SANDIA REPORT

SAND2006-7568
Unlimited Release
Printed January 2007

A Bayesian method for characterizing distributed micro-releases: II. Inference under model uncertainty with short time-series data

J. Ray, Y. M. Marzouk, M. Kraus and P. Fast

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94-AL85000.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone:	(865) 576-8401
Facsimile:	(865) 576-5728
E-Mail:	reports@adonis.osti.gov
Online ordering:	http://www.doe.gov/bridge

SAND2006-7568
Unlimited Release
Printed January 2007

A Bayesian method for characterizing distributed micro-releases: II. Inference under model uncertainty with short time-series data

J. Ray and Y. M. Marzouk
Sandia National Laboratories, P. O. Box 969, Livermore CA 94551
jairay,ymarzou@sandia.gov

M. Kraus
250 S. Peterson Blvd, Peterson AFB, CO 80914-3180
Mark.Kraus.ctr@northcom.mil

P. Fast
Lawrence Livermore National Laboratory, Livermore CA 94550
pfast1@llnl.gov

Abstract

Terrorist attacks using an aerosolized pathogen preparation have gained credibility as a national security concern after the anthrax attacks of 2001. The ability to characterize such attacks, i.e., to estimate the number of people infected, the time of infection, and the average dose received, is important when planning a medical response. We address this question of characterization by formulating a Bayesian inverse problem predicated on a short time-series of diagnosed patients exhibiting symptoms. To be of relevance to response planning, we limit ourselves to 3–5 days of data. In tests performed with anthrax as the pathogen, we find that these data are usually sufficient, especially if the model of the outbreak used in the inverse problem is an accurate one. In some cases the scarcity of data may initially support outbreak characterizations at odds with the true one, but with sufficient data the correct inferences are recovered; in other words, the inverse problem posed and its solution methodology are consistent. We also explore the effect of model error—situations for which the model used in the inverse problem is only a partially accurate representation of the outbreak; here, the model predictions and the observations differ by more than a random noise. We find that while there is a consistent discrepancy between the inferred and the true characterizations, they are also close enough to be of relevance when planning a response.

Acknowledgment

This work was supported Sandia National Laboratories' LDRD (Laboratory Directed Research and Development) funds, sponsored by the Computational and Information Sciences IAT.

Contents

Appendix

Figures

Tables

This page intentionally left blank.

1 Introduction

The anthrax attacks of 2001 [4] raised the credibility of aerosolized pathogens being used in bioterrorist attack. Early warning, either in the form of an anomalous increase in syndromes detected by public health monitoring networks [5] or via detection by environmental sensors, holds the highest potential to reduce casualties. However, syndromic surveillance can only provide heightened awareness—it results in neither definite evidence of an attack nor in the identification of the pathogen. Also, the introduction of an aerosolized pathogen into a population may not always be captured on environmental sensors. Examples include small releases that may not travel far, low quality formulations (coarse and heavy particulate matter) which precipitate easily, as well as releases in areas which are not well instrumented. In such a case, the first definitive diagnosis of a patient will be the first intimation of an attack, but by then the disease may have established itself in the population. Being able to infer the characteristics of the release (also referred to as the bioterrorist or BT attack)—i.e., the number N of the people infected, the time τ of infection, and the average dosage $\overline{D}$ received by the infected people—has important ramifications in planning a response [6]. The inferred characteristics can also serve as initial conditions for various epidemic models that can predict the evolution and spread of the disease in a population [7] and its ramification on society [6, 8].

Inferring the characteristics of the outbreak can be challenging. The observables on which inferences are based consist of the time the diagnosed patients turned symptomatic (typically expressed as a time interval during which they developed symptoms) and the location of their residence and place of work. In case of a very mobile population, e.g., a military force engaged in operations, the location of residence and/or work may be hard to define. The model that relates the time of exhibition of symptoms to the characteristics of the genesis of the outbreak is the incubation period distribution, which in many cases is dependent on the dosage received. To be relevant in an operational, consequence-management sense, these inferences have to be drawn early in the outbreak; a time-series obtained from a 3–4 day observation period may be considered representative. Apart from scarcity of observations, the incubation period distribution used in the inferences may be a poor model for the particular instance of the disease. Thus these inferred characteristics are expected to be rather approximate, and quantifying the uncertainty in the characterization becomes a key requirement of the inference process.

In our previous study [9], we developed a Bayesian inverse formulation for BT attack characteristics $(N, \tau, \overline{D})$ and demonstrated the feasibility of inference in simplified settings. We developed probability density functions (PDFs) for N, τ and $\overline{D}$ using a rather unsophisticated but easy-to-implement numerical approach. These results may be considered only qualitatively accurate. In the present study, as in the previous one, we will limit ourselves to temporal analysis; we will not take the location of diagnosed patients into consideration. Further, all tests will be performed with anthrax as the pathogen. Broad uniform priors will be used in the inference process. We will operate within a self-imposed limit of a 3–5 day observation period. We will adopt a more sophisticated

adaptive integration scheme so that the inferred PDFs are quantitatively correct. We will explore how the inferences are affected by the size of outbreak (N), the dosage received ($\overline{D}$), and the frequency with which the data is collected. Inferences will be conditioned on two alternative time series, one with a six hour temporal resolution, and the other with a twenty four hour resolution. We will also show, given suffi cient data, how our approach recovers the correct characterization of the outbreak even though shorter time series may support characterizations signifi cantly different from the true one. This may be considered a demonstration of the consistency of the method. In the interest of realism, we also consider cases where the anthrax model used to generate the observed data (via simulated outbreaks) is different from the one used in inference. We verify the method in settings where every infected patient gets an identical dose as well as one where there exists a large spectrum of doses received. We conclude with an application of this method to the Sverdlovsk outbreak of 1979 [3]. The results of this study will provide a measure of the accuracy and robustness of this Bayesian method, preparatory to extending this purely temporal analysis into a spatio-temporal one.

2 Previous Work

The question of inferring the characteristics of the genesis of an outbreak from a partially observed epidemic has not been extensively studied. Walden and Kaplan [10] developed a Bayesian formulation to estimate the size and time of a bioterrorist attack which they tested on a low-dose anthrax attack corresponding, approximately, to the Sverdlovsk outbreak [3] of 1979, using an incubation period model developed by Brookmeyer [11]. They also demonstrated the use of priors—prior belief regarding the size N of the outbreak—to develop a smooth PDF for N in spite of a small infected population ($N = 100$) and a short time-series (5 days long), with data collected on a daily basis. An alternative approach (maximum likelihood) was used by Brookmeyer and Blades [12] to infer the size of the 2001 anthrax attacks [4], before estimating the reduction of casualties by the timely administration of antibiotics. This inference process was diffi cult due to the small number of symptomatic patients (11 infectees in 3 separate attacks). They also used the anthrax incubation model in [11]. Both [10] and [11] developed similar expression for the likelihood function, i.e., the probability of observing a time series given an attack at time τ with N infected people. The incubation period distribution was not dose-dependent, and hence no dosages were inferred in the two studies.

Signifi cantly more effort has been spent in characterizing the incubation period of inhalational anthrax. The bulk of the work has been experimental, with non-human primates being subjected to anthrax challenges [13, 14, 15, 16, 17, 1]. Brookmeyer et al [11] developed a low-dose incubation period model applicable to the Sverdlovsk outbreak; their more recent work, based on a competing risks formulation, includes dose-dependence [18]. A more empirical study, but based on signifi - cantly more data, was done recently by Wilkening [2], where he compared four different models called Models A, B, C and D. Model D is a slight modifi cation of Brookmeyer's dose-dependent model described in [18]. While Wilkening's Model A agreed with Model D at the high-dose limit, their low-dose behavior was different. Further, Wilkening developed two variants of his Model A, A1 and A2. A1 is a simpler model but its comparison with experimental results is slightly worse than A2. In this study, we will use Wilkening's Model D for simulated BT attacks while Model A2 will be used in the inference scheme. A complete discussion is deferred till Sec. 3.2.

The issue of whether a person exposed to a number of spores will actually contract the disease is a separate question that will not be addressed in the study. We will concentrate on inferring the number of people who are actually infected, not merely exposed to the pathogen. The problem of estimating the probability of infection from D spores was addressed by Brookmeyer et al in [18] as well as by Glassman [19] and Druett et. al [20]. Haas [21] has established that even at low doses, the risk to *large* populations is *not* statistically insignifi cant.

An effort with aims similar to ours is the Bayesian Aerosol Release Detector (BARD) [22]. It poses an inverse problem to infer the location and height of an anthrax release (the approach is general but has only been tested with anthrax), the time of release and the quantity of material released. The observables are the number of respiratory visits to emergency departments collated in 24-hour intervals and by zip code - such information can be obtained from typical syndromic surveillance systems such as RODS [23, 24]. The model that relates the observables to the characteristics of the outbreak includes a Gaussian dispersion plume [25], Glassman's infection relation [19], and

a log-normal incubation period distribution; different means and standard deviations are used depending upon the dosage. The patient is assumed to visit the emergency department on completion of incubation. BARD develops a likelihood ratio - the ratio of the probability of obtaining the observables due to an anthrax attack to the probability of observing them in its absence - and if it is above a preset threshold, develops a posterior probability distribution for the location of the release, the quantity released and the time of release. It also supplies the zip codes of the location that contributed most to the result. BARD does not calculate the the number of people infected or the dosage - however, given PDFs for the location and quantity, the magnitude of the outbreak and the dosage may be trivially obtained by using the BARD inferences as the initial condition in a Gaussian plume to disperse the aerosol and using Glassman's (or Druett's [20]) model to decide the probability of infection.

3 The Inverse Problem

In this section we formulate the inverse problem to infer the characteristics of a BT attack. In particular, we infer $\{N, \tau, \overline{D}\}$, where N is the number of people infected, τ is the time of the infection (counted in terms of days before the first patient exhibits symptoms), and $\overline{D}$ is the average dosage received by each person. The formulation is Bayesian and we develop a joint probability distribution for $\{N, \tau, \overline{D}\}$, thus capturing uncertainty in the inference conditioned on the data. The observables (input to the inverse problem) are embodied in the time-series $\{t_i, n_i\}, i = 0 \ldots M$, where n_i is the number of people developing symptoms in the time-interval $(t_{i-1}, t_i]$. M is the length of the time-series and is expected to be small—we limit the observation period to 3–5 days to be of relevance in response planning and consequence management. t_0 corresponds to the time when the first symptomatic patient is identified. For lack of additional information, we will use broad uniform priors in our Bayesian inverse problem.

In Sec. 3.1 we derive the inverse problem and the expressions for the PDFs. In Sec. 3.2 we describe the dose-dependent anthrax incubation period models used. We use Wilkening's [2] Model A2 in the inverse problem and Models A2 and D for simulating BT attacks. In Sec. 3.3 we verify our inverse problem. In particular, we demonstrate how as we increase M, the inverse problem recovers the correct values of $\{N, \tau, \overline{D}\}$ even if shorter time-series support an alternative set of inferences for the BT attack's characteristics.

3.1 Formulation of the Problem

Consider an attack at time τ where N people are infected, with each of the N people receiving the same dose of D anthrax spores. The incubation period obeys a dose-dependent log-normal distribution; we refer to its cumulative distribution function (CDF) as $C(t, D)$. For a few days M (say 3–5 days) we can expect (1) a series t_i, $i = 0 \ldots M$, of times, perhaps the endpoints of 24-hr intervals, when patients' symptoms are observed and (2) the series n_i, $i = 0 \ldots M$, of new patients who turned symptomatic between $t_i - \Delta t$ and t_i where $t_i - t_{i-1} = \Delta t, i \neq 0$, and Δt is a constant. We define survival probability as $P_{surv}(t, D) = 1 - C(t, D)$.

We can state the problem as such: Given a time-series $(t_i, n_i), i = 0 \ldots M$, of patients showing symptoms over a few days, estimate (N, τ, D) from these data. n_i patients are assumed to have developed symptoms over the time interval between t_{i-1} and t_i. We will henceforth assume that the time-interval $t_i - t_{i-1} = \Delta t$ is a constant, typically 6- or 24-hours.

Let $L = \sum_{i=0}^{M} n_i$ be the total number of people who have developed symptoms by t_M. Thus $N - L$ infected people are still asymptomatic and the probability of such an event is $\{P_{surv}(t_M - \tau, D)\}^{N-L}$. The probability that n_i people will develop symptoms in the time interval between t_{i-1} and t_i is $\{C(t_i - \tau, D) - C(t_i - \Delta t - \tau, D)\}^{n_i}$. Since L symptomatic people may be chosen out of a total infected population of N in

$$\frac{N!}{(N-L)!L!}$$

ways and L symptomatic people can be divided into the sequence $n_0, n_1, \ldots n_M$ in

$$\frac{L!}{n_0! n_! \ldots n_M!}$$

different ways, the probability of observing the $\{t_i, n_i\}, i = 0 \ldots M$ time-series given a BT attack characterized by (N, τ, D), or equivalently, the likelihood function $\mathcal{L}$, is

$$\begin{aligned}
\mathcal{L}(N,\tau,D) &\equiv p(\{t_i,n_i\}, i=0\ldots M | N,\tau,D) \\
&= \frac{N!}{L!(N-L)!} \times \frac{L!}{n_0! n_1! \ldots n_M!} \times \{P_{surv}(t_M - \tau, D)\}^{N-L} \times \\
&\quad \prod_{i=0}^{M} \{C(t_i - \tau, D) - C(t_{i-1} - \tau, D)\}^{n_i} \\
&= \frac{N!}{(N-L)! \prod_{i=0}^{M} n_i!} \times \{P_{surv}(t_M - \tau, D)\}^{N-L} \times \\
&\quad \prod_{i=0}^{M} \{C(t_i - \tau, D) - C(t_{i-1} - \tau, D)\}^{n_i}. \qquad (1)
\end{aligned}$$

We postpone the discussion of the CDF $C(t, D)$ till Sec. 3.2.

Eq. 1 has to be incorporated into an expression that allows inference. Exploiting Bayes rule, we obtain

$$\pi(N,\tau,D|\{t_i,n_i\}, i=0\ldots M) = \frac{\mathcal{L}(N,\tau,D)\pi_N(N)\pi_\tau(\tau)\pi_D(D)}{\pi(\{t_i,n_i\}, i=0\ldots M)} \qquad (2)$$

where π_N, π_τ and π_D are the priors for N, τ and D and $\pi(\{t_i, n_i\}, i = 0 \ldots M)$ is the probability of observing a $\{t_i, n_i\}, i = 0 \ldots M$ time-series in any circumstance. Note that the denominator simply acts as a normalizing constant and is implicitly obtained by normalizing this expression in Eq. 3. In this study, we use broad uniform distributions as priors. The joint probability distribution $\pi(N, \tau, D|\{t_i, n_i\}, i = 0 \ldots M)$ is normalized as

$$\overline{P}(N,\tau,D) = \frac{\pi(N,\tau,D|\{t_i,n_i\}, i=0\ldots M)}{\int_{N=N_S}^{N=N_E} \int_{\tau=\tau_S}^{\tau=\tau_E} \int_{D=D_S}^{D=D_E} \pi(N,\tau,D|\{t_i,n_i\}, i=0\ldots M) dN d\tau dD} \qquad (3)$$

and marginalized to obtain individual PDFs g_N, g_τ and g_D for N, τ and D viz.

$$\begin{aligned}
g_N(N) &= \int_\tau \int_D \overline{P}(N,\tau,D) dD d\tau \\
g_\tau(\tau) &= \int_N \int_D \overline{P}(N,\tau,D) dD dN \\
g_D(D) &= \int_\tau \int_N \overline{P}(N,\tau,D) dN d\tau \qquad (4)
\end{aligned}$$

A more detailed derivation can be found in [9].

The multidimensional integrals in Eq. 4 are evaluated using an iterative adaptive Monte Carlo method, often referred to as the VEGAS method [26]. It was observed that when the $\{t_i, n_i\}, i = 0 \ldots M$ time-series was short, the integrand $\overline{P}$ could be noisy, mostly zero in the bulk of the domain with sharp peaks in certain localized regions. Adaptive quadrature approaches, though more accurate than Monte Carlo, were far slower, especially with short time-series data. We used the GSL [27] library's implementation of the VEGAS algorithm in this study.

3.2 Anthrax Incubation Models

This subsection tersely states mathematical models we use to predict the onset time of symptoms in a person exposed instantaneously to D anthrax spores. The onset time is a random variable, and therefore described by a cumulative distribution function (CDF).

The CDF for Wilkening's "Model D" is given by (see App. A)

$$C_{ModelD}(t,D) = \int_0^t F(t-s;D,\lambda,\theta)g(s;D)ds, \tag{5}$$

which is a convolution of $F(t;D)$, the probability that at least one spore out of a dose of D spores will germinate into a vegetative anthrax cell by time t and $g(s;D)$ which is the PDF of the time s taken, post-germination, to reach a bacterial load at which time symptoms appear. F and g are defi ned as

$$\begin{aligned} F(t;D,\lambda,\theta) &= \frac{1}{p}\left(1-\exp\left(-\frac{D\lambda}{\lambda+\theta}Q(t)\right)\right), \quad \text{where} \\ Q(t) &= 1-\exp(-(\lambda+\theta)t) \end{aligned} \tag{6}$$

and

$$g(s;D) = \frac{1}{\sqrt{2}\pi\sigma_s s}\exp\left(-\frac{1}{2}\frac{\log^2(s/M_s)}{\sigma_s^2}\right). \tag{7}$$

These distributions depend on a number of parameters:

- N_{thresh}, a threshold bacterial load in a person that causes symptoms
- t_2, the bacterial load doubling time in a given medium (e.g. mediastinal lymph nodes where the spores germinate), which can be obtained in *in vitro* laboratory experiments
- t_M, which is the time required to reach a bacterial load of N_{thresh} and is given by

$$t_M = t_{lag} + \frac{t_2}{\log(2)}\frac{N_{thresh}}{D}$$

- t_{lag}, a lag time in bacterial growth experiments (typically 1 hour).
- σ_s, the variance of the log of the time required to reach the symptomatic bacterial load
- θ, the probability rate of clearance of a spore (by the immune system), specified in terms probability of clearance per spore per day
- λ, the probability rate of germination of a spore, specified in terms of probability of germination per spore per day

M_s, the median time to symptoms, is set to t_M in Wilkening's models used here. The values of the parameters for Model D are $\theta = 0.109$ day^{-1}, $\lambda = 8.79 \times 10^{-6}$ day^{-1}, $t_{lag} = 1$ hour, $t_2 = 2.07$ hour, $N_{thresh} = 10^9$ and $\sigma_s = 0.544$ day^{-1}.

Sartwell [28] found that the incubation period for a number of diseases were log-normally distributed, which is at odds with Eq. 5. Wilkening's Model A2 assumes a log-normal distribution,

$$C_{ModelA2}(t,D) = \frac{1}{2}\left[1 + \mathrm{erf}\left(\frac{\ln(t/t_0)}{\sqrt{2}S}\right)\right], \quad S = 0.804 - 0.079\log_{10}(D) \ . \tag{8}$$

where t_0, the median incubation time, is obtained by solving the integral equation obtained from Eq. 5

$$0.5 = \int_0^{t_0} F(t_0 - s; D, \lambda, \theta) g(s; D) ds.$$

However, while solving for t_0, Wilkening used a slightly different set of parameters: $\theta = 0.11$ day^{-1}, $\lambda = 8.84 \times 10^{-6}$ day^{-1}, $t_{lag} = 1$ hour, $t_2 = 2.06$ hour, and $\sigma_s = 0.542$ day^{-1}. The reason for the slight change in parameters as well as the difference between Models A2 and D is discussed below. This completes the specification of the Model A2.

The parameters in Eq. 5 and 8 were obtained by fitting to the median incubation period as observed in experiments with non-human primates (performed by Henderson et al [13] and Friedlander et al [16]) and the data from the Sverdlovsk outbreak. However, the average dosage during the Sverdlovsk outbreak had to be inferred from atmospheric dispersion models and the probability of exhibiting symptoms (in infinite time) given a dose of D spores. This was done by Wilkening [2]. If one uses Glassman's model [19] for the probability of infection, one obtains an average dose of 2.4 spores. Alternatively, if one uses Eq. 11 (which is similar to Druett's [20] in form and was used by Brookmeyer in [18]; see App. A) one obtains a dose of 300 spores. Wilkening retained both the possibilities and incorporated them into separate models. Model D is based on a dose of 300 spores at Sverdlovsk while A2 assumes 2.4 spores.

In Fig. 1, we plot the median incubation period as predicted by Model A2 and D, as a function of dosage D. The dosage at Sverdlovsk, inferred as 2.4 spores (represented by •) is used to derive the parameters for Model A2 (solid line); the alternative inference of 300 spores (represented by a

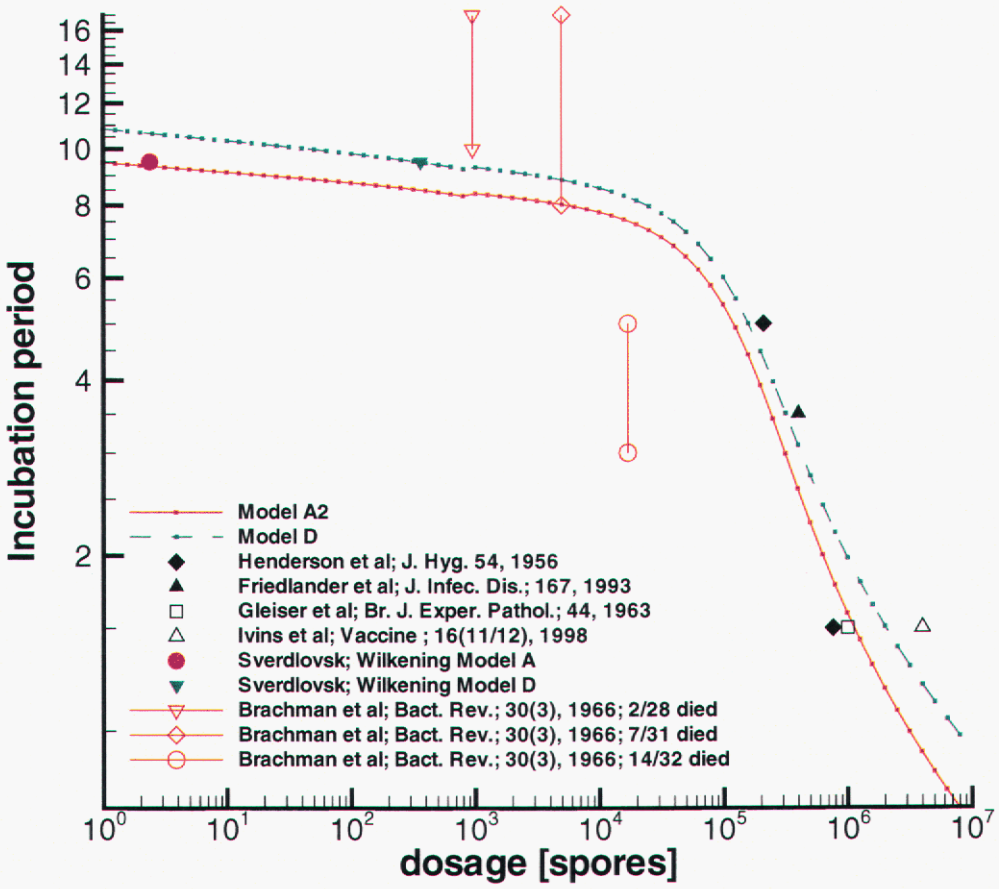

Figure 1. The median incubation period for anthrax as a function of dosage D. The solid line is Model A2 which assumes a dose of 2.4 spores at Sverdlovsk, while the dashed line is Model D, which assumes 300 spores. The solid symbols are median incubation periods which were obtained from experimental investigations or from the data from the Sverdlovsk outbreak. Symbols which are not filled denote experiments where the population of primates were too small to draw statistically meaningful results. The experiments by Brachman et al [1] are shown by vertical lines between symbols. In these tests, only the lower and upper bounds of the incubation period were provided. These were not used for determining model parameters and are only provided for reference.

fi lled ∇) is used for Model D (dashed line). Studies by Henderson [13] with 2.1×10^5 and 7.6×10^5 spores (represented as fi lled $\diamond$) and Friedlander with 3.5×10^5 spores (represented by fi lled $\triangle$) were used to derive the parameters of both the models. Studies by Ivins et al [17] (unfi lled $\triangle$) and Gleiser et al [15] (unfi lled $\square$) were conducted with very few primates and consequently are plotted only for reference. Brachman [1] conducted studies where he tried to simulate the effect of a low dose, received regularly over an extended period of time, as might be the case in a contaminated wool-sorting mill. The primates went through extended periods when they received no spores at all. The dose was calculated as the total number of spores breathed in and was generally low, between 1000 and 10,000 spores. We plot the ranges of incubation periods observed (only the range was provided) for various dosages for reference.

We see that the tests by Gleiser et al and Ivins et al agree with both the models, which in turn agree with each other, except at the low dose limit. Brachman's tests show median incubation periods which are at odds with the models' predictions; however the mode of infection, that approximating a continuous, low-level infection process spread over days or months was very different from the quick (timescale of an hour) challenge one would expect in a BT attack. Both the models show a kink at $D = 10^3$; this is because they are evaluated with a lower value of λ ($= 1.3 \times 10^{-6}$ day^{-1}) corresponding to a primate ID_{50} of 55,000 spores for comparison with primate anthrax challenge results at the high dose limit, while the low dose predictions were developed with a human ID_{50} of 8600 spores for comparison with the inferences of the dose received at Sverdlovsk. To the best of the authors' knowledge, this is the sum total of experimental data obtained from anthrax challenges of non-human primates where incubation times were measured. We have omitted a study by Klein et al [29] in which an incubation period increase was observed with increasing doses, because only one primate was subjected to a given dose, making the behavior statistically unreliable.

