Optimal Managed Competition Subsidies

Keaton Miller[*†] Univ. of Oregon
Amil Petrin Univ. of Minnesota, NBER
Robert Town Univ. of Texas-Austin, NBER
Michael Chernew Harvard Univ., NBER

November 2018: Draft Beta 6

Abstract

The Medicare Advantage program enables Medicare recipients to receive their health care benefits via private insurance plans instead of through the federal government. Insurers receive a per-enrollee payment from the government, and may add additional benefits and charge an additional premium – a differentiated products approach which mirrors many other goods provided via a government subsidy. Subsidies are pegged to a measure of the government's cost, but the optimal subsidy in different geographic markets – conditional on a fixed amount of government expenditures across all markets – depends on the interactions between consumer demand and supply-side responses to changes in the payments offered by the government. We calculate the optimal geographic subsidy distribution in Medicare Advantage by estimating a flexible oligopoly model that features heterogeneous demand with switching costs, and supply-side price-setting and benefit design behavior. We introduce an equilibrium approximation approach for calculating counterfactual outcomes when the curse of dimensionality prevents a direct search for equilibria. We find the optimal subsidy schedule increases mean consumer surplus significantly in part by taking advantage of locations where the extensive margin budgetary impact of moving individuals from traditional Medicare to MA outweighs the intensive margin impact of increasing payments for people already enrolled in MA, and in part by redistributing funds to areas where marginal government expenditures result in the greatest increases in welfare.

JEL Codes: I13, I18, D43, D61, L10.

Keywords: Managed Competition, Industrial Organization, Optimal Subsidies, Medicare Advantage.

[*]Miller (corresponding author): keatonm@uoregon.edu; Petrin: petrin@umn.edu; Town: robert.town@austin.utexas.edu; Chernew: chernew@hcp.med.harvard.edu

[†]The authors would like to thank Mark Colas, Clare Evans, David Evans, Gautam Gowrisankaran, Ben Hansen, Kendall Houghton, Van Kolpin, Scott Petty, Jeremy Piger, Christopher Snyder, Glen Waddell, Wes Wilson, and seminar and conference participants at the University of Minnesota, the University of Texas - Austin, the International Industrial Organization Conference, North American and European Meetings of the Econometric Society, and the Asian Workshop on Econometrics and Health Economics for helpful comments and feedback.

1 Introduction

When the private market's provision of a good deviates from the socially optimal outcome, welfare can be improved through government intervention. The many ways in which market outcomes can be suboptimal suggest a potentially broad and direct government role in producing goods. However, there is a long-standing concern that government production programs can be inefficient, as government bureaucracies may lack sufficient incentives to efficiently design and deliver goods and services. This concern has led to the idea that governments can either regulate private firms or procure goods and services directly from those firms. These strategies raise a host of strategic and informational issues that make efficient implementation challenging (Laffont and Tirole, 1993).

An alternative approach is for a government to provide subsidies to (potentially heterogeneous) consumers who purchase the good from competing firms under the rationale that profit maximization motives and competitive market pressures will push firms to provide the optimal quantity, product variety, and quality, at a price nearing marginal cost. This rationale is persuasive enough that the "managed competition" approach is widely employed by the US government to provide health insurance (Gruber, 2017). For example, under the "Medicare Advantage" program (MA), Medicare beneficiaries can forgo the traditional Medicare (TM) benefit structure and enroll in one of many differentiated health plans offered by private firms. The firm assumes the financial and logistic responsibilities of managing the enrollee's care and, in return, receives a risk-adjusted per-capita payment from the government, based on a county-specific "benchmark rate," and premium payments from the enrollee. A similar system is used by Medicare Part D and the individual insurance marketplaces created under the Patient Protection and Affordable Care Act (ACA) and is also related to charter school initiatives.

A fundamental challenge with significant welfare implications in these publicly-financed, privately-produced environments is how to best set the subsidies. In particular, supply and demand conditions often vary across geographically distinct markets. This implies that the optimal subsidy should also vary across geographies as a function of those conditions. In practice, geographic variation in subsidies is typically driven by measures of the government's average cost of providing the service. Indeed, the MA benchmark rate is a function of the county-level average cost of covering TM enrollees and varies in our data from \$7,761 to \$16,001 across markets. The optimal subsidy depends in part on the private firms' costs, which may differ from the government's cost. More broadly, the optimal subsidy also depends on the quality- and price-setting behaviors of firms, and

the interaction between those behaviors and consumer preferences.

In this paper, we develop an approach for determining the optimal subsidy level across many markets in managed competition settings and apply it to the MA program. We take the mechanism that links a county-level subsidy to payments to firms as fixed and focus on differences across markets. In our setup, MA insurers may set premiums, plan benefits, and other product characteristics in response to both the benchmark rates set by the Centers for Medicare and Medicaid Services (CMS) and other competitive conditions.[1] To our knowledge, we are the first to study the optimal subsidy level in differentiated products environments in which firms can adjust both price and non-price characteristics in response to changes in the subsidy.

We start by estimating the demand for MA plans with a model that allows for considerable preference heterogeneity across observable plan and consumer characteristics and unobservable plan-specific 'quality' differences. Due to the variation in provider networks offered by plans both within and across insurers, we include a flexible specification of switching costs borne by consumers. We estimate our demand system with individual-level panel data on demographics and plan choices from the Medicare Current Beneficiary Survey (MCBS) for the years 2008-2015 supplemented with data on plan premiums and benefit structures from CMS.

We find the mean price elasticity of demand for plans to be -2.66. MA plans are more valuable to beneficiaries who have lower income, lower educational attainment, are younger, and are in better health. We also find that consumers who wish to change plans (and in which case would likely need to switch providers) face significant but plausible switching costs. The average cost to switch between MA insurers is $523 and the cost to switch between TM and MA is $898. For comparison, the average annual plan premium is $509.

We find the MA program as-is creates $149 in consumer surplus per year per Medicare beneficiary. Plans capture an average of $119 in variable profits per beneficiary. Our estimates imply that in 2015, the MA program generated a total of $7.38 billion in consumer surplus and $4.75 billion in firm variable profit with $98 billion in payments to MA firms.

We then consider the impact of counterfactual benchmarks on outcomes. The traditional approach to computing counterfactual equilibria, a fixed-point search over firm best-response functions, is impractical due to the complexity of MA products – we capture over 40 product characteristics. Instead, we estimate policy functions for product characteristics from the data, use those

[1] We do not explore the role of subsidy levels in affecting entry and exit. While insurers and plans in our data do enter and exit, the median share captured by new entrants in our data is 0.18%. Decarolis et al. (2015) make a similar argument.

estimated functions to predict characteristics under counterfactual benchmarks and then solve the firms' first-order conditions for prices taking those characteristics as given. The counterfactual equilibria show that the rate at which increases in the benchmark are passed through to consumers varies across markets. We also find that, for several markets, the current benchmark rates are set in such a way that the government spends less on an MA enrollee than an equivalent TM enrollee. As a consequence, for small benchmark increases, the extensive margin effect of switching people from TM to MA can outweigh the intensive margin effect of spending more on current MA enrollees and total government expenditures can decrease.

Our primary exercise determines the optimal benchmarks that maximize average consumer surplus subject to maintaining current spending. The implied optimal benchmarks are meaningfully different than the CMS benchmarks and take advantage of the extensive/intensive margin effects mentioned above. The average difference between the optimal benchmark and the CMS benchmark is 3.6%. The change in benchmarks can be dramatic – 14% of the counties have a optimal benchmark that is 10% greater than the benchmark in our data. The estimates imply that the optimal benchmark will increase average Medicare beneficiary consumer surplus 30% to $193 per year.

Our estimates of the optimal benchmark imply significant variance in the consumer surplus across individuals. For this reason we also explore calculating the benchmarks under different CMS objective functions that take into account both the mean and the variance of consumer surplus across individuals. We find that there are benchmark schedules that increase the average consumer surplus and reduce the variance relative to the 2015 policy.

Our approach allows us to perform other counterfactual exercises. A country-wide benchmark increase of $1 leads to an increase in government expenditures of $16.2 million and a $5.19 million increase in consumer surplus – in other words, the marginal pass-through rate is 32%. Minimizing government expenditures results in savings of 0.5% though at the cost of 20% of consumer surplus. Maximizing consumer surplus with the constraint that benchmarks stay at least at their 2015 levels results in an increase in the mean to $190.

We contribute to an extensive MA literature – see McGuire et al. (2011) for a review. Our work is most related to Town and Liu (2003), Lustig (2010), Curto et al. (2015) and Aizawa and Kim (2018). Town and Liu (2003) estimate a nested logit demand system for MA plans and calculate that the program (then called Medicare+Choice) generated $113 in consumer surplus and $244 in profits per Medicare beneficiary in 2000 with significant geographic variation. Curto et al. (2015) also estimate a nested logit model of MA demand using more recent data and find that

the program generated approximately $600 in total annual surplus, with the majority captured by insurers. They also estimate that, on average, MA plan costs are 12% lower than TM costs, though in 47% of counties MA does not have a cost advantage over TM. Nosal (2012) estimates a dynamic demand model of MA plan choice and finds very large switching costs—$4,000 at the median. Aizawa and Kim (2018) estimate a demand model that is similar to ours in order to explore the role of advertising in equilibrium selection. They analyze a counterfactual price equilibrium when advertising is prohibited, but do not endogenize other product characteristics.

Several recent papers estimate the passthrough from benchmarks to premiums and benefits in MA. Using an unanticipated change in the benchmark in 2000, Cabral et al. (2018) find that 54% of a benchmark increase is passed to enrollees, while Duggan et al. (2016) use variation in the benchmark across urban and rural counties and estimate a much smaller passthrough. Song et al. (2013) calculate a passthrough from benchmarks to plan bids, which are a measure of premiums and the value of benefits, of 53%. We expand upon this literature by conducting counterfactual analyses that endogenize the ways in which firms adjust plan features, not just premiums, in response to changes in the benchmark rate. In this way, our work is related to Fan (2013).

We also relate to recent work on optimal subsidy structures in different health insurance contexts – much of this work focuses on the ACA exchanges. Tebaldi (2017), Jaffe and Shepard (2017) and Einav et al. (2018) examine the optimality of different subsidy and/or risk-adjustment strategies in different ACA insurance exchanges. Ericson and Starc (2015) examine the implications of age-based price regulation in an ACA-like insurance exchange. Bundorf et al. (2012) examine the impact of health-status-linked premiums in the employer-sponsored setting. More broadly, we relate to a literature that considers various risk-adjustment strategies designed to address adverse selection—see Geruso and Layton (2017) for a review. CMS currently risk-adjusts MA subsidies, apparently successfully (Newhouse et al., 2015), and therefore we do not examine that issue here. Decarolis et al. (2015) examine the optimality of using vouchers verses the current subsidy strategy in Medicare Part D and find that the two systems generate similar welfare. None of these papers examine the impact of subsidies on non-premium characteristics of plans.

We proceed in Section 2 with a discussion of the institutional details of the Medicare Advantage program. Section 3 presents our data on Medicare beneficiaries and MA plans. We present our model in Section 4 and estimate its parameters in Section 5. Section 6 details the specifics of our optimal subsidy exercise, describes our approach to counterfactual analysis, and presents the results of that exercise. We conclude in Section 7.

2 The Medicare Advantage Program and the Role of Subsidies

In this section, we provide institutional background on the Medicare Advantage program. We then briefly describe the intuition underlying why demand and supply factors should be taken into account when optimizing the subsidy offered to private insurance firms.

2.1 The Medicare Advantage Program

Medicare was enacted in 1965. In its original form, the program provided hospital and medical insurance benefits to seniors (age 65 or older) through its Part A (hospital insurance) and Part B (physician and outpatient insurance) programs, respectively. Under its traditional fee-for-service reimbursement structure, private providers treat Medicare beneficiaries, and the providers bill Medicare (and the beneficiary for appropriate co-pays/coinsurance) for the services provided. Medicare then pays the provider according to a pre-set reimbursement schedule. In 1972, a significant expansion to Social Security extended the Medicare program to those who are eligible for Social Security disability benefits as well as those with end-stage renal disease (ESRD) (Lyons, 1972).

Partly in response to Medicare's increasing costs,[2] Congress created the Medicare Part C program as part of the Tax Equity and Fiscal Responsibility Act of 1982. The Part C program allowed Medicare administrators to contract with private insurers for the management of its beneficiaries' health benefits. By 1985, the Health Care Financing Administration, the precursor to the modern-day Centers for Medicare and Medicaid Services (CMS), formalized the rules for the contract process and began a series of trial programs based in part on the ideas of Enthoven (1978). Under these early trials, the government handed over management of the medical care of select groups of Medicare enrollees in exchange for a fixed payment that did not vary with the realized medical expenditures of each individual. This program was brought to the entire country in the early 1990s under the name Medicare+Choice.[3]

The Medicare+Choice program struggled to attract plans in many areas and nationwide enrollment remained under 10%. Critics blamed both the low subsidy rates and the lack of meaningful risk adjustment in payments, which incentivized firms to cream-skim relatively healthy individuals from the Medicare risk pool. The Medicare Prescription Drug, Improvement, and Modernization Act of 2003 aimed to remove this incentive by adding a comprehensive risk adjustment component to the payment mechanism. Under the new system, firms submit detailed demographic and

[2] In 1970, Medicare composed about 0.5% of GDP. By 1980, Medicare had grown to 1.1% of GDP.

[3] See McGuire et al. (2011) for a comprehensive history of the Medicare Advantage program.

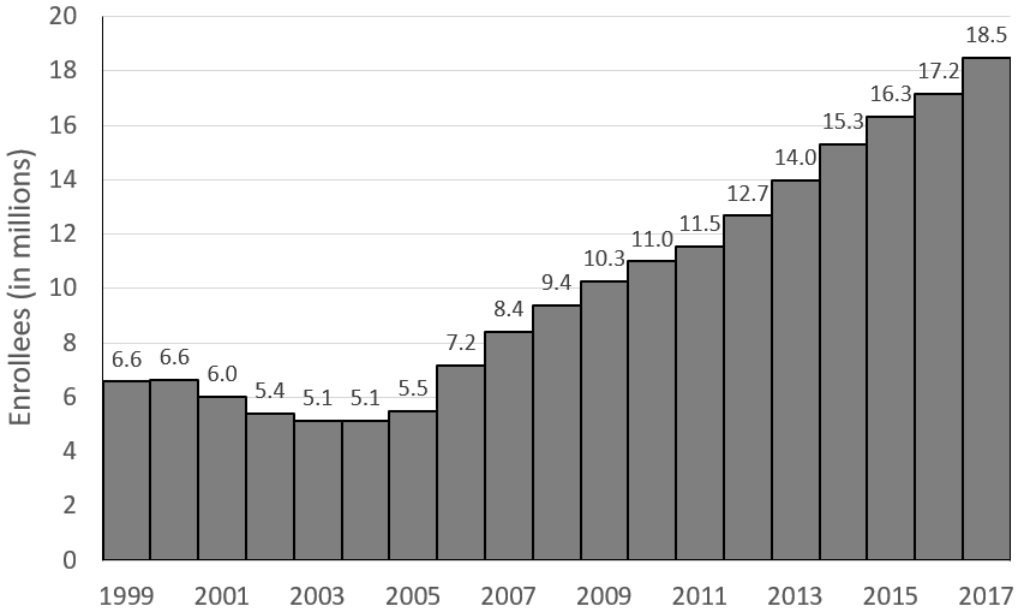
diagnosis data about beneficiaries to CMS at the time of enrollment. CMS assigns each enrollee a score based on its fee-for-service expenditures on similar individuals in TM. Payments to firms are then adjusted according to the average score of the individuals. Proponents argued that this mechanism would compensate firms for taking on risk without reimbursing specific procedures – thus maintaining the profit motive which would (in theory) lead to cost reductions. The updated program was renamed Medicare Advantage.

By 2015, 95% of Medicare beneficiaries had an MA plan operating in their county. A beneficiary who chooses to enroll in an MA plan forgoes TM benefits and instead receives medical benefits from the MA plan exclusively. MA enrollees must continue to pay the Medicare Part B premium, and may pay a private plan premium as well. MA plans generally provide a more generous benefit package than TM, including benefit categories not included under TM, such as dental and eye coverage. Many plans include a drug benefit. Plans also often offer lower out-of-pocket cost sharing and Part B premium reductions. These increases in benefit generosity generally come with the typical managed care tradeoffs of restrictive provider networks and referral requirements. Insurers compete along the dimensions of benefit design, premia, and provider networks, and often heavily market their plans (Aizawa and Kim, 2018).

