



Finanziato
dall'Unione europea
NextGenerationEU



Ministero
dell'Università
e della Ricerca



Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA



UNIVERSITÀ
DI SIENA 1240

Percorso dottorale sviluppato con il sostegno finanziario di NextGenerationEU:

Missione 4, Componente 2, Investimento 3.3, CUP B63C22000950008

Borsa MUR exDM352/2022, cofinanziata da Estra SpA

Doctoral Programme of the Tuscan Universities

PhD in Economics

XXXVIII° Cycle

Coordinator: Prof. Simone D'Alessandro

Information Frictions and Behavior in Green Markets

Scientific Sector: SECS-P/01

Candidate

Elena Sestini
University of Siena

Supervisor

Paolo Pin
University of Siena

Co-supervisor

Stefano Verde
University of Siena

Academic Year of the fulfilment of the title PhD in Economics

2024/25



University of Siena
PhD in Economics of the Tuscan Universities
XXXVIII° Cycle

Date of Thesis Defense:

January 2026

Committee:

Pierpaolo Parrotta (University of Siena)

Joachim Schleich (University of Grenoble)

Alessandro Tavoni (University of Bologna)

Contents

1	Voluntary Disclosure Under Bounded Rationality: Experimental Evidence	3
1.1	Introduction	4
1.2	Related Literature	5
1.2.1	Experimental Evidence of Unraveling	5
1.2.2	Beliefs and Behavioral Type Analysis	7
1.3	Experimental Design	7
1.3.1	The Revelation Game	7
1.3.2	Implementation	8
1.3.3	Research Hypotheses	12
1.4	Results	13
1.4.1	Descriptive Statistics	13
1.4.2	Sender Disclosures	14
1.4.3	Receiver Guesses	15
1.4.4	Analysis of Elicited Beliefs	17
1.4.5	Behavioral Analysis	19
1.4.6	CRT-2 Score	21
1.4.7	Regression Analysis	21
1.5	Welfare Analysis	23
1.6	Discussion and Conclusion	26
2	Walk the Talk? Greenwashing in the Electricity Market	30
2.1	Introduction	30
2.2	An overview of Greenwashing	33
2.2.1	Definitions	33
2.2.2	Historical Evolution	34
2.2.3	Regulations	34
2.3	Literature Review	35
2.3.1	Firm-level Greenwashing Measurement	35
2.3.2	Using AI to Detect Greenwashing	38

2.3.3	Greenwashing in the Energy Sector	40
2.4	Transparency and Accountability for Renewables in the Italian Electricity Market	41
2.4.1	Main Actors	41
2.4.2	Guarantees of Origin (GO)	42
2.4.3	Fuel Mix Composition	43
2.5	Data Collection	44
2.5.1	Web Scraping	44
2.5.2	Sustainability-related Keywords List	45
2.5.3	Fuel Mix Data	47
2.5.4	Data Preprocessing	47
2.6	Greenwashing Index	48
2.7	Variables	49
2.7.1	Dependent Variables	49
2.7.2	Independent Variables	50
2.8	Empirical Analysis	51
2.8.1	Descriptive Statistics	51
2.8.2	Regression Analysis	51
2.9	Discussion and Conclusion	55
3	Estimating Residential Electricity Demand under Time-of-Use Tariff: Evidence from Micro-Level Data	59
3.1	Introduction	59
3.2	Related Literature	62
3.2.1	Short- and Long-Run Electricity Demand Elasticity	62
3.2.2	Estimation of Electricity Demand Elasticity under ToU Tariffs	64
3.2.3	Italian Electricity Demand	65
3.3	Italian Electricity Market	67
3.3.1	Market structure	67
3.3.2	Time of Use (ToU) Tariff	68
3.3.3	An overview of demand and prices across time bands	69
3.4	Modelling Electricity Demand	71
3.4.1	Theoretical framework	71
3.4.2	Which Price to Use?	73
3.5	Data and Methods	74
3.5.1	Dependent Variables	75
3.5.2	Independent Variables and Controls	75
3.6	Results	76
3.6.1	Descriptive Statistics	76

3.6.2	Total Electricity Demand	78
3.6.3	Additional Regressions	79
3.6.4	Peak and Off-Peak Hours	81
3.6.5	Heterogeneity Analysis	84
3.7	Discussion and Conclusion	85
A	Appendix of Chapter 1	89
A.1	Between-Subject Variations	89
A.2	Payoff Tables (Within-Subject Variations)	90
A.3	Senders' Disclosure Rates	90
A.4	χ^2 Test Results	91
A.5	Receivers' Average Guesses across Treatments	92
A.6	Non-parametric Tests on Receivers' Guesses	92
A.7	Senders' Elicited Beliefs	92
A.8	Non-parametric Tests on Senders' Beliefs	93
A.9	A Comparison between Senders' Beliefs and Receivers' Guesses	94
A.10	Senders' Belief and Non-disclosure	95
A.11	A Comparison between Receivers' First Order Beliefs and Senders' Disclosure	96
A.12	Receivers' Average First Order Beliefs	97
A.13	Receivers' Average Second Order Beliefs	97
A.14	Receivers' Second Order Beliefs	98
A.15	Receiver Experiment: Receiver's Second Order Beliefs Non-parametric Tests	99
A.16	Regression Results	100
A.17	Mean Squared Error	102
B	Appendix of Chapter 2	103
B.1	Greenwashing Indices Distributions	103
B.2	Descriptive Statistics	105
B.3	Greenwashing Index	106
B.4	Ateco 2007 Classification	109
B.5	Sources for Words' List Creation	110
B.6	Additional Results	111
C	Appendix of Chapter 3	116
C.1	Market Structure	116
C.2	Tariff Example	118
C.3	Results	119
C.4	Additional Results: Heterogeneity Analysis	121
	List of Tables	125

List of Figures	127
Bibliography	129

Acknowledgment

I am deeply grateful to several people who, in different ways, have encouraged and supported me in reaching this goal. My first thanks go undoubtedly to my family. To my parents, *Gianna and Marco*, to whom I owe the values, determination, and education that have accompanied me throughout my entire academic journey. Thank you for always listening to me and encouraging me to find a solution even in the most difficult moments. A special thought goes to my grandfather, *Fidenzo*, to whom I have always shared updates about my research progress. His kind words and constant interest have been a precious source of motivation and strength along this demanding path.

I am sincerely grateful to my supervisor, *Paolo*, from whom I have learned so much both technically and personally. Thank you for the generosity, guidance, and constant support over these years. Huge thanks also go to *Stefano*, my co-supervisor, for his steady help and all I have learned from working with him. And finally, an enormous thank you to *Sébastien*, who has been not only a kind person and an excellent co-author, but also a brilliant supervisor during my first experience abroad. He made my stay in Switzerland a professional success and ensured that Lausanne will always hold a special place in my heart.

I am also truly thankful to all the *UNIL PhD students* who welcomed me as a colleague, and to my *EPFL PhD friends*, who made this experience unforgettable. I miss our lake dinners so much. A special thought goes to *Elena and Tanzhou*: it was such a joy to come to the office every morning knowing I would spend the day with two kind and special people who so quickly became real friends.

A heartfelt thank you goes to *Estra S.p.A.* for giving me the opportunity to work with unique and valuable material that greatly enriched my research. I am especially grateful to *Manuela*, who guided me throughout the process and worked with exceptional commitment to make the collaboration truly efficient. I also wish to thank all the skilled professionals I had the pleasure of meeting at the company, whose contribution made data sharing possible.

My thanks also go to my *fellow PhD colleagues*, among whom I have found not only collaborators but also friends who will remain part of my life well beyond academia, such as *Edo and Michi*. Looking back on these years, filled with moments of joy, challenges, and peaks of anxiety, I can now smile, precisely because I shared them with you.

I cannot forget my lifelong friends, *Ele and Bea*, with whom I share a deep and indescribable bond that always brings us back together, even when our paths seem to drift apart. They have always been by my side, through both good times and bad. And to all my other wonderful friends: you know who you are, thank you for being there. Each of you means more to me than words can say.

I wish to conclude these acknowledgments with the person I have chosen as my life partner, *Andrea*. He, together with my mother, is among the few who have truly seen the worst version of me during these three years, and to them I owe my most sincere thanks. With Andrea, there is no topic I cannot discuss, from personal matters to happy moments, from how to design our home or what to have for dinner, to my endless doubts about code or the best econometric method to use. If he still puts up with me after these three years, he probably can for a lifetime. Thank you for being a constant presence in the darkest moments of my life. Today, I finally dedicate to you a small but meaningful joy.

Introduction

This dissertation explores how information frictions and behavioral biases shape decision making and economic outcomes in markets where environmental and moral dimensions are crucial. Motivated by the growing importance of sustainability and transparency in recent decades, the goal is to investigate why asymmetric information still threatens efficiency, and how behavioral aspects and policy interventions can help mitigate these issues. It both investigates the household dimension, understanding whether market inefficiencies stem from unaware, careless, or biased choices, and corporate perspective, where image, revenues and insufficient monitoring come into play in presence of informational opacity.

The dissertation includes three chapters¹. The first chapter, *Voluntary Disclosure Under Bounded Rationality: Experimental Evidence* (co-authored with Sébastien Houde from the University of Lausanne) explores the strategic mechanism behind voluntary information disclosure. It employs a sender-receiver experiment framed in the housing market and tests whether bounded rationality distorts communication with asymmetry of information, compared to the theoretical benchmark of full unraveling (Milgrom, 1981).

The second chapter, *Walk the Talk? Greenwashing in the Electricity Market* (co-authored with Stefano Verde from the University of Siena), proposes a novel approach to measure greenwashing in the electricity industry. Combining numerical and text-based material, it analyzes heterogeneity in firms engaged in greenwashing and discusses which economic and demographic characteristics are more likely associated with misleading environmental communication.

The third chapter, *Estimating Residential Electricity Demand under Time-of-Use Tariff: Evidence from Micro-Level Data*, provides new evidence on residential electricity demand under ToU tariff, when the pricing system is not easily accessible. It employs a unique panel dataset from Estra S.p.A., the co-funder of my PhD scholarship with whom I have closely collaborated over the past three years, to estimate short-run demand elasticities, substitution patterns across peak and off-peak hours, and consumption behavior under different demographic and geographic characteristics.

¹Two are coauthored studies: while the analyses have been discussed with my coauthors, I am solely responsible for the final form as complete papers. Any remaining errors or typographical mistakes are my own.

My research employs a combination of methodologies, including experimental and behavioral analysis, cross-sectional and panel data econometrics, as well as web scraping and text-mining with Natural Language Processing (NLP).

Together, these chapters address several interconnected research questions: (i) under which conditions do individuals reveal private information, and how do cognitive biases distort strategic reasoning and communication? (ii) Do companies exploit private information to appear greener than they actually are, and which characteristics can help explain misleading behavior? (iii) Can price be used as an optimal policy tool to reduce electricity consumption or to diminish peak load, cutting costs and increasing efficiency?

By addressing these questions, this thesis contributes to three strands of the economic literature: behavioral and experimental economics, by identifying cognitive mechanisms underlying information disclosure; corporate environmental communication, by proposing an empirical approach to quantify greenwashing; and energy economics, by providing new micro-level evidence on demand responsiveness in the electricity market.

While Chapters 1 and 2 share a similar setting, focusing on strategic disclosure and communication under asymmetric information by individuals and firms, Chapter 3 does not take lack of price transparency as its primary object of analysis. Instead, it considers price opacity as a structural feature of the market that may affect households' responses to price fluctuations.

Overall, this research underscores the crucial role of adequate information in achieving economic efficiency, showing that not only accessibility but also transparency and credibility is essential to align private incentives with collective welfare. It contributes to a deeper understanding of how to effectively address imperfect and asymmetric information and how to promote responsible, efficient, and sustainable behavior among economic agents.

Chapter 1

Voluntary Disclosure Under Bounded Rationality: Experimental Evidence

Sébastien Houde¹, Elena Sestini²

¹ *School of Economics and Management, University of Lausanne*

² *Department of Economics and Statistics, University of Siena*

Abstract

We study how the structure of the decision environment, specifically, the common knowledge distribution of quality, affects voluntary disclosure behavior under bounded rationality. In an online sender-receiver game mimicking the rental housing market, we experimentally vary the granularity and shape of this distribution. Standard theory predicts no effect, but we find that these features significantly alter behavior. Participants systematically anchor on the mean of the distribution, a bias that impedes unraveling. This emerges in both senders and receivers and intensifies when the mean is more salient (e.g., under normal distributions) and the quality signal is more granular. Our findings reveal how the decision environment distorts strategic disclosure.

1.1 Introduction

In the presence of asymmetric information, voluntary disclosure is a simple policy instrument that can be a first-best solution when the right conditions are met. This result was first demonstrated by Milgrom (1981). He showed that sellers should be inclined to always report privately held quality information, regardless of the quality level. Otherwise, buyers would be skeptical of any withheld information, and infer that no news is bad news. In equilibrium, full information unraveling should then occur. In practice, it is, however, rarely the case that voluntary disclosure schemes fulfill their theoretical promises (Jin et al., 2021; Montero and Sheth, 2021).

Several frictions prevent the full unraveling of information. In this paper, we investigate the role of bounded rationality combined with the complexity of the decision environment as one particular type of friction. Our main hypothesis is that the nature of the quality distribution, in particular its shape and level of coarseness, can impede or facilitate individuals' higher-order strategic thinking, and, thus the well-functioning of a voluntary disclosure scheme.

We test this hypothesis using an online experiment resembling the housing rental market where landlords and potential tenants exchange information about property quality before concluding a transaction. The experimental design, inspired by recent experimental studies such as Jin et al. (2021), uses a simple sender-receiver game to replicate a real-world scenario where landlords are privately informed about their property's quality and must decide whether to disclose this information to prospective tenants. The quality of the house is initially only known to the landlord and cannot be manipulated. The tenants, in turn, try to estimate the property's quality when there is no disclosure.

In this setting, we manipulate two dimensions of the market environment: the shape and coarseness of the common knowledge quality distribution and investigate how it impacts unraveling and strategic thinking. We do not allow for interactions between senders and receivers. Instead, we conduct two separate experiments where participants act either as senders or receivers, always interacting with a computer. As we explain below, this design allows us to provide a strong test of limited strategic thinking abilities, and to rule out social preferences in the biases we document.

While canonical models suggest that full unraveling should happen irrespective of the nature of the probability distribution and the granularity of quality information, our findings demonstrate that both dimensions of quality affect the senders and receivers' behaviors and beliefs. Our study identifies two novel biases associated with players' limited strategic thinking abilities. First, we find that the probability distribution of quality significantly affects disclosure trend by senders and, in absence of disclosure, the guess by receivers. For example, players are more likely to focus on the mean if the quality distribution is normally distributed. On the other hand, if quality follows a uniform or a

fat-tail distribution, the bias-to-the-mean is much less pronounced. Second, while a large portion of senders believe that receivers will guess the mean in case of no disclosure, full disclosure happens at the maximum value. In addition, similar to Jin et al. (2021), we observe that an average of 15% of participants never disclose their quality score when they are endowed with the highest level of quality. This behavior can only be rationalized by confusion with the game's rules or random choice. Receivers' beliefs are consistent with such senders' actions and vice versa.

Our results have implications for several real-world markets, such as health care, finance, and housing, prone to deep asymmetric information problems. In this study, we purposefully framed our experimental game as a housing rental market, given that there are several dimensions of rental properties that may be either unobservable or costly to access for potential renters, and this could lead to large market inefficiencies. Environmental disamenities such as climate risks and exposure to pollution, for instance, are dimensions of quality that are typically not disclosed to renters, but that are rapidly having first-order welfare implications for housing markets. Each of those disamenities follows different types of distributions and could be communicated with different degrees of detail, which, as we show, affect the effectiveness of voluntary disclosure. For firms, the design of voluntary disclosure schemes with complex underlying information has also become a highly debated topic. In Europe, for instance, new regulations for the disclosure of ESG-related information have been put in place but with very broad guidelines. Translating these guidelines into effective schemes requires an understanding of how individuals can process such information. Our study aims to contribute to such understanding and identify biases that can be accounted for in policy design.

The remainder of this article is organized as follows. Next, we provide an overview of the literature. Section 1.3 describes the game's structure, main hypotheses, and experimental design. Section 1.5 presents the computation of expected payoffs and total welfare. Last, we discuss the empirical results and highlight the key findings.

1.2 Related Literature

Our study relates to two strands of literature in experimental economics. Studies investigating the unraveling of information in games of imperfect information and studies carrying behavioral analysis with elicited beliefs. We briefly discuss each of these strands, and highlight our contributions.

1.2.1 Experimental Evidence of Unraveling

Although strategic disclosure in games of imperfect information has been extensively studied from a theoretical perspective for over forty years (Grossman, 1981; Jovanovic,

1982), empirical evidence assessing information unraveling and the barriers to it remains limited. Forsythe et al. (1989) is one of the earliest experimental studies investigating disclosure behavior in the presence of adverse selection. They find that players' observed strategies are consistent with the game's predicted optimal sequential equilibrium. In particular, buyers tend to assume the worst quality in the absence of disclosure, which incentivizes sellers to disclose more frequently, ultimately improving market efficiency.

One important barrier to full information unraveling with voluntary disclosure is the presence of disclosure costs. When there exist such costs, the rate of unraveling depends on these costs and the distribution of signals. Benndorf et al. (2015) test the role of costly disclosure in a context where workers have to disclose their productivity and find evidence that these costs can be an important barrier to unraveling. Their results, however, suggest that additional mechanisms are at play. They also find that players fail to make the optimal equilibrium choices suggesting that limited strategic thinking abilities are also at play. Jin et al. (2021) conduct a laboratory experiment to directly test this hypothesis. They use a revelation game and incorporate the elicitation of beliefs from both senders and receivers to uncover players' strategic abilities. Their findings largely corroborate theoretical predictions, demonstrating almost complete disclosure for the most favorable values. However, deviations from the theory occur as some senders refrain from disclosing intermediate values. This concealment proves advantageous for senders, as receivers tend to overestimate on average. Nonetheless, these instances of non-disclosure are consistently explained by the elicited beliefs, which are, on average, accurate. Consistent with the findings of Jin et al. (2021), Montero and Sheth (2021) find that complete unraveling does not occur, both in context-specific treatments and in conditions allowing communication. They conclude that the failure of complete unraveling can be attributed to receivers' naivety and insufficient skepticism toward concealed information, leading to over-guessing.

Like us, Li and Schipper (2020) experimentally test how the coarseness of information impacts unraveling. Their experimental setting is, however, different and considers a persuasion game with verifiable but potentially vague information. Their findings are closely aligned with our hypotheses. Utilizing two treatment scenarios (e.g., a coarser and a more granular one), they demonstrate that unraveling is more frequent in a game with coarser quality levels than in games with four possible levels. This supports our prior that a coarser framework is associated with higher unraveling, as it eases players' burden to think strategically. Additionally, Li and Schipper (2020) also find a correlation between cognitive abilities and strategic reasoning.

Our paper contributes to this literature by providing experimental evidence that the nature of quality distribution can have a first-order impact on unraveling players' limited strategic ability. We adapt the one-shot revelation game and use within-subject and between-subject variations to isolate how different features of the quality distribution impact players' strategies. To our knowledge, this is the first experiment that tests

unraveling behavior under different common-knowledge probability distributions.

1.2.2 Beliefs and Behavioral Type Analysis

A central focus of our research is to investigate the role of bounded rationality. One of the most adopted behavioral models to explain strategic thinking is the level-k classification, first introduced by Stahl and Wilson (1995) and Nagel (1995). There is a large literature in experimental economics measuring level-k rationality, including in the context of revelation and communication games. Cai and Wang (2006) apply both a Quantal Response Equilibrium and the model proposed by Crawford (2003) and show that the level-k approach enables 75% of senders and receivers to be classified into distinct k-level types, revealing that lower levels of sophistication are associated with overcommunication. Similarly, Benndorf et al. (2015) show that experimental data align well with level-k predictions, indicating that the failure of complete unraveling in decision-making is driven by agents who face difficulties due to the complexity of level-k reasoning.

We elicit beliefs through a belief-elicitation procedure. We follow Sah and Read (2017) and Jin et al. (2021), who ask participants about their beliefs at the end of the experiment without offering any economic incentives. Some studies (Kugler et al., 2007; Montero and Sheth, 2021) provide monetary rewards for accurate predictions, which allows aligning participants' earnings with the accuracy of their expectations versus actual payoffs. Participants in both experiments report their beliefs about others' strategies.

1.3 Experimental Design

Our experimental design uses a version of the revelation game proposed by Dickhaut et al. (1995) with an implementation that follows Cai and Wang (2006) and Jin et al. (2021). First, we summarize the game's structure, and then we discuss the experimental scenarios and treatment variations. Next, we lay out the research hypotheses.

1.3.1 The Revelation Game

We design two experiments based on several rounds of the following one-shot revelation game. The game includes two players: a sender and a receiver. At the beginning of the game, Nature determines a state b from a probability distribution F over a finite set $B \subset \mathbb{R}$, where $b \in B$, which means that b can take any of the discrete values representing different levels of quality contained in B . Each sender is assigned a quality level b by Nature and can decide either to disclose or conceal the information to the receiver. Note that disclosure must be truthful and both players know that if the sender shows that quality value, this is the correct one. The game is sequential: the sender makes the first

move, and the receiver follows. This structure ensures that the receiver makes decisions only after the sender’s action. This is a typical game characterized by asymmetric and imperfect information. Before the game starts, only the sender knows the true state b , as it is determined by Nature. Although the receiver is aware of sender’s possible strategies, the receiver only obtains precise information about b if the sender chooses to disclose it. In addition, ex-ante both players are informed of the range of possible quality values and the probability distribution, together with a payoff table whose computations are described below.

The set of actions available to the sender is $M(b) = \{b, null\}$. He/she can either disclose the true state b or choose not to disclose any information ($null$). If the sender discloses the information, the receiver simply reads it on the screen and earns the respective payoff. If the sender does not disclose the quality, the receiver is asked to guess it taking into account the information he/she has received before. The set of possible guesses receiver can make is denoted G and depends on the true distribution of the set of possible states $b \in B$, with $B \subset G$. G contains also the half values to allow for computations¹.

The sender’s utility increases with the receiver’s action and does not depend on the quality value he/she is assigned. The receiver’s utility is concave in the receiver’s action and reaches the peak when the guess corresponds to the true value. Payoffs are computed and accumulated each round, depending on both players’ actions. The general payoff’s structure is taken from Jin et al. (2021) and Cai and Wang (2006)². The payoff function of the sender is $EP_S(b, \hat{g}) = 110 - p|\hat{b} - g|^{1.4}$, where \hat{b} corresponds to the max b with $b \in B$; p represents the penalty factor, that we decided to adapt to each within-scenario in order to preserve similar distance between payoffs related to similar actions in all the three settings³. The payoff of the sender is independent of the state b and strictly increasing with g . On the other hand, the payoff of the receiver is $EP_R(b, g) = 110 - p|b - g|^{1.4}$. It increases when the distance among the true value b and the current guess g decreases.

Payoffs shown to participants are collected in Figures A.4, A.3, A.2 in Appendix.

1.3.2 Implementation

The experimental design includes two parallel online sessions, each one tailored to investigate behaviors and beliefs of senders and receivers separately. The first session involves participants acting exclusively as senders (Sender Experiment), while the second session assigns participants to the role of receivers (Receiver Experiment). In both

¹ $B \subset G$, since G contains also half values. In particular $G=1,1.5,2,2.5,3,3.5,4,4.5,5$ in TA, $G=1,1.5,2,2.5,3$ in TB and $G=1,1.5,2$ in TC

²Since their function allowed for negative payoffs in a round, we adopted small changes to prevent distortions related to avoiding actions.

³ p assumes value 15 in the TA variation, 30 in the TB and 60 in the TC

experiments, participants are instructed to play against a computer, simulated to behave as a rational human⁴. The experiments have been constructed using Qualtrics and disseminated online via the Prolific platform, targeting an average American demographic without specific selection criteria. Before launching the experiment, we received the ethical approval by the University of Lausanne committee and the approval of the pre-registration plan, together with the pre-analysis plan, by the AEA RCT Registry.

Figure 2.1 illustrates the flow of the questionnaires⁵. The figure shows clearly the two randomization processes (within- and between-subject), which have been consistently applied across both experiments. At the beginning of the experiment each participant is randomly assigned to one out of three potential groups (between-subject randomization) and remains there until to the end of the study. Every group is assigned to a different probability distribution of b states (e.g., uniform, bell-shaped and fat-tailed distributions). Participants engage in multiple rounds of the game within each probabilistic scenario, facing decisions under varying levels of coarseness (e.g. they face decisions when quality values can take two, three and five different quality scores). All players experience to the three coarseness frameworks, with a random order. Also the sequence of play within each scenario is randomized (within-subject randomization). At the end of the game, participants in both studies are asked to answer to four logical open-ended questions corresponding to the Cognitive Reflection Test CRT-2 by Thomson and Oppenheimer (2016), in a random order.

Sender Experiment

The sender experiment is structured as follows.

Participants: there are two players competing in a one shot game. The sender, who is assigned by Nature a feasible state $b \in B$, can take one of the two possible actions: *disclose* and *not disclose* the state b . In this experimental framework the receiver is simulated by a computer programmed to act rationally.

Treatment variations: the experimental framework is built along two treatments directions, ***between-subject*** and ***within-subject*** variations. The *between-subject* treatment variation is adopted to control for potential behavioral distortions due to the knowledge of a specific distribution of quality values b . Along this line we create three independent cohorts. The first one, which represents the *control* unit (C), is told that b follows a **uniform distribution**, which means that all possible quality values have the same probability to manifest. The second group that represents the *first treatment* ($T1$), is acknowledged of a **bell-shaped distribution** of b values, with extreme quality values attached to a lower

⁴We clearly define rationality as maximizing own payoffs.

⁵The full set of experimental instructions is available in the following online material: <https://drive.google.com/file/d/1tHC1qZPT61wr0cdI59AvaazTHCbF0WfL/view?usp=sharing>.

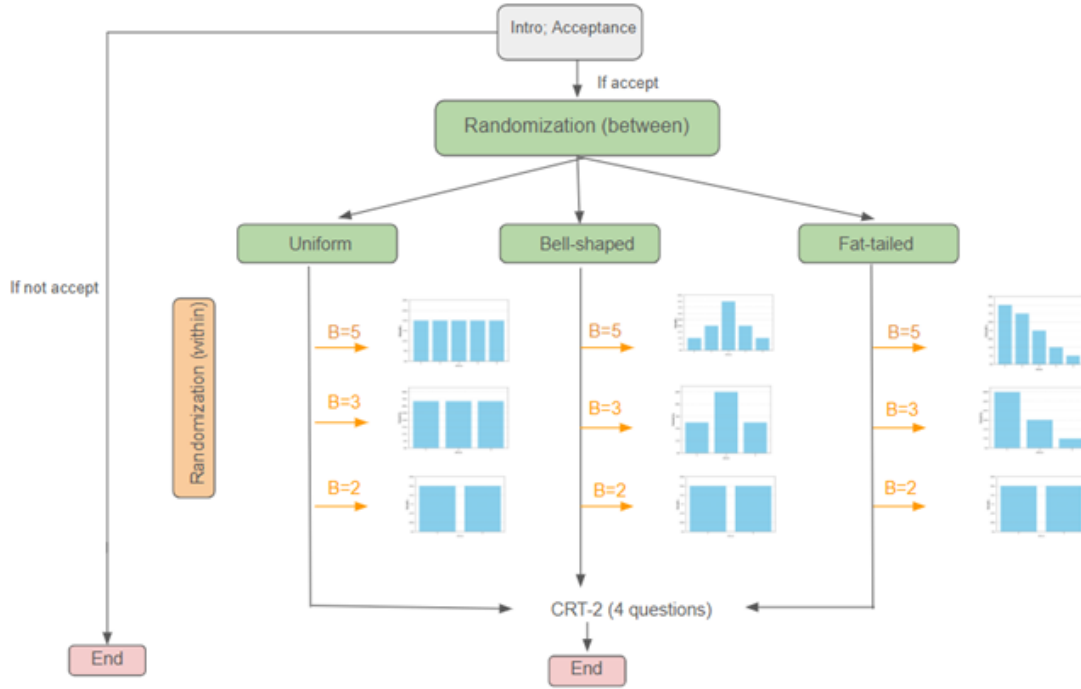


Figure 1.1: Randomization flow

probability than intermediate ones. The last group, labeled as *second treatment* ($T2$), is associated with the knowledge of a **fat-tailed distribution** of b , with smaller quality values holding higher probability than larger ones.

The second scenario presents a *within-subject* treatment variation. The objective of this scenario is to check, independently of the distribution, whether different complexity affects unraveling. We change the complexity of the game by manipulating the size of B , the number of possible qualities. We identify three within-subject scenarios. In the most complex version of the game, participants are told that $b \in \mathbf{B} = \mathbf{1,2,3,4,5}$, which means that the house can be assigned to one out of five quality levels, with higher numbers representing higher quality. This treatment is labeled as *first treatment* (TA); the second scenario includes three potential quality levels from 1 to 3, formally with $b \in \mathbf{B} = \mathbf{1,2,3}$. They are asked to play 3 rounds under these conditions. This represents the *second treatment* framework (TB) and corresponds to the average level of complexity we are going to test; in the third scenario they are told that $b \in \mathbf{B} = \mathbf{1,2}$ and asked to play 2 rounds. This is identified as a *control group* (TC) since it is a very simple version of the game which almost excludes strategic reasoning and which we expect to produce unraveling percentages close to full unraveling. In addition, in the case with $B = 1,2$ it is meaningless differentiating between several probability distributions and we depicted a uniform probability in all the treatment variations. To summarize, each individual is subjected to two randomization processes. At the beginning of the experiment he/she is attached to a specific probability distribution group; then the individual plays repeated rounds of the game across the three within-subject variations allowing for different degree

of complexity, whose order is randomized⁶.

By considering treatment interactions, we end up with nine treatment scenarios⁷: **C.A** which receives the information of *uniform distribution*, with b assuming one out of five values, formally $b \in B = 1,2,3,4,5$; **C.B** which receives the information of *uniform distribution*, with b assuming one out of three values, formally $b \in B = 1,2,3$; **C.C** which receives the information of *uniform distribution*, with b assuming one out of two values, formally $b \in B = 1,2$; **T1.A** which receives the information of *bell-shaped distribution*, with b assuming one out of five values, formally $b \in B = 1,2,3,4,5$; **T1.B** which receives the information of *bell-shaped distribution*, with b assuming one out of three values, formally $b \in B = 1,2,3$; **T1.C** which receives the information, with b assuming one out of two values, formally $b \in B = 1,2$ and the distribution follows a *uniform distribution, given that with only two possible outcomes we cannot generate a bell-shaped distribution*; **T2.A** which receives the information of *fat-tailed distribution*, with b assuming one out of five values, formally $b \in B = 1,2,3,4,5$; **T2.B** which receives the information of *fat-tailed distribution*, with b assuming one out of three values, formally $b \in B = 1,2,3$; **T2.C** which receives the information of *fat-tailed distribution*, with b assuming one out of two values, formally $b \in B = 1,2$.

Experimental Framework: After the initial randomization, each participant is asked to imagine being a property owner who wants to rent his house to a potential tenant, that is represented by a computer instructed as a rational agent. The rules of the game and each player possible actions are explained with very simple words. All possible payoffs depending on both players' strategies are collected in a very clear table and showed each round to keep the framework as transparent as possible. Every participant, regardless of the group to which has been assigned, is randomly attributed to one of the first *within-subject* variation, and is asked to play several rounds of the game accordingly⁸. Every time an individual faces a new treatment condition, he/she is carefully informed. In every round, each participant receives a different quality. At the end of the decision making phase in each coarseness scenario, before moving to the following variation, we elicit beliefs: we ask one time per each coarseness to report the belief about receiver's guess in absence of disclosure. We elicit beliefs three times in total.

⁶For example, individual i is previously exposed to the case of $B = 1,2,3$, then to the case $B = 1,2$ and lastly to $B = 1,2,3,4,5$, while individual j is first assigned to $B = 1,2$, then to $B = 1,2,3,4,5$ and to $B = 1,2$

⁷Note that each individual is subjected to only three of them.

⁸This is not a repeated game, in the sense that the same player does not play twice with the same quality. This prevent any form of learning. However, each player plays once with all possible qualities.

Receiver Experiment

The information related to *participants* and *treatment variations* remain unchanged, except that in this experiment each participant is asked to imagine being a potential tenant looking for a house, and is informed to play against a computer instructed as a rational homeowner. *Experimental Design*: at the beginning of the experiment each participant is randomly assigned to one of the *between-subject* group, which is either $T1, T2$ or C , and stays there until the end of the experiment. The information provided about the distribution type are exactly the same. At this point, each participant is introduced to the role of potential tenant. Again, the rules of the game, the set of actions and respective payoffs are described. The two randomization flows work exactly as described in the previous paragraph. At the beginning of each within-subject scenario, individuals are asked to play one round of the game in which the sender does not disclose the quality, and therefore they have to communicate their guess. We ask this question once for each within-subject setting with different coarseness levels, for a total of three times. They can select a value from 1 to the maximum possible quality value of the specific within-variation, including half values. Next, we elicit beliefs for each possible value. We first ask for first order beliefs: for each quality we ask if they expect the sender to reveal or not the value. Then we ask second order beliefs: for each quality we also ask which is the sender's belief they expect about the receiver guess if they do not disclose the value. Again, the order of questions is totally randomized.

1.3.3 Research Hypotheses

This experimental study aims to test several research hypotheses. With our treatment variations design, we focus on two main distortions of rational decision making: biases introduced by information on the probability distribution of quality, and distortions generated by the structural coarseness of the game. Within this framework, we attempt to measure how real world decision making deviates from the theoretical predictions of Milgrom (1981).

First, we expect that, as the level of coarseness increases (e. g., when the number of alternatives decreases from five to two), the rate of disclosure by senders will rise. A greater number of potential qualities foster strategic thinking for senders, who may benefit from concealing. Conversely, when the number of alternatives decreases, there is less room for strategic behavior. Similarly, on the receiver side, we expect that a coarser framework increases the accuracy of guesses.

Second, we vary the probability distributions of quality values to investigate whether different distributions significantly alter agents' reasoning. This is crucial in several real-world applications. Information on the general condition of housing in the area, as well as the exposure to environmental risks could alter expectations. For this reason, we expect

that the type of distribution (uniform, bell-shaped, fat-tailed) significantly affects players' strategic decision making under bounded rationality.

Specifically, we hypothesize that the fat-tailed distribution leads senders to disclose lower values more often than in the other two treatments. In parallel, the bell-shaped distribution is expected to induce senders to believe that receivers anticipate values clustered around the mean, making them more inclined to withhold information when assigned very low values perceived as unlikely. The higher probability of low values in the fat-tailed distribution increases the likelihood that senders disclose these values, since they expect receivers to guess lower values more frequently than higher ones.

Finally, consistently with Milgrom (1981); Crawford and Sobel (1982); Jin et al. (2021), we expect a greater tendency to disclose higher (i.e., more favorable) values than lower ones, given the specific payoff structure of the game. This is linked to a very low level of full unraveling, consistent with theory.

1.4 Results

1.4.1 Descriptive Statistics

Sender Experiment was conducted in June 2024. We gathered a total of 524 responses. We exclude responses that took less than two minutes to be completed, considering shorter answers too rapid to reflect thoughtful participation. The Sender Experiment's final dataset includes 519 participants, with 219 males, 291 females, and 9 who opted not to disclose their gender. The age of participants ranges from 18 to 80, with an average of 40.5 years.

Receiver Experiment was launched in July 2024, and the final sample includes 502 participants, with 244 males and 255 females, and an average age of 46 years. All participants are based in the US. Participants accumulate tokens during the game, which are converted to dollars at the end of the experiment. The exchange rate is set at 1 cent per 3.5 tokens for Sender Experiment and 1 cent per 1.65 tokens for Receiver Experiment. Each survey is designed to take a maximum of 15 minutes, with a guaranteed base pay of \$1.50. They also have the opportunity to earn additional compensation based on their performance, up to a maximum of 702 tokens in Sender Experiment and 330 in Receiver Experiment, equivalent to an additional \$2.00 in both cases. The average completion time for the survey is 8.58 minutes in Sender Experiment and 12.65 minutes in Receiver Experiment, with average extra earnings of \$1.65 and \$1.53, respectively.

Table 1.1: Descriptive Statistics

Variable	Sender Experiment			Receiver Experiment		
	Mean	Min	Max	Mean	Min	Max
Age	40.5	18	80	46	18	80
Duration (min)	8.58	2.06	58.12	12.65	2.47	48.18
Score (tokens)	576.55	223	702	239.99	87	330
Cents	165	64	200	153	53	200
Males	219	-	-	244	-	-
Females	291	-	-	255	-	-
Total	519	-	-	502	-	-

1.4.2 Sender Disclosures

We start by analyzing the reporting trends of the senders. Table A.1 in Appendix presents a summary of aggregated reporting rates at the level of the treatment groups (within- and between-subject), and Table A.2 provides more granular information differentiating for quality scores. The variable *Disclose* is a dummy that takes value 1 if quality score is disclosed and 0 otherwise. We first observe that senders reveal more favorable values than unfavorable ones. Surprisingly, we note that across the three within-group variations (when the level of coarseness changes) we do not record full unraveling for all players neither when the quality score is the maximum. In fact, between 13 and 18% of the participants chose not to disclose the highest value, a countervailing trend with respect to rational agents aiming at maximizing their payoffs, which can be explained either as misinterpretation of the rules, general confusion, or limitations in the use of rational decision-making frameworks (Jin et al., 2021). When the quality is the median (3 in the case with five potential quality scores and 2 in the case with three qualities), the reporting rates are 58% and 59%. For the most unfavorable quality scores, the rates drop to less than 26%. In the coarser framework (with two potential quality scores), the reporting rate is 82% for quality equal to 2, and 27% for quality equal to 1.

Figures 1.2a,1.2b,1.2c illustrates reporting trends across groups for different coarseness. The dashed lines represent unraveling prediction, with 100 percent probability of reporting the quality for all quality scores larger than the minimum and indifference between disclosing and hiding the minimum quality⁹. These plots help understand how the reporting trends in each scenario deviate from theory. We can easily note that full unraveling is not observed in any scenario. Thus we try to measure the fit between observed disclosure rates and the full-unraveling benchmark by computing the Mean Squared Error (MSE). Results show that recorded reporting trends are closer to the expected full unraveling predictions when there are only two possible qualities, which is the coarser scenario (MSE equal to 0.043). The divergence grows as the range of possible qualities widens (MSE equal to 0.087 and

⁹In the graphs we assigned a probability of 50 percent for quality 1, but note that it can be any point on the y-axis.

