ENTROPY:
Gaining Knowledge by Admitting Ignorance

Alexander Schekochihin
Rudolf Peierls Centre for Theoretical Physics
University of Oxford

with thanks to James Binney, who convinced me that
it's better to gamble than to count

What is a System?

What is a System?

A system is something that can be in one of a set of (micro)states

$$\alpha = 1,\ 2,\ 3,\ \ldots,\ \Omega \ggg 1$$

What is a System?

A system is something that can be in one of a set of (micro)states

$$\alpha = 1,\ 2,\ 3,\ \ldots,\ \Omega \gg\gg\gg 1$$

There is a probability for it to be in each of these states:

$$p_1,\ p_2,\ p_3,\ \ldots,\ p_\alpha,\ \ldots,\ p_\Omega, \quad \sum_{\alpha=1}^{\Omega} p_\alpha = 1.$$

What is a System?

A system is something that can be in one of a set of (micro)states

$$\alpha = 1,\ 2,\ 3,\ \dots,\ \Omega \gg\gg\gg 1$$

There is a probability for it to be in each of these states:

$$p_1,\ p_2,\ p_3,\ \dots,\ p_\alpha,\ \dots,\ p_\Omega,\quad \sum_{\alpha=1}^{\Omega} p_\alpha = 1.$$

Each state has properties (quantum numbers): e.g., energies

$$E_1,\ E_2,\ E_3,\ \dots,\ E_\alpha,\ \dots,\ E_\Omega$$

(or momenta, angular momenta, spins, etc.)

What is a System?

A system is something that can be in one of a set of (micro)states

$$\alpha = 1,\ 2,\ 3,\ \ldots,\ \Omega \ggg\gg 1$$

There is a probability for it to be in each of these states:

$$p_1,\ p_2,\ p_3,\ \ldots,\ p_\alpha,\ \ldots,\ p_\Omega, \quad \sum_{\alpha=1}^{\Omega} p_\alpha = 1.$$

Each state has properties (quantum numbers): e.g., energies

$$E_1,\ E_2,\ E_3,\ \ldots,\ E_\alpha,\ \ldots,\ E_\Omega$$

(or momenta, angular momenta, spins, etc.)

If we know p_α, we can calculate averages $\leftrightarrow$ predict measurement outcomes

Mean energy: $U = \langle E_\alpha \rangle = \sum_\alpha p_\alpha E_\alpha$

What is a System?

A system is something that can be in one of a set of (micro)states

$$\alpha = 1,\ 2,\ 3,\ \ldots,\ \Omega \ggg\gg 1$$

There is a probability for it to be in each of these states:

$$p_1,\ p_2,\ p_3,\ \ldots,\ p_\alpha,\ \ldots,\ p_\Omega, \quad \sum_{\alpha=1}^{\Omega} p_\alpha = 1.$$

Each state has properties (quantum numbers): e.g., energies

$$E_1,\ E_2,\ E_3,\ \ldots,\ E_\alpha,\ \ldots,\ E_\Omega$$

(or momenta, angular momenta, spins, etc.)

If we know p_α, we can calculate averages ↔ predict measurement outcomes

Mean energy: $U = \langle E_\alpha \rangle = \sum_\alpha p_\alpha E_\alpha$

Pressure: $P = -\left\langle \frac{\partial E_\alpha}{\partial V} \right\rangle = -\sum_\alpha p_\alpha \frac{\partial E_\alpha}{\partial V}$

external parameter (volume)

What is a System?

A system is something that can be in one of a set of (micro)states

$$\alpha = 1,\ 2,\ 3,\ \ldots,\ \Omega \ggg\gg 1$$

There is a probability for it to be in each of these states:

$$p_1,\ p_2,\ p_3,\ \ldots,\ p_\alpha,\ \ldots,\ p_\Omega,\quad \sum_{\alpha=1}^{\Omega} p_\alpha = 1.$$

Each state has properties (quantum numbers): e.g., energies

$$E_1,\ E_2,\ E_3,\ \ldots,\ E_\alpha,\ \ldots,\ E_\Omega$$

(or momenta, angular momenta, spins, etc.)

If we know p_α, we can calculate averages ↔ predict measurement outcomes

Mean energy: $U = \langle E_\alpha \rangle = \sum_\alpha p_\alpha E_\alpha$

More generally, $\text{Force} = -\sum_\alpha p_\alpha \frac{\partial E_\alpha}{\partial\, \text{Displacement}}$

What is a System?

A system is something that can be in one of a set of (micro)states

$$\alpha = 1,\ 2,\ 3,\ \ldots,\ \Omega \ggg\gg 1$$

There is a probability for it to be in each of these states:

$$p_1,\ p_2,\ p_3,\ \ldots,\ p_\alpha,\ \ldots,\ p_\Omega, \quad \sum_{\alpha=1}^{\Omega} p_\alpha = 1.$$

Each state has properties (quantum numbers): e.g., energies

$$E_1,\ E_2,\ E_3,\ \ldots,\ E_\alpha,\ \ldots,\ E_\Omega$$

(or momenta, angular momenta, spins, etc.)

If we know p_α, we can calculate averages ↔ predict measurement outcomes

Mean energy: $U = \langle E_\alpha \rangle = \sum_\alpha p_\alpha E_\alpha$

E.g., tension: $f = \sum_\alpha p_\alpha \frac{\partial E_\alpha}{\partial L}$, magnetisation: $M = -\sum_\alpha p_\alpha \frac{\partial E_\alpha}{\partial B}$

What Do We Want to Know?

Mean energy: $U = \langle E_\alpha \rangle = \sum_\alpha p_\alpha E_\alpha$

external parameter (volume)

Pressure: $P = -\left\langle \frac{\partial E_\alpha}{\partial V} \right\rangle = -\sum_\alpha p_\alpha \frac{\partial E_\alpha}{\partial V}$

What Do We Want to Know?

What can be?

We know these from Quantum Mechanics
(or, more broadly, physics)

$$E_\alpha = E_\alpha(V, \text{parameters}\ldots)$$

Mean energy: $U = \langle E_\alpha \rangle = \sum_\alpha p_\alpha E_\alpha$

external parameter
(volume)

Pressure: $P = -\left\langle \frac{\partial E_\alpha}{\partial V} \right\rangle = -\sum_\alpha p_\alpha \frac{\partial E_\alpha}{\partial V}$

What Do We Want to Know?

How likely is it to happen?
This is what we want to determine.

What can be?
We know these from Quantum Mechanics
(or, more broadly, physics)

$$E_\alpha = E_\alpha(V, \text{parameters} \ldots)$$

Mean energy: $U = \langle E_\alpha \rangle = \sum_\alpha p_\alpha E_\alpha$

external parameter
(volume)

Pressure: $P = -\left\langle \frac{\partial E_\alpha}{\partial V} \right\rangle = -\sum_\alpha p_\alpha \frac{\partial E_\alpha}{\partial V}$

The Philosophy of Good-Enough-ism

How likely is it to happen?
This is what we want to determine.

We do not need to know the p_α's in any "true" sense.
We just need to know such p_α's that the few macroscopic quantities that we are ever likely to measure come out right.

