Estimating the causal effect of an intervention in a time series setting: the C-ARIMA approach

Fiammetta Menchetti[†], Fabrizio Cipollini, Fabrizia Mealli

July 2, 2021

† DiSIA, fiammetta.menchetti@unifi.it

Outline

1. Motivation & research goals
2. Causal framework
3. C-ARIMA
4. Empirical application
5. Results

Motivation

- The **Rubin Causal Model** (RCM) allows to define the causal effect of a treatment as a contrast of potential outcomes and develop methods for its estimation under a previously discussed set of assumptions

Motivation

- The **Rubin Causal Model** (RCM) allows to define the causal effect of a treatment as a contrast of potential outcomes and develop methods for its estimation under a previously discussed set of assumptions
- In time series settings, identification and estimation of causal effects using potential outcomes have been formalized in the context of **randomized experiments** (Bojinov and Shephard; 2019; Rambachan and Shephard; 2019; Bojinov et al.; 2020)

Motivation

- The **Rubin Causal Model** (RCM) allows to define the causal effect of a treatment as a contrast of potential outcomes and develop methods for its estimation under a previously discussed set of assumptions
- In time series settings, identification and estimation of causal effects using potential outcomes have been formalized in the context of **randomized experiments** (Bojinov and Shephard; 2019; Rambachan and Shephard; 2019; Bojinov et al.; 2020)
- Methods exploiting logitudinal information abound in the literature, the most popular being **DiD** (e.g., Card and Krueger (1993); Meyer et al. (1995)) and **synthetic controls** (e.g,Abadie et al. (2010, 2015))

Motivation

- The **Rubin Causal Model** (RCM) allows to define the causal effect of a treatment as a contrast of potential outcomes and develop methods for its estimation under a previously discussed set of assumptions
- In time series settings, identification and estimation of causal effects using potential outcomes have been formalized in the context of **randomized experiments** (Bojinov and Shephard; 2019; Rambachan and Shephard; 2019; Bojinov et al.; 2020)
- Methods exploiting logitudinal information abound in the literature, the most popular being **DiD** (e.g., Card and Krueger (1993); Meyer et al. (1995)) and **synthetic controls** (e.g,Abadie et al. (2010, 2015))
- They usually require some control observations and often rely on functional form assumptions for the trend (such as the parellel trend assumption for DiD)

Motivation

- The **Rubin Causal Model** (RCM) allows to define the causal effect of a treatment as a contrast of potential outcomes and develop methods for its estimation under a previously discussed set of assumptions
- In time series settings, identification and estimation of causal effects using potential outcomes have been formalized in the context of **randomized experiments** (Bojinov and Shephard; 2019; Rambachan and Shephard; 2019; Bojinov et al.; 2020)
- Methods exploiting logitudinal information abound in the literature, the most popular being **DiD** (e.g., Card and Krueger (1993); Meyer et al. (1995)) and **synthetic controls** (e.g,Abadie et al. (2010, 2015))
- They usually require some control observations and often rely on functional form assumptions for the trend (such as the parellel trend assumption for DiD)
- In the econometrics literature, **intervention analysis** (Box and Tiao; 1975, 1976) employs ARIMA models to assess the impact of shocks on time series, but fails to define the causal estimands and to discuss the assumptions enabling the attribution of the effect to the intervention

Research goals

Closing the gap between causal inference under the RCM and intervention analysis, in this paper we propose a novel approach, C-ARIMA, to **estimate the causal effect of an intervention in observational time series settings under the RCM.**

Research goals

Closing the gap between causal inference under the RCM and intervention analysis, in this paper we propose a novel approach, C-ARIMA, to **estimate the causal effect of an intervention in observational time series settings under the RCM.**

The C-ARIMA approach has two pillars of equal importance:

Research goals

Closing the gap between causal inference under the RCM and intervention analysis, in this paper we propose a novel approach, C-ARIMA, to **estimate the causal effect of an intervention in observational time series settings under the RCM.**

The C-ARIMA approach has two pillars of equal importance:

1. A causal framework under the RCM, comprising: the **assumptions** allowing the postulation of potential outcomes and related causal contrasts (causal estimands); the assumptions for the identification and estimation of these causal estimands

Research goals

Closing the gap between causal inference under the RCM and intervention analysis, in this paper we propose a novel approach, C-ARIMA, to **estimate the causal effect of an intervention in observational time series settings under the RCM.**

The C-ARIMA approach has two pillars of equal importance:

