---

**Article:**
Andrianakis, I., Vernon, I., McCreesh, N. et al. (5 more authors) (2017) History matching of a complex epidemiological model of human immunodeficiency virus transmission by using variance emulation. Journal of the Royal Statistical Society: Series C (Applied Statistics), 66 (4). pp. 717-740. ISSN 0035-9254

https://doi.org/10.1111/rssc.12198

---

# History matching of a complex epidemiological model of HIV transmission using variance emulation.

I. Andrianakis†, I. Vernon, N. McCreesh, T.J. McKinley, J.E. Oakley, R.N. Nsubuga, M. Goldstein and R.G. White

## Summary

Complex stochastic models are commonplace in epidemiology, but their utility depends on their calibration to empirical data. History matching is a (pre-)calibration method that has been applied successfully to complex deterministic models. In this work, we adapt history matching to stochastic models, by emulating the variance in the model outputs, and therefore accounting for its dependence on the model's input values. The proposed method is applied to a real complex epidemiological model of HIV in Uganda with 22 inputs and 18 outputs, and is found to increase the efficiency of history matching, requiring 70% of the time and 43% fewer simulator evaluations compared to a previous variant of the method. The insight gained into the structure of the HIV model, and the constraints placed upon it, are then discussed.

**Keywords**: Calibration, Gaussian processes, Stochastic simulators, Inverse problems, Individual based models.

†Corresponding author: LSHTM, Keppel Street, London, WC1E 7HT
email: andrianakis@yahoo.com

## 1. Introduction

Mathematical modelling has played a large role in informing our understanding of infectious disease transmission and epidemiology. In the field of HIV, it has been used to investigate the role of partnership concurrency (overlapping sexual partnerships) on HIV transmission (McCreesh et al., 2012), estimate the contribution of acute, early stage infection to overall transmission (Powers et al., 2013), and estimate the proportion of transmission that occurs outside of cohabiting partnerships (Bellan et al., 2013). Modelling can also be used to inform policy, by allowing the effects of different control interventions to be estimated and compared, without the need for expensive and time-consuming randomised control trials. For instance, modelling has been used to predict the effects of making antiretroviral therapy (ART) universally available to people living with HIV, regardless of how far their disease has progressed (Granich et al., 2009), and estimating the effects of expanding access to ART and/or pre-exposure prophylaxis in men who have sex with men in the UK (Punyacharoensin et al., 2016).

In this study, we analyse a mathematical model of HIV transmission and partnership concurrency, called Mukwano, developed at the London School of Hygiene and Tropical Medicine. It is an individual based model with 22 inputs and 18 outputs, and is also stochastic, meaning that repeated evaluations for the same input parameters do not return the same output, but rather samples from a distribution with unknown characteristics. The usefulness of this and other models depends on our ability to calibrate them to measured empirical data (Grimm et al., 2006; May, 2004). Calibration is a type of inverse problem that attempts to estimate the input parameters of a system, such that its outputs are consistent with the available empirical data. Poor calibration results in a mathematical model that does not accurately reflect what we know about the current situation, greatly reducing our ability to make projections into the future. Poor calibration can also result in the amount of uncertainty in future projections being underestimated, leading to over-confident predictions being made, and potentially harmful policy decisions.

Calibration approaches range from simple least squares estimation techniques, to advanced probabilistic methodologies. Markov chain Monte Carlo (MCMC)-based techniques (Gibson and Renshaw, 1998; O'Neill and Roberts, 1999) are popular calibration methodologies. However, they tend to require the calculation of the likelihood function, which in the case of Mukwano is not available, while a data augmentation approach would require a numerical integration over a very large hidden state space. In smaller scale models, simulation based techniques, based on repeated evaluations of the simulator (Toni et al., 2009; McKinley et al., 2009; Andrieu et al., 2010) have been applied with some success. The simulator we are analysing in this work has a large number of inputs and outputs, and would require a large number of evaluations because of a) the high dimensionality of the input space, and b) the part of that space that matches the empirical data can be very small, due to the multiple constraints imposed by the large number of outputs. Furthermore, Mukwano is a stochastic

simulator, which requires multiple evaluations for each set of inputs in order to extract statistics about its output values, such as means and variances. Finally, the above methods attempt to make inferences over the entire input space using all available outputs at once. This can be an unnecessarily complicated task, as it requires simultaneously capturing the behaviour of all outputs in parts of the input space that are very far from the 'region of interest'.

History matching (Craig et al., 1997) is a form of calibration methodology, sometimes referred to as a pre-calibration step, that is designed to address the above problems. It is based on the use of an 'emulator', a statistical model of the simulator that is fast to evaluate (Sacks et al., 1989), and is therefore less disadvantaged by long simulator running times. It works by rejecting the input space where the simulator does not match the data, rather than the other way around. As a result, the entire set of outputs does not have to be taken into account at once, thus reducing the burden of analysing a large number of possibly complex outputs simultaneously, as is required by more traditional approaches. Finally, it focusses in on the region of interest in a series of iterations (waves), bypassing the need to model all of the simulator's outputs over all of its input space, and benefitting from the fact that Mukwano is expected to be 'well-behaved' in smaller input space regions. These characteristics make history matching particularly suitable as a pre-calibration step for simulators with large numbers of inputs and outputs, and long evaluation times which make the direct application of other calibration methodologies nearly impossible. Additionally, it may be viewed as an appropriate analysis methodology for simulators that are not considered accurate enough to warrant a full Bayesian analysis, which is much more computationally expensive.

History matching is not only useful for producing a large number of calibrated input samples. A careful study of the patterns that appear in the input and output spaces can be very informative about the way the simulator models various processes as well as the effect the constraints imposed by the empirical data have on the structure of the non-implausible space. These features of history matching illustrate the way Mukwano handles the HIV transmission process and how the empirical data shape the values that input parameters, such as contact rates and concurrency parameters, are allowed to take.

In previous work (Andrianakis et al., 2015), history matching was applied to Mukwano, but a rough approximation was used to account for the stochastic variability in its outputs, which was found to unnecessarily slow the method's convergence. In the present work, we refine the treatment of stochastic outputs by explicitly emulating their variance in addition to their mean and improving the overall efficiency of history matching. We also study correlation patterns between calibrated input and output samples, which provide useful insights into various processes of the simulator.

The structure of this paper is as follows: Section 2 describes the stochastic simulator studied, which is a dynamic, individual based HIV simulator, calibrated with epidemiological and behavioural data from a cohort in Uganda. Section 3 describes history matching in its standard form, and introduces

the proposed adaptation that allows it to handle stochastic models. Section 4 presents the results of history matching with the proposed adaptation, and Section 5 concludes the paper.

## 2. Model description and problem setup

The simulator we analyse in this work, known as Mukwano, is a dynamic, stochastic, individual based computer model that simulates heterosexual sexual partnerships and HIV transmission (McCreesh et al., 2012). Each individual is represented by a number of characteristics, of which some remain constant during simulated life (e.g. gender and date of birth), whereas others can change (e.g. HIV status). Changes in personal characteristics result from events such as the start and the end of sexual relationships. These events are stochastic: if and when an event occurs is determined by sampling from appropriate probability distributions. To generate model outcomes for a simulated population, the characteristics of the simulated individuals are aggregated.

Births, deaths, partnership formation and dissolution, and HIV transmission are modelled using time-dependent rates. At birth, simulated individuals are assigned to one of two sexual activity groups ('high activity' and 'low activity'), and to one of two partnership concurrency groups ('high concurrency' and 'low concurrency'). Each sexual activity group has associated male and female sexual contact rates, which determine the rate at which individuals form new partnerships. The duration of each new partnership is determined by the activity group of one of the partners, chosen at random. If their activity group is high, the partnership will have a short duration. If it is low, the partnership will have a long duration.

Seven different HIV stages are simulated, as shown in Figure 1. The natural history of HIV before antiretroviral therapy (ART) is represented by four stages: primary, CD4$\geq$ 200 cells/$\mu$l, CD4$<$ 200 cells/$\mu$l pre-AIDS, and AIDS. Infected people move sequentially through the four stages, and each stage has an associated HIV transmission probability. After 2004, when ART first became available in the population we are modelling, simulated people can be in an additional three stages: ART from pre-AIDS, ART from AIDS, and AIDS from ART. Possible routes of progression through the seven stages are shown in Figure 1. The probability of moving from an non-ART stage to an ART stage increased between 2004 and 2008, representing the increasing availability of ART in the population over time.

Twenty behavioural and two epidemiological inputs are varied, including a mixing parameter, which determines the tendency for individuals to preferentially form partnerships with people in their own activity group, and an input which determines the duration of the long and short duration partnerships. Many behavioural inputs are permitted to take different values in each of three calendar time periods. This enables sexual behaviour to vary over time, and allows the simulator to be fitted to trends in HIV prevalence in the population. A full list of the 22 simulator inputs and their original plausible ranges is shown in Table 1.

The simulator is calibrated to 18 demographic, behavioural and epidemiological outputs that include male and female population sizes in 2008, and male and female HIV prevalences at three time points. They also include a number of outputs that ensure that the prevalence and incidence of

monogamous and concurrent sexual partnerships in the simulator closely match estimates from the empirical population. The empirical data were collected from a rural general population cohort in South-West Uganda. The cohort was established in 1989, and currently consists of the residents of 25 villages (Mulder et al., 1994a,b; Seeley et al., 1991). Every year, demographic information on the cohort is updated, the population is tested for HIV, and a behavioural questionnaire is conducted. In 2008, this included questions that allowed the prevalence of monogamous and concurrent short duration and long duration partnerships to be estimated. All 18 simulator outputs and their empirical data are shown in Table 2. The intervals given for each of the outputs represent the limits for an acceptable match, and we considered them to be 95% confidence intervals for the purposes of the calibration. Their mean value was therefore used to define the value of the empirical data $\mathbf{z}$, and their difference was considered to represent 4 times the square root of the observation error variance $V_o$.

