Health Care Demand Among Low-Income Individuals

Angie Acquatella
Toulouse School of Economics *

October 4, 2024

Low-income individuals are typically the most price sensitive segment of the market, but this is not true in the market for health care. Using data from the RAND and Oregon experiments, I show that low-income individuals are less likely to *participate* in health care markets, relative to higher income counterparts, attenuating the average demand elasticity for this group. The key insight is that income effects may exclude low-income individuals from participating because, when marginal utility of consumption is high, forgoing non-medical consumption becomes too costly. These findings have implications for policy design and model specification.

Keywords: income effects, low-income health care demand, corner solution
JEL Codes: I12, I14, D11

*Contact: angelique.acquatella@tse-fr.eu. I am grateful to David Cutler, Marty Gaynor, Maggie Shi, Rena Conti, Jim Rebitzer, Matthew Zahn, Michael Ricks, Valentina Reig, Victoria Marone, Ben Vatter, Jean-Marie Lozachmeur, François Salanié, Christian Hellwig, Samuel Arenberg, Paolo Pertile, the seminar participants at the Toulouse School of Economics, Oxford, Verona, and Tilburg University for helpful comments, and to Lara Torka for her excellent research assistance.

1 Introduction

While low-income individuals are typically the most price sensitive segment of the market, this is not true in the market for health care. This is because this market has unique features that make low-income individuals less likely to *participate*, relative to their higher income counterparts. Low-income individuals use fewer health care resources, and recent literature suggests that the socioeconomic gradient in utilization would prevail even in the presence of generous public health care subsidies (Shepard, Baicker, and Skinner, 2020). I find that low-income individuals are also *less responsive* to health care prices than their higher income counterparts. This is because a disproportionate share consume zero health care, and survey evidence suggests that the main reasons involve *high perceived costs* of utilization.

This paper characterizes the health care demand profile of low-income individuals theoretically and empirically. Theoretically, accounting for income effects implies that it is optimal to consume no health care when income is low. The key insight is that forgoing non-medical consumption is costly when the marginal utility of (non-health) consumption is high. Empirically, I examine data from three different sources—the Medical Expenditure Panel Survey (MEPS), the RAND Health Insurance experiment, and the Oregon experiment—and find that low-income individuals (1) are less likely to use health care than their higher income counterparts, holding prices fixed; (2) have an attenuated response to price changes because a disproportionate share consume zero health care; and (3) perceive large utilization costs, even when out-of-pocket costs are zero.

Basic economics tells us that optimal health care consumption trades off the marginal benefit of care against the marginal cost of forgoing (non-health) consumption. If we account for income effects in health care demand, then the marginal costs of forgoing consumption will differ across individuals with different incomes. Even in cases where health care *is* highly valuable, individuals with very low incomes will find it costly to give up other consumption, e.g. an individual with chronic back pain will not forgo their monthly rent payment to go to physical therapy, even if they value pain reduction. In general, the cost of forgoing consumption corresponds to the relative price of care, weighted by the marginal utility of non-health consumption. If utility functions are concave, then low-income individuals may optimally consume no health care (i.e. the corner solution) because a high marginal utility of consumption means forgoing non-health consumption is too costly. A simple model suffices to show income effects may affect health care demand along both the extensive and intensive margin.

Indeed, evidence from the RAND experiment reveals that the reason low-income individuals appear less responsive to prices, on average, is because a disproportionate share consume zero health care. Low-income individuals who do choose to consume some positive amount of care are price sensitive, as in other markets. While dated, the RAND experiment allows us to study the relationship between income and health care demand without confounding supply-side factors that also affect low-income consumption of health care services (e.g. limited provider entry in low-income areas, insufficient Medicaid payments, or availability of unobserved alternatives such as

uncompensated or charitable care). These income effects in health *care* demand contrast the evidence on health *insurance* demand, where low-income individuals seem to be quite price sensitive to premiums (Krueger and Kuziemko, 2013; Cutler and Reber, 1998; Nyman, 2003).

Albeit a simple and foundational point, income effects in health care demand are largely ignored in empirical work, with one exception (Finkelstein, Hendren, and Luttmer, 2019).[1] These income effects are important for at least three reasons. First, ignoring income effects might lead to bias in the value of health care implied by observed demand. Second, if the regulator's objectives involve inducing some minimal level of health care utilization, then policy instruments which rely on demand-side incentives will not be very effective among the low-income population. Third, the distributional incidence of public health insurance policies might disproportionally favor the rich (McClellan and Skinner, 2006; Bhattacharya and Lakdawalla, 2006), and income effects highlight the importance of conditioning these policies on the basis of income.

The main contribution of this article is that it shows theoretically and empirically that income effects work primarily through the extensive margin channel. The price and income elasticities of health care demand have been extensively examined in the literature, and two large experiments stand out in this space. The RAND Health Insurance Experiment conducted by Newhouse and his team of researchers (1992) quantify the causal effect of price on health care spending, and they estimate a price elasticity of 0.2 (Manning et al., 1987). The Guaranteed Income experiment conducted by Miller et al. (2024) quantify the causal effect of income on health care spending, and their results suggest that the income elasticity of health care demand is small among low-income individuals (100% to 200% of the Federal Poverty Line). Acemoglu, Finkelstein, and Notowidigdo (2013) use income variation from oil price shocks and estimate an average income elasticity of 0.7 (among the entire income distribution). My article can be viewed as complementary in that it sheds light on the relationship between the price elasticity and income, and the key margins by which this relationship manifests.

This article proceeds as follows. Section 2, lays out a general model of health care spending and characterizes the tradeoffs that arise for individuals with different incomes (where the key insight comes from studying the corner solution). Section 3 empirically characterizes the demand profile of low-income individuals across three different empirical settings: the RAND Health Insurance Experiment, the Oregon Experiment, and the Medical Expenditure Panel Survey. The data allow me to characterize demand along the extensive and intensive margins across income groups. Section 4, discusses the implications of income effects for model misspecification and policy design. Section 5 concludes.

[1] Finkelstein, Hendren, and Luttmer (2019) quantify the welfare impact of Medicaid by estimating the willingness to pay (WTP) for the program. A key component of WTP is the marginal rate of substitution between health care and (non-health care) consumption; the authors account for the fact that the marginal utility of consumption is higher for low-income individuals by scaling down their calibrated marginal rate of substitution of health.

2 Tradeoffs in the Health Care Consumption Decision

To conceptually characterize the trade-offs faced by low-income individuals, I introduce a simple model of health care demand where utility is concave in consumption (which captures the idea that low-income individuals place a high value on cash). Individuals choose how much to spend on health care by trading off the marginal benefits of care against their marginal costs of forgoing consumption, which consist of their marginal out-of-pocket costs scaled by the marginal utility of non-medical consumption.

Intuitively, having a low level of income implies that forgoing other consumption is very costly: consider a low-income individual with chronic back pain. While she may unquestionably derive positive benefits from receiving care for pain management, forgoing her monthly rent payment may be too costly for her to choose any care. Her marginal cost of forgone consumption is large because her low level of income (consumption) puts her at the steepest part of her utility function. High marginal costs of forgone consumption imply that agents may not be able to equate marginal benefit to marginal cost, thus they may optimally choose to consume zero health care (i.e. the corner solution).

Even as insurance becomes more generous, the tradeoff for a low-income individual may hardly be altered. This is because marginal reductions in the out-of-pocket share are scaled by the (large) marginal utility of consumption; e.g., it is still not worthwhile to forgo the monthly rent payment when she has to pay 20% versus 15% of a physical therapy visit. Even as care becomes nominally free, the hassle costs associated with consuming health care (e.g., parking or taking the bus, taking time off from work, scheduling the appointment) get scaled up by the marginal utility of consumption, which implies that some low-income individuals may not be induced to consume any care when the price is zero.

For high-income individuals, however, the marginal utility of consumption is small and close to zero. This implies that they tradeoff (presumably positive) marginal health benefits of additional care against a very small marginal cost of forgone consumption and they will have higher levels of health care expenditures (Acemoglu, Finkelstein, and Notowidigdo, 2013). Moreover, the demand elasticity with respect to coinsurance for high-income individuals should be small: a $50 copay for chronic back pain management is essentially free for a high-income individual, and behavior is unlikely to differ with or without a copay.

2.1 A Model of Health Care Demand

Utility: An individual i derives utility over two goods: (non-medical) consumption and medical spending. The individual's utility function is:

$$u(x, m; l) = v(x) + H(m; l), \tag{1}$$

where where x is non-medical consumption, m is medical spending, and l denotes a health state. Utility of non-medical consumption is independent of the state, but utility derived from medical spending may vary with the state.

Assumptions: Utility is increasing and concave in non-medical consumption, $v'(x) \geq 0, v''(x) \leq 0$, utility from the first unit of non-medical consumption is infinite, $\lim_{x \to 0} v'(x) \to \infty$. Assume also that the marginal utility of medical spending is increasing and concave, $H_m(m; l) \geq 0$, $H_{mm}(m; l) \leq 0$, that the marginal utility from the first unit of medical spending is finite, $H_m(0; l) << \infty$, and that states l are ordered such that the marginal benefit of care is increasing in the state, $H_{ml}(m; l) \geq 0$.

Distribution of Incomes: Individuals choose (non-negative) medical and non-medical consumption to maximize their utility subject to their budget constraint, and each individual i is endowed with an income y_i. The distribution of incomes across individuals is given by $y_i \sim F(y_i)$, where $y_i \in [0, \overline{y})$.

Individuals make optimal choices about consumption and medical spending subject to a budget constraint. For individual i facing out-of-pocket costs $oop(m_i)$, the budget constraint is:

$$y_i \geq x_i + oop(m_i). \tag{2}$$

The optimal medical spending decision maximizes individual utility (1) subject to the budget constraint (2), a non-negative medical spending constraint $m \geq 0$, and a non-negative consumption constraint $c \geq 0$. The solution will either involve positive medical spending in which the consumer equates marginal benefit to marginal cost, or a corner solution of zero medical spending. The Kuhn-Tucker conditions tell us that the solution will satisfy either:

$$\begin{aligned} & & & & \overbrace{H'(m_i; l)}^{\text{marginal benefit of medical spending}} &= \overbrace{v'(x_i) \cdot oop'(m)}^{\text{marginal cost of forgone consumption}} & (3) \\ \text{or} \quad & m_i = 0 & \text{and} & & H'(m_i; l) &< v'(x_i) \cdot oop'(m). & (4) \end{aligned}$$

The optimality conditions reveal two important insights: first, that income plays a key role in determining the consumer's medical spending tradeoff via the marginal cost of forgone consumption. If our aim is to understand differential behavior across individuals with different incomes, we need to account for differences in the marginal utility of consumption across high and low-income individuals, and moving away from quasi-linear utility is a step in the right direction. Because the costs of forgoing consumption are scaled by the marginal utility of consumption, health care will appear more costly to a low-income individual, relative to a high-income individual.

Second, the zero medical spending choice can be informative about economic primitives. To illustrate, consider a setting where the marginal benefit of care is homogeneous across individuals (e.g. insulin for diabetics), but the marginal costs of forgoing consumption are heterogeneous (due to income and concavity differences). If we observe income and are willing to assume individuals

are making rational decisions, then the share of individuals who do not use any care will be informative about the marginal cost of forgone consumption, and thus about concavity of the utility function with respect to non-medical consumption. Moreover, if extensive margin utilization is monotonic in income such that there is a single income threshold above which individuals do not use any care, then there will be a unique concavity parameter that rationalizes observed choices.

2.2 Equilibrium Medical Spending for High- versus Low-income Individuals

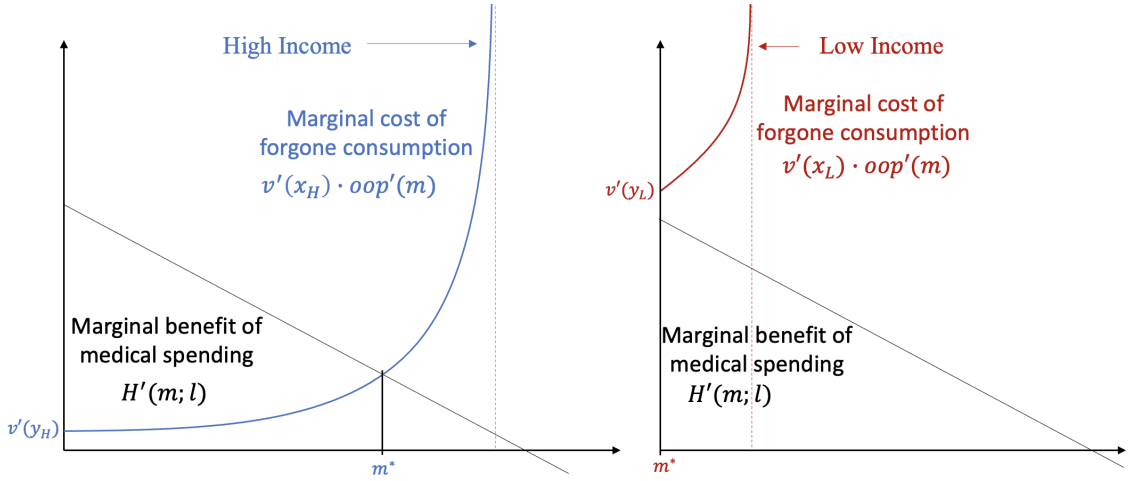
Figure 1 illustrates the tradeoffs in the medical spending decision for a high- versus low-income individual. Because medical spending reduces consumption, the marginal cost of forgone consumption is exactly the reflection of the marginal utility of consumption. The costs of forgoing consumption asymptote to infinity as the out-of-pocket costs approach income (and consumption approaches zero). This tradeoff is illustrated in Panel (a) of Figure 1. When medical spending is zero, the costs of forgoing consumption are equal to the marginal utility of consumption for an individual who consumes the entirety of their income. Thus, when utility is concave and the marginal benefit of care is bounded at zero, we should not expect to see low-income individuals utilizing any health care.

How does insurance (i.e. reduction in the marginal out-of-pocket price) change the tradeoffs for low-income individuals? In general, lowering the out-of-pocket costs will increase health care utilization along the extensive margin because the likelihood that the marginal benefit of care exceeds the marginal costs of forgone consumption will increase as out-of-pocket costs are reduced. However, the number of individuals who switch from consuming no care to consuming some positive amount of care will depend on the relative magnitude of the costs of forgoing consumption. If the marginal utility of consumption is large, reductions in the out-of-pocket cost will still be scaled up by a very large number and may not change behavior at all. This point is illustrated in Panel (b) of Figure 1.

