Long-term causal effects in policy analysis

Panagiotis (Panos) Toulis
Econometrics & Statistics, Booth School
University of Chicago
Chicago, IL, 60637
panos.toulis@chicagobooth.edu

David C. Parkes
Department of Computer Science
Harvard University
Cambridge, MA, 02138
parkes@eecs.harvard.edu

Abstract

In most causal problems we want to evaluate the long-term effects of policy changes but only have access to short-term experimental data. For example, for the long-term effects of minimum wage increase we may only have access to one-year worth of employment data. In this technical note we argue that such conceptual gap between what is to be estimated and what is in the data has not been adequately addressed. To make our criticism constructive we describe our approach in studying multiagent systems and the long-term effects of interventions in such systems. Central to our approach is behavioral game theory, where a behavioral model of how agents act conditional on their latent behaviors is combined with a temporal model of how behaviors evolve.

1 Introduction

Here is an old debate: are Democratic policies better than Republican ones? As in any election, voters in the impending 2016 U.S. presidential election have to decide based on data such as that in Figure 1, which show GDP growth for the U.S. economy over time, and political affiliation of the incumbent president [1].

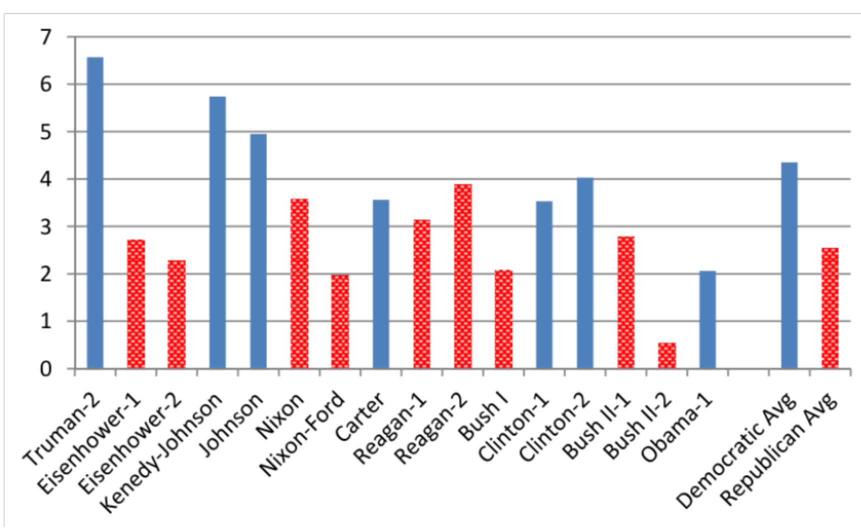

Figure 1: GDP growth (y-axis) by presidency and term (x-axis). Blue color (solid when printed in black & white) indicates affiliation with the Democratic Party and red color (checked pattern in black & white) indicates affiliation with the Republican Party. Data from Blinder and Watson [1].

30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.

To formalize the problem suppose Y is the random variable of GDP growth in some term, and Z is the binary indicator of party affiliation; if $Z = 1$ then the president in that term was a Democrat, and if $Z = 0$ the president was a Republican. For the observed values we use Y_k to denote the GDP growth observed in term k, and Z_k for the party affiliation of the incumbent president in that term. What we want to estimate here is the *Democrat-Republican performance gap* (D-R):

$$\tau = \mathbb{E}\left(Y|Z = 1\right) - \mathbb{E}\left(Y|Z = 0\right).$$

This quantity could be estimated—and very frequently is—through the sample averages,

$$\hat{\tau}(Z) = Ave(Y_k|Z_k = 1) - Ave(Y_k|Z_k = 0),$$

where Z now denotes the party affiliations observed in the data, and Ave denotes the sample conditional average, i.e., $Ave(Y_k|Z_k = a) = \sum_k Y_k \mathbb{I}\{Z_k = a\} / \sum_k \mathbb{I}\{Z_k = a\}$. In our problem this implies that the D-R gap is roughly $\hat{\tau} = +1.8\%$ (see also the right-most barplots of Figure 1), and a t-test can show that the difference is highly significant [1]. But what does this estimate really say about τ? The short answer is not much. One important problem, which unfortunately is rarely acknowledged in public discourse, is that policies do not change things instantaneously but have *long-term effects*.

Here, the long-term effect problem is related to whether we should attribute the GDP growth in one term to policies from the previous terms. If we introduce such complication the conclusions we make from Figure 1 can be drastically different. Let T denote the number of years it takes for a president to impact GDP growth. Here we consider only cases where T is a multiple of four so that the long-term horizon can be measured in presidential terms. Let $\mathcal{S}$ denote the right-shift operator, i.e., for $x = (x[1], x[2], \ldots, x[n])$ we define $\mathcal{S}x = (*, x[1], x[2], \ldots, x[n-1])$, where '$*$' denotes the null value; $\mathcal{S}^d$ denotes d successive applications of the operator, and $\mathcal{S}^0$ is the identity such that $\mathcal{S}^0 x = x$. Then the estimate $\hat{\tau}(\mathcal{S}^d Z)$ is the estimate of τ if we assume that a presidential policy requires $4d$ years to have an effect.

