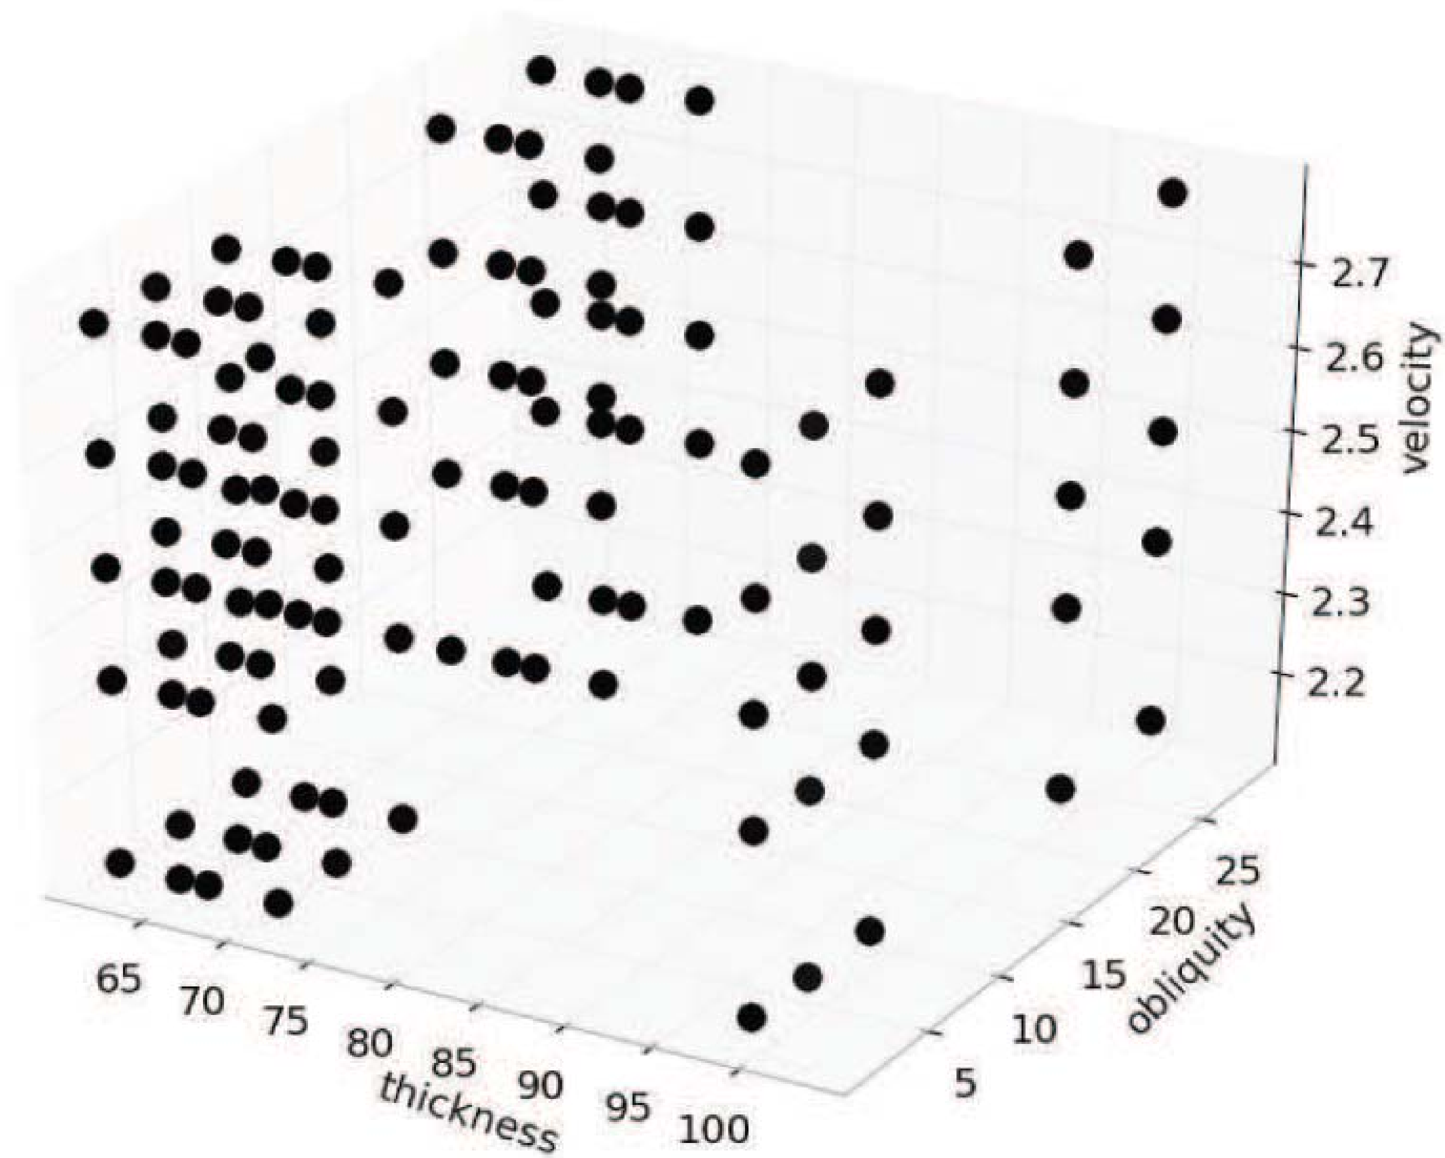

Support Points at iteration 0

Initialization with 5 support points per marginal

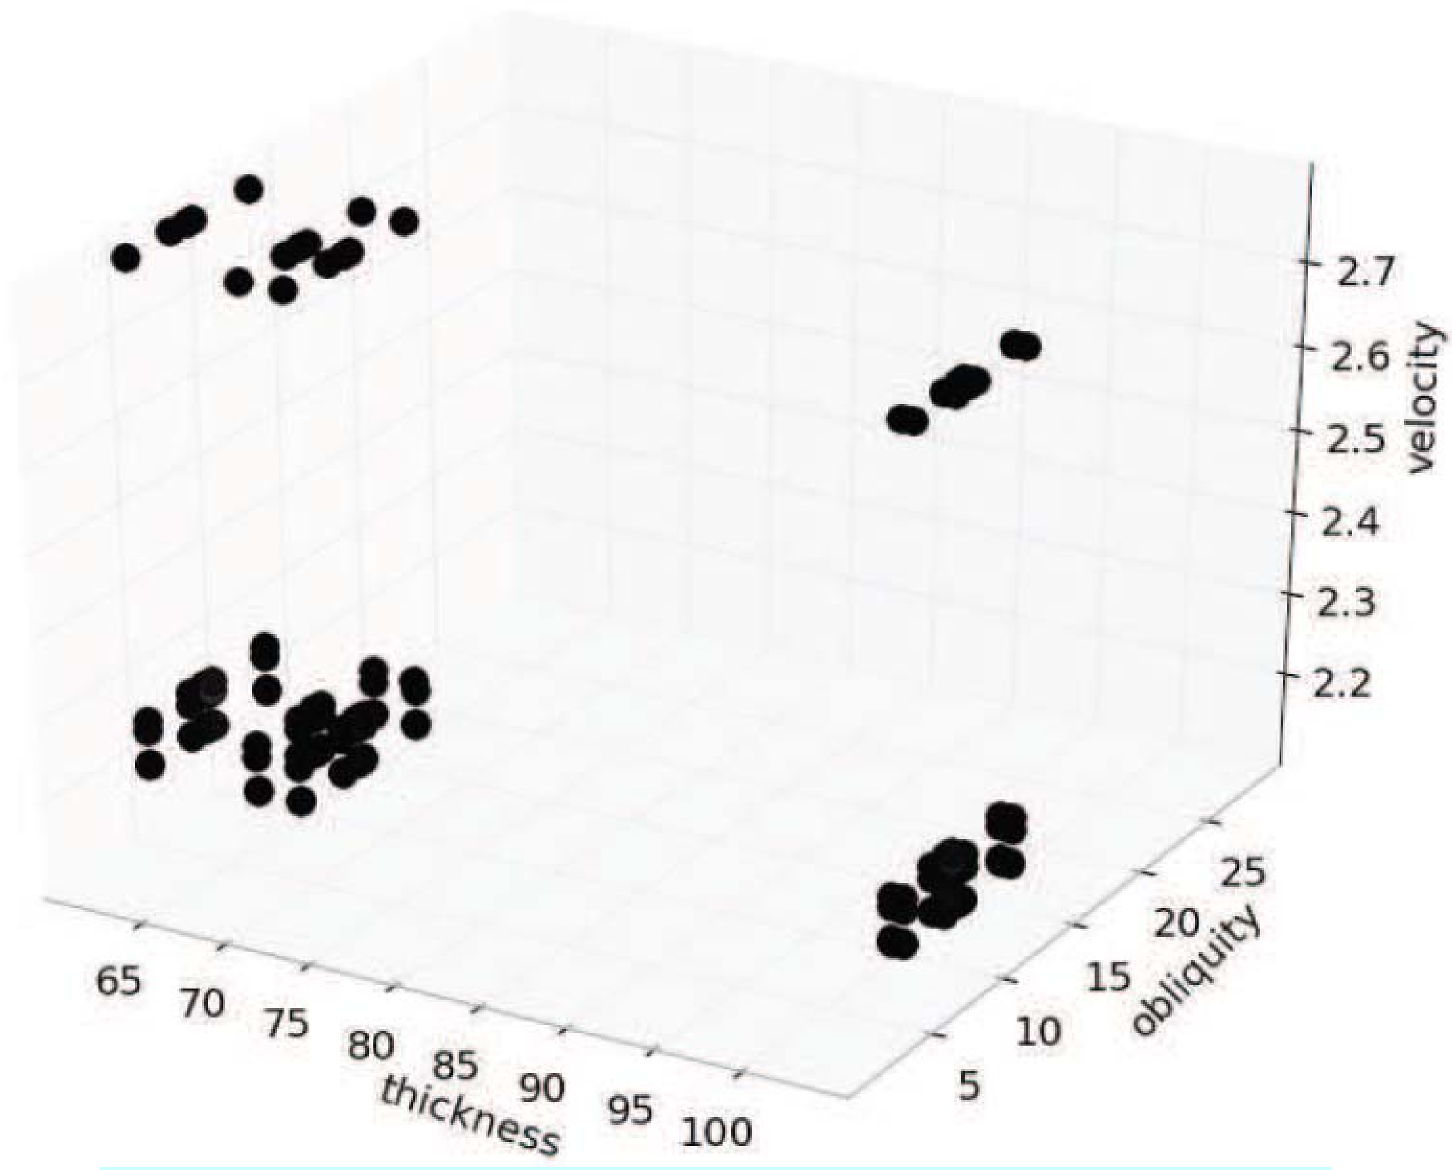

Support Points at iteration 1000

Initialization with 5 support points per marginal

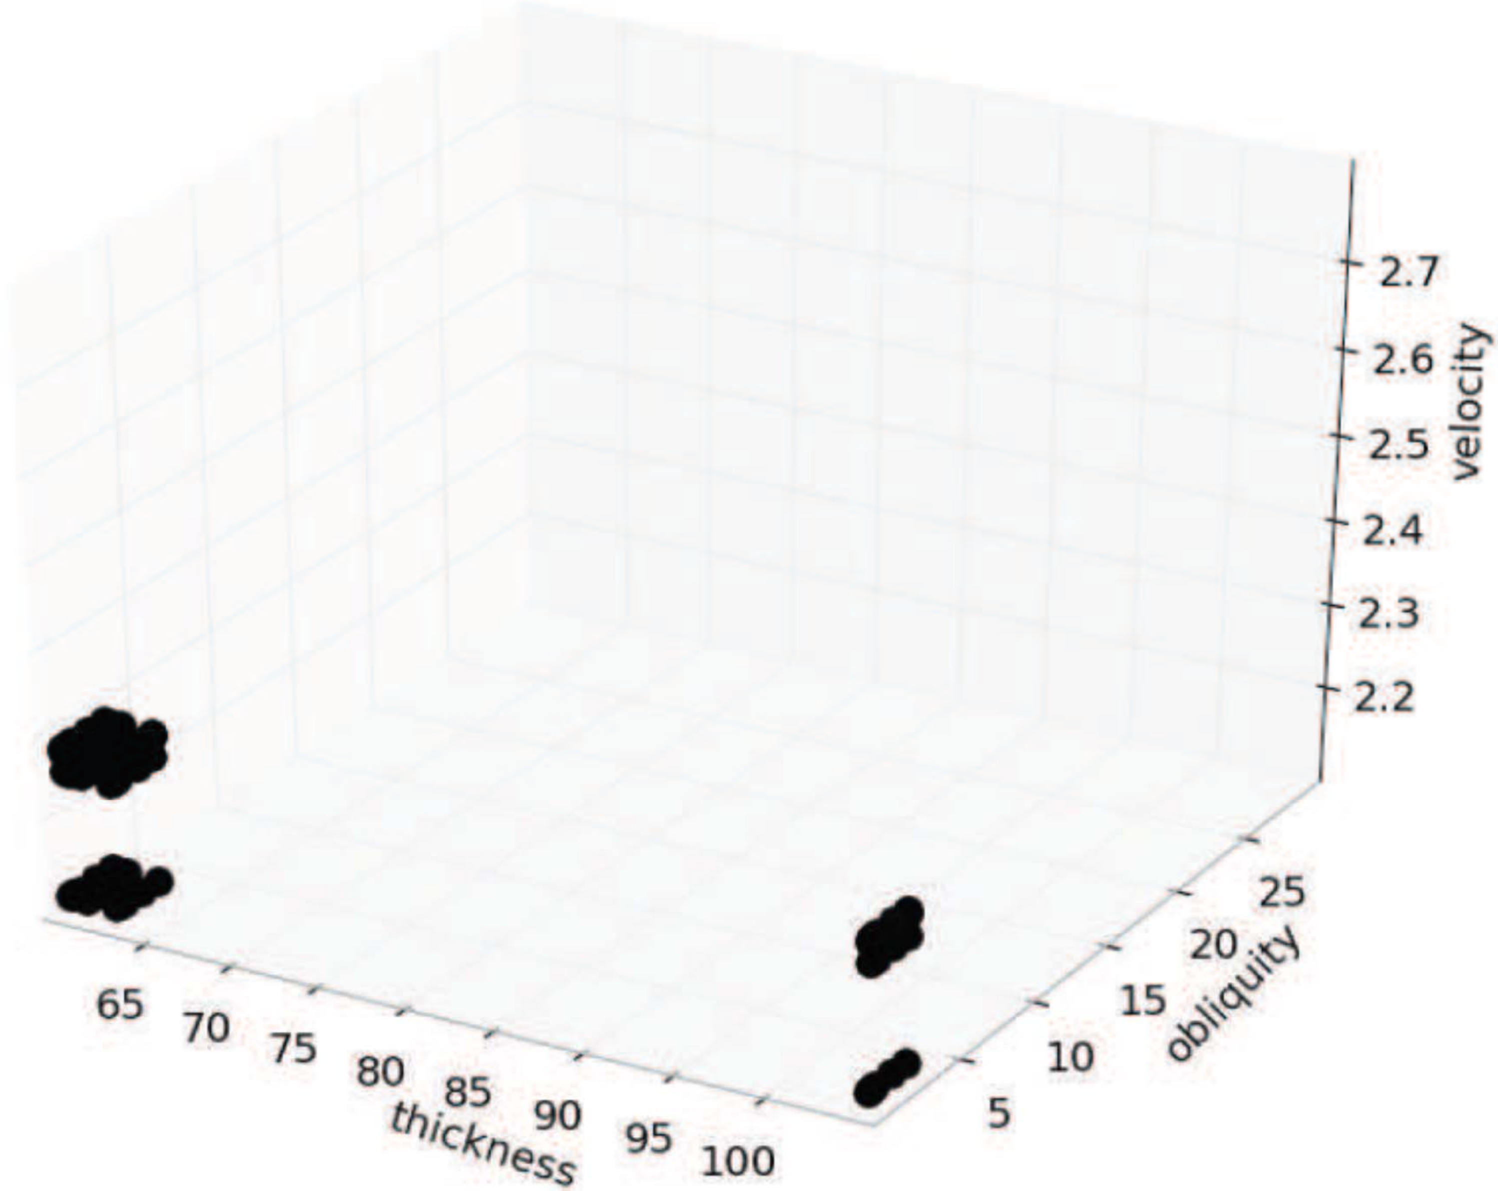

Support Points at iteration 3000

Initialization with 5 support points per marginal

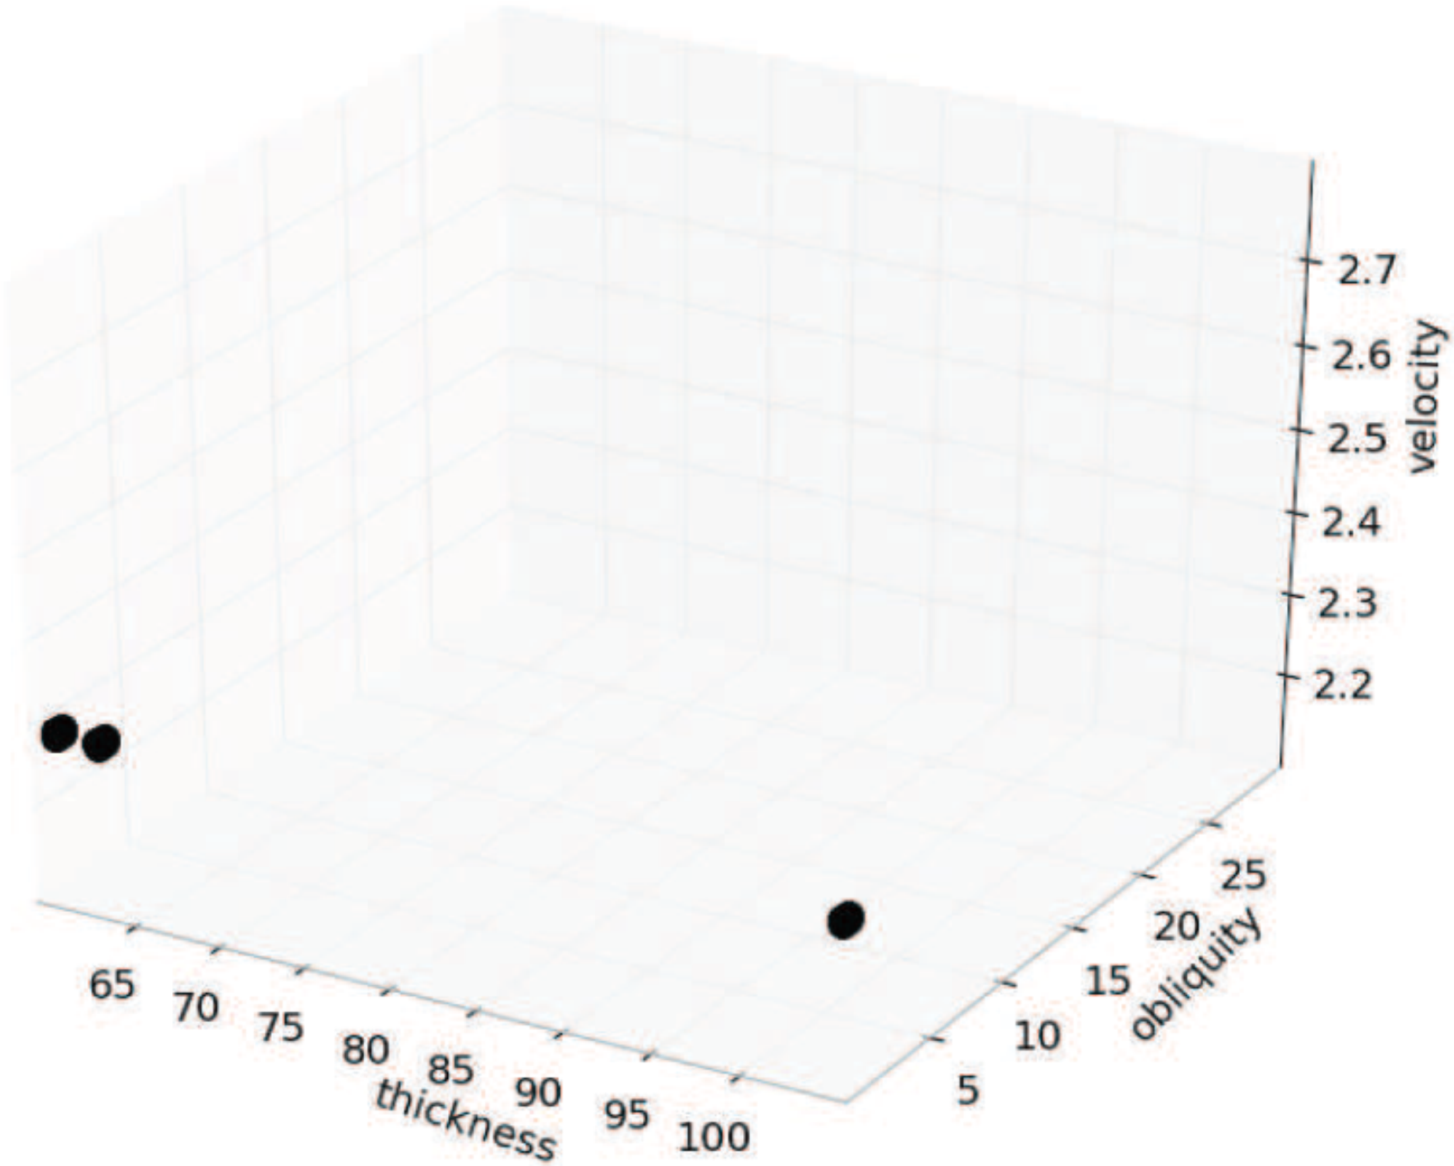

