RESEARCH ARTICLE

Do speed cameras reduce road traffic collisions?

Daniel J. Graham^{1*}, Cian Naik², Emma J. McCoy¹, Haojie Li³

1 Imperial College London, London, United Kingdom, **2** University of Oxford, Oxford, United Kingdom, **3** Southeast University, Nanjing, China

* d.j.graham@imperial.ac.uk

Abstract

This paper quantifies the effect of speed cameras on road traffic collisions using an approximate Bayesian doubly-robust (DR) causal inference estimation method. Previous empirical work on this topic, which shows a diverse range of estimated effects, is based largely on outcome regression (OR) models using the Empirical Bayes approach or on simple before and after comparisons. Issues of causality and confounding have received little formal attention. A causal DR approach combines propensity score (PS) and OR models to give an average treatment effect (ATE) estimator that is consistent and asymptotically normal under correct specification of either of the two component models. We develop this approach within a novel approximate Bayesian framework to derive posterior predictive distributions for the ATE of speed cameras on road traffic collisions. Our results for England indicate significant reductions in the number of collisions at speed cameras sites (mean ATE = -15%). Our proposed method offers a promising approach for evaluation of transport safety interventions.

OPEN ACCESS

Citation: Graham DJ, Naik C, McCoy EJ, Li H (2019) Do speed cameras reduce road traffic collisions? PLoS ONE 14(9): e0221267. <https://doi.org/10.1371/journal.pone.0221267>

Editor: Quan Yuan, Tsinghua University, CHINA

Received: May 13, 2019

Accepted: August 4, 2019

Published: September 16, 2019

Copyright: © 2019 Graham et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The road accident data are available from the UK Department of Transport at the following website: <https://data.gov.uk/dataset/cb7ae6f0-4be6-4935-9277-47e5ce24a11f/road-safety-data>.

Funding: The author(s) received no specific funding for this work.

Competing interests: The authors have declared that no competing interests exist.

Introduction

Fixed speed limit enforcement cameras are a common intervention used to encourage drivers to comply with maximum legal speed limits. The cameras are installed at sites on selected links in order to detect speed limit violations, which can subsequently be punished with monetary fines, driver licence disqualification points, or prosecution. Since the introduction of speed cameras (SCs) there has been considerable debate about their effects on road traffic collisions (RTCs). At various times claims have been made that SCs serve to reduce RTCs, that they have no effect, or even that they increase RTCs by encouraging more erratic driving behaviour.

The paper is structured as follows. Section two outlines broad trends in road traffic casualties for Britain and then sets out a formal causal modelling framework to estimate the effects of SCs on RTCs. Section three describes our approximate Bayesian DR approach and presents some simulations that demonstrate its properties. Section four describes the data available for estimation and outlines our chosen model specifications. Results are then presented in section five and conclusions are drawn in the final section.

Speed cameras, road traffic collisions and causality

A number of academic studies of the effect of speed cameras on RTCs have been undertaken [1]. Most studies find that speed cameras have led to a reduction in RTCs, but the range of estimated effects is large (from 0% to -55%). Variation in estimates is to be expected given that study results pertain to diverse empirical contexts, but it is also the case that a number of different methods have been applied which can have a critical influence on results obtained. In particular, since SCs are not randomly assigned, it is essential that any adopted method recognises that the observed relationship between SCs on RTCs may be subject to confounding. Confounding arises when the characteristics that influence treatment assignment (i.e. whether a site is 'treated' and 'untreated' with an SC) also matter for outcomes (i.e. RTCs). Regression to the mean (RTM), for instance, is a well known manifestation of confounding that arises via 'selection bias'.

The extent to which confounding has been recognised and addressed in existing studies varies considerably. Some studies have simply ignored it, using simple before-and-after methods with control groups [2–9]. Others have used the empirical Bayes (EB) method as suggested by [10], largely to adjust for effects of confounding that arise via RTM [11–16]. Finally, there are a small number of studies that have used time-series methods, either interrupted time-series analyses with control groups or ARIMA, to test for changes in outcome rates [17–19]. Where studies have attempted to address confounding this has been done via the inclusion of covariates in outcome regression (OR) models, typically using Poisson or negative binomial Generalised Linear Models (GLMs).

In a previous paper we adopted a propensity score (PS) matching approach to evaluate the effectiveness of speed cameras [1]. A key advantage of the PS over OR approach is that it provides an effective way of isolating a valid control group by ensuring that the distribution of pre-treatment covariates matches those of the treated group and that genuine overlap in the support of the covariates exists between the two groups. However, as with the OR approach, valid inference from PS models crucially depends on the unknown PS model being correctly specified.

In this paper we build on our previous work by developing and applying an estimation approach which we believe has much to offer in evaluating the effectiveness of road safety interventions. Our approach uses the principle of doubly-robust (DR) estimation, which provides robustness to model misspecification by combining both OR and PS models to derive an average treatment effect (ATE) estimator which is consistent and asymptotically normal under correct specification of just one of the two component models. The DR property offers a significant methodological advance for traffic safety analyses because it allows us evaluate interventions using combined inference from two key modelling standpoints: via a model for factors influencing assignment of road safety measures and via a model for the determinants of RTCs. A good specification of just one of these two models will yield valid inference.

The DR approach is attractive for our application because the PS and OR models we can construct make different assumptions about the nature of confounding. For the PS model, we are able to faithfully represent via measured covariates the formal criteria that exist for the assignment of speeds cameras to sites. For the OR model, we can difference our response variable before and after treatment to allow for the existence of site level time-invariant unobserved effects in addition to measured confounders.

