

G. Verdoolaege, R. Fischer, G. Van Oost
and JET EFDA contributors

Potential of a Bayesian Integrated Determination of the Ion Effective Charge via Bremsstrahlung and Charge Exchange Spectroscopy in Tokamak Plasmas

“This document is intended for publication in the open literature. It is made available on the understanding that it may not be further circulated and extracts or references may not be published prior to publication of the original when applicable, or without the consent of the Publications Officer, EFDA, Culham Science Centre, Abingdon, Oxon, OX14 3DB, UK.”

“Enquiries about Copyright and reproduction should be addressed to the Publications Officer, EFDA, Culham Science Centre, Abingdon, Oxon, OX14 3DB, UK.”

The contents of this preprint and all other JET EFDA Preprints and Conference Papers are available to view online free at www.iop.org/Jet. This site has full search facilities and e-mail alert options. The diagrams contained within the PDFs on this site are hyperlinked from the year 1996 onwards.

Potential of a Bayesian Integrated Determination of the Ion Effective Charge via Bremsstrahlung and Charge Exchange Spectroscopy in Tokamak Plasmas

G. Verdoolaege, R. Fischer, G. Van Oost
and JET EFDA contributors*

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

¹Association EURATOM-Belgian State Department of Data Analysis, Ghent University, 9000 Gent, Belgium

²Max-Planck-Institut für Plasmaphysik, EURATOM Association, 85748 Garching, Germany

³Association EURATOM-Belgian State Department of Applied Physics, Ghent University, 9000 Gent, Belgium.

* See annex of F. Romanelli et al, "Overview of JET Results",
(Proc. 22nd IAEA Fusion Energy Conference, Geneva, Switzerland (2008)).

Abstract

Reliable and accurate estimates of the ion effective charge Z_{eff} in tokamak plasmas are of key importance with respect to impurity transport studies and the establishment of thermonuclear burn criteria. These issues are of fundamental interest to ITER and reactor operational scenarios in general. However, Z_{eff} estimates derived from bremsstrahlung spectroscopy on the one hand and from the weighted summation of individual impurity concentrations obtained via Charge Exchange Spectroscopy (CXs) on the other hand, often are not compatible. This is a long-standing problem in fusion plasma diagnosis. A rigorous analysis of uncertainty sources and their propagation in the experimental determination of Z_{eff} can contribute significantly to the derivation of a Z_{eff} value with reduced uncertainty that is consistent with both the bremsstrahlung and CXs data sets. In the present work, Bayesian Probability Theory is used in an integrated approach as a powerful tool for an advanced error analysis in the derivation of Z_{eff} , even in the presence of systematic errors on the data. A simple probabilistic model is proposed for the estimation of Z_{eff} , first assuming only statistical uncertainty, next taking into account also systematic deviations. The obtained Z_{eff} estimates have smaller error bars than the Z_{eff} values derived from the individual bremsstrahlung and CXs measurements, approaching ITER requirements. The estimates are shown to be consistent with all available information. In addition, systematic errors on the data are quantized through the requirement of data consistency between different time slices in the acquired measurements.

Index Terms

Bayesian Probability Theory, Integrated Data Analysis, ion effective charge, tokamak plasmas.

I. INTRODUCTION

LARGE amounts of plasma diagnostic data are generated in magnetic confinement experiments. This information needs to be processed and validated in an efficient way in order to facilitate subsequent data interpretation. Although over the years tremendous efforts have gone in the development of increasingly sophisticated diagnostic techniques, the task of data analysis is often considered secondary. This is particularly the case for the systematic analysis of diagnostic uncertainties and the validation of measurements in the light of underlying physical models and complementary data sets. However, this consistent processing of diagnostic data can dramatically increase the reliability and robustness of physics results, and can contribute significantly to the optimization of diagnostic design. The process of the derivation of quantities of interest from the raw, unaltered measurements, is referred to as data validation and analysis. It can involve complex analysis techniques and consistency checks among diagnostics, both at the level of the raw and processed data. Since at both levels the data are affected by uncertainties of a diverse nature, a probabilistic analysis of all available information is appropriate.

The task of data validation involves, in addition, the formulation of error estimates on derived quantities of interest. The study of the propagation of uncertainty, or errors, is a well understood topic in statistics that is essential to the correct interpretation of measurements in all branches of experimental science. Significance is attached only to phenomena observed in the measured data—and derived quantities—that are unlikely to be caused by coincidence. In magnetic confinement fusion research, the diagnosis of the hot plasma poses many challenges regarding experimental design and measurement of the relevant plasma quantities. Not only the physics of the plasma itself is involved, but also the interoperation of numerous optical, mechanical and electronic components, and control and computational software (see e.g. [1]). Frequently, large error bars are associated with the physical plasma quantities of interest, indicating a corresponding substantial uncertainty. In quantizing the uncertainty affecting plasma quantities it is therefore essential to rigorously draw up an inventory of all root sources of uncertainty contributing to the total error and to study the error propagation. Here, by root sources of uncertainty we denote those that can be quantized directly, without requiring any prior analysis. However, there basically lie three difficulties in this approach. First, the root sources of uncertainty are often of a very heterogeneous nature, resulting from many different causes. We may roughly discern the following uncertainty classes:

- diagnostic design and measurement uncertainties,
- calibration uncertainties,
- physical model uncertainties,

each of which in turn can be subdivided in uncertainties of a very different origin. Model uncertainty can manifest itself, for example, in uncertainty on physical constants such as atomic data. The root uncertainties combine, often via a complex chain of calculations, in the compound errors on the quantities of interest. However, the

application of standard error propagation laws can in some cases become particularly difficult and tedious. A second difficulty is that most common studies of error propagation are conducted under the implicit assumption that not only the statistical root uncertainties, but also the errors on derived quantities are the result of Gaussian processes. An additional hypothesis of independence of the elementary observables then leads to standard Gaussian error propagation laws. However, many cases exist where the Gaussian and independence approximations are not justified [2]. A third problem in common error analysis that is frequently overlooked is the fact that usually only statistical sources of uncertainty are considered, whereas systematic uncertainties may contribute substantially to the overall error. This is an often made (silent) mistake that can sometimes be noticed implicitly in a case where the experimentalist applies simple Gaussian error propagation laws. The existence of systematic uncertainties in a particular data set is usually hard to discover. For instance, repeated measurements may help quantizing statistical uncertainties but yield no information on systematic discrepancies. In general, a more advanced error analysis is required that can handle statistical as well as systematic uncertainties, taking into account the complete error statistics through general (possibly non-Gaussian) probability distributions. This can be accomplished in a systematic and straightforward manner via Bayesian Probability Theory (BPT). In this paper we will use the BPT framework to calculate probabilities associated to physical quantities. The idea is to model all uncertainties entering the analysis in terms of a suitable probability distribution and from there construct a probability distribution of the desired physical quantities. In a Bayesian context, there exists also the opportunity to include expert prior knowledge. It is then possible to study the sensitivity of the physical quantity to the various uncertainties, and the most crucial uncertainties can be identified.

The process of data validation can be aided substantially by considering complementary data sets, originating from different diagnostics, each containing some information on the quantities of interest. Due to existing interdependencies between different diagnostics or via the requirement of data consistency at different time instants, this may contribute to the reduction of statistical uncertainty and the quantization of systematic uncertainties. This leads into the realm of (Bayesian) *Integrated Data Analysis*, or IDA, an emerging field in fusion data processing involving the combination of data from multiple diagnostics, possibly together with information on an underlying physical model [2]–[6]. It is particularly useful in the event of inconsistency between physical quantities derived from different diagnostic measurements. The redundant information present in the various measurements can help determine the physical quantities with a higher accuracy. Of further interest are more data from other diagnostics measuring quantities that bear a strong correlation with the quantities of interest. Then, through a detailed study of error quantization and propagation, using BPT as a probabilistic framework, a most plausible value can be deduced for the desired physical quantities in the light of all available information. This leads in many cases to more reliable, more complete and more robust information about this quantity. On the other hand, the available space for diagnostic set-ups at ITER and future fusion reactors will be restricted, and physical quantities will need to be assessed from a limited data set. As such, any type of available information will have to be exploited. The IDA concept provides an outstanding framework that can accomplish these tasks.

In this paper, BPT and IDA are applied to the estimation of the ion effective charge Z_{eff} in the plasma core from the combined data set comprising bremsstrahlung emissivity measurements as well as line intensity data for individual impurities obtained from Charge Exchange Spectroscopy (CX) [7], [8]. Bremsstrahlung and CX spectroscopy are two popular diagnostic techniques for determining Z_{eff} . Since the two methods rely on a different measurement principle, many of the sources of uncertainty on Z_{eff} are of a very different nature when comparing both methods. Unfortunately, this leads in most cases to a systematic inconsistency between the Z_{eff} value assessed from bremsstrahlung measurements (which we will call the *continuum* Z_{eff}) and the Z_{eff} calculated from CX impurity density measurements (which we will refer to as the *CX* Z_{eff}). This provides the motivation for the present work regarding the estimation of Z_{eff} using IDA.

The precision of the Z_{eff} estimates (continuum or CX), or any other experimentally determined plasma quantity, can be enhanced by reducing the sources of uncertainty, which in all cases should be the diagnostician's first approach. However, in many instances it is not obvious at all how to further increase one's knowledge about the involved uncertainties. For instance, current technological limits may prevent improving hardware or calibration accuracy and reliability. In some cases we might simply not be prepared to spend a—possibly disproportionate—amount of effort and time to decrease the error on the measurements. This may be the case, for example, when trying to specify and quantize hard to discover systematic uncertainties in the data. Indeed, often the experimentalist is very much aware of systematic uncertainties in his data, but he has no clear idea about where exactly this uncertainty

enters the data descriptive model (forward model), let alone that he is able to quantize the uncertainties. A way out of such a situation can be provided by the requirement of consistency between measurements performed in different time slices during data acquisition. This mode of operation will be demonstrated in the present work. In any case the diagnostician should realize that some root sources of uncertainty will always be present and that these should be dealt with effectively. At least he should be able to state a reliable error margin together with the value of the physical quantities of interest. But it would be even better to model the uncertainties in such a way that the overall error on the final quantity is reduced. Both of these possibilities lie within the reach of the IDA framework. Nevertheless, we would like to stress that the use of IDA should definitely not be regarded as a compensation for careless experimental work. Rather, the IDA approach is very useful for a proper error analysis in a well conceived experiment, considerably implemented, where a trustworthy quantitative estimate of the main root sources of statistical uncertainty is available. In addition, another important feature of the IDA framework is its transparency. IDA is by no means a black box tool that converts uncertain or inconsistent data to the ‘correct numbers’ in some obscure way, with a single push of a button. Instead, there is a clear recipe that needs to be followed, requiring well thought-out input on behalf of the diagnostician. Changes in the free parameters of the probabilistic model have a well-defined influence on the results of the analysis regarding the physical quantities of interest.

This paper aims at illustrating some of the basic IDA concepts and techniques through the estimation of a Z_{eff} value consistent with all available data. Thus we demonstrate the potential and the benefits of an IDA approach for increasing the reliability of a Z_{eff} estimate in tokamak plasmas. We advocate a stepwise approach, starting from a simple probabilistic model for the estimation of Z_{eff} (and the electron density) from continuum and CX measurements. We show that this model, together with some minimal assumptions about systematic uncertainties in the data, already allows useful inferences on a consistent Z_{eff} . In a later stage the uncertainties can be gradually modeled in more detail, which is likely to lead to a yet enhanced level of intrinsic data consistency. In our present analysis, for demonstration purposes we first do not take into account the influence of statistical nor systematic uncertainties, i.e. we perform a deterministic model inversion. Next, we model the statistical uncertainties and study the consistency of the total data set. Posterior distributions for the quantities of interest are derived via Bayes’ theorem. Finally, motivated by the observed important systematic deviation between the continuum and CX Z_{eff} , we provide a basic modeling of the systematic uncertainties as well. We present a method for inferring information on systematic discrepancy in the case where no informative prior information on the corresponding systematic uncertainties is available. This is accomplished through the demand of consistency between different time slices during a single or multiple discharges. The marginal posterior distributions are estimated through a Markov Chain Monte Carlo simulation. Our approach leads to a substantial increase of the intrinsic consistency of the data. Thus, we obtain values for Z_{eff} that are in agreement with both the bremsstrahlung and CX measurements, characterized by smaller error bars than those for either the continuum or CX Z_{eff} , while obtaining a quantitative estimate for the systematic errors as well. We hope that the illustration of IDA provided in this paper will contribute to the dissemination throughout the fusion diagnostics community of the application of IDA for enhanced data analysis.

The remainder of the paper is organized as follows. In Section II, a brief account is given of the experimental determination of Z_{eff} from continuum measurements and individual CX line intensity data. An overview of the main sources of uncertainty in the derivation of Z_{eff} is presented. Section III contains some preliminaries on Bayesian Probability Theory, Bayesian computation and Integrated Data Analysis. These are necessary for the understanding of the application of IDA to the consistent determination of Z_{eff} , as described in Section IV. Section V briefly discusses some considerations for the application of IDA at ITER, focusing on the estimation of Z_{eff} . Finally, Section VI presents some conclusions and a discussion on future work.

II. EXPERIMENTAL DETERMINATION OF Z_{eff}

The ion effective charge Z_{eff} provides a local measure for the impurity concentration in the plasma, averaged over all impurities. Z_{eff} is defined as follows:

$$Z_{\text{eff}} \equiv \frac{\sum_i n_i Z_i^2}{\sum_i n_i Z_i} = \frac{\sum_i n_i Z_i^2}{n_e}, \quad (1)$$

where the sum is over all ion species, charged Z_i with density n_i , while n_e is the electron density. There is a direct relation between Z_{eff} and the power radiated from the plasma through bremsstrahlung, as well as the dilution of

the fuel. In addition, Z_{eff} is an important quantity in the study of impurity transport. Several plasma diagnostic techniques exist, both passive and active, for the determination of Z_{eff} . Bremsstrahlung and Charge Exchange spectroscopy as well as measurements of the plasma resistivity, continuum soft-X-rays and neutron yields have been used to determine Z_{eff} [1]. In this work, we concentrate on the bremsstrahlung and CXS methods.

A. Visible bremsstrahlung spectroscopy

The most widely used method to assess Z_{eff} is by measuring the bremsstrahlung emissivity from the plasma, integrated along several lines of sight, usually in the visible part of the spectrum. Care should be taken that the continuum signal is not spoiled by atomic line radiation. This can be accomplished by selecting a narrow part of the spectrum free of line radiation through a Fabry-Pérot interference filter, a technique first applied by Kadota and co-workers on JIPP T-II in 1980 [9], see also [10], [11]. Alternatively, the continuum level may be deduced from the background of the CX spectrum. A further possibility for estimating the background bremsstrahlung emissivity from the spectrum, described in [12] and [13], is to use BPT in a mixture model approach, where part of the continuum spectrum is assumed to be spoiled by an (unknown) contribution from atomic line radiation. Once the bremsstrahlung background has been estimated, a radial bremsstrahlung emissivity profile $\epsilon_{\text{ff}}(r)$ can be reconstructed from the line-integrated measurements, e.g. through Abel inversion. When the local plasma electron density n_e and temperature T_e are known, a Z_{eff} profile can be determined from the bremsstrahlung profile through the following relation (see e.g. [14]):

$$Z_{\text{eff}} = C \frac{\sqrt{T_e} \epsilon_{\text{ff}}}{\bar{g}_{\text{ff}}(T_e) n_e^2}, \quad (2)$$

where C is a constant and \bar{g}_{ff} is the so-called Gaunt factor that includes all quantum mechanical effects and is slightly dependent on T_e .

B. Charge Exchange Spectroscopy

Charge Exchange Spectroscopy involves the observation of the (visible) radiation resulting from charge exchange reactions between plasma ions and a beam of highly energetic hydrogen or deuterium atoms. Several plasma quantities can be derived from the observation of the charge exchange radiation, which are essentially the impurity temperature and flow velocity, as well as the impurity density. The determination of local absolute impurity densities (or concentrations) is particularly difficult. The lines of sight of a CX diagnostic are directed such that they are more or less tangential to the magnetic flux surfaces at the crossing with the beam. In this case, the observed CX spectral intensity Φ_{CX} for the impurity species charged Z is related to the impurity density n_Z by

$$\Phi_{\text{CX}}(\lambda) \Delta\lambda = \frac{1}{4\pi} n_Z \frac{\Delta\lambda}{\sqrt{\pi} \lambda_D} \exp \left[-\frac{(\lambda - \lambda_0)^2}{\lambda_D^2} \sum_E \langle \sigma v(E) \rangle_{\text{CX}} \int n_b(E, s) ds \right], \quad (3)$$

where the integral is over the intersection of the line of sight with the beam profile. A summation over the beam energy E is necessary because a neutral beam generally consists of several energy components. In addition, λ is the observation wavelength, λ_D is the Doppler width of the spectral line and λ_0 is the wavelength at the line peak. $\langle \sigma v \rangle_{\text{CX}}$ is the effective emission rate for the CX line under study [15] and n_b is the local beam density, which is related to the vacuum density $n_{b,0}$ by

$$n_b = n_{b,0} \exp \left\{ - \int n_e(s) \left[\sigma_{s,e} + \sum_i \sigma_{s,Z_i} \frac{n_i(s)}{n_e(s)} \right] ds \right\}, \quad (4)$$

because the beam is attenuated through interaction with the plasma. This integral is along the beam path while $\sigma_{s,e}$ and σ_{s,Z_i} are the electron-induced and ion-induced (species Z_i) beam stopping cross-sections, respectively. Since the neutral beam is observed at the intersections with a set of lines of sight, plasma quantities that are derived through interaction with the beam are more or less localized at the intersection of the sight line with the beam volume. Therefore, in practice radial profiles can be calculated directly without the need for an inversion method. In order to derive from (3) an estimate of the density n_Z of impurity species Z , an iterative procedure is required because the local beam density depends on densities of all plasma ion species, including species Z [16]. As mentioned above, in addition to local impurity density estimates, also the bremsstrahlung emissivity can be deduced from the CX

spectra by measuring the background signal. This is a line-integrated quantity since the bremsstrahlung originates from the entire core plasma volume.

Both the bremsstrahlung emissivity data and the individual impurity density estimates employed in this work were obtained during several discharges at JET using one of the JET horizontal CX systems equipped with 12 lines of sight (see e.g. [16]–[18]). By using a periscope to view inside the plasma, the lines of sight can be made more or less tangential to the magnetic flux surfaces at the intersection with the neutral beam. On JET, the CX lines of sight are aligned on two PINIs of the neutral beam in octant 8. The spectra are obtained through a Czerny-Turner spectrometer and a two-dimensional CCD array with a typical time resolution of 50 ms. The spatial resolution is roughly 7 cm, depending amongst others on the number of PINIs used¹. Apart from the possibility to monitor carbon or beryllium and deuterium, other spectrometers have been introduced in order to measure CX lines from puffed neon and argon [19], [20].

