

Parameter inference for a stochastic kinetic model of expanded polyglutamine proteins

H. F. Fisher^{1,2}  | R. J. Boys¹ | C. S. Gillespie¹  | C. J. Proctor³ | A. Golightly¹ 

¹ School of Mathematics, Statistics & Physics, Newcastle University, Newcastle Upon Tyne, UK

² Population Health Sciences Institute, Newcastle University, Newcastle Upon Tyne, UK

³ Institute of Cellular Medicine, Newcastle University, Newcastle Upon Tyne, UK

Correspondence C. S. Gillespie, School of Mathematics, Statistics, and Physics, Newcastle University, Newcastle Upon Tyne, NE1 7RU, UK.

Email: colin.gillespie@newcastle.ac.uk

Funding information

Engineering and Physical Sciences Research Council

Abstract

The presence of protein aggregates in cells is a known feature of many human age-related diseases, such as Huntington's disease. Simulations using fixed parameter values in a model of the dynamic evolution of expanded polyglutamine (PolyQ) proteins in cells have been used to gain a better understanding of the biological system. However, there is considerable uncertainty about the values of some of the parameters governing the system. Currently, appropriate values are chosen by ad hoc attempts to tune the parameters so that the model output matches experimental data. The problem is further complicated by the fact that the data only offer a partial insight into the underlying biological process: the data consist only of the proportions of cell death and of cells with inclusion bodies at a few time points, corrupted by measurement error. Developing inference procedures to estimate the model parameters in this scenario is a significant task. The model probabilities corresponding to the observed proportions cannot be evaluated exactly, and so they are estimated within the inference algorithm by repeatedly simulating realizations from the model. In general such an approach is computationally very expensive, and we therefore construct Gaussian process emulators for the key quantities and reformulate our algorithm around these fast stochastic approximations. We conclude by highlighting appropriate values of the model parameters leading to new insights into the underlying biological processes.

KEYWORDS

Gaussian process emulator, history matching, Markov chain Monte Carlo (MCMC), optimal design, stochastic kinetic model

1 | INTRODUCTION

One of the main aims of modeling biological systems is to describe and understand the temporal evolution of the system taking account of the potentially complex inter-

relationships between components within the system. Models can also be used to facilitate *in silico* experiments in which virtual experiments are performed on a computer. These *in silico* experiments have an advantage over laboratory-based experiments as, in general, they are

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. *Biometrics* published by Wiley Periodicals, Inc. on behalf of International Biometric Society.

much cheaper and faster to conduct. This can lead to a better understanding of, for example, the biological system, how to focus drug development and how to construct more efficient designs of future laboratory-based in vitro experiments.

The accumulation of abnormal protein deposits within cells are hallmarks of neurodegenerative diseases affecting humans as they age. There are many such diseases, (e.g., Alzheimer's and Parkinson's disease), in which different parts of the brain are affected resulting in a range of clinical symptoms such as loss of motor function, dementia, and behavioral changes. Despite differences between symptoms, there are similarities in the underlying molecular mechanisms leading to the accumulation of protein aggregates and the neuronal cell death (Gan *et al.*, 2018). Although age is the greatest risk factor for neurodegeneration, there is a group of diseases that are caused by genetic mutations. In nine of these diseases, the mutation occurs in a gene that contains a repeat of a CAG nucleotide (base pair) triplet (Lieberman *et al.*, 2019). As CAG encodes for the amino acid glutamine, these diseases are known as polyglutamine diseases and the proteins they encode are referred to as PolyQ proteins, although different genes and proteins are involved in each disease, for example, Huntington's disease (HD), (Lieberman *et al.*, 2019). HD is an adult-onset, progressive disease characterized by loss of motor function, psychiatric disorders, and dementia (McColgan and Tabrizi, 2018). In this disease, the CAG triplet is found in the Huntingtin gene (HTT; HGNC:4581). In the normal HTT gene, the CAG triplet is repeated 10–35 times but in the mutated form the segment is expanded from 36 to over 120 repeats. This leads to the formation of an abnormally long protein which may be cleaved within the cell into small toxic fragments that enter the nucleus where they form aggregates, known as inclusion bodies (or inclusions) and sequester other proteins causing transcriptional dysregulation although the exact mechanisms of how this occurs are still unknown (McColgan and Tabrizi, 2018). In addition, Huntingtin fragments form aggregates in the cytoplasm and may impair cellular function, for example, impairment of proteostasis (McColgan and Tabrizi, 2018).

There are currently no effective interventions for the prevention, delay in onset or slowing down of disease progression in HD or in any other neurodegenerative disorder. This is mainly due to a lack of a full understanding of the underlying molecular mechanisms. In particular, although protein aggregation is a common feature of all these disorders, it is still not fully known how protein aggregates actually contribute to the disease process. There has been considerable controversy over which stage of the aggregation process is most toxic to cells, and it has been suggested that the formation of large inclusion bodies may be

protective as they sequester misfolded proteins and prevent the overload of protein degradation pathways (Ross and Poirier, 2004). This has been shown experimentally in cell culture (Arrasate *et al.*, 2004). However, this protective effect may only be short-term, as large aggregates also contain proteins involved in the removal and repair of damaged protein, and it has been suggested that they may induce quiescence and induction of cell death via necrosis (Ramdzan *et al.*, 2017), which can cause damage to neighboring cells. In addition, nuclear inclusions may impair redox signaling leading to an increase in oxidative stress and further damage (Paul and Snyder, 2019). The controversy regarding cytotoxicity of PolyQ proteins is largely due to insufficient understanding of the molecular mechanisms involved. This motivated our previous study which used live cell imaging with fluorescent reporter systems to examine the relationship between PolyQ protein, activation of the stress kinase p38MAPK (MAPK14; HGNC:6876), reactive oxygen species (ROS) generation, inhibition of the proteasome (a protein complex which degrades cellular proteins), and formation of PolyQ nuclear inclusions (Tang *et al.*, 2010). It was found that proteasome inhibition usually preceded formation of inclusion bodies and that p38MAPK inhibitors alleviated the inhibition of proteasomes and delayed the onset of inclusion formation. Conversely, the addition of proteasome inhibitors resulted in earlier formation of inclusions. This study also included a computational model which explored the complex interactions of PolyQ proteins with other elements of the cell, and they used computer simulations from the model (with fixed parameter values) to suggest ways to reduce the toxicity of PolyQ proteins on cells. As their model describes dynamic interactions at the single cell level, the number of different biochemical species vary discretely over time, often with low copy numbers (Gillespie, 1977). Also the interaction of the species is driven by Brownian motion and so accurate modeling requires that it take account of the inherent stochasticity present in the system.

Currently, plausible values for the parameters in the PolyQ model are determined by using model simulations for specified values of the parameters and then adjusting them in an attempt to match experimental data. Clearly this is a difficult task, and one with which statisticians can make a contribution. A key difficulty in conducting parameter inference for complex stochastic models such as the PolyQ model is that the experimental data are typically incomplete and subject to measurement error. This necessitates the use of computationally intensive schemes such as Gibbs sampling (Boys *et al.*, 2008), pseudo-marginal Metropolis–Hastings (Golightly and Wilkinson, 2011; Georgoulas *et al.*, 2016; Wilkinson, 2018), population Monte Carlo (Koblents and Míguez, 2015; Koblents *et al.*, 2019), and approximate Bayesian computation (Toni *et al.*,

2009; Liepe *et al.*, 2014; Wu *et al.*, 2014). Inference is particularly challenging for the PolyQ model since the data consist only of the proportion of cells which are dead and the proportion of cells which contain inclusion bodies. Since the corresponding model probabilities are intractable, these latent proportions must be estimated at each observation time by repeatedly forward simulating the model over the duration of the time course. Computational cost can be reduced by replacing the expensive and exact stochastic simulator with a cheap approximation. For example, approaches which replace the continuous-time stochastic model with a discrete-time approximation include tau-leaping (Gillespie, 2001; Cao *et al.*, 2006) and the chemical Langevin equation (Gillespie, 2000; van Kampen, 2001).

1.1 | Contributions and organization of the paper

The aim of this paper is to examine what insight the limited available experimental data provides about the PolyQ model parameters and to check whether the model gives a reasonable description of the dynamic variation observed in the experimental data.

