Funded by the European Union
Grant Agreement No 101095424

FLASH

D3.8 Policy report including the determination of redistribution parameters through in-kind health transfers for different countries and needs groups

(Version 1, 30/06/2025)

DELIVERABLE: Deliverable 3.8: Policy report including the determination of redistribution parameters through in-kind health transfers for different countries and needs groups
WORK PACKAGE: WP3: Resource allocation, unmet health needs and redistribution
AUTHOR(S): Francesco Andreoli, Lorenzo De Menech, Cristina Elisa Orso
DUE DATE: 30/06/2025
ACTUAL SUBMISSION DATE: 30/06/2025
DISSEMINATION LEVEL Select the proper one and delete the remaining [] CO: Confidential, only for members of the consortium (including the Commission Services) [] CI: Classified as referred to in Commission Decision 2001/844/EC [x] PU: Public (must be available on the website)
GRANT AGREEMENT No: 101095424
PROJECT STARTING DATE: 01/01/2023
PROJECT DURATION: 48 months
COORDINATOR UNIVR - University of Verona

Quality of information - Disclaimer according to the Art. 17.3 of GA Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or EUROPEAN HEALTH AND DIGITAL EXECUTIVE AGENCY (HADEA). Neither the European Union nor the granting authority can be held responsible for them.

Glossary of terms

- In-kind transfer: is a non-monetary benefit that is provided by public institutions to individuals, mostly in the form of services or goods rather than cash. In the context of healthcare provision, it includes public services such as hospital treatment, GP consultations or diagnostic procedures that are free of charge or subsidized by the government. These transfers are important to promote welfare and ensure access to care to the overall population.
- Redistribution: refers to the deliberate reallocation of economic resources across individuals or households, usually from higher-income to lower-income groups, through mechanisms such as progressive taxation and social transfers (both cash and in-kind). In health economics, redistribution plays a central role in promoting equity of access to care, correcting market failures, and providing financial protection. Public healthcare systems, for instance, are inherently redistributive when financed through general taxation and offer universal or means-tested services regardless of individuals' ability to pay.
- Out-of-pocket expenditure (OOP expenditure): refers to direct payments made by individuals at the time of receiving healthcare services, which are not reimbursed by any form of health insurance or public subsidy. This includes expenditures for consultations, medications, diagnostic tests, or hospital stays that are either partially or fully excluded from public or private coverage.
- Inpatient and outpatient care: inpatient care involves medical services that require admission to a hospital or healthcare facility for at least one overnight stay. It typically includes interventions for more severe health conditions, such as surgeries, intensive treatments, or acute care episodes, and is often used as a marker of significant health events and resource use. In contrast, outpatient care refers to medical services provided without an overnight stay. It includes GP visits, specialist consultations, diagnostic exams, preventive screenings, and minor procedures. Outpatient care represents the most frequent and accessible form of contact with the healthcare system and often serves as the entry point for patients.

1 Introduction

Understanding how public healthcare provision interacts with income redistribution is a central issue in public economics (Melli et al., 2025; Honekamp and Possenriede, 2008). Redistribution traditionally refers to changes in the income distribution due to fiscal policy, such that the Lorenz curve after taxes and transfers (post-fisc) lies entirely on or above the Lorenz curve before fiscal intervention (pre-fisc), and strictly above it at least at one point. In cash, this is achieved through progressive tax-transfer schemes (e.g. increasing average tax rates on higher incomes). However, redistribution can also be pursued through in-kind provision of merit goods such as healthcare and education (van Doorslaer et al., 1999). Universal public healthcare systems, in which all individuals are entitled to a common bundle of services regardless of income, represent a canonical form of in-kind redistribution. These systems guarantee coverage to the poor but also extend benefits to the non-poor, raising important questions about the progressivity and efficiency of such policies (Besley and Coate (1991); Currie and Gahvari (2008)). Unlike cash transfers, in-kind benefits may generate differentially valued consumption bundles across the income distribution, and can lead to voluntary opt-out behavior among high-income users, who may substitute toward private alternatives (Epple and Romano (1996)).

A key question is how such universal systems can effectively target resources toward those in need without means-testing. Economic theory has highlighted self-targeting mechanisms as a solution to this dilemma. In essence, when in-kind transfers are used as the redistributive tool, the public provision can be designed as a separating device between rich and poor – a property known as the self-targeting feature of public provision. By offering a baseline level of service that is accessible to all but imposes certain costs (monetary or non-monetary) or quality limitations, the program induces higher-income individuals to opt out or supplement their consumption, thereby implicitly targeting benefits to lower-income groups. This concept, discussed by Currie and Gahvari (2008) among others, provides a theoretical rationale for universal in-kind transfers under conditions of imperfect information. When the government cannot perfectly observe who is poor, targeting through means-testing may be costly, cumbersome and somewhat controversial. In this setting, universal provision can induce self-selection: individuals reveal their type through their consumption choices. Specifically, the design of the public service—through features like baseline quality, waiting times, or provider choice—can be calibrated such that higher-income individuals, who place a greater value on time and convenience, opt to supplement or exit the public system by purchasing private care or paying out-of-pocket for upgrades. In contrast, lower-income individuals, facing tighter budget constraints and lower opportunity costs, are more likely to consume the public option despite its frictions. In this way, universal systems can replicate the redistributive effects of targeted transfers without explicit means-testing, using the structure of choice and the implicit costs of public provision to sort users by income.

In the context of healthcare, out-of-pocket (OOP) expenditures play a pivotal role in this self-targeting mechanism. OOP health spending refers to direct payments by individuals for services or quality enhancements not fully covered by the public system. Even under comprehensive national health coverage, households often face choices about whether to use the standard publicly provided services or to incur OOP costs for additional or higher-quality care (for example, paying for private providers, faster access, or better amenities). These spending choices are inherently linked to income and to preference heterogeneity in the population. Higher-income households generally have both a greater ability and a greater incentive to incur OOP expenditures, since the utility cost of public system limitations (such as waiting time or lower service intensity) is higher for them. Likewise, individuals with a stronger preference for health service quality or convenience (regardless of income) are more inclined to pay extra to supplement the basic care. Therefore, analyzing OOP expenditure as a function of income and preferences is essential for

understanding how universal health systems might achieve implicit redistribution. If the system is working as intended, we would expect to see wealthier or more quality-sensitive individuals self-select into higher OOP spending, effectively financing a larger share of their healthcare privately, while lower-income individuals rely on the publicly provided baseline. The distribution of OOP expenditures thus serves as an empirical fingerprint of the redistributive mechanism: at the extreme ends, it indicates the extent to which the public provision is supporting more the most deprived individuals (with the better-off "opting out" by paying from their own pocket) and reveals how different income groups consume healthcare under a universal scheme. Understanding these patterns is not only important for equity – ensuring that financial burdens do not fall disproportionately on the poor – but also for efficiency, as it sheds light on whether public resources are being utilized by the intended beneficiaries or being crowded out by those who could afford private care.

1.1 Redistribution via in-kind transfers: a theoretical and structural perspective

This deliverable is grounded in a rigorous theoretical framework that formalizes the above intuition. We draw on the model of Currie and Gahvari (2008) and related literature on in-kind transfers to articulate the conditions under which self-targeting occurs in a universal healthcare system. For clarity, consider a stylized economy with two types of individuals: high-income ("rich") and low-income ("poor"). Both types have access to a publicly provided good (healthcare) and derive utility from consumption (c) and health service g (with g_h denoting the quantity/quality consumed by a high-income individual and g_l that consumed by a low-income individual). The public healthcare system offers a baseline level of service g_l that is intended for the poor, while the rich can obtain a higher level g_h by paying additional costs out of pocket (for simplicity, assume both g_l and g_h are provided within the public system, but g_h involves extra expense). The difference between high-quality care g_h and basic care g_l manifests itself as OOP expenditure by the rich. In a first-best world with no informational constraints, the social planner could redistribute income and provide each group their optimal healthcare consumption (possibly with the rich paying more in taxes to finance the poor). However, with imperfect information – specifically, when the government cannot directly observe who is poor or rich – the allocation must satisfy an incentive compatibility constraint: the rich must prefer consuming the higher-quality option g_h with its associated OOP cost over masquerading as poor and consuming g_l at no (or minimal) cost. This constraint implies a second-best solution in which the publicly provided quantity for the poor is deliberately set below the rich household's preferred level (and indeed lower than what the poor would get in a first-best scenario without the constraint). In other words, there is a downward distortion in g_l – the quality/quantity of free public care is limited – in order to deter wealthy individuals from utilizing the free care. The rich then willingly pay OOP for the higher g_h, while the poor consume the baseline g_l. This separating equilibrium, which is essential for redistribution in universal programs, has concrete implications for observable expenditure patterns.

