

Faculté des sciences

Causal Inference in A/B testing

In collaboration with Euroconsumers

Author: **Nathan DEVAUX**

Supervisors: **Olivier CAELEN, Dr. Jacopo DE STEFANI**

Reader: **Anna KIRILIOUK**

Academic year 2025–2026

Master [120] en science des données, orientation statistique

I would like to express my sincere gratitude to Jacopo De Stefani, Tom Verweijen, Gert Roosens, Luna Lopez Tomas Jesus, and the entire Euroconsumers team for welcoming me at their Brussels offices. Their responsiveness throughout our exchanges, combined with the monthly meetings and valuable advice they provided, were instrumental in bringing this project to completion.

Contents

1	Introduction	5
1.1	Focus on the Data	7
2	Introduction to the notion of causality	11
3	Random Controlled Trials	17
3.1	Why Randomized Controlled Trials?	17
3.2	The Ideal RCT with Full Compliance	18
3.3	Non-Compliance and the IV Framework	19
3.4	Stratification	21
3.5	Assumptions	23
3.5.1	Stable Unit Treatment Value Assumption (SUTVA) . . .	23
3.5.2	Relevance Assumption	26
3.5.3	Exclusion Restriction Assumption	27
3.5.4	Monotonicity Assumption	28
3.5.5	Effective Randomisation Assumption	28
3.6	Intention to treat (ITT)	30
3.6.1	CACE	33
4	Machine Learning Approach	35
4.1	Why Double Machine Learning ?	36
4.1.1	G-Computation	36
4.1.2	Inverse Probability Weighting	36
4.1.3	Double Machine Learning	37
4.2	Pseudo-Observations	37
4.2.1	Construction of Pseudo-Observations	37
4.3	ITT and CACE with censored information	40
4.4	Double Machine Learning	42
4.4.1	Step 1 : Nuisance Functions	43
4.4.2	Step 2 : Partialling Out and Final Estimation	43
4.4.3	Results on the Euroconsumers Data	44
4.4.4	Unconfoundedness Assumption	45
4.4.5	Nuisance Functions : Interpretability	47
4.4.6	Financial Analysis	47

5	Meta Learners	51
5.1	Individual Treatment Effect	52
5.2	S-Learner	52
5.3	T-Learner	53
5.4	X-Learner	53
	5.4.1 X-Learner on the data from Euroconsumers	55
6	Conclusion	57
6.1	Summary of Results	57
6.2	Critical Analysis	58
6.3	Recommendations for Euroconsumers	59

Chapter 1

Introduction

The master's thesis was conducted in collaboration with Euroconsumers. Euroconsumers describes itself as a cluster composed of consumer organizations operating in five different countries: Belgium, Spain, Italy, Portugal, and Brazil. Its objective is to actively promote and defend consumers' interests. The company is considered as one of the world's leading consumer clusters in terms of innovative information, personalized services, and the protection of consumer rights. The organization works to ensure that markets are fair, consumers are safe, and benefit from honest relations with businesses and authorities.

Euroconsumers owns six categories of magazine: Consumption & comparative testing, Legal, Health, Investment, High Tech and Real Estate. Subscribers also have access to support services in four different areas: Legal, Tax, Investment and Comparative Studies. Euroconsumers employs 1,069 people, and more than 30 percent of them work in Belgium.

During the collaboration, there was an opportunity to work with experts in their respective fields, including three data scientists, a data engineer, three dashboarding experts and one business analyst. Monthly meetings were held over the course of a year, during which the theories and results developed in this thesis were presented and discussed. Their feedback and advice were invaluable in shaping the direction of this work

Euroconsumers provided data on the survival time of their clients, collected in the context of a marketing campaign designed to evaluate how clients respond to different types of incentives and whether these incentives influence their retention time. In a business model based on subscriptions, understanding what drives client retention is central. Identifying which clients are most responsive to a given incentive allows the company to design more targeted and cost-efficient campaigns in the future, this specific aspect is developed in Section ??.

More precisely, the campaign consisted of sending emails to a subset of clients.

This subset was randomly divided into two groups: one called the action group, which received an email containing either a 30-euro voucher or a 50% discount on their subscription renewal for 6 months, and the other called the control group, which received nothing. The central question of this thesis is whether receiving such an email, and potentially using the incentive, has a measurable causal effect on how long a client remains subscribed.

To answer this question rigorously, the campaign was designed as an A/B test in which clients were randomly divided into the two groups described above. This random allocation is precisely the design of a Randomized Controlled Trial (RCT), a core framework of causal inference, and is why it will be the first framework discussed in this thesis. This two-group structure is the foundation of the analysis conducted throughout this thesis.

Thus, the two complementary objectives of this thesis are to study the causal effect of this marketing campaign using rigorous statistical methods, and to summarise the related theory in a way that supports Euroconsumers in designing and evaluating future campaigns. The goal is to draw concrete conclusions from both a business and a theoretical point of view.

To draw these conclusion different approaches are used in this thesis to evaluate this effect: Randomized Controlled Trials (RCTs) and Causal Machine Learning. These two approaches are quite different and each has its own advantages and limitations. The RCT framework estimates an average causal effect on a population (ITT, CACE), and it provides a global summary of the treatment effect. The Machine Learning approach, on the other hand, can estimate a heterogeneous effect of the treatment on the population conditionally on individual characteristics (ITE). Moreover, to draw meaningful conclusions regarding the marketing study, it is also interesting to have different methods with their own assumptions, and compare their conclusions to provide a more robust evaluation of the marketing campaign.

In the case of the randomized controlled trials, the main sources of information were the works Hernán and Robins (2020), Branson and Keele (2020) and Frangakis and Rubin (2002). The work of Hernán and Robins (2020) served as the source of the theoretical framework of the RCT, as well as the general overview of the subject. There are clear explanations of the different metrics we will use later for the data. In the work of Branson and Keele (2020), the assumptions used in RCTs were presented and the test to prove the validity of these assumptions, with particular emphasis on the testing methods for the assumption of randomization, which is the most difficult assumption to verify. In the work Frangakis and Rubin (2002), the concept of principal stratification is theorized and put to the test. This work is used as principal stratification is a core concept to understand the result in RCTs.

In the case of the machine learning approach, the main sources of information

were the CausalML (Chen et al. (2020)) and EconML (Battocchi et al. (2019) benchmarks), which present detailed documentation of the methods used. In fact, both are open-source Python packages providing implementations of causal machine learning methods. Different methods introduced in these two benchmarks are explained and applied throughout this thesis. In addition, the work of Hernán and Robins (2020) was also used in the machine learning part.

The two methods presented in this thesis are representative of two different approaches, and they are based on two different sets of assumptions. The RCT approach is based on strong assumptions such as the relevance, exclusion, monotonicity and randomisation assumptions, but it offers interesting metrics, all these assumptions are explained in Section 3.5. Within the RCT approach, some of the metrics are even possible to be estimated using Machine Learning techniques. Therefore, the idea was to include a Machine Learning approach through a discussion on Causal Machine Learning, which is a new area of research introduced in the work of Schölkopf (2019). The Machine Learning approach is based on the conditional unconfoundedness assumption (Section 4.4.4).

1.1 Focus on the Data

The Euroconsumers dataset contains 16,861 observations and 12 features. Each row corresponds to one client, and each column represents a feature. The first column contains an anonymized identifier for each individual. Thus, the client's name or his address never appears in the dataset.

The dataset includes a *Gender* variable, that takes the value *M* for male, *F* for female, and *U* for unknown, only one observation has unknown gender, thus, this observation is removed from the data ($n=16,860$). The second and third columns correspond to *Age* and *Age Group*. Age ranges from 19 to 97 years old, while the age group variable is divided into four categories: 18 to 29, 30 to 44, 45 to 59, and 60 to 110. Both variables contain 798 missing observations. To avoid redundancy between Age and Age Group, only the Age Group will be used in the analysis.

Table 1.1: Distribution of individuals by age group

Age group	N
18–29	353 (2.09%)
30–44	3,014 (17.88%)
45–59	6,704 (39.76%)
60–110	6,032 (35.77%)
Unknown	757 (4.50%)

The data were collected in two waves, meaning that all the data were not collected at the exact same time. There were a difference of 3 month between the

two. The first wave includes 8,326 individuals, while the second includes 8,534 individuals. The `Treatment` variable indicates whether the individual received the email, while `Treatment_category` specifies the treatment: a 30-euro voucher, a 50-euro discount, or a missing value for individuals in the control group who did not receive the email. There are 6,749 individuals in the 30-euro voucher group, 6,744 in the 50-euro discount group, and 3,368 individuals in the control group.

In addition, the dataset contains information on client characteristics and past activity. The `Seniority` variable measures the number of months the individual has been a client and ranges from 12 to 15 months. Two additional variables capture each individual's activity on the website, measured over the six months preceding the survey and over the year preceding the survey (number of connection to their account). From these two variables, a categorical variable called `WebVisit_Group` is constructed, taking three possible values: 0 for inactive clients (no visits in the past 12 month, it represents 8,083 observations, 48%), 1 indicating moderate activity (between 1 and 33 visits, representing 4,451 observations, 26.4%), and 2 indicating frequent activity (more than 33 visits, it represents 4,316 observations, 25.6%). The threshold of 33 visits corresponds to the median number of visits among active clients, chosen to ensure a roughly balanced distribution between categories 1 and 2. Another variable reports the number of months after the survey until the individual churned, it is our target variable. All the variable and their statistics are summarized in Table 1.2.

Table 1.2: Description of the variables

Variable	Type	Description	Summary statistics
<i>Demographic characteristics</i>			
Age (with missing values)	Numerical	Age of the client (years)	mean = 55.305, std = 13.05, min = 19, max = 97
Age_Group (with missing values)	Ordinal	Age group: [18,29], [30,44], [45,59], [60,110]	[18,29] = 353, [30,44] = 3,014, [45,59] = 6,704, [60,110] = 6,032, Unknown = 757
Age_Group (without missing values)	Ordinal	Age group: [18,29], [30,44], [45,59], [60,110]	[18,29] = 353, [30,44] = 3,014, [45,59] = 7,461, [60,110] = 6,032, Unknown = 757
Gender	Categorical	Gender of the client	Male = 9,925, Female=6,935
<i>Client activity</i>			
Seniority	Numerical	Number of months the individual has been a client	mean = 13.384, std = 1.055, min = 12, max = 15
WebVisit_Group	Ordinal	Web activity category derived from yearly visits: 0 = Inactive, 1 = Low activity, 2 = Active	Inactive = 8094, Low = 4451, Active = 4316
<i>Outcome</i>			
TimeToEvent_Surv	Numerical	Number of months after the campaign until the client churned (administratively censored at 14 months)	mean = 9.21, std = 3.92, min = 1, max = 14

Table 1.3 presents 5 randomly selected observations from the dataset of Euro-consumers with all the different variables we have introduced and that we use in our analysis.

Table 1.3: Sample of randomly selected observations from the data

Age_group	Age missing	Wave	Seniority	Web visit group	Gender_M	Treatment assignment	Treatment category	Accepted Treatment	TimeToEvent Surv	Censored
[45,59]	0	2	12.0	0	0	1	50_discount	0	9	1
[45,59]	0	1	13.0	0	1	1	50_discount	0	13	1
[45,59]	0	1	15.0	1	1	1	50_discount	0	10	0
[45,59]	0	1	14.0	1	1	1	30E_voucher	1	13	1
[45,59]	0	1	13.0	0	0	0	—	0	11	0

Chapter 2

Introduction to the notion of causality

The term Causal Inference is hard to define and encompasses a wide range of methods and frameworks. In this thesis, we follow the potential outcomes as presented in Hernán and Robins (2020). The book provides an unified foundation for the methods developed in the following chapters.

Causal reasoning is part of our nature. We are perpetually trying to understand whether a given action produces a particular outcome and especially whether it caused it. In fact, the challenge of causal inference is to dissociate association from causality. Some events in the day-to-day life could be correlated without one causing the other, and the goal in causal inference is to quantify this causing between two events.

A concrete example of this distinction is provided by Hernán and Robins (2020): consider a patient A who receives a treatment and after five days he is still alive. Imagine we are able to see the conclusion in the counterfactual scenario, meaning the patient A receives no treatment and he dies five days later. In this scenario, we would have a clear evidence of causal effect as the only difference between the two scenarios being the treatment itself. In reality, it is impossible to observe both scenarios for the same individual. This problem is known as the fundamental problem of causal inference, and it has been studied in multiple frameworks and in experimental design in order to recover the causal effects from observed data.

In this thesis, the situation is organized with three core random variables: the assignment variable, the treatment variable and the outcome variable. The assignment variable Z is a binary random variable that indicates the group assignment. As quickly explained in the Chapter 1, in causal inference, there are two groups: one receiving the treatment named the treatment group or action

group and one that do not receive the treatment named the control group. Formally, Z can take three values: $Z = 0$ if the individual is in the control group, $Z = v$ if the individual receives the email with the 30-euro voucher, and $Z = d$ if the individual receives the email with the 50% discount. For simplicity, if we don't want to make distinction between the two assignments, we use the notation $Z = 1$ whenever $Z \in \{v, d\}$. Similarly, the treatment variable D is a binary random variable that indicates the actual treatment received. It takes also three values: $D = 0$ if the individual did not use any incentive, $D = v$ if the individual used the 30-euro voucher, and $D = d$ if the individual used the 50% discount. We can also simplify the notation to $D = 1$ whenever $D \in \{v, d\}$.

