Data

Star

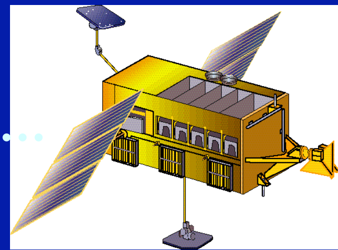

Rossi X-ray
Timing Explorer

10,000 light years

10^{46} X-rays/second

10^{3} X-rays/second

Poisson process (cf. Aneta):
photon counting noise
dominates!

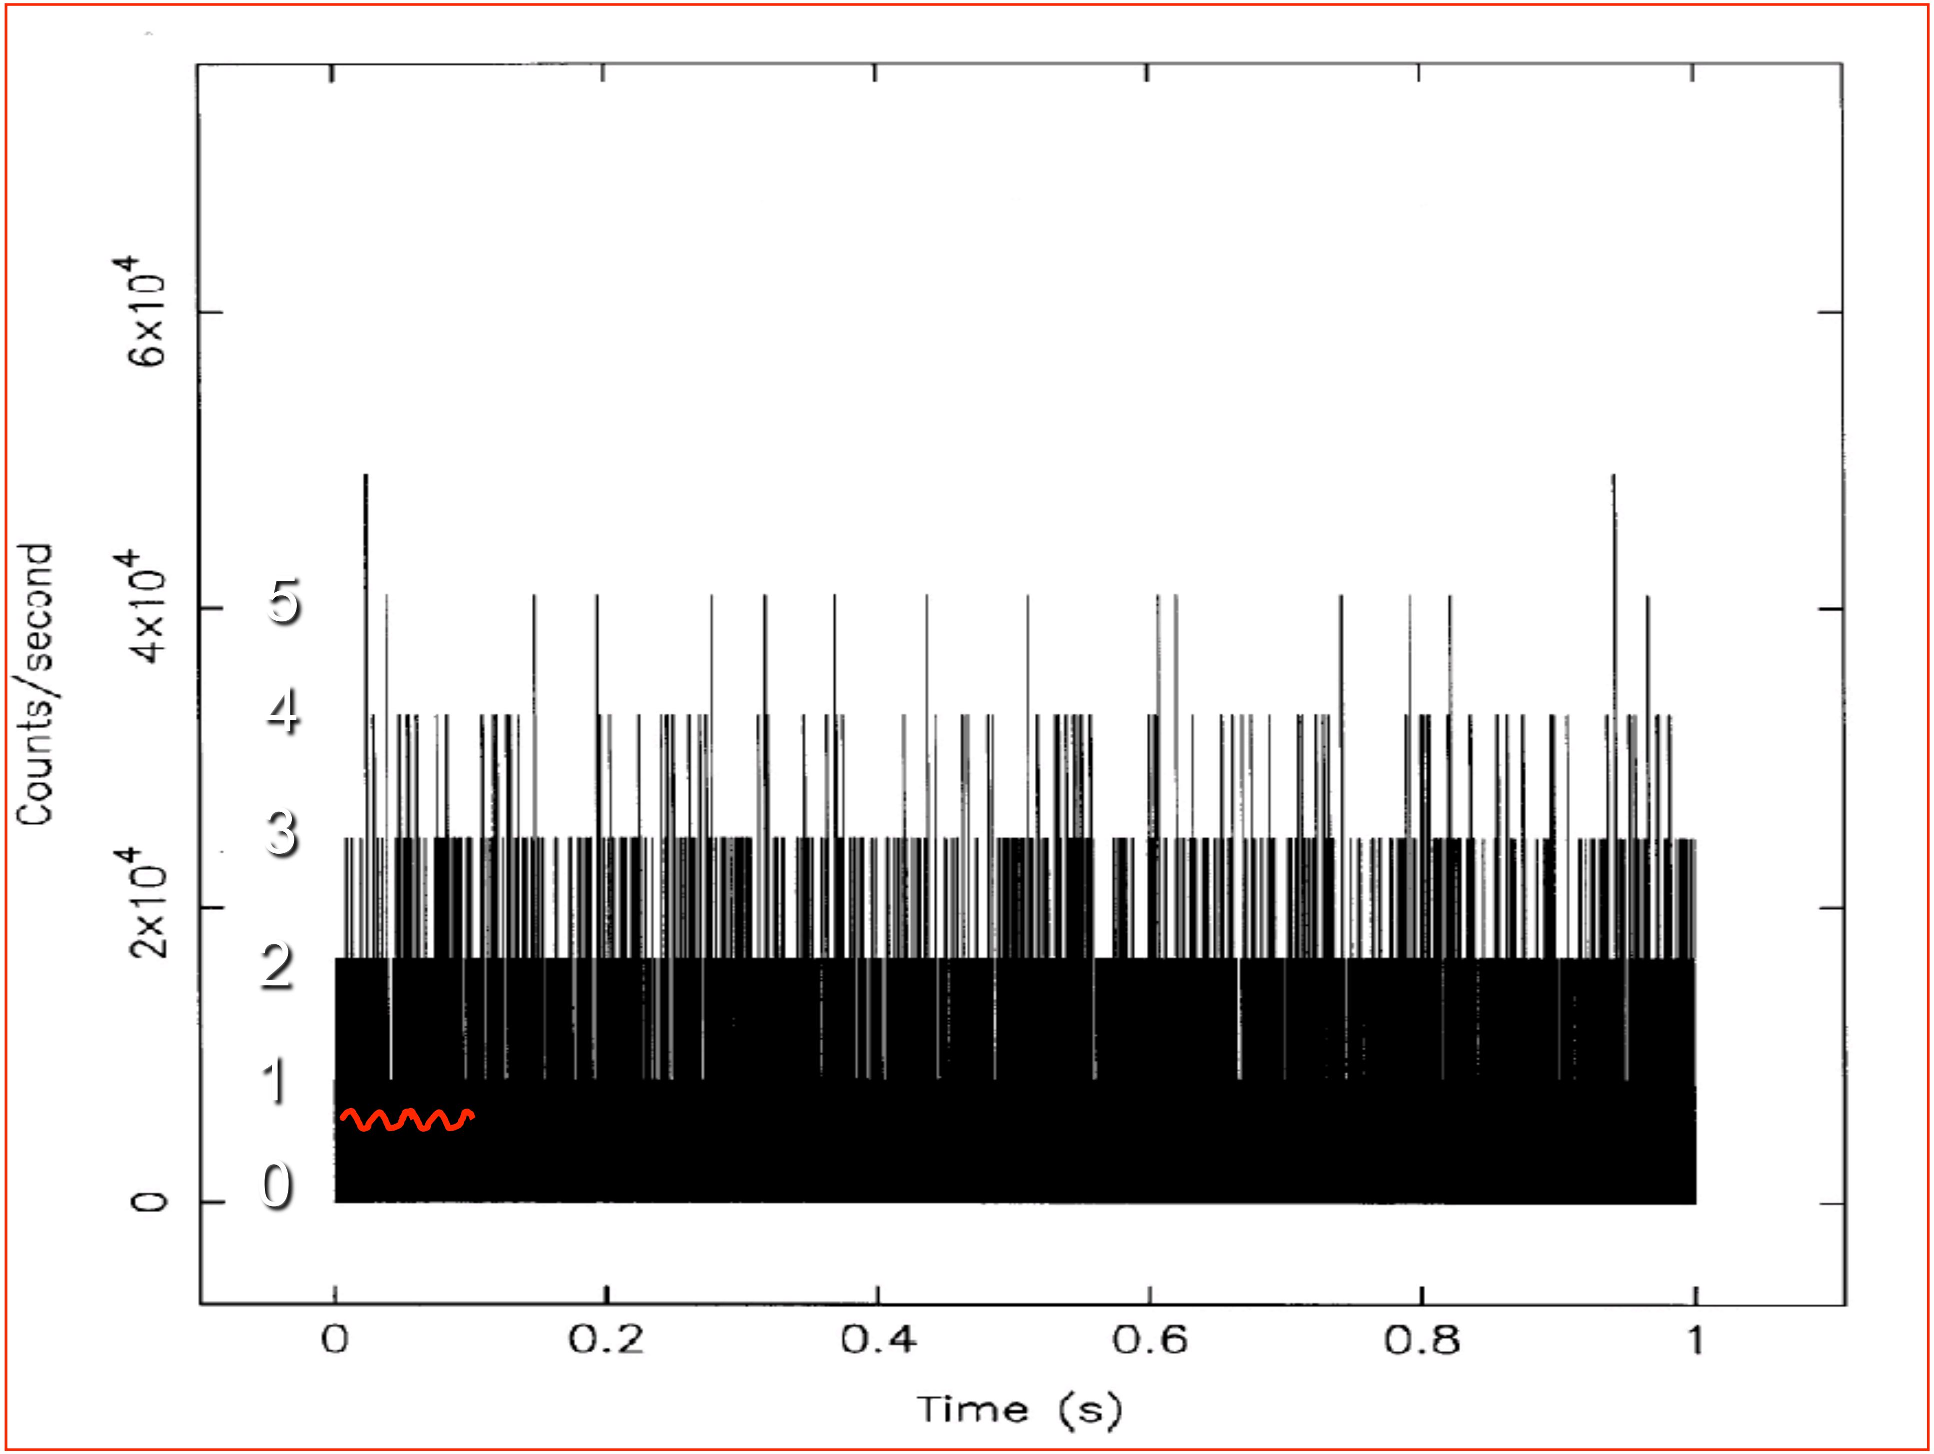

Counts/second
6x10^4
4x10^4
2x10^4
0
5
4
3
2
1
0
0
0.2
0.4
0.6
0.8
1
Time (s)

Examples of Fourier power spectra

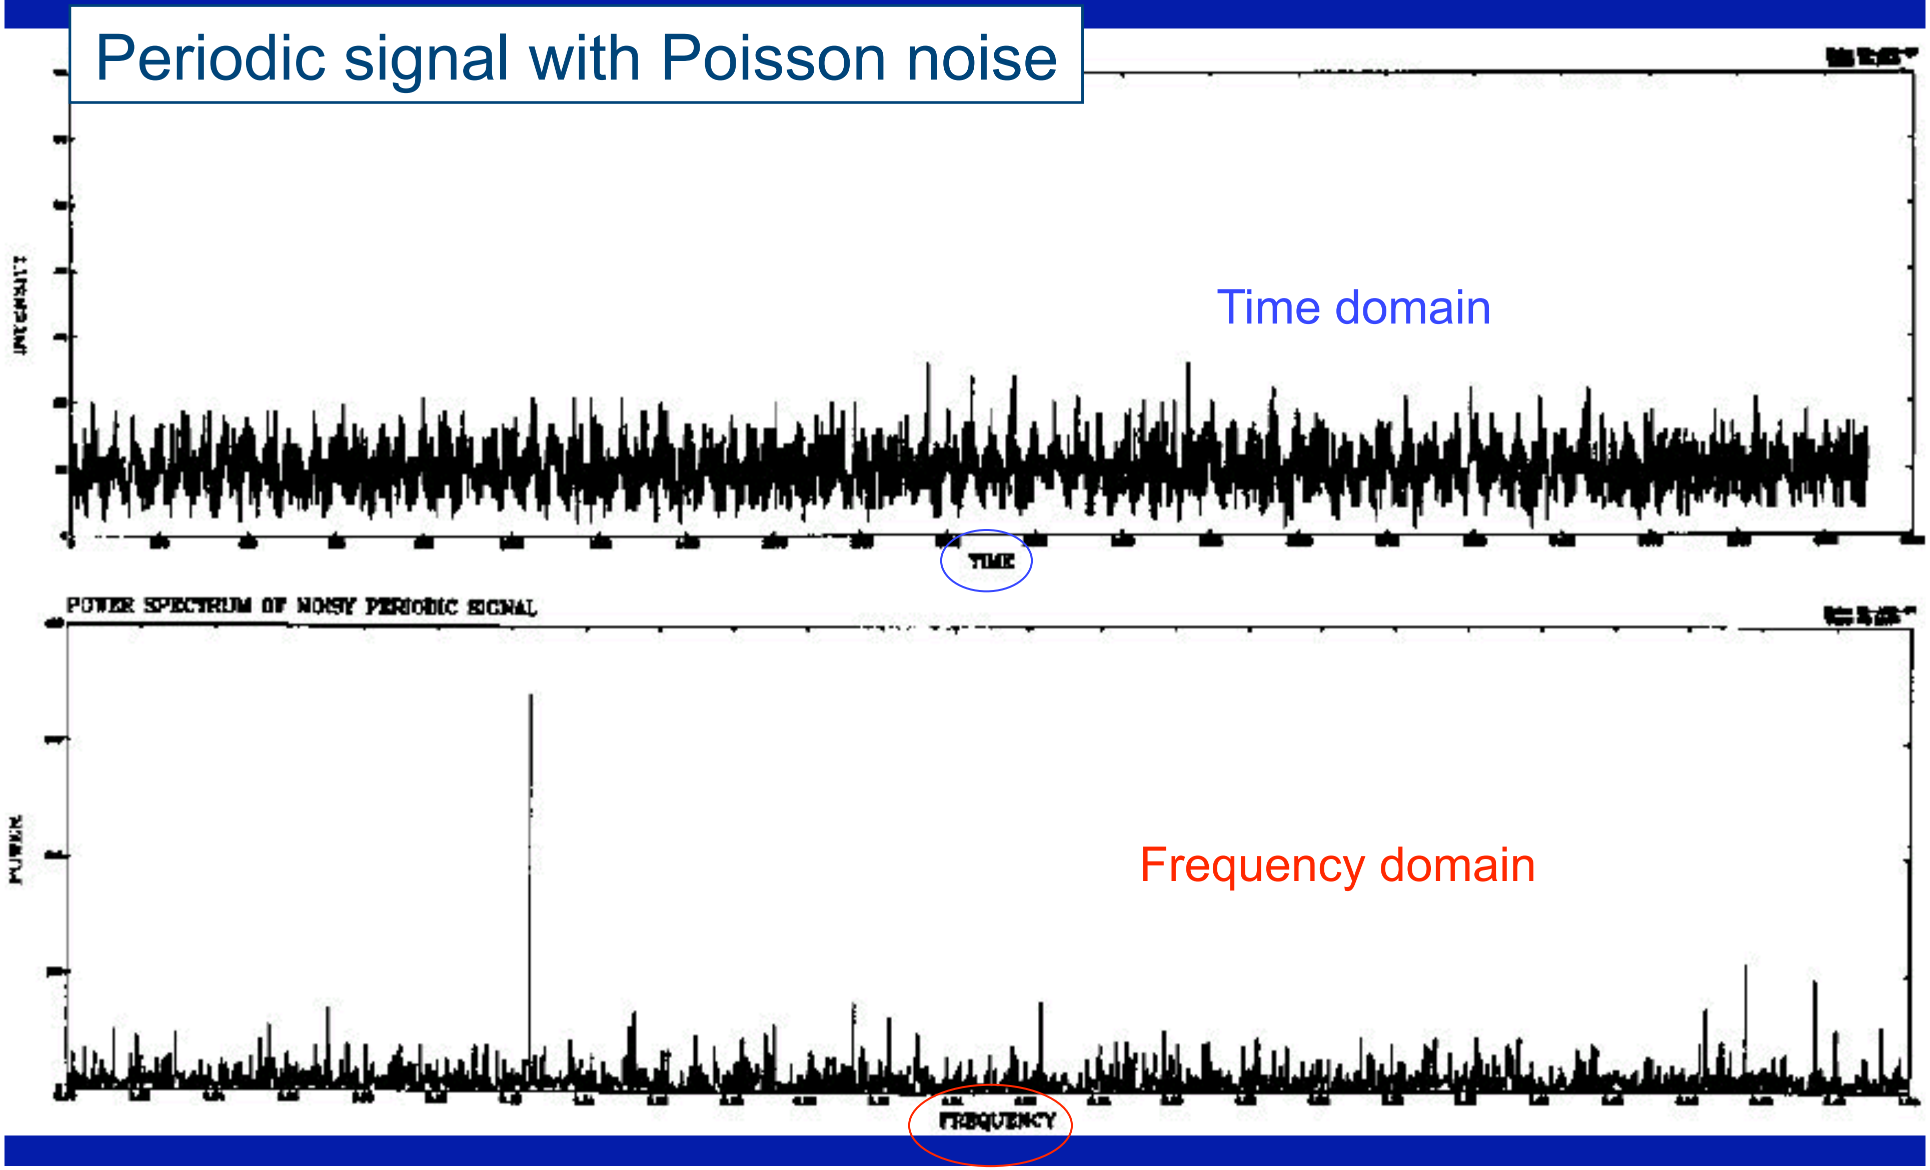

Quasi-periodic oscillation (QPO)

