EUROFUSION WPJET1-PR(14) 12309

A.J. Webster et al.

Distinguishing Cause from Correlation in Fusion Plasma Experiments

Preprint of Paper to be submitted for publication in Proceedings of the Royal Society A

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

Distinguishing Cause from Correlation in Fusion Plasma Experiments

Anthony J. Webster[1] and JET EFDA contributors*

JET-EFDA, Culham Science Centre, OX14 3DB, Abingdon, UK

[1]*EURATOM-CCFE Fusion Association, Culham Science Centre, OX14 3DB, Abingdon, OXON, UK*
** See annex of F. Romanelli et al, "Overview of JET Results",*
(24th IAEA Fusion Energy Conference, San Diego, USA (2012)).

Preprint of Paper to be submitted for publication in
The Journal of the Royal Statistical Association

ABSTRACT

The generic question is considered: How can we determine the probability of an otherwise quasi-random event, having been triggered by an external influence? A specific problem is the quantification of the success of techniques to trigger, and hence control, plasma instabilities in magnetically confined fusion (MCF) experiments. The successful development of such techniques is essential to ensure tolerable heat loads on components in large fusion devices, something that is necessary for the development of an economically successful MCF power plant. Bayesian probability theory is used to rigorously formulate the problem and to provide a formal solution. Accurate but pragmatic methods are developed to estimate triggering probabilities, and are illustrated with experimental data. These allow results from experiments to be quantitatively assessed, and rigorously quantified conclusions to be formed.

1. INTRODUCTION

Presently the most successful high-performance nuclear fusion experiments use a toroidal magnetic field to confine the plasma in a tokamak [1]. Unfortunately the majority of these plasmas produce edge-localised instabilities (“ELMs”) [2], that eject particles and energy onto the plasma-facing components. In future power-plant scale devices, the size of the ELMs must be limited to avoid components being damaged. One proposal is to deliberately and rapidly trigger ELMs, with the hope of having a larger number of smaller ELMs. A method that has been found successful at triggering ELMs is to apply a rapidly increasing radial magnetic field to the plasma, that “kicks” the plasma vertically to trigger ELMs [3–5]. To understand and optimise the method it is necessary to explore the threshold between when the kicks trigger ELMs, and when they simply perturb the plasma. For example figure 1 shows a scan of kick sizes, with a systematic variation of the amplitude and duration of the voltage applied to coils that produce the radial magnetic field that can trigger the ELMs. For these situations it is necessary to be able to quantify the success by which ELMs are triggered, so as to better understand physically what is triggering the ELM, and practically, what is required to successfully do so. A complication is that ELMs occur regularly without kicks, so how can we determine whether an ELM is due to a kick or is naturally occurring? A major step towards this is to understand the time-scales over which the kick influences the plasma most strongly. This then determines a narrower time window when a kick is able to trigger an ELM, but it does not rule out the possibility that some ELMs occur naturally within that time window also. Correctly accounting for this possibility is a purpose of this paper. In the next section we outline how Bayesian probability theory can be used to formulate a general solution to this problem. The generic problem however is much broader - for example, how can you tell if the increase in insurance claims due to lightning strikes are due to global warming or statistical chance? It might be possible to adapt some of the methods here to those broader questions. The following sections explore how the solution can be applied to the specific question of determining the success of “kick-triggering” experiments.

2. FORMULATION AND FORMAL SOLUTION OF THE PROBLEM

Bayesian probability theory starts from some simple but reasonable requirements about a theory of probability, and shows how the usual laws of probability theory can be derived from them [7], such as the basic product rule,

$$P(A, B|C) = P(A|B, C)P(B|C) = P(B|A, C)P(A|C) \tag{1}$$

and the basic sum rule,

$$P(A|C) + P(\bar{A}|C) = 1 \tag{2}$$

where P(A,B|C) reads as the probability of A and B being simultaneously true given that C is true, and P(A|B,C) reads as the probability of A being true given that B and C are true, and Ā is used to denote the negation of A (i.e. "not A"). These relations can be used to formulate a variety of mathematical expressions that relate the probability of events [7]. One useful expression that is derived in the appendix for convenience is,

$$P(A|C) = P(A|B, C)P(B|C) + P(A|\bar{B}, C)P(\bar{B}|C) \tag{3}$$

In the context of our subsequent discussion this will be used to obtain the probability of the event A occurring at time t, in terms of the probability of event A occurring at time t given that action B has (or has not) triggered the event, with C determining the time at which the triggering is attempted. Specifically we consider the probability density of observing the next ELM at time t to t + dt after a "vertical kick" starts to influence the plasma at $t = \tau_m$, with,

$$P(t|\tau_m) = P(t|K, \tau_m)P(K|\tau_m) + P(t|\bar{K}, \tau_m)P(\bar{K}|\tau_m) \tag{4}$$

where $P(t|\tau_m)$ is the probability density of observing an ELM at time t after a (vertical) kick starts to influence the plasma at time τ_m, $P(t|K, \tau_m)$ is the probability density of an ELM at time t after τ_m given that a kick at has successfully triggered an ELM, and $P(t|\bar{K}, \tau_m)$ is the probability density of an ELM at time t after τ_m if a kick has failed to trigger an ELM. $P(K|\tau_m)$ is the probability of a kick successfully triggering an ELM and $P(\bar{K}|\tau_m)$ is the probability of a kick not triggering an ELM, given that the kick starts to influence the plasma at time τ_m. Because of (2), $P(\bar{K}|\tau_m) = 1 - P(K|\tau_m)$, and we have,

$$P(t|\tau_m) = P(t|K, \tau_m)P(K|\tau_m) + P(t|\bar{K}, \tau_m)(1 - P(K|\tau_m)) \tag{5}$$