Given the prohibitive computational cost associated with fitting the stochastic polyQ model, we consider a computationally feasible approximation and perform exact (simulation-based) inference for the resulting model. We eschew the use of an approximate simulator in favor of *directly emulating* the empirical logit of the proportions of interest with a Gaussian process (GP) (Santner *et al.*, 2003; Rasmussen and Williams, 2006). In particular, combining the fitted GP emulator with a Gaussian measurement model (on the same scale as the observed data) allows a direct approximation of the observed data likelihood without recourse to further simulation from the stochastic PolyQ model.

The findings from this analysis give new information regarding the parameters of the model, which in turn leads to new insights into the underlying biological processes. For example, this analysis has shown that during the first stages of the aggregation process both aggregation and disaggregation probably occur much faster and that the disaggregation/aggregation ratio is likely to be a magnitude higher than was originally assumed. The biological implication of this is that it will take longer to reach the threshold size required for inclusion formation but that there will also be higher levels of small aggregates which will inhibit the proteasome.

The remainder of the paper is organized as follows. Section 2 details the experimental data available and outlines how it was collected. A complete description of the model is given in Section 3 which includes the underlying

stochastic model which describes the dynamic evolution of this biological process and the observation model, which links the data to the underlying process. Our prior assumptions about parameter values and initial levels are given in Section 4. A method for determining the posterior distribution for model quantities is described in Section 5 and, because a standard simulation-based Markov chain Monte Carlo (MCMC) solution is prohibitively expensive, we develop GP emulators in Section 6 which facilitate timely generation of realizations from the posterior distribution. As we have high prior uncertainty on model parameters, in Section 7 we employ a history matching technique which removes implausible training points in an attempt to fit the GPs within regions of nonnegligible posterior support. We also include an assessment of the quality of the final fitted GPs. We present our findings about the PolyQ model in Section 8 and point out new insights into this biological system.

2 | EXPERIMENTAL DATA

We have data from two different sets of experiments. These were carried out in the same laboratory and are described briefly below. A more comprehensive description of the experimental procedures can be found in Tang *et al.* (2010).

2.1 | Cell death

The first experiment begins with a large number of human U87MG glioblastoma cells maintained in a suitable medium at 37°C at 5% CO₂. The cells are transfected with a construct that contains the expanded PolyQ Huntington protein (see Section 2.2 for details). The survival of cells is monitored over time. Three different repeats of the experiment are carried out under the following experimental conditions:

- (i) *Control group*: No intervention.
- (ii) *Proteasome inhibition group*: Cells are treated with a proteasome inhibitor 24 h after transfection. The stochastic kinetic model captures this scenario by reducing the initial value $k_{\text{proteff}} = 1$ to $k_{\text{proteff}} = 0.05$ after 24 h.
- (iii) *p38 inhibition group*: Cells were pretreated for 2 h with a p38 inhibitor. The stochastic kinetic model captures this scenario by setting $k_{\text{p38act}} = 0.05$.

The changes to parameter values used to describe the different experimental conditions were determined by conversations with the experimentalists. It was felt that reducing the rates to zero in conditions (ii) and (iii) overstated

TABLE 1 Proportions of cell death and inclusion bodies observed under different experimental conditions

Condition	24 h	30 h	36 h	42 h	48 h
<i>Cell death</i>					
(i)	0.1503	–	0.1455	–	0.2608
	0.1788	–	0.1821	–	0.2846
(ii)	0.1897	–	0.1807	–	0.2250
	0.1640	–	0.1973	–	0.2998
(iii)	0.2168	–	0.2344	–	0.3644
	0.2436	–	0.2095	–	0.3872
<i>Inclusion bodies</i>					
(i)	0.0909	0.2857	0.3824	0.4286	0.7273
	0.0175	0.2131	0.3538	0.4304	0.5742
(ii)	0.0909	0.4247	0.6125	0.7403	0.8194
	0.1154	0.5075	0.7895	0.8667	0.9157
(iii)	0.0303	0.0444	0.0612	0.1373	0.2200
	0.0189	0.0938	0.1286	0.1613	0.1833

the effect of the experimental treatment, and so these rates are set to low but nonzero values.

The cells are stained with propidium iodide, which is a fluorescent dye with the property that it only binds to nonviable (dead) cells. This technique is called *propidium iodide exclusion* and is used to identify the viability of the cells over time. The fluorescent dye can be viewed under a microscope, and estimates of the proportion of cell death can be observed over time. These proportions of cell death form our experimental data which can be seen in Table 1 (top panel). Each row of the table corresponds to experiments carried out under the different experimental conditions outlined above. We have two repeats of each experiment.

Within the stochastic model, cell death can occur via two biological pathways, either via proteasome inhibition or by p38 activation. This process is monitored within the model by using the dummy species P_{Ideath} and $p38_{\text{death}}$. These species are both binary variables, with a zero value representing that the cell is alive.

2.2 | Inclusion body formation

Experimentalists monitor the number of cells with inclusion bodies by first creating a construct that encodes an expanded PolyQ Huntington protein which is tagged with yellow fluorescent protein (YFP). U87MG cells are transfected with the construct and then imaged at 10 min intervals between 24 and 48 h after transfection. Time lapse images reveal the formation of inclusions (by detecting YFP), and the percentage of cells with inclusions is

recorded every 6 h. The experimental results are given in Table 1 (bottom panel).

Within the stochastic model, the number of cells containing inclusions is counted via the number of cells containing SeqAggP. In the laboratory, image analysis suffers from a thresholding problem in detecting inclusions and we capture this aspect by defining cells with inclusions as those in which SeqAggP contains more than 10 PolyQ proteins (i.e., $\text{SeqAggP} > 10$). The value of this threshold is based on the biological modeler's best understanding of the mechanism, though in fact the actual value is not crucial since aggregates grow very rapidly, and there is only a very small timeframe when SeqAggP is less than 10.

3 | THE STOCHASTIC MODEL

3.1 | PolyQ mechanism

The stochastic model for the PolyQ mechanism we consider in this paper is a reduced form of that proposed by Tang *et al.* (2010). The original model developed by Tang *et al.* (2010) was constructed in the Systems Biology Markup Language (SBML) and is available from the Biodels database (Li *et al.* 2010) (model ID: BIOMD000000285). The model was developed to investigate the effects of PolyQ on proteasome function, oxidative stress, and cell death. It contained a pool of PolyQ which was continually turned over, being degraded by the proteasome. It was assumed that the PolyQ aggregation process started by two monomers interacting, which then interacts with further monomers causing the aggregate to grow (AggPolyQ1–5). At early stages, it was assumed that disaggregation could also take place but that when the aggregate reached a threshold size (denoted by SeqAggP in the model), disaggregation does not take place, so that the aggregate continues to grow and an inclusion forms. The model also included a generic pool of protein which could be damaged by ROS leading to misfolding and aggregation. It was assumed that small aggregates bind to the proteasome but cannot be degraded and so inhibit proteasome function. In addition, they may increase levels of ROS. The model also included the stress kinase p38MAPK (simply modeled as two pools: inactive [p38] and active [p38P]) with activation occurring more frequently when levels of ROS are high (Sato *et al.*, 2014), and that high levels of p38P activated a cell death pathway (denoted by p38death). Proteasomes bound by aggregates (AggPProteasome) could also activate a cell death pathway (denoted by PIdeath). The model also included the turnover of a fluorescent protein (mRFPu) as an increase in mRFPu levels indicates that the proteasome is inhibited, shown by live cell imaging in

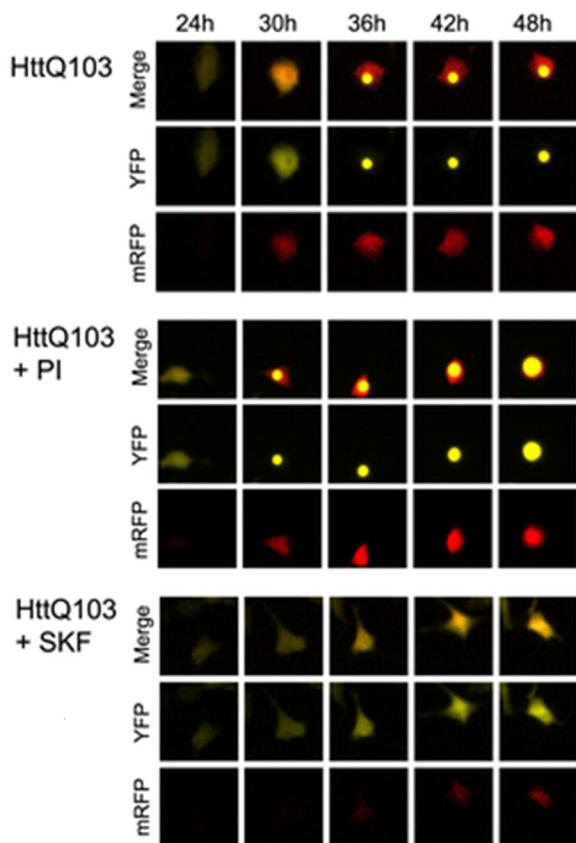


FIGURE 1 Example data from live cell imaging. This figure appears in color in the electronic version of this article, and any mention of color refers to that version

Tang *et al.* (2010). An example of this data is shown in Figure 1.