Quasi-periodic signal

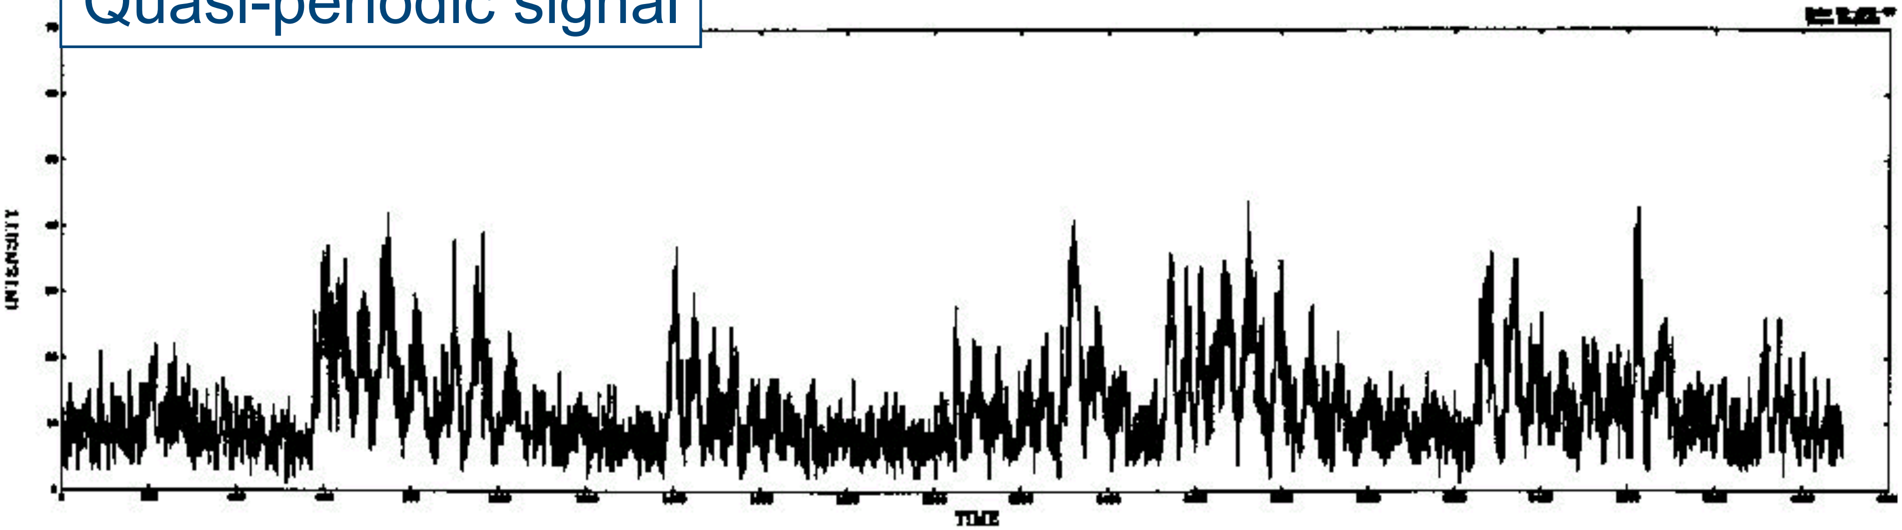

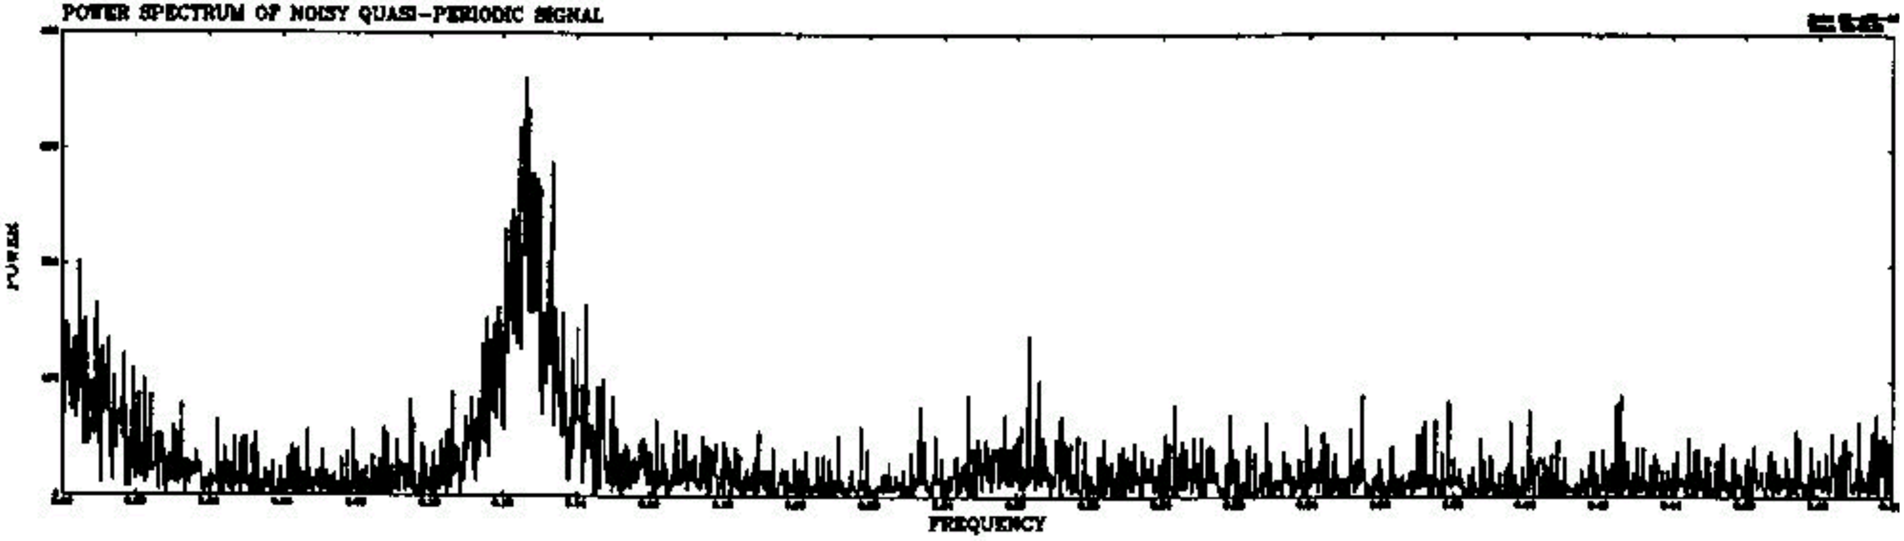

Black hole candidate XTE J1550-565

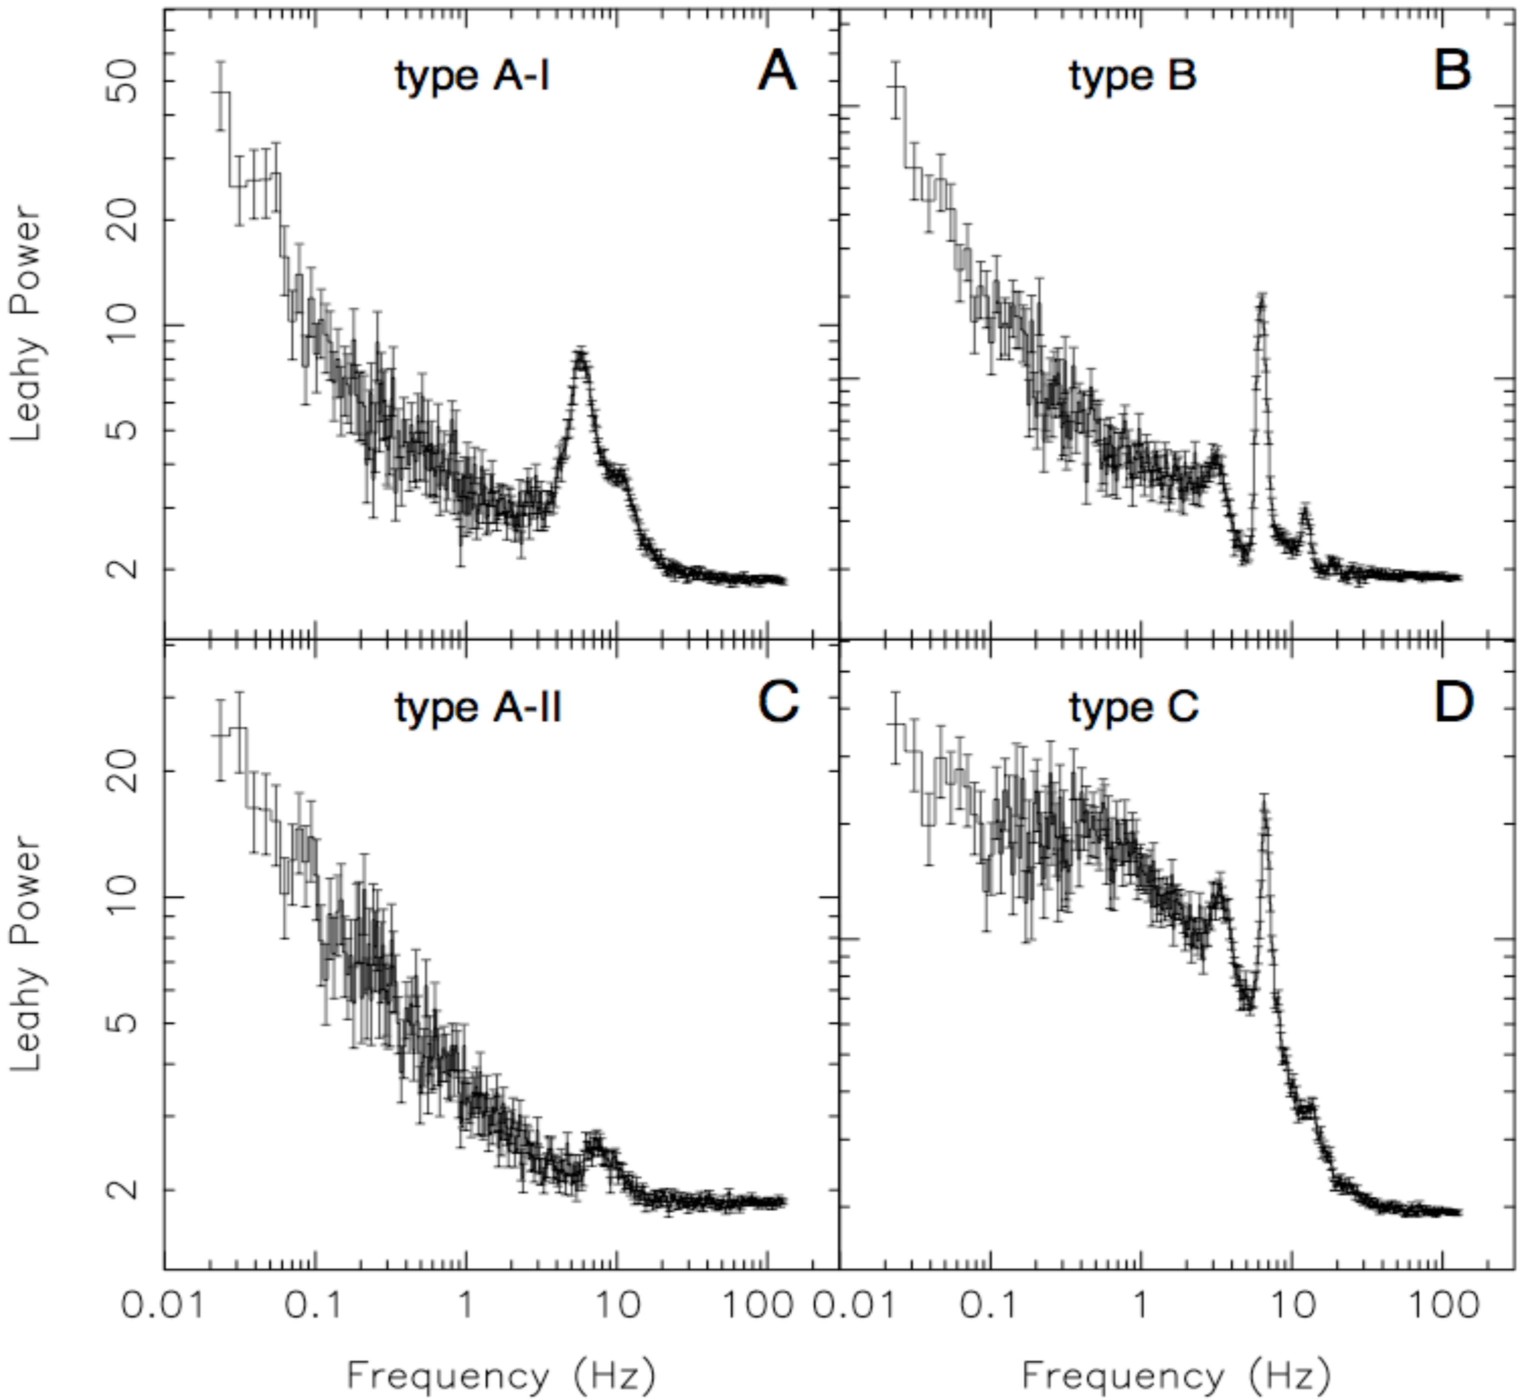

Neutron stars

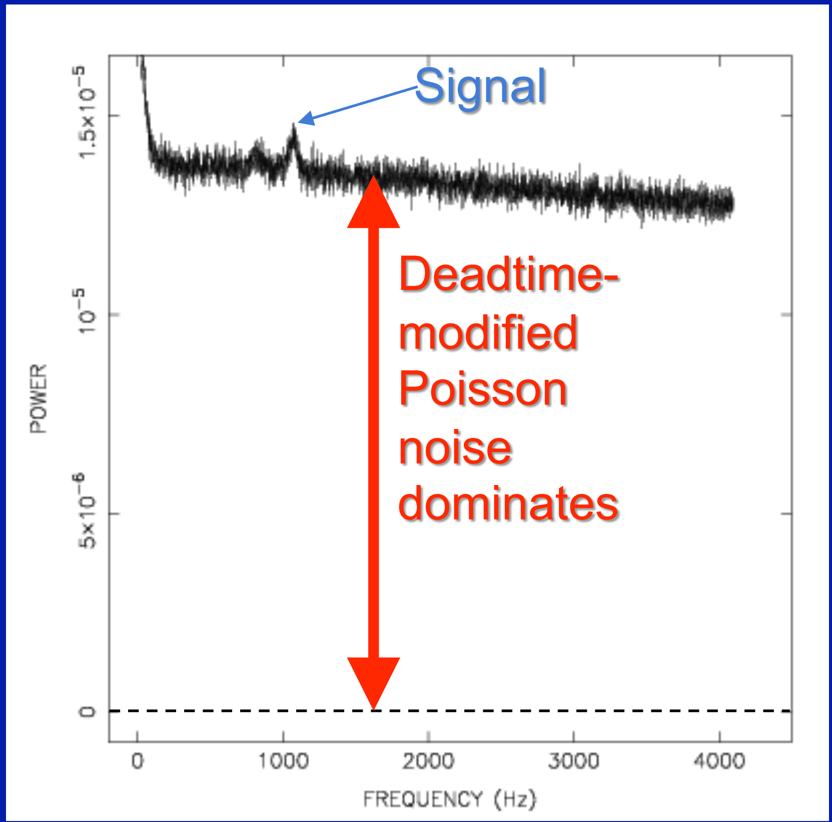

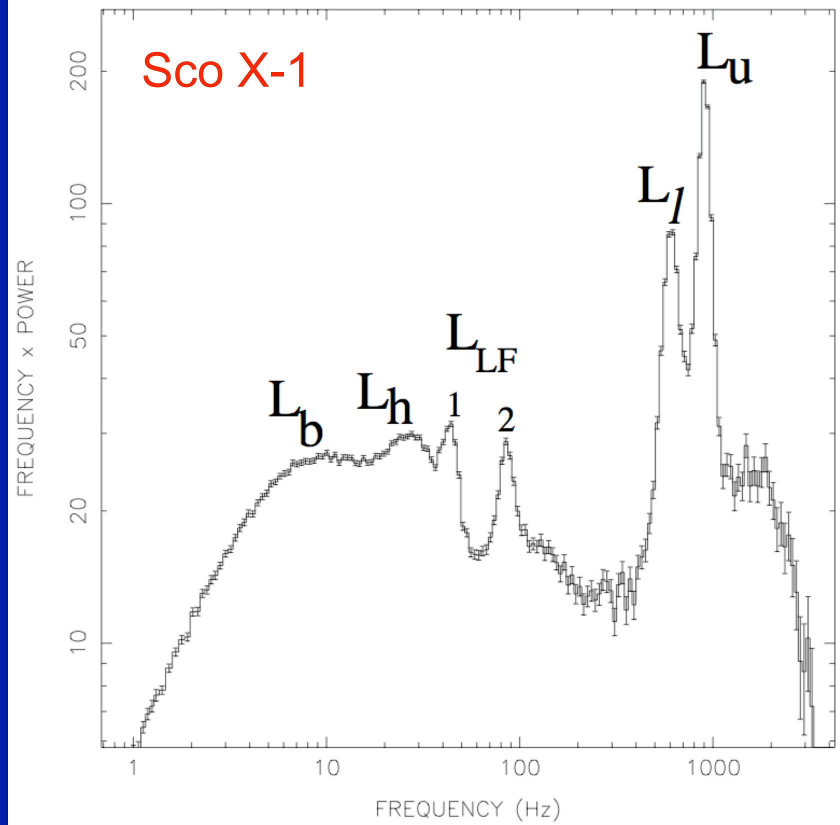

Averaging 1000's of seconds of data

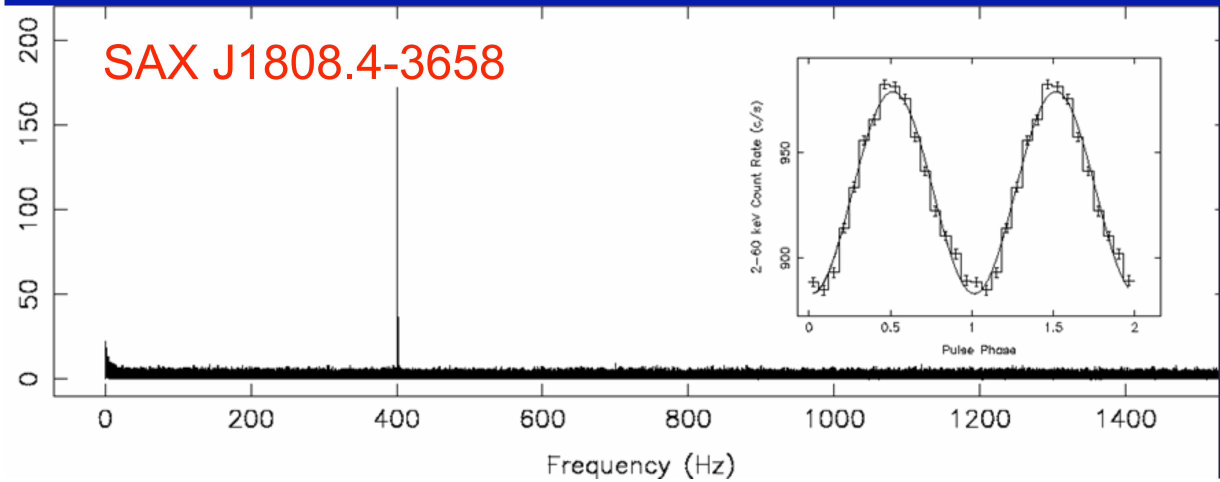

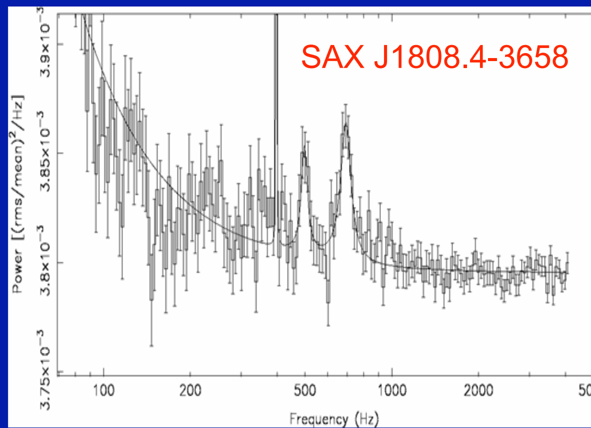

Orbital time scale in strong gravity region: milliseconds

$$\nu_K = \sqrt{GM/r^3}/2\pi \approx 1200\,\mathrm{Hz}\left(\frac{r}{15\,\mathrm{km}}\right)^{-3/2} m_{1.4}^{1/2}$$

$$\nu_{ISCO} = (6^{3/2}/2\pi)(c^3/GM) \approx (1580/m_{1.4})\,\mathrm{Hz}$$ (Schwarzschild)

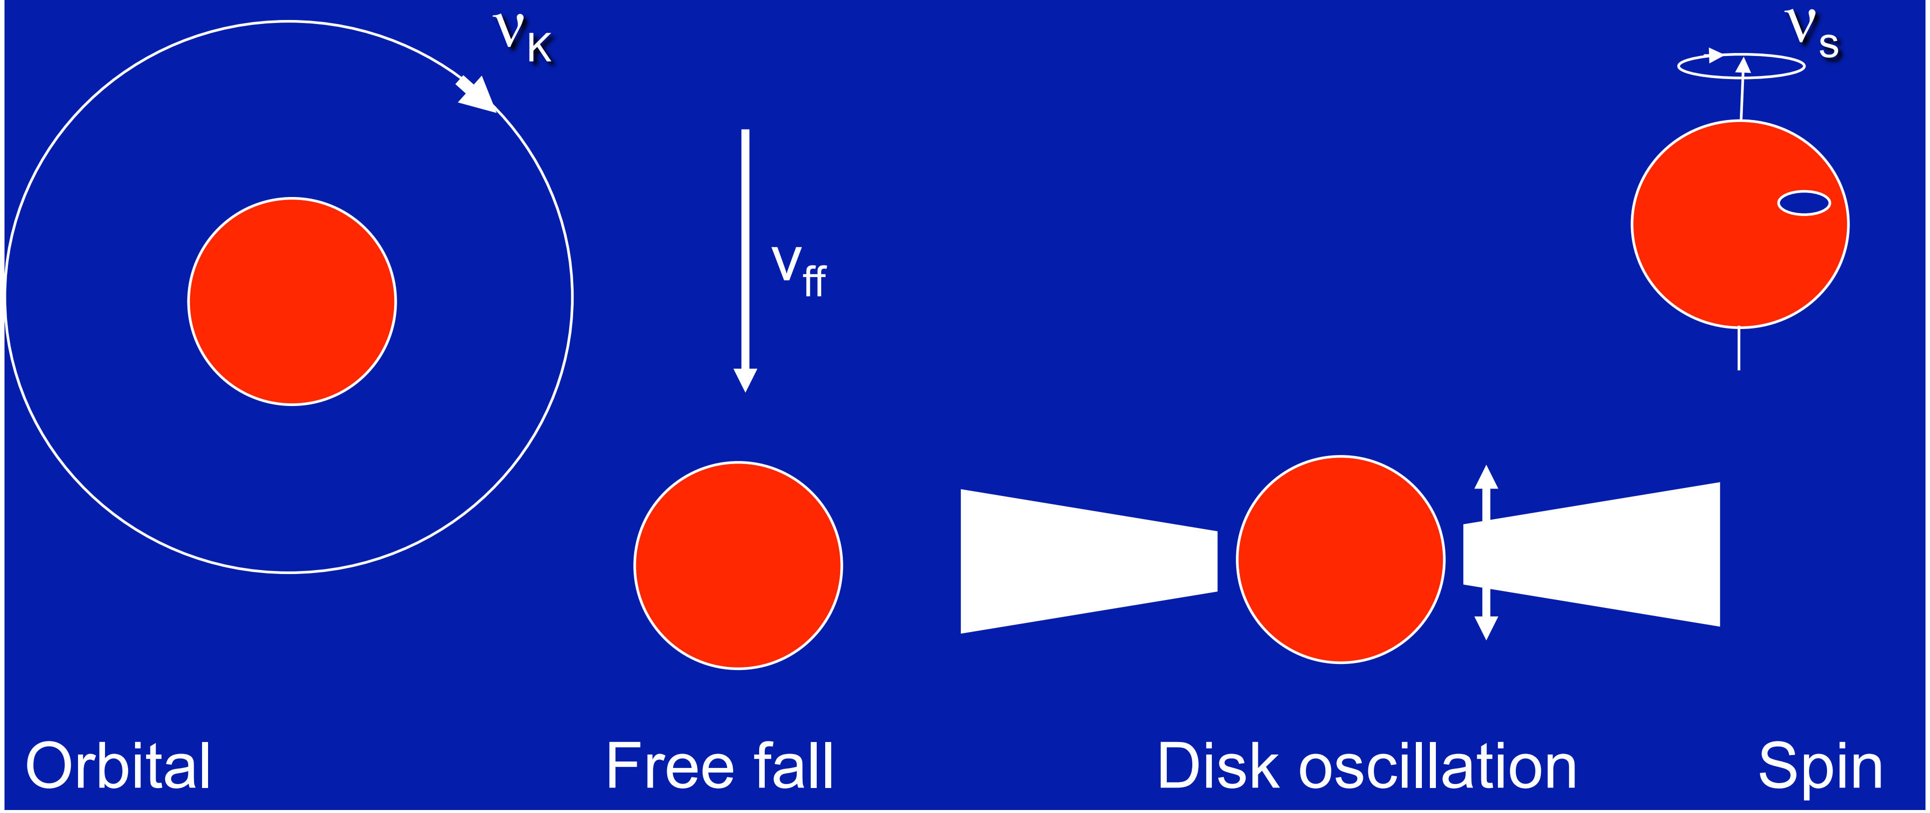

So we get these interesting power spectra . …

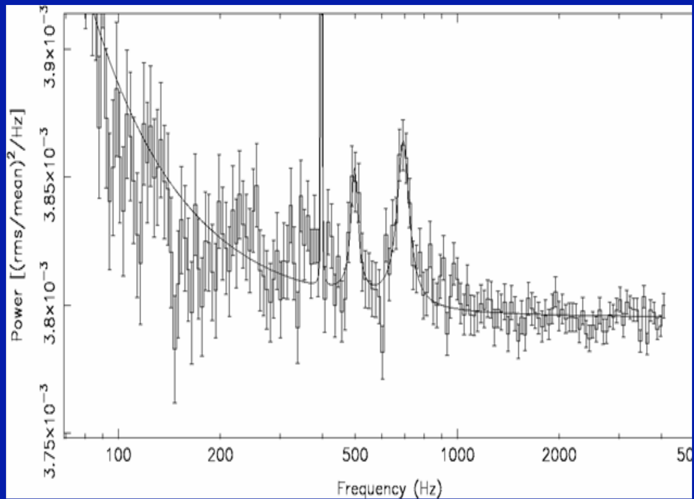

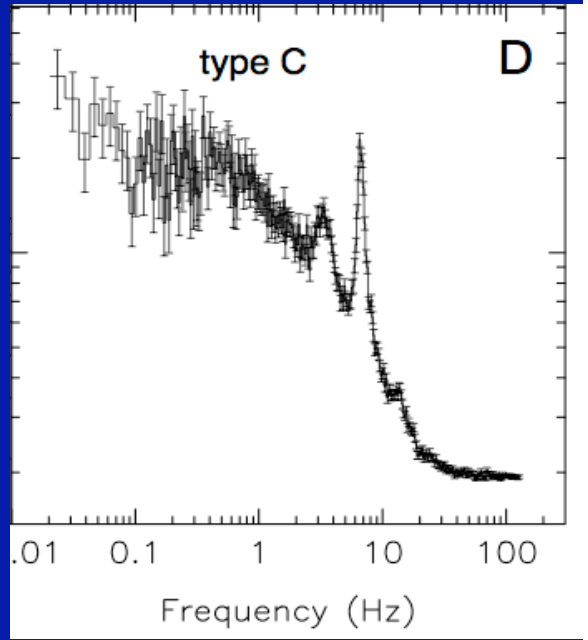

- What do the structures in the power spectra mean ?
- What is significant, what is not ?
- How to quantify what you can see ?

‘Noise’ (stochastic variability) in light curve causes broad components in the power spectrum

- Counting statistics noise (Poisson noise) → **white noise**
- Poisson noise modified by instrumental effects (e.g. deadtime)
 and other instrumental noise
- Noise that is stochastic source variability:
 QPO, band limited noise, red noise, etc.

All these can occur at the same time, possibly together with deterministic signals.

- They can be the background against which you are trying to detect something else.
- Or they can be the signal you are trying to detect!

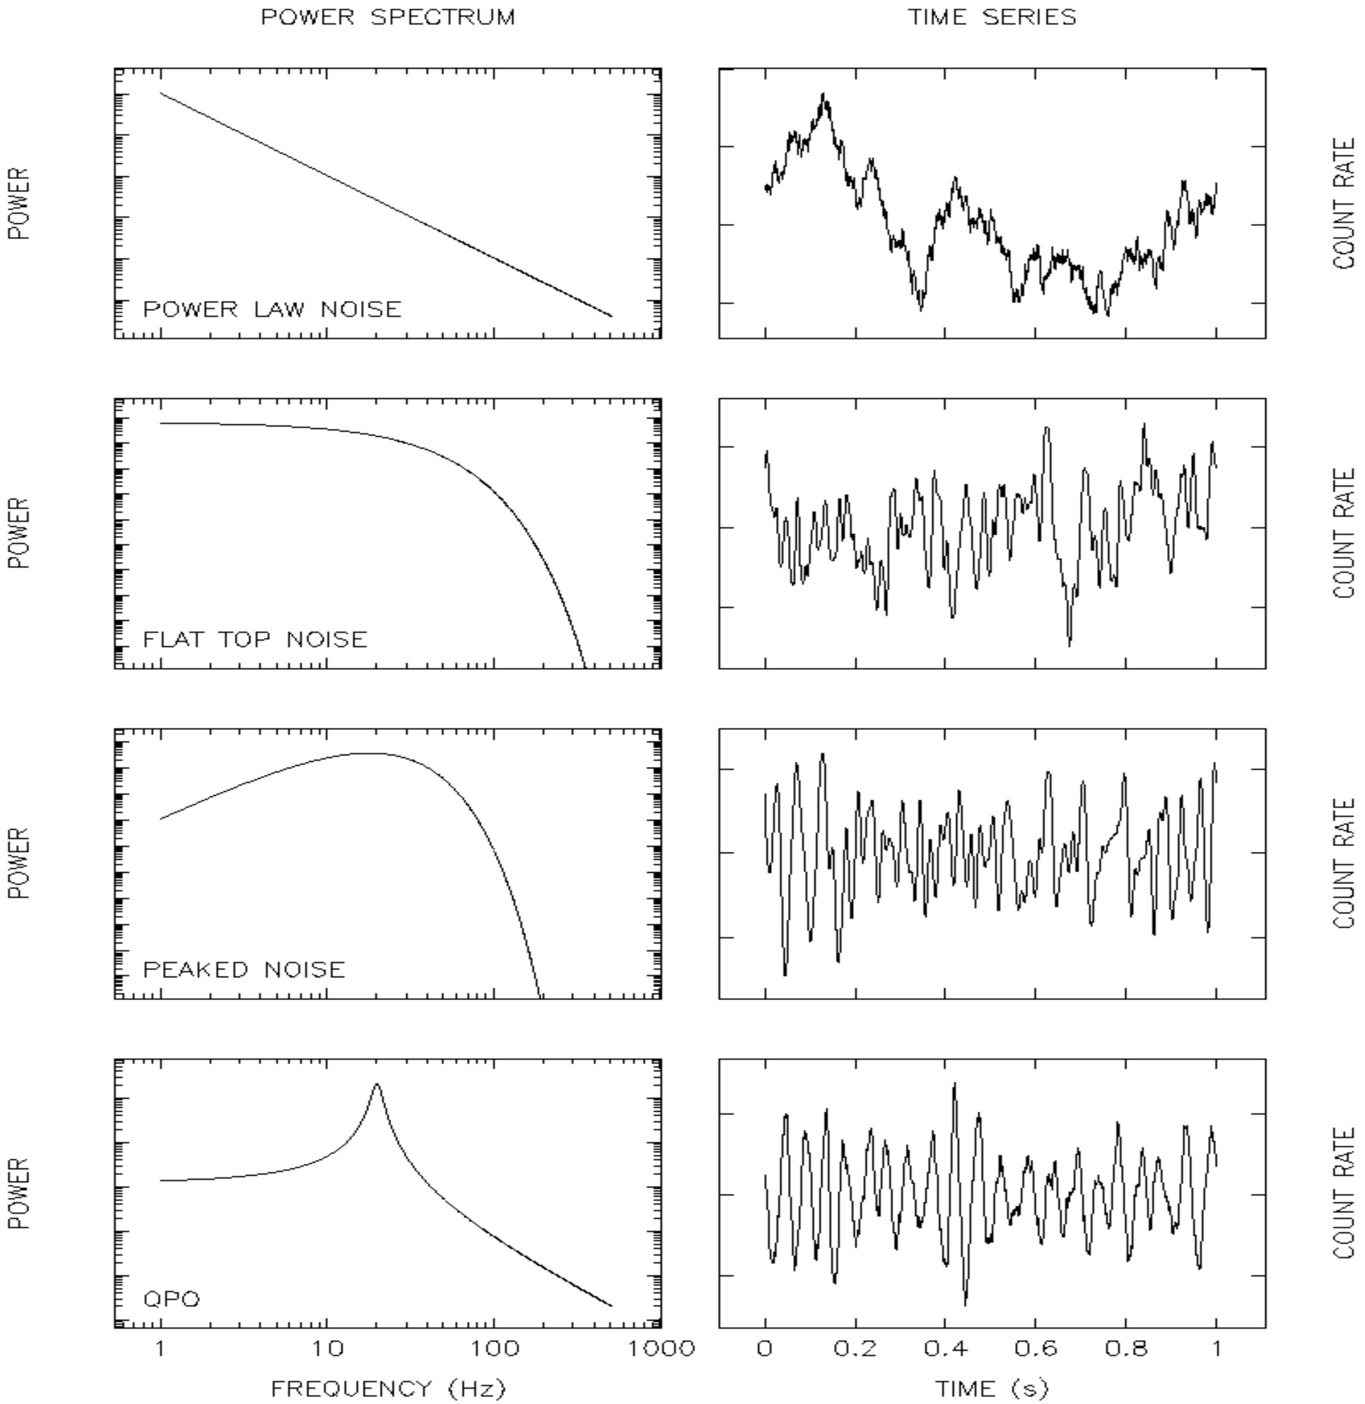
POWER SPECTRUM
TIME SERIES
POWER
COUNT RATE
POWER LAW NOISE
FLAT TOP NOISE
PEAKED NOISE
QPO
1
10
100
1000
FREQUENCY (Hz)
0
0.2
0.4
0.6
0.8
1
TIME (s)

Various possible QPO signals

Signal not directly observable in time domain.
Hence various possible time domain signals can underly the QPO peak we see in frequency domain

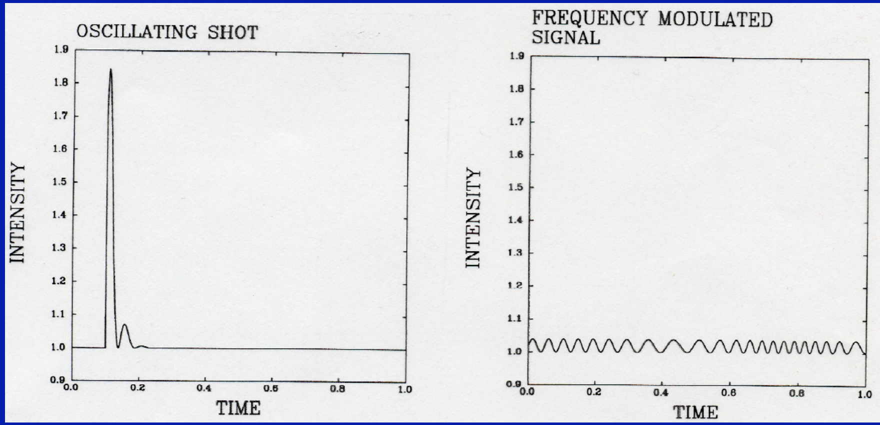

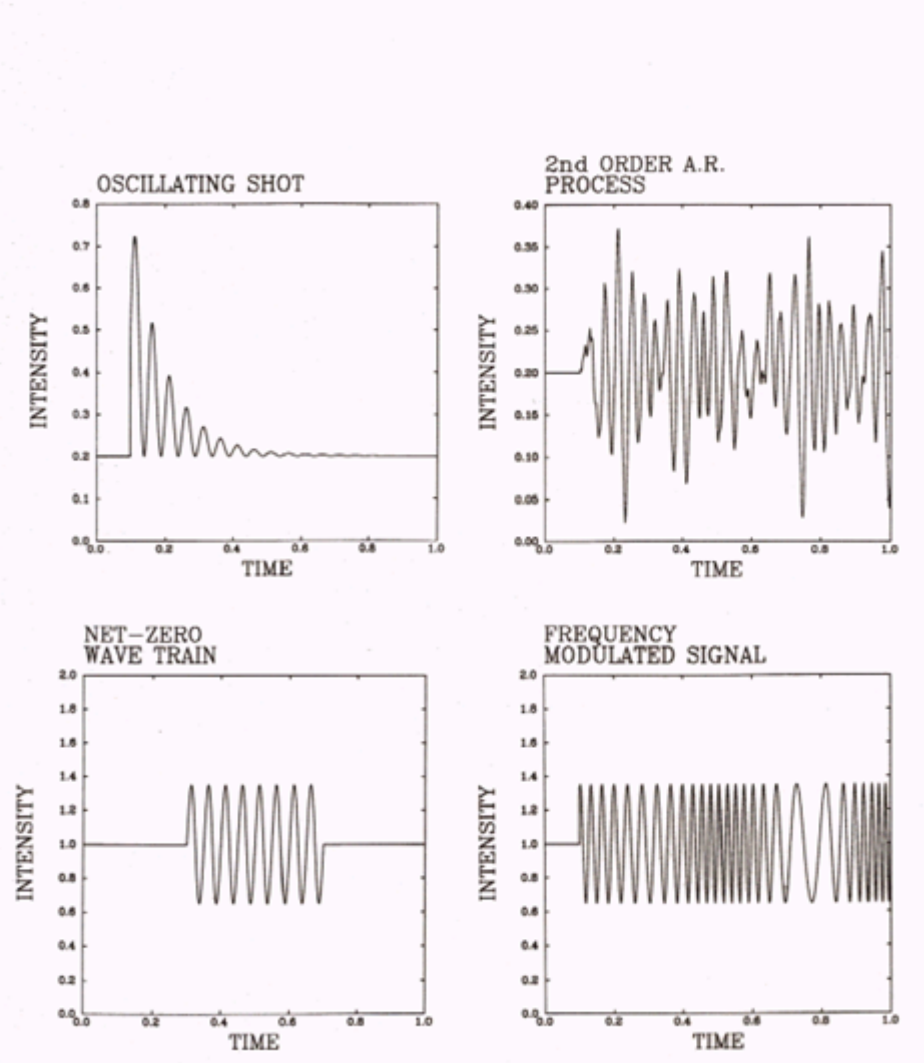

Selected literature

BOOKS:

