MISSIONE 4
ISTRUZIONE
RICERCA

FOSSR
Fostering Open Science in Social Science Research
Innovative tools and services to investigate economic and societal change

DATA DRIVEN POLICY LEARNING: THE ROLE OF FOSSR

FOSSR Policy Brief Series

Issue #1 | October 2023

Finanziato dall'Unione europea NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

Consiglio Nazionale delle Ricerche

OUTLINE

Presentation

This policy brief presents a succinct account of the **potentials and limitations of data-driven policy learning** emerging as an innovative paradigm able to harness the power of data and predictive analytics to enhance **ex-ante policy design** – such as the optimal individuals to support or the optimal treatment to provide – across diverse socio-economic domains.

By integrating advanced technologies such as **machine learning and causal inference techniques**, policy makers can assess the potential impacts of different policy options on a wide range of societal, economic, and environmental sectors.

This shift towards **evidence-based** ex-ante **policy design** has the potential to revolutionize decision-making processes.

The **FOSSR project** envisions the creation of a **Policy Learning Platform (PLP)** that will bridge the gap between recent theoretical developments in policy learning and their practical application in real-world policies.

Using this platform, policy makers can identify **potential risks and trade-offs** before policy implementation, thereby refining **policy execution** and basing decisions on empirical evidence, which promotes transparency and accountability.

1. Introduction and questions

In a rapidly evolving world, policy makers are confronted with increasingly complex and interconnected **policy challenges**. The traditional approach to policy-making, mainly reliant on historical precedents and expert judgment (Sutherland and Burgman, 2015), may fall short in addressing these dynamic issues. **Data-driven policy learning**, an innovative approach that leverages the **power of data** and predictive analytics to anticipate the effects of policies, is a promising research area that policy-makers can consider to bettering their ability to design future policies in various areas of **socio-economic intervention** (Athey and Wager, 2021).

This policy brief delves into the concept of data-driven policy learning, and presents its potential benefits and challenges.

By harnessing cutting-edge technologies such as ***machine learning*** and ***causal inference* techniques**, policy makers can assess the impact of different policy options on various sectors of society, the economy, and the environment. This shift towards **ex-ante evidence-based policy design** has the potential to revolutionize how decisions are made and policies are formulated.

Data-driven policy learning has the potential to enhance **policy effectiveness** by reducing uncertainty and minimizing unintended consequences. By testing various scenarios in a controlled environment, policy makers can identify potential risks and trade-offs before implementation (Cerulli, 2023). This approach not only enhances policy design and implementation but also promotes accountability and transparency by grounding decisions on **empirical evidence**.

Depending on the discipline and on the specific context of application, policy learning can assume different meanings. Loosely defined, policy learning can be seen as whatever **learning process** carried out by policy makers on the basis of past policy experience that are useful to improve future policy design. In this sense, policy learning can refer to the general operation of the policy, including stakeholder alignment, policy management, policy monitoring, etc. In this document, however, I consider policy learning within the perimeter of the **quantitative data-driven** (counterfactual) **policy evaluation cycle**, generally carried out using specialized statistical and econometric techniques. Yet, policy learning also has a **qualitative nature** impacting on the **design of the policy**, which may be of the utmost importance (May, 1992).

In general, **qualitative and quantitative policy learning** should complement each other and provide fruitful cross-fertilization. The impact of both is necessary for a real bettering of the policy design and operation and the boundary between their separate role is, most of the times, blurred.

This policy brief aims at responding to a series of challenging questions: (1) **what** is data-driven policy learning? (2) **How to embed data-driven** policy learning within the policy evaluation cycle? (3) **Which statistical model** can be devised for data-driven policy learning to promote welfare maximization? (4) **What is the usefulness** of data-driven policy learning within the FOSSR project? (5) **What** are main **challenges** and limitations of data-driven policy learning? (6) **What perspective** and future development can we envisage for data-driven policy learning? In what follows, I will try to respond to all these questions.

2. Topic addressed

Within the perimeter of quantitative (or data-driven) policy learning, Figure 1 shows a **project funding policy evaluation cycle** where ex-post and ex-ante quantitative policy evaluation are properly integrated (Cerulli, 2022). Let's describe this scheme.