The popularity of the Part C program has steadily increased since the Medicare Advantage reforms were implemented. Figure 1 shows the growth in MA enrollment over time. Enrollments remained near six million until the reforms passed in 2003 were implemented in 2007. Currently, over 33% of Medicare beneficiaries are enrolled in an MA plan.

The enrollee-specific subsidy from CMS to private firms is determined, in large part, by a fixed "benchmark" payment rate for each county, which varies significantly across geographies (Newhouse et al., 2012). CMS calculates the benchmark rate by starting with the average per-capita FFS Medicare spending within the county. This average is formed by adding the average Part A spending to the average Part B spending (as opposed to the average sum of Part A and Part B spending). Counties are then ranked by this average spending and placed into quartiles. The benchmark rate for counties in the top quartile of FFS spending is set to 95% of their average FFS spending. The benchmark rate for the second quartile is set to 100% of average FFS spending, the third quartile benchmark is 107.5% of average FFS spending, and the bottom quartile is set to 115% of average FFS spending. A cap and floor that varies by urban/rural status is applied to the benchmark rates across counties. Figure 2 illustrates the distribution of Medicare Advantage benchmarks across the counties in our sample for 2015.

Figure 1: Medicare Part C Enrollment, 1997-2017

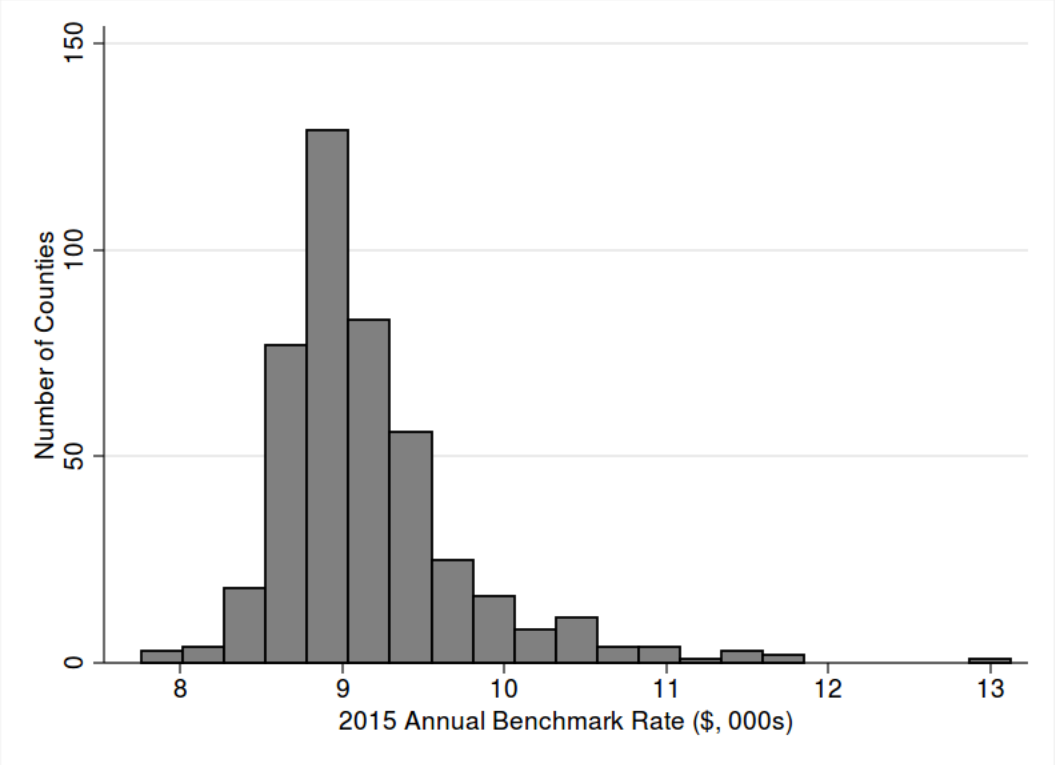


Source: CMS enrollment files.

In June of each year, insurers submit plan-level 'bids' for particular counties. The bid for each plan represents the insurer's offer to provide a stated set of benefits to a person of average risk for the next year in exchange for a particular price. The final bid amount must be related to the firm's revenue requirements and may be above or below the benchmark rate—CMS accepts all bids that conform to technical and financial solvency standards. Firms who bid above the benchmark rate must charge premiums to enrollees. Firms who bid below the benchmark receive some portion of the difference as a 'rebate' that must be passed on to consumers in the form of increased benefits or lower Medicare premiums. MA plans that offer a bundled prescription drug benefit submit a separate bid through the Part D process which maps in a similar way to a Part D premium which enrollees must pay.

At the beginning of our study period, the rebate payment was equal to 75% of the difference between the bid and the benchmark. In 2008, CMS introduced a system to measure insurer quality by assessing performance along many different dimensions and assigning a summary 'star rating' to each firm's plans. In 2012, after multiple iterations of the rating criteria, CMS began using the star rating to determine the rebate for each plan. Under this new approach, firms with at least 4.5 stars (out of five) earn 70% of the difference between the benchmark and the bid. New entrants and

Figure 2: Medicare Advantage Benchmarks, 2015



Note: Includes counties observed within the Medicare Current Beneficiary Survey.

those with 3.5 or 4 stars earn 65% of the difference. All others earn 50%. Additionally, firms with at least 4 stars have a 5% bonus applied to the benchmark rate itself (CMS Office of the Actuary, 2017).

Beneficiaries can enroll in plans during an Open Enrollment period that runs from October 15 to December 7 of each year, mirroring a commonly-observed private insurance system design. Beneficiaries are not allowed to enroll in plans outside of the Open Enrollment period except when they become newly Medicare eligible or in the case of "major" life changes (e.g. relocation, death in the family, etc.).[4] After the enrollment period closes, firms collect and transmit risk-adjustment information to CMS, which calculates the final subsidies and begins monthly payments in January of the following year.

To summarize, payments from CMS to insurers for an enrollee i living in county m enrolled in plan j in year t based on a benchmark B_{mt} can be represented by

$$Payment_{ijt} = \begin{cases} B_{mt} \times \phi_{jt} \times RiskAdjustment_{it} & \text{if } bid_{jt} \geq B_{mt}\phi_{jt} \\ (bid_{jt} + \lambda jt \times (B_{mt} \times \phi_{jt} - bid_{jt})) \times RiskAdjustment_{it} & \text{if } bid_{jt} < B_{mt}\phi_{jt} \end{cases} \quad (1)$$

where ϕ_{jt} and λ_{jt} capture the impact of the star rating system on the benchmark rate and the rebate, respectively. For notational simplicity, we will denote the market-level (i.e. county-year-level) base benchmark with B_m and denote plan-specific payments with B_j.

The MA program is a significant component of the federal budget. In 2015, payments to the plans in our data were $98 billion – traditional Medicare spending on the individuals in our data totaled $298 billion. The MA market is also relatively concentrated. The top four insurers nationwide, United Health Group, Humana, Kaiser, and Aetna, have 56% of total enrollment. This concentration is not a consequence of lack of choice – the average Medicare beneficiary has access to 10 plan options with 64% of beneficiaries having access to 5 or more plans. For a minority of beneficiaries the choice set is small – 25% of beneficiaries in our 2015 data have access to 3 or fewer plans. The average bid is roughly 90% of expected Traditional Medicare costs (MedPAC, 2017).

2.2 Optimal Subsidy Setting Across Markets

Much of the policy discussion surrounding Medicare payments focuses on setting the payments equal to the average cost of providing the service. For example, the Medicare Payment Advisory

[4] More recently, enrollees have been allowed to switch to a "5 star" plan at anytime during the year.

Commission (MedPAC) typically recommends that payments be cut if they determine the current payment is greater than the cost of provision (see MedPAC, 2017, chap. 9). However, under heterogeneous demand, setting the subsidy equal to a function of the measured average cost of production is only optimal under perfect competition. In the case of imperfect competition, the optimal allocation of the subsidy across markets (for a fixed subsidy budget) depends on the shape of demand and its heterogeneity, market structure, and equilibrium quality responses to changes in the subsidy. If there is meaningful heterogeneity in these factors across markets, the optimal subsidy will typically vary across markets in a way that may not be correlated with measured average costs. The goal of our analysis is to account for these different factors in determining the optimal subsidy.

In Appendix A, we present a simple model of price-setting with demand heterogeneity that makes this point. The model examines two geographically separated monopolistic markets, with the demand in one market yielding higher average willingness to pay for the product. Not surprisingly given this setup, the optimal subsidy will be zero in one market and the entire budget will be allocated to the market with the higher willingness to pay. In this example, the subsidy also increases total welfare (ignoring the welfare loss from tax collections) as it counteracts the distortionary impact of monopoly power on prices. It is also easy to construct models that focus on the production of quality heterogeneity that yield similar optimal subsidy distributions.

3 Data

To understand the demand and supply conditions in the MA market, we combine administrative data on plan characteristics and enrollment from CMS with micro-level data on consumer choices from the Medicare Current Beneficiary Survey. In the following subsections, we discuss each of these data sources in turn.

3.1 Medicare Advantage plans

We obtain the annual Plan Finder database from CMS for the years 2008 to 2015, which contains comprehensive information on plan benefits on an annual basis. We extract plan premiums, copays and coinsurance rates for primary care, specialist, and hospital visits, indicators for HMO, PPO, and FFS plan types, and indicators for vision, dental, and prescription drug coverage of any kind. We also extract the county-level geographic coverage for each plan. We collect star ratings, plan-

level bid amounts, and county-level benchmarks from other CMS sources. Finally, CMS releases enrollment counts each month for each county and plan. We average these monthly data over each plan year and combine them with CMS counts of the number of people singly-eligible for Medicare benefits (i.e. not also eligible for Medicaid) to form product shares at the plan-county-year level.

As part of the bidding and enrollment process, CMS calculates and publicizes the expected out-of-pocket costs (OOPC) for each plan for different demographic groups. CMS creates these estimates by forming a representative bundle of services used by TM enrollees. CMS then calculates the out-of-pocket costs for that bundle under each plan's benefit structure, which implicitly assumes that service-level consumption patterns do not change between TM and MA. We collect the OOPC estimate for each plan by age group (Under 65, 65-69, 70-74, 75-79, 80-84, and Over 85) and health status (Excellent, Very Good, Good, Fair, and Poor).

We focus on the market for individual insurance described in Section 2, and therefore drop plans sponsored by employers and plans designed for individuals who are "dual-eligible" for both Medicare and Medicaid, as plans in these categories operate under a different payment system and benefit structure. Due to CMS data restrictions, we drop plan-county observations with ten or fewer enrollees. While we obtain data on plans for every county, our micro-data comes from a sample of counties. Thus, for consistent presentation, we focus on those plans which fall within our micro-data sample area.

Summary statistics on the 50,593 plan-county-year observations in our data are reported in Table 1. The first column presents unweighted means of the product characteristics we collect. The average plan charges a monthly premium of $42 and requires a $94 deductible. Most plans offer at least one additional category of coverage beyond TM's offerings: 76% of plans offer prescription drug coverage, 60% offer dental coverage, and 90% offer vision coverage. Out-of-pocket costs vary substantially by age and health status: the average plan is expected to cost a 65 year-old in excellent health $1,384 in out-of-pocket costs during the plan year, while an 85 year-old in poor health is expected to spend $4,099 on out-of-pocket costs. On average, the plans we observe enroll 672 people per county-year and are offered in 15 counties. To capture differences between plan offerings and consumer choices, the second column collapses the data by plan-year and weights the resulting 10,751 observations by total enrollment. The average plan chosen by a consumer has a lower premium and deductible, is higher quality as measured by the star rating, has lower copays, and lower out-of-pocket costs.

As our focus is on the way in which differences in the benchmark lead to differences in consumer

surplus across markets, it is useful to examine the variation in benefit design and enrollment by the level of the benchmark in the data. Table 2 presents the mean characteristics of plans across benchmark quartiles. We calculate benchmark quartiles at the market level to illustrate the fact that markets vary in their level of competition. As benchmarks increase, observable plan benefits generally improve. The mean premium is $607 in the first quartile, and $443 in the fourth. Dental coverage is included in 55% of plans in the lowest quartile, and 69% of plans in the highest. Estimated out-of-pocket costs for a 75-79 year-old in good health vary from $2,715 in the first quartile to $2,336 in the fourth.

These patterns are not always monotonic: the average star rating of a plan decreases from 2.42 in the first quartile to 2.15 in the second quartile before increasing to 2.31 and 2.64 in the third and fourth quartiles, respectively. These patterns reflect the fact that benchmarks are assigned to markets non-randomly as a function of average TM costs in previous years. Though the costs faced by private firms are likely different than average TM costs, they are certainly correlated and therefore the benchmark alone is an insufficient statistic for understanding the behavior of firms.

The non-monotonicity notwithstanding, these differences in benefit design are correlated with differences in enrollment. At the plan level, average enrollment increases from 269 in the first quartile to 1,039 in the fourth. The total market-level MA share increases from .149 to .259.

3.2 Medicare Current Beneficiary Survey

Our individual-level data on Medicare beneficiaries come from the Medicare Current Beneficiary Survey (MCBS), a rolling-panel survey of a nationally representative sample of beneficiaries sponsored by CMS and produced by Westat. The survey is designed to obtain a complete picture of Medicare beneficiaries and their interactions with the Medicare system. Participants are interviewed multiple times per year over three years, and responses are linked to CMS administrative and claims data to ensure accuracy.

We obtain the MCBS survey responses for 2008-2015.[5] We observe detailed demographic information, including income, age, sex, race, and education. Survey respondents self-report their health status, choosing between Excellent, Very Good, Good, Fair, and Poor. We transform these variables into indicators for each demographic group that match the groups captured by the plan-level OOPC estimates. We also observe the respondent's county of residence and variables identifying

[5]The MCBS did not release data for 2014.

Table 1: MA Plan Summary Statistics

Variable	Unweighted mean	Enrollment-weighted mean
Annual premium ($)	509	395
Annual Part B premium reduction ($)	30.29	22.63
Deductible ($)	94.03	50.80
Out-of-pocket limit ($)	3,846	3,944
Star rating	2.40	3.31
Supplemental coverage indicators		
Prescription drug	.764	.954
Dental	.608	.659
Vision	.890	.940
Plan type indicators		
HMO	.450	.791
PPO	.276	.156
FFS	.273	.053
Copays		
Primary care ($)	11.31	9.00
Specialist ($)	27.81	25.55
Hospital stay ($)	237	207
Coinsurance		
Primary care (%)	.476	.132
Specialist (%)	.358	.122
Hospital stay (%)	.059	.050
Selected CMS-estimated out-of-pocket costs		
65-69 year-old, excellent health ($)	1,384	1,278
75-79 year-old, good health ($)	2,450	2,123
85+ year-old, poor health ($)	4,099	3,581
Enrollment	672	20,688
Number of counties covered	14.6	8.47
Obs. level	Plan-County-Year	Plan-Year
Obs.	50,593	10,751

Table 2: Mean plan characteristics by market-level benchmark quartile

Variable	0-25th	26-50th	51-75th	76-100th
Annual premium ($)	607	545	512	443
Annual Part B premium reduction ($)	8.01	21.45	31.79	43.83
Deductible ($)	97.40	87.30	102.39	89.79
Out-of-pocket limit ($)	4,334	4,208	3,446	3,766
Star rating	2.42	2.15	2.31	2.64
Supplemental coverage indicators				
Prescription drug	.735	.745	.749	.801
Dental	.554	.565	.571	.691
Vision	.852	.884	.869	.929
Plan type indicators				
HMO	.249	.394	.403	.610
PPO	.364	.298	.282	.221
FFS	.387	.307	.315	.168
Copays				
Primary care ($)	13.25	11.47	12.23	9.62
Specialist ($)	31.39	28.48	28.68	25.15
Hospital stay ($)	277	244	249	204.73
Coinsurance				
Primary care (%)	.658	.797	.388	.264
Specialist (%)	.417	.417	.411	.246
Hospital stay (%)	.057	.091	.062	.033
Selected CMS-estimated out-of-pocket costs				
65-69 year-old, excellent health ($)	1,549	1,452	1,313	1,336
75-79 year-old, good health ($)	2,715	2,612	2,339	2,336
85+ year-old, poor health ($)	4,699	4,324	3,799	3,977
Annual benchmark ($)	8,860	9,543	9,820	11,035
Plan-level enrollment	269	550	535	1,039
Market-level MA share	.149	.211	.235	.259
Obs.	6,697	11,658	15,268	16,970

Note: An observation is a county-year-plan. The benchmark quartile is defined at the market (county-year) level.

their choice of MA plan, if any.[6]

The MCBS does not cover every county in the country. Instead, it employs a multi-level clustered sampling procedure to construct a nationally representative sample of Medicare recipients. The clustering of the samples ensures that within a geographic area, there is considerable variation in demographics and plan decisions (though we use CMS enrollment data to form product shares). The data include sampling weights which we use throughout our analysis to transform our results into a nationally representative form.