Given the complexity of the original model (with 27 chemical species and 72 reactions), it was necessary to consider a simplification for this study. The changes to the Tang model have been to reduce the number of species and reactions by removing the generic pool of protein and its associated reactions of misfolding and aggregation, and by removing the fluorescent protein mRFPu, which had been originally included as a marker of proteasome function. In the process of model construction, it was found necessary to add in an effect of ROS on the aggregation process. The experimental biologists suggested that ROS increases the propensity of PolyQ proteins to aggregate due to ROS interfering with the ubiquitin-proteasome system (UPS), so that PolyQ proteins are more likely to aggregate than be degraded when ROS levels are high. The exact mechanism is not fully understood, and so we simply included ROS in the rate laws for PolyQ aggregation.

The resulting reduced model contains 14 chemical species and 33 reaction channels, which is still relatively large when compared to other stochastic kinetic models for which fully Bayesian inference is available in the literature (Boys *et al.*, 2008; Henderson *et al.*, 2010). A complete list of

the reactions and their stochastic rate laws describing the PolyQ model can be found in Web Appendix A of the Supporting Information. All reactions except those relating to aggregation follow the law of mass action kinetics. We used a Hill function for the effects of ROS as it is thought that low levels of ROS would have little effect on the UPS with a maximal effect when ROS levels are high. We assumed that the rate of aggregation would be half the maximal rate when ROS is at its basal level, which is achieved by using the value 10 in the denominator of the rate law.

The model is represented graphically in Figure 2, where oval shapes (nodes) represent chemical species and an arrow between two nodes represents a reaction which can take place involving the two chemical species. The figure clearly shows the complex interlinkages and complex feedback loops within the model.

3.2 | The observation model

We assume that the experimental data are noisy versions of the equivalent quantities in the stochastic model and adopt a simple model for this observation error which is additive independent normal noise on a logit scale (as the data are proportions). Other error models are possible, such as embedding the model probabilities within a Beta distribution, but this logistic normal model affords advantages in the subsequent analysis.

Suppose the stochastic model for the experiment run under condition c with parameter vector θ gives $p_{t,c}^D(\theta)$ and $p_{t,c}^I(\theta)$ as the probabilities of death and of the presence of inclusions at time t and the logits of their observed versions as $y_{t,c,r}^D$ and $y_{t,c,r}^I$ in experimental repeat r . We assume the observation model is, for $c = 1, 2, 3$, $r = 1, 2$

$$y_{t,c,r}^D = \text{logit } p_{t,c}^D(\theta) + \sigma_D \varepsilon_{i,c,r}, \quad i = 1, 2, 3$$

$$y_{t,c,r}^I = \text{logit } p_{t,c}^I(\theta) + \sigma_I \varepsilon_{i,c,r}^*, \quad i = 1, \dots, 5,$$

where the $\varepsilon_{i,c,r}$ and $\varepsilon_{i,c,r}^*$ are independently and identically distributed standard normal quantities and t_i denote the i th time point at which an observation is available. Note that this model assumes the same measurement error distribution for each experimental condition and repeat. The implications of this observation model on the probability scale is explored in Web Appendix B of the Supporting Information.

4 | PRIOR INFORMATION

Tables 2 and 3 contain a list of all parameters in the model. Some parameters are known quite accurately in the literature. For example, the synthesis and degradation rates

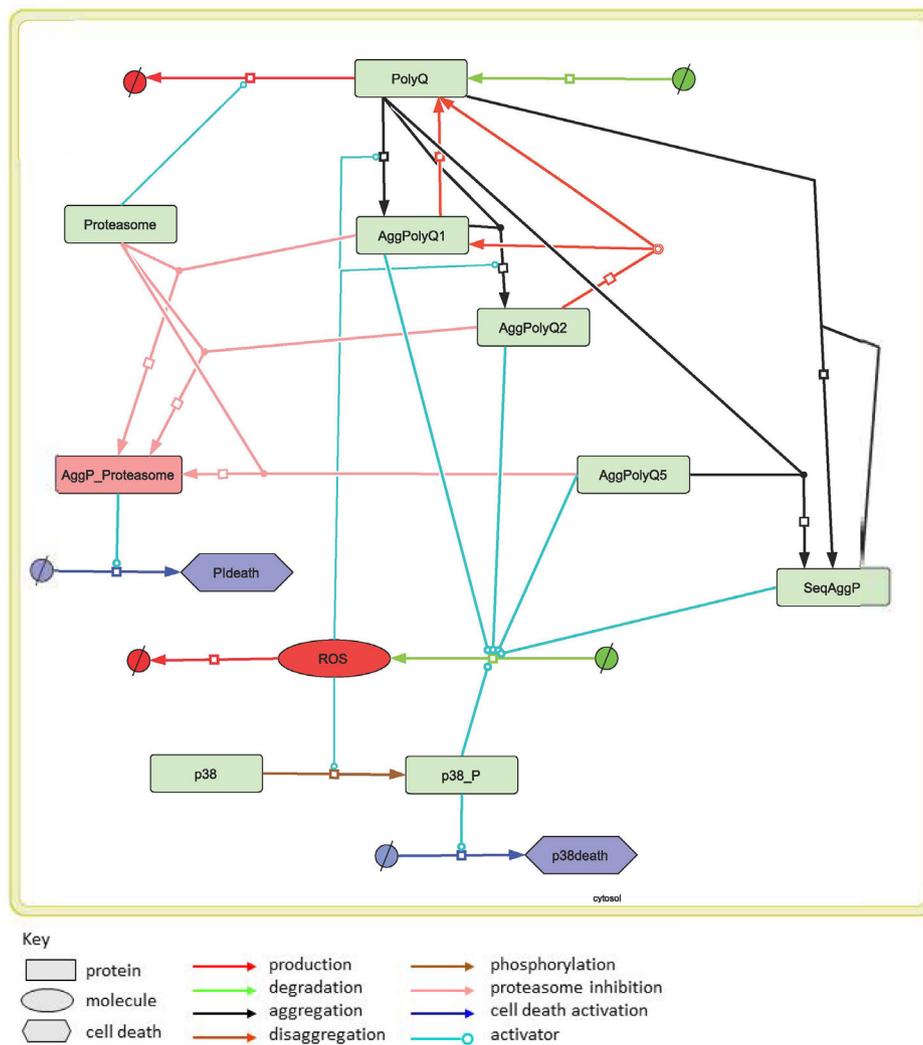


FIGURE 2 Network diagram of the PolyQ model. This figure appears in color in the electronic version of this article, and any mention of color refers to that version

TABLE 2 Parameters in the PolyQ model under condition (i) with their known values

Parameter	Value	Unit	Reference
k_{synPolyQ}	0.25	molecule s^{-1}	Calculated from value of k_{degPolyQ} to give steady-state level of 1000 molecules
k_{degPolyQ}	2.5×10^{-7}	molecule $^{-1}\text{s}^{-1}$	Half-life of 30 h (Persichetti <i>et al.</i> , 1996)
k_{actp38}	2×10^{-3}	molecule $^{-1}\text{s}^{-1}$	Phosphorylation occurs within (Aoki <i>et al.</i> 2013)
k_{p38act}	1.0	Dimensionless	Dummy variable (set to 0.05 for p38 inhibition)
k_{proteff}	1.0	Dimensionless	Dummy variable (set to 0.05 for proteasome inhibition)

of PolyQ proteins k_{synPolyQ} and k_{degPolyQ} can be obtained by considering the half-life of PolyQ proteins (Persichetti *et al.*, 1996).