Mean energy: $U = \langle E_\alpha \rangle = \sum_\alpha p_\alpha E_\alpha$

Pressure: $P = -\left\langle \frac{\partial E_\alpha}{\partial V} \right\rangle = -\sum_\alpha p_\alpha \frac{\partial E_\alpha}{\partial V}$

external parameter (volume)

Trainspotter's Statistics

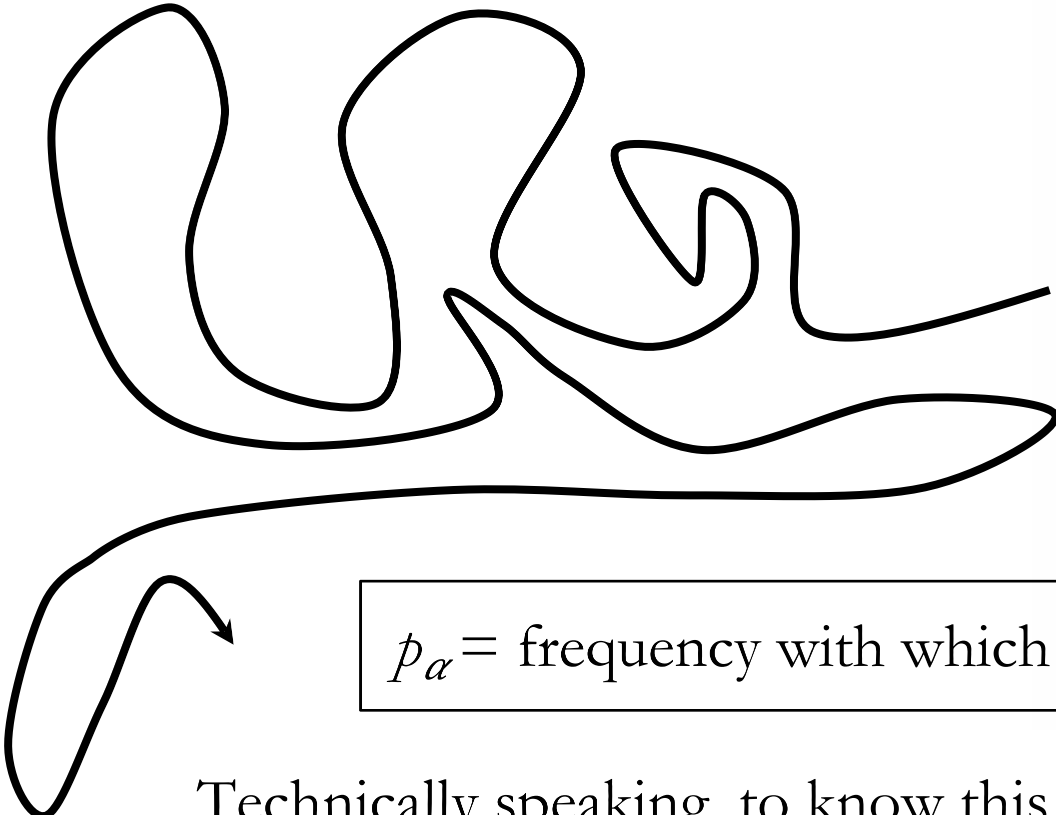

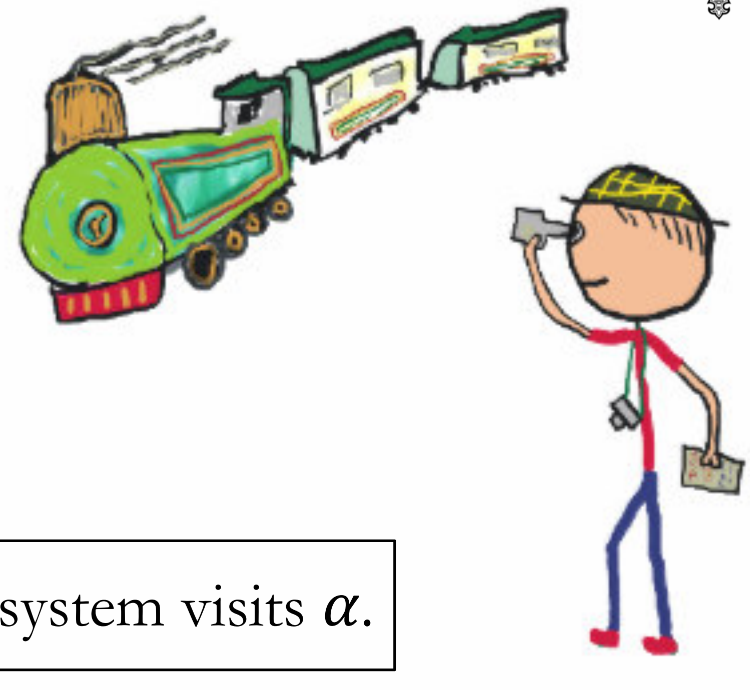

p_α = frequency with which system visits α.

Technically speaking, to know this, we need to understand (or assume something about) the system's dynamics.

Also, we must believe that the time necessary for the system to visit all of its phase space is a reasonable time for us to wait.

This is not, however, how forecasting works...

Gambler's Statistics

p_α = likelihood that the system is in α.

A fair and adequate assignment of likelihoods must be based on some information about the system, however incomplete.

Gambler's Statistics

p_α = likelihood that the system is in α.

A fair and adequate assignment of likelihoods must be based on some information about the system, however incomplete.

If we know nothing, the only adequate expectation is

$$p_\alpha = \frac{1}{\Omega}$$

(all states are equally probable)

Gambler's Statistics

Ludwig Boltzmann (1844-1906)

p_α = likelihood that the system is in α.

A fair and adequate assignment of likelihoods must be based on some information about the system, however incomplete.
If we know nothing, the only adequate expectation is

$$p_\alpha = \frac{1}{\Omega}$$

isolated system in equilibrium
(Boltzmann's *equal a priori probabilities postulate*)

Gambler's Statistics

p_α = likelihood that the system is in α.

A fair and adequate assignment of likelihoods must be based on some information about the system, however incomplete.

What if we know one thing, e.g., mean energy:

$$U = \langle E_\alpha \rangle = \sum_\alpha p_\alpha E_\alpha$$

This is one constraint on Ω>>>>1 likelihoods.

How do we work out p_α's using this and *only this* information?

A Fair Game

Principle: use knowledge available, admit ignorance of everything else.

A Fair Game

Principle: use knowledge available, admit ignorance of everything else.

1) Randomly assign $N >> \Omega$ "units of likelihood" to each of Ω states: N_α to state α. Then

$$p_\alpha = \frac{N_\alpha}{N}$$

A Fair Game

Principle: use knowledge available, admit ignorance of everything else.

1) Randomly assign $N >> \Omega$ "units of likelihood" to each of Ω states: N_α to state α. Then

$$p_\alpha = \frac{N_\alpha}{N}$$

2) Check that these p_α satisfy constraints imposed by knowledge. If not, keep trying until they do.

A Fair Game

Principle: use knowledge available, admit ignorance of everything else.

1) Randomly assign $N >> \Omega$ "units of likelihood" to each of Ω states: N_α to state α. Then

$$p_\alpha = \frac{N_\alpha}{N}$$

2) Check that these p_α satisfy constraints imposed by knowledge. If not, keep trying until they do.

The most likely outcome of this game is the one that can be obtained in the largest number of ways:

this maximum is very sharp at large N

$$W = \frac{N!}{N_1! \cdots N_\Omega!} \rightarrow \text{max (subject to constraints)}$$

(thus, we determine p_α's by maximising W in this way)

A Fair Game

Using Stirling's formula $\ln \mathcal{N}! = \mathcal{N}\ln\mathcal{N} - \mathcal{N} + O(\ln\mathcal{N})$, get:

$$\ln W = \underbrace{\mathcal{N}\ln\mathcal{N} - \mathcal{N}}_{\sum_\alpha(\mathcal{N}_\alpha \ln\mathcal{N} + \cancel{\mathcal{N}_\alpha})} + O(\ln\mathcal{N}) - \sum_\alpha \left[\mathcal{N}_\alpha \ln \mathcal{N}_\alpha - \cancel{\mathcal{N}_\alpha} + O(\ln\mathcal{N}_\alpha)\right]$$

$$= -\sum_\alpha \mathcal{N}_\alpha \ln\frac{\mathcal{N}_\alpha}{\mathcal{N}} + O(\ln\mathcal{N})$$

$$= -\mathcal{N}\left[\sum_\alpha p_\alpha \ln p_\alpha + O\left(\frac{\ln\mathcal{N}}{\mathcal{N}}\right)\right].$$

this maximum is very sharp at large N

$$W = \frac{\mathcal{N}!}{\mathcal{N}_1!\cdots\mathcal{N}_\Omega!} \to \text{max (subject to constraints)}$$

(thus, we determine p_α's by maximising W in this way)

