

Understanding the determinants of the attitude-behaviour gap in sustainable consumption at the household level

Master Thesis

MSc Economics and Business [Data Science and Marketing Analytics]

*Erasmus School of Economics
Erasmus University Rotterdam*

Palash Saxena | 504885

Supervisor: dr. Martijn G. de Jong
Second Assessor: dr.ir. Rommert Dekker

Date: 24-07-2020

The views stated in this thesis are those of the author and not necessarily those of the supervisor, second assessor, Erasmus School of Economics or Erasmus University Rotterdam.

ABSTRACT

There has been growing concern regarding rising levels of consumption as the threat of climate change grows larger. A lot of research and focus has been devoted to promoting sustainable production during the last decades, but the importance of meeting it with the same values on the consumption side is coming to the forefront. However, this task has proved to be challenging due to the phenomenon that is called the *attitude-behaviour* gap. It has been observed that there tends to be a disparity between what consumers believe and how they act. Therefore, despite their calls for more sustainable measures, consumers are more likely to not act on those measures themselves. This does not only impede the progress towards sustainable consumption but also affects production decisions and profits. The existing literature on the subject suggest a range of factors that help to explain the gap and the study is often concentrated on specific regions and industry. This paper is concerned with investigating the factors that determine the gap for consumers all over the world, based on survey data collected online from youths. Using machine learning methods on a host of social, demographic, and cognitive variables that were measured in the survey, the paper highlights the factors that influence the gap. It demonstrates how gender, country of residence, risk aversion, perception, experience with climate change, trust, and source of knowledge play a crucial role. Based on these findings, marketers can help to achieve sustainability goals by targeting consumers differently and pushing their behaviour closer to their attitude. The paper concludes with suggestions that can help to understand these factors in more detail, that may be necessary to make strides in this field of research.

Contents

1	Introduction.....	6
1.1	Managerial relevance.....	7
2	Brief theoretical overview.....	9
3	Data	12
3.1	Data overview	12
3.2	Data cleaning.....	13
3.2.1	Dependent variable.....	13
3.2.2	Independent variables	13
4	Methods	14
4.1	Data subsampling	14
4.1.1	Under-sampling.....	14
4.1.2	Over-sampling.....	14
4.1.3	Mixed sampling.....	15
4.2	Data resampling	15
4.3	Regularization methods	16
4.3.1	Ridge regression.....	16
4.3.2	Lasso regression	18
4.3.3	Elastic net regression	20
4.4	Black box methods.....	21
4.4.1	Support Vector Machines (SVM).....	22
4.4.2	Artificial Neural Network (ANN)	25
4.5	Model evaluation metrics	28
4.5.1	Sensitivity	28

4.5.2	Specificity	29
4.5.3	Balanced accuracy	29
4.5.4	ROC-AUC	29
5	Results	29
5.1	Descriptives	30
5.1.1	Demographic variables	30
5.1.2	Opinions on climate change.....	30
5.1.3	Source of knowledge	31
5.1.4	Trust in leaders	31
5.1.5	Perception of oneself and others	32
5.1.6	Experience and expectations regarding extreme weather	32
5.1.7	Annual methane emissions	33
5.1.8	Annual growth in household consumption expenditure	33
5.1.9	Annual per capita plastic waste	33
5.1.10	Annual municipal water withdrawal.....	34
5.1.11	Annual meat production.....	34
5.1.12	Annual meat consumption	34
5.2	Model comparison.....	34
5.2.1	Regularization methods	35
5.2.2	Black box methods	36
5.3	Model selection	37
5.4	Model interpretation.....	38
5.4.1	Gender	38
5.4.2	Country of residence.....	39
5.4.3	Trust in media	40
5.4.4	Experience with climate change.....	40

5.4.5	Opinions on climate change.....	40
5.4.6	Perception of oneself and others	40
5.4.7	Risk aversion.....	40
6	General Discussion	41
6.1	Theoretical implications	41
6.2	Research limitations and further research	42
6.3	Concluding remarks	43
	References.....	43
	Appendix	52

1 Introduction

In recent decades there has been growing focus towards sustainable production and consumption because of climate change. There has been an increased pressure on the natural resources available to us due to a surge in consumption worldwide. The expenditure on consumption represents 60% of the global GDP and is expected to double by 2030, the same year as the deadline to reach the Sustainable Development Goals (SDGs) set by the United Nations. At the same time there has been a rapidly growing middle class across the world, which means that global consumption keeps on increasing. This transition to the middle class is also where the carbon footprint of people starts to increase, by consuming more and as a result also creating more waste (Mortillaro, 2017).

While there has been significant research and movement in sustainable production, the expectation that consumers will follow up on their 'green' values has been unrealised. Rather, the consumption behaviour has been found to be quite different from the 'green' attitude of the people which is referred to as the attitude-behaviour gap (Terlau & Hirsch, 2015). In general, people's attitudes do not match their behaviour and a review of 47 empirical studies estimated the correlation between the two at below 0.30 (Wicker, 1969). This discrepancy can be due to people's social and cultural conditioning combined with certain barriers that prevent their actions to be aligned with their attitudes (Brown et al., 2007).

To understand and mitigate this attitude-behaviour gap, first it needs to be looked at what sustainable consumption entails. It implies consuming goods and services in a manner that does not put a strain on the natural resources of the planet and the needs of the present and future generations (Agenda 21, 1992; Belz & Peattie, 2009). On a deeper level, it is a multidimensional concept which involves the behavioural, attitudinal, and cognitive aspects of the human psyche. While behaviour drives the ultimate choices of the people, attitude predicts what choices people are expected to make in the future as it is unaffected by short term events and explains their consumption patterns. Further, cognitive aspects shape the determination of people towards indulging in sustainable behaviours (Quoquab & Mohammad, 2017).

The difference in the environmental intent of the people and the environmental impact it creates has raised a significant need to investigate the nature and determinants of people's beliefs. This will help to explain the attitude-behaviour gap and direct the consumption behaviour towards sustainability. This study aims to explain said gap based on people's demographic characteristics, opinions and willingness to spend to mitigate climate change, trust in leaders, source of knowledge of climate change, perception of oneself and other people in their country, experience and expectations with extreme weather, and personal values and habits. These decisions based on beliefs can be categorised as: purchase of goods with high environmental impact (e.g., consumption of meat), use of environmentally important goods (e.g., consumption of water in households), and household waste disposal (e.g., re-using plastic bags and

recycling of plastic bottles) (Stern, 2000). The specific examples mentioned here are looked at in this study to map and analyse people's beliefs with their engagement towards these activities.

This study aims to show that a positive 'green' outlook alone does not translate into sustainable consumption. Specifically, by considering a diverse range of factors and employing different methodologies, the study investigates the reasons that contribute to this disparity and how they interact with it. This knowledge could help in designing policies that nudge consumers towards a sustainable pattern of consumption. The paper is divided into various parts to gain better insights of the parameters that influence the presence or absence of attitude-behaviour gap in the consumers. Chapter 2 expands on the theoretical framework that supports the prevalence of the attitude-behaviour gap, and outlines features that could explain it. The data used for this study is detailed in chapter 3, and the machine learning models that were used to investigate any relationship between the gap and the features are explained and presented in chapter 4. Chapter 5 elaborates on the preliminary analysis performed and focusses on the findings from the machine learning algorithms by underlining the features that affected it and in what manner. The paper concludes with guidelines based on the results, the limitations of this study and possible measures to mitigate them in Chapter 6.

1.1 Managerial relevance

Consumption choices can greatly affect the greenhouse gas footprints of the companies, with Unilever estimating that as much as 70% of its footprint depends on how consumers choose and dispose of products. It should be noted that consumption does not only involve the action of purchasing but also the acknowledgement and consent to the firms' practices that are not always transparent to the consumer (Gonzalez, 2019). How consumers choose brands and products is highly dependent on the heuristics provided by the whole marketing environment (Hauser, 2011). Further, any perceived changes in the concern towards climate change by the consumers can lead to an over-reaction on the production side (Devinney et al., 2010). The strain of how resources are consumed thereby rests on both the individuals and the marketers.

According to Kahneman (2011), there are two mental systems which influence decision-making. Mental system 1 decides day to day behaviour based on heuristics whereas mental system 2 shapes slow, logical, and conscious behavioural change over a longer period. The pro-environmental survey responses are a result of the conscious thoughts and desires shaped by the latter while the actual purchasing behaviour depends on the former. At first, the goals of marketing may look at odds with those of sustainability; the former views resources as ever abundant and needs as ever-increasing while the latter stresses upon needs to be restrained as resources are limited (White et al., 2019). But as previously discussed how marketing

influences consumers' heuristics and how mental systems influence consumption, it seems that marketing has a huge role to play in sustainable use of resources.

There are many ways marketers can formulate strategies to initiate and even accelerate sustainable behaviour. Given how much social influence plays a role in society, campaigns that stress responsible behaviour by certain segments tend to influence similar behaviour by others. Such is the pattern seen with the installation of solar panels, which depends on whether the neighbours have them installed as well (Bollinger & Gillingham, 2012). The simpler method is to make sustainable behaviour as the default option such that taking the action becomes easier. One study suggested placing recycling bins nearby to make it easier for people to engage in recycling activities (Ludwig et al., 1998). As people form positive habits, it tends to have a positive spill over by encouraging them to indulge more. This was observed with IKEA's *Live Lagom* initiative in which people started with small steps like reducing household waste to having more efficient energy systems at home. Brands can also use emotional appeal in the form of praise or guilt to guide people into making better decisions. In one study it was seen that people who were publicly praised for their energy efficient activities saved more energy than people who were paid for the same (Handgraaf et al., 2011). Framing information rationally is also quite critical. It is seen that people are not themselves motivated to buy energy efficient products, but they can be if shown the long-term consequences in the form of monetary costs (Hardisty et al., 2016).

The environmental impact of any one individual is small, but quite significant if many individuals do the same thing. Most behaviours related to sustainability are driven by personal habits or routine, while others depend on factors like income and available infrastructure. Moreover, environmental impacts may be unknown to the consumers or seem insignificant compared to the choices they must make. At times, people make choices that are environmentally beneficial but that do not necessarily arise from environmental concerns. It is demonstrated that tendencies towards sustainable consumption vary remarkably depending on the context, the behaviour and the agent pursuing it (Stern, 2000). Such predispositions or the mismatch between beliefs and actions provides room for the creation of awareness amongst consumers which can bring out behavioural changes to defeat unsustainable habits accrued over time (Kahneman, 2011; Ruckriegel, 2014). Social marketing can be much more effective than traditional environmental awareness campaigns, as they first investigate behaviour and then build on a tactic targeting that behaviour. It has proven to be successful in bridging the gap and lead to many sustainability projects (McKenzie- Mohr, 1999). Therefore, with the ever expanding role of marketers in the present world in making individual lifestyle decisions, it can also take the shape of a tool to responsibly guide consumerism.

2 Brief theoretical overview

There is vast literature on the prevalence of attitude-behaviour gap amongst consumers and the factors that drive it. Various studies used a different mix of factors to explain the inconsistency in beliefs and consumption. Wall (1995) observed that it is more important to study how concern of climate change translates into behaviour than the level of concern and action. Maitney (2002) studied how experience can influence consumer behaviour and how valuable it is for shifting to a pro-environment behaviour. Tindall et al. (2003) argued that women in comparison to men are more pro-active at the individual and political level, but only found evidence for the former. While Kolmuss and Agyeman (2002) used a mix of internal factors like knowledge and external factors like political climate to explain the gap. Limited knowledge and financial resources were important factors that helped in explaining the difference in gap across blacks and whites (Taylor, 1989). Courtney-Hall and Rogers (2002) concluded that environmental education does not necessarily lead to pro-environment behaviour. On the other hand, Korheren and Lappalainen (2004) argued that environmental education is affective to propagate pro-environmental behaviour at the household level. This was supported by the research project by Devine-Wright et al. (2004) which concluded that pro-environmental behaviour is brought about at the household level through children when they learn more about the environment. Kennedy et al., 2009 investigated if consumers were even aware of this gap and what factors are important for them to explain it. They concluded that both a lack of knowledge and the prevalence of copious amounts of contradictory information can influence the gap. Thus, there is no definite answer or consensus that can be reached in the extant research, which is quite extensive and based on different models. But there is a common overall presence of factors regarding values, cognition, socio-cultural influences, and demographics which are considered.

