

The quality of maximum likelihood estimates of ion channel rate constants

D. Colquhoun*, C. J. Hatton* and A. G. Hawkes†

*Department of Pharmacology, University College London, Gower Street, London WC1E 6BT and †European Business Management School, University of Wales Swansea, Singleton Park, Swansea SA2 8PP, UK

Properties of maximum likelihood estimators of rate constants for channel mechanisms are investigated, to see what can and cannot be inferred from experimental results. The implementation of the HJCFIT method is described; it maximises the likelihood of an entire sequence of apparent open and shut times, with the rate constants in a specified reaction mechanism as free parameters. The exact method for missed brief events is used. Several methods for testing the quality of the fit are described. The distributions of rate constants, and correlations between them, are investigated by doing sets of 1000 fits to simulated experiments. In a standard nicotinic receptor mechanism, all nine free rate constants can be estimated even from one single channel recording, as long as the two binding sites are independent, even when the number of channels in the patch is not known. The estimates of rate constants that apply to diliganded channels are robust; good estimates can be obtained even with erroneous assumptions (e.g. about the value of a fixed rate constant or the independence of sites). Rate constants that require distinction between the two sites are less robust, and require that an EC_{50} be specified, or that records at two concentrations be fitted simultaneously. Despite the complexity of the problem, it appears that there exist two solutions with very similar likelihoods, as in the simplest case. The hazards that result from this, and from the strong positive correlation between estimates of opening and shutting rates, are discussed.

(Received 10 October 2002; accepted after revision 17 December 2002; first published online 24 January 2003)

Corresponding author D. Colquhoun: Department of Pharmacology, University College London, Gower Street, London WC1E 6BT, UK. Email: d.colquhoun@ucl.ac.uk

There are two good reasons for trying to identify kinetic mechanisms for receptors. Firstly, it is only by doing so that one can study sensibly the effect of structure changes in agonists (for example, does the change in structure alter the ability to bind, or the ability of the agonist to activate the receptor once bound?). Secondly, it is only by doing so that the effect of mutations in a receptor can be studied rationally (for example, does the mutated residue form part of the agonist binding site?). These questions have been reviewed by Colquhoun (1998).

In order to answer the questions of interest, two things must be done. First a qualitative reaction scheme must be postulated, and then values for the rate constants in the scheme must be found. In many ways the first step is the harder, because unless the reaction scheme is a sufficiently good description of actual physical structural reality, it cannot be expected that physically meaningful conclusions can be drawn from it.

The only sort of receptor for which it has so far been possible to achieve these aims are the agonist-activated ion channels, and then only by observation of single ion channels. In earlier studies (e.g. Colquhoun & Sakmann, 1981, 1985), rate constants in the mechanism could not be estimated directly. Rather, individual distributions (shut

times, open times, burst lengths etc.) were fitted separately with arbitrary mixtures of exponential distributions (e.g. Colquhoun & Sigworth, 1995), and correlations between these quantities were measured separately. It was not possible to take into account all of the information in the record simultaneously, so information from individual distributions had to be cobbled together in a rather arbitrary way to infer values for the rate constants in the mechanism. It was also not possible to make proper allowance for the inability to resolve brief events in a single channel record. Since that time, better methods of analysis have been devised, the most appealing of which is to maximise the likelihood of the entire sequence of open and shut times. The appeal of this method stems from the facts that (a) it provides estimates of the rate constants in a specified mechanism directly from the observations, (b) it is based on measurements of open and shut times (an 'idealisation' of the observed record), so the user has a chance to check the data going into the calculation, (c) the calculation can be carried out without having to choose arbitrarily which particular distributions to plot and (d) it takes into account correctly the fact that in most real records subsequent intervals are not independent of one another (it is common, for example, to find that long open times are followed on average by short shut times), and

uses all of the information in the observed record in a single calculation (Fredkin *et al.* 1985). Since, in the usual general treatment of ion channels, the rate constants for the connections between each pair of states are tabulated in the \mathbf{Q} matrix, it may be said that the method provides an estimate of the \mathbf{Q} matrix. The method was first proposed by Horn & Lange (1983), but at that time it was not possible to allow for the fact that brief events cannot be seen in experimental records. The implementation of the method by Ball & Sansom (1989) had the same problem. Since many brief events are missed in most experimental records, the method was not useable in practice until this problem had been solved. Ball & Sansom (1988*a,b*) gave the exact solution for the missed events problem in the form of its Laplace transform, and various approximate solutions have been proposed too, the best of which appears to be that of Crouzy & Sigworth (1990) (see Hawkes *et al.* 1990; Colquhoun & Hawkes, 1995*b*). However there is no longer any need for approximations because the exact solution to the problem has been found by Hawkes *et al.* (1990, 1992).

Two computer programs are available for doing direct maximum likelihood fitting of rate constants, MIL (Qin *et al.* 1996, 1997) and HJCFIT (Colquhoun *et al.* 1996). The MIL program is available at <http://www.qub.buffalo.edu/index.html> and HJCFIT from <http://www.ucl.ac.uk/Pharmacology/dc.html>. The former uses (a corrected form of) the approximate missed event method of Roux & Sauve (1985); the latter uses the exact solution.

It is the responsibility of anyone who proposes an estimation method to describe the properties of the estimators, and in this paper we describe some of the properties of estimates of rate constants found with HJCFIT. This provides the background for the method, and the necessary justification for the use of HJCFIT to analyse experimental results on nicotinic receptor channels in the accompanying paper (Hatton *et al.* 2003).

METHODS

Resolution and observed events

The durations of openings and shuttings that are measured from an experimental record are extended by the failure to detect brief openings and shuttings. These measured values will be referred to as apparent or extended durations. It is crucial when making allowance for missed brief events that the data should have a well-defined time-resolution (t_{res}) or dead-time, defined so that all events shorter than t_{res} are omitted, and all events longer than t_{res} are present in the record. This is easily achieved by retrospective imposition of a fixed resolution on the data, as described by Colquhoun & Sigworth (1995). This is also desirable even for fitting of empirical distributions (by programs such as our EKDIST), so it is surprising that most other programs that are available, both free and commercial, still do not incorporate this ability. An extended open time, or e -open time is defined for theoretical purposes as the duration of an event that (a) starts when an opening longer than t_{res} occurs and (b) continues until a

shutting longer than t_{res} is encountered. The e -opening can therefore contain any number of shuttings, as long as they are all shorter than t_{res} , separating openings that may be of any length (except for the first, which must be greater than t_{res}) (Ball & Sansom, 1988*a*; Hawkes *et al.* 1990, 1992). The method used in HJCFIT (and EKDIST) for imposition of a fixed resolution follows, as closely as possible, this theoretical definition, though in real records there will always be a few events that cannot be interpreted unambiguously even by time course fitting, in particular those that consist of several contiguous brief events in quick succession (e.g. Colquhoun & Sigworth, 1995).

Distributions of observed events

The theoretical distributions of extended open times etc. will be referred to as HJC distributions because they are calculated by the methods of Hawkes *et al.* (1990, 1992). In contrast, the ideal distributions would be calculated by the simpler methods of Colquhoun & Hawkes (1982) on the assumption that no events are missed.

The ideal ($t_{\text{res}} = 0$) distribution (probability density function, pdf) of an open time is:

$$f_o(t) = \phi_A \mathbf{G}_{\text{AF}}(t) \mathbf{u}_F, \quad (1)$$

where ϕ_A is a $1 \times k_A$ row vector giving the probabilities that the first opening starts in each of the open states (k_A = number of open states), \mathbf{u}_F is a $k_F \times 1$ unit column vector (k_F = number of shut states) and \mathbf{G}_{AF} is a $k_A \times k_F$ matrix defined by Colquhoun & Hawkes (1982); it can be calculated from the \mathbf{Q} matrix (see also Colquhoun & Hawkes, 1995*a,b*). The HJC distribution of the duration of an apparent opening can be written using an analogous notation as:

$$f_o(t) = \phi_A {}^e\mathbf{G}_{\text{AF}}(t) \mathbf{u}_F, \quad (2)$$

where ${}^e\mathbf{G}_{\text{AF}}$ is the HJC analogue of \mathbf{G}_{AF} (the simplicity of the notation disguises the fact that the calculation of the former is a good deal more complicated than calculation of the latter). Exactly analogous results hold for distribution of shut times.

The likelihood, l , of a whole sequence of observed (apparent) open and shut times can now be calculated, as described by Colquhoun *et al.* (1996), as:

$$l = \phi_A {}^e\mathbf{G}_{\text{AF}}(t_{o1}) {}^e\mathbf{G}_{\text{FA}}(t_{s1}) {}^e\mathbf{G}_{\text{AF}}(t_{o2}) {}^e\mathbf{G}_{\text{FA}}(t_{s2}) {}^e\mathbf{G}_{\text{AF}}(t_{o3}) \dots \mathbf{u}_F, \quad (3)$$

where t_{o1}, t_{o2}, \dots represent the first, second apparent open time and t_{s1}, t_{s2}, \dots first, second apparent shut time, etc. Note that openings and shuttings occur in this expression in the order in which they are observed. Thus ϕ_A is a $1 \times k_A$ row vector giving the probabilities that the first opening starts in each of the open states, $\phi_A {}^e\mathbf{G}_{\text{AF}}(t_{o1})$ is a $1 \times k_F$ row vector the elements of which give the probability density of the open time t_{o1} multiplied by the probabilities that the first shut time, t_{s1} , starts in each of the shut states given that it follows an opening of duration t_{o1} . Then $\phi_A {}^e\mathbf{G}_{\text{AF}}(t_{o1}) {}^e\mathbf{G}_{\text{FA}}(t_{s1})$ is a $1 \times k_A$ row vector, the elements of which give the probability density of (t_{o1}, t_{s1}) multiplied by the probabilities that the next open time, t_{o2} , starts in each of the open states given that it follows an opening of duration t_{o1} and a shut time of duration t_{s1} . This continues up to the end of the observations. The process of building up the product in eqn (3) gives, at each stage, the joint density of the time intervals recorded thus far multiplied by a vector that specifies probabilities for which state the next interval starts in, conditional on the durations of those intervals. This process uses all the information in the record about correlations between intervals.

Calculation of the likelihood in practice

The exact solution for ${}^e\mathbf{G}_{AF}(t)$ (Hawkes *et al.* 1990), has the form of a piecewise solution; one result is valid between t_{res} and $2t_{res}$, another between $2t_{res}$ and $3t_{res}$, and so on. Furthermore it is not, as in the ideal case, a sum of k_A exponentials (with time constants that are the reciprocals of the eigenvalues of $-\mathbf{Q}_{AA}$), but involves all of the $k - 1$ eigenvalues of \mathbf{Q} (where k is the total number of states, $k_A + k_F$), multiplied by polynomial functions of time, of higher degree in each time interval. For long intervals this expression becomes quite complicated (and eventually gets unstable), but luckily we are rescued from this complication by Hawkes and Jalali's beautiful asymptotic form (Hawkes *et al.* 1992), which is essentially exact above $2t_{res}$ or $3t_{res}$. The beauty of this solution lies in the fact that (a) it can be written in the usual form, as a mixture of simple exponentials, and (b) it has the 'right number' of exponentials, k_A , exactly the same as when events are not missed, though the time constants, and areas, of the k_A exponentials are, of course, not the same as for the ideal ($t_{res} = 0$) open time pdf – they have to be found numerically.

The program, HJCFIT, uses the exact solution for the first two dead times, i.e. for intervals of duration between t_{res} and $2t_{res}$, and for intervals between $2t_{res}$ and $3t_{res}$. For longer intervals, the asymptotic form is used. It is easily verified that this procedure is, for all practical purposes, exact, because the program allows both exact and asymptotic solutions to be plotted as superimposed graphs (e.g. Fig. 6B) and in all cases these curves became indistinguishable for durations well below $3t_{res}$.

Dealing with an unknown number of channels: choosing appropriate start and end vectors

The program uses a modified simplex algorithm to maximise the likelihood of the observed sequence of open and shut times.

Although the simulated experiments all contained one channel, the number of channels in the membrane patch is not known in real experiments. Therefore most of the simulations shown here were analysed, as were the real experiments, by methods that do not assume a number of channels. This is done by dividing the record into stretches (groups of openings) that are short enough that we can be almost sure that each group originates from only one channel. At low acetylcholine concentrations the groups are short – they consist of individual 'activations' of the channel (bursts) which are made up of one or more apparent openings. For the muscle type nicotinic receptor there would rarely be more than 14 openings per activation (mean about 4.7; Hatton *et al.* 2003), though the omission of brief shuttings means that the apparent number of openings per activation is smaller than the true number. At high agonist concentrations, long groups (clusters) of openings occur during which the probability of being open is high enough that we can be sure that the whole cluster originates from a single channel (Sakmann *et al.* 1980). In either case a suitable critical shut time for definition of bursts can be decided from the distribution of apparent shut times (see Colquhoun & Sigworth, 1995).

Since each group of openings is thought to originate from one channel only, a likelihood can be calculated for that group from:

$$l = \phi_b {}^e\mathbf{G}_{AF}(t_{o1}) {}^e\mathbf{G}_{FA}(t_{s1}) {}^e\mathbf{G}_{AF}(t_{o2}) {}^e\mathbf{G}_{FA}(t_{s2}) {}^e\mathbf{G}_{AF}(t_{o3}) \dots \mathbf{e}_F. \quad (4)$$

This is the same as eqn (3), apart from the initial and final vectors. For low concentration records, these initial and final vectors, ϕ_b and \mathbf{e}_F , were calculated as described by Colquhoun, Hawkes &

Srodzinski (1996; eqns (5.8), (5.11)), and so will be referred to as CHS vectors. This method is based on the fact that the long shut times that separate one channel activation from the next are known to be equal to or longer than the t_{crit} value that was used to define the bursts of openings. If there are several channels in the patch, rather than one, then the two consecutive activations may arise from different channels and in this case the true shut time (for one channel) between one activation and the next would be longer than the observed interval, so in such a case it must be longer than t_{crit} . The initial and final vectors are found on the basis that although we do not know the true (one channel) length of the shut times between bursts, we do know that they must be greater than t_{crit} . Therefore the relevant probabilities are integrated over all shut times greater than t_{crit} . This method is appropriate only for cases in which the shut times between bursts are spent in states that are represented by the mechanism being fitted. For records at high concentrations, the shut times between bursts of openings will be spent largely in desensitised state(s). We preferred to omit desensitised states from the mechanisms being fitted in most cases, because there is still uncertainty about how to represent them, and because it was not our aim to investigate desensitisation. Therefore for high concentration records the CHS vectors were not used, and the likelihood for each high concentration group (cluster) of openings was calculated from eqn (3). Although not exact, this procedure can be justified by the fact that the bursts observed at high agonist concentrations usually contain many openings, so the effect of the initial and final vectors will be quite small. The simulations described below test these procedures.

In HJCFIT, the likelihoods are calculated, from eqn (3) or eqn (4), for each of the groups of openings in the record, and the resulting burst log-likelihoods, $L = \log(l)$, are added to get an overall likelihood for the whole observed record. The simplex algorithm used by HJCFIT finds the values for the rate constants in the mechanism (the elements of the \mathbf{Q} matrix) that maximise this likelihood. The sum of all of these individual burst log-likelihoods gives a proper overall log-likelihood only if the bursts behave independently of each other. This is likely to be true even if only one channel is present (Colquhoun & Sakmann, 1985). We shall continue to refer to the sum as the overall log-likelihood, though pedantically it should perhaps be called a pseudo-log-likelihood.

Bad intervals

In real records it is not uncommon for bits of the record to have to be excised, for example because the occasional double opening occurs, or because the patch temporarily becomes noisy. Such events are marked as 'unusable' while the experiment is being measured in SCAN (details at end of Methods). When groups of openings are being constructed, the group must end at the last good interval that precedes an unusable interval. The conservative approach would then be to discard that entire group and start looking for a new group at the next good interval. However, the criterion for construction of a group is that all openings in a group come (almost certainly) from one individual channel. There is no requirement that they correspond to any physiological or mechanistic phenomenon, as there would be, for example, when constructing burst lengths that are intended to represent individual channel activations. Therefore in most cases it will be appropriate to choose the option in HJCFIT to treat bad intervals as a valid end-of-group. This procedure is followed even if the analysis specified that all data were from one channel so analysis in bursts would not be required normally. Simulated data, as used here except for Figs 12 and 13, contains no unusable intervals.

Constraints between rate constants and use of the EC₅₀

If the mechanism contains cycles then one of the rate constants in the cycle may be fixed by the constraint of microscopic reversibility (see, for example, Colquhoun & Hawkes, 1982) from the values of all of the others, thus reducing the number of free parameters by one. This was done in all the fits described here, though HJCFIT also allows irreversible mechanisms to be specified; the calculations assume a steady state, but not necessarily equilibrium.

In almost all the cases discussed here, the record is fitted in bursts, to circumvent the lack of knowledge of the number of channels in the patch, as we are forced to do when analysing real records. This means that, for low concentration records, we have no knowledge of how frequently the channel is activated, so at least one rate constant cannot be determined (at least if only a single concentration record is fitted). One way round this is to fix one rate constant at a plausible value. The effect of fixing one of the association rate constants, at either its correct value (that used to simulate the data) or at an incorrect value, is investigated below.

A better method in principle is to use information from another source. One option is to specify, from other experiments, the concentration of agonist that results, at equilibrium, in the probability of a channel being open (P_{open}) that is 50% of the maximum possible P_{open} , i.e. the EC₅₀. In HJCFIT there is an option to supply a value for the equilibrium EC₅₀, which, in conjunction with values of all the other rate constants, can be used to calculate the value of any specified rate constant at each stage during the fitting. This reduces the number of free parameters by one, though of course good estimates will be obtained only insofar as an accurate EC₅₀ can be specified.