A Fair Game

Using Stirling's formula $\ln \mathcal{N}! = \mathcal{N} \ln \mathcal{N} - \mathcal{N} + O(\ln \mathcal{N})$, get:

$$\ln W = \underbrace{\mathcal{N} \ln \mathcal{N} - \mathcal{N}}_{\sum_\alpha (\mathcal{N}_\alpha \ln \mathcal{N} + \cancel{\mathcal{N}_\alpha})} + O(\ln \mathcal{N}) - \sum_\alpha [\mathcal{N}_\alpha \ln \mathcal{N}_\alpha - \cancel{\mathcal{N}_\alpha} + O(\ln \mathcal{N}_\alpha)]$$

$$= -\sum_\alpha \mathcal{N}_\alpha \ln \frac{\mathcal{N}_\alpha}{\mathcal{N}} + O(\ln \mathcal{N})$$

$$= -\mathcal{N} \left[\sum_\alpha p_\alpha \ln p_\alpha + O\left(\frac{\ln \mathcal{N}}{\mathcal{N}} \right) \right].$$

I will call this Gibbs entropy

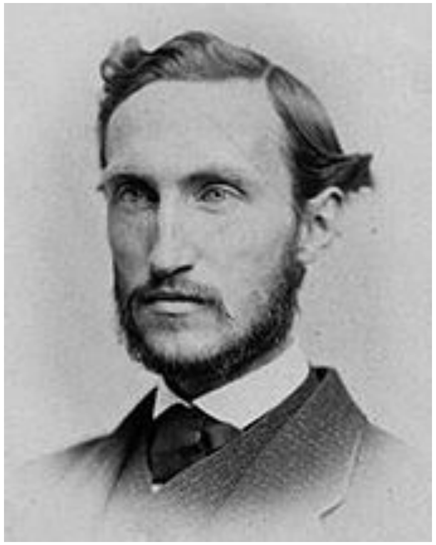

Josiah Willard Gibbs
(1839-1903)

$$S_G = -\sum_\alpha p_\alpha \ln p_\alpha \to \text{max (subject to constraints)}$$

(thus, we determine p_α's by maximising S_G in this way)

A Fair Game

The entropy of total ignorance is

$$p_\alpha = \frac{1}{\Omega} \quad \Rightarrow \quad S_G = -\sum_\alpha \frac{1}{\Omega} \ln \frac{1}{\Omega} = \ln \Omega$$

Boltzmann entropy

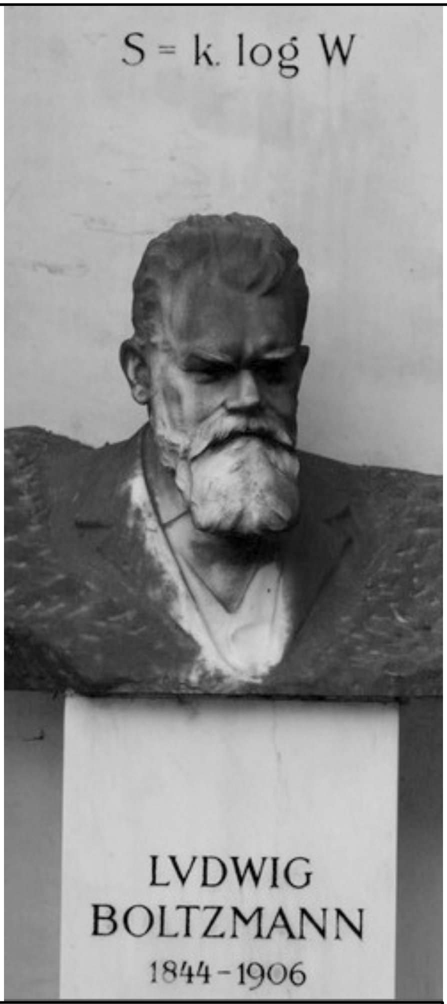

I will call this Gibbs entropy

$$S_G = -\sum_\alpha p_\alpha \ln p_\alpha \to \text{max (subject to constraints)}$$

(thus, we determine p_α's by maximising S_G in this way)

A Fair Game

The entropy of total ignorance is

$$p_\alpha = \frac{1}{\Omega} \quad \Rightarrow \quad S_\mathrm{G} = -\sum_\alpha \frac{1}{\Omega} \ln \frac{1}{\Omega} = \ln \Omega$$

Boltzmann entropy

It corresponds to the number of all possible distributions of N units of probability between Ω states: $W_\mathrm{max} = \Omega^N$, so

$$S_\mathrm{G,max} = \frac{1}{N} \ln W_\mathrm{max} = \ln \Omega$$

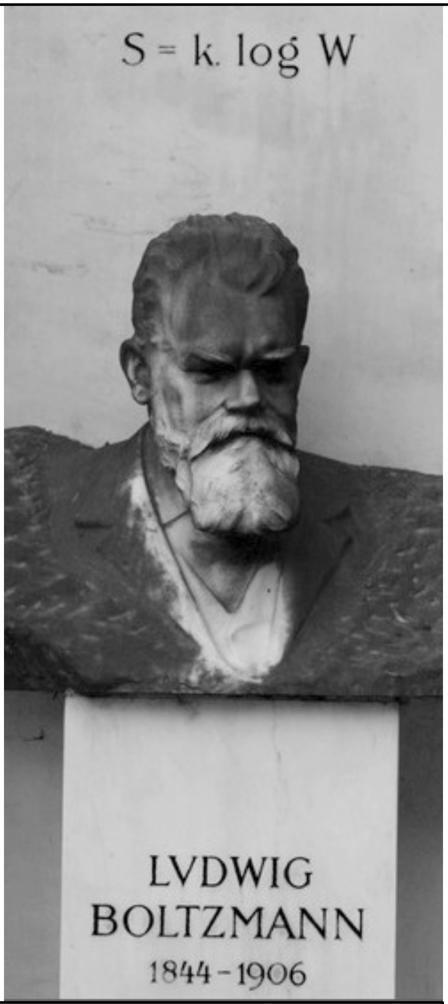

I will call this Gibbs entropy

$$S_\mathrm{G} = -\sum_\alpha p_\alpha \ln p_\alpha \rightarrow \text{max (subject to constraints)}$$

(thus, we determine p_α's by maximising S_G in this way)

Shannon's Theorem

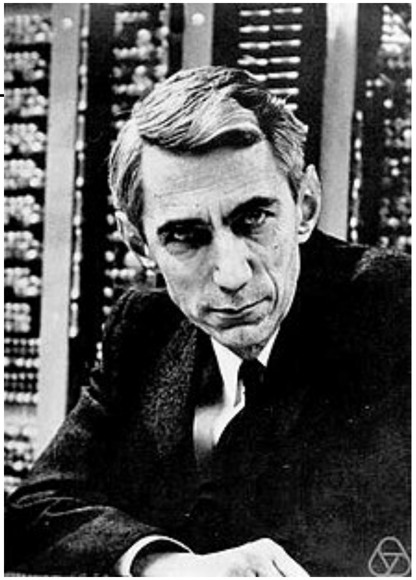

Claude Shannon
(1916-2001)

Was that the only possible fair game?

Shannon declared that a good measure of uncertainty of set of probabilities p_α, called $H(p_1, \ldots, p_\Omega)$, must be

1) a continuous function of p_α's

Shannon's Theorem

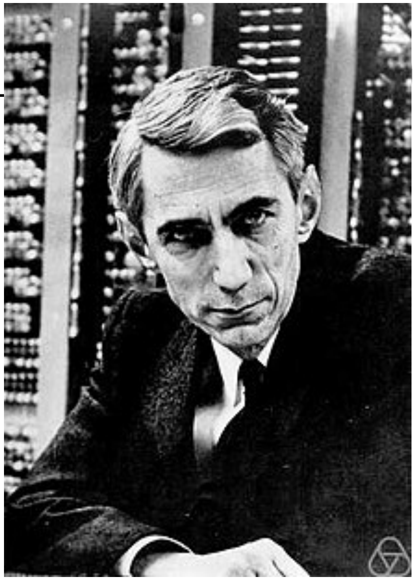
Claude Shannon
(1916-2001)

Was that the only possible fair game?

Shannon declared that a good measure of uncertainty of set of probabilities p_α, called $H(p_1, \ldots, p_\Omega)$, must be

1) a continuous function of p_α's,
2) a symmetric function of p_α's (list order irrelevant)

Shannon's Theorem

Was that the only possible fair game?