As mentioned previously, the simulator is stochastic, and the variance of the outputs changes, sometimes drastically, with changes in the values of the input parameters. In Andrianakis et al. (2015), it was observed that not accounting properly for the variance in the outputs, and in particular their dependence on the input values, reduced the efficiency of history matching and limited the insight gained into Mukwano's structure and the consequences of the observational constraints. The methodological developments proposed in this work address this issue.

## 3. Methods - History matching

### *3.1. Overview*

History matching is a method that attempts to identify the part of the simulator's input space that is likely to result in matches between the simulator's outputs and the empirical data (observations) . This part of the input space is referred to as *non-implausible* and has a high probability of containing the vast majority of the input parameters' posterior mass. This space's complement is known as *implausible*, where matches between the outputs and the observations are highly unlikely to be found. History matching was first developed in the field of oil reservoir simulations (Craig et al., 1997), but has since been applied to the calibration of computer models in fields ranging from galaxy formation, oceanography, systems biology, and epidemiology (Vernon et al., 2010a; Vernon and Goldstein, 2010; Vernon et al., 2014; Goldstein et al., 2013; Williamson et al., 2013; Andrianakis et al., 2015).

History matching works in iterations, known as *waves*, where the implausible space is first identified and then discarded. Each wave focusses the search for implausible space in the space that was characterised as non-implausible in all previous waves; thus, the non-implausible space shrinks with each iteration. The implausibility of the input space is determined with the *implausibility measure*, which is a measure of the distance between the observations and the simulator's output when evaluated at input $\mathbf{x}$.

Even though the implausibility measure could be calculated using the simulator directly, this turns out to be impractical even for simulators of moderate complexity. The reason is that the input space is high dimensional, and an exhaustive search would require a prohibitively large number of simulator evaluations. For this reason, fast surrogates of the simulator are used, which are known as *emulators*. An emulator is essentially a regression model, that predicts the simulator's output for a particular input $\mathbf{x}$, and is also capable of quantifying the uncertainty of these predictions. A key feature of an emulator is its almost negligible evaluation time. Gaussian processes (GP) are used to build the emulators in this work; however Bayes Linear models (Vernon and Goldstein, 2010) or simpler substitutes such as linear regression models could also be employed.

### *3.2. History matching of stochastic simulators (fixed variance)*

In this section, we present a summary of history matching of stochastic simulators, as presented in Andrianakis et al. (2015). The next section introduces the extensions to history matching that improve its efficiency on the calibration of Mukwano and stochastic simulators in general.

We suppose that the simulator has $P$ inputs denoted as $\mathbf{x} = [x_1, x_2, \ldots, x_P]^{\mathrm{T}}$, which are continuous and lie in a bounded subset $\mathcal{X} \subset \mathbb{R}^P$. The simulator also has $R$ outputs $f(\mathbf{x})$, the $r$-th of which is denoted by $f_r(\mathbf{x})$. Suppose also the existence of $\mathbf{z} = [z_1, z_2, \ldots, z_R]^{\mathrm{T}}$ observations, one for each simulator output, which typically come with their own uncertainty bounds (e.g. 95% confidence intervals).

Unlike deterministic simulators, which return the same value each time they are evaluated at the same input $\mathbf{x}$, stochastic simulators typically return draws from a distribution, which has a mean and a variance that we respectively denote by $g(\mathbf{x})$ and $s(\mathbf{x})$. We write the $r$-th output of the $k$-th evaluation of a stochastic simulator at input $\mathbf{x}$ as

$$f_{r,k}(\mathbf{x}) = g_r(\mathbf{x}) + \epsilon_{r,k}(\mathbf{x}), \tag{1}$$

where $g_r(\mathbf{x})$ is the mean and $\epsilon_{r,k}(\mathbf{x})$ is a zero mean noise term with variance $s_r(\mathbf{x})$.

At each wave, the simulator is evaluated $K$ times at each of the $N$ design points; an estimate of the mean simulator's response is

$$\hat{g}_r(\mathbf{x}_n) = \frac{1}{K}\sum_{k=1}^{K} f_{r,k}(\mathbf{x}_n). \tag{2}$$

An emulator of the simulator's mean output is then built using the training points $D = \{\mathbf{x}_n, \hat{g}_r(\mathbf{x}_n)\}$, for all the outputs that this is possible, which at wave $\eta$ are denoted as $r \in R_\eta$. We write the emulator's prediction for the mean output as $\mathrm{E}^*[g_r(\mathbf{x})]$ and the uncertainty of the prediction (variance) as $V_{c,r}(\mathbf{x})$.

An estimate of the simulator's variance at each of these points is

$$\hat{s}_r(\mathbf{x}_n) = \frac{1}{K-1}\sum_{k=1}^{K}(f_{r,k}(\mathbf{x}_n) - \hat{g}_r(\mathbf{x}_n))^2. \tag{3}$$

Andrianakis et al. (2015) used the 90-th percentile of the variances $\hat{s}_r(\mathbf{x}_n)$, $n = 1, \ldots, N$, which we denote by $V_{s,90}$, in the calculation of the implausibility measure.

The implausibility measure for the $r$-th output, is formulated as the distance between $z_r$ and $\mathrm{E}^*[g_r(\mathbf{x})]$, weighted with the uncertainty introduced by the error terms that link the two quantities.

We assume that $z_r$ is a noisy measurement from an underlying, unobservable physical process $y_r$, with their relationship described by

$$z_r = y_r + \phi_r, \tag{4}$$

where $\phi_r$ is a random variable that follows a unimodal distribution with zero mean and variance $V_{o,r}$. Its variance can be derived from considerations of the measurement process and as it links the (unobserved) physical process $y_r$ and the measurements $z_r$, has no dependence on the simulator's inputs $\mathbf{x}$.

The physical process $y_r$ is linked to a single realisation of the simulator $f_{r,k}(\mathbf{x})$ via the model discrepancy $\delta_r$, using

$$y_r = f_{k,r}(\mathbf{x}^*) + \delta_r, \tag{5}$$

where $\mathbf{x}^*$ is known as the 'best input'. This discrepancy arises because simulators are virtually always simplifications of the physical process $y_r$ (reality), either because we do not fully understand $y_r$, and therefore cannot model it exactly, or because some parts have been deliberately left out of the

modelling process. Accounting for model discrepancy can protect against overfitting the (potentially) wrong values of $\mathbf{x}$, and makes the simulator's predictions more robust. For more on this point the reader can consult Kennedy and O'Hagan (2001); Goldstein and Rougier (2009); Vernon et al. (2010a); Brynjarsdottir and O'Hagan (2010). The $\delta_r$ term represents the model expert's beliefs about the simulator's deficiencies, and as such is subjective and should be treated with some caution. Methods for a structured elicitation of model discrepancy are discussed in Goldstein and Rougier (2009); Goldstein et al. (2013). In Andrianakis et al. (2015), as well as in the present work, $\delta_r$ is assumed to follow a unimodal distribution with zero mean and variance $V_{m,r}$; as $\delta_r$ is defined as the difference between $y_r$ and the simulator evaluated at its 'best value' $\mathbf{x}^*$, it is also $\mathbf{x}$-invariant.

Finally, because we are using an emulator in place of the actual simulator, we need to take into account the error between the emulator's prediction $\mathrm{E}^*[g_r(\mathbf{x})]$ and the simulator's mean output $g_r(\mathbf{x})$, which we denote by $\zeta_r(\mathbf{x}) = g_r(\mathbf{x}) - \mathrm{E}^*[g_r(\mathbf{x})]$. $\zeta_r(\mathbf{x})$ is also assumed to be unimodal with zero mean and variance $V_{c,r}(\mathbf{x})$.

Combining the above with Equations 1, 4, 5 the link between the observed data $z_r$ and the emulator's prediction $\mathrm{E}^*[g_r(\mathbf{x})]$ is

$$z_r = \mathrm{E}^*[g_r(\mathbf{x}^*)] + \phi_r + \delta_r + \zeta_r(\mathbf{x}^*) + \epsilon_{k,r}(\mathbf{x}^*).$$

Based on the above analysis, the implausibility measure for a single output $r$ at a given value of $\mathbf{x}$ is given by

$$I_r(\mathbf{x}) = \frac{|z_r - \mathrm{E}^*[g_r(\mathbf{x})]|}{(V_{o,r} + V_{m,r} + V_{c,r}(\mathbf{x}) + s_r(\mathbf{x}))^{1/2}}. \tag{6}$$

Equation 6 is a measure of the distance between the observation $z_r$ and the emulator's posterior mean $\mathrm{E}^*[g_r(\mathbf{x})]$, weighted by the square root of the variances of all the uncertainties we have considered so far.

In general, $s_r(\mathbf{x})$ is unknown unless the simulator is evaluated at $\mathbf{x}$. In Andrianakis et al. (2015) this was approximated by $V_{s,90}$, the 90-th percentile of the observed variances, which could be seen as a conservative estimate in the absence of more detailed information. The approximation of $s_r(\mathbf{x})$ with $V_{s,90}$ is clearly a rough one, because it essentially assumes that the variance of the simulator's output is constant w.r.t. the input $\mathbf{x}$ (fixed variance), which is not necessarily true. As a result, the rejection of input space is not as efficient as it would have been if more accurate estimates of the variance were available. In the present work, we aim to refine this approximation via the use of variance emulators. This will be described in Section 3.3.

