The Social Determinants of Choice Quality: Evidence from Health Insurance in the Netherlands*

Benjamin Handel, UC Berkeley
Jonathan Kolstad, UC Berkeley
Thomas Minten, London School of Economics
Johannes Spinnewijn, London School of Economics

August 28, 2020

Abstract

Market provision of impure public goods such as insurance, retirement savings and education is common and growing as policy makers seek to offer more choice and gain efficiencies. This approach induces an important trade-off between improved surplus from matching individuals to products and misallocation due to well documented choice errors in these markets. We study this trade-off in the health insurance market in the Netherlands, with a specific focus on misallocation and inequality. We characterize choice quality as a function of predicted health risk and leverage rich administrative data to study how it depends on individual human capital, socioeconomic status and social and information networks. We find that choice quality is low on average, with many people foregoing options that deliver substantive value. We also find a strong choice quality gradient with respect to key socioeconomic variables. Individuals with higher education levels and more analytic degrees or professions make markedly better decisions. Social influence on choices further increases inequality in decision making. Using panel variation in exposure to peers we find strong within firm, location and family impacts on choice quality. Finally, we use our estimates to model the consumer surplus effects of different counterfactual scenarios. While smart default policies could improve welfare substantially, including the choice of a high-deductible option delivers little welfare gain, especially for low-income individuals who make lower quality choices and are in worse health.

*We thank the Central Bureau of Statistics of the Netherlands and especially Annemieke Redeman for help with the data. Chloé de Meulenaer, Miguel Fajardo Steinhäuser and William Parker provided excellent research assistance. We thank seminar participants at Amazon, Edinburgh University, IO^2 Stanford seminar series, Erasmus University Rotterdam, KULeuven, LSE, Statistics Netherlands, the KVS, the NBER Health Care and the Paris-London Public Economics meetings for excellent comments. We gratefully acknowledge funding by ERC (grant #716485), ESRC and STICERD.

I Introduction

Consumer choice is a central aspect of market function. Increasingly policy makers rely on market solutions that provide choice in the provision of products viewed as public goods, such as retirement investments (see, e.g., Hastings et al. (2013) and Chetty et al. (2014)), schooling (see, e.g., Neilsen (2017)), electricity (see, e.g., Ito (2015)), and health insurance (see, e.g., Enthoven, Garber and Singer (2001)). One important rationale for facilitating choice in such markets — rather than a uniform product, whether offered directly by the government or a regulated private firm — is the opportunity to match heterogeneous consumers with products that provide them with greater surplus. Whether consumers are matched with the best products for them, however, hinges on their ability to effectively choose among offerings.

In practice, if consumers make choice errors, as much prior work documents, the welfare gains from greater choice and competition are diminished, or even eliminated. Furthermore, if choice errors vary in the population we are concerned not only with the average consumer-product match but the distribution of choice quality and surplus. If consumers with lower socioeconomic status are less able to make complex decisions or have less opportunity to engage with those decisions, choice-based policies may increase inequality (e.g., Campbell (2016) and Mullainathan and Shafir (2013)).

In this paper we investigate consumer choice quality and welfare in the context of health insurance provision in the Netherlands. We use rich administrative data on the universe of the population of the Netherlands (approximately 12 million people) including detailed data on demographics, health status, education, employment and social and family networks. These data are linked to individual health insurance choices in the Dutch market in which private insurers offer products under a set of regulatory constraints on product attributes. The policy approach is similar in spirit to Affordable Care Act in the United States and many other managed competition approaches implemented or discussed in other countries. Beyond the specifics of health insurance, the choice environment shares many features with market based solutions for impure public goods more generally (e.g. retirement savings). Products are offered within limits set by a regulator or market designer. Consumers are expected to make choices over a variety of dimensions of which financial outcomes are a key aspect.

The Dutch setting is particularly well suited to studying choice quality because it is a relatively simple choice environment in which we can identify "good" and "bad" choices. The dimension of choice we focus on is the choice of deductible — the amount in each year a consumer must pay out-of-pocket before insurance payments kick in. All insurance contracts (i.e. for every plan design for every brand) have a baseline deductible (375 EUR in 2015). Consumers can also choose to switch to five deductible options, each with a higher deductibles in 100 EUR increments up to an additional 500 EUR (875 maximum total deductible in 2015). When consumers elect a higher incremental deductible, they get a premium rebate of about half of the incremental deductible amount.[1] Because all insurance brands offer the full range of deductibles for all products offered we can abstract away from models of brand or other plan characteristics that typically enter utility in choosing insurance products. Instead, we focus on a relatively simple model in which enrollees choose a deductible level based only on health risk and risk preference in the choice of deductible level.

To assess choice quality, we use tools from machine learning to predict health risk as a function of prior utilization and a rich set of observables. We demonstrate that the model performs well in predicting health

[1] The policy to offer a high-deductible option is a point of frequent debate in the Netherlands: proponents argue that it allows for improved matching of consumers to deductibles and makes consumers more cognizant of health costs while opponents argue that it causes unraveling in pooling based on health risk, hurting sick consumers. In a letter addressed to the Parliament in 2016 (see Schippers (2016)), the Minister of Health for the Netherlands wrote "The voluntary deductible plays an important role in our health care system. It offers choice to consumers. This I deem necessary for the support by healthy consumers for our solidarity health care system. Furthermore, it raises awareness of health care costs, which has a positive effect on the collective health care expenditures." Her successor has been using the same arguments recently in his exchanges with the Parliament.

spending and, in particular, whether someone is expected to have expenses exceeding the baseline deductible. Using these risk projections together with our choice model we show that, for values of risk aversion spanning the health insurance choice literature, we can classify individuals into those who should clearly be opting for a high deductible and those who should clearly be opting for a low deductible.

We demonstrate that approximately 60% of consumers would be better off choosing a higher deductible based on health risk. In contrast, only about 10% actually elect an incremental deductible above the 375 EUR baseline in practice. Even among those for whom we predict health spending to be almost certainly below 375 EUR, take-up of the deductible is only about 15%; a stark contrast with the 100% predicted by the standard model. This large gap between predicted and observed choices cannot be rationalized by risk preference estimates, the only free parameter in the standard model of insurance demand and deductible choice.[2]

What factors contribute to this extremely low take-up of financially beneficial deductible options? To better understand overall choices as well as sources of heterogeneity, we estimates a series of models of choice incorporating detailed observables characteristics on human capital, financial status and peer effects. Specifically, we model deductible choice — take-up of the 500 EUR deductible — as a function of health status, observable characteristics and the interaction of the two.

We find substantial heterogeneity in deductible choice as a function of educational background. For the level of education, we find that consumers with advanced degrees above college are most likely to elect the high deductible (16%), followed by consumers with college degree (13%), consumers with only a high school degree (8%) and high-school dropouts (4%). We also observe approximately 80 distinct education fields: the top three fields, in terms of proportion choosing the 500 EUR deductible, are statistics (29%), mathematics (21%), physics (21%), while the bottom are hair and beauty services (4%), computer use (4%), and child care and youth services (4%). Similarly, those in jobs in more analytic fields are more likely to choose the high deductible.

Differences by educational background and job type are almost entirely explained by the interaction with individuals' predicted health status. Individuals with higher education levels (college degree and above) *who are predictably healthy* are much more likely to elect the high deductible than those with high school education or lower (difference in take-up rate of up to 18%). Election of the 500 EUR deductible is similar and close to zero for consumers *who are predictably sick* in each group.

In contrast with the role played by human capital factors, we find a modest role for income and financial capital once we control for education. Descriptively, as household income increases, the proportion of people who elect the high deductible increases, from 4% for the bottom income quartile to 10% at the top income quartile. However, once we control for education and other socio-demographic factors, the take-up rate hardly increases with income. Similarly, those with higher liquid savings (more than 2000 EUR), lower debt (mortgage and overall) and higher net worth are statistically more likely to choose the 500 EUR deductible option when healthy but the magnitude of these effects are economically small. Taken together these results suggest that income and financial assets have only a small direct effect on choice quality.

One mechanism affecting decisions of individuals and, in turn, determining how different populations make choices, are social or information networks. The role of peer effects in insurance choice have been studied in specific settings (e.g., University of California employees in Sorenson (2006)). In a broader population peer effects may

[2] A number of other models could explain the observed choices including information frictions (see, e.g., Handel and Kolstad (2015*b*) and Bhargava, Loewenstein and Sydnor (2017)), limited attention and salience (see, e.g., Bordalo, Gennaioli and Shleifer (2012), first-order risk aversion (see, e.g., Sydnor (2010)) or liquidity constraints (see, e.g., Ericson and Sydnor (2018) and Finkelstein, Hendren and Luttmer (forthcoming)). While we are unable to differentiate between all of these mechanisms we can shed some light. We directly observe liquid assets allowing us to rule out liquidity constraints as a primary factor. Given that the baseline deductible is typically the default option for consumers, who have to actively elect an incremental deductible, default effects are likely to play an important role (see, e.g., Madrian and Shea (2001), Handel (2013), and Chetty et al. (2014)). We return to these mechanisms but note that our main analysis is agnostic to the specific drivers of choice errors.

affect the average choice quality but could also be an important contributor to inequality. For example, if local peers impact choices, we might expect heterogeneity by geography in choice quality (e.g., urban versus rural). If work peers have large affects on choices, this could further exacerbate the differences in choice quality by job type and firm. Finally, effects within families would be suggestive of important intergenerational transfer of choice capital, either good or bad.

To study peer effects, we leverage the detail of the data that allows us to identify workplace colleagues, neighbors, and family members and their choices. As has been well documented, estimating peer effects poses important empirical challenges, such as, e.g., separately identifying peer effects from correlated unobservable heterogeneity in a peer group (see, e.g., Manski (1999)). We address these challenges using a switcher-design similar in spirit to that described in Abowd, Kramarz and Margolis (1999) to identify the effects of workplace and geographic peers. We estimate a first-stage panel regression with individual fixed effects and firm (location) fixed effects, controlling for predictable differences in health. This framework leverages switchers moving across firms (location) to identify firm (location) fixed effects on deductible choices. In a second-stage, we project these fixed effects onto take-up within the firm (or location), to estimate the extent to which the fixed effects explain differences in take-up.

Our results show that within firm peers have a substantial impact on individual decisions. A 10% increase in the number of co-workers taking a high-deductible causes a 1.4% increases in high-deductible take-up for people switching into the firm. We find similar results for neighborhood peer effects. Peer effects are strongest for individuals who are predicted to have low health costs. For example, a 10% increase in the number of people taking up a high-deductible causes a 1.5 % increase in high-deductible take-up for healthy people switching into the postcode. Conversely, an individual who is predictably sick is .2 % more likely to take up a high-deductible. Thus, peer effects are strong and positive, but only when an individual is predictably healthy and *should* take up a higher deductible and not when they are predictably sick and should not take up the higher deductible.

Studying the impact of family requires a different design. We leverage an event-study design and find that when parents switch their choices, children under 30 living apart from the parent have a 25 percent chance of following their parents and switching. Children over 30 follow their parents' switches, but to a lesser degree, only increasing incremental deductible take-up by 10 percentage points after a parent switch to that deductible. Interestingly, this effect is driven by the take-up response by children who are predictably healthy, but does not differ with their parents' health. This suggests that the primary driver is learning about the parents' decision and considering it in the context of their own health but does not depend on whether the parent made an effective choice.

We next combine these disparate factors affecting choices into a measure of choice quality and study overall inequality in choices and outcomes. To do so, we use our regression estimates to predict consumer choices as a function of health status. Then, using our model of consumer surplus, we rank consumers in terms of choice quality, conditional on health.[3] The top 5% of decision makers choose the surplus maximizing deductible only 55% of the time, while the remaining 95% of the population make choices that are worse than choosing at random. The choice quality is low on average, but also subject to substantial heterogeneity. We assess the underlying correlates of consumers being good (top 5%) or very poor (bottom 5%) decision-makers.

The 5% best decision-makers have an average gross income of 105,000 EUR and net worth of about 250,000 EUR, relative to about an income of 40,000 EUR and net worth of 5,000 EUR for the 5% worst decision makers.

[3]We note that the model does not reflect the potential consumer surplus impacts of additional cost-sharing on health care utilization (see, e.g., Brot-Goldberg et al. (2017)), though we discuss this aspect of consumer surplus. A recent working paper by Remmerswaal, Boone and Douven (2019) shows evidence of limited moral hazard with respect to the deductible policy we investigate in the Netherlands.

Those with college education are 3.48 times more likely to be in the group of 5% best decision makers, relative to their proportion in the general population, while high school dropouts are only 0.3 times as likely to be in this group. In terms of education field, statisticians (19.66 times), philosophers (13.14 times), and economists (6.59 times) are much more likely to be present in the group of top decision-makers while those trained in hair and beauty services (0.64 times) and protection of persons (0.38 times) are much less likely to best present. In terms of profession, those in business services (2.77 times) and insurance (2.13 times) are more likely to be in the group of best decision-makers while those in cleaning (0.26 times) and those retired (.07) or on social insurance benefits (0.32 times) are less likely to be in this group. The group with the worst 5% of decision-makers is largely the opposite of the top 5% group, but is not an exact inverse.

Our results paint a detailed picture of the role of socio-demographic factors and peer effects in deductible choice. Taken together, the results show that a range of factors that are outside of the standard model of insurance choice not only are present, but have large impacts on choices. Moreover, choice quality is strongly correlated with level of education and job and displays important geographic, intra-firm and family peer effects. Not only are such effects outside of the standard model of decision making, they demonstrate that choice quality has an important role to play in inequality, particularly when policies provide impure public goods in environments that depend on choice.

A variety of policy options might be employed to address choice quality, and inequality in choices, in the context of the Dutch health insurance market. To shed light on those we also use our model estimates to study several counterfactual policies. First, we consider the consumer surplus gains from an optimal allocation of consumers to deductibles. This scenario offers useful first-best benchmark, but also relates to a plausible policy intervention if regulators were to use a smart default approach. That is, our counterfactual captures what would happen if our model were used to default people into plans and they took that advice. The other two counterfactual policies reflect policies in which the choice set is limited to either only the higher deductible option (875 EUR) or only the lower deductible option (375 EUR), essentially eliminating choice.

The average benefit from a smart default policy is an improvement in welfare per enrollee of between 58 and 69 EUR, where the lowest value in the range reflects a high assumed CARA coefficient of 10^{-3} and the largest value reflects risk neutrality. These are small in absolute terms, but high relative to the average money at stake of about 145 EUR. Eliminating choice reduce average welfare, but the impact is smaller. The offered option to take a high deductible increases consumer welfare only by 7 EUR to 8 EUR per person. Only offering the 875 EUR deductible would decrease consumer welfare by 26 to 45 EUR.

These results, like much of the prior literature, ignore the role of inequality in outcomes. To incorporate this, we weight outcomes as a function of income using parameters borrowed from the inequality literature (see, e.g., Atkinson (1970)). The value of of offering the option to take a high deductible further decreases when using income-dependent welfare weights, since individuals with lower income take worse decisions and have worse health on average. This negative correlation between income and health also reduces the appeal of mandating all individuals in the high-deductible option: for our high inequality aversion scenario the social surplus loss from only offering the high-deductible is between 134 and 149 EUR, much larger than the analysis that does not factor in inequality aversion.

Overall, the counterfactual analysis shows that the existence and magnitude of choice frictions dramatically reduce the value of offering the high-deductible option, especially for the low-income individuals, who are both less healthy and make worse choices. Instead, consumers would be much better off if their choices were better directed (e.g., through smart defaults a la Handel and Kolstad (2015*a*) or Abaluck and Adams (2019)). Even with such a directed policy, the high-income consumers have the most to gain, despite their higher quality decisions,

because they are healthier on average.

The analysis to this point, and the focus of the paper, has abstracted away from the underlying micro-foundations of choice errors. The simplicity of the choice setting allows us to gain some insights. We develop a series of simulations of alternative models of decision making in the literature and compare them to the observed relationships between health status and deductible choice. We find that, though there are several models one could write down that could come close to rationalizing the data, a model with high switching costs combined with imperfect information about health risk fits the data very well. In contrast, the predicted take-up patterns under random errors or rational inattention do worse jobs of matching our data and neither risk or loss aversion reduce the predicted take-up rates sufficiently to match the observed patterns. High switching costs would further decrease the welfare gains from offering deductible choice, but without imposing strong structural assumptions we cannot separate switching costs from other explanations for the strong default effect (e.g., limited attention).

This paper relates to several distinct literatures. It relates closely to prior work on insurance choice including papers without choice frictions (e.g., Cohen and Einav (2007), Cardon and Hendel (2001)) and many with choice frictions (e.g., Handel and Kolstad (2015*b*), Abaluck and Gruber (2016*a*), Abaluck and Gruber (2016*b*), Fang, Keane and Silverman (2008), Ketcham et al. (2012), Bhargava, Loewenstein and Sydnor (2017), Barseghyan et al. (2013)).[4] Recent work in the health insurance literature has looked at the interaction between choice frictions and adverse selection (e.g., Handel (2013), Spinnewijn (2017), Handel, Kolstad and Spinnewijn (2019) and Polyakova (2016)) and moral hazard (e.g., Brot-Goldberg et al. (2017), Baicker, Mullainathan and Schwartzstein (2015), Einav, Finkelstein and Schrimpf (2015)). It also relates to work by Sorenson (2006), who studies peer effects in the insurance choices of employees at a large firm.[5]

Our analysis also relates to papers that study choice quality and the incidence of consumer frictions in other domains (e.g., Allcott, Lockwood and Taubinsky (2019)). Most notably, there a number of papers that study choice quality and default effects in retirement savings. For example, Chetty et al. (2014) study retirement savings in Denmark using granular nationwide data and show that default effects are much more powerful than subsidies in how they impact consumers' savings portfolios. One key difference between our study and theirs is that we study a context where we can, to a large extent, determine whether consumers are allocated to a poor option for them. While Chetty et al. (2014) show that defaults and subsidies induce very different behavior changes, ultimately it is difficult to tell what savings levels consumers should have in their environment. In addition, we study the effects of peers on choices, which they do not. Cronqvist and Thaler (2004) discuss the privatization of social security in Sweden with a focus on how subtle design factors, such as default effects, can have important implications for consumer choices. Beshears et al. (2016) show that, in a retirement setting, consumers are more likely to switch away from the default option if they will benefit more from doing soon, which is similar to what we find in our environment.

Lastly, our paper relates to two papers on the voluntary deductible in the Netherlands. Van Winssen, Van Kleef and Van de Ven (2015) show that the overall voluntary deductible take-up in the Netherlands is low and a large share of individuals would have gained by taking a higher deductible. They use data from a single insurer to find that a voluntary deductible is most profitable for mostly young, male, and healthy individuals, but do not study individual choice quality. Van Winssen, Van Kleef and Van de Ven (2016) discuss several potential reasons from the behavioral economics literature to explain the low overall take-up. While these papers do study the general lack of take-up of the voluntary deductible, our paper links individual choices of the entire Dutch population with

[4] The literature on insurance choice with choice frictions / behavioral choice foundations is summarized in the chapter by Handel and Schwartzstein (2019).

[5] There are also quite a few papers studying peer effects in health behaviors more broadly (see, e.g., Fadlon and Nielsen (2019), Chen, Persson and Polyakova (2019))

granular data on socio-economic and educational registries and employer-employee links to provide an in-depth analysis of determinants of choice quality. Moreover, another key benefit of our paper is that we are able to calibrate take-up under different behavioral models and compare that take-up observed choices.

The rest of the paper proceeds as follows. Section II describes health insurance in the Netherlands and describes our data. Section III presents our choice framework and consumer cost risk prediction model. Section IV presents our empirical analysis of deductible choice and section V presents the resulting heterogeneity in choice quality. Section VI presents our analysis of counterfactual policies and consumer welfare, concluding with a discussion of range of behavioral models that could potentially explain our findings. Section VII concludes.

II Institutional Context and Data

We exploit a unique consumer choice setting in the health insurance market in the Netherlands and link data on health insurance choices to data from various administrative registers. We present the institutional context and data here.[6]

II.A Health Insurance in the Netherlands

All individuals in the Netherlands are obligated to directly buy health insurance from a private health insurance market.[7] The Health Insurance Act of 2006 introduced a managed competition model in which the government strictly regulates the contents of the basic package of health insurance. The regulation also (i) prohibits price discrimination, (ii) prohibits the rejection of individuals from purchasing insurance and (iii) mandates that all individuals purchase coverage.[8] Insurers compete for consumers on premiums, provider choice, and supplementary insurance.[9] In 2015, there were 25 health insurers that together offered 53 separate insurance contracts. As shown in the left panel of Figure 1, yearly premiums for the mandatory health insurance with the smallest possible deductible have a mean of 1195 EUR and a fairly compact distribution around this mean.

Consumers enroll between mid November and the end of December for the following year.[10] During that period, health insurers advertise their insurance packages through various media. If no action is taken by the consumer, she will automatically extend her current contract. Relatively few consumers switch insurers each year (only 6.8% of individuals in 2015).

Each individual faces a compulsory deductible (375 EUR in 2015), but can opt for an extra voluntary deductible of 100, 200, 300, 400 or 500 EUR on top of this compulsory deductible (maximum total deductible of 875 EUR in 2015).[11,12] The compulsory deductible, introduced in 2008, has gradually increased from 150 EUR

[6] A more comprehensive overview of the health system and changes to the health insurance model in the Netherlands can be found in Kroneman et al. (2016).

[7] Every adult individual is required to choose a plan. Legally, individuals can only make the insurance policy choice for other adults if they have been provided with written consent from that adult. Children aged below 18 years are also required to have an individual plan. In practice, most parents register their children with the same insurer.

[8] To limit incentives for selection of consumers based on their health, the government has installed a sophisticated risk adjustment system. Yet, van Kleef, Eijkenaar and van Vliet (2019) show it is still profitable for insurers to attract healthy consumers.

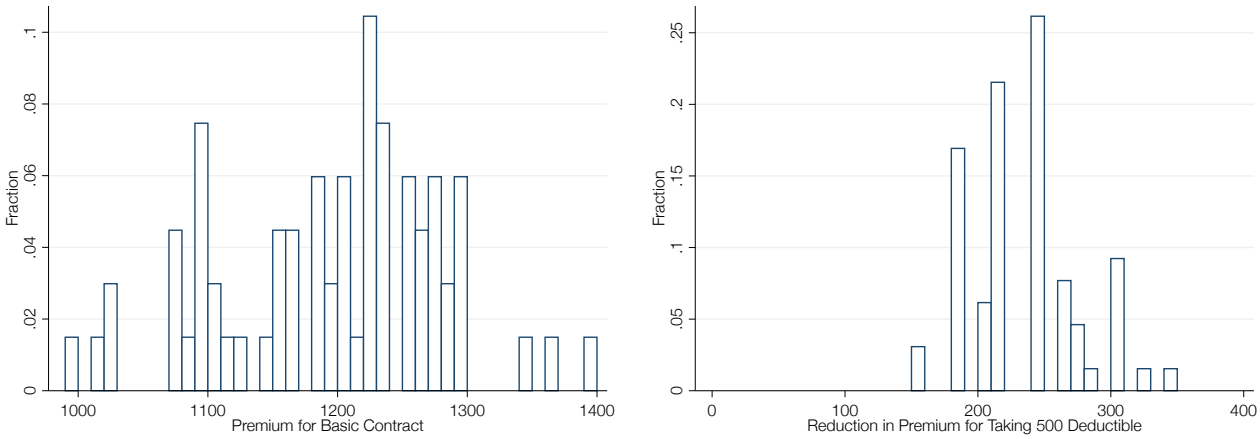
[9] The basic package covers drugs, doctor and hospital expenditures. Supplementary insurance covers dental care, additional physical therapy, alternative medicine, and other care. In 2015, approximately 90% of insurees bought supplementary insurance. The average premium for the supplementary insurance averaged 233 EUR in 2015.

[10] Generally, there are no brokers involved in the choice of health insurance in the Netherlands.

[11] The government does not mandate insurers to provide the choice of voluntary deductible. However, in practice, almost all insurers provide the option take the voluntary deductible. In 2015, of the 68 insurance contracts that have price information available on https://www.homefinance.nl, only one insurer (the new entrant ANNO12) does not provide a voluntary deductible option.

[12] Preventive, maternal and GP care is covered at zero cost by all insurers by law, and the deductible does not apply to the corresponding expenses. We exclude these preventive expenses from our cost prediction model.

FIGURE 1: DISTRIBUTION OF PREMIA AND PREMIUM REDUCTIONS



Notes: Histograms of yearly premiums in 2015 for basic coverage (left-hand side) and premium reductions for those contracts when electing a maximal voluntary deductible of 500 for a total deductible of 875 EUR (right-hand side). Data on prices are obtained from homefinance.nl.

in 2008 to 385 EUR in 2017[13], while the options for the extra voluntary deductible have remained the same. By opting for a higher deductible, consumers receive a premium reduction. Figure 1 shows the (unweighted) histogram of premium reductions consumers can get by electing the additional 500 EUR deductible across health plans offered in 2015. The distribution has a mean of 233 EUR and most of the mass lies between 200 and 300 EUR, making the deductible election a quite standardized decision across all insurance contracts.

