

Studies of Neurological Transmission Analysis using Hierarchical Bayesian Mixture Models

MIKE WEST

Institute of Statistics and Decision Sciences
Duke University, Durham NC 27708-0251

Abstract.

Hierarchically structured mixture models are studied in the context of data analysis and inference on neural synaptic transmission characteristics in mammalian, and other, central nervous systems. Mixture structures arise due to uncertainties about the stochastic mechanisms governing the responses to electro-chemical stimulation of individual neurotransmitter release sites at nerve junctions. Models attempt to capture scientific features such as the sensitivity of individual synaptic transmission sites to electro-chemical stimuli, and the extent of their electro-chemical responses when stimulated. This is done via suitably structured classes of prior distributions for parameters describing these features. Such priors may be structured to permit assessment of currently topical scientific hypotheses about fundamental neural function. Posterior analysis is implemented via stochastic simulation. Several data analyses are described to illustrate the approach, with resulting neurophysiological insights in some recently generated experimental contexts. Further developments and open questions, both neurophysiological and statistical, are noted.

Keywords: Bayesian Computations; Deconvolution; Mixture models; Neural Response Activity; Parameter Identification in Mixtures; Quantal Synaptic Transmission Hypothesis.

Research partially supported by the NSF under grants DMS-9024793, DMS-9305699 and DMS-9304250. This work represents part of a collaborative project with Dr Dennis A Turner, of Duke University Medical Center and Durham VA. Data was provided by Dr Turner and by Dr Howard V Wheal of Southampton University. A slightly revised version of this paper is published in the *Journal of the American Statistical Association* (vol 92, pp587-606), under the modified title *Hierarchical Mixture Models in Neurological Transmission Analysis*. The author is the recipient of the 1997 Mitchell Prize for “*the Bayesian analysis of a substantive and concrete problem*” based on the work reported in this paper.

1 SCIENTIFIC CONTEXT AND BACKGROUND

The last fifteen years or so has seen the development of a major field of study concerning the stochastic characteristics of electrochemical transmission between nerves at synaptic junctions in mammalian, and other, central nervous systems (Walmsley, Edwards and Tracy 1987). Fundamental debates in the discipline concern the scope of validity of specific hypotheses about synaptic communication, with variants of a basic, ‘quantal’ hypothesis being of central interest (Clements 1991; Martin 1966; Walmsley, Edwards and Tracy 1988). Much recent experimental and ensuing statistical analysis has been concerned with examining the evidence for or against this simple hypothesis in differing synaptic junctions and under differing environmental conditions (Kullman and Nicoll 1992; Redman 1990; Turner 1987; Turner and Wheal 1991). The relevant statistical issues have been addressed in derived models based on discrete mixture distributions for recorded neural signals, under ranges of partial constraints that impose characteristics either in conformity with quantal hypotheses or allowing specific kinds of departures from such hypotheses (Ling and Tolhurst 1983; Wong and Redman 1980). Our previous statistical work in the area has complemented existing approaches to mixture estimation, such as based on EM computations (Kullman 1989) and entropy formulations (Kullman 1992), by introducing general classes of Bayesian mixture models and their associated analyses (Escobar and West 1995; Turner and West 1993; West and Cao 1993; West and Turner 1994). It is characteristic of these approaches that they adopt a neutral scientific standpoint on the issue of mixture model composition; like previous, non-Bayesian approaches, the framework has been directed at modelling experimental data using general and flexible classes of mixture models, and then attempting to evaluate the scientific issues of mixture component structure, and such issues as the validity of the quantal hypothesis, on a less formal post-hoc basis. The current article, by contrast, reports on our recent work with new models that are motivated by the desire to more directly represent the underlying scientific structure, and hence the resulting data configurations. Hierarchical mixture models are structured to reflect ranges of possible, and scientifically plausible, deviations away from the central, and simplifying, quantal structure hypothesis, and so provide flexibility in modelling observed non-quantal structure as well as direct and formal assessment of the hypothesis. This allows exploration and reporting of a variety of observed departures from this basic hypothesis, and assists in some of our current work focused simply on cataloguing the diversity of EPSP phenomena. In developing these models and their analyses, we encounter several issues of wider interest involving, in particular, prior modelling, Bayesian computation, parameter identifiability and model comparisons; these issues, among others, are discussed in the context of both theoretical development and several analyses of neurophysiological data sets. Current research frontiers, and future prospects, are also noted. To begin, we provide some background on the experimental context and examples of current experimental data, leading into the basic modelling framework.

Current neurophysiological experimentation enables the isolation of rather small areas of nerve tissue on large nerve cells in mammalian central nervous systems, and in which synaptic connections consist of very small (though uncertain) numbers of transmission sites (Turner and Schlieckert 1990). Electro-chemical techniques are used to stimulate nerve tissue to induce neurotransmitter release across identified synapses. Electrical potentials are induced in the cell body of the receiving nerve cell, and estimates of the maximum levels of potential reached as a result of transmission are

recorded as summary EPSP (excitatory post-synaptic potential) signals; the units of measurement are millivolts, or mV. This is repeated to produce samples of EPSP levels, with sample sizes usually in the hundreds. We refer to these kinds of measurements as EPSP signal recordings. The signal level achieved on any occasion measures, indirectly, the level of chemical neurotransmitter released at synaptic sites on the stimulated neuron on that occasion. The measurements are subject to various sources of error: background variability in electrical potentials in the cell, typically including rather minor spontaneous transmission effects, synaptic noise and noise arising in the physical and electronic recording process. Interspersed between signal recordings are concomitant measurements made under the same experimental protocols but with no active stimulation of the nerve tissue, so producing similar samples of measurements on the background noise. We refer to these kinds of measurements simply as noise recordings; they provide data on noise corrupting whatever structure is evidenced in the EPSP signal data. From the signal and noise data sets together, we are interested in inferences on the neurophysiological structure underlying neural transmission, as evidenced in the levels of excitation in the (noisy) signal responses.

Mixture model structure arises naturally from the science and experimental context, as is elaborated and developed in Section 2 below. To introduce the basic framework, note that neural transmission stems from individual neurotransmitter release sites on the isolated area of the pre-synaptic neuron; when stimulated, any subset of these sites may release neurotransmitter, and the resulting post-synaptic potential level change induced is the sum of levels corresponding to each of these contributions. Suppose there are s sites, typically fewer than 7 or 8 in current experiments. Transmission at any site occurs with some probability, possibly site-specific. Writing μ_1, \dots, μ_s for the release levels for the s sites, it follows that the induced EPSP signal on any occasion reaches a level in a set $\theta_1, \dots, \theta_k$, say, where each θ_i is a sum of the subset of μ_j values corresponding to sites transmitting. Generally, there are therefore $k = 2^s$ possible levels, and EPSP signal measurements represent selected values from this set plus contributions from the background experimental and synaptic noise. The mixture structure arises as it is impossible to identify which sites transmit on any occasion. As a result, the signals arise from a mixture of $k = 2^s$ component distributions. The mixture components are located at the θ_i , themselves linear functions of subsets of the fundamental site release levels μ_j ; one key objective is inference about these site levels. The probabilities weighting components are functions of fundamental, site-specific probabilities that govern transmitter release on each occasion; these are also uncertain and to be estimated. The shape of each of the mixture components is that of the background noise distribution. As discussed further below, the noise distribution is, throughout this work, assumed to be normal so we have a problem of inference in a mixture of normals with a common component variance, that of the underlying noise distribution.

As we shall see, this basic model will be elaborated considerably by developing structured, hierarchical prior distributions for both the locations and the associated probabilities of the normal components of the mixture; this modelling is based on developments to incorporate qualitative aspects of the underlying science by appropriately describing relationships among the fundamental site-specific release levels and release probabilities. The number s of underlying release sites is also uncertain, so inducing uncertainty about the number k of mixture components; our work addresses this issue through prior modelling and sensitivity analyses. For now, in this introduction, the key point is that the mixture model arises from the scientific context; non-normal features evidenced

in signal data configurations are therefore attributable to the mixture parameters. In our models, we negate the original, negative, excitatory potential changes so that the μ_j are non-negative. Hence EPSP signal data configurations are expected to appear as positively skewed, possibly with multiple modes located near positive component means; the examples in Figures 1 and 2 are typical of current experiments.

We follow most previous works in assuming normally distributed noise. We have two main reasons. First, a primary focus here is to present novel and appropriate models and analyses for the signal recordings, extending previous approaches all based on normal noise distributions. Considering non-normal noise distributions would unduly complicate and cloud the key issues here, i.e., the objective of closer and more appropriate modelling of the underlying scientific structure of the signal data. Second, ranges of existing EPSP data sets, and others being generated by neurosurgeons, do have noise recordings that closely conform to normality assumptions. For the future, it should be noted that some EPSP experiments are quite evidently subject to non-normal errors, and others more subtly so; hence appropriately capturing the noise characteristics is clearly important in some EPSP signal analyses. We are currently developing extensions of the models here to flexibly incorporate non-normal noise distributions, extending work in previous signal models, reported in Cao and West (1996); such extensions will be reported elsewhere.

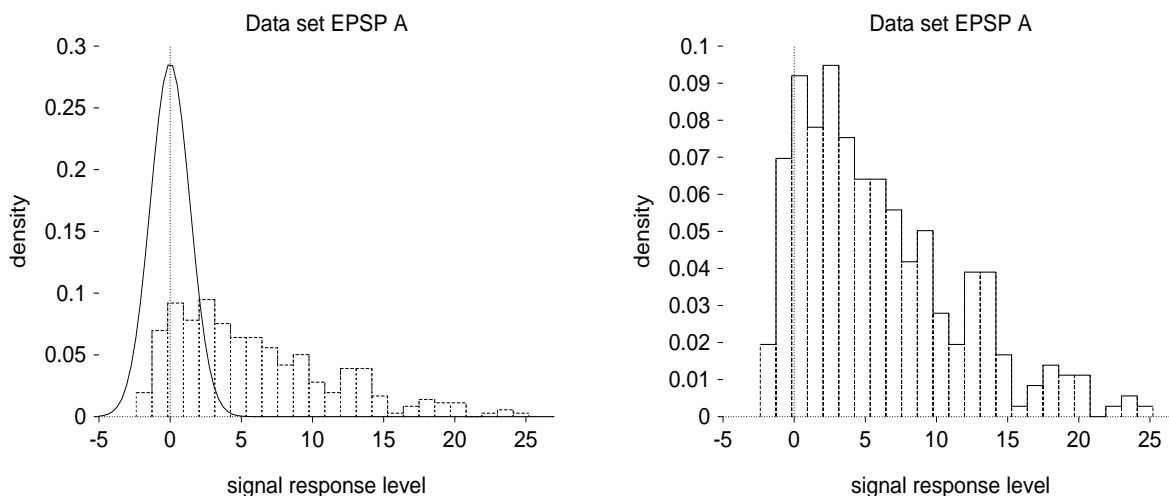


Figure 1. Data from EPSP A experiment. The first frame displays a histogram of the EPSP signal data with the estimated normal noise pdf superimposed on the same scale. For clarity, the signal data is redisplayed in the second frame, indicating signal structure consistent with several sites.

A first example EPSP data set appears in Figure 1. The range and shape of the signal histogram relative to the noise indicates an overall mixture of quite a few components with appreciable probabilities, consistent with more than one active release site. This is a typical data configuration and a typical sample size of $n = 325$.

Two additional example EPSP data sets appear in Figure 2. In each case, the normal pdf determined by the sample moments of the corresponding noise measurements is superimposed on the signal histogram. These were obtained in a paired-pulse experiment performed by Dr Howard

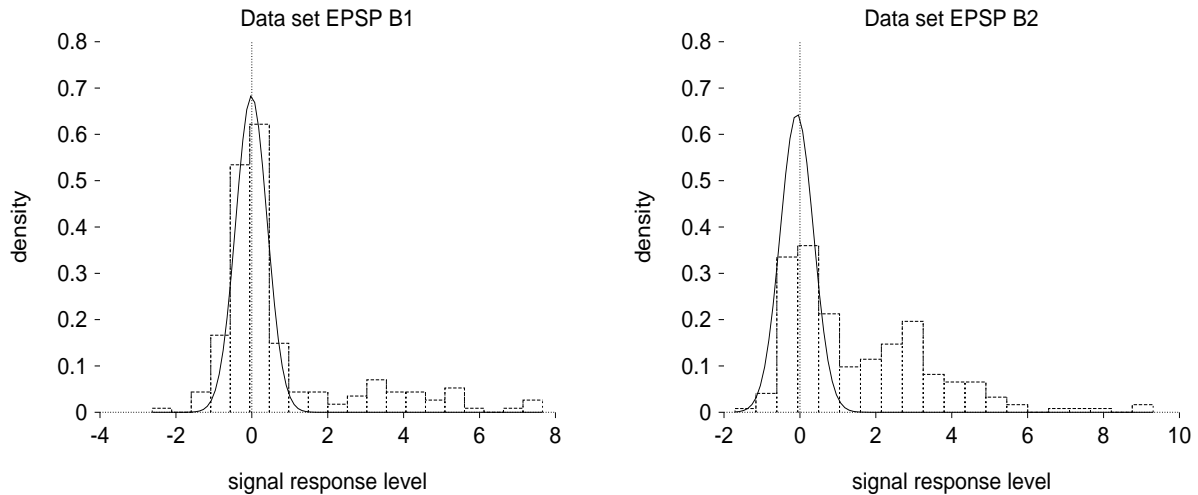


Figure 2. Histograms of neural signal data from two EPSP data sets obtained in a paired-pulse experiment on the same synaptic tissue. Each data set has $n = 222$ signal observations. The superimposed curves are normal densities representing the background noise underlying the signals, fitted by matching the mean and variance of the concomitant noise sample, in each of the two cases.

