

The Incidence of Adverse Selection: Theory and Evidence from Health Insurance Choices*

Michael Geruso[†]

Timothy J. Layton[‡]

Adam Leive[§]

Abstract

Existing research on selection in insurance markets focuses on how adverse selection distorts prices and misallocates products across people. This ignores the distributional consequences of *who* pays the higher prices. In this paper, we show that the distributional incidence depends on the correlations between income, expected costs, and insurance demand. We discuss the general implications for the design of subsidies and mandates when policymakers value both equity and efficiency. Then, in an empirical case study of a large employer, we show that the incidence of selection falls on higher-income employees, who are more likely to choose generous health insurance plans.

*We are grateful to helpful comments from Jason Abaluck, Zarek Brot-Goldberg, Colleen Carey, Amy Finkelstein, Josh Gottlieb, Jon Gruber, Jonathan Holmes, Vicky Marone, David Meltzer, Joe Newhouse, Mark Shepard, Johannes Spinnewijn, Seth Zimmerman, and seminar participants at Harvard Health Care Policy, Yale Public/Labor Economics Workshop, University of Chicago Harris, George Washington University Public Health, ASHEcon, NTA, and NBER Insurance meetings.

[†]UT-Austin and NBER. Email: mike.geruso@austin.utexas.edu

[‡]Harvard University and NBER. Email: layton@hcp.med.harvard.edu

[§]UC-Berkeley and NBER. Email: leive@berkeley.edu

1 Introduction

In both private insurance markets and social insurance programs that allow optional “top-up” coverage, adverse selection is a chronic problem. Typically, adverse selection distorts prices upwards, causing consumers who would generate positive social surplus from buying insurance to not buy it (Einav, Finkelstein and Cullen, 2010). Economists thus view adverse selection as a problem of inefficient sorting, a problem that can be fixed by shifting certain consumers from not purchasing to purchasing or by shifting certain consumers from one plan to another. Corrective policies focus on inducing such shifts in enrollment.

However, adverse selection does not just cause consumers to re-sort. It also induces implicit wealth transfers across consumers. Upward price distortions caused by adverse selection cause the same insurance product to cost more, leading to a wealth transfer away from (inframarginal) consumers who would choose to purchase insurance at both the distorted and undistorted prices. The prior literature on adverse selection has largely ignored these transfers, labeling them “welfare irrelevant,” as they have no effect on social surplus under predominant social welfare functions used in this literature, which put equal weight on a dollar given to any consumer. It is important to understand that the common form of the social welfare function used in the insurance literature stands in contrast to the evaluation of social welfare in other domains of the public finance literature, where the value of transfers is often allowed to depend on an individual’s marginal utility of consumption (Piketty and Saez, 2013a,b; Lieber and Lockwood, 2019) or on other factors (Saez and Stantcheva, 2016; Sher, 2024).

To fix ideas, consider a consumer whose willingness-to-pay (WTP) for generous health insurance is sufficiently high to make their choice to enroll in generous coverage inframarginal (invariant) to any relevant change in the price of generous coverage—because they are very risk averse, because they are wealthy, or for another reason. By construction, the consumption choices of a number of low-WTP, low-cost-to-insure consumers has no impact on this consumer’s enrollment decision. But the choices of consumers in the low-cost group may matter a great deal for the *price* the high-WTP consumer pays because these choices affect average plan costs: Average plan costs and therefore prices are lower when the low-cost group enrolls and higher when they do not. As a result, adverse selection will affect the private surplus for the high-WTP consumer, but will have no effect on the social surplus generated by *this* consumer. The price distortion induced by adverse selection is

merely a transfer away from the high-WTP consumer. Understanding these selection-induced transfers and their distributional impacts is important, as the key policy response to address selection distortions involves (often substantial) subsidies and mandates. Although these corrective policies could increase social welfare under standard formulations, they may also have important distributional consequences that conflict with societal goals around the distribution of resources. In other words, adverse selection and its corrective policies may involve classic equity-efficiency trade-offs.

In this paper, we analyze these trade-offs. In particular, we characterize the incidence of adverse selection, with a particular focus on the conditions under which the burden of selection will disproportionately fall on high-income households (progressive selection) or on low-income households (regressive selection). We then empirically assess whether these transfers are quantitatively important using administrative data on health insurance choices.

We begin by showing that the key factor determining whether selection is progressive or regressive is the (unconditional) correlation between income and demand for insurance. While a simple textbook treatment of risk averse agents typically yields the prediction that willingness-to-pay for insurance should decline with income or wealth, actual insurance market choices are complex and consumer choice is imperfect (see, e.g., [Handel and Kolstad, 2015](#); [Polyakova, 2016](#); [Ericson and Starc, 2016](#); [Abaluck and Gruber, 2023](#)). Indeed, there is evidence in some settings that higher income predicts higher health insurance demand (as in [Mahoney, 2015](#), [Finkelstein, Hendren and Luttmer, 2019](#), and [Finkelstein, Hendren and Shepard, 2019](#)). [Gropper and Kuhnen \(2025\)](#) find wealthier consumers also purchase more life and property insurance, in contrast to a standard model in which higher wealth or income reduces insurance demand. This correlation is therefore an empirical question.

We illustrate these concepts in a modified version of the graphical model of [Einav, Finkelstein and Cullen \(2010\)](#), in which we focus on consumer surplus. In that standard framework, it is straightforward to show that the inframarginal consumers, whose choices are *not* distorted by adverse selection, lose more surplus due to selection than the marginal consumers lose. Thus, while the incidence of selection with respect to *social* surplus falls primarily on the marginals, the incidence of selection with respect to *consumer* surplus falls primarily on the inframarginals. Similarly, just as the study of the welfare (social surplus) effects of selection typically ignores the *inframarginal consumers*, the study of the distributional effects of selection can typically ignore any effects on *marginal consumers*, as the consumer surplus of the marginals is unaffected by selection-induced price distortions by the

envelope theorem. Making this theoretical point about the selection-induced transfers to and from inframarginal consumers is the paper's key conceptual contribution.

We then provide evidence of a strong gradient of demand for insurance in income in a large employer setting where we can link information on consumer health insurance choices with consumer income. Employees have significant variation in income, but face the same insurance plan choice set. We find that in our setting around 67% of employees with annual salaries over \$120,000 choose the more generous health insurance plan, compared to only 44% of employees with salaries below \$35,000. These patterns hold among employees overall as well as among new employees who do not face inertia in their choices. Inframarginal enrollees in the generous plan are thus more likely to be higher-income, and, according to our framework, those higher-income enrollees bear more of the burden of selection (and benefit more from corrective subsidies).

To more precisely quantify the incidence of selection, we estimate a discrete choice model of insurance demand using the administrative data on consumer healthcare expenditures and insurance choices. We then simulate counterfactual changes in enrollment and prices under alternative subsidy regimes. We find that, relative to the setting where there is no incremental subsidy for the more generous option, the observed level of the corrective subsidy increases surplus of those in the highest income group by twice as much as for those in the lowest income group. Specifically, employees earning over \$120,000 receive \$710 in surplus from the subsidy compared to \$330 in surplus for employees earning less than \$35,000. In our setting, adverse selection combines with income differences in demand for more generous insurance to create an equity-efficiency trade-off: reducing the efficiency losses from adverse selection involves a subsidy that disproportionately benefits higher-income consumers.

Such a disparity in the incidence of subsidies used to correct selection was, to our knowledge, previously overlooked. It is not common for studies of insurance market selection to report the correlation between income (or wealth) and insurance choice. This omission is largely because the types of administrative or public data that contain information on insurance choice or insurance claims often lack links to income. When income data is available, any correlation is most often either ignored given the modeling assumptions (such as constant absolute risk aversion) or an unreported nuisance parameter. A small number of prior studies have reported this correlation, often incidentally. We provide a brief survey of these papers, which include settings in health insurance (primary and sup-

plemental), long-term care insurance, flood insurance, property insurance, and life insurance. Our review indicates that across a range of different markets and income levels, the incidence of adverse selection is often borne by higher-income consumers.

