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

January 2019

Abstract

When markets fail to provide socially optimal outcomes, governments often intervene through 'managed competition' where firms compete for per-consumer subsidies. Subsidies are generally set across geographies according to estimates of average consumer-level expenditures, a method which may not be welfare-maximizing. We introduce a framework for determining the optimal subsidy schedule that allows for differences in consumer preferences, switching costs, potential moral hazard, and differences in private costs. We apply it to the Medicare Advantage (MA) program, which offers Medicare recipients private insurance that replaces Traditional Medicare (TM). The optimal schedule increases mean consumer surplus by 30% over current policy. One-third of this improvement comes from encouraging individuals to switch from TM to MA where MA has a cost advantage. The rest comes from redistributing funds to areas where marginal expenditures create the largest welfare increases. These redistribution incentives are driven more by supply-side heterogeneity than demand-side heterogeneity.

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. We are grateful for support from the National Institute on Aging under grant P01 AG032952.

1 Introduction

When the private market's provision of a good deviates from the socially optimal outcome, welfare may be improved through government intervention. Often times the obvious intervention is direct government provision of the good or service. 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 approaches also 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 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 "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 government-operated 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 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 Affordable Care Act. Elements of the managed competition approach also appear in education, where public, charter, and private primary and secondary schools compete on program offerings, education quality, and productive efficiency (Hoxby, 2000, Levin, 2002), as well as low-income housing policy, where differences in tax credits by geography incentivize new construction (Baum-Snow and Marion, 2009).

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. Simple examples demonstrate that the optimal subsidy should also vary across geographies as a function of those conditions. Broadly speaking, the optimal subsidy depends on the costs of the private firms, the willingness-to-pay of consumers, and the way in which competition leads firms to transform additional subsidy dollars into consumer surplus. In practice, however, geographic variation in subsidies is typically driven solely by mea-

sures of the *government's* average cost of providing the service, which may vary substantially from private firms' costs and may be unrelated to demand conditions. For example, the MA benchmark rate is set solely as a function of county-level average TM costs.

In this paper, we develop an approach for determining the optimal subsidy level across many markets in managed competition settings conditional on a fixed level of government expenditures. In our setup, multiproduct firms choose equilibrium prices and product characteristics in response to the subsidy level set by the government and other competitive conditions.[1] Consumers exhibit preference heterogeneity across observable plan and consumer characteristics and unobservable product-specific 'quality' differences. We take the mechanism that links a county-level subsidy to payments to firms as fixed and focus on differences across markets. 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 apply our approach to the MA program. Using individual-level panel data on consumer demographics, option sets, and choices for the years 2008-2015, we provide descriptive evidence that that the current policy is suboptimal and that firms respond to different benchmark rates. The panel nature of our data allows us to estimate switching costs, which are relevant due to the prevalence of restrictive provider networks. Our demographic variables include a self-reported health status which allows us to flexibly capture demand behaviors which vary with health. We estimate the parameters of our model and find that MA plans are more valuable to those who have lower income, lower educational attainment, are younger, and are in better health. Switching costs are significant but plausible – the average cost to switch between MA insurers is \$523, which is comparable to the average annual plan premium of \$509. Firms costs differ both within and across markets. Our estimates imply that in 2015, the MA program generated a total of \$5.94 billion in consumer surplus and \$4.75 billion in variable profit with \$98 billion in payments to MA plans.

To find the optimal subsidy schedule, we must calculate market outcomes under counterfactual policies. The traditional approach to computing counterfactual equilibria, a fixed-point search over 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 estimated functions to predict characteristics under counterfactual benchmarks and then solve the firms' first-order conditions for prices taking those characteristics as given. We find

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

that, on average, an increase in government spending by $1 increases consumer surplus by $0.32, though this pass-through rate varies across markets. Market-level benchmark rates and government spending can have a non-monotonic relationship. For several markets, the current benchmark rates are set such that the government spends less on an MA enrollee than an equivalent TM enrollee. In these markets, for some benchmark rate 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.

We find the optimal benchmarks are meaningfully different than the 2015 policy – the average difference between the optimal benchmark and the 2015 benchmark is 3.6%. The change can be dramatic – 14% of the markets have a optimal benchmark that is 10% greater than the benchmark in our data. The estimates imply that the optimal benchmark will increase average consumer surplus 30% from $149 to $193 per Medicare-beneficiary-year though with increased variance across consumers. Minimizing government expenditures with the constraint that benchmarks stay at least at their 2015 levels – in other words taking full advantage of intensive/extensive margin cost savings opportunities – results in an increase in the mean consumer surplus to $164. We conclude that one-third of the improvement from the current policy to the optimal policy comes from these opportunities and two-thirds comes from redistributing spending across markets. By exploring the distribution of the derivative of total surplus across markets – the key metric which drives the objective function – we find that this redistribution is driven more by heterogeneity in firm costs than heterogeneity in consumer preferences. We explore other social welfare functions and find benchmark schedules that increase the average consumer surplus and reduce the variance relative to the 2015 policy.

We build upon 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 generated $113 in consumer surplus and $244 in profits per Medicare beneficiary in 2000 with significant geographic variation. Curto et al. (2015) estimate a similar model 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 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.

There is a literature examining the rate at which MA benchmark increases are passed through to consumer surplus. Using an unanticipated change in the benchmark in 2000, Cabral et al. (2018) estimate a passthrough rate of 54%, 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 considering firms' choices of plan features, not just premiums, in response to benchmark changes. 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. 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) study health-status-linked premiums in the employer-sponsored setting. More broadly, we relate to a literature that considers various strategies designed to address adverse selection—see Geruso and Layton (2017) for a review. MA subsidies are risk-adjusted, 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 plan characteristics.

We discuss the institutional details of the Medicare Advantage program in Section 2, and provide simple examples of the optimal subsidy problem in Section 3. We detail our data on Medicare beneficiaries and MA plans in Section 4. We present our full model in Section 5 and estimate its parameters in Section 6. Section 7 describes our approach to counterfactual analysis, and we present the results in Section 8. We conclude in Section 9.

2 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) and Part B (physician and outpatient) insurance programs, respectively. Under its fee-for-service (FFS) reimbursement

structure, private providers treat Medicare beneficiaries and Medicare pays the provider according to a pre-set reimbursement schedule while beneficiaries pay applicable copays and/or coinsurance. In 1972 the Medicare program was expanded 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 the increasing costs of Medicare,[2] in 1982 Congress authorized Medicare administrators to engage in a series of "Part C" trials in which the government handed over management of the medical care of select groups of Medicare enrollees to private insurers in exchange for a payment that did not vary with the realized medical expenditures of each individual.[3] This program was brought to the entire country in the early 1990s under the name Medicare+Choice.

The Medicare+Choice program struggled to attract plans and nationwide enrollment hovered near 5 million – less than 10% of the Medicare population. 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 Traditional Medicare (TM) 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 demographic and diagnostic data about enrollees to the Centers for Medicare and Medicaid Services (CMS) at the time of enrollment. CMS assigns each enrollee a score based on its FFS expenditures on similar individuals in TM. Payments to firms are then adjusted according to these risk scores. 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 (MA).[4]

By 2015, the most recent year in our data, 95% of Medicare beneficiaries had an MA plan operating in their county and enrollments had more than tripled to 16.3 million. A beneficiary who enrolls in an MA plan forgoes TM benefits and receives medical benefits from the MA plan exclusively. MA enrollees 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 may offer lower out-of-pocket cost shares and Part B premium reductions. These increases in benefit generosity come with the typical managed care trade-offs of restrictive provider networks and referral requirements. Insurers compete along the dimensions of benefit design,

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

[3] These trial contracts were based in part on the ideas of Enthoven (1978).

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

premia, and provider networks, and often heavily market their plans (Aizawa and Kim, 2018).

The enrollee-specific subsidy from CMS to insurers is based on a fixed "benchmark" rate for each county, which varies significantly across geographies (Newhouse et al., 2012). CMS calculates the benchmark rate by starting with the average risk-adjusted per-capita FFS Medicare spending within the county.[5] Counties are then ranked by 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.

Each year, after benchmarks are released, 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. Firms who bid above the benchmark rate must charge premiums to enrollees. Firms who bid below the benchmark receive a portion of the difference as a 'rebate' that must be passed on to consumers through increased benefits or lower Medicare premiums. MA plans that offer a prescription drug benefit submit a separate bid which maps in a similar way to a Part D premium.

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 multiple dimensions and assigning a summary 'star rating' to firms. In 2012, after iterating on the rating criteria, CMS began using the 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 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 in the fall prior to the plan year. Beneficiaries may also enroll in MA when they become newly Medicare eligible and after "major" life changes (e.g. relocation, death in the family, etc.).[6] After the enrollment period closes, firms collect and transmit risk-adjustment information to CMS, which calculates the final

[5]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).