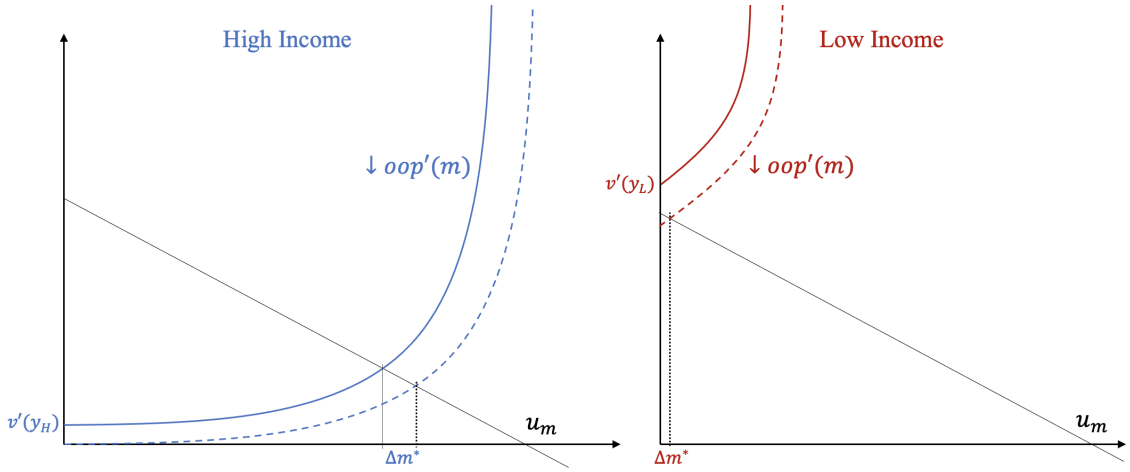
Implicitly, these arguments about low-income health care demand hold because I have assumed that the marginal utility of medical care is positive, but finite and bounded at zero. There are a number of health care services for which this assumption is reasonable (e.g. physical therapy, cataract surgery, mental health counseling, post-acute care, etc). There are also a number of health care services where this assumption is *not* reasonable, particularly acute care. While my empirical analysis does not distinguish across the two different types of care, it is important to maintain that income effects are likely prevalent among non-emergency health care services that involve consumer discretion, and unlikely for acute care.

The model of health care demand with concave consumption utility has two main empirical predictions. First, when health care benefits are positive but finite, low-income individuals should be less likely to consume any health care services, relative to their higher income counterparts. Second, the average response to price reductions should be smaller among low-income individuals

Figure 1: Optimal Medical Spending Choice for High v. Low-Income



(a) Determinants of Medical Spending



(b) Demand Response to Change in Marginal Out-of-Pocket Price

Notes: This figure shows the marginal benefit curve of medical spending and the marginal cost of forgone consumption for a high- and low-income individual, respectively. Both individuals derive the same marginal benefits from medical spending, but differ in their levels of income and thus their costs of forgoing consumption are different. In all panels, the marginal cost of forgoing consumption at $m = 0$ is equal to the marginal utility of consuming the entirety of income, and the cost asymptotes to infinity as $oop(m) \to y$. The panels on the left show in blue the costs for a high-income individual, and the panels on the right shows in red the costs for a low-income individual. Panel (a) on the top shows the optimal choice of medical spending for each individual, and Panel (b) shows the effect of a reduction in the marginal out-of-pocket price of medical care on the marginal cost of forgone consumption.

when these reductions fail to induce individuals to consume some positive amount of care. The (small) average response for low-income individuals will be driven by the inframarginal individuals for whom the marginal costs of forgoing consumption are large.

3 Empirical Differences in Health Care Demand Across Low- and High-Income Individuals

The main focus of my empirical analysis is to characterize the health care demand profile of low-income individuals, relative to that of higher income counterparts. I study the relationship between income and health care demand across three different empirical settings: one large cross-sectional database (Medical Expenditure Panel Survey), the RAND health insurance experiment, and the Oregon health insurance experiment. The three key findings are that, relative to their higher income counterparts, low-income individuals are the least likely use any health care services, less (or equally) responsive to changes in health care prices, and perceive health care as prohibitively costly even when covered by full insurance (Medicaid).

Identifying income effects in health care demand is challenging for at least two reasons. First, suppliers of health care services may be less willing to provide their services in low-income neighborhoods, introducing correlation between income and unobserved heterogeneity in access to health care services. Second, if high- and low-income individuals have different insurance plans, they may also face different prices when making health care utilization decisions. In light of these concerns, the RAND and Oregon experiments provide ideal empirical settings to study the health care demand profile of low-income individuals because these potentially confounding factors are held fixed. However, the samples of participants in each experiment are either dated, or not nationally representative. For this reason, I supplement my analysis with data from MEPS. While each empirical setting has its limitations, I find a consistent empirical narrative supporting the claim that income effects play a central role in health care demand, particularly through the extensive margin.

3.1 Low-income Individuals are *Less Likely to Use* Health Care Services

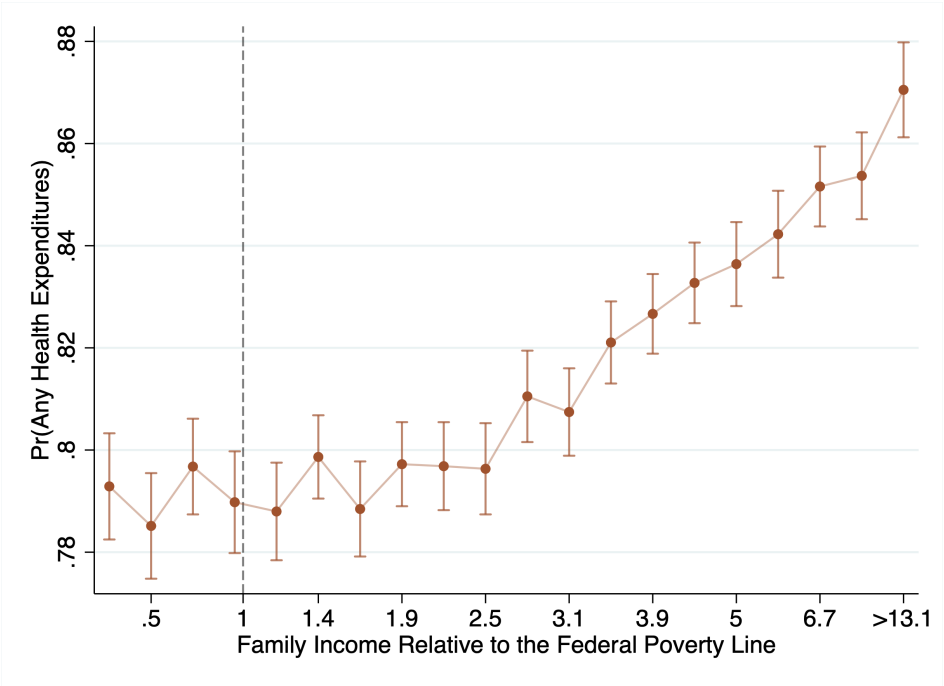
To begin, I provide descriptive evidence about extensive margin health care utilization choices across income groups using data from the Medical Expenditure Panel Survey (MEPS). The MEPS data are public survey data that contain information about medical expenditures and health conditions of individuals. A key advantage of the MEPS data is that it is recent, contains a large number of individuals, and rich health information for each individual.

Data: The data are provided by the Agency for Healthcare Research and Quality, and consist of annually collected, repeated cross-sections, where individuals are followed over the course of two years. I observe the individual's chronic conditions, health history, full disease profile at the International Classification of Diseases (ICD) Code level, impairments to physical and cognitive functioning, and detailed self-reported health questionnaires. This information allows me to (at least partially) control for heterogeneity in underlying health status, and analyze differential utilization while holding observable health characteristics fixed.

Sample construction: I use surveys collected between 2007 to 2020, which contain reliable family income information, as well as information on age, sex, race, education, employment, the

individual's health insurance type (private, public, or none), chronic conditions, health history, and impairments to physical and cognitive functioning.[2] Children and retirees are dropped from the sample, resulting in a final sample of 156,951 unique individuals. The data provider provides survey sampling weights to account for non-random sampling. I assign individuals to income ventiles at the yearly level and based on their family income as a percentage of the Federal Poverty Line (FPL), which adjusts for household size.

Figure 2: Probability of Having Any Health Care Expenditures by Income Ventile



Notes: Graph shows the predicted probability that an individual will consume a strictly positive amount on health care in a given year, conditional on their income. The probability is estimated in a probit regression where the dependent variable is a binary indicator for whether the individual consumed any care, and the independent variables are 20 categorical dummies for the individual's income ventile plus controls. Regressions are estimated at the person-year level and controls include insurance type (private, public, or none), age, sex, age squared, age by sex interactions, race, education, smoking habits, chronic conditions, impairments to physical and cognitive functioning, year fixed effects, and survey panel fixed effects. Standard errors are adjusted to account for survey design with non-random sampling. N=265,112

Results: An individual in the highest income ventile is 8 percentage points more likely to use any health care services than an individual in the lowest income ventile, controlling for observable characteristics about health and demographics. Figure 2 shows the estimated probability that an individual uses some positive amount of health care services conditional on their income, controlling for observables, and the dashed vertical line denotes the FPL. Holding observable health characteristics fixed, high-income individuals are more likely to use health care services than their low-income counterparts. See Appendix for additional specifications that describe the relationship between income, health care utilization, and health status.

[2]Chronic conditions and health history include indicator dummies for congestive heart failure, myocardial infarction (heart attack), peripheral vascular disease, cerebrovascular disease, chronic obstructive pulmonary disease, peptic ulcer, rheumatoid arthritis, diabetes, paralysis, skin malignancy, liver disease, cancer, and HIV/AIDS.

One important caveat in the interpretation of Figure 2 is that it describes the *equilibrium* relationship between health care consumption and income, which is a product of both demand and supply side factors. Thus, while the patterns in the data are consistent with income effects in health care demand, they could also be explained by supply side confounding factors that are correlated with income, such as limited provider entry in low-income neighborhoods. Moreover, individuals in this data face different prices for health care services, and the information about individual health insurance plans is coarse. The following two empirical settings are well suited to overcome these concerns.

3.2 Low-income Individuals are *Less Responsive* to Health Care Prices

Next, I characterize the health care demand *response* separately by income terciles using data from the RAND Health Insurance Experiment. The three takeaways from this analysis are that, relative to their higher income counterparts, low-income individuals: (1) are less likely to use any health care services, holding prices fixed; (2) are less price responsive at low levels of medical spending; and (3) do not spend more on health care when it is free versus when there is cost-sharing.

Using the exogenous variation in cost-sharing assigned by the RAND experiment, I estimate three objects related to income effects in health care demand: (1) extensive margin utilization by income tercile, holding prices fixed; (2) the average treatment effect of price on total health care spending, conditional on income tercile; and (3) the IV quantile treatment effects of price on total health care spending, using RAND health insurance plan assignment (random) as an instrument for marginal health care price at the end of the calendar year (endogenous).

While dated, the RAND experiment allows us to study the relationship between income and health care demand without confounding supply-side factors that also affect low-income consumption of health care services (e.g. limited provider entry in low-income areas, insufficient Medicaid payments, or availability of unobserved alternatives such as uncompensated or charitable care).[3] The RAND experiment thus provides a unique empirical setting to isolate income effects in the demand for health care services: a large-scale randomized control trial where individuals across a broad range of the income distribution were randomly assigned to health insurance plans. A key limitation from this data is that the sample size is relatively small.

3.2.1 RAND Health Insurance Experiment Overview and Data

RAND Experiment Overview: The experiment randomly assigned individuals to health insurance plans with varying levels of cost-sharing. It ran from 1974 to 1981, and enrolled approximately 1,900 families and 5,800 individuals. Individuals were assigned to one of six plans[4] with varying

[3]The RAND investigators paid health care providers for their submitted charges, which means providers did not have an (insurer driven) incentive to under-treat or dismiss patients. Moreover, the RAND experiment was conducted across six urban areas where physical access health care providers was similar across individuals with varying incomes.

[4]Six plans refers to six categories of cost-sharing which applied to either all services or sub-categories of services (e.g. dental or outpatient psychiatric care). Four plans had either 95%, 50%, 25%, or 0% coinsurance for all services;

levels of cost-sharing: 95%, 50%, 25%, or 0%. The experiment ran in six different locations, and assignment was random conditional on location and starting month. The investigators oversampled low-income families in the design: 17% of families in the experiment had incomes below the FPL, which was 6 percentage points greater than the national share of families below the poverty line in 1973.

The RAND Health Insurance Experiment data has the advantage of spanning a broad range of incomes: as a percent of the FPL, the poorest RAND family earned 19% FPL, while the richest family earned 11,200% FPL. In dollars and normalized to a family of four, the poorest family earned $860 while the richest family earned $25,000 in 1973 ($6,000 and $176,000 in 2023 dollars, respectively). The investigators collected income data in the two years prior to enrollment, and excluded individuals with incomes greater than $25,000 in 1973 dollars (which corresponds to $175,00 in 2023 dollars). Because the experiment was carried out in six different locations, income data was collected on different years. To discretize individuals into income terciles using comparable income data, I follow the RAND investigators and use inflation-adjusted family income in 1973 dollars, normalized by family size.

As a percentage of the FPL in 1973, the income cutoff for the bottom tercile was 203% FPL, for the middle tercile 336% FPL, and the top tercile had incomes ranging from 336% to 11,200% FPL. On average, individuals in the bottom income tercile are younger, have lower levels of employment and schooling, and larger families, relative to the middle and top income terciles.[5] Table 1 shows average spending levels across income terciles and across plans. For the lowest income tercile, medical spending levels were almost always smaller in magnitude than the spending levels of their higher income counterparts. The only exceptions were in the mixed coinsurance plan and the 95% coinsurance plan, where the middle income tercile had lower spending levels than the bottom tercile.

Data: The data have been originally published in the public domain by Newhouse and his team of researchers, and carefully revisited in Aron-Dine, Einav, and Finkelstein (2013). For my analyses, I use the data processed and made available by Aron-Dine, Einav, and Finkelstein (2013).

3.2.2 Effect of Cost-Sharing on Health Care Utilization and Spending Across Income Groups

Using the experimental variation in cost-sharing from the RAND experiment, I describe the extensive and intensive margin treatment effects of price (cost-sharing) on health care demand for different income terciles. The estimates are simple differences in means across the different RAND plans, estimated separately within each income tercile and with standard errors clustered at the family level. Results are shown in Figure 3, with the extensive margin differences shown in Panel

the "mixed" coinsurance plan had 50% cost sharing for mental health and dental care, and 25% for everything else; and the "individual deductible" plan had 0% coinsurance for inpatient services, and 95% coinsurance for outpatient services. Formally, there were twenty four unique plans which varied along two dimensions: cost-sharing and out-of-pocket maximum (referred to as maximum dollar expenditure, or MDE, limit by the experimenters).

[5]Detailed summary statistics about demographic characteristics of individuals across income terciles can be found in Appendix Table A.1.

Table 1: Average Spending by Income and Plan

	Overall	*Income Tercile* (1)	(2)	(3)
	Share	Average Spending (1973 $)		
0% Coinsurance	.34	$1,819	$2,294	$ 2,489
(Free Care)		(4,421)	(5,565)	(5,178)
25% Coinsurance	.12	1,392	1,488	1,806
		(3,764)	(5,321)	(4,004)
50% Coinsurance	.07	1,445	2,052	1,450
		(9,143)	(14,841)	(3,028)
95% Coinsurance	.18	1,326	1,008	1,756
		(4,916)	(2,525)	(4,849)
Mixed Coinsurance	.08	1,763	1,496	2,494
		(5,434)	(4,595)	(7,015)
Individual Deductible	.21	1,151	2,002	1,967
		(3,180)	(6,629)	(4,821)
Number of families	1,908	636	636	636

Notes: Table shows the shares of individuals enrolled across plans, and average individual spending by income tercile within each plan. Individuals were assigned to terciles at the family level. As a percentage of the FPL in 1973, the income cutoff for the bottom tercile was 203% FPL, for the middle tercile 336% FPL, and the top tercile had incomes ranging from 336% to 11,200% FPL. The 50% cost-sharing plan was discontinued half-way through the experiment, which explains why standard errors are larger for this plan.