To avoid common sources of misspecification error, we estimate both of our component models using semiparametric Generalized Additive Mixed Models (GAMMs) which make minimal a priori assumptions on the functional form of the relationships under study. We also use a matching algorithm prior to forming the DR model to establish a valid control group.

Thus, in our approach, potential biases from confounding are addressed by combining three compatible modelling tools: via matching to achieve comparability between treated and control sites and via model based adjustment for valid ATE estimation through the regression model for RTCs and the PS model for the treatment assignment mechanism.

DR estimators have been studied and applied extensively in the frequentist setting [20–26]. A further contribution of the paper is that we develop our binary DR estimator within the Bayesian paradigm. A Bayesian representation of the DR model has proven difficult to formulate in previous work because DR estimators are typically constructed as solutions to estimating equations based on a set of moment restrictions that do not imply fully specified likelihood functions. We choose the Bayesian paradigm for three main reasons. First, DR estimation of the ATE involves prediction and extrapolation over covariate distributions with underlying uncertainty in parameter estimates. Bayesian inference provides a suitable framework for prediction that explicitly addresses such uncertainty in the sense that both the predicted observations, and the relevant parameters for prediction, have the same random status. Second, by deriving a posterior predictive distribution for the ATE, rather than a fixed value, we can make probability statements about the causal quantity of interest allowing us to discuss findings in relation to specific hypotheses or in terms of credible intervals which can offer a more intuitive understanding of the effects of SCs for public policy formulation. Finally, we develop an approximate Bayesian approach that can utilise prior information about the parameters of interest, which could be useful in evaluating safety interventions when historical data or training data from other regions are available.

A causal inference framework to quantify the effects of speed cameras

Road traffic casualties in Britain

For the year zcolored spanning October 2015 to September 2016 the UK DfT recorded a total of 182,747 casualties on British roads of which 25,420 were classified as killed or seriously injured (KSIs) and 1,800 as fatalities [27]. Since 2010 the annual numbers of fatalities and KSIs have not changed significantly, following several years in which road safety was improving. The average number of fatal road traffic incidents over the period 2010 to 2016 is approximately 1,800. Since the volume of road traffic has continued to grow over this period, however, the number of fatalities per vehicle mile driven has been falling [28].

The DfT argue that there is good evidence to suggest that while the absolute number of fatalities on British roads now appears to be relatively static, overall absolute casualty numbers are continuing to fall. In short, levels of safety appear to be improving in relative terms and not deteriorating in absolute terms. Given the changes that have occurred in vehicle technology, medical care, and road safety interventions, however, the DfT also note that a comprehensive causal understanding of the factors underpinning casualty trends is currently out of reach. In this paper we attempt to contribute to such an understanding by quantifying the causal impact of one type of safety intervention: speed cameras (SCs).

ATE estimation within the potential outcomes framework

Our sample comprises n , $i = 1, \dots, n$, links on the road network. Some links have a SC other do not. We define $D_i \in \{1, 0\}$ as a binary random variable indicating the presence or otherwise of a SC and we refer to this as the treatment variable. We are interested in the effect of the treatment on an outcome Y_i , which measures collision frequency. We define $Y_i(1)$ and $Y_i(0)$ as the *potential outcomes* for unit i under treated and control status respectively. Recognising that

SCs are not assigned randomly, we also define X_i as a random vector of pre-treatment covariates that capture characteristics of links that are relevant to whether a SC was assigned or not, and are also relevant for outcomes. Thus, the data we observe for each link is a random vector, $z_i = (y_i, d_i, x_i)$, where y_i is the response or outcome, d_i is treatment status, and x_i a vector of pre-treatment (or baseline) covariates.

Ideally, we would assess the effects of SCs on each link by calculating the individual causal effect (ICE): $\tau_i = Y_i(1) - Y_i(0)$, but the data reveal only outcomes that have actually occurred not potential outcomes. Thus, the data reveal the random variable $Y_i = Y_i(1)I_1(D_i) + Y_i(0)(1 - I_1(D_i))$, where $I_1(D_i)$ is the indicator function for receiving the treatment. But they do not reveal the joint density, $f(Y_i(0), Y_i(1))$, since a SC cannot be both present and absent on a link simultaneously. Consequently, our focus inference on estimation of the ATE, defined as

$$\tau = \mathbb{E}[Y_i(1)] - \mathbb{E}[Y_i(0)], \quad (1)$$

which measures the difference in expected outcomes under treatment and control status. Note that sometimes in the safety literature the average effect of the treatment on the treated (ATT), i.e. $\mathbb{E}[Y_i(1) - Y_i(0)|D_i = 1]$, is estimated, but since in our case all matched units could potentially be exposed to the treatment ATE is a more appropriate estimand than ATT.

We can estimate the ATE without observing all potential outcomes if three key assumptions hold.

1. Conditional independence—the potential outcomes for each unit must be conditionally independent of treatment assignment given observed covariates X_i : $Y_i(0), Y_i(1) \perp\!\!\!\perp I_1(D_i)|X_i$.
2. Common Support—the support of the conditional distribution of assignment to treatment given covariates must overlap with that of the conditional distribution of assignment to control. This requires that $0 < \Pr(I_1(D_i) = 1|X_i = x) < 1, \forall x$.
3. Stable Unit Treatment Value Assumption (SUTVA)—observed and potential outcomes must satisfy the SUTVA [29–32], which requires that observed outcomes under a given treatment allocation are equivalent to potential outcome under that same treatment allocation: $Y_i = I_1(D_i)Y_i(1) + (1 - I_1(D_i))Y_i(0)$ for all $i = 1, \dots, n$.

