



**Canadian
Institute
of Actuaries**

**Institut
canadien
des actuaires**



**CASUALTY
ACTUARIAL
SOCIETY**

Assessing the Impact of Marijuana Decriminalization on Vehicle Accident Experience

Vyacheslav Lyubchich

December 2022

Document 222163

Table of contents

Executive summary	3
1 Introduction	4
2 Literature overview	5
Marijuana and decriminalization impacts	5
Broader view	5
3 Data overview	6
3.1 Canada	6
3.2 United States.....	8
4 Methods	10
4.1 Before-after comparisons	10
4.2 Propensity score matching	11
4.3 Regression random forest.....	12
5 Results	13
5.1 Canada	13
5.2 United States.....	20
6 Bibliography	23
7 About the author	25



Executive summary

This report addresses a request by the Canadian Institute of Actuaries (CIA) and Casualty Actuarial Society (CAS) for analysis of the impact of marijuana decriminalization on vehicle accident experience. The literature review shows that while marijuana impairment affects driving behaviour, the behaviour is not always riskier; for example, slower speeds and longer following distances of impaired drivers have been reported. The observational studies of road accidents report mixed results, most often not detecting significant effects, particularly in the long term.

This report utilizes technical advancements of machine learning algorithms for finding patterns in data, rigorous approaches for selecting control units from observational data, and recently developed pathways for a causal interpretation of black-box models.

Canadian and US data for 2016–2019 were used in the study, including official reports on collisions of private vehicles and losses in Canada, fatal accidents, and weather factors in the United States. For each data source, statistical and machine learning models were chosen to account for different sources of variability.

The analysis of 10 Canadian regions accounted for the regional differences and modelled a baseline linear trend that was also observed in the pre-legalization data alone. The analysis showed no statistically significant changes in the average cost per claim and claim frequency after marijuana legalization in Canada. The quarterly data available for Québec led to similar findings.

The US results varied by the state pair selected as the control in statistical comparisons. The tests for the decriminalization effect on fatalities failed to detect a statistically significant change. The machine learning techniques also allowed the author to account for other factors, including the weather, and annual and weekly patterns of fatalities.

Overall, this study overcame the disadvantages of earlier research on the effects of marijuana decriminalization by incorporating novel methodologies that do not rely on the linearity of relationships or parametric inference. The study did not detect statistically significant persistent impacts of decriminalization.



1 Introduction

Private cars are essential for many people living in the United States and Canada, and have been an attribute of material prosperity postulated in the idea of the American dream. The insurance risks related to cars have been affected in recent years by regional changes in the legal rules on using marijuana. Decriminalization of marijuana in several regions was associated with a higher number of drivers operating vehicles under the influence of the drug (Pollini et al. 2015) and more car crash fatalities (Cook et al. 2020). While the debates and decisions in the process of decriminalization of marijuana continue, insurance companies need to assess the associated changes in vehicle accident experience to price the risk and adjust the cost of insurance.

There have been mixed reports about the implications of decriminalization of marijuana (e.g., see a brief overview by Hall and Lynskey 2020), which could be explained in part by the fact that the studies consider only one region (data limitation) or rely on only one statistical method. Some other publications consider relative change without the overall impact (e.g., Pollini et al. 2015 studied the proportion of drivers who tested positive for delta-9-tetrahydrocannabinol [THC], not the changes in the number or severity of car accidents).

This study proposes data-driven and reliable tools for assessing the impact of changes in marijuana laws on vehicle accident experience, by harnessing the publicly available data with the power of modern statistical and machine learning techniques under the integrative framework for impact analysis and uncertainty quantification. The developed framework will be applicable for analysis of other policy changes, such as decriminalization of psilocybin or future adoption of self-driving vehicles. The study provides new insight into the impacts of marijuana decriminalization on the vehicle accident experience and the implications for the insurance industry, and develops a systematic approach for assessing similar types of changes in the future. This allows insurers to hedge the related risks and develop more effective pricing strategies for insurance products. Thus, the objectives of this study are to provide robust conclusions by analyzing as many regions in Canada and the United States as possible and by approaching the analysis using multiple modern nonparametric techniques from statistics and machine learning domains.

The key steps are summarized below:

1. Establish relevant “controls” for quantifying the effect of marijuana decriminalization. The control could be the same territory (before–after, if quality time series of sufficient length are available for both periods) or similar territories for concurrent comparisons.
2. Collect car accident data including frequency, severity, and type of accidents. Collect information about confounding conditions such as weather, time of the day and day of the week (e.g., see Bailey et al. 2020), and driver’s age, sex, and impairment.
3. Apply statistical techniques to quantify the effect of marijuana decriminalization while controlling for the confounding factors.
4. Expand the set of traditional linear and parametric models by harnessing the power of modern nonparametric statistical and machine learning techniques for regularizing the model coefficients, empirically calibrating p -values and confidence intervals using bootstrap procedures, and building flexible nonlinear models with random forests and artificial neural networks (e.g., see Peters et al. 2017, Zhao and Hastie 2021 for promising results of such models in causal inference).
5. Summarize the results across different methods and geographic regions.

The approaches used in the study are data-driven and can be used for assessing the impact of other types of step-wise changes, such as decriminalization of other substances or the adoption of self-driving cars.



2 Literature overview

Marijuana and decriminalization impacts

Some studies investigated (1) impacts of marijuana on the ability to drive (controlled and observational studies), (2) the impact of decriminalization on the proportion of drivers who tested positive for marijuana use (observational studies), and (3) the impact of decriminalization on public health and safety outcomes (observational studies). In the first group of studies, Gelmi et al. (2021) used a controlled experiment with 33 volunteers in Switzerland (all smokers or with smoking experience) and detected no significant impact on the driving ability of subjects who smoked cannabidiol (CBD)-rich marijuana, while all those participants exceeded the Driving Under the Influence of Drugs (DUID) legal limit for THC in blood after smoking CBD-rich marijuana. Hartman et al. (2016) reported the results of another experiment (19 volunteers in Iowa, USA), which showed slower speeds and longer following distances by the subjects who were impaired with THC, suggesting that drivers were aware of the impairment and were more cautious. On the other hand, Bédard et al. (2007) used Fatality Analysis Reporting System (FARS) data for 1993–2003 for drivers aged 20–50 and found that drivers who tested positive for cannabis (5% of drivers) had higher odds ratios for demonstrating “driver-related factors” (40 factors coded in police reports, including speeding, operating erratically, and failure to yield or obey traffic signs).

In the second group of studies, Chung et al. (2019) reported a higher proportion of marijuana detection in patients admitted to Colorado hospitals with a traumatic injury after marijuana decriminalization in the state. Pollini et al. (2015) used FARS data for California (2008–2012) and detected a statistically higher proportion of CBD’s prevalence in 2012 compared to a pre-decriminalization period of 2008–2010. Some studies were able to supplement FARS data with non-crash data. Romano et al. (2017) summarized the results of two studies (Li et al. 2013, Romano et al. 2014) that used very similar data, from FARS (fatal cases) and the 2007 National Roadside Survey (NRS; non-fatal controls), and show that while crude odds ratios show a significant impact of marijuana impairment, the adjusted ratios (adjusted for age, sex, and race) in all considered subsets of states no longer provide evidence of the contribution of marijuana to a crash.

