Applying particle filtering in complex compartmental models of pre-vaccination pertussis

Xiaoyan Li[1*], Nathaniel D. Osgood[1]

1 Department of Computer Science, University of Saskatchewan, Saskatoon, Saskatchewan, Canada

* xiaoyan.li@usask.ca

Abstract

Particle filtering is a modern state inference and identification methodology that allows filtering of general non-Gaussian and non-linear models in light of time series of empirical observations. Several previous lines of research have demonstrated the capacity to effectively apply particle filtering to low-dimensional compartmental transmission models. We demonstrate here implementation and evaluation of particle filtering to more complex compartmental transmission pertussis models – including application for aggregate, two-age-groups and 32-age-groups population structure with two different contact matrix, respectively – using over 35 years of monthly and annual pre-vaccination provincial data from the mid-western Canadian province. Such filtering supports estimation of (via sampling from) the entire state of the dynamic model, both latent and observable, for each point in time, thereby supporting estimation of proportion of susceptible individuals, and those at varying levels of immune protection associated within waning of pertussis immunity. Estimation of the latent state also supports capacity for model projection. Following evaluation of the predictive accuracy of these four particle filtering models, we then performed prediction and intervention experiments based on the most accurate model. Using that model, we contribute the first full-paper description of particle filter-informed intervention evaluation. We conclude that applying particle filtering with relatively high-dimensional pertussis transmission models, and incorporating time series of reported counts, is a valuable technique to assist public health authorities in estimating and predicting pertussis outbreaks in the context of incoming empirical data. Within this use, the particle filtering models can moreover perform counterfactual analysis of interventions to assist the public health authorities in intervention planning.

Author summary

Pertussis is a common and highly contagious childhood disease. In contrast to some other prevalent childhood diseases, immunity conferred by natural exposure or vaccination to pertussis is widely believed to wane relatively rapidly. Such waning and associated uncertainties raises challenges for accurate pertussis modelling and simulation, with the speed and current degree of such waning in the population remaining a notable point of uncertainty. In recent years, the sequential monte carlo method of particle filtering has been employed for incorporating empirical time series data so as to estimate underlying model state of simpler communicable disease transmission models, particularly for influenza and measles. This paper characterizes

the first time in which such methods have been applied to the more complex and higher dimensional models of pertussis, using over 35 years of monthly and yearly empirical time series from the pre-vaccination era in a mid-western Canadian province. In addition, we contribute the first full-paper description of particle filter-informed intervention evaluation for transmission models. Our findings suggest that applying particle filtering with relatively high-dimensional pertussis transmission models can serve as a valuable technique to assist public health authorities in estimating and predicting pertussis outbreaks in the context of incoming empirical data. Application of particle filtering can moreover support evaluation of intervention outcomes to assist public health authorities in intervention planning.

1 Introduction

Pertussis is a common childhood disease, which is a highly contagious disease of the respiratory tract that caused by the bacterium *Bordetella pertussis* [1]. It is most dangerous for infants, due to risks of severe complications, post-paroxysm apnia [1]. The most frequent complication is pneumonia, while seizures and encephalopathy occur more rarely [1]. Pertussis is a highly contagious disease only found in humans, and spreads from person to person by coughing, sneezing, and prolonged proximity [2]. Evidence indicates a secondary attack rate of 80% among susceptible household contacts [3]. In contrast to some other prevalent childhood diseases, immunity conferred by natural exposure or vaccination to pertussis is widely believed to wane relatively rapidly, leading to significant risks of infection even in adults who have been previously infected. It is notable that babies can be infected by adults, such as parents, older siblings, and caregivers who might not even know they have already contracted this disease [2]. Pertussis incidence shows no distinct seasonal pattern. However, it may increase in the summer and fall [3].

In the pre-vaccination era, pertussis was one of the most common childhood infectious diseases and a major cause of childhood mortality. In 1860, the mortality rate of all-age pertussis in Demark was 0.015% [4], but that burden fell heavily on infants and children. Research into historical mortality rates from pertussis indicate that the death rate in infancy is higher than in other groups [4]. In recent years globally, there were an estimated 24.1 million cases of pertussis, and about 160,700 deaths per year [5]. Since the 1980s, there has been a rising trend in the reported cases of pertussis in the United States [5]. The most recent peak year of the reported cases of pertussis in the United States is 2012, when the Centers for Disease Control and Prevention (CDC) reported 48,277 cases, but many more are believed to go undiagnosed and unreported [5]. Research aimed at estimating the level of population susceptibility and predicting the transmission dynamics of pertussis could aid outbreak prevention and control efforts by health agencies, such as performing intervention before the predicted next outbreak, and in targeted outbreak response immunization campaigns [6].

Dynamic modelling has long served as an important tool for understanding the spread of the infectious diseases in population [16], including pertussis, and for evaluating the impacts of interventions such as immunization and quarantine. In recent years, particle filtering as a machine learning method has been employed for incorporating empirical time series data (such as surveillance [7] and online communicational behavior data [8]) to ground the hypothesis as to the underlying model state models in some previous researches [9–12], especially for the infectious diseases of influenza [13, 14] and measles [7]. In this paper, we apply the particle filtering algorithm in a more complex and widely used compartmental model [15] of pertussis by incorporating the reported pertussis cases in Saskatchewan during the pre-vaccination era. Particle filtering for pertussis is different than for other pathogens

on account of the need for state estimation to estimate the population segments at varying levels of immunity. Another need concerns extends from the heterogenous nature of the mixing and incidence burden between different age groups. For this reason, age-structured models are examined here. Specifically, we have examined two categories of age-structured particle filtering models – with 2 age groups and with 32 age groups. Moreover, we have proposed and explored three methods for calculating the contact matrix, so as to reduce the degrees of freedom associated with characterization of the contact matrix. This contribution compares the results obtained from all the particle filtering models by incorporating the empirical data across the whole timeframe evaluating the predictive accuracy of the models. Finally, using the minimum discrepancy particle filtering model, we demonstrate how we can evaluate intervention effects in a fashion that leverages the capacity of particle filtering to perform state estimation.

2 Methods

2.1 The mathematical models

As noted above, the infectious dynamics of pertussis is more complex than the infectious diseases that confer lifelong immunity, such as measles, due temporary character of the immunity acquired by *Bordetella pertussis* infection. As the time since the most recent pertussis infection increases, the immunity of a person wanes [16]. People with lower immunity generally tend to be more easily infected, and exhibit a higher risk of transmitting the infection once infected.

In this paper, we have employed the structure of the popular pertussis mathematical model in [15]. To capture the characteristics of pertussis in waning of immunity and the different level of infectiousness and susceptibility involved with infection in light of pre-existing immunity, the compartmental model in [15] further divides the infectious population into three groups: weakly infectious (I_w), mild infectious (I_m), and fully infectious (I), and divided the recovered population into four groups: R_1, R_2, R_3 and R_4 of successively increasing immune system strength. Moreover, in reflection of the focus on the pre-vaccination era (pertussis reported cases in Saskatchewan from 1921 to 1956), the stocks related to vaccination (V_1, V_2, V_3 and V_4) in the original compartmental model [15] are not included in this research.

Fig 1. The transfer diagram for the pertussis model without vaccination. adapted from [15]

Fig 1 shows the mathematical structure of the compartmental pertussis model adapted from [15] . In this compartmental model of pertussis in the pre-vaccination era, the total population is divided into 8 distinct epidemiological classes. Newborns enter directly into class S of fully susceptible individuals. If a fully susceptible individual contacts an infective individual and is successfully transmitted pertussis, this previous susceptible person becomes infectious and enters the class (state) I of full infectives. Infective individuals in state I of have full cases of pertussis, with all of the usual symptoms. When individuals recover from the state I of infectives, they achieve full immunity and enter state R_4. In this state, they are fully protected and can not be infected by pertussis. However, as time goes by, their immunity wanes and they enter into a less strong immunity class of R_3. When individuals in class R_3 are exposed to an infective, they are assumed to return to the highest immunity class of R_4 without becoming infectious. Otherwise, their immunity keeps fading, and they enter to the relatively lower immunity class of R_2. When a person in the class of R_2 is sufficiently

(re-)exposed to an infective for transmission to occur, the infected individual enters the I_w state with weak infectivity. Individuals in the I_w class have the weakest infective capability to infect a susceptible. After they recover, the individuals in class of I_w then secure the highest immunity, (re-) entering the class of R_4 from which they originally waned. By contrast, if people in the class of R_2 are not re-exposed to the infectives, their immunity continues waning, and they enter the minimally immune class of R_1. Similarly, if a person in class R_1 is re-exposed to an infective, this person gets infected with mild infectivity and enters the class of I_m. Individuals in the class of I_m have a higher infectious capability compared with those in the class of the weak infective (I_w), but exhibit a lower infectious capability compared to the fully infective individuals in I. When recovered, the individuals in class I_m enter the class R_4 again. If the individuals in the class of R_1 are not re-exposed, they eventually lose all of their immunity and move back to the class of S whence they originated at birth. Given the presence of multiple infection states (I, I_w, I_m) as well as multiple levels of immunity (R_1, R_2, R_3, R_4), three invariants bears noting. Firstly, regardless of the pre-existing level of immunity, following recovery from an infection, that individual always returns to the full level of natural immunity (R_4). Secondly, in any of the recovered states (R_1, R_2, R_3, R_4), immunity continues to wane absent re-exposure. Thirdly, as the level of immunity is reduced, the severity of resulting infectiousness rises, with no infectiousness being possible at all from exposure in states R_3 and R_4.

It is notable that the model of Hethcote (1997) [15] makes use of a formulation in which each state variable is of unit dimension, representing a fraction of the population in different age groups of the same class. However, for the sake of comparison against empirical data, the model in this paper is represented in a re-dimensionalized fashion, with the state variables representing counts of persons. Based on this structure in Fig 1, four models are considered in this research. Using n to denote the total number of age groups in the models, we consider models of the aggregate population ($n = 1$), of two age groups ($n = 2$), 32 age groups model ($n = 32$) with the contact matrix introduced in the paper of Hethcote (1997) [15] and a further 32 age groups ($n = 32$) model with a re-balanced contact matrix. In the two age group model, the population is divided into child and adult categories, where the child age group includes individuals in their first 5 years of life, while the adult age group includes individuals aged 5 and above. In the 32 age group models, we follow Hethcote in assuming age groups as 0–1 month, 2–3 month, 4–5 month, 6–11 month, 1 year, 2 year, 3 year, 4 year, 5 year, 6 year, 7 year, 8 year, 9 year, 10 year, 11 year, 12 year, 13 year, 14 year, 15 year, 16 year, 17 year, 18 year, 19 year, 20–24 year, 25–29 year, 30–39 year, 40–49 year, 50–59 year, 60–69 year, 70–79 year, 80–89 year, 90 years and above [15]. It bears noting that while more detailed age structure can better capture both the effects of population aging and inter-group heterogeneity, in terms of particle filtering, it will require estimation of a larger underlying model state – potentially adversely affecting the accuracy of that estimation.

In addition to examining the effects of assuming different age structures, we have explored 3 different methods in formalizing the contact matrix that characterizes contacts between different age groups [7, 15] in transmission models. In the two age group model ($n = 2$), the contact matrix is consistent with the contact matrix introduced in a previously contributed application of particle filtering to measles [7]. This contact matrix used in [7] is powerful and effective in mathematical models with a small number of age groups. However, this method encounters problems with extension to larger number of age groups, given that the number of free parameters rises as the square of the total number of age groups in the model (readers interested in the mathematical proof are referred to S4 Appendix). Thus, in the 32-age-groups models, we have explored two alternative contact matrices exhibiting a reduced number of free parameters. In the first 32-age-groups model, we employed an identical contact matrix

to that employed in [15]. Specifically, to generate the contact matrix in [15], this research has assumed that the average number of people contacted a person in age group i per unit time are distributed among people in age group j in proportion to the fraction of all contacts per unit time received by people in the age group j [15]. However, this contact matrix used in [15] is un-balanced, which means the total contacts of age group i to age group j is not equal to the total contacts of age group j to age group i (see S1 Appendix). Thus, we have explored another contact matrix in the 32-age-groups model where the contact matrix is balanced, according to the methodology established in [17]. As a result of the balancing, the total contacts in age group i to age group j is taken as equal to the total contacts of age group j to age group i. The number of free parameters in the contact matrix correspondingly grows linearly with the total number of the age groups in the model. Readers interested in details of the mathematical models and the mathematical characterization of the alternative contact matrices are referred to S1 Appendix.

2.2 Particle filter implementation

Particle filtering is a modern state inference and identification methodology that allows filtering of general non-Gaussian and non-linear models in light of time series of empirical observations [7, 9, 10, 14, 18, 19]. At a given time, each state in the distribution (represented by "particles" according to the principles of sequential importance sampling) can be seen as representing a competing hypothesis concerning the underlying state of the system at that time. The particle filtering method can be viewed as undertaking a "survival of the fittest" of these hypotheses, with fitness of a given particle being determined by the consistency between the expectations of the hypothesis associated with that particle and the empirical observations.

Particle filtering offers value in the context of underlying state equation models exhibiting stochastic variability. In the implementation of the particle filtering algorithm, we extended the pertussis compartment model by adding "system noise" by letting some parameters evolve with time. Readers interested in additional details with the state space models implementation of the pertussis particle filtering models are referred to S2 Appendix. In performing the particle filtering, we followed several past contributions [7, 9, 11, 13] by employing the "Condensation Algorithm" [19, 20] as the weight update rule of the particles at each time after considering the arriving observation. By making the simplifying choice of the prior as the proposal distribution [18, 19], the Condensation Algorithm is straightforward to implement, and represents the simplest and widely used variant of particle filtering algorithm. We further followed [7, 9, 11, 13, 21] in assuming that the likelihood function in the particle filtering algorithm is based around the negative binomial distribution. Specifically, for the particle filtering model with an aggregate population, the likelihood function simply applies the negative binomial distribution to link the result of the model and observation data via monthly reported pertussis cases. By contrast, for the particle filtering models with an age-structured population, the likelihood function is given by the product of negative binomial density functions applied successively to the observation and corresponding model quantity for each empirical dataset [7]. It bears noting that the value of the dispersion parameter r [7, 9] associated with the negative distribution in all scenarios in this paper is chosen to be 10.

In the aggregate particle filtering model of pertussis, only one empirical dataset – the reported monthly population-wide incident pertussis cases in the province of Saskatchewan from the year 1921 to 1956 – is employed [22]. To assess model accuracy, the RMSE – Root Mean Square Error – between the monthly empirical data and corresponding item in the particle filtering model is calculated as the discrepancy among particles sampled by weight, according to importance sampling principles. In the

age-structured particle filtering model of pertussis using 2 age groups, three empirical datasets are used – monthly reported population-wide incident pertussis cases, the yearly pertussis reported cases among children less than five years old, and the yearly reported pertussis cases in those five or older. The weight update rule (likelihood function) is the same as for the aggregate model, except for at the close of each year, where the value of the likelihood function is the product of the negative binomial function applied to each of the three types of obesrvations – monthly population-wide pertussis reported cases, yearly pertussis reported cases for (separately) child and adult age groups. It bears remarking that this structure mirrors that used in a previous contribution using particle filtering for a measles model [7]. For assessing accuracy of the two age-grouop model, the discrepancy includes the monthly RMSE and the sum of yearly RMSE of each age group divided by 12 (so as to convert the unit to from year to month).

In the pertussis age-structured particle filtering models involving 32 age groups, 7 empirical datasets are employed – monthly population-wide reported pertussis cases, and yearly pertussis reported cases for each of six age groups, including less than 1 year, from 1 to 4 years, from 5 to 9 years, from 10 to 14 years, from 15 to 19 years and 20 years and older. The update weight rule is similar to that for the age structured pertussis model of two age groups, except that at the close of each year it is the product of likelihood functions applied for each of the 7 empirical datasets. When assessing the predictive accuracy of the model, the discrepancy also includes the monthly RMSE and the sum of yearly RMSE of each age group divided by 12 (so as to convert the unit from year to month).

Particle Filtering Algorithm

The particle filter algorithm that we employed in this paper is given as follows [7, 9, 18, 23]:

1. At time k=0:

(1) Sample $X_0^{N(i)}$ from $q_0(x_0^N)$;

(2) Compute a weight for each particle $w_0^{(i)} = \frac{1}{N_s}$. It indicates that the weight at initial time follows the uniform distribution.

2. At time $k \geq 1$, perform a recursive update as follows:

(1) Advance the sampled state by sampling $X_k^{N(i)} \sim q_k(x_k^N | y_k, X_{0:k-1}^{N(i)})$ and set $X_{0:k}^{N(i)} = (X_{0:k-1}^{N(i)}, X_k^{N(i)})$;

(2) Update the weights (by the Condensation Algorithm) to reflect the probabilistic and state update models as follows:

$$w_k^{(i)} = W_{k-1}^{(i)} p(y_k^M | X_k^{N(i)}).$$

Normalize the weights $W_k^{(i)} = \frac{w_k^{(i)}}{\sum_{i=1}^{N_s} w_k^{(i)}}$

(3) Calculate the S_{eff} : $\frac{1}{\sum_{i=1}^{N_s} (w_k^{(i)})^2}$

(4) If $S_{eff} < S_T$ (S_T is the minimum effective sample size – the threshold of resampling), perform resampling to get a new set of $X_k^{N(i)}$. And set the weight of the new particles as $\frac{1}{N_s}$.

3 Results

3.1 Results of models incorporating empirical datasets across all timeframe

Recall that to explore the predictive performance of particle filtering in different compartmental pertussis models, four particle filtering models have been built in this research – the aggregate particle filtering pertussis model (denoted as $PF_{aggregate}$), the age-structured particle filtering model of 2 age groups (denoted as PF_{age_2}), the age-structured particle filtering model of 32 age groups with the original Hethcote contact matrix (denoted as $PF_{age_32_Hethcote}$), and the age-structured particle filtering model of 32 age groups with the re-balanced contact matrix (denoted as $PF_{age_32_rebalanced}$). In all of the four particle filtering models, the total number of particles in the particle filtering algorithm is 3000; for clarity in exposition, we sampled the same number when generating the plots of the 2D histogram and for calculating the discrepancy. To compare the accuracy of a particle filtered model against that of a traditional model of pertussis calibrated against comparable data, we have further built a calibrated model of the aggregate population, henceforth denoted *Calibrated*. In this calibrated model, the values of the parameters obtained from calibration against the empirical dataset are listed as follows. The initial value estimated from the calibration process in class S, I and R1 are 19420, 500, 9960. The value of the infectious contact rate (β) is 56.692; it bears emphasis that this value incorporates both a rate of contact and the probability of transmission. The calibrated value of the reporting rate of pertussis is 0.01. The other parameters are the same as the particle filtering models.

Table 1. Comparison of the average discrepancy of the calibrated model and all four particle filtered pertussis models considering empirical data across all observation points; parentheses give the 95% confidence intervals.

Model	Monthly	Yearly in Month	Total
Calibrated	34.2	NONE	NONE
$PF_{aggregate}$	20.9 (20.0, 21.9)	NONE	NONE
PF_{age_2}	19.9 (18.8, 21.0)	21.0 (19.2, 22.7)	40.9 (38.1, 43.6)
$PF_{age_32_Hethcote}$	20.6 (20.1, 21.2)	25.8 (23.0, 28.6)	46.4 (43.1, 49.7)
$PF_{age_32_rebalanced}$	19.8 (19.5, 20.1)	28.1 (24.2, 31.9)	47.9 (43.9, 51.9)

Each of the five particle filtering models was run 5 times (the random seed generated from the same set). Shown here are the average and 95% confidence intervals (in parentheses) of the mean discrepancy for each model variant.

By comparing the discrepancy of these models, we sought to identify the model offering the greatest predictive validity. We then use the most favorable model to perform prediction and intervention analysis. To assess model results, each of the four particle filtering models was run 5 times with random seeds generated from the same set. We then calculate the average and 95% confidence intervals of the mean discrepancy. Table 1 displays the average discrepancies of the four pertussis particle filtering models and the calibrated deterministic pertussis compartmental model, where the discrepancy considers the entire timeframes. The results of table 1 suggests that particle filtering models significantly improves the predictive accuracy beyond what is achieved via calibration. It is notable that both the calibrated deterministic model and the aggregate particle filtering model only offer monthly average discrepancy, because the yearly observations are stratified by age, but no age stratification is implemented in these two models. Table 1 indicates that the particle filtering models are significantly more

Fig 2. Boxplot of monthly and yearly discrepancy of all models at monthly observation points, considering empirical data across all observation points. "Calibrate" indicates the calibration model with aggregate population structure; "PF_a1" indicates the particle filtering model with aggregate population structure; "PF_a2" indicates the particle filtering model with 2 age groups; "PF_a32H" indicates the particle filtering model with 32 age groups and the contact matrix introduced in [15]; "PF_a32R" indicates the particle filtering model with 32 age groups and the re-balanced contact matrix. "_M" indicates the discrepancy of the models compares with the monthly empirical data – the pertussis reported cases among all population; "_Y" indicates the sum of discrepancy (of each age group) of the models compares with the yearly empirical data – the pertussis reported cases divided in age groups – and adjust the unit to Month by dividing 12. It is also notable that the dot in the boxplot indicates the mean value, while the horizontal line indicates the median value.

accurate compared to the calibration model – the average discrepancies of the particle filtering models are significantly lower than the average discrepancy of the calibrated deterministic model. Moreover, although the monthly average discrepancies among the four particle filtering models with different population structure and contact matrix structure are quite close, the particle filtering models PF_{age_2} and $PF_{age_32_rebalanced}$ exhibit smaller average discrepancies. With respect to the yearly average discrepancies, table 1 shows that the age-structured model with two age groups offers better predictive performance than the model with 32 age groups; as noted, the aggregate model lacks the age stratification required to calculate yearly discrepancy. It is notable that the total number of the yearly empirical datasets against which the calibration is assessed is different between the age-structured models with 2 age groups (which is compared with 2 yearly empirical datasets) and that with with 32 age groups (which is compared with 6 empirical yearly datasets). The yearly average discrepancies listed in table 1 are the sum of the average discrepancy with across each empirical dataset. Thus, this difference may contribute to the result that the yearly average discrepancies of the model with 32 age groups are greater than the model with 2 age groups. However, we also employ the particle filtering model with 2 age groups as the minimum average discrepancy model to explore the performance of the outbreak prediction of pertussis below.

Fig 2 shows the boxplot of the distribution of the datasets of discrepancies among the calibrated model and the four particle filtering models, where the boxplot summarizes monthly discrepancy estimates for a given model at different times. Each of the particle filtering models was run 5 times (with the random seed being generated from same set). Then the average monthly and yearly discrepancy among these five runs at each time between the particle filtering models and the empirical data are plotted. Both the monthly and yearly (is adjusted to unit of Month by dividing 12) distribution of the discrepancies of each age structured models are plotted in Fig 2. This boxplot also indicates that when considered over time, the discrepancies of all the particle faltering models tend to be smaller than for the calibrated model, although there are similar median discrepancy values. More notable yet is the fact that the discrepancies associated with the calibrated model are significantly more variable than those for the particle filtered models. This suggests that particle filtering improves the consistency of the model's match against empirical data, when compared to a traditional deterministic model with calibrated parameters. Finally, it bears note that, the datasets of the discrepancy of the model PF_{age_2} has a particularly narrow distribution, especially for the dataset of the yearly discrepancy.

Fig 3 compares the output of the calibration model and the empirical data. It indicates that the deterministic model with calibrated parameters encounters difficulties in tracking the oscillation of the outbreak of pertussis almost across the entire model

time horizon, because the deterministic model of pertussis would approaches a stable equilibrium. These results indicate that the particle filtering models considered here can not only decrease the discrepancy between the model results and the empirical data , but can further track the oscillation of the outbreaks of pertussis – by comparing the results of particle filtering model and deterministic model with calibrated parameters. All of these results suggest that incorporating particle filtering in the compartmental model of pertussis could enhance the simulation accuracy and support more accurate outbreak tracking.

Fig 3. Reported pertussis cases of the calibration model (monthly).

Fig 4. 2D histogram posterior result over the total timeframe for the aggregate particle filtering pertussis model. The posterior result is sampled after the weight is updated when the empirical data coming at each time step.

Fig 5. 2D histogram prior result over the total timeframe for the aggregate particle filtering pertussis model. The prior result is sampled before the weight is updated when the empirical data coming at each time step.

Fig 6. 2D histogram posterior result over the total timeframe of the two-age stratified pertussis model. (a) the monthly particle filtering result summed for the entire population. (b) the yearly particle filtering result for the child (top) and adult (bottom) age groups.

Fig 7. 2D histogram posterior result over the total timeframe of the age structured model of 32 age groups with the Hethcote contact matrix. (a) the monthly particle filtering result summed for the entire population. (b) the yearly particle filtering results of each age group of empirical datasets; age groups are successively older from top to bottom.

Fig 4 presents the posterior results of the pertussis particle filtering model with aggregate population structure for the entire timeframe. For this diagram, the results of the particle filtering model are sampled according to the weight of all particles after the weights are updated by incorporating the empirical data from the current time point. Those time-specific values are then plotted; the values of empirical data points are shown in red, while the sampled posterior distribution of particle filtering model are shown in blue. The blue color saturation indicates the relative density of sampled points within a given 2D bin. Fig 4 demonstrates that most of the empirical data points are located in or near the high density location of the results of the particle filtering model. It further indicates that the particle filtering model has the capability to track the outbreak of pertussis over time, especially compared with the calibrated model whose results are shown in Fig 3. It is notable that the particle filtered results can follow the patterns of empirical data as they arrive; this capacity to update its estimate of model state – both latent and observed – in accordance as new data arrives is central to the function of particle filtering. By contrast, calibration lacks a means of updating the estimate of the model state over time, and is instead relegated to estimating parameter values, rather than the values of the state at varying points in time.