3.3 Verification of the Inverse Problem

In this section we verify the inverse problem formulated in Sec. 3.1. We use Wilkening's Model A2, as described in Sec. 3.2, to simulate inhalational anthrax outbreaks of different sizes. Since the model provides a PDF of the incubation period, the result (time-series) of a simulated outbreak is stochastic i.e. it can be thought of as the sum of an expected outbreak behavior and a noise term. The same model is used for inference i.e. there are no systematic errors between the inference and simulation models. Thus the uncertainty in the inference is due to incomplete observation (a short time-series) and the effect of the noise. We investigate how the quality of the inference varies with the size of the outbreak and the dosage received. We also investigate whether a higher-resolution time-series spanning the same observation period performs signifi cantly better than a low-resolution one.

In Table 1, we list time-series at 6-hour resolution; that is, the number of patients showing symptoms collected over 6-hour intervals obtained from 6 simulated outbreaks. All infected patients received an identical dose D. N indicates the number of people infected who will develop symptoms over time. τ is the time of attack, measured in days prior to the exhibition of symptoms in the fi rst diagnosed patient. Thus, τ is always a negative quantity.

In Figs. 2, 3, 4, 5, 6 and 7, we plot the resulting PDFs for $\{N, \tau, D\}$, for Cases A, B, C, D, E and F, per the time-series listed in Table 1. In Table 2, we summarize the MAP (maximum a posteriori) estimates and 90% confi dence intervals (CI) for N, τ and $\log_{10}(D)$ as obtained from the high-resolution time-series corresponding to 5 days of data. We see that the MAP estimate for N (the value of N corresponding to the peak of the PDF) is generally close to the correct value even with 3 days of data; increasing the length of the observation period to 5 days usually sharpens the PDF, thus showing a reduction in uncertainty with more data. This observation holds true for small attacks ($N = 10^2$) as well as for large ones ($N = 10^4$). An exception is Case F, which will be discussed later. Using a 6-hour resolution time-series, as opposed to a daily one, has the effect of sharpening the PDFs. The time of attack τ is also inferred quite easily, except for the small, low-dose Case A. Larger attacks (Case D, E and F) have sharp PDFs for τ, as compared to Cases A, B and C. This is a direct consequence of the number of symptomatics in the time-series. The PDF for τ is naturally affected by the temporal resolution of the time-series, since it decides the resolution with which PDF for τ can be developed. However, the resolution does not automatically translate into a lower degree of uncertainty. While the PDFs developed from the 6-hour time-series are smoother that those developed from the 24-hour time-series, they are not always sharper/narrower. The dosage D is the hardest to infer. Figs. 2, 3 and 4 show that it is virtually impossible to infer the dosage for small ($N = 10^2$) attacks - the PDFs spread over 5 orders of magnitude. Table 2 shows that the MAP estimates for dosage for the small attacks are incorrectly inferred. Large attacks ($N = 10^4$) allow the development of PDFs for D which are slightly more informative.

Figs. 5, 6 and 7 demonstrate how data from a partially observed outbreak (i.e. data from the fi rst 3-5 days of an attack) may support multiple hypotheses, and at times, support the "wrong" hypotheses more than the correct one. Case D shows peaks in the PDF at $N \sim 4 \times 10^3$ and $N \sim 10^4$. Peaks in the PDF for dosage (measured as $\log_{10}(D)$) occur at 10 spores and between 10^4 and 10^5 spores. These two PDFs, taken together, indicate that the data could be generated from a smaller, high-dose attack or a larger, low-dose attack. For this particular case, the PDFs overwhelmingly indicate a large, low-dose attack, which is the correct characterization. A similar ambiguity is observed in Fig. 6. In both these cases, the higher-resolution time-series did rule out the alternative hypotheses. In Case F (Fig. 7; $N = 10^4, \tau = -1.5, D = 10^4$), we observe a similar existence of 2 hypotheses, with data initially (i.e. up to Day 5) supporting the wrong hypothesis (a larger, low-dose attack) far more than the correct one. Increasing the temporal resolution makes matters worse as the peak at the higher end in the PDF for dosage actually shrinks. We characterize the large low-dose attack approximately as $N = 16,000, \tau = -1$ with a dosage of 10 spores. The evolution of this large, low-dose attack is plotted in Fig. 8. Over the fi rst 5 days of the outbreak's evolution, the inferred low-dose attack behaves very similar to the actual attack ($N = 10^4, \tau = -1.5, D = 10$ spores) though, as indicated by Fig. 7, the data supports the low-dose attack better. In Fig. 8, inset, we plot the evolution of both the attacks, along with the observed data (in symbols), over 10 days. It is clear that the large, low-dose attack, as inferred with 5 days of data, does not fi t the data at all, and *would not be supported if we had a time-series that spanned 7-8 days.* This is verifi ed by plotting the PDFs for N, τ and $\log_{10}(D)$, developed from time-series 6 and 7 days long, in Fig. 7. However, such a long observation period would be irrelevant for consequence planning purposes. We stress that while a Bayesian analysis may identify competing hypotheses, the degree of belief assigned to each is determined by data. In a partially observed attack, this characterization may be completely erroneous, especially if the data is scarce. The only possible remedy is the use of informative

priors for N, τ and/or D, not the broad uniform priors used in this study. In their absence, such ambiguities will remain and will have to be accounted for in the consequence management plans which use these inferences. How this might be done is beyond the scope of this study, but Sec. 5 outlines a few possibilities.

To summarize, solution of the inverse problem successfully provides N and τ for small and large attacks. D can be inferred for large attacks. The PDFs are sharper for large attacks and for high-dose attacks. The effect of higher temporal resolution, which captures the evolution of the outbreak better, is generally to reduce the uncertainty caused by noise and to sharpen the PDFs. When working with partially observed data, the Bayesian method may identify multiple hypotheses which are supported, to a larger or smaller extent, by the data. In some cases, e.g. Case F, the data might even support (initially) the wrong hypothesis, but the correct characterization is recovered as more data becomes available. This may be considered a proof of consistency of this inverse problem.

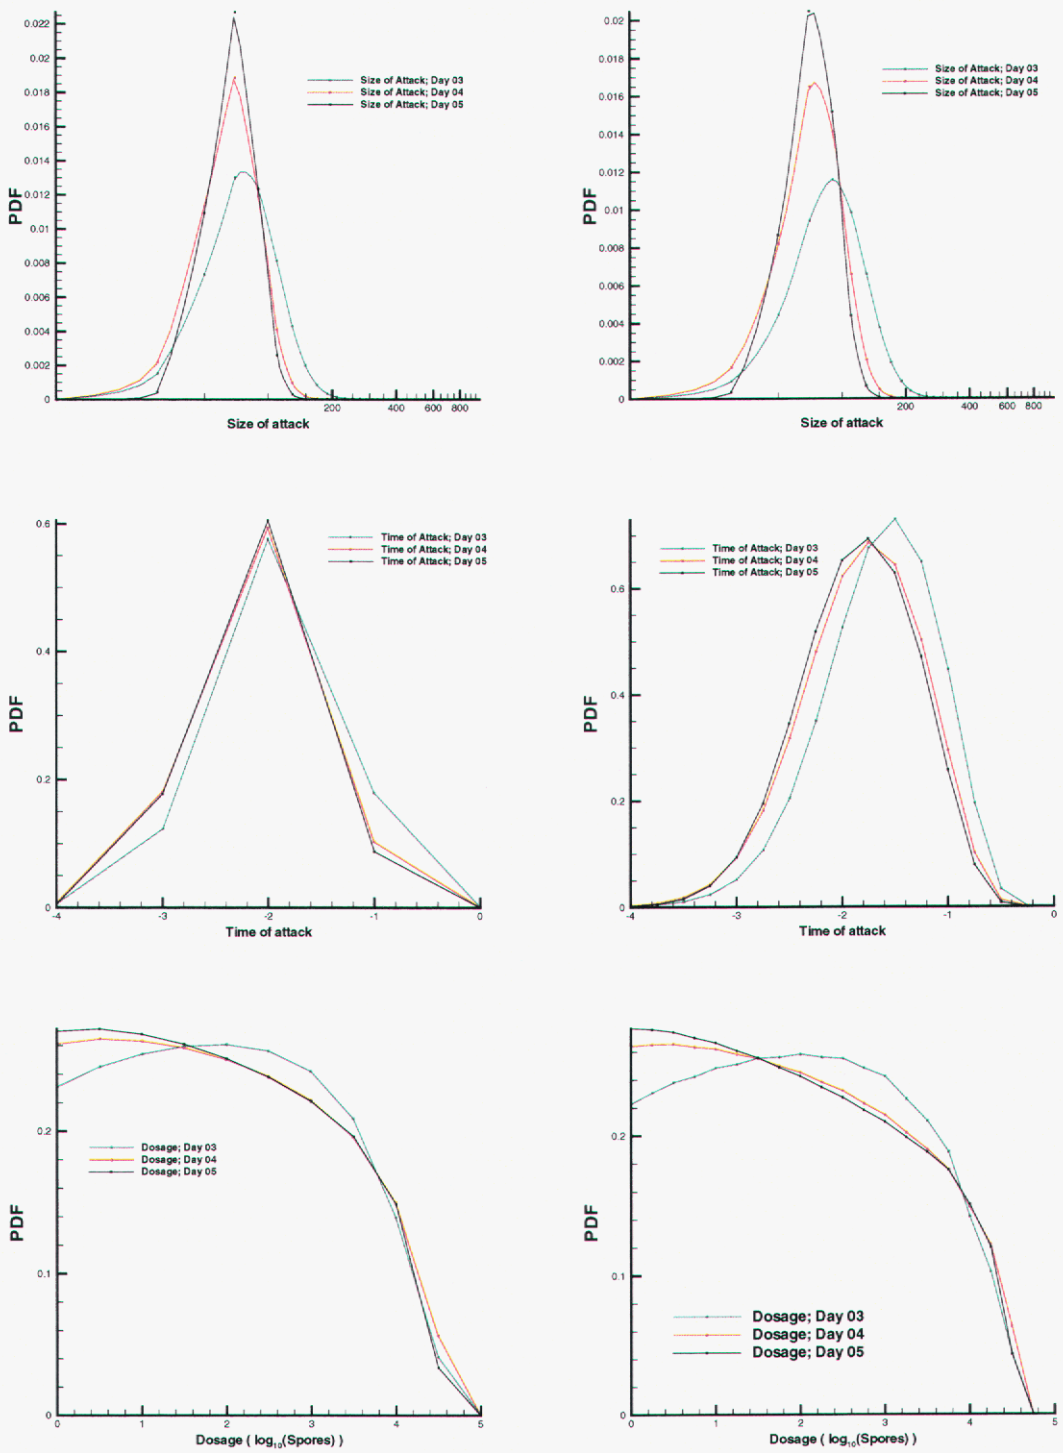

Figure 2. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case A, as tabulated in Table 1. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, D\}$ are $\{10^2, -0.75, 10^0\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

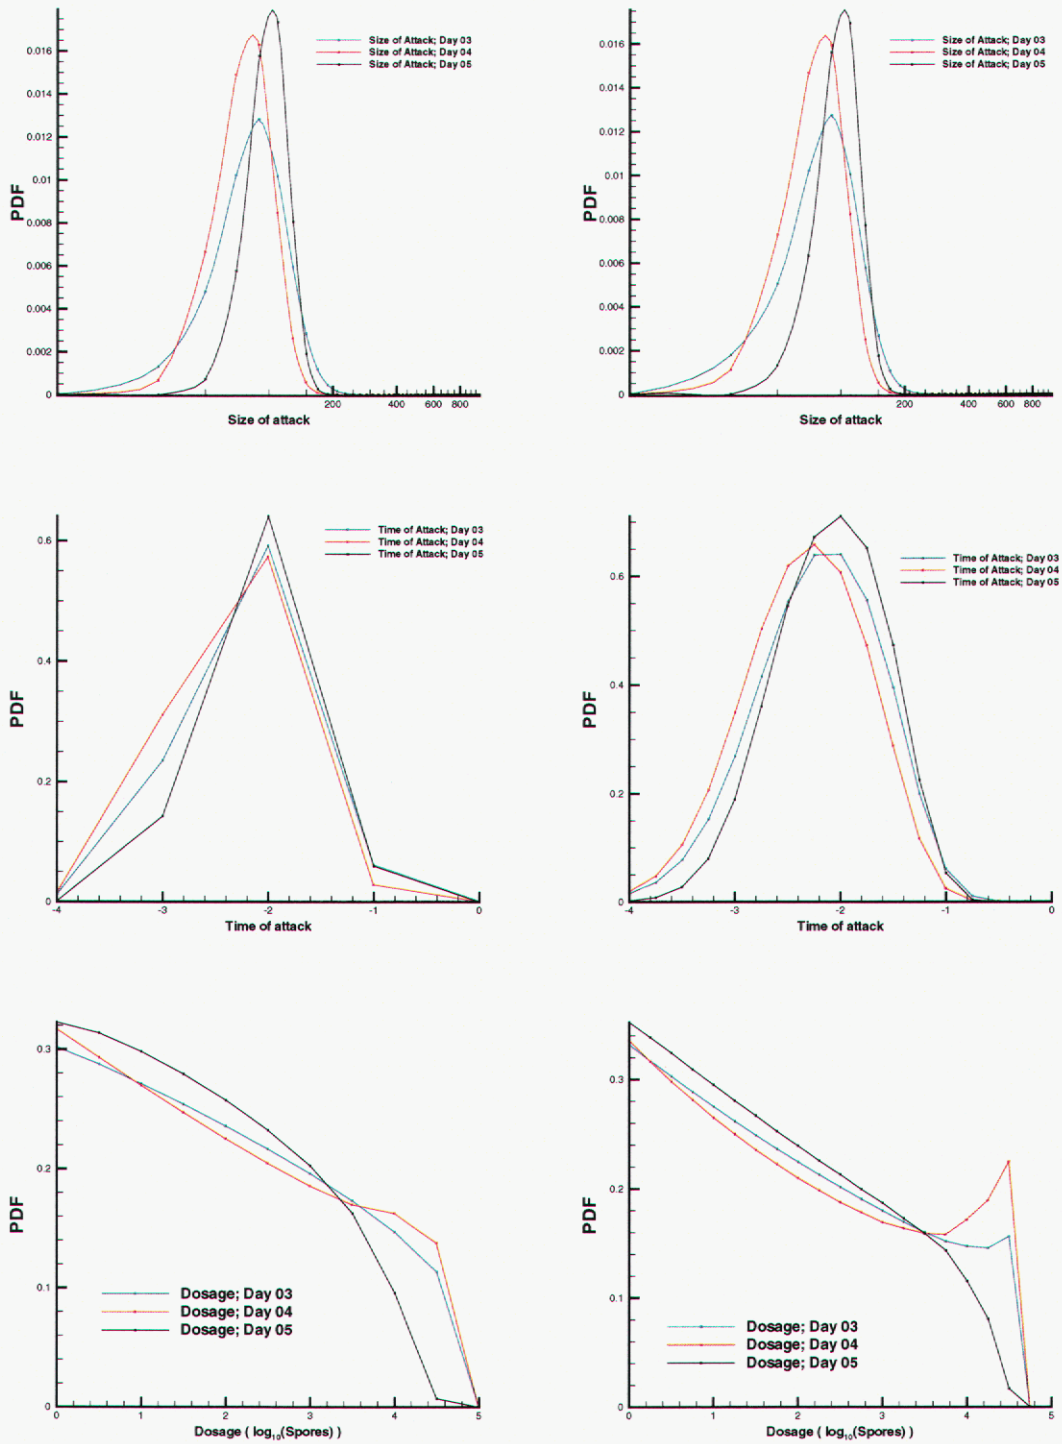

Figure 3. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case B, as tabulated in Table 1. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, D\}$ are $\{10^2, -2.25, 10^2\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

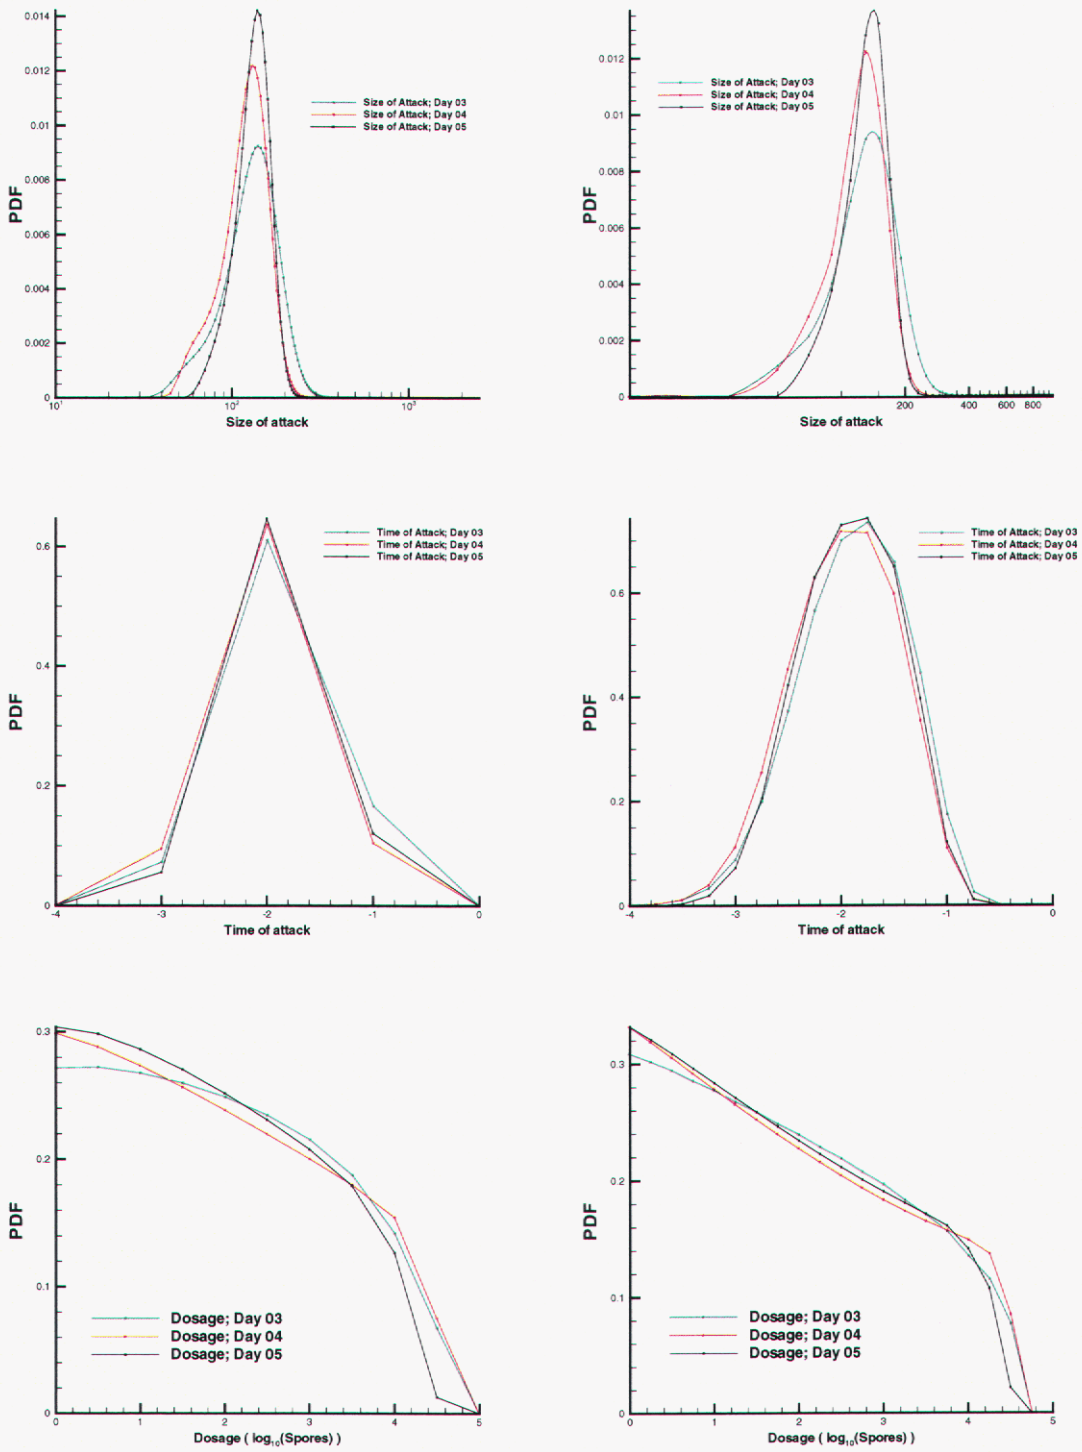

Figure 4. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case C, as tabulated in Table 1. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, D\}$ are $\{10^2, -2.25, 10^4\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

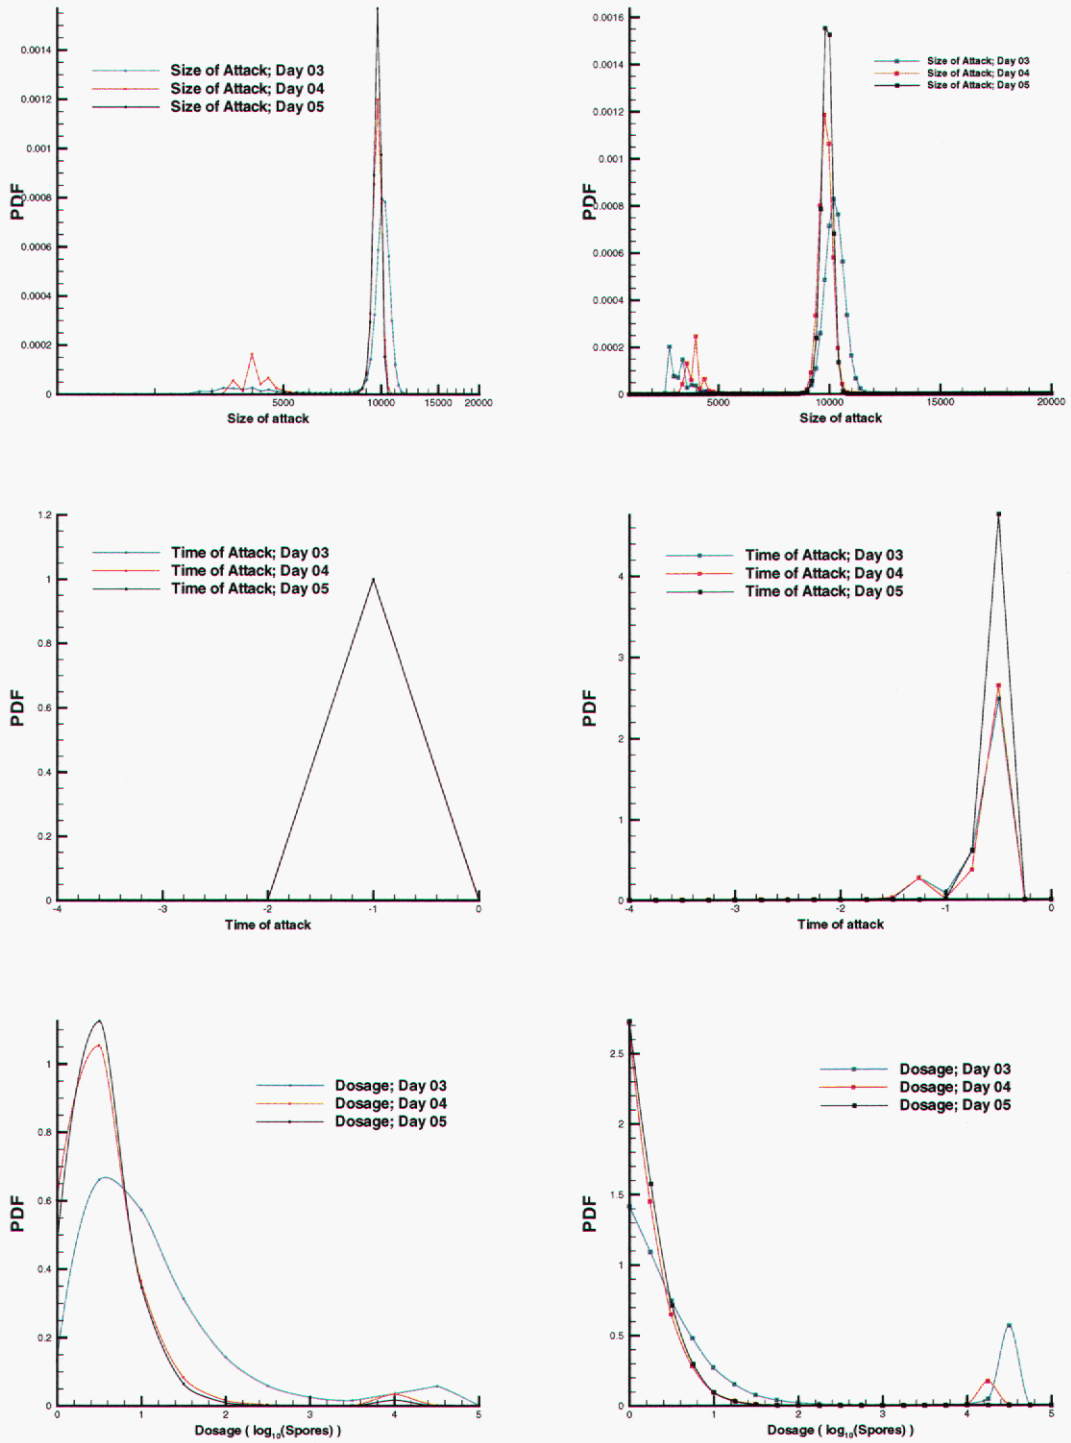

Figure 5. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case D, as tabulated in Table 1. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, D\}$ are $\{10^4, -0.5, 10^0\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