1. A causal framework under the RCM, comprising: the **assumptions** allowing the postulation of potential outcomes and related causal contrasts (causal estimands); the assumptions for the identification and estimation of these causal estimands
2. An inferential methodology based on **ARIMA models** that involves predicting the outcome in the absence of intervention and contrasting it with the observed outcome

Research goals

Closing the gap between causal inference under the RCM and intervention analysis, in this paper we propose a novel approach, C-ARIMA, to **estimate the causal effect of an intervention in observational time series settings under the RCM.**

The C-ARIMA approach has two pillars of equal importance:

1. A causal framework under the RCM, comprising: the **assumptions** allowing the postulation of potential outcomes and related causal contrasts (causal estimands); the assumptions for the identification and estimation of these causal estimands
2. An inferential methodology based on **ARIMA models** that involves predicting the outcome in the absence of intervention and contrasting it with the observed outcome

C-ARIMA shares many features with "CausalImpact" (Brodersen et al.; 2015), but it is based on ARIMA models and thus it can be used as an alternative by those that are not familiar to (or are not willing to adopt) Bayesian inference.

Assumptions

Let $W_{i,t} \in \{0,1\}$ be a random variable describing the treatment assignment of unit $i \in \{1,\ldots,N\}$ at time $t \in \{1,\ldots,T\}$, where 1 denotes that a "treatment" (or "intervention") has taken place and 0 denotes control. We maintain the following assumptions:

Assumptions

Let $\mathrm{W}_{i,t} \in \{0,1\}$ be a random variable describing the treatment assignment of unit $i \in \{1,\ldots,N\}$ at time $t \in \{1,\ldots,T\}$, where 1 denotes that a "treatment" (or "intervention") has taken place and 0 denotes control. We maintain the following assumptions:

(A1: Single persistent intervention)

$$\exists t^* \in \{1,\ldots,T\} \; s.t \; \mathrm{w}_{i,t} = 0 \;\; \forall t \leq t^* \; and \; \forall t > t^*, \mathrm{w}_{i,t} \in \{(1,\ldots,1),(0,\ldots,0)\}$$

Assumptions

Let $W_{i,t} \in \{0,1\}$ be a random variable describing the treatment assignment of unit $i \in \{1,\dots,N\}$ at time $t \in \{1,\dots,T\}$, where 1 denotes that a "treatment" (or "intervention") has taken place and 0 denotes control. We maintain the following assumptions:

(A1: Single persistent intervention)

$$\exists t^* \in \{1,\dots,T\} \; s.t \;\; w_{i,t} = 0 \;\; \forall t \leq t^* \; and \; \forall t > t^*, w_{i,t} \in \{(1,\dots,1),(0,\dots,0)\}$$

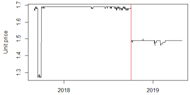

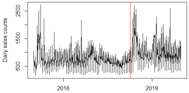

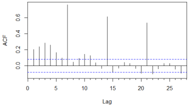

Assumptions

Let $W_{i,t} \in \{0,1\}$ be a random variable describing the treatment assignment of unit $i \in \{1,\ldots,N\}$ at time $t \in \{1,\ldots,T\}$, where 1 denotes that a "treatment" (or "intervention") has taken place and 0 denotes control. We maintain the following assumptions:

(A1: Single persistent intervention)

$$\exists t^* \in \{1,\ldots,T\} \; s.t \; w_{i,t} = 0 \;\; \forall t \leq t^* \; and \; \forall t > t^*, w_{i,t} \in \{(1,\ldots,1),(0,\ldots,0)\}$$

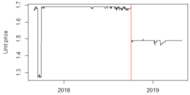

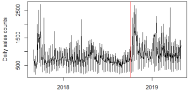

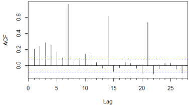

(A2: Temporal no-interference)

$$For \; all \; i \in \{1,\ldots,N\}, \; Y_{i,t}(w_{1:N,t^*+1:T}) = Y_{i,t}(w_{i,t^*+1:T})$$

Assumptions

(A3: Covariates-treatment independence)

$$X_{i,t}(w_{i,t^*+1:T}) = X_{i,t}(w'_{i,t^*+1:T}).$$

(A4: Non-anticipating individualistic treatment)

$$\Pr(W_{1:N,t^*+1} = w_{1:N,t^*+1} \,|\, W_{1:N,1:t^*}, Y_{1:N,1:T}(w_{1:N,1:T}), X_{1:N,1:T}) = \\ = \prod_{i=1}^{N} \Pr(W_{i,t^*+1} = w_{i,t^*+1} \,|\, Y_{i,1:t^*}(w_{i,1:t^*}), X_{i,1:t^*}).$$

Above assumptions are essential to define, estimate and attribute the causal effect to the intervention. Moreover, they allow us to ease notation: under Assumption 1, for all $t > t^*$ we can write $w_{i,t} = w_i$ and if Assumption 2 holds we can also drop the i subscript. From now on, we can use $Y_t(w)$ to denote the potential outcome of a generic unit at time $t > t^*$.