Fig 5 shows the prior results of the pertussis particle filtering model with aggregate population structure for the entire timeframe. For the prior diagram, the results are

sampled before the weight updated when the current empirical data come. Compared with the posterior results shown in Fig 4, the prior values of sampled particles of Fig 5 spread in a wider range. This difference in dispersion indicates that the weight update process of particle filtering algorithm in this paper has the capability to combine the empirical data to the particle filtering model to constrain the particles in a tighter range as suggested by the empirical data.

Fig 6 displays the 2D histogram plots comparing both the monthly and yearly empirical datasets (on the one hand) with the distributions of samples from the posterior distribution of incident cases from the pertussis age structured pertussis particle filtering model with 2 age groups (denoted as PF_{age_2}) (on the other). This figure demonstrates that the model PF_{age_2} is capable of tracking and simulating the outbreaks of pertussis, as evidenced by the fact that most of the monthly and yearly empirical data (shown in the red dashes) in each month are located in or near the high density area of the sampled distribution of the particle filtering model (shown in blue in the 2D histogram plots).

Fig 7 displays the 2D histogram plots comparing both the monthly and yearly empirical datasets (on the one hand) with the sampled posterior distribution of incident cases from the pertussis age structured pertussis particle filtering model with 32 age groups and the Hethcote contact matrix (denoted as $PF_{age_32_Hethcote}$) (on the other). It is notable that the total number of the yearly empirical datasets employed is 6. This figure also demonstrates that the model $PF_{age_32_Hethcote}$ is capable of tracking and simulating the outbreaks of pertussis, because most of the monthly and yearly empirical data in each time are located in or near the high density area of the results of the particle filtering model.

Fig 4, Fig 6 and Fig 7 represent the 2D histogram posterior result of all the particle filtering models, except for the age-structured model of 32 age groups with a re-balanced contact matrix. Results are omitted for this final model as they are highly similar to the 32-age-group model with Hethcote contact matrix, which is shown in Fig 7. The 2D histogram plots shown indicate that both the age-structured particle filtering models and the aggregate population particle filtering model have the capability to closely track the outbreak pattern of pertussis. The results of the models could match the empirical datasets quite well, including both monthly empirical dataset and yearly empirical datasets. In contrast to the calibrated model whose results are shown in Fig 3, the particle filtering models are capable of localizing the model's prediction of empirical data near the empirical data. Although the results in table 1, Fig 2 (for discrepancy) and Fig 4, Fig 6, Fig 7 (for posterior distribution) suggest that all four pertussis particle filtering models are capable of tracking and estimating the pertussis outbreaks, in the interest of brevity of exposition, we employ the minimum discrepancy model – the age-structured particle filtering pertussis model of 2 age groups – to perform the prediction and intervention analysis below.

The pertussis particle filtering models described here are particularly notable by virtue of their capacity to perform ongoing estimation of entire model state – including the latent and observable state of the dynamic models during the period of incorporating the empirical observations. In the investigation of pertussis considered here, the empirical data – monthly population-wide pertussis reported cases and yearly reported pertussis cases of different age groups – are only related to an aggregation of the latent states of weakly infectious (I_w), medium infectious (I_m), and fully infectious (I). However, by applying particle filtering to the compartmental models of pertussis and the empirical data, the pertussis particle filtering models are capable of estimating the entire model state across the whole model time horizon during which empirical datasets on the basis of incorporation of multiple lines of empirical data. Fig 8 shows the 2D histogram plots of samples from the distributions of latent stocks with the

minimum discrepancy model (the two age stratified model incorporating both the monthly and yearly empirical datasets) as an example. Figure 8 indicates that most of the fully infectious individuals are estimated to be located amongst the children (less than 5 years of age) age group, while most of the weak infectious and medium infectious individuals are located in the "adult" (equal and greater than 5 years) age group. Most of the recovered individuals are also located in the adult age group. Such estimates accord with empirical observations with respect to pertussis transmission and build confidence in the capacity of the model to meaningfully estimate the latent state. As noted below, estimation of latent state not only provides strong support for understanding the current epidemiological context, but can also be an important enabler for projecting results forward to periods where empirical data is not yet available, for understanding of the effects of interventions, as well as for providing retrospective insights into historical circumstances.

Fig 8. 2D histogram results for the latent states from particle filtering of the two age group stratified dynamic models. (a) the child age group (those within their first 5 years of life). (b) the adult age group (years 5 and up). The latent states in plots both (a) and (b) are "S", "I", "I_m", "R_1", "R_2", "R_3", "I_w", "R_4" in order.

3.2 Prediction with the minimum discrepancy model

To assess the predictive capacity of the pertussis particle filtering models in anticipating outbreaks, we performed out-of-sample prediction experiments. Informally, each such experiments examines the capacity of the model to project results into the future, having considered data only to some "current" time. That is, the model is particle filtered so as to incorporate data only to up to but not including a "Prediction Start Time" (T^*), and then begins projecting (predicting) forward, starting at T^*. More specifically, in this process, the weights of particles will cease updating in response to observations at time T^*, following that point, all of the particles run without new empirical data being considered. In this paper, all of the prediction experiments are run 4 years after the "Prediction Start Time" T^*. To evaluate the predictive capacity of the model, we examined the effects of changing the prediction start time T^* so as to pose different archetypal types of prediction challenges. It is notable that the minimum discrepancy model – the age structured model with 2 age groups where the child age group is up to 5 year and incorporating both the monthly and yearly empirical datasets, which is identified in the previous section – is employed to perform all of these experiments.

(1) Prediction started from the first or second time points of an outbreak.

(2) Prediction started before the next outbreak.

(3) Prediction started from the peak of an outbreak.

(4) Prediction started from the end of an outbreak.

Fig 9. 2D histogram depicting prediction using the minimum discrepancy model from the first or second time points of an outbreak. (a) prediction from month 190. (b) prediction from month 269.

Figs 9–12 display the prediction results of these situations with respect to the monthly 2D histogram of reported cases of the total population. In the 2D histogram

Fig 10. 2D histogram depicting prediction using the minimum discrepancy model from the peak of an outbreak. (a) prediction from month 176. (b) prediction from month 233.

Fig 11. 2D histogram depicting prediction using the minimum discrepancy model from the end of an outbreak. (a) prediction from month 209. (b) prediction from month 296.

Fig 12. 2D histogram depicting prediction using the minimum discrepancy model prior to the next outbreak. (a) prediction from month 99. (b) prediction from month 216.

plots of Figs 9–12, the empirical data having been considered in the particle filtering process (i.e., incorporated in training the models) are shown in red, while the empirical data not having been considered in the particle filtering process (and only displayed to compare with model results) are shown in black. The vertical straight line labels the "Prediction Start Time" (T^*) of each experiment.

These prediction results suggest that the pertussis particle filter model offers the capacity to probabilistically anticipate pertussis dynamics with a fair degree of accuracy over a year or so. From the 2D histogram plots, empirical data lying after the prediction start time – and thus not considered by the particle filtering machinery – mostly lie within the high-density range of the particles. While not formally evaluated here, it is notabe that in such examples the particle filter model appears to be able to accurately anticipate a high likelihood of a coming outbreak and non-outbreak. Such an ability could offer substantial value for informing the public health agencies with accurate predictions of the anticipated evolution of pertussis over the coming months.

3.3 Intervention with the minimum discrepancy model

Based on the previous discussion, the capacity to accurately estimate (sample from) the latent state of a model allows particle filtering used with pertussis models examined here to be capable of estimating the entire latent state and projecting outbreak occurrence and progression. The capacity for such state estimation also support particle filtering models in simulating intervention strategies.

In this section, we have implemented several experiments to simulate stylized public health intervention policies, based on the minimum discrepancy particle filtering pertussis model identified above. The stylized intervention strategies are characterized in an abstract way, and are typically performed before or at the very beginning of an outbreak. For simplicity, we examine them as a historical counterfactual – taking place at a certain historic context. Moreover, to support easy comparison with the baseline prediction results of the minimum discrepancy model absent any interventions, all of the intervention strategies are simulated starting at the start month of an outbreak (month 269) in this project. Moreover, we assume here that month 269 is the "current time" in the scenario – that we wish to examine the effects of that intervention using all data and only data available up to but not including month 269, and simulate the results of the intervention forward from that point. The baseline prediction result of the minimum discrepancy model without any interventions is shown in Fig 9 (b). We examine below the impact of two stylized intervention policies – quarantine and vaccination.

Fig 13. 2D histogram of simulating quarantine during a pertussis outbreak. This is realized by decreasing the contact rate by 20%.

Fig 14. 2D histogram of simulating quarantine during a pertussis outbreak. This is realized by decreasing the contact rate by 50%.

Fig 15. 2D histogram of simulating an immunization campaign during an outbreak. This is realized by characterizing a stylized vaccine-induced protection level among 20% of the population.

Fig 16. 2D histogram of simulating an immunization campaign during an outbreak. This is realized by characterizing a stylized vaccine-induced protection level among 50% of the population.

Fig 13 and Fig 14 display results from simulation of quarantine intervention strategies [24] whose effects are characterized as decreasing the contact rate parameter to be 20% and 50% less than its pre-intervention value, respectively. Similarly to the 2D histogram plot of the prediction result without any intervention shown in Fig 9 (b), the red dots represent the empirical data incorporated in the particle filtering model (here, up to just prior to the point of intervention), while the black dots represent the empirical data not incorporated in the model, but presented for comparison purposes. It is notable that because the interventions being characterized are counterfactual in character – i.e., did not in fact take place historically – the empirical data shown in black reflect the situation of absent any intervention (that is, in a baseline context that lacked an intervention of the sort simulated here). By comparing the quarantine intervention results (see Fig 13 and Fig 14) with the model result without intervention shown in Fig 9 (b) and the empirical data during the intervention period (that laying after the triggering of the intervention, and not incorporated in the particle filtering model), we can see that, although the interventions are implemented in a stylized fashion, by virtue of its ability to estimate the underlying epidemiological state at the point of intervention, the pertussis particle filtering model is capable of using the estimated state to serve as a tool for probabilistically evaluating pertussis related intervention policies.

To simulate an immunization intervention during a pertussis outbreak, a vaccination parameter is incorporated to represent the fraction of the population moved with respect to their immunity status. Specifically, recall that the pertussis model characterizes a chain of successively higher levels of vaccine-induced protection; this parameter moves a certain fraction of the population from the previous stock (before vaccination) to the stock representing the next higher level of vaccination (following vaccination). Fig 15 and Fig 16 show the results of the vaccination intervention. The layout of the 2D histogram plots of the vaccination interventions with Fig 15 and Fig 16 is same as the quarantine plots of Fig 13 and Fig 14. These results of pertussis interventions demonstrate that by virtue of its ability to estimate the underlying epidemiological state, the use of particle filtering with pertussis models supports evaluating public health intervention policies to reduce or even avoid the outbreak of pertussis, etc.

4 Discussion and conclusion

This paper presents a new method of tracking pertussis outbreak patterns by integrating a particle filtering algorithm with a pertussis compartmental model and empirical incidence data. This contribution represents the first time that the particle filtering algorithm has been incorporated in the research of pertussis transmission dynamics. The models examined demonstrated a notable degree of accuracy in predicting pertussis dynamics over multi-month timeframes – the 2D histogram plots comparing the empirical data and samples from the posterior distributions of the particle filtering

models' output data (the monthly and yearly reported cases of pertussis) indicates that the high probability density region of the model's prediction of empirical data encompasses or lies near those data. Moreover, the discrepancy of the pertussis particle filtering model's predictions vs. observed data is reduced by approximately 60% when compared with a traditional calibration model. This reduction indicates that the particle filtering algorithm is capable of enhancing model predictions when compared with predictions resulting from the traditional calibration process. Additionally, it bears emphasis that the calibrated deterministic model encounters marked difficulties in tracking the fluctuation of the outbreak pattern of measles, while the particle filtering model is capable to tracking the oscillation of the outbreak of pertussis.

Moreover, it is of great significance that the pertussis particle filtering models could support effective estimation of the entire state of the pertussis transmission models during the time incorporating the empirical datasets – including latent states of strong interest, including states associated with waning of natural immunity and differing levels of infection severity. Combined with the capability of performing outbreak projections, such particle filtering models can serve as powerful tools for understanding the current epidemiology of pertussis in the population, for projecting forward evolution of pertussis spread – including occurrence of outbreaks.

Beyond that, in a further contribution, this research further marks the first instance of research demonstrating the capacity to perform public health intervention experiments using particle filtering models.

Despite the strengths of these contributions, there remain a number of limitations of this work, and priorities for future research. We briefly comment on several.

Within this work, four particle filtering models have been researched, including an aggregate population model, a two age group-stratified population model, and 32 age group population models with both a contact matrix derived from Hethcote (1997) [15] and (separately) using a re-balanced contact matrix. Although the results of all these four particle filtering models could match the empirical data quite well, the minimum discrepancy model emerged from the 2 age group age-stratified particle filtering model, where the individuals in the child age group represent children in the first 5 years of life, and incorporating both monthly and yearly empirical datasets. In this regard, it is notable that according to the mathematical deduction of the age structured population model introduced in [7] – adapted to pertussis in this research – with more age groups considered in the age-structured model, the model can simulate the aging rate (c_i) more accurately. However, in this paper, the discrepancies associated with both of the 32 age group particle filtering models do not show evident decrease compared with the two age group particle filtered models. We provide here some possible reasons are listed as follows. Firstly, the stochastic processes considered in both the 32 age group age-structured model and the two age-groups age-structured model are different, especially in characterizing the stochastic evolution of the contact rate. Secondly, the likelihood functions employed in this project – which are captured as the product of negative binomial density functions across all empirical datasets and sharing a common dispersion parameter – may be too naïve to capture the difference between the age groups within the empirical datasets. Thirdly, as the number of age groups increase, the dimensions of the state space of the particle filtering models increase dramatically. This latter issue must be considered in light of the limitations of the particle filtering algorithm, particularly the fact that the particle filtering method employing a condensation algorithm may encounter problems in high-dimensional systems. In such systems, the probability density functions would be more complex, which may require high dispersion for representation using the likelihood functions employed, because it is difficult to represent the details of the multivariate likelihood function using the product of simple probability density functions. Research is needed into more effective

multivariate likelihood function design. The relationship between the the nominal state space dimensionality and the number of particles required for effective particle filtering also merits additional research, particularly in late of observed limitations in the benefit of particle filtering for high dimensional systems [12]. Finally, when comparing the discrepancy for distinct models, our lack of normalization for the count of datasets used may lead to artificially stacking the comparison against the 32 age group model; while the 32 age group does not exhibit markedly better discrepancy against the monthly aggregate observations than does the 2 age group model, this consideration may suggest that it is stronger than the yearly discrepancy numbers would suggest.

It also bears emphasizing the critical role of the stochastic process noise within the state space models plays within successful particle filtering, and the challenges associated with managing such noise. The stochastics associated with these factors represents a composite of two factors. Firstly, there is expected to both be stochastic variability in the measles infection processes and some evolution in the underlying transmission dynamics in terms of changes in reporting rate, and changes in mixing. Secondly, such stochastic variability allows characterization of uncertainty associated with respect to model dynamics – reflecting the fact that both the observations and the model dynamics share a high degree of fallability. Such stochastics impact the particle filtering model in distinct ways during the estimation and prediction period, but results in both periods are sensitive to the degree of stochastics involved. It was important that the noise in the particle filtering models in this paper should be controlled in a proper range, by tuning the parameters of diffusion coefficients in the stochastic processes related to the Brownian motion. The need to characterize and tune stochastic noise effectively can impose challenges in the speed with which particle filtering models can be prepared for a new area of application.

The initial values of the age-structured population models in this paper are estimated both manually and by the particle filtering algorithm. Specifically, the population distribution among the different age groups are tuned manually, while the population distribution among different stocks within a given age group is estimated by the particle filtering algorithm by setting the initial values of stocks in a proper range following a uniform distribution, but maintaining a total number of individuals for that age group across the stocks. Especially in building the 32-age-groups particle filtering models, much time and efforts was been used in estimating the population distribution among the different age groups.

Particularly in light of the growing risk of pertussis outbreaks triggered by combinations of vaccine hesitancy and waning dynamics from earlier generations of MMR vaccine, a keen need for future work involves application the models presented here to data in the vaccination era. While vaccination elements of the models discussed here were only glancingly tapped by this research (in the context of demonstrating capacity to reason about the effects of an immunization intervention), because of their incorporation into the existing model structure, extension of this work to vaccine-era dynamics should require only very limited changes to the models involved.

While application of particle filtering to pertussis dynamics is not without its challenges, the approach examined here demonstrates great promise for creating models that are automatically kept abreast of the latest evidence, for understanding the underlying epidemiology of pertussis in the population – including the balance of the population at varying levels of immunity – for projecting forward pertussis dynamics and outbreak prediction over a year's time, and for evaluation of counter-factual interventions. The results of this paper – which represents both the first application of particle filtering to pertussis and the first to use particle filtering to assess the tradeoff between public health applications – suggest that particle filtering may represent an important element in the arsenal of public health tools to address the increasingly

difficult challenge of controlling pertussis in the context of vaccine hesitancy and waning of both natural- and vaccine induced- immunity.

Supporting information

S1 Appendix. The introduction of mathematical models.

S2 Appendix. The introduction of Particle filter implementation.

S3 Appendix. Results of the comparison of error plot between calibration model and particle filtering model and the comparison of running time of particle filtering models.

S4 Appendix. Proof of the n-square grows of the unknown parameters.

S5 Appendix. Mathematical deduction of the parameter of death rate in the age-structured demographic model.

References

1. World Health Organization. Pertussis, Vaccines and Diseases, Immunization, Vaccines and Biologicals. 2018. `http://www.who.int/immunization/diseases/pertussis/en/`.
2. Centers for Disease Control and Prevention. Pertussis Causes and Transmission. 2017. `https://www.cdc.gov/pertussis/about/causes-transmission.html`.
3. Centers for Disease Control and Prevention. Epidemiology and Prevention of Vaccine-Preventable Diseases.2018. `https://www.cdc.gov/vaccines/pubs/pinkbook/pert.html`.
4. Chow MYK, Khandaker G, McIntyre P. Global Childhood Deaths From Pertussis: A Historical Review. Clin Infect Dis. 2016;63(suppl 4):S134–S141.
5. Centers for Disease Control and Prevention. Whooping Cough and the Vaccine (Shot) to Prevent It. 2015. `https://www.cdc.gov/vaccines/parents/diseases/child/pertussis.html`.
6. Doroshenko A, Qian W, Osgood ND. Evaluation of outbreak response immunization in the control of pertussis using agent-based modeling. PeerJ. 2016;4:e2337.
7. Li X, Doroshenko A, Osgood ND. Applying Particle Filtering in Both Aggregated and Age-structured Population Compartmental Models of Pre-vaccination Measles. PLoS ONE. 2018;13.
8. Safarishahrbijari A, Osgood N. Social Media Surveillance Improves Outbreak Projection via Transmission Models. JMIR Public Health and Surveillance. 2019;.
9. Osgood N, Liu J. Towards Closed Loop Modeling: Evaluatng the Prospects for Creating Recurrently Regrounded Aggregate Simulation Models Using Particle Filtering. In: Proceedings of the 2014 Winter Simulation Conference. WSC '14. Piscataway, NJ, USA: IEEE Press; 2014. p. 829–841.

10. Dukic V, Lopes HF, Polson NG. Tracking Epidemics With Google Flu Trends Data and a State-Space SEIR Model. Journal of the American Statistical Association. 2012;107(500):1410–1426.

11. Oraji R, Hoeppner VH, Safarishahrbijari A, Osgood ND. Combining Particle Filtering and Transmission Modeling for TB Control. In: 2016 IEEE International Conference on Healthcare Informatics (ICHI). ieeexplore.ieee.org; 2016. p. 392–398.

12. Kreuger K, Osgood N. Particle filtering using agent-based transmission models. In: 2015 Winter Simulation Conference (WSC); 2015. p. 737–747.

13. Safarishahrbijari A, Lawrence T, Lomotey R, Liu J, Waldner C, Osgood N. Particle filtering in a SEIRV simulation model of H1N1 influenza. In: 2015 Winter Simulation Conference (WSC). ieeexplore.ieee.org; 2015. p. 1240–1251.

14. Ong JBS, Chen MIC, Cook AR, Lee HC, Lee VJ, Lin RTP, et al. Real-time epidemic monitoring and forecasting of H1N1-2009 using influenza-like illness from general practice and family doctor clinics in Singapore. PloS one. 2010;5(4):e10036.

15. Hethcote HW. An age-structured model for pertussis transmission. Mathematical biosciences. 1997;145(2):89–136.

16. Hethcote H. The Mathematics of Infectious Diseases. SIAM Review. 2000;42(4):599–653.

17. Garnett GP, Bowden FJ. Epidemiology and control and curable sexually transmitted diseases: opportunities and problems. Sex Transm Dis. 2000;27(10):588–599.

18. Arulampalam MS, Maskell S, Gordon N, Clapp T. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. IEEE transactions on signal processing: a publication of the IEEE Signal Processing Society. 2002;50(2):174–188.

19. Murphy KP. Machine Learning: A Probabilistic Perspective. MIT Press; 2012.

20. Blake A, Isard M. The CONDENSATION algorithm-conditional density propagation and applications to visual tracking. In: Advances in Neural Information Processing Systems; 1997. p. 361–367.

21. Dorigatti I, Cauchemez S, Pugliese A, Ferguson NM. A new approach to characterising infectious disease transmission dynamics from sentinel surveillance: Application to the Italian 2009–2010 A/H1N1 influenza pandemic. Epidemics. 2012;4(1):9–21.

22. Annual Report of Department of Public Health in the Province of Saskatchewan; 1921-1956.

23. Doucet A, Johansen AM. A tutorial on particle filtering and smoothing: Fifteen years later. Handbook of nonlinear filtering. 2009;12(656-704):3.

24. Tulu TW, Tian B, Wu Z. Modeling the effect of quarantine and vaccination on Ebola disease. Adv Difference Equ. 2017;2017(1):178.

