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

May 2018

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 payment from the government for each individual enrolled, and may add additional benefits and charge an additional premium – an approach which mirrors many other goods provided via a government subsidy. The optimal subsidy in different 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. However, governments subsidies are typically pegged only to a measure of average cost. We study optimal subsidy design in Medicare Advantage by estimating a flexible supply and demand systems in an oligopoly setting that features demand-side heterogeneity and switching costs, and supply-side price-setting and benefit design behavior. We find the the optimal subsidy structure differs from the implemented one and significantly improves consumer surplus.

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

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

*Miller: keatonm@uoregon.edu; Petrin: petrin@umn.edu; Town: robert.town@austin.utexas.edu; Chernew: chernew@hcp.med.harvard.edu

1 Introduction

There is a long-standing concern that the government provision of services is inefficient. According to this view, government bureaucracies lack powerful enough incentives to efficiently design and deliver the good or service. An alternative approach to government provision is to let private firms compete for customers and subsidize those firms according to the number of consumers they obtain. The economic rationale for this approach is that profit maximization motives and competitive market pressures will push firms to provide the optimal quality at a price approaching marginal cost. This rationale is persuasive enough that this publicly financed/privately provided approach is widely employed by the US government. For example, in the Medicare system under the Medicare Advantage (hereafter MA) program, enrollees can forego traditional Medicare (TM) and choose one of many differentiated private plans who assume responsibility for managing the enrollee's health care benefits. The plan receives a per-capita payment from the government (based on the government set "benchmark rate") as well as premium payments from the enrollee. A similar system is used within the health exchanges created by the Affordable Care Act to subsidize the health insurance purchases of low-income individuals. This basic structure not only underlies the private components of government financed health care but is similar to Charter School initiatives.

A fundamental issue with significant welfare implications in these publicly financed, privately provided systems is how to best set the subsidies. In particular, the demand and imperfectly competitive supply conditions often vary across geographies. This implies that the optimal subsidy should also vary by geography. While the government often lets the subsidy vary by geography, typically that variation is principally driven by measures of average cost alone. For example, in MA, the benchmark subsidy rate is set via a function of the county-level average cost to the government of providing benefits to TM enrollees. In our data the annual benchmark rate varies from \$8,387 to \$17,010. However, the optimal subsidy depends not only on the cost to private firms of providing the service, which may differ from the government's cost and is often difficult for the government to observe, but also on the more general supply responses of firms, such as quality (which will impact costs) and price setting behaviors, and their interaction with consumer preferences.

We develop an approach to determining the optimal subsidy in settings in which the government relies on competing private firms to provide goods which are ultimately chosen by consumers. We then apply this approach to the Medicare Advantage market. Insurers in this market have the ability to set premiums and plan benefits in response to the benchmark rates set by the Centers for

Medicare and Medicaid Services (CMS).[1] To calculate the optimal subsidy, we start with a model of demand for MA plans, estimate the supply responses to changes in the benchmark rate, finally determine equilibrium outcomes under alternative policies. To the best of our knowledge, we are the first to study the optimal subsidy problem in differentiated product environments directly.

Our demand model, which we believe is of independent methodological interest, allows for considerable preference heterogeneity across observable plan and consumer characteristics and unobservable plan-specific 'quality' differences. Due to the differences 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 data on demographics and plan choices from the Medicare Current Beneficiary Survey (MCBS) for the years 2008-2010 supplemented with data on plan characteristics including premiums and benefit structures from CMS.

The results from the demand model are quite reasonable. We find the plan-level mean premium elasticity of demand to be -3.45. MA plans are more valuable to beneficiaries who have lower income, lower educational attainment, are younger and are in better health. We also find significant but plausible switching costs in MA. The within insurer / across plan switching cost for median-income individuals is \$182, the between MA insurer switching cost is \$618 and the switching cost between a MA plan and traditional Medicare is \$736. We find that the MA program as-is creates \$29.56 in consumer surplus per year per Medicare Beneficiary. Our estimates imply that in 2010, the MA program generated \$1.135B in total consumer surplus.

With the estimated perference parameters in hand, we then turn to estimating supply-side responses to changes in the benchmark. These responses capture both changes to product characteristics (which include benefit design) and to premiums. We employ a hybrid approach: we estimate changes to product characteristics directly from the data by estimating the policy responses of firm in the style of Bajari et al. (2007). The estimated supply responses vary across plans and geographies but, in general, imply that plan mean utility is increasing in the benchmark. Importantly, for the calculation of the optimal benchmark the relationship between benchmark and mean utility is often S-shaped. It is convex over the lower portion of the benchmark space and then becomes concave at higher benchmark levels.

With these estimated demand and supply responses, we then compute our primary counterfactuals of interest: determining the optimal benchmarks under different CMS objective functions. We

[1] While insurers may enter and exit, the median share captured by new entrants in our data is 1.5%.

first calculate the benchmark that maximizes total enrollee consumer surplus subject to a spending constraint. The implied optimal benchmarks are meaningfully different than the CMS benchmarks. The average differential between the optimal benchmark and the CMS benchmark is 12%. The change in benchmarks can often be dramatic – 3% of the counties have a optimal benchmark that is 20% greater than the benchmark in our data. The estimates imply that the optimal benchmark will increase Medicare beneficiary consumer surplus by $288 per year. Our estimates of the optimal benchmark imply significant variance in the average consumer surplus by county. 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 counties. We find that there are benchmarks that increase average consumer surplus and reduce the variance of consumer surplus across counties relative to the benchmarks actually employed by CMS.

Our approach allows us to perform other counterfactual exercises. Within our sample counties, an increase of the overall benchmark of $1, which leads to an increase in government expenditures of approximately $10 million, generates a $508,000 increase in consumer surplus. Finally, 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 1% increase in premiums, holding everything else constant results in a 4.9% shift in enrollees from MA to traditional Medicare.

Our paper contributes to an extensive literature studying the impact of Medicare Advantage. McGuire et al. (2011) provide an excellent review. Our work is most closely related to Town and Liu (2003), Lustig (2010) and Curto et al. (2014). Town and Liu (2003) estimate a nested logit demand for MA advantage plans and calculate that the program (then called Medicare+Choice) generated $113 in consumer surplus and $244 in profits per Medicare beneficiary in 2000. They also found significant geographic variation in consumer surplus. Curto et al. (2014) also estimate a nested logit model of MA demand and find that more recently the program generated approximately $600 in annual surplus with the insurers capturing the majority of the gains. Recently, several papers estimate the passthrough from benchmarks to premiums and benefits in MA. Using an unanticipated change in the benchmark in 2000, Cabral et al. (ming) finds that 54% of an increase in the benchmark is passed onto enrollees, while Duggan et al. (2016) uses discrete variation in the benchmark across urban and rural counties and estimates a much smaller passthrough.[2] Our

[2]Song et al. (2013) also find a passthrough from benchmarks to bids, which is a measure of premiums and the value of benefits, of 53%.

paper is also related to Nosal (2012) who estimates a dynamic demand model of MA plan choice and finds very large switching costs ($4000 at the median) in the MA program.

There is also recent related work on calculating the optimal subsidy structure in different contexts – much of it focused on Patient Protection and Affordable Care Act (ACA) exchanges. Tebaldi (2017) and Jaffe and Shepard (2017) both examine the optimality of price linked relative to a fixed voucher subsidy setting strategies in different ACA insurance exchanges. Tebaldi (2017) also examines the optimality of age based subsidies. Ericson and Starc (2015) examine the implications of age gradient premium regulation in an ACA-like exchange. Bundorf et al. (2012) examine the impact of allowing health status linked premiums in the employer-sponsored setting. Decarolis et al. (2016) examine the optimality of using vouchers verses the current subsidy setting strategy in Medicare Part D. None of these papers examine the impact of the subsidy on the non-premium characteristics of the plan.

We proceed in Section 2 with a discussion of the institutional details of Medicare Advantage program. In Section 3 we introduce our model of demand. We discuss our data on Medicare beneficiaries and MA plans in Section 4 and estimate the demand model in Section 3.1. Section 7 details the design of our optimal subsidy exercise and presents the results of that exercise along with other counterfactuals. We conclude in Section 8 with a discussion of the policy implications and directions for future research.

2 The Medicare Advantage Program and the Role of Subsidies

In this section we provide institutional background on the Medicare Advantage Program and then briefly discribe the intuition underlying why demand and supply factors should be taken into account in order to optimally set the private insurance subsidy.

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.[3] 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

[3] Medicare's currently covers the diasabled and those with end-stage renal disease. Those additions to the eligible population was added in 1972.

then pays the provider according to a pre-set reimbursement schedule.

