Reducing Medicare Costs by Accounting for Selection

Ira Abraham*

Click here for the latest version

November 14, 2025

Abstract

The U.S. Medicare Advantage (MA) program was designed to harness competition among private insurers and the cost-containment potential of managed care to reduce spending and improve care for beneficiaries. However, imperfect risk adjustment—where payments are linked to beneficiaries' health—fails to fully account for the selection of healthier individuals into Medicare Advantage. This generates a trade-off between higher costs from selection into MA and lower costs from managed-care efficiency. Using detailed administrative claims data combined with additional survey measures of health, I develop a structural model that jointly captures individuals' plan choices and insurers' equilibrium price setting. The model links these demand and supply behaviors to quantify how their interaction affects government expenditure. Counterfactual simulations show that eliminating MA would reduce government spending by $20.8 billion annually. Adjusting for selection beyond current risk adjustment into the subsidy design would realign the MA program with its cost-saving intent, reducing CMS expenditure by $7.56 billion while resulting in a combined loss of $0.585 billion in consumer surplus and producer profits.

*Boston College (email:ira.abraham@bc.edu). I am grateful to Michael Grubb, Kit Baum, and Diana Bowser for their invaluable guidance and support. I also thank Guillermo Nunez, David Escalante, and the Research Protections team at Boston College for assistance with data access. Financial support to purchase data from the Economics Department and the Connell School of Nursing at Boston College is gratefully acknowledged.

1 Introduction

Medicare is the U.S. public health insurance program for individuals aged 65 and older and for some younger individuals with disabilities, end-stage renal disease (ESRD), or amyotrophic lateral sclerosis (ALS). The program is administered by the Centers for Medicare & Medicaid Services (CMS), a federal agency within the Department of Health and Human Services. When Medicare began in 1965, the federal government acted as the direct insurer under what is now called Traditional Medicare (TM). Over time, a parallel option—Medicare Advantage (MA)—was introduced, in which private insurers contract with CMS to provide the same core benefits. Beneficiaries thus choose between a government-run plan (TM) and private managed-care plans (MA).

The economic logic of Medicare Advantage is to harness competition and the managed-care efficiency of private insurers to benefit both enrollees and the government. CMS sets a benchmark rate based on the average cost of beneficiaries in Traditional Medicare. Private insurers submit bids that indicate the cost at which they can provide equivalent coverage. If a plan's bid is below the benchmark, the difference is shared: CMS keeps part of the savings, and the rest becomes a rebate that plans must use to lower premiums or enhance benefits. Payments are also risk-adjusted so that plans enrolling sicker individuals receive higher compensation. Risk adjustment is carried out by assigning each enrollee a risk score, which is a composite index reflecting predicted health spending based on diagnoses and demographic characteristics. The benchmark is then multiplied by the risk score to account for the health status of beneficiaries. This design makes MA plans attractive to enrollees—who often receive more generous benefits at lower cost—while allowing the government to capture some of the savings generated by competition and efficient care management.

Despite its design, evidence indicates that Medicare Advantage (MA) may lead to overpayments through two main channels, as highlighted by Congressional Budget Office (2022). One source of overpayment is "cream-skimming" (or advantageous selection), in which MA plans attract healthier beneficiaries than those in Traditional Medicare (TM), even after accounting for risk adjustment. Aizawa and Kim (2018) study the mechanism through which MA plans engage in cream-skimming and show that it occurs primarily through targeted advertising. Another source of overpayment is "upcoding", the tendency of private plans to record more health conditions than TM, resulting in higher risk score for individuals enrolled in MA. This paper focuses on the welfare effects of cream-skimming. Congressional Budget Office (2022) note that the magnitude of budgetary losses due to advantageous selection conditional on risk scores in MA is uncertain, given the difficulty of predicting plan pricing adjustments and beneficiary enrollment responses to benchmark changes. Given this policy relevance, I develop a structural model of plan demand, insurer supply, and government expenditure to quantify the welfare effects and the extent of overpayment. Geruso and Layton (2015) examine the welfare impacts of upcoding, which serves as a complementary analysis to the policy question studied in this paper.

I use detailed individual-level administrative data to develop a random coefficients demand model that captures the health composition of individuals who choose among different Medicare Advantage (MA) plans and Traditional Medicare (TM). Health composition is measured using two dimensions: the individual's risk score and a survey-based health indicator. The survey-based poor health indicator captures selection conditional on the risk score. I construct this variable using the Medicare Current Beneficiary Survey (MCBS), following the approach in Aizawa and Kim (2018). On the supply side, insurers are assumed to compete à la Bertrand to determine equilibrium premiums. The key mechanism in the model is that when payments to insurance plans change, equilibrium prices adjust, which in turn alters the health composition of individuals selecting into different plans based on their health characteristics. The government expenditure model depends on the health composition of enrollees across plans, as fiscal costs vary systematically with beneficiaries' health status.

I find that if all current MA enrollees were to switch to Traditional Medicare (TM), government expenditure would decrease by \$20.8 billion. This amount represents the budgetary loss in government spending attributable to selection into MA. In a subsequent counterfactual analysis, I predict the TM cost of covering a population based on individuals' risk scores, poor health indicators, and other demographic characteristics. I refer to this predicted cost as the selection-corrected benchmark. I find that in the majority of counties, the selection-corrected benchmark is lower than the current benchmark level, while in the remaining counties it is higher.

In the second counterfactual, I set the new benchmark as the minimum of the selection-corrected benchmark and the current benchmark, and analyze the individuals who switch from MA to TM as well as those who remain in MA. I compare the expenditure on switchers under MA at the baseline benchmark with their expenditure under TM. I find that MA generates a budgetary gain of approximately \$2.3 billion, driven by relatively unhealthy individuals switching out of MA. I then perform the same decomposition for the stayers and find that if the stayers were to switch to TM, it would result in a budgetary gain of about \$20.92 billion. This shows that the marginal enrollee in MA is relatively unhealthy and the selection is leading to overpayments in MA through the stayers.

In the third counterfactual, I set the benchmark equal to the selection-corrected benchmark. The selection-corrected benchmark yields budgetary savings of about \$7.56 billion for

CMS, reflecting the government's share of the managed-care cost savings. This scenario also leads to a modest reduction in combined consumer surplus and producer profits of around \$0.585 billion. These findings suggest that the current MA market is not performing as intended, as advantageous selection is offsetting the managed-care cost savings. However, when the benchmark is adjusted to the selection-corrected level, the MA program becomes budget-saving while quantifying the associated costs in terms of reduced consumer welfare and producer surplus. Overall, this framework provides a foundation for evaluating reforms to MA payment policy and highlights the trade-offs between budgetary savings and welfare.