While the first two groups of studies lay down important groundwork for investigating marijuana impacts, this project and the third group of studies primarily focus on estimating aggregated effects on crashes and fatalities. The studies in this group report mixed results, most often not detecting the effects of marijuana decriminalization on road crashes (Hall and Lynskey 2020). Cook et al. (2020) studied semi-annual 2010–2017 FARS data from cities with a population greater than 100,000 located in states that did not enact medical marijuana laws by 2010. Using the difference-in-differences approach, this study identified a slight reduction (9%) in fatal crashes after the adoption of medical marijuana use in the states, but no effect of decriminalization was detected, except for young male drivers (13% increase in fatal crashes). Lane and Hall (2019) used interrupted monthly time-series analysis (2009–2016) and detected a short-term increase in FARS-reported fatalities in Colorado, Washington, Oregon and their neighbouring jurisdictions (an increase of 1 fatality per million residents in the year following marijuana legalization). Farmer et al. (2022) studied quarterly crash rates per mile of travel in 11 states during 2009–2019, while accounting for time effects, seat belt use rate, unemployment, and alcohol use rates. The authors concluded that estimates of marijuana legalization effects are not always statistically significant and vary by state. Windle et al. (2021a) and Windle et al. (2021b) extend to Canada their finding obtained from the US FARS database about effects of decriminalization on fatal collision rates and fatalities. Windle et al. (2021a) find evidence of increased fatal collision rates when the model adjusted for linear time trend is used; however, a comment on the paper posted online by Jiang et al. reports the used Poisson model as a poor fit. Finally, NHTSA DOT (2017) provides a great overview of the related research and existing practical questions related to drug testing processes, devices used, and training of personnel.

This study brings together the technical advancements of machine learning algorithms for finding patterns in data, rigorous approaches for selecting control units from observational data, and recently developed pathways for a causal interpretation of black-box models. A combination of all these steps has not appeared in a single study on the effects of marijuana decriminalization before. Specifically, this study attempts to overcome the disadvantages of earlier studies that relied on the assumed linearity (the author suggests using nonlinear models where appropriate, such as random forests) and parametric inference (the author suggests using bootstrapping to derive confidence intervals without relying on a particular distribution).

Broader view

There are other open questions regarding marijuana-related research, which may affect the way people perceive the drug, recognize different levels of tolerance, and even test people for being influenced by the drug. For example, Chapter 12 of Potter and Weinstock (2019) compares marijuana with the tobacco, alcohol, and drug (opioids) industry, while Ammerman et al. (2015) provide a broad overview of international experiences and argue that some effects of marijuana, such as addiction, are overlooked compared with alcohol and tobacco. Alexander (2003) argues

about the inadequacy of existing marijuana screening methods. Brands et al. (2021) provide a recent overview of related questions, regarding the dose–response relationship (concentration of THC in blood and actual driving behaviour; uncertainty about which THC concentration consistently results in impairment), tolerance after frequent use, effects of CBD:THC ratios, and different routes of taking the drug (smoke/vapor/orally).

3 Data overview

To avoid confounding with the impacts of the COVID-19 pandemic, data from 2020 or later are not used. More specifically, there are no provinces in Canada that can act as a control group for estimating the effect of marijuana decriminalization (see Section 3.1); hence the effect of the COVID-19 pandemic would be confounded with other effects. In the United States, states serve as the controls for estimating the decriminalization effects (see Section 3.2) and some researchers argue that one does not need to explicitly adjust for confounders when a control group is used (Elvik 2002). However, the impacts of COVID-19, such as the mobility changes due to imposed lockdowns of different stringency, have varied considerably within states (Barnes et al. 2020), across states (Lin et al. 2021), and globally (Gupta et al. 2021, Rahman et al. 2021, Yasin et al. 2021). Moreover, the recovery of human mobility after the relaxation of the lockdowns was also spatially heterogeneous. Such trends pose additional difficulties if included in the study (e.g., see Choi et al. 2018) since confounding factors may affect sites differently in time (Smith et al. 1993).

3.1 Canada

The legalization of recreational use of marijuana took effect across Canada on October 17, 2018.¹ Historically (before 2018), recreational use of marijuana in the province of British Columbia had been tolerated more than in the rest of the country;² however, due to substantial differences with other provinces, the other provinces cannot serve as good controls for studying the impacts of marijuana use in British Columbia.

The data used in this study come from annual reports of the General Insurance Statistical Agency³ (GISA), which collects insurance information from most Canadian regions (except British Columbia, Manitoba, Québec, and Saskatchewan; see Exhibits AUTO1003, AUTO1005, and AUTO1010 on the GISA website), and Groupement des Assureurs Automobiles⁴ (GAA), which provides the information for Québec (see Exhibit 1A.2 in the Summary Report on the GAA website). From these reports, information on the *collision of private vehicles* per accident year was extracted. The claims in this category are characterized by loss development factors close to 1; hence they are not expected to change substantially as the losses from accidents that happened in the most recent years get finalized. Dollar amounts were adjusted to the prices of 2019 to account for the changing value of money, using the Consumer Price Index (all item groups) by Statistics Canada.⁵

Figure 1 shows the dynamics of the studied variables by region. The regions differ in population; hence the annual number of earned vehicles was used to weight the observations in further analyses (Figure 2).

¹ <https://www.sencanada.ca/en/sencaplus/news/cannabis-act> (accessed July 11, 2021)

² https://en.wikipedia.org/wiki/Cannabis_in_British_Columbia (accessed July 11, 2021)

³ <https://www.gisa.ca/> (accessed July 1, 2021)

⁴ <https://gaa.qc.ca/en/statistics/at-a-glance/> (accessed July 1, 2021)

⁵ Statistics Canada. Table 18-10-0005-01 Consumer Price Index, annual average, not seasonally adjusted, <https://doi.org/10.25318/1810000501-eng> (release date January 20, 2021)




Figure 1: Annual average cost per claim (in 2019 dollars) and frequency of collisions. The dashed line denotes the year of marijuana decriminalization.




Figure 2: Number of earned vehicles during 2016–2019 used as the regional weights. The lower and upper hinges correspond to the first and third quartiles.