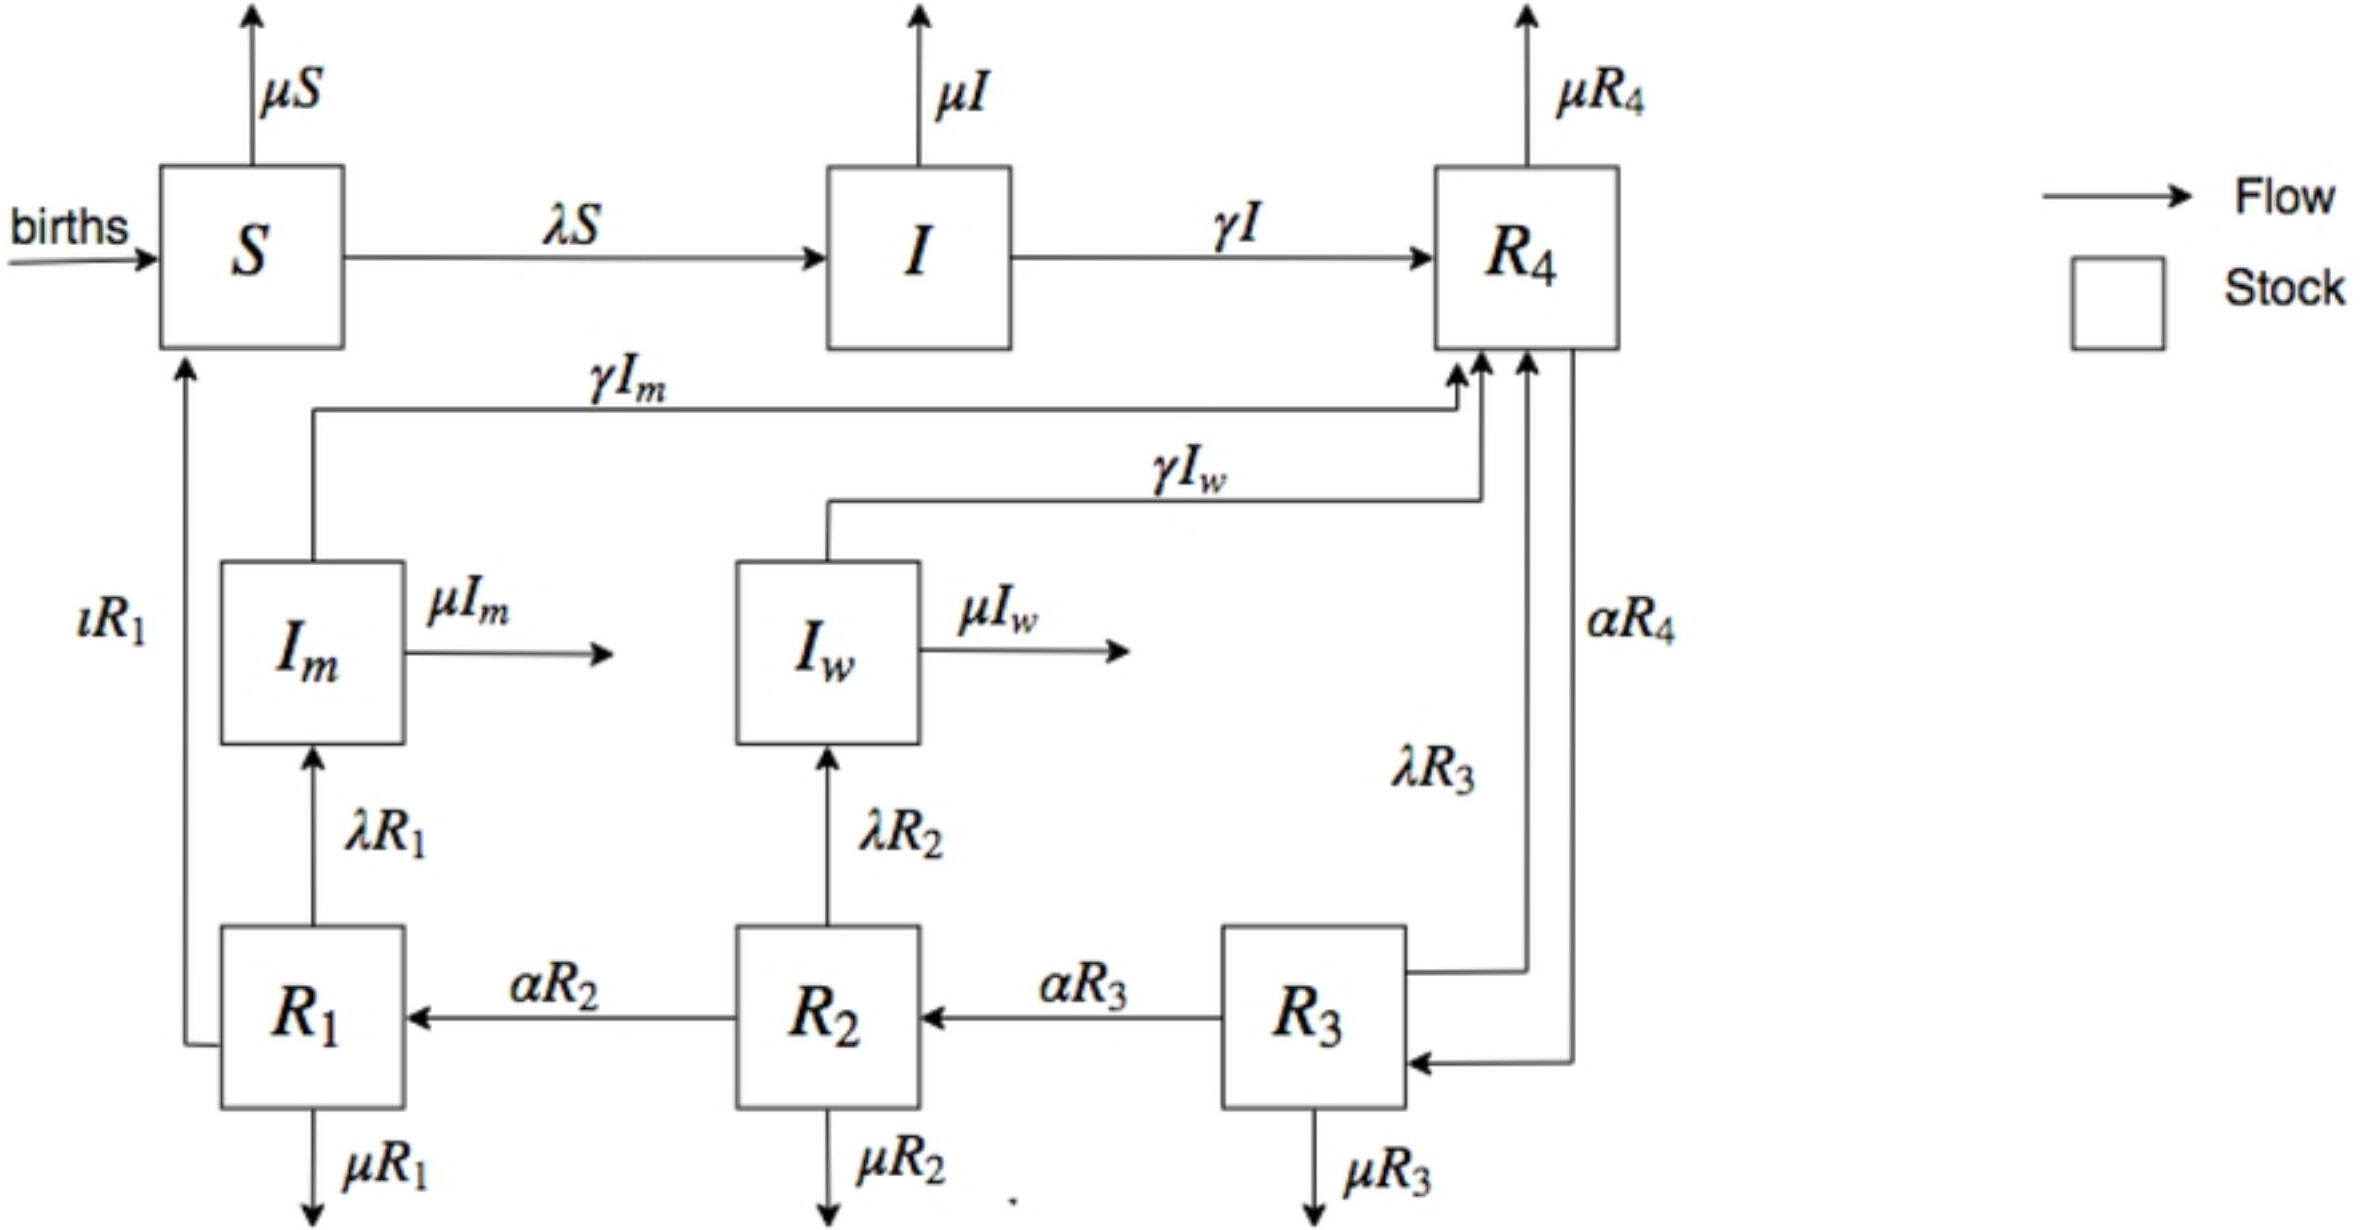
μS
μI
μR_4
births
S
λS
I
γI
R_4
Flow
Stock
γI_m
γI_w
ιR_1
I_m
μI_m
I_w
μI_w
αR_4
λR_3
λR_1
λR_2
R_1
αR_2
R_2
αR_3
R_3
μR_1
μR_2
μR_3

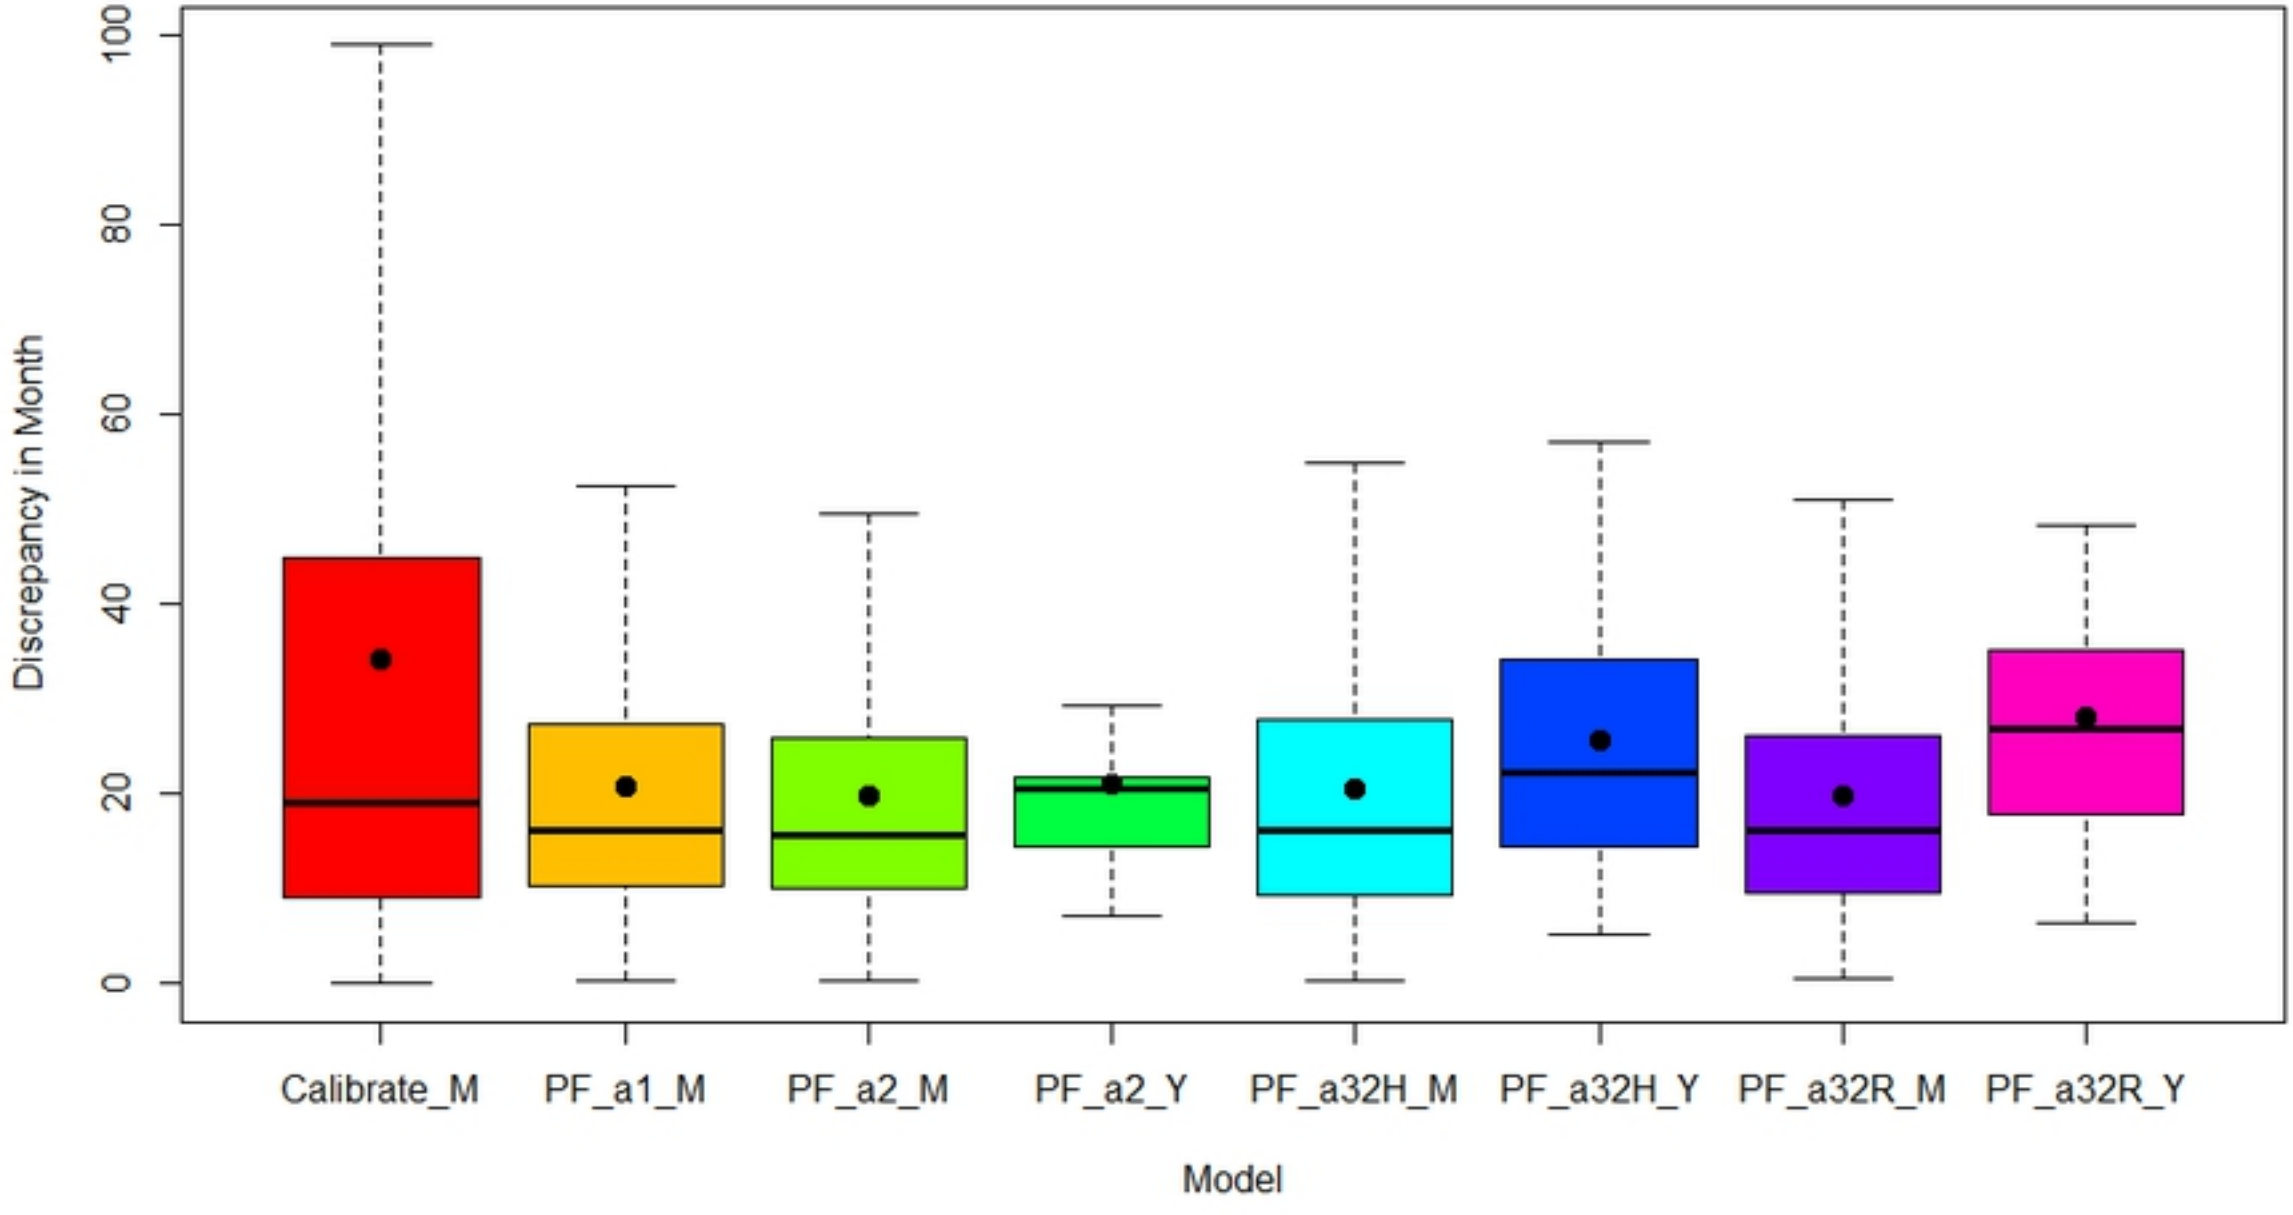
Discrepancy in Month
0
20
40
60
80
100
Calibrate_M
PF_a1_M
PF_a2_M
PF_a2_Y
PF_a32H_M
PF_a32H_Y
PF_a32R_M
PF_a32R_Y
Model

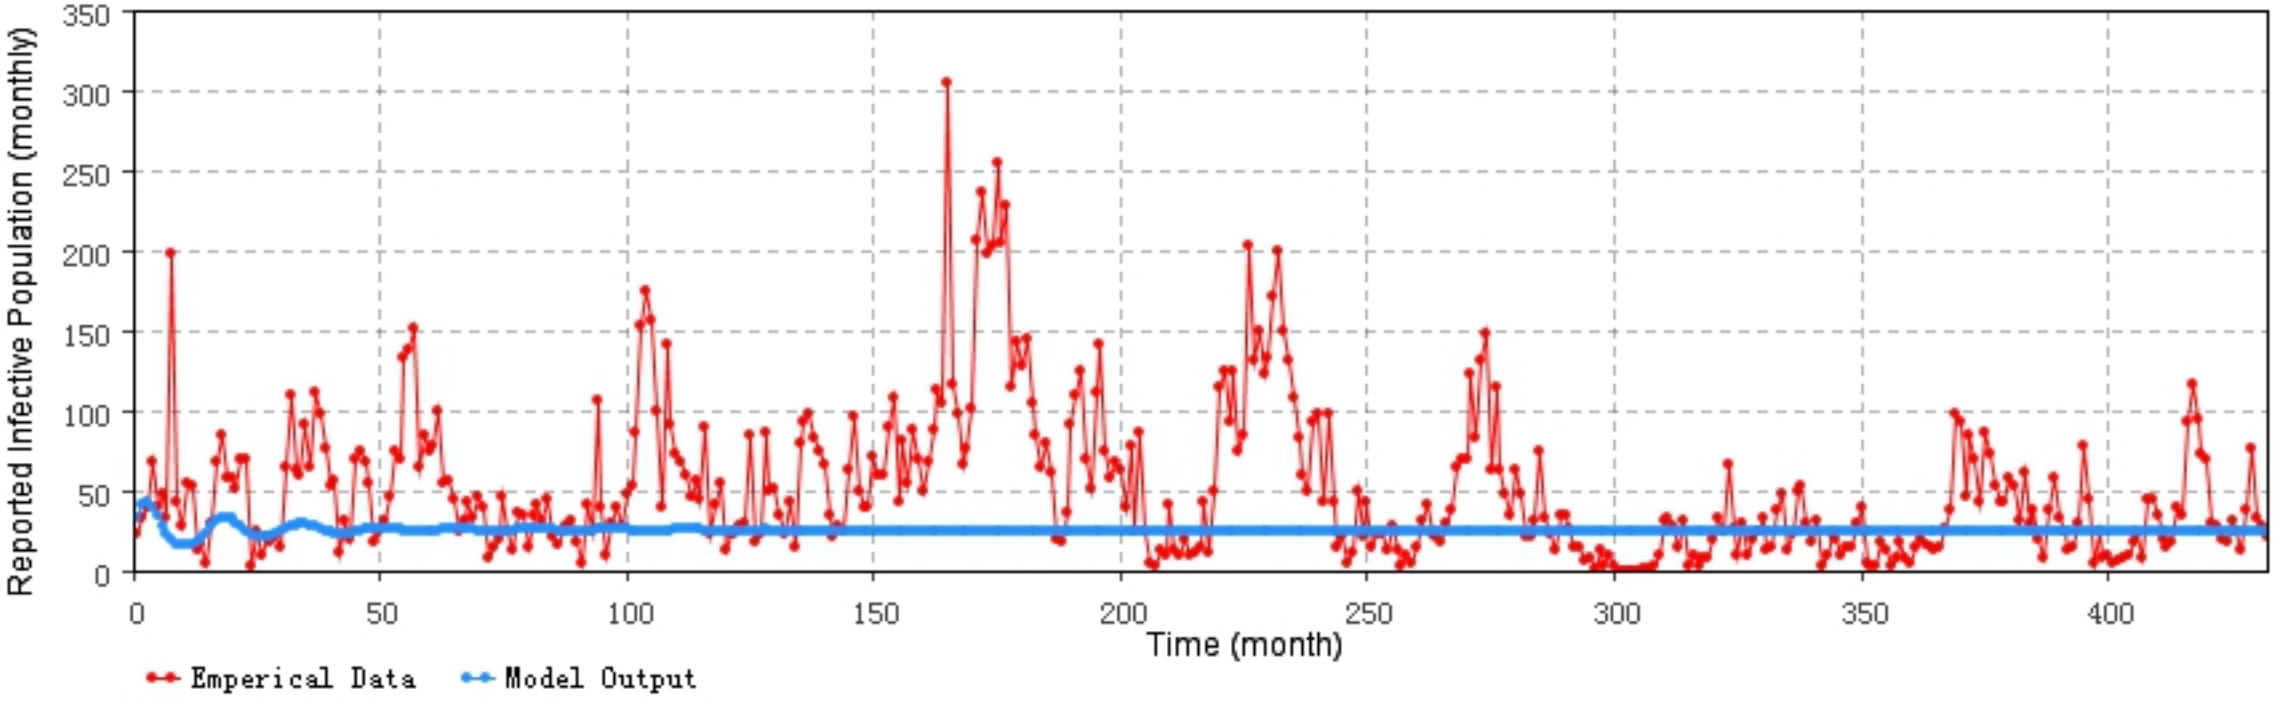
Reported Infective Population (monthly)
350
300
250
200
150
100
50
0
0
50
100
150
200
250
300
350
400
Time (month)
Emperical Data
Model Output

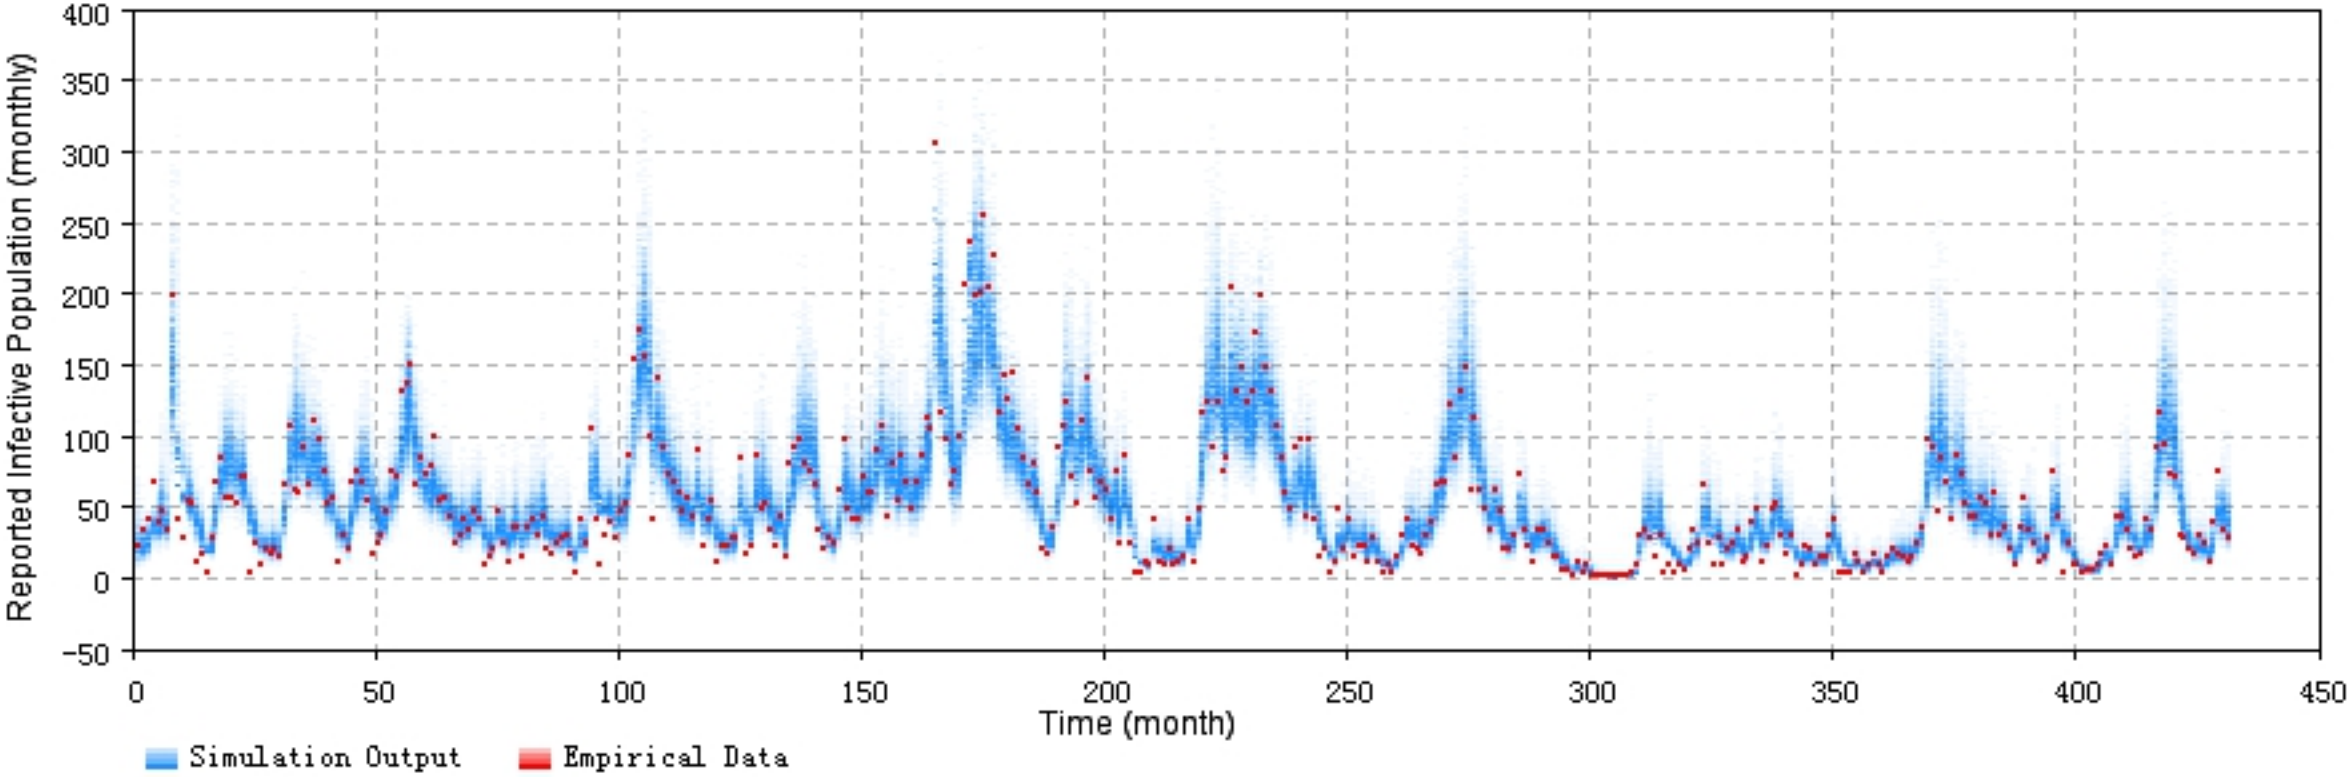
Reported Infective Population (monthly)
400
350
300
250
200
150
100
50
0
-50
0
50
100
150
200
250
300
350
400
450
Time (month)
Simulation Output
Empirical Data