Support Points at iteration 7100

Previous examples

NO DATA

What if you have data?

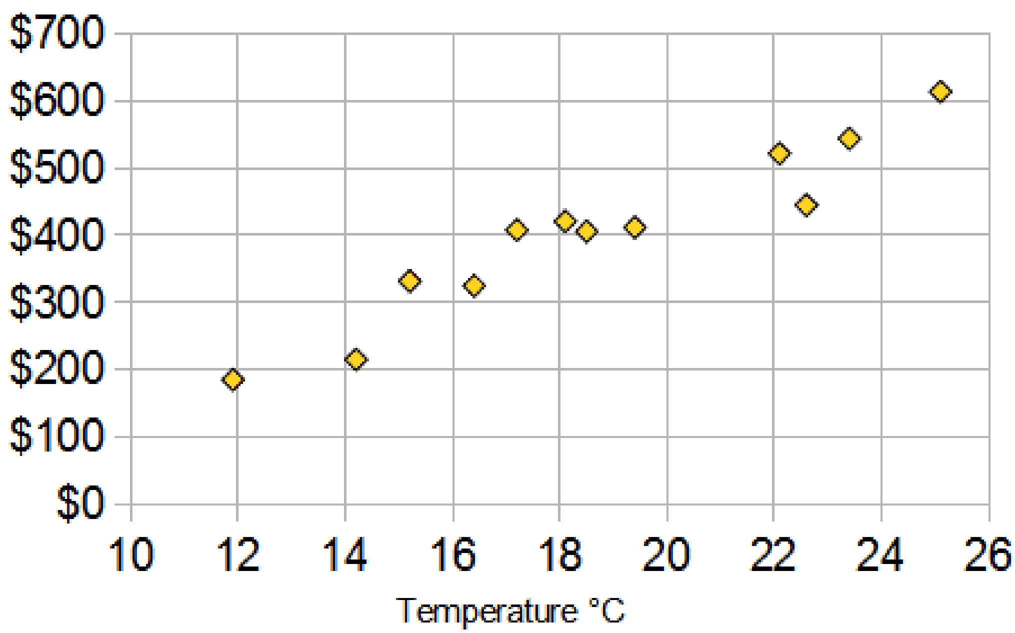

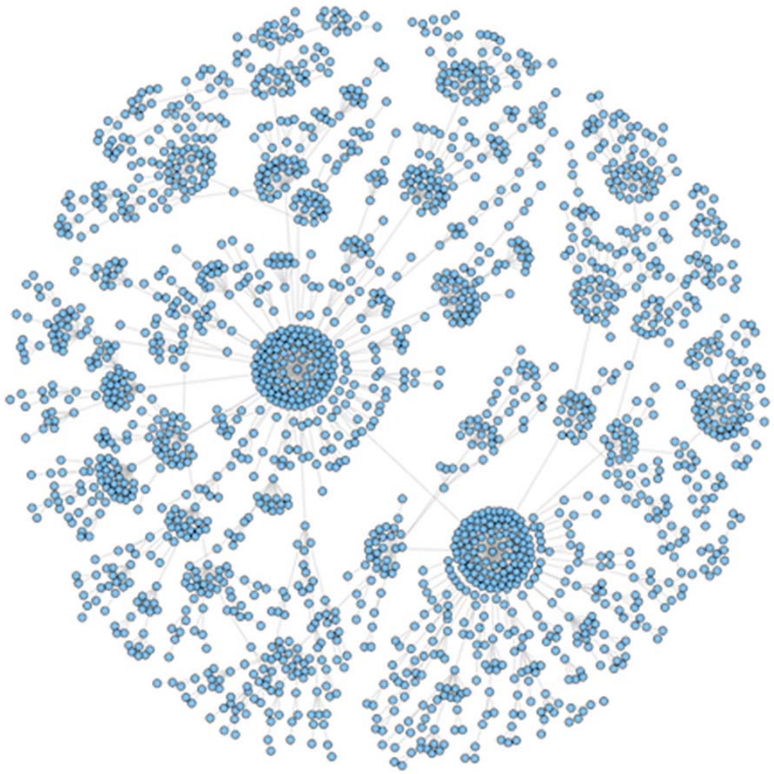

Previous examples

NO DATA

What if you have data?

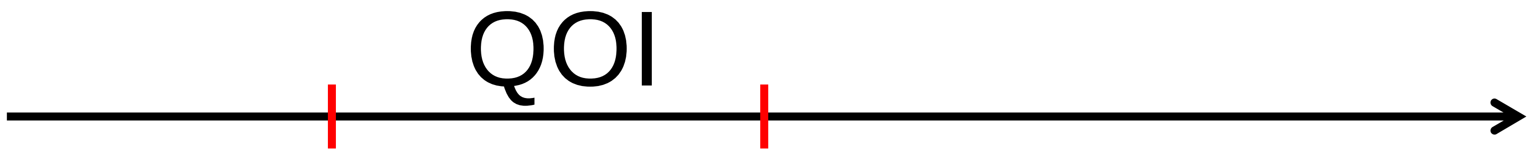

Optimal bounds become functions of the data (intervals of confidence)

How do we compute the best functions of the data?