Causal estimands

Definition

For the two treatment paths w, w', the point causal effect at time $t > t^*$ is,

$$\tau_t(\mathrm{w}; \mathrm{w}') = \mathrm{Y}_t(\mathrm{w}) - \mathrm{Y}_t(\mathrm{w}'). \tag{1}$$

Causal estimands

Definition

For the two treatment paths w, w', the point causal effect at time $t > t^*$ is,

$$\tau_t(\mathrm{w}; \mathrm{w}') = \mathrm{Y}_t(\mathrm{w}) - \mathrm{Y}_t(\mathrm{w}'). \tag{1}$$

The cumulative causal effect at time $t > t^*$ is

$$\Delta_t(\mathrm{w}; \mathrm{w}') = \sum_{s=t^*+1}^{t} \tau_s(\mathrm{w}; \mathrm{w}'). \tag{2}$$

Causal estimands

Definition

For the two treatment paths w, w', the point causal effect at time $t > t^*$ is,

$$\tau_t(\mathrm{w}; \mathrm{w}') = \mathrm{Y}_t(\mathrm{w}) - \mathrm{Y}_t(\mathrm{w}'). \tag{1}$$

The cumulative causal effect at time $t > t^*$ is

$$\Delta_t(\mathrm{w}; \mathrm{w}') = \sum_{s=t^*+1}^{t} \tau_s(\mathrm{w}; \mathrm{w}'). \tag{2}$$

The temporal average causal effect at time $t > t^*$

$$\bar{\tau}_t(\mathrm{w}; \mathrm{w}') = \frac{1}{t - t^*} \sum_{s=t^*+1}^{t} \tau_s(\mathrm{w}; \mathrm{w}') = \frac{\Delta_t(\mathrm{w}; \mathrm{w}')}{t - t^*}. \tag{3}$$

Causal estimands

Definition

For the two treatment paths w, w', the point causal effect at time $t > t^*$ is,

$$\tau_t(\mathrm{w}; \mathrm{w}') = \mathrm{Y}_t(\mathrm{w}) - \mathrm{Y}_t(\mathrm{w}'). \tag{1}$$

The cumulative causal effect at time $t > t^*$ is

$$\Delta_t(\mathrm{w}; \mathrm{w}') = \sum_{s=t^*+1}^{t} \tau_s(\mathrm{w}; \mathrm{w}'). \tag{2}$$

The **temporal average causal effect** at time $t > t^*$

$$\bar{\tau}_t(\mathrm{w}; \mathrm{w}') = \frac{1}{t - t^*} \sum_{s=t^*+1}^{t} \tau_s(\mathrm{w}; \mathrm{w}') = \frac{\Delta_t(\mathrm{w}; \mathrm{w}')}{t - t^*}.$$

C-ARIMA

let us assume $\{\mathrm{Y}_t(\mathrm{w})\}$ evolving as

$$Y_t(\mathrm{w}) = \frac{\theta_q(L)}{\phi_p(L)}\varepsilon_t + \tau_t \mathbb{1}_{\{\mathrm{w}=1\}} \tag{3}$$

where,

- $\phi_p(L)$ and $\theta_q(L)$ are lag polynomials having roots all outside the unit circle
- given this representation, the point causal effect at time $t > t^*$ is $\tau_t \equiv \mathrm{Y}_t(\mathrm{w}=1) - \mathrm{Y}_t(\mathrm{w}=0)$
- $\tau_t = 0\ \forall t \leq t^*$ and $\mathbb{1}_{\{\mathrm{w}=1\}}$ is an indicator function which is one if $\mathrm{w}=1$
- ε_t is white noise with mean 0 and variance σ^2_ε

C-ARIMA

let us assume $\{\mathrm{Y}_t(\mathrm{w})\}$ evolving as

$$Y_t(\mathrm{w}) = \frac{\theta_q(L)}{\phi_p(L)}\varepsilon_t + \tau_t \mathbb{1}_{\{\mathrm{w}=1\}} \tag{3}$$

where,

- $\phi_p(L)$ and $\theta_q(L)$ are lag polynomials having roots all outside the unit circle
- given this representation, the point causal effect at time $t > t^*$ is $\tau_t \equiv \mathrm{Y}_t(\mathrm{w}=1) - \mathrm{Y}_t(\mathrm{w}=0)$
- $\tau_t = 0 \; \forall t \leq t^*$ and $\mathbb{1}_{\{\mathrm{w}=1\}}$ is an indicator function which is one if $\mathrm{w}=1$
- ε_t is white noise with mean 0 and variance σ^2_ε