Insurers can make agreements with employers, municipalities and various associations to offer group plans. These group plans are selected packages of basic and supplemental insurance on which the insurers offer premium reductions (*collectiviteitskorting*) of up to 10%. This feature in the insurance market leaves the choice of voluntary deductible unaltered for a given insurance contract. An exception to this are collective agreements between some municipalities and insurers for low-income individuals (*gemeentepolissen*), with income thresholds below 130% of the minimum wage. These policies are subsidized by municipalities, sometimes by covering the mandatory deductible amount, in which case they would not involve a deductible choice.[14]

The design of the compulsory deductible combined with a voluntary deductible has been a central topic of the policy debate. The desirability of consumer deductible choice has repeatedly been discussed in the Dutch parliament. In 2016, the Minister of Health Affairs, Schippers (2016) argued that having the option of a voluntary deductible increases general support for the health care system by the healthy, and makes individuals more aware of their health costs. Similar arguments have been put forward in recent exchanges in the Parliament in 2018 and 2019.

II.B Data and Sample

We use data on health insurance choices and health expenditures for all individuals in the Netherlands. The data is linked at Statistics Netherlands to other administrative registers, which provide information on their income,

[13]The size of the compulsory deductible was 350, 360 and 375 EUR in 2013, 2014 and 2015 respectively, then 385 EUR from 2016 onwards.

[14]In 2015, just over half a million individuals, about 3% of the population, were covered by this type of contract. As these contracts are mostly tied to generous supplemental coverage, the premium remains high relative to basic plans with high-deductible option, which is still the better option for predictably health individuals (see Douven et al. (2019)).

Table 1: Distribution of Deductible Choices

Default Deductible	90.94%
Extra Deductible (+100 to +500EUR)	9.06%
Breakdown of Extra Deductible Choices	
+100EUR	10.64%
+200EUR	10.41%
+300EUR	6.02%
+400EUR	1.72%
+500EUR	71.21%

Notes: This table shows the breakdown of deductible choices in 2015. A large majority (90.94%) sticks to the default 375 EUR deductible. Of the 9.06% individuals that take an extra deductible, most individuals take the 500 EUR extra deductible.

wealth, education, employment and other demographic variables.

We restrict attention to all individuals who are at least 18 years old in January of the year in which they decide on their health insurance contract and deductible. We exclude from the sample adults who have incomplete or unreliable health data records in the two previous years.[15] The remaining sample consists of about 13.25 million adults in each year. As explained in Section III.B, we use a random sample of 1.25 million of these individuals to estimate and calibrate a cost prediction model, leaving approximately 12 million adults each year for the analyses, which we call our baseline sample.

Health Insurance Deductible Data on health insurance contract choices in the years between 2013 and 2017 are obtained from Vektis, an organization that is responsible for the collection of data from all health insurers. Our data include only information on an insurer and deductible choice. We do not observe the choice of provider network nor whether individual takes supplementary insurance, but these choice dimensions are orthogonal to the deductible choice except for minor price differences. Table 1 shows the take-up of different deductible amounts in 2015. The voluntary deductible take-up in our sample is 9.06% in 2015. More than 2 out 3 individuals opting for an extra deductible take the maximum extra deductible of 500 EUR.

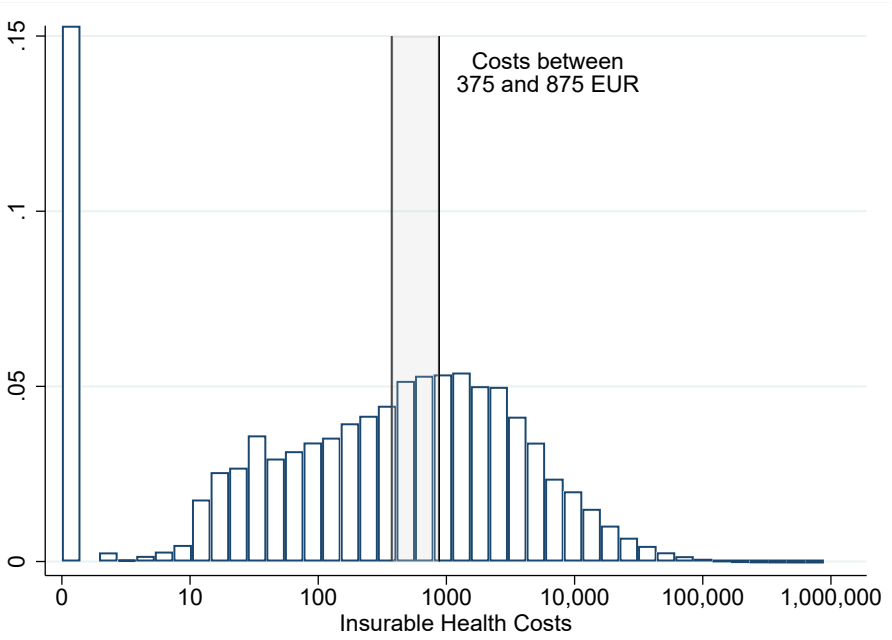
Health Care Costs Data on health care costs contain annual health care expenditures by category. The categories are medicines, hospital care, geriatric care, paramedical care and physiotherapy, mental health care, aids and tools for health, health care in foreign countries, health care transport, multidisciplinary care, sensory handicap care, and other care. In addition to these categories which are subject to the deductible, we also have data on neonatal and maternal care, care by GPs and home care, where cost sharing does not apply.[16]

Figure 2 presents the distribution of the (log) aggregate health care expenditures that are subject to cost-sharing in 2015. This aggregate distribution is skewed with about 19 percent of individuals making zero expenditures and more than 10 percent of individuals spending more than 5000 EUR. Table 2 presents the distributions of annual expenditures for the different categories of medical spending. These distributions are similarly skewed. Hospital expenditures (1,388 EUR), drugs expenditures (320 EUR) and mental health care (243 EUR) are the three categories with the highest mean spending.

[15] Insurers in the Netherlands are split in two categories: insurers who actually bear the risk and proxy insurers who only act as middleman. Vektis, the data provider, deems data from the about 10 proxy insurers over our sample period, covering approximately 4% of people ($\approx$ 500,000), to be unreliable. Hence, we do not use these observations in our analysis.

[16] Some miscellaneous items are also exempt from cost sharing. These include preventative care such as breast cancer screening and flu shots, as well as costs made for organ donation. We cannot separately identify these costs from hospital care, so our measured insurable costs will be slightly overestimated.

Figure 2: Distribution of Insurable Health Care Costs



Notes: This figure shows the distribution of the $\log_{10}$ of total yearly insurable health care costs in 2015, for all individuals in our baseline sample. 13.1% of individuals have health costs falling in the 375 to 875 EUR interval.

Other Data We obtain information on other variables from a number of administrative registers and link these to the health and insurance data. Our data includes standard demographics like age, gender and household status. We use third-party reported information from tax registers on household income and household wealth. The former includes pre-tax income from labor, self-employment and capital and government transfers. The latter includes information on net worth, liquid and other financial assets, mortgage and other debt. We also observe data on the highest formal education level attained for more than half of the sample. These data also include information on the specific field of study for individuals who proceed past high school. Finally, we use employer-employee data to link individuals at the firm level and identify their sector of employment. We provide more detail about the different registers and variables in the Data Appendix A. Table 3 provides some summary statistics for the year 2015.

III Choice and Risk Framework

In this section we develop a simple model of deductible choice that guides our analysis. We first develop a stylized model of choice. We then turn to a model to predict health care cost; the central input into a frictionless, rational choice model.

III.A Deductible Choice Model

Each individual is subject to a compulsory deductible of 375 and can choose a voluntary deductible d at corresponding premium p from menu $\Omega = \{(d, p_d)\}$. An individual draws health cost x from an individual-specific distribution $F_i(x)$. Depending on her deductible choice d, health cost translates into an out-of-pocket expense

Table 2: Distribution of Annual Health Care Costs

	Mean	p10	p50	p90	p99
All Care	2,695	86	495	6,032	35,974
Insurable Care	2,272	0	332	5,043	31,133
Hospital Care	1,388	0	85	2,829	21,575
Medicines	320	0	53	758	3,253
Mental Care	243	0	0	0	4,801
Tools and Medical Aid	107	0	0	145	2,284
Geriatric Care	53	0	0	0	0
Transport	45	0	0	0	1,081
Multidisciplinary Care	33	0	0	124	397
Physiotherapic Care	32	0	0	0	1,095
Dental Care	26	0	0	0	825
Other Care	7	0	0	0	151
Sensory Handicap Care	3	0	0	0	0
Always Insured Care	423	75	121	327	8,042
Nursing Care	228	0	0	0	7,587
GP Care	157	75	119	272	659
Maternal Care	37	0	0	0	1,796
Observations					11,991,629

Notes: This table shows the distribution of health expenditures by subcategory, for the full sample in 2015. Expenditures are divided into insurable expenditures, that are subject to cost sharing (and to which the deductible applies) versus always insured expenditures, that are not subject to cost sharing. All values are in EUR.

Table 3: Summary Statistics

	Mean		Mean
Demographics		**Household Financial Status**	
Male	48.8%	Gross Household Income	73,289
Age	50.3	*10th Percentile*	*20,077*
Has Children	69.2%	*Median*	*60,358*
Has a Partner	62.9%	*90th Percentile*	*135,981*
Education Level		Household Net Worth	166,890
Less than High School	13.2%	*10th Percentile*	*-28,918*
High School	24.1%	*Median*	*32,694*
College	16.8%	*90th Percentile*	*403,923*
Further Studies	0.6%	Mortgage Debt	54.1%
Unknown	45.4%	Other Debt	34.2%
Employment Status		Savings > 2000 EUR	80.4%
Employee	44.3%		
Self-Employed	9.9%		
Retired	24.2%		
Student	6.3%		
Other Not Working	15.3%		
Observations			11,991,628

Notes: This table shows summary statistics for the full sample in 2015.

$s = min\{d, x\}$. We denote by $G_{i,d}(s)$ the distribution of out-of-pocket spending, derived from $F_i(x)$ and the deductible choice d. Expected utility for a rational individual in a frictionless environment is defined, therefore, as:

$$U_{i,d} = \int u_i(W_i - p_d - s)G_{i,d}(s)ds. \tag{1}$$

Using this definition of expected utility, we can define an individual's certainty equivalent from choosing one contract as $CE_{i,d}$, where $U_{i,d} = u_i(W_i - CE_{i,d})$.

Using this benchmark for frictionless decision making, we define frictions as driving the wedge between the actually observed choice and the choice made by a rational individual in a frictionless environment. Denoting the certainty equivalent for individual i's observed choice by CE_i and for the utility-maximizing choice by CE_i^*, we define the welfare loss due to choice frictions expressed as a money-metric as:

$$\Delta w_i^* = CE_i^* - CE_i.$$

A central decision variable when considering to elect a higher deductible higher than the compulsory deductible of 375 EUR is the chance that expenditures stay below 375 EUR, which we denote by π_i. In theory, the optimal decision depends on the probability distribution of expenditures between 375 EUR and 875 EUR too, but the share of expenditures that fall in this range is small. We also simplify the decision to a binary choice between the baseline deductible of 375 EUR and adopting the full 875 EUR deductible while gaining the associated premium savings. Empirically, most individuals who elect a deductible higher than the compulsory deductible choose the maximum possible deductible. As we discuss below, interior choices between the two levels are not easily rationalized under standard preferences.

Under this simplified environment, we approximate expected utility by:

$$U_{i,d} \approx \pi_i u_i(W_i - p_d) + (1 - \pi_i)\, u_i(W_i - p_d - d), \tag{2}$$

and the contract space, including the following two contracts:

$$\Omega = \{(0, 0), (500, -250)\}.$$

This setup demonstrates the relative simplicity of the environment we study. In expected payoff terms, $\bar{\pi} = 0.5$ is the (approximate) threshold between optimally choosing the additional 500 EUR deductible and saving 250 EUR in premium. As we note, risk aversion does not increase this threshold by much. For a standard value of absolute risk aversion of 10^{-5} (e.g., Cohen and Einav (2007)), this threshold increases to 0.5006. Even for extreme risk aversion of 10^{-3}, this threshold is still only 0.56.[17]

Based on this framework, the central aspect of decision making necessary is an estimate of individuals' risk of spending more than 375 EUR. To do so, we develop a cost prediction model.

III.B Cost Prediction Model

For every individual, we generate yearly health risk predictions, with the explicit goal of evaluating the choice of the voluntary deductible. We set up our prediction model as a binary classification algorithm that predicts the probability (π_i) of having health expenditures below the compulsory deductible level of 375 EUR. Thus, our

[17] For a CARA utility function of the form $u(z) = -\sigma e^{-\sigma z}$, the cutoff value π^* for switching to the high deductible being optimal is given by $(1 - e^{\sigma 250})/(e^{-\sigma 250} - e^{\sigma 250})$.

prediction model accords with the underlying behavioral model.[18]

The yearly predictions of π_i are made using an ensemble learning model consisting of a random forest model, a boosted regression trees model and a LASSO model. Using such an ensemble learner is a standard technique to maximize prediction accuracy of a classification problem (Einav et al. (2018)). We only include predictors that are known at the time of choice of deductible (at the end of year $t-1$). The predictors that we include are: gender, year of birth, pre-tax household income in deciles ($t-2$, $t-1$), working status, education level, education field, and past health spending per category ($t-2, t-1$). In each year, there are approximately 20 variables for per-category health spending, so we have a fine level of detail with which to predict future medical spending. On average, we have approximately 50 predictors in our model every year.[19]

Our prediction algorithm follows four steps, similar to the prediction analysis in Einav et al. (2018). First, several key parameters of the random forest, boosted regression trees, and LASSO models are tuned. Second, these three separate prediction models are trained using a training sample. Third, the obtained predictions are combined into an ensemble predictor. Finally, the ensemble prediction is calibrated. We train the ensemble learner algorithm on a random sub-sample of 800,000 individuals. The training sample contains an additional 450,000 observations to combine the predictors and calibrate the ensemble predictor to observed data. All the results and plots in the analyses in this paper are then using only the hold-out sample of about 12 million observations each year. In Appendix B, we provide more information on the detail of each step of the prediction.

Figure 3 describes the precision and outcomes of the prediction model. Panel A shows a bin scatter plot of the share of low-cost realizations by the predicted low-cost probability. The relationship between *ex ante* probabilities and *ex post* realizations is very strong as all observed shares are close to the 45 degree line. Appendix Figure B.1 shows that the prediction model is similarly well-calibrated for subgroups of individuals with different ages, education levels and income quartiles as well as by 500 EUR deductible take-up.[20] The ROC curve in Panel B shows that the ensemble model performs best and improves on the individual models. Panels C and D illustrate the predictive value of the model, comparing the distribution of realized cost for the top 5% and bottom 5% in terms of predicted low-cost probabilities. The ex-post spending for the group that is predicted to be healthiest is much more skewed towards the low end of the distribution than the same distribution for the consumers predicted to be sickest, which is skewed towards the high end of the cost distribution.

Having established the predictive performance of the model, Figure 4 presents the histogram of the predictions for the *ex ante* probability of being in the low spending group. There is substantial dispersion in predicted risks over the full range of potential probabilities. The distribution is bi-modal, with a substantial share of individuals having either a very low probability or a very high probability of being low spenders. We include threshold measures for choosing the 500 EUR deductible to demonstrate that the distribution of risk places a significant share of the population well above and below the cutoffs respectively.

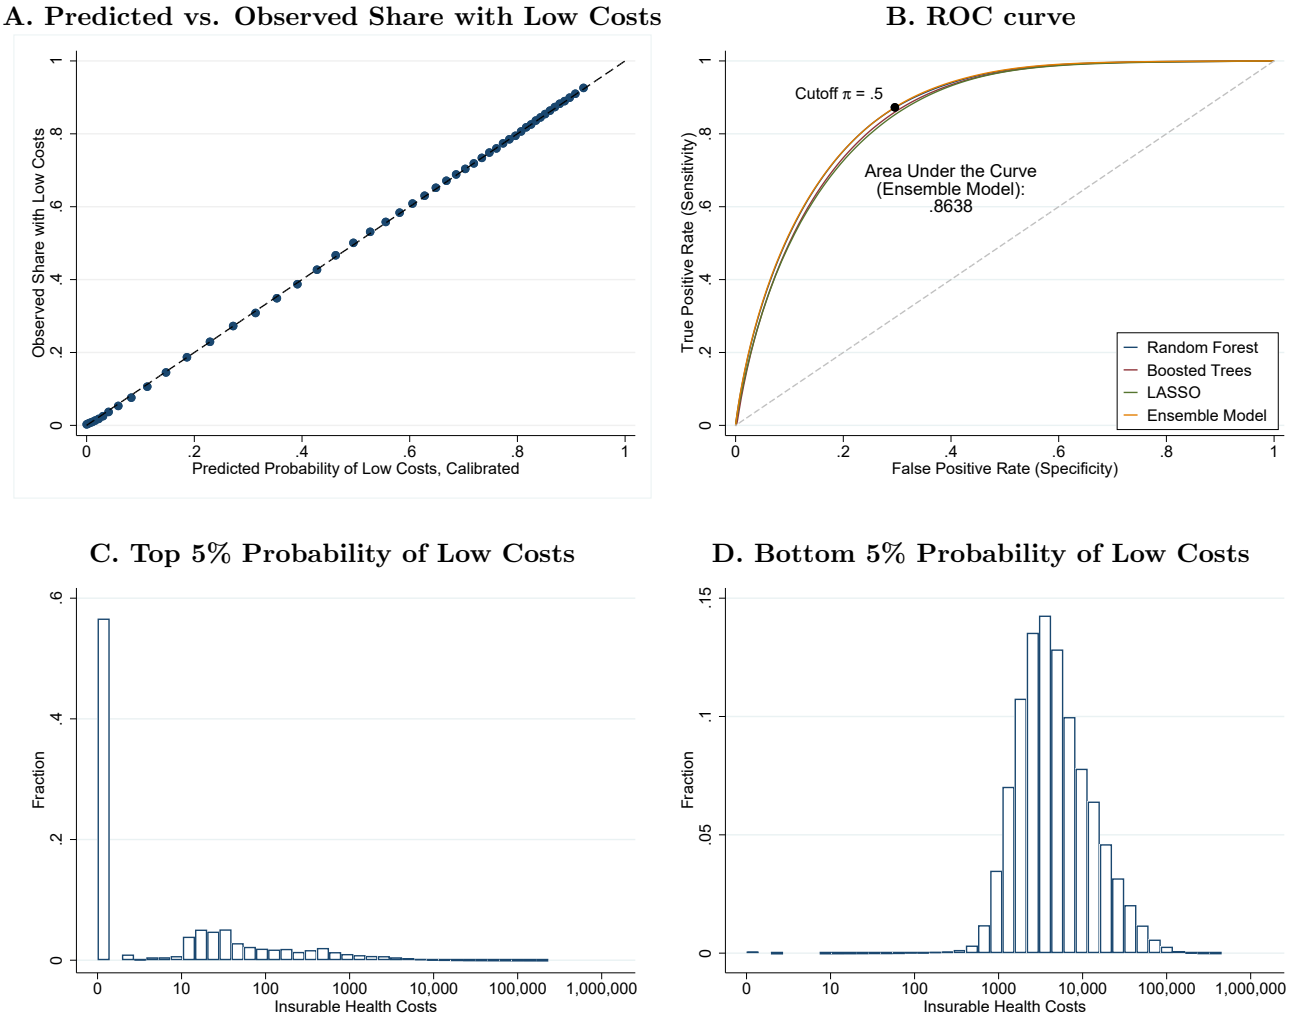
Taken together, these figures show that health expenses are, to a large extent, predictable and bi-modal in our population. These features allow us to assess, with a good degree of robustness, whether a given individual

[18] The empirical prediction model also underscores why intermediate ranges between 375 and 875 EUR are not useful choices. The distribution of health spending makes falling in that range of expenditures extremely unlikely. Therefore, predicting risk in this range is extremely difficult and choosing such a choice is almost never *ex ante* optimal. Figure 2 shows that the share of *ex post* realized expenditures that fall between 375 and 875 is 13.1%. When *ex ante* predicting which bracket individuals' costs would fall into, the sum of the raw predicted probabilities for the intermediary brackets is smaller than 1%. We provide further detail on this in Appendix B.3.

[19] We have a different number of predictors in some years, as the categorization of health costs changes slightly in our study period. Every year, we include all health cost categories in our data set as predictors.

[20] In Panel A of Appendix Figure B.1, the prediction accuracy is plotted for individuals who take the 500 EUR deductible, and individuals who do not. While individuals with a 500 EUR extra deductible have an *ex post* higher chance to make low costs than our predictions, differences in prediction accuracy between the two groups are small, suggesting a minor role for private information about health risk or moral hazard conditional on the predictors.

FIGURE 3: PREDICTED VS. REALIZED COSTS



Notes: Panel A presents a binned scatter plot of our predicted probability of having low costs against the realized share of individuals with low costs. Panel B plots the ROC curve of the different prediction methods used. The bottom figures present ex-post cost realizations of individuals with predicted low (Panel C) and predicted high (Panel D) costs. The year is 2015 for all Figures.

is better off electing a high or low deductible.

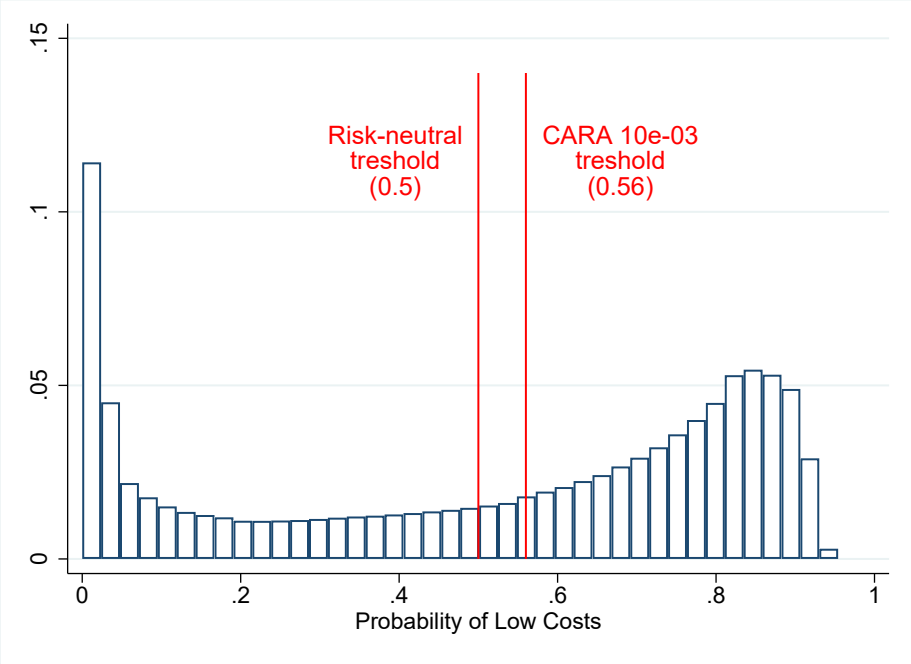
IV Empirical Analysis of Deductible Choice

In this section we study the determinants of deductible choice. Building on our simple model, we relate choice to health risk, the primary component of deductible choice in a frictionless, rational model. We then incorporate individual and environmental factors that capture choice elements outside of the standard model that can drive a wedge between the observed and optimal choices.

We begin by presenting non-parametric graphical evidence relating risk to deductible choice. These plots relate health risk to deductible take-up, overall and for specific subgroups. We then estimate a set of simple regressions, using variation across and within individuals. We rely on a simple OLS regression in a linear probability model:[21]

[21] As alternatives, we relax the linearity assumption: $Y = \alpha + \gamma X + [\beta + \nu X] \times 1[P(costs < 375) \geq .5] + \epsilon$, and also consider a

FIGURE 4: DISTRIBUTION OF COST PROBABILITY PREDICTIONS



Notes: This figure shows the distribution of the predicted probabilities of having health costs below 375 EUR. These probabilities are obtained when predicting the binary variable (having insurable health costs below 375) with the ensemble machine learner described in Section III.B, and further in Appendix B. The figure presents the risk-neutral threshold for someone to choose the 500 EUR incremental deductible if the incremental premium reduction is the modal incremental premium reduction of 250 EUR. It then presents the same threshold for extreme risk-aversion (CARA coefficient $1 * 10^{-3}$).

$$Y = \alpha + \gamma X + [\beta + \nu X]P(costs < 375) + \epsilon \tag{3}$$

where Y is an indicator variable taking the value of 1 when an individual takes the 500 voluntary deductible and 0 otherwise, $P(costs < 375)$ is the predicted probability of having costs lower than 375 EUR (π_i in the theoretical section), and X includes all variables of interest. The primary coefficients of interest are γ and ν. The former captures how different observables affect the intercept, i.e., the average take-up of the 500EUR deductible by individuals who are the sickest (with $\pi_i = 0$). The latter measures how different factors affect the relationship between risk and deductible choice. $\gamma + \nu$ captures the impact on average take-up by individuals who are the healthiest (with $\pi_i = 1$). Each regression also includes year and insurer fixed effects. The insurer fixed effects control for potential differences in insurer marketing / steering and/or differences in insurer incremental deductible premium, though as we showed earlier there is limited dispersion in the latter.