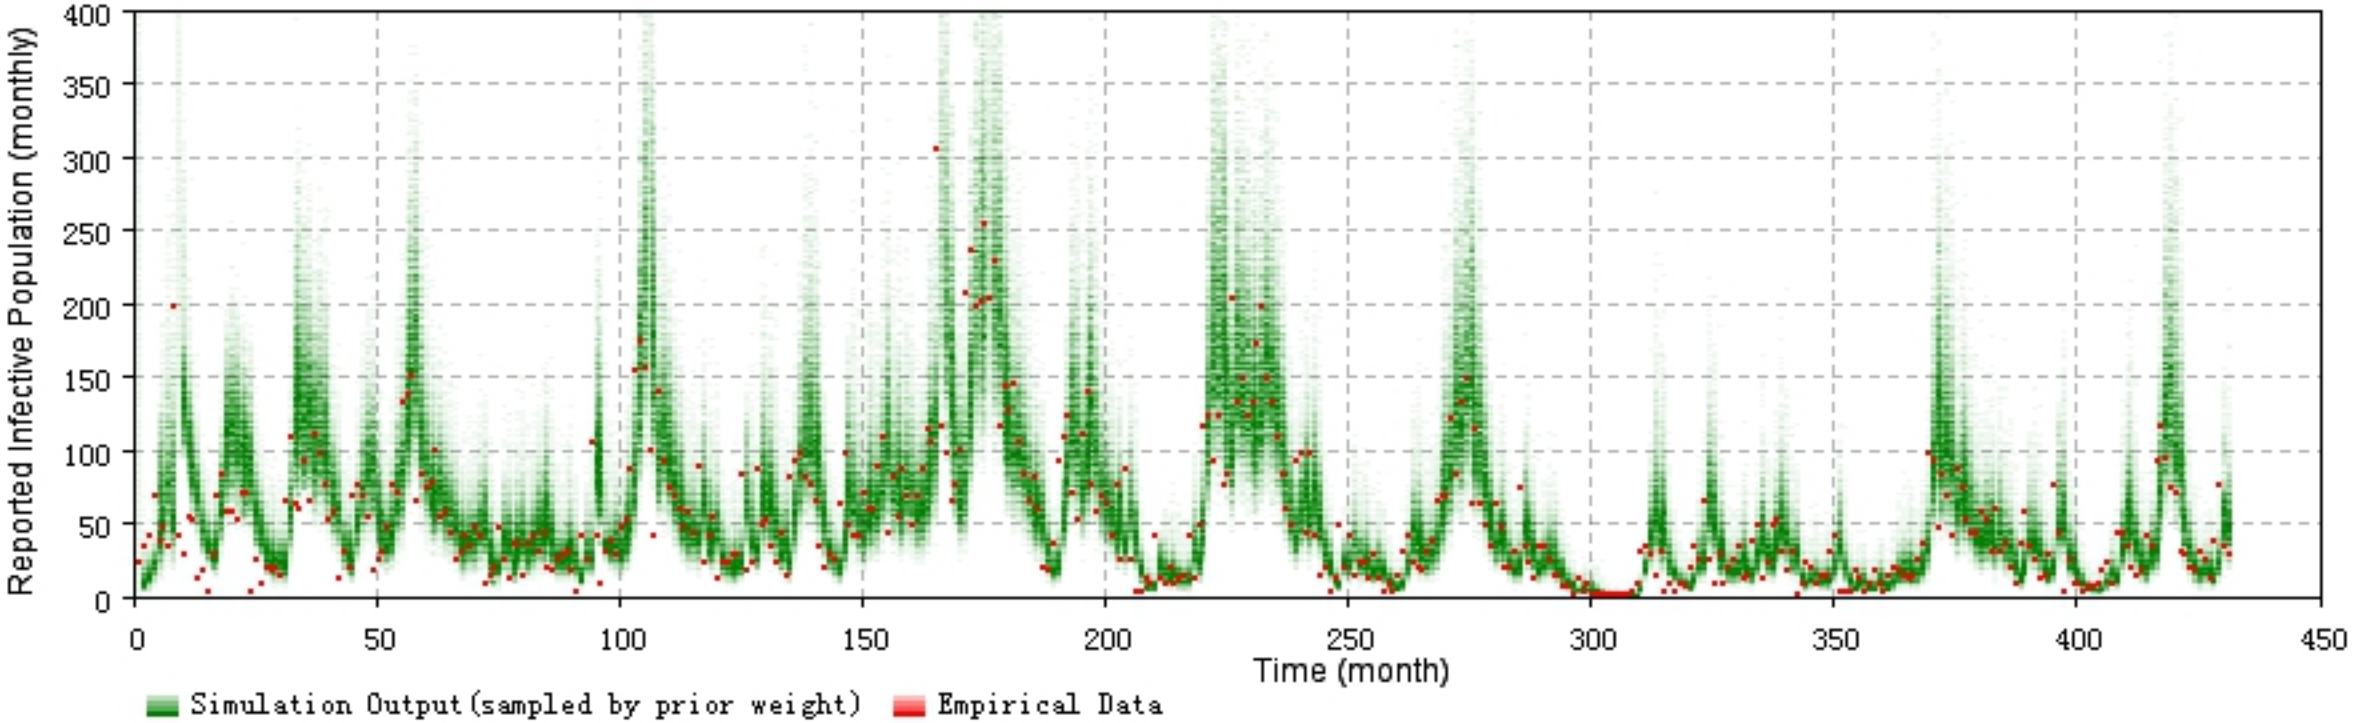
Reported Infective Population (monthly)
400
350
300
250
200
150
100
50
0
0
50
100
150
200
250
300
350
400
450
Time (month)
Simulation Output (sampled by prior weight)
Empirical Data

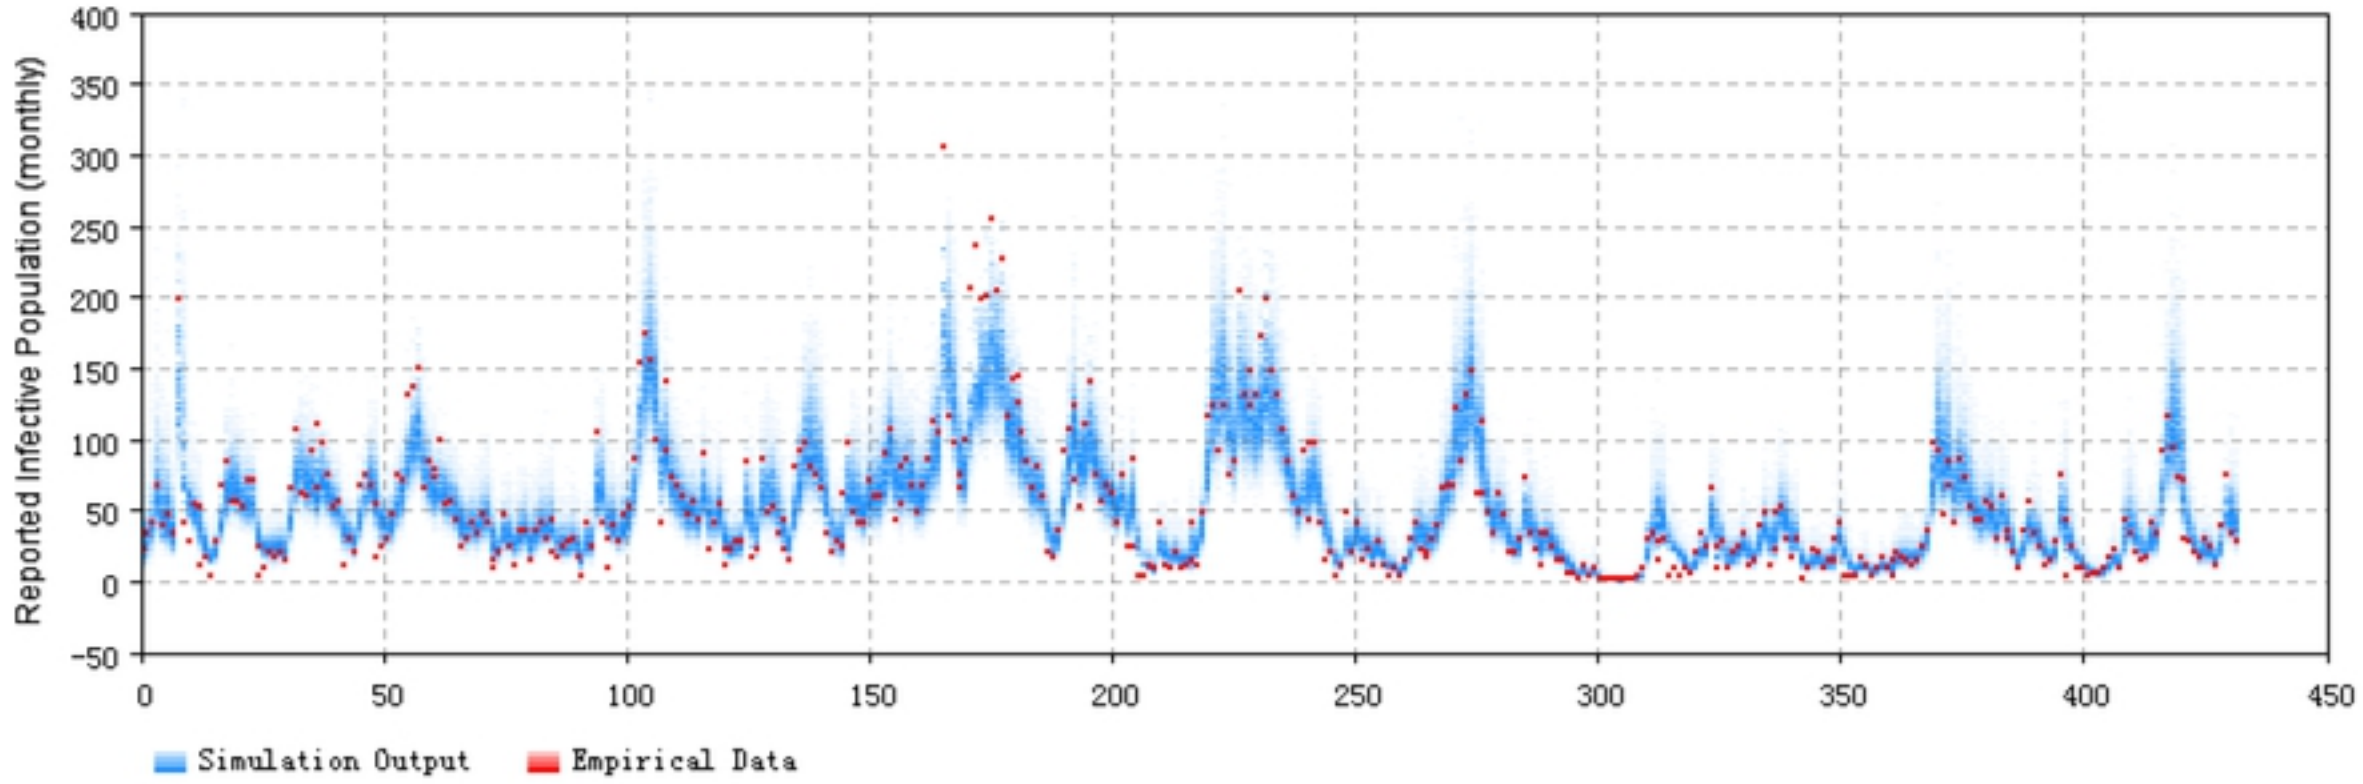
Reported Infective Population (monthly)
400
350
300
250
200
150
100
50
0
-50
0
50
100
150
200
250
300
350
400
450
Simulation Output
Empirical Data

(a)

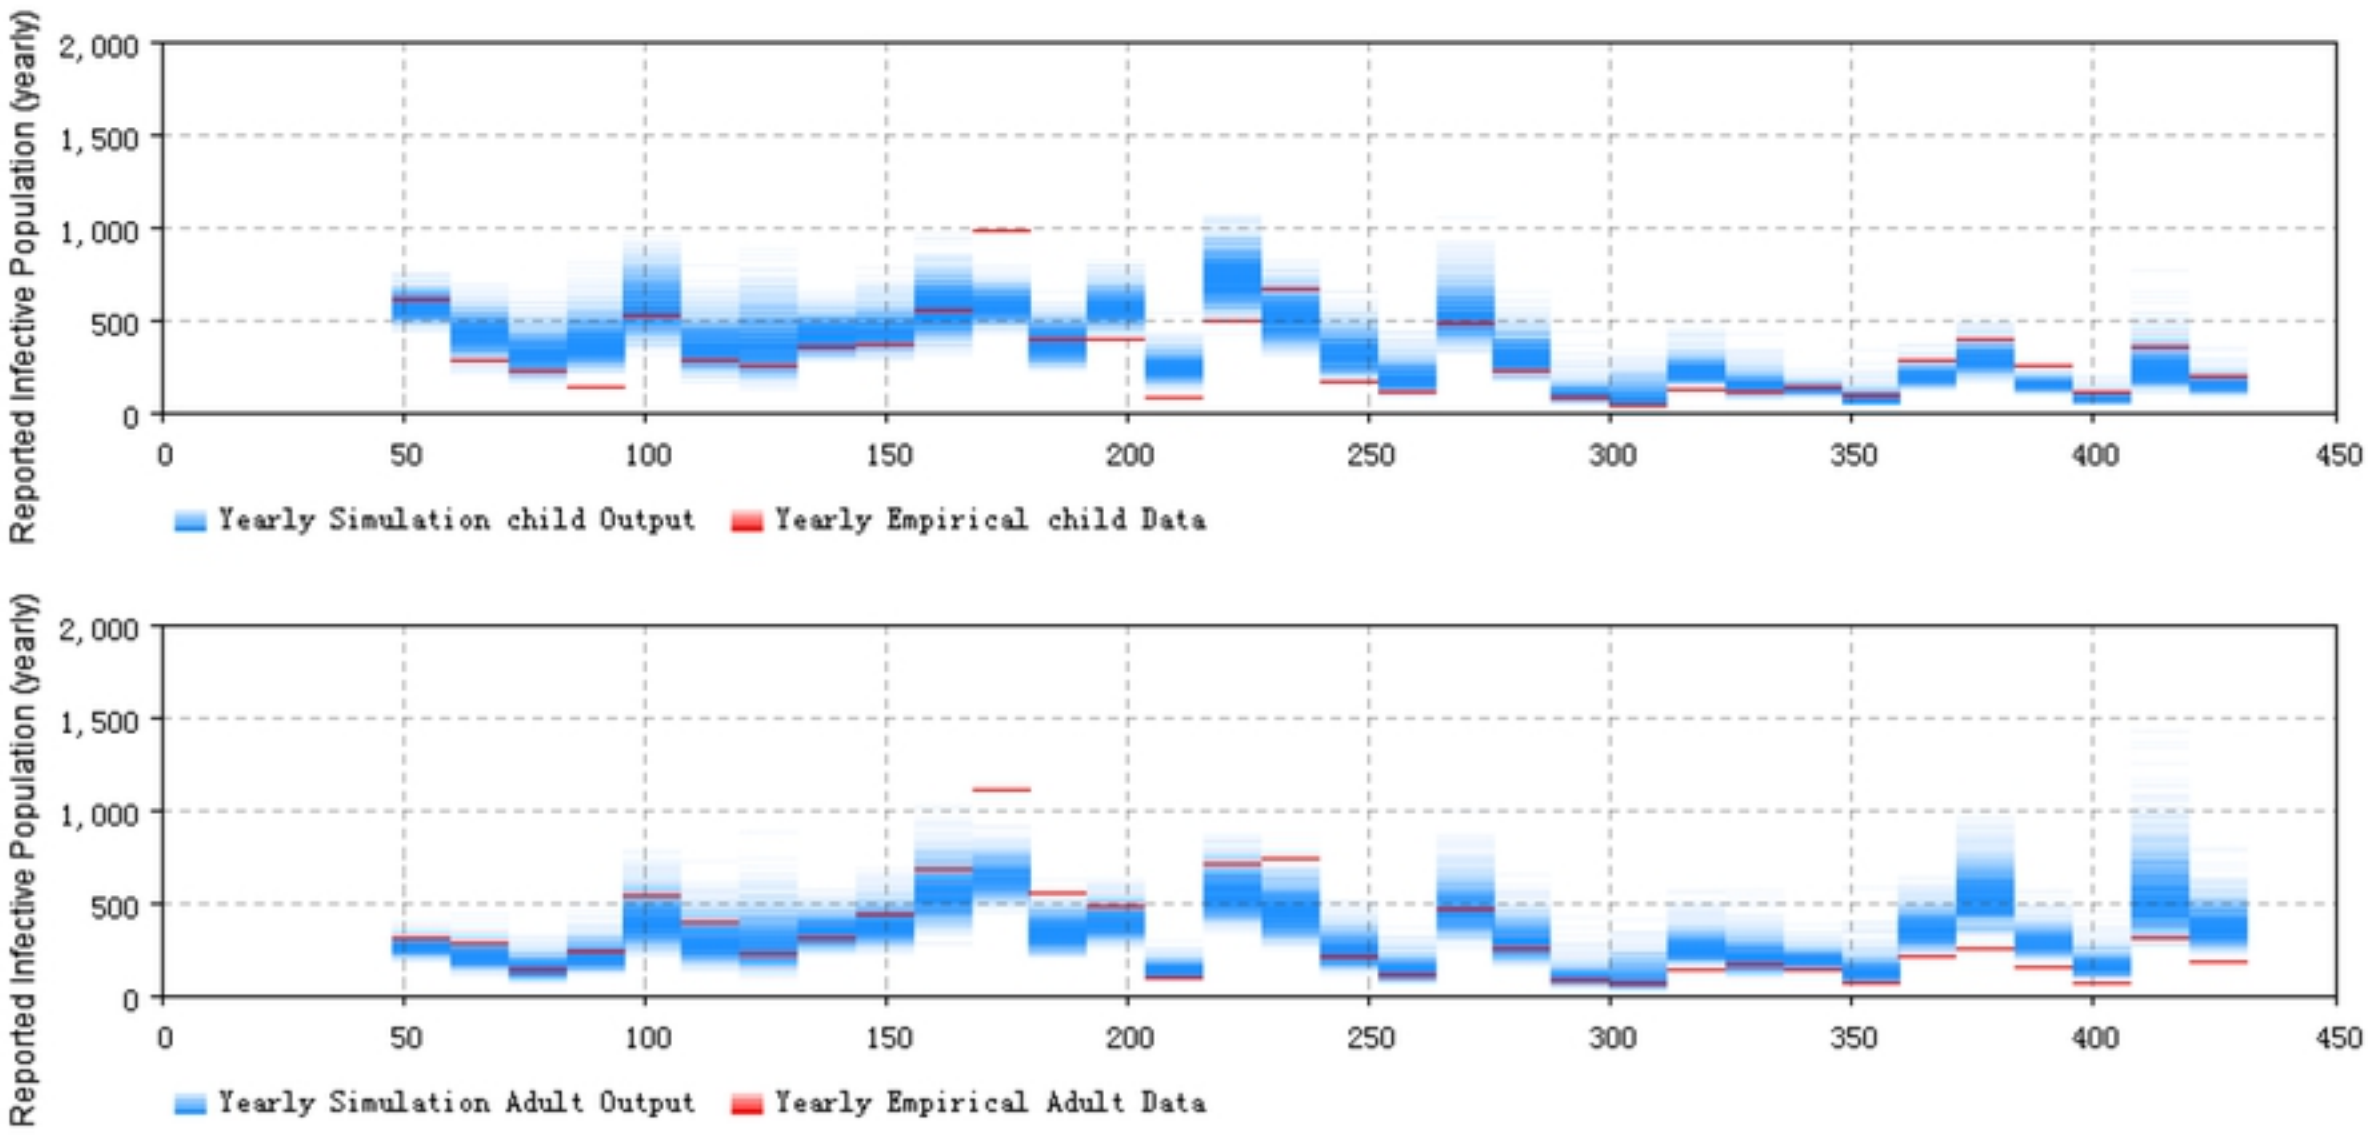
Reported Infective Population (yearly)
2,000
1,500
1,000
500
0
0
50
100
150
200
250
300
350
400
450
Yearly Simulation child Output
Yearly Empirical child Data
Reported Infective Population (yearly)
2,000
1,500
1,000
500
0
0
50
100
150
200
250
300
350
400
450
Yearly Simulation Adult Output
Yearly Empirical Adult Data

(b)

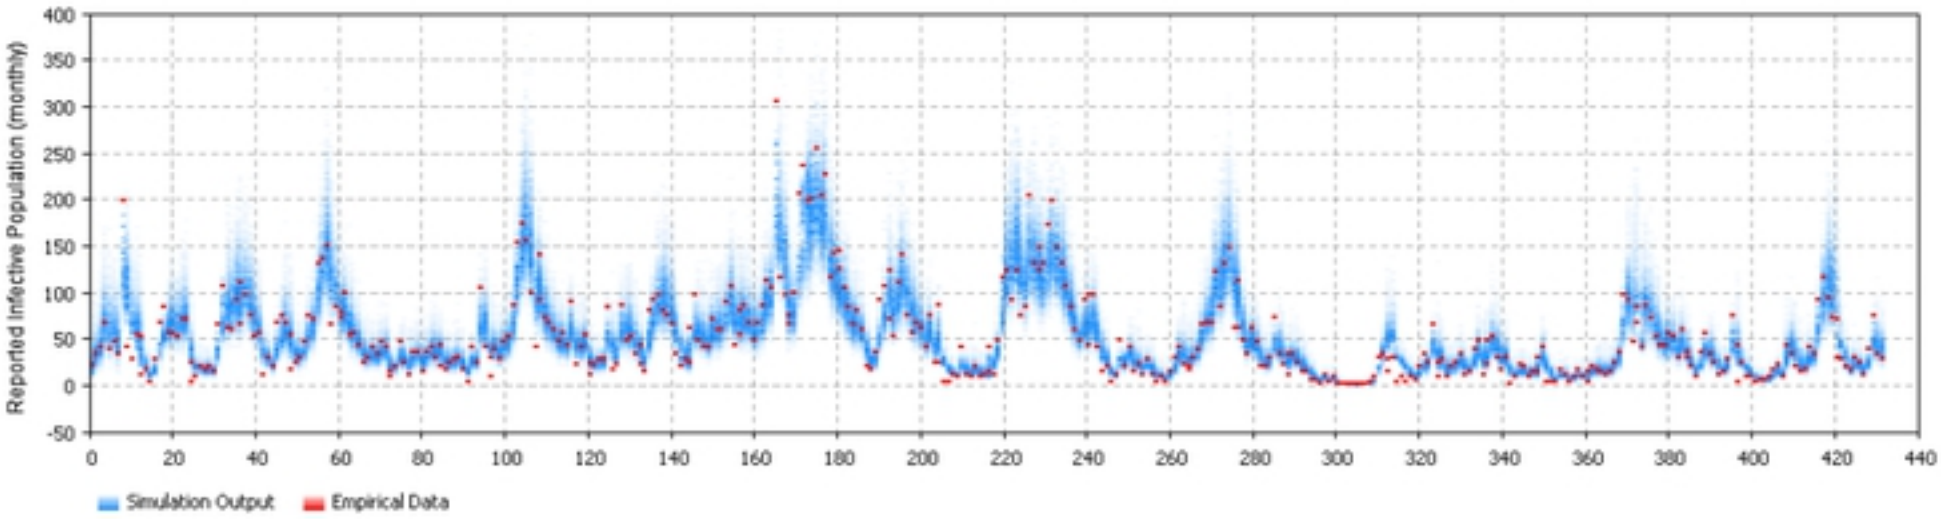
Reported Infective Population (monthly)
Simulation Output
Empirical Data

(a)