C-ARIMA

let us assume $\{Y_t(w)\}$ evolving as

$$(1 - L^s)^D(1 - L)^d Y_t(w) = \frac{\Theta_Q(L^s)\theta_q(L)}{\Phi_P(L^s)\phi_p(L)}\varepsilon_t + (1 - L^s)^D(1 - L)^d X_t'\beta + \tau_t \mathbb{1}_{\{w=1\}} \tag{3}$$

where,

- $\phi_p(L)$ and $\theta_q(L)$ are lag polynomials having roots all outside the unit circle
- given this representation, the point causal effect at time $t > t^*$ is $\tau_t \equiv Y_t(w = 1) - Y_t(w = 0)$
- $\tau_t = 0 \ \forall t \leq t^*$ and $\mathbb{1}_{\{w=1\}}$ is an indicator function which is one if $w = 1$
- ε_t is white noise with mean 0 and variance σ_ε^2
- $\Theta_Q(L^s)$, $\Phi_P(L^s)$ are the lag polynomials of the seasonal part of the model having roots all outside the unit circle
- $(1 - L^s)^D$ and $(1 - L)^d$ are the differencing operators to ensure stationarity

After some manipulation, Equation (3) becomes

$$S_t = z_t + \tau_t$$

where, $S_t = T(Y_t) - T(X_t)'\beta$ and $T(\cdot)$ is the transformation of Y_t needed to achieve stationarity, i.e. $T(Y_t) = (1 - L^s)^D(1 - L)^d Y_t$; z_t includes the stationary part of the model, namely,

$$z_t = \frac{\Theta_Q(L^s)\theta_q(L)}{\Phi_P(L^s)\phi_p(L)}\varepsilon_t$$

Denoting with H_0 the situation where the intervention has no effect, namely, $\tau_t = 0$ for all $t > t^*$, the k-step ahead forecast of S_t under H_0, conditionally on the information up to time t^* is

$$\hat{S}_{t^*+k} = E[S_{t^*+k}|\mathcal{I}_{t^*}, H_0] = \hat{z}_{t^*+k|t^*}$$

After some manipulation, Equation (3) becomes

$$S_t = z_t + \tau_t$$

where, $S_t = T(Y_t) - T(X_t)'\beta$ and $T(\cdot)$ is the transformation of Y_t needed to achieve stationarity, i.e. $T(Y_t) = (1 - L^s)^D(1 - L)^d Y_t$; z_t includes the stationary part of the model, namely,

$$z_t = \frac{\Theta_Q(L^s)\theta_q(L)}{\Phi_P(L^s)\phi_p(L)}\varepsilon_t$$

Denoting with H_0 the situation where the intervention has no effect, namely, $\tau_t = 0$ for all $t > t^*$, the k-step ahead forecast of S_t under H_0, conditionally on the information up to time t^* is

$$\hat{S}_{t^*+k} = E[S_{t^*+k} | \mathcal{I}_{t^*}, H_0] = \hat{z}_{t^*+k|t^*}$$

Inference

(Causal effect estimators)

For any integer k, let $S_{t^*+k}(\mathrm{w})$ be the observed potential outcome time series and let $\hat{S}_{t^*+k}(\mathrm{w}')$ be the corresponding estimate of the missing potential outcomes under model (3). Then, estimators of the point, cumulative and temporal average effects are, respectively,

$$\hat{\tau}_{t^*+k}(\mathrm{w};\mathrm{w}') = S_{t^*+k}(\mathrm{w}) - \hat{S}_{t^*+k}(\mathrm{w}')$$

$$\hat{\Delta}_{t^*+k}(\mathrm{w};\mathrm{w}') = \sum_{h=1}^{k} \hat{\tau}_{t^*+h}(\mathrm{w};\mathrm{w}')$$

$$\hat{\bar{\tau}}_{t^*+k}(\mathrm{w};\mathrm{w}') = \frac{1}{k}\sum_{h=1}^{k} \hat{\tau}_{t^*+h}(\mathrm{w};\mathrm{w}') = \frac{\hat{\Delta}_{t^*+k}(\mathrm{w};\mathrm{w}')}{k}.$$