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

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

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

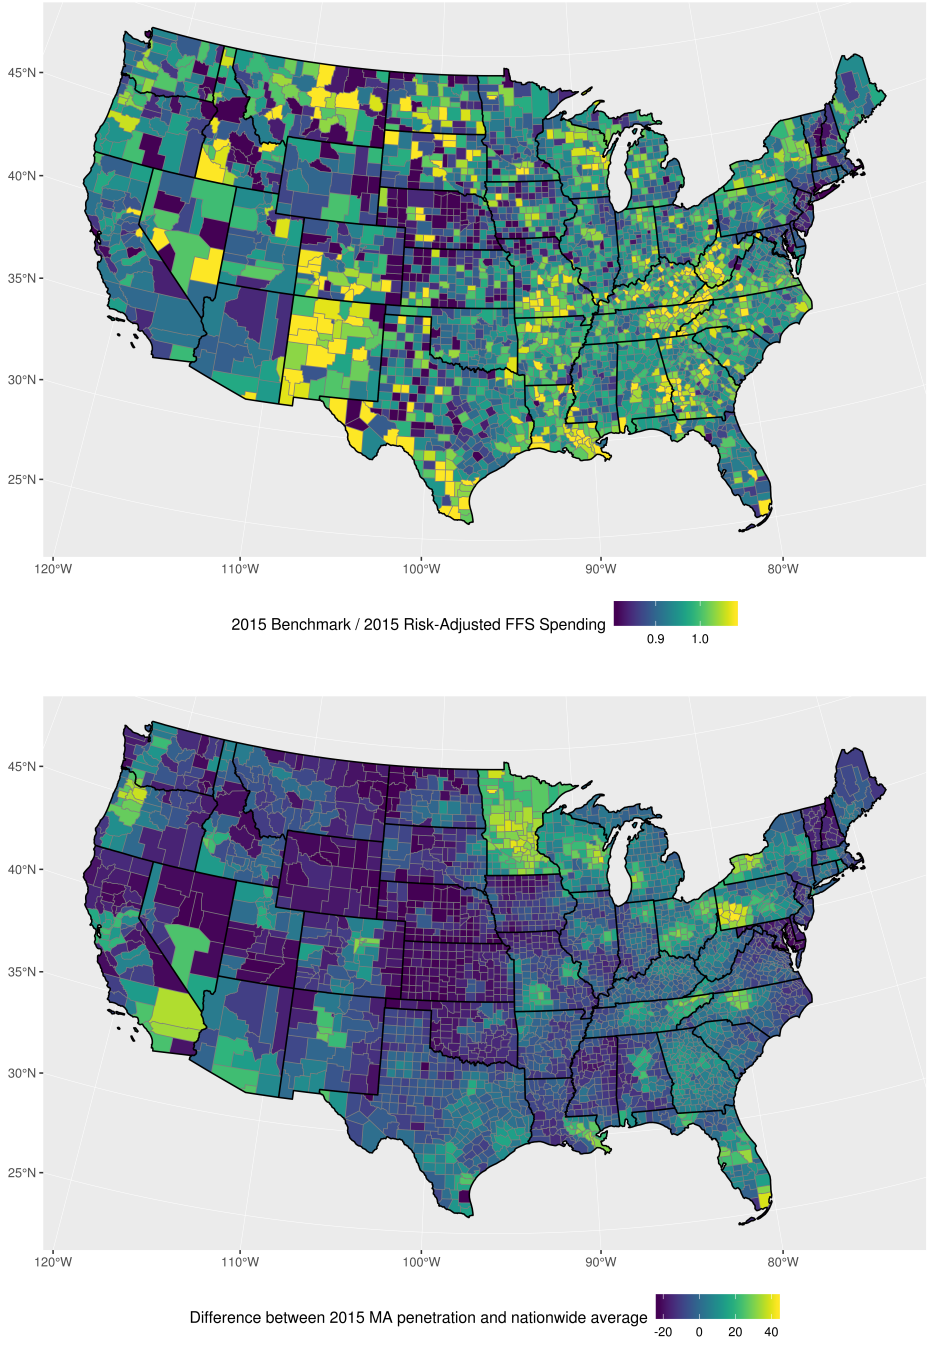
where ϕ_{jt} captures any bonus to the benchmark rate and λ_{jt} captures the impact of the star rating on the rebate percentage. For simplicity, we will denote the market-level (i.e. county-year level) base benchmark with B_m and denote risk-neutral plan-specific payments with B_j.

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 TM costs (MedPAC, 2017).

Figure 1 illustrates the 2015 policy and the resulting market outcomes with county-level maps of the US. The top map illustrates the ratio of the 2015 benchmark rate to the average FFS spending in 2015. The bottom map illustrates the difference between the MA enrollment percentage in each county and the national average. As consumer surplus is related to the total share of MA plans, these graphs offer a visual check of the optimality of the current government policy. If private costs are tightly linked to the government's costs, and the supply side is the only channel through which differences in policy lead to different outcomes, then we would expect those areas which had larger benchmarks (relative to FFS spending) to have greater enrollment (relative to the national average). Instead, we see a significant deviation from this pattern. Areas with high benchmarks, such as much of New Mexico, do not have particularly high enrollment. Furthermore, those areas with high enrollment, such as Minnesota and southwestern Pennsylvania, do not have particularly high benchmarks. We thus conclude that there are likely gains to be made by redistributing government funds across counties.[7]

[7] Appendix B presents other views of the benchmark distribution by county and across populations.

Figure 1: 2015 Medicare Advantage Benchmarks Relative to FFS Spending, and Relative Market Penetration, by County



Notes: Data from CMS benchmark and enrollment files. The top map illustrates the ratio of the 2015 benchmark rate to the 2015 risk-adjusted FFS spending in each county. To show detail, the data are windsorized at the 5th and 95th percentiles. The bottom graph shows the difference between the Medicare Advantage penetration rate in each county and the overall penetration. Penetration is defined as the total number of people enrolled in any MA plan divided by the number of people eligible for Medicare benefits. Positive numbers, therefore, indicate that the county had a higher percentage of Medicare eligibles enrolled in an MA plan than the national average.

3 Simple Examples of the Optimal Subsidy Problem

Before moving to the full model, we present simple examples of the optimal subsidy problem which illustrate the gains that can be obtained by considering the interactions between supply and demand conditions across markets. Consider a government which seeks to maximize consumer surplus by allocating a fixed subsidy budget $\bar{B}$ across M markets denoted by m. Under a managed competition framework, the government does not allocate spending directly, but instead chooses a schedule of market-level purchase subsidies $\{B_m\}$. Following the MA nomenclature, we refer to B_m as the market-level 'benchmark rate.' Let $CS_m(B_m)$ be the consumer surplus and $GovExp_m(B_m)$ be the government spending in market m as a function of the benchmark rate. The generic optimal subsidy problem is given by

$$\max_{\{B_m\}} \sum_{m=1}^{M} CS_m(B_m) \quad \text{s.t.} \quad \sum_{m=1}^{M} GovExp_m(B_m) = \bar{B} \tag{2}$$

The solution to this problem depends on the model of firm and consumer behavior from which the CS and $GovExp$ functions are derived. To illustrate the way in which heterogeneity in consumer preferences and firm costs can interact to generate different optimal policy schedules, we focus on a discrete choice demand framework under monopoly, a simplified version of the framework we will introduce in Section 5 to model MA.

In each market m, there exists a measure of consumers, denoted by i, and a single firm providing a single product at price p_m. Consumers make a discrete choice between the product offered by the firm and the outside option of no purchase. The utility of purchasing the good from the monopolist is given by $u_{im} = \alpha_m p_m + \beta_m + \epsilon_{im}$. In this equation, α_m is a market-specific price sensitivity, β_m is a market-specific taste for the good, and ϵ_{im} is an per-consumer idiosyncratic taste for the good. Let ϵ_{im} be drawn from the Type-I extreme value distribution, and normalize the utility of the outside good to zero. The probability that a consumer in market m purchases the good, which we will refer to as the inside share of the good, is then given by $s_m = \frac{\exp(\alpha_m p_m + \beta_m)}{1+\exp(\alpha_m p_m + \beta_m)}$. Each market hosts a different monopolist firm, which faces marginal costs c_m and chooses the price p_m knowing the consumer demand characteristics described above and the government policy B_m, which consists of a per-purchase payment. The firm's profit-maximization problem is therefore given by $\max_{p_m}(p_m - c_m + B_m)s_m(p_m)$.

If the consumer does not choose the private product, the government must provide the good at a cost of TM_m. Let $s_m(B_m)$ be the inside share of the good at the profit-maximizing price under

policy B_m. The government's expenditures are then $GovExp_m(B_m) = s_m(B_m) * B_m + (1 - s_m(B_m)) * TM_m$ and the consumer surplus is given by $CS_m(B_m) = \frac{1}{\alpha_m} \ln(1 + \exp(\alpha_m p_m^*(B_m) + \beta_m))$. This "inclusive value" formulation of consumer surplus is the amount consumers in market m would be willing to pay before knowing their idiosyncratic shock to have the choice of purchasing the good (Small and Rosen, 1981).

Table 1 provides numeric solutions to Equation 2 for several combinations of markets that mimic the MA policy environment. For each example, markets differ by parameters $\{\alpha_m, \beta_m, TM_m, c_m\}$. We apply the MA benchmark policy described in Section 2 to obtain an initial subsidy schedule B_m^{MA}. We calculate the government's budget constraint under the MA policy and then solve for the optimal subsidy schedule holding the total budget fixed. We compare the per-market benchmark rate, government expenditures, and consumer surplus under the MA approach to the optimal approach. Across examples, the second market is identical and has "average" characteristics taken across the other markets.

In Example (a) the market with the lowest TM_m has still lower c_m and price sensitivity is higher when costs are lower. The optimal policy takes advantage of the cost savings available in the lowest-cost market and redistributes expenditures to the other two markets. The total consumer surplus increases by 14.6%.

Example (b) changes the relationship between firm costs and the government's costs – while firm costs are still positively correlated with the government's costs, firm costs are less extreme. In this example, the optimal policy moves more money in percentage terms from the lowest-cost market to pay for increases in the other two. The change in firm costs has led to a difference in the rank ordering of market outcomes: while in Example (a) the lowest-cost market has the highest consumer surplus, in Example (b) the highest-cost market has the highest consumer surplus. The optimal policy changes the relative rank ordering of consumer surplus across markets and increases the total by 16%. Example (c) duplicates the cost structure of Example (b) but reverses the relationship between price sensitivity and costs. This change results in the highest changes in government expenditures in percentage terms, as well as the largest increase in total consumer surplus of 20.4%.

Taken together, these examples illustrate important features of the optimal policy problem. First, there is a disconnect between the benchmark rate and the realized government expenditures and consumer surplus. Indeed, in Example (c) the optimal policy increases the benchmark rate and government spending by the greatest percentage in the third market, yet consumer surplus improves

the most in percentage terms in the second market. As a corollary, the optimal policy depends not only on the interactions between supply and demand conditions in any particular market, but on the heterogeneity in those conditions across markets. While the second market is identical across all three examples, the optimal benchmark rate and corresponding expenditures and consumer surplus differ. Finally, all three examples illustrate significant improvements in consumer surplus over the MA policy, despite differences in the nature of the relationship between firm costs and the government's costs.

