Catalogue no. 11-522-X
ISSN 1709-8211

Proceedings of Statistics Canada Symposium 2022:
Data Disaggregation: building a more representative data portrait of society

Heterogeneous causal effects of labour market programs: A machine learning approach

by Andy Handouyahia, Tristan Rikhi, Georges Awad and Essolaba Aouli

Release date: June 28, 2024

Statistics Canada Statistique Canada

Canada

Proceedings of Statistics Canada Symposium 2022
Data Disaggregation: Building a more representative data portrait of society

Heterogeneous Causal Effects of Labour Market Programs: A Machine Learning Approach[1]

Andy Handouyahia, Tristan Rikhi, Georges Awad, Essolaba Aouli[2,3]

Abstract

In this paper, we look for presence of heterogeneity in conducting impact evaluations of the Skills Development intervention delivered under the Labour Market Development Agreements. We use linked longitudinal administrative data covering a sample of Skills Development participants from 2010 to 2017. We apply a causal machine-learning estimator as in Lechner (2019) to estimate the individualized program impacts at the finest aggregation level. These granular impacts reveal the distribution of net impacts facilitating further investigation as to what works for whom. The findings suggest statistically significant improvements in labour market outcomes for participants overall and for subgroups of policy interest.

Key Words: Program Evaluations; Causal Machine Learning; Modified Causal Forests; Individualized Treatment Effects.

1. Introduction

In evaluations of active labour market programs, propensity score matching is a well-known technique used to conduct incremental impact analyses of these programs to gauge their effectiveness at helping participating Canadians. This approach produces robust estimates of average population effects; however, it is not optimal for identifying how treatment effects vary across observable characteristics, i.e., treatment effect heterogeneity. This analysis explores newly developed causal machine learning techniques to uncover these characteristics systematically.

In general, machine learning methods are primarily used either for predictive or descriptive purposes. Unlike typical machine learning algorithms, causal machine learning is not trying to predict an outcome but to estimate a net impact which is the difference between the expectations of an outcome for participants and similar non-participants. Significant efforts have been made to connect causality to predictive machine learning in different fields, including statistics, social sciences, health, and econometrics. When sample sizes are sufficiently large, these methods can

---

[1] The views expressed in research papers are those of the authors and do not necessarily reflect the opinions of Employment and Social Development Canada (ESDC) or of the federal government. This paper is based on methodologies and analyses developed in the context of Labour Market Development Agreements (LMDAs) evaluations. In developing these methodologies, evaluators at ESDC benefited from advice and peer reviews from various academic experts. In particular, the evaluation team would like to thank Professors Michael Lechner and Jeff Smith for providing advice on these evaluation studies.

[2] Andy Handouyahia, Director Strategic Planning and Methodology; Georges Awad, Manager, Evaluation; Tristan Rikhi, Evaluation Officer; and Essolaba Aouli, Manager, Strategic Planning and Methodology; Evaluation Directorate, Employment and Social Development Canada, 140 Promenade du Portage, Gatineau, Quebec K1A 0J9 Canada.

[3] The authors would like to thank Jérôme Mercier Director General, Evaluation, who provided senior level input to support the project.

estimate net impacts at a fine-grained level, thereby enabling the systematic detection of groups with heterogeneous effects. The obtained results can shed light on potential improvements to consider for future program development and the delivery of employment benefits and support measures.

2. Data

This study uses data from the Labour Market Program Data Platform. The Platform transforms administrative data from Employment Insurance (EI) part I data on EI claims, EI part II data on the Labour Market Development Agreements (LMDAs), and income tax data from the Canada Revenue Agency (CRA) into a rich longitudinal and relational database. It was developed to promote quality impact evaluations as a means for evidence-based policy for improving the well-being of Canadians. The platform contains over 25 years of data on program participants and non-participants. It consists of a large number of variables reflecting the individuals' labour market experience including the socio-demographic characteristics (e.g. age, gender, marital status, disability), their economic region and province, their qualifications (e.g. occupational group, skill levels related to last job before opening their EI claim, industry codes) as well as their labour market history (e.g. use of EI benefits and weeks, employment/self-employment earnings, use of social assistance, incidence of employment in the five years preceding participation). The study targets active employment insurance claimants who began a Skills Development (SD) intervention under LMDAs from 2010 to 2012, and followed up to 2017. The comparison group consists of program-eligible individuals who did not participate but have similar characteristics as participants. The key labour market outcome indicators for this analysis are the incidence of employment and the employment earnings. The indicators are measured on a yearly basis.

