

# Bayesian Model Averaging to Account for Model Uncertainty in Estimates of a Vaccine's Effectiveness

Carlos R Oliveira<sup>1,2</sup>, Eugene D Shapiro<sup>1,3</sup>, Daniel M Weinberger<sup>3</sup>

<sup>1</sup>Department of Pediatrics, Yale University School of Medicine, New Haven, CT, USA; <sup>2</sup>Department of Biostatistics, Yale University School of Public Health, New Haven, CT, USA; <sup>3</sup>Department of Epidemiology of Microbial Diseases, Yale University School of Public Health, New Haven, CT, USA

Correspondence: Carlos R Oliveira, Yale University School of Medicine, 464 Congress Avenue, Suite 204, New Haven, CT, 06520, USA, Tel +1 203 785 5474, Email carlos.oliveira@yale.edu

**Purpose:** Vaccine effectiveness (VE) studies are often conducted after the introduction of new vaccines to ensure they provide protection in real-world settings. Control of confounding is often needed during the analyses, which is most efficiently done through multivariable modeling. When many confounders are being considered, it can be challenging to know which variables need to be included in the final model. We propose an intuitive Bayesian model averaging (BMA) framework for this task.

**Patients and Methods:** Data were used from a matched case-control study that aimed to assess the effectiveness of the Lyme vaccine post-licensure. Cases were residents of Connecticut, 15–70 years of age with confirmed Lyme disease. Up to 2 healthy controls were matched to each case subject by age. All participants were interviewed, and medical records were reviewed to ascertain immunization history and evaluate potential confounders. BMA was used to systematically search for potential models and calculate the weighted average VE estimate from the top subset of models. The performance of BMA was compared to three traditional single-best-model-selection methods: two-stage selection, stepwise elimination, and the leaps and bounds algorithm.

**Results:** The analysis included 358 cases and 554 matched controls. VE ranged between 56% and 73% and 95% confidence intervals crossed zero in <5% of all candidate models. Averaging across the top 15 models, the BMA VE was 69% (95% CI: 18–88%). The two-stage, stepwise, and leaps and bounds algorithm yielded VE of 71% (95% CI: 21–90%), 73% (95% CI: 26–90%), and 74% (95% CI: 27–91%), respectively.

**Conclusion:** This paper highlights how the BMA framework can be used to generate transparent and robust estimates of VE. The BMA-derived VE and confidence intervals were similar to those estimated using traditional methods. However, by incorporating model uncertainty into the parameter estimation, BMA can lend additional rigor and credibility to a well-designed study.

**Keywords:** Bayesian model averaging, model uncertainty, vaccine effectiveness, Lyme vaccine

## Plain Language Summary

How well a vaccine works in real-world settings depends on a lot of different factors. When conducting a study to assess how effective a vaccine is in routine clinical settings, researchers often gather information on a large number of possible influencing factors. When analyzing the data of these studies, it can be hard to know what variables to include in the analysis and what variables to ignore. The statistical efficiency of the study, or its ability to detect an effect with a minimum number of subjects, will be significantly reduced when unimportant variables are included in the analysis. Different methods have been suggested to help figure out which variables should be included in the analytic model. However, when you choose just one model to describe your data, readers will always wonder what effects the variables you did not include might have had on your results. In this article, we describe an easy-to-use algorithm that helps investigators systematically search for suitable statistical models across every possible combination of control variables in their data and test how well their final estimates hold up across different variable-selection assumptions. We illustrate how this method works by re-analyzing data from a recently published study assessing the effectiveness of the Lyme vaccine and comparing its performance relative to that of traditional variable-selection methods. We show how this method of analysis can add credence to a well-designed study by providing a more complete picture than methods that only look at one model at a time.