• Jenkins & Watts: Spectral Analysis and its Applications – Holden-Day 1968
• Bloomfield: Fourier Analysis of Time Series: an Introduction – Wiley 1976
• Bracewell: The Fourier transform and its Applications – McGraw-Hill 1986
• Bendat & Piersol: Random data : Analysis and Measurement Procedures – Wiley 1986

ARTICLES:

• Groth 1975, ApJS 29, 285
• Leahy et al. 1983, ApJL 266, 160

• §2 of Lewin, van Paradijs & van der Klis 1988, SSR 46, 273
• van der Klis 1989 (in NATO ASI 'Timing Neutron Stars') updated 1994: http://staff.science.uva.nl/~michiel/Fourier_techniques.pdf

• Vaughan et al. 1994, ApJ 421, 738; 435, 362
• Vaughan & Nowak 1997, ApJL 474, 43

FOURIER TRANSFORM

A Fourier transform gives **decomposition of signal into sine waves.**

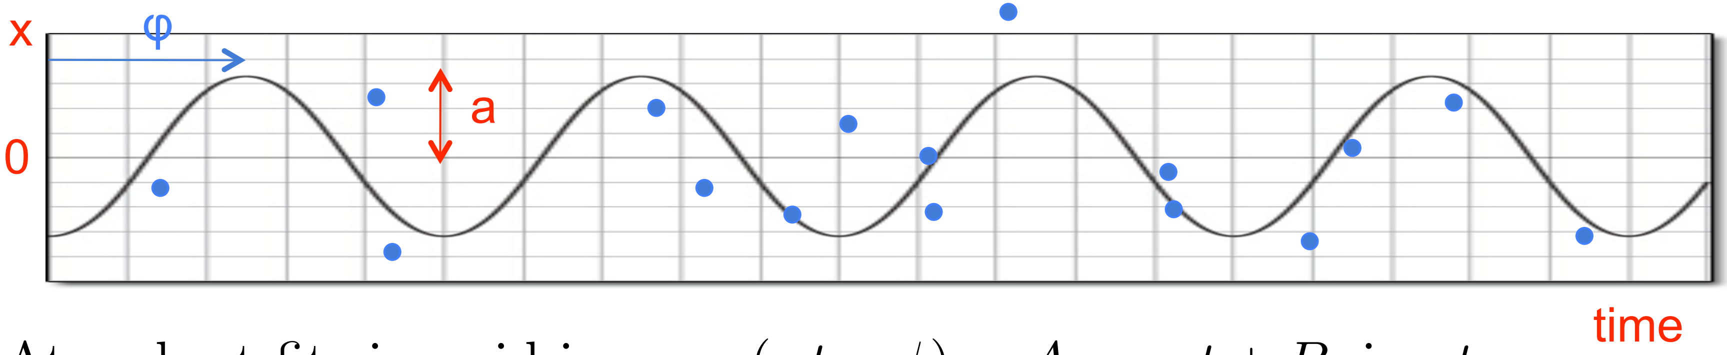

At ω, best-fit sinusoid is: $a\cos(\omega t - \phi) = A\cos\omega t + B\sin\omega t$
($a = \sqrt{A^2 + B^2}$ and $\tan\phi = -B/A$)

Do this at many frequencies ω_j, then

$$x(t) = \frac{1}{N}\sum_j a_j \cos(\omega_j t - \phi_j) = \frac{1}{N}\sum_j (A_j \cos\omega_j t + B_j \sin\omega_j t)$$

Fourier: $A_j = \sum_k x_k \cos\omega_j t_k$; $B_j = \sum_k x_k \sin\omega_j t_k$; $x_k \equiv x(t_k)$

So: **correlate** data with sine and cosine wave.
Good correlation: large A, B — bad correlation: small A, B

COMPLEX REPRESENTATION

A **way of handling** the two numbers (A, B or a, ϕ) you get at each ω.

$$a_j = \sum_k x_k e^{i\omega_j t_k}$$

$$x_k = \frac{1}{N} \sum_j a_j e^{-i\omega_j t_k}$$

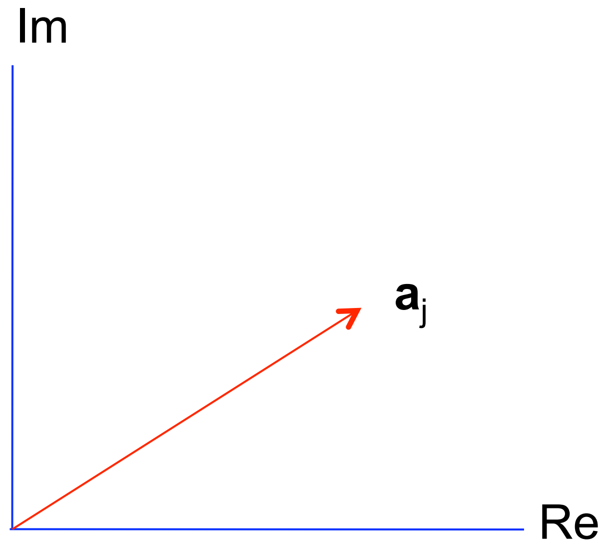

The **Fourier amplitudes** a_j are complex numbers:
$a_j = |a_j| e^{i\phi_j} = |a_j|(\cos\phi_j + i \sin\phi_j)$

If the signal x_k is real then imaginary terms at $+j$ and $-j$ cancel out in $\sum_j$,

to produce strictly real terms $2|a_j| \cos(\omega_j t_k - \phi_j)$

DISCRETE FOURIER TRANSFORM OF REAL TIME SERIES

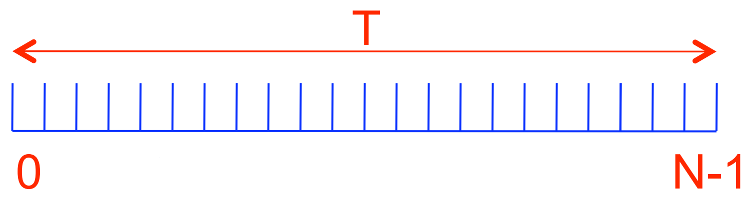

Time series: $x_k, \quad k = 0, \ldots, N-1$

Transform: $a_j, \quad j = -\frac{N}{2}+1, \ldots, \frac{N}{2}$

$$a_j = \sum_{k=0}^{N-1} x_k e^{2\pi i jk/N} \qquad j = -\frac{N}{2}+1, \ldots, \frac{N}{2}$$

$$x_k = \frac{1}{N} \sum_{j=-N/2+1}^{N/2} a_j e^{-2\pi i jk/N} \qquad k = 0, \ldots, N-1$$

Time step $\delta t = \frac{T}{N}$; Frequency step $\delta\nu = \frac{1}{T}$

x_k refers to time $t_k = \frac{kT}{N}$; a_j refers to frequency $\omega_j = 2\pi\nu_j = \frac{2\pi j}{T}$

So, for $e^{i\omega_j t_k}$ we have written $e^{2\pi i jk/N}$

DISCRETE FOURIER TRANSFORM OF REAL TIME SERIES - cont'd

- Fourier theorem: transform gives **complete** description of signal
- Highest frequency you need for this is the **Nyquist frequency**

$$\nu_{Ny} = \nu_{N/2} = \frac{N}{2T} = \text{half the sampling frequency } \frac{1}{\delta t} = \frac{N}{T}, \text{ as}$$

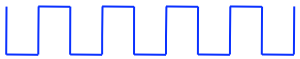

"up-down" is the fastest observable frequency.

$a_{N/2} = \sum_k x_k e^{i\pi k} = \sum_k x_k (-1)^k$ for real x_k is always real

- At zero frequency you get $a_0 = \sum_k x_k$, also always real for real x_k. (Called the **DC component**)
- At all frequencies in between you get complex Fourier amplitudes a_j, so:
- N, the number of input values $x_k \equiv$ number of output values; count them: a_0; $(|a_j|, \phi_j)$ pairs for $j = 1, \ldots, N/2 - 1$; $a_{N/2}$.
- Orthogonal, if the x_k are uncorrelated then the a_j are uncorrelated.

CONTINUOUS FOURIER TRANSFORM

Decomposes a **function** into an **infinite** number of sinusoidal waves.

Signal $x(t) \quad -\infty < t < \infty$

Transform $a(\nu) \quad -\infty < \nu < \infty$

$$a(\nu) = \int_{-\infty}^{\infty} x(t) e^{2\pi\nu i t}\, dt \qquad -\infty < \nu < \infty$$

$$x(t) = \int_{-\infty}^{\infty} a(\nu) e^{-2\pi\nu i t}\, d\nu \qquad -\infty < t < \infty$$

What is the relation of this 'ideal case' with the discrete Fourier transform **when we define** $x_k = x(t_k), \quad t_k = kT/N$?

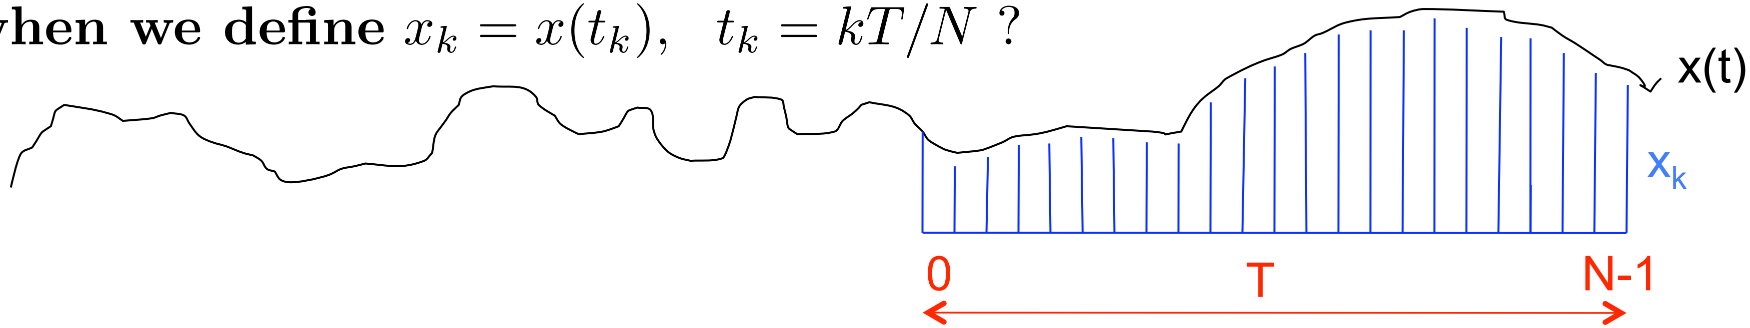

CONVOLUTION THEOREM

If $a(\nu)$ is the Fourier transform of $x(t)$ and
$b(\nu)$ is the Fourier transform of $y(t)$ then:

the transform of the product $x(t) \cdot y(t)$ is the convolution of $a(\nu)$ and $b(\nu)$:

$$a(\nu) \circledast b(\nu) \equiv \int_{-\infty}^{\infty} a(\nu')b(\nu - \nu')d\nu'$$

"the transform of the product is the convolution of the transforms" (and vv).
[Convolution denoted by $\circledast$]

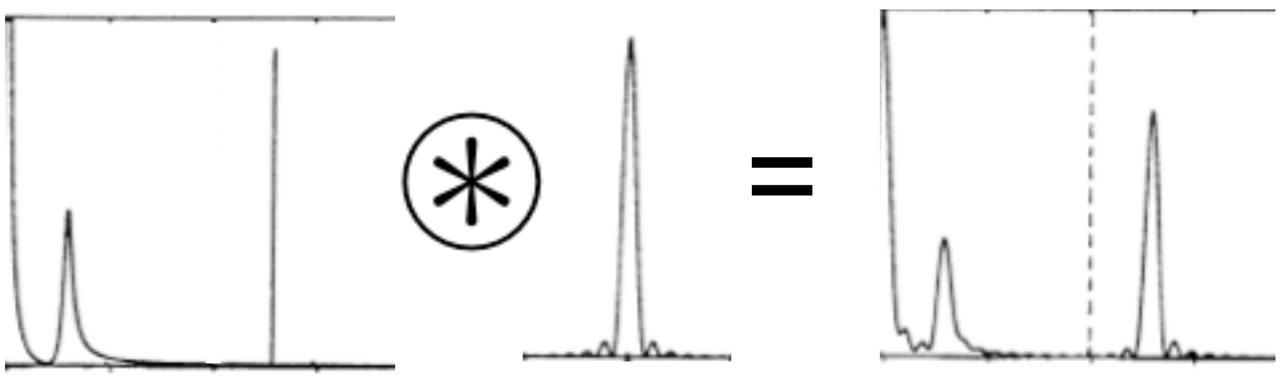

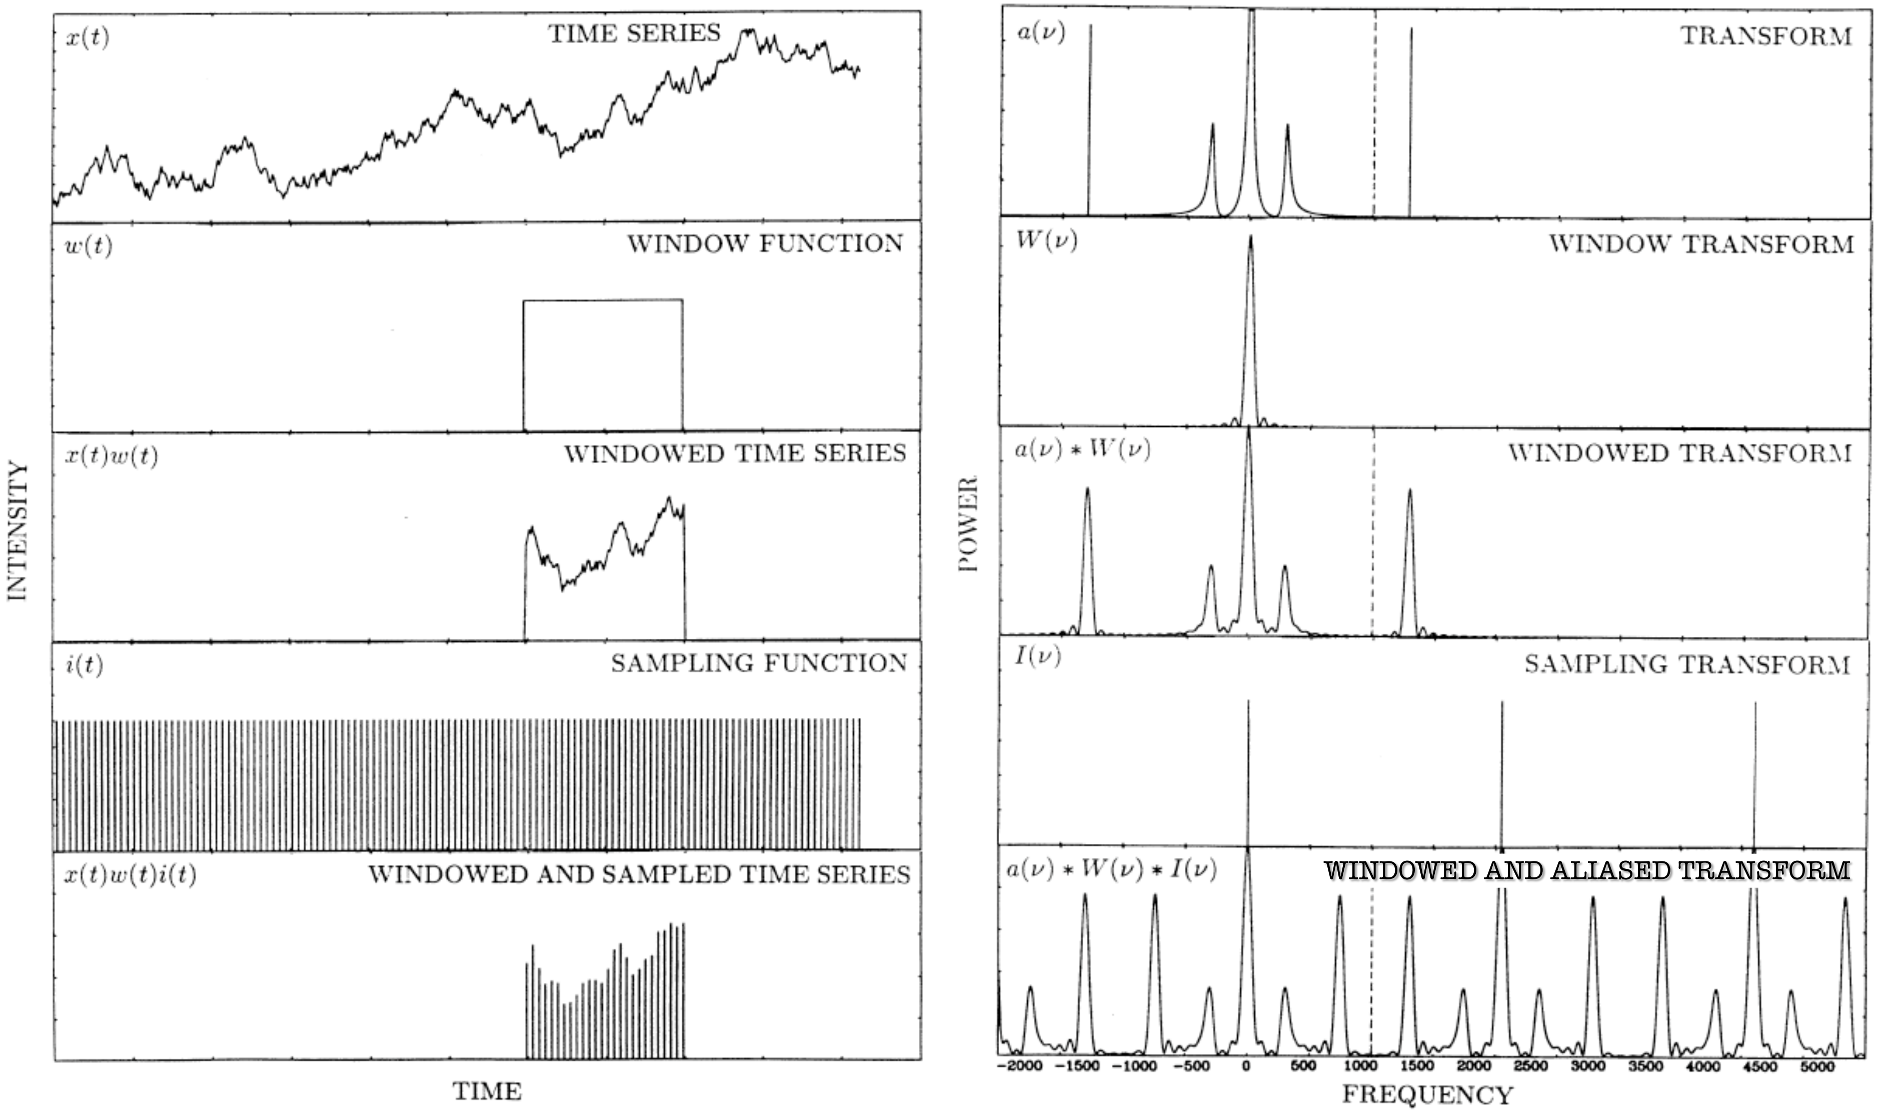

So: the discrete Fourier amplitudes are values **at the Fourier frequencies** of the **windowed** and **aliased** continuous Fourier transform.
Windowing: due to finite duration of the data convolve with window transform.
Aliasing: due to discrete sampling of data reflect around Nyquist frequency.

Is aliasing a problem?

Not so much as one might fear, as in practice, we do not really discretely sample the data, but rather **bin the data up!**

That means that before discrete sampling we **convolve** the x(t) with the bin width (we take a 'running average').

So, in the frequency domain, we **multiply** $a(\nu)$ with $B(\nu) = \dfrac{sin(\pi\nu T/N)}{\pi\nu T/N}$

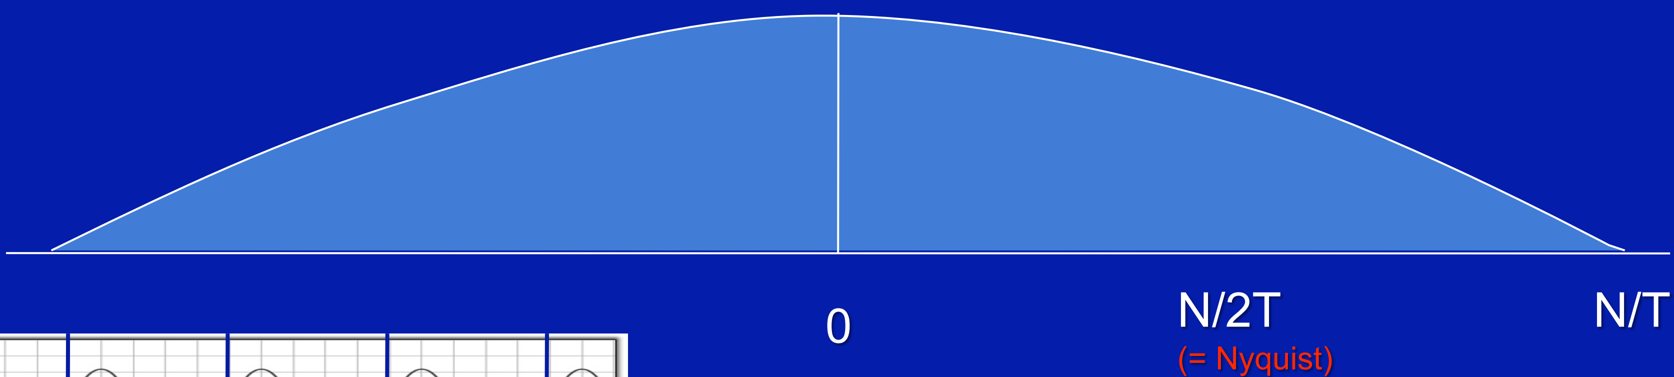

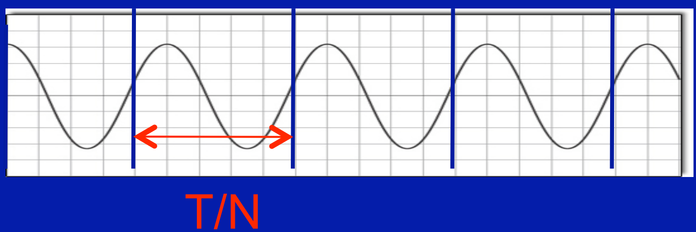

Is windowing a problem?

Yes, for steep spectra the **"leakage"** can be severe.

- Steep 'red noise' becomes less steep, limit ν^{-2}
- Delta functions become spread out