3. Methodology

The main objective of this study is to examine whether the incremental impacts from participating in labour market programs vary across participants' observable characteristics. We uncover these heterogeneous impacts using recently developed causal machine learning methods (see Athey 2019, and Athey and Imbens 2019). The causal machine learning literature developed estimators that combine the predictive power of machine learning with the microeconomic literature that are flexible enough to uncover heterogeneity while reliably estimating incremental impacts at a fine-grained level (Cockx et. al., 2019).

Causal Forests are composed of Causal Trees, introduced by Athey and Imbens (2016), and serve as the starting point of the Causal Forest literature. The trees sequentially split the data into smaller groups based on the values of the covariates until it obtains leaves (i.e., the point where no further splits can be made) with increasingly homogeneous characteristics. At the causal tree's leaves, the net impact is estimated as the difference in the average outcomes between participants and similar non-participants. Depending on the available number of observations, there may be many such splits, producing numerous leaves containing homogenous groups. Each leaf of the causal tree will have its own 'personalized' or 'individualized' treatment effect. The results from the leaves are averaged over the numerous trees from the causal forest leading to a singular estimate for the each individual instead of a single average treatment effect for the entire population.

Building on the causal forests by Athey and Wager (2018), Lechner (2019) proposes the Modified Causal Forests (MCF). MCF provides two main innovations over Causal Forests: 1) trees handle selection bias better and 2) computationally efficient and reliable way to estimate the precision of estimated net impacts at various aggregation levels using weight-based inference methods. In this study, we use the MCF because it provides robust estimates in the presence of selection bias, improving on Athey's Generalized Random Forests. For technical details on Casual Forests and Modified Causal Forests, please refer to Athey and Wager (2018) and Lechner (2019).

The MCF allows estimating net impacts at three levels of aggregation. The Individualized Average Treatment Effects (IATEs) measure the impact of participation at the finest aggregation level of the features available. On the other extreme, the Average Treatment Effects (ATEs) and ATET represent the population averages and participant population averages, respectively. The ATE and ATET are considered the classical parameters investigated in many econometric causal studies. The Group Average Treatment Effect (GATE) parameters are in between the IATEs and ATEs. It is similar to traditional sub-group analysis, where one preselects the variables before estimation and according to policy interest.

4. Results

This section reports some results for the outcome indicator 4-year annual average of incidence of employment and employment earnings using the Individualized Average Treatment Effects (IATEs) for participants. This indicator takes the average (over four years in our case) of the yearly incidence of employment which takes a value of one if an individual had reported employment earnings in a given year. The analysis will concentrate on the IATEs to detect whether the incremental impact from participating in a program varies among sub-populations of participants.

Figure 4-1 presents the distribution of estimated individualized average effects (IATEs) for participants only. Vertical continuous lines positioned at zero on the x-axis delimits regions of gain and loss of the program's net impact. The mean of the IATEs is equivalent to the ATET, which is 4 percentage points and highly significant, while the standard deviation is 8 percentage points. About 82% of the net impacts are positive, indicating that most participants benefits from the program relative to similar non-participants. These results point to substantial net impact heterogeneity. While mainly positive, there remains a non-trivial number of participants who had a lower incidence of employment than their comparable sub-group of non-participants over the four-year post participation period (18% with impacts less than 0 percentage points). Conversely, the distribution of IATEs shows that there are numerous participants with IATEs well above the ATET of 4 percentage points.