## Introduction

Observational studies are the most appropriate way to measure a vaccine's effectiveness (VE) post-licensure.<sup>1</sup> When the vaccine-preventable disease is rare, or when there is a prolonged latency between immunization and disease, the case-control design can

be particularly favorable, as smaller sample sizes are needed and subjects do not need to be followed for a prolonged period of time.<sup>2</sup> As an observational method, however, the case-control design is susceptible to the effects of confounding.<sup>3</sup> The ideal time to deal with known confounders is during the design phase of the study. However, assessment for and control of confounding is often also needed during the analyses of the study, which is most efficiently done using multivariate modeling.<sup>4</sup>

Multivariable modeling's primary goal is to determine whether the observed effect of vaccination on disease is confounded by one of the measured variables, and thus, requires some form of adjustment. The simplest approach assumes that all collected variables are important and uses the "full model", which includes all potential confounders as independent variables. However, the statistical efficiency of the study will be significantly reduced when unimportant or non-confounding variables are included in a model.<sup>5</sup> More importantly, certain variable combinations may introduce bias due to collider stratification,<sup>6</sup> which can lead to erroneous conclusions regarding statistical significance.<sup>7,8</sup>

To address the problems of overfitting, several strategies have been developed that help investigators select the "best" model among a set of candidate models and estimate the vaccine's effectiveness with the assumption that the chosen model is the best for their data.<sup>9</sup> While these methods are easy to implement, they can have several limitations. For example, restricting the analysis to a single model can lack transparency, and may lead to uncertainty about the potential effects of the variables ultimately not included in the final model.

An alternative approach to selecting a single model is Bayesian model averaging (BMA). Several studies have shown that in the presence of model uncertainty, BMA can provide better inferences than single-model selection methods.<sup>10,11</sup> However, few have explored the use of BMA in the context of case-control VE studies.<sup>12,13</sup> In this study, we explore the importance of model uncertainty by re-analyzing data from a Lyme VE study using various confounder selection techniques. The theoretical framework for the use of BMA in case-control studies has been previously described.<sup>14-16</sup> Here, we focus on the application of BMA towards the analysis of VE data and compare its performance to that of traditional single-best-model-selection methods.

## Materials and Methods

### Vaccine Effectiveness Data

A vaccine for Lyme disease (LYMERix<sup>TM</sup>) was available for adults from 1998 until 2002. To explore the importance of model uncertainty in VE studies, we re-analyzed data from a matched case-control study that aimed to assess the effectiveness of the Lyme vaccine post-licensure. The methods for the Lyme VE study have been previously described.<sup>17</sup> Briefly, cases were residents of Connecticut, 15–70 years of age, reported to the Department of Public Health as having Lyme disease. Up to 2 controls were identified and matched to each case by age. All participants were interviewed, and medical records were reviewed from primary care providers to ascertain immunization history, evaluate potential confounders, and verify the participant's case-control status. Study definitions are summarized in Table 1.

**Table 1** Study Definitions for Study on the Effectiveness of Lyme Vaccine

Variable	Definition
Definite Lyme disease	Physician diagnosed erythema migrans measuring >5 cm, or a positive laboratory test result for antibodies to <i>Borrelia burgdorferi</i> and objective evidence of either early disseminated or of late Lyme disease (eg, facial nerve palsy, meningitis, carditis, or arthritis).
Case	Residents of Connecticut, 15–70 years of age, diagnosed with first episode of "definite Lyme disease" between 2000 and 2013.
Control	Individuals without signs or symptoms of Lyme disease at focal time, matched to a case by age ( $\pm 5$ years) and geographic area (same telephone exchange).
Focal time	The common date between cases and controls. For controls, the focal time will be the date of telephone interview. For cases, the focal time will be the date of their diagnosis of Lyme disease.
Exclusion criteria	Significant immunosuppression (because of either an illness or a medication) or prior treatment for Lyme disease.
Fully immunized	Received $\geq 3$ doses of Lyme vaccine at least 1 month prior to focal time.