To ensure that we can model our MCBS micro-data as draws from the same process that generates our plan-level enrollment data, we must drop some responses. We exclude any individuals who were eligible for Medicaid at any point during the year, and those with missing address information. After applying these exclusions, the sum of the MCBS sample weights used in our analysis differs from the sum of eligibles nationwide in our CMS data by less than 1%.

There are a number of categories of individuals in the MCBS data who do not fall into the standard set of Medicare beneficiaries often studied in the literature (i.e. age 65-plus retirees without outside insurance) but nonetheless are eligible to purchase MA plans and are therefore included in CMS' enrollment files as such. These categories include individuals with employer-provided insurance plans, those whose original Medicare eligibility was not related to age, those diagnosed with end-stage renal disease (ESRD), and those who choose not to enroll in Part A and Part B for the full year. As these individuals do purchase MA plans with some frequency, we cannot exclude these groups without violating our model of the data-generating process. We instead include them throughout the analysis and create "administrative status" identifiers to capture their differential behavior. As the presence of these individuals in our analysis changes the distribution of consumer surplus, we also report results with reference to the "standard analysis" set, which excludes these groups, to ease comparisons with other analyses of the Medicare system.

MA is not the only option available to consumers who wish to obtain health insurance beyond what is available in the TM system. Variation in the price of other options may make MA plans more or less attractive. To capture this variation, we focus on Medicare supplemental insurance, more commonly referred to as Medigap. Medigap is often purchased by TM beneficiaries and pays for care that is not covered under the TM benefit design. For example, TM covers 80% of the cost of physician services, and Medigap coverage may be used to pay for the remaining 20%. The

[6] In some years, the MCBS does not report the plan choice directly and instead reports which firm the individual has chosen, along with information about plan features and any premiums paid by the individual. We match this data to the plan data to identify each individual's choice.

benefit structure of Medigap plans is standardized by CMS and indexed by letters: all Medigap "Plan A" policies have the same benefit structure regardless of the insurer selling the plan. For each person, we obtain the rate for Medigap Plan C offered by United Healthcare for their age and gender in their location that year from Weiss Ratings. Plan C covers most of the coinsurance and deductibles that beneficiaries are responsible for under TM and is the most popular Medigap plan.[7]

Summary statistics on our 58,444 individual-year observations covering 727 counties, 2,947 county-year markets, and 32,993 unique individuals are reported in Table 3. The first two columns report the mean and standard deviations for variables across all observations, weighted by the MCBS sample weights. The mean age of individuals in our data is 73. Slightly more than half of our observations are of females. Over 90% of individuals are coded as White, with 7.9% Black and 1.0% Hispanic, the largest categories defined by CMS. Over three-quarters of the individuals surveyed self-reported "Good" or better health. 24% of individuals report having college degrees and 18% did not graduate high school. 33% of Medicare beneficiaries receive some insurance from a current or previous employer, and 16% are eligible for reasons other than turning 65. The second set of columns splits the data into those enrolled in an MA plan and those enrolled in TM. On average, MA enrollees have lower income, are less likely to be White, have lower educational attainment, and are in slightly better health. The last set of rows shows that individuals with any employer-provided insurance and those with ESRD are less likely to be enrolled in MA.

The third and final set of columns of Table 3 illustrates the rolling-panel nature of our data and focuses on individual-year observations for which the individual was included in the previous year's survey and was enrolled in TM at that time. 23,278 observations fall into this category. The data are then split into those who switched from TM to an MA plan (4.3% of observations), and those who remained on TM. For most variables, those who switched look similar to the larger group of MA enrollees. The biggest differences come in the health status distribution – although those who are enrolled in an MA plan are healthier than those on TM, switchers from MA to TM are on average healthier still.

[7]Massachusetts, Minnesota, and Wisconsin have alternative plan definitions; in those states we use the rate for the plan closest to Plan C. Additionally, United Healthcare did not offer plans in New York during our study period. For those individuals, we averaged the rates offered by all other insurers.

Table 3: Medicare beneficiary micro-data summary statistics

	All observations		By MA enrollment		TM → MA status	
Variable	Mean	Std. dev.	MA	TM	Switchers	Non-switchers
MA enrollment indicator	.226	.418	1	0	1	0
Income	$46,535	68,736	38,168	48,982	39,707	44,387
Age	73.4	10.0	73.7	73.3	73.5	75.1
Outside Good Price	$2,390	590	2,487	2,361	2,398	2,382
Demographic indicators						
Female	.536	.499	.549	.532	.540	.539
Black	.079	.270	.093	.075	.084	.068
Hispanic	.010	.100	.016	.008	.013	.007
Education indicators						
Bachelor's degree or higher	.236	.425	.178	.253	.211	.232
Attended college	.304	.460	.310	.302	.304	.291
Graduated high school	.291	.454	.310	.286	.298	.301
Health status indicators						
Excellent	.173	.379	.169	.174	.175	.160
Very Good	.309	.462	.320	.306	.318	.303
Good	.302	.459	.302	.302	.313	.315
Fair	.155	.362	.158	.154	.145	.161
Poor	.060	.237	.051	.062	.049	.061
Administrative status indicators						
Employer-provided insurance	.330	.470	.011	.423	.000	.434
Non-aged eligibility	.159	.365	.164	.157	.188	.160
ESRD	.007	.081	.004	.007	.002	.007
Full-year Part A/B enrollee	.898	.302	.975	.876	.970	.898
Obs.	58,444		13,225	45,219	994	22,284

Note: An observation is a person-year. Statistics reported in this table are weighted according to sampling weights provided by the Medicare Current Beneficiary Survey. Demographic categories are defined by CMS administrative data. Observations are eligible for inclusion in the last column if they come from beneficiaries who were included in the MCBS in the previous period, and if they were enrolled in TM in that period.

4 A Model of Medicare Advantage

In this section, we present our model of supply and demand for Medicare Advantage which captures the key features of the institutional background described in Section 2. Our model of demand for MA plans is inspired by Goolsbee and Petrin (2004) and leverages the rich individual-level information in the MCBS data and the detailed plan characteristic data from CMS. The demand model captures beneficiary preference heterogeneity along multiple dimensions. In this sense, our model is similar to the demand model in Aizawa and Kim (2018) although we estimate the parameters using a different approach. On the supply side, we model firms as maximizing multi-product profits by choosing prices and product characteristics while taking into account preference heterogeneity. We abstract from the entry and exit decisions, and the decision over the number of plans offered. As we do not observe the risk scores used by CMS in our micro-data, we write our model of costs in terms of an average risk individual.[8]

Each period, denoted by t, begins with the government choosing a benchmark rate B_m for each market m and outside supplemental insurance firms determining a price schedule $\{p_0\}_m$. Firms $f \in \mathcal{F}_{mt}$ publicly observe costs and then simultaneously choose for each of their plans $j \in \mathcal{J}_f$ product characteristics $\{X, oopc, \xi\}_j$ and prices p_j. Consumers, denoted by i, then choose a single plan or the outside option of TM during an enrollment period. We present the details of each of these steps in reverse, starting with the consumer's demand decision and moving to the firm's product-level decisions.

4.1 Demand

Consumers have demographic characteristics z_i. They belong to an income bracket g and an age-and-health-status type h. For notational simplicity, let g_i and h_i be indicators which are equal to one if consumer i is a member of group g or h, respectively.

Consumers enter the enrollment period already a member of some plan k. We define three indicators S_{sij} to capture the way in which consumer behavior today might depend on their previous enrollment. Let S_{1ij} be equal to one if the consumer's previous choice k is the outside good – we refer to this as the *Medicare-to-MA* switch indicator. Let S_{2ij} – the *MA Interfirm* indicator – be one if the consumer's previous choice k is offered by a different firm than the plan j under consideration. Finally, let S_{3ij} – the *MA Intrafirm* indicator – be one if the consumer's previous

[8] Related work has found that the current implementation of the risk-adjustment system effectively reduces incentives to cream-skim for cost reasons (Newhouse et al., 2015).

choice k was offered by the same firm as j but k and j are different products.

Let u_{ijmt} denote the consumer's utility from enrolling in a particular plan j. Dropping the market and time subscripts, the choice specific utility for MA plans is given by:

$$\begin{aligned} u_{ij} = &\left(\alpha + \sum_g \alpha_g g_i\right) p_j + \beta_h \sum_h oopc_{hj} h_i + \sum_s \beta_s S_{sij} \\ &+ \sum_s \sum_h \beta_{sh} S_{sij} h_i + \beta_z z_i + \beta X_j + \xi_j + \epsilon_{ij} \end{aligned} \tag{2}$$

In this equation, α represents the price sensitivity common to all consumers, while the α_g parameters represent the individual-level variance in that price sensitivity by income group g. β_h captures the impact of expected out-of-pocket costs for that consumer's specific health group h_i. The impact of switching is parameterized by β_s, and is allowed to vary by health group through β_{sh}. The parameter β_z captures overall differences in the tastes of consumers with different demographic characteristics for MA versus TM. The vector β captures mean tastes for plan characteristics X_j.

As is standard in random utility demand systems, we decompose the unobservable portion of utility into two components. The term ξ_j represents the portion of unobserved (to the econometrician) utility for a plan that is common across individuals in a market. ϵ_{ij} represents the idiosyncratic taste that consumer i has for plan j which is assumed to be independently drawn according to a Type-I extreme value distribution.

Though we do not model the purchase of Medicare supplemental plans directly, we allow the utility of the outside good to vary with the price schedule of supplemental insurance $\{p_0\}$, which we denote at the individual level by p_{0i}, though the following equation:

$$u_{i0} = \beta_0 p_{i0} + \epsilon_{i0} \tag{3}$$

We normalize the expected utility of the outside good to 0 by subtracting $\beta_0 p_{i0}$ from each of the products j.

We include switching costs in our framework due to the evidence presented in Table 3 and the consistent empirical finding of significant inertia in plan enrollment (Nosal (2012) and Aizawa and Kim (2018)).[9] Enrollees in TM face a different set of rules, benefits, and provider networks than

[9]Like Aizawa and Kim (2018), we do not model the dynamic decision of the enrollee given the switching costs for several reasons. First, such analysis is computationally intensive. Second, it likely requires assuming that Medicare beneficiaries dynamically optimize in accordance with a standard model of neoclassical preferences with a discount

in MA and those rules, benefits, and networks vary across insurers and even across plans within an insurer. Switching between plans entails learning about the new plan's administrative procedures and network. In addition, Medicare beneficiaries are automatically re-enrolled in their previous plan (assuming that plan is still available) if they take no action during their open enrollment period. Thus, it is virtually costless to enroll in the previous year's choice.[10] Similar to Handel (2013), we model these costs with a direct utility impact through the $\sum_s \beta_s S_{sij}$ term.

Following Berry et al. (1995), it is useful to rewrite the utility u_{ij} into a mean

$$\delta_j = \alpha_0 p_j + \beta X_j + \xi_j$$

and an individual-specific deviation from that mean

$$\mu'_{ij} = \sum_g \alpha_g g_i p_j + \sum_h \beta_h oopc_{hj} h_i + \sum_s \beta_s S_{sij} + \sum_s \sum_h \beta_{sh} S_{sij} h_i + \beta_z z_i - \beta_0 p_{i0} + \epsilon_{ij}.$$

Let $\mu_{ij} = \mu'_{ij} - \epsilon_{ij}$. Then, given our distributional assumption on ϵ_{ij}, we can write the probability that consumer i chooses plan j (i.e. the share function) in a common logit form:

$$s_{ij} \equiv \Pr(i\text{chooses}j) = \frac{\exp(\delta_j + \mu_{ij})}{1 + \sum_{k \in \mathcal{J}_m} \exp(\delta_k + \mu_{ik})} \tag{4}$$

Given our utility specification, the consumer surplus for an individual i, given prices and product characteristics, can be written as

$$CS_i = E[\max_j u_{ij}]/\alpha_i = \frac{1}{\alpha_i} \ln \left(\sum_j \exp(\delta_j + \mu_{ij}) \right) \tag{5}$$

where α_i is the price sensitivity of individual i (Small and Rosen, 1981). This form of consumer surplus, also known as the "inclusive value" form, can be thought of as the amount individual i would be willing to pay to have the choice of MA plans available in their county, relative to a world in which TM was the only option. It does not include the benefit of TM itself versus no insurance.

factor close to one. However, recent work has shown that in different but related settings that model does not explain Medicare beneficiary behavior well (e.g Dalton et al. (2018)). Third, (Nosal, 2012) has estimated such a model and found extremely high (perhaps implausibly so) switching costs. Fourth, our estimation approach does a good job of capturing inertia that is salient for our counterfactual analysis.

[10] If an enrollee's plan is not available for the upcoming year and if the enrollee does not make an active plan selection, CMS allows the insurer to move the enrollee to the most similar plan the insurer will be offering. We obtain and incorporate these 'crosswalk' enrollments into our empirical analysis.

4.2 Supply

Firms start each period with common knowledge of the distribution of consumer demographics, including which consumers are enrolled in which plans, and each others costs, including cost shocks. Firms also have common knowledge of the per-enrollee subsidy they receive B_j, which is a function of the benchmark, the firm's prior performance, and the plan bid.[11] Each firm is assumed to choose a vector of prices and product characteristics for their products to maximize profit, taking as given the other prices and product characteristics in the market.

To form the profit function, we model marginal costs as log-linear in observed characteristics r_j (which includes X_j and $oopc_j$, but may also include other components) and some unobserved component ω_j following

$$\ln(mc_j) = \gamma_r r_j + \omega_j. \tag{6}$$

Given these marginal costs and the demand model presented in the previous section, we write the profit maximization problem for firm f as

$$\max_{p_f, X_f, oopc_f, \xi_f} \pi_f = \sum_{j \in \mathcal{J}_f} \left[(p_j + B_j - mc_j) \int_i s_{ij}(p, X, oopc, \xi; \theta) di \right]. \tag{7}$$

This maximization problem is written with respect to the set of vectors $\{p_f, X_f, oopc_f, \xi_f\}$ to indicate that firms choose these attributes for all of their products simultaneously.[12] The arguments to the s_{ij} function are written without subscripts to denote the dependence of s_{ij} on the characteristics of all of the products in the market; θ is the vector of demand system parameters.

5 Estimation

In this section, we take our model of competition in Medicare Advantage to the data and estimate its parameters. We start by estimating the demand system, which allows us to recover marginal cost and product-specific unobservable characteristics. We then regress our estimated marginal cost on product characteristics.

[11] We do not model the bidding process explicitly, though we account for it in our counterfactual approach. See Section 6.

[12] In practice, firms do not have direct control over each of the out-of-pocket cost measures. Instead, CMS calculates them as a function of the plan's benefit design, including features we do not observe. The plan characteristics we observe capture 38% of the variation in the out-of-pocket cost estimates, and therefore it is more practical to model these characteristics as separate choices made by firms rather than, say, as functions of other observables.

We estimate the parameters of the demand system following the approach of Goolsbee and Petrin (2004). We split the parameters to be estimated into two categories. First, we estimate the set of parameters which capture the impact of individual-level variation on choice-specific utility – those parameters that define μ_{ij} – with a maximum likelihood approach. We then estimate the parameters common to individuals – those parameters that define δ_j – with an instrumental variables approach.