These three assumptions are sometime referred to collectively as *strong ignorability*, and they allow the ATE to be estimated from observational data as follows.

$$\tau = \mathbb{E}_i(Y_i(1) - Y_i(0)) = \mathbb{E}_X[\mathbb{E}_i(Y_i(1)|X_i = x) - \mathbb{E}_i(Y_i(0)|X_i = x)] \quad (2)$$

$$= \mathbb{E}_X[\mathbb{E}_i(Y_i(1)|X_i = x, I_1(D_i) = 1) - \mathbb{E}_i(Y_i(0)|X_i = x, I_1(D_i) = 0)] \quad (3)$$

$$= \mathbb{E}_X[\mathbb{E}_i(Y_i|X_i = x, I_1(D_i) = 1) - \mathbb{E}_i(Y_i|X_i = x, I_1(D_i) = 0)]. \quad (4)$$

The equality of (2) and (3) is justified via conditional independence, the substitution of observed and potential outcomes implied by the SUTVA gives (4), and common support ensures that there are both treated and control units and thus that the population ATE in (4) is estimable.

Thus, if strong ignorability holds, the potential outcomes approach offers a route to obtaining valid causal estimates of the ATE of SCs. To proceed we need to estimate the relevant expectations in (4) above.

Causal estimators

Following [33], we can express our observed data as joint densities of the form

$$f_Z(z) = f_{Y|D,X}(y|d, x)f_{D|X}(d|x)f_X(x). \quad (5)$$

If strong ignorability holds, one of the following estimators can be used to estimate the ATE of SCs.

1. Outcome regression (OR) model— $f_{D|X}(d|x)$ and $f_X(x)$ are unspecified and instead we construct a model for $\mathbb{E}[Y_i|D_i, X_i]$; the mean of the conditional density of outcome given covariates. This is done via the OR model $\Psi^{-1}\{m(D_i, X_i; \beta)\}$, for given link function Ψ , regression function $m()$, and unknown parameter vector β . Correct specification of the OR model for the mean response provides a consistent estimates of the ATE using

$$\hat{\tau}_{OR} = \frac{1}{n} \sum_{i=1}^n \left[\Psi^{-1}\{m(1, X_i; \hat{\beta})\} - \Psi^{-1}\{m(0, X_i; \hat{\beta})\} \right]. \quad (6)$$

2. Propensity score (PS) model—under the PS approach we assume a model for $f_{D|X}(d|x)$; the conditional probability of assignment to treatment given covariates, and leave $f_{Y|D,X}(y|d, x)$ and $f_X(x)$ unspecified. The PS model, which we write $\pi(D_i|X_i; \alpha)$, can form the basis of several different nonparametric estimators, but of primary interest here is the weighting estimator proposed by [34]

$$\hat{\tau}_{IPW} = \frac{1}{n} \sum_{i=1}^n \left[\frac{I_1(D_i) \cdot Y_i}{\pi(D_i|X_i; \hat{\alpha})} - \frac{[1 - I_1(D_i)] \cdot Y_i}{1 - \pi(D_i|X_i; \hat{\alpha})} \right], \quad (7)$$

which is consistent under correct PS model specification due to the fact that $\mathbb{E}[Y_i(1)] = \mathbb{E}\{[Y_i(1) \cdot I_1(D_i)]/\pi(D_i|X_i; \alpha)\}$ and similar for control treatment status.

3. Doubly-robust (DR) model—specify both an OR and PS model and use them in a combined estimation to yield a DR estimator. This can be achieved by forming a function of the inverse estimated PSs and using that function to weight or augment the OR model. Here, we estimate the weighted model

$$e(D_i, X_i; \xi) = \Psi^{-1}\{m(D_i, X_i; \xi)\} \quad (8)$$

where the unknown parameter vector ξ is obtained by weighting the model with

$$\hat{\kappa}_i(D_i, X_i) = \frac{I_1(D_i)}{\hat{\pi}(D_i|X_i; \hat{\alpha})} + \frac{1 - I_1(D_i)}{1 - \hat{\pi}(D_i|X_i; \hat{\alpha})}. \quad (9)$$

This model will provide consistent estimates of $\mathbb{E}[Y_i|D_i, X_i]$ if the model $\Psi^{-1}\{m(I_1(D_i), X_i; \beta)\}$ is correct, since although weighting may reduce efficiency, it does not adversely affect the consistency and asymptotic normality of the OR model. If the OR model specification is incorrect, but not the PS, the model DR model will still be consistent because weighting yields estimating equations of the form

$$\sum_{i=1}^n \hat{\kappa}_i(D_i, X_i) \frac{1}{\phi} \frac{\partial e(d_i, x_i; \xi)}{\partial \xi^T} [y_i - e(d_i, x_i; \xi)] = 0, \quad (10)$$

in which $\phi_i \equiv \phi(D_i, X_i)$ is a working conditional variance for Y_i given (D_i, X_i) . These effectively correct for the bias in approximating $\mathbb{E}[Y_i|D_i, X_i]$ using $\Psi^{-1}\{m(D_i, X_i; \beta)\}$ [24].

We use estimates of ξ to form the DR estimator

$$\hat{\tau}_{DR} = \frac{1}{n} \sum_{i=1}^n \left[\Psi^{-1} \{m(1, X_i; \hat{\xi})\} - \Psi^{-1} \{m(0, X_i; \hat{\xi})\} \right]. \quad (11)$$

Approximate Bayesian doubly-robust estimation

So far our discussion of DR estimation has been conducted within the frequentist semiparametric paradigm. As alluded to previously, however, there are good reasons why a Bayesian inferential approach may be particularly beneficial for estimation of road safety interventions. But inference for DR estimators is less straightforward in the Bayesian than frequentist setting because DR models are based on moment restrictions that do not yield fully specified likelihood functions. Here, we make some improvements to the approach proposed by [35] in the context of continuous treatment. In contrast to that paper we focus on binary treatments using PS weighting rather than augmentation to achieve the DR model and we implement ways of incorporating prior information into the posterior distribution of the ATE. The basic theory underpinning approximate Bayesian inference in this context is covered comprehensively in [35] and so we provide only a brief summary here.