## Statistical Analysis

Unadjusted and adjusted odds ratios (OR), along with their 95% CI, were estimated using conditional logistic regression. Under the rare-disease assumption, the OR from a case-control study is roughly equivalent to the risk ratio, and VE can be calculated as  $(1 - \text{OR}) \times 100\%$ .<sup>18,19</sup> We first considered the unadjusted model, which contained only the dichotomous immunization status variable, and the full model, which adjusted for all potential confounders. We then tested 3 traditional, single-model-selection approaches. The first selection approach was the two-stage method. In this approach, we first considered each of the potential confounders separately by fitting numerous binary conditional logistic regression models. A multivariable model was then built by including all variables that had a p-value cutoff of  $<0.05$  on the binary models. The second selection approach we used employs backward selection. This approach began with the “full model” (all potential confounders), and then each potential confounder was eliminated, one-at-a-time (least significant first), based on whether it had a p-value cutoff of  $<0.2$ .<sup>20</sup> The third approach we used incorporated the leaps-and-bounds algorithm to search across the modeling space and select the model with the lowest Akaike information criterion (AIC) score.<sup>21</sup>

## Bayesian Model Averaging Framework

In the Bayesian framework, it is asserted that there is uncertainty about what the best model is. Hence, our approach begins by first assigning a prior probability to each model to create a probability distribution of all candidate models. The prior probabilities reflect how likely a given model is based on how well it fits the data and how many covariates exist in each model. To approximate this, we use the Bayesian information criterion (BIC), which is a measure of goodness of fit that penalizes the overfitting models (based on the number of parameters in the model) and minimizes the risk of multicollinearity.<sup>22,23</sup> BIC estimates can be obtained in most statistical packages or can easily be computed with two values from the logistic regression outputs (sample size and the maximum likelihood).

Our approach entails systematically searching for potential models using every combination without repetition of non-collinear control variables. After fitting each potential model, we assign a BIC score to each model and approximate the probability of each model by comparing its fit and complexity relative to the “best BIC-model” (the model with the lowest BIC). Model weights were then estimated using the relative difference in BICs, which is equivalent to the Bayesian posterior model probability given the candidate models.

The BIC-derived weights are then used to rank all models by their posterior probability of being the best approximation of the true model. We use this rank to ascertain how informative each candidate model is relative to the top model. Then, we use the estimated posterior model probabilities to reduce the modeling space to the most realistic models by dropping those that provide little to no improvement in fit (ie, those with a posterior probability of  $<1\%$  relative to the best model).

After reducing the modeling space to only the top subset of models, we recompute the BIC-weights for the newly defined subset of models and only use this top subset for model averaging procedures. Credibility intervals (95% CI) are calculated using the unconditional error. After calculating the pooled OR and 95% CI on the log scale, they are exponentiated to obtain the weighted average of the VE estimate.

This BMA approach also allows for comparisons to be made about the relative importance of each potential confounder based on how it may improve the fit of the model. To make these assessments, we extract the BIC-weights from the top subset of candidate models that included a confounder of interest and take the sum of them to estimate the given variable’s posterior inclusion probability (PIP). For additional details about our BMA approach, see [Supplemental Methods](#).

## Results

### Lyme Vaccine Effectiveness

A total of 912 subjects were included in the analysis (358 cases and 912 matched controls) of the parent Lyme VE study. Data from a total of 13 potential confounders were considered for inclusion in the multivariate models as potential covariates. The distribution of potential confounders by case status is shown in [Table 2](#). An initial set of 8192 distinct candidate models were fit and considered for the analyses, ranging from the simplest model (only the vaccine variable) to the most complex model (the vaccine variable plus all potential confounders). The unweighted VE point estimates across all models ranged between 56% and

**Table 2** Characteristics of Participants Enrolled in Lyme Vaccine Study by Case Status, N=912

Variables	Cases, N (%)		Matched Controls, N (%)		Standardized Mean Difference*
Total	358	39%	554	61%	–
Demographics					
Age (years), median (range)	46.0	(36.0–55.0)	47.0	(37.0–55.0)	0.025
Non-Hispanic White	348	(97.2%)	540	(97.5%)	0.017
Female	188	(52.5%)	365	(65.9%)	0.274
Some college education	249	(69.6%)	424	(76.5%)	0.158
Currently unemployed	41	(11.5%)	71	(12.8%)	0.042
Risk Factors					
Has an occupational exposure to ticks	52	(14.5%)	50	(9.0%)	0.171
Lives near wooded area	340	(95.0%)	512	(92.4%)	0.105
Engages in frequent outdoor activities	307	(85.8%)	448	(80.9%)	0.131
Has pets at home	222	(62.0%)	370	(66.8%)	0.100
Protective Measures					
Uses protective clothing while outdoors	157	(43.9%)	325	(58.7%)	0.299
Use of tick repellent on skin or clothing	98	(27.4%)	189	(34.1%)	0.146
Performs checks for ticks post-exposure	270	(75.4%)	419	(75.6%)	0.005
Sprays property with tick acaricides	13	(3.6%)	38	(6.9%)	0.145