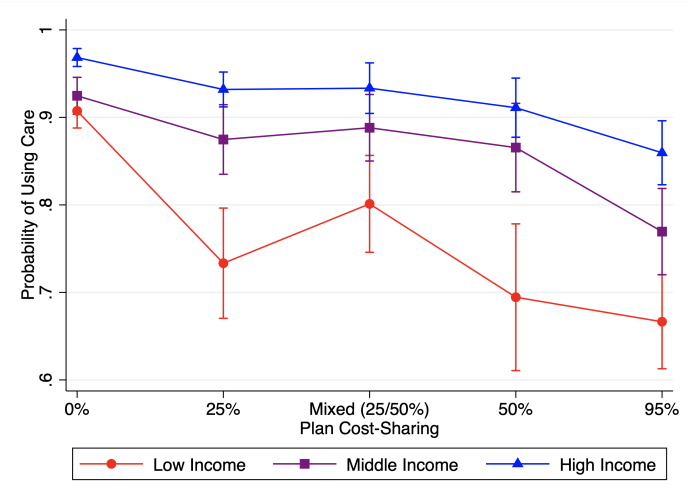
(a), and the intensive margin differences shown in Panels (b) and (c).

Regarding the extensive margin, individuals in the bottom income tercile (shown in red) were less likely to use any health care services than their higher income counterparts (purple and blue) when assigned to RAND plans with strictly positive cost-sharing. Panel (a) shows that the probability of using health care for low-income individuals within a given plan was always below the probabilities of the middle and highest income terciles, with the exception of the free care plan.

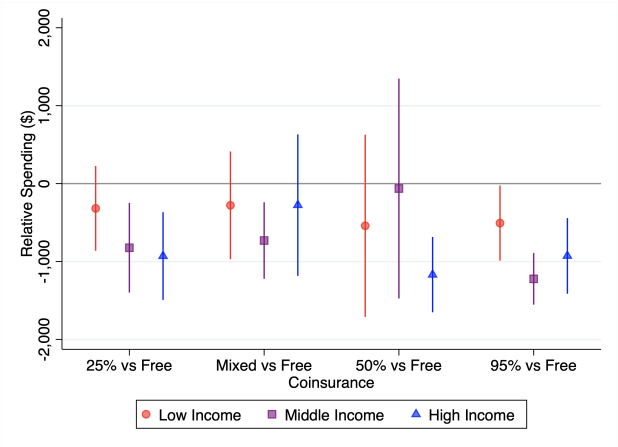
Regarding the intensive margin, low-income individuals had similar levels of yearly health care spending in the free care plan and in plans with strictly positive cost-sharing. Panel (b) shows the average difference in total health care spending across the free care plan and the 25%, 50%, and 95% cost-sharing plans. Estimates suggest that, relative to the free care plan, average yearly health care spending was only significantly lower among individuals in the middle and highest income terciles, but not among individuals in the bottom income terciles. However, average spending for the lowest income terciles is attenuated by the fact that low-income individuals were less likely to use health care in the first place, as seen in Panel (a). Nonetheless, Panel (c) shows that, even conditional on having any positive health care spending in a given year, low-income individuals did not have significant differences in spending across plans.

Taken together, Table 1 and Figure 3 show that, relative to their higher income counterparts, individuals in the lowest income tercile had lower levels of average spending, were less likely to use health care services (holding prices fixed), and had similar spending levels across plans with and without cost-sharing (conditional on utilization).

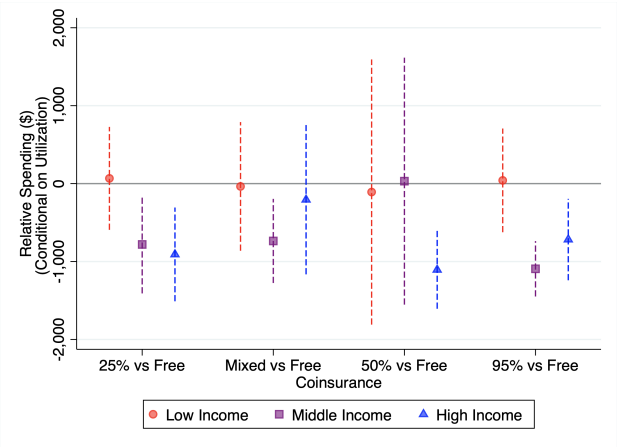
Figure 3: Health Care Utilization by RAND Plan Across Income Groups



(a) Extensive Margin Utilization by Plan



(b) Average Spending Across Plans



(c) Conditional Spending Across Plans

Notes: Panel (a) shows the predicted share of individuals who spent zero on health care over the RAND experiment period. The probability that an individual uses some positive amount of health care services conditional on income is estimated in a regression where the dependent variable is a binary indicator for whether the individual consumed non-zero care, and the independent variables are RAND plan group dummies and year plus location by month fixed effects. N=16,028. Panel (b) repeats the same regression but with inflation-adjusted spending as the dependent variable. The figure plots the regression coefficients of the plan dummies. The omitted category is the free care plan. N=16,028. Panel (c) is identical to (b), but conditions on the subsample of individuals with strictly positive health care spending within a given year. N=13,902. Regressions in (b) and (c) are estimated separately by income tercile subsample. The individual deductible plan is excluded from the analysis. Standard errors are clustered at the family level. All regressions are estimated at the person-year level.

3.2.3 From Plan Effect to Elasticities

From the analyses thus far, it is clear that price affects both the extensive and intensive margin in different ways across different income groups. To investigate the extensive and intensive margin effects jointly within a single empirical framework, I estimate quantile regressions of the health care demand response for each income group. I find that low-income individuals are less responsive to price at the bottom quantiles of the medical spending distribution, relative to the higher income counterparts.

While estimating an elasticity is the usual way to quantify a demand response, the extent to which one can transform the average treatment effects of the RAND experiment into elasticities is limited for two reasons: frequent zeroes in the outcome (spending) and the non-linearity of health insurance contracts (Aron-Dine, Einav, and Finkelstein 2013; Mullahy and Norton, 2022; Chen and Roth, 2023). When a substantial share of individuals consume zero health care, the concept of an elasticity is not well defined because the derivative of health care demand with respect to price is not well defined when either demand or price are equal to zero (Chen and Roth, 2023; Mullahy and Norton, 2022). When health insurance contracts are non-linear, the researcher has to make assumptions about which "price" is relevant for the intensive margin of health care demand, given that the marginal price changes throughout the year as individuals incur health expenditures (Aron-Dine, Einav, and Finkelstein, 2013; Lin and Sacks, 2019).

Health insurance plans in the RAND experiment were non-linear because of the expenditure limit feature, which reduced cost-sharing to zero after some threshold level of spending. This implies that, while the initial decision to use any health care services may have responded directly to the initial level of cost-sharing (exogenously assigned by the RAND plan), the marginal price of care varied throughout the year as the individual (endogenously) incurred health expenditures. Depending on the researcher's assumptions about consumer forward-looking behavior, the relevant price could be the realized marginal price at the end of the year, some weighted average of the price paid throughout the year, or the current "spot" price of care. Recent evidence by Lin and Sacks (2019) motivates the assumption that an individual's yearly spending responds to the *marginal price* at the end of the year: they found that individuals in the RAND experiment who hit their expenditure limit in a given year front-loaded their future health expenditures to that same calendar year.

For my main specification, I estimate quantiles of the (annual) health care demand response under the assumption that individuals respond to the *marginal* price, and include quantile estimates under the two alternative assumptions in the Appendix. Given that the marginal end-of-year price is endogenous to the individual's spending decisions throughout the year, I instrument for the end-of-year price by using initial (randomized) plan assignment, following Lin and Sacks (2019). The instrument leverages the fact that individuals who got assigned to plans with higher initial coinsurance (e.g. 95%) were more likely to hit their out-of-pocket maximum at the end of the year, resulting in a zero marginal price. Because a substantial share of individuals got assigned to the

free care plan (zero price), I estimate the semi-elasticity of health care spending with respect to price (cost-sharing), as opposed to the usual log-log elasticity, in order to avoid dropping individuals assigned to the free care plan.

Empirical Framework: For each income tercile, I estimate the semi-elasticity of health care spending with respect to cost-sharing using a smoothed quantile instrumental variable estimator (SIVQR), following Chernozhukov and Hansen (2005) and Kaplan and Sun (2017). The SIVQR estimator enables recovering the differential impact of price changes on spending along the quantiles of the full medical spending distribution within an instrumental variables framework. The SIVQR estimator relies on four identifying assumptions: (1) that potential outcomes can be expressed as quantile functions of cost-sharing, additional covariates, and uniformly distributed unobserved heterogeneity; (2) that conditional on the instrument, the conditional distribution of unobserved heterogeneity has a τ-quantile equal to zero; (3) that quantile functions are strictly increasing and left-continuous in the unobserved heterogeneity; (4) that there are no systematic differences in individuals' ranks across treatments (i.e. rank similarity is satisfied); and (5) that the instrument is exogenous conditional on covariates (i.e. the selection equation is satisfied).[6]

As mentioned above, I instrument for the end-of-year price by using initial (randomized) plan assignment. Denoting individuals by i, calendar year by t, month by μ, location by l, and the outcome variable (log medical spending) by y, the resulting objective function for estimating β (the semi-elasticity) can be written as:

$$q_n(\tau, \beta, \gamma) = \frac{1}{n} \sum_{i=1}^{n} \rho_\tau \left(y_{it} - {p_{it}}' \beta - \mathbf{z}_i' \gamma - \lambda_t - \lambda_{l\mu} \right) \tag{5}$$

where p_{it} is the marginal price of care at the end of the year, $\mathbf{z}_i$ the vector of initial RAND plan assignment dummies, λ_t and $\lambda_{l\mu}$ are year and location by month fixed effects, and ρ_τ denotes the estimator's weight on quantile τ.[7] I estimate this objective function for each income tercile, smoothing the underlying moment condition as suggested in Kaplan and Sun (2017). Our estimate of interest $\hat{\beta}_\tau$ solves the smoothed estimating equations:

$$\mathbf{0} = \frac{1}{n} \sum_{i=1}^{n} \mathbf{z}_i \left[\tilde{I} \left(\frac{y_{it} - p_{it}' \beta_\tau - \lambda_t - \lambda_{l\mu}}{h} \right) - \tau \right], \tag{6}$$

where $\tilde{I}(\cdot)$ smoothly decreases from 1 to 0.

Figure 4 displays the results.[8] The main takeaway from the IV quantile regression estimates is that low-income individuals at lower quantiles of the medical spending distribution are less responsive to price, relative to their higher income counterparts. Among the lowest income tercile, there is

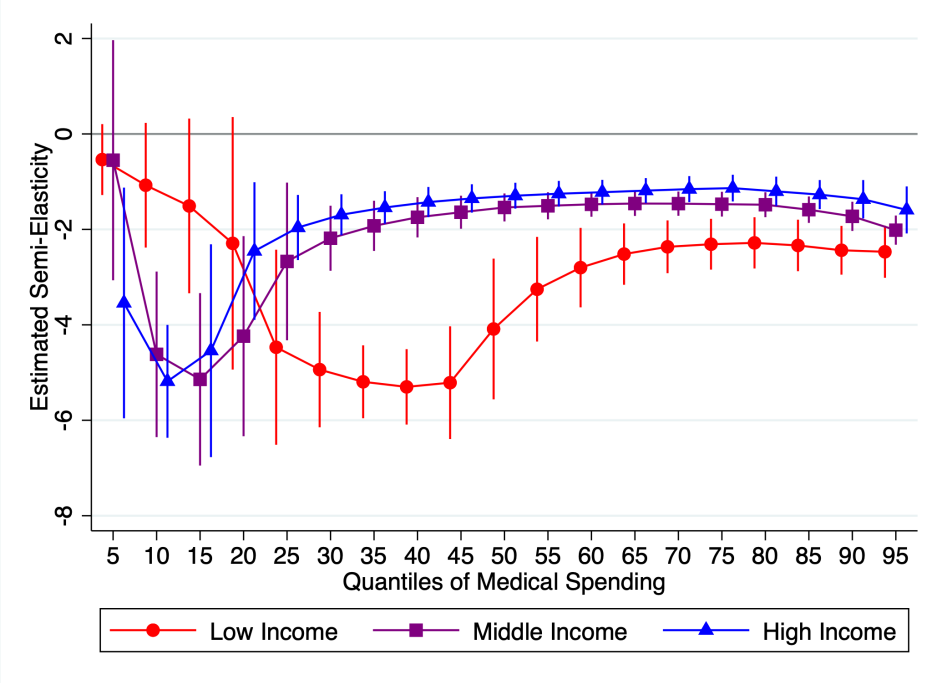
[6] Chernozhukov and Hansen (2005) and Kaplan and Sun (2017) show that quantile treatment effects can reveal causal effects of treatment when unobserved heterogeneity, in terms of covariates and treatment status, is constant at each quantile.

[7] Formally, $\rho_\tau \equiv \tau u^+ + (1-\tau) u^-$, where u^+ and u^- correspond to the unobserved heterogeneity above and below quantile τ, respectively.

[8] The full table of estimates is included in Appendix Table A.2.

a larger share of inframarginal individuals with zero medical spending (corner solution), but those who do incur medical expenditures (interior region) are as price sensitive as their higher income counterparts: the semi-elasticity estimates between the 30th and 45th medical spending quantiles for the lowest income tercile are similar to the semi-elasticity estimates between the 5th and 20th medical spending quantiles for the middle and highest income terciles.

Figure 4: IV Quantile Regression: Semi-Elasticities by Income Tercile



Notes: IV quantile regressions are estimated separately within each income tercile using the SIVQR smoothed estimation procedure in Kaplan and Sun (2017). Standard errors are block bootstrapped at the family level based on 500 replications. The dependent variable in each of these quantile regressions is log(1+medical spending). The independent variables are the continuous coinsurance rate associated with the RAND plan (e.g. sticker price), and year plus location by month fixed effects. N=16,028.

Potential Threats to Identification: Health insurance plans in the RAND experiment had out-of-pocket expenditure limits of either 5, 10, or 15% of income, or $1,000 (whichever amount was smaller), which may raise two potential concerns. First, this plan feature would invalidate the exclusion restriction (assumption 2) if expenditure limits were systematically correlated to the quantiles of the medical spending distribution in different ways for the different income terciles. However, the average expenditure limit was constant along the medical spending distribution, suggesting that plan incentives were not systematically different along different quantiles of the medical spending distribution within each tercile (see Appendix Figure A.3, Panel a). Second, this plan feature meant that low-income individuals were systematically more likely to have lower expenditure limits than their higher income counterparts. However, lower-income individuals were not significantly more likely to hit their expenditure limits for most plans, with the single exception of the 25% cost-sharing plan, suggesting that the expenditure limits are unlikely to be driving differential behavior at the lower quantiles of the medical spending distribution (see Appendix Figure A.3, Panel

b). Lastly, while this feature may be driving spending differences among high spenders, it could not explain differential behavior at the bottom quantiles of the income distribution, where price corresponded to the initial randomly assigned cost-sharing.