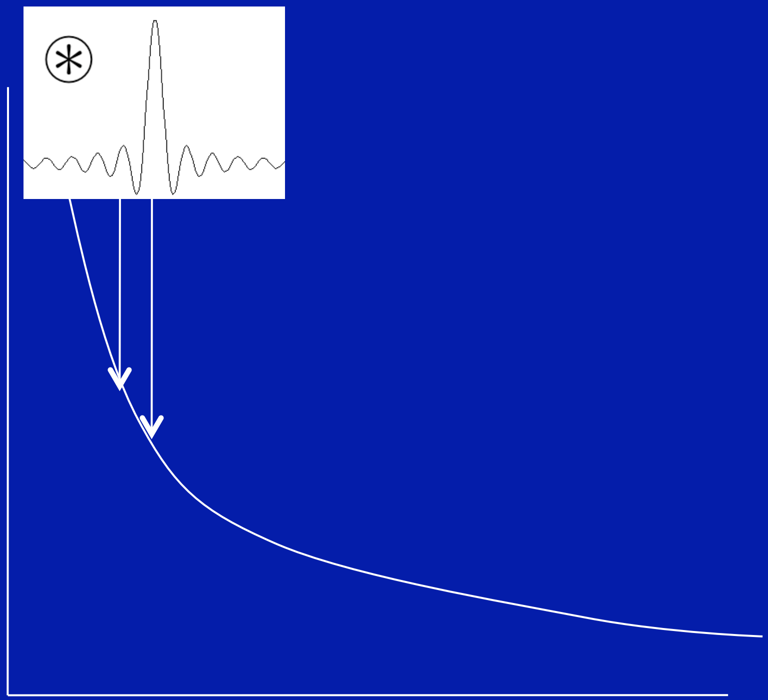

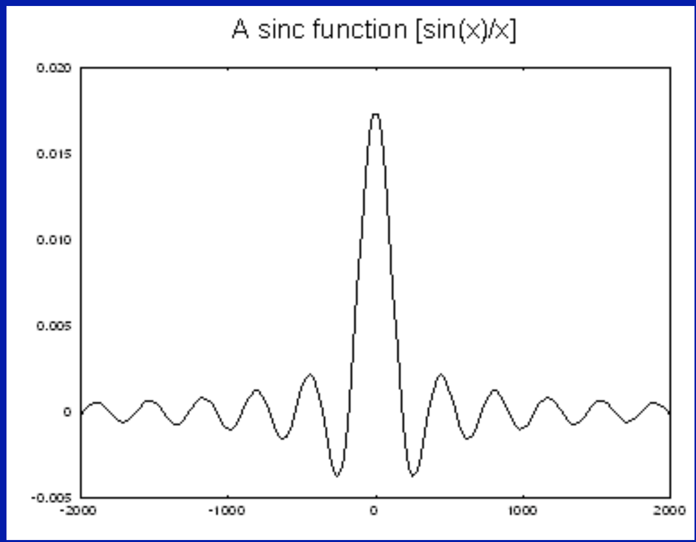

Fourier transform of a sinusoid

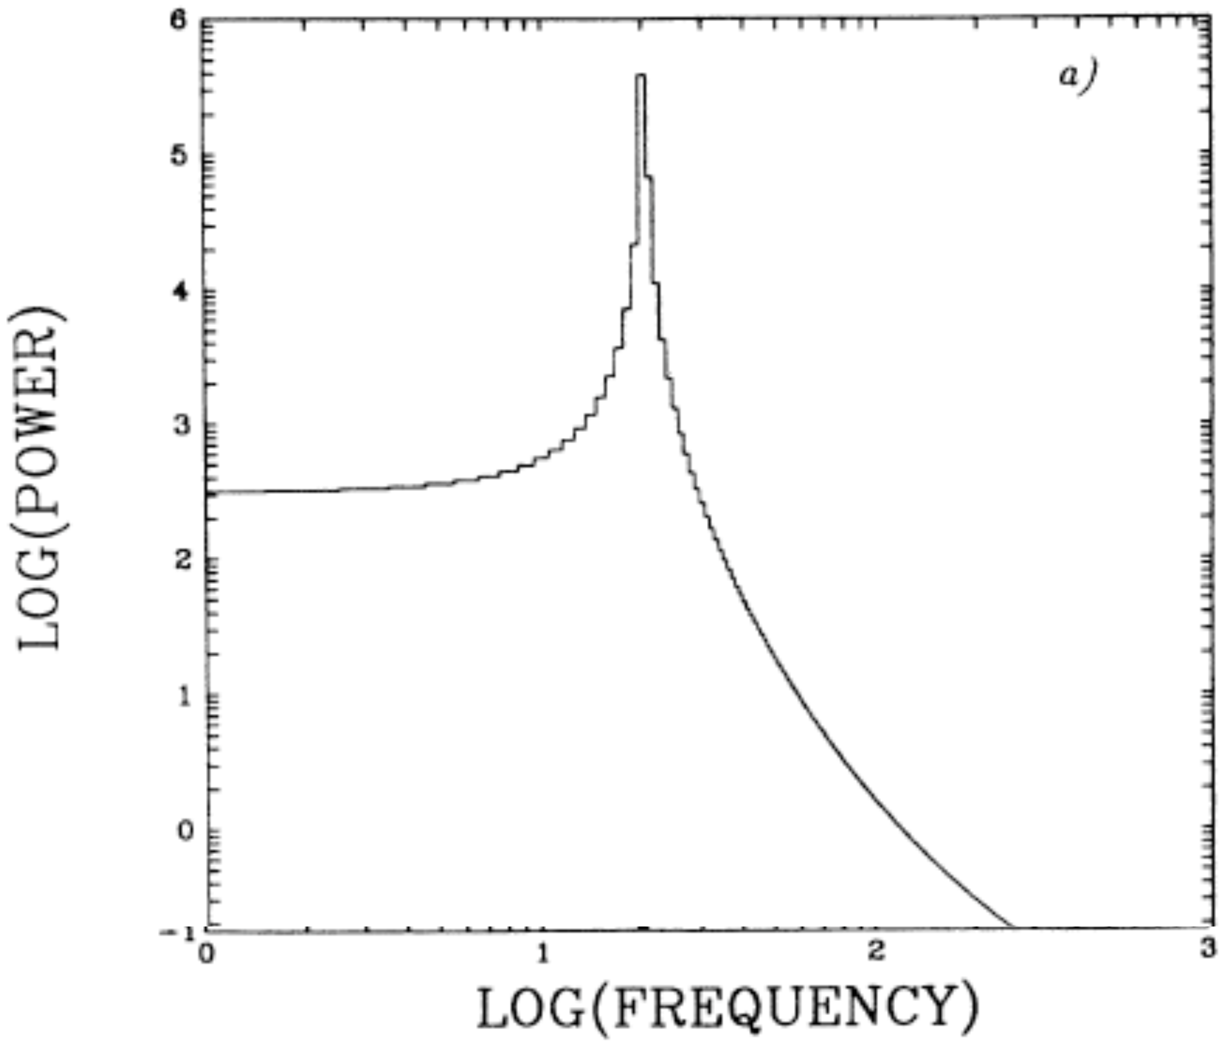

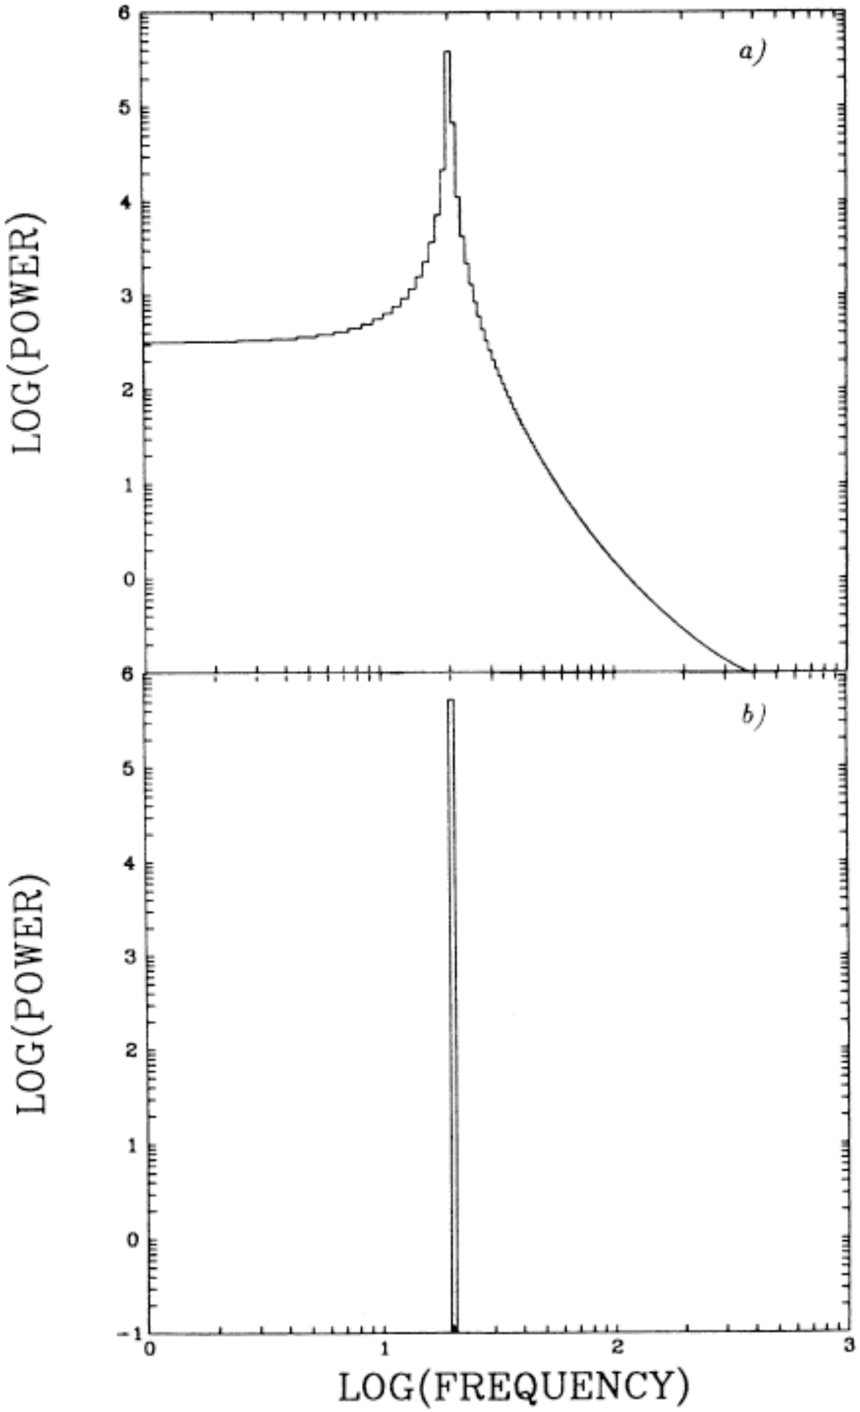

a)
b)
LOG(POWER)
LOG(POWER)
LOG(FREQUENCY)

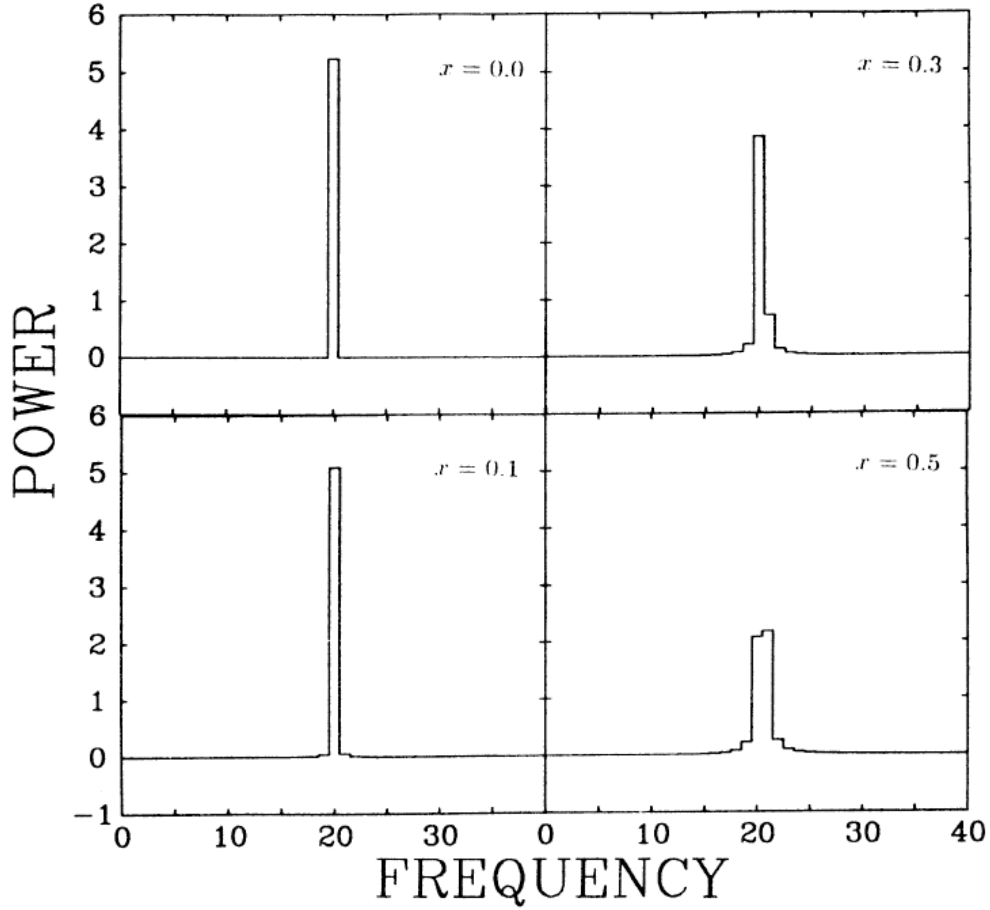

x = 0.0
x = 0.3
x = 0.1
x = 0.5
POWER
FREQUENCY

Fourier transform of a sinusoid

$$\boxed{|a_j|^2} = \tfrac{1}{4}A^2N^2\left(\frac{\sin \pi x}{\pi x}\right)^2\left[\left(\frac{\pi x/N}{\sin \pi x/N}\right)^2 + \left(\frac{\pi x/N}{\sin[\pi(2j+x)/N]}\right)^2 + \right.$$

$$\left. + 2\left(\frac{\pi x/N}{\sin \pi x/N}\right)\left(\frac{\pi x/N}{\sin[\pi(2j+x)/N]}\right)\cos[(N-1)(2\pi(j+x)/N)+2\phi]\right]$$

$$x = (\nu_{\mathrm{sine}} - \nu_j)T$$

$$\boxed{\approx \tfrac{1}{4}A^2N^2\left(\frac{\sin \pi x}{\pi x}\right)^2} \qquad x/N \ll 1 \text{ and } 0 \ll j/N \ll \tfrac{1}{2}$$

Lorentzians

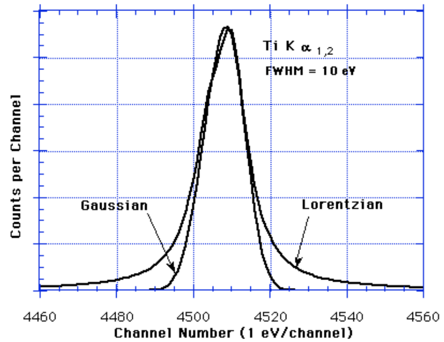

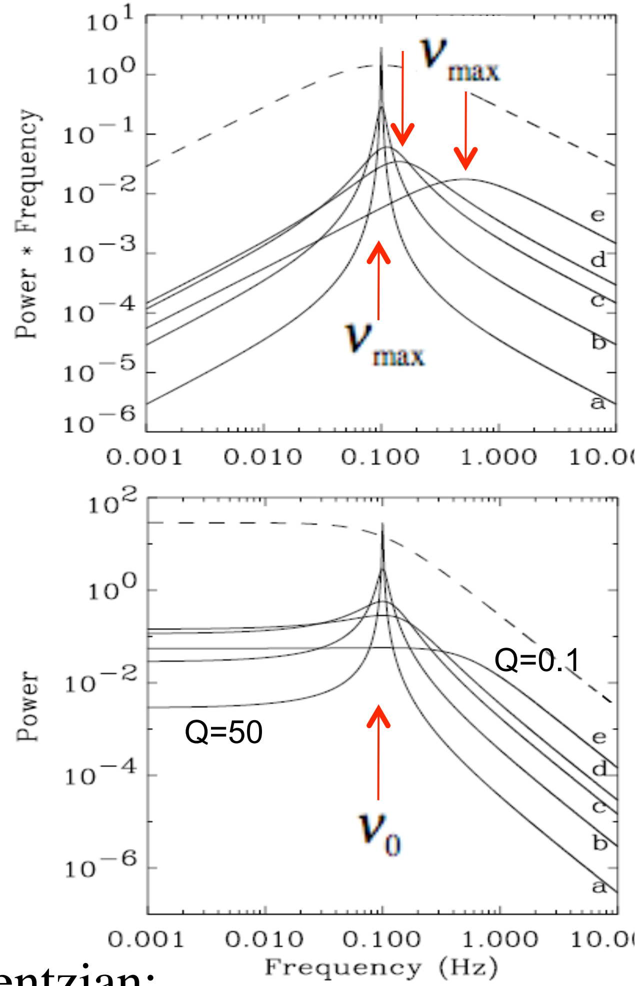

$$\text{Lorentzian}: P(\nu) = \frac{r^2 \Delta}{\pi} \frac{1}{\Delta^2 + (\nu - \nu_0)^2}$$

$$\text{Centroid frequency: } \nu_0 \;\; ; \;\; \text{Power}: r^2 = \int_{\nu=-\infty}^{\infty} P(\nu)\, d\nu$$

$$\text{Half-width}: \Delta \quad ; \qquad \text{Coherence}: Q \equiv \nu_0 / 2\Delta$$

$$\text{Characteristic frequency}: \nu_{\max} = \sqrt{\nu_0^2 + \Delta^2}$$

The power spectrum of an exponential $x(t) = e^{-t/\tau}$ is a Lorentzian:

$P(\nu) \propto \dfrac{1}{\Delta^2 + \nu^2}$. So the power spectrum of an exponentially damped sinusoid

$x(t) = e^{-t/\tau} \times \cos(2\pi\nu_0 t)$ is the convolution $P(\nu) \propto \dfrac{1}{\Delta^2 + \nu^2} \otimes \delta(\nu - \nu_0)$.

But: there are many other ways in which Lorentzians can be produced!

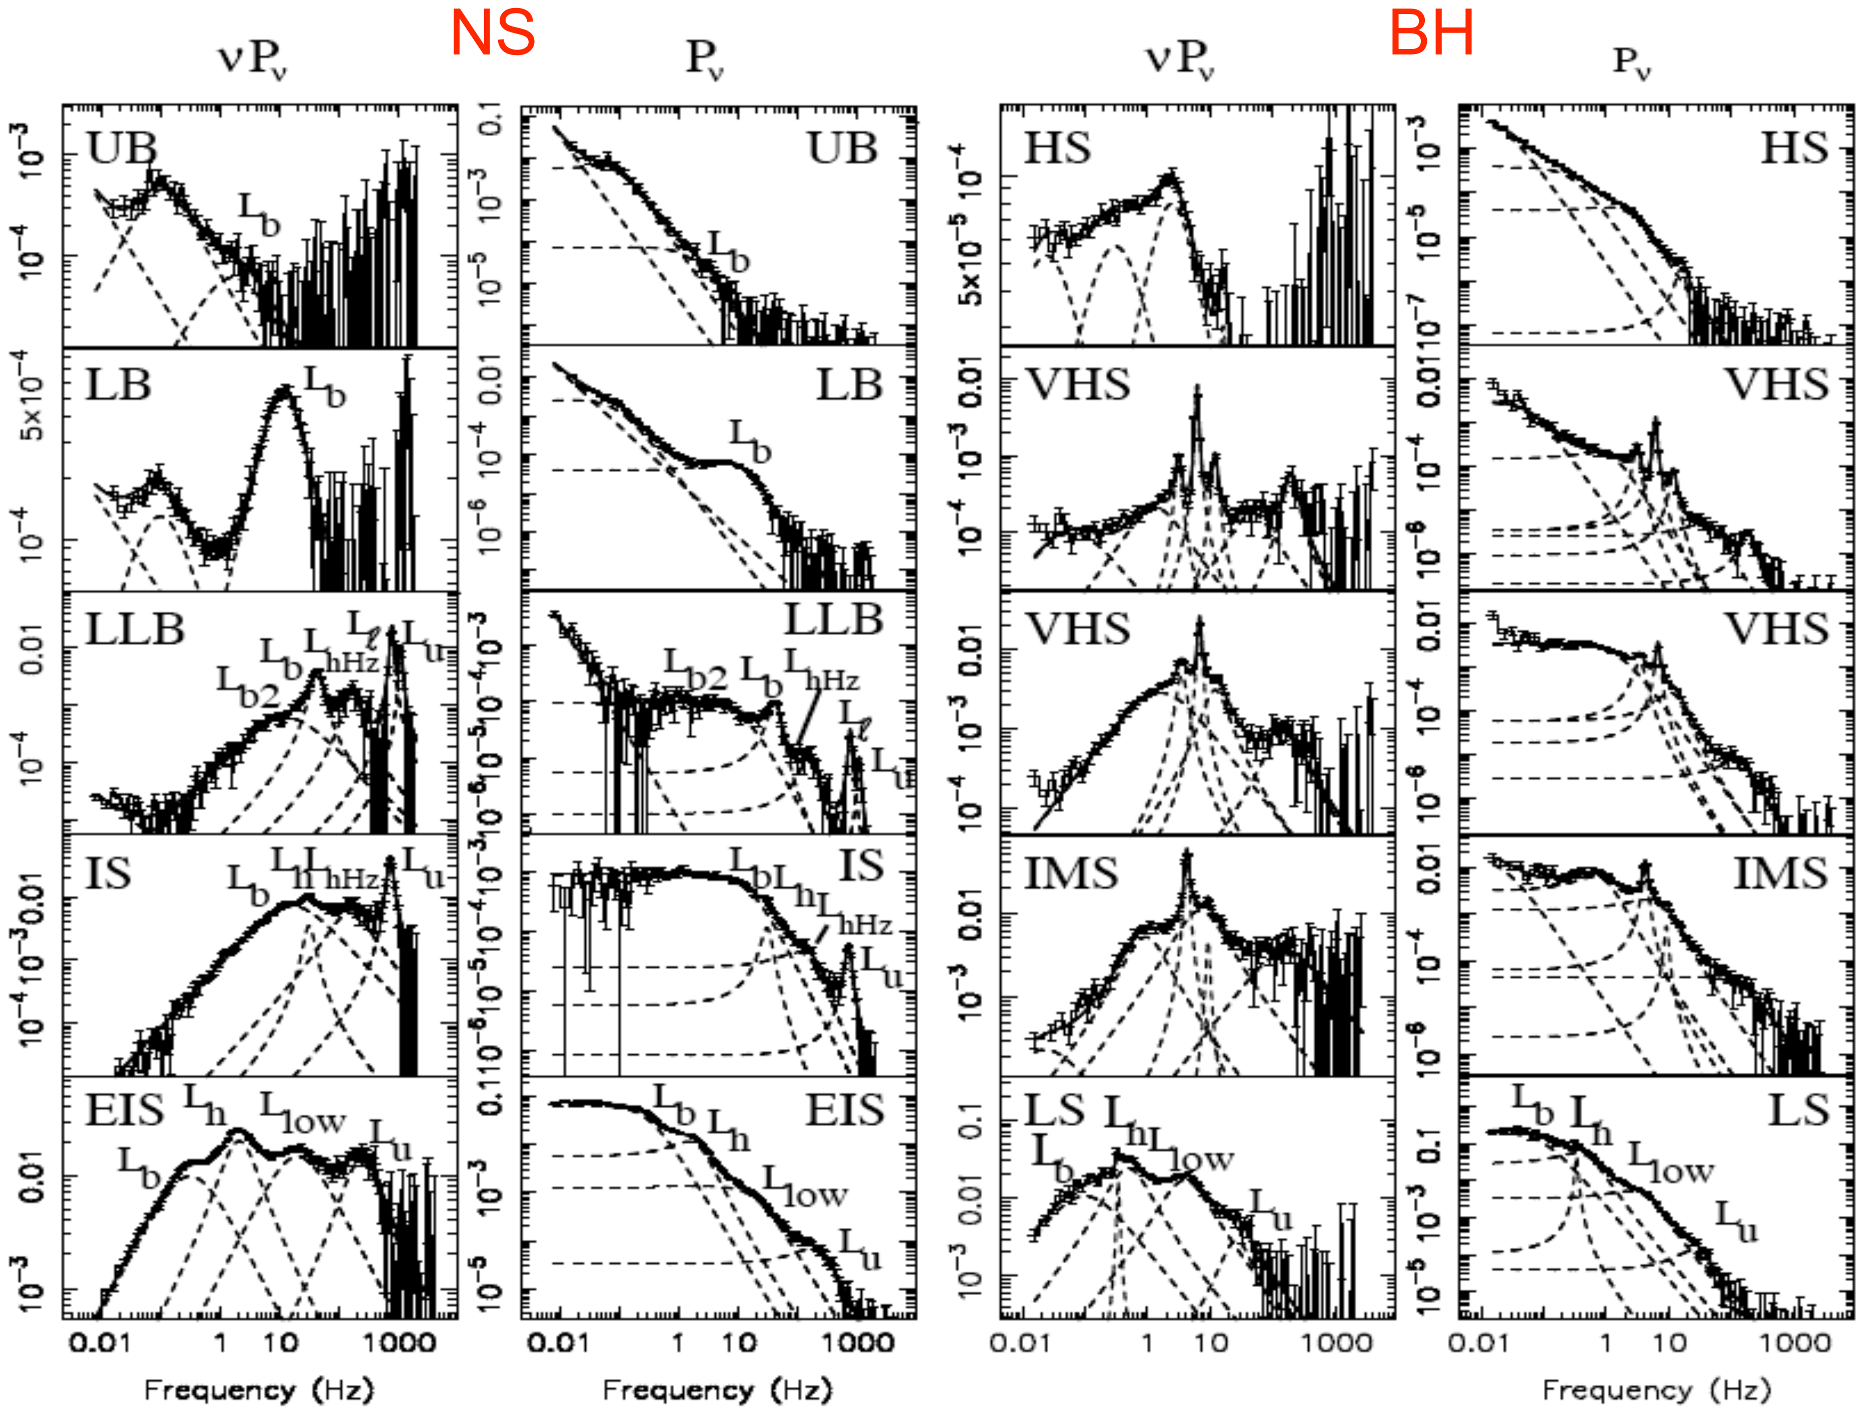

NS
BH
νPν
Pν
UB
LB
LLB
IS
EIS
HS
VHS
IMS
LS
Frequency (Hz)

POWER SPECTRUM – LEAHY NORMALIZATION

Parseval's theorem: $\sum_k x_k^2 = \frac{1}{N} \sum_j |a_j|^2$

Variance in the real time series x_k:

$$\mathrm{Var}(x_k) \equiv \sum_k (x_k - \overline{x})^2 = \sum_k x_k^2 - \frac{1}{N} \left(\sum_k x_k \right)^2 = \frac{1}{N} \sum_j |a_j|^2 - \frac{1}{N} a_0^2$$
$$= \frac{1}{N} \sum_{j \neq 0} |a_j|^2$$

Leahy normalized power spectrum (choice of normalization to be addressed):
$P_j \equiv \frac{2}{N_{ph}} |a_j|^2$; $\quad j = 0, \ldots, \frac{N}{2}$; $\quad$ where $N_{ph} = \sum_k x_k = a_0$

Then: $\mathrm{Var}(x_k) = \frac{N_{ph}}{N} \left(\sum_{j=1}^{N/2-1} P_j + \frac{1}{2} P_{N/2} \right)$: variance is sum of powers.

As a_j has the same dimension as x_k, the dimension of $P_j \propto |a_j|^2/a_0$ is also the same as x_k: $[P_j] = [a_j] = [x_k]$.

POWER DENSITY SPECTRUM

Power density gives power per unit of frequency (i.e., per Hz), so that integral over power density spectrum is sum of powers:

$$\int_{\nu_{j1}}^{\nu_{j2}} p(\nu) d\nu = \sum_{j=j1}^{j2} P_j$$

Now $\delta\nu = 1/T$, so the Leahy normalized power density at ν_j is: $p(\nu_j) \equiv P_j/\delta\nu = TP_j$. Dimension: $[p(\nu)] = [x_k/\nu]$

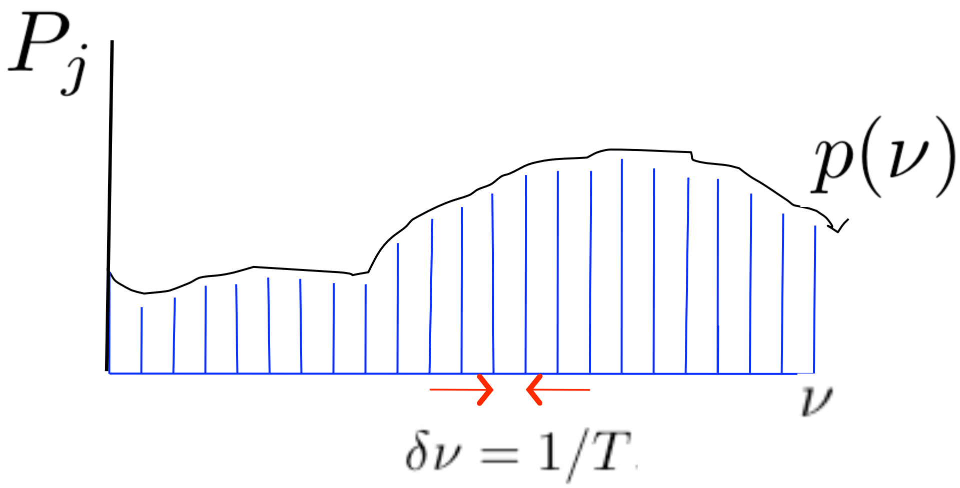

FRACTIONAL RMS AMPLITUDE

Fractional rms amplitude of a signal in a time series:

$$r \equiv \frac{\sqrt{\frac{1}{N}Var(x_k)}}{\overline{x}} = \frac{N}{N_{ph}}\sqrt{\frac{N_{ph}}{N^2}\left(\sum_{j=1}^{N/2-1} P_j + \frac{1}{2}P_{N/2}\right)} = \sqrt{\frac{1}{N_{ph}}\left(\sum_{j=1}^{N/2-1} P_j + \frac{1}{2}P_{N/2}\right)}$$

r is dimensionless and often expressed in %.

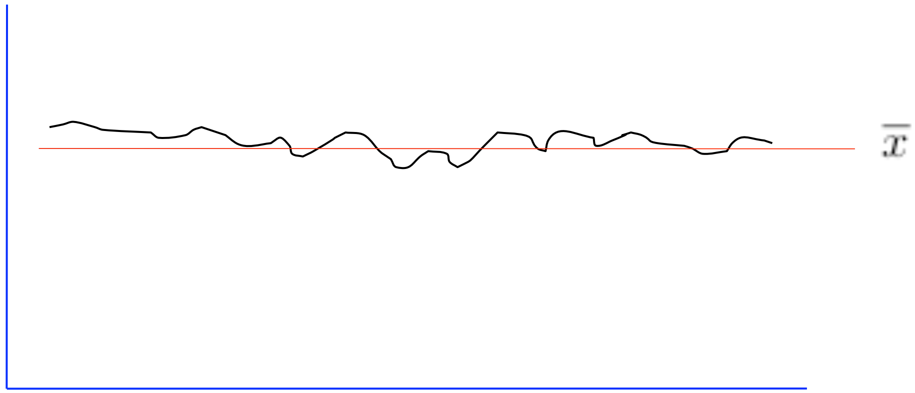

"Rms normalized" power density: $q(\nu_j) \equiv TP_j/N_{ph} = p_j/N_{ph}$

$q(\nu)$ has the nice property that fractional rms is just $r = \sqrt{\int q(\nu)d\nu}$.
Dimension of $q(\nu)$ is $[q] = [1/\nu] = [t]$; physical unit of $q(\nu)$ is $(\text{rms/mean})^2/\text{Hz}$.