In HJCFIT, the number of free parameters can be reduced by constraining any rate constant to be a fixed multiple of another. This is, of course, not desirable if the true rates do not obey the constraint. The effects of incorrectly applying such constraints is investigated below (see Fig. 14).

Other sorts of constraint can be applied. For example an upper limit can be set for the value of any association rate constant, to prevent physically unrealistic values being found. This is achieved simply by resetting the rate constant in question to its upper limit before evaluating the likelihood. Likewise if a value of a rate constant should go negative during the fitting process, it can be reset to a value near zero. It is a virtue of search methods like simplex, that any arbitrary constraint of this sort is easily incorporated.

A more effective way of preventing a rate constant from going negative is to do the fitting process with the logarithm of the rate constant (e.g. Ball & Sansom, 1989). This is now the default method in HJCFIT, because it not only prevents negative rates being fitted, but is also three or four times faster, presumably because the likelihood surface has a shape such that fewer changes of the search direction are needed (in simplex, a change of direction needs two function evaluations for each fitted parameter).

When a wide range of parameter values is explored, it is possible to get into regions where the likelihood cannot be calculated, and in such cases it is important that the program should not crash, but carry on looking for better parameter values. The main strategy for achieving this in HJCFIT is to keep a record during the fitting process of the best (highest likelihood) set of parameters so far. If

values are explored subsequently for which the likelihood cannot be calculated (e.g. a matrix becomes singular, or asymptotic roots cannot be found), then the parameters are replaced by the previous best values, to which a random perturbation (within a specified range) is applied, to prevent looping. Out of nearly 50 000 fits done for this work, only two crashes resulted from numerical problems during fitting.

Errors

After the fit is completed, internal estimates of errors are calculated by obtaining a numerical estimate of the 'observed information matrix', the Hessian matrix, \mathbf{H} , with elements defined as:

$$h_{ij} = - \left. \frac{\partial^2 L(\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} \right|_{\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}}, \quad (5)$$

where $L = \log(l)$ denotes the log(likelihood), $\boldsymbol{\theta}$ is the vector of free parameters, denoted $\hat{\boldsymbol{\theta}}$ at the point where the likelihood is maximised and θ_i is the i th free parameter.

This is inverted to obtain the covariance matrix, \mathbf{C} (with elements denoted c_{ij}), as:

$$\mathbf{C} = \mathbf{H}^{-1}. \quad (6)$$

The square roots of the diagonal elements of the covariance matrix provide approximate standard deviations for each parameter estimate. The off diagonals are used to calculate the correlation matrix, which is printed by HJCFIT, as:

$$r_{ij} = \frac{c_{ij}}{\sqrt{(c_{ii}c_{jj})}}. \quad (7)$$

To avoid rounding errors, before calculation of the numerical estimates of the second derivatives, it was first found what increment in each parameter was needed to decrease the log(likelihood) by a specified amount (by default, 0.1 log units). In cases where the fit is very insensitive to a parameter it may be impossible to find a suitable increment, and in this case the corresponding row and column of \mathbf{H} for that parameter are simply omitted from eqns (5) and (6) (this procedure is justifiable only insofar as the row and column that are omitted consist of elements that are all close to zero). In other words parameters that have little effect on the likelihood are treated, for the purposes of error estimation, as constants.

Estimatability

A vital problem for the analysis of experiments is to know how many free parameters (in this case, rate constants) can be estimated from experimental data. It was shown by Fredkin & Rice (1986) and Bauer *et al.* (1987) that a single record can provide estimates of up to $2 k_A k_F$ rate constants. For scheme 1 (Fig. 1), $k_A = 3$ and $k_F = 4$ so up to 24 rate constants could be estimated in principle (compared with 14 in scheme 1). However the simulations done here show that under our conditions (resolution 25 μ s and unknown number of channels in the patch), the practical limit is more like 8–10. Indeterminacy of parameters should be shown by lack of reproducibility in replicate experiments (as long as the fit is not simply leaving initial guesses little changed). An indication of indeterminacy should also be apparent in a single fit when the calculations in eqns (5), (6), (7) indicate large (or indeterminate) errors and/or high correlations. However the number of replicates will always be limited (especially

if several runs are fitted simultaneously), and simulations provide a valuable way of investigating determinacy in more detail, under realistic conditions.

Plotting of histograms

For the display of open times etc., the histograms show the distribution of log(time), with the frequency density plotted on a square root scale, as is conventional for experimental results (Sigworth & Sine, 1987). The histograms of estimates of rate constants are shown without any transformation, to make clear the real form of the distribution.

Checking the quality of the fit

The whole fitting process is done on the basis of the list of open and shut times produced by idealisation of the data; nothing need be plotted. After the fit it is, of course, very important to see that the results do in fact describe the observations. The quality of the fit is best judged by looking at the data with the fits superimposed. Although the whole record is fitted at once, in order to display the fit it is still necessary to show separate histograms. In HJCFIT, the following plots can be drawn at the end of the fit (examples are shown in Figs 6, 14, 16 and 18).

(1) Histogram of apparent open times, with the HJC open time pdf that is predicted by the fitted parameters with the resolution that was used for the analysis (e.g. Fig. 6A). If all is well, the HJC distribution should be a good fit to the observations in the histogram, even though the curve was not fitted to the histogram data. In addition the estimated 'true' distribution is shown as a dashed line, i.e. the distribution of open times with perfect resolution that is predicted by the fitted parameters, as calculated from the much simpler results of Colquhoun & Hawkes (1982).

(2) Histogram of apparent shut times, with fits superimposed as for open times (e.g. Fig. 6B). For apparent open and shut times, the asymptotic distribution, which is a mixture of k_A or k_F exponentials respectively, can also be plotted (with components if required) as a check that it becomes essentially identical with the exact distribution above $3t_{\text{res}}$ (e.g. Fig. 6B).

(3) The conditional pdf of apparent open times can be plotted to test whether correlations between open and shut times are described adequately by the mechanism and its fitted rate constants. The histogram includes only those openings that are preceded (or followed, or both) by shittings with lengths in a specified range. On the data is superimposed the appropriate conditional HJC distribution that is calculated from the fitted rate constants and resolution (calculated as described in Colquhoun *et al.* 1996). In addition the distribution of all apparent open times can be shown for comparison (as a dashed line). An example is shown in Fig. 6C. This display may be repeated for different shut time ranges.

(4) The mean apparent open times for openings that are preceded (or followed) by shittings with lengths in a specified series of ranges (e.g. 0.05–0.1 ms, 0.1–0.15 ms, 0.15–1 ms, 1–10 ms and 10–100 ms). The experimentally observed values are shown, together with the values that are calculated from the fit by HJC methods (see Colquhoun *et al.* 1996). This provides another test of how well the fit describes the observations. An example is shown in Fig. 6D, which also shows the theoretical continuous relationship between mean open time and adjacent shut time as a dashed line, though this cannot be compared directly with the data because of the need to bin data for display (Colquhoun *et al.* 1996).

(5) The three-dimensional distributions. The bivariate HJC distribution of apparent open time and adjacent apparent shut time, $f(t_o, t_s)$, (or of open times separated by specified lag) can be plotted. It is hard to make much of this visually, so Magleby & Song (1992) suggested that the useful information in it could be more helpfully plotted as a 'dependency plot' (see also Colquhoun & Hawkes, 1995b, for a brief account).

Dependency is defined as:

$$d(t_o, t_s) = \frac{f(t_o, t_s)}{f_o(t_o)f_s(t_s)} - 1, \quad (8)$$

where $f(t_o, t_s)$ is the bivariate HJC pdf, as above, and $f_o(t_o)$ and $f_s(t_s)$ are the normal unconditional HJC distributions of apparent open time and shut time, respectively. The dependency will therefore be zero if open and shut times are independent, above zero if there is an excess of open times at any specified shut time, and below zero if there is a deficiency. The observed dependency plot can be displayed as a three-dimensional graph (though a large number of observations is needed to get a smooth one), and the dependency predicted by the fit (calculated by HJC methods) can be similarly displayed, though there is no way to superimpose the fit on the data. Examples are shown in Fig. 6E and F.

Simulations

Experimental results are simulated by use of the high quality pseudo-random number generator of Wichmann & Hill (1985). This generates a uniformly-distributed number, u , between 0 and 1, which was used to generate an exponentially distributed interval, the duration of the sojourn in the current state, state i say. The mean time spent in the i th state is $\tau_i = -1/q_{ii}$ and the corresponding random duration is $-\tau_i \ln(u)$. If the current (i th) state is connected to more than one other state then another random number is generated to decide, with the appropriate probability, which state is visited next. Adjacent intervals of equal conductance are then concatenated to generate a simulated open or shut time.

In each simulated experiment, 20 000 intervals were generated (this may need something of the order of 0.5 million individual state transitions, because many transitions are between states of equal conductance). Next a fixed resolution was imposed on these 20 000 intervals. With a resolution of 25 μ s (as used in most experiments, Hatton *et al.* 2003), slightly over half of the 20 000 intervals were eliminated as being undetected (shorter than 25 μ s), leaving about 9000–10 000 resolved intervals that were used as the input for fitting.

On a 1.5 GHz PC, each fit took from about 40–60 s (for about 9500 transitions at a single concentration with 9 free parameters), to 4 or 5 min (for two concentrations with 13 free parameters, or 12 plus an EC₅₀ constraint).

The program that is used, HJCFIT, is available at <http://www.ucl.ac.uk/Pharmacology/dc.html>, together with program, SCAN, that is used for the fitting of durations and amplitudes that form the input to HJCFIT. The programs SIMAN (to inspect results of repeated simulations), SCBST and SCALCS (to calculate ideal burst properties and macroscopic currents, respectively) are also available there. The current version of HJCFIT, as well as allowing the input of experiments analysed with SCAN, also allows simulated experimental results to be generated internally, and fitted. This can be repeated any specified number of times. In this paper, 1000 sets of simulated data were generated and each set fitted to generate 1000 estimates of each rate constant.

RESULTS

The mechanism

The HJCFIT program allows essentially any sort of reaction mechanism to be specified, but the immediate aim of this paper was to investigate the properties of estimates from experiments on nicotinic acetylcholine (ACh) receptors (Hatton *et al.* 2003), so all simulations were done using the mechanisms shown in Fig. 1.

Scheme 1 is a mechanism that has been used in several studies of the ACh receptor (e.g. Colquhoun & Sakmann, 1985; Milone *et al.* 1997). It represents a molecule with two ACh binding sites that are different from each other, denoted in Fig. 1 as the *a* and *b* sites. Occupation of either site alone can produce mono-liganded openings, although these are rare and brief (e.g. Colquhoun & Sakmann, 1981; Jackson, 1988; Parzefall *et al.* 1998). The notation for the rate constants is such that the subscripts *a* and *b* denote which of the two sites is involved, and the subscripts 1 and 2 indicate whether the binding is the first (other site vacant) or second (other site occupied). Thus, for example, k_{+2a} denotes the association rate constant for binding to the *a* site when the *b* site is already occupied. This is, of course, not the only mechanism that can be envisaged, but it is the mechanism that is most appropriate

in the light of what is known about the structure of the receptor, and it can describe with quite good accuracy all the observations.

In general any such mechanism must allow for the possibility that the channel may open spontaneously when no ligand is bound. Such openings are thermodynamically inevitable, and have been reported for the embryonic form of the nicotinic receptor (mouse muscle cells in culture, Jackson, 1984), and occur with some mutant receptors (Grosman & Auerbach, 2000). However spontaneous openings seem to be either too infrequent (or too short) to be detected in the adult muscle receptor, and we have not been able to detect them. Since the purpose of scheme 1 is to fit data, we cannot include states that are not detectable in our observations.

Scheme 2 (Fig. 1B) is the same mechanism as scheme 1, but with a single desensitised state added. This mechanism was used to simulate experiments with high concentrations of agonist, when the records contain long desensitised periods. This mechanism is too simple to describe accurately the desensitisation process; that requires a cyclic mechanism (Katz & Thesleff, 1957) and many more desensitised states (Elenes & Auerbach, 2002). However it is not our intention to investigate desensitisation here, and in all cases scheme 1 was fitted to the simulated experiment. For higher agonist concentrations, scheme 2 was used to simulate the observations, but the results were then fitted in bursts (see Methods) with scheme 1. When scheme 2 was used to simulate observations, we took the rate constant for entry into the desensitised state $\beta_D = 5 \text{ s}^{-1}$, and for exit from the desensitised state $\alpha_D = 1.4 \text{ s}^{-1}$. These values gave rise to desensitised periods (spent in state A_2D) in the simulated record of mean length $1/\alpha_D = 714 \text{ ms}$, roughly as observed. At $10 \mu\text{M}$ ACh, inspection of the shut time distribution showed that using $t_{\text{crit}} = 5 \text{ ms}$, to define bursts when fitting scheme 1, eliminated essentially all of the desensitised intervals, while including most of the others. Scheme 2 can also be used, with a much shorter-lived 'desensitised' state ($1/\alpha_D$ about 1 ms) to describe the 'extra shut state' invoked by Salamone *et al.* (1999), as in Hatton *et al.* (2003).

Fitting records at a single concentration with constraints

Constraints. Scheme 1 (Fig. 1) has 14 rate constants, but one of them (k_{+1a}) was always determined by microscopic reversibility so there are 13 free parameters to be fitted. This can be reduced to 10 free parameters if it is assumed that the binding to site *a* is the same whether or not site *b* is occupied, and vice versa. This assumption implies that the two different binding sites behave independently of each other while the channel is shut. This is plausible, given the distance between the sites, but it is not inevitable. Nevertheless this assumption of independence has been




Figure 1. The two reaction schemes that were used for simulation of experiments

A = agonist, R = shut channel, R^* = open channel. R_a = the *a* binding site, R_b = the *b* binding site. D = desensitized channel.

Table 1. Rate constants used for simulation, and as initial guesses for fits

		True 1	Guess 1	Guess 2	True 2	Guess 3	Guess 4
α_2	s^{-1}	2000	1500	1500	2000	1500	1500
β_2	s^{-1}	52 000	50 000	50 000	50 000	50 000	50 000
α_{1a}	s^{-1}	6000	13 000	2000	4000	2000	10 000
β_{1a}	s^{-1}	50	50	20	80	20	50
α_{1b}	s^{-1}	50 000	15 000	80 000	40 000	80 000	20 000
β_{1b}	s^{-1}	150	10	300	10	30	30
k_{-2a}	s^{-1}	1500	6000	1000	12000	1000	20000
k_{+2a}	$M^{-1} s^{-1}$	2.0×10^8	1.0×10^8	1.0×10^8	0.5×10^8	1.0×10^8	1.0×10^8
k_{-2b}	s^{-1}	10 000	5000	20 000	2000	20 000	1000
k_{+2b}	$M^{-1} s^{-1}$	4.0×10^8	1.0×10^8	1.0×10^8	5.0×10^8	4.0×10^8	2.0×10^8
k_{-1a}	s^{-1}	1500	6000	1000	400	1000	1000
k_{+1a}	$M^{-1} s^{-1}$	2.0×10^8	1.0×10^8	1.0×10^8	0.2×10^8	1.0×10^8	0.5×10^7
k_{-1b}	s^{-1}	10 000	5000	20 000	100	20 000	1000
k_{+1b}	$M^{-1} s^{-1}$	4.0×10^8	1.0×10^8	1.0×10^8	3.0×10^8	4.0×10^8	2.0×10^8
E_2	—	26	—	—	25	—	—
E_{1a}	—	0.0083	—	—	0.02	—	—
E_{1b}	—	0.003	—	—	0.00025	—	—
K_{2a}	μM	7.5	—	—	240	—	—
K_{2b}	μM	25	—	—	4	—	—
K_{1a}	μM	7.5	—	—	20	—	—
K_{1b}	μM	25	—	—	0.333	—	—

Values for rate constants are shown in the first 14 rows. The last seven rows show, for the true rate constants, the corresponding values of the equilibrium constants. True 1: the true values for rate constants for scheme 1 (Fig. 1), that were used for simulation of experiments shown in Figs 2–10. Guess 1: the initial guesses used for fitting in Fig. 2 (the two sites are similar in the guesses). Guess 2: initial guess with the two sites very different as used for the results in Figs 3–10. In the first 3 columns of values, the sites are assumed to be independent so the values obey the constraints defined in eqns (9) and (10). True 2: the true rates used for simulation of non-independent binding sites (Figs 14, 15, 17 and 18). Guess 3: initial guesses used for Figs 14 and 15 (similar values used for Fig. 17) Guess 4: initial guesses used for Fig. 18.

made in earlier studies. It implies imposition of the following three constraints:

$$k_{-2a} = k_{-1a}, \quad k_{-2b} = k_{-1b}, \quad k_{+2b} = k_{+1b}. \quad (9)$$

These, together with the microscopic reversibility constraint, assure also that:

$$k_{+2a} = k_{+1a}. \quad (10)$$

When a single low-concentration record is fitted in bursts (see Methods), there is no information available about how frequently the channel is activated, so whether or not the above constraints are applied, it is necessary to supply more information in order to get a fit. This was done in two ways. Either (a) one of the rate constants (k_{+2a}) was fixed at an arbitrary value such as $10^8 M^{-1} s^{-1}$ (the effects of error in this value are investigated below), or (b) an EC_{50} value was specified, and used to calculate one of the rate constants (see Methods). In either case the number of free parameters is reduced to nine.

Initial guesses based on the two binding sites being similar. As with any iterative fitting method, initial guesses for the free parameters have to be supplied. It is always

important to check that the same estimates are obtained with different initial guesses. It is quite possible, if the fit is very insensitive to the value of one of the rate constants, for convergence to be obtained with the initial guess being hardly changed. This does not mean that it was a good guess, but merely that the data contain next to no information about that particular rate constant; it is easy to get spurious corroboration of one's prejudices. And in a complex problem like this it is quite possible that the likelihood surface will have more than one maximum; a bad guess may lead you to the wrong maximum. This problem can be illustrated by what happens when attempts are made to start the fit of scheme 1 with guesses that make the two binding sites almost the same, when in fact they are different. In general it seems like quite a good idea to start from a 'neutral' guess like this, but in practice it can give problems. (Note, too, that all the calculations assume that eigenvalues are distinct, so it is inadvisable to start with guesses that are identical.)