The theoretical framework yields three key predictions about OOP spending in a universal health system under a separating equilibrium:

- Income and OOP: If healthcare quality/quantity is a normal good (demand rises with income), then higher-income individuals will spend more out of pocket for additional care. Formally, the public provision level g chosen by an individual declines as income Y increases ($\partial g/\partial Y < 0$), so that an increase in income leads to greater OOP expenditure to reach a higher total consumption of healthcare. In plain terms, richer households supplement

the basic publicly provided care with privately financed services, increasing their OOP spending as their income grows.

- Differential responses by income level: since the baseline public provision g_l is set sub-optimally low for the poor (to maintain separation), an exogenous increase in income has a larger effect on OOP spending for a low-income household than for a high-income household. In other words, when a poor individual's income rises, they are likely to significantly increase their OOP health expenditures (since they can now afford to seek more or better care beyond the meager public baseline), whereas a rich individual is already consuming close to their desired level g_h and an extra dollar has a smaller marginal effect on their OOP spending. This prediction follows from the fact that at the margin, the poor face a steeper trade-off (higher marginal rate of substitution) in the separating equilibrium – they had been constrained by the low g_l, so additional income relaxes that constraint and induces a jump in private spending.

- Preference heterogeneity and OOP: Individuals with stronger preferences for health service quality (or lower tolerance for the public system's drawbacks) will exhibit larger increases in OOP spending for a given income change. If healthcare quality is a normal good, those with higher taste for quality are more inclined to spend out of pocket, so the effect of an income gain on OOP is amplified for these individuals. This point recognizes that even within income groups, heterogeneity in preferences (e.g. risk aversion, valuation of time and comfort) can lead to different OOP outcomes – a feature that is critical for understanding the dispersion of OOP spending in the data.

These theoretical predictions underscore the joint importance of income and preference heterogeneity in shaping how households utilize public versus private health services. They also illustrate the self-targeting mechanism in action: the rich and the quality-oriented self-select into higher private spending, leaving the publicly financed resources to be utilized mostly by the poor or less demanding users. The challenge, and the focus of this research, is to verify these patterns empirically and to quantify the relationship between income, preferences, and OOP expenditures within a universal health system. We aim to test whether the data reflect the predicted separation – for instance, does OOP spending indeed rise with income and disproportionately so for those with higher income potential or stronger quality preferences? Answering these questions is crucial for evaluating the equity of health financing and the effectiveness of implicit targeting in public healthcare.

To bring the theory to the data, we develop a structural econometric model that links individuals' income potential, healthcare demand, and out-of-pocket spending. The model is designed to capture the key elements of the theoretical framework while accounting for real-world complexities in the data. Specifically, we specify a system of equations for health service utilization and expenditures as follows: (i) a healthcare demand equation determining whether an individual utilizes any healthcare (public or otherwise), (ii) an OOP expenditure equation determining how much the individual spends out of pocket (if they have any demand), and (iii) an income equation capturing the determinants of the individual's income. These three components form a triangular system where income influences healthcare demand, and both income and demand influence OOP spending. A core feature of our approach is the treatment of unobserved heterogeneity. Factors such as innate health status, risk preferences, and human capital (income-earning ability) are not fully observable to the researcher, yet they affect both the ability to pay and the need or desire for additional healthcare. For example, unobserved human capital can raise a person's income while also affecting their health awareness or opportunity cost of time, thereby influencing healthcare utilization decisions. Similarly, unobserved preferences for quality (e.g.

how much someone values quicker access or better amenities) directly impact OOP spending independent of income. If not properly accounted for, these latent factors can confound the relationship between income and OOP – a classic endogeneity problem. High-income individuals might spend more out of pocket not only because of their income, but also because they possess certain traits (education, health knowledge, etc.) that correlate with earning ability and drive them to seek better care. Our structural model tackles this issue by explicitly modeling the covariance between unobserved determinants of income and health spending.

In this report, we proceed as follows: Section 2 and 3 introduce the data and outline the structural estimation model, detailing the identification strategy and the incorporation of heterogeneity through the SQTE approach. Section 4 discusses the estimation results, examining how income potential and preference heterogeneity quantitatively relate to healthcare utilization and OOP expenditure patterns. Section 5 concludes, highlighting the implications of our findings for the design of equitable health policies and the understanding of redistribution in public healthcare systems. Through this structured approach, the deliverable aims to provide both a theoretical and empirical contribution to the analysis of health sector redistribution, offering evidence on how universal health provision can achieve redistributive goals and what observable expenditure patterns emerge as a result. The analysis not only tests the theoretical predictions in real-world data but also quantifies the magnitude of these effects, thereby informing debates on public healthcare financing and the role of private payments in a universal system. Ultimately, the insights gained will be valuable for policymakers and researchers interested in the intersection of health, inequality, and public economics, shedding light on the conditions under which universal healthcare can serve as an effective instrument of redistribution.

2 Methods

2.1 Baseline OLS Model

We begin by estimating a baseline Ordinary Least Squares (OLS) regression to assess the relationship between individual income and out-of-pocket (OOP) healthcare expenditures. The model controls for a rich set of covariates and fixed effects to isolate the income–OOP gradient. Formally, the baseline specification can be written as:

$$(1) \quad \text{OOP}_{ict} = \alpha_0 + \beta y_{ict} + \mathbf{X}'_{ict}\gamma + \theta_c + \theta_t + \psi_{ct} + \varepsilon_{ict} \tag{1}$$

where OOP_{ict} is the OOP health expenditure for individual i in country c at time t, y_{ict} is the individual's total income (e.g. disposable household income, in PPP-adjusted terms), X_{ict} is a vector of observable characteristics (including age, gender, marital status, employment status, education, health need indicator, number of chronic conditions, government health expenditure (as % of total health exp.), θ_c and θ_t are country and wave (year) fixed effects, and ψ_{ct} captures any country-specific time trends or shocks. While this OLS model provides an initial estimate of the income effect, it may suffer from omitted-variable bias if there are unobserved determinants of healthcare demand correlated with income. For instance, risk preferences or differential access to healthcare could influence both income and OOP spending, violating the exogeneity of y_{ict} in Equation (1). As a result, the OLS estimate β might not reflect the true causal effect of income on OOP expenditures.

2.2 Instrumental Variable (IV) Approach

To address endogeneity concerns, we implement an instrumental variable strategy as our main empirical approach (Angrist and Krueger, 2001). In particular, we construct a simulated tax instrument that leverages cross-country and over-time variation in tax and benefit policies to obtain exogenous variation in disposable income. The intuition is that differences in fiscal policy (tax codes, pension rules, etc.) across countries and years will induce variation in individuals' net incomes that is unrelated to their personal characteristics or unobserved preferences. By using this policy-driven income variation as an instrument, we can isolate the causal impact of income on OOP spending. Concretely, we employ a two-stage least squares (2SLS) estimator. In the first stage, individual income is regressed on the instrument and controls, and in the second stage OOP spending is regressed on the predicted income from the first stage. The first-stage regression is:

$$(2) \quad y_{ict} = \delta_0 + \delta_1 Z_{ict} + \mathbf{X}'_{ict}\delta_2 + \theta_c + \theta_t + \mu_{ict} \tag{2}$$

where Z_{ict} is the constructed instrument (described in detail in next Section) that captures the component of income determined by tax-and-transfer policies, and the other terms are as defined above. The coefficient δ_1 measures the strength of the relationship between the instrument and actual income. In the second stage, we then estimate the causal effect of income on OOP expenditures as:

$$(3) \quad \text{OOP}_{ict} = \alpha_1 + \beta_{\text{IV}}\hat{y}_{ict} + \mathbf{X}'_{ict}\gamma_1 + \theta_c + \theta_t + \eta_{ict} \tag{3}$$

Here $\hat{y}_{ict}$ is the predicted income for individual i obtained from Equation (2), and β_{IV} is the IV estimate of the income effect on OOP spending. The control variables X_{ict} and fixed effects in (3) mirror those in (2) to ensure consistency across stages. By using only the policy-driven variation in income (through $\hat{y}_{ict}$) to explain OOP, this IV approach purges the effect of any idiosyncratic factors that jointly affect income and healthcare spending. In other words, it isolates exogenous income differences for identification.