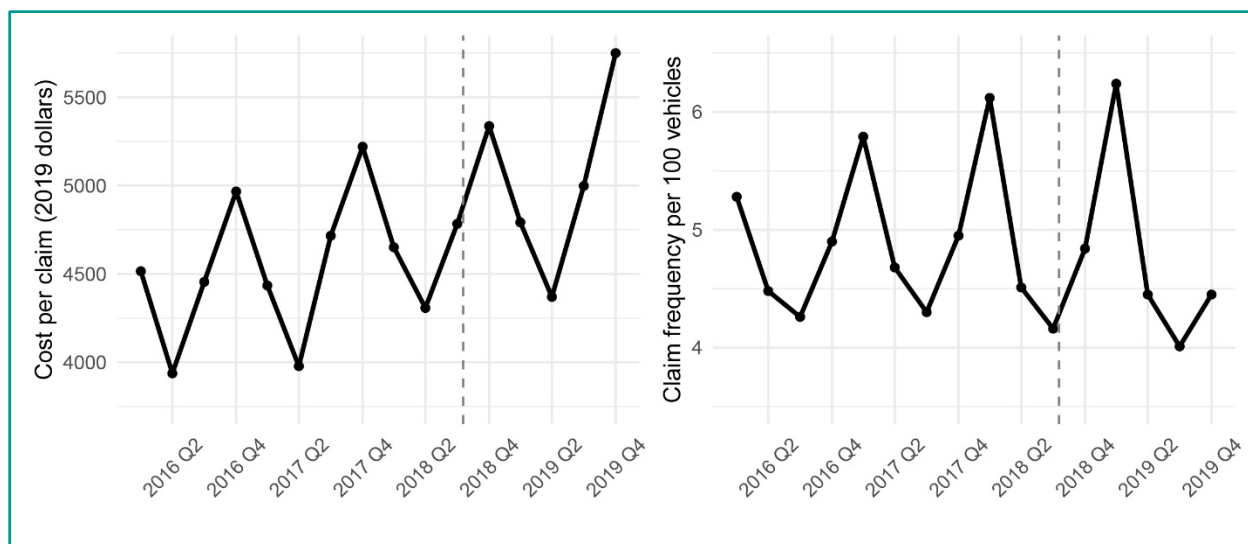


Figure 3: Quarterly average cost per claim (in 2019 dollars) and frequency of collisions in Québec. The dashed line denotes the date of marijuana decriminalization.

3.2 United States

In the United States, the legality of marijuana varies by jurisdiction. Table 1 shows a summary of legalization history by the state during the study period for which the accident data are available (see below). The following state statistics were used to select the control group states: urbanization (percentage of the total population),⁶ population,⁷ vehicle miles per licensed driver,⁷ road miles per 1000 persons (calculated as the total road and street mileage divided by the population),⁷ and vehicles per 1000 people.⁸ Number of licensed drivers was also considered, but correlated too strongly ($r = 0.99$) with population.

To study US auto accidents, the Countrywide Traffic Accident Dataset (CTAD) v. 4 is used (Moosavi et al. 2019). The CTAD covers the contiguous United States from February 2016 to December 2020 (about 4.2 million accidents). However, more detailed analysis suggests that spatiotemporal data are not equally well represented and some states lack data at the beginning of the sample period. Nevertheless, the advantages of this dataset include high spatial and temporal resolutions and universal data collection across the United States using application programming interfaces (APIs) of MapQuest Traffic and Microsoft Bing Map Traffic. A disadvantage of this dataset is that severity of accidents is measured in terms of how severely the traffic flow was affected, not in terms of property damage or personal injuries.

The CTAD also contains information on confounding variables such as weather conditions recorded at a nearby airport at the approximate time of the accident. However, to incorporate weather impacts into models, a complete weather history (not just at the time of accidents) is required. Air temperature and precipitation information was obtained from two alternative sources:

1. Hourly data were obtained from the US ASOS network of automated airport weather observations,⁹ using airport codes from the CTAD. In the cases when an airport from the CTAD could not be found, its location was approximated using latitudes and longitudes of reported accidents, and geodesic distances were used to replace the missing airport with a nearby airport represented in the ASOS database.
2. Daily gridded data from Daymet,¹⁰ where grid cells were selected based on coordinates of airport weather stations in the ASOS network.

The beginning of the timescale in Table 1 (early 2016) is defined by the availability of data in the CTAD. Since some states (including Colorado and Washington) had already legalized marijuana by 2016, they cannot be used for the

⁶ https://en.wikipedia.org/wiki/Urbanization_in_the_United_States (accessed November 1, 2021)

⁷ <https://www.fhwa.dot.gov/ohim/onh00/onh2p11.htm> (accessed December 25, 2021)

⁸ https://en.wikipedia.org/wiki/List_of_U.S._states_by_vehicles_per_capita (accessed November 1, 2021)

⁹ <https://mesonet.agron.iastate.edu/request/download.phtml?network=AWOS> (accessed January 2, 2022)

¹⁰ Daymet: Daily Surface Weather Data on a 1-km Grid for North America, Version 4 <https://doi.org/10.3334/ORNLDAAAC/1840> (accessed January 2, 2022)

before–after comparisons of the frequency of car accidents in this study. The end of the timescale in Table 1 (December 2019, same as for Canada) was chosen to avoid confounding the estimated effects with the impacts of the COVID-19 pandemic. State-licensed sales of recreational cannabis in Michigan began on December 1, 2019 (one month before the end of the study period). For all the Before and After periods to be of at least one year long for more reliable conclusions, Michigan was removed from consideration.

In addition to the CTAD, the FARS¹¹ of the National Highway Traffic Safety Administration (NHTSA) of the United States Department of Transportation (DOT) was used to extract data for this study.

The database was supplemented with such variables as Year (numeric), Month (1–12, categorical), Weekday (1–7, categorical), Weekend (0/1, categorical), Holiday (0/1, categorical, based on Wuertz et al.’s 2022 holiday calendar for the New York Stock Exchange).

The accident experience metrics were normalized by the number of registered vehicles by year and state, available from <https://www.fhwa.dot.gov>.

Table 1: Groups of states for studying cannabis legalization in the United States (from March 2016 to December 2019). The dates when the commercial sales started are used in analysis and are shown in parentheses.

Group	States
Legalized	CA (January 1, 2018), MA (November 20, 2018), NV (July 1, 2017)
Fully illegal	ID, KS, NE, NC, SC, TN, WY
Illegal recreational	AL, AR, FL, IN, IA, KY, OK, TX, UT, WV, WI
Legalized after the study period	AZ, MT, NJ, NY, SD, VA
Mixed (unused in the study)	AK, CO, CT, DE, DC, GA, HI, IL, LA, ME, MD, MI, MN, MS, MO, NH, NM, ND, OH, OR, PA, RI, VT, WA

¹¹ <https://www.nhtsa.gov/research-data/fatality-analysis-reporting-system-fars> (accessed July 1, 2021)

4 Methods

Use two-tailed tests with significance level $\alpha = 0.05$ unless noted otherwise ($1 - \alpha = 0.95$ or 95% corresponds to the confidence level). In other words, the results are statistically significant when the corresponding p -values are below α .

4.1 Before-after comparisons

For simple before-after comparisons when no concurrent control is available (e.g., for Canada), the following regression models with parametric inference and residual (semi-parametric) bootstrap are applied.