This equation is exact, and can be developed in various ways. The aim is to construct a mathematical expression whose terms can be accurately estimated from the available experimental data, allowing

$P(K | \tau_m)$ to be determined. Here (5) will be used to formalise the idea that triggered ELMs will appear in a short time interval after a kick, using the information to rigorously define the probability of an ELM having been triggered. Crucially, this is done in a way that allows the possibility that some of the observed ELMs will have occurred naturally, the resulting estimates quantitatively account for this.

Some background. Following a "kick" in the JET tokamak, there is a time period of order 1-1.5 milliseconds before the kick's magnetic field starts to strongly interact with the plasma, and another time period proportional to the kick's duration, before the kick reaches its maximum strength. The time delay is due to the time needed for the magnetic field that the control coils are trying to produce, to diffuse through JET's metal vacuum vessel and other conducting components. By observing ELMs that are clearly triggered and paced at the kick-frequency, it appears that triggered ELMs occur approximately within τ_m to $\tau_m + \Delta\tau$, where $\Delta\tau$ is the kick's duration. For example see figure 2, whose results are consistent with a resistive diffusion of the "kicked" magnetic field through JET's metal vacuum vessel over a 1.0–1.5ms timescale, and ELMs being triggered after the magnetic perturbation has grown towards its maximum value over a timescale of 1.5-4.5ms depending on the kick's duration and amplitude. However, we cannot infer the kick success simply by counting the fraction of ELMs that occur in this time period after a kick, because a fraction of those will be expected to occur in that time interval naturally. Eq.5 accounts for this through the term $P(t | \bar{K}, \tau_m)P(\bar{K} | \tau_m)$, the probability of observing ELMs at a time t since a kick, given that the ELM is not triggered. Formally all of this can be achieved by integrating both sides of (5) from τ_m to $\tau_m + \Delta\tau$, with respect to t. Then after integration we get,

$$P_{\Delta\tau} = P(K) + P_{\bar{K}}(1 - P(K)) \tag{6}$$

where $P_{\Delta\tau} \equiv \int_{\tau_m}^{\tau_m+\Delta\tau} P(t)dt$, $\int_{\tau_m}^{\tau_m+\Delta\tau} P(t|K)dt = 1$ by the definition that these kicks will successfully trigger ELMs, and $P_{\bar{K}} \equiv \int_{\tau_m}^{\tau_m+\Delta\tau} P(t|\bar{K})dt$ is the probability of ELMs occurring naturally and being observed within τ_m to $\tau_m + \Delta\tau$ without having been triggered by kicks. Consequently,

$$P(K) = \frac{P_{\Delta\tau} - P_{\bar{K}}}{1 - P_{\bar{K}}} \tag{7}$$

$P_{\Delta\tau}$ and $P_{\bar{K}}$ both depend on $\Delta\tau$ and τ_m, although for reasonably large numbers of kicks we will find that $P_{\bar{K}}$ (and consequently $P_{\Delta\tau}$), are approximately independent of τ_m. Notice that if we had ignored naturally occurring ELMs, with $P_{\bar{K}} = 0$, we would simply have $P(K) = P_{\Delta\tau}$, or in other words, we would estimate the probability of a kick successfully triggering an ELM simply as the probability of observing an ELM in the "kick-triggered" time interval. If $P_{\bar{K}}$ is sufficiently small then the estimates of $P(K) = P_{\Delta\tau}$ and (7) will be similar. Also if $P_{\bar{K}}$ is small, then,

$$P(K) = P_{\Delta\tau} - P_{\bar{K}}(1 - P_{\Delta\tau}) + O\left(P_{\bar{K}}^2\right) \tag{8}$$

which is useful for obtaining simple leading-order error estimates as the sum of errors in $P_{\Delta\tau}$ and $P_{\bar{K}}$, which are easy to calculate and record. For kick-triggering studies, $P_{\bar{K}}$ is usually small. $P_{\Delta\tau}$ can be estimated from the kick-ELM data by counting the number of ELMs that occur naturally within the set of time intervals τ_m to $\tau_m + \Delta\tau$; this will be considered in a later section. $P_{\bar{K}}$ can be estimated from equivalent ELMing plasma data without kicks, which we discuss next.

3. ESTIMATING $P_{\bar{K}}$

Next we consider how to estimate $P_{\bar{K}}$ from ELM time data that does not have kicks or other triggering mechanisms. Firstly we calculate the probability of an ELM being observed in a time interval τ_m to $\tau_m + \Delta\tau$, finding a simple estimate for the limit where $\tau_m \sim n\bar{t}$ has $n >> 1$. This is a useful approximation, so we explore when it is reasonable, and test the estimate by calculating the number of ELMs that we would expect to occur naturally in the time intervals $(\tau_m, \tau_m + \Delta\tau)$ that we might incorrectly think are "triggered" ELMs, and compare this estimate with observations from experimental data.