Also, to reduce the number of parameters in the model, it was thought reasonable to fix some rates to be identical to others or to be functions of other rates. For example, the disaggregation rates for different sized PolyQ aggregates have been fixed as known proportions of the rate for

single PolyQ aggregates: $k_{\text{disaggPolyQ2}} = 0.8 k_{\text{disaggPolyQ1}}$, $k_{\text{disaggPolyQ3}} = 0.6 k_{\text{disaggPolyQ1}}$, and so on. In each case, these identities or fixed proportions have been chosen according to the biological modeler's best understanding of the reaction system.

Only limited information is available for the remaining stochastic rate parameters, and we represent these fairly weak prior beliefs using independent log-normal

TABLE 3 Parameters in the PolyQ model under condition (i) with their prior distribution

Parameter	Prior distributions
k_{aggPolyQ}	$\theta_1 = \log k_{\text{aggPolyQ}} \sim N(\log 10^{-7}, 5)$
$k_{\text{disaggPolyQ}_1}$	$\theta_2 = \log k_{\text{disaggPolyQ}_1} \sim N(\log(5 \times 10^{-7}), 5)$
k_{seqPolyQ}	$\theta_3 = \log k_{\text{seqPolyQ}} \sim N(\log(8 \times 10^{-7}), 5)$
k_{inhprot}	$\theta_4 = \log k_{\text{inhprot}} \sim N(\log(5 \times 10^{-9}), 5)$
k_{remROS}	$\theta_5 = \log k_{\text{remROS}} \sim N(\log(2 \times 10^{-4}), 5)$
$k_{\text{genROSeqAggP}}$	$\theta_6 = \log k_{\text{genROSeqAggP}} \sim N(\log(10^{-7}), 5)$
$k_{\text{genROSAggP}}$	$\theta_7 = \log k_{\text{genROSAggP}} \sim N(\log(5 \times 10^{-6}), 5)$
k_{inactp38}	$\theta_8 = \log k_{\text{inactp38}} \sim N(\log 0.8, 5)$
$k_{\text{genROSp38}}$	$\theta_9 = \log k_{\text{genROSp38}} \sim N(\log(7 \times 10^{-7}), 5)$
k_{p38death}	$\theta_{10} = \log k_{\text{p38death}} \sim N(\log(9 \times 10^{-8}), 5)$
k_{PIdeath}	$\theta_{11} = \log k_{\text{PIdeath}} \sim N(\log(2.5 \times 10^{-8}), 5)$
σ_D	$\sigma_D \sim \text{InvChi}(2, 0.12)$
σ_I	$\sigma_I \sim \text{InvChi}(0.75, 0.05)$

distributions with medians set at our biological modeler's best guess. On the log scale, the kinetic rate parameters are denoted by $\theta_1, \dots, \theta_{11}$ and these parameters have independent normal prior distributions. The weak beliefs are represented by prior variances of five on the log scale as these correspond to a plausible range of values for the (untransformed) kinetic rates from 0.01 to 100 times their median value.

We have prior beliefs for the level of measurement error in the cell death experiments (σ_D) from time course data on two independent replicates from an independent (but similar) study measuring cell death proportions. The difference between these replicates on the logit scale forms a random sample from a $N(0, \sigma_D^2)$ distribution, and so these logit differences lead to an inverse gamma posterior for σ_D^2 under a vague prior. This posterior corresponds to an inverse chi distribution for σ_D , where this distribution is defined as that of $1/\sqrt{X}$, where X has a gamma distribution. Further, as these data arise from a similar but different study, we choose to construct our prior as a powered down version of this inverse chi posterior as this makes it more diffuse around an appropriate value (Ibrahim and Chen, 2000). Our knowledge about the level of measurement error in the inclusion body experiments is weaker, and so we use an even more powered down prior for σ_I .

There is also uncertainty about some of the initial levels of the different species. This uncertainty is captured by independent prior distributions elicited from the biological modeler. Other species levels are known due to the construction of the experiment, such as those for $\text{AggPolyQ}_1 - \text{AggPolyQ}_5$, others obey a conservation law ($\text{p38} + \text{p38P} = 100$), and some levels, such as that for this conservation level and for PolyQ , set at fairly arbitrary high values reflecting the relative abundance of these species within

TABLE 4 Value or prior distribution of the initial levels of the chemical species

Name	Value or prior distribution
PolyQ	1000
Proteasome	$DLN(6.9, 0.1^2)$
AggPolyQ_{1-5}	0
SeqAggP	0
AggPProteasome	0
ROS	$DLN(2.5, 0.25^2)$
p38P	$U\{0, 1, \dots, 5\}$
p38	100-p38P
PIdeath	0
p38death	0

the cell. The values or prior distribution of the initial levels are listed in Table 4. Note that two of these initial levels have a discrete log-normal DLN prior distribution, with probability function $\pi(i) = Pr(i - 0.5 < X < i + 0.5)$, where X is a log-normal random variable.

5 | POSTERIOR INFERENCE

Knowledge of the kinetic rates θ and the observation error levels $\sigma = (\sigma_D, \sigma_I)$ is summarized in the posterior distribution, which has density

$$\pi(\theta, \sigma | \mathbf{y}) \propto \pi(\theta)\pi(\sigma)\pi(\mathbf{y} | \theta, \sigma),$$

where \mathbf{y} represents the collection of all data on cell death and inclusions. Although the normal likelihood has a simple form, it does require calculation of the probabilities $p_{i,c}^D(\theta)$ and $p_{i,c}^I(\theta)$. Unfortunately, analytic expressions for these probabilities are not available due to the complexity of the stochastic model. Prior uncertainty around the initial species levels \mathbf{x}_0 is a further complicating factor as, for example, $p_{i,c}^D(\theta) = E_{\mathbf{x}_0}\{p_{i,c}^D(\theta, \mathbf{x}_0)\}$, where $E_{\mathbf{x}_0}\{\cdot\}$ denotes expectation with respect to \mathbf{x}_0 . Therefore, we estimate these probabilities using n independent realizations of the stochastic model, initialized according to the prior distribution on initial levels. Suppose that a typical probability is estimated by the proportion \hat{p}_n . For sufficiently large n , the empirical logit of such proportions $\text{elogit } \hat{p}_n = \log\{(\hat{p}_n + 0.5/n)/(1 - \hat{p}_n + 0.5/n)\}$ is unbiased for logit p and its sampling variability is closely described by a normal distribution with variance $1/\{n\hat{p}_n(1 - \hat{p}_n)\}$. Thus, taking an improper constant prior for logit p gives its posterior distribution as a normal distribution with mean $\text{elogit } \hat{p}_n$ and variance $1/\{n\hat{p}_n(1 - \hat{p}_n)\}$. Therefore, we can integrate out posterior uncertainty about logit p in the observation model, modifying it to have independent components of

the form

$$y \sim N(\text{elogit } \hat{p}_n(\theta), \sigma^2 + 1/[n \hat{p}_n(\theta)\{1 - \hat{p}_n(\theta)\}]).$$

Building an MCMC inference scheme around this normal likelihood is straightforward. We found that estimating the probabilities using $n = 1000$ model simulations gave empirical logits which fitted well to a normal distribution. Note that, in order for these estimates to be uncorrelated, they have each been calculated using independent realizations of the model.

Unfortunately, the MCMC scheme requires the generation of many millions (or even billions) of realizations from the stochastic model as it explores the posterior distribution. Clearly, this problem prohibits using such a scheme for obtaining realizations from the posterior distribution in a timely manner. In such situations, it is commonplace to use stochastic approximations for these deterministic model probabilities. We will use GP approximations (sometimes called emulators), which are popular in the deterministic computer model literature and elsewhere. Useful background articles in this area are Santner *et al.* (2003) and Bayarri *et al.* (2007). The suitability of GPs as emulators is highlighted in O'Hagan (2006). Also see Henderson *et al.* (2009) for an illustration of the utility of GP emulators in the analysis of complex biological models.