Our results thus provide a new perspective on the problem of adverse selection in social insurance programs. In markets and programs where demand for insurance increases in expected claims costs but also in income, we expect “progressive selection,” where adverse selection inefficiently distorts prices upward but the incidence of those price distortions falls disproportionately on higher-income consumers and households. In such settings, costly interventions to correct the distortions (such as subsidies) might, on the margin, be less desirable because they disproportionately benefit these higher-income groups.

2 Distributional Incidence of Selection

We briefly describe the incidence of adverse selection in insurance markets using a series of figures representing various cases of selection. All figures consider the simple setting where consumers choose whether or not to purchase an insurance product, as in [Einav, Finkelstein and Cullen \(2010\)](#).¹ The product consists of a fixed set of characteristics and a single uniform price charged to all consumers (i.e., community rating). That price is assumed to be set in a competitive equilibrium and is thus equal to the average cost of the set of consumers who opt to purchase insurance at that price. A consumer’s willingness-to-pay for insurance is assumed to reveal their valuation of insurance (i.e., the demand curve is equal to the consumer’s benefit curve). This assumption simplifies discussion of consumer surplus, but is not needed for evaluating the incidence of selection. There is adverse selection if the cost to insure higher willingness-to-pay consumers exceeds the cost to insure lower willingness-to-pay consumers.²

Panel A of [Figure 1](#) replicates Figure 1 from [Einav, Finkelstein and Cullen \(2010\)](#), showing the welfare loss due to adverse selection. Consumer types s are ordered on the x-axis according to their willingness-to-pay for insurance, yielding a unit demand curve. $c(s)$ represents the average cost of consumers of type s , i.e., all consumers who value insurance at $WTP(s)$. Consumers with $WTP(s) > P$ purchase insurance, while consumers with $WTP(s) < P$ do not. $s_m(P)$ represents the marginal

¹The same logic applies to settings where choice is between more and less generous coverage (such as our empirical setting), but it becomes more difficult to illustrate the key concepts graphically.

²See [Einav, Finkelstein and Cullen \(2010\)](#) for additional details of the model, including key regularity conditions.

consumer type, who is indifferent between purchasing or not purchasing insurance at price P , i.e., $WTP(s_m(P)) = P$. Finally, $AC(s_m(P))$ represents the average cost across all consumers who purchase insurance at price P , or, equivalently, when the marginal consumer type is $s_m(P)$.

The competitive equilibrium price occurs where $WTP(s)$ crosses $AC(s_m)$: the average cost of consumers purchasing insurance at price P^{eqm} is equal to that price, thus satisfying the zero profit condition. At that price, s_m^{eqm} is the marginal consumer type. However, this price does not maximize welfare. When insurers charge the equilibrium price P^{eqm} and s_m^{eqm} is the marginal consumer, consumers with $s > s_m^{eqm}$ value insurance more than the cost of providing it to them, as represented by $WTP(s)$ exceeding $c(s)$, but do not purchase it. Indeed, all consumers with $s < s_m^*$ value insurance more than the cost of providing it to them and thus would generate positive social surplus by purchasing insurance. The welfare loss equals the gap between $WTP(s)$ and $c(s)$ for the consumers with $s_m^{eqm} < s < s_m^*$, depicted by the red triangle in Panel A of Figure 1. This type of welfare loss has been the focus of much of the empirical literature on adverse selection: selection distorts the equilibrium price above the welfare-maximizing price, and so price inefficiently sorts some consumers out of insurance.

Panel B shifts the focus from *social* surplus to *consumer* surplus to show how selection affects different types of consumers. All curves remain the same, as do the equilibrium and welfare-maximizing prices. But now we illustrate forgone private/consumer surplus under the selection-induced equilibrium price relative to the welfare-maximizing price for each s -type. There is no effect of selection on private surplus for consumers with $s > s_m^*$: these inframarginal consumers would not purchase insurance at either the equilibrium price or the welfare-maximizing price. Consumers with $s_m^{eqm} < s < s_m^*$ are hurt somewhat by selection—they forgo surplus equal to the difference between their valuation of insurance $WTP(s)$ (which is always less than P^{eqm}) and the welfare-maximizing price P^* . Consumers with $s < s_m^{eqm}$, however, are hurt the most by selection, forgoing surplus equal to the full gap between P^{eqm} and P^* .

The key insight is that *inframarginals*, whose choices are not distorted by adverse selection, are hurt more by selection than marginals. This is true despite the fact that the *marginals* have been the main (implicit) focus of an adverse selection literature concerned with *social* surplus. More formally, the negative effects of adverse selection on consumer surplus are weakly monotonically increasing in the consumer's willingness-to-pay for insurance. This increase is strictly monotonic among the

marginal consumers but weakly monotonic overall because all inframarginal insured consumers experience an identical loss in surplus while all inframarginal consumers who never purchase H are unaffected.

The distributional implications of this insight thus depend on the relationship between willingness-to-pay and the consumer characteristic of interest. If high-willingness-to-pay consumers are disproportionately high income, then the burden of adverse selection falls on them. Appendix A also presents figures describing the incidence of selection for cases where (1) adverse selection causes the insurance contract to unravel completely and (2) the contract is advantageously selected. In the case of advantageous selection, the distributional incidence is flipped, with the benefits of advantageous selection accruing more to the highest-WTP consumers than to the lowest-WTP consumers. See Appendix A for a full description.³

These figures clarify that distributional incidence of selection depends on two factors: (1) the direction of selection (adverse versus advantageous) and (2) the correlation between willingness-to-pay for insurance and the stratifying characteristics of interest, such as income. If selection is adverse and income is positively correlated with willingness-to-pay, then selection is progressive: selection negatively affects higher-income consumers more than lower-income consumers. In practice, this seems plausible: If insurance is a normal good (i.e., WTP is higher for high-income consumers) conditional on health status, and health is correlated enough with WTP to cause adverse selection but not sufficiently correlated with WTP to offset the higher conditional demand for insurance among the rich, WTP will be positively correlated with both cost and income, and adverse selection will be progressive, hurting the rich more than the poor.

If, on the other hand, selection is adverse and income is *negatively* correlated with willingness-to-pay, then selection is regressive: selection negatively affects low-income consumers more than high-income consumers. If selection is advantageous, these results are flipped. These insights indicate that assessing the incidence of selection requires determining the joint distribution of willingness-to-pay for insurance, expected costs, and income.

In Appendix A and Table A.1, we survey the literature for (rare) cases in which empirical studies of insurance choice reported information sufficient to recover our correlation of interest. Though

³These insights do not necessarily hold if demand does not reflect consumer valuation. Frictions may cause marginal consumers to value insurance above cost. If the wedge between demand and valuation is large enough, marginal consumers could plausibly lose more surplus than inframarginal consumers due to selection-induced price distortions.

none of these studies aimed at assessing—or even addressing the possibility of—progressivity or regressivity, we show that in cases of health insurance (primary and supplemental), long-term care insurance, property insurance, life insurance, and flood insurance, there are economically meaningful correlations between income, willingness to pay, and market-level selection that show the relevance of our conceptual framework.

3 Evidence from a Large Employer

To empirically examine this phenomenon, we draw on the administrative records from a large public university that employs over 25,000 people for whom we link income and insurance market choices. The parameters we estimate from this employer are, of course, specific to it. Our purpose in the empirical exercise is to illustrate that the costs of selection can be disproportionately borne by the rich, and that the magnitude of this progressivity can be significant.

3.1 Data and Setting

The setting includes employees in a range of occupations with substantial variation in salary. The university includes faculty, administrators, scientists, physicians, nurses, medical technicians, and other staff. Appendix [Figure A.3](#) presents statistics on the salary distribution.

The administrative data span 2011 to 2017 and report salary, demographics, health insurance choices, and annual health care spending of each employee and dependent. Data on salary is collapsed into bins of \$5,000 intervals. Demographic information includes employee gender and age collapsed into bins (generally of five-year intervals). We also observe category of employment (faculty versus staff), division of the university (academic or medical), and the hiring date for each employee.