The probability of observing the nth ELM in $(\tau_m , \tau_m + \tau)$, is,

$$\int_{\tau_m}^{\tau_m+\tau} p(x|n)dx \tag{9}$$

where $x = \Sigma_{i=1}^{n} t_i$, and t_i are the waiting times between the (i − 1)th and ith ELM. The probability of any ELM being in $(\tau_m , \tau_m + \tau)$, is,

$$P_{\bar{K}}(\tau_m) = \sum_{n=1}^{\infty} \int_{\tau_m}^{\tau_m+\tau} p(x|n)dx \tag{10}$$

for the present we have decided to explicitly emphasise the dependence of $P_{\bar{K}}$ on τ_m by writing $P_{\bar{K}}(\tau_m)$. As is outlined in 7.16 of Jaynes [7], p(xln) can be written as,

$$p(x|n) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \psi(\alpha)^n e^{-i\alpha x}\, d\alpha \tag{11}$$

with,

$$\psi(\alpha) = \int_{-\infty}^{\infty} P(t) e^{i\alpha t}\, dt \tag{12}$$

where P(t) is the ELM waiting time pdf, that gives the probability of observing an ELM in time t to t + dt as P(t)dt. For many situations P(t) can be reasonably approximated by a Gaussian [8], and for those cases we exactly have,

$$p(x|n) = \frac{1}{\sqrt{2\pi n\sigma^2}}\, exp\left(-\frac{(x - n\bar{t})^2}{2n\sigma^2}\right) \tag{13}$$

where $\bar{t}$ is the average ELM waiting time and σ is the ELM waiting times' standard deviation. These two quantities $\bar{t}$ and σ can be estimated from the ELM time data with an error of order $1/\sqrt{M}$, where M is the total number of observed ELMs [9]. Equation (13) is exact [7] whenever P(t) is a Gaussian, however if n >> 1 then with only a small number of exceptions [7], the Central Limit Theorem ensures that (13) remains true independent of the specific form of ELM waiting time pdf [7]. Using (10) and (13) the probability of an ELM in time τ_m, $\tau_m + \tau$ is,

$$P_{\bar{K}}(\tau_m) = \sum_{n=1}^{\infty} \int_{\tau_m}^{\tau_m + \tau} \frac{dx}{\sqrt{2\pi n \sigma^2}} exp\left(- \frac{(x - n\bar{t})^2}{2n\sigma^2} \right) \tag{14}$$

For $n\bar{t} >> \bar{t}$ (n >> 1) , the sum can be approximated by an integral, with $y = n\bar{t}$ and $dy = \bar{t}$, giving,

$$P_{\bar{K}} = \int_0^{\infty} \frac{dy}{\bar{t}} \int_{\tau_m}^{\tau_m+\tau} \frac{dx}{\sqrt{2\pi y(\sigma^2/\bar{t})}} exp\left(-\frac{(x - y)^2}{2y(\sigma^2/\bar{t})}\right) \tag{15}$$

where the dependence of P_K on τ_m has been removed for reasons that will be explained shortly. Changing variables, with $x = \hat{x}\sigma^2/\bar{t}$ and $y = v^2\sigma^2/\bar{t}$, so that $dy/\sqrt{y} = 2dv(\sigma^2/\bar{t})$, and exchanging the order of integration, gives,

$$P_{\bar{K}} = \frac{1}{\bar{t}} \int_{\tau_m(\sigma^2/\bar{t})}^{(\tau_m+\tau)(\sigma^2/\bar{t})} d\hat{x} \left(\frac{\sigma^2}{t}\right) \frac{1}{\sqrt{2\pi}} \int_0^{\infty} dv\, exp\left(\frac{(x - v^2)^2}{2v^2}\right) \tag{16}$$

Using the integral [10],

$$\int_0^{\infty} dx\, exp\left(-ax^2 - \frac{b}{x^2}\right) = \sqrt{\frac{\pi}{4a}}\, exp\left(-2\sqrt{ab}\right) \tag{17}$$

it can easily be shown that,

$$\int_0^{\infty} dv\, exp\left(-\frac{(\hat{x} - v^2)^2}{2v^2}\right) = \sqrt{\frac{\pi}{2}} \tag{18}$$

a surprising result, because like a normalised probability distribution with a parameter $\hat{x}$, the result is independent of $\hat{x}$. Consequently, doing the final integral over $\hat{x}$ and cancelling terms, we get,

$$P_{\bar{K}} = \frac{\Delta\tau}{\bar{t}} \tag{19}$$

A surprising and disappointing aspect of (19) is that τ_m does not appear, and the influence of correlations when $\tau_m \simeq \bar{t}$ for example, are not captured by the calculation. The only approximation that is always present enters subtly, through the need of n >> 1 for the sum to be accurately approximated by an integral. The approximation n >> 1 removes the information about the discrete nature of ELMs and consequently is unable to describe the influence of correlations between τ_m and $\bar{t}$. Consider the example with τ_m small. Then (9) could be accurately approximated by keeping only the first few terms in the sum over n, and the approximation of (15) would be poor. The influence of correlations between the natural ELM frequency and the kick frequency is discussed in the Appendix, where it finds that provided $\Delta\tau/\sigma$ is reasonably small, where σ is the standard deviation in the natural ELM period, then any coherence between the natural ELMs and kick frequency will be rapidly lost within the first one or possibly two ELMs. Another caveat is that if P(t) is non-Gaussian, then n >> 1 is required for the central limit theorem to ensure that p(x|n) will tend to (13); but occasionally for some distributions even if n >> 1 the rate of convergence can be slow and the estimates less good than expected. In summary, for small $\tau_m/\bar{t}$, for which n ~ 1 is most relevant, then in principle equations (9), (11), and (12) must be used; but despite the caveats mentioned above, provided $\tau_m/\bar{t}$ >> 1 and $\Delta\tau/\sigma$ are reasonably small then (15) and (19) should in most cases provide an accurate approximation.