6 | EMULATION

We need to build GP emulators approximating the stochastic output of the proportion of cell death and that of inclusions under the three different conditions. Time (t) is an input variable in these six emulators, in addition to the unknown stochastic rate constants (θ) and observation error levels (σ). However, it can be problematic to specify appropriate covariance kernels over time. In any case, for Bayesian inference, all that is needed are emulators for the distributions of the probabilities under each experimental condition at the distinct time points occurring in the datasets. Therefore, we will build a total of 24 emulators: 9 for proportions of cell death and 15 for proportions of inclusion bodies. Using these time-condition-specific emulators also has the advantage that they can be built in parallel. Thus we build GP emulators for

$$x_{t_i,c}^D(\theta) = \text{elogit } \hat{p}_{t_i,c}^D(\theta) \quad i = 1, 2, 3 \quad \text{and}$$

$$x_{t_i,c}^I(\theta) = \text{elogit } \hat{p}_{t_i,c}^I(\theta) \quad i = 1, \dots, 5,$$

for each condition $c = 1, 2, 3$. The aim is to replace the empirical proportions in the observation model with GP

emulators, allowing for their uncertainty. Fortunately, this is straightforward as each emulator prediction is also normally distributed and so, as before, the additional uncertainty introduced by the GPs can be integrated out analytically and there is no need for the MCMC scheme to integrate over an additional latent layer within the model.

6.1 | Training data

In order to build the emulators for the $x_{t_i,c}^D(\theta)$ and $x_{t_i,c}^I(\theta)$, we first need to evaluate their values at some set of chosen values for θ . This amounts to choosing the size and values in an n_d -point design $\Theta = (\theta_1, \dots, \theta_{n_d})$. In common with other work in this area (Santner *et al.*, 2003; Henderson *et al.*, 2010), we choose to base our design on a Latin hypercube design (LHD) constructed using the *maximin* algorithm of Morris and Mitchell (1995). These designs are popular as they produce an effective and efficient coverage of a bounded space. We take the bounded space to be the central hypercube of the prior distribution defined by the marginal 99% intervals for the rates θ_j . Inevitably, determining an appropriate design from which to construct GP emulators is a sequential process. This is because it is not unusual to find that there are only a few design points in the region of high posterior support, particularly when the marginal priors have large uncertainty (as we have here). To help with this problem, as part of the sequential building of the emulators, we filter out design points which are implausible (inconsistent with the data) by using a history matching technique (see Section 7). Note that obtaining estimates at each point in the LHD for the proportions of cell death and of inclusions over repeated simulations of the stochastic model is easily parallelizable on a high-performance computer system.

6.2 | Mean function and covariance function

The mean function was taken as a linear predictor in the components of θ , with the least-squares estimates as the coefficients. Thus each emulator has a mean function at input $\theta = (\theta_1, \dots, \theta_{11})^T$ of the form

$$m(\theta) = \hat{\beta}_0 + \sum_{i=1}^{11} \hat{\beta}_i \theta_i.$$

This choice was taken to give a parsimonious yet reasonably well-fitting mean function and essentially leaves the residuals to this fit being modeled by a zero mean function GP. As is commonly used in the emulation literature, we

specified a squared exponential covariance function, with

$$K(\theta_i, \theta_j | a, \mathbf{r}) = a \exp \left\{ - \sum_{k=1}^{11} r_k^2 (\theta_{ik} - \theta_{jk})^2 \right\},$$

where, for example, θ_{ik} denotes the k th component of θ_i . The parameters of this function control the overall level of variability and smoothness of the process, with smaller values of the inverse length scales r_k giving smoother realizations.

6.3 | Hyperparameter estimation

The hyperparameters a and $\mathbf{r} = (r_i)$ for each emulator need to be estimated from the training data before we can use them as part of the MCMC inference scheme. When fitting a typical GP to training data $\mathbf{x}(\Theta) = (x(\theta_1), \dots, x(\theta_{n_d}))^T$, the likelihood for the hyperparameters results from $\mathbf{x}(\Theta) | a, \mathbf{r} \sim N_{n_d}(\mathbf{m}(\Theta), K(\Theta, \Theta | a, \mathbf{r}))$, where $N_{n_d}(\cdot, \cdot)$ denotes an n_d -dimensional Gaussian distribution and $K(\Theta, \Theta | a, \mathbf{r})_{ij} = K(\theta_i, \theta_j | a, \mathbf{r})$. In general, it is not possible to obtain a posterior distribution in closed form for a prior on (a, \mathbf{r}) . Here we take independent weak log-normal prior distributions for the hyperparameters, with $a \sim \text{LN}(0, 100)$ and $r_i \sim \text{LN}(0, 100)$, $i = 1, \dots, 11$. The GPs were fitted by using the “no-U-turn” Hamiltonian Monte Carlo sampler and implemented via the RStan package interface to the Stan package (Carpenter *et al.*, 2017; Stan Development Team, 2018).

6.4 | Modified observation model

For a typical GP, the prediction at a point θ^\dagger has distribution

$$x(\theta^\dagger) \sim N(m^*(\theta^\dagger), v^*(\theta^\dagger)), \quad (1)$$

where

$$m^*(\theta^\dagger) = m(\theta^\dagger) + K(\theta^\dagger, \Theta)K(\Theta, \Theta)^{-1}\{\mathbf{x}(\Theta) - \mathbf{m}(\Theta)\}$$

$$v^*(\theta^\dagger) = K(\theta^\dagger, \theta^\dagger) - K(\theta^\dagger, \Theta)K(\Theta, \Theta)^{-1}K(\theta^\dagger, \Theta)^T.$$

Strictly speaking, this emulator distribution is a function of the GP hyperparameters and so we should average over their posterior uncertainty. However, we have found that there is little gain in doing this over using a delta approximation in which the hyperparameters are fixed at their posterior mean.

We can incorporate the stochastic approximation of the GP emulator to elogit \hat{p}_n into our observation model and integrate it out to give a model with independent components of the form

$$y \sim N(m^*(\theta), \sigma^2 + v^*(\theta) + 1/[n p_n^*(\theta)\{1 - p_n^*(\theta)\}]),$$

where

$$p_n^*(\theta) = \text{eexpit}(\theta)$$

$$= \{e^{m^*(\theta)}(1 + 0.5/n) - 0.5/n\} / \{1 + e^{m^*(\theta)}\}$$

is the GP prediction at point θ on the probability scale.

7 | HISTORY MATCHING AND VALIDATION

It is quite likely that large parts of parameter space will give rise to model outputs, which are incompatible with the observed proportions. In practice, we need our GPs to be accurate over regions of parameter space with nonnegligible posterior support and so time spent fitting them in regions of negligible posterior support is wasted. We can target the building of our GPs within regions of nonnegligible posterior support using methods known as *history matching* (Cumming and Goldstein, 2010; Vernon *et al.*, 2014). These methods work by determining which design points are implausible before undertaking the computationally intensive task of evaluating the model output (here empirical logits) at these points. After fitting a GP over, implausible design points are determined by comparing data points y_i with their GP prediction. An implausibility measure at data point y_i takes the form

$$I_i(\theta) = \frac{|y_i - m^*(\theta)|}{\sqrt{\sigma^2 + v^*(\theta) + 1/[n p_n^*(\theta)\{1 - p_n^*(\theta)\}]}}.$$

Adopting a conservative strategy, the implausibility measure over the entire dataset is calculated as $I(\theta) = \max_i I_i(\theta)$. Large values of $I(\theta)$ indicate that θ is an implausible point in parameter space. It is common to declare points as implausible if $I(\theta) > 3$, this threshold being determined using Pukelsheim’s 3-sigma rule. Notice that the implausibility measure requires a value to be given for the measurement error level. We take the upper 1% value of its prior distribution ($\sigma_D = 0.58$) as this possibly overestimates σ_D and in doing so lowers the chance of incorrectly declaring a point as implausible while still ruling out a significant proportion of parameter space.