Figure 4-1
Incidence of employment - Distribution of estimated IATEs

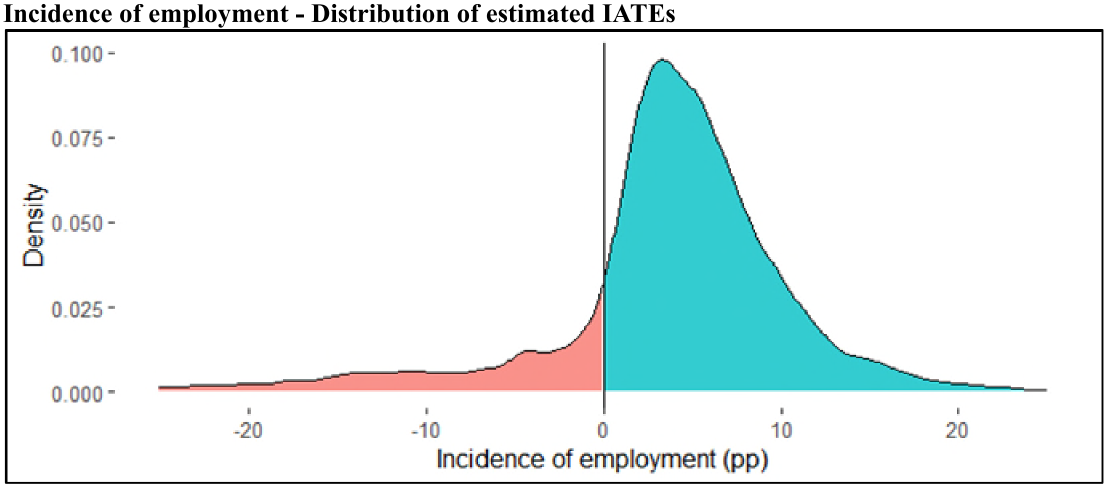

Note: *IATE = Individualized Average Treatment Effects (IATES)*
Source: Labour Market Program Data Platform of Employment and Social Development Canada

Figure 4-2 below shows the Group Average Treatment Effect by subgroups for incidence of employment. The results suggest that while older workers and females benefit the most in terms of incidence of employment, positive findings were also found for most subgroups of SD participants relative to similar non-participants.

Figure 4-2
Estimates of Group Average Treatment Effects for incidence of employment by subgroup