Figure 2 shows the distributions of 1000 estimates obtained from fitting a single record at a low ACh concentration, 30 nM, with the constraints in eqns (9) and

(10), and with $k_{+1a} = k_{+2a}$ fixed at $1 \times 10^8 \text{ M}^{-1} \text{ s}^{-1}$ (half its true value in this case). The resolution imposed before fitting was $25 \mu\text{s}$, as in most experiments. The fitting was done in bursts of openings that corresponded to individual activations of the channel, defined by $t_{\text{crit}} = 3.5 \text{ ms}$, and the likelihood calculation for each burst was started and ended with CHS vectors (see Methods). In this case the rate constants, rather than their logarithms, were the free parameters. The true rate constants (those used for the simulation) are shown in Table 1, and the initial guesses for the fitting are shown as 'guess 1' in column 3 of Table 1.

On each histogram of the 1000 estimates, the true value is marked with an arrow. The distribution of the estimates of

α_2 , the shutting rate for diliganded channels, in Fig. 2A has two peaks. One, shown enlarged in the inset, is close to the true value of $\alpha_2 = 2000 \text{ s}^{-1}$. This peak contains 73 % of all estimates and these have a mean of $2045 \pm 174 \text{ s}^{-1}$ (coefficient of variation 8.5 %), so these estimates have a slight positive bias but are quite good. The other 27 % of estimates of α_2 are much bigger, nowhere near the true value. A similar picture is seen with the estimates of β_2 shown in Fig. 2B. Again 73 % of estimates (the same 73 %) are near the right value, $\beta_2 = 52000 \text{ s}^{-1}$, and the other 27 % are much too big. The main peak has a mean of $52736 \pm 3692 \text{ s}^{-1}$, the coefficient of variation being 7.0 %, slightly lower than for α_2 .

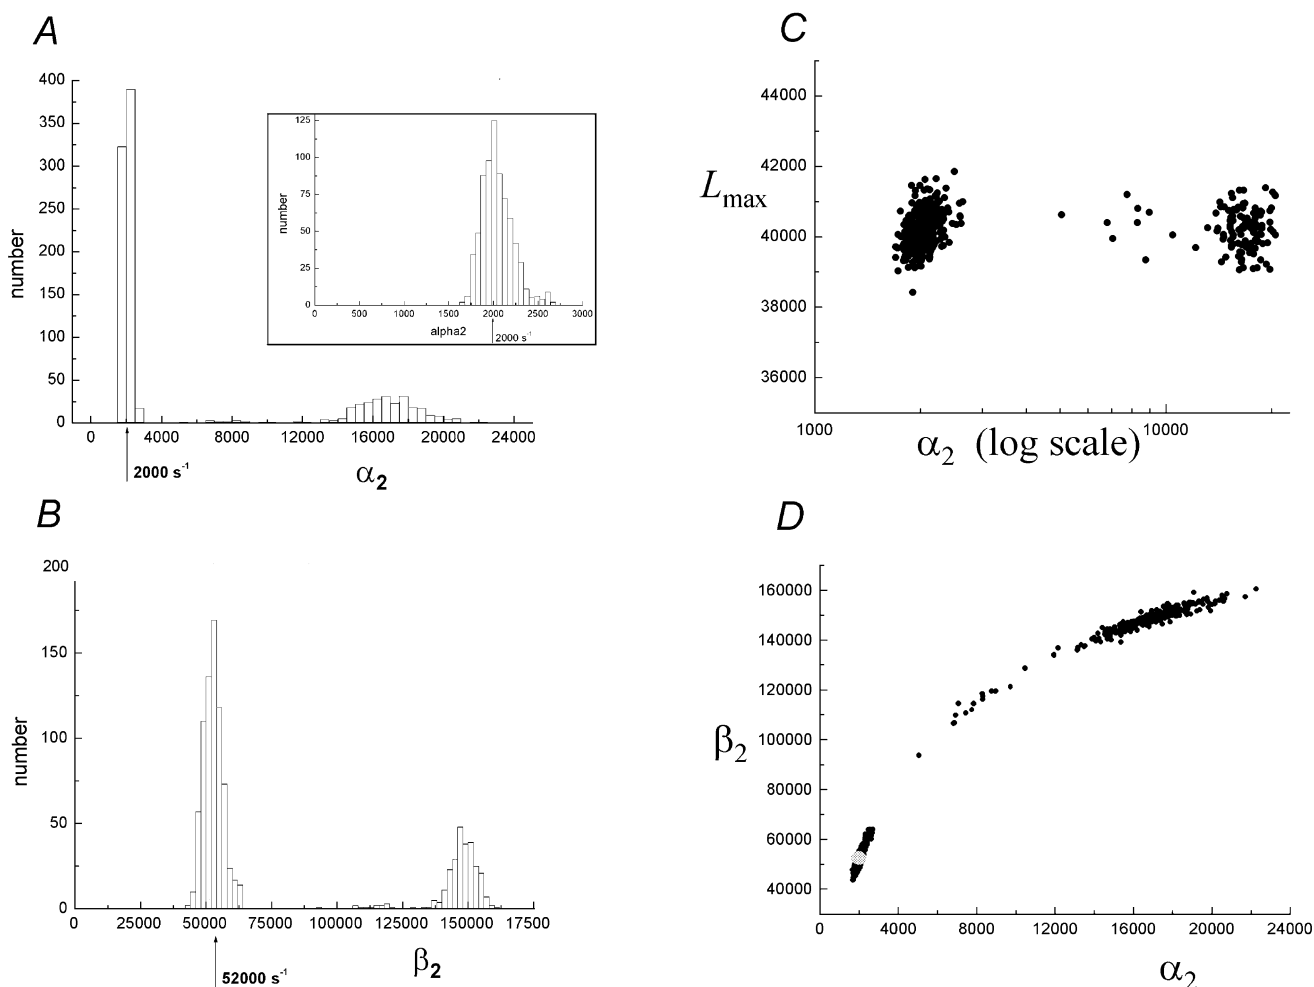


Figure 2. Distributions of the 1000 estimates of rate constants found by fitting, with HJCFIT, to 1000 simulated experiments

The arrows mark the true value of the rate constants. The simulation used the true rate constants given in Table 1 ('true 1'), and each fit started from 'guess 1' (Table 1). Each 'experiment' consisted of 20 000 transitions (about 9000 transitions after the resolution of $25 \mu\text{s}$ was imposed), at a single low concentration, 30 nM, of ACh. The fitting was done in bursts of openings defined by a $t_{\text{crit}} = 3.5 \text{ ms}$, with CHS vectors (see Methods). The sites were assumed to be independent (eqn (9)), and $k_{+1a} = k_{+2a}$ was fixed at $1 \times 10^8 \text{ M}^{-1} \text{ s}^{-1}$ (half its true value in this case). A, distribution of 1000 estimates of α_2 . The inset shows the region near 2000 s^{-1} enlarged. B, estimates of the 1000 estimates of β_2 from the same fits as used for A. C, the maximum likelihood (L_{max}) attained in each of the 1000 fits plotted against the value of α_2 from that fit D, the value of α_2 from each of the 1000 fits plotted against the value of β_2 from the same fit. The pale circle marks the true values.

Figure 2C shows that there is essentially no difference between the 'goodness of fit', as measured by the maximum value of the log-likelihood attained in the 'experiments' that gave good estimates, and those that gave estimates that were much too fast. All of the fits fall clearly into either the 'right solution' or into the 'fast solution' peaks, apart from 10 or so (1%) that are smeared in between the two main peaks. This behaviour resembles the very simplest version of the missed event problem, which is known to have two solutions (see Discussion).

When the estimate of α_2 is plotted against the value of β_2 from the same fit, in Fig. 2D, it is clear that the two values are very strongly correlated – the fits that give good estimates of α_2 also give good estimates of β_2 , and vice versa. This phenomenon will be discussed below.

Initial guesses based on the two binding sites being different. When similar experiments are simulated, but with initial guesses for the fit that start from the supposition that the two binding sites are not similar, these better guesses very rarely lead to the incorrect 'fast solution'. The guesses used for each of the 1000 fits are shown in column 4 of Table 1 ('guess 2'). The results are shown in Figs 3–5.

In this case none of the 1000 fits converged on the incorrect 'fast solution'. The mean of 1000 estimates of α_2 was $2016.5 \pm 146.4 \text{ s}^{-1}$, compared with a true value of 2000 s^{-1} (Fig. 3A). The coefficient of variation (CV) of the estimates is 7.3% and there is a very slight positive bias of +0.82% (calculated as a fraction of the true value). For β_2 the mean was $52285 \pm 3248 \text{ s}^{-1}$, compared with a true




Figure 3. Distribution of estimates of rate constants, and of quantities derived from them, for 1000 fits to experiments simulated as in Table 1 (starting from 'guess 2')

The arrows mark the true values of the rate constants. Each 'experiment' consisted of about 9000 transitions at a single low concentration, 30 nM, of ACh. The constraints in eqns (9) and (10) were true for the mechanism used for simulation, and were applied during the fit. In these fits $k_{+1a} = k_{+2a}$ was fixed arbitrarily at $1 \times 10^8 \text{ M}^{-1}\text{s}^{-1}$, half its true value. Distribution of 1000 estimates of A, α_2 ; B, β_2 ; C, $E_2 = \beta_2/\alpha_2$; D, $k_{-2a} + k_{-2b}$. The graph in E shows the 1000 pairs of α_2 and β_2 values plotted against each other to show the positive correlation between them ($r = +0.916$, see Table 3).

value of $52\,000\text{ s}^{-1}$ (Fig. 3B). The CV was 6.2%, and bias +0.55%. Again the estimates of α_2 and β_2 show a positive correlation (Fig. 3E), though over the narrower range of values found here it is much more nearly linear than seen in Fig. 2D. The ratio of these two quantities, $E_2 = \beta_2/\alpha_2$, represents the efficacy for diliganded openings (Colquhoun, 1998). Because of the strong positive correlation between α_2 and β_2 , this ratio is better defined than either rate constant separately. The 1000 estimates of E_2 shown in Fig. 3C have mean of 25.96 ± 32 (true value 26), so their CV is 2.9% with an insignificant bias of -0.16% . The total dissociation rate of agonist from diliganded receptors, $k_{-2a} + k_{-2b}$, was also well-defined. The distribution of 1000

estimates shown in Fig. 3D has a mean of $11\,463 \pm 573\text{ s}^{-1}$, compared with a true value of $11\,500\text{ s}^{-1}$. The CV was 5.0%, and bias -0.32% . This is somewhat more precise than the two separate values, $k_{-2a} = k_{-1a}$ (CV = 12%, bias = 1.8%) and $k_{-2b} = k_{-1b}$ (CV = 6.0%, bias = 0.1%) (see Fig. 4F,H). In this example the negative correlation between these two values was modest ($r = -0.274$) so their sum is more precise than their separate values to a correspondingly modest extent.

The parameters for singly liganded receptors are generally less precisely estimated than those for diliganded receptors, especially when unconstrained (see below), but

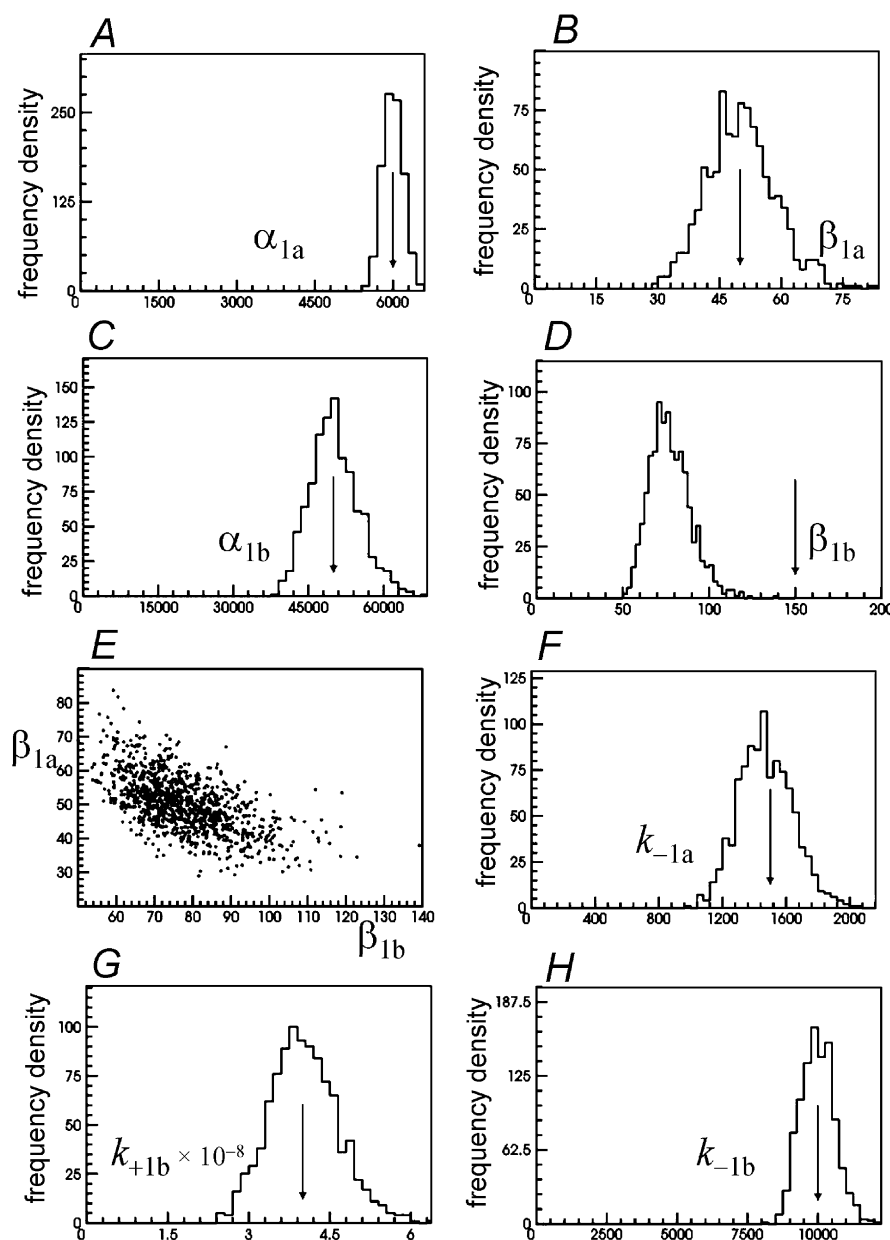


Figure 4. Distributions of the 'one site' rate constants for the same set of 1000 fits as in Fig. 3. The arrows mark the true values of the rate constants.

quite reasonable estimates can be found if the constraints in eqns (9) and (10) are true, as in the present case. Figure 4 shows the distributions of the estimates of the other free parameters for the same simulations as those shown in Fig. 3. These are the singly liganded opening and shutting rates, α_{1a} , β_{1a} , α_{1b} and β_{1b} , and the binding rate constants, k_{-1a} , k_{-1b} and k_{+1b} . In these fits $k_{+1a} = k_{+2a}$ was fixed arbitrarily at $1 \times 10^8 \text{ M}^{-1}\text{s}^{-1}$, half its true value. It can be seen that the estimates of all of these parameters are tolerably good, apart from β_{1b} , which is, on average about half of its true value. This happens because $k_{+1a} = k_{+2a}$ was fixed at half of its true value; if we fix k_{+2a} at its true value, $2 \times 10^8 \text{ M}^{-1}\text{s}^{-1}$, then good estimates of β_{1b} are found too. It is natural to ask, why it is primarily the estimate of β_{1b} that is affected by an error in the fixed value of $k_{+1a} = k_{+2a}$? There is a good intuitive reason for this happening. Inspection of the expressions for the equilibrium occupancies for scheme 1 (Fig. 1) shows that the relative

frequencies of the two sorts of singly-liganded openings is given by:

$$f_1 = \frac{\beta_{1a}k_{+2a}k_{-2b}}{\beta_{1b}k_{+2b}k_{-2a}} \tag{11}$$

Furthermore, the frequency of openings with both sites occupied, relative to the frequency with only the *a* site occupied is given by:

$$f_{2a} = \frac{\beta_2k_{+1b}c}{\beta_{1a}k_{-1b}} \tag{12}$$

where *c* is the agonist concentration, and the corresponding relative frequency when only the *b* site is occupied is given by:

$$f_{2b} = \frac{\beta_2k_{+1a}c}{\beta_{1b}k_{-1a}} \tag{13}$$




Figure 5. Distributions of equilibrium constants for the same fits as in Figs 3 and 4

A–D, distributions of the four equilibrium constants (E_{1a} , E_{1b} , K_a , K_b) that refer to the two separate sites, calculated from the fitted rate constants (as shown Figs 3 and 4). The arrows mark the true values of the rate constants. E, the negative correlation between k_{-2a} ($= k_{-1a}$) and β_{1b} ; F, the positive correlation between k_{+1b} ($= k_{+2b}$) and β_{1a} .