"SOURCE" FRACTIONAL RMS AMPLITUDE

If the x_k are the sum of source and background: $x_k = b_k + s_k$, then the rms amplitude as a fraction of just the s_k:

$r_s = r \cdot \frac{B+S}{S}$, where B and S are sums of the b_k and s_k, so $B+S = \sum x_k = N_{ph}$

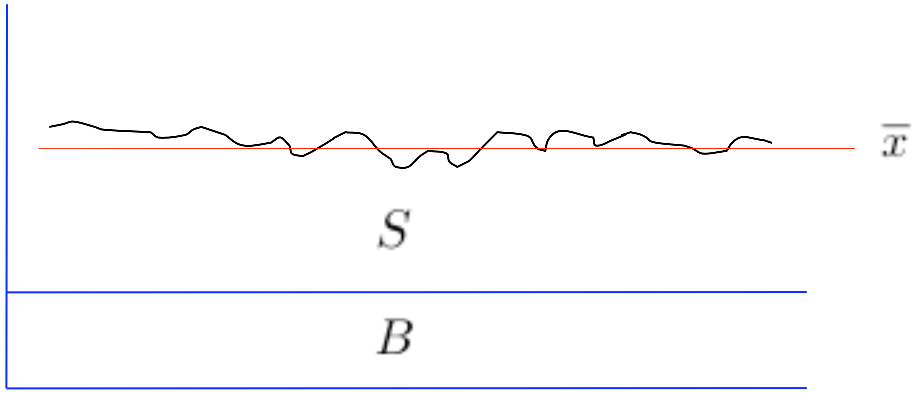

"Source rms normalized" power density: $q_s \equiv q \cdot \left(\frac{B+S}{S}\right)^2 = TP_j \cdot \frac{B+S}{S^2}$

Now $r_s = \sqrt{\int q_s(\nu) d\nu}$; q_s has the same unit as q: $(\text{rms/mean})^2/\text{Hz}$.

Detecting 'something' in a power spectrum

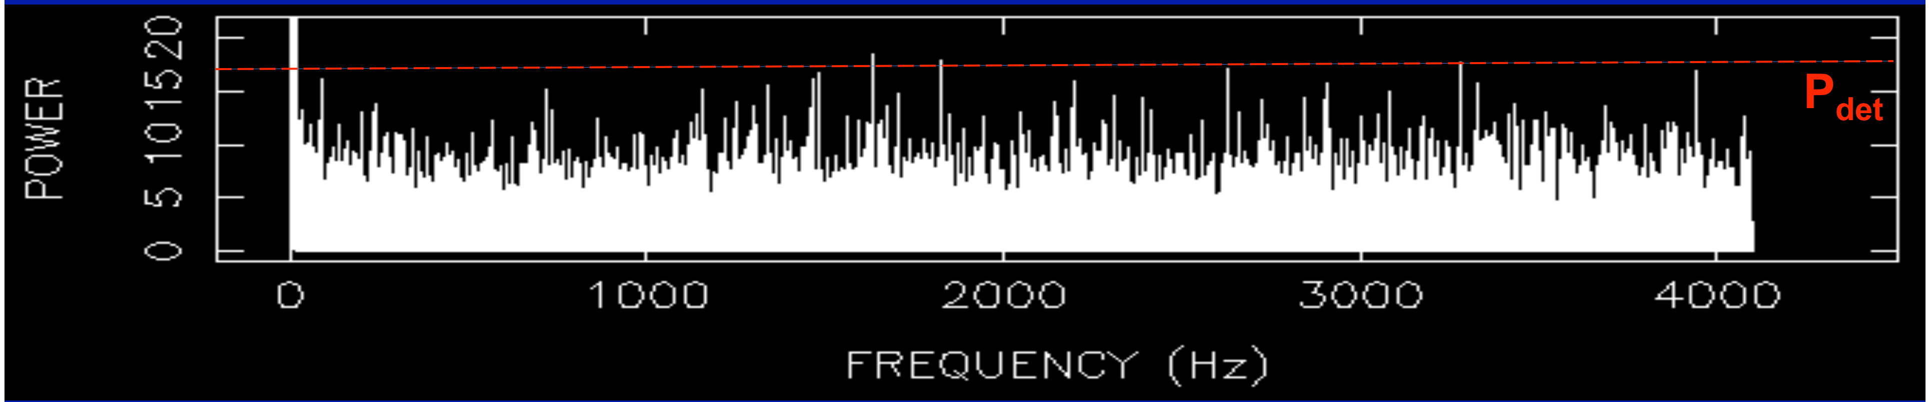

How big must a power be to constitute a **significant excess** over the noise?
The (1-ε) confidence **detection level** P_{det} is a level that has a **false alarm probability** of ε. If there is just noise, prob($P_j > P_{det}$) = ε.
Take ε small, e.g., ε=1% for 99% confidence.
If $P_j > P_{det}$ then with 99% confidence there is something else than just noise: a source signal.

1. **Detection**: to determine P_{det} you just need to know the noise power distribution
2. **Quantifying signal strength**: to also determine how strong the detected signal power is, or to set an upper limit, you also need to know what is the interaction between noise and signal powers.

NOISE POWER DISTRIBUTION

Noise powers follow a chi-squared distribution with 2 degrees of freedom (dof).

This can be seen as follows:
$P_j \propto A_j^2 + B_j^2$, where $A_j = \sum_k x_k \cos \omega_j t_k$, and $B_j = \sum_k x_k \sin \omega_j t_k$, $k = 0, \ldots, N-1$

So, each A_j and each B_j is a linear combination of the x_k. Hence if the x_k are normally distributed then the A_j and B_j are as well $\rightarrow$ P_j is $\propto \chi^2$ with 2 dof by definition.

If the x_k follow some other distribution (e.g. Poisson) then the central limit theorem ensures that A_j and B_j are still approximately normal (for large N) $\rightarrow$ the P_j are still approximately χ^2 with 2 dof.

Exact expressions depend on the normalization of the P_j.

LEAHY POWER DISTRIBUTION

The Leahy normalization is chosen such that if the x_k are Poisson distributed, then the P_j exactly follow the chi-squared distribution with 2 dof, χ^2_2.
This is actually an exponential distribution:

$$prob(P_j > P) = Q(P|2) = e^{-P/2}$$

Properties of this distribution:
mean $\langle P_j \rangle = 2$; standard deviation $\sigma_{P_j} = 2$; P_j uncorrelated

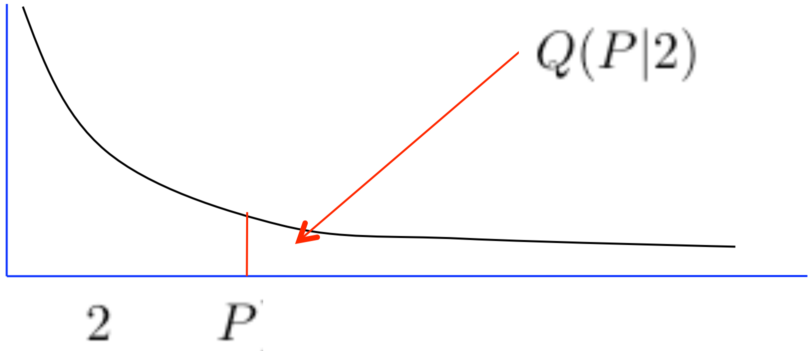

So, the power spectrum is **very noisy**. This does not improve with:
- longer observation — you just get more powers
- broader time bins — you just get a lower ν_{Ny}

  Solution: **smooth** the power spectrum.

Smoothing methods

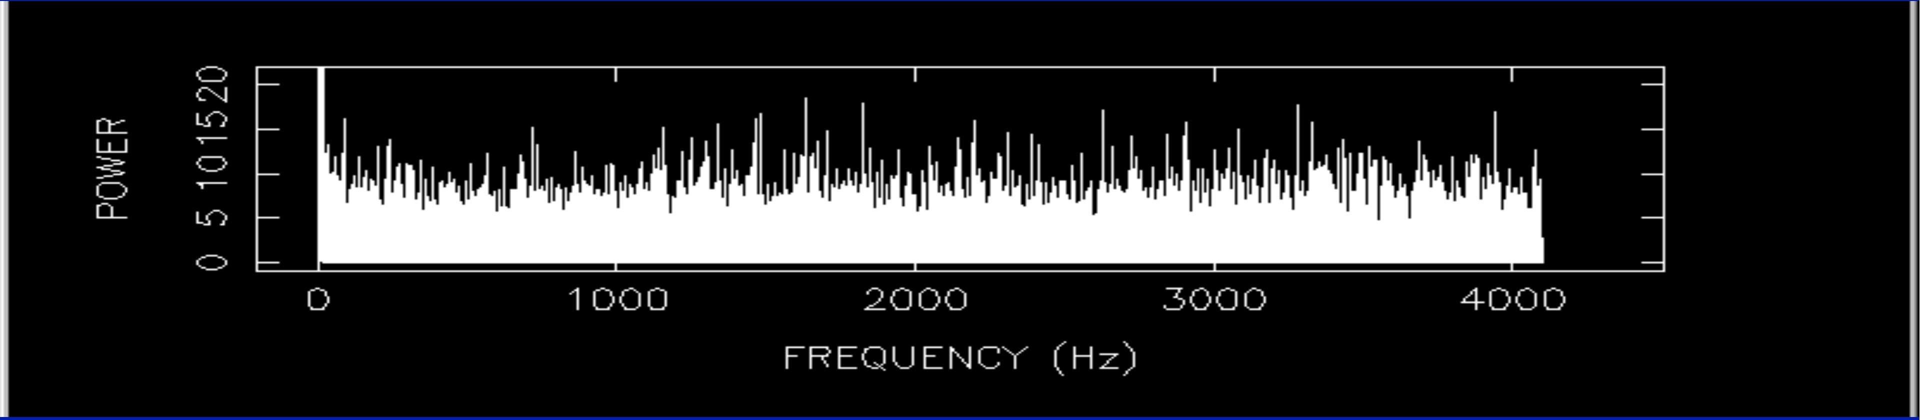

- Average several power spectra of subsegments of the time series
- Average adjacent bins in a power spectrum

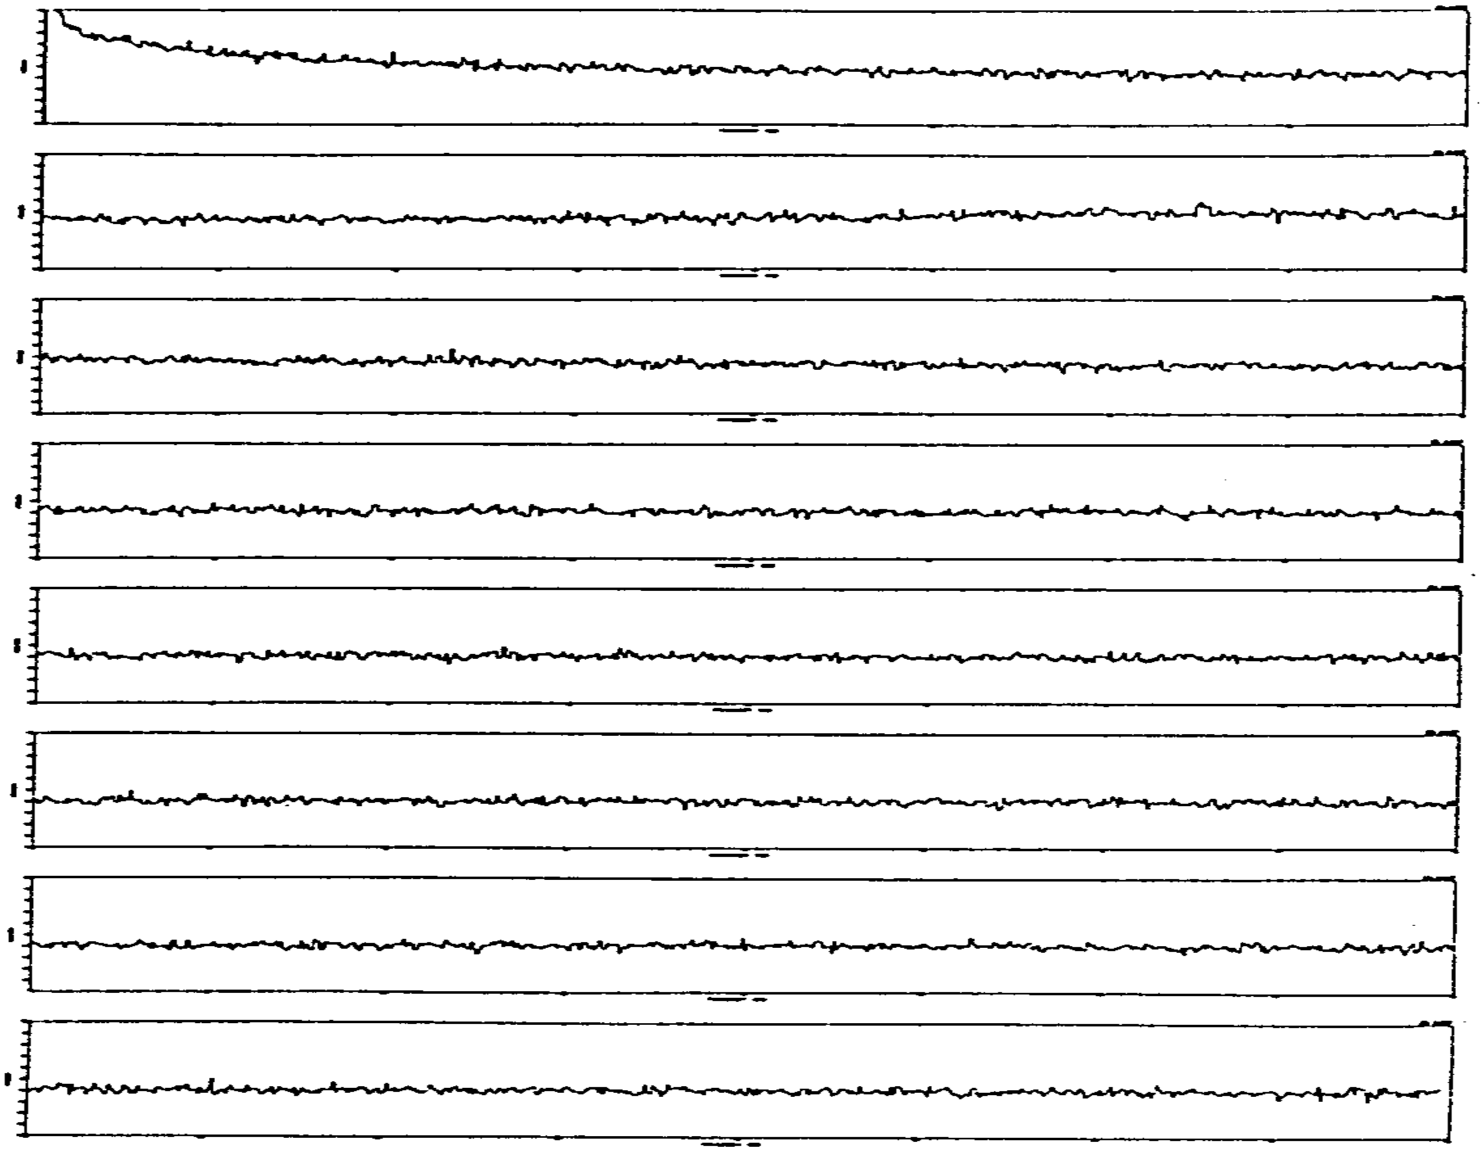

Example: average of thousands of power spectra of GX 5-1

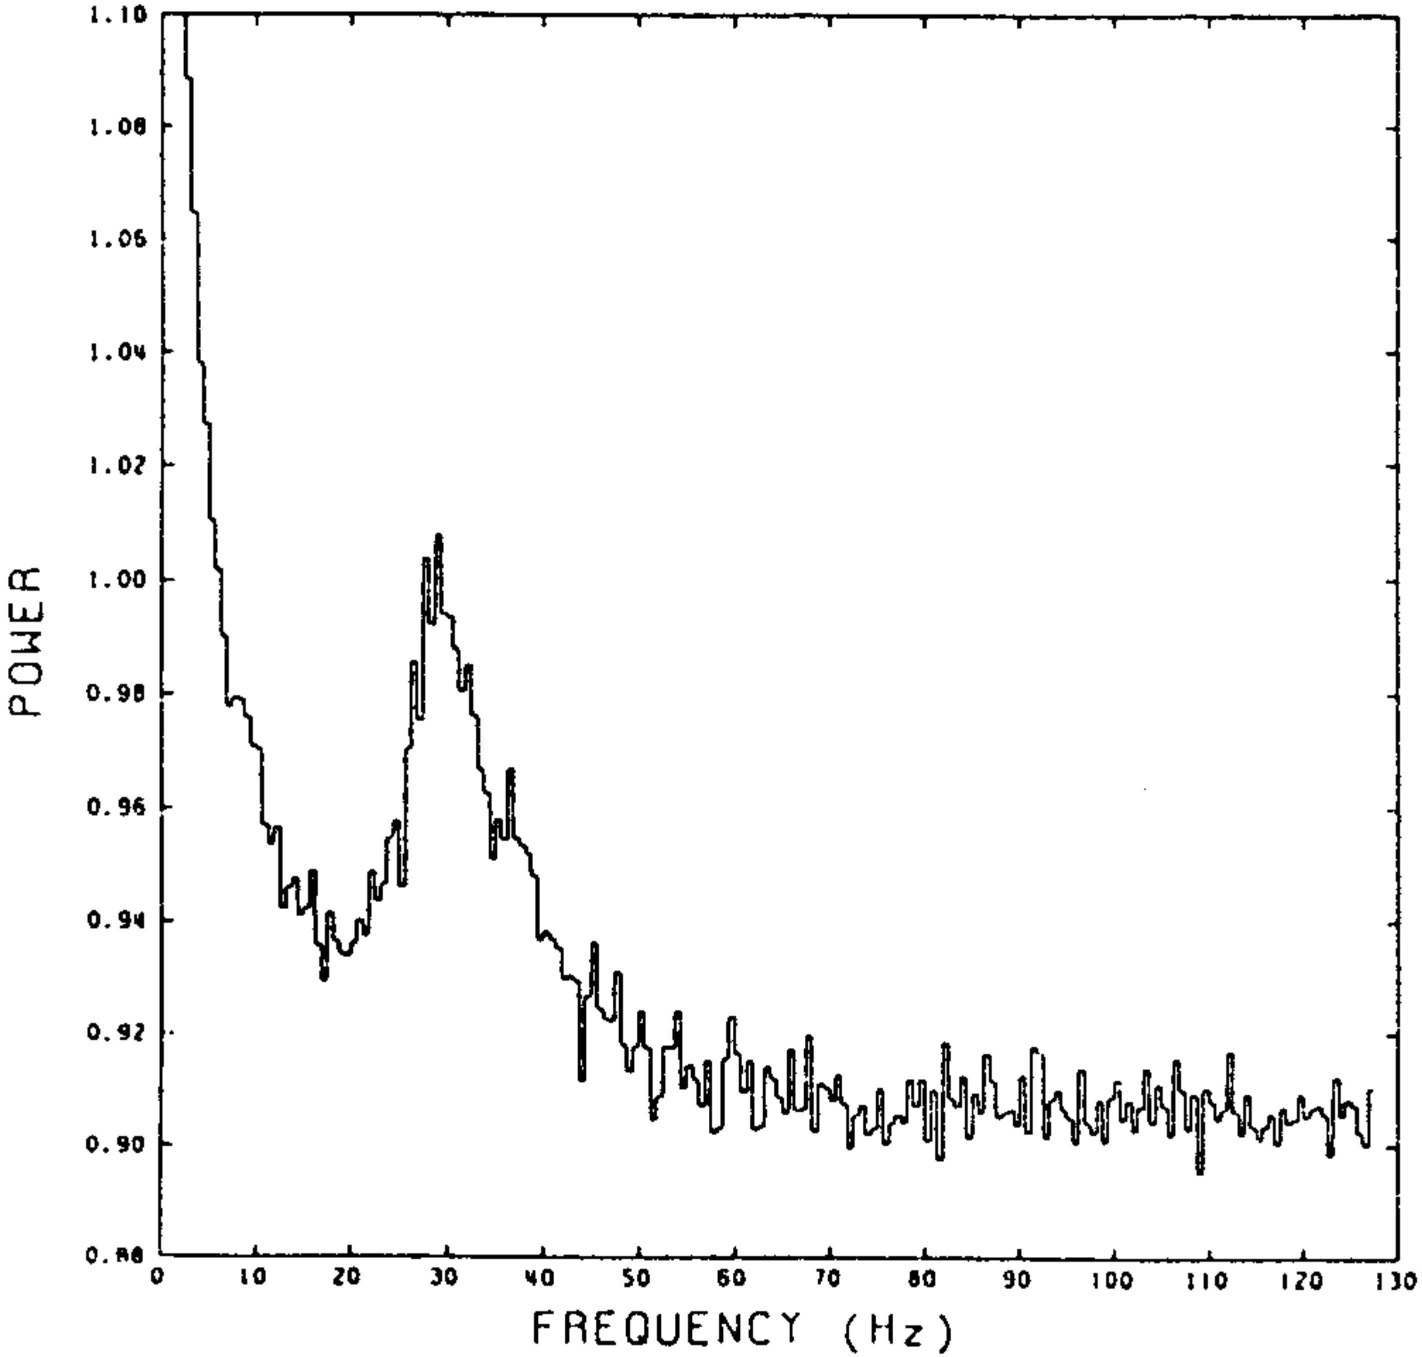
POWER
1.10
1.08
1.06
1.04
1.02
1.00
0.98
0.96
0.94
0.92
0.90
0.88
0
10
20
30
40
50
60
70
80
90
100
110
120
130
FREQUENCY (Hz)