The relaxed distributional assumptions we made for the uncertainty terms $\phi, \delta$ and $\zeta$, (i.e. simply that they have zero mean and are unimodal) allow the use of Pukelsheim's 3 sigma rule (Pukelsheim, 1994) to derive cutoff limits for the above implausibility measure, such that when the value of $I_r(\mathbf{x})$ is larger than a cutoff $I_c$, then the input $\mathbf{x}$ can be considered implausible. Pukelsheim's 3 sigma rule

is a powerful (and underused) result which states that any continuous, unimodal distribution has at least 95% of its probability mass within 3 standard deviations, regardless of any asymmetry in the distribution. That is, if $\mathbf{x}$ was indeed the best input $\mathbf{x}^*$, then $I(\mathbf{x})$ should be $< 3$ with probability 95% for even the most asymmetric (but still unimodal) distributions that contribute to $I(\mathbf{x})$, and with a probability much higher than 95% for more symmetric, less unusual cases. Therefore, if $I_r(\mathbf{x}) > 3$ it suggests that we would be unlikely to obtain an acceptable match between outputs and observed data, were we to run the simulator at $\mathbf{x}$ (see Vernon et al. (2010a) for details). We should note here, that the parts of the input space for which $I_r(\mathbf{x}) < 3$ do not necessarily lead to matches between $z_r$ and $g_r(\mathbf{x})$, and hence do not necessarily represent a 'good' input $\mathbf{x}$: the implausibility can be small either because $z_r$ and $\mathrm{E}^*[g_r(\mathbf{x})]$ are close, or because there is still a large amount of uncertainty regarding the simulator's behaviour at $\mathbf{x}$. In other words, the denominator of $I_r(\mathbf{x})$ is still large.

The above single output implausibility measure has natural extensions to several outputs. One such extension is the maximum implausibility defined as:

$$I_M(\mathbf{x}) = \max_{r \in R_\eta}(I_r(\mathbf{x})), \tag{7}$$

where $R_\eta$ is the set of outputs that we wish to consider in wave $\eta$. Further extensions and analysis on implausibility measures can be found in Vernon et al. (2010a). Note that the above definition only involves a subset of the outputs as represented by the set $R_\eta$. Often, at early waves of the history match we would only emulate and construct implausibility measures for a small subset of the outputs, as some outputs may be very badly behaved over the whole input space. This subset $R_\eta$ would usually increase in size in later waves as we narrow the search to a smaller region of input space. This should be compared to a standard fully Bayesian or likelihood based analysis, where one has the difficult task of modelling all outputs simultaneously from the outset. This is a major strength of history matching.

### *3.3. History matching with variance emulation*

In the previous section, we claimed that using a fixed estimate for the simulator's variance reduces the efficiency of history matching, because it assumes that $s_r(\mathbf{x})$ is constant w.r.t. $\mathbf{x}$, which is not true in general. In this section, we are proposing a method that mitigates this problem by providing better estimates of the variance via emulation. The proposed method is based on an independent emulation of the mean $g(\mathbf{x})$ and variance $s(\mathbf{x})$. At each wave, the simulator is evaluated at $N$ points and the training data $\{\mathbf{x}_n, \hat{g}(\mathbf{x}_n)\}$ and $\{\mathbf{x}_n, \hat{s}(\mathbf{x}_n)\}$ are calculated using Equations 2 and 3. We start with a description of the emulator of the mean.

A GP emulator is built by considering a Gaussian process as a prior for the simulator's $r$-th output:

$$g_r(\mathbf{x}) \sim \mathcal{N}(h(\mathbf{x})\boldsymbol{\beta}, \sigma^2 c(\mathbf{x}, \mathbf{x}')). \tag{8}$$

The GP has a mean function $h(\mathbf{x})\boldsymbol{\beta}$, with $h(\mathbf{x})$ being a vector of deterministic functions of $\mathbf{x}$ (regressors), $\boldsymbol{\beta}$ a vector of regression coefficients, and $\sigma^2$ the variance of the process. The correlation function $c(\mathbf{x}, \mathbf{x}')$ can be a kernel, such as a Gaussian or a Matérn, that determines the correlation between $g_r(\mathbf{x})$ and $g_r(\mathbf{x}')$.

A key point here is that the training data $D_m \equiv \{\mathbf{x}_n, \hat{g}(\mathbf{x}_n)\}$ are not the actual mean outputs of the simulator, but estimates, the accuracy of which depends on the variance $s(\mathbf{x}_n)$ and the number of repetitions $K$. For this reason, we use a heteroscedastic noise component in the emulators of the mean, i.e. a noise term that will be different for each training point. The GP model of the training data is

$$g_r(D_m) \sim \mathcal{N}(H\boldsymbol{\beta}, \sigma^2 A + \text{diag}([\nu_1, \nu_2, \ldots, \nu_n]^{\text{T}})).$$

In the above equation, $H$ is an $(n \times q)$ matrix, whose $n$-th row is the polynomial $h(\mathbf{x}_n)$ from Equation 8. $A$ is a symmetric correlation matrix with entries $A_{i,j} = c(\mathbf{x}_i, \mathbf{x}_j)$. The noise components $\nu$ are calculated as $\nu_n = \hat{s}_n(\mathbf{x}_n)/K$ and the operator diag(.) transforms the column vector to a diagonal matrix.

The hyperparameters in the above expression are estimated using maximum likelihood (e.g. see Andrianakis et al. (2015); Rasmussen and Williams (2006)) and the above emulator provides an estimate of the simulator's mean value at an untested point $\mathbf{x}$ and an associated variance of the estimate. We denote these two quantities as $\text{E}^*[g_r(\mathbf{x})]$ and $V_{c,r}(\mathbf{x})$. All emulators are validated using a separate validation set of simulator runs, following the methods described in Bastos and O'Hagan (2009).

A similar procedure is followed for emulating the variances. First, we log-transform the variance data defining $\hat{\xi}(\mathbf{x}) \equiv \ln(\hat{s}(\mathbf{x}))$, resulting in the training data set $D_v = \{\mathbf{x}_n, \hat{\xi}(\mathbf{x}_n)\}$. This transformation is helpful, because $\hat{\xi}(\mathbf{x})$ is closer to a Gaussian distribution than $\hat{s}(\mathbf{x})$, and therefore easier to model using a GP. The prior for the training data is

$$\hat{\xi}_r(D_v) \sim \mathcal{N}(H\boldsymbol{\beta}, \sigma^2 A + \mathcal{I}\nu).$$

Since the data $D_v$ are also estimates we could have used the same heteroscedastic model we used for the emulators of the mean. However, this would require estimating the $\nu$'s using a fourth order statistic of the simulator runs (variance of the variance), which could be unstable unless we had a very large number of repetitions per design point. For this reason, we take a simpler approach, and assume that the noise level in the variance data is constant and equal to $\nu$, a hyper-parameter that is estimated along with the other hyper-parameters of the GP, using maximum likelihood.

The variance emulators provide an estimate of the log variance for any untried input $\mathbf{x}$ in the current non-implausible space, which we denote by $\text{E}^*[\xi_r(\mathbf{x})]$ and will be used in the implausibility defined in Section 3.3.1. Variance emulation in a regression setting has also been discussed in Hender-

son et al. (2009); Vernon and Goldstein (2010); Ankenman et al. (2010); Boukouvalas et al. (2014), but its integration within the history matching framework is studied here for the first time.

### *3.3.1. The implausibility measure (emulated variance)*

The implausibility measure for one output is again given by Equation 6. In this case, the variance emulators provide an improved approximation to $s_r(\mathbf{x})$, which is $s_r(\mathbf{x}) \approx \exp(\mathrm{E}^*[\xi(\mathbf{x})])$. This term is a function of $\mathbf{x}$, and should therefore be more accurate than the fixed $V_{s,90}$ used previously. Furthermore, in most cases it should hold that $\exp(\mathrm{E}^*[\xi(\mathbf{x})]) < V_{s,90}$, which results in larger implausibility values for a given $\mathbf{x}$ and, therefore, a more efficient space rejection. The proposed implausibility measure for one output takes the form:

$$I_r(\mathbf{x}) = \frac{|z_r - \mathrm{E}^*[g_r(\mathbf{x})]|}{(V_{o,r} + V_{m,r} + V_{c,r}(\mathbf{x}) + \exp(\mathrm{E}^*[\xi_r(\mathbf{x})]))^{1/2}}. \tag{9}$$

The above argument is illustrated in Figure 2. The top panel shows 100 simulator evaluations at each of 8 different design points (grey dots). The output studied is the 2007 female HIV prevalence. The horizontal black lines show the estimated mean of each design point, and the red lines represent $\pm 2$ standard deviations calculated with the second largest variance of the design points shown (similar to $V_{s,90}$). The green lines show $\pm 2$ standard deviations, calculated with the variance estimated from the actual 100 repetitions at each design point. The observations are shown with the horizontal blue line. The bars in the bottom panel show a simplified form of the implausibility $I = |z - \hat{g}(\mathbf{x})|/\hat{s}(\mathbf{x})^{1/2}$, calculated with the respective variances of the top panel. The horizontal black line is the cutoff implausibility value, which is set at 3. The figure shows that the overestimation of the variance for points 1,2,4,5 by the use of $V_{s,90}$ reduces their implausibility, such that they are either accepted or rejected marginally ($I(\mathbf{x}) \approx 3$). Improving the estimate of $s(\mathbf{x})$ increases the implausibility and allows rejecting those points with greater confidence. This toy example conveys the essence of the method we are proposing in this work.

### *3.3.2. Procedure*

The procedure of history matching using the emulated variance is presented below:

(a) Define the initial $P$-dimensional non-implausible space $\mathcal{X}_{\eta=0}$.
(b) Select $N$ training and $N'$ validation points from the current non-implausible space $\mathcal{X}_\eta$, using a space filling design.
(c) Evaluate the simulator $K$ times at each of the training and validation points; calculate the training data $D_m$ for the mean and $D_v$ for the variance.
(d) Build and validate an emulator for as many of the $g_r(\mathbf{x})$ as possible.
(e) Build and validate an emulator for as many of the $\xi_r(\mathbf{x})$ as possible.

(f) Evaluate the implausibility measure $I_M(\mathbf{x})$ for a large number of $\mathbf{x} \in \mathcal{X}_\eta$ such that the complexity of $\mathcal{X}_\eta$ is represented with sufficient accuracy. Use the single output implausibility from Equation 9. $\mathcal{X}_{\eta+1}$ is the set of $\mathbf{x} \in \mathcal{X}_\eta$ for which $I_M(\mathbf{x})$ is less than the chosen threshold.