Partly in response to Medicare's increasing cost growth, Congress created the Part C program. As outlined in the Tax Equitiy and Fiscal Responsibility Act in 1982 (TEFRA), under Part C, Medicare can directly risk contract with private insurers for the managment of its beneficiaries health benefits.[4] By 1985, the Health Care Financing Administration, the precursor to the modern-day Centers for Medicare and Medicaid Services (CMS), and formalized the rules surrounding risk-based contracting and began a series of trial programs based in part on the ideas of Enthoven (1978). The government contracted with private insurers to manage the care of select groups of Medicare enrollees in exchange for a fixed payment that did not vary with the realized medical expenditures of each beneficiary. This program was brought to the entire country in the early 1990s under the name Medicare+Choice. The program renamed Medicare Advantage (MA) as part of the Medicare Prescription Drug, Improvement, and Modernization Act of 2003.[5]

Currently, the vast majority of Medicare beneficaries have the option of enrolling in a MA plan. If they do enroll in an MA plan, the enroll forgoes the traditional Medicare program having the plan manage the beneficary's medical (and generally drug) benefits. The MA enrollee may pay a premium on top of the Part B premium they must continue to pay. MA plans generally provides a more generous benefit package than TM and can include dental and eye coverage (which are not covered under TM) as well as lower out-of-pocket cost sharing and Part B rebates. While MA plans usually have more generous benefits they also typically employ more restrictive provider networks than TM.

The popularity of the program (with the exception of the period of the early 2000s where subsidies were reduced) has steadly increased over the last 20 years. Figure 1 shows the growth in MA enrollment over time. Enrollments remained near five million until the reforms passed in 2003 were implemented in 2007. Currently, 33% of Medicare beneficiaries are enrolleed in an MA plan.

While the particular features of the program have evolved over time, the basic financial approach underlying the MA system has not signficantly changed over its history. Private insurers compete to attract enrollees receiving a risk-adjusted, per enrollee payment from the government. Plans compete along multiple dimensions including premiums, benefit generousity, and provider network. Insurers often heavily market their plans.

The CMS subsidy is determined, in large part, by a fixed 'benchmark' payment rate for each

[4] The cost of the Medicare program has increased substantially. In 1970, Medicare composed about 0.5% of GDP. By 1980, Medicare had grown to 1.1% of GDP.

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

Figure 1: Medicare Advantage Enrollment, 1997-2017

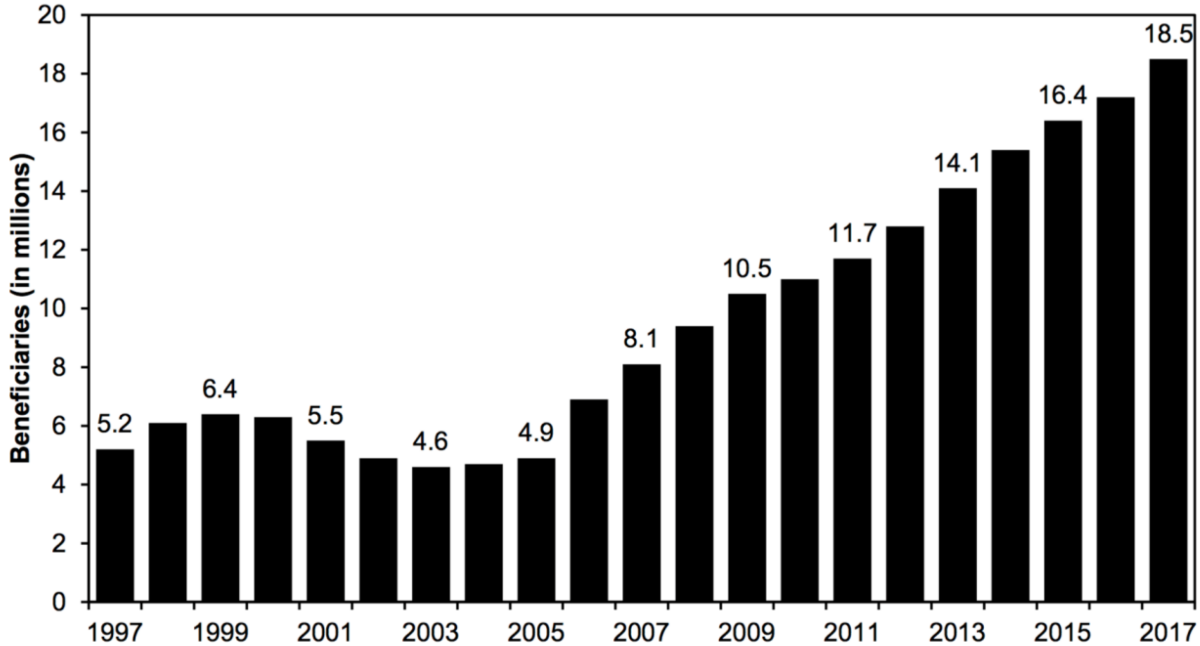

Note: MA (Medicare Advantage).

Source: Medicare managed care contract reports and monthly summary reports, CMS.

county, which varies signficantly across geographies ((Newhouse et al., 2012)). CMS calculates the benchmark rate by starting with the average per capita FFS Medicare spending within the county. The average spending 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 is applied to the benchmark rates across counties. Figure 2 illustrates the distribution of Medicare Advantage benchmarks across the counties in our sample for 2010.

Figure 2: Medicare Advantage benchmarks, 2010

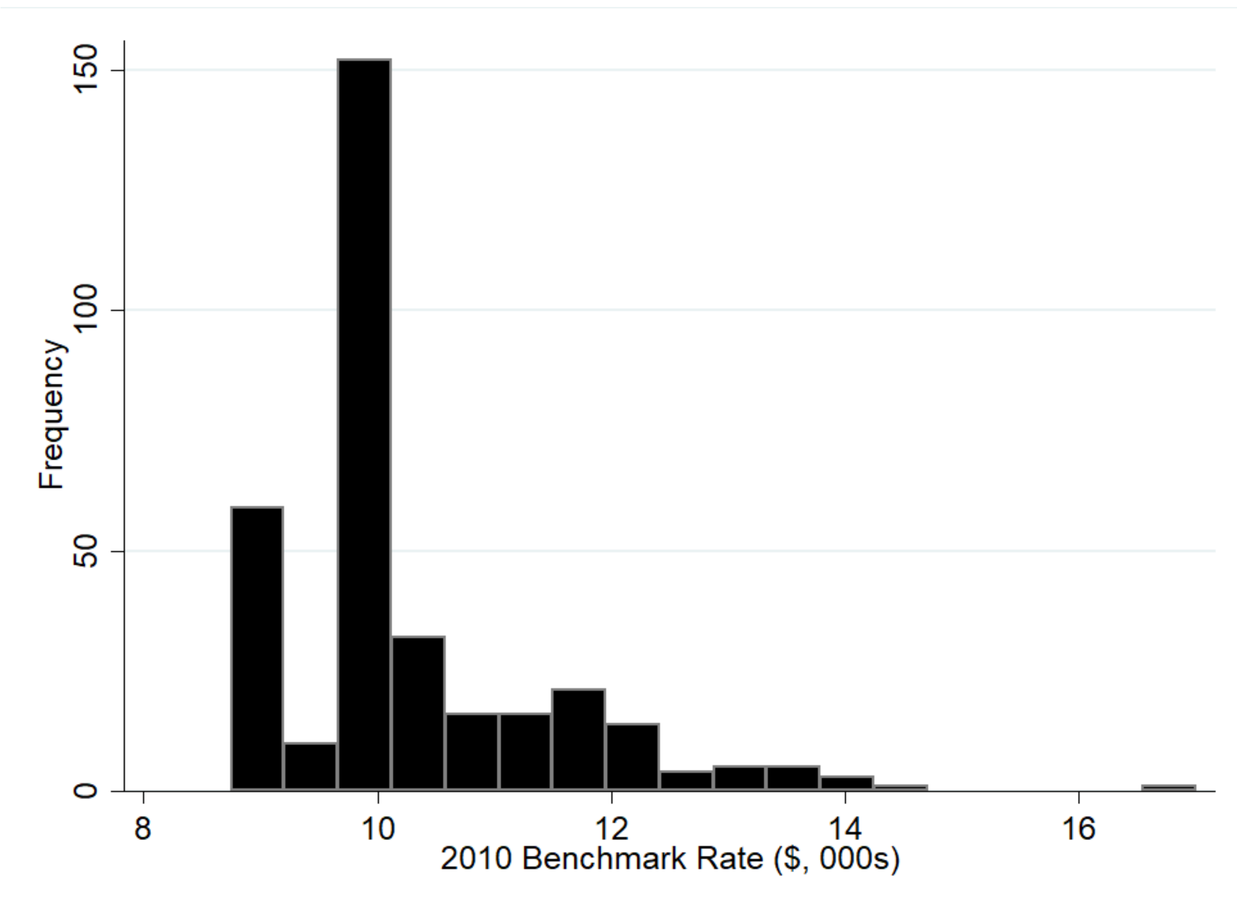

Note: Includes counties observed within the Medicare Current Beneficiary Survey.