According to Caruana et al. (2005) the existing research can be broadly divided into two groups. First, that focuses on how the data collection brings out flaws in the study being conducted and can lead to a gap. They claim that quantitative surveys only capture the rational responses of the consumers and not the heuristics that shape their consumption patterns. The responses can be biased by a social desirability which are in line with the socially accepted norms (Trudel et al., 2009; Carrington et al., 2010), and the sample can also be biased by consumers that are close to the issues regarding ethical consumption of resources (Auger & Devinney, 2007). Further, they also tend to focus on factors like insufficient money, lack of time, unwillingness to change or incomplete information about products (De Pelsmaker et al., 2005; Shaw & Clarke, 1999). The second group comprises studies that challenge the assumption that consumer behaviour is driven solely by rational forces and not by social, cultural, or historical context (Caruana, 2007). It recognizes that there are several cognitive limitations that prevent the individual to take a deliberative action, as several habits and heuristics automate the process of decision making (Jackson, 2005). Therefore, they incorporate rational choice theory based on calculative decisions by the individual,

and conflicts between their deontological and teleological assessments as well (Shaw et al., 2015). While culture and experiences form the slowly forming deontological assessments, the degree of self-reflection shapes the sensitivity of the teleological assessments. (Davies et al., 2012). The present study is positioned towards the latter group, as a range of cognitive factors were considered here to explain the gap. Although the data used in the study included variables that allow for measuring the social desirability bias, due to limited data and scope it was not covered here. Nevertheless, the same variables were used to study the impact on the gap and exploring how they might have influenced the responses.

Rajecki (1982) highlighted two problems with the incorporation of attitude and behaviour in studies, the difference in time between when data collected is for attitude and that for behaviour, and a broad measurement of attitude. These concerns regarding research methodology were accounted for in the present study, as data for both attitude and behaviours was collected at the same time, and attitude was measured by multiple questions specific to the research problem. This is because when comparing attitude towards climate change with behaviours that contribute to climate change, there is a higher explicability of attitude towards that behaviour (Fishbein & Ajzen, 1975; Ajzen & Fishbein, 1980). But, consumer behaviour is said to be quite sticky, as it is based not only on the barriers present in their environment, restricted choices or the inequality in access to sustainable products, but also habits, cultural environment, expectations and social norms (Jackson, 2005). The present study focused on a mix of factors that could explain the gap. One of the factors that may prevent consumers to not act sustainably is that they may not have experienced the consequences of their decisions (McDevitt et al., 2007), so considering their experience with extreme weather can be helpful. The medium through which people get informed about climate change and the amount of information they have on the subject can be crucial. For example, negative framing of information in news generates more reaction than otherwise positive information that is conveyed (Herr et al., 1991). Some studies suggested that people feel they do not have enough knowledge to make ethical decisions (Bray et al., 2010). Trust plays a huge role in determining whether people believe the information conveyed to them so that they act on it. With new technologies and spread of fake news even through the media, politicians and famous people, there is an abundance of contradictory information which can significantly limit sustainable behaviour (Kennedy et al., 2009). It is observed that even if people agree with the science of climate change, it is still unthinkable how it could affect at the individual/household level (Gibson et al., 2010). However, personal responsibility can drive up sustainable behaviour as people with a high sense of responsibility are shown to engage more in pro-environmental behaviour (Hungerford & Volk, 1990). Demographic factors like country of residence, gender and education were considered as well. There is significant difference in concern regarding climate change across regions of the world. Diekmann and Franzen (1999) showed that even though people in poorer countries were aware of the severity of climate change, they did not have enough economic resources to act on it. There needs to be

proper infrastructure for such things (e.g., recycling and waste management) such that people can avail them properly. Further, people with enough resources who can satisfy their personal needs are able to think beyond and about social issues (Borden & Francis, 1978). However, it is also suggested that people with more resources consume more and thereby end up polluting more, whereas being poor is the best way to reduce one's environmental impact (Gibson et al., 2010). The level of agreeableness and temperament difference between males and females could also be a deciding factor. According to Lehmann (1999), women are more concerned about climate change and willing to change their behaviours compared to men. Finally, it is also noted that the level of education does lead to a more extensive knowledge regarding environmental issues, but it does not necessitate pro-environmental action (Kolmuss & Agyeman, 2002).

A lot of studies have focused on determining the gap in purchasing behaviour of consumers towards organic products and on the purchasing barriers that complicate sustainable behaviour (e.g., Cottingham & Winkler, 2007; Naspetti & Zanolli, 2009; Hughner et al., 2007; Schöberl, 2012). However, not a lot of them focused on the analysis of attitude-behaviour gap at the household level. The few which do either focused on customer segmentation (e.g., Newton & Meyer, 2013; Gust, 2004; Anable, 2005) and tried to explain gap based on different types of customers or focused on a certain behaviour for a specific region or country (e.g., Gibson et al., 2010). This paper tried to conduct a study with responses from participants across the world and to include behaviours related to consumption at the household level. The importance of understanding the reasons behind how individuals and households consume and what factors can bring about a behavioural change was highlighted by Heiskanen and Pantzar (1997) who argued that without this understanding changes in consumption patterns by consumers, businesses, and governments cannot be brought about. Households represent not only places where people live but also where they share, form their identity, and invest in emotions (Blunt & Dowling, 2006). Gibson et al. (2010) mentioned that even though consumption decisions by the households are economic, they are in the background driven by cultural values, norms, and beliefs. For that reason, it is important to decipher how households consume based on the different contexts and networks than merely motivating households to consume more sustainably. Moreover, the impact of consumption on the environment will keep changing with the increase in the number of households per year (Caird & Roy, 2006). Household behaviours in this paper included avoiding/ reducing meat consumption, reusing or recycling plastic bags and reducing water consumption, and there are various reasons why it was important to study them. First, the production of meat-based food consumes significantly more land area, freshwater and fossil energy as compared to plant-based food. To produce 1 kg of meat protein about 100 times more water and 11 times more energy is used than to produce 1 kg of plant protein (Pimentel & Pimentel, 2003). By changing to a plant-based diet, the emissions of greenhouse gases can also be significantly reduced. Second, water is the most important and largest natural resource of which just 3% is freshwater and even less is accessible to the cities. According to the Global

Environment Facility (GEF) freshwater reserves are predicted to drop by 40% with the growing demand, which is further threatened by climate change. Even though the domestic consumption of water comprises just 10% of the total water use, it is growing fast and can have a huge potential impact (Danielsson, 2019). As per the WHO, minimum per capita use of 25 litres per day is needed for basic needs compared to the 262 litres per capita per day used in the US. It is estimated that just by behavioural changes this water consumption can be halved (Foerch, 2007). Third, the lightweight, durable, and inexpensive nature of plastics has contributed to its extensive use worldwide for more than 60 years. A lot of oil and energy is expended into its production and compounded with the fact that it is mostly non-biodegradable (Hopewell et al., 2009), its reuse and recycling become important issues for sustainability. It is estimated that 8 metric tonnes of plastic end up in the ocean each year. With a growing middle class, by 2050 the demand for plastic looks set to double over the 300 million metric tonne usage in 2017 (d'Ambrières, 2019). At the household level, plastics are used in thousands of products, but this study focuses on the reuse of plastic shopping bags and recycling of plastic bottles.

3 Data

3.1 Data overview

The data used in this survey was collected through Facebook as part of a survey conducted by The State of Youth which was founded by the KidsRights organization. State of Youth is a non-political and non-religious platform to make young people aware and give them voice against the most pressing issues in the world. KidsRights is an academic partner of the Erasmus University, and the survey was developed by the academics of the university. The survey options were in English and it was carried out through Qualtrics. Intermediate results of the survey were presented to the world leaders at the United Nations General Assembly on September 23.

The data was collected worldwide from youths between the ages 13-24 in 2019 between September 21 and November 25. Around 10,900 respondents all over the world answered questions regarding their socio-economic details, motivation, and opinions on climate change. 51.5% females and 47.5% males with an average age of 19 years comprised the respondents, with 40% of them having completed a university education at the time. The percentage of respondents who reported their mathematical ability and understanding of English to be between *good* and *very good* were 68% and 95% respectively. 98.63% had no problem understanding the questions. The survey consisted of two parts, one mandatory and one optional part. The compulsory part dealt with the general information of the participants, their opinions on climate change, their attitude and behaviour regarding it. While the optional part focused more on their willingness to tackle climate change and their perception regarding others and themselves.

3.2 Data cleaning

This study made use of both the mandatory and optional parts of the survey. As not all respondents chose to answer the optional part, there were a lot of missing values, and there were also missing values in the parts they opted to answer. Given that the methods used in this study do not work with missing values, and imputing values for survey responses are not meaningful, the data could not be used as is. Therefore, all missing values were removed from both parts of the dataset. As such, the resulting dataset contained only 3781 responses of the total number represented by respondents from 67 countries. The dependent variable was extracted from a couple of questions from the survey, so that it represented the Attitude-Behaviour Gap of the respondents. The remaining questions in the survey formed the independent variables that were used in this study.

3.2.1 Dependent variable

The two questions in the survey ‘*Do you think that climate change is a serious problem?*’ and ‘*Do you think that extra actions are needed to tackle climate change?*’ were used to map the attitude of the respondents. Any respondent answering above 4 (neutral) on the 7-point Likert scale on either of the questions was marked as having a sustainable attitude¹. The respondents were also asked if in the past 12 months they and/or their family members did any of the following: avoided eating meat or reduced their meat consumption, re-used plastic shopping bags and/or recycled plastic bottles or reduced water consumption. If any of the respondents reported either of these behaviours, they were marked as having a positive sustainable behaviour². In the end the difference between the attitude and behaviour was taken. For all the instances where the difference between the attitude and behaviour was 0, such that there was no attitude-behaviour gap, they were labelled as *No_Gap*. Instances with a difference of 1, such that there was an attitude-behaviour gap, were labelled as *Gap*. Few instances where the difference showed to be as -1³ were not considered and were removed from the dataset.

3.2.2 Independent variables

The remaining questions in the survey related to the respondents’ demographics (age, gender, country of residence, education, gender), opinions regarding climate change and willingness to spend, source of knowledge of climate change, trust in leaders, perceptions regarding themselves and youths in their country, experience and expectations with extreme weather, and personal habits and behaviours.

¹ Assigned a value of 1 in case of presence of sustainable attitude, 0 in case of absence of sustainable attitude.

² Assigned a value of 1 in case of presence of sustainable behaviour, 0 in case of absence of sustainable behaviour.

³ Cases with an absence of sustainable attitude, but with a presence of sustainable behaviour.

Some data external to the survey was also included in the analysis. These variables pertain to the most recent country-based data available on annual methane emissions (World Bank, 2012), annual growth in household consumption expenditure (World Bank, 2018), annual per capita plastic waste (Ritchie & Roser, 2018), annual municipal water withdrawal (Ritchie & Roser, 2017), annual meat production (Ritchie & Roser, 2017) and annual meat consumption intake (Ritchie & Roser, 2017).

Table 3.1 in Appendix presents an exhaustive list of all the independent variables that were considered in this study. It lists out all the survey questions along with their labels used in the analysis and their sub-categories/ levels on which they were measured.

4 Methods

4.1 Data subsampling

Datasets that have imbalanced classes in their dependent variables can lead to biased predictions and learning models when trained upon by machine learning algorithms that try to classify data. These classifiers tend to perform better for classes that are in the majority compared to minority classes. This can be attributed to the fact that most classifier learning algorithms like artificial neural network (ANN) and support vector machines (SVM) work on the assumption that the dependent variable is equally distributed in its classes. Therefore, assigning equal misclassification costs to imbalanced classes can lead to biased models which perform poorly for minority classes. As most real-life datasets encounter the problem of having huge class imbalances in the dependent variable, it is crucial to perform subsampling techniques which can balance the classes (Zheng & Jin, 2020).

4.1.1 Under-sampling

Upon under-sampling several instances of the majority class were removed at random from the dataset. This was performed until both the classes achieve parity in the counts of instances they contain (Japkowicz, 2000).

4.1.2 Over-sampling

In over-sampling, the minority class was increased in size by randomly re-sampling its instances with replacement until both the classes were of equal sizes (Japkowicz, 2000).

4.1.3 Mixed sampling

This method employs the techniques of both over-sampling the minority class and under-sampling the majority class. SMOTE⁴ was used in this study, which created synthetic instances for the minority class to over-sample rather than over-sampling with replacement. The majority class was under-sampled by randomly removing samples from it to achieve a specific degree of balance in the classes (Chawla et al., 2002).