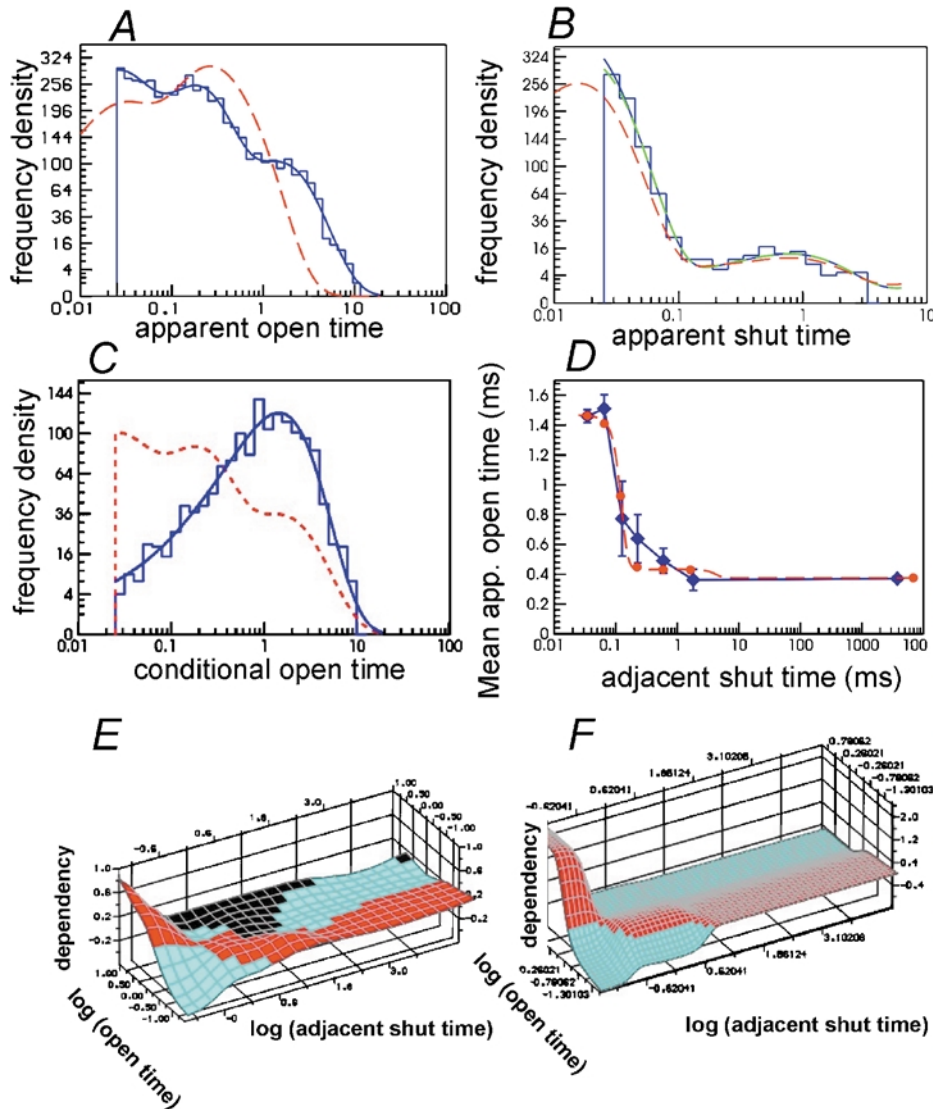


Figure 6. Tests of the quality of the fit produced in a single simulated experiment of the same sort as the 1000 simulations used to generate Figs 3–5

A, histogram of distribution of apparent ('observed') open times with a resolution of $25 \mu\text{s}$. The solid blue line shows the HJC open time distribution predicted by the fitted values of the rate constants with a resolution of $25 \mu\text{s}$. Dashed red line: the open time distribution predicted by the fitted values of the rate constants without allowance for limited resolution – the estimate of the true open time distribution. B, as in A but for apparent shut times (note that only shut times up to $t_{\text{crit}} = 3.5 \text{ ms}$ are used for fitting so only these appear in the histogram). The HJC distribution is, as always, calculated from the exact expressions up to $3t_{\text{res}}$ (i.e. up to $75 \mu\text{s}$ in this case), and thereafter from the asymptotic form. Green line: the asymptotic form plotted right down to $t_{\text{res}} = 25 \mu\text{s}$. It is seen to become completely indistinguishable from the exact form for intervals above about $40 \mu\text{s}$. C, histogram of distribution of 'observed' open times with a resolution of $25 \mu\text{s}$ for openings that are adjacent to short shuttings (durations between 25 and $100 \mu\text{s}$). Solid line: the corresponding HJC conditional open time distribution predicted by the fitted values of the rate constants with a resolution of $25 \mu\text{s}$. Dashed line: the HJC distribution of all open times (same as the solid line in A). D, conditional mean open time plot. The solid diamonds show the observed mean open time for apparent openings (resolution $25 \mu\text{s}$) that are adjacent to apparent shut times in each of seven shut time ranges (plotted against the mean of these shut times). The bars show the standard deviations for each mean. The shut time ranges used were (ms) $0.025\text{--}0.05$, $0.05\text{--}0.1$, $0.1\text{--}0.2$, $0.2\text{--}2$, $2\text{--}20$, $20\text{--}200$ and > 200 . Note, however, that shut time greater than t_{crit} (3.5 ms) will be shorter than predicted if there was more than one channel in the patch so the values on the abscissa above 3.5 ms are unreliable. The solid circles show the corresponding HJC conditional mean open times predicted by the fitted values of the rate constants with a resolution of $25 \mu\text{s}$, for each of the seven ranges. The dashed line shows the continuous HJC relationship between apparent mean open times conditional on being adjacent to the shut time specified on the abscissa. E, 'observed' dependency plot for apparent open times and adjacent shut times (resolution $25 \mu\text{s}$). Regions of

The fit is sensitive to the values of these ratios of opening frequencies (in this particular case the open states are not connected to each other, so they are simply the ratios of the areas of the three components of the open time distribution). All three ratios will be unaffected by a decrease in the value of $k_{+1a} = k_{+2a}$, if, at the same time, β_{1b} is reduced by the same factor. Attempting to compensate for a reduction in $k_{+1a} = k_{+2a}$ in other ways does not work. For example a concomitant increase in β_{1a} in eqn (11) can keep f_1 unchanged, but will result in changes in eqns (12) and (13). It is only by decreasing β_{1b} that the predicted relative frequencies of the three sorts of openings will be unchanged.

Figure 5A–D shows the distributions of the equilibrium constants, calculated for each of the 1000 fits from the rate constants shown in Figs 3 and 4. Figure 5A and B shows the two ‘efficacies’ for singly liganded openings, E_{1a} ($= \beta_{1a}/\alpha_{1a}$) for when only the *a* site is occupied, and E_{1b} ($= \beta_{1b}/\alpha_{1b}$) for when only the *b* site is occupied. The estimates are tolerable apart from the bias caused by specification of an incorrect value for $k_{+1a} = k_{+2a}$. The equilibrium constants for binding to *a* and *b* sites, K_a and K_b , are shown in Fig. 5C and D (the notation K_a can be used because the constraints imply that $K_{1a} = K_{2a}$ and similarly for the *b* site). Apart from the bias caused by specification of an incorrect value for $k_{+1a} = k_{+2a}$, the estimates are not too bad (CV = 11.6% for K_a but larger (CV = 17.5%) for K_b). The plot in Fig. 5E shows that there is quite a strong negative correlation ($r = -0.74$) between the estimates of $k_{-1a} = k_{-2a}$ and of β_{1b} . Figure 5F shows a stronger positive correlation ($r = +0.92$) between the estimates of $k_{+1b} = k_{+2b}$ and of β_{1a} . Correlations of this magnitude are a sign of ambiguity in the separate values of the parameters concerned.

The quality of the fit obtained in a single simulated experiment. Figures 3–5 showed the distributions of 1000 estimates of rate constants. In practice, experiments are analysed one at a time, and after the estimates of the rate constants have been obtained, the extent to which they describe the observations is checked. Figure 6 shows examples of these checks in the case of a single experiment that was simulated under exactly the same conditions as were used to generate Figs 3–5. More details of these plots are given in Methods (see Checking the quality of the fit).

Notice that the fit looks excellent despite the 2-fold error in the (fixed) value of $k_{+1a} = k_{+2a}$, and the consequent error in β_{1b} . Figure 6A–C shows the data as histograms, for (A) all open times, (B) all shut times and (C) open times that are adjacent to short (up to 100 μ s) shut times. On each of these

histograms, the solid line that is superimposed on (not fitted to) the data is the appropriate HJC distribution calculated from scheme 1 using the values of the rate constants that were obtained for the fit and the imposed time resolution of 25 μ s. The fitting was done as described for Figs 3–5. The HJC distributions (solid blue lines in Fig. 6A–C), were, as always, calculated from the exact expressions up to $3t_{res}$ (i.e. up to 75 μ s in this case), and thereafter from the asymptotic form. The green line in Fig. 6B shows the asymptotic form plotted right down to $t_{res} = 25 \mu$ s. It is seen to become completely indistinguishable from the exact value for intervals above about 40 μ s, thus justifying the claim that the calculations are essentially exact. For the apparent open times in Fig. 6, the exact and asymptotic were hardly distinguishable right down to 25 μ s (see Hawkes *et al.* 1992, for more details).

In Fig. 6A and B, the red dashed line shows the estimate of the ideal distribution (no missed events) calculated from the fitted rate constants (see Methods for details). It is clear from Fig. 6A that the apparent open times are greatly extended by the failure to detect many brief shuttings.

The conditional distribution in Fig. 6C shows that short openings very rarely occur adjacent to short shuttings (the dashed line shows the HJC distribution of all open times longer than 25 μ s: see Methods).

Figure 6D shows a conditional mean open time plot. The diamond symbols show the data. Each represents the mean apparent open time for openings that are adjacent to shut times within a specified range. Seven shut time ranges were specified (see legend) and the means of the open times (blue diamonds) are plotted with their standard deviations (bars). The HJC predictions (calculated from the fitted rate constants and a resolution of 25 μ s, as in Colquhoun *et al.* 1996) are shown, for the same ranges, as red circles. The dashed red line shows the theoretical continuous relationship between mean open time and apparent shut times, but this cannot be used directly as a test of fit, because shut time ranges must be used that are wide enough to encompass a sufficient number of observations.

Figure 6E and F shows the observed and the predicted dependency plot, respectively, for the same ‘experiment’ (see Methods). The dependency plot calculated from the fitted rate constants by the HJC method (Fig. 6F) shows that the shortest apparent shut times are much more likely to occur adjacent to long apparent openings than next to short openings, and that long apparent shut times are predicted to be rather more common adjacent to short shut times. The ‘observations’ (Fig. 6E) are qualitatively similar, but exact comparison is difficult with 3D plots,

Table 2. Approximate standard deviations obtained from the Hessian matrix in a single fit, compared with the values calculated directly from 1000 fits

	Estimate	SD from one run	SD from 1000 fits
α_2	2006.8	174	146
β_2	53 158.5	3891	3248
α_{1a}	5825.8	197	193
β_{1a}	39.8	6.6	8.4
α_{1b}	46 364.1	4031	4839
β_{1b}	99.1	14.4	11.9
k_{-1a}	1231.8	148	171
k_{-1b}	9917.7	632	596
k_{+1b}	3.26×10^8	0.53×10^8	0.64×10^8

Table 3. Approximate correlations obtained from the Hessian matrix in a single fit, compared with the values calculated directly from 1000 fits

	α_2	β_2	α_{1a}	β_{1a}	α_{1b}	β_{1b}	k_{-1a}	k_{-1b}	k_{+1b}
α_2									
β_2	0.937 0.916								
α_{1a}	-0.043 -0.006	-0.068 -0.049							
β_{1a}	0.013 0.042	0.040 0.075	-0.033 0.011						
α_{1b}	-0.055 0.054	-0.095 0.002	0.489 0.447	-0.018 0.000					
β_{1b}	-0.107 -0.087	-0.062 -0.079	-0.025 -0.023	-0.606 -0.596	0.281 0.399				
k_{-1a}	-0.007 0.028	0.011 0.055	0.018 0.025	0.763 0.799	0.047 0.029	-0.791 -0.736			
k_{-1b}	-0.542 -0.456	-0.525 -0.451	0.193 0.188	-0.192 -0.261	0.064 0.075	0.366 0.363	-0.219 -0.275		
k_{+1b}	-0.106 -0.070	-0.097 -0.064	-0.040 0.027	0.915 0.916	-0.106 -0.079	-0.588 -0.585	0.669 0.700	0.028 -0.035	

In each cell, the upper value for the correlation coefficient is calculated from the covariance matrix for a single fit, and the lower value is found from 1000 pairs of estimates.

and a large number of observations is needed to get a smooth 3D plot.

The quality of internal estimates of variance and correlation. In the last fit of the set of 1000 shown in Figs 3–5, the Hessian matrix was calculated as described in Methods. The approximate standard deviations for the parameter estimates, and the correlations between pairs of estimates, were compared with the values measured directly from the 1000 fits. The values are shown in Tables 2 and 3.

There is good general agreement between the errors and correlations that are predicted in this particular ‘experiment’ and the values actually found by repetition of the experiment 1000 times. The calculation of errors via the Hessian matrix thus produces, at least in this case, a good prediction of what the real errors and correlations

will be. Of course, in real life it is not so easy to repeat an experiment under exactly the same conditions. When experiments are repeated at different times, and with different batches of cells, we (Gibb *et al.* 1990; Hatton *et al.* 2003) and others (e.g. Milone *et al.* 1997; Bouzat *et al.* 2000) have often found quantitative differences between repeated experiments that are beyond what would be expected from experimental error.

Use of an EC_{50} value as a constraint. The fixing of a rate constant at an arbitrary value (as in Figs 3–5) is obviously an unsatisfactory solution to the problem of the patch containing an unknown number of channels. In real life we do not know the true value of a rate constant, and there are two ways to circumvent this problem. One is to fit simultaneously results at several different concentrations (see below). Another is to use an independently determined

EC_{50} value to constrain the missing rate constant, (see Methods). The EC_{50} for the true rates in Table 1 is $3.3 \mu M$. Rather than fixing $k_{+1a} = k_{+2a}$ at an arbitrary value, its value is calculated at each iteration from the specified EC_{50} plus the values of the other rate constants.

When 1000 fits were done, like those shown in Figs 3–5, but with k_{+2a} calculated from the (correct) EC_{50} ($3.3 \mu M$), reasonable estimates were obtained for all nine free rate constants, including β_{1b} , for which the mean of all 1000 estimates was $158.4 \pm 43.6 s^{-1}$ (true value $150 s^{-1}$). The results for all the rate constants, with the specified EC_{50} being the correct value, are shown in Figs 7 and 8.

As always, the rate constants for the diliganded receptor are better defined than those that refer to the two separate sites, but even the worst estimates are tolerable. This applies to the binding equilibrium constants for the two binding sites too, which are quite scattered. The distribution of K_a (Fig. 8E) has a CV = 21.8 % and bias = -1.4 %, and the distribution of K_b (Fig. 8F) has a CV = 17.3 % and

bias = -1.3 %. However these two quantities show a strong (though not linear) negative correlation (Fig. 8G). Therefore it is not surprising that their product, $K_a K_b$, is rather more precisely determined, as shown in Fig. 8H, which has a CV = 6.35 % and bias -1.3 %. It is this product that occurs in those terms that refer to diliganded receptors in the expressions for equilibrium state occupancies.

The success of this procedure depends, of course, on having an accurate value for the EC_{50} , undistorted by desensitisation (unless desensitisation is part of the mechanism to be fitted). In general it will be best if the EC_{50} can be determined from a one-channel Popen curve determined under conditions similar to those used for the HJCFIT data. To test the effects of using an incorrect EC_{50} , the simulations were repeated but using an EC_{50} that was half, or double, the correct value.

When an EC_{50} of $6.6 \mu M$ (twice its correct value) was used, most of the parameters were still estimated quite well. The exceptions were β_{1b} , and $k_{+1a} = k_{+2a}$, both of which were

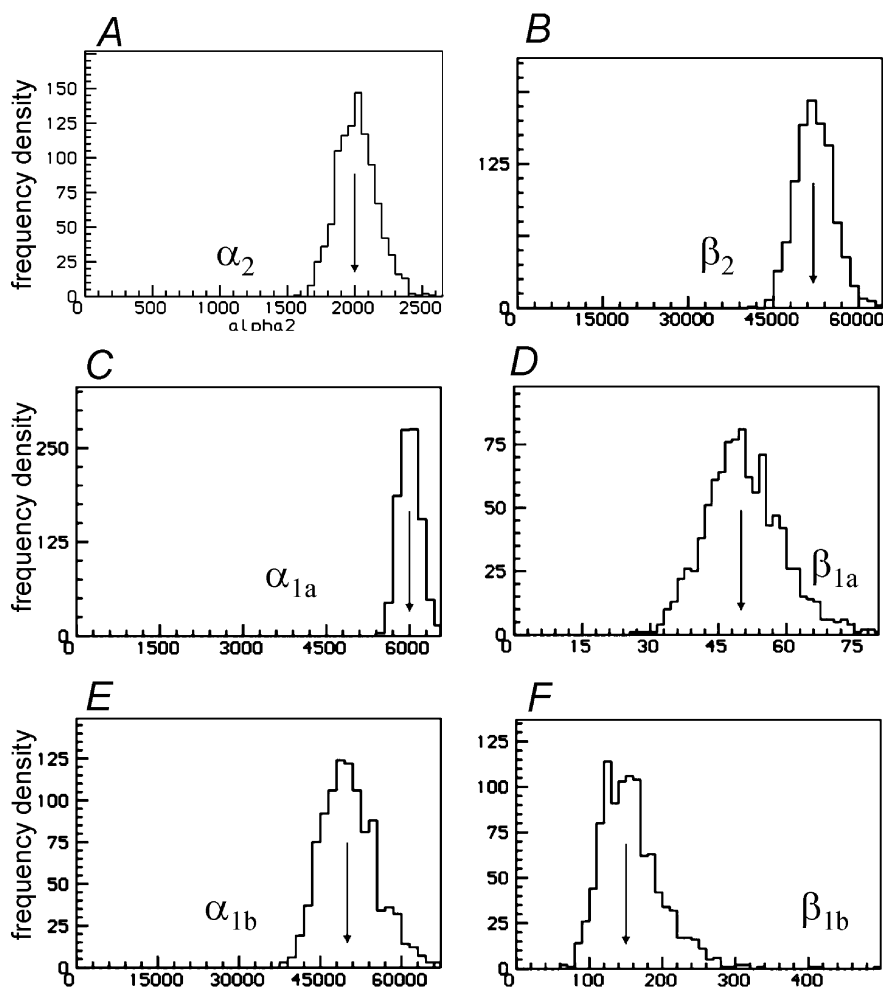


Figure 7. The six gating rate constants

Single low concentration of ACh (30 nM) fitted in bursts ($t_{\text{crit}} = 3.5 \text{ ms}$), with the two sites constrained to be independent. No parameters were fixed but $k_{+1a} = k_{+2a}$ was calculated from the other rates so as to give the specified (correct) EC_{50} . The arrows mark the true values of the rate constants.

too small, by factors of 4.0 and 4.3 respectively, as shown in Fig. 9.

