

Entry and Competition in Insurance Markets: Evidence from Medicare Advantage*

Matthew V. Zahn[†]

January 1, 2025

[Click here for latest version](#)

Abstract

Governments frequently turn to private markets to deliver public benefits. This structure can lower the government's costs if it designs a payment system that attracts competitive firms with cost controls the government lacks. In this paper, I analyze the implications of this system for the Medicare Advantage program. I use administrative data to develop and estimate a model of firm entry and product offering decisions that captures how firms endogenously respond to government policies as well as consumer sorting and utilization of health insurance. I then use the model to simulate other payment policies in Massachusetts. I find that under the current design, the government overpays firms for their participation and the enrollment they generate. Under a policy that lowers firm payments and transfers a portion of this money to consumers, the government can reduce spending by 1% (\$10 billion when applied nationally). This policy leads to similar firm participation and enrollment in Medicare Advantage, while reversing the historic positive selection into the program and the government's higher per-beneficiary payments for those in Medicare Advantage relative to Traditional Medicare.

Keywords: competition, endogenous plan menus, entry, Medicare Advantage

JEL Codes: D82, G22, H75, I11, I13, L11, L13

*First version: October 1, 2023. I am grateful to Elena Krasnokutskaya, Nicholas Papageorge, and Jonathan Elliott for their mentorship, guidance, and advice. I thank Angie Acquatella, Leila Agha, Sam Arenberg, Chris Carroll, Michael Chernew, Jeffrey Clemens, David Cutler, Vilsa Curto, Michael Darden, Liran Einav, Randall Ellis, Keith Ericson, Amy Finkelstein, Stelios Fourakis, Barton Hamilton, Silin Huang, Adam Isen, Jihye Jeon, Ed Kong, Tim Layton, Robin Lee, Lixiong Li, Luca Maini, Robert Moffitt, Ariel Pakes, Dan Polsky, Maria Polyakova, Anna Russo, Stephen Ryan, Marc Rysman, Eric Schulman, Mark Shepard, Maggie Shi, Nicholas Tillipman, Justin Trogdon, Benjamin Vatter, Myles Wagner, Sara Walters, Zenon Zabinski and seminar participants at Johns Hopkins University, Harvard Medical School, Boston University, the Annual Health Econometrics Workshop, International Industrial Organization Conference, DC IO Day, and the NBER Aging and Health Fellowship for insightful comments and discussions. This research was supported by the National Institute on Aging, Grant Number T32-AG000186 and the Hopkins Business of Health Initiative. Mark Drozd provided excellent research assistance. Any errors are my own.

[†]Hopkins Business of Health Initiative, Johns Hopkins University; matthew.zahn@jhu.edu.

I Introduction

As of 2021, the U.S. government spends nearly \$830 billion—10% of all spending—each year on healthcare for seniors in the Medicare program (Cubanski and Neuman, 2023). The majority of beneficiaries receive these benefits through Traditional Medicare (TM), the public insurance option. The remainder receive coverage through Medicare Advantage (MA), which are private insurance plans that are subsidized by the government. There are at least three policy rationales for subsidizing a private market for Medicare benefits. First, private firms have developed expertise in limiting moral hazard healthcare utilization, which allows these companies to deliver benefits at lower cost. Second, competition lowers premiums and offers products with extra services not covered by TM (e.g., vision, dental, hearing, etc.) or financial coverage to attract enrollees. Third, competition also gives firms an additional incentive to lower their costs, which generates additional savings for the government. This structure for using private markets to deliver public goods appears in other settings, including education and housing (see e.g., Baum-Snow and Marion, 2009; Hoxby, 2000; Neilson, 2021; Poterba, 1996).

Promoting entry and robust participation in insurance markets faces several challenges. Chief among these are adverse selection—the tendency for sicker people to prefer more generous insurance plans—and moral hazard—the propensity to consume additional healthcare because it is cheaper—are the most salient. Concerns about selection may lead to firms offering plans with less generous coverage in markets with sicker patients or failing to enter these markets altogether—a behavior typically referred to as “cream skimming.” Firms must also contend with the existence of the public option. Traditional Medicare offers baseline coverage—which has gaps and higher out-of-pocket costs—at a relatively low premium. The private market must be competitive on both of these dimensions—coverage and premium—to attract enrollment. These forces create challenges for the design of government policies (i.e., subsidies) to support the private market. The policy must not only incentivize partici-

pation, it also needs to address selection and consumer price sensitivities. This is a difficult problem to solve because evaluating counterfactual policies requires a model that captures the complex interplay between government policy, firm entry, and product offering decisions, as well as consumer plan and healthcare utilization choices. Prior work has captured some of these features in isolation, but we lack a unified framework that incorporates each of these components.

In this paper, I develop a model of firm entry and product offering decisions in health insurance markets. The framework captures how firms endogenously modify their participation decisions in response to changes in policy, competitive conditions, and consumer demand. I estimate the model using administrative data from the Medicare program. Since these resources includes the new Medicare Advantage encounter data, I am able to measure healthcare utilization in these private plans.¹ I then use the model to run counterfactual simulations to evaluate the impacts of policies to promote participation in the Medicare Advantage market. There are three key findings. First, subsidies are necessary to support the existence of this private health insurance market. If no subsidy is provided, firms are unable to effectively compete against the public option to attract enrollees. Second, under its current form Medicare Advantage is not fully taking advantage of the private market's expertise since healthier individuals tend to enroll in the private market more than sicker ones. Third, current policies used in Medicare Advantage overpay for the outcomes they deliver. The overpayment is driven by subsidization levels that are too high relative to an insurer's fixed costs of participating in the market and the cost of providing coverage to those that enroll in Medicare Advantage plans. An alternative policy that lowers these payments along with a subsidy targeted at low-income seniors—who tend to have greater health needs—can deliver comparable outcomes to current policies, while lowering government costs by 1% (approximately \$125 per enrollee). These values were calculated for a single state but when projected out nationally imply savings of \$10 billion for the government. This policy also reverses the

¹To my knowledge, this is the first paper in economics to leverage these data.

historic positive selection into Medicare Advantage and leads to lower per-beneficiary costs for MA enrollees than those that opt into TM.

Medicare Advantage is an attractive setting to study the supply side of insurance markets. Private insurers administer and operate insurance plans that receive a subsidy to account for the health status of each beneficiary they enroll. Plans that report costs below benchmarks for the government's costs of providing TM also receive additional payments to fund extra services or better cost sharing benefits for their enrollees. In two descriptive analyses, I show how this subsidy policy creates variation in plan choice set generosity and how this variation separately identifies adverse selection and moral hazard. This feature enables me to design a rich model of the supply side of this insurance market that captures how healthcare utilization driven by adverse selection and moral hazard impacts firm decisions.

The model has two stages. In the first stage, firms choose which markets to enter and which insurance products to offer. These choices are made to optimize expected net profits, taking into account the actions of their rivals, subsidies from the government, expected consumer demand, the healthcare utilization of their enrollees, and the fixed costs of entry. Demand and healthcare utilization are then realized in the second stage of the model. Access to administrative data allows me to capture rich levels of observable and unobservable heterogeneity in the estimation of consumer preferences and healthcare utilization. These components of the model capture how consumer selection across plans not only responds to, but also influences, the entry and product offering decisions of firms in equilibrium. I estimate firm fixed costs using moment inequalities derived from revealed preference assumptions to rationalize observed entry and product offerings. As a result, my model can characterize equilibria resulting from different subsidy policies.

Model estimates indicate consumers are price sensitive and value their expected utility from healthcare consumption. Consistent with the incentives of private health plans to control costs, I find that MA plans have significant utilization costs to reduce the amount of healthcare their beneficiaries consume relative to TM. These utilization costs are also

effective at limiting the amount of moral hazard utilization their enrollees consume. Individuals in this market display a modest risk aversion, consistent with the high level of financial generosity of Medicare Advantage plans in terms of cost sharing (i.e., coinsurance and out-of-pocket maximums). My estimate for the identified set of firm fixed costs captures the costs of establishing a provider network, the efficiency of entering markets with existing networks, and per-plan regulatory costs. Finally, I perform a series of exercises that demonstrate how allowing for selection and endogenous participation impacts predictions from my model. My results illustrate how models without these features may overstate the effects of counterfactual policies on consumer welfare and government spending.

I then use the model to weigh the tradeoffs of promoting firm participation in Medicare Advantage markets. To preserve tractability, I simulate outcomes for a single state (Massachusetts) and restrict firm strategies to enter groups of counties and offer products at the network type-financial generosity level (i.e., HMO or PPO and low or high generosity). The current policy in this market is to subsidize firms for each beneficiary they enroll. I start by assessing whether the government needs to subsidize a private market for Medicare benefits. Absent subsidies, premium competition with the public option—Traditional Medicare—leads the private market to unravel. Medicare Advantage plans offer a more generous product relative to TM and absent a subsidy, these firms cannot profitably attract enrollment because of the premium competition with TM. In other words, subsidies are necessary to sustain the existence of the private market. While eliminating subsidies results in the lowest amount of government spending, it ignores many of the other benefits MA has for consumers in the form of insurance products that are more financially generous than TM.

Having established that subsidies are necessary to sustain the private market, I consider whether alternative systems can deliver better outcomes for the program. In particular, I focus on whether alternative systems can improve how consumers sort between TM and MA as well as deliver cost savings to the government. In the first scenario, I eliminate the supply subsidy and transfer the average observed firm subsidy to consumers. This policy

expands the size of the private market by 52% on average as healthier beneficiaries leave TM. Mechanically, the demand subsidy allows plans with low costs to effectively have a negative premium, which is not permissible under the supply subsidy. Government spending increases by almost 3% (over \$219 million) due to the expansion of the private market. Thus, subsidies play a significant role in terms of influencing who enrolls in Medicare Advantage. Absent a targeted policy that induces sicker people to enroll in Medicare Advantage before the healthier ones, the government is unlikely to realize any of the potential savings promised by the program.

In the second scenario, I simulate a targeted policy that addresses deficiencies in the observed policy and untargeted simulated policy. The targeted policy has two components: a reduction in supply side subsidy benchmarks and a means tested demand subsidy. This targeted approach delivers market outcomes—firm entry and enrollment—similar to the observed policy. Government spending under this policy falls by close to \$100 million (1%) on average, which is approximately \$125 per enrollee. If this figure is applied to total government Medicare spending it amounts to roughly \$10 billion in government cost savings. Finally, this policy reverses the historic positive selection into MA and makes the program more cost effective than TM on a risk adjusted per-beneficiary basis. This exercise shows that there is room to improve how this private insurance market is regulated and may be able to deliver on its promised cost savings.

This paper contributes to our understanding of promoting choice in health insurance markets by rigorously capturing the role of the supply side of the market. Prior work in this space has weighed the value of offering choice based on an analysis of consumer demand. Prominent examples are Marone and Sabety (2022) and Ho and Lee (2022). Both extend the framework of Einav et al. (2013), which allows consumers to adjust their health spending based on their insurance coverage (i.e., moral hazard) to understand when consumer choice over insurance products with different levels of coverage is desirable. Both find there are limited gains to offering choice over different levels of financial coverage if a sufficient

baseline level is offered.² My contribution extends these analyses by adding a demand model of comparable richness to a complete model of health plan supply—one that not only captures decisions about entry but also product variety. These features allow my framework to determine what entry and product offering decisions will arise endogenously under different policy regimes, taking account of the demand response. As a result, I can expand our understanding of the tradeoffs associated with incentivizing firm participation in competitive insurance markets.

My analysis also extends prior work on endogenous participation in insurance markets. Kong et al. (2022) and Geddes (2022) study how policies to mitigate adverse selection can induce greater insurer entry into markets and allow enhanced competition to improve consumer welfare. Miller et al. (2021) focus on how firms endogenously alter their plan characteristics in response to subsidization policies, while holding participation fixed. My model builds on this work by capturing both margins—firm participation and plan offering decisions are endogenous within my framework. These features are necessary to fully quantify how counterfactual policies may alter firm decisions and their impacts on consumers. For example, while a model that allows firms to endogenously reposition their product offerings to changes in policy, they rule out equilibria where it is optimal for the firm to exit the market altogether. This action may carry different implications for consumer welfare than the change in product characteristics induced by the policy. A contribution of my analysis is to simulate a model that captures both of these margins for supply to respond.

My work also contributes to the literature studying the equilibrium effects of adverse selection and the design of health insurance markets. Examples include Einav et al. (2019), which develops a framework to weigh the tradeoffs between demand subsidies and risk adjustment in a joint framework. Tebaldi (2022) assesses the ability of targeted subsidies to alter selection patterns to improve market outcomes for consumers, and Polyakova and Ryan

²Ho and Lee (2022) note that the gains from choice can improve if choice over financial and non-financial characteristics are offered. Wagner (2022) explores the conditions under which it is optimal to offer plan menus with plans differentiated in terms of their financial coverage and network types.

(2020) document how imperfect competition can distort the efficiency of targeted demand subsidies. Closely related to my analysis, Curto et al. (2021) studies the current regulatory framework used in MA—sometimes referred to as “managed competition”—as a model for insurance markets.³ I extend these analyses by studying how managed competition in MA impacts firm participation and product offering decisions. As a result, my model can answer whether managed competition generates sufficient entry or product offerings that are valuable to consumers and whether alternative regulatory schemes perform better at achieving these outcomes.

Finally, this paper relates to prior studies of product repositioning and firm entry. A common challenge for papers in these literatures is handling multiple equilibria. While Berry (1992) opted to model an outcome common to all equilibria, recent work has looked to partial identification methods to estimate the set of parameters consistent with multiple model equilibria (e.g., Ciliberto and Tamer, 2009; Eizenberg, 2014; Fan and Yang, 2020, 2022; Wollmann, 2018; and Ciliberto et al., 2021). My own analysis relies on partial identification based on moment inequalities generated by revealed preference to account for multiple equilibria in the spirit of Pakes et al. (2015). Methodologically, I combine models of entry and product repositioning by capturing how firms choose to offer different types of products in different geographic markets. Moreover, my findings highlight the importance of accounting for endogenous participation when performing counterfactual analyses that alter firm entry incentives.

The paper proceeds as follows. In Section II, I present the empirical setting with a description of the Medicare Advantage program and the data. Section III presents a descriptive analysis that highlights how variation in government subsidy policies induces

³There is an extensive literature on Medicare Advantage in economics. Examples include how insurers invest and compete over non-premium characteristics captured by quality measures (Vatter, 2022); overpayments associated with the risk adjustment system (Geruso and Layton, 2020); whether risk adjustment has attenuated the incidence of risk selection between MA and TM (Brown et al., 2014 and Newhouse et al., 2015); the pass-through of plan subsidies to consumers (Cabral et al., 2018 and Duggan et al., 2016); and the impact of plan quality on mortality (Abaluck et al., 2021).

variation in plan entry and financial generosity capable of separately identifying adverse selection and moral hazard. The model is presented in Section [IV](#). I then discuss estimation and identification in Section [V](#) followed by results and model fit in Section [VI](#). In Section [VII](#), I simulate how alternative policies impact firm entry and product offering decisions as well as their associated welfare benefits and costs. Section [VIII](#) concludes.

II Empirical Setting

This section provides an overview of Medicare Advantage’s institutional background and the data I use in my analysis. Each year, beneficiaries eligible for Medicare must choose between Traditional Medicare and Medicare Advantage to receive healthcare coverage. Traditional Medicare, composed of Medicare Part A and Part B, covers inpatient and outpatient services (e.g., hospital visits, doctor appointments, lab tests, etc.). Since Traditional Medicare is provided by the government, most healthcare providers accept it as payment under a fee-for-service system. Medicare Advantage (originally called Medicare Part C) are health insurance plans administered by private firms and subsidized by the government. The plans are required to cover the same services as Traditional Medicare at a minimum, but typically include additional services not covered by Traditional Medicare like vision, dental, and prescription drugs.⁴ Since Medicare Advantage is private insurance, enrollees must navigate a network of acceptable providers. Traditional Medicare beneficiaries do not have to navigate these restrictions. While both Traditional Medicare and Medicare Advantage have out-of-pocket (OOP) costs for enrollees (e.g., premiums, deductibles, copays, etc.), they tend to be lower for Medicare Advantage plans.⁵ Appendix Figure [E.1](#) provides detailed breakdown of the Medicare program and the coverage options available to seniors.

⁴Traditional Medicare enrollees may supplement their coverage with a Medicare Part D plan, which covers the costs of prescription drugs.

⁵Traditional Medicare enrollees may purchase Medigap policies to cover some of these costs.

II.A Medicare Advantage

Medicare Advantage dates back to the early 1980s. The goal of the program was to use private firms to deliver Medicare services to tap into two potential benefits. The first benefit stems from the expertise of private firms. Health insurance companies have developed strategies and mechanisms that can reduce the amount of healthcare enrollees consume as well as increase the services offered to consumers. The government is unable to accomplish these goals under Traditional Medicare in its current form and could realize significant cost savings by relying on these private firms to deliver Medicare benefits. The second benefit relates to competitive markets. Competition creates incentives for these firms to further lower their costs, which generates additional savings for the government. These forces also lower premiums, which allow consumers to more readily access products with additional services.

The initial design of the program was unable to deliver these benefits. The primary issue stemmed from selective firm participation. Historically, the Centers for Medicare and Medicaid Services (CMS) set payment rates for MA plans. Insurers tended to participate in years when CMS offered higher payments or in specific geographies where the payments were greater or had healthier patients (“cream-skimming”). These behaviors hampered the ability of Medicare Advantage to deliver its potential benefits to the government. These circumstances motivated a series of reforms to the program that created the regulatory structure currently in place.

To address concerns about firm participation, Congress authorized a new system for determining subsidies paid to Medicare Advantage plans.⁶ The system is organized around benchmarks that reflect the government’s costs of providing TM benefits to a typical beneficiary. CMS sets these rates annually at the county-level and they are observed by insurers. CMS considers each county a distinct market and limits enrollees to choose among plans offered in their county of residence. Insurers submit estimates for their costs of providing

⁶While CMS uses the term “bidding system” and “bid” when discussing this process, they do not resemble auctions and I avoid using these terms when possible to prevent confusion.

Medicare coverage to that population for each plan they offer.⁷ Let b_j and B_j denote the requested subsidy and government cost benchmark for plan j , respectively. The government will pay plan j $\min\{b_j, B_j\}$ for each individual the plan enrolls. If $b_j < B_j$, the plan also receives a “rebate” payment the plan must use to fund additional benefits. Alternately if $b_j > B_j$, then the difference between the subsidy and the benchmark is passed along to consumers as part of the plan’s premium. Medicare Advantage plans can also charge premiums if they offer additional benefits relative to TM.

Risk adjustments were also introduced by Congress to address concerns about Medicare Advantage targeting healthier populations. The purpose of risk adjustment is to scale the subsidies paid to plans based on the health of each enrolled beneficiary. These transfers to plans are adjusted linearly based on a beneficiary’s risk score which is calculated by CMS (i.e., the subsidy for a beneficiary with a risk score 1.1 is 10% larger). Given this adjustment structure, enrollment in MA plans is typically weighted by beneficiary risk scores. The base risk score is the output of a CMS model that takes beneficiary demographics (i.e., age, gender, Medicare eligibility, and Medicaid status) and specific types of diagnoses from the prior year.⁸ The base scores are then normalized by a factor based on TM costs such that the typical TM beneficiary has a risk score equal to one. Finally, risk scores for MA beneficiaries are scaled down to account for more intense coding of diagnoses for MA beneficiaries.⁹

⁷Insurers generally submit a single bid for each offered plan. Insurers are allowed to breakup a plan’s footprint into multiple segments and submit separate bids for each segment. In practice the use of multiple segments is rare and I abstract from them in this paper.

⁸The diagnoses that are included in the risk score calculation come from inpatient and outpatient hospital stays, physicians, and clinically trained non-physicians (e.g., psychologist, podiatrist, etc.). New beneficiaries that do not have recorded diagnoses from the prior year use a different CMS model to calculate their base risk score.

⁹This pattern is referred to as “upcoding” and is pervasive among MA plans. This behavior costs the government more than \$650 per-enrollee each year and is too large to be offset by the current adjustments used by CMS (Geruso and Layton, 2020).

II.B Data

My analysis uses information from 2016–2018 and primarily relies on three types of administrative data from the Medicare program. First, for every beneficiary eligible for Medicare, I observe their demographic information and choice of MA plan or TM. The second are medical claims for beneficiaries that enroll in TM. For a 20% random sample of TM beneficiaries each year, I observe their inpatient, outpatient, and physician claims. I also have access to inpatient discharge records for 100% of the Medicare population. The third are records of encounters between MA beneficiaries and medical providers, which CMS recently made available for research. These files contain information similar to medical claims except for service payments. The MA encounter data cover 100% of inpatient and outpatient records and physician encounters for a cohort of over 12 million beneficiaries, which covers roughly 52% of MA beneficiaries in my analysis sample. These data allow me to construct choice probabilities, risk scores, and county-level demographics for the Medicare population.

I supplement the administrative data with four additional sources. The first are characteristics for every MA plan offered including the plan’s premium, network type, and financial generosity as measured by expected out-of-pocket costs. The second are worksheets that firms complete to receive their subsidies from the government. In particular, these files contain the specific subsidy amount the firm requested for the plan, how the plan’s premium is broken down between the base and supplemental premium, how much supplemental revenue is required to fund extra benefits, and the allocation of rebate payments to cover these benefits. This paper appears among the first in economics to leverage both the MA encounter data alongside plan-level subsidies, which are both essential for my analysis of Medicare Advantage. Third, from DRG InterStudy I observe whether a firm offers other insurance products (i.e., commercial group, commercial individual, Medicaid managed care, etc.) at the county-level. Finally, I obtain information on provider supply and market characteristics from the Health Resources Services Administration, American Hospital Association,

and Census Bureau. Appendix A provides a detailed summary of every data set and its use within this paper.

I restrict my analysis to beneficiaries eligible for Medicare due to age (i.e., non-disabled and non-ESRD) and are enrolled in TM or a MA HMO or Local PPO plan.¹⁰ I also exclude employer sponsored, special needs plans, and Part B only plans. I drop a small number of individuals because they are missing information necessary to calculate risk scores or are enrolled in a MA plan with missing characteristic information. See Appendix A and Appendix Table E.1 for a detailed discussion of the sample criteria. After using the 2016 data to construct risk scores, my full sample for 2017–2018 contains 73,941,784 beneficiary-year observations and 40,141,182 unique beneficiaries. The utilization sample contains 4,424,824 beneficiary-years (2,410,546 beneficiaries). The full sample contains 3,702 plan-year observations for 2,263 unique MA plans.

Table 1 contains summary statistics for the Medicare Advantage markets in my sample. Beneficiaries typically face a monthly premium of \$20 for MA plans, nearly all of which is used to fund supplemental benefits. MA plans are heavily subsidized by the government—the typical subsidy and rebate payments are approximately \$750 and \$66 per-beneficiary-per-month, respectively—consistent with the benchmarks CMS sets for each market. CMS estimates that the average MA beneficiary will have \$140 per-month (\$1,680 annually) in out-of-pocket costs. The average market has seven plans offered by three firms. The majority of these plans are HMOs, which tend to have lower costs, narrower networks, and more cost controls relative to Local PPOs. Roughly three of the plans in a market are considered “high generosity” based on monthly out-of-pocket cost estimates. Despite having several plans, most markets are highly concentrated, which suggests plans may have considerable power in these markets.

¹⁰“ESRD” refers to end-stage renal disease. Citizens in the United States diagnosed with ESRD are eligible for Medicare benefits regardless of their age. HMO stands for health maintenance organization and PPO stands for preferred provider organization.

III Descriptive Analysis

This section uses reduced form methods to highlight data variation in my setting that is critical for the identification of my model of health insurance supply and demand. First, I demonstrate how government policies influence the number of MA plans in a local market and their characteristics. These policies act as a plausibly exogenous source of variation that induces plans to offer different levels of financial generosity. As a result, changes in these policies create variations in the average generosity of the plan choice sets facing consumers. I then demonstrate how this variation in plan choice set generosity allows me to separately identify healthcare utilization driven by private health information from moral hazard—an essential feature to identify the model of healthcare utilization and demand.

III.A How policy influences plan entry and characteristics

I demonstrate how firms respond to changes in their payments following the implementation of the Affordable Care Act (ACA). As discussed in Section II, firms offering Medicare Advantage plans receive two payments from the government. The first is a subsidy for every beneficiary they enroll and the second is a rebate that is paid to plans that request subsidies below the government’s TM cost benchmarks. Rebates must be used to provide more generous benefits to enrollees.

In an effort to control costs, the ACA took steps to reduce payments to Medicare Advantage plans. This law transitioned county cost benchmarks to a new system that aligned them more closely with Traditional Medicare costs.¹¹ Plans face a weighted average of the county-specific benchmarks where they entered when choosing their subsidies. This structure allows consumer sorting and healthcare utilization from other geographies to influence the products that are available in local markets. Thus, variation in these county-level benchmarks across markets can act as a source of plausibly exogenous variation in plan subsidy

¹¹A detailed discussion of the ACA reforms is available in Appendix B.

and rebate payments, which affects entry incentives and the characteristics of the insurance products presented to consumers.

To test whether this is a valid source of policy variation, I empirically assess whether cross-market variation in CMS benchmarks predicts entry and the generosity of insurance plans in a market. For each plan j in county m in year t , I construct the plan’s leave-one-out benchmark $B_{jt \setminus m}$ as:

$$B_{jt \setminus m} = \sum_{k \in A_{jt} \setminus m} w_{jkt} B_{mt} \quad (1)$$

where A_{jt} denotes the set of counties where plan j entered in year t and w_{jkt} are weights based on the number of people plan j enrolled in market k such that $\sum_{k \in A_{jt} \setminus m} w_{jkt} = 1$. The notation $A_{jt} \setminus m$ denotes the set of counties plan excluding market m .