The university offered employees a choice between health plans. There were two traditional plans—a higher and lower coverage option. Starting in 2014, the employer added a high-deductible plan (HDHP) with a Health Savings Account (HSA). We label these plans as high-coverage (H), medium coverage (M), and low-coverage for the HDHP (L). Each plan had the same provider network and were financially differentiated based on premiums and cost-sharing parameters. The plans were all relatively generous: The actuarial value was 88% for H , 85% for M , and 77% for L .

Employee contributions to premiums in H increased substantially during the 7 years of our study period (2011-2017). Employee premiums for employee-only coverage rose from \$588 to \$1,275, and employee premiums for family coverage rose from \$4,584 to \$6,066 over this period. There were also minor increases in deductibles and out-of-pocket limits, but the rise in premiums was the main change to insurance contracts over time.

The employer aggregated the claims data to the annual level for each employee and dependent to protect confidentiality. We observe the component of annual health spending paid by insurance and the component paid out-of-pocket by employees. We also observe an indicator for whether an employee has one of several chronic conditions as recorded on the insurance claims, which enables us to construct measures of employee health.⁴ Summary statistics of the key variables are provided in [Table A.2](#).

3.2 Case Study Descriptive Statistics

The key correlations from Section 2 are between demand for more generous insurance coverage, spending risk, and income. Panel A of [Figure 2](#) plots enrollment in the H plan for each income bin. Income and enrollment in H are positively correlated: While only 57% of individuals with incomes below \$35k enrolled in H , almost 75% of individuals with incomes above \$120k enrolled in H . Panel B of [Figure 2](#) shows this pattern also holds when restricting to new enrollees, indicating that this correlation is not due to some artifact of different defaults or choice sets over time combined with differential tenure and inertia. In short, Panels A and B of [Figure 2](#) show that there is a strong correlation between income and demand for H .

Panel C shows that enrollment in H is also higher among individuals with a chronic health condition: 80% of individuals with a chronic condition choose H compared to only 57% of individuals with no chronic health condition. Further stratification by both health status and income indicates that much of the correlation between income and demand for H is driven by healthy individuals: Among individuals with a chronic condition, there is little difference in enrollment in H by income, but among individuals with no chronic condition there is a large difference in enrollment. More-generous insurance thus seems to be a normal good, conditional on health. Broadly, two types of employees choose H : the sick and the rich.

⁴The chronic conditions are hypertension, diabetes, hyperlipidemia, asthma, rheumatoid arthritis, ischemic heart disease, and chronic obstructive pulmonary disease (COPD).

In addition to differences in unit demand for (enrollment in) H by income, there is evidence of differences in price sensitivity by income. Panel A of [Figure 3](#) shows that H 's incremental premium increased markedly between 2011 and 2017, from around \$1,325 to around \$1,680. Over this period, the figure shows H 's overall market share declined significantly, from 80% to below 50%, but unequally by income. As H 's market share decreased, the average income of the remaining (infra-marginal) H enrollees increased significantly, from about \$68,100 to \$84,400. Over this same period, the average claims cost of the remaining H enrollees also increased significantly, from about \$7,700 to \$11,700. The increase in average cost across *all* employees—from about \$7,000 to \$8,800—was less than half the increase among H enrollees. These patterns imply that H was adversely selected (overall and on the price margin). Finally, the last plot shows an increase in the average age of H enrollees, suggesting that at least part of the change in income and health status may be due to changes in the age composition of H enrollees.

Panel B of [Figure 3](#) further clarifies the relationship between income and demand for more generous insurance. Here, we plot changes in H 's market share over time by income bin. Willingness-to-pay is almost monotonic in income, with the lowest income group always exhibiting the lowest levels of enrollment in H and the highest income group always exhibiting the highest levels of enrollment in H . Low- and high-income groups' price-enrollment curves differ in slope in addition to level. While H 's market share for the lowest income group dropped by 37 percentage points, H 's market share for the highest income group dropped by 25 percentage points, both in response to the same increase in the incremental premium. The right plot in Panel B shows the same statistics but restricting to new enrollees. These results are noisier (as there are fewer new enrollees), but the same pattern holds—the highest income group always exhibits higher demand for H than the lower income groups.

These results collectively indicate that there is strong adverse selection against H in this market and there is a strong positive correlation between income and demand for H . Therefore, in this setting the burden of selection is likely to fall disproportionately on higher-income consumers: Selection pushes up the price of more generous coverage, and this price distortion is disproportionately borne by higher willingness-to-pay inframarginal consumers, who are more likely to be higher income.

Of course, higher-income employees may also be sicker and have more healthcare costs; thus, they need not have lower *marginal utility* than lower income employees. Progressivity or regressivity

are typically defined relative to income rather than (unobservable) marginal utility, but the distributional impacts across patients of varying healthcare needs is of independent interest. In Appendix A, we show that the positive correlation between income and demand for H holds conditional on prior spending and age. Controlling for prior healthcare spending actually has little influence on the positive correlation between income and demand (Table A.3). By contrast, controlling for age substantially reduces the correlation, though it remains statistically significant. Since the positive correlation between demand for H and income persists when controlling for both prior spending and age, selection is both unconditionally progressive *and* progressive conditional on health status and age.

3.3 Estimating the Distributional Consequences of Selection

Exactly how progressive is adverse selection in this setting? To provide a quantitative answer to this question, we consider for each income group how consumer surplus under the equilibrium price P^{eqm} differs from consumer surplus under the efficient price P^* . Calculating surplus requires knowing P^{eqm} , P^* , $WTP(s)$, and $c(s)$ (see Figure 1). In this section, we estimate a model of demand for insurance as in Abaluck and Gruber (2011, 2016), Heiss et al. (2013), and Ericson and Sydnor (2026) to recover these parameters.

Demand Model: We estimate a conditional logit model of plan choice that specifies utility as a linear function of salary, premiums, expected out of pocket payments, plan characteristics, and individual-level characteristics:

$$U_{ijt} = \delta_j \cdot x_{it} + \beta_0 \cdot \pi_{jt} \cdot f(y_{it}) + \beta_1 \cdot \mu_{ijt} \cdot f(y_{it}) + \beta_2 \cdot \sigma_{ijt}^2 + \xi \cdot z_{jt} + \eta \cdot \mathbb{1}(j = j_{t-1}^*) + \epsilon_{ijt} \quad (1)$$

where i indexes employees, j indexes plan, and t indexes years. Employee characteristics x_{it} include indicators of \$20,000 salary bins, 5-year age bins, above-median tenure with the employer, gender, academic division, faculty, and employee-only coverage. These characteristics may shift demand for each plan as denoted by δ_j . We then include a number of plan-specific variables: π_{jt} denotes premiums in plan j in year t , μ_{ijt} denotes expected out-of-pocket payments for employee i in plan j in year t and σ_{ijt}^2 denotes the variance of those payments. To flexibly model demand by income, we include interactions between these plan characteristics and a polynomial in income, denoted by

$f(y_{it})$. We use a second-order polynomial, though our results are not sensitive to the degree used. z_{jt} includes additional plan characteristics—the deductible and out-of-pocket maximum—that may influence demand even conditional on expected out-of-pocket payments. These terms may capture liquidity constraints in a reduced-form way, in contrast to a more structural setup that explicitly models the dynamics of payments throughout the year and borrowing constraints (Ericson and Sydnor 2026). To capture the role of inertia in plan choices, $\mathbb{1}(j = j_{t-1}^*)$ is an indicator for employees choosing plan j in the previous year. Finally, ϵ_{ijt} is an i.i.d. error term with a type I extreme value distribution that captures unmodeled shocks, such as employee misperceptions of contract features or errors in forecasting spending risk.

We calculate the mean and variance of out-of-pocket payments in each plan by estimating the distribution of expected out-of-pocket costs for each plan $F_i(OOP_{ij})$ in the standard way (e.g., as in Abaluck and Gruber, 2011; Handel, Hendel and Whinston, 2015; Handel and Kolstad, 2015). We divide the full population into prior spending-by-gender-by-age cells, applying the non-linear cost-sharing schedule for the plan to each individual’s total costs to get the out-of-pocket cost under the plan, and assume that for each individual in a given cell, $F_i(OOP_{ij})$ is the ex ante distribution of out-of-pocket costs they face. Additional details about the construction of out-of-pocket spending distributions are presented in Appendix A and the regression results from estimating Equation 1 are presented in Table A.4. The model predictions of H ’s market share in each year match the data closely (Figure A.6).