4.2 Data resampling

In case of machine learning algorithms, the problem of over-fitting can easily occur, where the model generalizes on the data so well that it even starts to incorporate the noise in the data. In such cases, even though the model performs well on the data it trained on, it performs poorly on any new data that is fed to it. For tackling such a problem, a technique of cross-validation was applied, such that a subset of the training data was kept separately to evaluate the model.

To make the models more robust, the cross-validation procedure can be applied on many subsets of the data with different model hyperparameters, which is called a k-fold cross validation. As seen in Figure 4.1, the data was split into k subsets or folds, and the model was then trained on $k-1$ parts of the data (in blue). The remaining fold (in green) was used to evaluate the model, and this process was repeated k




Figure 4.1: k-fold cross validation.

times such that every time a different subset was held out for evaluation. The performances of the model across different folds were then compared for the different hyperparameter values. If there were multiple hyperparameters that could be tweaked, then a grid of all the values of all the hyperparameters was tested as part of the cross-validation. The procedure then went through each of the possible combinations of those values in different folds and calculated how well each one of them performed. Model performances were compared between the different folds, and based on how they fit the data, the best hyperparameter configuration was chosen (Grimm et al., 2016).

⁴ Synthetic Minority Over-sampling TEchnique

4.3 Regularization methods

Datasets with high dimensionality perform poorly on a simple model like the linear regression, which is also not very well suited for classification tasks. Independent variables can be highly correlated in high dimensional data, while at the same time it is not necessary that all variables are important to explain the dependent variable. For these reasons, regularization methods which can result in a sparse solution by reducing the dimensionality were preferred with a dataset with numerous variables.

$$RSS = \sum_{i=1}^n \left(y_i - \beta_0 - \sum_{j=1}^p \beta_j x_{ij} \right)^2 \quad (4.1)$$

Equation (4.1) represents a simple linear regression model, where n is the number of observations in the dataset and y_i is the response variable. β_0 is the intercept and β_j denote the coefficients for p number of predictors. RSS is the residual sum of squares which minimises the distance between the response and the predictors. A penalty was applied over this ordinary least squares (OLS) method of finding the best fit between the response and its predictors. There are three different ways of applying this penalty which were used in this study. These are namely, the L2 regularization or ridge regression, the L1 regularization or lasso regression, and the combination of L1 and L2 regularization or the elastic net regression.

While regression techniques work for numerical dependent variables only, they can be adapted to classification tasks. In a binary classification task, the threshold, or the point of separation of the classes of the dependent variable can be set at between its values⁵. The threshold can be changed depending on the task at hand and decides whether a given observation belongs to one class or the other. Therefore, all regression methods were adapted for performing classification tasks.

4.3.1 Ridge regression

This method added a shrinkage penalty to the least squares method to limit the problem of high dimensionality in the data. The coefficient estimates β_j of the model were compressed by the penalty such that coefficients of variables that were not significant for the model shrank to near zero.

$$\sum_{i=1}^n \left(y_i - \beta_0 - \sum_{j=1}^p \beta_j x_{ij} \right)^2 + \lambda \sum_{j=1}^p \beta_j^2 \quad (4.2)$$

Equation (4.2) shows the penalty β_j^2 that was added, where λ controls the extent to which it was applied. When $\lambda=0$, the method gets reduced to a simple regression as then there is no penalty. As its value increases, the impact of the penalty increases and shrinks the coefficients towards zero. The intercept is not

⁵ For instance, for a categorical dependent variable labelled as 0 and 1, the default threshold will be at 0.5.

affected, as the objective is to only compress the relationship between the response and the predictors. This can be seen in the left-hand panel of Figure 4.2, where coefficients slowly move towards zero as λ is increased.




Figure 4.2: Ridge trace plot with coefficients displayed against different values of λ and $\|\hat{\beta}_\lambda^R\|_2 / \|\hat{\beta}\|_2$.

Adapted source: James et al., 2017.

The amount to which the estimates are compressed can be seen in the right-hand panel, where $\hat{\beta}_\lambda^R$ represents the set of ridge coefficients at different values of λ and $\|\hat{\beta}\|_2$ is the L2 penalty. Therefore, with every increase in λ , the value of $\|\hat{\beta}_\lambda^R\|_2 / \|\hat{\beta}\|_2$ decreases until ultimately falling to near zero (James et al., 2017). The y-axis on both the panels of Figure 4.2 mentions that standardized coefficients are used in this method unlike the least squares. As different variables can be on different scales and the coefficients are regularized by the square of the coefficients, having absolute values can lead to larger penalties for larger values. This made it crucial to standardize variables before applying ridge regression.

Another feature of ridge regression is that it tackles the problem of high variance encountered in least squares coefficients, where even a small change in training data can have a large effect on their estimates. At the cost of introducing a small bias β_j^2 , L2 regularization can bring a huge drop in variance which makes the model and its estimates more robust (Kutner et al., 2018).

4.3.2 Lasso regression

With the ridge regression even though the coefficients could be reduced in size and brought to near zero values, it always contains all the predictors in the final model. Therefore, it does not lead to variable selection and the interpretation of the model depends on all predictors which can be complicated if number of predictors are large. This problem was solved by introducing L1 regularization to the least squares method.

$$\sum_{i=1}^n \left(y_i - \beta_0 - \sum_{j=1}^p \beta_j x_{ij} \right)^2 + \lambda \sum_{j=1}^p |\beta_j| \quad (4.3)$$

In eq. (4.3), the coefficients were compressed similarly as in the ridge regression but with the L1 penalty of $|\beta_j|$. λ played the same role of controlling the extent of compression as earlier. Left-hand panel of Figure 4.3 shows that an increase in the value of λ can ultimately drive the coefficients to zero. Similarly, in the right-hand panel of Figure 4.3, the set of lasso coefficients $\hat{\beta}_\lambda^L$ at increasing levels of λ can be seen shrinking by $\|\hat{\beta}_\lambda^L\|_1 / \|\hat{\beta}\|_1$. Every coefficient reduces to zero at some point, which can help to decide the number of predictors that can be kept in the model at different values of λ . This leads to a sparse solution as fewer number of predictors are retained in the final model, and its interpretation becomes much easier after lasso regression. The L1 penalty like the L2 penalty results in robust estimates of the coefficient as it reduces their variance by introducing some bias to the model.




Figure 4.3: Lasso trace plot with coefficients displayed against different values of λ and $\|\hat{\beta}_\lambda^L\|_1 / \|\hat{\beta}\|_1$.

Adapted source: James et al., 2017.

As the penalty of the model is $|\beta_j|$ instead of β_j^2 , the coefficients can at times be completely offset by the penalty. In such cases the variables are deemed to be not as important to explain the model, their

coefficients reduce to zero and lead to variable selection. This can be further explained by forming constraint functions for both lasso and ridge regressions. Both lasso and ridge coefficients are minimised

$$\sum_{j=1}^p |\beta_j| \leq s \quad (4.4)$$

$$\sum_{j=1}^p \beta_j^2 \leq s \quad (4.5)$$

subject to their respective constraint functions in eq. (4.4) and eq. (4.5). At every level of λ , s defines the boundary or the constraint that gives the same coefficients as eq. (4.3) for lasso and as eq. (4.2) for ridge regression. Small values of s imply that the coefficients will be small, while large values do not shrink the coefficients much. As seen in Figure 4.4, when $p = 2$ predictors, the smallest RSS is achieved when the ridge coefficients are inside a circle which can be denoted by $\beta_1^2 + \beta_2^2 \leq s$, and the lasso coefficients are inside a diamond which can be denoted by $|\beta_1| + |\beta_2| \leq s$. There is constant RSS at each red ellipse which gets bigger as RSS increases, while $\hat{\beta}$ denotes the least squares RSS. For sufficiently large values of s , the constraints regions include the least squares region. But when least square method is subject to the ridge




Figure 4.4: Constrain regions of ridge and lasso regressions (in blue), with the least square RSS (in red).

and lasso constraints, the respective coefficients are achieved at where the outmost ellipse meets the respective constraint regions. For ridge regression, the ellipse never meets the rounded edges of the circle on the y-axis. While the sharp edges of the diamond can intersect the ellipse at the y-axis, where $x = 0$. It can be seen how the lasso shrinks down the coefficients completely in comparison to the near zero approach

of ridge regression. The same reasoning also holds when $p > 2$ predictors with constraint regions in higher dimensions.

4.3.3 Elastic net regression

There can be many correlated variables in a dataset which could lead to the problem of multicollinearity in the model. The earlier methods solved this problem by introducing a little bias in all the coefficient estimates to offset the high variances or leaving out most of them completely. Even though lasso regression leads to variable selection, it picks out only one of all the correlated predictors (Desboulets, 2018). Figure 4.5 shows how small perturbations in the data can lead to unstable results with lasso regression even with a dataset that has been standardized. The elongated red ellipses represent the RSS in case of two correlated predictors, and the left-hand panel shows that the model retains β_1 and leaves out β_2 . But with a slight change in data as seen in the right-hand panel, the model ends up picking β_2 while leaving out β_1 in this instance.




Figure 4.5: Unstable lasso results in case of correlated features in the dataset

To solve for both high variances and retaining all relevant predictors whether correlated or not, elastic net regression was also used (Mol et al., 2009). As seen in eq. (4.6), it employs a mixture of both L1 and L2 penalties of the lasso and ridge regressions seen before. Where α decides the extent to which a

$$\sum_{i=1}^n \left(y_i - \beta_0 - \sum_{j=1}^p \beta_j x_{ij} \right)^2 + \lambda \left(\alpha \sum_{j=1}^p |\beta_j| + (1 - \alpha) \sum_{j=1}^p \beta_j^2 \right) \quad (4.6)$$

mixture of both the techniques are used. When $\alpha = 0$, it becomes a ridge regression while when $\alpha = 1$, it is just a lasso regression. For $0 \leq \alpha \leq 1$, the elastic net uses a specified mix of both the algorithms. This

changes the constraint region of the elastic net (Figure 4.6) resulting in a more convex boundary compared to the lasso. As a result, slight perturbations in the data with correlated variables keep the intersection of the ellipse and the constraint region at the centre. This ensures that all significant and correlated variables are selected in the final model and the results are stable unlike the lasso (Zou & Hastie, 2005).




Figure 4.6: Constraint region of the elastic net (in grey)

Intuitively, ridge regression can be used when there are many predictors which seem important to explain the response. Lasso regression can be used when there are many predictors which might not be important to explain the response. But when there are huge number of predictors and it is not clear if most of them should be retained or left out, elastic net can be used as a middle path between the two extremes. It leaves out predictors not significant for the model, while also retaining those that explain the model together.

Finally, for all regularization models mentioned above, the best value of λ was ascertained while fitting the model. By using the k-fold cross validation technique discussed before, a range of values were tested in a grid. This is referred to as a *grid search*, and the value of λ for which the lowest cross-validation error was achieved was chosen to be used in the final model (James et al., 2017).

4.4 Black box methods

There are certain machine learning algorithms which are not only powerful in their implementation but are also as complicated to comprehend (Lantz, 2015). The inner workings of these methods can be so complex that it blinds the user of exactly how it reaches upon a certain prediction or decision. This opaque process in between giving some inputs to the model and getting outputs from it is referred to as a black box. Black

box methods that were discussed and implemented in this paper are the Support Vector Machines (SVM) and the Artificial Neural Network (ANN).

4.4.1 Support Vector Machines (SVM)

The SVM is one of the most popular and high performing algorithms in machine learning. It is built on many underlying concepts that need to be discussed in the run up to the final workings of the method. The algorithm is fundamentally based on a decision boundary that can separate the classes⁶ of the response variable. This decision boundary is also called a *hyperplane*, which can be defined as in eq. (4.7). Where

$$\beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_p X_p = 0 \quad (4.7)$$

$X = (X_1, X_2, \dots, X_p)^T$ represents a point in a p -dimensional setting. If a point lies on the hyperplane then it satisfies the eq. (4.7). Otherwise, it will be greater or lower than 0, in which case the point belongs to one or the other class. If a point is far away from either side of the hyperplane, then it can be classified with a high confidence that it belongs to the particular class in which it is present.