After constructing the leave-one-out benchmarks for each plan, I aggregate benchmarks, entrants, and plan characteristics—weighting by plan enrollment—to the market level. Then I run regressions of the following form:

$$Y_{mt} = \beta_0 + \beta_1 B_{t \setminus m} + \beta_m + \beta_t + \epsilon_{mt} \quad (2)$$

where Y_{mt} is the market-level outcome (i.e., total entrants or average plan characteristic), $B_{t \setminus m}$ is the market average leave-one-out benchmark for the plans active in market m in year t , and β_m and β_t are county and year fixed effects respectively. The sample for these regressions is counties with Medicare Advantage plans in 2017–2018. This time period includes the first year when all counties completed their transition to the ACA payment system. Similar results are obtained when using publicly available Medicare Advantage enrollment data that include additional years.

Table 2 reports the estimated effects of cross-market benchmark variation on firm participation and the financial generosity of Medicare Advantage plans in a market. Participation is measured by the number of firms or plans in a county. My estimates indicate that

the leave-one-out benchmark does not predict the number of firms active in a market, but does have a positive and significant relationship with the number of plans in the market. In other words, higher benchmarks are associated with more plan entry. One way to interpret this pattern is that benchmarks can influence the intensive participation margin yet firms are likely to enter these markets even when benchmarks are low. This interpretation is arguably consistent with the fact that most firms offering MA plans are active in other insurance segments and have already paid the sunk costs of entry.

I considered three measures of the financial generosity of Medicare Advantage plans: supplemental premiums which plans only charge if they provide additional benefits relative to TM, rebate payments used to fund additional benefits, and the rebate dollars specifically allocated toward improved cost sharing. Each of these variables are directly observable in CMS data. For each outcome, the estimated coefficient on the leave-one-out benchmark is significant and has an interpretation consistent with more generous insurance. Higher average benchmarks predict MA plans receive larger rebate payments and allocate these dollars toward providing more generous cost sharing. Higher average benchmarks also predict significantly lower supplemental premiums. This pattern is consistent with plans earning higher rebates, which can offset the costs plans would otherwise charge consumers for offering additional benefits.

Taken together, this analysis highlights how variation in CMS benchmarks can induce plausibly exogenous variation in firm participation and the generosity of the health insurance choice sets presented to consumers in their local market. As the cost benchmarks change each year, firms update their plan offerings, subsidy requests, and collect rebate payments. These rebates are reinvested by the plans to provide additional benefits relative to TM. In [Appendix B](#), I provide further evidence of how firms respond to changes in benchmarks with event studies documenting their responses to the ACA reforms.

III.B Elasticity of healthcare consumption

The previous section illustrated how government policy creates plausibly exogenous variation in the number of plans available to consumers and their financial generosity across markets. This section demonstrates how this variation in choice set generosity can identify the elasticity of healthcare consumption (i.e., moral hazard) based on a similar approach used by Marone and Sabety (2022). The ability to identify and quantify this elasticity is important for motivating the structure of the equilibrium model of health plan demand that can account for selection on health information as well as moral hazard. Both components are necessary to fully capture the policy environment where firms leverage their expertise to control these costs when offering more generous insurance products relative to the public option.

To get a sense for the incidence of selection in Medicare, Figure 1 plots average annual healthcare utilization along two margins. The left panel compares utilization among TM and MA beneficiaries unconditionally and conditional on the six most common groupings of observable characteristics.¹² TM beneficiaries tend to utilize more healthcare than MA beneficiaries unconditionally and conditional on observable characteristics. This pattern could be explained by either unobserved health differences (selection) or steps MA plans take to manage the amount of healthcare their enrollees consume (impacting moral hazard). The right panel compares utilization among MA beneficiaries across plans with different levels of financial generosity. Utilization tends to be greater in MA plans with a high level of financial generosity unconditionally and conditional on observed characteristics. Greater health needs or moral hazard could rationalize the higher utilization in more financially generous MA plans. These patterns highlight the two key empirical challenges for quantifying the elasticity of healthcare consumption—endogenous sorting into plans and endogenous plan generosity.

¹²These groupings summarize a beneficiary’s risk score, age, gender, income, and their county’s Medicare mortality and Medicaid eligibility rates.

To address selection into plans, I estimate a nested logit model of consumer demand for Medicare Advantage plans. All MA plans belong to a common nest $g = 1$ while Traditional Medicare is captured as the outside option $g = 0$ with a utility normalized to zero. The utility individual i receives from MA plan j in county m in year t is denoted as:

$$u_{ijmt} = \delta_{jmt} + \xi_{jmt} + p_{jmt}\alpha_1 + \text{SR}_{jmt}\alpha_2 + \zeta_{ig} + (1 - \sigma)\epsilon_{ijmt} \quad (3)$$

where δ captures measures of the plan’s quality rating, provider network, and time trends, ξ are unobserved plan characteristics, p is the plan premium, SR is the dollar value of extra benefits the plan provides relative to TM, and $\zeta_{ig} + (1 - \sigma)\epsilon_{ijmt}$ are individual-level unobservable determinants of demand that are assumed to follow a Type I extreme-value distribution. Estimation follows Berry (1994), where I use the policy and demographic instruments to address the endogeneity of premiums and supplemental revenue with respect to the unobserved demand shifters ξ . These instruments are discussed in more detail in Section V.C. Parameter estimates are presented in Appendix Table E.2.

Given valid instruments, this model produces choice probabilities that are unbiased estimates for the probability that a typical individual will enroll in a Medicare Advantage plan or Traditional Medicare as a function of plan characteristics (i.e., price, provider network, and additional services). I rely on these predicted probabilities in two ways. First, I use them to construct the probability an individual enrolls in Medicare Advantage contract k in their market. Second, I use them to generate a measure of the financial generosity of the plans in a market. Formally:

$$\begin{aligned} Z_{kmt} &= \sum_{j \in \mathcal{J}_{kmt}} s_{jmt} \\ \mathbb{E}[OOP_{mt}] &= \sum_{j \in \mathcal{J}_{mt}} s_{jmt} \cdot OOP_{jmt} \end{aligned} \quad (4)$$

where \mathcal{J}_{kmt} is the set of plans included in contract k in market mt , s_{jmt} is the model implied

probability plan j is chosen in market mt and OOP_{jmt} are CMS produced estimates for the average out-of-pocket costs for plan j (i.e., Medicare Advantage or Traditional Medicare).

To empirically quantify the elasticity of healthcare consumption, I estimate the following model:

$$\log(1 + Q_{it}) = \beta_1 \log(1 + \mathbb{E}[OOP_{mt}]) + \beta_2 \mathbf{X}_{ijmt} + \beta_{k(j)} + v_{it} \quad (5)$$

where Q_{it} measures the healthcare utilization of beneficiary i during year t and \mathbf{X}_{ijmt} denotes a vector of individual-, market-, and plan-level characteristics. To capture measures Medicare Advantage plans put in place to impact healthcare utilization, I include contract fixed effects denoted by $\beta_{k(j)}$. Unobservable individual characteristics influencing healthcare utilization are captured by v_{it} . The coefficient of interest in this model is β_1 , which represents the elasticity of healthcare utilization with respect to its expected out-of-pocket costs.

I estimate equation (5) using two-stage least squares. The two endogenous parameters are the contract fixed effects $\beta_{k(j)}$ and the market-level plan generosity measure $\mathbb{E}[OOP_{mt}]$. I use the probabilities an individual enrolls in Medicare Advantage contract k Z_{kmt} constructed from the nested logit model as instruments for the contract fixed effects $\beta_{k(j)}$. For the market-level plan generosity measure, I use the leave-one-out plan benchmarks $B_{jt \setminus m}$. The validity of this design requires individual-level unobserved determinants of healthcare consumption to be conditionally independent of plan menu generosity and the probability of enrolling in contract k .

My estimation sample is a subset of the full utilization sample for 2017–2018. Specifically, it is restricted to individuals that were not enrolled in TM or MA in the prior year and as a result had to make an active plan choice. Table 3 reports the estimates for the elasticity of healthcare consumption. The first column reports the OLS estimate of -0.45. When instrumenting for sorting into contracts and plan menu generosity at the market-level this falls to -0.37. These estimates suggest most of the observed relationship between healthcare utilization and choice set generosity is attributable to moral hazard. The estimated elasticity

of -0.37 is greater than the benchmark of -0.2 from the RAND health insurance experiment (Manning et al., 1987). These differences may be attributable to the RAND elasticity being with respect to a plan’s actuarial value while the elasticity I estimate is with respect to monthly expected out-of-pocket costs. Column 3 uses an alternative measure for generosity at the contract-level within a market which produces an estimated elasticity of -0.18 that is marginally significant.¹³

This analysis highlights how the policy variation in the Medicare program generates choice set variation that can identify the elasticity of healthcare consumption. This margin is an important feature for the structural model of this market to capture as it influences consumer enrollment decisions as well as firm entry and product offering choices. A benefit of the structural model I developed is the ability to capture the feedback between these channels and their impact on market equilibria.

IV Empirical Model

IV.A Overview

This section provides an overview of the model. It begins with a description of individuals and their role within the model, followed by a similar treatment for firms and the government. The summary concludes with a discussion of timing and equilibrium.

Individuals. The model captures the decision of a senior eligible for Medicare, denoted by i , about their health insurance coverage for year t . These individuals are characterized by groupings of observed demographic characteristics (i.e., combinations of age, gender, low-income status, pre-existing health diagnostics, etc.) that are indexed by c , risk aversion ψ , and a propensity to consume additional healthcare when its price falls ω . These characteris-

¹³Formally this measure of generosity is defined as $\mathbb{E}[OOP_{kmt} = \frac{s_{jmt}}{\sum_{\ell \in \mathcal{J}_{kmt}} s_{\ell mt}} \cdot OOP_{jmt}]$.

tics are the private information of individuals and not observed by firms, which may create a selection problem from the health insurer’s perspective.

Individuals face a series of choices in the model. First, the senior must decide whether to enroll in a Medicare Advantage plan or Traditional Medicare. At the time of this choice, they do not know the realization of their health state for the year h_{it} . As a result, individuals form expectations about their health state and healthcare consumption. This expected healthcare utilization along with risk aversion and preferences for other plan characteristics factor into an individual’s health insurance coverage choice.

After choosing a health plan, individuals realize their health state and must decide how much healthcare to consume. An individual chooses the optimal amount of healthcare to consume Q_{ijt}^* by weighing the benefits of utilizing healthcare and their associated costs. These costs include administrative measures firms implement to limit healthcare consumption ϕ_{ijt} and the out-of-pocket costs paid by an individual given their chosen plan’s cost structure $OOP_{ijt}(Q_{ijt}^*)$. More financially generous health plans have lower out-of-pocket costs, which may induce some individuals to consume extra healthcare—sometimes referred to as “moral hazard.”

Firms. The model also captures the decisions of firms that may participate in Medicare Advantage markets. The set of potential entrants are firms that are endowed with CMS contracts to offer Medicare Advantage plans within a specified service area A —typically a state. These firms possess expertise and employ practices that allows them to offer health insurance benefits more efficiently (i.e., managed care), which is among the reasons why the government wants to tap into this private market to deliver these benefits. The model captures these efficiencies with the utilization cost parameter ϕ_{ijt} that appears when individuals choose how much healthcare to consume.

In the model, firms decide what counties within a service area to enter and which products to offer. In this setting, products are health insurance plans—indexed by j —which

have two key dimensions. The network type (i.e., HMO or Local PPO) is the first dimension. The type of network influences the form and strength of the utilization costs plans can use to control the amount of healthcare their enrollees consume. HMO plans are generally more restrictive than PPO plans. The second dimension is whether the plan has a high or low level of financial generosity. This decision impacts the amount of out-of-pocket costs enrollees pay for consuming healthcare. Both of these characteristics impact the amount of healthcare individuals expect to consume, which enters their health plan decision. After making these choices, firms set premiums for their plans based on the subsidies they request from the government.

Firm participation decisions are made on the basis of expected net profits, which are the difference between expected variable profits and fixed costs. Firms form expectations of what their variable profits will be given their own market entry and product offering decisions as well as those of their rivals. Variable profits will depend on who enrolls in each plan and how much healthcare those individuals consume. Fixed costs are a function of the products and markets the firm chooses to enter. These fixed costs represent the costs associated with provider networks, regulatory compliance, and market research. The optimal participation decision for a firm maximizes expected net profits given the participation decisions of rival firms.

Government. The model captures the government's role in setting policies that impact the functioning of this market. The first is Traditional Medicare's cost sharing, which determines the amount TM beneficiaries pay out-of-pocket for their healthcare consumption. The second is the subsidy scheme used to pay to Medicare Advantage plans. Under the current system the government sets county level cost benchmarks B_{mt} each year that reflects the historic costs the government has paid to provide TM benefits to individuals in county m . Firms observe these benchmarks and submit subsidy requests that reflect their costs for providing TM benefits to this population b_{jt} . These requests are evaluated against cost benchmarks to

determine whether the plan receives a rebate payment to fund additional benefits or if any costs are passed along to consumers as part of the plan’s premium p_{jt} .

Equilibrium. This model captures the strategic interaction between firms—indexed by n —that decide to enter Medicare Advantage markets. The model is set up as a two stage game and is summarized in Figure 2. During Stage 1, firms observe their fixed costs and the distribution of shocks they will face in Stage 2. Given this information and the cost benchmarks B_{mt} , firms simultaneously decide which plans to offer in each market within a service area. In Stage 2, firms choose their subsidies which determines the premiums for their plans. A firm’s strategy is a bundle of $(\mathcal{J}_{nt}, \mathbf{b}_{nt})$, where $\mathcal{J}_{nA} = \bigcup_{m \in A} \mathcal{J}_{nmt}$ is the set of plan offerings (i.e., network type and generosity level) the firm chooses in Stage 1 to offer in each market within service area A and \mathbf{b}_{nt} is the vector of subsidies the firm chooses in Stage 2 for each plan. The set of markets where firm n offers products is defined as $A_{nt} = \{m \mid m \in A \text{ where } \mathcal{J}_{nmt} \neq \emptyset\}$.

The model has a subgame perfect equilibrium (SPE).¹⁴ For a given set of Stage 1 strategies \mathcal{J}_t , the firm subsidy choices \mathbf{b}_t constitute a Nash equilibrium. When choosing these strategies, firms internalize how consumers will sort across plans offered to them and how they will consume healthcare given those plan choices. Formally, firms make their participation and subsidy decisions for service area A to maximize net profits:

$$\max_{(\mathcal{J}_{nt}, \mathbf{b}_{nt})} \Pi_{nt}(\mathcal{J}_{nt}, \mathcal{J}_{-nt}, \mathbf{b}_{nt}, \mathbf{b}_{-nt}) - F_{nt}(\mathcal{J}_{nt}) \quad (6)$$

where Π and F are firm n ’s variable profits and fixed costs respectively and $-n$ denotes the strategies of firm n ’s rivals. A strategy $(\mathcal{J}_{nt}^*, \mathbf{b}_{nt}^*)$ is a SPE if it maximizes firm n ’s net profits given the strategies played by rivals $(\mathcal{J}_{-nt}^*, \mathbf{b}_{-nt}^*)$

The model may have multiple equilibria. This multiplicity arises from different real-

¹⁴I assume the existence of the subgame perfect equilibrium for this model. Proving the existence of the equilibrium is beyond the scope of this paper.

izations of unobservable fixed costs for firms that can alter the set of markets the firm enters or products that are offered in those markets. Thus multiple SPE are possible where firms may optimally choose different \mathcal{J}_t^* 's that result in a unique Nash equilibrium for the subsidy choices \mathbf{b}_t^* . The following sections present the details of the model and its components. Consistent with solving for SPEs these components are presented in reverse order.

IV.B Demand

Healthcare utilization. This component of the model captures how an individual chooses how much healthcare to consume given their health insurance plan and realization of their health state h_{it} . The optimal amount of healthcare for an individual to utilize Q_{ijt}^* maximizes their utility given its associated costs, which depend on the type of plan the individual chose.¹⁵ Formally, an individual chooses Q_{ijt}^* to solve:

$$\max_{Q_{ijt}} u(Q_{ijt}; h_{it}, \omega_i, j) = v(Q_{ijt}, h_{it}, \omega_i) - \phi_{ijt}1[Q_{ijt} > 0] - OOP_{ijt}(Q_{ijt}) \quad (7)$$

where

$$v(Q_{ijt}, h_{it}, \omega_i) = Q_{ijt} - h_{it} - \frac{1}{2\omega_i h_{it}}(Q_{ijt} - h_{it})^2 \quad (8)$$

$$\phi_{ijt} = \exp(\mathbf{X}_{ijt}^\phi \boldsymbol{\beta}^\phi) \quad (9)$$

Following Einav et al. (2013) the value of healthcare utilization in Equation (8) is quadratic in the difference between the individual's healthcare utilization and health state. Intuitively, an individual aims to align their healthcare consumption with the need implied by their health state. The parameter ω_i captures how responsive an individual's healthcare utilization decision is to its costs and is typically interpreted as the their elasticity of demand

¹⁵Healthcare utilization is composed of inpatient, outpatient, physician, and hospice services. I do not model choices for prescription drug coverage and do not include it in my measure healthcare utilization.

for healthcare or moral hazard. Like Ho and Lee (2022), the moral hazard parameter is interacted with an individual’s health state, which implies that the effect of moral hazard is increasing in an individual’s health need. Individuals face two costs associated with healthcare utilization. The first is a “utilization cost” captured by ϕ_{ijmt} , which was first introduced by Ho and Lee (2022). This term captures the barriers individuals navigate to access care (e.g., provider networks, referrals, prior authorization, etc.). As show in Equation (9), cost varies with the network type of the plan an individual has chosen (i.e., TM, MA-HMO, or MA-PPO).¹⁶ The second cost of utilization is out-of-pocket costs, which are represented by $OOP_{ijt}(\cdot)$ and varies by plan type (i.e., network type and generosity level). Details about these cost structures and the solution to the utilization problem are available in Appendix C.

Health state distribution. The health state of individuals follows a log normal distribution $F_{it}(h)$:

$$\log h_{it} \sim \mathcal{N}(\mu_{it}, \sigma_{h,it}^2) \quad (10)$$

As noted in the literature, this distribution assumption captures the right skew in healthcare utilization. Variation in the parameters μ_{it} and $\sigma_{h,it}$ generates selection based on health need in the model by altering the amount of healthcare an individual chooses to consume. This selection is allowed to arise from both observable and unobservable characteristics.

The mean of an individual’s health μ_{it} and moral hazard ω_i are jointly normally distributed as follows:

$$\begin{bmatrix} \mu_{it} \\ \log \omega_i \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \mathbf{X}_{it}^\mu \boldsymbol{\beta}^\mu \\ \mathbf{X}_i^\omega \boldsymbol{\beta}^\omega \end{bmatrix}, \begin{bmatrix} \sigma_\mu^2 & \\ \sigma_{\mu,\omega} & \sigma_\omega^2 \end{bmatrix} \right) \quad (11)$$

where the means are a function of observable characteristics \mathbf{X}_{it}^μ for the health state mean and a constant for the moral hazard mean \mathbf{X}_i^ω . Unobserved heterogeneity in μ_{it} and $\log \omega_i$

¹⁶Utilization costs may also depend on individual characteristics. The current version of the model limits utilization costs to depend on network type but this can be relaxed.

arise from the distribution's variance and covariance parameters. The variance of the health state distribution $\sigma_{h,it}$ is modeled as a projection onto observable characteristics \mathbf{X}_{it}^σ :

$$\sigma_{h,it} = \mathbf{X}_{it}^\sigma \boldsymbol{\beta}^\sigma \quad (12)$$

At the time of their plan choices, individuals know the parameters of their health state distribution μ_{it} and $\sigma_{h,it}$ as well as their elasticity of healthcare consumption ω_i . This information influences their plan choice, which is how selection on both health and moral hazard arises within the model. Let θ_1 summarize the parameters of the health state distribution and utilization costs to estimate. This vector includes the mean shifters $\{\boldsymbol{\beta}^\mu, \boldsymbol{\beta}^\omega, \boldsymbol{\beta}^\sigma, \boldsymbol{\beta}^\phi\}$ and the variance-covariance parameters $\{\sigma_\mu, \sigma_\omega, \sigma_{\mu,\omega}\}$.

Plan choice. Individuals must choose among the health insurance plans available in their market \mathcal{J}_{mt} . Markets are defined as a county-year pair, where counties are indexed by m . Plans are classified into two groups indexed by g . The first group is Traditional Medicare $g = 0$ and the second group $g = 1$ contains all Medicare Advantage plans. An individual chooses the health insurance plan $j \in \mathcal{J}_{mt}$ that maximizes their expected utility over their health state distribution.

$$\max_{j \in \mathcal{J}_{mt}} U_{ijmt} = \int -\exp(-\psi \times l_{ijmt}(h, \omega, j)) dF_{it}(h) \quad (13)$$

where

$$l_{ijmt} = \delta_{jmt} + \gamma SR_{jt} + \alpha_{it} p_{jt}(b_{jt}) + \beta_{it} u(Q_{ijt}^*(h, \omega, j); h_{it}, \omega_i, j) + l_{ijmt} + \zeta_{ig} + (1 - \sigma) \epsilon_{ijmt} \quad (14)$$

$$\delta_{jmt} = \theta_2 X_{jmt} + \xi_{jmt} \quad \alpha_{it} = \alpha_0 + \alpha_1 y_{it} \quad \beta_{it} = \beta_0 + \beta_1 y_{it}$$

The coefficient of absolute risk aversion is denoted as ψ and is common to all beneficiaries. To rule out risk-loving preferences, I constrain the CARA coefficient to be non-negative

$\psi = \exp(\beta_\psi)$. The term l_{ijmt} summarizes the utility individual i receives from plan j . The outside option is Traditional Medicare ($j = 0$) whose expected utility is normalized to one.

The factors that enter l_{ijmt} are noted in Equation (14). The first term is the mean utility of the plan common to all individuals in the market. This mean utility may depend on observable characteristics like the plan's star rating or provider network (X_{jmt}) and an unobservable demand shock ξ_{jmt} . The second is the supplemental revenue the plan requires to fund additional benefits beyond what is covered by TM SR_{jt} . The next term is the plan premium $p_{jt}(b_{jt})$, which may differentially impact beneficiaries with low-incomes (y_{it}).¹⁷ The third component is the individual's utility they will receive from the plan given their health state realization and the amount of healthcare they expect to utilize. These quantities depend on the amount of out-of-pocket costs and the utilization costs the individual will incur, which depend on the network type and generosity of plan j . The fourth component ι_{ijmt} captures the switching costs of changing between TM or MA, capturing the inertia in plan choices. These costs are not incurred by new beneficiaries or those that switch among MA plans. The final components are the unobservables ζ_{ig} and ϵ_{ijmt} , which are assumed to jointly follow a generalized extreme value distribution.

As noted previously, individuals are classified into categories based on their observable characteristics, which are indexed by c . Let $s_{c jmt}$ denote the probability that individuals in group c in market mt choose plan j .¹⁸ The plan's market share s_{jmt} is obtained by integrating these choice probabilities over the distribution of observable types within the market. Finally, let θ_3 denote the parameters in the utility function that are independent of mean utility $\{\beta_\psi, \gamma, \alpha_0, \alpha_1, \beta_0, \beta_1, \iota\}$.

¹⁷I do not observe a continuous measure for income in my data. The administrative data contain indicators for low-income status.

¹⁸Individual of the same type have the same amount of expected healthcare utilization.

IV.C Supply

Subsidy choice and plan premiums. Continuing backwards, the next action within the model is how plans choose their subsidy payments from the government. A firm chooses the subsidy for each plan by maximizing their expected profits across all markets the plan entered within a service area.¹⁹ Let A_{jt} denote the set of counties plan j entered within service area A in year t . Given the set of plan offering decisions the firm made in Stage 1 \mathcal{J}_{nAt} , the firm chooses the subsidy vector \mathbf{b}_{nt} by solving:

$$\max_{\mathbf{b}_{nt}} \Pi_{nt} = \sum_{m \in A_{nt}} \sum_{j \in \mathcal{J}_{nmt}} \sum_{c \in C} \int [\text{MR}_{cjmt} - \text{MC}_{cjmt}] s_{cjmt}(\mathbf{b}; \Theta) M_{cmt} dF_{ct}(h) \quad (15)$$

where M_{cmt} is the number of type c beneficiaries in market mt and MR_{cjmt} and MC_{cjmt} are specified as:

$$\text{MR}_{cjmt} = \bar{r}_{cmt} \min\{b_{jt}, B_{jt}\} + p_{jt}(b_{jt}) \quad \text{MC}_{cjmt} = \Lambda_{jt} Q_{cjt}^*(h, \omega, j) + \lambda_{jt} \quad (16)$$

The marginal revenue for enrolling an individual of observed type c is denoted by MR_{cjmt} . Two items contribute to marginal revenue. The first is the subsidy payment the plan receives from the government, which is equal to the requested b_{jt} if the request is below the plan's cost benchmark $B_{jt} = \sum_{m \in A_{jt}} B_{mt} w_{mt}$, where w_{mt} are market size weights.²⁰ If the requested subsidy is above the benchmark, the plan's subsidy payment is equal to the benchmark. These payments from the government are risk adjusted based on the average risk score of beneficiaries of observed type c in the market \bar{r}_{cmt} . The second and third components of marginal revenue are the plan premium, which depends on the subsidy request for the plan. Premiums are discussed later in this section.

The marginal cost for a type c individual denoted by MC_{cjmt} . This cost is broken

¹⁹In general service areas are states. A more detailed discussion of service areas is provided in Appendix C.