Using these parameter estimates, Figure 4 simulates an overall demand curve (in blue) for H versus M ,⁵ as well as demand curves for each income group. These simulations include all employees but remove the effect of inertia, simulating demand when all employees make an active choice. These curves show that our model captures differences in demand by income, with the demand curve for the highest income group lying everywhere above the curve for the lowest income group. All income groups include some consumers with very high willingness-to-pay, such that the demand curves converge at low levels of s . The same qualitative pattern between demand and salary is observed for each coverage type (Figure A.5).⁶

⁵For simplicity, we restrict our simulation and counterfactuals to the choice between the two traditional plans, corresponding to the pre-2014 menu, because plan L initially had very low enrollment. Measuring the overall welfare effects in a model with L is also complicated by frictions that many consumers face in acquiring and processing information about the HSA see (see Leive, Friedberg and Davis, forthcoming). Moreover, our focus is on the incidence of selection-induced transfers by income, which for inframarginal consumers does not depend on demand revealing valuation.

⁶Some of the demand curves also cross below the x-axis, indicating negative WTP for H versus M . These are found for

Equilibrium Subsidy: In order to perform counterfactual simulations where we remove the employer’s incremental subsidy to H versus M , R_H , we need to first define and estimate that subsidy. From the data, we directly observe the net-of-subsidy price of H , P_H^c , but we do not know the gross price, P_H , i.e., the price that employees would face if the employer did not differentially subsidize H versus M . Both this price and the incremental subsidy to H are equilibrium objects that depend on how consumers would sort in the absence of the subsidy. To determine the gross price and the incremental subsidy, we follow the prior literature in assuming that prices will be set to satisfy the zero-profit condition (here, a break-even condition for the employer plan). Price will equal the average cost of the employees choosing the plan at that price. This makes the price an equilibrium parameter, determined by the demand $W_H(s)$ and the cost curves $c_H(s)$. In Appendix A we provide a detailed derivation and description of how we use the observed price and the demand and cost curves to find the incremental subsidy, R_H . We find that $R_H = \$467$ per year. With this subsidy, the model implies that H ’s overall market share at the equilibrium price is around 60%. This market share ranges from 45% for the lowest income group to above 80% for the highest income group. These match the observed shares fairly well.

Counterfactual Simulation of Subsidy Removal: We are now set up to perform a counterfactual simulation of removing the incremental subsidy for H (R_H), and instead providing a single fixed subsidy R that is constant across insurance choices. Adverse selection distorts the equilibrium price upward, resulting in consumers who place high value on H not enrolling in it. The purpose of the incremental subsidy R_H is to “correct” the price of H for adverse selection and induce (closer to) efficient sorting of employees across plans. To assess the effects of adverse selection on consumer surplus, we thus compare surplus with no subsidy (where the effects of adverse selection are unrestrained) versus surplus with the corrective subsidy provided by the employer (where the effects of adverse selection are weakened by the corrective subsidy). This comparison reveals the overall effects of selection on consumer surplus, and how those effects are distributed across income groups.

To perform this counterfactual simulation, we set R_H to zero and find the price where $P_H^e = AC_H(s_H^e) - R$, where R is the fixed subsidy for purchasing any insurance. We find that without the incremental subsidy R_H , the average cost curve $AC_H(s_H)$ intersects the demand curve at a market

lower-income levels and reflect the influence of the quadratic polynomial in income and the absence of an error term in the simulation.

share of about 40%, implying that the incremental subsidy increases H 's market share by 20 percentage points.⁷ This share again varies by income: While 28% of employees earning less than \$35,000 would still buy H without the subsidy, that share rises to 55% for employees earning between \$75,000 to \$120,000 and 65% for employees earning above \$120,000. Higher-income consumers are much more likely to be inframarginal than lower-income consumers.

To fully assess the distributional consequences of adverse selection and the corrective subsidy, we also need to know who the marginal enrollees are and how much surplus they gain from enrolling in H at the subsidized price. Panel A of [Figure 5](#) plots H 's market shares without the subsidy (dark gray bars) and with the subsidy (light gray bars) for each income group. The difference in the height of the bars reflects the sizes of the marginal group for each income group. The proportion of marginal enrollees is similar across income groups, but smaller among those earning the highest salaries. The large differences in the size of the price-insensitive inframarginal enrollees therefore already suggests that higher income groups are hurt more by adverse selection and benefit more from the corrective subsidy. Further, the similarity of the size of the marginal groups across income groups indicates that it is unlikely that accounting for the surplus gained by the marginals reverses much of the disproportionate harm of selection on the rich. Additionally, for any individual marginal enrollee, the maximum harm of selection is bounded above by the magnitude of the harm to each inframarginal enrollee, as [Figure 1](#) shows.

To determine the total forgone surplus due to adverse selection for each income group, we use estimates from the demand model to determine how much surplus is lost by the marginals when moving from the case with the corrective incremental subsidy to the case without it. Panel B of [Figure 5](#) shows the total forgone surplus (inframarginals + marginals) for each income group, averaged across everyone in the group. As expected, higher income groups are hurt more by selection. On average, an employee in the highest income group loses \$710 in surplus when removing the corrective subsidy, while an employee in the lowest income group loses \$330 of surplus. This means that the average high-income employee loses more than *twice* as much surplus due to adverse selection, compared to the average low-income employee. By comparison, the average change in *social* surplus is \$322 ([Figure A.7](#)).

⁷The subsidy reduces the incremental price consumers face from \$2,604 to \$1,704.

Counterfactual Simulation of No Subsidy versus Optimal Subsidy: We can also assess the full forgone surplus comparing the “no subsidy” case to the “optimal subsidy” case. The optimal subsidy is the amount that leads aggregate demand to intersect the marginal cost curve. We assume that there is no moral hazard, so the marginal cost of improved coverage is zero for all consumers. In our case, marginal cost is below demand even when the incremental price is zero. The optimal subsidy therefore induces enrollment in H by income at the shares depicted along the horizontal axis in [Figure 4](#). Forgone surplus from the optimal subsidy versus the observed subsidy varies across income groups for two reasons: (1) the size of the marginal group differs across income groups and (2) the valuation of H versus M of the marginal enrollees differs across income groups.

Panel C of [Figure 5](#) shows the average forgone surplus by income group in moving from no subsidy to the optimal subsidy. Here, the forgone surplus by income differs by more in dollars than with the observed subsidy, though it is smaller in percentage terms. The subsidy increases social surplus by \$568, on average, but induces increases in private surplus of \$2,287 for the highest income group and \$1,357 for the lowest income group. The highest income group thus loses over 68% more consumer surplus than the lowest income group due to adverse selection. Adverse selection is therefore significantly progressive in this setting. Correcting the price distortion using the typical policy remedy—a subsidy—would be regressive: The rich would benefit more.

[Appendix A](#) evaluates whether our conclusion about incidence is sensitive to key modeling assumptions. First, the two subsidy counterfactuals require us to simulate price changes out of sample. When considering a counterfactual based on in-sample price variation only, we again find that higher-income employees lose twice as much surplus from adverse selection as low-income employees—\$307 vs. \$157 ([Figure A.8](#)). Next, we allow for a wedge between valuation and demand that may differ across income groups. For the direction of incidence to be reversed, marginal consumers with the lowest incomes would have to act as if they under-value insurance (relative to their true valuation) by about \$2,200 more than marginal consumers with the highest incomes. The result that high-income consumers are harmed more by selection therefore appears robust.

4 Conclusion

The incidence of adverse selection falls disproportionately on the rich in the case study we examine here. The distributional consequences of removing the subsidy for the adversely selected plan are

large: high-income employees are hurt more than twice as much as low income employees. More generally, there is little reason to expect no correlation between willingness-to-pay and income in any market. This implies that selection—and the policies aimed at addressing it—will, in general, be redistributive.