A common issue in A/B testing is named contamination, whereby individuals in the control group ($Z = 0$) are still able to receive the treatment and using it ($D = 1$). In our problem, this situation does not arise as we don't observe the pair $Z = 0, D = 1$ indicating that the client does not receive the email with the treatment but used the treatment anyway. This observation simplifies our problem as developed in Section 3.4.

The last core variable is the outcome variable Y , which represents the quantity we seek to causally attribute to the treatment D . It could be binary as in the example with the patient A introduced at the beginning of the chapter or it could be a continuous variable and could measure the length of time of a certain event. In our analysis, we use it as continuous outcome that measuring the number of months a client remains subscribed. Finally, X denotes the set of pre-treatment covariates available in the data (e.g., Age_Group, Seniority, WebVisitGroup,...), and are described in Section 1.1.



Figure 2.1: Visualisation of the problem

As illustrated in Figure 2.1, we clearly see the two groups: the action group ($Z = 1$) and the control group ($Z = 0$). As the individuals in the data from Euroconsumers are not fully compliant, the individual in the action group could have used the discount $D = d$, the voucher $D = v$ or they don't use the incentive

($D = 0$). The full compliance situation ($Z = D$) will be seen in Chapter 3 but it rarely holds in practice.

Table 2.1: Assignment and incentive take-up

	Count	Percentage
<i>Assignment (Z)</i>		
Email with discount ($[Z = d]$)	6,744	40.00%
Email with voucher ($[Z = v]$)	6,749	40.03%
Control group ($[Z = 0]$)	3,367	19.97%
<i>Incentive take-up (D) — Action group only</i>		
Used the discount ($[D = d Z = d]$)	636	9.43%
Used the voucher ($[D = v Z = v]$)	984	14.58%
Used a treatment ($[D = 1 Z = 1]$)	1,620	12.00%
Non-compliant ($[D = 0 Z = 1]$)	11,873	88.00%

The Table 2.1 represents the distribution of the observations in each group for the data of Euroconsumers. We see that 984 individuals used the 30-euro voucher and 636 used the 50-euro discount. We also see that 11,873 individuals are considered as non-compliant, meaning that despite being in the treatment group, they did not use the incentive. Furthermore, we have 1,620 individuals that use the treatment when they are in the action group, thus, a compliance rate of 12% without distinguishing between the two incentives. This can be separated into a compliance rate of 9.43% in the discount assignment group, and 14.58% in the voucher assignment group. These compliance rates are very important for the next chapters.

In many causal studies, the observed variables Z , D , X and Y do not capture the full picture. We need to introduce a new notation: let U_i denote the complete set of characteristics of individual i , regrouping observed and unobserved characteristics. The observed pre-treatment covariates X_i can be seen as the observable part of U_i , that is $X_i \subset U_i$, while the remainder $U_i \setminus X_i$ represents unobserved characteristics. These unobserved components are problematic as they may influence treatment decisions and outcomes.

The presence of such unobserved components is hard to represent and to handle. In this thesis, two different strategies are used to address it, corresponding of the two approaches used in this thesis. In the RCT framework, the randomisation of the assignment variable Z_i guarantees that $Z_i \perp U_i$ by design, which allows us to recover valid causal estimates despite the endogeneity of D_i , it is further developed in Chapter 3. The Double Machine Learning requires the conditional unconfoundedness assumption, meaning that X_i is sufficient to account for all confounding which is a strong assumption discussed in Section 4.4.4. The Meta-Learner framework requires the Instrumental variables assumption and

the SUTVA assumption (Section 3.5).

Following Hernán and Robins (2020), we define the potential treatment variables $D(Z = 0)$ and $D(Z = 1)$, which we summarize as $D(0)$ and $D(1)$, representing the treatment that would be received under the control and action assignment respectively. Since we don't have contamination in our data, we have $D(0) = 0$ for all individuals, meaning that no client in the control group can use the incentive without receiving the email. In addition, as reported in Table 2.1, 12% of the individuals in the action group have $D(1) = 1$, meaning they used the incentive, whereas the remaining 88% have $D(1) = 0$ and did not use it.

We can visualise the relationships between the variables using a Directed Acyclic Graph (DAG) (Pearl, 2009). In a DAG, each node represents a variable and each directed arrow $A \rightarrow B$ indicates that A has a direct causal effect on B , it indicate that A meaning that a change in A directly produces a change in B . In our case, with non-compliance ($Z \neq D$), the causal pathway of interest is:

$$Z \rightarrow D \rightarrow Y$$

This pathway means that Z has no direct effect on Y , in other words, all the effect of Z on Y is transmitted through D . To formalize this, we can introduce the notation $Y(z, d)$, where z represents the assignment and d the actual treatment receipt. The absence of a direct effect of Z on Y is written in Equation 2.1.

$$Y(z, d) = Y(z', d) \quad \forall z, z', d \quad (2.1)$$

Intuitively, Equation 2.1 means that for the same individual i , by fixing the treatment status d , the potential outcome does not depend on the assignment z . In our data, it implies $Y_i(1, 0) = Y_i(0, 0)$: a client that does not use the incentive will have the same retention time regardless of they received the email or not. In addition, the pair $Y_i(0, 1)$ does not appear in our data since we do not observe contamination.

Under these conditions, we can simplify the notation $Y(z, d)$ to $Y(d)$ with $d \in \{0, 1\}$. This simplification and this causal pathway is known as the Exclusion Restriction assumption and it is further developed in Section 3.5.3. We use $Y(1)$ to denote the potential outcome under treatment and $Y(0)$ the potential outcome without treatment through this thesis.




Figure 2.2: Kaplan-Meier Curves by Wave

Finally, as mentioned in Section 1.1, the two waves did not start at the same time and did not cover the same observation period. Figure 2.2 shows the Kaplan-Meier survival curves (Kaplan and Meier, 1958) separated by wave and by treatment assignment. The curves represent the estimated probability of remaining subscribed beyond time t :

$$\hat{S}(t) = \prod_{j: \tilde{Y}_{(j)} \leq t} \left(1 - \frac{d_{(j)}}{r_{(j)}} \right) \quad (2.2)$$

where $\tilde{Y}_{(j)}$ is the j -th ordered event time, $d_{(j)}$ is the number of individuals who churned at time $\tilde{Y}_{(j)}$, and $r_{(j)}$ is the number of individuals still at risk just before $\tilde{Y}_{(j)}$.

The Figure 2.2 suggests that the two curves are close to each other in both waves. In addition, the two waves do not cover the same observation period: Wave 1 provides up to 13 months, while Wave 2 only provides up to 9 months. At the end of these two waves some individuals have not yet churned and their exact survival time is therefore unknown. We call this administratively censored, it means that because of how the study was designed it is not possible to have information beyond a certain point. In the data, this distinction is captured by the binary variable *Censored* where 0 means the individual has churned and 1 otherwise (Table 1.3). Furthermore, in our data, it is not possible to be retained (*Censored*=1) in the middle of a wave.

Chapter 3

Random Controlled Trials

A well-suited framework for causal causal inference in our settings is the Randomized Controlled Trial (RCT), as the action and control group were already randomly constructed in the Euroconsumers data, it motivates the use of the RCT. The RCT is grounded in the potential outcome framework developed by Rubin (1980), which formalizes the notion of counterfactual outcomes and defines causal effects as comparisons between potential outcomes under different treatment conditions, following Hernán and Robins (2020).

3.1 Why Randomized Controlled Trials?

It exists other frameworks than the Randomized Controlled Trials for causal inference, but the problem is that it requires a specific data structure as the RCT that we may not have in the data from Euroconsumers. Thus, the choice of the RCT framework in this thesis is motivated by the structure of the Euroconsumers data.

For example, in the Difference-in-Differences (DiD) approach (Angrist and Pischke, 2009), it requires panel data with two observations on the same individuals both before and after the treatment. The metric we want to estimate is the Average Treatment Effect on the Treated (ATT), defined as:

$$ATT = (Y_{treated,after} - Y_{treated,before}) - (Y_{control,after} - Y_{control,before})$$

The framework also relies on his own assumption called the parallel trends assumption. The assumption requires that in the absence of treatment, the two groups would have evolved in the same way. This approach is not applicable here as each observation in the Euroconsumers data corresponds to a single client and all the result are reported after the beginning of the study. Thus, we don't have pre-treatment measurements and measurements during the study as

it is required in the DiD.

Another framework used in causal inference is the Regression Discontinuity Design (RDD) (Shaw, 2026). It exploits a continuous threshold in a running variable that determines treatment assignment, comparing individuals just above and below the threshold. The same problem arises as this framework is not applicable as the assignment Z in our data was purely determined by randomisation, and we don't have a variable or threshold that drives the treatment assignment.

Therefore, the RCT framework is the most suited approach because of the origin of the data: the assignment is determined randomly, a control group is available, and each client represents one observations after the beginning of the campaign. The RCT provides a basis for causal inference, but its validity is not unconditional, as it imposes a set of assumptions that must hold for plausible causal interpretation, all the assumptions are developed in Section 3.5.

3.2 The Ideal RCT with Full Compliance

The RCT with full compliance is a good introduction, and it will be easier to add one problem at a time and his consequences. The situation represents the simplest configuration for causal identification. In this setting, every individual follows their assigned condition exactly. Those assigned to the treatment group receive the treatment and those assigned to the control group do not.

$$D_i = Z_i \quad \forall i \quad (3.1)$$

It means that randomized assignment determines entirely the treatment receipt. And under full compliance, the causal estimand, meaning the target we seek to estimate, is the Average Treatment Effect (ATE). Formally,

$$\text{ATE} = E[Y \mid Z = 1] - E[Y \mid Z = 0] = E[Y \mid D = 1] - E[Y \mid D = 0] \quad (3.2)$$

In fact, under full compliance, the analysis is considerably simplified. Since Z is assigned by chance and $D = Z$, the ATE is reduced to a difference in means between the action and control groups, as shown in Equation 3.2. In addition, since treatment receipt D is determined solely by random assignment, it is independent of all individual characteristics, observed and unobserved, as expressed in Equation 3.3. This independence guarantees that no confounding is present, and the difference in outcomes has a clear causal interpretation.

$$D_i \perp (U_i \setminus X_i) \quad (3.3)$$

In conclusion, the ideal RCT with full compliance provides clear interpretation for causal identification. In fact, in this specific situation, all the unobserved and observed characteristics don't have an impact on the probability of using the treatment, and the causal effect is directly identify in the ATE. However, full

compliance rarely holds. In many cases, individuals may refuse their assigned treatment, or fail to adhere to it. The equality $D = Z$ no longer holds. The following section formalizes this complication and introduces the instrumental variable framework.

3.3 Non-Compliance and the IV Framework

In the ideal RCT presented in Section 3.2, the equality $D = Z$ held by construction, ensuring that randomized assignment is entirely determined by treatment receipt. In practice, the individuals may refuse their assigned treatment (e.g. $D = 0, Z = 1$) or receive the treatment despite being assigned to the control group ($D = 1, Z = 0$) (not treated in the thesis). When some individuals refuse the action treatment ($Z = 1$), we call it a problem of non-compliance, and it breaks the fundamental properties of the ideal RCT meaning that D is no longer independent of unobserved characteristics anymore. In fact, when non-compliance is present, treatment receipt is no longer determined solely by randomization but is also driven by individual characteristics (U_i). This reasoning motivates Equation 3.4. The equation shows that the use of the treatment is not independent from characteristics anymore.

$$D_i \not\perp (U_i \setminus X_i) \quad (3.4)$$

In the Euroconsumers data, the probability of using the voucher or the discount can also be impacted by individual characteristics. A client particularly motivated to save money will have more chance to use the voucher or the discount than a client who is indifferent to money. We can have the same reasoning for the retention time Y . In fact, a client attached to Euroconsumers and that will never end his subscription does not have the same expected retention time as a client who is struggling financially and no longer wishes to maintain his subscription. All these characteristics could be observed (X) or unobserved ($U \setminus X$) in the data.

When non-compliance is present, the independence between the use of the incentive D and characteristics U is broken, making it impossible to directly attribute differences in outcomes to the treatment alone.

$$E[Y \mid D = 1] - E[Y \mid D = 0] \quad (3.5)$$

For example, Equation 3.5 makes a biased estimation of the causal effect as the characteristics of the individual represented in the first term are not the same as those in the second term. Formally,

$$P(X \mid D = 1) \neq P(X \mid D = 0)$$

It is also shown in Figure 3.1, both observed and unobserved characteristics influence the decision to use the incentive D , making D endogenous (defined

later). For example, a client facing financial problems is more likely to use the discount to reduce the cost of his subscription, but it will also affect his expected retention time. Consequently, the Equation 3.5 does not isolate the causal effect of D on Y as it captures the systematic differences in retention that would have existed between treated and not treated. This problem is known as confounding meaning that the effect of D on Y is confounded by the characteristics U that simultaneously impacts use of the incentive D and the outcome Y .