How do we use computers to extract as much juice as possible from the data?

Scientific Computation of Optimal Statistical Estimators

Solving PDEs: Two centuries ago

$$\Delta u = f$$

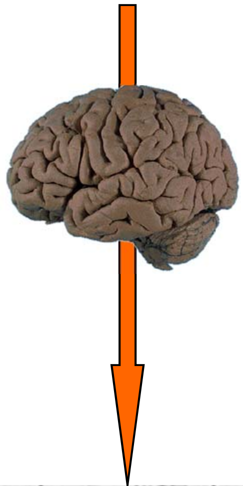

A. L. Cauchy
(1789-1857)

S. D. Poisson
(1781-1840)

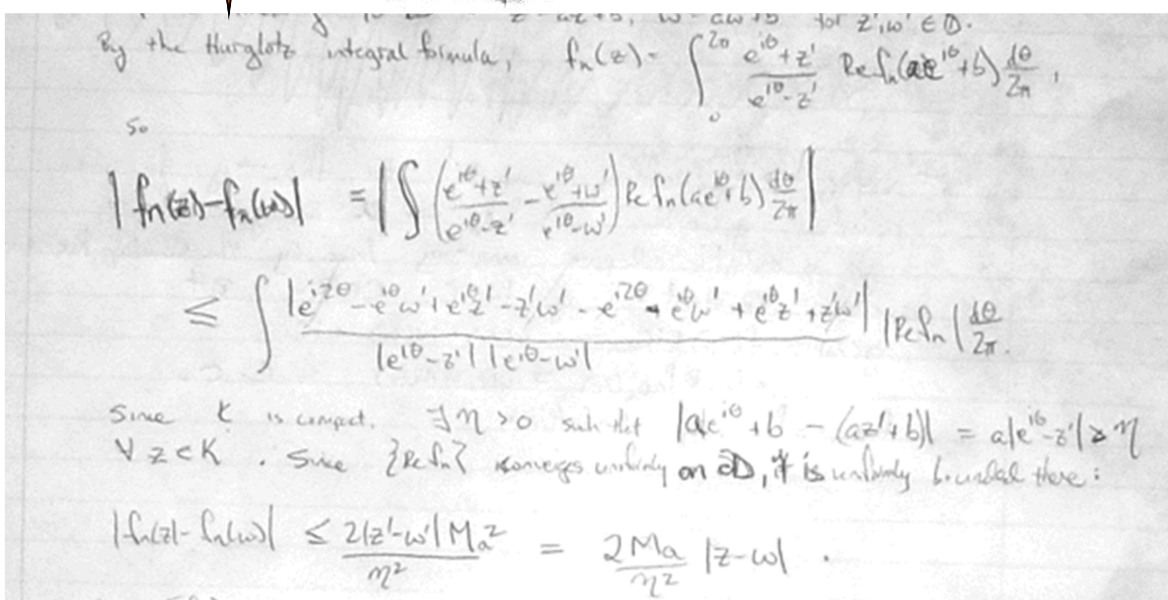

Solving PDEs: Now.

$$\Delta u = f$$

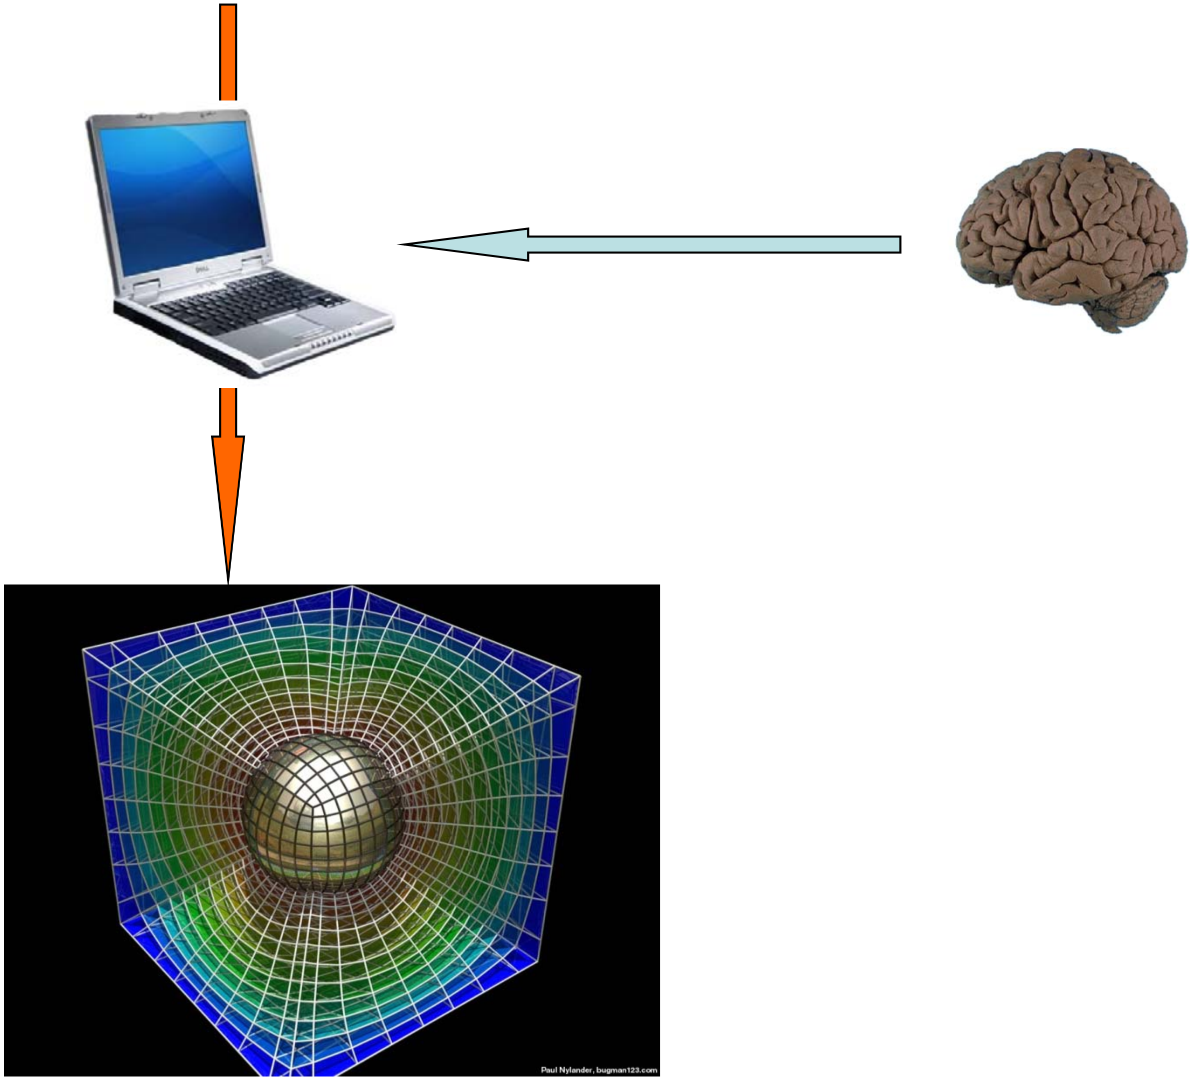

Paradigm shift

$$\Delta u = f$$

J. V. Neumann
(1903-1957)

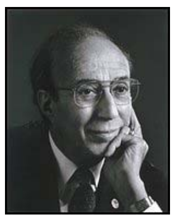

H. Goldstine
(1913-2004)

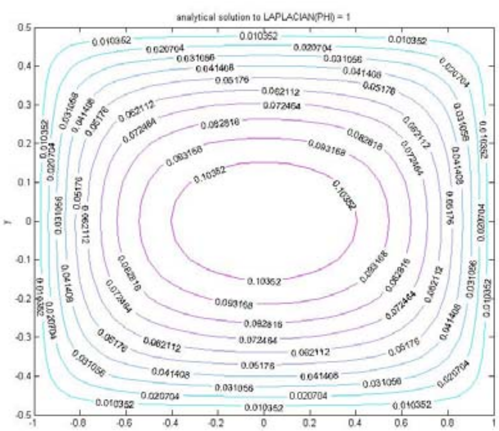

Where are we at in finding statistical estimators?

$$X^2 = \sum \frac{(o-e)^2}{e}$$

where

X^2 is Chi-squared,
$\sum$ stands for summation,
o is the observed values,
e is the expected values.

Percentage Points of the Chi-Square Distribution

Degrees of Freedom	Probability of a larger value of x^2							
	0.99	0.95	0.90	0.75	0.50	0.25	0.10	0.05
1	0.000	0.004	0.016	0.102	0.455	1.32	2.71	3.84
2	0.020	0.103	0.211	0.575	1.386	2.77	4.61	5.99
3	0.115	0.352	0.584	1.212	2.366	4.11	6.25	7.81
4	0.297	0.711	1.064	1.923	3.357	5.39	7.78	9.49
5	0.554	1.145	1.610	2.675	4.351	6.63	9.24	11.07
6	0.872	1.635	2.204	3.455	5.348	7.84	10.64	12.59
7	1.239	2.167	2.833	4.255	6.346	9.04	12.02	14.07
8	1.647	2.733	3.490	5.071	7.344	10.22	13.36	15.51
9	2.088	3.325	4.168	5.899	8.343	11.39	14.68	16.92
10	2.558	3.940	4.865	6.737	9.342	12.55	15.99	18.31
11	3.053	4.575	5.578	7.584	10.341	13.70	17.28	19.68
12	3.571	5.226	6.304	8.438	11.340	14.85	18.55	21.03
13	4.107	5.892	7.042	9.299	12.340	15.98	19.81	22.36
14	4.660	6.571	7.790	10.165	13.339	17.12	21.06	23.68
15	5.229	7.261	8.547	11.037	14.339	18.25	22.31	25.00
16	5.812	7.962	9.312	11.912	15.338	19.37	23.54	26.30
17	6.408	8.672	10.085	12.792	16.338	20.49	24.77	27.59
18	7.015	9.390	10.865	13.675	17.338	21.60	25.99	28.87
19	7.633	10.117	11.651	14.562	18.338	22.72	27.20	30.14
20	8.260	10.851	12.443	15.452	19.337	23.83	28.41	31.41
22	9.542	12.338	14.041	17.240	21.337	26.04	30.81	33.92
24	10.856	13.848	15.659	19.037	23.337	28.24	33.20	36.42
26	12.198	15.379	17.292	20.843	25.336	30.43	35.56	38.89
28	13.565	16.928	18.939	22.657	27.336	32.62	37.92	41.34
30	14.953	18.493	20.599	24.478	29.336	34.80	40.26	43.77
40	22.164	26.509	29.051	33.660	39.335	45.62	51.80	55.76
50	27.707	34.764	37.689	42.942	49.335	56.33	63.17	67.50
60	37.485	43.188	46.459	52.294	59.335	66.98	74.40	79.08

Find the best climate model given current information

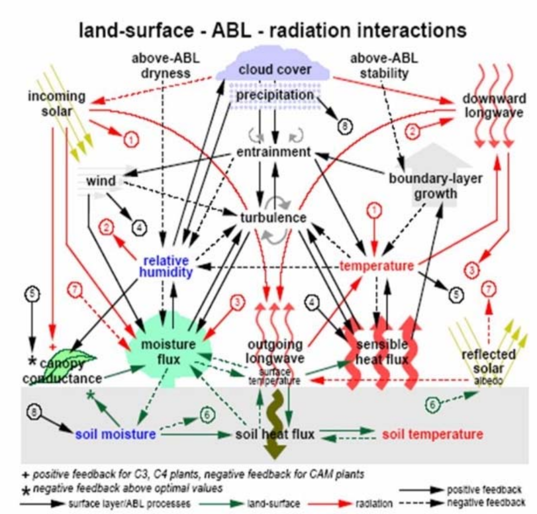

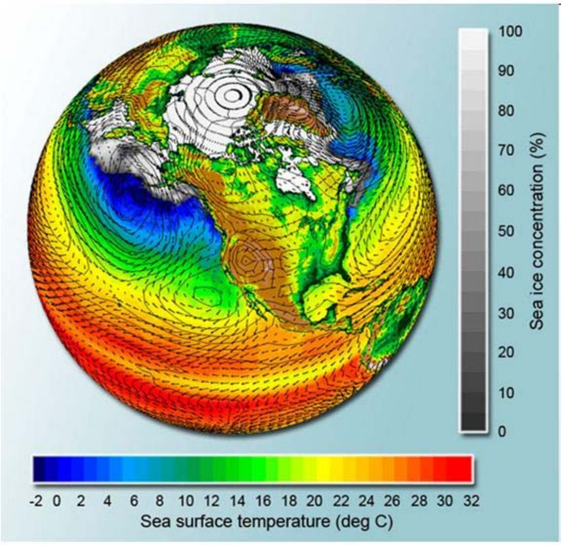

Exascale Co-Design Center for Materials in Extreme Environments

ExMatEx

Ab-initio Methods	Molecular Dynamics	Phase-Field Modeling	Continuum Methods
Inter-atomic force model, equation of state,	**Defect and interface mobility, nucleation**	**Direct numerical simulation of multi-phase evolution**	**Multi-phase material response, experimental observables**
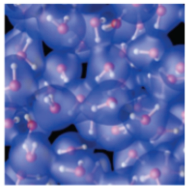	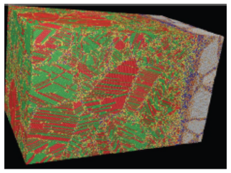	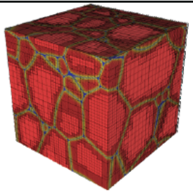	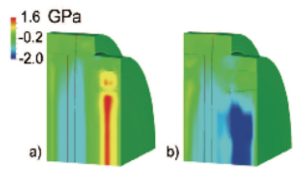
Code: Qbox/LATTE Motif: Particles and wavefunctions, plane wave DFT with nonlocal norm-conserving, ScaLAPACK, BLACS, and custom parallel 3D FFTs Prog. Model: MPI	Code: SPaSM/ddcMD Motif: Particles, domain decomposition, explicit time integration, neighbor and linked lists, dynamic load balancing, parity error recovery, and *in situ* visualization Prog. Model: MPI + Threads	Code: AMPE/GL Motif: Regular and adaptive grids, implicit time integration, real-space and spectral methods, complex order parameter (phase, crystal, species) Prog. Model: MPI	Code: VP-FFT/ALE3d Motif: Regular and irregular grids, implicit time integration, 3D FFTs, polycrystal and simgle crystal plasticity, Prog. Model: MPI

Where are we at in finding statistical estimators?