2 Related Literature

I build on an extensive literature studying the Medicare Advantage (MA) market. Several recent papers examine how the program's subsidy design affects efficiency and welfare. Miller et al. (2023) analyzes how benchmarks should be reallocated across geographic regions to maximize consumer welfare under the current budget constraint. Curto et al. (2021) develops a structural model of plan bidding and finds that the MA program generates substantial surplus for participants, roughly two-thirds of which is captured by insurers. Mathew Zahn (2023) models firm entry and product-offering decisions in the MA market to study a counterfactual subsidy design in which consumers, rather than firms, receive the subsidy; he finds that this policy encourages greater plan entry. Relative to these studies, my paper differs in two key ways. First, I account for selection conditional on risk adjustment by incorporating survey data to quantify overpayments arising from such selection. Second, I capture consumer preferences over provider networks through a hospital-choice model. This feature is important because a central trade-off for MA enrollees is choosing a narrower private-plan network versus the broad network available in traditional Medicare. Finally, my framework builds on Aizawa and Kim (2018), who estimate a similar model to study advertising's role in equilibrium selection, but I additionally model network utility and analyze counterfactual subsidy designs aimed at mitigating overpayment from advantageous selection.

More broadly, my study relates to a growing literature on the Medicare Advantage (MA) market that study different features of the market. Brown et al. (2014) examines the effect of risk adjustment on risk selection using a reduced-form model and finds that MA plans tend to enroll individuals with higher risk scores but lower medical costs conditional on those scores. Geruso and Layton (2015) document that MA insurers engage in upcoding—strategically manipulating patient diagnoses to increase payments from Medicare. Stockley et al. (2014) show that when demand is unresponsive to lower premiums (the Part B buy-down), plans have an incentive to pass cost subsidies to consumers through more generous benefits that consumers may not value at cost, leading to inefficiently high benefit generosity. They attribute the limited premium responsiveness to the way premiums are displayed on the Medicare online portal. Curto et al. (2019) compare health care spending in MA and traditional Medicare (TM) and find that MA insurer revenues are roughly 30 percent higher than their health care spending. They show that lower spending in MA relative to TM primarily reflects lower utilization of services rather than lower prices for the same services. Duggan et al. (2016) exploit a policy-induced increase in MA reimbursement and find that higher reimbursements lead to greater private firm entry and higher MA enrollment, but that only about one-fifth of the additional reimbursement is passed through to consumers in the form of better coverage.

Song et al. (2013) show that a \$1 increase in the benchmark raises plan bids by about \$0.53, implying that plans possess meaningful market power. Pelech et al. (2019) analyze the decrease in benchmark payments to MA plans following the Affordable Care Act (ACA) and find corresponding declines in plan bids and beneficiary rebates-specifically, a \$1 decrease in the benchmark led to roughly a 60-cent decrease in bids. Chernew et al. (2023) study how changes in benchmarks affect plan characteristics and estimate that a \$1,000 annual reduction in benchmarks would increase annual premiums by about \$60 and annual deductibles by about \$27. Finally, Decarolis et al. (2020) examines subsidy design in Medicare Part D.

3 Data

I use administrative data from the Centers for Medicare & Medicaid Services (CMS) for the year 2019. The first data source is the Medicare Advantage (MA) Plan Characteristics File, which includes information on plan attributes such as premiums and the counties in which each plan operates. I supplement these data with publicly available CMS data on MA plans to obtain additional plan characteristics, including copayments, deductibles, out-of-pocket (OOP) maxima, quality ratings, market size (measured as the number of Medicare-eligible individuals in a county), plan market shares, and county level benchmark rates.

I also use the Master Beneficiary Summary File (MBSF), which contains individual-level data on a 5% random sample of Medicare beneficiaries. The MBSF records enrollment information for each beneficiary, including whether they are enrolled in Traditional Medicare or a Medicare Advantage (MA) plan. For those enrolled in TM, MBSF contains information on the total cost to Medicare and the beneficiary from different claims files like inpatient, outpatient, and other claims. This data is used to predict the TM cost of individuals in the supply side. The MBSF also includes detailed beneficiary characteristics such as age, gender, race, income measures, and county of residence. I use these data to construct the county-level distribution of beneficiary demographics for the random-coefficients estimation in the demand model.

In 2019, total Medicare enrollment was approximately 61 million, so the 5% sample includes about 3 million individuals. Roughly 37% of beneficiaries were enrolled in MA plans. I exclude ESRD patients from the analysis and also special needs plans, pace plans, employer-sponsored plans or PFFS plans which have a different reimbursement scheme from CMS which is different from the standard plans. At this time, only 2014 CMS risk score files are available for research use, so I proxy beneficiaries' 2019 risk scores using the 2014 distribution. To capture health heterogeneity beyond risk scores, I incorporate data from the Medicare Current Beneficiary Survey (MCBS), which is a nationally representative survey of approximately 15,000 Medicare beneficiaries conducted annually. Following Aizawa and Kim (2018), I use MCBS to account for selection conditional on risk score. I use the Inpatient Encounter Data, which contain MA claims for hospital stays, including ICD-10 diagnosis codes and the hospital identifiers where patients were treated. These data are used to estimate a hospital choice model and compute the network utility component of consumer preferences. Further details on the construction of the hospital-choice data are provided in Appendix 12.1 . In addition, I use the Inpatient, Outpatient, Skilled Nursing Facility (SNF), Home Health Agency (HHA), Durable Medical Equipment (DME), and Carrier Encounter files to calculate the average cost of care for Medicare beneficiaries in the counterfactual analysis.

4 Institutional Details

In the MA market insurers compete along the dimensions of benefit design, premiums, and provider networks, and often heavily market their plans. Plans generally offer a set of 'in-network' providers which enrollees may utilize with lower cost-sharing than 'out-of-network' providers. MA plans generally provide a more generous benefit package than TM, such as including dental, vision, and/or hearing coverage (DVH). Many plans include a drug benefit. Plans may also offer a reduction in the Part B premium.

The enrollee-specific subsidy from CMS to insurers is based on a "benchmark" rate for each county, which varies across geographies and over time and is not influenced by MA firms. CMS calculates the benchmark schedule each year using the average spending in the TM market within a county. Counties are ranked by average spending and placed into quartiles. The benchmark for counties in the top quartile is set to 95% of their FFS spending. The benchmark for the second quartile is 100% of FFS spending, the third quartile benchmark is 107.5% of spending, and the bottom quartile benchmark is 115% of spending. A floor that varies by urban/rural status applies. The higher benchmarks are intended to encourage health plan participation. Lower benchmarks in the higher spending areas are intended to contain spending.

Each year, after benchmarks are published by CMS, insurers submit detailed proposals to offer MA plans. That is, each insurer submits a 'bid' which include benefit and cost-sharing designs, and detailed information about the insurer's expected revenues and expenses for both TM-covered and non-TM-covered services. Ultimately, the 'bid amount' represents the insurer's offer to provide all services covered by TM to a person of average risk in the plan's coverage area in exchange for a particular level of revenue. The bid amount must be related to the plan's projected costs and may be above or below the benchmark rate. Insurers that bid above the benchmark must charge premiums to enrollees. Firms that bid below the benchmark receive a portion of the difference as a 'rebate' that must be passed on to consumers through decreases in cost-sharing (e.g. reductions in copays) or by offerings of services not covered under TM (e.g. dental). Supplemental benefits may also be paid for by an additional premium.