Given the policy eligibility rules, a unit can or cannot be eligible for the policy in question (**step 1**). If eligible, she can decide whether or not to apply for receiving the policy benefit (**step 2**). If she decides to apply (self-selection), then she can or cannot be selected by a technical committee (**step 3**) that will generally evaluate both the CV of the candidate and her proposed project. If the candidate passes also this step, she will be finally selected for support (**step 4**), thus receiving the treatment and exhibiting a certain behavior (**step 5**).

At **step 6**, the policy maker can quantitatively assess – in an ex-post manner – the effect of the policy by contrasting – in a counterfactual way – the performance of treated and untreated individuals. This is the so-called ex-post evaluation process, carried out using specialized causal inference techniques.

In this framework, policy learning comes up at **step 7**, where the policy maker can use results from ex-post evaluation for bettering the selection process (**step 3**) with the purpose of increasing policy's welfare effects.

Figure 1. Project funding Policy Evaluation Cycle

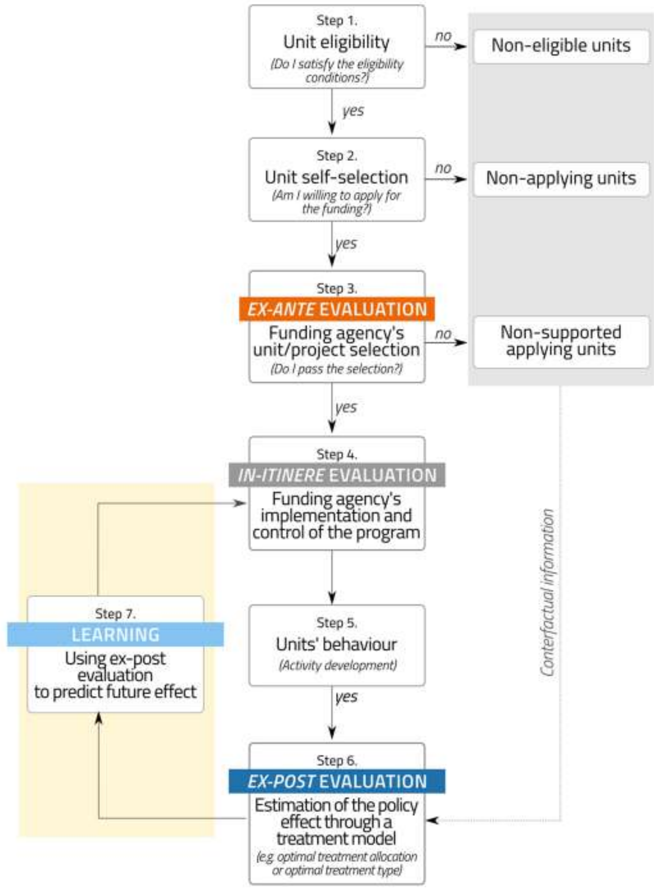

In order to appreciate how a learning process aimed at **welfare maximization** takes place, let's do two examples. The first is an example about ***optimal treatment allocation*** (Kitagawa and Tetenov 2018; Cerulli 2023); the second about ***optimal treatment type*** (Zhou et al., 2023; Atan et al., 2018).

Optimal treatment allocation

Suppose to have a setting where a policy maker has to decide who should be assigned a specific benefit (such as a grant). Given a certain monetary budget C and some eligibility constrains, the policy maker has to decide the subset of the population who should be "treated". This selection process can be carried out randomly, but the policy maker hardly will do it randomly as she wants to make welfare effects as larger as possible. If there exists an **allocation of the funds that provides larger welfare effects** than random assignment, the policy maker will be willing to implement it.

In this specific setting, **a policy can be defined as a relationship** (or mapping) between *Xi* – **the set of characteristics of individual** *i* (such as her age, education, social status, etc.) – and the **treatment variable** Ti indicating whether or not the individual should be treated:

Policy: *Xi* –> *Ti*

Every individual *i* has a potential policy effect on a given target variable *Y* (for example, personal income) that we can define as *e(Xi)*. This effect entails a **counterfactual logic**, as it can be defined as the expected outcome of the individual when treated minus the expected outcome of the same individual when untreated. Therefore, computing *e(Xi)* requires to estimate a counterfactual component, something that can be carried out using ***causal inference techniques*** (Cerulli, 2022). Under specific assumptions (e.g., selection on observable), the policy maker can thus estimate a mapping between *Xi* and *e(Xi)*. This mapping can be learnt from data obtained from past policy rounds. This can be done using specialized predictive techniques based on machine learning (Hastie et al., 2001; Cerulli, 2023a).