Wheal of Southampton University (Isaac, Turner and Wheal, 1994). A paired-pulse experiment is designed to investigate changes or adaptations in the neural response characteristics at a single synaptic junction due to previous potentiation (e.g., Kullman and Nicoll 1992). In this case, the tissue is stimulated, as usual, to generate the first EPSP data set, both noise and signal recordings; this is EPSP B1 in the Figure. Then, following potentiation, the experiment is rapidly repeated, generating EPSP B2 in this case. Inference about differences in release levels, and in the stochastic nature of neurotransmitter release more generally, are relevant to physiological theories of adaptation and evolution of nervous systems, and of particular interest in connection with growth (or deterioration) in function in mammalian nervous systems and brains (Kullman and Nicoll 1992). So it is of interest here to compare inferences made for the two cases, asking questions about how the stochastic mechanisms governing transmission may have changed between the two cases as a result of recent potentiation. The data sets displayed are atypical in the small sample sizes, which will evidently be a limiting factor in inferences about site parameters, but typical of recent paired pulse experiments of Wheal which are designed to stimulate rather few neural sites with low release probabilities.

Section 2 describes the basic modelling framework introduced in this paper, and interesting special cases. This includes detailed and quite intricate structure for prior distributions for parameters representing individual transmission site characteristics. Section 3 describes Bayesian analyses via stochastic simulation, discussing various computational issues and critical parameter identification questions; the appendix contains summary technical details for the simulation analysis. Section 4 summarises some analyses of the EPSP A data, exploring modelling and inference issues of general concern as well as discussing the specific neural data set. Section 5 discusses analyses of the paired pulse data sets with more of a focus on the scientific issues. Section 6 explores a scientifically very important, and recent, model extension that introduces significant complications into the analysis,

and raises some issues of analysis in the context of the EPSP B2 data. Section 7 concludes with some summary comments and indications of current and expected modelling developments.

2 MODELS FOR SITE SPECIFIC RELEASE CHARACTERISTICS

2.1 Neural Signal and Noise Models

Replicate measures of maximum EPSP signal levels are assumed independently generated according to an underlying stochastic, synaptic site-specific mechanism described as follows. At the identified synaptic connection, the experiment isolates a very small region on the axon branch of a pre-synaptic neuron, and stimulus of the neuron leads to chemical neurotransmitter release at some or all of a few minute synaptic transmission sites. The sum across sites of the neurotransmitter released on any occasion induces the potential change in the receiving neuron, and this is measured, subject to noise. Each of the $s > 0$ release sites transmits to a site-specific maximum level independently of all other sites. There is no identification of individual sites (e.g., such as might arise were we able to physically mark sites in the synaptic tissue, and identify some individual transmissions) so we arbitrarily label the sites $1, \dots, s$. Site by site we assume that, on any trial,

- sites ‘fire’ independently, with individual, site-specific chances of firing on any occasion;
- a transmitting site produces a site-specific, fixed packet or ‘quantum’ of neurotransmitter;
- recorded maximum EPSP levels represent the sums of potentials induced by the sites transmitting, with additive synaptic and experimental noise.

In any experiment we assume individual transmitter release probabilities and transmission levels to be fixed. Changes in these characteristics over time due to forced experimental and environmental changes is of later interest, and evaluating such changes is one of the guiding motivations in developing these models; see Section 5 for illustration. It is stressed that these assumptions incorporate and relax standard, consensus assumptions about neural function underlying EPSP activity. Symbolically, we represent recorded EPSP levels as $y = \{y_1, \dots, y_n\}$ where y_i is the level on trial i ; the sample size n is typically a few hundreds. Given s sites assumed, site j transmits on any trial with probability π_j , independently of other sites and independently across trials. The level of transmission on firing is the site-specific quantity μ_j . Physically, EPSP levels are measured via intra-cellular probes as electrical potentials induced on cell bodies, and responses are in negative potentials in units of millivolts (mV). By convention, we negate the readings so that the μ_j must be non-negative values. The recordings are subject to background synaptic and experimental noise, including both electronic distortions in recording the cellular potential and errors induced in computing the estimated maximum potential levels. These noise sources combine to produce additive errors, assumed to be approximately normally distributed independently across trials, as discussed in the introduction; write v for the variance of these normal errors. Concomitant noise measurements are available to assess this and to provide prior information relevant to v ; differences in cell body potentials measured without experimental stimulus provide these measurements. It is stressed that the noise affecting signal recordings has the same distribution, of variance v , as the raw noise recordings, there being no scientific or experimental reason to assume otherwise. We also model a systematic bias in the signal recordings; although the associated noise measures are typically close to zero mean, the signal recordings are often subject to a small but non-negligible

potential shift induced by the experimental procedures; call this shift m . We then have the following basic signal model.

For $i = 1, \dots, n$, the signal observations are conditionally independent,

$$y_i \sim N(y_i|\theta_i, v) \quad \text{with} \quad \theta_i = m + \sum_{j=1}^s z_{ij}\mu_j$$

and where the z_{ij} are latent, unobserved, site transmission indicators,

$$z_{ij} = \begin{cases} 1, & \text{if site } j \text{ transmits on trial } i; \\ 0, & \text{otherwise.} \end{cases}$$

Under the model assumptions, these indicators are conditionally independent Bernoulli quantities with $P(z_{ij} = 1) = \pi_j$ for each site j and all trials i .

Thus EPSP data are drawn from a discrete mixture of (at most) $k = 2^s$ normal components of common variance, induced by averaging normals determined by all possible combinations of the z_{ij} and weighted by the corresponding chances of those combinations. On any trial, the normal component selected is determined by the column s -vector $z_i = (z_{i1}, \dots, z_{is})'$ realised on that trial by individual sites firing or not firing; the selected component has mean $\theta_i = m + z_i'\mu$ where $\mu' = (\mu_1, \dots, \mu_s)$, and is selected with chance $\prod_{j=1}^s \pi_j^{z_{ij}} (1 - \pi_j)^{1-z_{ij}}$; specifically, and conditional on all model assumptions and parameters, the y_i are conditionally independently drawn from the distribution with density

$$p(y_i|\mu, \pi, m, v) = \sum_{z_i} p(z_i|\pi) p(y_i|\mu, z_i, m, v) = \sum_{z_i} \left\{ \prod_{j=1}^s \pi_j^{z_{ij}} (1 - \pi_j)^{1-z_{ij}} \right\} N(y_i|m + z_i'\mu, v), \quad (1)$$

where the s -vector z_i ranges over all 2^s possible values. If there are common values among the site levels μ_j , then the mixture distribution function will be equivalent to one with fewer than 2^s components. Also, if either $\mu_j = 0$ or $\pi_j = 0$ then site j disappears from the model and the distribution reduces to a mixture of at most 2^{s-1} components. Additional cases of special interest include:

- a) So-called compound binomial models, in which the site response levels are equal, all at a basic 'quantum' level μ_0 , but release probabilities differ. The aggregate response levels then run between m and $m + s\mu_0$ and the mixture of 2^s normal components effectively reduces to a mixture of just $s + 1$, with associated weights determined by the discrete compound binomial resulting from the distinct chances π_1, \dots, π_s across sites. These kinds of special cases have received considerable attention in the literature (Walmsley, Edwards and Tracey 1988; Kullman 1992).
- b) Precise quantal-binomial models, in which the site levels are equal and the release probabilities are constant too, $\pi_j = \pi_0$. Now the mixture reduces to one with distinct normal means $m + j\mu_0$ ($j = 0, \dots, s$), and corresponding binomial weights $\binom{s}{j} \pi_0^j (1 - \pi_0)^{s-j}$ (Martin 1966; Redman 1990; Clements 1991).

Our previous work, following others, modelled EPSP data in the framework of 'standard' mixture models for density estimation, from a particular Bayesian viewpoint; see West and Cao (1993), West

and Turner (1994), for example, and Escobar and West (1995) for statistical background. In these models, the component means and weights are essentially unrestricted, rather than being modelled directly as functions of the underlying synaptic parameters μ and π . A significant drawback is that it is then, in general, very difficult to translate posterior inferences about unrestricted normal mixture model parameters to the underlying μ and π parameters of scientific interest, especially in the context of uncertainty about s . One simple example in West and Turner (1994) shows how this can be done; that is a rare example in which the data appear to be consistent with the full quantal hypothesis, so that the convoluted process of backtracking from the mixture model to the underlying site-specific parameters is accessible. However, we have encountered very few data sets in which this is the case, and this inversion process is generally difficult. Hence the direct approach, inferring the neural parameters directly, is developed.

2.2 Prior distributions for synaptic parameters

Model completion requires specification of classes of prior distributions for the determining parameters μ , π , m and v for any given s . Assessment of reasonable values of s , as well as values of the μ and π quantities for any given s , is part of the statistical inference problem. From a technical viewpoint, uncertainty about s can be formally included in the prior distribution, so as to provide posterior assessment of plausible values. The models below allow this. From the viewpoint of scientific interpretation, however, inference about the site parameters are best made conditional on posited values of s , and then issues of sensitivity to the number of sites explored. Development and examples explore these issues in detail.

Begin by conditioning on a supposed number of sites s . Current implementation assumes priors for all model parameters which, though rather intricately structured in certain dimensions to represent key qualitative features of the the scientific context, are nevertheless inherently ‘uniform’ in appropriate ways, providing ‘reference’ initial distributions. Note that other classes of priors may be used, and some obvious variations are mentioned below; however, those used here are believed to be appropriately vague or uninformative in order that resulting posterior distributions provide benchmark or reference inferences that may be directly compared with supposedly objective non-Bayesian approaches (e.g., Kullman 1989, 1992; comparative analyses will be reported elsewhere in a further article).

In addressing uncertainty about s alone, we make the observation that models with fewer than s sites are implicitly nested within a model having s sites. To see this, constrain some $r > 0$ (arbitrarily labelled) site levels μ_j to be zero; then (1) reduces to precisely the same form based on $s - r$ sites, whatever the value of the π_j corresponding to the zeroed μ_j . Zeroing some μ_j simply confounds those sites with the noise component in the mixture—a site transmitting a zero level, with any probability, simply cannot be identified. Hence assessing whether or not one or more of the site levels are zero provides assessment of whether or not the data support fewer sites than assumed. This provides a natural approach to inference on the number of active sites, assuming that the model value s is chosen as an upper bound.

Note that a similar conclusion arises by considering $\pi_j = 0$ for some indices j ; i.e., an inactive site may have a zero release probability rather than (or as well as) a zero release level. However, the zeroing of a site probability induces a degeneracy in structure of the model and obviates its use as a technical device for inducing a nesting of models with fewer than s sites in the overall model.

Hence our models define inactive sites through zeros among the μ_j , restricting the π_j to non-zero values, however small. Thus $\mu_j = 0$ for one or more sites j is the one and only way that the number of active sites may be smaller than the specific s .

2.2.1 General structure

For a specified s , analyses reported below are based on priors with the following structure.

Quantities μ , π , (m, v) are mutually independent. We will describe classes of marginal priors for μ and π separately; each will involve certain hyper-parameters that will themselves be subject to uncertainty described through hyper-priors, in a typical hierarchical modelling framework, and these hyper-parameters will also be assumed mutually independent. To anticipate development below, hyper-parameters denoted q and a are associated with the prior for μ , a single quantity b determines the prior for π , and the joint prior is then of the form

$$p(\mu, \pi, m, v, q, a, b) = p(\mu, \pi, m, v | q, a, b) p(q, a, b) = p(\mu | q, a) p(\pi | b) p(m, v) p(q) p(a) p(b). \quad (2)$$

The component densities here are now described in detail. Two comments on notation: first, conditioning statements in density functions include only those quantities that are required to determine the density, implicitly indicating conditional independence of omitted quantities; second, for any vector of h quantities $x = (x_1, \dots, x_h)$, for any $j \leq h$ the notation x_{-j} represents the vector x with x_j removed, i.e., $x_{-j} = x - \{x_j\} = (x_1, \dots, x_{j-1}, x_{j+1}, \dots, x_h)$.

2.2.2 Priors for μ and associated hyper-parameters q, a

We develop a general classes of priors for μ , and comment on various special cases. The class has the following features:

- (a) a component baseline uniform distribution for each μ_j over a prespecified range $(0, u)$, with a specified upper bound u ;
- (b) components introducing positive prior probabilities at zero for each of the μ_j in order to permit assessment of hypotheses that fewer than the chosen (upper bound) s are actually non-zero, and hence to infer values of the number of active sites;
- (c) components permitting exact common values among the elements of μ to allow for the various special cases of quantal transmission, and specifically the questions of whether or not pairs or subsets of sites share essentially the same ‘quantal’ transmission level.

These are developed as follows. Let $F(\cdot)$ be distribution on $(0, u)$, having density $f(\cdot)$, for some specified upper bound u . Write $\delta_0(x)$ for the Dirac delta function at $x = 0$, and $U(\cdot | a, b)$ for the continuous uniform density over (a, b) . Then suppose the μ_j are *conditionally independently* drawn from the model

$$(\mu_j | F, q) \sim q \delta_0(\mu_j) + (1 - q) f(\mu_j)$$

for some probability q . Under this prior, the number h of non-zero values among the μ_j —the number of active sites—is binomial, $(h | s, q) \sim Bn(s, 1 - q)$, with mean $s(1 - q)$, independently of F . If F were uniform, say $U(\cdot | 0, u)$, this prior neatly embodies the first two desirable features (a) and (b) described above. Note the two distinct cases:

- Setting $q = 0$ implies that $h = s$ is the assumed number of active sites, so then inference proceeds conditional on $\mu_j > 0$ for each j .