Figure 3.1: DAGs on Euroconsumers Data

The solution to correct for the confounding effect of individual characteristics is to use the Instrumental Variable (IV) framework. In this framework, a variable is said to be endogenous when it is correlated with unobserved characteristics $U \setminus X$. This means that part of its variation is driven by factors outside the model, making it impossible to isolate the causal effect of that variable on Y . In our setting, D is endogenous because the decision to use the voucher or discount is partly driven by unobserved individual motivation, as illustrated in Figure 3.1 and discussed just before.

Conversely, a variable is exogenous if it is independent of $U \setminus X$, meaning its variation is not driven by any unobserved factor. In our setting, Z is exogenous by design since it is randomly assigned, regardless of any individual characteristic. We can visualise it in Figure 3.1 where the characteristics don't have direct effect on Z .

The goal of the instrumental variable (IV) is to use this exogenous variable Z to recover a valid causal effect and delete the effect of the characteristics on D . In our analysis, our variable Z is simultaneously the assignment variable and serves as an instrument meaning it is used as a source of exogenous variation in D to correct for its endogeneity.

$$E[D \mid Z = 1] - E[D \mid Z = 0] \quad (3.6)$$

As Z is randomly assigned and the characteristics are thus balanced in the ac-

tion and control group in expectation. Equation 3.6 captures only exogenous variation in the use of the incentive D induced by the assignment, not the part driven by individual characteristics. For example, a client with financial difficulties (more likely to use the incentive and more likely to cancel his subscription) is expected to be equally represented in the two assignment groups. Therefore, their characteristics cancel out in comparison, and it isolates the effect of Z on D from any confounding influence.

Similarly, by making the difference of expected value between the action group and control group represented in Equation 3.7, we also cancel the effect of the characteristics (U) on the retention time (Y).

$$E[Y \mid Z = 1] - E[Y \mid Z = 0] \quad (3.7)$$

All the reasoning holds if the assignment variable Z is a valid instrument. In RCT, the assumption holds as Z normally satisfy three conditions. First, it must be uncorrelated with the unobserved characteristics. Secondly, it must be correlated with D . Thirdly, it must affect Y only through D . All these conditions are formalized and developed in Section 3.5.

3.4 Stratification

As explained in Section 3.3, non-compliance breaks the similarity between the randomized assignment Z and the treatment actually received D , and the instrumental variable framework recovers a valid causal effect despite the endogeneity of D . However, this causal effect cannot be identified for the entire population. The framework of principal stratification, introduced by Frangakis and Rubin (2002), provides a rigorous framework to identify the subpopulations for which a causal effect can be meaningfully estimated.

The approach proposed by Frangakis and Rubin (2002) is to partition the population into four latent subgroups based on their potential treatment behavior under each possible assignment. These subgroups correspond to the four possible pairs of values $D(0), D(1)$. Thus, individuals are classified according to how they would behave under each possible assignment.

When the post-treatment variable is binary, four principal strata can be defined:

$$(D(0), D(1)) \in \{(0, 0), (0, 1), (1, 1), (1, 0)\}.$$

These strata are named as follows:

- **Never-takers:** $D(0) = 0$ and $D(1) = 0$,
- **Compliers:** $D(0) = 0$ and $D(1) = 1$,
- **Always-takers:** $D(0) = 1$ and $D(1) = 1$,

- **Defiers:** $D(0) = 1$ and $D(1) = 0$.




Figure 3.2: Stratification

Figure 3.2 illustrates the four principal strata and their relationship to the assignment Z and treatment D . The first stratum named the Never-Takers is the group of individuals who never take the treatment regardless of their assignment. The second stratum called the Compliers are the only group that follow their corresponding assignment, if they are in the control group they won't use treatment if they are in the action group they will use the treatment. The third stratum is the always takers, and independently from the assignment, they will use the treatment, in our case, we don't have Always-takers as it is not possible to use the treatment when the individuals are in the control group. The last stratum, the Defiers do the opposite of their assignment treatment by using the treatment when they are in the control group and by not using the treatment when they are in the action group, this group is also not present in our study.

Here, the difficulty is that only one element of the pair can be observed for all the individuals, the unobserved element is called the counterfactual potential outcome. Indeed, if an individual is assigned to the control group, only $D(0)$ could be observed, and if the individual is assigned to the treatment group, only $D(1)$ is observed. As a result, one element always remains latent characteristic in the pair $D(0), D(1)$. In our case, without contamination, we are able to separate the individual in the action group into Never-takers and Compliers, but in the control group we are not able to set which individual is in which group.

In order to make an example, by returning to the five random observations in Table 1.3, we can verify how our data are concretely distributed. The first three observations are in the action group as Treatment assignment = 1 but they

did not use the incentive (Accepted Treatment = 0), the observed component is therefore $D(1) = 0$ and we don't have the counterfactual ($D(0) = ?$), it cannot be observed. However, since neither Always-takers nor Defiers exist in our data, $D(0) = 0$ for all individuals in the action group, we can therefore conclude that the first three observations in Table 1.3 are Never-Takers. Applying the same reasoning, the fourth observation used the incentive, so $D(1) = 1$ and we know that the other pair is $D(0) = 0$ as we don't observe Defiers or Always-takers ($D(0) = 1$ is not possible). The fourth observation is therefore a Complier. The last individual belongs to the control group $D(0) = 0$. By construction, we cannot observe what would have been the reaction in the action group ($D(1) = ?$), we cannot determine whether this individual is a Never-taker or a Complier.

Among individuals assigned to the action group, 12% used the treatment, meaning that 12% of the action group are Compliers and 88% are Never-takers. Among these strata, the most relevant group to study the causal inference is the group of Compliers whose treatment status is actually determined by the instrument Z . Never-takers and always-takers would behave the same way regardless of their assignment, so the instrument carries no information about the treatment effect for them. Defiers are ruled out by the Monotonicity Assumption in Section 3.5. Therefore, Compliers are the only subgroup that truly reacts to the intervention, and consequently the only group for which the IV framework can identify a causal effect, as formalized through the CACE in Section 3.6.1.

3.5 Assumptions

The Instrumental Variable (IV) assumptions of the RCT framework must be satisfied for the assignment variable Z to serve as a valid instrument. When these assumptions are fulfilled and empirically supported, the Instrumental Variable ensures that the confounding effect of all individual characteristics cancels out in the comparison between the two groups, allowing the estimations of the Intention-to-Treat (ITT) and the Complier Average Causal Effect (CACE) to have a valid causal interpretation.

3.5.1 Stable Unit Treatment Value Assumption (SUTVA)

Before starting the core assumptions of the IV framework, a preliminary condition must hold for causal inference to be coherent. The condition is known as the Stable Unit Treatment Value Assumption or SUTVA. It was formally introduced by Rubin (1980) and it is discussed in both Hernán and Robins (2020) and Shaw (2026).

The SUTVA assumption is decomposed into two distinct conditions.

- **No interference:** the potential outcome of unit i depends only on its own treatment assignment, thus, the potential outcome does not depend on the treatment received by any other unit $j \neq i$. For instance, if an individual received an email with a voucher and shared it with another client, this could influence the retention decision of that other client regardless of their own assignment, which would constitute a violation of this condition.
- **Treatment-version irrelevance:** it exists only one version of each treatment level. It means that the assignment $D = 1$ must have the same meaning for all units, so that the potential outcome $Y_i(d)$ is uniquely defined regardless of how the treatment was delivered.

If the two conditions hold, the potential outcomes for individual i if he is using the incentive $Y_i(1)$ and if he is not using the incentive $Y_i(0)$ are well defined and the observed outcome can be written as:

$$Y_i = D_i Y_i(1) + (1 - D_i) Y_i(0)$$

Where D_i represents whether individual i used the incentive and Y_i represents their observed outcome. Without at least one of these conditions, this equality breaks down. Either the potential outcomes are no longer independent across units, making individual comparisons meaningless. Or the treatment does not have a unique definition, preventing any consistent causal interpretation of $Y_i(d)$. In the case of RCT setting, SUTVA is generally considered plausible. However, as Shaw (2026) notes, violations are frequent and difficult to detect in field settings where spillovers, word-of-mouth effects, or competitive reactions are present.

In our setting, we make the assumption that clients are unlikely to influence each other's retention based on their assignment, and no spillover, word-of-mouth exists. With this the no interference assumption holds. In our data, the assignment variable for the discount ($Z=d$) or voucher ($Z=v$) do not have the same impact on retention. Treating them as a single version of the treatment would violate the SUTVA assumption, as $Y(1)$ would not have a unique meaning across all individuals.

To restore the Treatment-version irrelevance condition, we have to separate the data into two subgroups: one for the voucher and one for the discount, and we can compare it against the same control group ($Z = 0$) as a baseline, Figure 3.3 represents the new two independent sub-studies. With this new configuration the treatment version is unique and well-defined in each study and the SUTVA assumption is satisfied.




Figure 3.3: Visualisation of the splitting of the study into 2 new study to respect SUTVA assumption

The separation does not fundamentally change the conclusions as we simply identify two distinct causal effects rather than a single one. However, some results in this thesis violate the treatment-version irrelevance condition. Those results will have to be interpreted with caution.

Other solutions exist to solve this heterogeneity of treatment without separating the study such as the Causal Forest proposed by Wager and Athey (2018). It estimates a Conditional Average Treatment Effect that varies across individuals based on their observed covariates and it requires conditional unconfoundedness assumption $Y_i(d) \perp D_i \mid X_i$, which is stronger than the IV assumptions used in this thesis, and it does not resolve the SUTVA violation in a strict sense but it rather sidesteps it by modelling heterogeneity directly. For these reasons, this approach falls beyond the scope of this thesis.

3.5.2 Relevance Assumption

The next assumption is named the Relevance Assumption. It requires that the instrument Z has a non-null effect on treatment uptake D , conditionally on covariates X . Formally,

$$E[D|X, Z = 1] \neq E[D|X, Z = 0] \quad (3.8)$$

As noted in Hernán and Robins (2020), it captures the non-null association between Z and D . In the RCT setting, this assumption is expected to hold by construction: since the instrument Z corresponds to the randomized assignment itself, it is specifically designed to influence whether individuals receive the treatment, making a non-null first-stage relationship structurally guaranteed. A way to conceptually see this is by the use of the Wald estimator (equivalent to the CACE):

$$\hat{\beta}_{IV} = \frac{E[Y | Z = 1] - E[Y | Z = 0]}{E[D | Z = 1] - E[D | Z = 0]} \quad (3.9)$$

The estimator is the simplest form of the IV estimator. The numerator captures the difference in expected outcome across instrument values, while the denominator scales it by the difference in treatment uptake knowing the instrument Z . If the relevance assumption is violated, the denominator is not valid anymore as it goes to zero, making the estimator undefined. It also intuitively prove the causal pathway between the assignment variable Z and the use of the incentive D in the DAG represented in Figure 3.1.

More formally, the failure of relevance goes beyond an undefined estimator. Consider the first-stage regression:

$$D_i = \alpha + \gamma Z_i + \varepsilon_i \quad (3.10)$$

The goal in this regression is to find the probability for an individual to use the effective treatment. And the relevance assumption requires $\gamma \neq 0$ as Z_i has to have an impact on D_i .

Thus, to verify the Relevance Assumption, we estimate the equation 3.10 and test the null hypothesis: $H_0: \gamma = 0$ against $H_1: \gamma \neq 0$. The results are summarised in Table 3.1. The result shows that $\hat{\gamma} = 0.12$ and is highly significant ($p < 0.001$) and the F-Statistics is above 10. It confirms that Z is a strong instrument for D and the Relevance Assumption is therefore empirically supported.

Table 3.1: First-Stage Regression Results: $D_i = \alpha + \gamma Z_i + \varepsilon_i$

Variable	Coefficient	Std. Error	t-stat	p-value
Intercept	0.000	0.005	0.000	1.000
Z_i (Treatment assignment)	0.120	0.006	21.436	< 0.001
<i>Model diagnostics</i>				
R^2	0.027			
F-statistic	459.5			
Prob(F-stat)	< 0.001			
Observations	16,861			

3.5.3 Exclusion Restriction Assumption

The third assumption is named the Exclusion Restriction assumption. It requires that the instrument Z affects the outcome Y only through the effect of the treatment D . Thus, the instrument Z has no direct effect on the outcome Y . As introduced in Chapter 2 the Equation 3.11 holds.

$$Y_i(z, d) = Y_i(z', d) \quad \forall, z, z', d \quad (3.11)$$

Where z and z' denote two arbitrary values of the assignment variable Z , it could be 0 or 1, and d denotes a fixed value of the treatment D , meaning it could be 0 or 1. Intuitively, if $d = 0$, it could be any value of Z at the end it will be the same outcome. Conversely, if $d = 1$, it also have the same outcome for every value of Z . Thus, this condition gives the opportunity to write the potential outcomes to be written simply as $Y_i(d)$ and it remove any dependence on Z . In fact, the instrument Z must operate as a nudge, it shifts who receives the treatment, but at the end, this is the treatment status D that influence the outcome Y . The only pathway from Z to Y must go through D :

$$Z \rightarrow D \rightarrow Y \quad (3.12)$$

As explained in Hernán and Robins (2020), in the RCT setting, this assumption is expected to hold in a double blind randomized situation, meaning that individuals do not know their assignment. In our setting, if the study is not double-blind: clients who receive the email are aware that they have been assigned to the action group, and conversely, clients who do not receive any email may become aware that a marketing campaign is ongoing, this could bias our estimation, for instance, if the clients in the control group feel excluded or by actively seeking an alternative subscription, it will have an impact on the retention time. Such a violation would bias both the ITT and the CACE estimates, as the difference $E[Y \mid Z = 1] - E[Y \mid Z = 0]$ would no longer delete the effect of the confounding through the assignment variable Z , but also a direct awareness effect specific to one of the group ($Z = 0$ or $Z = 1$). In our setting, we assume this violation is unlikely given the nature of the campaign but we know that it is a potential threat to the validity of the Exclusion Restriction.