2.3 Structural quantile treatment effect

Identification of the causal effect of income on health expenditures is achieved through an instrumental variable (IV) strategy embedded within a structural model of healthcare behavior. Specifically, we exploit exogenous variation in income that arises from policy-driven differences in tax-benefit systems across countries, age groups, and time periods. These institutional variations (captured by the instrument Z_{ict}) affect individual income potential without being directly correlated with unobserved health needs or preferences. This strategy mimics a natural experiment and allows us to recover causal effects net of endogeneity concerns such as reverse causality or omitted variables.

Beyond estimating average effects, our goal is to understand how income impacts out-of-pocket (OOP) health expenditures across the distribution of individuals, especially when such impacts are shaped by heterogeneity in human capital, preferences, and health shocks. This motivates the use of the Structural Quantile Treatment Effects (SQTE) approach Ma and Koenker (2006)), yields a Structural Quantile Treatment Effects (SQTE) model for OOP. It enables us to quantify how the elasticity of OOP with respect to income may differ, for example, between an individual at the lower end of the income-potential distribution versus one at the higher end, or between those with weaker vs stronger preferences for quality. By doing so, the model closely mirrors the theory: it can capture the predicted larger OOP response among low-income (but rising-income) individuals and among more preference-driven individuals. In summary, our empirical strategy is tailored to disentangle the interplay between income, self-selection, and health expenditures, providing a rigorous test of the self-targeting mechanism in a universal healthcare setting.

The following section formalizes a structural model to analyze the causal relationship between individual income, out-of-pocket (OOP) healthcare expenditure, and the demand for healthcare services. We focus in particular on the role of unobserved heterogeneity in shaping these relationships. For instance, latent traits such as human capital, family background, or health preferences may affect both income and healthcare behavior. These considerations motivate the implementation of the Structural Quantile Treatment Effect (SQTE) approach (Ma and Koenker (2006), which extends standard IV estimation to account for distributional heterogeneity and nonlinear interactions between income, preferences, and expenditure behavior.

We formalize the problem using the following system of structural equations:

$$OOP_{ict} = \alpha_1 Y_{ict} + \mathbf{X'i}\alpha_2 + uict \quad \text{if } D_{ict} = 1 \tag{4}$$

$$D_{ict} = \nu_1 Y_{ict} + \mathbf{X'i}\nu_2 + vict \tag{5}$$

$$Y_{ict} = \beta_1 Z_{ict} + \mathbf{X'i}\beta_2 + Uict \tag{6}$$

Equation (4) captures OOP spending for individuals with positive healthcare demand, allowing us to model how the effect of income varies across the OOP distribution. Equation (5) models latent healthcare demand, and equation (6) describes income as a function of observed covariates and the simulated tax instrument.

To address endogeneity and heterogeneity, we explicitly model the structure of the error terms in each equation 7, 8 and 9. Here, ϵ_Y represents unobserved human capital or income potential, which affects both earnings and downstream healthcare behavior. ϵ_D captures idiosyncratic health shocks, orthogonal to income potential, while ϵ_O reflects unobserved preferences for healthcare quality, such as willingness to pay for convenience or aversion to public queues. Crucially,

ϵ_O is assumed to be uncorrelated with both ϵ_Y and ϵ_D, ensuring that heterogeneity in preferences can be distinguished from differences in ability or health risk.

$$U_{ict} = \psi_{act} + \epsilon_Y \tag{7}$$

$$v_{ict} = \phi_{acy} + \pi\epsilon_Y + \epsilon_D \tag{8}$$

$$u_{ict} = \omega_{act} + (\gamma\epsilon_Y + \epsilon_O)(\delta Y_{ict} + 1) \tag{9}$$

This error structure implies that unobserved traits can simultaneously influence income, healthcare care demand, and OOP spending, creating endogeneity through shared unobservables, particularly through ϵ_Y. The interaction term $(\gamma\epsilon_Y + \epsilon_O)(\delta Y_{ict} + 1)$ in equation (9) allows income to amplify the effect of preference heterogeneity: for example, richer individuals may value time more and thus exhibit a stronger willingness to supplement public care, even at higher cost. This reflects real-world frictions in universal healthcare systems where non-price rationing (e.g., waiting times) creates incentives for income-based opt-out behavior.

Traditional IV estimators, by contrast, assume homogeneous treatment effects and recover only average impacts of income on OOP. This is a limitation in settings like healthcare, where budget constraints, insurance coverage, and time costs differ widely across individuals. The SQTE approach addresses this by estimating the causal effect of income at different quantiles of the outcome distribution, providing a richer picture of how responsiveness to income varies across individuals. In particular, the model captures the theoretically predicted pattern that income effects on OOP are strongest among individuals with low income but rising income potential, and among those with stronger preferences for quality or convenience.

In summary, our structural approach combines IV identification with a quantile framework to disentangle the role of income, self-selection, and unobserved heterogeneity in shaping healthcare behavior. It provides a flexible and policy-relevant tool for testing key predictions of in-kind redistribution theory—especially those related to self-targeting, opt-out behavior, and the redistributive effectiveness of universal public healthcare systems.

3 Data

The Survey of Health, Ageing and Retirement in Europe (SHARE) is a multidisciplinary and cross-national panel database that collects microdata on individuals aged 50 and over across numerous European countries. The survey is designed to support research on the interplay between health, socio-economic status, social and family networks in ageing populations. It has been developed following studies structured similarly such as the Health and Retirement Study (HRS) in the United States and the English Longitudinal Study of Ageing (ELSA), allowing user to compute also comparative analyses [1].

For the purpose of the project, we leverage survy information, with a particular interest in:

- Health needs, including self-reported health status, diagnosed conditions, limitations in activities of daily living (ADLs and IADLs), chronic diseases, and mental health indicators. Objective measures such as grip strength and walking.
- Healthcare utilization, detailing the frequency of visits to general practitioners, specialists, hospital stays, and prescribed treatments.
- Out-of-pocket (OOP) expenditures, including direct payments made by individuals for healthcare services and medications not fully covered by public or private insurance. These include spending on doctor visits, hospital care, dental care, medications, and medical tests.

Importantly, SHARE follows individuals over multiple waves starting from 2004 to 2022. Indeed, the longitudinal nature of the dataset facilitates the study of dynamics in health deterioration, healthcare access, and expenditure patterns over time, making it particularly well-suited for evaluating the redistributive effects of public healthcare provision and the role of self-targeting mechanisms.

To complement individual-level information from SHARE, we rely on EUROMOD, the European Union's tax-benefit microsimulation model, to compute simulated tax burdens at the individual level [2]. EUROMOD models the detailed structure of national tax systems across European countries in a harmonized framework, allowing us to estimate individuals' tax liabilities based on their reported characteristics. This approach enables us to exploit institutional variation in taxation across countries, while ensuring consistency in the way tax burdens are measured. Importantly, EUROMOD is maintained and regularly updated by national experts in collaboration with the Joint Research Centre (JRC), and its outputs are validated against external data sources, making it a reliable tool for comparative fiscal analysis. In our study, simulated tax burdens from EUROMOD are used to approximate the formal fiscal pressure faced by individuals, independent of behavioral responses, providing a consistent measure of taxation across institutional contexts.

3.1 Sample preparation

As we aim to investigate whether healthcare systems exhibit redistributive features across European countries, our main criterion for sample selection is the ability to calculate healthcare expenditure at the individual level. We specifically focus on out-of-pocket expenditure as reliable proxy of healthcare demand within the private healthcare provision sector. Subsequently, we include those waves that have clean and consistent measures of this variable. Among the nine available SHARE waves, only five met this standard, as others lacked sufficient information on out-of-pocket spending. Furthermore, we restricted the sample to countries with data from at

[1] For further information, see: https://share-eric.eu/
[2] For detailed information: https://euromod-web.jrc.ec.europa.eu/

least three waves, allowing to exploit cross-country and over-time variation in fiscal policy to the maximum extent. Table 1 reports stepwise sample refinement process, leading to an unbalanced panel of 222,496 observations.