Find the best estimator or model

$$\theta$$

+ (sample) data

$$\theta(data)$$

Can we turn model design into a computation?

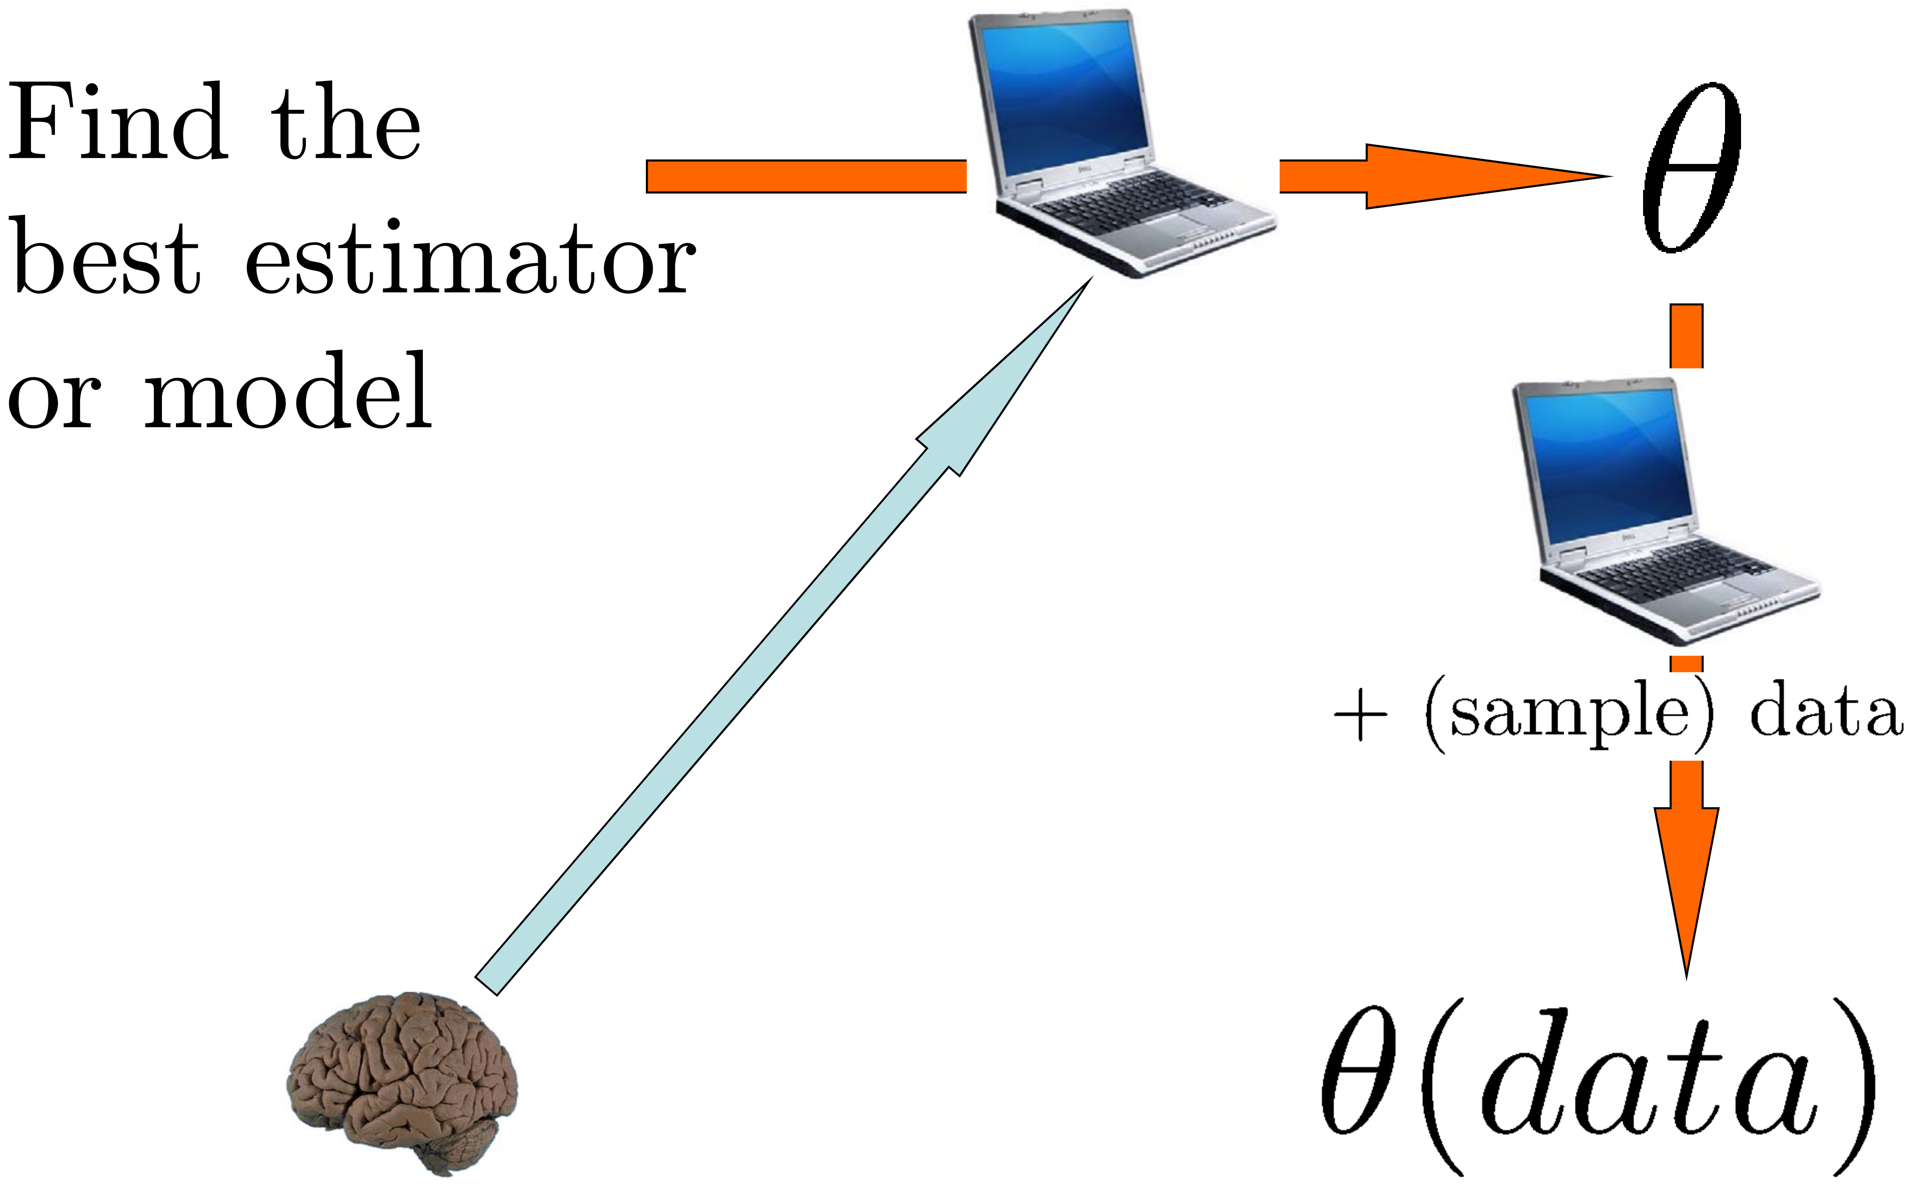

The UQ Problem with sample data

We want to estimate $\Phi(\mu^\dagger) = \mu^\dagger[X \geq a]$

$\mu^\dagger$: Unknown or partially known measure of probability on $\mathbb{R}$

You know $\mu^\dagger \in \mathcal{A}$

We observe $d = (d_1, \ldots, d_n) \in \mathbb{R}^n$

n i.i.d samples from $\mu^\dagger$

Your estimation: function of the data $\theta(d)$

Estimation error $\theta(d) - \Phi(\mu^\dagger)$

Statistical Error

$$\mathcal{E}(\theta, \mu^\dagger) = \mathbb{E}_{d \sim (\mu^\dagger)^n}\left[\left[\theta(d) - \Phi(\mu^\dagger)\right]^2\right]$$

Optimal bound on the statistical error

$$\max_{\mu \in \mathcal{A}} \mathcal{E}(\theta, \mu)$$

Optimal statistical estimators

$$\min_{\theta} \max_{\mu \in \mathcal{A}} \mathcal{E}(\theta, \mu)$$

Game theory and statistical decision theory

John Von Neumann

Abraham Wald

You

The universe

Estimator

θ

Measure of probability

μ

Loss/Statistical Error

$$\mathcal{E}(\theta, \mu)$$

Minimize

Maximize

Computer

Estimator

θ

The universe

Measure of probability

μ

Loss/Statistical Error

$$\mathcal{E}(\theta, \mu)$$

Minimize

Maximize

The space of admissible scenarios along with the space of relevant information, assumptions, beliefs and models tend to be infinite dimensional, whereas calculus on a computer is necessarily discrete and finite

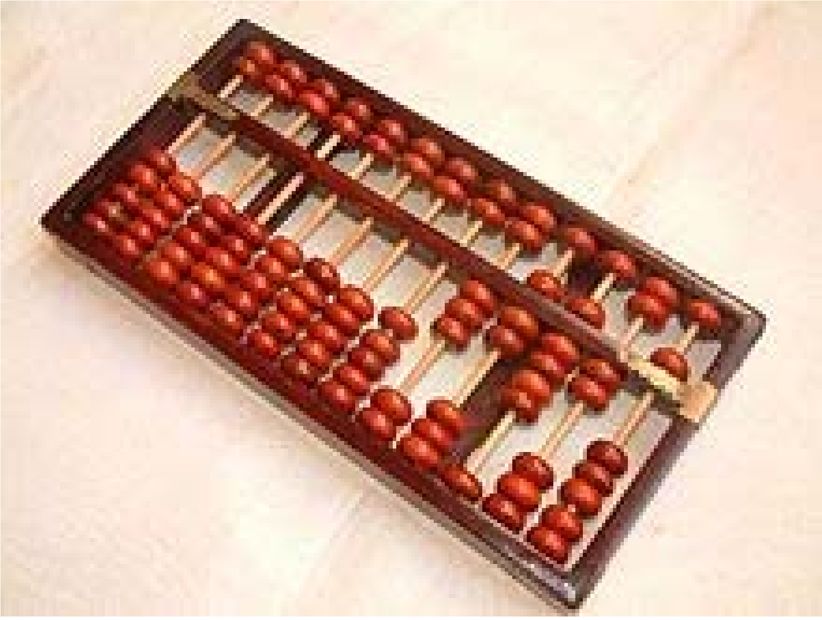

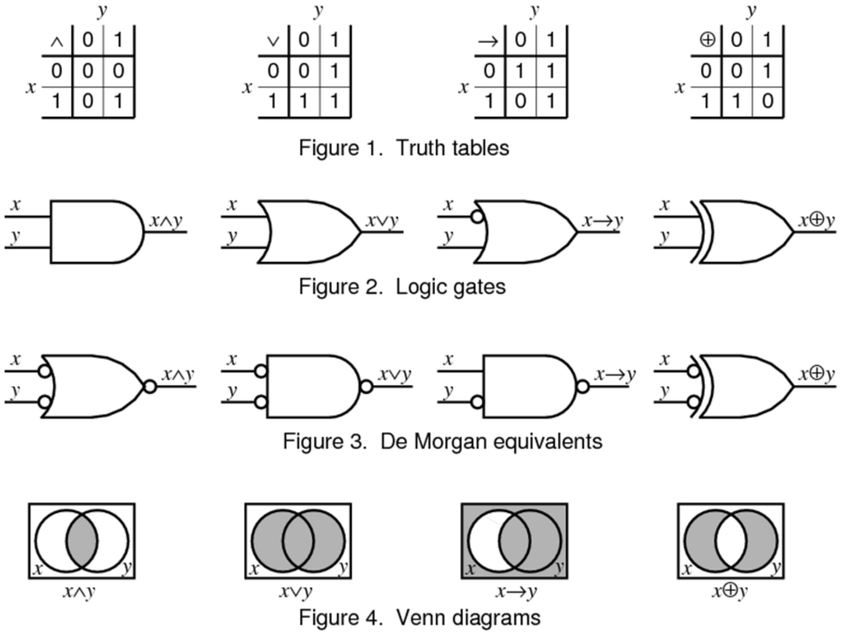

Arithmetic and Boolean logic

We need a form of calculus allowing us to manipulate infinite dimensional information structures