Table 1: Example solutions to simple optimal subsidy problems

α_m	β_m	TM_m	c_m	B_m			$GovExp_m$			CS_m		
				MA	Opt.	Δ%	MA	Opt.	Δ%	MA	Opt.	Δ%
(a): Firm costs are positively correlated with the government's cost												
-3	-1	10	8	11.5	11.116	-3.34	11.303	10.947	-3.15	0.677	0.630	-6.87
-2.5	0	12.5	12.5	12.5	13.105	4.84	12.500	12.763	2.11	0.098	0.229	133
-2	1	15	18	14.25	16.625	16.7	15.000	15.092	0.618	0.000	0.029	10,500
									Total:	0.775	0.888	14.6
(b): Firm costs are less positively correlated with the government's cost												
-3	-1	10	11	11.5	10.850	-5.65	10.432	10.063	-3.54	0.113	0.026	-77.4
-2.5	0	12.5	12.5	12.5	12.868	2.95	12.500	12.629	1.03	0.098	0.172	75.5
-2	1	15	13.5	14.25	14.725	3.33	14.581	14.822	1.65	0.409	0.522	27.7
									Total:	0.620	0.720	16
(c): Price sensitivity is higher when costs are higher												
-2	1	10	11	11.5	10.769	-6.36	10.750	10.226	-4.88	0.347	0.174	-49.9
-2.5	0	12.5	12.5	12.5	13.053	4.42	12.500	12.730	1.84	0.098	0.216	120
-3	-1	15	13.5	14.25	14.992	5.21	14.701	14.995	2	0.169	0.350	107
									Total:	0.614	0.740	20.4

Notes: This table illustrates the solution to the optimal subsidy problem described by Equation 2 for example sets of markets. In each market, a measure of consumer has the option to purchase a product from a profit-maximizing monopolist. α_m, β_m, TM_m, and c_m are the market-level price sensitivity, taste for the good, government's marginal cost, and the monopolist's marginal cost, respectively. B_m is the benchmark rate set by the government, and CS_m and $GovExp_m$ are the consumer surplus and government expenditures in each market. In each example, we calculate B_m using the MA policy described in Section 2 and find the government expenditures and consumer surplus. We use the total government expenditures under the MA policy as the budget constraint for Equation 2. Δ% columns indicate the percentage change in the outcome from the MA policy to the optimal policy.

4 Data

We combine administrative data on plan characteristics and enrollment from CMS with micro-level data on consumer choices from the Medicare Current Beneficiary Survey.

4.1 Medicare Advantage plans

We obtain the annual Plan Finder database from CMS for the years 2008 to 2015. 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 for each plan-year. 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.[8]

As part of the bidding and enrollment process, CMS calculates and publicizes out-of-pocket cost (OOPC) estimates for each plan. CMS creates these estimates by forming a representative bundle of services used by TM enrollees in different demographic groups. 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 drop plans sponsored by employers and plans designed for individuals who are "dual-eligible" for 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. 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 2. The first column presents unweighted means of the plan characteristics. The average plan charges a monthly premium of $42 and requires a $94 deductible. Most plans offer at least one additional category of coverage: 76% of plans offer prescription drug coverage, 60% offer dental coverage, and 90% offer vision coverage. Out-of-pocket costs vary by age and health status: a 65 year-old in excellent health is expected to spend $1,384 on out-of-pocket costs for the average plan, while an 85 year-old in poor health is expected to spend $4,099. On average, the plans we observe enroll 672 people per county-year and are offered in 15 counties. To capture differences between

[8]Plans may submit a single bid for multiple counties that have different benchmark rates. In these cases, we use CMS rules to construct a county-level bid for each plan that can be compared to that county's benchmark.

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, a higher star rating, and lower out-of-pocket costs.

Table 3 presents the mean characteristics of plans across benchmark quartiles calculated at the market level. 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 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.

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.

4.2 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. 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 responses for 2008-2015.[9] We observe detailed demographic information, including income, age, sex, race, and education. 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 their MA plan, if any.[10]

The MCBS does not cover every county in the country. Instead, it employs a multi-level

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

[10] 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.

Table 2: 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

Notes: This table reports summary statistics for the plans in county-year markets in which we observe micro-data through the Medicare Advantage Beneficiary Survey (MCBS). The first column reports raw means across all plan-county-year observations. Many plans are available across multiple counties, and so the second column collapses the data to the plan-year level and weights the resulting observations by total enrollment. All prices are in 2015 dollars. The Star rating is calculated by CMS and ranges from zero to five. We obtain 30 out-of-pocket cost estimates for each plan, though report only three here for brevity.

Table 3: Mean plan characteristics by market-level benchmark quartile

	Benchmark quartile			
Variable	**1st**	**2st**	**3rd**	**4th**
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

Notes: An observation is a county-year-plan; all reported figures are unweighted means across observations in the relevant quartile. The benchmark quartile is defined at the market (county-year) level.

clustered sampling procedure that 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). We use sampling weights to ensure our results are nationally representative.

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

There are 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 who are eligible to purchase MA plans and are therefore included in CMS' enrollment files. These include individuals with employer-provided insurance plans, those whose original Medicare eligibility was not age-related, those with ESRD, and those who are not full-year enrollees in Part A and Part B. As these individuals purchase MA plans, we cannot exclude them without violating our model of the data-generating process. We instead include them and create "administrative status" identifiers to capture their behavior. As the presence of these individuals in our analysis changes the distribution of consumer surplus, we report some results with reference to the "standard analysis" set, which excludes these groups, for comparisons with other Medicare analyses.

Those who do not enroll in MA have access to other options, and variation in the price of those options may make MA plans more or less attractive. We focus on Medicare supplemental insurance (a.k.a. Medigap) which pays for care not covered by TM. For example, TM covers 80% of the cost of physician visits, and a Medigap plan may pay for the rest. Medigap plans are standardized by CMS and indexed by letters: all "Plan A" policies have the same structure across insurers. For each person, we obtain the rate for Medigap Plan C offered by United Healthcare 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.[11]

Summary statistics on our 58,444 individual-year observations covering 2,947 county-year markets and 32,993 unique individuals are reported in Table 4. The first two columns report means and standard deviations across all observations. 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 by CMS as White, with 7.9% Black and 1.0% Hispanic. Over 75% self-report "Good" or better health. 24%

[11] 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.

report having college degrees and 18% did not graduate high school. 33% receive some insurance from a current or previous employer, and 16% are Medicare-eligible for reasons other than turning 65. The second set of columns splits the data by MA enrollment. On average, MA enrollees have lower income, are less likely to be White, and have lower educational attainment. Individuals with employer-provided insurance or with ESRD are less likely to be enrolled in MA.

The third and final set of columns of Table 4 illustrates the panel nature of our data and focuses on panel observations for which the individual was enrolled in TM in the previous year – 23,278 observations total. We split the data into those who switched from TM to MA (4.3% of observations), and those who remained on TM. Those who switched are generally similar to the larger group of MA enrollees. Some differences are seen in health status – switchers from TM to MA are slightly healthier on average than the broader MA population.

Table 4: 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

Notes: An observation is a person-year. Statistics reported here are weighted according to sampling weights provided by the Medicare Current Beneficiary Survey (MCBS). Income and prices are in 2015 dollars. The outside good price is the United Healthcare premium for Medigap Plan C (see text for details). Demographic categories are defined by CMS administrative data. The first set of two columns reports means and standard deviations for all observations in the microdata. The third and fourth columns split the observations into those enrolled in MA and those enrolled in TM. The last two columns split the observations by switching behavior. Only those person-year observations for which we observed the individual enrolled in TM in last year's MCBS are included in the these last columns.

5 A Model of Medicare Advantage

Our model of MA builds upon the example introduced in Section 3 and leverages the rich individual-level information in the MCBS and the detailed plan data from CMS. The demand side is inspired by Goolsbee and Petrin (2004) and captures beneficiary heterogeneity along multiple dimensions. In this sense, our demand model is similar to that of Aizawa and Kim (2018) although we estimate parameters with a different approach. On the supply side, we model multi-product insurers as profit-maximizers that choose prices and product characteristics. We take the number of firms and products as given. As we do not observe individual-level risk scores, we model marginal costs in terms of an average risk individual.[12]

Each period t begins with the government choosing benchmarks B_m for each market m (defined as a county) and Medigap insurers determining a price schedule $\{p_0\}_m$. B_m maps to plan-specific subsidies B_j via Equation 1. ϕ and λ are taken to be exogenous. 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.

5.1 Demand

Consumers have demographic characteristics z_i which include income bracket g and age-and-health-status type h. 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 period enrolled in plan k. We define three indicators S_{sij} to capture the effects of this previous enrollment. Let S_{1ij} be equal to one if k is the outside good – we call this the *Medicare-to-MA* indicator. Let S_{2ij} – the *MA Interfirm* indicator – be one if k is offered by a different firm than j. Finally, let S_{3ij} – the *MA Intrafirm* indicator – be one if 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

[12] 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).

market and time subscripts, the choice specific utility for MA plans is given by:

$$u_{ij} = \left(\alpha + \sum_g \alpha_g g_i\right) p_j + \beta_h \sum_h oopc_{hj} h_i + \sum_s \beta_s S_{sij} + \sum_s \sum_h \beta_{sh} S_{sij} h_i + \beta_z z_i + \beta X_j + \xi_j + \epsilon_{ij} \tag{3}$$

In this equation, α and α_g together capture heterogeneity in 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. β_{sh} captures switching costs and varies by health group. β_z captures heterogeneous tastes for MA versus TM. β captures mean tastes for plan characteristics X_j.