Let $\theta_I = \{\alpha_g, \beta_h, \beta_s, \beta_{sh}, \beta_z, \beta_0\}$ be the set of parameters which comprise the individual-specific portion of demand. For a given candidate value of those parameters, $\tilde{\theta}$, we use the Berry (1994) contraction to find the unique set of product fixed effects $\delta_j(\tilde{\theta})$ that match predicted shares to observed market shares. Let C_{ij} be an indicator variable that is equal to one if person i chose product j.[13] We form the likelihood function

$$L_i(C_{ij}; \tilde{\theta}, \delta(\tilde{\theta})) = \Pi_j s_{ij}^{C_{ij}}, \tag{8}$$

where s_{ij} is given by Equation 4. In the first stage of our estimation procedure, we maximize the weighted log likelihood function

$$l(C; \tilde{\theta}) = \sum_i ln(L_i) w_i, \tag{9}$$

where w_i is the sample weight of each observation taken from the MCBS.

At the point estimate, $\hat{\theta}$, we store the unique $\hat{\delta}_j$ recovered by the Berry contraction mapping. In the second step, we regress these $\hat{\delta}_j$ on observable product characteristics according to the terms in our demand model

$$\delta_j = \alpha_o P_j + \beta X_j + \xi_j \tag{10}$$

where ξ_j is a product-specific unobservable.

5.1 Identification

Since ξ_j is observed by consumers but not by us, it is likely to be correlated with the plan's product share and potentially the price and other characteristics of the plan. To identify the coefficients

[13] Implementing this variable leads to two complications. First, the plan choice variables in the MCBS vary from year to year and we are not always able to recover a unique choice. Second, we model utility as a function of past enrollment, and the rolling panel design ensures that in each year, a significant fraction of our observations have no data about past enrollment. We solve these problems by eliminating as many plans as possible given what we do observe and taking draws from the conditional distribution of product shares.

α_0 and β, we make the common assumption that the unobservable product characteristic ξ_j is uncorrelated with the observed product characteristics X_j. This implies both that β is identified and that instruments formed as functions of the X_j are valid for price — we use the summation instruments of Berry et al. (1995).

The biggest threat to this assumption in our model is the fact that firms choose ξ_j and X_j simultaneously as a function of the benchmark, and thus these characteristics may be correlated through mutual dependency upon the benchmark. In practice, however, our data captures a wealth of detail about the financial characteristics of the plan benefits, and thus the largest components of ξ are likely to be network breadth, advertising, and unobservable quality. These components are unlikely to change from year-to-year in response to small changes in the benchmark, and thus it is potentially reasonable to assume that ξ_j is uncorrelated with B_j. As we do not use the benchmark directly in our estimation procedure, we can test this assumption by estimating $\hat{\xi}_j$ and computing its empirical correlation with B_j. Additionally, some of these components of ξ_j are likely to be constant across plans within a given firm[14], and so we include firm fixed effects in our primary specification to capture these components and reduce the extent to which ξ_j might be correlated with X_j. These fixed effects have the additional benefit of avoiding the need to model changes in the definition of the star rating used by CMS that are unrelated to the underlying true quality level of the firm.

The conditions covered by TM, and thus the coverage requirements for MA plans, change from year to year. This suggests that there may be some component of marginal cost which changes from year to year but is constant across firms. Additionally, state-level policy differences suggest that firms face some state-specific costs. Following the logic of Hausman et al. (1994), we use the average prices in other counties as additional price instruments. However, since a single plan is often offered at the same price in several geographically contiguous counties, we calculate this instrument using the set of "non-contiguous counties" – that is, for a plan in a county, we instrument the price with the average price of all plans in counties of the same state which do not share borders with the county under consideration.

5.2 Demand results

Parameter estimates for the first stage of the demand estimation procedure are in Table 4 and Table 5. Table 4 presents parameters representing the impact of various demographic characteristics

[14]Indeed, Aizawa and Kim (2018) make the assumption $\xi_j = \xi_f$ for all plans offered by the same firm.

Table 4: Demand parameter estimates for income groups, out-of-pocket costs, and demographic characteristics

Variable	Coefficient	Std. Err.
Income-level price effects (per $1000)		
Medium income	0.153	0.039
High income	0.089	0.042
Estimated out-of-pocket costs (per $1000)	-0.038	0.024
Medigap price (per $1000)	0.218	0.046
MA × Demographic characteristics		
Age	1.704	0.153
Age^2	-0.118	0.010
Female indicator	-0.035	0.028
Black indicator	0.317	0.055
Hispanic indicator	0.130	0.132
Graduated high school	-0.078	0.043
Attended college	-0.165	0.043
College degree or higher	-0.514	0.047
Administrative indicators		
Has employer-provided insurance	-3.649	0.089
Non-aged eligibility	0.318	0.048
ESRD diagnosis	-0.933	0.196
Full year Medicare enrollment	1.778	0.061
Weighted Log Likelihood	-58,148	
Obs.	58,444	

on the demand for MA plans, while Table 5 presents the switching cost parameter estimates. In the first panel of Table 4 the estimates indicate that high and medium income consumers are less price sensitive than low income consumers. As expected, out-of-pocket costs create disutility, while an increase in the cost of Medigap insurance drives consumers towards MA plans. The estimates on our demographic parameters indicate that non-white individuals have a stronger preference for MA plans, as do individuals with lower levels of education. Our administrative indicators enter with appropriate signs and reasonable magnitudes.

The three panels of Table 5 display the switching costs parameters for Traditional Medicare to MA switching costs, inter-insurer switching costs and intra-insurer switching costs, respectively. We interact each switching cost term with indicators for self-reported health status, with "Poor" as the excluded group. The highest switching costs are incurred by consumers switching from TM to MA, with slightly lower costs incurred by those switching between firms and still lower costs for those switching between individual plans offered by the same insurer. These results suggest

Table 5: Switching cost parameter estimates

Variable	Coefficient	Std. Err.
TM-to-MA switch ×		
Constant	-2.892	0.088
Excellent health	0.193	0.102
Very good health	0.258	0.094
Good health	0.146	0.084
Fair health	0.132	0.078
Inter-firm switch ×		
Constant	-1.627	0.129
Excellent health	0.011	0.152
Very good health	0.113	0.140
Good health	0.032	0.139
Fair health	-0.007	0.147
Intra-firm switch ×		
Constant	-0.515	0.152
Excellent	-0.089	0.177
Very good	0.006	0.165
Good	-0.005	0.165
Fair	0.033	0.173
Weighted Log Likelihood	-58,148	
Obs.	58,444	

Note: The omitted group is those in "Poor" health.

the primary component of switching costs is the disutility of changing providers. Under TM, beneficiaries have access to virtually any provider while MA plans generally use a tightly defined network of providers. Of course, different MA plans include different providers in their networks. However, different plans under the same contract generally deploy the same network of providers. While health status – particularly being in better than "Poor" health – appears to have an impact on the costs of switching from TM to MA, the interaction indicators do not enter with significance for either of the other switching cost types.

The second stage results are in Table 6. The first column presents the OLS estimates and the instrumental variable results are displayed in column (2). Consistent with OLS estimate on price being biased towards zero, the IV price coefficient is much large in magnitude than the OLS. For this reason, we concentrate our attention on the IV specification, which also includes firm-level fixed effects to control for those components of plan quality which are constant across a firm.[15]

[15]Though we collect star ratings, which are a measure of plan quality, we do not include them in this specification for two reasons. First, they are generally calculated at the firm level, and so are already picked up by firm fixed effects. Second, CMS changed the definition of the ratings several times throughout the sample period, substantially increasing the noise in the measure.

The second stage parameter estimates are, in general, sensible. Most premiums are equal to zero, but for the plans with a positive premium the average plan elasticity is -2.66.[16] The average semi-elasticity of increasing premiums $1 is .3 percent. The implied price sensitivity of the demand estimates are similar to estimates from the literature. For example, using an earlier sample period, Aizawa and Kim (2018) estimate an average MA semi-elasticity of .75 percent. We also test the null hypothesis that ξ_j is uncorrelated with B_j. The empirical correlation coefficient is -0.0025, and so we cannot reject the null hypothesis.

The presence of CMS-estimated out-of-pocket costs in the first stage, which are in part a function of the variables included here, changes the interpretation of the parameters on these plan characteristics. The parameters estimated here are not the direct effect of these features *per se*, but rather their impact on utility holding the total expected out-of-cost expenditures constant. In that light, our results indicate that consumers prefer to incur out-of-pocket costs on copays and coinsurance, rather than deductibles. Consumers prefer expenditures on plan benefits to reductions in Medicare Part B premiums. While prescription drug coverage is valued by consumers, vision and dental coverage isn't. PPOs, which generally have broader networks than HMOs, are preferred to the latter plan type.

With the first and second stage estimates in hand, we can transform the switching cost parameter estimates into monetized values of the implied switching costs. These monetized values are shown in Table 7. The cost incurred by an average medium-income individual switching from TM to MA is $918, which is almost twice the mean annual premium in our data, while the same average individual switching between plans within an MA firm incurs a cost of only $177, about a third of the average annual premium.

To illustrate the importance of these switching costs on the utility of MA, and therefore the surplus received by consumers, we simulate, for each market, a female beneficiary with a high-school education, medium income, and Good health entering the Medicare program at age 65. Since she has not previously selected TM or a MA plan, she starts with no switching costs. Each year, she faces the set same of plan choices (taken to be the set of choices available in 2015, the most recent year in our data), though the particular product characteristics (i.e. the expected out-of-pocket costs) she faces change over time as she ages. We simulate her choice probabilities and calculate

[16] An important issue in the recent Aetna-Humana proposed merger was the rates of substitution between MA and traditional Medicare if the carriers were to increase premiums. We find that a $100 increase in all plan premiums (or $8.33 per month), holding everything else constant, results in a 5.6% shift in enrollees from MA to traditional Medicare.

Table 6: Second stage demand parameter estimates

Variable	OLS	IV
Annual premium (per $1000)	-.116 (.016)	-3.11 (.452)
Part B reduction (per $1000)	-.000 (.000)	-.008 (.003)
Deductible (per $1000)	-.055 (.323)	-1.50 (.157)
Out-of-pocket limit (per $1000)	.020 (.004)	-.094 (.028)
Coverage indicators		
Prescription drugs	.539 (.023)	2.59 (.177)
Vision	-.199 (.032)	-.952 (.341)
Dental	-.133 (.020)	-.207 (.177)
Copays		
Primary doctor	-.014 (.002)	.109 (.017)
Specialist	.016 (.001)	.067 (.007)
Hospital stay (per $1000)	.040 (.060)	.432 (.584)
Coinsurance		
Primary doctor	.015 (.005)	.055 (.012)
Specialist	-.014 (.006)	.145 (.021)
Hospital stay	.015 (.011)	-.040 (.023)
HMO indicator	.256 (.026)	1.72 (.577)
PPO indicator	.433 (.032)	4.22 (.389)
Fixed effects	None	Firm-level
Mean implied elasticity (if < 0)	-.018 (.036)	-2.66 (1.78)
Mean ds_j/dp_j	-.000 (.001)	-.033 (.060)
Observations	50,593	50,593

Note: Robust standard errors are in parentheses. Observations are at the year-plan-county level.

Table 7: Implied dollar switching costs

Switch type	Low income	Medium income	High income
Medicare-to-MA	$877	918	894
Between-firms	510	534	521
Within-firm	167	177	176

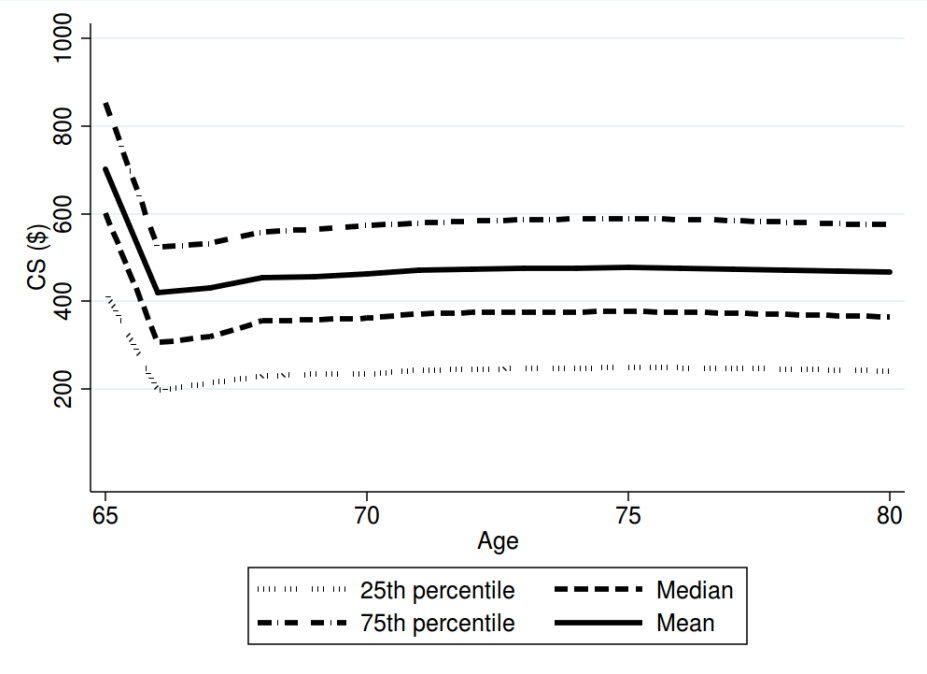
Note: These dollar switching costs are calculated by combining the our estimated price and switching cost parameters for each person and then averaging the resulting per-individual cost across income groups.

the expected consumer surplus she receives each year per Equation 5. Figure 3 shows the results of the simulations by illustrating the mean, median, 25th and 75th percentiles of surplus taken across all markets in 2015. The surplus cost of switching is seen in the drop from the first year of Medicare eligibility to subsequent years. In the first year, with no switching costs, the average annual consumer surplus across markets is $701. By age 70, the average consumer surplus has decreased to $462. The effect of switching costs swamps the small variations in the distribution of consumer surplus after age 66. As our simulated consumer moves in and out of out-of-pocket-cost groups, her consumer surplus varies slightly.

Given a set of demand parameters, we can estimate plan-level marginal costs from the first-order conditions for profit maximization (Equation 7). Just as ξ_j, the product-specific unobservable component of demand, is likely to be correlated with p_j, it is also likely to be correlated with ω_j, the product-specific unobservable cost component. However, as mentioned above, given our estimates of the demand parameters and the product shares in the data, we can estimate a $\hat{\xi}_j$ for each plan. We estimate the parameters of Equation 6 by assuming that any remaining unobservable components of cost are uncorrelated with ξ_j and our observables. As the out-of-pocket cost measures for different age and health groups are highly collinear we include the median out-of-pocket cost measure, the measure for a 65-69 year old in Good health – the results are nearly identical when using alternative out-of-pocket cost measures.

The results are presented in Table 8. We include firm and county fixed effects to concentrate on the contribution of different plan-level design parameters to the cost structure and to account for differences in costs which stem from geography alone. The major cost sharing plan design parameters, including the out-of-pocket costs, deductibles, and out-of-pocket limits, enter with the correct sign. Offering prescription drugs increases costs by 3.8%. The coefficient on the demand unobservable is positive and significant at the 1% level, suggesting that specifications which did not

Figure 3: Simulated MA consumer surplus for a new Medicare beneficiary



This graph depicts the distribution of consumer surplus received by a simulated female Medicare beneficiary with a high school education, medium income, and Good health across all markets.

Table 8: Marginal cost parameter estimates

Variable	$\ln(mc_j)$
Deductible (per $1000)	-.016 (.001)
OOP limit (per $1000)	-.003 (.000)
OOP cost estimate (per $1000, 65-69, Good)	-.011 (.000)
Demand unobservable (ξ_j)	.011 (.000)
Coverage indicators	
Prescription drugs	.038 (.001)
Vision	-.005 (.001)
Dental	.007 (.001)
Copays	
Primary doctor	.002 (.000)
Specialist	.000 (.000)
Hospital stay (per $1000)	-.034 (.002)
Coinsurance	
Primary doctor	-.000 (.000)
Specialist	.003 (.000)
Hospital stay	.003 (.002)
HMO indicator	-.038 (.001)
PPO indicator	.036 (.001)
Fixed effects	Firm, county
Observations	50,593

Note: OOP stands for "out-of-pocket." Robust standard errors are in parentheses. Observations are at the year-plan-county level.

take the correlation between demand and cost unobservables into account would be biased. As with the demand system, the copay and coinsurance parameters must be interpreted in the context of the overall out-of-pocket cost parameter. HMOs face costs approximately 7.4% lower than PPOs.