Given estimates of *e(Xi)* for each individual, the total welfare *W* can be defined as the sum of these estimates over all the treated units. The optimal treatment assignment *T** can be thus defined as the one producing the largest welfare. Given a unitary cost of treatment, a simple way to obtain the largest welfare is to rank individual according to their expected effect, from the largest to the smallest, providing grants sequentially to the first, the second, the third, etc. until the monetary budget runs out.

This approach is workable, but **neglects possible constrains** arising in funds allocation. For example, the policy maker would like to treat only young people, or people with a certain degree of physical disability. In these cases, the optimal policy can vary according to the level of age or disability chosen as treatment threshold. In general, this entails to reformulate the welfare maximization within specific *policy classes* such as *threshold* policies, *linear combination* policies, or *fixed-depth decision* tree policies. These **policy rules** are **context-specific** as the policy maker can decide which one to use depending on the **specific policy environment** she has to deal with (Manski, 2004). However, once a policy class has been selected, the policy maker can come up with an optimal assignment rule mapping – in an ex-ante manner – the characteristics of the individual and her expected treatment status (treated vs. untreated) able to maximize the welfare.

Optimal treatment type

Consider a scenario where three types of treatment are available: A, B, C. Based on the individual characteristics *Xi* and a corresponding measured outcome *Yi*, we are interested in providing new individuals with the treatment that is expected to generate the largest *Y*. If *Y* is personal income and A, B, and C respectively a 30-hour, a 60-hour and a 90-hour training, we would like to provide each individual with the type of training generating the largest impact on income. Similar to the previous example, this decision is embedded within a counterfactual setting, as we have to compute – given *Xi* – the level of income in three different states of the world, i.e. when the treatment is A, B, and C respectively. Under specific assumptions, we can use the data for this purpose and apply machine learning techniques to generate estimates of the **three counterfactuals**. Each individual will be assigned to the treatment providing the largest expected income, and the **social welfare** will be in this way **fully maximized**. The policy-maker can thus have a mapping between individual characteristics and treatment types which assures welfare maximiza-