As is standard in random utility demand systems, we decompose the unobservable (to the econometrician) portion of utility into two components. ξ_j represents the portion of unobserved plan utility that is common across individuals. ϵ_{ij} represents the idiosyncratic taste of consumer i for plan j which is assumed to be drawn independently from the Type-I extreme value distribution.[13]

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}, through

$$u_{i0} = \beta_0 p_{i0} + \epsilon_{i0}. \tag{4}$$

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

We include switching costs due to the evidence of Table 4 and the consistent finding of inertia in plan enrollment (Nosal, 2012, Aizawa and Kim, 2018).[14] Enrollees in TM face a different set of benefits and provider networks than in MA and those details vary across insurers and plans. Switching between plans entails learning about administrative procedures and provider networks. In addition, Medicare beneficiaries are automatically re-enrolled in their previous plan if they take no action during their open enrollment period – it is virtually costless to re-enroll. Similar to Handel (2013), we model these costs with a direct utility impact through the $\sum_s \beta_s S_{sij}$ term.

[13] We have also explored demand systems with persistent idiosyncratic tastes and found similar results.

[14] Like Aizawa and Kim (2018), we do not model consumers as dynamic for several reasons. First, such analysis is computationally intensive. Second, it likely requires assuming that individuals choose according to a model of neoclassical preferences with a discount factor close to one. However, recent work has shown that in related settings that model does not explain Medicare beneficiary behavior well (e.g. Dalton et al., 2018). Third, (Nosal, 2012) estimates such a model and finds extremely high (perhaps implausibly so) switching costs. Fourth, our estimation approach captures the inertia that is salient for our counterfactual analysis.

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

$$\delta_j = \alpha_0 p_j + \beta X_j + \xi_j \tag{5}$$

and an individual-specific deviation from that mean

$$\mu'_{ij} = \sum_g \alpha_g g_i p_j + \sum_h \beta_h oopc_{hj} h_i + \sum_s \beta_s S_{sij} + \sum_s \sum_h \beta_{sh} S_{sij} h_i + \beta_z z_i - \beta_0 p_{i0} + \epsilon_{ij}. \tag{6}$$

Let $\mu_{ij} = \mu'_{ij} - \epsilon_{ij}$. Given our distributional assumption on ϵ_{ij}, the probability that consumer i chooses plan j (i.e. the share function) is a logit form

$$s_{ij} \equiv \Pr(i \text{chooses} j) = \frac{\exp(\delta_j + \mu_{ij})}{1 + \sum_{k \in \mathcal{J}_m} \exp(\delta_k + \mu_{ik})}. \tag{7}$$

The consumer surplus for an individual i given prices and product characteristics is

$$CS_i = E[\max_j u_{ij}]/\alpha_i = \frac{1}{\alpha_i} \ln \left(\sum_j \exp(\delta_j + \mu_{ij}) \right) \tag{8}$$

where α_i is the price sensitivity of i. As with the example in Section 3, this is the amount i would be willing to pay to have the choice of MA plans, relative to a world in which TM was the only option – it does not include the benefit of TM itself.

5.2 Supply

Insurers start each period with common knowledge of the distribution of consumer demographics including enrollment status and each others costs including cost shocks. Firms also have common knowledge of the per-enrollee subsidy they receive B_j.[15]

We model marginal costs as log-linear in observed characteristics r_j (which includes X_j and $oopc_j$ and possibly other observables) and some unobserved component ω_j with

$$\ln(mc_j) = \gamma_r r_j + \omega_j. \tag{9}$$

[15] We do not explicitly model bidding though we account for it in our counterfactual approach. See Section 7.

Given these costs and the demand model above, the profit maximization problem for firm f is

$$\max_{\{p_f, X_f, oopc_f, \xi_f\}} \pi_f = \sum_{j \in \mathcal{J}_f} \left[(p_j + B_j - mc_j) \int_i s_{ij}(p, X, oopc, \xi; \theta) di \right]. \tag{10}$$

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.[16] s_{ij} depends on the characteristics of all products in the market; θ is the vector of demand system parameters.

6 Estimation

We estimate the parameters of the demand system following the approach of Goolsbee and Petrin (2004). First, we estimate parameters which capture individual-level variation – 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 μ'_{ij}. For a given candidate value $\tilde{\theta}$ 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.[17] We form the likelihood function

$$L_i(C_{ij}; \tilde{\theta}, \delta(\tilde{\theta})) = \Pi_j s_{ij}^{C_{ij}}, \tag{11}$$

where s_{ij} is given by Equation 7. In the first stage of our estimation procedure, we apply the MCBS sample weights w_i and maximize the weighted log likelihood function

$$l(C; \tilde{\theta}) = \sum_i ln(L_i) w_i, \tag{12}$$

At the point estimate $\hat{\theta}$ we store the unique $\hat{\delta}_j$ recovered by the Berry contraction. We then regress $\hat{\delta}_j$ on observable product characteristics according to Equation 5. We recover marginal cost parameters by inverting the firms' first-order conditions and estimating Equation 9.

[16] In practice, firms do not directly choose each of the OOPC measures as they are calculated by CMS. The plan characteristics we observe capture 38% of the variation in the OOPC measures and therefore it is practical to model these as separate choices made by firms rather than as functions of other observables.

[17] Calculating C_{ij} involves two complications. First, the 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 implies that a fraction of our observations have no past enrollment data. We address this by eliminating plans given what we do observe and taking draws from the conditional shares of the remaining products.

6.1 Identification

Since ξ_j is observed by firms but not by us, it is likely to be correlated with the price and other characteristics of the plan. To identify the coefficients 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).

One potential concern is the fact that firms choose ξ_j and X_j simultaneously as a function of the benchmark, which implies that these characteristics may be correlated through this mutual dependency. Our data captures a wealth of detail about plan benefits, and thus the largest components of ξ are likely to be network breadth, advertising, and unobservable quality. These are unlikely to change in response to small changes in the benchmark, and thus it is reasonable to assume that ξ_j is uncorrelated with B_j which we can test by estimating $\hat{\xi}_j$ and computing its empirical correlation with B_j. Some of the components of ξ_j are likely to be constant across plans within a given firm[18], and so we include firm fixed effects to reduce the extent to which ξ_j might be correlated with X_j.

Firms face state-level regulations and federal-level coverage requirements for MA plans which change over time which suggests that there may be some component of marginal cost which is constant across firms within a state-year. Following the logic of Hausman et al. (1994) we use average prices in other counties as additional price instruments. Since plans are often offered at the same price in geographically contiguous counties, we calculate this instrument using "non-contiguous counties" – that is, for a plan in a county, we calculate the average price of plans in counties of the same state which do not share borders with the county under consideration.

Just as ξ_j 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, after estimating the demand parameters, we can calculate a $\hat{\xi}_j$ for each plan. We estimate Equation 9 by assuming that any remaining unobservable components of cost are uncorrelated with ξ_j and our observables.

6.2 Demand results

Maximum likelihood parameter estimates are reported in Table 5. High and medium income consumers are less price sensitive than low income consumers. Out-of-pocket costs create disutility, while an increase in the cost of Medigap drives consumers to MA plans. Non-White individuals

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

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 bottom three panels of Table 5 report switching cost parameters. We interact each switching cost 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. Inter-firm switches are less costly and intra-firm switches are cheaper still. These results suggest the primary component of switching costs is the disutility of changing 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.

Table 6 reports estimates of Equation 5. The first column presents the OLS estimates and the second presents IV results. Consistent with OLS estimates on price being biased towards zero, the IV price coefficient is larger in magnitude than the OLS coefficient. For this reason, we concentrate our attention on the IV specification.[19] These estimates in general correspond to sensible priors. For the plans with a positive premium the average plan elasticity is -2.66.[20] The average semi-elasticity of increasing premiums $1 is .3 percent, 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 OOPC estimates in the first stage, which are in part a function of these variables, changes the interpretation of these parameters. The parameters estimated here are not the direct effect of these features *per se*, but rather their impact on utility holding expected OOPC 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 are preferred to HMOs.

Table 7 reports the mean switching cost by type and income group calculated by dividing the switching cost coefficients by the price coefficient for each person. The cost incurred by an average medium-income individual switching from TM to MA is $918, which is almost twice the mean

[19] We do not include star ratings 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, increasing the noise in the measure.

[20] An important issue in the proposed Aetna-Humana merger was the rate of substitution between MA and TM if insurers 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 5: Maximum likelihood estimates of individual-specific preferences

Variable	Coefficient	Std. Err.
Income-price interaction (per $1000)		
Medium income	0.153	0.039
High income	0.089	0.042
Estimated OOPC (per $1000)	-0.038	0.024
Medigap price (per $1000)	0.218	0.046
MA × Demographics		
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
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-Insurer 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-Insurer 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	

Notes: This table reports maximum likelihood estimates of the individual-specific components of Equation 3. In the estimation, each individual receives the estimated out-of-pocket cost (OOPC) for their age-health demographic group (e.g. 70-74 Good). Income groups are defined by terciles of the income distribution for the entire MCBS sample; low income is the omitted group. The omitted group for the switching cost interactions is 'Poor' health. We weight the likelihood function using the MCBS sample weights to obtain nationally representative estimates.

Table 6: Estimates of mean preferences for plan characteristics

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

Notes: This table reports estimates of the plan-specific components of Equation 3. To form these estimates, we first use maximum likelihood estimation to recover the parameters in Table 5, calculate plan fixed effects δ_j, and then regress these fixed effects on plan observables via Equation 5. Observations are thus at the market-plan level. The IV specification uses the summation instruments of Berry et al. (1995) and our 'non-contiguous county' variation of Hausman instruments. Robust standard errors are in parentheses.

Table 7: Switching costs in dollars by switch type and income group

	Income group		
	Low	Medium	High
Switch type			
Medicare-to-MA	$877	918	894
Between-Insurers	510	534	521
Within-Insurer	167	177	176

Notes: We transform our switching cost parameters into dollars by dividing them by the coefficient on price. To form the statistics reported here, we calculate the effective switching cost for each person based on their health status and income group, and then take the mean by switch type and income group. Observations are weighted by MCBS sampling weights.

annual premium in our data, while the same average individual switching between plans within an MA insurer incurs a cost of only $177, about a third of the average annual premium.