(g) Increase wave counter $\eta$ by 1 and repeat steps (b) to (f), until:

(i) The emulator uncertainty $V_c$ is smaller than the other uncertainties (e.g. $V_o$ or $V_m$), so more waves would not reduce $\mathcal{X}_\eta$ further, or

(ii) A large number of simulator runs from the final wave's non-implausible space are sufficiently close to the observations for the needs of the application, or

(iii) All $\mathcal{X}_\eta$ has been characterised as implausible.

In the above sequence of non-implausible spaces, it holds that $\mathcal{X}_\eta \subset \mathcal{X}_{\eta-1} \subset \ldots \subset \mathcal{X}_0$. During this process of space reduction, $\mathcal{X}_\eta$ might lose properties such as convexity or connectivity. In general, history matching can handle non-convex spaces, as it can identify, for example, disconnected regions. Ideally, if disconnected regions were to be found in $\mathcal{X}_\eta$, they could be emulated separately. In most cases however, identifying such regions in high dimensional spaces is far from trivial. As for the GP emulators, these can be thought of as defined over the wider (convex) region, but we only choose to evaluate them, for the purpose of history matching, within the non-implausible space. For more on this point the reader can consult Vernon et al. (2010b).

History matching is also very efficient in dealing with models that are unidentifiable. Correlation ridges and multiple modes in the posterior, typical manifestations of identifiability issues, pose no problem to history matching, whereas they can plague other methods, including MCMC based ones. Finally, if the simulator is incapable of matching the observations, history matching will reject all the input space as implausible, flagging this condition, while other likelihood and simulation based methods will always attempt to return a posterior distribution, regardless of how poorly the simulator fits the data.

### *3.3.3. Further points - extensions*

The value of $N$ can be determined by the available computational resources, but a very rough rule of thumb suggests setting $N = 10P$, where $P$ is the number of inputs (Loeppky et al., 2009). The number of validation runs can be chosen as $N' \approx N/10$. The training and validation data are best selected using some space filling method, e.g. by maximising the minimum distance between points, such that they fill the entire input space. This type of design generally leads to emulators that can more accurately describe the simulators over most parts of the non-implausible space. Furthermore, simulator runs from previous waves can be used as training points for the present wave if they fall within or close to the current non-implausible region.

The number of repeated evaluations $K$ of the simulator at each design point is considered fixed

throughout this work. The value of $K$ can be chosen such that the variance of the estimator $\hat{g}(\mathbf{x})$, which is given by $\hat{s}(\mathbf{x})/K$, is smaller than the observation error, and also by considering the computational budget that is available for running the simulator. Another approach would be to use a variable number of repetitions $K(\mathbf{x})$, such that the simulator is evaluated more times at the $\mathbf{x}$'s where $s(\mathbf{x})$ is expected to be large, and vice versa. The emulators of the variance can provide some guidance in this direction, as they can predict $s(\mathbf{x})$ at a new location $\mathbf{x}$. This possibility has been explored in a Kriging setting in Fedorov and Hackl (1997); Ankenman et al. (2010) and in an optimisation setting in Picheny et al. (2013). Although using a fixed number of repetitions is a simple and robust approach to the problem, a $K$ that varies with the input $\mathbf{x}$ could further increase the efficiency of history matching by reducing the total number of simulator evaluations.

The simulator's 18 outputs were modelled with independent univariate emulators. Note that it is the 'residual processes' $g(\mathbf{x}) - h(\mathbf{x})\boldsymbol{\beta}$ that are assumed to be independent between the different outputs. The independent emulators can therefore still capture strong correlations between outputs via the trend term $h(\mathbf{x})\boldsymbol{\beta}$, which often justifies putting more detail into it, as is discussed in Vernon et al. (2010a, 2014). Nevertheless, it is possible that the efficiency of history matching could be improved by using multivariate emulators if there are correlations between these residual processes. Emulators with separable covariance functions such as those in Rougier (2008) or Conti and O'Hagan (2010) could be used, although these require the same set of correlation length parameters to be used for each output. Multivariate emulators with non-separable covariance functions are discussed in Fricker et al. (2013), but these are more computationally demanding to fit.

## 4. History matching the Mukwano simulator

### *4.1. Comparison with the fixed variance approach*

In order to evaluate the benefits of including the variance emulators in history matching, we compare it to the history match shown in Andrianakis et al. (2015). In that work, the *fixed variance* approach was used. That is, the variance of the simulator's output $s(\mathbf{x})$ was not emulated, but was rather fixed to the 90th quantile of the estimated variances $\{\hat{s}(\mathbf{x}_n) : n = 1, \ldots, N\}$. In the proposed *emulated variance* methodology, the variance in the output of the stochastic model is emulated, and is therefore a function of the input $\mathbf{x}$.

To facilitate the comparison, both history matches were designed to be the same in terms of the number of simulator runs per wave, the number of simulator inputs and outputs, and the empirical data. The observation variance terms $V_o$ were also identical, and the model discrepancy was set in both cases equal to 10% of the variance of $\{\hat{g}(\mathbf{x}_n)\}$, $n = 1, \ldots, N$. Since the emulated variance approach essentially estimates $s_r(\mathbf{x})$ as a function of $\mathbf{x}$, instead of using a fixed and relatively large value, we expect that the efficiency of history matching will increase. In the following, we show that this is indeed the case, and demonstrate that the benefits gained by the inclusion of the variance emulators outweigh the extra computational effort of building them. A comparison of the costs in terms of CPU and user time is given in section 4.3

To quantify the closeness of an actual simulator run to the empirical data $\mathbf{z}$, we define the *simulator run implausibility* for a single output as

$$I_{\mathcal{R},r}(\mathbf{x}) = \frac{|z_r - \hat{g}_r(\mathbf{x})|}{(V_{o,r} + V_{m,r} + \hat{s}_r(\mathbf{x}))^{1/2}}. \tag{10}$$

Note that no emulators are involved in this metric, and it is not part of the history matching algorithm; it is only a metric that quantifies how close the simulator's output is to the empirical data, when evaluated at input $\mathbf{x}$. The *overall* simulator run implausibility is defined as the maximum of $I_{\mathcal{R},r}(\mathbf{x})$ across all outputs, i.e. $I_{\mathcal{R}}(\mathbf{x}) = \max_r(I_{\mathcal{R},r}(\mathbf{x}))$.

Figure 3 shows the empirical cumulative distribution of the simulator run implausibility at each wave. This figure can also be interpreted as the proportion of each wave's simulator runs with an implausibility $I_{\mathcal{R}}(\mathbf{x})$ smaller than the value indicated in the horizontal axis. Panel (a) shows the runs from Andrianakis et al. (2015) (using the fixed variance approach), and panel (b) the runs from the proposed methodology. The figure shows that in Andrianakis et al. (2015) the non-implausible region contained 65% of non-implausible runs after 9 waves. Under the current methodology, the same target was reached after 6 waves, a substantial improvement.

Figure 4 shows the (log10) proportion of the original input space that is calculated as non-implausible after each wave, using both methodologies. This figure again shows that the addition of the variance emulators causes the non-implausible space to shrink by a larger amount at each

wave. Especially in later waves, the rate of space reduction is much higher under the proposed approach. The main reason for this is that $s_r(\mathbf{x})$ dominates the uncertainties in these stages, and having improved estimates from the use of variance emulators allows the space to shrink faster.

### *4.2. Results and insights into the Mukwano simulator*

We now discuss the insights generated by our analysis of the Mukwano simulator. A practical way of visualising the reduction of the non-implausible space over consecutive history matching waves is via the minimum implausibility and optical depth plots. The first type of plot shows, for a grid of values for 2 selected inputs $x_1$ and $x_2$, an estimate of minimum implausibility for an input $\mathbf{x}$ if we were to fix $x_1$ and $x_2$ to a specific value and vary the remaining components $x_i$ $i = 3, \ldots, p$. The optical depth plots show an estimate of the probability of obtaining a non-implausible value were we to fix $x_1$ and $x_2$ to specific values and sample from the remaining elements of $\mathbf{x}$ (see Vernon et al. (2010a) for details).

Figure 5 shows an example of these plots for the high activity contact rate (*hacr1*) and the proportion of men in the high activity group (*mhag*) inputs. These plots show that if both inputs take a large value, it is very unlikely that the outputs will match the empirical data, as indicated by the high minimum implausibility value and the low optical depth on the upper right corners of both figures. This is consistent with behaviour data from the study population in rural Uganda. When both parameters have large values, a high proportion of men will be in the high activity group, and these men will form partnerships at a high rate. This will result in there being too many partnerships in the simulator, and the proportion of men and women with one and/or 2+ partnerships will be above the plausible ranges for the associated outputs (Outputs 10-18).

The optical depth plots indicate the regions where most of the non-implausible input space can be found (essentially the depth of the non-implausible space conditioned on the two inputs used for the axes of the plot). In this case, as is shown in the right panel of Figure 5, this occurs where both *hacr1* and *mhag* are low. This can often be due to a large number of mediocre input points, and therefore a naive search of the input space may be more likely to find solutions within this region. The minimum implausibility plots show the two-dimensional projection of the regions of input space that can be discarded by different cutoffs, and give an indication, especially in later waves, of where the most promising input points may lie. In this example, this is when *hacr1* is low and *mhag* takes intermediate values, as shown in the left panel of Figure 5.

Figure 6 shows the combined minimum implausibility and optical depth plots for 10 out of the 22 inputs whose range was reduced the most after the history match. Panel (a) shows the rejected space after 9 waves using the fixed variance approach. Panel (b) shows the space reduction after 6 waves of the proposed methodology. The two dimensional projections of the non-implausible space are very similar, implying that the proposed methodology (emulated variance) achieved a similar input space

reduction as the fixed variance approach of Andrianakis et al. (2015), in 3 fewer waves.

Figure 6 also shows that the male concurrency parameter in the high concurrency group in the third risk period (*mchc3*) can take higher values when the high activity contact rate in the third risk period (*hacr3*) is lower. The values of these parameters are constrained by the need to fit the simulator to data on the point prevalence of men with 2+ short duration partnerships in 2008 (output 14), which suggest that no more than 2.1% of men have concurrent short duration partnerships at a given point in time. When the high activity contact rate is high, the probability of a man forming a second, concurrent partnership needs to be low, to prevent the point prevalence of concurrent short duration partnerships in the simulator being too high. At lower contact rates, this is relaxed slightly, and the probability of men forming additional partnerships can be higher.