Parametric models

Consider a mixed-effects model (Zuur et al. 2009)

$$y_{tp} = a + bX_t + ct + \beta_p + e_{tp}, \quad (1)$$

where y is the response variable (claim frequency or average cost per claim), t is the numeric time index (year), p is the space index (province), a is the grand intercept, b and c are the fixed-effect coefficients, X is the indicator whether t is before the marijuana decriminalization date or after, $\beta_p \sim N(0, \delta^2)$ are random intercepts accounting for different average levels of the response variable by province, and $e_{tp} \sim N(0, \sigma^2)$ are the model residuals. Weights w_{tp} can be applied to the residual variances for weighted estimation, such as $e_{tp} \sim N(0, \sigma^2 \times w_{tp})$ (Section 4.1.2 in Zuur et al. 2009). In particular, $w_{tp} = (\text{Number of earned vehicles})_{tp}^{-1}$ can be used to assign bigger weights to observations corresponding to more earned vehicles. We are interested in estimating and testing the statistical significance of the coefficient b that summarizes the effects of marijuana decriminalization.

The inclusion of a time trend in the regression model is one of the ways to account for the trend and hence avoid spurious results (e.g., see Section 10.5 in Wooldridge 2013). This is possible due to the conditional interpretation of coefficients in multiple regression that might differ from marginal relationships (e.g., see Section 1.3.1 in Chatterjee and Simonoff 2013). For example, coefficient b in the model (1) can be interpreted as the expected change in the claim frequency with the marijuana decriminalization, holding everything else in the model fixed. In this analysis, the linear trend is appropriate for capturing the trending behaviour of time series (based on Figure 1) and hence is a preferred option of regression analysis of nonstationary series, compared with using in the regression model detrended differenced series (e.g., $\Delta y_{t,p} = y_{t,p} - y_{t-1,p}$) or deviations from an estimated trend \tilde{y}_{tp} (i.e., $\tilde{e}_{tp} = y_{tp} - \tilde{y}_{tp}$).

For single-province seasonal data, consider the model

$$y_t = a + bX_t + ct + \sum_{i=1}^{s-1} d_i S_{it} + e_t, \quad (2)$$

where S_{it} is an indicator variable that takes on the value of 1 if time t belongs to the season i and the value of 0 otherwise ($i = 1, \dots, 4$ for quarterly data; for identifiability reasons only 3 coefficients d_i are estimated). In other words, model (2) includes a linear trend with additive seasonality modelled using a categorical variable for the seasons (quarters).

Bootstrap

To make the inference based on the mixed-effects models robust to non-normality, implement the semi-parametric residual bootstrap algorithm (Carpenter et al. 2003, Loy et al. 2022):

1. Estimate parameters and error terms for the model.
2. Sample independently with replacement from the sets of error terms.
3. Rescale the bootstrapped error terms so that their empirical variance is equal to the model estimates.
4. Obtain bootstrap samples by combining the bootstrapped samples via the fitted model equation.
5. Refit the model and extract the statistics of interest.
6. Repeat steps 2–5 B times.

From the bootstrapped distributions, obtain confidence intervals for the desired confidence level $1 - \alpha$, as defined in Chapter 5 in Davison and Hinkley (1997). For the estimate \hat{b} and its bootstrap counterparts \hat{b}^* , the basic bootstrap confidence limits are

$$2\hat{b} - \hat{b}_{[(B+1)(1-\alpha/2)]}^*, \quad 2\hat{b} - \hat{b}_{[(B+1)\alpha/2]}^* \quad (3)$$

the percentile limits are

$$\hat{b}_{[(B+1)\alpha/2]}^*, \quad \hat{b}_{[(B+1)(1-\alpha/2)]}^* \quad (4)$$

and the limits with normal approximation are

$$\hat{b} - v^{1/2}z_{1-\alpha/2}, \quad \hat{b} + v^{1/2}z_{1-\alpha/2} \quad (5)$$

where v is the approximate variance and z is the quantile of $N(0,1)$ distribution.

4.2 Propensity score matching

When randomized controlled trials cannot be implemented, propensity score matching (PSM) is used to match study subjects by estimated probability for them to receive the treatment, based on some baseline covariates X_1, X_2, \dots, X_k that are not affected by the treatment (Austin 2011, Olmos and Govindasamy 2015, Zhao et al. 2021). Hence, PSM attempts to create homogeneous groups to potentially avoid confounding the studied effects with the effects of the baseline covariates. As a result, given similar propensity scores, the assignment of subjects to the treatment can be treated as random, which is important for satisfying assumptions of many statistical models and causal inference (Damrongplasit et al. 2010). For example, it is desirable to compare the state with decriminalized marijuana with the state(s) that have similar baseline characteristics such as population and urbanization, to minimize the effect of these characteristics on the differences in vehicle accident experience between the states. Relevant examples of implementing PSM include Lopez Bernal et al. (2018), Damrongplasit et al. (2010), and Stuart et al. (2014).

The propensity scores are the conditional probabilities of being assigned a treatment, which are estimated from the model

$$Treatment \sim X_1 + X_2 + \dots + X_k,$$

which is, essentially, a classification model with the assigned treatment (decriminalized or not) being the response variable.¹² In this study, the covariates X_1, X_2, \dots, X_k are represented by

1. urbanization,
2. population,
3. vehicle miles per licensed driver,
4. road miles per 1000 persons, and
5. vehicles per 1000 people.

Binary logistic regression is the most widely used form of this classifier, but machine learning techniques allow more flexibility in modelling nonlinear relationships and interactions (combined effects) of several covariates. From the tree-based methods and neural networks, random forest is a decision-tree-based method that delivers competitive performance and, at the same time, has only a few hyperparameters that the user has to set (Hastie et al. 2009). Hence, the random forest classifier is used in this study to predict the probability of a state being assigned the treatment.

In the implemented random forest algorithm, the original dataset is resampled with replacement 500 times, and a classification tree is grown on each such bootstrap sample. Each tree uses the covariates X_1, X_2, \dots, X_k to split the data into smaller subsets, homogeneous based on values of the response variable. For additional robustness (decorrelation of trees), each split of the tree uses only a random subset of \sqrt{k} covariates from the total of k covariates. Each tree of the forest is then used for predictions, and the proportion of trees that classify the state as decriminalized is used as that treatment probability or the propensity score. For more details on the random forest algorithm, see Breiman (2001) and Chapter 15 in Hastie et al. (2009).

At the matching step, the propensity scores are used to select appropriate control subjects from a sufficient number of untreated units. The matching techniques vary from identifying the nearest neighbours to optimization and

¹² Another way of thinking about PSM is that it is a dimensionality reduction technique allowing one to employ k covariates at once for finding subjects that are close to the treated subjects in this k -dimensional space (Olmos and Govindasamy 2015).

clustering algorithms. Matching with replacement allows a subject from the untreated group to be matched with more than one subject from the treated group. Several matching methods are usually applied since the ultimate results might be sensitive to the selection of control subjects (Zhao et al. 2021). For example, Damrongplasit et al. (2010) used stratified matching where a range of propensity scores with both treated and untreated subjects was first selected, then separated into smaller intervals (each such interval also contained both types of subjects). In each such stratum, the average treatment effect was estimated to assess the variability of the estimates. In this study, the nearest-neighbour selection is used with varying numbers of control subjects to assign to each treated subject. With 1:1 matching, the selected matches are most similar in terms of their propensity scores, but larger proportions can also be used for better estimates for the counterfactual (Olmos and Govindasamy 2015). The standardized mean differences are typically used to assess the balance of the covariates in the adjusted treatment and control groups, with the rule-of-thumb threshold of 0.1 or 0.25 (Zhao et al. 2021).