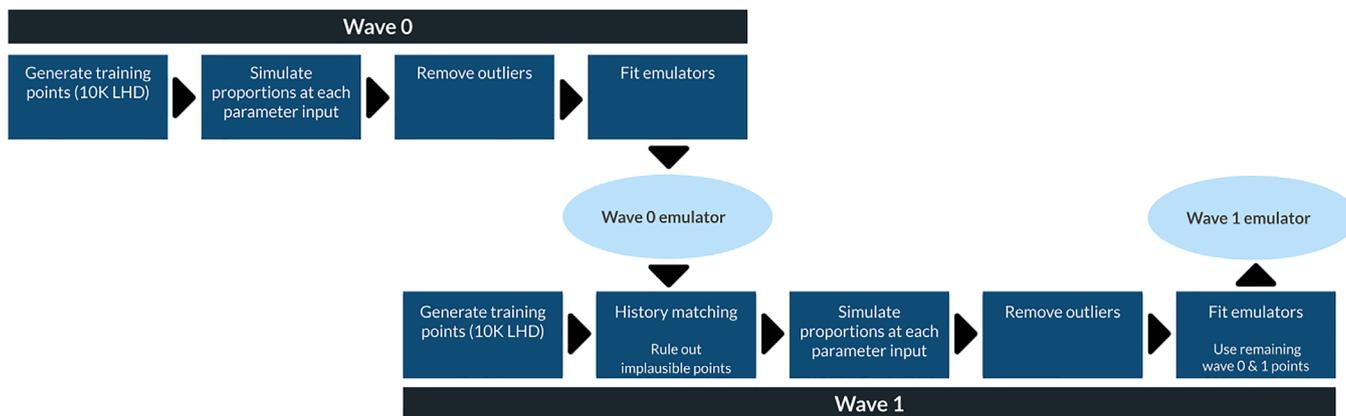


FIGURE 3 Flow chart showing the work flow involved in iteratively fitting Gaussian process emulators. This figure appears in color in the electronic version of this article, and any mention of color refers to that version

7.1 | Iterative fitting of Gaussian processes

We used a sequential procedure for fitting the GP in order to make best use of the limited computing resources available to simulate realizations from the stochastic model (Figure 3). In what follows, we provide a brief description of this procedure and refer the reader to Web Appendix C of the Supporting Information for further details.

We began in Wave 0 of the fitting procedure by constructing a 10K-point LHD within the cuboid defined by the 11 univariate marginal 99% prior equi-tailed intervals for each parameter θ_j . Multiple realizations ($n = 1000$) were then obtained from the stochastic model at these design points to determine the empirical logit proportions of cell death and of inclusion bodies. Next we eliminated those design points which gave proportions (of cell death or of inclusion bodies) outside the range (0.01,0.99), as such values look to be inconsistent with the data and they take considerable computing time to obtain. Finally, in Wave 0, we fitted a Gaussian process (GP_0) to the remaining 415 training points. Next, in Wave 1, we constructed an additional 10K-point LHD and removed any design points which had predicted probabilities (of the form $\text{eexpit}\{m_{GP_0}^*(\theta_i)\}$) outside (0.01,0.99). The points surviving Wave 0 were then added and any points with implausibilities $I_{GP_0}(\theta_i) > 3$ removed. Note that these implausibilities were calculated using the Gaussian processes GP_0 fitted at the end of Wave 0. Multiple realizations from the stochastic model were then simulated at the remaining 429 design points, empirical logits calculated, and Gaussian processes GP_1 fitted to this output. In practice, this process could be repeated to give many more waves of history matching, and thereby narrow down the plausible parameter space even further. However, due to our limited computer resources, we terminated our process with GP_1 .

7.2 | Emulator validation

Before using the GP_1 emulators as part of the inference scheme, it is important to verify that they provide an adequate description of model realizations of the empirical logits. There are a variety of diagnostic checks available based on comparisons of realizations from the stochastic model at a new set of design points (θ_i^\dagger) and comparing these with predictions made from the emulators; see, for example, Bastos and O'Hagan (2009) and Gneiting *et al.* (2007).

We constructed a validation design of a similar size to that used for fitting the GP_1 emulators by repeating the Wave 0–1 procedure described in the previous section. This approach gives an independent design, but still contains points in regions of nonnegligible posterior support, where we most need the emulators to fit well. For each GP, we calculated individual prediction errors (IPEs) at each design point as $D(\theta_i^\dagger) = \{x(\theta_i^\dagger) - m^*(\theta_i^\dagger)\} / \sqrt{v^*(\theta_i^\dagger)}$, where $x(\theta_i^\dagger)$ is an empirical logit (for a particular time-condition proportion) calculated from $n = 1000$ runs of the stochastic model with rates $\theta = \theta_i^\dagger$. Individual values of $D(\theta_i^\dagger)$ are informative about emulator fit, but it is particularly instructive to look at their overall distribution. If the assumptions underpinning the emulator are appropriate then the $D(\theta_i^\dagger)$ values should be a random sample from a standard normal distribution. In particular, we would expect that roughly 95% of the IPE values are within the interval $(-2, 2)$. If the magnitude of the IPE is too large, this indicates that, at this input point, the emulator either fits poorly or its variability is underestimated. Conversely, too many very small values are indicative of an inflated emulator variance. Gneiting *et al.* (2007) suggest that checks also be made using the probability integral transform (PIT) to assess the (standard) normality of the $D(\theta_i^\dagger)$ as, if this holds, values of $\Phi\{D(\theta_i^\dagger)\}$ should form

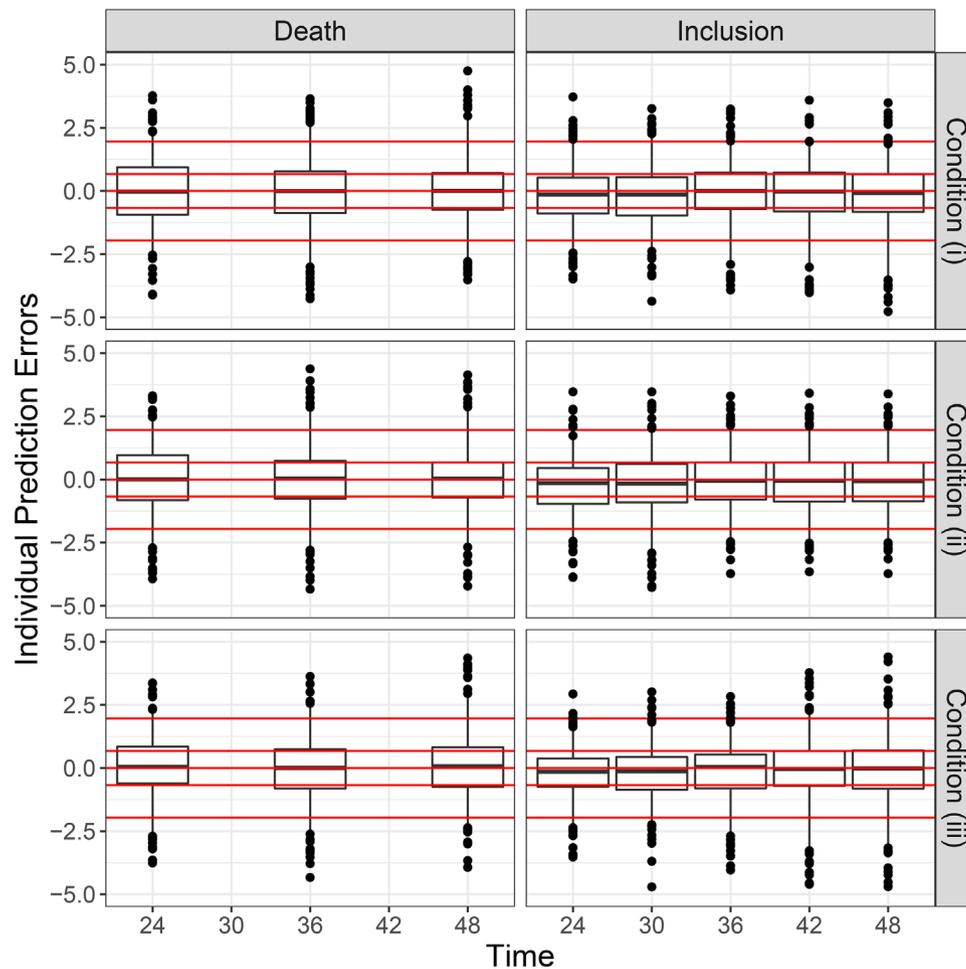


FIGURE 4 Boxplots of the individual prediction errors for each emulator together with horizontal lines showing the position of the quartiles and upper and lower 2.5% points of the standard normal distribution. This figure appears in color in the electronic version of this article, and any mention of color refers to that version

a random sample from a standard uniform distribution, where $\Phi(\cdot)$ is the standard normal distribution function. The IPEs for all 24 time-condition emulators are shown in Figure 4. Rather than produce additional PIT plots we give boxplot summaries for each emulator and include horizontal lines showing the positions of the quartiles and upper and lower 2.5% points of the standard normal distribution. These and other plots (see Web Appendix D of the Supporting Information) suggest that there are no strong departures from the univariate normal assumptions (1) used to form the likelihood function.