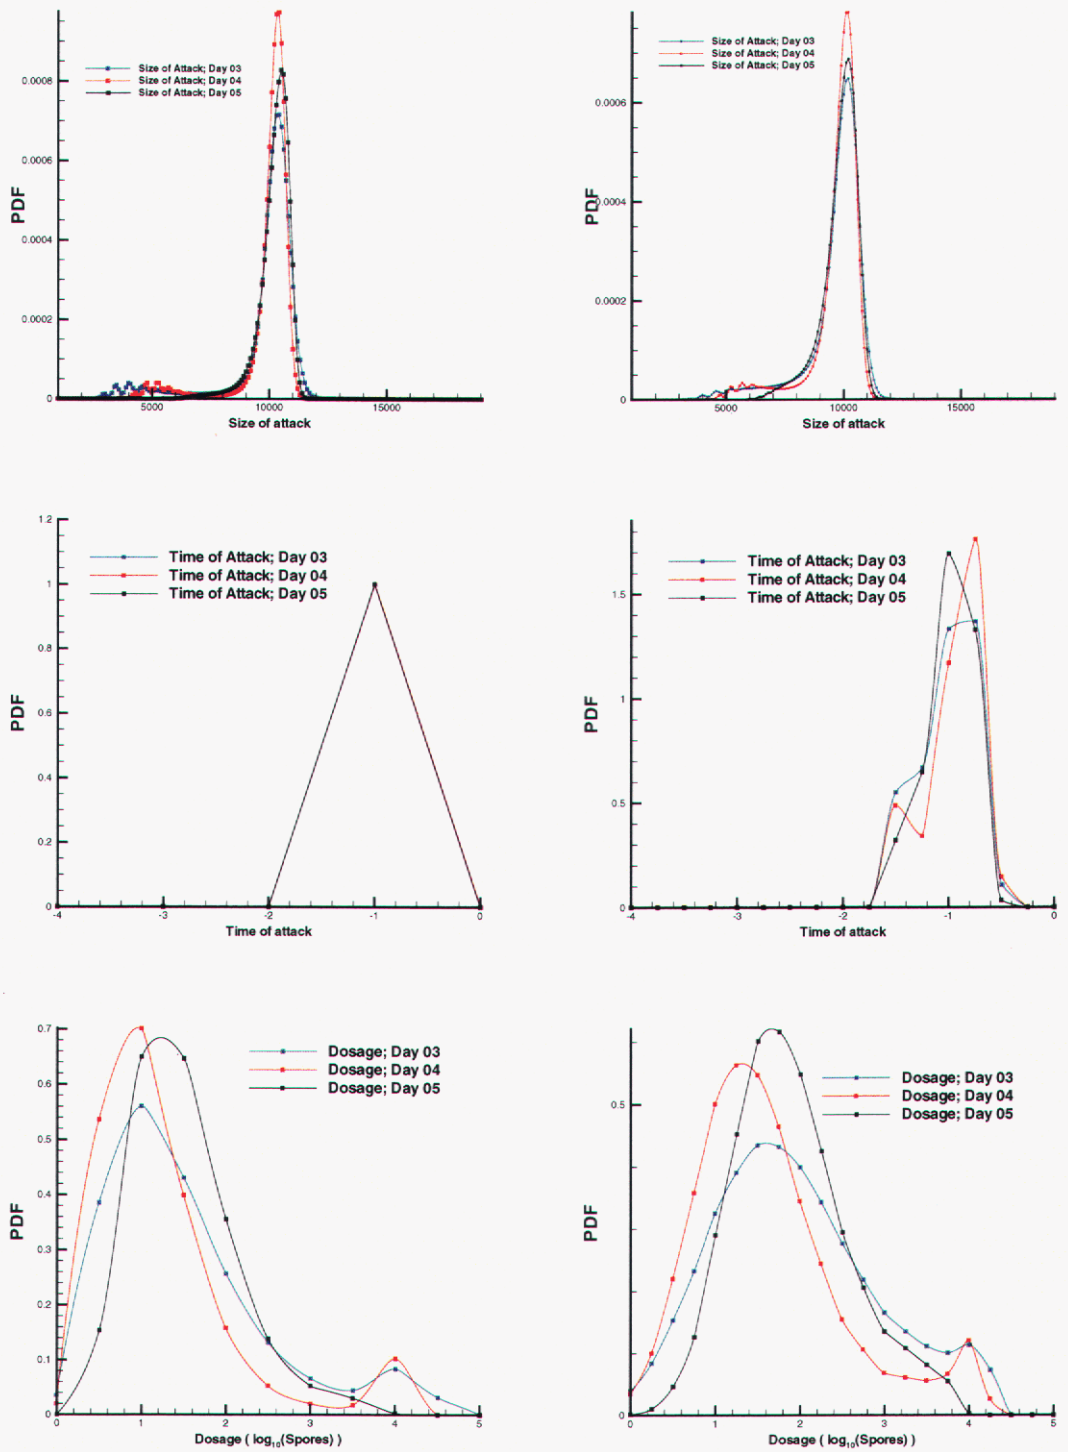

Figure 6. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case E, as tabulated in Table 1. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, D\}$ are $\{10^4, -1.0, 10^2\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

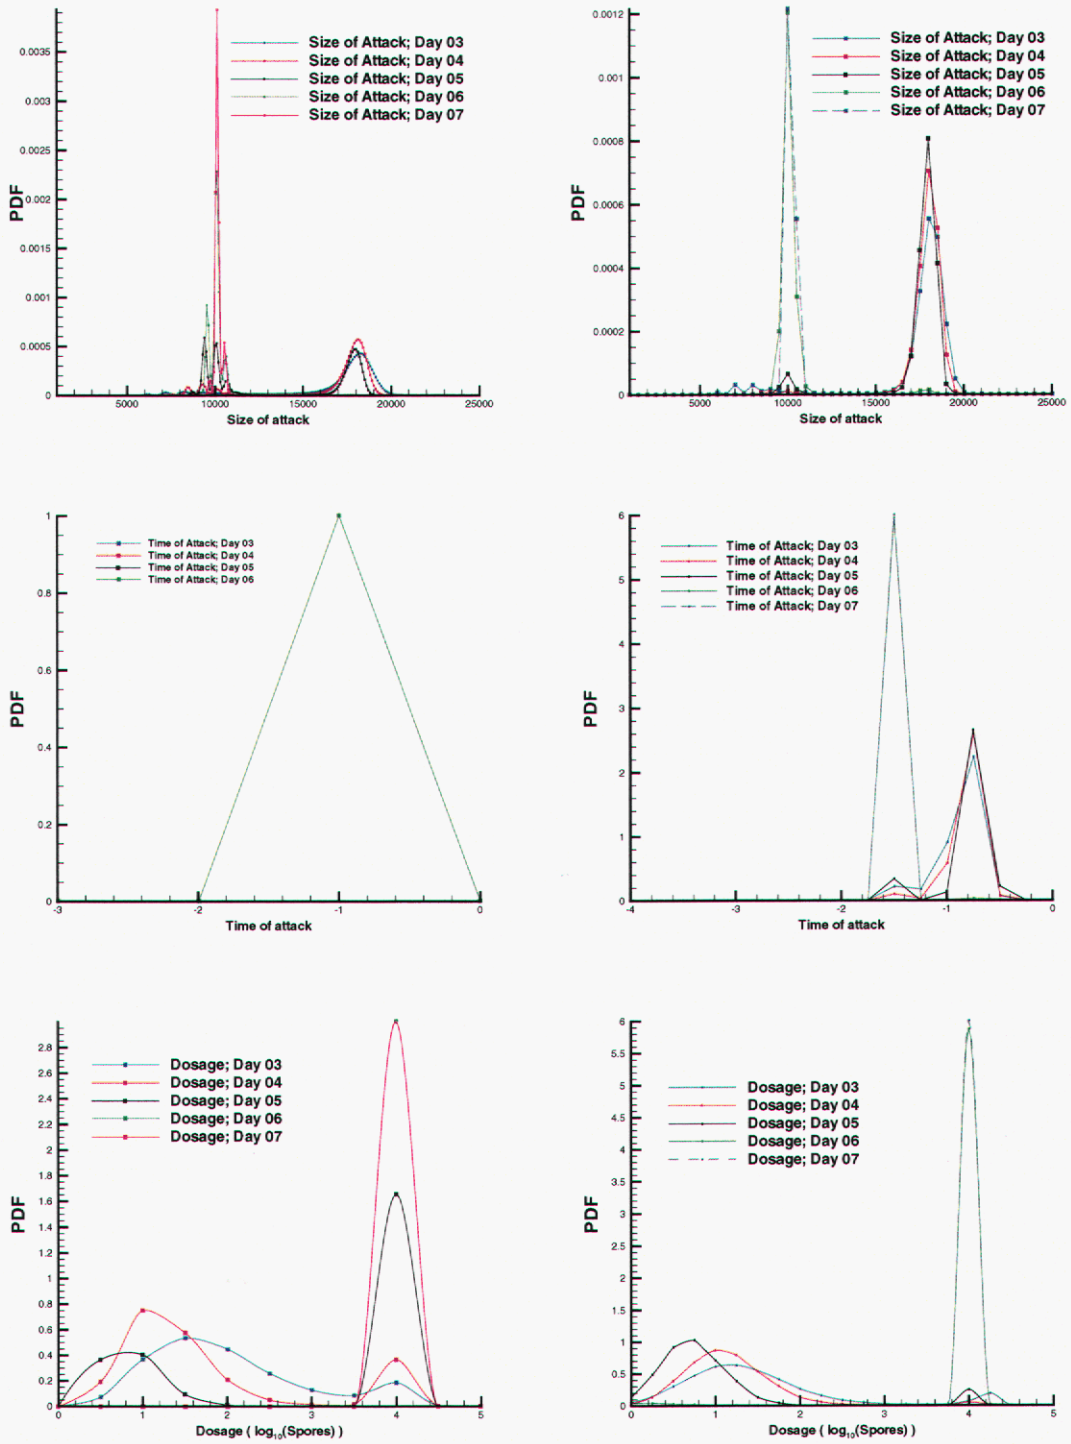

Figure 7. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case F, as tabulated in Table 1. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, D\}$ are $\{10^4, -1.5, 10^4\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

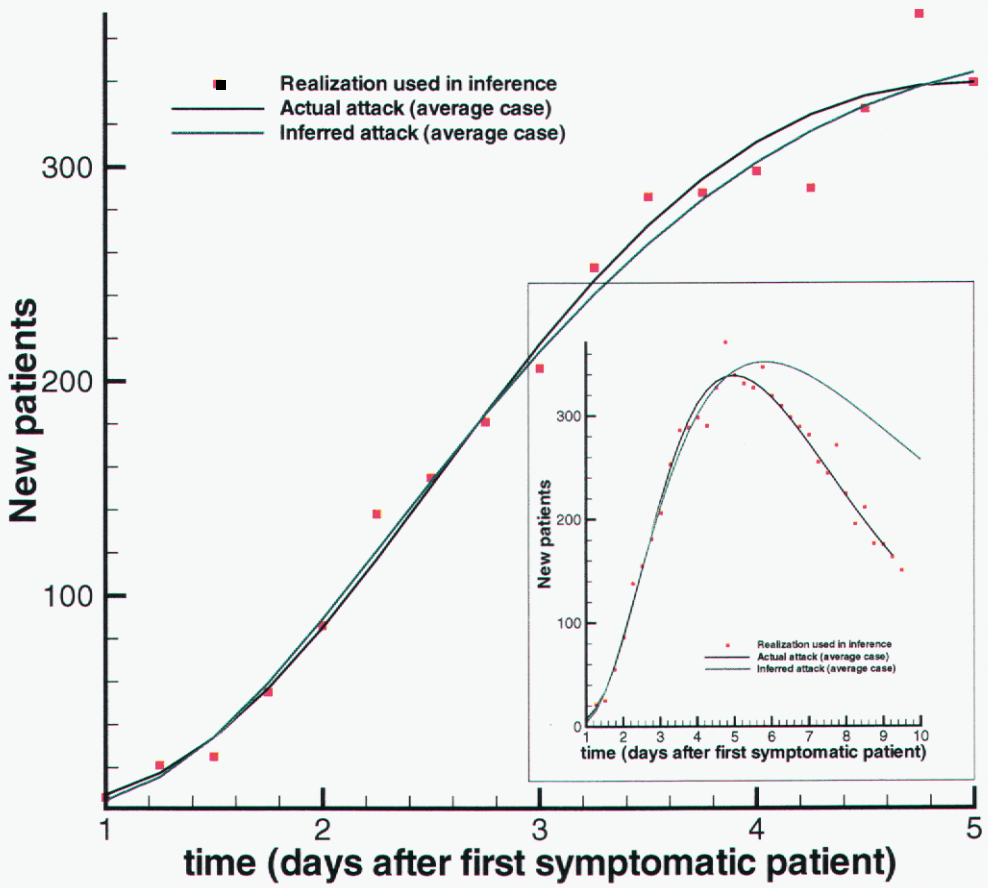

Figure 8. The time evolution of Case F plotted over 5 days. The observed data is plotted as symbols. We plot the number of patients showing symptoms, as collected over 6-hour intervals, on the vertical axis. The black line shows the evolution of the average case of the $N = 10^4, \tau = -1.5, D = 10^4$ spores attack while the blue one plots the evolution of a large, low-dose attack, inferred from the PDFs in Fig. 7 as $N = 16,00, \tau = -1, D = 10$ spores. We see that both the attacks behave very similarly over the first 5 days. Inset: We plot the evolution over 10 days. It is clear that the correct characterization ($N = 10^4, \tau = -1.5, D = 10^4$ spores) fits the observed data far better than the large, low-dose attack.

Table 1. The time-series as obtained from 6 different outbreaks, simulated with the parameters $\{N, \tau, D\}$ as mentioned at the bottom of the table. The table has been divided into 24-hour sections, where the components of the time-series are summed up to produce the low-resolution time-series (24-hour resolution) used to investigate the effect of temporal resolution.

time (days)	Case A	Case B	Case C	Case D	Case E	Case F
0.0	1	1	1	7	2	6
0.25	0	2	2	12	13	21
0.50	0	1	1	39	18	25
0.75	1	1	1	50	39	55
1.00	2	2	2	77	38	86
1.25	0	3	3	77	62	138
1.50	1	2	3	98	81	155
1.75	2	1	1	126	115	181
2.0	1	1	2	162	129	206
2.25	1	1	2	146	137	253
2.50	2	3	3	148	140	286
2.75	3	1	4	149	160	288
3.0	2	1	2	163	190	298
3.25	1	3	3	181	174	290
3.50	1	1	2	162	181	327
3.75	1	1	3	165	200	371
4.00	2	1	2	177	199	339
4.25	1	5	5	169	236	331
4.50	1	4	5	217	201	327
4.75	3	2	2	167	215	347
5.0	1	1	1	182	217	319
5.25	1	3	3	163	236	309
5.50	1	1	5	188	207	298
N	100	100	100	10,000	10,000	10,000
τ (days)	-0.75	-2.25	-2.25	-0.5	-1.0	-1.5
D (spores)	1	100	10,000	1	100	10,000

Table 2. The MAP estimate and the 90% confidence intervals (in parentheses) for N, τ and $log_{10}(D)$ as calculated from Figs. 2, 3, 4, 5, 6 and 7. Only the high-resolution time-series from Day 5 have been used to calculate the estimates and CI separately. The number in the curly brackets {} is the correct value.

Case	N	τ	$log_{10}(D)$
A	70, (40.35 – 122.9) {100}	-1.75, (-2.92 – -1.044) {−0.75}	0.0, (0.18 – 4.16) {0}
B	110, (65.8 – 148.5) {100}	-2.0, (-3.12 – -1.33) {−2.25}	0.00, (0.145 – 4.02) {2}
C	150, (91.1 – 196) {100}	-1.75, (-2.86 – -1.21) {−2.25}	0.0, (0.153 – 4.11) {4}
D	9800, (9460 – 10,370) {10,000}	-0.50, (-0.78 – -0.27) {−0.50}	0.00, (0.023 – 1.02) {0}
E	10,200, (8263 – 10,849) {10,000}	-1.0, (-1.567 – -0.60) {−1.00}	1.75, (0.921– 3.33) {2}
F	18,000, (10,770 – 18,930) {10,000}	-0.75, (-1.72 – -0.523) {−1.50}	0.75, (0.158 – 2.02) {4}

4 Results

In this section, we apply the inverse problem methodology developed in Sec. 3 to more realistic situations. We begin with tests in which we relax the assumption of people receiving a uniform dose; instead we allow N infected geographically-distributed people to receive dosages consistent with atmospheric dispersion, and focus on inferring a representative (average) dose $\overline{D}$. In Sec. 4.1, we use Wilkening's Model A2 to simulate BT attacks and then infer its characteristics after an incomplete observation. An average dosage is inferred for the infected population. Then, in Sec. 4.2 we use Wilkening's Model D (see Sec. 3.2) to simulate the outbreak while Model A2 is used for inference. Infected people are assumed to receive the same dose. This introduces a systematic error in the inference which is unrelated to the incomplete nature of the observations/time-series; that is, even in the limit of infi nite data, the errors will not tend to zero. In Sec. 4.3 we relax the assumption of equal dosage and repeat the analysis in Sec. 4.1 while using Model D for conducting the BT attack simulations. In Sec. 4.4 we apply this inference technique to the Sverdlovsk outbreak of 1979.

4.1 Variable Dosage Cases with Model A2

In this subsection we relax the assumption that all the infected people receive the same dosage, but conduct the simulated attacks with Wilkening's Model A2 as in Sec. 3.3. In order to obtain a realistic distribution of dosages in a geographically distributed population, we simulate an explosive point release of a quantity of spores at a height of 100 meters with a Gaussian plume model. This has the net effect of exposing different numbers of people to different dosage levels, as described in Appendix B. We see, from Fig. B.1, that given a quantity of spores, the number of people infected depends on the total population in the domain, the orientation of the plume and the population distribution in the domain. A release does not lead to many infected people if the high concentration isopleths of the plume miss the localized regions of high population densities.

In Table 3 and 4 we list the time-series obtained from Cases RA, RB, RC and RD. We also list the cumulative dosage levels $D_{0.01}, D_{0.25}, D_{0.50}, D_{0.75}$ and $D_{0.99}$. Thus 1 % of the population receives a dosage of $D_{0.01}$ spores or less, 75% of the population less than $D_{0.75}$ spores and $D_{0.99}$ is very close to the maximum dosage. This is a measure of the width of the spectrum of dosages. In Fig. 9, we plot the inverse of the cumulative dose distribution i.e. the abscissa is the fraction of the population which receives a dosage equal or less than the the mantissa. Inset, we plot a histogram of the number of infected people (who will ultimately exhibit symptoms) in each dosage bin. We note that while the spectrum may easily span two orders of magnitude, about 80% of the infected people lie within a range of doses that span a decade. We will essentially try to estimate a representative number for this span, by fi tting the inverse model developed in Sec. 3.1, which assumes a constant dosage. This is a source of model error, and adds to the uncertainty caused by incomplete observation and the stochastic nature of the data. This model error is not expected to diminish with the availability of data and thus the one of the aims of this investigation is to quantify it.

In Figs. 10, 11, 12 and 13 we plot the PDFs for $\{N, \tau, \log_{10}(\overline{D})\}$ as inferred from the time-series in

Table 3 and 4. In Table 5, we summarize the MAP estimates and the 90% CI for $\{N, \tau, \log_{10}(\overline{D})\}$, as obtained from the PDF shown in the fi gures. Only the PDFs from the time-series from Day 5 have been used to calculate the MAP estimates and CIs. We see many of the behaviors observed in Sec. 3.3. Dosage $\overline{D}$ is diffi cult to infer for all attacks (large confi dence intervals), while the time τ is easy for all cases. We can bound the size N of the attack quite accurately for all the cases, regardless of resolution of the time-series. We also plot the PDFs obtained from seven days of data for reference. In all cases the MAP estimates for N, obtained from 5 days of data are within 20% of the correct characterization. Further, the 90% CIs invariably bracket the true characterization. The MAP estimates obtained from the low and high-resolution time-series are similar. The 90% CIs obtained from the high-resolution time-series are generally slightly tighter, as a result of reduction in uncertainty due to stochastic noise - however, since this reduction is rather muted, the contribution of stochastic noise to uncertainty in the inference is rather modest, compared to model error. Thus while the model errors incurred by using an inverse problem derived under the assumption of constant dosage are not negligibly small, the current formulation may be suffi cient for placing bounds on the characterization of the BT attack.

Table 3. The time-series as obtained from four outbreaks simulated using Wilkening's Model A2, with the parameters $\{N, \tau, \overline{D}\}$ as mentioned at the bottom of the table. The table has been divided into 24-hour sections, where the components of the time-series are summed up to produce the low-resolution time-series (24-hour resolution) used to investigate the effect of temporal resolution. The infected patients receive a spectrum of dosages, denoted by the mean dosage $\overline{D}$ and $D_{0.01}, D_{0.25}, D_{0.50}, D_{0.75}$ and $D_{0.99}$. x % of the population receives a dosage of D_x or less. The time-series are continued past Day 5 to Day 8 in Table 4.

Time (days)	Case RA	Case RB	Case RC	Case RD
0.0	1	3	1	9
0.25	1	2	2	8
0.50	0	5	2	8
0.75	3	12	0	23
1.00	1	13	5	34
1.25	2	26	6	46
1.50	2	28	10	69
1.75	5	48	6	84
2.0	6	54	10	118
2.25	4	64	8	115
2.50	8	69	17	129
2.75	6	52	11	153
3.0	11	81	16	125
3.25	6	61	15	163
3.50	9	88	13	134
3.75	8	105	19	156
4.00	17	68	14	175
4.25	13	65	20	129
4.50	9	88	22	144
4.75	3	87	7	154
5.0	6	76	8	129
N	318	2989	454	4537
τ	-1.5	-1.5	-1.25	-1.25
$\overline{D}$	2912.8	2776.8	13,870.5	13,150.4
$D_{0.01}$	5.3×10^1	6.5×10^1	1.39×10^2	1.32×10^2
$D_{0.25}$	1.23×10^3	1.15×10^3	3.96×10^3	3.47×10^3
$D_{0.50}$	2.91×10^4	2.59×10^4	1.34×10^4	1.24×10^4
$D_{0.75}$	4.12×10^4	3.95×10^4	1.91×10^4	1.87×10^4
$D_{0.99}$	8.28×10^4	8.69×10^4	5.79×10^4	5.91×10^4

Table 4. Continuation of Table 3 beyond Day 5. The time-series as obtained from four outbreaks simulated using Wilkening's Model A2, with the parameters $\{N, \tau, \overline{D}\}$ as mentioned at the bottom of the table. The table has been divided into 24-hour sections, where the components of the time-series are summed up to produce the low-resolution time-series (24-hour resolution) used to investigate the effect of temporal resolution. The infected patients receive a spectrum of dosages, denoted by the mean dosage $\overline{D}$ and $D_{0.01}, D_{0.25}, D_{0.50}, D_{0.75}$ and $D_{0.99}$. x % of the population receives a dosage of D_x or less.