Average of M power spectra

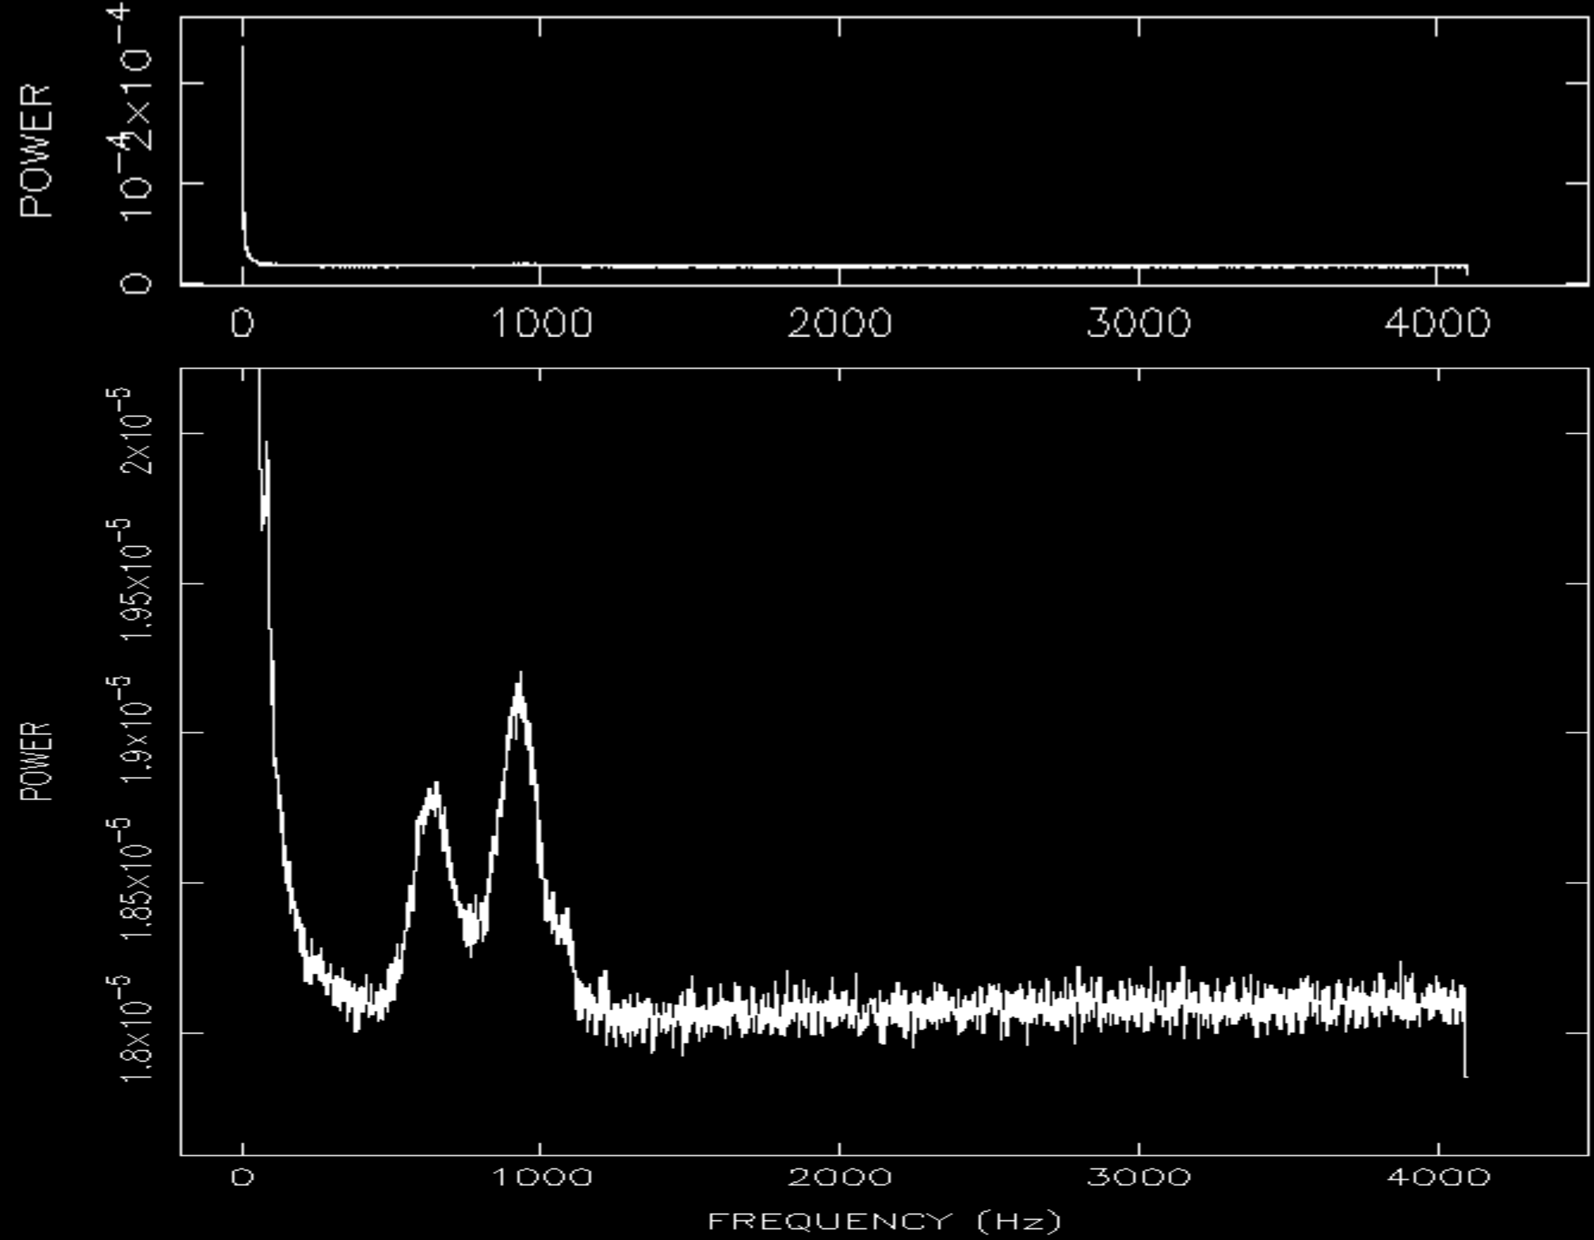

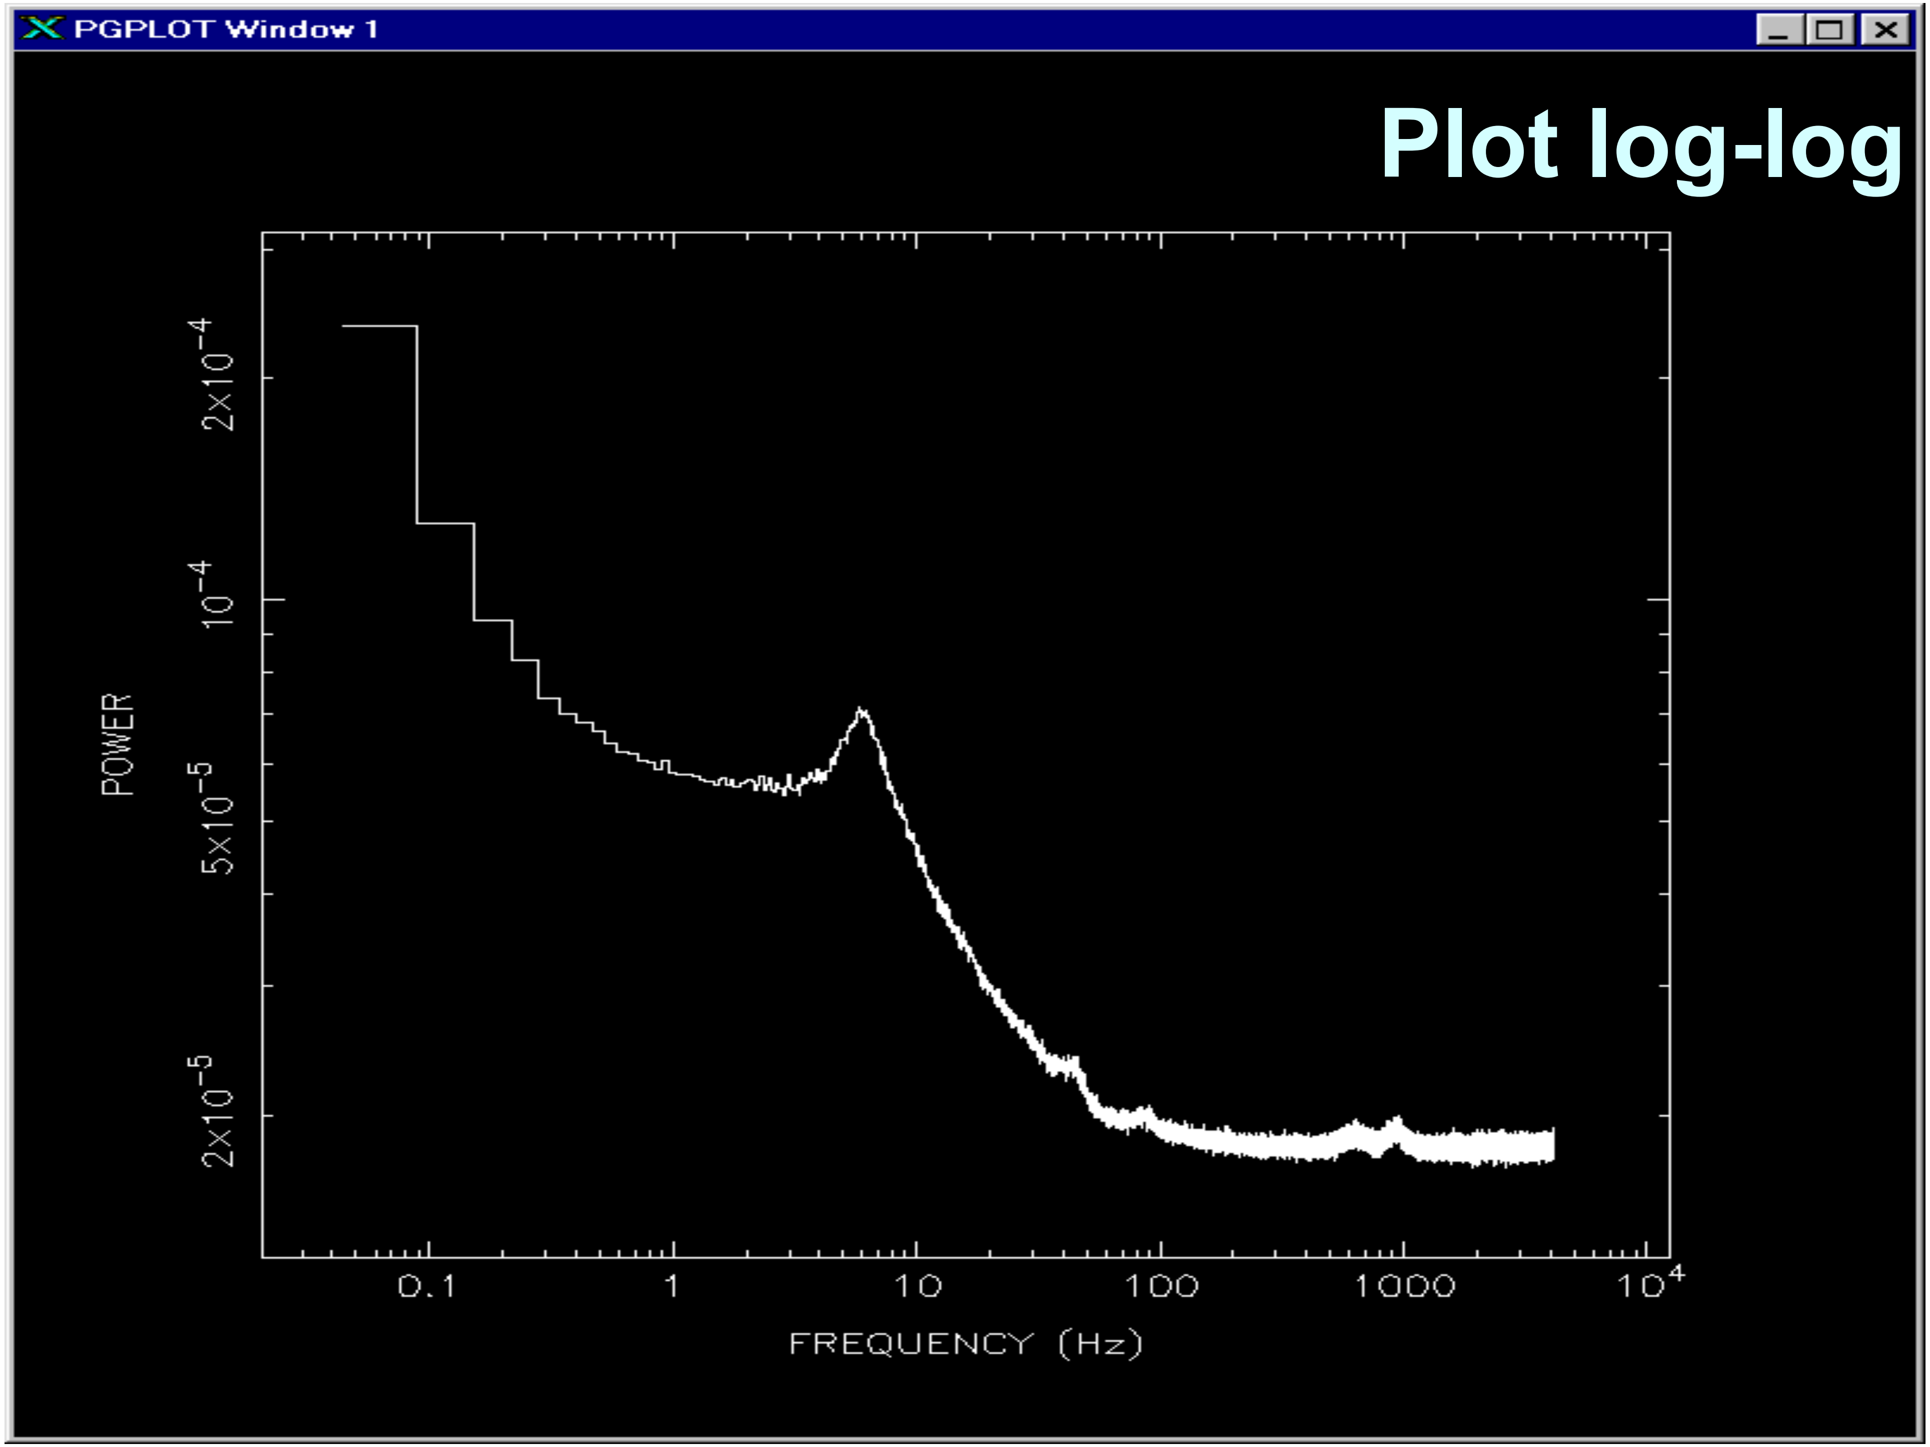
PGPLOT Window 1
Plot log-log
POWER
2×10^{-4}
10^{-4}
5×10^{-5}
2×10^{-5}
0.1
1
10
100
1000
10^4
FREQUENCY (Hz)