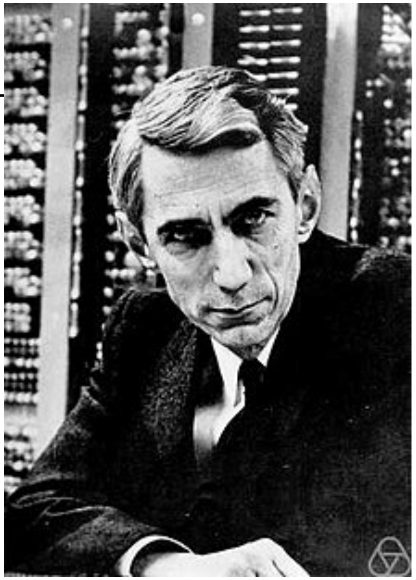

Claude Shannon
(1916-2001)

Shannon declared that a good measure of uncertainty of set of probabilities p_α, called $H(p_1, \ldots, p_\Omega)$, must be

1) a continuous function of p_α's,
2) a symmetric function of p_α's (list order irrelevant),
3) equal probabilities represent maximum uncertainty:

$$H(p_1, \ldots, p_\Omega) < H\left(\frac{1}{\Omega}, \ldots, \frac{1}{\Omega}\right) \equiv H_\Omega$$

Shannon's Theorem

Was that the only possible fair game?

Shannon declared that a good measure of uncertainty of set of probabilities p_α, called $H(p_1, \ldots, p_\Omega)$, must be

1) a continuous function of p_α's,
2) a symmetric function of p_α's (list order irrelevant),
3) equal probabilities represent maximum uncertainty:

$$H(p_1, \ldots, p_\Omega) < H\left(\frac{1}{\Omega}, \ldots, \frac{1}{\Omega}\right) \equiv H_\Omega$$

4) more equiprobable states → more uncertainty: if $\Omega' > \Omega$, $H_{\Omega'} > H_\Omega$

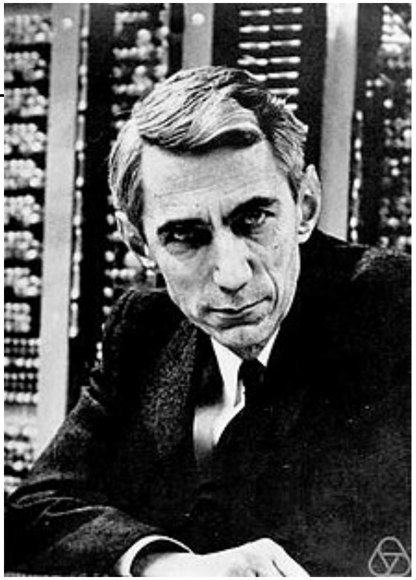

Claude Shannon
(1916-2001)

Shannon's Theorem

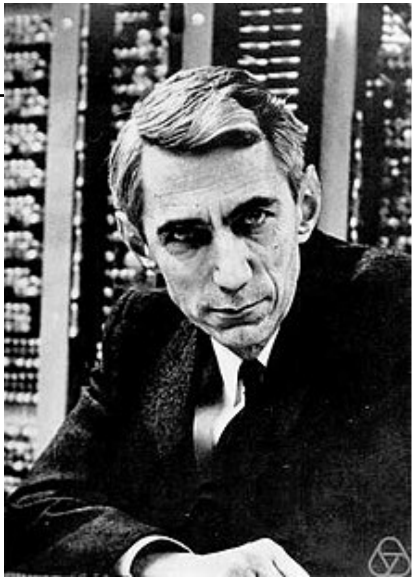

Claude Shannon
(1916-2001)

Was that the only possible fair game?

Shannon declared that a good measure of uncertainty of set of probabilities p_α, called $H(p_1, \ldots, p_\Omega)$, must be

1) a continuous function of p_α's,
2) a symmetric function of p_α's (list order irrelevant),
3) equal probabilities represent maximum uncertainty:

$$H(p_1, \ldots, p_\Omega) < H\left(\frac{1}{\Omega}, \ldots, \frac{1}{\Omega}\right) \equiv H_\Omega$$

4) more equiprobable states → more uncertainty: if $\Omega' > \Omega$, $H_{\Omega'} > H_\Omega$,
5) additive and independent of counting scheme: if we split all states into subgroups with probabilities w_i, and calculate H_i for each subgroup (with conditional probabilities p_α/w_i), then

$$H(p_1, \ldots, p_\Omega) = H(w_1, w_2, \ldots) + \sum_i w_i H_i$$

total uncertainty — uncertainty in choosing subgroup — uncertainty within subgroup

Shannon's Theorem

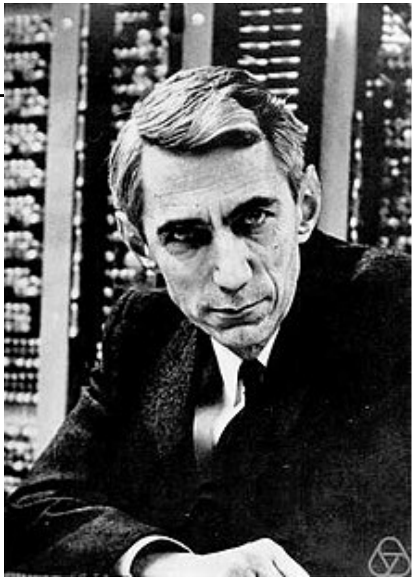

Claude Shannon
(1916-2001)

Was that the only possible fair game?

Shannon declared that a good measure of uncertainty of set of probabilities p_α, called $H(p_1, \ldots, p_\Omega)$, must be

1) a continuous function of p_α's,
2) a symmetric function of p_α's (list order irrelevant),
3) equal probabilities represent maximum uncertainty:

$$H(p_1, \ldots, p_\Omega) < H\left(\frac{1}{\Omega}, \ldots, \frac{1}{\Omega}\right) \equiv H_\Omega$$

4) more equiprobable states → more uncertainty: if $\Omega' > \Omega$, $H_{\Omega'} > H_\Omega$,
5) additive and independent of counting scheme: if we split all states into subgroups with probabilities w_i, and calculate H_i for each subgroup (with conditional probabilities p_α / w_i), then

$$H(p_1, \ldots, p_\Omega) = H(w_1, w_2, \ldots) + \sum_i w_i H_i$$

total uncertainty — uncertainty in choosing subgroup — uncertainty within subgroup

He then proved that the only function with these properties is

$$H(p_1, \ldots, p_\Omega) = -k \sum_\alpha p_\alpha \ln p_\alpha = k S_G$$

How to Maximise Entropy (with strings attached)

$$\Lambda = S_{\mathrm{G}}(p_1, \ldots, p_\Omega) - \sum_i \lambda_i F_i(p_1, \ldots, p_\Omega) \to \max$$

Entropy

$$-\sum_\alpha p_\alpha \ln p_\alpha$$

"Lagrange multipliers"

constraints in the form

$$F_i(p_1, \ldots, p_\Omega) = 0$$

How to Maximise Entropy (with strings attached)

$$\Lambda = S_G(p_1, \ldots, p_\Omega) - \sum_i \lambda_i F_i(p_1, \ldots, p_\Omega) \to \max$$

Entropy

$$-\sum_\alpha p_\alpha \ln p_\alpha$$

"Lagrange multipliers"

constraints in the form

$$F_i(p_1, \ldots, p_\Omega) = 0$$

Always: $\sum_\alpha p_\alpha - 1 = 0$

Mean energy known: $\sum_\alpha p_\alpha E_\alpha - U = 0$

...etc.

Total Ignorance (Isolated System)

$$\Lambda = -\sum_\alpha p_\alpha \ln p_\alpha - \lambda\left(\sum_\alpha p_\alpha - 1\right) \to \max$$

Always: $$\sum_\alpha p_\alpha - 1 = 0$$

Mean energy known: $$\sum_\alpha p_\alpha E_\alpha - U = 0$$

...etc.

E.g., in the case of **total ignorance**, the only constraint is

$$\frac{\partial \Lambda}{\partial p_\alpha} = -\ln p_\alpha - 1 - \lambda = 0 \quad \Rightarrow \quad p_\alpha = e^{-(1+\lambda)}$$ all equal

Total Ignorance (Isolated System)

$$\Lambda = -\sum_\alpha p_\alpha \ln p_\alpha - \lambda \left(\sum_\alpha p_\alpha - 1 \right) \to \max$$

Always: $\sum_\alpha p_\alpha - 1 = 0$

Mean energy known: $\sum_\alpha p_\alpha E_\alpha - U = 0$

...etc.