The rebate payment varies across carriers and over time based on the CMS 'star rating' measure of insurer quality. Under current policy, insurers with at least four stars (out of five) also receive a 5% bonus to the benchmark rate. The payment from CMS to insurers for an enrollee i living in county m enrolled in plan j in year t based on a benchmark B_{mt} and a plan bid b_{jt} is calculated as the product of the 'base' payment and the risk-scores as shown below.

$$base_{ijt} = \begin{cases} B_{mt} \times \phi_{jt} & \text{if } b_{jt} \geq B_{mt}\phi_{jt} \\ (b_{jt} + \lambda_{jt} \times (B_{mt} \times \phi_{jt} - b_{jt})) & \text{if } b_{jt} < B_{mt}\phi_{jt} \end{cases} \tag{1}$$

$$Payment_{ijt} = Risk_{it} \times base_{ijt}$$

where ϕ_{jt} captures any bonus to the benchmark rate and λ_{jt} is the rebate percentage. $Risk_{it}$ is the risk-score that is assigned to each individual and the payment made to insurer for an individual i's 'risk-adjusted' to reduce the incentive for selection on the basis of expected cost of covering an individual. The risk-scores is calculated using a comprehensive risk adjustment model called Hierarchical Conditions Categories (HCC). It is a predicted relative cost of beneficiary based on demographics (age, sex, Medicaid/ESRD/disability status) and past year diagnoses. In the Medicare program, a risk score of 1.0 represents the average expected cost for a beneficiary. Individuals with risk scores greater than 1.0 are expected to have higher-than-average medical spending, while those with scores below 1.0 are expected to have lower-than-average spending. These risk scores are computed using an additive model that combines demographic characteristics with Hierarchical Condition Category (HCC) coefficients, both of which are calibrated using TM data. However, even within a given risk score, there exists a distribution of actual medical spending, reflecting variation in health status that is not fully captured by the risk adjustment model. This residual variation gives rise to advantageous selection-a phenomenon in which beneficiaries with lower actual spending than predicted by their risk score are more likely to enroll in Medicare Advantage (MA) plans (Brown et al. (2014)).

5 Reduced form evidence on Selection

From the MCBS data, I construct a poor health indicator. This variable takes the value 1 if an individual reports both vision or hearing difficulties and physical or cognitive limitations.[1] The poor health indicator variable is used to capture residual health heterogeneity and account for selection conditional on risk score. To validate this measure, I estimate an OLS regression of total TM spending on demographic and health characteristics, including the risk score and the poor health indicator variable. Table 1 shows that conditional on risk score, individuals who take the value 1 for poor health indicator variable have a $2859.57 higher FFS cost ceteris paribus. Table 2 shows individuals who choose MA have higher risk

Table 1: Determinants of TM Cost

	Coefficient	Std. Error
risk score	5442.05***	(678.73)
poorhealth	2859.57**	(1227.40)
age	-67.63**	(29.63)
male	-177.56	(843.09)
white	961.27	(1104.02)
Constant	7393.84***	(2335.38)
Observations	1,966	
Adjusted R^2	0.03	

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

score but better health with respect to vision, hearing or cognitive abilities. The correlation

[1] These indicators take the value 1 if individuals report difficulty with vision or hearing, using the telephone, shopping, paying bills, memory loss that interferes with daily activities, trouble concentrating, or problems making decisions that interfere with daily activities.

coefficient between risk score and poor health indicator is 0.1887.

Table 2: Health Measures across TM and MA

	TM	MA	p-value*
risk score (full sample)	0.78	0.83	0.0000
risk score (survey sample)	0.77	0.85	0.0025
poorhealth (survey sample, %)	14.24	11.61	0.0410

* One-sided p-value for difference test.

To test for selection into MA, I run a linear probability model on whether individuals choose TM or MA on their demographic and health characteristics. Table 3 shows that individuals with higher risk scores are more likely to choose MA, and, conditional on risk score, those in better health (i.e., poor health indicator equal to zero) are more likely to enroll in MA. Consistent with the literature, I find advantageous selection into MA conditional on risk score.

Table 3: Selection into Medicare Advantage

	Coefficient	Std. Error
poorhealth	-0.0835**	(0.0325)
risk score	0.0505***	(0.0186)
age	0.0011	(0.0008)
male	-0.0402*	(0.0223)
white	-0.1085***	(0.0291)
constant	0.4429***	(0.0616)
Observations	2,032	
Adjusted R^2	0.016	

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

6 Demand Model

To estimate consumer preferences over Medicare Advantage plans, I specify the following random coefficient logit plan choice model. (The market subscript is not explicitly shown for ease of notation.) The outside option to the individuals is Traditional Medicare.

$$u_{ij} = \alpha(\mathbf{z_i}, N_{ij})1 + \beta(\mathbf{z_i}, N_{ij}) oop\,max_j + \gamma \mathbf{x_j} + +\xi_j + \epsilon_{ij},$$

where $oop\,max_j$ is the out-of-pocket maximum for plan j, which is assumed to be the only endogenous variable in the model, representing the plan's price component. This assumption is motivated by the fact that most of the variation in cost-sharing across plans is concentrated in the out-of-pocket maximum. While other papers in the literature typically treat premiums as endogenous, in the context of Medicare Advantage, most plans bid below the benchmark, resulting in a zero premium. The vector $\mathbf{x_j}$ includes exogenous plan characteristics such as premium, deductibles, copay, and quality rating. The constant has a random coefficient that captures heterogeneity in preferences for Medicare Advantage relative to Traditional Medicare. Observed individual characteristics $\mathbf{z_i}$ include age, gender, race, risk score, and a poor health indicator derived from survey data. N_{ij} is the network utility of individual i from plan j, which is computed from a separate hospital-choice demand estimation as described in Subsection 6.1

To address the endogeneity issue of the price variable, I use two instrumental variables. The first instrument is the average number of insurance plans offered by the same insurer in other markets, following Decarolis et al. (2020). The intuition is that if a firm offers more plans in other markets, it likely reflects broader firm-level characteristics—such as a larger customer base, greater administrative capacity, or a more aggressive pricing strategy—that influence the firm's overall pricing behavior. These factors can affect the price a firm sets in the market of interest. Importantly, market-specific demand or cost shocks are assumed to be uncorrelated across markets, ensuring that plan availability in other markets provides a valid source of exogenous variation in price.

The second instrument that I use following Mathew Zahn (2023) is the leave-one-out instrument which is defined as the following:

$$B_{j\setminus m} = \sum_{k \in A_j \setminus m} \omega_{jk} B_m$$

where A_j denotes the set of counties where plan j entered and ω_{jk} are weights based on the number of people enrolled in the plan. Benchmark that any plan faces is average of benchmark rates in counties served by plan weighted by projected plan enrollment. Each plan submits single sealed bid. So the plan generosity depends on benchmarks of all the counties the plan decides to enter and therefore the OOP Max will be correlated to the instrument of choice. The exogeniety is based on the assumption that the demand shocks in one market is uncorrelated with the demand shocks in the other market. The instrument will be uncorrelated to the demand shocks of county m since the benchmark of county m is excluded from the instrument. However, if the plans are choosing counties where the demand shock are correlated, then the exogeneity doesn't hold. This could be a possible explanation if counties are geographically near each other. As a robustness, I create the leave-one-out benchmark using the counties that are in a different state as compared to the rest of the counties. The results are robust under such a definition.