The Bayesian bootstrap was first introduced by [36] and applied in weighted likelihood models by [37]. The basic idea is to create new datasets by repeatedly re-weighting the original data in order to obtain the posterior distribution for some parameter of interest. If we treat our observed data, z_i say, as effectively coming from a multinomial distribution with distinct values a_k , $k = (1, \dots, K)$, and attach a probability to each distinct value $\theta = (\theta_1, \dots, \theta_K)$, then by placing an improper Dirichlet prior on θ

$$\pi(\theta) \propto \prod_{k=1}^K \theta_k^{-1}. \quad (12)$$

the posterior density also has a Dirichlet distribution

$$p(\theta|v) \propto \prod_{k=1}^K \theta_k^{n_k-1}. \quad (13)$$

with parameter n_k . This posterior can be stimulated via the weighted likelihood

$$\tilde{L}(\theta) = \prod_{i=1}^n f(z_i; \theta)^{w_i}, \quad (14)$$

in which the weights $w = (w_1, \dots, w_n)$ are distributed according to the uniform Dirichlet distribution and simulated as n independent standard exponential (i.e. gamma(1,1)) variates and standardised. The weighted likelihood reduces to

$$\tilde{L}(\theta) = \prod_{i=1}^n \left\{ \prod_{k=1}^K \theta_k^{I_k(z_i)} \right\}^{w_i} = \prod_{k=1}^K \theta_k^{\sum_{i=1}^n w_i I_k(z_i)} = \prod_{k=1}^K \theta_k^{n\gamma_k}, \quad (15)$$

say, where $n\gamma_k$ is the sum of the weights w_i for which $z_i = a_k$. Since the vector $\gamma = (\gamma_1, \dots, \gamma_K)$ has a Dirichlet distribution with parameters $n_k = (n_1, \dots, n_K)$,

$$p(\gamma) \propto \prod_{k=1}^K \gamma_k^{n_k-1} \quad (16)$$

and since at the point of maximisation of $\tilde{L}(\theta)$ is $\tilde{\theta} = \gamma$, then the solutions to the maximised weighted likelihood function with repeatedly sampled uniform Dirichlet weights $w^{(l)}$ represent a sample from the posterior of θ under the improper prior $\prod_k \theta_k^{-1}$.

To apply the Bayesian bootstrap to our DR model we estimate

$$e(D_i, X_i; \xi) = \Psi^{-1}\{m_A(D_i, X_i; \xi)\} \quad (17)$$

with weights

$$w_i^{(l)} \cdot \hat{\kappa}_i(D_i, X_i). \quad (18)$$

The maximiser of $\tilde{L}(\xi)$, which we denote $\tilde{\xi}$, implies a solution to

$$\sum_{i=1}^n w_i^{(l)} \cdot \hat{\kappa}_i(D_i, X_i) \cdot \frac{1}{\phi} \frac{\partial e(d_i, x_i; \xi)}{\partial \xi^T} [y_i - e(d_i, x_i; \xi)] = 0, \quad (19)$$

which as noted above has the DR property. We repeatedly draw sets of random weights $\{w_i^{(l)}\}_{i=1}^n$ as n standardised independent standard exponential variates and solve (19) to build up an empirical posterior density of $\tilde{\xi}$, denoted $p_n(\tilde{\xi})$, from which the sampled values $\tilde{\xi}^{(l)}$ are consistent with the DR estimating equations.

[37] apply sampling-importance resampling (SIR) to improve accuracy of the weighted bootstrap approach, but this improvement requires a fully specified likelihood function. Instead, for our restricted moment model, we use the resampling scheme proposed by [38] which extends Rubin's bootstrap in a general Bayesian nonparametric context. Two attractive features of Muliere and Secchi's approach for causal modelling are that it ensures that predictive distributions are not constrained to be concentrated on observed values and it allows us to take prior opinions into account. The posterior predictive distribution of the ATE, incorporating prior information, is obtained in the following way.

- i. Estimate the PS model $\pi(D_i|X_i; \alpha)$, and form

$$\hat{\kappa}_i(d_i, x_i; \hat{\alpha}) = \frac{I_1(d_i)}{\hat{\pi}(d_i|x_i; \hat{\alpha})} + \frac{1 - I_1(d_i)}{1 - \hat{\pi}(d_i|x_i; \hat{\alpha})}. \quad (20)$$

- ii. Draw a single set of random weights $\{w_i^{(l)}\}_{i=1}^n$ and form the combined weights $w_i^{(l)} \cdot \hat{\kappa}_i(d_i|x_i; \hat{\alpha})$ and estimate the weighted model

$$\Psi^{-1}\{m_A(d_i, x_i; \xi^{(l)})\}. \quad (21)$$

- iii. Repeatedly compute (ii) using new weights $\{w_i^{(l)}\}_{i=1}^n$ to obtain the empirical posterior distribution $p_n(\tilde{\xi})$.
- iv. Introduce a prior distribution p_0 for ξ and a positive number k , the 'measure of faith' that we have in this prior. This can range anywhere from 1 to a size comparable to the number of samples of ξ .
- v. Generate m observations $x_1^*, \dots, x_m^*$ from $\frac{kp_0 + Lp_n}{k+L}$, where p_n is as above. We choose $m = L$ in our case.
- vi. For $i = 1, \dots, m$ generate v_i from a $\Gamma(\frac{L+k}{m}, 1)$ distribution.