²⁰In practice county-level benchmarks are weighted by the plan's projected enrollment. Market size weights ease the burdens for computing the model's solution. Market size is also highly correlated with realized enrollment.

down into components. The first term captures how plan costs depend on the amount of healthcare they expect beneficiaries will consume. The term Λ_{jt} represents the price that MA plan j pays providers for the healthcare utilization of the beneficiaries in their plan. Prior empirical work has documented that MA plans tend to pay similar prices to healthcare providers as TM.²¹ Consistent with these fact patterns I assume $\Lambda_{jt} = 1$. The second term λ_{jt} is an unobserved cost that captures non-utilization contributions to marginal costs and rationalizes the observed pricing and entry equilibrium.

The policy environment makes Medicare Advantage plan premiums a function of the plan's subsidy b_{jt} and whether the plan offers additional benefits relative to TM. These features are clear when looking at the two components of the plan premium paid by consumers:

$$p_{jt}(b_{jt}) = \underbrace{\max\{b_{jt} - B_{jt}, 0\}}_{\text{base}} + \underbrace{\max\{\text{SR}_{jt} - \text{Rebate}_{jt}, 0\}}_{\text{supplemental}} \quad (17)$$

where $\text{Rebate}_{jt} = \max\{\kappa_{jt}(B_{jt} - b_{jt}), 0\}$ and the size of κ_{jt} depends on the star rating of the plan.²² The supplemental revenue MA plans need to provide additional benefits is denoted as SR_{jt} . Plans can offset these costs if they receive a rebate payment.

I model the amount of supplemental revenue a plan needs to fund additional benefits relative to TM as function of plan characteristics W_{jt} that includes the plan's network type, quality rating, and generosity level as well as the CMS benchmark the plan faces. The variable ε_{jt} denotes an efficiency shock the plan receives to the amount of revenue required to fund these extra benefits.

$$\text{SR}_{jt} = \theta_4 W_{jt} + \varepsilon_{jt} \quad (18)$$

Fixed costs of entry. Firms are endowed with CMS contracts that define the set of possible plans they may offer within a service area A . Each year, firms decide which plans

²¹See e.g., Curto et al. (2019), Pelech (2020), and Trish et al. (2017).

²²The levels of κ_{jt} are 0.50 if the plans has 3 stars or fewer, 0.65 if the plan has 3.5 or 4 stars, and 0.70 if the plan has 4.5 or 5 stars.

they will offer in each market within the service area. The primary fixed cost of entry into a market is establishing a new or updating an existing network of providers enrollees may use to receive healthcare services. A Medicare Advantage plan’s provider network must annually certify that it meets network adequacy and access criteria established by CMS. Given this institutional setting, it is useful to think of the entry decision as reoccurring each year, which abstracts from distinctions between sunk vs fixed costs of entry.

I assume that a firm’s fixed cost for offering MA plans is additively separable across markets and has an observable and unobservable component. The fixed cost for insurer n to offer MA plans in year t is:

$$F_{nt} = \sum_{m \in A_{nt}} \sum_{j \in \mathcal{J}_{nmt}} [F_{nmt} + \nu_{2jnmt}] \quad (19)$$

The observable component of the fixed cost of entering market m has three parts. The first captures the number of plans the firm has chosen to enter into market m . The second and third components are measures of provider supply. Specifically, H_{mt} denotes the number of hospital systems in the market and P_{mt} denotes the number of primary care physicians active in the market. These terms are intended to capture—in a reduced form manner—the costs of bargaining with providers to join the firm’s network. I allow the parameters on these terms to vary based on whether firm n has an existing provider network in the market from another insurance segment (e.g., commercial group, individual, exchange etc.). This feature captures efficiencies some insurers may have that eases entry into Medicare Advantage.

$$F_{nmt} = \rho_1[\text{Number of plans}]_{nmt} + \rho_{2n}H_{mt} + \rho_{3n}P_{mt} \quad (20)$$

where

$$\rho_{\{2,3\}n} = \rho_{\{2,3\}\text{net}}1[\text{Other presence}]_{nmt-1} + \rho_{\{2,3\}\text{none}}(1 - 1[\text{Other presence}]_{nmt-1}) \quad (21)$$

The unobserved component of fixed costs are denoted by $\nu_{2jnm t}$, which are independent over time. These costs are observed by firms when making their Stage 1 decisions and the selection problem they create is discussed more in Section V.

After observing $\nu_{2jnm t}$, firms simultaneously choose which plans to enter into a market by weighing their expected profits against the fixed costs of entry. Firms calculate their expected profits over the joint distribution of the Stage 2 unobservables $e = (\xi, \varepsilon_1, \varepsilon_2)$ and distribution of health states in the population. I assume that firms know the form of this distribution but not the realizations they will face. The unobservable $\nu_{1jnm t}$ denotes a mean zero expectation error, which implies firms on average accurately predict their variable profits. Based on these factors, there are three entry conditions which govern the actions of firms.

Condition 1. *Firm participation decisions must result in positive net profits:*

$$\sum_{m \in A_{nt}} \sum_{j \in \mathcal{J}_{nm t}} \underbrace{\mathbb{E}[\Pi_{jnm t}(\mathcal{J}_t)] + \nu_{1jnm t}}_{\text{expected variable profits}} - \underbrace{(F_{nm t} + \nu_{2jnm t})}_{\text{fixed costs}} \geq 0 \quad (\text{EC.1})$$

where \mathcal{J}_t denotes the vector of strategies played by all firms in the service area.

Condition 2. *Firms choose the markets to enter and products to offer optimally:*

$$\sum_{m \in A_{nt}} \sum_{j \in \mathcal{J}_{nm t}} \mathbb{E}[\Pi_{jnm t}(\mathcal{J}_t)] + \nu_{1jnm t} - (F_{nm t} + \nu_{2jnm t}) \geq \sum_{m \in A'_{nt}} \sum_{j \in \mathcal{J}'_{nm t}} \mathbb{E}[\Pi_{jnm t}(\mathcal{J}_t)] + \nu_{1jnm t} - (F_{nm t} + \nu_{2jnm t}) \quad (\text{EC.2})$$

$$\forall A'_{nt} \neq A_{nt} \text{ and } \mathcal{J}'_{nm t} \neq \mathcal{J}_{nm t}$$

Condition 3. *Conditional on entering a market, firms choose to offer the set of products*

that yields the highest net profits:

$$\sum_{j \in \mathcal{J}_{nmt}} \mathbb{E}[\Pi_{jnmt}(\mathcal{J}_t)] + \nu_{1jnmt} - (F_{nmt} + \nu_{2jnmt}) \geq \sum_{j \in \mathcal{J}'_{nmt}} \mathbb{E}[\Pi_{jnmt}(\mathcal{J}_t)] + \nu_{1jnmt} - (F_{nmt} + \nu_{2jnmt}) \quad (\text{EC.3})$$

$$\forall \mathcal{J}'_{nmt} \neq \mathcal{J}_{nmt} \text{ and } \forall m \in A$$

Notice these conditions capture how firm decisions are interconnected across markets. A plan's cost benchmark are weighted averages of the county specific benchmarks where a plan is offered. These benchmarks play a critical role in determining a plan's marginal revenue and additional benefits. As a result, firms must consider how entering or exiting a particular market impacts their overall net profits.

V Estimation and Identification

This section describes how the model is estimated. I follow the generalized method of moments to estimate the health state, consumer preference, and fixed cost parameters. The first subsection focuses on the Stage 2 parameters—health states and consumer preferences—and relevant implementation details. The second subsection describes how I derive the moment inequalities to recover the identified set of fixed cost parameters in Stage 1. This discussion includes a description of how the moment inequalities are used for inference. The section concludes with an overview of the how the model's parameters are identified from the data.

V.A Health state and consumer preferences

Moments. To estimate the health state distribution parameters, I match moments based on healthcare utilization patterns. Specifically, I target the unconditional mean and vari-

ance of healthcare utilization as well as the mean and variance of utilization conditional on observables such as risk score quantiles. These moments help the model replicate the relationship between observable characteristics and healthcare utilization seen in the MA encounter and TM claims data. To capture the propensity to consume healthcare as its cost decreases (moral hazard), I include the mean and variance of the utilization distribution across quantiles of plan choice set generosity and risk scores. As discussed in the prior section, variation in healthcare utilization across markets with different levels of financial generosity captures utilization not driven by health need. Choice set generosity is measured by the average rebate payment paid to MA plans in the market. To further capture “moral hazard” spending, I also target the average healthcare utilization conditional on being in the coinsurance region. I match the utilization cost parameters with the average probability of consuming no health care conditional on plan type.

I also include moments based on plan choices to estimate consumer preferences for health plans. To capture risk aversion and consumer sorting across plans, I target plan choice probabilities conditional on plan and consumer types. As discussed in Appendix C, I constrain plan-level market shares implied by the model to match their observed analogs. I also use moments based on IV restrictions in the demand model to address endogenous premiums and supplemental benefits. This condition requires that the unobserved demand shock ξ_{jmt} is uncorrelated with a vector of instruments Z_{jmt} . I describe the types of instruments I use and the intuition they bring to the identification argument later in this section. The specific instruments that I use include the plan’s marginal revenue around the benchmarks; the number of hospitals, hospital beds, and primary care physicians active in a plan’s footprint in the previous year; and the average characteristics of non-overlap rival counties (i.e., share rural, share with college degree, median income, share female, share white, share of all Medicare beneficiaries that died, and the share of Medicare beneficiaries eligible for Medicaid).

Implementation details. My analysis relies on the MA encounter data to measure healthcare utilization among MA beneficiaries. There are two challenges to working with these data. The first is the absence of payment information. I overcome this shortcoming by using a measure of healthcare utilization based on TM prices that was proposed by Jung et al. (2022) specifically for MA encounter data. I follow their implementation for deriving these standardized prices using all of the claims and encounter data available to me. I then merge these utilization metrics onto the MA encounter and TM claims data for consistency.

The second challenge relates to the completeness of the encounter data that private insurers report to CMS.²³ To attenuate this concern, I follow the procedures in Jung et al. (2022) to assess the completeness of the encounter data. These authors propose a test based on comparing the encounter data to other sources that contain information about MA healthcare utilization (i.e., the Medicare Provider Analysis and Review (MedPAR) and the Healthcare Effectiveness Data Information System (HEDIS)). MA contracts have a high level of data completeness if they meet minimum thresholds for enrollment and the difference between the number of hospitalizations, ambulatory, or emergency department visits recorded in the encounter data and MedPAR or HEDIS. Appendix Table E.4 highlights that there are no systematic differences between MA beneficiaries enrolled in plans with a high degree of data completeness relative to those that are not.²⁴ Additionally, the utilization patterns I observe across TM and MA beneficiaries are consistent with other studies that do not rely on encounter data, which further mitigates concerns about encounter data completeness (Curto et al., 2019).²⁵

²³“Completeness” is the notation that all encounter records for a plan’s beneficiaries appear in the data provided by CMS.

²⁴A similar exercise is presented in Appendix Table E.5 for TM beneficiaries. Individuals in the TM claims data are marginally more likely to be female or low income but the size of the difference is modest.

²⁵Curto et al. (2019) find in 2010 for three MA insurers covering 40% of MA enrollees that the unadjusted difference in utilization in MA was 30% lower than TM. Since they also found that MA plans paid prices similar to TM, this gap can be directly attributed to reduced utilization of healthcare services by MA beneficiaries. Once controls are added this gap becomes 9–25% lower than TM. Due to the large growth in MA penetration since 2010, it is intuitive that this gap has gotten smaller over time as more TM beneficiaries enroll into MA plans.

Risk scores play an important role in my analysis. The risk scores that CMS calculates for each Medicare beneficiary are generally not produced in the files made available to researchers. However, CMS does provide the algorithms to generate these risk scores based on the demographic and diagnosis information that is made available. I lack the data to fully replicate the CMS risk scores because I do not have utilization data for all Medicare beneficiaries. I address this challenge by approximating the CMS risk score using their published formula and the diagnoses available to me from inpatient claims and discharges, which I have for the universe of Medicare beneficiaries.²⁶ I generate the base risk score using the CMS algorithm for the appropriate year with beneficiary demographics and prior year inpatient diagnoses. These base scores are then normalized by the average base score for all TM beneficiaries that year. Finally, risk scores for beneficiaries that were in a MA plan the previous year are deflated by the coding pattern adjustment reported by CMS.²⁷

To assess the quality of my approximated risk score, I aggregate my scores to levels where CMS reports average risk scores. Figure 3 plots the distribution of average risk scores at the county-level from 2017–2018 that I calculated against the values CMS reported for those counties. The means of the two distributions are nearly identical. The variance of the approximated risk score distribution is smaller relative to the true variance, consistent with the missing diagnoses. For a model of endogenous plan participation, approximating the risk score distribution well on average captures the first-order effects that impact whether a firm chooses to operate in a market. Any bias in the estimates of a firm’s entry incentives from the reduced variance of the distribution is likely small.

V.B Moment inequality derivation and inference

Derivation. To derive the moment inequalities for estimating the identified set of fixed costs I need the distribution of Stage 2 shocks and resolve the selection bias introduced by

²⁶Since I have 100% of TM inpatient discharges, I have all diagnoses recorded in the inpatient claims that I do not possess.

²⁷These adjustments were 5.66% and 5.91% in 2017 and 2018, respectively.

the unobserved fixed costs ν_{2jnm} . The Stage 2 distribution of unobservables $e = (\xi, \varepsilon)$ is required to calculate a plan's expected variable profits. I recover this empirical distribution given estimates for the Stage 2 model parameters $\Theta = \{\theta_1, \theta_2, \theta_3, \theta_4\}$. The unobserved fixed costs ν_{2jnm} create a selection problem because firms observe these costs when making their entry decisions. The following assumption allows me to address this bias.

Assumption 1. *A plan offered in adjacent markets within a service area—typically a state—has the same unobserved fixed cost ν_2 .*

Assumption 1 is supportable when viewing the unobserved fixed costs as regulatory compliance, business intelligence, and marketing, which are unlikely to vary meaningfully across markets. Firms likely rely on common personnel for these tasks and the amount of resources devoted to them likely scales with the number of markets a particular plan enters.

Unbiased moment inequalities are derived based on revealed preference, the separability of fixed costs, and Assumption 1. Revealed preference requires that the entry and product offering decisions observed in data are optimal relative to the other choices that the firm *could* have made. Let A_{nt} and \mathcal{J}_{nt} denote the observed market and product offerings decisions firm n made in service area A and A'_{nt} and \mathcal{J}'_{nt} denote their unobserved analogs. Revealed preference implies:

$$\begin{aligned} \sum_{m \in A_{nt}} \sum_{j \in \mathcal{J}_{nm}} \mathbb{E}[\Pi_{jnm}(\{\mathcal{J}_{nt}, \mathbf{b}_{nt}\}_{\forall n \in A})] + \nu_{1jnm} - (F_{nm} + \nu_{2jnm}) \geq \\ \sum_{m \in A'_{nt}} \sum_{j \in \mathcal{J}'_{nm}} \mathbb{E}[\Pi_{jnm}(\{\mathcal{J}_{nt}, \mathbf{b}_{nt}\}_{\forall n \in A})] + \nu_{1jnm} - (F_{nm} + \nu_{2jnm}) \end{aligned} \quad (22)$$

Suppose the firm removes plan j from market m such that $A'_{nt} = A_{nt} \setminus m$. I rearrange the terms in Equation (22) such that:

$$\sum_{m \in A} \sum_{j \in \mathcal{J}_{nm}} \Delta \mathbb{E}[\Pi_{jnm}(A_{nt}, A'_{nt})] + \Delta \nu_{1jnm}(A_{nt}, A'_{nt}) - F_{nm} - \nu_{2jnm} \geq 0 \quad (23)$$

where $\Delta X(A_{nt}, A'_{nt}) = X(A_{nt}) - X(A'_{nt})$.

I can derive a similar inequality by adding market m' to plan j 's observed footprint such that $\hat{A}_n = A_n + m'$. Rearranging terms yields:

$$\sum_{m \in A} \sum_{j \in \mathcal{J}_{nmt}} \Delta \mathbb{E}[\Pi_{njmt}(A_{nt}, \hat{A}_{nt})] + \Delta \nu_{1jnm}(A_{nt}, \hat{A}_{nt}) + F_{njm't} + \nu_{2jnm't} \geq 0 \quad (24)$$

The separability of unobserved fixed costs allows me to isolate a specific plan's ν_{2njm} shock for each perturbed market using Equations (23) and (24). By Assumption 1 $\nu_{2jnm't} = \nu_{2jnm't}$ if m and m' are adjacent. This allows me to bound $\nu_{2jnm't}$ and combine these equations such that:

$$\sum_A \sum_{\mathcal{J}_{nmt}} \Delta^+ \mathbb{E}[\Pi(m, m')] + \Delta^+ \nu_1(m, m') - \Delta^- F(m, m') - \underbrace{(\nu_{2jnm't} - \nu_{2jnm't})}_{\approx 0} \geq 0 \quad (25)$$

where $\Delta^+ X(m, m') = \Delta X(A_{nt}, A'_{nt}) + \Delta X(A_{nt}, \hat{A}_{nt})$ and $\Delta^- X(m, m') = \Delta X(A_{nt}, A'_{nt}) - \Delta X(A_{nt}, \hat{A}_{nt})$.

It remains to address the approximation errors ν_1 . Recall that these errors are mean zero across all markets within a service area. This error is eliminated by averaging over all the pairwise combinations of Equation (25) for each market within a service area. This procedure yields a set of unbiased moment inequalities for plan j .

$$\mathbb{E}[m_j(\theta)] = \sum_A \sum_{\mathcal{J}_{nmt}} \mathbb{E}[\Delta^- F(m, m') - \Delta^+ \mathbb{E} \Pi(m, m') - \Delta^+ \nu_1(m, m')] \leq 0 \quad (26)$$

where the expectation is taken over adjacent market combinations within a service area.

I generate additional inequalities by interacting each plan inequality with a set of “instruments” that are independent of the unobservable ν_2 terms. Specifically, I leverage the independence over time assumption and use lagged counts of markets with existing provider networks and provider supply counts as instruments. These instruments are valid because the

unobservable fixed costs are independent over time. These two types of moment inequalities form the null hypothesis for the inference procedure I use to construct an estimate for the identified set of fixed cost parameters.

Inference. I follow the inference procedure proposed by Chernozhukov et al. (2019), which is well-suited for models with many moment inequalities. Their procedures are built around a studentized test statistic that detects violations of the moment inequalities.

$$T = \max_{1 \leq k \leq K} \frac{\sqrt{D}\varphi_k}{\varsigma_k} \quad (27)$$

where k indexes the moment inequalities, K denotes the total number of inequalities, φ and ς are the mean and standard deviation of the moment inequalities, and D is the total number adjacent market pairs for a plans.

I implement the self-normalized one step procedure, which has a closed form for its critical values. This feature lowers the procedure’s computational burden relative to multi-step or bootstrap alternatives. The tradeoff is that the identified sets may be more conservative. Additional details related to the computation of the moment inequalities and the inference procedure are presented in Appendix C.

V.C Identification

The objects to identify in the model are the joint distribution of individual health states, moral hazard, utilization costs, and consumer preferences for differentiated health insurance plans. The ideal data set for this exercise has two key characteristics. First, it would track individuals over time and measure their health states. Second, the data would contain variation in how individuals are exposed to different choice sets of health insurance plans with alternate levels of financial coverage. The source of this variation in plan choice sets is driven by exogenous changes in a policy instrument. This data set would capture how plan

enrollment and healthcare utilization change as variation in the policy alters the average level of financial generosity of the health plan choice set. These data could facilitate non-parametric identification of the model's parameters.

In most practical applications the ideal data set is not obtainable and additional assumptions are required. Relative to the ideal data set, my administrative data has substantial cross-sectional variation in health insurance plan choice sets—every county in the United States—but relatively short panel variation—two years after constructing ex ante risk scores. Given these realities, parametric assumptions are necessary to assist identification. The benchmarks CMS sets at the market-level each year are a source of variation in the size and generosity of health insurance plan choice sets available to consumers. These market-level benchmarks form the plan-level benchmarks firms face when making their product offering decisions. Thus, variation in the benchmarks in other markets provides a plausibly exogenous source of variation in the number of plans in a market's choice set as well as their financial generosity.

Given these parametric assumptions and plausible exogenous variation in choice sets, the parameters associated with healthcare utilization are identified. The extent to which consumers make similar healthcare utilization choices when facing similar choice sets over time identifies the persistent component of the health state distribution. Variation in healthcare utilization over time among similar individuals, aided by the distributional assumption, identifies unobserved heterogeneity in these decisions. The parameters that influence moral hazard—the propensity to consume more healthcare when it is less expensive—are identified by variation in healthcare utilization as the generosity of choice sets respond to variation in CMS benchmarks. Deviations from trends in healthcare utilization by network type induced by variation in benchmarks identifies changes in the threshold health state needed for healthcare consumption (i.e., utilization costs).

These sources of variation also identify consumer preferences for health insurance plans. Risk aversion is identified by how consumers choose health plans as variation in benchmarks

alters the generosity of plans within consumer choice sets. The extent to which consumers with comparable health needs pick more generous plans captures a measure of their tolerance for uncertainty about the out-of-pocket costs associated with their expected health need. Preferences for plan characteristics that are common across its footprint are identified by the extent to which consumers opt into the plan across markets and over time. Switching costs between TM to MA and vice-versa are identified by the extent to which beneficiaries remain within each program over time.

To address the potential correlation between unobserved plan-market level demand shocks and the premiums and extra benefits chosen by firms, I rely on instrumental variables. These instruments must be correlated with a plan’s premium and extra benefits but independent of the plan-market shock. Many instruments are possible in this setting including variations of the widely used Berry et al. (1995) and Hausman (1997) instruments. I use two types of instruments. The first set is based on CMS policies which are exogenous to firm pricing decisions yet correlated with a plan’s subsidy choice and premium. Specifically, I use a plan’s marginal revenue around the benchmark—determined by the κ parameter. The second set of instruments are demographics from non-overlapping markets of rival plans. The intuition for the market demographics instruments follows Fan (2013). Healthcare utilization is correlated with observable characteristics. Thus, the demographics of the Medicare population in a county influence the costs of offering a MA plan. Suppose there are two plans A and B , which overlap in market 1, while only plan B is present in market 2. The demographics of market 2 directly impact plan B ’s choices and indirectly impact plan A ’s choices through the competition channel in market 1. Thus, the demographics from market 2 can serve as an instrument for plan A ’s choices in its markets.

Finally, the fixed costs of entry within the model are partially identified. These parameters cannot be point identified without imposing the assumption of an equilibrium selection mechanism. I use a revealed preference approach in the spirit of Pakes et al. (2015) to derive moment inequalities that are consistent with these multiple equilibria. Revealed preference

is based on the assumption that firms are making optimal decisions based on the information available to them at the time of their action. This condition allows me to determine that other choices the firm could have made—yet did not—must be weakly less profitable. In prior section, I illustrated how I use this assumption to derive unbiased moment inequalities to recover the identified set of fixed cost parameters. I further leverage an exclusion restriction based on the independence of the unobservables in the firm’s entry problem over time to provide additional bounds on the identified set.

VI Results

VI.A Estimates

The top panel of Table 4 contains the demand parameter estimates. Parameter estimates for the health state distribution are available in Appendix Table E.6.²⁸ In general, parameter estimates have the correct sign and are significant. Demand slopes down in premiums with lower income beneficiaries having a higher degree of sensitivity and slopes up in the value of the supplemental benefits offered by MA plans. Consumer preferences also depend on the value of healthcare they expect to consume net of utilization and out-of-pocket costs. Enrollment choices respond more to upfront costs represented by premiums than the utility from healthcare utilization. The estimate of the nesting parameter is significant and indicates MA plans are closer substitutes to each other than TM. The estimated switching costs between the programs are modest and consumer risk aversion is more consistent with risk neutral behavior.²⁹ More risk neutral behavior in this setting could reflect the low financial risk seniors face in this market. MA plans provide generous cost sharing and out-of-pocket

²⁸Estimates for the supplemental revenue regression are available in Appendix Table E.7.

²⁹Switching costs range from \$80–120. The estimated CARA coefficient implies an individual would be indifferent between earning nothing and a 50-50 gamble where they win \$100 or lose \$96.37. The literature has produced similar estimates for the average level of risk aversion: Dickstein et al. (2023) \$99.32 and \$97.40; Ho and Lee (2022) \$99.97; Marone and Sabety (2022) \$91.70; Handel (2013) \$91; and Einav et al. (2013) \$84.

maxima, which makes choosing among them akin to short term gambles over relatively small amounts of money.

The bottom panel of Table 4 presents quantities implied by the demand and utilization model. I start by evaluating the amount of moral hazard estimated by the model. In this context, moral hazard captures the propensity for individuals to consume more healthcare as the cost utilization falls. To measure this force, I simulate how healthcare utilization changes as the coinsurance rate moves from 100% to 0% holding deductibles and out-of-pocket maximums fixed. The changes in utilization are greatest for beneficiaries enrolled in TM (49%) relative to individuals in a MA plan (4%–8%).³⁰ The different magnitude of the moral hazard effects between TM and MA is consistent with utilization costs. MA plans take measures to limit the amount of healthcare their enrollees utilize, which are not present in TM. These effects are driven by the utilization cost parameters, whose implied dollar values are about \$140 for TM, \$1,180 for MA HMOs, and \$880 for MA PPOs. These implied MA utilization costs are in line with estimates from Ho and Lee (2022) which ranged from \$550–\$1,710.

Table 5 contains estimates for the identified set of fixed cost parameters. For computational reasons, I use a subset of moment inequalities from 20% of service areas. The identified set does not contain zero for any of the fixed cost parameters and their signs have intuitive interpretations. For example, fixed costs increase with the number of plans offered within a market. This estimate appears consistent with the significant amount of regulatory compliance Medicare Advantage plans must satisfy and complete before entering the marketplace. My estimates suggest that fixed costs are substantially lower (roughly 75% based on the median of the intervals) in markets where the firm has an existing provider network. This finding is consistent with firms having to devote fewer resources to establish a provider network for their Medicare Advantage offerings.