6.1 Hospital Choice Demand Model

The network-utility measure used in the demand estimation in Section 6 is derived from a separately estimated hospital-choice demand model. Following the approach established in Capps et al. (2003) and subsequent work, the network value measure NV_{ij} is obtained by estimating a hospital-choice model using the inpatient hospitalization dataset. The utility of patient i with diagnosis d for hospital h :

$$U_{idh} = \eta_h + \gamma X_h Z_i + \delta Dist_{ih} + \epsilon_{idh}$$

where η_h is a hospital fixed effect, $Dist_{ih}$ is the distance of hospital h from the centroid of the zip code of individual i and also includes distance-squared and and distance interacted with patient age, gender, and income group. Z_i are patient observable characteristics and X_h are hospital characteristics. Hospital characteristics variables are from American Hospital Association (AHA) Annual Survey Database. The patient observable x hospital characteristics variables is patient diagnosis category (using the top-level CCS category) interacted with hospital's service offerings (e.g., cancer patient $\times$ hospital has oncology services). The data creation for the estimation is described in the Appendix 12.1. Conditional on needing to be hospitalized for diagnosis d, the expected utility of plan j's hospital network (N_j) is calculated as follows:

$$\begin{aligned} HospEU_{idj}(N_j) &\equiv E\Big[\max_{h \in N_j}\{\hat{u}_{idh}(N_j) + \epsilon_{idh}\}\Big] \\ &= log\Big[\sum_h exp(\hat{u}_{idh}(N_j))\Big] \end{aligned} \tag{2}$$

where $\hat{u}_{idh}(N_j)$ is the estimated hospital utility excluding the logit error term. Many covariates that enter hospital utility are known at the time of plan choice (e.g., distance, and demographics). However, other variables are not realized until later: notably diagnosis category. I follow the literature and assume that consumers have expectations over these variables based on observed patterns in the data. Consumers have expectations for their hospital use frequency for each diagnosis d. I estimate these frequencies using a Poisson regression of the number of hospitalizations in the data (for a given diagnosis $\{d\}$) on age-sex and income groups. The predicted values ($\hat{p}_{id}$) from the regression is used as the probability weights to calculate the network utility as follows:

$$\begin{aligned} Network\ Utility &= Illness\ probability * HospEU \\ NV_{ij} &= \sum_d \hat{p}_{id}.HospEU_{id}(N_j) \end{aligned} \tag{3}$$

The estimation results for the model can be found in the Appendix 12.3.

7 Supply Model

The supply model follows the standard multi-product firm framework in the literature (Berry et al., 1993). Insurer f chooses a menu of prices (out-of-pocket maxima) for each plan j in the set J_{fm} that it offers in market m. Insurers' geographic locations and other product characteristics are taken as exogenous to simplify the analysis. Insurer f's profit in market m is the sum of the profits from all plans it offers in that market and is given by:[2]

$$\Pi_{fm} = N_m \sum_{j \in J_{fm}} \int_{h_i} [p_{fjm} + sub_{fjm}(r_i) - c_{fjm}(\mathbf{z_i})].s_{fjm}(\mathbf{z_i}) dF_m(\mathbf{z_i})$$

N_m is the number of Medicare beneficiary in market m, p_{fjm} is the premium charged by plan j in market m. The vector $\mathbf{z_i}$ contains all the individual characteristics including the risk score and the poor health indicator variable. sub_{fjm} is the subsidy received by the plan and defined as:

$$sub_{fjm} = min\{\underbrace{b_{fj}}_{bid}\bar{r}_j + \underbrace{\lambda_{fj}(B_{fjm} - b_{fj})\bar{r}_j}_{rebate}, B_{fjm}.\bar{r}_j\}$$

where b_{fjm} is plan's bid. B_{fjm} is the benchmark rate faced by the plan (adjusted for the quality rating of the plan). λ is the rebate percentage (which is also dependent in the plan's quality rating). r_i average risk score of individuals enrolled in plan j. Lastly, $s_{fjm}(\mathbf{z_i})$ is the probability of choosing insurer f and plan j by an individual having characteristics $\mathbf{z_i}$. CMS requires that this rebate rate is used to fund additional benefits in terms of cost-sharing reductions and supplemental benefits. I assume that this constraint is binding (following Miller et al. (2023)) in the sense that the rebate the firm receives is exactly equal to the incremental cost of providing the plan's cost-sharing benefits over the cost of providing a 'base' TM-equivalent plan. So the incremental cost only depends on the cost sharing

[2]Each plan has a "service area" that may span multiple counties. In such cases, the plan's benchmark is the service-area average of county benchmarks, weighted by projected plan enrollment. Despite this, following Curto et al. (2021), I model competition at the county level.

characteristics of the plan. That is: $Cincr(CS_j) = rebate_j$. This implies that :

$$\lambda(B_{fjm} - b_{fjm})\overline{r}_j = Cinc_{fjm} \tag{4}$$

Substituting for the bid:

$$\Pi_{fc} = N_m \sum_{j=1}^{L_{fc}} \int_{h_i} [p_{fjm} + r_i.B_{fjm} - \frac{1-\lambda}{\lambda} Cinc_{fjm} - c_{fjm}(\mathbf{z_i})].s_{fjm}(\mathbf{z_i}) dF_c(\mathbf{z_i})$$

The incremental cost has the following functional form $Cinc_{fjm} = -\gamma_x.cost\ sharing_{fjm}$. The marginal cost of the firm

$$c_{fjm}(\mathbf{z_i}) = \gamma_{TM}.\overline{TM}_m(\mathbf{z_i}) - \gamma_x.cost\ sharing_{fjm} + \delta_f + \eta_{fjm}$$

where $\overline{TM}_m(\mathbf{z_i})$ is the predicted average of of individuals $\mathbf{z_i}$ choosing plan j, δ_f is firm fixed effects. The unobserved η_{fjm} shock is possibly correlated with the observables related to the marginal consumer's health care cost. Following Aizawa and Kim (2018), I use the average reimbursement cost in TM in market m, $TMavg_m$, as an instrument. Further details of supply side are in Appendix 12.4.

8 Government Expenditure

The government expenditure consists of the subsidy paid to MA insurers and the cost incurred in the TM market. Government expenditure in market m is given below (the market subscript is omitted for ease of exposition):

$$\mathrm{GE}(B_j) = \underbrace{\sum_{j \in J_m} s_j \left[\overline{r}_j \cdot B_j - \frac{(1-\lambda)}{\lambda} \mathrm{Cinc}_j \right] N_m}_{\text{MA subsidy}} + \underbrace{\left(1 - \sum_{j \in J_m} s_j \right) N_m \, \overline{\mathrm{TM}}_m(z_i)}_{\text{TM cost}}.$$

Here, B_j is the plan-level benchmark rate, $\overline{r}_j$ is the average risk score of individuals enrolled in plan j, Cinc_j is the incremental cost (as discussed in Section 7), s_j is the share of individuals enrolled in plan j from the demand model, and N_m is the number of Medicare-eligible individuals in market m.

The predicted average TM cost, $\overline{\mathrm{TM}}_m(z_i)$, is obtained by calculating the probability that an individual chooses TM (the outside good) from the random-coefficients demand model and then computing the weighted average of individual characteristics to determine the mean z_i

among individuals who select the outside good. This average is combined with the estimated coefficients from Table 1 to predict the cost in the TM market.