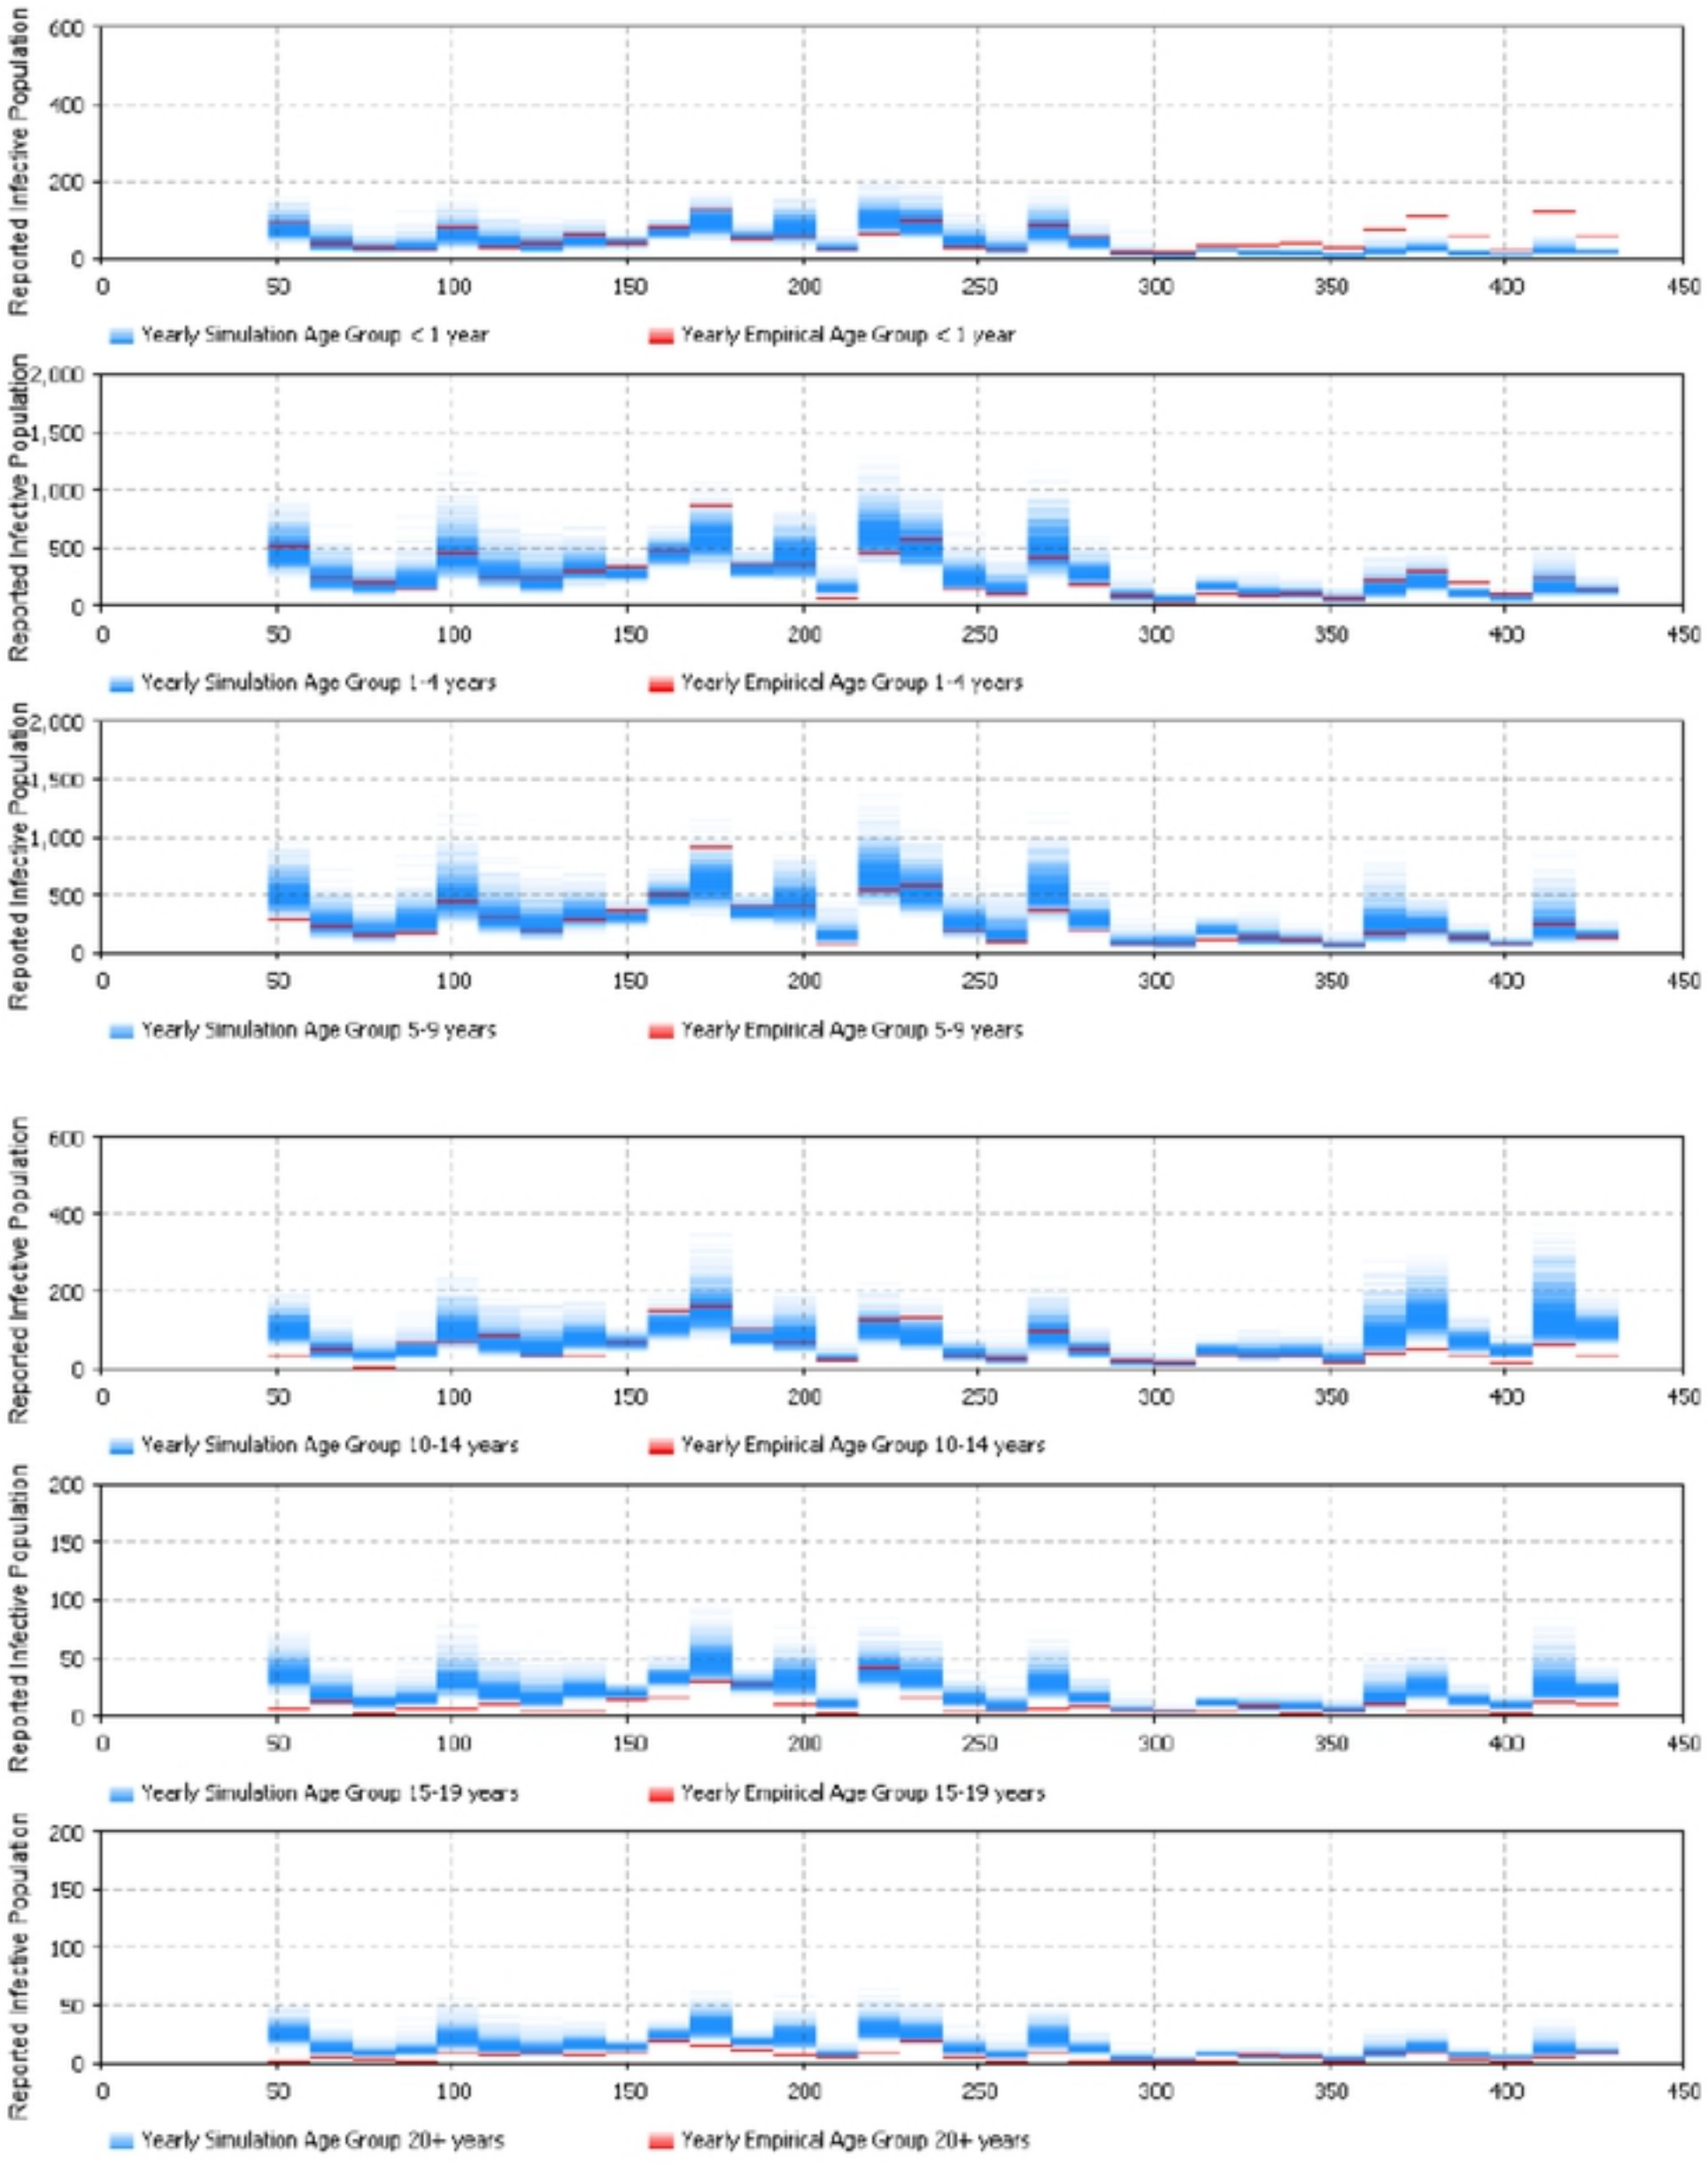
Reported Infective Population
Yearly Simulation Age Group < 1 year
Yearly Empirical Age Group < 1 year
Yearly Simulation Age Group 1-4 years
Yearly Empirical Age Group 1-4 years
Yearly Simulation Age Group 5-9 years
Yearly Empirical Age Group 5-9 years
Yearly Simulation Age Group 10-14 years
Yearly Empirical Age Group 10-14 years
Yearly Simulation Age Group 15-19 years
Yearly Empirical Age Group 15-19 years
Yearly Simulation Age Group 20+ years
Yearly Empirical Age Group 20+ years

(b)

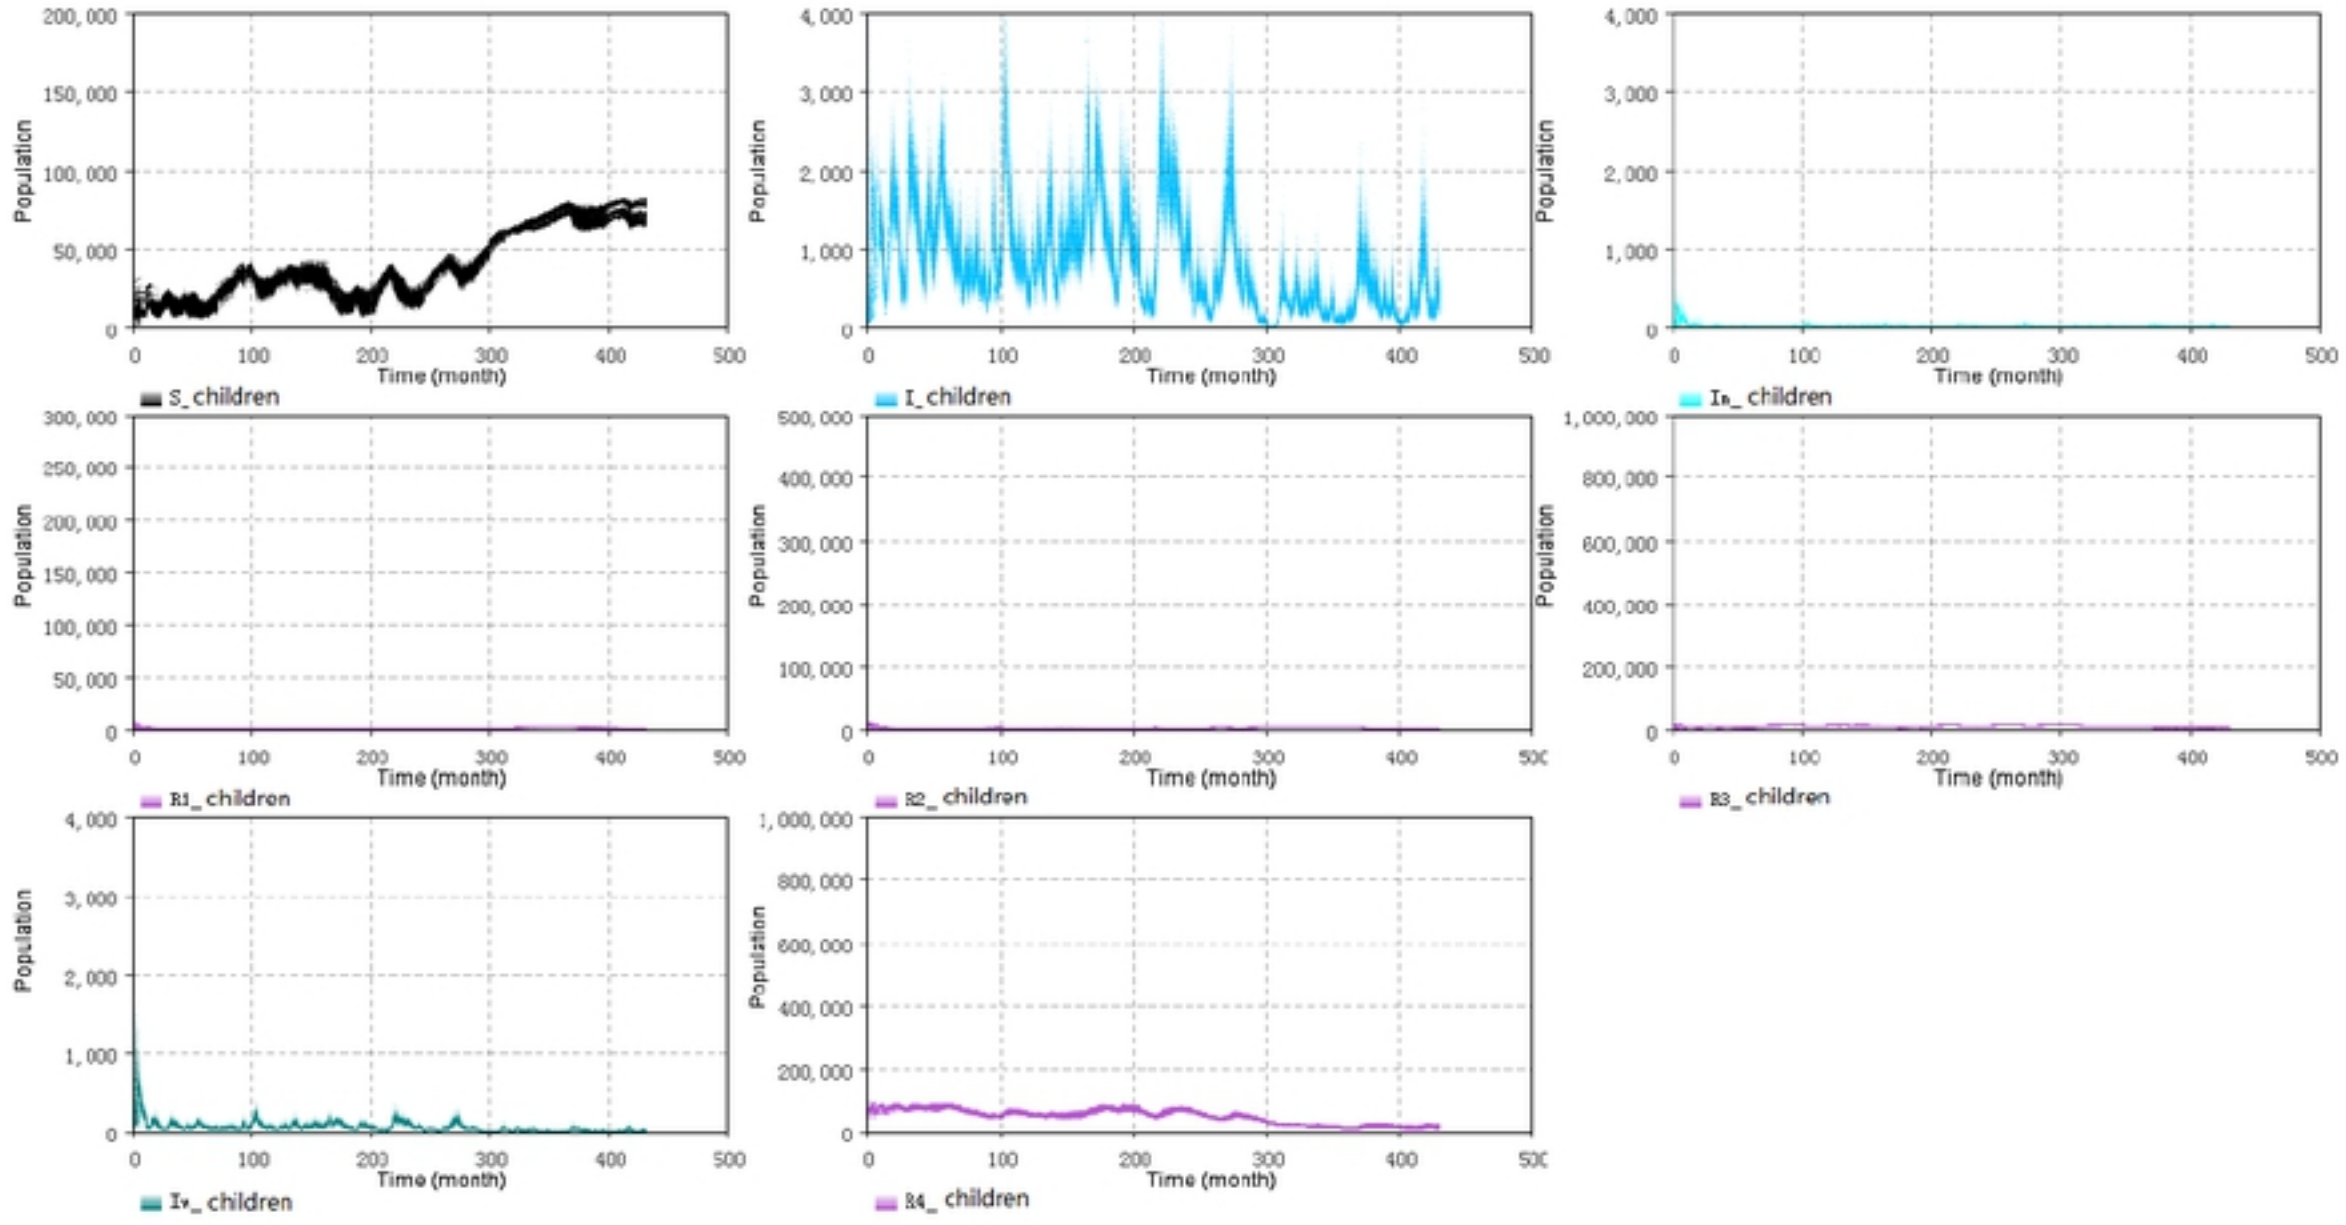
Population
Time (month)
S_children
I_children
Im_children
R1_children
R2_children
R3_children
Iw_children
R4_children

(a)

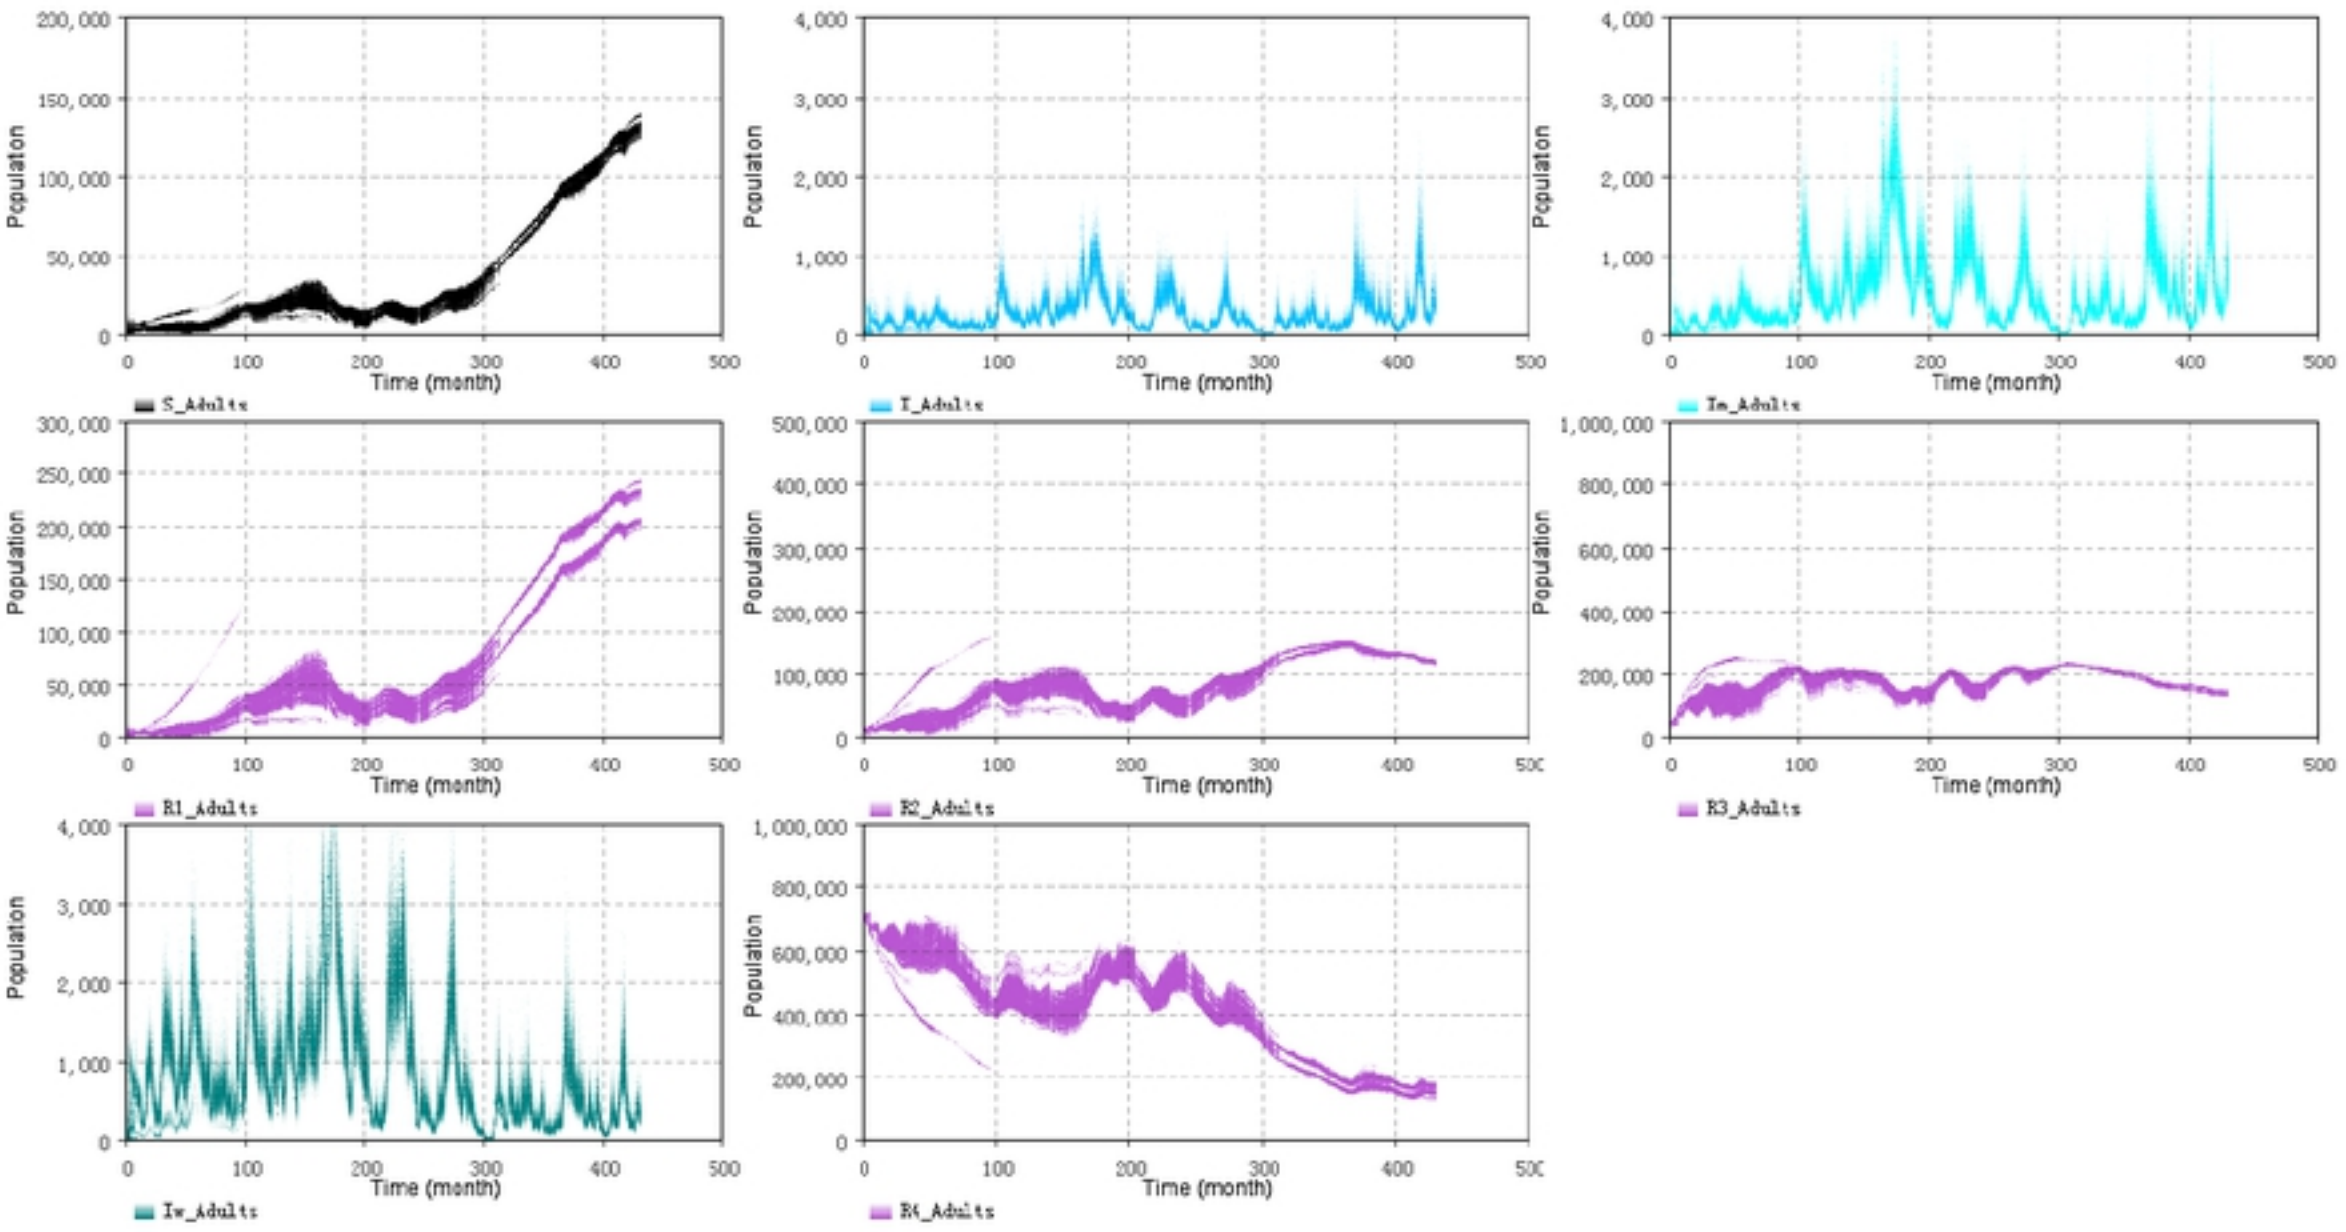
Population
Time (month)
S_Adults
I_Adults
Im_Adults
R1_Adults
R2_Adults
R3_Adults
Iw_Adults
R4_Adults