Table 1: Sample construction table

	Observations
SHARE complete dataset (waves 1,2,5,6,7,8,9)	401,964
Delete countries with < 3 waves	
Dataset	318,926
Delete waves with missing out-of-pocket expenditure	
Dataset	231,521
Missing observation in key demographic controls	
Dataset	222,496

3.2 Sample descriptives

The final analytical sample comprises 222,496 observations, drawn from five waves of the SHARE dataset and covering 15 European countries.[3] On average, the dataset includes approximately 14,900 individuals per country and around 37,200 individuals per wave. As shown in Figure 1, the sample spans a broad spectrum of healthcare systems and macro-regional contexts: from Northern European countries with predominantly tax-funded, universal models, to Central European systems based on social health insurance, and hybrid frameworks in Southern and Eastern Europe, where private out-of-pocket spending plays a more substantial role. This diversity allows for meaningful cross-country comparisons in the context of healthcare financing and redistribution.

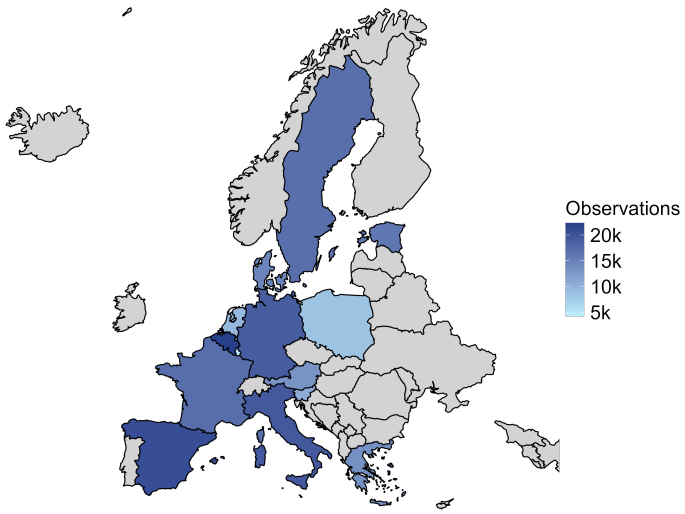

Figure 1: Distribution of observations by country

[3]Austria, Germany, Sweden, Netherlands, Spain, Italy, France, Denmark, Greece, Belgium, Czech Republic, Poland, Luxembourg, Slovenia, Estonia.

Table 2 presents key summary statistics for the variables used in the analysis. The average age of individuals in the sample is approximately 67 years, with a standard deviation of about 10 years. Financial variables, including income and out-of-pocket (OOP) expenditure, have been winsorized at the 99th percentile to mitigate the influence of extreme values. For income, although the SHARE dataset provides both raw and imputed values, we rely on the imputed income variables (De Luca et al., 2015), which undergo additional consistency checks and offer more reliable information. Income is classified into three main streams: income from employment and self-employment; pension income[4]; and other sources of income[5]. Our individual income measure includes all these components except for partner income, in order to maintain a consistent individual-level perspective. Although adjusted for purchasing power parity to enable meaningful cross-country comparisons, the income distribution remains notably right-skewed[6].

The average out-of-pocket (OOP) healthcare expenditure in the sample is approximately 271 euros, but values can rise to as much as 6,000 euros in certain regions, with prescription medications accounting for the largest share of these expenses. On the healthcare demand side, individuals report a relatively high number of doctor visits on average, while the hospitalization rate remains lower, at around 15%. The average number of chronic conditions per individual is moderate; however, over half of the sample (52%) reports experiencing at least one limitation in daily functioning, including activities of daily living (ADL), instrumental activities of daily living (IADL), or mobility.

Table 2: Summary statistics of main variables of interest

	Mean	SD	Min	Max
Age	67.22	10.02	50	106
Female	0.55		0	1
In hospital last 12 months	0.15		0	1
Seen/Talked to medical doctor	6.60	9.01	0	98
Number of chronic diseases	1.78	1.58	0	14
Needs (dummy)	0.52		0	1
Limitations with ADL	0.24	0.86	0	6
Limitations with IADLs	0.45	1.33	0	9
Mobility limitations	1.64	2.34	0	10
Total income	11,715.39	12,109.45	0	119,171
Income from earnings	4,494.64	10,468.81	0	73,111
Pension income	6,382.70	8,748.49	0	67,812
Other sources of income	1,815.04	5,336.08	0	51,514
Total OOP spending	270.98	469.18	0	6,002
OOP - inpatient care	7.87	45.16	0	697
OOP - outpatient care	117.76	300.10	0	3,060
OOP - prescribed drugs	108.14	193.48	0	1,912

Figures 2 and 3 provide a more detailed view of private healthcare spending patterns across Europe. Figure 2 highlights substantial cross-country differences in average out-of-pocket (OOP) expenditure per capita, even after adjusting for purchasing power parity. Countries such as Belgium, Italy, and Greece display considerably higher spending levels, while others like Slovenia, the Netherlands, and the Czech Republic report substantially lower expenditures. This heterogene-

[4]Including all forms of pension: old-age, survivor, disability insurance, early retirement, and social pension

[5]Such as income from rent, stocks, interest, or financial support from a partner

[6]Figure 8 reports the CDF and histogram distribution of income for all countries included in the analysis

ity may reflect differences in national health system design—such as cost-sharing mechanisms, the scope of public coverage, and access to services—as well as population health needs.

To gain an initial sense of how private healthcare spending relates to underlying health needs, we consider the number of chronic conditions across countries. As shown in Figure 3, out-of-pocket (OOP) expenditure does not exhibit a clear or consistent relationship with the burden of chronic conditions. This suggests that while health needs are certainly relevant, they may not fully explain the observed variation in private spending, pointing to the importance of exploring additional factors such as institutional settings or access-related differences. For instance, Slovenia and the Czech Republic display relatively high average levels of chronic conditions but low OOP expenditure, while Belgium reports among the highest OOP spending despite comparatively low morbidity. Italy and Greece occupy an intermediate position, with both elevated health needs and higher-than-average OOP costs. In contrast, countries such as France and the Netherlands exhibit lower levels of OOP expenditure even at moderate or low levels of chronic illness.

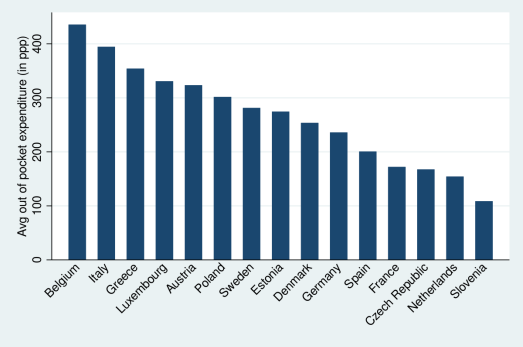

Figure 2: Average out-of-pocket expenditure by country

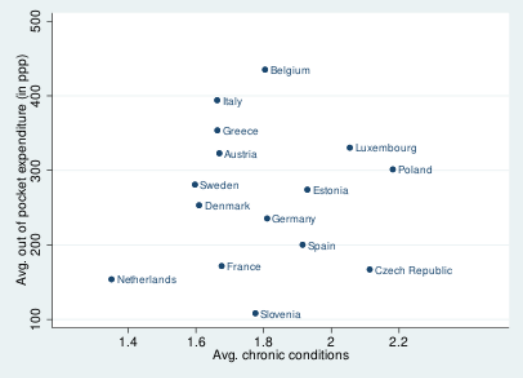

Figure 3: Dispersion graph of mean chronic conditions and out-of-pocket expenditure by country

These descriptive patterns suggest that individuals with comparable healthcare needs may face substantially different financial burdens depending on their country of residence. This variation motivates a deeper investigation into the redistributive role of health systems—specifically, how income conditions exposure to out-of-pocket (OOP) costs when health needs are held constant. To account for institutional and economic heterogeneity across countries, our analysis incorporates several macro-level indicators, including the share of government health expenditure relative to GDP, the proportion of publicly funded health spending within total health expenditure, as well as broader economic trends captured by variables such as the GDP deflator. These dimensions provide a more comprehensive view of the structural factors shaping the interaction between need, income, and private healthcare spending.

3.3 Construction of the simulated tax instrument

The core of our IV approach is the construction of Z_{ict}, the simulated-income instrument derived using the EUROMOD microsimulation model. We follow a three-step procedure to compute an expected tax liability for each individual in the survey data, using synthetic populations and probabilistic matching.

The goal is to obtain the portion of income variation driven solely by policy rules and is summarized with the workflow in figure 4.