Robustness: When estimates involve a log-like transformation of the outcome variable (medical spending), the magnitude of the estimates will be sensitive to the specific choice of transformation (Chen and Roth, 2023). I re-estimate the IV quantile regressions under two alternative transformations of the outcome variable. While the magnitude of the point-estimates does differ across alternative specifications, the overall shape of the quantile semi-elasticity estimates is reassuringly similar (see Appendix Figure A.4). Quantile regression estimates of the semi-elasticity under two alternative assumptions about "price" (initial spot price of care or average out-of-pocket share) also look similar to the main IV quantile regression estimates (see Appendix Figure A.5).

Alternative empirical frameworks to quantify the demand response also paint a similar picture. For instance, one could summarize the average effect of a price change on health care spending by computing pairwise elasticities, which are always defined because these rely on the point-slope formula across two points. Alternatively, one could estimate a Poisson regression, where estimates can be directly interpreted as average percentage differences in spending across plans. Across both of these alternative empirical specifications, low-income individuals appear less responsive to the price of health care than their higher income counterparts. See Appendix Figure A.6 and Appendix Figure A.7 for pairwise elasticities and Poisson regressions estimates, respectively.

In sum, evidence from the RAND experiment suggests that the health care demand response to cost-sharing among the lowest income tercile differs from that of higher income counterparts in three ways. First, given the same insurance contract, low-income individuals are less likely to use health care services along the extensive margin, relative to their higher income counterparts. Second, among low-income individuals who did use health care services, average spending was similar across the free care plan and plans with positive cost-sharing. Third, the combined extensive and intensive margin effects estimated in a quantile regression framework suggest that low-income individuals are less responsive to price at the bottom quantiles of the medical spending distribution.

This might seem surprising initially because, while low-income individuals tend to be the most price responsive segment of the market for most goods, they appear to be the most reluctant group to use any health care services at all. However, it is costly for individuals with high marginal utility of consumption to forgo non-health care consumption. Thus, one can rationalize these empirical patterns in a framework where income effects induce some individuals to choose the corner solution of no health care. Note also that these findings are consistent with prior literature. Chandra, Gruber, and McKnight (2014) estimate an elasticity of -0.16 for a low-income population, and Lavetti, DeLeire, and Ziebarth (2023) estimate an elasticity of -0.1 for low-income enrollees in the Affordable Care Act marketplaces. Both estimates are within the range of the confidence interval for the -0.2 estimate from the RAND experiment (Manning et al., 1987), yet it is nonetheless surprising that the elasticity is not larger than the average. Brot-Goldberg et al. (2017) similarly find an income gradient in empirical responses to cost-sharing, where the average response among

middle-income individuals is smaller than the response of higher income individuals.

3.3 Low-income Individuals are *Likely to Say* They Didn't Use Health Care Because it Cost Too Much

The extent to which we can extrapolate conclusions from the RAND experiment to behavior of low-income individuals today may be limited because the experiment took place nearly five decades ago. For that reason, I also turn to more recent evidence from the Oregon Health Insurance Experiment, which provided low-income individuals the opportunity to enroll in free health insurance (Medicaid) in 2008. The main takeaway from this analysis is that, among individuals who reported having unmet health care needs (i.e. needed health care and did not receive it), the most frequently cited reason was that health care *costs were too high*. This remains true even among the set of individuals who enrolled in Medicaid as a result of the Oregon lottery.[9]

3.3.1 Oregon Experiment Overview and Data

Oregon Experiment Overview: In 2008, Oregon used a lottery to give low-income individuals the possibility to enroll in Medicaid. Lottery winners were drawn from a waiting list of people who signed up and were given the opportunity to enroll themselves and any household member, provided they met the eligibility criteria. The Medicaid program in the state of Oregon is called the Oregon Health Plan (OHP), and is composed of two distinct programs: OHP Standard and OHP Plus. The lottery allowed enrollment in OHP Standard for low-income individuals who were non-elderly, non-disabled, and non-pregnant who were not otherwise eligible for OHP Plus.

Actual enrollment among lottery winners into the Medicaid (OHP Standard) program was notoriously low, with under 30% of lottery winners actually enrolling in OHP Standard (Allen et al., 2010). Among those who did enroll in response to winning the lottery, there was substantial attrition in insurance coverage because enrolled individuals had to recertify their eligibility every six months, resulting in a non-trivial share of lottery winners (and treatment compliers) who did not have insurance for every month of the year (Finkelstein et al., 2012). Lottery losers (control group) were also able to get insurance through OHP Plus or other sources in the year following the lottery (Finkelstein et al., 2012).

Data: For my analyses, I use the administrative data on lottery winners and losers and the mail survey data collected by Finkelstein et al. (2012), both of which are publicly available on the National Bureau of Economic Research Oregon data repository. The administrative data contain individual information on lottery outcomes, whether they submitted an application for OHP standard, if their application was approved and on what date, their household size, birth year, and gender. The mail survey was conducted in the first year after the experiment and in three waves: upon notification of the lottery outcome, six months later, and one year later. The data contain

[9] Oregon's Medicaid program has zero cost-sharing for the set of covered services, but beneficiaries face the full costs of care for uncovered services.

a rich set of outcomes at the individual level, including health care utilization measures, whether they were enrolled in any form of health insurance (including private, employer sponsored, or Medicare), the number of months for which they had an active insurance policy (out of the previous six months), whether they had medical or prescription needs that went unmet, whether they had particular chronic diseases, self-reported health, household income, and demographic information.

There are two main limitations with the Oregon survey data. First, it is known that individuals overstated inpatient hospital use in the survey data, relative to observed hospital use in the administrative data (Finkelstein et al., 2012). Second, lottery winners who did enroll in Medicaid were adversely selected, with worse health than the general waiting list population, and higher levels of pre-lottery emergency care utilization (Kowalski, 2023). It is also the case that individuals who signed up for the initial waiting list were older and had worse health, relative to the general low-income population (Allen et al., 2010). Therefore, the utilization measures in the survey data may be biased upward due to reporting error, selection, or both.

3.3.2 Unmet Health Care Needs and Reasons Why

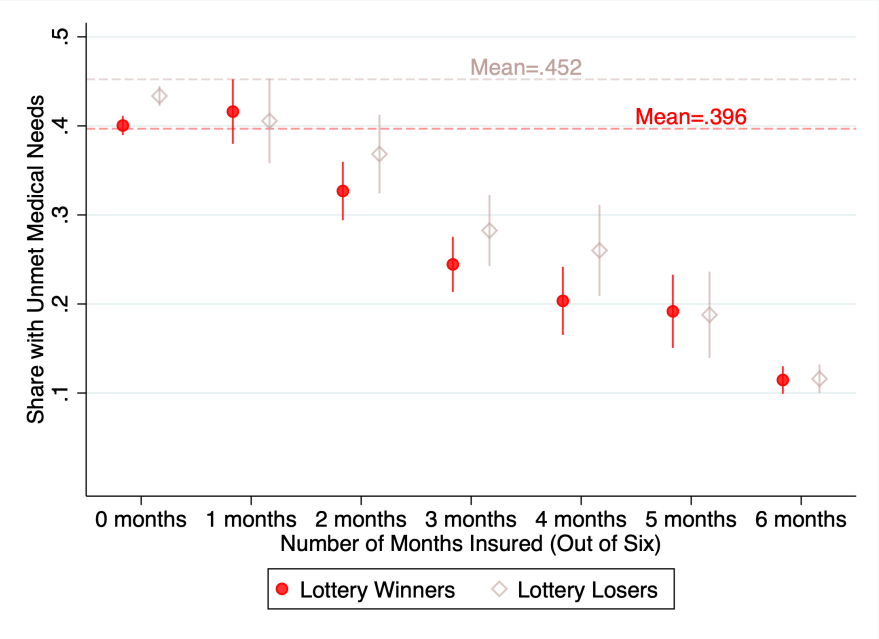
Unmet Health Care Needs: Individuals were asked to report whether they had unmet medical or prescription needs in the previous six months. Given that not all lottery winners who enrolled in Medicaid recertified their eligibility, and that lottery losers had other avenues to get insurance, there is substantial variation in the number of months for which an individual was actually insured, both within the set of lottery winners and the set of lottery losers. Figure 5 shows the share of individuals who reported having unmet medical needs, by lottery outcome and number of months for which they had an active insurance policy.

From Figure 5, it is evident that the likelihood that an individual reports having unmet health care needs depends more on whether they had an active insurance policy than it does on whether they won the Oregon lottery. Nonetheless, lottery winners were on average less likely to report having unmet needs than lottery losers; 27% of lottery winners reported having unmet medical needs, versus 46% of lottery losers.

Next, I turn to the reasons *why* individuals had unmet health care needs (which are only available conditional on reporting unmet needs). If an individual stated that they had unmet needs, the survey then asked them to describe the reasons why their needs were unmet. Respondents were given a list of potential reasons, including not being able to get an appointment, unwillingness from doctors to accept them as patients, whether the office was closed, whether they felt impaired by outstanding medical debt, whether they thought seeking care would be too costly, etc. I classify reasons as being either 'supply side' or 'demand side.'

Insurance Status Classification: I divide the sample into three groups: uninsured, partially insured, and fully insured. An individual is categorized as "uninsured" if they reported not having insurance for the previous six months (at the time they took the survey); "partially insured" if they reported having insurance for at least one month but for less than six months; and "fully insured" if they reported having insurance for six out of the previous six months.

Figure 5: Share of Individuals Who Reported Having Unmet Medical Needs By Lottery Outcome



Notes: Figure shows the share of individuals who reported having unmet medical needs. The share of individuals with unmet needs is calculated in a regression where the dependent variable is a binary indicator for whether the reported having unmet needs, and the independent variables are indicators for the number of months insured preceding the survey, plus survey wave dummies. The three survey waves are stacked and regressions are weighted using the sampling weights provided by Finkelstein et al (2012). N=53,403.

Table 2: Reasons Why Did Not Receive Medical Care (Although Needed)

	Insurance Status		
	Uninsured	Partially insured	Fully insured
A. Supply Side (%)			
Doc would not accept	0.69	5.01	10.39
Could not get appointment	5.89	9.51	16.93
Office was closed	1.53	2.06	4.47
Did not have doctor	19.49	12.85	9.64
B. Demand Side (%)			
Cost too much	72.45	56.17	36.68
Did not think had insurance	76.14	62.34	12.04
Outstanding debt	15.57	12.34	10.05
Total N with Unmet Needs	6,353	778	1,453

Notes: This table summarizes the number of individuals who reported having unmet medical care needs during the past 6 months, in the sample of respondents to the third survey wave (12 months after the lottery). Insurance status indicates level of coverage over the past 6 months: partially insured individuals reported having at least one month of insurance, but less than six months; fully insured individuals were insured for the full six month period over which they were asked about. Reasons are categorized by supply and demand side. Individuals could state multiple reasons for going without medical care.

Table 2 shows the frequency of reported reasons for why individuals went without needed medical care by health insurance status.[10] Among the fully insured group, the most frequently cited reason for having unmet medical needs was that care was too costly. Across all three groups, supply-side reasons are less frequently cited than demand-side reasons. The likelihood that an individual reported that care cost too much declines with insurance coverage.

The fact that individuals with Medicaid reported that care was "too costly" might seem surprising because there is no associated cost-sharing in the OHP Standard plan. There are at least three possible interpretations of this evidence: misunderstanding of benefits, uncovered services, and transaction costs of getting care. First, it is possible that individuals who won the Oregon Medicaid lottery did not fully understand the plan benefits, perceiving (incorrectly) that meeting their health care needs would result in high out-of-pocket costs. Second, it is possible that meeting their health care needs involved using services that were not covered by the Oregon Medicaid plan, such as mental health services, dental care, and durable medical equipment.[11] For these set of services, Medicaid beneficiaries face the full costs of care, and individuals who reported that care was too costly had a higher prevalence of diagnoses associated with uncovered medical services (see Appendix A.3 for more detail). Third, obtaining care that is fully subsidized still involves transaction costs for the individual, such as taking a day off of work, or finding transportation to the health care facility. Thus, it is possible that respondents were referring to these transaction costs as "too costly." Therefore, it would be misleading to conclude that the lower utilization of health care services by low-income individuals is solely due to behavioral frictions (e.g. misunderstanding of plan benefits).

To contextualize the results, the extensive margin health care utilization among this group of low-income individuals was 74% across all services (see Table 2 in Finkelstein, Hendren, and Luttmer 2019), which is similar in magnitude to extensive margin utilization in the MEPS data. Thus, while the finding that health care spending increased in response to insurance is robust across a number of studies (Finkelstein et al., 2012; Baicker et al., 2013; Taubman et al., 2014; Finkelstein et al., 2016), the extensive margin utilization of low-income individuals *with* insurance did not 'catch-up' to utilization levels of higher income individuals (87% in the MEPS data). This is all the more surprising given that Oregon enrollees were sicker, relative to the average low-income population (Kowalski, 2023; Allen et al., 2010); thus, we would have expected average health care utilization to be higher within this sample of low-income individuals simply due to the fact that they were sicker.

[10]The main text includes the results from the third survey wave, which includes the mix of insured lottery winners and uninsured lottery losers. Similar tables can be constructed using data from the two earlier survey waves, leveraging variation in reported insurance status. For completeness, Appendix Table A.4 summarizes the reasons across all three survey waves.

[11]Covered benefits under Medicaid are determined by the "Prioritized List of Health Services and Treatments" (Oregon Health Services Commission, 2009). Covered services are described in Appendix Table A.3.1.

4 Discussion

Whether the observed demand patterns capture income effects, transaction costs, behavioral frictions, or some other phenomenon is beyond the scope of what the RAND, Oregon, and MEPS data can answer. The model presented here builds upon the simple idea that the marginal utility of cash is large for low-income individuals,[12] and this may push them towards the corner solution.

4.1 Income Effects and Alternative Mechanisms

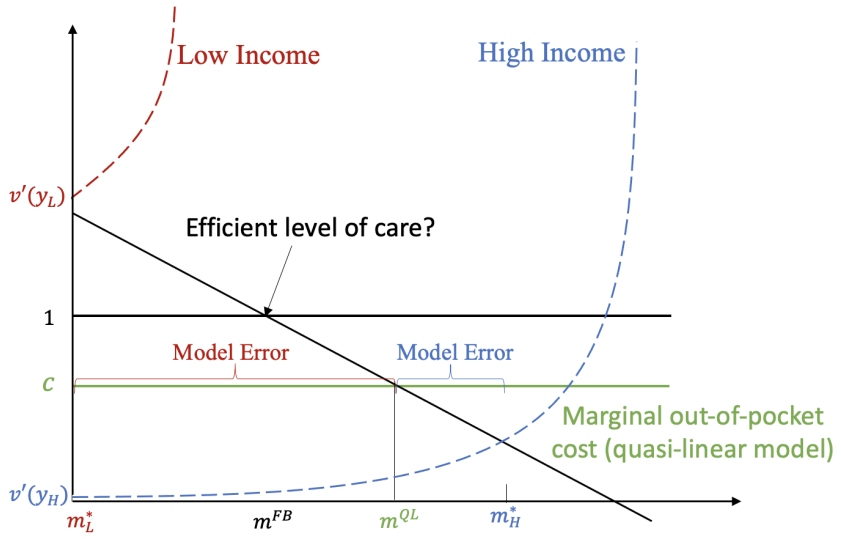
Alternative models could generate similar demand predictions, such as a model with transaction costs: taking time off of work to go to the doctor, having limited transportation means to access a health care facility, or living far away from a health care facility, are all factors that affect the decision to use health care services. If income is correlated to the size of transaction costs, then we should expect similar demand patterns as the ones described here. Empirically quantifying these transaction costs and understanding their relationship to socioeconomic status could be a fruitful area for future research. Alternatively, one could consider a behavioral frictions mechanism, such as misperceived benefits of using care (as in Baicker, Mullainathan, and Schwartzstein, 2015). If behavioral error is correlated with educational attainment and income, then utilization patterns among low-income individuals will not vary with price, but rather with informational campaigns.