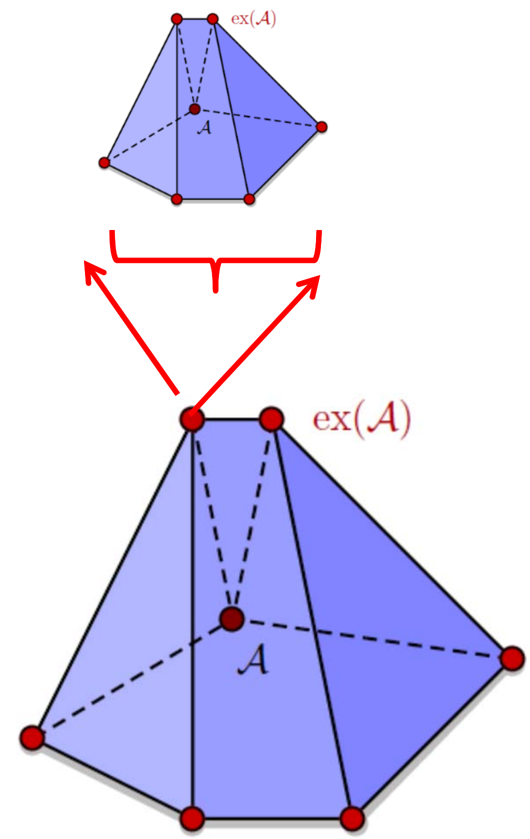

New form of reduction calculus
A simple example

Paul is given one pound of play-doh.
What can you say about how much mass he is putting above a if all you have is the belief that he is keeping the seesaw balanced around m?

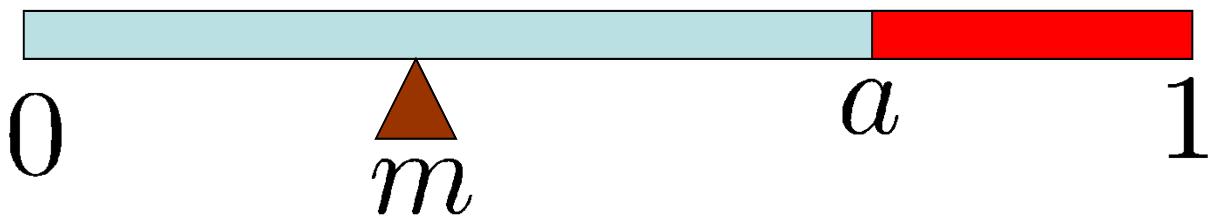

10,000 children are given one pound of play-doh. On average, how much mass can they put above a While, on average, keeping the seesaw balanced around m?

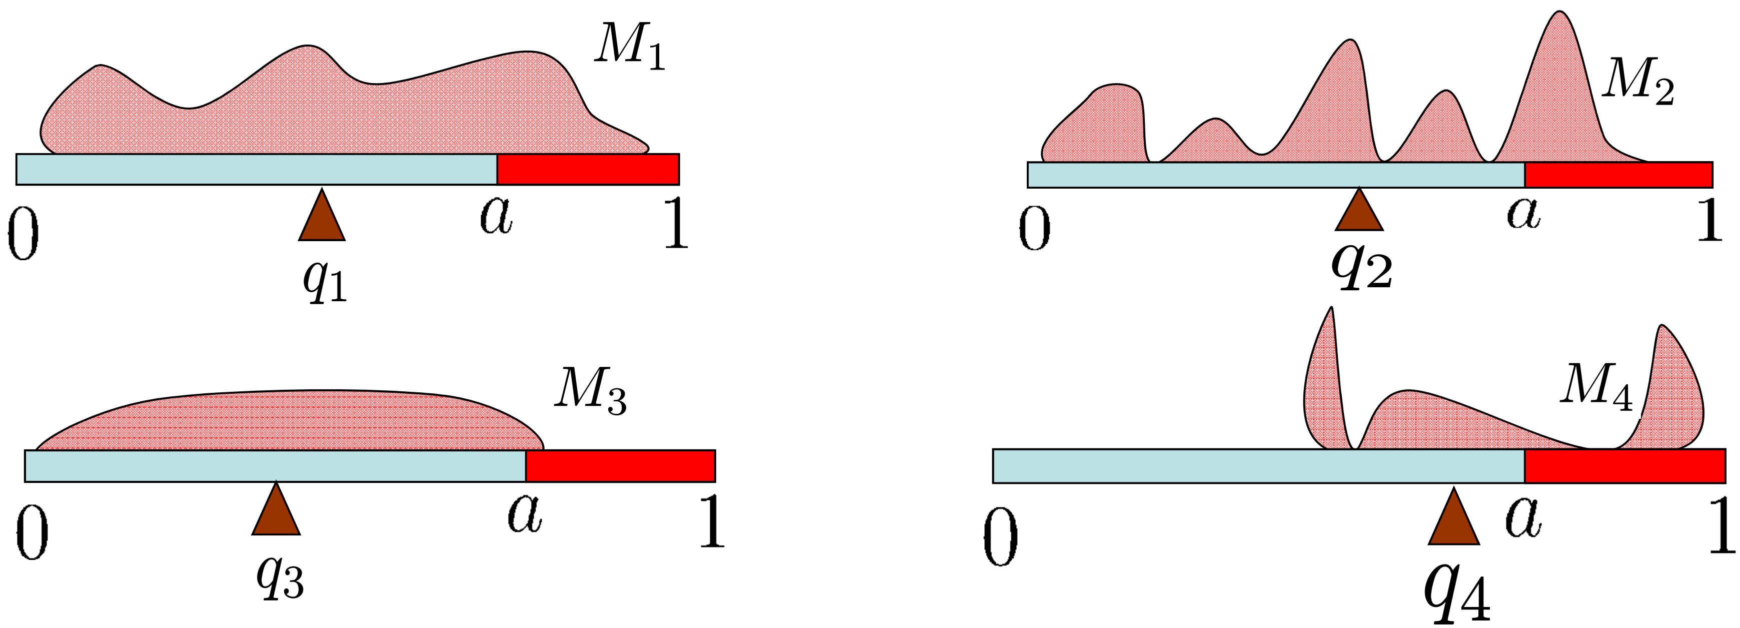

$$\frac{q_1 + \cdots + q_{10,000}}{10,000} = m$$

$$\max \frac{M_1 + \cdots + M_{10,000}}{10,000} \quad ?$$

What is the least upper bound on

$$\mathbb{E}_{\mu\sim\pi}\big[\mu[X \geq a]\big]$$

If all you know is $\mathbb{E}_{\mu\sim\pi}\big[\mathbb{E}_\mu[X]\big] = m$?

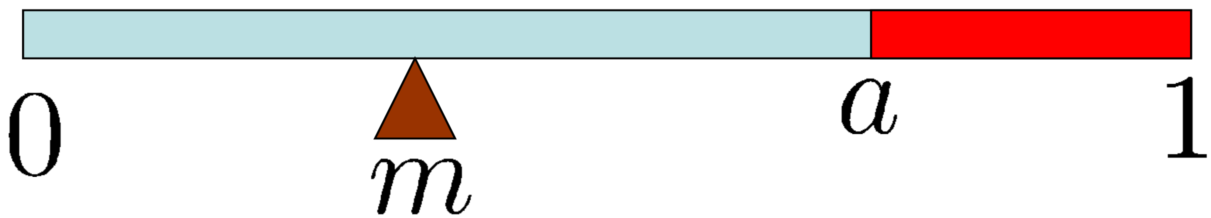

$$\mu \in \mathcal{A} := \mathcal{M}\big([0,1]\big)$$

Answer

$$\sup_{\pi\in\Pi} \mathbb{E}_{\mu\sim\pi}\big[\mu[X \geq a]\big]$$

$$\Pi := \Big\{\pi \in \mathcal{M}(\mathcal{A}) \,:\, \mathbb{E}_{\mu\sim\pi}\big[\mathbb{E}_\mu[X]\big] = m\Big\}$$

$$\sup_{\pi \in \Pi} \mathbb{E}_{\mu\sim\pi}\left[\mu[X \geq a]\right]$$

$$\Pi := \left\{\pi \in \mathcal{M}\big(\mathcal{M}([0,1])\big) \ : \ \mathbb{E}_{\mu\sim\pi}\left[\mathbb{E}_{\mu}[X]\right] = m\right\}$$

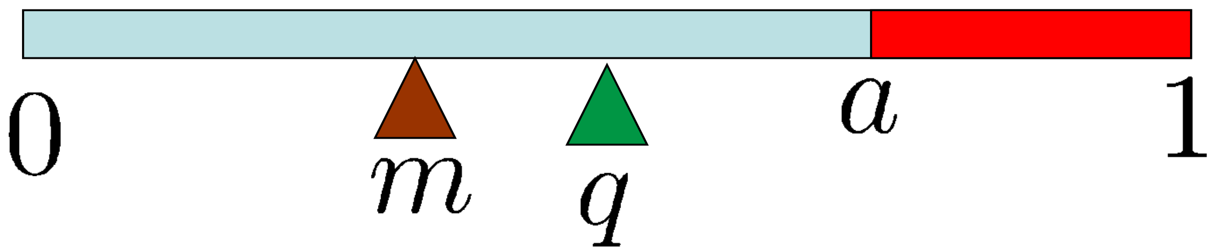

$$\sup_{\pi \in \Pi} \mathbb{E}_{\mu\sim\pi}\left[\mu[X \geq a]\right] = \sup_{\mathbb{Q}\in\mathcal{M}([0,1])\,:\,\mathbb{E}_{\mathbb{Q}}[q]=m} \mathbb{E}_{q\sim\mathbb{Q}}\left[\sup_{\mu\in\mathcal{M}([0,1])\,:\,\mathbb{E}_{\mu}[X]=q} \mu[X \geq a]\right]$$

$$\sup_{\pi \in \Pi} \mathbb{E}_{\mu \sim \pi}\left[\mu[X \geq a]\right]$$

$$\Pi := \left\{\pi \in \mathcal{M}\big(\mathcal{M}([0,1])\big) \;:\; \mathbb{E}_{\mu\sim\pi}\left[\mathbb{E}_{\mu}[X]\right] = m\right\}$$

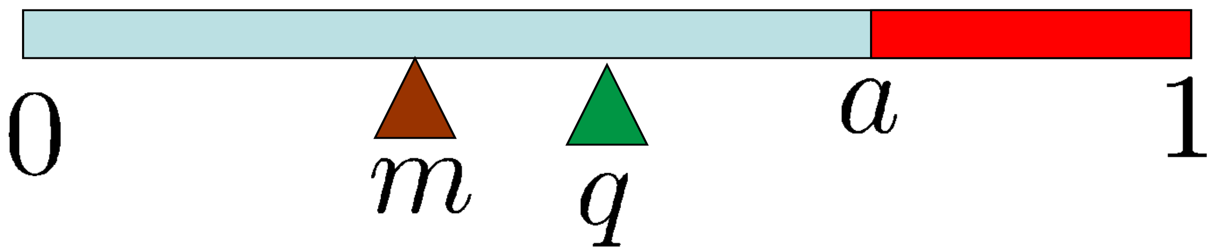

$$\sup_{\pi \in \Pi} \mathbb{E}_{\mu \sim \pi}\left[\mu[X \geq a]\right] = \sup_{\mathbb{Q} \in \mathcal{M}([0,1])\,:\,\mathbb{E}_{\mathbb{Q}}[q] = m} \mathbb{E}_{q \sim \mathbb{Q}}\left[\min\left(\frac{q}{a}, 1\right)\right]$$

$$\sup_{\pi \in \Pi} \mathbb{E}_{\mu \sim \pi} \left[\mu[X \geq a] \right]$$

$$\Pi := \left\{ \pi \in \mathcal{M}(\mathcal{M}([0,1])) \ : \ \mathbb{E}_{\mu \sim \pi} \left[\mathbb{E}_{\mu}[X] \right] = m \right\}$$

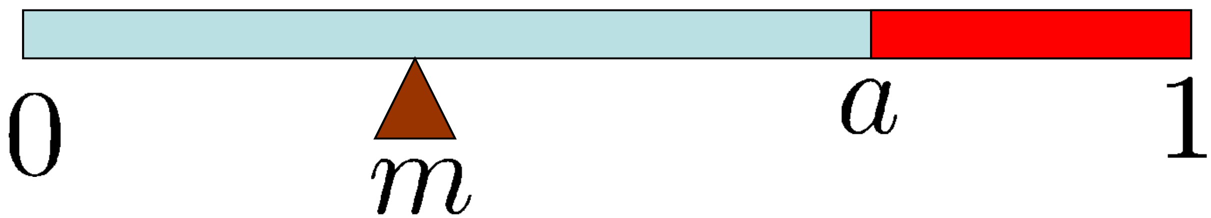

$$\sup_{\pi \in \Pi} \mathbb{E}_{\mu \sim \pi} \left[\mu[X \geq a] \right] = \frac{m}{a}$$

Why develop this form of calculus? What else could we do?