IV.A Relationship Between Risk and Deductible Choice

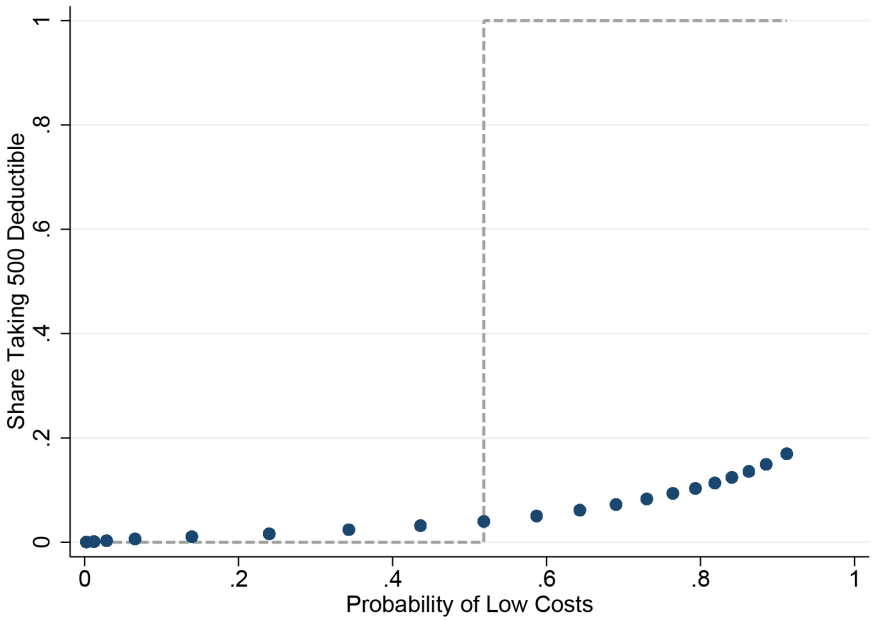
We begin by focusing on the relationship between predicted health and deductible choice in the whole population. Predicted health is the only driver of deductible choice from the standard model over most ranges of risk (recall that risk aversion has little impact on choices as captured by the threshold cutoff to choose the high deductible).

Figure 5 plots the empirical relationship between predicted health risk and deductible choice.[22] For com-

probit model: $Pr(Y = 1) = \Phi(\alpha + \gamma X + [\beta + \nu X] \times P(costs < 375) + \epsilon)$. The findings are unchanged, as shown in Appendix Table C.2.

[22] Appendix Figure C.2 shows the same plot disaggregated for the different years in our sample period.

Figure 5: Take-up of Voluntary Deductible as Function of Predicted Health Costs



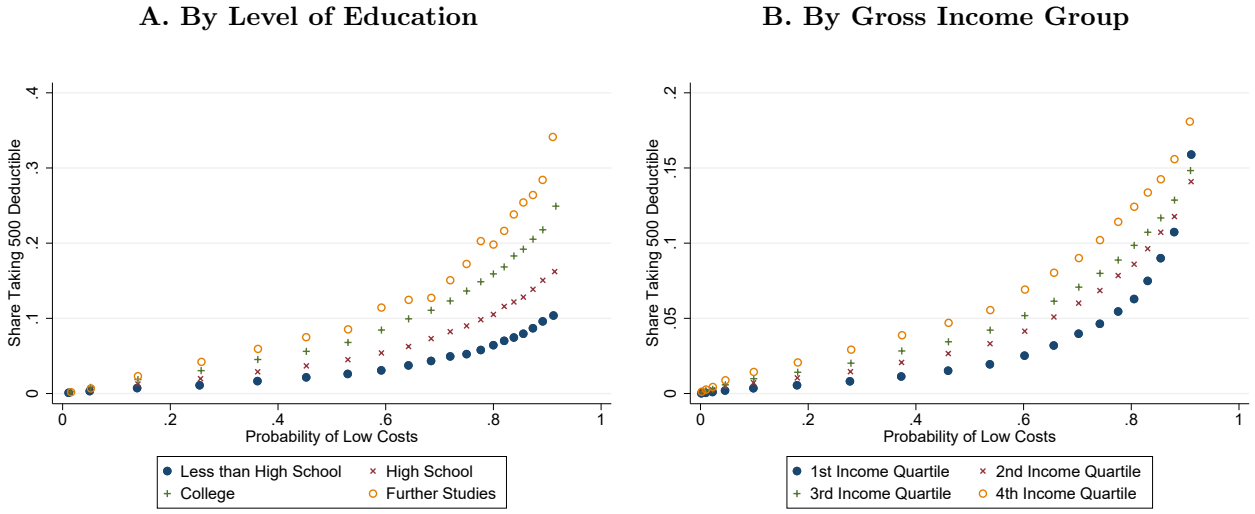
Notes: This figure shows a binned scatterplot of the relationship between the predicted probability of having costs below 375 EUR (the compulsory baseline deductible) and the take-up of the voluntary 500 EUR extra deductible.

parison, the grey dashed line presents the pattern predicted by the rational, full information benchmark. Two key facts emerge. First, as expected, people who are healthier are more likely to elect the higher incremental 500 EUR deductible. Second, the relationship between risk and deductible choice is substantially weaker than one would expect if consumers were making utility-maximizing choices in the standard model. For example, the share of consumers in the healthiest predicted health bin electing the high deductible is only 17%, despite the fact that 100% would gain *ex ante* from taking the high deductible. Indeed, the simple choice exercise conducted in Section III suggests that, assuming individuals know their predicted health risk, the take-up rate should jump from 0 to 100% around a low-cost probability of .5. The relationship in Figure 5 is in sharp contrast to those predictions.[23]

The same two key facts are confirmed when using only within-individual variation in predicted health risk (see Appendix Table C.1). Increases in health risk lead to statistically significant decreases in extra deductible take up, but the effect size of the response is small. Our estimates indicate that the response rate is proportional to the size of the change, but individuals are more responsive to negative (relative to positive) changes in predictable health risk.

[23]We note that moral hazard could underlie some of the positive correlation between deductible choice and health risk making, if anything, the selection on risk is *smaller* than what we observe. The confounding effect is arguably small, though. We use *ex ante* predicted health risks rather than *ex post* cost realizations limiting scope for actual spending to impact the relationship. We also find that the difference in predicted and realized risk for individuals who do take the extra deductible is very small (see Panel A in Appendix Figure B.1). In the literature, the moral hazard effects of health coverage are generally estimated to be small relative to the effects we find. This is confirmed for the specific context by Remmerswaal, Boone and Douven (2019) using an age-discontinuity in the deductible choice at 18 years old.

FIGURE 6: DEDUCTIBLE TAKE-UP BY EDUCATION AND BY INCOME



Notes: These figures show binned scatter plots of the relationship between the predicted probability of having costs below 375 EUR (staying under the voluntary deductible range) and the take-up of the voluntary 500 EUR extra deductible, by education level in Panel A and by household gross income quartile in Panel B. In Panel B, we have excluded the group of individuals with gross income below minimum social assistance, which mostly consists of students, self-employed and households with negative capital income.

IV.B Subgroup Specific Relationship Between Risk and Deductible Choice

We now turn to understanding how different observable factors change deductible choice with respect to health risk. We do so both with simple plots, following Figure 5 but dividing the population by observable characteristics, and by formalizing these results with our simple regression framework.

Figure 6 plots the relationship between health and deductible take up by education level and income. Panel A shows a large difference in the relationship by education level. Those in the healthiest predicted risk decile with a college degree (i.e., bachelor or master) elect the higher deductible about 23% and even 30% with an advanced degree. In contrast, those with less than high school education in the healthiest predicted decile elect the higher deductible only 10% of the time and those with high school education only approximately 15% of the time. For all of these education levels, when people are predicted to be sick they almost never elect the higher deductible.

Panel B of Figure 6 shows the same relationship by quartiles of gross income (including capital income and government transfers). The relations are similar as with education, where higher levels of income are associated with better deductible choices.[24]

We next turn to estimating regression models following equation 3. Table 4 presents the intercept coefficients α and γ and the slope coefficients β and ν.[25]

There is relatively little variation in the demographic-specific intercepts. For those in worst health — *ex ante* probability of zero of having cost less than 375 EUR — higher education is associated with a lower rate of take-up of the higher deductible. The effect of income, however, is the opposite. As can be expected from the graphical evidence, some of these differences change when relaxing the linearity assumption on the relation between take-up and risk, but they are consistently small (see Appendix Table C.2).

There is much more variation in demographic-specific slopes. The interaction with the predicted health risk is

[24] We note that income effects may directly increase the take-up of deductibles, but their interaction with the predicted low-cost probability is ambiguous.

[25] For completeness, Appendix Table C.3 shows the estimates when not including interaction terms between controls and predicted risk.

Table 4: Deductible Take-Up: Baseline Regression Estimates

	Take-up of 500 Deductible	
	intercept	*slope*
High School	-0.011***	0.057***
College Degree	-0.034***	0.165***
Further Studies	-0.047***	0.226***
2nd Income Quartile	0.004***	-0.007***
3rd Income Quartile	0.004***	0.007***
4th Income Quartile	0.002***	0.039***
36 to 50 years old	0.020***	-0.045***
51 to 65 years old	0.029***	-0.047***
65+ years old	0.034***	-0.082***
Male	-0.004***	0.025***
Has Partner	-0.002***	0.013***
Has Children	0.004***	-0.028***
Self-employed	-0.006***	0.026***
Constant	-0.041***	
Prob. Low Costs		0.098***
Year and Insurer FE	YES	
Observations	57,100,388	

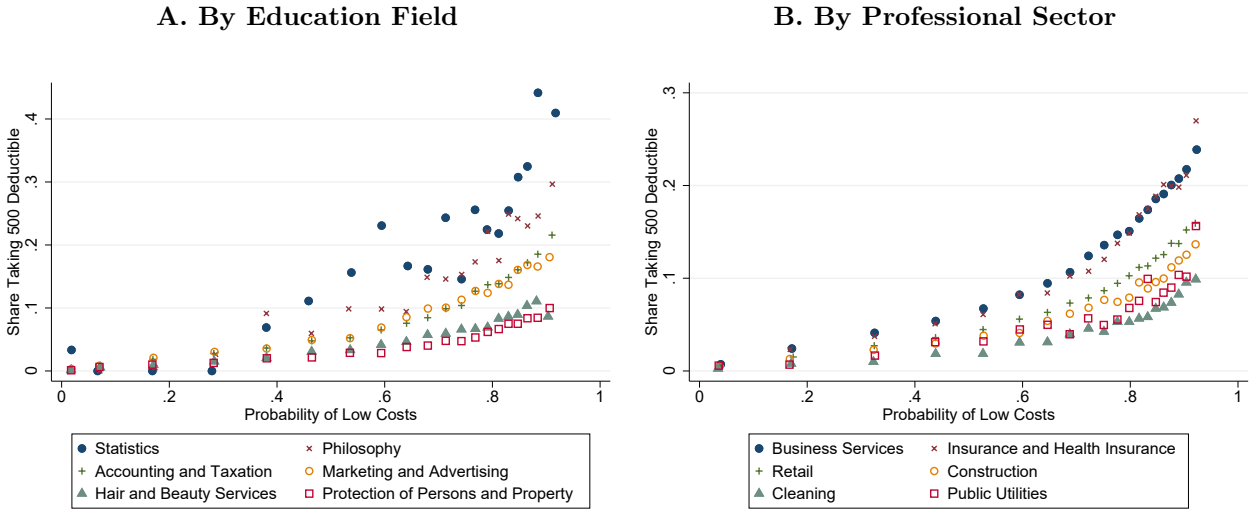
Notes: This table plots coefficients from our regressions studying deductible choice, as explained in Section IV. Each variable is interacted with the probability of having low health expenses; the impact on the intercept is reported in the first column, and the impact on the slope in the second column. The dependent variable in all specifications is a dummy that takes value of 1 when the individual takes up the voluntary 500 EUR extra deductible. The prob. costs < 375 EUR variable is obtained from our prediction algorithm. The reference groups for the different demographic categories are: 1st income quartile, education lower than high school or unknown, and age between 18 and 35. *** p<0.01, ** p<0.05, * p<0.1 with robust standard errors.

larger for those with higher education reflecting the fact that individuals are more responsive to their health status in selecting the higher deductible with higher education levels. The interaction with education is both statistically significant and economically meaningful. An individual in good health — *ex ante* very high probability of being low cost — who has completed graduate studies beyond college is 23% more likely to take up the high deductible that an equivalent person with less than a high school education. The interaction of income and the gradient of take-up has the same pattern though the effects are modest once we include controls; the highest income quartile is slightly less than 4% more likely to take up the high deductible if they are in good health compared to the lowest income quartile.[26]

Table 4 also presents the effects of age, gender and household composition on deductible choice, controlling for health risk, income and education level. There are statistically significant differences in responsiveness to underlying health risks, though the magnitude of the effects are relatively small. We also note that, despite

[26] Appendix Figure C.1 plots the relation between take-up and household income. The income gradient is as important in magnitude as for predicted health risks, but once we control for predicted health and other variables capturing socio-economic status, the gradient becomes nearly flat. We further confirm the limited role of income effects by investigating how deductible choice responds to income shocks. Table C.1 indicates that deductible choices do not vary with within-individual income variation.

FIGURE 7: DEDUCTIBLE TAKE-UP BY EDUCATION FIELD AND BY PROFESSIONAL SECTOR



Notes: This figure shows for 6 fields of study and 6 professional sectors a binned scatterplot of the relationship between the predicted probability of having costs below 375 EUR (the compulsory baseline deductible) and the take-up of the voluntary 500 EUR extra deductible. Refer to Tables C.5 and C.6 for an overview of the deductible take-up in all fields and sectors, respectively.

the relative simplicity of the models we estimate, these effects are very robust to alternative specifications. For brevity, we present those results in Tables C.2.

Overall, Table 4 demonstrates that the strongest relationship between deductible take-up and observable characteristics is for education level. This is indicative of the potential role of expertise, cognitive ability and information frictions in insurance choices. To shed more light on the role these effects may play we perform the same analysis as above but use richer data on the specific field of education and professional sector of employment.

Figure 7 plots the relationship between deductible choice and predicted health risk by education field and professional sector. Since there are many education fields and professional sectors, we present only 6 specific fields and sectors that are indicative of the broader patterns. Statistics majors are the most responsive to predicted health risk: they choose the additional deductible approximately 43% of the time when they are in the healthiest predicted health bin and choose the additional deductible almost never when they are in the sickest predicted bin. The effect stands in stark contrast to those with training in "Protection of Persons and Property" or "Hair and Beauty Services." Even for the healthiest group in those fields take-up of the higher deductible is only approximately 10%. Similarly, for professions where people are likely to be more analytical in nature and professions that require more advanced schooling, deductible choice is also higher for those with low risk — the prediction of the standard, rational model. Despite that, however, we also see that even for those in the insurance industry, take-up is only around 30% for those in the best health.

Table 5 presents the corresponding regression analysis, including baseline controls for predicted health risk, income, education level, age, gender and household structure. Even controlling for these other factors, more quantitative / analytic professions (e.g., statistics) are more responsive to predicted health when making deductible choices (column 1). For example, among the predictably healthy, someone with statistics training is 28.2% more likely to choose a higher deductible, controlling for age, income, gender, and education level than someone with hair and beauty training. Column 2 of Table 5 also confirms that more analytic professions (e.g., business services and insurance) are more responsive to predicted health when making deductible choices. For example, someone in the insurance sector who is predictably healthy is approximately 8% more likely to choose

Table 5: Deductible Take-Up Regression by Subgroup

	(1) Education Field		(2) Professional Sector		(3) Liquidity and Financials		(4) Environment	
	intercept	*slope*	*intercept*	*slope*	*intercept*	*slope*	*intercept*	*slope*
Statistics	-0.042**	0.247***						
Philosophy	-0.003	0.046***						
Accounting and Taxation	-0.003***	0.024***						
Marketing and Advertising	-0.000	-0.004						
Hair and Beauty	0.007***	-0.035***						
Protection of Persons	0.008***	-0.068***						
Business Services			-0.012***	0.045***				
Insurance			-0.025***	0.078***				
Retail			-0.002***	-0.002*				
Construction			-0.001	-0.018***				
Cleaning			0.003***	-0.033***				
Public Utilities			0.006***	-0.008*				
2nd Net Worth Quartile					0.003***	-0.004***		
3rd Net Worth Quartile					0.000*	0.021***		
4th Net Worth Quartile					-0.002***	0.061***		
Has Savings > 2000EUR					-0.006***	0.028***		
Has Mortgage Debt					-0.000	0.005***		
Has Other Debt					0.005***	-0.023***		
Share of Colleagues with 500 Ded.							-0.105***	0.459***
Share in Postcode with 500 Ded.							-0.329***	1.055***
Father With 500 Deductible							-0.029***	0.288***
Mother With 500 Deductible							0.015***	0.294***
Constant	-0.043***		-0.050***		-0.042***		0.010***	
Prob. Low Costs		0.101***		0.117***		0.094***		-0.060***
Baseline Controls	YES		YES		YES		YES	
Year and Insurer FE	YES		YES		YES		YES	
Observations	30,799,129		32,299,835		57,013,765		16,938,401	

Notes: The regressions follows our baseline specification (see Table 4). Additional controls are: in Column (1), dummies for six selected educational fields of study, as well as their interactions with health risk. The reference category for field of study is all other fields of study; in Column (2) dummies for six selected professional sectors, as well as their interactions with health risk. The reference category is all other sectors; in Column (3), a dummy for liquidity (household savings¿2000EUR), a dummy for having household mortgage debt and other household debt, household net worth quartiles, as well as their interactions with predicted health risk; and in Column (4), the fraction of individuals taking up an extra 500 EUR deductible in firm and neighborhood, and dummies for whether the father or mother is taking up an extra 500 EUR deductible. Note that the shares are calculated excluding the individual for which the share is calculated (i.e. the person's take-up is excluded from both numerator and denominator), and shares are calculated only if there are more than 10 individuals that firm or neighborhood. ***p<0.01, ** p<0.05, * p<0.1 with robust standard errors.

a higher deductible, controlling for age, income, gender, and education level, than someone in the public utilities sector.

To shed further light on the relationship between the specific field of study and choice quality we report key take-up measures for a selection of fields in Table 6. Columns 1 and 2 present the share taking up the high

Table 6: Deductible Take-up and Field of Study

Education Field	(1) Take-up of 500 Deductible	(2) Probability Low Costs	(3) Take-up of 500 Ded. | Being Predictably Healthy
1 **Statistics**	29%	87%	34%
2 Mathematics	21%	85%	27%
3 Physics	21%	91%	26%
4 Architecture and town planning	18%	88%	21%
5 Physical science	18%	82%	22%
6 Earth science	18%	88%	21%
7 **Philosophy and ethics**	17%	82%	21%
8 Medicine	17%	83%	20%
16 Sociology and cultural studies	14%	82%	18%
17 Mining and extraction	14%	91%	17%
18 Economics	14%	84%	17%
19 Humanities and Arts	14%	84%	18%
41 **Accounting and taxation**	11%	78%	14%
42 Agriculture, forestry and fishery	10%	81%	13%
43 **Marketing and advertising**	10%	80%	13%
83 Secretarial and office work	5%	65%	7%
84 **Protection of persons and property**	4%	78%	6%
85 Child care and youth services	4%	66%	6%
86 Computer use	4%	65%	6%
87 **Hair and beauty services**	4%	65%	5%
90 Literacy and numeracy	2%	62%	4%

Notes: For a selection of fields of study, this table shows: in Column (1), the fraction of individuals who take-up the 500 EUR extra deductible, in Column (2), the fraction of individuals with a probability of low costs < 375 EUR, and in Column (3), the fraction of individuals who take-up the 500 EUR extra deductible, conditional on having predicted health costs < 375 EUR. The full list of fields is provided in Appendix Table C.5.

deductible and the predicted low-cost probability respectively. The primary results of interest are presented in column 3, which shows the rate of take-up of the high deductible among those with a high probability of having low cost — the group for which we expect high adoption under the standard model. The table shows that quantitative fields are grouped at the top of the table, exhibiting greater responsiveness to predicted health risk when making deductible choices, while those in less quantitative fields are grouped at the bottom of the table, exhibiting lower responsiveness. An exhaustive list of education fields is presented in Appendix Table C.5. We present a similar analysis of myriad professions in Table C.6 in the appendix, showing a very similar gradient by professional sector. More analytical sectors exhibit greater responsiveness to predicted health risk when making deductible choices while those in less analytical sectors exhibit lower responsiveness.

In addition to an individual's human capital and income, we observe a range of additional variables related to a household's financial capital. We study the relationship between deductible choices and wealth, measured by the household's net worth, debt (mortgage or any other debt) and a measure of liquidity that takes on a value of 1 if a household has more than 2000 EUR in liquid savings and 0 otherwise.

Table 5 presents the results of a regression examining the association between these financial variables and

incremental deductible take up, controlling for predicted health spending and our baseline controls. We find that household liquid savings are positively correlated with deductible take up: having liquid savings of greater than 2000 EUR is associated with a 0.8 percentage point increase in deductible take up. Note that in theory, liquidity and debt constraints could either increase the demand for insurance (to avoid large expenditures) or reduce the demand for insurance (to avoid paying the premium) (see Ericson and Sydnor (2018)). The sign of the effect we find is consistent with the former explanation. In line with this, we also find that households who are in debt (excluding mortgage debt) are also less likely to take-up the deductible. The effects, however, are small in both cases. Finally, we find that take-up rate for wealthier individuals is higher and this effect is fully driven by wealthier individuals with better health. That is, wealthier individuals are more responsive to taking the incremental deductible as they become healthier. Hence, rather than capturing wealth effects on insurance choices, this result could be indicative of information frictions for people with fewer financial resources.

IV.C Peer Effects on Choice Quality

Thus far we have shown the impact of socio-demographic characteristics and, especially, long-run human capital on deductible choice quality. In this section we study how environmental factors, measured by exposure to peers' choices, impact choice quality. Specifically, we investigate the impacts of the deductible choices by (i) co-workers, (ii) neighbors and (iii) parents.

Table 5 presents the cross-sectional regression estimates for the association between individuals' take-up of the deductible and the take-up rates by their respective peers, closely following our main regression equation and controlling for health risk, baseline demographics, education level and income. For the firm take-up rate, we calculate the proportion of individuals taking the 500 EUR deductible in an individual's firm, defined at the establishment level, excluding herself. For the location take-up rate, we calculate the proportion of individuals in an individual's 6-digit postcode taking the 500 EUR deductible, excluding herself. For parental deductible choice, we use a variable that is one if a given parent elects the 500 EUR deductible.

The cross-sectional associations between these environmental factors and deductible choice quality are very strong. For example, these regressions find that when the share of colleagues choosing a high deductible in a firm is 10% higher then the probability a given individual chooses an extra deductible is 1.0% lower when predictably healthy, but 3.5% higher when predictably healthy. For location, an increase in the local take-up rate by 10% increases the take-up probability by 7.0% for predictably healthy individuals. For intra-family deductible choices we find that if an individual's father (mother) chose the 500 EUR deductible, and that individual has good predicted health, then that individual is 25.7% (31.0%) more likely to elect the high deductible themselves.

While these cross-sectional correlations are instructive, there is a long literature discussing the reflection problem in analysis of peer effects, where it is easy to confound underlying correlated unobservables for a peer group (see, e.g., Manski (1993)). We now turn to panel analyses that aim to quantify the causal implications of these peer effects for deductible choice quality.

IV.C.1 Co-workers and Neighbors: Movers Design

We use the deductible choices by firm switchers and location movers in a two-part framework to quantify the causal impact of place of work or home on deductible choice quality. Note that this causal impact could be a combination of both (i) peer effects and (ii) firm or location-specific unobserved heterogeneity (e.g. the firm promotes a certain kind of deductible choice).

The first part of our framework obtains individual fixed effects and firm or location fixed effects from an linear OLS framework, similar in spirit to Abowd, Kramarz and Margolis (1999):

Table 7: Switchers Design: Firm and Postcode Effects

	Firms		Postcodes	
	> 100 employees	*> 500 employees*	*> 500 inhabitants*	*> 2000 inhabitants*
Baseline case: split sample	0.135***	0.145***	0.101***	0.151***
	(0.010)	(0.015)	(0.009)	(0.017)
Including switchers	0.208***	0.169***	0.120***	0.166***
	(0.008)	(0.012)	(0.008)	(0.016)
Excluding switchers	0.099***	0.149***	0.057***	0.088***
	(0.008)	(0.014)	(0.009)	(0.018)

Notes: This table displays the results of an AKM-style regression capturing peer effects at the firm and the postcode level. In a first step, firm and postcode fixed effects are obtained from regressing individual take-up of the 500 deductible on household gross income and probability of low costs in deciles, with individual and time fixed effects. In a second step, firm and postcode fixed effects are regressed on the share of take-up in an individual's firm or postcode. The results of this regression are displayed here for different minimum sizes of firms and postcodes, and different identification methods for the fixed effect. In the first row, fixed effects are computed off one randomly selected half of the sample, and the second step regression is computed off the other half. In the second row, both the first and the second step are performed on the entire sample. In the third row, the first step is performed on the entire sample, but the second one excludes firm or postcode switchers. *** p<0.01, ** p<0.05, * p<0.1 with robust standard errors.

$$y_{i,x,t} = \alpha_i + \gamma_t + \theta_x + \beta_1 w_{i,t} + \beta_2 \pi_{i,t} + \epsilon_{i,t}$$

Here, α_i is an individual fixed effect, γ_t is a time period fixed effect and θ_x is a firm fixed effect. $w_{i,t}$ and $\pi_{i,t}$ are the individual household's gross income and health level (in deciles).