(b)

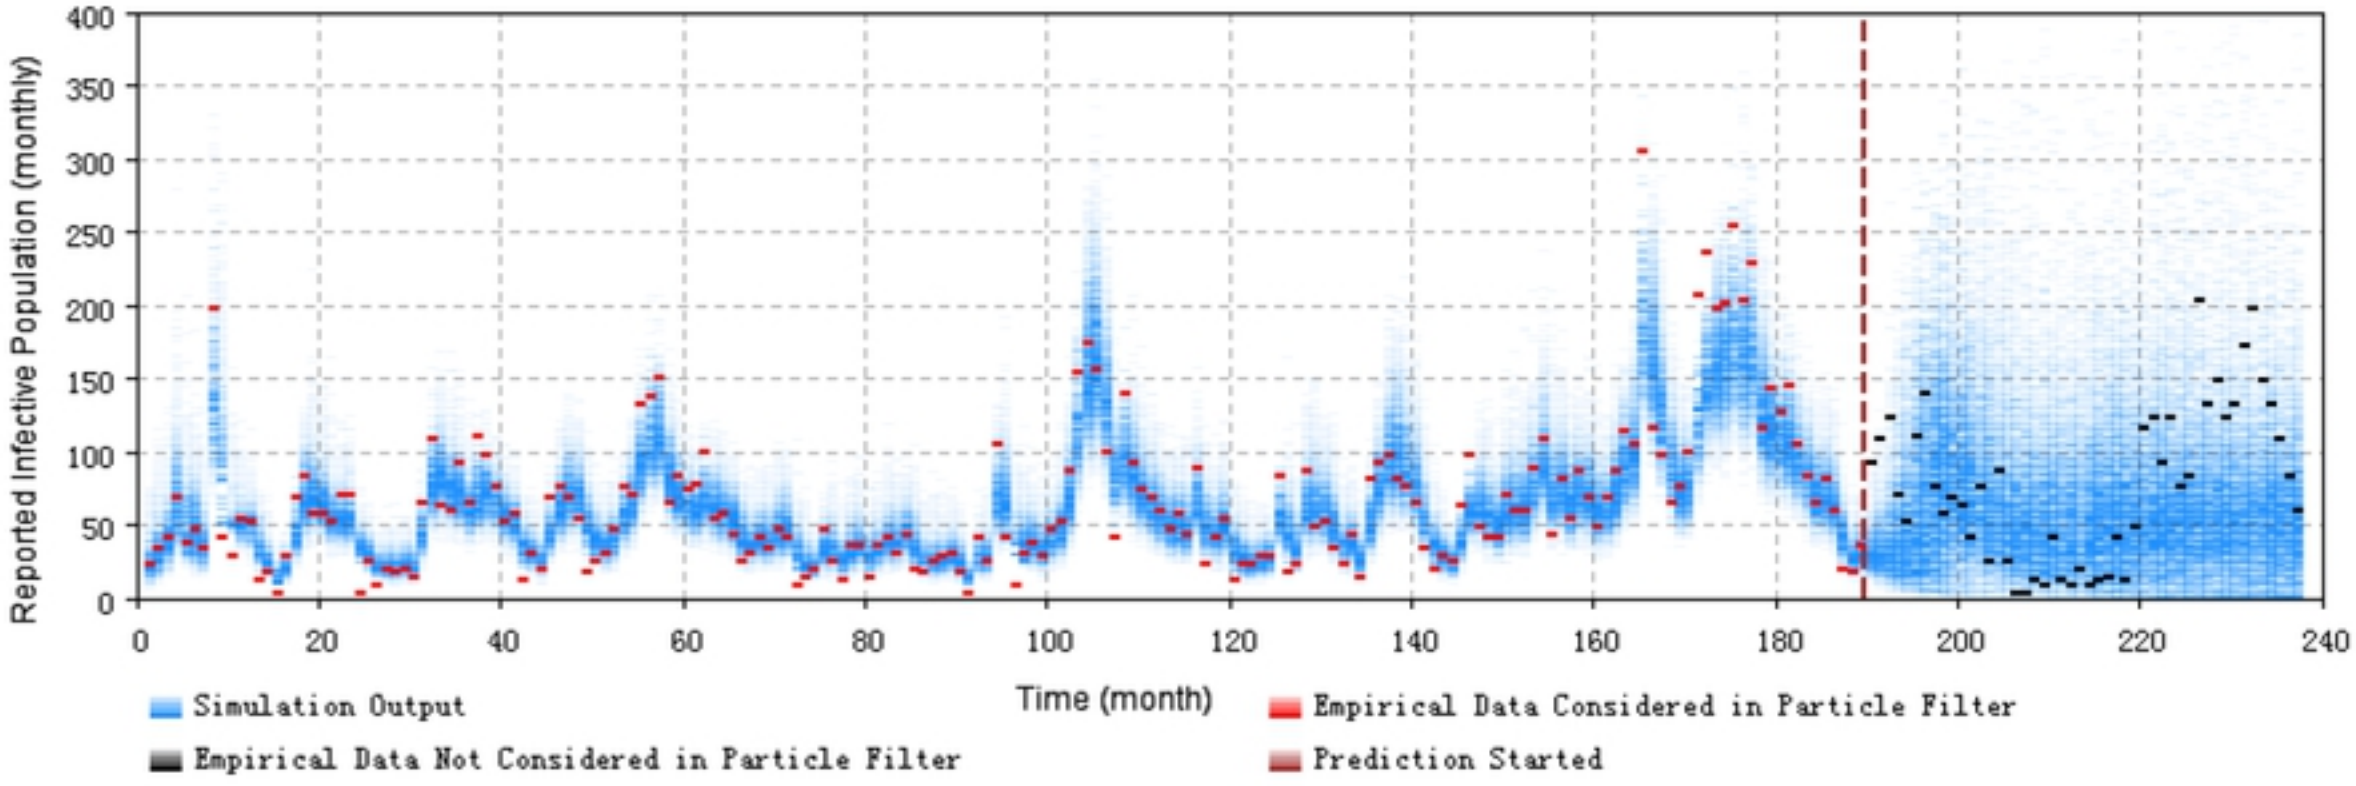
Reported Infective Population (monthly)
Time (month)
Simulation Output
Empirical Data Considered in Particle Filter
Empirical Data Not Considered in Particle Filter
Prediction Started

(a)

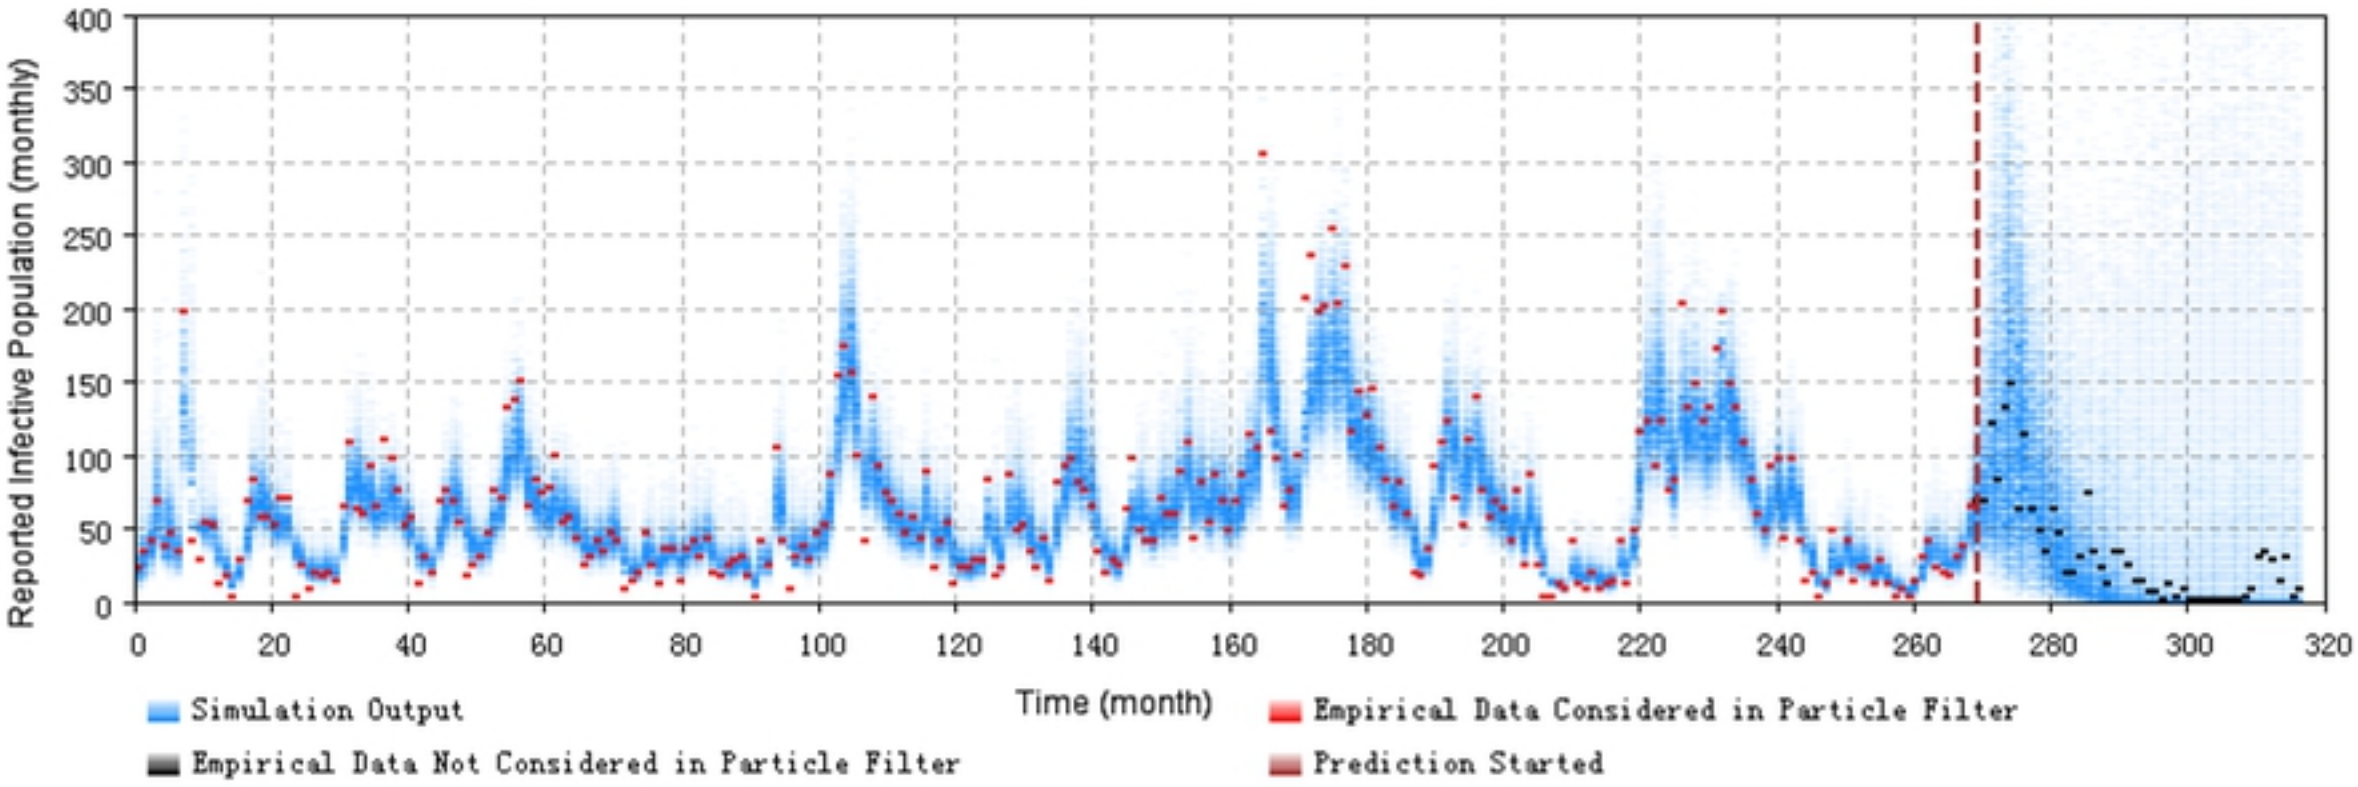
Reported Infective Population (monthly)
Time (month)
Simulation Output
Empirical Data Considered in Particle Filter
Empirical Data Not Considered in Particle Filter
Prediction Started

(b)

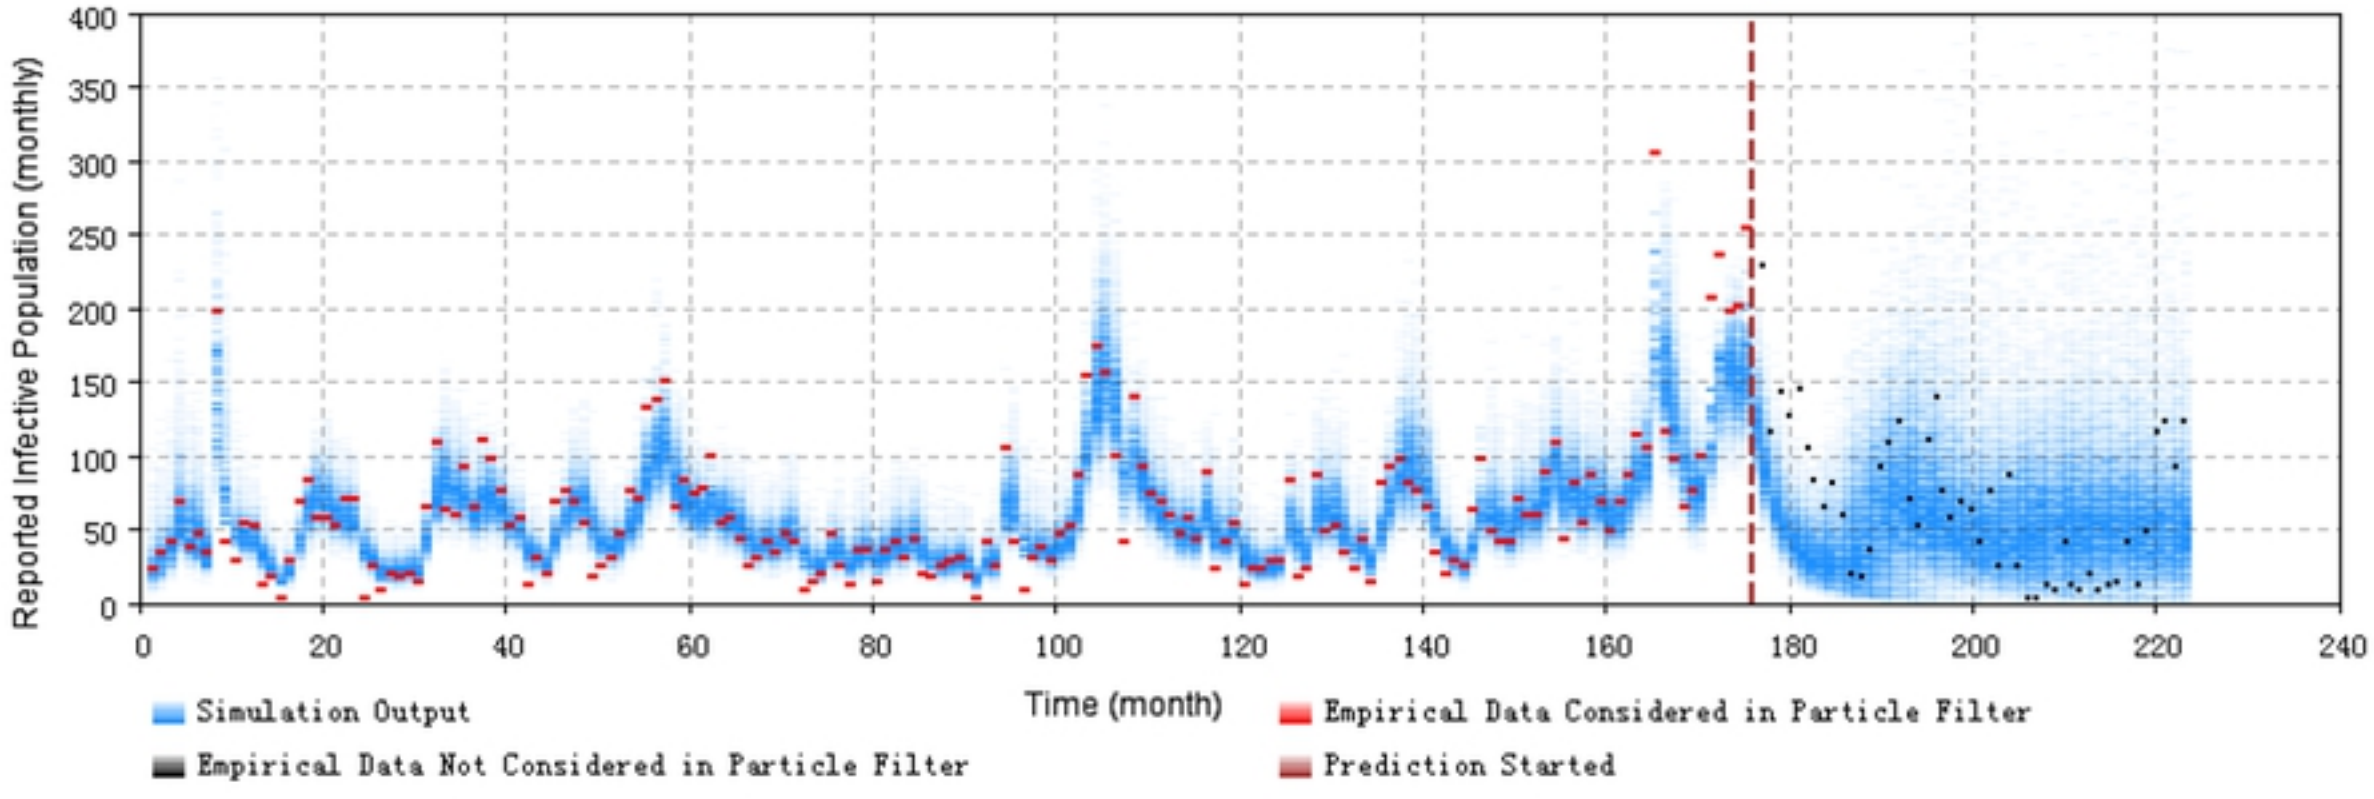
Reported Infective Population (monthly)
Time (month)
Simulation Output
Empirical Data Considered in Particle Filter
Empirical Data Not Considered in Particle Filter
Prediction Started

(a)

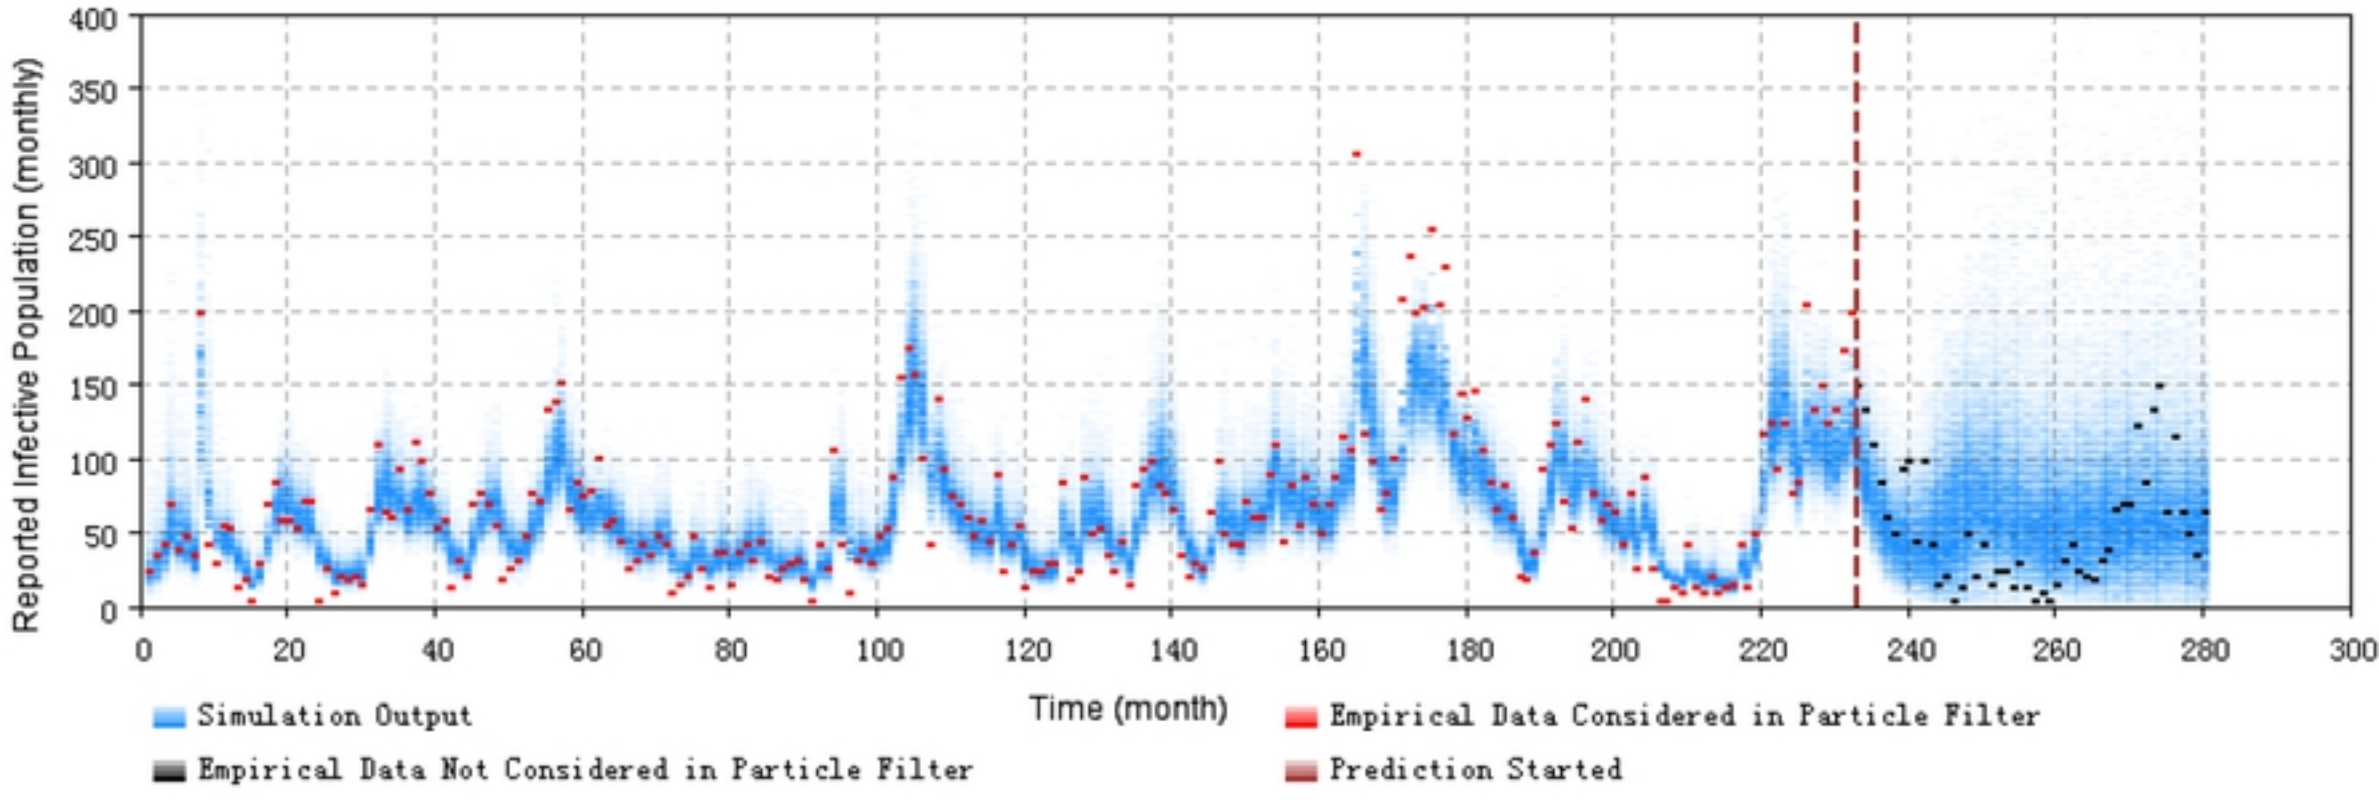
Reported Infective Population (monthly)
Time (month)
Simulation Output
Empirical Data Considered in Particle Filter
Empirical Data Not Considered in Particle Filter
Prediction Started