0.168 with three and five qualities)¹⁰. In addition, as we can observe from Figure 1.2a, the shift to 100 percent reporting is delayed with respect to theory: while full unraveling expects full disclosure from quality 2, the line gets steeper at quality 3 and reach the highest probability of disclosure at quality 5. A clear pattern emerges: as the span of potential qualities increases from 2 to 5, the average disclosing rates diverge more from the theoretical predictions. Different probability distributions also plays a crucial role in strategic behavior, as results highlight differences in average reporting between control and treatment groups.

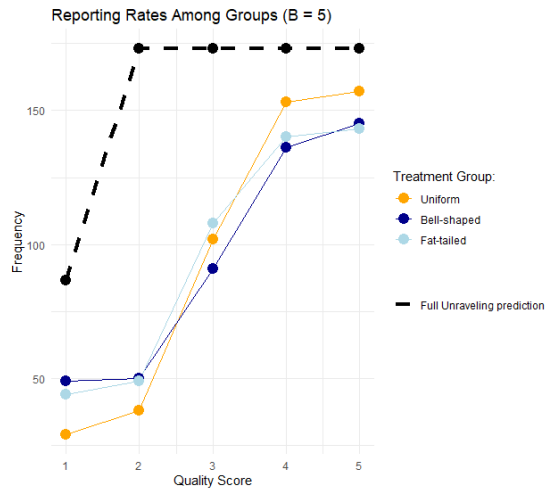
Individuals in the control group (uniform probability distribution) on average are more likely to hide unfavorable qualities and disclose more frequently favorable ones, compared to both treatment groups (bell shaped and fat-tailed distribution). Indeed, the median quality value acts as a pivotal point where the disclosure trends among different groups reverse.

Table A.3 in Appendix summarizes the χ^2 tests results across treatment groups. In the coarsest scenario (only two possible qualities) we observe a weak difference between the control and the group assigned to bell-shaped probability distribution (p-value = 0.0982). When the number of alternatives increases from two to three, we observe notable differences between the control and the treatment groups for quality equal to 3 (with p-values 0.0009 and 0.0588 for bell-shaped and fat-tailed distributions respectively) and for median quality 2 (with p-values 0.0757 and 0.0873 for bell-shaped and fat-tailed distributions respectively). Last, when the span of potential qualities increases to five, significant differences emerge between uniform and fat-tailed groups when quality is 5 (p-value 0.0433), between uniform and bell-shaped (p-value 0.0489), and uniform and fat-tailed when quality is 4 (p-value 0.0778); there is also a significant difference between control and treatments for quality 1 (with p-values 0.0113 and 0.0652 for bell-shaped and fat-tailed distributions respectively) and for quality 4 (with p-values 0.0489 and 0.0778 for bell-shaped and fat-tailed distributions respectively); for quality 5, the difference is significant only between the control and the fat-tailed distribution group (p-value 0.0433).


1.4.3 Receiver Guesses

Moving to Receiver Experiment, we propose a similar analysis to control for potential distortions introduced by the experimental framework on receiver guesses. On receiver's side, the theoretical benchmark of full unraveling expects all guesses be equal to the minimum quality, given that the absence of disclosure by the sender is a signal of poor quality, as to say *no news, bad news* (Milgrom, 1981). Also in Receiver experiment, observed behavior is far from theoretical predictions. Again, the coarsest scenario shows

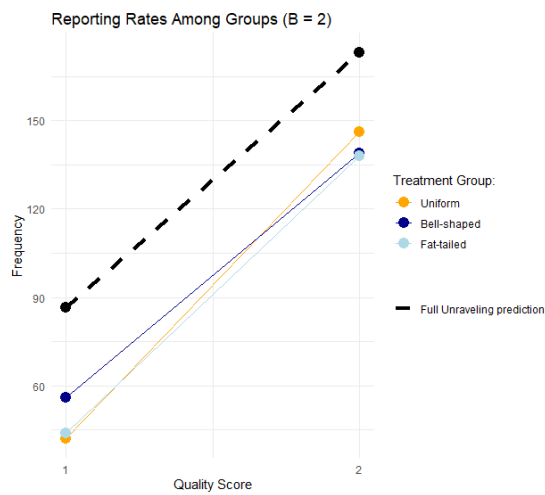
¹⁰See Appendix A.17 for an overview of the method used.



(a) Disclosure $B=5$



(b) Disclosure $B=3$



(c) Disclosure $B=2$

Figure 1.2: Plots of disclosures at different granularity levels across treatment groups.

the highest frequency for guesses equal to 1 (compatible with tested theory) compared to the others.

Table A.4 in Appendix presents the average guesses. The average guess is 2.83 with five qualities and 1.92 and 1.44 with three and two qualities respectively. Some distortions are introduced by providing information on a specific probability distribution. Figure 1.3 graphically represents the average guesses. At first glance, we note that the frequency distribution of receiver guesses across groups mimics the probability distribution to which each group is assigned. The dark blue dots mimics a bell shaped distribution, peaking at the mean and median values. Similarly, the light blue dots are close to the shape of the fat-tailed distribution we provide to respondent in the instructions, with higher frequency attached to smaller qualities. Last, the yellow dots also peak at the mean value, with frequencies diminishing as moving to extreme qualities. These trends are evident in scenario with five qualities. We adopt non parametric tests to compare average guesses between groups¹¹. Significant differences exist between control and fat-tailed group when we have five potential qualities (p-values equal to $2.2e - 06$ and $1.6e - 06$) and similarly with three potential qualities (p-values $5.4e - 05$ and $4.3e - 06$). The same differences between groups are observed with three potential qualities (p-values $5.4e - 05$ and $4.3e - 06$). No differences between groups emerge with two qualities. Detailed results of the tests are collected in Table A.5.

We can conclude that individuals are susceptible to probability distribution information and it plays a crucial role in strategic behavior.

1.4.4 Analysis of Elicited Beliefs

In this section we report the results of the belief elicitation analysis. In particular, we elicit first- and second-order beliefs to probe participants' strategic behavior, assess the accuracy of their expectations about opponents' strategies across nine treatment conditions, and examine the extent to which their actions align with those beliefs.

Sender's Beliefs

After playing all the rounds for each coarseness scenario, the sender is asked to report the guess of the receiver in case of non-disclosure. Average beliefs are summarized in Table A.6. In scenario with five possible qualities the average belief is 2.77. For the uniform and bell-shaped distribution groups the average belief is around 2.90, and for fat-tailed group is 2.47. These differences are statistically significant (p-values equal to $4.1e - 07$ between uniform and fat-tailed groups and $2.5e - 07$ between bell-shaped and fat-tailed

¹¹We used the Kruskal-Wallis test to determine if there were any differences in the medians of the three groups simultaneously. If the test indicated significant differences, we proceeded with post-hoc analyses with Mann-Whitney U tests including Bonferroni correction, conducting pairwise comparisons across groups


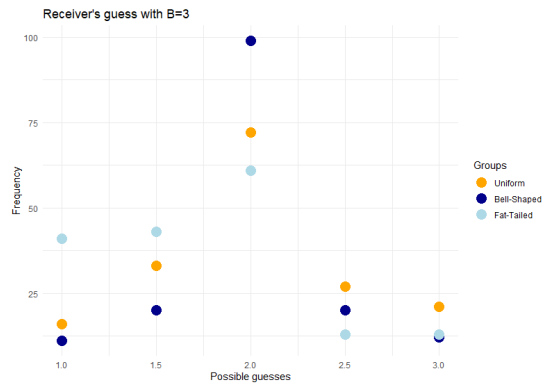
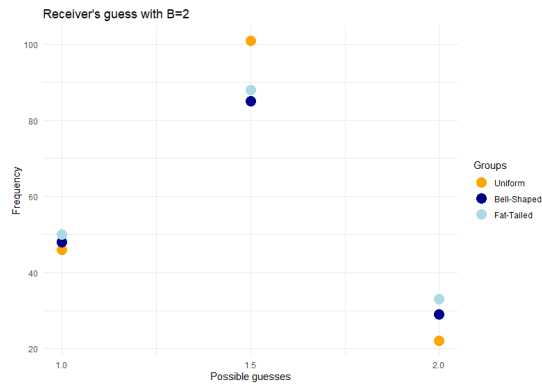
(a) Guess $B=5$ (b) Guess $B=3$ (c) Guess $B=2$

Figure 1.3: Receiver Experiment: plots of receiver's guess at different granularity levels across treatment groups.

groups)¹². Moving to the case with only three qualities, the average guess is 1.88. We still observe the same differences across groups. The average guesses are 1.88 for uniform and 2.0 for bell-shaped arms, compared to 1.69 for fat-tailed (p-values equal to $5.7e - 06$ and $9.6e - 09$). In the coarsest scenario (with only two qualities), the average guess is 1.46.

On average, sender's beliefs on receiver's guesses are correct, as emerges from Figure

¹²Non parametric Kruskal-Wallis and post-hoc Mann-Whitney U test with Bonferroni correction results are collected in A.7

A.5 in Appendix. It is easy to note that Figures (a), (c), (e), that plot elicited beliefs of sender, nearly completely overlap figures (b), (d), (f), that represent observed receiver’s guesses.

Table A.8 offers an overview of how sender’s observed disclosure align with sender’s beliefs, given that we elicited beliefs at the end of all the rounds in each scenario. First we count how many respondent report a specific belief, then we count how many respondent disclose or not each quality. Participants whose belief concerning receiver’s guess in absence of disclosure is 1, are those coherent with full unraveling. According to theory if the expectation of the guess is the minimum quality, the sender should always disclose qualities higher than the minimum (i.e. higher or equal to 2). The column labeled as Nd indicates the number of non-disclosure for the quality identified as b . Surprisingly we find that when there are 5 qualities, the 65% of them does not report quality 2, the 35% does not report quality 3 and the 14% does not report neither quality 4 nor quality 5. Similarly, with three qualities the 32% and the 16% do not report qualities 2 and 3 respectively. When quality is 2, the 14% does not disclose the quality equal to 2.

Receivers’ Beliefs

In Receiver Experiment we elicit first and second order beliefs. Given a certain quality, we first ask receiver what he/she expects the sender’s action will be: report or not the quality. Figure A.6 compares receiver’s first order belief with observed sender’s action. Even in this case, we clearly observe that expectations closely mirrors the actual reporting by senders in Sender Experiment, confirming the match between expectations and real outcomes without any type of interaction.

In addition, the receiver is asked to report his guess concerning the belief of the sender about receiver’s belief when he/she decides to hide the quality. We identify these as second order beliefs. Table A.10 summarizes average second order beliefs across groups, which are graphically presented in Figure A.7. On average we observe that these beliefs gradually increase with the quality score and respondents in fat-tailed group nearly always report beliefs significantly lower than those in other groups in all scenarios, except when qualities are two (all results of statistical tests are in Table A.11).

1.4.5 Behavioral Analysis

Considering prior findings by Jin et al. (2021), we expected ex-ante that full unraveling is pretty unlikely to occur in reality. To get further insights into strategic decision making, we identify behavioral types of senders and receivers observing actions and beliefs. We aim to recognize individual profiles reflecting defined strategies to measure the extent of full unraveling, against naivety and confusion. Previous studies have explored distortions in bounded rationality by adopting behavioral models to measure naivety, sophistication,

confusion, and social preferences¹³. Given that our participants are told to play with a computer, we exclude in advance distortions due to other-regarding preferences¹⁴ (Ferraro et al., 2003), reputation effect and any other social signal. Although multiple studies have shown that prosocial behavior might occur in nonsocial games when participants interact with computers, those are typically attributed to confusion (Ferraro et al., 2003; Burton-Chellew et al., 2016). We do not use level-k type analysis as Stahl and Wilson (1995); Nagel (1995); Crawford (2003); Cai and Wang (2006). Indeed, we propose the following type classification to understand at which quality senders shift from hiding to reporting and to check whether elicited beliefs can help understand how they mediate between hide and reveal. Understanding the cutoff value tells about strategic forces that diverge from full unraveling.

Table 1.2: Sender and Receiver Experiment: behavioral type analysis

Panel A: Sender	B=5		B=3		B=2		Pooled	
	Count	%	Count	%	Count	%	Count	%
Full Unravel	84	16.18%	275	52.99%	423	81.50%	63	12.14%
FU with rational beliefs	12	2.31%	47	9.06%	95	18.30%	6	1.16%
Bias to the Mean	166	31.98%	203	39.11%	–	–	91	17.53%
Others	269	51.83%	29	5.59%	96	18.50%	365	70.33%
Panel B: Receiver	B=5		B=3		B=2		Pooled	
	Count	%	Count	%	Count	%	Count	%
$g = 1$	29	5.78%	68	13.55%	144	28.69%	11	2.19%
$g = 1$ and belief FU	4	0.08%	29	5.78%	124	24.70%	0	0.00%
$1 < g < g_m$	183	36.45%	96	19.12%	–	–	56	11.16%
$g = g_m$	147	29.28%	232	46.22%	274	54.58%	64	12.75%
$g > g_m$	143	28.49%	106	21.12%	84	16.73%	17	3.39%

Table 1.2 presents a summary of our type classification for senders (Panel A) and receivers (Panel B). A sender type s is identified with the disclosure rule $D_s : B \rightarrow \{0,1\}$, with $D_s(b) = 1$ if disclose at b and $D_s(b) = 0$ otherwise; receiver type r , is one of the following $r \in \mathcal{R} = \{g = 1, 1 < g < g_m, g = g_m, g > g_m\}$ ¹⁵. We identify four sender types. (i) *Full Unravel (FU)* comprises individuals who always disclose for qualities above the minimum and either disclose or not at the minimum. We report the share of each type at different levels of signal granularity (columns 1 to 3). The column 'Pooled' aggregates behavior across all rounds and all granularity levels and provides an overall classification of each individual's disclosure strategy. The FU rate rises as information becomes coarser: 16.18% with five qualities, 52.99% with three, and 81.50% with two. This pattern is consistent with prior findings. Considering beliefs, we isolate a subsample of full unravelers labeled *FU with rational beliefs*, those who always disclose above the minimum and state

¹³See section 1.2.2

¹⁴Other-regarding preferences refer to situations where individuals care not only about their own outcomes but also about the payoffs or well-being of others.

¹⁵Where g_m indicates the mean

that the receiver guesses the minimum under non-disclosure. This group is small, but its frequency increases with coarser information (from 2.31% with five qualities to 18.30% with two). (ii) We define *Bias to the mean (BTM)* senders who disclose values at and above the mean and conceal lower values. This profile is common, representing 32% with five qualities, 39% in the intermediate case, and 18% in the pooled scenario. It cannot be separately identified in the coarsest case. The remaining profiles are grouped as (iii) *Others*. We attempted to distinguish degrees of naïveté, but with few quality levels this was not informative. This category therefore includes all patterns that are neither *FU* nor *BTM*, and likely reflects additional naïveté, confusion, or random behavior.

Panel B classifies receiver types by their guesses at different levels of granularity. The likelihood of (i) $g = 1$ increases with coarseness: only 5.78% guess 1 under non-disclosure with five qualities, rising to 13.55% with three and 28.69% with two. Similarly, $g = 1$ and *belief FU* identifies respondents who always guess the minimum and expect full unraveling (senders disclose all qualities above the minimum to which they are indifferent). We then count receivers who guess between the minimum and the mean (ii) $1 < g < g_m$, exactly at the mean (iii) $g = g_m$, and above the mean (iv) $g > g_m$. With five qualities there is no pronounced difference in the frequency of these three categories, whereas with coarser scenarios more participants tend to select the mean. Full unraveling also increases as the granularity of information decreases. This suggests that coarser environments are not only more likely to produce higher rates of unraveling, but also more likely to reach complete unraveling.

1.4.6 CRT-2 Score

At the end of the experiment, respondents are asked to complete four open ended questions which serve as a check for individual cognitive ability. We decide to employ the *CRT-2 test* proposed by Thomson and Oppenheimer (2016), a variation of the most common CRT by Frederick (2005), to which subjects affiliated to recruiting companies might have been exposed several times. We build the *CRT-2 score* as a four-items aggregate variable. Questions and outputs of the test are collected in Table 1.3. Each question represents a different item. For each question, we assign 1 if the answer is correct and 0 otherwise, and we compute the average to get the total score. *CRT-2 score* ranges from 0 (e.g. when none of the answers is correct) to 1 (e.g. when all the answers are right). The average *CRT-2 score* is 0.6127 in Sender Experiment and 0.7634 in Receiver Experiment.

1.4.7 Regression Analysis

In this section, we present the outcomes of additional analyses that account for demographic characteristics and cognitive abilities. Table 1.4 reports the results of linear regressions with the *sender's beliefs* as the dependent variable, estimated separately for

Table 1.3: CRT-2 results for Sender and Receiver’s experiment

Question:	Sender		Receiver	
	Correct	% Correct	Correct	% Correct
Q1: <i>If you’re running a race and you pass the person in second place, what place are you in?</i>	347	67.33%	338	65.13%
Q2: <i>A farmer had 15 sheep, and all but 8 died. How many are left?</i>	448	86.32%	402	80.08%
Q3: <i>Emily’s father has three daughters. The first two are named April and May. What is the third daughter’s name?</i>	396	76.30%	373	74.30%
Q4: <i>How many cubic feet of dirt are there in a hole that is 3’ deep, 3’ wide, and 3’ long?</i>	81	15.61%	85	16.93%
CRT-2 score:	Avg score	0.6127	Avg score	0.7634

different levels of granularity. Column (1) aggregates all beliefs across scenarios, while Columns (2)–(4) display the results for each level of granularity. At first sight, we observe a strong negative effect of Treatment 2 (fat-tailed distribution) on average beliefs regarding the receiver’s guess. This effect becomes stronger as granularity increases, moving from two to five quality levels, and persists in the aggregated case shown in Column (1). In addition, we find a negative and significant effect of the CRT-2 score on beliefs, with the magnitude of the effect increasing with granularity.

We further estimate linear regression models using the *average disclosure rate* as the dependent variable, with results reported in Table ?? in the Appendix. These results are mostly insignificant: the only exception is a negative effect of Treatment 1 (bell-shaped distribution) relative to the control in the scenario with three quality levels, although this finding does not allow us to draw additional conclusions.

Finally, we estimate a logit model with the dependent variables emerged from behavioral type analysis. Table A.13 in Appendix show the results. In panel A we use as dependent *full unraveling type*, a dummy equal to 1 if disclosure decisions are consistent with the full unraveling prediction, and 0 otherwise. In this case, we find only a negative effect of Treatment 1 (bell-shaped distribution) in the scenario with three qualities, but not in the other scenarios. Moreover, CRT-2 score significantly reduces the probability of behaving as a full-unraveling type, with significant effects in both the $B = 2$ and $B = 5$ scenarios. In panel B we use *bias to the mean type*, a dummy equal to one when sender discloses based on the expected value computation, as previously explained, and 0 otherwise. These specifications offer some relevant insights. First we note that when the number of alternatives is five or three both treatments reduced the probability of being biased to the mean compared to the uniform group. This is true a also for the general case with lower

significance. In addition, higher score in the CRT-2 test is associated to higher probability of being of this type.

Turning to the Receiver Experiment, we estimate linear models with the receiver's *guess* as the dependent variable. Results are shown in Table A.14. Panel A reports the models including treatments and demographics as covariates, while Panel B extends the specification by adding the receiver's average first- and second-order beliefs. Results are consistent across the two panels. We note that being in Treatment 2 (Fat-Tailed distribution) significantly reduces guesses as granularity increases. No significant effects of CRT-2 score are detected. As expected, in Panel B we observe that higher second-order beliefs are positively associated with higher guesses, which confirms the internal consistency between beliefs and actual behavior, although it does not provide additional substantive insights.

Table 1.4: Sender Experiment: linear regression with belief as outcome variable

	(1) Belief (All)		(2) Belief ($B=5$)		(3) Belief ($B=3$)		(4) Belief ($B=2$)	
	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
Intercept	2.184***	0.086	3.010***	0.161	2.013***	0.099	1.528***	0.057
Bell-shaped (T1)	0.040	0.046	0.023	0.086	0.058	0.053	0.039	0.031
Fat-tailed (T2)	-0.249***	0.046	-0.436***	0.086	-0.250***	0.053	-0.060*	0.030
Male	-0.044	0.039	-0.071	0.072	-0.034	0.045	-0.028	0.026
Other / NA	0.006	0.145	-0.094	0.270	0.019	0.167	0.094	0.096
Age	0.002	0.001	0.002	0.003	0.002	0.002	0.001	0.001
CRT score	-0.205***	0.073	-0.243*	0.137	-0.237***	0.085	-0.136***	0.049
Observations	519		519		519		519	
R ²	0.105		0.078		0.091		0.043	
Residual Std. Error	0.427 (df = 512)		0.796 (df = 512)		0.492 (df = 512)		0.283 (df = 512)	
F Statistic	10.00*** (df = 6; 512)		7.23*** (df = 6; 512)		8.55*** (df = 6; 512)		3.84*** (df = 6; 512)	

Notes: * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$. Reference categories: Control = Uniform, Sex = Female.

1.5 Welfare Analysis

In this section we compute expected payoffs for both players' types and the total welfare of the game. The simulation exercise is performed on the uniform distribution scenario. The objective is to compare empirical distribution of players' types total welfare with theoretical predictions. As presented in Section 1.4.5, we identify three types of sender and four types of receiver. We recall that sender's pure strategies are *FU* (full unravel), *BTM* (bias-to-mean), and *Others*¹⁶. Receiver is one of the following types $r \in \mathcal{R} = \{g = 1, 1 < g < g_m, g = g_m, g > g_m\}$ ¹⁷. To compute expected payoffs we first fix types (s, r) and state b . If $D_s(b) = 1$ (sender discloses a certain b), expected payoffs are¹⁸:

¹⁶See Section 1.4.5.

¹⁷Where g_m indicates the mean.

¹⁸For more details see Section 1.3.1.

$$\begin{aligned}
 EP_S^D(\hat{b}; g|g = b) &= U_S(\hat{b}; g|g = b) = 110 - p|\hat{b} - g|^{1.4}, \\
 EP_R^D(b; g|g = b) &= U_R(b; g|g = b) = 110 - p|b - g|^{1.4}.
 \end{aligned}$$

where \hat{b} is the max $b \in B$, and is conditional to the guess be equal to the true value as the revelation is always truthful. If $D_S(b) = 0$ (sender does not disclose a fixed b), the receiver guesses g according to a certain probability distribution q_r and expected payoffs for the sender and receiver are:

$$EP_S^{ND}(\hat{b}; g) = \sum_g q_r(g) U_S(\hat{b}, g), \quad EP_R^{ND}(b, g) = \sum_g q_r(g) U_R(b, g)$$

In the following analysis we compute the total expected payoff for the sender (EPS) and for the receiver (EPR) for the case $B = 1, \dots, 5$ aggregating over the five states, and Welfare (W) as the sum of expected payoffs of both players:

$$\begin{aligned}
 EPS(\hat{b}; g) &= \sum_{b \in B} \left[D_S(\hat{b}; g|g = b) U_S(b) + (1 - D_S(b)) \sum_g q_R(g) U_S(\hat{b}, g) \right], \\
 EPR(b, g) &= \sum_{b \in B} \left[D_S(b, g|g = b) U_R(b, b) + (1 - D_S(b)) \sum_g q_R(g) U_R(b, g) \right], \\
 W(\hat{b}, b, g) &= EPS_{\hat{b}, g} + EPR_{b, g}.
 \end{aligned}$$

In the end, we compute the Deadweight Loss (DWL) for each player as the deviation between the equilibrium benchmark ($FU, g = 1$) and the welfare obtained for each combination of strategies:

$$DWL(\hat{b}, b, g) \equiv W_{FU, g=1} - W(\hat{b}, b, g)$$

Table 1.5 reports the 3×4 pure strategies payoff matrix for the sender (rows) and the receiver (columns). First we note that if the sender plays FU , the payoffs are constant across all receiver types: the sender obtains 321 and the receiver 550. When the sender reveals the quality, the receiver simply gets the maximum payoff as the disclosure is always correct (with $b = g$). Second, for any receiver type other than $g = 1$, the sender would strictly gain by switching from FU to BTM . The last row *Others* includes all possible combinations of disclosure that do not represent previous types¹⁹. The last row contains the payoffs of the best response (BR) of the sender among the possible combinations of

¹⁹We find 2⁵ possible combinations of disclosure from which we remove FU and BTM profiles. The payoffs for this profile of strategies with respect to each pure strategy of the sender is computed as the mean of all possible combinations given a precise strategy of the receiver (e.g., we compute the payoff for each combination of *Others* given $g = 1$ and we report the average in the table).

disclosure included in *Others* against each possible pure strategy of the receiver²⁰.

BRs confirm what anticipated: *BTM* is the sender's BR against any receiver type except $g = 1$, against which the sender's BR is *FU*. Conversely, the receiver's BR is $g = 1$ against *BTM*, $g = g_m$ against *Others*, and all four receiver types are payoff-equivalent against *FU*. This means that, if the sender deviates to *BTM*, the receiver best-responds with $g = 1$, which lowers the sender's payoff by about 35 tokens relative to *FU* (from 321 to 286) and induces a return to *FU*. In addition, we find a unique pure-strategy Nash equilibrium (*FU*, $g = 1$) as already introduced by Milgrom (1981). While some off-equilibrium profiles (e.g., *BTM* against $g = g_m$ or $g > g_m$) show higher sender payoffs and even higher total welfare, they are not Nash-stable: they push the receiver's best response back to $g = 1$, which in turn restores *FU* as the sender's best reply.

Table 1.5: Payoff matrix ($B = 5$).

Sender \ Receiver	$g = 1$	$1 < g < g_m$	$g = g_m$	$g > g_m$
<i>FU</i>	(321 ; 550)	(321 ; 550)	(321 ; 550)	(321 ; 550)
<i>BTM</i>	(286 ; 535)	(355 ; 531)	(416 ; 495)	(472 ; 424)
<i>Others</i>	(160 ; 424)	(251 ; 471)	(333 ; 493)	(406 ; 465)
<i>Others (BR for each g)</i>	(286 ; 535)	(325 ; 519)	(416 ; 495)	(503 ; 395)

Table 1.6 summarizes total welfare, defined as the sum of sender and receiver payoffs, and the associated deadweight loss (DWL) relative to the benchmark (*FU*, $g = 1$), which is 871. Consistent with the BR logic, welfare exceeds the benchmark for (*BTM*, $g = g_m$) by 41, for (*BTM*, $g > g_m$) by 25 and for (*BTM*, $1 < g < g_m$) by 15; by contrast, (*BTM*, $g = 1$) reports welfare 821, and a DWL of 50. In addition, if we consider the BR among the *Others* for each receiver's type, we note that there is an increment in welfare compared to the equilibrium if the receiver overguesses.

Table 1.6: Welfare e Deadweight Loss per profilo ($B = 5$)

Sender	Welfare				DWL			
	$g = 1$	$1 < g < g_m$	$g = g_m$	$g > g_m$	$g = 1$	$1 < g < g_m$	$g = g_m$	$g > g_m$
<i>FU</i>	871	871	871	871	0	0	0	0
<i>BTM</i>	821	886	912	896	50	-14	-41	-24
<i>Others</i>	585	722	826	872	286	149	45	-1
<i>Others (BR for each g)</i>	821	843	886	898	50	28	-15	-27

Note. DWL computed with respect to the baseline *FU* (871).

Next, we include the empirical distributions of types in the analysis. Based on observed behavior in the experiments, we estimate the type distributions for senders and receivers

²⁰For each combination, BR of the senders are the following, where 1 indicates disclosure and 0 no-disclosure: BR to $g = 1$: 10111; BR to $1 < g < g_m$: 01011; BR to $g = g_m$: 00011; BR to $g > g_m$: 00001

as

$$s_t = \{FU = 0.14, BTM = 0.35, Others = 0.51\},$$

and

$$r_t = \{g = 1 = 0.04, 1 < g < g_m = 0.33, g = g_m = 0.28, g > g_m = 0.35\}.$$

These shares are interpreted as the probabilities that a randomly drawn participant behaves according to each type, and they are used to construct mixed strategies for each side. Table 1.7 reports expected payoffs, total welfare, and DWL for four scenarios under random matching²¹. With the observed type mix (1), total welfare is 815, implying a DWL of about 4% relative to the benchmark ($FU \times g = 1$) in (2). If senders best-respond to the observed receiver mix by switching to *BTM* (3), welfare rises of about 3% above the benchmark. If instead receivers best-respond to the observed sender mix by choosing $g = g_m$ (4), welfare equals 835, which means that a small positive DWL of about 4%. Overall, the equilibrium prediction ($FU, g = 1$) works only with fully rational players; in real world scenarios where many receivers over-guess in the absence of disclosure, senders have incentives to act strategically and move toward *BTM*, which raises the sender’s payoff and, in these simulations, often raises total welfare as well. Figure 1.4 compares total welfare across different scenarios. We conclude that the Nash Equilibrium is informationally efficient, but given empirical evidence of overguessing by receivers, it is not welfare efficient.

Table 1.7: Expected payoffs, welfare and DWL across scenarios ($B = 5$)

Scenario	E_{Sender}	E_{Receiver}	Welfare	DWL	DWL (%)
(1) Obs. S \times Obs. R	342	473	815	56	6
(2) FU S \times $g=1$ R (benchmark)	321	550	871	0	0
(3) BR_s (<i>BTM</i>) to Obs. R	410	483	894	-23	-3
(4) BR_r ($g = g_{\text{mean}}$) to Obs. S	348	487	835	36	4

1.6 Discussion and Conclusion

Our analysis provides critical insights into strategic behavior and reasoning under asymmetric information and truthful revelation. We implement a controlled experimental design to test the theoretical predictions of full unraveling theory, while accounting for the bounded rationality of human beings. The empirical evidence is consistent with previous work and supports the notion of bounded rationality in strategic contexts. While we observe a considerable degree of unraveling, it falls short of the levels reported by Jin et al. (2021), confirming that full unraveling is highly context dependent.

²¹In the experiments, sender and receiver act independently; here we simulate random matching given the specified type distributions.

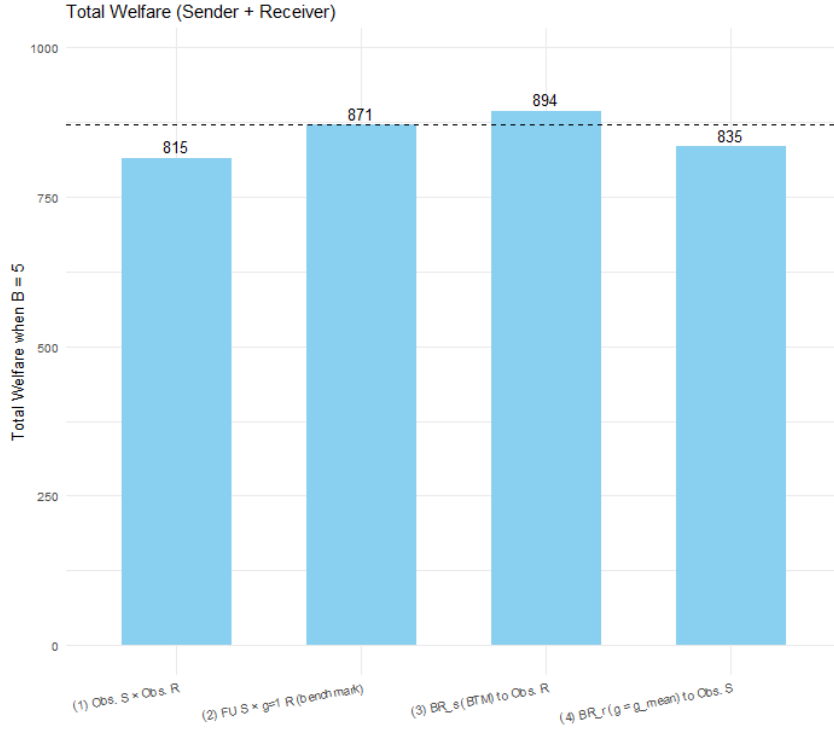


Figure 1.4: Total welfare across scenarios ($B = 5$). The dashed line indicates the $(FU \times g=1)$ benchmark.

We introduce two treatment variations: one within-subject and one between-subject. In the within-subject design, participants play the same game under three levels of complexity (two, three, and five possible quality values) presented in random order. We find that the share of senders classified as *full unravelers* rises sharply as the framework becomes coarser: from around 17% in the high-granularity case (five alternatives) to almost 82% in the simplest case (two alternatives). This confirms one of our central hypotheses: a larger number of alternatives increases the scope for strategic thinking, as senders anticipate that receivers may be confused by non-disclosure. Correspondingly, receivers exhibit higher guesses, but as the number of alternatives shrinks, their guesses in cases of non-disclosure converge toward 1, the minimum value predicted by theory. From a policy perspective, this result is highly relevant. It suggests that in markets marked by adverse selection (e.g., housing, insurance, etc.), providing too much information may foster strategic concealment and discouraging voluntary disclosure. Our findings support the idea that presenting a restricted number of labels can serve as a compromise: sufficient to sustain unraveling while avoiding the loss of information observed in the coarsest scenario. In line with Jin et al. (2021), we also confirm that more favorable values (i.e., higher qualities) are disclosed more frequently than lower ones across all scenarios.

The second innovation in our design concerns the between-subject variation, where participants are randomly assigned to one of three groups differing in the probability distribution of quality values. While unraveling theory predicts that such distributions

should not affect the final outcome, our evidence indicates the opposite: we call it *follow-the-distribution* bias. Both senders' first-order beliefs and receivers' guesses are sensitive to the assigned probability distribution, and their distributions closely mirror the shape of the assigned treatment information. Although average disclosure rates by senders do not differ significantly across groups, the treatment significantly affects receivers' guesses as well as senders' first-order beliefs about receivers' expectations. This finding has strong implications for policy making. In contexts such as the housing market, where, for instance, hidden information on climate risk plays a crucial role, the general conditions of the surrounding area (e.g., exposure to floods, wildfires, or earthquakes) can substantially influence expectations and the value attached to disclosure. Making such probability public information could therefore help this type of market interaction, though the context-dependent nature of this information should be carefully interpreted.

As explained earlier, rational senders should disclose all values and being indifferent to the lowest one, and expect that in the case of no disclosure, the receiver should infer that the hidden value is the minimum possible. When this does not occur, we expect players to compute the expected value and disclose all values equal and above the mean, assuming the receiver will guess the mean in the case of non-disclosure. Indeed, a large share of participants behaves as *bias-to-the-mean* types. However, we also observe that full unraveling does not center on the mean but instead happens at the maximum values (4–5), which reveals a systematic *confusion-at-the-mean* bias.

Another important result is the alignment between beliefs and behavior across sender and receiver experiments. Sender's first-order beliefs are consistent with observed receiver's actions, and second-order beliefs in the receiver experiment also match their actual guess. This consistency across different experiments indicates robustness, especially given the absence of feedback, learning, or direct interaction between participants.

At the same time, we find some inconsistency between sender's beliefs about receiver's guesses and sender's disclosure choices, and we observe that some sender fails to unravel even at maximum values. This is an unexpected finding, though also reported by Jin et al. (2021) at lower frequency. The authors interpret this behavior as other-regarding preferences or social motives. We aimed to rule out such effects, following the virtual player method tested by Ferraro et al. (2003); March (2019). This setup is supposed to isolate payoff-maximizing, belief-driven strategies and eliminate social-preference and reputation confounds. Thus, any apparent pro sociality in our experiment is easily explained by confusion or misunderstanding. Indeed, Nielsen et al. (2022) notes that some participants report acting as if they were playing against humans, even when the opponent is a computer. A limitation of our study is that we did not test comprehension of the instructions nor collect feedback on perceptions of playing against a computer. These omissions make difficult explain whether confusion is due to a misunderstanding the rules or a lack of rationality.

Additional limitations include the absence of sender–receiver interactions, which does not reproduce the dynamic feedback present in real markets, and the fact that we document only a restricted set of probability distributions. Additional external validity may require exploring empirically grounded distributions, such as those observed in real housing markets.

Future research should expand the complexity of the experimental setting, for instance, by increasing the number of alternatives, incorporating richer descriptions or images of the goods, and allowing real senders and receivers to interact. Following our research schedule, we aim at reproducing a similar framework in the field.

In summary, our findings provide several takeaways for the study of voluntary, costless disclosure under asymmetric information. We demonstrate that while coarser frameworks reduce informational precision, they also mitigate strategic concealment and encourage voluntary disclosure. Policymakers designing disclosure regulations should therefore carefully balance information granularity against the risk of discouraging transparency. Furthermore, the distribution of quality levels within specific contexts should be considered and, where possible, made publicly available. Once again, our results highlight that well-designed interventions can help address information asymmetries and correct market distortions.

Chapter 2

Walk the Talk? Greenwashing in the Electricity Market

Elena Sestini¹, Stefano F. Verde¹

¹ *Department of Economics and Statistics, University of Siena*

Abstract

Greenwashing is a widespread phenomenon rooted in information gaps and lack of transparency at the business level. In the electricity sector, firms hold private information on their actual environmental performance, while external stakeholders must rely on disclosed claims. This paper contributes to the literature in two main ways. First, it introduces a new way to quantify greenwashing in the electricity sector by combining textual content from company websites with data on the share of renewable sources in their fuel mix. Second, it identifies the firm-level characteristics associated with greenwashing behavior, with a particular focus on the link between economic performance and greenwashing activity. The results show that, while greenwashing is not widespread, it remains present among a subset of electricity providers. Companies that primarily operate in the electric sector are less prone to engage in greenwashing, while smaller firms show a higher likelihood of doing so. These findings deliver valuable insights to academic audience, policymakers and relevant stakeholders.

2.1 Introduction

Economic failures driven by the persistent problem of information asymmetry continue to threaten real world markets. The green market represents a typical asymmetric

information scenario where the public often lacks clear evidence on whether companies are truly meeting their environmental responsibilities (Roulet and Touboul, 2015). At the corporate level, greenwashing has become increasingly common, as companies exploit the knowledge gap between management and stakeholders to build a strong green image for economic gain. This trend is partly a response to the climate change crisis, which has pressured companies to improve their ESG performance. Financial incentives to appear environmentally friendly have led to misleading sustainability claims (European Securities and Markets Authority (ESMA), 2023), including the use of vague terminology, unrelated assertions, and improper certifications (Matejek and Gössling, 2014; Dempere et al., 2024). Companies often selectively disclose positive environmental information while concealing negative aspects, favoring qualitative soft information over quantitative, monetized hard information (Rupley et al., 2012).