- vii. Sample new parameters $\tilde{\xi}_{MS}$ from $x_1^*, \dots, x_m^*$ using the weights $v_1, \dots, v_m$ to form the posterior $p_m(\tilde{\xi})$.
- viii. Resample V values of the covariate vector uniformly over the observed values and a single vector $\xi^{(m)}$ from $p_m(\tilde{\xi})$.
- ix. Form a sampled value of the ATE random variable as

$$\tau_{BDR}^{(m)} = \frac{1}{V} \sum_{v=1}^V \left[\Psi^{-1} \{m_A(1, x_v; \tilde{\xi}^{(m)})\} - \Psi^{-1} \{m_A(0, x_v; \tilde{\xi}^{(m)})\} \right]. \quad (22)$$

- x. Repeat this procedure M times, $m = (1, \dots, M)$, to obtain the posterior predictive distribution.

Simulations

In this subsection we present some simulation to demonstrate the DR properties of our approximate Bayesian approach. The simulations are based on the following data generating process: a binary treatment D is assigned as a function of covariate X , and the outcome of interest Y depends on both treatment D and covariate X

$$X \sim \text{Normal}(0, 10)$$

$$D \sim \text{Bernoulli}(\text{expit}(\alpha_0 + \alpha_1 X))$$

$$Y \sim \text{Normal}(\beta_0 + \beta_1 D + \beta_2 X, 5)$$

where $\alpha_0 = 2, \alpha_1 = 0.2, \beta_0 = 10, \beta_1 = 5, \beta_2 = 0.2$. The true ATE is given by parameter β_1 , that is $\tau = 5.0$.

The following models are tested:

1. $\hat{\tau}_{BOR1}$ —an approximate Bayesian OR model based on the correctly specified model: $\mathbb{E}[Y|D, X] = \beta_0 + \beta_1 D + \beta_2 X$. The point estimate reported in the simulations is the mean value of the ATE posterior predictive distribution, i.e.

$$\hat{\tau}_{BOR} = \frac{1}{L} \sum_{l=1}^L \left[\frac{1}{V} \sum_{v=1}^V \left[\Psi^{-1} \{m(1, x_v; \tilde{\beta}^{(l)})\} - \Psi^{-1} \{m(0, x_v; \tilde{\beta}^{(l)})\} \right] \right]. \quad (23)$$

2. $\hat{\tau}_{BOR2}$ —same as [1.] except based on an incorrectly specified OR model with covariate X excluded.
3. $\hat{\tau}_{PS1}$ —an approximate Bayesian inverse PS weighted model based on the correctly specified PS model

$$\hat{\tau}_{PS1} = \frac{1}{L} \sum_{l=1}^L \left[\frac{1}{V} \sum_{v=1}^V \left[y_v \cdot \frac{d_v - \hat{\pi}(d_v|x_v; \tilde{\alpha}^{(l)})}{\hat{\pi}(d_v|x_v; \tilde{\alpha}^{(l)})(1 - \hat{\pi}(d_v|x_v; \tilde{\alpha}^{(l)}))} \right] \right] \quad (24)$$

4. $\hat{\tau}_{PS2}$ —an approximate Bayesian inverse PS weighted model based on an incorrectly specified PS model, in which the PS is generated randomly from the continuous uniform distribution: $\hat{\pi}(D|X) \sim \text{Uniform}(0, 1)$.

5. $\hat{\tau}_{BDR1}$ —an approximate Bayesian DR model based on an incorrectly specified OR model (X excluded) but with weights based on the correct PS model

$$\hat{\tau}_{BDR} = \frac{1}{L} \sum_{l=1}^L \left[\frac{1}{V} \sum_{v=1}^V \left[\Psi^{-1} \left\{ m_A \left(1, x_v; \tilde{\xi}^{(l)} \right) \right\} - \Psi^{-1} \left\{ m_A \left(0, x_v; \tilde{\xi}^{(l)} \right) \right\} \right] \right]. \quad (25)$$

6. $\hat{\tau}_{BDR2}$ —an approximate Bayesian DR model based on a correctly specified OR model but with weights based on the incorrect PS model.
7. $\hat{\tau}_{BDR3}$ —an approximate Bayesian DR model based on the incorrectly specified OR model weighted with weights based on the incorrect PS model.

The simulations are based on 1000 runs on generated datasets of size 1,000. In each case, we place a Normal prior on the treatment coefficient β_1 , with mean equal to the true value (5 in this case). We set the measure of faith k to be relatively low so as not to overly affect the results. Table 1 shows our simulation results. Mean values and variances of the point estimates obtained (i.e. means and variances of the ATE distributions) and the mean squared error (MSE) are reported.

The mean of the posterior distribution for the ATE from the correctly specified OR model, $\hat{\tau}_{BOR1}$, provides a good approximation to the true value of τ . The incorrectly specified OR model, BOR2, fails to address confounding and consequently $\hat{\tau}_{BOR2}$ provides a poor approximation to the true ATE. A good estimate of τ is achieved via the correctly specified PS model ($\hat{\tau}_{PS1}$), but when the PS is model is misspecified ($\hat{\tau}_{PS1}$) the estimate of the ATE is far away from the true value. In our simulations the PS model is severely misspecified, or simply wrong, having being generated randomly. This tendency of the inverse PS model to fail quite considerably under severe misspecification is well known in the literature [26]. Weighting the incorrectly specified OR model with weights $\hat{\kappa}(D, X)$, based on a correctly specified PS model, as in the BDR1 model, provides correction for misspecification bias with an average point estimate very close to the true value, but slightly larger posterior variances relative to the correctly specified OR model. The BDR2 model simulation also produces valid point estimates because weighting by weights based on an incorrectly specified PS model does not induce bias, but it does increase variance. Finally, if both the OR and PS models are wrongly specified as in BDR3, the model fails to produce a good point estimate of the mean ATE.