Figure 2. Learning process exemples

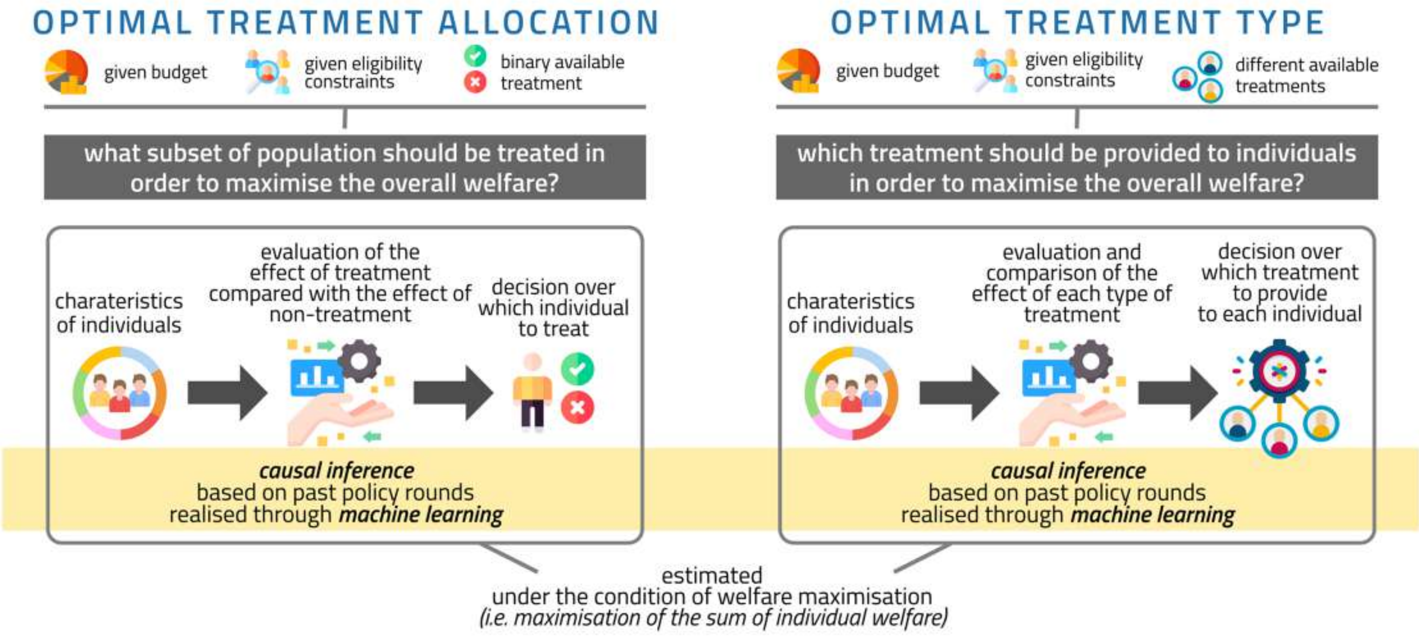

3. How FOSSR can address

The **FOSSR project** has the potential to be an exceptional resource for acquiring knowledge and expertise to enhance the effectiveness of data-driven policy learning. One of the primary objectives of the FOSSR project is to create **innovative tools and services** for gathering and analyzing data. This includes **assessing policies** both retrospectively and prospectively, with the aim of aiding users in adopting **advanced quantitative methods** and software solutions for research in the social sciences.

Through the integration of diverse data sources, infrastructures, and skill sets, and by bringing together a diverse group of experienced and young researchers from various fields such as sociology, political science, economics, statistics, and computer science, FOSSR positions itself as a valuable learning platform. This platform supports collaborative, **interdisciplinary studies** and practical projects, maintaining a focus **on data-driven policy learning**.

In the context of this policy brief, the FOSSR project envisions the creation of a **Policy Learning Platform (PLP)** that will bridge the gap between recent theoretical developments in policy learning and their practical application in real-world policies. Consistent with FOSSR's primary objectives, this platform will be **open-source** and developed using **three software tools**: two open-source options (Python and R) and one commercial tool (Stata). By harnessing and i**ntegrating data from the Cessda, Share, and RISIS** social sciences data infrastructures, this platform will enable researchers, practitioners, and policymakers **to anticipate the impacts of policies** and design targeted policy scenarios across a wide range of social domains, including immigration, aging, employment, innovation, education, and more.

4. Policy evidence

The **FOSSR's Policy Learning Platform** is expected to have a significant impact as a **tool supporting policy decisions made by policymakers**. The PLP will offer the following features:

- **Open access for policymakers**, enabling them to construct policy impact scenarios tailored to their specific areas of intervention.
- A **user-friendly web interface** for the platform's data repository, facilitating the importation of external data in standard formats for direct online processing and analysis.
- An intuitive desktop **interface that presents results**, allowing users to easily explore various policy impact scenarios by adjusting specific parameters.
- A straightforward and **user-friendly tutorial**, possibly in the form of documents and videos, accessible to all users.

Additionally, FOSSR envisions hosting **training events aimed at making data-driven policy learning** concepts and tools accessible to practitioners and individuals without technical backgrounds.

Lastly, but certainly not least, FOSSR will actively promote the communication and dissemination of the PLP through various means and institutional channels, with the goal of expanding the platform's user base.

Data-driven policy learning is not without its **limitations**. Below, I list the most relevant:

Data availability. Learning from data requires having access to a comprehensive and abundant information set. In this context, it is crucial not only to gather data but also to ensure that the data encompasses all pertinent scenarios necessary for achieving optimal policy learning. This helps prevent issues stemming from insufficient data coverage. For instance, when choosing the best treatment type, one must rely on a dataset that includes details about all available treatments. If, due to any reason, information about – let's say – treatment C is absent, it would render the computation of the optimal policy unfeasible. In this event, we can come across problems of data sparseness.