In the second step, we regress the obtained fixed effects θ_x on the average share of the of high-deductible take-up in firm or location x over time:

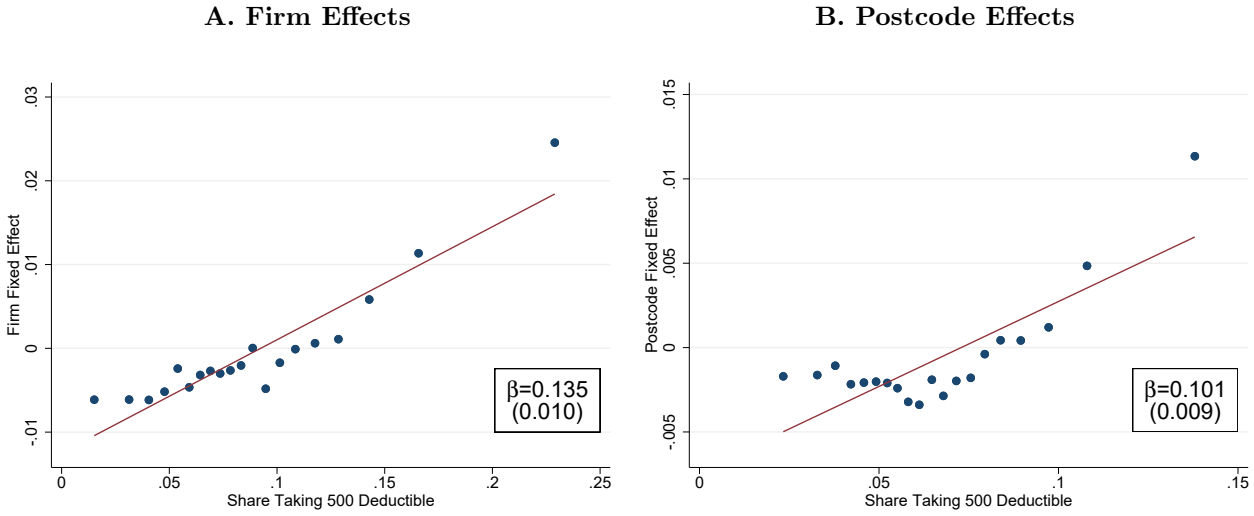
$$\theta_x = \beta \bar{h}_x + \epsilon_x$$

Crucially, there are several possible ways to implement this two-step framework based on how we include switchers in the second step. The argument against including switchers in the second step is that these are the same individuals identifying the fixed effects in step one, so that if there are a lot of switchers at a firm/location regressing the fixed effect on $\bar{h}_x$ becomes closer to regressing a variable on itself. The argument for including switchers is that they are likely more dynamic than others in the firm/location, since they may have re-optimized things about their lives, including deductible choice, more recently. If we exclude these switchers in step two, but they are more influential on the choices of peers, then we may mis-estimate the impact of $\bar{h}_x$ on θ_x.

To deal with this issue, we split the sample in half, and run step one on half the sample and run step two on the other half of the sample, using step one fixed effect estimates on the left-hand side of the step two regression. This approach allows for switchers to be included in both step one and step two without having step two regress fixed effects on observations that directly identified those fixed effects. Our large sample size allows us to have strong statistical power despite only using half the sample in each regression, though we do focus on larger firms to mitigate issues related to estimating noisy fixed effects for smaller firms. While we use this split-sample approach as our primary approach, we also show results for this two step framework where we don't split the sample and we either (i) exclude all switchers from step two or (ii) include all switchers in step two.

Table 7 presents our AKM-style results for firms. We show results for all firms with more than 100 employees and for all firms with greater than 500 employees. The split-sample results show that, when the proportion of individuals in a firm choosing the high deductible is 10% higher, the firm fixed effect is around 1.4% higher.

FIGURE 8: AKM RESULTS: FIRMS AND POSTCODES



Notes: This figure shows the relationship between firm (Panel A) and postcode (Panel B) fixed effects, and the share of take-up of the 500 deductible in the firm or postcode. Fixed effects are obtained from regressing individual take-up of the 500 deductible on household gross income and probability of low costs in deciles, with individual and time fixed effects. The share of take-up is then computed for each individual as the share of colleagues or neighbors who chose the high deductible (i.e., excluding herself), averaged over employees and over the five years in our sample. In Panel A, we include all firms employing 100 people or more; in Panel B, all postcodes with a population of 500 or more.

Panel A of Figure 8 shows the relation between the firm take-up rates and the fixed effects to highlight the strong fit. Thus, there is a meaningful causal effect: someone who switches to a firm is more likely to choose a high deductible if others in the firm are doing so, controlling for health and income. The causal estimate also explains more than half of the cross-sectional correlation shown in Appendix Table C.3.

When we include switchers in step two, and don't use the split-sample approach, these estimated coefficients are higher, implying an effect of 2.1% for firms with more than 100 individuals. Thus, not surprisingly, including the same individuals in step two who we used to identify fixed effects in step one biases our coefficients upwards. Conversely, when we don't use the split-sample approach but exclude switchers from step two, our coefficients are biased downward (1.0% for firms with more than 100 people). This suggests that these switchers may ultimately have more influences on the choices of peers than other, more static, employees at the firm.

We use the same two-step AKM-style approach to investigate the impact of neighbors / postcode on deductible choice. Table 7 presents the results for postcodes with more than 500 individuals and postcodes with more than 2000 individuals. Our primary split-sample approach shows that for a 10% increase in postcode high-deductible take-up, 1.0 % more individuals causally take-up the high deductible in neighborhoods with more than 500 people, and 1.5 % more do so in neighborhoods with more than 2000 people. Interestingly, for neighborhoods, these numbers are very similar when we implement the full sample specification including movers, perhaps because movers are a lower proportion of people in the postcode relative to firms. When movers are excluded, the estimates are much lower (.57%) for postcodes with more than 500 people) suggesting perhaps that the movers do have an out-sized influence on neighborhood peer effects. Figure 8 plots the regression fit for our primary estimates (split-sample), highlighting the strong fit and the differences in results across these two specifications.

While these results are informative about the causal effects of firms and locations on deductible choice, we also want to quantify their impact on deductible choice quality. To shed further light on this, we also re-run our AKM approach on each of two samples: (i) individuals who are predictably healthy (with predicted low cost

probability above 50% - such that the high deductible is the right choice - in all 5 years) and (ii) individuals who are predictably unhealthy (with predicted low cost probability below 50% in all 5 years). The left panels in Figure 9 present the results for the firm fixed effects using our primary split-sample approach. The results are clear: when an individual is predictably healthy, the firm effect is strong and positive, with a 10% increase in the number of healthy people taking up a high-deductible causing a 1.7 % increase in high-deductible take-up for healthy people switching into the firm, holding all else equal. Conversely, an individual who is predictably sick is *less* likely to take up a high deductible if more people in the firm do take up that deductible, though this relationship is relatively flat. The right panels in Figure 9 present the results for the location fixed effects. The results are very similar: when an individual is predictably healthy, the postcode effect is strong and positive, with a 10% increase in the number of healthy people taking up a high-deductible causing a 1.5% increase in high-deductible take-up for healthy people switching into the postcode, holding all else equal. Conversely, an individual who is predictably sick is not more likely to take up a high-deductible. This relationship is now basically flat.

In the same spirit, Panels A and B of Appendix Figure C.4 plot the relationship in the data between predicted health and deductible choice for individuals grouped in quartiles of the firm and location fixed effects. The Figure shows that the difference in take-up rates across firms and locations are again larger for individuals who are predictably healthy, while the take-up rates by individuals who are predictably unhealthy are consistently low. As the dispersion in firm and location take-up rates is relatively small, the overall differences in take-up rates are less pronounced than our earlier results, for example comparing individuals with different education and income in Figure 6.

Taken together, these results suggest that both firm and location effects are strong and positive, but only when an individual is predictably healthy and *should* take up a higher deductible and not when they are predictably sick and should not take up the higher deductible.

IV.C.2 Parents: Event-study Design

For the effects of parents' choices on their children, we cannot use the AKM design, but instead rely on an event-study design to investigate the causal linkage between parents' and adult children's decisions. In particular, we study the deductible choice of adult children when one parent switches from not taking any voluntary deductible to the 500 EUR deductible. We estimate the following specification:

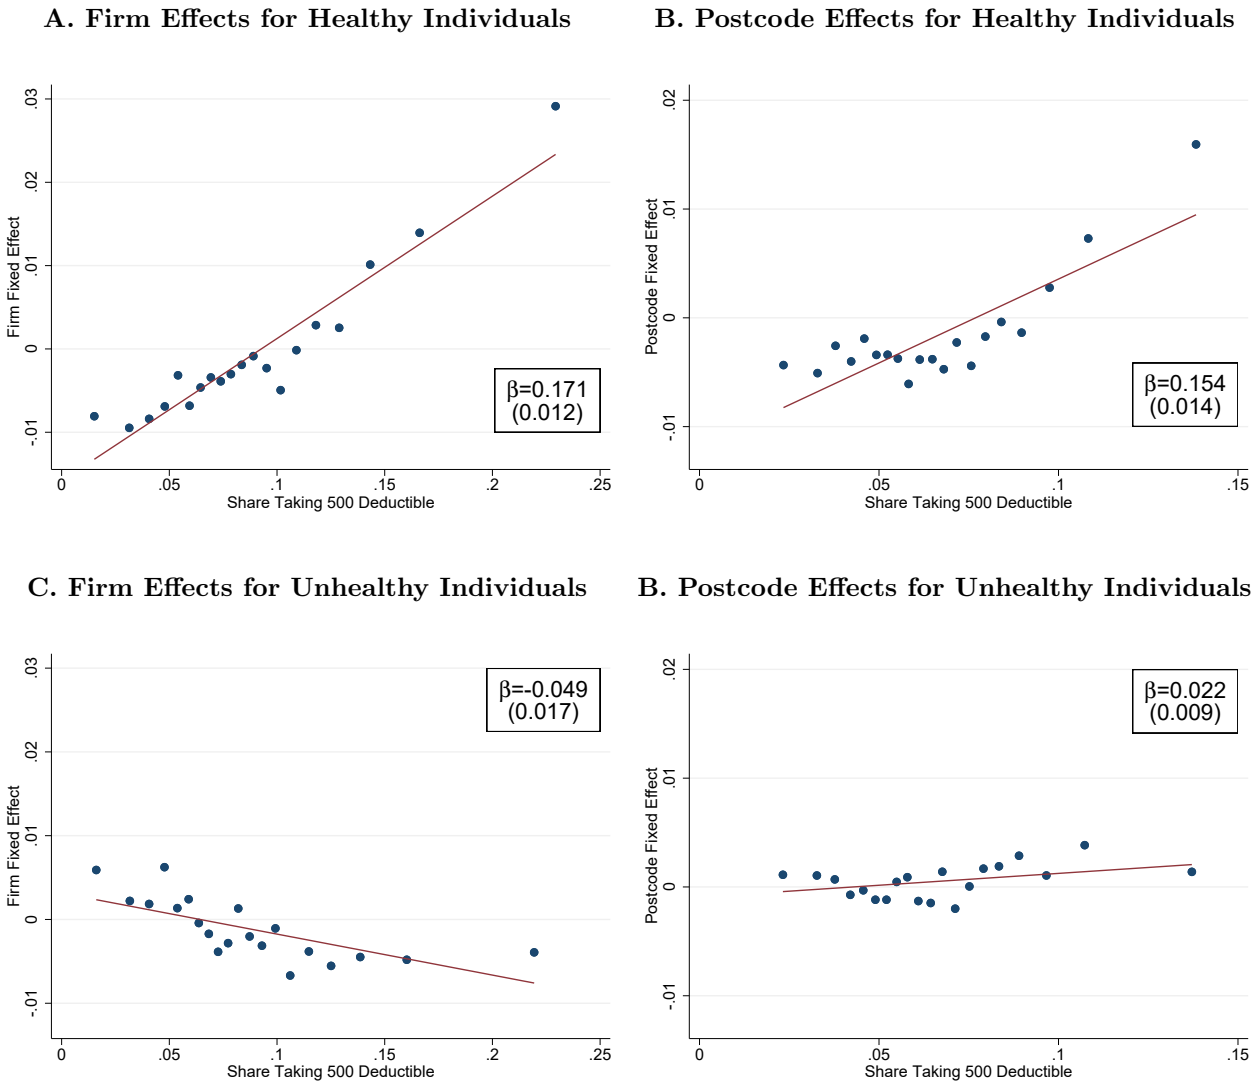
$$d_{it} = \gamma_t + \sum_{j=-N_0}^{N_1} \beta_j \cdot \mathbf{1}[J_{it} = j] + X_{it}\beta + \epsilon_{it}.$$

Here, γ_t is a time fixed effect, $J_{it} = t - E_i$ denotes event time, that is the time in years relative to the moment that the parent switched, and $[-N_0; N_1]$ is the window of dynamic effects around the event. We restrict our sample to "stable" changes, i.e., we exclude individuals whose parent's deductible was not always zero before the switch and is not always 500 after the switch, during the five year window we consider. The causal impact could be a combination of peer effects - either from the parent on the child or vice versa - and some unobserved heterogeneity in the family. In particular, the parents may make the actual deductible decision for their adult children. To mitigate the latter, our main specification presented here excludes families where the parents and adult children are still living together and we report the estimates for children who are younger and older than thirty.

Figure 10 shows the dynamic impact of a father's deductible switch on his children's decisions.[27] The estimates

[27] We focus on one parent here to abstract away from number of parents who switch. Results for mothers are similar to results for fathers.

FIGURE 9: AKM RESULTS: SPLITTING BY HEALTH STATUS

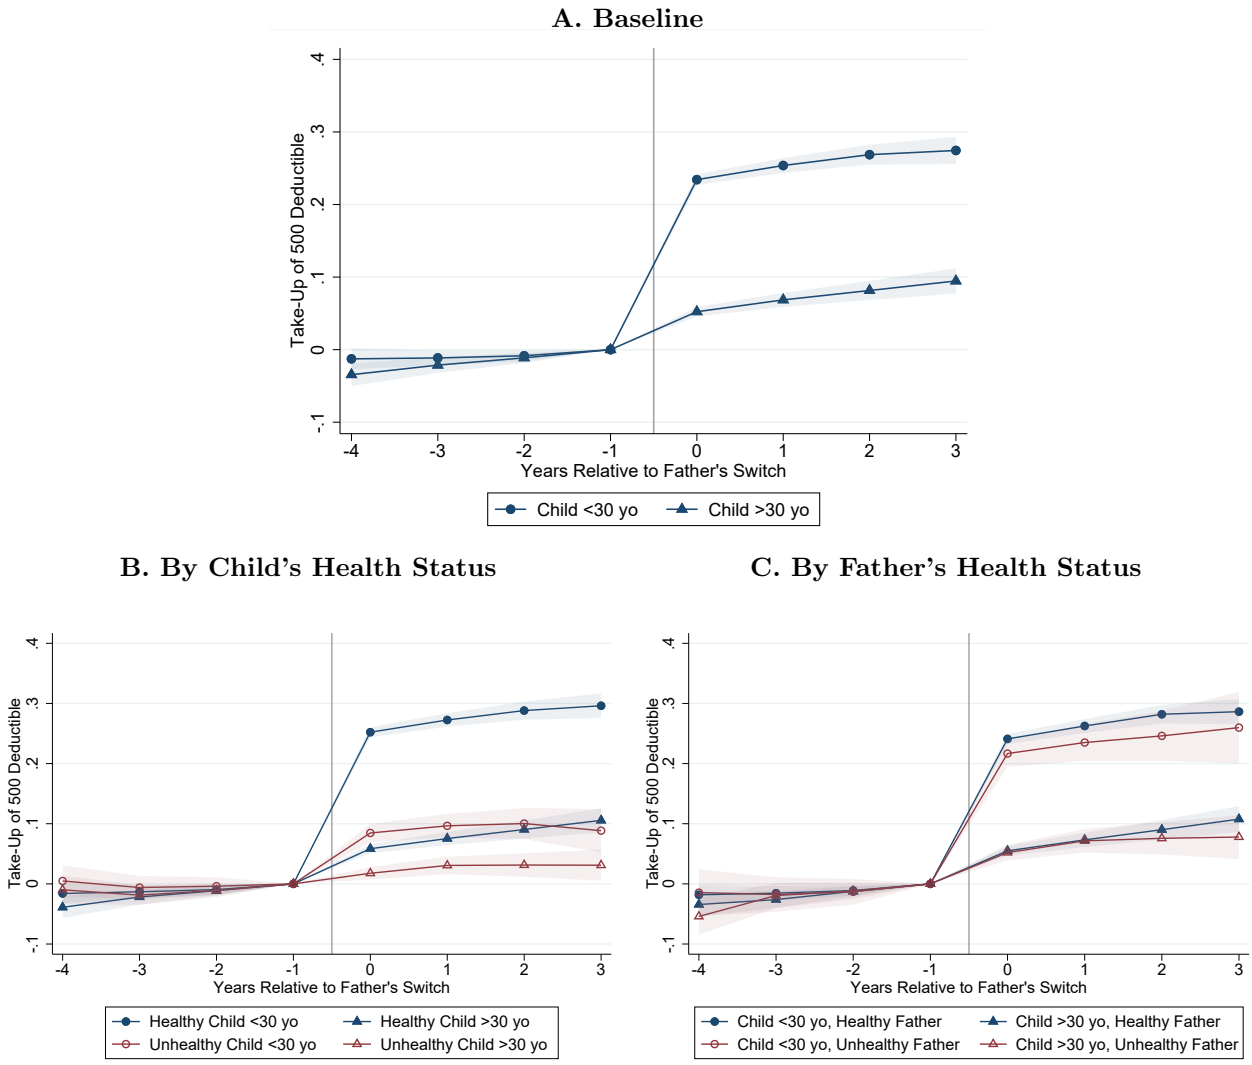


Notes: Notes from Figure 8 apply; but here the relationship between the fixed effects and the share of take-up is plotted separately for individuals who are predictably healthy (i.e., with a probability of low costs greater than .5 in all five years in our sample), in Panel A and B, and predictably unhealthy in Panel C and D (for whom the probability of low costs is below .5 for all five years).

show a clear discontinuous increase in the take-up of the deductible in the year the father switches. Children over 30 are, not surprisingly, less likely to follow their father's lead, though their is still a meaningful effect. The increase is 23 percentage points for children under 30 and 6 percentage points for children above 30. These causal estimates are respectively above and below the cross-sectional estimate of 18 percentage points reported in Table 5. In both cases, there is little anticipation in the take-up rate in the years before and the effect persists in the years after.

We also investigate the heterogeneous event impacts as a function of children's health status and also as a function of parent's health status. We abstract away from health status changes that occur in the five-year window and assign individuals to healthy or unhealthy based on their average predicted health over this time period. The impact is significantly larger for children who are in good health, as shown in Panel B of Figure

FIGURE 10: PARENT EFFECT ON DEDUCTIBLE CHOICE



Notes: The Figure shows the estimates of the dynamic effects using an event-study design of the impact of a parent (here, father) switch from a 0 to a 500 deductible on a child's take-up, excluding all children who still live with their parents. The baseline regression displays the estimates, split between children who are younger or older than 30 years old. The two bottom figures split the impact between predicably healthy/unhealthy children (left) and between predictably healthy/unhealthy fathers (right). Years considered are 2013 to 2017.

10. For the children under 30, there is a 30-40 % higher chance that they also switch to a high-deductible when in good health. This increase is only about 15% when they are in bad health. Panel C of Figure 10 shows the same analysis, but as a function of the father's health status instead of the child's health status. Interestingly, effect heterogeneity as a function of father's health is much lower than heterogeneity as a function of child's health, as children are similarly likely to switch regardless of whether their father took the 'right' decision by switching to the high deductible or not. The overall relation between childrens' predicted health and deductible choice grouped by the take-up of their parents is shown in Panel C of Appendix Figure C.4. Unlike for firms and locations, we cannot rank individuals by the causal effect their parents take-up may have.

Taken together, these results suggest that parental effects are strong and positive, but only when a child is predictably healthy and *should* take up a higher deductible.

V Understanding High and Low Quality Decision Makers

The analysis thus far has covered a variety of observable characteristics that affect choice quality. To provide more insight into how these different features factor into overall decision quality we rely on our model of decision making to assess the quality of choice at the individual level. We then use this measure to form a distribution of decision making quality. Finally, we use this distribution to characterize the observable characteristics most associated with being a high quality decision maker (leaving little surplus on the table) versus a low quality decision maker (those making choices providing little value).

To do so, we first calculate the potential savings for each individual by comparing her deductible choice to the choice that minimizes her expected out-of-pocket expenditures. We follow our stylized model in Section III.A, where an individual minimizes expected payments by choosing the extra 500 EUR deductible at a premium of 250 EUR if her predicted probability of achieving low costs (less than 375 EUR) exceeds 50%. The cost savings simply correspond to the difference in certainty equivalents for risk-neutral preferences:

$$\Delta w_i^{*,\sigma=0} = CE_i^{*,\sigma=0} - CE_i^{\sigma=0}. \tag{4}$$

As shown in Figure 4, allowing for risk aversion makes only small differences to the value of different choices.[28]

V.A Heterogeneity in Choice Quality by Health

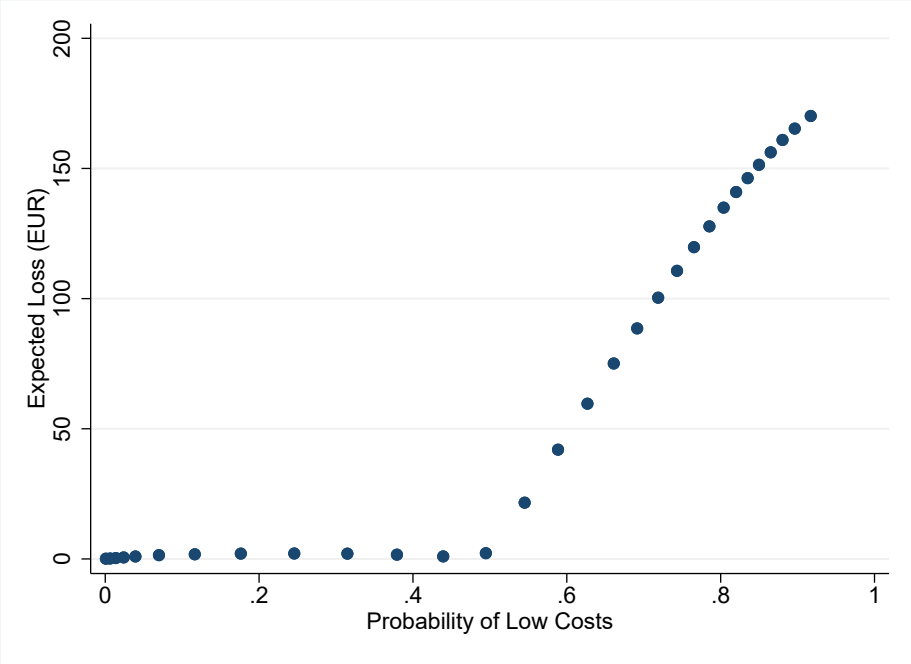
Using the expected cost savings as measure of consumer welfare, we find that approximately 52% of consumers would have been better off with the 500 EUR voluntary deductible in 2015, but less than 7% of consumers took it. Of the population of the Netherlands, only 54.4% of individuals chose the cost-minimizing deductible. The average amount of money left on the table per individual is 66.2 EUR. While small in absolute value, these savings are roughly half of the total surplus at stake in the decision, which is 145 EUR on average.[29] The bin-scatter plot in Figure 11 shows how the average cost savings vary with the predicted probability π. The overall costs savings combine the expected loss from over-insurance for low-cost individuals ($\pi \geq .5$) who do not take the extra deductible and from under-insurance for high-cost individuals ($\pi < .5$) who do. This graph is the result of the combination of the mechanical relationship between π and the potential cost savings, which are V-shaped around $\pi = 0$, as well as the actual distribution of choices made conditional on π. Very few individuals under-insure: most individuals with high predicted risk stick to coverage without extra deductible, as they should. On the other hand, relatively few individuals opt for the deductible when they should and the expected loss from over-insurance increases as the predicted risk is lower.

Figure 11 demonstrates an important feature of our setting. Choice error is strongly correlated with health risk. In particular, those in the best health tend to leave the most money on the table. For the purposes of measuring welfare and equity this correlation will affect our results if an individual's health is also correlated with socio-economic background. Our earlier findings, however, demonstrate that socio-economic factors affect choice directly, even conditional on health, so there is ample opportunity for policy options that are mediated through choice to change welfare.

