Estimating online ad effectiveness

Kay H. Brodersen
Hal R. Varian
Google,Inc.

January 7, 2015

A note about the paper

- Created using `knitr`
- Allows you to merge LaTeXand **R** commands
- Runs **R** each time you change paper and caches results
- Inserts **R** code and output into paper
- Voila: reproducible research

Estimating online ad effectiveness

1. Apply treatment: change ad spend, bid, budget, etc.
    1.1 Treat in some places but not others
    1.2 Treat in some times but not others
2. Compare to counterfactual: what would have happened without experiment?
3. Of course, counterfactual can not be observed, so it must be estimated

Procedure

Train a model to predict metric of interest such as website visits or revenue.

Test the model on out-of-sample data to evaluate performance.

Treat Apply model to periods or places where treatment was applied to estimate counterfactual

Compare actual and (estimated) counterfactual outcomes.

Hypothetical example of train-test-treat

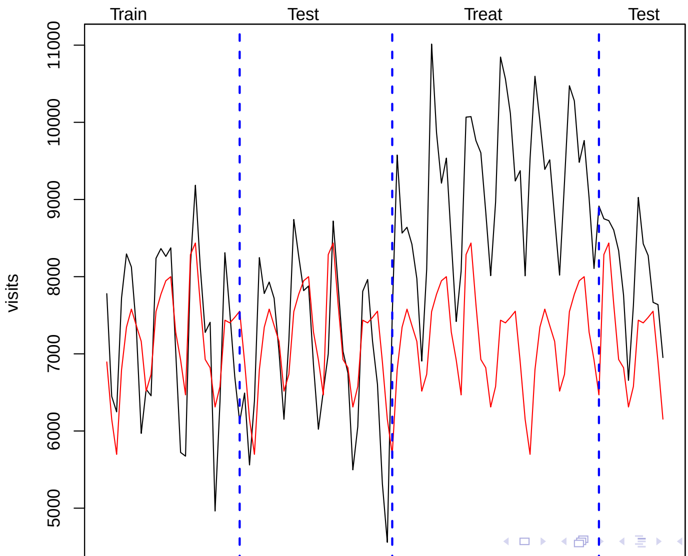

Bayesian Structural Time Series

We will do this in a time series context using BSTS (available from CRAN.) BSTS combines:

Kalman filter. Accounts for seasonality and trend

Spike-and-slab regression. Automated selection of predictors

Bayesian model averaging. Avoids overfitting.

Described in Scott-Varian [2013,2014], Brodersen et. al. [2013]. Related to “interrupted regression”, “synthetic controls”.

Predictors selected by BSTS

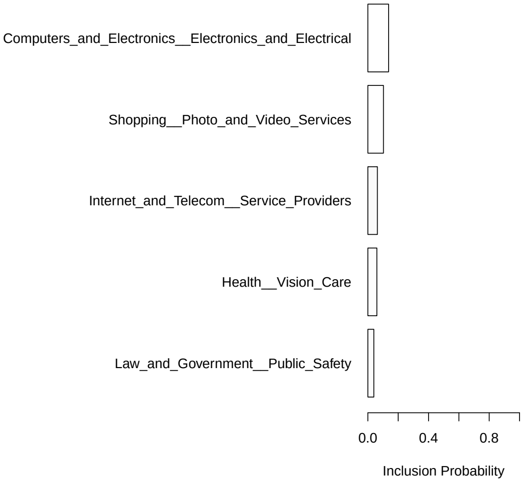

Prediction from BSTS

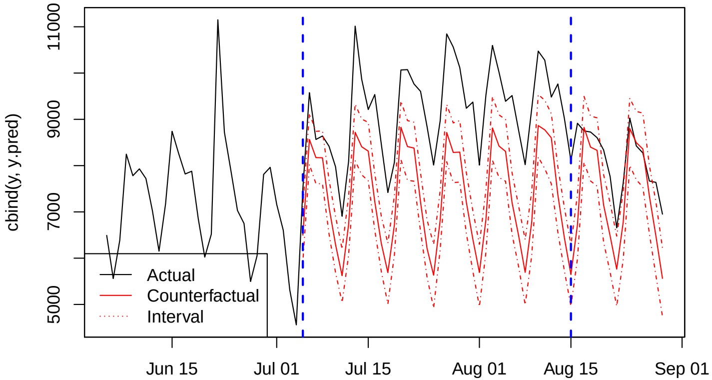

Research strategy?