The aim of this study is to measure the presence of greenwashing activity within the electricity sector, responsible for the green energy transition. We focus on firm-level heterogeneity as it may help explain and predict greenwashing. While numerous peer-reviewed studies have proposed different methods for measuring greenwashing, relatively little empirical evidence exists for the electricity industry, despite it represents one of the most environmentally sensitive sectors. Moreover, although many contributions examine the impact of greenwashing on firm performance and financial outcomes, there remains a significant gap in the literature concerning its potential drivers. This study aims to address that gap. In recent years, Italy has moved toward full liberalization of the energy market, completed in 2024 for both electricity and gas. This shift from a protected market to a free one for private users has heightened competition and attracted many new entrants. As shown in Table B.5 in the Appendix, the number of electricity providers rose from 276 in 2018 to 520 in 2022, with an exponential trajectory. In this increasingly competitive environment, it becomes even more critical to monitor fraudulent practices like greenwashing, which can deceptively enhance a company’s reputation at minimal cost if not detected. This paper contributes to the existing literature in several ways. First, it presents a novel tool to quantify greenwashing in the electricity sector. Building on the idea that greenwashing represents a decoupling (Yu et al., 2020), we measure greenwashing as the distance between what a company says and what it does: the first component is derived from quantifying the sustainability-related content on company websites, while the second component is based on the actual fuel mix data. We test two versions of the same index: the first, a peer-based version, assesses a company’s performance relative to its competitors and is calculated using percentiles, allowing each company’s score to be influenced by its peers’ behavior. The second, a normalized version, measures greenwashing in absolute terms, and is unaffected by the level of greenwashing among other companies. This is the first tool that combines qualitative data from websites with quantitative fuel mix data. The final dataset that we adopt for regression analysis has been created combining

web scraping and Natural Language Processing (NLP) analysis, with fuel mix data and balance sheet information at company level. This integrated approach not only offers an overview of greenwashing but also shares valuable insights into the determinants and consequences of this activity. We focus on the electricity market given its substantial environmental impact not only through the production of electricity but mostly via the quality of electricity sold. Some companies might choose to greenwash to mask underlying environmental shortcomings while building an eco-friendly image. Our research addresses the following questions: (i) How prevalent is the phenomenon of greenwashing among electricity companies? (ii) What are the common demographic and financial characteristics of firms that engage in greenwashing?

Our findings indicate that although the overall prevalence of greenwashing in the electricity sector is relatively low, a distinct subset of firms are engaged in activities that significantly misrepresent their environmental performance. Our regression analyses shows that smaller firms are more inclined to engage in greenwashing than larger ones and that sectorial heterogeneity plays a role. The relationship between economic performance and greenwashing appears to be positive, although with weak effects. We use lagged Profit Margin and Returns on Assets (ROA) as proxies for economic performance. This effect disappears for larger businesses.

Preventing the activity of greenwashing is crucial for several reasons. Deliberate distortion of information not only misleads consumers and investors by sending false signals about a firm's true economic and environmental performance, but it might also damage a company's financial standing and reputation, potentially resulting in investment losses and declining stock prices. In addition, unchecked greenwashing can produce a "market for lemons". As Akerlof (1978) shows, when buyers cannot distinguish high quality from low quality goods, the market unravels into lemons. In the green-energy context, unverifiable environmental claims play the role of unobserved quality. As a consequence, without credible signals, green suppliers compete at very similar prices as greenwashers, and may lose the incentive to invest in green electricity. It also carries legal risks and erodes trust in sustainable finance markets, undermining their integrity (European Securities and Markets Authority (ESMA), 2023), with stakeholders penalizing firms suspected of greenwashing (Weber, 2018).

Our study contributes to academic research on sustainability and greenwashing in the electric sector, and provides relevant policy insights to governmental regulatory bodies and corporate managers. For instance, the EU's recent Directive (EU) 2024/825, published in March 2024, represents a major step forward by imposing strict rules on business communications, targeting deceptive marketing practices, and ensuring that environmental claims are reliable and verifiable. However, mandatory sustainability disclosure currently applies only to larger companies meeting certain criteria, leaving many firms operating in various sectors with largely voluntary disclosure. As a result, the risk of misleading

claims in the green markets persists, eroding consumer confidence and distorting market perceptions. While environmental content is expected to be a significant part of a sustainability disclosure, the degree of emphasis and space devoted to environmental issues remain discretionary and may reflect managerial decisions (Gorovaia and Makrominas, 2024). Our perspective aligns with Dempere et al. (2024), who emphasize transparency and accountability, highlight the need for more stringent regulations, international cooperation, and increased public awareness. Our work demonstrates that companies still have considerable freedom to mix true and false information. We underscore the importance of vigilant oversight - especially on formal communication channels like websites and social media - to help restore trust in sustainable markets and ensure a fair distribution of resources.

The remainder of this paper is organized as follows. In Section 2, we provide an overview of the concept of greenwashing and its evolution over time; in Section 3 we present an extensive literature review, while in Section 4 we describe data collection process; in Section 5, we explain the computation and the interpretation of the novel greenwashing index and in Section 6 we describe the main variables. Last, Section 7 illustrates our empirical analysis and results, properly discussed in Section 8.

2.2 An overview of Greenwashing

2.2.1 Definitions

Greenwashing is generally identified as deceptive behavior involving false and misleading promotion of environmentally responsible practices or products to enhance corporate image. Over recent decades, the concept of greenwashing has evolved and been defined in multiple ways, reflecting its multifaceted nature. For this reason, there is no single, universally accepted definition (Lyon and Montgomery, 2015). In this literature review, we identify three common interpretations of greenwashing. The first is selective disclosure, referring to firms that communicate only favorable environmental information while concealing less flattering details (Lublóy et al., 2025; De Freitas Netto et al., 2020). The second interpretation is decoupling, which defines greenwashing as the gap between a company's actual environmental activities and its public claims (Lublóy et al., 2025; De Freitas Netto et al., 2020). This view emphasizes the prevalence of symbolic actions that generate "green talk" to meet stakeholder demands for sustainability without the actual support of substantive actions, thus failing to fulfill the commitments described verbally (Siano et al., 2017). Finally, greenwashing can also be understood through the lens of corporate legitimacy theory, which views firms as seeking social approval and legitimacy to operate within society (De Freitas Netto et al., 2020).

Delmas and Burbano (2011) differentiate between firm-level greenwashing, which engages firms cheating about their environmental performance and is central in this research,

and product-level greenwashing, that involves products are marketed as environmentally friendly to boost sales despite being environmentally detrimental.

2.2.2 Historical Evolution

The term *greenwashing* was first coined in 1986 by environmentalist Jay Westervelt, in reference to deceptive practices in the hotel industry aimed at promoting towel reuse (Orange and Cohen (2010)). The concept gained widespread attention in the early 1990s, and Gourier and Mathurin (2024) identified three relevant waves of greenwashing coverage. The first phase, from 1990 to 1992, and the second wave, from 2006 to 2010, were characterized by accusations against many large companies, primarily in the oil, gas and consumer goods industry, for delivering false information about their environmental action. The third wave, beginning in 2018, has been characterized by greenwashing accusations directed at investment firms. According to Gourier and Mathurin (2024), since 2018, the greenwashing index has consistently exceeded 8%, significantly influencing the discourse on climate risk in an unprecedented manner.

The threats posed by misleading environmental claims, combined with rising concerns over environmental degradation, prompted the European Commission to conduct an online screening of green claims in 2020. The investigation revealed that about half of the businesses analyzed failed to provide sufficient information or accessible evidence of genuine environmental commitment, with many companies relying on vague or generic terminology to obscure their actual environmental impact (European Commission, 2021). In response, the European Commission has taken decisive steps in recent years, introducing proposals to strengthen consumer protection and promote transparent communication, with the final adoption of Directive (EU) 2024/825.

2.2.3 Regulations

In recent decades, regulatory frameworks have played a crucial role in controlling greenwashing by establishing stringent guidelines for business practices. This activity is recognized as a serious threat worldwide.

In the United States, the Federal Trade Commission (FTC) has long been striving to combat misleading environmental claims. The FTC first introduced the Green Guides in 1992, which provide marketers with guidelines to avoid environmental statements that could deceive consumers. These guides offer advice on general principles applicable to all environmental marketing, interpretation of claims by consumers, substantiation of claims, and ways to avoid misleading consumers (Federal Trade Commission, 2012). More recently, in March 2024, the Securities and Exchange Commission (SEC) finalized the "Enhancement and Standardization of Climate-Related Disclosures for Investors" rule. Effective from 2025, this rule requires the largest publicly traded companies in the U.S.

to annually report their greenhouse gas (GHG) emissions and material climate risks UN Principles for Responsible Investment (2023).

In the European Union, several regulatory measures have been implemented to enhance transparency and consumer protection. In March 2023, the European Commission proposed the Directive on Green Claims to complement and execute the earlier proposal for a Directive on Empowering Consumers for the Green Transition. This Directive aims to ensure that green claims are reliable, comparable, and verifiable across the EU, thus protecting consumers from greenwashing and facilitating informed purchasing decisions European Commission (2024). It specifically targets voluntary claims made by businesses to consumers and is a key component of the European Green Deal and the Circular Economy Action Plan. The accepted proposal was published as Directive (EU) 2024/825 in April 2024. This directive also specifies guidelines for identifying misleading environmental claims, such as generic terms like “eco-friendly” and “biodegradable,” unsubstantiated future performance claims, and claims of carbon neutrality that rely solely on carbon credits outside the production’s supply chain¹. The Directive amends the earlier Directive 2005/29/EC on unfair business-to-consumer practices and Directive 2011/83/EU on consumer rights, emphasizing enhanced protection against unfair practices and improved informational transparency.

In the United Kingdom, the Financial Conduct Authority (FCA) has introduced a new anti-greenwashing rule with accompanying guidance in 2024. This rule mandates that financial service businesses ensure that sustainability claims about their products or services are fair, clear, and not misleading, thus protecting consumers by accurately representing sustainable products and services.

Lastly, Australia’s Consumer Law includes provisions that prohibit false, misleading, and deceptive conduct, extending to environmental claims. These regulatory efforts highlight a global commitment to mitigating greenwashing and fostering environmental integrity.

2.3 Literature Review

2.3.1 Firm-level Greenwashing Measurement

The first branch of the literature relevant to this study concerns methods for quantifying greenwashing at the company level. Early contributions in this field relied on qualitative assessments and content analyses of sustainability reports and corporate communications. Initially, these evaluations were conducted manually, and only later were they automated to facilitate the conversion of qualitative information into quantitative indicators. Typically,

¹The Directive is available in the EU website at this link

researchers employ ESG disclosure scores, ESG data, or the textual content of corporate reports and press releases to assess the extent of firms' environmental communication. A valuable contribution is provided by Lublóy et al. (2025), who summarize various approaches adopted in academic studies to measure greenwashing at the firm level. Their review highlights two key conceptual foundations underpinning much of the existing literature: selective disclosure and decoupling.

Studies building on the concept of selective disclosure define greenwashing as the tendency of firms to disclose only favorable environmental information while concealing unfavorable aspects, creating a misleading impression of environmental responsibility. One of the most influential measures developed within this framework is proposed by Marquis et al. (2016). The authors analyze a panel of more than 4,000 companies from 45 countries to construct a greenwashing index, defined as the difference between the absolute disclosure ratio (which represents the share of key environmental indicators addressed in corporate documents) and the weighted disclosure ratio (which measures the extent to which these indicators are substantively disclosed in a given year). On the same line, Huang and Huang (2020) introduce a novel measure that combines selective disclosure with the concept of expressive manipulation, defined as the use of symbolic or rhetorical language to create a green image without verifiable environmental performance. Using environmental reports from Chinese listed companies, the authors compute selective disclosure as the proportion of disclosed items relative to those required by regulation, while expressive manipulation captures the extent to which disclosures are symbolic rather than substantive. The greenwashing index is then derived as the geometric mean of these two components².

The concept of decoupling remains the most widely adopted framework in the greenwashing literature. It defines greenwashing as the divergence between corporate communications and actual practices, commonly referred to as symbolic and substantive actions respectively. In this view, firms seek to maintain a favorable public image through symbolic actions, while avoiding the costs and administrative burdens associated with sustainability efforts that would be reflected in substantive disclosures (Uyar et al., 2020). The most prominent and widely applied measure based on the decoupling framework has been introduced by Yu et al. (2020). The authors construct a greenwashing score using Bloomberg ESG disclosure data and Thomson Reuters Asset4 ESG performance data for an international sample of firms from 47 countries. Greenwashing is quantified as the difference between ESG disclosure scores and ESG performance scores, measuring the gap between what firms say and what they do in their environmental and sustainability practices.

Several approaches have been employed to quantify symbolic action. Some studies

²For related approaches, see also Zhang et al. (2023) and Jia and Li (2023), who extend and refine this framework.

focus on the volume or content of firms' official communications. For instance, Ruiz-Blanco et al. (2022) construct a greenwashing index for the S&P 100 companies in the United States by comparing the content of sustainability reports, evaluated using the sentence-level weighting system proposed by Beretta and Bozzolan (2004), with firms' actual environmental performance as measured by Bloomberg ESG scores. In line with this literature, our study follows Zhou and Wang (2024) and Li et al. (2023b), who proxy symbolic corporate action by computing the proportion of sustainability-related words contained in the textual corpus. Different approaches have also been proposed for measuring the substantive component of greenwashing. For example, Grewal et al. (2022) use total carbon emissions as an indicator of environmental performance, while Tang et al. (2023) and Zhou and Wang (2024) rely on firms' sustainability-related investments as a proxy for substantive action.

Building on this framework, several recent studies have proposed new approaches to measuring greenwashing. For instance, Ghitti et al. (2024) analyze a sample of the 500 largest U.S. publicly listed companies and introduce two novel indicators. The first measure incorporates environmental regulation violations as a key variable, identifying greenwashing when a firm's ESG rating exceeds that of its peers despite a higher number of recorded violations. The second indicator captures the discrepancy between projected and subsequently validated ESG ratings. In contrast, Gull et al. (2023) develop a greenwashing index based on lagged environmental performance scores. Using Bloomberg's current environmental disclosure scores for listed firms across 47 countries, the authors compute greenwashing as the difference between the current disclosure score and the lagged environmental performance score, normalized by the logarithm of total assets.

Relevant insights come from the work of Zhang (2023), who demonstrate that firms in cleaner production industries, in regions with highly developed green finance, in highly environmentally regulated areas, and state-owned enterprises are more restrained in greenwashing. They analyze a panel of Chinese listed firms during the period from 2011 to 2021 and measure greenwashing as the peer-relative difference between a normalized measure of a firm's Bloomberg ESG disclosure score relative to its peers in the distribution and a normalized measure of a firm's ESG real-performance score relative to its peers, which is a very similar approach to the one of our paper.

Another popular method which has been widely used to measure the gap between stated commitments and actual environmental practices is content analysis. Although already used in qualitative research, content analysis has recently become popular in quantitative projects. Bazillier and Vauday (2009), for example, quantify greenwashing by counting sustainable documents issued by companies and analyzing the length of these reports. Similarly, Grewal et al. (2022) develop an index based on the difference between "talk" and "walk" variables. To assess the talk variables, the authors measure the length of reports, the tone of content, and the specificity of language used, by employing a fine

tuned FinBERT algorithm³. To measure the actual performance they summed Scope 1 and Scope 2 GHG multiplied by -1. The final step involves computing the gap between rhetoric and tangible action, averaging the three indicators to derive a compound greenwashing index.

Li et al. (2023a) focus on heavily polluting industries and analyze a sample of firms listed on the Shanghai and Shenzhen stock exchanges. They construct a greenwashing variable defined as the difference between normalized measures of a firm’s Environmental Impact Disclosure (EID) and Environmental Performance (EP). A key contribution of this study lies in the inclusion of substantive actions from both input and output perspectives. Environmental investments serve as proxies for firms’ environmental commitment (inputs), while green patent applications capture environmental innovation outcomes. Both measures are normalized to reflect each firm’s relative standing within its industry. Following Hu et al. (2022), environmental investments are considered a reliable indicator of a firm’s environmental commitment and are log-transformed to correct for skewness. Similarly, the second component (that is the number of green patent applications) is log-transformed, in line with Quan et al. (2021), to evaluate firms’ environmental performance and innovation quality.

Further studies explore the extent to which firms discuss environmental compliance without substantive evidence (Zhang et al., 2023; Xu et al., 2023), or by analyzing the tone of the documents, focusing on the specificity of words used in sustainability disclosures and employing sentiment analysis (Grewal et al., 2022). Zhang et al. (2023) evaluates greenwashing as the peer-relative difference between a normalized measure of a firm’s ESG disclosure score relative to its peers in the distribution and a normalized measure of a firm’s ESG real-performance score relative to its peers. It uses Bloomberg data. A similar approach has been used by Liao et al. (2023). Peng and Kong (2024) develop an indicator system based on the prior research cited, focusing on four dimensions of corporate environmental governance: governance and structure, processes and control, inputs and outputs, and compliance and legality.

2.3.2 Using AI to Detect Greenwashing

The second body of literature we review focuses on adopting artificial intelligence (AI) to extract quantitative information from the qualitative content in company reports. Traditionally, content analysis was manually performed by experts, but recent advancements in AI and machine learning (ML) have introduced new methodologies. These technologies,

³FinBERT is a domain-adapted variant of BERT for finance, that is further pre-trained on large financial corpora and fine-tuned on task-specific labeled data.

particularly ML and natural language processing (NLP)⁴, are now crucial in evaluating the extensive and complex textual data found in sustainability reports Huang et al. (2023).

Several recent studies integrate artificial intelligence (AI) and machine learning (ML) techniques as tools to analyze corporate sustainability materials. For instance, using natural language processing (NLP), Kang and Kim (2022) and Luccioni et al. (2020) examine how sustainability is communicated in corporate disclosures, while Ning et al. (2021) applies ML-based topic modeling to explore online sustainability reports. Smeuninx et al. (2020) focus on the readability of such reports, and Amel-Zadeh et al. (2021) assess how well corporate narratives align with the United Nations Sustainable Development Goals (SDGs). Gorovaia and Makrominas (2024) employ standard NLP techniques to quantify the environmental content of CSR reports from U.S. public firms over the period 2008–2022, analyzing both tone and readability. They use a pre-established lexicon of environmental terms (DiCoEnviro, which is developed by linguists for NLP applications) to construct an environmental content score. Similarly, Lee and Kim (2021) implement a keyword-matching approach to develop an SDG social index, measuring the presence of SDG-related terms in social media textual data to capture public perceptions of corporate SDG commitments. The keywords are drawn from the Compiled Keywords for SDG Mapping file developed by the SDSN (Sustainable Development Solutions Network) Australia/Pacific and ACTS. However, as noted by Kang and Kim (2022), this method is limited by its reliance on exact word matching, which may overlook semantically related expressions. To overcome this limitation, Kang and Kim (2022) employ a sentence similarity approach based on a pre-trained language model, allowing for a more nuanced mapping of corporate disclosures to SDGs, complemented by sentiment analysis to differentiate between positive and negative information trends. Along the same line, Bingler et al. (2022) introduce ClimateBERT, a context-aware language model specifically designed to identify climate-related financial information in TCFD (Task Force on Climate-related Financial Disclosures) reports, distinguishing substantive environmental disclosures from purely rhetorical statements.

Other studies adopt ML to extract relevant topics from corporate documents. Fiandrino and Tonelli (2021) apply Latent Dirichlet Allocation (LDA) probabilistic topic modeling to examine documents related to the review of the Non-Financial Reporting Directive (NFRD), identifying the main themes and arguments emerging from the debate. Similarly, Ning et al. (2021) investigate the role of CSR in enhancing stakeholder engagement and firm performance, using LDA for topic modeling and a Latent Markov Model to explore the dynamic relationship between a firm’s strategic intent and its maturity over time. Smeuninx et al. (2020) apply NLP to assess text complexity, comparing the effectiveness of standard readability formulas against advanced NLP tools, which examine passive

⁴Natural Language Processing is the area of AI that enables computers to analyze, understand, and generate human language.

structures, syntactic depth, and lexical density.

Recently, a working paper by Gourier and Mathurin (2024) develops a news-implied greenwashing index using an advanced NLP algorithm on historical Wall Street Journal articles, identifying trends in the reporting of greenwashing relative to climate risk even if they do not contain the specific word. This index, built as the percentage of climate risk-related articles that are greenwashing-related over time, measures the proportion of articles that discuss greenwashing in the context of climate-related discussions, showing a steady increase since 2018. Another interesting working paper by Lagasio (2024) outlines a methodology for evaluating greenwashing practices in corporate sustainability reports. This method combines cutting-edge NLP techniques, ESG Focus Scores, and a Greenwashing Severity Index (GSI) to assess the depth of companies' commitments to environmental, social, and governance issues.

Our research relates to this area of literature for two primary reasons. Firstly, we utilize NLP to propose a novel list of sustainability-related terms, which helps in quantifying the presence of green-related terms on company websites. Secondly, we apply NLP to analyze material extracted from company websites and, when available, from their sustainability reports. This approach helps us understand not only the extent of companies' communication about their sustainability efforts but also the manner in which they release this information.

2.3.3 Greenwashing in the Energy Sector

The final section of our literature review examines the measurement of greenwashing within the energy sector. While a substantial body of research documents firms' greenwashing activities across various industries, relatively few studies have investigated this phenomenon within the energy sector, and even fewer within the electricity market.

Karaman et al. (2021) analyze CSR practices among energy firms listed in the Thomson Reuters Eikon database between 2012 and 2018 and find a positive correlation between CSR reporting, assurance practices, and actual sustainability performance. Their results suggest that, in general, companies in the energy sector are engaged in sustainability initiatives rather than using CSR as a pure marketing tool. Furthermore, firms with higher CSR performance tend to publish reports aligned with the GRI guidelines and to verify the reliability of the disclosed information. Jin et al. (2024) adopt a content analysis approach to measure the degree of greenwashing among Chinese energy companies, developing a scientifically grounded evaluation index system encompassing five dimensions: source control, business processes, governance, environmental management, and social reputation. The framework is based on 22 disclosure items representing key aspects that CSR reports should include to effectively assess greenwashing behavior. The authors compute a total greenwashing score by combining two metrics: selective disclosure and

declarative manipulation. Selective disclosure is calculated as the shortfall in required disclosure items multiplied by 100, while declarative manipulation is obtained by assessing the proportion of superficial disclosures relative to the total number of disclosed items, also multiplied by 100. The overall score, ranging from 0 to 100, captures the degree of corporate greenwashing, with higher values indicating stronger misrepresentation.

More recently, Boedijanto and Delina (2024) review the potential of AI in detecting greenwashing in the energy sector. They advocate integrating web scraping, NLP tools, and Life Cycle Assessment (LCA) databases to construct a more comprehensive and objective assessment of firms' environmental claims. The study concludes by recommending mandatory disclosure standards to enhance the accuracy and transparency of environmental reporting, noting that reliance on voluntary disclosure often results in distorted or incomplete information.

To the best of our knowledge, our research is the first to provide a comprehensive analysis of greenwashing using data on Italian electricity utilities and introducing an innovative measurement tool tailored to this context that be adopted in similar markets.

2.4 Transparency and Accountability for Renewables in the Italian Electricity Market

2.4.1 Main Actors

In the Italian electricity market, several key actors contribute to its functioning. At the top level, the Italian Government, mainly through the Ministry of the Environment and Energy Security (MASE) and the Ministry of Economy and Finance (MEF), sets the strategic policy direction, defining objectives such as decarbonization, the expansion of renewable energy, and the liberalization of the retail segment.

The independent regulatory authority, ARERA (Autorità di Regolazione per Energia Reti e Ambiente), translates these strategic goals into specific regulations and obligations. Although independent from the Government, ARERA operates within the framework of national and EU legislation. It oversees the functioning of energy markets, sets tariff methodologies, protects consumers, and establishes technical and operational standards for both electricity networks and market activities.

Three main market and system operators operate under ARERA's supervision: Terna, Gestore dei Mercati Energetici (GME), and Gestore dei Servizi Energetici (GSE). Terna acts as the Transmission System Operator (TSO), managing the national high-voltage grid and ensuring real-time balance between electricity generation and demand. GME functions as the market operator, organizing wholesale electricity markets—including the day-ahead, intraday, and ancillary services markets—as well as environmental markets such as the

trading of Guarantees of Origin (GO). GSE, a state-owned enterprise, issues Guarantees of Origin for renewable electricity and manages incentive mechanisms to promote renewable energy and energy efficiency.

Finally, there are electricity producers, traders, suppliers, and final consumers that continuously operate in the the market.

2.4.2 Guarantees of Origin (GO)

Guarantees of Origin (GOs) are electronic certificates adopted at European Union level to quantify the amount of electricity produced from renewable sources. The system was first standardized under Directive 2009/28/EC (RED I), which established that each GO corresponds to 1 MWh of electricity generated from renewable sources, and introduced the requirement for the computation of the Fuel Mix. Further refinements were made with Directive (EU) 2018/2001 (RED II), which extended the scope of GOs to cover all renewable sources and standardized the calculation of the Residual Mix at the European level⁵.

Today GOs represent the main accounting mechanism for renewable electricity in Europe. They are essential for electricity suppliers to prove that their green offers effectively sell renewable electricity, and they contribute to transparency and trust building. In Italy, GOs are issued by the Gestore dei Servizi Energetici (GSE) to producers of renewable electricity. For every MWh of renewable electricity injected into the grid, a corresponding GO is issued. These certificates can be traded independently from the physical electricity, meaning that a supplier can purchase GOs to prove to final consumers that part, or all, of their electricity supply comes from renewable sources. The trading and registration of GOs in Italy is managed by the Gestore dei Mercati Energetici (GME), which organizes three main platforms: (i) the M-GO (GO Market), the central marketplace for GO trading⁶. GME acts here as central counterparty, ensuring liquidity and transparency, with all prices publicly available; (ii) the PB-GO (Bilateral Platform for GOs), an electronic platform for the mandatory registration of bilateral GO transactions. Operators must register both freely negotiated bilateral trades and GOs allocated by GSE through competitive procedures; (iii) the GO Bulletin Board⁷, which promotes long-term contracting of GOs. On this platform, operators can anonymously post non-binding offers to buy or sell GOs, allowing potential counterparties to express interest.

Thanks to the transparent accountability ensured by the GO system, companies can reliably offer Green Tariffs. These contracts guarantee to final customers that the electricity

⁵Access to mentioned EU Directives here: Directive 2009/28/EC; Directive 2018/2001.

⁶Established under the 2011 AEEGSI Resolution ARG/elt 104/11

⁷Introduced by Ministerial Decree of 14 July 2023, no. 224

supplied comes from renewable sources. In order for a contract to qualify as a renewable electricity supply contract, the supplier must hold a volume of GOs equal to the amount of electricity sold as renewable, and the same GO cannot be allocated to more than one contract. Only by the cancellation of GOs can energy suppliers certify the quantity of electricity sold as renewable.

2.4.3 Fuel Mix Composition

The Fuel Mix refers to the set of primary energy sources used to generate the electricity supplied by retailers to final customers. The procedure that determines the composition of the fuel mix used by producers to generate electricity, and by suppliers to deliver it to customers, is called *Fuel Mix Disclosure*. This mechanism tracks the sources of electricity production, taking into account imported quantities and Guarantees of Origin data (GOs)⁸.

According to the Ministerial Decree of 31 July 2009⁹, the GSE is mandated to publish each year both the National Fuel Mix and the company-specific mixes. In turn, each electricity supplier is required to disclose its own mix composition to customers, through its website (when available) and on electricity bills. Subsequent updates, in particular Legislative Decree 28/2011 and Ministerial Decree of 14 July 2023, no. 224, have further refined the system and established the current platforms for the trading and registration of Guarantees of Origin.

In Italy, the GSE is responsible for computing and publishing the national and company-specific fuel mixes. GSE uses information received from producers and suppliers, together with other official data. It also verifies that each supplier has cancelled a sufficient number of GOs to prove to final customers that electricity sold under renewable contracts is genuinely renewable. By 31 March of each year, all electricity suppliers must submit to GSE the required information for the previous year. GSE, in collaboration with Acquirente Unico (AU), checks the accuracy of the data. In parallel, Terna provides information on the electricity fed into the grid during the previous year, broken down by network and producer, and including both national production and total imports. GSE computes three different fuel mixes in separated steps:

1. National Production Mix: GSE first uses Terna's data to compute the national production mix, expressed in percentage terms by source (e.g. renewables, coal, natural gas, etc.).
2. National Residual Mix: GSE then integrates this with GO data (issued and cancelled certificates) to calculate the national residual mix. An example of the national fuel

⁸For further information and official documentation, please refer to the GSE documentation

⁹Full text available at the following link.

mix is reported in Appendix, Figure B.6.

3. Company-Specific Fuel Mix: finally, GSE computes the mix of primary energy sources for each supplier. This is done by combining: (i) the information submitted by each supplier through the national online platform, (ii) the national residual mix, and (iii) the volume of GOs cancelled by the supplier in favor of its customers.

2.5 Data Collection

The data collection process involved multiple steps and integrated materials from various sources, as illustrated in Figure 2.1. To construct the greenwashing index, we require two types of information: what companies communicate and what they actually do. The former is obtained by scraping publicly available content from each company’s website. The latter is sourced from Fuel Mix data delivered by the Italian authority GSE. We complement this material with economic and demographic information downloaded from the Aida database¹⁰ with each company’s economic ID (partita IVA). Further details on economic and demographic variables are provided in Section 2.7.2.

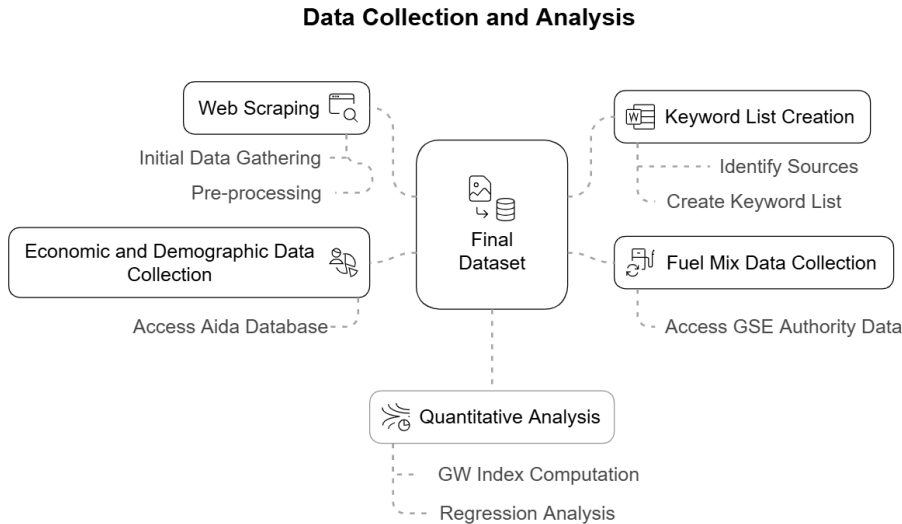


Figure 2.1: Data collection process

2.5.1 Web Scraping

Web scraping of electricity company websites was conducted in February 2024. We chose to rely on website content for several reasons. First, given the structure of the Italian

¹⁰Aida provides comprehensive economic, financial, and managerial public information on companies operating in Italy. Access to the website is under subscription.

electricity market, only a minority of electricity providers publish Non-Financial Disclosures (NFDs) or sustainability reports on a regular basis. Such documents are typically produced by listed companies or large enterprises subject to mandatory reporting obligations under national or EU regulations. Firms not covered by these obligations, particularly small and medium-sized enterprises (SMEs), rarely issue such reports voluntarily. Since a substantial share of our sample consists of SMEs, relying only on official sustainability disclosures would have significantly restricted the scope and representativeness of our dataset. Moreover, while official documents must comply with standardized formats and formal language to meet regulatory requirements, company websites offer greater communicative flexibility. On their websites, firms can present information in a way that best aligns with their branding and public image. In contrast to sustainability reports, often audited and directed toward investors and institutional stakeholders, website content primarily targets consumers seeking contract information, who may possess limited prior knowledge of the company. It reflects how they strategically communicate environmental responsibility to the general public. We therefore consider this material a reliable indicator of their intentions to engage in greenwashing.

To collect the web-based material, we begin by manually compiling the URLs of the companies' websites. We set a limit of two clicks from the homepage, given the heterogeneous structure of these websites. The manual collection is necessary to avoid errors that might arise from automated systems, particularly in handling pop-ups or missing critical links. Once the URLs were gathered, we employ Python packages for data extraction. The HTTP library `requests` is used to retrieve content from static pages, while `selenium` is employed for pages requiring interaction or dynamic loading. The HTML content is parsed using the `BeautifulSoup` module from the `bs4` package, which enable us to extract the desired information. To simplify the management of browser drivers needed by `selenium`, we utilize the `webdriver manager` package, which automatically detects the browser version, downloads the appropriate driver, and configures it.

We gathered data from 397 firms¹¹. The material for each company is saved as a text file and used for keywords frequency analysis.

2.5.2 Sustainability-related Keywords List

The next step involves compiling a list of sustainability-related keywords to measure their frequency within the scraped website content of electricity companies. The presence and intensity of these terms serve as indicators of the extent to which firms engage in pro-environmental communication and promote their sustainability efforts to the public.

¹¹This is the number of companies who had accessible websites on February 2024, given the list for which we had fuel mix information for 2022. The 2023 fuel mix data have been made available at the end of 2024 by the authority.

The methodology for constructing the keyword list partly follows Blasi et al. (2021), who extract a set of relevant sustainability-related terms performing a content analysis of a well-known book on the circular economy. We have identified seven relevant papers plus the SDGs official document on which we perform keywords extraction¹². Each selected document is preprocessed¹³ and tokenized¹⁴. To ensure robust extraction of relevant keywords from each document, we employ both the Bag of Words (BoW)¹⁵ and TF-IDF¹⁶ methods. Using these two approaches, we extract a total of top 100 words from each document (50 with each method) and merge them into a single comprehensive list. This complementary approach allows for the inclusion of both the most frequently used words and those that are relatively unique across documents, an important step to ensure broad coverage of relevant sustainability-related terms.

The first extraction phase, based on sectoral research, primarily captures energy and economics oriented sustainability terms. The second phase involves extracting additional keywords from three online sustainability dictionaries¹⁷, which provide a broader and more technical vocabulary not limited to economic contexts. The collected terms are then integrated into a final list.¹⁸ This final list is translated into Italian and manually filtered to maintain only words explicitly related to sustainability. The final list counts 217 stemmed terms used for frequency computations.

For text extraction from PDF documents, we use Python’s PyMuPDF library (`fitz`). Each document is tokenized using the `nltk` library, and standard preprocessing procedures, such as the removal of special characters, numbers, and stopwords¹⁹, are applied to ensure data consistency and reliability across sources.

¹²See Table B.4.

¹³Text preprocessing refers to the preparation of raw text for analysis or modeling. This typically includes steps such as removing punctuation, converting all text to lowercase, and eliminating stopwords (e.g., prepositions, articles, and conjunctions) to simplify the input.

¹⁴Tokenization is the process of splitting cleaned text into smaller units, or tokens, such as individual words or phrases.

¹⁵The Bag of Words (BoW) model represents a text by counting the frequency of each word in the vocabulary, disregarding grammar and word order.

¹⁶TF-IDF (Term Frequency–Inverse Document Frequency) assigns a weight to each word based on its frequency in a document, adjusted for how rare it is across all documents.

¹⁷See Table B.5.

¹⁸Unlike academic sources, sustainability dictionaries tend to include several technical and domain-specific terms, enriching the overall coverage.

¹⁹Stopwords are commonly used words that are filtered out during preprocessing to reduce noise in NLP tasks, as they contribute little semantic value.

2.5.3 Fuel Mix Data

Fuel Mix composition data is provided by the GSE, the Italian authority responsible for promoting renewable electricity production and electricity efficiency. This dataset includes the percentage share of renewable energy sources, which constitutes one of the two main components of the greenwashing variable. In addition, it contains percentage breakdowns for carbon-based sources, natural gas, oil derivatives, nuclear, and other energy sources, along with the total energy sold measured in MWh. The complete dataset for 2024 counted 520 Italian electricity providers. After excluding those companies for which we didn't collect material from websites²⁰, the final sample on which we are able to compute the greenwashing index is made of 352 firms. Descriptive statistics of fuel mix components are collected in Table B.1 in Appendix.

2.5.4 Data Preprocessing

After downloading the material from the web, we initiate the data cleaning process. Initially, we count the total number of words by removing punctuation and tokenizing the text. Standard preprocessing steps are applied, including punctuation removal, tokenization, stopword removal, and the elimination of tokens shorter than three characters. Subsequently, we standardize the tokens by applying both lemmatization and stemming to facilitate frequency analysis. For each company, we calculate the frequency of sustainability-related terms by dividing the number of keyword occurrences within the scraped website content, by the total number of tokens.²¹ The analysis is conducted using both lemmatized and stemmed versions of the keyword list, as shown in Figure B.3. Although lemmatization is generally preferred for words similarity computation (Pramana et al., 2022), we decide to use stemmed words as we are not just comparing the precise words matching but we want to gather any nuances of green messages, which with more precise words identification might be disperse. Figure B.10 shows the word cloud of the most frequent terms from the keyword list, based on the entire text corpus.

Descriptive statistics are collected in Table B.1. On average we find 5,228 words per website and 270 words belonging to the sustainability-related keywords list, which account for an average of 5% of the entire corpus.

²⁰Companies excluded from the scraping process are those that either did not have an active website at the time of data collection or were not included in the fuel mix dataset published by GSE for the year 2022. Conducting additional scraping in subsequent years would not yield comparable results.

²¹This variable represents the *Disclosure* component of the greenwashing index, as discussed in the following section.

2.6 Greenwashing Index

Building on the literature on decoupling, we propose a novel greenwashing index tailored to electricity firms and grounded in new data sources. Following Yu et al. (2020), we compute Greenwashing as the difference between two components:

$$GW_i = Disclosure_i - Performance_i \quad (2.1)$$

where $Disclosure_i$ identifies how much a company talks about sustainability in its website: here it is measured as the presence of sustainability-related terms on company websites for company i . It is calculated by counting the number of occurrences of green terms from the keywords list previously defined, in the total number of words extracted from company i 's websites. On the other side, $Performance_i$ reflects the company's actual environmental performance: this component represents the percentage of renewable electricity for company i in its fuel mix.

The index does not contrast renewable-specific claims with renewable-specific performance metrics. Instead, it measures the symbolic component of sustainability communication by considering the general intensity of sustainability-related content on corporate websites, and compares it with a standardized indicator of environmental performance. This reflects the nature of environmental communication in the electricity sector, where firms primarily present their environmental positioning using broader symbolic language; similarly, the renewable component of the fuel mix represents the most direct, transparent, and comparable indicator of firm-level environmental performance for electricity retailers. By combining a qualitative disclosure measure with a quantitative performance benchmark, the index does not measure the accuracy of specific statements to specific actions. Rather, it quantifies whether firms that place greater emphasis on sustainability communication online also exhibit similar environmental performance.

Positive values suggest the presence of greenwashing, where the higher the score, the stronger the greenwashing activity by the company. Negative values indicate the opposite phenomenon, also known as *greenmuting*, with companies acting more sustainably than they claim. A score near zero means a company's actions align closely with their words, indicating no greenwashing.

Given this framework, we promote two versions of the index, computed using different normalization techniques. The first is a peer-based index and delivers a clear picture of each company's performance relative to its competitors: we compute percentiles of both components and we use them in Equation 2.1. The distribution of the percentiles of the two components, $Performance$ and $Disclosure$, are plotted in Figures B.1a and B.2a. The distribution of the Peer-based Greenwashing Index (GW Peer) is depicted in Figure 2.2a.