Nonetheless, a model which ignores income effects may bias estimates of primitives in the individual's objective function. Because shutting down income effects implies that the marginal cost of forgone consumption is fully determined by the marginal out-of-pocket costs (regardless of whether the individual is high- or low-income), any two individuals facing the same out-of-pocket price and sickness profile should spend the same amount on health care. Any unexplained differences will be absorbed by the empirical error term. To illustrate, consider a linear cost sharing contract, with out-of-pocket costs cm for $c \in [0, 1]$. Suppose the true model of utility were $u^{True}(x, m) = v(x) + H(m)$ for $v(x)$ concave, and consider a misspecified quasi-linear utility model, $u^{QL}(x, m) = x + H(m)$. The marginal cost of forgoing consumption will equal c under the quasilinear utility model (for both high- and low-income individuals).

Figure 6 shows the equilibrium choices of medical spending for a low- and high-income individual, plotted against the equilibrium choice of m implied by the quasi-linear utility model. Low-income individuals will appear to underconsume high-value care, i.e. care for which marginal benefit exceeds the marginal cost. If the efficient benchmark constituted the case without insurance ($p = 1$), low- and high- income individuals would still appear to deviate from the efficient allocation, and ignoring income effects might lead one to conclude that either (1) the marginal benefit of care was different across these two individuals, or (2) that low-income individuals underestimated the benefits of care (and that high-income individuals overestimated them).

[12] This consideration is largely consistent with Miller et al. (2024), who found that giving low-income individuals a recurrent unconditional cash transfer of $1,000 a month for three years (ranging from 30% to 60% of participant income) only moderately raised discretionary (dental) health care spending.

Figure 6: Quasi-linear Misspecification



Notes: This figure shows the marginal benefit and marginal cost curves implied by the quasi-linear utility model, overlaid against the equilibrium medical spending choices for a low- and high-income individual, m_L^* and m_H^* with the same marginal benefit curve. In equilibrium, each individual equates their marginal benefit of care against their marginal cost of forgoing consumption. The quasi-linear model implies that the equilibrium choice is m^{QL} for both high- and low- income individuals. If the efficient benchmark constituted the no insurance case, the efficient level of medical spending m^{FB} equates the marginal benefit of care to one, the marginal cost of care without insurance.

4.2 Policy Implications

If the regulator's objectives include inducing some minimal amount of health care utilization among the low-income population (e.g. primary care check ups), then designing policies that rely on demand-side incentives may not be very cost-effective.[13] This is because low-income individuals derive a higher marginal utility from cash than they do from health care services, and even with a full subsidy of out-of-pocket health care costs, they may not see a health care provider absent additional utilization-contingent cash-incentives (Bradley, Neumark, and Saxe Walker, 2018).

To induce low-income utilization, the regulator could opt for supply-side policy interventions. For example, government-run Community Health Centers (CHCs) provide low cost primary care services in low-income communities. Work by Bailey and Goodman-Bacon (2015) suggests that CHCs are highly effective at inducing utilization among the low-income population, particularly through providing beneficiaries a regular source of care.[14] Moreover, the rollout of CHCs substantially reduced mortality among beneficiaries. Demand-side interventions, such as the Medicaid program, can also lead to mortality reductions among low-income beneficiaries (Wyse and Meyer,

[13]There are at least two reasons for why it may be socially desirable to induce some minimal amount of utilization: externalities and internalities. If the costs associated with an acute health emergency exceed an individual's ability to pay for care, and those costs are ultimately born by health care providers (uncompensated care), or the state, then inducing individuals to use primary care or preventative care will be less costly to society. Similarly, if individuals are myopic or liquidity constrained such that they are unable to optimally invest in their long-term health, then it will be socially efficient to induce some lower bound of health investments through the use of health care services.

[14]Note that, if observed utilization patterns were purely driven by a correlation between behavioral error and income, then supply-side and demand-side policies should not differentially induce utilization among the low-income population.

2023), but at a greater cost per life-year. Wyse and Meyer (2023) estimate that Medicaid costs $179,000 per life-year saved (2019 dollars). In contrast, Bailey and Goodman-Bacon's (2015) estimates of the cost per life-year saved through government-run primary care centers corresponds to $60,000 per life-year in 2019 dollars, a third of the Medicaid cost per life-year.

If policy instruments are limited in their ability to target low-income populations, then my results also have distributional implications. For example, a uniform health care subsidy will primarily constitute a transfer to middle- and high-income individuals, while having a small effect on low-income individuals. This may be particularly undesirable if the regulator's preferences put a higher social welfare weight on the poor, relative to the rich. In general, any (second-best) optimal insurance policy which cannot condition on income will involve more generous cost-sharing if income effects are taken into account. Very generous cost-sharing is needed to get low-income individuals to move into the interior, and this disproportionately benefits middle- and high-income individuals. If means-tested cost-sharing schemes are possible, then optimal insurance will become less generous as income rises (Nyman 2003, 2008; Cremer and Lozachmeur, 2024).

5 Conclusion

In a fairly general model, I show that the marginal cost for forgone consumption is an important determinant of the medical spending decision across high- and low-income individuals. Empirically, I show that low-income individuals are less likely to consume health care services, less responsive to health care prices, and consider high health care costs to be a major deterrent to obtaining necessary medical care. Thus, while low-income individuals are typically the most price sensitive segment in markets *where they participate*, they are less likely to participate in health care markets.

Taken together, my findings suggest that income effects are an important aspect of health care demand, and they manifest primarily through pushing low-income individuals towards the corner solution of zero medical spending. These effects have important implications for policy design and demand estimation in health care markets.

References

Acemoglu, Daron, Amy Finkelstein, and Matthew J. Notowidigdo, "Income and Health Spending: Evidence from Oil Price Shocks," *The Review of Economics and Statistics*, October 2013, *95* (4), 1079–1095.

Agency for Healthcare Research and Quality, "Medical Expenditure Panel Survey (MEPS). Content last reviewed March 2023," Agency for Healthcare Research and Quality, Rockville, MD 2023.

Allen, Heidi, Katherine Baicker, Amy Finkelstein, Sarah Taubman, and Bill J. Wright, "What The Oregon Health Study Can Tell Us About Expanding Medicaid," *Health Affairs*, August 2010, *29* (8), 1498–1506.

Aron-Dine, Aviva, Liran Einav, and Amy Finkelstein, "The RAND Health Insurance Experiment, Three Decades Later," *Journal of Economic Perspectives*, February 2013, *27* (1), 197–222.

Baicker, Katherine, Sarah L. Taubman, Heidi L. Allen, Mira Bernstein, Jonathan H. Gruber, Joseph P. Newhouse, Eric C. Schneider, Bill J. Wright, Alan M. Zaslavsky, and Amy N. Finkelstein, "The Oregon Experiment Effects of Medicaid on Clinical Outcomes," *New England Journal of Medicine*, May 2013, *368* (18), 1713–1722.

_ , Sendhil Mullainathan, and Joshua Schwartzstein, "Behavioral Hazard in Health Insurance," *The Quarterly Journal of Economics*, November 2015, *130* (4), 1623–1667.

Bhattacharya, Jay and Darius Lakdawalla, "Does Medicare benefit the poor?," *Journal of Public Economics*, 2006, *90* (1-2), 277–292.

Bradley, Cathy J., David Neumark, and Lauryn Saxe Walker, "The effect of primary care visits on other health care utilization: A randomized controlled trial of cash incentives offered to low income, uninsured adults in Virginia," *Journal of Health Economics*, November 2018, *62*, 121–133.

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

Chandra, Amitabh, Jonathan Gruber, and Robin McKnight, "The impact of patient cost-sharing on low-income populations: Evidence from Massachusetts," *Journal of Health Economics*, January 2014, *33*, 57–66.

Chen, Jiafeng and Jonathan Roth, "Logs with Zeros? Some Problems and Solutions," *The Quarterly Journal of Economics*, May 2024, *139* (2), 891936.

Chernozhukov, Victor and Christian Hansen, "An IV Model of Quantile Treatment Effects," *Econometrica*, 2005, *73* (1), 245–261.

Chetty, Raj, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler, "The Association Between Income and Life Expectancy in the United States, 2001-2014," *JAMA*, April 2016, *315* (16), 1750–1766.

Commission, Oregon Health Services, "Prioritization of health services. A report to the Governor and the 75th Oregon Legislative Assembly," Technical Report 75th Oregon Legislative Assembly October 2009.

Cremer, Helmuth and Jean-Marie Lozachmeur, "Nonlinear reimbursement rules for preventive and curative medical care," *TSE Working Paper*, 2024, (24-152).

Cutler, David M. and Sarah J. Reber, "Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection," *The Quarterly Journal of Economics*, May 1998, *113* (2), 433–466.

Ellis, R. P. and T. G. McGuire, "Provider behavior under prospective reimbursement. Cost sharing and supply," *Journal of Health Economics*, June 1986, *5* (2), 129–151.

Finkelstein, Amy, "Oregon Health Insurance Experiment Public Use Data," 2013.

Finkelstein, Amy N., Sarah L. Taubman, Heidi L. Allen, Bill J. Wright, and Katherine Baicker, "Effect of Medicaid Coverage on ED Use Further Evidence from Oregons Experiment," *New England Journal of Medicine*, October 2016, *375* (16), 1505–1507.

Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer, "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment," *Journal of Political Economy*, December 2019, *127* (6), 2836–2874.

_ , _ , and Mark Shepard, "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts," *American Economic Review*, April 2019, *109* (4), 1530–1567.

_ , Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P. Newhouse, Heidi Allen, Katherine Baicker, and Oregon Health Study Group, "The Oregon Health Insurance Experiment: Evidence from the First Year," *The Quarterly Journal of Economics*, August 2012, *127* (3), 1057–1106.

Kaplan, David M. and Yixiao Sun, "Smoothed Estimating Equations For Instrumental Variables Quantile Regression," *Econometric Theory*, 2017, *33* (1), 105–157.

Kowalski, Amanda E., "Reconciling Seemingly Contradictory Results from the Oregon Health Insurance Experiment and the Massachusetts Health Reform," *The Review of Economics and Statistics*, May 2023, *105* (3), 646–664.

Krueger, Alan B. and Ilyana Kuziemko, "The demand for health insurance among uninsured Americans: Results of a survey experiment and implications for policy," *Journal of Health Economics*, September 2013, *32* (5), 780–793.

Lavetti, Kurt, Thomas DeLeire, and Nicolas R. Ziebarth, "How do lowincome enrollees in the Affordable Care Act marketplaces respond to costsharing?," *Journal of Risk and Insurance*, June 2022, *90* (1), 155–183.

Lin, Haizhen and Daniel W. Sacks, "Intertemporal substitution in health care demand: Evidence from the RAND Health Insurance Experiment," *Journal of Public Economics*, July 2019, *175*, 29–43.

Manning, Willard G., Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, and Arleen Leibowitz, "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment," *The American Economic Review*, 1987, *77* (3), 251–277.

McClellan, Mark and Jonathan Skinner, "The incidence of Medicare," *Journal of Public Economics*, January 2006, *90* (1), 257–276.

Miller, Sarah, Elizabeth Rhodes, Alexander W. Bartik, David E. Broockman, Patrick K. Krause, and Eva Vivalt, "Does Income Affect Health? Evidence from a Randomized Controlled Trial of a Guaranteed Income," July 2024.

Mullahy, John and Edward C. Norton, "Why Transform Y? The Pitfalls of Transformed Regressions with a Mass at Zero," *Oxford Bulletin of Economics and Statistics*, 2024, *86* (2), 417–447.

Newhouse, Joseph P., *Free for All?*, Harvard University Press, 1993.

Nyman, John A., *The Theory of Demand for Health Insurance*, Stanford University Press, 2003.

_ , "Health Insurance Theory: The Case of the Missing Welfare Gain," *The European Journal of Health Economics*, 2008, *9* (4), 369–380.

Office, Congressional Budget, "How CBO Analyzes Proposals for a Single-Payer Health Care System | Congressional Budget Office," December 2020.

Shepard, Mark, Katherine Baicker, and Jonathan Skinner, "Does One Medicare Fit All? The Economics of Uniform Health Insurance Benefits," *Tax Policy and the Economy*, May 2020, *34*, 1–41.

Taubman, Sarah L., Heidi L. Allen, Bill J. Wright, Katherine Baicker, and Amy N. Finkelstein, "Medicaid Increases Emergency-Department Use: Evidence from Oregons Health Insurance Experiment," *Science (New York, N.Y.)*, January 2014, *343* (6168), 263–268.

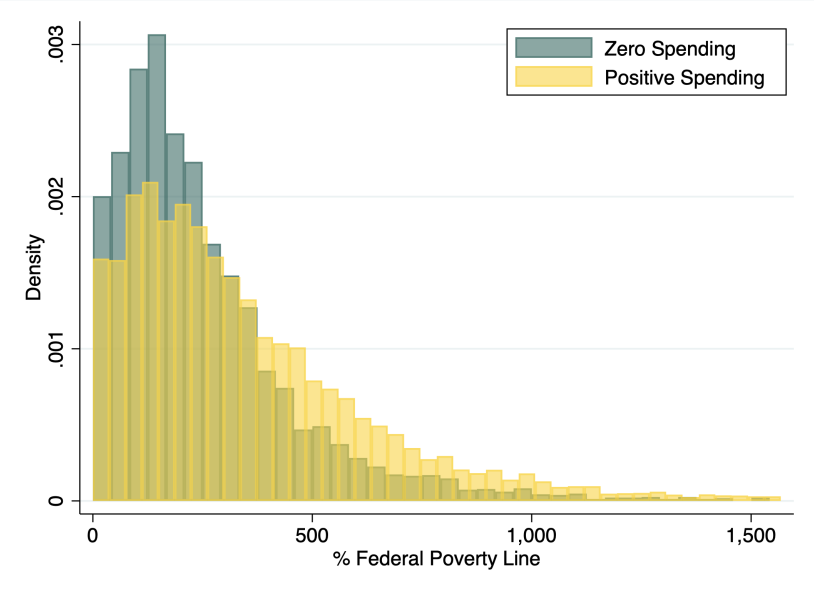
A Empirical Appendix

A.1 Medical Expenditure Panel Survey (MEPS) Analysis

A.1.1 Additional Descriptive Evidence on Income and Extensive Margin Utilization

Figure 2 in the main text showed a positive relationship between income and the likelihood of having any health care expenditures, holding observable characteristics about health status fixed. Figure A.1 shows the distribution of incomes conditional on whether individuals had zero or any positive amount of health care spending. Among those who consume zero health care, the likelihood of having a lower income is higher than among individuals who consume any positive amount of care.