Unlike the Relevance Assumption, the Exclusion Restriction can't be testable from the data. With only one instrument, it doesn't exist statistical test that can confirm the absence of the direct pathway (3.12). And if the assumption doesn't hold, as cited in Shaw (2026), it could produce biases in the IV estimates.

3.5.4 Monotonicity Assumption

The fourth assumption is called the Monotonicity assumption. It requires that the instrument Z affects the probability of using the treatment D in the same direction for all individuals in the population. Formally,

$$D_i(1) \geq D_i(0) \quad \forall, i \quad (3.13)$$

This condition refers to the group defiers introduced in the stratification section 3.4. As a reminder, the defiers are the individuals who would take the treatment if assigned to the control group and would not take it if assigned to the treatment group. It means that the defiers are systematically reacting to the instrument in the opposite direction and would have the opposite effect of Equation 3.13:

$$D_i(1) < D_i(0) \quad \forall, i \quad (3.14)$$

In our setting, a defier would be an individual who would use the incentive if he is not receiving the email ($Z = 0$) but would not use it if he is receiving the email ($Z = 1$). It is not observed in our data: since individuals in the control group never receive the email, they cannot access the incentive, so $D_i(0) = 0$ for all individuals.

Like the Exclusion Restriction, monotonicity cannot be formally tested. However, in our case, it is a structural property of the experimental design that guarantees the absence of defiers.

3.5.5 Effective Randomisation Assumption

The last assumption in the IV framework is the Effective Random Assignment assumption. It is the strongest assumption and it requires that the instrument Z is independent of the potential treatments $D(z)$, and the potential outcome $Y(z, d)$, conditionally on covariates X . It means that neither the latent propensity to comply nor the potential retention time of an individual should influence who receives the email. Formally,

$$Z \perp \{D(z), Y(z, d)\} \mid X \quad (3.15)$$

Since the covariates X are pre-treatment variables, and are determined before the assignment Z , it implies a stronger marginal independence that assumption have to test:

$$Z \perp X \quad (3.16)$$

This independence implies the following balancing condition, which can be verified empirically:

$$E[X | Z = 1] = E[X | Z = 0] \quad (3.17)$$

Intuitively, the distribution of observed covariates should be identical in the two assignment groups. And this is the condition verified empirically by the two tests below. In our setting, the assumption normally holds by construction since the assignment was determined randomly of any individual characteristic.

The first test is called the Love Plot and it shows how the covariates are balanced across the two groups, if the mean differences defined by :

$$\text{SMD}_k = \frac{E[X_k | Z = 1] - E[X_k | Z = 0]}{\sqrt{\frac{\text{Var}(X_k | Z=1) + \text{Var}(X_k | Z=0)}{2}}} \quad (3.18)$$

is below the conventional threshold of 0.1 (Austin, 2009), the assumption is considered satisfied. k indexes each covariate in X and the threshold is tested for each covariate one by one. In Figure 3.4, we see that each distribution of the covariates is well balanced across the two groups, supporting the plausibility of the assumption.




Figure 3.4: Love Plot

The second test was introduced by Branson and Keele (2020). The intuition behind the test is to simulate how the covariates would be balanced if the two groups were rebuilt under true randomization. If Z is truly as-if randomly assigned, the observed covariate balance should be consistent with what we would

expect from a completely randomized experiment. It is important to note that the test cannot prove that Z is truly random, the goal is rather to falsify the assumption if the observed balance is inconsistent with randomisation and prove that Z is as-if random.

The test exploits the exchangeability of Z under the null hypothesis: $H_0 : Z \perp X$. It means that if the assignment was determined purely by chance than any permutation of the observed assignment is equally possible. We can generate $M = 10,000$ random permutations of Z , compute the Mahalanobis distance for each permutation, and it will form the reference distribution of balance under true randomization. In this distribution, the original randomisation p-value is defined as the proportion of permutations yielding a Mahalanobis distance at least as large as the observed one. We interpret a large p-value as an observed balance that is consistent with the true randomisation, controversially, a small p-value means an imbalance under H_0 .

In our data, the test yields a p-value below 0.05 when all covariates are included, leading to a rejection of H_0 . However, this result is driven by the Seniority variable, whose pooled standard deviation is very small due to its restricted range of 12 to 15 months. As a consequence, even negligible mean differences between groups produce inflated SMD values, defined as:

$$SMD_k = \frac{\bar{X}_{k,Z=1} - \bar{X}_{k,Z=0}}{S_k} \quad (3.19)$$

In our data, the mean Seniority is 13.335 months for the control group and 13.396, a difference of less than two days which is negligible. When Seniority is excluded from the test, the p-value is above 0.05 and H_0 is no longer rejected, confirming that the Randomisation Assumption is plausible for all remaining covariates.

3.6 Intention to treat (ITT)

All the different assumptions are fulfilled, we can now detect the effect of the treatment on the retention time. The first metric that we identify in the Instrument Variable (IV) Framework is the Intention-To-Treat:

$$ITT = E[Y \mid Z = 1] - E[Y \mid Z = 0] \quad (3.20)$$

The ITT represents the average difference in outcomes between the assigned treatment group ($E[Y|Z = 1]$) and the assigned control group ($E[Y|Z = 0]$). The metric could be seen as conservative as it captures only the effect of being assigned to the treatment group and does not capture the effect of the treatment. As a result, the true effect of the treatment is diluted across the entire population, and it includes effects for whom the assignment had no behavioral impact. In fact, as explained in the stratification section, and under the Monotonicity assumption, the population could be seen as three strata: never-takers,

always-takers and compliers. (Defiers are not take into account because of the Monotonicity assumption, Section 3.5.4)

1. **Never-takers** ($D(0) = 0, D(1) = 0$): Under the Exclusion Restriction, the instrument Z affects the outcome Y only through the treatment D . Meaning that if the never-takers never use the treatment than their outcome Y is the same in the control and in the action group:

$$\begin{aligned} E[Y(D(1)) - Y(D(0)) \mid D(1) = 0, D(0) = 0] \\ &= E[Y(0) - Y(0) \mid D(1) = 0, D(0) = 0] \\ &= E[0 \mid D(1) = 0, D(0) = 0] \\ &= 0 \end{aligned} \tag{3.21}$$

where $Y(D(1))$ represents the potential outcome in action group and $Y(D(0))$ represents the potential outcome in control group.

2. **Always-takers** ($D(0) = 1, D(1) = 1$): We know that we don't have always-takers in our data. Nevertheless, the same reasoning applies, the assignment treatment is not affecting their outcome as they don't change their treatment status, thus, their contribution to the ITT is also null:

$$\begin{aligned} E[Y(D(1)) - Y(D(0)) \mid D(1) = 1, D(0) = 1] \\ &= E[Y(1) - Y(1) \mid D(1) = 1, D(0) = 1] \\ &= E[0 \mid D(1) = 1, D(0) = 1] \\ &= 0 \end{aligned} \tag{3.22}$$

3. **Compliers** ($D(0) = 0, D(1) = 1$): Their treatment status is determined by the instrument Z meaning that their retention time may change. It is precisely within this stratum that the assignment induces a change in behavior, and therefore a change in outcome. Denoting the causal effect for compliers as τ :

$$\begin{aligned} E[Y(D(1)) - Y(D(0)) \mid D(1) = 1, D(0) = 0] \\ &= E[Y(1) - Y(0) \mid D(1) = 1, D(0) = 0] \\ &= \tau \end{aligned} \tag{3.23}$$

Knowing that the impact of the never-takers and always-takers on the ITT is null, the only stratum that result in an observed difference between the two group is the compliers. If π_C, π_N and π_A , represents the proportion of Compliers, Never-takers and Always-takers respectively in the population, we can decomposed the Equation 3.6 as:

$$\begin{aligned} E[Y \mid Z = 1] &= E[Y(D(1)) \mid D(0) = 1] \cdot \pi_A \\ &\quad + E[Y(D(1)) \mid D(0) = 0] \cdot (\pi_N + \pi_C) \end{aligned} \tag{3.24}$$

$$\begin{aligned}
E[Y \mid Z = 1] &= E[Y(D(1)) \mid D(1) = 1, D(0) = 0] \cdot \pi_C \\
&\quad + E[Y(D(1)) \mid D(1) = 1, D(0) = 1] \cdot \pi_A \\
&\quad + E[Y(D(1)) \mid D(1) = 0, D(0) = 0] \cdot \pi_N
\end{aligned} \tag{3.25}$$

$$\begin{aligned}
E[Y \mid Z = 0] &= E[Y(D(0)) \mid D(1) = 1, D(0) = 0] \cdot \pi_C \\
&\quad + E[Y(D(0)) \mid D(1) = 1, D(0) = 1] \cdot \pi_A \\
&\quad + E[Y(D(0)) \mid D(1) = 0, D(0) = 0] \cdot \pi_N
\end{aligned} \tag{3.26}$$

$$\begin{aligned}
\text{ITT} &= E[Y(D(1)) - Y(D(0)) \mid D(1) = 1, D(0) = 0] \cdot \pi_C \\
&\quad + E[Y(D(1)) - Y(D(0)) \mid D(1) = 1, D(0) = 1] \cdot \pi_A \\
&\quad + E[Y(D(1)) - Y(D(0)) \mid D(1) = 0, D(0) = 0] \cdot \pi_N
\end{aligned} \tag{3.27}$$

We can simplify the notation with the Equation 3.21, Equation 3.22 and Equation 3.23;

$$\text{ITT} = \tau \cdot \pi_C + 0 \cdot \pi_A + 0 \cdot \pi_N = \pi_C \cdot \tau \tag{3.28}$$

In this decomposition we understand that the ITT is a diluted version of the causal effect τ . The next section introduce the CACE that avoid this dilution.

In Figure 3.5, we see the survival curves separated by the assignment variable. This visualisation is directly related to the ITT. In fact, the vertical gap between the blue and red curve at a certain time t is visual analog of the ITT effect on the survival probability. It suggest a small but positive ITT effect as the two curves are close to each other.




Figure 3.5: KM-Curve by assignment variable Z

The ITT cannot be computed directly from raw survival times, as some observations are censored. A naive difference in means would therefore yield a biased estimate. We have to use the Restricted Mean Survival Time (RMST) that will be introduced in the next chapter.

3.6.1 CACE

The Compliers Average Causal Effect is a good method to find the causal effect of the treatment on the outcome:

$$CACE = \frac{E[Y|Z = 1] - E[Y|Z = 0]}{E[D|Z = 1] - E[D|Z = 0]} \quad (3.29)$$

In reality, the CACE is the ITT divided by the proportion of compliers. Since there are no always-takers ($\pi_A = 0$) and no contamination ($E[D | Z = 0] = 0$), we have:

$$\begin{aligned} E[D | Z = 1] &= 1 \cdot \pi_C + 1 \cdot \pi_A + 0 \cdot \pi_N \\ &= \pi_C \end{aligned} \quad (3.30)$$

Therefore,

$$E[D | Z = 1] - E[D | Z = 0] = \pi_C - 0 = \pi_C \quad (3.31)$$

And,

$$CACE = \frac{ITT}{\pi_C - 0} = \frac{ITT}{\pi_C} \quad (3.32)$$

It is important to note that the CACE is not biased and it is not affected by the confounders effect. The resulting CACE cannot be computed yet, because of the censored information. But if we observed a CACE of 1, it means that, on average, compliers assigned to the action group ($Z=1$) have a retention time that is 1 month longer than compliers assigned to the control group. Similar problem as in the ITT, we cannot compute the CACE as we are ignoring the censored observations. The next chapter addresses this limitation by introducing the RMST framework, which properly accounts for censored information.




Figure 3.6: KM Curves — Compliers vs Non-Compliers vs Control

As discussed in Section 3.2, the ITT and CACE estimates are already causally valid thanks to the randomization of Z . However, the covariates X also influence both the probability of using the treatment $E[D|X]$ and the outcome $(E[Y|X])$, as illustrated in the DAG, Figure 3.7. With the help of the covariate adjustment in the next chapter, we will not be able to correct for bias, but we will reduce the residual variance of the estimates, and it will lead to more precise inference.




Figure 3.7: DAGs Data Euroconsumers

The goal in the next chapter will be to capture this effect of the covariates on the probability of using the treatment $(E[D|X])$ and on the outcome $(E[Y|X])$.

Chapter 4

Machine Learning Approach

The ITT and CACE estimators presented in the previous chapter are intuitive, but they suffer from two limitations. The first problem is that the outcome variable Y_i is censored for a subset of individuals, meaning that their exact survival time is not fully observed. In the previous section, we did not make the distinction between a result of $Y = 13$ with *censored* = 1 and $Y = 13$ with *censored* = 0 which is problematic because it does not give the same information. Based on Andersen and Perme (2004), the pseudo-observation technique can solve this problem and addresses this first limitation by transforming censored survival times into uncensored quantities $\hat{\theta}_i$ that can be used as outcomes in any standard regression model, at multiple time horizons τ .