The subsidy paid to MA plans consists of two components: the risk-adjusted benchmark rate, $\bar{r}_j \cdot B_j$, and a second term, $\frac{(1-\lambda)}{\lambda}\,\text{Cinc}_j$. If risk adjustment perfectly accounted for all health measures—so that there were no advantageous selection into MA—then $\bar{r}_j \cdot B_j$ would represent the cost of covering MA enrollees had they instead chosen TM. The term $\frac{(1-\lambda)}{\lambda}\,\text{Cinc}_j$ would then represent the savings CMS receives from MA, which I refer to as the *rebate savings* to CMS.

In the first counterfactual, I calculate the change in government expenditure if the MA program were shut down. This has two effects: first, there are savings because risk adjustment currently overpays the MA market; and second, CMS would lose the rebate savings. The results in Section 9 show that the savings outweigh the loss to CMS.

In the next counterfactual, I vary the benchmark level and analyze government expenditure, consumer surplus, and insurer profits. As the benchmark changes, the cost in the MA market changes directly, which alters equilibrium prices and shifts the demographic composition (including health) and share of consumers choosing MA versus TM. When the benchmark decreases:

$$\boxed{\Delta G = \underbrace{\sum_{\text{stayers}} \Delta\text{capitation}}_{\text{Intensive margin}} + \underbrace{\sum_{\text{switchers}} (\text{TM cost} - \text{MA capitation})}_{\text{Extensive margin}}}$$

Since we have established that, conditional on risk score, the average MA enrollee is healthier than the average TM enrollee, one might expect healthier individuals to switch back to TM when benchmarks fall. However, the results show that the marginal individuals who switch back are, in fact, less healthy, and most of the savings from lowering the benchmark come from the stayers.

9 Estimation Results

9.1 Demand estimates

Table 4 shows the demographic coefficient estimates on the coefficient and price. The random coefficient on the constant shows the preference of individuals on the inside good (MA plans) vs the outside good (TM). The table shows that older, men, non-white individuals prefer the MA plans over TM. Individuals with higher risk scores and a poor health indicator equal to zero prefer MA plans. The descriptive evidence showed that higher risk score individuals choose MA plans. This choice pattern can be rationalized if individuals with higher risk score might also prefer lower copay or deductible which MA plans offer but are not captured by the random coefficient. Individuals with higher risk score have a lower price sensitivity. Since the price variable here is out of pocket maxima, the sign of the price coefficient is not obvious. On the one hand we would expect unhealthy individuals (with higher risk score) to be less price sensitive. But since out of pocket maxima is modeled as the price, we would expect it to be more binding for individuals with higher risk score. The coefficient shows that the first channel is more prominent. Poor health indicator variable is more price sensitive than risk score. Older individuals and women are more price sensitive. The coefficients show that as network utility increases, individuals prefer inside good and individuals with higher network utility have lower price sensitivity.

Table 4: Demographic Coefficient Estimates

	Constant		Price	
Variable	Coef.	SE	Coef.	SE
Age	0.411***	(0.048)	-0.211***	(0.032)
Male	0.027***	(0.007)	0.355***	(0.004)
White	-0.100*	(0.047)	0.079	(0.237)
Risk score	-0.090**	(0.032)	0.892***	(0.027)
Poorhealth	-0.483***	(0.026)	0.039	(0.134)
Network utility	0.464***	(0.038)	0.123*	(0.058)

Robust standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 5 shows the estimates of the mean utility. Prices, premium coefficients are negative and significant as expected. Higher quality rating leads to higher mean utility as expected. Copay has a positive coefficient which is unexpected which may be due endogeneity is the variable which is assumed away in the model.

9.2 Supply estimates

The table below shows the marginal cost estimates. The coefficient of predicted TM cost implies that as the predicted TM cost rises rises by a dollar, the cost to the MA insurer rises by 0.70. This implies that MA plans have a 30 percent cost saving as compared to TM conditional on health. This estimate is consistent with the literature. Aizawa and Kim (2018) find that marginal cost of MA plan is about 35 percent less, on average, than the TM

Table 5: Mean Utility Estimates

Variable	Coefficient	SE
Constant***	-0.286	(0.011)
Prices (oop max)***	-1.518	(0.025)
Copay (Primary)***	0.258	(0.068)
Quality rating**	0.362	(0.129)
Deductible	-0.018	(0.011)
Premium***	-1.323	(0.106)

Robust standard errors.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

cost.

Table 6: Marginal cost estimates

Characteristic	**Estimate**	**SE**
$\overline{TM}_m(\mathbf{h_i})$	0.70	(0.53)
Copay (Primary)	-118.26	(80.14)
Deductible*	77.08	(37.43)
Constant*	-17853.47	(8664.08)
Firm FE	Yes	

Number of observations = 7,516
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

10 Welfare and Counterfactual

I estimate consumer surplus, insurer profits, and government expenditure under the current benchmark setting, which serve as the baseline measures. Table 7 shows that the baseline consumer surplus is $3.488 billion, insurer profits are $10.177 billion, and total estimated government expenditure is $672.8 billion. The actual government expenditure in 2019 was $799 billion. My estimate is lower because I exclude ESRD beneficiaries, dual-eligible individuals, and certain plans with different payment structures, as described in Section 3. Of the total estimated government expenditure, $206 billion is spent on MA, and $466.8 billion is spent on TM.

Figure 1 presents the results of the first counterfactual, which shuts down the MA market. The results show that if all MA enrollees were instead covered by TM, the cost to CMS would be $185.2 billion—representing a decline of $20.8 billion from the current level of government expenditure of $206 billion. As described in Section 8, the savings from eliminating overpayments to MA plans due to advantageous selection exceed the loss in rebate savings. The MA cost that CMS would expect in the absence of advantageous selection equals the TM cost of covering the same beneficiaries minus the rebate savings. This expected MA cost, shown in Figure 1, amounts to $181.4 billion. Thus, the net increase in government expenditure attributable to the presence of MA is $24.6 billion.

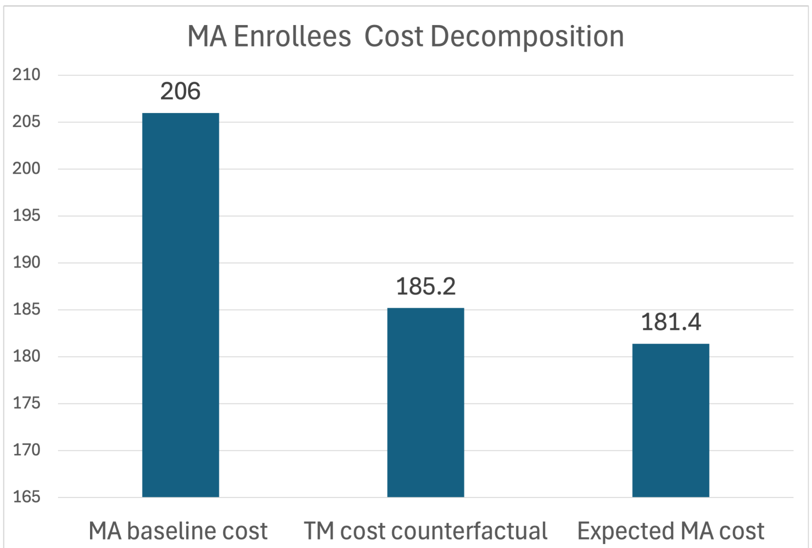

Figure 1: MA shutdown counterfactual

I run two additional counterfactual analyses in which I adjust the benchmark rate and then perform a decomposition of government expenditure, as above, separately for switchers and stayers. *Switchers* are defined as individuals who move from MA to TM, while *stayers* are those who remain in MA. The mathematical details of this decomposition are provided in Appendix 12.5.