After the benchmark rate is set for each county, firms submit 'bids' to offer plans in particular counties in exchange for a particular payment rate, which may be above or below the benchmark rate. Insurers often offer multiple plans in a given county and submit different bids for each plan. MA plans typically also offer a bundled prescription drug plan (Part D) and submit a seperate bid for that product that maps into the part D premium which enrollees in the plans will also have to pay. During our study period, insurers who submit bids that are under the benchmark must translate 75% of that difference into a 'rebate' to consumers. This rebate principally takes the form of more generous benefits but can also include reductions in Part B or Part D premiums.[6]

Firms offer plans to beneficiaries during an annual Open Enrollment period in the fall 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 in the case of "major" life changes (e.g. relocation, death in the family, etc.).[7] After the enrollment period closes, firms send information about their patients' demographics and diagnoses to CMS. CMS uses this data to calculate a risk-adjusted payment to the firm at the patient level based on Medicare expenditures on similar individuals relative to the average Medicare enrollee. Payments from CMS to firms therefore follow

$$payment = \begin{cases} bid + RiskAdjustment & \text{if } bid \geq benchmark \\ bid + \theta \times (benchmark - bid) + RiskAdjustment & \text{if } bid < benchmark \end{cases}$$

Today, 31% of Medicare beneficiaries are in an MA plan and payments to plans exceed $170 billion each year. The top four insurers nationwide, United Health Group, Humana, Kaiser, and Aetna, have 56% of total enrollment. While the average Medicare beneficiary has access to 10 plan options, 25% have access to 3 or fewer plans. 64% of beneficiaries have access to 5 or more 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 cost of providing the service. For example, Medicare Payment Advisory Commission (MedPAC) typically recommend that payments be cut if they determine the payment is greater than

[6] Beginning in 2012, CMS evaluates the quality level of the plan proposals and awards a "star rating" according to pre-announced criteria. These star ratings determine the amount of the difference between the bid and the benchmark that is given to firms.

[7] Recently, enrollees are now allowed to switch to a "5 star" plan at anytime during the year. This program was initiated after the end of our sample period.

the cost of provision.[8] If their is demand heterogeneity, setting the subsidy equal to a function of the measured average cost of production will only be optimal under perfect competition. However, 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 hetergeneity in these factors across markets, the optimal subsidy will typically vary across markets that may not be very 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 with demand heterogeneity that makes this point. The model examines two geographically seperated 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 set such that it is zero in one market and the entire subsidy will be allocated to the market with a higher willingness to pay. In this example, the subsidy is also increases total welfare (ignoring the welfare loss from tax collections) as it counteracts the distortive impact of monponopoly 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 A Model of Demand for Medicare Advantage Plans

Our model of demand for MA plans is inspired by Goolsbee and Petrin (2004) and leverages the rich individual-level heterogeneity in the MCBS data and the detailed plan characteristic data available from CMS. In the model, consumers select plans based upon the premiums, plan features, their expected out-of-pocket expenses and costs of switching away from the plan in which they are currently enrolled. Because we have detailed individual data, preferences over a given plan can vary across consumers (and therefore geography) and, correspondingly, consumer responses to changes in premiums and plan characteristics can vary across individuals (and geography). The ability to accurately capture these demand responses is important for assessing the optimality of a given benchmark.

Consumers, denoted by i, live in a market (i.e. county) m. Time is denoted by t. Consumers choose between one of several plans, denoted by j, with an outside option representing automatic enrollment in traditional Medicare. Consumers have demographic characteristics z_i, belong to an

[8] See, for example, (MedPAC, 2017), Chapter 9.

income bracket denoted by g and an age-and-health-status type h. The indicator variable d_{gi} is 1 if consumer i is in income bracket g, and the indicator variable d_{hi} is 1 if consumer i has age-and-health-status type h. Consumers observe plan premiums, p_{mj} and product characteristics X_{mj} (such as whether the plan is an HMO plan or a PPO plan). Importantly, consumers also observe expected plan out-of-pocket costs, C_{hj}, that vary by age-and-health-status.[9]

Dropping the market and time subscripts, a consumer's choice-specific utility for a particular plan j is given by:

$$\begin{aligned} u_{ij} =& \alpha_0 p_j + \sum_g \alpha_g p_j d_{gi} + \sum_h \beta_h C_{hj} d_{hi} + \sum_s \beta_s S_{sij} \\ &+ \sum_s \sum_l \beta_{sz} z_{il} S_{sij} + \beta_z z_i + \beta X_j + \xi_j + \epsilon_{ij} \end{aligned} \tag{1}$$

In this equation, α_0 represents the impact of changes in premium that common to all beneficiaries. We allow price sensitivity to vary by income category through the α_g parameters. The parameters β_{kh} captures the impact of changes in expected out-of-pocket costs on plan valuation for those with health status h for out-of-pocket-costs of type k. The impact of switching costs are captured by S_{sij}, β_s, and β_{sz}, which are discussed in more detail below. The parameter β_z captures the tastes of consumers with different demographic characteristics for Medicare Advantage plans relative to the outside good. The vector β captures mean tastes for plan characteristics, X_j.

In our framework, beneficiaries can incur switching costs for moving between TM and MA, between different MA insurers and different plans within an insurer. Empirically, there is significant inertia in plan enrollment (Nosal, 2012). Enrollees in TM face a different set of rules and provider networks than in MA and those rule and networks vary across insurers and even across plans within an insurer. Switching a plan entails learning about the plan's administrative procedures and network structures. 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 enrollee in the previous year's choice. Similar to Handel (2013), we include the $\sum_s \beta_s S_{sij}$ term to capture the impact on utility of switching from one option to another.

We include three types of switching costs to fully capture the range of switching scenarios. First, consumers currently enrolled in Traditional Medicare face a *Medicare-to-MA* switching cost.

[9] This information is available to consumers through Medicare's Plan Compare web site.

Second, consumers who are currently enrolled in a MA plan face a *Different Insurer* switching cost if they switch to a plan offered by a different MA firm. Finally, consumers who are currently enrolled in a MA plan who switch to a different plan offered by the same insurer incur a *Same Firm* switching cost. We allow these switching costs to vary by demographic characteristics through the $\sum_s \sum_l \beta_{sz} z_{il} s_{sij}$ term.

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

Consumers who do not enroll in a Medicare Advantage plan may choose to purchase supplemental insurance, often called Medigap insurance. Although we do not model the choice of Medigap insurance directly, we incorporate this feature of the market by allowing the utility of the outside good to vary with the price of a benchmark Medigap insurance plan P_{i0}, which may vary by age and gender:

$$u_{i0} = \beta_0 p_{i0} + \epsilon_{i0} \tag{2}$$

We normalize the expected utility of the outside good to 0 by subtracting $\beta_0 p_{i0}$ from each of the products j. Following Berry et al. (1995), we rewrite the utility obtained from good j into a mean:

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

and an individual specific deviation:

$$\begin{aligned}\mu_{ij}^{'} = \sum_g \alpha_g p_j d_{gi} + \sum_k \sum_a \sum_h \beta_{kh} OOPC_{kahj} d_{ahi} \\ + \sum_s \beta_s switch_{sj} + \sum_s \sum_l \beta_{sz} Z_{il} switch_{sj} + \beta_z Z_i - \beta_0 p_{i0} + \epsilon_{ij}\end{aligned}$$

Given a set of plans J_f for a market m, the distributional assumption on ϵ_{ij}, and letting $\mu_{ij} = \mu_{ij}^{'} - \epsilon_{ij}$ the probability that consumer i chooses plan j (i.e. the share function) can be written as:

$$Pr(i \text{ chooses } j) = s_{ij} = \frac{\exp(\delta_j + \mu_{ij})}{1 + \sum_{j' \in J_m} \exp(\delta_{j'} + \mu_{ij'})} \tag{3}$$

3.1 Estimation

We broadly follow the estimation 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 indivdiual variation on plan valuation. These are the parameters that are included in μ_{ij}. With μ_{ij} in hand, we then estimate the parameters of δ_j.