The second limitation is, as presented in the previous section, that the covariates X_i simultaneously influence both the probability of using the treatment D_i and the outcome Y_i . This problem of confounding prevents a direct OLS interpretation of β_1 in the structural model:

$$E(Y \mid D = d, X = x) = \beta_0 + \beta_1 d + \beta_2 x \quad (4.1)$$

The Double Machine Learning framework, introduced by Chernozhukov et al. (2018), solves this second limitation by explicitly split out the effect of X_i on both D_i and $\hat{\theta}_i$ before estimating $\hat{\tau}$.

Therefore, the goal of this chapter is to introduce the pseudo-observation technique to resolve the censoring problem and to apply the Double Machine Learning (DML) framework to estimate the causal effect of the treatment knowing the covariates. The main resources used are Andersen and Perme (2004), Chernozhukov et al. (2018), and the CausalML and EconML frameworks.

4.1 Why Double Machine Learning ?

Before introducing the Pseudo-observation approach and the Double Machine Learning framework, it is useful to recall that it exists multiple framework for causal estimation, each with their own interpretation and limitations. The goal of the section is to motivate the choice of the Double Machine Learning method by introducing two alternatives: G-Computation and Inverse Probability Weighting.

4.1.1 G-Computation

The G-computation method (Hernán and Robins, 2020) is an outcome-based approach to covariate adjustment. The idea can be model in two different steps. The first step consists of modelling the outcome directly as a function of both the treatment and the covariates:

$$\hat{\mu}(d, X_i) = E[Y_i \mid D_i = d, X_i]$$

$\hat{\mu}(d, X_i)$ can in principle be estimated by any machine learning algorithm, including non-parametric methods. The average treatment effect is then recovered in the second step by standardisation over the covariate distribution:

$$\hat{\tau}_G = \frac{1}{n} \sum_{i=1}^n [\hat{\mu}(1, X_i) - \hat{\mu}(0, X_i)]$$

By fixing the treatment status and compute the treated and untreated potential outcomes using the fitted model, the model gives the estimation of the average causal effect across the population. However, as noted by Hernán and Robins (2020), G-Computation does not naturally propagate the uncertainty arising from the estimation of $\hat{\mu}$, which can lead to overconfident confidence intervals.

4.1.2 Inverse Probability Weighting

An other method, also discussed in Hernán and Robins (2020), is Inverse Probability Weighting (IPW). In this method, rather than modelling the outcome as in the G-Computation method, the goal is to reweight each observed outcome by the inverse of the individual's probability of receiving the treatment, known as the propensity score $e(X_i) = P(D_i = 1 \mid X_i)$. The final estimation of the effect is defined as:

$$\hat{\tau}_{IPW} = \frac{1}{n} \sum_{i=1}^n \left[\frac{D_i Y_i}{\hat{e}(X_i)} - \frac{(1 - D_i) Y_i}{1 - \hat{e}(X_i)} \right]$$

The intuition is that by putting more weight on individuals who received a treatment that was unlikely given their covariates, and putting less weight on those for who it was likely, and the IPW reconstructs a pseudo-population in which D_i is independent of X_i , restoring a valid comparison between the two

groups.

A known limitation from this model is its sensitivity to extreme propensity scores. In fact, the estimation becomes unstable when the propensity scores is close to 0 or 1. In our data, the global compliance rate is 12%, meaning that some individuals in the action group have a low estimated probability of using the treatment given their covariates. It makes the IPW estimator not perfectly suited to our setting.

4.1.3 Double Machine Learning

Where the G-computation fits a model to explain the outcome Y and IPW fits a model to predict the probability of treatment uptake D , the Double Machine Learning method, introduced by Chernozhukov et al. (2018), does both simultaneously. It corrects the influence of the covariates X_i through two mechanisms: the partialling-out and the cross-fitting procedure. Under these conditions, Chernozhukov et al. (2018) establish that the resulting estimator is $\sqrt{n}$ -consistent and asymptotically normal, even when flexible machine learning models are used for the nuisance functions. The Double Machine Learning method is explained in details in Section 4.4.

For these reasons, the Double Machine Learning framework is the most appropriate choice in our setting. It improves the robustness of the G-Computation and the potential instability of IPW under low compliance, while providing valid inference without requiring parametric assumptions on the functional form of the nuisance functions.

4.2 Pseudo-Observations

The pseudo-observation approach, based on Andersen and Perme (2004), solves the problem of censored information by transforming the censored survival times into a set of uncensored quantities that can be used in any standard regression model.

4.2.1 Construction of Pseudo-Observations

The key challenge in survival analysis is that the outcome Y_i cannot be fully observed for a subset of individuals. In fact, because of the administrative censoring problem, some of the clients are still retained at the end of the observation window. To help us in the formalisation, we first have to define the two core quantities of survival analysis.

The first quantity is the survival function $S(t)$ that represents the probability of an individual to survive beyond time t . Formally,

$$S(t) = P(Y > t), \quad t \geq 0 \quad (4.2)$$

By construction, $S(0) = P(Y > 0) = 1$ meaning that all clients are retained at time 0 and $S(t)$ is a non-increasing function of t . In our setting, $S(t)$ represents the probability that a client is still retained at month t .

The second quantity is the Restricted Mean Survival Time (RMST) representing the expected survival time truncated at a fixed horizon τ . Formally,

$$\mu(\tau) = E[\min(Y_i, \tau)] = E[Y_i \wedge \tau] = \int_0^\tau S(t) dt \quad (4.3)$$

Therefore, in our data, $E[Y_i \wedge \tau]$ represents the expected number of months a client remains subscribed, truncated at horizon τ . It means that $\min(Y_i, \tau)$ equals Y_i if the client churned before τ , and equals τ if the client is still retained at the end of the observation window. The equality $E[Y_i \wedge \tau] = \int_0^\tau S(t) dt$ follows from the standard result for non-negative random variables. Starting from the definition:

$$\begin{aligned} E[Y_i \wedge \tau] &= \int_0^\tau P(Y_i \wedge \tau > t) dt \\ &= \int_0^\tau P(\min(Y_i, \tau) > t) dt \\ &= \int_0^\tau P(Y_i > t) dt \\ &= \int_0^\tau S(t) dt \end{aligned} \quad (4.4)$$

Where the third equality holds because for $t \in [0, \tau]$, the event $\{\min(Y_i, \tau) > t\}$ is equivalent to $\{Y_i > t\}$, since $\tau > t$ is always satisfied on this interval.

In our data, $S(t)$ and $\mu(\tau)$ are unknown and must be estimated from the data $(\hat{S}(t), \hat{\mu}(\tau))$. $\hat{S}(t)$ has already been defined in Chapter 2 as:

$$\hat{S}(t) = \prod_{j: \tilde{Y}_{(j)} \leq t} \left(1 - \frac{d_{(j)}}{r_{(j)}}\right) \quad (4.5)$$

Where at each observed churn time, the estimator is updating the survival probability by multiplying the current estimate by the fraction of individuals who did not churn yet. Censored individuals contribute to the risk set $r_{(j)}$ up to their censoring time, after which they are removed from the at-risk pool.

We can illustrate the situation by considering the five observations in Table 4.1.

Table 4.1: Sample of randomly selected observations from the data

Age_group	Age missing	Wave	Seniority	Web visit group	Gender_M	Treatment assignment	Treatment category	Accepted Treatment	TimeToEvent Surv	Censored
[45,59]	1	1	13.0	1	0	1	30E_voucher	0	13	1
[45,59]	1	1	12.0	0	1	1	30E_voucher	1	3	0
[45,59]	0	1	13.0	0	0	1	30E_voucher	0	5	0
[45,59]	0	2	13.0	0	1	0	—	0	9	1
[60,110]	0	1	14.0	1	0	1	30E_voucher	0	7	0

In this table, two individuals have in the columns Censored: Censored = 1 (still retained at the end of the survey) and the three others have Censored = 0 (churned). If we use the naive approach by treating the five observations as if they churned at their observation time, it would underestimate the true retention time. The solution is to use the pseudo-observation approach that avoids this by reweighting each individual's contribution to the RMST. In fact, we can obtain the RMST by integrating the KM curve as follows:

$$\hat{\mu}(\tau) = \int_0^\tau \hat{S}(t) dt \quad (4.6)$$

The pseudo-observations are then obtained for individual i at time horizon τ by using a jackknife procedure:

$$\hat{\theta}_i = n \cdot \hat{\mu}(\tau) - (n-1) \cdot \hat{\mu}^{-i}(\tau) \quad (4.7)$$

where $\hat{\mu}^{-i}(\tau)$ is the RMST estimated on the observations without individual i , and n represents the number of observations. Intuitively, $\hat{\theta}_i$ measures the influence of individual i on the overall RMST.

For example, if we consider the five observations from Table 4.1:

- The first individual with $Y = 13$, Censored = 1 is still retained at the end of the observation window. Removing this individual eliminates a long-surviving observation from the risk set $r_{(j)}$ at every time point, which decreases $\hat{S}(t)$ and therefore decreases $\hat{\mu}^{-i}(\tau)$. This produces a large pseudo-observation $\hat{\theta}_i$, reflecting a long retention time.
- Conversely, the second individual with $Y = 3$, Censored = 0 churned very early. Removing this individual eliminates an early churn event, which increases $\hat{S}(t)$ at $t = 3$ and therefore increases $\hat{\mu}^{-i}(\tau)$. This produces a small pseudo-observation $\hat{\theta}_i$, reflecting a short retention time.
- The individuals with $Y \in \{5, 7\}$, Censored = 0 churned at intermediate times and have a moderate impact on $\hat{\mu}^{-i}(\tau)$, yielding intermediate pseudo-observations.
- The individual with $Y = 9$, Censored = 1 is still retained at month 9. Similarly to the first individual, removing them decreases $\hat{\mu}^{-i}(\tau)$, producing a relatively large pseudo-observation, though smaller than for $Y = 13$ as their censoring occurs earlier.

Table 4.2: Pseudo-observations at $\tau \in \{6, 9, 13\}$ months for the sample observations

Age_group	Wave	Treatment assignment	Accepted Treatment	TimeToEvent Surv	Censored	$\hat{\theta}_i(\tau = 6)$	$\hat{\theta}_i(\tau = 9)$	$\hat{\theta}_i(\tau = 13)$
[45,59]	1	1	0	13	1	6.0	9.0	13.0
[45,59]	1	1	1	3	0	2.5	2.5	2.5
[45,59]	1	1	0	5	0	4.5	4.5	4.5
[45,59]	2	0	0	9	1	6.0	9.0	—
[60,110]	1	1	0	7	0	6.0	6.5	6.5

— indicates that $\hat{\theta}_i(\tau = 13)$ is not defined for Wave 2 as the observation window is limited to 9 months.

In Table 4.2, we see the resulting pseudo-observations at $\tau \in \{6, 9, 13\}$ month for the five observations of Table 4.1. Table 4.2 illustrates how censored observations are handled through reweighting rather than being discarded or naively imputed. The pseudo-observations are computed separately for each wave to account for the difference in observation windows. For Wave 1, we compute pseudo-observations at $\tau \in \{6, 9, 13\}$ months as we have information until time 13. For Wave 2, the observation window is limited to 9 months and pseudo-observations are therefore computed at $\tau \in \{6, 9\}$. Thus, for the individuals in Wave 2, the pseudo-observations at $\tau = 13$ are undefined.

The fundamental property of pseudo-observations, established by Andersen and Perme (2004), is that under the assumption of independent censoring:

$$E(\hat{\theta}_i \mid X_i) \approx E(Y_i \wedge \tau \mid X_i) \quad (4.8)$$

where $Y_i \wedge \tau = \min(Y_i, \tau)$. This means that $\hat{\theta}_i$ is an approximately unbiased estimator of the individual's expected restricted survival time, conditionally on their covariates. As a result, pseudo-observations can be used as a substitute for the unobserved uncensored outcome in any standard regression or machine learning model, without requiring parametric assumptions on the survival distribution.

4.3 ITT and CACE with censored information

Having resolved the censoring problem through pseudo-observations, we can compute the true ITT and CACE while properly accounting for censored observations.

The ITT, as we have seen, is defined as:

$$\text{ITT} = E[Y \mid Z = 1] - E[Y \mid Z = 0] \quad (4.9)$$

And the problem is that, in our case, Y_i is not fully observed for censored individuals, so we cannot directly compute $E[Y \mid Z = z]$. However, by the fundamental property of pseudo-observations established in Equation 4.8:

$$E(\hat{\theta}_i \mid X_i) \approx E(Y_i \wedge \tau \mid X_i) \quad (4.10)$$

So now, if we take expectations over X_i and conditioning on $Z = z$, this implies that the group-level RMST can be consistently estimated by the sample mean of pseudo-observations within each group:

$$E[Y_i \wedge \tau \mid Z = z] \approx \frac{1}{n_z} \sum_{i: Z_i = z} \hat{\theta}_i = \hat{\mu}_z(\tau) \quad (4.11)$$

We can write the pseudo-ITT as:

$$\widehat{\text{ITT}}_{RMST}(\tau) = \hat{\mu}_1(\tau) - \hat{\mu}_0(\tau) = \int_0^\tau \hat{S}_1(t) dt - \int_0^\tau \hat{S}_0(t) dt \quad (4.12)$$

where $\hat{S}_1(t)$ and $\hat{S}_0(t)$ are the probabilities of surviving on time t for the action and control groups respectively. We evaluate $\widehat{\text{ITT}}_{RMST}(\tau)$ at $\tau \in \{6, 9, 13\}$ months to capture the short-term, medium-term, and long-term effect of Z on retention.