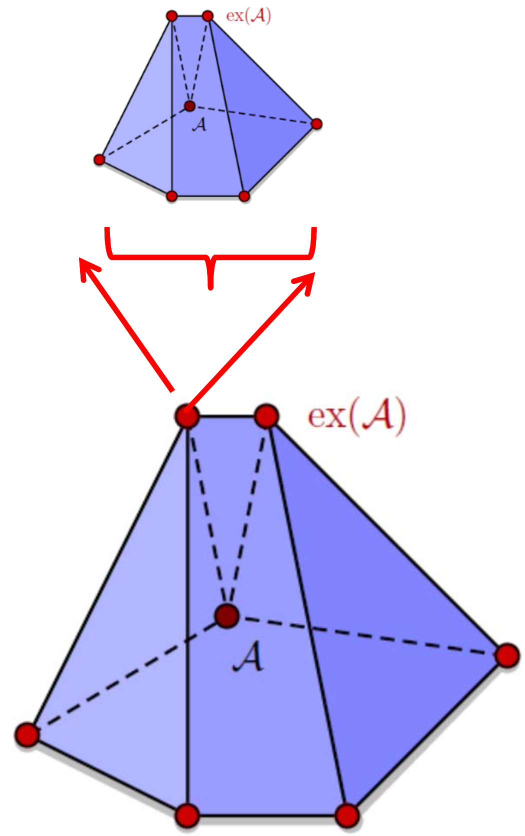

Computer

Estimator

θ

The universe

Measure of probability

μ

Loss/Statistical Error

$$\mathcal{E}(\theta, \mu)$$

Minimize

Maximize

Min/Max Tree

Allows you to design optimal experimental campaigns
and turn the process of scientific discovery into a computation

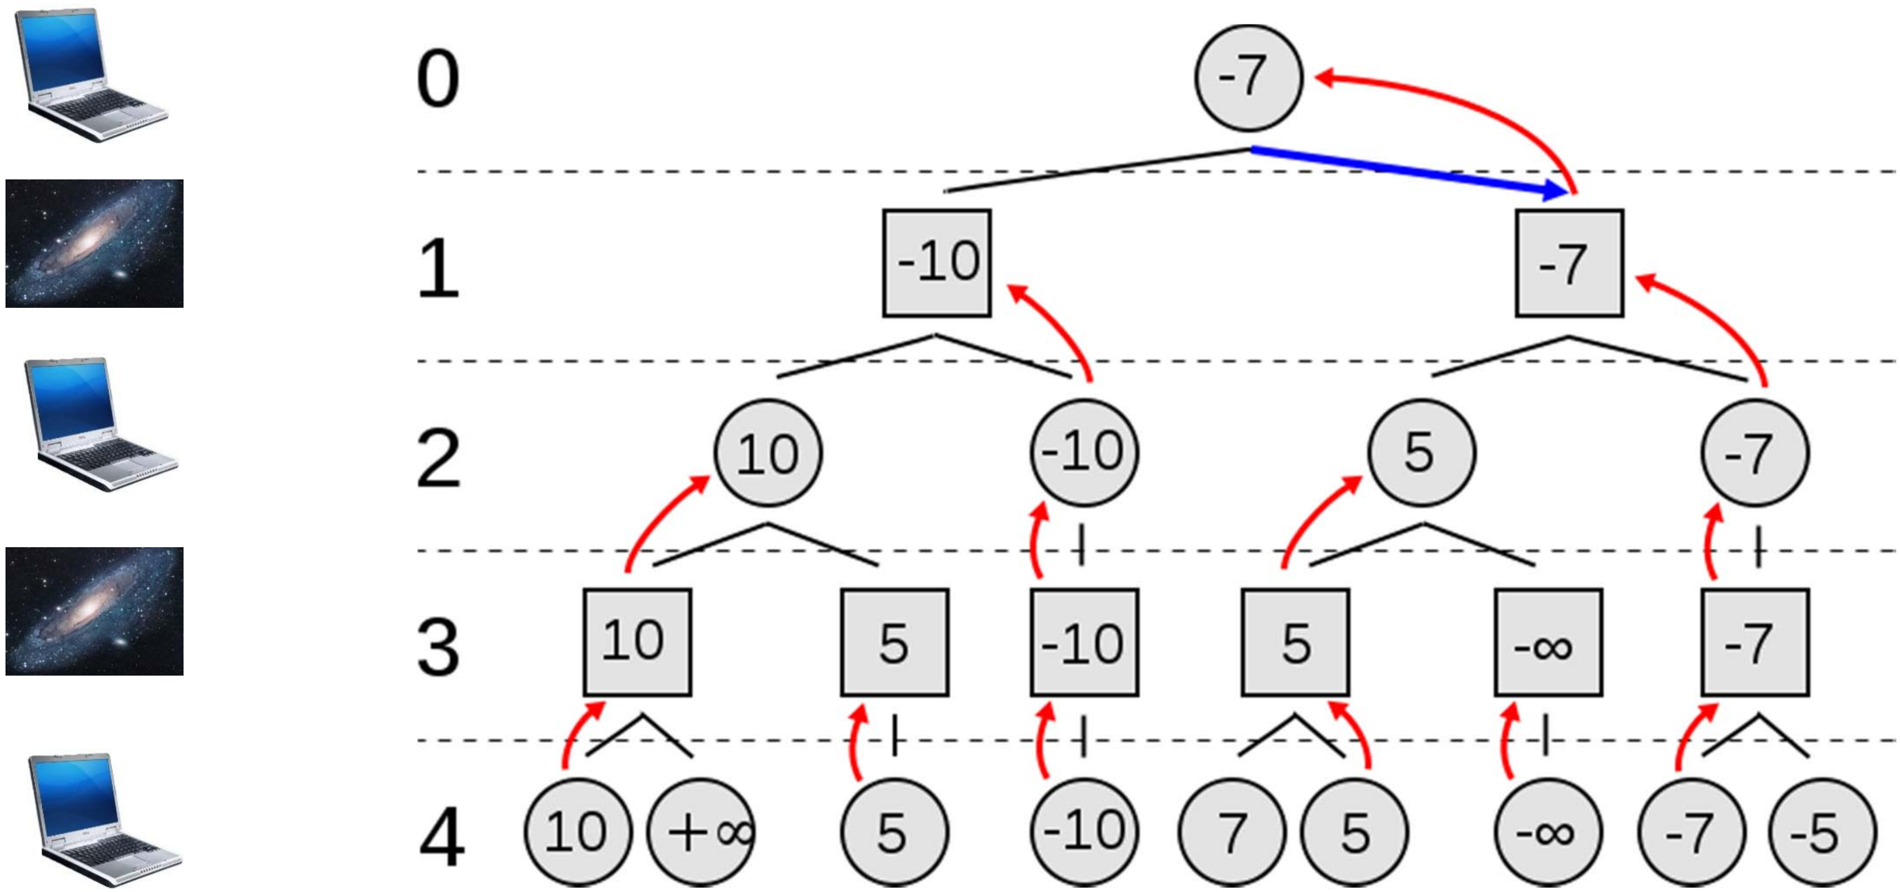

Machine learning

Develop the best model of reality given available information

Act based on That model

Gather new information

Provide the ability to compute optimal strategies in information games/wars

Provide the ability to make optimal decisions with regards to information available to and assumptions made by participants in a conflict.

MDA STTR solicitation

Application: MDA
Missile Defense Strategies

Quantifying MDA's confidence in M&S-based predictions of BMDS performance

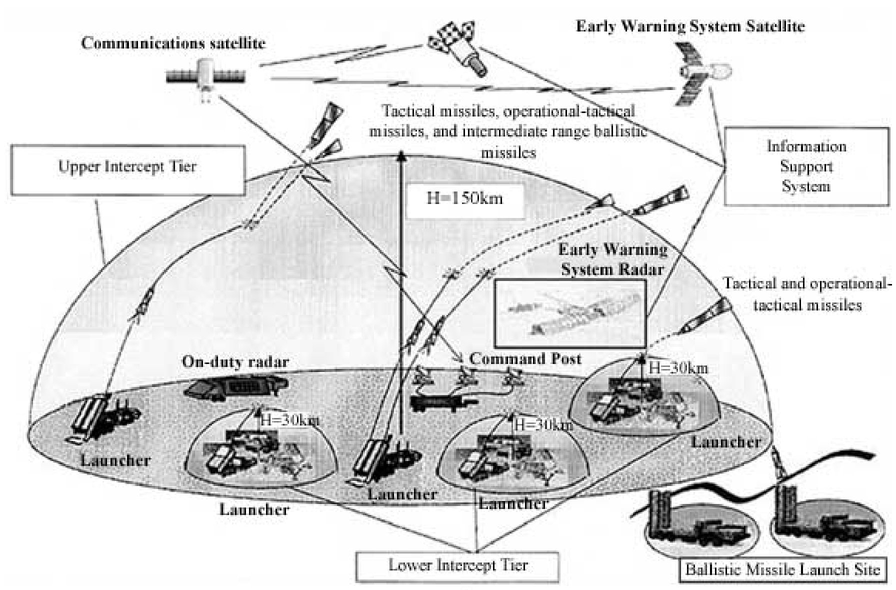

Well posed problem

UQ is the business of computing optimal bounds or making optimal predictions on quantities of interest given available information

It is about the optimal processing of information

Why optimal? 2010 Iceland ash cloud
Under-estimate risk = Loss of Life
Over-estimate risk = Economic Loss

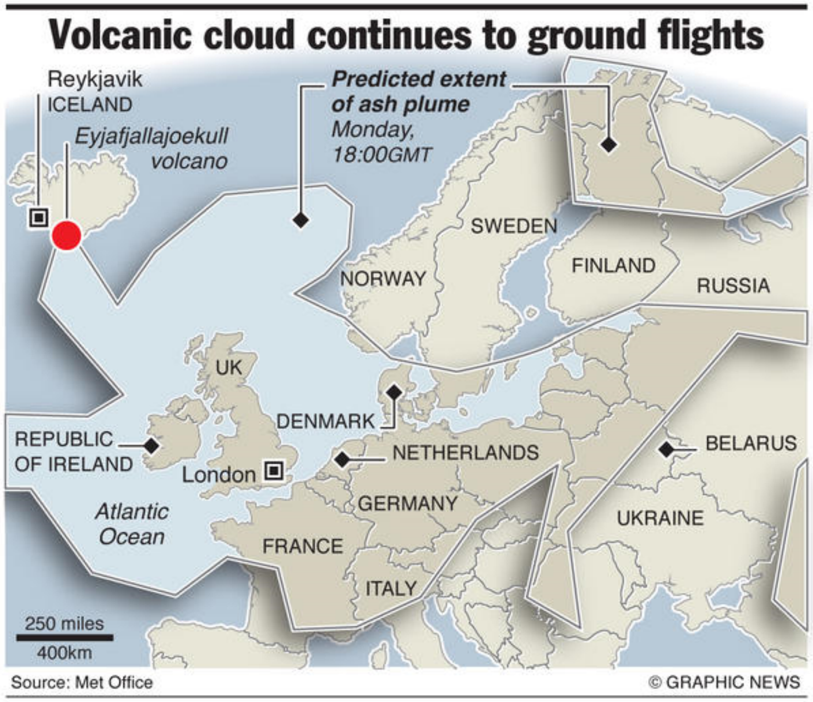

Airlines want payouts for 'overreaction'

KATHERINE HADDON

PUBLISHED: 2010/04/23 07:38:09 AM

AS EUROPE's airspace reopened and weary passengers boarded long- delayed flights home, airline executives pressed for government compensation to cover the industry's huge losses for what some deemed an overreaction by governments.

"Airspace was being closed based on theoretical models, not on facts"

Why turn the UQ challenge into a well posed mathematical question

What is the meaning of life?