- Otherwise, restricting to $q > 0$ allows for assessment of the number of active sites, subject to the specified upper bound s . In practice we will assign a hyperprior to q in the case $q > 0$. The class of beta distributions is conditionally conjugate, and the uniform prior suggests itself as a reference, $q \sim U(q|0, 1)$. One immediate, and nice, consequence of a uniform prior is that the resulting prior for the number of non-zero values among the μ_j has (averaging the binomial with respect to q) a discrete uniform prior over $0, 1, \dots, s$. This is a suitably vague and unbiased initial viewpoint with respect to the number of active sites. Other beta priors may be explored, of course. One specific choice we use in current work is $(q|s) \sim Be(s-1, 1)$; note the explicit recognition of dependence on the specified values of s . The reasoning behind this choice is as follows. First, we are currently focused on EPSP experiments designed isolate rather small numbers of sites, down to just a few, say 1 – 4 from the viewpoint of scientific intent and expectation. So s values up to 7 or 8 may be explored, but lower values are typical. Whatever value of s is chosen, h is expected to be in the low integers, so guiding choice of the prior for q to induce a prior for h favouring smaller values. Consider values of s in the relevant range $3 \leq s \leq 8$, or so. Then, integrating $p(h|s, q)$ with respect to the specific prior $(q|s) \sim Be(s-1, 1)$ we obtain a distribution $p(h|s)$ that is almost completely insensitive to s , having a diffuse and decreasing form as h increases, and with $E(h|s) = 1$ for any such s . Thus, this specific beta prior for q has the attractive feature of consistency with a scientifically plausible prior $p(h|s) \approx p(h)$, incorporating the scientific view of likely small numbers of active sites, and almost independently of the upper bound s specified.

This structure provides a baseline uniform prior for release levels, together with the option for allowing a smaller number of sites than the s specified. So far, however, there is no explicit recognition of the special status, in the scientific area, of quantal hypotheses as represented through common values among the μ_j . If $F(\cdot)$ is a continuous distribution, then the prior implies that the non-zero μ_j are distinct. Though this might allow arbitrarily close μ_j values, it is desirable to have the opportunity to directly assess questions about common values, and perhaps subgroups of common values, in terms of posterior probabilities. There is also a significant technical reason, discussed in Section 3 below in connection with issues of parameter identification, that calls for a prior that gives positive probability to exact equality of collections of the non-zero μ_j . We therefore extend the prior structure so far discussed to provide this. We do this using a standard Dirichlet model, fully described in Appendix A. Specifically, a Dirichlet process prior for F induces a discrete structure which gives positive prior probability to essentially arbitrary groupings of the set of non-zero μ_j into subsets of common values; this is a general framework that permits varying degrees of partial quantal structure, from the one extreme of completely distinct values to the other of one common value. This structure is most easily appreciated through the resulting set of complete conditional posterior distributions for each of the μ_j given μ_{-j} . As detailed in Appendix A, these are defined by

$$p(\mu_j|\mu_{-j}, q, a) = q\delta_0(\mu_j) + (1-q)\{r_j U(\mu_j|0, u) + (1-r_j)h_j^{-1} \sum_{i \in N_j} \delta_{\mu_i}(\mu_j)\}, \quad (3)$$

where h_j is the number of non-zero elements of μ_{-j} and N_j is the corresponding set of indices, $N_j = \{i|\mu_i > 0, i = 1, \dots, s; i \neq j\}$, and $r_j = a/(a + h_j)$. The hyper-parameter a is subject to uncertainty and is included in the analysis using existing approaches for inference on precision

parameters in Dirichlet models, developed in West (1992) and illustrated in Escobar and West (1995). As shown there, gamma priors for a , or mixtures of gamma priors, are natural choices, and our application currently uses diffuse gammas models.

In summary, (3) shows explicitly how site level μ_j may be zero, implying an inactive site, or take a new non-zero value, or be equal to one of the non-zero values of other sites. The roles of hyperparameters q and a are evident in this equation. With this structure we have defined prior components $p(\mu|q, a)p(q)p(a)$ of the full joint prior in equation (2).

2.2.3 Priors for π and associated hyperparameter b

The structure of the prior for π parallels, in part, that of μ in allowing common values. The detailed development for μ can be followed through with the same reasoning about a baseline prior and a structure to induce positive probabilities over subsets of common values. We restrict to positive release probabilities, as discussed earlier, so that the prior structure for the π_j is simpler in this respect. Specifically, we assume the π_j independently drawn from a distribution on $(0, 1)$ that is assigned a Dirichlet process prior. We take Dirichlet precision $b > 0$ and base measure to be $bU(\cdot|0, 1)$. Then, as in the development for μ the full joint prior $p(\pi|b)$ is defined by its conditionals

$$(\pi_j|\pi_{-j}, b) \sim wU(\pi_j|0, 1) + (1 - w)(s - 1)^{-1} \sum_{i=1, i \neq j}^s \delta_{\pi_i}(\pi_j) \quad (4)$$

where $w = b/(b + s - 1)$, for each $j = 1, \dots, s$. As for the site levels, the induced posteriors will now allow inference on which sites may have common release probabilities, as well as on the precise values of such probabilities.

Non-uniform beta distributions might replace the uniform baseline in application, and are particularly relevant in analyses of more recent data sets that are generated with a view to engendering rather low levels of excitation consistent with expected low levels of release probabilities. Our analyses reported below retain the uniform prior, for illustration and so as to avoid question of overtly biasing towards lower or higher values. Note also that we explicitly exclude the possibility of non-zero release probabilities, though some or all may be very small, and so rule out the use of $\pi_j = 0$ as a device for reducing the number of active sites; that is completely determined by zero values among the μ_j , as discussed earlier.

As with a in the model for the site levels, the hyper-parameter b will be assigned a prior, and again a diffuse gamma prior is natural. With this structure, we have defined the additional prior components $p(\pi|b)p(b)$ of the full joint prior in equation (2).

2.2.4 Priors for noise moments m and v

In most experiments we anticipate a possible systematic bias m in both noise and signal recordings, induced by the direct measurement of cell membrane potential. This is typically very small relative to induced EPSP levels, and the raw noise recordings provide data on which to assess this. In addition, the noise measurements inform on background variability v , i.e., they are drawn from the $N(\cdot|m, v)$ distribution. In our analyses, our prior for m and v as input to the signal analysis is simply the posterior from a reference analysis of the noise data alone, i.e., a standard conjugate normal, inverse gamma distribution based on the noise sample mean, variance and sample size.

With this structure, we have defined the final component $p(v, m)$ of the full joint prior in equation (2).

3 POSTERIOR DISTRIBUTIONS AND COMPUTATION

Calculation of posterior distributions is feasible, as might be expected, via variants of Gibbs sampling (e.g., Gelfand and Smith 1990, Smith and Roberts 1993). The specific collection and sequence of conditional posterior distributions used to develop simulation algorithms are described in some mathematical detail in Appendix B, and briefly summarised here.

A key step is to augment the data with the latent site transmission indicators z ; thus the sampling model is expanded as

$$\begin{aligned} p(y, z|\mu, \pi, m, v, q, a, b) &= p(y|\mu, z, m, v)p(z|\pi) \\ &= \left\{ \prod_{i=1}^n N(y_i|m + z'_i\mu, v) \right\} \left\{ \prod_{j=1}^s \prod_{i=1}^n \pi_j^{z_{ij}} (1 - \pi_j)^{1-z_{ij}} \right\}, \end{aligned}$$

providing various conditional likelihood functions that are neatly factorised into simple components. The full joint posterior density, for all model parameters together with the uncertain indicators z , has the product form

$$p(\mu, \pi, z, m, v, q, a, b|y) \propto p(\mu|q, a)p(q)p(a)p(z|\pi)p(\pi|b)p(b)p(m, v)p(y|\mu, z, m, v) \quad (5)$$

where the component conditional prior terms are as described in the previous section.

3.1 Conditional posteriors and MCMC

The posterior (5) yields a tractable set of complete conditional distributions characterising the joint posterior, and hence leads to implementation of Gibbs sampling based on this structure. Each iteration of this MCMC sampling scheme draws a new set of all parameters and latent variables by sequencing through the conditionals now noted; Appendix B gives full details and derivations.

- (1) Sampling site levels μ proceeds by sequencing through $j = 1, \dots, s$, at each step generating a new value of μ_j given the latest sampled values of μ_{-j} and all other conditioning quantities. For each j , this involves sampling a posterior that is a simple mixture of several point masses with a truncated normal distribution. Sampling efficiency is improved using variations of so-called configuration sampling for the discrete components, as introduced by MacEachern (1994), and developed in West, Müller and Escobar (1994).
- (2) Sampling site release probabilities π similarly proceeds by sequencing through $j = 1, \dots, s$, at each step generating a new value of π_j given the latest sampled values of π_{-j} and all other conditioning quantities. For each j , this involves sampling a mixture of discrete components with a beta component. Again, configuration sampling improves simulation efficiency.
- (3) Sampling indicators transmission z involves a set of n independent draws from conditional multinomial posteriors for the n individual binary 2^s -vectors z_i , $i = 1, \dots, n$. These simulations are easily performed.

- (4) Sampling the systematic bias quantity m involves a simple normal draw.
- (5) Sampling the noise variance v involves sampling an inverse gamma posterior.
- (6) Sampling the hyper-parameter q is also trivial, simply involving a draw from an appropriate beta distribution.
- (7) Sampling the hyper-parameters a and b follows West (1992), and Escobar and West (1995), and involves a minor ‘augmentation’ of the parameter space, with simple beta and gamma variate generations.

Iterating this process produces sequences of simulated values of the full set of quantities $\phi \stackrel{\text{def}}{=} \{\mu, \pi, z, m, v, q, a, b\}$ which represent realisations of a Markov chain in ϕ space whose stationary distribution is the joint posterior (5) characterised by the conditionals. The precise sequence (1)–(7) of conditional posteriors detailed above (and more fully in Appendix B) is the sequence currently used in simulation; at each step, new values of the quantities in question are sampled based on current values of all quantities required to determine the conditional in question. Convergence is assured by appeal to rather general results, such as Tierney (1994, especially Theorem 1 and Corollary 2) relevant to the models here; this ensures that successive realisations of $\phi = \{\mu, \pi, z, v, q, a, b\}$ generated by this Gibbs sampling set-up eventually resemble samples from the exact posterior in (5). Irreducibility of the resulting chain is a consequence of the fact that all full conditionals defining the chain are everywhere positive. Practical issues of convergence are addressed by experiments with several short runs from various starting values. Important practical issues of convergence is discussed in the next section and in the EPSP data analyses to follow.

3.2 Parameter Identification and Relabelling Transmission Sites

The model structure as described is not identifiable, in the traditional sense of parameter identification especially associated with mixture models (Titterton, Smith and Makov, 1985, chapter 3). This is obvious in view of the lack of any physical identification of the neural transmission sites we have arbitrarily labelled $j = 1, \dots, s$; for example, π_1 is the release probability for the first site in our labelling order, but this could be any of the actual release sites. As a result, the model is invariant under arbitrary permutations of the labels on sites, and likelihood functions are similarly invariant under permutations of the indices of the sets of parameters (μ_j, π_j) . Also, as the priors for the μ_j and π_j separately are exchangeable, corresponding posteriors are also invariant under labelling permutations.

To impose identification we need a physically meaningful restriction on the parameters. Perhaps the simplest, and most obvious, constraint is to order either site levels or site release chances. We can do this with the site release levels. In posterior inferences, relabel so that site j corresponds to release level μ_j , the j^{th} largest value in μ . As the priors, hence posteriors, include positive probabilities on common values, we break ties by imposing a subsidiary ordering on the release probabilities for (ordered) sites with common release levels. Thus, if μ_3, μ_4 and μ_5 happened to be equal, the sites relabelled 3–5 would be so chosen in order of increasing values of their chances π_j ; were these equal too, as permitted under the priors and hence posteriors here, then the labelling

is arbitrary and we have (in this example) three indistinguishable sites: common release levels and common release probabilities. The posterior analysis required to compute corresponding posterior inferences is extremely simple in the context of posterior simulations; each iteration of the Gibbs sampling analysis produces a draw of (μ, π) (among other things). Given a draw, we first create a separate vector μ^* containing the *ordered* values of μ , and a second vector π^* containing the elements of π rearranged to correspond to the ordering of μ to μ^* . If subsets of contiguous values in μ^* are common, the corresponding elements in π^* are then rearranged in increasing order themselves. Through iterations, repeat draws of (μ^*, π^*) are saved, so building up samples from the required posterior distribution that incorporates the identification of sites.

Note that we could alternatively identify the model by a primary ordering on the π_j , rather than the μ_j . In data analyses we routinely explore posterior summaries identified each each – i.e., ordered by μ_j and ordered by π_j – as each is only a partial summary of the full posterior. This is relevant particularly in cases when it appears that the data are in conformity, at least partially, with quantal structure. Some discussion follows in data analyses in later sections.

In connection with this ordering issue, the quantal structure of the priors for μ and π is very relevant, as we alluded to earlier. Consider an application in which the data support common, or very close, values among the μ_j . Suppose that the prior is simply uniform, with no chance that consecutive values of the ordered μ_j are equal. Then imposing the ordering for identification results in a distortion of inferences – values we should judge to be close, or equal, are ‘pushed apart’ by the imposed ordering. This undesirable effect is ameliorated by the priors we use, allowing exact equality of neighbouring levels. The same applies to the π_j .