6 Optimal Geographic Variation in Medicare Advantage Benchmarks

To find the optimal subsidy schedule across geographies, we must calculate the impact of changes in the benchmark on consumer welfare and government expenditures in different markets. This requires capturing the way in which changes in the benchmark flow through the imperfectly competitive insurers' market to changes in prices and product characteristics. In this section, we present our optimal subsidy problem and then discuss our equilibrium approximation approach.

6.1 The optimal subsidy problem

Let B_m be the benchmark rate for a particular market, $CS_{im}(B)$ be the equilibrium consumer surplus for consumer i in market m when the benchmark rate is B, and $GovExp_m(B)$ be the equilibrium total government expenditures on the Medicare program (including both TM and MA). We consider three types of maximization problems given by

$$\max_{\{B_m\}} \sum_m \int_i CS_{im}(B_m)di \quad \text{s.t. } \sum_m GovExp_m(B_m) = \bar{M} \tag{11}$$

$$\max_{\{B_m\}} \alpha \sum_m \int_i CS_{im}(B_m)di - (1-\alpha)Var(CS) \quad \text{s.t. } \sum_m GovExp_m(B_m) = \bar{M} \tag{12}$$

$$\max_{\{B_m\}} \sum_m \int_i CS_{im}(B_m)di \text{ s.t. } B_m > \bar{B}_m \forall m \text{ and } \sum_m GovExp_m(B_m) = \bar{M} \tag{13}$$

The first equation maximizes consumer surplus subject to a government budget constraint. As we document below, the solution to this problem results in significant variation in the optimal benchmarks and the resulting consumer surplus. This variation may not be politically feasible to implement. For this reason, Equation 12 includes a penalty for the variance in consumer surplus across individuals with $\alpha \in [0,1]$ determining the relative weight that is placed on consumer surplus versus the variance in consumer surplus. Currently, the CMS benchmark formula includes floors on the benchmark amount suggesting that is the preferred policy solution in areas where the FFS spending formula results in low benchmarks. We therefore also solve for the optimal benchmarks using the objective function in Equation 13, which includes the constraint that all benchmarks must be set above some floor level.

Solving these problems requires calculating both the CS and $GovExp$ functions. Given prices and product characteristics as a function of B_m, CS as a function of the benchmark as used here can be derived from Equation 5. The government's expenditures can be derived from the plan-level share function (Equation 4) with

$$GovExp(B_m) = \sum_j \int_i \left(s_{ij} Payment_{jm} + (1 - s_{ij}) TM_m\right) di \tag{14}$$

where $Payment_{jm}$ is the plan-level payment of Equation 1, s_{ij} is the share function of Equation 4, and TM_m is average TM spending in the market in 2015, which we obtain from CMS. To account

for differences in observed risk across counties which may come from the MCBS sampling process, as opposed to true differences in risk across geographies, we normalize all spending to the average risk level — in other words, we set $RiskAdjustment_{it} = 1$ in Equation 1 for all individuals in all markets. Since the benchmark and rebate quality adjustment terms, ϕ_{jt} and λ_{jt}, are determined by previous firm performance, we take them as given.

6.2 An equilibrium approximation approach to counterfactual analysis

The key inputs to both CS and $GovExp$ are the set of prices and product characteristics in each market as a function of that market's benchmark. Given our estimates of the supply and demand parameters, a traditional approach to counterfactual analysis would proceed by solving for counterfactual Nash equilibria by varying the benchmark and searching for a fixed point in firms' best response functions, usually through analyzing first-order maximization conditions. For example, Fan (2013) uses this approach to solve for a new, post-merger equilibrium in prices and five product characteristics in a daily newspaper market with five competitors – a total of 30 first-order conditions. However, as each plan in our data has 47 observable characteristics and the average market has 14 plans, solving for even a single alternative equilibrium in one market involves a fixed point search with $14 * 47 = 658$ first-order conditions – which requires significantly greater computational effort (Daskalakis et al., 2009). In 25% of our markets, there are more than 1,000 first-order conditions to consider. To find an optimal policy, we must search over hundreds of equilibria in each of the 445 markets we observe in the 2015 data. As a consequence, a direct search for equilibria is computationally infeasible.

Our alternative "equilibrium approximation" approach combines policy function estimation with a reduced-dimension fixed-point problem and begins with the observation that our data contain substantial variation in the benchmark rate, which is the key primitive we wish to vary in our counterfactual analysis. Since the benchmark rate is set as a function of TM spending, it is exogenous from the perspective of an individual firm's actions. As seen in Table 2, variation in the benchmark rate comes with observed variation in product characteristics. In other words, the observed data trace out firms' policy functions for product characteristics. To predict counterfactual product characteristics, we regress observed characteristics on the benchmark and other regressors designed to capture relevant features of firm behavior, and use the fitted values at counterfactual benchmarks to update product characteristics. We then solve for prices in a Bertrand Nash game where firms face the demand system estimated in Section 5.

Given the dimensionality of product differentiation in this market, multiple equilibria are possible. Our approach pools data from different markets and so we adopt the Equilibrium Selection (ES) assumption of Bajari et al. (2007): we assume that firms across markets play the same equilibrium strategies with respect to the benchmark. While this assumption may be violated when markets are geographically isolated (Noel, 2007), our data include several large firms covering overlapping geographies. Indeed, our data is characterized by plans that cover multiple counties. We therefore conclude that the existence of isolated markets with widely disparate equilibrium behavior is unlikely.

Our approach builds on past efforts to use policy function estimation for counterfactual analysis. Goolsbee and Petrin (2004) study competition between cable and satellite television systems and use estimated functions for prices and product characteristics to calculate the welfare improvement brought about by the introduction of satellite TV. Benkard et al. (2018) build upon the policy function estimation approach of Bajari et al. (2007) to estimate strategic entry and exit behavior in the airline industry to simulate industry outcomes under counterfactual merger scenarios. We extend these efforts by combining our estimated policy functions with our demand model to solve for prices and calculate the welfare effects of changes in the government policy.

We start by simplifying the problem and define a 'price-adjusted' δ for each plan given by

$$\delta'_j = \delta_j - \alpha_0 P_j. \tag{15}$$

δ'_j includes the effects of both observables X_j and unobservables ξ_j and is easily calculated with our demand estimator. We regress δ' on the benchmark, the 'rank' of the plan with respect to the δ' within the firm, an indicator for whether or not the firm has the highest share in the market, the total number of plans in the market, and firm fixed effects. We repeat this exercise with each of the out-of-pocket cost terms on the left-hand side.

Given a benchmark for a market, we use our estimate of the policy functions to calculate the new δ and updated out-of-pocket costs for each plan in that market. We do not use the estimated policy at the proposed benchmark directly. Instead, we take the difference between the estimated policy at the proposed benchmark and the estimated policy at the observed benchmark and apply that difference to the observed policy. This increases the flexibility of our approximation and maintains plan-level heterogeneity without adding additional terms to our regressions. Given our result that ξ_j is uncorrelated with the benchmark, it also has the effect of keeping ξ_j constant across

counterfactual runs and independent of the benchmark, consistent with our identifying assumptions. We apply a similar approach to the other primitives and estimate marginal cost as a function of δ' and out-of-pocket costs, consistent with our modeling approach, and estimate the plan bid as a function of marginal cost, consistent with CMS bidding rules.

With the set of new product characteristics, marginal costs, and bids in hand, we must still calculate prices. As this reduces the number of first-order conditions to one per plan, it is straightforward to solve for equilibrium prices by iterating over the $\frac{\partial s_j}{\partial p_j}$ matrix. We implement this solver with a non-negativity restriction on prices to match the observed behavior in the data.

To summarize, given a proposed benchmark for a given county, we calculate CS_m and $GovExp_m$ as follows:

1. Use the estimated policy functions to calculate δ'_j and $oopc_j$ for each plan.
2. Use the estimated marginal cost and bid functions with the new product characteristics to calculate mc_j and bid_j.
3. Taking as given these new characteristics and costs, solve for the new equilibrium in prices.
4. Given the new prices and characteristics, calculate consumer surplus with Equation 5 and government spending with Equation 14.

This process can be used to calculate the impact of alternative policies across the maximization problems represented by Equations 11, 12, and 13. This is made more difficult in practice by the fact that both our objective function and our constraint are non-linear – and, in the case of our constraint, non-monotonic as well. As a consequence, we cannot easily solve for the constraint surface and instead satisfy the constraint numerically using a penalty function approach. We address the possibility of multiple local maxima by implementing a multistart procedure inspired by Rinnooy Kan and Timmer (1987). To reduce the computational burden, we take advantage of the separable nature of our problem (i.e. each market's outcome depends only on that market's benchmark) and calculate outcomes over a grid of benchmarks for each market and form spline approximations. We use these spline approximations to find candidate solutions quickly and then refine the most promising candidates using the numerical equilibrium solver.

We do not model the impact of a change in the benchmark on plan entry and exit due to the significantly increased computational requirement. However, incorporating entry and exit into the analysis should not significantly impact the results as the plans that enter and exit during our

sample period comprise a small average share of the total MA market and thus contribute very little to consumer surplus. The average share of exiters in the year before they exit is 0.62%. The average share of new entrants in the year they enter is 0.72% and the average share two years later is 1.1%. As our counterfactual benchmarks are within the support of the current benchmark distribution, the role (or lack thereof) of past entry and exit on consumer surplus should be informative about its impact on consumer surplus in the counterfactuals. Decarolis et al. (2015) use a similar logic in their analysis of the Medicare Part D market.

6.3 Counterfactual results

We apply our approach to the set of counties observed in the MCBS in 2015. Throughout our exercise, we use the sample weights provided by the MCBS to calculate average CS, GS, and the variance of CS when applicable. We proceed by presenting the results of our policy function estimation exercise and then detail the components of our CS_m and $GovExp_m$ functions that are generated by those policy functions for a selection of the counties in our data. We then present our primary exercise in which we maximize average consumer surplus, and finally discuss alternative objective functions.

6.3.1 Policy function estimation results

Table 9 reports the results of our policy function estimation exercise for the δ' term. Column (1) is a simple regression of δ' on the benchmark alone. This specification leads to a counterintuitive negative sign on the benchmark, though the fit is poor. Column (2) adds firm fixed effects. The R^2 rises from .034 to .307, though the negative sign on the benchmark remains. Column (3), our preferred specification, adds other regressors designed to capture the diversity of plans offered by firms. After controlling for the rank of the plan within its firm, an indicator for the firm that is leading in market share, and the number of plans in the market, the benchmark enters with a positive sign, and R^2 increases to .514.[17]

Table 10 repeats the exercise for each of the 30 out-of-pocket costs in our data. In this table, the rows correspond to individual out-of-pocket cost terms and the columns report the estimates and R^2 for a specification analogous to that of Column (3) of Table 9. Across all demographic categories, an increase in the benchmark leads to a decrease in the expected out-of-pocket costs for

[17]For comparison, Goolsbee and Petrin (2004) report an R^2 of .187 for their analogous exercise (Table 9), and Benkard et al. (2018) report pseudo-R^2s ranging from .125 to .999 across a range of similar exercises.

Table 9: Policy function estimation results for adjusted-δ term

Variable	(1)	(2)	(3)
Benchmark	-.541	-.248	.030
	(.014)	(.015)	(.013)
Plan rank			-.982
			(.013)
Leading firm indicator			1.30
			(.031)
# plans in market			-.031
			(.001)
Fixed Effects	None	Firm	Firm
Observations	50,593	50,593	50,593
R-squared	.034	.307	.514

enrollees. Within an age group, the largest reductions go to those in the worst health categories. R^2s range from .210 to .312.

To complete our set of firm primitives, we regress marginal cost on product characteristics and the plan bid on the marginal cost. For both, we include firm and county fixed effects. For consistency with our other policy function estimation exercises, the left-hand side variables are in levels.[18] The first column of Table 11 reports the parameter estimates for marginal cost. As with our previous marginal cost estimation exercise, we include only one out-of-pocket cost term for simplicity. Increasing the δ' leads to an increase in the marginal cost, and increasing the out-of-pocket costs decreases the marginal cost. The second column reports the parameter estimate for the plan bid. Firm bids are slightly lower than the marginal cost. Given the rules surrounding the bidding process – in particular, the requirement that bids largely reflect the cost of TM benefits – this result is sensible.

6.3.2 The components of surplus and expenditures

Figure 4 illustrates several components used in constructing the CS_m function for Travis County, TX (which contains Austin), Cook County, IL (which contains Chicago), and Worth County, GA (which is a rural county near Albany). We chose these counties for illustration due to their different sizes and the relatively typical nature of their counterfactual equilibria outcomes. The first graph of Figure 4 illustrates how the price-adjusted δ'_j, weighted by product share, changes as the benchmark

[18] Similar regressions in logs instead of levels produced nearly identical results.

Table 10: Policy function estimation results for out-of-pocket costs

Demographic Cat.	Benchmark	Plan rank	Leading firm	# plans	R-squared
Ages 0-64					
Excellent	-.050 (.003)	.113 (.002)	-.071 (.007)	-.004 (.000)	.279
Very good	-.018 (.004)	.142 (.003)	-.106 (.010)	-.008 (.000)	.231
Good	-.055 (.005)	.173 (.003)	-.140 (.013)	-.011 (.000)	.215
Fair	-.073 (.007)	.244 (.005)	-.195 (.018)	-.016 (.000)	.210
Poor	-.093 (.010)	.336 (.007)	-.268 (.025)	-.022 (.001)	.211
Ages 65-69					
Excellent	-.019 (.002)	.071 (.002)	-.062 (.006)	-.006 (.000)	.286
Very good	-.017 (.003)	.098 (.002)	-.083 (.008)	-.008 (.000)	.256
Good	-.063 (.005)	.145 (.003)	-.121 (.011)	-.010 (.000)	.230
Fair	-.087 (.007)	.237 (.005)	-.185 (.018)	-.014 (.000)	.222
Poor	-.162 (.010)	.341 (.006)	-.226 (.024)	-.012 (.001)	.245
Ages 70-74					
Excellent	-.028 (.002)	.079 (.002)	-.066 (.006)	-.006 (.000)	.285
Very good	-.031 (.003)	.104 (.002)	-.085 (.008)	-.007 (.000)	.253
Good	-.059 (.005)	.148 (.003)	-.126 (.011)	-.010 (.000)	.240
Fair	-.114 (.007)	.239 (.005)	-.181 (.018)	-.012 (.000)	.226
Poor	-.209 (.010)	.328 (.006)	-.207 (.024)	-.010 (.001)	.253
Ages 75-79					
Excellent	-.032 (.002)	.083 (.002)	-.056 (.006)	-.004 (.000)	.294
Very good	-.037 (.003)	.109 (.002)	-.085 (.008)	-.007 (.000)	.258
Good	-.060 (.005)	.156 (.003)	-.123 (.011)	-.009 (.000)	.240
Fair	-.091 (.007)	.236 (.005)	-.179 (.017)	-.013 (.000)	.228
Poor	-.088 (.010)	.314 (.006)	-.267 (.024)	-.023 (.001)	.240
Ages 80-84					
Excellent	-.039 (.002)	.085 (.002)	-.055 (.006)	-.004 (.000)	.300
Very good	-.054 (.003)	.111 (.002)	-.082 (.008)	-.006 (.000)	.270
Good	-.071 (.005)	.158 (.003)	-.130 (.011)	-.009 (.000)	.245
Fair	-.101 (.007)	.240 (.005)	-.174 (.017)	-.012 (.000)	.239
Poor	-.102 (.010)	.319 (.006)	-.259 (.024)	-.021 (.001)	.250
Ages 85+					
Excellent	-.054 (.003)	.105 (.002)	-.066 (.007)	-.002 (.000)	.312
Very good	-.062 (.003)	.122 (.002)	-.085 (.008)	-.004 (.000)	.285
Good	-.076 (.005)	.164 (.003)	-.134 (.012)	-.008 (.000)	.249
Fair	-.101 (.007)	.240 (.005)	-.181 (.017)	-.013 (.000)	.241
Poor	-.083 (.010)	.307 (.006)	-.281 (.024)	-.022 (.001)	.248