³⁰Other studies have estimated similar amounts of moral hazard: Dickstein et al. (2023) 22% and 11%; Ho and Lee (2022) 26.3% and 3.5%; Marone and Sabety (2022) 24% and 14%; and Einav et al. (2013) 30%.

VI.B Model Fit

Figure 4 presents a subset of the data moments targeted in estimation alongside their model predicted counterparts. The top panel contains the unconditional mean and variance of the utilization distribution. The model almost perfectly matches the mean and slightly under predicts the variance. The bottom panel presents the mean and variance of the utilization distribution conditional on risk score, which is a strong predictor of utilization. Again, the model closely fits the mean moments and the variance moments for the second and third quartiles of risk score distribution. Consistent with the unconditional variance moment, the model under fits the variance in the tails of the risk score distribution.

Figure 5 presents non-targeted data moments alongside their model analogs. Overall, my model successfully captures how consumers sort across health insurance plans and utilize healthcare. The left panel shows average utilization by plan type and the right panel shows the variance of utilization by plan type. The model slightly over-predicts average utilization in MA plans and under-predicts average TM utilization. The model does replicate relative differences in observed utilization by plan types. For example, utilization is higher in TM than MA—consistent with utilization costs—and utilization is higher in more generous MA plans—consistent with more generous cost sharing. Capturing these patterns is important as they allow the model to reflect the selection patterns observed in Medicare and the supply side considerations firms face when deciding which markets to enter and types of products to offer.

Overall, the model fit is reasonable. On average the model closely matches observed utilization patterns by demographic characteristics. It slightly over-predicts the mean utilization in MA plans while under-predicting these quantities for TM and the variance of healthcare utilization. That said, the model accurately reflects sorting and utilization dynamics by plan types, which are key features for the supply side of the model.

VII Counterfactuals

In this section I use my estimated model to quantify the tradeoffs of promoting private firms to participate in the market for Medicare benefits. I describe the simulation setting and details in Section VII.A. To build intuition for later results, in Section VII.B I demonstrate how my modeling assumptions impact equilibrium outcomes. Then in Section VII.C, I simulate the effects of four distinct subsidy policies for promoting entry and participation in this competitive insurance market and quantify their tradeoffs.

VII.A Simulation setup

My simulations focus on the 2018 Massachusetts service area, which is summarized in Table 6. Like most MA markets, Massachusetts is highly concentrated. The top two firms—Blue Cross Blue Shield of Massachusetts (BCBS) and Tufts Health Plan (Tufts)—controlled over 56% of all MA enrollment in 2018. Tufts is the market leader and offers HMO plans of high and low generosity in 8 markets. BCBS primarily offers PPO plans of high and low generosity in 11 markets. The remaining share of the market is spread across five firms which primarily offer HMO plans.³¹

For the simulations, I make two assumptions for tractability. These assumptions can be relaxed as computational resources allow and do not alter the underlying model. First, I assume that BCBS and Tufts are strategic players that choose which markets to enter and products to offer. Each firm is restricted to offering plan types that align with their observed network offering (i.e., HMO or PPO) but can choose the financial generosity of the plans. The other firms are treated as a competitive fringe whose product offering choices are taken as exogenous. The fringe plans set their prices as a function of the average strategic price. Thus, the choice set for an endogenous firm in each market is to offer no plan, a low generosity plan, a high generosity plan, or both. Second, I assume that entry decisions

³¹One of these firms offers PPO and HMO plans but has a state-wide market share of 1%.

are made at the Combined Statistical Area (CSA) level. CSAs are groupings of counties used by the U.S. government for adjacent communities that demonstrate economic or social linkages.³² Massachusetts has two CSAs; based around Boston and Springfield. I group all other counties in Massachusetts into a third pseudo-CSA. This assumption is supported by observed entry patterns in Massachusetts. Firms that enter one of the markets within a CSA typically enter the others as well. Thus, a strategic firm must decide for each plan they offer whether to enter no markets, Boston-area markets, Springfield-area markets, other markets, or a combination of these markets. Given the number of players in the game, the size of their choice sets, and draws necessary to calculate expected profits, I need to compute 40,960 pricing equilibria for each counterfactual.

To solve for the equilibria of the model, I follow the procedure proposed by Lee and Pakes (2009). This method has been used by other papers that solve models with multiple equilibria (see e.g., Wollmann, 2018). The procedure uses a best response iteration approach to find the entry and product offering equilibria that are consistent with the Stage 1 necessary conditions in Equation (EC.1). I compute fixed costs similar to other moment inequality papers in the literature (see e.g., Geddes, 2022; Wollmann, 2018). I evaluate the observed fixed costs at the median values from the estimated identified set. Given these estimates, I recover ranges for the unobserved fixed cost ν_2 that are consistent with the moment inequalities for the strategic firms. I take 100 random draws from a normal distribution with a mean and variance calibrated from these ranges and recover the pure strategy equilibria associated with each fixed cost realization. Additional details on how I compute equilibria of the model are available in Appendix D.

I define consumer surplus as an individual’s expected certainty equivalent utility from enrolling in a plan. The literature has used similar measures of consumer welfare (see e.g.,

³²In practice CSAs can span states. The service areas defined in model do not span states. As a result, I focus on CSA groupings of counties within a state if the CSA spans multiple states.

Einav et al., 2013; Ho and Lee, 2022). Thus the consumer surplus for individual i is:

$$CS_{it} = \int \frac{1}{-\alpha_i} \log \left[1 + \sum_{j \in \mathcal{J}_{mt}} \exp(U_{ijmt}^{CE}) \right] dF \quad (28)$$

where dF denotes the distribution of unobserved heterogeneity in the health states and moral hazard and α_i is the marginal utility of income. The certainty equivalent utility U_{ijmt}^{CE} is discussed in Appendix C.

I define net welfare (NW) as the sum of consumer surplus (CS) and firm profits (Π) net of government spending on Traditional Medicare and Medicare Advantage (G_{TM} and G_{MA} respectively):

$$NW = CS + \Pi - (G_{TM} + G_{MA}) \quad (29)$$

VII.B Model assumptions

The results presented in this section inform how features of my model impact equilibrium outcomes. To do this, I simulate the effect of lowering payments made to Medicare Advantage plans, a common counterfactual exercise in the literature. I start by simulating the equilibria of the model under status quo policies. I then simulate the effects of a policy change for models with different levels of restrictiveness to see how the model’s predictions change. The full model allows firms to alter their entry and product offering decisions. The most restrictive model holds these decisions fixed at the values observed in the data. Intermediate versions of the model allow firms to either alter their entry or product offering decisions while holding the other margin fixed at the observed values.

Table 7 presents results for a simulation that lowers Medicare Advantage benchmarks by \$1,200 annually (\$100 per beneficiary-month). The first column reports the average predictions across all equilibria from a model with endogenous firm entry and product offerings under the current policy environment. The remaining columns predict how these quantities

will change once MA benchmarks are lowered by \$1,200 under different models. The key predictions that change across models relate to the likelihood plans enter and which markets they enter. For example, models that restrict product repositioning rule out the presence of the high generosity HMO plans in any equilibria. Similarly, models that restrict market entry miss how strategic plans exit markets where relatively sicker seniors reside. Each of the simulated models predict that MA plans increase their premiums in response to the benchmark reduction, which causes relatively sicker beneficiaries—who are more price sensitive—to dis-enroll from those plans. This leads each model to make similar predictions about the size of the MA market following the policy change. Consistent with prior work, I find that the increase in premiums induced by the benchmark cut leads to a large reduction in MA enrollment. The individuals that leave MA for TM are relatively healthier than the TM population and this sorting leads to TM becoming less adversely selected than it is under the baseline policy.

Figure 6 plots the average change in welfare and government spending due to the policy change predicted by each model. In general, firm profits and consumer surplus fall under this benchmark reduction policy. Government spending falls with the contraction of the MA market. The reduction in government spending is large enough to offset the declines in producer and consumer surplus such that net welfare increases. All of the models make similar predictions for firm profits, consistent with the decline in MA enrollment. There are notable differences in estimated changes in consumer surplus. Models that do not capture endogenous product variety will overstate the reductions in consumer surplus relative to those that capture this margin. An implication of this pattern is that consumers value product variety in these markets and models that fail to capture the product repositioning margin may under estimate the welfare implications of alternative policies.

Another observation is that entry is only marginally impacted by changes to the supply subsidy. This finding suggests that firm fixed costs are not substantial barriers for participation. This effect may be driven in part by the set of potential entrants containing firms

already in the health insurance market. It is possible that firms not already in the health insurance business may enter and have higher fixed costs. However, this is not typically observed and not a first-order concern. Moreover, the government is seeking to leverage private sector expertise in this market, so entry of a non-existing health insurance firm is unlikely to deliver these benefits for the government.

These simulations underscore the close connection between selection and consumer price sensitivity in my model. MA plans can attract more enrollment when they are able to operate profitably without charging consumers high premiums. These conditions are satisfied when the government heavily subsidizes plans or MA plans target markets where the plan can increase its subsidy payments. These channels are important for understanding the mechanisms behind the results in the next section that explore the effects of alternate subsidy systems.

VII.C Alternative subsidy policies

In this section, I use the model to explore the effects of alternative subsidy policies that could be used to regulate a competitive insurance market. These exercises assess the strength of the private market and whether these policies can improve its functioning in terms of lowering costs to the government. The latter is especially important since a key motivation for the governments use of a private market to deliver public benefits is to leverage the expertise of private firms to deliver these benefits more efficiently. The results are presented in Table 8. The first column simulates the effects of the current system where firms are directly compensated for each beneficiary they enroll and the size of the payment is adjusted by the beneficiary's risk score. These values are the same as what appeared in the previous table.

The first counterfactual I consider is a system where no subsidy is provided to firms. Under this scenario, the government allows private firms to offer insurance plans that meet quality thresholds but the plans receive no other assistance. I find that absent subsidies,

the private market unravels. This counterfactual predicts the lowest amount of government spending and consumer surplus. Relative to the average outcome under the status quo policy, total government spending falls by roughly 15% and consumer surplus falls 100%. This is driven by the normalization for the utility of enrolling in TM. These results illustrate the strength of the competition private firms face from TM. Despite offering more generous insurance plans than TM, absent a subsidy there is no premium private firms could charge that would cover their costs and attract any enrollment. Thus, the government must provide some support for this market to sustain its existence.

The second alternative system is presented in column 3 of Table 8. Under this policy, consumers receive an untargeted subsidy to purchase MA plans if they opt out of TM. The only revenue firms collect is in the form of premiums paid by beneficiaries that enroll in their plans. The size of the untargeted subsidy is equal to the average subsidy paid to firms observed in the data—roughly \$786 (\$9,430) per beneficiary-month (-year). Relative to the baseline scenario, this policy encourages more entry by the strategic plans and enrollment in MA. Average MA penetration grew from roughly 31% in the baseline scenario to about 47%. This growth in the market is driven by low generosity plans that have low costs of providing supplemental benefits. By changing the incidence of the subsidy to the demand side of the market, these plans become “free” for consumers. Relative to the baseline, this policy brings sicker seniors into the MA market, which remains more positively selected than TM.

These results indicate how Medicare Advantage is not fully capturing the expertise of private firms in controlling healthcare utilization. Under the baseline and untargeted subsidy policies, MA attracts healthier enrollees than TM. This sorting arises in part since healthier seniors are less sensitive to premiums. The healthier types then benefit from utilizing healthcare at lower costs than what they would pay under TM. Neither policy is wholly effective at offsetting consumer price sensitivity or attempting to target sicker populations to enroll in MA. However, each policy has distinct advantages—the supply subsidy gives firms a primary revenue source independent of premiums, while direct subsidies

give consumers a clear incentive to enroll in MA. These patterns motivate designing a subsidy system that gives sicker seniors an incentive to enroll in MA, while preserving the incentives of private firms to participate in the market.

The last policy I simulate aims to accomplish this objective. This targeted policy has two components. The first it preserves supply subsidies but lowers cost benchmarks by \$1,200. The second piece passes along some of these savings to consumers in the form of a means tested subsidy. Low-income seniors receive a \$600 payment if they enroll in a MA plan, while all other seniors get a \$300 payment for enrolling in MA. The motivation for this means testing is twofold. First, low-income seniors tends to have greater health needs and consume more healthcare. Thus, from the government's perspective is more advantageous to have them covered in the private market which delivers benefits more efficiently than the government. Second, as a group with greater health needs, they may also benefit from having access to more generous insurance than what they'd have access to under the public option. This policy produces entry probabilities that are similar to marginally higher than the baseline scenario. Consistent with benchmarks getting cut, plan premiums increase under this policy but these are partially offset by the demand subsidies. The overall size of the MA market declines slightly, which is driven by the premium increases not offset by the demand subsidy. The means testing also appears to effectively induce sicker seniors to enroll in MA. This sorting is noticeable from the increase in the average risk of seniors enrolled the low generosity strategic plans and the fact that this policy leads to MA being more adversely selected than TM. this is the only simulated policy that induces this sorting pattern and leads MA costing the government less money on a per-beneficiary basis.

Figure 7 evaluates the effects of the untargeted and targeted subsidy policies on government spending and welfare. The untargeted demand subsidy increases welfare and government spending. Relative to the status quo policy, consumer surplus grows by nearly 75%, consistent with the large cash transfer individuals receive from the government. Firm profits also grow under this policy, which is consistent with the growth of the MA market. The ex-

pansion of the MA market is associated with a roughly 3% increase in government spending, which offsets the welfare gains under the policy and leads to a decline in net welfare. The targeted policy tends to increase welfare and lower government spending. Firm profits grow by 12% under this system, which is attributable to larger risk adjustment payments from insuring sicker seniors. Since the increase in premiums are only partially offset by demand subsidies, this policy is associated with a modest decline in consumer surplus (approx. 15%). However, the reduction in consumer surplus is offset by the growth in firm profits and a 1% reduction in government spending to create a nearly 2% increase in net welfare. If this figure is applied to total government spending on the Medicare program, it amounts to roughly \$10 billion in savings for the government.

Finally, I explore the distributional consequences of these policies in Table 9. The table reports consumer surplus under each counterfactual policy by observable beneficiary characteristics. Under the current policy, the oldest seniors, low-income seniors, and individuals with risk scores in the first and third quartile have the average highest surplus. This pattern likely reflects the value MA plans can deliver for these groups. For individuals with moderate health needs, MA provides lower costs in the form of more generous cost sharing. For the sickest individuals—who tend to be older—MA plans limit their medical bills by offering out-of-pocket maximums that do not exist in TM. The untargeted demand subsidy raises welfare for everybody, consistent with the uniform cash transfer. However, the healthiest individuals benefit the most as these individuals tend to be the least price sensitive and shift from TM to MA first. Under the targeted policy, low-income beneficiaries benefit the most—consistent with means tested demand subsidy.

VIII Conclusion

This paper studies competition and participation in Medicare Advantage insurance markets. I develop and estimate an equilibrium model of health plan supply and demand that cap-

tures the feedback among government policy, firm entry and product offering decisions, and consumer sorting and utilization of health insurance plans. My model accounts for multiple equilibria that may arise in firm decisions about which markets to enter and products to offer. I then use this model to assess the optimality of different government policies to regulate choice and competition in these markets.

My findings indicate policymakers face tradeoffs when deciding how to use private markets to deliver a public good like health insurance benefits. Subsidies are necessary to sustain the private market. However, current subsidy policy is not letting the program deliver on its stated goals. Medicare Advantage plans tend to attract healthier individuals to enroll and the current risk adjustment structure distorts the government's costs of subsidizing the market. Supply subsidies give firms more flexibility to set lower premiums and attract relatively sicker individuals into MA. Moreover, they are unable to bring in the sickest TM beneficiaries into the program, which may reflect MA plans still charging non-zero prices at current subsidy levels or switching costs seniors face when attempting to move out of TM. An untargeted demand subsidy increases MA enrollment but this policy disproportionately benefits healthier seniors that are not price sensitive. I find that a targeted policy that leverages the benefits of supply and demand subsidies is capable of delivering similar firm participation and MA enrollment. However, this policy successfully attracts sicker seniors to the private Medicare market and lowers total government and government spending per beneficiary in MA. My estimates imply roughly \$10 billion in cost savings for the government when projected out nationally. Policymakers must weigh these considerations when deciding which framework enables this private market to deliver value for consumers and state policy objectives.

References

- Abaluck, Jason, Mauricio Caceres Bravo, Peter Hull, and Amanda Starc (May 2021). “Mortality Effects and Choice Across Private Health Insurance Plans”. *The Quarterly Journal of Economics* 136.3, pp. 1557–1610.
- Baum-Snow, Nathaniel and Justin Marion (2009). “The Effects of Low Income Housing Tax Credit Developments on Neighborhoods”. *Journal of Public Economics* 93.5-6, pp. 654–666.
- Berry, Steven T. (1992). “Estimation of a Model of Entry in the Airline Industry”. *Econometrica* 60.4, pp. 889–917.
- (1994). “Estimating Discrete-Choice Models of Product Differentiation”. *The RAND Journal of Economics* 25.2, pp. 242–262.
- Berry, Steven T., James Levinsohn, and Ariel Pakes (1995). “Automobile Prices in Market Equilibrium”. *Econometrica* 63.4, pp. 841–890.
- Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston (2014). “How Does Risk Selection Respond to Risk Adjustment? New Evidence from the Medicare Advantage Program”. *American Economic Review* 104.10, pp. 3335–64.
- Cabral, Marika, Michael Geruso, and Neale Mahoney (2018). “Do Larger Health Insurance Subsidies Benefit Patients or Producers? Evidence from Medicare Advantage”. *American Economic Review* 108.8, pp. 2048–87.
- Chernozhukov, Victor, Denis Chetverikov, and Kengo Kato (2019). “Inference on Causal and Structural Parameters Using Many Moment Inequalities”. *The Review of Economic Studies* 86.5, pp. 1867–1900.
- Ciliberto, Federico, Charles Murry, and Elie Tamer (2021). “Market Structure and Competition in Airline Markets”. *Journal of Political Economy* 129.11, pp. 2995–3038.
- Ciliberto, Federico and Elie Tamer (2009). “Market Structure and Multiple Equilibria in Airline Markets”. *Econometrica* 77.6, pp. 1791–1828.
- Cohen, Alma and Liran Einav (June 2007). “Estimating Risk Preferences from Deductible Choice”. *American Economic Review* 97.3, pp. 745–788.
- Conlon, Christopher and Jeff Gortmaker (Aug. 2023). *Incorporating Micro Data into Differentiated Products Demand Estimation with PyBLP*. Working Paper 20208. National Bureau of Economic Research.
- Cubanski, Juliette and Tricia Neuman (Jan. 2023). “What to Know about Medicare Spending and Financing”. *Kaiser Family Foundation*. Available at <https://www.kff.org/medicare/issue-brief/what-to-know-about-medicare-spending-and-financing/>.

- Curto, Vilsa, Liran Einav, Amy Finkelstein, Jonathan Levin, and Jay Bhattacharya (2019). “Healthcare Spending and Utilization in Public and Private Medicare”. *American Economic Journal: Applied Economics* 11.2, pp. 1–31.
- Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya (2021). “Can Health Insurance Competition Work? Evidence from Medicare Advantage”. *Journal of Political Economy* 129.2, pp. 570–606.
- Dickstein, Michael, Kate Ho, and Nathaniel Mark (2023). “Market Segmentation and Competition in Health Insurance”. *Journal of Political Economy* Forthcoming.
- Duggan, Mark, Amanda Starc, and Boris Vabson (2016). “Who Benefits When the Government Pays More? Pass-Through in the Medicare Advantage Program”. *Journal of Public Economics* 141, pp. 50–67.
- Einav, Liran, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen (2013). “Selection on Moral Hazard in Health Insurance”. *American Economic Review* 103.1, pp. 178–219.
- Einav, Liran, Amy Finkelstein, and Pietro Tebaldi (2019). *Market Design in Regulated Health Insurance Markets: Risk Adjustment vs. Subsidies*. Working Paper.
- Eizenberg, Alon (2014). “Upstream Innovation and Product Variety in the U.S. Home PC Market”. *Review of Economic Studies* 81.3, pp. 1003–1045.
- Fan, Ying (2013). “Ownership Consolidation and Product Characteristics: A Study of the U.S. Daily Newspaper Market”. *American Economic Review* 103.5, pp. 1598–1628.
- Fan, Ying and Chenyu Yang (2020). “Competition, Product Proliferation, and Welfare: A Study of the US Smartphone Market”. *American Economic Journal: Microeconomics* 12.2, pp. 99–134.
- (June 2022). *Merger, Product Variety and Firm Entry: the Retail Craft Beer Market in California*. Working Paper 30146. National Bureau of Economic Research.
- Geddes, Eilidh (2022). *The Effects of Price Regulation in Markets with Strategic Entry: Evidence from Health Insurance Markets*. Working Paper.
- Geruso, Michael and Timothy J. Layton (2020). “Upcoding: Evidence from Medicare on Squishy Risk Adjustment”. *Journal of Political Economy* 128.3, pp. 984–1026.
- Handel, Benjamin R. (Dec. 2013). “Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts”. *American Economic Review* 103.7, pp. 2643–82.
- Hausman, Jerry A. (1997). “Valuation of New Goods Under Perfect and Imperfect Competition”. *The Economics of New Goods*. Ed. by Timothy Bresnahan and Robert Gordon. University of Chicago Press, pp. 207–248.
- Ho, Kate and Robin S. Lee (2022). “Health Insurance Menu Design for Large Employers”. *The RAND Journal of Economics* Forthcoming.

- Hoxby, Caroline M. (2000). “Does Competition Among Public Schools Benefit Students and Taxpayers?” *American Economic Review* 90.5, pp. 1209–1238.
- Jung, Jeah, Caroline Carlin, Roger Feldman, and Linh Tran (2022). “Implementation of Resource Use Measures in Medicare Advantage”. *Health Services Research* 57.4, pp. 957–962.
- Kong, Edward, Timothy J. Layton, and Mark Shepard (2022). *Adverse Selection Pricing and Unraveling of Competition in Insurance Markets*. Working Paper.
- Lee, Robin S. and Ariel Pakes (2009). “Multiple Equilibria and Selection by Learning in an Applied Setting”. *Economic Letters* 104.1, pp. 13–16.
- Manning, Willard G., Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, and Arleen Leibowitz (1987). “Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment”. *The American Economic Review* 77.3, pp. 251–277.
- Marone, Victoria R. and Adrienne Sabety (2022). “When Should There Be Vertical Choice in Health Insurance Markets?” *American Economic Review* 112.1, pp. 304–42.
- Miller, Keaton, Amil Petrin, Robert Town, and Michael E. Chernow (2021). *Optimal Managed Competition Subsidies*. Working Paper.
- Neilson, Christopher (2021). *Targeted Vouchers, Competition Among Schools, and the Academic Achievement of Poor Students*. Working Paper.
- Newhouse, Joseph P., Mary Price, J. Michael McWilliams, John Hsu, and Thomas G. McGuire (2015). “How Much Favorable Selection Is Left in Medicare Advantage?” *American Journal of Health Economics* 1.1, pp. 1–26.
- Pakes, Ariel, Jack Porter, Kate Ho, and Joy Ishii (2015). “Moment Inequalities and Their Application”. *Econometrica* 83.1, pp. 315–334.
- Pelech, Daria M. (2020). “Prices for Physicians’ Services in Medicare Advantage and Commercial Plans”. *Medical Care Research and Review* 77.3, pp. 236–248.
- Polyakova, Maria and Stephen P. Ryan (Dec. 2020). *Subsidy Targeting with Market Power*. Working Paper 26367. National Bureau of Economic Research.
- Poterba, James M. (1996). “Government Intervention in the Markets for Education and Health Care: How and Why?” *Individual and Social Responsibility: Child Care, Education, Medical Care, and Long-term Care in America*. University of Chicago Press, pp. 277–308.
- Tebaldi, Pietro (Mar. 2022). *Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA*. Working Paper 29869. National Bureau of Economic Research.

- Trish, Erin, Paul Ginsburg, Laura Gascue, and Geoffrey Joyce (Sept. 2017). “Physician Reimbursement in Medicare Advantage Compared With Traditional Medicare and Commercial Health Insurance”. *JAMA Internal Medicine* 177.9, pp. 1287–1295.
- Varadhan, Ravi and Christophe Roland (2008). “Simple and Globally Convergent Methods for Accelerating the Convergence of Any EM Algorithm”. *Scandinavian Journal of Statistics* 35.2, pp. 335–353.
- Vatter, Benjamin (2022). *Quality Disclosure and Regulation: Scoring Design in Medicare Advantage*. Working Paper.
- Wagner, Myles (2022). *Health Insurance Menu Design with Multi-Dimensional Plans: Evidence from ACA Insurance Exchanges*. Working Paper.
- Wollmann, Thomas G. (2018). “Trucks Without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles”. *American Economic Review* 108.6, pp. 1364–1406.

Tables and Figures

Table 1: Market level summary statistics, 2017–2018

	Mean	SD	P10	P90
Monetary characteristics				
Premium	16.7	18.4	0.5	43.2
Supplemental premium	15.8	17.2	0.5	40.6
Base premium	0.9	2.7	0.0	2.3
Subsidy	742.4	47.5	688.5	798.2
Rebate	86.9	44.0	42.4	154.2
Benchmark	861.5	45.8	806.9	912.5
Average OOPC	132.2	22.8	101.8	155.0
Plan menus				
Firms	6.4	3.2	3.0	10.0
Plans	14.4	9.1	4.0	27.0
High generosity plans	6.0	6.5	0.0	15.0
HMOs	10.0	8.2	1.0	21.0
Market size	106,763.5	168,766.8	6,032.0	238,268.0
Plan enrollment	7,113.9	12,538.9	202.2	23,143.8
Market				
MA penetration	25.5	15.9	5.2	46.8
Market share	6.3	4.8	1.6	11.6
Market share MA	15.4	19.3	3.8	33.3
HHI	334.6	362.8	15.8	842.6
HHI MA	4,577.4	2,241.5	2,322.9	8,407.4

Notes: This table contain market-level summary statistics for the 4,241 MA markets in the analysis sample. Markets are defined as county-year pairs. Plan characteristics are weighted by within market enrollment. “Average OOPC” measures the average expected monthly out-of-pocket costs in a Medicare Advantage plan across health states. The “high generosity plans” earn this designation based on this cost measure.