42

The UQ challenge in the prediction context (Deepwater Horizon Disaster)

You want to find a 95% interval of confidence on the spill rate

6 different techniques lead to 6 distinct predictions with non-overlapping 95% confidence intervals

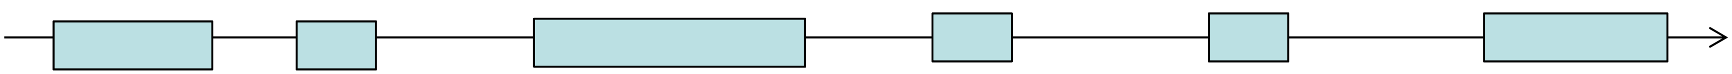

You may as well ask Paul the octopus

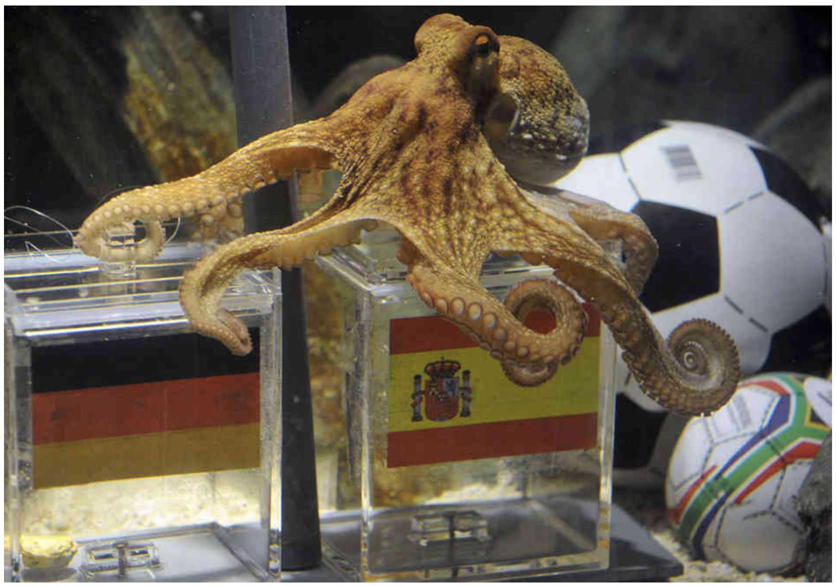

Opponent	Tournament	Stage	Date	Prediction	Result	Outcome
Poland	Euro 2008	group stage	8 June 2008	Germany	2–0	Correct
Croatia	Euro 2008	group stage	12 June 2008	Germany[2][21]	1–2	Incorrect
Austria	Euro 2008	group stage	16 June 2008	Germany	1–0	Correct
Portugal	Euro 2008	quarter-finals	19 June 2008	Germany	3–2	Correct
Turkey	Euro 2008	semi-finals	25 June 2008	Germany	3–2	Correct
Spain	Euro 2008	final	29 June 2008	Germany[2]	0–1	Incorrect
Australia	World Cup 2010	group stage	13 June 2010	Germany[30]	4–0	Correct
Serbia	World Cup 2010	group stage	18 June 2010	Serbia[30]	0–1	Correct
Ghana	World Cup 2010	group stage	23 June 2010	Germany[30]	1–0	Correct
England	World Cup 2010	round of 16	27 June 2010	Germany[31]	4–1	Correct
Argentina	World Cup 2010	quarter-finals	3 July 2010	Germany[24]	4–0	Correct
Spain	World Cup 2010	semi-finals	7 July 2010	Spain[32]	0–1	Correct
Uruguay	World Cup 2010	3rd place play-off	10 July 2010	Germany	3–2	Correct

Paul the Octopus (hatched in 2008, died October 2010) came to worldwide attention with his accurate predictions in the 2010 World Soccer Cup.

Available Information defines the (optimization) problem to solve

You don't want to ignore possibly relevant information

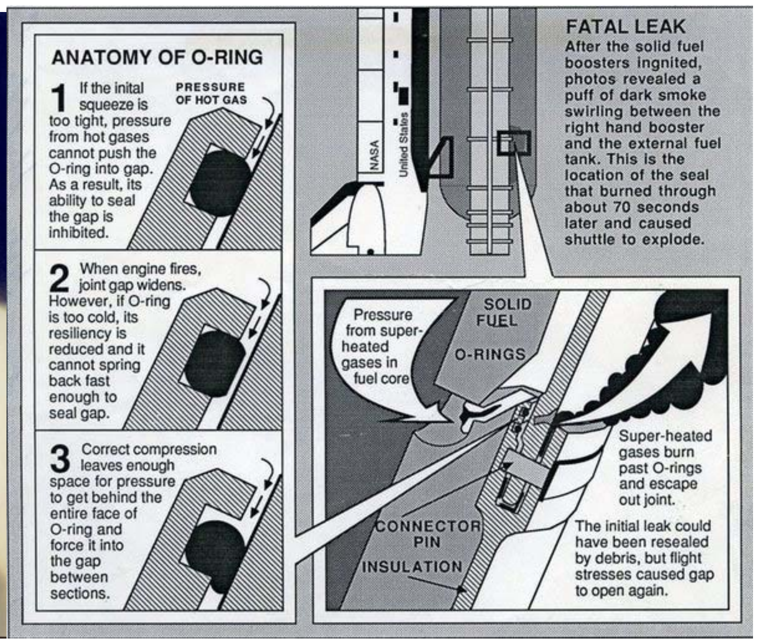

2003 Space Shuttle Columbia disaster

Why you don’t want to add possibly false assumptions

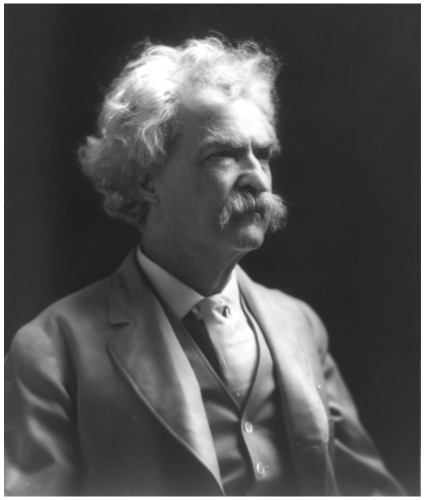

What gets us into trouble *is not what we don't know*.
It's what we know for sure that just ain't so.

--Mark Twain

In particular be careful about assumptions concerning the occurrence and impact of rare events.

The design of most of our nuclear power plants is based on the assumption of the availability of a steady supply of electricity to power the cooling system pumps for both the reactor cores as well as nearby “spent fuel ponds” where decommissioned reactor fuels rods are stored.

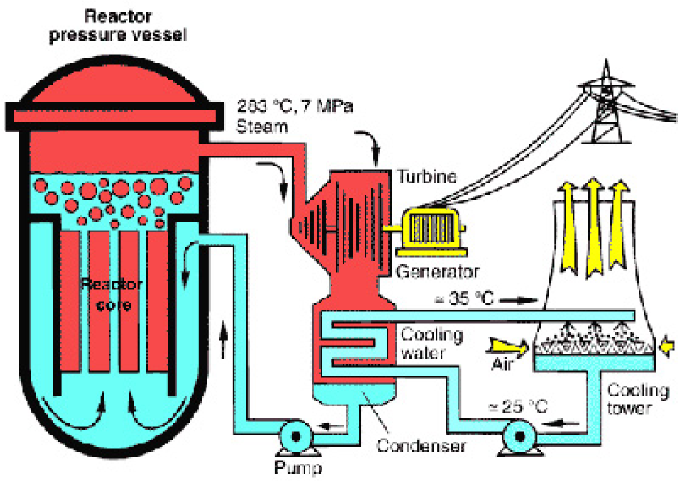

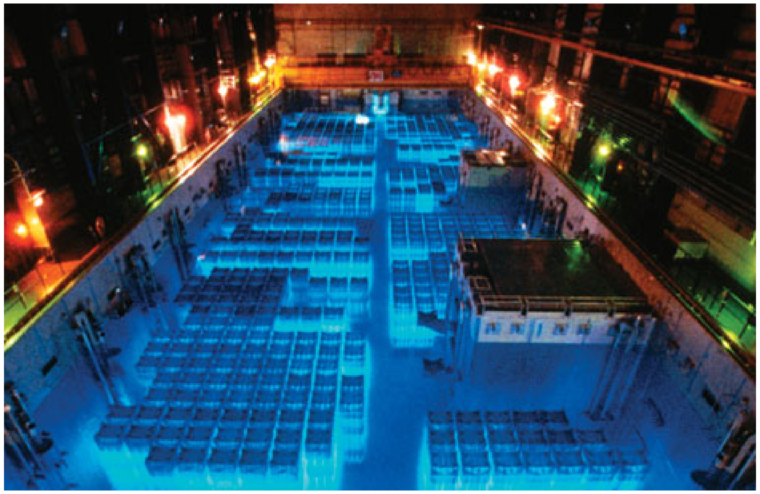

Our nuclear power plants are only required to store enough fuel on hand to keep the backup generators running for **one week**.

2008 NASA-funded study by the National Academy of Sciences ``Severe Space Weather Events—Understanding Societal and Economic Impacts.''

Carrington event (*solar super-storm of 1859)*

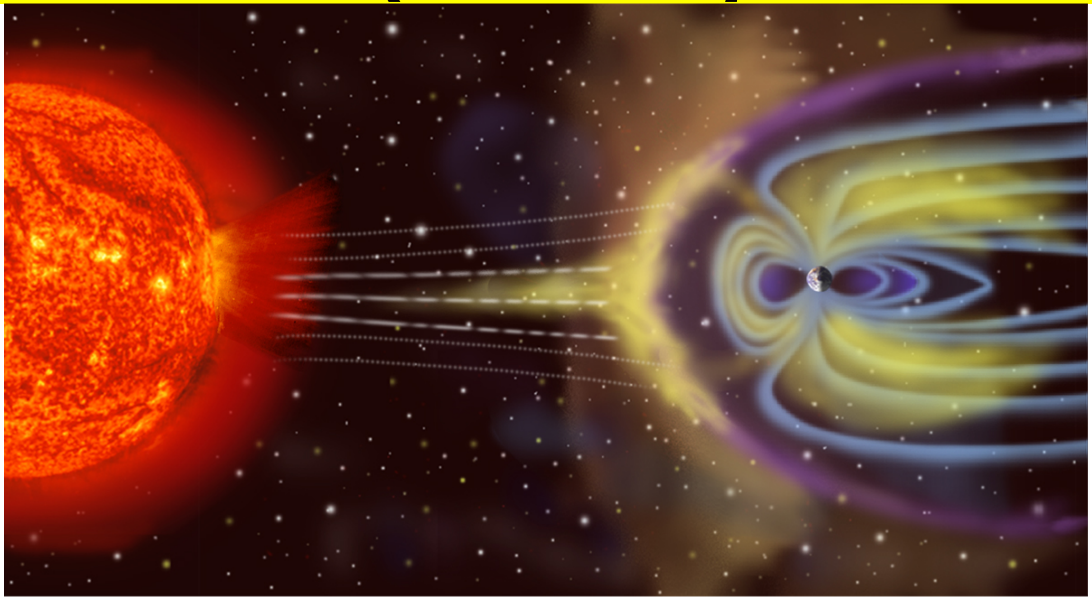

"Telegraph systems all over Europe and North America failed, in some cases giving telegraph operators electric shocks. Telegraph pylons threw sparks. Some telegraph systems continued to send and receive messages despite having been disconnected from their power supplies. Compasses and other sensitive instruments reeled as if struck by a massive magnetic fist."

2008 NASA-funded study by the National Academy of Sciences
``Severe Space Weather Events—
Understanding Societal and Economic Impacts.''

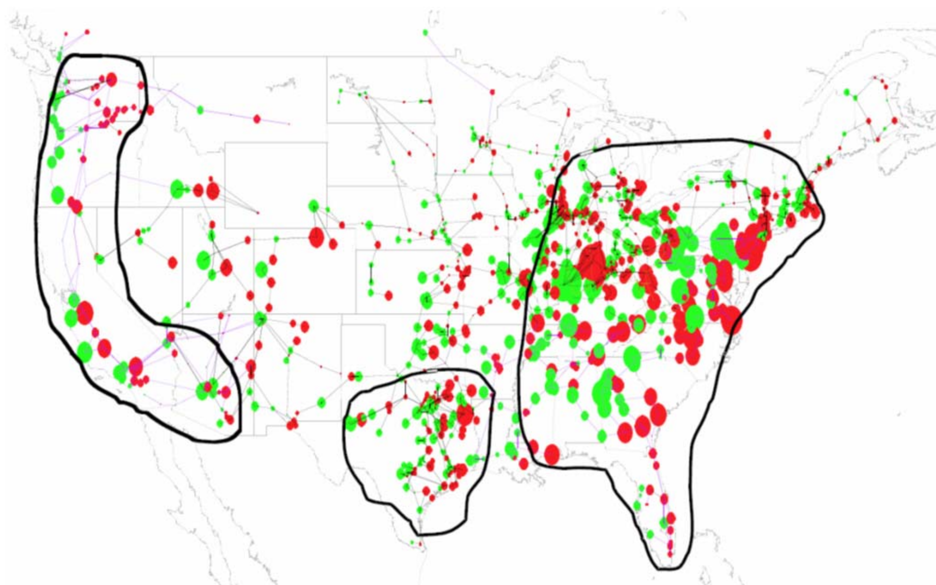

Regions susceptible to system collapse due to the effects of extreme geomagnetic disturbance.

Location nuclear power plants

THE HUFFINGTON POST

400 Chernobyls

Posted: 01/03/11 11:23 AM ET

MIT engineer warns of nuclear Armageddon, urges preventative measures

There are nearly 450 nuclear reactors in the world, with hundreds more either under construction or in the planning stages. Imagine what havoc it would wreak on our civilization, and the planet's ecosystems, if we were to suddenly experience not just one or two nuclear meltdowns, but 400. In this article, you will come to understand that unless we take significant preventative measures, this Apocalyptic scenario is not only possible, but probable.