Our paper advances the study of selection markets by providing a framework to assess the incidence of selection, which has largely been ignored in studies of selection markets. Rather than a mere theoretical possibility or curiosity, our estimates imply that the magnitude of progressivity can be substantial. Costly corrective actions may therefore not be as socially desirable as the prior literature suggests: Correcting the distortions of selection markets may introduce a difficult trade-off between reducing inefficiency and increasing inequality.


References

- Abaluck, Jason, and Jonathan Gruber.** 2011. "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program." *American Economic Review*, 101(4): 1180–1210.
- Abaluck, Jason, and Jonathan Gruber.** 2016. "Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance." *American Economic Review*, 106(8): 2145–2184.
- Abaluck, Jason, and Jonathan Gruber.** 2023. "When Less Is More: Improving Choices in Health Insurance Markets." *Review of Economic Studies*, 90(3): 1011–1040.
- Bhargava, Saurabh, George Loewenstein, and Justin Sydnor.** 2017. "Choose to lose: Health plan choices from a menu with dominated option." *The Quarterly Journal of Economics*, 132(3): 1319–1372.
- Cutler, David M, and Sarah J Reber.** 1998. "Paying for health insurance: the trade-off between competition and adverse selection." *The Quarterly Journal of Economics*, 113(2): 433–466.
- Einav, Liran, Amy Finkelstein, and Jonathan Levin.** 2010. "Beyond testing: Empirical models of insurance markets." *Annual Review of Economics*, 2(1): 311–336.
- Einav, Liran, Amy Finkelstein, and Mark R. Cullen.** 2010. "Estimating Welfare in Insurance Markets Using Variation in Prices." *The Quarterly Journal of Economics*, 125(3): 877–921.
- Einav, Liran, Amy Finkelstein, and Neale Mahoney.** 2021. "The IO of selection markets." In *Handbook of Public Economics*. Vol. 5, 389–426. Elsevier.
- Einav, Liran, and Amy Finkelstein.** 2011. "Selection in insurance markets: Theory and empirics in pictures." *Journal of Economic Perspectives*, 25(1): 115–38.
- Ericson, Keith Marzilli, and Justin R Sydnor.** 2026. "Liquidity constraints and the value of insurance." *American Economic Journal: Microeconomics*, 18: 426–453.
- Ericson, Keith M Marzilli, and Amanda Starc.** 2016. "How product standardization affects choice: Evidence from the Massachusetts Health Insurance Exchange." *Journal of Health Economics*, 50: 71–85.
- Fang, Hanming, Michael P Keane, and Dan Silverman.** 2008. "Sources of advantageous selection: Evidence from the Medigap insurance market." *Journal of political Economy*, 116(2): 303–350.
- Finkelstein, Amy, and Kathleen McGarry.** 2006. "Multiple dimensions of private information: evidence from the long-term care insurance market." *American Economic Review*, 96(4): 938–958.
- Finkelstein, Amy, Nathaniel Hendren, and Erzo FP Luttmer.** 2019. "The value of medicaid: Interpreting results from the oregon health insurance experiment." *Journal of Political Economy*, 127(6): 2836–2874.
- Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard.** 2019. "Subsidizing health insurance for low-income adults: Evidence from Massachusetts." *American Economic Review*, 109(4): 1530–67.
- Gropper, Michael J, and Camelia M Kuhnen.** 2025. "Wealth and Insurance Choices: Evidence from US Households." *Journal of Finance*, 80: 1127–1170.
- Handel, Ben, Igal Hendel, and Michael D Whinston.** 2015. "Equilibria in health exchanges: Adverse selection versus reclassification risk." *Econometrica*, 83(4): 1261–1313.


- Handel, Benjamin R, and Jonathan T Kolstad.** 2015. "Health insurance for "humans": Information frictions, plan choice, and consumer welfare." *American Economic Review*, 105(8): 2449–2500.
- Heiss, Florian, Adam Leive, Daniel McFadden, and Joachim Winter.** 2013. "Plan Selection in Medicare Part D: Evidence from Administrative Data." *Journal of Health Economics*, 32(6): 1325–1344.
- Leive, Adam, Leora Friedberg, and Brent Davis.** forthcoming. "Links between puzzles in household finance: evidence from employee benefit choices." *American Economic Journal: Economic Policy*.
- Lieber, Ethan, and Lee Lockwood.** 2019. "Targeting with In-Kind Transfers: Evidence from Medicaid Home Care." *American Economic Review*, 109(4): 1461–1485.
- Mahoney, Neale.** 2015. "Bankruptcy as implicit health insurance." *American Economic Review*, 105(2): 710–46.
- Piketty, Thomas, and Emmanuel Saez.** 2013a. "Optimal labor income taxation." In *Handbook of Public Economics*. Vol. 5, 391–474. Elsevier.
- Piketty, Thomas, and Emmanuel Saez.** 2013b. "A theory of optimal inheritance taxation." *Econometrica*, 81(5): 1851–1886.
- Polyakova, Maria.** 2016. "Regulation of insurance with adverse selection and switching costs: Evidence from Medicare Part D." *American Economic Journal: Applied Economics*, 8(3): 165–95.
- Saez, Emmanuel, and Stefanie Stantcheva.** 2016. "Generalized social marginal welfare weights for optimal tax theory." *American Economic Review*, 106(1): 24–45.
- Sher, Itai.** 2024. "Generalized Social Marginal Welfare Weights Imply Inconsistent Comparisons of Tax Policies." *American Economic Review*, 114: 3551–3577.
- Wagner, Katherine R. H.** 2022. "Adaptation and Adverse Selection in Markets for Natural Disaster Insurance." *American Economic Journal: Economic Policy*, 14(3): 380–421.

Figure 1: Social and Private (Consumer) Surplus Under Adverse Selection

(a) Social Surplus




(b) Private (Consumer) Surplus



Notes: Figures plots demand and cost curves for insurance. Consumer types s are ordered along the x-axis according to their willingness-to-pay. Panel A shows the social and consumer surplus under adverse selection. The efficient allocation is at s_m^* where demand $WTP(s_m)$ intersects marginal cost $c(s)$, but the equilibrium is at s_m^{eqm} where demand intersects average cost $AC(s_m)$. The shaded triangle in Panel A denotes the standard efficiency loss from adverse selection. Panel B shows the consumer surplus that is lost in moving from the efficient allocation to the equilibrium and distinguishes between foregone surplus for inframarginal consumers and foregone surplus for marginal consumers.


Figure 2: Enrollment in *H* by Income




Notes: Figure plots the percentage of employees choosing *H* by income level. Panel (a) shows choices of all employees during the sample period. Panel (b) shows initial choices for new employees hired during the sample period. Panel (c) plots the percentage of employees choosing *H* split by salary level (above/below \$60,000) and diagnosis of a chronic condition (yes/no). Health status, as measured by chronic conditions, is strongly predictive of *H*. Those with one or more chronic conditions are more likely to choose *H* than another plan, regardless of their income. Chronic conditions include hypertension, diabetes, hyperlipidemia, asthma, rheumatoid arthritis, ischemic heart disease, and chronic obstructive pulmonary disease (COPD). By contrast, those without a chronic condition are about much more likely to choose *H* if they earn above \$60,000: the difference in plan shares is about 10 percentage points. This graph suggests that demand for *H* is positively related to income, even conditional on health status.

Figure 3: Trends in H

(a) Trends Across All Enrollees




(b) Trends by Income




Notes: Figure plots variation in employee premiums, market share for H , income, costs, and age among people who choose H over time. Incremental premiums reflect annual differences between the H and M premiums. The bottom panel splits trends in market share for H by income level, both for all employees and for new employees separately. Means are calculated by averaging across all types of coverage (employee-only, employee plus child, employee plus spouse, family coverage).