Inference

Theorem

Let $\{Y_t\}$ follow the regression model with ARIMA errors defined in Equation (3). Under the null hypothesis that the intervention has no effect, $H_0 : \tau_t(\mathrm{w};\mathrm{w}') = 0$ for all $t > t^*$, the estimators of the point, cumulative and temporal average effects are distributed as follows,

$$\hat{\tau}_{t^*+k}(\mathrm{w};\mathrm{w}')|H_0 \sim N\left[0, \sigma_\varepsilon^2 \sum_{i=0}^{k-1} \psi_i^2\right] \tag{4}$$

$$\widehat{\Delta}_{t^*+k}(\mathrm{w};\mathrm{w}')|H_0 \sim N\left[0, \sigma_\varepsilon^2 \sum_{h=1}^{k} \left(\sum_{i=0}^{k-h} \psi_i\right)^2\right] \tag{5}$$

$$\hat{\bar{\tau}}_{t^*+k}(\mathrm{w};\mathrm{w}')|H_0 \sim N\left[0, \frac{1}{k^2}\sigma_\varepsilon^2 \sum_{h=1}^{k} \left(\sum_{i=0}^{k-h} \psi_i\right)^2\right] \tag{6}$$

where, the ψ_i's are the coefficients of a moving average of order $k-1$ whose values are functions of the ARMA parameters in Equation (3).

Summarizing, to estimate the effect of an intervention with C-ARIMA we need to follow these steps:

1. estimate the ARIMA model only in the pre-intervention period, so as to learn the dynamics of the dependent variable and the links with the covariates without being influenced by the treatment
2. based on the process learned in the pre-intervention period, perform a prediction step and obtain an estimate of the counterfactual outcome during the post-intervention period
3. by comparing the observations with the corresponding forecasts at any time point after the intervention, evaluate the resulting differences, which represent the estimated point causal effects

Comparison with REG-ARIMA

Fitting a linear regression with ARIMA errors (REG-ARIMA) is another widely used approach to estimate the effect of an interventions on time series. In its simplest formulation, such a model can be written as,

$$\mathsf{Y}_t = c + D_t\beta_0 + z_t$$
$$z_t = \frac{\theta_q(L)}{\phi_p(L)}\varepsilon_t$$

where,

- z_t is a stationary ARMA(p, q)
- D_t is a dummy variable taking value 1 after the intervention and 0 otherwise and β_0 is its regression coefficient, which gives the size of the "effect"

Comparison with REG-ARIMA

REG-ARIMA is a standard intervention analysis approach that is used when the intervention is supposed to have produced a level shift on the outcome.

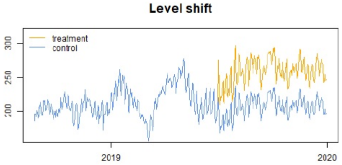

Two main differences between C-ARIMA and REG-ARIMA:

- without a critical discussion of the **assumptions**, the effect grasped by β_0 can not be attributed to the intervention
- REG-ARIMA is fitted on the entire time series and thus it may require the estimation of many models to learn the structure of the effect; C-ARIMA can estimate **any form of effect** (level shift, slope change and irregular time-varying effects) in only **one step**, since it assumes no structure on τ_t.

Empirical application

- On October 4th, 2018, the Florence branch of a supermarket chain store permanently lowered the price of 707 store brands in several product categories

Empirical application

- On October 4th, 2018, the Florence branch of a supermarket chain store permanently lowered the price of 707 store brands in several product categories
- Among the 284 items in the "cookies" category, 11 store brands were selected for a permanent price reduction ranging from -5.7% to -23.2%; for each of them, the supermarket chain identified 11 direct competitors.

Empirical application

- On October 4th, 2018, the Florence branch of a supermarket chain store permanently lowered the price of 707 store brands in several product categories
- Among the 284 items in the "cookies" category, 11 store brands were selected for a permanent price reduction ranging from −5.7% to −23.2%; for each of them, the supermarket chain identified 11 direct competitors.

Figure: Store brands (first row) and direct competitor brands (second row).

Empirical application

- We performed separate analyses on the two subgroups of store brands and competitor brands under two different definitions of intervention: permanent price discount (store brands) and relative price increase (competitors)
- We considered each product to be treated and under Assumption 2 we estimated 11 models for the store brands and 10 models for the competitor brands
- Analysis period: September 1, 2017, until April 30, 2019
- Dependent variable: Daily sales counts (log scale)
- Covariates: i) Day-of-the-week dummies; ii) holiday dummies; iii) unit price (actual price for the competitor brands and, to respect Assumption 3, the price prior to the intervention for the store brands)

Empirical application results

Figure: Causal effect estimates of the permanent price reduction on sales of **store-brand cookies** at three different time horizons. In this table, $\hat{\bar{\tau}}_t$ is the estimated temporal average effect, while $\hat{\beta}_0$ is the coefficient estimate of the intervention dummy of REG-ARIMA.