In order to derive Z_{eff} from the individual impurity concentrations measured through CXS, we note from (1) that

$$Z_{\text{eff}} = 1 + \sum_i Z_i(Z_i - 1) \frac{n_i}{n_e}. \quad (5)$$

In practice, Z_{eff} can be approximated by limiting the sum to the most abundant impurity species. At JET, the limiter surfaces and divertor target plates are fabricated from graphite, Carbon Fiber Composites or beryllium. The by far dominant impurity in many JET discharges is therefore carbon [15]. Beryllium usually need not be considered, while oxygen contents are reduced significantly by gettering. Finally, concentrations of metallic impurities are generally of the order of 0.01%. Moreover, at the typical temperatures in most of the plasma cross-section at JET, only fully stripped carbon needs to be taken into account, while the other charge states can be neglected. If also other impurities yield an appreciable contribution to Z_{eff} , their influence can still be modeled through a systematic uncertainty, as described in Section IV. In any case, considering only fully stripped carbon suffices for illustrating the key points in this work. The active CX component due to interaction of the beam with fully stripped carbon is usually monitored by CX at the strongest CVI transition in the visible, namely $n = 8 \rightarrow 7$ at 5290.5 Å. A schematic of the JET CXS system is shown in Fig. 1, with the instrument providing the data labelled ‘KS5a’ (observation port in octant 7). Since JET is not equipped with a dedicated diagnostic for space-resolved visible bremsstrahlung measurements (by means of an interference filter), we used bremsstrahlung data from the background level of the CX spectra in the neighbourhood of 529 nm. An example CX spectrum of carbon VI from the horizontal JET CXS diagnostic is shown in Fig. 2.

C. Data set

We selected a few JET discharges to demonstrate the potential of IDA with respect to the integrated estimation of Z_{eff} . These are pulses #61346, #61347 and #61348, all part of the same physical program studying D and T fueling for ITB plasmas. For all measured quantities that are of importance in this work, error estimates are proposed by the respective diagnosticians. However, and this is the case for many measured plasma quantities in general, often the error estimates were, necessarily, established by empirical arguments, requiring a lot of assumptions, and through rough approximation. The rigorous estimation of uncertainty levels for the quantities of interest here would require a dedicated approach. Therefore, we started our analysis by relying on the error estimates that are in common use among diagnosticians [17]. However, it is not always clear whether cited error bars are meant to represent statistical uncertainty, or rather whether they are of a combined statistical and systematic nature. Furthermore, if the errors should be interpreted as statistical errors, do they signify a single standard deviation, three standard deviations, or still something different? In most cases, the errors should be understood to be caused both by statistical and systematic effects. But also in that case one can pose the question with what probability the real quantity should lie within the indicated error bars.

We nevertheless would like to mention here a few commonly used relative errors on several local plasma quantities that are of interest in this work, without worrying for the moment about any of the issues posed above. The errors are indicated in Table I, applying to JET conditions and on the magnetic axis. Taking into account only fully stripped carbon as an impurity, the relative error for the CX Z_{eff} depends on the ratio of the carbon concentration

¹Recently, an upgrade was performed of the JET core CXS system, increasing sensitivity as well as spatial and temporal resolution [18]. However, the data used in the present work were acquired before this upgrade.

to the CX Z_{eff} itself, but we took a typical value here. It should also be mentioned that the error bars can in general increase above the values stated here in the vicinity of the plasma periphery.

D. Inconsistency between Z_{eff} measurements

A long-standing issue in plasma diagnosis is the inconsistency between the estimates of Z_{eff} obtained via the various diagnostic techniques. In particular, the values for the continuum and CX Z_{eff} are often not compatible. This appears to be a general problem observed at various machines [3], [22], [23]. The value of both Z_{eff} estimates may react very differently to changing plasma conditions and also the absolute value of the continuum and CX Z_{eff} are hardly ever in good agreement, being persistently plagued by strong systematic discrepancies. We present some examples from various discharges at JET. Fig. 3a shows the Z_{eff} time trace from visible bremsstrahlung, line-averaged along a centrally viewing chord, as well as its equivalent CX line integral, reconstructed from contributions of C^{6+} , Be^{4+} and He^{2+} [3]. It can be seen that the CX Z_{eff} lies systematically below the continuum Z_{eff} . Apart from a discrepancy in absolute number between the two Z_{eff} signals, there is also a certain qualitative difference. Indeed, whereas the continuum Z_{eff} has a continuously rising tendency, the CX Z_{eff} remains basically constant after $t = 14$ s. Another example from JET is shown in Fig. 3b, where the continuum and CX Z_{eff} time traces on the magnetic axis are displayed. The continuum Z_{eff} was calculated from the inversion of the bremsstrahlung line integrals obtained from the baseline level of the CX spectra. In this discharge, only fully stripped carbon (C^{6+}) was taken into account as an impurity. There is some qualitative resemblance between both signals, but the continuum Z_{eff} on the magnetic axis is on the average a factor of 2.5 higher than the CX Z_{eff} . Since in this instance the continuum Z_{eff} is generally rather high, reaching peak values of over 9, it is clear that the Z_{eff} discrepancy here is at least partly due to a considerable overestimation of the continuum Z_{eff} . This trend will be confirmed by the analysis in Section IV. Fig. 3c shows similar behavior for the axial Z_{eff} traces (based on C^{6+}) in JET pulse #61346, which we will analyze in more detail in Section IV. Clearly, again there is a substantial systematic discrepancy between the two Z_{eff} estimates, the average ratio of the continuum Z_{eff} to the CX Z_{eff} being about 1.5. Finally, Fig. 3d shows Z_{eff} profiles for the same discharge #61346 at about 5.7 s. Especially on the magnetic axis the continuum and CX Z_{eff} are inconsistent, differing by more than a factor of two. At a normalized minor radius ρ of about 0.55 the Z_{eff} values turn out to be systematically inconsistent as well, although the ratio is different from that on the magnetic axis. Note that the continuum Z_{eff} reaches values below 1 here, which is possibly an artifact of the inversion process from the line-integrated measurements. In addition, the continuum Z_{eff} profile is clearly peaked whereas the CX Z_{eff} profile is rather flat. We wish to stress, however, that the data shown in Figs. 3b, 3c and 3d (together with all other data used in the analysis in Section IV) are the result of a routine processing of the raw measurements. It is very well possible that a dedicated analysis of the measurements, using established methods, succeeds in mitigating to a certain extent the here observed inconsistencies, by manually reducing uncertainties in the raw data (see below). Nevertheless, in order to demonstrate the power of an IDA analysis, we will continue working with the routinely processed data.

The reason for the observed inconsistency of Z_{eff} values lies in the various uncertainties that enter the derivation of Z_{eff} , for both the continuum and CX values. As regards the continuum Z_{eff} (profile), the following most important sources of uncertainty in its calculation from the raw bremsstrahlung signals can be mentioned, listed here in descending order of estimated impact on the accuracy of the Z_{eff} measurement.

- The reconstruction of radial emissivity profiles can introduce large uncertainties, particularly near the plasma periphery. It is an ill-posed inverse problem with only a few line-integrated measurements available. For regularization purposes the profile is usually expressed in terms of a set of smooth basis functions (polynomials, splines, ...) but this may yield unrealistically smooth profiles. Furthermore, uncertainty on the edge profile may propagate towards the central profile [24]. In addition, there may be poloidal asymmetries in the plasma, so that the emissivity can no longer be assumed to be constant on magnetic flux surfaces. Moreover, the reconstruction critically rests on knowledge of the magnetic equilibrium, which depends on the accuracy of the assumed equilibrium model.
- Uncertainty on the n_e and T_e profiles that are used for the calculation of the Z_{eff} profile, introduces further errors. Due to the quadratic dependence of ϵ_{ff} on n_e , especially the n_e profile should be known relatively accurately. This can be an issue, particularly near the plasma boundary. The difference in toroidal location of the respective diagnostics for the measurement of ϵ_{ff} , n_e and T_e can cause even more inaccuracies.

- A long-standing issue regarding the measurement of bremsstrahlung in the visible is the reflection of the plasma light on various surfaces inside the tokamak vessel. Reflections can be minimized by mounting a viewing dump on the wall that is seen by the diagnostic, but it is not feasible to completely eliminate the effect of reflections through adjustments of the hardware. However, by assuming a reflection model for the bremsstrahlung on the vessel surfaces it is possible to eliminate most of the influence of reflections on a derived Z_{eff} profile through appropriate manipulation of the measured signals [25].
- The spectral window that is used for the determination of the bremsstrahlung continuum should be free of atomic line radiation. This should be verified regularly, but even then it is still possible that under certain plasma conditions this requirement is no longer fulfilled. If the bremsstrahlung emissivity is determined from the background of the CX spectrum, then errors are introduced because one has to separate the background from the rest of the CX spectrum. Furthermore, the contribution of non-bremsstrahlung edge components in the continuum can greatly influence the reconstructed profile, an influence that can propagate towards the center of the profile during the inversion process. Examples of these components are recombination radiation, molecular bands and black-body radiation from hot material components inside the tokamak vessel.
- The relative channel-to-channel calibration and the absolute calibration of the bremsstrahlung system also represent a source of error.
- There may be several issues related to a suboptimal or inaccurate design of the involved diagnostic hardware. In addition, detector electronics inevitably introduce a certain amount of measurement noise.
- In the derivation of the bremsstrahlung emissivity an approximation has to be chosen for the Gaunt factor. This represents an uncertainty in the physical model, although it is estimated that the effect is modest [11].

As far as the calculation of the CX Z_{eff} is concerned, the following list provides the most critical associated uncertainties, again in descending order of possible impact.

- One of the main issues is the determination of a correct fit to the active CX spectral lines that are used to monitor the impurity species of interest.
- Again, the analysis depends on the knowledge of the magnetic equilibrium, in order to assess the geometry of the lines of sight with respect to the magnetic flux surfaces.
- Although in many plasmas, especially in the core and in particular at JET, fully stripped carbon yields by far the dominant contribution to Z_{eff} , the influence of other impurity species might not be negligible. If a certain impurity species, occurring in the plasma with a significant abundance, is not monitored by the CX system, then this may result in a Z_{eff} underestimation.
- The n_e profile enters the derivation of the CX Z_{eff} through calculation of the neutral beam attenuation, see (4). Due to the exponential dependence, uncertainties in the n_e profile will introduce uncertainty in the beam attenuation and hence in the calculated Z_{eff} . The dependence of the CX Z_{eff} on n_e turns out to be more or less linear, so the influence of n_e uncertainties on the CX Z_{eff} is in any case less important than for the case of the continuum Z_{eff} . The influence of T_e profiles is not very important.
- The relative and absolute calibration represent a further source of error.
- As is the case for the bremsstrahlung system, there may be important issues related to the design of the CX diagnostic hardware (including detector noise).
- The derivation of the CX Z_{eff} rests on the determination of atomic data (cross-sections, rate coefficients). It is however estimated that these atomic data are accurate enough to introduce only minor uncertainty in the eventual CX Z_{eff} , at least compared to the other sources of uncertainty.

We may conclude that, due to the many important uncertainties influencing both the continuum and CX Z_{eff} estimates, the determination of a consistent global Z_{eff} value, and *a fortiori* a consistent Z_{eff} profile that is reliable over the whole plasma cross-section, is at present a real challenge. The remainder of this work concerns the modeling of the uncertainties through probability distributions, in order to derive a probability distribution for Z_{eff} from which further inferences can be drawn.

III. BAYESIAN PROBABILITY THEORY

A. Bayesian probabilities

The mathematical framework that we will be using for our analysis is Bayesian Probability Theory. BPT has proven to be very effective in the solution of problems on the basis of uncertain information. Applied to experimental

physics, usually one has a set of measurements and a physical model linking the measurements with the physical quantities of interest, which are in fact model parameters. Since the measurements are uncertain, it is natural to describe them through a probability distribution, parametrized by the physical model parameters. Assessing the probability of the data given some model parameters is rather straightforward, but the inverse problem of estimating the model parameters given the uncertain data was solved only in 1763 by the reverend Thomas Bayes [26] and again in a more general form in 1812 by Laplace [27]. From this developed Bayesian Probability Theory and a clear rationale for BPT was provided by Cox [28] and Jaynes [29], who described the theory as *an extension of deductive logic to cases where there is uncertainty* [29]. BPT has many advantages over and is much more general than the traditional frequentist approach [29]–[32]. The Bayesian paradigm defines probability as a degree of plausibility of a certain proposition, not as a long-run frequency. As a consequence, whereas the frequentist definition of probability necessitates the field of statistics to infer the truth of an hypothesis, Bayesians can directly assess the probability of the hypothesis, given the available information. This renders many complex inference problems simpler to solve using BPT than via the frequentist approach.

In our current application, a physical model is assumed providing the so-called *forward model* or *data descriptive model* that links the model parameters to the measured data. The formulation of the forward model is the start of any Bayesian analysis. In its most complete form, the forward model predicts the raw, uncalibrated data as recorded by the detector(s). Our aim is to make inferences about the model parameters of interest, e.g. the effective charge, given the measured data. This is the inverse problem, which is usually much harder to solve than it is to formulate the forward model. The propositions of parameters of our physical model taking on a certain value, constitute the relevant hypotheses. The probability (density) of such propositions can be derived from the data at hand through Bayes' theorem:

$$p(\boldsymbol{\theta}|\boldsymbol{x}, I) = \frac{p(\boldsymbol{x}|\boldsymbol{\theta}, I)p(\boldsymbol{\theta}|I)}{p(\boldsymbol{x}|I)}, \quad (6)$$

where in general $\boldsymbol{\theta}$ is a vector of parameters that characterize the physical model, while \boldsymbol{x} is a vector of measurements. Furthermore, I represents any additional information that is available. This comprises, for instance, additional information on the physical system under study, i.e. the physical model that is used, or in the present application it includes the knowledge of positivity of the signals (since we are studying positive physical quantities). The factor $p(\boldsymbol{x}|\boldsymbol{\theta}, I)$ in the nominator of the right-hand side of (6) is called the *likelihood* of the parameter vector $\boldsymbol{\theta}$. To emphasize that it is a function of $\boldsymbol{\theta}$, it is sometimes written as $L(\boldsymbol{\theta})$. The likelihood is the probability (density) to observe the data vector \boldsymbol{x} , *given* a set of model parameters $\boldsymbol{\theta}$. It typically involves a difference between the observed data \boldsymbol{x} and the data calculated from the given set of parameters $\boldsymbol{\theta}$, via the assumed physical model. Thus, the likelihood describes the data misfit. The factor $p(\boldsymbol{\theta}|I)$ is called the *prior* probability (density) of the parameter vector $\boldsymbol{\theta}$, and it embodies everything we are willing to assume about the model parameters, *before* gathering the data. The latter is imperative, and an often made mistake is to use the data for constructing a prior PDF. This does not prevent, however, to employ information on data that were gathered in a *previous* experiment for defining a prior. The quantity on the left-hand side of (6) is named the (joint) *posterior* PDF. It is the PDF of interest, as it provides the probability (density) of the parameters of interest $\boldsymbol{\theta}$ of the physical model, *given* the observed data. The factor $p(\boldsymbol{x}|I)$ in the denominator of (6) is named the *evidence*. It does not depend on the parameters $\boldsymbol{\theta}$, and it is often ignored since it merely normalizes the posterior. However, this factor is important when using BPT for choosing between different models that should explain a given data set (model selection) [33]. Bayes' theorem represents the process of learning, since it tells us how to update our prior beliefs on a certain proposition as soon as new information (data) becomes available.

The choice of a suitable prior probability distribution for a certain physical parameter that accurately—and objectively—represents the *a priori* available information, is sometimes a difficult problem. As the amount of measured data increases, the influence of the prior on the posterior decreases gradually (see e.g. [32]). Hence, when a large amount of data is available, the choice of prior is not really important. However, in cases where we have a limited data set or when the data quality is poor, the prior will be all the more important. Especially in such instances, one should be cautious not to encode more information into the prior PDF than is in fact supported by the available prior knowledge. In addition, care has to be taken to construct a prior that does not lead to an *improper* posterior, i.e. one that can not be normalized because its integral over the parameter space diverges. When little or no prior information is available, this is encoded in a so-called *uninformative prior*. Many principles exist to construct uninformative priors, all leading to (slightly) different results. However, in many cases the uniform

distribution is a good choice for an uninformative prior and we will employ it in this work as well. Finally, a prior distribution may also depend on other unknown parameters, called *hyperparameters*. In turn, these hyperparameters can be assigned a prior distribution (hyperprior), possibly parameterized by hyperhyperparameters, and so on. Thus, one creates a so-called *hierarchical* probabilistic model, possibly consisting of several levels [32].

As noted above, $p(\boldsymbol{\theta}|\mathbf{x}, I)$, where $\boldsymbol{\theta} \equiv (\theta_1, \dots, \theta_p)^T$, is the *joint* posterior density for all parameters of the physical model. However, we typically want to make inferences about individual parameters $\theta_k, k = 1, \dots, p$. For instance, the posterior that we will use in Section IV for the estimation of Z_{eff} is a joint distribution for both Z_{eff} and n_e . Hence, to deduce the distribution for Z_{eff} alone, we will have to marginalize the joint posterior. In fact, the posterior often depends also on parameters we are not even interested in, but that necessarily enter the forward model. These are called *nuisance parameters*. In order to obtain the PDF for an individual parameter θ_i , we have to *marginalize* the joint PDF, i.e. integrate out the other parameters:

$$\begin{aligned} p(\theta_k|\mathbf{x}, I) &= \int p(\boldsymbol{\theta}|\mathbf{x}, I) d\theta_1 \dots d\theta_{k-1} d\theta_{k+1} \dots d\theta_p \\ &\sim \int p(\mathbf{x}|\boldsymbol{\theta}, I) p(\boldsymbol{\theta}|I) d\theta_1 \dots d\theta_{k-1} d\theta_{k+1} \dots d\theta_p. \end{aligned}$$

The integral is over the complete parameter space spanned by the parameters $\theta_1, \dots, \theta_{k-1}, \theta_{k+1}, \dots, \theta_p$ and it may also be a sum in the case of discrete parameters. In addition, in order to characterize the posterior PDF we might want to calculate its moments. To do this, we also need to integrate the posterior over the model parameters. The integrations can be very difficult to perform, which leads into the field of Bayesian computational methods.

B. Bayesian computational methods

As mentioned in the previous subsection, in BPT often integrals need to be calculated of the joint posterior distribution of all model parameters. This is necessary to obtain the marginal posterior probability distribution for the individual quantities of interest or to integrate out the nuisance parameters. Furthermore, the posterior represents the most complete description of the quantities of interest. However, usually when mentioning experimentally obtained information on a physical quantity, one is not interested in the full associated probability distribution but rather in some summary of the distribution. Often this summary comes in the form of a ‘best estimate’ for the quantity of interest together with some measure of uncertainty, but it may include any kind of summarizing item such as skewness or kurtosis. The decision on which quantities to use as a best estimate and uncertainty measure, depends on the problem at hand. However, when the distribution is unimodal (has a single local maximum) and is more or less symmetric, the mode of the distribution (i.e. the value of the variable where the probability density reaches a maximum) can be taken as a best estimate. The mode of the posterior distribution is referred to as the *Maximum a Posteriori* (MAP) estimate for the parameters of interest (for uniform priors, this comes down to the maximum likelihood estimate). If the distribution is highly skewed, the mean or the median usually provide better summarizing items for the posterior. Uncertainty bounds (often iso-probability contours) are best defined based on the amount of probability that they include, i.e. one calculates credible intervals (the Bayesian analogue of classical confidence intervals). Depending on the skewness of the distribution, the iso-probability contours may be asymmetrically distributed around the mode. Alternatively, for more symmetrical distributions the covariance of the joint posterior or the variance of the marginal posteriors may also be quoted as uncertainty measures. The point here is that usually moments of a distribution have to be derived in order to summarize the distribution, which again involves integrating the distribution.

Now, integration of the probability distributions can be difficult when the distribution is intractable, in which case one needs to resort to numerical integration methods. However, as the number of model parameters—hence the dimensionality of the integration space—rises, deterministic numerical integration becomes very slow compared to stochastic integration techniques. This is where the use of Bayesian computational methods comes in. Simulation techniques such as *Markov Chain Monte Carlo* (MCMC) are particularly popular.