Similarly, the pseudo-CACE is defined as:

$$\widehat{\text{CACE}}_{RMST}(\tau) = \frac{\widehat{\text{ITT}}_{RMST}(\tau)}{\hat{\pi}_c} = \frac{\hat{\mu}_1(\tau) - \hat{\mu}_0(\tau)}{E[D \mid Z = 1] - E[D \mid Z = 0]} \quad (4.13)$$

where $E[D \mid Z = 0] = 0$ as there is no contamination in the Euroconsumers data.

Confidence Intervals. We can compute the 95% confidence interval for the $\widehat{\text{ITT}}_{RMST}$ using the sandwich variance estimator. In fact, since $\widehat{\text{ITT}}_{RMST}(\tau)$ is a difference of two independent sample means of pseudo-observations, its variance is:

$$\widehat{\text{Var}}\left(\widehat{\text{ITT}}_{RMST}\right) = \frac{\hat{\sigma}_1^2}{n_1} + \frac{\hat{\sigma}_0^2}{n_0} \quad (4.14)$$

Where $\hat{\sigma}_1^2$ and $\hat{\sigma}_0^2$ represent the variance in the action and control group, n_1 and n_0 represent the number of observation in the action and control group. With the help of the Central Limit Theorem, since the pseudo-observations $\hat{\theta}_i$ are independent across individuals and their variance is finite, the sample mean within each group converges in distribution to a Gaussian:

$$\sqrt{n_z}(\hat{\mu}_z(\tau) - \mu_z(\tau)) \xrightarrow{d} \mathcal{N}(0, \sigma_z^2) \quad (4.15)$$

For the $\widehat{\text{CACE}}_{RMST}$, since it is a ratio of two estimators, the variance is obtained by the Delta Method:

$$\widehat{\text{Var}}\left(\widehat{\text{CACE}}_{RMST}\right) = \frac{1}{\hat{\pi}_c^2} \left(\widehat{\text{Var}}\left(\widehat{\text{ITT}}_{RMST}\right) + \widehat{\text{CACE}}_{RMST}^2 \cdot \frac{\hat{\pi}_c(1 - \hat{\pi}_c)}{n_1} \right) \quad (4.16)$$

where the second term captures the uncertainty in the estimated compliance rate $\hat{\pi}_c$, modelled as a Bernoulli proportion. The covariance between $\widehat{\text{ITT}}_{RMST}$ and $\hat{\pi}_c$ is zero by randomisation, since Z_i is independent of all individual characteristics.

Table 4.3: Summary of $\widehat{\text{ITT}}_{RMST}(\tau)$ and $\widehat{\text{CACE}}_{RMST}(\tau)$ Estimates at different Horizons (τ)

Horizon / Group	$\widehat{\text{ITT}}_{RMST}$ (months)	95% CI	$\widehat{\text{CACE}}_{RMST}$ (months)	95% CI
$\tau = 6$ months — Short-Term				
Voucher ($Z = v$)	0.0488	[−0.0130, 0.1106]	0.3345	[−0.0897, 0.7587]
Discount ($Z = d$)	0.0539	[−0.0080, 0.1157]	0.5712	[−0.0863, 1.2287]
$\tau = 9$ months — Medium-Term				
Voucher ($Z = v$)	0.1153	[0.0050, 0.2256]	0.7906	[0.0328, 1.5485]
Discount ($Z = d$)	0.1307	[0.0204, 0.2410]	1.3860	[0.2117, 2.5603]
$\tau = 13$ months — Long-Term (Wave 1 only)				
Voucher ($Z = v$)	0.1832	[−0.0717, 0.4381]	1.1900	[−0.4680, 2.8480]
Discount ($Z = d$)	0.1830	[−0.0719, 0.4379]	1.8121	[−0.7187, 4.3429]

Table 4.3 summarises the $\widehat{\text{ITT}}_{RMST}$ and $\widehat{\text{CACE}}_{RMST}(\tau)$ in the three time horizons. A clear pattern emerges: the $\widehat{\text{ITT}}_{RMST}$ effect grows monotonically with the time horizon, suggesting that the effect of being assigned to the action group accumulates over time rather than being concentrated in the short term. The uncertainty also increases with the time horizon, producing wider confidence intervals. The only statistically significant estimates are found at $\tau = 9$ months, for both the $\widehat{\text{ITT}}_{RMST}$ and the $\widehat{\text{CACE}}_{RMST}$. It is in fact the larger τ with all the observations and it tends to reduce the variance as seen in Equation 4.15. This instability can also be explained by a small compliance rate that makes the variance of the CACE increased as $\hat{\pi}_c^2$ is small in the two treatments.

As discussed in Section 3.2, the results in Table 4.3 are already causally valid thanks to the randomization of Z . However, the covariates X_i also influence both the probability of using the treatment and the outcome. Adding covariate adjustment through the Double Machine Learning framework will not correct for bias but will reduce the residual variance of the estimates, leading to more precise inference. However, as discussed in Section 4.4.4, this comes at the cost of the conditional unconfoundedness assumption.

4.4 Double Machine Learning

In this section, we address the second limitation: the covariates X_i explain part of the variation in both the treatment receipt D_i and the outcome $\hat{\theta}_i$. The Double Machine Learning framework, introduced by Chernozhukov et al. (2018), explicitly removes this variation, reducing the residual variance and yielding more precise estimates of the causal effect. It is important to note that this improvement in the variance relies on the conditional unconfoundedness assumption, developed in subsection 4.4.4.

4.4.1 Step 1 : Nuisance Functions

The first step consists of estimating two nuisance functions that capture the effect of the covariates X_i on the outcome and on the treatment respectively.

We call them Outcome Regressions and Propensity Score:

$$\hat{m}_0(X_i) = E(\hat{\theta}_i | X_i) \quad (\text{Outcome Regressions})$$

$$\hat{p}(X_i) = P(D_i = 1 | X_i, Z_i = 1) \quad (\text{Propensity score})$$

where $\hat{\theta}_i$ represents the RMST pseudo-observation for individual i . The goal of the Outcome Regression ($\hat{m}_0(X_i)$) is to use any suitable machine learning technique to capture the variation in the pseudo-observed retention time $\hat{\theta}_i$ that is explained by the covariates alone. It is estimated on all individuals who did not use the treatment ($D_i = 0$) by including the control group ($Z = 0$) and the never-takers ($Z = 1, D = 0$), in order to avoid that the treatment effect contaminate the estimation of the baseline outcome.

The propensity score $\hat{p}(X_i)$ estimates the probability of using the treatment given the observed characteristics of individual i , the propensity score could also be estimated by any suitable machine learning model. In our case, we use a Linear Regression for $\hat{m}_0$ and a Logistic Regression for $\hat{p}$. Other models, such as Random Forest or Gradient Boosting, were also tested but did not lead to significantly different estimates of $\hat{\tau}$. Linear and Logistic Regression were therefore retained for their interpretability, as their coefficients can be directly read and discussed in Section 4.4.5.

To avoid overfitting, the DML framework uses a cross-fitting procedure (Chernozhukov et al., 2018). Meaning that the data are split into K folds. For each fold k , $\hat{m}_0$ and $\hat{p}$ are estimated on the remaining $K - 1$ folds, and the residuals $\tilde{Y}_i$ and $\tilde{D}_i$ are computed on fold k .

4.4.2 Step 2 : Partialling Out and Final Estimation

The second step consists in removing the variation explained by the covariates from the outcome and the treatment, by computing the residuals:

$$\tilde{Y}_i = \hat{\theta}_i - \hat{m}_0(X_i) \quad (4.17)$$

$$\tilde{D}_i = D_i - \hat{p}(X_i) \quad (4.18)$$

We see in equation 4.17 and equation 4.18 that the goal is to have the resulting residuals with $\tilde{Y}_i$ and $\tilde{D}_i$ that represent the outcome not explained by the covariates and the probability of being treated not explained by the covariates

respectively. By construction, both residuals are orthogonal to X_i , which eliminates the confounding effect of the covariates.

The causal effect $\hat{\tau}$ can now be estimated via :

$$\tilde{Y}_i = \hat{\tau} \cdot \tilde{D}_i + \varepsilon_i \quad (4.19)$$

Since both residuals are orthogonal to X_i , the OLS estimator of $\hat{\tau}$ in Equation 4.19 is consistent and asymptotically unbiased under the regularity conditions of Chernozhukov et al. (2018).

$$\hat{\tau} = \frac{\sum_{i=1}^n \tilde{D}_i \tilde{Y}_i}{\sum_{i=1}^n \tilde{D}_i^2} \quad (4.20)$$

The estimator represents the average effect of using the treatment on the restricted survival time $\hat{\theta}_i$, adjusted for the confounding influence of the covariates X_i .

To quantify the uncertainty around the estimation we can build 95% confidence interval. In fact, under the regularity conditions of Chernozhukov et al. (2018), $\hat{\tau}$ is asymptotically normal:

$$\sqrt{n}(\hat{\tau} - \tau) \xrightarrow{d} \mathcal{N}(0, V) \quad (4.21)$$

where V is estimated by the heteroscedasticity-robust sandwich variance estimator, which remains valid without any parametric assumption on the distribution of ε_i . It means that this asymptotic result does not require the data to follow a normal distribution and it relies on the large sample size and the independence of observations, allowing the construction of valid 95% confidence intervals without any parametric assumption on the distribution of Y or D .

4.4.3 Results on the Euroconsumers Data

The survival time and censoring columns are first transformed into pseudo-observations $\hat{\theta}_i$ in three different horizons $\tau \in \{6, 9, 13\}$ months, corresponding to short-term, medium-term, and long-term retention effects. As discussed in Section 4.2, pseudo-observations at $\tau = 13$ months are computed on Wave 1 only, as Wave 2 covers only 9 months of follow-up.

We use a Linear Regression for the outcome regression $\hat{m}_0$ and a Logistic Regression for the propensity score $\hat{p}$, with $K = 5$ cross-fitting folds. We first estimate the effect of the treatment without making distinction between the voucher and the discount ($D_i = 1$). By combining the two treatments into a single binary variables, it violates the SUTVA assumptions, thus, the results are biased, as discussed in Section 3.5. Then, we separate the data by treatment type to restore the SUTVA assumption and the result are not biased.

Overall Treatment Effect

Table 4.4: DML Estimates of the Overall Treatment Effect $\hat{\tau}$ at Three Horizons

Horizon	$\hat{\tau}$ (months)	SE	CI 95%	p-value
$\tau = 6$	0.5117	0.0011	[0.5095, 0.5139]	< 0.001
$\tau = 9$	0.9796	0.0019	[0.9759, 0.9833]	< 0.001
$\tau = 13$	1.6289	0.0038	[1.6215, 1.6363]	< 0.001

$\tau = 13$ restricted to Wave 1. $p < 0.001$.

In a first time, we identify the effect of the treatment on the retained without making distinction between the two treatment and violates the SUTVA assumption. And then, we identify the effect of the two treatment on the retained time. Table 4.4 summarises the Double Machine Learning estimations of the causal effect of treatment uptake on restricted survival time at the three horizons. The estimation grow monotonically with the time horizon. And all estimates are highly statistically significant ($p < 0.001$) with narrow confidence intervals, reflecting the precision gained from the large sample size and the covariate adjustment.

Treatment Effect by Type

Table 4.5: DML Estimates by Treatment Type at Three Horizons

Horizon	Treatment	$\hat{\tau}$ (months)	SE	95% CI	p-value
$\tau = 6$	Voucher ($Z = v$)	0.4148	0.0012	[0.4124, 0.4172]	< 0.001
	Discount ($Z = d$)	0.6951	0.0021	[0.6909, 0.6993]	< 0.001
$\tau = 9$	Voucher ($Z = v$)	0.7341	0.0021	[0.7300, 0.7381]	< 0.001
	Discount ($Z = d$)	1.4534	0.0040	[1.4457, 1.4612]	< 0.001
$\tau = 13$	Voucher ($Z = v$)	1.2659	0.0042	[1.2578, 1.2740]	< 0.001
	Discount ($Z = d$)	2.3822	0.0109	[2.3608, 2.4035]	< 0.001

$\tau = 13$ restricted to Wave 1. *** $p < 0.001$.

Table 4.5 shows the estimations of the causal effect across the three horizons. Similarly, of the previous identification of the ITT, the effects of the voucher and the discount are not similar in all the different horizons it gradually increasing, the use of the discount suggests a bigger impact on the retention time of clients.