Table 8 reports estimates of Equation 9. As the OOPC measures for different groups are highly collinear we focus on the median OOPC, the measure for a 65-69 year old in Good health – the results are similar when using alternative OOPC measures. We include firm and county fixed effects to focus on the costs of different plan designs and to account for costs which stem from geography. The major cost sharing plan design parameters, including the OOPC, deductibles, and out-of-pocket limits, enter with the correct sign. Offering prescription drugs increases costs by 3.9%. The coefficient on the demand unobservable is positive and significant at the 1% level, suggesting that specifications which did not 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 OOPC parameter. HMOs face lower costs than PPOs.

There has been some discussion about the usefulness of structural techniques for estimating costs and counterfactual outcomes (e.g. Angrist and Pischke, 2010, Nevo and Whinston, 2010). Criticism has focused on the use of demand elasticities and first-order conditions to calculate marginal costs. Some have explored this question by comparing post-merger price changes to pre-merger analyses of elasticities and costs (Peters, 2006, Weinberg, 2011, Bjornerstedt and Verboven, 2016). In our setting, CMS reports the actual risk-adjusted per-capita TM expenditures in each market, which, if our estimation approach is consistent, are likely to be correlated with our estimated marginal costs. We mimic the spirit of an exercise in Curto et al. (2015, fig. 6) and compare the share-weighted estimated MA cost to TM costs at the county level for 2015 in Appendix Figure B.3. The two cost measures are positively correlated with a coefficient of .527; on average, estimated MA costs are

lower than FFS costs. We return to this last finding in our counterfactual analysis.

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
R^2	.749

Notes: This table reports the results of estimating Equation 9, where the left-hand side variable comes from inverting the first-order conditions of Equation 10 at the demand parameters presented in Tables 5 and 6. Estimates are formed via OLS. Robust standard errors are in parentheses. OOP stands for "out-of-pocket." Observations are at the plan-market level.

7 Optimal Geographic Variation in Medicare Advantage Subsidies

We now turn to the problem of setting the benchmark rates to maximize the consumer surplus generated by MA keeping government expenditures constant. To update the notation of Equation 2 to our full model, 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 in market m including both TM and MA. $\bar{B}$ is the total budget available to the government. We consider three types of

optimal subsidy problems given by

$$\max_{\{B_m\}} \sum_m \int_i CS_{im}(B_m)di \quad \text{s.t.} \sum_m GovExp_m(B_m) = \bar{B} \tag{13}$$

$$\max_{\{B_m\}} \alpha \sum_m \int_i CS_{im}(B_m)di - (1-\alpha)Var(CS) \quad \text{s.t.} \sum_m GovExp_m(B_m) = \bar{B} \tag{14}$$

$$\max_{\{B_m\}} \sum_m \int_i CS_{im}(B_m)di \text{ s.t. } B_m > \bar{B}_m \forall m \text{ and } \sum_m GovExp_m(B_m) = \bar{B} \tag{15}$$

The first equation maximizes consumer surplus subject to a government budget constraint and is analogous to Equation 2. As we document below, the solution increases the variance in consumer surplus, which may not be politically feasible to implement. For this reason, Equation 14 penalizes the variance in consumer surplus across individuals – $\alpha \in [0,1]$ determines the relative weight on consumer surplus versus the variance in consumer surplus. Currently, the CMS benchmark formula includes floors suggesting that may be the preferred policy solution in areas where the FFS spending formula results in low benchmarks. Equation 15 therefore includes a floor constraint.

Solving these problems requires calculating the CS and $GovExp$ functions. Given prices and product characteristics as a function of B_m, $CS(B_m)$ comes from Equation 8. Expenditures are

$$GovExp(B_m) = \sum_j \int_i (s_{ij} Payment_{jm} + (1 - s_{ij}) TM_m)\, di, \tag{16}$$

where $Payment_{jm}$ is the plan-level payment of Equation 1, s_{ij} is the share function of Equation 7, and TM_m is per-enrollee TM spending in the market. As we do not observe individual-specific risk scores, we calculate spending using the average risk level in the market — in other words, we set $RiskAdjustment_{it} = RiskAdjustment_{mt}$ in Equation 1 for all i. Since the benchmark and rebate adjustment terms, ϕ_{jt} and λ_{jt}, are determined by previous performance, we take them as given.

7.1 An equilibrium approximation approach to counterfactual analysis

The inputs to CS and $GovExp$ are prices and product characteristics in each market as a function of the benchmark. A traditional counterfactual approach would derive policy functions by varying the benchmark and searching for Nash equilibria, usually through analyzing firms' first-order conditions for profit maximization. For example, Fan (2013) uses this approach to find a post-merger equilibrium in prices and five product characteristics in a daily newspaper market with five com-

petitors – 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 a single equilibrium in one market involves a fixed point search across 658 first-order conditions requiring significantly greater computational effort (Daskalakis et al., 2009). To find an optimal policy, we must search over hundreds of equilibria in each of the 445 markets in the 2015 data, a computationally infeasible task.

Our alternative "equilibrium approximation" approach combines policy function estimation with a reduced-dimension fixed-point problem and begins with the observation that we observe variation in the benchmark rate, the key primitive we wish to vary. Table 3 shows this variation comes with variation in observed product characteristics – the data trace out policy functions. To predict counterfactual characteristics, we regress observed characteristics on the benchmark (and other covariates) and use the fitted values at counterfactual benchmarks to update product characteristics. We then solve for prices in a Bertrand Nash game.

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 concerning when markets are geographically isolated (Noel, 2007), our data is characterized by firms offering products in multiple overlapping geographies. 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 pay television systems and use estimated functions for prices and product characteristics to calculate the welfare improvement caused by the introduction of satellite TV. Benkard et al. (2018) estimate strategic entry and exit behavior in the airline industry and 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 simplify the problem by defining a 'price-adjusted quality' δ' for each plan given by

$$\delta'_j = \delta_j - \alpha_0 P_j. \tag{17}$$

δ'_j includes the effects of both observables X_j and unobservables ξ_j and is easily calculated. We regress δ'_j on the benchmark, the 'rank' of δ'_j with respect to the firm's other plans, an indicator which is one if the insurer has the highest share in the market, the total number of plans in the

market, and insurer fixed effects. We repeat this exercise with each of the OOPC terms.

Given a benchmark for a market, we use our estimated policy functions to calculate the new δ and updated OOPC for each plan in that market. We take the difference between the estimated policy at the new benchmark and the estimated policy at the observed benchmark and apply that difference to the observed policy. This adds flexibility to our approximation and maintains plan-level heterogeneity without adding terms to our regressions. We apply a similar approach to other primitives by estimating marginal cost as a function of δ' and OOPC, consistent with our model, and estimating the bid as a function of marginal cost, consistent with CMS bidding rules.

With the set of product characteristics, costs, and bids in hand, we calculate 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.

These rules – using the differences in the policy functions and the non-negativity restriction on prices – make our counterfactual and estimation approaches consistent in the sense that when we calculate equilibria at the benchmarks in the data, the outcomes match the data precisely.

To summarize, given a benchmark for a given county, we calculate CS_m and $GovExp_m$ by:

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 8 and government spending with Equation 16.

We now have the ingredients to the maximization problems represented by Equations 13, 14, and 15. As we find equilibrium outcomes numerically, we cannot solve for the constraint surface analytically and instead we satisfy the constraint numerically using a penalty function. Across these problems, both the objective function and the constraint are non-linear – the constraint is also non-monotonic. We address the possibility of multiple local maxima with 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—each market's outcome depends only on that market's benchmark—and calculate outcomes for each market over a grid of benchmarks. We

use spline approximations to find candidate solutions and then refine the most promising using the numerical equilibrium solver.

We do not model the impact of the benchmark on plan entry and exit due to the computational burden required. However, incorporating entry and exit into the analysis should not significantly impact the results as the plans that enter and exit in our sample comprise a small share of the total MA market and thus contribute 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 Medicare Part D. Maruyama (2011) studies entry and pricing behavior under counterfactual subsidy policies in Medicare+Choice but does not endogenize product characteristics.

8 Counterfactual results

We apply our approach to solve optimal subsidy problems for our 2015 sample. We present our estimated policy functions, explore the local behavior of total surplus near the 2015 policy, and then illustrate the global behavior of CS_m and $GovExp_m$. We then present our primary exercise in which we maximize average consumer surplus and discuss alternative objective functions.

8.1 Policy function estimates

Table 9 reports estimates of our policy functions. The first three columns focus on the δ' policy function. Column (1) fits δ' on the benchmark alone, and Column (2) adds insurer fixed effects. Column (3), our preferred specification, controls for the rank of the plan, an firm-level indicator for high market share, and the number of plans in the market. The benchmark enters with a positive sign, and the R^2 is .514.[21] Columns (4)-(6) report estimates for three OOPC policy functions using an analogous specification; Appendix Table B.1 reports estimates for all thirty OOPC. In general, increases in the benchmark lead to decreases in estimated OOPC, and these decreases are larger for those who are older and those in worse health. R^2s range from .210 to .312.

[21]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: Estimates of the price-adjusted quality and selected out-of-pocket cost policy functions

	(1)	(2)	(3)	(4)	(5)	(6)
				OOPC	OOPC	OOPC
Variable	δ'	δ'	δ'	65-69 V. Good	70-74 Good	75-79 Fair
Benchmark	-.541	-.248	.030	-.017	-.059	-.091
	(.014)	(.015)	(.013)	(.003)	(.005)	(.007)
Plan rank			-.982	.098	.148	.236
			(.013)	(.002)	(.003)	(.005)
Leading insurer ind.			1.30	-.083	-.126	-.179
			(.031)	(.008)	(.011)	(.017)
# plans in market			-.031	-.008	-.010	-.013
			(.001)	(.000)	(.000)	(.000)
Fixed Effects	None	Firm	Firm	Firm	Firm	Firm
Observations	50,593	50,593	50,593	50,593	50,593	50,593
R^2	.034	.307	.514	.256	.240	.228

Notes: The independent variables are the δ' defined by Equation 17 and the CMS OOPC estimates by age and health status. The benchmark is at the market level. The plan rank variable is the rank of the plan within its insurer, ordered by the value of the δ', with 1 as the top-ranked plan. The leading insurer indicator is equal to 1 if the firm has the highest total share in the market. Estimates are obtained via OLS. Robust standard errors are in parentheses.