Conventional matching

As an alternative to propensity scores, the proximity of states in the k -dimensional space of covariates can be used to select similar states, without relating the covariates to the decriminalization outcomes. Sometimes this approach is called “conventional matching” (Zhao et al. 2021). In such an unsupervised case, a random forest can still be used to extract a proximity matrix (Shi and Horvath 2006) to identify neighbors in the k -dimensional space. Proximity in a random forest is quantified as follows: if a pair of out-of-bag observations (not used to construct a tree, since due to sampling with replacement some observations are left out) end up in the same terminal node of the tree, their proximity is increased by 1 (see Section 15.3 in Hastie et al. 2009).

4.3 Regression random forest

The random forest method has been introduced in the section on PSM as a classifier (to predict the assignment of a unit to treatment; the response variable is binary) and as a tool to find distances in high-dimensional space for a more conventional matching of states based on their proximity in that space (Breiman 2001, Hastie et al. 2009). Random forests can be used also in a regression setting when the modelled response variable is numeric, such as the car accident rate (e.g., see Bailey et al. 2020).

Random forests have just a few hyperparameters including the number of trees, number of variables randomly selected to select the best split, and minimal size of the subsets (aka terminal nodes of a regression tree) that do not get split anymore. The number of trees affects the stability of results but as long as it is relatively large the estimation error stabilizes. In this study, 500 trees are used in each random forest, and the two other hyperparameters are tuned using ten-fold cross-validation minimizing the root mean squared error (RMSE).

Since random forests consist of a large number of regression trees, the direct interpretability of the model is lost (the method becomes a black box). However, post-processing of the trained random forest allows one to obtain valuable insights including those similar to statistical models. Thus, Altmann et al. (2010) proposed permutations of covariates to derive p -values assessing the statistical significance of the covariate contribution to the predictive accuracy of a random forest. To visualize the relationships learned by the random forest, partial dependence plots (PDPs) are constructed by varying the values of the selected covariate(s) while keeping other covariates fixed (Chapter 10 in Hastie et al. 2009). This method applies to many other types of black-box models. The resulting PDPs show marginal effects of the variables, and the way they are estimated corresponds to the back-door adjustment of Pearl (1993) used to identify causal effects from observational data (Zhao and Hastie 2021). More specifically, when the predictor is a binary variable such as marijuana decriminalization, the PDP shows the average treatment effect (Zhao and Hastie 2021).



5 Results

5.1 Canada

In this section, the number of earned vehicles per region (Figure 2) was used as the weights.

5.1.1 Average cost per claim

Results for the average cost per claim are as follows. The p -value for the effect of marijuana decriminalization in Canada in the mixed-effects model is above the significance level (Table 2), which implies there is not enough evidence of the decriminalization effect. The bootstrap distribution for this parameter has its centre close to zero (Figure 4) and all versions of the confidence intervals contain zero (Table 3), implying no significant effects.




Figure 4: Density plot of the bootstrap distribution of coefficient b in the mixed-effects model of the effect of marijuana decriminalization on the average cost per claim in Canada. The horizontal lines correspond to 66% and 95% confidence.

Table 2: Parametric estimates from the mixed-effects model of the effect of marijuana decriminalization on the average cost per claim in Canada

	Value	Std.Error	DF	t-value	p-value
XAfter	376.912	380.70	18	0.99005	0.33527
Year	-12.204	144.01	18	-0.08475	0.93340

Table 3: Bootstrap confidence intervals for the coefficient b in the mixed-effects model of the effect of marijuana decriminalization on the average cost per claim in Canada

	Term	Estimate	Lower	Upper	Type	Level
1	XAfter	376.91	-384.09	1,134.1	norm	0.95
2	XAfter	376.91	-338.37	1,185.4	basic	0.95
3	XAfter	376.91	-431.57	1,092.2	perc	0.95

5.1.2 Claim frequency

Results for the claim frequency per 100 earned vehicles lead to the same conclusions as do the results for the average cost per claim presented in the previous section. Specifically, the p -value for the effect of marijuana decriminalization in Canada in the mixed-effects model is above the significance level (Table 4), which implies there is not enough evidence of the decriminalization effect. The bootstrap distribution for this parameter has its centre close to zero (Figure 5) and all versions of the confidence intervals contain zero (Table 5), implying no significant effects.




Figure 5: Density plot of the bootstrap distribution of coefficient b in the mixed-effects model of the effect of marijuana decriminalization on the claim frequency in Canada. The horizontal lines correspond to 66% and 95% confidence.

Table 4: Parametric estimates from the mixed-effects model of the effect of marijuana decriminalization on the claim frequency in Canada

	Value	Std.Error	DF	t-value	p-value
XAfter	-0.19125	0.13038	18	-1.4668	0.159684
Year	0.10657	0.04930	18	2.1616	0.044364

Table 5: Bootstrap confidence intervals for the coefficient b in the mixed-effects model of the effect of marijuana decriminalization on the claim frequency in Canada

	Term	Estimate	Lower	Upper	Type	Level
1	XAfter	-0.19125	-0.40626	0.0160142	norm	0.95
2	XAfter	-0.19125	-0.39958	0.0071157	basic	0.95
3	XAfter	-0.19125	-0.38961	0.0170906	perc	0.95

5.1.3 Québec average cost per claim

Here model (2) is applied to quarterly average costs per claim in Québec. With the observed trend and seasonality, there is not enough evidence of changing average costs after the marijuana decriminalization (Table 6, p -value above 0.05). See the fit of this model in Figure 6; all observed values belong to the 95% prediction interval.

Table 6: Results of estimating the effect of marijuana decriminalization on average cost per claim in Québec within a trend-seasonal model

	Estimate	Std. Error	t-value	p-value
(Intercept)	-331,924.566	79,125.587	-4.19491	0.0018435344
XAfter	12.824	97.316	0.13177	0.8977778234
Year	166.800	39.229	4.25193	0.0016843695
Quarter 2	-450.174	73.391	-6.13388	0.0001106417
Quarter 3	139.939	73.391	1.90675	0.0856636605
Quarter 4	716.961	77.319	9.27280	0.0000031612




Figure 6: Quarterly average cost per claim (in 2019 dollars) in Québec. Black are the observed values, blue are the fitted values and 95% prediction intervals. The dashed line denotes the date of marijuana decriminalization.




Figure 7: Density plot of the bootstrap distribution of coefficient b in the linear model of the effect of marijuana decriminalization on the average cost per claim in Québec.

Figure 7 shows the bootstrap distribution of the parameter with a percentile interval. All versions of the 95% intervals contain zero; hence there is not enough evidence of the decriminalization effect: normal interval (-170.50, 190.57), basic interval (-160.59, 193.20), and percentile interval (-167.55, 186.24).