4.4.4 Unconfoundedness Assumption

It is important to note that the Double Machine Learning framework presented in this section relies on the conditional unconfoundedness assumption:

$$Y_i(d) \perp D_i \mid X_i \quad (4.22)$$

The assumption means that the observed covariates X_i (Age Group, Seniority, WebVisit Group, Gender) are sufficient to fully account for all confounding between treatment receipt D_i and the outcome Y_i . Intuitively, it means that conditional on X_i , there are no unobserved characteristics $U_i \setminus X_i$ that simultaneously influence the probability of using the incentive D_i and the retention time Y_i . This assumption cannot be tested from the data and is stronger than the IV framework developed in Chapter 3, which explicitly corrected for unobserved confounding by using the randomized assignment Z_i as an instrument.

If the conditional unconfoundedness assumption is violated, $\hat{\tau}$ cannot be given a causal interpretation, but the nuisance function coefficients retain their interpretability of the coefficients. In fact, the outcome regression $\hat{m}_0$ describes how the observed covariates X_i are associated with retention time, and the propensity score $\hat{p}(X_i)$ describes which client profiles are more likely to use the incentive.

An extension that does not require this assumption is the Double Machine Learning IV (DMLIV) framework, also introduced by Chernozhukov et al. (2018). The main difference is in the final estimation step. In the classical DML, the causal effect is estimated by regressing the outcome residuals $\tilde{Y}_i$ on the treatment residuals $\tilde{D}_i$:

$$\hat{\tau}_{DML} = \frac{\sum_{i=1}^n \tilde{D}_i \tilde{Y}_i}{\sum_{i=1}^n \tilde{D}_i^2} \quad (4.23)$$

In the DMLIV, the instrument Z_i is additionally residualized with respect to X_i to obtain $\tilde{Z}_i$, and the final estimator becomes:

$$\hat{\tau}_{DMLIV} = \frac{\sum_{i=1}^n \tilde{Z}_i \tilde{Y}_i}{\sum_{i=1}^n \tilde{Z}_i \tilde{D}_i} \quad (4.24)$$

By using $\tilde{Z}_i$ in place of $\tilde{D}_i$ in the numerator, the DMLIV exploits the exogenous variation in D_i induced by the instrument Z_i , correcting for the endogeneity of D_i .

In our situation with the data from Euroconsumers, the conditional unconfoundedness is difficult to fully justify given the limited number of observed covariates available. For example, a client facing financial difficulties may be simultaneously more likely to use a monetary incentive and more likely to cancel their subscription, a relation that the covariates X do not capture. Similarly, some client may have subscribed to Euroconsumers for specific reasons, time-limited purpose, which would also affect their retention independently of their observed characteristics. The goal of these two example is to illustrate that a lot of possible unobserved confounding is plausible in this setting, and that the

DML estimations should therefore be interpreted with caution. The goal of this model is to be seen as a complement of the IV-based estimation in Chapter 3. In fact, the advantage of the DML approach in this context is the interpretability it provides through the coefficients of outcome regression and the propensity score as discussed in next Section.

4.4.5 Nuisance Functions : Interpretability

With the help of our two linear models the estimated coefficients are directly interpretable for the Outcome Regression ($\hat{m}_0$) and the Propensity Score $\hat{p}$:

Table 4.6: Coefficients of the Nuisance Functions

Covariate	Outcome Regression ($\hat{m}_0$)	Propensity Score (OR) $\hat{p}$
Intercept	10.504	—
Seniority	− 0.201	0.871
WebVisit_Group	−0.033	2.251
Gender_M	+ 0.365	0.956
Wave	0.000	0.573
Age_missing	− 1.143	0.623
Age_[30,44]	+ 1.280	1.233
Age_[45,59]	+ 1.621	1.116
Age_[60,110]	+ 1.940	1.407

Bold indicates significance at the 1% level ($p < 0.01$).

Table 4.6 shows the coefficients of the two nuisance functions, estimated on Wave 1 at $\tau = 13$. Bold values indicate significance at the 1% level. In the outcome regressions $\hat{m}_0$, we can identify that the variables Age have a large coefficient and a positive effect on retention than when Age is between 18 and 29 and an Age missing has a negative impact on the prediction of the retention. In addition, the seniority has a small but significant negative effect (−0.201 months per additional month of Seniority). Being a male has also a significant positive impact, while WebVisit_Group and Wave are not significant.

4.4.6 Financial Analysis

The goal of this subsection is to quantify the financial impact of the campaign on Euroconsumers. The analysis is separated by in two fort the treatments and also separated by the two sets of assumptions, corresponding to the two causal frameworks developed. We assume that each client generates a revenue of €20 per month for Euroconsumers.

Financial Analysis under the IV Framework

The first analysis uses the $\widehat{CACE}_{RMST}(\tau)$ estimates from Section ??, which are causally valid under the IV assumptions verified in Chapter 3 and it also captures the effect of the observed and unobserved characteristics (U). In this first case, since the only significant estimates are found at $\tau = 9$ months, the financial analysis is conducted at this horizon.

Voucher The cost of a voucher is assumed to be €30. Although Euroconsumers may obtain vouchers at a lower cost through partnerships, we retain the face value as a conservative assumption. The $\widehat{CACE}_{RMST}(9)$ for the voucher is 0.791 months, and there are 984 compliers.

$$\begin{aligned}\text{Cost} &= 30\text{€} \times 984 = 29,520\text{€} \\ \text{Revenue} &= 20\text{€} \times 0.791 \times 984 = 15,559\text{€} \\ \text{Net} &= 15,559\text{€} - 29,520\text{€} = -13,961\text{€}\end{aligned}$$

Discount. The discount corresponds to a 50% reduction on the monthly subscription for six month. Since the monthly revenue per client is €20, the cost per complier is €10. The $\widehat{CACE}_{RMST}(9)$ for the discount is 1.386 months, and there are 636 compliers.

$$\begin{aligned}\text{Cost} &= 10\text{€} \times 636 \times 6 = 38,160\text{€} \\ \text{Revenue} &= 20\text{€} \times 1.386 \times 636 = 17,630\text{€} \\ \text{Net} &= 17,630\text{€} - 38,160\text{€} = -20,530\text{€}\end{aligned}$$

To account for the uncertainty in the $\widehat{CACE}_{RMST}(9)$ estimates, we also propagate the 95% confidence intervals into the financial results:

Table 4.7: Financial Analysis under the IV Framework ($\tau = 9$ months)

Treatment	Cost (€)	Revenue (€)	Net (€)	Net 95% CI (€)
Voucher	29,520	15,559	-13,961	[-28,874, +954]
Discount	38,160	17,630	-20,530	[-35,467, -5,593]

The final results for the voucher campaign, under the IV framework, yields a result that oscillates between -28,874€ and +954€ while the discount campaign yields between -35,467€ and -5,593€. The voucher treatment suggests a possible profitability where the discount treatment suggests significant loss.

Financial Analysis under the DML Framework

The second analysis uses the DML estimates from Section 4.4, which rely on the conditional unconfoundedness assumption (Discussed in Section 4.4.4). The estimates are statistically significant at all three horizons that is why we conduct

the analysis at $\tau = 9$ months for comparability with the IV framework, and at $\tau = 13$ months for a longer-term perspective.

By using the same costs as for the IV framework and by using the confidence interval find in Table 4.5, we obtain the results:

Table 4.8: Financial Analysis under the DML Framework

Horizon	Treatment	Cost (€)	Revenue 95% CI (€)	Net 95% CI (€)
$\tau = 9$	Voucher	29,520	[14,366, 14,526]	[-15,154, -14,994]
	Discount	38,160	[18,395, 18,586]	[-19,771, -19,574]
$\tau = 13$	Voucher	29,520	[24,754, 25,072]	[-4,766, -4,448]
	Discount	38,160	[30,029, 30,573]	[-8,131, -7,587]

Revenue CI = $\text{€}20 \times [lb_{\tau}, ub_{\tau}] \times n_C$. Net CI = Revenue CI - Cost. $\tau = 13$ restricted to Wave 1.

In the both framework, the two campaigns are significantly not profitable, because of their high cost relatively to their positive retention effect it generates. The DML estimation at $\tau = 13$ months suggest a more optimistic picture for the discount and the voucher as the effect of the incentive is increasing through the time but it keeps being not profitable. The results should by interpreted with caution knowing the conditional unconfoundedness assumption is considered true.

In the IV framework and in the DML, the financial analysis suggests that neither campaign is profitable. However, the voucher appears to yield smaller net loss than the discount, thus, replacing the discounts by the vouchers would reduce financial losses.

Chapter 5

Meta Learners

The goal of this chapter is to identify the most responsive clients to the campaign and to see how this information can improve the profitability of future campaigns. The goal of the DML and the IV framework were to identify the effect of using the incentive on the retention of the client by providing a single average causal effect $\hat{\tau}$. The DML helps to identify which characteristics could influence the outcome retention and the propensity score but it does not identify the profiles that respond the most strongly to the treatment. If Euroconsumers wants to target only the most responsive clients to maximise his profitability rather than sending incentives to all clients. To adress this approach, we introduce the Meta-Learner framework, which estimates heterogeneous treatment effect at the individual level by using machine learning techniques.

Change of notation In this new framework, the treatment variable D_i is abandon and we use directly the randomized assignment Z_i directly. Therefore the notation is redefined through the chapter as:

- $Y_i(1)$ denotes the potential retention time of individual i if assigned to the action group ($Z_i = 1$)
- $Y_i(0)$ denotes the potential retention time of individual i if assigned to the control group ($Z_i = 0$)
- The treatment effect is the effect of the email, not of using the incentive

As discussed in Chapter 3, since Z_i is randomly assigned, $Z_i \perp U_i$ holds by design and we don't have endogeneity issue and conditional unconfoundedness assumption is automatically satisfied. The estimated heterogeneous effect therefore has a direct causal interpretation.

The Meta-Learner regroups different models such as the S-Learner, the T-Learner and the X-Learner that we will defined in this thesis. The theoretical framework used in this section follow the documentation of the CausalML

(Chen et al., 2020) and EconML (Battocchi et al., 2019) packages, which provide open-source Python implementations of all three methods. The X-Learner is the most robust method of the three meta learners and it is therefore the only one fitted on the Euroconsumers data. Before defining the three meta-learners, we introduce a new metric called the Individual Treatment Effect (ITE).

5.1 Individual Treatment Effect

A new metric we have to introduce is called the Individual Treatment Effect (ITE) for individual i and is defined as:

$$ITE_i = Y(1)_i - Y(0)_i \quad (5.1)$$

Where $Y(1)_i$ represents the potential outcome in the action group for individual i and $Y(0)_i$ the potential outcome in the control group for individual i . As shown in Table 5.1, it is impossible to identify the ITE_i by direct observations: for each individual, one of the two potential outcomes is observable and the other is not (counterfactual).

Table 5.1: Sample of randomly selected observations — Fundamental problem of causal inference

Age.group	Wave	Seniority	Web visit group	Treatment assignment	Treatment category	Accepted voucher	TimeToEvent Surv	Censored	Y(1)	Y(0)	Y(1)-Y(0) ITE
[60,110]	1	15.0	2	0	—	0	13	1	?	13	?
[45,59]	1	13.0	2	1	50_discount	0	5	0	5	?	?
[30,44]	1	13.0	2	1	30E_voucher	1	14	0	14	?	?
[60,110]	2	14.0	2	1	30E_voucher	0	7	0	7	?	?
[30,44]	2	13.0	2	0	—	0	9	1	?	9	?

What can be estimated instead is the Conditional Average Treatment Effect (CATE):

$$CATE_i = E[Y_i(1) - Y_i(0) \mid X = X_i] \quad (5.2)$$

It represents the average treatment effect for individuals that share the same profile X_i , and it is the best estimable approximation of the individual-level effect.

5.2 S-Learner

The S-Learner is the simplest meta-learner as it fits a single model $M(X, Z)$ on the entire dataset, using both the covariates X_i and the treatment indicator Z_i as features.

$$M(X, Z) = E[\hat{\theta}_i \mid X_i, Z_i] \quad (5.3)$$

We can compute the CATE by prediction the outcome twice for each individual by fixing the assignment $Z = 1$ and $Z = 0$, and by taking the difference:

$$\widehat{CATE}_i^S = M(X_i, 1) - M(X_i, 0) \quad (5.4)$$

For the three meta learners any regression algorithm can be used for the outcome models M . However, a known limitation of the S-Learner is that some machine learning models apply to avoid overfitting, which can shrink the coefficient of Z_i . In that case, the difference $M(X_i, 1) - M(X_i, 0)$ is shrunk towards zero and the CATE is biased.

5.3 T-Learner

The T-Learner addresses the limitation of the S-Learner by fitting two separate models: one is fitted on the control group and one for the action group:

$$M_0(X) = E[\hat{\theta}_i \mid X_i, Z_i = 0] \quad (5.5)$$

$$M_1(X) = E[\hat{\theta}_i \mid X_i, Z_i = 1] \quad (5.6)$$

The CATE is then estimated as the difference between the two predictions:

$$\widehat{CATE}_i^T = M_1(X_i) - M_0(X_i) \quad (5.7)$$

A limitation of the T-Learner is when the two groups are very unbalanced in size, the model trained on the smaller group will have higher variance.

5.4 X-Learner

The most complete meta-learner is the X-learner, it estimates the CATE in two different stages. The X-Learner exploits the observed outcome of each group to estimate the counterfactual outcome of the other group, which can lead to more efficiency in the estimates, especially when the two groups are unbalanced in size, which is the case the data from Euroconsumers with 3,367 observations in the control group and 6,744 in the action group of the discount and 6,749 observation in the action group with voucher.