Table 11: Policy function estimation results for marginal cost and bids

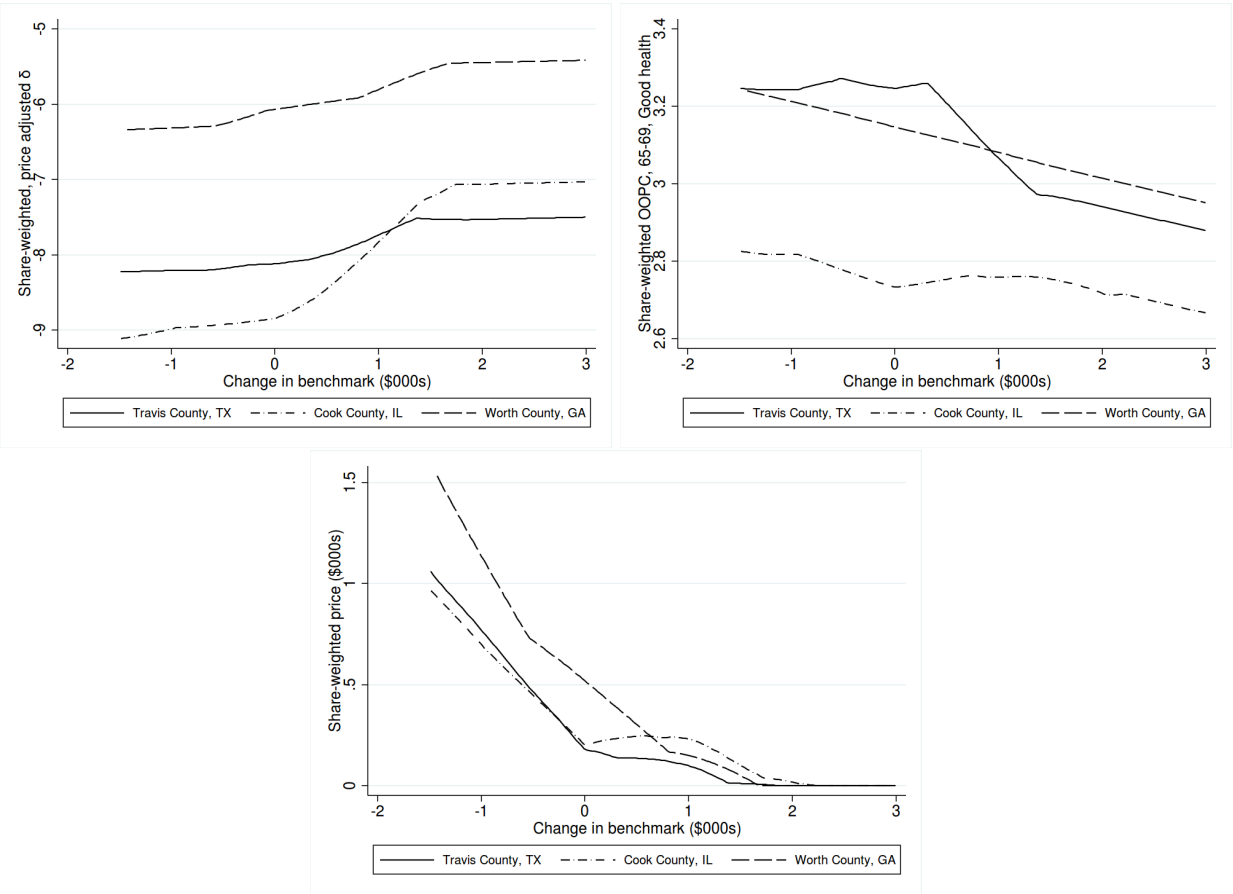
Variable	Marginal cost	Bid
δ'	.146	
	(.001)	
65-69 Good	-.112	
	(.003)	
Marginal cost		.918
		(.006)
Fixed Effects	Firm, County	Firm, County
Observations	50,593	50,593
R-squared	.779	.489

changes from its 2015 level. All three counties show increases as the benchmark increases. However, Cook County, in which 26 plans are offered, shows a larger increase than Travis County, which has 13 plans, or Worth County, which has 3 plans. The second graph, which depicts the share-weighted average out-of-pocket cost for an individual of age 65-69 in Good health, shows more diversity in the path of the product characteristic as the benchmark increases, illustrating the relative weight of the out-of-pocket costs in the utility function and firms' responses. Finally, the third graph shows the share-weighted average price. While the price increases slightly in Cook County for some benchmark increases, this increase is more than offset by the increase in the average δ'_j. Eventually, as the benchmark increases, plans across all counties hit the zero lower bound on price.

Figure 5 illustrates components of the $GovExp_m$ function for the same three counties. The left-hand graph illustrates the total inside share of MA as the benchmark increases, which determines the number of people for which the government is paying a MA subsidy (as opposed to TM costs directly) in our expenditure calculation. Cook County, which experiences the largest increase in average plan δ and also has the most plans available, sees the highest increases in inside-share. These increases come at a cost, however, as the right-hand graph shows the share-weighted plan bid. Cook County's bids increase nearly 1-for-1 with an increase in the benchmark, whereas the average bid in Travis County increases with a shallower slope, and Worth County's bids remain almost flat despite increasing benchmarks.

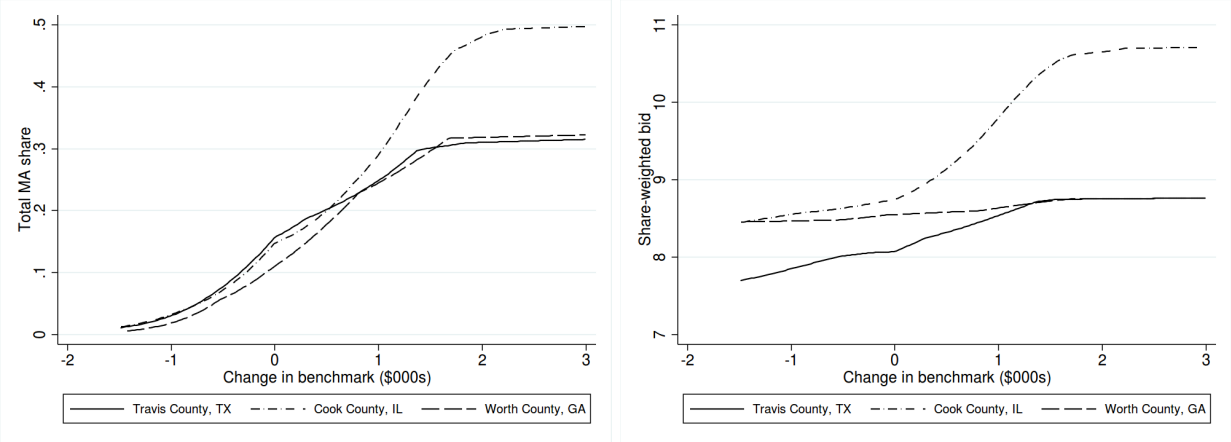
The results of combining these components into the CS_m and $GovExp_m$ functions are shown for our selected counties in Figure 6. The first graph shows per-capita consumer surplus. Under the current policy (where the change in the benchmark is equal to zero), the three counties receive similar surpluses on average. As the benchmark in each county is increased, the average surplus

Figure 4: Product characteristics under counterfactual benchmarks, selected counties



These graphs illustrate three outputs from our counterfactual simulations: the price-adjusted $delta_j$, OOP costs for a 65-69 year old in Good health, and the product's premium. These outputs are shown for Travis County, TX (which contains Austin), Cook County, IL (which contains Chicago), and Worth County, GA (which is a rural county near Albany). All outputs are averages across plans weighted by the market share of each plan.

Figure 5: Average plan bids and MA share under counterfactual benchmarks, selected counties



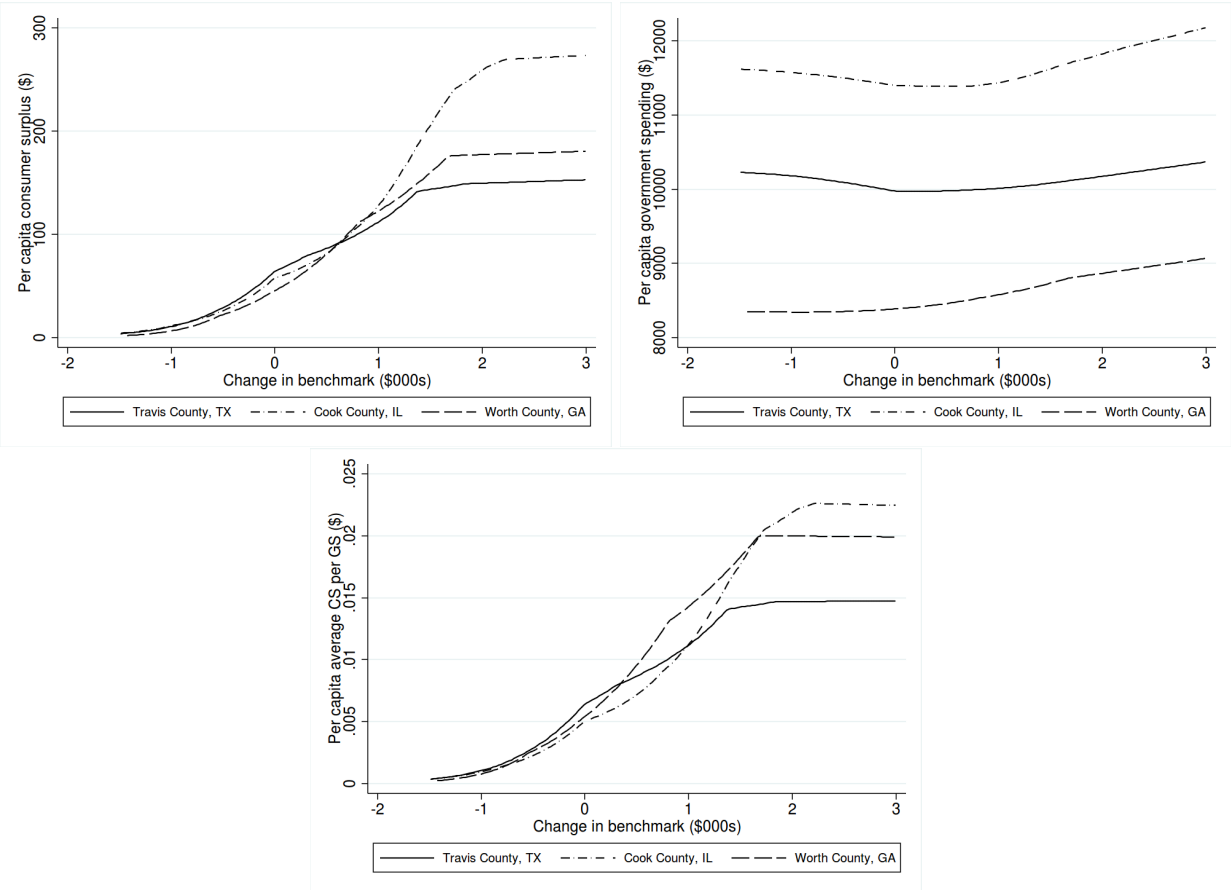
These graphs illustrate two outputs from our counterfactual simulation: the total share of MA (as opposed to TM) and the share-weighted average plan bid. These outputs are shown for Travis County, TX (which contains Austin), Cook County, IL (which contains Chicago), and Worth County, GA (which is a rural county near Albany).

in Cook County grows faster than the others and eventually overtakes them after an increase of $1,000.

The second graph of Figure 6 illustrates per-capita government expenditures on the Medicare program. First, note that since we include spending on Traditional Medicare, even as benchmarks decrease, the per-capita government expenditure in each county remains positive. Cook and Travis Counties illustrate a common pattern in the data: a flat or even decreasing level of government expenditures for modest increases in the benchmark rate. This is a result of the two effects of increasing the benchmark. First, increasing the benchmark leads firms to increase the value offered by their products, which affects the extensive margin of the number of people enrolling in any MA plan. Since, in these counties, the current benchmark is set below the average TM expenditure, any switch from TM to MA saves the government the difference. This is offset by the intensive margin effect of spending more money on people who would have remained on MA even with a lower benchmark. In these counties, the extensive margin effect initially dominates, but it is eventually overcome by the intensive margin effect. In other words, we find potential gains to consumer surplus with flat or even decreasing government expenditures. This occurs in 169 out of 446 markets in our 2015 sample representing 46% of the Medicare population.

The second graph also illustrates the effect of the MA payment system as defined in Equation 1, which rebates a percentage of the difference between the bid and the benchmark back to firms. De-

Figure 6: Per-capita consumer surplus and government expenditures under counterfactual benchmarks, selected counties



These graphs illustrate two outputs from our counterfactual simulation: the per-capita consumer surplus and the per-capita government spending. These outputs are shown for Travis County, TX (which contains Austin), Cook County, IL (which contains Chicago), and Worth County, GA (which is a rural county near Albany).

spite the fact that bids in Worth County stay nearly constant, per-capita government expenditures in the county increase as the benchmark increases.

The final graph of Figure 6 combines the two functions to show the average MA surplus delivered to consumers per dollar spent by the government on the Medicare program. The slope of this line is related to the marginal impact of spending an extra dollar in a particular county through the MA benchmark mechanism, which is the key margin explored by the constrained maximization algorithm of our optimal policy search. Over small increases in the benchmark, Worth County experiences the largest gains in surplus per expenditures. However, Cook County offers the largest possible overall gains of these counties. In other words, the optimal allocation of funds to these counties depends in part on the level of funds available for allocation.

6.3.3 Maximizing consumer surplus

Table 12 reports the results of the optimal policy that results from solving the maximization problem of Equation 11. Column (1) reports consumer surplus under the 2015 policy across a number of dimensions. The mean consumer surplus is \$148.65, and the interquartile range extends from \$9.3633 to \$176.66. The average surplus gained by White Medicare recipients, \$145.67, is lower than either Black recipients (\$167.90) or Hispanic recipients (\$262.13). Low income individuals receive higher surpluses than either medium or high income individuals. The mean variable profit per Medicare beneficiary captured by MA plans, calculated by summing the variable profit for all plans and dividing by the number of Medicare beneficiaries, is \$118.87.

Column (2) reports the results of solving Equation 11. The mean consumer surplus increases 29.6% to \$192.65, and the interquartile range extends from \$16.98 to \$268.69. Under this policy, both Whites (\$191.01) and Blacks (\$203.90) receive increases, though average Hispanic surplus decreases (\$249.83). Low income individuals continue to receive the highest surpluses. Mean variable profit increases to \$160.14 per Medicare beneficiary. The last row shows that this outcome satisfies the non-linear government expenditure constraint to within less than \$0.50 per person.

These changes are presented in more detail in Figure 7. The left-hand histogram illustrates the distribution of benchmarks under the optimal policy, and can be compared to Figure 2. The right-hand histogram shows the distribution of changes in the benchmark. Most counties receive modest increases – the interquartile range of the difference between the optimal and 2015 benchmarks is from \$0 to \$658. The mean change is \$330 and the median is \$267. In percentage terms, 75% of counties get an increase, and 81% of those increases are of less than 10%.