Figure 2 shows how the data would look like for $d = 0, 1, 2,$ and 3. When $d = 0$ we assume that policies have instantaneous effects and so we use $\hat{\tau}(Z) = +1.78\%$ to estimate τ. When $d = 1$ we assume that policies have an effect after $T = 4$ years and so we use $\hat{\tau}(\mathcal{S}Z) = +.70\%$ to estimate τ—notice that the D-R gap estimate was reduced significantly. When $d = 2$ we assume that policies have an effect after $T = 8$ years and so we use $\hat{\tau}(\mathcal{S}^2 Z) = -1.07\%$ to estimate τ, showing a negative D-R gap. This is a drastic change from our initial estimate. When $d = 3$ we have to conclude as before that Republican policies are better than Democrat policies since $\hat{\tau}(\mathcal{S}^3 Z) = -1.08\%$. For larger d (not shown here) the estimate reverses back to a positive D-R gap!

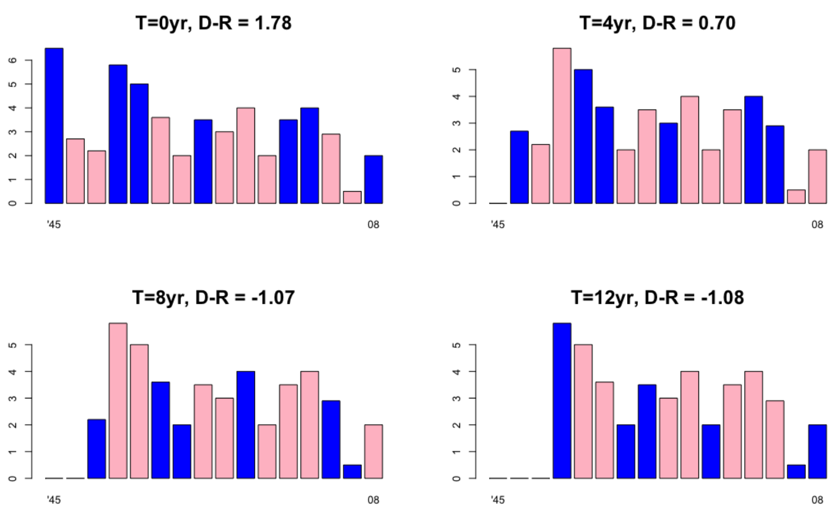

Figure 2: D-R gap estimates, indicated by "D-R" in the plots, when $T = 0, 4, 8, 12$. The colors indicating party affiliation (dark=D, light=R in black & white) are shifted to the right for every increase in T. The estimates are drastically different when $T = 0$ compared to $T = 4$ or $T = 8$.

So which analysis is correct? In such complex problems we can't really tell since we first need to understand how the underlying complex system responds to the intervention, which may heavily depend on the context; for instance, an interest rate hike by the Federal Reserve has visible short-term effects on financial indices, but an educational intervention may take generations to show any effects on educational or social indices. In this paper we aim to describe our approach to long-term policy effects by focusing on multiagent economies. This is a more tangible goal because in such economies we can leverage microeconomic information and behavioral game theory to model how agents make decisions, and combine them with temporal models of how these decisions evolve over time to estimate long-term effects.

1.1 Long-term causal effects in multiagent economies

A multiagent economy is comprised of agents interacting under specific economic rules. A common problem of interest is to experimentally evaluate changes to such rules, also known as *treatments*, on an objective of interest. For example, an online ad auction platform is a multiagent economy, where one problem is to estimate the effect of raising the reserve price on the platform's revenue. As mentioned earlier, assessing causality of such effects is a challenging problem because there is a conceptual discrepancy between what needs to be estimated and what is available in the data, as illustrated in Figure 3.

What needs to be estimated is the *causal effect* of a policy change, which is defined as the difference between the objective value when the economy is treated, i.e., when *all* agents interact under the new rules, relative to when the same economy is in control, i.e., when *all* agents interact under the baseline rules. Such definition of causal effects is logically necessitated from the designer's task, which is to select either the treatment or the control policy based on their estimated revenues, and then apply such policy to all agents in the economy. The *long-term causal effect* is the causal effect defined after the system has stabilized, and is more representative of the value of policy changes in dynamical systems. Thus, in Figure 3 the long-term causal effect is the difference between the objective values at the top and bottom endpoints, marked as the "targets of inference".

What is available in the experimental data, however, typically comes from designs such as the so-called A/B test, where we randomly assign *some* agents to the treated economy (new rules B) and the others to the control economy (baseline rules A), and then compare the outcomes. In Figure 3 the data are depicted as the solid time-series in the middle of the plot, marked as the "observed data".