Regarding the simulator's outputs Figure 7 shows 18 panels, one for each output given in Table 2. Each panel shows a scatter plot of the mean against the variance of the simulator runs at each wave. The vertical bands show the empirical data and the associated 2 standard deviations arising from the observation error and model discrepancy at wave 6, and hence represent the target of the history match. Two key conclusions can be extracted from this figure: first, as the history match progresses, the mean output of the simulator converges to the empirical data, as can be seen by the green dots that are centered around the empirical data patches. Second, the variance of the output is far from being constant, and varies not only in the first, but also in the later waves. This shows that the HIV transmission model has a non trivial variance dependence on $\mathbf{x}$, even across the tiny region of input space where good matches are to be found. This further justifies the need for estimating the variance, instead of using a fixed and crudely large estimate, as in Andrianakis et al. (2015).

Figure 8 shows scatter plots between simulator outputs at wave 7, which provide insight into how the simulator handles the HIV transmission process. Panel (a) shows a strong negative correlation (r=-0.96) between the prevalence of HIV in women in 1992, and female population size in 2008. This occurs because a high prevalence of HIV in 1992 greatly increases the mortality rate in the simulated population, decreasing population size. As the rate at which new people are born in the simulation is a function of the number of women in the simulated population, this also reduces population growth. Strong positive correlations are present between male and female HIV prevalences in the same year (panel (b), r=0.96). This occurs because heterosexual sex is the major route of HIV transmission in Uganda, and is the only route that is included in the simulator. As simulated men can therefore only be infected by women, and vice versa, men and female HIV prevalences are necessarily highly correlated. There are also strong correlations between HIV prevalences in different years (panel (c), r=0.94), reflecting the fact that HIV is infectious. This means that the rate of new infections is higher at higher HIV prevalences, and that the prevalence of HIV in any given year is likely to be higher if the prevalence of HIV was also high in earlier years.

### *4.3. Computational cost of the proposed methodology*

In order to allow a comparison of the computational costs between the proposed methodology and the fixed variance approach, the number of simulator runs was kept the same between the two methods at each wave. The number of design points per wave are shown in Table 3. The number of design points in the initial waves were approximately 250, following the recommendation of 10 design points per input (Andrianakis et al., 2015), a number that was doubled from wave 5 onwards in an attempt to improve the emulators and increase the space rejection. The simulator was run $K = 100$ times at each design point for estimating its mean and variance, the same number as in Andrianakis et al. (2015) to ensure consistency. The number of emulator evaluations per wave and in total can therefore be extracted by Table 3 by multiplying the quantities by $K = 100$.

We recorded the time it took to complete the simulator runs, shown in Table 4 in the row 'Simulator running time'. Note that these calculations assume a 100 core cluster (which is the average number of cores we had at our disposal) - running the simulator on a single core, but otherwise identical machine, would have taken 100 times as long. Runs in wave 1 took up to 4 times longer to complete than the runs in subsequent waves. This was because the simulator run times were longer in some very implausible areas of the input space. This is a common feature of computer models, but is more prominent in the stochastic case: we may see vastly different run times in different parts of the input space. Once we were aware of this, say after analysing the wave 1 runs, we could create designs that exploit this feature, but we leave this for future work.

We also recorded the time required for training the emulators, noting that to avoid local minima in the estimation of the emulator's hyperparameters, the optimisation routine was initialised from 20 different starting points (see Andrianakis et al. (2015) for more details). At each wave, the optimisation scheme was run $20R$ times where $R = 18$ is the number of outputs, and therefore the number of emulators we built. Table 4 also shows the time it took the cluster to run the optimisation routines, again assuming the existence of 100 cores. The second row of Table 4 refers to the training of the emulators of the mean, and the third row to the emulators of the variance.

Currently, history matching is not a fully automated procedure, and manual intervention is required from the user at two stages of the process at each wave. The first is to collect the data from the cluster and set up the emulators to be trained. The second involves using the 'newly built' emulators to identify the non-implausible space, and select the design points where the simulator is next to be run. Although large parts of this process could be further automated, we believe that a number of manual checks can ensure that the history match is converging to the right values, and can save time in the long run. We estimated that approximately 3 hours of staff (user) time were required at each of the two stages mentioned earlier. For the fixed variance approach, this time was approximately 2 hours, as variance emulators did not need to be built.

In summary, and as shown in the last row of Table 4, the use of variance emulators allowed

completing the history match around 70% of the time, taking approximately 9.4 days instead of 13.2. The total number of simulator evaluations was also brought down to 201700 from 351700, a 43% reduction.

## 5. Conclusion

In this work, we history matched Mukwano, an individual based stochastic simulator that models the transmission of HIV in the presence of concurrent relationships. The simulator is used for assessing the impact of concurrent relationships on the incidence and prevalence of the disease and to evaluate this effect relative to other changes in sexual behaviour. History matching allows calibrating Mukwano to empirical data, a step that is necessary before using the simulator to infer any epidemiological parameters or make predictions about the evolution of the disease.

The present paper addressed a shortcoming of history matching when applied to stochastic models, by explicitly emulating the variance of the outputs, an approach that increased the overall efficiency of the method. After 6 waves, history matching produced samples that matched all Mukwano's 18 outputs 70% of the time, while the parameter space was reduced by a factor of $10^{11}$. Furthermore, the variance emulation proposed here reduced the time required by the method to 70% compared to previous work and also reduced the number of simulator evaluations by 43%.

A study of the non-implausible space and the calibrated outputs provided useful insights into the simulator's structure. The constraints, imposed by the empirical data, were traced back to the inputs, finding that they created distinctive correlation patterns in the non-implausible space. For example, constraints in the number of concurrent partnerships meant that contact rates and concurrency parameters could not be simultaneously large or small. Additionally, strong correlations between the simulator's outputs at the last wave, illuminated the way the simulator handles specific aspects of the HIV transmission process. In particular, correlations between HIV prevalence and population size across time and genders revealed the major route of HIV transmission in the simulator and links between HIV prevalence and mortality. This type of analysis enhanced our understanding of Mukwano and can be very helpful in its further development.

Although history matching is a methodology for calibrating slow and high dimensional simulators, such as Mukwano, it is not geared towards making probabilistic statements about the posterior of the simulator's parameters, and instead should be viewed as a) a useful pre-calibration step to identify a small region of input space where the posterior will reside (while simultaneously checking that the simulator is fit for purpose, and hence that such a calibration is meaningful) or b) the appropriate analysis for model development and checking of a model not thought to be sufficiently accurate to warrant a full Bayesian analysis (see Vernon et al. (2010a) and the associated discussions for more details). In this way, history matching should not be thought of as a direct competitor to other calibration methods, but rather as a procedure that will help improve the efficiency of whatever subsequent technique one wishes to employ. An extension of this method would be to combine it with probabilistic calibration methods, which would typically be computationally infeasible if applied to the original input space of a simulator of Mukwano's complexity, but may be successful if they are applied to the greatly reduced non-implausible space that results from history matching.

## Acknowledgements

The authors would like to thank the editor and two reviewers for their helpful comments, which helped improve the presentation of this work.

## Funding

This work was funded by the UK Medical Research Council (MRC) and the UK Department for International Development (DFID) under the MRC/DFID Concordat agreement that is also part of the EDCTP2 programme supported by the European Union (MR/J005088/1). RGW is additionally funded by the Bill and Melinda Gates Foundation (TB Modelling and Analysis Consortium: OPP1084276) and UNITAID (4214-LSHTM-Sept15; PO #8477-0-600).

## References

Andrianakis, I., Vernon, I., McCreesh, N., McKinley, T. J., Oakley, J. E., Nsubuga, R., Goldstein, M. and White, R. G. (2015) Bayesian history matching and calibration of complex infectious disease models using emulation: a tutorial and a case study on HIV in Uganda. *PLoS Computational Biology*, **11**, 1–18.

Andrieu, C., Doucet, A. and Holenstein, R. (2010) Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society, Series B (Methodological)*, **72**, 269–342.

Ankenman, B., Nelson, B. and Staum, J. (2010) Stochastic kriging for simulation metamodeling. *Operations research*, **58**, 371–382.

Bastos, L. S. and O'Hagan, A. (2009) Diagnostics for Gaussian process emulators. *Technometrics*, **51**, 425–438.

Bellan, S., Fiorella, K. J., Melesse, D., Getz, W., Williams, B. and Dushoff, J. (2013) Extra-couple HIV transmission in sub-Saharan Africa: a mathematical modelling study of survey data. *The Lancet*, **381**, 1561–1569.

Boukouvalas, A., Cornford, D. and Stehlik, M. (2014) Optimal design for correlated processes with input-dependent noise. *Computational Statistics & Data Analysis*, **71**, 1088–1102.

Brynjarsdottir, J. and O'Hagan, A. (2010) Learning about physical parameters: The importance of model discrepancy. *Tech. rep.*, http://www.tonyogahan.co.uk/academic/pub.html.

Conti, S. and O'Hagan, A. (2010) Bayesian emulation of complex multi-output and dynamic computer models. *Journal of Statistical Planning and Inference*, **140**, 640–651.

Craig, P. S., Goldstein, M., Seheult, A. H. and Smith, J. A. (1997) Pressure matching for hydrocarbon reservoirs: a case study in the use of Bayes linear strategies for large computer experiments. *(with discussion) in Case Studies in Bayesian Statistics, eds. C. Gastonis et al. Springer-Verlag*, **III**, 37–93.

Fedorov, V. and Hackl, P. (1997) *Model-oriented design of experiments*. Berlin: Springer.

Fricker, T. E., Oakley, J. E. and Urban, N. M. (2013) Multivariate Gaussian process emulators with nonseparable covariance structures. *Technometrics*, **55**, 47–56.