To calculate the counterfactual benchmark, I compute the average predicted TM cost for each MA plan given the demographic composition of its enrollees. I refer to this as the *selection-corrected benchmark*. I find that the selection-corrected benchmark lowers the benchmark for most plans but raises it for some. In the second counterfactual (CF2 in Table 7), I set the benchmark rate to the minimum of the current benchmark and the selection-corrected benchmark; I refer to this as the *capped selection-corrected benchmark*. In the third counterfactual (CF3 in Table 7), I set the benchmark equal to the selection-corrected benchmark itself and report the resulting welfare changes.

Under the capped selection-corrected benchmark, approximately 4.15 million individuals switch from MA to TM, resulting in a reduction in consumer surplus of \$3.072 billion. When the benchmark falls, firms' costs increase, equilibrium prices rise, and profits decline by \$1.99 billion. Government expenditure decreases by \$32.1 billion. Because this counterfactual only lowers the status quo benchmark, individuals switch from MA to TM but not in the opposite direction. Consequently, government expenditure in TM increases due to the additional cost incurred for the new TM enrollees (the switchers). Government expenditure in the MA market decreases both because of the exit of switchers (the extensive margin) and because of the reduction in expenditure on stayers resulting from the lower benchmark (the intensive margin). Below, I decompose the total change in government expenditure into the components attributable to switchers and stayers.

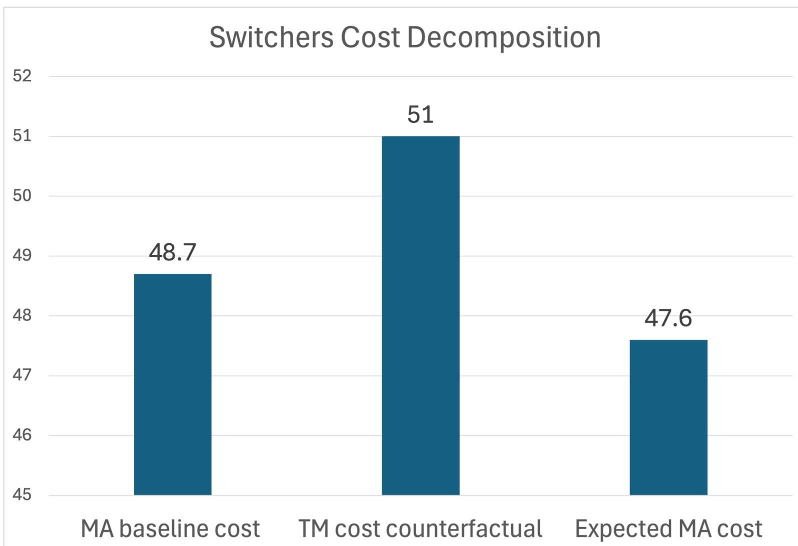

Figure 2: Cost decomposition of switchers

Figure 2 shows that the switchers cost TM \$51.0 billion, which corresponds to the increase in government expenditure on TM reported in Table 7. If these switchers had remained in MA under the baseline benchmark, it would have cost CMS \$48.7 billion. This indicates that there is less advantageous selection among the marginal enrollees, as the composition of switchers consists of relatively *unhealthy* individuals. For these switchers, MA still generates gross savings to CMS of about \$2.3 billion. However, the expected MA cost after accounting for rebate savings is \$47.6 billion, implying that selection results in a net loss of \$1.1 billion.

Figure 3 presents the cost decomposition for individuals who remained in MA even after the reduction in the benchmark. This group consists of approximately 10.3 million individuals. If these individuals were to switch to TM, government expenditure under TM would fall to \$136.08 billion. Covering these individuals under MA at the baseline benchmark rate

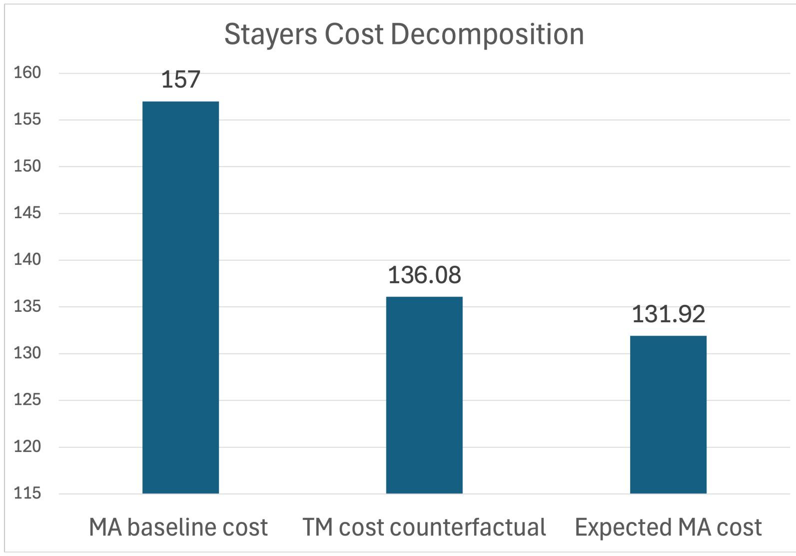

Figure 3: Cost decomposition of stayers

would cost CMS \$157 billion. Because the stayers constitute a healthier pool, it is more cost-effective for CMS to cover them under fee-for-service given the current benchmark setting. The expected MA cost after accounting for rebate savings is \$131.92 billion, implying that selection results in a net loss of \$25.08 billion attributable to MA.

Table 7: Baseline and Counterfactual Results

	Baseline	CF 2: Capped	Change (Capped)	CF 3: Uncapped	Change (Uncapped)
Consumer Surplus	3.488	0.416	–3.072	3.035	–0.453
Profits	10.177	8.184	–1.99	10.045	–0.132
Government Expenditure	672.8	640.7	–32.1	655.83	–16.9
Decomposition					
Government Expenditure – MA	206.0	122.9	–83.1	232.47	+26.47
Government Expenditure – TM	466.8	517.9	+51.1	423.36	–43.44

All values are in billions of dollars. Numbers are scaled upto the national level. Positive sign indicates an increase in the counterfactual values as compared to the baseline values; negative indicates a decrease in the counterfactual values as compared to the baseline values.

In the next counterfactual, I set the benchmark equal to the predicted MA cost of MA enrollees, without capping. As mentioned earlier, the benchmark increases for most plans and decreases for some, resulting in switching of enrollees both from MA to TM and from TM to MA. As shown in the table above, the selection-adjusted benchmark leads to budgetary savings of \$16.9 billion at a cost of \$0.585 billion in reduced consumer welfare and profits. Furthermore, the calculated rebate savings amount to \$7.56 billion, which reflects the intended program outcome of the MA design.