Under the assumption that a hyperplane perfectly separates the two classes, there could be multiple hyperplanes which are able to do so. The farther away a hyperplane is from the observations, the better the classification will be, and the region of minimum distance between the observations and a hyperplane is called as *margin*. The *maximal margin hyperplane* is the boundary where the margin is the largest such that the minimum distance from the observations is the highest (James et al., 2017). In Figure 4.7 for two classes




Figure 4.7: Maximal margin hyperplane and the margin with two classes

(red and blue), the maximal margin hyperplane (black solid line) and the margin (region between the black solid line to either of the red dotted line or the blue dotted line) can be seen. This can be confirmed

⁶ Only problems with binary classes are considered here.

mathematically by the constraints given in eq. (4.8) and eq. (4.9), where y_i corresponds to the class of the response variable, and $x_{i1}, x_{i2}, \dots, x_{ip}$ to the instances in the predictors. The margin is maximized based on the constraints that the observations are at least M distance from the hyperplane and are on the correct side of the hyperplane. Given that the maximal margin hyperplane seems to be decided

$$\sum_{j=1}^p \beta_j^2 = 1 \quad (4.8)$$

$$y_i(\beta_0 + \beta_1 x_{i1} + \beta_2 x_{i2} + \dots + \beta_p x_{ip}) \geq M \quad \forall i = 1, \dots, n. \quad (4.9)$$

by just three observations (two blue and one red) in Figure 4.7, these points are called as *support vectors*. The position of the hyperplane would change with any change in the position of these observations, but not with changes in the remaining observations unless they disturb the margin (Lam et al., 2012).

However, classes are mostly not perfectly separable and even when they are, having a hyperplane that perfectly separates classes can be susceptible to overfitting. Any change in the training data that changes the support vectors can drastically change the hyperplane. Therefore, a *soft margin* is preferred instead, such that some observations can end up on the wrong side of the margin or even the hyperplane. This approach adds a little bias to decrease the variance in the results, such that the model becomes more robust. Having simple decision boundaries can also lead to higher accuracy of classifying test observations to the correct class than having complicated boundaries that try to get everything right on the training observations (Hamel, 2009). This underlines the concept of *support vector classifier*, which maximizes the margin with constraints given in eq. (4.10), eq. (4.11) and eq. (4.12). Where, ε_i refer to the *slack variables*

$$\sum_{j=1}^p \beta_j^2 = 1 \quad (4.10)$$

$$y_i(\beta_0 + \beta_1 x_{i1} + \beta_2 x_{i2} + \dots + \beta_p x_{ip}) \geq M(1 - \varepsilon_i) \quad (4.11)$$

$$\varepsilon_i \geq 0, \sum_{i=1}^n \varepsilon_i < C \quad (4.12)$$

and C is the budget parameter. The slack variables provide the flexibility for some observations to be on the incorrect side of the margin ($\varepsilon_i > 0$) or the incorrect side of the hyperplane ($\varepsilon_i > 1$). For $\varepsilon_i = 0$, the observations are on the correct side of the margin. The parameter C controls the amount of violations that can be made, such that for $C = 0$, no observations can be on the incorrect side. When $C > 0$, no more than C number of violations can be made. As the value of C increases and a greater number of observations can

trespass the margin, the number of support vectors that decide the hyperplane also increases⁷. Therefore, it controls the bias-variance trade-off discussed before, and its optimum value can be estimated by the k-fold cross validation procedure.

Another assumption that needs to be relaxed above is that the classes are always linearly separable. The response and predictors may be non-linearly related to each other and that must be accounted for. As seen in Figure 4.8, the input space of the predictors in the left-hand panel cannot be classified by a linear boundary. Although, by projecting their polynomial terms it could be possible to separate the classes linearly in a higher dimensional feature space as in the right-hand panel. However, with a large set of




Figure 4.8: Non-linearity in the input space and linear separation in the feature space

features or predictors this can become computationally intensive. To build an efficient support vector machine the *kernel trick* is used. It uses the dot product of the features to obtain their coordinates in higher dimensional space (Lantz, 2015). For a linear kernel, the dot product can be seen in eq. (4.13), which also represents a support vector classifier. K represents the kernel and the right-hand side is the dot product of

$$K(x_i, x_{i'}) = \sum_{j=1}^p x_{ij} x_{i'j} \quad (4.13)$$

$$K(x_i, x_{i'}) = (1 + \sum_{j=1}^p x_{ij} x_{i'j})^d \quad (4.14)$$

⁷Sometimes C is also denoted as a cost parameter, such that a higher value leads to a harder margin and lower violations, while a lower value gives a soft margin which is more open to misclassification errors.

two features. By extending this idea to polynomials with higher degree, the kernel would become as shown in eq. (4.14), where d is the degree of polynomial to be used. Combining this polynomial kernel with a support vector classifier leads to a support vector machine that can have a flexible decision boundary in a higher dimensional space. The optimum value of d used as part of SVM in this study was obtained through the k-fold cross validation process.

4.4.2 Artificial Neural Network (ANN)

Neural networks are a complex series of algorithms that are modelled on the workings of the brain. They consist of *neurons* that form a multi-layer network of nodes, which also makes its process in machine learning tasks hard to understand. A neuron can be depicted as taking in input signals (x), which are then




Figure 4.9: A neuron and its components

weighted with specific values (w) and passed through an activation function (f) to create an output signal (y) (Figure 4.9). Similarly, it can also be understood mathematically with the eq. (4.15), where n is the number of input signals and their corresponding weights. The weights measure how much each input should contribute to the sum of input values such that if it passes a certain condition then the activation function sends an output signal. If not, then the neuron does not pass any signal. Different activation functions are

$$y(x) = f\left(\sum_{i=1}^n w_i x_i\right) \quad (4.15)$$

meant for different purposes based on their different characteristics and applications. The range of output they give can be quite different and specific to the machine learning task at hand. The most common activation function that is used is the *logistic sigmoid*, shown in Figure 4.10. The feature of this function




Figure 4.10: Logistic sigmoid activation function

is that it squishes the output signal between 0 to 1 as can be seen in eq. (4.16), where e is the natural logarithm base. The results obtained with this function are quite compressed in the high and end tails of the curve, such that any sum below -5 or above +5 always result in 0 or 1, respectively. Therefore, it is

$$f(x) = \frac{1}{1 + e^{-x}} \quad (4.16)$$

crucial to standardize the independent variables before feeding to the algorithm, so that predictors with large values do not end up dominating the others.

Many neurons connected in multiple layers, taking inputs from the preceding layer, and passing on to the next layer until a final output is reached builds up the neural network. Figure 4.11 displays a neural network with one additional layer which is called the *hidden layer*, between the *input layer* and the *output layer*. The number of hidden layers to use in the network depends on the complexity of the task at hand, as each layer solves a part of the task. For tasks that are not so complex, only one hidden layer proves to be enough. As its value depends on the type of problem being solved, the k-fold cross-validation technique can be used to get its optimum value. Each of the nodes here represents a neuron tasked with functions depending on their layer, using their respective activation functions, and passing that value forward. The




Figure 4.11: Structure of a neural network with one hidden layer

input layer takes in the inputs and sums them with different weights. Then, a bias (b) term is added before passing it to the activation function (σ) as in eq. (4.17). The bias decides that only if the weighted sum is greater/smaller than the bias, then only can a neuron in the next layer be activated. Like the weights, the bias term is different for each neuron and this process is repeated for each of the neurons in the hidden layer. Therefore, the activations/inputs of the input layer ($a^{(0)}$) influence the activations in the hidden layer ($a^{(1)}$) depending on the different weights and biases. This process keeps on repeating through all the hidden layers and to the output layer in the end. When a value is passed on to the output node, it activates it for the last time and generates a value. This one cycle through the training data is called an *epoch*.

$$a^{(1)} = \sigma(Wa^{(0)} + b) \quad (4.17)$$

The first time a neural network starts training on the data, it chooses the weights and biases randomly. Then based on the output it produces compared to the actual output/response in the data, it starts tweaking the weights and biases so that it gets closer to the actual response. It goes through many epochs before it can extract the best setting to be able to solve the task. When it gets the output for a single training observation, the cost is measured for each output it gets wrong, which is the squared distance between the output and the response. The average cost after going through all the training observations is an indicator of how the network performs. By taking the negative gradient ($-\nabla$) of the cost function (C), it can be estimated in which direction the cost decreases the most. Therefore, by adjusting the weights and biases in the \vec{W} matrix in eq. (4.18) by the negative descent the cost can be minimised so that the network gets closer

$$-\nabla C(\vec{W}) \quad (4.18)$$

to the response. This process of minimising cost by taking a multiple of the negative gradient and nudging the model is also called *gradient descent*.

Once the network has reached the output layer, it can only retrace its steps and go backwards to make any necessary changes. Therefore, weights and biases are nudged in the preceding layer and then the one before that and so on. This is how one training observation goes on to change all the weights and biases in the model. It ensures that only the neurons which lead up to the correct result are activated and the neurons which do not are deactivated. The whole process of going backwards to adjust the model to get closer to the best setting of weights and biases is called *backpropagation* (Figure 4.12). This is repeated for all training observations and in the end all the changes to the parameters are averaged, which gives the value of the negative gradient of the cost function. As the entire mechanism can be quite computationally intensive, the data is usually sub-divided into mini batches and adjustments are made after calculating the gradient descent for each batch.




Figure 4.12: Backpropagation in a neural network to adjust weights and biases.

The amount by which the weights are adjusted after the cost is estimated by the gradient descent is known as *learning rate*. It decides how quickly the neural network will be trained, as large values mean large adjustments and thereby large decreases in cost. Small values, on the other hand, make small adjustments and decrease the cost slowly. If the learning rate is too large, it can overshoot the local minimum value of cost while if the value is too small it may never reach the minimum and get stuck. The value that reduces the learning rate is called the *decay*. As the learning rate and so the decay depends on the nature of problem the neural network is tackling, there is no one value that fits all. Therefore, the optimum value of decay was measured by the k-fold cross validation method.

4.5 Model evaluation metrics

This paper used methods which are completely different in their mechanisms. Therefore, some common evaluation metrics that help to compare how each of them perform are discussed. With imbalanced datasets, most of the classifiers perform well for the majority class and abysmally for the minority class. In such cases, the overall performance tends to be high because it is biased by the underlying distribution of the classes, which is known as *accuracy paradox*. Therefore, evaluation metrics that account for class imbalance were used (Sokolova et al., 2006). *Sensitivity* and *specificity* measured the accuracies of the respective positive⁸ class and negative⁹ class in the data. Similarly, *balanced accuracy* considered the performance of the classifier on both the classes, and not the biased overall accuracy. Lastly, the *ROC* curve was helpful to display the relationship between sensitivity and specificity.

4.5.1 Sensitivity

It measures the fraction of correctly classified response of the positive class (True Positive) out of the total responses of the positive class. The total responses of the positive class comprise of both correctly classified

⁸ Class that is of interest in the study which is labelled as 1

⁹ The remaining class, which is labelled as 0 or -1

positive responses and positive responses which were incorrectly classified as negative responses (False Negative).

$$Sensitivity = \frac{True\ Positive}{True\ Positive + False\ Negative} \quad (4.19)$$

4.5.2 Specificity

It measures the fraction of correctly classified response of the negative class (True Negative) out of the total responses of the negative class. The total responses of the negative class comprise of both correctly classified negative responses and negative responses which were incorrectly classified as positive responses (False Positive).

$$Specificity = \frac{True\ Negative}{True\ Negative + False\ Positive} \quad (4.20)$$

4.5.3 Balanced accuracy

The balanced accuracy measures the performance of the classifier based on how well it performs on both the classes. It is an average of the sensitivity and specificity as they measure the performance on the separate classes.

$$Balanced\ Accuracy = \frac{1}{2}(Sensitivity + Specificity) \quad (4.21)$$

4.5.4 ROC-AUC

The ROC curve and area under the curve (AUC) indicate how well the classifier can discriminate between the two classes. The ROC curve indicates the relationship between the true positive rate (sensitivity) and false positive rate (1-specificity). The AUC value measures the performance, with a higher value indicating a higher ability to separate classes and vice-versa.

5 Results

The variables in the dataset had to be adapted in a manner that fulfilled the input requirements of the machine learning models and aid them in making meaningful predictions. First, all the categorical variables in the dataset were converted to factor variables and then one-hot encoded. One-hot encoding transformed each category in these variables into individual features and noted their presence/absence by a value of 1/0. Second, the dataset was randomly divided into train and test subsets in a 75:25 ratio. This was done to fit the model on the train data and then to evaluate it on new data points in the test data. Third, all the numerical variables in the train data were scaled between 0 and 1. This is achieved by subtracting each value of a variable by its mean value and then dividing by its standard deviation. The same mean and variance is then

used to scale the test data to maintain the same scale. This was also done for the features that represented the Likert scales, as there was a mix of 7-point and 11-point scales in the data. As all the Likert scales were labelled with the numbers they were represented by on the rating scale, treating them as numeric variables and scaling them preserved their ordering. Lastly, the techniques of over-sampling, under-sampling and mixed sampling were applied on the train dataset to solve the problem of class imbalance in the dependent variable. After these steps to pre-process the data, machine learning methods were applied.