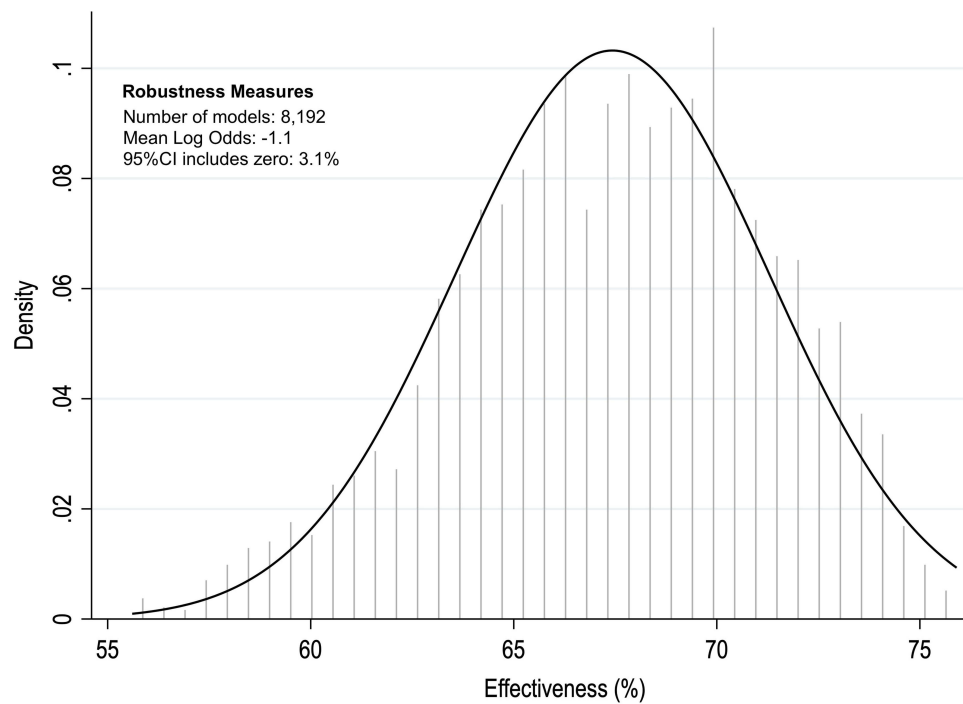
**Note:** \*The standardized mean difference (SMD) is the difference in means between cases and controls in units of the pooled standard deviation.

76% (Figure 1). The statistical significance was robust to model specification, with confidence intervals crossing zero in <5% of all candidate models (n=8192).


## Bayesian Model Averaging

After the exclusion of the non-informative models (those with a probability of being the best model <0.01), the top subset of candidate models was selected (n=15) and weights for each model in the top subset were re-normalized for model averaging procedures. The VE estimates for the top 15 models and the consensus BMA estimates are shown in Figure 2. For comparison, the estimate using a model that only controlled for race, which was excluded from model averaging procedures due to its poor fit, is also shown in Figure 2.