Figure 4: Simulated Demand and Cost Curves

(a) Demand Curves by Income




(b) Overall Demand and Cost



Notes: Panel A plots simulated demand curves for H separately by income level (shaded gray lines), and for the full sample (in blue). Price reflects the incremental price between H and M . Panel B plots the demand curve (solid line) and average cost curve (dotted line) for H using the model estimates. The curves intersect at the equilibrium market share of 60% for H . The demand curve is inclusive of the employer's subsidy as described in the text.

Figure 5: Counterfactual Simulations



Notes: Figure presents consumer surplus from the subsidy for H under the existing employer subsidy in panel (b) or the optimal subsidy in panel (c). The optimal subsidy induces enrollment in H by income at the shares depicted in [Figure 4](#) when the incremental price is zero. Panel (a) plots the simulated market shares from the model to illustrate differences in take-up by income level at the equilibrium price under the existing subsidy.

A Supplementary Analysis [Online Appendix]

Unraveling: Figure A.1 shows the effects of adverse selection on social and consumer surplus for the case where the market fully unravels. As discussed in Einav and Finkelstein (2011), full unraveling occurs when the average cost curve lies everywhere above the demand curve. In Figure A.1 we present the unraveling case where it would be efficient for everyone to be enrolled in insurance, i.e., where the demand curve is everywhere above the marginal cost curve. Panel A shows the welfare loss due to adverse selection, which is represented by the gap between the demand curve and the marginal cost curve for all consumers. Panel B shows the lost private/consumer surplus. Here, because all consumers generate positive social surplus from enrolling in insurance, we define P^* as the highest price at which all consumers choose to enroll, or $P^* = WTP(s = 1)$. In this case, all consumers are marginal consumers, and a consumer of type s forgoes surplus of the amount $WTP(s) - P^* = WTP(s) - WTP(s = 1)$ due to adverse selection.


In both Figure 1 and Figure A.1, the negative effects of adverse selection on consumer surplus are (weakly) increasing in the consumer's willingness-to-pay for insurance. In Figure A.1 the monotonicity is strict, as all consumers are marginal. It is useful to contrast this pattern to the impacts of selection on social surplus, which is highest for the lowest willingness-to-pay types and monotonically decreasing in willingness-to-pay, reaching zero for the highest willingness-to-pay type ($s = 0$).

Advantageous selection: Figure A.2 illustrates social and private surplus under advantageous selection. As in the main text, consumer types s are ordered on the x-axis according to their willingness-to-pay for insurance. $WTP(s)$ represents the demand curve, which reveals consumer valuation of insurance in dollars. $c(s)$ represents the average cost of consumers of type s , i.e., all consumers who value insurance at $WTP(s)$. Consumers with $WTP(s) > P$ purchase insurance, while consumers with $WTP(s) < P$ do not. $s_m(P)$ represents the marginal consumer type, who is indifferent between purchasing or not purchasing insurance at price P , i.e., $WTP(s_m(P)) = P$. Finally, $AC(s_m(P))$ represents the average cost across all consumers who purchase insurance at price P , or, equivalently, when the marginal consumer type is $s_m(P)$.

Panel A shows that advantageous selection distorts prices downward instead of upward ($P^{eqm} < P^*$), hurting overall social welfare due to *over*-consumption of insurance. Panel B shows that the benefits of advantageous selection are weakly monotonically increasing in willingness-to-pay for insurance, meaning that the consumers with the highest willingness-to-pay experience the largest benefits. Thus, the incidence of both adverse selection and advantageous selection is largest for the highest willingness-to-pay consumers, but that incidence is negative in the case of adverse selection (high WTP consumers are hurt the most) and positive in the case of advantageous selection (high WTP consumers benefit the most).

Figure A.1: Social and Private Surplus Under Adverse Selection - Unraveling

(a) Social Surplus



(b) Private Surplus






Figure A.2: Social and Private Surplus Under Advantageous Selection

(a) Social Surplus



(b) Private Surplus



Evidence from the Literature on Distributional Incidence of Selection: Table A.1 lists papers that have reported a correlation between income and insurance choice. To construct the table, we began with the set of papers cited in either of two review articles that summarize the empirical literature on selection markets: “Beyond Testing: Empirical Models of Insurance Markets” (Einav, Finkelstein and Levin, 2010) and “The IO of Selection Markets” chapter in the Handbook of Industrial Organization (Einav, Finkelstein and Mahoney, 2021). We supplemented these sets with other recent empirical studies of selection markets. From this frame, we narrowed attention to papers in which income (or wealth) data and product choice were apparently available to the researchers. If a sign and significance of the income-insurance purchase correlation at the individual level was available in the published paper or appendix, we included it in the table.

Table A.1: Correlation Between Income and Insurance Choice in the Literature

Study (1)	Context (2)	Income Range in Sample (3)	Correlation: Income and Takeup of Adversely Selected Option (4)	Correlation: Income and Takeup of Higher- Priced Option (5)	Distributional Incidence of Subsidy to Higher Price Option (6)
Cutler & Reber (1998)	Health - ESHI	(<45k, >75k)	+	+	Regressive
Finkelstein & McGarry (2006)	Long Term Care	Rep. US sample	–	+	Regressive
Fang, Keane, & Silverman (2008)	Medigap	(5k, >50k)	–	+	Regressive
Bhargava, Loewenstein, & Sydnor (2017)	Health - ESHI	(<20k, >100k)	–	–	Progressive
Finkelstein, Hendren, & Shepard (2019)	Health - Nongroup	(16k, 27k)	+	+	Regressive
Wagner (2022)	Flood	(<50k, >300k) HH	+	+	Regressive
Gropper & Kuhnen (2025)	Property	(<19k, >130k)	n/a	+	Regressive
Gropper & Kuhnen (2025)	Life	(<19k, >130k)	n/a	+	Regressive

Notes: Table lists studies of adverse selection containing information on the relationship between income and insurance demand across a variety of insurance domains. Incomes are at the individual level except for Wagner 2022, which is household income. The income range listed is imprecise, but conservative. For example, Gropper & Kuhnen (2025) do not list the full income range but report the 25th and 90th percentiles of earned income in their data, respectively, as \$18,916 and \$132,081, which we report as <\$19,000 to >\$130,000.

Table A.1 includes research covering a wide variety of settings, including health insurance (primary and supplemental), long-term care insurance, property insurance, life insurance, and flood insurance. Column 3 makes clear that these studies are not limited to narrow or unrepresentative income ranges. Column 4 indicates whether the adversely selected contract options are also the options preferred by higher income consumers. Column 5 indicates whether higher income consumers in these settings exhibit higher willingness-to-pay for “more” insurance—either on the extensive margin of taking up insurance, or on the intensive margin of taking up a more expensive insurance option within the insurance domain.

Overall, Table A.1 lends credence to the possibility that adverse selection is likely to be progressive in some settings, in the sense that the burden of the price distortions it creates fall disproportionately on higher income consumers (who disproportionately take-up the costlier insurance products). This survey of the empirical literature suggests that what determines whether adverse selection is progressive or regressive is whether the higher-priced insurance option in a market is the adversely selected option (as in the case of flood insurance) or not (as in the case of Medigap). This re-examination of the literature corroborates recent work by Gropper and Kuhnen (2025), which conducts a deep examination of the statistical correlation between income and wealth and the demand for life and property insurance. In column 6 of Table A.1, we highlight that subsidizing the higher-priced insurance option would be regressive in almost all of the surveyed studies.

Descriptive statistics of sample: Figure A.3 displays the number of employees into the 8 groups of salary ranges that we use throughout the analysis. Most of these categories have at least 4,000 employees. These 8 groups are constructed by collapsing narrower bins of \$5,000 increments of salary to aid interpretation.

Figure A.3: Number of Employees by Income Bin




Table A.2 presents summary statistics of the analytic sample. Mean income is \$72,313, which is higher than the U.S. average, and the standard deviation is \$44,851. There is substantial variation in salary within the university. Given the academic setting, both age (45.8 years) and tenure with the university (10.3 years) are also higher than the average of the U.S. workforce. Nearly 59 percent of employees are female. There are slightly more employees in the academic division compared to the medical division (56.7 percent vs. 43.3 percent). Over the entire sample period, most employees choose *H*, and the least popular option is *L* (the HDHP/HSA). The final two rows list the number of unique employees (25,056) and the number of employee-years (101,672).