Time (days)	Case RA	Case RB	Case RC	Case RD
5.25	9	69	7	128
5.50	8	90	16	108
5.75	10	78	8	147
6.00	8	85	11	121
6.25	8	83	12	106
6.50	7	55	14	113
6.75	7	71	10	90
7.0	6	75	7	100
7.25	7	59	7	95
7.50	4	67	6	85
7.75	6	63	6	75
8.00	3	59	10	76
N	318	2989	454	4537
τ	-1.5	-1.5	-1.25	-1.25
$\overline{D}$	2912.8	2776.8	13,870.5	13,150.4
$D_{0.01}$	5.3×10^1	6.5×10^1	1.39×10^2	1.32×10^2
$D_{0.25}$	1.23×10^3	1.15×10^3	3.96×10^3	3.47×10^3
$D_{0.50}$	2.91×10^4	2.59×10^4	1.34×10^4	1.24×10^4
$D_{0.75}$	4.12×10^4	3.95×10^4	1.91×10^4	1.87×10^4
$D_{0.99}$	8.28×10^4	8.69×10^4	5.79×10^4	5.91×10^4

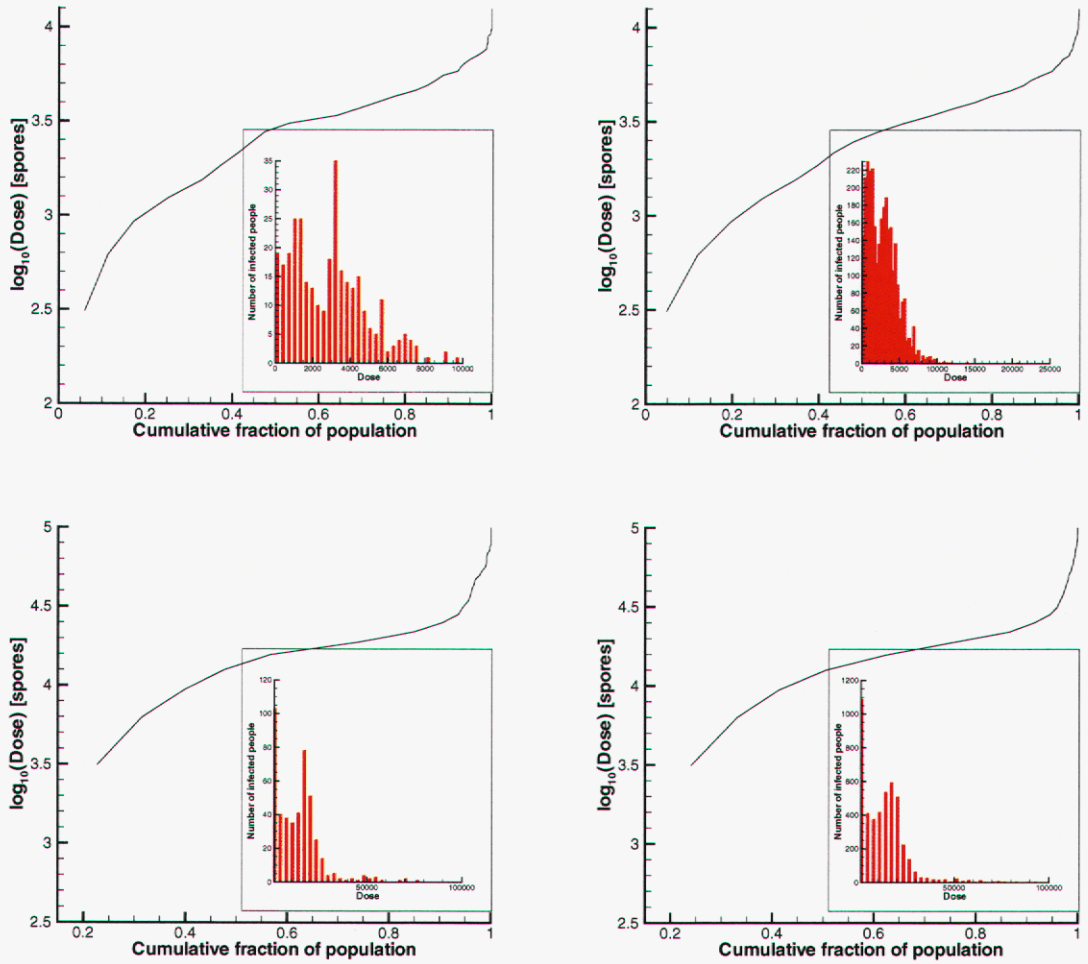

Figure 9. The inverse cumulative distribution of the dosages. The abscissa is the fraction of the population which receives a dosage equal or less than the ordinate. On top we plot the low-dose attacks, Case RA and RB, where the plume missed most of the population centers. Below, are the high-dose attacks, Case RC and RD. The attacks on the left are the smaller attacks, conducted with a lower population density. Inset: we plot the histogram containing the number of people in each dosage bin. We see that while the histogram has a long tail, the bulk of the population receives doses which span a spectrum a decade wide.

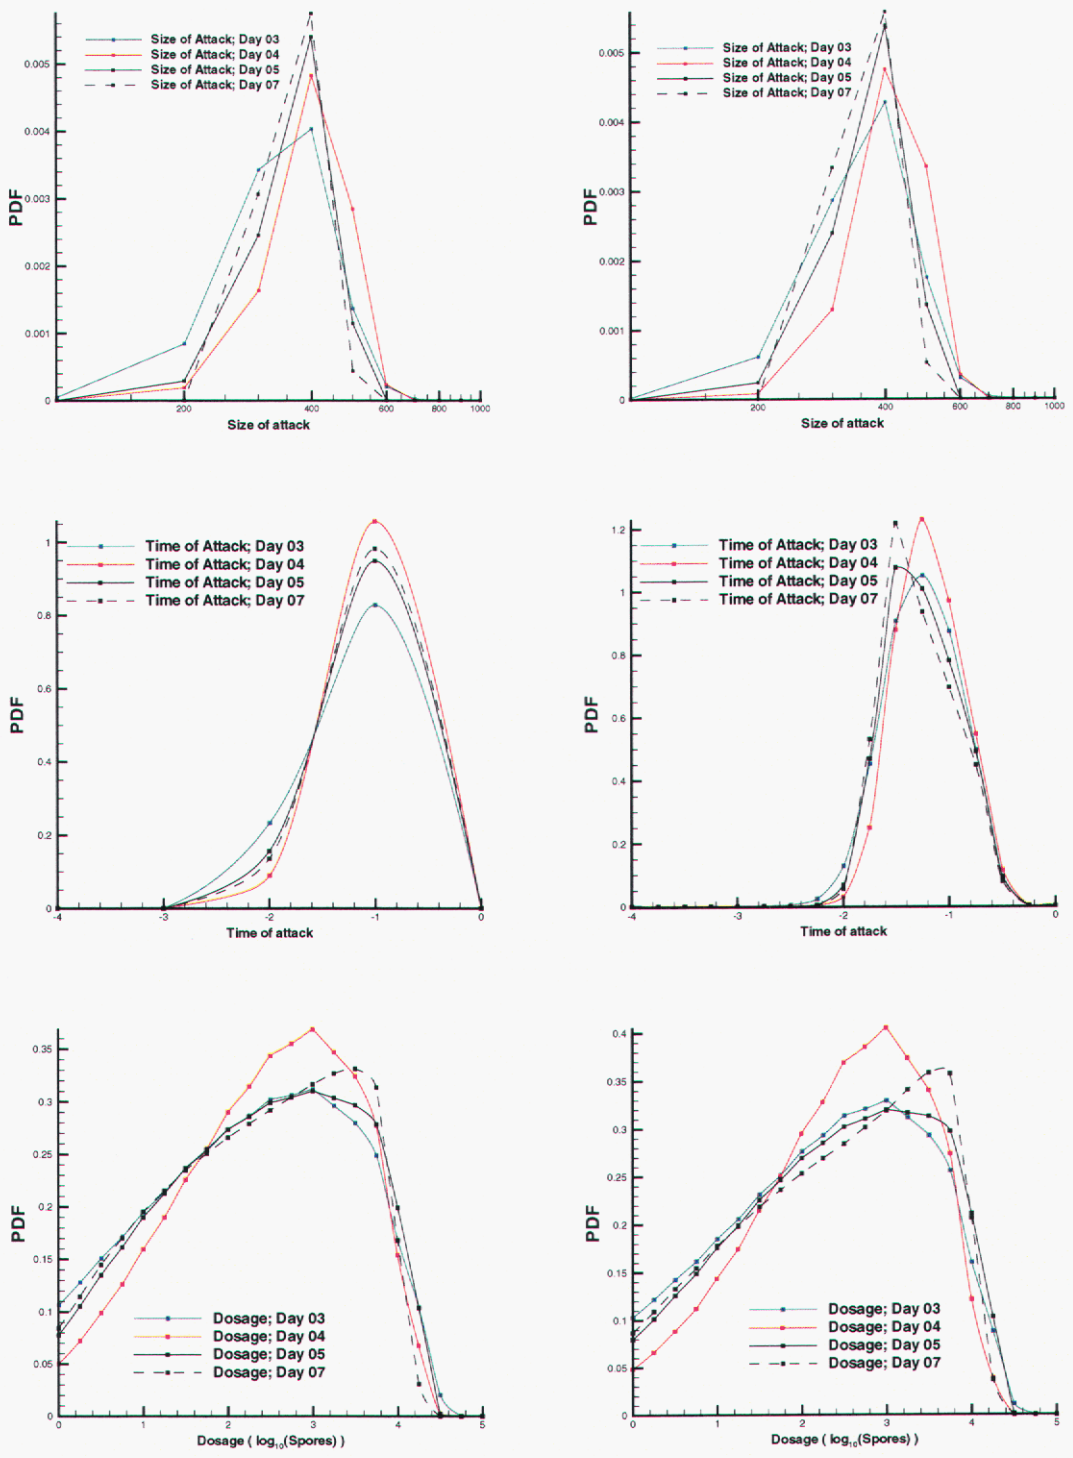

Figure 10. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case RA, as tabulated in Tables 3 and 4. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, \log_{10}(\overline{D})\}$ are $\{161, -0.75, 3.56\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively). PDFs develop from 8 days of data are also plotted for reference.

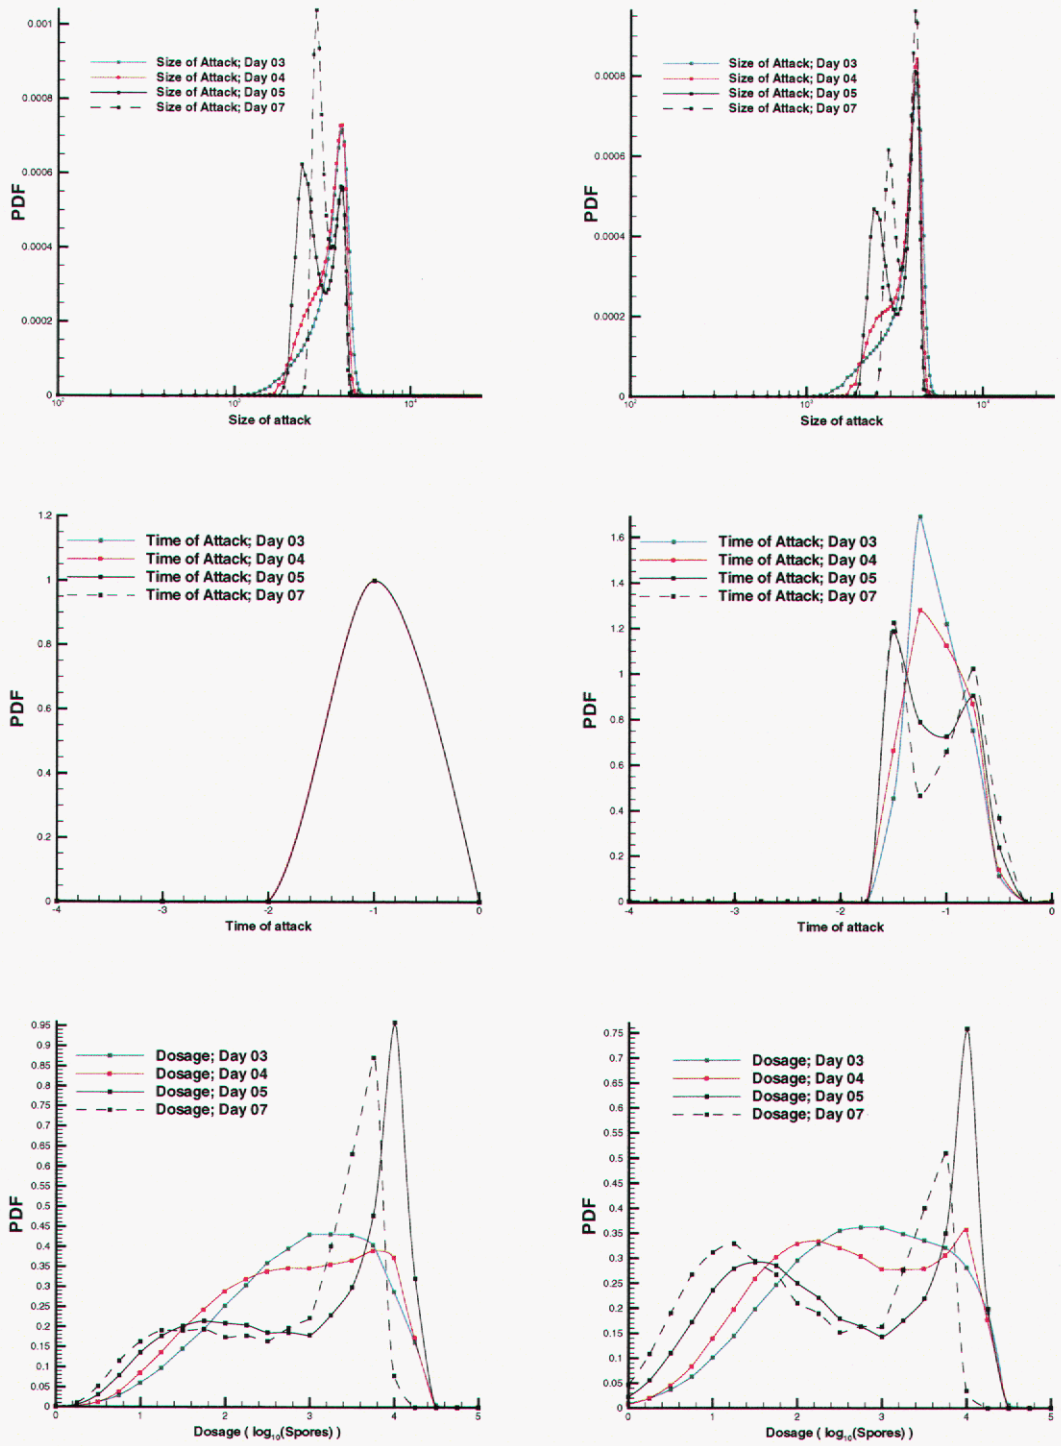

Figure 11. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case RB, as tabulated in Tables 3 and 4. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, \log_{10}(\overline{D})\}$ are $\{1453, -0.75, 3.563\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively). PDFs develop from 8 days of data are also plotted for reference.

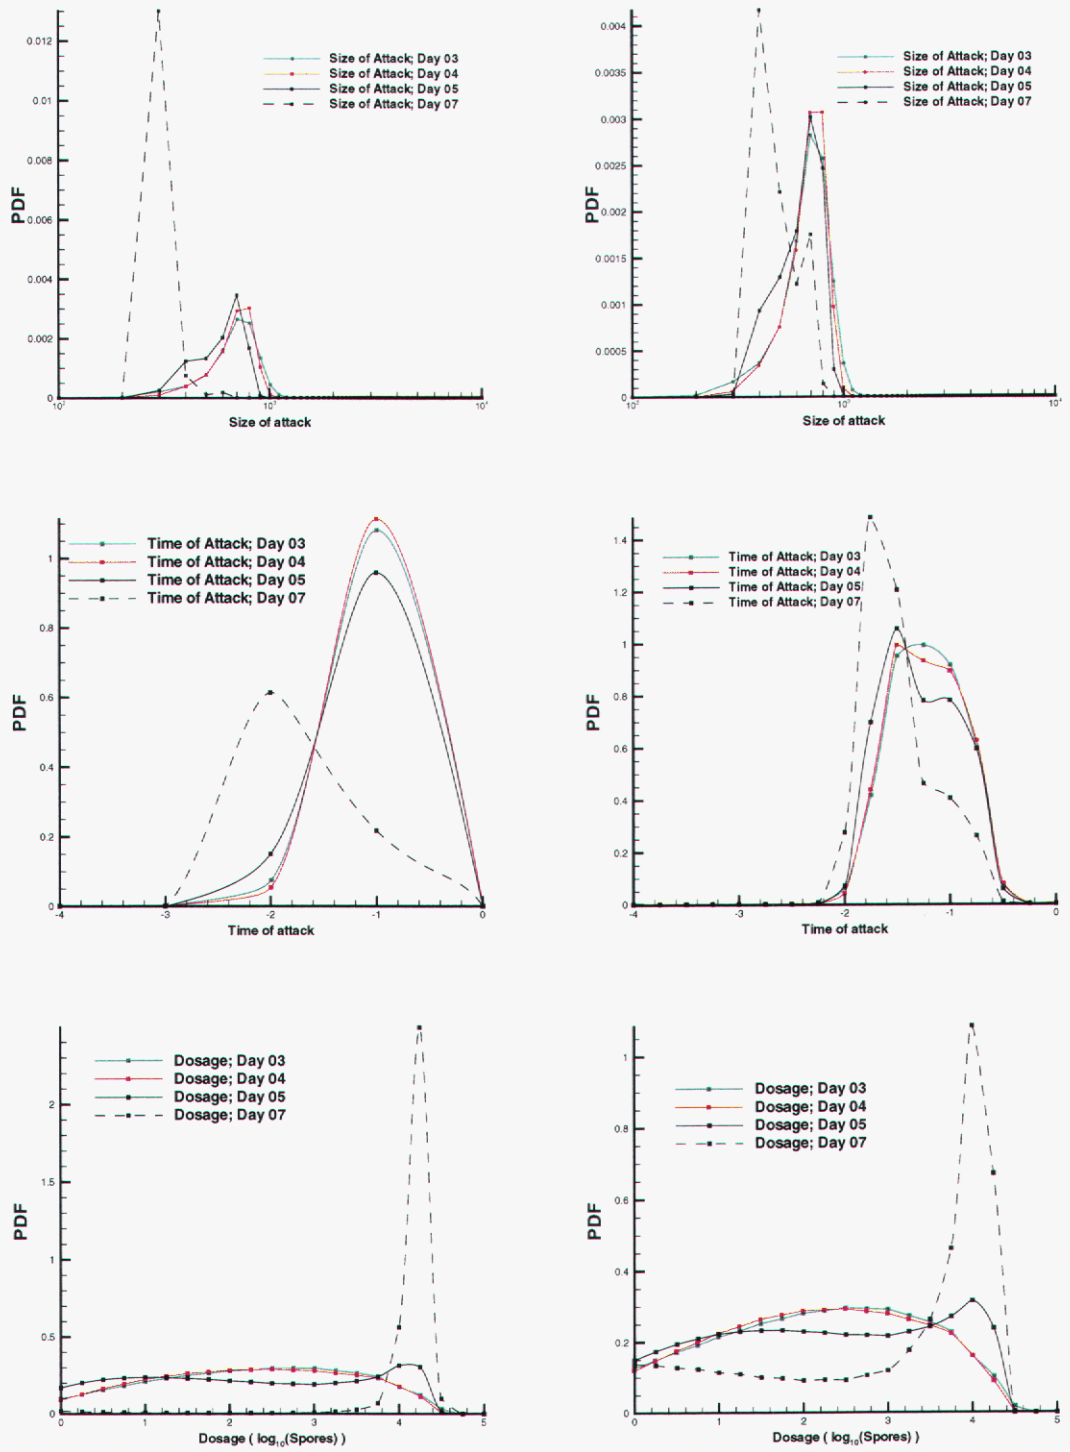

Figure 12. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case RC, as tabulated in Tables 3 and 4. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, \log_{10}(\overline{D})\}$ are $\{453, -0.75, 4.229\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively). PDFs develop from 8 days of data are also plotted for reference.

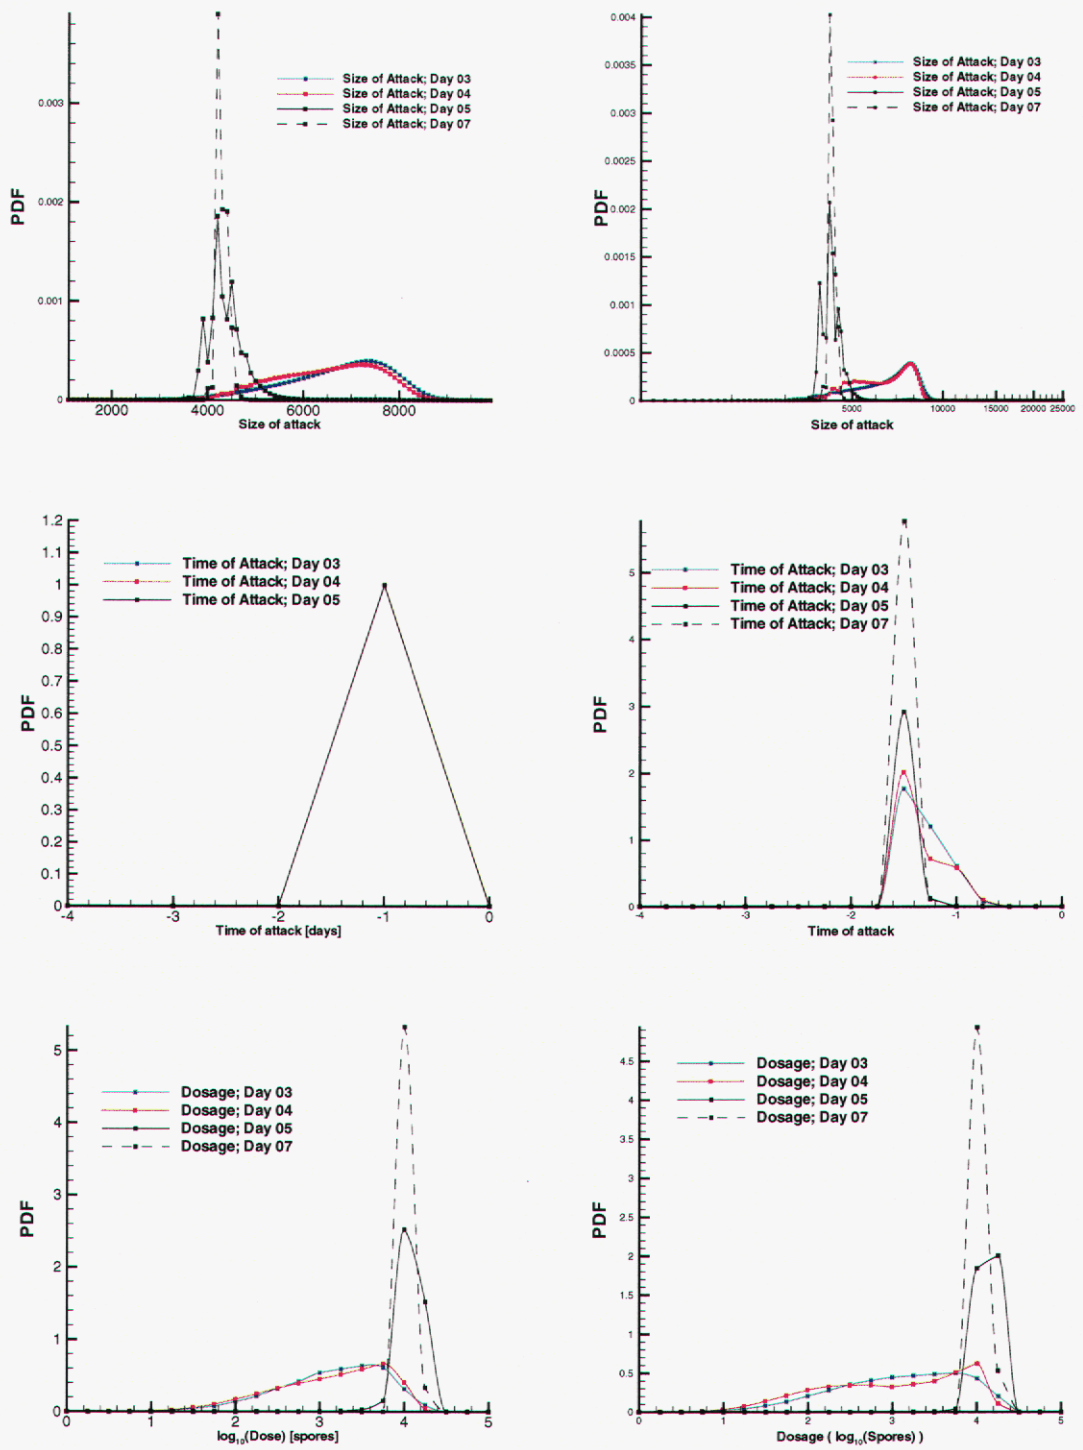

Figure 13. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case RD, as tabulated in Tables 3 and 4. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, \log_{10}(\overline{D})\}$ are $\{4453, -0.5, 4.22\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively). PDFs develop from 8 days of data are also plotted for reference.

Table 5. The MAP estimate and the 90% confidence intervals (in parentheses) for N, τ and $log_{10}(\overline{D})$ as calculated from Figs. 10, 11, 12 and 13. The time-series from Day 5 have been used to calculate the estimates and CI. The number in the curly brackets {} is the correct value.

Case	N	τ	$log_{10}(\overline{D})$
RA (6-hr resolution)	400, (228.8 – 581)	-1.5, (-1.93 – -0.714)	3.0, (0.49 – 4.12)
RA (24-hr resolution)	400, (225.8 – 580.2)	-1.0, (-1.86 – -0.105)	3.0, (0.48 – 4.12)
	{318}	{−1.5}	{3.46}
RB (6-hr resolution)	4100, (2288 – 4444)	-1.5, (-1.84 – -0.559)	4.00, (0.714 – 4.249)
RB (24-hr resolution)	2400, (2226 – 4304)	-1.0, (-2.80 – -0.133)	4.00, (1.098 – 4.439)
	{2989}	{−1.5}	{3.44}
RC (6-hr resolution)	700, (422 – 925)	-1.5, (-1.95 – -0.706)	4.0, (0.3 – 4.33)
RC (24-hr resolution)	700, (381.7 – 869.4)	-1.0, (-1.85 – -0.104)	4.0, (0.27 – 4.43)
	{454}	{−1.25}	{4.14}
RD (6-hr resolution)	4200, (3867 – 4895)	-1.5, (-1.97 – -1.28)	4.25, (3.8 – 4.7)
RD (24-hr resolution)	4200, (3894 – 5143)	-1.0, (-1.58 – -0.083)	4.00, (3.8 – 4.7)
	{4453}	{−1.25}	{4.11}

4.2 Uniform Dosage Cases with Model D

Given the poor characterization of inhalational anthrax (see Sec. 3.2), it is unlikely that a particular model of anthrax (for instance, the one used for inferring the attack characteristics) will be an accurate representation of its actual behavior in an infected population. Thus in the the interest of realism, one must consider the possibility of there existing a systematic difference (perhaps small) between the actual evolution of the attack and the model used to interpret the observed data. Since experimental results support Wilkening's Model D (Eq. 15) to the same extent as Model A2 (Eq. 16), in this subsection we will conduct BT attacks with Model D while inferring them using Wilkening's Model A2. In Table 6 we list the time-series obtained from 6 different simulations, conducted using Wilkening's Model D. As in Table 1, the time-series have a 6-hour resolution, with 24-hour demarcations to indicate the components that have to be summed to obtain time-series with 24-hour resolution. The values of N, τ and D are mentioned at the bottom of the table.