When an EC_{50} of $1.65 \mu\text{M}$ (half its correct value) was used, the errors were worse. The estimates of the 'diliganded parameters', α_2 , β_2 and total dissociation rate, $k_{-2a} + k_{-2b}$, were still very good, as were the estimates of α_{1a} and α_{1b} (data not shown). The distributions of the estimates of the other parameters were all centred on means that were more or less incorrect. The largest errors were again in β_{1b} ,

and $k_{+1a} = k_{+2a}$, both of which were too big on average, by factors of 3.4 and 4.4 respectively. The means for the other rate constants were too big on average by factors that varied from 0.95 for $k_{-2b} = k_{-1b}$, to 1.35 for β_{1a} . Some of the results are shown in Fig. 10.

Correlations between parameter estimates. The correlation between estimates of two different parameters is a purely statistical phenomenon. It has already been illustrated in Figs 2D, 3E, 4E, 5E, 5F and 8G. If the

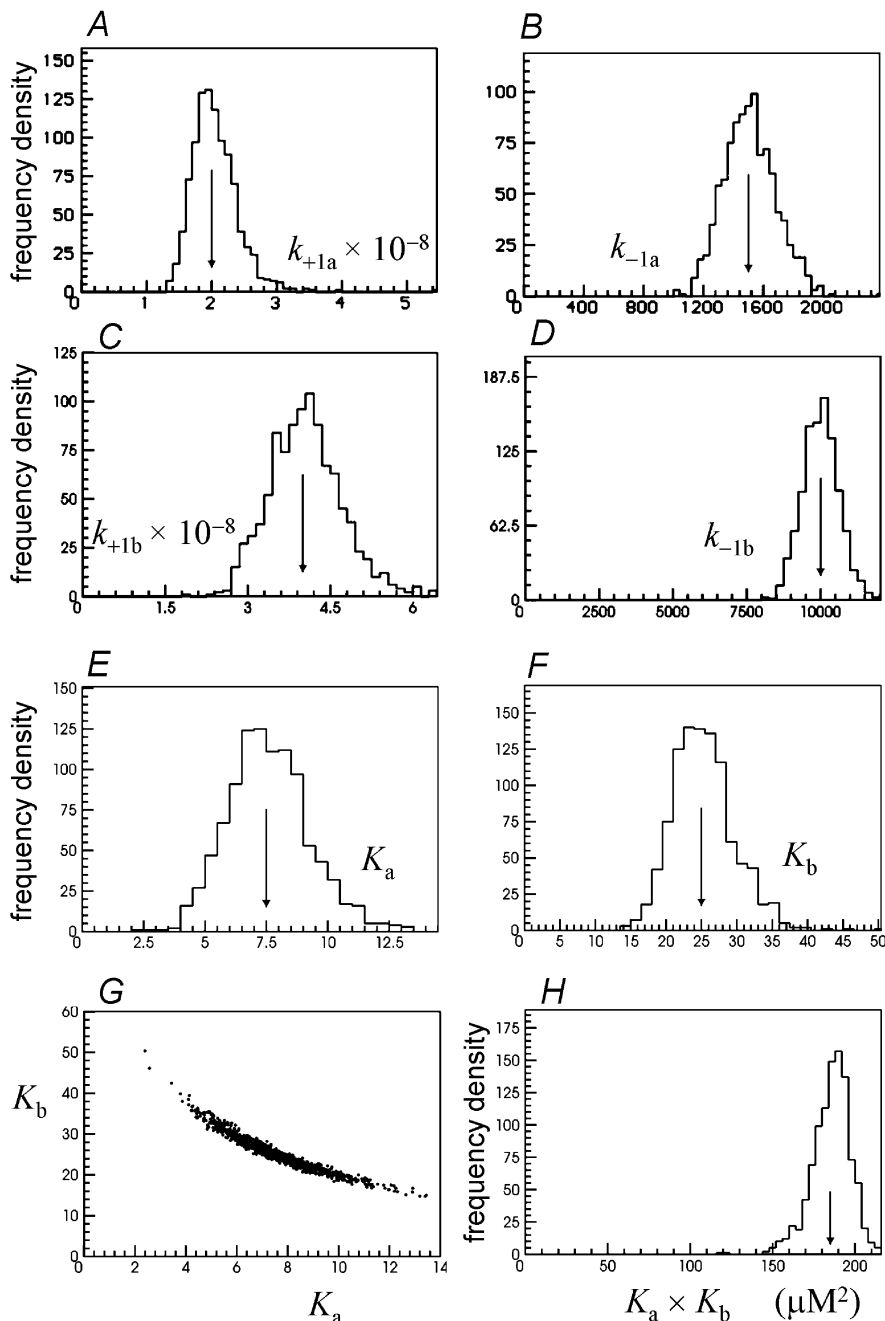


Figure 8. The same set of fits, with constrained EC_{50} , as shown in Fig. 7

A–D, distributions of the four binding rate constants, and E, F, two derived equilibrium constants, K_a and K_b . G, plot of K_b against K_a for each of the 1000 fits; H, distribution of the product $K_a K_b$. The arrows mark the true values of the rate constants.

estimates are precise enough the correlations vanish. It is quite distinct from the correlation between, for example, adjacent open and shut times (see Fig. 6C–F) which is a physical property of the mechanism, and gives interesting information about it (e.g. Fredkin *et al.* 1985; Colquhoun & Hawkes, 1987). The statistical correlation between parameter estimates resembles the negative correlation seen between repeated estimates of the slope and intercept of a straight line, or the positive correlation seen between the EC_{50} and maximum when fitting a Langmuir binding curve. It is merely a nuisance that limits the speed and accuracy of the fitting process. The correlation can be seen, in the form of a correlation coefficient, from the calculation of the covariance matrix (see Methods), as illustrated in Table 3.

Figure 11 shows in graphical form the correlations between all possible pairs of parameters, for the set of simulated fits shown in Figs 7 and 8. They are arranged as in the correlation matrix shown in Table 3.

The effect of the strong correlation between the estimates of α_2 and β_2 on the fitting process is illustrated in Fig. 12, for an experiment on wild type human receptor (30 nM ACh, see Hatton *et al.* (2003). In this case the correlation coefficient between estimates of α_2 and β_2 was $r = 0.915$, a typical value. The likelihood surface is in 10-dimensional space, and so cannot be represented. Fig. 12A shows a 3D 'cross section' of the actual likelihood surface that was constructed by calculating the likelihood for various

values of α_2 and β_2 , with the seven other free parameters fixed at their maximum likelihood values.

The correlation appears as a diagonal ridge (coloured pink). Along this ridge, the values of α_2 and β_2 change roughly in parallel (so the efficacy, $E_2 = \beta_2/\alpha_2$, does not change much), and the likelihood increases only slowly towards its maximum (marked red). Figure 12B shows a contour representation of the same surface near its maximum. Dashed lines show the coordinates of the maximum point, the maximum likelihood estimates being $\alpha_2 = 1524 \text{ s}^{-1}$ and $\beta_2 = 50\,830 \text{ s}^{-1}$. The contours are shown also for $\log(\text{likelihood})$ values of $L = L_{\text{max}} - 0.5$ and $L = L_{\text{max}} - 2.0$. The tangents to these contours provide 0.5- and 2.0-unit likelihood intervals for the estimates of α_2 and β_2 (these correspond roughly to one and two standard deviations, but being asymmetrical they provide better error estimates: see Colquhoun & Sigworth, 1995).

The effect of this correlation on the fitting process is illustrated in Fig. 13.

The vertical axis gives the likelihood that corresponds to the values of α_2 and β_2 that are reached at various stages during the fitting process. The initial guess is marked at the bottom of the graph, and the likelihood increases during the course of the fit. At first the increase is rapid but there is a long final crawl along a diagonal ridge near the maximum. This involves many changes of direction and slows the fitting process considerably, not least with the simplex method employed in HJCFIT. In this case the rate

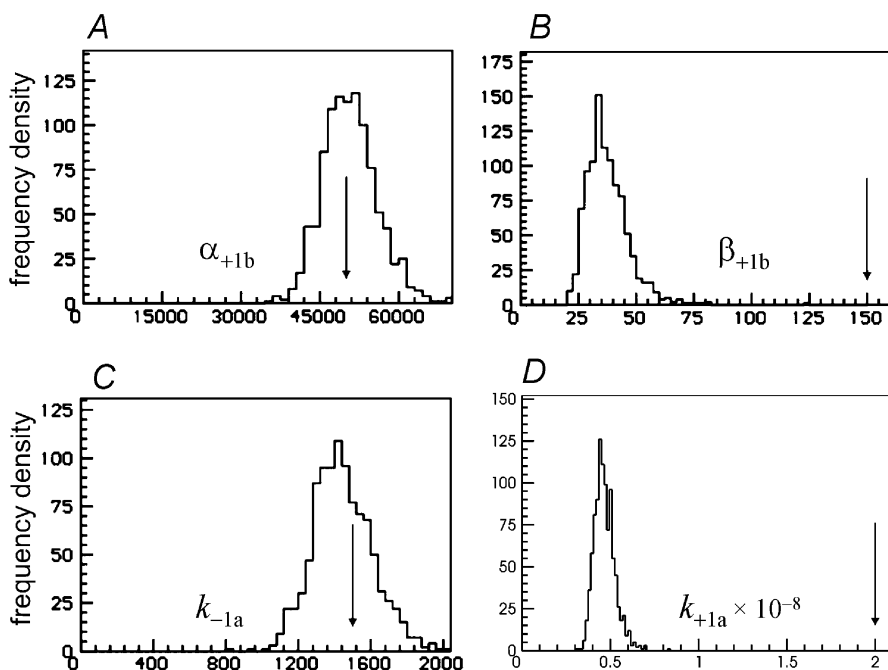


Figure 9. Single low concentration of ACh (30 nM) fitted as in Figs 7 and 8, apart from specification of an incorrect value for the EC_{50} , twice its correct value

The arrows mark the true values of the rate constants.

constants, not their logarithms, were used as the free parameters. However fitting the logarithms of the rates (see Methods) speeds up the fit and speed is not a problem in practice.

The effects of fitting as though the binding sites were independent when they are not. It is quite possible to obtain good fits to low-concentration data even if it is assumed incorrectly that the binding sites are independent. The rate constants in Table 1 (labelled 'true 2') were used to simulate 1000 experiments. These rates represent sites that interact (see Hatton *et al.* 2003). The microscopic equilibrium constant for binding to the *a* site when the *b* site is vacant, $K_{1a} = k_{-1a}/k_{+1a} = 20 \mu\text{M}$, but for binding to the *a* site when the *b* site is occupied $K_{2a} = k_{-2a}/k_{+2a} = 240 \mu\text{M}$, so binding at the *a* site has a lower affinity if the *b* site is occupied; there is negative cooperativity in the binding of agonist to the shut channel (see Jackson, 1989 and Hatton *et al.* 2003). Likewise for binding to the *b* site $K_{1b} = k_{-1b}/k_{+1b} = 0.33 \mu\text{M}$, but when the *a* site is occupied $K_{2b} = k_{-2b}/k_{+2b} = 4 \mu\text{M}$. Again there is negative cooperativity in the binding of agonist to the shut channel.

These values were used to simulate the experiments, but during the fit, the (inappropriate) constraints in eqns (9) and (10) were applied. The initial guesses shown in Table 1 ('guess 3') also obeyed these constraints. A single low (30 nM) concentration was used and k_{+1a} (assumed, incorrectly, to be the same as k_{+2a} was constrained to give the specified EC_{50} (9.697 μM , its correct value). The results were fitted in bursts ($t_{\text{crit}} = 3.5 \text{ ms}$), with CHS vectors (see Methods). Although good fits could be obtained to the distributions of apparent open and shut times, many of the parameter estimates were quite wrong, as shown in Figs 14 and 15.

Figure 14A and B shows that the estimates of rates constants in a single fit (actually the last of the 1000 fits) predict well the distributions of apparent open time, and apparent shut time. Figure 14C shows that the conditional open time distribution, for openings that are adjacent to the shortest shut times (25–100 μs), is also predicted well. However the fact that something is wrong is shown, in this case, by the dependence of mean open time on adjacent shut time (Fig. 14D). Although the prediction of the fit is

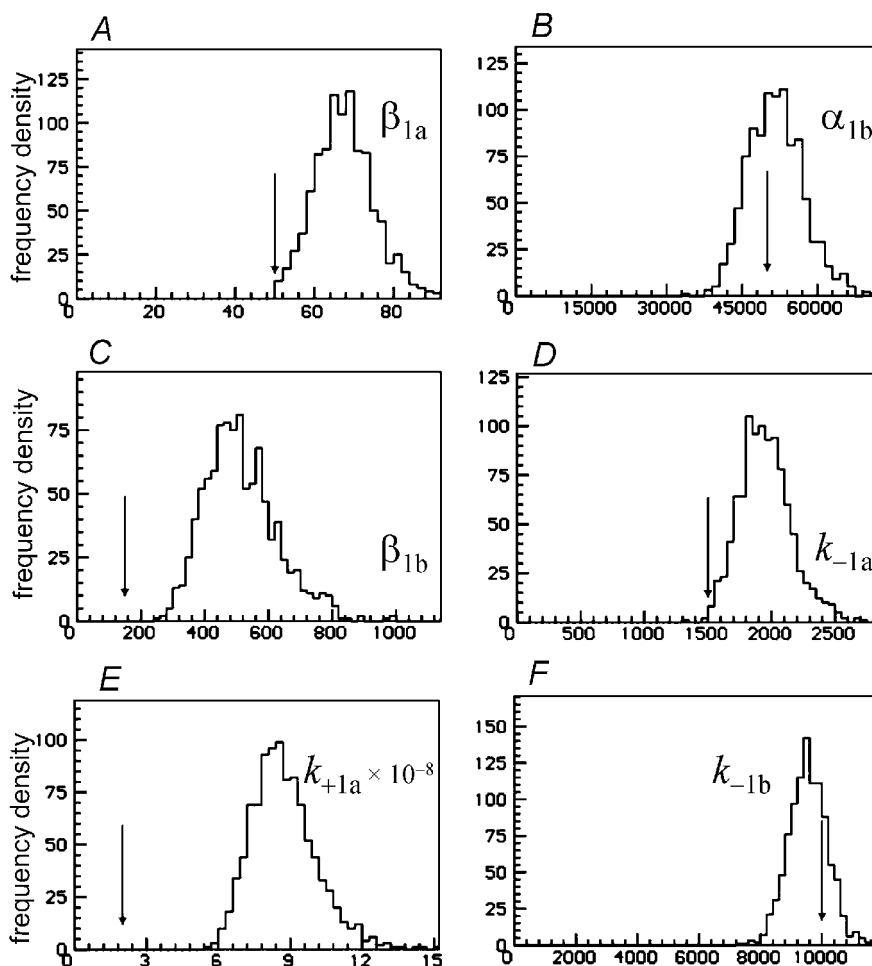


Figure 10. Single low concentration of ACh (30 nM) fitted as in Figs 7 and 8, apart from specification of an incorrect value for the EC_{50} , half its correct value

The arrows mark the true values of the rate constants.

quite good for the shortest shut times (as shown also in Fig. 14C), and for the longest shut times, the prediction is quite bad for shut times between about 0.3 and 30 ms. This is also visible in the conditional apparent open time distribution shown in Fig. 14E. This shows the distribution of apparent open times that are adjacent to shut times in the range 0.5–10 ms, and the predicted fit is bad.