[28] Note that we over-estimate the cost savings for those who do not take the 500 EUR deductible, but do take an intermediate deductible. However, we under-estimate the gain for those who do not take a voluntary deductible with predicted probability just below 50%.

[29] We define the stake as $|250 - (1 - \hat{\pi})500|$ EUR, which is at most 250 EUR and equal to 0 for individuals with $\pi = .5$.

Figure 11: Expected Loss and Health Cost Probability



Notes: This figure is a binned scatterplot of the relationship between the predicted probability of health costs below 375 EUR and the expected loss due to over- or under-insurance. For individuals with a predicted probability of low costs below 0.5, the expected losses due to under-insurance are very small (on average close to zero), as a very low fraction of people under-insures by taking the 500 EUR extra deductible. For individuals with a predicted probability of low costs above 0.5, expected losses due to over-insurance increase with this probability, and reach almost 170 EUR for people with a very high chance (0.9+) of low costs, as most people leave money on the table by over-insuring for costs that happen with a very low probability.

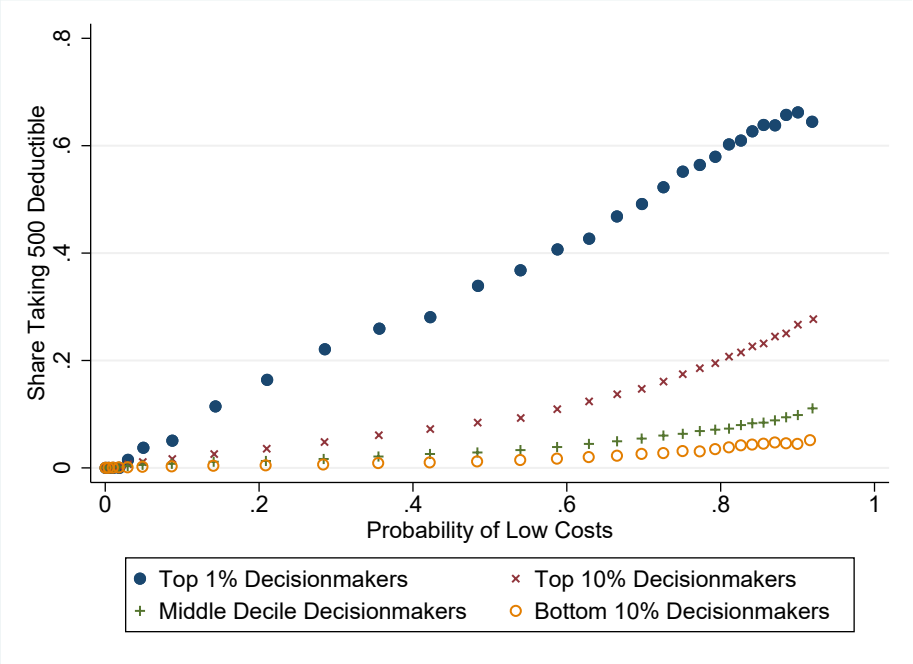
V.B Heterogeneity in Choice Quality conditional on Health

We now revisit our earlier empirical analysis of deductible choice to assess the overall distribution of choice quality, conditional on health, and which kinds of consumers are the best and worst choosers. To control for differences in health, we predict consumers' choices as a function of their underlying health risk π_{it} and observable characteristics X_{it}, allowing for interactions between the two, in a first step. We thus get predicted deductible choice probabilities $d(X_{it}, \pi_{it})$, which we then translate into consumer welfare $\Delta w^{*,\sigma=0}(X_{it}, \pi_{it})$ based on equation 4 in a second step. In a final step, we average the cost savings over the different health risks using the population distribution of predicted health risks, $\Delta w^{*,\sigma=0}_{\pi_{pop}}(X_{it})$. We provide more detail on this procedure in Appendix D.1.

We find significant heterogeneity in choice quality, even when controlling for differences in health risk. We rank individuals from worst to best decision makers based on how much value they are predicted to leave on the table on average across a representative distribution of population health. The very best decision makers (the top .1%) choose the cost-minimizing deductible 73% of times, conditional on some health risk drawn from the population distribution (see Panel A of Appendix Figure D.2). The top 5% decision makers have a probability of 55% to make the right choice. All other decision makers are predicted to make worse choices than an individual choosing randomly.

Figure 12 shows the responsiveness of deductible choices to health risk for different quantiles of choice quality. The performance of the very best decision makers is striking relative to the others. The take-up rate of the top 1% of decision makers is much steeper, coming close to the 45-degree line. The median quality decision-maker, on the other hand, essentially sticks to the compulsory deductible regardless of the underlying health risk.

Figure 12: Heterogeneity in Choice Quality



Notes: This figure illustrates dispersion in choice quality, by showing a binned scatter plot of the relationship between the predicted probability of having costs below 375 EUR (staying under the voluntary deductible range) and the take-up of the voluntary 500 EUR deductible for four selected subgroups that differ in their expected loss. The bottom 1% expected loss group comes close to a rational consumer, with high take-up of the deductible for low expected costs. The top 10% expected loss group has losses that are due almost entirely to over-insurance.

Table 8 shows how observable characteristics change with good and bad choices. The entries in the left panel give the average values in each group of choice quality. The entries in the right panel give the ratio of the proportion of consumers with that characteristic in each group relative to the proportion of consumers with that characteristic in the population overall. Within each panel, the left-hand column describes the 5% of best decision-makers while the right-hand column describes the 5% of worst decision-makers.

The results paint a telling picture of who is making the best choices in our context. The best decision-makers have an average gross income of 105K EUR and net worth of about 250K EUR. The worst decision makers, though, only have an average income of 40k EUR and net worth of 5K EUR. The massive difference in income and wealth are complemented with substantial differences in education. Those with college education are 3.48 times more likely to be in the best decision making group and with further education are 15.57 more likely. In terms of education field, statisticians (19.66 times), philosophers (13.14 times), and economists (6.95 times) are much more likely to be present in the group of top decision-makers while those trained in hair and beauty services (0.64 times) and protection of persons (0.38 times) are much less likely to best present. In addition, those in business services (2.77 times) and insurance (2.13 times) are more likely to be in the group of best decision-makers while those in cleaning (0.26 times) and those not working (0.32 times) are less likely to be in this group.[30]

We also see that high quality decision makers are in peer settings where decision making quality is higher, both in terms of where they work and where they live. The average firm and postcode fixed effects decile for the top 5% decision makers is 6.41 and 6.07 respectively. The differences in parental take-up across the different

[30]Note that a zero value in the right panel of Table 8 does not mean that no single individual with the respective characteristic can be in that group. Instead, it means that given the predicted choices based on observable characteristics for all individuals, no individual with that specific characteristic (and his/her other respective characteristics) is predicted to end up in that group.

Table 8: Best and Worst Decision Makers

	Mean			Over/underrepresentation	
	Top 5% decisionmakers	*Bottom 5% decisionmakers*		*Top 5% decisionmakers*	*Bottom 5% decisionmakers*
Demographics			**Education level**		
Gender (male)	62%	28%	Less than high school	0.30	2.99
Age	36	63	High school	0.82	0.33
Has children	59%	34%	College	3.48	0.00
Has a partner	46%	90%	Further Studies	15.57	0.00
Financials			Unknown	0.08	1.05
Gross income	105,801	39,347	**Education field**		
Net worth	250,632	4,969	Statistics	19.66	0.00
Has Mortgage Debt	64%	19%	Philosophy	13.14	0.00
Has Other Debt	27%	53%	Economics	6.95	0.01
Has Savings >2000EUR	91%	38%	Tax and administration	3.30	0.01
Peer Effects			Marketing and advertising	1.91	0.06
Firm FE decile	6.41	4.09	Hair and beauty services	0.64	1.79
Postcode FE decile	6.07	5.47	Protection of persons	0.38	2.24
Mother With 500 Deductible	37%	0%	**Work Status**		
Father With 500 Deductible	45%	0%	Student	2.80	0.16
			Retired	0.07	2.47
			Self-employed	2.07	0.05
			Employee	1.16	0.31
			On Benefits	0.32	1.94
			Professional sector		
			Business services	2.77	0.09
			Insurance	2.13	0.07
			Retail	1.10	0.34
			Construction	0.75	0.24
			Cleaning	0.26	1.40
			Public utilities	1.51	0.11
Observations					11,369,800

Notes: This table presents observable characteristics for the groups that our model considers to be the top 5% and the bottom 5% decision makers. The entries in the left panel give the average value of the variable in each group. The entries in the right panel give the ratio of the proportion of consumers with that characteristic in each group relative to the proportion of consumers with that characteristic in the population overall. For example, the group of best decision makers has 6.95 time more economics majors, proportionally, than the population overall.

groups is particularly striking. For those in the top group of decision makers, the average take-up rate among father and mothers is 45% and 37% respectively, while it is zero for those in the bottom group. Finally, we also find that better decision-makers are significantly younger on average (36 year old vs. 63 in bottom 5 %), more likely to be male and more likely to have children.

This section has painted a stark picture of the differences in choice quality by socio-demographic characteristics. Most starkly, poorer individuals, individuals with lower education levels, and individuals in less quantitative fields are far more likely to make worse decisions. While we study only one specific setting, this accords with the results on 401(k) choices in Chetty et al. (2014), suggesting a broader pattern of how choice quality links to socio-demographic characteristics and, ultimately, inequality in outcomes under choice-based policies.

VI Policies and Welfare

In this section we study the efficiency and equity implications of alternative policies such as (i) using predicted risk to default enrollees into a plan and (ii) limiting the choices available. Accounting for how the incidence of choice frictions falls on individuals with different observable characteristics, we explore not only the efficiency, but also the equity implications of the different policy options.

Before we begin we note some important assumptions about this exercise. We focus our policy analysis on consumer welfare without accounting for the potential implications of moral hazard and adverse selection. In the presence of moral hazard, the reduction in health expenditures in response to an extra deductible could benefit the insurer. From a social welfare perspective, however, the evidence suggests that higher deductibles are relatively inefficient at reducing low value care at the margin (e.g., Brot-Goldberg et al. (2017)) suggesting a limited role for moral hazard in optimal policy. In the presence of adverse selection, we also expect equilibrium prices to respond to the regulation of choice, which would further affect sorting and consumer welfare. In particular, the option to buy less comprehensive coverage allows individuals with good health to contribute less to the health insurance system. We ignore the pricing repercussions this may have.[31] Finally, we also ignore any direct welfare effects of choice frictions beyond the misallocation to plans (e.g., search or switching costs), but we complement the welfare analysis by characterizing some micro-foundations that can rationalize the observed choices.

VI.A Counterfactual Policies

We study the impacts of three counterfactual choice policies to assess the potential to improve consumer welfare in our context. First, we consider how much better off consumers would be if everyone were allocated to the best option for them *ex ante* (according to our estimates and welfare model). This is useful as a first-best benchmark given the current choice architecture. It is also a measure of the impact of policy interventions that improve consumer decision-making or use predictive models to establish "smart defaults" (Handel and Kolstad (2015*a*), Gruber et al. (2020)). Next, we consider the impact of two alternative policies that limit choice; one that offers only the high deductible option and one where only the low deductible option is offered. These policies are clearly feasible and also reflect the underlying trade-off between offering greater choice and exacerbating choice errors.

We study both the efficiency and equity implications of these alternative policies. In assessing the efficiency implications for each policy — the surplus generated by the plans chosen — we allow for four different values of risk aversion (assumed to be homogeneous in each implementation) including (i) risk neutral (ii) CARA $= 10^{-5}$ (iii) CARA $= 10^{-4}$ and (iv) CARA $= 10^{-3}$. To assess the equity implications we rely on income as the measure of inequality and consider alterternative welfare weights for deciles of the income distribution. Following Atkinson (1970), the welfare of an individual in income decile y_δ is weighted by $y_\delta^{-\epsilon}/\left(\Sigma y_\delta^{-\epsilon}/10\right)$ for $\epsilon = .5$ and $\epsilon = 1.5$.[32] In our primary analysis, we rely on the observed correlations between health and socio-demographic status in the data. In the appendix, we also perform an analysis that assumes identical health distributions conditional on non-health X_{it}, using the predicted choice probabilities $d(X_{it}, \pi_{it})$ as in subsection V.B.

Table 9 presents the average welfare impact per person (in EUR) for the three different policies we consider.

[31] By removing choice frictions, we may expect adverse selection to become worse (e.g., Handel, Kolstad and Spinnewijn (2019)). Interestingly, comparing the average predicted low-cost probability for workers taking the extra deductible (.763) and for those who should take the extra deductible (.760) suggests that in this context the pricing repercussions from reducing choice frictions may be limited.

[32] The Atkinson index of inequality uses a social welfare function of the form $y_i^{1-\epsilon}$ with $\epsilon \geq 0$ a measure of inequality aversion. Here, we weigh the welfare gain for each individual depending on income decile they are in by $y_\delta^{-\epsilon}/\left(\Sigma y_\delta^{-\epsilon})/10\right)$, which ensures comparability with the unweighted case. We could model equity concerns more broadly by differentially weighting outcomes for individuals as a function of their predicted health π_i and characteristics X_i.

Table 9: Welfare Impact of Alternative Policies

	Optimal Deductible	High Deductible Only (875 EUR)	Low Deductible Only (375 EUR)
Risk Neutral			
Unweighted	68.8	-26.2	-8.3
Low Inequality Aversion	56.9	-64.4	-6.3
High Inequality Aversion	37.4	-133.6	-3.4
σ=.0001			
Unweighted	67.8	-28.1	-8.2
Low Inequality Aversion	56.0	-66.1	-6.2
High Inequality Aversion	36.8	-135.1	-3.3
σ=.001			
Unweighted	58.0	-44.6	-7.0
Low Inequality Aversion	47.7	-81.6	-5.3
High Inequality Aversion	30.9	-148.7	-2.7

Notes: This table shows the average welfare impact (in EUR per person) of three alternative policies concerning the extra deductible: optimal deductible (all individuals taking the optimal deductible given their health risk), high deductible only (only the 500 EUR extra deductible is available), and low deductible only (the low deductible is the only option). The welfare impact is calculated with equal weights for all income deciles, low inequality aversion or high inequality aversion. Weights y_δ are computed as $y_\delta^{-\epsilon} / \left(\Sigma y_\delta^{-\epsilon}/10\right)$ for $\epsilon = .5$ and $\epsilon = 1.5$. The welfare impact is calculated not controlling for health. The corresponding welfare impact when assigning each individual the population health distribution is in Appendix Table D.1. Our sample contains the choices of $9,415,666$ individuals in 2015 (out of $11,991,629$ individuals for which the probability of low costs and the deductible choice are both non missing), excluding students, self-employed people, individuals with a gross income below the social assistance level and individuals with missing observables.

Column 1 presents the results for the scenario where individuals are allocated to their *ex ante* optimal deductible choice in the current environment. The average consumer welfare gain, when not weighted for inequality, is 68.8 EUR for risk neutral individuals. This gain decreases only slightly when introducing reasonable levels of risk aversion and is still 58 EUR for individuals assuming our highest level of risk aversion. When we weight for equity as a function of income the gain of *ex ante* optimal allocation is reduced. With high inequality aversion the average benefit of this policy is 37.4 EUR for a risk neutral consumer. The decline results from the fact that lower income individuals are less likely to be healthy and, thus, more likely to have the default option of a low deductible be the correct choice for them. Because most choice errors result from not actively choosing the higher deductible, there is less to be gained if many low income enrollees are better off in the low deductible plan. Appendix Table D.1 shows how this relationship is reversed when controlling for differences in health, reflecting the higher incidence of choice frictions among low-income individuals.

Columns 2 and 3 show the consumer welfare impacts when consumers are offered only the high deductible (with the corresponding premium reduction) or the low deductible, respectively. Neither policy that limits the choice offerings is welfare-increasing, even relative to the status quo where consumers are making poor choices in general. Mandating the extra 500 EUR deductible leads to a welfare losses with no inequality aversion of 22.6 EUR when risk neutral and 44.6 EUR with high risk aversion. With high inequality aversion, however, this policy is much worse, with welfare losses of 133.6 EUR when risk neutral and 148.7 EUR with high risk aversion. This policy is especially bad because it is forcing sick, lower income consumers into what would have been the wrong choice for them. Mandating a low deductible, on the other hand, has a much smaller impact due to the fact that, in practice, most people already choose that deductible. The small impact ranges between 0 and 10

EUR on average across the range of risk aversion and inequality aversion parameters we investigate.

Our counterfactual analysis allows to draw some important conclusions for choice-based policies more generally and for the specific implementation in the Netherlands, using a low baseline deductible with the option to take a higher deductible. While a policy that is able to move people to plans based on *ex ante* risk could substantially increase welfare, the welfare gain from the offered deductible choice is small. Moreover, due to both the correlation between income and health and the correlation between income and choice quality, accounting for higher inequality aversion actually reduces the welfare loss of this policy. The option to select a higher deductible increases welfare mostly for the high-income individuals, who are healthier and make better choices. The value of this option is very limited for low-income individuals and may well become negative when factoring in equilibrium price changes.

Importantly, in our setting we do not have good measures of potential costs associated with decision making. If making a decision imposes a cost on enrollees — as has been shown in a number of other settings (see e.g. Handel and Kolstad (2015*b*) in health insurance) — these costs may exceed the relatively small gains we find from offering the option to take a higher deductible.

VI.B Structural Choice Foundations

While it is not the focus of this paper to test different decision-making models, it is useful to assess what kinds of micro-foundations can rationalize the decision-making patterns that we document as this allows for a further refinement of the welfare analysis and policy recommendations.

To shed some light on this, we simulate choice patterns under a range of distinct micro-foundations and compare the predictions of those simulations to our observed data. We study a number of potential models of decision making that are proposed in the literature, including switching costs, loss aversion, imperfect information, rational inattention and mistakes.

Switching costs occur when consumers with a default plan option must pay some cost c_s to switch plans. This could be, e.g., a paperwork / transaction cost or reflect some reduced form of a multi-stage model with search and search costs.[33] In our setup, agents with one deductible level face this cost c_s when switching to the other deductible level. Specifically, if the consumer had originally chosen the low deductible, she now chooses the high deductible if:

$$250 - (1 - \pi)500 - c_s > 0 \tag{5}$$

This assumes the model premium reduction of 250 for the 500 EUR deductible. We consider heterogeneous population switching costs $c_s \sim U(0, 2 \times \bar{c}_s)$ for different average switching costs $\bar{c}_s$.

Loss aversion occurs when losses loom larger than gains. In contrast with standard risk aversion, loss aversion can reduce the take-up of a deductible even when financial stakes are small (e.g., Sydnor (2010)). Following Kőszegi and Rabin (2007), we assume that realized payoffs are evaluated relative to expected payoffs and losses receive a relative weight λ. In our setup, agents will choose the high deductible if:

$$250 - (1 - \pi)500 - (\lambda - 1)\pi(1 - \pi)500 > 0. \tag{6}$$

Decisions could be made based on imperfect information. In our model, imperfect information enters by allowing consumers to receive a noisy signal $\hat{\pi}$ about their health, where $\hat{\pi} = \pi + \epsilon$ and $\epsilon \sim N(0, \sigma_\epsilon)$. They make

[33] See a discussion of potential inputs into switching costs in Handel (2013).

a decision based on that noisy signal and choose the high deductible (for the model premium reduction of 250) if and only if

$$250 - (1 - \hat{\pi})500 > 0. \tag{7}$$

where the signal-to-noise ratio equals $\sigma_\pi/\sigma_\epsilon$.

Alternatively, individuals may decide rationally whether to pay attention and acquire information. In our model, rational inattention means that consumers, again, receive a noisy signal about their health, but then decide whether or not to pay a cost c_r to learn the true value of his/her health risk. Upon receiving the signal, agents face an expected choice value that integrates over the probability distribution of their potential true health statuses.[34] The result of our rational inattention setup is that, if a consumer starts with the low deductible, they will choose the high deductible if and only if one of the following conditions holds:

$$\hat{\pi} > 0.5 \text{ and } \int_0^{0.5} [-250 + (1 - \pi)500] f(\pi|\hat{\pi}) \, \mathrm{d}\pi < c_r \tag{8}$$

$$\hat{\pi} > 0.5 \text{ and } \int_0^{0.5} [-250 + (1 - \pi)500] f(\pi|\hat{\pi}) \, \mathrm{d}\pi > c_r \text{ and } \pi > 0.5 \tag{9}$$

$$\hat{\pi} \leq 0.5 \text{ and } \int_{0.5}^{1} [250 - (1 - \pi)500] f(\pi|\hat{\pi}) \, \mathrm{d}\pi > c_r \text{ and } \pi > 0.5 \tag{10}$$

The first condition results when consumers are so confident they are low risk that they don't find it worthwhile to pay the cost of precisely determining their health status, instead just electing to choose the high deductible right away. The second and third conditions occur when consumers decide to pay the cost to obtain a more precise signal, and are differentiated only by whether the initial signal value is bigger or smaller than the risk-neutral threshold of $\pi = 0.5$ for high deductible choice under the modal premium reduction.

Finally, consumers may simply make mistakes. In our model, we assume a share $1-\alpha$ of agents make rational, frictionless choices, while share α of agents make random choices.

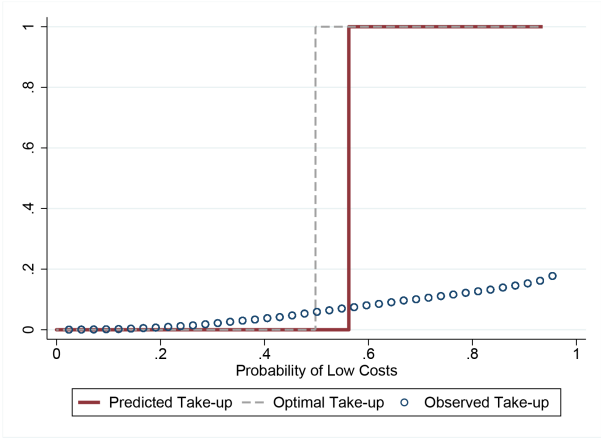
Figure 13 presents simulations of the deductible take-up rate as a function of health risk for the alternative decision models. For comparison, each panel plots the observed take-up rates and the deductible choice for the case where consumers are rational, frictionless, and risk-neutral, as in Figure 5. As discussed before, in a frictionless world, all consumers below a 50% probability of clearing the low deductible will elect the high-deductible, which looks starkly different from the observed low take-up rates. Risk-aversion only slightly alters this threshold, moving it to a marginally higher probability of low spending for the case where consumers are risk-averse with CARA coefficient of $1 * 10^{-4}$ (Panel A). Loss-aversion helps to reduce the take-up rate of individuals around the 50% threshold, but the simulated take-up rates remain too high for a loss-aversion parameter of $lambda = 2.25$ (Sydnor (2010)). Moreover, even with such strong loss aversion, individuals in very good health are predicted to always take up the deductible as the variance in financial payoffs they would get exposed to converges to zero.

We then turn to the simulations for a decision models with switching costs (panels C and D). With a homogeneous switching cost of 119 EUR, about 10 percent of the population would take up the high deductible, which corresponds to the observed take up rate. With heterogeneous switching costs uniformly distributed around the same mean of 119 EUR, we still predict meaningfully more high deductible purchases than we observe in the data, especially as consumers become predictably healthier and healthier. Heterogeneous switching costs with a higher mean of 650 EUR (panel C) look much more similar to observed purchases as a function of health status.

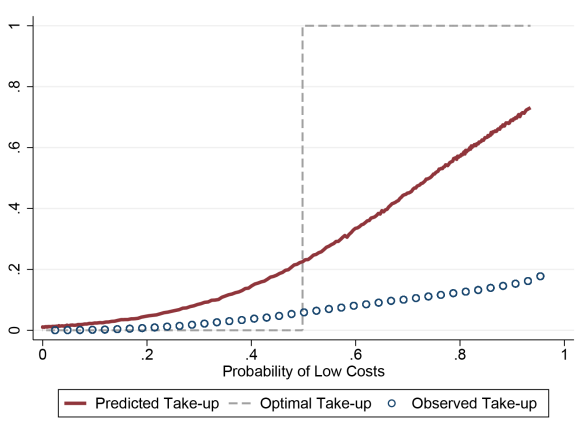
[34] Section D.2 presents the model in detail. Our model is similar in spirit to that laid out in Ho, Hogan and Scott Morton (2017), though there consumers obtain signals about plan characteristics while here they about signals about their own health status. We could recast our model as related to uncertainty about plan characteristics, likely with similar results.