The second version is a Normalized Greenwashing index (GW Norm), computed by applying the min-max normalization for both $Disclosure_i$ and $Performance_i$, before

integrating them into Equation 2.1. For the component *Disclosure* the normalization is as follows:

$$\text{Disclosure}_{norm} = \frac{\text{Disclosure}_i - \text{Disclosure}_{min}}{\text{Disclosure}_{max} - \text{Disclosure}_{min}} \quad (2.2)$$

where Disclosure_{min} and Disclosure_{max} represent the minimum and maximum values found in the dataset, respectively. We apply the same normalization process to the *Performance* variable, labeled as $\text{Performance}_{norm}$. The resulting score, GW_i , ranges from -1 to 1 . The distribution of the GW Norm index is depicted in Figure 2.2b. This version of the index attributes a score which is independent of the performance of peers, and compute the final score relative to the broadest possible spectrum of observed corporate behaviors.




Figure 2.2: Different GW indices

Following previous studies, and as a robust check, we compute also a standard normalized version index, whose description is left in Appendix B and its distribution is in Figure B.7. All the three versions are robust and highly correlated, as shown in Table B.9 in Appendix.

2.7 Variables

2.7.1 Dependent Variables

The dependent variables used in regression analysis are *Peer-based* and *Normalized Greenwashing Indices*, which are based on the same data but differ in the normalization of components. Both versions are continuous and range from -1 to 1 , with the maximum value identifying greenwashing and the minimum one indicating the opposite phenomenon, also known as greenmuting. Values close to 0 indicate absence of greenwashing and a balanced tradeoff between environmental performance and disclosure efforts. This means that both green companies, documenting their engagement towards the environment, and

brown companies, not mentioning green items in their website, can be around 0. Further details are provided in Section 2.6.

2.7.2 Independent Variables

Demographic and financial variables have been computed or downloaded from the Italian Aida database. Table B.8 provides a detail description of all variables. Note that all the economic variables have been used as regressors at their t-1 lags. First we identify two variables that can be used to measure the economic performance of a company: (i) *Net Profit Margin (NPM)* has been computed as $NPM = (\text{Net Income}/\text{Revenues}) \times 100$ and it measures the percentage of revenues that remain as profit after accounting for all expenses (operating, interest, and taxes); it represents a standard indicator of profitability in financial analysis (Nariswari and Nugraha, 2020). The second (ii) is Return on Assets (ROA) and is calculated as the ratio of EBIT²² to Total Assets. This profitability index is commonly used as a proxy for economic performance (Liao et al., 2018; Uyar et al., 2020; Cariola et al., 2020) and represents the return on the capital invested in the company. Then we include other economic variables to get insights into the financial and strategic structure of businesses. We include (iii) *Financial Leverage*, calculated as the sum of long-term and short-term debts divided by total assets (Cariola et al., 2020) and (iv) *risk management*, computed as the ratio of capital expenditure to equity. This measure serves as a proxy for the risk associated with management’s investment strategy. We also include the continuous variable (v) *Market Share*, expressed in KWh, that measures the share that each company holds in the market.

To check for sector heterogeneity we use the ATECO 2007 code, the Italian classification of economic activities²³. Based on the ATECO code we define three dummy variables: (vi) *Provider* assumes value 1 if the company has code 35.14 and 0 otherwise, indicating those companies whose core activity is trading of electricity; (vii) *Other* assumes value 1 if ATECO code is either 35.11 (production of electricity), 35.12 (transmission of electricity) or 35.13 (distribution of electricity), and 0 otherwise, indicating all the business in the electric market but selling electricity; (viii) *Electric* assumes 1 if the code is one of the previous mentioned, indicating all companies whose main business relates to electricity market. The complete list of ATECO 2007 codes is presented in Table B.3.

Next, we control for the (ix) *age* of a company: using the foundation year provided for each company, we computed the number of years the company has been operating,

²²Earnings Before Interest and Taxes (EBIT) is obtained by subtracting operating expenses (including depreciation) from ordinary revenues.

²³It is regulated and updated by ISTAT and represents the Italian version of the European statistical classification of economic activities (NACE), which provides a standardized way to categorize industries across the European Union

following Blasi et al. (2021). To reduce right skewness, since there are many more young companies than older ones, we use the logarithm of this variable. Then we include a geographical dummy, (x) *North*, that assumes value 1 if the company is legally based in a region from the north of Italy and 0 otherwise. Another relevant information is held in the variable (xi) *Foreign*, a dummy equal to 1 if the ultimate owner of the company is not from Italy and 0 otherwise. Last, we check for the (xiii) *size* component using the logarithm of the firm’s total assets (Blasi et al., 2021).

2.8 Empirical Analysis

2.8.1 Descriptive Statistics

After cleaning the data and removing missing values, the final sample used for the regression analysis consists of 337 observations. Table B.6 reports descriptive statistics for the main variables employed in the analysis. It is also worth briefly describing the structure of the Italian electricity market in terms of numbers. In recent years, particularly close to the mandatory liberalization in 2024, the number of electricity providers has increased substantially. While in 2018 there were around 250 providers, by 2022 this figure had risen to more than 500. This sharp growth of market entries reflects a huge increase in competition. As a result, the market is now highly heterogeneous, characterized by a large number of small and medium-sized enterprises alongside a few dominant players.

Table 2.1: Business Analysis by Size

Size	n. Providers (electr.)	n. Other (electr.)	n. Electric Sector	n. Non-Electric Sector
Micro	51	3	55	15
Small	47	7	54	39
Medium	38	5	46	50
Large	29	7	37	40

To illustrate this heterogeneity, Table 2.1 reports the distribution of firms by size category, while Table B.7 in the Appendix presents the average values of key variables across these groups. The results show that the majority of firms whose core business lies in electricity provision belong to the micro and small size categories, indicating that the electricity sector in Italy remains largely composed of smaller firms. By contrast, larger firms are less likely to operate exclusively within the electricity market, as evidenced by the patterns reported in the last column of Table 2.1.

2.8.2 Regression Analysis

The aim of this section is to explore whether firm-level demographic and economic characteristics can explain differences in greenwashing behavior. We estimate the following

linear model:

$$GW_i = \alpha + \mathbf{\Gamma}^\top \mathbf{E}_i + \mathbf{\Delta}^\top \mathbf{D}_i + \varepsilon_i, \quad (2.3)$$

where GW_i is the greenwashing index, defined in Section 2.6, and captures the gap between a company's stated environmental claims and its actual environmental performance²⁴. \mathbf{E}_i contains economic variables, while \mathbf{D}_i includes demographic controls, and ε_i is the error term. All economic variables are lagged by one year to reduce simultaneity concerns.

We estimate several model specifications by varying the dependent variable (GW Peer vs GW Norm) and the set of regressors. Table 2.2 presents the OLS estimates. Panel A uses lagged net profit margin (NPM_{t-1}) as the key economic indicator (columns (1)–(4)), while Panel B offers a robust analysis replacing it with lagged return on assets (ROA_{t-1}) (columns (5)–(8)). All models control for sector affiliation (Electric Provider, Other Electricity-related, or a broader Electricity Sector dummy), firm age, firm size, leverage $_{t-1}$, risk management $_{t-1}$, market share, geographical provenience and if the ultimate holder is Italian or not.

At first glance, results using the GW Peer index as the dependent variable appear relatively weak, with few statistically significant effects. In contrast, more meaningful insights emerge when using the normalized greenwashing measure. Note that the first measure, by construction, neglects differences in absolute values between two companies. Results are broadly consistent across both panels.

We find a clear sectoral pattern. Firms whose core business is the electricity market show significantly lower greenwashing scores, particularly when measured using the normalized index. In Panel A, the Provider and Other Electricity-related dummies are both negative and significant at the 5% level in column (3). In Panel B, similar results hold in column (7). When using the broader Electricity Sector dummy (columns (2), (4), (6), (8)), the coefficient is consistently negative and highly significant for the normalized measure (−0.0950, p-value 0.0035), and although still negative for the peer measure, it does not reach conventional significance levels. Overall, firms whose main activity is within the electricity sector are less likely to engage in greenwashing.

In all specifications, we observe a persistent, negative, and strongly significant effect of firm size on greenwashing: larger firms tend to greenwash less. Figure B.11 in the Appendix provides a visual representation of the index, with the disclosure component on the y-axis and the performance component on the x-axis. The 45-degree line represents equality between disclosure and performance, corresponding to the absence of greenwashing.

²⁴We use both the Peer-based index and the Normalized index, as they provides similar but still different interpretation of greenwashing among firms.

Table 2.2: OLS: Greenwashing

Variable	Panel A (NPM)				Panel B (ROA)			
	GW Peer		GW Norm		GW Peer		GW Norm	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Provider (electr. sect.)	-0.021 (0.041)	–	-0.076** (0.034)	–	-0.021 (0.041)	–	-0.077** (0.034)	–
Other (electr. sect.)	-0.076 (0.083)	–	-0.198** (0.101)	–	-0.078 (0.082)	–	-0.201** (0.100)	–
Electric sect.	–	-0.031 (0.037)	–	-0.099*** (0.033)	–	-0.030 (0.037)	–	-0.098*** (0.033)
Age (log years)	0.027 (0.031)	0.022 (0.029)	0.028 (0.028)	0.015 (0.028)	0.032 (0.031)	0.026 (0.029)	0.035 (0.027)	0.022 (0.028)
Size _{t-1} (log assets)	-0.030*** (0.010)	-0.030*** (0.010)	-0.024** (0.010)	-0.022** (0.010)	-0.030*** (0.010)	-0.029*** (0.010)	-0.023** (0.010)	-0.021** (0.010)
Financial leverage _{t-1}	-0.037 (0.111)	-0.026 (0.107)	0.089 (0.112)	0.112 (0.110)	-0.043 (0.105)	-0.032 (0.102)	0.085 (0.104)	0.107 (0.103)
NPM _{t-1}	0.005 (0.003)	0.005 (0.003)	0.007** (0.003)	0.007** (0.003)	–	–	–	–
ROA _{t-1}	–	–	–	–	0.003** (0.002)	0.003** (0.002)	0.005*** (0.002)	0.005*** (0.002)
Risk management _{t-1}	0.0003 (0.002)	0.0003 (0.002)	0.001 (0.001)	0.001 (0.001)	0.0002 (0.002)	0.0001 (0.002)	0.001 (0.001)	0.001 (0.001)
Market share (KWh)	0.001 (0.002)	0.001 (0.002)	0.001 (0.002)	0.001 (0.002)	0.001 (0.002)	0.001 (0.002)	0.001 (0.001)	0.001 (0.001)
North	-0.092** (0.038)	-0.094** (0.037)	-0.049 (0.033)	-0.052 (0.033)	-0.091** (0.037)	-0.093** (0.037)	-0.047 (0.032)	-0.050 (0.033)
Foreign	0.114 (0.071)	0.108 (0.070)	0.063 (0.075)	0.050 (0.077)	0.120* (0.071)	0.114 (0.071)	0.075 (0.076)	0.061 (0.079)
Constant	0.299*** (0.108)	0.300*** (0.108)	0.255** (0.116)	0.259** (0.117)	0.278** (0.110)	0.279** (0.110)	0.220* (0.114)	0.224* (0.115)
Observations	336	336	336	336	336	336	336	336
R ²	0.060	0.059	0.065	0.059	0.063	0.062	0.076	0.070
F Statistic	2.064** (df = 10;325)	2.253** (df = 9;326)	2.249** (df = 10;325)	2.255** (df = 9;326)	2.188** (10;325)	2.386** (9;326)	2.686*** (10;325)	2.706*** (9;326)

Notes: Robust standard errors in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Each dot corresponds to a firm, with colors varying by size. While no perfectly clear pattern emerges, we can observe a prevalence of darker red dots, indicating larger firms, clustered around the 45-degree line, suggesting a higher concentration of larger firms with greenwashing levels close to zero.

In addition, we find a positive relationship between past year profitability and current greenwashing behavior, although the magnitude and significance of the effect are modest. Profit margin_{t-1} is statistically significant only when using the normalized GW index: a 1% increase in profit margin corresponds to a 0.006 increase in greenwashing (p-values 0.0575 and 0.0503 in (3) and (4)). ROA_{t-1} is positive and significant in models (6), (7), and (8), with a 1% increase in ROA associated with an increase in greenwashing of 0.004 (p-values 0.0641, 0.0071 and 0.0061 in (6), (7) and (8)). While statistically significant in some models, these coefficients are very small in size and should be interpreted with caution. In addition, firms with legal provenience from the North exhibit less greenwashing compared to the rest, although this is true only in specifications (1),(2),(5),(6). Other

controls such as firm age, risk management, financial leverage, market share and ultimate owners provenience do not appear to be significantly associated with greenwashing in any specification. Due to the cross-sectional nature of firm-level data, the R^2 and adjusted R^2 values are relatively low, indicating that much of the variation in greenwashing is explained by unobserved factors. However, all models pass joint F-tests.

For additional analysis, we examine the relationship between publishing a sustainability report and the propensity to greenwash. We re-estimate the models replacing the variable *size* with the dummy *report* to mitigate collinearity concerns, given that firm size is a primary driver of mandatory sustainability reporting. The results, reported in Appendix, Table B.10, are fully consistent with the baseline specifications. Moreover, the coefficient on *report* is negative and statistically significant across all models, indicating that publishing a sustainability report is associated with a lower propensity to greenwash.

To improve model selection and avoid overfitting, we also estimate a LASSO (L1-penalized) regression with 5-fold cross-validation. The optimal penalty parameter λ is chosen to minimize out-of-sample MSE. Predictors with nonzero coefficients are then re-estimated using OLS with robust standard errors. Variance inflation factors (VIFs) are computed post-LASSO and reveal no multicollinearity concerns. Table B.11 reports the post-LASSO results. In Model (1) (GW Peer as dependent variable), only ROA_{t-1} and firm size are selected. In Model (2) (GW Norm), two sector dummies (Electric Sector and Other Electricity related) are also maintained. In both models, higher ROA is associated with more greenwashing (coefficients of 0.0030 and 0.0040, and p-values 0.0887 and 0.0129 respectively). Larger firms are significantly less likely to greenwash (-0.022 and -0.015 , p-values 0.0028 and 0.0343), with stronger significance. Electric sector dummy continue to show negative estimates in Model (2) (-0.0841 , p-value 0.0074). Both models pass joint significance tests, and model fit improves slightly compared to the baseline OLS. Overall, we confirm previous findings.

Given the strong effect of firm's size, we estimate another model that includes an interaction term between firm size and profit margin. Table 2.3 shows the results for two specifications using the GW Norm index. The key findings are robust: the sector variables remain negative and significant, and firm size continues to be negatively associated with greenwashing. Profit margin maintains its positive coefficient, while the interaction term between size and profit margin is strongly negative and statistically significant (-0.003 , p-values of 0.0085 and 0.0084). This suggests that the positive association between profitability and greenwashing becomes weaker for larger firms.

As a final step, we conduct size-based subsample regressions. Following the European Commission classification (European Commission, 2003), we divide firms into four groups: micro (turnover ≤ 2 million of euros), small (turnover ≤ 10 million of euros), medium (≤ 50 million of euros), and large ($>$ turnover ≤ 50 million of euros). Table B.12 in the Appendix reports results for SME and large firms. In line with previous results, the

Table 2.3: Models with interaction term: Profit-margin effect by firm's size

	Dependent variable: Greenwashing Norm	
	(1)	(2)
Provider (electric)	-0.078** (0.033)	—
Other (electric)	-0.195** (0.099)	—
Electric sector	—	-0.099*** (0.032)
Age (log years)	0.025 (0.027)	0.013 (0.028)
Size _{t-1} (log assets)	-0.018* (0.010)	-0.016* (0.010)
Profit Margin _{t-1}	0.035*** (0.011)	0.036*** (0.011)
Financial Leverage _{t-1}	0.086 (0.108)	0.108 (0.106)
Risk Management _{t-1}	0.001 (0.001)	0.001 (0.001)
North	-0.058* (0.032)	-0.061* (0.032)
Foreign	0.052 (0.074)	0.038 (0.076)
Size × Profit Margin _{t-1}	-0.003*** (0.001)	-0.003*** (0.001)
Constant	0.218** (0.110)	0.220** (0.110)
Observations	336	336
R ²	0.082	0.076
Adjusted R ²	0.054	0.051
F Statistic	2.900*** (df = 10; 325)	2.997*** (df = 9; 326)

Note: Robust (HC1) standard errors in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

association between profitability and greenwashing is significant only for medium-sized firms (specifications 2 and 4). Models for large firms (specifications 1 and 3) do not pass joint significance tests, suggesting weaker or no reliable relationship in that group. Interestingly, these subsample regressions also reveal that the electric sector effect is concentrated among small and medium-sized firms: firms in the electricity sector are consistently less likely to greenwash, and this is strongest in the SME segment. These findings suggest that profitability contributes to greenwashing primarily among smaller firms, and this effect diminishes as firm size increases. Meanwhile, electricity sector affiliation is robustly associated with lower greenwashing intensity, particularly among SMEs.

2.9 Discussion and Conclusion

Given growing global environmental concerns, stringent environmental regulations, and the influence of diverse stakeholders, adopting environmentally sustainable practices has become not only desirable but essential for businesses. Companies today have to address

an important challenge: pursuing economic performance while fulfilling environmental responsibility.

Public authorities and regulatory bodies typically focus on enforcing environmental standards and monitoring corporate compliance, thus emphasizing environmental performance. In particular, electricity firms are under growing pressure to develop sustainable offerings and actively contribute to the clean energy transition (Vollero et al., 2016). This may incentivize firms to adopt sustainable practices to meet market expectations and preserve their reputation. Also customers show a growing demand for environmentally friendly products and positive attitude toward green brands and sustainable energy solutions which might push companies to act responsively (Chen, 2008). However, asymmetric information and lack of verifiability, may allow companies “talking the environmental talk, without really walking the walk” (Vos, 2009). Other stakeholders, such as suppliers and industry peers, are mainly interested on economic performance. This divergence in stakeholders pressure can lead to strategic misalignment, where firms engage in symbolic environmental actions or greenwashing to preserve competitiveness and legitimacy at a lower cost (Ali et al., 2025). In such contexts, distinguishing between authentic environmental activities and deceptive communication becomes particularly critical.

To address this challenge, we develop a novel metric to detect greenwashing in the electric sector, combining real-world environmental performance data with sustainability claims extracted from companies’ websites. Our findings indicate that while greenwashing is not pervasive across the electricity market, it remains a non-negligible activity, highlighting the importance of continued scrutiny and the development of robust tools to differentiate substantive from symbolic action. These results are in line with Zhang (2023), who showed that companies operating in cleaner industries or regions with developed green finance and stringent environmental regulation, are more restrained in greenwashing behavior.

Interestingly, our findings further indicate that greenwashing is prevalent among SMEs, rather than among large firms. This misaligns with Zhang (2022), who identify large size as a significant determinant of greenwashing intensity. However, several reasons may explain this pattern. SMEs often lack the technical capabilities and financial resources required to implement meaningful environmental change. For these firms, even small reputational gains from appearing environmentally responsible can translate into relatively large competitive advantages. In contrast, large companies typically face greater reputational risks and are more likely to have established internal governance mechanisms that ensure consistency between sustainability communication and operational behavior. The absence of such systems in smaller firms may contribute to unintentional inconsistencies or facilitate opportunistic communication strategies.

Our results also reveal sectoral heterogeneity: firms whose core business is in electricity (i.e. generation, distribution, or retail) appear significantly less likely to engage in greenwashing. This evidence aligns with Ruiz-Blanco et al. (2022), who demonstrates

that companies in environmentally sensitive industries are less involved in greenwashing. Electricity firms are also more frequently subject to regulatory oversight and benefit from incentive schemes (such as feed-in tariffs or capacity markets, which reward verifiable environmental performance (Lin and Xie, 2024)).

Our findings report a positive association between previous year profit margins and current greenwashing behavior, particularly among smaller firms. The estimated effect is modest and only marginally significant. Because our data are cross-sectional, we interpret this result with caution. The absence of greenwashing measures at $t - 2$ prevents us from assessing the persistence or pre-trends of this relationship. Contrary to Zhang (2022), who identify financial constraints, particularly leverage, as one of the main drivers of greenwashing, our results do not provide any evidence related to financial leverage nor risk management.

Our regression analysis presents some limitations. First, since our data on corporate environmental communication are obtained through web scraping, the sample is limited to companies that hold an official website. This may introduce unavoidable selection bias, particularly for smaller or local electricity providers, whose online presence is often limited or poorly maintained. Our results might underestimate the effect. Second, the analysis is based on cross-sectional data, which restricts our ability to draw conclusions about temporal dynamics or causality. A longitudinal design would allow for a more robust assessment of the evolution of greenwashing practices and their relationship with firm characteristics over time. Third, in order to maintain a consistent scraping procedure, we applied a uniform rule to retrieve content within two clicks from the homepage. While this ensures comparability, it may overlook some material for few companies' websites which have complex and deeply nested structures. Last, although we plan to develop an integrated version of the greenwashing index that incorporates information from sustainability reports, this is currently feasible only for a limited subsample of large firms, as small and medium-sized enterprises rarely publish such documentation. To address these weaknesses, future work might focus on developing an automated web scraping system able of collecting data over multiple years, enabling a panel data structure. For further research, we aim to expand the scope of the analysis by evaluating how greenwashing and environmental performance jointly affect subsequent economic outcomes, assessing not only the determinants but also the consequences of corporate environmental communication strategies.

Our findings have important implications for government and regulatory authorities, as well as for the management of electricity companies and their key stakeholders, including investors and consumers. The main policy implication of this study is that the incidence of greenwashing can be reduced through greater transparency of information provided to energy users. For instance, publicly available data such as the fuel mix disclosure, which is already subject to regulatory oversight, plays a key role in reducing information asymmetries and limiting the scope for misleading communication. Our results therefore underscore

the crucial role of regulators in promoting transparency and accountability across the electricity sector. In line with Dempere et al. (2024), we support the call for stricter environmental reporting standards, stronger international cooperation, and enhanced public awareness to counter deceptive environmental communication. Corporate Social Responsibility (CSR) reporting, in this regard, should serve as an essential mechanism to legitimize corporate actions and facilitate access to environmentally responsible markets (Karaman et al., 2021; Zhang et al., 2023).

Recent developments at the European level, particularly the 2024 EU Directive 2024/825 aimed at protecting consumers from greenwashing, represent an important step forward. Nonetheless, regulatory gaps persist, especially in digital environments and online platforms that fall outside the current legal framework. In this light, we argue that vigilance, transparency, and accountability must be extended across all communication channels to ensure that firms not only talk the environmental talk but also truly walk the walk.

Chapter 3

Estimating Residential Electricity Demand under Time-of-Use Tariff: Evidence from Micro-Level Data

Elena Sestini¹

¹ *Department of Economics and Statistics, University of Siena*

Abstract

Residential electricity demand is shaped by a variety of behavioral, structural, geographic, and demographic factors and, as such, it is highly context dependent. This study provides a novel contribution by estimating the price elasticity of electricity demand for Italian households under time-of-use (ToU) tariffs. We first compute the total short-run elasticity of demand and then further investigate substitution patterns across hours of the day. Our results suggest that while electricity consumption is relatively inelastic to price changes, several contextual factors help explain consumption patterns. We also find some price responsiveness across peak and off-peak hours. Based on preliminary evidence, we believe that raising consumer awareness and providing clearer information on prices may be effective in promoting electricity savings and shifting demand away from peak periods.

3.1 Introduction

Electricity is a basic consumption good in daily life, as households rely on it for heating, cooling, cooking, lighting, food storage, and transportation (International Energy

Agency, 2018). Understanding the structural and behavioral dynamics behind electricity consumption is therefore essential to identify strategies for reducing demand and mitigating environmental impacts. This issue has become even more pressing in recent years, as extreme weather events have increased the need for cooling and heating services, while energy savings are widely recognized as a key component of climate change mitigation (IPCC, 2014). In addition to climate pressures, recent global developments have significantly reshaped electricity demand. Market liberalization reforms, geopolitical shocks such as the war in Ukraine, which altered prices of primary energy sources, and policy initiatives aimed at decarbonization have all renewed academic and policy interest in household electricity demand (Labandeira et al., 2017). Although a vast literature exists on residential electricity demand (Romero-Jordán et al., 2016), household responses to price and non-price factors remain uncertain due to significant heterogeneity across consumers and contexts.

Accurately assessing the mitigation potential of demand reductions is crucial for multiple stakeholders. Policymakers need reliable evidence to design pricing schemes that encourage efficient consumption. Utilities can adapt tariff structures to shift peak loads to off-peak periods, thus improving system efficiency. Consumers, if provided with transparent incentives, may adjust their behavior and achieve savings. In this regard, Demand Management Systems (DMS)¹ and time-varying tariffs play a central role. Recent proposals include dynamic pricing models, such as multi-TOU or critical peak pricing (CPP)², designed to redistribute consumption and prevent rebound peaks (Muratori and Rizzoni, 2015).

In addition to price incentives, behavioral factors are of particular importance for achieving energy savings (Zhu et al., 2018). Social scientists have offered contrasting perspectives: environmental psychologists emphasize the role of individual choices that often deviate from rational price–consumption responses (Pepper et al., 2009), while environmental sociologists highlight practices shaped by broader social norms and constraints (Gram-Hanssen, 2011). Electricity consumption reflects both fixed needs, which are satisfied at specific times of the day, and flexible uses, which may be shifted across hours. Investigating the substitutability between peak and off-peak consumption is therefore key to understanding household behavior (Bigerna and Bollino, 2015).

Elasticity estimates are a powerful tool for understanding the economic, environmental, and distributional implications of changing electricity prices. However, because electricity is typically sold under nonlinear tariffs and households differ greatly, the impact of price changes is far from uniform (Reiss and White, 2005). As tariff designs continue to evolve,

¹It refers to the set of tools, policies, and technologies adopted to shift or reduce electricity consumption by consumers, in order to increase energy efficiency.

²It represents a dynamic electricity tariff where the utility charges a much higher price during a limited number of critical peak events in a year. Outside those events, the consumer pays a more stable rate similar to ToU tariffs.

understanding their effects on consumer behavior remains a priority. Yet, empirical evidence for Italy is scarce, particularly under Time-of-Use (ToU) tariffs.

This paper addresses this gap by providing new evidence on residential electricity demand in Italy, exploiting a unique dataset of bi-monthly consumption data, obtained thanks to a collaboration with an Italian utility, covering around 2,400 households over 2023–2024. Building on household production theory and the empirical frameworks of Filippini (1995b, 2011), we estimate short-run residential demand using log-log specifications, adopting Fixed Effects (FE) and Random Effects (RE) models. The study addresses three main research questions: (i) How responsive is total household electricity consumption to price fluctuations under ToU tariffs? (ii) Which household and geographical characteristics are most relevant for policies aimed at reducing demand? (iii) To what extent do Italian households adjust their consumption across different time bands in response to price changes? Answering these questions allows us to anticipate the effect of a substantial change in price, such as an additional tax (e.g., a new tax). In particular, we measure the extent to which price measures are efficient in reducing peak load consumption.

Preliminary results indicate that, consistent with previous findings, total residential electricity demand in Italy exhibits limited sensitivity to price fluctuations. By contrast, several geographical and demographic characteristics appear to explain consumption patterns more effectively than price signals alone. Additional analysis across ToU band reveals asymmetric substitution patterns between peak and off-peak hours, with off-peak consumption being relatively more responsive to price changes.

This paper contributes to the literature in three distinct ways. First, it provides novel elasticity estimates for Italian households under ToU contracts, in a relatively scarcely studied context, using a unique panel dataset. Second, it contributes to non-linear price design literature, investigating substitution patterns across time bands and shedding light on the potential for peak load shifting. Third, it integrates consumption data with geographical (e.g., urbanization, altitude, etc.) and contractual characteristics (e.g., green tariff) to document heterogeneity across households, locations, and seasons. Together, these findings offer further understanding to the broader literature on electricity demand and residential consumption habits by offering evidence from a distinct institutional setting where ToU tariffs depends on national prices, monthly updated by the Gestore dei Mercati Energetici (GME)³.

The paper is structured as follows. We begin reviewing the relevant literature and illustrating the characteristics of the Italian electricity market, with particular attention to the institutional background, the structure of the ToU tariff under analysis, and the

³GME is controlled by Gestore dei Servizi Energetici (GSE S.p.A.), a company owned by the Ministry of Economy and Finance. It operates power, gas and environmental markets, in accordance with the guidelines given by the Ministry of Environment and Energy Security (MASE) and the regulatory provisions issued by ARERA.

mechanisms of price formation. Next, we describe data, models, and empirical methods employed in the analysis. Finally, we discuss the results and key takeaways from the study.

3.2 Related Literature

Starting from the 1970s, a large body of research has focused on estimating the price and income elasticity of electricity demand, using a wide variety of econometric techniques and data sources. Energy demand elasticities have been estimated for electricity, gas, and other energy products, typically treating household income and energy prices as the main drivers of consumption behavior (Silva et al., 2017). Despite this sizable empirical evidence, elasticity estimates remain highly sensitive to the dataset, geography, contract type, market structure and estimation method. Reported values range from close to zero, suggesting almost no consumer reaction, to above one, indicating strong sensitivity to price changes. This variation reflects the heterogeneity of energy markets and the fact that specific study settings often limit the generalization of results for policy decisions (Zhu et al., 2018). We place this paper at the intersection of three strands of research. First, we review previous contributions that estimate electricity demand in the short and long run in various countries using different types of data. Second, we focus the attention on studies related to the impact of ToU tariffs on residential electricity demand. Finally, we consider the literature on electricity demand estimation in the Italian market.

3.2.1 Short- and Long-Run Electricity Demand Elasticity

Empirical work estimating electricity demand elasticity to price and income can be broadly divided into two strands: studies using aggregated data at state, regional, or city level (Alberini and Filippini, 2011; Filippini, 2011; Azevedo et al., 2011; Blázquez et al., 2013; Boogen et al., 2017), and studies relying on disaggregated household-level data (Ito, 2014; Burns and Mountain, 2021; Boogen et al., 2021). Most contributions use panel data, though time series approaches (Pellini, 2021) and cross-sectional analyses (Burns and Mountain, 2021) also appear in the literature. Common methods include panel cointegration models, structural time series models (STSM), and other panel estimators. Control variables differ across studies, but typically include household income, electricity price, substitute energy prices as variables of interest; climate indicators, population, urbanization level, and, when micro data are available, household characteristics (size, age composition, education), and dwelling features, as additional variables and controls. Other contributions use large-scale simulations to analyze and quantitatively assess the impact of demand response programs using different electricity price structures like ToU, Flat price and Critical Peak Pricing (Muratori and Rizzoni, 2015).

A number of meta-analyses summarize elasticity estimates. For instance, Brons et al.

(2008) and Havranek et al. (2012) report short-run price elasticities for energy products in the range of -0.09 to -0.76 , and long-run elasticities between -0.31 and -1.16 . Zhu et al. (2018), collecting studies from 1950 to 2017, find short-run price elasticities ranging from -0.94 to 0.61 , with a mean of -0.22 , and long-run elasticities averaging -0.57 . They highlight that elasticities differ depending on data frequency: monthly data typically yield more elastic responses than quarterly or annual data, likely due to seasonality, climate effects, and holiday patterns. Similarly, Labandeira et al. (2017) confirm a short-run elasticity of about -0.21 and a long-run elasticity of about -0.61 . They show that results depend strongly on model choice, consumer type, data period (before or after energy crises), and country context.

Panel data analyses often highlight seasonal differences. For example, Hung and Huang (2015) use county-level monthly data for Taiwan (2007–2013) and find that both demand and speed of adjustment differ between summer and non-summer months, with lower elasticities during hot summers. Likewise, Paul et al. (2009), using a partial adjustment model for U.S. regions and seasons, show higher elasticities in summer than in winter (short-run -0.15 vs. -0.11 ; long-run -0.52 vs. -0.32).

Other studies emphasize heterogeneity across households or regions. For instance, Bernard et al. (2011) use four independent household surveys from Quebec (1989–2002) and find a short-run elasticity of -0.51 and a long-run elasticity of -1.32 . Reiss and White (2005), with household-level data for California (1993–1997), estimate -0.39 . By contrast, Holtedahl and Joutz (2004), using Taiwanese annual aggregates (1955–1995), report almost constant elasticities of about -0.15 in both the short and long run.

Similar to Boogen et al. (2017) we build on household production theory. Boogen et al. (2017) uses a static model of short run elasticity and a dynamic model of residential electricity consumption incorporating a correction introduced by Kiviet for long-run elasticity in Switzerland. They estimate a short run elasticity is around -0.30 while the latter is around -0.60 . They use a unique survey of households in Switzerland that includes detailed information on a household’s annual electricity consumption, residential and socio-demographic characteristics, its stock of appliances, and its use of these appliances. They also adopt an IV approach to correct for endogeneity introduced by average prices.

We also carefully reviewed studies that examine which type of price should be used. A well known contribution comes from Ito (2014), that, analyzing household-level administrative panel data from Japan, show that consumers respond mainly to average prices rather than marginal or expected marginal prices, reflecting the complexity of non-linear electricity tariffs.

Other approaches use panel cointegration techniques. Nakajima (2010), for Japan, apply panel unit root tests and cointegration analysis on data from 46 prefectures (1975–2005). They find that electricity is income inelastic (not a necessity good) but price elastic, with estimates above one. The same result is showed by Narayan et al. (2007), that also apply

panel unit root and cointegration tests to G7 countries. Miller and Alberini (2016) show that while elasticities appear relatively stable over time, cross-sectional splits suggest that higher average prices are associated with lower elasticities.

Earlier seminal studies, like Bohi and Zimmerman (1984), report a wide range of short-run elasticities (-0.07 ; -0.88) and long-run elasticities (-0.18 ; -4.56). More recent contributions include Schulte and Heindl (2017), who apply a quadratic expenditure system to German expenditure data (1993–2008) and find electricity price elasticities around -0.43 to -0.50 .

3.2.2 Estimation of Electricity Demand Elasticity under ToU Tariffs

A body of research relevant to our study estimates demand elasticity under ToU electricity tariffs. These contracts are designed within DMS to shift consumption away from congested peak hours toward off-peak periods, thus improving system efficiency and reducing stress on the grid. However, as Muratori and Rizzoni (2015) highlight, finding an effective balance in tariff design is not straightforward: while price signals can influence consumption behavior, poorly calibrated tariffs risk creating new rebound peaks in other time slots. Several studies conducted on ToU tariffs rely on experimental settings (Faruqui and Malko, 1983; Aubin et al., 1995; Gyamfi et al., 2013; Kim et al., 2022; Schittekatte et al., 2024).

Given the nature of our study, we decided to focus mainly on previous works based on observational data. Two particularly relevant contributions for our methodological framework are Filippini (1995b) and Filippini (2011), both analyzing residential electricity demand in Switzerland by time of day. Filippini (1995b), uses micro-level household data and a log-linear system of stochastic demand equations for peak and off-peak periods and estimate short-run own-price elasticities of approximately -0.60 during peak and -0.80 during off-peak, and long-run elasticities of -0.71 during peak and -1.92 during off-peak. In a subsequent study, Filippini (2011) employs aggregate data from 22 Swiss cities over the period 2000–2006. Building on household production theory and applying a log-log specification estimated with Random Effects (RE) and LSDVC⁴ to address endogeneity through lagged consumption, they find short-run own-price elasticities between -0.77 and -0.85 (peak) and between -0.65 and -0.75 (off-peak) while long-run elasticities range from -1.60 to -2.26 (peak) and -1.27 to -1.65 (off-peak). Their results provide evidence of substitution between peak and off-peak consumption. They also estimate

⁴LSDVC is a bias-corrected FE estimator for dynamic panel data models with small time dimension T and large cross-section N . It applies analytical or bootstrap corrections to LSDV to reduce the finite-sample bias of the lagged dependent variable coefficient. It is an alternative to GMM methods.

short-run cross-price elasticities between 0.79 and 0.91 (peak) and between 0.40 and 0.36 (off-peak); long-run estimates vary between 1.76 and 2.31 (peak) and between 0.68 (peak) and 0.91 (off-peak). A complementary line of research adopts demand system approaches, particularly the Almost Ideal Demand System (AIDS)⁵ model (Deaton and Muellbauer, 1980; Rossi, 1988). Filippini (1995a) apply this framework to residential electricity demand in Switzerland, finding not only significant own- and cross-price elasticities (with peak and off-peak consumption acting as substitutes) but also evidence that household characteristics matter: the presence of children increases consumption during peak hours, while the use of electric boilers reduces it, as they can decide to turn on the electric boiler during the off-peak period and turn it off during the peak period.

More recently, Burns and Mountain (2021) estimate the elasticity of substitution between peak and off-peak periods using Australian household data. They exactly study how the variation of peak price relative to off-peak price affects consumption during peak hours with respect to consumption in off-peak hours (and vice versa), which is a slightly different concept compared to cross price elasticity. Based on cross-sectional evidence from 6,957 households on ToU tariffs, they compute substitution elasticity as the ratio of peak to off-peak prices. Their results suggest that a 1% increase in peak prices relative to off-peak prices induces only a 0.2% shift in consumption from peak to off-peak hours. The authors conclude that despite retail market liberalization, the installation of smart meters, and policy support for ToU tariffs, there has been little measurable improvement in consumers' responsiveness to time-varying electricity prices in Australia.

3.2.3 Italian Electricity Demand

One of the major contribution of this paper is to explore Italian residential electricity demand across different hours of the day. While several studies exist for both European (Filippini, 2011; Blázquez et al., 2013) and non-European countries Filippini and Pachauri (2004); Alberini and Filippini (2011); Hung and Huang (2015); Burns and Mountain (2021), evidence for the Italian market remains scarce, particularly in the context of ToU tariffs. One of the plausible reasons is the difficulty of accessing residential data, even at aggregated levels, due to the limited availability of public statistics.

An important contribution is provided by Faiella and Lavecchia (2021), who investigate households' energy demand and expenditure using survey-based microdata covering Italian households over 1997–2018. They estimate average price elasticities of -0.28 for electricity, -0.43 for heating and for transport. In simulation exercises, they show that increases in

⁵It is a consumer demand model used to study consumer behavior;it gives an arbitrary second-order approximation while satisfying key theoretical restrictions (adding-up, homogeneity, and symmetry).

energy prices driven by carbon taxation would reduce energy demand across Italian households, but with regressive distributional impacts, as poorer households would experience a stronger compression in consumption while facing relatively larger expenditure increases. Similarly, Bardazzi and Pazienza (2020) estimate national-level elasticities of -0.70 for electricity and -0.62 for natural gas, and highlight significant regional heterogeneity, with stronger responses in Central and Southern Italy. They also find that electricity demand peaks when the household head is around 50 years old, while natural gas demand continues to increase with age, reflecting more time spent at home.