For a given candidate value of the individual-specific parameters $\theta = \{\alpha_g, \beta_l, F_s\}$, we use the Berry (1994) contraction to find the unique set of product fixed effects $\delta_j(\theta)$ that match predicted shares to observed market shares. We estimate the μ_{ij} parameters using maximum likelihood. For an individual i who choses plan j, the likelihood function is

$$L(C_i; \theta, \delta(\theta)) = \Pi_j s_{ij}^{1\{C_i = j\}} \tag{4}$$

where $s_i j$ is given by Equation 3.

The first step of our estimate procedure maximizes the log-likelihood function over the space of θ. 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:

$$\delta_j = \alpha_o P_j + \beta X_j + \xi_j \tag{5}$$

where ξ_j is a product-specific unobservable.

If ξ_j is observed by firms and consumers, it is likely to be correlated with the premium and any rebate. We instrument for price using average MA plan premiums in other counties. Since the same plan is offered at the same premiums in many counties, we calculate this instrument using the set of "non-contiguous counties" – that is, for a plan in a county, we average all plans in counties which do not share borders with the county under consideration. We also employ instruments formed by summing product characteristics at the firm and market levels, per Berry et al. (1995), and the "differentiation" instruments of Gandhi and Houde (2016).

4 Data

We estimate our model with administrative data on product characteristics from CMS and micro-level data on consumer choices from the Medicare Current Beneficiary Survey (MCBS). In the following subsections, we discuss each of these data sources in turn.

4.1 Medicare Advantage Plans

We obtain the annual Plan Finder database from CMS, which contains comprehensive information on plan benefits on an annual basis.[10] We observe plan premiums, copays for doctor 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 observe the county-level geographic coverage for each plan segment. Finally, we separately obtain the bids made by each firm and the benchmark rate for each county-year.

The CMS data also contains estimates of 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 Traditional Medicare enrollees. CMS then calculates the out-of-pocket costs for that bundle under each plan's benefit structure.[11] We observe two OOPC estimates for each plan and demographic group: the total OOPC for the plan, and the OOPC for prescription drugs alone. We observe these two estimates by age group (65-69, 70-74, 75-79, 80-84, and 85+) and health status (Excellent, Very Good, Good, Fair, and Poor).

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 eligible for Medicare benefits to form product shares at the plan-county-year level. We obtain the benchmark subsidy rates from the U.S. Health Resources and Services Administration.[12]

Summary statistics on the plans in our data are reported in Table 1. More than a third of the plans offered have zero premium, and the vast majority of plans offer some form of supplementary coverage (most commonly prescription drug or vision coverage). CMS-estimated out-of-pocket costs vary considerably across age and health status groups, reflecting expected utilization of services.

[10] Recent years can be downloaded from `https://www.medicare.gov/Download/DownloadDB.asp`

[11] The underlying assumption is that enrollment in an MA plan does not significantly change the bundle of services consumed by beneficiaries.

[12] `https://datawarehouse.hrsa.gov/topics/ahrf.aspx`

Table 1: Summary statistics for MA plans

Variable	Mean	Std. dev.
Annual premium	\$571	\$598
Deductible	52.5	229
Out-of-pocket limit	2,855	1,766
Copays		
Primary physician visit	\$12.54	6.96
Specialist	26.9	11.5
Hospital stay	208	177
Supplemental coverage indicators		
Prescription drug	.621	.485
Dental	.488	.500
Vision	.878	.328
Plan type indicators		
HMO	.268	.443
PPO	.096	.294
FFS	.637	.481
Selected CMS-estimated out-of-pocket costs		
65-69 year-old, excellent health	\$2,334	398
75-79 year-old, good health	3,376	752
85+ year-old, poor health	4,322	1,171
Unique plans	5,043	
Year-county markets	1,034	
Obs.	31,454	

Note: An observation is a plan-county-year.

4.2 Medicare Beneficaries

Our data on individual participants in MA markets come from the Medicare Current Beneficiary Survey Cost and Use File, a rolling-panel survey of a nationally representative sample of Medicare recipients sponsored by CMS and produced by Westat. The survey is designed to obtain a complete picture of the Medicare system from the beneficiary perspective, including expenditures and payments for all medical services including services which aren't covered under the beneficiary's plan, changes in health status, prescription and hospitalization events, as well as satisfaction with care. Participants are interviewed multiple times per year over four years, and responses are linked to CMS claims data to ensure accuracy.

We obtain the MCBS survey responses from 2007-2010. We observe demographic information, including income, age, sex, race, education. Individuals self-report health status, choosing between Excellent, Very Good, Good, Fair, and Poor. We transform these variables into indicators for each demographic group. We also observe the individual's home county and plan choice.[13] We exclude any individuals with missing address information, individuals who are living in Puerto Rico, individuals who have insurance provided by an employer, individuals under 65, or obtained their Medicare eligibility for some reason other than age, and individuals who were eligible for Medicaid at any point during the year.

The MCBS does not cover every county in the country. Instead, it employs a tri-level sampling procedure to construct a nationally representative sample of Medicare recipients. The data include sampling weights which we use throughout our estimation procedure to obtain national-level parameter estimates, where appropriate.

Summary statistics on our 20,565 individual-year observations are reported in Table 2. The mean age of individuals in our data is 77. Slightly more than half of our observations are of females. Over three-quarters of the individuals surveyed self-reported "Good" or better health. Finally, 23% of individuals report having college degrees. 18% did not graduate high school.

To capture potential variation in the value of the outside good across geographies, we obtain Medicare supplemental insurance rates from Weiss Ratings. Medicare supplemental insurance, more commonly referred to as Medigap, is coverage that sits on top of TM coverage and pays for care that is not covered under the TM benefit design. For example, TM covers 80% of the cost

[13] The MCBS does not report the plan choice directly. Instead, it reports which firm the individual has chosen, along with information about the premiums paid by the individual and the plan features. We match this data to the plan data to identify each individual's choice.

Table 2: Summary statistics for individuals included in Medicare Current Beneficiary Survey

Variable	Mean	Std. dev.
MA enrollment indicator	.08	.462
Income	$38,192	44,615
Age	75.1	7.56
Outside Good Price	$2,199	503
Demographic indicators		
Female	.559	.497
Black	.070	.255
Hispanic	.010	.100
Education indicators		
Bachelor's degree or higher	.198	.398
Attended college	.289	.453
Graduated high school	.309	.462
Health status indicators		
Excellent	.189	.391
Very Good	.337	.473
Good	.315	.465
Fair	.122	.327
Poor	.037	.189
Obs.	12,091	

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.

of physician services and Medigap coverage may cover the remaining 20%. The benefit structure of Medigap plans is standardized 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. Plan C covers most of the coinsurance and deductibles that beneficiaries are responsible for under TM and is the most popular Medigap plan.[14]

5 Demand Parameter Estimates

Parameter estimates for the first stage of the demand estimation procedure are in Table 3 and Table 4. Table 3 presents the impact of the interactions of demographic and health status variables with plan characteristics, while Table 4 presents the switching cost parameter estimates. In the

[14] 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: 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	0.205	0.082
High Income	0.162	0.083
Total Out-of-Pocket costs (per $1000)		
Excellent health	-0.220	0.156
Very good health	-0.171	0.141
Good health	-0.185	0.121
Fair health	-0.167	0.101
Poor health	-0.194	0.081
MA × Demographic characteristics		
Age	-0.165	0.078
Female	-0.075	0.100
Black	-0.070	0.188
Hispanic	-0.071	0.415
Graduated high school	-0.003	0.140
Some college	-0.119	0.138
Bachelor's degree	-0.220	0.157
Medigap price (per $1000)	0.739	0.107
Weighted Log Likelihood	-14,365	
Observations	12,091	

first panel of Table 3 the estimates indicate that high and medium income consumers are less price sensitive than low income consumers. Out-of-pocket parameters are all negative but the parameter is only significantly different from zero for the poorest health beneficiaries. The only significant demographic variable is age – older beneficiaries recieve less value from MA enrollment than their younger counterparts. As expected, an increase in Medigap premiums increases the value of MA plans.

The three panels of Table 4 display the switching costs parameters for Traditional Medicare to MA switching costs, inter-insurer switching costs and intra-insurer switching cost, respectively. The highest switching costs are incurred by consumers switching from Medicare to Medicare Advantage, with slightly lower costs incurred by those switching between firms offering plans under different Medicare Advantage contracts and still lower costs for those switching between individual plans offered under the same contract. These results suggest the primary component of switching costs is the disutility of changing providers. Under Traditional Medicare beneficiaries have access to virtually any providers while MA plans generally uses more a tightly defined network of providers.