To complete our set of plan primitives, we regress marginal cost on product characteristics and the plan bid on the marginal cost. For both, we include insurer and county fixed effects. For consistency with our other estimates, the left-hand side variables are in levels.[22] The first column of Table 10 reports the parameter estimates for marginal cost. As with Table 8, we include only one OOPC for simplicity. Increasing the δ' leads to an increase in the marginal cost, and increasing the OOPC decreases the marginal cost. The second column reports estimates for the bid. Bids are slightly lower than marginal costs. Given CMS bidding rules – in particular the requirement that bids reflect the cost of providing TM-equivalent benefits – this result is sensible.

8.2 Local behavior around the 2015 policy

The construction of our optimal subsidy problem implies that the optimal policy re-allocates subsidy dollars from places where the marginal impact is low to places where the marginal impact is high. These opportunities are represented by differences in the derivatives of total surplus (here defined as consumer surplus minus government expenditures) at the 2015 policy. Understanding these derivatives can explain both how additional spending is passed on to consumers and why some

[22] We repeated the counterfactual exercise with these functions in logs instead of levels and obtained similar results.

Table 10: Estimates of the marginal cost and bid functions

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

Notes: Each column of this table reports the results of an OLS regression. The independent variable in the first column is the marginal cost estimate obtained through inverting the first-order conditions of Equation 10. The independent variable in the second column is the plan bid for an individual of 1.0 risk from the CMS data. δ' is the price-adjusted δ of Equation 17. Robust standard errors are in parentheses.

Table 11: Distribution of surplus derivatives across markets at 2015 benchmarks

Derivative w. r. t. benchmark	Mean	25th %-ile	Median	75th %-ile
Total surplus	-.015	-.116	.004	.084
Consumer surplus	.099	.043	.080	.135
Government expenditures	.114	-.013	.096	.203
TM spending	-1.46	-2.02	-1.35	-.833
MA spending	1.58	.914	1.48	2.11

Notes: This table presents the distribution of market-level derivatives of the per-capita total surplus function with respect to the benchmark, where total surplus is defined as per-capita consumer surplus minus per-capita government expenditures. The units here are dollars: a $1 increase in the benchmark rate increase the per-capita consumer surplus in the average market by 9.9 cents. Derivatives are calculated numerically for each of the 445 markets in our 2015 data by simulating a $1 increase in the benchmark rate for each market and comparing the outcomes to those at the 2015 policy. Markets are weighted evenly.

counties are more desirable targets for subsidies than others.

Table 11 reports the mean, median, and 25th and 75th percentiles of derivatives of per-capita surplus functions with respect to the benchmark at the 2015 policy. On average a $1 increase in the benchmark leads to a $.11 increase in expenditures and a $.10 increase in consumer surplus. While the mean expenditure derivative is positive, it is negative for 126 out of the 445 markets in our sample. The last two rows decompose expenditures into TM and MA spending; a $1 increase to the benchmark rate increases MA spending by $1.58 but decreases TM spending by $1.46 as individuals switch from TM to MA.

To explore *how* changes in the benchmark lead to these changes, we decompose the constituent

Table 12: Decomposing the components of government expenditures

	Mean	25th %-ile	Median	75th %-ile
Average payment per MA enrollee ($)	9,069	8,713	9,022	9,322
Total MA share	.258	.138	.229	.357
Expenditures per TM enrollee ($)	9,782	9,131	9,591	10,417
∂ MA share / ∂B ($\times$1000)	.153	.087	.147	.214
∂ Avg. payment per MA enrollee / ∂B	.799	.701	.783	.868
∂ Share-weighted bid / ∂B	.417	.154	.350	.567

Notes: The top panel reports the levels of variables used to calculate government expenditures; the middle panel reports relevant derivatives. The bottom panel illustrates the linearity of our bid policy function and the way in which substitution between plans increases the share-weighted bid. Statistics are reported at the market level; all markets are evenly weighted. Derivatives are calculated numerically by simulating a $1 increase in the benchmark rate for each market and comparing the outcomes to those at the 2015 policy.

components of the surplus functions. Table 12 reports distributions of components of government expenditures. The top panel reports the levels of the average payment per MA enrollee, the total MA share, and the expenditures per TM enrollee. The middle panel reports derivatives that can be used to reconstruct the derivatives reported in Table 11: $\frac{\partial MA}{\partial B} \approx \frac{\partial S}{\partial B} \times Payment + \frac{\partial Payment}{\partial B} * S$ and $\frac{\partial TM}{\partial B} \approx -\frac{\partial S}{\partial B} \times TMExp$. The bottom panel shows changes in plan bids.[23] The rebate payment rules imply that average payments increase by more than the share-weighted bid. As a result, increasing the benchmark by $1 leads to an increase in the average share-weighted bid of $.42 and an increase in the average MA payment of $.80.

These results point to a mechanism for surplus gains. While increasing the benchmark increases the payment per MA enrollee everywhere, in some counties the intensive margin payment increase (spending more on current MA enrollees) is outweighed by the extensive margin share increase (moving beneficiaries from TM to MA) and the total spending decreases. The extent to which this mechanism can generate gains depends upon non-local behavior; we return to this when we discuss the global behavior of the surplus functions.

Consumer surplus is a non-linear function of demographics and the prices and characteristics of all the plans. As a consequence, a mechanical decomposition of consumer surplus is not straightforward. To explore the mechanisms through which consumer surplus is affected by changes in the benchmark, we turn off components of the policy function and re-calculate surplus.

The first panel of Table 13 reports the share-weighted δ', OOPC, firm costs, and prices, as well

[23]The change in mean bids is constant across counties thanks to our linear policy function specification. However, substitution between plans means that the change in the share-weighted mean bid is higher than the change in the unweighted mean.

as the derivative of the total MA share and of consumer surplus for the full model at the 2015 policy. The second panel turns off changes to δ' and OOPC. An increase in the benchmark still leads firms to lower prices, encouraging consumers to switch from TM to MA and switch from lower quality plans to better plans. Since the δ' and OOPC are not changing at the plan level, the changes seen in this panel come from substitution and new enrollment. The resulting change in consumer surplus is similar to the change seen in the full model. The third and fourth panel, which turn on changes in δ' and OOPC respectively, obtain similar results. We conclude that lower prices leading to new enrollments and substitution is the key mechanism that creates gains for consumers.

These decompositions do not explain *why* the derivatives of the surplus functions with respect to the benchmark vary across counties. Insofar that the optimal policy reallocates money to places where these derivatives are large, answering this question can help identify the underlying features of the data which lead to the optimal benchmarks we calculate below.

Our model captures heterogeneity across six factors: market-level average risk, firm and government marginal costs, product characteristics, consumer demographics, switching costs, and the number of products. Each of these factors may contribute to the total surplus derivative in different ways for different markets – one market may have a large derivative generated by very low firm costs, and another may have a similar derivative generated by very favorable demographics. We characterize the relative importance of these factors by 'turning off' heterogeneity in the data.

For the first five factors, we remove heterogeneity by constructing a weighted average and setting the factor equal to the average everywhere.[24] We calculate the absolute value of the difference between the derivative of the surplus function taken with the full dataset to the derivative with the altered dataset. We define the mean of this number across markets as the "mean absolute difference" (MAD). The MAD is a single-number measure of the extent to which the derivatives are driven by the different features of the data – a higher MAD means that the heterogeneity we see in the data along that dimension leads to greater differences in the derivatives. With a slight abuse of notation, if $\frac{\partial TS_m^f}{\partial B_m^f}$ is the derivative of the total surplus for market m using data with full heterogeneity, and $\frac{\partial TS_m^a}{\partial B_m^a}$ is the derivative using altered data, we calculate

$$MAD_a = \frac{1}{M} \sum_m \left| \frac{\partial TS_m^f}{\partial B_m^f} - \frac{\partial TS_m^a}{\partial B_m^a} \right|. \tag{18}$$

[24]For example, to remove heterogeneity in firm and government marginal costs, we calculate the average plan costs weighted by the number of enrollees, and the average government FFS costs weighted by the number of Medicare beneficiaries, and then set the marginal costs for all plans and the FFS costs in all markets equal to those averages.

Table 13: Mechanisms behind the pass-through of benchmark increases to consumer surplus

	Mean	25th %-ile	Median	75th %-ile
Full model				
$\partial\delta'/\partial B$ ($\times$1000)	.730	.286	.606	1.02
∂ OOP / ∂B	-.153	-.243	-.144	-.047
∂ Cost / ∂B	.379	.179	.334	.53
∂ Price / ∂B	-.155	-.393	-.086	.052
∂ MA share / ∂B ($\times$1000)	.153	.087	.147	.214
∂ CS / ∂B	.099	.043	.08	.135
Prices only				
$\partial\delta'/\partial B$ ($\times$1000)	.710	.258	.584	1.00
∂ OOP / ∂B	-.092	-.182	-.082	.016
∂ Cost / ∂B	.372	.170	.326	.525
∂ Price / ∂B	-.157	-.398	-.087	.052
∂ MA share / ∂B ($\times$1000)	.151	.086	.144	.212
∂ CS / ∂B	.098	.041	.079	.133
δ' and prices				
$\partial\delta'/\partial B$ ($\times$1000)	.735	.287	.610	1.03
∂ OOP / ∂B	-.091	-.181	-.082	.015
∂ Cost / ∂B	.374	.174	.329	.527
∂ Price / ∂B	-.156	-.396	-.087	.052
∂ MA share / ∂B ($\times$1000)	.153	.088	.148	.215
∂ CS / ∂B	.099	.043	.081	.135
OOP and prices				
$\partial\delta'/\partial B$ ($\times$1000)	.705	.256	.580	.996
∂ OOP / ∂B	-.154	-.244	-.144	-.047
∂ Cost / ∂B	.376	.175	.331	.529
∂ Price / ∂B	-.156	-.395	-.086	.052
∂ MA share / ∂B ($\times$1000)	.150	.086	.143	.211
∂ CS / ∂B	.097	.041	.079	.133

Notes: This table reports the derivatives of share-weighted product characteristics, costs, and prices, as well as the derivative of the total MA share and of consumer surplus, all at the 2015 policy. The top panel reports these derivatives for the full model. The second panel reports results for a model in which product characteristics are fixed at their 2015 levels and prices are allowed to vary with the increase in the benchmark. The third and fourth panels turn on changes to price-adjusted-δ and out-of-pocket costs, respectively. All derivatives are calculated numerically for each of the 445 markets in our 2015 data by simulating a $1 increase in the benchmark rate for each market and comparing the outcomes to those at the 2015 policy.