Data quality. The accuracy and reliability of predictions are contingent on the quality and relevance of the data being used. Inaccurate data can lead to biased or misleading results. Additionally, the complexity of real-world systems presents a formidable challenge. Policies can have ripple effects that extend beyond immediate areas of concern, making it difficult to account for all potential variables and interactions. This opens up challenging problems, as how to deal with policy learning in the presence of spillover effects.

Ethical considerations. Privacy concerns, data security, and potential misuse of personal information must be carefully considered to protect individuals' rights and uphold ethical standards. The selection of data sources and the transparency of methodologies used for analysis are essential to maintain public trust and legitimacy.

While **data-driven policy learning** is currently in its early stages, there exists a pressing and highly **achievable potential** for rapid expansion in both the methodological and practical aspects. Emerging methodological advancements are being documented in scholarly works, addressing areas that have not been explored thus far, such as continuous treatment scenarios or scenarios with policy spillover effects. Initiatives like **FOSSR can serve to accelerate this progress** by offering fresh perspectives to the field, fostering a fertile ground for refining methodologies, creating software solutions, and facilitating real-world applications.

References

Atan, O., Zame W.R. & van der Schaar, M. (2018). Learning Optimal Policies from Observational Data. Proceedings of the 35th International Conference on Machine Learning, Stockholm, Sweden.

Athey, S., and Wager, S. (2021). Policy Learning with Observational Data. Econometrica, 89(1), 133–161.

Cerulli, G. (2022). Econometric Evaluation of Socio-Economic Programs: Theory and Applications. 2nd edition, Springer.

Cerulli, G. (2023). Optimal treatment assignment of a threshold-based policy: empirical protocol and related issues. Applied Economics Letters, 30(8), 1010-1017.

Cerulli, G. (2023a). Fundamentals of Supervised Machine Learning: With Applications in Python, R, and Stata. Springer Series in Statistics and Computation. Springer.

Hastie T., Tibshirani, R., & Friedman, J. (2001). The Elements of Statistical Learning. Springer Series in Statistics. Springer, New York.

Kitagawa, T., and Tetenov, A. (2018). Who Should Be Treated? Empirical Welfare Maximization Methods for Treatment Choice. Econometrica, 86(2), 591–616.

Manski, C.F. (2004). Statistical Treatment Rules for Heterogeneous Populations. Econometrica, 72(4), 1221–1246.

May, P.J. (1992). Policy Learning and Failure. Journal of Public Policy, 12(4), 331-354

Sutherland, W., Burgman, M. (2015). Policy advice: Use experts wisely. Nature, 526, 317–318.

Zhou, Z., Athey, S., and Wager, S. (2023). Offline Multi-Action Policy Learning: Generalization and Optimization. Operations Research, 71(1), 148–183.

FOSSR - Fostering Open Science in Social Science Research aims to become an Italian Open Science Cloud, along the lines of the European Open Science Cloud project, in which to integrate innovative services developed by the project for data collection, data curation and Fairness, and data analysis on economic and societal change. FOSSR fosters the building of an integrated knowledge sharing platform, a single point of access to all the tools and services made available by the Italian nodes of social science infrastructures: CESSDA, SHARE and RISIS adopt the common theme of the development of Open Science in the Italian context with the goal of creating a framework of tools and services for the social science scholar community.

FOSSR wants to promote toward multiple audiences, a widespread knowledge and awareness of the data and methodologies employed in empirical social science, fostering the growth of a broad societal environment favourable to further thriving of social science research in Italy, providing easy, open, streamlined access to social science data through innovative interfaces. The integration of this pool of resources shall concretely contribute to the realization of open science for scholars in social sciences, going with an important program of scientific training on methods and instruments for social science research based on FAIR empirical data.

FOSSR Policy Brief Series aim at communicate key findings from the FOSSR thematic network to a non-specialized audience with a strong emphasis on the demonstration of usage cases of FOSSR resources. The series can accomplish two goals: improving the use of data for evidence-based policymaking and assisting the stakeholders in making informed decisions.

AUTHOR OF THE CURRENT ISSUE:

Giovanni Cerulli, CNR-IRCrES

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
Consiglio Nazionale delle Ricerche