Notice that we retain the original, unidentified parametrisation for model specification and the simulation analysis operates on the unconstrained posteriors. This has a theoretical benefit in allowing easy and direct application of standard convergence results (Tierney 1994). More practically, it has been our experience that imposing an identifying ordering on parameters directly through the prior distribution hinders convergence of the resulting MCMC on the constrained parameter space. This arises naturally as the ordering induces a highly structured dependence between parameters, relative to the unconstrained (albeit unidentified) parameter space. In addition, one rather intriguing practical payoff from operating the MCMC on the unrestricted space is an aid in assessing convergence of the simulation. This is due to the exact symmetries in posterior distributions resulting from permutation invariance, partly evidenced through the fact that marginal posteriors $p(\mu_j, \pi_j | y)$ are the same for all j . Hence summaries of marginal posteriors for, say, the μ_j should be indicative of the same margins for all j ; the same if true of the π_j . However, in many applications convergence will be slow. One reason for this is that, in some models and with some observed signal data sets, the posterior will heavily concentrate in widely separate regions of the parameter space, possibly being multi-modal with well separated modes. Such posteriors are notorious in the Markov chain Monte Carlo literature, and the Gibbs sampling routine may get ‘stuck’ in iterations in just one (or typically several) region of high posterior probability, or around a subset of modes. Some of the examples below illustrate this phenomenon, and there is certainly need for further algorithmic research to provide faster alternatives to the raw Gibbs sampling developed here (possibly based on some of the methods mentioned in Besag and Green 1993, or Smith and Roberts 1993, for example). However, this is not a problem as far as the ordered site levels, and the correspondingly reordered release probabilities, are concerned, due to the symmetries in the unconstrained posterior. This is

because posterior samples from one region of the (μ, π) space may be reflected to other regions by permuting sites indices, while the resulting ordered values remain unchanged. Hence the unconstrained sampling algorithm may converge slowly as far as sampling the posterior for (μ, π) , but the derived sequence for (μ^*, π^*) converges much faster. This is discussed further in data analysis in Section 4.

Due to the lack of identification, actual output streams from MCMC simulations exhibit switching effects, as parameter draws jump between the modes representing the identification issue. For example, in a study supporting two distinct site release levels with values near 1 and 2meV, respectively, the simulated series of the unidentified μ_1 (μ_2) will tend to vary around 1meV (2meV) for a while, then randomly switch to near 2meV (1meV) for a while, and so on. Our experience is that, across several analyses of different EPSP data sets, convergence is generally ‘clean’ in the sense that, following what is usually a rapid initial, burn-in period, the output streams remain stable and apparently stationary between points of switching between these posterior modes. Informal diagnostics based on a small number of short repeat runs from different starting values verify summarised analyses from one final, longer run. In addition, the theoretical convergence diagnostic based on the symmetry of the unidentified posterior is used in final analysis, as illustrated below.

4 ANALYSES OF EPSP A DATA SET

Some summaries of initial analyses of the EPSP A are displayed in Figures 3 and 4. First, three separate analyses were performed using models with $s = 3, 4$ and 5 , respectively; to do this we simply set $q = 0$ in the prior $p(\mu|q, a)$ to fix the number of sites at the specified s value in each case. The prior for the μ_j has $u = 20.0$. Each MCMC simulation was run for 45,000 iterations; summaries of the last 40,000 are reported. Informal assessment of convergence was performed as discussed at the end of Section 3.

The first frame of Figure 3 displays posterior predictive density functions from the three analyses – Bayesian density estimates, of use in graphical assessment of model fit. These are simply computed as Monte Carlo averages of the density function (1), across a range of y values, averaging with respect to the 40,000 posterior draws of all parameters. There are sharp differences between the case $s = 3$ and the others, but $s = 4$ and 5 produce rather similar density estimates. Increasing s beyond 5 makes little difference to the estimated density function. Corresponding informal comparisons of the cumulative distributions plotted against the empirical cdf of the signal data indicate close conformity of the data with cases $s = 4$ or 5 , but apparent discrepancies with models based on less than 4 sites. These graphical diagnostics and comparisons agree with more formal analyses in the model allowing $q > 0$ with s specified as the maximum number of active sites, and using the $Be(s - 1, 1)$ prior for q , discussed earlier (similar results arise with a uniform prior). Approximate posterior probabilities on the number h of active sites are as follows: at $s = 3$, we have $P(h = 3|y, s = 3) \approx 1$; at $s = 4$, we have $P(h = 3|y, s = 4) \approx 0.14$ and $P(h = 4|y, s = 4) \approx 0.86$; at $s = 5$, $P(h = 4|y, s = 5) \approx 0.49$ and $P(h = 5|y, s = 5) \approx 0.51$; finally, at $s = 6$ we have $P(h = 4|y, s = 6) \approx 0.30$, $P(h = 5|y, s = 6) \approx 0.43$ and $P(h = 6|y, s = 6) \approx 0.27$. These summaries reflect the inherent uncertainties in estimating the number of active sites, and the problem of overfitting as higher numbers are permitted. Though the prior for q marginally favours lower values of h , increasing s allows more sites and the model naturally tends to overfit, attempting to tailor itself to increasing minor features in the data as is common in other varieties of mixture models.

On the basis of parsimony, $h = 4$ is identified as a minimally acceptable number of active sites, and $h = 5$ seems relatively favoured. Models with $s = h = 4$ and $s = h = 5$ are discussed further.

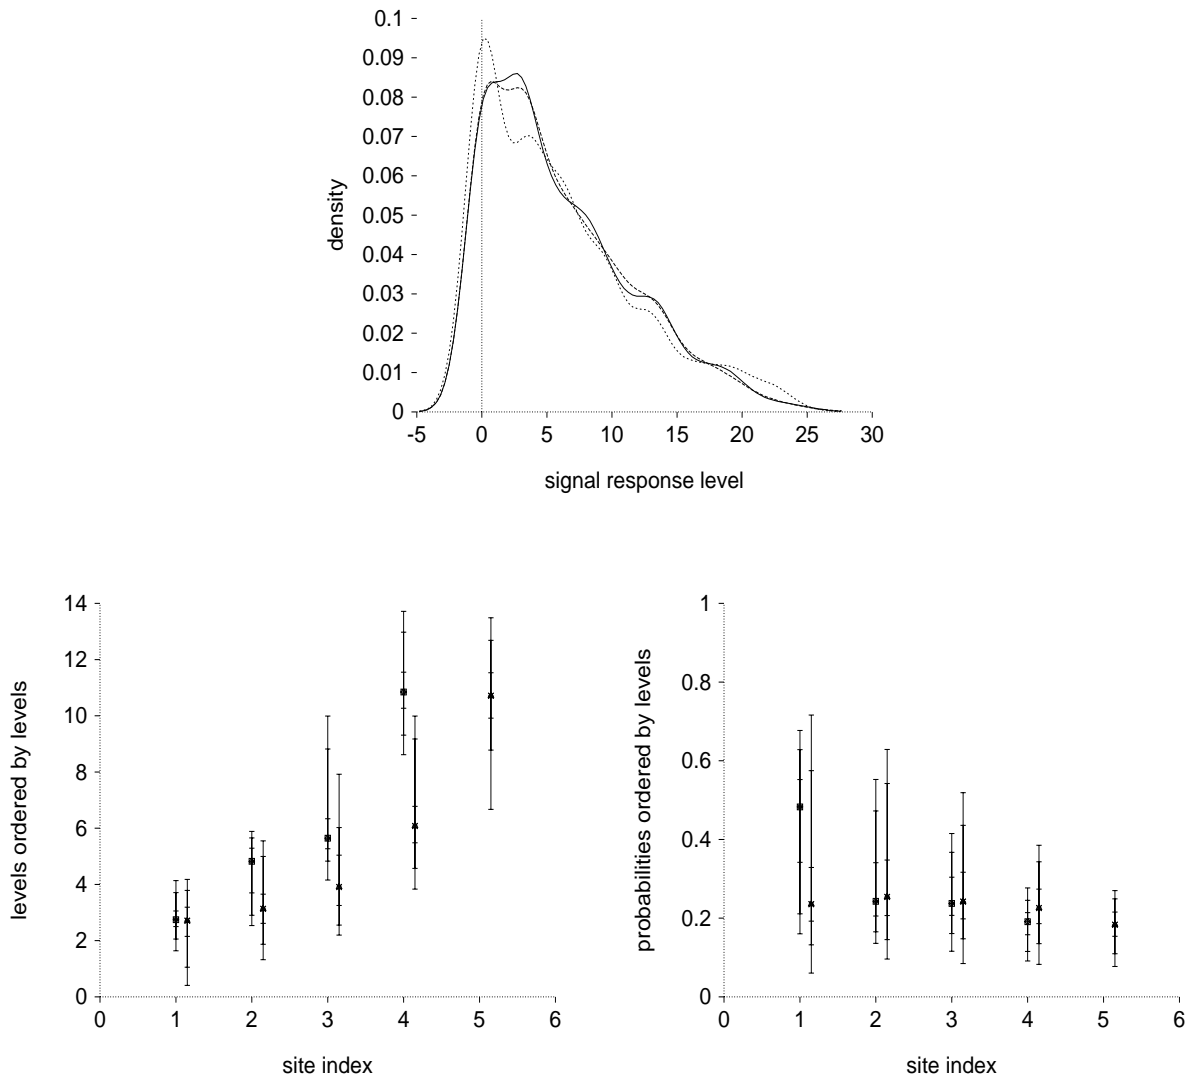


Figure 3. Summaries from analyses of EPSP A data. The first frame displays predictive densities from analyses in models with 3 sites (dotted line), 4 sites (dashed line), and 5 sites (full line). The second frame displays, at each site index j , posterior intervals for the ordered μ_j in analysis with 4 sites (left hand interval) and 5 sites (right hand interval). These are 98% equal tails intervals with the medians, quartiles, 5% and 95% points marked. The third frame displays similar intervals for the π_j , in order of increasing values of the μ_j in each case, for each analysis.

Additional frames in Figure 3 summarise the marginal posteriors for μ and π in each of the analyses with 4 and 5 active sites. The first frame displays approximate 98% equal tails posterior intervals for μ_1, \dots, μ_s , for each of $h = s = 4$ and 5, with the levels ordered for identification; the vertical bars in each case run from the lower 1% to the upper 99% posterior quantile, and have medians, quartiles, 5% and 95% points marked. At site index $j = 1$, for example, the two vertical bars correspond to posterior intervals for μ_1 , the smallest release level, from each of the

two separate analyses. The vertical scale is meV. At $j = 5$ we naturally have only one interval, for μ_5 from the second analysis. There are apparent differences; analysis with just 4 sites naturally favours higher μ_j values at higher levels in order to accommodate the skewed upper tail of the signal data. In each analysis, the graphs indicate differences among the site release levels. This is supported by approximate posterior probabilities on equal/quantal values in the two analyses, which indicate strong support for at least three distinct values among the levels of active sites, in each case; in analysis with $s = h = 4$ we have approximate posteriors $P(h^* = 3|y, h = 4) = 0.64$, $P(h^* = 4|y, h = 4) = 0.32$, whereas in analysis with $s = h = 5$ we have $P(h^* = 3|y, h = 5) = 0.58$, $P(h^* = 4|y, h = 5) = 0.28$ and $P(h^* = 5|y, h = 5) = 0.03$. For release probabilities π_j in order of increasing μ_j , intervals appear in a similar display in the third frame of Figure 3. Now the differences are less apparent, though still evident. In analysis here with $s = h = 4$ we have approximate posteriors $P(k^* = 2|y, h = 4) = 0.53$ and $P(k^* = 3|y, h = 4) = 0.37$, and when $s = 5$ we have $P(k^* = 2|y, h = 5) = 0.45$, $P(k^* = 3|y, h = 5) = 0.30$, for example. This does suggest more disparity among the release levels than among the release chances, though there are differences and considerable uncertainty. What is clear is that a strict quantal structure, with a single common release level and a single common release probability, is quite untenable. This data set is also typical of many in that there is the appearance of more disparity among the release levels than among release probabilities.

Further summaries of the analyses with fixed $s = h = 4$ appear in Figure 4. The predictive density is superimposed on a data histogram, with some indications of uncertainty in terms of pointwise intervals representing ± 1.65 approximate posterior standard deviations about the density estimate. The other two frames display, site by site, approximate 98% intervals, with quantiles marked as in Figure 3, but now for the *unidentified* μ_j and π_j . Recall that the μ_j share a common marginal posterior, as do the π_j . These frames indicate that the MC output summarised is in close conformity to this theoretical exactness, so acting as a confirmatory convergence diagnostic in this case.

Some comments about convergence of the Gibbs sampling algorithms, from a practical perspective, are in order. The predictive density function and posteriors for both ordered release levels and release probabilities displayed in the figure are stable over iterations, and in fact converge quite rapidly; the general features displayed appear after a few thousand, sometimes a few hundred iterations of the simulation algorithm. In this example, as in most real analyses, the true posterior margins are multimodal, reflecting the lack of identification of the unconstrained parameters; in the full parameter space there are multiple modes in the joint posterior. MCMC iterations tend to ‘stick’ around a posterior mode, sometimes for a long while, and this is exacerbated when there are evidently distinct and separated release levels and/or probabilities, so that the modes are widely separated. Though lack of convergence, and lack of identification, is ‘problematic’, it works to advantage here, as discussed earlier. Mapping to the ordered/identified parameters induces physical interpretation and convergence is not such an issue since the regions of unconstrained parameter space ‘explored’ by the simulation are simply reflected under all parameter permutations. Second, the extent to which the margins in the unconstrained parameters differ is related to the extent of real differences in underlying site levels and/or chances; bigger differences induce more widely separated concentrations of posterior mass, and usually posterior modes, resulting in poor convergence for the unconstrained parameters.