(b)

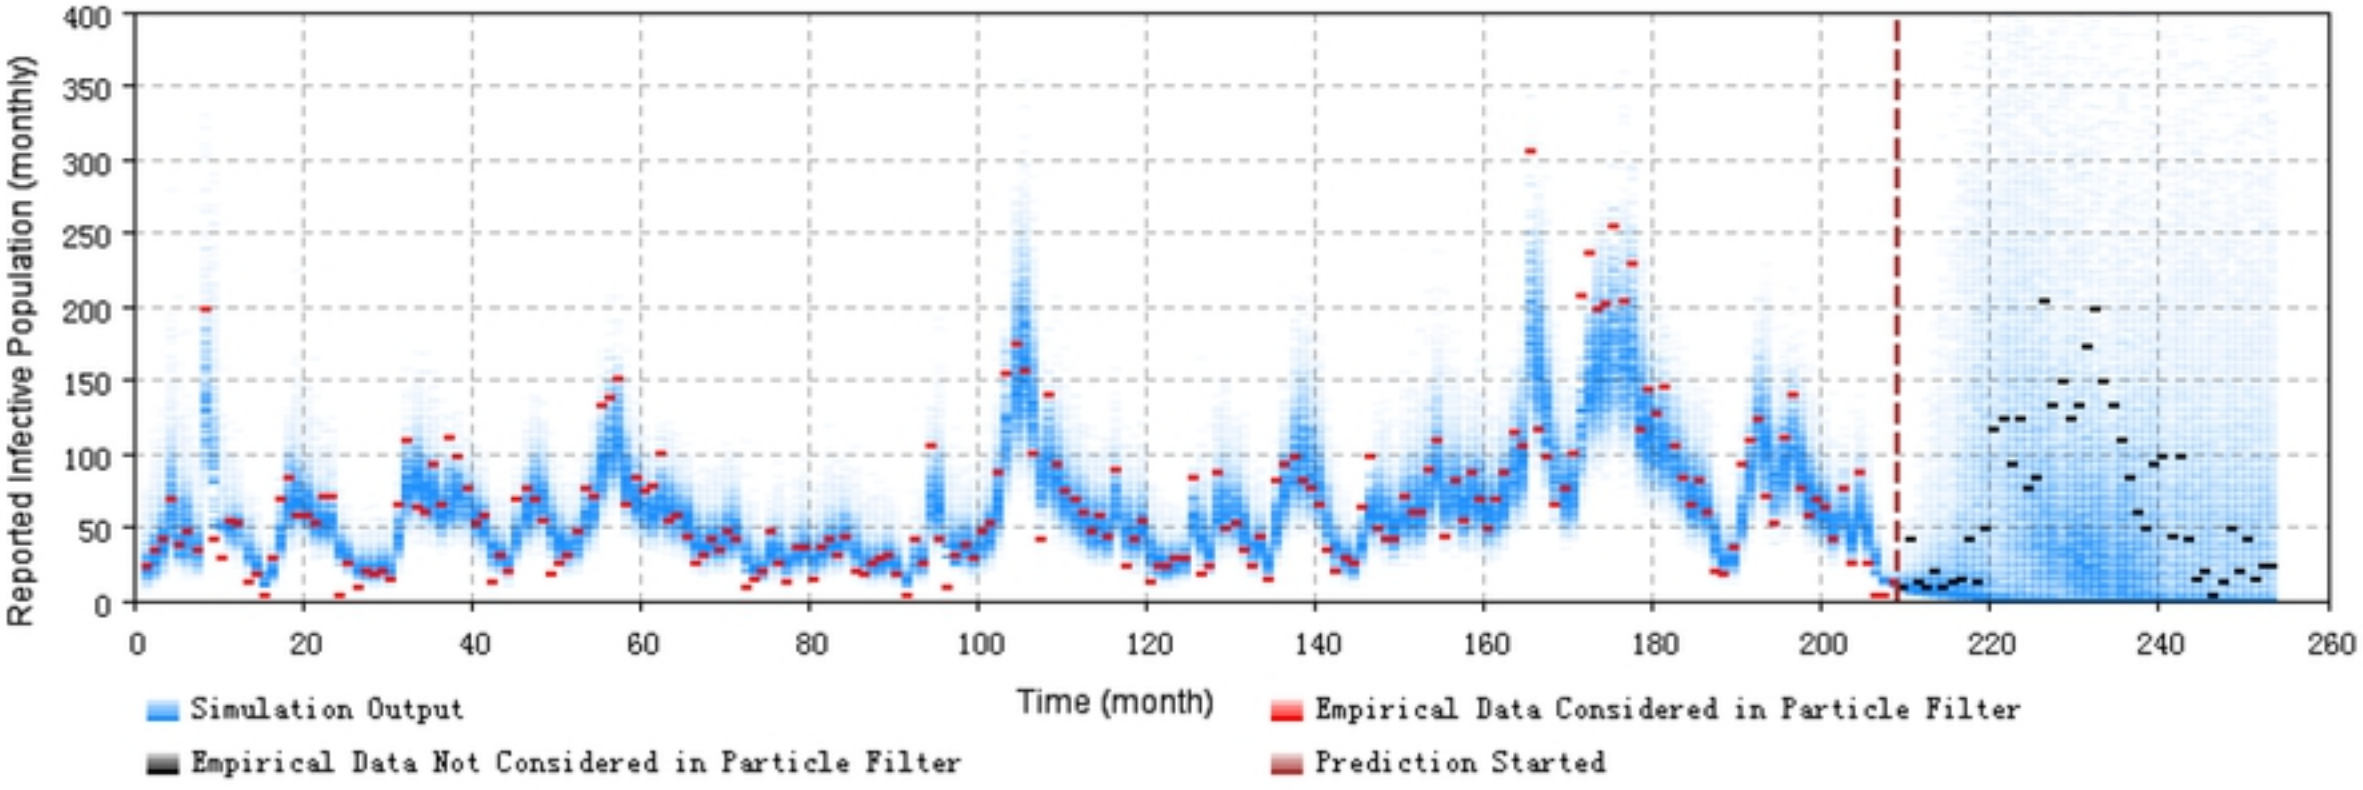
Reported Infective Population (monthly)
Time (month)
Simulation Output
Empirical Data Considered in Particle Filter
Empirical Data Not Considered in Particle Filter
Prediction Started

(a)

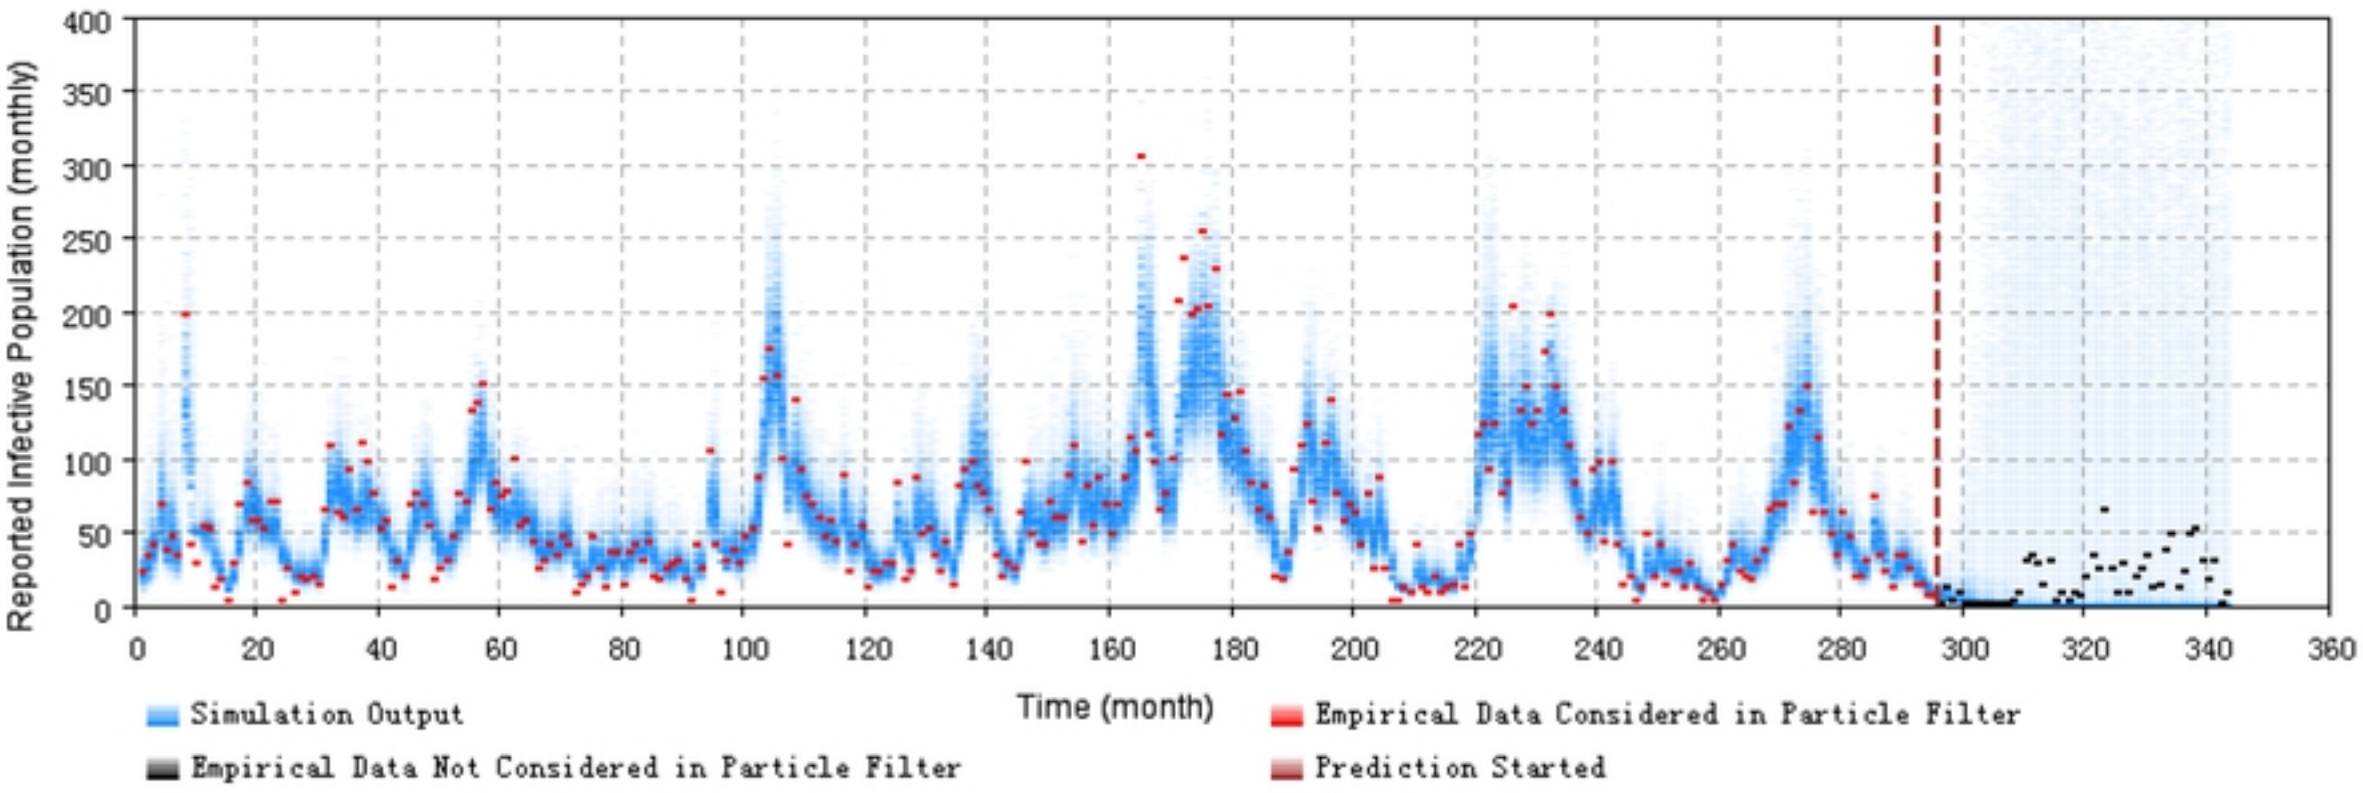
Reported Infective Population (monthly)
Time (month)
Simulation Output
Empirical Data Considered in Particle Filter
Empirical Data Not Considered in Particle Filter
Prediction Started

(b)

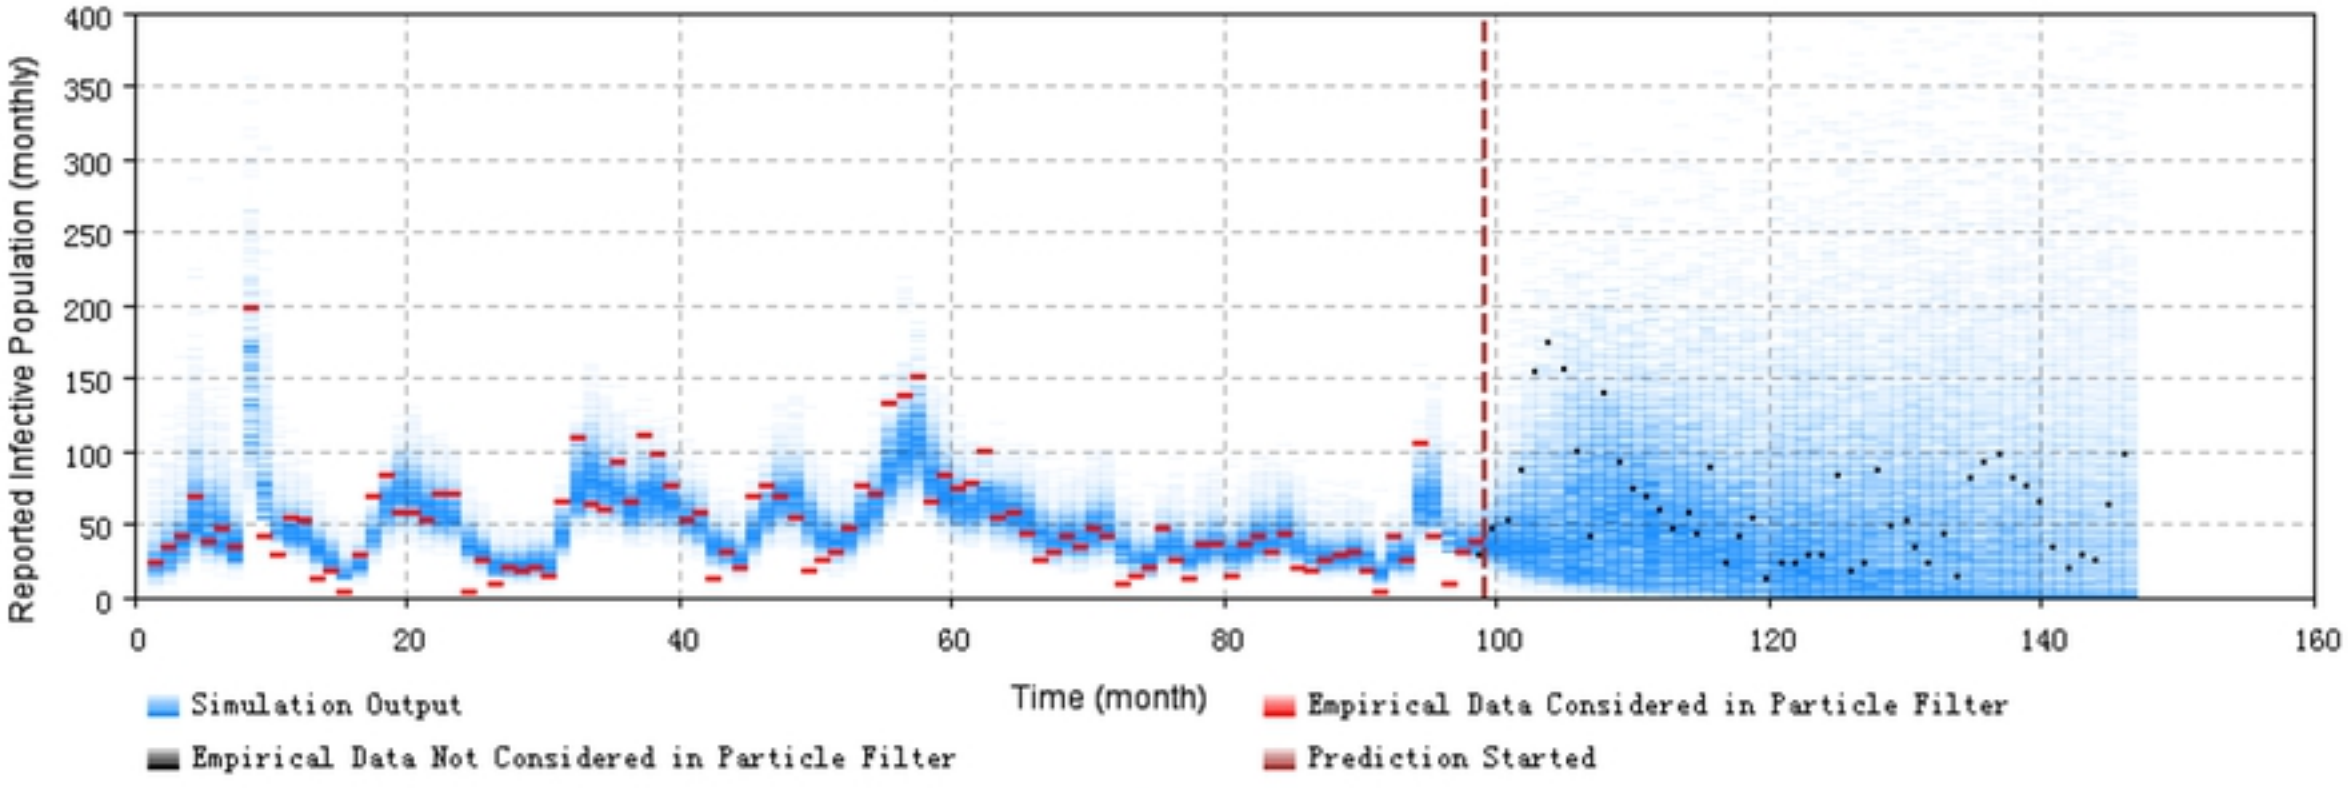
Reported Infective Population (monthly)
Time (month)
Simulation Output
Empirical Data Considered in Particle Filter
Empirical Data Not Considered in Particle Filter
Prediction Started

(a)

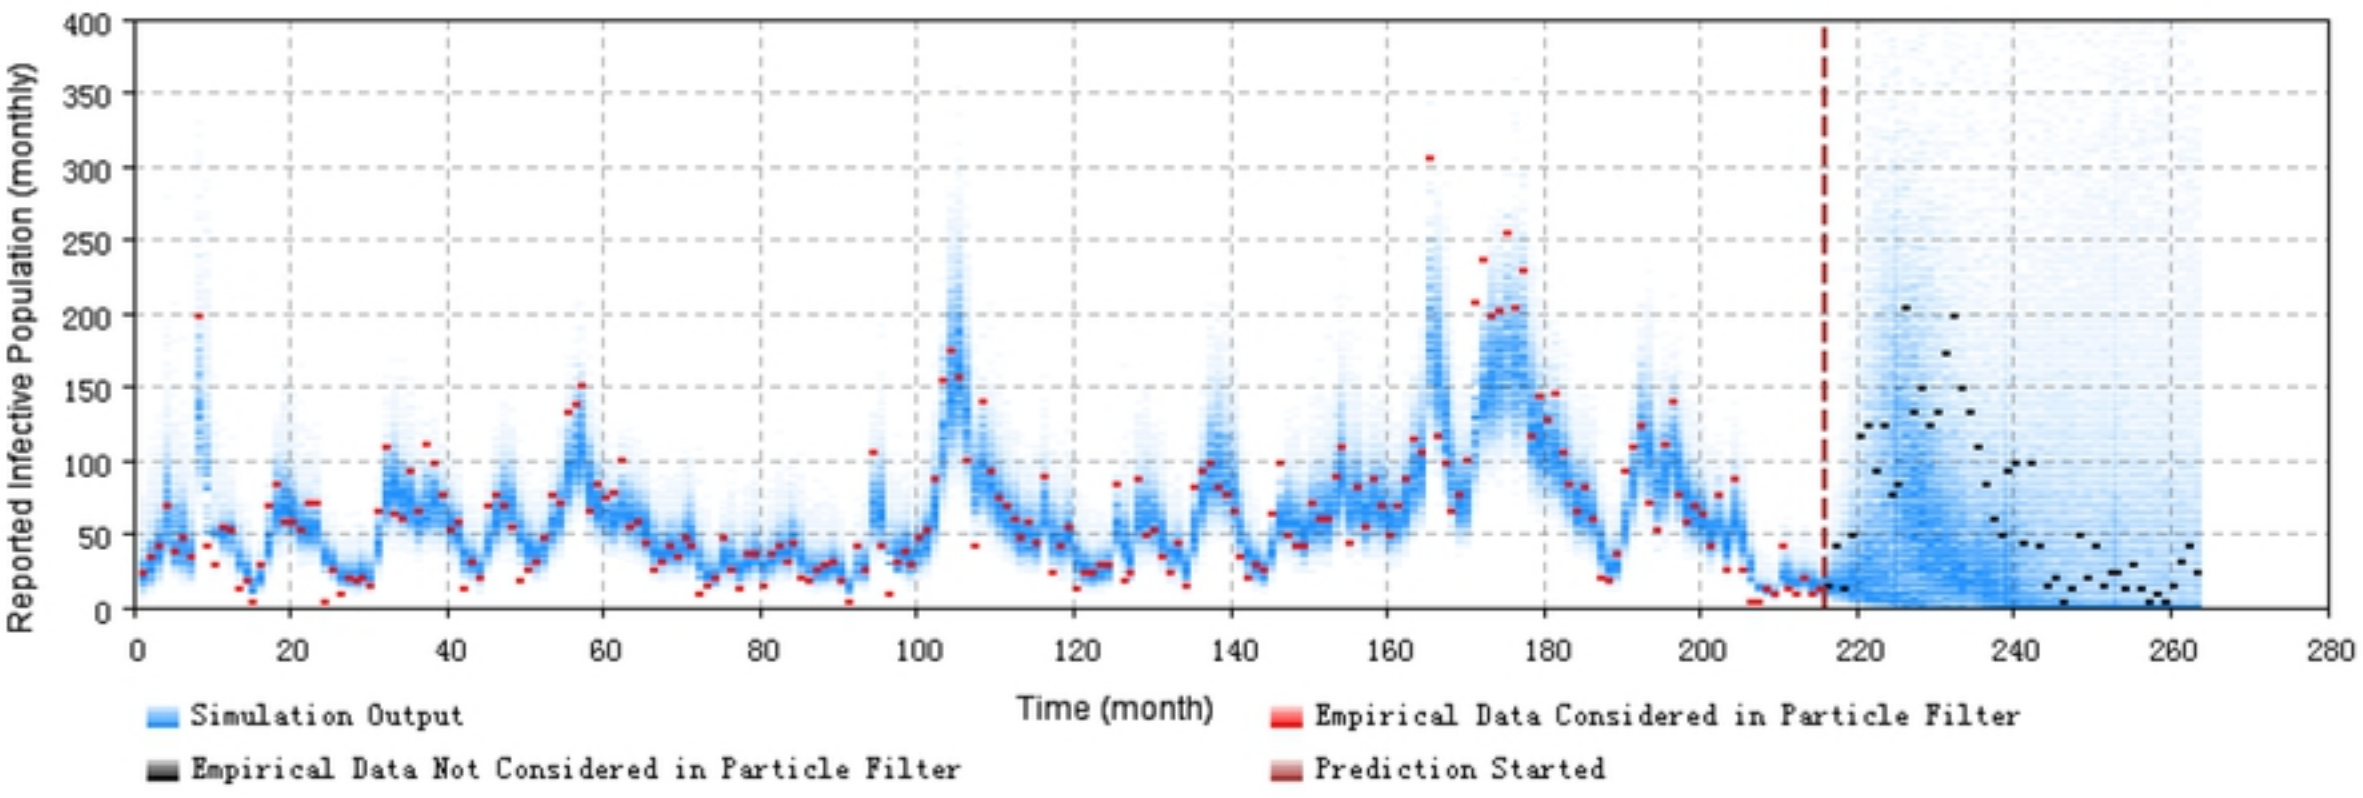
Reported Infective Population (monthly)
Time (month)
Simulation Output
Empirical Data Considered in Particle Filter
Empirical Data Not Considered in Particle Filter
Prediction Started