More closely related to our focus are studies that measure intraday electricity demand under ToU contracts. Torriti (2012) conduct a quasi-experimental before and after analysis of Italian ToU tariffs, using smart-meter data with 15-minute frequency for 1,446 households in the Province of Trento (Northern Italy). They find significant load shifting in the morning peak: overall demand increased by 13.7% under ToU compared to flat tariffs, although average expenditure decreased by 2.2%. The morning peak shifted earlier, from 8:00 to 6:45), with lower intensity and a smoother profile, while a new, higher evening peak emerged after 9:00 pm. Interestingly, consumption during dinner and TV hours remained largely unresponsive to price signals. Thus, given these outcomes, ToU tariffs did not reduce total demand or solve peak-load problems: they provided modest savings but increased overall consumption.

Bigerna (2012) also show that the price effect on electricity demand varies across the day and across geographical zones, with estimated elasticities ranging between -0.03 and -0.10 . A more structural approach is taken by Bigerna and Bollino (2015), who estimate a complete system of hourly electricity demand using market bid data from January 2010 to December 2011. They adopt a two-stage maximization framework: in the first stage, electricity demand is aggregated into day versus night, while in the second stage it is disaggregated into 12 hourly blocks for day and night. Their results indicate that daily hours are moderately a luxury good, while night-time demand behaves as a normal good. Price responsiveness is limited but varies by period: conditional own-price elasticities range from -0.05 to -0.12 at night and from -0.09 to -0.12 during the day, while unconditional elasticities are somewhat smaller in absolute terms (from -0.04 to -0.07 at night and -0.07 to -0.10 during the day). Substitution patterns reveal that daily hours are substitutes for night-time demand but also complements to some morning hours, while late-night hours show substitutability within themselves.

Overall, the Italian evidence suggests that residential electricity demand is generally inelastic, though with relevant variation across hours, regions, and household characteristics. ToU tariffs appear to induce some load shifting, but without reducing total demand or peak stress on the grid.

3.3 Italian Electricity Market

3.3.1 Market structure

The Italian electricity market was established following the so-called Bersani Decree (Legislative Decree No. 79 of March 16, 1999), which transposed the first EU Directive on the creation of an internal energy market and marked the start of electricity sector liberalization in Italy. Trading on the electricity exchange, and consequently the launch of the first operational phase of the market, began five years later. The primary objectives were to promote competition in electricity generation and trading activities, based on neutrality, transparency, and objectivity, through the creation of a centralized marketplace, and to ensure the economic management of an adequate supply of dispatching services.

The responsibility for managing the Italian electricity market in line with these objectives lies with the Gestore dei Mercati Energetici (GME), pursuant to Article 5 of Legislative Decree 79/99. GME is a company wholly owned by the Gestore dei Servizi Energetici (GSE) and operates under a specific set of market rules, which define market functioning and participation procedures. These rules are drafted by GME and approved by the Ministry of Economic Development, subject to the opinion of the Regulatory Authority for Energy, Networks and Environment (ARERA).

From a temporal perspective, the Italian wholesale electricity market is divided into sections and subsections. The first distinction is between:

- Day-Ahead Market (MGP): This market hosts the largest number of transactions for the purchase and sale of electricity. It is called “day-ahead” because trading must take place the day before, for each hour of the following day, relative to the physical production of electricity.
- Intra-Day Market (MI): The Intra-Day Market allows market participants to adjust the positions defined in the Day-Ahead Market in case of excess or shortage of electricity production compared to consumption.
- Daily Products Market (MPEG): This market is dedicated to the trading of daily products with an obligation for physical delivery of electricity.
- Ancillary Services and Balancing Market (MSD): This is the tool available to Terna S.p.A. to manage the balance of electricity flows on the national grid, where Terna acts as the central counterparty⁶.

Additional information about the market structure are included in Appendix ???. In addition to the Spot and Forward Markets, the Italian electricity market includes the

⁶See Appendix C.1 for further details.

Piattaforma per la Consegna Derivati Energia (CDE), which records transactions of energy financial derivatives traded on the IDEX.

Furthermore, the national markets, MGP and MI, are integrated respectively with European initiatives established for the implementation of single cross-border market coupling throughout Europe.

3.3.2 Time of Use (ToU) Tariff

The type of contract under analysis is a ToU tariff offered by the reference company. This contract is based on three distinct time bands. Under such pricing scheme, the electricity price varies depending on the time of day at which the household consumes electricity, according to the following schedule:

- F1 (peak hours): 8 a.m. - 7 p.m. on weekdays;
- F2 (mid-peak hours): 7 a.m. - 8 a.m. and 7 p.m.-11 p.m. on weekdays, and 7 a.m. - 11 p.m. on Saturdays;
- F3 (off-peak hours): 11 p.m. - 7 a.m. every day and all hours on Sundays.

A small subset of contracts in our dataset uses a bi-hourly scheme, where F2 and F3 are combined into a single off-peak block. For the purposes of time-band analysis, we restrict the sample to households with tariffs that separate the three bands (F1, F2, and F3).

As shown in Figure C.1 in Appendix, the price per kilowatt-hour under the ToU scheme we are studying consists of a fixed and a variable component:

$$\text{Price}_t = \underbrace{c_{\text{fixed}}}_{\approx 0.03 \text{ €/kWh}} + \underbrace{c_{\text{variable},t}}_{\text{based on PUN}_t} .$$

The final bill also includes other fixed charges (e.g., for commercialization, metering, and activation). For our purpose, we ignore any fixed expense⁷ as we assume that consumption is not affected by fixed amounts which are common to every type of contract. We focus on the variable component, that in this type of contract represents the National Unique Price (PUN in Italian).

The PUN is the average of seven zonal prices weighted by electricity volumes traded in each geographical zone. It is determined in the Italian wholesale Day-Ahead Market, a centralized auction where electricity for the following day is traded by hourly intervals. It

⁷This information is not available in our dataset for all previous tariffs, however we checked on the website and it was the same for all available contracts at date.

is published daily by the GME and expressed in €/MWh⁸. Formally, the PUN at hour h is calculated as:

$$\text{PUN}_h = \frac{\sum_z P_{z,h} \cdot Q_{z,h}}{\sum_z Q_{z,h}},$$

where $P_{z,h}$ is the zonal price in zone z at hour h , and $Q_{z,h}$ is the corresponding traded quantity.

The GME publishes monthly average PUN values for each time band at the end of every month. Consumers do not observe the current month's price in real time, as it is only computed once the month has ended. Therefore, we assume that households react to the previous month's prices, which they can observe in their latest electricity bill. Although daily prices are continuously updated on the GME website, it is uncommon for average consumers to monitor them regularly. Moreover, even if they did, these values would not correspond to the final price on which their bill is calculated.

3.3.3 An overview of demand and prices across time bands

ToU electricity contracts have been introduced to provide incentives to consumers to shift demand away from peak hours, in order to reduce costs and increase efficiency. Peak demand periods put pressure on the grid and might require activating expensive plants. By increasing the prices during peak hours and lower prices off-peak, consumers are encouraged to move flexible consumption to off-peak times.

Time band F1 typically corresponds to peak hours and is generally expected to have the highest demand charges. F3, by contrast, represents off-peak periods, while F2 falls in the mid-peak range. Prices across these bands are typically set to reflect peak and off-peak patterns. Figure 3.1 graphically represents price trends over the two years of interest. Surprisingly, there are periods in which the national price PUN is higher in F2 than in F1. The reason is that prices are determined by the interaction between electricity demand and supply in the wholesale market. During daylight hours, the large contribution of renewable generation, especially photovoltaic, lowers prices in F1 by shifting the aggregate supply curve downward, as solar power has a nearly zero marginal cost. In spring and summer, F2 prices can even exceed F1 prices due to the longer daylight period⁹. In addition, higher and inelastic demand may occur in F2, often linked to activities such as cooking, washing,

⁸By 9 a.m. on the day before delivery, all market participants, including producers on the supply side and wholesalers, traders, and retailers on the demand side, submit price-quantity offers and bids for each of the 24 hours of the next day. These are organized by bidding zones (six in Italy). At 11 a.m., the EUPHEMIA algorithm simultaneously matches supply and demand across all coupled European markets. For each hour and zone, it determines the marginal price: the price of the last accepted offer that balances supply and demand. The hourly Italian PUN is then computed as a volume-weighted average of the zonal prices.

⁹In the spring and summer months, F2 prices can exceed F1 prices, reflecting the higher number of daylight hours compared to winter.

and cleaning for a large portion of households at the same time, driving prices above those in F1 in certain periods.

Electricity demand's trends are showed in Figure 3.2. Demand increases in both the hottest and coldest months, with major peaks in summer (July and August). This pattern can be explained by two main factors. First, self-production through photovoltaic systems reduces households' reliance on the grid during peak hours. Second, off-peak bands include weekends, when a larger share of the population stays at home, increasing residential electricity use.

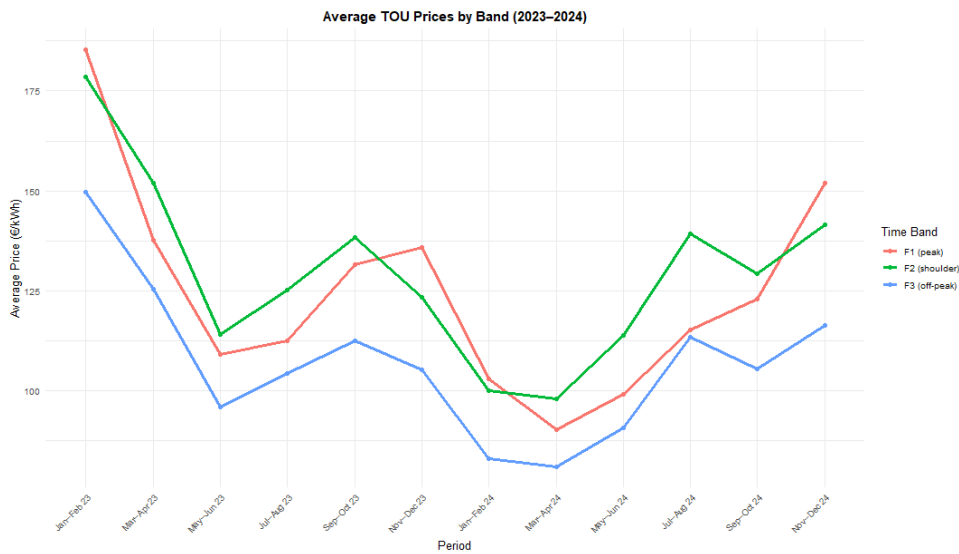


Figure 3.1: Bi-monthly PUN by time band 2023-2024

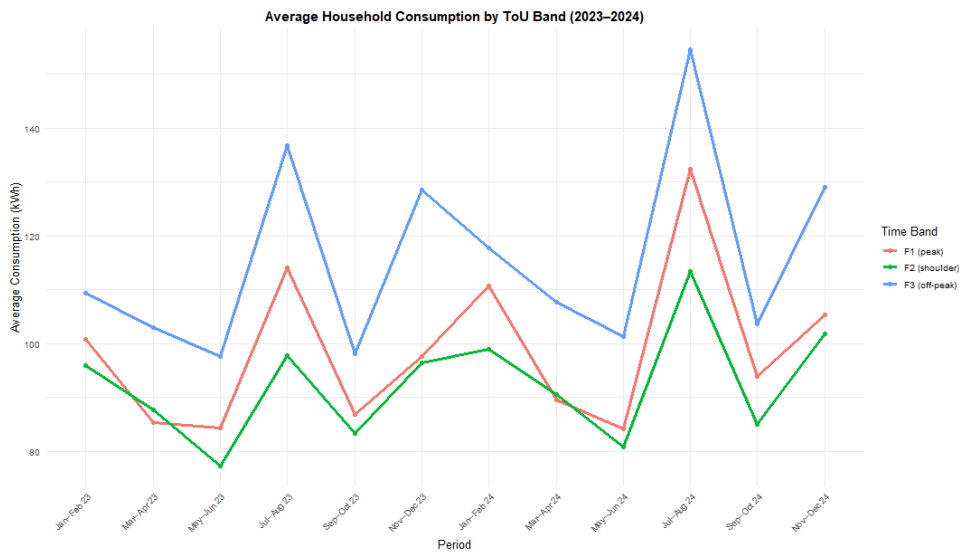


Figure 3.2: Bi-monthly electricity demand by time band 2023-2024

3.4 Modelling Electricity Demand

3.4.1 Theoretical framework

We model the demand for residential electricity using Household Production Theory, which is common in literature (Becker, 1965; Flaig, 1990; Filippini, 1995b; Hung and Huang, 2015; Boogen et al., 2017). We assume that households are not only direct consumers of electricity but they buy a certain amount of it to produce other services from which they perceive a certain utility. Households are then represented as producers of “commodities” starting from an initial amount of electricity they decide to acquire, such as thermal comfort, lighting, or food preservation. For this reason, electricity is not consumed directly, rather, it is combined with other input, such as appliances, dwelling characteristics, and household time, to produce these commodities.

From a theoretical point of view, we assume Z_{jt} denote the quantity of commodity j produced at time t by the household according to a certain household production function:

$$Z_{jt} = f(Q_{jt}, X_{jt}, T_{jt}, I_{jt}),$$

The production function is based on multiple inputs. Q_{jt} is the electricity used for commodity j at time t , X_{jt} is a vector of other inputs, T_{jt} denotes household time and I_{jt} is the average income¹⁰. In the first stage, the household chooses the cost-minimizing combination of inputs for each commodity given input prices (including the electricity tariff) and the available technology. In the second stage, the household allocates income across commodities to maximize utility, subject to the budget constraint, where the “shadow price” of electricity reflects both the market tariff and the household’s technology.

Within this framework, the derived demand for electricity in each time band depends on: (i) the own price of electricity, which affects the marginal cost of producing energy-intensive services (own-price elasticity); (ii) the prices in alternative time bands, which influence substitution patterns (cross-price elasticities); (iii) weather conditions, which shift the marginal utility of comfort-related services; (iv) household and geographical characteristics, which determine the production technology; (v) average household income at municipality level.

To model our empirical application we use a log-log model, as it is common to estimate demand elasticity. Following Boogen et al. (2017) we identify the electricity demand for household i at time t as follows:

¹⁰In similar studies they include house’s characteristics, like the number of rooms, appliances, members of the family. In this study we only have the availability of contract type information and geographical characteristics

$$\ln Q_{it} = \beta_0 + \beta_1 \ln P_{t-1} + \beta_2 G_i + \beta_3 D_{it} + \beta_4 T_{it} + \beta_5 I_{it} + \varepsilon_{it} \quad (3.1)$$

where the dependent variable is the logarithm of total electricity demand¹¹ for household i at time t (here bi-monthly consumption). The coefficient of interest is β_1 , that represents the price elasticity, G_i is a vector of geographical characteristics, D_{it} represents contract characteristics, T_{it} is a vector of weather controls, and I_{it} is the average income computed at municipality level for individual i at time t . Given the monthly frequency of the data, we can examine seasonal patterns in electricity consumption. Since electricity prices and weather conditions vary throughout the year, it is reasonable to expect that residential consumption differs between summer and non-summer months, and we check for seasonality (Hung and Huang, 2015). Note that, as price component, we use the lagged value of the price variable (the PUN described earlier) for each time band, rather than the contemporaneous price at time t . This choice reflects the fact that the monthly average PUN is computed only at the end of each month. Given the structure of the electricity market, current prices are not known to households when they make their consumption decisions at time t . Therefore, we assume that residential demand responds to the price of the previous month ($t-1$), which is reported in the most recent electricity bill¹².

In the second part of our analysis, we aim to provide a more detailed interpretation of elasticity of demand across time bands. Following Filippini (1995b) and Filippini (2011), we specify the two following models for each time band¹³:

$$\ln Q_{it,F1} = \beta_0 + \beta_{F1} \ln P_{t-1,F1} + \beta_{F2} \ln P_{t-1,F2} + \beta_3 G_i + \beta_4 D_{it} + \beta_5 T_{it} + \varepsilon_{it} \quad (3.2)$$

$$\ln Q_{it,F1} = \beta_0 + \beta_{F1} \ln P_{t-1,F1} + \beta_{F3} \ln P_{t-1,F3} + \beta_3 G_i + \beta_4 D_{it} + \beta_5 T_{it} + \varepsilon_{it} \quad (3.3)$$

where the logarithm of electricity consumption in each time band (either $\ln Q_{it,F1}$, or $\ln Q_{it,F2}$, or $\ln Q_{it,F3}$) is regressed on the logarithm of the corresponding price variable. Here β_{F1} represents the own-price elasticity for electricity consumption in F1 time band, and β_{F2} is the cross price elasticity of demand in F2 band. We do not add the second cross price elasticity (here the logarithm of price at $t-1$ of F3) and we prefer to estimate two separated models because the high correlation between prices distorts the estimates. For this reason we keep one of the three as a baseline. Next, following Hung and Huang (2015), we also

¹¹Here there is not distinction across time bands.

¹²Bigerna and Bollino (2015) support that average Italian residential consumers are supposed to be informed of different prices across different hours, due to massive information since the beginning of the market liberalization in 2004, and secondly, the Energy Authority ruled that the bi-monthly electricity bill for the final consumer must explicitly highlight the amount of the unbalancing costs.

¹³This specification holds for demand in F1; we test the same model for F2 and F3 respectively.

insert temperature variables ($\ln Temp$, $\ln TempMin$, $\ln TempMax$)¹⁴, contract info D_{it} , temporal dummies T_{it} , geographic controls G_i . In addition, we test another specification including the income I_{it} ¹⁵. This specification captures the short-run responsiveness of residential electricity demand to price signals under this specific tariff structure.

We can also express price elasticity, ε_E , as follows:

$$\varepsilon_E = \frac{\partial Q_i^E}{\partial P^E} \cdot \frac{P^E}{Q_i^E}$$

It represents the percentage change in electricity demand for a 1% change in price.

3.4.2 Which Price to Use?

The definition of the relevant price variable is a central issue in the estimation of electricity demand elasticities, and it is highly context-dependent. A long debate in the literature concerns whether to use the *marginal price*, particularly relevant under two-part tariffs¹⁶, where the marginal rate influences consumption decisions independently of total usage (Nordin, 1976), or the *average price*, typically computed as the ratio between the total electricity bill (including fixed charges) and total consumption (Shin, 1985).

The main advantage of using the marginal price lies in its exogeneity: while the marginal price can influence consumption, consumption itself does not affect the marginal rate. On the other hand, using the average price introduces potential endogeneity, since the average price is mechanically affected by the quantity consumed, also includes a fixed fee that produces biased estimates, and introduce simultaneity issues. Several studies use both prices (Taylor et al., 2004; Boogen et al., 2021). This bias is highlighted, for instance, in Taylor et al. (2004), who show that average-price specifications in residential water demand lead to biased elasticity estimates.

By contrast, several studies argue that the average price may better capture consumer decision making, as it is the price which can be easily retrieved from the bill, and thus, salient to households. Electricity bills usually contains several information, making it difficult for consumers to understand the nonlinear structure of marginal prices, which might change by time of the day or by the amount consumed. Shin (1985) and Borenstein (2009) emphasize the behavioral relevance of average price, and Ito (2014), using household-level panel data from administrative records, provides evidence that consumers respond primarily to average rather than marginal or expected marginal prices (under non-linear

¹⁴In the literature, cooling degree days (CDD) and heating degree days (HDD) are two commonly applied weather variables. However, these two variables are not available to us.

¹⁵At this time we proxy this variable with its lag $I_{i,t-1}$, since it is not yet available for year 2024

¹⁶It is a pricing scheme consisting of a fixed charge plus a variable charge based on actual consumption, that, contrary to ToU tariffs, remains the same across different hours.

pricing structure, where the price changes depending on the amount consumed).

In the Italian context, we adopt as our best measure the *Prezzo Unico Nazionale* (PUN), a national wholesale price component that directly affects the variable part of retail tariffs. Fixed charges are ignored, as they are largely homogeneous across households. We also control for the presence of green contracts, which may involve slightly higher charges that are not directly observable to us. Consistent with the regulatory structure of the Italian electricity market, we treat households as price-takers and interpret the PUN as the relevant marginal price. However, since households do not know *ex ante* the precise price for each time band, we use the lagged value of the PUN as a proxy for the price signal actually perceived by consumers through their previous bills.

Given these considerations, we expect limited consumer responsiveness to price changes. This is due both to the lack of transparent knowledge about how tariffs are calculated and to the strong correlation between prices across time bands, which reduces the scope for meaningful substitution of consumption from peak to off-peak hours.

3.5 Data and Methods

The data come from multiple sources. The primary dataset was obtained through a collaboration, and under a confidentiality agreement, with the Italian energy utility Estra S.p.A., which operates on a national scale. It contains bi-monthly billing and electricity consumption records over a two-year period, from January 2023 to December 2024. The dataset includes (i) total electricity consumption and billed amounts; (ii) consumption disaggregated by time slots, corresponding to the tariff bands F1, F2, and F3; (iii) contract characteristics, such as green versus non-green contracts, fixed price versus ToU tariffs, the number of tariff bands, and the power allocated to each contract.

All data are anonymized and do not contain households' demographic details¹⁷. The only available geographic information refers to the province and municipality of each family. The initial sample consists of approximately 2,400 unique IDs, each corresponding to a different household. This is a representative sample of the Italian population customer base, randomly drawn from a broader pool of over 90,000 customers of Estra S.p.A. The random selection followed a two-stage stratified sampling design: stratification was first applied based on contract type (fixed vs. ToU, green vs. non-green), and then by geographic distribution at the municipal level according to ISTAT's regional classification¹⁸.

To complement this dataset, we add price data from the GME. Specifically, we collect the PUN, which represents the variable component of electricity prices. We include both

¹⁷Due to privacy issues we are not allowed to get access to additional households informations like stock of appliances, number of members of the family, number of rooms, etc.

¹⁸The stratification procedure is available upon request.

(i) the average monthly PUN (undifferentiated across time slots), and (ii) the average PUN by time band, already computed and available on the GME website.¹⁹

We further enrich the dataset with municipality-level geographical characteristics from ISTAT²⁰, including (i) coastal location (dummy); (ii) altitude zone; (iii) urbanization level; (iv) territorial surface area; and (v) legal and resident population.

Next, we incorporate weather and temperature data at the provincial level, obtained from the Italian weather forecasting provider 3B Meteo²¹. We collected monthly averages for (i) average temperature, (ii) minimum temperature, and (iii) maximum temperature. Last, we include annual average income at municipality level from the Ministry of Economics and Finance (MEF)²²

3.5.1 Dependent Variables

To measure the elasticity of total residential electricity demand, we use the logarithm of electricity consumption as the dependent variable. Specifically, $\ln Q_{i,t}$ denotes the logarithm of total electricity consumption by household i at time t , without distinguishing between time bands. This specification is employed when estimating the elasticity of overall bi-monthly electricity demand in Equation 3.1.

For a more granular analysis, we also estimate the elasticity of electricity demand to price across time bands. In this case, $\ln Q_{F1,i,t}$ represents the logarithm of electricity consumption by household i at time t during time band F1. Similarly, $\ln Q_{F2,i,t}$ and $\ln Q_{F3,i,t}$ correspond to consumption during bands F2 and F3, respectively. This is the response variable used in specifications similar to Equations 3.2 and 3.3.

3.5.2 Independent Variables and Controls

The main variable of interest is the logarithm of price, which is used to compute price elasticity. We adopt the *Prezzo Unico Nazionale* (PUN) as the reference price, representing the variable component of the electricity tariff. The term $\ln P_{i,t-1}$ denotes the logarithm of the average PUN over the previous two months for household i at time t . We use the lagged price rather than the contemporaneous one, since the monthly average PUN is computed and published by the GME at the end of each month. Consequently, it is not available to households at the time of consumption. Therefore, we assume that household consumption decisions are responsive to the price observed in the previous electricity bill.

¹⁹Monthly PUN values were downloaded from the following link.

²⁰Data retrieved from ISTAT webpage: see this link.

²¹You can check the website at this link.

²²See the MEF webpage at this link.

In line with the disaggregation of consumption by time band, we also compute the logarithm of the average price for each time slot. Specifically, $\ln P_{F1,i,t-1}$ represents the average lagged price during peak hours (F1), while $\ln P_{F2,i,t-1}$ and $\ln P_{F3,i,t-1}$ correspond to the lagged average prices in the F2 and F3 bands, respectively.

Regarding contract characteristics, we consider: (i) $green_{i,t}$, a dummy variable equal to 1 if the household has a green electricity contract, and 0 otherwise; (ii) $power_{i,t}$, representing the contractual power, acts as a proxy for both household size and the degree of building electrification (e.g., a greater number of electric appliances implies a higher contractual power).

Weather controls are represented by: (iii) $Temp_{i,t}$, the average monthly temperature (in degrees Celsius) in municipality i at time t ; (iv) $TempMin_{i,t}$ and $TempMax_{i,t}$, the average minimum and maximum monthly temperatures, respectively.

We also include several geographic variables available at the municipality level, which remain constant over time: (v) $Altitude\ Zone_i$, a categorical variable with five levels derived from altitude; (vi) $Population_i$, the number of residents in municipality i ; (vii) $Urbanization_i$, a three-level categorical variable indicating the degree of urbanization; (viii) $Coastal_i$, a dummy equal to 1 if the municipality is located on the coast; (ix) $Island_i$, equal to 1 if the municipality is on an island; (x) $Surface_i$, the surface area of municipality i (in km^2); (xi) $Geographical\ Area_i$, a categorical variable that takes value 0 for municipalities in Northern Italy, 1 for Central Italy, and 2 for the South or Islands. Last, we include (xii) $Average\ Income_{i,t-1}$, that represents annual average income at municipality level. It is computed as the ratio between taxable income and the total number of taxpayers.

3.6 Results

3.6.1 Descriptive Statistics

Table 3.1 reports summary statistics of the main variables. The initial sample included about 2,400 unique IDs but we retained only those households enrolled in a ToU tariff with three time bands. Subsequently, we filtered out vacant dwellings. Identifying active households is a common challenge in studies based on raw household-level electricity data (Reiss and White, 2005). To avoid including uninhabited homes, we first removed all IDs with total consumption equal to zero in at least eight out of twelve bimonthly billing periods, and then we excluded remaining observations with zero consumption. The final sample consists of 1,738 households.

On average, households consume approximately 331 kWh of electricity over a two-month billing cycle, with substantial variation across the sample. When disaggregating consumption by time band, average demand is 107 kWh in F1 (peak), 99 kWh in F2, and 124 kWh in F3. Off-peak consumption (F2+F3) therefore accounts for about 223 kWh,

indicating that the majority of electricity use occurs outside of peak hours. This pattern is consistent with the fact that part of daytime residential demand is satisfied by renewable generation, and that off-peak periods include weekends, when households typically spend more time at home.

Turning to electricity prices, the average national wholesale price (PUN) is 119 €/MWh, with notable variability across bands. Peak prices (F1) are substantially higher, averaging 140 €/MWh, while F2 is slightly lower (142 €/MWh), and F3 shows the lowest price (116 €/MWh). Geographical characteristics exhibit considerable heterogeneity. Households are distributed across municipalities located at an average altitude of 249 meters, ranging from coastal towns to mountain areas above 2000 meters. Approximately 38% of households are located in coastal municipalities, while 5% live on islands. Urbanization levels are well distributed across the three categories used in the analysis, with an average of 1.97 suggesting that many households reside in medium-density areas. Population size varies widely across municipalities, ranging from small rural areas with fewer than 200 residents to large cities with more than 850,000.

Table 3.1: Descriptive Statistics

Variable	Mean	SD	Min	Max	N
Consumption F1	107.164	91.271	1.002	2,909.591	19,428
Consumption F2	99.252	80.471	1.004	2,112.554	19,428
Consumption F3	124.515	112.382	1.014	2,986.038	19,428
Consumption F2+F3	223.767	188.016	2.256	4,945.680	19,428
Total Consumption	331.287	272.025	3.849	7,855.271	19,428
Total Amount Paid	121.324	68.255	10.020	791.890	18,329
Average Income _{t-1} (municipality)	20,326.638	3,933.114	11,344.903	36,252.595	19,428
P _{F1,t-1}	140.410	68.411	90.250	360.730	19,428
P _{F2,t-1}	142.373	53.288	97.995	309.960	19,428
P _{F3,t-1}	116.720	41.442	80.930	244.940	19,428
P _{t-1}	119.554	25.657	87.833	190.356	19,428
Green	0.523	0.499	0.000	1.000	19,428
N. Residents	79,834.077	139,343.499	102.000	851,199.000	19,390
Altitude Zone	3.135	1.459	1.000	5.000	19,390
Altitude (m)	248.032	278.589	0.000	2,035.000	19,390
Coastal	0.379	0.485	0.000	1.000	19,390
Island	0.045	0.207	0.000	1.000	19,390
Urbanization	1.950	0.752	1.000	3.000	19,390
Avg. Temperature	17.215	6.531	3.300	29.750	19,428
Min Avg. Temperature	13.389	6.211	-0.650	25.600	19,428
Max Avg. Temperature	21.516	6.865	8.900	37.100	19,428

Climate variables highlight significant seasonal variation, with an average temperature of 17.2 °C across the sample period. Minimum monthly averages can fall below zero, while maximum averages exceed 37 °C, reflecting the diverse climatic conditions across Italy. These fluctuations are expected to play an important role in driving seasonal electricity demand, particularly for heating and cooling.

Finally, households pay on average €121 per billing cycle, though with considerable variation.

Overall, the descriptive statistics confirm the strong heterogeneity in consumption

patterns, geographical settings, and climatic conditions across Italian households.

3.6.2 Total Electricity Demand

In this section, we present the estimation results of the model described in Equation 3.1. Overall, residential electricity demand in Italy appears highly inelastic to price changes. Across both the FE and RE specifications reported in Table 3.2, the elasticity with respect to the lagged price is approximately -0.16 , implying that a 1% increase in price leads to an average reduction of about 0.16% in electricity consumption.

Weather variables exhibit heterogeneous effects. Higher average temperatures are associated with greater electricity use, while increases in minimum daily temperatures reduce consumption by around 0.1%. Conversely, higher maximum daily temperatures are linked to lower demand, suggesting that extreme heat does not translate into proportionally higher residential electricity use.

Table 3.2: FE/RE Estimates – Total Electricity Consumption Model

Variable	FE	RE
<i>Intercept</i>	—	6.079*** (0.299)
<i>Price Elasticity</i>		
$\ln PUN_{t-1}$	-0.156*** (0.035)	-0.142*** (0.036)
<i>Weather Controls</i>		
$\ln(\text{Temp})$	0.410** (0.200)	0.487** (0.201)
$\ln(\text{TempMin})$	-0.097* (0.057)	-0.114* (0.058)
$\ln(\text{TempMax})$	-0.568*** (0.176)	-0.638*** (0.176)
<i>Geographic Controls (RE only)</i>		
Altitude Zone (level 2)	—	0.219** (0.107)
Altitude Zone (level 3)	—	0.405*** (0.065)
Altitude Zone (level 4)	—	0.430*** (0.078)
Altitude Zone (level 5)	—	0.557*** (0.077)
Urbanization (level 2)	—	-0.018 (0.053)
Urbanization (level 3)	—	-0.340*** (0.059)
Coastal (dummy)	—	-0.299*** (0.066)
Island (dummy)	—	0.033 (0.102)
Power (kW)	—	0.166*** (0.016)
Green contract	—	0.072*** (0.021)
Area Geographical (1)	—	0.155** (0.063)
Area Geographical (2)	—	0.235*** (0.053)
<i>Time Controls</i>		
Mar 2023	-0.046** (0.017)	-0.046** (0.017)
May 2023	-0.069** (0.027)	-0.067** (0.027)
Jul 2023	0.269*** (0.036)	0.272*** (0.036)
Sep 2023	-0.033 (0.029)	-0.030 (0.029)
Nov 2023	-0.017 (0.015)	-0.018 (0.015)
Jan 2024	-0.054*** (0.015)	-0.053*** (0.015)
Mar 2024	-0.113*** (0.020)	-0.108*** (0.021)
May 2024	-0.103*** (0.029)	-0.097*** (0.029)
Jul 2024	0.352*** (0.037)	0.355*** (0.037)
Sep 2024	-0.022 (0.026)	-0.023 (0.026)
Observations	19,428	19,390
R^2_{within}	0.076	—
R^2_{overall}	—	0.067

Standard errors in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The RE specification includes additional geographic and demographic controls. Households in high-altitude areas consume significantly more electricity than those in lowland municipalities, consistent with higher climatic needs. The effect is statistically significant at the 1% level. The degree of urbanization is also positively correlated with demand: medium- and low-urbanized municipalities consume less electricity compared to large urban areas, with a strongly significant effect. Coastal municipalities display lower consumption levels, possibly reflecting smaller or more energy-efficient dwellings and milder climatic conditions.

As expected, higher contractual power is positively associated with electricity consumption, with a one-unit increase leading to approximately a 0.16 increase in total consumption. Surprisingly, households subscribed to green contracts consume roughly 7% more electricity than comparable non-green households. Although modest in magnitude, this finding can be interpreted in two ways. On the one hand, it may indicate a rebound effect, whereby adopting a green tariff reduces the perceived moral cost of electricity consumption. On the other hand, households opting for a green contract may deliberately shift their energy use toward electricity, for example by replacing gas-based appliances with electric ones, such as heat pumps or induction cookers. In this case, higher electricity demand would reflect a conscious substitution away from other energy sources perceived as less environmentally friendly. Seasonal dummies confirm pronounced temporal patterns. Electricity demand peaks in summer, particularly in July, when cooling needs are highest, while it declines markedly during winter months, reflecting the limited role of electricity in residential heating, given that most Italian households rely on natural gas systems. The Hausman test favors the RE specification, although coefficient magnitudes remain broadly consistent across models. Overall, residential electricity demand under ToU contracts responds only weakly to price signals but is more strongly shaped by climatic, geographic, and demographic factors.

3.6.3 Additional Regressions

Table C.7 summarizes the RE estimations, this time including the logarithm of lagged income²³ as an explanatory variable to capture income elasticity. Across all specifications, residential electricity demand remains consistently inelastic with respect to price, while income elasticities are consistently positive and significant around 0.7. This aligns with theoretical expectations and indicates that households in wealthier municipalities systematically consume more electricity.

Column (1) shows that the price elasticity of total electricity consumption is around

²³We recall that this variable serves as a proxy for income, since contemporaneous income data are not yet available.

-0.07 and statistically significant at 10% level. The coefficient on municipal average income is 0.75 and is highly statistically significant. Columns (2)–(5) report results separately for the three time periods. As in the previous case, price elasticity during peak hours is not significant, while income elasticity presents a coefficient of 0.79 and is highly significant across all specifications. Demand is more price-sensitive during off-peak hours (F2 and F3), with elasticities around -0.50 that are highly significant. Cross-price effects are positive and significant when off-peak consumption is the dependent variable, with coefficients around 0.38 in absolute value.

To explore heterogeneity in price responsiveness, households are divided into quartiles according to their baseline consumption. Quartiles are constructed ex-ante, prior to model estimation, using two alternative criteria as robustness checks. In the first specification (within-period quartiles) as represented in Figure 3.3a, consumption quartiles are defined within each bimonthly billing period, allowing households' relative positions to vary over time. In the second (seasonally adjusted quartiles) as shown in Figure 3.3b, quartiles are based on the residuals from a regression of total consumption on month fixed effects, removing seasonality and capturing deviations from expected monthly usage.

For each quartile, we estimate a FE model controlling for weather conditions and time dummies, with standard errors clustered at the household level. Results are highly consistent across both approaches, as shown in Table C.2. Figure 3.3 presents the estimated elasticities graphically. In both specifications, estimated price elasticities remain negative and statistically significant, confirming the robustness of the findings. Low-consumption households (represented in the first quartile) display the highest price sensitivity, around -0.24 , whereas high-consumption households exhibit much lower elasticities (between -0.01 and -0.03). The similarity of results across the two definitions of quartiles suggests that the observed heterogeneity is not driven by seasonality or by the grouping criterion, but rather reflects behavioral differences: those demanding less electricity are more responsive to price changes, while high-consumption households show more rigid demand patterns.

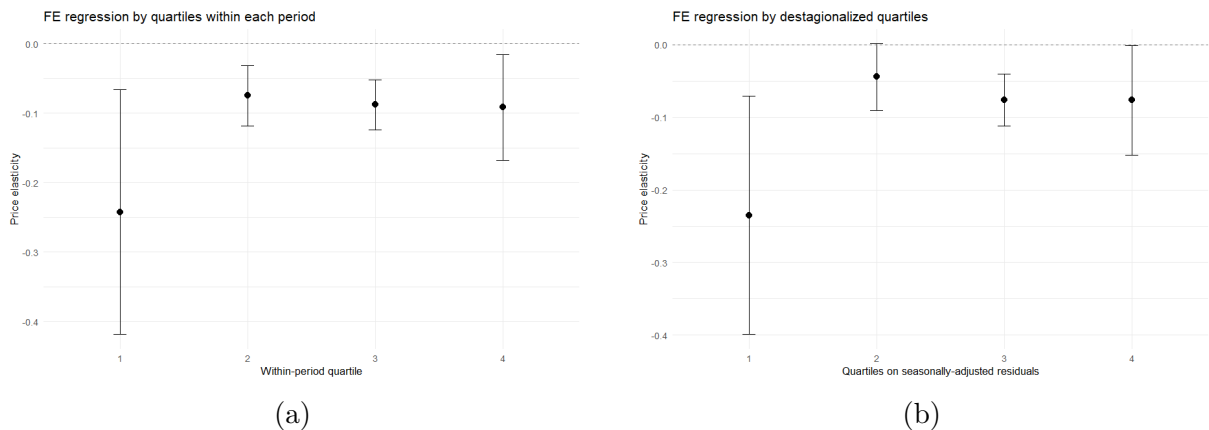


Figure 3.3: Differences in elasticity across levels of consumption.

3.6.4 Peak and Off-Peak Hours

This section presents the main findings on the substitutability of electricity consumption across time bands. Our objective is to assess how Italian households respond to ToU price variations and whether they move consumption between peak and off-peak periods. The ToU tariff is structured into three bands (F1, F2, and F3), where the variable component of the retail price depends on the wholesale PUN price, which differs across bands. We estimate own- and cross-price elasticities of demand using a log-log specification.

First, we observe that prices across the three bands are highly correlated: the correlation between lagged F1 and F2 prices is 0.98, between F1 and F3 is 0.986, and between F2 and F3 is 0.99. This strong co-movement of prices implies that including all three prices simultaneously in the same regression would lead to multicollinearity and biased estimates. To address this issue, we estimate separate models in which consumption in one band is regressed on its own lagged price and the lagged price of one additional band, holding the third band constant.