	Time horizon:					
	1 month		3 months		6 months	
Item	$\hat{\bar{\tau}}_t$	$\hat{\beta}_0$	$\hat{\bar{\tau}}_t$	$\hat{\beta}_0$	$\hat{\bar{\tau}}_t$	$\hat{\beta}_0$
1	0.14	0.14	$0.15^{\cdot}$	$0.12^{\cdot}$	0.18^{***}	0.16^{**}
	(0.12)	(0.09)	(0.08)	(0.07)	(0.06)	(0.06)
2	0.14	0.10	$0.13^{\cdot}$	$0.12^{\cdot}$	0.14^{**}	0.13^{**}
	(0.12)	(0.14)	(0.08)	(0.07)	(0.05)	(0.05)
3	$0.19^{\cdot}$	$0.15^{\cdot}$	0.21^{**}	0.15^{*}	0.25^{***}	0.24^{***}
	(0.11)	(0.08)	(0.07)	(0.07)	(0.05)	(0.04)
4	0.49^{***}	0.00	0.30^{***}	0.19^{*}	0.32^{***}	0.28^{***}
	(0.09)	(0.13)	(0.06)	(0.08)	(0.04)	(0.05)
5	−0.02	−0.06	0.07	−0.07	$0.11^{\cdot}$	−0.06
	(0.12)	(0.12)	(0.08)	(0.11)	(0.06)	(0.10)
6	0.24^{*}	$0.26^{\cdot}$	0.34^{***}	0.24^{*}	0.37^{***}	0.23^{*}
	(0.12)	(0.14)	(0.08)	(0.12)	(0.06)	(0.11)
7	0.55^{***}	0.75^{***}	0.34^{***}	0.70^{***}	0.30^{***}	0.77^{***}
	(0.10)	(0.11)	(0.07)	(0.10)	(0.05)	(0.11)
8	0.26^{***}	0.29^{**}	0.25^{***}	0.29^{**}	0.14^{**}	0.28^{**}
	(0.08)	(0.09)	(0.07)	(0.10)	(0.05)	(0.09)
9	0.47^{***}	0.70^{***}	0.20^{***}	0.29^{***}	0.21^{***}	0.26^{***}
	(0.06)	(0.10)	(0.04)	(0.09)	(0.03)	(0.06)
10	0.66^{***}	0.85^{***}	0.57^{***}	0.82^{***}	0.33^{***}	0.85^{***}
	(0.11)	(0.14)	(0.08)	(0.15)	(0.06)	(0.15)
11	0.12^{*}	0.02	0.16^{**}	0.04	0.14^{***}	0.08
	(0.06)	(0.11)	(0.05)	(0.11)	(0.04)	(0.13)

Note: $^{\cdot}p<0.1$; $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$

Empirical application results

Figure: Causal effect estimates of the permanent price reduction on sales of **store-brand cookies** at three different time horizons. In this table, $\hat{\bar{\tau}}_t$ is the estimated temporal average effect, while $\hat{\beta}_0$ is the coefficient estimate of the intervention dummy of REG-ARIMA.

	Time horizon:					
	1 month		3 months		6 months	
Item	$\hat{\bar{\tau}}_t$	$\hat{\beta}_0$	$\hat{\bar{\tau}}_t$	$\hat{\beta}_0$	$\hat{\bar{\tau}}_t$	$\hat{\beta}_0$
1	0.14	0.14	0.15·	0.12·	0.18***	0.16**
	(0.12)	(0.09)	(0.08)	(0.07)	[illegible]	[illegible]
2	0.14	0.10	0.13·	0.12·	[illegible]	[illegible]
	(0.12)	(0.14)	(0.08)	(0.07)	[illegible]	[illegible]
3	0.19·	0.15·	0.21**	0.15*	[illegible]	[illegible]
	(0.11)	(0.08)	(0.07)	(0.07)	[illegible]	[illegible]
4	0.49***	0.00	0.30***	0.19*	[illegible]	[illegible]
	(0.09)	(0.13)	(0.06)	(0.08)	[illegible]	[illegible]
5	−0.02	−0.06	0.07	−0.07	[illegible]	[illegible]
	(0.12)	(0.12)	(0.08)	(0.11)	[illegible]	[illegible]
6	0.24*	0.26·	0.34***	0.24*	[illegible]	[illegible]
	(0.12)	(0.14)	(0.08)	(0.12)	[illegible]	[illegible]
7	0.55***	0.75***	0.34***	0.70***	[illegible]	[illegible]
	(0.10)	(0.11)	(0.07)	(0.10)	[illegible]	[illegible]
8	0.26***	0.29**	0.25***	0.29**	[illegible]	[illegible]
	(0.08)	(0.09)	(0.07)	(0.10)	[illegible]	[illegible]
9	0.47***	0.70***	0.20***	0.29***	[illegible]	[illegible]
	(0.06)	(0.10)	(0.04)	(0.09)	[illegible]	[illegible]
10	0.66***	0.85***	0.57***	0.82***	0.33***	0.85***
	(0.11)	(0.14)	(0.08)	(0.15)	(0.06)	(0.15)
11	0.12*	0.02	0.16**	0.04	0.14***	0.08
	(0.06)	(0.11)	(0.05)	(0.11)	(0.04)	(0.13)