Table 14 reports this calculation for the components listed above. For each derivative, we report the mean, the interquartile range, and the MAD. The first column presents the full model, and matches the statistics reported in Table 11. Columns (2)-(6) remove one component of heterogeneity at a time. A consistent pattern emerges: firm factors (costs and product characteristics) have higher MADs than consumer factors (demographics and switching costs), which are higher than the MADs generated by heterogeneity in risk.[25] On the firm side, costs matter more than product characteristics; on the consumer side, demographics matter more than switching costs.

The other source of heterogeneity across markets is in the number of products available. The firm first-order conditions imply that differences in the number of products available, all other things equal, will lead to differences in prices and therefore surplus. However, removing heterogeneity in this dimension poses challenges. First, the number of products is a discrete measure – the other factors are continuous. Second, setting the number of products to a constant everywhere entails removing products in some markets and adding products in others, which in turn requires choosing which products to remove, choosing the new products, and assigning switching costs. However, if we start from a dataset that does not have any heterogeneity in these factors, these questions become moot as the products are identical and the assignment of consumers does not matter. Column (7) presents results from a dataset with heterogeneity in the marginal cost, product characteristics, and switching costs removed, and Column (8) additionally sets the number of products equal to the population-weighted average. The MADs presented in Column (8) are calculated with respect to Column (7) and are in general higher than the MADs for the demographic and switching cost factors, but are lower than the cost and characteristic factors.[26]

These factors are not orthogonal. For example, turning off heterogeneity in risk increases the interquartile range of the total surplus derivative which implies that the heterogeneity in risk is negatively correlated with heterogeneity along other dimensions. Turning off the heterogeneity in marginal costs leads to a reduction in the interquartile range, implying that the heterogeneity in that dimension is positively correlated with other sources of heterogeneity. Column (7) illustrates this correlation through the MAD – though the sum of the MADs coming from turning off differences in marginal cost, product characteristics, and switching costs one at a time is .342, the MAD from turning off all three of these features simultaneously is only .178.

To summarize, our exploration of the behavior of firms and consumers near the 2015 policy

[25] Risk does not enter the consumer surplus equation, and therefore the MAD is zero.

[26] This is in a sense reassuring – if we found that this factor was significantly more important than any other, our decision to abstract from entry and exit would be more concerning.

indicate that prices are the key mechanism through which changes in the benchmark rate are passed on to consumers. Furthermore, heterogeneity in firm costs is the most important factor driving differences in the marginal effect of additional government spending across markets. As a consequence, firms costs provide the largest redistribution incentives.

8.3 The non-local behavior of surplus and expenditures

The local behaviors of the surplus and expenditure functions at the 2015 policy point to possible avenues for consumer surplus gains but do not provide sufficient information to calculate the optimal policy. Though the policy function approximations we use are linear, the pricing and share functions are not – indeed, the logit share function has both concave and convex portions – and so we should expect the CS and $GovExp$ functions to be non-linear as well.

Figure 2 illustrates components of the CS_m function for Travis County, TX (containing Austin), Cook County, IL (containing Chicago), and Worth County, GA (a rural county near Albany). We chose these counties due to their different sizes and the typical nature of their counterfactual equilibria. The first graph illustrates the share-weighted δ'_j as a function of the benchmark. 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 depicts the share-weighted average OOPC for the 65-69 Good category and shows non-monotonicity as the benchmark increases, illustrating the relative weight of the out-of-pocket costs in the utility function. The third graph shows the share-weighted average price. The price increases slightly in Cook County for some benchmark increases due to substitution. As the benchmark increases, prices hit the zero lower bound.

Figure 3 illustrates components of $GovExp_m$ for the same counties. The left-hand graph depicts the total inside share of MA. Cook County, which experiences the largest increase in average plan δ and also has the most plans available, sees the highest increases. The right-hand graph shows the share-weighted average 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.

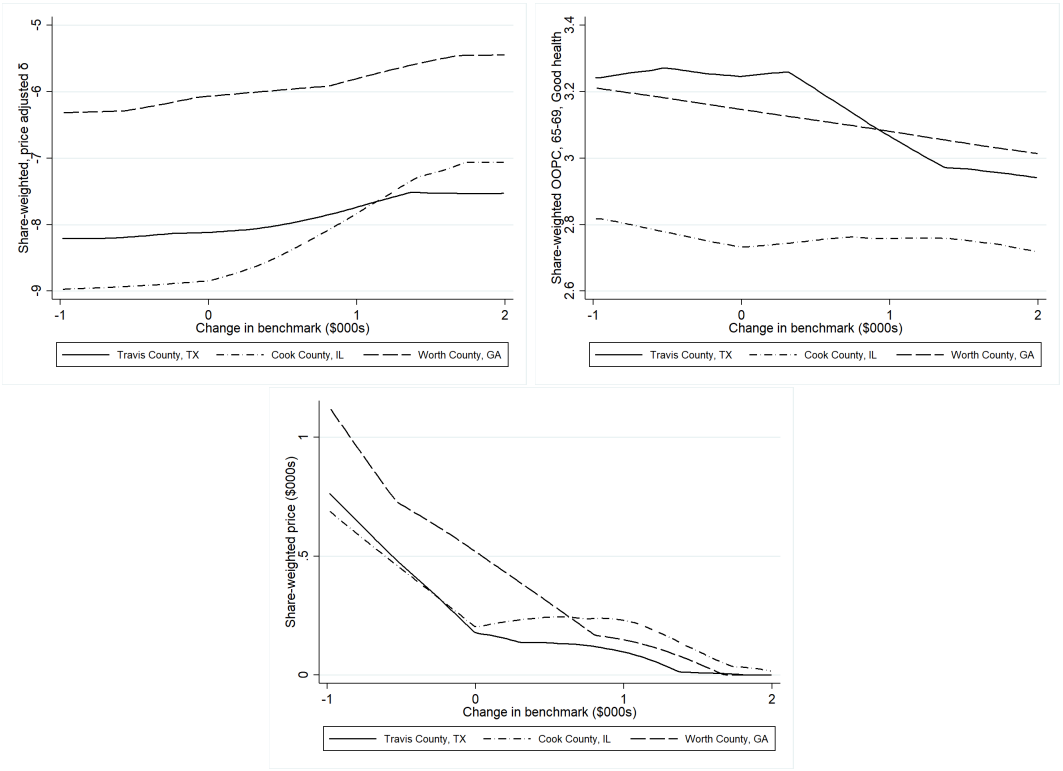
Figure 4 combines these components into the CS_m and $GovExp_m$ functions. The first graph shows per-capita consumer surplus. Under the current policy, the three counties receive similar surplus. As the benchmark in each county is increased, the average surplus in Cook County grows faster than the others, eventually overtaking them after an increase of $1,000. In the prior section,

Table 14: The determinants of differences in surplus derivatives

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\partial TS/\partial B$								
Mean	-.015	-.013	.089	.103	.023	.016	.127	.234
IQR	.200	.209	.089	.247	.151	.222	.046	.073
MAD		.011	.168	.139	.066	.035	.178	.108*
$\partial CS/\partial B$								
Mean	.099	.099	.306	.216	.104	.105	.295	.314
IQR	.092	.092	.218	.170	.091	.085	.195	.138
MAD		0.0	.207	.137	.059	.011	.205	.072*
$\partial GS/\partial B$								
Mean	.114	.112	.217	.112	.081	.089	.168	.079
IQR	.216	.216	.289	.271	.166	.224	.183	.076
MAD		.011	.171	.094	.076	.030	.145	.104*
$\partial TM/\partial B$								
Mean	-1.46	-1.48	-2.39	-2.20	-1.50	-1.79	-2.67	-3.02
IQR	1.19	1.22	1.18	1.18	1.53	1.35	.871	1.10
MAD		.094	1.37	1.04	.591	.340	1.35	.436*
$\partial MA/\partial B$								
Mean	1.58	1.60	2.61	2.31	1.58	1.88	2.84	3.10
IQR	1.20	1.24	1.10	1.14	1.58	1.34	.882	1.14
MAD		.102	1.36	1.06	.613	.315	1.40	.356*
Input spec.								
Risk	D	I	D	D	D	D	D	D
Marginal costs	D	D	I	D	D	D	I	I
Product chars.	D	D	D	I	D	D	I	I
Demographics	D	D	D	D	I	D	D	D
Switching costs	D	D	D	D	D	I	I	I
Number of products	D	D	D	D	D	D	D	I

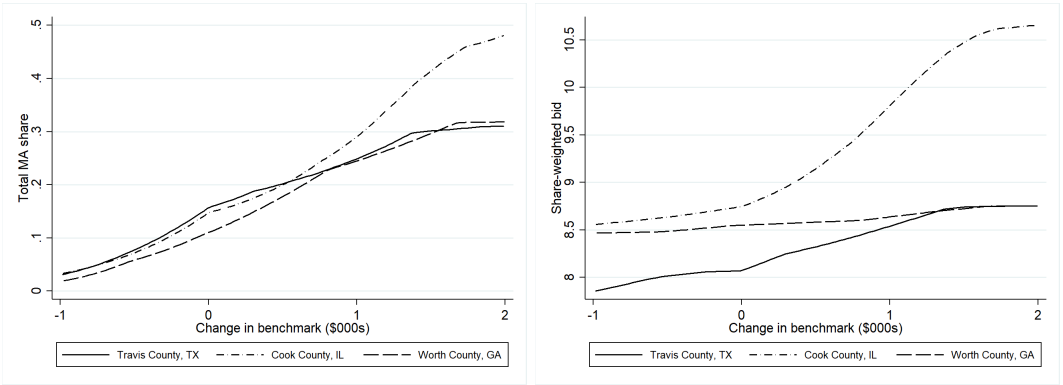
Notes: This table reports the mean and inter-quartile range ('IQR') of the derivatives of per-capita total surplus, consumer surplus, and government spending across markets in our 2015 data. Columns indicate different specifications of model inputs, which are detailed in the lower panel. 'I' indicates that the relevant feature of the input is identical for all markets, 'D' indicates that the feature is taken from the data. The first column reports the full model and matches the reported values in Table 11. Subsequent columns 'turn off' heterogeneity in the input. In the later columns, 'MAD' is the mean absolute distance between the value of the derivative under that set of input conditions and the value of the derivative in the full model per Equation 18. For Column (8), the MAD (denoted by *) is calculated with respect to Column (7), to illustrate the marginal effect of turning off heterogeneity in the number of products offered. See text for details. All derivatives are calculated numerically by simulating a $1 increase in the benchmark rate for each market and comparing the outcomes to those at the 2015 policy.