Table 3.3 reports results for peak-band consumption²⁴. We estimate both FE and RE specifications, using F2 as the cross-price elasticity in columns (1)–(2) and F3 in columns (3)–(4). The FE model exploits within-household variation only, whereas the RE model combines within- and between-household variation under the assumption that unobserved heterogeneity is uncorrelated with the regressors.

Results indicate no significant elasticity estimates (-0.30 for own-price and 0.33 for cross prices) for demand during peak hours, confirming that it remains sticky independently of own and cross prices. but statistically significant own-price response.

The RE specification includes geographic and demographic controls, which are consistent across models. Households located in mountainous areas consume substantially more electricity during peak hours (about 56% more than those in flat areas) likely reflecting greater climatic needs. In contrast, households in less urbanized areas and coastal municipalities consume significantly less electricity than those in large urban centers and inland areas, respectively. Households in Southern regions also display higher peak-hour consumption relative to those in the North²⁵.

Seasonal effects are evident: electricity demand peaks in summer, particularly in July, reflecting cooling use, and decreases in spring and autumn relative to the winter baseline. These findings highlight the combined influence of climatic and structural factors in shaping Italian households' electricity consumption. The Hausman test favors the RE specification, however coefficient magnitudes are broadly consistent across both models.

²⁴Recall that F1 band corresponds to weekdays from 8 a.m. to 7 p.m.

²⁵Replacing the geographic variable with a South dummy (1 = South, 0 = North-Center) confirms a strong positive association between living in the South and higher electricity consumption.

Table 3.3: FE/RE Estimates – Consumption in Peak Band (F1)

Variable	Consumption in F1			
	FE (F1–F2)	RE (F1–F2)	FE (F1–F3)	RE (F1–F3)
<i>Intercept</i>	—	4.170*** (0.250)	—	4.199*** (0.241)
<i>Price Elasticities</i>				
$\ln P_{F1,t-1}$	-0.296 (0.186)	-0.280 (0.188)	-0.286 (0.180)	-0.270 (0.181)
$\ln P_{F2,t-1}$	0.334 (0.228)	0.322 (0.230)	—	—
$\ln P_{F3,t-1}$	—	—	0.332 (0.227)	0.320 (0.229)
<i>Weather Controls</i>				
$\ln(\text{Temp})$	0.399* (0.221)	0.472** (0.222)	0.399* (0.221)	0.472** (0.222)
$\ln(\text{TempMin})$	-0.096 (0.064)	-0.110* (0.065)	-0.096 (0.064)	-0.110* (0.065)
$\ln(\text{TempMax})$	-0.584*** (0.195)	-0.651*** (0.195)	-0.584*** (0.195)	-0.651*** (0.195)
<i>Geographic Controls (RE only)</i>				
Altitude Zone (level 2)	—	0.249** (0.115)	—	0.249** (0.115)
Altitude Zone (level 3)	—	0.414*** (0.070)	—	0.414*** (0.070)
Altitude Zone (level 4)	—	0.442*** (0.084)	—	0.442*** (0.084)
Altitude Zone (level 5)	—	0.569*** (0.083)	—	0.569*** (0.083)
Urbanization (level 2)	—	-0.016 (0.057)	—	-0.016 (0.057)
Urbanization (level 3)	—	-0.392*** (0.063)	—	-0.392*** (0.063)
Coastal (dummy)	—	-0.320*** (0.071)	—	-0.320*** (0.071)
Island (dummy)	—	0.042 (0.110)	—	0.042 (0.110)
Area Geographical (1)	—	0.156** (0.068)	—	0.156** (0.068)
Area Geographical (2)	—	0.283*** (0.057)	—	0.283*** (0.057)
Power (kW)	—	0.129*** (0.017)	—	0.129*** (0.017)
Green contract	—	0.070** (0.023)	—	0.070** (0.023)
<i>Time Controls</i>				
Mar 2023	-0.090*** (0.020)	-0.090*** (0.020)	-0.104*** (0.015)	-0.103*** (0.015)
May 2023	-0.056*** (0.019)	-0.055*** (0.019)	-0.062*** (0.018)	-0.061*** (0.018)
Jul 2023	0.328*** (0.046)	0.327*** (0.047)	0.317*** (0.040)	0.317*** (0.040)
Sep 2023	-0.002 (0.025)	-0.0002 (0.026)	-0.010 (0.022)	-0.008 (0.022)
Nov 2023	-0.067*** (0.016)	-0.068*** (0.016)	-0.068*** (0.016)	-0.069*** (0.016)
Jan 2024	0.069* (0.036)	0.067* (0.036)	0.051** (0.026)	0.051* (0.026)
Mar 2024	-0.053 (0.043)	-0.053 (0.043)	-0.059 (0.039)	-0.059 (0.040)
May 2024	-0.051 (0.037)	-0.049 (0.037)	-0.054 (0.035)	-0.052 (0.036)
Jul 2024	0.420*** (0.033)	0.421*** (0.034)	0.428*** (0.038)	0.428*** (0.038)
Observations	19,428	19,390	19,428	19,390
R^2_{within}	0.075	—	0.075	—
R^2	—	0.062	—	0.062

Standard errors in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Hausman test: p -value = 0.999.

Similarly, Table 3.4 reports estimates for the F2 and F3 bands²⁶. Columns (1)–(2) show results for F2 consumption. Households display significant own-price elasticity under both FE and RE models: a 1% increase in the F2 price reduces consumption by about 0.51%. Columns (3)–(4) report results for F3 consumption, where own-price elasticity is of similar magnitude (–0.59) and statistically significant. Cross-price elasticities with respect to F1 prices are small and positive but only weakly significant, suggesting limited and asymmetric substitution effects. Owing to the very high correlation between F2 and F3 prices, estimates of cross elasticities between these two bands are not replaced.

As in the peak-band model, geographic and seasonal controls confirm that demand is higher in mountainous areas and lower in less urbanized or coastal municipalities. Seasonal

²⁶Recall that F2 corresponds to 7 a.m.–8 a.m and 7 p.m.–11 p.m. on weekdays and 7 a.m.–11 p.m. on Saturdays; F3 corresponds to 11 p.m.–7 a.m. every day and all hours on Sundays

dummies consistently show higher summer demand and reduced spring and autumn consumption. Also contractual power (kW) and green contract adoption have a positive effect on demand.

Table 3.4: FE/RE Estimates for ln Consumption in Off-Peak Bands (F2, F3)

Variable	Consumption in F2		Consumption in F3	
	FE (F2-F1)	RE (F2-F1)	FE (F3-F1)	RE (F3-F1)
<i>Intercept</i>	—	4.915*** (0.246)	—	5.434*** (0.222)
<i>Price Elasticities</i>				
$\ln P_{F2,t-1}$	-0.511** (0.224)	-0.519** (0.226)	—	—
$\ln P_{F3,t-1}$	—	—	-0.591** (0.208)	-0.590** (0.210)
$\ln P_{F1,t-1}$	0.358* (0.183)	0.371** (0.185)	0.322* (0.165)	0.329** (0.167)
<i>Weather Controls</i>				
$\ln(\text{Temp})$	0.590*** (0.218)	0.670*** (0.219)	0.248 (0.203)	0.334 (0.204)
$\ln(\text{TempMin})$	-0.120* (0.063)	-0.137** (0.063)	-0.073 (0.058)	-0.090 (0.059)
$\ln(\text{TempMax})$	-0.710*** (0.191)	-0.783*** (0.192)	-0.408** (0.179)	-0.490** (0.179)
<i>Geographic Controls (RE only)</i>				
Altitude Zone (level 2)	—	0.178 (0.115)	—	0.222** (0.105)
Altitude Zone (level 3)	—	0.415*** (0.070)	—	0.387*** (0.064)
Altitude Zone (level 4)	—	0.413*** (0.083)	—	0.438*** (0.076)
Altitude Zone (level 5)	—	0.537*** (0.082)	—	0.560*** (0.076)
Urbanization (level 2)	—	-0.011 (0.057)	—	-0.024 (0.052)
Urbanization (level 3)	—	-0.352*** (0.063)	—	-0.319*** (0.058)
Coastal (dummy)	—	-0.290*** (0.070)	—	-0.295*** (0.065)
Island (dummy)	—	0.039 (0.109)	—	0.025 (0.100)
Area Geographical (1)	—	0.170** (0.067)	—	0.139** (0.062)
Area Geographical (2)	—	0.208*** (0.057)	—	0.213*** (0.052)
Power (kW)	—	0.168*** (0.017)	—	0.195*** (0.016)
Green contract	—	0.079*** (0.023)	—	0.083*** (0.021)
<i>Time Controls</i>				
Mar 2023	-0.045** (0.020)	-0.044** (0.020)	-0.030** (0.013)	-0.028** (0.014)
May 2023	-0.078*** (0.019)	-0.076*** (0.019)	-0.022 (0.016)	-0.019 (0.016)
Jul 2023	0.173*** (0.045)	0.175*** (0.046)	0.272*** (0.037)	0.276*** (0.037)
Sep 2023	-0.016 (0.025)	-0.013 (0.025)	-0.056** (0.020)	-0.052** (0.020)
Nov 2023	-0.046** (0.016)	-0.046** (0.016)	0.037** (0.014)	0.037** (0.014)
Jan 2024	-0.137*** (0.035)	-0.138*** (0.036)	-0.157*** (0.024)	-0.156*** (0.024)
Mar 2024	-0.129** (0.042)	-0.128** (0.043)	-0.181*** (0.036)	-0.178*** (0.037)
May 2024	-0.111** (0.036)	-0.108** (0.037)	-0.119*** (0.032)	-0.114*** (0.033)
Jul 2024	0.314*** (0.033)	0.316*** (0.033)	0.347*** (0.035)	0.351*** (0.035)
Observations	19,428	19,390	19,428	19,390
R^2_{within}	0.049	—	0.089	—
R^2	—	0.039	—	0.075
Hausman p -value		0.855		0.053

Standard errors in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

To further mitigate multicollinearity, we test an alternative specification that aggregates F2 and F3 into a single off-peak category, consistent with practices in other countries (Filippini, 2011; Burns and Mountain, 2021). In this specification, F1 represents peak time, while F2 and F3 are combined into off-peak hours. Off-peak consumption is measured as the sum of F2 and F3, and the corresponding price is computed as a simple average.²⁷ Results, reported in Table C.1 in Appendix, confirm the main findings: peak demand is not significantly responsive to prices (coefficients around -0.3), while off-peak demand

²⁷Alternative constructions, such as weighted averages, will be considered in further analysis.

displays stronger negative own-price elasticity and a modest positive cross-price effect with respect to peak prices.

3.6.5 Heterogeneity Analysis

We now turn to heterogeneity analysis, exploring whether price responsiveness varies across different characteristics. To this end, we extend equations 3.1, 3.2, and 3.3 by including interaction terms between electricity prices and available geographic and demographic characteristics.

Table C.3 presents RE estimates of total electricity consumption with interaction terms. Each column shows a different specification with one interaction included. Across different specifications, the own-price elasticity of total demand remains negative, with magnitudes ranging from -0.02 to -0.19 , and is statistically significant in most cases. Notably, households in higher-altitude zones (levels 2, 4 and 5) exhibit stronger price responsiveness (elasticities around -0.30), while responsiveness is weaker at the highest altitude level. Less urbanized areas display significant but small sensitivity to price, while coastal households show a stronger reaction (around -0.27). In contrast, households holding green electricity contracts are associated with a positive and significant interaction term, although the size of the effect is negligible. Interaction effects by macro-areas (North/South) are not statistically significant.

Table C.4 examines heterogeneity in peak-hour consumption (F1). Own- and cross-price elasticities are significant only in specification (1) when we include the interaction with geographical area. We observe substantial heterogeneity. Again, households in altitude zones 2 and 4 exhibit significantly stronger reductions in F1 demand in response to own-price increases (elasticities around -0.56) and an increase in consumption when F2 price increases. Households in less urbanized and coastal areas are more price responsive, consistent with the findings from the total consumption model. Green contract holders again show contrasting patterns.

Last, Tables C.5 and C.6 report heterogeneity estimates for mid-peak (F2) and off-peak (F3) consumption. For both time bands, own-price elasticities are negative and significant, while cross-price elasticities with respect to F1 prices are positive and significant, indicating substitution from peak to off-peak consumption when peak prices increase. Heterogeneity is again evident: Panels A and B show that green contract holders display a positive cross-price elasticity (around 0.18) and a negative own-price elasticity (around -0.15). This suggests that the positive coefficient observed for green households in the total consumption model (Table C.3) is mainly driven by stickiness in peak-hour consumption.

Overall, these findings highlight heterogeneity in household electricity demand elasticities across geographic and contract characteristics. Households in less urbanized and coastal areas appear more responsive to price signals. Moreover, price responsiveness is

consistently higher in off-peak (F2/F3) hours than in peak (F1), suggesting that households have greater flexibility to adjust consumption during evening and night compared to daytime hours.

Table 3.5: Summary of significant own- and cross-price elasticities by time band

	F1 consumption	F2 consumption	F3 consumption
Own-price elasticity	-0.29	-0.51**	-0.59**
Cross-price with F1	—	+0.36**	+0.32*

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Elasticities based on FE/RE models by band.

3.7 Discussion and Conclusion

This study provides new evidence on residential electricity consumption behavior in Italy under a ToU tariff structure. It addresses two main questions: (i) how elastic is total electricity demand with respect to national price variations? (ii) To what extent do behavioral patterns emerge across time bands, indicating potential shifts in consumption from peak to off-peak periods? Especially for flexible activities such as laundry or dishwashing.

Two central findings emerge. First, total residential electricity demand is resilient to price changes. Second, substitution across time bands exists, although it is limited and asymmetric. As summarized in Table 3.5, peak consumption (F1) reacts weakly to its own price (elasticity around ~ 0.30) and the estimates are not significant, whereas off-peak demand (F2 and F3) is more responsive, with own-price elasticities around ~ 0.51 and ~ 0.59 . Cross-price effects show that increases in peak-hour prices slightly raise off-peak consumption, while F1 demand does not respond significantly to F2 or F3 prices.

Using the lagged PUN as the price variable, we estimate an elasticity of total demand of about ~ 0.16 , consistent with previous studies that find short-run elasticities around ~ 0.2 (Labandeira et al., 2017). Heterogeneity across consumption levels reveals that low-consuming households are more sensitive to price changes, whereas high-consuming households exhibit limited behavioral response.

This study leverages detailed geographical and contractual information. Electricity demand increases systematically with altitude, suggesting higher consumption in mountainous areas, while households in coastal and less urbanized municipalities consume less. Regional differences are also evident, with average electricity consumption being higher in Southern Italy. A possible explanation is that southern regions are characterized by higher unemployment rates, meaning that more people spend time at home. Since time spent at home is often associated with higher residential energy use, unemployment may contribute to greater electricity consumption in these areas (Cicala, 2023). An interesting finding is that households subscribed to green tariffs consume more electricity than others. Although this may appear counterintuitive, since such consumers are often assumed to be

more environmentally conscious, it could reflect socioeconomic differences among green adopters that are not captured in our data. Two alternative explanations can be considered. First, the pattern could be associated with a rebound effect: households holding green contracts may perceive their environmental footprint as lower and consequently pay less attention to their electricity consumption. This interpretation is concerning and warrants further investigation. Second, the result might be related to a voluntary preference for adopting a greater number of electricity-based appliances, which are perceived as more environmentally friendly than alternative energy sources. Additional data are required to discriminate among the two interpretations.

Seasonal patterns are also pronounced. Relative to the January–February 2023 baseline, electricity consumption declines in spring and peaks in July and August, driven by cooling needs. The effect is particularly strong during the exceptionally hot and humid summer of 2024. The contrast between summer and winter demand is less marked, as low temperatures also increase electricity use for heating purposes. This pattern is not surprising, given that many heating systems in Italy rely on natural gas, whereas cooling systems are predominantly electricity-based.

In the second part of the analysis, we examine price responsiveness across time bands. This is particularly relevant for policy design, as ToU tariffs can reduce inefficiencies and costs related to peak load. The Italian ToU scheme under study is structured either into three time bands (F1, F2, F3) or two aggregated bands (F1 vs. off-peak = F2+F3), consistent with international practice (Filippini, 2011; Burns and Mountain, 2021). Descriptive statistics show that average monthly consumption is highest in F3, followed by F2 and F1. While this may seem surprising, it aligns with the current generation mix: renewable capacity has expanded rapidly, and many households self-produce electricity during daytime hours via rooftop solar panels. Unfortunately, our dataset does not include direct information on self-generation capacity.

Estimation results confirm limited but asymmetric substitution between time bands. Peak-hour (F1) consumption is relatively unresponsive to both own and cross-price changes, whereas off-peak consumption (F2 and F3) shows stronger own-price responsiveness and partial substitution when F1 prices rise. These results are consistent with previous evidence (Filippini, 1995a, 2011) and remain robust when F2 and F3 are aggregated into a single off-peak category. Additional analyses highlight substantial heterogeneity in price responsiveness. Households in high-altitude and coastal areas are systematically more elastic across all bands, possibly reflecting differences in heating and cooling needs or housing characteristics.

This study also presents some limitations. Unlike many survey-based analyses, our dataset does not include demographic characteristics (e.g., house’s information, number of households, etc.), which could provide additional behavioral insights and help control for household habits. Moreover, the structure of the electricity market constrains the

identification strategy: prices are correlated across time bands, and we rely on the lagged PUN as the reference price. Finally, the voluntary choice between flat and ToU tariffs implies that self-selection cannot be entirely ruled out. Future experimental work could further reinforce the causal interpretation of our findings. Further research could address these limitations by collecting additional data to extend the time span of the analysis and consider additional tariffs. In particular, combining observational data with demographic information will really enrich the behavioral dimension of the study.

Taken together, the findings have clear policy implications. It is worth noting that under the ToU tariff currently in place, the pricing system lacks full transparency. Prices are communicated ex post, and households must check the GME website to obtain indicative real-time price information, which may differ from the final billed rate. Nonetheless, ToU pricing offers a valuable framework for assessing whether price-based mechanisms can effectively influence electricity use and how they can be leveraged by regulators. The overall inelasticity of residential demand suggests that a uniform increase in prices, for instance, through the introduction of a national carbon tax, would not substantially reduce overall consumption. By contrast, dynamic pricing mechanisms appear more promising. For example, a time-varying energy tax (with higher prices during peak or high-emission hours and lower prices during off-peak periods) could shift consumption toward cleaner and less costly times, reducing pressure on the grid. Furthermore, the heterogeneity analysis and geographic differences highlight that uniform price policies may yield uneven behavioral responses. Tailored policies adjusted to local climatic, economic, and infrastructural conditions, are likely to be more effective. Combining targeted information campaigns to improve consumer awareness with well-designed dynamic pricing schemes could meaningfully contribute to managing peak loads and promoting more efficient electricity use.

Conclusions

The thesis contributes to a deeper understanding of how to effectively address imperfect and asymmetric information and how to foster responsible, efficient, and sustainable behavior among different economic agents.

Chapter 1 provides empirical evidence of systematic deviations from the theoretical prediction of full unraveling. Treatment variations significantly influence the strategic thinking and beliefs of both informed and uninformed individuals, confirming the presence of bounded rationality. Although full unraveling represents an optimal policy benchmark in theory, it is rarely observed in real-world settings. The informed party often benefits from withholding information, as a large share of receivers tend to anchor their beliefs around the mean value of the probability distribution.

Chapter 2 investigates strategic misreporting of environmental information in the Italian electricity market. Promoting a customized Greenwashing Index, the study documents heterogeneity in firms' greenwashing behavior. Firm size, profitability, and sectoral characteristics emerge as key determinants of greenwashing. The chapter contributes both methodologically and conceptually to the existing literature, offering a novel framework to quantify and analyze misleading environmental communication.


Chapter 3 estimates short-run elasticities of household electricity demand under ToU tariffs, filling an important gap in the Italian literature. Although price changes do not affect total consumption, price differences across time bands partially influence demand shift to off-peak hours. These results suggest that ToU pricing can be considered as an effective tool to increase efficiency and reduce peak load by encouraging consumption away from peak periods.

The main message of the thesis is that information must be accessible, verifiable and clear to all market participants. If these conditions are not met, several factors like cognitive biases, strategic misrepresentation and simple unawareness, can distort the optimal economic behavior. In conclusion, it calls for active monitoring of business environmental communications and the design of tailored instruments for households, enabling individuals to recognize that their actions generate both environmental and economic externalities that shape the sustainability of the community they belong to.


Appendix A

Appendix of Chapter 1


A.1 Between-Subject Variations




(a) Control: TA




(b) Control: TB




(c) Control: TC




(d) Treatment 1: TA




(e) Treatment 1: TB




(f) Treatment 1: TC



(g) Treatment 2: TA



(h) Treatment 2: TB



(i) Treatment 2: TC

Figure A.1: (a),(b),(c): Probability distributions in control group, (d),(e),(f): Probability distributions in Treatment 1, (g),(h),(i): Probability distributions in Treatment 2

A.2 Payoff Tables (Within-Subject Variations)

		RENTER							
OWNER	1	1.5	2	2.5	3	3.5	4	4.5	5
1	6,110	23,104	40,95	56,84	70,70	84,56	95,40	104,23	110,6
2	6,95	23,104	40,110	56,104	70,95	84,84	95,70	104,56	110,40
3	6,70	23,84	40,95	56,104	70,110	84,104	95,95	104,84	110,70
4	6,40	23,56	40,70	56,84	70,95	84,104	95,110	104,104	110,95
5	6,6	23,23	40,40	56,56	70,70	84,84	95,95	104,104	110,110

Figure A.2: Payoff Table in TA

		RENTER				
OWNER	1	1.5	2	2.5	3	
1	31,110	57,99	80,80	99,57	110,31	
2	31,80	57,99	80,110	99,99	110,80	
3	31,31	57,57	80,80	99,99	110,110	

Figure A.3: Payoff Table in TB

		RENTER		
OWNER	1	1.5	2	
1	50,110	87,87	110,50	
2	50,50	87,87	110,110	

Figure A.4: Payoff Table in TC

A.3 Senders' Disclosure Rates

Table A.1: Sender Experiment: average disclosure rates

Value	B=5		B=3		B=2	
	N	Mean	N	Mean	N	Mean
Disclose (=5)	445	0.857				
Disclose (=4)	429	0.827				
Disclose (=3)	301	0.580	441	0.850		
Disclose (=2)	137	0.264	306	0.590	423	0.815
Disclose (=1)	122	0.235	121	0.233	142	0.274
Distribution	%		%		%	
Uniform	55.06%		59.00%		54.02%	
Bell-shaped	55.09%		52.63%		57.02%	
Fat-tailed	55.63%		55.56%		52.30%	
No group distinction	55.26%		55.75%		54.43%	

Table A.2: Sender Experiment: average disclosure rates (between-subject)

	Uniform		Bell-shaped		Fat-tailed	
	N	Mean	N	Mean	N	Mean
B=5						
Disclose (b = 5)	157	0.902	145	0.848	143	0.822
Disclose (b = 4)	153	0.879	136	0.795	140	0.805
Disclose (b = 3)	102	0.586	91	0.532	108	0.621
Disclose (b = 2)	38	0.218	50	0.292	49	0.282
Disclose (b = 1)	29	0.167	49	0.287	44	0.253
B=3	N	Mean	N	Mean	N	Mean
Disclose (b = 3)	156	0.897	141	0.825	144	0.828
Disclose (b = 2)	119	0.684	86	0.503	101	0.580
Disclose (b = 1)	33	0.190	43	0.251	45	0.259
B=2	N	Mean	N	Mean	N	Mean
Disclose (b = 2)	146	0.839	139	0.813	138	0.793
Disclose (b = 1)	42	0.241	56	0.327	44	0.253

A.4 χ^2 Test Results

Table A.3: Sender Experiment: χ^2 test results on disclosure rates across different treatments

Comparison	B=2		B=3		B=5	
	χ^2	p-value	χ^2	p-value	χ^2	p-value
b=5 C vs T1	-	-	-	-	1.8631	0.1723
C vs T2	-	-	-	-	4.0842	0.04329
T1 vs T2	-	-	-	-	0.25811	0.6114
b=4 C vs T1	-	-	-	-	3.8779	0.04893
C vs T2	-	-	-	-	3.1096	0.07783
T1 vs T2	-	-	-	-	0.0065222	0.9356
b=3 C vs T1	-	-	3.1549	0.0757	0.81448	0.3668
C vs T2	-	-	2.9242	0.08726	0.30021	0.5838
T1 vs T2	-	-	2.8346e-30	1	2.4181	0.1199
b=2 C vs T1	0.25025	0.6169	10.977	0.0009225	2.1117	0.1462
C vs T2	0.93816	0.3328	3.5714	0.05878	1.5326	0.2157
T1 vs T2	0.10627	0.7444	1.788	0.1812	0.010499	0.9184
b=1 C vs T1	2.735	0.09817	1.5751	0.2095	6.4153	0.01131
C vs T2	0.015445	0.9011	1.9994	0.1574	3.3977	0.06529
T1 vs T2	1.9841	0.159	0.00084095	0.9769	0.34043	0.5596

A.5 Receivers' Average Guesses across Treatments

Table A.4: Receiver Experiment: average guesses per group

Within-groups	Average Guess	Range	Between-groups		
			Uniform	Bell-shaped	Fat-tailed
B=5	2.83	1-5	3.02	2.98	2.50
B=3	1.92	1-3	2.01	2.01	1.75
B=2	1.44	1-2	1.43	1.44	1.45

A.6 Non-parametric Tests on Receivers' Guesses

Table A.5: Receiver Experiment: Kruskal-Wallis Test and Post hoc Mann-Whitney U Results for Guesses: Receiver

Case	Kruskal-Wallis		Post hoc Mann-Whitney U and Cliff's Delta			
	Chi-squared	P-value	Comparison	P-value	Cliff's Delta	Confidence Interval (95%)
B=5	33.564	5.147e-08	C vs T1	1	0.0069	[-0.1145, 0.1281]
			C vs T2	2.2e-06	0.3072	[0.1872, 0.4181]
			T1 vs T2	1.6e-06	0.3123	[0.1896, 0.4254]
B=3	28.098	7.918e-07	C vs T1	1	-0.0041	[-0.1200, 0.1119]
			C vs T2	5.4e-05	0.2581	[0.1420, 0.3673]
			T1 vs T2	4.3e-06	0.2862	[0.1709, 0.3938]
B=2	0.24612	0.8842	C vs T1	1	-0.0146	[-0.1259, 0.0971]
			C vs T2	1	-0.0281	[-0.1379, 0.0823]
			T1 vs T2	1	-0.0130	[-0.1257, 0.1000]

A.7 Senders' Elicited Beliefs

Table A.6: Sender Experiment: Average Guesses per Group

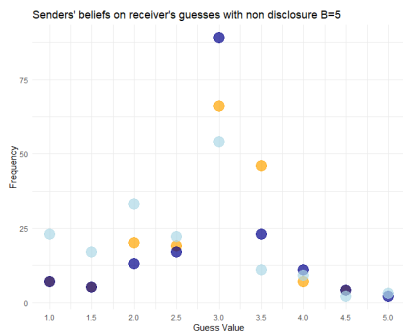
Within-groups	Average Guess	Range	Between-groups		
			Uniform	Bell-shaped	Fat-tailed
B=5	2.77	1-5	2.91	2.93	2.47
B=3	1.88	1-3	1.95	2.0	1.69
B=2	1.46	1-2	1.46	1.50	1.40

A.8 Non-parametric Tests on Senders' Beliefs

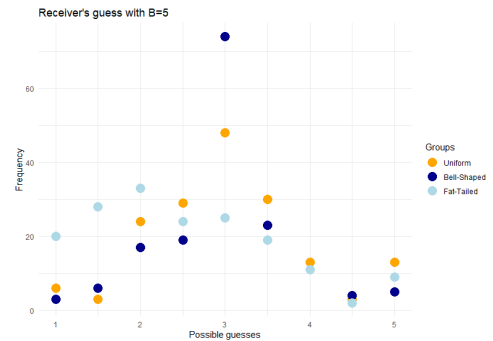
Table A.7: Sender Experiment: Kruskal-Wallis and Post-Hoc Mann-Whitney U Test on sender's belief

Within-group	Kruskal-Wallis χ^2	p-value	Post-Hoc	p-value (Bonferroni)
B=5	38.14	5.225e-09	C vs T2	4.1e-07
			T1 vs T2	2.5e-07
B=3	40.617	1.514e-09	C vs T2	5.7e-06
			T1 vs T2	9.6e-09
B=2	10.963	0.004163	C vs T2	0.1201
			T1 vs T2	0.0045

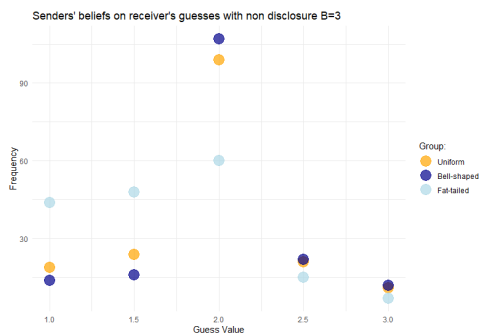
A.9 A Comparison between Senders' Beliefs and Receivers' Guesses




(a) Sender's belief $B = 5$ (orig.)



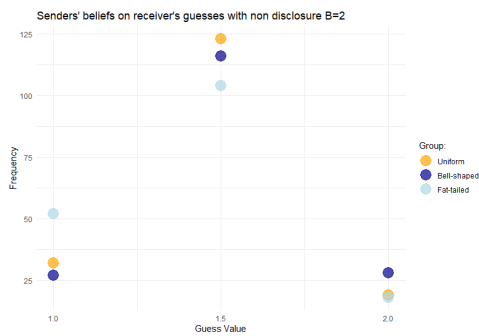
(b) Receiver's guess $B = 5$ (new)




(c) Sender's belief $B = 3$ (orig.)



(d) Receiver's guess $B = 3$ (new)



(e) Sender's belief $B = 2$ (orig.)



(f) Receiver's guess $B = 2$ (new)

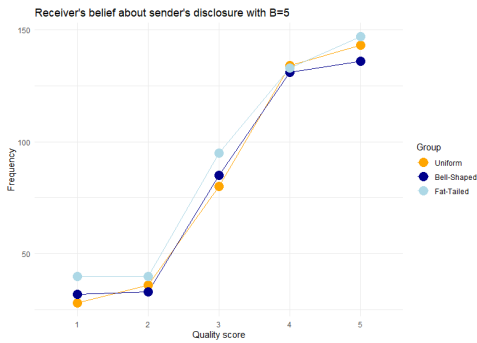
Figure A.5: (a), (c), (e) represents sender's elicited beliefs about receiver's guess; (b), (d), (f) represents the observed guess of receiver with no disclosure by sender.

A.10 Senders' Belief and Non-disclosure


Table A.8: Sender Experiment: Sender's belief and non-disclosure decision across different coarseness

Belief	Panel A					Panel B					Panel C				
	N	Avg	b	Nd	Mean	N	Avg	b	Nd	Mean	N	Avg	b	Nd	Mean
1.0	37	0.07	1	25	0.68	77	0.15	1	52	0.68	111	0.21	1	81	0.73
			2	24	0.65			2	25	0.32			2	16	0.14
			3	13	0.35			3	12	0.16					
			4	5	0.14										
			5	5	0.14										
1.5	27	0.05	1	20	0.74	88	0.17	1	75	0.85	343	0.66	1	256	0.75
			2	16	0.59			2	22	0.25			2	62	0.18
			3	7	0.26			3	10	0.11					
			4	4	0.15										
			5	2	0.07										
2.0	66	0.13	1	47	0.71	266	0.51	1	206	0.77	65	0.13	1	40	0.62
			2	46	0.70			2	115	0.43			2	18	0.28
			3	15	0.23			3	40	0.15					
			4	8	0.12										
			5	8	0.12										
2.5	58	0.11	1	51	0.88	58	0.11	1	45	0.78					
			2	49	0.84			2	35	0.60					
			3	22	0.38			3	5	0.09					
			4	9	0.16										
			5	9	0.16										
3.0	209	0.40	1	165	0.79	30	0.06	1	20	0.67					
			2	165	0.79			2	16	0.53					
			3	96	0.46			3	11	0.37					
			4	28	0.13										
			5	25	0.12										
3.5	80	0.15	1	62	0.78										
			2	57	0.71										
			3	42	0.53										
			4	21	0.26										
			5	15	0.19										
4.0	27	0.05	1	15	0.56										
			2	14	0.52										
			3	15	0.56										
			4	8	0.30										
			5	9	0.33										
4.5	10	0.02	1	9	0.90										
			2	9	0.90										
			3	6	0.60										
			4	6	0.60										
			5	1	0.10										
5.0	5	0.01	1	3	0.60										
			2	2	0.40										
			3	2	0.40										
			4	1	0.20										
			5	0	0.00										

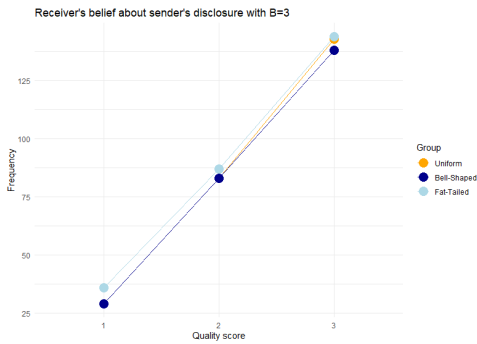
A.11 A Comparison between Receivers' First Order Beliefs and Senders' Disclosure



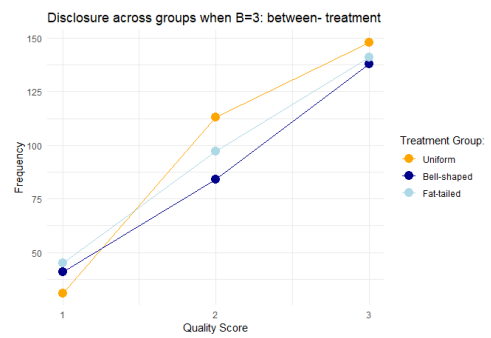
(a) Receiver's belief $B = 5$



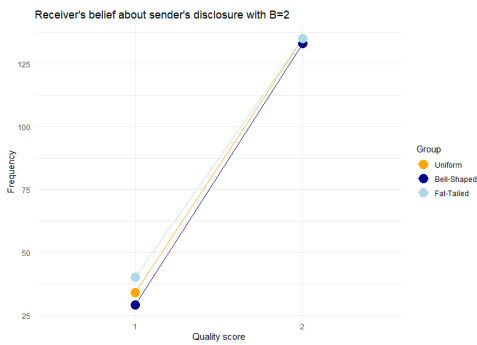
(b) Sender's disclosure $B = 5$



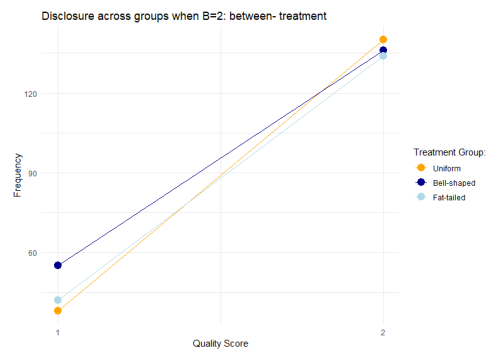
(c) Receiver's belief $B = 3$



(d) Sender's disclosure $B = 3$



(e) Receiver's belief $B = 2$



(f) Sender's disclosure $B = 2$

Figure A.6: (a), (c), (e) represents receiver's elicited beliefs about sender's reporting decision; (b), (d), (f) represents the observed decision of the sender.