Over the past 152 years the Earth has been struck by at least two naturally occurring severe geomagnetic solar storms of such a magnitude that if they were to occur today, in all likelihood would initiate a chain of events leading to catastrophic failures at most of our world's nuclear reactors. During the Great Geomagnetic Storm of May 14-15, 1921, brilliant aurora displays were reported in the Northern Hemisphere as far south as Mexico and Puerto Rico, and in the Southern Hemisphere as far north as Samoa. Just 62 years earlier, an even more powerful solar storm, referred to as "The Carrington Event," raged from August 28 to September 4, 1859.

The UQ challenge in the prediction context (Deepwater Horizon Disaster)

You want to find a 95% interval of confidence on the spill rate

6 different techniques lead to 6 distinct predictions with non-overlapping 95% confidence intervals

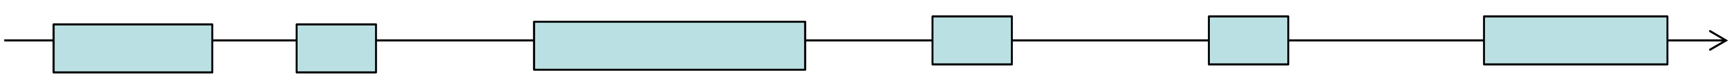

In developing this calculus we have

Uncovered extreme brittleness of Bayesian Inference.

Bayesian Brittleness. H. Owhadi, C. Scovel, T. Sullivan. 2013. arXiv:1304.6772

Discovered new Selberg Integral formulas.

Brittleness of Bayesian inference and new Selberg formulas. H. Owhadi, C. Scovel. 2013. arXiv:1304.7046

$$\mathbb{P}[A|B] = \mathbb{P}[B|A]\frac{\mathbb{P}[A]}{\mathbb{P}[B]}$$

$$\text{Posterior}[\theta|\text{data}] = \text{Likelihood}[\text{data}|\theta]\frac{\text{prior}[\theta]}{\text{prior}[\text{data}]}$$

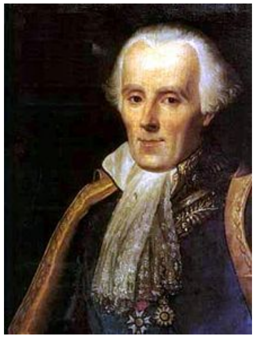

Reverend Thomas Bayes
1701-1761

Pierre Simon Laplace
1749-1827

Application of Bayes theorem in absence of genuine prior information has fueled a 250 years old debate with practical consequences in science, industry, medicine and law

When the prior **is** the data generating distribution

No controversy. Bayesian estimators are optimal.

When the prior **may not be** the data generating distribution

The controversy starts when Bayesian estimators are used without rigorous performance analysis.

A warm-up problem

You have a bag containing 100 coins

99 coins are fair

1 always land on head

You pick one coin at random from the bag

You flip it 10 times and 10 times you get head

What is the probability that the coin that you have picked is the unfair one?

Answer

$$\overbrace{}^{(1)}$$

$$\mathbb{P}[A|B] = \mathbb{P}[B|A]\frac{\mathbb{P}[A]}{\mathbb{P}[B]} = \overbrace{\frac{1}{1 + 99 \cdot (0.5)^{10}}}^{(1)} \approx 0.91$$

A: The coin is unfair

B: You observe 10 heads

Robustness If

bag contains 101 coins

and

fair coins are slightly unbalanced: probability of a head is 0.51

Then

(1) still a good approximation of correct answer

What if random outcomes are not head or tail but decimal numbers, perhaps given to finite precision?

Problem 2

We want to estimate $\Phi(\mu^\dagger) = \mu^\dagger[X \geq a]$

$\mu^\dagger$: Unknown or partially known measure of probability on $\mathbb{R}$

We observe $d = (d_1, \ldots, d_n) \in \mathbb{R}^n$

n i.i.d samples from $\mu^\dagger$

$$d \in B_\delta^n := \prod_{i=1}^n B_\delta(x_i)$$

$B_\delta(x)$: open ball of radius δ centered on x

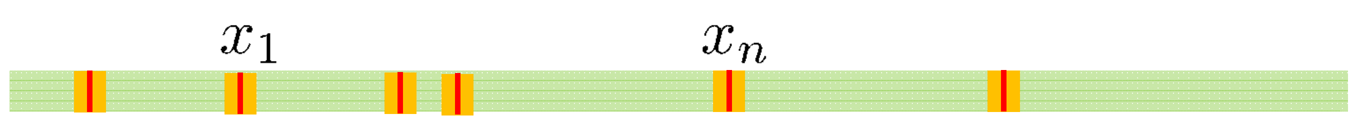

Bayesian Answer

Bayesian model class

$\{\mu(\theta) \mid \theta \in \Theta\}$

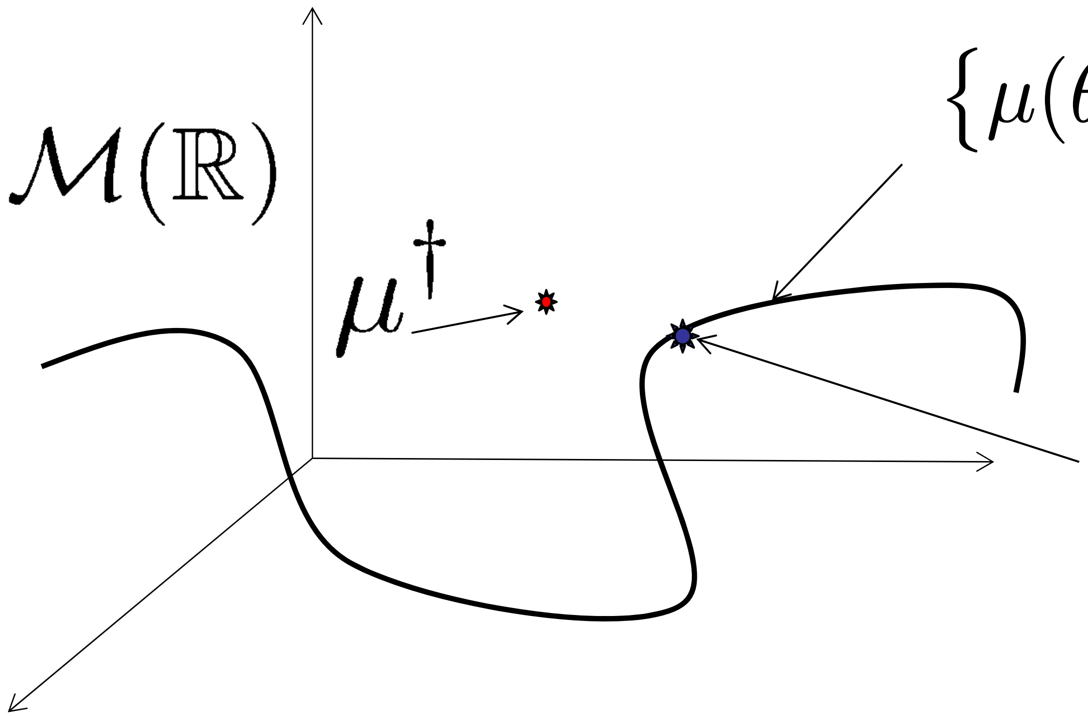

Assume

$\mu^\dagger$: Random element of

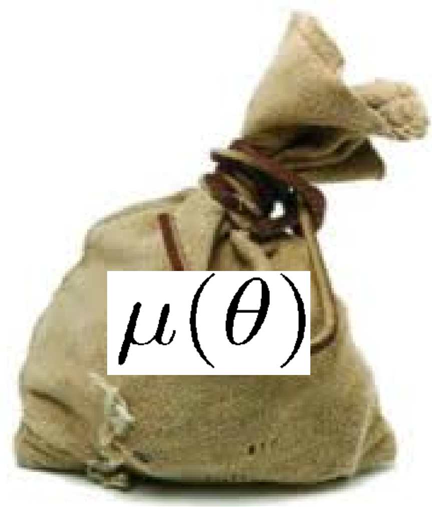

$$\mu^\dagger \in \{\mu(\theta) \mid \theta \in \Theta\}$$

Model is well specified

$$\mu^\dagger \notin \{\mu(\theta) \mid \theta \in \Theta\}$$

Model is misspecified

Questions

What happens to posterior values if our Bayesian model is a little bit wrong?

How sensitive is Bayesian Inference to local misspecification?

G. E. P. Box

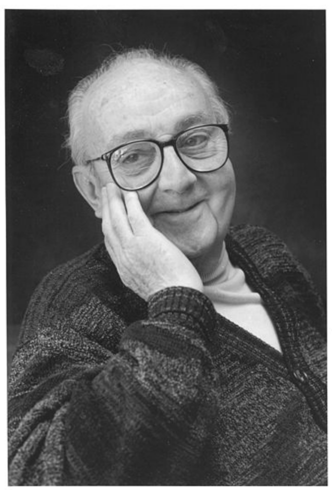

"Essentially, all models are wrong, but some are useful"

"Remember that all models are wrong; the practical question is how wrong do they have to be to not be useful?"

Answer

If you perturb the model (prior) just a little and if the resolution of your measurements is fine enough, then no matter the size of the data your posterior values can be anything you want

Figure. As measurement resolution $\delta \to 0$, the smooth dependence of the prior value on the prior (top-left) shatters into a patchwork of diametrically opposed posterior values.

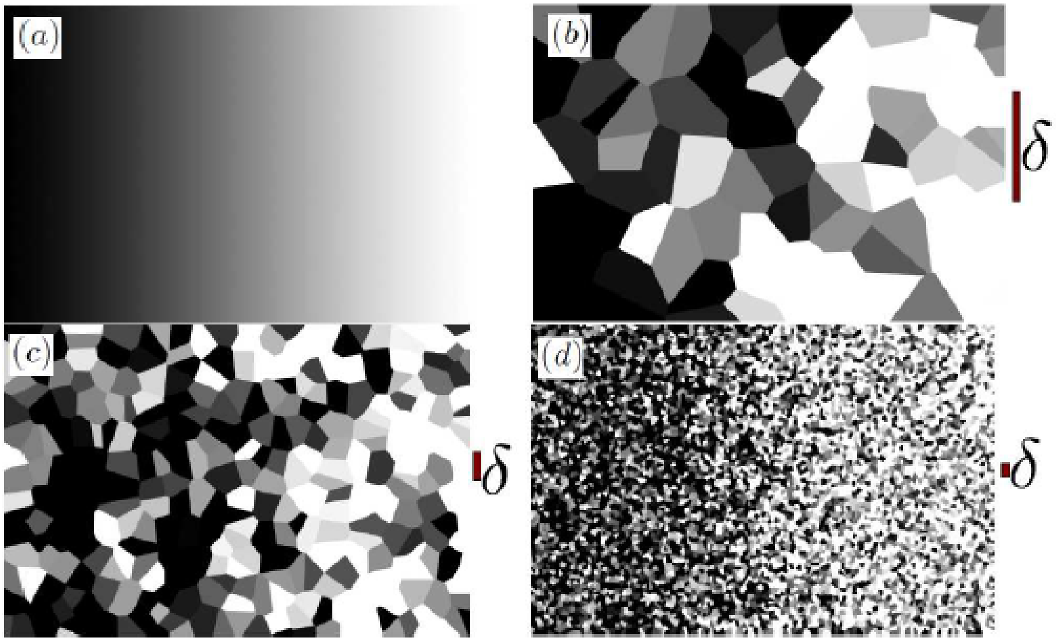

Are these results compatible with classical Robust Bayesian Inference?

Perform posterior Sensitivity Analysis over classes of priors

Box (1953) Huber (1964) Wasserman(1991)

Classical Robust Bayesian Inference:

What you do not know is finite → Robustness

Our brittleness results:

What you know is finite → Brittleness

Is Bayesian Inference Brittle?
Where do we go from here?

Robust Bayesian Inference as it currently stands leads to

Brittleness under finite information or local misspecification

Why? Robust Bayesian Inference as it currently stands is based on estimates posterior to the observation of data

Can we fix it?

Perhaps: compute robustness and accuracy estimates **prior** to the observation of the data

Need to compute optimal priors

Difficulty Need a new form of reduction calculus allowing us to solve optimization problems over spaces of measures over spaces of measures and functions

Papers:

- Optimal Uncertainty Quantification. H. Owhadi, Clint Scovel, T. Sullivan, M. McKerns and M. Ortiz. SIAM Review Vol. 55, No. 2 : pp. 271-345, 2013 (Expository Research Papers)
- Bayesian Brittleness.
  H. Owhadi, C. Scovel, T. Sullivan. 2013. arXiv:1304.6772
- Brittleness of Bayesian inference and new Selberg formulas.
  H. Owhadi, C. Scovel. 2013. arXiv:1304.7046

Grants:

- AFOSR. Grant number FA9550-12-1-0389. Scientific Computation of Optimal Statistical Estimators, 2012-2015.
- DOE/LANL. Exascale Co-Design Center for Materials in Extreme Environments. 2012-…
- DOE/NNSA. PSAAP, Uncertainty Quantification, ASC Predictive Science Academic Alliance Program, 2008-2013.

Collaborators: C. Scovel, T. Sullivan, M. McKerns, M. Ortiz, S. Han, R. Murray, M. Tao, U. Topcu, F. Theil, D. Meyer, A.A Kidane, A. Lashgari, B. Li, G. Ravichandran, M. Stalzer, M. Adams, J. Mihaly and A. J. Rosakis.