5.1 Descriptives

To obtain an overview of the data before relying on the estimates of the methods alone, descriptive statistics of seemingly important variables were calculated. Further, to verify if there was a significant difference in mean between *Gap* and *No_Gap* categories within variables, unpaired t-tests were conducted. Equation (5.1) describes the null hypothesis of equal means between the categories, and the alternate hypothesis of unequal means between the categories that is tested here.

$$\begin{aligned} H_0: \mu_{Gap} &= \mu_{No_Gap} \\ H_1: \mu_{Gap} &\neq \mu_{No_Gap} \end{aligned} \quad (5.1)$$

In the case of unequal variances between the categories, Welch two sample t-test was preferred. The tables in this section display the descriptive statistics, along with the F-test value for variances, t-test values, and their corresponding p-values at 95% significance level.

5.1.1 Demographic variables

On an average, the age of the respondents in the reduced dataset was between 18 to 19 years and they had completed their education up to age 18. The participants which showed a greater count of *Gap* were from India, Nigeria, Indonesia, South Africa, Brazil, United Kingdom, United States of America, Egypt, and Morocco. Given the difference in levels of agreeableness of males and females, a larger proportion of males showed a *Gap* compared to females and other genders.

5.1.2 Opinions on climate change

When it came to pick the best option to control climate change, most youths who showed a *Gap* were more in support of community focussed actions like building more awareness, having litter clean-ups, less recycling or declaring a climate change emergency. While almost none of them showed proclivity towards having high taxes on activities that harm the environment or being better prepared for natural disasters. The ones who showed *No_Gap* advocated mostly for clean production by the companies or having strict laws to prevent pollution.

5.1.3 Source of knowledge

When asked to rate the knowledge they gained from various sources about climate change, the mean rating for those who showed a *Gap* was always slightly higher than those who showed *No_Gap*. The largest differences were in case of movies and teachers (Table 5.1), which shows that *Gap* responses attributed their learning about climate change considerably more to these sources. Only in case of teachers and movies there was found to a significant difference in average values of *Gap* and *No_Gap*.

Table 5.1: Source of knowledge¹⁰

Variable	Category	Mean	St. Dev.	F	t	p
Teachers	<i>Gap</i>	4.71	2.07	1.13	2.27	0.02
	<i>No_Gap</i>	4.35	1.94			
Parents	<i>Gap</i>	4.22	2.23	1.27	1.46	0.15
	<i>No_Gap</i>	3.96	1.97			
Friends	<i>Gap</i>	4.56	1.97	1.12	0.49	0.62
	<i>No_Gap</i>	4.45	1.86			
Other People	<i>Gap</i>	4.93	1.77	0.97	0.16	0.87
	<i>No_Gap</i>	4.91	1.80			
News	<i>Gap</i>	5.79	1.47	0.92	0.42	0.67
	<i>No_Gap</i>	5.74	1.53			
Movies	<i>Gap</i>	5.13	1.79	0.81	3.89	<0.01
	<i>No_Gap</i>	4.50	2.00			

5.1.4 Trust in leaders

When asked to rate the extent to which they trust different leaders, the mean rating for those who showed a *Gap* was always slightly higher than those who showed *No_Gap*, except in the case for trust in scientists. On average, youths slightly to strongly disagreed that they trust the politicians, media, and famous people. The largest differences were in case of media and politicians (Table 5.2), which shows that *Gap* responses tend to trust them more relative to *No_Gap* responses. There was found to be a significant difference in the mean of the categories for politicians and media.

¹⁰ Measured on a 7-point Likert Scale

Table 5.2: Trust in leaders¹¹

Variable	Category	Mean	St. Dev.	F	t	p
Politicians	<i>Gap</i>	2.60	1.84	1.71	2.06	0.04
	<i>No_Gap</i>	2.30	1.40			
Media	<i>Gap</i>	3.80	1.94	1.34	2.44	0.01
	<i>No_Gap</i>	3.42	1.68			
Scientists	<i>Gap</i>	5.80	1.50	1.66	-1.78	0.07
	<i>No_Gap</i>	6.00	1.15			
Famous People	<i>Gap</i>	3.68	1.94	1.43	1.71	0.08
	<i>No_Gap</i>	3.40	1.62			

5.1.5 Perception of oneself and others

On an average, respondents perceived that a little over 50% of the youths in their country considered climate change to be a serious problem, and that only around 34% of the youths in their country would be willing to spend on climate change (Table 5.3). They also perceive that on an average they are slightly better informed than others about climate change. The results were similar across *Gap* and *No_Gap* respondents, and none of the variables showed a significant difference in the average values of the two categories.

Table 5.3: Perception of oneself and others

Variable	Category	Mean	St. Dev.	F	t	p
Percentage of young people who think climate change is a problem ¹²	<i>Gap</i>	51.70	24.54	1.20	-1.38	0.16
	<i>No_Gap</i>	54.27	22.42			
Percentage of young people are willing to spend on climate change ¹³	<i>Gap</i>	34.07	24.13	1.30	-0.39	0.69
	<i>No_Gap</i>	34.84	21.10			
Compared to other young people in your country, how much do you know about climate change? ¹⁴	<i>Gap</i>	5.41	1.17	1.12	-1.40	0.16
	<i>No_Gap</i>	5.53	1.10			

5.1.6 Experience and expectations regarding extreme weather

Relatively less proportion of *Gap* respondents reported to have experienced extreme weather in the past 12 months than *No_Gap* respondents. At the same time, more than 50% of the *Gap* respondents expected the

¹¹ Measured on a 7-point Likert Scale

¹² Measured on a 0-100 scale with breaks of 10

¹³ Measured on a 0-100 scale with breaks of 10

¹⁴ Measured on a 7-point Likert Scale

number of days with extreme weather to either remain the same or fewer in the next 12 months, while more than 50% people who showed *No_Gap* expected more days with extreme weather.

5.1.7 Annual methane emissions

It represents the amount of methane that was released into the atmosphere by a country in one year, and therefore the contribution to air pollution and sustainability. It could be related to the beef production in the countries and might influence how consumers respond to it. No significant difference was found between the mean values of methane emissions between the categories (Table 5.4).

Table 5.4: Annual methane emissions (kilotons)

Variable	Category	Min	Mean	Median	Max	St. Dev.	F	t	p
Annual methane emission	<i>Gap</i>	43.27	345927	223315.5	636395.8	270139.6	1.20	-1.38	0.16
	<i>No_Gap</i>	7.86	343621	223315.5	636395.8	263183.4			

5.1.8 Annual growth in household consumption expenditure

Table 5.5 shows the amount by which the annual growth in household consumption expenditure increases. This factor could have implications on how respondents from different countries contribute to sustainable consumption. With a p-value of less than 0.01, there was a significant difference in the average amount of household consumption between the categories (Table 5.5).

Table 5.5: Annual household consumption (% growth)

Variable	Category	Min	Mean	Median	Max	St. Dev.	F	t	p
Annual household consumption	<i>Gap</i>	-2.37	5.92	5.22	11.02	2.71	0.87	3.85	<0.01
	<i>No_Gap</i>	-2.37	5.90	5.13	13.26	2.90			

5.1.9 Annual per capita plastic waste

Per capital plastic waste might be able to explain respondents' behaviour regarding recycling and re-using plastic bags and bottles. There was also a significant difference in the amount of plastic waste generated per capita between the categories on an average (Table 5.6).

Table 5.6: Annual per capital plastic waste (kg)

Variable	Category	Min	Mean	Median	Max	St. Dev.	F	t	p
Annual per capita plastic waste	<i>Gap</i>	0.01	0.09	0.06	0.66	0.11	0.87	-2.41	0.01
	<i>No_Gap</i>	0.01	0.11	0.06	0.57	0.12			

5.1.10 Annual municipal water withdrawal

Annual municipal water withdrawal is an indicator of the amount of water that is consumed by municipalities and its activities in each country. It might be able to explain respondents' behaviour regarding household water consumption. The average amount of water withdrawal by both categories did not show a significant difference (Table 5.7).

Table 5.7: Annual water withdrawal (million m³)

Variable	Category	Min	Mean	Median	Max	St. Dev.	F	t	p
Annual water withdrawal	<i>Gap</i>	7.20	28880	13990	62090	24999	1.06	-0.07	0.94
	<i>No_Gap</i>	20.00	29030	13990	62090	24282			

5.1.11 Annual meat production

This factor was included in the model to study if the annual supply of meat has an effect over the respondents' behaviour regarding household meat consumption. However, no significant difference in the average values of the categories could be proved (Table 5.8).

Table 5.8: Annual meat production (thousand tonnes)

Variable	Category	Min	Mean	Median	Max	St. Dev.	F	t	p
Annual meat production	<i>Gap</i>	0.19	390.21	1980.84	12219.20	2563.66	0.80	-1.57	0.16
	<i>No_Gap</i>	0.47	922.00	1980.84	12219.20	2851.40			

5.1.12 Annual meat consumption

This factor was included in the model to study if the annual demand of meat has an effect over the respondents' behaviour regarding household meat consumption. As shown in Table 5.9, there seemed to be a significant difference between the average values of the two categories for the amount of meat consumed.

Table 5.9: Annual meat consumption (kilo calorie per person per day)

Variable	Category	Min	Mean	Median	Max	St. Dev.	F	t	p
Annual meat consumption	<i>Gap</i>	15	112.80	42	604	151.25	0.66	-5.94	<0.01
	<i>No_Gap</i>	15	187.70	71	604	186.00			

5.2 Model comparison

The machine learning methods mentioned in this paper were applied on the dataset and compared with each other. Each one of them were primarily measured based on the extent to which they were correctly able to predict the *Gap* (sensitivity), and then the accuracy of being able to predict both *Gap* and *No_Gap* (balanced

accuracy). The methods were applied to three different training datasets: under-sampled, over-sampled and mix sampled. K-fold cross validation with 10 folds was applied to tune the hyperparameters while training all the models, and the process was repeated 5 times.

5.2.1 Regularization methods

Table 5.10 shows the performance metrics for ridge regression, and so do Table 5.11 and Table 5.12 for lasso and elastic net regression, respectively. The tables mention the metrics balanced accuracy, sensitivity, and specificity for all three training samples that were used while training.

The under-sampled training dataset gave the best performance results out of all three models. It could be attributed to the fact that only original instances are preserved in the under-sampled data, as opposed to the creation of synthetic data in the over-sampled and mix sampled data. This could be crucial when working with survey data where unique and relevant information provides more information to the model than synthetically compiled instances. Models with over-sampled and mix sampled data always performed better at predicting the negative class (*No_Gap*) as seen with the specificity scores.

Table 5.10: Performance metrics for ridge regression

Sample	Balanced Accuracy	Sensitivity	Specificity
Under-sample	62%	60%	64%
Mixed sample	55%	28%	82%
Over-sample	63%	53%	74%
Positive Class: <i>Gap</i>			

Table 5.11: Performance metrics for lasso regression

Sample	Balanced Accuracy	Sensitivity	Specificity
Under-sample	65%	70%	60%
Mixed sample	56%	30%	82%
Over-sample	61%	58%	65%
Positive Class: <i>Gap</i>			

Table 5.12: Performance metrics for elastic net regression

Sample	Balanced Accuracy	Sensitivity	Specificity
Under-sample	66%	70%	62%
Mixed sample	56%	30%	81%
Over-sample	61%	53%	70%
Positive Class: <i>Gap</i>			

Sensitivity of 70% for both lasso and elastic net regressions, and 60% for ridge regression was achieved with the under-sampled data. As the ridge regression tried to retain maximum variables in the model, the results suggest that not all of them were relevant to the model which affected the performance. The elastic net managed to score slightly higher on the balanced accuracy (66%) metric compared to the lasso (65%). Even though the difference was not huge, it could have meant that variable(s) important to the model that might be left out by the lasso were included in the elastic net.