Table 2: Benchmarks impact plan entry and choice set generosity, 2017–2018

	(1) Firms	(2) Plans	(3) Supplemental Permium	(4) Rebate	(5) Rebate Allocation Cost Sharing
Avg plan benchmark LOO	-0.0001 (0.0006)	0.0057*** (0.0020)	-0.0677*** (0.0153)	0.0523** (0.0203)	0.0618*** (0.0214)
County FEs	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓
Mean of Dep. Var.	2.85	5.91	15.37	59.54	30.56
R^2	0.96	0.98	0.92	0.93	0.92
Obs	3,958	3,958	3,958	3,958	3,958

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Standard errors clustered at the market-level. This table reports estimates from OLS regressions of the average plan benchmark onto market-level measures of plan generosity. An observation is a county-year. Monetary values are converted into 2008\$. Benchmarks for each plan are constructed by taking the enrollment weighted average across all markets where a plan is present, leaving out the focal market. These plan level benchmarks are then aggregated to the market level as an enrollment weighted average “Avg Benchmark (LOO).” “LOO” stands for leave-one-out. Outcomes are calculated as within market enrollment weighted averages.

Table 3: Elasticity of healthcare consumption estimates, 2017–2018

	(1)	(2)	(3)
	OLS	IV	IV
$\log(1 + \mathbb{E}[OOP_{mt}])$	-0.45*** (0.16)	-0.37** (0.18)	
$\log(1 + \mathbb{E}[OOP_{kmt}])$			-0.18* (0.10)
Year FE	✓	✓	✓
Contract FE	✓	✓	✓
Individual Controls	✓	✓	✓
Market Controls	✓	✓	✓
Plan Controls	✓	✓	✓
Observations	1,095,683	1,095,683	1,095,683
R^2	0.146	0.080	.

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Standard errors clustered at the market-level. This table reports OLS and IV estimates for the elasticity of healthcare consumption. Healthcare utilization is measured in standardized dollar units proposed by Jung et al. (2022). The estimation sample include individuals from the utilization sample that were not enrolled in TM or MA during the prior year. It contains individuals that enrolled in TM or a MA plan. The coefficients of interest are measures of plan choice set generosity at the market or market-contract level. Individual controls include their age and indicators for risk score quantiles, female, and low income status. Market controls include measures of how rural the county is, the share of the population with a college degree, the mortality rate of the Medicare population, and the share of the Medicare population that is eligible for Medicaid. Plan controls include indicators for star ratings (at half star intervals) and whether the plan is a HMO, Local PPO, or TM.

Table 4: Demand estimates and model implied quantities

			Estimate	95% CI	
Demand	Premium (α_i)	Mean	-6.403	[-6.772, -6.033]	
		Low income	-3.190	[-3.205, -3.176]	
	Utilization utility (β_i)	Mean	-1.144	[-1.147, -1.142]	
		Low income	0.342	[0.341, 0.343]	
	Supplemental revenue (γ)	Coefficient	1.175	[0.353, 1.996]	
	Nesting parameter (σ)	Coefficient	0.528	[0.527, 0.530]	
	Fixed effects (θ_2)	Contract		✓	
		Year		✓	
		Star rating		✓	
Quantities	Moral hazard (ω_i)	TM	0.49	[0.48, 0.49]	
	Pct. change in utilization from 100% to 0% coins.	PPO-Low	0.08	[0.08, 0.08]	
		PPO-High	0.05	[0.05, 0.05]	
		HMO-Low	0.08	[0.07, 0.08]	
		HMO-High	0.04	[0.04, 0.04]	
		Utilization costs (ϕ) (\$1,000)	TM	0.14	[0.14, 0.14]
		HMO	1.18	[1.18, 1.19]	
		PPO	0.88	[0.87, 0.88]	
	Switching costs (ι) (\$1,000)	Coefficient	-0.76	[-0.76, -0.76]	
		Mean	0.12	[0.11, 0.13]	
		Low income	0.08	[0.08, 0.08]	
	Risk aversion (ψ) (\$)	CARA coefficient ($\times 10^{-4}$)	3.77	[3.70, 3.84]	
		Cohen and Einav (2007) gamble	96.37	[96.34, 96.39]	
	Beneficiary-year observations			73,396,892	
	Plan-year observations			3,624	

Notes: This table reports estimates for demand parameters and quantities implied by the demand and healthcare utilization model. Estimates are obtained from a two-stage GMM procedure that targets observed utilization and plan choice decisions and IV restrictions. Confidence intervals are constructed from standard errors obtained from the variance-covariance matrix of the GMM estimator. Detailed parameter estimates and standard errors are available in Appendix Table E.6.

Table 5: Fixed cost identified set estimates

	Identified set
Number of plans	[612.7, 1,478.1]
Existing network	
Total hospital systems	[110.5, 252.0]
Total doctors	[1.5, 2.4]
No network	
Total hospital systems	[453.3, 981.6]
Total doctors	[4.1, 9.0]
Moment inequalities	124

Notes: This table reports the estimated identified set for the fixed cost parameters. Costs are reported in \$1,000 units. Sets are constructed by inverting the test statistics from Chernozhukov et al. (2019). The self-normalized one step procedure is used with $\alpha = 0.05$.

Table 6: Summary of Massachusetts Medicare Market, 2018

Firm	Offered plans	Markets	Market share	
			All	MA only
Tufts Health Plan	HMO (L-H)	8	3.40	32.99
Blue Cross-Blue Shield of Mass.	PPO (L-H), HMO-H	11	2.39	23.20
United Health	HMO (L-H)	7	2.04	19.75
Baystate Health	HMO (L-H)	4	1.03	9.99
Harvard Pilgrim	HMO (L-H)	7	0.69	6.66
Fallon Community	HMO (L-H)	4	0.66	6.40
Aetna	PPO-L, HMO-L	7	0.11	1.03
Medicare Advantage		12	10.30	–
Traditional Medicare		14	89.70	–
Total markets/beneficiaries		14	784,249	80,811

Notes: This table reports the observed market structures for Massachusetts in 2018.

Table 7: Impact of modeling choices on equilibrium outcomes

	Baseline	Cut benchmarks \$1,200			
	Entry + Repos	Entry + Repos	No Entry or Repos	Repos only	Entry only
Probability offered					
PPO-L	0.68	0.55	1.00	0.85	0.56
PPO-H	0.39	0.37	1.00	0.76	0.33
HMO-L	0.98	0.85	1.00	0.85	0.79
HMO-H	0.77	0.52	NA	0.43	NA
Benchmark (\$1,000)					
PPO-L	10.45	9.18	9.46	9.19	9.17
PPO-H	10.48	9.25	9.37	9.23	9.25
HMO-L	10.58	9.34	9.46	9.46	9.32
HMO-H	10.40	9.18	NA	9.46	NA
Premium (\$1,000)					
PPO-L	0.47	0.94	0.71	0.95	0.91
PPO-H	8.62	8.33	3.14	5.08	7.07
HMO-L	0.00	0.08	0.08	0.08	0.07
HMO-H	1.75	2.04	NA	1.99	NA
Average risk score					
PPO-L	0.80	0.75	0.77	0.76	0.75
PPO-H	0.85	0.81	0.75	0.78	0.80
HMO-L	0.94	0.89	0.89	0.88	0.90
HMO-H	0.95	0.91	NA	0.93	NA
Market size and selection					
MA share (%)	0.31	0.17	0.18	0.15	0.20
MA risk score	0.91	0.86	0.88	0.87	0.87
TM risk score	1.13	1.10	1.10	1.09	1.11
Per-bene risk adjusted government cost (\$1,000)					
MA	10.19	9.02	9.06	9.10	9.00
TM	9.84	9.43	9.47	9.40	9.48

Notes: This table reports how simulated quantities from the model change as different features are added. Each column reports the average value across all equilibria of the model. Appendix Table E.8 reports the range of quantities across equilibria. The top panel reports quantities for the strategic plans. The bottom panel reports service area quantities. The first two columns report predicted outcomes from a model where firms are allowed to alter their market entry and product offering decisions (“Entry + Repos”). The third column restricts firms to the market entry and product offering decisions observed in the data (“No Entry or Repos”). The fourth column allows the strategic firms to alter their product offerings in the markets where these firms are observed to have entered in the data (“Repos only”). The final column allows the strategic firms to alter the markets where they enter but they can only offer products that are observed in the data (“Entry only”).

Table 8: Equilibrium outcomes under alternative subsidy systems

	Baseline	No subsidy	Untargeted subsidy	Targeted subsidy
Probability offered				
PPO-L	0.68	0	0.97	0.67
PPO-H	0.39	0	0.67	0.43
HMO-L	0.98	0	0.99	0.98
HMO-H	0.77	0	0.59	0.80
Benchmark (\$1,000)				
PPO-L	10.45	NA	10.49	9.25
PPO-H	10.47	NA	10.44	9.26
HMO-L	10.58	NA	10.57	9.37
HMO-H	10.40	NA	10.40	9.18
Premium (\$1,000)				
PPO-L	0.46	NA	9.75	0.83
PPO-H	8.62	NA	16.65	8.75
HMO-L	0.00	NA	9.32	0.38
HMO-H	1.75	NA	15.82	3.96
Average risk score				
PPO-L	0.80	NA	0.84	0.91
PPO-H	0.85	NA	0.87	0.81
HMO-L	0.94	NA	0.99	1.13
HMO-H	0.95	NA	0.99	0.95
Market size and selection				
MA share (%)	0.31	0.00	0.47	0.27
MA risk score	0.91	NA	0.95	1.10
TM risk score	1.13	1.07	1.16	1.05
Per-bene risk adjusted government cost (\$1,000)				
MA	10.19	0.00	9.90	9.45
TM	9.84	9.01	10.47	9.91

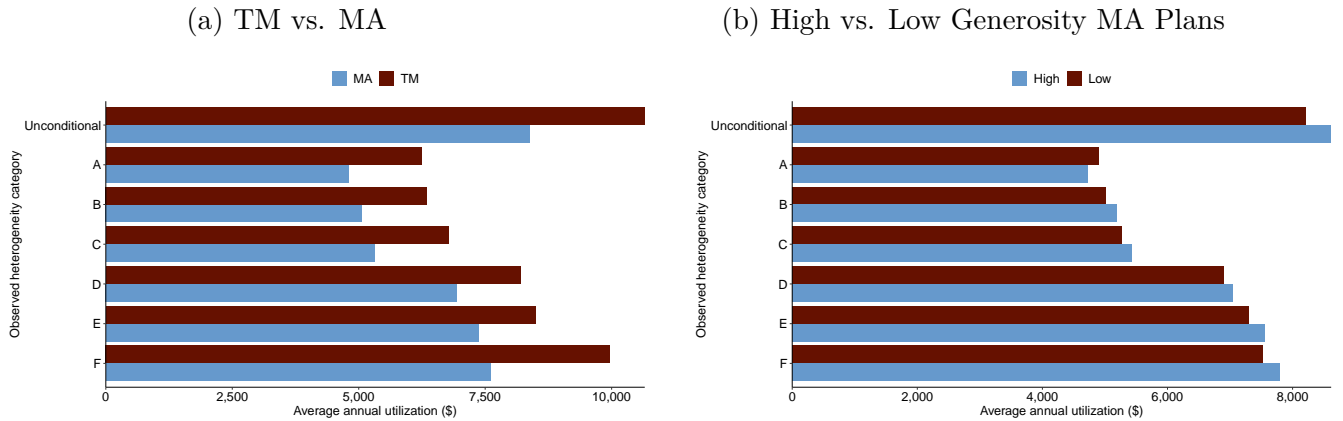
Notes: This table reports how simulated equilibrium outcomes change as the delivery system for Medicare Advantage subsidies changes. Each column reports the average value across all equilibria of the model. Appendix Table E.9 reports the range of quantities across equilibria. The top panel reports quantities for the strategic plans. The bottom panel reports service area quantities. “Baseline” refers to the current system, which is a supply side subsidy that is scaled by a beneficiary’s risk score. “No subsidy” refers to a system where the government does not provide any subsidy to Medicare Advantage plans but regulates the plans meet their minimal coverage standards. “Untargeted subsidy” simulates a system that gives the observed enrollment weighted average risk adjusted pre-beneficiary subsidy for Massachusetts (approximately \$9,432 per year) to consumers to offset the costs of a Medicare Advantage plan. “Targeted subsidy” cuts CMS benchmarks by \$1,200 and offers a demand subsidy of \$600 to low income beneficiaries that enroll in Medicare Advantage plans and \$300 for non-low income MA enrollees.

Table 9: Consumer surplus by observables under alternative subsidy systems

	Baseline	Untargeted subsidy	Targeted subsidy
Age ≥ 86	0.35	0.61	0.29
Female	0.11	0.20	0.10
Low income	0.47	0.82	0.40
High Medicaid county	0.08	0.14	0.07
Risk score Q1	0.29	0.51	0.25
Risk score Q2	0.25	0.44	0.21
Risk score Q3	0.27	0.47	0.23
Risk score Q4	0.24	0.41	0.20

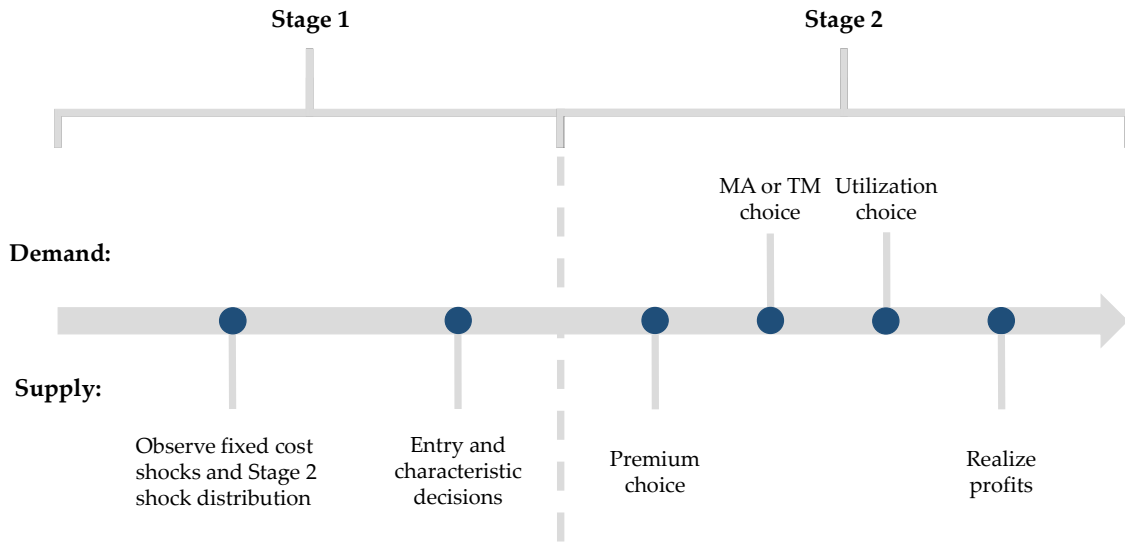
Notes: This table reports how consumer surplus is distributed across observable characteristics under different subsidy policies. Each column reports the average value across all equilibria of the model. Appendix Table E.10 reports the range of quantities across equilibria. “Baseline” refers to the current system, which is a supply side subsidy that is scaled by a beneficiary’s risk score. “Untargeted” simulates a system that gives the observed enrollment weighted average risk adjusted pre-beneficiary subsidy for Massachusetts (approximately \$9,432 per year) to consumers to offset the costs of a Medicare Advantage plan. “Targeted” cuts CMS cost benchmarks by \$1,200 and offers a demand subsidy of \$600 to low income beneficiaries that enroll in Medicare Advantage plans and \$300 for non-low income MA enrollees.

Figure 1: Average Healthcare Utilization, 2017–2018




Notes: This figure compares the average annual healthcare utilization of Medicare beneficiaries. The averages are presented unconditionally and for the six most common groupings observable heterogeneity. Observable categories summarize a beneficiary’s risk score, age, gender, income, and their county’s Medicare mortality and Medicaid eligibility rates. These groupings are constructed by converting risk scores into quantiles and defining all possible combinations of these characteristics.

Figure 2: Model summary




Notes: This figure summarizes the timing and decisions made in the model. Firm decisions are below the central line and correspond to the supply side of the model. Beneficiary decisions are above the central line and correspond to the demand side of the model.

Figure 3: Average market risk scores 2017–2018




Notes: This figure compares the distribution of risk scores at the market-level. Markets are defined as a county-year pair. The red distribution is the true risk score reported by CMS. The blue distribution comes from the risk scores that I calculate using only inpatient diagnoses. These individual risk scores are averaged across all beneficiaries in the market to construct the distribution.

Figure 4: Targeted moment fit




Notes: This figure plots a subset of the targeted moments used to estimate the health state distribution and demand parameters. The targeted moments included in the figure are the unconditional mean and variance of the utilization distribution and the mean and variance of the utilization distribution conditional on risk score quartiles.

Figure 5: Untargeted moment fit



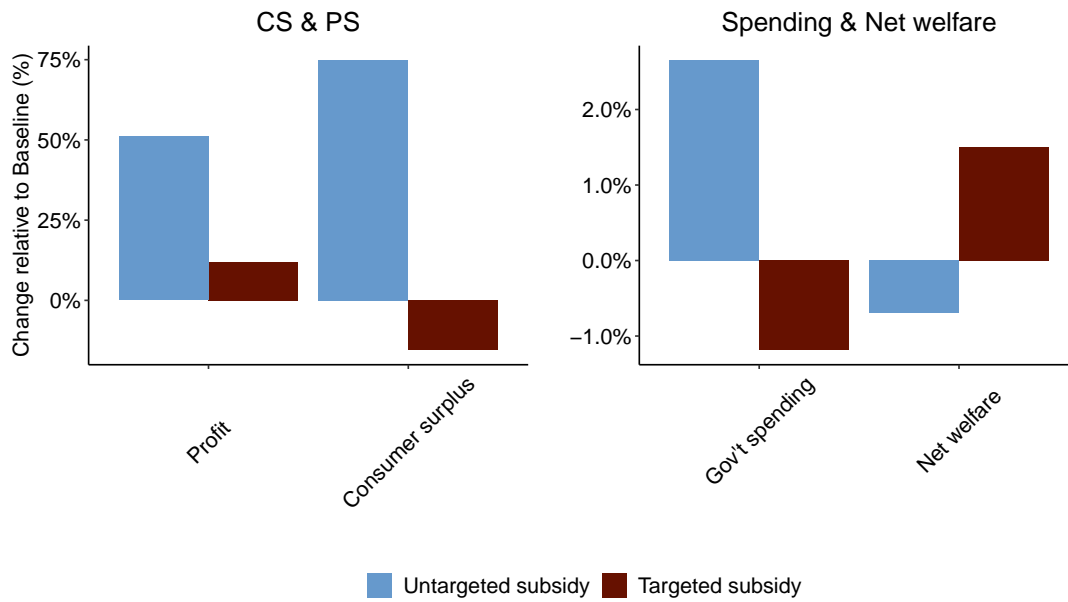
Notes: This figure plots untargeted data moments and their model analogs. The moments included in the figure are the mean and variance of healthcare utilization conditional on plan type.

Figure 6: Welfare changes from \$1,200 benchmark cut by modeling choices



Notes: This figure plots the average changes in welfare and spending metrics across model equilibria stemming from a \$1,200 benchmark cut. Each bar represents a different modeling assumption. “Entry + Repos” is a model where strategic firms can alter their market entry and product offering decisions. “No Entry or Repos” is a model that restricts the strategic firms to their observed market entry and product offerings. “Repos only” is a model that restricts the strategic firms to their market entry choices but allows them to alter their product offerings in those markets. “Entry only” is a model that restricts strategic firms to their observed product offerings but allows them to alter which markets they enter. The left panel reports firm profits and consumer surplus. The right panel reports total government spending on TM and MA and net welfare, which is the sum of consumer surplus and firm profits net of government spending.

Figure 7: Welfare changes from different subsidy systems



Notes: This figure plots the average changes in welfare and spending metrics across model equilibria stemming from two alternative subsidy systems. ‘Untargeted subsidy’ simulates a system that gives the observed enrollment weighted average risk adjusted pre-beneficiary subsidy for Massachusetts (approximately \$9,432 per year) to consumers to offset the costs of a Medicare Advantage plan. ‘Targeted subsidy’ cuts CMS benchmarks by \$1,200 and offers a demand subsidy of \$600 to low income beneficiaries that enroll in Medicare Advantage plans and \$300 for non-low income MA enrollees. The left panel reports firm profits and consumer surplus. The right panel reports total government spending on TM and MA and net welfare, which is the sum of consumer surplus and firm profits net of government spending.

SUPPLEMENTAL APPENDIX

Entry and Competition in Insurance Markets: Evidence from Medicare Advantage

Matthew V. Zahn*

January 1, 2025

A Data and Sample Construction

In this section, I provide detailed descriptions of the data sets I use in my analysis and how the analysis samples for the demand and utilization estimates are formed.

A.1 Data Sources

My analysis relies on 11 data sources. A description of each data source and how it is used within my analysis appear below.

Medicare Beneficiary Summary File. This data set contains individual level information on all beneficiaries in the Medicare program. I observe the beneficiary’s demographics such as age, sex, dual eligible status, reason for Medicare eligibility, and date of death. I can also track the beneficiary’s county of residence in each month they were enrolled in Medicare. I also observe how the beneficiary opted to receive Medicare benefits (i.e., through Traditional Medicare or Medicare Advantage). If the beneficiary enrolled in Medicare Advantage, I observe the contract and plan identifiers for their chosen plan. I can also observe information about Medicare Part D plans but I do not use this information as part of my main analysis. I have access to these data from 2014–2019. The Beneficiary Summary File is used to construct market shares and demographics, as well as provide the observable characteristics of individuals in the demand and utilization models.

This data set also contains aggregate information about healthcare utilization and spending by category (e.g., inpatient, outpatient, etc.) for Traditional Medicare beneficiaries. I opt not to use this information because I am unable to construct the standardized utilization metric for this roll up of each beneficiaries claims. As a result, I would not have a consistent utilization metric for Traditional Medicare and Medicare Advantage beneficiaries. I do use this information to inform my calibration of the cost structure for Traditional Medicare that

*Hopkins Business of Health Initiative, Johns Hopkins University; matthew.zahn@jhu.edu.

appears in the utilization model.

Traditional Medicare Claims. This data set contains information about the utilization of healthcare among Traditional Medicare beneficiaries at the claim level. I have access to TM claims and discharges for inpatient, outpatient, carrier, hospice, and Part D services with differing levels of coverage. I observe 100% of inpatient and hospice claims as well as inpatient, outpatient, carrier, and Part D claims for a 20% random sample of TM beneficiaries each year. I have access to these data from 2014–2019. My analysis focuses on inpatient, outpatient, and carrier claims. These claims data are used for three purposes. First, I use them to recover diagnoses for the risk score calculation. Second, they are used to construct the standardized price measure developed by Jung et al. (2022). Third, I use them as part of the utilization moments to estimate the parameters of the health state distribution and hassle costs of healthcare utilization.

Medicare Advantage Encounter Data. This data set contains information about the utilization of healthcare among Medicare Advantage beneficiaries at the encounter level. Unlike traditional claims data sets, the encounter data contain no payment information but do contain most other fields found in these sources. I have access to MA encounter data for inpatient (hospitals and SNFs), outpatient, carrier, hospice, and Part D services with differing levels of coverage. I observe 100% of inpatient, outpatient, and hospice encounters; all encounters for a cohort of 12 million MA beneficiaries (roughly 50–60% of the entire MA population depending on the year) which covers roughly 52% of MA beneficiaries in my analysis sample; and 20% of Part D encounters. I have access to these data from 2016–2018. The encounter data are used for four purposes. First, I use them to recover diagnoses for the risk score calculation. Second, I apply the standardized price measure developed by Jung et al. (2022), which I discuss in more detail below. Third, I use them as part of the utilization moments to estimate the parameters of the health state distribution and hassle costs of healthcare utilization. Fourth, I use average plan level utilization to recover the plan’s negotiated prices along with the inversion of the plan’s first order condition.

Medicare Advantage Bid Templates. This data set contains the information MA plans provide to CMS as part of the regulatory process that determines their subsidy and rebate payments. I have access to these submissions for every MA plan from 2006–2018 and they are publicly available on the CMS website. From this data source I recover the subsidy amount the plan requested, the size of its rebate payment, how its rebate was allocated, and the amount of revenue the plan needs to fund extra benefits relative to Traditional Medicare.

They also report how plan’s premium is broken out between the base and supplemental premium. The bid templates also detail the numerical values of the cost sharing characteristics of the plan as well as their projected allowed amounts for medical claims. These data are used in three places within my analysis. First, I rely on them as part of inverting the plan’s first order condition to recover the plan’s negotiated prices. Second, I use them when estimating the size of a plan’s supplemental premium. These data are also used to inform my calibration of the plan out-of-pocket cost functions that are used in the utilization model.

Medicare Advantage Enrollment. This data set tracks monthly county-level enrollment for all Medicare Advantage plans. The data also contain information about plan characteristics including network type and whether the plan is a special needs plan. I have access to these data from 2006–2019 and they are publicly available on the CMS website. These data provide characteristics of Medicare Advantage plans that appear as part of the demand, utilization, and fixed cost models and are used to determine the analysis sample.