Examples are shown in Fig. 15 of the distributions of rate constants obtained in 1000 fits that were done under the same conditions as the single fit shown in Fig. 14. Despite the grossly incorrect assumptions (and the somewhat subtle indication of imperfect fit shown in Fig. 14D and E), the estimates of the ‘diliganded’ rate constants, α_2 and β_2 are nevertheless quite good (Fig. 15A and B). The estimates of the total dissociation rate from diliganded receptors,

$k_{-2a} + k_{-2b}$, was estimated reasonably well too (Fig. 15C), though with some bias (true value, $14\,000\text{ s}^{-1}$, mean of 1000 estimates $14\,900\text{ s}^{-1}$ with a CV of 6.3% and bias +6.4%). However, as might be expected, the rate constants that refer to the two separate sites are not well-estimated, being anything from poor to execrable. The estimates for α_{1a} and α_{1b} were poor (bias +61% and +12% respectively), but the estimates of β_{1a} (shown in Fig. 15D) and β_{1b} were worse (bias -63%, CV 30% for β_{1a} ; bias +135%, CV 12.1% for β_{1b}), and the estimates of the association and dissociation rates were inevitably very poor. For example the estimates of k_{-1a} (true value 400 s^{-1}) and of k_{-2a} (true value 2000 s^{-1}) were constrained by the fit to be the same, and had a mean slightly below either true value, 367 s^{-1} ; this distribution is shown in Fig. 15E and F (on two

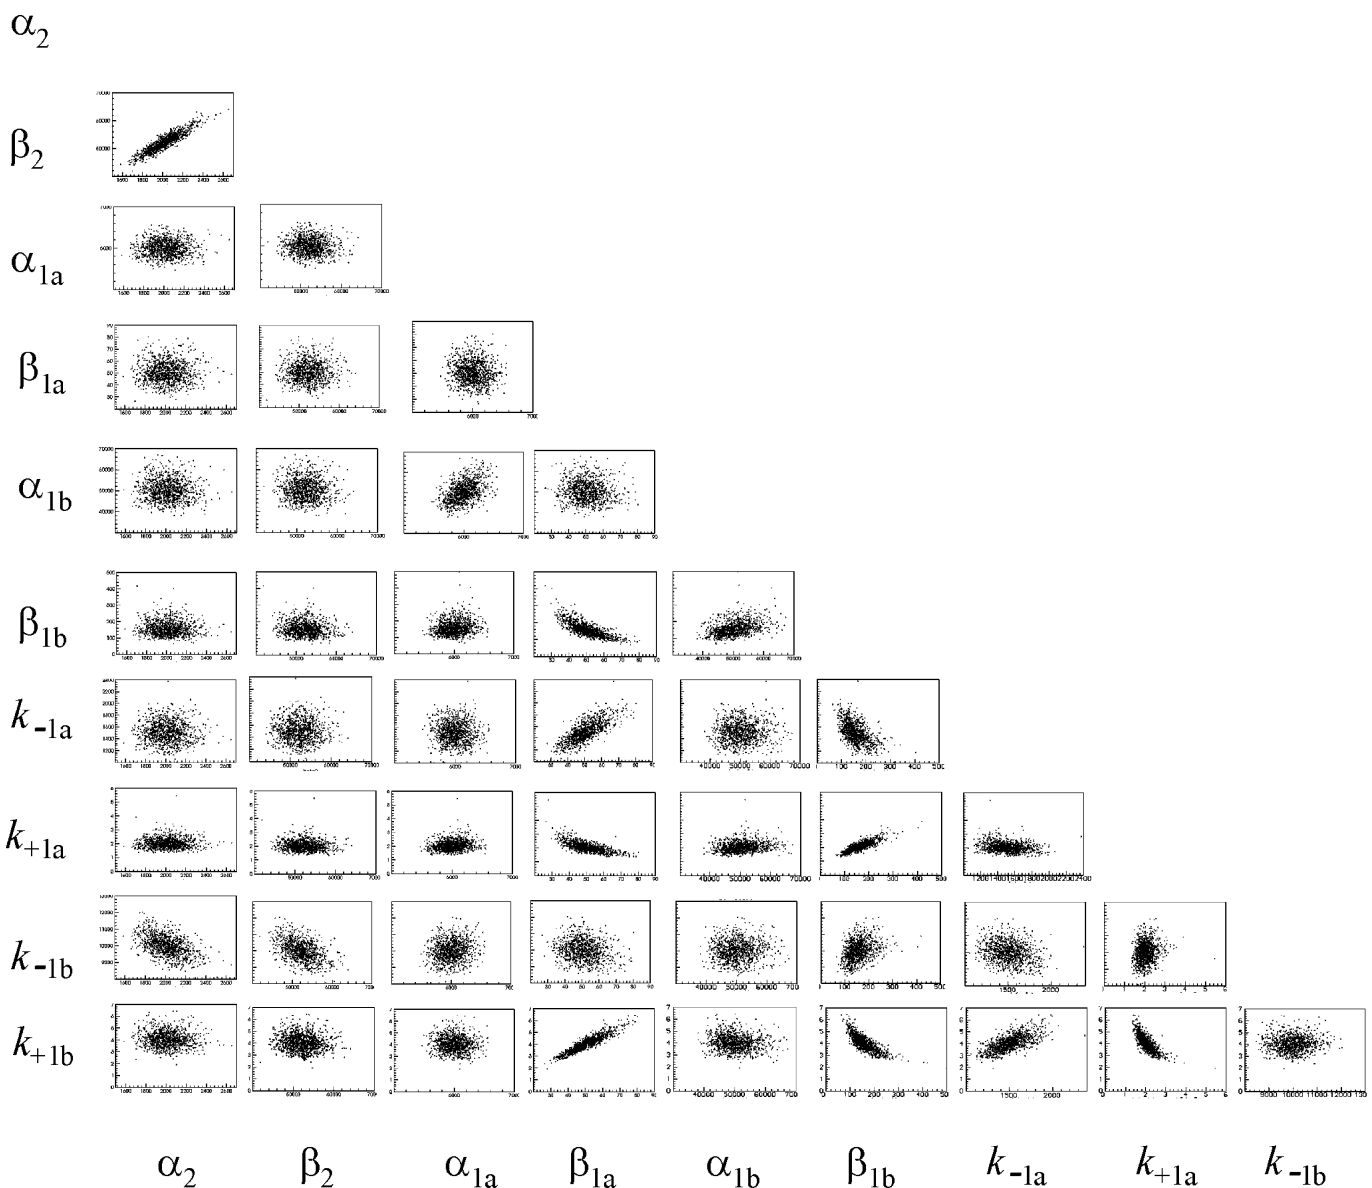


Figure 11. Correlation matrix shown graphically for the 1000 fits illustrated in Figs 7 and 8
Fitted values are plotted for the 45 possible pairs of parameters.

different scales, to allow display of the arrow that indicates the true values of k_{-1a} (Fig. 15E) and of k_{-2a} (Fig. 15F).

Simultaneous fits of records at more than one concentration

It has been shown above, that when the two binding sites are independent, all of the rate constants can be estimated quite well from a single low concentration experiment. Since in such experiments it will usually not be known how many channels were present in the patch it is necessary to

fit the record in bursts, and this means that information about the absolute frequency of channel activations is missing. Nevertheless all of the rate constants can be found if an EC_{50} value can be specified, as shown in Figs 7–10. Another way to obtain information about the absolute frequency of channel activations is to use high agonist concentrations (see Methods); indeed this will usually be the best way to obtain information about the EC_{50} . The method (Sakmann *et al.* 1980) works only if the channel shows the right amount of desensitisation to allow

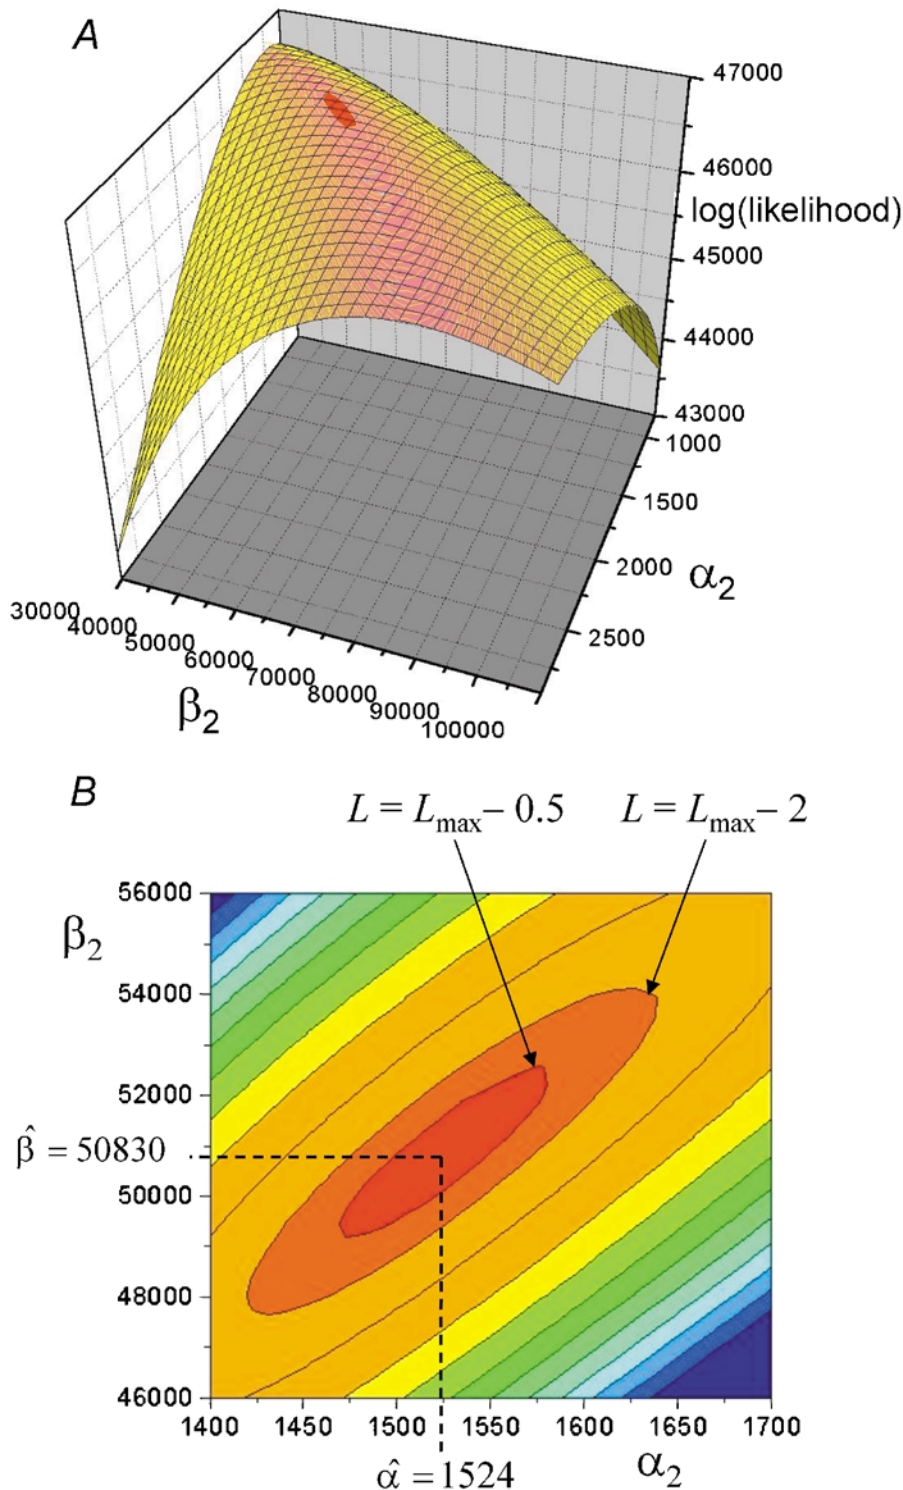


Figure 12. The shape of the likelihood surface near its maximum

A, the surface shows the likelihood for various values of the shutting rate α_2 , and the opening rate β_2 , for the doubly occupied receptor. In order to plot this surface, the seven other free parameters were fixed at their maximum likelihood values, and the likelihood was calculated for various values of α_2 and β_2 . B, contour representation of the surface shown in A, near the maximum. Dashed lines show the coordinates of the maximum point, the maximum likelihood estimates being $\alpha_2 = 1524 \text{ s}^{-1}$ and $\beta_2 = 50830 \text{ s}^{-1}$. The contours are shown also for log(likelihood) values of $L = L_{\max} - 0.5$ and $L = L_{\max} - 2.0$.

definition of clear clusters of openings that all originate from the same channel. When this is the case, the upper part of a concentration– P_{open} curve can be constructed (e.g. Sine & Steinbach, 1987; Colquhoun & Ogden, 1988). At high concentrations there will be few singly liganded openings, so the rate constants for their opening and shutting cannot be determined if high concentrations only are used. However simultaneous fit of a low concentration (fitted in bursts) and a high concentration (fitted on the assumption that only one channel is active) allows good estimates to be obtained for all ten free constants, without the need to specify an EC_{50} . The program HJCFIT is designed for such simultaneous fitting.

The case where the binding sites are independent

A set of 1000 fits was done in which the likelihood was maximised simultaneously for two separate simulated experiments, one at a low concentration, 30 nM, and one at a high concentration, 10 μM (not far from the EC_{50}). The high concentration data were simulated using scheme 2 (Fig. 1), with the true rate constants shown in Table 1 and desensitisation rate constants set to $\beta_D = 5 \text{ s}^{-1}$ and $\alpha_D = 1.4 \text{ s}^{-1}$. This produces a record with long desensitised periods (mean $1/\alpha_D = 714 \text{ ms}$) and clusters of openings that last 240 ms on average, and contain 400 openings on average (though there will appear to be fewer after the resolution of 25 μs has been imposed). This is similar to what is observed in experimental records, though desensitisation is more complex than suggested in scheme 2 (Elenes & Auerbach, 2002). The low concentration was analysed in bursts ($t_{\text{crit}} = 3.5 \text{ ms}$) as above, each burst corresponding with a single activation, and CHS vectors were used (see Methods). The high concentration was also analysed in

groups ($t_{\text{crit}} = 5 \text{ ms}$, no CHS vectors), but in this case the groups represented the long ‘desensitisation clusters’. Both sets of simulated data were fitted with scheme 1 (Fig. 1). This was not the same as the scheme 2 that was used to simulate the data, but the procedure of analysis in groups (bursts, clusters) excises the desensitised periods. Thus this simulation also provides a test of the common procedure of excising desensitised periods from a high concentration record and then fitting what remains with a mechanism that does not include desensitised states.

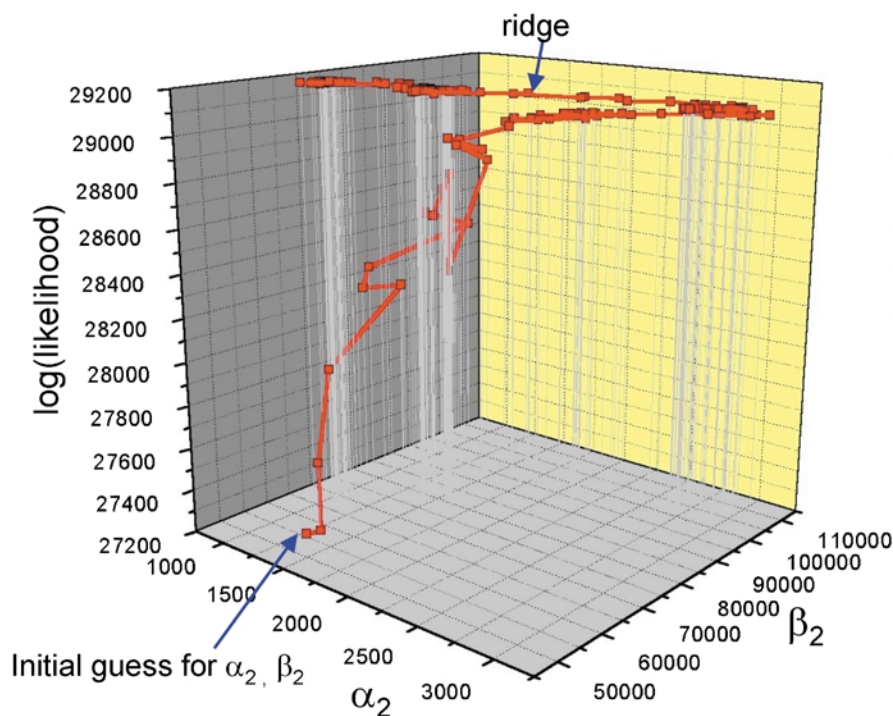
The results showed that the simultaneous fitting of high and low concentration records gave good estimates of all 10 free rate constants, without having to resort to fixing one or to providing an EC_{50} (distributions not shown). It also showed that excision of desensitised periods is a satisfactory procedure. The estimates of the ‘diliganded rate constants’ were excellent. The coefficient of variation (CV) for α_2, β_2 , and for the total dissociation rate of agonist from diliganded receptors, $k_{-2a} + k_{-2b}$, were all below 5% (and for $E_2 = \beta_2/\alpha_2$ the CV was 1.7%), all with bias below 0.8%. For the ‘monoliganded rate constants’ the CV varied from 3.3% for α_{1a} and 3.6% for $k_{-1b} = k_{-2b}$, up to 21% for β_{1b} . After exclusion of 16/1000 ‘experiments’ that gave outlying estimates of β_{1b} , the estimates for all 10 free rate constants had CVs between 3.2% (for β_2) and 14.1% (for β_{1b}), and bias between about 0.2% (for $\beta_2, \alpha_{1a}, k_{-1b} = k_{-2b}$, and $k_{+1b} = k_{+2b}$), up to 2.7% for β_{1b} .

Can the two sites be distinguished using high concentrations alone?

The advantage of using a high concentration is that quite long stretches of record can be obtained that originate from one channel only, at least for the muscle nicotinic

Figure 13. The progress of fitting in one example

The likelihood (vertical axis) that corresponds with the values of α_2 and β_2 that are reached at various stages during the fitting process – notice the final crawl along a diagonal ridge near the maximum. Along the ridge, the values of α_2 and β_2 change roughly in parallel (so the efficacy, $E_2 = \beta_2/\alpha_2$, does not change much), and the likelihood increases only slowly.



and glycine receptors which have the right amount of desensitisation. The disadvantage is that there are few singly liganded openings. It is quite common for channel properties to be analysed by use of high concentrations only, so it is natural to ask what can and cannot be inferred from such records. To do this, 1000 fits were done on simulated data obtained with an agonist concentration of $10 \mu\text{M}$, which is of the order of the EC_{50} . The simulations were done with the mechanism in scheme 2 (Fig. 1), with $\beta_D = 5 \text{ s}^{-1}$ and $\alpha_D = 1.4 \text{ s}^{-1}$, as before, and the true values of the other rate constants shown in Table 1 ('true 1'). The results were fitted with scheme 1 (Fig. 1) after removing the desensitised periods by fitting in bursts ($t_{\text{crit}} = 5 \text{ ms}$), without CHS vectors. The two sites were independent. As usual the 'diliganded' parameters were well defined: for α_2 , CV = 6.8% and bias = 2.3%; for β_2 , CV = 7.0% and bias = 1.1%; and for total dissociation rate from diliganded receptors, $k_{-2a} + k_{-2b}$, CV = 4.7% and bias = -0.6%. In a

single 'experiment' out of the 1000, the predicted fit of the apparent open and shut times was good, as shown in Fig. 16A and B. The apparent open time distribution (Fig. 16A) was essentially a single exponential, as expected from the lack of singly liganded openings at high concentration: over 99% of the area was in a component with a time constant of 1.43 ms in the asymptotic HJC distribution, and this fits the histogram. In the predicted ideal distribution (dashed line in Fig. 16A and B), 99% of the area is in a component with time constant of 0.342 ms, much shorter because of missed brief shuttings (the true value would be 0.5 ms, but the fitted value of α_2 was unusually high in this particular experiment, 2920 s^{-1} , rather than 2000 s^{-1}).