Gibson, G. J. and Renshaw, E. (1998) Estimating parameters in stochastic compartmental models using Markov chain methods. *IMA Journal of Mathematics Applied in Medicine and Biology*, **15**, 19–40.

Goldstein, M. and Rougier, J. (2009) Reified Bayesian modelling and inference for physical systems. *Journal of Statistical Planning and Inference*, **139**, 1221–1239.

Goldstein, M., Seheult, A. and Vernon, I. (2013) Assessing model adequacy. In *Environmental Modelling: Finding Simplicity in Complexity, Second Edition* (eds. J. Wainwright and M. Mulligan). John Wiley & Sons, Ltd, Chichester, UK: Wiley-Blackwell.

Granich, R., Gilks, C., Dye, C., De Cock, K. M. and Williams, B. G. (2009) Universal voluntary HIV testing with immediate antiretroviral therapy as a strategy for elimination of HIV transmission: a mathematical model. *The Lancet*, **373**, 48–57.

Grimm, V., Berger, U., Bastiansen, F., Eliassen, S. and Ginot, V. (2006) A standard protocol for describing individual-based and agent-based models. *Ecological Modelling*, **198**, 115–126.

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.

Kennedy, M. C. and O'Hagan, A. (2001) Bayesian calibration of computer models. *Journal of the Royal Statistical Society. Series B*, **63**, 425–464.

Loeppky, J. L., Sacks, J. and Welch, W. J. (2009) Choosing the sample size of a computer experiment: a practical guide. *Technometrics*, **51**, 366–376.

May, R. M. (2004) Uses and abuses of mathematics in biology. *Science*, **303**, 790–793.

McCreesh, N., O'Brien, K., Nsubuga, R. N., Shafer, L. A., Bakker, R., Seeley, J., Hayes, R. J. and White, R. G. (2012) Exploring the potential impact of a reduction in partnership concurrency on HIV incidence in rural Uganda: a modeling study. *Sexually Transmitted diseases*, **39**, 407–413.

McKinley, T. J., Cook, A. R. and Deardon, R. (2009) Inference in epidemic models without likelihoods. *The International Journal of Biostatistics*, **5**.

Mulder, D. W., Nunn, A. J., Kamali, A., Nakiyingi, J., Wagner, H. U. and Kengeya-Kayondo, J. F. (1994a) Two-year HIV-1-associated mortality in a Ugandan rural population. *Lancet*, **343**, 1021–1023.

Mulder, D. W., Nunn, A. J., Kamali, A., Wagner, H. U. and Kengeya-Kayondo, J. F. (1994b) HIV-1 incidence and HIV-1-associated mortality in a rural Ugandan population cohort. *AIDS*, **8**, 87–92.

O'Neill, P. D. and Roberts, G. O. (1999) Bayesian inference for partially observed stochastic epidemics. *Journal of the Royal Statistical Society. Series A (General)*, **162**, 121–129.

Picheny, V., Ginsbourger, D., Richet, Y. and Caplin, G. (2013) Quantile-based optimization of noisy computer experiments with tunable precision. *Technometrics*, **55(1)**, 2–13.

Powers, K. A., Ghani, A., Miller, W., Hoffman, I., Pettifor, A. and Kamanga, G. (2013) The role of acute and early HIV infection in the spread of HIV and implications for transmission prevention strategies in Lilongwe, Malawi: a modelling study. *The Lancet*, **378**, 256–268.

Pukelsheim, F. (1994) The three sigma rule. *The American Statistician*, **48**, 88–91.

Punyacharoensin, N., Edmunds, W., De Angelis, D., Delpech, V., Hart, G. and Elford, J. (2016) Effect of pre-exposure prophylaxis and combination hiv prevention for men who have sex with men in the uk: a mathematical modelling study. *The Lancet HIV*, **3**, 94–104.

Rasmussen, C. E. and Williams, C. K. I. (2006) *Gaussian Processes for Machine Learning*. The MIT press.

Rougier, J. (2008) Efficient emulators for multivariate deterministic functions. *Journal of Computational and Graphical Statistics*, **17**, 827–843.

Sacks, J., Welch, W. J., Mitchell, T. J. and Wynn, H. P. (1989) Design and analysis of computer experiments. *Statistical Science*, **4**, 409–435.

Seeley, J., Wagner, U., Mulemwa, J., Kengeya-Kayondo, J. and Mulder, D. (1991) The development of a community-based HIV/AIDS counselling service in a rural area in Uganda. *AIDS Care*, **3**, 207–217.

Toni, T., Welch, D., Strelkowa, N., Ipsen, A. and Strumpf, M. P. H. (2009) Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society Interface*, **6**, 187–202.

Vernon, I. and Goldstein, M. (2010) A Bayes linear approach to systems biology. *Tech. rep.*, MUCM Technical Report.

Vernon, I., Goldstein, M. and Bower, R. G. (2010a) Galaxy formation: a Bayesian uncertainty analysis. *Bayesian Analysis*, **5**, 619–670.

— (2010b) Rejoinder for galaxy formation: a Bayesian uncertainty analysis. *Bayesian analysis*, **5**, 697–708.

— (2014) Galaxy formation: Bayesian history matching for the observable universe. *Statistical science*, **29**, 80–90.

Williamson, D., Goldstein, M., Allison, L., Blaker, A., Challenor, P., Jackson, L. and Yamazaki, K. (2013) History matching for exploring and reducing climate model parameter space using observations and a large perturbed physics ensemble. *Climate Dynamics*, **41**, 1703–1729.

**Table 1.** Simulator input parameter description and ranges. These define the input parameter space over which the history match search is performed.

| Number | Input description | Abbr. | Min. | Max. |
|---|---|---|---|---|
| 1 | Proportion of men in the high sexual activity group | *mhag* | 0.01 | 0.5 |
| 2 | Proportion of women in the high sexual activity group | *whag* | 0.01 | 0.5 |
| 3 | Mixing by activity group [$\epsilon$] | *mag* | 0 | 1 |
| 4 | High activity contact rate (risk behaviour 1) [partners/yr]* | *hacr1* | 0 | 10 |
| 5 | Low activity contact rate (risk behaviour 1) [partners/yr]* | *lacr1* | 0 | 2 |
| 6 | Start year for risk behaviour 2 | *sy2* | 1986 | 1992 |
| 7 | High activity contact rate (risk behaviour 2) [partners/yr]* | *hacr2* | 0 | 10 |
| 8 | Low activity contact rate (risk behaviour 2) [partners/yr]* | *lacr2* | 0 | 2 |
| 9 | Start year for risk behaviour 3 | *sy3* | 1998 | 2002 |
| 10 | High activity contact rate (risk behaviour 3) [partners/yr]* | *hacr3* | 0 | 10 |
| 11 | Low activity contact rate (risk behaviour 3) [partners/yr]* | *lacr3* | 0 | 2 |
| 12 | Mean HIV transmission probability per sex act during primary stage of infection (mean of male to female and female to male transmission probabilities) | *atp* | 0 | 1 |
| 13 | Ratio of male to female/female to male transmission probabilities | *rtp* | 1 | 3 |
| 14 | Proportion of low activity men in high concurrency group | *lmhc* | 0 | 1 |
| 15 | Proportion of low activity women in high concurrency group | *lwhc* | 0 | 1 |
| 16 | Male concurrency parameter in high concurrency group (risk behaviour 1) | *mchc1* | 0 | 1 |
| 17 | Female concurrency parameter in high concurrency group (risk behaviour 1) | *fchc1* | 0 | 1 |
| 18 | Male concurrency parameter in high concurrency group (risk behaviour 2) | *mchc2* | 0 | 1 |
| 19 | Female concurrency parameter in high concurrency group (risk behaviour 2) | *fchc2* | 0 | 1 |
| 20 | Male concurrency parameter in high concurrency group (risk behaviour 3) | *mchc3* | 0 | 1 |
| 21 | Female concurrency parameter in high concurrency group (risk behaviour 3) | *fchc3* | 0 | 1 |
| 22 | Duration of long-duration partnerships [years] | *dlp* | 5 | 20 |

(*) simulator input parameters that codetermine partnership formation. The actual rate of partnership formation in the simulator will vary from this due to adjustment for concurrency and partnership balancing.

**Table 2.** Description of simulator outputs and the limits defined as an acceptable match.

| Number | Output description | Abbr. | Min. | Max. |
|---|---|---|---|---|
| 1 | Population size in 2008 (male) | *psm* | 2986 | 3650 |
| 2 | Population size in 2008 (female) | *psf* | 3374 | 4124 |
| 3 | Average male partnership incidence in 2008 (partners/year) | *ampi* | 0.4 | 0.489 |
| 4 | HIV prevalence in 1992 (male) | *p92m* | 0.084 | 0.112 |
| 5 | HIV prevalence in 1992 (female) | *p92f* | 0.096 | 0.124 |
| 6 | HIV prevalence in 2001 (male) | *p01m* | 0.07 | 0.09 |
| 7 | HIV prevalence in 2001 (female) | *p01f* | 0.083 | 0.107 |
| 8 | HIV prevalence in 2007 (male) | *p07m* | 0.06 | 0.084 |
| 9 | HIV prevalence in 2007 (female) | *p07f* | 0.093 | 0.119 |
| 10 | Point prevalence of men with 1 long duration partnership in 2008 (%) | *m1l* | 34.62 | 42.31 |
| 11 | Point prevalence of men with 1 short duration partnership in 2008 (%) | *m1s* | 10.86 | 13.27 |
| 12 | Point prevalence of men with 1 partnership (either type) in 2008 (%) | *m1* | 37.83 | 46.24 |
| 13 | Point prevalence of men with 2+ long duration partnerships in 2008 (%) | *m2l* | 3.38 | 4.13 |
| 14 | Point prevalence of men with 2+ short duration partnerships in 2008 (%) | *m2s* | 1.69 | 2.07 |
| 15 | Point prevalence of men with 2+ partnerships (any combination) in 2008 (%) | *m2* | 8.66 | 10.59 |
| 16 | Point prevalence of women with 2+ long duration partnerships in 2008 (%) | *w2l* | 0.85 | 1.03 |
| 17 | Point prevalence of women with 2+ short duration partnerships in 2008 (%) | *w2s* | 0.42 | 0.52 |
| 18 | Point prevalence of women with 2+ partnerships (any combination) in 2008 (%) | *w2* | 2.17 | 2.65 |