The steps are as follows:

- **Synthetic population simulation**: We generate a synthetic population that enumerates all combinations of key individual characteristics and discrete income levels. In practice, we consider combinations of five traits – age group, gender, marital status, education level, and employment status – and stratify these "types" by income deciles (for earnings) and pension deciles. This yields a grid of hypothetical individuals covering the full spectrum of demographics and income levels. We then run this synthetic dataset through EUROMOD, a detailed tax-benefit microsimulation model for European countries. For each synthetic individual type in each country-year, EUROMOD computes the applicable tax liabilities (and transfers) under that country-year's tax and pension rules. This produces a table of simulated tax outcomes $\psi_{ct}(O, y)$, where O denotes a particular combination of observable characteristics (our five traits) and y represents a given income/pension level (e.g. the threshold income for a decile). Essentially, this step yields the policy-determined tax for every cell in the [characteristics] $\times$ [income level] matrix, under each country and year's policy environment.
- **Probability of income decile assignment:** Next, for each actual respondent in our data (from SHARE), we estimate the probability that the individual falls into each income/pension decile category, conditional on their observed characteristics. Rather than using the individual's reported income (which could be endogenous), we predict their income distribution based on their demographic and socio-economic profile. Specifically, we fit a multinomial logit (MNL) model for each country c to estimate $P(D_{it} = j \mid X_i)$ – the probability that individual i at time t belongs to income/pension decile j (where $j = 1, \ldots, 10$) given their traits X_i. Formally, for a given country and year we model:

$$P(D_{it} = j \mid \mathbf{X}_i) = \frac{\exp\left(\mathbf{X}_i'\beta_j\right)}{\sum_{k=1}^{J} \exp\left(\mathbf{X}_i'\beta_k\right)}, \quad \text{for } j = 1, \ldots, J \tag{10}$$

  for $j = 1, \ldots, J$ (with $J = 10$ deciles). In this formulation, D_{it} is a categorical variable indicating the joint earnings–pension decile of individual i, and X_i is the vector of the five individual characteristics (age group, gender, marital status, education, employment) used in the synthetic data. We estimate this MNL separately for each country, allowing the income distribution to differ across countries in a flexible way. The predicted probabilities $p_{ijct} \equiv P(D_{it} = j)$ represent how likely an individual with traits X_i is to fall into each income bracket j in country c at time t. Intuitively, this step says: based on observable traits, what portion of the income distribution would we expect this person to occupy? (For example, a younger, highly-educated, employed individual might have a high probability of being in an upper income decile, whereas an older retired person might have a higher probability of a lower income decile.)
- **Expected tax calculation:** finally, we combine the information from the first two steps to compute each individual's expected tax liability under the prevailing policy. For individual i in country c at time t, the instrument Z_{ict} is defined as the expected tax (or tax-induced income reduction) given their characteristics. Using the simulated tax function from step 1 and the decile probabilities from step 2, we calculate:

$$z_{ict} = E[\text{Tax}_{ict}] = \sum_{j=1}^{J} \psi_{ct}(\mathbf{O}_i, y_j) \cdot P(D_{it} = j \mid \mathbf{X}_i) \tag{11}$$

where $\psi_{ct}(O_i, y_j)$ is the simulated tax liability for a synthetic individual with observable profile O_i (the traits of person i) and income level y_j (the representative income for decile j in country c, time t), and $P(D_{it} = j \mid X_i) = p_{ijct}$ is the probability that person i falls in decile j as estimated above. In other words, we take a weighted average of the simulated tax burdens across all income deciles, using the individual-specific probabilities as weights. The resulting z_{ict} is an imputed tax (a continuous dollar/euro amount, or its log) for each individual, which reflects the policy-driven component of that person's income given their characteristics. If tax policies are more progressive or heavier in country c at time t, z_{ict} will be higher (indicating that, say, a high-income type would face a large tax), whereas in a lower-tax environment z_{ict} will be lower. Crucially, z_{ict} does not incorporate person i's actual realized income—only the exogenous policy structure and the broad position (decile) they are likely to be in. By construction, this simulated tax instrument varies across countries and over time with changes in tax laws and benefit rules, but for a given individual it is fixed by their demographics and the policy environment (not by their idiosyncratic income shocks). This z_{ict} serves as the instrument in Equation (2) above, providing a proxy for disposable income that is purged of endogenous individual variation. Essentially, we are using the predicted fiscal burden dictated by policy as an instrument for actual income. This approach follows the tradition of "simulated instruments" in public economics, ensuring that only external (policy-induced) income variation identifies the effect.

Using this instrument in the first stage (2), we obtain the portion of income variation driven by tax/pension policies, and then in the second stage (3) we estimate how that exogenous income component affects OOP expenditures. The inclusion of extensive controls X_{ict} and fixed effects in both stages helps account for observable heterogeneity and common shocks. Under the assumptions discussed, β_{IV} in Equation (3) can be interpreted as the causal elasticity of OOP spending with respect to income, i.e. how much extra health spending is induced by higher income, holding other factors constant. This IV estimate forms the core empirical finding of our study, as it reveals the redistributive interplay between income and healthcare spending after netting out confounders.

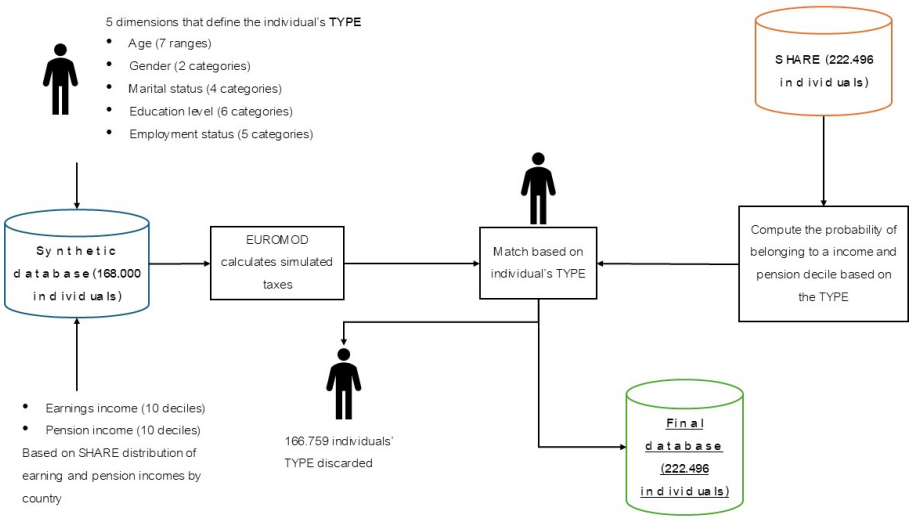

Figure 4: Schematic representation of dataset and instrument construction

Before moving to the results, it should be important to validate the outcome given by the construction of the instrument of simulated taxes. Figure 5 and 6 help us identify cross-country differences in the distribution of the simulated tax values from EUROMOD and also to ensure internal validity after assigning each individual a simulated tax and benefit following equation 11.

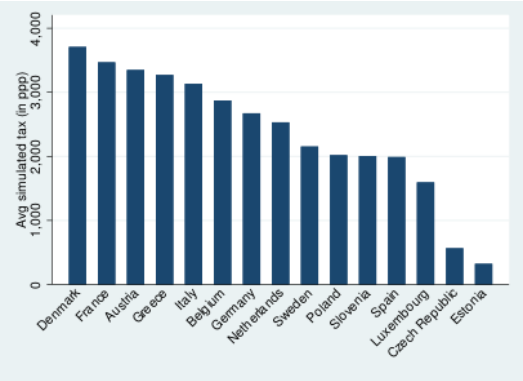

Figure 5: Average simulated tax values by country in ppp

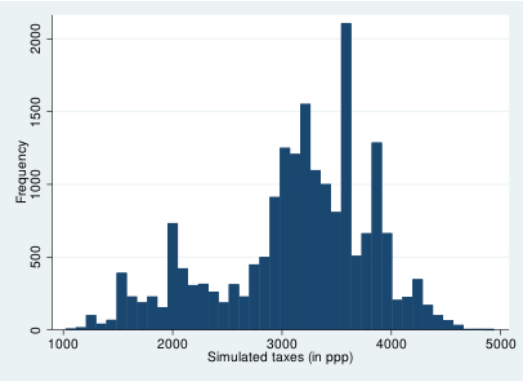

Figure 6: Summary statistics of individual simulated tax in Italy

Average simulated tax values across European countries appear broadly consistent with historical trends reported by the European Commission, even after adjusting for differences in purchasing power parity (PPP)[7]. As expected, Denmark emerges as one of the countries with the highest average tax burdens. Figure 6 presents the distribution of individual-level simulated

[7]See Taxation Trends for a detailed comparison.

tax values for Italy. The distribution is notably non-symmetric, mirroring the underlying income distribution. It exhibits both dispersion and clustering, which likely reflect the structure of the Italian tax-benefit system, including tax brackets, means-tested benefits, and eligibility thresholds.