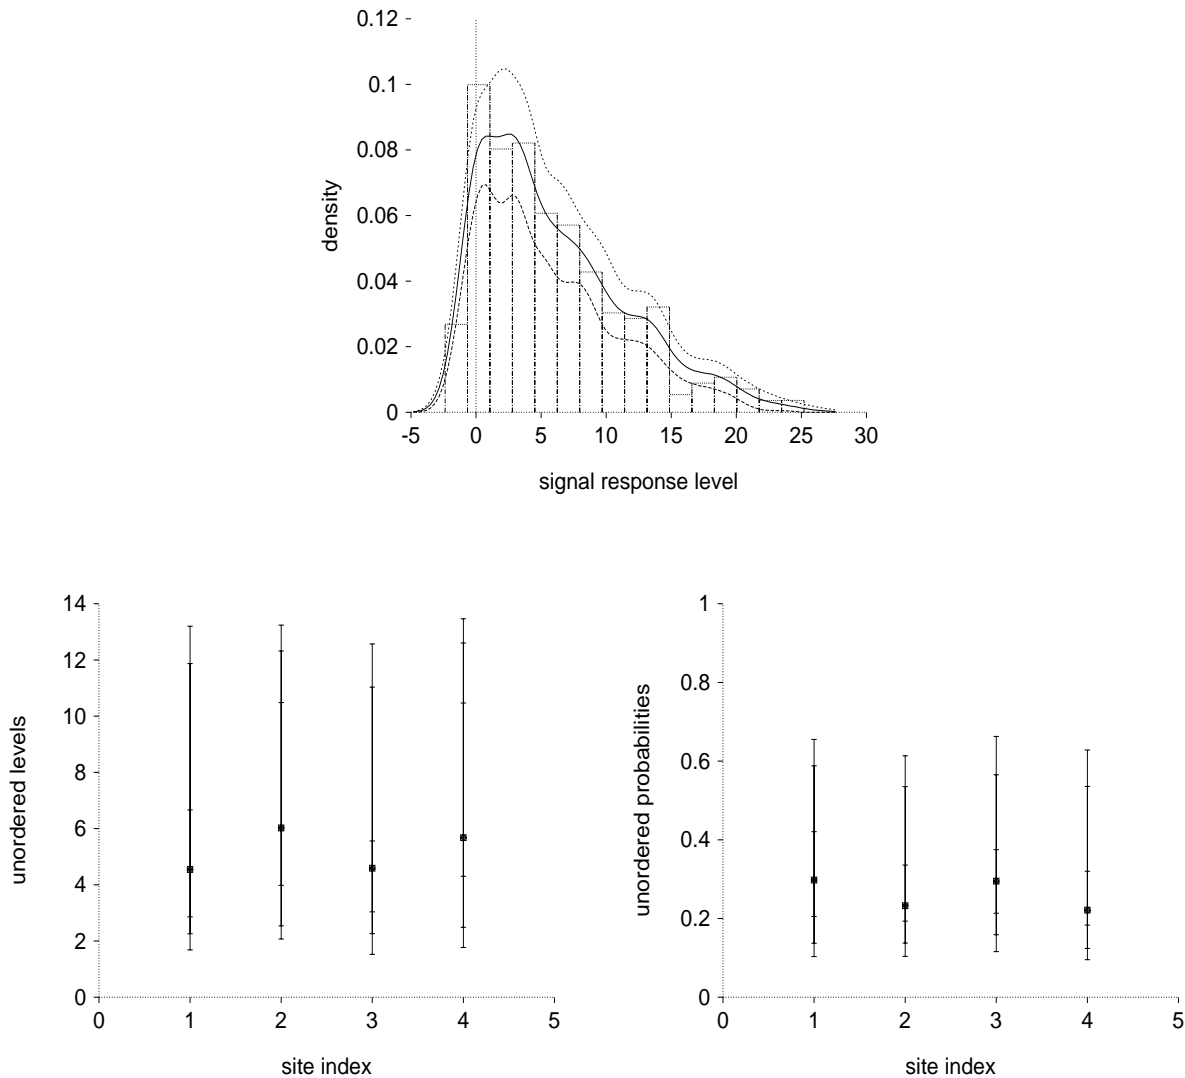


Figure 4. Summaries from analysis of EPSP A data with $s = h = 4$ active sites. The first frame displays the data histogram and the predictive density from the analysis. Uncertainty about the predictive density is crudely indicated via pointwise bands representing ± 1.65 posterior standard deviations about the density estimate. The other two frames display posterior intervals for the *unidentified* release levels μ_j and probabilities π_j , in a format similar to that of Figure 3. See text for discussion.

5 ANALYSIS OF PAIRED PULSE EPSP DATA SETS

Figure 5 provides outputs from an analysis of the first of the paired pulse data sets EPSP B1 in Figure 2, in a format similar to that just described for the EPSP A analyses. In analysis with at most $s = 4$ sites and the $Be(s - 1, 1)$ prior on q , we find approximate probabilities $P(h = 2|y, s = 4) \approx 0.61$, $P(h = 3|y, s = 4) \approx 0.24$ and $P(h = 4|y, s = 4) \approx 0.16$, suggesting just two or, at most, three active sites; the issues of overfitting and parsimony considerations lead us to discount the 0.16 probability on 4 active sites. The graphs display posterior intervals for the release levels and chances under models with fixed numbers of sites $h = s = 2$ and $h = s = 3$; recall that these intervals have medians, quartiles, 1,5,95 and 99% points marked.

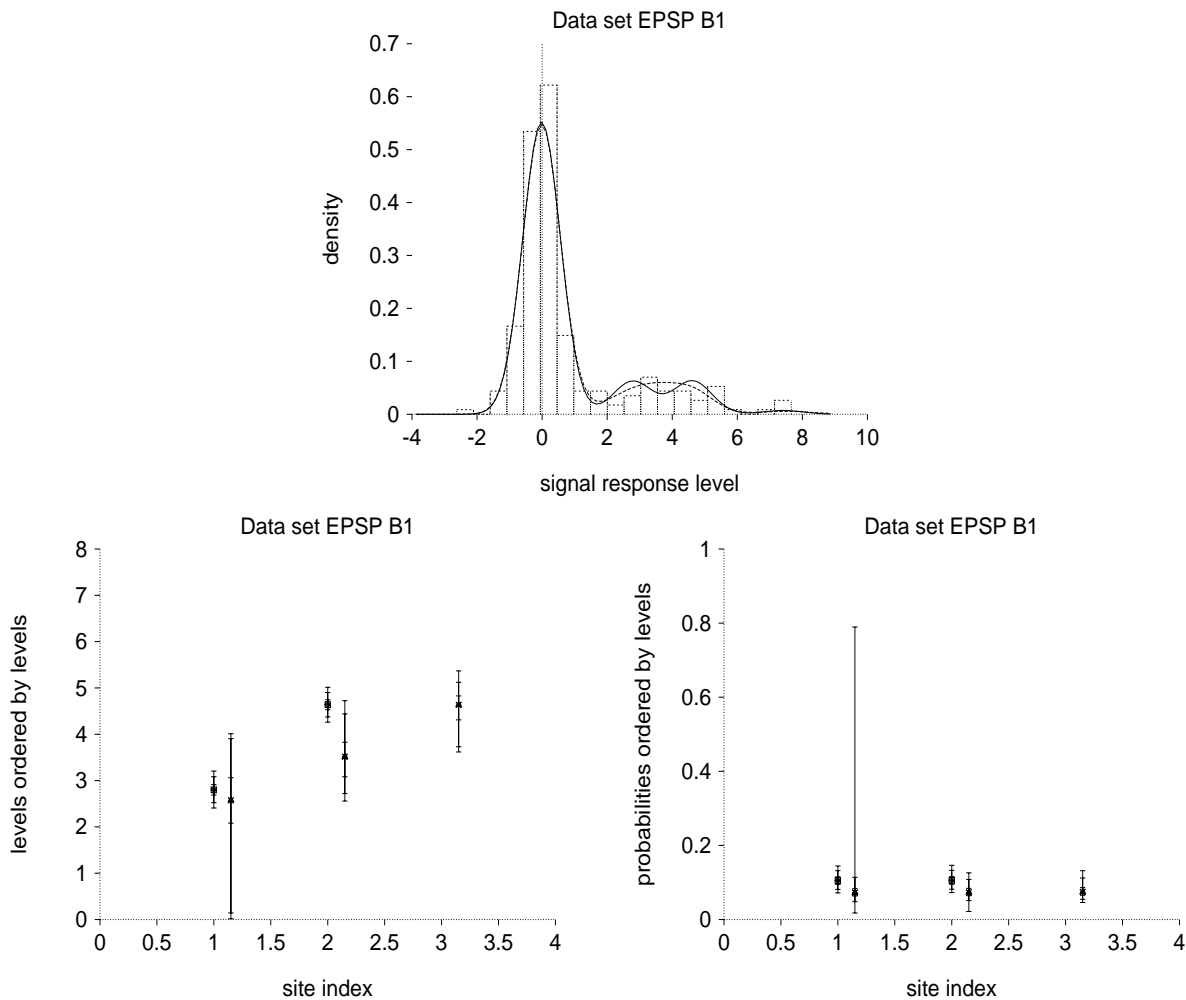


Figure 5. Outputs from analysis of the EPSP B1 data from separate analysis with $h = s = 2$ and $h = s = 3$. The top frame plots the posterior predictive densities from the two analyses together with a histogram of the data. The density curves are from the model with $h = 2$ (full line) and $h = 3$ (dashed line). The other frames display the usual posterior intervals for the μ_j and π_j in order of increasing μ_j . At each site index, the left hand interval corresponds to analysis with $h = s = 2$ sites, the right hand one to $h = s = 3$ sites.

Note the close similarity of predictive density functions, the main difference being that the model with 2 sites leads to more evident minor modes around 3meV and 5meV, respectively. From the graph of intervals for the π_j , ordered by levels, it appears that the two or three sites have rather low release probabilities around 0.07 – 0.12. The data strongly support quantal structure in these probabilities, with $P(k^* = 1|y, h = 2) = 0.94$ and $P(k^* = 1|y, h = 3) = 0.84$. Release levels are likely distinct, with $P(h^* = 2|y, h = 2) \approx 1$ and $P(h^* = 3|y, h = 3) = 0.55$, and with estimated values near 2.5 – 3meV and 4.5 – 5meV for a 2 site model, and around 2 – 3meV, 3 – 4meV and 4 – 5meV for a 3 site model. Note the significant uncertainties about release levels under the 3 site model, partly due to the small sample size. This, coupled with small release probabilities, means that there are very few, though heavily influential, observations really informing about site levels; in particular, the data around 7.5meV in the histogram are influential in determining the sum of

two or three of the μ_j . In the final frame of Figure 5, note the relatively high posterior precision in estimating release probabilities, with the exception of π_1 in the 3 site model. Though much posterior mass for π_1 is concentrated in the 0.07 – 0.12 range, the marginal posterior is diffuse over a range of values consistent with the baseline uniform prior; this reflects the fact that the existence of a third site is dubious, related to the high degree of uncertainty about μ_1 .

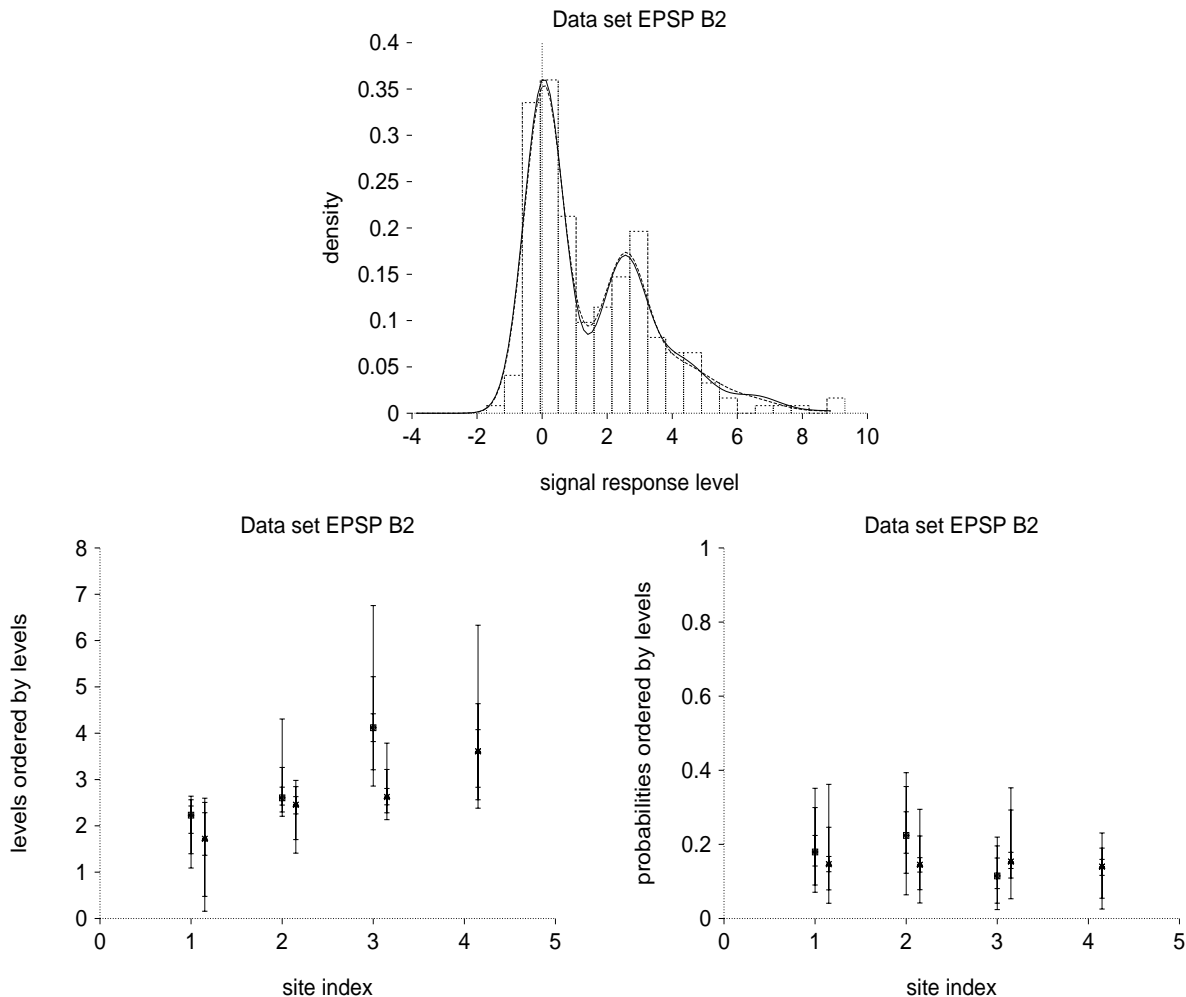


Figure 6. Outputs from analysis of the EPSP B2 data from separate analysis with $h = s = 3$ and $h = s = 4$. The top frame plots the posterior predictive densities from the two analyses together with a histogram of the data. The density curves are from the model with $h = 3$ (full line) and $h = 4$ (dashed line). The other frames display the usual posterior intervals for the μ_j and π_j in order of increasing μ_j . At each site index, the left hand interval corresponds to analysis with $h = s = 3$ sites, the right hand one to $h = s = 4$ sites.