In view of the absence of singly liganded openings at $10 \mu\text{M}$ it is not surprising that the estimates of α_{1a} , β_{1a} , α_{1b} and β_{1b} are all undefined in this case – the values are all over the

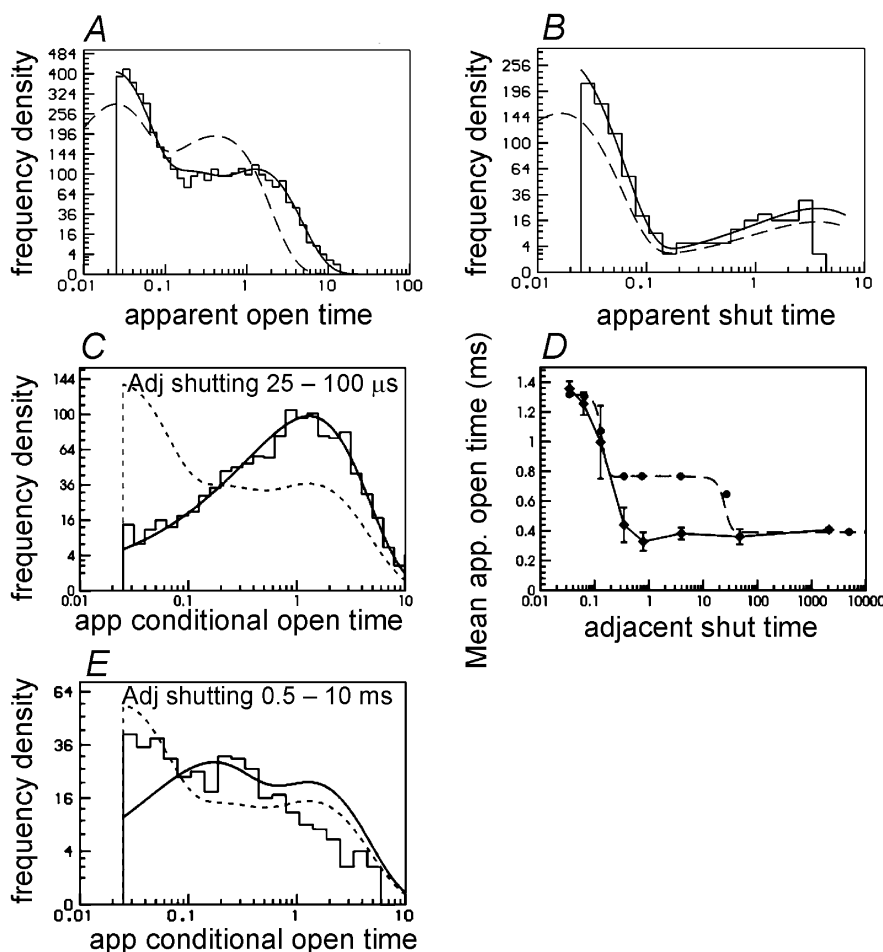


Figure 14. Predicted fits to a simulated experiment, when the fit assumed, incorrectly, that the two sites were independent

A single low concentration of ACh (30 nM) was used, with EC_{50} constraint and fitted in bursts. The rate constants used for simulation are in Table 1 (labelled 'true 2'), but the fit applied the constraints in eqns (9) and (10). A and B, predicted fits to apparent open and shut time, as in Fig. 6. C, conditional distribution of open times for apparent openings that are adjacent to the shortest shut times ($25 - 100 \mu\text{s}$). D, observed and predicted conditional mean open time plot (as in Fig. 6D and Methods). E, as in C but for apparent open times that are adjacent to shut times in the range $0.5 - 10 \text{ ms}$.

place. On the other hand, the association and dissociation rate constants for the two sites separately are defined, though the estimates are not very precise (Fig. 16C–E). For k_{-1a} (Fig. 16C), CV = 34%, bias = +10%. For k_{+1a} (Fig. 16D), CV = 9.6%, bias = +2.5%. For k_{-1b} (Fig. 16E), CV = 6.3%, bias = -2.2% and for k_{+1b} (Fig. 16F), CV = 15%, bias = +2.3%. The equilibrium constants calculated from these rates are K_a (Fig. 16G, CV = 32%, bias = +6.7%) and K_b (Fig. 16H, CV = 16%, bias = -2.3%). With this amount of scatter a 2-fold difference in equilibrium affinities of the two sites would be barely detectable.

Non-independent binding sites

Up to now, the two binding sites have always been assumed to be different from each other, but independent of one another. If it is allowed that the binding of the agonist to one site can affect the binding to another site, so the constraints in eqns (9) and (10) can no longer be applied, there are 13 free rate constants to be estimated, rather than 10. In Figs 14 and 15, the effects were investigated of fitting as though the two sites were

independent when they are not. We now describe attempts to fit all 13 rate constants in the case where the sites are not independent.

In every case that has been investigated so far it has proved impossible to get good estimates of all 13 rate constants under conditions that can be realised in practice (this is true, at least, for the values of the rate constants used here). In particular, they cannot be obtained under conditions where the number of channels that are present in a low-concentration record is unknown. At high concentrations it is possible to obtain long stretches of record that are known to contain only one channel, but high concentration records alone do not contain enough information about singly liganded states to allow estimation of all 13 rate constants. The same is true if high concentration records are fitted simultaneously with low concentration records, the latter being fitted in bursts because of the unknown number of channels in the patch. Fixing one of the rate constants, or determining one of them from a known EC_{50} , does not help either, and improving the resolution from 25 μ s to 10 μ s




Figure 15. Distributions of some of the rate constant estimates obtained from 1000 fits under the same conditions as Fig. 14

The fits assumed, incorrectly, that the two sites were independent. The arrows mark the true values of the rate constants.

does not solve the problem. In most of these simulations, the problem lay in separating the values of k_{+1a} and k_{+1b} , the estimates of all the other rate constants being good or at least acceptable. The estimates of these two rate constants were very smeared, with a tendency to approach unreasonably large values, no doubt as a result of the strong positive correlation seen between them (their ratio was better determined). Values for k_{-1b} were poor too. Despite the near-useless estimates of at least two of the rate constants, good predictions of the data (such as those shown in Figs 6 and 14)

could be obtained, as might be expected from the large number of free parameters.

Examples of poorly estimated parameters are shown in Fig. 17. These are from an attempt to fit all 13 parameters with a large amount of high-resolution data, three low concentrations (10 nM, 30 nM and 100 nM) fitted simultaneously, with a resolution of 10 μ s. However all three records were fitted in bursts ($t_{\text{crit}} = 3.5$ ms), with CHS vectors, to avoid any assumption about the number of channels in the patch. All the parameters were reasonably

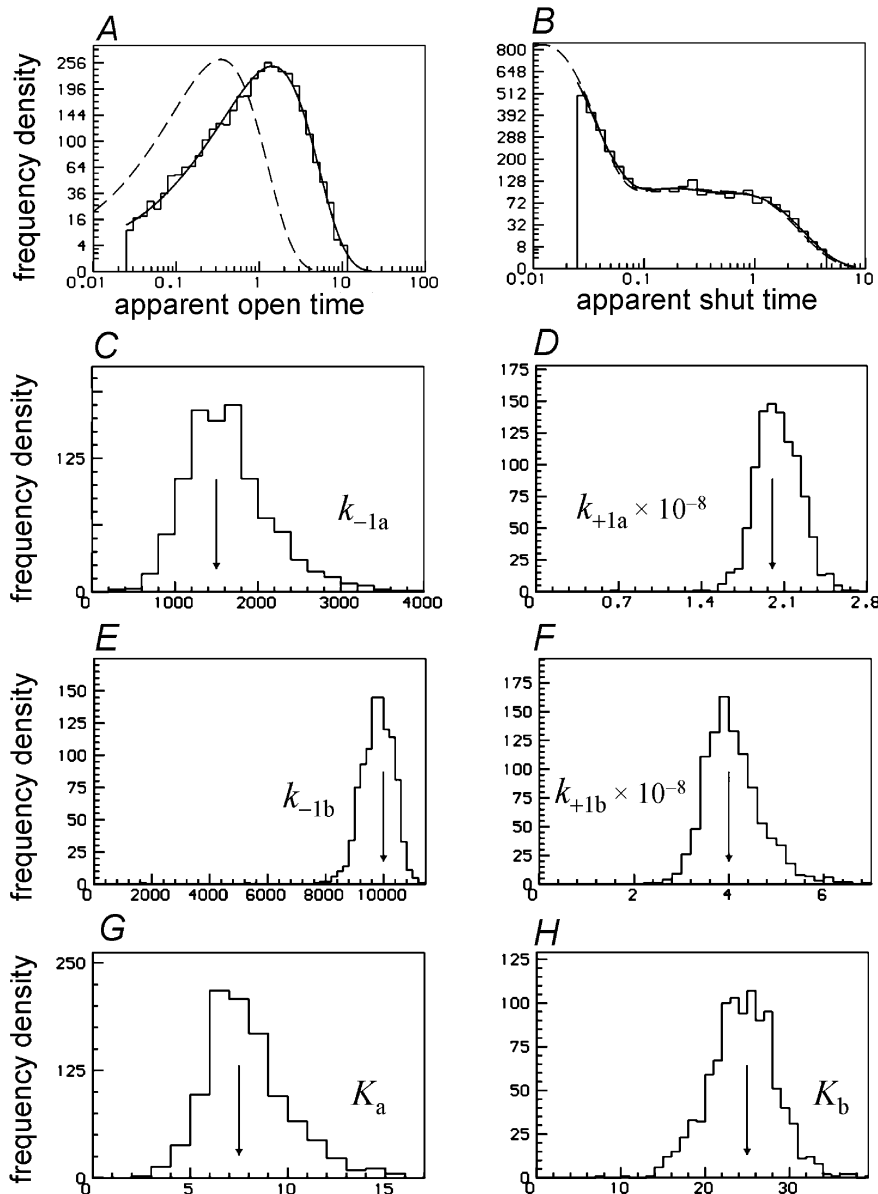


Figure 16. Analysis of 1000 fits to simulated experiments at a single high (10 μ M) concentration of ACh

The arrows mark the true values of the rate constants. A, apparent open times from a single simulated experiment, with the HJC distribution predicted by the fit superimposed on the histogram (resolution 25 μ s); the dashed line is the predicted true distribution. B, apparent shut times (up to $t_{\text{crit}} = 5$ ms) as in A. C–F, distributions of 1000 estimates of the association and dissociation rate constants. G–H, distributions of the 1000 estimates of the microscopic equilibrium constants, K_a and K_b , for binding to the two sites (calculated from the fitted rate constants).

estimated apart from the three shown: k_{+1a} (Fig. 17A) and k_{+1b} (Fig. 17C) are very smeared, and k_{-1b} (Fig. 17B) is poor. This is a reflection of the strong positive correlation between k_{+1a} and k_{+1b} , (Fig. 17D); the correlation coefficient was +0.61 in this plot (which is curtailed by an upper limit of $10^{10} \text{ M}^{-1} \text{ s}^{-1}$ placed on any association rate constant during the fit. The correlation coefficient between estimates of the equilibrium constants for the first bindings, K_{1a} and K_{1b} , was very strong indeed (+0.99995 in this case).

The only way in which it has proved possible, so far, to get good estimates of all 13 rate constants, is by simultaneous fit of records at two concentrations, with the assumption that only one channel is present at the low concentration(s) as well as at the high concentration(s). In other words the low concentration records are not fitted in bursts, but the likelihood is calculated from the entire sequence, including all shut times. Good estimates of all 13 free rate constants could be obtained by simultaneous fit of two low concentrations (10 nM and 100 nM), or by simultaneous fit of a low concentration (30 nM) and a high concentration (10 μM), as long as it was supposed that the low concentration record(s) originated from one channel. In this case the entire shut time distribution is predicted by the fit, not only shut times up to t_{crit} . Fits predicted from one such simulated experiment are shown in Fig. 18A–D.

The distributions of both apparent open times (Fig. 18A and C) and of apparent shut time (Fig. 18B and D) are predicted well by a single set of rate constants at both concentrations. The shut time distributions (Fig. 18B and D) include all shut times (above t_{res}), and the longer time between activations at the lower concentration is obvious. The estimates of all 13 rate constants found in 1000 such fits were good, but the only ones shown are the distributions of k_{+1a} , k_{-2b} and k_{+1b} (Fig. 18E–G). The estimates of these are now quite good, whereas the estimates of these rates shown in Fig. 17A–C (with more and better data, but fitted in bursts) were bad. The estimates of k_{+1a} and k_{+1b} are now essentially independent (Fig. 18H), rather than strongly correlated (Fig. 17D). The set of 1000 fits exemplified in Fig. 18 gave CVs of about 5% or less for α_2 , β_2 , α_{1a} , α_{1b} , k_{-2a} , ($k_{-2a} + k_{-2b}$), and for the corresponding equilibrium constant E_2 (all with bias less than 0.4%); CVs of about 5–10% were found for β_{1b} , k_{+2a} , k_{-1a} , k_{+1a} , k_{+1b} , and for the corresponding equilibrium constants E_{1b} , K_{2a} , K_{1b} (all with bias less than 1%). The least precise estimates were for β_{1a} , (CV = 11.5%, bias = +0.84%), k_{-1b} (CV = 11.4%, bias = +0.96%), k_{-2b} (CV = 18.7%, bias = -1.2%), k_{+2b} (CV = 17.6%, bias = -0.7%), and for E_{1a} (CV = 12.8%, bias = +1.3%), K_{2b} (CV = 13.1%, bias = +0.1%), K_{1a} (CV = 12.6%, bias = +1.2%).




Figure 17. Non-independent sites

Estimates of k_{+1a} (A), k_{-1b} (B) and k_{+1b} (C) from 1000 simulated fits using simultaneous fit of three concentrations of ACh (10 nM, 30 nM and 100 nM), with a resolution of 10 μs . All three records were fitted in bursts ($t_{\text{crit}} = 3.5 \text{ ms}$), using CHS vectors. D, plot of the estimate of k_{+1b} against that of k_{+1a} from the same fit (an upper limit of $10^{10} \text{ M}^{-1} \text{ s}^{-1}$ was placed on association rate constants in these fits). The arrows mark the true values of the rate constants.

DISCUSSION

Diliganded openings

The first major conclusion from this study is that good estimates of the rate constants that relate to diliganded openings can be obtained under all the circumstances that

have been tested. This is true if an association rate or EC_{50} is fixed at an incorrect value, if desensitisation is present in the data but is not fitted, and it is true even if the fitting process assumes that the two binding sites are independent when they are not (at least for the examples tested here). The analyses all assume a resolution of 25 μ s, and they




Figure 18. Non-independent sites

Illustration of the prediction of the distribution of apparent open and shut times by the rate constants estimated by a single simultaneous fit to two low concentration of ACh (10 nM and 100 nM) simulated experiments. In this case it was assumed that one channel was present throughout both records so the entire shut time distribution is fitted (B and D). The two sites were not assumed to be independent in this case, and good estimates were obtained for all 13 free rate constants (see text). In particular the three rate constants which could not be estimated when the number of channels was not assumed (Fig. 17A–C) are now estimated well (E–G), and the estimates of k_{+1a} and k_{+1b} are now essentially uncorrelated (H). The arrows mark the true values of the rate constants.

make no assumptions about the number of channels in the patch. Our results suggest that the estimate of the opening rate constant for the diliganded channel, β_2 , is at least as good as that for its shutting rate constant, α_2 , so there is, using the HJCFIT method, no need to estimate the former by extrapolation as is usual with the MIL method (see, for example, Salamone *et al.* 1999). Estimates of the total dissociation rate from diliganded receptors ($k_{-2a} + k_{-2b}$) are of similarly good quality. We thus cannot agree with Akk & Steinbach (2000) when they describe β_2 and binding rates as being so fast that it is 'close to impossible to evaluate them independently'.

These three quantities alone are sufficient to describe the characteristics of the receptor that are of physiological importance. In particular, they allow calculation of a good approximation to the mean length of the diliganded burst of openings, and it is these that carry most of the current through the endplate membrane at the neuromuscular junction. The current carried by monoliganded openings is very much smaller and quite negligible from the physiological point of view. For diliganded channel activations, the mean number of channel openings per activation will be, to a good approximation (neglecting returns from a singly-liganded state):

$$\mu_{\text{ob}} = 1 + \frac{\beta_2}{k_{-2a} + k_{-2b}}. \quad (14)$$

The mean open time is $1/\alpha_2$, so the mean open time per burst will be:

$$\mu_{\text{tob}} = \frac{\mu_{\text{ob}}}{\alpha_2}. \quad (15)$$

Because the large majority of shut times within an activation are very short sojourns in A_2R (mean length $1/(\beta_2 + k_{-2a} + k_{-2b}) \approx 15 \mu\text{s}$), the total shut time per burst is:

$$\mu_{\text{tsb}} = \frac{\mu_{\text{ob}} - 1}{\beta_2 + k_{-2a} + k_{-2b}}, \quad (16)$$

and the approximate mean activation length, $\mu_{\text{tob}} + \mu_{\text{tsb}} \approx \mu_{\text{ob}}$ should be a good estimate of the time constant for decay of synaptic currents. For a more complete account of the relationship between single channel events and the time course of macroscopic currents, see Colquhoun *et al.* (1997) and Wyllie *et al.* (1998).