Taken together, the outputs shown in Figures 5 and 6 align closely with known institutional patterns across European tax systems. This correspondence suggests that EUROMOD is successfully capturing the salient features of national tax policies and can be considered a credible tool for comparative analysis[8].

4 Results

Table 3 reports the results from our instrumental-variable (IV) estimation strategy, which addresses potential endogeneity in the relationship between income and out-of-pocket (OOP) healthcare spending. Column (1) presents the naïve OLS estimates, column (2) shows the first stage of the IV approach, and column (3) reports the second-stage IV estimates. We use a simulated tax burden—constructed using EUROMOD—as an instrument for log income. This strategy isolates exogenous variation in income due to national tax system features, while holding actual earnings and consumption decisions constant.

The instrument performs strongly in the first stage, with an F-statistic of 216.6, well above conventional thresholds for weak instrument concerns. In the second stage, we find that income has a strong and statistically significant effect on OOP spending: a 1% increase in income leads to an estimated 0.26% increase in OOP expenditures. This elasticity is slightly smaller than the naïve OLS estimate (0.287), but remains highly significant, suggesting that endogeneity likely biases OLS estimates only modestly downward.

Crucially, this relationship persists even after controlling for a rich set of covariates, including functional needs, number of chronic diseases, gender, marital status, and employment. As expected, indicators of health need show strong positive associations with OOP spending, confirming the role of medical demand. However, the fact that income remains a significant determinant—even net of health status—highlights the role of financial capacity and preferences for quality, convenience, or access in shaping healthcare consumption within universal systems.

Differences in OOP spending by marital and employment status also point to heterogeneity in financial burden across population subgroups. Compared to married individuals, widowed, divorced, and never-married individuals tend to spend less, potentially reflecting differences in household composition, support structures, or healthcare-seeking behavior. Employment status reveals further variation: while employed individuals spend significantly less than retirees, groups like the unemployed, permanently sick, and homemakers show a positive association with OOP spending. For homemakers in particular—who are a small and low-income subgroup—the large coefficient may partly reflect the high sensitivity of OOP to income constraints rather than a clear behavioral pattern (such as opting out of public provision).

Finally, though not shown in the table, education also appears to play an important role: higher levels of education are positively correlated with OOP spending, even after controlling for income and health. This may suggest a preference for higher-quality or supplementary care among more educated individuals, or greater awareness and access to healthcare options.

In sum, while health needs are clearly a strong predictor of OOP spending, they do not fully account for the observed variation across income and socio-demographic groups. This underlines the importance of considering financial and structural factors—beyond health status alone—when assessing equity in healthcare financing.

[8] In appendix, figure 9 refers to the distribution of simulated taxes for each country in our sample

Table 3: IV estimates of the effect of individual's income on healthcare expenditure, captured by out-of-pocket expenditure

	(1) OLS	(2) First stage	(3) IV
Total income (log)	0.287***		0.261***
	(0.002)		(0.059)
Predicted taxes (log)		0.229***	
		(0.016)	
Needs (dummy)	0.378***	−0.012	0.377***
	(0.010)	(0.012)	(0.010)
Number of chronic diseases	0.266***	0.006	0.266***
	(0.003)	(0.004)	(0.003)
Female	0.366***	−0.285***	0.358***
	(0.009)	(0.011)	(0.020)
Registered partnership	−0.021	0.054	−0.020
	(0.035)	(0.041)	(0.035)
Never married	−0.143***	0.090***	−0.141***
	(0.020)	(0.024)	(0.020)
Divorced	−0.124***	0.048**	−0.123***
	(0.016)	(0.019)	(0.016)
Widowed	−0.229***	0.624***	−0.212***
	(0.014)	(0.017)	(0.040)
Employed or self-employed	−0.133***	0.031*	−0.131***
	(0.015)	(0.019)	(0.016)
Unemployed	0.202***	−1.797***	0.155
	(0.031)	(0.046)	(0.112)
Permanently sick	0.185***	−0.510***	0.168***
	(0.027)	(0.035)	(0.047)
Homemaker	0.640***	−3.033***	0.558***
	(0.020)	(0.028)	(0.186)
Country FE	*Yes*	*Yes*	*Yes*
Interview year FE	*Yes*	*Yes*	*Yes*
R-squared	0.44	0.59	0.18
Obs.	222,496	222,496	222,496

Standard errors in parentheses

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

Marital status variables are estimated with respect to being married; employment status variables are estimated with respect to being retired.

The fact that higher-income individuals spend more in absolute terms on healthcare through private out-of-pocket (OOP) payments indicates that healthcare spending is not proportional across income groups. This finding is consistent with our theoretical framework centered on in-kind redistribution and opt-out behavior. In systems with universal public healthcare provision, individuals with greater financial resources may selectively "opt out" of public provision by purchasing supplementary or higher-quality care in the private sector. This self-selection mechanism implies that the public service—while available to all—is used more intensively by lower-income individuals, whereas wealthier individuals contribute to the system's financing but rely relatively less on its services. In this way, in-kind transfers can produce redistribution not by equalizing consumption, but by enabling differential use tied to financial capacity. Our empirical finding of a positive income elasticity of OOP spending provides evidence consistent with this mechanism: income is a strong predictor of private healthcare spending, even after controlling for health status and demographic characteristics.

To further assess whether this income gradient reflects discretionary spending conditional on health need, we stratify the sample by chronic health burden in Table 4. This allows us to examine whether income plays a more pronounced role when individuals face ongoing or complex health needs. Among individuals with fewer chronic conditions, income is not significantly associated with OOP expenditure, suggesting limited room—or perhaps limited motivation—for opting out of public services when healthcare needs are minimal. Conversely, among individuals with multiple chronic conditions, a clear and significant positive income–OOP gradient emerges. This pattern supports the view that in the presence of higher health needs, wealthier individuals are more likely to supplement or replace public care with privately financed services, potentially to gain faster access, higher quality, or greater flexibility. The combination of universal coverage and scope for private supplementation thus creates a system where redistribution is achieved not by limiting the rich's access, but by allowing the poor to rely more heavily on the publicly funded core.

Table 4: Heterogeneous income effects on out-of-pocket spending by chronic condition intensity

	(1)	(2)
	≤ 2 chronic cond.	> 2 chronic cond.
Total income (log)	0.068	0.408***
	(0.095)	(0.070)
Needs (dummy)	0.291***	0.258***
	(0.013)	(0.020)
List of controls	*Yes*	*Yes*
Country FE	*Yes*	*Yes*
Interview year FE	*Yes*	*Yes*
R-squared	0.07	0.21
Obs.	112,926	60,189

Standard errors in parentheses

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

To complement this perspective, table 5 focuses on healthcare demand behavior, rather than health status per-se. Here, we distinguish between individuals with positive healthcare demand (i.e., those who seek care) and those with no recorded utilization. By doing so, we explicitly condition on the extensive margin of healthcare use and assess how income affects OOP spending among users. The results show that the income gradient in OOP becomes stronger when we restrict the sample to individuals who are already utilizing healthcare services. This suggests

that income-related differences are not solely about whether individuals seek care in response to health needs (the extensive margin), but also about how much they spend once they enter the healthcare system (the intensive margin). Among users, higher-income individuals appear more likely to supplement public provision with privately financed services. This pattern is consistent with our theoretical model of opt-out behavior, where wealthier individuals, conditional on health needs, allocate more resources to private care options.