E.g., in the case of **total ignorance**, the only constraint is

$$\frac{\partial \Lambda}{\partial p_\alpha} = -\ln p_\alpha - 1 - \lambda = 0 \quad \Rightarrow \quad p_\alpha = e^{-(1+\lambda)}$$ all equal

$$\frac{\partial \Lambda}{\partial \lambda} = \sum_\alpha p_\alpha - 1 = 0 \quad \Rightarrow \quad p_\alpha = \frac{1}{\Omega}$$ **method works!**

Gibbs' Ensemble

$$\Lambda = -\sum_\alpha p_\alpha \ln p_\alpha - \lambda\left(\sum_\alpha p_\alpha - 1\right) - \beta\left(\sum_\alpha p_\alpha E_\alpha - U\right) \rightarrow \max$$

Always: $$\sum_\alpha p_\alpha - 1 = 0$$

Mean energy known: $$\sum_\alpha p_\alpha E_\alpha - U = 0$$

...etc.

If we add an extra constraint,

$$\frac{\partial \Lambda}{\partial p_\alpha} = -\ln p_\alpha - 1 - \lambda - \beta E_\alpha = 0 \quad \Rightarrow \quad p_\alpha = e^{-(1+\lambda)-\beta E_\alpha}$$

Gibbs' Ensemble

$$\Lambda = -\sum_\alpha p_\alpha \ln p_\alpha - \lambda\left(\sum_\alpha p_\alpha - 1\right) - \beta\left(\sum_\alpha p_\alpha E_\alpha - U\right) \to \max$$

Always: $\sum_\alpha p_\alpha - 1 = 0$

Mean energy known: $\sum_\alpha p_\alpha E_\alpha - U = 0$

...etc.

If we add an extra constraint,

$$\frac{\partial \Lambda}{\partial p_\alpha} = -\ln p_\alpha - 1 - \lambda - \beta E_\alpha = 0 \quad \Rightarrow \quad p_\alpha = e^{-(1+\lambda)-\beta E_\alpha}$$

$$\frac{\partial \Lambda}{\partial \lambda} = \sum_\alpha p_\alpha - 1 = 0 \quad \Rightarrow \quad e^{-(1+\lambda)}\sum_\alpha e^{-\beta E_\alpha} = 1$$

Gibbs' Ensemble

$$\Lambda = -\sum_\alpha p_\alpha \ln p_\alpha - \lambda\left(\sum_\alpha p_\alpha - 1\right) - \beta\left(\sum_\alpha p_\alpha E_\alpha - U\right) \to \max$$

Always: $\sum_\alpha p_\alpha - 1 = 0$

Mean energy known: $\sum_\alpha p_\alpha E_\alpha - U = 0$

...etc.

If we add an extra constraint,

$$\frac{\partial \Lambda}{\partial p_\alpha} = -\ln p_\alpha - 1 - \lambda - \beta E_\alpha = 0 \quad \Rightarrow \quad p_\alpha = e^{-(1+\lambda)-\beta E_\alpha}$$

$$\frac{\partial \Lambda}{\partial \lambda} = \sum_\alpha p_\alpha - 1 = 0 \quad \Rightarrow \quad e^{-(1+\lambda)} \sum_\alpha e^{-\beta E_\alpha} = 1$$

$\sum_\alpha e^{-\beta E_\alpha} = Z(\beta)$ "partition function"

Gibbs' Ensemble

$$\Lambda = -\sum_\alpha p_\alpha \ln p_\alpha - \lambda\left(\sum_\alpha p_\alpha - 1\right) - \beta\left(\sum_\alpha p_\alpha E_\alpha - U\right) \to \max$$

Always: $\sum_\alpha p_\alpha - 1 = 0$

Mean energy known: $\sum_\alpha p_\alpha E_\alpha - U = 0$

...etc.

If we add an extra constraint,

$$\frac{\partial \Lambda}{\partial p_\alpha} = -\ln p_\alpha - 1 - \lambda - \beta E_\alpha = 0 \quad \Rightarrow \quad p_\alpha = e^{-(1+\lambda)-\beta E_\alpha} = \frac{1}{Z(\beta)} e^{-\beta E_\alpha}$$

$$\frac{\partial \Lambda}{\partial \lambda} = \sum_\alpha p_\alpha - 1 = 0 \quad \Rightarrow \quad e^{-(1+\lambda)} \sum_\alpha e^{-\beta E_\alpha} = 1$$

$\sum_\alpha e^{-\beta E_\alpha} = Z(\beta)$ "partition function"

Gibbs' Ensemble

$$\Lambda = -\sum_\alpha p_\alpha \ln p_\alpha - \lambda\left(\sum_\alpha p_\alpha - 1\right) - \beta\left(\sum_\alpha p_\alpha E_\alpha - U\right) \to \max$$

Always: $\sum_\alpha p_\alpha - 1 = 0$

Mean energy known: $\sum_\alpha p_\alpha E_\alpha - U = 0$

...etc.

If we add an extra constraint,

$$\frac{\partial \Lambda}{\partial p_\alpha} = -\ln p_\alpha - 1 - \lambda - \beta E_\alpha = 0 \quad \Rightarrow \quad p_\alpha = e^{-(1+\lambda)-\beta E_\alpha} = \frac{1}{Z(\beta)} e^{-\beta E_\alpha}$$

$$\frac{\partial \Lambda}{\partial \lambda} = \sum_\alpha p_\alpha - 1 = 0 \quad \Rightarrow \quad e^{-(1+\lambda)} \sum_\alpha e^{-\beta E_\alpha} = 1$$

$\sum_\alpha e^{-\beta E_\alpha} = Z(\beta)$ "partition function"

$$\frac{\partial \Lambda}{\partial \beta} = \sum_\alpha p_\alpha E_\alpha - U = 0 \quad \Rightarrow \quad U = \sum_\alpha \frac{E_\alpha e^{-\beta E_\alpha}}{Z(\beta)} = -\frac{1}{Z}\frac{\partial Z}{\partial \beta}$$

Gibbs' Ensemble

$$p_\alpha = \frac{e^{-\beta E_\alpha}}{Z(\beta)}$$ likelihoods based of knowledge of U

$$Z(\beta) = \sum_\alpha e^{-\beta E_\alpha}$$ normalisation constant

$$U = -\frac{1}{Z}\frac{\partial Z}{\partial \beta}$$ implicit equation for mysterious β

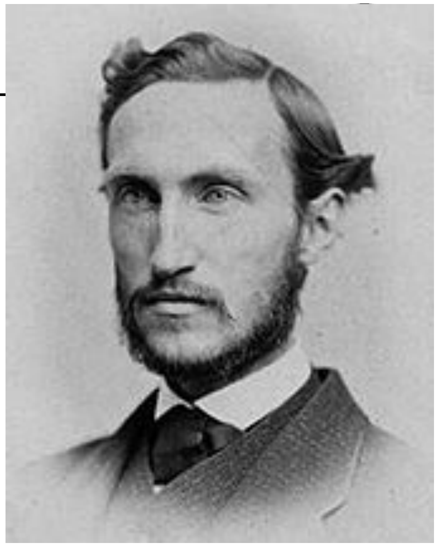

Josiah Willard Gibbs
(1839-1903)

It looks like we have calculated something, but what does it mean?