Estimates for $\bar{t}$ and σ must be found from equivalent ELM data for natural (untriggered) ELMs. They can be estimated from a set of M ELM waiting times $\{t_i\}$, of $\bar{t} \simeq \frac{1}{M}\Sigma_{i=1}^{M} t_i$ and $\sigma^2 \simeq 1\frac{1}{M-1}\Sigma_{i=1}^{M}(\bar{t}_i - t)^2$ both of which have errors of order $\overline{1/}\sqrt{M}$ [9]. Ideally there will be plenty of data for natural ELMs, so that M >> 1, and we can neglect the errors in $\bar{t}$ and σ compared with the other errors in (23). Otherwise the leading order errors give,

$$P_{\bar{K}} \simeq \frac{\Delta\tau}{\bar{t}}\left(1 \pm \frac{\sigma}{\bar{t}\sqrt{M}}\right) \tag{20}$$

with the error estimated from $\bar{t} \simeq (1/M)\,\Sigma_{i=1}^{M} t_i \pm \sigma/\sqrt{M}$. The greatest source of systematic error is likely to be through unintended changes to the plasma's properties between pulses, leading to incorrect estimates for $\bar{t}$ and σ. This risk can be reduced by using a reference pulse from the same session, or a reference phase from within a pulse, although longer pulses will give greater statistical accuracy. In practice $P_{\bar{K}}$ is comparatively small, and unless the number of kicks are much larger than the ~15 in the example considered here, then the errors are dominated by those in $P_{\Delta\tau}$ that are statistical in origin and can only be reduced by increasing the number of kicks in the study.

We can test Eq. 19 by calculating the number of natural (untriggered) ELMs we would expect to observe in the "kick-triggered" time intervals, and compare this with the number that are observed in equivalent time intervals of experimental data (without kicks). Assuming that (19) is a reasonable approximation for $P_{\bar{K}}$, as it very often will be, then we can estimate $N_{\bar{K}}$ as follows. The probability of observing an unordered sequence of $N_{\bar{K}}$ time intervals $(\tau_1, \tau_1 + \tau), (\tau_2, \tau_2 + \tau), \ldots, (\tau_{N_K}, \tau_{N_K} + \tau)$, with $N_{\bar{K}}$ ELMs spread among them is given by,

$$P(N_{\bar{K}}, N_K) = \binom{N_K}{N_{\bar{K}}} (P_{\bar{K}})^{N_{\bar{K}}} (1 - P_{\bar{K}})^{N_K - N_{\bar{K}}} \tag{21}$$

i.e. the Binomial distribution (because the order in which the $N_{\bar{K}}$ natural ELMs coincide with an interval τ_i to $\tau_i + \Delta\tau$ is not important). Consequently the expected number of (untriggered) ELMs to be observed by chance in the "triggered ELMs" time interval can be calculated from

$\langle N_{\bar{K}} \rangle = \Sigma_{N_{\bar{K}}=1}^{N_K} N_{\bar{K}} P(N_{\bar{K}}, N_K)$, and is found using (19) for $P_{\bar{K}}$ to be,

$$\langle N_{\bar{K}} \rangle = N_K \frac{\Delta \tau}{\bar{t}} \tag{22}$$

The standard deviation is $\sqrt{\langle N_{\bar{K}}^2 \rangle - \langle N_{\bar{K}} \rangle^2} = \sqrt{N_K \frac{\Delta \tau}{\bar{t}} \left(1 - \frac{\Delta \tau}{\bar{t}}\right)}$ which allows an estimate for the error in (22). Therefore taking $N_{\bar{K}} \simeq \langle N_{\bar{K}} \rangle$, we obtain an estimate of,

$$N_{\bar{K}} \simeq N_{\bar{K}} \frac{\Delta \tau}{\bar{t}} \pm \sqrt{N_K \frac{\Delta \tau}{\bar{t}}} \sqrt{1 - \frac{\Delta \tau}{\bar{t}}} \tag{23}$$

Eq. 23 allows us to estimate $N_{\bar{K}}$, given the number of kicks N_K, and a reference phase of normal ELMs from which the ELMs' average waiting time $\bar{t}$ and standard deviation σ can be estimated.