Each component of the top 15 models, their BIC-derived weights (ie, their model probability), and the PIP of each potential confounder are summarized in Figure 3. The PIP can be used to assess how much the data support adding a variable to a regression. When interpreting the PIP, a cutoff of 0.5 is generally used to determine which regressors have evidence of an effect, and values >0.95 are suggestive of strong evidence for an effect.<sup>14</sup> From Figure 3, we can see that 6 out of the 14 potential confounders considered had very low PIP's (<0.01), thus including them in the final model would merely have resulted in a reduction of statistical efficiency without adding to the model's explanatory power. In the original analysis, traditional 2-stage selection was used, and 7 covariates were retained in their final model. Of these, only 2 covariates (use of protective clothing while outdoors and female sex) exhibited strong evidence of an effect on BMA (PIP >0.95). Using BMA, the model-averaged aOR was 0.31, corresponding to a VE of 69% (95% CI: 18–88%). This BMA estimate takes into account the effect of all potential confounders included in the most plausible models, weighted by the probability of each model.



**Figure 1** Distribution of the vaccine's effectiveness across all possible models.



**Figure 2** Vaccine effectiveness for model-averaged and the top subset of candidate models. \*Excluded from BMA if PIP <0.01.

**Abbreviations:** AOR, adjusted odds ratio; VE, vaccine effectiveness; BMA, Bayesian model averaging; PIP, posterior inclusion probability.

Variables	Model Rank															PIP*
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
Lyme vaccine	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	1.00
Age (yrs.), median (range)																<0.01
Non-Hispanic White																<0.01
Female	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	1.00
Some college education			■	■				■			■	■	■			0.37
Currently unemployed																<0.01
Has an occupational exposure to ticks					■	■							■	■		0.29
Lives near wooded area																<0.01
Engages in frequent outdoor activities																0.02
Has pets at home							■	■	■	■	■	■	■	■		0.23
Uses protective clothing while outdoors	■	■	■	■	■	■	■	■	■	■	■	■	■	■	■	1.00
Use of tick repellent on skin or clothing																<0.01
Performs checks for ticks post-exposure																<0.01
Sprays property with tick acaricides		■	■					■		■			■			0.50
Model Probability	0.15	0.15	0.12	0.12	0.08	0.08	0.05	0.04	0.04	0.04	0.03	0.03	0.03	0.03	0.02	1.00

\*PIP= posterior inclusion probability

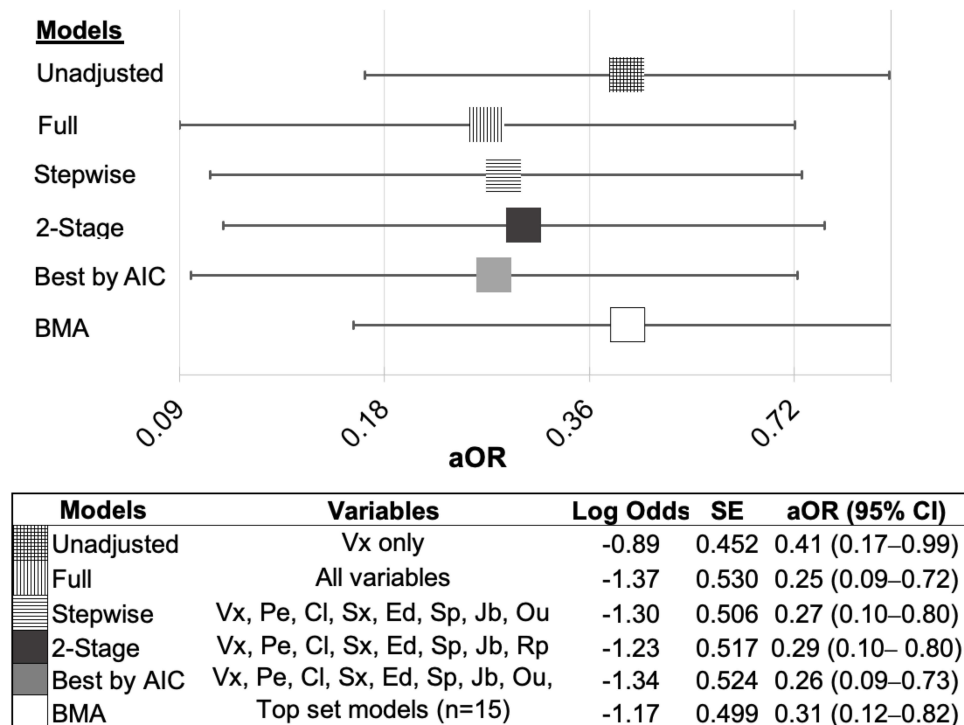
Probability	
Highest	■
Lowest	■
Not Included	□
Forced in Model	■