1. Use a parametric model for impact of ad campaign?
   - Benefit: Can use all the data to estimate
   - Cost: Restrictive functional form: a parallel shift
2. Use alternative estimation technique?
3. Use different models for seasonality and trend

1. Use parallel shift for ad impact

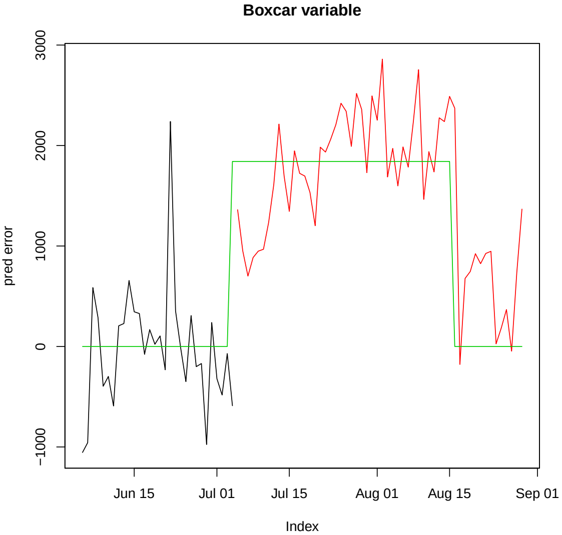

Boxcar indicator variable for campaign

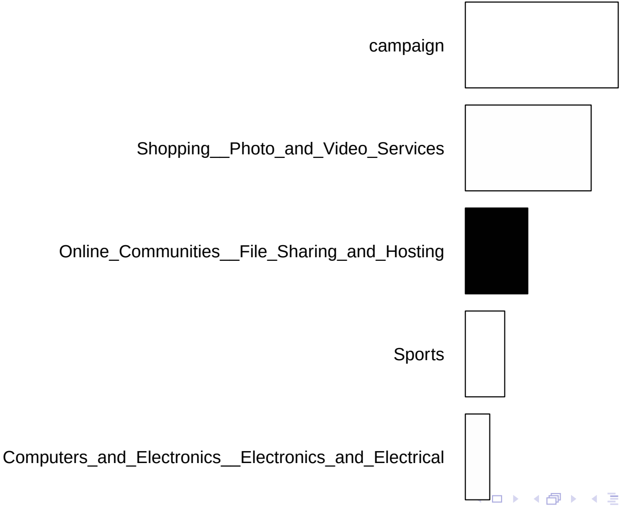

2. Alternative estimation: linear model

Drop Kalman filter, just use simple linear model

- July 4 holiday dummy
- Top two categories from Google Trends as regressors

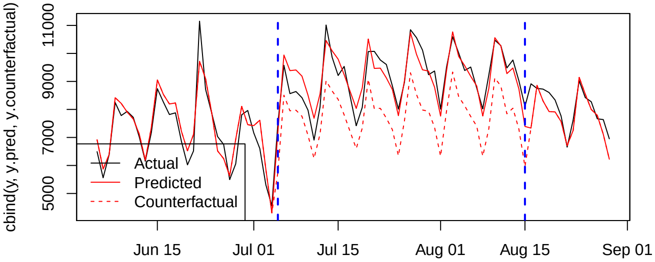

Deseasonalized the data first

Deseasonalize by fitting model with holiday regressor + day-of-week dummies. (Explain spike.)

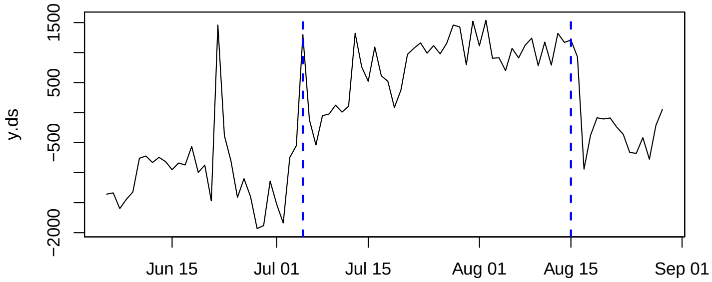

Alternative approaches

1. Make no adjustment for seasonality (since predictor already has appropriate seasonality)
2. Deseasonalize both predictor and outcome
   - Use boxcar regressor
   - Use extrapolation

3. Alternative seasonality: detrend first

Deseasonalize by fitting model with holiday regressor + day of week dummies. (Explain spike.)

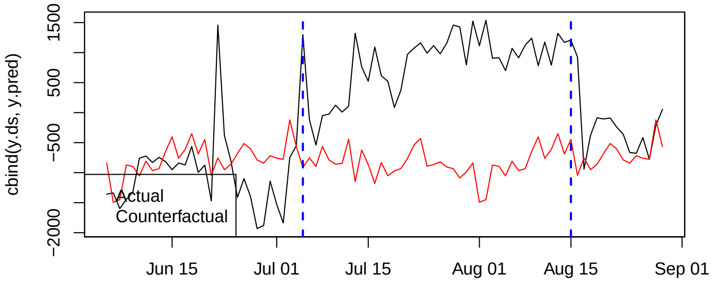

Use weekly data

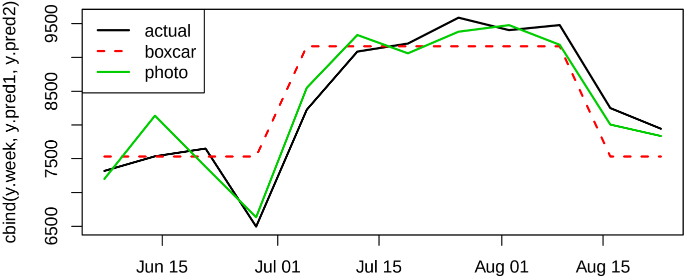

Summary

	method	estimate
1	bsts-extrap	1830.43
2	bsts-boxcar	1362.88
3	bsts-boxcar-all-predictors	1279.05
4	bsts-boxcar-top-predictors	1327.06
5	lm-boxcar	1434.57
6	lm-extrap	1289.19
7	not deseasonalized	1393.41
8	deseasonalized-boxcar	1300.67
9	deseasonalized-extrap	1298.37
10	week-boxcar	1248.61

What about revenue?