Figure 6 provides outputs from analysis of the EPSP B2 data set, the second of the paired pulse data sets, which are summarised and compared with the inferences for EPSP B1 just described. It is apparent from the signal histogram that, working with the same synaptic tissue, there is more ‘action’ in EPSP B2 than in EPSP B1 in the sense of higher EPSP levels. Neurologists are interested in whether this potentiation-induced activity is due to increased release probabilities or increased transmission levels, or both.

Analysis with $q > 0$ to assess the number of active sites indicates 3 or 4 sites. For a maximum $s = 4$, as in analysis of the first pulse, we obtain approximate posteriors $P(h = 3|y, s = 4) = 0.32$ and $P(h = 4|y, s = 4) = 0.68$. At a maximum of $s = 5$, we obtain $P(h = 3|y, s = 5) = 0.13$, $P(h = 4|y, s = 5) = 0.55$ and $P(h = 5|y, s = 5) = 0.32$. This rather strongly indicates 4 active sites, particularly in view of the over-fitting issue, which is an increased number relative to the first pulse, pre-potentiation. The estimated predictive densities from 3, 4 and 5 site models are almost indistinguishable (see Figure 7). Conditional on the likely 4 active sites, the marginal posterior on the number of distinct release levels have approximate masses of $P(h^* = 2, 3, 4|y, h = s = 4) \approx 0.44, 0.37, 0.15$ respectively; that for the number of distinct release probabilities is $P(k^* = 1, 2, 3|y, h = s = 4) \approx 0.63, 0.29, 0.08$, respectively. In summary, of the likely 4 active sites it appears reasonable to infer that the sites share a common release probability, but differ in release levels. This is, again, a consistent features across a variety of experimental data sets. Release probabilities across sites are small; on a 4 site model, the margins displayed in Figure 6 suggest the π_j lie in the region of 0.1 – 0.2, with posterior medians around 0.14 – 0.16.

Combining and comparing these inferences with the analysis of the EPSP B1 data, we come to the following conclusions about the nature of the synaptic site being explored and the effects of the potentiation experiment. It seems plausible that that there are four transmission sites (n.b., it should be noted that analysis based on $s = 5$ sites supports this conclusion). Initially, sites are activated with a common but very low release probability, around 0.07 – 0.1. In the first pulse, only 2, perhaps (but less likely) 3 sites are active, with differing transmission levels in the 3 – 5meV range; if just 2 sites, the levels are fairly precisely estimated at about 2.5 – 3meV and 4.5meV respectively. Following potentiation, the second pulse data set is suggestive of the following effects. First, the site release probabilities are still constant across sites, but have increased to around 0.14 – 0.16 or higher, perhaps a doubling of the common release probability in the first pulse. This ties in with a leading hypothesis about the physiological effects of potentiation (e.g., Kullman and Nicoll 1992), that, following potentiation, release sites are in a state of more immediate readiness to transmit, and will therefore be more likely to do so at the rapid repeat stimulation in a paired pulse experiment. Indeed, the support for more active sites in the second pulse suggest the existence of 4 sites, simply not evident in the first pulse experiment pre-potentiation. The effects on transmission levels seem to work in the opposite direction. On a 4 site model, the second pulse has release levels that are generally lower than those of the 2 or 3 sites evident in the first pulse. Though a tentative conclusion in view of the high levels of posterior uncertainty about the ordered release levels in the second pulse, this is also consistent with physiological theory suggesting that release sites, having transmitted on one stimulus occasion, require a waiting period in order to transmit again to similar maximum levels, i.e., a period of replenishment.

6 ADDITIONAL VARIANCE COMPONENTS AND A MODEL EXTENSION

Recent attention has focussed on questions of additional variability in EPSP outcomes due to so-called *intrinsic variability* in release levels of individual synapses (Turner *et al* 1995); we describe this concept here, and define an elaborated class of models incorporating it. This concept gives rise to model modifications in which components of the normal mixture representation have variances that increase with level, and this leads to considerable complications, both substantive and technical. The technical complications have to do with developing appropriate model extensions, and

associated MCMC techniques to analyse the resulting models; some brief development is mentioned here, with examples. The substantive complication is essentially that competing models with and without intrinsic variance components cannot be readily distinguished on the basis of the observed EPSP data alone; an observed data configuration might arise from just one or two sites with significant intrinsic variance, or it might arise from a greater number of sites with low or zero intrinsic variance. In such cases, and especially when inferences about site characteristics are heavily dependent on the number of sites and levels of intrinsic variance, we are left reliant on the opinions of expert neurophysiologists to judge between the models. Unfortunately, in its current state, the field is represented by widely varying expert opinions, from the one extreme of complete disregard of the notion of intrinsic variability, to the other of belief in high levels of intrinsic variability as the norm. In some examples we have studied, including the EPSP A data set of this paper, this issue is relatively benign, as inferences about release levels and probabilities are relatively insensitive to intrinsic variances. In others cases, it is evidently highly relevant. Further collaborative research to refine knowledge of intrinsic variability effects is part of the current frontiers of the field. Here we give a short discussion of the notion and our current approach to modelling intrinsic variability.

Intrinsic variability refers to variation in levels of neurotransmitter release at specific sites. As developed above, site j has a fixed release level μ_j , and, while these levels may differ across sites, they are assumed fixed for the duration of the experiment under the controlled conditions. However, the mechanism of electro-chemical transmission suggests that this may be an over-simplification. A site transmits by releasing a packet of (many) molecules of a chemical transmitter, and these molecules move across the synaptic ‘cleft’ to the receiving cell. The induced potential response is proportional to the number of molecules received. So, the assumption of fixed μ_j implies that (a) the number of molecules transmitted is constant across occasions, and (b) all ejected molecules are received by the post-synaptic cell. Each of these is questionable, and the issue has given rise to the notion of intrinsic variance, i.e., variability in the site-specific level of release across occasions. Whatever the actual structure of variability, it is evidenced in EPSP data sets through mixture components having higher variance at higher release levels.

The basic extension of our models to admit the possibility of intrinsic variability involves remodelling the data y_i as coming from the conditional normal model

$$y_i \sim N(y_i | m + \sum_{j=1}^s z_{ij} \gamma_{ij}, v) \quad (6)$$

where the γ_{ij} are new site-, and case-, specific release levels and v is the noise variance. Retaining independence across sites and trials, we assume that the γ_{ij} are distributed about an underlying expected release level μ_j – the same site-specific level as before, but now representing underlying average levels about which releases vary. Then the form of the distribution of the γ_{ij} about μ_j represents the intrinsic variability in induced responses due to variation in amounts of neurotransmitter released by site j and also to variation in the success rate in moving the transmitter across the synaptic cleft. Various parametric forms might be considered; our preliminary work, to date, is based on exploration of models in which

$$p(\gamma_{ij} | \mu_j, \tau_j) \propto N(\gamma_{ij} | \mu_j, \tau_j^2 \mu_j^2) I(0 < \gamma_{ij} < u) \quad (7)$$

for all trials i and each site j ; here u is the earlier specific upper bound on release levels. The new parameters $\tau_j > 0$ measure intrinsic variability of sites $j = 1, \dots, s$; ignoring the truncation in (7), τ_j is an effective constant (though site-specific) coefficient of variation in release levels about the underlying μ_j . This model has been implemented, extending the prior structure detailed in Section 2 to incorporate the full set of site-, and trial-, specific release levels $\{\gamma_{ij}\}$ together with the new parameters τ_1, \dots, τ_s . Note that, as our model allows $\mu_j = 0$ for inactive sites, we use (7) only for $\mu_j > 0$; otherwise, $\mu_j = 0$ implies $\gamma_{ij} = 0$ for each $i = 1, \dots, n$.

We can see the effects of quantal variability by integrating the data density (6) with respect to (7). This is complicated due to the truncation to positive values; as an approximation, for the purposes of illustrating structure here, assume this is not binding, i.e., that μ_j and τ_j are such that the mass of the basic normal distribution in (7) lies well within the interval $(0, u)$. Then (6) and (7) combine and marginalise over γ_{ij} to give

$$y_i \sim N(y_i | m + \sum_{j=1}^s z_{ij} \mu_j, v + \sum_{j=1}^s z_{ij} \tau_j^2 \mu_j^2). \quad (8)$$

The mean here is as in the original formulation, equation (1), with active sites contributing the expected release levels μ_j . But now the variance of y_i is not simply the noise variance v ; it is inflated by adding in factors $\tau_j^2 \mu_j^2$ for each of the active sites. Hence the feature relevant to modelling increased spread of EPSP data configurations at higher response levels.

It should be clear that this extended model can be managed computationally with direct extensions of the MCMC algorithms of Section 3 and Appendix B. We are now interested in the full posterior $p(\mu, \pi, z, \gamma, \tau, m, v, q, a, b | y)$, extending the original posterior in (5) to include the new quantities $\gamma = \{\gamma_{ij}, i = 1, \dots, n; j = 1, \dots, s\}$ and $\tau = \{\tau_1, \dots, \tau_s\}$. There are difficulties in the posterior MCMC analysis due to the complicated form of (7) as a function of μ_j , and also due to the truncation of the basic normal model. These issues destroy part of the nice, conditionally conjugate sampling structure, and are currently handled using some direct analytic approximations and Metropolis-Hasting accept/reject steps in the extended simulation analysis. Details are very briefly reviewed in Appendix C. It is beyond the scope of the current paper to develop the technical aspects of this fully here, but our discussion would be incomplete without raising the issues of intrinsic variability, currently becoming vogue in the field, and without providing some exploratory data analysis. Appendix C briefly details the major technical modifications to the analysis already spelled out in Section 3 and Appendix B. After further technical refinements and practical experience, full modelling and technical details will be reported elsewhere. Here, for completeness, we briefly summarise results of analyses of some of data set EPSP B2 already discussed.

We take the τ_j to be a *priori* independent $\tau_j \sim U(\tau_j | 0, 0.5)$ in our analysis, a maximum 50% level for the intrinsic coefficient of variation. Analysis of the EPSP B2 data with a maximum of $s = 4$ active sites, and $q > 0$ to assess inactive sites, yields approximate probabilities $P(h = 2 | y, s = 4) = 0.19$, $P(h = 3 | y, s = 4) = 0.60$ and $P(h = 4 | y, s = 4) = 0.20$. Compare these with the earlier probabilities in the model with $\tau_j = 0$; that analysis led to a 68% probability of 4 active sites, and a 32% chance of 3 active sites. So, allowing intrinsic variability reduces the most likely number of active sites from 4 to 2. Figure 7 provides some summary graphics from two separate analysis