**Figure 3** Composition and posterior inclusion probabilities of confounders among the top subset of models.

**Notes:** In this heat map, each row corresponds to a potential confounder, and each column corresponds to a particular model. The best model is on the left, and the worst model is on the right. The shade of each square represents the posterior probability of each potential confounder. The darker shades represent higher posterior probabilities, and unshaded squares represent variables that were not included in a given model. Since the immunization status was included in all models a priori, it is checked, and the PIP for this variable is set to 1.

## Comparison of Model Selection Approaches

Figure 4 summarizes the performance of the various single-best model-selection techniques examined. The unadjusted effectiveness of 3 doses of Lyme vaccine was 59% (95% CI: 1% to 83%). Controlling for the confounders that were statistically significant on bivariate analysis (ie, two-stage method), the estimated adjusted effectiveness of 3 doses of the vaccine was 71% (95% CI: 21–90%). The estimate of VE using the backward stepwise elimination approach was 73%



**Figure 4** Comparison of single-model selection and Bayesian model-averaging approaches.

**Abbreviations:** Vx, vaccine status; Pe, has pets at home; Cl, wears protective clothing while outdoors; Sx, male or female sex; Ed, at least some college education; Sp, sprays property with acaricide; Jb, has job exposure to ticks; Ou, engages in frequent outdoor activities; Rp, use of tick repellent on skin or clothing; Tk, use of tick repellent on skin or clothing; SE, standard error; aOR, adjusted odds ratio; BMA, Bayesian model averaging; AIC, Akaike information criterion.



(95% CI: 26–90%), slightly higher than the two-stage approach. The estimates of VE using the leaps and bounds algorithm and selecting the best model based on AIC was 74% (95% CI: 27–91%), which most closely resembled the estimate of the full model (VE=75%; 95% CI: 28–91%). Between all of the confounder selection approaches tested, the BMA approach resulted in the lowest standard error (SE=0.499).

## Discussion

Although case–control VE studies have been conducted for decades, little prior work has been done to define the optimal approaches for variable selection in this context. In this article, we explore the value of using BMA in the analysis of case–control VE data. Several key insights can be gained from this study. First, we show how the selection of a poorly specified model could significantly affect the interpretation of the VE. For example, when re-analyzing the Lyme VE data, had the model that only controlled for race been selected as the final model (excluded model in Figure 2), the VE would have been significantly lower (59%), and the confidence intervals would have been bloated to suggest that it was not statistically significant (95% CI: –1% to 83%). This demonstrates the potential advantage of BMA, as it is systematic in the consideration of alternative models and can facilitate assessments of the effect confounders have when included in the model.<sup>14</sup>

Second, we show how in the context of uncertainty, BMA methods are likely to provide a more complete picture than single-model-selection methods. We found that the top-ranked model in the Lyme VE study had a weight of 0.15, which means that there is a 15% probability that it is the best model given the alternative models. Had this model clearly stood apart from all candidate models as the one with the best fit (ie, if it is >90% better than all other candidate models), then further BMA would have been unnecessary, and the single best model could have been used to estimate the VE. However, the models ranked 2–4 had a similar fit and competed for the top spot (probabilities of 0.12–0.15). This variability in fit underlines another advantage of using the BMA estimate, as this uncertainty in model selection is then accounted for in the credibility intervals across several models.<sup>24</sup>

Last, it is worth noting that while BMA outperformed traditional single-best model selection methods in terms of standard errors, the resulting point estimates and confidence intervals were not substantially different. In practice, researchers do not know a priori if the model selection process will result in heterogeneity. However, even if no substantial heterogeneity exists, there can still be advantages to adopting the BMA framework. Investigators can enhance transparency by reporting the VE distribution under a variety of modeling scenarios, as well as by describing what conditions are needed to reverse their conclusions. Investigators can also use the BMA framework to test the robustness of their findings to alternative modeling configurations. In our analysis, for example, after systematically evaluating all candidate models, we found that >95% of the models supported the conclusion that the effectiveness of 3 doses of the Lyme vaccine was statistically significant. This consistency suggests robustness and reduces uncertainty about the possibility of false-positive conclusions (ie, classifying the vaccine's effectiveness as statistically significant when it is not).