To test the above Eq. 23, we have counted the number of ELMs within time intervals of n×(0.02) to n×(0.02)+0.004 seconds, with n = 0,1,2, ... , during the time period of 9–13.5 seconds of ten equivalent JET plasmas that do not have kicks. Details of the pulses are in [11], we consider the subset of: 83630, 83629, 83628, 83627, 83626, 83625, 83624, 83640, 83641, and 83642, for which the average ELM waiting time across all pulses was $\bar{t} = 0.030$ seconds. Figure 3 plots the number of ELMs that are counted in the time intervals that would usually be presumed to be "kick-triggered", divided by the number of "kick-triggered" time intervals (estimated as the time divided by the period between the intervals), along with our estimate for $N_{\bar{K}}/N_K$ using Eq. 23. The agreement is surprisingly good. If the results were Gaussianly distributed then we would expect approximately 68% of the data sets to be within one standard deviation, indicated on the plots by the dashed lines. The results are reassuringly consistent with this.

4. ESTIMATING $P_{\Delta\tau}$

The estimation of $P_{\Delta\tau}$ given the observation of m ELMs in n "kick-triggered" time intervals, is a classic and well-known problem in Bayesian probability theory that we will briefly review next. The probability of observing m successes in n trials, with a probability p of a success at each trial, is given by the binomial distribution,

$$P(m, n | p) = \binom{n}{m} p^m (1 - p)^{n-m} \tag{24}$$

Bayes theorem gives [7, 9],

$$P(p|m, n) = \frac{P(m, n | p) P(p)}{\int_0^1 P(m, n | p) P(p) dp} \tag{25}$$

which allows us to obtain a probability distribution for p given the observed m successes from n trials, and a suitable prior distribution P(p). Present models for ELMs have ELMs being triggered once a threshold of pressure gradient or current is exceeded [12]. If we regard kick-triggering in a similar way, with a probability of success that rapidly steps between zero and one - either a kick is strong enough to trigger an ELM or it is not, then we might follow the reasoning of Jaynes (pages 382-386 of Ref. [7]), that leads to the "Haldane" prior P(p) with,

$$P(p) = \frac{C}{p(1-p)} \tag{26}$$

with C a constant. Using this prior in (25) along with (24) for P(m, n|p), we get (after doing the integral in the denominator) [7],

$$P(p|m,n)dp = \frac{(n-1)!}{(m-1)!(n-m-1)!}p^{m-1}(1-p)^{n-m-1}dp \tag{27}$$

This gives an expected value for p of,

$$\langle p \rangle = \frac{m}{n} \tag{28}$$

with a standard deviation σ of,

$$\sigma = \sqrt{\frac{m(n-m)}{n^2(n+1)}} \tag{29}$$

Alternately, we might regard kick-triggering as a statistical process as in [8], with a success probability that increases gradually with the size of the kick. Then it would seem reasonable to take a uniform prior, with P(p) a constant. With this assumption we get,

$$P(p|m,n)dp = \frac{(n+1)!}{m!(n-m)!}p^{m}(1-p)^{n-m}dp \tag{30}$$

This gives an expected value for p of,

$$\langle p \rangle = \frac{m+1}{n+2} \tag{31}$$

with a standard deviation σ of,

$$\sigma = \sqrt{\frac{(m+1)(n-m+1)}{(n+2)^2(n+3)}} \tag{32}$$

Whichever prior we regard as being most reasonable, we can estimate $P_{\Delta\tau}$ from the N ELMs observed during the total of N_K kicked time intervals, as,

$$P_{\Delta\tau} \simeq \langle p \rangle \pm \sigma \tag{33}$$

with ⟨p⟩ and σ estimated from either 28 and 29, or 31 and 32. We will evaluate $P_{\Delta\tau}$ with both priors to provide examples of how the results can become modified.

Equations 7, 20, and 33 allow us to rigorously estimate the probability of kicks triggering ELMs, given the number of kicks N_K, a suitable estimate for τ_m and $\Delta\tau$, the observed number of ELMs N in the intervals τ_m to $\tau_m + \Delta\tau$, and an equivalent phase of plasma with natural (unkicked) ELMs from which to estimate $N_{\bar{K}}$ from (20).

5. EXAMPLES: QUANTIFYING KICK-TRIGGERING SUCCESS

A kick consists of a step in the voltage that is applied to the vertical control coils, that has a duration in time and an amplitude [6]. In a recent JET experiment (Pulse No: 83440), the duration and amplitude of the kick was deliberately varied to explore how kick-triggering success depended on them. To analyse the kick-triggering probability it is necessary to explore a bit more of the physics involved in a vertical kick. The diffusion time of JET's vacuum vessel is of order 1.3ms, so a kick of duration τ_{kick} has negligible influence on the plasma for 1.3ms, has a maximum perturbation over τ_{kick} to τ_{kick} + 1.3ms, and falls to less than half of its value in of order $2 \times \tau_{kick}$. Based on past experience, ELMs in JET are not expected to be triggered with kicks of less than 10Wb (10 Volt seconds), so it seems safe to assume that a kick cannot modify plasma stability until it has produced at least 5Wb. After ~1.3ms the kick's constant voltage produces a perturbation to the radial magnetic field that grows approximately linearly with time, at least for kick durations of less than 4–5ms or so. So for a kick with voltage V_{kick} it seems reasonable to define a minimum time τ_m for which $(\tau_m - 1.3)V_{kick} = 5$, i.e. $\tau_m = 5/V_{kick} + 1.3$, below which the kick has negligible influence. If we take $2\tau_{kick}$ as an upper time limit over which the kick significantly perturbs the plasma, then $\Delta\tau = 2\tau_{kick} - 5/V_{kick} - 1.3$.