1) *The origins of MCMC*: MCMC methods are very popular for drawing samples from (marginal) posterior distributions. The origins of MCMC lie in an illustrious paper by Metropolis *et al.* on the computation of high-dimensional integrals in statistical mechanics [34]. The method was later generalized and improved by Hastings [35], who viewed the Metropolis algorithm chiefly as a way to sample from high-dimensional probability distributions, which is precisely its primary modern use. Hastings noted that the Metropolis method involved the transition matrix

of a Markov chain. He expressed the target distribution to sample from in terms of the invariant distribution π of the Markov chain. This scheme is presently known as the *Metropolis-Hastings* (M-H) algorithm, one of several existing MCMC methods.

In the early 1990s, another popular MCMC method, known as *Gibbs sampling*, was proposed, based on the work of Gelfand and Smith [36], who built on the paper of Geman and Geman [37]. The Gibbs sampler is a related simulation algorithm that is especially useful for sampling multivariate distributions, particularly when the univariate conditional densities are known, or are easy to sample from. Later, it was shown that the Gibbs sampling algorithm is in fact a special case of Metropolis-Hastings [38]. The history of the M-H algorithm is summarized in [39]. Further details on MCMC methods can be found in [40]–[44].

2) *The principle of Monte Carlo integration*: The principle of Monte Carlo integration is very simple. For simplicity, we work in one dimension where we want to evaluate the following integral over some domain \mathcal{D} :

$$\int_{\mathcal{D}} h(\theta)f(\theta)d\theta. \quad (7)$$

To do this, we can view $f(\theta)$ as a (not necessarily normalized) probability density, so that (7) becomes the expected value of h under f : $\mathbb{E}_f(h)$. We now generate a (large) sample $\theta_1, \dots, \theta_n$ from f . Then, we can approximate (7) by the empirical average

$$\bar{h}_n = \frac{1}{n} \sum_{j=1}^n h(\theta_j). \quad (8)$$

Indeed \bar{h}_n converges almost surely to (7) due to the Strong Law of Large Numbers. As a special case, by setting $h(\theta) = \theta, \theta^2$, etc., we can approximate the moments of the distribution f , explaining the relevance of Monte Carlo integration in the current context.

3) *The principles of MCMC*: The key idea behind MCMC for sampling from arbitrary probability distributions, is to set up an ergodic Markov chain of random variables $(\Theta^{(i)})$, whose limiting distribution is the target distribution of interest π . Here, $(\Theta^{(i)})$ refers to state i of the Markov chain². If π is a multivariate distribution, $(\Theta^{(i)})$ consists of multiple components $\Theta_k^{(i)}$ and the Markov chain moves in a multidimensional space. The chain's ergodicity ensures various properties of convergence and stability [40]. Inferences drawn from running an MCMC simulation are usually summarized in terms of *ergodic averages* of the form

$$\bar{g}_n \equiv \frac{1}{n} \sum_{i=1}^n g(\Theta^{(i)}), \quad (9)$$

where g is a—possibly vector-valued—function of interest of the chain states and n is the current number of samples drawn. This represents the Monte Carlo aspect of MCMC. The so-called *Ergodic Theorem* is the Law of Large Numbers for Markov chains and guarantees the convergence of ergodic averages [40]. More precisely, the Ergodic Theorem states that if the Markov chain $(\Theta^{(i)})$ is ergodic, then

$$P \left(\bar{g}_n \rightarrow \int g(\theta)\pi(\theta)d\theta \right) = 1,$$

where P denotes probability. For example, if we take $g(\theta) \equiv \theta$, then \bar{g}_n converges surely to the mean of the stationary (target) distribution π , and similarly for other moments of the distribution. In practice, for an arbitrary starting value $\theta^{(0)}$, a chain $(\Theta^{(i)})$ is generated using a transition kernel with stationary distribution π , which ensures the convergence in distribution of the chain to a random variable from π . If the chain is ergodic, the starting value $\theta^{(0)}$ is (in principle) not important.

The question remains how to set up the Markov chain. One way to do this is through the M-H algorithm that draws samples from approximate distributions, and then corrects those draws to better approximate the target distribution. The samples (state vectors) are drawn sequentially and the distribution of the draws depends only on the last value drawn, forming a Markov chain. In the M-H algorithm, the transition from a previous state $\Theta^{(t-1)}$ to the current state $\Theta^{(t)}$ of the Markov chain may occur as follows (although there are other possibilities). Consecutively for all components k of the state vector a new value ϕ_k is proposed. The candidate value of the

²It is common practice in probability theory to denote random variables by capital letters and their realizations by lower-case letters.

component is sampled from a univariate conditional distribution, called the *proposal* or *instrumental distribution* $q(\phi_k | \theta_1^{(t)}, \dots, \theta_{k-1}^{(t)}, \theta_k^{(t-1)}, \theta_{k+1}^{(t-1)}, \dots, \theta_p^{(t-1)})$, which is easy to sample from. Note that the new values $\theta_j^{(t)}$ for $j < k$ are immediately used in the current step of the chain for sampling the value ϕ_k , the proposal for $\theta_k^{(t)}$. Henceforth we will use the notation for the current parameter vector: $\theta_c \equiv (\theta_1^{(t)}, \dots, \theta_{k-1}^{(t)}, \theta_k^{(t-1)}, \theta_{k+1}^{(t-1)}, \dots, \theta_p^{(t-1)})$. Now, the new value ϕ_k is accepted or rejected according to a specific criterion, given below. The target density π must also be available up to some extent: a general requirement is that the ratio

$$\frac{\pi(\phi)}{q(\phi_k | \theta)}, \quad \text{with } \phi \equiv (\theta_1^{(t)}, \dots, \theta_{k-1}^{(t)}, \phi_k, \theta_{k+1}^{(t-1)}, \dots, \theta_p^{(t-1)}),$$

is known up to a constant *independent* of θ , for all k . This is crucial in Bayesian simulation, since we often only know the target posterior distribution up to a multiplicative factor (namely the evidence). It is important to note that knowledge of the functional form of the target density π does *not* automatically allow to generate samples from π ; hence the use of MCMC.

The M-H algorithm can be summarized in pseudocode as follows.

Algorithm 1 (Metropolis-Hastings)

- 1: Given $\theta^{(t-1)} = (\theta_1^{(t-1)}, \theta_2^{(t-1)}, \dots, \theta_p^{(t-1)})$
- 2: **repeat**
- 3: **for** $k = 1 \rightarrow p$ **do**
- 4: Generate
 $\Phi_k \sim q(\phi_k | \theta_1^{(t)}, \dots, \theta_{k-1}^{(t)}, \theta_{k+1}^{(t-1)}, \dots, \theta_p^{(t-1)})$
- 5: Let $\Phi \equiv (\theta_1^{(t)}, \dots, \theta_{k-1}^{(t)}, \Phi_k, \theta_{k+1}^{(t-1)}, \dots, \theta_p^{(t-1)})$
- 6: Take

$$\Theta_k^{(t)} = \begin{cases} \Phi_k & \text{with probability } \rho(\theta_c, \phi) \\ \theta_k^{(t-1)} & \text{with probability } 1 - \rho(\theta_c, \phi) \end{cases},$$

where

$$\rho(\theta_c, \phi) \equiv \min \left\{ \frac{\pi(\phi)}{\pi(\theta_c)} \frac{q(\theta_k | \phi)}{q(\phi_k | \theta_c)}, 1 \right\}$$

- 7: **end for**
- 8: **until** Satisfied

The probability $\rho(\theta_c, \phi)$ is called the *Metropolis-Hastings acceptance probability*. Only in the symmetric case, where $q(\theta_k | \phi) \equiv q(\phi_k | \theta_c)$ does the acceptance depend solely on $\pi(\phi)/\pi(\theta_c)$, in which case the algorithm is simply called the *Metropolis algorithm*. In the general case, if the ratio $\pi(\phi)/q(\phi_k | \theta_c)$ is increased by the proposed value ϕ_k , the proposal will be accepted. Hence, the algorithm preferentially samples in regions of probability space where most of the probability density of π is concentrated.

Thus, instead of sampling directly from the target distribution, which might be very difficult, or, indeed, analytically impossible, the M-H algorithm samples from the proposal distribution *and then introduces a correction*, such that the Markov chain's distribution eventually converges to the actual distribution of interest. In order to allow the Markov chain to efficiently converge to the target distribution, one has to run the chain for an initial period, called the *burn-in*. Once convergence has been reached, the M-H algorithm primarily generates samples from the full, possibly multivariate, target density π . In addition, and most relevant for the application we have in mind, it is not difficult to prove that every component subsequence $\Theta_k^{(i)}$, for all k , also forms a Markov chain with limiting stationary distribution the corresponding marginal distribution

$$\pi_k(\theta_k) \equiv \int \pi(\theta) d\theta_1 \dots d\theta_{k-1} d\theta_{k+1} \dots d\theta_p.$$

This means that, in effect, MCMC is also a stochastic integration method by providing samples from the respective marginal distributions [40]. The ergodic averages (9) thus enable to calculate moments of the marginal distributions

as well. Only the post-convergence samples are used for the Monte Carlo estimation of quantities. To monitor convergence of the chain, one studies the acceptance probability and time traces of the quantities of interest.

The transitions of the chain depend on q , which in turn depends on the current parameter vector θ_c . However, in many applications the proposal distribution only depends on $\theta_k^{(t-1)}$, not on the other component values. Transitions should not be taken too small, because we want the support of the target distribution to be explored efficiently. On the other hand, neither should the transitions be taken too large, because then the chain's acceptance probability will become too small.

The samples that are generated by the M-H algorithm are not mutually independent. For one thing, there may be repeated occurrences of the same value. However, this poses no real problem since we are merely interested in the convergence of empirical averages (9) to the quantity $\mathbb{E}_\pi[\mathbf{f}(\Theta)]$, which is ensured by the Ergodic Theorem since the Metropolis chain is ergodic.

For an M-H simulation to be efficient, on the one hand the acceptance probability throughout the simulation should be sufficiently high, but on the other hand the Markov chain still has to be able to move through the entire support of the target distribution π . A *random walk* chain is a common choice that satisfies these requirements. A Markov chain performs a random walk when candidate components Φ_k are given by

$$\Phi_k = \Theta_k^{(t-1)} + \epsilon_t,$$

where ϵ_t is a (small) perturbation with distribution g , independent of $\Theta_k^{(t-1)}$. The random walk Metropolis-Hastings algorithm thus samples from the proposal distribution centered at the current value of the Markov chain: $q(\Phi_k | \theta_k^{(t-1)})$ is of the form $g(\Phi_k - \theta_k^{(t-1)})$. Common distributions for g include the uniform, the normal, Student's t and the Cauchy distributions. The Gibbs sampler arises as a special case of the M-H algorithm when samples from the univariate conditional distributions $\pi_k(\theta_k | \theta_1^{(t)}, \dots, \theta_{k-1}^{(t)}, \theta_{k+1}^{(t-1)}, \dots, \theta_p^{(t-1)})$ can be readily obtained at every time step t [40]. No proposal distribution is needed and none of the samples is rejected, resulting in an acceptance probability uniformly equal to unity.

C. Integrated Data Analysis

The idea behind IDA is to consider information from several heterogeneous sources, be it theory, calibration data or measurements from experiments that are either necessary for the calculation of the physical quantities of interest, or complementary to already available measurements. The latter can be very useful in situations where individual experiments provide diffuse or even conflicting information on the physical quantities. Combined with a proper analysis using BPT, this can result in much more accurate results that are compatible with all available data. At the same time, the integrated analysis permits to quantify systematic uncertainties in the data and to study their effect on the parameters of interest. This is exactly the situation one faces when comparing Z_{eff} estimates from bremsstrahlung spectroscopy with the results from CXS. The IDA framework has proven its usefulness in the diagnosis of fusion plasmas in [2], [4], where it was used for calculating electron density and electron temperature profiles from Thomson scattering data and interferometry measurements. In [45] measurements from a Thomson scattering system, interferometer, diamagnetic loop and neutral particle analyzer were combined with an equilibrium reconstruction, modeling the diagnostic interdependencies using Bayesian graphical models. Recently, lithium beam emission spectroscopy was combined with interferometry data to obtain full electron density profiles. Together with hardware improvements, this allowed to increase the temporal resolution by a factor of 400 compared to classical inversion techniques [46]–[48].

An inherent feature of the IDA approach is that it can be conducted through several stages, gradually increasing the complexity of the probabilistic model. Hence, depending on the desired level of sophistication, one can start the analysis with a simple model, describing the uncertainties only roughly. One can then compare the estimates for the quantities of interest and their uncertainty with the classical approach. As such, it is imperative to first identify the major sources of uncertainty in the quantities of interest, given the measured data, and rank these uncertainties in an order of decreasing importance. Then, the uncertainties have to be formally quantified in terms of a suitable probability distribution. Statistical uncertainties in the data (measured data and calibration data) have to be modeled through likelihood PDFs, describing the error statistics of the measurement. Systematic uncertainties can be modeled by introducing nuisance parameters, possibly in a hierarchical model using hyperparameters. The nuisance parameters subsequently have to be integrated out (marginalized). Uncertainties in the physical model

itself can be taken into account by introducing and marginalizing proper model flexibility. Although also single experimental set-ups can be analyzed in this way, we here consider mainly the integration of multiple diagnostics.

1) *The IDA recipe*: We now outline the basic recipe that can be used for the Bayesian integrated analysis of (tokamak) diagnostic data. First, estimates have to be determined for the statistical measurement errors. The effect of the statistical uncertainties is then modeled into an appropriate likelihood PDF. Often the normal distribution is selected here; a consequence of the Central Limit Theorem since the error is assumed to arise as a result of many small, independent effects of a stochastic nature. The sum of the associated random variables then approaches a Gaussian distribution. In a first stage, when little is known about the exact nature of the various uncertainties, their contribution to measurement uncertainty can be gathered into only a few likelihood PDFs. Likewise, only a few general nuisance parameters can be used to describe the effect of all systematic uncertainties. In a later stage, the contribution of each individual uncertainty can be modeled in more detail. This is the approach taken in the present work, and it is a natural way of working, if one considers that virtually *any* uncertainty is made up of multiple ‘sub-sources’ of uncertainty.

The next step is to formalize any additional available information into prior PDFs. This could involve, for instance, the formulation of a prior distribution on the nuisance parameters describing the systematic uncertainties. The information required to do this can often be obtained from additional (e.g. laboratory) measurements, see e.g. [4]. If one wants to make the least assumptions possible, one has to use uninformative priors.

Then all PDFs must be combined according to Bayes’ theorem. It specifies how to integrate all relevant information in a probabilistic framework in order to infer the quantities of interest. The nuisance parameters, if any, have to be marginalized, integrating out systematic effects and physical model uncertainties.

Eventually, one ends up with the joint posterior PDF of the quantities of interest. One can then sample from the marginal distributions for each of the parameters, or calculate the marginal moments. In the present context, this (together with the marginalization of nuisance parameters) is mainly done using MCMC methods. Characteristics such as the mode or mean of the marginal distributions can then be used as best estimates for the respective parameters, and credible intervals can be used to construct statistical error bars.

The shape of posterior distributions allows to detect inconsistencies between different diagnostics. This can be signaled, for example, by a multimodal posterior distribution (with multiple local maxima). On the other hand, it is clear that if the marginal posterior distribution for a specific physical quantity, derived from measurements by diagnostic A, is sufficiently similar to the marginal distribution of the same quantity derived from measurements by diagnostic B, then both diagnostics can be regarded as mutually consistent with respect to the information they yield on that specific quantity of interest. If this is not the case, then the diagnostics produce inconsistent results. Thus, by measuring the similarity of the distributions, one may study the compatibility of diagnostic data with respect to a certain physical quantity. A popular similarity measure for probability distributions that may be used in this context is the Kullback-Leibler divergence [49]. The incompatibility of diagnostic data is usually caused by not taking into account a systematic uncertainty or by an oversimplification of the probabilistic model.

Finally, it is possible to study the impact of various experimental and model uncertainties on the marginal posteriors for the physical quantities, hence on their most plausible value and error estimates. This process is termed a *sensitivity analysis*. In practice, it involves reducing, or entirely switching off, individual root uncertainties and observing the effect on the posterior PDF. By manipulating statistical root uncertainties, one can study their impact on the final physical quantities and their uncertainty. In addition, once the foregoing IDA analysis has revealed estimates for the involved systematic uncertainties, one may investigate their influence on the (in)compatibility of the various diagnostic data sets, by measuring the Kullback-Leibler divergence between the marginals resulting from the individual measurement sets, while varying the systematic uncertainties. This way, it is possible to detect the most crucial sources of uncertainty a posteriori. Once the most critical uncertainties have been identified, one can attempt to reduce them. Thus, IDA permits the optimization of experimental set-ups. Here, also the techniques of *Bayesian Experimental Design* can be of assistance, enabling even the optimization of future diagnostics. The idea is to maximize a utility function that quantifies the information gain on the quantities of interest as measurements are being taken. More information on this scheme can be found e.g. in [50].

2) *Unambiguous forward model invertibility*: An obvious but often instructive question that can be posed at the start of any Bayesian error analysis, is the following: what results does one obtain on the quantities of interest when no uncertainty on the measurements is assumed? Mathematically, this amounts to solving the forward model for the quantities of interest and we will demonstrate this initial approach when we will apply IDA to the estimation

of Z_{eff} in the next section. Naturally, for a unique solution to exist for the inverse problem, the forward model needs to obey certain regularity conditions. Let us illustrate with a simple example what happens when trying to make inferences in systems that are not that well-behaved. Suppose the forward model has the following form:

$$x = \theta_1 \theta_2,$$

where x represents a single, infinitely precise measurement depending on two parameters θ_1 and θ_2 , assumed to be positive. Clearly it is not possible to make useful inferences on, say, θ_1 without additional information on θ_2 . Rather, there is an infinite number of solutions, depicted in Fig. 4a. Next, suppose there is some uncertainty on the measurement x , represented, for example, by an additive Gaussian noise term ν with mean zero and standard deviation σ_ν :

$$x = \theta_1 \theta_2 + \nu, \quad \text{with } \nu \sim \mathcal{N}(0, \sigma_\nu^2).$$

As a consequence, the distribution of x given θ_1 and θ_2 , i.e. the likelihood of θ_1 and θ_2 , is also a normal distribution, with mean $\theta_1 \theta_2$ and standard deviation σ_ν :

$$L(\theta_1, \theta_2) = p(x|\theta_1, \theta_2) = \frac{1}{\sqrt{2\pi}\sigma_\nu} \exp\left[-\frac{(x - \theta_1 \theta_2)^2}{2\sigma_\nu^2}\right].$$

This likelihood, a function of θ_1 and θ_2 , is plotted in Fig. 4b for arbitrary values $x = 5$ and $\sigma_\nu = 1$. Naturally, since the forward model does not permit unambiguous determination of both θ_1 and θ_2 , also the case with an uncertain measurement will not result in useful inferences on θ_1 and θ_2 simultaneously. To solve this problem, the forward model could be expanded to contain a measurement yielding additional information on θ_1 and θ_2 , e.g. through a second measurement y , possibly given by $y = \theta_1/\theta_2$. Another possibility is to include useful prior information on θ_1 or θ_2 . In any case, before starting the IDA analysis it is always a good idea to first check whether the forward model is invertible (e.g. assuming infinitely precise measurements) such that the parameters of interest can be obtained unambiguously, possibly aided by the available prior knowledge. In the simple example we provided here, it is obvious that the available information is not sufficient to deduce the parameters of interest, but in the case of a more complicated forward model it might be less discernible. Worse still, when the inverse problem can not be solved unambiguously, if the analysis is conducted further, deriving the joint and respective marginal posterior distributions for the parameters of interest, the problem is often not readily apparent. Let us illustrate the danger by pursuing the Bayesian analysis using the present example. Suppose we know that $0 \leq \theta_1 \leq 10$ and $1 \leq \theta_2 \leq 10$, but we possess no further (prior) information. Then we might choose a uniform prior distribution that is non-zero between 0 and 10 for θ_1 and non-zero between 1 and 10 for θ_2 . The marginal posterior distribution for θ_1 is plotted in Fig. 4c. Nothing in this plot reveals that still the system under study is indeterminate. In fact, as we will experience in the next section, the curve resembles the shape of a distribution one often gets in cases where the inverse problem is solvable unambiguously. However, although the prior information provides the ranges of possible values for θ_1 and θ_2 , it is, for example, not possible to suggest any single most probable value for θ_1 and θ_2 at the same time. Nevertheless, the marginal distribution in Fig. 4c creates the impression that, on the contrary, there *is* a single ‘most probable’ value for θ_1 (where the probability density is the highest), given by the mode of its marginal posterior at about $\theta_1 = 0.65$. At the same time, a similar argument holds for θ_2 . This is of course nonsense and just an artifact of our choice of integration boundaries during marginalization.