Plan Benefit Packages. This data set tracks characteristics for Medicare Advantage plans. The tracked characteristics include the plan’s premium, the counties included in the plan’s footprint, and how the counties within a plan’s footprint map to segment identifiers specific to the plan. I have access to these data from 2006–2019 and they are publicly available on the CMS website. These data provide characteristics of Medicare Advantage plans that appear as part of the demand model and are used to determine the analysis sample.

Out-of-Pocket Cost Estimates. This data set provides estimates for a beneficiary’s expected out-of-pocket costs in Medicare Advantage plans and Traditional Medicare. These estimates are produced annually for every MA plan and TM and are typically featured on the Medicare plan finder application. The estimates are available for discrete health statuses ranging from “Poor” to “Excellent.” The estimates are generated from a CMS developed model that takes the characteristics of MA plans, behavioral assumptions about how care is received (i.e., in-network), and utilization patterns from TM data for the plan’s enrollee population. Cost estimates are produced for specific services (e.g., inpatient hospital acute care, eye exams, hearing exams, etc.) and may be aggregated up accordingly. I have access to these data from 2007–2020. I obtained these materials through a Freedom of Information Act request and direct correspondence with CMS staff. These data provide characteristics of MA plans and TM that are relevant for the utilization and demand models as well as estimating the size of a plan’s supplemental premium.

Plan Ratings. This data set provides the star ratings used to denote the quality of a MA plan. I have access to these data from 2007–2020 and they are publicly available on the CMS website. These data provide characteristics of MA plans that are relevant for the demand model and estimating the size of a plan’s supplemental premium.

Plan Payments and Ratebooks. These data sets contains information on plan level payments, rebates, and risk scores as well as the benchmarks set by CMS. I have access to these materials from 2006–2019 and they are publicly available on the CMS website. These data are primarily used when solving the model for counterfactual entry patterns and assessing the validity of the risk scores I calculate.

Medicare Geographic Variation. These data contain information on the Medicare program and its beneficiaries at the county-level. I have access to these materials from 2007–2019 and they are publicly available on the CMS website. These data are primarily used as a diagnostic to test the validity of the risk scores I calculate.

DRG InterStudy. This data set contains estimated enrollment for all insurance companies at the county level. The enrollment estimates are broken out by insurance product type (i.e., commercial-HMO, commercial-PPO, Medicare Advantage, Medicaid managed care, etc.). I have access to these materials for 2015, 2017, and 2019. These data are used to estimate the identified set of parameters in firm fixed costs.

AHA Annual Survey and Area Health Resources Files. These data sets contains information about the number of providers (e.g., hospitals, hospital systems, doctors, etc.) and utilization of healthcare services at the county level. These data are available with different time coverage but cover the period from 2007–2018. The Area Health Resource Files are publicly available on the Health Resource Service Administration. These data are used to estimate the identified set of parameters in the firm fixed costs. I obtained the AHA data from the Wharton Research Data Services.

American Community Survey. This data set contains demographic information at the county level. Specially, I use these data to measure mean and median income, household size, educational attainment, and what percentage of a county is rural. These data are publicly available on the Census website. These materials are used within the demand model.

A.2 Demand Sample

The sample used to estimate the demand model combines most of the data sets described in the previous section. The main file is the Medicare Beneficiary Summary File, which is then supplemented with data sets that contain the characteristics of Medicare Advantage plans and local markets. The end result is a panel of Medicare beneficiaries from 2017–2018. The sample also relies on information from the 2016. The sample restrictions based on individual characteristics are detailed below.

1. Individuals that do not qualify for Medicare because of their age. This condition means that beneficiaries that were not 65 by end of the sample year or were eligible for Medicare due to disability status or having End Stage Renal Disease are dropped.
2. Individuals that were enrolled in Medicare Part A for a different number of months within a year than they were enrolled in Medicare Part B. This pattern primarily arises because enrollment in Medicare Part A is automatic while beneficiaries must opt into Part B. A beneficiary may delay enrolling in Part B if they are still working and have employer sponsored coverage.
3. The beneficiary resides in Alaska, Guam, Puerto Rico, or the Virgin Islands. The Medicare program has idiosyncratic differences in these geographies.
4. The beneficiary has an invalid or missing geographic identifier.
5. The beneficiary is missing data needed to calculate their risk score.

I further restrict the sample based on Medicare Advantage enrollees and plans.

1. The beneficiary is enrolled in a MA plan with missing characteristic information (i.e., bids, out-of-pocket costs, payments, etc.).
2. The beneficiary is enrolled in an employer sponsored, special needs, or Part B only MA plan.
3. The beneficiary is enrolled in a plan outside of the plan's official footprint. This pattern can occur if an individual previously resided in a plan's footprint but relocated to a new geography and retained their MA plan.

4. The individual is enrolled in a plan type other than a HMO or Local PPO. Other types of MA plans in the data include Private Fee-For-Service (PFFS) or Regional PPOs, which either have different subsidy regulations, small enrollment, or distinct cost structures. HMOs and Local PPOs enroll the vast majority of MA beneficiaries.

The net result of these restrictions is a sample that contains 73,941,784 beneficiary-year observations and 40,141,182 unique beneficiaries. The sample contains 3,702 plan-year observations of 2,263 unique MA plans. See Appendix Table E.1 for a detailed breakdown of the number of observations that dropped due to each sample restriction.

A.3 Utilization Sample

This section describes the utilization sample. This discussion includes how I construct the utilization metric applied to the Medicare Advantage encounter data and check them for data completeness. I conclude by describing precisely how the encounter data are used to estimate the model.

Utilization Measure Construction. I implement the algorithm proposed by Jung et al. (2022) to generate the standardized price utilization metric. At a high level this procedure generates these standardized prices based on Traditional Medicare claims data by netting out price differences attributable to geographic variation and applies them to services that appear in the Medicare Advantage encounter data. As part of their publication, the authors provide SAS code and an implementation guide that other users can modify to implement the algorithm based on the data they have available from CMS. I make two adjustments to the procedure proposed by Jung et al. (2022). First, I define the MA cohort to include all beneficiaries. Second, I use data from all available Traditional Medicare beneficiaries to construct the standardized prices. In both instances the written procedure used randomly drawn sub-samples to ease computation burdens. I relax these requirements to make use of all available data resources.

Data Completeness. The implementation in Jung et al. (2022) provides methods to assess the completeness of the Medicare Advantage encounter data. The first compares the number of hospitalizations that appear in the inpatient encounter files against those that appear in the MedPAR files. The second compares the number of emergency department and ambulatory care visits that appear in the encounter outpatient and carrier files against information that appears in the Healthcare Effectiveness Data Information System (HEDIS).

I consider a Medicare Advantage contract to have a high degree of data completeness if it has at least 2,500 enrollees, the difference between the number of hospitalizations in the encounter and MedPAR data is less than 10%, and the number of ambulatory or ED visits in the encounter and HEDIS data are within 20%.

The contacts that I identify as having a high degree of completeness overlaps with the list reported in Jung et al. (2022). I have fewer contracts than they do because I only have access to a cohort of the carrier encounter data. Thus, the utilization sample is composed of Traditional Medicare beneficiaries included in the 20% random sample defined by CMS and all Medicare Advantage beneficiaries enrolled in a plan associated with a contract that has a high level of data completeness. Beneficiaries in the random sample or a MA plan with high data completeness that are not observed in the claims or encounter data are assumed to have utilized no healthcare in that year.

Use in Estimation. The utilization sample is used to define the moments to target the parameters of the health state distribution and plan effects on individual utilization patterns. The model predicted utilization is also used to quantify the marginal costs of plans. This modeling choice is supported by evidence that documents Medicare Advantage plans paying similar prices as Traditional Medicare. Since utilization is measured in terms of standardized Traditional Medicare dollars, the model predicted utilization for a beneficiary also represents their marginal costs. I rely on these estimates when deriving the moment inequalities to recover the identified set of firm fixed costs.

A.4 Risk Score Calculation

CMS calculates risk scores for each beneficiary in the Medicare program. The general formula used in this calculation has three components and is reproduced below.

$$r_{it} = \underbrace{[R_{it}(\text{HCC Model}_t)]}_{\text{Base score}} / \underbrace{NF_t}_{\text{TM normalization}} \cdot 1\{\text{MA bene } t - 1\} \underbrace{CPA_t}_{\text{Coding pattern adjustment}} \quad (\text{A.30})$$

The first component is the base score, which is the output of the HCC models developed, maintained, and updated by CMS. Each version of the HCC model is publicly available on the CMS website. The HCC model takes a beneficiary’s demographics (i.e., age, sex, Medicare eligibility, Medicaid eligibility, etc.) and diagnoses from the prior year as inputs. The diagnoses must be recorded from inpatient or outpatient hospital visits, physicians, or

clinically trained non-physicians (e.g., psychologist, podiatrist). The HCC models return different base scores for different types of beneficiaries (e.g., new beneficiaries, dual eligibles, etc.).

The remaining parts of the formula modify the base score. The second component is a normalization factor. This adjustment is defined based on the costs and diagnoses of the Traditional Medicare population for a rolling reference period. The factor is calculated such that after it is applied to the base score, the average Traditional Medicare will have a risk score equal to one. The final component is a coding pattern adjustment that is intended to correct for “upcoding” among Medicare Advantage plans. The normalization factors and coding pattern adjustments used by CMS are published as part of their ordinary course.

As discussed in the main text, these risk scores are generally not made available in the data sets usable for researchers. I approximate the CMS risk scores with the data available to me based on Equation (A.30). To calculate the base scores, I gather diagnoses from the TM claims and MA encounter data for the years 2016–2018.¹ I then feed these into the HCC models for the years in my analysis sample along with the beneficiary demographics from the Medicare Beneficiary Summary File. I define the average TM base score within each sample year as the formalization factor. After applying the normalization factors to the base scores, I apply the reported coding pattern adjustments to Medicare Advantage beneficiaries. I compute two versions of these risk scores: one that uses only inpatient diagnoses (which I have for all beneficiaries) and another that uses inpatient, outpatient, and carrier diagnoses in the data available to me.

B Additional Descriptive Analyses

B.1 Quantifying the impact of the ACA on benchmarks

The ACA directly altered both payments to Medicare Advantage plans. These changes were motivated in part to address concerns about over-payments to plans participating in the program. The ACA aimed to lower subsidy payments by lowering the TM cost benchmarks to better align them with realized TM costs and limit how they could grow over time. The ACA sought to lower rebate payments by reducing the allowable fraction of the difference between the subsidy and the cost benchmark.² In general, these reforms were successful in

¹I exclude MA diagnoses generated from chart reviews.

²Despite aiming to reduce these payments, the ACA also introduced quality adjustments that increased benchmarks and allowable rebate fractions for plans with higher star ratings. These limited the size of the payment reductions for these types of plans.

lowering the payments MA plans received, which I quantify later in this section.

The ACA reforms to benchmarks were phased in from 2012–2016. Counties were given 2, 4, or 6 year transitions based on how far their current benchmarks were from the targets mandated by the ACA.³ Following the transition, a county’s cost benchmark was equal to the government’s projected TM costs for the county in the prior year. These projections were binned into quartiles and scaled by an adjustment factor.⁴ The benchmarks plans face are enrollment weighted averages of each market specific benchmark across its footprint. These footprints are plan specific and do not perfectly overlap with the footprints of rival plans. Thus the variation in the benchmarks for other markets can act as a source of plausibly exogenous variation in plan subsidy and rebate payments that induces plans to offer more or less financial coverage in a particular market.

To quantify the size of these reductions I estimate models of the following form:

$$B_{mt} = \alpha_0 + \alpha_1 \text{Post-ACA}_{mt} + \alpha_2 X_{mt} + \epsilon_{mt} \quad (\text{B.31})$$

where B_{mt} is the benchmark for county m in year t , Post-ACA_{mt} is an indicator for whether county m completed its transition ACA benchmarks, X_{mt} is a vector of characteristics for county m in year t and ϵ_{mt} is a county-year unobserved characteristics. I exclude observations for county-years that were mid-transition. Benchmarks are inflation adjusted to 2008 dollars.

My estimates of Equation (B.31) are presented in Table E.11. I consider three different measures of county benchmarks: the average benchmark weighted by observed plan enrollment, the non-quality adjusted benchmark, and the quality adjusted benchmark. For each benchmark I estimated two versions of Equation (B.31): one that used only county and year fixed effects in X_{mt} and another that used county characteristics that are likely correlated with healthcare utilization.⁵ Estimates from both specifications were similar.

Consistent with the legislation’s objectives my estimates indicate Medicare Advantage benchmarks fell significantly following the ACA’s implementation. The ACA lowered average benchmarks by approximately \$43–58 per-beneficiary-month or \$516–696 per-beneficiary-year. The benchmark reductions are much larger for plans that did not receive quality bonuses (approximately \$59–71 per-beneficiary-month or \$696–856 per-beneficiary-year) and

³Changes to the rebate fraction were phased in from 2012–2013 and did not vary by county or plan type.

⁴Counties in the first quartile (lowest projected TM costs) were adjusted up by 15% and 7.5% in the second. Third quartile counties received no adjustment and the fourth quartile (highest projected TM costs) was adjusted down by 5%.

⁵These characteristics include the share of the county classified as rural, the share of the county population with a college degree, median income in the county, the average age of the Medicare population in the county, the share of the population in the county that is white, the share of the county that is eligible for Medicaid, and realized per-capita TM costs.

much lower for plans that do receive quality bonuses (\$32–48 per-beneficiary-month or \$384–576 per-beneficiary-year). These estimates quantify the variation in MA benchmarks the ACA introduced. In the main text I use this variation to quantify the connection between these benchmarks and plan characteristics. While the equilibrium model cannot use this same ACA shock to benchmark due to data limitations, variation in benchmarks over time and across markets is used to separately identify private health information from moral hazard. This variation is also demonstrated in the main text of the paper.

B.2 ACA event studies

In the main text I describe how the ACA lowered cost benchmarks for Medicare Advantage plans. Counties were given 2, 4, or 6 years to transition to the new ACA benchmarks based on how far their current benchmarks were from the ACA targets. The phase in occurred from 2012–2016. I estimate a series of event studies that capture the impact of these changes to plan benchmarks altered Medicare Advantage market outcomes. The estimated event studies are of the following form:

$$Y_{mt} = \beta_m + \beta_t + \beta_0 B_{mt} + \sum_{i \in \{[2008, 2017] \setminus 2011\}} \beta_i 1[t = i] B_{mt} + \epsilon_{mt} \quad (\text{B.32})$$

Figures E.2–E.5 present the event study plots for models estimated by transition groups. Pooled estimates are similar and available upon request. The plots indicate that the benchmark reductions led to significant reductions in the payments to HMO and PPO plans. This was driven by a significant decrease in the subsidies plans requested following the benchmark reductions. These declines in requested subsidies were associated with significant increases in MA plan premiums in some markets relative to before the year before the benchmark reductions came into effect. Despite these cuts in benchmarks, MA plans tended to increase the amount of extra coverage they offer relative to TM in each type of transition county. The reductions in benchmarks led to modest but statistically significant reductions in the number of MA firms, contracts, and plans offered in most markets. Finally, Medicare Advantage enrollment increased despite the benchmark reductions. This growth is most pronounced in 2 year transition counties.

C Model and Estimation

In this section I provide additional details about components of the model and its estimation that are not covered in the main text.

C.1 Healthcare Utilization

Plan Cost Structures and Utilization Solution. The amount of healthcare agents choose to utilize in my model depends on the out-of-pocket costs associated with that level of utilization in their chosen health plan. While the insurance products examined in this paper are complex and have many idiosyncrasies, I make two simplifying assumptions to preserve tractability. First, I assume that the amount of money a beneficiary in a MA plan or TM can be expressed as a function of the amount of healthcare the beneficiary chooses to consume Q and (at most) three characteristics of the insurance contract: a deductible D , a coinsurance rate C , and an out-of-pocket maximum M . Second, I assume that there are only four out-of-pocket cost structures for Medicare Advantage plans—one for each network type and financial generosity category. I calibrate the cost structures for each Medicare Advantage plan and Traditional Medicare. The calibration for Medicare Advantage plans is informed by information included in the plan’s bid template that is submitted to CMS. Among the information included in these materials are estimates for the dollar value of total cost sharing and allowed amounts for each beneficiary the plan enrolls. I take the ratio of these values to generate a pseudo-coinsurance rate for the plan. These templates also report the plan’s out-of-pocket maxima and deductibles. The calibration for Traditional Medicare is informed by statues.⁶ Table C.1 reports the calibrated cost functions as well as the analytical expression for the optimal amount of healthcare to consume within each plan.

These calibrations align with stylized facts about Traditional Medicare and Medicare Advantage plans. In general, Traditional Medicare tends to have higher costs because of coverage gaps and no out-of-pocket maximum. This pattern is what drives many Traditional Medicare beneficiaries to supplement their coverage with additional insurance policies like Medigap. Part of Medicare Advantage’s value proposition is that it tends to have lower out-of-pocket costs relative to Traditional Medicare because it fills those coverage gaps. HMOs tend to have lower costs relative to PPOs, which is reflected in the calibration. However,

⁶For 2017–2018 the TM deductible for outpatient care was \$183 and a 20% coinsurance. For inpatient care, TM charges a per-hospitalization deductible which was approximately \$1,300 dollars for 2017 and 2018. An examination of the Cost and Use component of the Medicare Beneficiary Summary File for Traditional Medicare beneficiaries during this time period indicates that the average TM beneficiary that utilized inpatient care paid about this amount out-of-pocket.

Table C.1: Calibrated out-of-pocket cost functions and predicted healthcare utilization

	HMO		Local PPO		TM
	High	Low	High	Low	
Deductible D	\$0	\$0	\$200	\$300	\$1,500
Coinsurance C	6%	10%	8%	10%	20%
Out-of-pocket maximum M	\$3,500	\$6,000	\$5,000	\$7,000	NA
$Q^* > 0$	$h > \bar{h}$				
$Q^* = h$	NA	$h \leq \min\{\bar{h}_1, \bar{h}_2\}$		$h \leq \bar{h}_1$	
$Q^* = h(1 + \omega(1 - C))$	$h \leq \bar{h}_2$	$h \in (\bar{h}_1, \bar{h}_2) \ \& \ \bar{h}_1 < \bar{h}_2$		$h > \bar{h}_1$	
$Q^* = h(1 + \omega)$	$h > \bar{h}_2$	$h \geq \max\{\bar{h}_1, \bar{h}_2\}$		NA	
\bar{h}_1	$2D/(2 + \omega(1 - C))$				
\bar{h}_2	$2(M - D(1 - C))/(2C(1 + \omega) - C^2\omega)$				
$\bar{h} = \begin{cases} \bar{h}_{01} & \text{if } \bar{h}_{01} < \bar{h}_1 \text{ else} \\ \bar{h}_{02} & \text{if } \bar{h}_{02} < \bar{h}_2 \text{ else} \\ \bar{h}_{03} & \text{else} \end{cases}$	$\bar{h}_{01} = 2\omega\phi$		$\bar{h}_{02} = 2\omega(D(1 - C) + \phi)/(1 + \omega(1 - C)^2)$		
	$\bar{h}_{03} = 2\omega(M + \phi)/(1 + \omega)^2$				

Notes: This table summarizes the calibration of the out-of-pocket cost functions and the analytical solution for healthcare utilization for each plan type within the model.

HMO plans tend to have stricter measures in place that enrollees have to clear before utilizing care the plan will cover (i.e., referrals and prior authorization). These additional steps Medicare Advantage plans take to reduce utilization among their enrollees is captured by the plan-type component included in the hassle cost of utilizing care.

The middle panel of Table C.1 reports the analytic solution for the optimal amount of healthcare for a beneficiary to consume. These expressions depend on an individual's health state h_{it} , moral hazard parameter ω_i , and plan choice. These expressions have intuitive interpretations. Given the hassle costs of utilizing care ϕ_{ijt} an individual must have a sufficiently large health need to justify consuming a positive amount of healthcare. These hassle costs also capture measures MA insurers may use to limit the amount of care their beneficiaries consume. Once this health threshold is met, individuals in plans with a deductible face a marginal cost of one and will consume healthcare at that rate. As health needs grow and the beneficiary approaches their deductible amount, their utilization will jump beyond their deductible in anticipation of the lower marginal cost of consuming care due to the cost sharing with coinsurance. This behavior induces them to consume healthcare above their health state, which is traditionally interpreted as moral hazard spending and is partially mitigated by cost sharing. Similar logic applies for the discontinuity MA beneficiaries face as they

approach their plan’s out-of-pocket maximum. After reaching \bar{h}_2 spending discontinuously jumps to consume the full amount of care informed by their health state and moral hazard parameters, consistent with the fact that the marginal cost of care at this point is zero. The final item to note is that if the size of the coinsurance region for a plan is small relative to a beneficiary’s moral hazard parameter, it is optimal for them to immediately jump from the deductible region to the out-of-pocket maximum region.

Computing Q_{ijt}^* for a given set of model parameters requires integrating over the unobserved heterogeneity in the health state distribution. I employ quadrature to handle this integration in a relatively simple manner. I use nine nodes (n_s) to approximate the joint distribution of the observable component of health state distribution mean and the moral hazard parameter $(\bar{\mu}, \log \omega)$. These nodes and associated weighting matrix are denoted by d_s and W_s respectively.

For a given s node, I can evaluate draws from the health state distribution. Notice:

$$\begin{bmatrix} \bar{\mu}_{its} \\ \log \omega_{is} \end{bmatrix} = \begin{bmatrix} \mathbf{X}_{it}^\mu \boldsymbol{\beta}^\mu \\ \mathbf{X}_i^\omega \boldsymbol{\beta}^\omega \end{bmatrix} + d_s \cdot \text{chol} \left(\begin{bmatrix} \sigma_\mu^2 & \\ & \sigma_\omega^2 \end{bmatrix} \right) \quad (\text{C.33})$$

where “chol” denotes the Cholesky decomposition of the variance covariance matrix. I use 27 nodes (n_b) to approximate the health state distribution, whose nodes and weighting matrix are denoted by d_b and W_b . Thus for a given set of model parameters and s and b nodes the health state for an individual is:

$$h_{itsb} = \exp(\bar{\mu}_{its} + d_b \underbrace{\mathbf{X}_{it}^\sigma \boldsymbol{\beta}^\sigma}_{=\sigma_{h,it}}) \quad (\text{C.34})$$

From here it is straightforward to compute the node specific optimal healthcare utilization \hat{Q}_{ijtsb}^* and apply the quadrature weights to integrate over the health state distribution:

$$\hat{Q}_{ijts}^* = \sum_{b=1}^{n_b} W_b \cdot \hat{Q}_{ijtsb}^* \quad (\text{C.35})$$

Agents make the healthcare utilization decision conditional on their plan choice. Thus, the node specific optimal healthcare utilization \hat{Q}_{ijts}^* must be weighted by the node specific probability the individual enrolled in plan j , which is denoted by s_{ijmts} . After weighting \hat{Q}_{ijts}^* by the choice probabilities, I apply quadrature to integrate out the remaining unobserved heterogeneity and recover \hat{Q}_{ijt}^* :

$$\hat{Q}_{ijt}^* = \sum_{s=1}^{n_s} W_s \cdot s_{ijmts} \cdot \hat{Q}_{ijts}^* \quad (\text{C.36})$$

C.2 Plan Choice

Agents in the model pick the Medicare Advantage plan (or Traditional Medicare) from their plan menu \mathcal{J}_{mt} that maximizes their expected utility. The expectation is taken over the distribution of their future health state. Calculating choice probabilities from this model present two challenges. The first is the double exponentiation introduced by the CARA utility function and the second is integrating over the unobserved heterogeneity in the health state distributions. To address the former and avoid numerical issues, I follow Marone and Sabety (2022) and use certainty equivalent utility to construct choice probabilities, while quadrature is used to integrate the unobserved heterogeneity. Thus, for a given set of model parameters and s node an individual certainty equivalent utility for plan j is (U_{ijmts}^{CE}):

$$U_{ijmts}^{CE} = \bar{l}_{ijmts} - \frac{1}{\psi} \log \left(\sum_{b=1}^{n_b} W_b \cdot \exp[-\psi(l_{ijmts}(h_{itsb}) - \bar{l}_{ijmts})] \right) \quad (\text{C.37})$$

where $\bar{l}_{ijmts} = \mathbb{E}_h[l_{ijmts}(h_{itsb})]$. Given the assumptions on $\zeta_{ig} + (1 - \sigma)\epsilon_{ijmt}$, the node specific choice probabilities take the nested logit form. Applying quadrature integrates out the unobserved heterogeneity:

$$\bar{s}_{ijmts} = \frac{\exp(U_{ijmts}^{CE}/(1 - \sigma))}{1 + \underbrace{\sum_{\ell \in \mathcal{J}_{mt}} \exp(U_{i\ell mts}^{CE}/(1 - \sigma))}_{\equiv D_{igmt}}} \quad (\text{C.38})$$

$$s_{ijmts} = \bar{s}_{ijmts} \frac{D_{igmt}^{(1-\sigma)}}{1 + D_{igmt}^{(1-\sigma)}} \quad (\text{C.39})$$

$$s_{ijmt} = \sum_{s=1}^{n_s} W_s \cdot s_{ijmts} \quad (\text{C.40})$$

Finally, market shares s_{jmt} are obtained by integrating over the population of individuals within the market. Let W_{imt} and M_{mt} denote the weight on each individual in market and the market size. Market shares are computed as:

$$s_{jmt} = \sum_{i=1}^{M_{mt}} W_{imt} \cdot s_{ijmt} \quad (\text{C.41})$$

C.3 Subsidy Choice and Unobserved Costs

In this section, I provide additional details about service areas and how I recover unobserved MA plan costs. Defining service areas is important to determining the set of direct and indirect competitors for MA plans. Unobserved plan costs are important to capture as my healthcare utilization metric does not include all potential claim/encounter types and does not capture non-utilization based costs associated with enrollment.