A.12 Receivers' Average First Order Beliefs

Table A.9: Receiver Experiment: Average first order beliefs

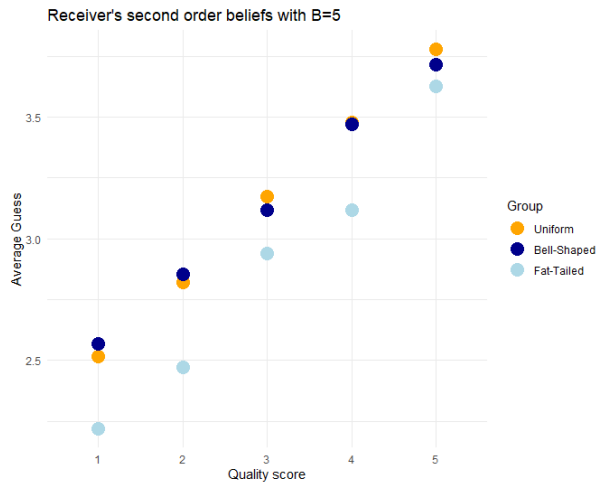
	Overall	Uniform	Bell-shaped	Fat-tailed
B=5				
Disclose (b = 5)	0.85	0.85	0.84	0.86
Disclose (b = 4)	0.79	0.79	0.81	0.78
Disclose (b = 3)	0.52	0.47	0.52	0.56
Disclose (b = 2)	0.22	0.21	0.20	0.23
Disclose (b = 1)	0.20	0.17	0.20	0.23
B=3				
Disclose (b = 3)	0.85	0.85	0.85	0.84
Disclose (b = 2)	0.50	0.49	0.51	0.51
Disclose (b = 1)	0.19	0.17	0.18	0.21
B=2				
Disclose (b = 2)	0.80	0.80	0.82	0.79
Disclose (b = 1)	0.21	0.20	0.18	0.23

A.13 Receivers' Average Second Order Beliefs

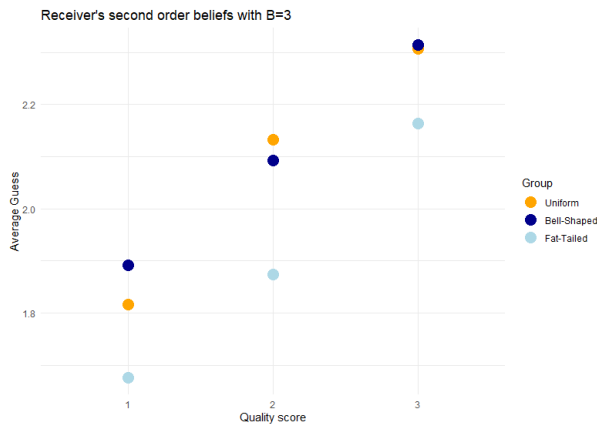
Table A.10: Receiver Experiment: Average second-order beliefs

	Overall	Uniform	Bell-shaped	Fat-tailed
B=5				
Second-order belief (b = 5)	3.71	3.78	3.72	3.63
Second-order belief (b = 4)	3.35	3.48	3.47	3.12
Second-order belief (b = 3)	3.08	3.17	3.12	2.94
Second-order belief (b = 2)	2.71	2.82	2.85	2.47
Second-order belief (b = 1)	2.43	2.52	2.57	2.22
B=3				
Second-order belief (b = 3)	2.26	2.31	2.31	2.16
Second-order belief (b = 2)	2.03	2.13	2.09	1.87
Second-order belief (b = 1)	1.79	1.82	1.89	1.68
B=2				
Second-order belief (b = 2)	1.61	1.64	1.62	1.58
Second-order belief (b = 1)	1.43	1.40	1.47	1.42


A.14 Receivers' Second Order Beliefs



(a) Belief $B=5$



(b) Belief $B=3$



(c) Belief $B=2$

Figure A.7: Receiver Experiment: Receiver's beliefs about sender's expectation of receiver's guess

A.15 Receiver Experiment: Receiver's Second Order Beliefs Non-parametric Tests

Table A.11: Kruskal-Wallis Test and Post hoc Mann-Whitney U Results for HOB: Receiver

Case	Kruskal-Wallis		Post hoc Mann-Whitney U and Cliff's Delta			
	Chi-squared	P-value	Comparison	P-value	Cliff's Delta	Confidence Interval (95%)
B=5						
b=5	0.5638	0.7544	–	–	–	–
b=4	10.5800	0.0050	C vs T2	0.012	0.1774	[0.0555, 0.2941]
			T1 vs T2	0.019	0.1700	[0.0472, 0.2877]
b=3	4.8933	0.0866	–	–	–	–
b=2	21.3150	0.0000	C vs T2	0.00061	0.2281	[0.1084, 0.3413]
			T1 vs T2	0.00008	0.2599	[0.1398, -0.3724]
b=1	17.5790	0.0002	C vs T2	0.00552	0.1916	[0.0708, 0.3070]
			T1 vs T2	0.00018	0.2485	[0.1275, 0.3623]
B=3						
b=3	5.6677	0.0588	C vs T2	0.120	0.1240	[0.0046, 0.2399]
			T1 vs T2	0.120	0.1257	[0.0051, 0.2426]
b=2	21.7080	0.0000	C vs T1	0.8752	0.0632	[-0.0548, 0.1794]
			C vs T2	4.5e-05	0.2598	[0.1446, 0.3681]
			T1 vs T2	0.0014	0.2106	[0.0938, 0.3217]
b=1	17.5290	0.0002	C vs T1	0.304	-0.0959	[-0.2085, 0.0193]
			C vs T2	0.042	0.1467	[0.0291, 0.2602]
			T1 vs T2	8.8e-05	0.2497	[0.1332, 0.3593]
B=2						
b=2	2.6207	0.2697	–	–	–	–
b=1	3.4555	0.1777	–	–	–	–

A.16 Regression Results

Table A.12: Sender Experiment: OLS with average disclosures as outcome (including Age)

	Avg Discl. (All)		Avg Discl. (B = 5)		Avg Discl. (B = 3)		Avg Discl. (B = 2)	
	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
Intercept	0.554***	0.042	0.569***	0.047	0.587***	0.053	0.465***	0.061
Bell-shaped (T1)	-0.012	0.022	0.001	0.025	-0.063**	0.028	0.030	0.032
Fat-tailed (T2)	-0.010	0.022	0.006	0.025	-0.034	0.028	-0.017	0.032
Male (Sex=1)	0.014	0.019	0.006	0.021	0.010	0.024	0.038	0.027
Other / No answer (Sex=2)	-0.103	0.070	-0.109	0.079	-0.107	0.089	-0.081	0.102
Age	0.000	0.001	0.000	0.001	0.000	0.001	0.001	0.001
CRT-2 score	-0.015	0.036	-0.037	0.040	-0.015	0.045	0.042	0.052
Observations	519		519		519		519	
R ²	0.007		0.006		0.014		0.012	
Residual Std. Error	0.207 (df = 512)		0.234 (df = 512)		0.262 (df = 512)		0.301 (df = 512)	
F Statistic	0.60 (df = 6; 512)		0.50 (df = 6; 512)		1.18 (df = 6; 512)		1.06 (df = 6; 512)	

Note: * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$. Reference category: $t_between = T0$; Sex=0 (Female).

Table A.13: Sender Experiment: Logit regressions with behavioral types as outcome

	All		B = 5		B = 3		B = 2	
	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
Panel A: Full Unraveling type (FU)								
Intercept	-2.042***	0.611	-1.638**	0.539	0.070	0.410	1.380**	0.514
Bell-shaped (T1)	0.027	0.346	0.208	0.299	-0.839***	0.221	-0.171	0.289
Fat-tailed (T2)	0.320	0.328	0.238	0.298	-0.342	0.220	-0.305	0.283
Male (Sex=1)	0.277	0.276	0.064	0.246	0.056	0.185	-0.005	0.240
Other / No answer (Sex=2)	-14.521	793.541	-14.950	793.215	-0.768	0.728	-1.278	0.701
Age	0.009	0.010	0.009	0.009	0.005	0.007	-0.010	0.009
CRT-2 score	-0.913	0.514	-0.872	0.458	0.367	0.351	1.233**	0.436
Observations	519		519		519		519	
AIC	389.5		464.9		713.9		496.8	
Residual deviance	375.5		450.9		699.9		482.8	
Panel B: Bias-to-the-Mean type (BTM)								
Intercept	-2.563***	0.596	-0.905*	0.451	-0.319	0.422	0.586	0.414
Bell-shaped (T1)	-0.785*	0.325	-0.471	0.241	-0.840***	0.228	-0.264	0.224
Fat-tailed (T2)	0.206	0.271	0.020	0.229	-0.401	0.219	-0.153	0.224
Male (Sex=1)	0.518*	0.244	0.103	0.197	0.047	0.189	-0.254	0.188
Other / No answer (Sex=2)	0.544	0.856	0.115	0.738	-0.190	0.734	-1.183	0.731
Age	-0.008	0.010	-0.013	0.008	-0.005	0.007	-0.016*	0.007
CRT-2 score	1.851***	0.520	1.229**	0.394	0.747*	0.364	1.115**	0.357
Observations	519		519		519		519	
AIC	462.4		644.1		689.3		697.9	
Residual deviance	448.4		630.1		675.3		683.9	

Notes: * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$. Coefficients are log-odds from logit regressions.

Table A.14: Receiver Experiment: OLS regressions with guess as outcome

	All (Average Guess)		B = 5		B = 3		B = 2	
	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
Panel A								
Intercept	2.235***	0.083	3.003***	0.175	2.211***	0.099	1.490***	0.062
Bell-shaped (T1)	-0.018	0.049	-0.053	0.103	-0.008	0.058	0.008	0.037
Fat-tailed (T2)	-0.254***	0.048	-0.523***	0.101	-0.263***	0.057	0.023	0.036
Age	-0.00007	0.001	-0.00078	0.003	0.00052	0.001	0.00005	0.001
Male (Sex=1)	-0.047	0.040	0.006	0.084	-0.112**	0.047	-0.034	0.030
CRT-2 score	-0.092	0.074	0.086	0.156	-0.283***	0.088	-0.080	0.056
Observations	499		499		499		499	
R ²	0.071		0.062		0.084		0.008	
Residual Std. Error	0.443 (df = 493)		0.929 (df = 493)		0.525 (df = 493)		0.331 (df = 493)	
F Statistic	7.567*** (df = 5; 493)		6.534*** (df = 5; 493)		9.041*** (df = 5; 493)		0.831 (df = 5; 493)	
Panel B								
Intercept	1.430***	0.143	0.985***	0.282	0.904***	0.157	0.755***	0.105
Bell-shaped (T1)	-0.027	0.047	-0.039	0.093	-0.015	0.053	-0.001	0.034
Fat-tailed (T2)	-0.242***	0.046	-0.356***	0.093	-0.168***	0.053	0.033	0.034
Age	0.00002	0.001	0.00044	0.002	0.00113	0.001	0.00018	0.001
Male (Sex=1)	-0.023	0.038	0.012	0.076	-0.104**	0.043	-0.012	0.028
CRT-2 score	-0.056	0.071	0.218	0.143	-0.232***	0.080	-0.049	0.052
Average FOB	0.021	0.071	0.0088	0.179	0.223**	0.087	0.073	0.052
Average SOB	0.497***	0.073	0.595***	0.059	0.543***	0.054	0.435***	0.054
Observations	499		499		499		499	
R ²	0.153		0.226		0.246		0.134	
Residual Std. Error	0.424 (df = 491)		0.845 (df = 491)		0.478 (df = 491)		0.310 (df = 491)	
F Statistic	12.63*** (df = 7; 491)		20.48*** (df = 7; 491)		22.86*** (df = 7; 491)		10.85*** (df = 7; 491)	

Note: Robust SE non richieste/fornite; ***, **, * indicano $p < 0.01$, $p < 0.05$, $p < 0.10$.

A.17 Mean Squared Error

We quantify the gap between observed reporting trend and the expected theoretical prediction to check whether the number of potential quality values affect the disclosing rates using the Mean Squared Error (MSE).

We analyse three disclosure scenarios ($B = 2,3,5$). In each scenario we observe, for each respondent i and each bin j , a binary indicator:

$$d_{i,j} = \begin{cases} 1 & \text{if respondent } i \text{ disclosed at quality level } j \\ 0 & \text{otherwise} \end{cases}$$

Under full-unraveling theory predicts that every respondent discloses at all but the lowest-quality value, to which is indifferent between disclosing or not. For simplicity we attach a probability of 50% to the lowest quality, but it can be any number between 0 and 100. The vector of expected disclosure probabilities is $p^*=(0.5,1,1,1,1)$ for $B = 5$, $p^*=(0.5,1,1)$ for $B = 3$ and $p^*=(0.5,1)$ for $B = 2$.

In each scenario we compute the sample proportion of disclosures for each quality j :

$$\hat{p}_j = \frac{1}{N} \sum_{i=1}^N d_{i,j} \quad (\text{A.1})$$

so that $\hat{p} = (\hat{p}_1, \dots, \hat{p}_B)$. Then, to quantify how close is the observed disclosure profile to the theoretical benchmark, we compute the MSE as follows:

$$\text{MSE} = \frac{1}{B} \sum_{j=1}^B (\hat{p}_j - p_j^*)^2. \quad (\text{A.2})$$

with lower values indicating a better fit.

Appendix B

Appendix of Chapter 2

B.1 Greenwashing Indices Distributions




Figure B.1: Different Disclosure Metrics




Figure B.2: Different Performance Metrics

Table B.1: Descriptive Statistics for Greenwashing Analysis (N = 354)

Variable	Min	1st Qu.	Median	Mean	3rd Qu.	Max
Text Mining						
N. green words (lemmatized) (1)	0.00	43.75	111.50	183.60	200.50	5050.00
N. green words (stemmed) (2)	0.00	72.50	162.00	270.50	305.50	6692.00
Total Lemmatized (3)	5.00	902.50	2006.00	3160.30	3659.20	61771.00
Total Stemmed (4)	5.00	874.20	1941.50	3041.10	3504.00	59944.00
Total Tokens (5)	10.00	1499.00	3330.00	5228.00	6079.00	102134.00
Frequency on total tokens ((1)/(5))	0.00	0.02	0.03	0.03	0.04	0.11
Frequency on total tokens ((2)/(5))	0.00	0.04	0.05	0.05	0.06	0.15
Fuel Mix Composition						
Renewable %	0.00	7.02	7.06	24.98	27.10	100.00
Carbon %	0.00	13.92	18.94	15.13	18.96	18.97
Natural Gas %	0.00	45.77	62.25	49.72	62.32	62.32
Oil Product %	0.00	1.218	1.660	1.324	1.660	1.670
Nuclear Power %	0.00	2.192	2.990	2.385	2.990	2.990
Other %	0.00	5.178	7.040	5.624	7.050	7.060
Indicator's Components						
Performance Peer	0.00565	0.24859	0.50141	0.50141	0.75071	0.95763
Performance Stand	-0.81419	-0.58534	-0.58403	0.00000	0.06934	2.44577
Performance Norm	0.0000	0.0702	0.0706	0.2498	0.2710	1.0000
Disclosure Peer	0.004237	0.252119	0.501412	0.501412	0.750706	1.000000
Disclosure Stand	-2.4590	-0.6611	-0.2005	0.0000	0.4349	4.5580
Disclosure Norm	0.0000	0.2562	0.3219	0.3504	0.4124	1.0000
Greenwashing Peer	-0.929378	-0.202684	0.002825	0.000000	0.205509	0.966102
Greenwashing Stand	-3.8812	-0.4838	0.0431	0.0000	0.5976	5.1434
Greenwashing Norm	-0.85414	0.04765	0.17757	0.10068	0.26533	0.92980




Figure B.3: Comparison of keywords frequency: lemmatized vs stemmed keywords.




Figure B.4: Average Fuel Mix composition in 2023 for selected electricity companies.

B.2 Descriptive Statistics




Figure B.5: Number of Electricity Providers (Source: GSE)

Composizione del mix iniziale nazionale utilizzato per la produzione dell'energia elettrica immessa nel sistema elettrico italiano nel 2022*		Composizione del mix iniziale nazionale utilizzato per la produzione dell'energia elettrica immessa nel sistema elettrico italiano nel 2023**	
Fonti primarie utilizzate	%	Fonti primarie utilizzate	%
- Fonti rinnovabili	36,95%	- Fonti rinnovabili	46,31%
- Carbone	8,34%	- Carbone	5,27%
- Gas naturale	48,66%	- Gas naturale	42,99%
- Prodotti petroliferi	1,16%	- Prodotti petroliferi	0,90%
- Nucleare	0%	- Nucleare	0%
- Altre fonti	4,89%	- Altre fonti	4,53%

*dato a consuntivo **dato a pre-consuntivo

Figure B.6: National Fuel Mix 2022-2023. Souce GSE.

B.3 Greenwashing Index

Standardized GW

As a robust check to the normalized index, we decided to implement an alternative approach, which is very common in similar literature. We standardized each component in the following way:

$$\text{Disclosure}_{stand} = \frac{\text{Disclosure}_i - \text{Disclosure}_{mean}}{\text{Disclosure}_{sd}} \quad (\text{B.1})$$

This equation adjusts each company’s disclosure score by the mean and standard deviation of all scores, allowing to factor in the variability across companies. The performance score is adjusted similarly. This standardized index is graphically represented in Figure B.7.




Figure B.7: Standardized GW index distribution

A comparison of the GW Indices

Table B.1 compares the descriptive statistics of the Greenwashing indices. Both the Peer-based and the Normalized indices range from -1 to 1. The range and standard deviation are slightly broader for the Standardized Index. Examination of the distributional characteristics suggests that we cannot presume normality. We conducted both Kolmogorov-Smirnov and Jarque-Bera. As indicated in the results presented in Table B.2, only the Percentile-based GW index appears to exhibit normal distribution traits. However, this finding is not robust as it is only supported by the Kolmogorov-Smirnov test and not confirmed by the Jarque-Bera test. We conclude that the Percentile-based index shows the most favorable statistical properties, although none of the indices is normally distributed.

When comparing mean values, the peer-based and the standardized indices have mean 0, while the normalized version has mean equal to 0.1, depicting on average higher tendency

to greenwash. Our focus, however, is particularly on the positive range of values (greater than 0) to determine the extent and frequency of greenwashing among companies. This broad range of values suggests significant variation among individual cases. Specifically, we aim to analyze the frequency with which companies register positive values to discern whether they tend to exhibit low-level greenwashing, closer to 0, or whether they are more prone to significant greenwashing, with scores nearing 1. Figure B.8 illustrates the plots of positive greenwashing measure for each scenario. The Percentile-based index identifies 174 companies engaged in greenwashing; this result is robust when considering the Standardized version of the index. However, the Normalized version is more sensitive to detecting greenwashing behaviors, identifying 254 firms with positive greenwashing scores. Notably, the majority of these are very close to 0, indicating minimal greenwashing behavior.

Each approach to measuring greenwashing has its advantages and disadvantages. The Peer-based index benefits from robust statistical properties and is not sensitive to outliers. However, it relies on the current sample, and each company’s performance is directly compared to that of its competitors. This means it reflects relative rather than absolute differences between companies, and adding or removing companies from the sample could significantly affect each company’s performance assessment.

This limitation is addressed by the Normalized and Standardized versions of the index, which incorporate absolute scores in their calculations of the Greenwashing measure. These methods provide a measure that reflects each company’s performance in isolation rather than in comparison. However, these approaches are susceptible to the influence of outliers, which can significantly skew the overall distribution of the score.


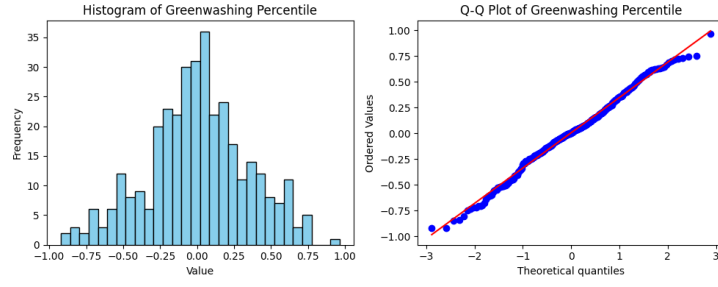




Figure B.8: Plot of GW indices: only positive values



(a) Percentile Greenwashing



(b) Standardized Greenwashing



(c) Normalized Greenwashing

Figure B.9: Distribution of Greenwashing Scores

Table B.2: Summary of Normality Tests for Greenwashing Scores

Index	Test Type	Statistic	P-value	H0: Normality
GW Percentile Scores	Jarque-Bera	6.39	0.041	Not accepted
	Kolmogorov-Smirnov	0.04563	0.43947	Accepted
GW Standardized Scores	Jarque-Bera	10.25	0.006	Not accepted
	Kolmogorov-Smirnov	0.09664	0.00250	Not accepted
GW Normalized Scores	Jarque-Bera	99.09	3.04e-22	Not accepted
	Kolmogorov-Smirnov	0.18560	3.79e-11	Not accepted

B.4 Ateco 2007 Classification

ATECO 2007	N.	Business description (English)
192030	1	Mixing of liquefied petroleum gases (LPG) and their bottling.
279003	1	Manufacture of electrical equipment for motor vehicles.
332009	1	Repair of other machinery and equipment n.e.c.
351000	6	Production, transmission and distribution of electricity (general).
351100	18	Production of electricity.
351200	1	Transmission of electricity.
351300	3	Distribution of electricity.
351400	165	Trade of electricity.
352000	3	Production of gas; distribution of gaseous fuels through mains (general).
352200	2	Distribution of gaseous fuels through mains.
352300	97	Trade of gas conveyed by mains.
412000	2	Construction of residential and non-residential buildings.
432101	3	Electrical installation.
467100	9	Wholesale of solid, liquid and gaseous fuels and related products.
473000	1	Retail sale of fuel in specialised stores.
477840	3	Retail sale of electrical household appliances in specialised stores.
479910	1	Retail sale via mail order houses or via Internet (any kind of products).
591100	1	Motion picture production activities.
619010	2	Internet service provider activities.
619099	1	Other information technology service activities n.e.c.
620100	2	Computer programming activities.
631111	1	Data processing, hosting and related activities: Web portals.
631119	1	Data processing, hosting and related activities n.e.c.
641930	1	Financial leasing.
642000	1	Activities of holding companies.
682001	1	Real estate activities with own or leased property.
711000	1	Architectural and engineering activities and related technical consultancy.
711210	1	Engineering related to building construction.
721909	1	Other research and experimental development on natural sciences and engineering.
749000	1	Other professional, scientific and technical activities n.e.c.
749032	1	Industrial chemical and physical analysis laboratories.
749090	1	Other professional activities n.e.c. (e.g., translation and interpretation).
749093	1	Technical testing and analysis.
829999	2	Other business support service activities n.e.c.

Table B.3: List of ATECO 2007 codes (source: Italian ATECO 2007 official list).

B.5 Sources for Words' List Creation

Title	Author	Journal
A review of renewable energy sources, sustainability issues and climate change mitigation	Owusu et al.(2016)	Cogent Engineering
Climate change and sustainability in the energy sector	Batruch (2017)	The Journal of World Energy Law & Business
Challenges in the decarbonization of the energy sector	Papadis and Tsatsaronis (2020)	Energy
Sustainability—Concept and its application in the energy sector.	Khan and Sahabuddin (2022)	Renewable Energy and Sustainability
Environmental sustainability: A definition for environmental professionals.	Morelli, J. (2011)	Journal of environmental sustainability
The Concept of Environmental Sustainability	Goodland (1995)	Annual Review of Ecology and Systematics
The sustainability concept: A review focusing on energy	Muniz et al. (2023)	Sustainability
Sustainable Development Goals	ONU (2015)	https://www.un.org/sustainabledevelopment/inequality

Table B.4: Articles for sustainability-related keywords

Table B.5: Glossari sulla Sostenibilità

Dictionary	Webpage
(i)	https://greeninstitute.ng/greenglossary
(ii)	https://relatedwords.io/sustainability
(iii)	https://www.british-business-bank.co.uk/business-guidance/guidancearticles/sustainability/glossary-of-terms#230548828-2052734356
(iv)	https://quantis.com/what-we-think/resources/abcs-of-sustainability-glossary/

Table B.8: Variable Definitions, Formulas, Sources and References

Variable	Formula / Computation	Source	Reference
GW Index	Disclosure - Performance (Section 5)	GSE Fuel Mix + Websites	Yu et al. (2020) (main)
Profit Margin _{t-1}	(Net Income / Revenues) × 100	AIDA	Nariswari and Nugraha (2020)
ROA _{t-1}	EBIT / Total Assets	AIDA	Liao et al. (2018); Uyar et al. (2020); Cariola et al. (2020)
Leverage _{t-1}	(Short-term Debt + Long-term Debt) / Total Assets	AIDA	Cariola et al. (2020)
Risk Management _{t-1}	Capital Expenditure / Equity	AIDA	—
Years of Activity	2023 - Year of Foundation	AIDA	Blasi et al. (2021)
Log Age	Log(Years of Activity)	AIDA	Blasi et al. (2021)
Size	Log(Total Assets)	AIDA	Blasi et al. (2021)
Provider	1 if ATECO = 35.14, 0 otherwise	AIDA	—
Other	1 if ATECO in {35.11, 35.12, 35.13}, 0 otherwise	AIDA	—
Electric	1 if ATECO in {35.11, 35.12, 35.13, 35.14}, 0 otherwise	AIDA	—

Table B.9: Correlation Matrix

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Greenwashing Peer	1.00	0.87***	0.74***	0.07	0.10	-0.05	-0.16**	-0.04	0.03	-0.05	-0.01	0.09	-0.01	0.02
Greenwashing Norm	0.87***	1.00	0.90***	0.09	0.12*	-0.02	-0.05	-0.02	0.03	-0.00	-0.03	0.06	-0.03	0.02
Greenwashing Min-Max	0.74***	0.90***	1.00	0.09	0.14**	-0.03	-0.11*	-0.07	0.01	0.02	0.01	0.02	-0.02	-0.03
Profit Margin _{t-1}	0.07	0.09	0.09	1.00	0.79***	-0.00	-0.00	-0.02	0.15**	-0.36***	-0.07	-0.07	-0.01	-0.08
ROA _{t-1}	0.10	0.12*	0.14**	0.79***	1.00	-0.14**	-0.12*	-0.15**	0.09	-0.23***	-0.04	0.01	-0.04	-0.14**
Age	-0.05	-0.02	-0.03	-0.00	-0.14**	1.00	0.56***	0.21***	0.01	-0.23***	0.07	-0.04	0.07	0.08
Size _{t-1}	-0.16**	-0.05	-0.11*	-0.00	-0.12*	0.56***	1.00	0.55***	0.09	-0.00	0.10	-0.25***	0.24***	0.23***
Revenues _{t-1}	-0.04	-0.02	-0.07	-0.02	-0.15**	0.21***	0.55***	1.00	-0.09	-0.05	0.20***	-0.12*	0.57***	0.26***
Net Income _{t-1}	0.03	0.03	0.01	0.15**	0.09	0.01	0.09	-0.09	1.00	-0.14*	-0.01	-0.02	-0.29***	0.12*
Financial Leverage _{t-1}	-0.05	-0.00	0.02	-0.36***	-0.23***	-0.23***	-0.00	-0.05	-0.14*	1.00	0.07	-0.12*	0.03	-0.11*
Risk Management _{t-1}	-0.01	-0.03	0.01	-0.07	-0.04	0.07	0.10	0.20***	-0.01	0.07	1.00	-0.02	0.12*	-0.06
Market Share (kWh)	-0.01	-0.03	-0.02	-0.01	-0.04	0.07	0.24***	0.57***	-0.29***	0.03	0.12*	1.00	0.16**	0.03
Report (dummy)	-0.18**	-0.09	-0.14*	-0.05	-0.13*	0.19***	0.46***	0.41***	0.02	-0.01	0.10	0.16**	1.00	0.11*
Foreign (dummy)	0.02	0.02	-0.03	-0.08	-0.14**	0.08	0.23***	0.26***	0.12*	-0.11*	-0.06	0.03	0.11*	1.00

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Variable ordering: (1) Greenwashing Peer; (2) Greenwashing Norm; (3) Greenwashing Min-Max; (4) Profit Margin_{t-1}; (5) ROA_{t-1}; (6) Age (log years); (7) Size_{t-1}; (8) Revenues_{t-1}; (9) Net Income_{t-1}; (10) Financial Leverage_{t-1}; (11) Risk Management_{t-1}; (12) Market Share (kWh); (13) Report; (14) Foreign.

B.6. ADDITIONAL RESULTS

Table B.10: OLS: Greenwashing (with Report as covariate)

Variable	Panel A (NPM, with Report)				Panel B (ROA, with Report)			
	GW Peer		GW Norm		GW Peer		GW Norm	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Provider (electr. sect.)	-0.009 (0.041)	–	-0.067* (0.034)	–	-0.010 (0.041)	–	-0.069** (0.034)	–
Other (electr. sect.)	-0.057 (0.081)	–	-0.181* (0.099)	–	-0.058 (0.080)	–	-0.184* (0.097)	–
Electric sect.	–	-0.016 (0.037)	–	-0.088*** (0.033)	–	-0.017 (0.037)	–	-0.088*** (0.033)
Age (log years)	-0.004 (0.024)	-0.008 (0.023)	0.002 (0.023)	-0.008 (0.023)	-0.0001 (0.025)	-0.004 (0.024)	0.009 (0.023)	-0.001 (0.023)
Report	-0.147*** (0.047)	-0.144*** (0.047)	-0.104** (0.048)	-0.098** (0.048)	-0.142*** (0.047)	-0.139*** (0.046)	-0.097** (0.048)	-0.090* (0.048)
Financial leverage _{t-1}	-0.085 (0.108)	-0.074 (0.103)	0.051 (0.113)	0.076 (0.110)	-0.086 (0.103)	-0.075 (0.098)	0.051 (0.104)	0.075 (0.101)
NPM _{t-1}	0.004 (0.003)	0.004 (0.003)	0.006* (0.004)	0.006* (0.004)	–	–	–	–
ROA _{t-1}	–	–	–	–	0.003 (0.002)	0.003* (0.002)	0.005** (0.002)	0.005** (0.002)
Risk management _{t-1}	0.0003 (0.002)	0.0003 (0.002)	0.001 (0.001)	0.001 (0.001)	0.0002 (0.002)	0.0002 (0.002)	0.001 (0.001)	0.001 (0.001)
Market share (kWh)	0.001 (0.002)	0.001 (0.002)	0.00003 (0.001)	0.0003 (0.001)	0.001 (0.002)	0.001 (0.002)	0.0001 (0.001)	0.0004 (0.001)
North	-0.099*** (0.038)	-0.101*** (0.037)	-0.055* (0.033)	-0.058* (0.033)	-0.098*** (0.037)	-0.100*** (0.037)	-0.054* (0.032)	-0.057* (0.032)
Foreign	0.078 (0.068)	0.074 (0.068)	0.034 (0.070)	0.023 (0.072)	0.085 (0.069)	0.081 (0.068)	0.046 (0.071)	0.035 (0.073)
Constant	0.149 (0.109)	0.152 (0.110)	0.139 (0.115)	0.150 (0.117)	0.133 (0.109)	0.136 (0.110)	0.109 (0.111)	0.121 (0.114)
Observations	336	336	336	336	336	336	336	336
R ²	0.061	0.060	0.062	0.057	0.064	0.063	0.073	0.067
Adjusted R ²	0.032	0.034	0.033	0.031	0.035	0.037	0.044	0.041
F Statistic	2.110** (df = 10;325)	2.313** (df = 9;326)	2.161** (df = 10;325)	2.173** (df = 9;326)	2.211** (10;325)	2.422** (9;326)	2.560*** (10;325)	2.591*** (9;326)

Notes: Robust standard errors in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B.11: Post-Lasso OLS

	Greenwashing Peer (1)	Greenwashing Norm (2)
ROA _{t-1}	0.003* (0.001)	0.004** (0.002)
Size	-0.022*** (0.007)	-0.015** (0.007)
Electric sect.	—	-0.084*** (0.031)
Other (electric sect.)	—	-0.112 (0.095)
Constant	0.194*** (0.073)	0.281*** (0.072)
Observations	336	336
R ²	0.031	0.060
Adjusted R ²	0.025	0.048
F Statistic	5.334*** (df = 2; 333)	5.259*** (df = 4; 331)

Note: Robust standard errors in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B.12: OLS: Greenwashing by firm size and model specification

	Specification A		Specification B	
	Large (1)	Small/Medium (2)	Large (3)	Small/Medium (4)
Provider (electric)	0.028 (0.068)	-0.099*** (0.035)	—	—
Other (electric)	-0.087 (0.167)	-0.233** (0.113)	—	—
Electric sector	—	—	0.015 (0.066)	-0.120*** (0.034)
Profit margin _{t-1}	-0.008* (0.004)	0.009*** (0.003)	-0.008* (0.004)	0.008*** (0.003)
Constant	0.064* (0.037)	0.163*** (0.020)	0.059 (0.039)	0.169*** (0.019)
Observations	79	257	79	257
R ²	0.040	0.065	0.029	0.059
Adjusted R ²	0.001	0.054	0.003	0.052
Residual Std. Error	0.290 (df = 75)	0.297 (df = 253)	0.290 (df = 76)	0.297 (df = 254)
F Statistic	1.034 (df = 3;75)	5.886*** (df = 3;253)	1.134 (df = 2;76)	8.033*** (df = 2;254)

Note: Robust (HC1) standard errors in parentheses.
Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.




Figure B.11: Scatterplot of Greenwashing Norm Index by Size

Appendix C

Appendix of Chapter 3

C.1 Market Structure

The Day-Ahead Market (MGP) is divided into several sub-markets. All of these markets are operated by the Gestore dei Mercati Energetici (GME), which is also responsible for the Forward Market (MTE) and the Energy Accounts Platform (PCE).

- MGP – Day-Ahead Market (Mercato del Giorno Prima): The MGP hosts the majority of electricity trading transactions. Energy blocks for each hourly market period of the following day are traded here. Participants submit bids specifying the quantity and maximum/minimum price at which they are willing to buy or sell. After market closure, bids are accepted on the basis of economic merit, subject to transmission capacity limits between zones. The MGP operates as an auction market, not a continuous trading platform.

All accepted buy and sell bids—whether related to injection portfolios or withdrawal portfolios—are settled at the zonal marginal equilibrium price. This price is determined, for each hourly market period, by the intersection of the demand and supply curves and may differ between zones if transmission capacity limits are saturated.

For every accepted purchase bid relating to a withdrawal portfolio in a given market zone, the GME calculates a compensation component equal to the product of the accepted quantity and the difference between the zonal marginal price and the PUN (Italian Single National Price). This can be positive (credit) or negative (debit). GME acts as the central counterparty.

- MI – Intra-Day Market (Mercato Infragiornaliero): The MI allows participants to adjust the commercial positions defined in the MGP through additional buy and sell bids. It consists of three MI-A auction sessions and one continuous trading session (MI-XBID).

In the MI-A auctions, intraday interconnection capacity between all Italian market zones and interconnected foreign areas participating in market coupling is allocated simultaneously with the matching of buy and sell offers. The MI-XBID continuous session is divided into three non-overlapping phases. GME acts as the central counterparty.

- **MPEG – Daily Products Market (Mercato dei Prodotti Giornalieri):** The MPEG is the platform for trading daily electricity products with delivery obligations. All electricity market participants are automatically admitted. Trading is continuous, and GME acts as the central counterparty. The products currently tradable on MPEG are “unit price differential” contracts, with Baseload and Peak Load delivery profiles.
- **MSD – Ancillary Services and Balancing Market (Mercato per il Servizio di Bilanciamento e Redispatching):** Because electricity cannot be stored, the quantity generated must exactly match demand. The task of maintaining this balance (dispacciamento) is the responsibility of Terna S.p.A., which uses the MSD as its main tool. In this market, Terna acts as the central counterparty. When demand deviates from the scheduled level, Terna performs balancing operations—recording the difference between scheduled and actual consumption and sending updated dispatching orders to generation plants. MSD offers are accepted based on economic merit and remunerated according to the pay-as-bid principle (i.e., at the offered price). The MSD is structured into: the Scheduling Phase (MSD ex-ante), and the Balancing Market (MB), both conducted through separate sessions.
- **PN – Nomination Platform (Piattaforma di Nomina):** The PN is the platform used for the submission and management of nominations, i.e., schedules of electricity injections and withdrawals, in accordance with market outcomes and bilateral contracts.

The national markets, MGP and MI, are integrated respectively with the SDAC (PCR) and SIDC (XBID & IDA). These are European initiatives established under the CACM Regulation to implement single cross-border market coupling throughout Europe. Under CACM, GME is designated as the Nominated Electricity Market Operator (NEMO) for Italy, responsible for developing and managing the SDAC and SIDC together with other NEMOs and Transmission System Operators (TSOs) of EU Member States.

C.2 Tariff Example

Facciamo luce sui prezzi	
Stima spesa	67,79 €/mese**
Componente Energia	F1 0,097996 €/kWh*
	F2 0,121699 €/kWh*
	F3 0,095825 €/kWh*
<p>Determinata in base ai valori variabili mensilmente assunti dall'indice PUN Index GME, pari alla media dei prezzi zonali ponderata per le quantità acquistate in ciascuna zona di mercato calcolato direttamente dal Gestore dei Mercati Energetici S.p.A. (GME) pubblicato sul sito www.mercatoelettrico.org. I prezzi della Componente Energia fanno riferimento ai valori assunti dall'indice PUN Index GME nel mese di Maggio 2025.</p> <p>F1: lun-ven dalle 8 alle 19 F2: lun-ven dalle 7 alle 8 e dalle 19 alle 23 + sabato dalle 7 alle 23 F3: lun-sab dalle 23 alle 7 + domenica e giorni festivi 24/24h</p>	
Spread contrattuale	0,033 €/kWh*
Componente Commercializzazione	132 €/POD/anno

Figure C.1: Electricity ToU tariff example

C.3 Results

Table C.1: FE/RE Estimates for Peak (F1) vs Off-Peak (F2+F3) Consumption

Variable	Peak Consumption (F1)		Off-Peak Consumption (F2+F3)	
	FE	RE	FE	RE
<i>Intercept</i>	—	4.185*** (0.245)	—	5.881*** (0.226)
<i>Price Elasticities</i>				
$\ln P_{F1,t-1}$	-0.291 (0.183)	-0.276 (0.185)	0.339** (0.169)	0.348** (0.170)
$\ln P_{\text{offpeak}}$	0.333 (0.227)	0.321 (0.229)	-0.551** (0.209)	-0.553** (0.211)
<i>Contract Control</i>				
Power (kW)	—	0.129*** (0.017)	—	0.185*** (0.016)
<i>Weather Controls</i>				
$\ln(\text{Temp})$	0.399* (0.221)	0.472** (0.222)	0.399* (0.203)	0.480** (0.204)
$\ln(\text{TempMin})$	-0.096 (0.064)	-0.110* (0.065)	-0.095 (0.058)	-0.112* (0.059)
$\ln(\text{TempMax})$	-0.584** (0.195)	-0.651*** (0.195)	-0.540** (0.179)	-0.615*** (0.179)
<i>Geographic Controls (RE only)</i>				
Altitude Zone (2)	—	0.249** (0.115)	—	0.199* (0.107)
Altitude Zone (3)	—	0.414*** (0.070)	—	0.398*** (0.065)
Altitude Zone (4)	—	0.442*** (0.084)	—	0.424*** (0.078)
Altitude Zone (5)	—	0.569*** (0.083)	—	0.548*** (0.077)
Urbanization (2)	—	-0.016 (0.057)	—	-0.019 (0.053)
Urbanization (3)	—	-0.392*** (0.063)	—	-0.330*** (0.059)
Coastal	—	-0.320*** (0.071)	—	-0.290*** (0.066)
Island	—	0.042 (0.110)	—	0.030 (0.102)
Area Geogr. (1)	—	0.156** (0.068)	—	0.152** (0.063)
Area Geogr. (2)	—	0.283*** (0.057)	—	0.212*** (0.053)
Green contract	—	0.070** (0.023)	—	0.079*** (0.021)
<i>Time Controls</i>				
<i>(baseline: Jan 2023)</i>				
Mar 2023	-0.097*** (0.017)	-0.096*** (0.017)	-0.038** (0.016)	-0.036** (0.016)
May 2023	-0.059*** (0.018)	-0.058*** (0.019)	-0.050*** (0.017)	-0.047*** (0.017)
Jul 2023	0.323*** (0.043)	0.323*** (0.044)	0.225*** (0.040)	0.228*** (0.040)
Sep 2023	-0.005 (0.024)	-0.004 (0.024)	-0.041* (0.022)	-0.037* (0.022)
Nov 2023	-0.068*** (0.016)	-0.069*** (0.016)	-0.000 (0.014)	-0.0003 (0.015)
Jan 2024	0.061* (0.031)	0.060* (0.032)	-0.149*** (0.029)	-0.148*** (0.029)
Mar 2024	-0.056 (0.041)	-0.056 (0.042)	-0.157*** (0.038)	-0.154*** (0.038)
May 2024	-0.053 (0.036)	-0.050 (0.037)	-0.115*** (0.033)	-0.111*** (0.034)
Jul 2024	0.423*** (0.035)	0.424*** (0.036)	0.334*** (0.033)	0.337*** (0.033)
Observations	19,428	19,390	19,428	19,390
R^2_{within}	0.075	—	0.071	—
R^2	—	0.062	—	0.061
Hausman p -value		0.997		0.9371

Standard errors in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C.2: Fixed-Effects Regressions by Quartiles of Electricity Consumption

	Obs.	Estimate	Std. Error	95% CI	p-value
Panel A. Within-period quartiles					
Quartile 1	4,860	-0.243	0.090	[-0.419, -0.066]	0.007
Quartile 2	4,859	-0.075	0.022	[-0.118, -0.032]	<0.001
Quartile 3	4,856	-0.088	0.018	[-0.124, -0.052]	<0.001
Quartile 4	4,853	-0.092	0.039	[-0.168, -0.016]	0.018
Panel B. Quartiles of seasonally adjusted residuals					
Quartile 1	4,857	-0.235	0.084	[-0.399, -0.071]	0.005
Quartile 2	4,857	-0.045	0.023	[-0.090, 0.001]	0.057
Quartile 3	4,857	-0.076	0.018	[-0.112, -0.041]	<0.001
Quartile 4	4,857	-0.077	0.038	[-0.152, -0.001]	0.046

Note: Each model is estimated separately by quartile of household electricity consumption using fixed effects. Panel A reports results based on quartiles of within-period consumption, while Panel B is based on quartiles of seasonally adjusted residuals.