FIGURE 13: DEDUCTIBLE TAKE-UP FOR DIFFERENT BEHAVIORAL MODELS

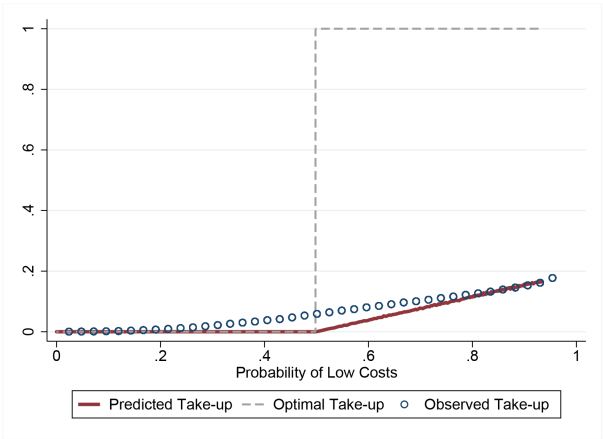
A. Optimal Choice



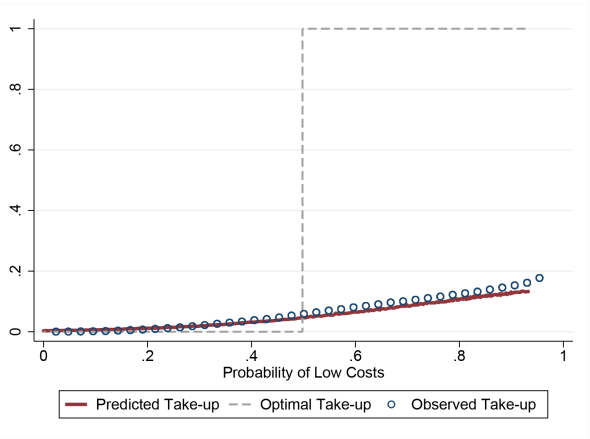
B. Loss Aversion



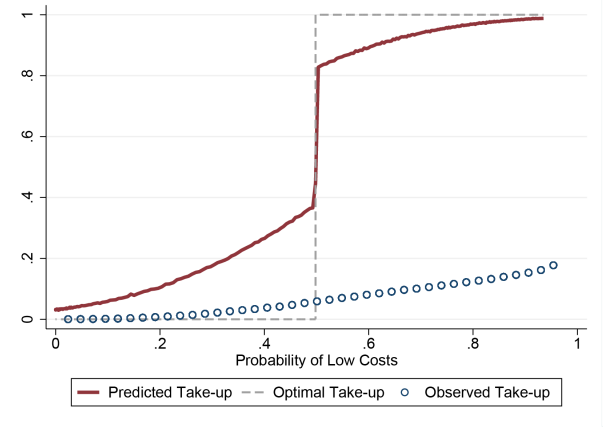
C. Heterogeneous Switching Costs



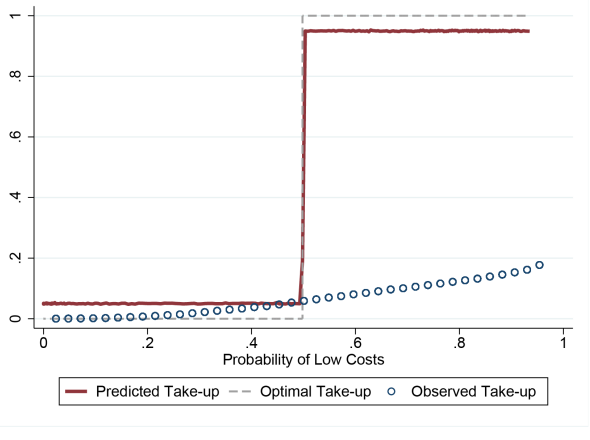
D. Het. Switching Costs and Imperfect Info



E. Rational Inattention



F. Mistakes



Notes: This figure presents the results from decision-making simulations for the various models discussed in detail in the text. For each model, we contrast the predicted take-up rate with both the observed take-up rate and the take-up rate by a rational consumer in a frictionless world.

But this specification still predicts no purchasing of a high deductible for consumers with higher predicted probabilities of higher health spending. However, when we combine our model of high switching costs with our model of imperfect information about health status (with an assumed signal-to-noise ratio of 1), the simulated choices as a function of health status map very closely to observed choices (panel D).

Figure 13 also presents results for the rational inattention model (panel E) and the random mistakes model (panel F). The simulations for the rational inattention model use an information acquisition cost of $c_r = 25$ (for much higher values, no one pays this cost to learn about their true health status, making the model's predictions the same as the imperfect information model). We see that the take-up rate becomes more responsive to health risk around the threshold value, since individuals have to have probabilistic signals close to the marginal thresholds to acquire information, even with a reasonably small cost of 25 EUR. Furthermore, consumers with larger probabilities of being healthy are predicted to purchase the higher deductible much more than they actually do in practice. So we would need to combine the model of rational inattention with high switching costs to obtain predictions that are closer to observed choices. The simulations for the random mistakes model assume that a random 10% of consumers make mistakes. Clearly, the overall take-up rate is too high, so we again need an extra force to lower the take-up rate. Moreover, in the random mistakes model, the take-up rate is now also too high for individuals who are predicted to have high costs. This would not be resolved by combining the mistakes model with the imperfect information model.

This section illustrates how simulations based on different choice models compare with our data. Though there are a plethora of models one could write down that could help rationalizing the data (e.g., inertia, limited attention), a model of high switching costs combined with imperfect information fits the data very well. Importantly, high switching costs would further decrease the welfare gains from offering deductible choice. While we don't structurally estimate these models in our current context, these simulations give a sense of what models might make sense to estimate, and potentially test formally vs. one another, to implement a more detailed investigation of the mechanisms underlying the choice patterns we have documented.

VII Conclusion

Many policy makers rely on market-based solutions to supply products, from health insurance managed competition to private retirement benefits and beyond. The rationale for these approaches is that regulated, market-based provision of impure public goods can deliver greater product variety and improved efficiency; In essence, getting the returns we expect from a market while still accounting for the public nature of these goods and services.

An important limitation to the effectiveness of these policies is the ability of consumers to choose between the available options and maximize their surplus. Ineffective decision-making undermines the gains from such policy approaches and can, in principle, be large enough to entirely undercut reliance on market-based provision for these products.

Using granular data from the Netherlands, we characterized nationwide quality in deductible choices and found that (i) these choices were poor on average and (ii) higher SES consumers make better choices than lower SES consumers, with a meaningful impact on realized surplus. Most notably, highly educated individuals who have more quantitative training make better choices than their counterparts. We also find that the peers one has at work and in the neighborhood and family members — a measure of the information or social networks — are important determinants of choice quality. A variety of other socio-economic factors have more limited impact on choice quality, including household income and household liquidity, once we control for educational background.

We bring these estimates together with a model of consumer surplus and study what factors are most predictive

of the best and worst decision-makers. Higher socioeconomic status, in particular human capital associated with higher income, is highly predictive of being in the group of best-decision makers, and vice-versa for the group of worst decision-makers. Similarly, being a part of peer groups who make good decisions is predictive of being in the group of best decision-makers, and vice-versa.

We also use our framework and estimates to study the consumer welfare impacts of several counterfactual choice policies. We show that if consumers could be allocated on an *ex ante* basis to the best deductible options for them (e.g. providing a smart default), they would capture roughly 40% more of the total surplus at stake in the current deductible choice environment. Limiting the offering to only a high deductible makes consumers meaningfully worse off on average and worse yet once we account inequality, since poor health and income are negatively correlated. Limiting the offering to only the low-deductible option also lowers welfare, though to a far smaller degree on average and even less when accounting for inequality. Insofar as establishing the market is costly (an issue in general but particularly in a selection market such as insurance) or individuals find choice making costly, those costs may exceed these small welfare losses from offering only one product.

Given the policy importance of our results, both for choice quality overall and for the choice quality - SES gradient, we believe that there are several fruitful directions for future research. At a micro level, it will be valuable to assess how different policy and technology solutions can improve choices in different market and regulatory environments, both overall and for lower SES consumers specifically. For example, a field experiment distinguishing between distinct behavioral foundations and/or distinct behaviorally-motivated policies could provide valuable additional insights, especially if linked to data similar to what we use in this study. Finally, it is important to explore policy options that account not only for low quality decision making on average but the distributional consequences of such decision-making issues. For example, one could design the choice menu to combat the regressive nature of choice quality by matching the default option closer to the typical low SES consumer than to the typical high SES consumer. Targeted defaults, as discussed in Handel and Kolstad (2015*a*) and Abaluck and Gruber (2016*a*), are another interesting path forward from a policy design standpoint. While the evidence for the importance of choice frictions and their unequal incidence in the population seems strong in our context, it will be valuable in future work to more generally characterize and measure the trade-offs between potentially regressive mistakes and efficiency gains due to choice-based policies.

References

Abaluck, Jason, and Abigail Adams. 2019. "What Do Consumers Consider Before They Choose? Identification from Asymmetric Demand Responses." Yale Working Paper. 4

Abaluck, Jason, and Jonathan Gruber. 2016*a*. "Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance." *American Economic Review*, 106(8): 2145–84. 5, 37

Abaluck, Jason, and Jonathan Gruber. 2016*b*. "Improving the Quality of Choices in Health Insurance Markets." NBER Working Paper No. 22917. 5

Abowd, John, Francis Kramarz, and David Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica*, 67(2): 251–333. 3, 21

Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky. 2019. "Regressive sin taxes, with an application to the optimal soda tax." *The Quarterly Journal of Economics*, 134(3): 1557–1626. 5

Atkinson, Anthony. 1970. "On the Measurement of Inequality." *Journal of Economic Theory*, 2: 244–263. 4, 31

Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein. 2015. "Behavioral Hazard in Health Insurance." *Quarterly Journal of Economics*, 130(4): 1623–1667. 5

Barseghyan, Levon, Francesca Molinari, Ted O'Donoghue, and Joshua Teitelbaum. 2013. "The Nature of Risk Preferences: Evidence from Insurance Choices." *American Economic Review*, 103(6): 2499–2529. 5

Beshears, John, James Choi, David Laibson, Brigitte Madrian, and Sean Wang. 2016. "Who is Easier to Nudge?" Harvard University Working Paper. 5

Bhargava, Saurabh, George Loewenstein, and Justin Sydnor. 2017. "Choose to Lose: Health Plan Choices from a Menu with Dominated Option." *The Quarterly Journal of Economics*, 132(3): 1319–1372. 2, 5

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience theory of choice under risk." *The Quarterly journal of economics*, 127(3): 1243–1285. 2

Brot-Goldberg, Zarek C, Amitabh Chandra, Benjamin R Handel, and Jonathan T Kolstad. 2017. "What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics." *The Quarterly Journal of Economics*, 132(3): 1261–1318. 3, 5, 31

Campbell, John Y. 2016. "Restoring Rational Choice: The Challenge of Consumer Financial Regulation." *American Economic Review*, 106(5): 1–30. 1

Cardon, James, and Igal Hendel. 2001. "Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey." *The RAND Journal of Economics*, 32(3): 408–427. 5

Chen, Yiqun, Petra Persson, and Maria Polyakova. 2019. "The Roots of Health Inequality and The Value of Intra-Family Expertise." National Bureau of Economic Research Working Paper 25618. 5

Chetty, Raj, John N. Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen. 2014. "Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark." *The Quarterly Journal of Economics*, 129(3): 1141–1219. 1, 2, 5, 30

Cohen, Alma, and Liran Einav. 2007. "Estimating Risk Preferences from Deductible Choice." *The American Economic Review*, 97(3): 745–788. 5, 11

Cronqvist, Henrik, and Richard Thaler. 2004. "Design Choices in Privatized Social Security Systems: Learning from the Swedish Experience." *American Economic Review*, 94(2): 424–428. 5

Douven, Rudy, Annette Zeilstra, Anne-Fleur Roos, and Annemiek Verrips. 2019. "Verbeteropties Gemeentepolis." CPB Netherlands Bureau for Economic Policy Analysis CPB Policy Brief. 7

Einav, Liran, Amy Finkelstein, and Paul Schrimpf. 2015. "The Response of Drug Expenditure to Nonlinear Contract Design: Evidence from Medicare Part D." *Quarterly Journal of Economics*, 130(2): 841–899. 5

Einav, Liran, Amy Finkelstein, Sendhil Mullainathan, and Ziad Obermeyer. 2018. "Predictive Modeling of U.S. Health Spending in Late Life." *Science*, 360(6396): 1462–1465. 12, 43

Enthoven, Alain, Alan Garber, and Sara Singer. 2001. "Near-Universal Coverage Through Health Plan Competition: An Insurance Exchange Approach." *Covering America: Real Remedies for the Uninsured, Editors J. Meyer and E. Wicks*, Washington DC: 155–172. 1

Ericson, Keith Marzilli, and Justin R Sydnor. 2018. "Liquidity Constraints and the Value of Insurance." Working Paper. 2, 21

Fadlon, Itzik, and Torben Heien Nielsen. 2019. "Family Health Behaviors." *American Economic Review*, 109(9): 3162–91. 5

Fang, Hanming, Michael Keane, and Daniel Silverman. 2008. "Sources of Advantageous Selection: Evidence from the Medigap Insurance Market." *Journal of Political Economy*, 116(2): 303–350. 5

Finkelstein, Amy, Nathan Hendren, and Erzo Luttmer. forthcoming. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." *Journal of Political Economy*. 2

Gruber, Jonathan, Benjamin R Handel, Samuel H Kina, and Jonathan T Kolstad. 2020. "Managing Intelligence: Skilled Experts and AI in Markets for Complex Products." National Bureau of Economic Research. 31

Handel, Benjamin. 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review*, 103(7): 2643–2682. 2, 5, 33

Handel, Benjamin, and Jonathan Kolstad. 2015*a*. "Getting the Most From Marketplaces: Smart Policies on Health Insurance Choice." *Brookings Hamilton Project Discussion Paper 2015-08*. 4, 31, 37

Handel, Benjamin, and Jonathan Kolstad. 2015*b*. "Health Insurance For Humans: Information Frictions, Plan Choice, and Consumer Welfare." *American Economic Review*, 105(8): 2449–2500. 2, 5, 33

Handel, Benjamin, and Joshua Schwartzstein. 2019. "Behavioral Economics and Health Care Markets." *Handbook of Behavioral Economics*. 5

Handel, Benjamin R, Jonathan T Kolstad, and Johannes Spinnewijn. 2019. "Information frictions and adverse selection: Policy interventions in health insurance markets." *Review of Economics and Statistics*, 101(2): 326–340. 5, 31

Hastings, Justine, Chris Neilson, Anely Ramirez, Unika Shrestha, and Seth Zimmerman. 2013. "(Un)informed College Choice: Evidence From Chile." Brown University, Working Paper. 1

Ho, Kate, Joseph Hogan, and Fiona Scott Morton. 2017. "The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program." *The RAND Journal of Economics*, 48(4): 877–905. 34

Ito, Koichiro. 2015. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review*, 104(2): 537–563. 1

Ketcham, Jonathan, Claudio Lucarelli, Eugenio Miravete, and Christopher Roebuck. 2012. "Sinking, Swimming, or Learning to Swim in Medicare Part D?" *The American Economic Review*, 102(6): 2639–2673. 5

Kőszegi, Botond, and Matthew Rabin. 2007. "Reference-dependent risk attitudes." *American Economic Review*, 97(4): 1047–1073. 33

Kroneman, Madelon, Wienke Boerma, M van den Berg, P Groenewegen, J de Jong, and E. van Ginneken. 2016. "Netherlands: Health System Review." *Health Systems in Transition*, 18(2): 1–240. 6

Madrian, B., and D. Shea. 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics*, 141(4): 1149–1187. 2

Manski, Charles. 1999. *Identification Problems in the Social Sciences.* Harvard University Press. 3

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies*, 60(3): 531–542. 21

Mullainathan, Sendhil, and Eldar Shafir. 2013. *Scarcity: Why having too little means so much.* Macmillan. 1

Neilsen, Christopher. 2017. "Targeted Vouchers, Competition Among Schools, and the Academic Achievements of Poor Students." Princeton Working Paper. 1

Polyakova, Maria. 2016. "Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D." *American Economic Journal: Applied Economics*, 8(3): 165–195. 5

Remmerswaal, Minke, Jan Boone, and Rudy Douven. 2019. "Selection and Moral Hazard Effects in Healthcare." CPB Netherlands Bureau for Economic Policy Analysis Working Paper. 3, 15

Schippers, Edith. 2016. "Kamerbrief van 24 November 2016 over Onderzoek naar het Vrijwillig Eigen Risico." Ministerie van Volksgezondheid, Welzijn en Sport. 1, 7

Sorenson, Alan. 2006. "Social Learning and Health Plan Choice." *The RAND Journal of Economics*, 37(4): 929–945. 2, 5

Spinnewijn, Johannes. 2017. "Heterogeneity, Demand for Insurance and Adverse Selection." *American Economic Journal: Economic Policy*, 9(1): 308–343. 5

Sydnor, Justin. 2010. "Over(insuring) Modest Risks." *American Economic Journal: Applied Economics*, 2(4): 177–199. 2, 33, 34

van Kleef, Richard C, Frank Eijkenaar, and René CJA van Vliet. 2019. "Selection incentives for health insurers in the presence of sophisticated risk adjustment." *Medical Care Research and Review.* 6

Van Winssen, KPM, RC Van Kleef, and WPMM Van de Ven. 2015. "How profitable is a voluntary deductible in health insurance for the consumer?" *Health Policy*, 119(5): 688–695. 5

Van Winssen, KPM, RC Van Kleef, and WPMM Van de Ven. 2016. "Potential determinants of deductible uptake in health insurance: How to increase uptake in The Netherlands?" *The European Journal of Health Economics*, 17(9): 1059–1072. 5

A Data Appendix

This Data Appendix provides information on the additional datasets we linked to our health cost and insurance data at Statistics Netherlands. Datasets are linked at the individual level based on anonymized individual identifiers.

Age and gender Dataset *Gbapersoontab* provides an overview of all people registered living in the Netherlands at any point since 1995. These registers form a basis for the administrative records of all individuals in the Netherlands. For our purposes, *Gbapersoontab* is used to obtain age and gender, and we use this person registry as the primary dataset to match all other datasets with.

Family and household links Family links come from the dataset *Kindoudertab*, which contains all known legal child-parent links. Household identifiers as well as family status variables in *Ipi* and *Inpatab* allow us to identify partners and other household links. Partnerships consist of all partners who are living in the same household.[35]

Education *Hoogsteopltab* is a dataset that includes the highest attained educational course for each individual, and originates from several educational registers and survey data. We link each educational course to its relevant International Standard Classification of Education (ISCED) level and field of education. There is almost universal coverage for the youngest cohorts, but educational information is missing for many individuals aged over 40. Overall, we observe highest education obtained for 54.6% of our full sample.

Income and Employment Status Datasets *Ihi* and *Inhatab* contain information on households' income, and originates from tax authorities. Our main definition of income used in the analysis is household gross income (called *bruto inkomen* by Statistics Netherlands). Gross income includes all labor income and capital income, as well as government transfers (e.g., UI, DI, pensions), and other transfers and income. We also use a socio-economic classification variable *seccoal1*, which classifies each individual based on where the majority of his or her personal income comes from. This variable is obtained from datasets *Ipi* and *Inpatab*.

Wealth Dataset *Vehtab* contains information from tax authorities on households' assets and debts. This information is partly self-reported (on tax forms) and third-party reported. Assets include financial assets (savings, stocks, bonds, and other participations), real estate and other assets (such as cash and movable assets). Debts include mortgages, study debt and other debt. The net wealth variable in the main text equals household assets minus household debts.

Employee-Employer links We use the dataset *Spolisbus* to link individuals to their firms, colleagues and sector. *Spolisbus* is a highly detailed dataset with monthly information on all employment contracts in the Netherlands, collected by the tax authorities based on third-party reported data. We adopt the same definition of a firm as in the firm registry (*Algemeen Bedrijfsregister*) of Statistics Netherlands. We sum each individual's total hours worked by year by firm. For each individual, we then select the firm at which that individual has worked the most hours in each year. The colleagues that we identify are thus all individuals who work the majority of their hours at the same firm. The sector categorization that we adopt is made by the authorities based on the collective labor agreements.

Location We match every individual with their yearly 6-digit postcode based on their registered residence. For this, we use datasets *Gbaadresobjectbus* and *Vslgwbtab*. Postcodes are obtained for each year on 1 October, as this is close to the period of deciding on their health insurance contract. 6-digit postcode information is at a neighbourhood level, and there are 12'116 distinct postcodes in 2015.

[35] This includes married partners, registered partners, but also partners who have not registered their partnership but are living in the same household.

B Health Cost Predictions

In this Appendix, we describe the binary prediction algorithm that we use to obtain risk probabilities, and discuss its accuracy across different subgroups, and the most important predictors. We also discuss an alternative non-binary prediction algorithm and argue why the binary predictions are preferable for the analysis in this paper.

B.1 Prediction Algorithm

We use an ensemble machine learning algorithm to predict the probability that an individual's health costs will not exceed the mandatory deductible of 375 EUR in any given year. The prediction algorithm we use is a standard machine learning method for binary classification, an ensemble learner that consists in our case of a random forest model, gradient boosted regression trees and LASSO model. To avoid overfitting, we train and calibrate the prediction algorithm on a training sample of 1.25 million individuals. We then use this trained prediction algorithm to obtain predictions for a hold-out sample of about 12 million individuals. All the analyses and statistics in the paper are developed use only this hold-out sample.

The prediction method we use follows four steps, which closely resemble the steps used in Einav et al. (2018). First, we follow standard practice in machine learning by tuning key parameters that govern the prediction models by 3-fold cross-validation. Second, we train the three resulting prediction models separately. Third, we combine the three obtained predictions into one using a linear combination that we calibrate in the data. Finally, we calibrate the resulting final ensemble predictions using a linear spline. As there is some variation in the number and definition of predictors that we have across time, we repeat these four steps for all years of study (2013-2017). We describe each of the four steps in more detail here.

Parameter Tuning As the three machine learning models that we use have parameters that are choosable by the researcher, we follow standard practice and tune these parameters using 3-fold cross validation. More specifically, we tune the following parameters using 100,000 observations: minimal node size (mid.node.size), number of variables used at each node (mtry) for the random forest model, learning rate (eta) for the boosted regression trees, and the shrinkage parameter (lambda) for the LASSO.[36] For each of these parameters, we optimize among 5 alternatives. We tune these parameters using 3-fold cross validation, where we are optimizing the area under the receiver operating characteristic curve (AUC).[37] Thus, for each of the parameter values we want to test, the model is trained on 2 folds (subsets of the training sample), and then the performance is measured in the 3rd fold. The parameter values for which the AUC in the hold-out sample is highest for each prediction algorithm are: mtry = 10, min.node.size = 10, eta = 0.2, lambda = 0.0001.

Estimating the Models Using these tuned parameter values, all models are estimated using a training sample of 800,000 individuals.

Obtaining Ensemble Predictor We combine the predictions from the random forest, gradient boosting regression trees, and LASSO into one ensemble prediction. Following Einav et al. (2018), we construct the ensemble prediction to be the linear combination $p_{ensemble} = \hat{\beta}_{rf}\hat{p}_{rf} + \hat{\beta}_{gb}\hat{p}_{gb} + \hat{\beta}_{lasso}\hat{p}_{lasso}$, where $\hat{p}_x$ is the prediction from algorithm x and $\hat{\beta}_x$ is the associated weight.

[36] We use the package CARET in R that provides a standardized way to tune parameters. The prediction models we use are RANGER (random forest), XGBLINEAR (boosted regression trees), and GLMNET (LASSO).

[37] This is a common metric used in the machine learning literature to measure the performance of a prediction model.

We obtain estimates for the weights from a constrained linear regression (with no constant and the weights summing to one) of the dummy for having costs below 375 EUR on the three individual predicted probabilities. For this step, we use 100,000 observations that we did not use in either step of parameter tuning nor the estimation of the models. We find associated weights in 2015 that are $\hat{\beta}_{rf} = 0.67$, $\hat{\beta}_{gb} = 0.08$ and $\hat{\beta}_{lasso} = 0.25$.

Calibrating Probabilities Finally, the raw probability predictions we get from the ensemble step are calibrated to the actual observed probabilities by estimating a linear spline. This calibration is done using 350,000 observations that are used in none of the previous steps. 10 equal sized bins are created based on the ranked predicted probability. In every bin the mean probability is calibrated to the observed mean probability for these observations. The piece-wise linear spline that follows from linearly interpolating all intermediary points serves as the last step in the prediction mechanism.

B.2 Additional Discussion of Prediction

While Figure 3 shows a calibration plot for the entire sample, Figure B.1 shows a calibration plot for certain subgroups of the sample. We see from Panel B, C and D that probabilities are well calibrated for distinct groups of education level, income quartile and age group. This makes us comfortable that the observed differences in choice quality across these different groups are not due to differential prediction accuracy of our ensemble predictor. Moreover, panel A of Figure B.1 shows that individuals who choose a 500 EUR deductible are more likely to have low costs than individuals who choose no extra deductible, conditional on the prediction of our model. However, the difference in *ex post* realized low cost fraction is small, leading us to conclude that the private information and moral hazard, conditional on our predictors, is small. More specifically, the average gap across probability bins between individuals who choose and who do not choose an extra deductible is 6.667%. Taking into account that across probability bins, the average share with low costs among people without extra deductible is 51.215%, we find that individuals who take a deductible are on average are 13.017% more likely to have low costs than our model predicts.