Gibbs' Ensemble

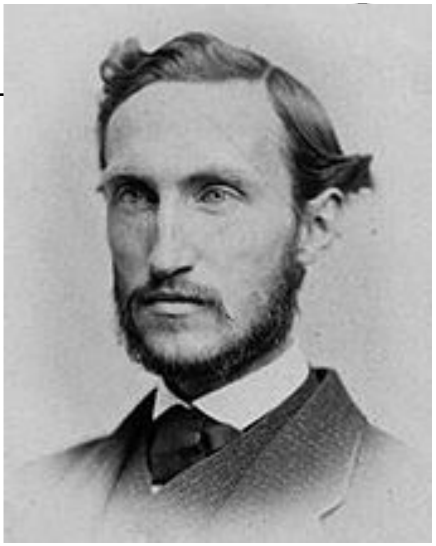

Josiah Willard Gibbs
(1839-1903)

$$\boxed{p_\alpha = \frac{e^{-\beta E_\alpha}}{Z(\beta)}}$$ likelihoods based of knowledge of U

$$Z(\beta) = \sum_\alpha e^{-\beta E_\alpha}$$ normalisation constant

$$U = -\frac{1}{Z}\frac{\partial Z}{\partial \beta}$$ implicit equation for mysterious β

It looks like we have calculated something, but what does it mean?

$$S_{\rm G} = -\sum_\alpha p_\alpha \ln p_\alpha = -\sum_\alpha p_\alpha(-\beta E_\alpha - \ln Z) = \beta U + \ln Z.$$

$$\mathrm{d}S_{\rm G} = \beta \mathrm{d}U + U\mathrm{d}\beta + \frac{\mathrm{d}Z}{Z}$$

"d" here means we are varying some parameter of the equilibrium, say V

Gibbs' Ensemble

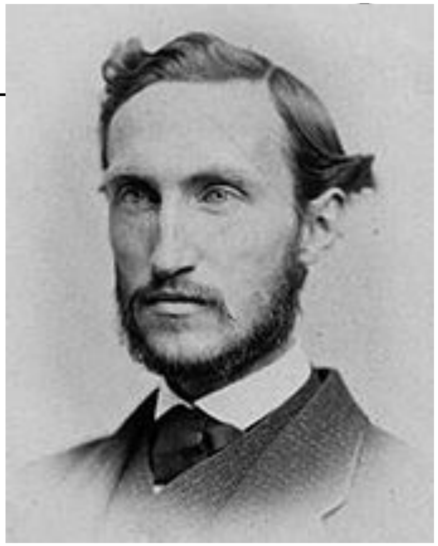

Josiah Willard Gibbs
(1839-1903)

$$\boxed{p_\alpha = \frac{e^{-\beta E_\alpha}}{Z(\beta)}}$$ likelihoods based of knowledge of U

$$Z(\beta) = \sum_\alpha e^{-\beta E_\alpha}$$ normalisation constant

$$U = -\frac{1}{Z}\frac{\partial Z}{\partial \beta}$$ implicit equation for mysterious β

It looks like we have calculated something, but what does it mean?

$$S_G = -\sum_\alpha p_\alpha \ln p_\alpha = -\sum_\alpha p_\alpha(-\beta E_\alpha - \ln Z) = \beta U + \ln Z.$$

$$\mathrm{d}S_G = \beta \mathrm{d}U + U\mathrm{d}\beta + \frac{\mathrm{d}Z}{Z}$$

"d" here means we are varying some parameter of the equilibrium, say V

$$= \beta \mathrm{d}U + \cancel{U\mathrm{d}\beta} + \sum_\alpha \underbrace{\frac{e^{-\beta E_\alpha}}{Z}}_{=p_\alpha}(-\beta \mathrm{d}E_\alpha - \cancel{E_\alpha \mathrm{d}\beta})$$

Gibbs' Ensemble

$$p_\alpha = \frac{e^{-\beta E_\alpha}}{Z(\beta)}$$ likelihoods based of knowledge of U

$$Z(\beta) = \sum_\alpha e^{-\beta E_\alpha}$$ normalisation constant

$$U = -\frac{1}{Z}\frac{\partial Z}{\partial \beta}$$ implicit equation for mysterious β

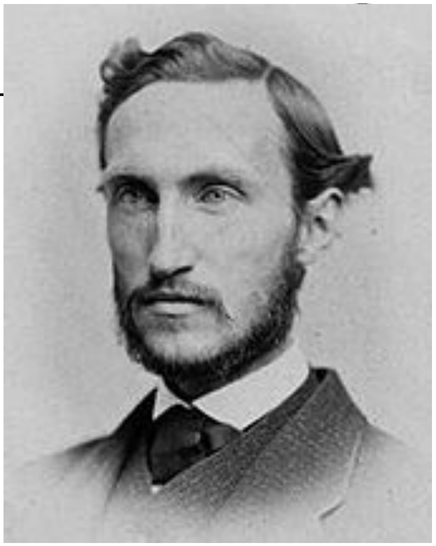

Josiah Willard Gibbs
(1839-1903)

It looks like we have calculated something, but what does it mean?

$$S_G = -\sum_\alpha p_\alpha \ln p_\alpha = -\sum_\alpha p_\alpha(-\beta E_\alpha - \ln Z) = \beta U + \ln Z.$$

$$dS_G = \beta dU + U d\beta + \frac{dZ}{Z}$$

$$dE_\alpha = \frac{\partial E_\alpha}{\partial V} dV$$

$$= \beta dU + \cancel{U d\beta} + \sum_\alpha \underbrace{\frac{e^{-\beta E_\alpha}}{Z}}_{= p_\alpha}(-\beta dE_\alpha - \cancel{E_\alpha d\beta})$$

Gibbs' Ensemble

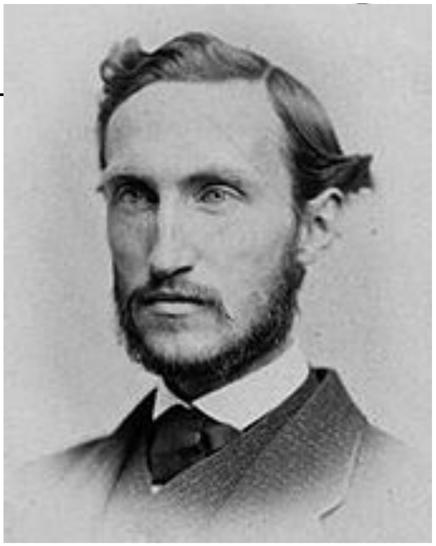

Josiah Willard Gibbs (1839-1903)

$$\boxed{p_\alpha = \frac{e^{-\beta E_\alpha}}{Z(\beta)}}$$ likelihoods based of knowledge of U

$$Z(\beta) = \sum_\alpha e^{-\beta E_\alpha}$$ normalisation constant

$$U = -\frac{1}{Z}\frac{\partial Z}{\partial \beta}$$ implicit equation for mysterious β

It looks like we have calculated something, but what does it mean?

$$S_G = -\sum_\alpha p_\alpha \ln p_\alpha = -\sum_\alpha p_\alpha(-\beta E_\alpha - \ln Z) = \beta U + \ln Z.$$

$$\mathrm{d}S_G = \beta \mathrm{d}U + U\mathrm{d}\beta + \frac{\mathrm{d}Z}{Z}$$

$$= \beta \mathrm{d}U + \cancel{U\mathrm{d}\beta} + \sum_\alpha \underbrace{\frac{e^{-\beta E_\alpha}}{Z}}_{=p_\alpha}(-\beta \mathrm{d}E_\alpha - \cancel{E_\alpha \mathrm{d}\beta})$$

$$= \beta(\mathrm{d}U + P\mathrm{d}V)$$

$$\mathrm{d}E_\alpha = \frac{\partial E_\alpha}{\partial V}\mathrm{d}V$$

pressure

$$P = -\sum_\alpha p_\alpha \frac{\partial E_\alpha}{\partial V}$$

Gibbs' Ensemble

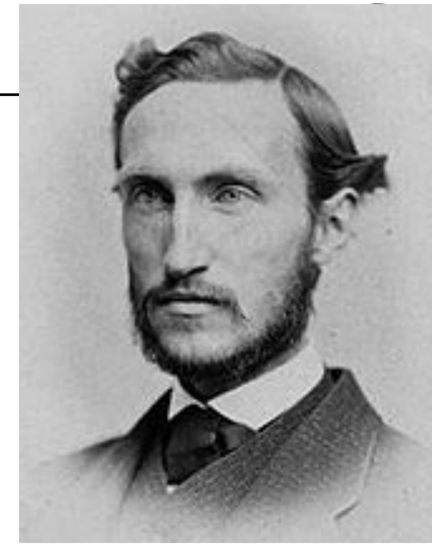

Josiah Willard Gibbs (1839-1903)

$$p_\alpha = \frac{e^{-\beta E_\alpha}}{Z(\beta)}$$ likelihoods based of knowledge of U

$$Z(\beta) = \sum_\alpha e^{-\beta E_\alpha}$$ normalisation constant

$$U = -\frac{1}{Z}\frac{\partial Z}{\partial \beta}$$ implicit equation for mysterious β