Figure A.1: Income Distribution Conditional on Health Care Utilization



Notes: This figure shows the distribution of incomes among the same individuals, conditional on whether they had or did not have any health care spending. The unique number of individuals is N=156,951.

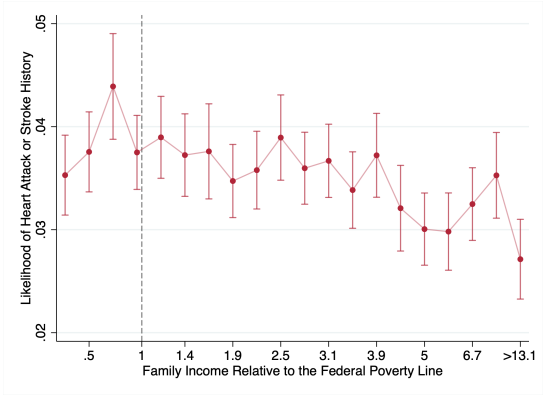
A.1.2 Descriptive Evidence on Income and Health Status

One potential explanation for this relationship could be that lower income individuals are in better health, and their overall health status is not captured by observables. While this explanation cannot be ruled out, concretely, we can look at factors correlated with health status, such as the likelihood of having a heart attack or stroke in the individual's medical history. The MEPS survey enumerators collect information about the individual's disease profile by first asking individuals about their own medical history, and then verifying with their health care provider whether or not the provided information is accurate.

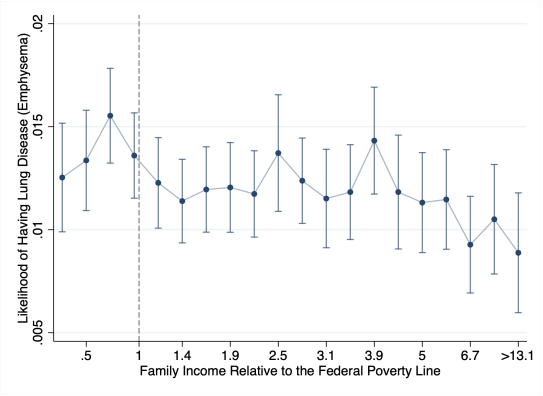
Figure A.2 shows the estimated likelihood that an individual had a particular acute or chronic health condition conditional on income, and controlling for demographics. The dashed vertical line

denotes the FPL (standardized across years). Relative to their higher income counterparts, panel (a) shows that lower income individuals are more likely to have had a heart attack or stroke; panel (b) shows that they are slightly more likely to suffer from emphysema, which is a lung condition that causes shortness of breath; and panels (c) and (d) show that they are equally likely to suffer from arthritis or diabetes. If incidence of these acute or chronic health conditions is correlated to underlying health status, then this descriptive evidence is at the very least inconsistent with the explanation that lower income individuals are in better health.

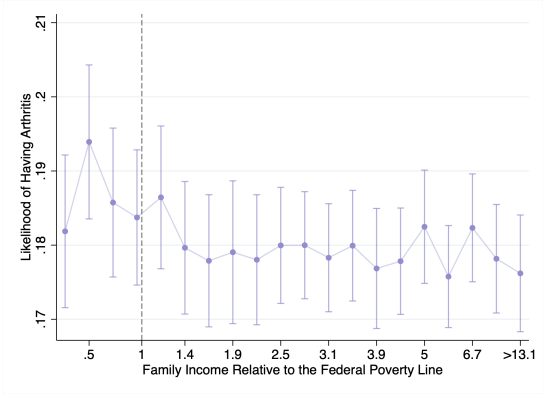
Figure A.2: Income and Health Status



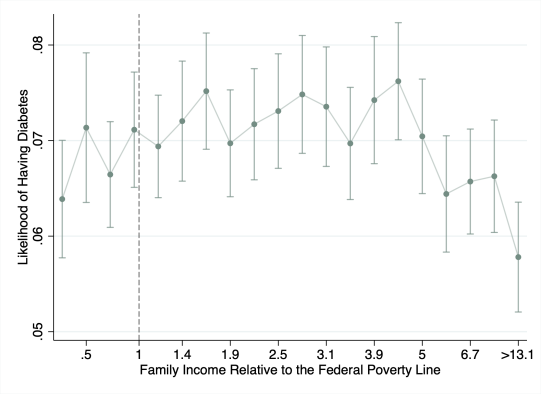
(a) Likelihood of Heart Attack or Stroke History



(b) Likelihood of Lung Disease (Emphysema)



(c) Likelihood of Arthritis



(d) Likelihood of Diabetes

Notes: All figures show the predicted probability that an individual had a specific health condition, conditional on their income. Each probability is estimated in a probit regression where the dependent variable is a binary indicator for whether the individual had a health condition or history thereof, and the independent variables are 20 categorical dummies for the individual's income ventile plus controls. Regressions are estimated at the person-year level and controls include insurance type (private, public, or none), age, sex, age squared, age by sex interactions, smoking habits, race, education, impairments to physical and cognitive functioning, year fixed effects, and survey panel fixed effects. Standard errors are adjusted to account for survey design with non-random sampling. For all figures, the unique number of individuals is N=156,951.

A.2 RAND Experiment

A.2.1 Demographic Covariates by Income Terciles

Table A.1 shows descriptive statistics about the individuals in the experiment, conditional on income tercile. On average, individuals in the bottom income tercile are younger, have lower levels of employment and schooling, and larger families, relative to the middle and top income terciles. Relative to their higher income counterparts, individuals in the bottom income tercile were also less likely to have a full time job. The share of females was similar across all three groups.

Table A.1: RAND Experiment: Descriptives by Income Tercile

	Mean	S.d.	Min	Max
Bottom Income Tercile				
Income ($)	$5,826	$2,126	$892	$9,216
Income (% FPL)	128%	47%	20 %	203%
Age	20.4	15.3	0	61
Yrs Schooling	10.9	3.0	0	21
Yrs Employed	9.4	9.3	0	46
Family Size	4.1	2.2	1	14
Middle Income Tercile				
Income ($)	$12,172	$1,717	$9,222	$15,250
Income (% FPL)	268%	38%	203%	336%
Age	23.4	15.4	0	61
Yrs Schooling	12.0	2.7	0	25
Yrs Employed	12.1	9.7	0	45
Family Size	3.9	1.6	1	10
Top Income Tercile				
Income ($)	$20,992	$5,110	$15,253	$50,847
Income (% FPL)	463%	113%	336%	1,120%
Age	30.7	17.2	0	61
Years of Schooling	12.9	2.8	1	24
Yrs Employed	15.6	11.3	0	45
Family Size	3.3	1.5	1	14

%	*Bottom*	*Middle*	*Top*
Employed Full-Time	19.1	25.7	29.3
Self-Employed	7.9	8.1	5.8
Female	55.5	50.7	49.4
White	69.0	91.1	92.7
N	1,804	1,967	2,144

Notes: N denotes the number of unique individuals in each income tercile. The number of unique families in each income tercile is 636, constant across all three bins.

Table A.2: Smoothed IV Quantile Regression Estimates of the Semi-Elasticity of Medical Spending by Income Tercile

	Quantiles									
	5	10	15	20	25	30	35	40	45	50
***A.* Income Tercile 1**										
	-0.5371	-1.0745	-1.5085	-2.2919	-4.4694	-4.9377	-5.1932	-5.2997	-5.2118	-4.0867
	(0.3792)	(0.6664)	(0.9338)	(1.3500)	(1.0427)	(0.6167)	(0.3898)	(0.4030)	(0.6026)	(0.7525)
	[0.1567]	[0.1069]	[0.1062]	[0.0896]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
***B.* Income Tercile 2**										
	-0.5502	-4.6184	-5.1418	-4.2377	-2.6716	-2.1856	-1.9256	-1.7464	-1.6401	-1.5403
	(1.2834)	(0.8849)	(0.9217)	(1.0697)	(0.8427)	(0.3478)	(0.2684)	(0.2152)	(0.1766)	(0.1496)
	[0.6681]	[0.0000]	[0.0000]	[0.0001]	[0.0015]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
***C.* Income Tercile 3**										
	-3.5417	-5.1838	-4.5420	-2.4536	-1.9608	-1.6929	-1.5414	-1.4281	-1.3514	-1.2941
	(1.2331)	(0.6034)	(1.1377)	(0.7356)	(0.3482)	(0.2187)	(0.1743)	(0.1616)	(0.1516)	(0.1384)
	[0.0041]	[0.0000]	[0.0001]	[0.0009]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]

	Quantiles								
	55	60	65	70	75	80	85	90	95
***A.* Income Tercile 1**									
	-3.2537	-2.8011	-2.5176	-2.3650	-2.3106	-2.2822	-2.3366	-2.4374	-2.4689
	(0.5592)	(0.4248)	(0.3278)	(0.2822)	(0.2716)	(0.2739)	(0.2753)	(0.2607)	(0.2781)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
***B.* Income Tercile 2**									
	-1.5058	-1.4743	-1.4583	-1.4608	-1.4718	-1.4827	-1.5870	-1.7268	-2.0157
	(0.1447)	(0.1330)	(0.1326)	(0.1296)	(0.1322)	(0.1333)	(0.1414)	(0.1561)	(0.1560)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
***C.* Income Tercile 3**									
	-1.2524	-1.2207	-1.1874	-1.1561	-1.1347	-1.2027	-1.2710	-1.3686	-1.5923
	(0.1368)	(0.1319)	(0.1335)	(0.1391)	(0.1428)	(0.1573)	(0.1556)	(0.2052)	(0.2513)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]

Notes: Smoothed quantile IV regressions are estimated seperately within each income tercile. Standard errors are displayed in parentheses and p-values in brackets. The dependent variable in each smoothed quantile IV regression is $\log(1 +$ medical spending$)$. The continuous coinsurance rate associated with each RAND plan is instrumented by the RAND plan indicator, controlling for year and location by month effects. Standard errors are cluster-bootstrapped at the family level, using 500 replications.

A.2.2 Medical Expenditure Limits Across Plans and Across Income Tercile

The RAND investigators designed plans that varied along two dimensions: cost-sharing and the out-of-pocket expenditure limit. Health insurance plans in the RAND experiment had out-of-pocket expenditure limits of either 5, 10, or 15% of income, or $1,000 (whichever amount was smaller). This plan feature could be problematic for the IV quantile regression estimates because, if the expenditure limits were systematically correlated to the quantiles of the medical spending distribution in different ways for the different income terciles, the exclusion restriction would be violated. In other words, if individuals at the lower quantiles of the medical spending distribution also got assigned higher expenditure limits, this plan feature might dissuade them from incurring medical expenditures in the first place.

Panel (a) of Figure A.3 shows the average expenditure limit assigned by the RAND plan for all individuals who did not get assigned to the free care plan, plotted against the quantile of the medical spending distribution within income tercile. While the lowest income tercile did have systematically lower levels of the expenditure limit, the average expenditure limit level was constant along the medical spending distribution within an income tercile. Moreover, we cannot statistically reject that the relationship between medical spending quantiles and the expenditure limit is equivalent across income terciles when tested in a stacked regression where the dependent variable is the expenditure limit level, and the independent variables are interaction terms between income tercile dummies with medical spending quantiles, plus income tercile dummies, plan assignment dummies, and location by month and year fixed effects. This suggests that plan incentives were not systematically different along different quantiles of the medical spending distribution within income terciles.

The fact that the lowest-income tercile got systematically lower expenditure limit could be a problem because they might have been more likely to hit their expenditure limit and face a zero marginal price for the rest of the year. Note that, while this feature should drive spending differences among high spenders (i.e. those at the upper range of medical spending quantiles), it could not explain differential behavior among the low spenders. Nonetheless, Figure A.3 Panel (b) shows that lower-income individuals were not significantly more likely to hit their expenditure limits for most plans, with the single exception of the 25% cost-sharing plan, suggesting that the expenditure limits are unlikely to be driving differential behavior in the initial decision to consume any health care services.

A.2.3 Alternative log-like transformations of the outcome variable in IV quantile regression estimates

An additional concern regarding the magnitude of the semi-elasticity estimates when there is many zero values in the outcome variable (medical spending) is that magnitudes are sensitive to the log-like transformation of the outcome variable (Chen and Roth, 2023). Figure A.4 displays IV quantile regression estimates of the semi-elasticity of health care spending with respect to cost-sharing with

two alternative log transformations of the outcomes variables: $\log(m + 0.1)$ and $\log(m + 0.01)$. While the magnitudes of the point-estimates does differ across specifications, the overall shape of the quantile semi-elasticity estimates is similar across alternative specifications. This evidence suggests that the claim that individuals from the bottom income terciles are less responsive to price at lower quantiles of the medical spending distribution, relative to their higher income counterparts, is fairly robust.

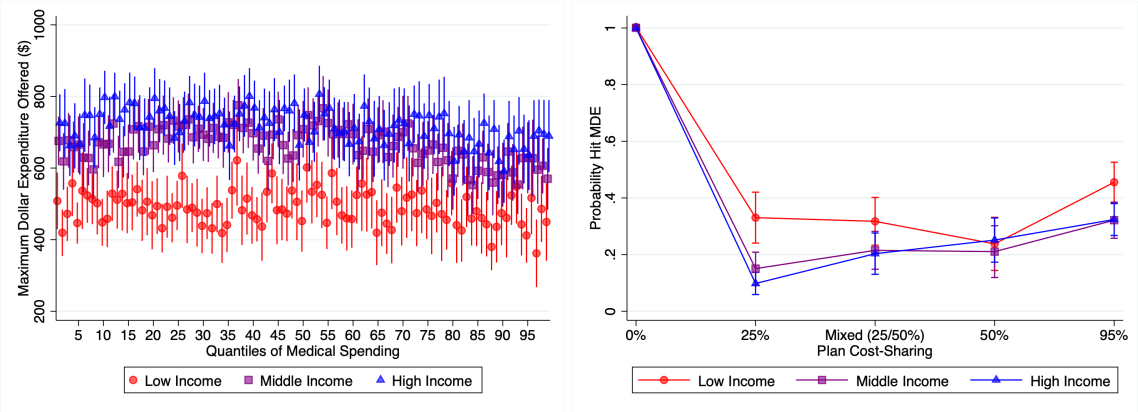
A.2.4 Alternative quantile regression estimates under alternative price assumptions

Depending on the researcher's assumptions about consumer forward-looking behavior, the relevant price could be the realized marginal price at the end of the year, some weighted average of the price paid throughout the year, or the current "spot" price of care. For my main specification, I have assumed that consumers respond to their marginal price at the end of the year, which is endogenous to their spending decisions throughout the year.

One could consider making an alternative assumption that consumers behave myopically. In this case, the relevant price that should be used to estimate the semi-elasticity should be the initial "spot" price of care (i.e. initial RAND plan cost-sharing). Moreover, there is no need to estimate quantile regressions within an IV framework because RAND plans were randomly assigned. Figure A.5 Panel (a) shows the quantile regression estimates using the spot price as the price of care.