Note: ·p<0.1; *p<0.05; **p<0.01; ***p<0.001

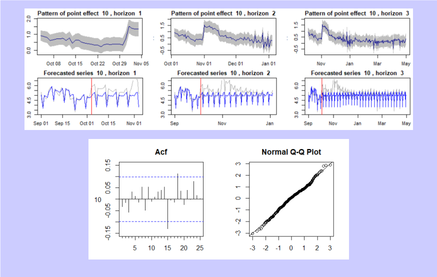

FIGURE: Causal effect estimates of the permanent price reduction on sales of **competitor-brand cookies** at three different time horizons. In this table, $\hat{\bar{\tau}}_t$ is the estimated temporal average effect, while $\hat{\beta}_0$ is the coefficient estimate of the intervention dummy of REG-ARIMA.

	Time horizon:					
	1 month		3 months		6 months	
Item	$\hat{\bar{\tau}}_t$	$\hat{\beta}_0$	$\hat{\bar{\tau}}_t$	$\hat{\beta}_0$	$\hat{\bar{\tau}}_t$	$\hat{\beta}_0$
1	−0.03	0.02	0.02	−0.16	0.04	−0.12
	(0.55)	(0.19)	(0.46)	(0.25)	(0.34)	(0.22)
2	−0.13	−0.18	−0.07	−0.13	−0.15	−0.13
	(0.50)	(0.22)	(0.47)	(0.20)	(0.36)	(0.19)
3	0.04	−0.06	0.09	−0.03	0.17	0.03
	(0.38)	(0.22)	(0.23)	(0.20)	(0.11)	(0.17)
4	0.00	0.08	−0.13	0.02	−0.04	0.01
	(0.29)	(0.21)	(0.21)	(0.22)	(0.14)	(0.13)
5	−0.03	−0.01	0.05	0.06	0.12**	0.12*
	(0.10)	(0.10)	(0.06)	(0.06)	(0.04)	(0.05)
6	−0.05	−0.01	0.03	0.06	0.09	0.10*
	(0.12)	(0.10)	(0.09)	(0.06)	(0.07)	(0.05)
7	0.04	−0.11	0.11	−0.05	0.40·	0.02
	(0.54)	(0.29)	(0.33)	(0.26)	(0.23)	(0.23)
8	−0.09	−0.02	−0.06	−0.10	−0.08*	−0.12
	(0.07)	(0.07)	(0.05)	(0.10)	(0.04)	(0.10)
9	−0.09	−0.08	−0.11	−0.11	−0.10	−0.09
	(0.13)	(0.13)	(0.09)	(0.08)	(0.06)	(0.06)
10	−0.03	−0.02	−0.12**	−0.09*	−0.11***	−0.08*
	(0.06)	(0.05)	(0.04)	(0.04)	(0.03)	(0.04)

Note: ·p<0.1; *p<0.05; **p<0.01; ***p<0.001

Figure: Causal effect estimates of the permanent price reduction on sales of **competitor-brand cookies** at three different time horizons. In this table, $\hat{\bar{\tau}}_t$ is the estimated temporal average effect, while $\hat{\beta}_0$ is the coefficient estimate of the intervention dummy of REG-ARIMA.