Monoliganded channels

The second major conclusion is that all the rate constants that are specific for one site or the other are harder to estimate, but good estimates can be obtained as long as the two sites are independent. This can be done without knowing the number of channels in the patch either by (a) using a single low concentration of agonist in conjunction with an EC_{50} value, or (b) using simultaneous fitting of a

high and a low concentration record. Of course these conclusions are dependent, as in any simulation study, on the values for the true rate constants (those that are used to simulate data) being close to the actual values that hold for real receptors. The 'true rates' used here are similar to our current best estimates (Hatton *et al.* 2003), but those estimates are, for the monoliganded rates, not yet very precise. In fact the estimates of the monoliganded rates from real data showed rather more variability than one might have expected from the simulation results. The reason for this is not yet known. It might be that there are genuine differences in these rates from one experiment to another, or it might be that the 'true rates' used here are not sufficiently close to the real values.

Independence of the two binding sites

If the two binding sites are not independent, then simulations suggest that it is no longer possible to estimate all 13 free rate constants, in the absence of knowledge of the number of channels in the patch. The only circumstance under which we found it possible to estimate all 13 rate constants was by simultaneous fit of either two low concentration records (Fig. 18), or a high and a low concentration record, for which it was known that the low concentration record(s) contained only one channel in the patch. There is usually no reliable way to know that a low concentration record contains only one channel, so it will not usually be possible to do this in practice.

Furthermore, when simulations were done of a channel with non-independent sites, it was found (Fig. 14) that good predictions of the data could be obtained when the fitting process assumed, incorrectly, that the sites were independent. Although such fits still gave good estimates of the 'diliganded' rates, at least two of the 'monoliganded' rates were in serious error, though this was not apparent from the analysis. This means that detection of non-independence is likely to be quite hard in practice, a fact that might account for the many conflicting results in the literature. However in the example given, the conditional mean open time plot (Fig. 14D) does show that there is a problem with the fit, as do conditional open time distributions over certain ranges of adjacent shut times (Fig. 14E), though not over others (Fig. 14C). This illustrates the value of looking at things other than the predicted fit to apparent open and shut times only.

The performance of the HJCFIT method

The only other implementation of the maximum likelihood fitting of an entire sequence is the MIL program (Qin *et al.* 1996; see Methods). We have made no direct comparison between the results of fitting with HJCFIT and with MIL, but some general observations can be made.

(1) The distributions of estimates from HJCFIT have been investigated (this paper), whereas those from MIL have not.

(2) The HJCFIT program incorporates the ability to do repeated simulations of the sort described here for essentially any mechanism and values of rate constants. A utility program, SIMAN, is available to inspect the outcome of such simulations.

(3) The missed event correction used by HJCFIT is exact, whereas that used by MIL is approximate. Strictly speaking the asymptotic form used in HJCFIT for intervals longer than $3t_{\text{res}}$ is an approximation but HJCFIT allows a visual check that it is essentially identical with the exact form in the region that is used (e.g. Fig. 6B). The asymptotic form has the advantage that it has the form of (the right number of) exponential components, the times constants and areas of which are printed. The practical importance of this difference in method for making the missed event correction has not been investigated.

(4) In HJCFIT the start and end vectors for each group of openings can be calculated from the exact HJC theory. This does not matter much if the groups of openings are long (as is usually the case in high concentration records) but it is shown here to be of great importance when individual activations are fitted as a group at low agonist concentrations. This feature allows records to be used when they contain an unknown number of channels.

(5) In HJCFIT it is possible to use a specified value of an EC_{50} , determined independently, to calculate the value of any specified rate constant from the values of all the others.

This reduces by one the number of free parameters that have to be estimated.

(6) In HJCFIT, the whole covariance matrix is calculated at the end of a fit, so as well as approximate standard deviations for the estimates, the correlation coefficients between all possible pairs of estimates are also printed. These can be very useful for detection of poorly defined parameters.

(7) In HJCFIT there are more ways to test the quality of the fit, once it has been done. Whereas MIL displays only open and shut time distributions at the end of a fit, HJCFIT can display also the following tests.

(a) The conditional distribution of apparent open times, for open times that are before, after or adjacent to shut times in a specified range. The ability to plot separately 'before' and 'after' is potentially useful for mechanisms that behave irreversibly (e.g. Wyllie *et al.* 1998), though it is not used in this paper.

(b) The relationship between the mean of the conditional apparent open time distribution and the adjacent shut time (range).

(c) The dependency plot.

In all of these cases the prediction of the observed distributions is calculated from the fitted rate constants (and the resolution) by exact HJC methods.

Correlations

Two entirely different sorts of correlation are relevant to this work.

(1) The correlation between open and shut times. This is inherent in the mechanism, it tells us about how states are connected (Fredkin *et al.* 1985; Colquhoun & Hawkes, 1987; Blatz & Magleby, 1989), and the maximum likelihood takes into account fully the information from such correlations in the experimental record. The ability of the mechanism, and the fitted rate constants to describe correctly this sort of correlation is what is tested by the conditional distributions, the conditional mean plot, and the dependency plot (Methods and Figs 6 and 14).

(2) Statistical correlation between estimates of parameters. This sort of correlation is not a property of the mechanism, but is a property of the estimation method. It imparts no interesting information, and it disappears if the data are sufficiently precise. It is this sort of correlation that can be estimated from the covariance matrix after fitting a single experiment, and shown as a graph in the case of repeated simulations (see Figs 2D, 3E, 4E, 5E, 5F, 8G and 11, and Table 3). There is a simple intuitive explanation for the strong positive correlation between the estimates of α_2 and β_2 (Figs 2D, 3E and 12A and B). The basis of the explanation is that the length of the activations (bursts) is




Figure 19. Two solutions for missed events problem in the simplest case

The two curves are plots of eqns (121) and (122) in Colquhoun & Hawkes (1995b, p. 456), for the example cited there in which the apparent mean open and shut times are 0.2 ms and 2.0 ms respectively, and the resolution was $t_{\text{res}} = 0.2$ ms. The intersections show that these simultaneous equations are satisfied either by true open and shut times of 0.299 and 0.879 ms respectively (the 'slow solution'), and equally by true open and shut times of 0.106 and 0.215 ms respectively (the 'fast solution').

quite well defined by the data. But the brevity of the shuttings that separate the individual openings in the burst means that many of them are not detected (with a resolution of 25 μ s, and a mean short shut time of 15 μ s, more than 80 % will be missed). This means that the length of individual openings, and the number of them in a burst, are rather poorly defined, though defining them is the essential core of the binding–gating problem. A large value of α_2 means that individual openings are short, so the observed mean burst length can be fitted only by postulating a large number of openings per burst, something that can be achieved (see eqn (14)) by having a large value for β_2 also. This argument suggests that experiments that happen to produce a large value of α_2 will also produce a large value of β_2 , as was observed in every case.

Initial guesses and the existence of two solutions

It is known that the simplest possible case of the missed event problem (one open state and one shut state) has two exact solutions in general (Colquhoun & Sigworth, 1983; see Colquhoun & Hawkes, 1995*b*). This is illustrated in Fig. 19.

The graph shows plots of eqns (121) and (122) in Colquhoun & Hawkes (1995*b*, p. 456), for the example cited there in which the apparent mean open and shut times are 0.2 ms and 2.0 ms respectively, and the resolution was $t_{\text{res}} = 0.2$ ms. The intersections show that these simultaneous equations are satisfied either by true open and shut times of 0.299 and 0.879 ms respectively (the ‘slow solution’), and equally by true open and shut times of 0.106 and 0.215 ms respectively (the ‘fast solution’). It has been shown (Ball *et al.* 1990) that these two solutions have near identical likelihoods. Something very similar appears to happen in the far more complex schemes analysed here. In Fig. 2 it was shown that, with some initial guesses, the fit could converge in some cases to give a fit with the correct values of α_2 and β_2 , and sometimes to a fit in which both values were much larger. Both fits had very similar likelihoods (Fig. 2*C*). Figure 2*D* shows that the strong positive correlation between the estimates of α_2 and β_2 that was seen in all experiments, extends over the entire range of α_2 and β_2 values in such a case (though it is rather non-linear in this case), but that the α_2, β_2 values almost all fall into one or the other of two clusters, one of which (the slow solution) corresponds to the correct values of α_2 and β_2 , and the other of which is analogous to the other (fast) solution. One may conjecture that missed event problems for any mechanism will always have, in this sense, two solutions.

In summary, the HJCFIT method, as implemented here, can provide good estimates of the main ‘diliganded’ rate constants that are needed to solve the physiological binding–gating problem, without assuming anything

about the number of channels in the patch, and with data of the usual resolution of single channel recordings. These estimates are surprisingly immune to various sorts of errors in the assumptions. The rate constants for each of the two binding sites separately can also be obtained, though only if it is assumed that the two sites are independent. These methods are applied to experimental data in the accompanying paper (Hatton *et al.* 2003).

REFERENCES

- Akk G & Steinbach JH (2000). Structural elements near the C-terminus are responsible for changes in nicotinic receptor gating kinetics following patch excision. *J Physiol* **527**, 405–417.
- Ball FG, Davies SS & Sansom MSP (1988*a*). Aggregated Markov processes incorporating time interval omission. *Adv Appl Prob* **20**, 546–572.
- Ball FG, Davies SS & Sansom MS (1990). Single-channel data and missed events: analysis of a two-state Markov model. *Proc R Soc Lond B Biol Sci* **242**, 61–67.
- Ball FG & Sansom MSP (1988*b*). Single channel autocorrelation functions. The effects of time interval omission. *Biophys J* **53**, 819–832.
- Ball FG & Sansom MSP (1989). Ion channel gating mechanisms: model identification and parameter estimation from single channel recordings. *Proc R Soc Lond B Biol Sci* **236**, 385–416.
- Bauer RJ, Bowman BF, & Kenyon JL (1987). Theory of the kinetic analysis of patch-clamp data. *Biophys J* **52**, 961–978.
- Blatz AL & Magleby KL (1989). Adjacent interval analysis distinguishes among gating mechanisms for the fast chloride channel from rat skeletal muscle. *J Physiol* **410**, 561–585.
- Bouzat C, Barrantes F & Sine S (2000). Nicotinic receptor fourth transmembrane domain. Hydrogen bonding by conserved threonine contributes to channel gating kinetics. *J Gen Physiol* **115**, 663–672.
- Colquhoun D (1998). Binding, gating, affinity and efficacy. The interpretation of structure–activity relationships for agonists and of the effects of mutating receptors. *Br J Pharmacol* **125**, 923–948.
- Colquhoun D & Hawkes AG (1982). On the stochastic properties of bursts of single ion channel openings and of clusters of bursts. *Philos Trans R Soc Lond B Biol Sci* **300**, 1–59.
- Colquhoun D & Hawkes AG (1987). A note on correlations in single ion channel records. *Proc R Soc Lond B Biol Sci* **230**, 15–52.
- Colquhoun D & Hawkes AG (1995*a*). A Q-matrix cookbook. In *Single Channel Recording*, ed Sakmann B & Neher E, pp. 589–633. Plenum Press, New York.
- Colquhoun D & Hawkes AG (1995*b*). The principles of the stochastic interpretation of ion channel mechanisms. In *Single Channel Recording*, ed Sakmann B & Neher E, pp. 397–482. Plenum Press, New York.
- Colquhoun D, Hawkes AG, Merlushkin A & Edmonds B (1997). Properties of single ion channel currents elicited by a pulse of agonist concentration or voltage. *Philos Trans R Soc Lond A* **355**, 1743–1786.
- Colquhoun D, Hawkes AG & Srodzinski K (1996). Joint distributions of apparent open times and shut times of single ion channels and the maximum likelihood fitting of mechanisms. *Philos Trans R Soc Lond A* **354**, 2555–2590.
- Colquhoun D & Ogden DC (1988). Activation of ion channels in the frog end-plate by high concentrations of acetylcholine. *J Physiol* **395**, 131–159.

- Colquhoun D & Sakmann B (1981). Fluctuations in the microsecond time range of the current through single acetylcholine receptor ion channels. *Nature* **294**, 464–466.
- Colquhoun D & Sakmann B (1985). Fast events in single-channel currents activated by acetylcholine and its analogues at the frog muscle end-plate. *J Physiol* **369**, 501–557.
- Colquhoun D & Sigworth FJ (1983). Analysis of single ion channel data. In *Single Channel Recording*, ed Sakmann B & Neher E, pp. 191–263. Plenum Press, New York.
- Colquhoun D & Sigworth FJ (1995). Fitting and statistical analysis of single-channel records. In *Single Channel Recording*, ed Sakmann B & Neher E, pp. 483–587. Plenum Press, New York.
- Crouzy SC & Sigworth FJ (1990). Yet another approach to the dwell-time omission problem of single-channel analysis. *Biophys J* **58**, 731–743.
- Elenes S & Auerbach A (2002). Desensitization of diliganded mouse muscle nicotinic acetylcholine receptor channels. *J Physiol* **541**, 367–383.
- Fredkin DR, Montal M & Rice JA (1985). Identification of aggregated Markovian models: application to the nicotinic acetylcholine receptor. In *Proceedings of the Berkeley Conference in Honor of Jerzy Neyman and Jack Kiefer*, ed Le Cam LM, pp. 269–289. Wadsworth, Monterey.
- Fredkin DR & Rice JA (1986). On aggregated Markov processes. *J Appl Prob* **23**, 208–214.
- Gibb AJ, Kojima H, Carr JA & Colquhoun D (1990). Expression of cloned receptor subunits produces multiple receptors. *Proc R Soc Lond B Biol Sci* **242**, 108–112.
- Grosman C & Auerbach A (2000). Kinetic, mechanistic, and structural aspects of unliganded gating of acetylcholine receptor channels. A single-channel study of second transmembrane segment 12' mutants. *J Gen Physiol* **115**, 621–635.
- Hatton C, Shelley C, Brydson M, Beeson D & Colquhoun D (2003). Properties of the human muscle nicotinic receptor, and of the slow-channel syndrome mutant ϵ L221F, inferred from maximum likelihood fits. *J Physiol* **547**, 000–000.
- Hawkes AG, Jalali A & Colquhoun D (1990). The distributions of the apparent open times and shut times in a single channel record when brief events can not be detected. *Philos Trans R Soc Lond A* **332**, 511–538.
- Hawkes AG, Jalali A & Colquhoun D (1992). Asymptotic distributions of apparent open times and shut times in a single channel record allowing for the omission of brief events. *Philos Trans R Soc Lond B Biol Sci* **337**, 383–404.
- Horn R & Lange K (1983). Estimating kinetic constants from single channel data. *Biophys J* **43**, 207–223.
- Jackson MB (1984). Spontaneous openings of the acetylcholine receptor channel. *Proc Natl Acad Sci U S A* **81**, 3901–3904.
- Jackson MB (1988). Dependence of acetylcholine receptor channel kinetics on agonist concentration in cultured mouse muscle fibres. *J Physiol* **397**, 555–583.
- Jackson MB (1989). Perfection of a synaptic receptor: kinetics and energetics of the acetylcholine receptor. *Proc Natl Acad Sci U S A* **86**, 1–6.
- Katz B & Thesleff S (1957). A study of the 'desensitization' produced by acetylcholine at the motor end-plate. *J Physiol* **138**, 63–80.
- Magleby KL & Song L (1992). Dependency plots suggest the kinetic structure of ion channels. *Proc R Soc Lond B Biol Sci* **249**, 133–142.
- Milone M, Wang HL, Ohno K, Fukudome T, Pruitt JN, Bren N, Sine SM & Engel AG (1997). Slow-channel myasthenic syndrome caused by enhanced activation, desensitization, and agonist binding affinity attributable to mutation in the M2 domain of the acetylcholine receptor α subunit. *J Neurosci* **17**, 5651–5665.
- Parzefall F, Wilhelm R, Heckmann M & Dudel J (1998). Single channel currents at six microsecond resolution elicited by acetylcholine in mouse myoballs. *J Physiol* **512**, 181–188.
- Qin F, Auerbach A & Sachs F (1996). Estimating single-channel kinetic parameters from idealized patch-clamp data containing missed events. *Biophys J* **70**, 264–280.
- Qin F, Auerbach A & Sachs F (1997). Maximum likelihood estimation of aggregated Markov processes. *Proc R Soc Lond B Biol Sci* **264**, 375–383.
- Roux B & Sauve R (1985). A general solution to the time interval omission problem applied to single channel analysis. *Biophys J* **48**, 149–158.
- Sakmann B, Patlak J & Neher E (1980). Single acetylcholine-activated channels show burst-kinetics in presence of desensitizing concentrations of agonist. *Nature* **286**, 71–73.
- Salamone FN, Zhou M & Auerbach A (1999). A re-examination of adult mouse nicotinic acetylcholine receptor channel activation kinetics. *J Physiol* **516**, 315–330.
- Sigworth FJ & Sine SM (1987). Data transformations for improved display and fitting of single-channel dwell time histograms. *Biophys J* **52**, 1047–1054.
- Sine SM & Steinbach JH (1987). Activation of acetylcholine receptors on clonal mammalian BC3H-1 cells by high concentrations of agonist. *J Physiol* **385**, 325–359.
- Wichmann BA & Hill ID (1985). An efficient and portable pseudo-random number generator. In *Applied Statistics Algorithms*, ed Griffiths P, pp. 238–242. Ellis Horwood, Chichester.
- Wyllie DJA, Béhé P & Colquhoun D (1998). Single-channel activations and concentration jumps: comparison of recombinant NR1/NR2A and NR1/NR2D NMDA receptors. *J Physiol* **510**, 1–18.

Acknowledgements

We are grateful to the Medical Research Council and the Wellcome Trust for supporting this work, and to Chris Shelley, Marco Beato and Lucia Sivilotti, for helpful comments.