5.2.2 Black box methods

The support vector machines had mixed results with the different samples that were used to train the model (Table 5.13 and Table 5.14). The linear kernel achieved the highest performance for the over-sampled data with a sensitivity of 58% and balanced accuracy of 64%. The optimum level of cost parameter C found through 10-fold cross-validation was at 5 with 3121 support vectors, so a harder margin was preferred to make classification decisions. While the polynomial kernel performed best for the under-sampled data, it could only get 45% sensitivity, with the optimal cost at 0.1 with degree (d) at 3. This indicates that the

Table 5.13: Performance metrics for SVM (linear kernel)

Sample	Balanced Accuracy	Sensitivity	Specificity
Under-sample	62%	53%	71%
Mixed sample	55%	30%	80%
Over-sample	64%	58%	71%
Positive Class: <i>Gap</i>			

Table 5.14: Performance metrics for SVM (polynomial kernel)

Sample	Balanced Accuracy	Sensitivity	Specificity
Under-sample	59%	45%	73%
Mixed sample	60%	38%	83%
Over-sample	49%	0%	99%
Positive Class: <i>Gap</i>			

assumption of non-linear relationship between the variables did not improve the performance of the model in predicting for *Gap*.

The other black-box method that was used is the artificial neural network. The metrics for this method (Table 5.15) suggest that the under-sampled data gave the best results in terms of both sensitivity

(63%) and balanced accuracy (63%). The optimal number of hidden layers for which this model performed best were 7, so the network required much complexity to solve this task. Further, the optimal decay used was 0.2, which means the learning rate was decreased gradually to find the local minimum.

Table 5.15: Performance metrics for artificial neural network (ANN)

Sample	Balanced Accuracy	Sensitivity	Specificity
Under-sample	63%	63%	64%
Mixed sample	55%	33%	77%
Over-sample	55%	20%	91%
Positive Class: <i>Gap</i>			

5.3 Model selection

The elastic net regression model with the under-sampled data was ultimately picked to explain and interpret the dataset in this study. Not only was it able to achieve a superior sensitivity (70%), it was also able to separate the two classes better than the others with a balanced accuracy of 66%. This can be further seen in




Figure 5.1: AUC-ROC curve for elastic net regression with under-sampled data

Figure 5.1 which shows the AUC-ROC curve. It achieved a score of 0.71 which was higher than all other models, indicating that it was better able to make a distinction between *Gap* and *No_Gap* in the data. Further, it had the same values for sensitivity and balanced accuracy on the training data as well, which means there was no under-fitting or over-fitting in the model. The value of λ for which the model produced

the best results was at 0.099 (Figure 5.2), where it achieved the highest score for sensitivity. This was estimated through a 10-fold cross validation process which was repeated 5 times.

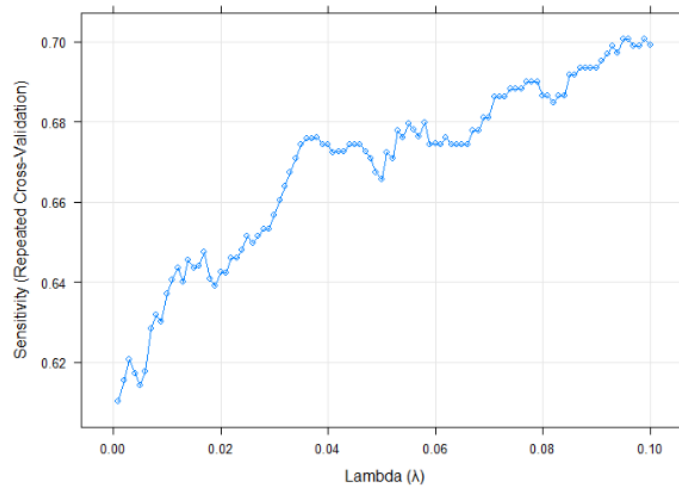


Figure 5.2: λ estimation with 10-fold cross validation (repeated 5 times) against sensitivity scores.

Further, the elastic net not only performed better but was also transparent with its results. Even though the neural network was quite close in terms of performance, it does not explain the model as clearly as the elastic net. Therefore, the elastic net was preferred in terms of both performance and interpretability.

5.4 Model interpretation

The importance of the variables used in this study to explain the attitude-behaviour gap can be seen in Figure 5.3. Only 8 sub-categories of all the variables used in this study were retained by the model, and the remaining were discarded. Gender differences seemed to play the most significant role, followed by the country the respondent resided in. The other factors that featured were based on the respondents' experience with climate change, trust in media, opinion on climate change, knowledge regarding the subject, and their risk aversion.

5.4.1 Gender

Females had the highest impact on the model. The higher the likelihood that a given respondent was female, the lower the likelihood that there was a gap between their attitude and behaviour. Almost all females considered climate change as a problem and thought extra actions are needed. 90% of the females followed up with re-using and recycling plastic bags.

Males on the other hand were predicted to increase the likelihood that there will be a gap between attitude and behaviour. While almost all males considered climate change a serious problem and advocated

for extra actions, only 43% admitted to reducing meat consumption, 64% to reducing water consumption and 80% to reusing and recycling plastic products.

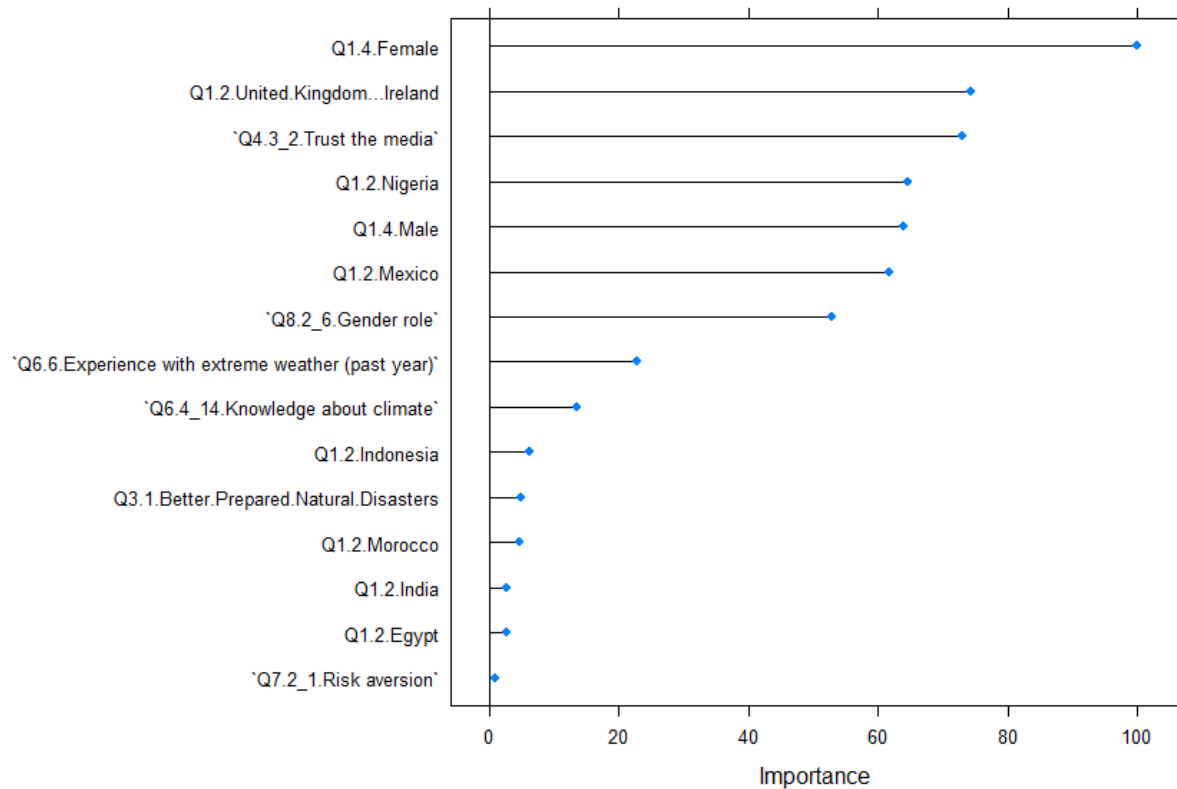


Figure 5.3: Variable importance from elastic net regression with down-sampled data

5.4.2 Country of residence¹⁵

With respondents from Indonesia, an emerging lower-middle income country, there was a higher probability of there being a gap. Almost all the residents of these countries showed sustainable attitude, but 88% reported reusing or recycling plastic products, 33% for a reduction in meat consumption and 73% for water consumption. While with respondents from developed and developing countries like United Kingdom of Great Britain and Northern Ireland, Mexico, and India it was predicted that they were most likely to not show a gap. Between 46%-62% reduced meat consumption, 81%-97% reused or recycled plastic products, and 56%-74% reduced water consumption.

The country level data with variables related to meat and water consumption, and plastic use and recycling did not prove significant to be able to explain the model.

¹⁵ The countries mentioned here represented at least 5% of the responses in the data.

5.4.3 Trust in media

Those respondents that showed a higher level of trust in media had a higher probability of not following up on their sustainable attitude. Less than half of the respondents who trusted the media reduced meat consumption, 62% reduced consumption water, and 82% reused or recycled plastic products. On the other hand, people who did not trust the media, 50% reduced meat consumption, 68% reduced water consumption, and 90% reused or recycled plastic bags.

5.4.4 Experience with climate change

The more experience respondents had with extreme weather in the past 12 months, the more likely were they to show no gap in their attitude and behaviour. Of the people who reported zero days of experience 44% reduced meat consumption, 77% reused or recycled and 55% reduced water consumption, while for people with experience 45% reduced meat consumption, 85% reused or recycled and 64% reduced water consumption.

5.4.5 Opinions on climate change

When asked about the best option to tackle climate change, people who felt that we should better prepared for natural disasters such as flooding, and drought were predicted to show less gap in behaviour. 37% and 57% of them reduced meat and water consumption, and 88% reused and recycled plastics.

5.4.6 Perception of oneself and others

When respondents felt that they knew more about climate change than other youths in their country, they were also more likely to not show a gap in behaviour. Compared to people who felt they knew less, around 10%-15% more people showed no gap when they reported knowing more than others.

5.4.7 Risk aversion

Respondents who were more willing to take risks were also more likely to have no gap in their behaviour. This was reflected in their behaviour where they had higher participation in all three behaviours than people who were less willing to take risks.

6 General Discussion

6.1 Theoretical implications

This study set out to ascertain the factors that can explain the attitude-behaviour gap by using survey data filled in by youths all around the world. Even though there were a lot of variables that were considered, there were only a few of them which were significant to have a final say in the model.

Gender differences proved to be quite important to understand the gap between attitude and behaviour concerning sustainable consumption. Women are known to be more pro-active than men at the individual level in converting the attitude into behaviour (Tindall et al., 2003). This was confirmed in this study, where on average a female respondent had a lower probability of showing the gap. Aspects of the personalities of both men and women can be further studied in detail to ascertain what separates them. Based on the findings, it would be possible to target men and women based on their respective personality traits to promote green behaviour.

The factor country of residence of the respondents, made a distinction between lower income, and developed and developing countries. Respondents from the former group were more probable to show the gap than the ones in the latter group. Therefore, income levels of countries and how it affects their citizens in making sustainable choices should be investigated. Further, other country level factors such as political environment, commitment to sustainable development and level of sustainable development could also be researched. Having a scale of factors determining the sustainability levels within countries could be helpful for governments to focus on areas most affecting the sustainable capabilities of their citizens.

The level of trust in media showed a negative relationship with the respondents following up on their sustainable attitude with the same sustainable behaviour. In other words, higher levels of trust in media increased the probability of there being the gap. This could be attributed to the information broadcasted through media channels regarding climate change and sustainability. As a result, it can be interesting to delve further into the kind of information that is transmitted, and how much of it assumes that climate change exists and should be tackled. Further, the reason behind people with low levels of trust to act more sustainably should also be researched. These questions may help to throw light on the credibility of information shared by the media and how people react to it.

It is conceivable that people who directly suffer from the effects of climate change the most, would also be more willing to take steps to mitigate them. This was confirmed by the model used in this study, as respondents who had experienced extreme weather in the past year showed to be less probable to show a gap in sustainable attitude and behaviour. Further study into whether regions highly affected by climate change show more sustainable behaviour (while also keeping other regional factors under consideration)

can be helpful. Providing a more direct taste of the effects of climate change to customers through innovative ways could be one of the solutions to drive pro-active sustainable behaviour.

Respondents who felt there was a greater need on being better prepared for natural disasters like floods and droughts had less probability to show the gap. Natural disasters around the world are increasing more than ever around the world and are believed to be brought about by climate change. If acknowledgement of this issue and of the need to act on it shows a relation with more sustainable behaviour, it could prove helpful to make people more aware of the issues close to their mind and environment. There can be widely different issues that affect different regions of the world and bringing more awareness about it to the local people could help build a diverse and significant sustainability drive.