IV. APPLICATION OF IDA TO THE DERIVATION OF A CONSISTENT Z_{eff}

We will now apply the IDA concepts outlined in the previous section to the estimation of Z_{eff} from the integrated data set consisting of measurements of the bremsstrahlung emissivity on the one hand and individual impurity concentrations obtained via CXS on the other hand. In the present approach to the problem of the integrated determination of Z_{eff} , we estimate the value of Z_{eff} locally, usually on the magnetic axis of the plasma. Hence, we employ the local bremsstrahlung emissivity, obtained via Abel inversion of line-integrated emissivities, as well as the local concentration of fully stripped carbon C^{6+} deduced from CXS reactions. We should note, however, that within the IDA framework it is entirely possible to include the line integration of local quantities into the forward model and infer estimates of local quantities (radial profiles) from the line-integrated data. This makes the use of an Abel inversion superfluous. For reasons of simplicity we will not follow this route at present. The influence from impurity species other than fully stripped carbon will be neglected, as motivated previously, or taken

into account in the form of a systematic uncertainty (Section IV-C). We propose a very simple forward model and we work through three stages of sophistication. First, we assume that the data are not affected by any statistical nor systematic uncertainty. Next, we include some statistical uncertainty and finally we introduce both statistical uncertainty and a rough systematic uncertainty.

In a more complete study of the uncertainties involved in the integrated estimation of Z_{eff} , one should start from the raw measurements as registered by the respective detectors. In this paper we will conduct a less rigorous analysis by considering as the measured data two artificially constructed plasma quantities, ϵ and δ , defined below. Both ϵ and δ are the result of a long chain of calculations, involving many uncertainties. We will summarize these uncertainties in only a few statistical and systematic parameters. Concretely, suppose we have calculated a value, at a certain time in a certain discharge on the magnetic axis, for each of the quantities ϵ and δ , defined as follows:

$$\epsilon = \frac{\epsilon_{\text{ff}} \sqrt{T_e}}{C \bar{g}_{\text{ff}}}, \quad (10)$$

$$\delta = \sum_i Z_i (Z_i - 1) n_i. \quad (11)$$

The sum is over all impurity species i . Taking into account only C^{6+} as the dominant impurity species in the core plasma, we get $\delta = Z_C (Z_C - 1) n_C = 30 n_C$. This means that ϵ is related to the measured local (magnetic axis) bremsstrahlung emissivity, while δ can be derived from the measured local CXS C^{6+} density. Now, from (2) and (5), at every time instant ϵ and δ are connected to the local n_e and Z_{eff} through

$$\begin{cases} \epsilon = n_e^2 Z_{\text{eff}} \\ \delta = n_e (Z_{\text{eff}} - 1) \end{cases}. \quad (12)$$

This constitutes our forward model. Hence, apart from Z_{eff} , we will also consider n_e as a quantity of interest. Indeed, the density is affected by some amount of uncertainty as well. Thus, due to the strong correlation of n_e with Z_{eff} , both quantities should be estimated jointly. From the ‘measurements’ ϵ and δ , possibly considered to be affected by statistical or systematic uncertainty, we want to derive estimates for n_e and Z_{eff} . Note that the quantity δ has an implicit dependency on n_e through the beam attenuation. Therefore, part of the uncertainty on δ is made up of uncertainty on whatever measurement we use for the electron density. In addition, we can compare the electron density estimate from the model (12) to the density measured by other diagnostics, for example through Thomson scattering or interferometry. At a later stage, we will also explicitly incorporate electron density measurements from these diagnostics, where we will cover the statistical uncertainties on the associated density measurements. Because of the correlation structure between Z_{eff} and n_e , this is likely to increase the accuracy of not only the density estimates, but also the Z_{eff} values.

A. Deterministic model inversion

The model (12) is a system of two equations, with two data (ϵ and δ) and two unknowns (n_e and Z_{eff}). If we assume that no measurement error was made, it turns out we can solve this system exactly. Hence, we presume also a zero error on n_e and Z_{eff} . The solution can already teach us a lot about the more realistic problem, where measurement error is present. After some simple algebra, we find the following solutions of (12):

$$\begin{cases} n_{e,\text{inv}} = \sqrt{\frac{\epsilon}{Z_{\text{eff},\text{inv}}}} \\ Z_{\text{eff},\text{inv}} = \frac{2\epsilon + \delta^2 \pm \delta \sqrt{4\epsilon + \delta^2}}{2\epsilon} \end{cases}. \quad (13)$$

It is easy to see that the solution for Z_{eff} with the minus sign in (13) is always smaller than 1, so we will no longer consider it, nor its associated solution for n_e .

We will calculate the solutions (13) from measurements of ϵ and δ for two JET discharges where only fully stripped carbon was monitored by the CX system. In Section IV-C we will describe the influence from other impurity species in terms of systematic uncertainties on the data. In order to better visualize the estimation results, we work with *time traces* of measurements ϵ and δ , and we also estimate n_e and Z_{eff} time traces. It should be noted that still the number of input data is two, since we will process all time points independently.

We selected the JET discharges #61346 and #61348, the former for which the time evolution of the on-axis continuum and CX Z_{eff} was already shown in Fig. 3c. For completeness, Fig. 5a shows the temporal evolution for discharge #61346 of the on-axis n_e , T_e , ϵ_{eff} and n_C , during operation of the neutral beam in JET octant 8. The signals for the derived parameters ϵ and δ are also depicted. Fig. 5b gives a similar overview for pulse #61348. The error bars shown correspond to the relative errors given in Table I. The error bars on ϵ and δ were calculated, via (10), assuming that the uncertainties on the other quantities are of a pure statistical nature, corresponding to a single standard deviation. We will have to come back to this assumption later on. Simple Gaussian error propagation was used as a first approximation. This yields 20% for the relative error on ϵ and 35% for δ . Fig. 6a now shows for pulse #61346 the solution $n_{e,\text{inv}}$ obtained from (13), and the comparison with the measured n_e signal. The density measurements at JET we use in our work are derived from LIDAR Thomson scattering, cross-checked with interferometry measurements. In short, we will refer to this measurement as the ‘LIDAR’ measurement. The signal for $n_{e,\text{inv}}$ is systematically higher than the LIDAR density signal. Fig. 6b shows the continuum and CX Z_{eff} signals, as well as the solution $Z_{\text{eff,inv}}$. The solution lies below the continuum and CX Z_{eff} , hence closer to the CX Z_{eff} . This seems rather unlikely and the reason will turn out to be that our model does not take into account systematic uncertainties (see Section IV-C). Similar remarks are valid for JET pulse #61348. Observing Figs. 6c and 6d, it can be noticed that the solution $n_{e,\text{inv}}$ for the density again does not match the LIDAR density, whereas the solution for Z_{eff} is relatively close to the CX Z_{eff} . The figures also show the error bars for the signals, the width of which we define again as a single standard deviation. Once more, the error bars were calculated by Gaussian error propagation analysis from the errors on ϵ and δ . The relative error for the solution $n_{e,\text{inv}}$ is about 17%: considerably larger than for the LIDAR density (5%). The error bars for the two n_e estimates in general do not overlap two-by-two, indicating a consistency problem (see Section IV-B3). The relative error for the calculated Z_{eff} from model inversion is about 25%, the same as for the continuum Z_{eff} , and slightly higher than the error bar on the CX Z_{eff} . The error bars for the CX and inverted Z_{eff} estimates overlap two-by-two, but generally not with the error bars on the continuum Z_{eff} . In any case, because of the clear failure of the model to predict the LIDAR n_e measurements (due to the systematic discrepancies), we cannot accept the present solutions of both n_e and Z_{eff} .

Concluding, the direct inversion of the system of equations (12) is an instructive exercise, but we have no measure indicating that these estimates are any better than either the continuum Z_{eff} or the CX Z_{eff} values we had before. To improve on this, we need to carry out a better modeling of the uncertainties on the data.

B. Statistical uncertainties

1) *Model definition:* As a next step, in order to make the model somewhat more realistic, we will consider some uncertainty on the data *before* inverting the model. We again took two measurements ϵ and δ , and we wish to estimate from this n_e and Z_{eff} , consistent with both measurements. To begin with, we assume that all uncertainties on the data are of a statistical nature, not systematic. We model the statistical uncertainties entering the data through two uncertainty terms ν_ϵ and ν_δ :

$$\begin{cases} \epsilon = n_e^2 Z_{\text{eff}} + \nu_\epsilon \\ \delta = n_e(Z_{\text{eff}} - 1) + \nu_\delta \end{cases}, \quad (14)$$

where $\nu = [\nu_\epsilon, \nu_\delta]^T$, which, by the Central Limit Theorem, we assume to be distributed independently identically Gaussian:

$$\nu \sim \mathcal{N}(0, \Sigma_\nu), \quad \Sigma_\nu \equiv \begin{bmatrix} \sigma_\epsilon^2 & 0 \\ 0 & \sigma_\delta^2 \end{bmatrix}.$$

Here, σ_ϵ and σ_δ denote the respective standard deviations. The fact that we take the noise on ϵ to be independent of the noise on δ , means that the measurements should have been performed with two different instruments. In our case, this is not really true, as we used bremsstrahlung measurements from the background of the CX spectrum. However, since both ϵ and δ are not raw data, we expect the error statistics to be quite different. Therefore, we will assume no noise correlation as yet, an assumption that can be relaxed in the future. The standard deviations σ_ϵ and σ_δ describe only the statistical errors on the data. As such, to define their values we should not accept the previously quoted relative errors (20% on ϵ , 35% on δ), since these take into account any systematic discrepancies as well. Instead, to estimate the statistical error on ϵ and δ , assumed to be normally distributed, we fit a Gaussian to a sample of ϵ (δ) values using Maximum Likelihood. These values were obtained from a plateau region in a

discharge, where both ϵ and δ were assumed to remain stationary. Hence, the variability of ϵ (δ) in this sample was presumed to be due only to statistical fluctuation (fixed mean). The thus acquired sample standard deviations resulted in a relative error roughly given by 15% for ϵ and 10% for delta. Henceforth, we will use these numbers, which we checked for several discharges, for the relative errors on ϵ and δ . A similar procedure gave a relative statistical error of 5% on n_e . This is the same value as mentioned in Table I, suggesting that the uncertainty on n_e is entirely due to statistical fluctuations. We will adopt this point of view for the remainder of our analysis.

Following the IDA recipe, we now construct the forward model in terms of appropriate likelihood functions. Since the noise terms ν_ϵ and ν_δ both have a normal distribution and are independent, we can write the likelihood for the two unknown parameters n_e and Z_{eff} as:

$$\begin{aligned} L(n_e, Z_{\text{eff}}) &= p(\epsilon, \delta | n_e, Z_{\text{eff}}, I) \\ &= \frac{1}{\sqrt{2\pi}\sigma_\epsilon} \exp\left[-\frac{(\epsilon - n_e^2 Z_{\text{eff}})^2}{2\sigma_\epsilon^2}\right] \\ &\quad \times \frac{1}{\sqrt{2\pi}\sigma_\delta} \exp\left[-\frac{(\delta - n_e[Z_{\text{eff}} - 1])^2}{2\sigma_\delta^2}\right], \end{aligned} \quad (15)$$

where I represents any additional information we might possess, apart from the prior information on n_e and Z_{eff} . We will call the first likelihood (containing ϵ) the *continuum likelihood*, and the second likelihood (containing δ) the *CX likelihood*. Next, we need to decide on the prior information that we are willing to trust. Initially, we do not wish to assume a lot of prior information and we only impose some boundaries between which n_e and Z_{eff} are known to lie with ‘certainty’. On JET the following safe intervals can be proposed:

$$\begin{aligned} n_e &\in [0, 15] \times 10^{13} \text{cm}^{-3}, \\ Z_{\text{eff}} &\in [1, 5]. \end{aligned} \quad (16)$$

We model this prior information in two, relatively uninformative, uniform priors U_{n_e} and $U_{Z_{\text{eff}}}$, cut off at the respective boundaries. Following Bayes’ theorem, Equation (6), we can now write down a posterior distribution (for each time point), which matters only up to a proportionality factor:

$$\begin{aligned} p(n_e, Z_{\text{eff}} | \epsilon, \delta, I) &\sim p(\epsilon, \delta | n_e, Z_{\text{eff}}, I) p(n_e, Z_{\text{eff}}, I) \\ &\sim \exp\left[-\frac{(\epsilon - n_e^2 Z_{\text{eff}})^2}{2\sigma_\epsilon^2} - \frac{(\delta - n_e[Z_{\text{eff}} - 1])^2}{2\sigma_\delta^2}\right] U_{n_e} U_{Z_{\text{eff}}}. \end{aligned} \quad (17)$$

Furthermore, we can include relevant information on n_e through additional LIDAR measurements. We assume a statistical uncertainty on the LIDAR density measurements and we can encode this information in an extra independent Gaussian likelihood factor with a standard deviation σ_{n_e} given by the relative error of 5% on the n_e measurement. This likelihood is to be multiplied with (17), so the posterior becomes:

$$\begin{aligned} p(n_e, Z_{\text{eff}} | \epsilon, \delta, n_{e,L}, I) &\sim p(\epsilon, \delta, n_{e,L} | n_e, Z_{\text{eff}}, I) p(n_e, Z_{\text{eff}} | I) \\ &\sim \exp\left[-\frac{(\epsilon - n_e^2 Z_{\text{eff}})^2}{2\sigma_\epsilon^2} - \frac{(\delta - n_e[Z_{\text{eff}} - 1])^2}{2\sigma_\delta^2}\right] \\ &\quad \times \exp\left[-\frac{(n_{e,L} - n_e)^2}{2\sigma_{n_e}^2}\right] U_{n_e} U_{Z_{\text{eff}}}, \end{aligned} \quad (18)$$

where $n_{e,L}$ denotes the available LIDAR n_e measurement.

2) Posterior calculations:

a) *Measurements of ϵ and δ* : Suppose we have made a (series of) measurement(s) of ϵ and δ , excluding LIDAR measurements. Since we essentially search for the two unknown parameters n_e and Z_{eff} , we work in a two-dimensional parameter space. For problems of such low dimensionality, it is still computationally feasible to calculate the posterior density (17) on a (sufficiently large) grid of values for n_e and Z_{eff} . Fig. 7 gives a contour plot of the (normalized) posterior density resulting from such a procedure. In order to make inferences about n_e or Z_{eff} individually, we have to marginalize the joint posterior. No closed-form expression could be found for the integral over n_e , so in order to calculate the marginal posterior for Z_{eff} , numerical integration is required. In the

present simple case, we did not need MCMC to accomplish this as we could just sum over the parameter values on the (sufficiently fine) grid. We took an ϵ and δ value from pulse #61346 at 5.7 s and we calculated the marginal posterior distributions; the results are displayed in Fig. 8. We have chosen uniform priors, therefore it is correct to say that the modes of the marginal posteriors (i.e. the MAP estimates) equal the respective solutions of the deterministic model (12). A measure for the uncertainty at the maximum of the joint posterior is given by the covariance matrix of the distribution. The covariance can be approximated by the inverse Hessian at the maximum. In our case, if p denotes the posterior, the Hessian is given by

$$\begin{bmatrix} \frac{\partial^2 \log(p)}{\partial n_e^2} & \frac{\partial^2 \log(p)}{\partial n_e \partial Z_{\text{eff}}} \\ \frac{\partial^2 \log(p)}{\partial Z_{\text{eff}} \partial n_e} & \frac{\partial^2 \log(p)}{\partial Z_{\text{eff}}^2} \end{bmatrix}.$$

In effect, this comes down to approximating the joint posterior at the maximum by a Gaussian. The diagonal elements of the covariance matrix give the posterior variances on n_e , respectively Z_{eff} . In addition, the form of the joint posterior, Fig. 7, reveals the dependence between Z_{eff} and n_e ³. The off-diagonal elements of the covariance matrix give the cross-correlation coefficients between Z_{eff} and n_e . However, as mentioned before, the joint and marginal distributions contain the most information on the parameters n_e and Z_{eff} , given the model, the data and the error bars on the data. Another measure of uncertainty is the standard deviation of the marginals, particularly suited when the skewness of the marginals is limited. We calculated the marginal posterior mean and standard deviation for n_e and Z_{eff} at each time point. The results, shown in Figs. 9a and 9b, are not much different from the signals obtained via deterministic inversion of the model (i.e. the MAP estimates). The reason is the almost zero skewness of the marginal distributions (visible in Fig. 8), such that the marginal posterior mean is about the same as the mode. Therefore, again the LIDAR measurement is not well approximated. Incidentally, it is needless to say that the marginal posterior distributions are not necessarily as symmetric around their respective modes as in the present example, and in many cases they can be considerably skewed.

The error bars on the estimated density and Z_{eff} are much smaller than in the case of deterministic inversion (about 10% for n_e , 5% for Z_{eff}). The main reason is, of course, the fact that in the present case for computing the relative error on ϵ and δ , we have taken into account only the purely statistical uncertainty on the data. Nevertheless, for completeness we repeated the analysis using the same error bars on ϵ and δ as assumed in Section IV-A (20% on ϵ , 35% on δ). This resulted in an average 15% relative error on the estimated n_e and 16% on Z_{eff} . Both are still smaller than the errors found using the Gaussian error propagation in deterministic model inversion. We can conclude that in the present example the Bayesian analysis results in smaller error bars than the Gaussian error propagation. This is not always the case and there are situations where the advantage of a Bayesian analysis is marginal and where a Gaussian procedure suffices.

b) Measurements of ϵ , δ and $n_{e,L}$: We now turn to the case where not only ϵ and δ , but also a LIDAR electron density measurement $n_{e,L}$ is given, each with its respective statistical error. We perform the same analysis as in the previous subsection, again starting from the measurements of JET pulse #61346 at 5.7 s. However, we now work with the posterior given in (18). The marginal posteriors for n_e and Z_{eff} are shown in Fig. 10. According to our expectations, and comparing to Fig. 8, the characteristic width of the marginal posterior for n_e has shrunk considerably with the inclusion of the extra density measurement. Again, we calculated the marginal posterior mean and standard deviation for n_e and Z_{eff} at each time point. The results are displayed in Figs. 11a and 11b. The estimate for n_e now matches the LIDAR density much better than in the case where only the measurements ϵ and δ were used, with still slightly smaller error bars (about 5%). Of course, this is largely due to the high trust we have chosen to put in the LIDAR (5% statistical error as well). The estimate for Z_{eff} is now about the same as the CX Z_{eff} ; hence it is closer to the CX Z_{eff} compared to the case without extra LIDAR measurements. The relative error on the Z_{eff} estimates is about 4%. Similar results were observed for other discharges.