**Table 3.** Number of design points the simulator was run at each wave. The number of simulator evaluations is given by the numbers shown in the table multiplied by $K = 100$.

| Wave | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | Total |
|---|---|---|---|---|---|---|---|---|---|---|
| Fixed variance | 240 | 242 | 249 | 250 | 516 | 520 | 500 | 500 | 500 | 3517 |
| Emulated variance | 240 | 242 | 249 | 250 | 516 | 520 | - | - | - | 2017 |

**Table 4.** Total time and breakdown of the tasks involved in history matching Mukwano using the two approaches. The tasks in the first 4 rows are parallelisable, and figures assume the usage of a 100 core cluster (2.5GHz, 8GB RAM).

| | Fixed variance | Emulated variance | % Reduction |
|---|---|---|---|
| Simulator running time [days] | 10.5 | 7.1 | 32 |
| Emulator training time [days] | 1.19 | 0.41 | - |
| Var. emulator training [days] | - | 0.38 | - |
| Total emulator training [days] | 1.19 | 0.79 | 34 |
| Staff time per wave [hours] | 4 | 6 | - |
| Total staff time [days] | 1.5 | 1.5 | 0 |
| Total time [days] | 13.2 | 9.4 | 29 |

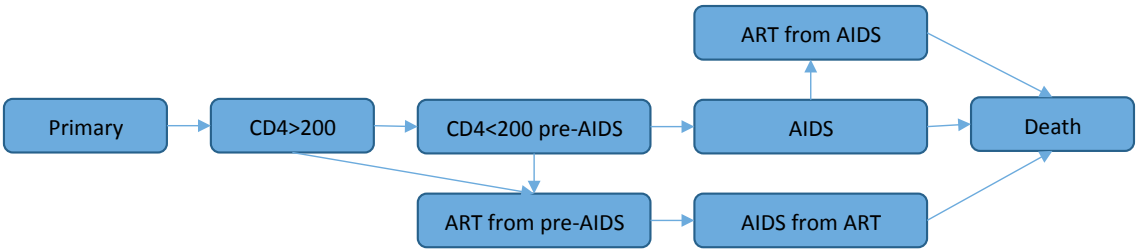


**Fig. 1.** **Schematic of simulated HIV natural history and antiretroviral treatment.**

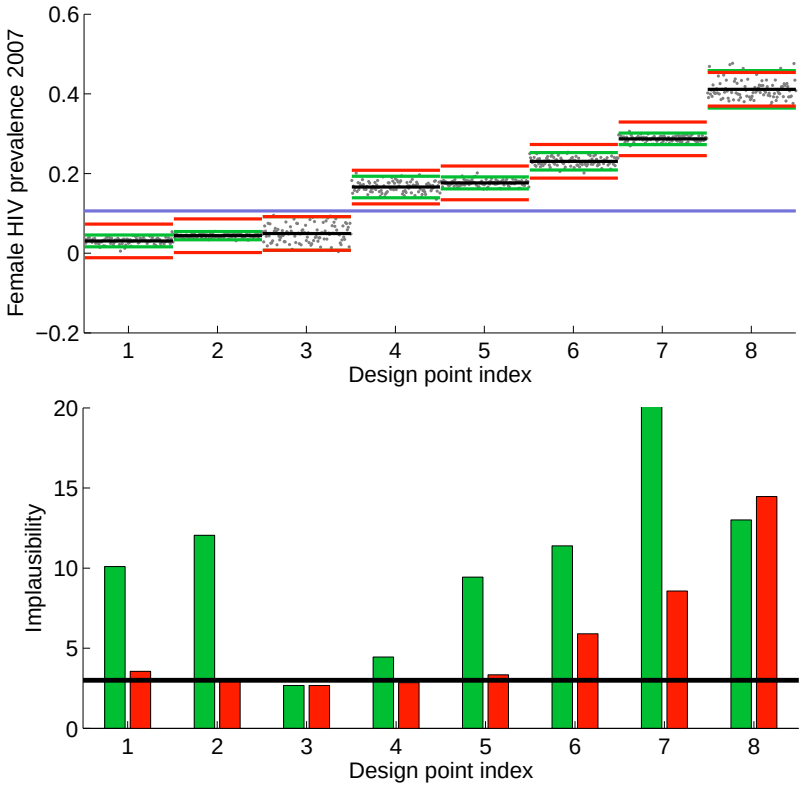


**Fig. 2.** **Improvements due to the emulated variance method.** Top panel: 100 simulator evaluations in 8 different design points (grey dots), their mean (horizontal black lines) and $\pm 2$ standard deviations calculated with $V_{s,90}$ or with the actual variance calculated by the 100 repetitions at each design point (green lines). The horizontal blue line represents the mean of the empirical data. Bottom panel: a simplified form of the implausibility calculated with the respective variances of the top panel.

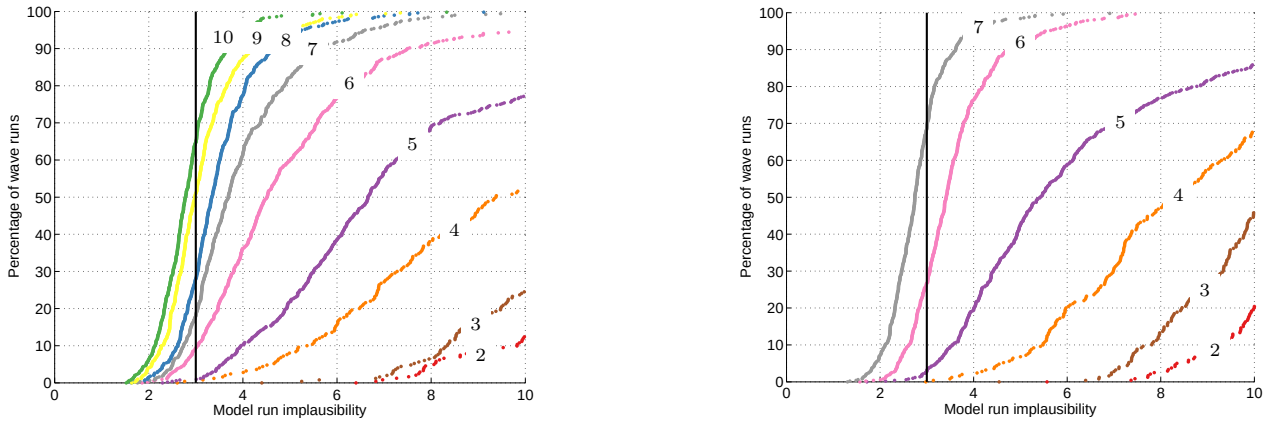


(a) Fixed variance (b) Emulated variance

**Fig. 3. Cumulative distribution function of simulator run implausibility $I_{\mathcal{R}}(\mathbf{x})$, by wave.** Each line represents the percentage of each wave's simulator runs with an $I_{\mathcal{R}}(\mathbf{x})$ less than the value indicated by the horizontal axis. The numbers on the curves indicate the wave number.

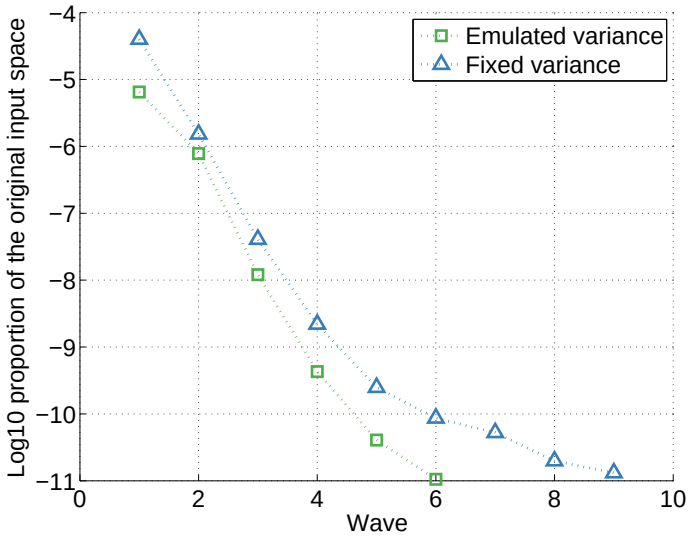


**Fig. 4. Proportion of the samples drawn at random in the original simulator input space, that are non-implausible after $k$ waves of history matching.** The proposed emulated variance methodology achieved the same reduction of non-implausible space in 3 fewer waves.

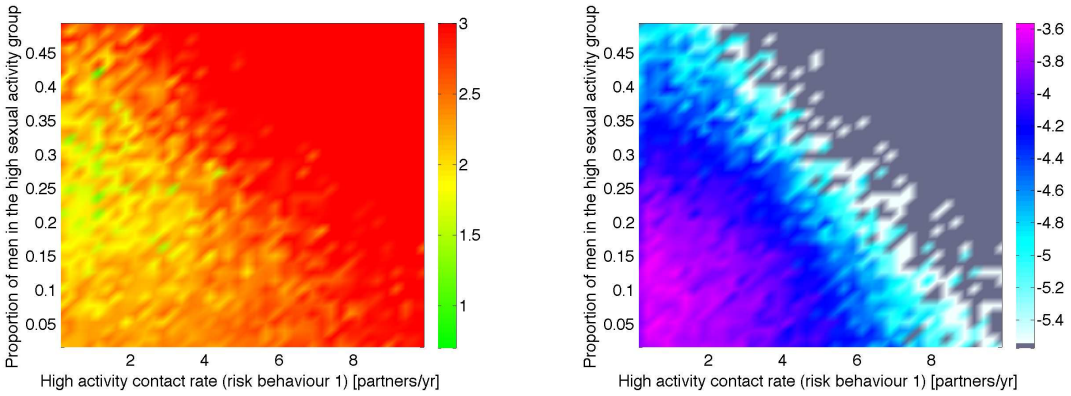


(a) Minimum implausibility (b) Optical depth

**Fig. 5. Examples of minimum implausibility and optical depth plots.** Minimum implausibility plots show an estimate of the minimum implausibility achievable by varying the remaining inputs for different values of the inputs shown along the x and y axes. Optical depth plots provide an estimate of the $\log_{10}$ probability of finding a non-implausible point once the two selected inputs are fixed to a certain value, and hence give the depth of the non-implausible region at each point.