Logarithmic rebin

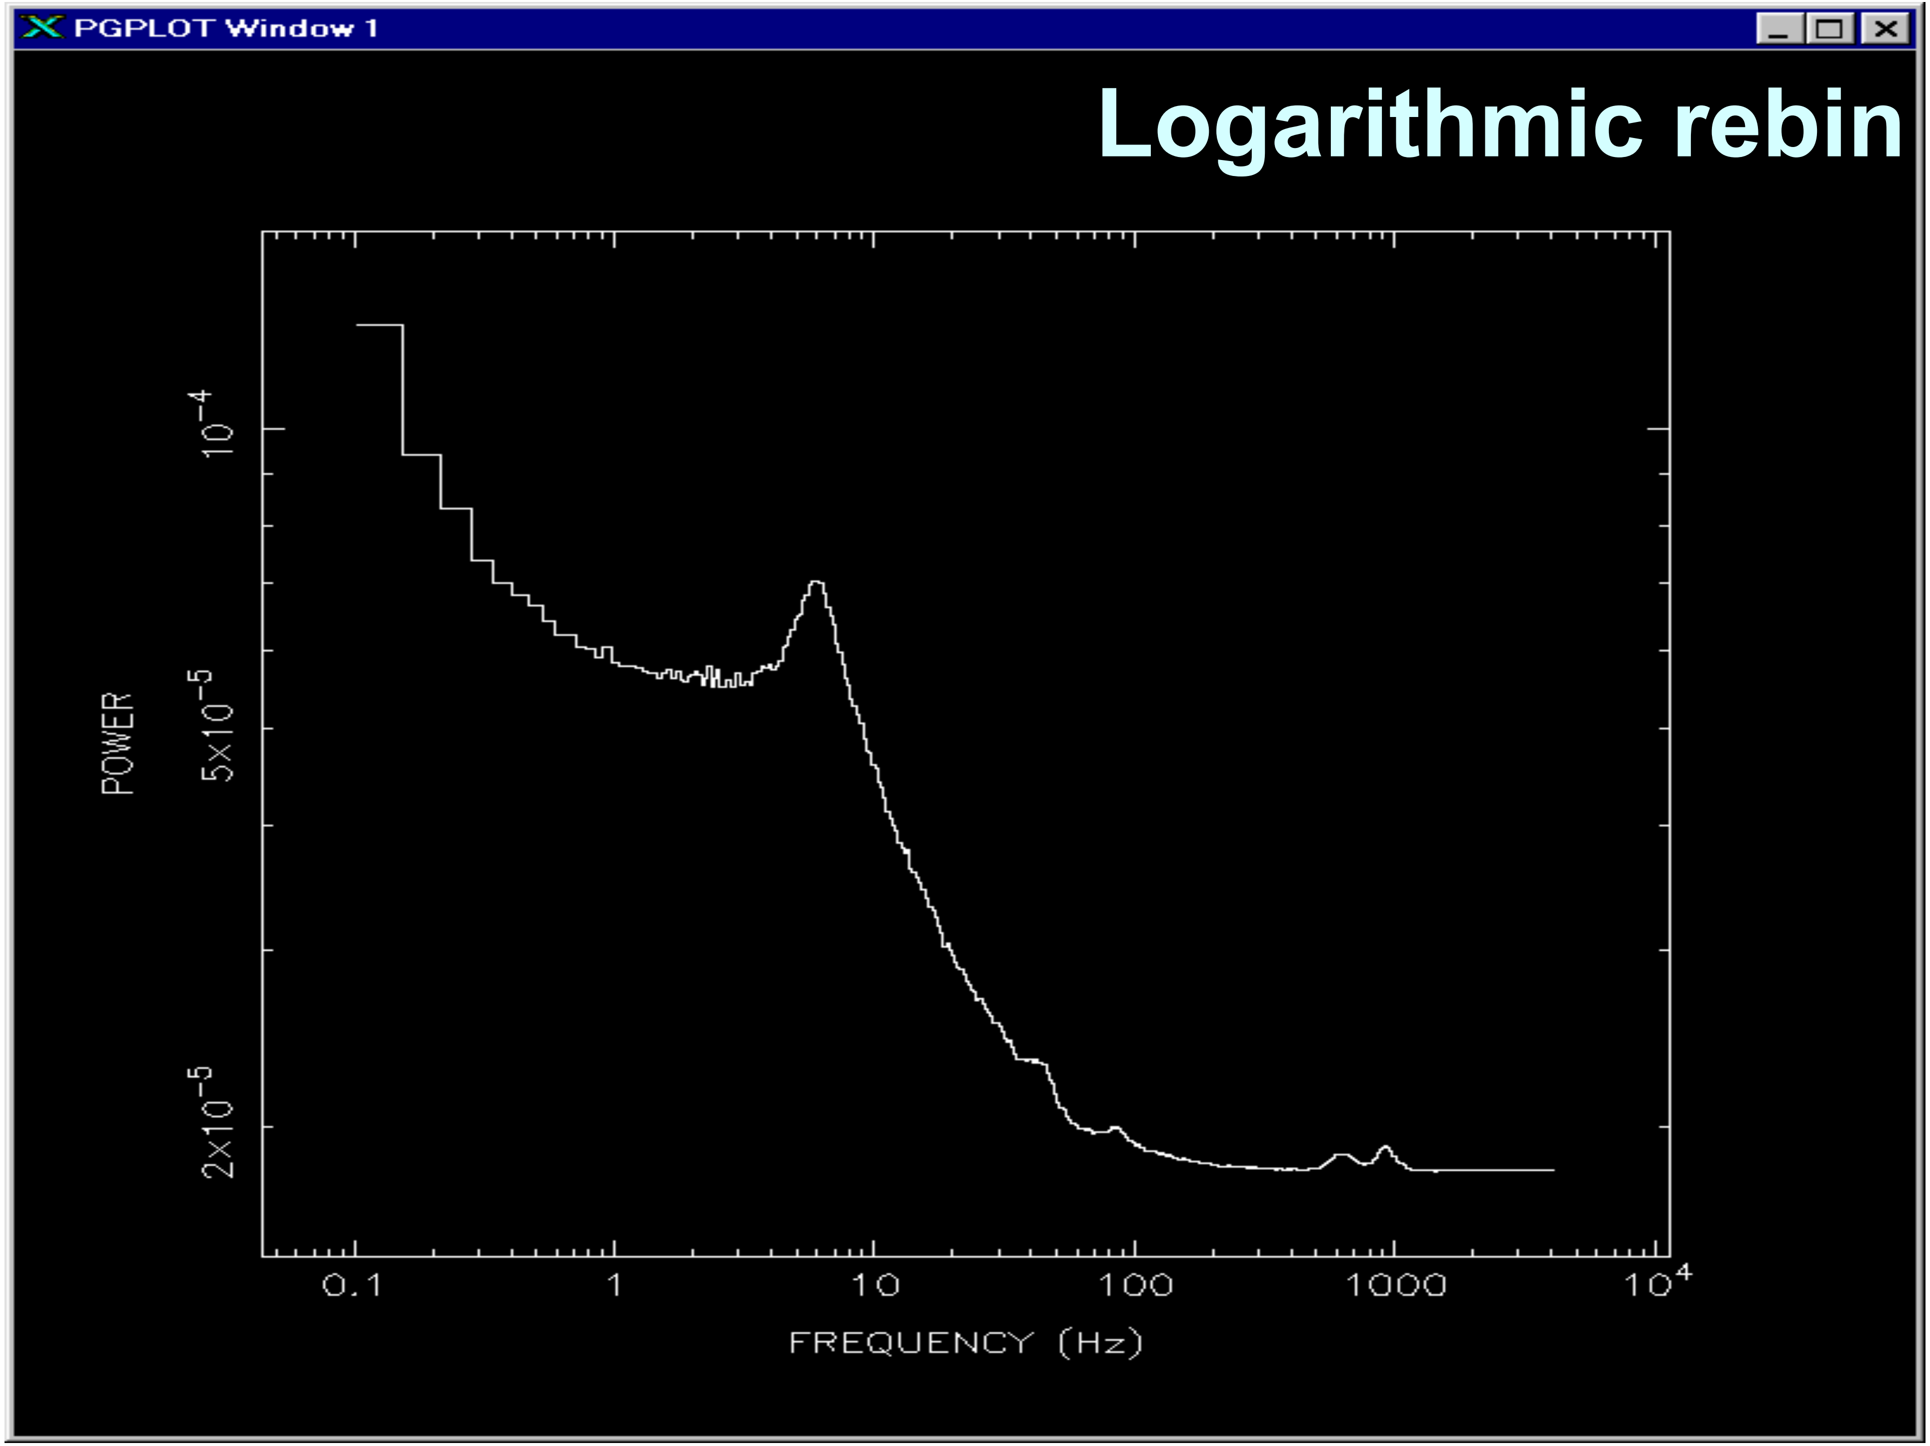

Subtract Poisson (counting) noise

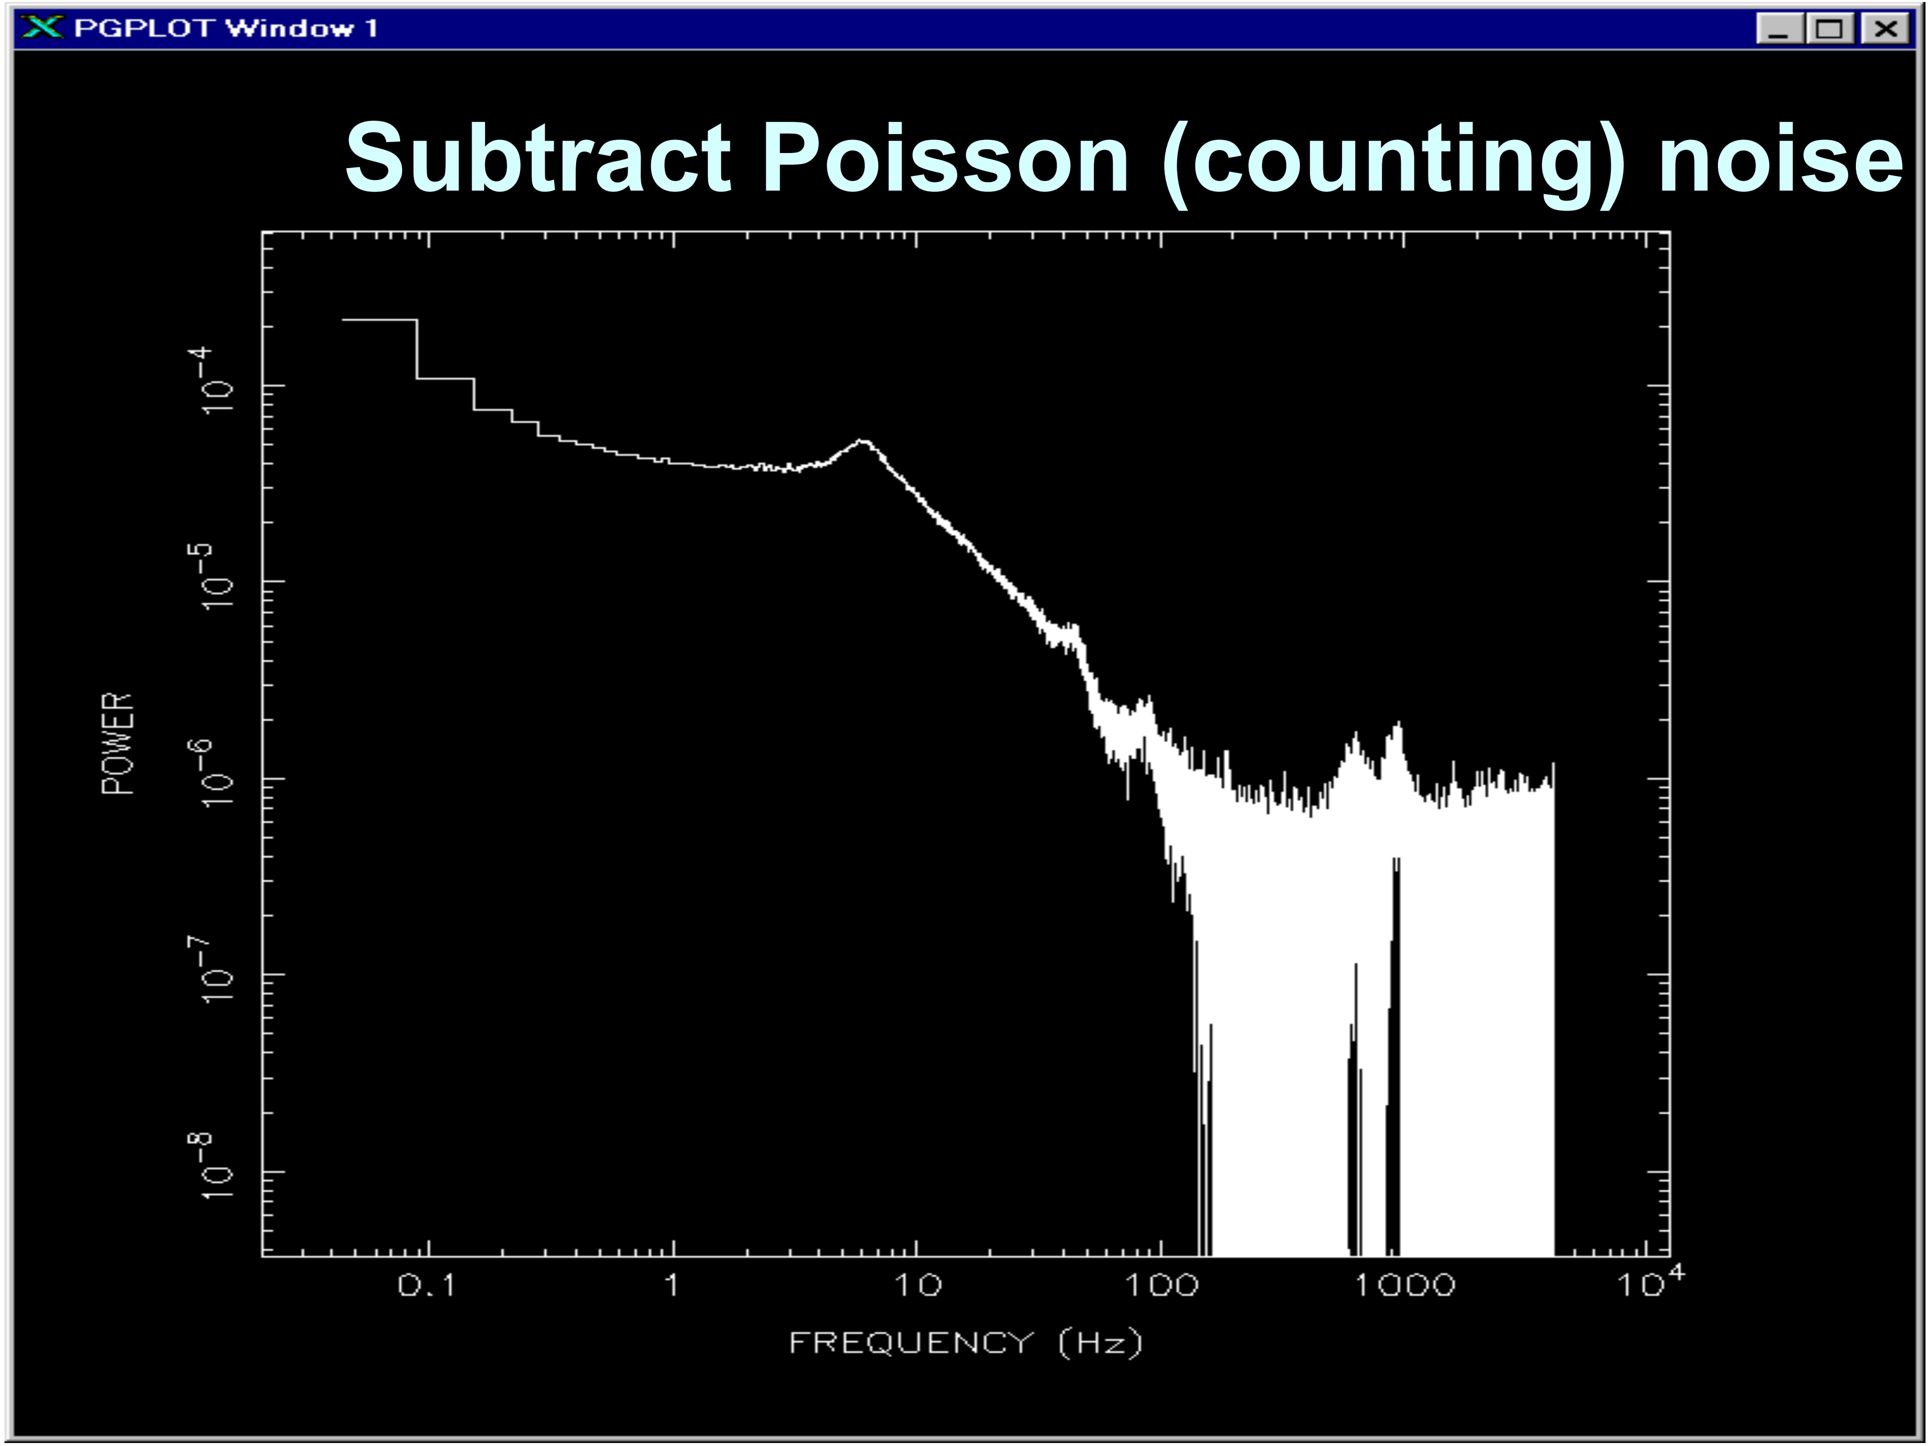

Logarithmic rebin

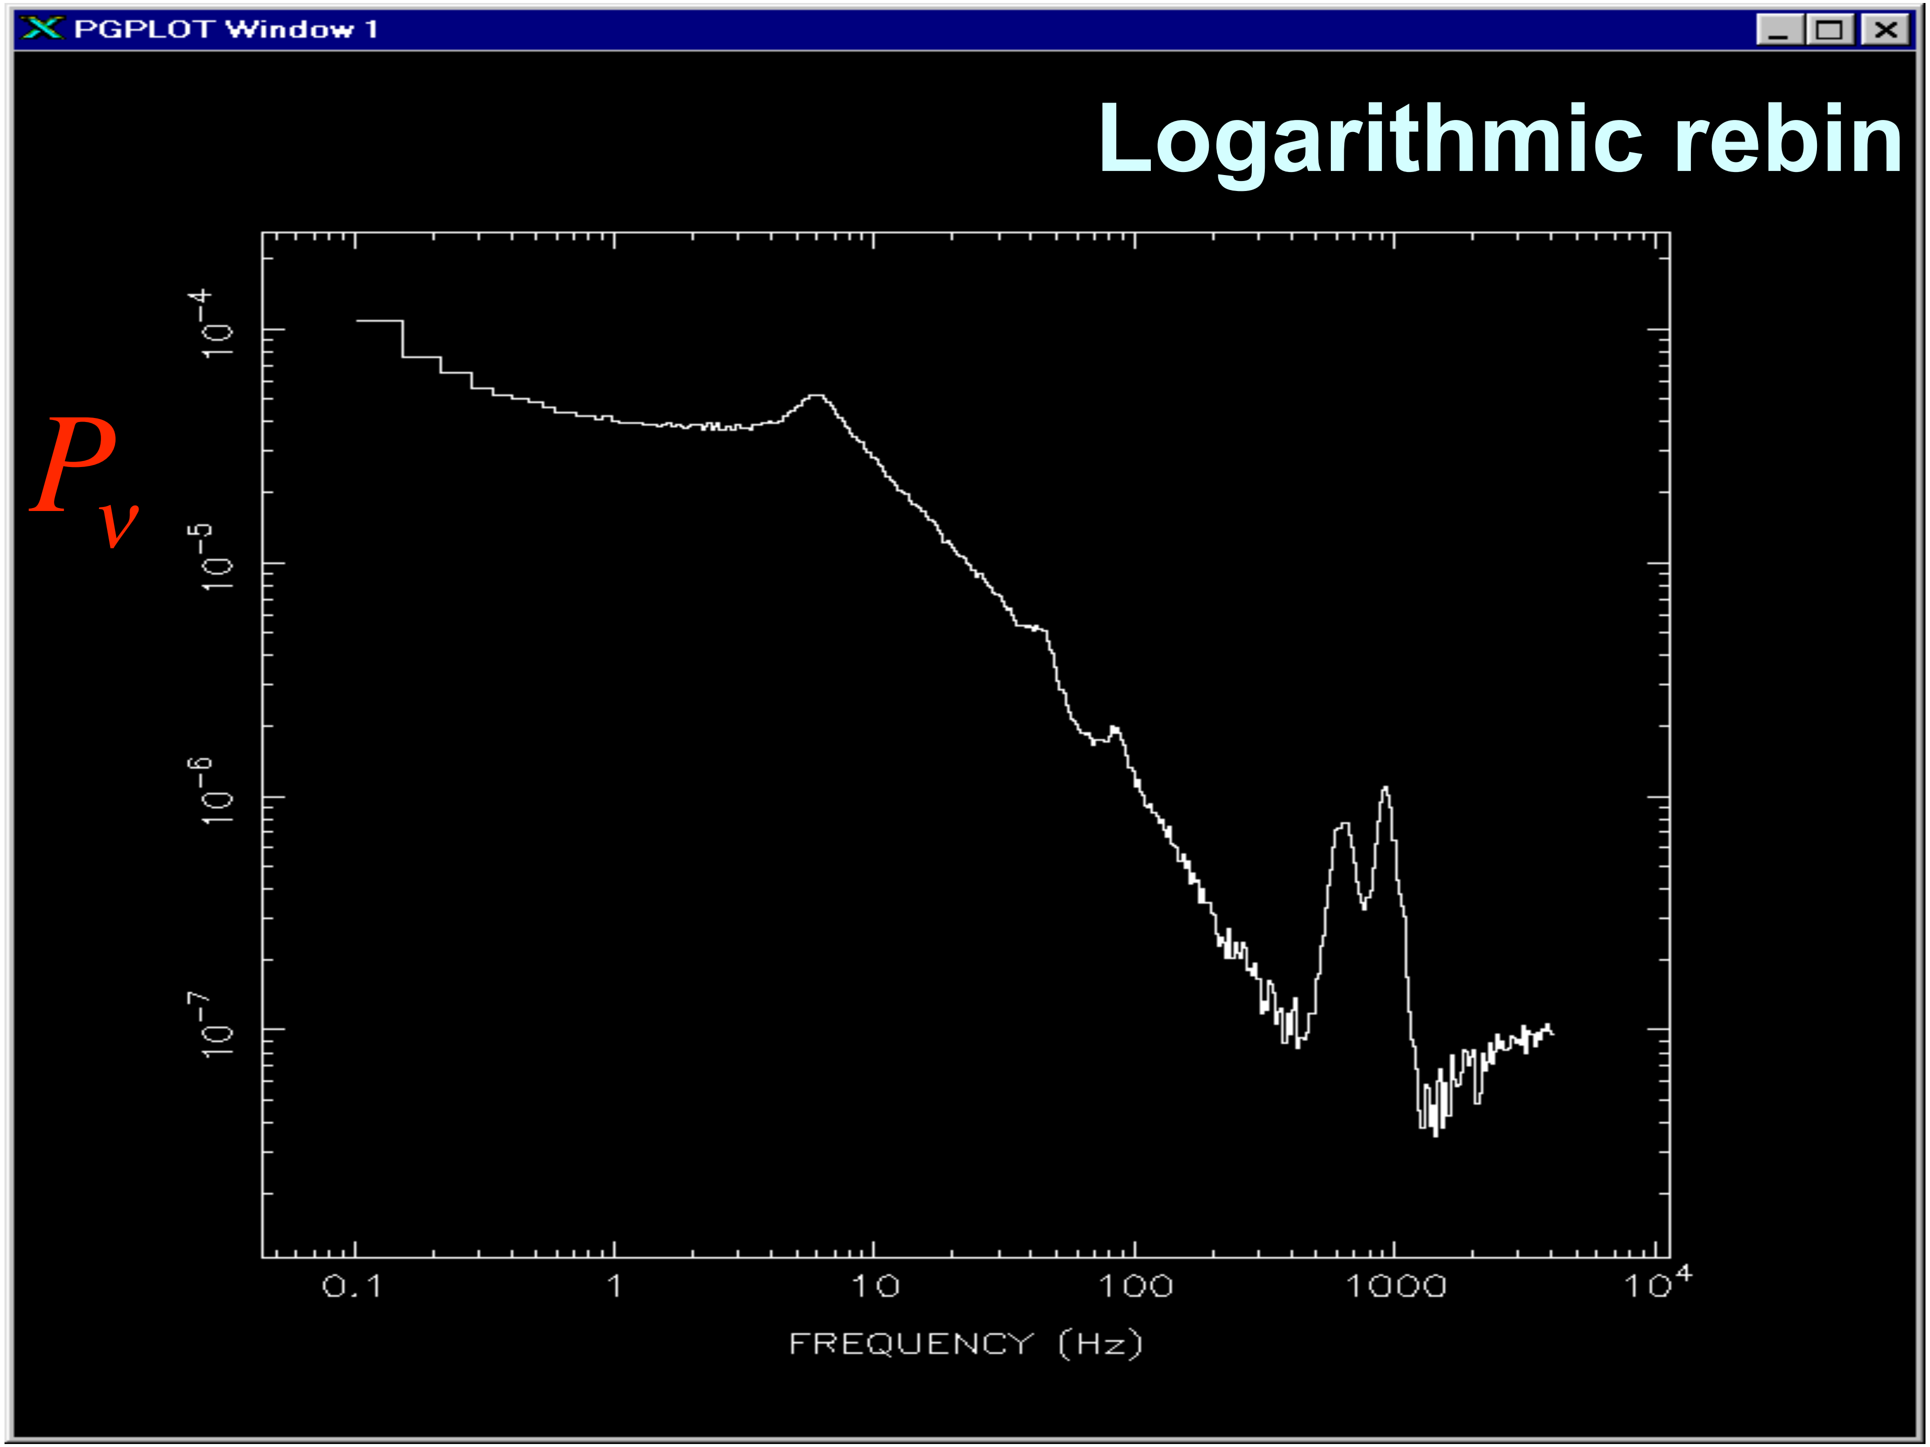

Multiply power with Fourier frequency

νP_ν

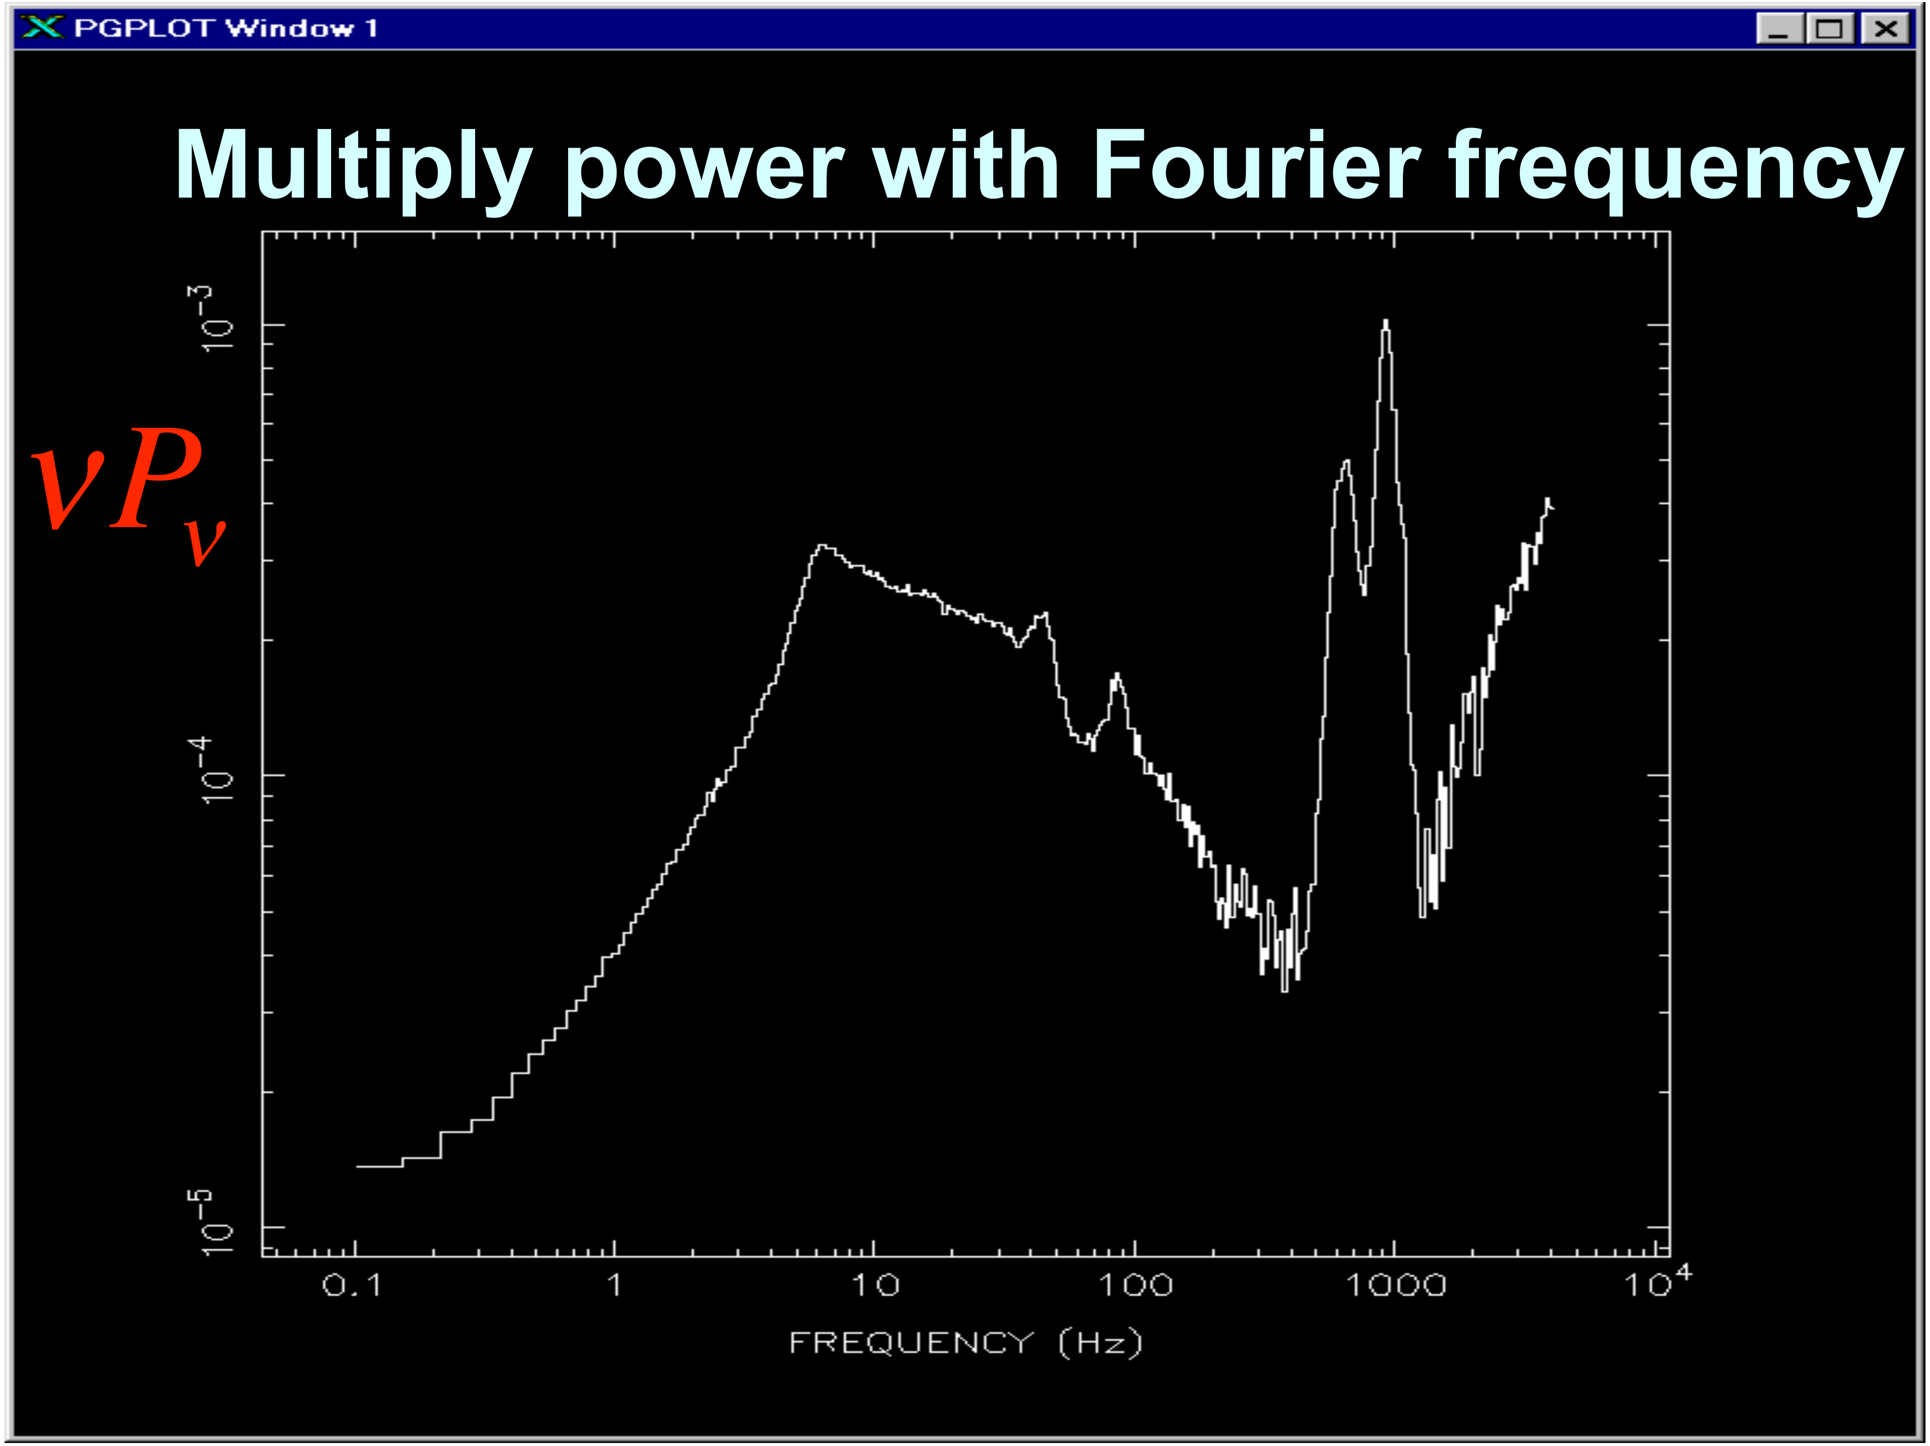

AVERAGED POWER DISTRIBUTION

Individual P_j follow ${\chi_2}^2$, the chi-squared distribution with 2 dof.

What is the distribution of the average of M powers $\frac{1}{M}\sum_M P_j \equiv \overline{P_M}$?

$\boxed{\text{Additive property of } \chi^2 \text{ distribution: sum of } M \text{ powers is distributed as } \chi^2_{2M}.}$

So $\overline{P_M}$ is distributed as χ^2_{2M}/M, and hence the probability for $\overline{P_M}$ to exceed some threshold P is:

$$prob(\overline{P_M} > P) = Q(MP|2M)$$

Properties of this distribution: average = 2; standard deviation = $2/\sqrt{M}$, as:

$$\langle \sum_M P_j \rangle = 2M \Longrightarrow \langle \overline{P_M} \rangle = 2 \text{ and } \sigma^2_{\sum P_j} = 4M \Longrightarrow \sigma_{\overline{P_M}} = \frac{2\sqrt{M}}{M} = \frac{2}{\sqrt{M}}$$

$\boxed{\text{Central limit theorem: for large } M \text{ the distribution of } \overline{P_M} \text{ tends to normal (Gaussian), with mean 2 and standard deviation } 2/\sqrt{M}.}$

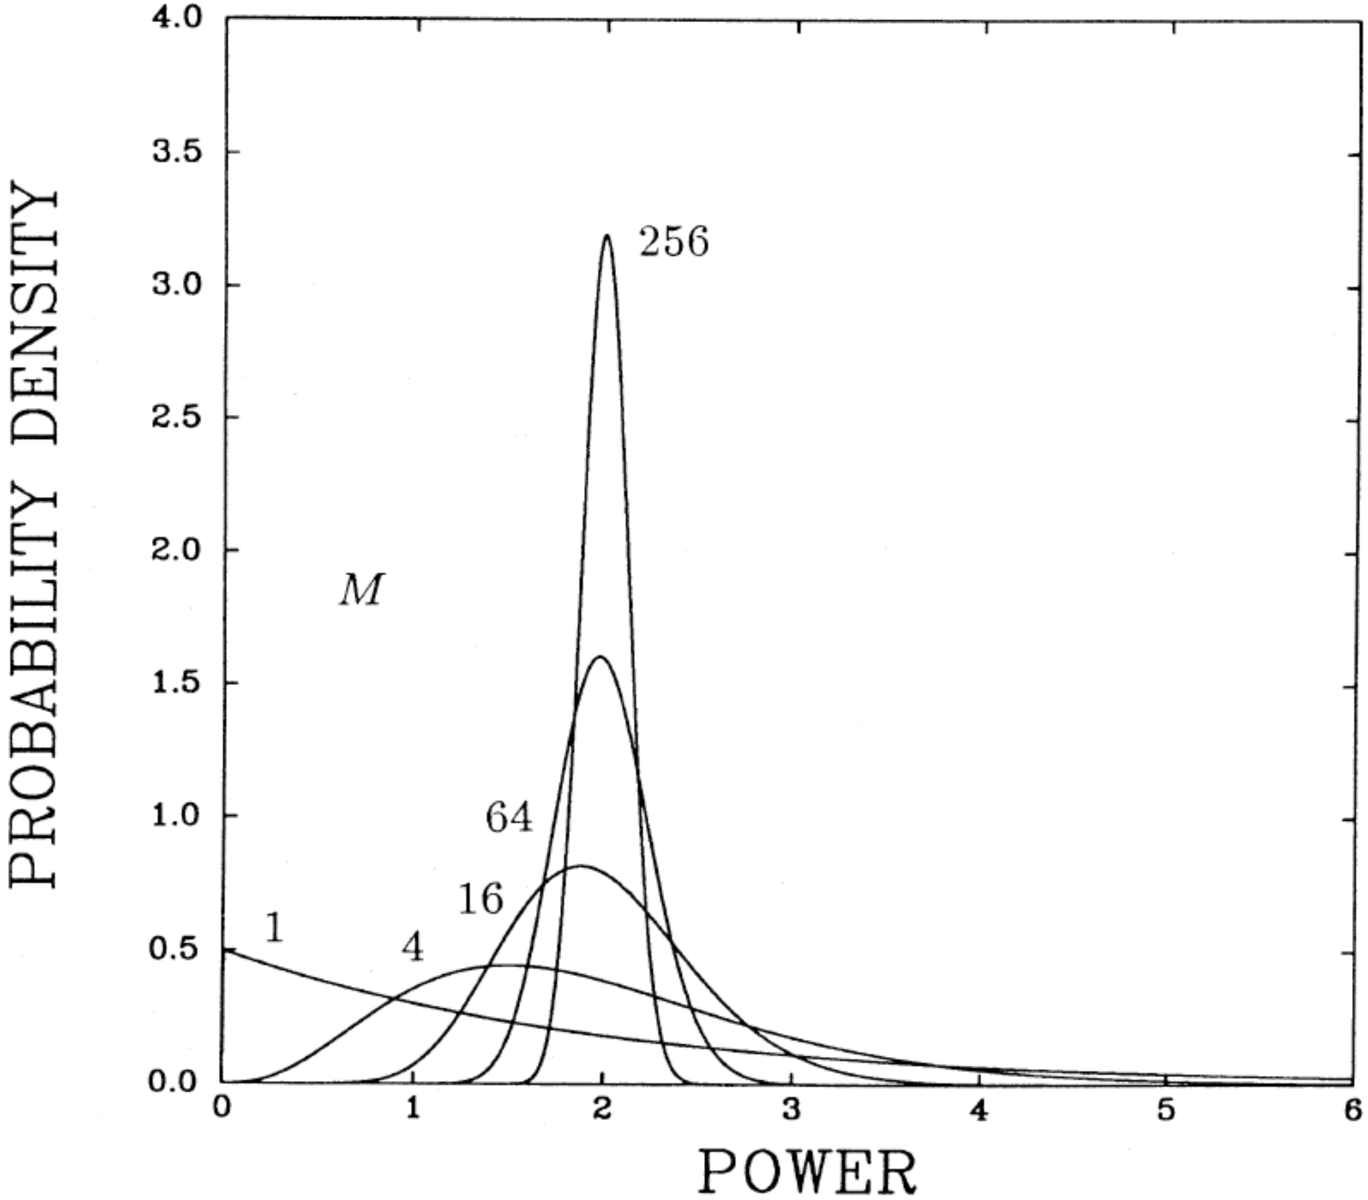

PROBABILITY DENSITY
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
256
M
64
16
1
4
0
1
2
3
4
5
6
POWER

Detection level - the number of trials !

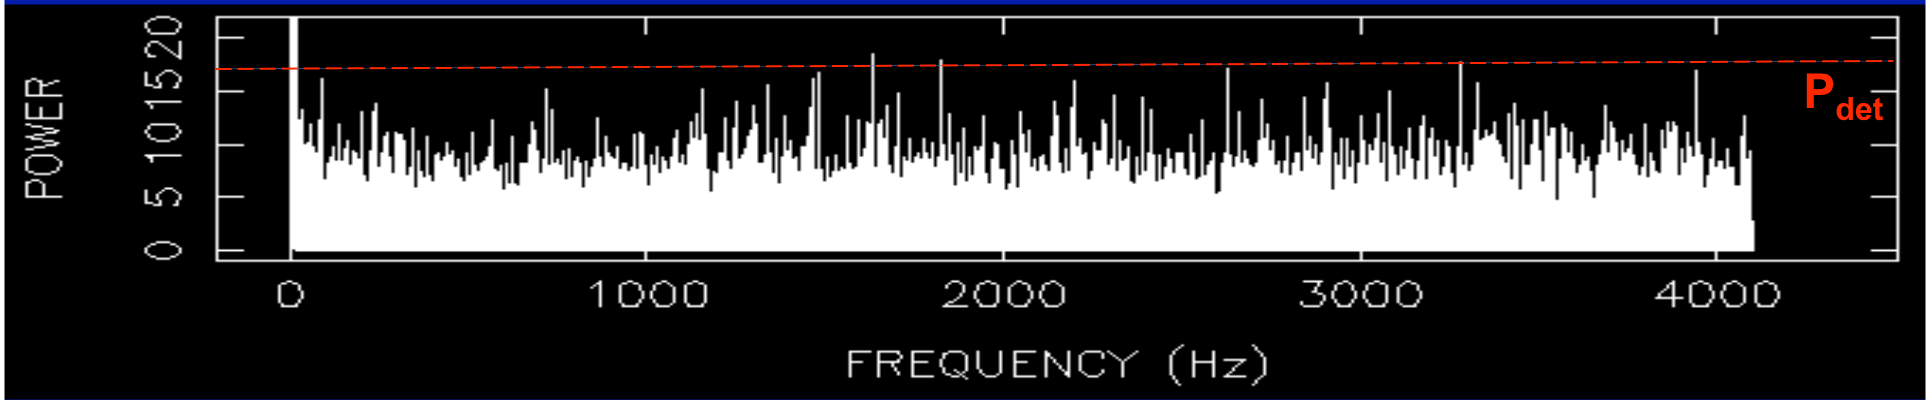

The (1-ε) confidence **detection level** P_{det} is a level that has a small **false alarm probability** of ε. The probability to exceed P_{det} by noise should be ε for **all** powers in the frequency range of interest **together !**

If you consider N_{trial} values P_j, then the probability per trial should be much smaller than ε, namely about ε/N_{trial}.

Hence detection level is given by $\boxed{\varepsilon \,/\, N_{trial} = Q(MP_{det}|2M)}$

It depends on – desired confidence level 1-ε

– number of powers averaged M

– number of trials N_{trial} !!!