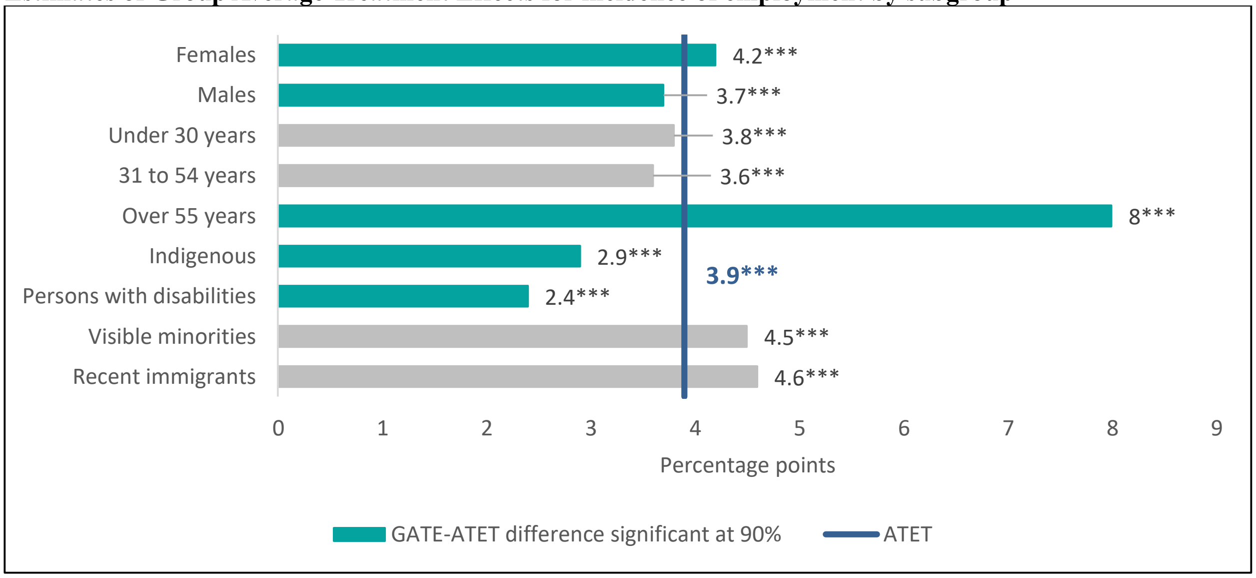

Statistical significance level *** 1%; ** 5%; * 10%
Note: *GATE = Group Average Treatment Effects (GATEs); ATET = Average Treatment Effects on Treated (ATET)*
Source: Labour Market Program Data Platform of Employment and Social Development Canada

Figure 4-3 below presents the distribution of estimated IATEs on employment earnings for participants only. The mean of the effects is the ATET which is approximately $1,997, while the standard deviation is approximately $4,880. About 68% of the net impacts are positive, indicating that the program benefits most SD participants positively. Again, we see substantial variation in IATEs and turn to Group Average Treatment Effect to get an informal characterization of the sub-groups who are benefitting or not from the program relative to non-participants.

Figure 4-3
Employment earnings - Distribution of estimated IATEs

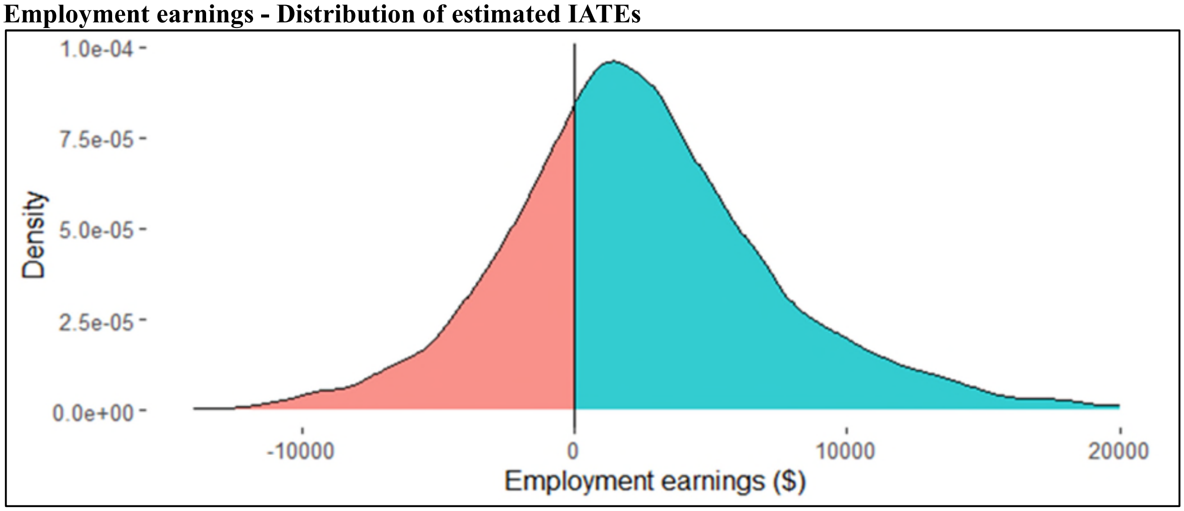

Note: *IATE = Individualized Average Treatment Effects (IATEs)*
Source: Labour Market Program Data Platform of Employment and Social Development Canada

Figure 4-4 below shows the Group Average Treatment Effect by subgroups for employment earnings. The results suggest that all subgroups experienced positive and statistically significant improvements in employment earnings following SD intervention. It was found that SD is most effective in improving the employment earnings of youth and male participants.