First Stage: Outcome Models

The first stage consists of estimating two outcome models, one for each group, using any suitable machine learning algorithm:

$$M_0(X) = E[\hat{\theta}_i \mid X_i, Z_i = 0] \quad (5.8)$$

$$M_1(X) = E[\hat{\theta}_i \mid X_i, Z_i = 1] \quad (5.9)$$

where $M_0(X)$ is trained on the control group ($Z_i = 0$) and $M_1(X)$ is trained on the action group ($Z_i = 1$). Each model learns the relationship between the covariates X_i and the pseudo-observed retention time $\hat{\theta}_i$ within its respective group.

Second Stage: Imputed Treatment Effects

In the second stage, the X-Learner uses the predictions of $M_0(X)$ and $M_1(X)$ to impute the individual treatment effects for each group separately:

$$\hat{D}_i^1 = \hat{\theta}_i - M_0(X_i) \quad \text{for } i \in \{Z_i = 1\} \quad (5.10)$$

$$\hat{D}_i^0 = M_1(X_i) - \hat{\theta}_i \quad \text{for } i \in \{Z_i = 0\} \quad (5.11)$$

Intuitively, $\hat{D}_i^1$ represents the estimated treatment effect for individuals in the action group: it compares their observed outcome $\hat{\theta}_i$ to what $M_0(X)$ predicts they would have obtained had they been in the control group. Similarly, $\hat{D}_i^0$ represents the estimated treatment effect for individuals in the control group: it compares what $M_1(X)$ predicts they would have obtained in the action group to their observed outcome $\hat{\theta}_i$.

Two CATE models are then trained on these imputed effects:

$$\tau_0(X) = M_3(\hat{D}^0 \sim X^0) \quad (5.12)$$

$$\tau_1(X) = M_4(\hat{D}^1 \sim X^1) \quad (5.13)$$

Final Aggregation

The final CATE estimate is obtained as a propensity-weighted combination of the two CATE models:

$$\hat{\tau}(X) = g(X) \cdot \tau_0(X) + (1 - g(X)) \cdot \tau_1(X) \quad (5.14)$$

where $g(X) = P(Z_i = 1 \mid X_i)$ is the propensity score, estimated from the data. The propensity score serves as a weighting function: when $g(X)$ is large (i.e. most individuals with profile X are in the action group), more weight is given to $\tau_0(X)$, which is estimated more reliably on the control group. Conversely, when $g(X)$ is small, more weight is given to $\tau_1(X)$.

The full procedure is illustrated in Figure 5.1 (It does not have the same notation as in the thesis but it offers visualisation of the method used).




Figure 5.1: X-Learner Procedure

5.4.1 X-Learner on the data from Euroconsumers

We have to build two distinct analyses with the X-Learner to respect the SUTVA assumption, as if we don't make the distinction $Y_i(1)$ could represent the potential retention time if receiving the email with the voucher and with the discount which is not the same treatment. In both the studies, the control group is set as the control group and we separate the action group in two with $Z = v$ and $Z = d$, as in Figure 3.3.

The mean CATE for the voucher (0.130 months) and the discount (0.147 months) are consistent with the $\widehat{ITT}_{RMST}$ estimates of 0.115 and 0.131 months reported in Table 4.3, which validates the X-Learner estimation at the population level. We also observe a large deviation indicating the heterogeneous individual responses to the assignment variable. Some individuals have negative CATE values, suggesting that for some client, receiving the email reduces the retention of the clients.

In the voucher study the most responsive profile is a female aged above or equals to 60, with the maximum seniority of 15, and a high web activity and in wave 1. The model predicts a CATE of 0.589 months for this individual suggesting that her different characteristics positively affect the CATE. Similarly, in the discount study, the most responsive profile is a male aged between 30 and 44, with a seniority of 13 and a high web activity, with an estimated CATE of 0.720 months. It is also interesting to note that the five most responsive observations to the discount study are all the same suggesting that this combination of covariates has a positive impact on the CATE.

In Figure 5.2, we see that the two uplift curves of the two assignments. In both cases, the curves lie above the random baseline, confirming that targeting the

most responsive clients identified by the model yields a much higher cumulative gain than random targeting. It suggests that Euroconsumers could improve the profitability of its later campaigns by targeting approximately the top 50 or 60% (Starts decreasing at this point) of clients ranked by predicted CATE, rather than sending the incentive to the entire population.

Table 5.2: X-Learner CATE Summary Statistics at $\tau = 9$ months

Treatment	n total	n treated	n control	Mean	Std	Min / Max
Voucher	10,116	6,749	3,367	0.130	0.181	$[-0.505, 0.589]$
Discount	10,111	6,744	3,367	0.147	0.239	$[-0.554, 0.720]$

Table 5.3: Most Responsive Client Profile — Voucher ($\tau = 9$ months)

Age Group	Seniority	WebVisit	Gender_M	Age missing	Wave	CATE
[60,110]	15.0	2	0	0	1	0.589

The top 5 individuals all share the same covariate profile

Table 5.4: Most Responsive Client Profile — Discount ($\tau = 9$ months)

Age Group	Seniority	WebVisit	Gender_M	Age missing	Wave	CATE
[30,44]	13.0	2	1	0	1	0.720

The top 5 individuals all share the same covariate profile




Figure 5.2: Uplift Curves

Chapter 6

Conclusion

This thesis identified the causal effect of a marketing campaign conducted by Euroconsumers on client retention, using two different frameworks: the Randomized Controlled Trial (RCT) framework that we have used with the Instrumental Variable framework, and the Causal Machine Learning framework with the Double Machine Learning technique and Meta-Learners. The campaign was randomly divided into three groups of clients: one group received an email containing a 30-euro voucher, another group received an email containing a 50% discount on their subscription for six months, and the last group received nothing and served as a baseline. The central question was to quantify the effect of receiving and using such an email on the expected retention time of the clients.

6.1 Summary of Results

Table 6.1 shows the main estimations obtained across the different method for both treatment types as $\tau = 9$ months (with the two waves) and at $\tau = 13$ (only Wave 1). It is important to add that the $\widehat{ITT}_{RMST}$, $\widehat{CACE}_{RMST}$ and the X-Learner are estimated under the IV assumptions and the SUTVA assumption (see Section 3.5), whereas the DML requires the conditional unconfoundedness assumption. In addition, the $\widehat{ITT}_{RMST}$ and the X-Learner estimate the effect of being assigned to the action group, whereas the $\widehat{CACE}_{RMST}$ and the DML method identify the effect of using the incentive D . The X-Learner approach reports range of minimum and maximum and not confidence interval in Table 6.1. Furthermore, the X-Learner should not be used for direct causal interpretation as it is better suited for campaign targeting. The results of the DML framework should be interpreted with caution as it requires the conditional unconfoundedness assumption, though the model produces narrow confidence intervals compared to the other methods.

Table 6.1: Summary of causal effect estimates — Voucher vs. Discount

Method	Estimand	Horizon	CI Voucher	CI Discount
$\widehat{ITT}_{\text{RMST}}$	Effect of Z on Y	$\tau = 9$	[0.005, 0.226]	[0.020, 0.241]
$\widehat{CACE}_{\text{RMST}}$	Effect of D on Y (compliers)	$\tau = 9$	[0.033, 1.549]	[0.212, 2.560]
DML	Effect of D on Y (adjusted)	$\tau = 9$	[0.730, 0.738]	[1.446, 1.461]
DML	Effect of D on Y (adjusted)	$\tau = 13$	[1.258, 1.274]	[2.361, 2.404]
X-Learner	Mean CATE (heterogeneous)	$\tau = 9$	[−0.505, 0.589]	[−0.554, 0.720]

Table 6.2 shows the financial results of the IV framework and the DML framework. We see that none of the campaign is significantly profitable in the observation window. The only possible profitable survey is under the IV framework at $\tau = 9$ while the other scenarios are clearly making net loss. Under the DML framework, by increasing the observations window, we see that the effect of the incentive accumulates over time and the net loss is decreasing. We may go to the profitable point beyond the observation window, but it is not possible to know how the effect of the incentive will be on the long run.

Table 6.2: Summary of financial results — Voucher vs. Discount

Framework		Horizon	CI Voucher (€)	CI Discount (€)
IV framework	Net (€)	$\tau = 9$	[−28,874; +954]	[−35,467; −5,593]
DML framework	Net CI (€)	$\tau = 9$	[−15,154; −14,994]	[−19,771; −19,574]
DML framework	Net CI (€)	$\tau = 13$	[−4,766; −4,448]	[−8,131; −7,587]

6.2 Critical Analysis

The different methods used in this thesis comes with its own set of assumptions, which must be considered when interpreting the result. At the end, the most credible results in the thesis are the $\widehat{ITT}_{\text{RMST}}$ and $\widehat{CACE}_{\text{RMST}}$, as they rely on the IV assumptions that were empirically verified when it's possible in Chapter 3. With the help of the Random Controlled Trials and the Instrumental variables Z_i , it guarantees that the characteristics are correctly balanced between the two groups, and the instrumental variable isolates the causal effect of using the incentive through the CACE. However, the resulting confidence interval are large which makes concrete interpretation difficult. These large intervals are driven by different factors as the low compliance rate (12% overall), and the restriction to Wave 1 at $\tau = 13$ months, which reduces the effective sample size.

The DML framework provides a covariate adjustment to estimate the causal effect and interpretable nuisance function coefficients. Even if the conditional unconfoundedness assumption is violated and $\hat{\tau}$ cannot be given a causal interpretation, the nuisance function coefficients retain their descriptive value: the

outcome regression $\hat{m}_0$ describes how the observed covariates are associated with retention time, and the propensity score $\hat{p}(X_i)$ describes which client profiles are more likely to use the incentive. Thus, these associations give information but they don't require the conditional unconfoundedness assumption. However, as discussed in Section 4.4.4, the conditional unconfoundedness assumption is difficult to justify in our setting given the limited number of observed covariates and the possibilities of unobserved characteristics. For example, the financial situation or the subscription motivation could influence both the decision to use the incentive and the retention time and it is not captured in our model.

The X-Learner is the last model used in our analysis, and it adds complementary perspective by estimating individual-level heterogeneous treatment effects and help to build uplift curves that confirm the potential improvement of profitability. Since Z_i is used as the treatment variables and we have proved that it respects the IV assumptions, it is known that the conditional unconfoundedness assumption is satisfied. However, the X-Learner cannot directly identify the effect of using the incentive D_i , as it will introduce endogeneity and bias our estimations.

6.3 Recommendations for Euroconsumers

Based on the results of the different approaches, different recommendations can be made to Euroconsumers.

A first recommendation is to keep this randomisation assignment between a control group and an action group as it deletes the confounding effect of observed and unobserved client characteristics. The treatment assignment variables if it is randomly assigned can be used as an instrumental variable and it enables rigorous causal evaluation.

Neither campaign is significantly profitable within the observation window, the voucher is preferred compared to the discount. A higher compliance rate is preferred (14.58% versus 9.43%), and in our setting, a complier in the discount campaign cost the half of a complier in the voucher campaign for similar effect on $\widehat{CACE}_{RMST}$.

The X-Learner reveals the heterogeneous reaction of the client to the use of the incentive. Clients with higher web activity are consistently among the most responsive profiles for both treatments, suggesting that web engagement is a strong predictor of campaign responsiveness. By targeting the campaign to the most active clients, rather than randomly sending emails, it will improve the profitability of the campaigns.

Bibliography

- Andersen, P. K. and Perme, M. P. (2004). Inference for outcome probabilities in multi-state models. *Lifetime Data Analysis*, 10:469–494.
- Angrist, J. D. and Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press.
- Austin, P. C. (2009). Balance diagnostics for comparing the distribution of baseline covariates between treatment groups in propensity-score matched samples. *Statistics in Medicine*, 28(25):3083–3107.
- Battocchi, K., Dillon, E., Hei, M., Lewis, G., Oka, P., Oprescu, M., and Syrgkanis, V. (2019). *EconML: A Python Package for ML-Based Heterogeneous Treatment Effects Estimation*.
- Branson, Z. and Keele, L. (2020). Evaluating a key instrumental variable assumption using randomization tests. *American Journal of Epidemiology*, 189(10):1192–1199.
- Chen, H., Harinen, T., Lee, J.-Y., Magamamp, M., Ng, P., and Paliwal, A. (2020). CausalML: A python package for causal machine learning.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68.
- Frangakis, C. E. and Rubin, D. B. (2002). Principal stratification in causal inference. *Biometrics*, 58(1):21–29.
- Hernán, M. A. and Robins, J. M. (2020). *Causal Inference: What If*. Chapman Hall/CRC.
- Kaplan, E. L. and Meier, P. (1958). Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53(282):457–481.
- Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition.

- Rubin, D. B. (1980). Randomization analysis of experimental data: The fisher randomization test comment. *Journal of the American Statistical Association*, 75(371):591–593.
- Schölkopf, B. (2019). Causality for machine learning. *arXiv preprint arXiv:1911.10500*.
- Shaw, C. (2026). *Causal Inference in Marketing: Panel Data and Machine Learning Methods*. Community Review Edition.
- Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242.