Of course, different MA plans included different providers in their networks. However, different plans under the same contract generally deploy the same network of providers. In general, the demographic variables are not correleted with switching costs.

The second stage results are in Table 5. 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 coeffient is much large in magnitude than the OLS. For this reason, we concentrate our attention on the IV specification.

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 -4.35. While lower than is typically estimated using market share level data, we believe this elasticity to be sensative given that these are a select group of plans with generally small market shares. Increases in the deductible, mean out-of-pocket costs, and copays reduce mean utility. Beneficiaries value prescription drug, vision and dental coverage. The one parameter estimate that is, at first glance, counter-intuitive is on the rebate. Recall that an increase in the rebate implies that the plan can offer more generous benefits. However, we are controlling for the most valued benefits so that coefficient on rebate reflects the value of increasing the rebate conditional. Nevertheless, we do not have a iron clad explanation for this coefficient. One possibility is that high rebate and conditional on benefits correlate with more nitch plans targetted to specific populations.

With the first and second stages in hand, we can transform the switching cost parameter estimates into monetized values of the implied switching costs. These monitized values are shown in Table 6. Sicker beneficiaries are less sensitive to expected out-of-pocket costs, though it's important to note that the expected costs are higher for individuals in poor health than individuals in excellent health.

To illustrate the importance switching costs on the utility of MA, we simulate an individual entering the Medicare program at age 65 since they have not previously selected TM or a MA plan, they have no switching costs, and calculate the consumer surplus they receive each year. Each year they face the same set of plan choices, though the particular product characteristics they face change over time as they age. Figure 3 shows the result of the simulation averaged across all markets. In the first year, with no switching costs, the average annual consumer surplus across markets is nearly \$1,000. By age 70, the annual consumer surplus has decreased to only \$231. As consumers move in and out of out-of-pocket-cost groups, the consumer surplus varies slightly.

Table 4: Demand parameter estimates for switching costs with demographic interactions

Variable	Coefficient	Std. Err.
Traditional Medicare-to-MA switch ×		
Constant	-4.607	0.697
Age	0.125	0.087
Female	0.023	0.112
Black	0.340	0.211
Hispanic	0.040	0.449
Graduated high school	0.055	0.160
Some college	-0.008	0.158
Bachelor's degree	-0.207	0.178
Inter-firm MA switch ×		
Constant	-1.872	1.218
Age	-0.215	0.157
Female	-0.064	0.191
Black	0.147	0.314
Hispanic	-0.864	1.445
Graduated high school	-0.090	0.256
Some college	-0.168	0.261
Bachelor's degree	-0.504	0.347
Intra-firm MA switch ×		
Constant	-1.834	0.842
Age	0.065	0.106
Female	0.248	0.137
Black	0.606	0.247
Hispanic	0.389	0.442
Graduated high school	-0.052	0.185
Some college	0.073	0.193
Bachelor's degree	0.157	0.212
Weighted Log Likelihood	-14,365	
Observations	12,091	

Table 5: Second stage demand parameter estimates

Variable	**OLS**	**IV**
Annual premium (per \$1000)	-.785 (.025)	-1.70 (.163)
Rebate (per \$1000)	-451 (.034)	-2.96 (.187)
Deductible (per \$1000)	-.143 (.052)	-.745 (.084)
Out-of-pocket limit (per \$1000)	-.010 (.007)	-.026 (.013)
Coverage flags		
Prescription drugs	.582 (.025)	1.03 (.067)
Vision	-.615 (.036)	.114 (.094)
Dental	-.204 (.023)	.154 (.080)
Copays		
Primary doctor	-.015 (.002)	-.063 (.006)
Specialist	-.016 (.001)	-.033 (.004)
Hospital stay (per \$1000)	.664 (.069)	-.039 (.102)
HMO flag	1.15 (.030)	2.01 (.072)
PPO flag	1.47 (.043)	1.24 (1.23)
Fixed effects	None	Firm-level
Mean implied elasticity (if < 0)	-.560 (.363)	-1.33 (.860)
Mean ds_j/dp_j	-.0028 (.0076)	-.0067 (.0179)
Observations	31,412	31,412

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

Table 6: Implied dollar switching costs

Switch type	Low income	Medium income	High income
Medicare-to-MA	\$712	736	754
Between-firms	598	618	633
Within-firm	176	182	187

Figure 3: Average MA consumer surplus by year for a new Medicare beneficiary

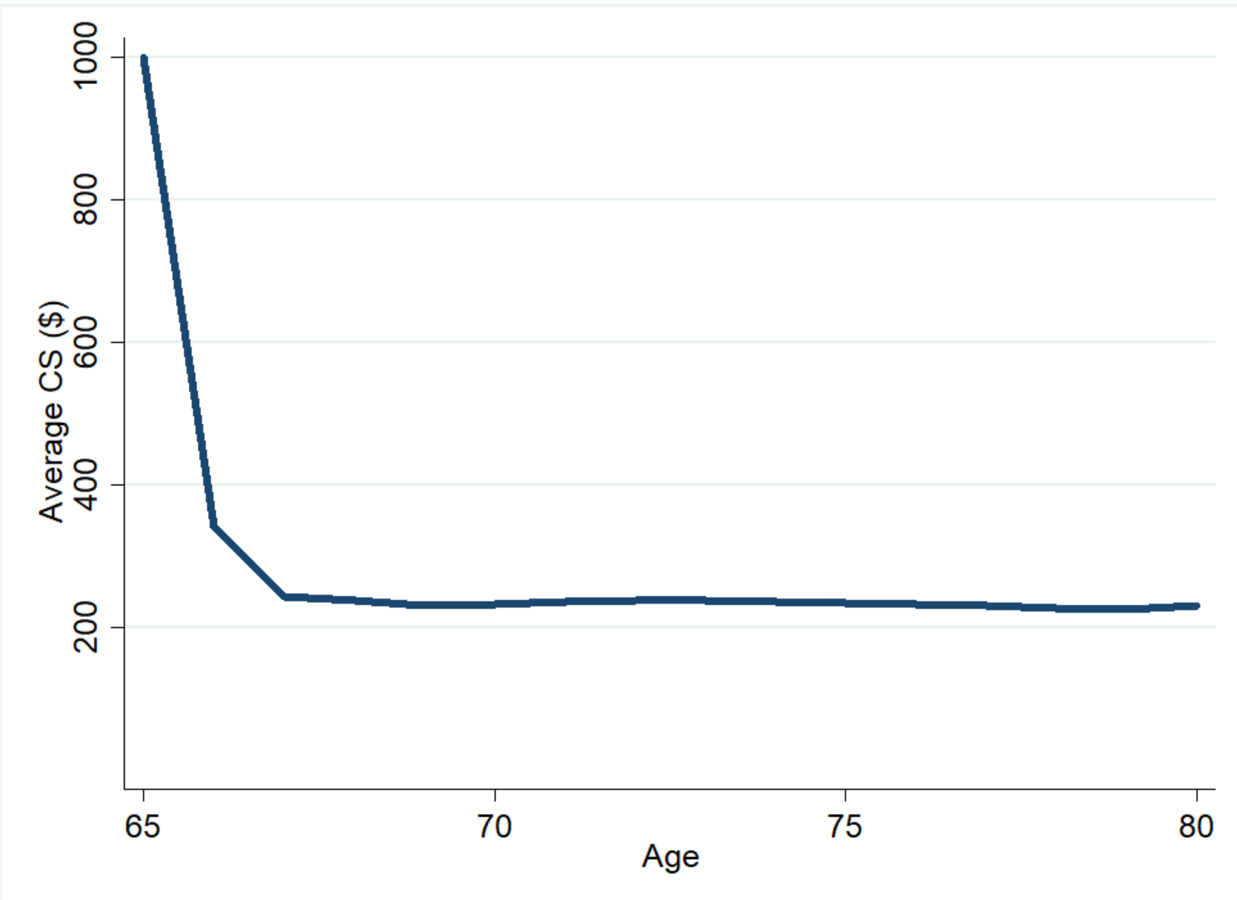

6 Supply-Side Responses to Alternative Benchmarks

In order to estimate the impact of changes in the benchmark on consumer welfare we need to model the insurer's response to changes in the benchmark. The obvious approach would be to construct a model in which firms set premiums and plan characteristics, use the estimated demand elasticities to back out the underlying cost parameters and solve for equilibrium premium and plan characteristics under different benchmarks. There is a growing literature estimating models with endogenous product characteristics and several papers take the approach described above.[15] For example, Starc and Town (2018) use this approach to study benefit design in Medicare Part D, and Fan (2013) uses set-up to study the impact of newspaper mergers on product attributes. With the demand parameters in hand, we could then measure the consumer surplus impact of the plans responses to different benchmarks.