The more the respondents perceived themselves as having a higher knowledge about climate change than their peers, the less probable they were to show a gap. This is based on people's tendency to uphold their social image, which ends up motivating them to be involved in acts that may elevate their image further. In this case, it was about having more knowledge about climate change than others. Therefore, one solution could just be to make people more aware and impart more climate change related knowledge which could lead to positive spill overs on how people choose to consume.

Finally, risk aversion was showed to have a negative relationship with sustainable behaviour as well. Respondents who were more inclined to take risks were also the ones to be less probable to show the gap. In this respect it could be useful to focus on the factors and traits that affect the level of risk aversion in people and to compare them with the traits of people who undertake sustainable consumption. This can help to decipher common and uncommon personality traits between the two sets of people and understand the dynamism behind it.

6.2 Research limitations and further research

There were various aspects that were either not considered in this study or were beyond its scope. One of the limitations of this research was the missing values in the survey data. Even though the survey was taken by many people, due to the copious amount of responses that were left unanswered in the survey, only a fraction of it could be used in this study. With more complete responses, more sophisticated machine learning methods could be used to analyse the data. Further, a lot of crucial responses belonging to those who showed the attitude-behaviour gap (minor class) were lost, which could have added more information in the analysis. Lastly, it limited the scope of the study to fewer countries and also resulted in a smaller representation of the countries that were retained.

One of the most common problems with using survey data is that of the social desirability bias. It has been shown that respondents at times tend to answer questions in line with what would be socially accepted. This tends to bring some bias when analysing data and is something common in marketing

studies. The current study did not account for the bias, as it involved removing more observations (biased responses) from the dataset. Therefore, the accuracy of the results presented here may be affected if there is considerable bias present in the data.

This research only scratches the surface of a multi-faceted problem, which must be tackled from various contexts. The different factors that affect the attitude-behaviour gap in consumers have been discussed in length, but each of them can be derived from their own set of contexts. How these factors influence the gap has also been established but the reasoning behind it has not been debated here. Only by understanding the latter can a more holistic model be achieved, which will not only inform us of what needs to be instructed but also how it can be propagated.

6.3 Concluding remarks

Several factors seem to contribute to the discrepancy between the consumers' beliefs and attitudes towards sustainable consumption and their behaviour. Gender, country, trust, experience, opinion, knowledge, and risk play their part in increasing or decreasing the gap. However, it only marks the initial foray into what is rather a much more comprehensive topic. Further opportunities exist for exploration of how these factors piece together and come into play, and for the refinement of the existing study.

References

Ajzen I. & Fishbein, M. (1980) *Understanding Attitudes and Predicting Social Behavior* (Englewood Cliffs, NJ, Prentice Hall).

Agenda 21. (1992). UN Konferenz Umwelt und Entwicklung. Rio de Janeiro.

Anable, J. (2005) —Complacent car addicts or —aspiring environmentalists? Identifying travel behaviour segments using attitude theory. *Transport Policy* 2005, 12, 65–78.

Auger, P., & Devinney, T. M. (2007). Do what consumers say matter? The misalignment of preferences with unconstrained ethical intentions. *Journal of Business Ethics*, 76, 361–383.

Belz, Frank-Martin & Peattie, Ken. (2009). *Sustainability Marketing: A Global Perspective*.

Blunt, A. and Dowling, R. (2006) *Home* Routledge

- Bollinger, B., & Gillingham, K. (2012). Peer Effects in the Diffusion of Solar Photovoltaic Panels. *Marketing Science*, 31(6), 900-912. <https://doi.org/10.1287/mksc.1120.0727>
- Borden, D. & Francis J.L. (1978) Who cares about ecology? Personality and sex difference in environmental concern, *Journal of Personality*, 46, pp. 190–203.
- Bray, J., Johns, N., & Kilburn, D. (2010). An Exploratory Study into the Factors Impeding Ethical Consumption. *Journal of Business Ethics*, 98(4), 597–608. <https://doi.org/10.1007/s10551-010-0640-9>
- Browne, A.L., Tucker, D.I, Johnston C.S. & Leviston, Z. (2007). Sustainable Water Use at the Household Level: A Scoping Report. CSIRO Water for a Healthy Country National Research Flagship Perth.
- Caird, S., & Roy, R. (2006): Household ecological footprints – demographics and sustainability. *Journal of Environmental Assessment Policy and Management* 8, 407–29.
- Carrington, M., Neville, B., & Whitwell, G. (2010). Why Ethical Consumers Don't Walk Their Talk: Towards a Framework for Understanding the Gap Between the Ethical Purchase Intentions and Actual Buying Behaviour of Ethically Minded Consumers. *Journal Of Business Ethics*, 97(1), 139-158. <https://doi.org/10.1007/s10551-010-0501-6>
- Caruana, R., Carrington, M. J., & Chatzidakis, A. (2015). “Beyond the Attitude-Behaviour Gap: Novel Perspectives in Consumer Ethics”: Introduction to the Thematic Symposium. *Journal of Business Ethics*, 136(2), 215–218. <https://doi.org/10.1007/s10551-014-2444-9>
- Chawla, Nitesh & Bowyer, Kevin & Hall, Lawrence & Kegelmeyer, W.. (2002). SMOTE: Synthetic Minority Over-sampling Technique. *J. Artif. Intell. Res. (JAIR)*. 16. 321-357. 10.1613/jair.953.
- Cottingham, M., & Winkler, E. (2007). The Organic Consumer. In S. Wright und D. McCrea (Ed.) *The Handbook of Organic and Fair Trade Food Marketing*: 29-53.
- Courtney-Hall, P., & Rogers, L. (2002). Gaps in Mind: Problems in environmental knowledge behavior modeling research. *Environmental Education Research*, 8(3), 283-297. doi: 10.1080/13504620220145438.

d'Ambrières, Woldemar (2009). Plastics recycling worldwide: current overview and desirable changes, *Field Actions Science Reports*, Special Issue 19 | 12-21.

Danielsson, M. (2019). Reduce Water Consumption at Home: SSWM - Find tools for sustainable sanitation and water management! Retrieved from <https://sswm.info/water-nutrient-cycle/water-use/hardwares/optimisation-water-use-home/reduce-water-consumption-at-home>.

David Hardisty, Yoonji Shim, Daniel Sun, & Dale Griffin (2016) , "Encouraging Energy Efficiency: Product Labels Activate Temporal Tradeoffs", in *NA - Advances in Consumer Research Volume 44*, eds. Page Moreau, Stefano Puntoni, and , Duluth, MN : Association for Consumer Research, Pages: 15-20.

Davies, I. A., Lee, Z., & Ahonhhan, I. (2012). Do consumers care about ethical-luxury? *Journal of Business Ethics*, 106(1), 37–51.

De Pelsmaker, P., Driesen, L., & Rayp, G. (2005). Do consumers care about ethics? Willingness to pay for fair-trade coffee. *Journal of Consumer Affairs*, 39(2), 363–385.

Desboulets, L. (2018). A Review on Variable Selection in Regression Analysis. *Econometrics*, 6(4), 45. <https://doi.org/10.3390/econometrics6040045>

Devine-Wright, Patrick & Devine-Wright, Hannah & Fleming, Paul. (2004). Situational influences upon children's beliefs about global warming and energy. *Environmental Education Research*. 10. 493-506. 10.1080/1350462042000291029.

Devinney, T. M., Auger, P., & Eckhardt, G. M. (2010). *The myth of the ethical consumer*. Cambridge: Cambridge University Press.

Diekmann, Andreas & Franzen, Axel. (1999). The Wealth of Nations and Environmental Concern. *Environment and Behavior - ENVIRON BEHAV*. 31. 540-549. 10.1177/00139169921972227.

Fishbein, M., & Ajzen, I. (1975). *Belief, Attitude, Intention, and Behavior: An Introduction to Theory and Research*. Reading, MA: Addison-Wesley

Foerch, W. (2007): Water Conservation. In: ROBBINS, P. (ed.). *Encyclopedia of Environment and Society*. Thousand Oaks, California (USA): Sage Publications, pp.1919-1922.

Gibson, C., Head, L., Gill, N., & Waitt, G. (2010). Climate change and household dynamics: beyond consumption, unbounding sustainability. *Transactions of the Institute of British Geographers*, 36(1), 3–8. <https://doi.org/10.1111/j.1475-5661.2010.00403.x>

Gonzalez, B. (2019). Why responsible consumption is everyone's business. *World Economic Forum*. Retrieved 31 March 2020, from <https://www.weforum.org/agenda/2019/09/why-responsible-consumption-is-everyone-s-business/>.

Grimm, K. J., Mazza, G. L., & Davoudzadeh, P. (2016). Model Selection in Finite Mixture Models: A k-Fold Cross-Validation Approach. *Structural Equation Modeling: A Multidisciplinary Journal*, 24(2), 246–256. <https://doi.org/10.1080/10705511.2016.1250638>

Gust, I. (2004). Strategies to promote sustainable consumer behaviour: the use of the lifestyle approach. MSc Thesis, Lund University, Sweden, 2004.

Hamel, L. (2009). *Knowledge discovery with support vector machines*. John Wiley & Sons.

Handgraaf, M., Van Lidth de Jeude, M., & Appelt, K. (2011). Public Praise vs. Private Pay: Effects of Rewards on Energy Conservation in the Workplace. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.1955898>.

Hauser, John. (2011). "A marketing science perspective on recognition based heuristics (and the fast-and-frugal paradigm)." *Judgment and Decision Making*, Vol. 6, No. 5, July 2011, pp. 396–408.

Heiskanen, E., & Pantzar, M. (1997). Toward Sustainable Consumption: Two New Perspectives. *Journal of Consumer Policy*, 20(4), 409–442. <https://doi.org/10.1023/a:1006862631698>

Herr, P. M., F. R. Kardes & J. Kim. (1991). 'Effects of Word-of-Mouth and Product-Attribute Information on Persuasion: An Accessibility-Diagnosticity Perspective', *Journal of Consumer Research* 17(4), 454-62.

Hopewell Jefferson, Dvorak Robert & Kosior Edward (2009). Plastics recycling: challenges and opportunities Phil. Trans. R. Soc. B3642115–2126 <http://doi.org/10.1098/rstb.2008.0311>.

Hughner, R., McDonagh, P., Prothero, A., Shultz, C., & Stanton, J. (2007). Who are Organic Food Consumers? A Compilation and Review of Why People Purchase Organic Food. *Journal of Cons. Behavior*, 6: 94–110.

Hungerford, H.R. & Volk, T.L. (1990) Changing learner behavior through environmental education, *The Journal of Environmental Education*, 21(3), pp. 8–21.

Jackson, T. (2005). Motivating Sustainable Consumption: a review of evidence on consumer behaviour and behavioural change. Sustainable Development Research Network.

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2017). An introduction to statistical learning: with applications in R. Springer.

Japkowicz. (2000). Learning from Imbalanced Data Sets: A Comparison of Various Strategies. AAAI Technical Report.

Kahneman, D., & Egan, P. (2011). Thinking, fast and slow. Random House Audio.

Kennedy, E. H., Beckley, T. M., McFarlane, B. L., & Nadeau, S. (2009). Why We Don't "Walk the Talk": Understanding the Environmental Values/Behaviour Gap in Canada. *Society for Human Ecology*.
<https://doi.org/https://www.jstor.org/stable/24707539>

Kollmuss, A., & Agyeman, J. (2002). Mind the Gap: Why do people act environmentally and what are the barriers to pro-environmental behavior? *Environmental Education Research*, 8(3), 239–260.
<https://doi.org/10.1080/13504620220145401>

Korhonen, K., & Lappalainen, A. (2004). Examining the environmental awareness of children and adolescents in the Ranomafana region, Madagascar. *Environmental Education Research*, 10(2), 195–216.
<https://doi.org/10.1080/13504620242000198177>

Kutner, M. H., Nachtsheim, C., & Neter, J. (2018). *Applied linear regression models*. McGraw-Hill Education/Asia.

Lehmann, J. (1999) *Befunde empirischer Forschung zu Umweltbildung und Umweltbewusstsein* (Opladen, Leske und Budrich).

Lam, H.-K., Ling, S. H., & Nguyen, H. T. (2012). *Computational intelligence and its applications: evolutionary computation, fuzzy logic, neural network and support vector machine techniques*. World Scientific.

Lantz, B. (2015). *Machine Learning with R - Second Edition*. Packt Publishing.

Ludwig, T., Gray, T., & Rowell, A. (1998). Increasing Recycling in Academic Buildings: A Systematic Replication. *Journal Of Applied Behavior Analysis*, 31(4), 683-686. <https://doi.org/10.1901/jaba.1998.31-683>

McKenzie-Mohr, D. & Smith, W. (1999) *Fostering Sustainable Behavior: an introduction to community-based social marketing* (Gabriola Island, Canada, New Society Publishers).