11 Conclusion

This paper develops an integrated model of demand, supply, and government expenditure in the Medicare Advantage (MA) market to evaluate the welfare effects of benchmark policy. By combining administrative risk scores with survey-based health measures, I provide evidence of selection into MA that persists even after conditioning on risk score. This selection increases government spending relative to Traditional Medicare (TM) and leads to systematic overpayment to MA plans.

Using this framework, I construct a selection-corrected benchmark that accounts for both risk scores and additional health dimensions. Counterfactual simulations show that benchmark adjustments could generate substantial budgetary savings for CMS at a modest cost to consumers and insurers. Specifically, I find that an selection-corrected benchmark reduces government spending by $10 billion at the cost of only $0.27 billion in consumer surplus and producer profits. These results underscore the importance of incorporating richer measures of beneficiary health into MA payment policy. This paper speaks directly to the policy discussion of overpayment to MA a tractable approach to evaluating reforms that can both contain federal spending and improve the efficiency of resource allocation in Medicare.

References

Naoki Aizawa and You Suk Kim. Advertising and risk selection in health insurance markets. *American Economic Review*, 108(3):828–867, 2018.

Steven T Berry, James A Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium: Part i and ii, 1993.

Jason Brown, Mark Duggan, Ilyana Kuziemko, and William Woolston. How does risk selection respond to risk adjustment? new evidence from the medicare advantage program. *American Economic Review*, 104(10):3335–3364, 2014.

Cory Capps, David Dranove, and Mark Satterthwaite. Competition and market power in option demand markets. *RAND Journal of Economics*, pages 737–763, 2003.

Michael E Chernew, Keaton Miller, Amil Petrin, and Robert J Town. Reducing medicare advantage benchmarks will decrease plan generosity, but those effects will likely be modest: Study examines the possible effects of reducing medicare advantage benchmark payments. *Health Affairs*, 42(4):479–487, 2023.

Congressional Budget Office. Options for reducing the deficit: 2023 to 2032. https://www.cbo.gov/budget-options/58626, 2022. Accessed October 13, 2025.

Vilsa Curto, Liran Einav, Amy Finkelstein, Jonathan Levin, and Jay Bhattacharya. Health care spending and utilization in public and private medicare. *American Economic Journal: Applied Economics*, 11(2):302–332, 2019.

Vilsa Curto, Liran Einav, Jonathan Levin, and Jay Bhattacharya. Can health insurance competition work? evidence from medicare advantage. *Journal of Political Economy*, 129 (2):570–606, 2021.

Francesco Decarolis, Maria Polyakova, and Stephen P Ryan. Subsidy design in privately provided social insurance: Lessons from medicare part d. *Journal of Political Economy*, 128(5):1712–1752, 2020.

Mark Duggan, Amanda Starc, and Boris Vabson. Who benefits when the government pays more? pass-through in the medicare advantage program. *Journal of Public Economics*, 141:50–67, 2016.

Michael Geruso and Timothy Layton. Upcoding or selection? evidence from medicare on squishy risk adjustment. *NBER Working Paper*, 21222, 2015.

Keaton Miller, Amil Petrin, Robert Town, and Michael Chernew. The optimal geographic distribution of managed competition subsidies. 2023.

Daria Pelech, Zirui Song, Joseph Newhouse, Andrew Ryan, Karen Stockley, Jacob Wallace, et al. Pricing and pass-through in response to subsidies and competition: evidence from medicare advantage before and after the affordable care act, 2019.

Zirui Song, Mary Beth Landrum, and Michael E Chernew. Competitive bidding in medicare advantage: Effect of benchmark changes on plan bids. *Journal of health economics*, 32(6): 1301–1312, 2013.

Karen Stockley, Thomas McGuire, Christopher Afendulis, and Michael E Chernew. Premium transparency in the medicare advantage market: Implications for premiums, benefits, and efficiency. Technical report, National Bureau of Economic Research, 2014.

12 Appendix

12.1 Hospital Data Creation

I used the following steps were to create the data set that I used to estimate the network utility. The base data set for the hospital choice model is the inpatient encounter claims data. This data set contains information of hospital visits by beneficiaries for different diagnosis. There are 180,383 beneficiaries in the inpatient data set accounting for 470,418 claims. There are a total of 8,706 principal diagnoses codes. I categorize principal diagnoses into Clinical Classifications Software Refined (CCSR) codes – a useful grouping defined by the U.S. Agency for Healthcare Research and Quality (AHRQ) that collapses detailed ICD-10 codes into about 443 clinically meaningful categories. I use the National Plan and Provider Enumeration System (NPPES) data to get the address and the taxonomy codes for each hospital. This data is publicly available in a downloadable file. Other hospital characteristics are taken from the American Hospital Association (AHA) Annual Survey Database which was accessed through the Wharton Research Data Services portal. Individual level charaterictics is derived by merging the data to CMS provided Master Beneficiary Summary Files (MBSF) which contains data on beneficiary age , gender, zip code, income category. I use claims from hospitals that were situated in states where the plan didn't operate. Additionally, I dropped claims if a beneficiary used a hospital more than 100 miles away (standard restriction in empirical hospital choice models) or if the patient-level zip code is missing/invalid or hospital location is missing.

Table 8: Patient Characteristics

Variable	**Mean**
No. of Hospitalizations	360,475
Age	73
Male	45%
White	74.93%
Emergency Department	68%
Principal Diagnosis	
Cardiovascular diseases	25.39%
Endocrine/Nutritional/Metabolic	10.68%
Respiratory diseases	9.51%
Digestive	8.91%
Injury/Poisoning	6.83%
Infectious & parasitic	6.75%
Musculoskeletal	5.90%
Mental/Behavioral	4.93%
Nervous System Disorders	3.59%
Cancer & Tumors	3.55%
All Other Diagnoses	13.97%

12.2 Plan Network Identification

There is no direct data on the network of hospitals in each plan. I back out the network of hospitals using the inpatient data as described. From the inpatient data, I can observe all the hospitals visited by individuals who opted for different plans. I assume that the chosen hospital is in the plan's network for every individual (except if it is more than 100 miles away in which case the observation is dropped as described in the previous section). A plan's network is backed out at an individual level based on location. There are few plans that are very popular and therefore many hospitals show up in the inpatient data set for those plans. There are a total of 3534 unique plans. 68% if beneficiaries are enrolled in 18% of the plans and the rest of the beneficiaries are enrolled in the remaining 82% of the plans. The network of hospitals are created such each individual in the same plan can have a different network of hospital based on his or her location. The hospitals in an individual network are the hospitals that are within 100 miles of the individuals location (zip code centroid). After creating this network, if the number of hospitals in an individuals network is greater than 50, then I keep the closest 50 hospitals to keep the choice set manageable. Finally, I dropped hospital networks from the choice set if the hospital did not have the service for the particular diagnosis (2% of the sample).