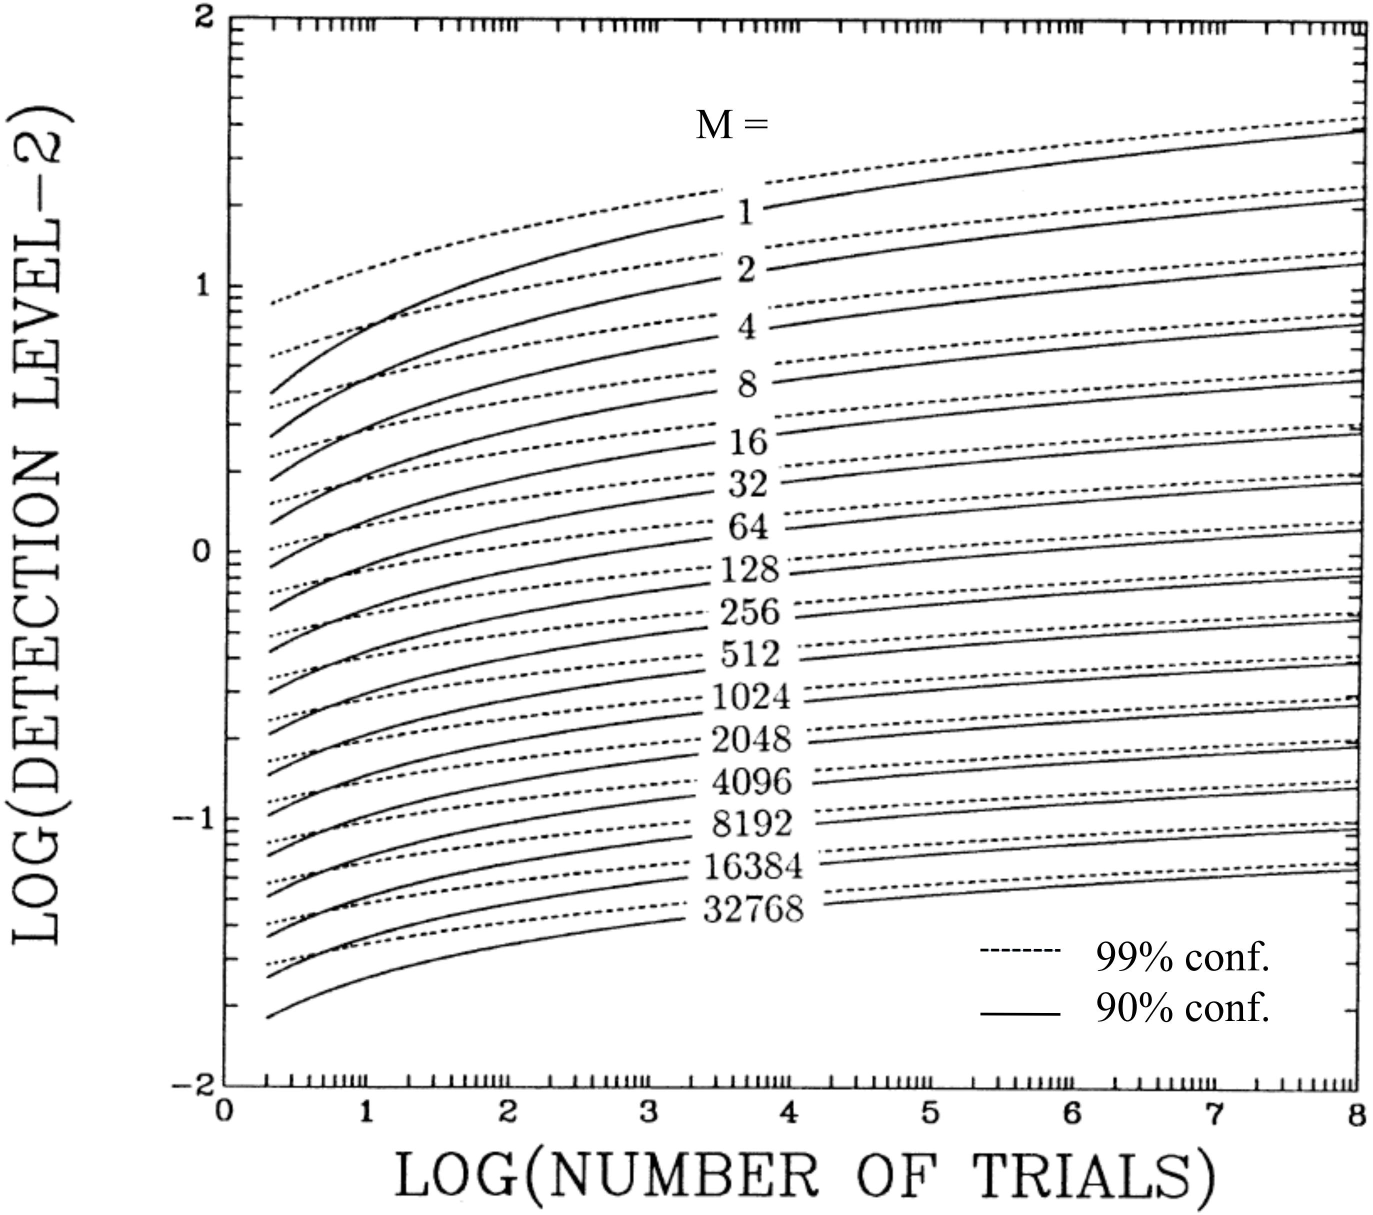
LOG(DETECTION LEVEL−2)
LOG(NUMBER OF TRIALS)
M =
1
2
4
8
16
32
64
128
256
512
1024
2048
4096
8192
16384
32768
99% conf.
90% conf.
2
1
0
−1
−2
0
1
2
3
4
5
6
7
8

Significance - the number of trials !

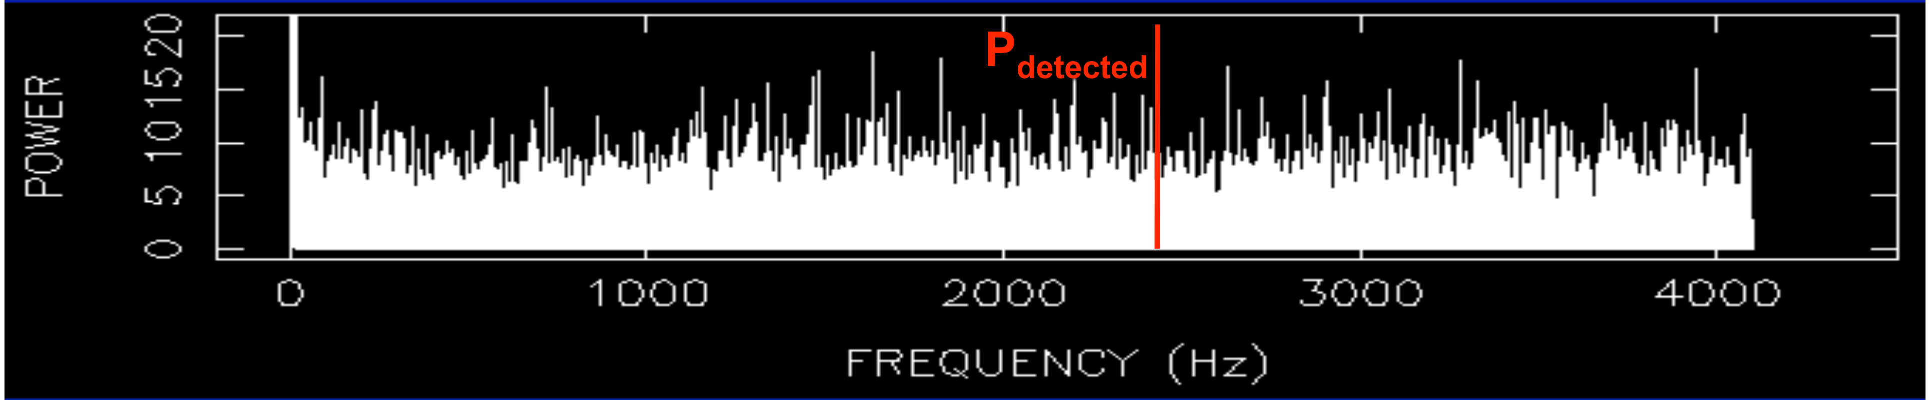

Turning this around, if we have detected a large power $P_{detected}$, then the confidence $1-\varepsilon$ at which we have detected it is

$$\boxed{1 - \varepsilon = 1 - N_{trial} Q(MP_{detected}|2M)}$$

(this probability is often expressed in sigma's and called significance, so 3σ for 99.7% -- note that while nice for ATels, this entire procedure is formally not considered proper statistics). In any case, once again:

It depends on – detected power $P_{detected}$

– number of powers averaged M

– number of trials N_{trial} !!!

Quantifying signal power

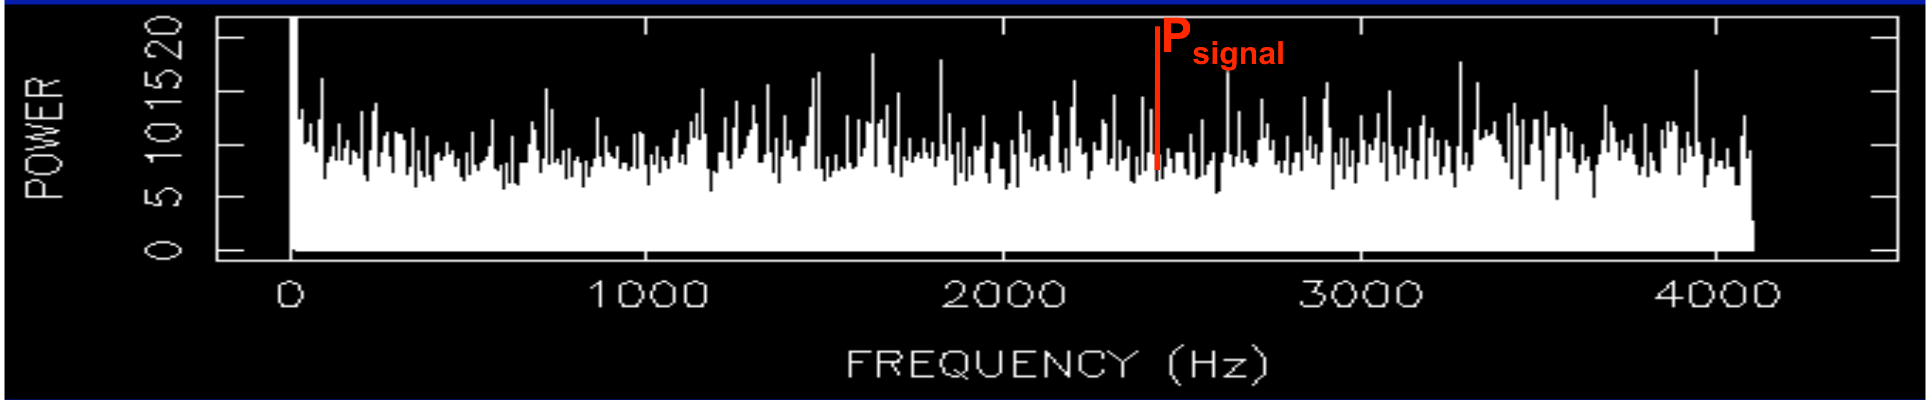

To detect signal power
(= to reject the null hypothesis 'just noise'),
and to quote the confidence of the detection,
you just need to know the noise power distribution [done].

To **quantify the signal power** you also need to know what is the interaction between noise and signal powers.

If you see a total power P_{tot}, then how much of that is P_{signal}?

SUPERPOSITION

Superposition theorem: Transform of the sum is sum of the transforms.

Suppose you have two signals x_k and y_k added together in one time series, then if

$$a_j = \sum_k x_k e^{i\omega_j t_k/N} \text{ and } b_j = \sum_k y_k e^{i\omega_j t_k/N} \Rightarrow a_j + b_j = \sum_k (x_k + y_k) e^{i\omega_j t_k/N}$$

So this is **not** in general true for the power spectrum!
It depends on relative phase how the two signals combine:
$|a_j + b_j|^2 = |a_j|^2 + |b_j|^2 + \text{ cross-terms} \Longrightarrow$
$P_{tot} = P_a + P_b + \text{ cross-terms}$

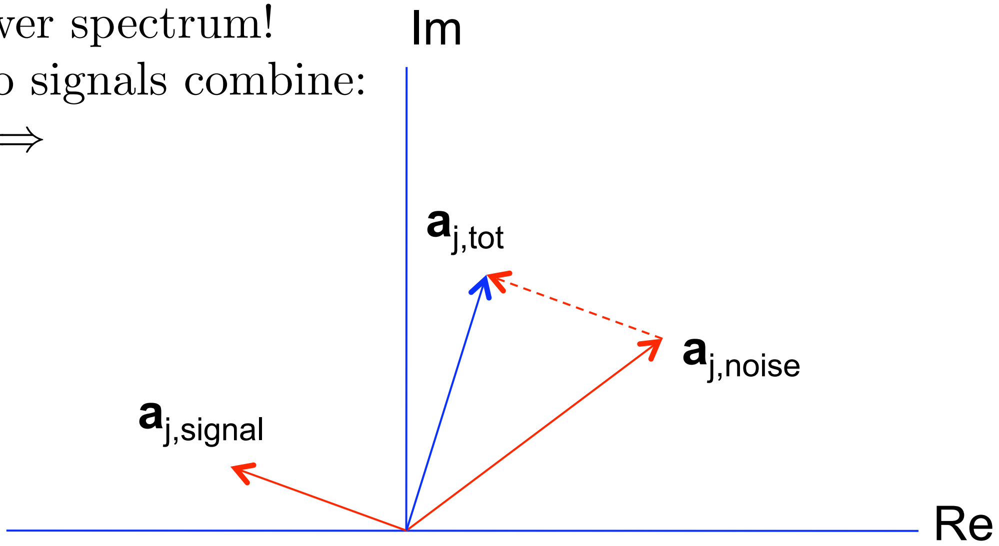

- Groth 1975 (ApJSupp 29, 285)
- Vaughan et al. 1994 (ApJ 435, 362)

discuss the distribution of P_{tot} given stochastic P_a and deterministic P_b.

If x_k and y_k are both uncorrelated noise: central limit theorem:
P_{tot} follows χ^2 distribution scaled to local mean power.
For large M: cross terms average out to zero and $\overline{P_M}$ follows normal distribution with standard deviation $\overline{P_M}/\sqrt{M}$

Powers are chi squared distributed around local mean power

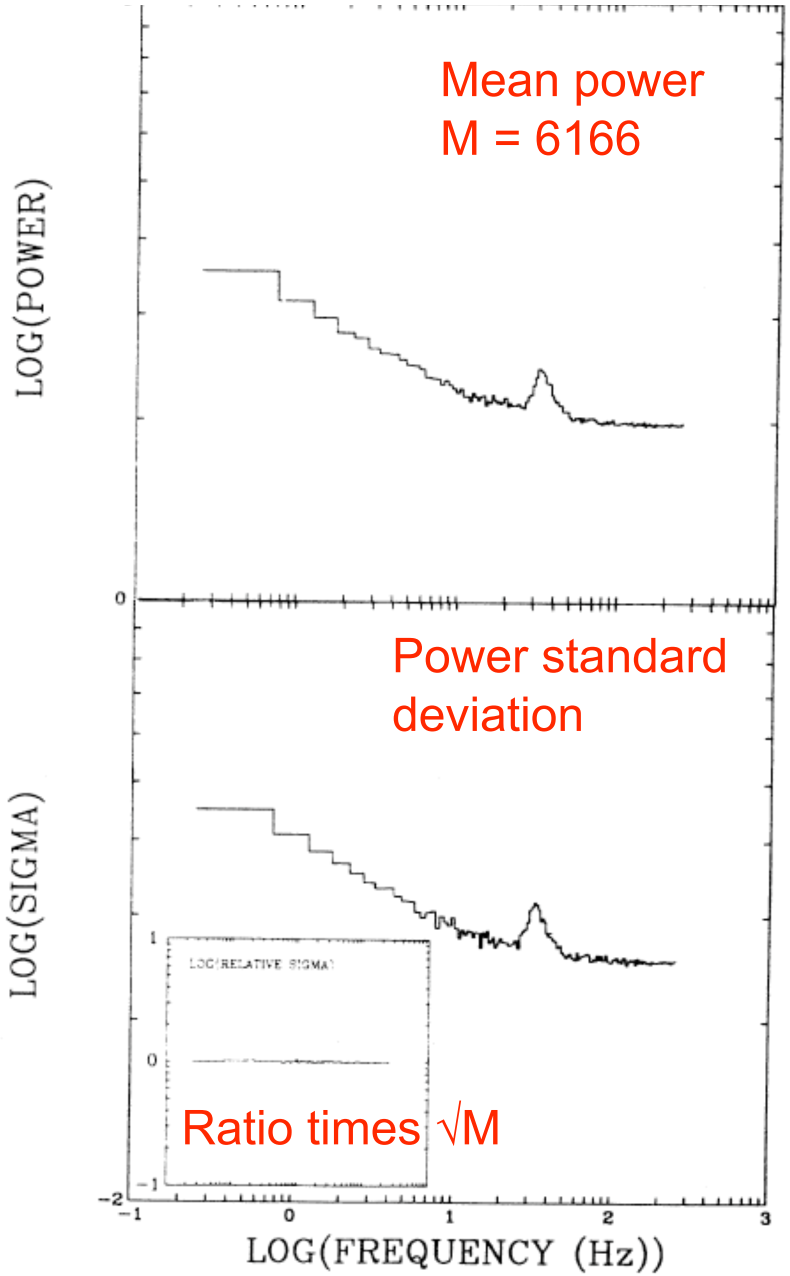

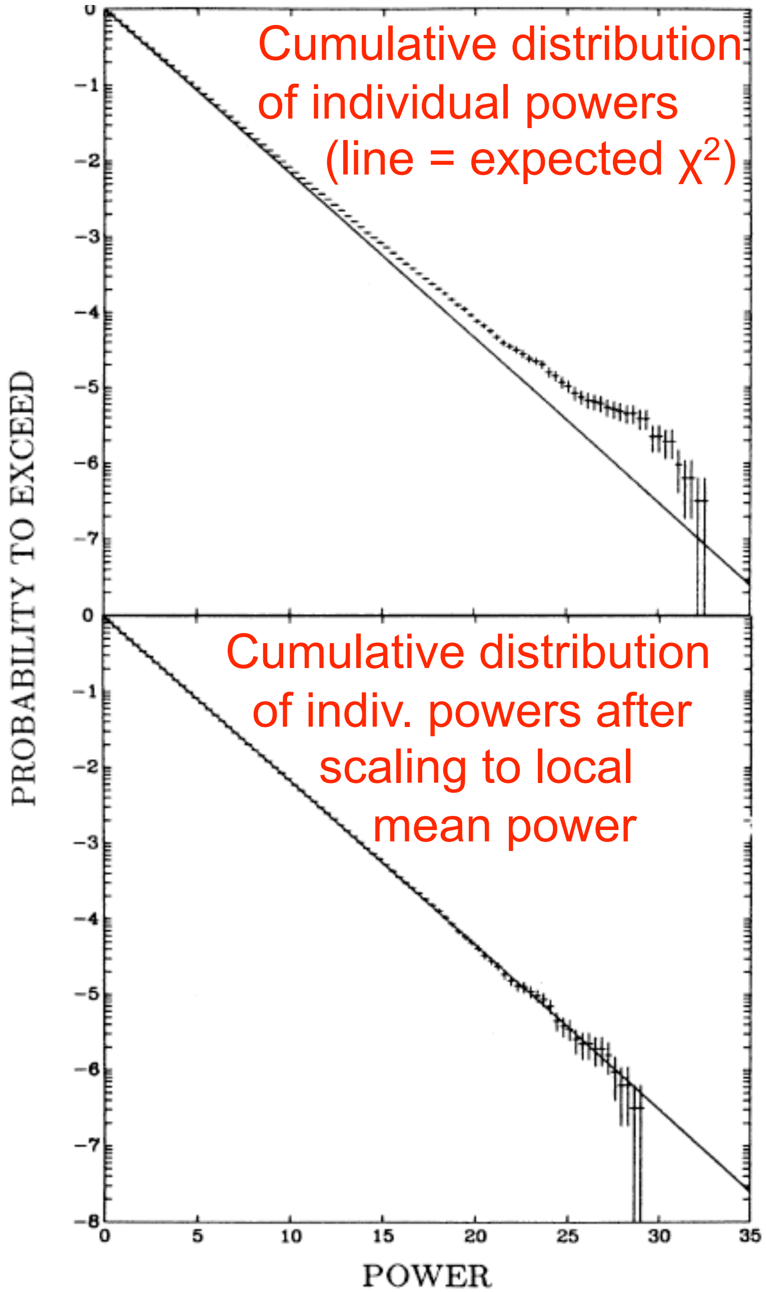

Detection against 'non-Poisson noise powers'

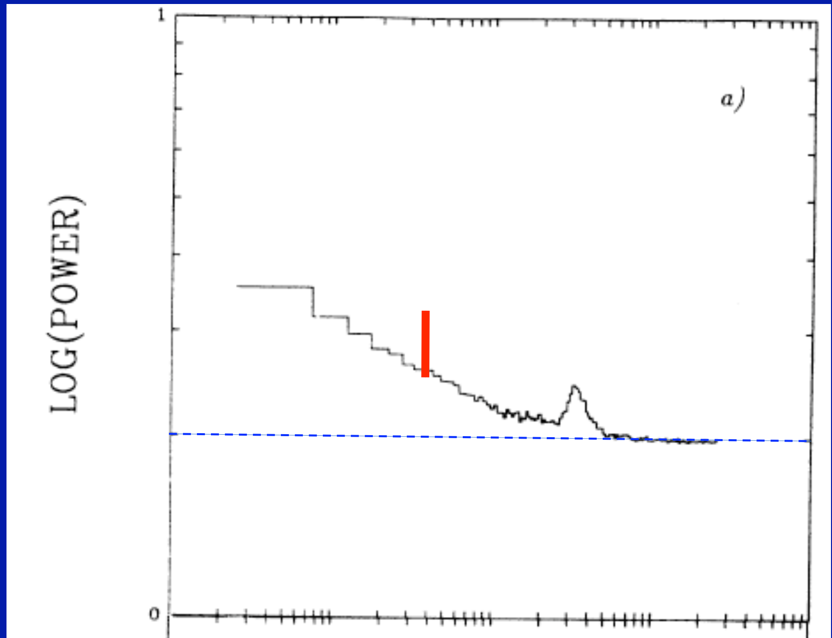

- When searching for a signal, always check what is the distribution of the 'background noise powers'
- Use χ^2 distribution scaled to local mean power to set detection level and evaluate significances

Upper limits and sensitivity when

$P_{tot} = P_{noise} + P_{signal}$

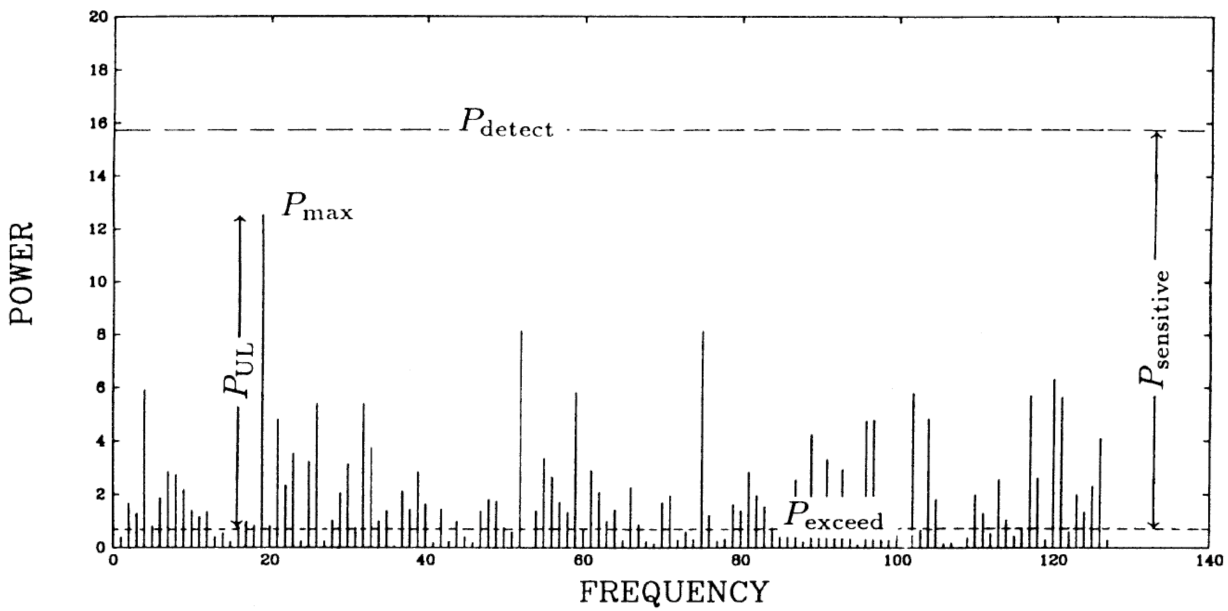

P_{detect} is level **unlikely** (prob ε) to be exceeded by noise in all trials
P_{exceed} is level **likely** (prob 1–δ) to be exceeded by noise in one trial
P_{max} is largest observed power.

- If we have a $P_j > P_{detect}$, a signal was detected at (1-ε) confidence: all P_j's together had a small probability ε to exceed this level
- If not, the (1–δ) confidence upper limit on P_{signal} is $P_{UL} = P_{max} - P_{exceed}$: if such a power would have been present at one given j, then it would with large probability (1–δ) have exceeded P_{max}, but this did not happen
- The (1–δ), (1-ε) sensitivity is $P_{sensitive} = P_{detect} - P_{exceed}$: if a signal power as large as $P_{sensitive}$ occurs at one given j, then it will with large probability (1–δ) exceed P_{detect}, and be detected at (1-ε) confidence

Signal to noise (Gaussian limit) — "single-trial significance"

So the sum of M Leahy-normalized powers of a Poisson-noise time series for large M is Gaussian with mean$= 2M$ and $\sigma = 2\sqrt{M}$, and $P = P_{\rm noise} + P_{\rm signal}$.

- Photon counting time series of length T and count rate $I_x = N_{ph}/T$
- Signal with fractional rms amplitude r, producing broad feature of width $\Delta\nu$
  Feature will contain $M = T\Delta\nu$ individual powers

Using $r = \sqrt{\sum P_j / N_{ph}} \Longrightarrow \sum P_j = N_{ph} r^2$, we find for the signal-to-noise:

$$\boxed{n_\sigma = \frac{N_{ph} r^2}{2\sqrt{M}}} = \frac{1}{2}\frac{T I_x r^2}{\sqrt{T\Delta\nu}} = \boxed{\frac{1}{2} I_x r^2 \left(\frac{T}{\Delta\nu}\right)^{1/2}}$$

In terms of source fractional rms $r_s = \dfrac{B+S}{S} r$, using $I_x = B + S$:

$$\boxed{n_\sigma = \frac{1}{2} I_x r_s^2 \left(\frac{S}{B+S}\right)^2 \left(\frac{T}{\Delta\nu}\right)^{1/2} = \frac{1}{2} r_s^2 \left(\frac{T}{\Delta\nu}\right)^{1/2} \frac{S^2}{B+S}}$$

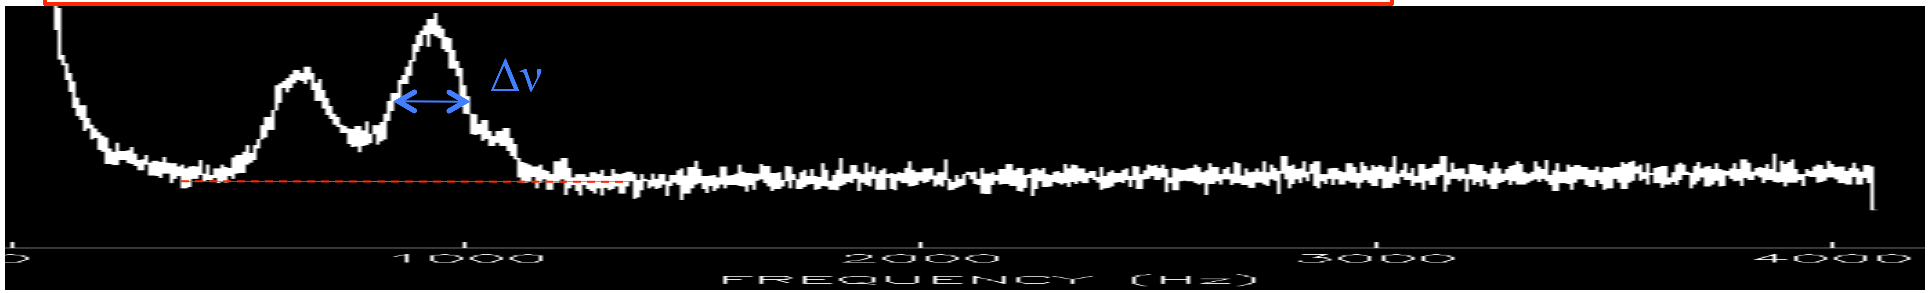

A practical procedure

- Segment data
- Fourier transform the segments
- Calculate power
- Leahy normalize → powers X^2 distributed
  → ML method (Aneta)

or …

- Average to reach Gaussian regime
- Rms normalize to r_s
- Set errors to local mean power / $\sqrt{M}$
- Analyze using standard chi-squared fitting techniques (Levenberg-Marquard) using multi-component models (e.g. Lorentzians)
- Characterize components by their rms and characteristic frequencies
- Method works fine for all broad features (= stochastic variability)
- Can easily be generalized to cross-spectral analysis
- Instrumental deadtime effects need to be carefully accounted for

Copyright Jeremy Drake

END