Table 5: Instrumental variable estimates on the extensive margin of healthcare demand

	(1)
	Positive healthcare demand
Total income (log)	0.264***
	(0.057)
Needs (dummy)	0.328***
	(0.010)
Number of chronic diseases	0.208***
	(0.003)
List of controls	*Yes*
Country FE	*Yes*
Interview year FE	*Yes*
R-squared	0.17
Obs.	198,210

Standard errors in parentheses

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

Taken together, the results from tables 4 and 5 suggest a pattern consistent with our theoretical expectations: income appears to influence healthcare consumption through private supplementation (i.e. out of pocket expenditure), especially among individuals with greater health needs or those actively engaging with the system. This interplay between health status and demand behavior seems to point toward a form of self-selection that may underpin how in-kind redistribution functions in practice, allowing for some degree of differentiation in care while preserving broad access.

4.1 Structural Quantile Treatment Effects

Building on the results reported in Table 5, which estimate the effect of income and health needs on the probability of any healthcare use (the extensive margin), we now turn to the second stage of the model: out-of-pocket (OOP) spending conditional on positive demand. The SQTE framework applies precisely to this subpopulation—those who have accessed the healthcare system—and allows us to explore how the causal impact of income on OOP varies across the distribution of expenditures, rather than focusing solely on average effects.

While the IV approach yields the average causal effect of income on OOP spending, policymakers may also be interested in how this effect varies across the distribution of OOP expenditures. For instance, does additional income primarily increase healthcare spending for those already spending a lot out-of-pocket (perhaps due to severe health needs), or does it affect lower spenders differently? To explore such heterogeneous effects, we extend our analysis using a Structural Quantile Treatment Effects (SQTE) framework. This is essentially an instrumental variable

quantile regression approach that allows the treatment effect of income to differ at different quantiles of the OOP outcome distribution. Implementing a structural quantile model enables us to explicitly account for unobserved heterogeneity in both the income-generating process and the OOP spending determination. In practical terms, the SQTE model introduces additional latent structures (e.g. through a two-stage quantile regression technique) to handle endogeneity at each quantile, modeling two types of unobservables: (i) factors that affect income (such as innate ability or social background) and (ii) factors that affect the propensity to spend on healthcare (such as risk preferences or health severity). By doing so, we can estimate the quantile-specific impact of income on OOP – for example, the effect of income on the 90th percentile of OOP spending versus the 10th percentile.

Methodologically, the structural quantile approach leverages the same instrument Z_{ict} for identification at each quantile, ensuring that the estimated heterogeneous effects are causally interpreted. We effectively obtain a series of β_τ estimates, where β_τ is the income effect at the τ-th quantile of the OOP distribution (for $\tau \in (0, 1)$). This provides a more nuanced picture of redistribution through public healthcare: it tells us whether increases in income disproportionately benefit those with high out-of-pocket burdens or those with lower burdens. For example, a larger β_τ at higher τ would suggest that extra income primarily enables those with very high medical expenses to spend even more (perhaps upgrading care or covering expensive treatments), whereas a uniform β_τ across τ would indicate a proportional shift across the board. Our SQTE results indeed allow us to probe whether the income gradient in OOP is steeper for certain subgroups or expenditure levels, helping to distinguish between behavioral differences and structural inequalities. In summary, this extension moves beyond average effects and incorporates distributional considerations, offering a comprehensive understanding of how public healthcare financing interacts with income. By combining the baseline OLS, the IV causal analysis, and the structural quantile treatment effects, our methodological approach provides robust evidence on the redistributive role of public healthcare systems in Europe, ensuring a comprehensive and credible analytical foundation.

Figure 7 presents the estimated Structural Quantile Treatment Effects (SQTE) of income on out-of-pocket (OOP) health expenditure, evaluated across quantiles of the unobserved distribution of potential OOP spending. The SQTE framework allows us to investigate how the causal effect of income varies across individuals who, due to latent characteristics—such as unobserved preferences for healthcare, health status, or income-related constraints—differ in their position within the potential outcome distribution.

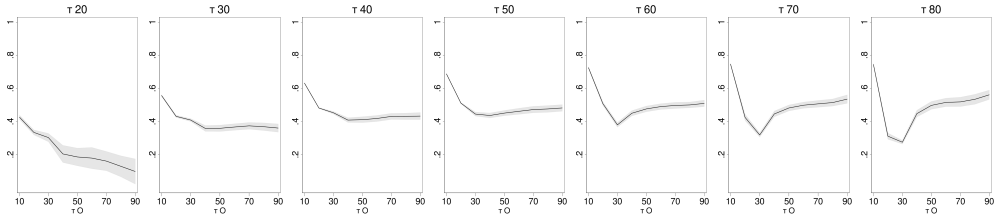

Figure 7: Structural quantile treatment effects (SQTE) across the distribution of out-of-pocket expenditure

In this setting, τ indexes heterogeneity in the potential OOP expenditure, and the estimated SQTEs capture the marginal effect of income conditional on this latent heterogeneity. Hence, we can interpret the results as reflecting how income affects OOP spending across individuals with different unobserved propensities to spend on healthcare.

At lower quantiles ($\tau = 20$) of the OOP potential outcome distribution, the SQTE is negative: increases in income are associated with a decline in OOP expenditure. This suggests that individuals with low latent OOP expenditure—potentially those who are either healthier, more reliant on public healthcare, or more financially constrained—use additional income to meet other needs or reduce reliance on private spending. It may also reflect a threshold effect, where a minimum level of necessary expenditure is required, beyond which preferences shift away from additional healthcare spending. At intermediate quantiles ($\tau = 30$–50), the SQTE is close to zero. In this segment, individuals exhibit an average latent inclination toward OOP spending. Here, income does not seem to significantly alter OOP spending, possibly because these individuals are operating within tight trade-offs: their healthcare utilization may be primarily driven by need (e.g., chronic conditions), with limited room for discretionary adjustment. This plateau in the SQTE may reflect a corner solution, where preferences and constraints offset each other.

At higher quantiles ($\tau = 60$–80), the estimated effect becomes increasingly positive. For individuals with a high latent propensity to spend on healthcare— with stronger preferences for private care—income increases lead to higher OOP expenditure. These individuals likely exhibit greater willingness and ability to opt out of public provision, responding elastically to income by increasing consumption of out of pocket expenditure. This non-linear pattern—negative effects at the bottom, flat in the middle, and positive at the top—suggests that the causal effect of income on OOP expenditure is highly heterogeneous.

The SQTE framework reveals distinct mechanisms across the distribution: budget constraints and reallocation effects at the bottom; medical need-driven stability in the center; and preference-driven, elastic responses at the top. These findings underscore the importance of accounting for latent heterogeneity in both preferences and constraints when analyzing the income-healthcare expenditure relationship.

5 Conclusion & Policy Implications

This study provides robust empirical evidence on the redistributive role of public healthcare systems in Europe, with a focus on the interplay between income and out-of-pocket (OOP) healthcare expenditures. By integrating baseline OLS estimates, instrumental variable (IV) methods, and structural quantile treatment effects (SQTE), our analysis goes beyond average effects and reveals important heterogeneity in how income influences healthcare spending behavior across the population.

Our results confirm that income has a strong and statistically significant impact on OOP healthcare spending, even in the context of universal public coverage. Higher-income individuals spend more privately, consistent with theoretical predictions of voluntary opt-out behavior in systems based on in-kind transfers. This supports the view that public healthcare provision functions as a form of self-targeting redistribution: wealthier individuals supplement or replace public services with private alternatives, thereby implicitly concentrating public resources on those with fewer means.

The SQTE analysis adds a crucial dimension by demonstrating that the causal effect of income on OOP spending is not uniform across the distribution. At lower quantiles, increases in income are associated with a reduction in OOP expenditure, suggesting that individuals with lower latent demand for private care—due to financial constraints, greater reliance on public services, or lower perceived need—may allocate additional income to other priorities or reduce financial stress. In the middle of the distribution, income has little or no effect, indicating that these individuals may already be operating at their preferred consumption level, where healthcare decisions are largely driven by medical need rather than financial capacity. At higher quantiles, however, income increases lead to significantly greater OOP spending, particularly among individuals with stronger preferences for quality, convenience, or faster access to care. These heterogeneous effects highlight the importance of unobserved traits—such as health preferences, risk aversion, or human capital—in shaping both income potential and healthcare behavior.