3) Consistency analysis: We would now like to visualize the degree of consistency of the various measured data sets. To this end, we made an overview of the marginal posterior distributions for n_e and Z_{eff} , taking into account various (combinations of) data sets. As an example, we took again JET pulse #61346 at 5.7 s. Fig. 12a presents the normalized marginal distributions for n_e for a posterior containing, respectively, a combination of continuum and CX likelihoods, continuum and LIDAR, CX and LIDAR and finally a combination of all three likelihood factors. The posterior densities for Z_{eff} are shown in Fig. 12b. The marginals for n_e and Z_{eff} for the combination of all

³A multimachine database yielded the following scaling relation between Z_{eff} and the line-averaged density \bar{n}_e : $Z_{\text{eff}} - 1 \sim \bar{n}_e^2$ [51], [52].

three likelihoods are symmetric, bell-shaped curves, so their maximum corresponds to the value of the estimated n_e , respectively Z_{eff} at 5.7 s, displayed in Figs. 11a and 11b.

The most striking features of this overview of marginal posteriors are the following. The density estimated from the combination of continuum and CX likelihoods (i.e. approximately the maximum of the solid curve), is considerably higher than the LIDAR density. This is a result already obtained in Fig. 9a. In addition, judging from the Z_{eff} marginals in Fig. 12b, the consistency between the continuum and CX data, each in combination with LIDAR measurements, is found to be negligible. This can be concluded from the small degree of overlap of the respective probability distributions. An integrated analysis of these data sets gives rise to the posterior distribution represented by the dotted curve. This marginal mostly overlaps with the CX-LIDAR density, indicating almost perfect consistency. However, the Z_{eff} marginal obtained from all three data sets is not at all consistent with the continuum-LIDAR marginal. In addition, note that the continuum-LIDAR marginal is substantially broader than the CX-LIDAR marginal. Thus, the Z_{eff} estimated from the continuum data has a considerably larger statistical error bar than the Z_{eff} derived from CXS.

We anticipate that a more sophisticated IDA scheme, where more uncertainties are modeled, or where the uncertainties are modeled in more detail, will reveal the sources of uncertainty that are responsible for the here observed data inconsistency. In the next subsection, we will see that this is actually the case. Indeed, it will become clear that in general one should not neglect the effect of *systematic* uncertainties on the input data. If we do not model all sources of uncertainty to a sufficient extent, in order to obtain intrinsic data consistency, we should not expect that our estimated Z_{eff} brings us much closer to the real, physical Z_{eff} value. Thus, the degree of consistency of the total data set in an IDA model can be seen as a measure for the quality of the corresponding parameter estimates obtained via the IDA analysis.

C. Statistical and systematic uncertainties

1) *Model definition:* We will now include some basic systematic uncertainties on the ϵ and δ measurements. One of the critical questions is whether this will increase the overall data consistency, compared to the case where only statistical uncertainties are modeled. We will also explicitly incorporate LIDAR density measurements, but we will not consider any systematic uncertainty on the LIDAR density. The statistical uncertainties on the input data ϵ and δ were taken of the same magnitude as in the previous subsection. In any case, the results concerning best estimates are only weakly dependent on the exact value of the statistical uncertainties [7].

As a first approximation we will assume that the measured data ϵ and δ may be off the real physical value by a scale factor. This scale factor may or may not be of interest, thus determining its status as a parameter of interest, or rather as a nuisance parameter, respectively. Let s_ϵ and s_δ stand for the scale factors belonging to the ϵ and δ measurements, respectively. Then the forward model is given by

$$\begin{cases} \epsilon = s_\epsilon n_e^2 Z_{\text{eff}} \\ \delta = s_\delta n_e (Z_{\text{eff}} - 1) . \\ n_{e,L} = n_e \end{cases} \quad (19)$$

For example, the factor s_δ may, amongst others, roughly account for impurity species other than fully stripped carbon, which could influence the value of Z_{eff} . If we suppose no measurement error, the density is assumed to be given exactly by $n_{e,L}$ and the forward model essentially reduces to a system of two equations (measurements ϵ and δ) with three unknowns: Z_{eff} , s_ϵ and s_δ . Clearly, without sufficient additional information on the three unknowns, it does not make much sense to try to derive a most probable or mean value for Z_{eff} , see the discussion in Section III-C2.

There are two possibilities now. Either we seek to deduce some useful prior information on Z_{eff} , s_ϵ and s_δ , or we need to gather more data. Instructive prior information could be the result of external measurements using other diagnostics. In a more detailed model, a systematic discrepancy on any quantity that enters the forward model could be quantified by additional laboratory measurements. Since such measurements are themselves affected by statistical uncertainty, this leads to a prior distribution on the corresponding parameter that describes the particular systematic uncertainty under study. A similar procedure was carried out in e.g. [4]. However, at this point we do not possess any objective informative prior information on the real Z_{eff} or the scale factors. In fact we are in a situation that arises rather frequently in the analysis of diagnostic data. Indeed, in many cases the diagnostician is

conscious of systematic uncertainty in the data, but it might be difficult to pinpoint where and how exactly this uncertainty enters the forward model. It may also be, as mentioned in the introduction, that trying to locate and quantify systematic uncertainties would require a disproportionate amount of effort and time. In such cases the procedure followed in the present work can be of assistance. The idea is, as outlined above, to include at different places in the forward model some very general parameters characterizing the systematic uncertainty, represented here by the scale factors s_ϵ and s_δ . Next, we are going to demand consistency between measurements performed in different time slices within a single discharge or within several discharges. Meanwhile, we assume that the parameters describing the systematic uncertainty remain constant over the considered time slices. This condition is reasonable in many situations and at least often can be presumed to be met to a sufficient extent. In the present case, we are able to acquire entire time traces for our measurements and this provides us with a lot more information. Hence, suppose we collect measurements at two different time instants t_1 and t_2 during a discharge, or in different discharges. We demand the plasma conditions at the two time instants to be sufficiently *different*, such that the variability of the data between t_1 and t_2 is not only caused by statistical fluctuations. Let us, as just discussed, also assume that the scale factor s_ϵ is the same at both t_1 and t_2 and similar for s_δ . At the time instants t_1 and t_2 we then have

$$\begin{cases} \epsilon_1 = s_\epsilon n_{e,1}^2 Z_{\text{eff},1} \\ \delta_1 = s_\delta n_{e,1} (Z_{\text{eff},1} - 1) \\ n_{e,L,1} = n_{e,1} \\ \epsilon_2 = s_\epsilon n_{e,2}^2 Z_{\text{eff},2} \\ \delta_2 = s_\delta n_{e,2} (Z_{\text{eff},2} - 1) \\ n_{e,L,2} = n_{e,2} \end{cases}, \quad (20)$$

with the subscripts 1 and 2 referring to t_1 and t_2 , respectively. For zero measurement error, this is equivalent to a system of four equations with four unknowns ($Z_{\text{eff},1}$, $Z_{\text{eff},2}$, s_ϵ and s_δ). This system can be solved exactly, as can be (and should always be!) readily checked. Therefore we can conduct a well-behaved IDA analysis to infer estimates of $n_{e,1}$, $n_{e,2}$, $Z_{\text{eff},1}$, $Z_{\text{eff},2}$, s_ϵ and s_δ . For two time instants the joint posterior density is given by

$$\begin{aligned} & p(n_{e,1}, Z_{\text{eff},1}, n_{e,2}, Z_{\text{eff},2}, s_\epsilon, s_\delta | \epsilon_1, \delta_1, n_{e,L,1}, \epsilon_2, \delta_2, n_{e,L,2}, I) \\ & \sim \exp \left[-\frac{(\epsilon_1 - s_\epsilon n_{e,1}^2 Z_{\text{eff},1})^2}{2\sigma_{\epsilon,1}^2} - \frac{(\delta_1 - s_\delta n_{e,1} [Z_{\text{eff},1} - 1])^2}{2\sigma_{\delta,1}^2} - \frac{(n_{e,L,1} - n_{e,1})^2}{2\sigma_{n_{e,1}}^2} \right] \\ & \times \exp \left[-\frac{(\epsilon_2 - s_\epsilon n_{e,2}^2 Z_{\text{eff},2})^2}{2\sigma_{\epsilon,2}^2} - \frac{(\delta_2 - s_\delta n_{e,2} [Z_{\text{eff},2} - 1])^2}{2\sigma_{\delta,2}^2} - \frac{(n_{e,L,2} - n_{e,2})^2}{2\sigma_{n_{e,2}}^2} \right] \\ & \times U_{n_{e,1}} U_{n_{e,2}} U_{Z_{\text{eff},1}} U_{Z_{\text{eff},2}} U_{s_\epsilon} U_{s_\delta}. \end{aligned} \quad (21)$$

Now, although the forward model (20) is invertible, there is still a difficulty associated with the posterior density (21). The problem is that the joint density, hence also the marginal distributions, are found to be very broad, i.e. the density is smeared out over a very large volume in the parameter space. As a result, not only is it hard to estimate such a distribution, but it also does not make much sense to summarize a similar distribution by point estimates for the quantities of interest, which would anyhow have very large statistical error bars attached to them. The reason is mainly the fact that we have introduced, through the uniform priors, an enormous uncertainty in both the values of Z_{eff} and the systematic uncertainties s_ϵ and s_δ . We chose to remain extremely ignorant about Z_{eff} and the scale factors, and the broad posterior is the price that is paid. In the present situation, we solved this problem as follows. We improved the statistics of the data by selecting a phase in the discharge wherein the plasma conditions are assumed to remain *stationary*. Referring to the model (20), this should be a region corresponding to time point t_1 . The constancy of the relevant plasma conditions during this phase may be checked, for example, by monitoring mainly the electron density and perhaps also the local continuum and CX Z_{eff} time traces throughout the discharge. Next, we considered a second phase (corresponding to t_2), again characterized by stationary plasma conditions, which should be, however, systematically different from the conditions in the first region. This way we can include in the posterior density the statistics of all data points within these two regions. Hence, the joint

posterior density becomes

$$\begin{aligned}
& p(n_{e,1}, Z_{\text{eff},1}, n_{e,2}, Z_{\text{eff},2}, s_\epsilon, s_\delta | \epsilon_1, \delta_1, \mathbf{n}_{e,L,1}, \epsilon_2, \delta_2, \mathbf{n}_{e,L,2}, I) \\
& \sim \exp \left\{ - \sum_{i=1}^{n_1} \left[\frac{(\epsilon_{1,i} - s_\epsilon n_{e,1} Z_{\text{eff},1})^2}{2\sigma_{\epsilon,1,i}^2} + \frac{(\delta_{1,i} - s_\delta n_{e,1} [Z_{\text{eff},1} - 1])^2}{2\sigma_{\delta,1,i}^2} + \frac{(n_{e,L,1,i} - n_{e,1})^2}{2\sigma_{n_{e,1,i}}^2} \right] \right\} \\
& \times \exp \left\{ - \sum_{j=1}^{n_2} \left[\frac{(\epsilon_{2,j} - s_\epsilon n_{e,2} Z_{\text{eff},2})^2}{2\sigma_{\epsilon,2,j}^2} + \frac{(\delta_{2,j} - s_\delta n_{e,2} [Z_{\text{eff},2} - 1])^2}{2\sigma_{\delta,2,j}^2} + \frac{(n_{e,L,2,j} - n_{e,2})^2}{2\sigma_{n_{e,2,j}}^2} \right] \right\} \\
& \times U_{n_{e,1}} U_{n_{e,2}} U_{Z_{\text{eff},1}} U_{Z_{\text{eff},2}} U_{s_\epsilon} U_{s_\delta}.
\end{aligned} \tag{22}$$

Here, the subscript 1 refers to the conditions in phase 1; likewise for the subscript 2. We now observe vectors of ϵ , δ and $n_{e,L}$ measurements, but note that we estimate only a single $n_{e,1}$, $n_{e,2}$, $Z_{\text{eff},1}$ and $Z_{\text{eff},2}$, characterizing the conditions in the respective plateau phases. n_1 and n_2 are the number of data points considered in phase 1, resp. phase 2. Finally, we note that phase 2 may—and often will—be part of a different discharge. The only requirement is that the sources of systematic uncertainty should be not too different in the two phases, such that s_ϵ and s_δ can be deemed constant between the phases. Ideally, discharges should be selected that are part of the same physical program.

The prior distributions for $n_{e,1}$, $n_{e,2}$, $Z_{\text{eff},1}$ and $Z_{\text{eff},2}$ are chosen similarly as in the case of only statistical uncertainty. The prior distributions for the scale factors should be sufficiently uninformative. Several (relatively) uninformative distributions for the scale parameters were tested, but eventually simulations on an artificial data set pointed out that a simple uniform distribution performs best. The uniform prior distributions U_{s_ϵ} and U_{s_δ} for s_ϵ , respectively s_δ just express the required positivity of both scale factors.

2) Posterior calculations:

a) *Artificial data set:* To sample from and marginalize the joint posterior distribution, we used the Metropolis-Hastings algorithm. This MCMC algorithm is a convenient tool in the present situation because we are dealing with six parameters to be estimated ($n_{e,1}$, $n_{e,2}$, $Z_{\text{eff},1}$, $Z_{\text{eff},2}$, s_ϵ and s_δ) and since the density (22) does not correspond to any known density. We first checked the estimation of the posterior density in the case of an artificial data set. Known values for n_e and Z_{eff} were assumed, at two different phases in an artificial discharge, with 15 data points per phase. From this, values for ϵ and δ were calculated, according to (20), using artificial values for s_ϵ ($= 0.5$) and s_δ ($= 1.8$). Next, Gaussian noise was added so as to obtain the relative statistical errors cited above (15% on ϵ , 10% on δ and 5% on $n_{e,L}$). A random walk Markov chain was run, accepting or rejecting samples according to the M-H acceptance probability and using the Cauchy distribution as a proposal distribution (in physics also known as the Lorentz distribution). This distribution has the advantage of being rather peaked on the one hand, but still having broad tails, so that the probability of sampling far away from its mode (i.e. the current value in the Markov chain) is sufficiently high. This ensures that the support of the posterior is scanned effectively. The Cauchy distribution for θ with location parameter x_0 and scale parameter κ has the following probability density:

$$C(\theta|x_0, \kappa) = \frac{1}{\pi} \frac{\kappa}{(x - x_0)^2 + \kappa^2}, \quad \kappa > 0.$$

Since the Cauchy distribution is symmetric in its argument versus its location parameter, it drops out of the acceptance probability, and we effectively use the Metropolis algorithm. The Cauchy inverse cumulative distribution function for a probability p is given by

$$F^{-1}(p|x_0, \kappa) = x_0 + \kappa \tan[\pi(p - 1/2)],$$

which can be used to sample a proposal value for the parameters of interest, based on the current value x_0 and given a probability p that is sampled from a uniform distribution. This way, the Markov chain becomes a random walk chain, by the definition given in Section III-B3. The scale parameter κ has to be chosen so as to maintain a post-convergence acceptance probability in the range 40% – 60%. This is a level recommended by Roberts, Gelman and Gilks [53] based on empirical studies, providing good convergence properties of the chain. The burn-in period was chosen to be 3000 samples, while the Monte Carlo period was 50,000 samples long.

The *a priori* chosen values for the parameters of interest are given in Table II. Estimates for the parameters are also mentioned, together with the respective statistical error bars. Since several of the marginal distributions

are often relatively skewed (see below), we chose the marginal posterior means as estimates for the parameters. Again because of the skewness of the marginal distributions, the error bars were obtained from the distributions by constructing a $100(1 - \alpha)\%$ Bayesian credible interval around the respective means of the marginal posteriors. The interval was created such that $100(1 - \alpha)/2\%$ of probability was situated toward the left of the mean and $100(1 - \alpha)/2\%$ toward the right. By choosing $\alpha = 0.317$ we get a 68.3% interval, which is the same probability contained within the interval $[\mu - \sigma, \mu + \sigma]$ for a normal distribution with mean μ and standard deviation σ . Thus, we obtain a similarity between the credible interval and the usual 1σ intervals quoted in Gaussian error analysis as error bars on a measurement. Moreover, if one of the marginal posterior distributions is skewed e.g. towards the right, the corresponding credible interval will be asymmetric around the posterior mean and so will the associated error bar. This indicates a larger uncertainty on the measurement toward the right of the mean than toward the left, corresponding to our intuition. The true parameter of interest will lie with a probability of $(1 - \alpha)$ within the specified credible interval. Note that in a Bayesian setting, the latter is a valid statement, in contrast to the case of a confidence interval in frequentist statistics. Indeed, in frequentist statistics the interval itself is random and one can only speak of the relative frequency of the true parameter lying within a long series of realizations of the interval, obtained from a series of data samples. We will denote the difference between the marginal posterior means and the upper or lower boundary of the corresponding credible interval by the common term ‘absolute error’, from which a relative error can be deduced as well. If a (post-convergence) MCMC sample $\{\theta_i\}$ of the marginal posteriors is available, as in the present case, a credible interval as described above is constructed easily by calculating sample quantiles as follows. Sort the samples in a rising order, denoted by $\{\theta_j\}$, $j = 1, \dots, n$, and locate the sample, say with index j_m , that is closest to the mean of the distribution. Then, in the sorted MCMC sample find the samples with index $j_l = j_m - \lfloor \frac{1-\alpha}{2}n \rfloor$ and $j_r = j_m + \lfloor \frac{1-\alpha}{2}n \rfloor$, where $\lfloor \cdot \rfloor$ denotes the integer part of its argument. The interval $[\theta_{j_l}, \theta_{j_r}]$ is the sought-after credible interval (see [54] for an explanation and [55] for a derivation of properties such as consistency and asymptotic normality of the sample quantiles). On the basis of the obtained post-convergence MCMC samples, we constructed normalized histograms for the estimated marginal posterior densities for the parameters of interest. These are shown in Fig. 13, with the mean denoted by a vertical line and the (approximate) 68.3% credible interval indicated by a darker shade in the histogram. The histograms for $n_{e,1}$ and $n_{e,2}$ are well approximated by a Gaussian, while the histograms for $Z_{\text{eff},1}$, $Z_{\text{eff},2}$, s_ϵ and s_δ have substantial non-zero skewness. This is reflected by the symmetric or asymmetric error bars in Table II. The traces⁴ (burn-in plus Monte-Carlo phase) of the quantities of interest during the MCMC sampling indicate good convergence of the chain; they are given in Fig. 14. A zoomed display of the trace for $n_{e,1}$ is given in Fig. 15, showing a clear burn-in period of some 100 samples, starting from an arbitrary initial value, after which the Markov chain converges to a stable distribution. Note also how, according to the M-H scheme, not every move of the chain is accepted, resulting in an acceptance probability of about 50%.

We can conclude from this analysis using an artificial data set that the quantities of interest are well estimated. Fig. 13 clearly shows that the marginal posterior mean provides a better point estimate for the quantities of interest compared to the marginal posterior mode in case the marginal posterior is skewed. This can only be obtained through the estimation of the full posterior distribution—a clear advantage over the estimation of the modes using Maximum a Posteriori. Our analysis also allows us to construct the credible intervals and we should note that the error bars on the values of Z_{eff} and the scale factors may become relatively large. Finally, we stress that deterministic inversion of the forward model would not be helpful here, since it does not permit us to make use of the improved statistics provided by including multiple data points per phase.

b) Real data set: We now turn to a real data set obtained from a set of discharges at JET, mentioned in Section II-C. We will select data points, on the magnetic axis, for phase 1 from a period in a first discharge where the LIDAR density assumes a plateau phase, in the sense that the remaining variability in the measurements is expected to stem from statistical fluctuation only. This means that ϵ and δ also remain in an approximate plateau phase during this period and this is a first assumption. Next, in a similar fashion we obtain data for phase 2 from a second discharge. However, the mean value of at least n_e in the second discharge should be systematically different from the values in the first discharge, i.e. the difference should *not* be due to statistical fluctuation only. Otherwise, the forward model theoretically becomes not invertible, while in practice we will obtain impractically

⁴An MCMC trace for a certain parameter of interest is the sequence of values adopted by the Markov chain in sampling from the associated marginal distribution. It should not be confused with a time trace of a plasma quantity during a discharge.

broad posteriors. The data for phase 2 may also originate from the same discharge, again provided the above conditions are fulfilled, but discharges with two plateau phases at different levels of the plasma quantities are less common. In addition, the data should be modeled reasonably well by the forward model, meaning in particular that the scale factors s_ϵ and s_δ should be more or less the same in the two phases. We expect this to be true at least for discharges belonging to the same physical program and this is a second assumption.