In Figs. 14, 15, 16, 17, 18 and 19 we plot the PDFs for $\{N, \tau, D\}$ for Cases G, H, I J, K and L per the time-series in Table 6. In Table 7, we summarize the MAP estimates and 90% confi dence intervals (CI) for N, τ and $\log_{10}(D)$ as obtained from the high-resolution time-series corresponding to 5 days of data. We see that many of the characteristics of the inferences in Sec. 3.3, in particular the non-unique/ambiguous characterizations, are reproduced here too. MAP estimates of N are within a factor of 2 of the correct answer, regardless of the size of the BT attack. MAP estimates of τ are consistently off by about a day. The dosage D is impossible to infer for small attacks. The net effect of higher-resolution time-series is to generate smoother PDFs. The inferences for Cases J and K show support for two different characterizations—a small, high-dose attack and a larger low-dose attack. In both cases the dual characterization disappears by the time the time-series incorporates Day 5 data. We also plot the PDFs for Case K obtained from 8 days of data using the 6-hour resolution time-series; they are close to the true characterization, except in case of dosage.

The systematic difference between Wilkening's Model D (used for simulated attacks) and Model A2 (used for inference) explain some of the mis-characterizations. In Fig. 1 we see that Model D predicts longer median incubation periods. Thus, given identical initial conditions, Model D predicts a slower increase in the epidemic curve vis-a-vis Model A2. Consequently, when data generated by Model D is interpreted using Model A2 (during the inference process), the slower increase is accommodated by reducing the number of index cases; Model A2 will consistently underestimate the size of the outbreak. The dosage and time of attack will also be "adjusted" by Model A2 to fi t the data. This is quite clear in the large attacks, Cases K and L. The low-resolution time-series supports two characterizations for Case K—a small, high-dose attack or a larger low-dose attack, both about 1 day before the fi rst exhibition of symptoms. The high resolution time-series also supports a dual characterization—a larger attack, 0.75 days before the fi rst exhibition of symptoms or a smaller attack, 1.5 days prior to symptoms, both at a very high dose. However, when a time-series 8 days long is used, we recover the correct characterization. In Fig. 20, we plot the evolution of the outbreak as predicted by Model A2 for various values of $\{N, \tau, D\}$. With $\{N, \tau, D\} = \{10^4, -0.75, 10^2\}$, the correct values, we obtain the solid black line. It is signifi cantly different from the time-series (plotted with symbols) which was generated using Model D, especially if the long term evolution of the outbreak is considered. Thus it is unlikely that Model A2 will correctly characterize the attack. With $\{N, \tau, D\} = \{5000, -1.6, 15,848\}$, the

approximate MAP estimate after 5 days of data, we obtain the blue line. The fi t to the observed time-series is moderate, though once a longer time-series (8 days long) is considered, it becomes untenable. Using $\{N, \tau, D\} = \{10, 500, -0.75, 1\}$, the MAP estimate after 8 days, we obtain the dashed line which shows a good fi t only if the latter part of the outbreak in considered. While the MAP estimates for N and τ are very close to the true value, D is not and the consequently the closeness of the MAP estimates of N and τ to the true values must be considered fortuitous. This discrepancy is due to the difference between Models A2 and D and cannot be overcome by more data (i.e., a longer time-series). We believe that a similar argument explains the discrepancy between the characterization and simulation as seen in Fig. 19, where the inferred attack size is a factor of 2 smaller than the true one and the time of attack is a day too late.

Model A2 and Model D are perhaps the two most accurate models of the evolution of inhalational anthrax. Even under model uncertainty, time-series of 3–5 day observation periods generate inferences which may be of use in consequence management. Regardless of the size of the attack, the MAP estimate for N is within a factor of 2 of the true fi gure. τ is usually inferred a day too late. This particular behavior is true only when Model D is used for the simulation and Model A2 is used for inference. In case the models were switched i.e., the BT attacks were simulated by Model A2 and inferred by Model D, the inferences of the attack's size would be overestimates (though still within a factor of 2) and τ would be a day too early. Dosages pose the stiffest challenge and are impossible to infer for small attacks ($N = 100$). These general observations hold true if all the infected people receive the same dosage D (as might be expected for an aerosol release in a closed environment). In the next subsection, we consider variable dosages.

Table 6. The time-series as obtained from six outbreaks simulated using Wilkening's Model D, with the parameters $\{N, \tau, D\}$ as mentioned at the bottom of the table. The table has been divided into 24-hour sections, where the components of the time-series are summed up to produce the low-resolution time-series (24-hour resolution) used to investigate the effect of temporal resolution. All infected patients receive the same dosage of D spores. The time-series for Case K, beyond Day 5, is $\{173, 187, 170, 181, 170, 166, 164, 152, 158, 150, 129, 139\}$. This is used later in the text.

time (days)	Case G	Case H	Case I	Case J	Case K	Case L
0.0	1	1	1	1	6	2
0.25	0	0	1	2	10	22
0.50	1	0	3	20	42	66
0.75	2	1	1	34	53	128
1.00	1	3	3	65	88	153
1.25	1	0	1	63	100	211
1.50	2	1	1	84	116	224
1.75	3	2	4	114	151	237
2.0	2	1	3	148	204	266
2.25	1	1	3	143	183	292
2.50	1	2	4	145	181	274
2.75	1	3	3	151	182	273
3.0	3	2	3	164	189	263
3.25	1	1	1	185	207	251
3.50	1	1	1	165	177	258
3.75	3	1	1	167	181	235
4.00	1	2	1	182	192	214
4.25	1	1	2	181	185	210
4.50	1	1	1	217	215	205
4.75	3	3	3	168	174	222
5.0	3	1	1	187	187	200
5.25	1	1	2	168	173	196
5.50	1	1	0	191	187	196
N	100	100	100	10,000	10,000	10,000
τ	-1.75	-0.75	-0.75	-0.75	-0.75	-0.75
D	1	100	10,000	1	100	10,000

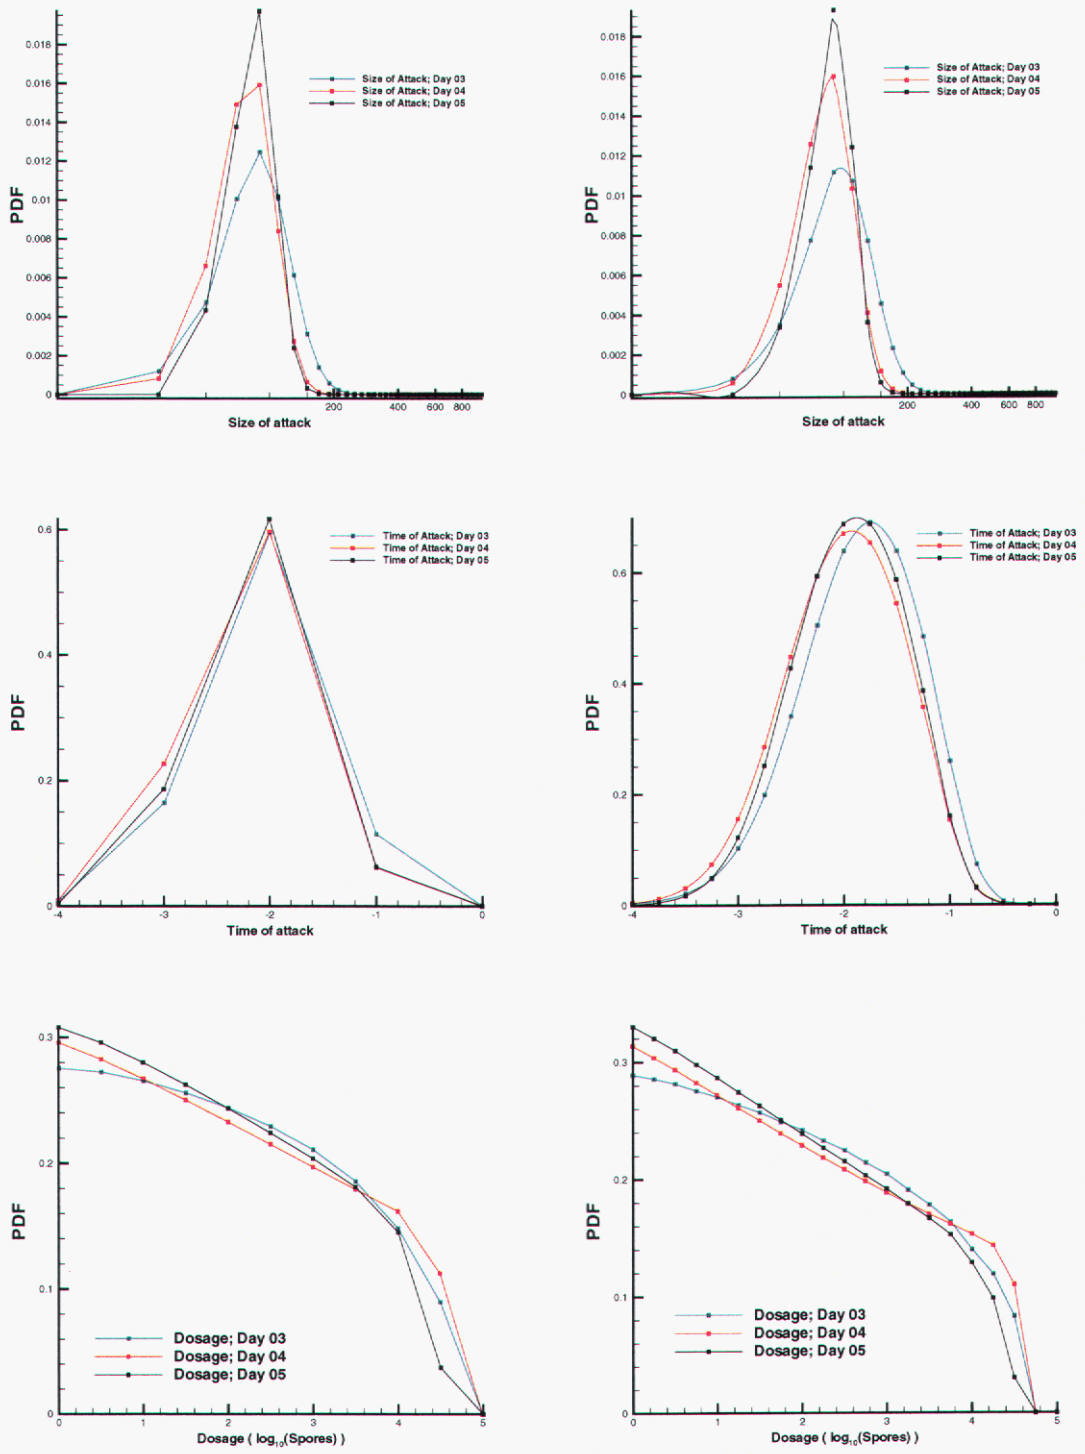

Figure 14. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case G, as tabulated in Table 6. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, D\}$ are $\{10^2, -1.75, 10^0\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

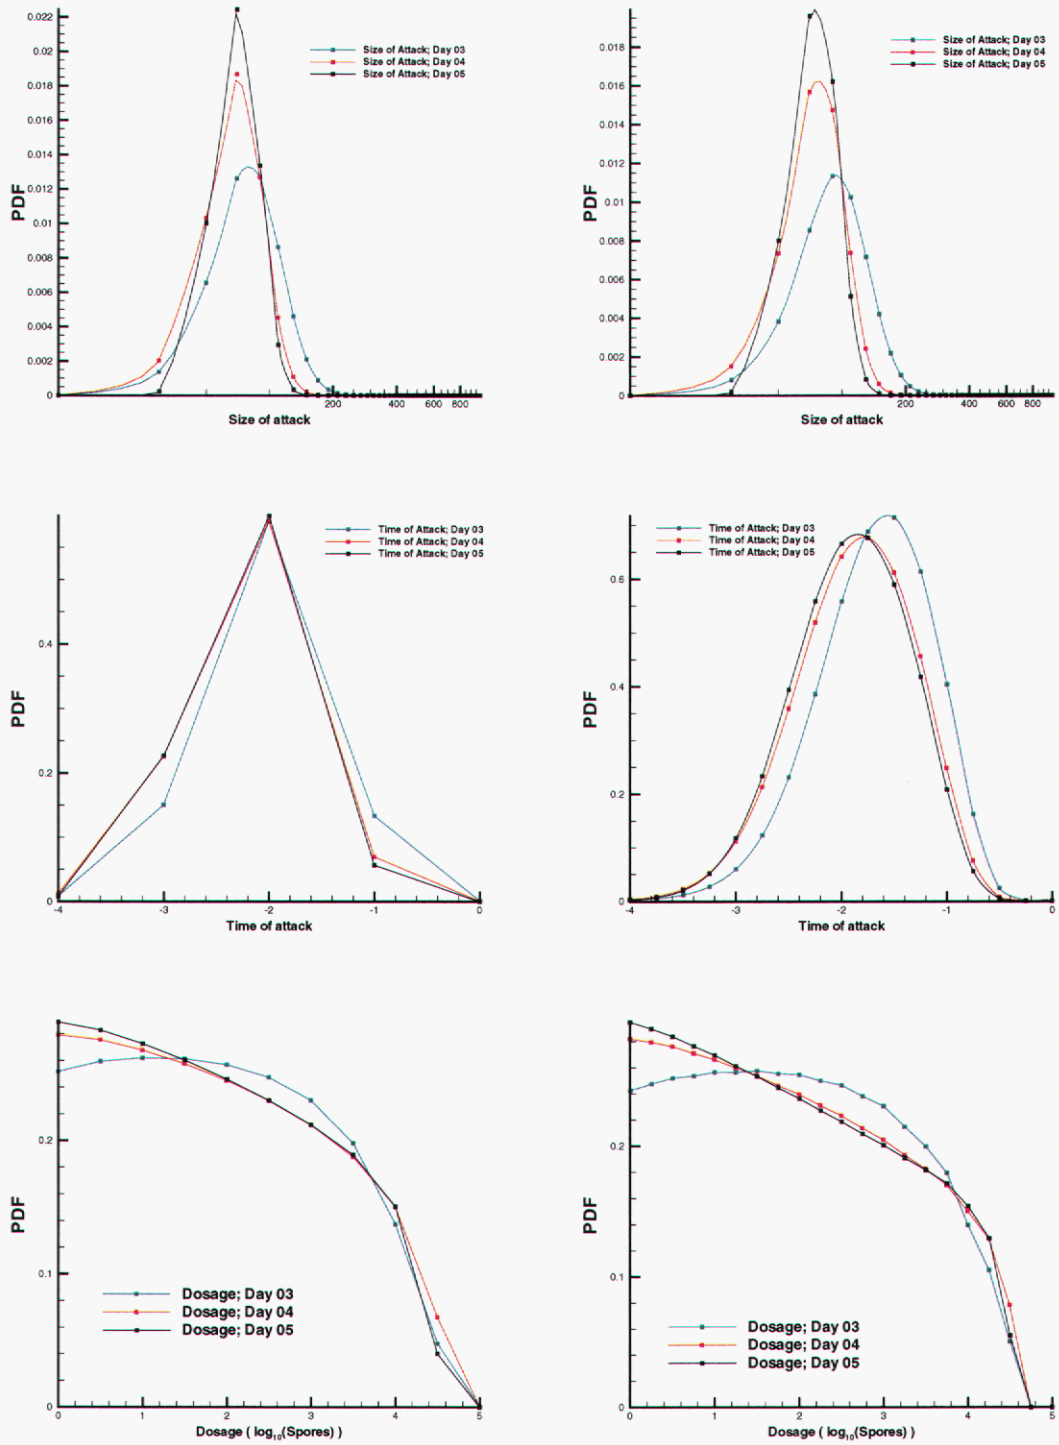

Figure 15. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case H, as tabulated in Table 6. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, D\}$ are $\{10^2, -0.75, 10^2\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

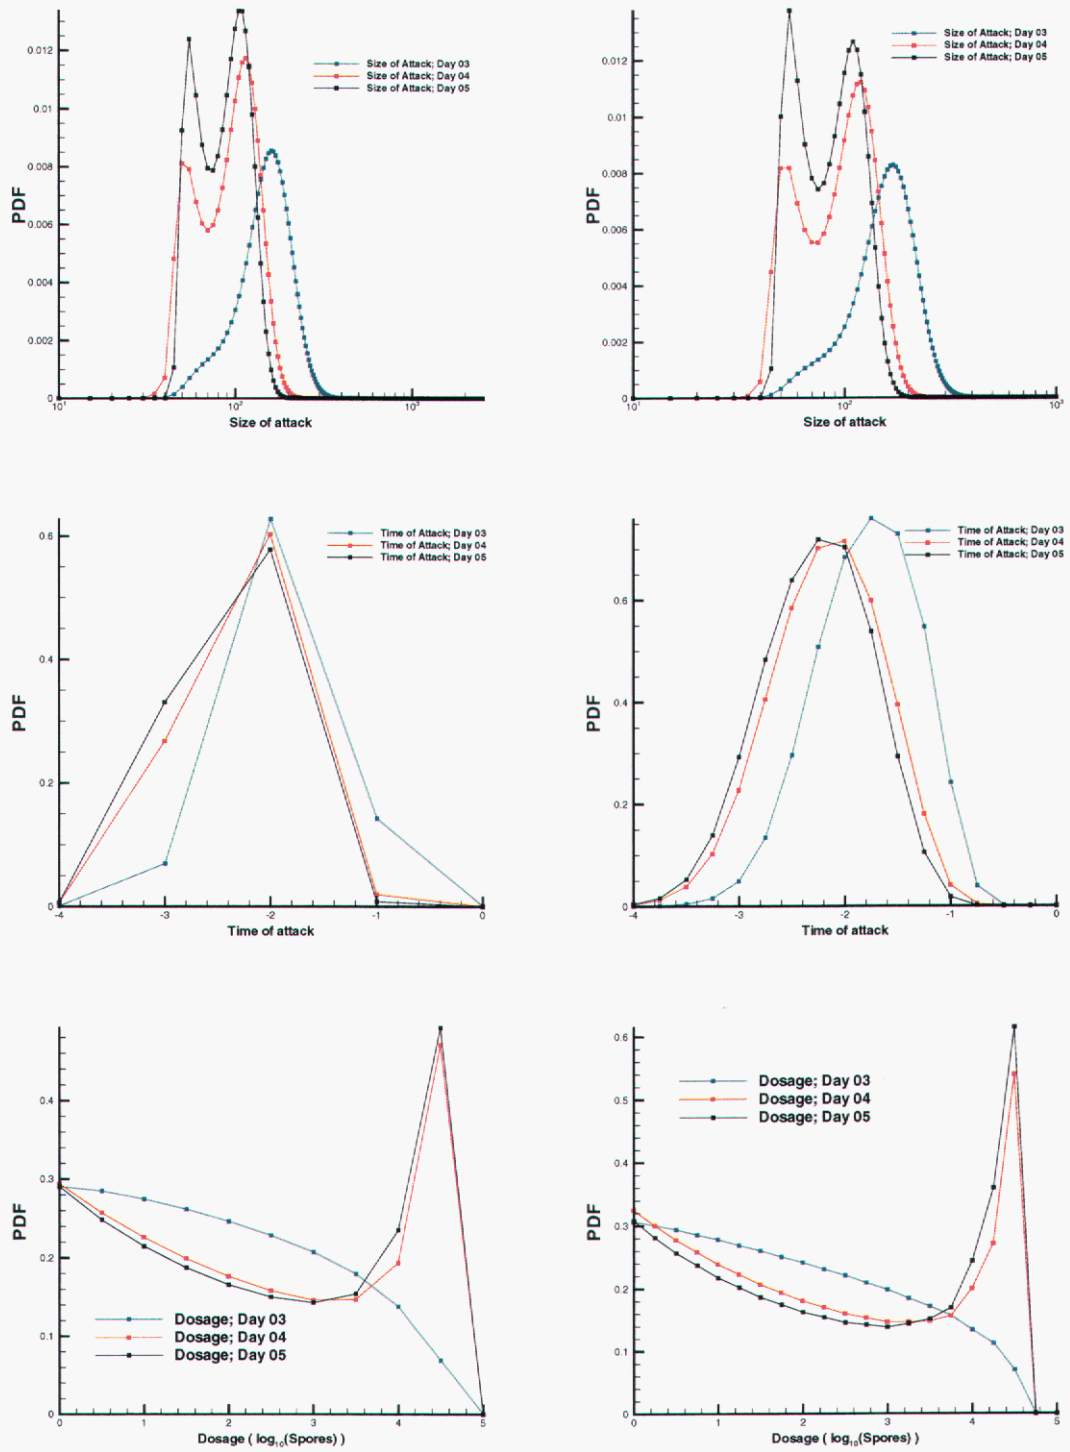

Figure 16. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case I, as tabulated in Table 6. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, D\}$ are $\{10^2, -0.75, 10^4\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

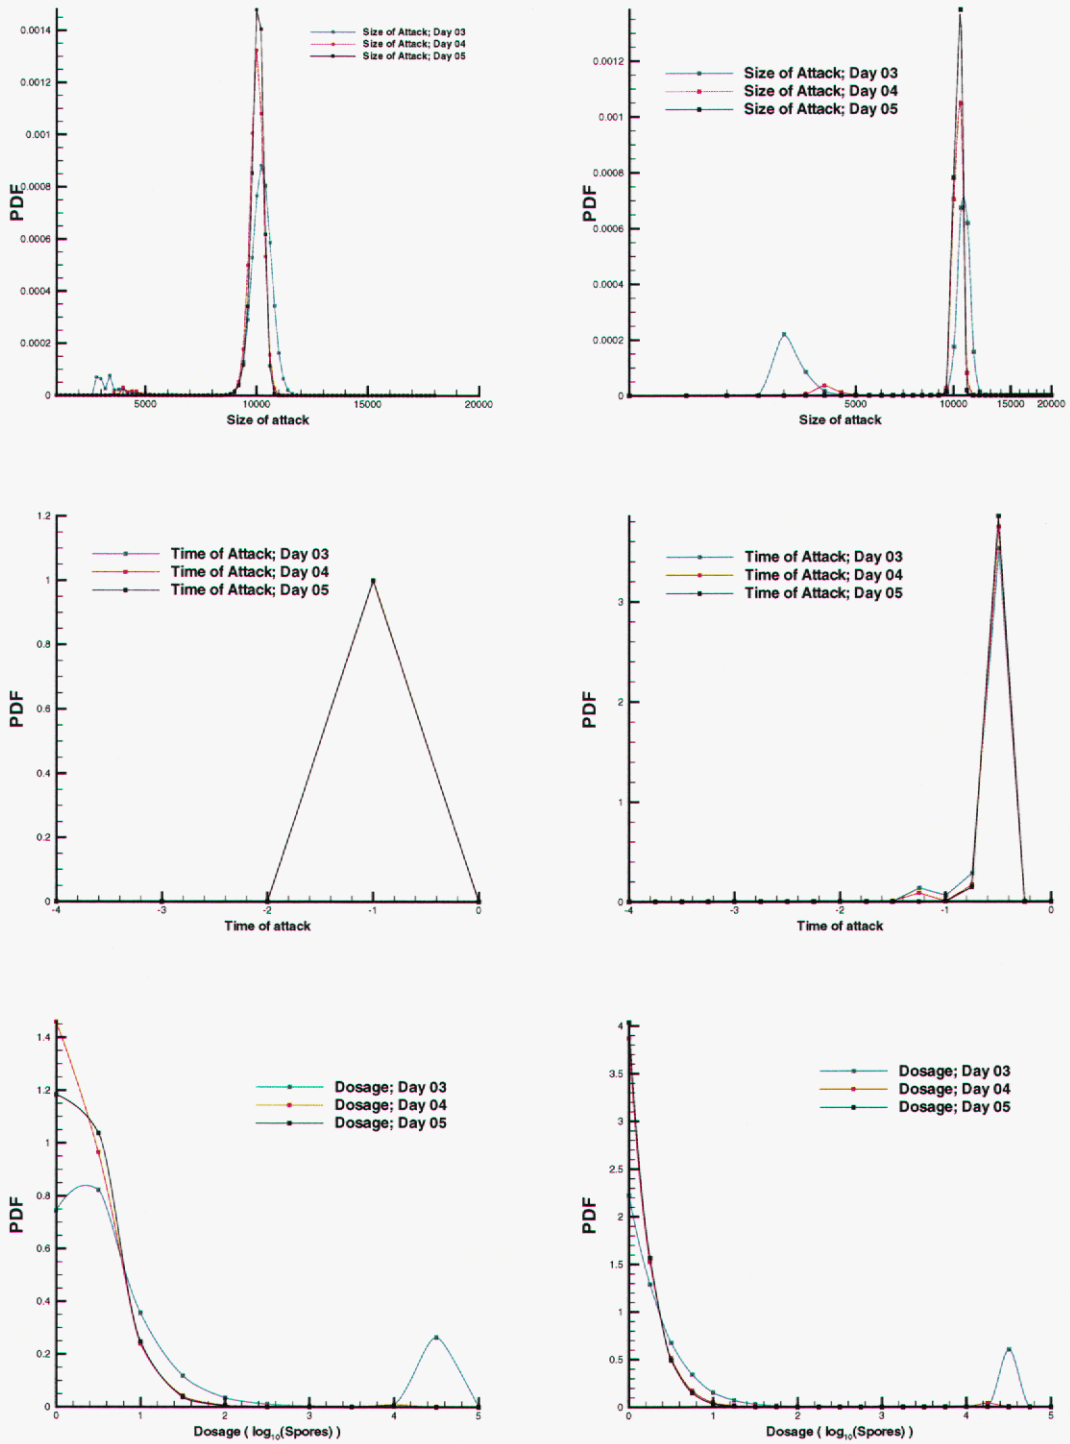

Figure 17. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case J, as tabulated in Table 6. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, D\}$ are $\{10^4, -0.75, 10^0\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