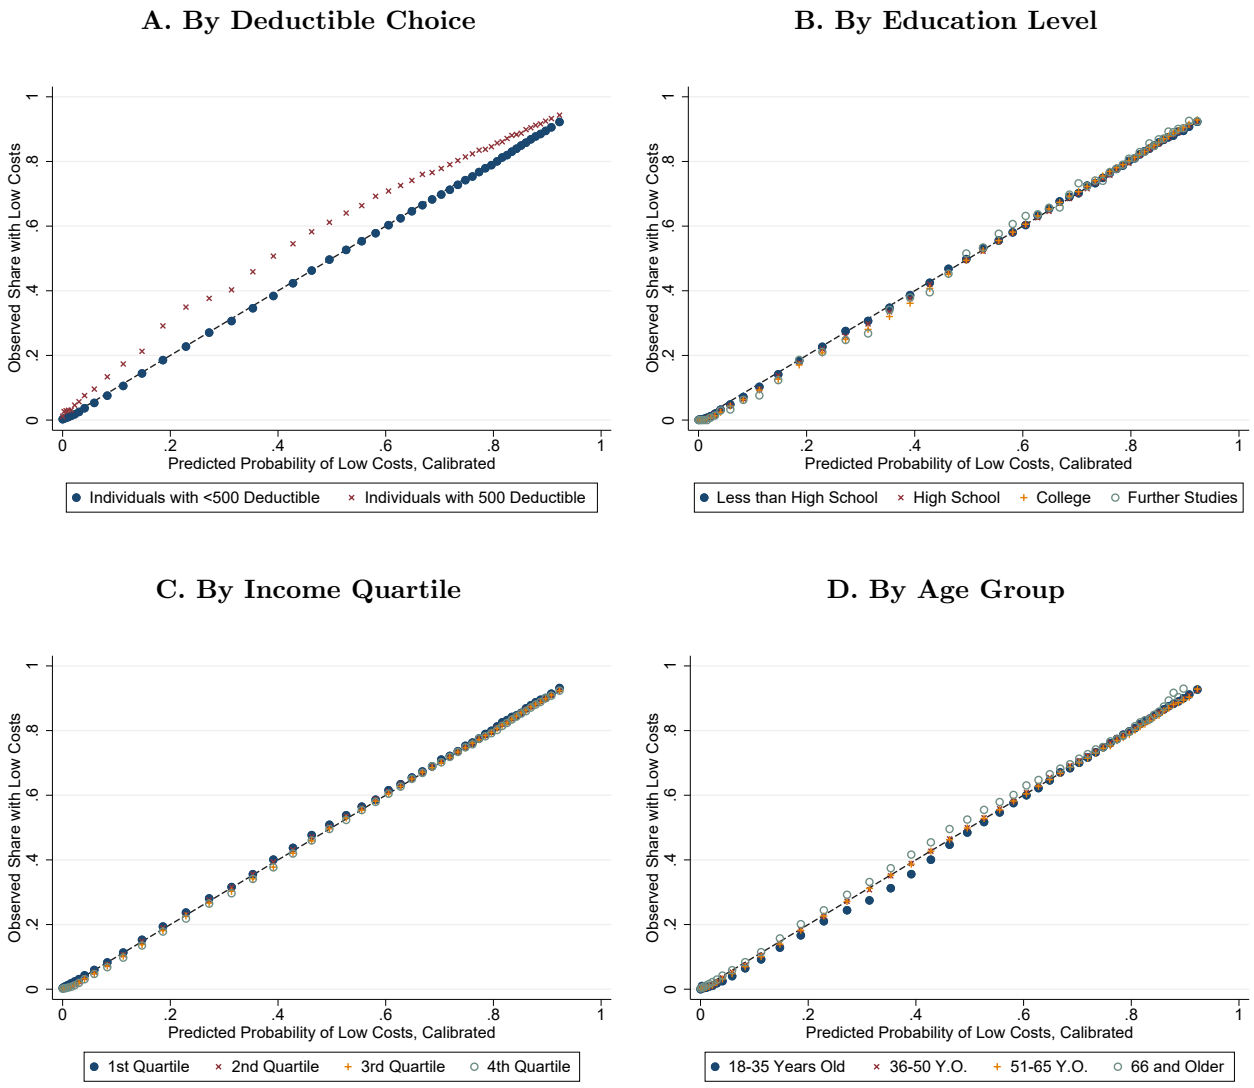
Figure B.2 presents the importance of different predictors in the random forest model, which is the model with the highest weight in our ensemble prediction. Not surprisingly, the most important predictors are different categories of past pharmaceutical spending, with $t-1$ values being more important than $t-2$ values. Hospital costs, costs to primary care visits and age are other important variables in the random forest prediction.

B.3 Non-Binary Prediction

In Section III.B, we simplified the deductible choice problem in the Netherlands into a binary choice between selecting a 875 EUR deductible, or the mandatory 375 EUR deductible. This is a simplification, as in fact there are 6 different deductible choices possible, which apply to different brackets with cutoffs at 375, 475, 575, 675, 775 and 875 EUR. However, two pieces of evidence show that reducing the problem to a binary one is appropriate for our context.

First, Panel A of Figure B.3 shows that the ex-post observed shares within each intermediary deductible bracket are small. This means that only a small fraction of individuals fall into the intermediary deductible ranges, which decreases the likelihood that the intermediary deductibles are optimal choices. Second, we find that when using a machine learning classifier to predict which individuals are going to fall into the intermediary brackets, the predicted mass in these intermediary brackets is small. Panel B of Figure B.3 shows that ex-ante, a random forest model trained on an unbalanced sample will give less than 1% probability mass to the intermediate

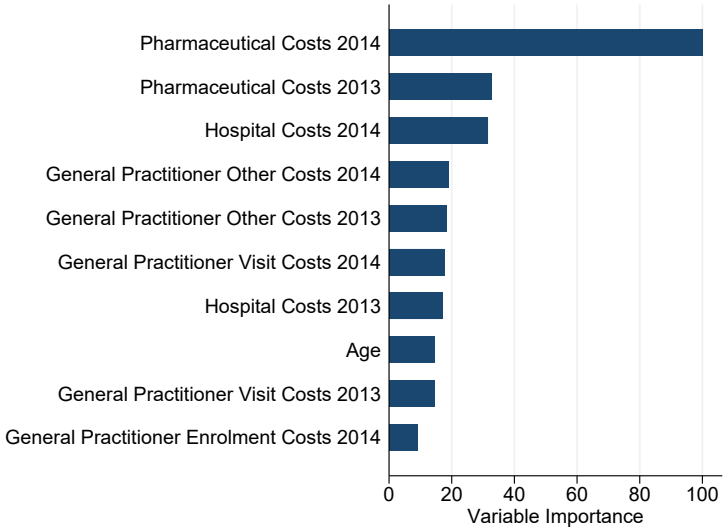
FIGURE B.1: PREDICTED VS. OBSERVED SHARE OF LOW COSTS, BY SUBGROUPS



Notes: This figure shows the calibration plot of the predicted probability of low costs for various subgroups of the sample. Panel A plots our prediction against the observed share of people with health costs below 375 EUR, separately for people having chosen the 500 deductible and people who have not. Panel B does the same exercise splitting the sample by education level. In Panel C, the sample in split by income quartile, and in Panel D, by age group.

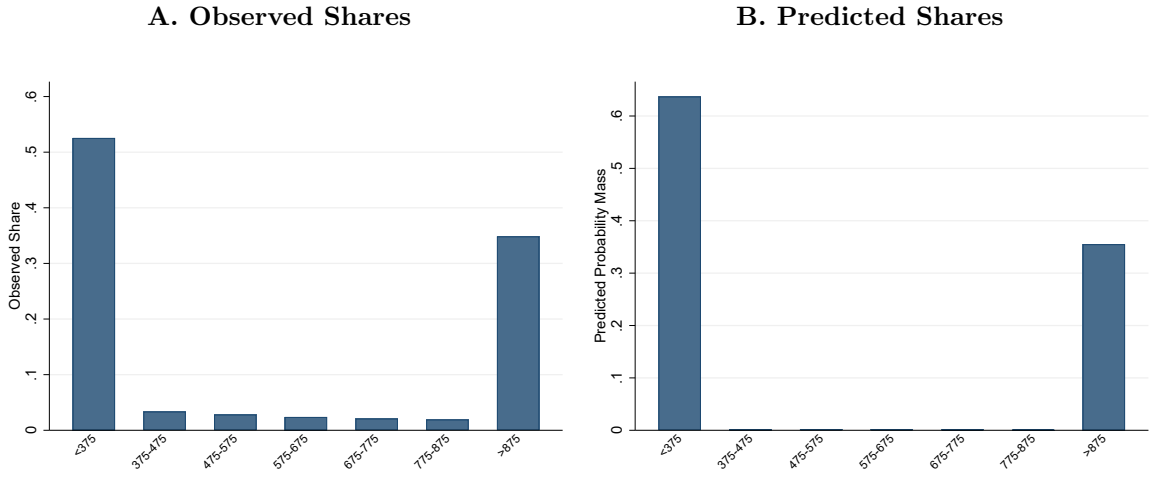
categories. This is largely due to the unbalanced classes, where the majority of individuals fall into the lowest or highest bracket. However, insofar as we cannot expect individuals to predict their future costs more accurately, the low probability with which most individuals are predicted to be in the intermediary deductible brackets further strengthens the case for a binary decision rule.

FIGURE B.2: VARIABLE IMPORTANCE IN PREDICTION WITH RANDOM FOREST



Notes: This figure shows the importance of selected variables in the prediction of health cost risk using only a random forest model. Variable importance is measured by the mean decrease in gini, ie. the average of a variable's total decrease in node impurity, weighted by the proportion of samples reaching that node in each individual decision tree in the random forest.

FIGURE B.3: COST PREDICTIONS WITH MULTIPLE DEDUCTIBLE CATEGORIES



Notes: Panel A plots the observed share of individuals with health costs in all the deductible health cost brackets in 2015. Panel B plots the predicted shares of individuals in all deductible health cost brackets, where the prediction is from a random forest with the same predictors as described in Section III.B.

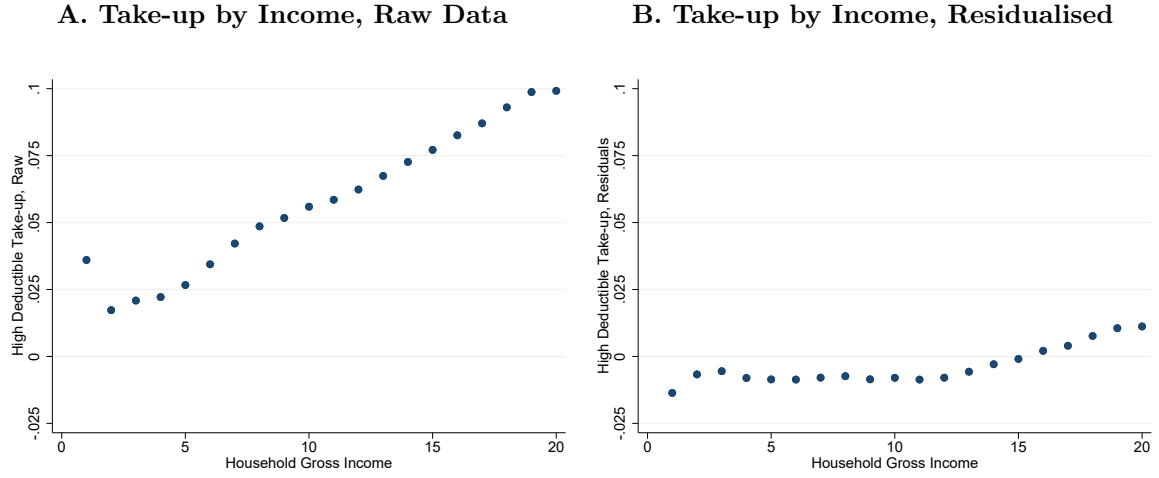
C Deductible Choice: Appendix Figures and Tables

Table C.1: Deductible Take-up: Impact of Health and Income Changes

	(1) No FE	(2) Individual FE	(3) First difference	(4) First difference	(5) First difference
Probability of Low Costs	0.115***	0.0570***	0.0422***		
Prob. Low Costs, Positive Δ				0.00691***	
Prob. Low Costs, Negative Δ				-0.0670***	
Δ Prob. Low Costs > +2 Deciles					0.0102***
Δ Prob. Low Costs = +2 Deciles					0.00685***
Δ Prob. Low Costs = +1 Decile					0.00342***
Δ Prob. Low Costs = -1 Decile					-0.00277***
Δ Prob. Low Costs = -2 Deciles					-0.00636***
Δ Prob. Low Costs < -2 Deciles					-0.0202***
Income ('000 EUR)	6.06e-05***	1.57e-05***	6.63e-06***	6.65e-06***	6.85e-06***
Number of Individuals	12,317,248	12,317,248	12,074,444	12,058,624	12,074,444
Observations	47,685,794	47,685,794	35,368,540	35,216,196	35,368,540

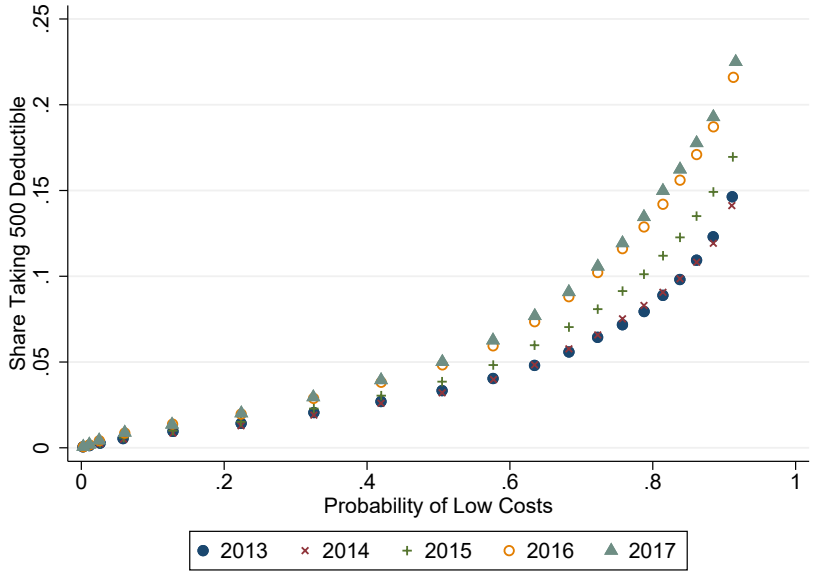
Notes: This table presents the result of an OLS regression of take-up of the 500 EUR extra deductible on probability of low costs, changes in probability of low costs, income, and changes in income. In column (1), take-up of the high deductible is regressed on the probability to have health costs lower than 375 EUR, and on income in thousands of EUR. Column (2) adds individual fixed effects. Column (3) regresses the first difference of deductible take-up on the first difference of the probability of low costs and the first difference of income. Column (4) splits the first difference in two distinct variables, one containing only positive shocks, the other only negative shocks. Column (5) creates six dummies capturing shocks of various magnitudes: positive and negative shocks of one, two, and strictly more than two deciles. In Columns (4) and (5), income first difference remains unchanged compared to Column (3). All regressions include year fixed effects. *** p<0.01, ** p<0.05, * p<0.1 with robust standard errors.

Figure C.1: Deductible Take-up as a Function of Income



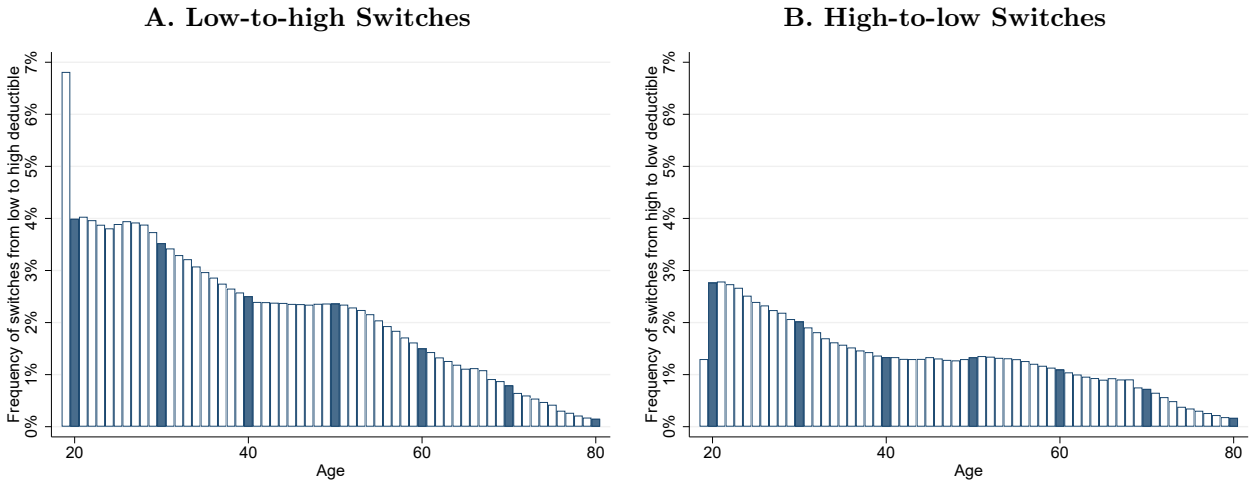
Notes: These figures plot the relationship between household gross income and the take-up of the 500 EUR extra deductible. Panel A plots take-up of 500 deductible by household income percentile. Panel B plots the residuals of an OLS regression of take-up of 500 EUR extra deductible on risk probability, four levels of education dummies, four age dummies, and indicators for gender, having a partner, and having children.

FIGURE C.2: DEDUCTIBLE CHOICE GRADIENT BY YEAR



Notes: This figure displays the relationship between take-up of the 500 deductible and the predicted probability of low costs, separately for the five years included in our final sample.

FIGURE C.3: FREQUENCY OF DEDUCTIBLE SWITCHES BY AGE



Notes: This figure displays the frequency of deductible switches by age, in years 2014 to 2017. Panel A displays only switches to a higher deductible, and Panel B to a lower deductible.

Table C.2: Robustness Check

	Baseline			500 vs. 0 Deductible			0 vs. >0 Deductible		
	Without	With Interaction		Without	With Interaction		Without	With Interaction	
	Interaction	*intercept*	*slope*	Interaction	*intercept*	*slope*	Interaction	*intercept*	*slope*
High School	0.017***	-0.011***	0.057***	0.018***	-0.012***	0.061***	0.025***	-0.014***	0.077***
College Degree	0.065***	-0.034***	0.165***	0.071***	-0.038***	0.181***	0.089***	-0.037***	0.210***
Further Studies	0.091***	-0.047***	0.226***	0.099***	-0.052***	0.250***	0.123***	-0.044***	0.275***
2nd Income Quartile	-0.003***	0.004***	-0.007***	-0.003***	0.004***	-0.007***	0.002***	0.009***	-0.005***
3rd Income Quartile	0.004***	0.004***	0.007***	0.005***	0.003***	0.009***	0.014***	0.011***	0.013***
4th Income Quartile	0.024***	0.002***	0.039***	0.026***	0.001***	0.045***	0.041***	0.015***	0.048***
36 to 50 years old	-0.011***	0.020***	-0.045***	-0.010***	0.022***	-0.046***	-0.006***	0.024***	-0.042***
51 to 65 years old	-0.004***	0.029***	-0.047***	-0.004***	0.030***	-0.048***	0.003***	0.036***	-0.045***
65+ years old	-0.001***	0.034***	-0.082***	0.000**	0.036***	-0.085***	0.007***	0.043***	-0.092***
Male	0.011***	-0.004***	0.025***	0.012***	-0.004***	0.028***	0.017***	-0.001***	0.030***
Has Partner	0.003***	-0.002***	0.013***	0.003***	-0.002***	0.014***	0.002***	-0.005***	0.018***
Has Children	-0.010***	0.004***	-0.028***	-0.011***	0.004***	-0.031***	-0.014***	0.004***	-0.035***
Self-employed	0.009***	-0.006***	0.026***	0.009***	-0.007***	0.028***	0.013***	0.000	0.023***
Constant	-0.042***	-0.041***		-0.045***	-0.043***		-0.055***	-0.044***	
Prob. Low Costs	0.122***		0.098***	0.129***		0.100***	0.169***		0.124***
Year and Insurer FE	YES	YES		YES	YES		YES	YES	
Observations	57,100,388	57,100,388		55,335,880	55,335,880		57,100,388	57,100,388	

	Baseline			Probit			Binary Pred. Low Costs		
	Without	With Interaction		Without	With Interaction		Without	With Interaction	
	Interaction	*intercept*	*slope*	Interaction	*intercept*	*slope*	Interaction	*intercept*	*slope*
High School	0.017***	-0.011***	0.057***	0.022***	0.006***	0.023***	0.019***	0.002***	0.032***
College Degree	0.065***	-0.034***	0.165***	0.051***	0.014***	0.051***	0.068***	0.013***	0.081***
Further Studies	0.091***	-0.047***	0.226***	0.063***	0.005**	0.081***	0.093***	0.019***	0.105***
2nd Income Quartile	-0.003***	0.004***	-0.007***	0.003***	0.030***	-0.040***	-0.001***	0.003***	-0.005***
3rd Income Quartile	0.004***	0.004***	0.007***	0.010***	0.041***	-0.044***	0.007***	0.008***	0.002***
4th Income Quartile	0.024***	0.002***	0.039***	0.024***	0.057***	-0.048***	0.027***	0.017***	0.016***
36 to 50 years old	-0.011***	0.020***	-0.045***	-0.008***	-0.007***	-0.001	-0.013***	-0.005***	-0.009***
51 to 65 years old	-0.004***	0.029***	-0.047***	0.000	0.001**	-0.002**	-0.012***	-0.005***	-0.006***
65+ years old	-0.001***	0.034***	-0.082***	-0.011***	-0.008***	-0.004***	-0.017***	-0.008***	-0.025***
Male	0.011***	-0.004***	0.025***	0.006***	0.007***	-0.002***	0.015***	0.001***	0.023***
Has Partner	0.003***	-0.002***	0.013***	0.005***	0.007***	-0.003***	0.003***	0.000***	0.008***
Has Children	-0.010***	0.004***	-0.028***	-0.009***	-0.000	-0.011***	-0.011***	-0.000	-0.020***
Self-employed	0.009***	-0.006***	0.026***	0.008***	0.018***	-0.013***	0.011***	0.005***	0.009***
Constant	-0.042***	-0.041***					-0.014***	-0.003***	
Prob. Low Costs	0.122***		0.098***	0.169***		0.191***			
Pred. Costs <375							0.062***		0.034***
Year and Insurer FE	YES	YES		YES	YES		YES	YES	
Observations	57,100,388	57,100,388		57,100,388	57,100,388		57,100,388	57,100,388	

Notes: This table performs a range of robustness checks on our baseline results. In the top panel, we compare our baseline regression with alternative definition of take-up of the high deductible. In the baseline, we define take-up as choosing the 500 deductible, as opposed to choosing any other deductible. In the second top panel, we keep only choices that are the 500 or the 0 deductible, and drop intermediate choices. In the third top panel, we instead define take-up as choosing any deductible strictly greater than 0. In the second bottom panel, we compare our baseline OLS regression with a probit specification. Finally, in the third bottom panel, we replace our linear probability of low costs with a binary indicator taking value one if the individual is predicted to have health costs lower than 375 EUR. In each panel, we present a regression with and without interacting our regressors with the probability of low costs.

Table C.3: Deductible Take-up Regression, Non Interacted

	(1) Baseline	(2) Education Field	(3) Professional Sector	(4) Liquidity and Financials	(5) Environment
High School	0.017***	0.016***	0.017***	0.015***	0.014***
College Degree	0.065***	0.062***	0.066***	0.062***	0.056***
Further Studies	0.091***	0.089***	0.097***	0.089***	0.088***
2nd Income Quartile	-0.003***	-0.008***	-0.011***	-0.009***	-0.006***
3rd Income Quartile	0.004***	0.000	-0.002***	-0.004***	-0.000
4th Income Quartile	0.024***	0.019***	0.017***	0.011***	0.014***
36 to 50 years old	-0.011***	-0.008***	-0.010***	-0.012***	0.007***
51 to 65 years old	-0.004***	-0.001***	-0.002***	-0.012***	0.027***
65+ years old	-0.001***	0.005***	0.003***	-0.016***	0.020***
Male	0.011***	0.015***	0.014***	0.012***	0.017***
Has Partner	0.003***	0.006***	0.004***	0.003***	0.008***
Has Children	-0.010***	-0.013***	-0.012***	-0.007***	-0.006***
Self-employed	0.009***	0.007***	0.008***	0.005***	0.007***
Statistics		0.139***			
Philosophy		0.024***			
Accounting and Taxation		0.012***			
Marketing and Advertising		-0.004***			
Hair and Beauty		-0.012***			
Protection of Persons		-0.033***			
Business Services			0.022***		
Insurance			0.027***		
Retail			-0.003***		
Construction			-0.013***		
Cleaning			-0.012***		
Public Utilities			0.001		
2nd Net Worth Quartile				0.004***	
3rd Net Worth Quartile				0.012***	
4th Net Worth Quartile				0.029***	
Has Savings > 2000EUR				0.008***	
Has Mortgage Debt				0.002***	
Has Other Debt				-0.009***	
Share of Colleagues with 500 Ded.					0.226***
Share in Postcode with 500 Ded.					0.404***
Father With 500 Deductible					0.181***
Mother With 500 Deductible					0.237***
Constant	-0.042***	-0.056***	-0.063***	-0.043***	-0.135***
Prob. Low Costs	0.122***	0.145***	0.148***	0.119***	0.160***
Year and Insurer FE	YES	YES	YES	YES	YES
Observations	57,100,388	30,799,129	32,299,835	57,013,765	16,938,401

Notes: Notes from Table 4 and 5 apply; this table displays the same regressions without interacting the regressors with the probability of low costs.