Figure 4-4
Estimates of Group Average Treatment Effects for employment earnings by subgroup

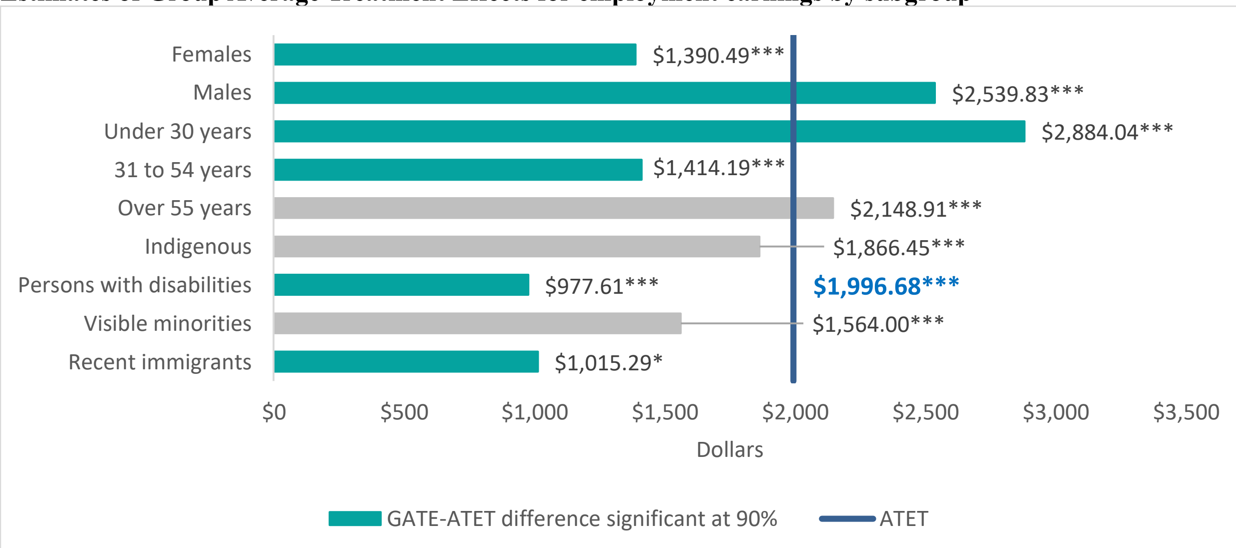

Statistical significance level *** 1%; ** 5%; * 10%
Note: *GATE = Group Average Treatment Effects (GATEs); ATET = Average Treatment Effects on Treated (ATET)*
Source: Labour Market Program Data Platform of Employment and Social Development Canada

5. Conclusion

This study implemented a causal machine learning approach, namely a modified random forest algorithm, to uncover heterogeneous effects on active employment insurance claimants who received SD intervention under LMDAs. At the granular level, the study produced individual average treatment effects, distributions of the IATEs of incidence of employment and employment earnings. The estimated standard deviations were about twice the value of ATET for all interventions indicating substantial heterogeneity in the IATEs for all outcomes.

Overall, incremental impacts demonstrate that participation in SD improves labour market attachment compared to similar non-participants. The GATEs shows that all subgroups of SD participants had, on average, positive and statistically significant improvements in incidence of employment and employment earnings.

When sample sizes are sufficiently large, this study can be extended to examine the profiles of participants with differential effects using either GATEs for socio-economic groups of interest or clustering IATEs to characterize groups who benefitted more or less from the interventions. These extended analyses will be able to, for the first time, shed light on which groups of participants are causally benefiting the most or least from program participation.

References

Athey, S., Tibshirani, J., & Wager, S. (2019), "Generalized Random Forests", *The Annals of Statistics*, 47(2), pp. 1148-

1178.

Athey, S., & Wager, S. (2018), "Estimation and Inference of Heterogeneous Treatment Effects using Random Forests", *Journal of the American Statistical Association*, 113(523), pp. 1228-1242.

Cockx, B., Lechner, M., Bollens, J. (2020), "Priority to Unemployed Immigrants? A Causal Machine Learning Evaluation of Training in Belgium", *IZA Discussion Papers*, ZA DP No. 12875.

Lechner, M. (2019), "Modified Causal Forests for Estimating Heterogeneous Causal Effects", unpublished report, St. Gallen, Switzerland: Swiss Institute for Empirical Economic Research.