Given the number of product characteristics observed for MA plans, this approach is not practical to implement without making a number of unrealistic assumptions. Instead, we use a hybrid approach. We use firms' first-order conditions for price-setting conditional on product characteristics to recover marginal costs as a function of those characteristics. We then estimate a flexible policy function that maps benchmarks into δ_j net of premiums. For a given benchmark, we then compute the new δ_j (net of premiums) and the new marginal costs. Given the new δ_j and marginal costs, we solve for the equilibrium in premiums.[16]

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

$$\max_{\{B_m\}} \sum_m \sum_i CS_{im}(B_m) \text{ s.t. } \sum_m GovExp_m(B_m) = \bar{M} \tag{6}$$

$$\max_{\{B_m\}} \alpha \sum_m \sum_i CS_{im}(B_m) - (1-\alpha) Var(CS) \text{ s.t. } \sum_m GovExp_m(B_m) = \bar{M} \tag{7}$$

[15] Crawford (2012) provides a review of this literature.

[16] In Appendix B we provide Monte Carlo evidence on the accuracy of our approach compared to explicitly solving for both prices and product characterstics in a static oliopoliy framework. In our simulations, we find that the approximation approach we use yields solutions very close to the actual equilibrium solutions.

$$\max_{\{B_m\}} \sum_m \sum_i CS_{im}(B_m) \text{ s.t. } B_m > \bar{B} \forall m \text{ and } \sum_m GovExp_m(B_m) = \bar{M} \tag{8}$$

The first equation simply maximizes consumer surplus subject to a government budget constraint. As we document below, the solution to this problem results in signficant variation in the optimal benchmarks and the resulting consumer surplus. This variation may not be politically feasible to impliment. For this reason, Equation 7 includes a penalty for the variance in consumer surplus across individuals with $\alpha \in [0, 1]$ determining the relative weight that is place on consumer surplus versus the variance in consumer surplus. Currently, the CMS benchmark formula includes floors on the benchmark amount suggesting that is the perfered policy solution to addressing a low implied benchmark. We therefore also solve for the optimal benchmarks using the objective function in Equation 8. This objective function includes a constraint that all benchmarks must be set above some floor level. Solving these problems requires calculating both the CS and $GovExp$ functions.

Given our assumptions on the utility function, the consumer surplus for an individual i, given prices and product characteristics, is calculated as (Small and Rosen (1981)),

$$CS_i = E[\max_j U_{ij}]/\alpha_i = \frac{1}{\alpha_i} \ln \left(\sum_j \exp U_{ij}(\delta_j, p_j, x_{ij}) \right) \tag{9}$$

where $U_{ij} = \delta_j + \mu_{ij}$ and α_i is the price sensitivity of individual i.

If individual i is enrolled in Traditional Medicare, the government will incur medical costs TM_i for that beneficiary. If, instead, individual i is enrolled in a MA plan, the government will pay the flat rate B_m times some risk adjustment factor R_i for that individual. Let $s_i^{MA} = \sum_j s_{ij}$ be the probability that individual i purchases an MA plan, and $\mathcal{A}_m$ be the set of Medicare beneficairies in m, the government expenditure function for a given market can be written as

$$GovExp_m(B_m) = \sum_{i \in \mathcal{A}_m} \left(s_i^{MA} R_i B_m + (1 - s_i) TM_i \right). \tag{10}$$

The MCBS data includes actual government spending for each Traditional Medicare enrollee (but not for MA enrollees) and a rich set of observables for all beneficiaries. We calculate TM_i for each individual in our data by regressing TM expenditures on our observables and projecting those estimates onto everyone in the data.[17] We form R_i for each individual by taking the ratio of the

[17] The TM cost regression assumes that there is no selection on unobservables. We have a very rich set demo-

estimate for each individual to the average estimated TM_i across all individuals.

In our approach we allow insurers to respond to changes in the government's benchmark rate by adjusting both plan premiums and characteristics. Since U_{ij} and s_{ij} are functions of prices and characteristics, we incorporate firm responses into our calculation of Equations 9 and 10. For computational reasons, we do not simulate equilibrium in this stage of the game directly. Instead, we separate the pricing decision from determining the other product characteristics. Let δ' be the 'price-adjusted- for each plan given by

$$\delta'_j = \delta_j - \alpha_0 P_j \tag{11}$$

Given our demand results, we can calculate $\hat{\delta'}_j$ for each plan. We also allow for the out-of-pocket costs that vary by age and health status to be adjusted by the plans when they face a different benchmark. We then estimate the following functions

$$\delta'_{jm} = g^{\delta}(j, f, B_m, x_m) + e^{\delta}_{jm} \tag{12}$$

$$OOPC_{hj} = g^{O}\delta(j, f, B_m, x_m) + e^{O}_{jm} \tag{13}$$

where f is the firm providing the plan, x_m is a vector of market characteristics, and e_{jm} is an unobservable.

Given a benchmark for a market, we use our estimate of the g functions to calculate price adjusted δ's and the updated OOPC's for each plan in that market. However, a firm that changes its product characteristics likely faces different costs as well, which will affect their pricing decision. Under the assumptions that firms face constant marginal costs and play a Nash pricing game, we recover marginal costs for each plan, mc_j using the approach of Berry et al. (1995). We can then estimate the implied marginal costs as a function of the δ''s as

$$mc_{jm} = h(j, f, \delta'_{jm}) + v_{jm}. \tag{14}$$

We estimate the g function using OLS allowing the benchmark to nonlinearly and flexiblely interact with the other arguments of g. This regressions include with plan and county fixed effects. In estimating the h function we also allowed for nonlinearieties with respect to δ' but the estimates

graphic, location and health status variables that we use as predictors so we believe the importance of selection on unobservables is largely mitigated. In this analysis we are also ignoring any spillovers from the MA sector to TM (Chernew et al. (2008), Baicker et al. (2013).

Table 7: Traditional Medicare spending regression

Variables	(1) logmedexp
Annual income	6.76e-07***
	(2.41e-07)
Female indicator	0.0612**
	(0.0256)
Bachelor's degree indicator	0.160***
	(0.0392)
Some college indicator	0.176***
	(0.0360)
Graduated high school indicator	0.116***
	(0.0348)
Black indicator	-0.338***
	(0.0528)
Hispanic indicator	-0.194
	(0.162)
Age	0.0282***
	(0.00156)
Inpatient stays	0.945***
	(0.0262)
Inpatient covered days	0.0168***
	(0.00390)
Inpatient coinsured days	-0.0940***
	(0.0186)
Ever smoked indicator	-0.0153
	(0.0234)
Health status indicators[1]:	
Excellent	-1.003***
	(0.0665)
Very good	-0.748***
	(0.0632)
Good	-0.487***
	(0.0623)
Fair	-0.166**
	(0.0661)
Diagnosis indicators[2]	Incl.
Fixed effects	County
Observations	15,197
R-squared	0.403

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

[1]Health status is self-reported and is relative to "Poor."

[2]The MCBS asks questions of the form "Have you ever been told you have *condition*?" We include indicators for affirmative responses to this question for the following conditions: high cholesterol, hardened arteries, high blood pressure, heart attack, angina, congenital heart defect, valve conditions, heart rhythm disorders, other heart conditions, stroke, cancer, rheumatoid arthritis, other arthritis, Alzheimer's disease, depression, other psychological disorders, hip replacement, Parkinson's disease, chronic obstructive pulmonary disorder, and paralysis.

indicated that a linear specification is a good approximation so that is the specification we use. We also explored estimating the functions using splines and the results from these specifications were very similar to our base specifications.

With all of these functions in hand, given a proposed benchmark for a given county, we calculate CS_m and $GovExp_m$ as follows:

1. Use the estimated g and h functions to calculate δ'_{jm} and mc_{jm} for each plan.
2. Taking as given these new product characteristics and costs, solve for the new equilibrium in prices.
3. Given the new prices and characteristics, use Equations 9 and 10 to calculate the consumer surplus and government spending for the county.