Data and model specifications

Treatment and outcome variable

We have data on the location of fixed speed cameras for 771 camera sites in the following English administrative districts: Cheshire, Dorset, Greater Manchester, Lancashire, Leicester, Merseyside, Sussex and the West Midlands. These sites form our group of treated units. To select

Table 1. Simulation results for posterior predictive distributions ($\tau = 5.0$).

	Av. Est.	Emp. Var.	MSE
BOR1	5.004	0.036	0.036
BOR2	5.350	0.036	0.157
PS1	4.998	0.118	0.119
PS2	276.740	1.84E+07	1.85E+07
BDR1	5.008	0.046	0.046
BDR2	5.018	0.862	0.862
BDR3	5.360	0.946	1.074

<https://doi.org/10.1371/journal.pone.0221267.t001>

potential control sites we randomly sampled a total of 4787 points on the network within our eight administrative districts. The large ratio of potential control to treated units is adopted to ensure that we have a sufficient number of control units after we apply a matching algorithm.

Our outcome variable is the number of personal injury collisions (PICs) per kilometre as recorded from the location of the speed cameras, or in the case of control groups, from the randomly selected point. The PIC data are taken from records completed by police officers each time that an incident is reported to them. The individual police records are collated and processed by the UK Department for Transport as the 'STATS 19' data. The location of each PIC is recorded using the British National Grid coordinate system and can be located on a map using Geographical Information System (GIS) software. Because the established dates of speed cameras vary from 2002 to 2004, the period of analysis is from 1999 to 2007 to ensure the availability of collision data for the years before and after the camera installation for every camera site.

Covariates

To adequately adjust for confounding we require a set of measured covariates that adequately represent the characteristics of units that simultaneously determine treatment assignment and outcome. For the UK there exists a formal set of site selection guidelines for fixed speed cameras [5] that are extremely valuable in choosing covariates. The criteria are as follows

1. Site length: between 400-1500 m.
2. Number of fatal and serious collisions (FSCs): at least 4 FSCs per km in last three calendar years.
3. Number of personal injury collisions (PICs): at least 8 PICs per km in last three calendar years.
4. 85th percentile speed at collision hot spots: 85th percentile speed at least 10% above speed limit.
5. Percentage over the speed limit: at least 20% of drivers are exceeding the speed limit.

Criteria one to three are primary guidelines for site selection and criteria four and five are of secondary importance. There are sites that do not meet the above criteria that will still be selected as enforcement sites, mainly for reasons such as community concern and engineering factors.

Selection of the speed camera sites was primarily based on collision history. collision data can be obtained from the STATS 19 database and located on the map using GIS. However, secondary criteria such as the 85th percentile speed and percentages of vehicles over the speed limit are normally unavailable for all sites on UK roads. If speed distributions differ between the treated and untreated groups, then the failure to include the speed data could bias the estimation, an issue discussed in previous research [5, 15]. For untreated sites with the speed limit of 30 mph and 40 mph, the national average mean speed and percentages of speeding are similar to the data for the camera sites. The focus groups for this study are sites with the speed limit of 30 mph and 40 mph throughout the UK. It is reasonable to assume that there is no significant difference in the speed distribution between the treated and untreated groups and hence exclusion of the speed data will not affect the accuracy of the propensity score model.

It is also possible that drivers may choose alternative routes to avoid speed cameras sites. collision reduction at camera sites may include the effect induced by a reduced traffic flow. The benefits of speed cameras will therefore be overestimated without controlling for the change in traffic flow. The annual average daily flow (AADF) is available for both treated and

untreated roads and the effect due to traffic flow is controlled for in this study by including the AADF in the propensity score model.

In addition to the criteria that strongly influence the treatment assignment, factors that affect the outcomes should also be taken into account when the propensity score model is specified. We further include road characteristics such as: road types, speed limit, and the number of minor junctions within site length, which are suggested as important factors when estimating the safety impact of speed cameras [2, 6].

Component model specifications

The outcome variable of interest is the number of collisions per site. For the OR model the response is specified in differenced form, i.e. the number of collisions in the post-treatment period minus the number of collisions in the pre-treatment period. Differencing allows for the existence of unit level time-invariant effects, which could be random or fixed. The PS model is estimated using a logit Generalized Additive Mixed Model (GAMM) specification. Matching and overlap is achieved using nearest neighbour matching via the `MatchIt` package in R. The weighted OR model is then estimated on the trimmed dataset, which satisfies matching and overlap conditions, using a Gaussian GAMM specification. We use GAMMs for both models to avoid making a-priori assumptions on the functional form of the relationships under study.

As mentioned in the introduction, the DR approach is particularly attractive for our application because of the differences inherent in our PS and OR model specifications. Due to the existence of formal criteria for SC assignment we have a high degree of confidence in the ability of our covariates to eliminate confounding via the PS model. For the OR model, differencing of the response variable before and after treatment allows for the existence of site level time-invariant unobserved effects in addition to measured confounders. The DR model comprises the differenced OR model weighted by the estimated PS model. Thus, there are subtle differences in the way we model the ATE via the PS or OR approaches. A degree of robustness is offered using a DR approach since we will obtain a consistent estimate of the ATE if just one of the component models is well specified.

Results

The objective of our application is to estimate the *marginal* effect of SCs on RTCs, having adjusted for baseline confounders. We estimate the following models: an OR model, an inverse PS weighted model, a DR model comprising an OR model weighted with the inverse PS covariate (DR), and a naïve model which is simply the OR model without covariates. For the naïve model we report results using the matched and full samples. All models are repeatedly estimated using the approximate Bayesian approach outlined above. In addition to the posterior predictive distribution for the ATE we report point estimates at the mean of the posterior. For comparison, we also report Frequentist results.