8 | POSTERIOR SUMMARIES AND CONCLUSIONS

To construct the emulators, training data (consisting of stochastic realizations from the model) were generated in parallel on a high-performance computer (HPC), with 64 compute nodes, where each node had at least 128 GB of

RAM and a minimum processor speed of 2.5 GHz. Hyperparameter inference was performed via the RStan interface to Stan, on a different HPC (with 23 cores and a minimum processor speed of 2.70 GHz). The total CPU time to obtain the fitted emulators was approximately 1 day.

Realizations from the joint posterior distribution of all parameters of interest were obtained by using the fitted emulators and again implemented via the RStan interface to Stan.

This MCMC algorithm uses the “no-U-Turn” Hamiltonian Monte Carlo sampler which efficiently explores the parameter space to maximize mixing. We ran three chains, each for 100K iterations, and using different initializations. For each chain, the first half (50K iterates) were discarded as warm-up/adaptation and the remaining half were thinned to give 1K (almost) unautocorrelated realizations from the posterior distribution. Convergence was assessed by a variety of methods, including graphical methods; see Web Appendix E in the Supporting Information for example traceplots, autocorrelation plots, and an

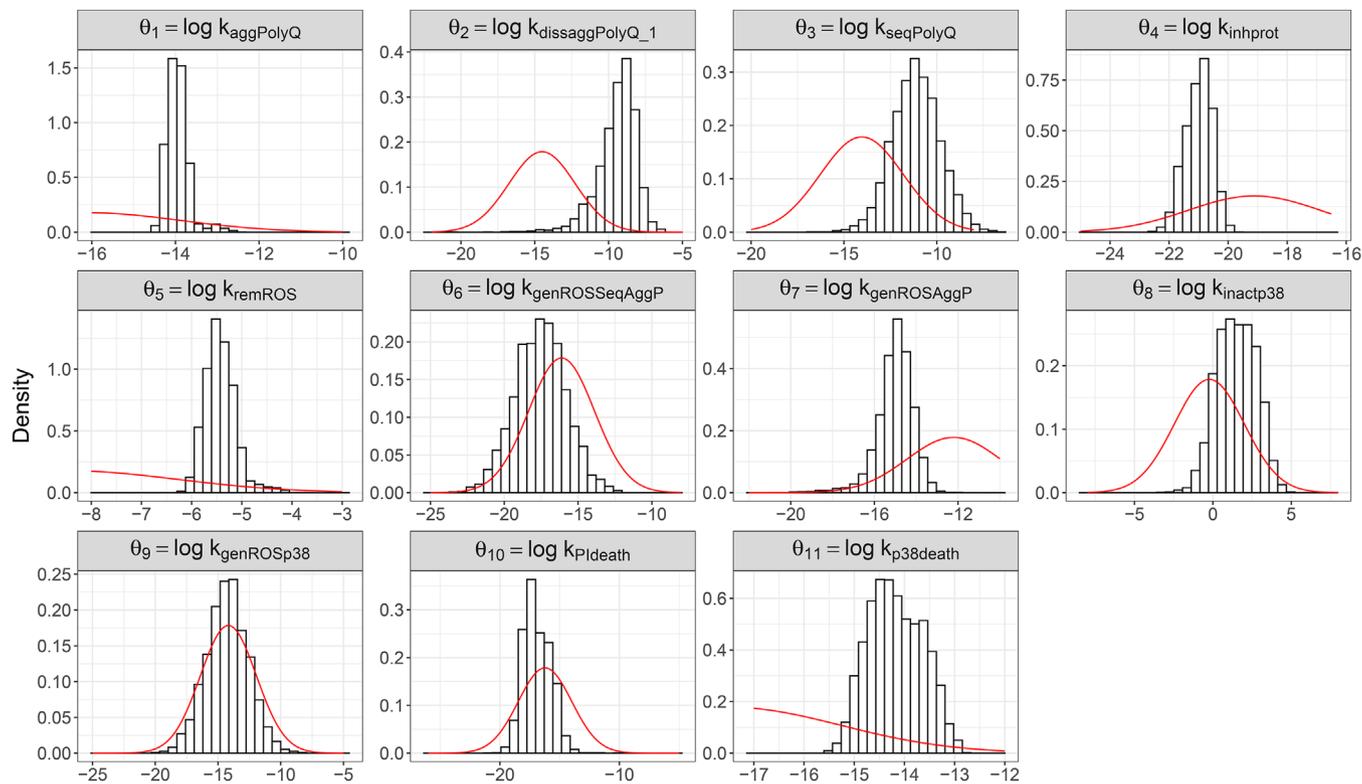


FIGURE 5 Marginal prior and posterior distributions for the (logged) stochastic kinetic rate parameters (histogram – posterior; red curve – prior). This figure appears in color in the electronic version of this article, and any mention of color refers to that version

assessment of model fit. We also examined diagnostics provided by the *shinystan* application (Gabry, 2018). In particular, we looked at the potential scale reduction statistic \hat{R} , which is the ratio of the within-chain variance and between-chain variance of the posterior sample (Gelman and Rubin, 1992). Here an \hat{R} -value close to one indicates that all chains have reached equilibrium, and in our posterior sample, all parameters had values less than 1.1. We also found that, for all parameters, the Monte Carlo standard errors was less than 10% of the posterior standard deviation and the effective sample sizes greater than 10% of the total sample size. Running the MCMC scheme for 100K iterations took approximately 7 days. Hence, the total computational cost (of obtaining the fitted emulators and performing the final calibration task) is of the order of 8 days. For comparison, consider use of the simulator directly inside an MCMC scheme. If we assume that a single iteration takes 60 seconds, then 100K iterations would take 70 days. We note that this estimate is conservative, as in reality, the mixing of such a scheme may necessitate many more iterations. Therefore, we expect that our use of emulators gives between one and two orders of magnitude increase in computational savings.

Graphical summaries of the marginal posterior distributions of the (logged) kinetic rates are given in Figure 5. The figure shows that, despite the experimental data being

thought to perhaps give only a very partial insight into the biological mechanism, it has in fact been very informative about many parameters. The analysis has confirmed values, with very much increased precision, for the rates of cell death due to activation of p38 ($k_{p38death}$) and due to inhibition of the proteasome by aggregates ($k_{PIdeath}$), which were also observed but with no information on cause of death. It is of particular note that the posterior distribution suggests that the early stages of the aggregation process, when it is reversible, probably occur at a much faster rate than had previously been assumed. This suggests that both the disaggregation of small aggregates and the formation and early growth of aggregates are more rapid. Also the ratio of disaggregation and aggregation rates ($k_{dissaggPolyQ_1}/k_{aggPolyQ}$) is around an order of magnitude higher than was thought. This suggests on the one hand that it may take longer to reach the threshold size required for inclusion formation but also that there will be more small aggregates present which will inhibit the proteasome. The decrease in posterior mean also suggests that proteasome inhibition ($k_{inhprot}$) is likely to occur at a slower rate than first thought; probably due to originally underestimating the production of small aggregates.

Looking at the rates involved in ROS turnover, the posterior suggests that the rate at which ROS is removed (kremROS) is faster and that less ROS is generated by

small aggregates (kgenROSaggP) than was thought, which means that the generation of ROS via the p38 pathway probably plays a larger contribution. This confirms the suggestion by Tang *et al.* (2010) that p38 inhibitors, an experimental intervention that they tested, have therapeutic potential to reduce the detrimental effects of the aggregation process.

ACKNOWLEDGMENTS

This paper arose from work in the PhD project of Holly Fisher (née Ainsworth) who was funded by the Engineering and Physical Sciences Research Council. We thank Doug Gray (University of Ottawa) for supplying us with the data and for discussions leading to the calibration of our prior distribution.