Additionally, a restricted model (without the indicator variable representing the periods Before/After),

$$y_{t'} = a + ct' + \sum_{i=1}^{s-1} d_i S_{it'} + e_{t'} \tag{6}$$

was estimated on t' before the decriminalization date and similar outputs were obtained from this model (Table 7, Figure 8). Figure 8 shows that before-decriminalization trends extended in the future can accurately predict the average costs per claim, without modelling the effect of decriminalization.

Table 7: Results of estimating the restricted trend-seasonal model for average cost per claim in Québec, using the data before decriminalization

	Estimate	Std. Error	t-value	p-value
(Intercept)	-293,708.61	76,830.157	-3.8228	0.00873184
Year	147.86	38.091	3.8818	0.00815333
Quarter 2	-459.57	79.293	-5.7958	0.00115583
Quarter 3	117.92	79.293	1.4871	0.18754294
Quarter 4	633.56	90.675	6.9871	0.00042774




Figure 8: Quarterly average cost per claim (in 2019 dollars) in Québec. Black are the observed values, blue are the fitted values (restricted trend-seasonal model based on the data before decriminalization) and 95% prediction intervals. The dashed line denotes the date of marijuana decriminalization.

5.1.4 Québec claim frequency

Here model (2) is applied to quarterly claim frequency in Québec. The results mirror those shown in the previous section for the average cost per claim.

With the observed trend and seasonality, there is not enough evidence of changing claim frequencies after the marijuana decriminalization (Table 8, p -value above 0.05). See the fit of this model in Figure 9; all observed values belong to the 95% prediction interval.

Table 8: Results of estimating the effect of marijuana decriminalization on claim frequency in Québec within a trend-seasonal model

	Estimate	Std. Error	t-value	p-value
(Intercept)	-167.016676	206.32901	-0.80947	0.437082930
XAfter	-0.217582	0.25376	-0.85743	0.411287857
Year	0.085714	0.10229	0.83791	0.421654958
Quarter 2	-1.327500	0.19138	-6.93659	0.000040119
Quarter 3	-1.675000	0.19138	-8.75239	0.000005314
Quarter 4	-1.018104	0.20162	-5.04968	0.000499480




Figure 9: Quarterly claim frequency per 100 earned vehicles in Québec. Black are the observed values, blue are the fitted values and 95% prediction intervals. The dashed line denotes the date of marijuana decriminalization.




Figure 10: Density plot of the bootstrap distribution of coefficient b in the linear model of the effect of marijuana decriminalization on the average cost per claim in Québec.

Figure 10 shows the bootstrap distribution of the parameter with a percentile interval. All versions of the 95% intervals contain zero; hence there is not enough evidence of the decriminalization effect: normal interval $(-0.7272, 0.2924)$, basic interval $(-0.7148, 0.2828)$, and percentile interval $(-0.7179, 0.2797)$.

Additionally, similar results were obtained from a restricted model (6) without the indicator variable representing the periods Before/After (Table 9, Figure 11). Figure 11 shows that before-decriminalization trends extended in the future can accurately predict the claims.

Table 9: Results of estimating the restricted trend-seasonal model for claim frequency per 100 earned vehicles in Québec, using the data before decriminalization

	Estimate	Std. Error	t-value	p-value
(Intercept)	-240.96462	175.684328	-1.3716	0.21927370
Year	0.12231	0.087102	1.4042	0.20985112
Quarter 2	-1.17333	0.181317	-6.4712	0.00064639
Quarter 3	-1.49000	0.181317	-8.2177	0.00017525
Quarter 4	-0.74385	0.207344	-3.5875	0.01153847




Figure 11: Quarterly claim frequency per 100 earned vehicles in Québec. Black are the observed values, blue are the fitted values (restricted trend-seasonal model based on the data before decriminalization) and 95% prediction intervals. The dashed line denotes the date of marijuana decriminalization.

5.2 United States

5.2.1 Propensity score matching

The PSM was done for several alternative ratios of the number of selected control states to the number of decriminalized states. Based on the absolute standardized mean differences in Figure 12, the best matching is achieved at the ratio 1:1 (the adjusted sample differences, corresponding to the selected controls and decriminalized states, are the smallest). The effect of matching is yet noticeable at the ratio 3:1, but at the higher ratio of 5:1 the selected control states do not achieve a desirable similarity in the baseline variables (Figure 12). The distribution of propensity scores provides the background information for the observed behaviours. The propensity scores of the treated states (states with decriminalized marijuana) vary from about 0.15 to 0.30, and there are only a few untreated states with propensity scores in this range. With the PSM algorithm tasked to select more matches for each treated state (as the ratio increases), more states with out-of-range propensity scores around 0 are selected as controls. Below are the matches for each decriminalized state:

- for the state of CA:
 - WV
 - FL, MT, WV
 - FL, MT, SD, TN, WV
- for the state of MA:
 - UT
 - KS, UT, VA
 - ID, IN, KS, UT, VA
- for the state of NV:
 - NY
 - AZ, NY, TX
 - AZ, NC, NE, NY, TX

Therefore, the 1:1 matching appears to be optimal in this study because it provides substantially more homogeneous samples (Figure 12A) and happens to follow the approach of stratified matching (Damrongplisit et al. 2010) where the control units were selected from the propensity score range of the treated units.

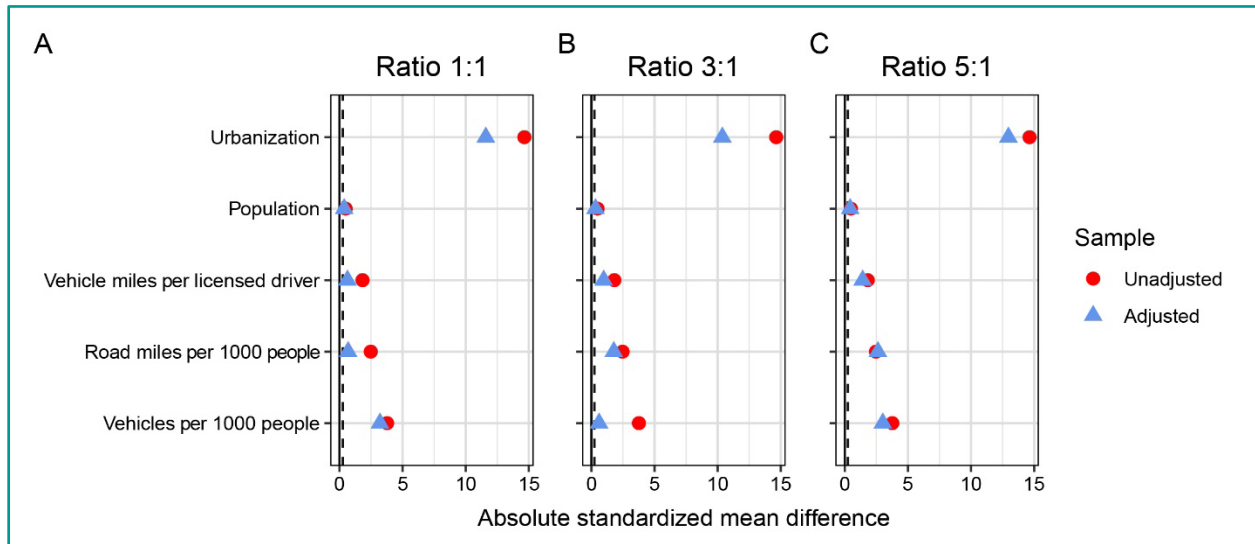


Figure 12: Standardized mean differences for PSM with different ratios of selected control:treated states. The dashed line corresponds to the threshold of 0.25.