Table A.2: Summary Statistics

	Mean	S.D.
Income (\$)	72,313	44,851
Age (years)	45.82	12.68
Academic division	56.7%	49.5%
Faculty	21.7%	41.3%
Medical division	43.3%	49.5%
Tenure (years)	10.26	10.16
Female	58.7%	49.2%
Household size	2.02	1.28
Total health spending (\$)	8,081	29,202
Plan <i>L</i>	6.9%	25.3%
Plan <i>M</i>	30.8%	46.2%
Plan <i>H</i>	64.7%	47.8%
<i>N</i>	25,056	
<i>NT</i>	101,672	

Plan Choice Regressions: Table A.3 shows results of linear probability models (LPMs) of choosing *H* against indicators for income levels. The estimates reveal a positive correlation between income and demand for *H*. Without controlling for other covariates, employees earning \$35,000 to \$45,000 are 8.1 percentage points more likely to choose *H* compared to employees earning less than \$35,000 (column 1). Estimates are of similar magnitudes for employees earning between \$45,000 and \$75,000. The predicted probability of choosing *H* are 14.2 percentage points higher for those earning \$75,000 to \$95,000 compared to those earning less than \$35,000. Estimates are nearly identical for those earning \$95,000 to \$120,000. The probability of choosing *H* is highest among those earning over \$120,000, at 19.7 percentage points above rates for those with the lowest incomes. Standard errors clustered by employee are in parentheses, and are small relative to the point estimate.


Column (2) adds the expected difference between *H* and *M* as a control, and estimates coefficients slightly smaller in magnitude but still statistically significant. Estimates are similar when also adding the standard deviation of the difference in out-of-pocket payments between *H* and *M* (column 3). The estimates drop substantially when adding age (column 4), with the only statistically significant difference between the highest and lowest income groups. The positive correlation is again strong when including lagged health spending instead of age (column 5) and when also adding indicators for faculty, academic division, and chronic condition status (column 6).

Plots of Plan Choices, Income, and Health Status: Figure A.4 shows a contour plot that builds on Figure A.3 idea by showing the percentage of employees choosing *H* by income group and decile of expected total health spending. The importance of income in plan choices can be seen by comparing plan shares in top of the graph (higher incomes) to those in bottom of the graph (lower incomes). For those earning less than \$35,000, it is not until the sixth decile of expected health spending that over 60 percent of employees choose *H*. Among those earning over \$120,000, by contrast, over 60 percent of employees in the first decile of expected spending choose *H*.

Table A.3: Linear Probability Models of High Coverage Plan Choice

	(1)	(2)	(3)	(4)	(5)	(6)
Income bin (relative to < \$35k)						
\$35k-\$45k	0.081 (0.009)	0.070 (0.009)	0.068 (0.009)	0.010 (0.009)	0.067 (0.010)	0.019 (0.010)
\$45k-\$55k	0.064 (0.009)	0.065 (0.010)	0.063 (0.010)	0.006 (0.009)	0.061 (0.011)	0.018 (0.011)
\$55k-\$65k	0.036 (0.010)	0.050 (0.010)	0.047 (0.010)	0.005 (0.010)	0.050 (0.011)	0.020 (0.011)
\$65k-\$75k	0.080 (0.010)	0.069 (0.011)	0.066 (0.011)	0.003 (0.010)	0.067 (0.012)	0.015 (0.012)
\$75k-\$95k	0.142 (0.010)	0.122 (0.010)	0.118 (0.010)	0.016 (0.010)	0.117 (0.011)	0.027 (0.011)
\$95k-\$120k	0.143 (0.011)	0.122 (0.011)	0.117 (0.011)	0.009 (0.011)	0.119 (0.012)	0.022 (0.012)
> \$120k	0.197 (0.011)	0.174 (0.011)	0.169 (0.011)	0.032 (0.012)	0.167 (0.012)	0.047 (0.014)
Expected OOP Difference (\$1,000s)		0.319 (0.011)	0.300 (0.010)	0.298 (0.010)	0.262 (0.011)	0.250 (0.010)
Std. dev of OOP Difference (\$1,000s)			0.021 (0.004)	-0.028 (0.004)	0.048 (0.005)	-0.002 (0.005)
Lagged health spending	No	No	No	No	Yes	Yes
Age bins	No	No	No	Yes	No	Yes
Other demographics	No	No	No	No	No	Yes
Observations	101522	90803	90803	90803	70924	70924
R^2	0.028	0.046	0.046	0.122	0.056	0.132

Figure A.4: Contour Plot of H by income and expected health spending



Construction of Out-of-Pocket Cost Distributions: This section details the procedure for constructing distributions of out-of-pocket costs for each employee and dependents. The approach is based on grouping people into “risk groups” according to demographics and previous health spending, and then to use the empirical distribution of out-of-pocket (OOP) payments among people in each risk group as a measure of beliefs.

We first divide each insured individual according to five discrete age bins (younger than 30, 30–39, 40–49, 50–59.5, 59.5 and older) and gender (male, female). Within these groups, we further split into terciles based on 1-year lags of total health spending, combining both plan paid spending and OOP spending. We classify people with the same grouping of age, gender, and cost tercile as being in the same risk group. For new hires, we do not observe their lagged spending, so we assign them to the tercile with median spending that is closest to their observed spending.

To construct the distribution of out-of-pocket spending under plan j for people in risk group g , we take the distribution of observed spending of people within risk group g who chose plan j . We assign this distribution to people in risk group g who chose a different plan $j' \neq j$.

To give an example, we group women aged 30–39 together, rank them by their total health spending in year $t - 1$, and divide them evenly into terciles based on year $t - 1$ spending. Within each tercile, we further split them based on their observed plan choice in year t . The empirical distribution of OOP for each of the coverage levels is taken as the OOP distribution for each woman in that sub-group if she had chosen that coverage level.

The final step is to combine OOP distributions of each member of the family. We implement this by taking 100 draws for each employee or dependent from their group-specific OOP distribution under each plan, and sum each of the 100 draws across all family members to arrive at a distribution of OOP costs for the family. If the sum of OOP within families for any draw exceeds the plan’s OOP max, we replace the OOP for that draw as the OOP max. This distribution of 100 OOP draws represents the family’s belief about OOP risk under each available plan.

In constructing each OOP distribution, we pool multiple years together. Doing so ensures that each risk group based on age, gender, lagged cost tercile, and plan choice has a sufficiently large number of individuals.

For simplicity, we assume draws are independent within families. Draws might be positively correlated if family members have similar tastes for health care consumption that we do not model. On the other hand, OOP draws (not necessarily spending draws) might be negatively correlated due to the non-linear nature of the insurance contract. We believe modeling these correlations would introduce unnecessary complexity into this calculation without providing meaningfully different results.

Imputation of marginal tax rates: The empirical analysis accounts for the tax deductibility of employer-sponsored insurance premiums. Our administrative records lack several pieces of information required for a direct calculation of the employee’s marginal tax rate, including information about spousal earnings, children, other sources of income, home ownership, and relevant deductions. In addition, marital status is reported incompletely and salary is recorded in bands to protect data confidentiality. Our approach is therefore to calculate marginal tax rates for respondents of the American Community Survey (ACS) using the National Bureau of Economic Research’s TAXSIM, and then to use hot-deck imputation to assign a marginal tax rate for the employees in our sample by matching on income, age, and gender.


Step 1: ACS data We use ACS surveys between 2011 and 2017, which record relatively comprehensive information that helps us calculate marginal tax rates. In particular, we use the following information from the survey: wage and salary income of respondent and spouse, interest received, retirement income and social security benefits, supplemental security income and public assistance income, state, marital status, age, number of dependents, and number of children under 13.

Step 2: Marginal tax rate calculation For each ACS observation, we use NBER TAXSIM to estimate the federal and state marginal tax rates based on the variables in the list above.