From a policy perspective, these findings carry several implications. First, while private supplementation can relieve pressure on public resources and accommodate individual preferences, there is a clear need to safeguard against inequities in access. The existence of a steep income gradient in OOP spending—even among those with similar health needs—suggests that ability to pay can influence the quality or timeliness of care received. This raises the risk that low-income individuals may experience unmet health needs or delay treatment due to cost, despite the presence of universal coverage. Second, policymakers should consider reinforcing financial protection by revisiting the design of cost-sharing arrangements. Services associated with high OOP burdens—such as outpatient care and prescription drugs, which account for the bulk of spending in the SHARE data—may require greater public coverage or income-adjusted co-payments to ensure affordability for the most vulnerable. This is especially relevant given the finding that, among individuals with multiple chronic conditions, income plays a larger role in determining OOP spending. Improving coverage for these high-need groups would enhance equity while maintaining the redistributive logic of universal systems. Third, systematic monitoring of equity indicators—such as the incidence of catastrophic health spending and reported unmet medical needs—should complement analyses of aggregate OOP levels. The results suggest that averages may mask important disparities across the income distribution and among subgroups with different health burdens or demand patterns. Incorporating distributional metrics into routine policy evaluation could help identify emerging inequities before they deepen. Finally, the evidence supports the need for nuanced policy tools that respond not only to observable income but also to unobserved drivers of healthcare behavior, such as health preferences, knowledge, and opportunity costs. While income remains a central axis of redistribution, the SQTE results highlight that

effective targeting and protection may also require attention to preference-sensitive dimensions of care, such as waiting times, access to providers, and service quality.

In conclusion, public healthcare systems in Europe continue to fulfill a core redistributive function, but the coexistence of private supplementation introduces both efficiency gains and equity risks. The structural and distributional analysis presented here clarifies where these tensions are most pronounced and provides actionable evidence to guide policy reforms. Ensuring that no individual is deterred from seeking necessary care due to financial barriers remains essential to the integrity and sustainability of universal health systems.

6 Appendix

The figure displays the empirical cumulative distribution function (ECDF, red line) and the kernel-smoothed histogram (blue bars) of the log of individual income across 15 countries.

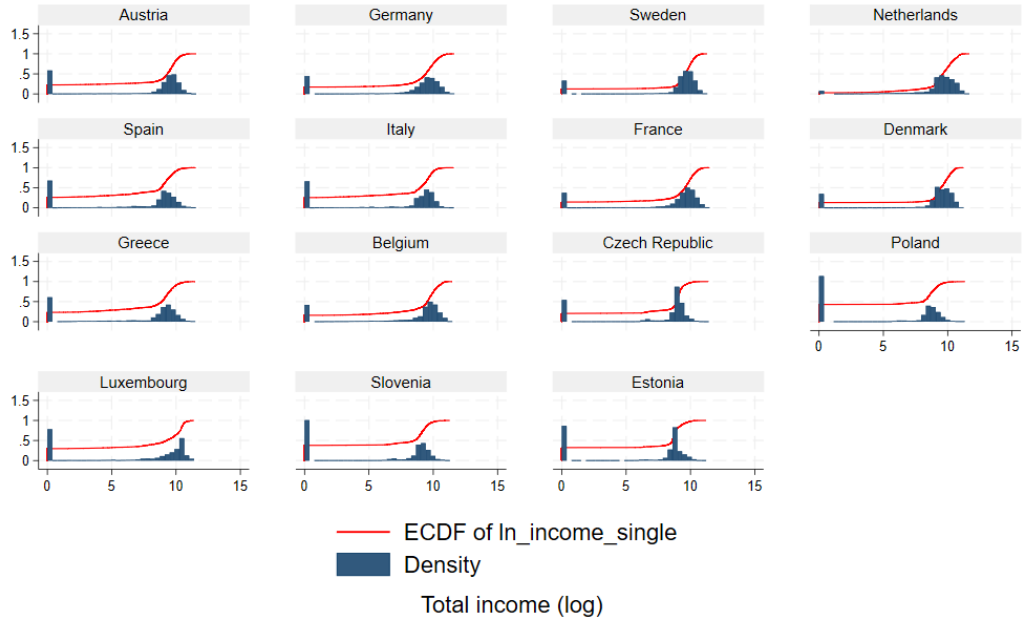

Figure 8: Distribution and cumulative distribution function (CDF) of individual total income by country

Figure 9 shows the distribution of simulated individual tax burdens across countries in our sample. Tax values are computed using EUROMOD and reflect cross-country differences in tax system design and income composition.

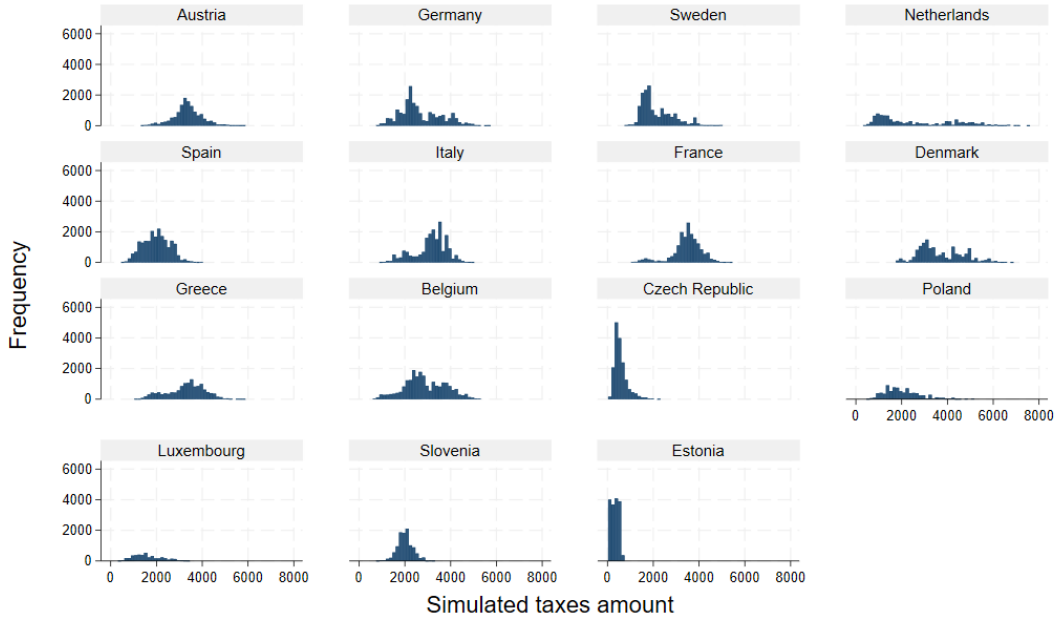

Figure 9: Distribution of simulated tax values by country

References

Angrist, J. D. and Krueger, A. B. (2001). Instrumental variables and the search for identification: From supply and demand to natural experiments. *Journal of Economic Perspectives*, 15(4):69–85.

Besley, T. and Coate, S. (1991). Public provision of private goods and the redistribution of income. *The American Economic Review*, 81(4):979–984.

Currie, J. and Gahvari, F. (2008). Transfers in cash and in-kind: Theory meets the data. *Journal of economic literature*, 46(2):333–383.

De Luca, G., Celidoni, M., and Trevisan, E. (2015). Item nonresponse and imputation strategies in share wave 5. In Malter, F. and Börsch-Supan, A., editors, *SHARE Wave 5: Innovations & Methodology*, pages 85–100. SHARE-ERIC / Max Planck Institute for Social Law and Social Policy, Munich.

Epple, D. and Romano, R. E. (1996). Ends against the middle: Determining public service provision when there are private alternatives. *Journal of public economics*, 62(3):297–325.

Honekamp, I. and Possenriede, D. (2008). Redistributive effects in public health care financing. *The European Journal of Health Economics*, 9(4):405–416.

Ma, L. and Koenker, R. (2006). Quantile regression methods for recursive structural equation models. *Journal of Econometrics*, 134(2):471–506.

Melli, G., Azzollini, L., and and, F. B. B. (2025). Healthcare provision and attitudes towards redistribution. a regional analysis across europe. *Journal of European Public Policy*, 0(0):1–29.

van Doorslaer, E., Wagstaff, A., van der Burg, H., Christiansen, T., Citoni, G., Di Biase, R., Gerdtham, U.-G., Gerfin, M., Gross, L., Häkinnen, U., John, J., Johnson, P., Klavus, J., Lachaud, C., Lauritsen, J., Leu, R., Nolan, B., Pereira, J., Propper, C., Puffer, F., Rochaix, L., Schellhorn, M., Sundberg, G., and Winkelhake, O. (1999). The redistributive effect of health care finance in twelve oecd countries. *Journal of Health Economics*, 18(3):291–313.