JET plasma Pulse No: 83440 was based on the plasma in Pulse No: 82630 over the period 18.5–20s, that was extended in 83440 to allow kicks to be tested over an extended pulse duration. During the 18.5–20s period in 82630 there were M = 24 ELMs, with an average period $\bar{t} = 0.059$ and standard deviation of $\sigma = 0.011$. This gives $\sigma/(\bar{t}\sqrt{M}) = 0.038$. Within the Pulse No: 83440 the kick amplitude and duration are varied, as shown in figure 1, with results as in figure 2, and as recorded in table I. The different kick durations lead to different values of $\Delta\tau$ that are also recorded in table I, along with the number of kicks in each set N_K, the number of ELMs that occur in the time interval τ_m to $\tau_m + \Delta\tau$, and the values estimated from that data for $P_{\bar{K}}$, $P_{\Delta\tau}$, P(K), and their errors. The quoted error for P(K) is the largest of the upper and lower errors as calculated using the estimates for $P_{\Delta\tau}$ and $P_{\bar{K}}$. The Haldane prior biases the results towards kicks being one of either successful or unsuccessful, and consequently because estimates for P(K) are all greater than 0.5, the Haldane prior gives higher kick-triggering probabilities than the uniform prior. Intuitively, the author feels more comfortable with the more conservative, lower estimates for P(K) that are obtained with a uniform prior, although most present models for ELM triggering favour a clear threshold for an ELM being triggered, or not. Independent of the choice of prior, from a physical point of view the results are surprising, because they indicate that kicks with quite small amplitude can still trigger ELMs if their duration is long enough. This has never previously been demonstrated, and can only be claimed with confidence because of the analysis presented here. When applied to larger data sets, a similar analysis will allow us to quantitatively assess how kick-triggering of ELMs depends on the kick properties such as its amplitude, duration, total Webers (Volt seconds), and maximum magnetic field perturbation. This will help us to understand how kicks trigger ELMs, and to optimise the kick properties for triggering ELMs with the "softest" possible kick with lowest voltage for example.

CONCLUSIONS

We have considered the generic problem of assessing whether an otherwise quasi-random event such as an edge-localised plasma instability (an "ELM"), has been triggered by an external influence. The specific problem of assessing the success of experiments to deliberately trigger ELMs by rapid vertical displacements of the plasma were considered, leading to a simple set of rigorously derived formula to estimate the kick-triggering probability and its accuracy. The advantage of the method is its rigour - we can now assert with confidence that vertical kicks with a voltage of only 3kV were in some cases triggering ELMs. This was unexpected, and has never previously been demonstrated. The method allows the success of kicks to be quantitatively assessed with unprecedented accuracy, allowing systematic studies to quantitatively determine how kick-triggering success depends on properties such as their duration, amplitude, and the peak magnetic field perturbation it produces. More generally, we believe that analogous calculations and direct adaptations of the method presented here can be used to help distinguish cause from accidental correlation of quasi-random events in other circumstances.

ACKNOWLEDGEMENTS

AJW thanks Steve Fitzgerald for interesting past discussions about path integral calculations that proved useful in this paper, and both Greg Colyer and Richard Kemp for helpful comments and discussions about different choices of prior P(p). The kicktriggering experiments discussed in this paper were possible due to the help and scientific coordination of Thomas Eich, Rory Scannel, Peter Lang, Peter Lomas, and Fernanda Rimini. This work was supported by EURATOM and carried out within the framework of the European Fusion Development Agreement. To obtain further information on the data and models underlying this paper please contact PublicationsManager@ccfe.ac.uk. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

THE SUM AND PRODUCT RULES

The basic sum and product rules of Bayesian probability theory [7] can be used to derive the key formulas used in this paper. Firstly consider Eq. (3), noting that,

$$\begin{aligned} P(A,B|C)+P(A,\bar{B}|C) &= \left[P(B|A,C)+P(\bar{B}|A,C)\right]P(A|C) \\ &= P(A|C) \end{aligned} \tag{34}$$

where the 1st line uses the basic product rule Eq. (1), and the 1st-2nd line uses the basic sum rule Eq. (2). Next using the product rule to expand $P(A,B|C) = P(A|B,C)P(B|C)$ and $P(A,\bar{B}|C) = P(A|\bar{B},C)P(\bar{B}|C)$, then (34) expands to gives,

$$P(A|C) = P(A|B,C)P(B|C) + P(A|\bar{B},C)P(\bar{B}|C) \tag{35}$$

which is Eq. (3). If we replace A with t, B with K, and C with τ_m, then we have,

$$P(t|\tau_m) = P(t|K,\tau_m)P(K|\tau_m) + P(t|\bar{K},\tau_m)P(\bar{K}|\tau_m) \tag{36}$$

which is Eq. 4, and gives $P(t|\tau_m)$ the probability density of observing an ELM at time t after a kick at time τ_m, in terms of $P(t|K,\tau_m)$ the probability density of observing an ELM at time t after a kick at time τ_m successfully triggers an ELM, and $P(t|\bar{K},\tau_m)$ the probability density of observing an ELM at time t after a kick at time τ_m fails to trigger an ELM. The quantity $P(K|\tau_m)$ is the probability of a kick being successful, given that the kick was at time τ_m. As noted in the main text, the sum rule requires that $P(\bar{K}|\tau_m) = 1 - P(K,\tau_m)$.