This process can be used to calculate the impact of alternative policies across the maximization problems represented by Equations 6, 7, and 8.[18]

Table 7 reports the results of our regression of the log of Traditional Medicare spending on individual observables for the individuals in our data who are enrolled in Traditional Medicare. In general, the results are in line with expectations. Older beneficiaries are more expensive; each year of a beneficiary's age is associated with a 2.8% increase in spending. Each inpatient episode nearly doubles spending, and each covered day in an inpatient facility increases spending by 1.7%. Our indicators for self-reported health status capture a wide dispersion in spending that follows the natural ordering of the health status categories.

Table 8 details our specification of g, which captures firms' decisions about product characteristics (in the form of our price-adjusted-delta measure) as a function of the benchmark rate. We regress the price-adjusted-delta on the quadratic form of the benchmark and several geographic measures, including the population in the county, the percentage of the county's population which is eligible for Medicare, the percentage below the poverty line, the numbers of hospital beds and MDs per capita, the population per square mile, and the income per capita. We interact each of these measures with the benchmark and the benchmark squared. Finally, we include firm fixed

[18] Given the number of firms and plans in our data, solving for the Nash Equilibrium in prices is computationally intensive. Since there are not closed form solutions for these equilibrium prices, derivative information is not available and gradient-free optimization algorithms must be used, increasing the number of points that must be evaluated. We take advantage of the separable nature of the maximization problems discussed in this section and calculate CS_m and $GovExp_m$ over a grid of benchmarks for each county to form a spline approximation of each function. To solve the maximization problem, we use these spline approximations to quickly find candidate solutions, which can be done with gradient-based solutions. We refine these candidate solutions using the more precise Nash Equilibrium solutions.

Table 8: Price adjusted-delta regression

Variables	(1) Adjusted delta
Benchmark	-0.0167
	(0.0669)
Benchmark^2	1.84e-06
	(7.22e-06)
County Population	1.88e-06
	(1.55e-06)
% Medicare Eligible	103.1***
	(33.50)
% Below Poverty Line	-0.884***
	(0.341)
Hospital Beds per Capita	216.4
	(991.1)
MDs per Capita	1,094
	(1,344)
Population per Square Mile	0.000303
	(0.000540)
Income per Capita	-0.000723***
	(0.000234)
County Population * Benchmark	-4.55e-10*
	(2.76e-10)
% Medicare Eligible * Benchmark	-0.0197***
	(0.00643)
% Below Poverty Line * Benchmark	0.000192***
	(6.51e-05)
Hospital Beds per Capita * Benchmark	-0.0329
	(0.197)
MDs per Capita * Benchmark	-0.349
	(0.255)
Population per Square Mile * Benchmark	-1.07e-07
	(8.88e-08)
Income per Capita * Benchmark	1.54e-07***
	(4.37e-08)
County Population * Benchmark^2	2.32e-14*
	(1.23e-14)
% Medicare Eligible * Benchmark^2	9.59e-07***
	(3.07e-07)
% Below Poverty Line * Benchmark^2	-1.00e-08***
	(3.07e-09)
Hospital Beds per Capita * Benchmark^2	1.03e-06
	(9.70e-06)
MDs per Capita * Benchmark^2	2.26e-05*
	(1.20e-05)
Population per Square Mile * Benchmark^2	6.34e-12*
	(3.63e-12)
Income per Capita * Benchmark^2	-7.94e-12***
	(2.01e-12)
Fixed effects	Firm
Interactions	Firm * Benchmark, Firm * Benchmark^2
Observations	34,648
R-squared	0.434

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

Table 9: Marginal cost regression

Variable	Marginal Cost
Adjusted Delta	140.4***
	(0.553)
Fixed effects	County, Firm
Observations	34,648
R-squared	0.958

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

effects and interact those effects with the benchmark and the benchmark squared. The result is a flexible function that captures heterogeneity in product characteristics across firms and counties.

Table 9 reports the results of our estimation of h, which captures how firms decisions about product characteristics flow through to marginal costs, without an explicit reference to the benchmark rate. We regress the estimated marginal cost on the price-adjusted-delta and county and firm fixed effects. These two functions are used to generate the CS_m and $GovExp_m$ functions for each county.

7 Counterfactuals

We begin by presenting the results of the regressions that lead to our construction of the CS_m and $GovExp_m$ functions. We then describe the results of solving variations of Equations 6, 7, and 8.

It is first useful to understand some of the variation in counties that leads to the results of the policies. Figure 4 presents share-weighted adjusted deltas as a function of the benchmark for three counties in our data. To create each line, we simulate market outcomes across a range of benchmarks, and then weight the adjusted delta product characteristic for each plan by that plan's market share. The results show significant variation in both the quality of the plans offered across different markets, as well as differences in the market-level responses to changes in the benchmark.

Table 10 reports the results of the optimal policies that result from solving the maximization problems outlined in Section . Column (1) reports consumer surplus across a number of dimensions under the current policy. Consumer surplus from MA is higher for Hispanic and Black beneficiaries than for White beneficiaries. Low income people benefit more from the program than do high income people. Column (2) reports the results of solving Equation 6. The mean consumer surplus