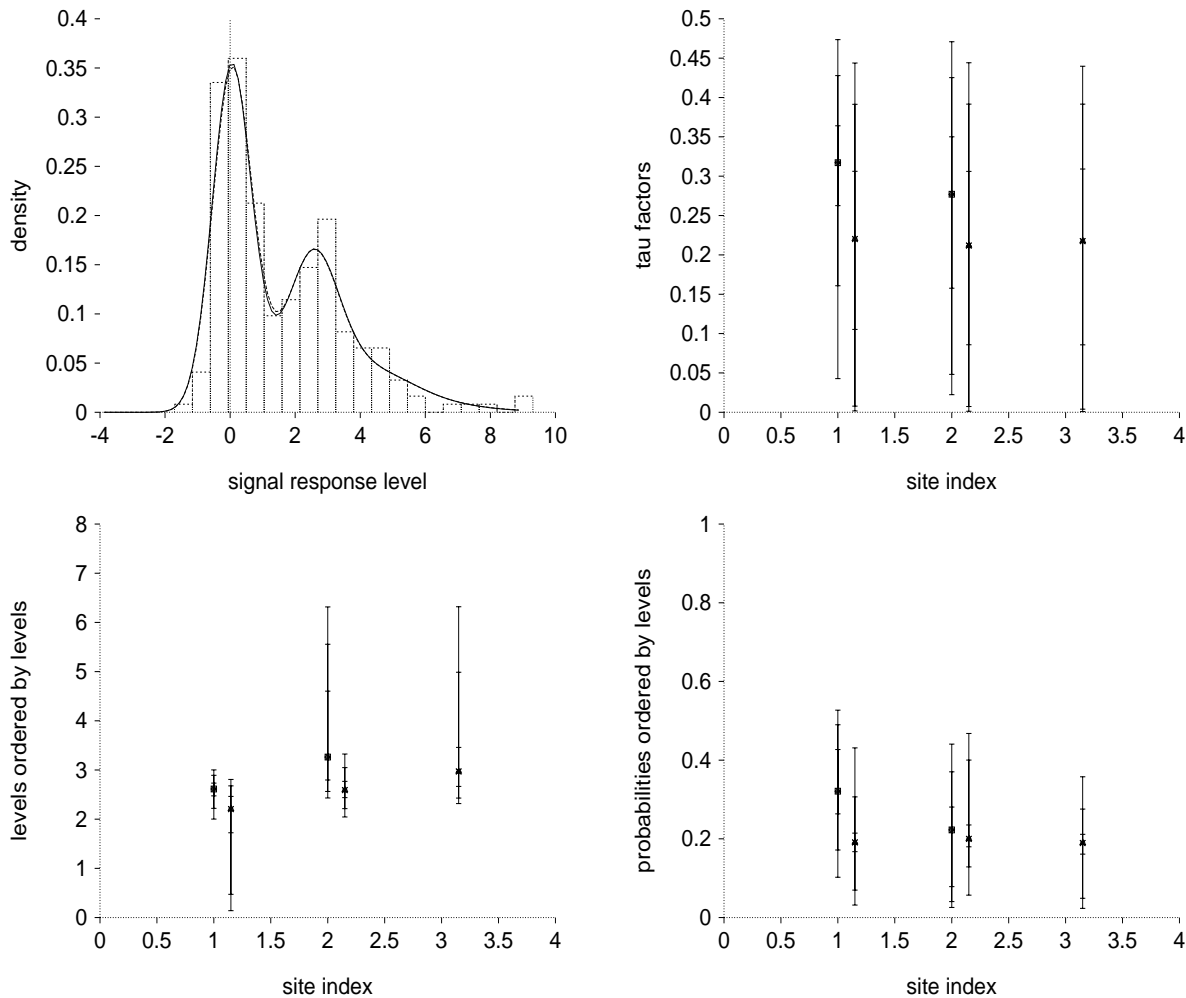


Figure 7. Outputs from analysis of the EPSP B2 data, from separate analysis with $h = s = 2$ and $h = s = 3$, in the model extension to incorporate site-specific measures of intrinsic variability in release levels. The top frame plots the posterior predictive densities from the two analyses together with a histogram of the data. The density curves are from the model with $h = 2$ (full line) and $h = 3$ (dashed line). The other frames display the usual posterior intervals for, in order, the τ_j , μ_j and π_j , with site indices representing increasing levels μ_j . At each site index, the left hand interval corresponds to analysis with $h = s = 2$ sites, the right hand one to $h = s = 3$ sites.

with fixed numbers of sites $h = s = 2$ and $h = s = 3$, respectively, in format as in Figure 6. Note, again, the close similarity of predictive density functions from the two models. Apart from impacting inferences on the number of active sites, note that the inferred average release levels μ_j are somewhat lower than in the earlier analysis with no intrinsic variance; this is to be expected, as intrinsic variability now accounts for some of the higher observed levels of response. Figure 7 also includes display of posterior intervals for the intrinsic variance parameters τ_j , in the same format as for release levels and chances. We note the inference of comparable levels across the 2 (possibly 3) active sites, at appreciable levels of 10 – 40% on the coefficient of variance scale, though with high uncertainty.

Clearly, the two competing explanations – increased, intrinsic variability at higher response

levels versus more sites with fixed release levels – are in significant conflict in this example. The intrinsic variance model supports essentially just 2 sites, attributing the data at higher levels of response to 2 sites with inherently high levels of intrinsic variability. The earlier analysis, constrained to no intrinsic variance, needs 4 active sites to adequately cater for the observed data. Choosing between models involves committing to one of the extremes – yes or no in favour of intrinsic variability. Current physiological research, and a keener focus on the molecular levels of neurotransmission, may shed light on the issue. We note also there is no imperative to accept either model on the standard bases of simplicity and parsimony. Though an intrinsic variance model leads to fewer active sites, and hence is apparently a simpler description of the data generating mechanism, it in no way represents a simpler model in terms of numbers of parameters and model complexity, as it requires many more latent μ_j parameters to describe the data.

7 CONCLUDING COMMENTS

In addition to the scientifically and statistically live issues arising in connection with investigations of intrinsic variability, several other areas are of current interest and are under initial investigation. As presented, our models make no explicit allowance for outliers in signal recordings, nor for systematic features such as arise when synaptic tissue exhibits inhibitory as well as excitatory responses to stimuli. An example of the latter is discussed in West and Turner (1994), for example. Another, related feature of some EPSP experiments is occasional, though rare, ‘blocking’ of the responses of all or a subset of transmission sites, leading to more signal records than expected located around the zero level. There are obvious ways to extend the current models to allow for these kinds of events and considerations; one such is to add mixture components explicitly representing such phenomena. Perhaps the most needed generalisation, currently, is to relax the assumption of normally distributed experimental noise. This is currently in development, using noise models based on work in Bayesian density estimation following Escobar and West (1995), and West, Müller and Escobar (1994), and preliminary results will be reported shortly, elsewhere. See Cao and West (1996) for related developments.

One beginning area of research involves issues of dependence in release characteristics over time and across sites. On the first, synaptic responses in general may be viewed as a stream where, at each site, the response may depend on previous occurrences. Current analyses assume independence. The very nature of the paired pulse potentiation experiments, however, explicitly recognise dependence effects, both in terms of response/no response and in terms of response levels, and of the importance of time elapsed between successive stimuli generating the responses. Statistical developments of time series issues are of interest here to explore and, ultimately, model time variation in patterns of synaptic transmission. Interest in transmission dependencies also extends across sites, i.e., to investigations of ‘connectivities’ between sites that may have physiological interpretation and importance, though serious inroads into study of this issue is some way in the future.

It will be apparent that the technical aspects of simulation based analysis in the novel mixture models here bear further study, from a practical perspective. It is certainly desirable to investigate variations on the basic Gibbs sampling schemes that enhance and improve convergence characteristics. It is possible that hybrid schemes based, in part, on the use of auxiliary variables (e.g., Besag and Green, 1993) will prove useful here; the need for improved computational tools will be more strongly felt in more complex models incorporating non-normal noise components. That said, the

current approach described and illustrated here provide a useful basis for exploring and assessing EPSP data in attempts to more directly evaluating the structure and stochastic characteristics of synaptic transmission as it is currently viewed in the neurophysiological community. Our current work on the applications side is focused on refining this basic framework and in exploring ranges of EPSP data sets to isolate and summarise the diversity of observed response mechanisms in mammalian central nervous systems.

APPENDIX A

The Dirichlet structure of the component prior distribution $F(\cdot)$ in the prior modelling for non-zero site levels μ_j in Section 2.2.2. is detailed here. Write $f(\cdot)$ for the density of $F(\cdot)$. Recall that, for some probability q , the μ_j are conditionally independent and drawn from

$$(\mu_j|F, q) \sim q\delta_0(\mu_j) + (1 - q)f(\mu_j)$$

for $j = 1, \dots, s$. Now $F(\cdot)$ is assigned a Dirichlet process prior with base measure $aU(\mu_j|0, u)$ for some precision $a > 0$ and the specified range $u > 0$. Marginalising over F , the full joint prior for all s elements of μ is now implicitly defined by

$$\begin{aligned} p(\mu|q, a) &= \int p(\mu|F, q)dP(F|a) = \int \left\{ \prod_{j=1}^s p(\mu_j|F, q) \right\} dP(F|a) \\ &= \int \left\{ \prod_{j=1}^s q\delta_0(\mu_j) + (1 - q)f(\mu_j) \right\} dP(F|a). \end{aligned}$$

This is a rather complicated expression whose structure is most easily, and usefully, understood by exploring the corresponding conditional priors; for each $j = 1, \dots, s$ we have

$$\begin{aligned} p(\mu_j|\mu_{-j}, q, a) &= \int p(\mu_j|F, \mu_{-j}, q)dP(F|a, \mu_{-j}) = \int p(\mu_j|F, q)dP(F|a, \mu_{-j}) \\ &= q\delta_0(\mu_j) + (1 - q)E(dF(\mu_j)|a, \mu_{-j}). \end{aligned}$$

Under the Dirichlet prior, the function $E(dF(\mu_j)|a, \mu_{-j})$ is trivially deduced; write h_j for the number of non-zero elements of μ_{-j} , and N_j for the corresponding set of indices, $N_j = \{i|\mu_i > 0, i = 1, \dots, s; i \neq j\}$. Then, using results in Ferguson (1973), or West, Müller and Escobar (1994), we have

$$E(dF(\mu_j)|a, \mu_{-j}) = r_j U(\mu_j|0, u) + (1 - r_j)h_j^{-1} \sum_{i \in N_j} \delta_{\mu_i}(\mu_j)$$

where $r_j = a/(a + h_j)$. We now have the full joint prior $p(\mu|q, a)$ defined via the set of conditionals displayed in equation (3) of the text.

APPENDIX B

More detail on the conditional posterior distributions underlying the MCMC sampling algorithm of Section 3 is provided here. Each of the seven conditional distributions is detailed, with comments about sampling at each step in the MCMC iterations.

(1) *Conditional posteriors for site levels μ*

For each j , we have the conditional prior (3) for μ_j , and the associated conditional likelihood function

$$p(y|\mu, z, m, v) = \prod_{i=1}^n N(y_i|m + z'_i\mu, v).$$

Write $M_j = \sum_{i=1}^n z_{ij}$. Now $M_j = 0$ if, and only if, $z_{ij} = 0$ for all i , so that the conditioning indicator vector z_i corresponds to no transmissions at all from the synaptic site labelled j . In this case, $p(y|\mu, z, m, v)$ does not depend on μ_j and hence the conditional posterior for μ_j collapses to the prior (3). Otherwise, site j is identified as transmitting for at least some of the signal observations, hence $M_j > 0$ and the likelihood reduces to a function of μ_j proportional to the normal density $N(\mu_j|m_j, v/M_j)$ where $m_j = M_j^{-1} \sum_{i=1}^n z_{ij}e_{ij}$ based on ‘residuals’ $e_{ij} \stackrel{\text{def}}{=} y_i - m - \sum_{r=1, r \neq j}^s z_{ir}\mu_r$ for $i = 1, \dots, n$.

Then, based on the conditional prior (3), we deduce (for this case of $M_j > 0$)

$$p(\mu_j|\mu_{-j}, z, m, v, q, a, y) \propto q_{j0}\delta_0(\mu_j) + q_{jj}g_j(\mu_j) + \sum_{i \in N_j} q_{ji}\delta_{\mu_i}(\mu_j)$$

with the following components:

- $q_{j0} = qN(0|m_j, v/M_j)$;
- $q_{jj} = (1 - q)r_j c_j/u$ with $c_j = \Phi((u - m_j)\sqrt{M_j}) - \Phi(-m_j\sqrt{M_j})$, and where $\Phi(\cdot)$ is the standard normal cdf;
- $q_{ji} = (1 - q)(1 - r_j)h_j^{-1}N(\mu_i|m_j, v/M_j)$ for $i \in N_j$;
- $g_j(\mu_j) = c_j^{-1}N(\mu_j|m_j, v/M_j)I(0 < \mu_j < u)$, the density of the normal distribution truncated to $0 < \mu_j < u$.

At each step on the MCMC iterations, we sample μ sequencing through $j = 1, \dots, s$, at each step generating a new value of μ_j given the latest sampled values of μ_{-j} and all other conditioning quantities. For each j , this involves computing the $2 + h_j$ positive constants q_{ji} , ($i = 0, j; i \in N_j$), then making a single multinomial draw with chances proportional to these q_{ji} to identify the selected mixture component. If $i = 0$ is selected, then $\mu_j = 0$; if $i > 0$ and $i \neq j$ is selected, then $\mu_j = \mu_i$; otherwise, $i = j$ and μ_j is drawn from $g_j(\cdot)$, simply the $N(\mu_j|m_j, v/M_j)$ truncated to $(0, u)$. Sampling efficiency is improved using variations on this basic routine involving configuration sampling, as introduced by MacEachern (1994), and developed in West, Müller and Escobar (1994).

(2) *Conditional posteriors for site release probabilities π*

As a function of any single π_j , the posterior (5) trivially implies

$$p(\pi_j|\pi_{-j}, z, b) \propto \pi_j^{t_j}(1 - \pi_j)^{n-t_j}p(\pi_j|\pi_{-j}, b)$$

where $t_j = \sum_{i=1}^n z_{ij}$. Under the prior (4), this becomes the mixture

$$p(\pi_j|\pi_{-j}, z, b) \propto s_{jj}Be(\pi_j|t_j + 1, n - t_j + 1) + \sum_{i=1, i \neq j}^s s_{ji}\delta_{\pi_i}(\pi_j),$$

where $Be(\cdot|\cdot, \cdot)$ represents the beta density, and the s_{ji} are positive constants given by $s_{jj} = w t_j!(n - t_j)!/(n + 1)!$, and $s_{ji} = (1 - w)\pi_i^{t_j}(1 - \pi_i)^{n-t_j}/(s - 1)$ for each $i \neq j$.

At each step on the MCMC iterations, we sample π sequencing through $j = 1, \dots, s$, at each step generating a new value of π_j given the latest sampled values of π_{-j} and all other conditioning quantities. For each j , this involves computing the s positive constants s_{ji} , ($i = 1, \dots, s$), then making a single multinomial draw with chances proportional to these s_{ji} to identify the selected mixture component. If $i \neq j$ is selected, then $\pi_j = \pi_i$; otherwise, $i = j$ and π_j is drawn from the beta distribution $Be(\pi_j|t_j + 1, n - t_j + 1)$. Here, again, configuration sampling improves efficiency.