- Ad clicks may cannibalize search clicks
- May want to look at total number of clicks (i.e., visitors)
- But ad clicks may be worth more or less than search clicks, so really want revenue (or profit)
- Can model ad revenue, search revenue separately or together

Examine a different advertiser using a different set of possible predictors

BSTS: predictor selection for visits

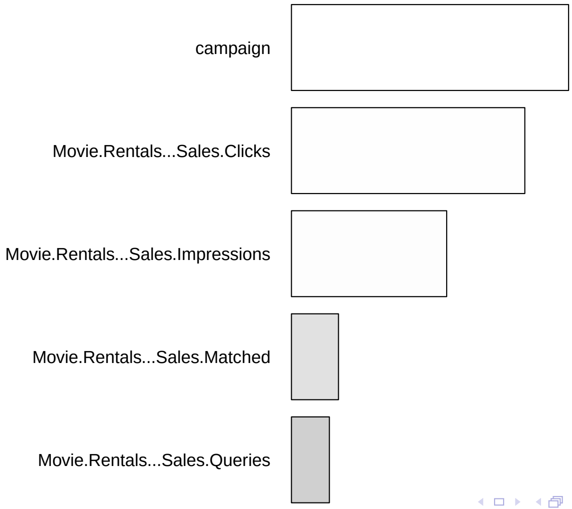

BSTS: Visits actual and counterfactual

Uses the BSTS extrapolation model and a linear regression

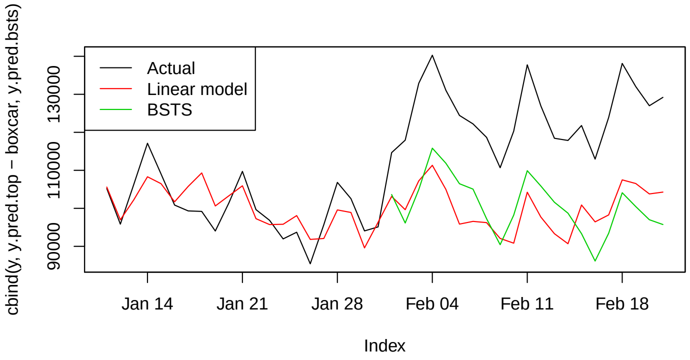

BSTS: predictor selection for revenue

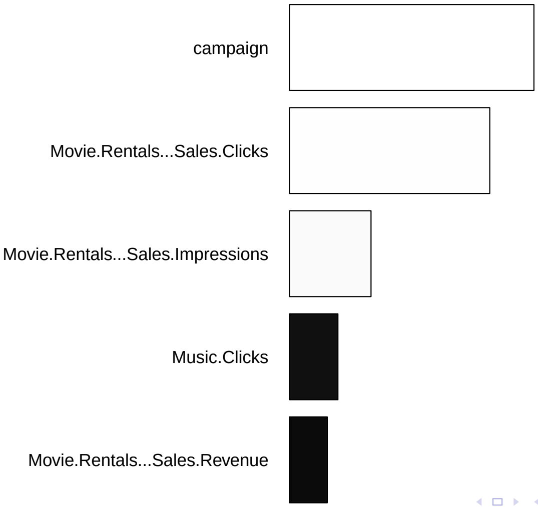

BSTS: Revenue actual and counterfactual

Uses the BSTS extrapolation model and a linear regression

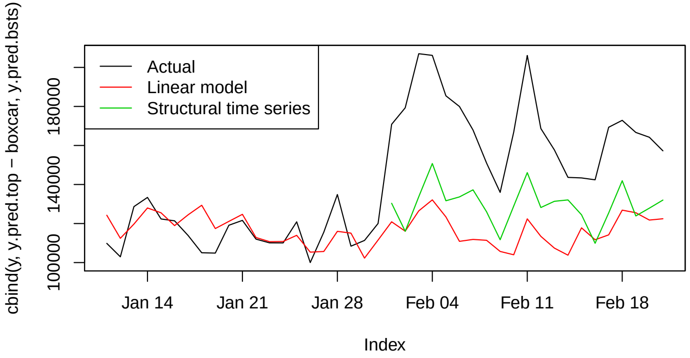

Revenue cannibalization

When campaign begins ad revenue increases significantly, organic revenue falls somewhat.