Figure 4: Simulated share-weighted adjusted delta for three markets

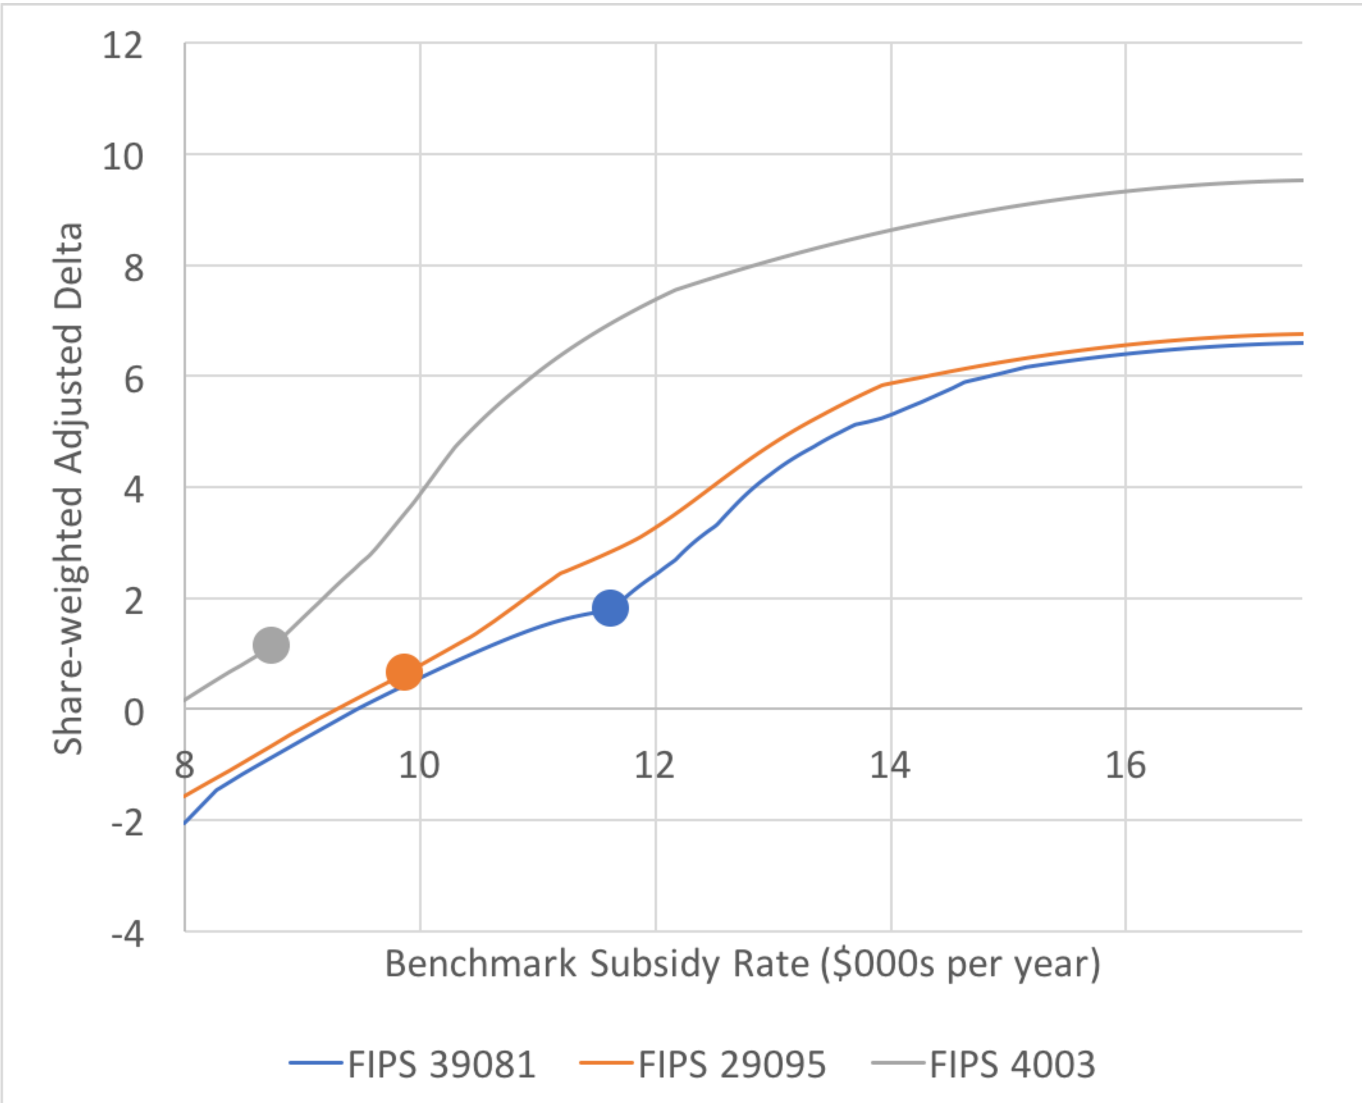

The data in this figure is calculated by solving for market outcomes for each value of the benchmark rate, and weighting the adjusted delta product characteristic for each firm by the market share of that firm. The large dots on each line indicate the actual policy in 2010.

Table 10: Counterfactual results for two objective functions

Variable	(1) Current Policy	(2) Max CS	(3) Max 1/2 CS − 1/2 Var
Mean CS	141.04	506.82	18.87
Variance of CS	25,906	923,652	546.56
Median CS	91.68	1.24	11.51
25th Percentile CS	48.51	0.06	4.49
75th Percentile CS	160.69	382.61	25.44
Mean White CS	139.22	529.42	18.30
Mean Black CS	146.93	365.91	23.50
Mean Hispanic CS	213.37	145.39	19.15
Mean Low Income CS	133.69	460.78	20.77
Mean Medium Income CS	170.39	575.33	20.11
Mean High Income CS	121.47	486.28	16.12

increases, from $26.74 under the baseline policy to $92.58. These gains come with a large increase in the variance, from 339 to 28,950. In particular, the optimal policy under scenario (2) sets the optimal benchmarks close to zero in many counties, and uses the resulting savings to increase benchmarks, and thus surplus, in counties where the market process more effectively transforms benchmark payments into consumer surplus. Whites and Blacks gain substantially under this policy, while Hispanic consumer surplus is lowered.

Column (3) reports the results of solving Equation 7, which adds a penalty term for the variance in consumer surplus. The resulting policy lowers mean consumer surplus slightly, from $26.74 to $25.60, but decreases the variance to 46.31. Again, Whites and Blacks gain under this policy, at the cost of Hispanic consumer surplus. Finally, Column (4) reports the results of solving Equation 8, which returns to the maximum CS objective function but keeps the current floor in benchmark rates. The results are almost identical to Column (2), though mean consumer surplus is slightly smaller, as is the variance.

The differences in the polices generated by the different social welfare objective functions can be see in Figure 5. The left and right histograms show the distribution of the change in benchmark rates for the specification in which CS is maximized and for the specification in which a linear combination of CS and the variance of CS is maximized, respectively. The results are striking. On the left hand side, the counterfactual policy reduces benchmarks dramatically, by over $5,000 in some cases, to fund equally large increases in several counties. When the variance in consumer surplus is taken into account, the policy reduces the benchmark rate a smaller amount in a smaller

Table 11: The effects of alternative welfare weights

Variable	Current Policy	(1) Max 1/2 CS − 1/2 Var	(2) Max 0.99 CS − 0.01 Var	(3) Max 0.995 CS − 0.005 Var
Mean CS	141.04	18.87	101.16	151.62
Variance of CS	25,906	546.56	5,859	12,947
Median CS	91.68	11.51	83.40	123.81
25th Percentile CS	48.51	4.49	38.70	62.85
75th Percentile CS	160.69	25.44	150.67	221.05
Mean White CS	139.22	18.30	99.09	148.77
Mean Black CS	146.93	23.50	119.03	177.44
Mean Hispanic CS	213.37	19.15	94.35	133.07
Mean Low Income CS	133.69	20.77	103.05	151.19
Mean Medium Income CS	170.39	20.11	111.96	168.92
Mean High Income CS	121.47	16.12	89.96	136.70

number of counties to fund more modest increases across the remaining counties.

Figure 5: Changes in benchmark rates for two social welfare functions

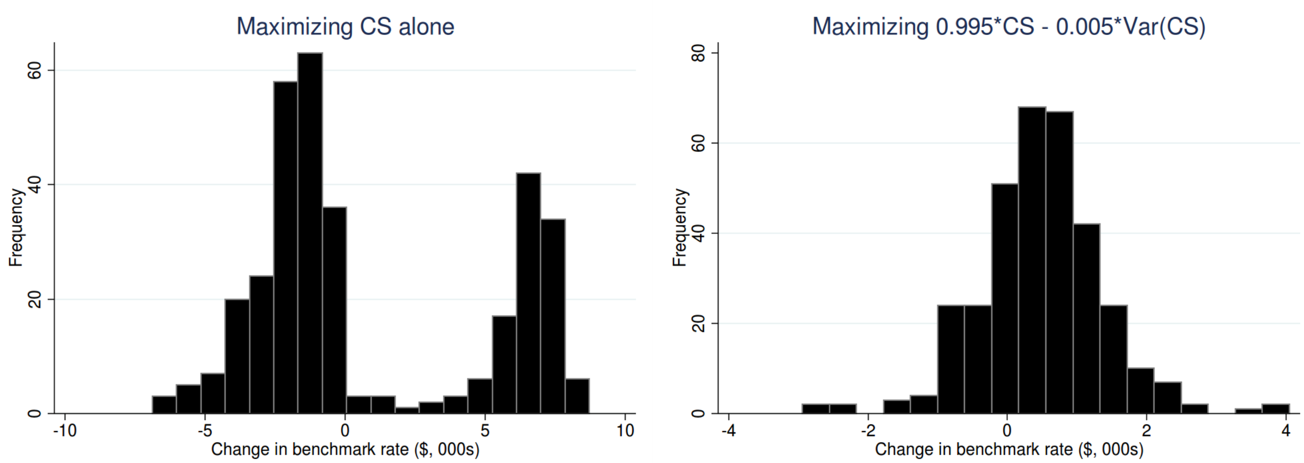

8 Conclusion

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 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 a hybrid approach for calculating the supply decisions. We apply our framework to the Medicare Advantage program in the United States, through which approximately one-third of U.S. 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 \$125.19 in consumer surplus per person per year. By maximizing a linear combination of consumer surplus and the variance of consumer surplus, we find an alternative policy that results in \$338.79 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

Baicker, K., Chernew, M. E., and Robbins, J. A. (2013). The spillover effects of medicare managed care: Medicare advantage and hospital utilization. *Journal of health economics*, 32(6):1289–1300.

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

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. (forthcoming). Does privatized medicare benefit patients or producers? evidence from the benefits improvement and protection act. *American Economic Review*.

Chernew, M., DeCicca, P., and Town, R. (2008). Managed care and medical expenditures of medicare beneficiaries. *Journal of Health Economics*, 27(6):1451–1461.

Crawford, G. S. (2012). Endogenous product choice: A progress report. *International Journal of Industrial Organization*, 30(3):315–320.

Curto, V., Einav, L., Levin, J., and Bhattacharya, J. (2014). Can health insurance competition work? evidence from medicare advantage. Working Paper 20818, National Bureau of Economic Research.

Decarolis, F., Polyakova, M., and Ryan, S. P. (2016). The welfare effects of supply-side regulations in 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.

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.

Gandhi, A. and Houde, J.-F. (2016). Measuring substitution patterns in differentiated products industries. *Working Paper*.

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

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

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

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

McGuire, T. G., Newhouse, J. P., and Sinaiko, A. D. (2011). An economic history of Medicare part C. *Milbank Q*, 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., 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.

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

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.

Starc, A. and Town, R. J. (2018). Externalities and benefit design in health insurance. Technical report.

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, $m\hat{c}_{jmt}$. We then regress the implied δ_{jm1} on $m\hat{c}_{jmt}$ 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 A.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 A.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 A.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 A.1: Monte Carlo Results for Seven Firms