The LIDAR density for a first couple of JET discharges (#61346 and #61348, part of the physical program mentioned in Section II-C) from which data were used, is shown in Fig. 16. From each discharge, 15 samples were extracted for constructing the plateau phases, indicated in the figure as well. Next, the posterior distributions were estimated using MCMC, similar to the analysis using artificial data. Again, ϵ , δ and $n_{e,L}$ were taken to be normally distributed with relative errors quoted above; this is a third assumption. We found the estimates of the marginal posterior means to be rather independent of these error levels, although the credible intervals are, of course, more strongly influenced by these error bars. The normalized histograms obtained from the MCMC samples are given in Fig. 17, while the best estimates of the parameters of interest in terms of the marginal posterior means, together with the statistical error bars based on the credible intervals, are summarized in Table III. The marginals for Z_{eff} , s_ϵ and s_δ are relatively skewed (non-Gaussian) distributions. Therefore, estimation of the parameters of interest through the posterior mean yields rather different results than via MAP.

Let us compare for discharge #61346 the estimated value of Z_{eff} to the value, averaged over all samples in plateau 1, found in the case where only statistical uncertainties were modeled. Plateau 1 is defined by the region depicted in Fig. 16 (#61346) and the mean value of the Z_{eff} signal obtained in Section IV-B (including LIDAR measurements) within this region is about 1.6. This is significantly lower than the value of about 2.2, estimated for the case where the systematic uncertainties are taken into account. Analysis of other discharges points out similar discrepancies. Thus, systematic uncertainties definitely play a role in the present situation and they should be taken into the analysis. This will be confirmed by the consistency analysis below.

A similar procedure was carried out on data from JET discharges #61347–#61348 and the results are given in Table IV. Comparing with Table III, first we note that the estimated Z_{eff} values are, well within the 68.3% error bars, consistent between the two analyses, i.e. about the same Z_{eff} value is found in Tables III and IV for pulse #61348, as it should be. Second, we see that the estimated density corresponds very well to the LIDAR density, which is a consequence of the low statistical uncertainty that we assumed on the LIDAR measurement. Third, the estimated values for s_ϵ correspond very well between both cases. Similarly, the values for s_δ are comparable and at least can be deemed mutually compatible in view of the error bounds on the estimates (which may become rather large). This means that, as anticipated, the systematic uncertainties in the discharges we analyzed appear to be characterized by roughly the same values of s_ϵ and s_δ . Fourth, comparing the relative errors on the estimated values for Z_{eff} with the ones given in Table I, we see that in several cases we have succeeded in reducing the Z_{eff} error bars considerably, which are now, by the forward model, assumed to be of a pure statistical nature. The width of the error bars is, however, strongly dependent on the statistics of the data, i.e. the number of samples included in each of the plateau regions. Therefore, further reduction of the residual statistical error is possible by selecting broader plateau regions. Finally, the estimated factor s_ϵ is in both cases well above one, while s_δ is well below one. Therefore, in the discharges analysed, the ϵ measurements, i.e. the continuum data, appear to be overestimated. By the forward model (20) and since the on-axis n_e is assumed to be measured (through LIDAR Thomson scattering and interferometry) with sufficient precision, this would mean that the continuum Z_{eff} is an overestimate of the real physical Z_{eff} . Similarly, the real Z_{eff} appears to have a tendency to be higher than the CX Z_{eff} . One possible reason is the fact that in the derivation of the CX Z_{eff} only C^{6+} has been considered as an impurity, whereas the influence of other impurities may not be negligible in the present case.

Next, we applied the same scheme to data from the discharges #61346–#61348, but now at a normalized minor radius ρ of about 0.55. Similar statistical errors were found on the data as those obtained on the magnetic axis. As noted before from Fig. 3d, at $\rho = 0.55$ the continuum and CX Z_{eff} are systematically inconsistent. However, since the average ratio of the continuum to the CX Z_{eff} is different from the typical ratio on the magnetic axis, the scale factors s_ϵ and s_δ modeling the systematic errors at $\rho = 0.55$ are expected to be different from those on the magnetic axis. This is confirmed by the results of the analysis given in Table V. The estimated Z_{eff} now corresponds in absolute value very well to the CX Z_{eff} . The nonphysical values (below 1) of the continuum Z_{eff} are clearly rejected. The estimated Z_{eff} is for both discharges lower than the corresponding on-axis estimates. This suggests a peaked Z_{eff} profile due to impurity accumulation on the magnetic axis. It confirms a similar tendency

reflected in the continuum Z_{eff} profile in Fig. 3d, although not seen in the CX Z_{eff} profile.

3) *Consistency analysis:* Finally, let us again perform a consistency analysis on the same data as used in Section IV-B, namely pulse #61346 at 5.7 s (magnetic axis). This time point falls inside the interval given in Fig. 16. However, this time we include in the posterior the values of the scale factors $s_\epsilon = 1.35$ and $s_\delta = 0.63$ obtained in the above combined analysis of discharges #61346 and #61348, as stated in Table III (i.e. ϵ and δ are scale-corrected). Figs. 18a and 18b depict the marginal posterior distributions for n_e and Z_{eff} , again taking into account a combination of continuum and CX likelihoods, continuum and LIDAR, CX and LIDAR and finally a combination of all three likelihood factors. It is apparent that the respective marginals, both for n_e and Z_{eff} , overlap for a large part. This indicates a considerable increase of consistency between the scale-corrected data ϵ and δ , compared to the case in Section IV-B, where only statistical uncertainties were taken into account (this could be quantified e.g. by calculating Kullback-Leibler divergences). In particular, in contrast to the analysis considering only statistical uncertainties, the Z_{eff} estimate is consistent with both the continuum (+ LIDAR) and CX (+ LIDAR) data. Thus, by including the systematic uncertainties into the analysis, even through the simple model we proposed, we have succeeded in rendering the data set intrinsically consistent. Therefore, since we took into our analysis both the continuum and CX data, modeling the associated systematic uncertainties, resulting in intrinsic data consistency, we put more trust in the presently estimated Z_{eff} compared to either the continuum or the CX Z_{eff} . Given the current forward model and the data, it is clearly the best estimate of Z_{eff} we are able to give. Should more—and possibly contradictory—information become available, e.g. expert spectroscopist knowledge, then ideally it should be included into the analysis: either in the forward model (likelihood) or as prior information.

We should note, however, that still a more detailed analysis of the uncertainties involved, carefully following the chain of calculations from the raw measurements to the final quantities of interest, is likely to still increase the accuracy of the estimates as well as the intrinsic data consistency. This can be accomplished by listing the various sources of statistical and systematic uncertainties and by inserting some useful prior information for these uncertainties into the IDA analysis. Moreover, such a procedure could result in a sufficiently narrow posterior, without the need for including the statistics of multiple samples. This would allow to process the data and derive estimates on a point-by-point basis. Nevertheless, we would like to stress again the value of the present mode of operation, demanding consistency between measurements in multiple time slices, in situations where very little is known about the nature of the systematic uncertainties.

V. CONSIDERATIONS FOR ITER AND DEMO

The framework of Integrated Data Analysis will have an even more important role to play in the analysis of diagnostic data for ITER and other next-step devices such as DEMO. To begin with, as already mentioned in the introduction, there will not be enough space available (e.g. through access ports)—even more so than at present-generation devices—to accommodate every diagnostic useful for studying the physical properties of the plasma. Some economy will have to be exercised in terms of number of diagnostics. Therefore, diagnostic data analysis will benefit from an integrated approach where interdependencies between redundant data sets enable a reliable assessment of plasma quantities, in case each of the individual data sets provide only limited information. In addition, access to diagnostic hardware will be reduced. This means that it will be difficult or impossible to correct any hardware failures, possibly introducing important systematic effects in the data. As demonstrated in the present work, IDA permits to efficiently deal with systematic uncertainties, either through the inclusion of relevant prior information, or by demanding consistency with other (laboratory) measurements or time frames. Furthermore, many diagnostic instruments will be exposed to severe physical conditions and there will be a critical need for *in situ* calibration. To do this, again data from different diagnostic sources will have to be combined, which can be accomplished most successfully via IDA.

With regard to the determination of Z_{eff} at ITER, we should note that the accuracy of the measurement required to allow a safe operation of the machine is stated to be 20% for line-integrated values and 10% for profiles [56], [57]. Concerning the profiles, this is typically more stringent than what can be accomplished with the current diagnostic capabilities, either using continuum or CX measurements (see also Table I). However, note that when modeling both statistical and systematic uncertainties in Section IV-C, we have succeeded in meeting this requirement in several instances; refer to Tables III, IV and V. As mentioned above, an even higher accuracy could be achieved by considering more samples per plateau region. Alternatively, the systematic uncertainties could be modeled at a more

detailed level. In any case, the required accuracy for Z_{eff} profiles at ITER comes into reach using an integrated analysis on continuum and CX measurements.

VI. CONCLUSION

The purpose of this work was, first, to address the long-standing difficulty of the inconsistency between the derivation of the ion effective charge Z_{eff} from bremsstrahlung measurements on the one hand, and from the weighted summation of individual impurity densities obtained from CXS on the other hand. Second, in doing so, we discussed the value and many key principles of the framework of Integrated Data Analysis by means of Bayesian Probability Theory. Third, we demonstrated the practical benefits of a gradual approach, starting from an elementary forward model. This already allows informative conclusions to be drawn, eliminating the necessity for a fully detailed model right from the beginning. In turn, inferences from a less sophisticated model may further guide a more detailed analysis of the involved uncertainties.

The systematic discrepancy between the continuum and CX Z_{eff} was shown to be apparent in both line-integrated and local estimates in discharges at JET. We gave an overview of the main sources of uncertainty, responsible for this discrepancy, in the assessment of Z_{eff} from bremsstrahlung and CX measurements individually. In this work, the uncertainties were modeled via the IDA framework. The main advantages are, to begin with, the possibility of integration of multiple data sets of a heterogeneous nature in order to derive more trustworthy estimates for the quantities of interest, with smaller error bars than obtained via a traditional analysis using measurements from individual diagnostics. In addition, IDA enables, in conjunction with BPT, a straightforward modeling of all sources of uncertainty, tackling non-Gaussian error propagation as well as systematic errors without difficulty. We briefly outlined some basic concepts of BPT, including the core idea represented by Bayes' theorem, expressing the concept of learning from the prior information, through the acquisition of measurements, to a joint posterior distribution for the quantities of interest. The likelihood compares the measurements to the outcome of the physical (forward) model that depends on some underlying parameters, which are to be estimated. Many real-world problems in Bayesian analysis involve non-standard posteriors or a large amount of parameters. Therefore, in order to render practical the marginalization of the posterior and the derivation of best estimates as well as error bars for the quantities of interest, stochastic integration methods such as Markov Chain Monte Carlo are appropriate. This allows to sample from the (marginal) posterior distribution by setting up a Markov Chain in a multidimensional space. A move of the chain is determined by sampling from a proposal distribution, accepting or rather rejecting the proposal with a certain probability depending on both the posterior and the proposal distribution. Monte Carlo averages over the thus obtained samples permit the estimation of e.g. the marginal posterior means of the quantities of interest. Statistical error bars can be obtained by constructing Bayesian credible intervals that encompass a certain amount of posterior probability. Next, the IDA recipe was given, where statistical uncertainty is modeled through likelihood distributions, while systematic uncertainty is covered by introducing nuisance parameters.

This framework was applied for the integrated estimation of the local Z_{eff} and the electron density, since the two quantities are strongly correlated. The analysis was conducted at three stages of sophistication. First, no *a priori* measurement error was assumed, permitting the deterministic inversion of the forward model including both continuum and CX measurements. Carrying out this procedure on time traces of measurements, we found unrealistically low values for Z_{eff} as well as a serious systematic overestimation of the electron density, compared to the well trusted value obtained from LIDAR Thomson scattering and interferometry. Error bars, calculated via Gaussian error propagation, were fairly large. In a second step, statistical measurement error was allowed and a posterior distribution for Z_{eff} and n_e was obtained. Provided additional measurements of the electron density were included into the analysis, the estimated and measured n_e were in good agreement, while Z_{eff} was estimated to equal the CX Z_{eff} . The statistical error bars, derived from the marginal posterior standard deviation, were found to be substantially smaller than obtained via the Gaussian analysis, indicating the non-Gaussian character of the error propagation. However, a consistency analysis pointed out the lack of intrinsic data consistency and, as such, the need for the modeling of systematic uncertainty. This was remedied to a large extent by introducing two scale factors in the continuum and CX data sets, acting as nuisance parameters. However, very little prior information was assumed about Z_{eff} and the scale factors, and it was pointed out that in such a case care should be taken to ensure invertibility of the forward model. This difficulty was solved by the demand of consistency of the measurements between two time slices, in two different discharges, during data acquisition. It is a way of working that is most useful in the frequent situation where little is known about the systematic uncertainties in the data. In addition,

the statistics of the data were improved by considering multiple measurements in phases characterized by roughly stationary plasma conditions. For marginalization of the posterior and inference of marginal posterior means, the Metropolis (MCMC) algorithm was applied. This also enabled the construction of 68.3% credible intervals, used to formulate statistical error bars. The estimated electron density was more or less equal to the density obtained from Thomson scattering. The Z_{eff} estimates on the magnetic axis were found to lie in between the CX and continuum Z_{eff} , suggesting an overestimation of the continuum data and an underestimation of the CX measurements in the data set under study. Further towards the plasma periphery, the estimated Z_{eff} turned out to lie very close to the CX Z_{eff} , dismissing unrealistically low levels of the continuum Z_{eff} . A consistency analysis revealed a substantially improved intrinsic data consistency compared to the case where only statistical uncertainties were modeled. Since we used data from two diagnostics, modeling the systematic uncertainty, this implies that the obtained Z_{eff} values are more reliable than either the continuum or CX Z_{eff} . The remaining statistical error bars on Z_{eff} suggest an accuracy that approaches requirements for Z_{eff} profiles at ITER. Even higher accuracy can be obtained by considering more data points per plateau phase in the forward model. Nevertheless, it should be stressed that through our analysis we already succeeded in mitigating the important and *a priori* entirely unknown systematic uncertainty in the data.

A next step in the present work would involve listing in more detail the various sources of uncertainty in the chain of data acquisition and subsequent calculations for both the continuum and CX systems. Additional measurements would have to provide prior information for the various nuisance parameters describing systematic uncertainties, possibly aided by consistency requirements between data in different time slices. This would enable the formulation of a posterior distribution, from which inferences could be drawn, resulting in still more accurate and consistent Z_{eff} estimates. If carried out to a full extent for both diagnostics, this would represent a considerable amount of work. Therefore, the strategy should be adopted where the analysis is carried out at increasingly higher levels of model sophistication. We hope that the present paper has laid the foundations for a possible route towards a more reliable derivation of Z_{eff} that closes the long-standing gap between the traditional Z_{eff} estimates. At the same time, we have endeavored to present an accessible overview of the methods of Integrated Data Analysis by means of Bayesian Probability Theory, which we hope will stimulate other researchers in the field of plasma diagnosis to put to use these important techniques for the improved validation and analysis of their diagnostic data.

ACKNOWLEDGMENT

This work, supported by the European Communities under the contract of Association between EURATOM and the Belgian State, was carried out within the framework of the European Fusion Development Agreement. The views and opinions expressed herein do not necessarily reflect those of the European Commission. The authors would like to thank M. von Hellermann for providing a wealth of information on (JET) Charge Exchange Spectroscopy and its uncertainties, as well as for several stimulating and helpful discussions.

REFERENCES

- [1] I. Hutchinson, *Principles of Plasma Diagnostics*, 2nd ed. Cambridge: Cambridge University Press, 2002.
- [2] R. Fischer, A. Dinklage, and E. Pasch, "Bayesian modelling of fusion diagnostics," *Plasma Phys. Control. Fusion*, vol. 45, no. 7, pp. 1095–1111, 2003.
- [3] S. Arshad *et al.*, "Chapter 11: Data validation, analysis, and applications for fusion plasmas," *Fusion Sci. Technol.*, vol. 53, no. 2, pp. 667–698, 2008.
- [4] R. Fischer, C. Wendland, A. Dinklage, S. Gori, V. Dose, and the W7-AS team, "Thomson scattering analysis with the Bayesian probability theory," *Plasma Phys. Control. Fusion*, vol. 44, no. 8, pp. 1501–1519, 2002.
- [5] R. Fischer and A. Dinklage, "Integrated data analysis of fusion diagnostics by means of the Bayesian probability theory," *Rev. Sci. Instrum.*, vol. 75, no. 10, pp. 4237–4239, 2004.
- [6] A. Dinklage, R. Fischer, and J. Svensson, "Topics and methods for data validation by means of Bayesian probability theory," *Fusion Sci. Technol.*, vol. 46, no. 2, pp. 355–364, 2004.
- [7] G. Verdoolaege, "Bayesian integrated determination of the ion effective charge via bremsstrahlung and Charge Exchange Spectroscopy in tokamak plasmas," Ph.D. dissertation, Ghent University, 2006.
- [8] G. Verdoolaege, M. von Hellermann, R. Jaspers, M. Ichir, and G. Van Oost, "Integrated Bayesian estimation of Z_{eff} in the TEXTOR tokamak from bremsstrahlung and CX impurity density measurements," in *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, ser. AIP conference proceedings, vol. 872. Melville, NY: American Institute of Physics, 2006, pp. 541–548.
- [9] K. Kadota, M. Otsuka, and J. Fujita, "Space- and time-resolved study of impurities by visible spectroscopy in the high-density regime of JIPP T-II tokamak plasma," *Nucl. Fusion*, vol. 20, no. 2, pp. 209–212, 1980.
- [10] G. Verdoolaege, G. Telesca, E. Delabie, G. Van Oost, and the TEXTOR team, "Design and commissioning of a new diagnostic for routine determination of Z_{eff} from visible bremsstrahlung measurements on TEXTOR," *Rev. Sci. Instrum.*, vol. 77, no. 10, p. 10F310, 2006.