DATA AVAILABILITY STATEMENT

The data and code that support the findings in this paper are available in the Supporting Information of this article

ORCID

H. F. Fisher  <https://orcid.org/0000-0002-8389-9140>

C. S. Gillespie  <https://orcid.org/0000-0003-1787-0275>

A. Golightly  <https://orcid.org/0000-0001-6730-1279>

REFERENCES

- Aoki, K., Takahashi, T., Tabata, S., Kurata, M., Matsushita, A., Nagai, K. and Ishikawa, T. (2013) Efficacy and tolerability of reduced-dose 21-day cycle rituximab and cyclophosphamide, doxorubicin, vincristine and prednisolone therapy for elderly patients with diffuse large B-cell lymphoma. *Leukemia & Lymphoma*, 54, 11, 2441–2447.
- Arrasate, M., Mitra, S., Schweitzer, E.S., Segal, M.R. and Finkbeiner, S. (2004) Inclusion body formation reduces levels of mutant huntingtin and the risk of neuronal death. *Nature*, 431, 805–810.
- Bastos, L.S. and O’Hagan, A. (2009) Diagnostics for Gaussian process emulators. *Technometrics*, 51, 425–438.
- Bayarri, M.J., Berger, J.O., Paulo, R., Sacks, J., Cafeo, J.A., Cavendish, J., Lin, C.-H. and Tu, J. (2007) A framework for validation of computer models. *Technometrics*, 49, 138–154.
- Boys, R.J., Wilkinson, D.J. and Kirkwood, T.B.L. (2008) Bayesian inference for a discretely observed stochastic kinetic model. *Statistics and Computing*, 18, 125–135.
- Cao, Y., Gillespie, D.T. and Petzold, L.R. (2006) Efficient step size selection for the tau-leaping simulation method. *Journal of Chemical Physics*, 124, 044109.
- Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P. and Riddell, A. (2017) Stan: A probabilistic programming language. *Journal of Statistical Software, Articles*, 76, 1–32.
- Cumming, J.A. and Goldstein, M. (2010) Bayes linear uncertainty analysis for oil reservoirs based on multiscale computer experiments. In: O’Hagan, A. and West, M. (Eds.) *The Oxford Handbook of Applied Bayesian Analysis*. Oxford: Oxford University Press, pp. 241–270.
- Gabry, J. (2018). *shinystan: Interactive Visual and Numerical Diagnostics and Posterior Analysis for Bayesian Models*. R package version 2.5.0.
- Gan, L., Cookson, M.R., Petrucelli, L. and Spada, A.R.L. (2018) Converging pathways in neurodegeneration, from genetics to mechanisms, *Nature Neuroscience*, 21, 1300–1309.
- Gelman, A. and Rubin, D.B. (1992) Inference from iterative simulation using multiple sequences. *Statistical Science*, 7, 457–472.
- Georgoulas, A., Hillston, J. and Sanguinetti, G. (2016) Unbiased Bayesian inference for population Markov jump processes via random truncations. *Statistics and Computing*, 27, 991–1002.
- Gillespie, D.T. (1977) Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, 81, 2340–2361.
- Gillespie, D.T. (2000) The chemical Langevin equation. *Journal of Chemical Physics*, 113, 297–306.
- Gillespie, D.T. (2001) Approximate accelerated stochastic simulation of chemically reacting systems. *Journal of Chemical Physics*, 115, 1716.
- Gneiting, T., Balabdaoui, F. and Raftery, A.E. (2007) Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69, 243–268.
- Golightly, A. and Wilkinson, D.J. (2011) Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface Focus*, 1, 807–820.
- Henderson, D.A., Boys, R.J., Krishnan, K.J., Lawless, C. and Wilkinson, D.J. (2009) Bayesian emulation and calibration of a stochastic computer model of mitochondrial DNA deletions in substantia nigra neurons. *Journal of the American Statistical Association*, 104, 76–87.
- Henderson, D.A., Boys, R.J., Proctor, C.J. and Wilkinson, D.J. (2010) Linking systems biology models to data: a stochastic kinetic model of p53 oscillations. In: O’Hagan, A. and West, M. (Eds.) *The Oxford Handbook of Applied Bayesian Analysis*. Oxford: Oxford University Press, pp. 155–187.
- Ibrahim, J.G. and Chen, M.-H. (2000) Power prior distributions for regression models. *Statistical Science*, 15, 46–60.
- Koblets, E., Mariño, I.P. and Míguez, J. (2019) Bayesian computation methods for inference in stochastic kinetic models. *Complexity*, 2019. Article ID 7160934, 15 pages.
- Koblets, E. and Míguez, J. (2015) A population Monte Carlo scheme with transformed weights and its application to stochastic kinetic models. *Statistics and Computing*, 25, 407–425.
- Li, C., Donizelli, M., Rodriguez, N., Dharuri, H., Endler, L., Chelliah, V., Li, L., He, E., Henry, A., Stefan, M.I., Snoep, J.L., Hucka, M., Novère, N.L. and Laibe, C. (2010) BioModels Database: An enhanced, curated and annotated resource for published quantitative kinetic models. *BMC Systems Biology*, 4, Article number: 92
- Lieberman, A.P., Shakkottai, V.G. and Albin, R.L. (2019) Polyglutamine repeats in neurodegenerative diseases, *Annual Review of Pathology: Mechanisms of Disease*, 14, 1–27.
- Liepe, J., Kirk, P., Filippi, S., Toni, T., Barnes, C.P. and Stumpf, M.P.H. (2014) A framework for parameter estimation and model selection from experimental data in systems biology using approximate Bayesian computation. *Nature Protocols*, 9, 439–456.
- McColgana, P. and Tabrizi, S. J. (2018): Huntington’s disease: a clinical review. *European Journal of Neurology*, 25, 24–34.
- Morris, M.D. and Mitchell, T.J. (1995) Explanatory designs for computational experiments. *Journal of Statistical Planning and Inference*, 43, 381–402.

- O'Hagan, A. (2006) Bayesian analysis of computer code outputs: a tutorial. *Reliability Engineering and System Safety*, 91, 1290–1300.
- Persichetti, F., Carlee, L., Faber, P.W., McNeil, S.M., Ambrose, C.M., Srinidhi, J., Anderson, M., Barnes, G.T., Gusella, J.F. and MacDonald, M.E. (1996) Differential expression of normal and mutant Huntington's disease gene alleles. *Neurobiology of Disease*, 3, 183–190.
- Rasmussen, C.E. and Williams, C.K.I. (2006). *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press.
- Ross, C.A. and Poirier, M.A. (2004) Protein aggregation and neurodegenerative disease. *Nature Medicine*, 10, S10–S17.
- Santner, T.J., Williams, B.J. and Notz, W.I. (2003). *The Design and Analysis of Computer Experiments*. New York: Springer.
- Sato, A., Okada, M., Shibuya, K., Watanabe, E., Seino, S., Narita, Y., Shibui, S., Kayama, T., Kitanaka, C. (2014) Pivotal role for ROS activation of p38 MAPK in the control of differentiation and tumor-initiating capacity of glioma-initiating cells. *Stem Cell Research*, 12(1), 119–131.
- Stan Development Team (2018) RStan: the R interface to Stan. R package version 2.18.2.
- Tang, M.Y., Proctor, C.J., Woulfe, J. and Gray, D.A. (2010) Experimental and computational analysis of polyglutamine-mediated cytotoxicity. *PLoS Computational Biology*, 6, 1–14.
- Toni, T., Welch, D., Strelkowa, N., Ipsen, A. and Stumpf, M.P.H. (2009) Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems authors. *Journal of the Royal Society Interface*, 6, 187–202.
- van Kampen, N.G. (2001). *Stochastic Processes in Physics and Chemistry*. Amsterdam: North-Holland.
- Vernon, I., Goldstein, M. and Bower, R. (2014) Galaxy formation: Bayesian history matching for the observable universe. *Statistical Science*, 29, 81–90.
- Wilkinson, D.J. (2018). *Stochastic Modelling for Systems Biology*, Third edition. London: Chapman & Hall/CRC Mathematical & Computational Biology. Taylor & Francis.
- Wu, Q., Smith-Miles, K. and Tian, T. (2014) Approximate Bayesian computation schemes for parameter inference of discrete stochastic models using simulated likelihood density. *BMC Bioinformatics*, 15, S3.

SUPPORTING INFORMATION

Web Appendix 1 referenced in Sections 3, 7 and 8 and all associated computer codes are available with this paper at the Biometrics website on Wiley Online Library.

How to cite this article: Fisher HF, Boys RJ, Gillespie CS, Proctor CJ, Golightly A. (2022) Parameter inference for a stochastic kinetic model of expanded polyglutamine proteins. *Biometrics*, 78, 1195–1208. <https://doi.org/10.1111/biom.13467>