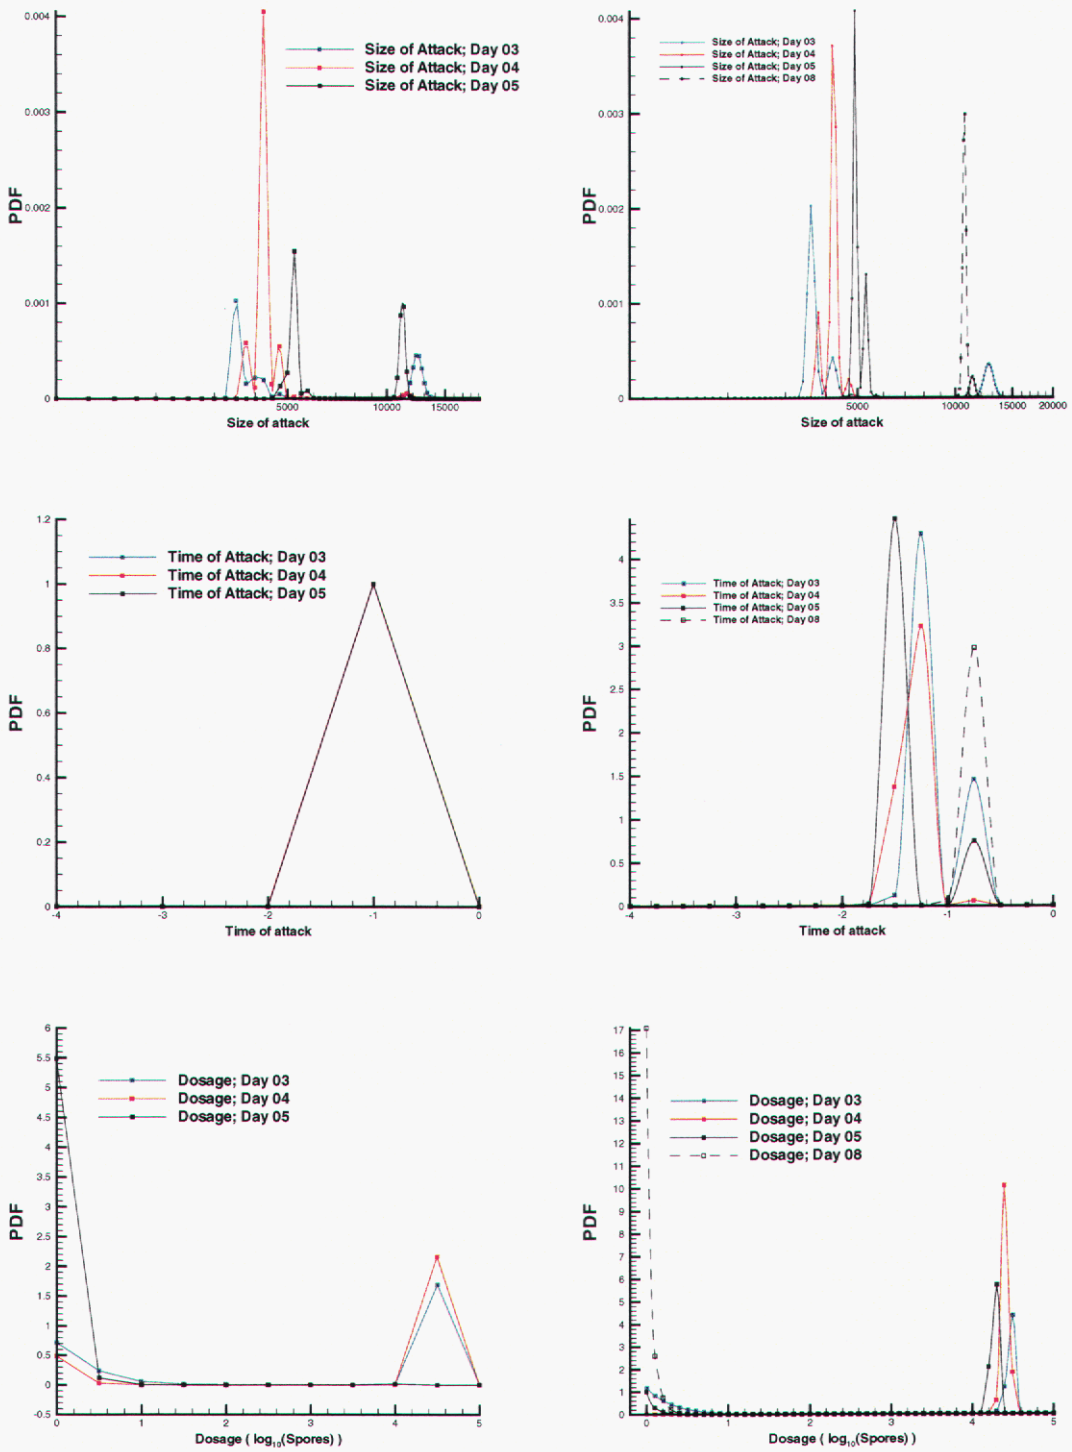

Figure 18. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case K, as tabulated in Table 6. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, D\}$ are $\{10^4, -0.75, 10^2\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

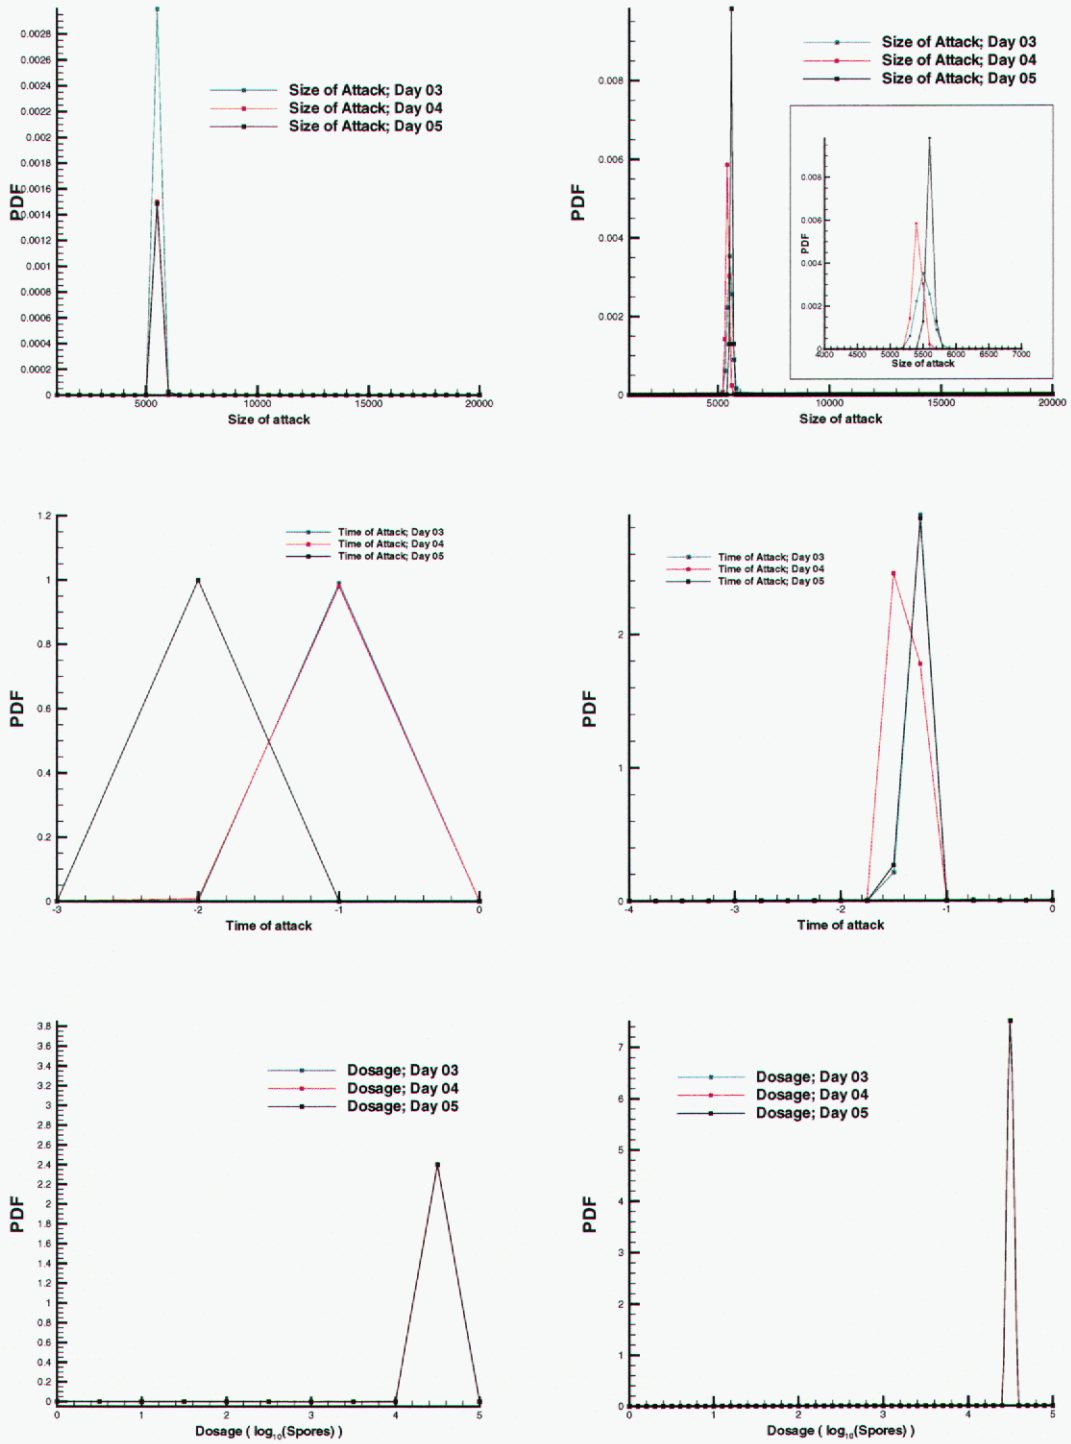

Figure 19. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case L, as tabulated in Table 6. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, D\}$ are $\{10^4, -0.75, 10^4\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

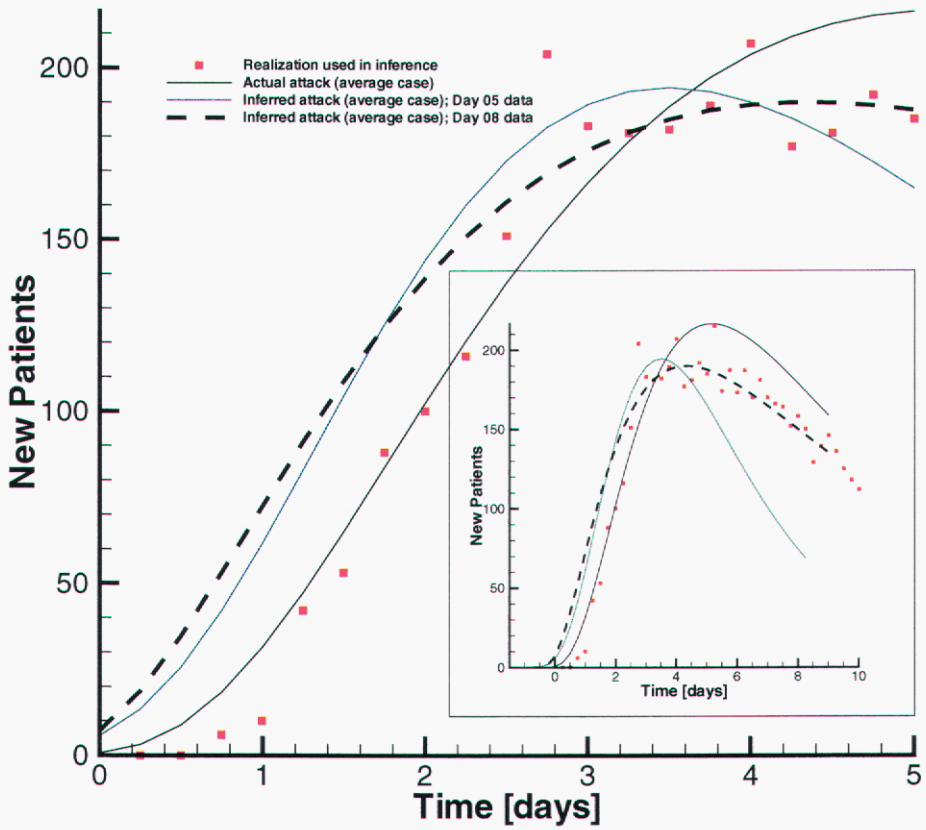

Figure 20. Evolution of Case K as predicted by Model A2 (used in inference) with various initial values of $\{N, \tau, D\}$. The symbols are the data that constitute the time-series Case K in Table 6. Using $\{N, \tau, D\} = \{10^4, -0.75, 10^2\}$, we obtain the solid black line. Using $\{N, \tau, D\} = \{5000, -1.6, 15, 848\}$, the MAP estimate after 5 days of data, we obtain the blue line. Inset: Using $\{N, \tau, D\} = \{10, 500, -0.75, 1\}$, the MAP estimate after 8 days, we obtain the dashed line which shows a good fit only if the latter part of the outbreak in considered.

Table 7. The MAP estimate and the 90% confidence intervals (in parentheses) for N, τ and $log_{10}(D)$ as calculated from Figs. 14, 15, 16, 17, 18 and 19. Only the high-resolution time-series from Day 5 have been used to calculate the estimates and CI separately. The number in the curly brackets {} is the correct value.

Case	N	τ	$log_{10}(D)$
G	90, (53.75 – 132.5) {100}	-1.75, (-2.98 – -1.15) {−1.75}	0.0, (0.15 – 4.09) {0}
H	70, (41.75 – 123.2) {100}	-1.75, (-2.98 – -1.09) {−0.75}	0.00, (0.17 – 4.195) {2}
I	55, (53.46 – 147.7) {100}	-2.25, (-3.25 – -1.51) {−0.75}	4.5, (0.169 – 4.61) {4}
J	10,500, (9678 – 10,960) {10,000}	-0.50, (-0.75 – -0.27) {−0.75}	0.00, (0.018 – 0.49) {0}
K	4,900, (4807 – 5478) {10,000}	-1.5, (-1.73 – -0.70) {−0.75}	4.3, (0.077– 4.483) {2}
L	5,600, (5,502 – 5,713) {10,000}	-1.25, (-1.725 – -1.05) {−0.75}	4.5, (4.42 – 4.687) {4}

4.3 Variable Dosage Cases with Model D

In this subsection we relax the assumption that all the infected persons receive the same dosage but conduct the BT attacks in the same manner as in Sec. 4.2 i.e., with Wilkening's Model D with the incubation period given by Eq. 15. A realistic distribution of spores among an infected population is obtained in a manner identical to the one described in Sec. 4.1. Inferences are done with Wilkening's Model A2 i.e., with an incubation period given by Eq. 16..

In Table 8, we list the time-series obtained from Cases M, N, O and P. We also list the cumulative dosage levels $D_{0.01}, D_{0.25}, D_{0.50}, D_{0.75}$ and $D_{0.99}$. In Figs. 22, 23, 24 and 25 we plot the PDFs for $\{N, \tau, \log_{10}(\overline{D})\}$ as inferred from the time-series in Table 8. In Table 9, we summarize the MAP estimates and the 90% CI for $\{N, \tau, \log_{10}(\overline{D})\}$, as obtained from the PDF shown in the figures. Only the PDFs from the time-series from Day 5 have been used. We see much the same behavior as in Sec. 4.2: the MAP estimate for N is within a factor of two of the true result and τ is estimated a day too late. Further, N is underestimated, per the explanation in Sec. 4.2. The effect of the higher-resolution time-series is again rather muted; the systematic model errors and the approximation caused by fitting a constant dosage model to variable dosage data overwhelm the reduction in uncertainty (due to noise) that the higher-resolution data achieves. As observed before, dosage is difficult to estimate for the small lower-dose attack. In Cases N and O, we also see dual characterizations of the attack.

Table 8. The time-series as obtained from four outbreaks simulated using Wilkening's Model D, with the parameters $\{N, \tau, \overline{D}\}$ as mentioned at the bottom of the table. The table has been divided into 24-hour sections, where the components of the time-series are summed up to produce the low-resolution time-series (24-hour resolution) used to investigate the effect of temporal resolution. The infected patients receive a spectrum of dosages, denoted by the mean dosage $\overline{D}$ and $D_{0.01}, D_{0.25}, D_{0.50}, D_{0.75}$ and $D_{0.99}$. x % of the population receives a dosage of D_x or less.

Time (days)	Case M	Case N	Case O	Case P
0.0	1	1	1	3
0.25	1	8	5	14
0.50	0	20	6	36
0.75	1	16	13	81
1.00	3	9	12	77
1.25	2	18	14	94
1.50	2	28	13	123
1.75	1	30	13	132
2.0	2	37	17	129
2.25	2	27	15	159
2.50	4	41	17	126
2.75	2	39	14	149
3.0	3	34	9	131
3.25	2	32	14	129
3.50	3	25	16	136
3.75	6	33	12	100
4.00	4	27	14	125
4.25	5	33	11	104
4.50	3	33	6	110
4.75	5	23	11	106
5.0	6	23	15	90
5.25	4	25	7	80
5.50	0	25	6	101
N	161	1453	453	4453
τ	-0.75	-0.75	-0.75	-0.5
$\overline{D}$	3603.5	3660.77	16,941	16,532
$D_{0.01}$	3.41×10^2	2.65×10^2	3.1×10^2	3.0×10^2
$D_{0.25}$	1.99×10^3	2.13×10^3	9.8×10^3	9.45×10^3
$D_{0.50}$	3.34×10^3	3.5×10^3	1.65×10^4	1.57×10^4
$D_{0.75}$	4.79×10^3	4.8×10^3	2.09×10^4	2.07×10^4
$D_{0.99}$	9.17×10^3	9.51×10^3	6.74×10^4	6.51×10^4

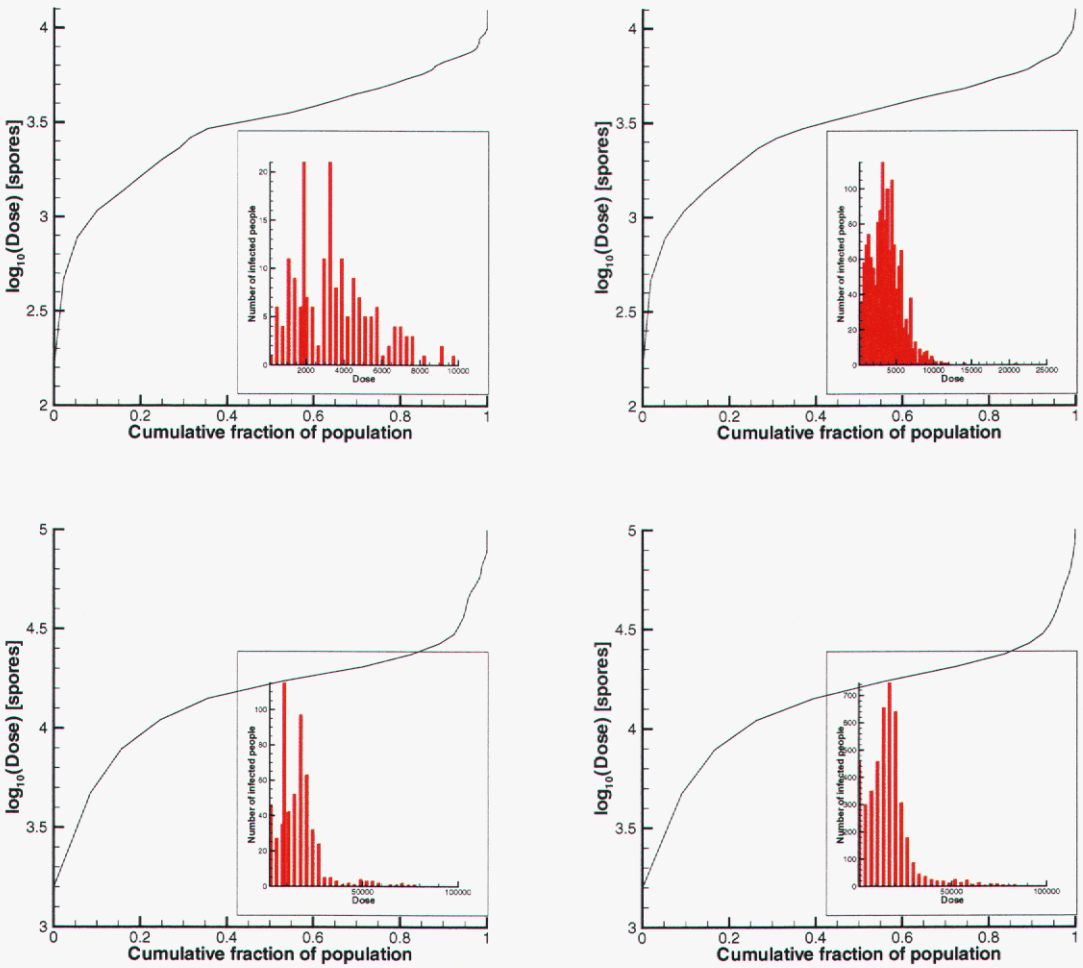

Figure 21. The inverse cumulative distribution of the dosages. The abscissa is the fraction of the population which receives a dosage equal or less than the the ordinate. On top we plot the low-dose attacks, Case M and N, where the plume missed most of the population centers. Below, are the high-dose attacks, Case O and P. The attacks on the left are the smaller attacks, conducted with a lower population density. Inset: we plot the histogram containing the number of people in each dosage bin. We see that while the histogram has a long tail, the bulk of the population receives doses which span a spectrum a decade wide.

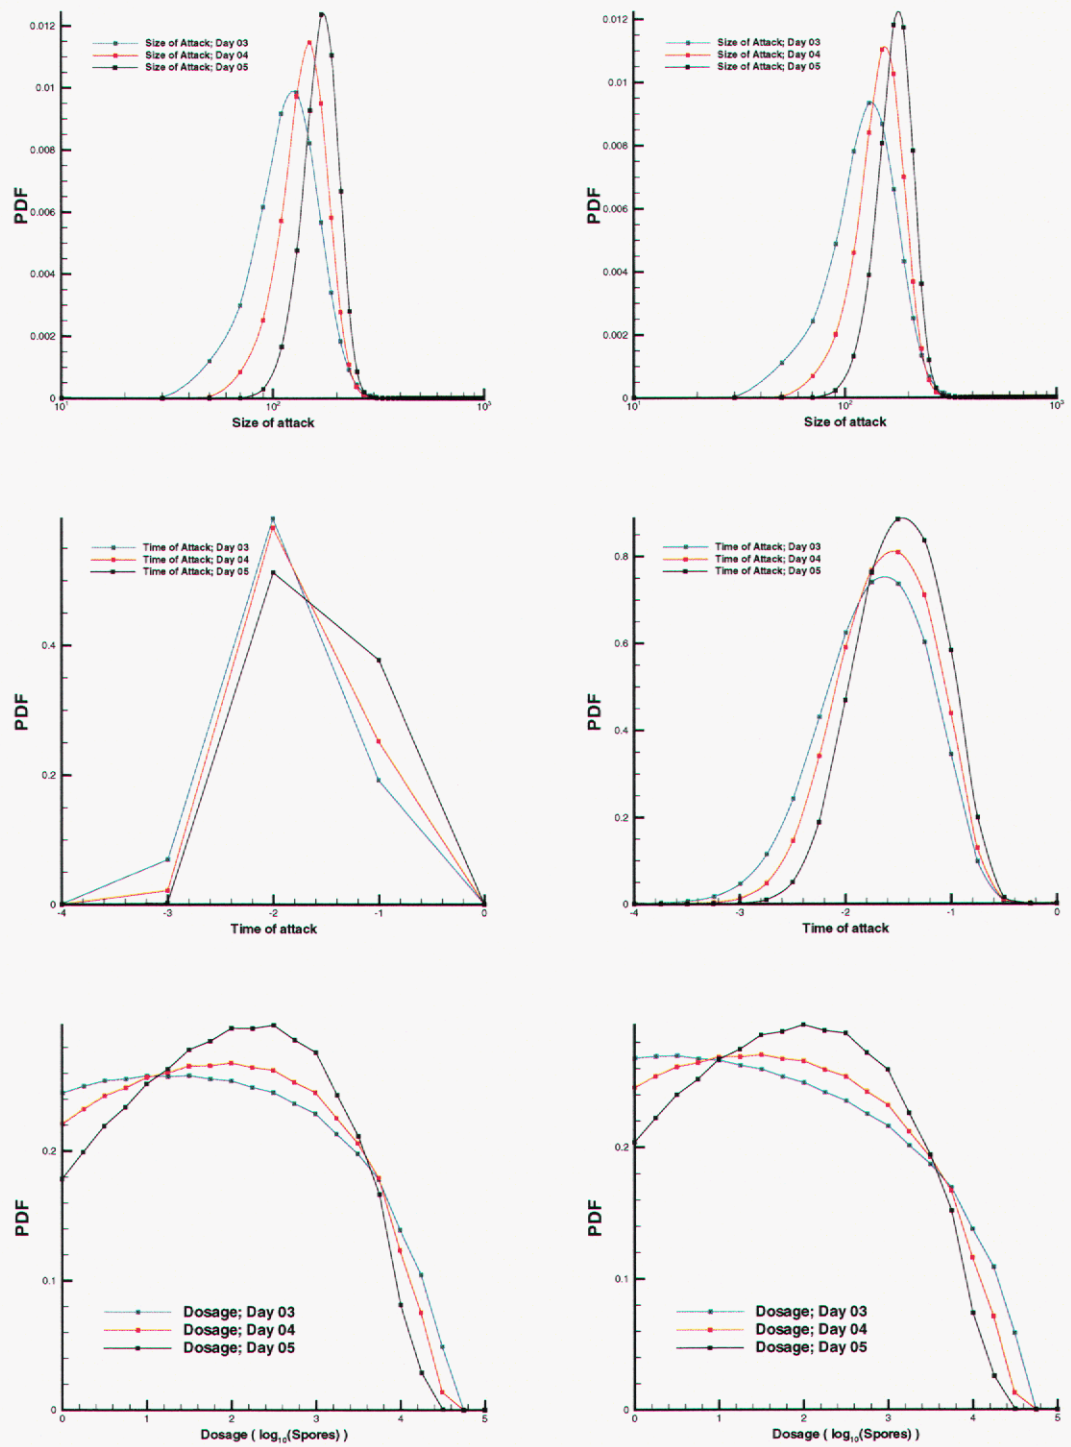

Figure 22. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case M, as tabulated in Table 8. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, \log_{10}(\overline{D})\}$ are $\{161, -0.75, 3.56\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