Maiteny, Paul. (2002). Mind in the Gap: Summary of research exploring 'inner' influences on pro-sustainability learning and behaviour. *Environmental Education Research - ENVIRON EDUC RES*. 8. 299-306. 10.1080/13504620220145447.

Mcdevitt, Roselie & Giapponi, Catherine & Tromley, Cheryl. (2007). A Model of Ethical Decision Making: The Integration of Process and Content. *Journal of Business Ethics*. 73. 219-229. 10.1007/s10551-006-9202-6.

Mol, C. D., Vito, E. D., & Rosasco, L. (2009). Elastic-net regularization in learning theory. *Journal of Complexity*, 25(2), 201–230. <https://doi.org/10.1016/j.jco.2009.01.002>

Mortillaro, N. (2017, November 14). More than 15,000 scientists from 184 countries issue 'warning to humanity' | CBC News. Retrieved from <https://www.cbc.ca/news/technology/15000-scientists-warning-to-humanity-1.4395767>

Naspetti, S., Zanolì, R. (2009). Organic Food Quality and Safety Perception Throughout Europe. *Journal of Food Products Marketing*, 15: 249-266.

Newton, P., & Meyer, D. (2013). Exploring the Attitudes-Action Gap in Household Resource Consumption: Does “Environmental Lifestyle” Segmentation Align with Consumer Behaviour? *Sustainability*, 5(3), 1211–1233. <https://doi.org/10.3390/su5031211>

Pimentel, D., & Pimentel, M. (2003). Sustainability of meat-based and plant-based diets and the environment. *The American Journal of Clinical Nutrition*, 78(3). doi: 10.1093/ajcn/78.3.660s.

Quoquab, F. & Mohammad, J., (2017). Managing Sustainable Consumption: Is It a Problem or Panacea?. Leal Filho W., Pociovalisteanu DM., Al-Amin A. (eds) Sustainable Economic Development. World Sustainability Series. Springer, Cham,.

Rajecki. D. W. (1982) *Attitudes: themes and advances* (Sunderland, MA, Sinauer).

Ritchie, H., & Roser, M. (2017, August 25). Meat and Dairy Production. Retrieved from <https://ourworldindata.org/meat-production>

Ritchie, H., & Roser, M. (2017, November 20). Water Use and Stress. Retrieved from <https://ourworldindata.org/water-use-stress>

Ritchie, H., & Roser, M. (2018, September 1). Plastic Pollution. Retrieved from <https://ourworldindata.org/plastic-pollution#citation>

Ruckriegel, K. (2014). Abschied vom Homo Oeconomicus, Stand 1.10.2014, "Wie nützlich sind die Wirtschaftswissenschaften - wie sind Wirtschaftswissenschaften nützlich?", Rahmenthema der 5. Lindauer Tagung der Nobelpreisträger für Wirtschaftswissenschaften vom 20. - 23.8.2014, available at: <http://www.ruckriegel.org.org>

Schöberl, S. (2012). Verbraucherverhalten bei Bio-Lebensmitteln: Analyse des Zusammenhangs zwischen Einstellungen, Moralischen Normen, Verhaltensabsichten und tatsächlichem Kaufverhalten. Doctoral disseration, Technische Universität München, 1-247. Available at: <https://mediatum.ub.tum.de/doc/1098882/1098882.pdf>

Shaw, D., & Clarke, I. (1999). Belief formation in ethical consumer groups: an exploratory study. *Marketing Intelligence and Planning*, 17(2), 109–120.

Shaw, D., Chatzidakis, A., Goworek, H., & Carrington, M. (2016). Explaining the Attitude-Behaviour Gap: why consumers say one thing but do another | Ethical Trading Initiative. *Ethicaltrade.org*. Retrieved 31 March 2020, from <https://www.ethicaltrade.org/blog/explaining-attitude-behaviour-gap-why-consumers-say-one-thing-do-another>.

Shaw, D., McMaster, R., & Newholm, T. (2015). Care and Commitment in Ethical Consumption: An Exploration of the ‘Attitude–Behaviour Gap.’ *Business Ethics*. <https://doi.org/10.1007/s10551-014-2442-y>

Sokolova, M., Japkowicz, N., & Szpakowicz, S. (2006). Beyond Accuracy, F-Score and ROC: A Family of Discriminant Measures for Performance Evaluation. *Lecture Notes in Computer Science AI 2006: Advances in Artificial Intelligence*, 1015–1021. https://doi.org/10.1007/11941439_114

Stern, P. C. (2000). New Environmental Theories: Toward a Coherent Theory of Environmentally Significant Behavior. *Journal of Social Issues*, 56(3), 407–424. doi: 10.1111/0022-4537.00175.

Taylor, D.E. (1989). Blacks and the environment: Toward an explanation of the concern and action gap between Blacks and Whites. *Environment and Behavior* 21, 2, 175-205.

Terlau, W. and Hirsch, D., (2015). Sustainable Consumption and the Attitude-Behaviour-Gap Phenomenon - Causes and Measurements towards a Sustainable Development. *International Journal on Food System Dynamics*, International Center for Management, Communication, and Research, [online] 6(3), pp.1-16. Available at: <<https://ideas.repec.org/a/ags/ijofsd/208880.html>>.

Tindall, D. B., Davies, S., & Mauboule`s, C. (2003). Activism and conservation behavior in an environmental movement: The contradictory effects of gender. *Society & Natural Resources*, 16(10), 909-932. <https://doi.org/10.1080/716100620>

Trudel, R., & Cotte, J. (2009). Does it pay to be good. *MIT Sloan Management Review*, 50(2), 61–68.

Wall, G. (1995). General versus specific environmental concern: A Western Canadian case. *Environment and Behavior* 27, 3, 294-316.

White, K., Habib, R., & Hardisty, D. (2019). How to SHIFT Consumer Behaviors to be More Sustainable: A Literature Review and Guiding Framework. *Journal Of Marketing*, 83(3), 22-49. <https://doi.org/10.1177/0022242919825649>.

Wicker, A.W. (1969). Attitudes versus actions: The relationship of verbal and overt behavioural response to attitude objects. *Journal of Social Issues*, 25, 41-79.

World Bank, World Development Indicators. (2012). Methane emissions (kt of CO2 equivalent) [Data File]. Retrieved from <https://data.worldbank.org/indicator/EN.ATM.METH.KT.CE>

World Bank, World Development Indicators. (2018). Households and NPISHs Final consumption expenditure (constant LCU) [Data file]. Retrieved from <https://data.worldbank.org/indicator/NE.CON.PRVT.KN>

Zheng, W., & Jin, M. (2020). The Effects of Class Imbalance and Training Data Size on Classifier Learning: An Empirical Study. *SN Computer Science*, 1(2). <https://doi.org/10.1007/s42979-020-0074-0>

Zheng, W., & Jin, M. (2020). The Effects of Class Imbalance and Training Data Size on Classifier Learning: An Empirical Study. *SN Computer Science*, 1(2). <https://doi.org/10.1007/s42979-020-0074-0>

Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2), 301–320. <https://doi.org/10.1111/j.1467-9868.2005.00503.x>

Appendix

Table 3.1: Variables and Categories

Label	Questions / Description	Levels
Dependent Variable		
<i>Gap</i>		‘Gap’, ‘No_Gap’
Independent Variables		
<i>Q1.1</i>	What is your age (in years)?	13, 14, 15, ..., older than 24
<i>Q1.2</i>	In which country do you currently live?	Afghanistan ..., Vietnam
<i>Q1.3</i>	Were you born in this country?	Yes, No
<i>Q1.4</i>	What is your gender?	Female, Male, Other
<i>Q1.5</i>	What is your highest completed education?	No formal education, Education up to age 12, Education up to age 14, Education up to age 16, Education up to age 18, Education after age 18, University

Table 3.1: Variables and Categories

<i>Q3.1</i>	Which option do you think is the best?	Climate emergency, High taxes, Strict laws, Planting trees, Litter clean-ups, Demonstrations, Consume responsibly, More awareness, Focus on wellbeing, Carbon footprint, Less waste & more recycling, Sustainable travel, Better technology, Clean production, Accepting, Be prepared
How much did you learn about climate change from the following sources?		
<i>Q4.2_1</i>	Teachers	Likert Scale 1-7
<i>Q4.2_2</i>	Parents	Likert Scale 1-7
<i>Q4.2_3</i>	Friends	Likert Scale 1-7
<i>Q4.2_4</i>	Other People	Likert Scale 1-7
<i>Q4.2_5</i>	News	Likert Scale 1-7
<i>Q4.2_6</i>	Movies	Likert Scale 1-7
To what extent do you disagree or agree with the following statements?		
<i>Q4.3_1</i>	I trust politicians	Likert Scale 1-7
<i>Q4.3_2</i>	I trust the media	Likert Scale 1-7

Table 3.1: Variables and Categories

<i>Q4.3_3</i>	I trust scientists	Likert Scale 1-7
<i>Q4.3_4</i>	I trust famous people	Likert Scale 1-7
<i>Q5.1_1</i>	How bad or good are you at mathematics?	Likert Scale 0-10
<i>Q5.1_1.0</i>	How good or bad is your understanding of the English language?	Likert Scale 0-10
To what extent do you disagree or agree with the following statements?		
<i>Q6.1_1</i>	I am willing to spend money to reduce climate change	Likert Scale 1-7
<i>Q6.1_2</i>	I am willing to spend money to reduce climate change, even if others don't	Likert Scale 1-7
<i>Q6.2_1</i>	According to you, what percentage of young people in your country think climate change is a serious problem?	0%-100% (intervals of 10)
<i>Q6.3_1</i>	According to you, what percentage of young people in your country are willing to spend money to tackle climate change?	0%-100% (intervals of 10)
<i>Q6.4_14</i>	Compared to other young people in your country, how much do you know about climate change?	Likert Scale 1-7
<i>Q6.6</i>	How many days with extreme weather did you experience in the past 12 months?	zero days, 1-5 days, ..., more than 20 days
<i>Q6.6.0</i>	Compared to this, how many days with extreme do you expect for the next 12 months?	Fewer days, Same, More days
<i>Q7.1_1</i>	In general, how willing are you to give up something today so you can benefit more in the future?	Likert Scale 0-10
<i>Q7.2_1</i>	In general, how willing are you to take risks?	Likert Scale 0-10

Table 3.1: Variables and Categories

<i>Q7.3_1</i>	In general, how willing are you to help others without experiencing anything in return?	Likert Scale 0-10
To what extent do you disagree or agree with the following statements?		
<i>Q8.1_1</i>	In my family we have enough money to buy things I want	Likert Scale 1-7
<i>Q8.1_2</i>	In my family we don't need to worry too much about paying our bills	Likert Scale 1-7
<i>Q8.1_3</i>	I don't think my family will have to worry about money in the future	Likert Scale 1-7
<i>Q8.2_1</i>	I never hide my mistakes	Likert Scale 1-7
<i>Q8.2_2</i>	I never take things that don't belong to me	Likert Scale 1-7
<i>Q8.2_3</i>	I don't gossip about other people's business	Likert Scale 1-7
<i>Q8.2_4</i>	I always obey laws, even if I am unlikely to get caught	Likert Scale 1-7
<i>Q8.2_5</i>	When I hear people talking privately, I avoid listening	Likert Scale 1-7
<i>Q8.2_6</i>	When jobs are scarce, men should have more rights to a job than women	Likert Scale 1-7
Please indicate for the following sentences whether you disagree or agree		
<i>Q8.3_1</i>	I work very hard most of the time	Disagree, Agree
<i>Q8.3_2</i>	I eat more than I should	Disagree, Agree
<i>Q8.3_3</i>	I like to visit new places	Disagree, Agree
<i>Q8.3_4</i>	Advertisements are silly	Disagree, Agree
<i>Q8.3_5</i>	My days follow a routine	Disagree, Agree
<i>methane</i>	Methane emissions (kt. of CO2 equivalent) - 2012	

Table 3.1: Variables and Categories

<i>house_cons</i>	Household consumption growth (% growth) - 2018
<i>per_capita_plastic</i>	Per capita plastic waste (kg) - 2010
<i>beef_prod</i>	Beef production (tonnes) - 2018
<i>water_withdrawn</i>	Municipal water withdrawal (m3/year) - 2015
<i>meat_cons</i>	Meat consumption (kilocalorie/person/day) - 2013