(b)

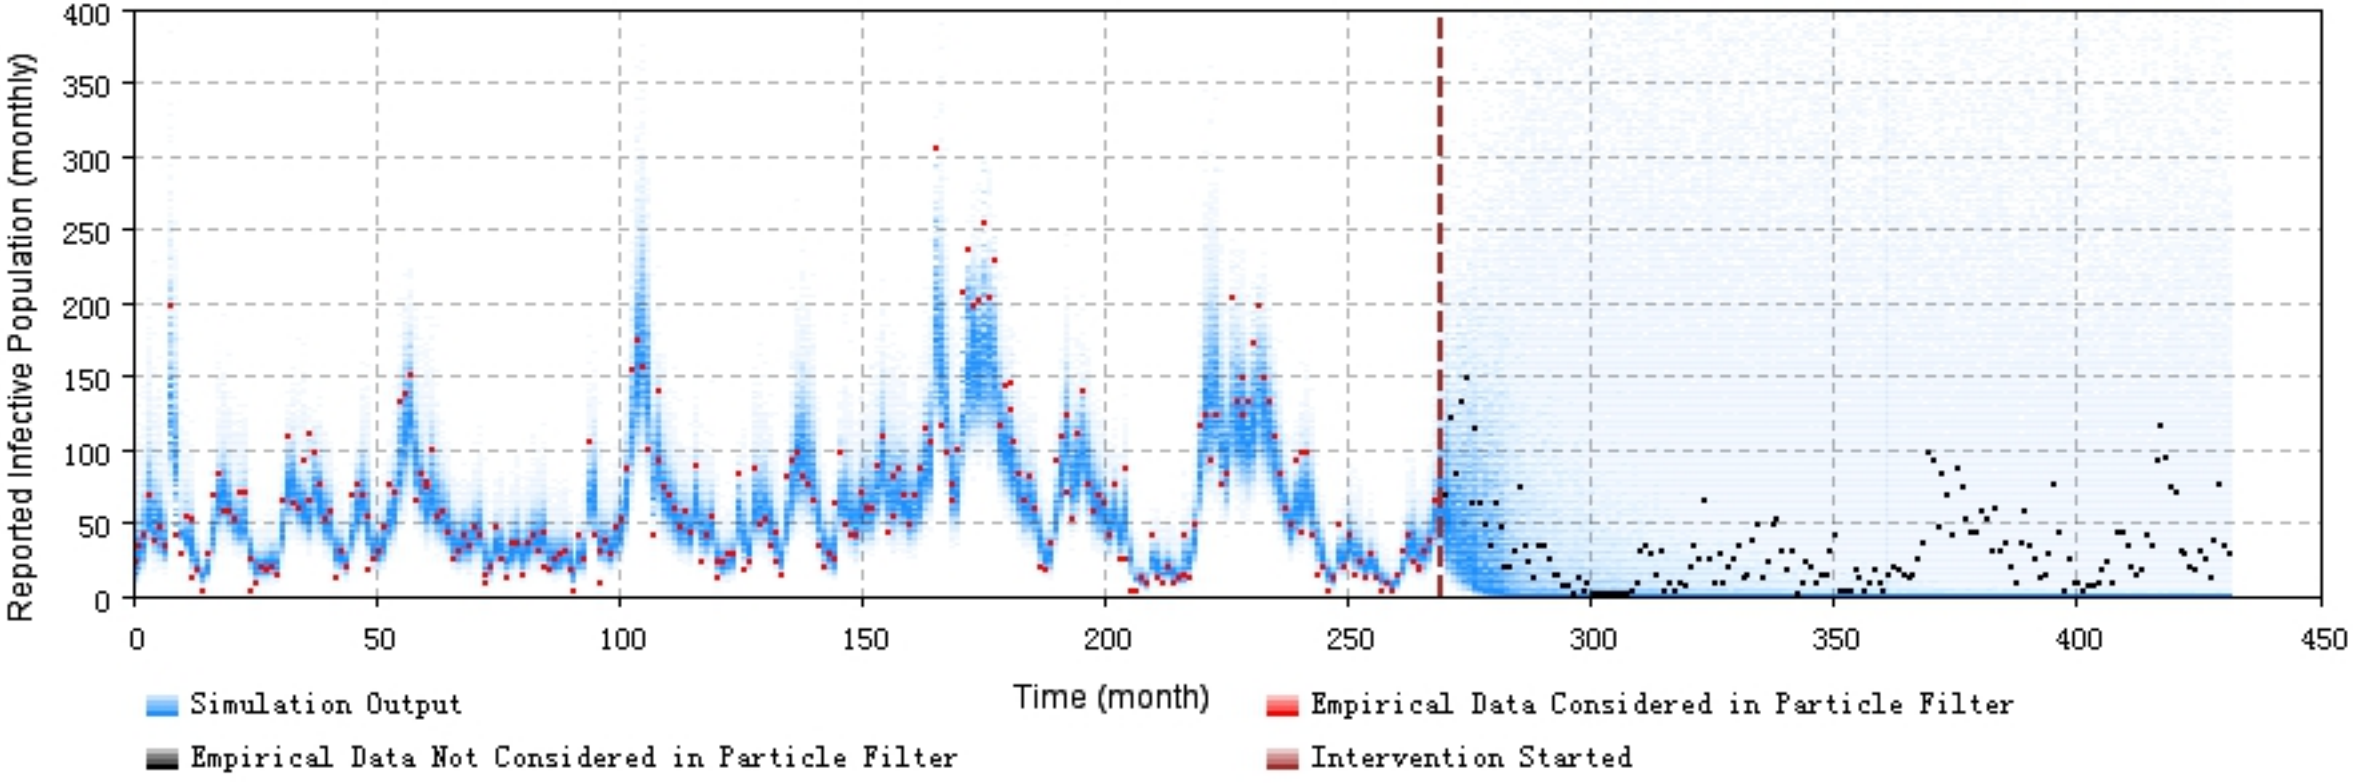
Reported Infective Population (monthly)
400
350
300
250
200
150
100
50
0
0
50
100
150
200
250
300
350
400
450
Time (month)
Simulation Output
Empirical Data Considered in Particle Filter
Empirical Data Not Considered in Particle Filter
Intervention Started

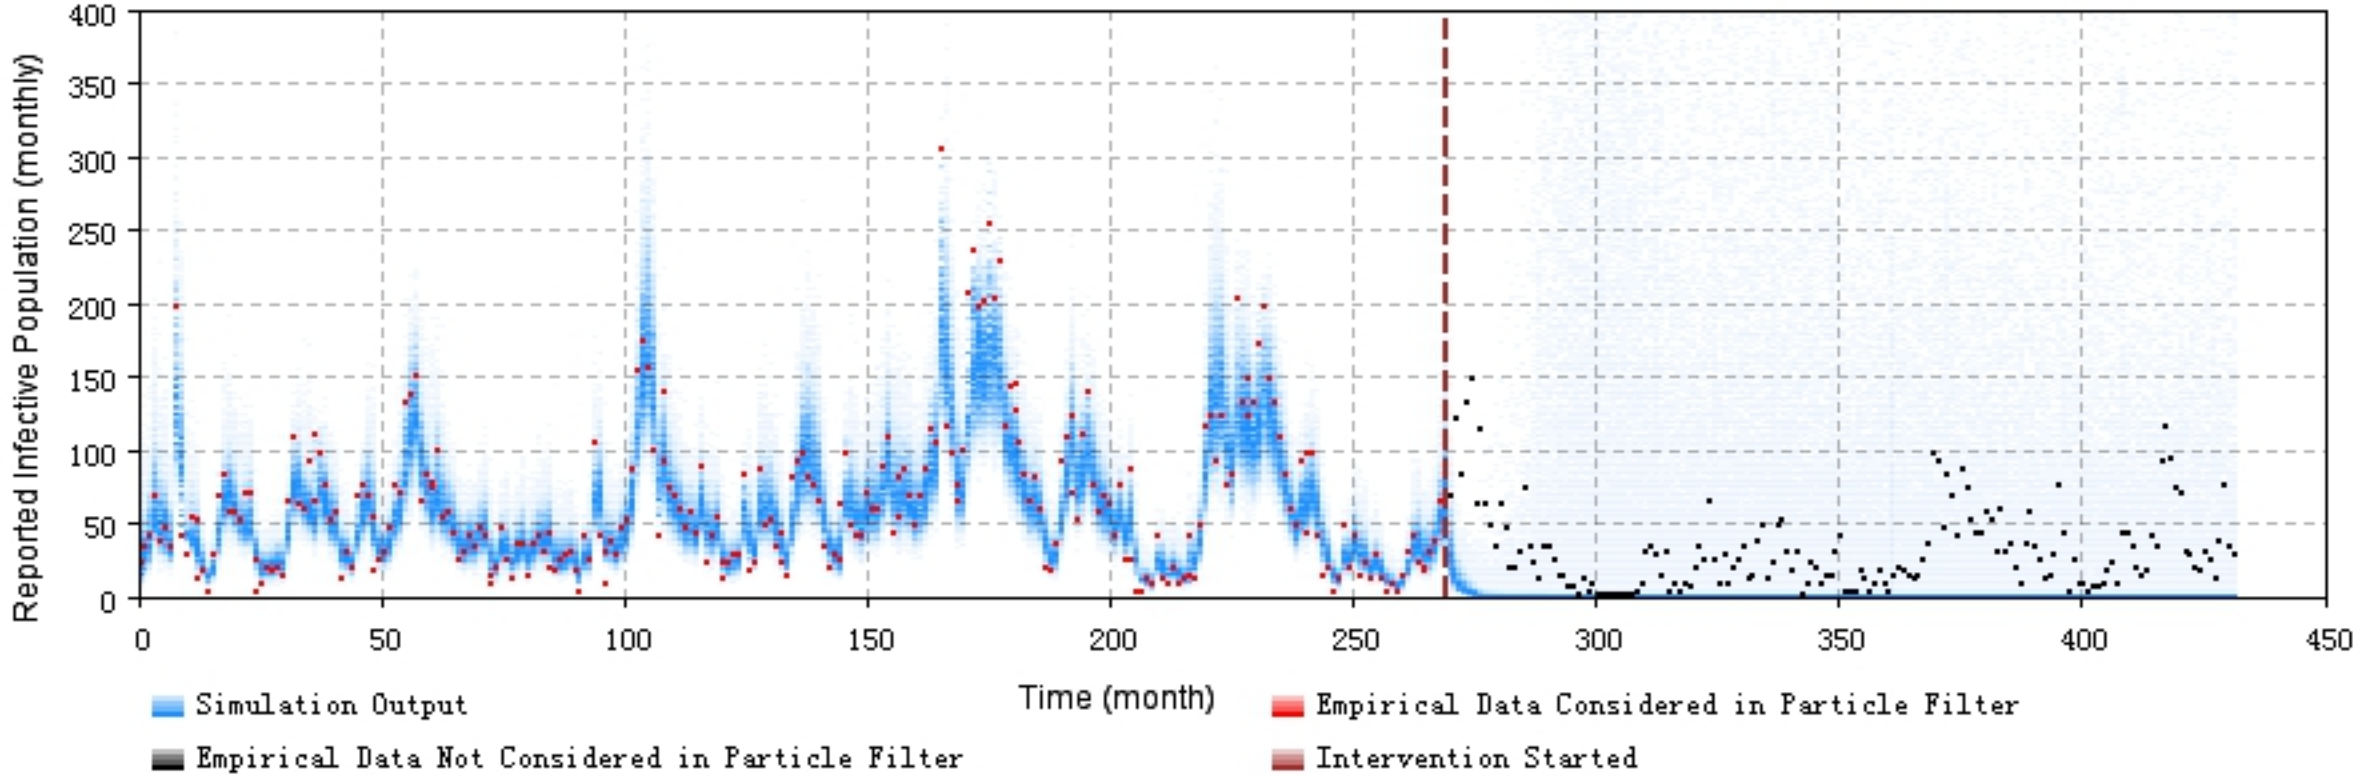
Reported Infective Population (monthly)
400
350
300
250
200
150
100
50
0
0
50
100
150
200
250
300
350
400
450
Time (month)
Simulation Output
Empirical Data Considered in Particle Filter
Empirical Data Not Considered in Particle Filter
Intervention Started

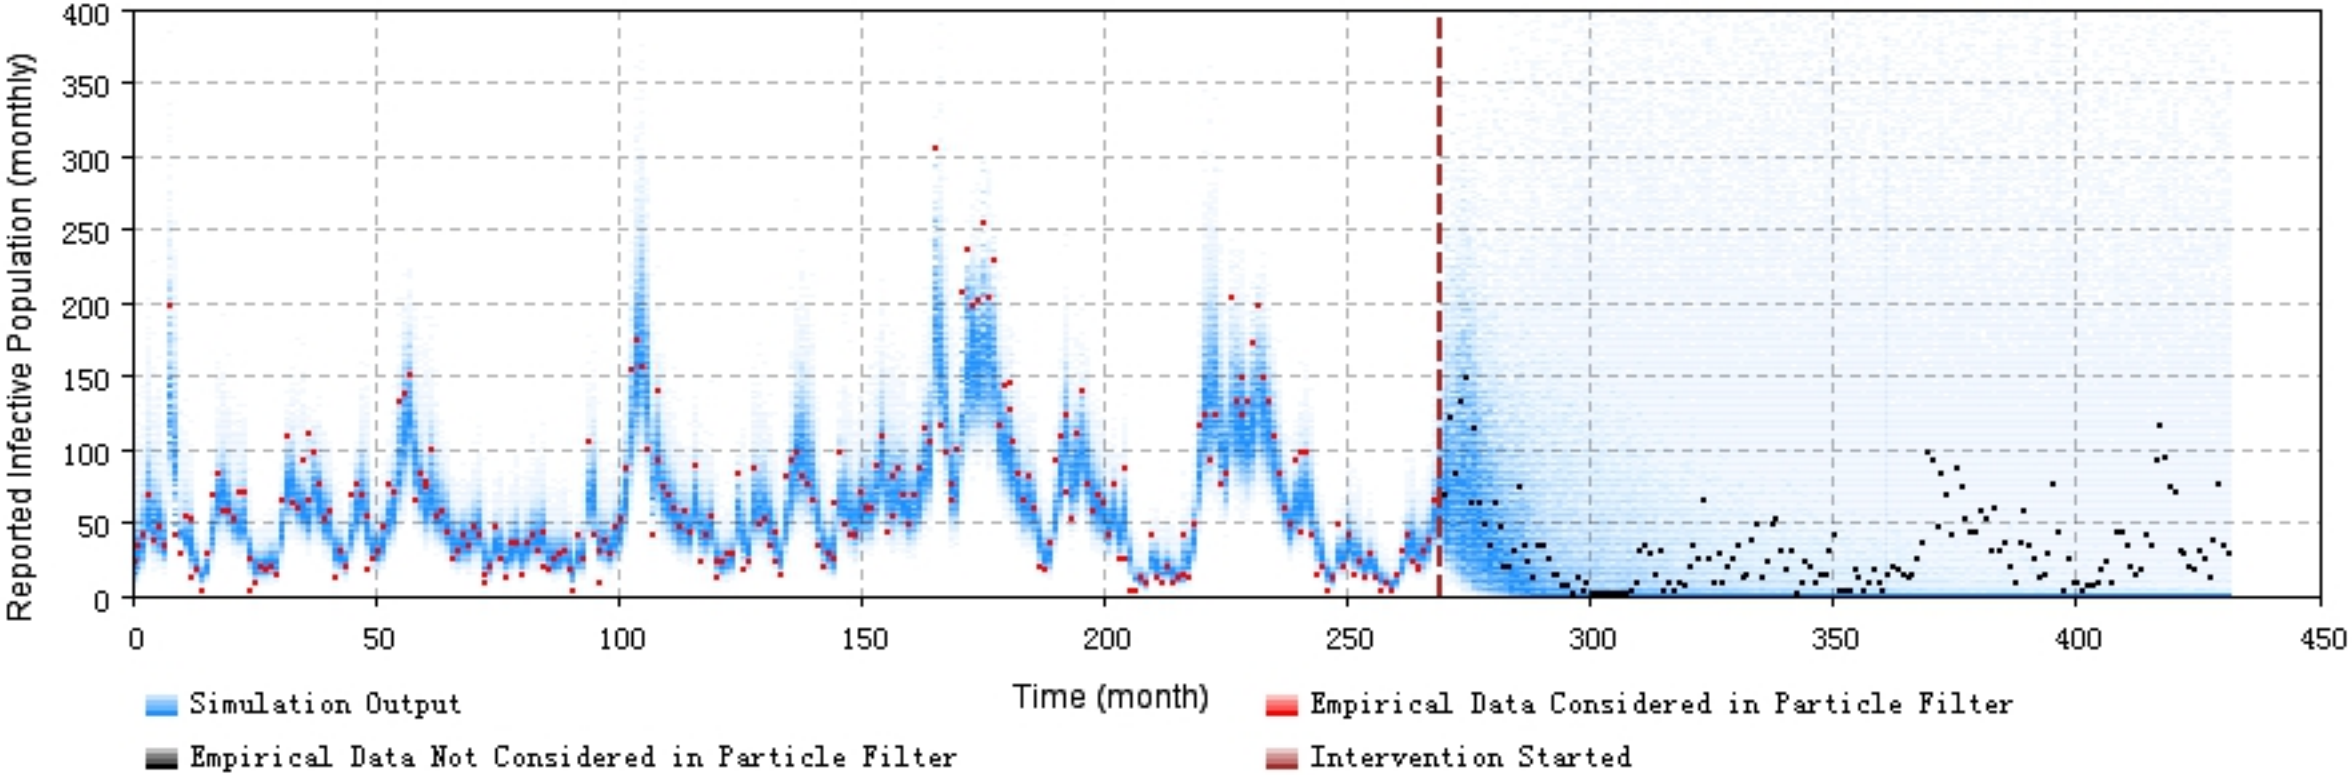
Reported Infective Population (monthly)
400
350
300
250
200
150
100
50
0
0
50
100
150
200
250
300
350
400
450
Time (month)
Simulation Output
Empirical Data Considered in Particle Filter
Empirical Data Not Considered in Particle Filter
Intervention Started

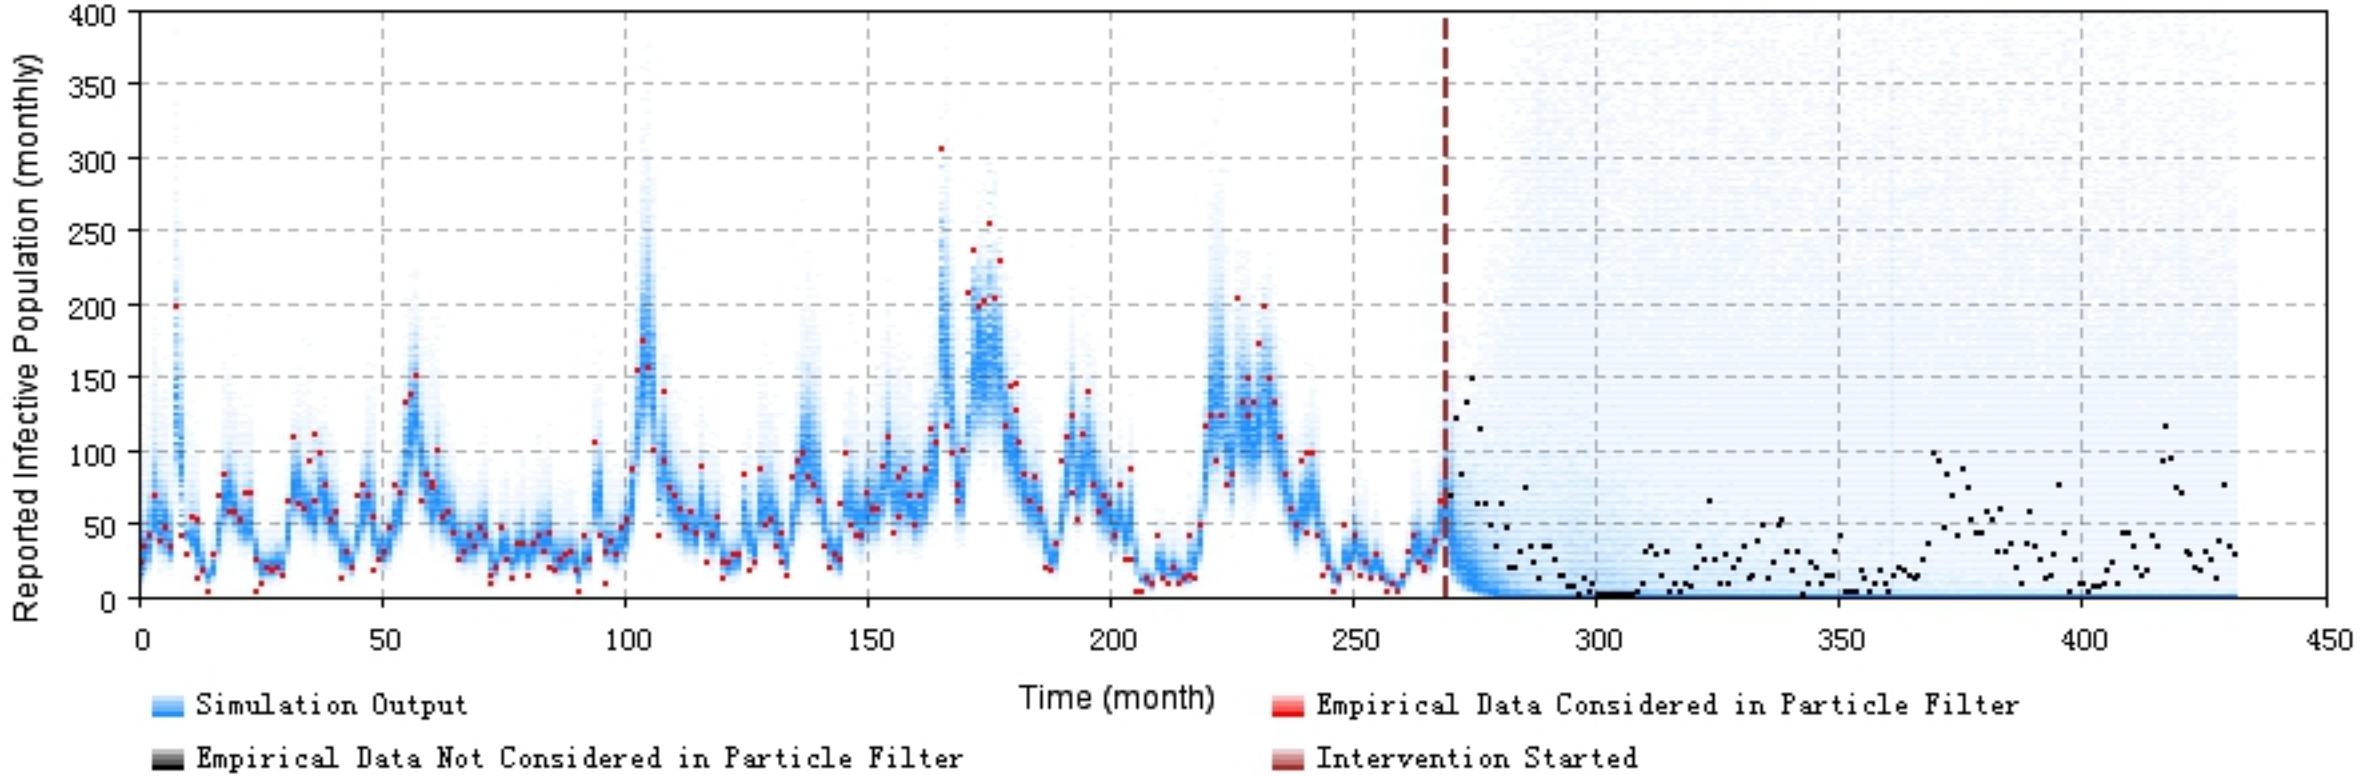
Reported Infective Population (monthly)
400
350
300
250
200
150
100
50
0
0
50
100
150
200
250
300
350
400
450
Time (month)
Simulation Output
Empirical Data Considered in Particle Filter
Empirical Data Not Considered in Particle Filter
Intervention Started