CORRELATIONS BETWEEN ELMS AND KICKS

To assess how much correlations between the ELMs' natural frequency and the kick frequency have the potential to influence the estimate, we consider a situation where an ELM has naturally occurred within the time window τ_m to $\tau_m + \Delta\tau$ that is usually associated with kick-triggered ELMs, and where the average ELM waiting time $\bar{t}$ equals the time between kicks $\tau_m - \tau_{m-1}$. Assuming the distribution of ELM waiting times is reasonably approximated by a Gaussian distribution with $P(t) = (1/\sqrt{2\pi\sigma^2})\exp\{-(t-\bar{t})^2/2\sigma^2\}$, that can often be the case for type I ELMs [8], then the probability p of an ELM between $\bar{t}$ to $\bar{t} + \Delta\tau$ is,

$$p \simeq \Delta\tau P(\bar{t}) = \frac{1}{\sqrt{2\pi}}\frac{\Delta\tau}{\sigma} \tag{37}$$

The probability of a sequence of m ELMs within: $\bar{t}$ to $(\bar{t} + \Delta\tau)$, $2\bar{t}$ to $(2\bar{t} + \Delta\tau)$, ..., $m\bar{t}$ to $(m\bar{t} + \Delta\tau)$, is then simply p^m, giving the number of correlated ELMs that we would expect to observe as,

$$\langle m \rangle = \frac{\sum_{m=1}^{\infty} m p^m}{\sum_{m=1}^{\infty} p^m} = \frac{p\frac{\partial}{\partial p}\sum_{m=1}^{\infty} p^m}{\sum_{m=1}^{\infty} p^m} = \frac{1}{1-p} \tag{38}$$

where we used $\Sigma_{m=1}^{\infty} p^m = p/(1-p)$ to evaluate the sums. $\langle m \rangle$ tends to 1 as $\Delta\tau/\sigma \to 0$, or infinity as $\Delta\tau/\sigma \to 1$. Most of the cases we are interested in have $\Delta\tau/\sigma << 1$. For those cases, even if $\bar{t} = \tau_m - \tau_{m-1}$, then the number of ELMs to be incorrectly counted as "triggered" by the kicks will be small, with $\langle m \rangle < 2$ for $\Delta\tau/\sigma < 0.5$, and of order 1 for $\Delta\tau/\sigma << 1$. In most cases, $\bar{t}$ will be substantially different to $\tau_m - \tau_{m-1}$, and the potential influence of correlations can be neglected entirely.

REFERENCES

[1]. J. Wesson Tokamaks (Oxford University Press, Oxford, 1997).

[2]. H. Zohm, Plasma Physics and Controlled Fusion **38**, 105, (1996).

[3]. A.W. Degeling, Y.R. Martin, J.B. Lister, L. Villard, A.N. Dokouka, V.E. Lukash, R.R. Khayrutdinov, Physics and Controlled Fusion, **45**, 1637, (2003).

[4]. P.T. Lang, A.W. Degeling, J.B. Lister, et al. Physics and Controlled Fusion, **46**, L31, (2004).

[5]. Y. Liang, Fusion Science and Technology **59**, 586, (2011).

[6]. F.G. Rimini, F. Crisanti, R. Albanese, et al. Fusion Engineering and Design, **86**, 539-543, (2011).

[7]. E.T. Jaynes "Probability Theory the Logic of Science", Cambridge University Press, (2003).

[8]. A.J. Webster and R.O. Dendy, Physical Review Letters **110**, 155004, (2013).

[9]. D.S. Sivia "Data Analysis a Bayesian Tutorial" Oxford University Press, Great Clarendon Street, Oxford, (1996).

[10]. L.S. Schulman "Techniques and Applications of Path Integration" John Wilely & Sons, Inc. (1981).

[11]. S. Brezinsek, T. Loarer, V. Phillips, et al. Nuclear Fusion, **53**, 083023, (2013).

[12]. A.J. Webster, Nuclear Fusion, **52**, 114023, (2012).

Kick Properties of JET Pulse No: 83440								
V_{kick} (kV)	12	9	6	9	6	3	6	3
τ_{kick} (ms)	2.5	2.5	2.5	3.3	3.3	3.3	4.5	4.5
$\Delta\tau$ (ms)	3.3	3.1	2.9	4.7	4.5	3.6	6.9	6.0
$P_{\bar{K}}$	0.056 ± 0.002	0.053 ± 0.002	0.049 ± 0.002	0.080 ± 0.003	0.076 ± 0.003	0.062 ± 0.002	0.116 ± 0.004	0.102 ± 0.004
N_K	15	15	15	14	14	14	14	14
N	14	13	9	14	14	9	14	11
With uniform prior for $P(p)$								
$P_{\Delta\tau}$	0.882 ± 0.076	0.824 ± 0.090	0.588 ± 0.116	0.938 ± 0.059	0.938 ± 0.059	0.625 ± 0.117	0.938 ± 0.059	0.750 ± 0.105
$P(K)$	0.88 ± 0.09	0.81 ± 0.10	0.57 ± 0.12	0.93 ± 0.07	0.93 ± 0.07	0.60 ± 0.13	0.93 ± 0.07	0.72 ± 0.12
With Haldane prior $P(p) = 1/p(1-p)$								
$P_{\Delta\tau}$	0.933 ± 0.062	0.867 ± 0.085	0.600 ± 0.122	1.000 ± 0.000	1.000 ± 0.000	0.643 ± 0.124	1.000 ± 0.000	0.786 ± 0.106
$P(K)$	0.93 ± 0.06	0.86 ± 0.09	0.58 ± 0.12	1.0 ± 0.00	1.00 ± 0.00	0.62 ± 0.12	1.00 ± 0.00	0.76 ± 0.11