- [11] N. Schoon, J. Ongena, G. Van Oost, J. Schwarz, G. Telesca, and A. Cosler, "Measurement of the bremsstrahlung radiation in the visible and determination of Z_{eff} on TEXTOR," LPP – Royal Military Academy, Brussels, Belgium, Tech. Rep. 92, 1990.
- [12] M. Krychowiak, R. König, and T. Klinger, "Bayesian analysis of the effective charge from spectroscopic bremsstrahlung measurements in fusion plasmas," *J. Appl. Phys.*, vol. 96, no. 9, pp. 4784–4792, 2004.
- [13] H. Meister, R. Fischer, L. Horton, C. Maggi, D. Nishijima, C. Giroud, K.-D. Zastrow, B. Zaniol, the ASDEX Upgrade Team, and JET-EFDA Contributors, " Z_{eff} from spectroscopic bremsstrahlung measurements at ASDEX Upgrade and JET," *Rev. Sci. Instrum.*, vol. 75, no. 10, pp. 4097–4099, 2004.
- [14] I. Hutchinson, *Principles of Plasma Diagnostics*, 2nd ed. Cambridge: Cambridge University Press, 2002, ch. 5.
- [15] A. Boileau, M. von Hellermann, L. Horton, J. Spence, and H. Summers, "The deduction of low- Z ion temperature and densities in the JET tokamak using Charge Exchange Recombination Spectroscopy," *Plasma Phys. Control. Fusion*, vol. 31, no. 5, pp. 779–804, 1989.
- [16] M. von Hellermann and H. Summers, *Active Beam Spectroscopy at JET*. London: Elsevier Science Publishers, 1993, p. 135.
- [17] M. von Hellermann, W. Mandl, H. Summers, H. Weisen, A. Boileau, P. Morgan, H. Morsi, R. Koenig, M. Stamp, and R. Wolf, "Visible Charge Exchange Spectroscopy at JET," *Rev. Sci. Instrum.*, vol. 61, no. 11, pp. 3479–3486, 1990.
- [18] C. Negus, C. Giroud, A. Meigs, K.-D. Zastrow, and D. Hillis, "Enhanced core charge exchange recombination spectroscopy system on Joint European Torus," *Rev. Sci. Instrum.*, vol. 77, no. 10, p. 10F102, 2006.
- [19] A. Boileau, M. von Hellermann, L. Horton, H. Summers, and P. Morgan, "Analysis of low- Z impurity behavior in jet by charge-exchange spectroscopy measurements," *Nucl. Fusion*, vol. 29, no. 9, pp. 1449–1458, 1989.
- [20] H. Weisen, M. von Hellermann, A. Boileau, L. Horton, W. Mandl, and H. Summers, "Charge Exchange Spectroscopy measurements of ion temperature and toroidal rotation in JET," *Nucl. Fusion*, vol. 29, no. 12, pp. 2187–2197, 1989.
- [21] M. Tunklev, P. Breger, K. Günther, M. von Hellermann, R. König, M. O'Mullane, and K.-D. Zastrow, "Modelling of passive charge exchange emission and neutral background density deduction in JET," *Plasma Phys. Control. Fusion*, vol. 41, no. 8, pp. 985–1004, 1999.
- [22] M. von Hellermann, "A Charge Exchange analysis package for ITER," in *Proceedings of the 3rd Workshop on Fusion Data Processing, Validation and Analysis*, Cadarache, 2004.
- [23] —, 2006, Private communication.
- [24] G. Telesca, G. Van Oost, J. Ongena, C. Leys, and N. Schoon, "Investigation of the sensitivity of visible continuum $Z_{\text{eff}}(0)$ to edge radiation," in *Proceedings of the International Conference on Plasma Physics*, vol. II, Innsbruck, 1992, pp. 1147–1149.
- [25] B. Schunke, G. Huysmans, and P. Thomas, "Evidence of the influence of reflections on the Z_{eff} profile measurements in the Tore Supra tokamak and their mitigation," *Rev. Sci. Instrum.*, vol. 76, no. 7, p. 083501, 2005.
- [26] T. Bayes, "An essay towards solving a problem in the doctrine of chances," *Phil. Trans. R. Soc. London*, vol. 53, pp. 370–418, 1763.
- [27] P. Laplace, *Théorie Analytique des Probabilités*, 1st ed. Paris: Courcier, 1812.
- [28] R. Cox, "Probability, frequency and reasonable expectation," *Am. J. Phys.*, vol. 14, no. 1, pp. 1–13, 1946.
- [29] E. Jaynes, *Probability theory: the logic of science*. Cambridge: Cambridge University Press, 2003.
- [30] D. Sivia, *Data Analysis: a Bayesian tutorial*. New York: Oxford University Press, 1996.
- [31] T. Loredo, "From Laplace to supernova SN 1987A," in *Maximum Entropy and Bayesian Methods*. Dordrecht: Kluwer Academic, 1990, pp. 81–142.
- [32] A. Gelman, J. Carlin, H. Stern, and D. Rubin, *Bayesian data analysis*, 2nd ed. Boca Raton: Chapman & Hall/CRC, 2004.
- [33] D. MacKay, "Bayesian interpolation," *Neural Comput.*, vol. 4, no. 3, pp. 415–447, 1992.
- [34] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller, "Equations of state calculations by fast computing machines," *J. Chem. Phys.*, vol. 21, no. 6, pp. 1087–1092, 1953.
- [35] W. Hastings, "Monte Carlo sampling methods using Markov chains and their applications," *Biometrika*, vol. 57, no. 1, pp. 97–109, 1970.
- [36] A. Gelfand and A. Smith, "Sampling based approaches to calculating marginal densities," *J. Am. Stat. Assoc.*, vol. 85, no. 410, pp. 398–409, 1990.
- [37] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 6, no. 6, pp. 721–741, 1984.
- [38] A. Gelman, "Iterative and non-iterative simulation algorithms," *Computing science and statistics*, vol. 24, pp. 433–438, 1992.
- [39] L. Tierney, "Markov chains for exploring posterior distributions," *Ann. Stat.*, vol. 22, no. 4, pp. 1701–1728, 1994.
- [40] C. Robert and G. Casella, *Monte Carlo statistical methods*, 2nd ed. New York: Springer, 2004.
- [41] R. Neal, "Probabilistic inference using Markov Chain Monte Carlo methods," 1993, <http://omega.albany.edu:8008/neal.pdf>.
- [42] W. Gilks, S. Richardson, and D. Spiegelhalter, *Monte Carlo Statistical Methods*. London: Chapman & Hall, 1996.
- [43] D. Gamerman, *Markov Chain Monte Carlo: stochastic simulation for Bayesian inference*. London: Chapman & Hall, 1997.
- [44] P. Gregory, *Bayesian logical data analysis for the physical sciences*. Cambridge: Cambridge University Press, 2005.
- [45] J. Svensson, A. Dinklage, J. Geiger, A. Werner, and R. Fischer, "Integrating diagnostic data analysis for W7-AS using Bayesian graphical models," *Rev. Sci. Instrum.*, vol. 75, no. 10, pp. 4219–4221, 2004.
- [46] R. Fischer, E. Wolfrum, J. Schweinzer, and the ASDEX Upgrade Team, "Probabilistic lithium beam data analysis," *Plasma Phys. Control. Fusion*, vol. 50, no. 8, p. 085009, 2008.
- [47] —, "Probabilistic lithium beam data analysis," in *Proceedings of the 34th EPS Conference on Plasma Physics*, P-5.077, Warsaw, 2007.
- [48] R. Fischer, E. Wolfrum, W. Suttrop, and A. Kallenbach, "Integrated density profile analysis in ASDEX Upgrade H-modes," in *Proceedings of the 35th EPS Conference on Plasma Physics*, P-4.010, Hersonissos, 2008.
- [49] S. Kullback, *Information theory and statistics*. New York: Dover Publications, 1968.
- [50] H. Dreier, A. Dinklage, R. Fischer, M. Hirsch, P. Kornejew, and E. Pasch, "Design of diagnostics: Case studies for Wendelstein 7-X," *Fusion Sci. Technol.*, vol. 50, no. 2, pp. 262–267, 2006.
- [51] G. Matthews *et al.*, "Scaling radiative plasmas to ITER," *J. Nucl. Mater.*, vol. 241, pp. 450–455, 1997.

- [52] G. Telesca, U. Samm, B. Unterberg, R. Koch, A. Messiaen, and G. Van Oost, "Radiated power and ionic effective charge during neon injection experiments on TEXTOR," *Nucl. Fusion*, vol. 36, no. 3, pp. 347–358, 1996.
- [53] G. Roberts, A. Gelman, and W. Gilks, "Weak convergence and optimal scaling of random walk Metropolis algorithms," *Ann. Appl. Probab.*, vol. 7, no. 1, pp. 110–120, 1997.
- [54] M. Chen and Q. Shao, "Monte Carlo estimation of Bayesian credible and HPD intervals," *J. Comput. Graph. Stat.*, vol. 8, no. 1, pp. 69–92, 1999.
- [55] P. Sen and J. Singer, *Large sample methods in statistics*. New York: Chapman & Hall, 1994.
- [56] A. Donné *et al.*, "Chapter 7: Diagnostics," *Nucl. Fusion*, vol. 47, no. 6, pp. S337–S384, 2007.
- [57] K. Young, "Chapter 1: Plasma measurements: an overview of requirements and status," *Fusion Sci. Technol.*, vol. 53, no. 2, pp. 281–303, 2008.

TABLE I

ESTIMATED RELATIVE ERRORS ON SEVERAL PLASMA QUANTITIES IN JET PLASMAS ON THE MAGNETIC AXIS. THE ERRORS ARE GENERALLY BOTH OF A STATISTICAL AND SYSTEMATIC NATURE.

Quantity	Relative error
n_e	5%
T_e	10%
ϵ_{ff}	20%
n_C	35%
Continuum Z_{eff}	25%
CX Z_{eff}	20%

TABLE II

A PRIORI CHOSEN AND MCMC ESTIMATED VALUES FOR THE PARAMETERS OF INTEREST IN AN EXPERIMENT FOR THE ESTIMATION OF n_e AND Z_{eff} USING ARTIFICIAL DATA AND INCLUDING SYSTEMATIC UNCERTAINTIES. THE ERROR BARS (ABSOLUTE AND RELATIVE ERROR) REPRESENT THE LENGTH OF A 68.3% CREDIBLE INTERVAL, WITH THE PROBABILITY EQUALLY DIVIDED AROUND THE RESPECTIVE MARGINAL POSTERIOR MEANS.

Parameter	Original value	Estimated value	Absolute error	Relative error
$n_{e,1}$ ($\times 10^{13} \text{cm}^{-3}$)	3.0	3.00	± 0.04	$\pm 1\%$
$n_{e,2}$ ($\times 10^{13} \text{cm}^{-3}$)	4.0	4.03	± 0.05	$\pm 1\%$
$Z_{\text{eff},1}$	2.0	1.96	+0.22/ -0.30	+11%/ -15%
$Z_{\text{eff},2}$	1.7	1.64	+0.27/ -0.33	+17%/ -20%
s_ϵ	0.5	0.53	+0.13/ -0.07	+24%/ -13%
s_δ	1.8	1.80	+0.53/ -0.36	+30%/ -20%

TABLE III

MEASURED AND MCMC ESTIMATED VALUES FOR THE PARAMETERS OF INTEREST IN AN EXPERIMENT FOR THE ESTIMATION OF n_e AND Z_{eff} USING DATA FROM JET DISCHARGES #61346 AND #61348 ON THE MAGNETIC AXIS AND INCLUDING SYSTEMATIC UNCERTAINTIES. THE TRADITIONALLY MEASURED VALUES WERE AVERAGED OVER THE RESPECTIVE PLATEAU REGIONS.

Parameter	Traditional mean value	Estimated value	Absolute error	Relative error
$n_{e,1}$ ($\times 10^{13} \text{cm}^{-3}$)	3.26	3.24	± 0.04	$\pm 1\%$
$n_{e,2}$ ($\times 10^{13} \text{cm}^{-3}$)	4.51	4.50	± 0.05	$\pm 1\%$
$Z_{\text{eff},1}$	Cont. 2.96	2.19	+0.28/ -0.43	+13%/ -20%
	CX 1.63			
$Z_{\text{eff},2}$	Cont. 3.22	2.36	+0.22/ -0.37	+9%/ -16%
	CX 1.90			
s_ϵ	n/a	1.35	+0.33/ -0.14	+25%/ -11%
s_δ	n/a	0.63	+0.46/ -0.11	+73%/ -17%

TABLE IV
 MEASURED AND MCMC ESTIMATED VALUES FOR THE PARAMETERS OF INTEREST IN AN EXPERIMENT FOR THE ESTIMATION OF n_e
 AND Z_{eff} USING DATA FROM JET DISCHARGES #61347 AND #61348 ON THE MAGNETIC AXIS AND INCLUDING SYSTEMATIC
 UNCERTAINTIES.

Parameter	Traditional mean value	Estimated value	Absolute error	Relative error
$n_{e,1}$ ($\times 10^{13} \text{cm}^{-3}$)	3.41	3.39	± 0.04	$\pm 1\%$
$n_{e,2}$ ($\times 10^{13} \text{cm}^{-3}$)	4.51	4.51	± 0.06	$\pm 1\%$
$Z_{\text{eff},1}$	Cont. 2.82	2.20	+0.26/	+12%/
	CX 1.55		-0.55	-25%
$Z_{\text{eff},2}$	Cont. 3.22	2.41	+0.20/	+8%/
	CX 1.90		-0.45	-19%
s_ϵ	n/a	1.32	+0.64/ -0.13	+49%/ -10%
s_δ	n/a	0.55	+0.31/ -0.08	+56%/ -14%

TABLE V
 MEASURED AND MCMC ESTIMATED VALUES FOR THE PARAMETERS OF INTEREST IN AN EXPERIMENT FOR THE ESTIMATION OF n_e
 AND Z_{eff} USING DATA FROM JET DISCHARGES #61346 AND #61348 AT $\rho = 0.55$ AND INCLUDING SYSTEMATIC UNCERTAINTIES.

Parameter	Traditional mean value	Estimated value	Absolute error	Relative error
$n_{e,1}$ ($\times 10^{13} \text{cm}^{-3}$)	1.56	1.56	± 0.02	$\pm 1\%$
$n_{e,2}$ ($\times 10^{13} \text{cm}^{-3}$)	1.96	1.95	± 0.02	$\pm 1\%$
$Z_{\text{eff},1}$	Cont. 0.97	1.56	+0.26/	+17%/
	CX 1.55		-0.31	-20%
$Z_{\text{eff},2}$	Cont. 1.09	1.85	+0.21/	+11%/
	CX 1.84		-0.29	-16%
s_ϵ	n/a	0.57	+0.15/ -0.07	+25%/ -12%
s_δ	n/a	1.14	+0.71/ -0.24	+62%/ -21%




Fig. 1. Schematic of the JET CXS diagnostic. Both the vertical system (spectrometer 1 and 2) and horizontal system (spectrometers 3, 4 and 5) are shown [21].

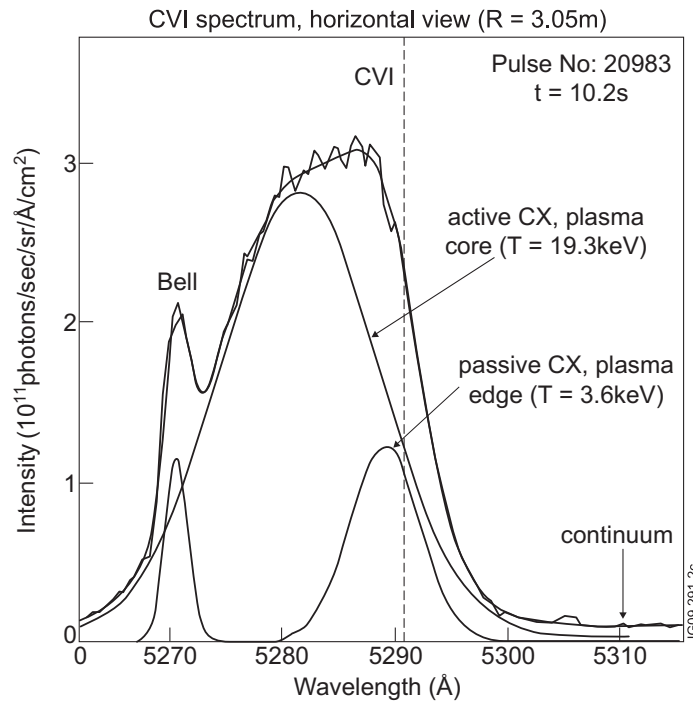
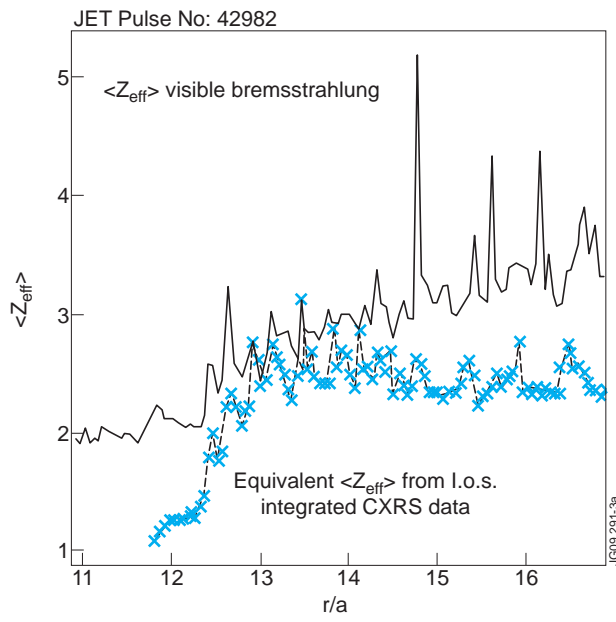
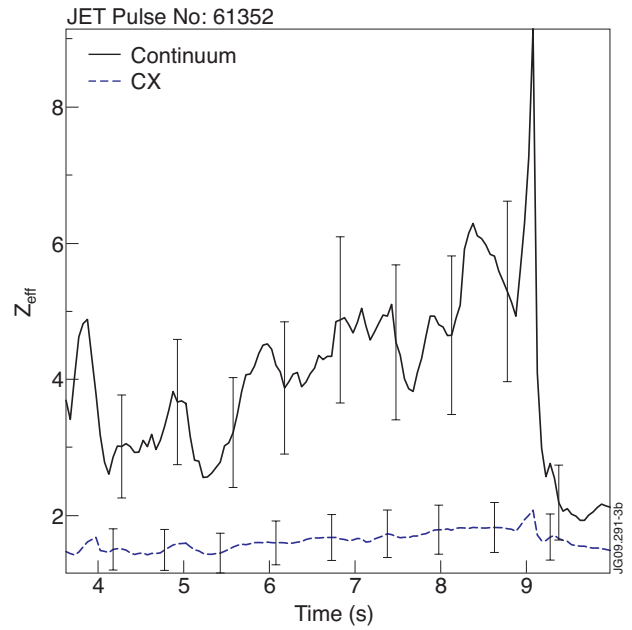


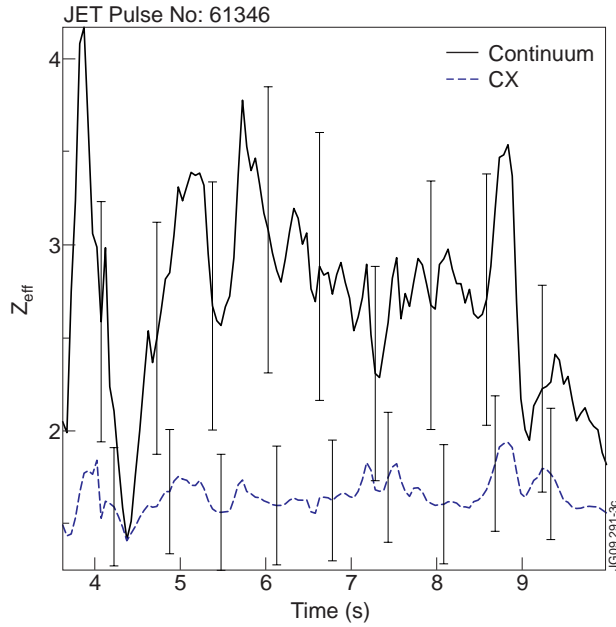
Fig. 2. Charge exchange spectrum of CVI at 5290.5 Å ($T_e = 19.3$ keV) from the horizontal JET CXS diagnostic. Note the active (beam-related) and passive contributions to the spectrum ($T_e = 3.6$ keV) [16].