It looks like we have calculated something, but what does it mean?

$$S_G = -\sum_\alpha p_\alpha \ln p_\alpha = -\sum_\alpha p_\alpha(-\beta E_\alpha - \ln Z) = \beta U + \ln Z.$$

$$\begin{aligned} dS_G &= \beta dU + U d\beta + \frac{dZ}{Z} \\ &= \beta dU + \cancel{U d\beta} + \sum_\alpha \underbrace{\frac{e^{-\beta E_\alpha}}{Z}}_{= p_\alpha}(-\beta dE_\alpha - \cancel{E_\alpha d\beta}) \\ &= \beta(dU + PdV) = \beta(dU - dW) = \beta\, dQ_{rev} \end{aligned}$$

heat in a reversible process = energy - work

Gibbs' Ensemble

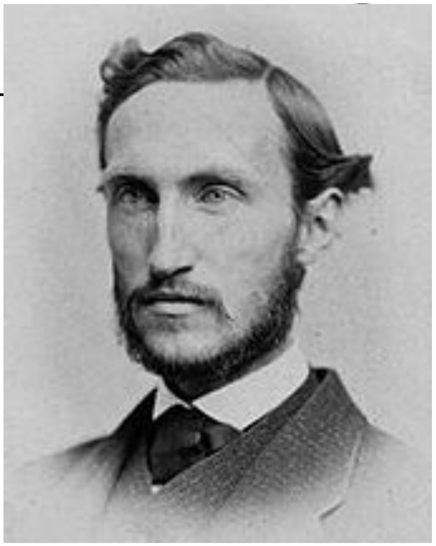

Josiah Willard Gibbs
(1839-1903)

$$\boxed{p_\alpha = \frac{e^{-\beta E_\alpha}}{Z(\beta)}}$$ likelihoods based of knowledge of U

$$Z(\beta) = \sum_\alpha e^{-\beta E_\alpha}$$ normalisation constant

$$U = -\frac{1}{Z}\frac{\partial Z}{\partial \beta}$$ implicit equation for mysterious β

It looks like we have calculated something, but what does it mean?

$$\mathrm{d}S_\mathrm{G} = \beta\, \mathrm{d}Q_\mathrm{rev}$$

Thus, change of entropy is proportional to reversible heat, and the const of proportionality is β.

Entropy Is "Real"!

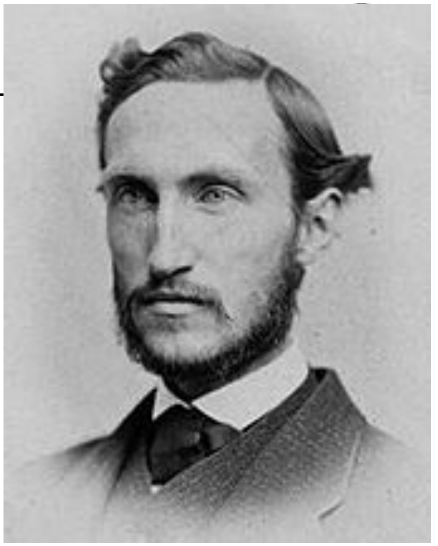

Josiah Willard Gibbs
(1839-1903)

$$p_\alpha = \frac{e^{-\beta E_\alpha}}{Z(\beta)}$$ likelihoods based of knowledge of U

$$Z(\beta) = \sum_\alpha e^{-\beta E_\alpha}$$ normalisation constant

$$U = -\frac{1}{Z}\frac{\partial Z}{\partial \beta}$$ implicit equation for mysterious β

It looks like we have calculated something, but what does it mean?

$$\mathrm{d}S_\mathrm{G} = \beta\, \mathrm{d}Q_\mathrm{rev}$$

Thus, change of entropy is proportional to reversible heat, and the const of proportionality is β. Therefore,

$$\mathrm{d}S = \frac{\mathrm{d}Q_\mathrm{rev}}{T}$$

$$\beta = \frac{1}{k_\mathrm{B} T}$$ and $$S = k_\mathrm{B} S_\mathrm{G}$$

temperature

Gibbs entropy is thermodynamic entropy!

System in a Bath

$$p_\alpha = \frac{e^{-\beta E_\alpha}}{Z(\beta)}$$

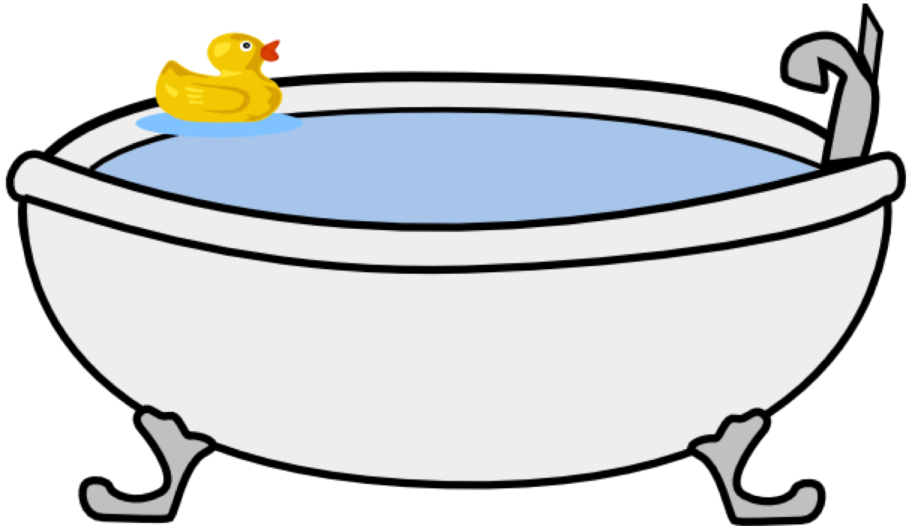

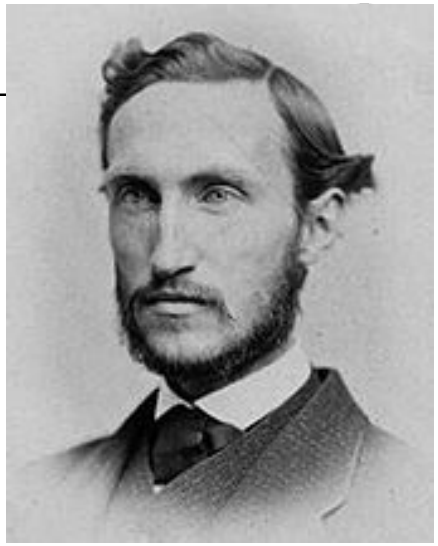

Josiah Willard Gibbs
(1839-1903)

It looks like we have calculated something, but what does it mean?

Physically, a system at fixed temperature is a system embedded in a "thermal bath",
i.e., a much larger environment with which it can exchange energy

$$\beta = \frac{1}{k_B T} \quad \text{and} \quad S = k_B S_G$$

temperature

Gibbs entropy is thermodynamic entropy!