Table I: The table summarises the kick properties and successes for JET PulseNo: 83440, recording the voltage V_{kick} and duration τ_{kick} of the kick. V_{kick} and τ_{kick} determine the duration $\Delta\tau$ over which the kick influences the plasma, allowing an estimate for $P_{\bar{K}}$ to be calculated. The number of kicks and the number of ELMs within a time interval $\Delta\tau$ allows an estimate for $P_{\Delta\tau}$ to be calculated, and subsequently also for P(K) the probability of the kick triggering an ELM. Estimates are calculated with a uniform prior P(p), and with the Haldane prior $P(p) = 1/p(1-p)$.

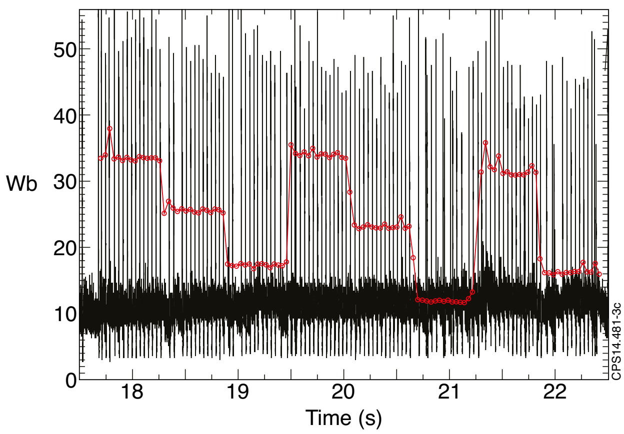

Figure 1: JET Pulse No: 83440. The measured Volt seconds (Wb) of each kick is plotted (red circles) versus the time in seconds during the pulse, with the strong Be II light emissions due to ELMs superimposed (not to scale). For kicks with large amplitude in Wb, ELMs appear to be clearly synchronised with the kicks, but for small kicks such as those between 20.6s and 21.4s their influence is much less clear.

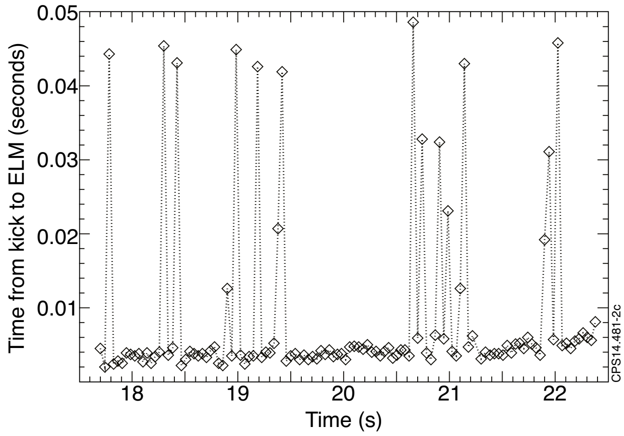

Figure 2: JET Pulse No: 83440. The time from the start of a kick to the next ELM is plotted versus time during the plasma pulse. Due to the limited kick duration and the time for the kicked magnetic field to diffuse through the plasma, kicks can only influence the plasma for of order 1-8ms after they start. Most of the ELMs bunched within this time interval will have been triggered by a kick, although some may have occurred naturally.

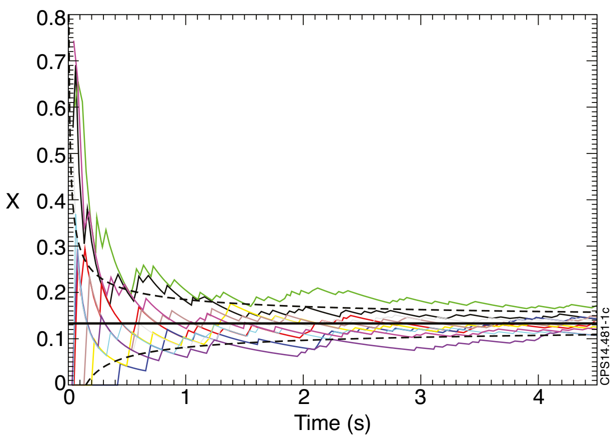

Figure 3: Comparison between calculated estimates for $N_{\overline{K}}$, and observed from experimental data. For ten equivalent JET plasmas, the number of ELMs are counted that occur by chance within the set of intervals τ_m to $\tau_m + \Delta\tau$, that if kicks were present would be presumed to be kick-triggered. Theoretical lines are: the estimated average (thick black line), the average ± the standard deviation (dashed black lines). Assuming the data is Gaussianly distribution then we would expect 68% of results to be within the dashed lines. The results are reassuringly consistent with this.