One could also consider making the assumption that consumers respond to the average price paid throughout the year within their plan, under the assumption that they do not know their specific

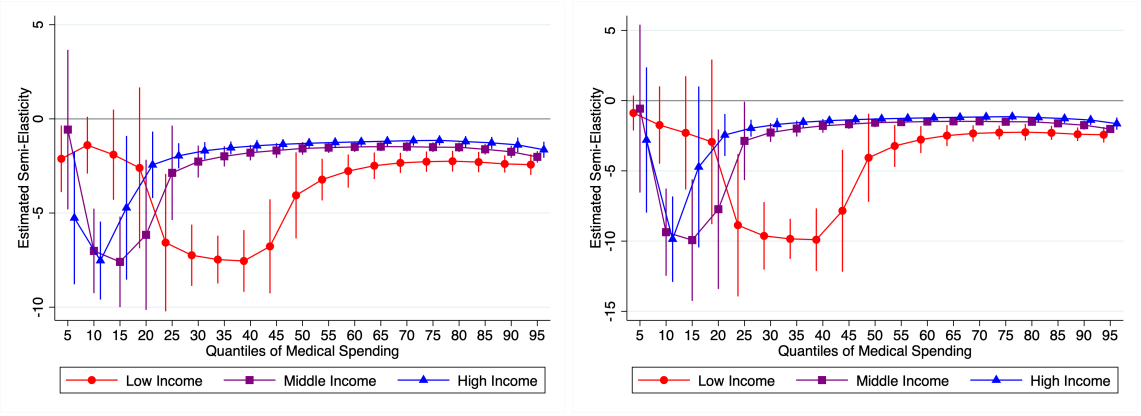
Figure A.3: Plan Maximum Dollar Expenditure and Income



(a) MDE Offered Conditional on Spending Quantile (b) Probability of Hitting MDE by Plan

Notes: Panel (a) shows the average maximum dollar expenditure limit across individuals at different quantiles of the medical spending distribution by income terciles. To break ties within the zero mass bottom quantiles, mean zero normal noise is added to the medical spending variable. Panel (b) shows the share of individuals who hit their maximum dollar expenditure limit within a year. The probability that an individual hit their expenditure limit conditional on income is estimated in a regression where the dependent variable is a binary indicator for whether the individual hit their expenditure limit, and the independent variables are categorical dummies for the individual's cost-sharing plan assignment. Standard errors are clustered on family. Regressions are estimated at the person-year level. The individual deductible plan is excluded from the analysis. N=16,028.

Figure A.4: Alternative Specifications of the Outcome Variable for IV quantile regression

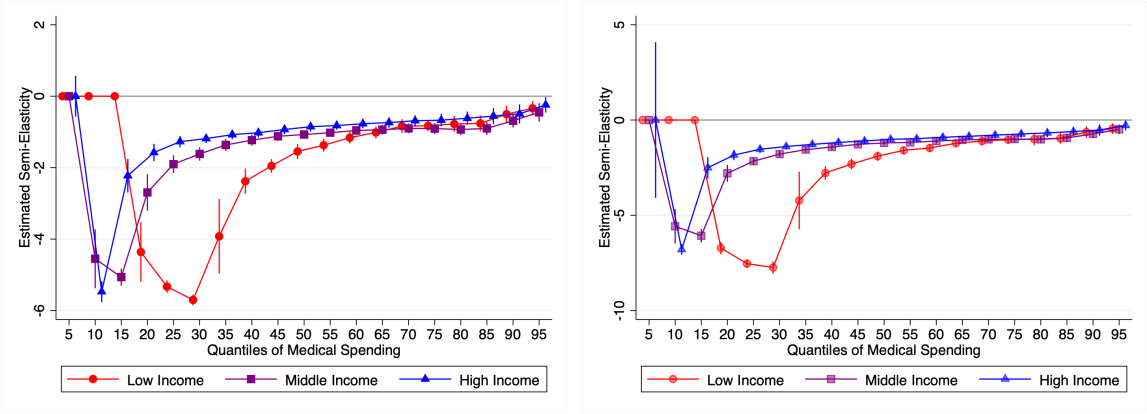


(a) Outcome Variable: $log(m_{it} + 0.1)$ (b) Outcome Variable: $log(m_{it} + 0.01)$

Notes: IV quantile regressions are estimated separately within each income tercile using the IVQR smoothed estimation procedure in Kaplan and Sun (2017). Standard errors are block bootstrapped at the family level based on 200 replications. The dependent variable in these quantile regressions is $\log(m + 0.1)$ in Panel (a) and $\log(m + 0.01)$ in Panel (b). The independent variables are the continuous coinsurance rate associated with the RAND plan, and year plus location by month fixed effects. N=16,028.

type and make decisions based on the average consumer within their plan. Figure A.5 Panel (b) shows the quantile regression estimates using the average out-of-pocket share within plan (and within income tercile) as the price of care. Under both specifications, the spending response across income groups looks very similar to the response in the IV quantile regression estimates in Figure 4: low-income individuals are less responsive to price at the bottom quantiles of the medical spending distribution.

Figure A.5: Additional Quantile Regression Estimates of the Semi-Elasticities by Income Tercile



(a) Price = Initial Spot Price (b) Price = Average Out of Pocket Share

Notes: Quantile regressions are estimated separately within each income tercile, block bootstrapped at the family level based on 200 replications. The dependent variable in each of these quantile regressions shown in Panels (a) and (b) is log(1+medical spending). The independent variables in Panel (a) include the (continuous) cost-sharing rate associated with the individual's assigned RAND plan (e.g. sticker price), and year plus location by month fixed effects. The independent variables in Panel (b) include the average out-of-pocket share faced by all individuals within the RAND plan and income tercile, and year plus location by month fixed effects.

A.2.5 Pairwise Elasticities

To summarize the average treatment effect of price on intensive margin health care demand, one could compute pairwise elasticities. When a substantial share of individuals consume zero health care, the concept of an elasticity is not well defined because the derivative of health care demand with respect to price is not well defined when either demand or price are equal to zero (Chen and Roth, 2023; Mullahy and Norton, 2022). However, one can summarize the average effect of a price change on health care spending by computing a pairwise elasticity, which is always well defined.

The pairwise elasticity describes the percent change in total spending divided by the percent change in price, computed only across two points (i.e. it is the linear approximation of the slope of the demand curve, calculated using the point-slope formula). For any two prices, p and p', the pairwise elasticity $\eta(p, p')$ is defined as:

$$\eta(p, p') \equiv \frac{\mathbb{E}[m_i|p'] - \mathbb{E}[m_i|p]}{\mathbb{E}[m_i|p]} \cdot \frac{p}{p' - p} \tag{A.1}$$

Within each income tercile, I separately compute pairwise elasticities by taking pairwise differences in average health care spending across plans, divided by the pairwise difference in coinsurance rate. Given the set plans offered, I estimate pairwise elasticities across the four coinsurance rates offered across plans, 0%, 25%, 50%, 95%. Following the RAND investigators and Aron-Dine, Einav, and Finkelstein (2013), I use person-year as the primary unit of analysis. I exclude the mixed coinsurance and individual deductible plans from this analysis. The average conditional spending component of the pairwise elasticity, $\mathbb{E}[m_i|p]$, is estimated based on regressing medical spending on plan dummies, λ_p, and year plus location by month fixed effects.[15] Pairwise elasticities are then computed by taking $\lambda_p = \mathbb{E}[m_i|p]$. Standard errors are bootstrapped based on 2,000 replications and clustered on family.

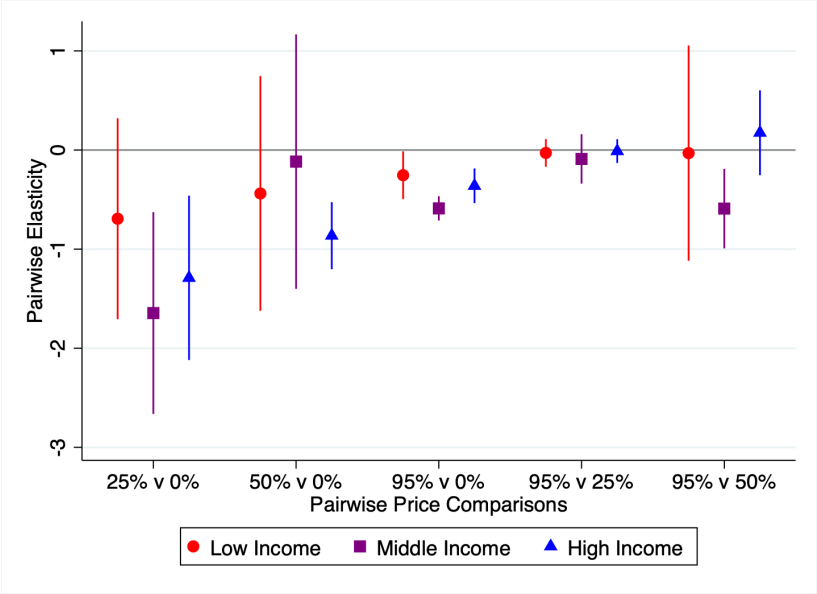
Denoting individuals by i, calendar year by y, medical spending by m, the insurance plan by p, month by t, and location by l, the empirical specification is:

$$m_{i,y} = \lambda_p + \lambda_y + \lambda_{l,t} + \epsilon_{i,t} \tag{A.2}$$

Figure A.6 shows the estimated pairwise elasticities across plans and income tercile. The lowest income tercile were the least responsive to cost-sharing, and the middle income individuals were the most responsive. For the lowest income tercile, the proportional spending reduction relative to the free care plan was not statistically different from zero among any of the cost-sharing plans. The highest income tercile had a smaller response to cost-sharing, relative to the middle income group.

[15] Plan assignment was random only conditional on enrollment month and site (see Appendix B of Newhouse et al. 1993).

Figure A.6: Implied Pairwise Elasticity



Notes: Figure shows the estimated pairwise elasticities across plan groups. The dependent variable is inflation-adjusted medical spending and the independent variables are RAND plan group dummies and year plus location by month fixed effects. Pairwise elasticities are computed by equation (1) using as inputs the plan dummy coefficients from the regression. Standard errors are bootstrapped based on 2,000 replications and clustered on family. All estimates are conducted separately by income tercile subsample, and all regressions are estimated at the person-year level. The mixed coinsurance and individual deductible plans are excluded from the analysis. N= 13,323.

A.2.6 Poisson Regression Estimate for the Intensive Margin Effect of Cost-Sharing

When the outcome of interest (health care demand) can equal zero, the concern with estimating an average treatment effect (ATE) using log-like transformations of the outcome variable is that the estimates cannot be interpreted as percentage effects because, unlike percentage units, the estimates will depend on the units of the outcome variable (Chen and Roth, 2023; Mullahy and Norton, 2022). They suggest conducting a Poisson regression in the event that the researcher is interested in expressing the ATE in percentage terms. For this reason, I estimate a the average treatment effect of cost-sharing in a Poisson regression.

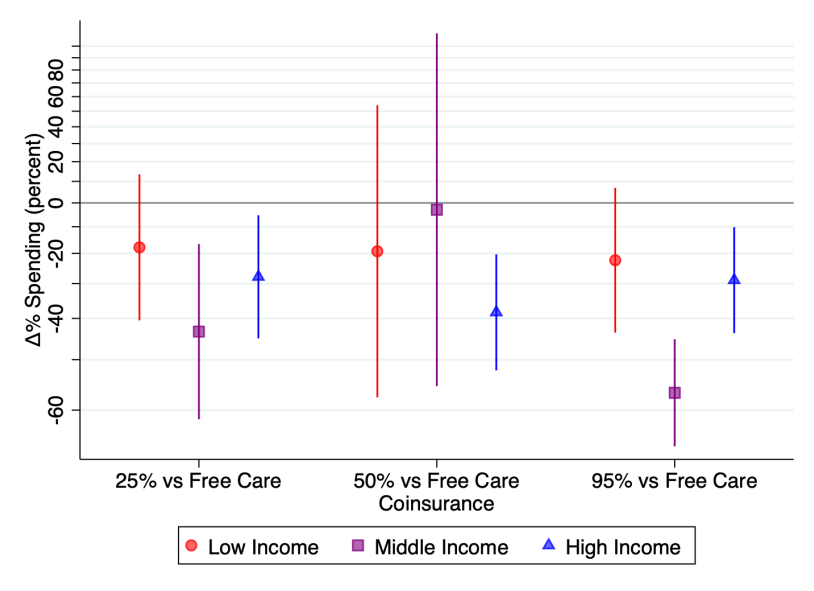
Empirical Framework: I use person-year as the primary unit of analysis, and exclude the mixed coinsurance and individual deductible plans from the analysis. Denoting individuals by i, calendar year by y, medical spending by m, the insurance plan by p, month by t, and location by l, the empirical specification is:

$$\mathbb{E}[m_{i,y}|X] = \exp\left(\beta_p + \beta_y + \beta_{l,t}\right), \tag{A.3}$$

and the main coefficient of interest is β_p. The omitted category is the free care plan, and thus coefficients on the categorical variable dummies for plans with 25%, 50%, and 95% coinsurance can all be interpreted as the exponentiated percent difference in medical spending, relative to the free plan. Standard errors are clustered on family. Figure A.7 shows the results from this regression. Labels on the y-axis have been transformed to correspond to the percentage difference in spending: if the coefficient on regressor x is equal to β, it implies that a unit increase in x corresponds to a

multiplicative increase of $\exp(\beta)$ on y (i.e. a 10% increase in the outcome would correspond to a coefficient $\beta = \log(1 + .10)$).

Figure A.7: Percentage Differences in Spending Across Plans Relative to Free Care



Notes: Figure shows coefficients from a Poisson regression in which the dependent variable is inflation-adjusted medical spending and the independent variables are RAND plan group dummies and year plus location by month fixed effects, as specified in by equation (8). Standard errors are clustered on family. All estimates are conducted separately by income tercile subsample, and all regressions are estimated at the person-year level. The mixed coinsurance and individual deductible plans are excluded from the analysis. N= 14,326

A.3 Oregon Health Insurance Experiment

The Oregon experiment has been extensively studied, as it provided a exogenous variation conducive to study the effects of health insurance on a outcomes, using the lottery as an instrument for insurance enrollment (Allen et al., 2010; Finkelstein et al., 2012; Baicker et al., 2013; Taubman et al., 2014; Finkelstein et al., 2016). Overall conclusions from these studies include that health insurance (Medicaid) increased overall health care utilization and self-reported health, but had little impact on clinical outcomes (Finkelstein et al., 2012; Baicker et al., 2013). Increases in utilization happened primarily through greater use of emergency services (Taubman et al., 2014; Finkelstein et al., 2016), as opposed to other types of care.

The results from the Oregon experiment are not directly comparable to the RAND experiment for a number of reasons.

A.3.1 Covered and Not Covered Services Under Oregon Medicaid

Individuals who were covered by the Oregon Medicaid programs (OHP Standard and OHP Plus) face no cost-sharing for the set of covered services. However, there are a number of uncovered

services for which beneficiaries face the full costs of care. The set of covered services are determined by the Oregon legislature and revised annually. The Oregon Health Services Commission produces a yearly report containing the "Prioritized List of Health Services and Treatments." Table A.3.1 broadly summarizes the set of covered services listed in this document. Importantly, the OHP Standard plan excludes a number of mental health services and has a limited hospital benefit.