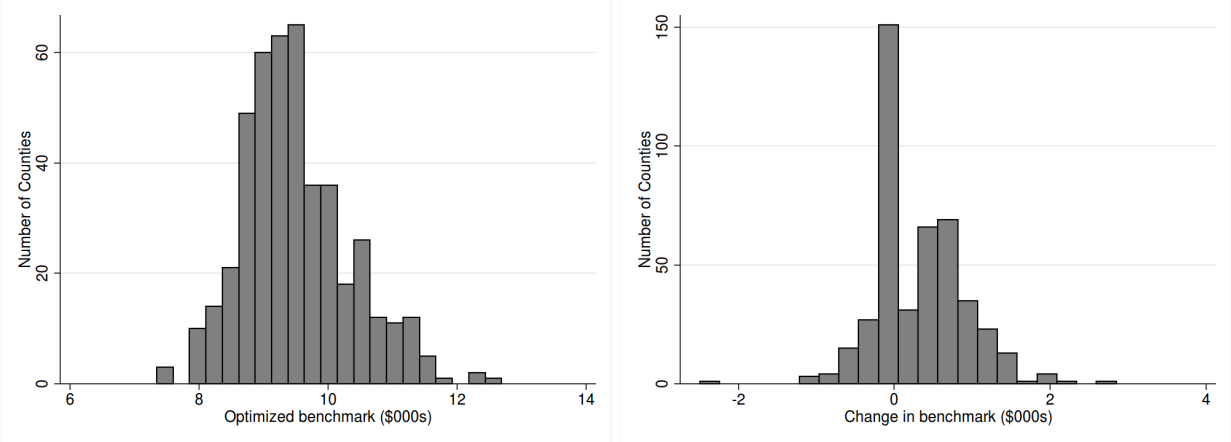
Given the components of the objective function explored above, a natural question is the extent to which the gains we identify come from intensive/extensive margin-related savings opportunities within the MA system and how much comes from redistribution after those savings have been achieved. To answer this question, we identify the degree of savings available by re-optimizing the benchmarks with the objective of minimizing government expenditures. The results of this exercise are in Column (3) of Table 12. The last row reports a decrease in mean government expenditures per person (which includes both TM and MA expenditures) from \$9,797 under 2015 policy to \$9,747. This savings comes with a reduction in mean consumer surplus of 19.9% to \$119.11, and corresponding reductions from across the demographic distribution. Since this exercise includes decreases in the benchmark for many counties, Column (4) repeats the exercise with the constraint that the benchmarks (and thus CS) must not decrease from 2015 levels. In this scenario, the government spends \$9,786, only \$11 less than under the 2015 policy, but mean consumer surplus increases to \$163.58. We conclude that taking advantage of savings opportunities contributes to roughly one third of the gains we find, with the rest coming from redistribution between counties once those savings have been achieved.

Table 12: Counterfactual surplus results

Consumer surplus	(1) 2015 Policy	(2) Max CS	(3) Min GS	(4) Min GS s.t. Floor
Mean	148.65	192.65	119.11	163.58
Variance	47,920	61,504	29,378	49,692
Median	65.28	101.89	49.40	82.98
25th Percentile	9.63	16.98	6.72	12.56
75th Percentile	176.66	268.69	155.28	202.53
Mean "Standard Analysis"	229.99	294.73	185.98	252.95
Mean White	145.67	191.01	116.35	160.82
Mean Black	167.90	203.90	142.65	180.77
Mean Hispanic	262.13	249.86	178.89	273.91
Mean Low Income	166.59	213.74	134.65	184.05
Mean Medium Income	155.37	203.57	124.53	171.00
Mean High Income	125.64	162.80	99.57	137.55
Mean Variable Profit	118.87	160.14	102.48	134.15
Mean Government Spending	9,797	9,797	9,747	9,786

Note: All figures are per Medicare beneficiary. The "Standard Analysis" set excludes individuals with employer-provided plans, those with ESRD, those whose original Medicare eligibility did not come from age, and those who did not purchase Part A and Part B for the full year. Mean Variable Profit is calculated by summing the profit for all plans and dividing by the total number of Medicare beneficiaries. Government spending includes spending on Traditional Medicare.

Figure 7: Optimized benchmarks by county when maximizing average consumer surplus



These graphs illustrate the distribution of benchmarks and the change in benchmarks when maximizing average consumer surplus.

Our approach to determining the optimal benchmarks relies on detailed information on firm and consumer behavior, and explicitly considers welfare tradeoffs across markets mitigated through the government budget constraint. To understand the difference in the implications of this approach relative to a more standard approach that uses county characteristics to determine benchmarks, we obtained county-level data from the U.S. Department of Health and Human Services' Area Health Resources File (AHRF). The AHRF combines data from a number of U.S. government and non-profit sources to provide a detailed view of health-relevant information at the county-year level. We collect median household income, the percentage of seniors (age 65+) in severe poverty, the unemployment rate, the population density, and per-capita counts of MDs, Medicare-certified hopsitals, skilled nursing facilities, and hospice facilities. We also collect the hospital readmission rate for Medicare patients and the "preventable" hospital admission rate.

Table 13 reports the mean values of these and other market level variables from our CMS and MCBS data by benchmark quartile for both the 2015 policy and the optimal policy. The top panel shows that current policy, which assigns benchmarks on TM costs and urban/rural status alone, ends up sorting counties along a number of other dimensions, including income, unemployment rate, and MDs and hospitals per capita.[19] The bottom panel shows that the optimal benchmarks create greater dispersion in TM spending across quartiles. The benchmarks themselves track TM spending more closely. The optimal policy also breaks some of the sorting of the current policy – the number

[19]The benchmark reported here is the "base" benchmark, before any quality bonuses have been applied, and is much lower than the reported average TM costs. As nearly 25% of plans qualify for a 5% quality bonus, the gains available by moving beneficiaries from TM to MA are smaller than they may appear in this table.

of beneficiaries, the unemployment rate, and the number of hospitals change non-monotonically across the quartiles. Furthermore, the optimal policy breaks some of the monotonicity in the number of MA firms and plans in each county as the benchmark increases. We interpret these results as providing evidence that our model captures effects beyond the number of competitors in the insurance market, or the presence of increased competition in the service-delivery sector, but instead is able to pick up on market-level differences in the distribution of consumer preferences and firm-level differences in performance insofar that performance is captured by our estimates of δ and marginal cost, and moves money to those places where firms most efficiently transform government subsidies into consumer surplus.

This result raises questions of implementability. To understand the degree to which our optimal benchmarks could be reproduced with market-level data, we regressed the logarithm of the optimal benchmarks on the logarithm of TM costs and the other variables in Table 13, except for the number of plans and firms. The full results are presented in Appendix Table C.1. These variables explain 82% of the variation in the optimal benchmarks. The TM costs, log beneficiaries, the unemployment rate, and the number of doctors per capita enter with significance at the 99% level. After converting the fit of the model from logs to levels, the standard deviation of the difference between the optimal benchmarks and the regression-model predicted benchmarks is \$351, slightly more than half of the size of the interquartile range of the difference between the optimal benchmarks and current policy.

Appendix Table C.2 compares the surplus generated by the optimal benchmarks to the regression-fit benchmark. Relative to the optimal benchmarks, the regression-fit benchmarks decrease mean surplus 2.8% to \$187.17. Mean variable profit increases 6.5% to \$173.20. Mean government spending increases 0.24% to \$9,821. These results suggest that our approach, which takes advantages of observed and inferred information about the firms present in different markets, adds significant value, but that improvements are possible by considering market conditions other than the TM costs alone.

6.3.4 Alternative social welfare functions

The redistributive nature of the optimal policy implies that these increases in the mean also come with an increase in the variance in the consumer surplus, from 47,920 under the 2015 policy to 61,504 under the policy presented in Column (2) of Table 12. As a consequence, we also explore the family of objective functions represented by Equation 12 in Table 14. Column (1) repeats the 2015 policy outcomes. Column (2) applies an objective function with 99% of the weight on average

Table 13: Mean county characteristics by benchmark quartile, 2015 policy versus optimal policy

2015 policy	**0-25th**	**26-50th**	**51-75th**	**76-100th**
Risk-adj. TM costs per capita	$9,128	9,373	9,908	10,726
Average risk score	.981	.974	.989	1.03
Beneficiaries	31,199	34,711	59,852	110,096
Median household income	53,405	50,899	58,534	62,810
Percent in deep poverty, 65+	2.71	2.49	2.59	2.78
Unemployment rate	5.69	5.64	5.38	5.38
Population density (per mi^2)	481	344	649	3,051
Resources per 10,000 people				
MDs	18.5	20.1	21.1	27.2
Medicare hospitals	.027	.023	.039	.034
Skilled nursing facilities	.600	.689	.673	.456
Hospice facilities	.151	.144	.116	.113
Medicare hospital readmission rate	17.0	17.2	17.6	18.3
Preventable hospital admission rate	51.2	55.8	53.4	52.7
2015 base benchmark	8,610	8,916	9,183	9,962
Number of MA plans	12.2	12.8	14.2	17.9
Number of MA firms	6.49	6.91	7.7	9.0
Obs.	112	111	111	111
Optimal policy	**0-25th**	**26-50th**	**51-75th**	**76-100th**
Risk-adj. TM costs per capita	$8,822	9,433	9,865	11,019
Average risk score	.989	.992	.983	1.01
Beneficiaries	31,323	51,297	48,772	104,465
Median household income	49,827	52,693	55,992	67,169
Percent in deep poverty, 65+	2.80	2.65	2.49	2.63
Unemployment rate	5.84	5.58	5.29	5.37
Population density (per mi^2)	331	585	799	2,813
Resources per 10,000 people				
MDs	16.3	20.0	22.0	28.6
Medicare hospitals	.027	.031	.032	.033
Skilled nursing facilities	.577	.599	.649	.593
Hospice facilities	.154	.117	.143	.111
Medicare hospital readmission rate	16.9	17.2	17.7	18.2
Preventable hospital admission rate	51.6	52.2	57.5	51.8
Base benchmark	8,555	9,166	9,624	10,650
Number of MA plans	13.0	14.8	13.5	15.9
Number of MA firms	7.06	8.13	7.00	7.96
Obs.	112	111	111	111

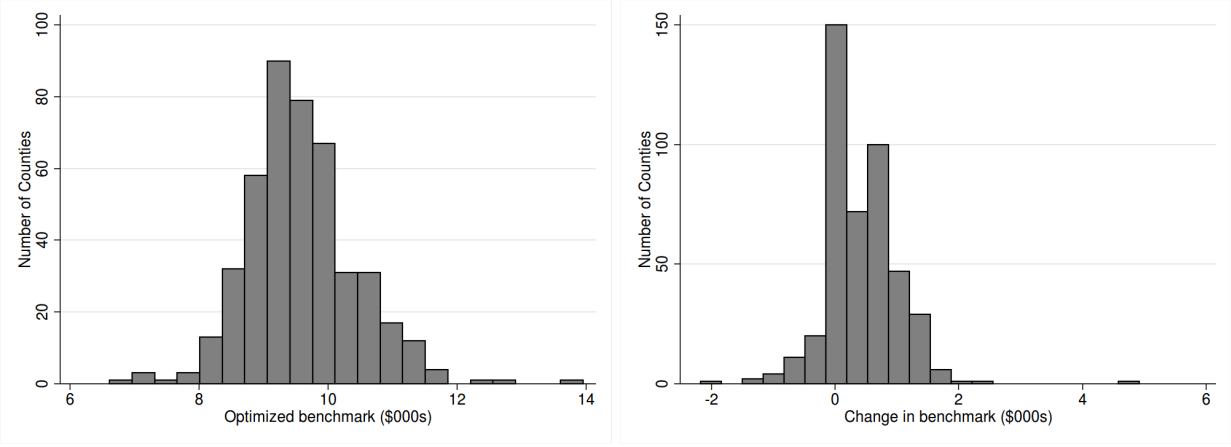
Note: The benchmark quartile is defined by the set of benchmarks under consideration. Risk-adjusted TM costs per capita are formed by taking actual per-capita expenditures reported by CMS and dividing by the average risk score in the county.

consumer surplus, and 1% of the weight on the variance of consumer surplus. Given the relative size of the mean surplus to the variance under the 2015 policy, it is perhaps not surprising that the optimal policy under this weighted objective function reduces benchmarks in many markets in order to fund increases in a few very expensive places. The resulting policy achieves a variance of 11,915, though with a mean of only $72.93.

Column (3) places 99.9% of the weight on consumer surplus and 0.1% on the variance, and achieves a mean of $185.60 while simultaneously decreasing the variance relative to the 2015 policy to 45,016. The interquartile range of $17.30 to $272.24 is higher than that of the pure maximization outcome. The other demographic patterns remain largely similar. Firms capture, on average, $162.76 in variable profit per Medicare beneficiary. Figure 8 illustrates the optimal benchmark distribution under this scenario. 83% of counties receive increases, and the interquartile range extends from $5 to $763. In percentage terms, the increases are similar – 75% are of less than 10%.

Under the scenario illustrated in Figure 8, some counties still receive decreases. Column (4) of Table 14 searches for a "consumer-Pareto" optimal policy by maximizing consumer surplus with the constraint that benchmarks remain at least at their 2015 levels. The result is very similar to the scenario in Column (3), with the except of variable profile, which is 5.3% lower.

Figure 8: Optimized benchmarks by county with penalized objective function



These graphs illustrate the distribution of benchmarks and the change in benchmarks when maximizing 0.999×average consumer surplus −0.001×the variance of consumer surplus.

7 Conclusion

Seeking to reduce the perceived inefficiency of government-provided services, policy makers in a number of contexts have implemented public-private partnerships in which the government provides

Table 14: Counterfactual surplus results alternative social welfare functions

Consumer surplus	(1) 2015 Policy	(2) Max 0.99 CS − 0.01 Var	(3) Max 0.999 CS − 0.001 Var	(4) Max CS s.t. Floor
Mean	148.65	72.93	185.60	186.22
Variance	47,920	11,915	45,016	57,581
Median	65.28	42.39	116.97	97.89
25th Percentile	9.63	3.63	17.30	16.98
75th Percentile	176.66	85.01	272.24	254.55
Mean "Standard Analysis"	229.99	115.22	286.53	286.72
Mean White	145.67	71.04	183.33	183.83
Mean Black	167.90	89.58	203.81	200.15
Mean Hispanic	262.13	108.58	243.59	290.14
Mean Low Income	166.59	83.47	210.03	208.08
Mean Medium Income	155.37	76.37	193.71	195.44
Mean High Income	125.64	59.87	155.23	157.23
Mean Variable Profit	118.87	73.22	162.76	154.12

Note: All figures are per Medicare beneficiary. The "Standard Analysis" set excludes individuals with employer-provided plans, those with ESRD, those whose original Medicare eligibility did not come from age, and those who did not purchase Part A and Part B for the full year. Mean Variable Profit is calculated by summing the profit for all plans and dividing by the total number of Medicare beneficiaries. Government spending includes spending on Traditional Medicare.

subsidies to private firms that are tied to the choices of end-users. The firms are then free to compete with each other – with competition and market forces working to bring down the total cost and increase the benefits of providing the service over time. In many cases, the services provided by firms have differentiated characteristics which are relevant to consumers. Additionally, these services may be offered in geographies with consumers who have substantially different preferences.

We provide a framework for calculating the optimal subsidies to provide to firms that takes into account both the supply and demand responses to alternative subsidy rates. We model demand with a discrete-choice system and avoid the curse of dimensionality in product characteristics by using an approximation approach for calculating the supply decisions. We rely on the observation of variation in the key parameters we wish to vary in the counterfactual analysis, and combine policy function estimation with a first-order condition solver over prices.

We apply our framework to the Medicare Advantage program in the United States, through which approximately one-third of U.S. seniors obtain Medicare benefits, and estimate our model using a combination of micro- and market-level data. We find that the optimal subsidies differ substantially from those currently employed by the government. Once switching costs are taken into account, the current policy generates an average of $148.65 in consumer surplus per person per year. By maximizing the mean consumer surplus, we find an alternative policy that results in an average of $192.65 in benefits per person per year.

Our framework can be adopted to any market in which subsidized firms offer differentiated products. For example, many charter schools offer specialized curricula which may appeal to different sets of parents. With data on family characteristics and choices, the benefits created by these schools and the outcomes of alternative voucher-style policies could be calculated.

References

Aizawa, N. and Kim, Y. S. (2018). Advertising and risk selection in health insurance markets. *American Economic Review*, 108(3):828–67.

Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370.

Benkard, C. L., Bodoh-Creed, A., and Lazarev, J. (2018). Simulating the dynamic effects of horizontal mergers: U.s. airlines. *Working Paper*.

Berry, S. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):pp. 841–890.

Bundorf, M. K., Levin, J., and Mahoney, N. (2012). Pricing and welfare in health plan choice. *American Economic Review*, 102(7):3214–48.

Cabral, M., Geruso, M., and Mahoney, N. (2018). Does privatized medicare benefit patients or producers? evidence from the benefits improvement and protection act. *American Economic Review*.

CMS Office of the Actuary (2017). Actuarial bid training–2017.

Curto, V., Einav, L., Levin, J., and Bhattacharya, J. (2015). Can health insurance competition work? evidence from medicare advantage. Working Paper 20818, National Bureau of Economic Research.

Dalton, C. M., Gowrisankaran, G., and Town, R. (2018). Salience, myopia, and complex dynamic incentives: Evidence from Medicare Part D. Technical report, National Bureau of Economic Research.