Service area definition. The geographic space where Medicare Advantage plans compete are called service areas. Service areas are defined at the state level. For larger states like California, Texas, and Florida, service areas are subsets of counties within the state based on commonly understood geographic boundaries (i.e., South Florida, West Texas, Southern California, etc.).

Observed entry patterns of plans largely align with these service area definitions. For the plans with an observed footprint that spans multiple service areas, I assign them to their primary service area where the plurality of their enrollees are located. For purposes of estimating the model, these plans make endogenous decisions within their primary service area but are taken as exogenous players in the other service areas where they are present.

Recovering unobserved costs. I used data on MA plan margins to recover unobserved costs. Given these data and my parameter estimates for the health state distribution and consumer demand, I solve Equation (15) analytically for λ_{jt} .

C.4 Stage 2 Estimation

This section describes the moments used to estimate the Stage 2 parameters of the model as well as the estimation algorithm.

To estimate the Stage 2 parameters I use the general method of moments. The overall procedure resembles a micro-BLP application and follow many of the best practices recommended by Conlon and Gortmaker (2023).

Let $\mathcal{M}(\theta)$ denote the vector of moment equalities that target healthcare utilization patterns and the IV restriction and depends on the model's parameters. I search for the parameter vector $\theta = \{\theta_1, \theta_3\}$ that solves:

$$\hat{\theta} = \arg \min_{\theta} \mathcal{M}(\theta)' \mathcal{W} \mathcal{M}(\theta) \tag{C.42}$$

where \mathcal{W} is a positive definite weighting matrix.

I first obtain an initial estimate for the optimal weighting matrix $\hat{\mathcal{W}}$ based on initial guesses for θ that fits the moments reasonably well. Given this estimate for $\hat{\mathcal{W}}$ I search for the parameter vector $\hat{\theta}$ which solves Equation (C.42). Once this process converges, I update my estimate for the optimal weighting matrix and repeat the search process. After the two-step estimation procedure is complete I obtain standard errors using the standard formula for the variance-covariance matrix of the GMM estimator.

Below is a description of the steps in the estimation algorithm for a candidate θ .

1. Compute the health state realizations h_{itsb} given the candidate parameter vector.
2. Compute the relevant quantities from the health state distribution to construct the model moments. These calculations are done for each category of observable heterogeneity c in each plan network-generosity type and the outside option.
3. Compute the utilization stage utility (see Equation (7)) for each health state realization. This requires recovering the out-of-pocket costs associated with the model implied Q_{ijtsb}^* for each plan choice type in the model. Hassle costs are recovered given the a candidate parameter vector.
4. For each market m :
 - (a) Recover the mean utility parameter δ_{jmt} using the Berry et al. (1995) contraction mapping that allows model predicted plan-level market shares to match their data analogs (i.e., $\hat{s}_{jmt}(\delta, \theta) = s_{jmt}$). I use the SQUAREM algorithm proposed by Varadhan and Roland (2008) to speed up the convergence of this fixed point.
 - (b) Use the model choice probabilities to construct the model predicted healthcare utilization and plan choice moments for the individuals in the market.
5. Recover the demand residual ξ_{jmt} for the IV moment using the 2SLS formula.
6. Compute the moments in $\mathcal{M}(\theta)$ and evaluate the objective function in Equation (C.42).

The estimates for θ_2 are recovered post-estimation using the formula for the 2SLS estimator with the values for δ_{jmt} associated with the $\hat{\theta}$ estimates as the dependent variable. Estimates for θ_4 are recovered from the auxiliary regression in Equation (18). Given these parameters estimates, I can recover the empirical distribution of the demand and efficiency shocks $e = (\xi, \varepsilon)$, which are used when deriving the moment inequalities.

C.5 Stage 1 Moment Inequality Derivation and Inference Details

Derivation. This section provides additional technical details related to the derivation of the moment inequalities used to estimate the parameters in Stage 1 of the model. As discussed in the main text, firms are endowed with CMS contracts that determine all possible plans the firm may offer in a service area. These contracts are network type specific and all plans offered under the contract have the same provider network and quality rating. Given this structure deviations from the observed decisions have a product characteristic and geographic component.

Let's first consider the characteristic deviations within a single market. To fix ideas, suppose we observe a firm with an HMO contract that entered plan j in market m as a low generosity HMO. There are two possible deviations to consider: plan j could have entered as a high generosity HMO or the firm could have also offered a second plan k as a high generosity HMO in the market alongside j .⁷ If a firm is observed to hold both HMO and PPO contracts within the service area, then same logic generates 14 possible deviations relative to the observed equilibrium.⁸

Now we can add the geographic component of the deviations. Let's further suppose that the service area in question has only four counties. For the firm with only an HMO contract there are 4,094 possible deviations where they enter at least one market and offer at least one product.⁹ By the same logic, for a firm with an HMO and PPO contract there are over $1.15e18$ possible deviations. Thus it is necessary to place restrictions on the types of deviations that are permissible to maintain tractability.

I start this process by defining the competitively relevant firms within a service area. A firm falls into this category if the share of MA beneficiaries it enrolls within its primary service area is greater than 5%. Firms that do not meet this threshold comprise the competitive fringe. These firms are not considered as part of the deviation sets and their decisions are taken as exogenous when solving the counterfactual equilibria. Next I define similar plan pairs among the competitively relevant firms. Two plans are considered similar if they are offered in the same service area, have the same network type and generosity level, star ratings within half a point, and a premium within a single standard deviation. For each plan in the similar plan pair, I iteratively simulate adding or removing the plan for each market within

⁷In cases where the firm offers the high generosity plan HMO k in markets other than m this deviation is equivalent to saying that plan k also enters m .

⁸The 14 deviations arises from the $2^4 - 2$ possible configurations of 4 possible plan types where at least one plan is offered and one of the possible configurations is observed in data.

⁹This number arises from the fact that there are 4 possible markets with 3 possible plan offerings in each market. Thus there are $2^{12} - 1$ possible entry configurations where at least one market is entered and one of these configurations is observed, leaving 4,094 deviations.

the service area holding fixed decisions about other markets and the choices of other firms. This process involves computing a firm's expected profits over the distribution of the demand and efficiency shocks $e = (\xi, \varepsilon)$. I take draws from this empirical distribution, compute the equilibrium given these draws, and average over the draws to compute the firm's expected profits.

After simulating the observed and counterfactual equilibria for the competitively relevant plans, I account for selection bias. As discussed in the main text, I leverage assumptions on the structural shocks ν_2 to employ a two level differencing strategy. The first difference is within firm and isolates the change in variable profits from adding or removing a market from a plan's observed footprint (see Equation (23)). The second difference is across similar plan pairs, where the isolated variable profit deviations involving adjacent markets are subtracted (see Equation (25)). I obtain unbiased moment inequalities for estimation by averaging over all adjacent market deviations within a plan.

Inference. I construct estimates for the identified set of fixed costs parameters by inverting the test statistic in Chernozhukov et al. (2019) for their SN1 subvector inference procedure. This method is attractive because it requires no tuning parameters and has a closed form for critical values, which reduces its computational burden. As described in the main text, the test statistic is based on studentization of the moment inequalities. To illustrate how the test statistic is constructed, let D denote the total number adjacent market pairs for a plan. Let $m_j(\theta)$ denote the inequality that eliminated the selection bias for plan pair j (i.e., Equation (25)):

$$m_j(\theta) = \sum_A \sum_{\mathcal{J}_{nmt}} \mathbb{E}[\Delta^- F(m, m') - \Delta^+ \mathbb{E} \Pi(m, m') - \Delta^+ \nu_1(m, m')] \leq 0 \quad (\text{C.43})$$

The mean and standard deviations for moment k are:

$$\varphi_k = \frac{1}{D} \sum_{d=1}^D m_{kd}(\theta) \quad \varsigma_k = \sqrt{\frac{1}{D} \sum_{d=1}^D (m_{kd} - \varphi_k)^2} \quad (\text{C.44})$$

These values for each moment are used to compute the test statistics:

$$T = \max_{1 \leq k \leq K} \frac{\sqrt{D} \varphi_k}{\varsigma_k} \quad (\text{C.45})$$

which are then assessed against the critical value for significance level α :

$$c(\alpha) = \frac{\Phi^{-1}(1 - \alpha/K)}{\sqrt{1 - \Phi^{-1}(1 - \alpha/K)^2/D}} \quad (\text{C.46})$$

I use the following procedure to invert the test statistics and construct the estimates for each subvector of the identified set.

1. Define a grid of 1,000 starting values for each parameter.
2. For each starting value in the grid minimize the test statistic until it falls just below the critical value.
3. Repeat for the entire grid of starting values for the parameter of interest.
4. Results from the optimization for each parameter represent the $1 - \alpha$ confidence set of the identified set of fixed cost parameters.

D Counterfactual Analyses

To compute the counterfactual equilibria of the model in a tractable way, I follow the procedure proposed by Lee and Pakes (2009). This approach has been used by other papers that solve models with multiple equilibria (see e.g., Wollmann, 2018). The method is based on an iterative best response. The procedure for solving for the equilibrium plan menu in year t proceeds as follows:

1. Set the initial plan menu in each market to what was observed in year $t - 1$ and endow the firms with a move order.
2. The first firm in the order best responds to $t - 1$ plan menu.
3. The second firm best responds to the $t - 1$ plan menu that includes the first firm's best response. This process continues for each firm in the move order.
4. After all firms have play their best responses, the process returns to the first firm. The algorithm stops when all firms have played without changing their best response.

An equilibrium in this procedure will satisfy the three Stage 1 entry conditions in Equations (EC.1)–(EC.3) that was used to derive the moment inequalities. As a result the

procedure will yield an equilibrium consistent with the simultaneous moves of firms in the model. The move order is determined by service area market shares in $t - 1$.

E Additional Tables and Figures

Table E.1: Summary of sample restrictions, 2017–2018

	Traditional Medicare		Medicare Advantage		Overall	
	N	Share	N	Share	N	Share
Individual criteria						
Initial sample	81,710,363	100	42,626,265	100	124,336,628	100
Age < 65	12,802,560	15.7	5,564,770	13.0	18,367,330	14.8
Months Part A \neq months Part B	11,197,972	13.7	341,081	0.8	11,539,053	9.3
ESRD or disabled	917,720	1.1	565,679	1.3	1,483,399	1.2
Invalid county ID	139,523	0.2	6,921	0.0	146,444	0.1
Alaska, Guam, Puerto Rico, or Virgin Islands	240,568	0.3	909,508	2.1	1,150,076	0.9
Missing risk score input	296,473	0.4	13,578	0.0	310,051	0.3
CMS aggregation criteria	246,745	0.3	298,147	0.7	544,892	0.4
MA criteria						
SNP, ESP, Part B only, or outside footprint			11,821,593	27.7	11,821,593	9.5
Missing plan characteristics			327	0.0	327	0.0
Non HMO or local PPO			3,000,316	7.0	3,000,316	2.4
Multiple segments			2,576,255	6.0	2,576,255	2.1
Analysis sample	55,817,400		17,579,492		73,396,892	
Unique beneficiaries	30,462,460		9,438,819		39,901,279	
Plan-year observations					3,624	
Unique plans					2,207	

Notes: This table summarizes the criteria used to isolate the analysis sample. These are based on individual and Medicare Advantage characteristics. Each row reports the number of beneficiaries impacted by each restrictions. The “N” column reports the number of beneficiaries and the “Share” column reports this value as a share of the initial sample of all Medicare beneficiaries.

Table E.2: Nested logit demand estimates, 2017–2018

	(1)
Premium	-0.54*** (0.06)
Supplemental revenue	0.59*** (0.16)
log MA share	0.52*** (0.01)
Year FE	✓
Contract FE	✓
Star rating	✓
Observations	24,572

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Standard errors clustered at the market-level. This table reports estimates for a nested logit demand model. These estimates are used in the reduced form analysis presented in Section III.

Table E.3: Individual level summary statistics, 2017–2018

	MA					TM				
	Mean	SD	P10	P90	N (1,000s)	Mean	SD	P10	P90	N (1,000s)
Demographics										
Age	74.9	7.3	67.0	86.0	17,579.5	75.4	8.0	66.0	87.0	55,817.4
Female	56.6	49.6	0.0	100.0	9,951.7	56.3	49.6	0.0	100.0	31,401.3
Low Income	10.2	30.2	0.0	100.0	1,788.3	12.8	33.4	0.0	100.0	7,161.2
New Medicare	3.9	19.3	0.0	0.0	680.0	4.8	21.5	0.0	0.0	2,707.1
New Medicaid	0.4	6.4	0.0	0.0	72.6	0.6	7.5	0.0	0.0	311.8
Died	3.1	17.3	0.0	0.0	541.5	4.1	19.9	0.0	0.0	2,295.1
Active Choice	21.6	41.2	0.0	100.0	3,803.3	5.4	22.6	0.0	0.0	3,012.5
Risk score (IP)	0.9	0.8	0.5	1.2	17,579.5	1.0	1.1	0.5	1.7	55,817.4
Risk score (IP-OP-CAR)	1.2	1.1	0.4	2.4	17,579.5	1.0	1.2	0.4	2.1	55,817.4
Util (Std. \$)	8.4	24.2	0.0	21.7	6,216.3	10.7	20.1	0.5	28.8	7,383.2
Util (Std. \$) Use	9.7	25.8	0.5	24.6	5,391.3	11.0	20.3	0.7	29.4	7,175.1
Markets										
White	77.8	15.4	55.8	95.0	17,579.5	80.8	15.0	59.0	96.1	55,817.4
Median Income	23,588.0	5,490.8	17,666.4	31,051.6	17,576.5	23,619.0	6,146.5	16,871.4	32,529.9	55,807.9
College	31.0	9.5	18.7	43.8	17,579.5	29.7	11.1	15.8	45.5	55,817.4
Average Normalized Risk Score (IP only)	1.0	0.1	0.9	1.0	17,579.5	1.0	0.1	0.9	1.0	55,817.4
Share New Beneficiaries	6.2	0.6	5.5	6.9	17,579.5	6.1	0.6	5.3	6.9	55,817.4
Share New Medicaid	1.1	0.6	0.5	2.1	17,579.5	1.0	0.5	0.5	1.6	55,817.4
Share Medicaid	20.5	8.4	11.6	33.6	17,579.5	19.1	7.7	11.0	30.0	55,817.4
Medicare Death Rate	3.7	0.4	3.3	4.2	17,579.5	3.7	0.4	3.3	4.3	55,817.4
Sample size										
Beneficiary-Years					17,579.5					55,817.4
Beneficiaries					9,438.8					30,462.5
Panel Sample					15,734.5					51,256.7

Notes: This table compares Medicare Advantage and Traditional Medicare beneficiaries in our analysis sample based. “Active Choice” measures whether a beneficiary changed their coverage option relative to the prior year or if they were new to the Medicare program. Healthcare utilization is measured in terms of standardized dollars. All market demographics except the rural share, college degree, and median income are measured for the Medicare population.

Table E.4: Summary statistics of MA sample, 2017–2018

	Other MA					Utilization sample				
	Mean	SD	P10	P90	N (1,000s)	Mean	SD	P10	P90	N (1,000s)
Demographics										
Age	74.9	7.3	67.0	86.0	11,363.2	74.9	7.3	67.0	86.0	6,216.3
Female	56.7	49.6	0.0	100.0	6,438.7	56.5	49.6	0.0	100.0	3,513.0
Low Income	10.0	30.0	0.0	100.0	1,137.3	10.5	30.6	0.0	100.0	651.0
New Medicare	3.9	19.4	0.0	0.0	444.6	3.8	19.1	0.0	0.0	235.4
New Medicaid	0.4	6.5	0.0	0.0	48.7	0.4	6.2	0.0	0.0	23.9
Died	3.1	17.2	0.0	0.0	348.7	3.1	17.3	0.0	0.0	192.8
Active Choice	21.0	40.7	0.0	100.0	2,382.8	22.9	42.0	0.0	100.0	1,420.5
Risk score (IP)	0.9	0.8	0.5	1.2	11,363.2	0.9	0.8	0.5	1.2	6,216.3
Risk score (IP-OP-CAR)	1.1	1.1	0.4	2.2	11,363.2	1.3	1.2	0.4	2.6	6,216.3
Util (Std. \$)					0.0	8.4	24.2	0.0	21.7	6,216.3
Util (Std. \$) Use					0.0	9.7	25.8	0.5	24.6	5,391.3
Markets										
White	77.4	15.6	55.8	95.0	11,363.2	78.6	15.0	55.8	94.8	6,216.3
Median Income	23,827.7	5,653.8	17,877.5	32,067.9	11,363.2	23,149.8	5,150.6	17,512.1	30,145.4	6,214.5
College	31.2	9.6	18.9	44.1	11,363.2	30.6	9.4	18.6	42.9	6,216.3
Average Normalized Risk Score (IP only)	1.0	0.1	0.9	1.0	11,363.2	1.0	0.1	0.9	1.0	6,216.3
Share New Beneficiaries	6.2	0.6	5.5	6.9	11,363.2	6.2	0.6	5.4	6.9	6,216.3
Share New Medicaid	1.1	0.6	0.5	2.1	11,363.2	1.0	0.5	0.5	1.7	6,216.3
Share Medicaid	20.8	8.4	12.0	33.6	11,363.2	19.9	8.2	11.4	31.3	6,216.3
Medicare Death Rate	3.6	0.4	3.3	4.1	11,363.2	3.7	0.4	3.3	4.2	6,216.3
Sample size										
Beneficiary-Years					11,363.2					6,216.3
Beneficiaries					6,443.5					3,542.2
Panel Sample					9,860.4					5,327.3

Notes: This table compares Medicare Advantage beneficiaries in our analysis sample based on whether they were enrolled in a contract with a high degree of data completeness. All beneficiaries in one of these contracts enter the utilization sample and are used to estimate the health state parameters. “Active Choice” measures whether a beneficiary changed their coverage option relative to the prior year or if they were new to the Medicare program. Healthcare utilization is measured in terms of standardized dollars. All market demographics except the rural share, college degree, and median income are measured for the Medicare population.

Table E.5: Summary statistics of TM sample, 2017–2018

	Other TM					Utilization sample				
	Mean	SD	P10	P90	N (1,000s)	Mean	SD	P10	P90	N (1,000s)
Demographics										
Age	75.4	8.0	66.0	87.0	48,434.2	75.4	7.8	67.0	87.0	7,383.2
Female	55.8	49.7	0.0	100.0	27,020.9	59.3	49.1	0.0	100.0	4,380.4
Low Income	12.1	32.6	0.0	100.0	5,862.9	17.6	38.1	0.0	100.0	1,298.2
New Medicare	5.0	21.8	0.0	0.0	2,411.8	4.0	19.6	0.0	0.0	295.4
New Medicaid	0.5	7.4	0.0	0.0	263.4	0.7	8.1	0.0	0.0	48.5
Died	4.1	19.9	0.0	0.0	1,990.5	4.1	19.9	0.0	0.0	304.6
Active Choice	5.5	22.9	0.0	0.0	2,681.3	4.5	20.7	0.0	0.0	331.2
Risk score (IP)	1.0	1.1	0.5	1.6	48,434.2	1.1	1.2	0.5	1.8	7,383.2
Risk score (IP-OP-CAR)	0.9	1.1	0.4	1.7	48,434.2	1.7	1.6	0.4	3.6	7,383.2
Util (Std. \$)					0.0	10.7	20.1	0.5	28.8	7,383.2
Util (Std. \$) Use					0.0	11.0	20.3	0.7	29.4	7,175.1
Markets										
White	80.8	15.0	59.1	96.1	48,434.2	81.0	15.0	58.5	96.1	7,383.2
Median Income	23,603.2	6,142.1	16,871.4	32,529.9	48,425.7	23,722.4	6,174.5	16,889.6	32,734.7	7,382.3
College	29.7	11.1	15.8	45.0	48,434.2	29.8	11.1	15.8	45.7	7,383.2
Average Normalized Risk Score (IP only)	1.0	0.1	0.9	1.0	48,434.2	1.0	0.1	0.9	1.0	7,383.2
Share New Beneficiaries	6.1	0.6	5.3	6.9	48,434.2	6.1	0.6	5.3	6.9	7,383.2
Share New Medicaid	1.0	0.5	0.5	1.6	48,434.2	1.0	0.5	0.5	1.6	7,383.2
Share Medicaid	19.1	7.6	11.0	29.9	48,434.2	19.2	7.8	11.0	30.3	7,383.2
Medicare Death Rate	3.7	0.4	3.3	4.3	48,434.2	3.7	0.4	3.3	4.3	7,383.2
Sample size										
Beneficiary-Years					48,434.2					7,383.2
Beneficiaries					26,459.9					4,002.6
Panel Sample					43,948.9					6,761.0

Notes: This table compares Traditional Medicare beneficiaries in our analysis sample based on whether they appear in the claims data. Beneficiaries with claims data enter the utilization sample and are used to estimate the health state parameters. “Active choice” measures whether a beneficiary changed their coverage option relative to the prior year or if they were new to the Medicare program. Healthcare utilization is measured in terms of standardized dollars. All market demographics except the rural share, college degree, and median income are measured for the Medicare population.

Table E.6: Parameter estimates

Variable		Parameter	SE
Health state distribution			
Mean μ_h	Risk score Q_1	0.591	0.0009
	Risk score Q_2	0.862	0.0008
	Risk score Q_3	1.025	0.0007
	Risk score Q_4	1.824	0.0007
	Female	-0.187	0.0002
	Low income	-0.024	0.0002
	Age > 84	0.090	0.0003
	Market mortality rate	0.037	0.0002
	Market Medicaid eligibility	-0.067	0.0002
	Variance σ_h	Risk score Q_1	1.148
Risk score Q_2		1.176	0.0004
Risk score Q_3		1.160	0.0003
Risk score Q_4		0.976	0.0004
Hassle cost ϕ	TM	-1.985	0.0051
	MA HMO	0.168	0.0014
	MA PPO	-0.133	0.0013
Mean moral hazard $\log \omega$	Constant	-0.699	0.0018
Unobs het $\sigma_\mu, \sigma_\omega$	Health state mean	0.689	0.0010
	Moral hazard	0.553	0.0028
	Corr($\mu_h, \log \omega$)	-0.259	0.0021
Demand			
Premium α	Mean	-6.403	0.1886
	Low income	-3.190	0.0075
Utilization utility β	Mean	-1.144	0.0012
	Low income	0.342	0.0006
Supplemental revenue γ		1.175	0.4190
Nesting parameter σ		0.528	0.0008
TM-MA switching cost ι		-0.763	0.0002
CARA ψ		-0.975	0.0037
Contract FEs		✓	
Year FEs		✓	
Star rating FEs		✓	
Beneficiary-year observations		73,396,892	
Plan-year observations		3,624	

Notes: This table reports estimates for the health state distribution and demand parameters. Estimates are obtained from a two-stage GMM procedure that targets observed utilization and plan choice decisions and IV restrictions. Confidence intervals are constructed from standard errors obtained from the variance-covariance matrix of the GMM estimator.

Table E.7: Supplemental revenue as a function on plan characteristics, 2017–2018

	Supplemental revenue
High generosity plan	0.46*** (0.01)
HMO	0.18*** (0.01)
Benchmark	0.12*** (0.01)
Year FE	✓
Plan star rating	✓
Mean of Dep Var	0.85
Observations	3,622

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Robust standard errors are in parentheses. This table reports estimates from an OLS regression of the supplemental revenue an MA plan needs to fund additional benefits relative to TM onto MA plan characteristics. Supplemental revenue is measured in thousands of dollars annually per-beneficiary. The unit of analysis is at the plan level.

Table E.8: Impact of modeling choices on equilibrium outcomes (equilibrium ranges)

	Baseline	Cut benchmarks \$1,200			
	Entry + Repos	Entry + Repos	No Entry or Repos	Repos only	Entry only
Probability offered					
PPO-L	[0, 1]	[0, 1]	1	[0, 1]	[0, 1]
PPO-H	[0, 1]	[0, 1]	1	[0, 1]	[0, 1]
HMO-L	[0, 1]	[0, 1]	1	[0, 1]	[0, 1]
HMO-H	[0, 1]	[0, 1]	NA	[0, 1]	NA
Benchmark (\$1,000)					
PPO-L	[9.96, 10.66]	[8.76, 9.46]	9.46	[8.76, 9.46]	[8.76, 9.46]
PPO-H	[9.96, 10.66]	[8.76, 9.46]	781	[8.76, 9.46]	[8.76, 9.46]
HMO-L	[9.96, 10.66]	[8.76, 9.46]	9.46	[9.46, 9.46]	[8.76, 9.46]
HMO-H	[9.96, 10.66]	[8.76, 9.46]	NA	[9.46, 9.46]	NA
Premium (\$1,000)					
PPO-L	[0.38, 0.61]	[0.70, 1.30]	0.71	[0.71, 1.30]	[0.70, 1.30]
PPO-H	[2.70, 12.38]	[3.14, 1710]	3.13	[3.07, 9.56]	[3.14, 11.26]
HMO-L	[0.00, 0.04]	[0.01, 0.23]	0.08	[0.08, 0.10]	[0.00, 0.23]
HMO-H	[0.91, 3.73]	[1.18, 6.46]	NA	[1.16, 3.52]	NA
Average risk score					
PPO-L	[0.77, 0.85]	[0.72, 0.78]	0.77	[0.74, 0.77]	[0.72, 0.78]
PPO-H	[0.74, 0.9]	[0.74, 0.86]	0.75	[0.75, 0.85]	[0.74, 0.86]
HMO-L	[0.9, 0.99]	[0.82, 0.98]	0.89	[0.88, 0.89]	[0.82, 0.98]
HMO-H	[0.8, 1.01]	[0.76, 0.96]	NA	[0.81, 0.96]	NA
Market size and selection					
MA share (%)	[0.05, 0.39]	[0, 0.41]	0.18	[0, 0.18]	[0, 0.75]
MA risk score	[0.83, 0.94]	[0.74, 0.94]	0.88	[0.76, 0.88]	[0.74, 1.02]
TM risk score	[1.08, 1.15]	[1.06, 1.19]	1.1	[1.06, 1.1]	[1.07, 1.27]
Per-bene risk adjusted government cost (\$1,000)					
MA	[9.9, 10.46]	[8.14, 9.56]	9.06	[9.01, 9.43]	[8.42, 9.43]
TM	[9.11, 10.00]	[9.01, 9.98]	9.47	[9.05, 9.47]	[9.01, 10.91]

Notes: This table reports how simulated quantities from the model change as different features are added. Each column reports the range of values across all equilibria of the model. The top panel reports quantities for the strategic plans. The bottom panel reports service area quantities. The first two columns report predicted outcomes from a model where firms are allowed to alter their market entry and product offering decisions (“Entry + Repos”). The third column restricts firms to the market entry and product offering decisions observed in the data (“No Entry or Repos”). The fourth column allows the strategic firms to alter their product offerings in the markets where these firms are observed to have entered in the data (“Repos only”). The final column allows the strategic firms to alter the markets where they enter but they can only offer products that are observed in the data (“Entry only”).