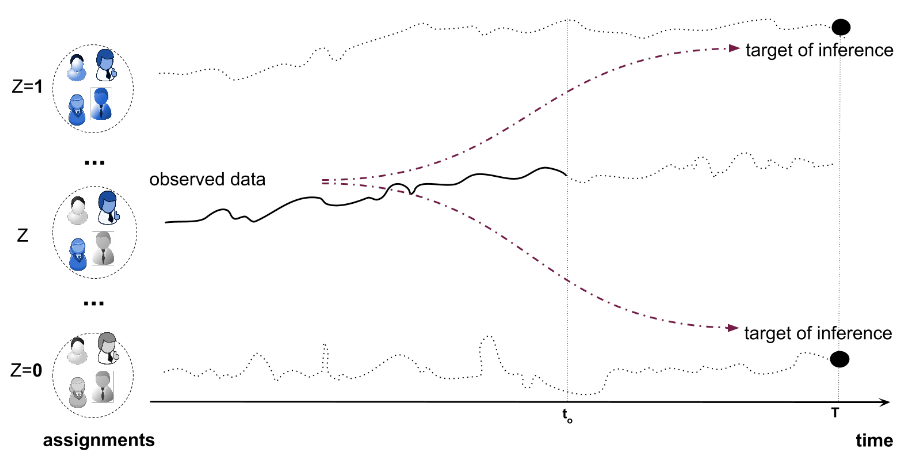

Figure 3: The two inferential tasks for causal inference in multiagent economies. First, infer agent actions across treatment assignments (y-axis), particularly, the assignment where all agents are in the treated economy (top assignment, $Z = 1$), and the assignment where all agents are in the control economy (bottom assignment, $Z = 0$). Second, infer across time, from t_0 (last observation time) to long-term T. What we seek in order to evaluate the causal effect of the new treatment is the difference between the objectives (e.g., revenue) at the two inferential target endpoints.

Therefore the challenge in estimating long-term causal effects is that we generally need to perform two inferential tasks simultaneously, namely, (i) infer outcomes across possible experimental policy assignments (y-axis in Figure 3), and (ii) infer long-term outcomes from short-term experimental data (x-axis in Figure 3). The following algorithm describes on a high-level our approach in estimating long-term causal effects. It relies on the concept of *behavior*, which defines a *behavioral model* of how agents act, and thus a mapping from the behavioral space to the space distributions over actions. The algorithm also relies on a set of assumptions that guarantee stability of certain quantities under the experimental assignment, which enables extrapolation over the y-axis in Figure 3. All details and theoretical results, such as Theorem 1 described next, are developed in the full version of this paper [5].

Algorithm 1 Estimation of long-term causal effects

1: Assume a temporal model of how population behavior evolves, and a behavior model of how population behavior predicts population action.
2: **for** $iter = 1, 2, \dots$ **do**
3:     Sample model parameters from prior.
4:     Sample initial population behavior $B_{0:0}$ at $t = 0$, which is assumed fixed but unknown.
5:     **for** both assignments $Z = \mathbf{1}$ and $Z = \mathbf{0}$: **do**
6:         Use the temporal model to sample population behaviors $B_{1:T}$, for $t = 1, \dots, T$.
7:         Set $W =$ likelihood of observed population actions (from 0 to t_0) given $B_{0:t_0}$.
8:         Sample long-term population action at $t = T$ conditional on population behavior $B_{T:T}$.
9:         Store the objective value for sampled population action at $t = T$, discounted by W.
10:     **end for**
11: **end for**

Theorem 1 (Sketch) *Suppose that the temporal and behavioral models are well-specified in Step 1, and the assumptions underlying Step 4 hold. Then, Algorithm 1 unbiasedly estimates the long-term causal effect depicted in Figure 3.*

Methodologically, our approach is aligned with the idea that for long-term causal effects we need a model for outcomes that leverages structural information pertaining to how outcomes are generated and how they evolve. In our application such structural information is the microeconomic information that dictates what agent behaviors are successful in a given policy and how these behaviors evolve over time.

A lot of burden is therefore placed on the behavioral game-theoretic model to predict agent actions, and the accuracy of such models is still not settled [2]. However, it is not necessary that such prediction is completely accurate, but rather that the behavioral model can pull relevant information from data that are inaccessible without game theory, thereby improving over classical methods. A formal assessment of such improvement, e.g., using information theory, is open for future work. An empirical assessment can be supported by the extensive literature in behavioral game theory [4, 3], which has been successful in predicting human actions in real-world experiments [6].

References

[1] Alan S Blinder and Mark W Watson. Presidents and the us economy: An econometric exploration. *The American Economic Review*, 106(4):1015–1045, 2016.

[2] P Richard Hahn, Indranil Goswami, and Carl F Mela. A bayesian hierarchical model for inferring player strategy types in a number guessing game. *The Annals of Applied Statistics*, 9(3):1459–1483, 2015.

[3] Richard D McKelvey and Thomas R Palfrey. Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38, 1995.

[4] Dale O Stahl and Paul W Wilson. Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization*, 25(3):309–327, 1994.

[5] Panos Toulis and C. David Parkes. Long-term causal effects via behavioral game theory. In *Proceedings of the 30th NIPS conference*, 2016.

[6] James R Wright and Kevin Leyton-Brown. Beyond equilibrium: Predicting human behavior in normal-form games. In *Proc. 24th AAAI Conf. on Artificial Intelligence*, 2010.