	Time horizon:					
	1 month		3 months		6 months	
Item	$\hat{\bar{\tau}}_t$	$\hat{\beta}_0$	$\hat{\bar{\tau}}_t$	$\hat{\beta}_0$	$\hat{\bar{\tau}}_t$	$\hat{\beta}_0$
1	−0.03	0.02	0.02	−0.16	0.04	−0.12
	(0.55)	(0.19)	(0.46)	(0.25)	(0.34)	[illegible]
2	−0.13	−0.18	−0.07	−0.13	−0.15	[illegible]
	(0.50)	(0.22)	(0.47)	(0.20)	(0.36)	[illegible]
3	0.04	−0.06	0.09	−0.03	0.17	[illegible]
	(0.38)	(0.22)	(0.23)	(0.20)	(0.11)	[illegible]
4	0.00	0.08	−0.13	0.02	−0.04	[illegible]
	(0.29)	(0.21)	(0.21)	(0.22)	(0.14)	[illegible]
5	−0.03	−0.01	0.05	0.06	0.12^{**}	[illegible]
	(0.10)	(0.10)	(0.06)	(0.06)	(0.04)	[illegible]
6	−0.05	−0.01	0.03	0.06	0.09	[illegible]
	(0.12)	(0.10)	(0.09)	(0.06)	(0.07)	[illegible]
7	0.04	−0.11	0.11	−0.05	$0.40^{\cdot}$	[illegible]
	(0.54)	(0.29)	(0.33)	(0.26)	(0.23)	[illegible]
8	−0.09	−0.02	−0.06	−0.10	-0.08^{*}	[illegible]
	(0.07)	(0.07)	(0.05)	(0.10)	(0.04)	[illegible]
9	−0.09	−0.08	−0.11	−0.11	−0.10	[illegible]
	(0.13)	(0.13)	(0.09)	(0.08)	(0.06)	(0.06)
10	−0.03	−0.02	-0.12^{**}	-0.09^{*}	-0.11^{***}	-0.08^{*}
	(0.06)	(0.05)	(0.04)	(0.04)	(0.03)	(0.04)

Note: $^{\cdot}$p<0.1; *p<0.05; **p<0.01; ***p<0.001

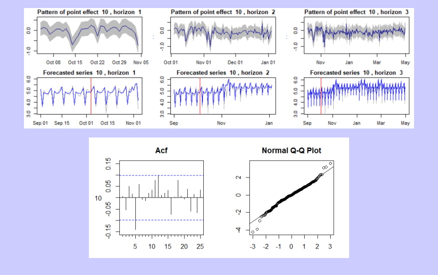

Discussion and further developments

- We proposed a novel approach, C-ARIMA, to estimate the effect of interventions in a time series setting under the RCM
- We believe that C-ARIMA can successfully be used as the frequentist alternative to CausalImpact to estimate the effect of interventions on a single time series and on multiple non-interfering series, meanwhile providing several improvements over the standard intervention analysis approach
- However, when used in our empirical context, this approach suffers from some limitations as we are not able to control for possible interactions beyond those stemming from price. Considering a setting where each store-competitor pair is modeled jointly might overcome this limitation, as discussed in (Menchetti and Bojinov; 2020).

References

Abadie, A., Diamond, A. and Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program, *Journal of the American Statistical Association* **105**(490): 493–505.

Abadie, A., Diamond, A. and Hainmueller, J. (2015). Comparative politics and the synthetic control method, *American Journal of Political Science* **59**(2): 495–510.

Bojinov, I., Rambachan, A. and Shephard, N. (2020). Panel experiments and dynamic causal effects: A finite population perspective. Preprint. Available at arXiv:2003.09915.

Bojinov, I. and Shephard, N. (2019). Time series experiments and causal estimands: exact randomization tests and trading, *Journal of the American Statistical Association* **114**(528): 1665–1682.

Box, G. E. and Tiao, G. C. (1975). Intervention analysis with applications to economic and environmental problems, *Journal of the American Statistical Association* **70**(349): 70–79.

Box, G. E. and Tiao, G. C. (1976). Comparison of forecast and actuality, *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **25**(3): 195–200.

Brodersen, K. H., Gallusser, F., Koehler, J., Remy, N. and Scott, S. L. (2015). Inferring causal impact using Bayesian structural time-series models, *The Annals of Applied Statistics* **9**(1): 247–274.

Card, D. and Krueger, A. B. (1993). Minimum wages and employment: A case study of the fast food industry in new jersey and pennsylvania, *Technical report*, National Bureau of Economic Research.

Menchetti, F. and Bojinov, I. (2020). Estimating causal effects in the presence of partial interference using multivariate bayesian structural time series models, *Preprint. Available at arXiv: 2006. 12269* .

Meyer, B. D., Viscusi, W. K. and Durbin, D. L. (1995). Workers' compensation and injury duration: evidence from a natural experiment, *The American economic review* **85**(3): 322–340.

Rambachan, A. and Shephard, N. (2019). A nonparametric dynamic causal model for macroeconometrics, *Preprint. Available at arXivpreprintarXiv: 1903. 01637* .