System in a Bath

$$p_\alpha = \frac{e^{-\beta E_\alpha}}{Z(\beta)}$$

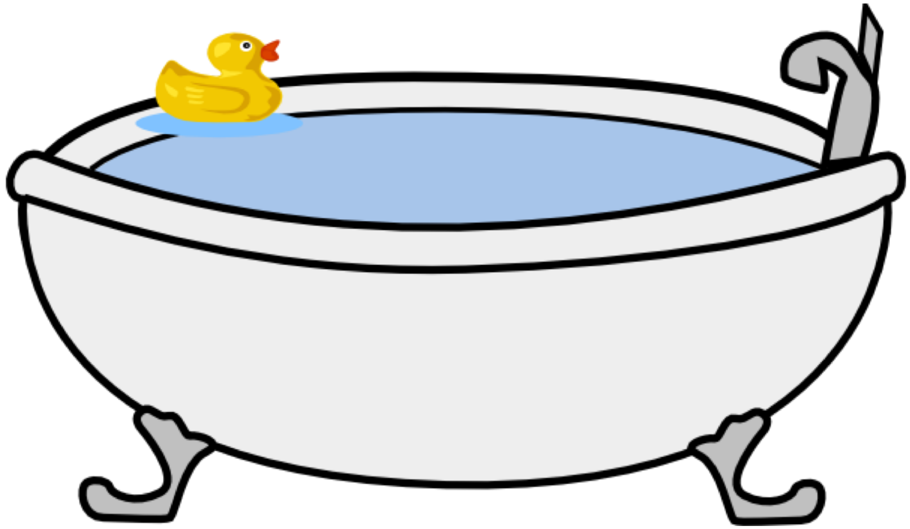

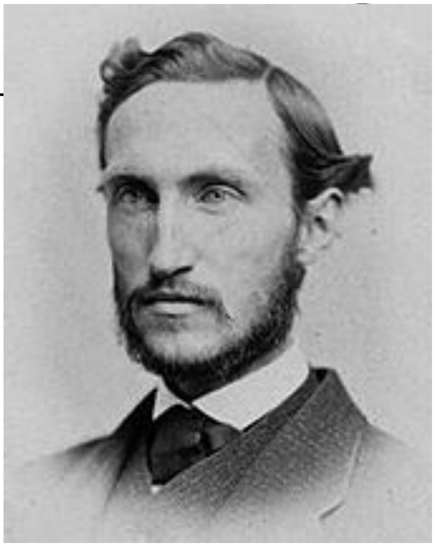

Josiah Willard Gibbs
(1839-1903)

It looks like we have calculated something, but what does it mean?

Physically, a system at fixed temperature is a system embedded in a "thermal bath",

i.e., a much larger environment with which it can exchange energy

more fun and games
with thermodynamics
in J. Chalker's lecture

The Second Law

The thermodynamic entropy of the world (closed system) increases with time.

The Second Law

The thermodynamic entropy of the world (closed system) increases with time.

Edwin Jaynes
(1922-1998)

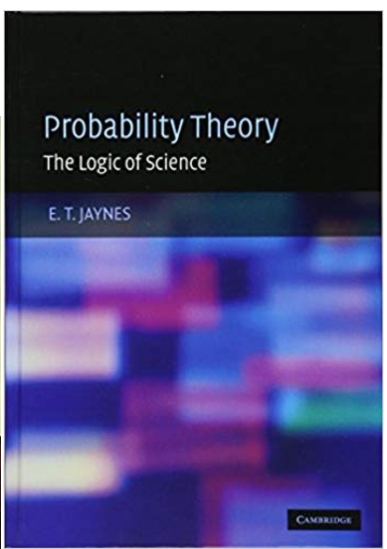

The Second Law

The thermodynamic entropy of the world (closed system) increases with time.

Time t: Measure something about some part of the world, maximise S_G subject to that info, get a set of p_α's and

$$S(t) = k_B S_{G,\max}(t)$$

Edwin Jaynes (1922-1998)

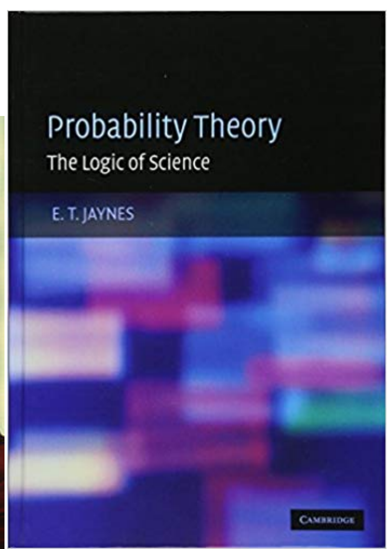

The Second Law

The thermodynamic entropy of the world (closed system) increases with time.

Time *t* : Measure something about some part of the world, maximise S_G subject to that info, get a set of p_α's and

$$S(t) = k_B S_{G,\max}(t)$$

Time *t'*>*t* : Consider the evolution of the system since t :
states evolve: $\alpha(t) \to \alpha(t')$,
but their probabilities stay the same, so

$$S_G(t') = -\sum_\alpha p_\alpha \ln p_\alpha = S_G(t) = S_{G,\max}(t) = \frac{1}{k_B} S(t)$$

The Second Law

The thermodynamic entropy of the world (closed system) increases with time.

Time *t*: Measure something about some part of the world, maximise S_G subject to that info, get a set of p_α's and

$$S(t) = k_B S_{G,\max}(t)$$

Time *t'*>*t*: Consider the evolution of the system since t:

states evolve: $\alpha(t) \to \alpha(t')$,

but their probabilities stay the same, so

$$S_G(t') = -\sum_\alpha p_\alpha \ln p_\alpha = S_G(t) = S_{G,\max}(t) = \frac{1}{k_B} S(t)$$

Now forget all you knew, make a new measurement, and maximise S_G again:

$$S(t') = \underbrace{k_B S_{G,\max}(t')}_{\substack{\text{the new } S_G\text{,}\\ \text{maximised at}\\ \text{time } t'}} \geqslant \underbrace{k_B S_G(t')}_{\substack{\text{the "true"}\\ S_G\text{, evolved}\\ \text{from time } t}} = k_B S_G(t) = S(t)$$

The Second Law

The thermodynamic entropy of the world (closed system) increases with time.

Time t: Measure something about some part of the world, maximise S_G subject to that info, get a set of p_α's and

$$S(t) = k_B S_{G,\max}(t)$$

Time $t' > t$: Consider the evolution of the system since t:
states evolve: $\alpha(t) \to \alpha(t')$,
but their probabilities stay the same, so

$$S_G(t') = -\sum_\alpha p_\alpha \ln p_\alpha = S_G(t) = S_{G,\max}(t) = \frac{1}{k_B} S(t)$$

Now forget all you knew, make a new measurement, and maximise S_G again:

$$S(t') = k_B S_{G,\max}(t') \geqslant k_B S_G(t') = k_B S_G(t) = S(t)$$

Thus, $\boxed{S(t') > S(t)}$ at $t' > t$, q.e.d.

The Second Law

The thermodynamic entropy of the world (closed system) increases with time.

The increase of S reflects our insistence to forget the detailed knowledge that we possess as a result of evolving in time any earlier state (even if based on earlier statistical inference) and re-apply at every later time the rules of statistical inference based on knowledge obtained in new measurements.

The Second Law

The thermodynamic entropy of the world (closed system) increases with time.

The increase of S reflects our insistence to forget the detailed knowledge that we possess as a result of evolving in time any earlier state (even if based on earlier statistical inference) and re-apply at every later time the rules of statistical inference based on knowledge obtained in new measurements.

- Clearly, the observer and her apparatus must be inside the system, otherwise it is not closed and observer will destroy the correspondence $\alpha(t) \rightarrow \alpha(t')$

The Second Law

The thermodynamic entropy of the world (closed system) increases with time.

The increase of S reflects our insistence to forget the detailed knowledge that we possess as a result of evolving in time any earlier state (even if based on earlier statistical inference) and re-apply at every later time the rules of statistical inference based on knowledge obtained in new measurements.

- Clearly, the observer and her apparatus must be inside the system, otherwise it is not closed and observer will destroy the correspondence $\alpha(t) \to \alpha(t')$

- S = entropy of the object of measurement + entropy of observer & her kit. A very precise measurement will imply massive increase in the latter. (More generally, the system gets entangled with its environment and when we measure it, our ignorance of the latter pollutes our knowledge of the former.)

The Second Law

The thermodynamic entropy of the world (closed system) increases with time.

The increase of S reflects our insistence to forget the detailed knowledge that we possess as a result of evolving in time any earlier state (even if based on earlier statistical inference) and re-apply at every later time the rules of statistical inference based on knowledge obtained in new measurements.

- Clearly, the observer and her apparatus must be inside the system, otherwise it is not closed and observer will destroy the correspondence $\alpha(t) \to \alpha(t')$

- S = entropy of the object of measurement + entropy of observer & her kit. A very precise measurement will imply massive increase in the latter. (More generally, the system gets entangled with its environment and when we measure it, our ignorance of the latter pollutes our knowledge of the former.)

go seriously quantum with Sid Parameswaran