Table A.3: List of Covered Health Care services by Oregon Health Plan (Medicaid) in 2009

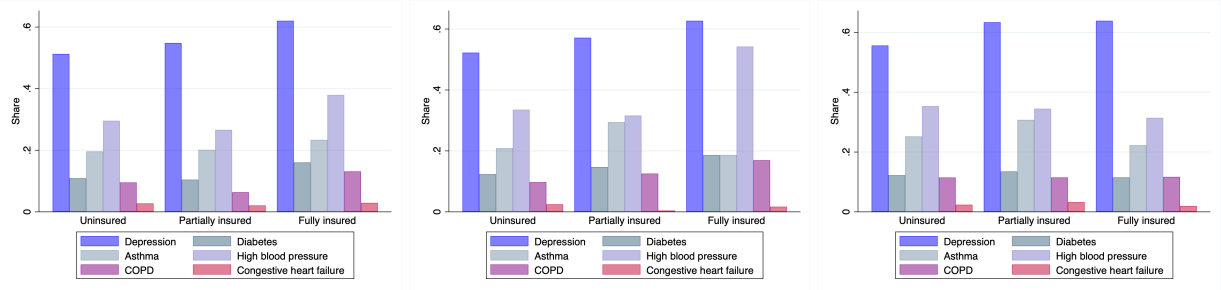
	Service covered	
	OHP Standard	OHP Plus
Physician services	✓	✓
Prescription drugs	✓	✓
Hospice care	✓	✓
Mental health services		✓
Chemical dependency services		✓
Durable medical equipment		✓
• diabetic supplies		✓
• respiratory/oxygen equipment		✓
• ventilators		✓
• suction pumps		✓
• tracheostomy/urology/ostomy supplies		✓
Dental services		
Emergency dental services	✓	✓
Hospital benefit	85% of OHP Plus	✓
• diagnostic tests	✓	✓
• all emergency services	✓	✓
• all preventing life threatening health deterioration	✓	✓

Figure A.8 describes the health care diagnoses within the subsample of individuals who reported having unmet medical needs across all three survey wave. The main takeaway from Figure A.8 is that, among insured low-income individuals who reported having unmet medical needs *because* health care cost too much, the most prevalent diagnosis was depression. While only suggestive, the data are consistent with the possibility that the Medicaid enrollees who had an active insurance policy for the full time period over which they were questioned sought out services that did have associated out-of-pocket costs in the OHP Standard plan. Therefore, it could have been reasonable for a fully insured individual to report that health care cost too much, despite being covered by OHP Standard.

A.3.2 Individuals Who Reported Unmet Needs Across Different Survey Waves

Individuals were asked in the survey to report whether they had unmet medical or prescription needs in the previous six months. Figure A.9 shows the share of individuals who reported having unmet

Figure A.8: Health Diagnoses for Subsample Reporting Medical Care "Cost Too Much" By Insurance Status

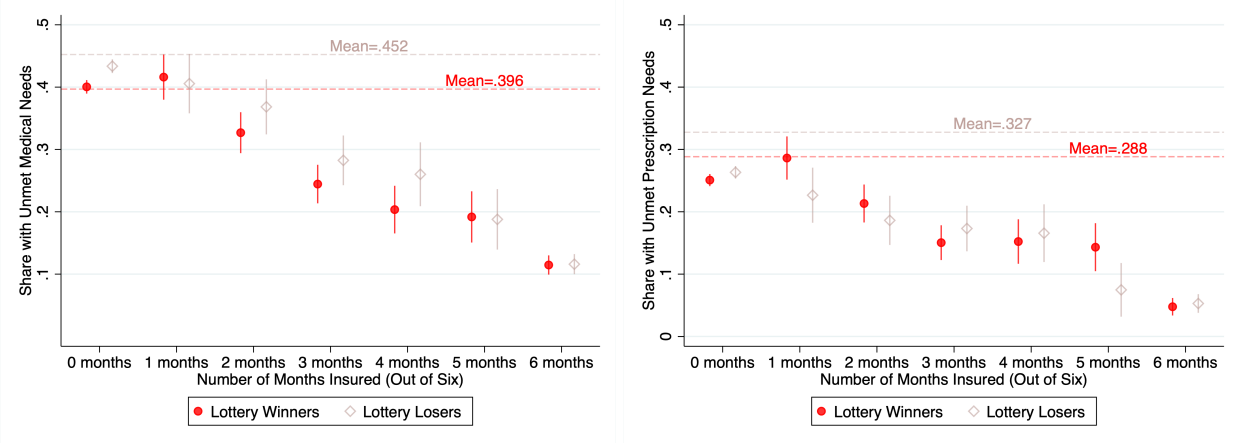


(a) Initial Survey Wave (b) Interim Survey Wave (6 months) (c) Final Survey Wave (1 year after)

Notes: Figure shows the share of individuals with a given chronic condition diagnosis, within the subsample of individuals who reported having unmet medical needs within a survey wave, conditional on insurance status. Panel (a) uses survey responses to the initial survey which was conducted at the time of the lottery (N= 9,518). Panel (b) uses survey responses to the interim survey which was conducted six months after the lottery notification date (N=1,963). Panel (c) uses survey responses to the final survey which was conducted 12 months after the lottery notification date (N=5,751).

medical needs, by lottery outcome and number of months for which they had an active insurance policy.

Figure A.9: Share of Individuals Who Reported Having Unmet Needs By Lottery Outcome



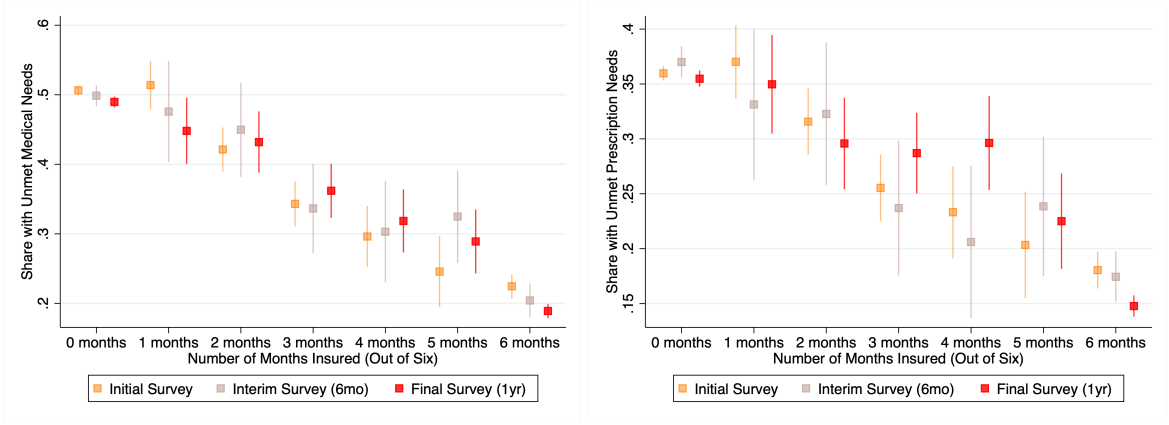
(a) Share with Unmet Medical Needs (b) Share with Unmet Prescription Needs

Notes: Figure shows the share of individuals who reported having unmet medical needs (Panel a) and prescription needs (Panel b). The share of individuals with unmet needs is calculated in a regression where the dependent variable is a binary indicator for whether the reported having unmet needs, and the independent variables are indicators for the number of months insured preceding the survey, plus survey wave dummies. The three survey waves are stacked and regressions are weighted using the sampling weights provided by Finkelstein et al (2012). N=53,403.

Figure A.10 shows the share of individuals who reported having unmet medical needs, by survey wave and number of months for which they had an active insurance policy. Across all three survey waves, individuals who were covered for a greater amount of time were less likely to report having unmet needs.

In each of the three surveys, individuals who reported having unmet health care needs were

Figure A.10: Share of Individuals Who Reported Having Unmet Needs By Survey Wave



(a) Share with Unmet Medical Needs

(b) Share with Unmet Prescription Needs

Notes: Figure shows the share of individuals who reported having unmet medical needs (Panel a) and prescription needs (Panel b). The share of individuals with unmet needs is calculated in a regression where the dependent variable is a binary indicator for whether the reported having unmet needs, and the independent variables are indicators for the number of months insured preceding the survey, separately within each survey subsample. The initial survey (0m) had 25,855 respondents, 11,846 (8,528) of whom reported having unmet medical (prescription) needs. The interim survey (6m) had 6,264 respondents, 2,576 (1,933) of whom reported having unmet medical (prescription) needs. The final survey (12 m) had 23,434 respondents, 8,862 (6,553) of whom reported having unmet medical (prescription) needs.

then asked to report the reasons why. Tables A.4 and A.5 show the reasons that respondents stated for why they had unmet medical and prescription needs, respectively.

Table A.4: Reasons Why Did Not Receive Medical Care (Although Needed)

	Insurance Status		
	Uninsured	Partially insured	Fully insured
A. Supply Side (%)			
0m survey			
Doc would not accept	1.10	5.38	13.19
Could not get appointment	6.35	9.29	14.54
Office was closed	2.26	2.69	5.40
Did not have doctor	22.71	19.12	10.79
Total N with Unmet Needs	9877	1302	667
6m survey			
Doc would not accept	0.62	8.31	15.61
Could not get appointment	6.33	10.89	15.61
Office was closed	2.08	3.15	3.65
Did not have doctor	23.42	18.91	8.31
Total N with Unmet Needs	1926	349	301
12m survey			
Doc would not accept	0.71	5.20	10.26
Could not get appointment	5.78	9.52	16.67
Office was closed	1.53	2.03	4.39
Did not have doctor	19.36	12.82	9.58
Total N with Unmet Needs	6592	788	1482
B. ***Demand Side*** (%)			
0m survey			
Cost too much	81.41	72.58	64.02
Did not think had insurance	85.64	77.34	20.84
Outstanding debt	17.01	15.28	18.14
Total N with Unmet Needs	9877	1302	667
6m survey			
Cost too much	81.98	59.89	53.49
Did not think had insurance	85.25	67.62	19.60
Outstanding debt	18.95	13.47	12.96
Total N with Unmet Needs	1926	349	301
12m survey			
Cost too much	71.74	55.96	36.44
Did not think had insurance	75.67	62.06	12.28
Outstanding debt	15.25	12.18	9.92
Total N with Unmet Needs	6592	788	1482

Notes: This table summarizes the number of individuals who reported having unmet medical care needs during the past 6 months for each of the three survey waves. Insurance status indicates number of months with insurance in the precedent 6 month window. Individuals could state multiple reasons for going without medical care.

Table A.5: Reasons Why Did Not Receive Drug Prescription (Although Needed)

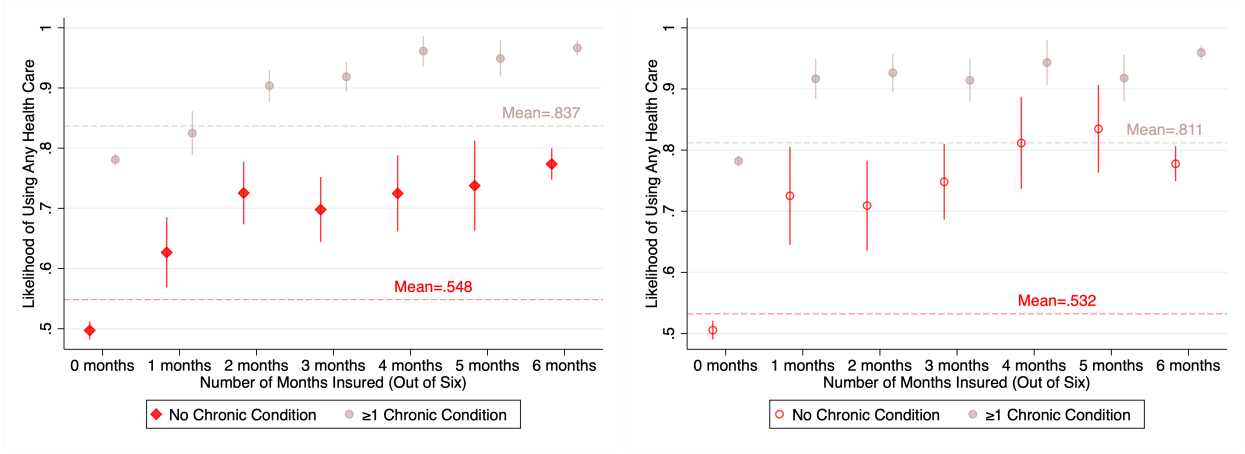
	Insurance Status		
	Uninsured	Partially insured	Fully insured
A. Supply Side (%)			
0m survey			
Doc would not accept	1.01	4.41	8.21
Could not get appointment	5.52	6.97	8.21
Office was closed	1.95	2.36	3.73
Did not have doctor	20.45	16.21	7.65
Total N with Unmet Needs	7017	975	536
6m survey			
Doc would not accept	0.63	7.69	10.12
Could not get appointment	5.88	7.29	11.28
Office was closed	2.10	1.62	1.95
Did not have doctor	21.55	14.17	5.45
Total N with Unmet Needs	1429	247	257
12m survey			
Doc would not accept	0.67	3.88	5.70
Could not get appointment	5.42	6.95	8.89
Office was closed	1.38	1.94	3.28
Did not have doctor	18.30	10.18	6.13
Total N with Unmet Needs	4776	619	1158
B. ***Demand Side*** (%)			
0m survey			
Cost too much	67.21	56.82	40.11
Did not think had insurance	70.76	59.59	14.18
Outstanding debt	17.29	14.77	13.25
Total N with Unmet Needs	7017	975	536
6m survey			
Cost too much	69.28	48.18	30.35
Did not think had insurance	71.94	54.66	10.12
Outstanding debt	18.68	11.74	7.78
Total N with Unmet Needs	1429	247	257
12m survey			
Cost too much	59.67	41.52	22.28
Did not think had insurance	63.32	47.66	6.99
Outstanding debt	15.26	10.99	6.82
Total N with Unmet Needs	4776	619	1158

Notes: This table summarizes the number of individuals who reported having unmet prescription needs during the past 6 months for each of the three survey waves. Insurance status indicates number of months with insurance in the precedent 6 month window. Individuals could state multiple reasons for going without prescriptions.

A.3.3 Extensive Margin Utilization in Oregon Sample by Health Type

Figure A.11 shows the likelihood that individuals in the Oregon experiment used any health care services, split by those at least one chronic condition versus those without. As one would expect, the individuals with at least one chronic condition are *more likely* to use health care services along the extensive margin, relative to those without. However, among those without chronic conditions who may have temporary health care needs, the likelihood of using any health care services was under 80%, similar in magnitude to utilization in the MEPS data.

Figure A.11: Extensive Margin Utilization in Oregon Experiment by Chronic Conditions



(a) Lottery Winners

(b) Lottery Losers

Notes: Figure shows the average likelihood that an individual in the Oregon experiment used any health care services, conditional on lottery outcome and on whether they reported having at least one of six chronic condition from the list: depression, diabetes, asthma, high blood pressure, emphysema, acute myocardial infarction, congestive heart failure, depression, high cholesterol, or kidney disease. The likelihood that an individual uses some positive amount of health care is estimated using a linear probability model in a regression where the dependent variable is a binary indicator for whether the individual consumed non-zero care, and the independent variables are indicators for the number of months insured preceding the survey, plus survey wave dummies. The three survey waves are stacked and regressions are weighted using the sampling weights provided by Finkelstein et al (2012). N=53,403.