12.3 Results-Hospital Choice Model

Table 9: Hospital choice demand estimates

Dependent Variable:	Estimates	Standard Errors
distance	-0.1531***	(0.0040)
distance squared	0.0012***	(3.31×10^{-5})
distance X age	-0.0003***	(2.41×10^{-5})
distance X $1\{male\}$	0.0021***	(0.0004)
distance X $1\{white\}$	0.0094***	(0.0011)
distance X $1\{poor\}$	6.98×10^{-5}	(0.0006)
distance X $1\{emergency\}$	-0.0112***	(0.0010)
teaching X cancer diagnosis	0.1505***	(0.0386)
teaching X cardio diagnosis	0.0137	(0.0189)
teaching X endocrine diagnosis	-0.0491***	(0.0160)
teaching X injury diagnosis	0.0107	(0.0185)
teaching X mental diagnosis	-0.9278***	(0.0649)
teaching X musculoskeletal diagnosis	-0.1907***	(0.0309)
teaching X nervous diagnosis	-0.1243***	(0.0221)
teaching X other diagnosis	-0.0433***	(0.0135)
teaching X respiratory	-0.0467***	(0.0152)
teaching X infectious	-0.0015	(0.0172)
psychiatric hospital X mental diagnosis	0.3398***	(0.0638)
heart hospital X cardio diagnosis	-0.0020	(0.0169)
orhtopedic hospital X musculoskeletal diagnosis	0.1492***	(0.0321)
Fixed-effects		
hospital	Yes	

Clustered (hospital) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 7 shows the estimates of the hospital choice model. Because all covariates are observed, I estimate the model by maximum likelihood. The above table show results. Consistent with previous papers' estimates, patients dislike traveling to more distant hospitals, with each extra mile of distance reducing a hospital's share by 14.2% on average. There is substantial heterogeneity in distance sensitivity, with higher dis-utility from distance for older, female, non-white and emergency patients. The other coefficient show how patients with different diagnosis category prefer different hospital characteristics like teaching hospital status.

12.4 Supply Side Estimation

As shown in the paper the profit function for a firm is given by:

$$\Pi_{fc} = N_m \sum_{j=1}^{L_{fc}} \int_{h_i} [p_{fjm} + r_i.B_{fjm} - \underbrace{\frac{1-\lambda}{\lambda} Cinc_{fjm} - c_{fjm}(\mathbf{z_i})}_{MC_j(oop\ max)}].s_{fjm}(\mathbf{z_i})dF_c(\mathbf{z_i})$$

Let both the incremental cost and the plan's cost constitute the total marginal cost of the plan. Substituting the functional form:

$$MC_j(oop\ max) = \gamma_{TM}.\overline{TM}_m(\mathbf{z_i}) - \frac{\gamma_x}{\lambda}.cost\_sharing_{fjm} + \delta_f + \eta_{fjm}$$

Out of pocket maximum enters linearly in the marginal cost, so $MC'(oop\ max)$ is $\frac{\gamma_x}{\lambda}$. For ease of exposition, consider each firm has only one plan. Then the first order condition for the plan is (ignoring the market subscript):

$$\frac{\partial \Pi_j}{\partial oop\ max_j} = [p_j + r_i.B_j - MC_j(oop\, max_j)] \frac{\partial s_j(\cdot)}{\partial prices_j} N_m - s_j(\cdot)MC'(prices_j)N_m$$

Equating to zero and solving,

$$MC = p_j + r_i.B_j - \frac{oop\ max_j}{\epsilon_d}.MC'(oop\ max)$$

Substituting for marginal cost , taking out *oop max* from the cost sharing and solving gives:

$$\gamma_{TM}.\overline{TM}_m(\mathbf{z_i}) - \frac{\gamma_x}{\lambda}.(copay_j + deductible) + \delta_f - p_j - r_i.B_j + \eta_j = [\frac{oop\ max_j}{\epsilon_d} + oop\ max_j]\frac{\gamma_p}{\lambda}$$

So the marginal cost estimation in pyblp (Python programming used for demand estimation), is estimated upto $\frac{\gamma_p}{\lambda}$ scale. I use the delta method rescale the results that are reported.

12.5 Decomposition of Government Expenditure on MA across stayers and switchers

In the counterfactual analysis, I decompose government expenditure into expenditure on *stayers* and *switchers*.

At the baseline benchmark, each individual i in market m chooses plan j with probability P_{ijm}. At the counterfactual benchmark, the model computes this probability as P'_{ijm}.
A switcher is defined probabilistically as an individual who switches from the inside good (MA) to the outside good (TM). For a given market m (suppressing the market subscript for convenience), the probability that individual i switches is given by

$$P(i \text{ switches}) = \max\left(\sum_j P_{ij} - \sum_j P'_{ij},\, 0\right).$$

Since the analysis is conducted for the capped selection-corrected benchmark, it follows that $\sum_j P_{ij} - \sum_j P'_{ij} \geq 0$. This switching probability is defined at the individual level.

The plan-level contribution to the switchers is defined as

$$\Delta P_{ij} = P_{ij} - P'_{ij}.$$

Summing these plan-level contributions across all plans returns the individual-level probability:

$$\sum_j (P_{ij} - P'_{ij}) = \sum_j P_{ij} - \sum_j P'_{ij}.$$

The government expenditure (baseline) on the MA market is defined as

$$GE_{\text{MA}} = \sum_j \left[R_j \cdot B_j - \frac{(1-\lambda)}{\lambda}\,\text{Cinc}(CS_j)\right] s_j \times N_m,$$

where R_j denotes the risk composition of individuals who choose plan j, defined as

$$R_j = \frac{\sum_i P_{ij} R_i}{\sum_i P_{ij}},$$

and R_i is the individual risk score. The plan share is given by

$$s_j = \frac{\sum_i P_{ij}}{I_m},$$

where I_m is the number of simulated individuals in market m.

In the counterfactual, when the benchmark changes, the risk composition of each plan changes because the choice probabilities P_{ij} change to P'_{ij}. The share of plan j also changes accordingly. The government expenditure on MA for the *stayers* (i.e., those who remain in

MA) is given by

$$GE_{\text{MA,stayers}} = \sum_j \left[R'_j \cdot B_j - \frac{(1-\lambda)}{\lambda} \text{ Cinc}(CS_j) \right] s'_j \times N_m,$$

where

$$R'_j = \frac{\sum_i P'_{ij} R_i}{\sum_i P'_{ij}}, \qquad s'_j = \frac{\sum_i P'_{ij}}{I_m}.$$

I also compute the government expenditure on the *switchers* to TM, had they stayed in MA at the baseline benchmark rate, by defining the plan-level probability of switching as $(P_{ij} - P'_{ij})$. Formally,

$$GE_{\text{MA,switchers}} = \sum_j \left[R_j^{\text{sw}} \cdot B_j - \frac{(1-\lambda)}{\lambda} \text{ Cinc}(CS_j) \right] s_j^{\text{sw}} \times N_m,$$

where

$$R_j^{\text{sw}} = \frac{\sum_i (P_{ij} - P'_{ij}) R_i}{\sum_i (P_{ij} - P'_{ij})}, \qquad s_j^{\text{sw}} = \frac{\sum_i (P_{ij} - P'_{ij})}{I_m}.$$

By construction,

$$GE_{\text{MA,baseline}} = GE_{\text{MA,switchers}} + GE_{\text{MA,stayers}}.$$