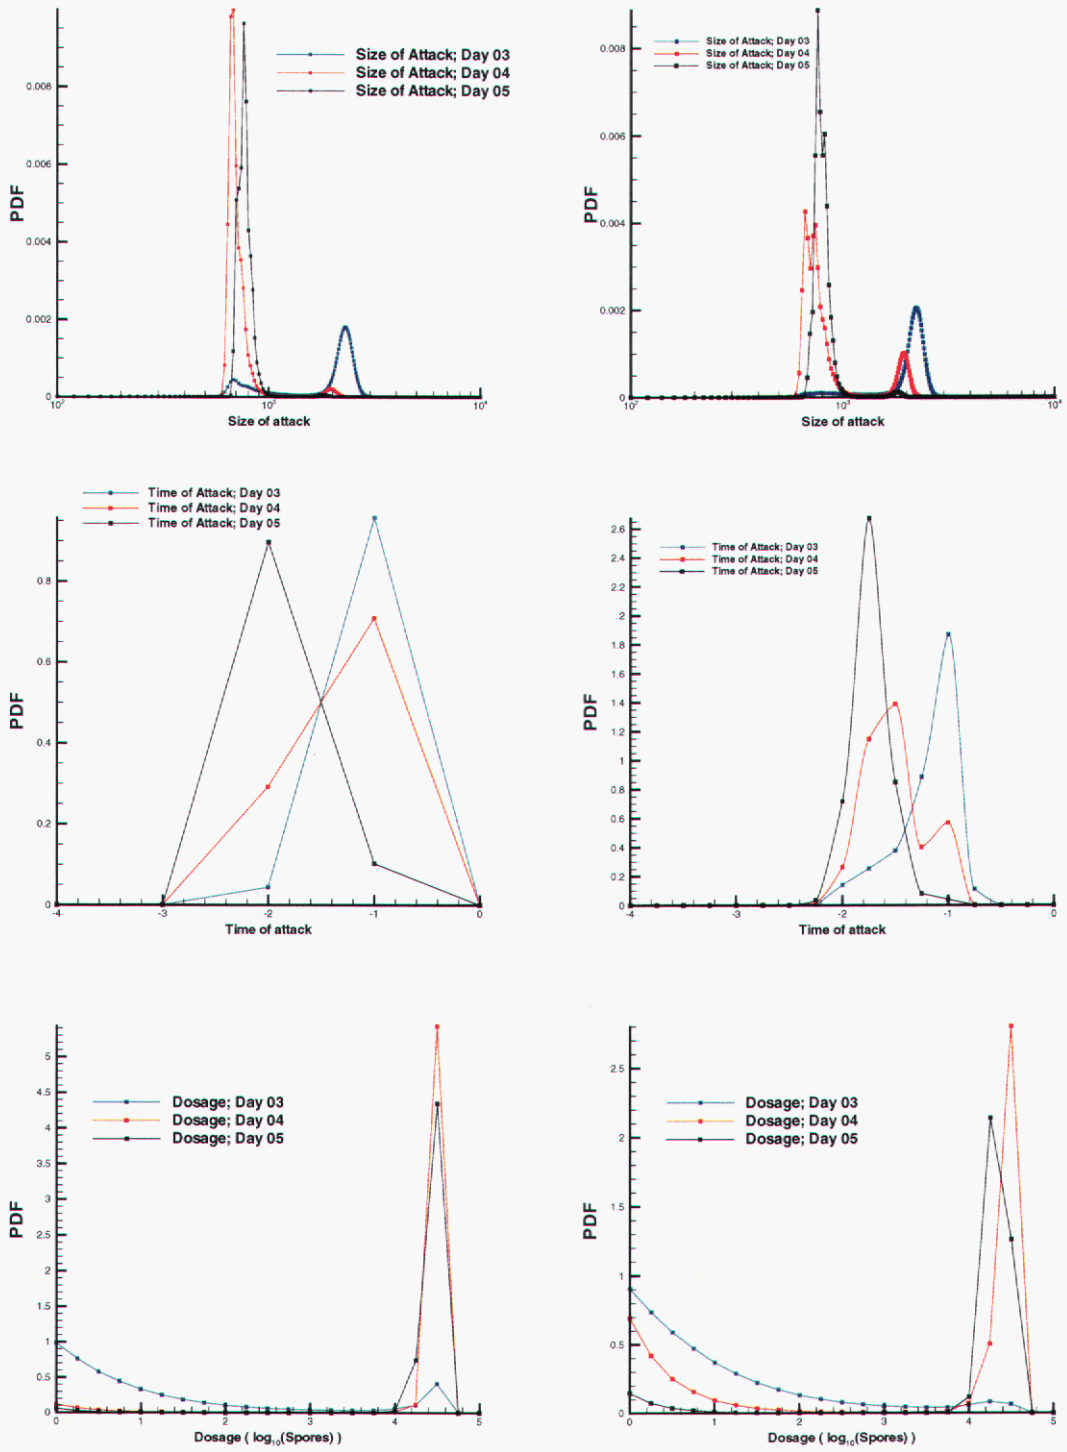

Figure 23. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case N, as tabulated in Table 8. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, \log_{10}(\overline{D})\}$ are $\{1453, -0.75, 3.563\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

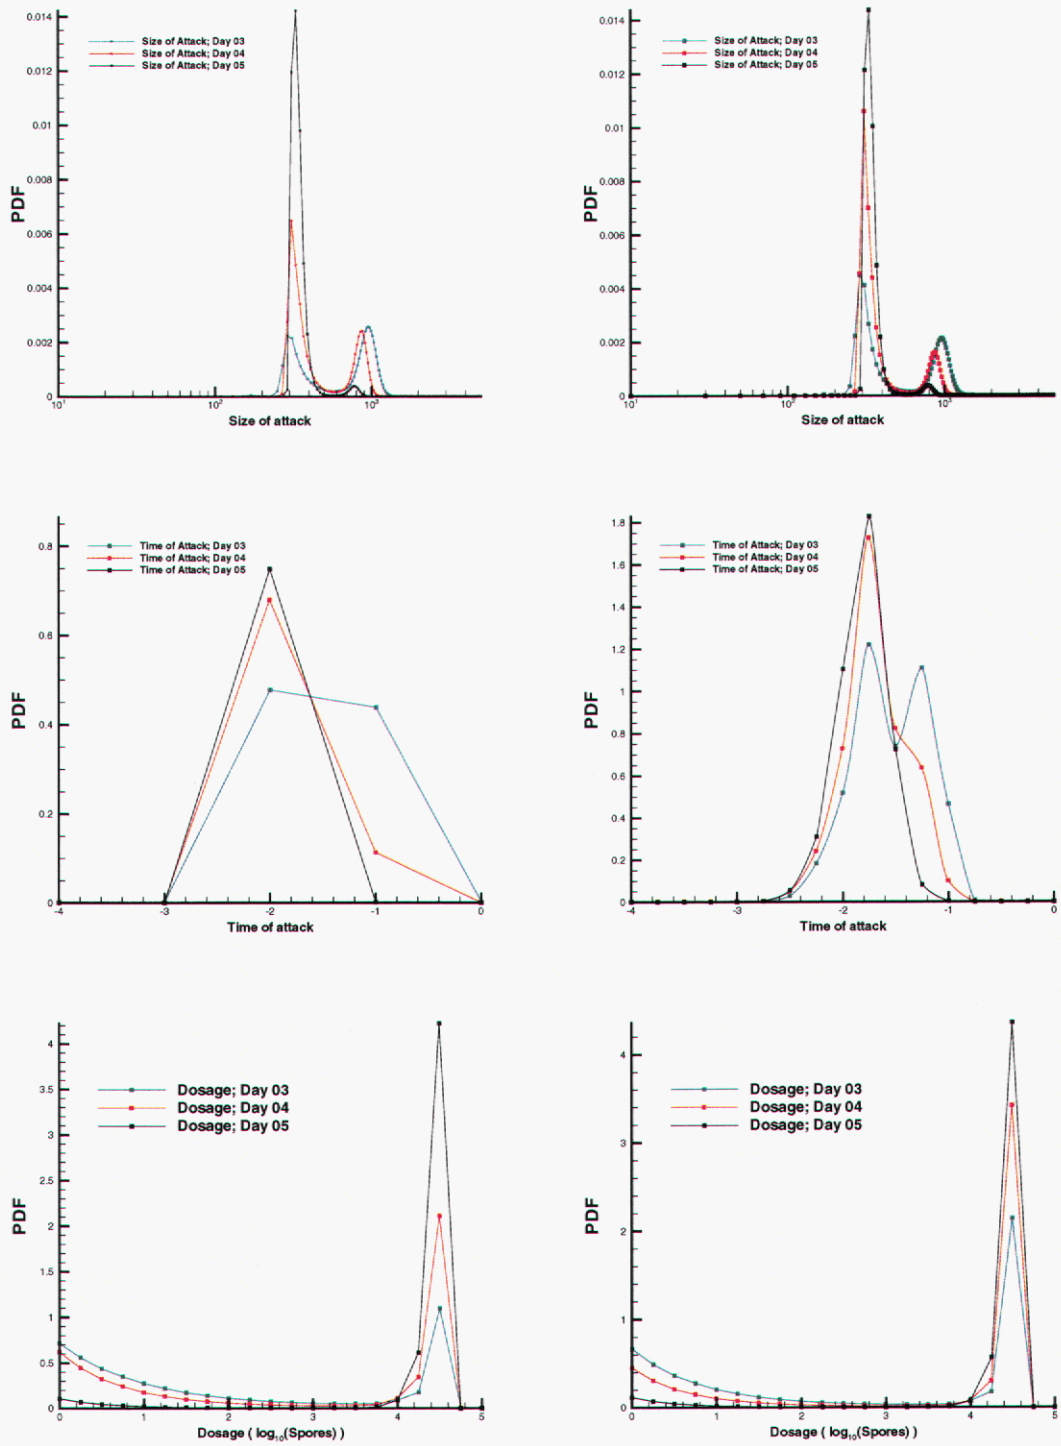

Figure 24. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case O, as tabulated in Table 8. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, \log_{10}(\overline{D})\}$ are $\{453, -0.75, 4.229\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

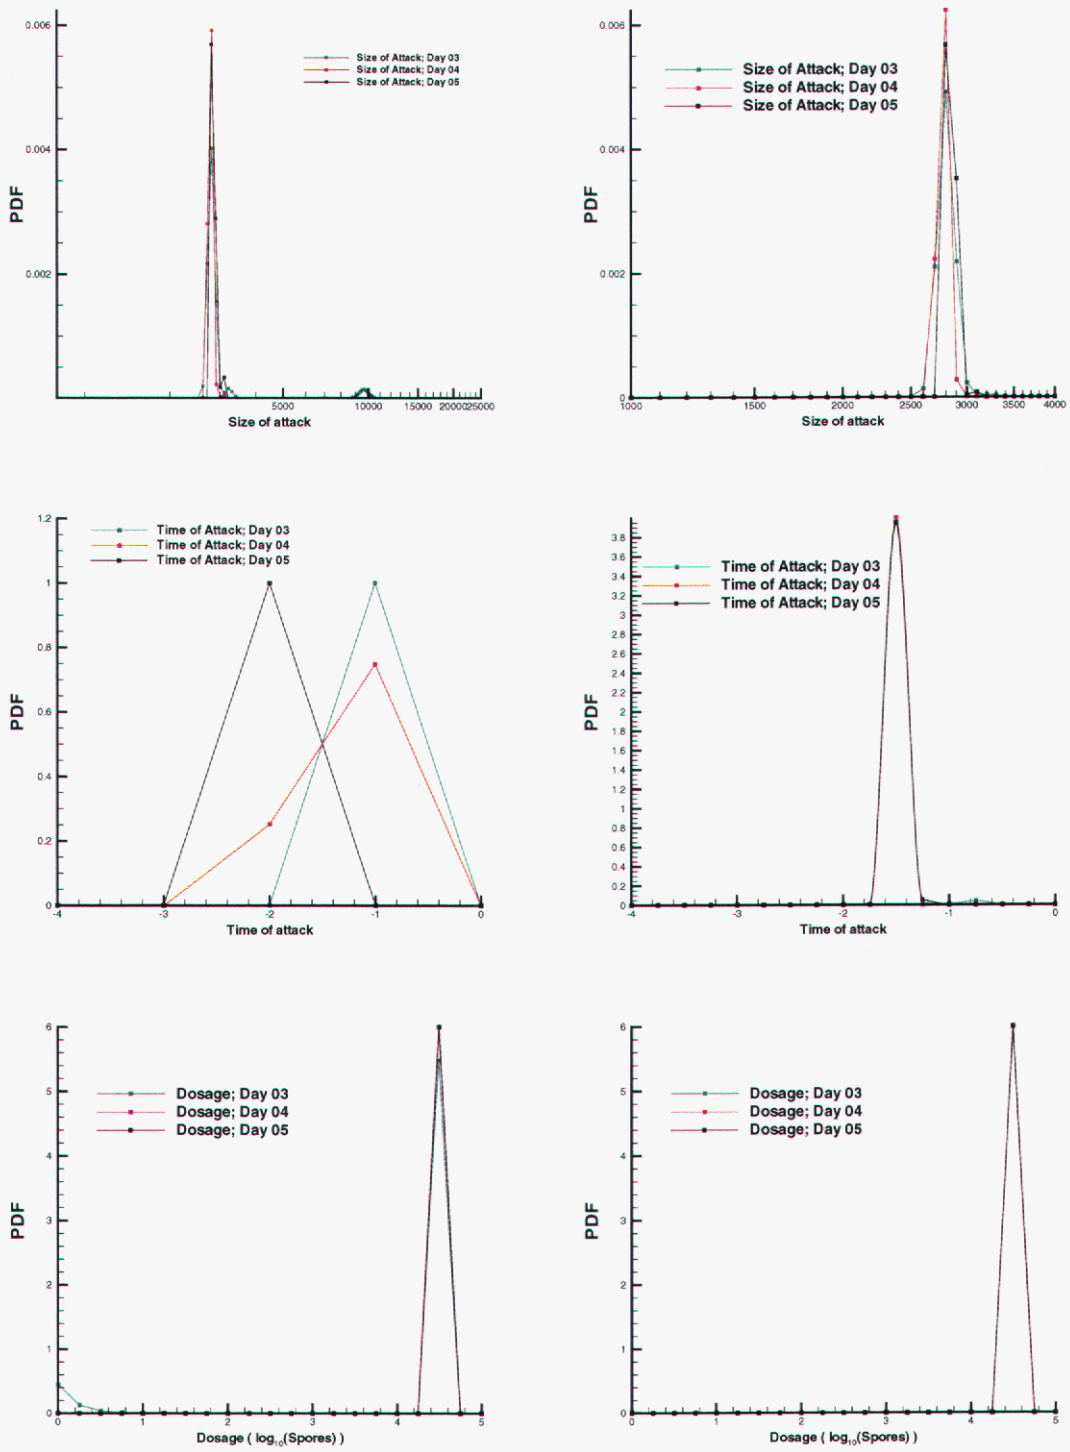

Figure 25. PDFs for N (top), τ (middle) and D (bottom) inferred from the time-series for Case P, as tabulated in Table 8. The low-resolution time-series (24-hour resolution) is used for the PDFs plotted on the left, while the original one is on the right. The correct values for $\{N, \tau, \log_{10}(\overline{D})\}$ are $\{4453, -0.5, 4.22\}$. Inferences are drawn, in both cases, with observation periods of 3, 4 and 5 days' length (blue, red and black lines respectively).

Table 9. The MAP estimate and the 90% confidence intervals (in parentheses) for N, τ and $log_{10}(\overline{D})$ as calculated from Figs. 22, 23, 24 and 25. Only the time-series from Day 5 have been used to calculate the estimates and CI. The number in the curly brackets {} is the correct value.

Case	N	τ	$log_{10}(\overline{D})$
M (6-hr resolution)	170, (130.1 – 243.9)	-1.5, (-2.3 – -0.864)	2.0, (0.235 – 3.82)
M (24-hr resolution	170, (124.9 – 238.9)	-2.0, (-3.7 – -0.265)	2.5, (0.263 – 3.87)
	{161}	{−0.75}	{3.56}
N (6-hr resolution)	760, (722 – 945.5)	-1.75, (-2.09 – -1.359)	4.25, (1.106 – 4.722)
N (24-hr resolution	760, (699.6 – 887.8)	-2.00, (-3.883 – -1.026)	4.50, (4.096 – 4.724)
	{1453}	{−0.75}	{3.56}
O (6-hr resolution)	330, (301 – 709.7)	-1.75, (-2.24 – -1.47)	4.5, (0.99 – 4.72)
O (24-hr resolution)	330, (301.9 – 710)	-2.00, (-3.90 – -1.13)	4.5, (1.03 – 4.72)
	{453}	{−0.75}	{4.23}
P (6-hr resolution)	2800, (2726 – 2992)	-1.5, (-1.72 – -1.272)	4.5, (4.275 – 4.725)
P (24-hr resolution)	2800, (2726 – 2998)	-2.0, (-3.90 – -1.13)	4.5, (4.275 – 4.725)
	{4453}	{−0.75}	{4.22}

4.4 The Sverdlovsk Incident of 1979

It is suspected that on April 2nd, 1979, there was an accidental release of a high-grade anthrax formulation from a military facility in Sverdlovsk, Russia [3]. 70 people are believed to have died [3, 11] and it has been estimated that 80 were infected [11]. This estimate was obtained under the assumption that all the fatalities were due to inhalational anthrax. The Sverdlovsk outbreak provides a good real-world test case for our inference procedure. Wilkening [2] estimates that the average dosage was either around 2-3 spores, based on his Model A, or around 300 spores based on his Model D, which is similar to the competing risks model of Brookmeyer [18]. Meselson [3] estimates 100-2000 spores as the likely dosage.

The Sverdlovsk case presents signifi cant challenges. It was a low-dose attack infecting fewer than a hundred people. The fi rst patient was detected on April 4th, 1979. The time-series of symptom onset is available on a day-by-day basis in [30]. Around April 12th, tetracycline was administered around Sverdlovsk; around the middle of April people were vaccinated. These prophylactic measures probably cured a few and increased the incubation period in others. Further, the data we work with almost certainly contains some recording errors. Also, the data was reconstructed from a variety of sources; public health records had been confi scated by the KGB [3]. Noisiness of the data, the effect of prophylaxis (which is not modeled in our inference process), and the small size of the infected population are expected to stress our inference process.

In Fig. 26 we plot the inferences for N, τ and $\log_{10}(\overline{D})$, based on the data in [30]. Model A2 is used for inference. The data was collected on a daily basis for 42 days, the duration of the outbreak. The time-series used in this study is { 2, 0, 4, 3, 2, 5, 7, 3, 4, 3, 4, 5, 0, 1, 0, 0, 1, 2, 1, 1, 0, 1, 0, 2, 0, 1, 0, 0, 0, 0, 2, 0, 0, 2, 1, 1, 0, 1, 2, 0, 0, 1}. The inference for N centers around 50 consistently, though the earlier inferences underestimate N. The time of release τ was easy to infer. The PDFs for dosage are unclear. The data in [30] shows a long tail after April 18th, roughly after 5-6 days of the start of prophylaxis, when the antibiotics might be expected to affect the progress of the disease. Since we do not model the effect of prophylaxis, we do not include data beyond April 18th. However, we are certainly within a factor of two of the correct value of N even with 9 days of data. In Table 10, we summarize the MAP estimates and confi dence intervals developed from 9 days of data.

Table 10. MAP estimates and 90 % CI for the inferences for the Sverdlovsk incident, after 9 days of data. The average dosage at Sverdlovsk is unknown. Wilkening [2] calculated it to be 2.4 spores or 300 spores, while Meselson et al [3] estimated it to be 100-2000 spores.

Variable	MAP estimate	90 % confi dence interval	True fi gure
N	50	(41.15 – 66.49)	75 – 80
τ	-2	(-7.81 – -1.09)	-2
$\log_{10}(\overline{D})$	1.3	(0.178 – 3.5)	2.4, 300, 100-2000

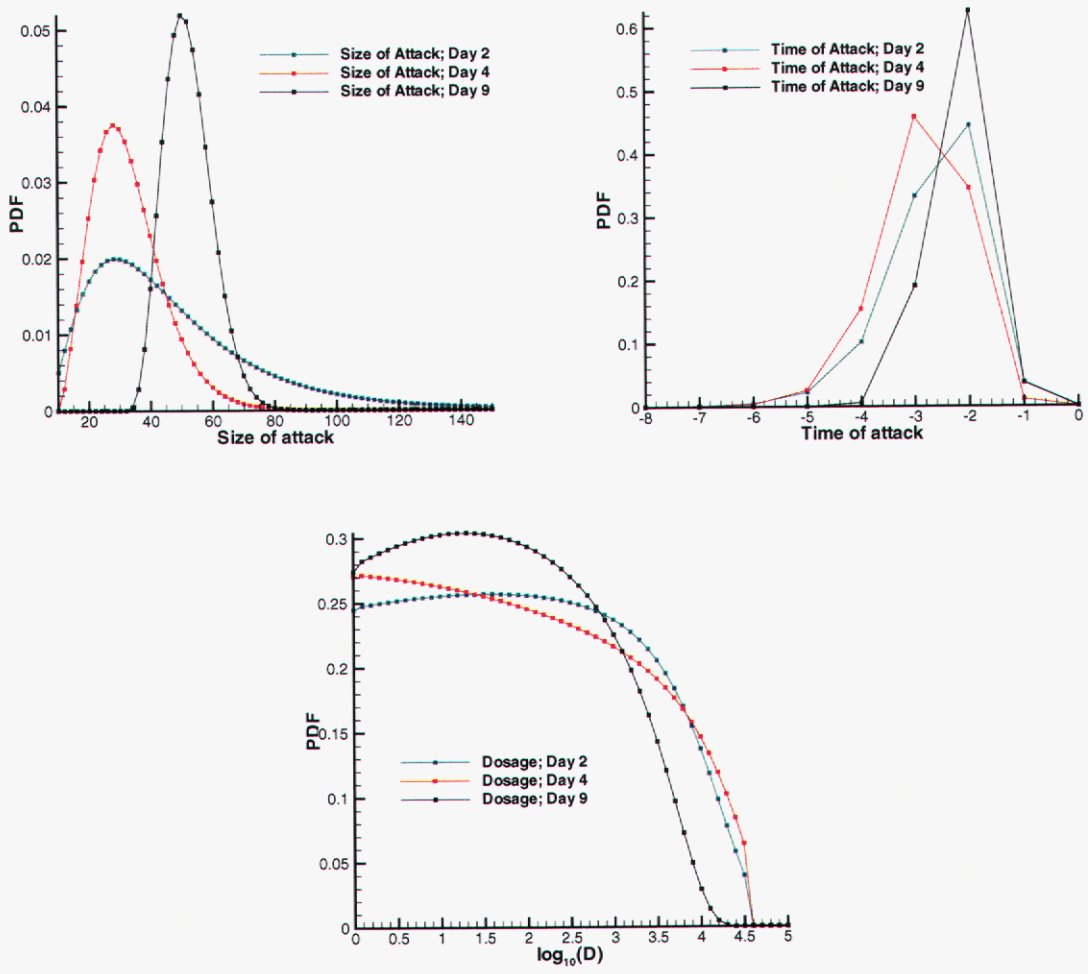

Figure 26. PDFs for the N, τ (above) and $\log_{10}$(dose) (below), for the Sverdlovsk attack. Blue, red and black lines denote inferences that use 3, 5 and 9 days of data.

5 Conclusions

We have developed a promising and extensible approach to characterizing a BT attack, i.e., estimating the number of people infected N, the time of the attack τ, and the average dosage received $\overline{D}$ based on a time-series of people developing symptoms. In this particular study, we have used aerosolized anthrax as the pathogen. The time-series used in this study have a resolution of 6 hours and 24 hours. The problem of characterizing the attack is posed as a Bayesian inverse problem and the inferred characteristics are provided as PDFs for $\{N, \tau, \log_{10}(\overline{D})\}$. The forward model used in the inversion was Wilkening's Model A2 [2], which is a dose-dependent incubation period distribution for inhalational anthrax. We have empirically demonstrated the consistency of this method—i.e., if the observed time-series and the model prediction differ only by a stochastic noise, then, given increasing amounts of data, the PDFs sharpen and the correct values of $\{N, \tau, \log_{10}(\overline{D})\}$ are recovered. This was verified by conducting simulated attacks with Wilkening's Model A2 wherein all infected people were given the same dosage. It was observed that in case of a partially observed outbreak, i.e., with time-series 3–5 days long, the MAP estimates for N and τ are largely correct, regardless of the size of the attack. However, it is difficult to infer dosage for small ($N = 100$) attacks. The use of a higher-resolution time-series (6-hour resolution) generally led to smoother and mostly narrower PDFs; thus, better capturing the structure of the data results in a reduction of uncertainty due to noise and the short time-series. We also notice that partial observation can lead to non-unique characterizations—multimodal PDFs that indicate that the scarce data may support very different characterizations. The inverse problem, formulated under the assumption of constant dosage, was used to characterize BT attacks where the infected population received a spectrum of doses, as obtained from an atmospheric dispersion model. This incurs model errors in addition to errors caused by incomplete observation and stochastic noise. However, the 90% CI obtained from 5 days of data invariably bracketed the true characterization. The effect of a high-resolution time-series on the reduction of uncertainty in the inferences was rather muted. Thus while the model errors incurred by applying a constant-dose model to a variable dose case were not negligible, the inverse problem was completely successful in bounding the characterization.