Daskalakis, C., Goldberg, P., and Papadimitriou, C. (2009). The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259.

Decarolis, F., Polyakova, M., and Ryan, S. P. (2015). Subsidy design in privately-provided social insurance: Lessons from medicare part d. Working Paper 21298, National Bureau of Economic Research.

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

Einav, L., Finkelstein, A., and Tebaldi, P. (2018). Market design in regulated health insurance markets: Risk adjustment vs. subsidies. *MIT Working Paper*.

Enthoven, A. C. (1978). Consumer-Choice Health plan (second of two parts). A national-health-insurance proposal based on regulated competition in the private sector. *N. Engl. J. Med.*, 298(13):709–720.

Ericson, K. M. and Starc, A. (2015). Measuring consumer valuation of limited provider networks. *American Economic Review*, 105(5):115–19.

Fan, Y. (2013). Ownership consolidation and product characteristics: A study of the us daily newspaper market. *American Economic Review*, 103(5):1598–1628.

Geruso, M. and Layton, T. J. (2017). Selection in health insurance markets and its policy remedies. *Journal of Economic Perspectives*, 31(4):23–50.

Goolsbee, A. and Petrin, A. (2004). The consumer gains from direct broadcast satellites and the competition with cable tv. *Econ*, 72(2):351–381.

Gruber, J. (2017). Delivering public health insurance through private plan choice in the united states. *Journal of Economic Perspectives*, 31(4):3–22.

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 103(7):2643–82.

Hausman, J., Leonard, G., and Zona, J. D. (1994). Competitive analysis with differenciated products. *Annales d'conomie et de Statistique*, (34):159–180.

Jaffe, S. P. and Shepard, M. (2017). Price-linked subsidies and health insurance markups. Working Paper 23104, National Bureau of Economic Research.

Laffont, J.-J. and Tirole, J. (1993). *A theory of incentives in procurement and regulation*. MIT press.

Lustig, J. (2010). Measuring Welfare Losses from Adverse Selection and Imperfect Competition in Privatized Medicare. Working paper.

Lyons, R. D. (1972). Nixon signs $5-billion bill expanding social security. *The New York Times*.

McGuire, T. G., Newhouse, J. P., and Sinaiko, A. D. (2011). An economic history of Medicare part C. *Milbank Quarterly*, 89(2):289–332.

MedPAC (2017). Status report on the medicare advantage program. In *Report to the Congress: Medicare Payment Policy*, chapter 13, pages 345–379.

Newhouse, J. P., Price, M., Hsu, J., McWilliams, J. M., and McGuire, T. G. (2015). How much favorable selection is left in medicare advantage? *American journal of health economics*, 1(1):1–26.

Newhouse, J. P., Price, M., Huang, J., McWilliams, J. M., and Hsu, J. (2012). Steps to reduce favorable risk selection in medicare advantage largely succeeded, boding well for health insurance exchanges. *Health Affairs*, 31(12):2618–2628. PMID: 23213145.

Noel, M. D. (2007). Edgeworth price cycles, cost-based pricing, and sticky pricing in retail gasoline markets. *The Review of Economics and Statistics*, 89(2):324–334.

Nosal, K. (2012). Estimating switching costs for medicare advantage plans. Working paper.

Rinnooy Kan, A. and Timmer, G. (1987). Stochastic global optimization methods. part ii: multi level methods. *Mathematical Programming*, 39(1):57–78.

Small, K. A. and Rosen, H. S. (1981). Applied welfare economics with discrete choice models. *Econometrica: Journal of the Econometric Society*, pages 105–130.

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

Tebaldi, P. (2017). Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca. Working paper.

Town, R. and Liu, S. (2003). The welfare impact of Medicare HMOs. *Rand J Econ*, 34(4):719–736.

Appendices

A A Simple Model of the Impact of Subsidies

It is useful to provide some simple economic intuition behind why demand conditions and the interaction with firms behavioral responses to those demand conditions imply that optimal subsidy allocations may differ across markets even when marginal costs are identical. Consider the following model in which the Government is allocating a fixed budget for health care services across two markets, A and B. In each market, enrollees can choose a publicly provided plan or a privately provided plan. Each market has the same number of individuals eligible for the program and the marginal and average cost for the government of providing the public plan is zero. The cost to eligibles of enrolling in the public plan is zero while the private plan may have a non-zero price. In each market, the private plan is offered by a monopolist with an average and marginal cost of zero. Thus, the subsidy for the private plan is relative to the implicit subsidy for the publicly provided plan.

Suppose the demand for the private plan in Market A is $P_A = 100 - Q_A$ while in Market B the demand is $P_B = 200 - 2 \times Q_B$. Let B_A and B_B be the per-enrollee subsidy that the firms in each market receive from the government. The firm in Market A solves $\max_{P_A}(P_A + B_A) \times (100 - P_A)$ and sets $P_A = 50 - B_A/2$ with corresponding quantity $Q_A = 50 + B_A/2$. The firm in Market B solves $\max_{P_B}(P_B + B_B) \times (200 - 2 \times P_B)$ and sets $P_B = 100 - B_B/2$ with corresponding quantity $Q_B = 50 + B_B/4$. Even though the cost structures are the same across markets, government subsidies lead to different equilibrium outcomes across the markets due to the interaction between the demand elasticities and the imperfectly competitive environment.

Consider a government which seeks to maximize the sum of consumer surplus across markets by allocating a fixed subsidy budget. The equilibrium outcomes above can be used to derive consumer surplus and government expenditures in each market as a function of the per-enrollee subsidy. In Market A, the consumer surplus is $CS_A = 1/2 \times (50 + B_A/2)^2$ and expenditures are $EXP_A = 50B_A + B_A^2/2$. In Market B, the consumer surplus is $CS_B = (50 + B_B^2/4)^2$ and expenditures are $EXP_B = 50B_B + B_B^2/4$.

Suppose the government's budget is \$50.5. As an extreme example, if the government allocates

the entire budget to Market A, $B_A = 1$ and the increase in consumer surplus over a baseline of $B_A = 0$ is $\Delta CS_A = 25.125$. If the government allocates the entire budget to Market B, $B_B \approx 1.005$ and the increase in consumer surplus over a baseline of $B_B = 0$ is $\Delta CS_B \approx 25.187$. Thus, in this stylized example, it is optimal for the government to set the subsidy structure to spend the entire budget in Market B: $B_A = 0$ and $B_B \approx 1.005$. Under this policy, the increase in profit for the monopolist in Market B over a baseline of $B_B = 0$ is approximately \$50.37. As a consequence, the subsidy improves total welfare (ignoring the tax distortion) by reducing the monopoly pricing distortion.

In this model, consumer surplus can be written as a function of government expenditures by solving for the subsidy rate associated with a given level of government spending. The derivative of this function is then the rate at which government spending is passed on to consumers. For Market A, $\frac{\partial CS_A}{\partial EXP_A} = \frac{1}{4}\left(1 + \frac{\sqrt{1250}}{\sqrt{EXP_A + 1250}}\right)$, and for Market B, $\frac{\partial CS_B}{\partial EXP_B} = \frac{1}{4}\left(1 + \frac{50}{\sqrt{EXP_B + 2500}}\right)$. Thus, for any positive level of government expenditures, it is optimal to set a higher subsidy rate for Market B than for market A.

In our analysis in the main body of the paper, we extend this simple model to allow for multiple firms to offer differentiated plans in an oligopolistic setting. Firms compete by choosing plan prices and non-plan characteristics. Consumers are heterogeneous along multiple dimensions and choose a single plan from the available options.

B Monte Carlo Analysis of Supply Approximations

Our approach approximates the responses of firms to changes in the benchmark instead of solving for the full equilibrium response – in particular, we approximate the change in plan characteristics and then solve for the equilibrium in prices. An obvious question is how well does such approximations in practice. In this appendix, we run several Monte Carlo experiments to examine this issue.

In these experiments, we simulate market-level data for two periods. We use the results in the first period as the basis for making predictions of the second period market outcomes. We explicitly solve the firm's problem in the second period, and compare the exact solution to an approximation created using the first period data and the approach we employ in the paper. For simplicity, we specify utility as $u_{ijmt} = \delta_{jmt} - \alpha p_{jmt}^2 + e_{ijmt}$ where e_{ij} is an iid Type I Extreme Value error term.

We allow price to affect utility in a nonlinear way as it allows for greater concavity in the plans objective function – alternatively we could have allowed for convexity in the cost specification. Plan costs are given by $c_{jmt} = exp(.2 + .2\delta_j + \nu_j)$ where ν_j is drawn from a $N(0, .1)$. In each period, we simulate M markets where M is 50, 100 or 200. In each market, the insurers receive a subsidy, z_{mt}. In the first period, the market-level subsidy is sequentially ordered and ranges from .1 to 1.1 with an interval of .1. In the second period we reverse the order of the subsidies so that the highest subsidy market is now the lowest and vice versa. In the simulations we allow the number of market participants to range from 3 to 7.

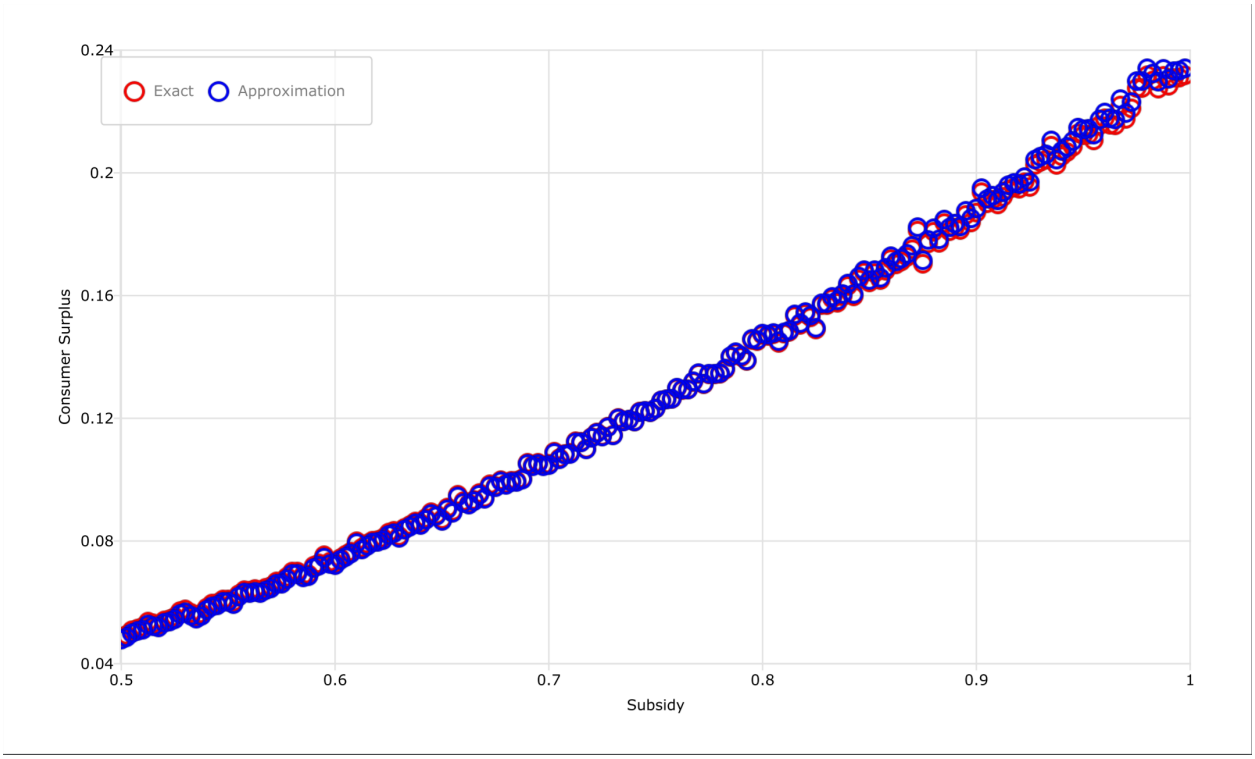
In order to forecast the equilibrium with new benchmarks, we follow a procedure that mimics our empirical approach. Specifically, we use the first period equilibrium results to estimate demand and given those results, invert the premium setting first-order condition to recover implied marginal cost, $\hat{mc_{jmt}}$. We then regress the implied δ_{jm1} on $\hat{mc_{jmt}}$ and its square to recover the cost function parameters. Importantly, we estimate the relationship between the implied δ_{jm1} and the benchmarks z_{m1} and use the results from this regression to forecast δ_{jm1} given the updated subsidies, z_{m2}. We then calculate the new expected marginal cost and then solve for the new premiums given the subsidies, the updated δ_{jm2} and c_{jm2}. We compare these results to the actual impact of changing the benchmarks where the actual impact is calculated by explicitly solving for premiums and δ_{jm2} using the plans' first-order conditions for both premiums and δ_{jm2} given the parameters of the model.

We define the logarithm error as $Err^{CS} = Log(CS_{Exact}) - Log(CS_{approx})$. Table B.1 presents some of the key statistics from our Monte Carlo experiments. The results displayed in the table indicate that our approximation approach generates estimates of consumer surplus, our key outcome that we need to match, that are very close to an approach using an exact solution. Figure B.1 illustrates this for a run with 7 firms and 200 markets. Red circles indicate the consumer surplus calculated when the firms' problem is solved explicitly, and the blue circles indicate the consumer surplus found using our approximation approach. The Figure shows our approach works well across the range of subsidies, and works particularly well near the middle of the subsidy range.

Table B.1: Monte Carlo Evidence

Number of Markets	Number of Plans	Mean Err^{CS}	Mean Absolute Err^{CS}
50	3	.0014	.0094
100	3	.0011	.0076
200	3	.00077	.0037
50	4	.00093	.0051
100	4	.00064	.0025
200	4	.00095	.0051
50	5	.00065	.0027
100	5	.0011	.0060
200	5	.0013	.0080
50	6	.0011	.0050
100	6	.0016	.0091
200	6	.0012	.0064
50	7	.0011	.0055
100	7	.0013	.0063
200	7	.0013	.0070

Figure B.1: Monte Carlo Results for Seven Firms



C Approximating the optimal benchmarks with market-level information

Table C.1: New benchmark regression exercise

VARIABLES	(1) ln_new_bench
ln Risk-adj. per-cap. TM costs	0.874***
	(0.0251)
Average risk score	0.00420
	(0.0324)
ln Beneficiaries	-0.0117***
	(0.00234)
ln Median household income	0.00766
	(0.0115)
Percent in deep poverty, 65+	-0.00403**
	(0.00193)
Unemployment rate	0.00424***
	(0.00141)
Population density (per mi^2)	-3.39e-07
	(4.72e-07)
MDs per 10,000	0.000324***
	(0.000120)
Hospitals per 10,000	-0.00331
	(0.0302)
Skilled nursing facilities per 10,000	0.00287
	(0.00501)
Hospice facilities per 10,000	-0.000588
	(0.0110)
Medicare hospital readmission rate	0.000779
	(0.00123)
Preventable hospital admission rate	-0.000165
	(0.000137)
Constant	0.264**
	(0.123)
Observations	443
R-squared	0.823

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table C.2: Comparing optimal benchmarks to regression benchmarks

Consumer surplus	(1) 2015 Policy	(2) Max CS	(3) Predicted
Mean	148.65	192.65	187.17
Variance	47,920	61,504	67,071
Median	65.28	101.89	86.09
25th Percentile	9.63	16.98	14.93
75th Percentile	176.66	268.69	250.62
Mean "Standard Analysis"	229.99	294.73	287.37
Mean White	145.67	191.01	186.66
Mean Black	167.90	203.90	190.47
Mean Hispanic	262.13	249.86	206.05
Mean Low Income	166.59	213.74	207.02
Mean Medium Income	155.37	203.57	197.07
Mean High Income	125.64	162.80	159.42
Mean Variable Profit	118.87	160.14	170.62
Mean Government Spending	9,797	9,797	9,821

Note: All figures are per Medicare beneficiary. The "Standard Analysis" set excludes individuals with employer-provided plans, those with ESRD, those whose original Medicare eligibility did not come from age, and those who did not purchase Part A and Part B for the full year. Mean Variable Profit is calculated by summing the profit for all plans and dividing by the total number of Medicare beneficiaries. Government spending includes spending on Traditional Medicare.