## Conclusion

In this paper, we show how a BMA analytic approach can provide greater transparency in publication by offering an avenue for the reporting of the VE's distribution under numerous hypothetical scenarios. The BMA procedures outlined in this study are intuitive and, by using BIC-derived weights, are easy to implement with most statistical software. Although this framework does not rule out the possibility that uncontrolled confounding or systematic biases could have affected the results of the study, by incorporating model uncertainty into the parameter estimation and through multi-model inferences, BMA can be used to lend an additional measure of rigor and credibility to a well-designed study.

## Ethics Statement

The study was conducted according to the guidelines of the Declaration of Helsinki and approved by the Institutional Review Board of Yale University, the United States Centers for Disease Control and Prevention, and the Connecticut Department of Public Health. Written informed consent was obtained from all subjects involved in the Lyme vaccine effectiveness study.

Owing to data privacy regulations, the raw Lyme vaccine data for this study cannot be shared.

## Acknowledgments

We thank Marietta Vazquez for assistance with this project. These analyses were supported, in part, by the following grants from the National Institutes of Health: UL1TR000142 (E.D.S.) K23AI159518 (C.R.O.), and KL2-TR001862 (E.D. S.), from the National Center for Advancing Translational Science (NCATS) at the National Institutes of Health and NIH Roadmap for Medical Research. Its contents are solely the authors' responsibility and do not necessarily represent the official views of NIH. Sponsors had no role in study design; in the collection, analysis and interpretation of data; in the writing of the report; or in the decision to submit the article for publication.

## Disclosure

Dr Shapiro received an honorarium from Sanofi for a conference call about Lyme vaccine; he also receives royalties from UpToDate and has served as an expert witness in cases related to Lyme disease. Dr Weinberger has received consulting fees from Pfizer, Merck, GSK, and Affinivax for work unrelated to this manuscript. Drs. Weinberger and Oliveira have served as investigators on research grants from Pfizer and Merck to Yale, which are unrelated to this manuscript. The authors report no other conflicts of interest in this work.