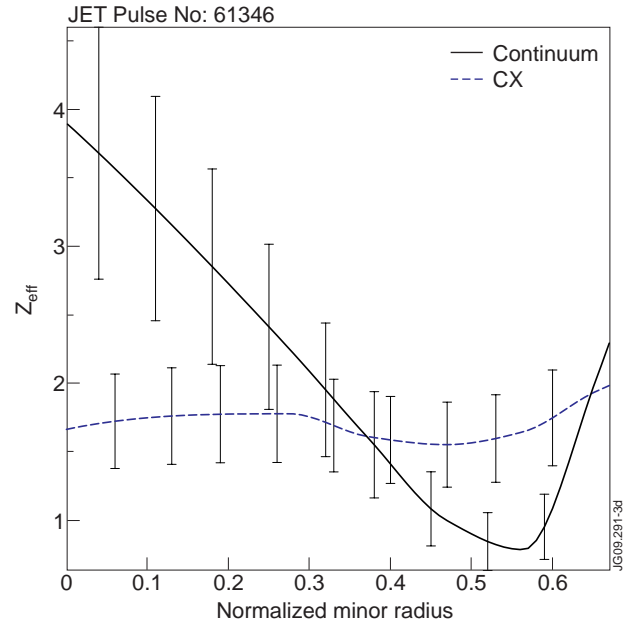
(a)



(b)



(c)



(d)

Fig. 3. A general inconsistency of the continuum and CX Z_{eff} , illustrated here by some JET pulses. The error bars are derived from the relative errors given in Table I. (a) Line-averaged continuum Z_{eff} and equivalent CX Z_{eff} line integral for JET #42982 [3]. (b) On-axis continuum Z_{eff} and CX Z_{eff} for JET #61352. (c) Similar, for JET #61346. (d) Z_{eff} profiles for JET #61346 at 5.7 s obtained from both continuum and CX measurements.




Fig. 4. Distributions illustrating the danger of an indeterminate inverse problem. (a) The curve $\theta_2 = x/\theta_1$, with $x = 5$. (b) Gaussian likelihood for θ_1 and θ_2 with the data x given by $x = \theta_1\theta_2 + \nu$ for $x = 5$ and $\sigma_\nu = 1$. (c) Marginal posterior distribution for θ_1 , assuming $0 \leq \theta_1 \leq 10$ and $1 \leq \theta_2 \leq 10$.




Fig. 5. Time traces for the on-axis n_e , respectively T_e , ϵ_{ff} , n_C , ϵ and δ for two JET pulses. The units for ϵ_{ff} have been abbreviated, and are $\text{W}/(\text{cm}^2 \text{sr } \text{\AA})$. (a) #61346. (b) #61348.


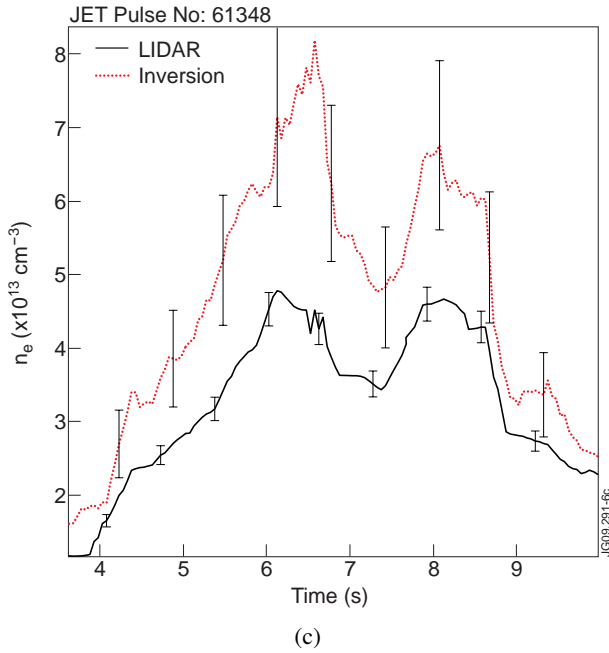




Fig. 6. Time traces for the on-axis n_e and Z_{eff} through deterministic model inversion for JET pulses #61346 and #61348. The traces are compared to the measured n_e and Z_{eff} signals. (a) #61346: n_e . (b) #61346: Z_{eff} . (c) #61348: n_e . (d) #61348: Z_{eff} .




Fig. 7. Contour plot for the (normalized) joint posterior density of the on-axis n_e and Z_{eff} for JET #61346 at 5.7 s. The numbers indicate probability density.




Fig. 8. Marginal posterior distributions for the on-axis n_e and Z_{eff} for JET #61346 at 5.7 s.

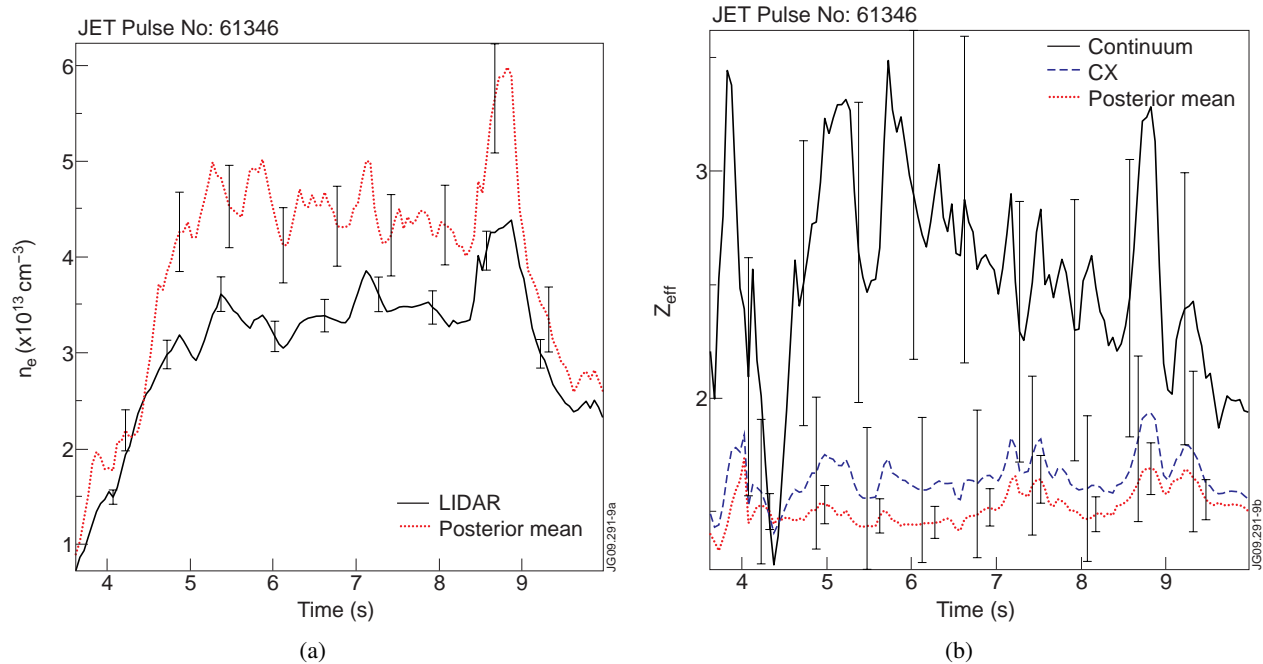


Fig. 9. Marginal posterior means of the on-axis n_e ((a)) and Z_{eff} ((b)) for JET #61346. The error bars on the posterior means correspond to a single posterior standard deviation.




Fig. 10. Marginal posterior distributions for the on-axis n_e and Z_{eff} for JET #61346 at 5.7 s. An additional independent LIDAR density measurement is considered.


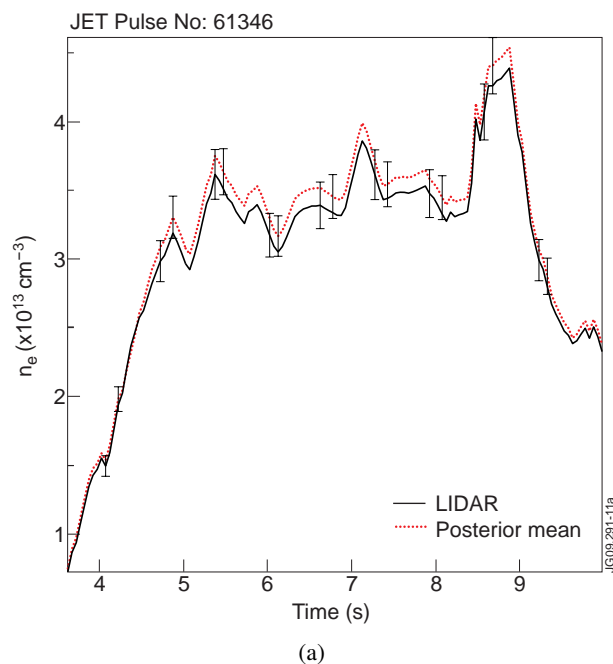


Fig. 11. Marginal posterior means of the on-axis n_e ((a)) and Z_{eff} ((b)) for JET #61346, including an independent density measurement. The Z_{eff} estimate (smallest error bars) overlaps with the CX Z_{eff} .





Fig. 12. Normalized marginal posterior distributions for the on-axis n_e ((a)) and Z_{eff} ((b)) for JET #61346 at 5.7 s, taking into account different data sets.




Fig. 13. MCMC estimated histograms for the quantities of interest in the estimation of n_e and Z_{eff} using artificial data and including systematic uncertainties. The respective means (vertical line) and 68.3% credible intervals (darker shade) are also indicated.

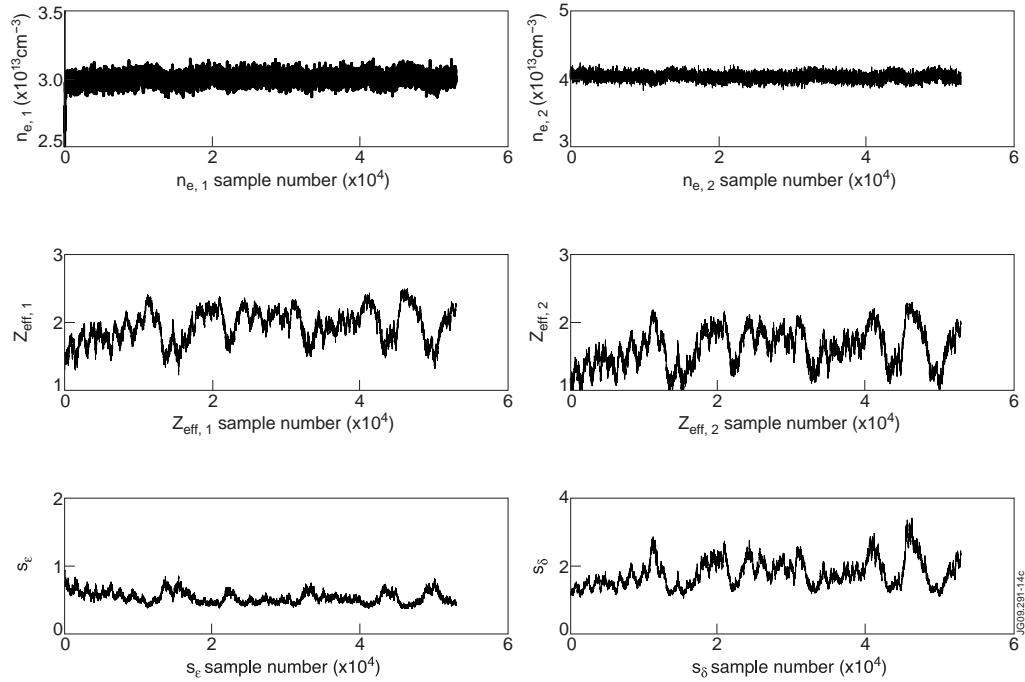


Fig. 14. MCMC traces for the quantities of interest in the estimation of n_e and Z_{eff} using artificial data and including systematic uncertainties.

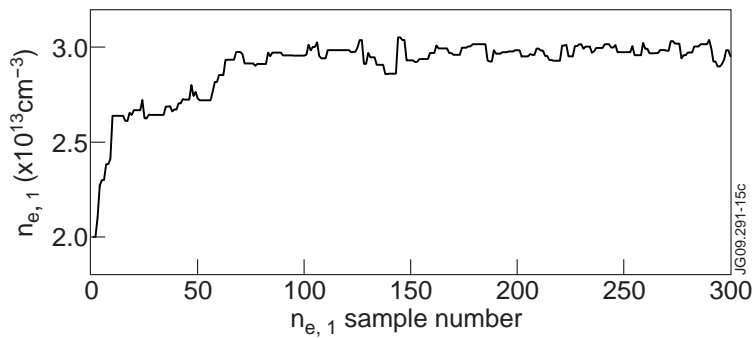


Fig. 15. Zoomed display of the MCMC trace for $n_{e,1}$, showing the initial burn-in period before convergence is established.

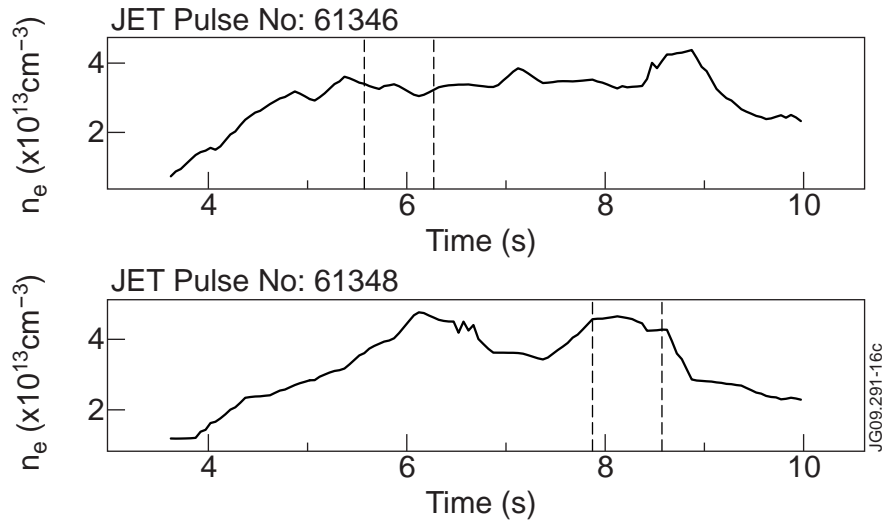


Fig. 16. LIDAR n_e for JET discharges #61346 and #61348, showing the plateau regions delimited by the dashed lines.

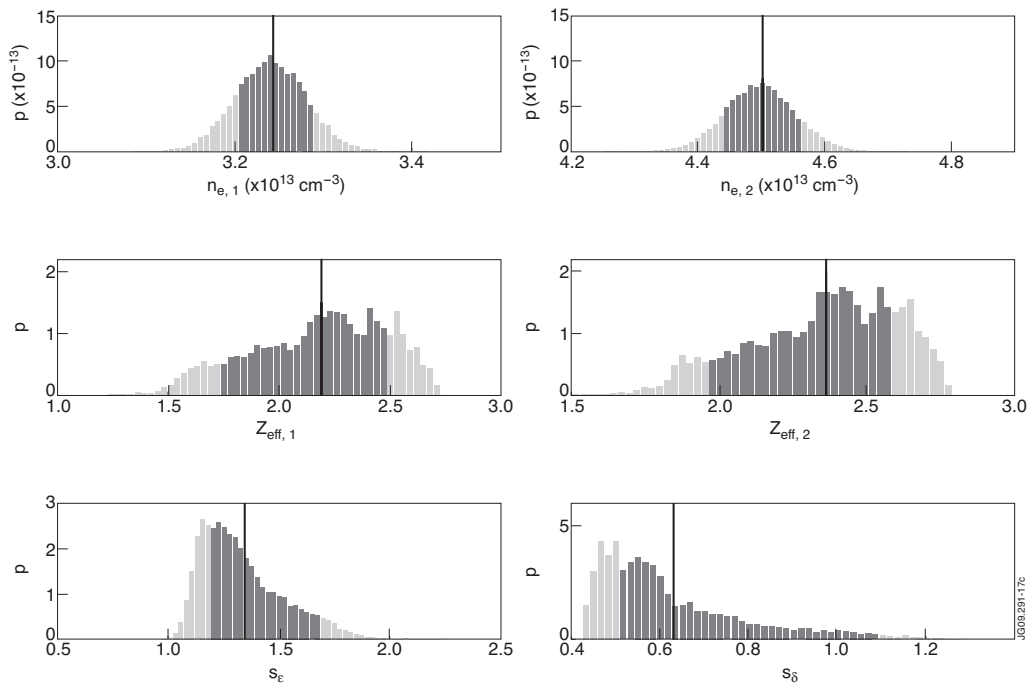


Fig. 17. MCMC estimated histograms for the marginal posterior means for the quantities of interest in the estimation of n_e and Z_{eff} using real data (JET #61346 and #61348) and including systematic uncertainties. The respective means (vertical line) and 68.3% credible intervals (darker shade) are also indicated.

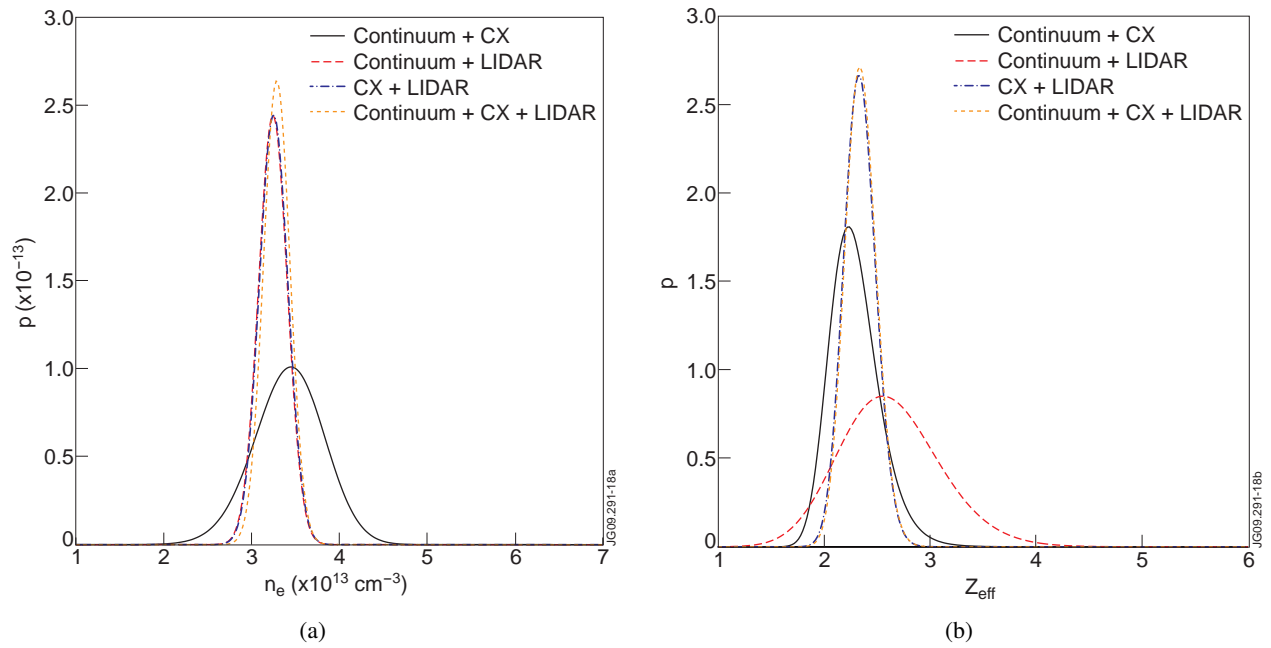


Fig. 18. Normalized marginal posterior distributions for the on-axis n_e ((a)) and Z_{eff} ((b)) for JET #61346 at 5.7 s, taking into account different data sets. The values obtained above for the scale factors s_ϵ and s_δ were included into the posterior.