The results are shown in Table 2 below including means and credible intervals of the ATE distributions. Our causal models (OR, IPW and DR) indicate that the presence of speed cameras corresponds with an average change in the number of RTCs of -14.4% to -15.5%. Note that the approximate Bayesian and Frequentist point estimates are very similar, which is what we would expect for linear models with uninformative priors. In comparison, the Naïve model which does not adjust for confounding, finds a higher ATE of -17.6% using the matched sample and -33.6% using the unmatched sample. Fig 1. below shows the posterior predictive distribution derived from the DR model.

Table 2. Bayesian and frequentist bootstrapped estimates of the average treatment effect.

	Bayesian bootstrap			Frequentist bootstrap	
	posterior mean	s.d.	95% cred. int.	Est.	s.e.
OR	-14.429	3.536	(-21.473, -7.350)	-14.825	3.419
IPW	-15.476	5.480	(-26.367, -4.808)	-15.504	4.935
DR	-14.359	3.605	(-21.841, -7.352)	-14.370	3.494
Naïve (matched sample)	-17.617	5.199	(-28.214, -7.800)	-17.887	5.118
Naïve (full sample)	-33.684	5.911	(-42.133, -13.624)	-34.682	3.573

<https://doi.org/10.1371/journal.pone.0221267.t002>

Predictive Posterior Distribution of ATE

Fig 1. Predictive posterior distribution of the average treatment effect from the doubly-robust model.

<https://doi.org/10.1371/journal.pone.0221267.g001>

Thus, it would appear that correcting for potential sources of confounding serves to reduce the magnitude of our ATE estimates, but we still find a substantial reduction in RTCs associated with presence of speed cameras. The difference in estimated ATE between the naïve and causal models makes sense given that the formal criteria used to assign SCs favours sites that have exhibited high rates of collisions in the past. Crucially, our causal models imply that SCs do make a real difference to RTCs over and above the modelled effect of confounding from non random assignment.

Conclusions

In this paper we have quantified the causal effect of speed cameras on road traffic collisions via an approximate Bayesian doubly robust approach. This is the first time such an approach has been applied to study road safety outcomes. The method we propose could be used more generally for estimation of crash modification factor (CMF) distributions. Simulations demonstrate that the approach is doubly-robust for average treatment effect estimation. Our results indicate that speed cameras do cause a significant reduction in road traffic collisions, by as much as 15% on average for treated sites. This is an important result that could help inform public policy debates on appropriate measures to reduce RTCs. The adoption of evidence based approaches by public authorities, based on clear principles of causal inference, could vastly improve their ability to evaluate different courses of action and better understand the consequences of intervention.

There are thus two important implications of our study that could ultimately improve highway safety. First, is that such inference could be employed to achieve a more effective assignment of SCs and consequent reduction of RTCs. Second, the approach outlined above could be used to continually monitor SC effectiveness as baseline conditions (e.g. related to road traffic and wider demographic and social characteristics) change, thus providing a mean of monitoring the effectiveness of road safety interventions dynamically.

Author Contributions

Conceptualization: Daniel J. Graham.

Data curation: Daniel J. Graham, Haojie Li.

Formal analysis: Daniel J. Graham, Cian Naik, Emma J. McCoy, Haojie Li.

Investigation: Daniel J. Graham, Cian Naik.

Methodology: Daniel J. Graham, Cian Naik, Emma J. McCoy.

Supervision: Daniel J. Graham.

Writing – original draft: Daniel J. Graham.

Writing – review & editing: Daniel J. Graham.

References

1. Li H, Graham D, Majumdar A. The impacts of speed cameras on road accidents: an application of propensity score matching methods. *Accident Analysis & Prevention*. 2013; 60:148–157. <https://doi.org/10.1016/j.aap.2013.08.003>
2. Christie S, Lyons R, Dunstan F, Jones S. Are mobile speed cameras effective? A controlled before and after study. *Injury Prevention*. 2003; 9:302–306. <https://doi.org/10.1136/ip.9.4.302> PMID: 14693888
3. Cunningham C, Hummer J, Moon J. Analysis of Automated Speed Enforcement Cameras in Charlotte, North Carolina. *Transportation Research Record*. 2000; 2078:127–134. <https://doi.org/10.3141/2078-17>