## References

1. Burnham KA. *Model Selection and Multimodel Inference*. 2nd ed. Springer-Verlag; 2002.
2. Clemens JD, Shapiro ED. Resolving the pneumococcal vaccine controversy: are there alternatives to randomized clinical trials? *Rev Infect Dis*. 1984;6(5):589–600. doi:10.1093/clinids/6.5.589
3. Shapiro ED. Case-control studies of the effectiveness of vaccines: validity and assessment of potential bias. *Pediatr Infect Dis J*. 2004;23(2):127–131. doi:10.1097/01.inf.0000109248.32907.1d
4. Breslow NE. Statistics in epidemiology: the case-control study. *J Am Stat Assoc*. 1996;91(433):14–28. doi:10.1080/01621459.1996.10476660
5. Robinson LJ, Jewell NP. Some surprising results about covariate adjustment in logistic regression models. *Int Stat Rev*. 1991;59(2):227–240. doi:10.2307/1403444
6. Greenland S. Quantifying biases in causal models: classical confounding vs collider-stratification bias. *Epidemiology*. 2003;14(3):300–306. doi:10.1097/01.EDE.0000042804.12056.6C
7. Kevin CKB, Rueda R, Rueda MR. Omitted variables, countervailing effects, and the possibility of overadjustment. *Political Sci Res Methods*. 2018;6(2):343–354. doi:10.1017/psrm.2016.46
8. Schisterman EF, Cole SR, Platt RW. Overadjustment bias and unnecessary adjustment in epidemiologic studies. *Epidemiology*. 2009;20(4):488–495. doi:10.1097/EDE.0b013e3181a819a1
9. Derksen SK, Keselman HJ. Backward, forward and stepwise automated subset selection algorithms: frequency of obtaining authentic and noise variables. *Br J Math Stat Psychol*. 1992;45:465. doi:10.1111/j.2044-8317.1992.tb00992.x
10. Wu Z, Zhao H. Statistical power of model selection strategies for genome-wide association studies. *PLoS Genet*. 2009;5(7):e1000582. doi:10.1371/journal.pgen.1000582
11. Hoeting JA, Davis RA, Merton AA, Thompson SE. Model selection for geostatistical models. *Ecol Appl*. 2006;16(1):87–98. doi:10.1890/04-0576
12. Kurum E, Warren JL, Schuck-Paim C, et al. Bayesian model averaging with change points to assess the impact of vaccination and public health interventions. *Epidemiology*. 2017;28(6):889–897. doi:10.1097/EDE.0000000000000719
13. Jackson ML, Ferdinands J, Nowalk MP, et al. Differences between Frequentist and Bayesian inference in routine surveillance for influenza vaccine effectiveness: a test-negative case-control study. *BMC Public Health*. 2021;21(1):516. doi:10.1186/s12889-021-10543-z
14. Viallefont V, Raftery AE, Richardson S. Variable selection and Bayesian model averaging in case-control studies. *Stat Med*. 2001;20(21):3215–3230. doi:10.1002/sim.976
15. Wang D, Zhang W, Bakhai A. Comparison of Bayesian model averaging and stepwise methods for model selection in logistic regression. *Stat Med*. 2004;23(22):3451–3467. doi:10.1002/sim.1930
16. Raftery AE. Approximate Bayes factors and accounting for model uncertainty in generalised linear models. *Biometrika*. 1996;83(2):251–266. doi:10.1093/biomet/83.2.251
17. Oliveira CM, Vazquez M. Clinical effectiveness of Lyme vaccine: a matched case-control study. *Open Forum Infect Dis*. 2021;21:319. doi:10.1093/ofid/ofab142
18. Niccolai LM, Ogden LG, Muehlenbein CE, Dziura JD, Vazquez M, Shapiro ED. Methodological issues in design and analysis of a matched case-control study of a vaccine's effectiveness. *J Clin Epidemiol*. 2007;60(11):1127–1131. doi:10.1016/j.jclinepi.2007.02.009
19. Oliveira CR, Niccolai LM, Sheikh H, et al. Assessment of clinical effectiveness of BNT162b2 COVID-19 vaccine in US adolescents. *JAMA Netw Open*. 2022;5(3):e220935. doi:10.1001/jamanetworkopen.2022.0935
20. Williams R. Step we gaily, on we go. *Stata J*. 2007;7(2):272–274. doi:10.1177/1536867X0700700212
21. Lindsey CS, Sheather S. Variable selection in linear regression. *Stata J*. 2010;10(4):650–669. doi:10.1177/1536867X10101000407
22. Raftery AE. Use and communication of probabilistic forecasts. *Stat Anal Data Min*. 2016;9(6):397–410. doi:10.1002/sam.11302
23. Raftery AE. Bayesian model selection in social research. *Sociol Methodol*. 1995;25:111–163. doi:10.2307/271063
24. Mu Y, See I, Edwards JR. Bayesian model averaging: improved variable selection for matched case-control studies. *Epidemiol Biostat Public Health*. 2019;16(2). doi:10.2427/13048



**Clinical Epidemiology**

Dovepress

**Publish your work in this journal**

Clinical Epidemiology is an international, peer-reviewed, open access, online journal focusing on disease and drug epidemiology, identification of risk factors and screening procedures to develop optimal preventative initiatives and programs. Specific topics include: diagnosis, prognosis, treatment, screening, prevention, risk factor modification, systematic reviews, risk & safety of medical interventions, epidemiology & biostatistical methods, and evaluation of guidelines, translational medicine, health policies & economic evaluations. The manuscript management system is completely online and includes a very quick and fair peer-review system, which is all easy to use.

Submit your manuscript here: <https://www.dovepress.com/clinical-epidemiology-journal>