Table E.9: Equilibrium outcomes under alternative subsidy systems (equilibrium ranges)

	Baseline	No subsidy	Untargeted subsidy	Targeted subsidy
Probability offered				
PPO-L	[0, 1]	NA	[0, 1]	[0, 1]
PPO-H	[0, 1]	NA	[0, 1]	[0, 1]
HMO-L	[0, 1]	NA	[0, 1]	[0, 1]
HMO-H	[0, 1]	NA	[0, 1]	[0, 1]
Benchmark (\$1,000)				
PPO-L	[9.96, 10.65]	NA	[9.96, 10.65]	[8.76, 9.45]
PPO-H	[9.96, 10.65]	NA	[9.96, 10.65]	[8.76, 9.45]
HMO-L	[9.96, 10.65]	NA	[9.96, 10.65]	[8.76, 9.45]
HMO-H	[9.96, 10.65]	NA	[9.96, 10.65]	[8.76, 9.45]
Premium (\$1,000)				
PPO-L	[0.38, 0.61]	NA	[9.68, 9.84]	[0.71, 1.27]
PPO-H	[2.7, 12.38]	NA	[10.49, 25.16]	[3.59, 17.72]
HMO-L	[0.00, 0.04]	NA	[9.26, 9.38]	[0.32, 0.47]
HMO-H	[0.92, 3.73]	NA	[9.72, 23.15]	[0.78, 15.96]
Average risk score				
PPO-L	[0.77, 0.85]	NA	[0.82, 0.84]	[0.76, 1.04]
PPO-H	[0.74, 0.90]	NA	[0.74, 0.94]	[0.00, 0.88]
HMO-L	[0.90, 0.99]	NA	[0.93, 1.03]	[1.10, 1.18]
HMO-H	[0.80, 1.01]	NA	[0.85, 1.06]	[0.90, 1.10]
Market size and selection				
MA share (%)	[0.05, 0.39]	[0, 0]	[0.25, 0.58]	[0.1, 0.36]
MA risk score	[0.83, 0.94]	NA	[0.85, 0.99]	[1.02, 1.15]
TM risk score	[1.08, 1.15]	[1.07, 1.07]	[1.13, 1.18]	[1.04, 1.07]
Per-bene risk adjusted government cost (\$1,000)				
MA	[9.90, 10.46]	[0, 0]	[9.49, 11.04]	[9.00, 9.76]
TM	[9.11, 10.00]	[9.01, 9.01]	[9.63, 10.99]	[9.33, 10.21]

Notes: This table reports how simulated equilibrium outcomes change as the delivery system for Medicare Advantage subsidies changes. Each column reports the range of values across all equilibria of the model. The top panel reports quantities for the strategic plans. The bottom panel reports service area quantities. “Baseline” refers to the current system, which is a supply side subsidy that is scaled by a beneficiary’s risk score. “No subsidy” refers to a system where the government does not provide any subsidy to Medicare Advantage plans but regulates the plans meet their minimal coverage standards. “Untargeted subsidy” simulates a system that gives the observed enrollment weighted average risk adjusted pre-beneficiary subsidy for Massachusetts (approximately \$9,432 per year) to consumers to offset the costs of a Medicare Advantage plan. “Targeted subsidy” cuts CMS benchmarks by \$1,200 and offers a demand subsidy of \$600 to low income beneficiaries that enroll in Medicare Advantage plans and \$300 for non-low income MA enrollees.

Table E.10: Consumer surplus by observables under alternative subsidy systems

	Baseline	Untargeted subsidy	Targeted subsidy
Age \geq 86	[0.05, 0.51]	[0.27, 0.83]	[0.11, 0.42]
Female	[0.02, 0.17]	[0.09, 0.27]	[0.03, 0.14]
Low income	[0.07, 0.69]	[0.37, 1.12]	[0.14, 0.57]
High Medicaid county	[0.01, 0.12]	[0.06, 0.19]	[0.02, 0.1]
Risk score Q1	[0.04, 0.43]	[0.23, 0.69]	[0.09, 0.35]
Risk score Q2	[0.04, 0.37]	[0.2, 0.59]	[0.08, 0.3]
Risk score Q3	[0.04, 0.39]	[0.21, 0.64]	[0.08, 0.32]
Risk score Q4	[0.04, 0.35]	[0.18, 0.56]	[0.07, 0.29]

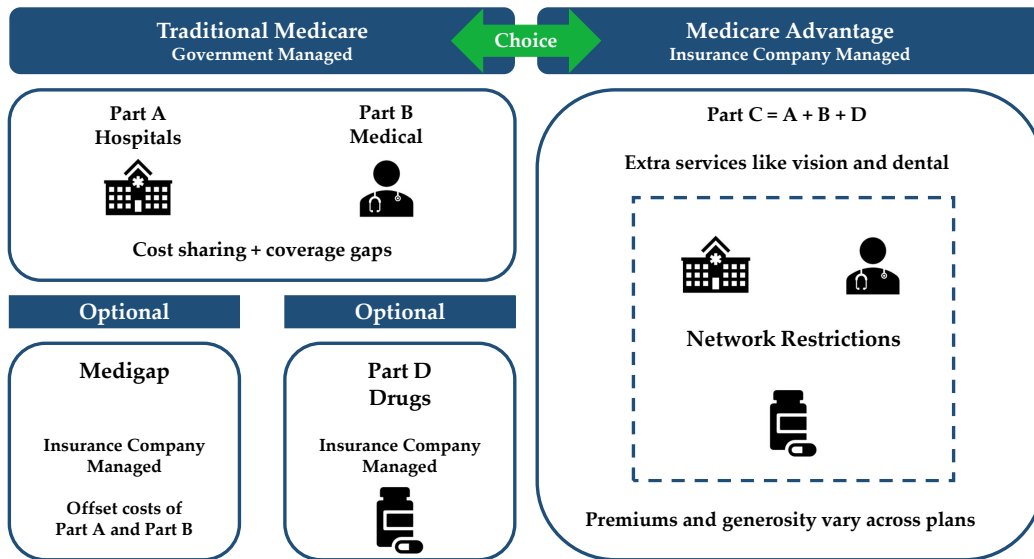
Notes: This table reports how consumer surplus is distributed across observable characteristics under different subsidy policies. Each column reports the range of values across all equilibria of the model. “Baseline” refers to the current system, which is a supply side subsidy that is scaled by a beneficiary’s risk score. “Untargeted” simulates a system that gives the observed enrollment weighted average risk adjusted pre-beneficiary subsidy for Massachusetts (approximately \$9,432 per year) to consumers to offset the costs of a Medicare Advantage plan. “Targeted” cuts CMS cost benchmarks by \$1,200 and offers a demand subsidy of \$600 to low income beneficiaries that enroll in Medicare Advantage plans and \$300 for non-low income MA enrollees.

Table E.11: Impact of ACA reform on county benchmarks, 2008–2017

	(1) Average Benchmark	(2) Average Benchmark	(3) NQB Benchmark	(4) NQB Benchmark	(5) QB Benchmark	(6) QB Benchmark
Post-ACA Transition	-43.10*** (1.38)	-58.46*** (1.71)	-58.06*** (1.50)	-71.44*** (1.79)	-32.01*** (1.29)	-47.50*** (1.66)
Year FEs	✓	✓	✓	✓	✓	✓
County FEs	✓		✓		✓	
Mean of dependent variable	746.21	746.21	737.21	737.21	749.76	749.76
F	4,444.8	915.7	4,204.7	980.1	5,470.3	874.1
Observations	15,056	15,056	15,056	15,056	15,056	15,056

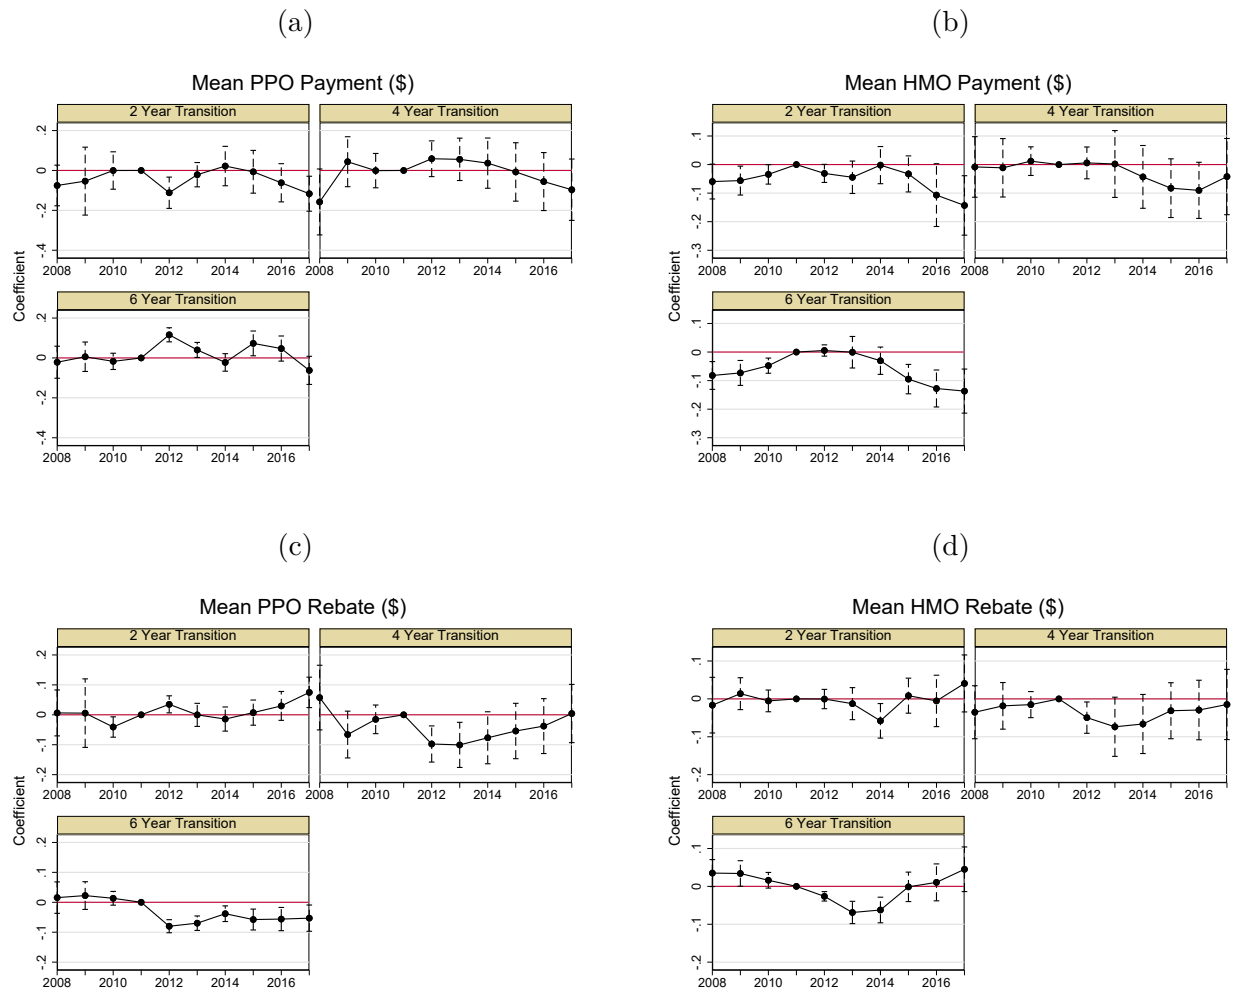
Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Standard errors clustered at the county-level. This table reports estimates for the impact of the ACA on MA county benchmarks. An observation is a county-year. Monetary values are converted into 2008\$. The sample contains counties in my analysis sample prior to 2012 and counties that completed their transition to the complete post-ACA benchmarks (counties had either 2, 4, or 6 years to transition). “NQB” stands for no quality bonus and “QB” stands for quality bonus. “Average Benchmark” measures the most ACA enrollment weighted average benchmark.

Figure E.1: Coverage choices in Medicare



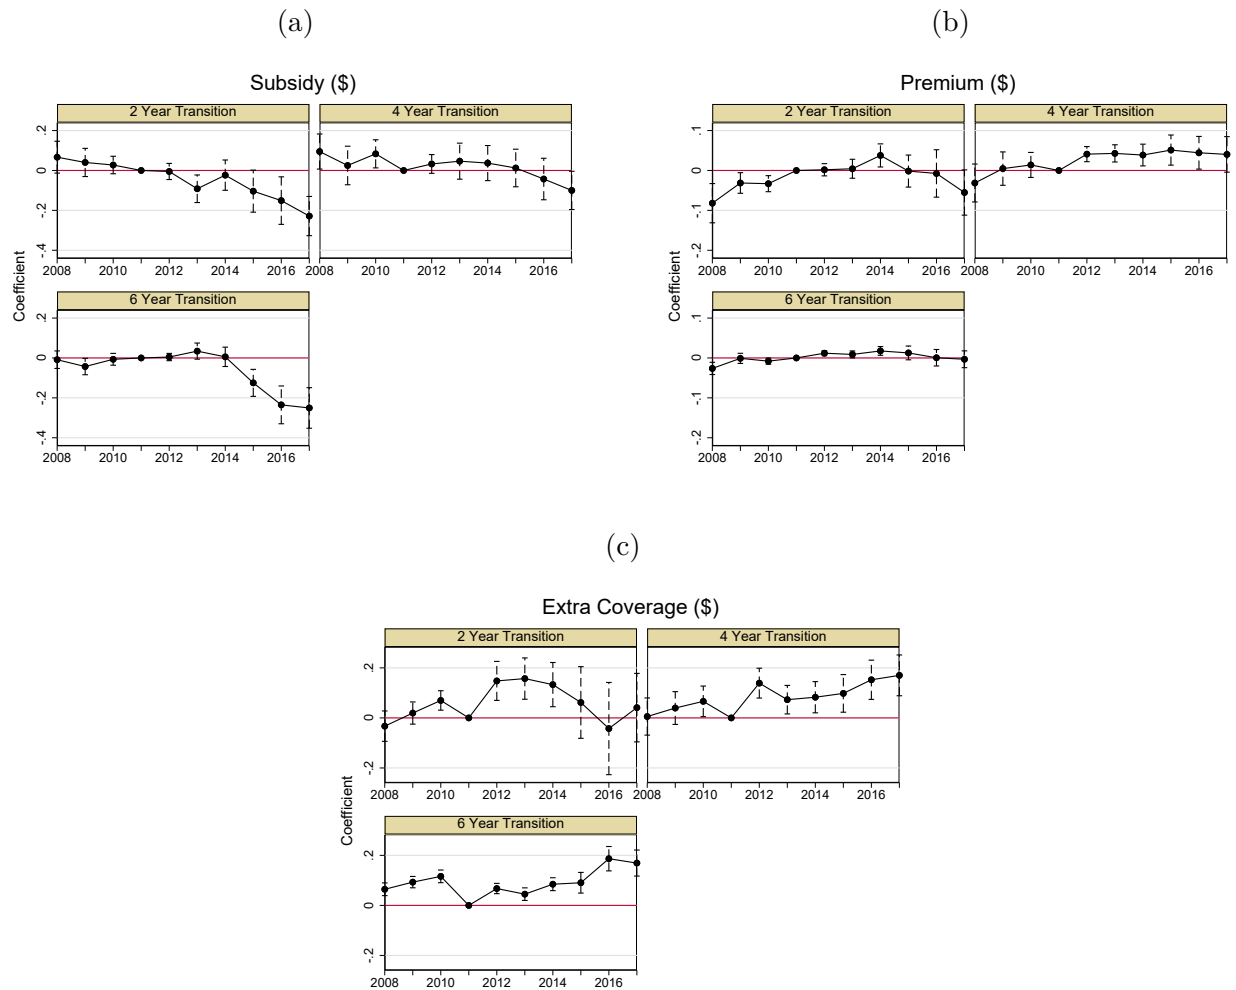
Notes: This figure summarizes the choices and tradeoffs Medicare beneficiaries face when making their annual health insurance coverage decisions.

Figure E.2: ACA benchmark event studies: average plan payments, 2008–2017



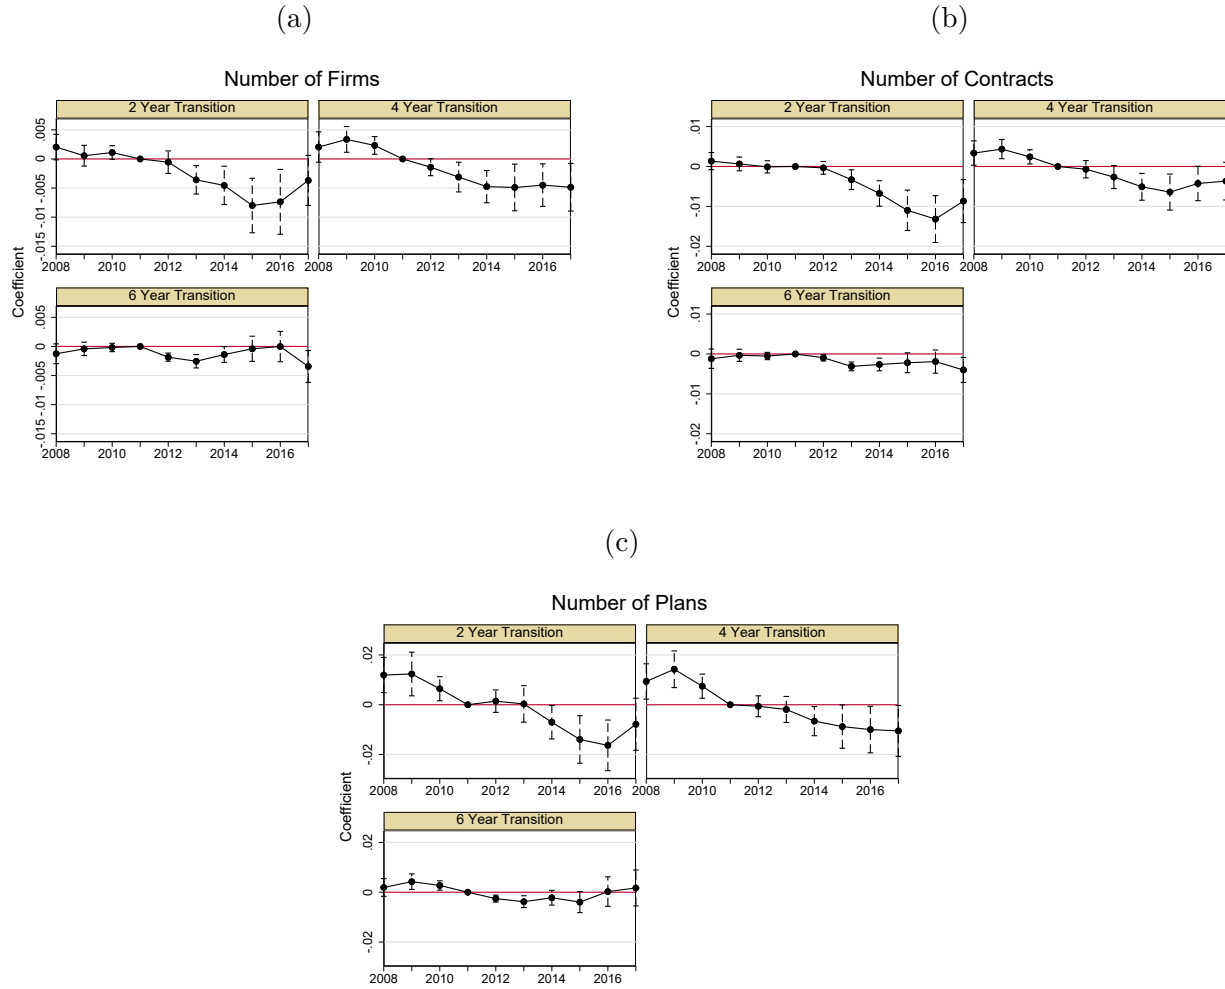
Notes: This figure plots the event study coefficients for the impact of ACA reductions to Medicare Advantage benchmarks on plan subsidy and rebate payments for HMO and PPO plans. Estimates are done separately by transition groups. An observation is a county-year. Dollar values are converted into 2008\$.

Figure E.3: ACA benchmark event studies: average plan characteristics, 2008–2017



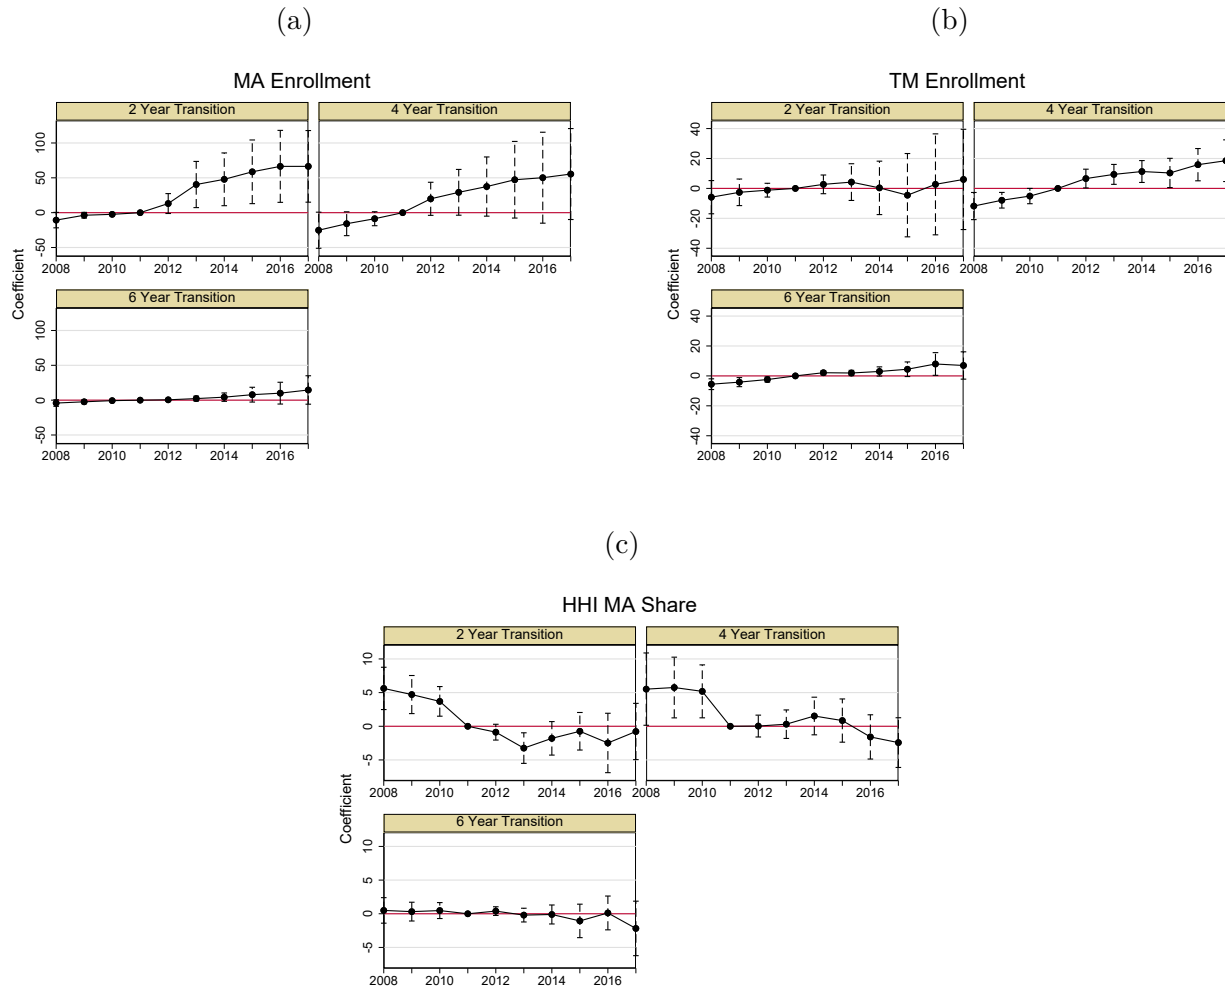
Notes: This figure plots the event study coefficients for the impact of ACA reductions to Medicare Advantage benchmarks on plan characteristics. Estimates are done separately by transition groups. An observation is a county-year. Dollar values are converted into 2008\$.

Figure E.4: ACA benchmark event studies: average entry, 2008–2017



Notes: This figure plots the event study coefficients for the impact of ACA reductions to Medicare Advantage benchmarks on firm entry. Estimates are done separately by transition groups. An observation is a county-year. Dollar values are converted into 2008\$.

Figure E.5: ACA benchmark event studies: average enrollment, 2008–2017



Notes: This figure plots the event study coefficients for the impact of ACA reductions to Medicare Advantage benchmarks on firm entry. Estimates are done separately by transition groups. An observation is a county-year. Dollar values are converted into 2008\$.