(3) *Conditional posteriors for indicators z*

As a function of the 2^s -vectors z_i , (5) simplifies as $\prod_{i=1}^n p(z_i|\mu, \pi, m, v, y_i)$ where

$$p(z_i|\mu, \pi, m, v, y_i) \propto p(z_i|\pi)p(y_i|\mu, z_i, m, v) \propto \left\{ \prod_{j=1}^s \pi_j^{z_{ij}} (1 - \pi_j)^{1 - z_{ij}} \right\} N(y_i|m + z'_i\mu, v).$$

These 2^s unnormalised probabilities are easily computed to provide a conditional multinomial posterior for the binary vector z_i that is easily sampled; a complete draw of a new value for all indicators z simply involves running through independent simulations for each $i = 1, \dots, n$.

(4) *Conditional posterior for systematic bias m*

The conditional likelihood function in m is a normal form, resulting in $p(m|\mu, z, v, y) \propto p(m|v)\exp(-\sum_{i=1}^n (e_i - m)^2/(2v))$ based on residuals $e_i = y_i - z'_i\mu$; this reduces to a density proportional to $p(m|v)N(m|\bar{e}, v/n)$, which is normal as $p(m|v)$ is normal. This posterior is trivially sampled.

(5) *Conditional posterior for noise variance v*

The conditional likelihood function in v is a standard inverse gamma form, resulting in the posterior $p(v|\mu, z, m, y) \propto p(v|m)v^{-n/2}\exp(-S/(2v))$ where $S = \sum_{i=1}^n (y_i - m - z'_i\mu)^2$ for each i . Under the inverse gamma prior $p(v|m)$ from the noise data analysis, this conditional posterior is also inverse gamma, and so easily sampled.

(6) *Conditional posterior for hyper-parameter q*

Knowledge of μ implies knowledge of $h \leq s$, the number of non-zero elements in μ , and a binomial likelihood function for q . Under the specific beta prior $q \sim Be(s - 1, 1)$, the resulting posterior is $(q|h) \sim Be(q|2s - h + 1, h + 1)$. Under an alternative uniform prior, this would change to $(q|h) \sim Be(q|s - h + 1, h + 1)$. Other beta priors yield appropriate beta posteriors.

(7) *Conditional posterior for hyper-parameter a*

The treatment here follows West (1992), and Escobar and West (1995). Assume a specified gamma prior $G(a|a_0, a_1)$ with density $p(a) \propto a^{a_0-1}e^{-a_1a}$ for $a > 0$. Data augmentation ideas show that the relevant conditional posterior for a given μ (and everything else) depends only on distinct, non-zero values among the elements of μ . So, for any μ , write $h \leq s$ for the number of non-zero values, and $h^* \leq h$ for the number of distinct non-zero values. Then the required conditional density $p(a|\mu)$ is the margin of a joint density $p(a, x|\mu)$ with the following conditionals:

- given x in $(0, 1)$,

$$(a|x, \mu) \sim p_x G(a|a_0 + h^*, a_1 - \log(x)) + (1 - p_x) G(a|a_0 + h^* - 1, a_1 - \log(x))$$

with weights p_x defined by $p_x/(1 - p_x) = (a_0 + h^* - 1)/\{h(a_1 - \log(x))\}$;

- given $a > 0$, $(x|a, \mu) \sim Be(x|a + 1, h)$.

Though the corresponding density $p(a|\mu)$ is not directly manageable, the expression through conditionals is natural in the context of iterative posterior simulation analyses, as described in the above references. We simply augment the parameter space with the latent variable x introduced to determine a joint posterior $p(a, x|\mu)$ with manageable conditionals. Write a_- for the value of a available from the latest iteration, so that the current values of h and h^* have been generated based on that value a_- . To sample a new value of a at the current iteration, generate a value of x from the beta distribution $p(x|a_-, \mu)$ and then sample from the gamma mixture $p(a|x, \mu)$.

(8) *Conditional posterior for hyper-parameter b*

This is treated essentially as is a . Assume a specified gamma prior $b \sim G(b|b_0, b_1)$. Given a vector π , write $s^* \leq s$ for the number of distinct values (all are non-zero, under the assumed prior). Then values of b are iteratively sampled in a fashion similar to a , now based on linked conditional distributions as follows; for some w in $(0, 1)$, the required conditional density $p(b|\pi)$ is the margin of a joint density $p(b, w|\pi)$ with the following conditionals:

- given w ,

$$(b|w, \pi) \sim p_w G(b|b_0 + s^*, b_1 - \log(w)) + (1 - p_w) G(b|b_0 + s^* - 1, b_1 - \log(w))$$

with weights p_w defined by $p_w/(1 - p_w) = (b_0 + s^* - 1)/\{s(b_1 - \log(w))\}$;

- given $b > 0$, $(w|b, \pi) \sim Be(w|b + 1, s)$.

The sampling scheme generates a new value of b is by augmenting the parameter space with w , in a way directly analogous to sampling a in step (7).

APPENDIX C

We briefly detail the posterior structure and new components of the MCMC analysis required to incorporate the new sets of parameters γ and τ in the model extension for intrinsic variability of Section 6. With reference to the original framework in Section 3, detailed in Appendix B above, there are several changes to the simulation of conditional posteriors, and two additional steps. The set of components labelled (1)–(8) in Appendix B are modified and extended as follows.

(1) *Conditional posteriors for site levels μ*

The μ_j are sampled from conditional posteriors that are marginalised over γ . For each site index j in turn, we have the usual conditional prior (3) for μ_j , and the associated conditional likelihood functions from (8); the latter are now more complicated as functions of μ_j and result in posteriors that are not directly simulated. Instead, we use analytic normal approximations as Metropolis-Hasting proposal distributions and embed an M-H independence chain step within the overall MCMC loop. Write m_j for the ‘current’ value of μ_j in the iterations, and fix the variances of (8) by plugging-in m_j for μ_j , for each j . This leads to approximate likelihood components that are

normal, as in the original scheme of Appendix B, though now with differing but known variances. Step (1) of Appendix B can now be completed as written to obtain a candidate draw for μ_j , and the associated M-H accept/reject ratio is easily computed and used to accept or reject the candidate value. Sequencing through $j = 1, \dots, k$ using this strategy produces a new vector μ .

(*) *Conditional posteriors for γ and τ*

The new MCMC steps draw release levels γ_{ij} and then intrinsic coefficients of variance τ_j from appropriate conditional posteriors.

First, the elements of γ . For a fixed i and j , note that, conditional on all other parameters and the observed data, the conditional posterior for the individual release level γ_{ij} is proportional to the product of (6) and (7). If $z_{ij} = 0$, this reduces to the prior (7). Otherwise, (6) and (7) trivially combine to give a normal posterior multiplied by the truncation indicator in (7). Either way, γ_{ij} is trivially sampled from the resulting truncated normal distribution.

Second, the elements of τ . Suppose the prior has the τ_j independent with common density $g(\tau_j)$ (a uniform prior in our current implementation). Then, conditional on γ and μ , the prior combines with the conditional likelihood components arising from (7) to yield independent conditional posteriors for the τ_j . In case $\mu_j = 0$, then (7) is vacuous and the posterior for τ_j is simply the prior, so easily sampled. Otherwise,

$$p(\tau_j|\gamma, \mu) \propto g(\tau_j)c(\tau_j)^{-n} \exp(-G_j/2\tau_j^2)/\tau_j^n$$

where $G_j = \sum_{i=1}^n (\gamma_{ij}/\mu_j - 1)^2$ and $c(\tau_j)$ is the probability on $(0, u)$ under the $N(\cdot|\mu_j, \tau_j^2\mu_j^2)$ distribution. This has the form of an inverse gamma density multiplied by $g(\tau_j)c(\tau_j)^{-n}$, suggesting a Metropolis-Hasting/independence chain step for sampling. Draw from the inverse gamma so determined as an M-H proposal density, and then accept/reject according to the resulting, and easily computed, M-H ratio.

(2) *Conditional posteriors for site release probabilities π*

As in the original model formulation in Appendix B.

Steps (3)–(5) inclusive in Appendix B remain essentially the same, but with levels γ_{ij} replacing the μ_j in the means of the conditional normal distributions for the y_i throughout. The remaining steps (6)–(8) are unchanged.

REFERENCES

- Besag, J. and Green, P.J. (1993) Spatial statistics and Bayesian computation, *J. Roy. Statist. Soc., Ser B*, **55**, 25-37.
- Clements, J.D. (1991) Quantal synaptic transmission?, *Nature*, **353**, 396.
- Cao, G. and West, M. (1994) Bayesian analysis of mixtures of mixtures, *Biometrics*, **52**, 221-227. (An earlier version of this article appeared in the *Proceedings of the 1994 ASA Winter Meetings, ASA Publications*)
- Escobar, D.M. and West, M. (1995) Bayesian density estimation and inference using mixtures, *J. Amer. Statistic. Assoc.*, **90**, 577-588.
- Ferguson, T.S. (1973) A Bayesian analysis of some nonparametric problems, *Ann. Statist.*, **1**, 209-230.
- Gelfand, A.E., and Smith, A.F.M. (1990) Sampling based approaches to calculating marginal densities, *J. Amer. Statistic. Assoc.*, **85**, 398-409.

- Isaac, J., Turner, D.A. and Wheal, H.V. (1994) Components of minimal EPSCs in rat hippocampal CA1 pyramidal cells exhibit varying shape parameters and ‘quantal’ separation, submitted for publication to *J. Physiology*.
- Kullman, D.M. (1989) Application of the expectation-maximization algorithm to quantal analysis of post-synaptic potentials, *J. Neurosci. Methods*, **30**, 231-245.
- Kullman, D.M. (1992) Quantal analysis using maximum entropy noise deconvolution, *J. Neurosci. Methods*, **44**, 47-57.
- Kullman, D.M. and Nicoll, R.A. (1992) Long-term potentiation is associated with increases in quantal content and quantal amplitude, *Nature*, **357**, 240-244.
- Ling, L., and Tollhurst, D.J. (1983) Recovering the parameters of finite mixtures of normal distributions from a noisy record: an empirical comparison of differing estimating procedures, *J. Neurosci. Methods*, **8**, 309-333.
- MacEachern, S.N. (1994) Estimating normal means with a conjugate style Dirichlet process prior, *Comm. Statist.: Simulation and Computation*, **23**, 727-741.
- Martin, A.R. (1966) Quantal nature of synaptic transmission, *Physiol. Rev.*, **46**, 51-66.
- Redman, S. (1990) Quantal analysis of synaptic potentials in neurones of the central nervous system, *Physiol. Rev.*, **70**, 165-198.
- Smith, A.F.M., and Roberts, G.O. (1993) Bayesian computation via the Gibbs sampler and related Markov chain Monte Carlo methods, *J. Roy. Statist. Soc., Ser B*, **55**, 3-23.
- Tierney, L.J. (1994) Markov chains for exploring posterior distributions, *Ann. Statist.*, **22**, 1701-1728.
- Titterton, D.M., Smith, A.F.M., and Makov, U.E. (1985) *Statistical Analysis of Finite Mixture Distributions*, Chichester: John Wiley and Sons.
- Turner, D.A. (1987) Identification of components of evoked composite EPSPs in CA1 hippocampal pyramidal neurons, *Neurosci. Abstr.*, **13**, 155.
- Turner, D.A. and Schlieckert, M. (1990) Data acquisition and analysis system for intracellular neuronal signals, *J. Neurosci. Methods*, **35**, 241-251.
- Turner, D.A. and West, M. (1993) Statistical analysis of mixtures applied to postsynaptic potential fluctuations, *J. Neurosci. Methods*, **47**, 1-23.
- Turner D.A. and Wheal H.V. (1991) Excitatory synaptic potentials in kainic-acid denervated rat CA1 pyramidal neurons, *J. Neurosci.*, **11**, 2786-2794.
- Turner D.A., Chen, Y., Isaac, J., West, M. and Wheal, H.V. (1995) Heterogeneity between non-NMDA synaptic sites in paired-pulse plasticity of CA1 pyramidal neurons in the hippocampus, *J. Physiology*, (to appear)
- Walmsley, B., Edwards, F.R., and Tracey, D.J. (1987) The probabilistic nature of synaptic transmission at a mammalian excitatory central synapse, *J. Neurosci.*, **7**, 1037-1046.
- Walmsley, B., Edwards, F.R. and Tracey, D.J. (1988) Nonuniform release probabilities underlie quantal synaptic transmission at a mammalian excitatory central synapse, *J. Neurosci.*, **60**, 889-908.
- West (1992) Hyperparameter estimation in Dirichlet process mixture models, *ISDS Discussion Paper #92-A03*, Duke University.
- West, M. and Cao, G. (1993) Assessing mechanisms of neural synaptic activity, in *Bayesian Statistics in Science and Technology: Case Studies*, eds: C. Gatsonis, J. Hodges, R. Kass and N. Singpurwalla, Springer Verlag: New York.
- West, M., Müller, P. and Escobar, M.D. (1994) Hierarchical priors and mixture models, with applications in regression and density estimation, in *Aspects of Uncertainty: A Tribute to D.V. Lindley*, eds: A.F.M. Smith and P. Freeman, Oxford University Press.
- West, M. and Turner, D.A. (1994) Deconvolution of mixtures in analysis of neural synaptic transmission, *The Statistician*, **43**, 31-43.
- Wong, K., and Redman, S.J. (1980) The recovery of a random variable from a noisy record with application to the study of fluctuations in synaptic potentials, *J. Neurosci. Methods*, **2**, 389-409.