4. De Pauw E, Daniels S, Brijs T, Hermans E, Wets G. An evaluation of the traffic safety effect of fixed speed cameras. *Safety Science*. 2014; 62:168–174. <https://doi.org/10.1016/j.ssci.2013.07.028>
5. Gains A, Heydecker B, Shrewsbury J, Robertson S. The national safety camera programme 3-year evaluation report. UK Department for Transport; 2004.
6. Gains A, Heydecker B, Shrewsbury J, Robertson S. The national safety camera programme 4-year evaluation report. UK Department for Transport; 2005.
7. Goldenbeld C, van Schagen I. The effects of speed enforcement with mobile radar on speed and accidents. An evaluation study on rural roads in the Dutch province Friesland. *Accident Analysis & Prevention*. 2005; 37:1135–1144. <https://doi.org/10.1016/j.aap.2005.06.011>
8. Jones AP, Sauerzapf V, Haynes R. The effects of mobile speed camera introduction on road traffic crashes and casualties in a rural county of England. *Journal of Safety Research*. 2008; 39:101–110. <https://doi.org/10.1016/j.jsr.2007.10.011> PMID: 18325421
9. Maher M. A note on the modelling of TfL fixed speed camera data. University College London; 2015.
10. Hauer E, Harwood DW, Council FM, Griffith MS. Estimating safety by the empirical Bayes method: a tutorial. *Transportation Research Record*. 2002; 1784:126–131. <https://doi.org/10.3141/1784-16>
11. Chen G, Meckle W, Wilson J. Speed and safety effect of photo radar enforcement on a highway corridor in British Columbia. *Accident Analysis & Prevention*. 2002; 34:129–138. [https://doi.org/10.1016/S0001-4575\(01\)00006-9](https://doi.org/10.1016/S0001-4575(01)00006-9)
12. Elvik R. Effects on accidents of automatic speed enforcement in Norway. *Transportation Research Record*. 1997; 1595:14–19. <https://doi.org/10.3141/1595-03>
13. Hoyer A. Safety effects of fixed speed cameras—an empirical Bayes evaluation. *Accident Analysis & Prevention*. 2015; 82:263–269. <https://doi.org/10.1016/j.aap.2015.06.001>
14. Mountain LJ, Hirst WM, Mahar MJ. Costing lives or saving lives: a detailed evaluation of the impact of speed cameras. *Traffic Engineering & Control*. 2004; 45:280–287.
15. Mountain LJ, Hirst WM, Mahar MJ. Are speed enforcement cameras more effective than other speed management measures? The impact of speed management schemes on 30 mph roads. *Accident Analysis & Prevention*. 2005; 37:742–754. <https://doi.org/10.1016/j.aap.2005.03.017>
16. Shin K, Washington SP, van Schalkwyk I. Evaluation of the Scottsdale Loop 101 automated speed enforcement demonstration program. *Accident Analysis & Prevention*. 2009; 41:393–403. <https://doi.org/10.1016/j.aap.2008.12.011>
17. Carnis L, Blais E. An assessment of the safety effects of the French speed camera program. *Accident Analysis & Prevention*. 2013; 51:301–309. <https://doi.org/10.1016/j.aap.2012.11.022>
18. Hess S, Polak J. Effects of speed limit enforcement cameras on accident rates. *Transportation Research Record*. 2003; 1830:25–34. <https://doi.org/10.3141/1830-04>
19. Keall MD, Povey LJ, Frith WJ. The relative effectiveness of a hidden versus a visible speed camera programme. *Accident Analysis & Prevention*. 2001; 33:277–284.
20. Robins JM. Robust estimation in sequentially ignorable missing data and causal inference models. In: *Proceedings of the American Statistical Association, Section on Bayesian Statistical Science*. Alexandria, VA: American Statistical Association; 2000. p. 6–10.
21. Robins JM, Rotnitzky A, van der Laan MJ. Comment on the Murphy and Van der Vaart article “On profile likelihood”. *Journal of the American Statistical Association*. 2000; 95:431–435.
22. Robins JM, Rotnitzky A. Comment on “Inference for semiparametric models: some questions and an answer”. *Statistical sinica*. 2001; 11:920–936.
23. van der Laan M, Robins JM. *Unified methods for censored longitudinal data and causality*. Berlin: Springer; 2003.
24. Lunceford JK, Davidian M. Stratification and weighting via the propensity score in estimation of causal treatment effects: a comparative study. *Statistics in Medicine*. 2004; 23:2937–2960. <https://doi.org/10.1002/sim.1903> PMID: 15351954
25. Bang H, Robins JM. Doubly robust estimation in missing data and causal inference models. *Biometrics*. 2005; 61:962–972. <https://doi.org/10.1111/j.1541-0420.2005.00377.x> PMID: 16401269
26. Kang JDY, Schafer JL. Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data. *Statistical Science*. 2007; 22(4):523–539. <https://doi.org/10.1214/07-STS227>
27. DfT. Reported road casualties in Great Britain: quarterly provisional estimates. London: UK Department for Transport; 2017.
28. DfT. Transport Statistics Great Britain: 2016. London: UK Department for Transport; 2016.

29. Rubin DB. Bayesian inference for causal effects: the role of randomization. *Annals of Statistics*. 1978; 6(1):34–58. <https://doi.org/10.1214/aos/1176344064>
30. Rubin DB. Comment on 'Randomization analysis of experimental data in the Fisher randomization test' by Basu. *Journal of the American Statistical Association*. 1980; 75(371):591–593.
31. Rubin DB. Comment: which ifs have causal answers? *Journal of the American Statistical Association*. 1986; 81(396):961–962.
32. Rubin DB. Neyman (1923) and causal inference in experiments and observational studies. *Statistical Science*. 1990; 5(4):472–480. <https://doi.org/10.1214/ss/1177012032>
33. Tsiatis AA, Davidian M. Comment: Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data. *Statistical Science*. 2007; 22(4):569–573. <https://doi.org/10.1214/07-STS227> PMID: 18516239
34. Horvitz DG, Thompson DJ. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*. 1952; 47:663–685. <https://doi.org/10.1080/01621459.1952.10483446>
35. Graham DJ, McCoy EJ, Stephens DA. Approximate Bayesian Inference for Doubly Robust Estimation. *Bayesian Anal*. 2016; 11(1):47–69. <https://doi.org/10.1214/14-BA928>
36. Rubin DB. The Bayesian Bootstrap. *The Annals of Statistics*. 1981; 9(1):130–134. <https://doi.org/10.1214/aos/1176345338>
37. Newton MA, Raftery AE. Approximate Bayesian Inference with the Weighted Likelihood Bootstrap (with discussion). *Journal of the Royal Statistical Society Series B (Methodological)*. 1994; 56(1):pp. 3–48.
38. Muliere P, Secchi P. Bayesian nonparametric predictive inference and bootstrap techniques. *Annals of the Institute of Statistical Mathematics*. 1996; 48(4):663–673. <https://doi.org/10.1007/BF00052326>