C.4 Additional Results: Heterogeneity Analysis

Table C.3: RE estimates with interaction terms – Total Consumption

Variable	Dependent variable: ln Total Consumption				
	(1) Area	(2) Altitude	(3) Urbanization	(4) Coastal	(5) Green
<i>Price Elasticity</i>					
$\ln PUN_{t-1}$	-0.155*** (0.047)	0.000 (0.047)	-0.098** (0.043)	-0.022 (0.039)	-0.185*** (0.038)
<i>Interaction terms</i>					
$\ln PUN \times$ Area North	0.043 (0.046)	—	—	—	—
$\ln PUN \times$ Area South	-0.005 (0.041)	—	—	—	—
$\ln PUN \times$ Altitude 2	—	-0.329*** (0.074)	—	—	—
$\ln PUN \times$ Altitude 3	—	0.017 (0.046)	—	—	—
$\ln PUN \times$ Altitude 4	—	-0.373*** (0.049)	—	—	—
$\ln PUN \times$ Altitude 5	—	-0.156*** (0.047)	—	—	—
$\ln PUN \times$ Urbanization 2	—	—	-0.047 (0.037)	—	—
$\ln PUN \times$ Urbanization 3	—	—	-0.141*** (0.042)	—	—
$\ln PUN \times$ Coastal	—	—	—	-0.266*** (0.034)	—
$\ln PUN \times$ Green contract	—	—	—	—	0.082** (0.032)
<i>Controls</i>					
Weather controls	Yes	Yes	Yes	Yes	Yes
Geographic controls	Yes	Yes	Yes	Yes	Yes
Seasonal dummies	Yes	Yes	Yes	Yes	Yes
Observations	19,390	19,390	19,390	19,390	19,390
R^2	0.062	0.067	0.063	0.064	0.063

Standard errors in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C.4: RE estimates with interaction terms – F1 Consumption

	Dependent variable: $\ln C_{F1}$ (Peak consumption)				
	(1) Area	(2) Altitude	(3) Urbanization	(4) Coastal	(5) Green
Panel A: Cross-price = P_{F2}					
Own price $\ln P_{F1,t-1}$	-0.392*	0.036	-0.228	0.085	-0.314*
	(0.220)	(0.203)	(0.204)	(0.195)	(0.191)
Cross price $\ln P_{F2,t-1}$	0.444*	0.024	0.279	-0.055	0.331
	(0.266)	(0.247)	(0.248)	(0.237)	(0.234)
$\ln P_{F1} \times$ Area 1	0.177	—	—	—	—
$\ln P_{F1} \times$ Area 2	0.094	—	—	—	—
$\ln P_{F2} \times$ Area 1	-0.197	—	—	—	—
$\ln P_{F2} \times$ Area 2	-0.104	—	—	—	—
$\ln P_{F1} \times$ Altitude 2	—	-0.560**	—	—	—
$\ln P_{F1} \times$ Altitude 3	—	0.070	—	—	—
$\ln P_{F1} \times$ Altitude 4	—	-0.562***	—	—	—
$\ln P_{F1} \times$ Altitude 5	—	-0.061	—	—	—
$\ln P_{F2} \times$ Altitude 2	—	0.420*	—	—	—
$\ln P_{F2} \times$ Altitude 3	—	-0.081	—	—	—
$\ln P_{F2} \times$ Altitude 4	—	0.451**	—	—	—
$\ln P_{F2} \times$ Altitude 5	—	-0.008	—	—	—
$\ln P_{F1} \times$ Urb. 2	—	—	-0.045	—	—
$\ln P_{F1} \times$ Urb. 3	—	—	-0.220*	—	—
$\ln P_{F2} \times$ Urb. 2	—	—	0.039	—	—
$\ln P_{F2} \times$ Urb. 3	—	—	0.196	—	—
$\ln P_{F1} \times$ Coastal	—	—	—	-0.428***	—
$\ln P_{F2} \times$ Coastal	—	—	—	0.357***	—
$\ln P_{F1} \times$ Green	—	—	—	—	0.173*
$\ln P_{F2} \times$ Green	—	—	—	—	-0.147
Controls (climate, geo, season)	Yes	Yes	Yes	Yes	Yes
Observations	19,390	19,390	19,390	19,390	19,390
R^2	0.060	0.064	0.060	0.062	0.060

Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C.5: RE estimates with interaction terms – F2 Consumption

	(1) Area	(2) Altitude	(3) Urbanization	(4) Coastal	(5) Green
Own price $\ln P_{F2,t-1}$	-0.29	-0.79***	-0.53**	-0.87***	-0.49**
Cross price $\ln P_{F1,t-1}$	0.17	0.65***	0.42**	0.71***	0.32*
$\ln P_{F2} \times$ Area 1	-0.25*	—	—	—	—
$\ln P_{F2} \times$ Area 2	-0.22*	—	—	—	—
$\ln P_{F1} \times$ Area 1	0.23*	—	—	—	—
$\ln P_{F1} \times$ Area 2	0.19*	—	—	—	—
$\ln P_{F2} \times$ Altitude 2	—	0.71***	—	—	—
$\ln P_{F2} \times$ Altitude 3	—	-0.18	—	—	—
$\ln P_{F2} \times$ Altitude 4	—	0.32*	—	—	—
$\ln P_{F2} \times$ Altitude 5	—	-0.11	—	—	—
$\ln P_{F1} \times$ Altitude 2	—	-0.81***	—	—	—
$\ln P_{F1} \times$ Altitude 3	—	0.18	—	—	—
$\ln P_{F1} \times$ Altitude 4	—	-0.45***	—	—	—
$\ln P_{F1} \times$ Altitude 5	—	0.04	—	—	—
$\ln P_{F2} \times$ Urb. 2	—	—	0.03	—	—
$\ln P_{F2} \times$ Urb. 3	—	—	0.27*	—	—
$\ln P_{F1} \times$ Urb. 2	—	—	-0.04	—	—
$\ln P_{F1} \times$ Urb. 3	—	—	-0.34***	—	—
$\ln P_{F2} \times$ Coastal	—	—	—	0.30**	—
$\ln P_{F1} \times$ Coastal	—	—	—	-0.38***	—
$\ln P_{F2} \times$ Green	—	—	—	—	-0.18*
$\ln P_{F1} \times$ Green	—	—	—	—	0.19*
Controls (climate, geo, season)	Yes	Yes	Yes	Yes	Yes
Observations	19,390	19,390	19,390	19,390	19,390
R^2	0.035	0.040	0.037	0.038	0.035

Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C.6: RE estimates with interaction terms – F3 Consumption

	(1) Area	(2) Altitude	(3) Urbanization	(4) Coastal	(5) Green
Own price $\ln P_{F3,t-1}$	-0.65**	-0.92***	-0.63**	-1.07***	-0.57**
Cross price $\ln P_{F1,t-1}$	0.38*	0.68***	0.39*	0.78***	0.29*
$\ln P_{F3} \times$ Area 1	-0.23	—	—	—	—
$\ln P_{F3} \times$ Area 2	0.03	—	—	—	—
$\ln P_{F1} \times$ Area 1	0.20	—	—	—	—
$\ln P_{F1} \times$ Area 2	-0.05	—	—	—	—
$\ln P_{F3} \times$ Altitude 2	—	0.66***	—	—	—
$\ln P_{F3} \times$ Altitude 3	—	-0.14	—	—	—
$\ln P_{F3} \times$ Altitude 4	—	0.46***	—	—	—
$\ln P_{F3} \times$ Altitude 5	—	0.05	—	—	—
$\ln P_{F1} \times$ Altitude 2	—	-0.68***	—	—	—
$\ln P_{F1} \times$ Altitude 3	—	0.10	—	—	—
$\ln P_{F1} \times$ Altitude 4	—	-0.60***	—	—	—
$\ln P_{F1} \times$ Altitude 5	—	-0.14	—	—	—
$\ln P_{F3} \times$ Urb. 2	—	—	0.03	—	—
$\ln P_{F3} \times$ Urb. 3	—	—	0.11	—	—
$\ln P_{F1} \times$ Urb. 2	—	—	-0.07	—	—
$\ln P_{F1} \times$ Urb. 3	—	—	-0.17	—	—
$\ln P_{F3} \times$ Coastal	—	—	—	0.49***	—
$\ln P_{F1} \times$ Coastal	—	—	—	-0.55***	—
$\ln P_{F3} \times$ Green	—	—	—	—	-0.15
$\ln P_{F1} \times$ Green	—	—	—	—	0.16*
Controls (climate, geo, season)	Yes	Yes	Yes	Yes	Yes
Observations	19,390	19,390	19,390	19,390	19,390
R^2	0.035	0.075	0.070	0.074	0.070

Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C.7: FE and RE estimates with Income lag

	(1) RE Total	(2) RE F1 (F2)	(3) RE F1 (F3)	(4) RE F2 (F1)	(5) RE F3 (F1)
<i>Price elasticities</i>					
$\ln PUN_{t-1}$ (tot. price)	-0.075* (0.039)	—	—	—	—
Own price lag	—	-0.267 (0.188)	-0.256 (0.181)	-0.495** (0.226)	-0.567*** (0.210)
Cross price lag	—	0.346 (0.230)	0.344 (0.229)	0.383** (0.185)	0.340** (0.167)
<i>Income effect</i>					
$\ln Income_{t-1}$	0.752*** (0.157)	0.799*** (0.168)	0.799*** (0.168)	0.785*** (0.167)	0.745*** (0.155)
<i>Controls</i>					
Weather controls	Yes	Yes	Yes	Yes	Yes
Geographic/demographic	Yes	Yes	Yes	Yes	Yes
Seasonal dummies	Yes	Yes	Yes	Yes	Yes
Observations	19,390	19,390	19,390	19,390	19,390
R^2	0.063	0.061	0.061	0.036	0.070
Model	RE	RE	RE	RE	RE

Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

List of Tables

- 1.1 Descriptive Statistics 14
- 1.2 Sender and Receiver Experiment: behavioral type analysis 20
- 1.3 CRT-2 results for Sender and Receiver’s experiment 22
- 1.4 Sender Experiment: linear regression with belief as outcome variable 23
- 1.5 Payoff matrix (B = 5). 25
- 1.6 Welfare e Deadweight Loss per profilo (B = 5) 25
- 1.7 Expected payoffs, welfare and DWL across scenarios (B = 5) 26

- 2.1 Business Analysis by Size 51
- 2.2 OLS: Greenwashing 53
- 2.3 Models with interaction term: Profit-margin effect by firm’s size 55

- 3.1 Descriptive Statistics 77
- 3.2 FE/RE Estimates – Total Electricity Consumption Model 78
- 3.3 FE/RE Estimates – Consumption in Peak Band (F1) 82
- 3.4 FE/RE Estimates for ln Consumption in Off-Peak Bands (F2, F3) 83
- 3.5 Summary of significant own- and cross-price elasticities by time band 85

- A.1 Sender Experiment: average disclosure rates 90
- A.2 Sender Experiment: average disclosure rates (between-subject) 91
- A.3 Sender Experiment: χ^2 test results on disclosure rates across different treatments 91
- A.4 Receiver Experiment: average guesses per group 92
- A.5 Receiver Experiment: Kruskal-Wallis Test and Post hoc Mann-Whitney U Results for Guesses: Receiver 92
- A.6 Sender Experiment: Average Guesses per Group 92
- A.7 Sender Experiment: Kruskal-Wallis and Post-Hoc Mann-Whitney U Test on sender’s belief 93
- A.8 Sender Experiment: Sender’s belief and non-disclosure decision across different coarseness 95
- A.9 Receiver Experiment: Average first order beliefs 97

A.10 Receiver Experiment: Average second-order beliefs	97
A.11 Kruskal-Wallis Test and Post hoc Mann-Whitney U Results for HOB: Receiver	99
A.12 Sender Experiment: OLS with average disclosures as outcome (including Age)	100
A.13 Sender Experiment: Logit regressions with behavioral types as outcome . .	100
A.14 Receiver Experiment: OLS regressions with guess as outcome	101
B.1 Descriptive Statistics for Greenwashing Analysis (N = 354)	104
B.2 Summary of Normality Tests for Greenwashing Scores	108
B.3 List of ATECO 2007 codes (source: Italian ATECO 2007 official list). . . .	109
B.4 Articles for sustainability-related keywords	110
B.5 Glossari sulla Sostenibilità	110
B.6 Descriptive Statistics	111
B.7 Mean Values by Firm Size	111
B.8 Variable Definitions, Formulas, Sources and References	112
B.9 Correlation Matrix	112
B.10 OLS: Greenwashing (with Report as covariate)	113
B.11 Post-Lasso OLS	114
B.12 OLS: Greenwashing by firm size and model specification	114
C.1 FE/RE Estimates for Peak (F1) vs Off-Peak (F2+F3) Consumption	119
C.2 Fixed-Effects Regressions by Quartiles of Electricity Consumption	120
C.3 RE estimates with interaction terms – Total Consumption	121
C.4 RE estimates with interaction terms – F1 Consumption	122
C.5 RE estimates with interaction terms – F2 Consumption	123
C.6 RE estimates with interaction terms – F3 Consumption	124
C.7 FE and RE estimates with Income lag	124

List of Figures

- 1.1 Randomization flow 10
- 1.2 Plots of disclosures at different granularity levels across treatment groups. . 16
- 1.3 Receiver Experiment: plots of receiver’s guess at different granularity levels across treatment groups. 18
- 1.4 Total welfare across scenarios ($B = 5$). The dashed line indicates the ($FU \times g=1$) benchmark. 27

- 2.1 Data collection process 44
- 2.2 Different GW indices 49

- 3.1 Bi-monthly PUN by time band 2023-2024 70
- 3.2 Bi-monthly electricity demand by time band 2023-2024 70
- 3.3 Differences in elasticity across levels of consumption. 80

- A.1 (a),(b),(c): Probability distributions in control group, (d),(e),(f): Probability distributions in Treatment 1, (g),(h),(i): Probability distributions in Treatment 2 89
- A.2 Payoff Table in TA 90
- A.3 Payoff Table in TB 90
- A.4 Payoff Table in TC 90
- A.5 (a), (c), (e) represents sender’s elicited beliefs about receiver’s guess; (b), (d), (f) represents the observed guess of receiver with no disclosure by sender. 94
- A.6 (a), (c), (e) represents receiver’s elicited beliefs about sender’s reporting decision; (b), (d), (f) represents the observed decision of the sender. 96
- A.7 Receiver Experiment: Receiver’s beliefs about sender’s expectation of receiver’s guess 98

- B.1 Different Disclosure Metrics 103
- B.2 Different Performance Metrics 103
- B.3 Comparison of keywords frequency: lemmatized vs stemmed keywords. . . 104
- B.4 Average Fuel Mix composition in 2023 for selected electricity companies. . 105
- B.5 Number of Electricity Providers (Source: GSE) 105

B.6	National Fuel Mix 2022-2023. Souce GSE.	105
B.7	Standardized GW index distribution	106
B.8	Plot of GW indices: only positive values	107
B.9	Distribution of Greenwashing Scores	108
B.10	Words Cloud of keywords extracted (English version).	111
B.11	Scatterplot of Greenwashing Norm Index by Size	115
C.1	Electricity ToU tariff example	118

Bibliography

- Akerlof, G. A. (1978). The market for “lemons”: Quality uncertainty and the market mechanism. In *Uncertainty in economics*, pages 235–251. Elsevier.
- Alberini, A. and Filippini, M. (2011). Response of residential electricity demand to price: The effect of measurement error. *Energy economics*, 33(5):889–895.
- Ali, S., Zahid, R. A., Battisti, E., and Vincurova, Z. (2025). Environmental innovation and corporate financial performance: the moderating role of greenwashing. *International Journal of Emerging Markets*.
- Amel-Zadeh, A., Chen, M., Mussalli, G., and Weinberg, M. (2021). Nlp for sdgs: measuring corporate alignment with the sustainable development goals. *Columbia Business School Research Paper*.
- Aubin, C., Fougere, D., Husson, E., and Ivaldi, M. (1995). Real-time pricing of electricity for residential customers: Econometric analysis of an experiment. *Journal of Applied Econometrics*, 10(S1):S171–S191.
- Azevedo, I. M. L., Morgan, M. G., and Lave, L. (2011). Residential and regional electricity consumption in the us and eu: How much will higher prices reduce co2 emissions? *The Electricity Journal*, 24(1):21–29.
- Bardazzi, R. and Pazienza, M. G. (2020). When i was your age: Generational effects on long-run residential energy consumption in italy. *Energy Research & Social Science*, 70:101611.
- Bazillier, R. and Vauday, J. (2009). The greenwashing machine: Is csr more than communication.
- Becker, G. S. (1965). A theory of the allocation of time. *The economic journal*, 75(299):493–517.
- Benndorf, V., Kübler, D., and Normann, H.-T. (2015). Privacy concerns, voluntary disclosure of information, and unraveling: An experiment. *European Economic Review*, 75:43–59.

- Beretta, S. and Bozzolan, S. (2004). A framework for the analysis of firm risk communication. *The International Journal of Accounting*, 39(3):265–288.
- Bernard, J.-T., Bolduc, D., and Yameogo, N.-D. (2011). A pseudo-panel data model of household electricity demand. *Resource and Energy Economics*, 33(1):315–325.
- Bigerna, S. (2012). Electricity demand elasticity in italy. *Atlantic Economic Journal*, 40(4):439–440.
- Bigerna, S. and Bollino, C. A. (2015). A system of hourly demand in the italian electricity market. *The Energy Journal*, 36(4):129–148.
- Bingler, J. A., Kraus, M., Leippold, M., and Webersinke, N. (2022). Cheap talk and cherry-picking: What climatebert has to say on corporate climate risk disclosures. *Finance Research Letters*, 47:102776.
- Blasi, S., Crisafulli, B., and Sedita, S. R. (2021). Selling circularity: Understanding the relationship between circularity promotion and the performance of manufacturing smes in italy. *Journal of Cleaner Production*, 303:127035.
- Blázquez, L., Boogen, N., and Filippini, M. (2013). Residential electricity demand in spain: New empirical evidence using aggregate data. *Energy economics*, 36:648–657.
- Boedijanto, F. J. O. and Delina, L. L. (2024). Potentials and challenges of artificial intelligence-supported greenwashing detection in the energy sector. *Energy Research & Social Science*, 115:103638.
- Bohi, D. R. and Zimmerman, M. B. (1984). An update on econometric studies of energy demand behavior. *Annu. Rev. Energy;(United States)*, 9.
- Boogen, N., Datta, S., and Filippini, M. (2017). Dynamic models of residential electricity demand: Evidence from switzerland. *Energy Strategy Reviews*, 18:85–92.
- Boogen, N., Datta, S., and Filippini, M. (2021). Estimating residential electricity demand: New empirical evidence. *Energy Policy*, 158:112561.
- Borenstein, S. (2009). To what electricity price do consumers respond? residential demand elasticity under increasing-block pricing. *Preliminary Draft April*, 30:95.
- Brons, M., Nijkamp, P., Pels, E., and Rietveld, P. (2008). A meta-analysis of the price elasticity of gasoline demand. a sur approach. *Energy economics*, 30(5):2105–2122.
- Burns, K. and Mountain, B. (2021). Do households respond to time-of-use tariffs? evidence from australia. *Energy Economics*, 95:105070.

- Burton-Chellew, M. N., El Mouden, C., and West, S. A. (2016). Conditional cooperation and confusion in public-goods experiments. *Proceedings of the National Academy of Sciences*, 113(5):1291–1296.
- Cai, H. and Wang, J. T.-Y. (2006). Overcommunication in strategic information transmission games. *Games and Economic Behavior*, 56(1):7–36.
- Cariola, A., Fasano, F., La Rocca, M., and Skatova, E. (2020). Environmental sustainability policies and the value of debt in eu smes: Empirical evidence from the energy sector. *Journal of Cleaner Production*, 275:123133.
- Chen, Y.-S. (2008). The positive effect of green intellectual capital on competitive advantages of firms. *Journal of business ethics*, 77:271–286.
- Cicala, S. (2023). Jue insight: Powering work from home. *Journal of Urban Economics*, 133:103474.
- Crawford, V. P. (2003). Lying for strategic advantage: Rational and boundedly rational misrepresentation of intentions. *American Economic Review*, 93(1):133–149.
- Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451.
- De Freitas Netto, S. V., Sobral, M. F. F., Ribeiro, A. R. B., and Soares, G. R. d. L. (2020). Concepts and forms of greenwashing: A systematic review. *Environmental Sciences Europe*, 32:1–12.
- Deaton, A. and Muellbauer, J. (1980). An almost ideal demand system. *The American economic review*, 70(3):312–326.
- Delmas, M. A. and Burbano, V. C. (2011). The drivers of greenwashing. *California management review*, 54(1):64–87.
- Dempere, J., Alamash, E., and Mattos, P. (2024). Unveiling the truth: greenwashing in sustainable finance. *Frontiers in Sustainability*, 5:1362051.
- Dickhaut, J. W., McCabe, K. A., and Mukherji, A. (1995). An experimental study of strategic information transmission. *Economic Theory*, 6:389–403.
- European Commission (2003). Commission recommendation 2003/361/ec of 6 may 2003 concerning the definition of micro, small and medium-sized enterprises. Official Journal of the European Union, L 124, 20 May 2003, 36–41.
- European Commission (2021). Screening of websites for ‘greenwashing’: Half of green claims lack evidence. Available online. Accessed: January 21, 2024, https://ec.europa.eu/commission/presscorner/detail/en/ip_21_269.

- European Commission (2024). Green claims. European Commission Environment. Accessed: January 22, 2025, https://environment.ec.europa.eu/topics/circular-economy/green-claims_en.
- European Securities and Markets Authority (ESMA) (2023). Progress report on greenwashing. Available from: [urlhttps://www.esma.europa.eu/sites/default/files/2023-06/ESMA30-1668416927-2498PROGRESSREPORT_ESMA_response_tOCOM_RfI_on_greenwashing_risks.pdf](https://www.esma.europa.eu/sites/default/files/2023-06/ESMA30-1668416927-2498PROGRESSREPORT_ESMA_response_tOCOM_RfI_on_greenwashing_risks.pdf).
- Faiella, I. and Lavecchia, L. (2021). Households' energy demand and the effects of carbon pricing in italy. *Bank of Italy Occasional Paper*, (614).
- Faruqui, A. and Malko, J. R. (1983). The residential demand for electricity by time-of-use: a survey of twelve experiments with peak load pricing. *Energy*, 8(10):781–795.
- Federal Trade Commission (2012). Ftc issues revised green guides. Federal Trade Commission. Accessed: January 23, 2025, <https://www.ftc.gov/news-events/news/press-releases/2012/10/ftc-issues-revised-green-guides>.
- Ferraro, P. J., Rondeau, D., and Poe, G. L. (2003). Detecting other-regarding behavior with virtual players. *Journal of Economic Behavior & Organization*, 51(1):99–109.
- Fiandrino, S. and Tonelli, A. (2021). A text-mining analysis on the review of the non-financial reporting directive: bringing value creation for stakeholders into accounting. *Sustainability*, 13(2):763.
- Filippini, M. (1995a). Electricity demand by time of use an application of the household aids model. *Energy Economics*, 17(3):197–204.
- Filippini, M. (1995b). Swiss residential demand for electricity by time-of-use. *Resource and Energy Economics*, 17(3):281–290.
- Filippini, M. (2011). Short-and long-run time-of-use price elasticities in swiss residential electricity demand. *Energy policy*, 39(10):5811–5817.
- Filippini, M. and Pachauri, S. (2004). Elasticities of electricity demand in urban indian households. *Energy policy*, 32(3):429–436.
- Flaig, G. (1990). Household production and the short-and long-run demand for electricity. *Energy economics*, 12(2):116–121.
- Forsythe, R., Isaac, R. M., and Palfrey, T. R. (1989). Theories and tests of " blind bidding" in sealed-bid auctions. *The Rand Journal of Economics*, pages 214–238.

- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic perspectives*, 19(4):25–42.
- Ghitti, M., Gianfrate, G., and Palma, L. (2024). The agency of greenwashing. *Journal of Management and Governance*, 28(3):905–941.
- Gorovaia, N. and Makrominas, M. (2024). Identifying greenwashing in corporate-social responsibility reports using natural-language processing. *European Financial Management*.
- Gourier, E. and Mathurin, H. (2024). A greenwashing index. *Available at SSRN 4715053*.
- Gram-Hanssen, K. (2011). Understanding change and continuity in residential energy consumption. *Journal of consumer culture*, 11(1):61–78.
- Grewal, J., Richardson, G. D., and Wang, J. (2022). Effects of mandatory carbon reporting on unrepresentative environmental disclosures. *Available at SSRN 4166184*.
- Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *The Journal of law and Economics*, 24(3):461–483.
- Gull, A. A., Hussain, N., Khan, S. A., Khan, Z., and Saeed, A. (2023). Governing corporate social responsibility decoupling: The effect of the governance committee on corporate social responsibility decoupling. *Journal of Business Ethics*, 185(2):349–374.
- Gyamfi, S., Krumdieck, S., and Urmee, T. (2013). Residential peak electricity demand response—highlights of some behavioural issues. *Renewable and Sustainable Energy Reviews*, 25:71–77.
- Havranek, T., Irsova, Z., and Janda, K. (2012). Demand for gasoline is more price-inelastic than commonly thought. *Energy Economics*, 34(1):201–207.
- Holtedahl, P. and Joutz, F. L. (2004). Residential electricity demand in taiwan. *Energy economics*, 26(2):201–224.
- Hu, J., Wu, H., and Ying, S. X. (2022). Environmental regulation, market forces, and corporate environmental responsibility: Evidence from the implementation of cleaner production standards in china. *Journal of Business Research*, 150:606–622.
- Huang, A. H., Wang, H., and Yang, Y. (2023). Finbert: A large language model for extracting information from financial text. *Contemporary Accounting Research*, 40(2):806–841.
- Huang, R. and Huang, Y. (2020). Does internal control contribute to a firm’s green information disclosure? evidence from china. *Sustainability*, 12(8):3197.

- Hung, M.-F. and Huang, T.-H. (2015). Dynamic demand for residential electricity in taiwan under seasonality and increasing-block pricing. *Energy Economics*, 48:168–177.
- International Energy Agency (2018). *The Future of Cooling*. International Energy Agency, Paris.
- IPCC (2014). Climate change 2014: Mitigation of climate change.
- Ito, K. (2014). Do consumers respond to marginal or average price? evidence from nonlinear electricity pricing. *American Economic Review*, 104(2):537–563.
- Jia, F. and Li, Z. (2023). Development of green finance, tourism, and corporate access to bank loans in china. *Environmental Science and Pollution Research*, 30(30):75532–75548.
- Jin, G. Z., Luca, M., and Martin, D. (2021). Is no news (perceived as) bad news? an experimental investigation of information disclosure. *American Economic Journal: Microeconomics*, 13(2):141–173.
- Jin, Y., Li, H., and Lu, J. (2024). How does energy policy uncertainty perception affect corporate greenwashing behaviors? evidence from china’s energy companies. *Energy*, 306:132403.
- Jovanovic, B. (1982). Truthful disclosure of information. *The Bell Journal of Economics*, pages 36–44.
- Kang, H. and Kim, J. (2022). Analyzing and visualizing text information in corporate sustainability reports using natural language processing methods. *Applied Sciences*, 12(11):5614.
- Karaman, A. S., Orazalin, N., Uyar, A., and Shahbaz, M. (2021). Csr achievement, reporting, and assurance in the energy sector: Does economic development matter? *Energy Policy*, 149:112007.
- Kim, J., Lee, S., and Jang, H. (2022). Lessons from residential electricity demand analysis on the time of use pricing experiment in south korea. *Energy Economics*, 113:106224.
- Kugler, T., Bornstein, G., Kocher, M. G., and Sutter, M. (2007). Trust between individuals and groups: Groups are less trusting than individuals but just as trustworthy. *Journal of Economic psychology*, 28(6):646–657.
- Labandeira, X., Labeaga, J. M., and López-Otero, X. (2017). A meta-analysis on the price elasticity of energy demand. *Energy policy*, 102:549–568.
- Lagasio, V. (2024). Measuring greenwashing: The greenwashing severity index. *Available at SSRN 4582917*.

- Lee, R. and Kim, J. (2021). Developing a social index for measuring the public opinion regarding the attainment of sustainable development goals. *Social Indicators Research*, 156:201–221.
- Li, R., Zhong, M., and Shahab, Y. (2023a). Can penalties for environmental violations deter firms from engaging in greenwashing? *Business Ethics, the Environment & Responsibility*.
- Li, Y. X. and Schipper, B. C. (2020). Strategic reasoning in persuasion games: An experiment. *Games and Economic Behavior*, 121:329–367.
- Li, Z., Ling, Z., and Xu, S. (2023b). When firms talk, do they act? the impact of environmental strategies and actions on executive promotion in china. *China Economic Review*, 82:102061.
- Liao, F., Sun, Y., and Xu, S. (2023). Financial report comment letters and greenwashing in environmental, social and governance disclosures: Evidence from china. *Energy Economics*, 127:107122.
- Liao, L., Lin, T., and Zhang, Y. (2018). Corporate board and corporate social responsibility assurance: Evidence from china. *Journal of business ethics*, 150:211–225.
- Lin, B. and Xie, Y. (2024). How feed-in-tariff subsidies affect renewable energy investments in china? new evidence from firm-level data. *Energy*, 294:130853.
- Lublóy, Á., Keresztúri, J. L., and Berlinger, E. (2025). Quantifying firm-level greenwashing: A systematic literature review. *Journal of Environmental Management*, 373:123399.
- Luccioni, A., Baylor, E., and Duchene, N. (2020). Analyzing sustainability reports using natural language processing. *arXiv preprint arXiv:2011.08073*.
- Lyon, T. P. and Montgomery, A. W. (2015). The means and end of greenwash. *Organization & environment*, 28(2):223–249.
- March, C. (2019). *The behavioral economics of artificial intelligence: Lessons from experiments with computer players*. Number 154. BERG Working Paper Series.
- Marquis, C., Toffel, M. W., and Zhou, Y. (2016). Scrutiny, norms, and selective disclosure: A global study of greenwashing. *Organization Science*, 27(2):483–504.
- Matejek, S. and Gössling, T. (2014). Beyond legitimacy: A case study in bp’s “green lashing”. *Journal of Business Ethics*, 120:571–584.
- Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, pages 380–391.

- Miller, M. and Alberini, A. (2016). Sensitivity of price elasticity of demand to aggregation, unobserved heterogeneity, price trends, and price endogeneity: Evidence from us data. *Energy Policy*, 97:235–249.
- Montero, M. and Sheth, J. D. (2021). Naivety about hidden information: An experimental investigation. *Journal of Economic Behavior & Organization*, 192:92–116.
- Muratori, M. and Rizzoni, G. (2015). Residential demand response: Dynamic energy management and time-varying electricity pricing. *IEEE Transactions on Power systems*, 31(2):1108–1117.
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American economic review*, 85(5):1313–1326.
- Nakajima, T. (2010). The residential demand for electricity in japan: an examination using empirical panel analysis techniques. *Journal of Asian Economics*, 21(4):412–420.
- Narayan, P. K., Smyth, R., and Prasad, A. (2007). Electricity consumption in g7 countries: A panel cointegration analysis of residential demand elasticities. *Energy policy*, 35(9):4485–4494.
- Nariswari, T. N. and Nugraha, N. M. (2020). Profit growth: Impact of net profit margin, gross profit margin and total assets turnover. *International Journal of Finance & Banking Studies*, 9(4):87–96.
- Nielsen, Y. A., Thielmann, I., Zettler, I., and Pfattheicher, S. (2022). Sharing money with humans versus computers: On the role of honesty-humility and (non-) social preferences. *Social Psychological and Personality Science*, 13(6):1058–1068.
- Ning, X., Yim, D., and Khuntia, J. (2021). Online sustainability reporting and firm performance: Lessons learned from text mining. *Sustainability*, 13(3):1069.
- Nordin, J. A. (1976). A proposed modification of taylor’s demand analysis: comment. *The Bell Journal of Economics*, pages 719–721.
- Orange, E. and Cohen, A. M. (2010). From eco-friendly to eco-intelligent. *The Futurist*, 44(5):28.
- Paul, A. C., Myers, E. C., and Palmer, K. L. (2009). A partial adjustment model of us electricity demand by region, season, and sector.
- Pellini, E. (2021). Estimating income and price elasticities of residential electricity demand with autometrics. *Energy Economics*, 101:105411.

- Peng, D. and Kong, Q. (2024). Corporate green innovation under environmental regulation: The role of esg ratings and greenwashing. *Energy Economics*, 140:107971.
- Pepper, M., Jackson, T., and Uzzell, D. (2009). An examination of the values that motivate socially conscious and frugal consumer behaviours. *International journal of consumer studies*, 33(2):126–136.
- Pramana, R., Subroto, J. J., Gunawan, A. A. S., et al. (2022). Systematic literature review of stemming and lemmatization performance for sentence similarity. In *2022 IEEE 7th international conference on information technology and digital applications (ICITDA)*, pages 1–6. IEEE.
- Quan, X., Ke, Y., Qian, Y., and Zhang, Y. (2021). Ceo foreign experience and green innovation: Evidence from china. *Journal of Business Ethics*, pages 1–23.
- Reiss, P. C. and White, M. W. (2005). Household electricity demand, revisited. *The Review of Economic Studies*, 72(3):853–883.
- Romero-Jordán, D., Del Río, P., and Peñasco, C. (2016). An analysis of the welfare and distributive implications of factors influencing household electricity consumption. *Energy Policy*, 88:361–370.
- Rossi, N. (1988). Budget share demographic translation and the aggregate almost ideal demand system. *European Economic Review*, 32(6):1301–1318.
- Roulet, T. J. and Touboul, S. (2015). The intentions with which the road is paved: Attitudes to liberalism as determinants of greenwashing. *Journal of Business Ethics*, 128:305–320.
- Ruiz-Blanco, S., Romero, S., and Fernandez-Feijoo, B. (2022). Green, blue or black, but washing—what company characteristics determine greenwashing? *Environment, Development and Sustainability*, 24(3):4024–4045.
- Rupley, K. H., Brown, D., and Marshall, R. S. (2012). Governance, media and the quality of environmental disclosure. *Journal of Accounting and Public Policy*, 31(6):610–640.
- Sah, S. and Read, D. (2017). Disclosure and the dog that didn’t bark: Consumers are too forgiving of missing information. In *Academy of Management Proceedings*, volume 2017, page 12839. Academy of Management Briarcliff Manor, NY 10510.
- Schittekatte, T., Mallapragada, D., Joskow, P. L., and Schmalensee, R. (2024). Electricity retail rate design in a decarbonizing economy: an analysis of time-of-use and critical peak pricing. *The Energy Journal*, 45(3):25–56.

- Schulte, I. and Heindl, P. (2017). Price and income elasticities of residential energy demand in germany. *Energy Policy*, 102:512–528.
- Shin, J.-S. (1985). Perception of price when price information is costly: evidence from residential electricity demand. *The review of economics and statistics*, pages 591–598.
- Siano, A., Vollero, A., Conte, F., and Amabile, S. (2017). “more than words”: Expanding the taxonomy of greenwashing after the volkswagen scandal. *Journal of business research*, 71:27–37.
- Silva, S., Soares, I., and Pinho, C. (2017). Electricity demand response to price changes: The portuguese case taking into account income differences. *Energy Economics*, 65:335–342.
- Smeuninx, N., De Clerck, B., and Aerts, W. (2020). Measuring the readability of sustainability reports: A corpus-based analysis through standard formulae and nlp. *International Journal of Business Communication*, 57(1):52–85.
- Stahl, D. O. and Wilson, P. W. (1995). On players models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254.
- Tang, P., Wang, C., Jiang, Q., Liu, X., and Wang, J. (2023). Symbol or substance? environmental regulations and corporate environmental actions decoupling. *Journal of Environmental Management*, 346:118950.
- Taylor, R. G., McKean, J. R., and Young, R. A. (2004). Alternate price specifications for estimating residential water demand with fixed fees. *Land Economics*, 80(3):463–475.
- Thomson, K. S. and Oppenheimer, D. M. (2016). Investigating an alternate form of the cognitive reflection test. *Judgment and Decision making*, 11(1):99–113.
- Torrìti, J. (2012). Price-based demand side management: Assessing the impacts of time-of-use tariffs on residential electricity demand and peak shifting in northern italy. *Energy*, 44(1):576–583.
- UN Principles for Responsible Investment (2023). Us policy: Sec esg-related disclosure. UN Principles for Responsible Investment. Accessed: January 23, 2025, <https://www.unpri.org/policy/us-policy/sec-esg-related-disclosure>.
- Uyar, A., Karaman, A. S., and Kilic, M. (2020). Is corporate social responsibility reporting a tool of signaling or greenwashing? evidence from the worldwide logistics sector. *Journal of Cleaner Production*, 253:119997.

- Vollero, A., Palazzo, M., Siano, A., and Elving, W. J. (2016). Avoiding the greenwashing trap: between csr communication and stakeholder engagement. *International journal of innovation and sustainable development*, 10(2):120–140.
- Vos, J. (2009). Actions speak louder than words: Greenwashing in corporate america. *Notre Dame JL Ethics & Pub. Pol’y*, 23:673.
- Weber, J. L. (2018). Corporate social responsibility disclosure level, external assurance and cost of equity capital. *Journal of Financial Reporting and Accounting*, 16(4):694–724.
- Xu, W., Li, M., and Xu, S. (2023). Unveiling the “veil” of information disclosure: Sustainability reporting “greenwashing” and “shared value”. *PloS one*, 18(1):e0279904.
- Yu, E. P.-y., Van Luu, B., and Chen, C. H. (2020). Greenwashing in environmental, social and governance disclosures. *Research in International Business and Finance*, 52:101192.
- Zhang, D. (2022). Are firms motivated to greenwash by financial constraints? evidence from global firms’ data. *Journal of international financial management & accounting*, 33(3):459–479.
- Zhang, D. (2023). Can environmental monitoring power transition curb corporate greenwashing behavior? *Journal of Economic Behavior & Organization*, 212:199–218.
- Zhang, K., Pan, Z., Janardhanan, M., and Patel, I. (2023). Relationship analysis between greenwashing and environmental performance. *Environment, Development and Sustainability*, 25(8):7927–7957.
- Zhou, B. and Wang, Q. (2024). Fintech matters in sustainable finance: Does it redistribute the supply of financial services? *Journal of International Financial Markets, Institutions and Money*, 91:101913.
- Zhu, X., Li, L., Zhou, K., Zhang, X., and Yang, S. (2018). A meta-analysis on the price elasticity and income elasticity of residential electricity demand. *Journal of Cleaner Production*, 201:169–177.