Table C.4: Predicted Health Risk by Observed and Optimal Deductible Choice

	2013	2014	2015	2016	2017
Probability of Low Costs	0.512	0.516	0.516	0.504	0.496
Healthy Individuals	0.752	0.758	0.760	0.759	0.762
Unhealthy Individuals	0.176	0.169	0.169	0.160	0.159
Individuals with 500 Deductible	0.748	0.760	0.763	0.762	0.763
Individuals with <500 Deductible	0.499	0.502	0.499	0.482	0.472
Share of Healthy Individuals	58.2%	58.9%	58.7%	57.4%	56.0%
Share of Individuals with the 500 Deductible	5.1%	5.3%	6.5%	8.0%	8.2%

Notes: This table displays, for the five years in our sample, the share of predictably healthy individuals and the share of individuals who took up the high deductible. It then shows the average probability of low costs for predictably healthy people (i.e., with a probability of low costs greater than .5), predictably unhealthy people, people who have taken up the 500 deductible and those who have not.

Table C.5: Deductible Take-up and Predicted Health by Field

	Education Field	(1) Take-up of 500 Deductible	(2) Probability Low Costs	(3) Take-up of 500 Ded. | Being Predictably Healthy
1	**Statistics**	29%	87%	34%
2	Mathematics	21%	85%	27%
3	Physics	21%	91%	26%
4	Architecture and town planning	18%	88%	21%
5	Physical science	18%	82%	22%
6	Earth science	18%	88%	21%
7	**Philosophy and ethics**	17%	82%	21%
8	Medicine	17%	83%	20%
9	Chemistry	16%	87%	20%
10	Biology and biochemistry	16%	83%	20%
11	Science, Mathematics and Computing	16%	85%	19%
12	Computer science	15%	87%	18%
13	Environmental protection	15%	86%	18%
14	Political science and civics	15%	85%	18%
15	Design	15%	85%	18%
16	Sociology and cultural studies	14%	82%	18%
17	Mining and extraction	14%	91%	17%
18	Economics	14%	84%	17%
19	Humanities and Arts	14%	84%	18%
20	Dental studies	14%	76%	18%
21	History and archaeology	13%	82%	16%
22	Business and administration	13%	82%	16%
23	Pharmacy	13%	73%	17%
24	Health	13%	79%	16%
25	Environmental protection technology	13%	84%	15%
26	Medical diagnostic and treatment technology	13%	81%	16%
27	Religion	13%	80%	17%
28	Law	13%	80%	16%
29	Psychology	12%	77%	16%
30	Management and administration	12%	81%	16%
31	Engineering and engineering trades	12%	87%	15%
32	Forestry	12%	86%	14%
33	Therapy and rehabilitation	12%	78%	15%
34	Finance, banking, insurance	12%	80%	15%
35	Social and behavioural science	12%	79%	15%
36	Health and Welfare	12%	80%	15%
37	Fisheries	12%	94%	15%
38	Journalism and reporting	12%	80%	14%
39	Training for teachers w. subject specialisation	11%	79%	14%
40	Education science	11%	75%	14%
41	**Accounting and taxation**	11%	78%	14%
42	Agriculture, forestry and fishery	10%	81%	13%
43	**Marketing and advertising**	10%	80%	13%
44	Chemical and process	10%	85%	12%
45	Arts	10%	80%	13%
46	Electronics and automation	10%	86%	12%

Table C.5: Deductible Take-up and Predicted Health by Field (cont'd)

47	Music and performing arts	10%	81%	12%
48	Training for teachers of vocational subjects	10%	81%	12%
49	Fine arts	10%	82%	12%
50	Humanities	10%	76%	12%
51	Library, information, archive	9%	78%	12%
52	Travel, tourism and leisure	9%	77%	12%
53	Electricity and energy	9%	88%	11%
54	Veterinary	9%	75%	12%
55	Mother tongue	9%	74%	12%
56	Audio-visual techniques and media production	9%	83%	10%
57	Building and civil engineering	9%	86%	10%
58	Life science	9%	79%	11%
59	Crop and livestock production	9%	79%	11%
60	Mechanics and metal work	9%	85%	10%
61	Wholesale and retail sales	8%	79%	11%
62	Foreign languages	8%	74%	11%
63	Motor vehicles, ships and aircraft	8%	87%	10%
64	Training for teachers at basic levels	8%	75%	10%
65	Materials (wood, paper, plastic, glass)	8%	86%	9%
66	Sports	8%	83%	10%
67	Teacher training and education science	8%	74%	10%
68	Military and defence	7%	81%	9%
69	Transport services	7%	83%	9%
70	Food processing	7%	78%	9%
72	Natural environments and wildlife	6%	86%	7%
73	Hotel, restaurant and catering	6%	77%	8%
74	Basic / broad, general programmes	6%	72%	9%
75	Social work and counselling	6%	70%	8%
77	Personal skills	6%	68%	8%
78	Textiles, clothes, footwear, leather	5%	70%	7%
79	Horticulture	5%	80%	6%
80	General Programmes	5%	71%	7%
81	Nursing and caring	5%	66%	7%
82	Domestic services	5%	66%	7%
83	Secretarial and office work	5%	65%	7%
84	**Protection of persons and property**	4%	78%	6%
85	Child care and youth services	4%	66%	6%
86	Computer use	4%	65%	6%
87	**Hair and beauty services**	4%	65%	5%
88	Occupational health and safety	4%	75%	5%
89	Training for pre-school teachers	3%	62%	0%
90	Literacy and numeracy	2%	62%	4%

Notes: For each field of study, this table shows: in Column (1), the fraction of individuals who take-up the 500 EUR extra deductible, in Column (2), the fraction of individuals with a probability of low costs < 375 EUR, and in Column (3), the fraction of individuals who take-up the 500 EUR extra deductible, conditional on having predicted health costs < 375 EUR.

Table C.6: Deductible Take-up and Predicted Health by Professional Sector

	Professional Sector	(1) Take-up of 500 Deductible	(2) Probability Low Costs	(3) Take-up of 500 Ded. | Being Predictably Healthy
1	**Business Services II**	13%	84%	16%
2	**Insurance and Health Insurance Firms**	12%	79%	15%
3	Business Services I	12%	82%	15%
4	Dairy Industry	12%	82%	14%
5	Banks	10%	81%	12%
6	Other Passenger Transport Land and Air	10%	79%	13%
7	Business Services III	10%	79%	13%
8	Agriculture	10%	85%	11%
9	Stoneware	9%	83%	11%
10	Publishers	9%	79%	11%
11	Cultural Institutions	9%	80%	11%
12	Telecommunications	9%	81%	12%
13	Government, Education and Science	9%	75%	12%
14	Food Industry	9%	80%	11%
15	Catering Industry I	9%	84%	10%
16	Tobacco Processing Industry	9%	76%	11%
17	Wholesale I	8%	82%	11%
18	Wholesale II	8%	81%	10%
20	Government, Police and Judiciary	8%	74%	11%
21	Wholesale of Wood	8%	82%	10%
22	Electronic Industry	8%	81%	13%
23	Carpentry	8%	83%	9%
24	Furniture and Organ Building	8%	83%	9%
25	Rail Construction	8%	78%	11%
26	NS Transport	8%	74%	10%
27	Sugar Processing Industry	7%	78%	10%
28	Chain Stores	7%	80%	9%
29	**Retail**	7%	79%	9%
30	Lending Industry	7%	81%	9%
31	Other Branches of Business	7%	79%	9%
32	Postal Transport	7%	72%	10%
33	Metal Industry	7%	80%	10%
34	**Construction**	7%	83%	9%
35	Merchant	7%	89%	8%
36	Mortar	7%	72%	9%
37	KLM Transport	7%	77%	9%
38	Bakeries	7%	79%	9%
39	Metal and Technical Industry	7%	82%	8%
40	Port Companies	7%	82%	9%
41	Chemical Industry	7%	79%	9%
42	General Industry	7%	81%	9%
43	Stone, Cement, Glass and Ceramic Industry	7%	77%	9%
44	Butchers Other	7%	80%	8%
45	Health, Mental and Social Industry	7%	71%	9%
46	Printing Industry	7%	80%	8%
47	Textiles Industry	7%	77%	9%
48	Inland Shipping	7%	83%	8%
49	Private Bus Transport	6%	70%	9%

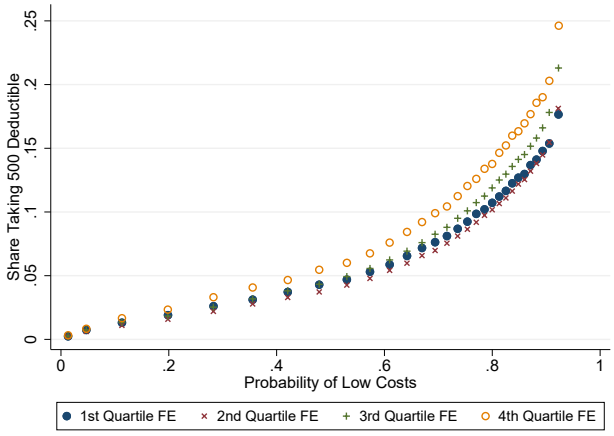
Table C.6: Deductible Take-up and Predicted Health by Professional Sector (cont'd)

50 Government, Local Government	6%	70%	9%
51 Butchers	6%	79%	8%
52 Wood, Brush and Packaging Industry	6%	82%	8%
53 Other Goods Transport Land and Air	6%	80%	8%
54 Government, Defense	6%	82%	11%
55 **Government, Public Utilities**	6%	77%	7%
56 Public Transport	5%	65%	8%
57 Security	5%	75%	7%
58 Plastering	5%	85%	6%
59 Taxi and Ambulance	5%	65%	8%
60 Catering Industry II	5%	70%	7%
61 Painting Industry	5%	81%	6%
62 Port Classifiers	5%	79%	6%
63 Fishing	4%	81%	6%
64 Work and Integration	4%	64%	6%
65 Dredging Industry	4%	85%	9%
66 Government, Other Institutions	4%	60%	7%
67 Roofing	4%	82%	5%
68 **Cleaning**	3%	70%	5%

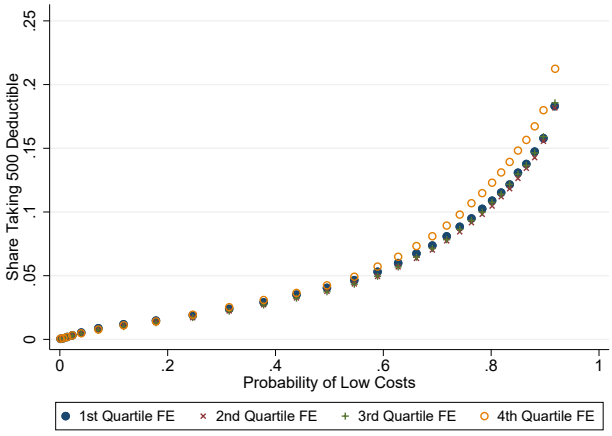
Notes: For each professional sector, this table shows: in Column (1), the fraction of individuals who take-up the 500 EUR extra deductible, in Column (2), the fraction of individuals with a probability of low costs < 375 EUR, and in Column (3), the fraction of individuals who take-up the 500 EUR extra deductible, conditional on having predicted health costs < 375 EUR.

FIGURE C.4: TAKE-UP VS. PROBABILITY OF LOW COSTS BY PEER EFFECTS

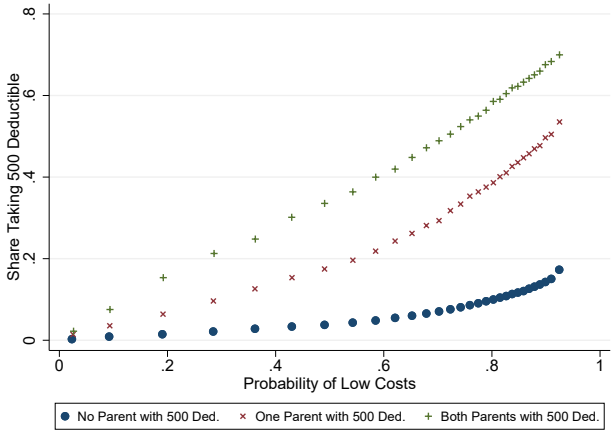
A. By Firm Fixed Effects



B. By Location Fixed Effects



C. By Parents' Choice



Notes: This figures shows the relationship between the probability of low costs and take-up of the high deductible for different subgroups. In Panel A, individuals are split in quartiles according to the fixed effect of the firm they are employed by. Those fixed effects are computed as detailed in Section IV.C.1. In Panel B, individuals are split in quartiles of postcode fixed effects, computed following the same method. In Panel C, individuals are split according to whether none of their parents, one of their parents, or both parents have taken up the 500 deductible.

D Choice Frictions, Welfare and Policies: Further Details

This Appendix Section provides further details underlying our analysis of choice quality, the counterfactual analysis and the microfoundations of choice frictions.

D.1 Predicted Choice Model

For our analysis of choice quality in Section V, we start by predicting the deductible take-up rate $d(X_{it}, \pi_{it})$ as a function of their predicted health π_{it}, observables X_{it} and their interaction by running the regression:

$$Y = \alpha + \sum \beta_\delta 1[\pi = \delta] + \gamma X + \sum \nu_\delta 1[\pi = \delta] X + \epsilon$$

Here, Y is a binary variable that is 1 when an individual takes the 500 voluntary deductible and X is a rich set of controls, including demographics (gender, age, having children, living with a partner), financial variables (household gross income in deciles, net worth in quartiles, a dummy for having savings $>$ 2000 EUR, for having a mortgage debt, for having another type of debt), education level and field, professional sector, and environment variables (firm and location fixed effect identified in Section IV.C.1 in deciles, mother and father take-up of the high deductible).

We then define

$$d_{\pi_{pop}}(X_{it}) = \sum_\delta d(X_{it}, \delta)\, dF_\delta,$$

which gives us the predicted deductible take-up rate for each observed X_{it} combination but as if there were a population of individuals with that X_{it} with the same health distribution as the overall population. In the same way, we predict the choice quality for individuals with demographic vector X_{it}, as captured by the probability to choose the contract that minimizes expected expenditures, $d^*_{\pi_{pop}}(X_{it})$, and the corresponding average financial loss $\Delta w^{*,\sigma}_{\pi_{pop}}(X_{it})$. That is,[38]

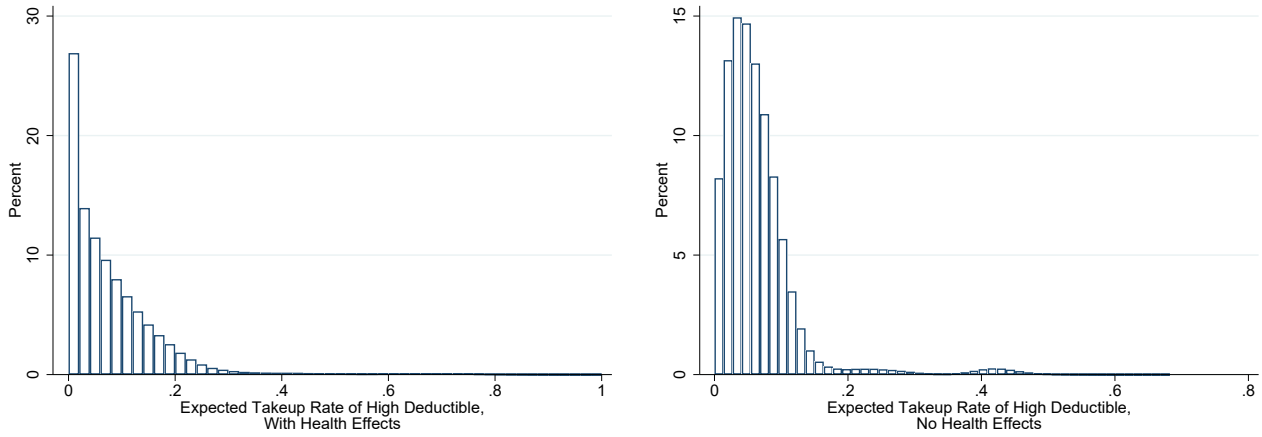
$$\begin{aligned} d^{*,\sigma}_{\pi_{pop}}(X_i) &= \sum_\delta \{1[\pi_\delta \leq .5][1 - d(X_{it}, \delta)] + 1[\pi_\delta > .5]\, d(X_{it}, \delta)\} dF_\delta, \\ \Delta w^*_{\pi_{pop}}(X_{it}) &= \sum \{1[\pi_\delta \leq .5]\, d(X_{it}, \delta)\left[CE^\sigma_{\pi_\delta,0} - CE^\sigma_{\pi_\delta,500}\right] + 1[\pi_\delta > .5][1 - d(X_{it}, \delta)]\left[CE^\sigma_{\pi_\delta,500} - CE^\sigma_{\pi_\delta,0}\right]\} dF_\delta. \end{aligned}$$

The choice quality varies through the deductible choice predicted by the set of demographics X_i for different health risks, but again reflects the population distribution of health risks.

Figure D.1 compares the distribution of predicted deductible choice, with and without the effect of healthcare cost risk. These are denoted in previous equations as $d(X_{it}, \pi_{it})$ and $d_{\pi_{pop}}(X_{it})$ respectively. As shown before, health has a meaningful impact on deductible choice, but there is substantial heterogeneity in likelihood of choosing a deductible just as a function of X_{it}, netting out health effects. While losses range up to 200 EUR when factoring health risk into choices, when assuming the population distribution of health for a given X_i the expected loss ranges between 50 and 80 as a function of X_i. Panel A of Figure D.2 ranks individuals according to the quality of their choice first, as discussed in the text, and then shows the distribution of the probability to make the right decision for the different groups of quality choice. Panel B of Figure D.2 shows the probability of making the right decisions for different income groups.

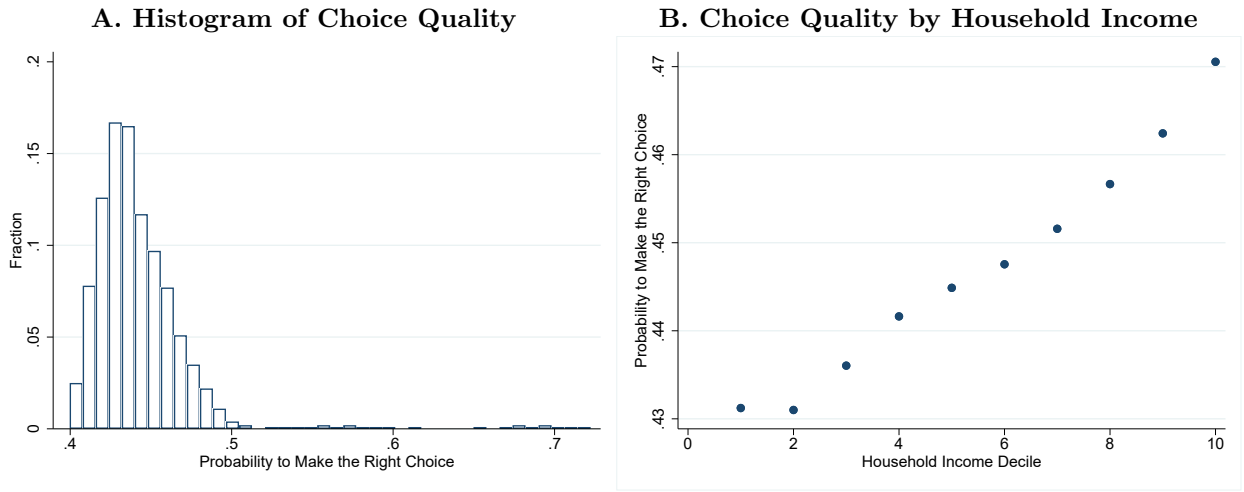
[38] Note that we use the average predicted risk for the different health deciles to calculate the certainty equivalents and to determine whether one should take up the deductible or not.

FIGURE D.1: PREDICTED DEDUCTIBLE CHOICE



Notes: This figure shows the distribution of predicted 500 EUR extra deductible take-up rate. Panel A shows the predicted 500 EUR deductible take-up with health effects, while Panel B shows the take-up without the health effects.

FIGURE D.2: HETEROGENEITY IN CHOICE QUALITY



Notes: Panel A shows the distribution of probabilities that consumers make the right deductible choice for a given set of socio-demographic characteristics $X_{i,t}$. The right choice is defined as the choice a rational consumer would make, as explained in Section III.A: to take the 500 EUR extra deductible if she expects her costs to be below 375 EUR with a probability larger than 0.5; to choose the low deductible otherwise. Individuals are binned in 1000 quantiles of choice quality; the variable displayed in this histogram is the binned average of the individual probability to make the right choice. Panel B shows the probability to make the right choice by income decile.

D.2 Models of Decision-Making: Rational Inattention

This model considers an agent who receives a noisy signal about her health, then decides whether or not to buy accurate information at price c_i. The noisy signal is defined as $\hat{\pi} = \pi + \epsilon$, where $\epsilon \sim N(0, \sigma_\epsilon)$, where we use $\sigma_\epsilon = \sigma_\pi$. The value of acquiring the accurate information depends on whether the information would change her deductible choice and thus on the condition density $f(\pi|\hat{\pi})$ for $\pi > .5$ and $\hat{\pi} < .5$ and vice versa. We simulate the conditional density taking random draws from the empirical distribution of π and the normal distribution of ϵ. We then group the resulting π and $\hat{\pi}$ in ten bins of length 0.1, indexing them from 1 to 10. Then for each bin

j of $\hat{\pi}$, we approximate the conditional density using:

$$p(\pi \in \pi_k | \hat{\pi} \in \hat{\pi}_j) = \frac{\#\text{individuals} \in \{\hat{\pi}_j \cap \pi_k\}}{\#\text{individuals} \in \hat{\pi}_j}$$

where π_k is bin k of π, and $\hat{\pi}_j$ is bin j of $\hat{\pi}$.

- Upon receiving a healthy signal ($\hat{\pi} > 0.5$), the agent only changes her decision when her true risk π is actually lower than 0.5. The expected loss from underinsuring when not acquiring the information equals:

$$L(\pi|\hat{\pi}) = \sum_{k=1}^{5} [-250 + (1 - \bar{\pi}_k) * 500] * p(\pi \in \pi_k | \hat{\pi} \in \hat{\pi}_j), \text{ for } \hat{\pi}_j = 6, 7, ...10.$$

  where $\bar{\pi}_k$ is the middle value of bin k of π (e.g., $\bar{\pi}_6 = .55$).

- Upon receiving a sick signal ($\hat{\pi} < 0.5$), the agent only changes her decision when her true risk is actually higher than 0.5. Her expected loss from overinsuring, given the noisy signal, is equal to

$$L(\pi|\hat{\pi}) = \sum_{k=6}^{10} [250 - (1 - \bar{\pi}_k) * 500] * p(\pi \in \pi_k | \hat{\pi} \in \hat{\pi}_j), \text{ for } \hat{\pi}_j = 1, 2, ...5.$$

Given that expected loss, the agent will decide to buy extra information if and only if the expected loss is greater than the information cost. Therefore the agent chooses the high deductible if and only if:

$$\hat{\pi} > 0.5 \text{ and } L(\pi|\hat{\pi}) < c_i \tag{11}$$

$$\hat{\pi} > 0.5 \text{ and } L(\pi|\hat{\pi}) > c_i \text{ and } \pi > 0.5 \tag{12}$$

$$\hat{\pi} \leq 0.5 \text{ and } L(\pi|\hat{\pi}) > c_i \text{ and } \pi > 0.5. \tag{13}$$

Panel E in Figure 13 shows the resulting take-up rates as a function of the true health risk π.

D.3 Counterfactual Policies

TABLE D.1: Counterfactual Policies, Controlling for Health Effects

	Optimal Deductible	High Deductible Only (875 EUR)	Low Deductible Only (375 EUR)
Risk Neutral			
Unweighted	63.7	-11.1	-5.3
Low Inequality Aversion	64.2	-10.6	-4.8
High Inequality Aversion	65.0	-9.8	-4.0
σ=.0001			
Unweighted	62.8	-12.8	-5.2
Low Inequality Aversion	63.2	-12.3	-4.7
High Inequality Aversion	64.0	-11.6	-3.9
σ=.001			
Unweighted	53.6	-28.6	-4.3
Low Inequality Aversion	53.9	-28.2	-3.9
High Inequality Aversion	54.5	-27.7	-3.3

Notes: Notes from Table 9 apply. This table performs the same exercise, except that each individual is attributed the population's health distribution, such that the correlation between income and health is controlled for.