Conventional matching




Figure 13: Random forest similarity matrix based on the five state variables.

Similarities between states derived from a random forest (Figure 13) were used to find alternative matches. Each state in the “Legalized” group was matched with the most similar state in the control group if the latter has not been used as a match:

- for the state of CA: NJ
- for the state of MA: FL
- for the state of NV: UT

5.2.2 Estimates of the effects

This section summarizes the results of the analysis of the daily average fatality rate and the number of fatal accidents (from the FARS dataset) and the car accident rate (from the CTAD dataset). All rates are calculated using the annual number of registered vehicles in the studied states. The following predictors were considered in the models: Legal (binary variable for marijuana decriminalization), State (categorical), Month (categorical), Weekday (categorical), Holiday (categorical), Average daily temperature, Average daily temperature range, and Precipitation.

The models were estimated separately for the matched states (mimicking the stratified PSM approach of Damrongplait et al. 2010) and by combining all selected states, then repeated for the manually matched states. Table 10 summarizes the estimated effects of the decriminalization; Figure 14 contains PDPs for the fatality rate (PDPs for the number of fatal accidents are omitted for the sake of space due to the very similar results).

Table 10: Estimates of the decriminalization effects based on regression random forests; Altmann’s p-values are in parentheses. The last column of results excludes UT and uses subsets of ±1 year from the decriminalization date for other states

States	Fatality rate	Rate of fatal accidents	Rate of road accidents
CA, WV	-0.032 (1.00)	-0.015 (0.58)	2.444 (0.01)
MA, UT	-0.026 (0.88)	-0.020 (0.89)	–
NV, NY	-0.015 (0.24)	-0.018 (0.99)	-0.387 (0.02)
CA, WV, MA, UT, NV, NY	-0.011 (0.95)	-0.009 (0.99)	0.642 (0.01)
CA, NJ	-0.014 (0.97)	-0.015 (0.98)	1.196 (0.01)
MA, FL	-0.066 (0.01)	-0.060 (0.01)	-0.943 (0.01)
NV, UT	-0.015 (0.95)	-0.012 (0.81)	–
CA, NJ, MA, FL, NV, UT	-0.012 (0.50)	-0.012 (0.54)	0.321 (0.01)