Step 3: Hot-deck imputation We match individuals between our administrative data and the ACS by year, age band, income band, and gender. We then use hot-deck imputation to assign a marginal tax rate to the matched employees in our sample. The imputation is repeated five times and we take the average to construct our estimate of the employee's marginal tax rate.


Demand curves by coverage type: Figure A.5 presents separate demand curves by salary for each coverage type. Similar to the pattern shown in Figure 4, demand varies monotonically with income for employees across coverage types, ranging from employee-only to family coverage. Price reflects the incremental price between H and M .

Figure A.5: Demand Curves in H by Income by Coverage Type



Model fit: Figure A.6 evaluates the fit of the model by plotting the share of employees predicted to choose H on the y-axis against the observed share choosing H on the x-axis. Each point represents the average for a particular year, as labeled, and the scatterplots are split by income for visual clarity. The solid line denotes the 45-degree line, which denotes the benchmark of perfect model fit. In general, the fit is quite good across salary groups and years.

Figure A.6: Model Fit by Salary and Year



Conditional logit results: Below we present a subset of the parameter estimates from estimating Equation 1. We present estimates on the variables that differ by plan (characteristics of the choices). The regression also includes characteristics of the individual: indicators for salary bins (\$20,000), age (5-years), gender, academic vs. medical division, faculty, above-median tenure, and lags of previous plan choices that shift choices in each plan (characteristics that vary by individual). For ease of interpretation, we divide premiums, deductibles, income and expected spending by \$1,000. Coefficients estimates reported are the parameters of the utility function, not marginal effects. Standard errors clustered by employee reported in parentheses.

Table A.4: Conditional Logit Results

	(1)
Employee Premium	-3.131 (0.718)
Expected out-of-pocket costs	-0.486 (0.174)
Variance of out-of-pocket costs	-2.972 (22.652)
Out-of-pocket limit	-0.263 (0.057)
Deductible	-2.991 (0.141)
Income $\times$ premium	-0.000 (0.000)
Income $\times$ (premium) ²	-0.008 (0.012)
Income $\times$ expected out-of-pocket	0.005 (0.003)
Income $\times$ (expected out-of-pocket) ²	-0.000 (0.000)
Employer HSA contribution	0.160 (0.182)
Observations	210,860


Derivation of Incremental Subsidy: Start by letting P_H and P_M represent the gross premiums for H and M , respectively. Also let $W_H(s)$ represent incremental willingness-to-pay for H vs. M from the demand model. As before, s is an index that orders employees from 0 to 1 according to their willingness-to-pay for H versus M , and $W_H(s)$ is a function that gives the willingness-to-pay (the price at which the employee is indifferent between purchasing and not purchasing H versus M) for a type- s employee. Let $c_H(s)$ and $c_M(s)$ represent the cost curves for H and M , or the average cost across employees of type s in plan H and M , respectively. R represents a subsidy for purchasing either plan, and R_H represents an additional incremental subsidy for choosing H .

We make a number of assumptions. First, we assume that P_M is fixed and equal to $P_M = AC_M(s = 1)$, where $AC_M(s = 1)$ represents the average cost in M if all employees enroll in M . Further, we let $R = P_M = AC_M(s = 1)$ so that the net-of-subsidy price of M is $P_M^c = P_M - R = 0$. This assumption is an approximation of how the pricing of the lower-generosity option works at employers. R_H is the incremental subsidy for H that is set by the employer. It is this incremental subsidy that we wish to estimate. P_H is assumed to be set endogenously to equal the average cost of the employees who choose H at a given net-of-subsidy price $P_H^c(P_H, R_H) = P_H - (R + R_H)$, making P_H an equilibrium parameter. At a given net-of-subsidy price $P_H^c(P_H, R_H)$, employees with $W_H(s) > P_H^c(P_H, R_H)$ choose H and all other employees choose M . Define $s_H(P_H^c(P_H, R_H))$ as the marginal s -type such that $W_H(s_H) > P_H^c(P_H, R_H)$. The average cost of the employees choosing H given net-of-subsidy price $P_H^c(P_H, R_H)$ is thus given by $AC_H(s_H) = \int_0^{s_H} c_H(s) ds$. And the equilibrium price of H is given by $P_H^e = AC_H(s_H^e) - (R + R_H)$ where s_H^e is the equilibrium marginal s -type such that $s_H^e = s_H(P_H^c(P_H^e, R_H))$.

In practice, we know P_H^e , but we do not know R_H . Backing out R_H requires additional assumptions. First, there is no moral hazard. This assumption is common in the structural literature on adverse selection, and it is likely to be a reasonable approximation. This assumption implies that $c_H(s) = c_M(s) = c(s)$, and it allows us to compute $R = P_M = AC_M(s = 1)$ and $AC_H(s_H)$ for all possible values of s_H (i.e., it allows us to draw the average cost curve). Now, given values for R , P_H^e , the demand curve $W_H(s)$, and the average cost curve $AC_H(s_H)$, we can find the value for R_H that makes the equilibrium price expression $P_H^e = AC_H(s_H^e) - (R + R_H)$ hold. To do so, we find the price where the demand curve plus the overall subsidy $W_H(s) + R + R_H$ crosses the average cost curve $AC_H(s_H^e)$. This is shown in Panel B of Figure 4 in the main text.


Change in social surplus from counterfactuals: Figure A.7 evaluates the foregone social surplus by income group for the two counterfactuals considered in the main text. Panel (a) shows the foregone social surplus in moving from the employer subsidy to no subsidy and Panel (b) shows the foregone social surplus in moving from the optimal subsidy to no subsidy. Unlike the case with consumer surplus, the changes in social surplus are flatter and not monotonic by income group.

Figure A.7: Change in social surplus for counterfactuals



Counterfactual using in-sample price variation: Figure A.8 evaluates the change in consumer surplus under a scenario in which the (incremental) price for H increases by the observed increase between 2011 and 2017. We apply this counterfactual to employees in 2013, the same year as the counterfactual simulations in the main text. The price increase is \$372, averaged across all employees. High-income employees lose about twice as much surplus as lower-income employees from this price increase: the estimated change in consumer surplus is \$307 for employees earning over \$120,000 compared to \$156 for employees earning below \$35,000. The patterns by income are therefore qualitatively similar to the changes under the two subsidy counterfactuals that increase prices using out-of-sample variation.

Figure A.8: Change in consumer surplus for counterfactual using in-sample price variation



Wedge between Valuation and Demand for Insurance: The theoretical framework and empirical analysis assume that demand reflects a consumer's valuation of insurance. While it is possible that frictions could lead consumers to value insurance more than the cost, any such wedge would have to differ across income levels to affect our results regarding the incidence of adverse selection. We assess the sensitivity of our result to this assumption by calculating the size of the differential wedge between valuation and demand by income that would be needed for lower-income consumers to be harmed more by adverse selection than higher-income consumers. We do so by calculating the amount that the demand curve would have to be shifted up for low-income marginal consumers (to reflect their true valuation), while keeping demand for high-income marginal consumers fixed. Marginal consumers are those who do not choose *H* at the unsubsidized price (\$2,640) but choose *H* at the subsidized price (\$1,704). [Figure A.9](#) plots the demand for these marginal consumers by income group, with the x-axis showing the fraction of marginal consumers among that particular income group.

For adverse selection to harm low-income employees more than high-income employees, marginal consumers earning below \$35,000 would need to under-value insurance by \$2,223 more than marginal consumers earning above \$120,000 do. This amount is quite large in magnitude, and is in between the subsidized and unsubsidized prices as shown on the graph. The amount is so large because marginal consumers contribute a small amount to private surplus for both income groups. High-income employees lose \$379 more in surplus than low-income employees, and over 90% of this difference is due to differences in surplus among inframarginal consumers. Any wedge between valuation and demand for the marginals must therefore differ by income to such a great extent to outweigh the much larger number of high-income inframarginal consumers. The qualitative result that higher-income employees are harmed more by adverse selection is robust to plausible differences by income group in any wedge between the valuation and demand for insurance.

Figure A.9: Demand of marginal consumers between prices with and without employer subsidy