Given the paucity of experimental data on the incubation period of anthrax, it would be unrealistic to expect that the forward model in any inference procedure would be an accurate model for the particular strain/preparation of the pathogen, should an anthrax-based BT attack occur. We incorporate this reality in this study. We address the issue of "model uncertainty"—that of a systematic, as opposed to stochastic, difference between the forward model and the phenomenon that produces the time-series—by simulating the BT attacks using Wilkening's Model D. Models A2 and D reproduce experimental results in non-human primates best and thus can be considered to be closest approximations to the true behavior of inhalational anthrax. The introduction of the systematic difference degrades the inferences. Model A2 predicts slightly shorter incubation periods; when "fitted" to data generated by Model D (which predicts an epidemic curve that rises in time slower than A2's predictions), the inference method compensates by underestimating N. The time of attack is estimated to be about a day too late. Dosage is difficult to estimate for small attacks. These general observations hold true when the assumption of uniform dosage is relaxed. The MAP estimates of N are within a factor of two of the true figure. Here again the effect of a higher-resolution time-series is rather muted—the errors introduced by the difference in models

Table 11. A summary of the effect of an attack's properties and the data collection methodology on the quality of inferences of attacks where infected people receive a spectrum of doses i.e. Cases RA, RB, RC RD, M, N, O, P and the Sverdlovsk incident of 1979.

Attack's and time-series properties	Inferred characteristics		
	N	τ	$\log_{10}(\overline{D})$
Size of attack	Bounded successfully regardless of size.	MAP estimate within time-series resolution, irrespective of size.	MAP accurately estimated but CIs are large, irrespective of size of attack.
Resolution of observed data	Neither MAP estimates nor CIs much affected by higher resolution.	MAP estimates within time-series resolution; consequently, higher resolution has a favorable effect.	MAP estimates more accurate with higher resolution; CIs not signifi cantly affected
Discrepancy between attack and inference models	MAP estimate within a factor of 2 (underestimate)	MAP estimate a day too late (consistently)	MAP estimates close, except for small, lower-dose attacks; CIs are wide.

and the uniform dose assumption in the forward model completely dominate the noise that could be reduced/rectifi ed by higher-resolution data. Non-unique characterizations were also observed; with the availability of more data, the PDFs became unimodal. However, due to the systematic errors introduced by "model uncertainty," the correct characterization was not recovered exactly. We summarize these in Table 11.

We fi nally applied the Bayesian approach to the outbreak at Sverdlovsk. The outbreak was small (about 80 were infected, and 70 died) and due to prophylaxis during the outbreak, the course of the disease was modifi ed. Using pre-prophylaxis data, our MAP estimate of N was 50 after 9 days of (daily) data. The time of attack was estimated to be -2, which agrees with the generally accepted reconstruction [3]. The dosage, as might be expected from the small size, was indeterminate.

Due to the Bayesian nature of the approach, the method is extensible; i.e., non-temporal aspects of the phenomenon being inferred can be incorporated via prior distributions. In the absence of such additional information, the present inferences were drawn using broad uniform priors. One of the easiest ways of improving the effi ciency of this method is to construct informative priors for N and τ from data collected by syndromic surveillance systems [5]. Priors drawn from such a source will invariably favor values of N higher than the true one and may serve as a counterbalance to our Bayesian methods that tend to underestimate N. Further, our approach is purely temporal;

while it exploits the rise of the epidemic curve, it completely ignores the spatial distribution of the symptomatic patients, which may be reliably expected to align with the wind direction. Work on this extension to the present approach will commence shortly.

Ultimately, characterization of a BT attack is of little value if not coupled to a response plan. Given the rarity of human inhalational anthrax, any appreciable number of anthrax cases (e.g., more than 3) over a 1–2 day period will probably lead to the implementation of response plans and consequence management. Some of the responses—for instance, city-wide prophylaxis—do not need much characterization beyond the identifi cation of the pathogen. Yet if this response is triggered by the diagnosis of patients, rather than by early warning sensors, one must expect that prophylaxis would arrive too late to be of help to many infected people. In such a case provisions would have to be made for medical personnel and material, which in certain instances may overwhelm local resources. While a broad and rapid response has its attractions, one could also respond in a measured manner by placing requisitions for certain resources while alerting others. Our ability to describe attack parameters probabilistically provides a simple means of estimating what the requisitions should be, based on an acceptable level of risk; what is less clear is the extent of the alerts that would need to be issued. Alerting of resources, with a view towards rapid requisitioning and delivery in the near future, can be thought of as a means of hedging against the risk of an inadequate initial requisition. Thus, we begin to see the outlines of a risk-based approach that designs a response commensurate with the BT attack at hand. Apart from being sustainable, such an approach will be smaller and quicker; it may also be safer, since a broad and rapid response can be vulnerable to feints.

References

[1] Philip S. Brachman, Arnold F. Kaufmann, and Frederic G. Dalldorf. Industrial inhalational anthrax. *Bacteriological Reviews*, 30(3):646–657, 1966.

[2] D. Wilkening. Sverdlovsk revisted : Modeling human inhalational anthrax. *Proceedings of the National Academy of Science*, 103(20):7589–7594, May 2006.

[3] Matthew Meselson, Jeanne Guillemin, Martin Hugh-Jones, Alexander Langmuir, Ilone Popova, Alexis Shelokov, and Olga Yampolskaya. The Sverdlovsk anthrax outbreak of 1979. *Science*, 266:1202–1208, 1994.

[4] John A. Jernigan et al. Bioterrorism-related innhalational anthrax: The fi rst 10 cases reported in the United States. *Emerging Infectious Diseases*, 7(6):933–944, 2001.

[5] Anna Goldenberg, Galit Shmueli, Richard A. Caruana, and Stephen E. Fienberg. Early statistical detection of anthrax outbreaks by tracking over-the-counter medication sales. *Proceedings of the National Academy of Sciences*, 99:5237–5240, 2002.

[6] Tara O'Toole, Michael Mair, and Thomas V. Inglesby. Shining light on Dark Winter. *Clinical Infectious Diseases*, 34:972–983, 2002.

[7] Stephen Eubank, Hasan Guclu, V. S. Anil Kumar, Madhav V. Marathe, Aravind Srinivasan, Zoltan Toroczkai, and Nan Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 429:180–184, 2004.

[8] Lynn Yang, Ben Wu, Dawn Manley, Julie Freutel, Marilyn Hawley, Keith Vanderveen, Jaideep Ray, Zach Heath, Christine Yang, Donna Djordjevich, Han Lin, John Jungels, Luis Hernandez, Bill Wilcox, Andrew Rothfuss, Michael Chen, Stephen Mueller, Howard Hirano, and Michael Johnson. BioNet Systems Modeling and Analysis. Technical Report SAND2005-7842, Sandia National Laboratories, Albuquerque, NM, 2005. Unclassifi ed limited release.

[9] J. Ray, Y. M. Marzouk, H. N. Najm, M. Kraus, and P. Fast. A Bayesian method for characterizing distributed micro-releases: I. The single-source case for non-contagious diseases. SAND Report SAND2006-1491, Sandia National Laboratories, Livermore, CA 94551-0969, March 2006. Unclassifi ed unlimited release.

[10] J. Walden and E. H. Kaplan. Estimating time and size of bioterror attack. *Emerging Infectious Diseases*, 10(7):1202–1205, 2004.

[11] R. Brookmeyer, N. Blades, M. Hugh-Jones, and D. A. Henderson. The statistical analysis of truncated data: application to the Sverdlovsk anthrax outbreak. *Biostatistics*, 2:233–247, 2001.

[12] Ron Brookmeyer and Natalie Blades. Statistical models and bioterrorism : Application to the U.S. anthrax attacks. *Journal of the American Statistical Association*, 98(464):781–788, 2003.

[13] D. W. Henderson, S. Peacock, and F. C. Belton. Observations on the prophylaxis of experimental pulmonary anthrax in the monkey. *Journal of Hygiene*, 54:28–36, 1956.

[14] Wilhelm S. Albrink and Robert J. Goodlow. Experimental inhalational anthrax in the chimpanzee. *The American Journal of Pathology*, 35(5):1055–1065, 1959.

[15] C. A. Gleiser, C. C. Berdjis, H. A. Hartman, and W. S. Gochenour. Pathology of experimental respiratory anthrax in *Macaca Mulatta*. *British Journal of Experimental Pathology*, 44:416–426, 1963.

[16] M. M. Friedlander, S. L. Welkos, M. L. Pitt, J. W. Ezzell, P. L. Worsham, K. J. Rose, B. E. Ivins, J. R. Lowe, G. B Howe, P. Mikesell, and Wade B. Lawrence. Postexposure prophylaxis against experimental inhalation anthrax. *Journal of Infectious Disease*, 167(5):1239–1243, 1993.

[17] B. E. Ivins, M. L. M. Pitt, P. F. Fellows, J. W. Farchaus, G. E. Benner, D. M. Waag, S. F. Little, G. W. Anderson, P. H. Gibbs, and A. M. Friedlander. Comparative effi cacy of experimental anthrax vaccine candidates against inhalational anthrax in rhesus macaques. *Vaccine*, 16(11/12):1141–1148, 1998.

[18] R. Brookmeyer, E. Johnson, and S. Barry. Modelling the incubation period of anthrax. *Statistics in Medicine*, 24:531–542, 2005.

[19] H. N. Glassman. Discussion on industrial inhalational anthrax. *Bacteriological Review*, 30:657–659, 1966.

[20] H. A. Druett, D. W. Henderson, L. Packman, and S. Peacock. Studies on respiratory infection. I.The influence of particle size on respiratory infection with anthrax spores. *Journal of Hygiene*, 51:359–371, 1953.

[21] C. N. Haas. On the risk of mortality to primates exposed to anthrax spores. *Risk Analysis*, 22(2):189–193.

[22] W. R. Hogan, G. Cooper, G. L Wallstrom, and M. Wagner. An Improved Bayesian Aerosol Release Detector. Technical report, The RODS Laboratory, 550 Cellomics Building, 100 Technology Drive, Pittsburgh, PA 15219, 2005.

[23] Real-time Outbreak and Disease Surveillance. http://rods.health.pitt.edu/default.htm.

[24] Fu-chang Tsui, Jeremy U. Espino, Virginia M. Dato, Per H. Gesteland, Judith Hutman, and Michael M. Wagner. Technical Description of RODS : A Real-time Public Health Surveillance System. *Journal of the Americal Medical Informatics Association*, 10(5):399–408, 2003.

[25] D. Bruce Turner. *Workbook of Atmospheric Dispersion Estimates: An Introduction to Dispersion Modeling*. Lewis Publishers, CRC Press LLC, 2000 N.W. Corporate Blvd. Boca Raton, FL 33431, 1994.

[26] G. Peter Lepage. A new algorithm for adaptive multidimensional integration. *Journal of Computational Physics*, 27:192–203, 1978.

[27] GNU Scientifi c Library. http://www.gnu.org/software/gsl/.

[28] P. E. Sartwell. The Distribution of Incubation Periods of Infectious Diseases. *American Journal of Hygiene*, 51:310–318, 1950.

[29] Frederick Klein, Jerry S. Walker, David F. Fitzpatrick, Ralph E. Lincoln, Bill G. Mahlandt, William L. Jones, James P. Dobbs, and Kenneth J. Hendrix. Pathophysiology of anthrax. *Journal of Infectious Diseases*, 116(2):123–138, 1966.

[30] Thomas V. Inglesby, Donald A. Henderson, John G. Barteltt, Michael S. Ascher, Edward Eitzen, Arthur M. Friedlander, Jerome Hauer, Joseph McDade, Michael T. Osterholm, Tara O'Toole, Gerald Parker, Trish M. Perl, Philip K. Russel, and Kevin Tomcat. Anthrax as a biological weapon - Medical and public health management. *J. Am. Med. Assoc.*, 281(18):1735–1745, 1999.

A Derivation of Anthrax Incubation Models

In this Appendix, we discuss in detail the derivation of two mathematical models from prior literature for predicting the onset time of symptoms in patients exposed to anthrax. In particular, we consider Wilkening's Model D and Model A2[2] which extend Brookmeyer *et al.*[18].

Anthrax spores are "inert" - they do not reproduce. Under favorable conditions, for instance in human lymph nodes or in a blood agar medium, they first germinate into vegetative anthrax bacterial cells and proceed to replicate. While inside a living being, anthrax spores are acted upon simultaneously by two competing processes (1) clearance (e.g. by the immune system) at rate θ per spore per day and (2) germination into vegetative cells at rate λ per spore per day. Thus given a person with D spores (initially), the governing equation for the remaining number of spores Δ is $\dot{\Delta} = (\lambda+\theta)\Delta$. The number of spores that germinate Δ_g are given by the relation $\dot{\Delta}_g = \lambda\Delta$. The solutions for these equations are

$$\Delta = D\exp(-(\lambda+\theta)t) \qquad \text{and} \qquad \Delta_g = \frac{D\lambda}{\lambda+\theta}\exp(-(\lambda+\theta)t).$$

The probability that a spore will germinate is $q = \Delta_g/\Delta = \lambda/(\lambda+\theta)$. Then, given a person with D spores (initially), the number of spores that will germinate in infinite time is $\nu = D\lambda/(\lambda+\theta)$.

In reality, the number of spores that germinate will vary. The probability that X spores germinate, $0 \le X \le D$, is given by the binomial distribution:

$$\begin{aligned}\frac{D!}{X!(D-X)!}q^X(1-q)^{D-X} &= \frac{D!}{X!(D-X)!}\frac{\nu^X}{D}\left(1-\frac{\nu}{D}\right)^{D-X} \\ &= \frac{D(D-1)\ldots(D-X+1)}{D^X}\frac{\nu^X}{X!}\left(1-\frac{\nu}{D}\right)^{D}\left(1-\frac{\nu}{D}\right)^{-X}.\end{aligned} \tag{9}$$

In the limit that $\nu \ll 1$ or $X/D \ll 1$, the probability $P(X)$ that X spores will germinate is given by

$$P(X) = \lim_{D\to\infty}\frac{D!}{X!(D-X)!}q^X(1-q)^{D-X} = 1.\frac{\nu^X}{X!}.\exp(-\nu).1. \tag{10}$$

The probability that at least 1 spore will germinate (also called the attack rate) is

$$p = 1 - P(X=0) = 1-\exp(-\nu) = 1-\exp\left(-\frac{D\lambda}{\lambda+\theta}\right). \tag{11}$$

Given D spores initially, the number of spores that will germinate by time t is

$$\int_0^t \dot{\Delta}_g du = \int_0^t D\lambda\exp(-(\lambda+\theta)u)\,du = \frac{\lambda D}{\lambda+\theta}(1-\exp(-(\lambda+\theta)t)).$$

Thus the probability of 1 spore germinating by time t is

$$\frac{\lambda}{\lambda+\theta}(1-\exp(-(\lambda+\theta)t)) = \frac{\lambda}{\lambda+\theta}Q(t).$$

Note that $\lim_{t\to\infty} Q(t) \to 1$.

The probability that at least 1 spore will germinate by time t is

$$1-\exp\left(-\frac{D\lambda}{\lambda+\theta}Q(t)\right) = \overline{F}(t;D,\lambda,\theta).$$

Note that $\lim_{t\to\infty}\overline{F} \to p$. We defi ne

$$F(t;D,\lambda,\theta) = \frac{\overline{F}(t;D,\lambda,\theta)}{p}. \tag{12}$$

F is the cumulative distribution function for a spore to germinate into a vegetative cell. Vegetative cells multiply, though the rate of multiplication varies for each cell. However, given a colony of vegetative anthrax cells, there is a growth rate c and one may fi nd a time t_2 in which the population A of the colony doubles. Further, we assume that people exhibit symptoms when the bacterial load reaches N_{thresh}, which takes, on average, time t_M.

$$\dot{A} = cA, \quad A = D\exp(ct), \quad c = \frac{\log(2)}{t_2} \quad \text{and} \quad t_M = \frac{t_2}{\log(2)}\log\left(\frac{N_{thresh}}{D}\right),$$

where log denotes the natural logarithm. Typically, the expression for t_M is written with a lag time

$$t_M = t_{lag} + \frac{t_2}{\log(2)}\log\left(\frac{N_{thresh}}{D}\right). \tag{13}$$

In reality, the time s to symptoms, post-germination, will follow a distribution. Since $s \geq 0$, Wilkening [2] assumed it to be log-normal (Brookmeyer, in [18], assumed it to be exponential)

$$g(s;D) = \frac{1}{\sqrt{2}\pi\sigma_s s}\exp\left(-\frac{1}{2}\frac{\log^2(s/M_s)}{\sigma_s^2}\right). \tag{14}$$

M_s is the median time to symptoms, post-germination, and one can set it to be equal to t_M. Thus the probability that at least a single spore germinates and the bacterial load reaches N_{thresh} so as to generate symptoms by time t, i.e. the cumulative distribution function for the incubation period is

$$C_{ModelD}(t,D) = \int_0^t F(t-s;D,\lambda,\theta)g(s;D)ds. \tag{15}$$

This is Wilkening's Model D. The values of the parameters are $\theta = 0.109$ day^{-1}, $\lambda = 8.79 \times 10^{-6}$ day^{-1}, $t_{lag} = 1$ hour, $t_2 = 2.07$ hour, $N_{thresh} = 10^9$ and $\sigma_s = 0.544$ day^{-1}.

Sartwell [28] found that the incubation period for a number of diseases were log-normally distributed, which is at odds with Eq. 15. Wilkening's Model A2 assumes a log-normal distribution,

$$C_{ModelA2}(t,D) = \frac{1}{2}\left[1 + \text{erf}\left(\frac{\ln(t/t_0)}{\sqrt{2}S}\right)\right], \quad S = 0.804 - 0.079\log_{10}(D) \ . \tag{16}$$

where t_0, the median incubation time, is obtained by solving the integral equation obtained from Eq. 15

$$0.5 = \int_0^{t_0} F(t-s;D,\lambda,\theta)g(s;D)ds.$$

However, while solving for t_0, Wilkening used a slightly different set of parameters: $\theta = 0.11$ day^{-1}, $\lambda = 8.84 \times 10^{-6}$ day^{-1}, $t_{lag} = 1$ hour, $t_2 = 2.06$ hour, and $\sigma_s = 0.542$ day^{-1}.

This page intentionally left blank.

B Methodology for Obtaining a Dosage Distribution Consistent with Atmospheric Dispersion over a Geographically Distributed Population

The distribution of dosages due to an atmospheric release of an aerosol can be modeled using a simple Gaussian plume model [25]. Further, if this release occurs over a domain with a non-uniform population distribution, one obtains a distribution of exposed people over a dosage spectrum. In this section, we describe a simple way to obtain a population-dosage distribution.

We consider a square domain, L km on each side. In this study $L = 10$ km. The domain is divided into N blocks per side; in this study $N = 100$. 25 population clusters are chosen in the form of Gaussian kernels $A\exp(-r^2/R^2)$, where $r^2 = |\mathbf{x} - \mathbf{x_0}|^2$. The strength of the kernel A, its center $\mathbf{x_0}$, and its lengthscale R are chosen randomly. The population density in any block, with its center at $\mathbf{x}$, is a sum of the strengths of all the 25 kernels. The strengths of the kernels are scaled to obtain a total population (in the domain) of P_{domain}. The population in a given block is obtained by multiplying the population density with the block area. This creates a geographically distributed population distribution.

The number of people exposed (i.e., people who breathed in the aerosol, but may or may not develop symptoms) and infected (i.e., people who will develop symptoms) is dependent on the location and size of the release and direction of the wind. We release 10^{13} spores at the origin, at a height of 100 meters. A wind speed of 4 m/s, and a Pasquill stability class of "B" is assumed. Pasquill stability classes indicate atmospheric stability; this particular case (class B) indicates a moderately unstable atmosphere with a 4 m/s wind and strong daytime insolation. Details of Pasquill stability classes and atmospheric dispersion are in [25]. In our study, wind directions are measured in degrees from due north; that is, a wind direction of zero degrees is a wind from due north, 90 degrees is a westerly wind and a direction of 180 degrees is a wind from due south. The release is assumed to be an explosive point release, and the concentration of the aerosol at any point (x,y) on the ground and any time t is given by [25]

$$\chi(x,y,t) = \frac{2Q_T}{(2\pi)^{3/2}\sigma_{x'}\sigma_{y'}\sigma_{z'}} \exp\left(-\frac{(x'-ut)^2}{2\sigma_{x'}^2}\right) \exp\left(-\frac{(y')^2}{2\sigma_{y'}^2}\right) \exp\left(-\frac{(H')^2}{2\sigma_{z'}^2}\right) \tag{17}$$

where (x',y') are the coordinates of (x,y) in a frame of reference where the x-axis is aligned with the wind. $\sigma_{x'}, \sigma_{y'}$ and $\sigma_{z'}$ are coeffi cients dependent on x' and the Pasquill stability class. H is the height of release and χ is the concentration of the aerosol in spores per unit volume. u is the wind velocity. Q_T is the total number of spores released. The relation between $\mathbf{x}'$ and $\mathbf{x}$ is given by

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} \cos(\pi-\theta) & -\sin(\pi-\theta) \\ \sin(\pi-\theta) & \cos(\pi-\theta) \end{pmatrix} \begin{pmatrix} x' \\ y' \end{pmatrix}$$

where θ is the wind direction. Assuming a breathing rate β of 30 liters a minute, one can obtain an expression for the number of spores breathed in per unit time. Integrating over infi nite time,

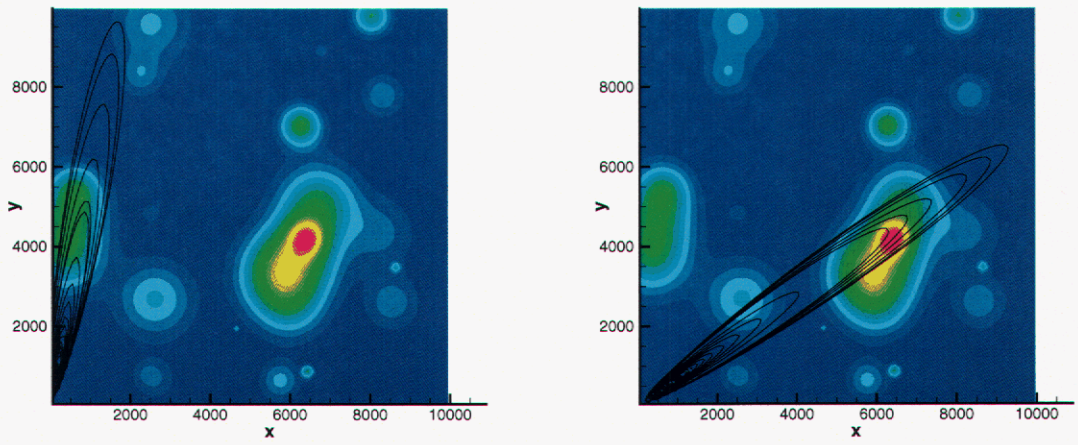

Figure B.1. Dosage plumes plotted over the population distribution for $\theta = 170^o$ (left) and 125^o (right). We see on the right that the extremities of the plume extend into a high population density region. Thus we may expect a substantial number of high-dosage cases, resulting in a higher average dosage $\overline{D}$.

one obtains the total number of spores D breathed in by a person positioned at (x,y) (alternatively, (x',y'))

$$D = \frac{Q_T \beta}{2\pi\sigma_{x'}\sigma_{y'}\sigma_{z'}} \exp\left(-\frac{(y')^2}{2\sigma_{y'}^2}\right) \exp\left(-\frac{(H')^2}{2\sigma_{z'}^2}\right) \left(1 + erf(x')\right).$$

The dosage in a given block is decided by the location of its center. Using Glassman's formula to judge the probability of showing symptoms (in infi nite time) given a dosage D, we get the probability density $P(D)$

$$P(D) = \frac{1}{2\pi DS} \exp\left[-\frac{1}{2}\left(\frac{\ln(D/D_0)}{S}\right)^2\right]$$

with $D_0 = 8000$ and $S = 3.2888$. Glassman's relation is used if we employ Model A2 to simulate the BT attack; if Model D is used, we employ Eq. 11 to determine the probability of infection, given a dose D. Since the population in a block is known, we can use the cumulative distribution determined from the probability distribution above to calculate the number of people in the block who will become infected and proceed to develop symptoms over time, per the incubation period model.

In this study, we employ $P_{domain} = 3 \times 10^5$ and 3×10^6; $\theta = 170^o$ and 125^o. The dosage plumes, superimposed on the population distribution, are plotted in Fig. B.1.

Distribution:

1	MS 9056	Youssef Marzouk	8351
4	MS 9159	Jaideep Ray	8964
2	MS 9018	Central Technical Files	8944
2	MS 0899	Technical Library	4536