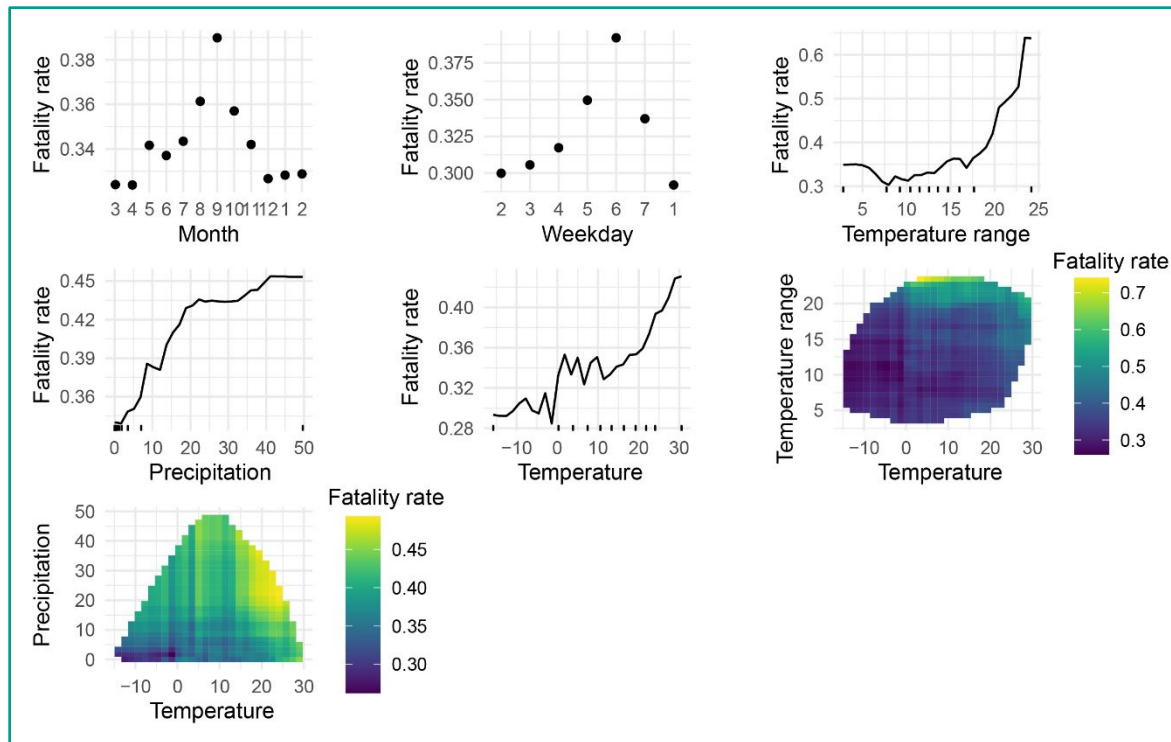


Figure 14: PDPs <https://doi.org/10.15288/jsad.2022.83.494>

6 Bibliography

- Gelmi, T.J., Weinmann, W., Pfäffli, M. (2021) Impact of smoking cannabidiol (CBD)-rich marijuana on driving ability. *Forensic Sciences Research* 1–13. <https://doi.org/10.1080/20961790.2021.1946924>
- Gupta, M., Pawar, N.M., Velaga, N.R. (2021) Impact of lockdown and change in mobility patterns on road fatalities during COVID-19 pandemic. *Transportation Letters* 13:447–460. <https://doi.org/10.1080/19427867.2021.1892937>
- Hall, W., Lynskey, M. (2020) Assessing the public health impacts of legalizing recreational cannabis use: The US experience. *World Psychiatry* 19:179–186. <https://doi.org/10.1002/wps.20735>
- Hartman, R.L., Brown, T.L., Milavetz, G., et al. (2016) Cannabis effects on driving longitudinal control with and without alcohol. *Journal of Applied Toxicology* 36:1418–1429. <https://doi.org/10.1002/jat.3295>
- Hastie, T.J., Tibshirani, R.J., Friedman, J.H. (2009) *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd edn. Springer, New York, NY.
- Lane, T.J., Hall, W. (2019) Traffic fatalities within US states that have legalized recreational cannabis sales and their neighbours. *Addiction* 114:847–856. <https://doi.org/10.1111/add.14536>
- Li, G., Brady, J.E., Chen, Q. (2013) Drug use and fatal motor vehicle crashes: A case-control study. *Accident Analysis & Prevention* 60:205–210. <https://doi.org/10.1016/j.aap.2013.09.001>
- Lin, L., Shi, F., Li, W. (2021) Assessing inequality, irregularity, and severity regarding road traffic safety during COVID-19. *Scientific Reports* 11:1–7. <https://doi.org/10.1038/s41598-021-91392-z>
- Lopez Bernal, J., Cummins S., Gasparrini, A. (2018) The use of controls in interrupted time series studies of public health interventions. *International Journal of Epidemiology* 47:2082–2093. <https://pubmed.ncbi.nlm.nih.gov/29982445/>
- Loy, A., Steele, S., Korobova, J. (2022) Lmeresampler: Bootstrap methods for nested linear mixed-effects models. R package version 0.2.2. <https://github.com/aloy/lmeresampler>
- Moosavi, S., Samavatian, M.H., Parthasarathy, S., et al. (2019) Accident risk prediction based on heterogeneous sparse data: New dataset and insights. In: *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pp. 33–42
- NHTSA DOT (2017) *Marijuana-Impaired Driving*. Report to Congress. National Highway Traffic Safety Administration, United States Department of Transportation, Washington, DC.
- Olmos, A., Govindasamy, P. (2015) Propensity scores: A practical introduction using R. *Journal of MultiDisciplinary Evaluation* 11:68–88.
- Pearl, J. (1993) [Bayesian analysis in expert systems]: Comment – Graphical models, causality and intervention. *Statistical Science* 8:266–269.
- Peters, J., Janzing, D., Schölkopf, B. (2017) *Elements of Causal Inference*. MIT Press, Cambridge, MA.
- Pollini, R.A., Romano, E., Johnson, M.B., Lacey, J.H. (2015) The impact of marijuana decriminalization on California drivers. *Drug and Alcohol Dependence* 150:135–140. <https://doi.org/10.1016/j.drugalcdep.2015.02.024>
- Potter, A., Weinstock, D.M. (2019) *High Time: The Legalization and Regulation of Cannabis in Canada*. McGill-Queen's Press University Press, Montréal.
- Rahman, M., Paul, K.C., Hossain, A., et al. (2021) Machine learning on the COVID-19 pandemic, human mobility and air quality: A review. *IEEE Access* 9:72420–72450. <https://doi.org/10.1109/ACCESS.2021.3079121>



- Romano, E., Torres-Saavedra, P., Voas, R.B., Lacey, J.H. (2014) Drugs and alcohol: Their relative crash risk. *Journal of Studies on Alcohol and Drugs* 75:56–64. <https://doi.org/10.15288/jsad.2014.75.56>
- Romano, E., Torres-Saavedra, P., Voas, R.B., Lacey, J.H. (2017) Marijuana and the risk of fatal car crashes: What can we learn from FARS and NRS data? *Journal of Primary Prevention* 38:315–328. <https://doi.org/10.1007/s10935-017-0478-3>
- Shi, T., Horvath, S. (2006) Unsupervised learning with random forest predictors. *Journal of Computational and Graphical Statistics* 15:118–138. <https://doi.org/10.1198/106186006X94072>
- Smith, E.P., Orvos, D.R., Cairns, J., Jr (1993) Impact assessment using the before-after-control-impact (BACI) model: Concerns and comments. *Canadian Journal of Fisheries and Aquatic Sciences* 50:627–637. <https://doi.org/10.1139/f93-072>
- Stuart, E.A., Huskamp, H.A., Duckworth, K., et al. (2014) Using propensity scores in difference-in-differences models to estimate the effects of a policy change. *Health Services and Outcomes Research Methodology* 14:166–182. <https://doi.org/10.1007/s10742-014-0123-z>
- Windle, S.B., Eisenberg, M.J., Reynier, P., et al. (2021a) Association between legalization of recreational cannabis and fatal motor vehicle collisions in the United States: An ecologic study. *Canadian Medical Association Open Access Journal* 9:E233–E241. <https://doi.org/10.9778/cmajo.20200155>
- Windle, S.B., Sequeira, C., Filion, K.B., et al. (2021b) Impaired driving and legalization of recreational cannabis. *CMAJ* 193:E481–E485. <https://doi.org/10.1503/cmaj.191032>
- Wooldridge, J.M. (2013) *Introductory Econometrics: A Modern Approach*, 5th edn. Cengage Learning, Mason, OH.
- Wuertz, D., Setz, T., Chalabi, Y. (2022) timeDate: Rmetrics – chronological and calendar objects. R package version 4021.104. <https://CRAN.R-project.org/package=timeDate>
- Yasin, Y.J., Grivna, M., Abu-Zidan, F.M. (2021) Global impact of COVID-19 pandemic on road traffic collisions. *World Journal of Emergency Surgery* 16:1–14. <https://doi.org/10.1186/s13017-021-00395-8>
- Zhao, Q., Hastie, T. (2021) Causal interpretations of black-box models. *Journal of Business & Economic Statistics* 39:272–281. <https://doi.org/10.1080/07350015.2019.1624293>
- Zhao, Q.-Y., Luo, J.-C., Su, Y., et al. (2021) Propensity score matching with R: Conventional methods and new features. *Annals of Translational Medicine* 9(9). <https://doi.org/10.21037/atm-20-3998>
- Zuur, A., Ieno, E.N., Walker, N.J., et al. (2009) *Mixed Effects Models and Extensions in Ecology with R*. Springer, New York, NY.



7 About the author

Dr. Vyacheslav Lyubchich is an Associate Research Professor at the University of Maryland Center for Environmental Science (UMCES). He received his Ph.D. in Statistics from the Orenburg State University, Russia in 2011. In the same year, he was awarded the Government of Canada postdoctoral fellowship to continue his research in time series methodology at the Department of Statistics and Actuarial Science of the University of Waterloo, Canada. Since 2015, V. Lyubchich is a research faculty and a founding member of the Environmental Statistical Collaborative at the UMCES.



© 2022 Canadian Institute of Actuaries and Casualty Actuarial Society

Canadian Institute of Actuaries

360 Albert Street, Suite 1740

Ottawa, ON K1R 7X7

613-236-8196

head.office@cia-ica.ca

cia-ica.ca

seeingbeyondrisk.ca



The Canadian Institute of Actuaries (CIA) is the qualifying and governing body of the actuarial profession in Canada. We develop and uphold rigorous standards, share our risk management expertise, and advance actuarial science to improve lives in Canada and around the world. Our more than 6,000 members apply their knowledge of math, statistics, data analytics, and business in providing services and advice of the highest quality to help Canadian people and organizations face the future with confidence.

The Casualty Actuarial Society

4350 N. Fairfax Drive, Suite 250

Arlington, VA 22203

<https://www.casact.org>

The Casualty Actuarial Society (CAS) is a leading international organization for credentialing and professional education. Founded in 1914, the CAS is the world's only actuarial organization focused exclusively on property and casualty risks and serves over 9,500 members worldwide. CAS members are experts in property and casualty insurance, reinsurance, finance, risk management and enterprise risk management. Professionals educated by the CAS empower business and government to make well-informed strategic, financial and operational decisions.