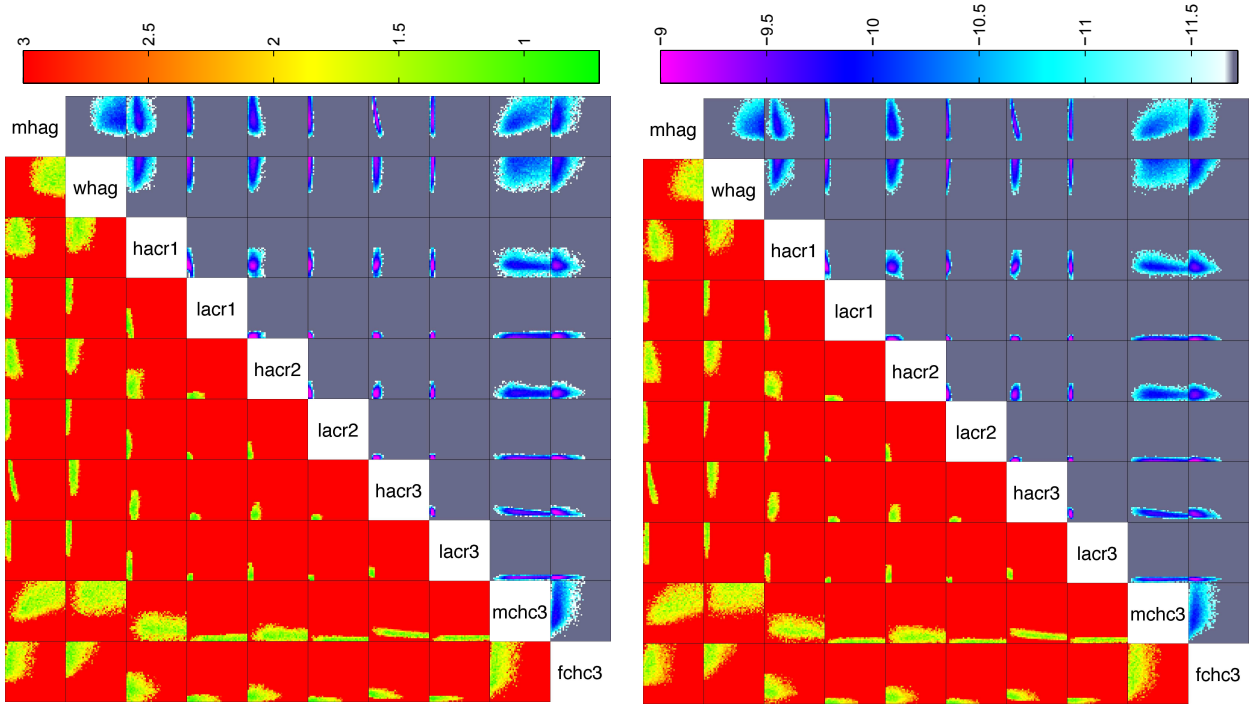


(a) Wave 9 (fixed variance) (b) Wave 6 (emulated variance)

**Fig. 6. A comparison of the minimum implausibility (below and left of diagonal) and optical depth plots (above and right of diagonal) for 10 key inputs.** All axes vary between 0 and 1 (normalised). For the minimum implausibility plots, the inputs that appear across the main diagonal vary along the horizontal axis for the plots that appear to the left of the input names, and along the vertical axis for those that appear below. For the optical depth plots, the inputs vary across the horizontal axis for the plots that appear above the input names and across the vertical axis for those that appear to the right.

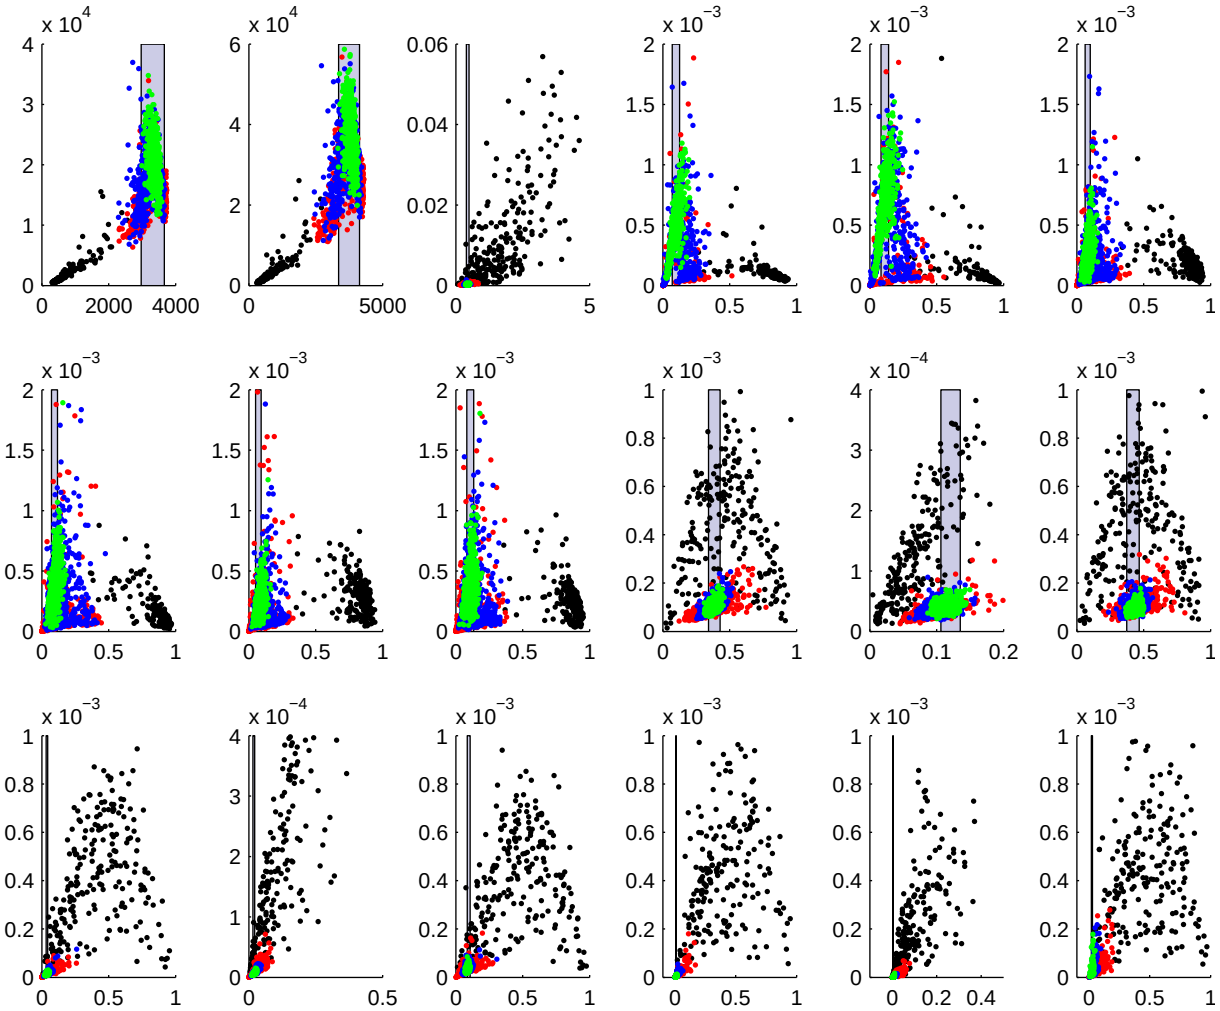

**Fig. 7.** **Scatter plots of simulator's mean output (horizontal axis) against its variance (vertical axis) for waves 1 (black), 3 (red) 5 (blue) and 7 (green).** The 18 outputs are arranged first from left to right and then from top to bottom. The vertical patches in each panel show the empirical data on the mean outputs with $\pm$ 2 standard deviations derived from the observation error and model discrepancy at wave 6.

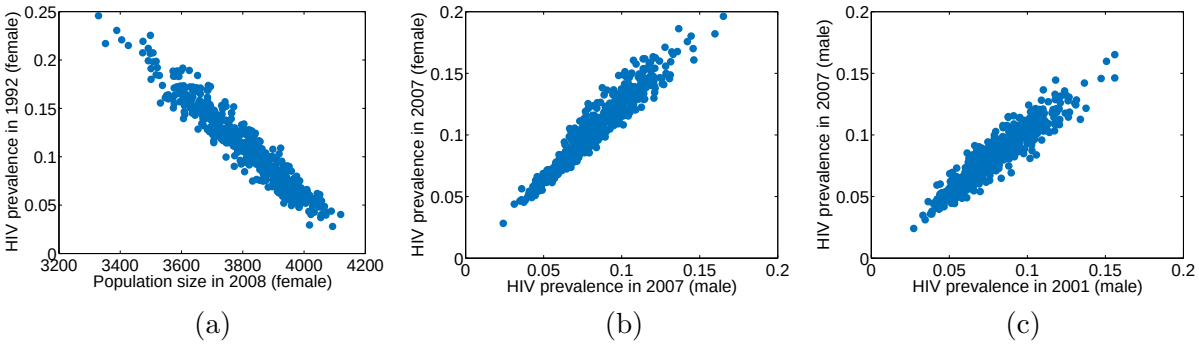


**Fig. 8.** **Scatter plots of mean output values from simulator runs at wave 7.** The large correlation between outputs is an indication of the way the simulator models the HIV transmission across sexes and across time (see text).