Figure 2: Product characteristics under counterfactual benchmarks, selected counties



Notes: These graphs illustrate three outputs from our counterfactual simulations: the price-adjusted δ_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). The horizontal axis is the change in the benchmark relative to the 2015 level. All outputs are averages across plans weighted by the market share of each plan.

Figure 3: Average plan bids and MA share under counterfactual benchmarks, selected counties



Notes: 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). The horizontal axis is the change in the benchmark relative to the 2015 level.

we found that the pricing mechanism was the biggest contributor to gains near the 2015 policy. This graph illustrates that this remains true even at some distance from the 2015 policy; though increases in CS still occur after all plans are offered with zero price, the slope is closer to zero. The second graph illustrates per-capita government expenditures. This graph illustrates the potential gains noted in the previous section: Cook and Travis Counties have a flat or even decreasing level of government expenditures for modest increases in the benchmark rate. These graphs suggest that significant gains are possible in some markets simply by incentivizing switches from TM to MA.

The second graph also illustrates the effect of the MA payment system (Equation 1), which rebates a portion of the difference between the bid and the benchmark back to firms. Despite the fact that bids in Worth County stay nearly constant, expenditures increase with the benchmark.

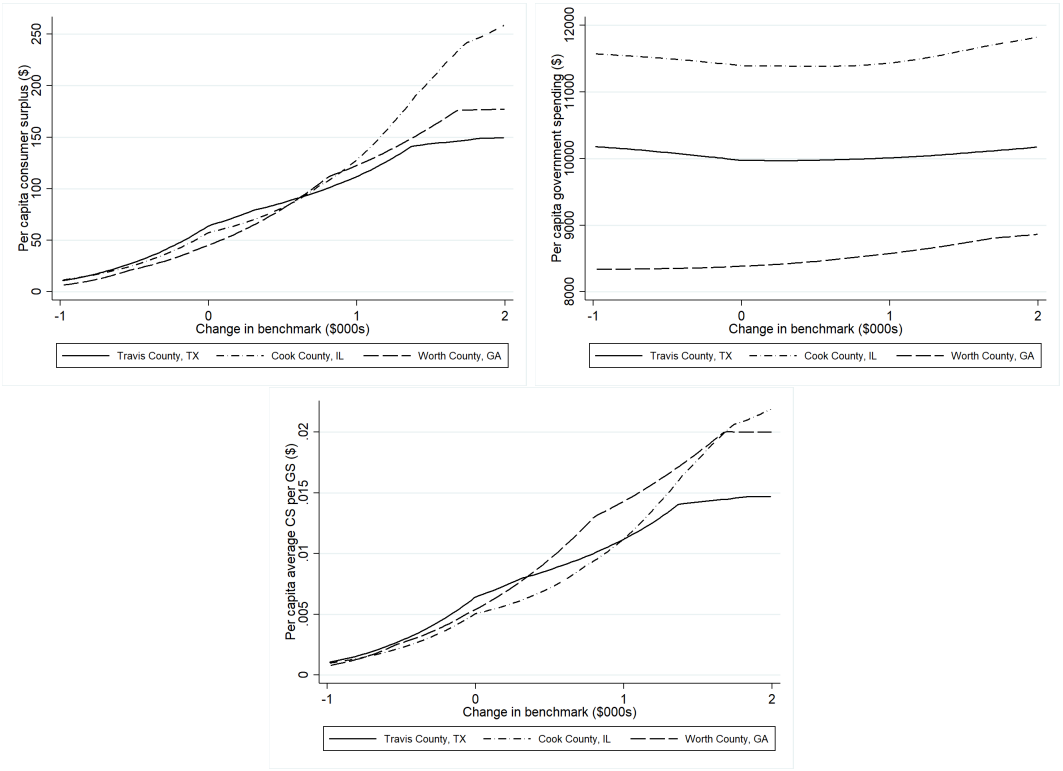
The final graph of Figure 4 combines the two functions to show the average MA surplus delivered to consumers per dollar spent by the government on the Medicare program. The slope of this line is related to the marginal impact of spending an extra dollar in a particular county through the MA benchmark 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.

8.4 Maximizing consumer surplus

Table 15 reports the results of our counterfactual exercise. Column (1) reports surplus under the 2015 policy across a number of dimensions. The mean consumer surplus is \$148.65, and the interquartile range extends from \$9.63 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 more surplus than either medium or high income individuals. The mean variable MA profit per Medicare beneficiary, 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 13. The mean consumer surplus increases 29.6% to \$192.65; the interquartile range extends from \$16.98 to \$268.69. Both Whites (\$191.01) and Blacks (\$203.90) receive increases, though 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 expenditure constraint is satisfied to within \$0.50 per person.

Figure 4: Per-capita consumer surplus and government expenditures under counterfactual benchmarks, selected counties



Notes: These graphs illustrate two outputs from our counterfactual simulation: the per-capita consumer surplus and the per-capita government spending. The third graph illustrates the ratio of consumer surplus to 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). The horizontal axis is the change in the benchmark relative to the 2015 level.

Table 15: Surplus under the 2015 policy and the optimal policy

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

Notes: This table reports results of solutions to the optimal subsidy problem presented in Section 7. The first column reports data for the 2015 policy. The second column reports the results of maximizing consumer surplus per Equation 13. The third column reports the results of minimizing government expenditures, and the fourth column minimizes government expenditures with the constraint that no benchmark is decreased. 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.

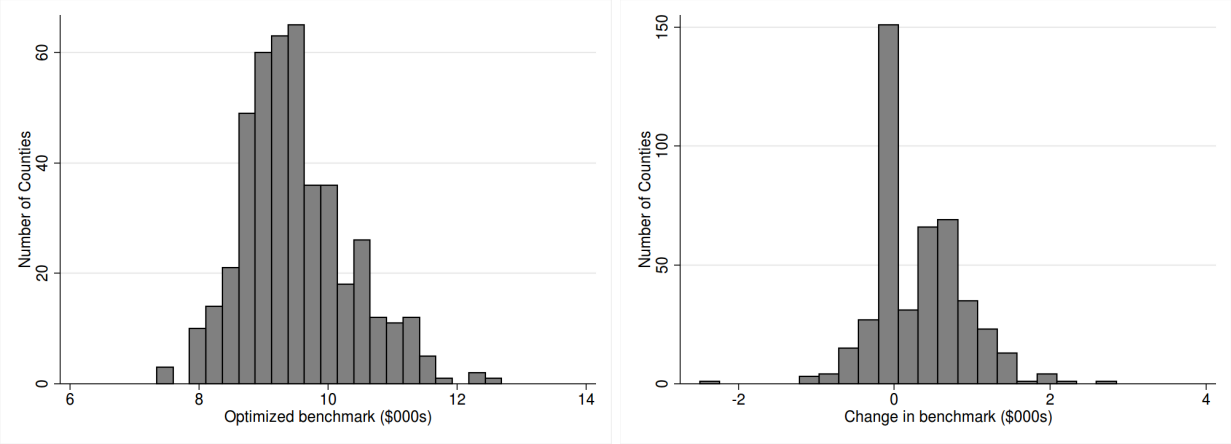
The benchmark changes are detailed in Figure 5. The left-hand histogram illustrates the distribution of benchmarks under the optimal policy, and can be compared to Appendix Figure B.1. The right-hand histogram shows the distribution of changes in the benchmark. Most counties receive modest increases – 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%.

The differences between the optimal policy and the 2015 policy may have political economy implications if the changes result in a large-scale redistribution of government expenditures and consumer surplus dollars from states with one political alignment to states with an opposing alignment. To explore these issues, we summarize the total consumer surplus and government expenditures by state in Appendix Table B.4. The changes in expenditures are small; the largest magnitude change is in Louisiana, where expenditures drop by 3.71%. Of the 41 states (plus Washington D.C.) included in the MCBS, 35 receive increases in consumer surplus; only Arkansas, Florida, Louisiana, Oklahoma, South Dakota, and Wisconsin experience decreases. Other southern states including Alabama, Georgia, North and South Carolina, Texas, and Virginia receive surplus increases. These results suggest the optimal policy does not split along political divisions.

A natural question is how much of the gains come from intensive/extensive margin-related savings opportunities and how much comes from redistribution after those savings have been achieved. To answer this question, we re-optimize the benchmarks with the objective of minimizing government expenditures. Column (3) of Table 15 reports the results. The last row reports a decrease in expenditures per person from \$9,797 to \$9,747. This savings comes with a reduction in mean consumer surplus of 20% 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.

To understand the differences in the results of our approach relative to an approach that uses county characteristics to determine benchmarks, we obtain 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-

Figure 5: Optimized benchmarks by county when maximizing average consumer surplus



Notes: These graphs illustrate the solution of Equation 13. The left-hand graph shows the distribution of benchmarks under the optimal policy and the right-hand graph illustrates the distribution of the change in benchmarks from the 2015 policy to the optimal policy.

relevant information at the county-year level. We collect median household income, the percentage of seniors in severe poverty, the unemployment rate, population density, and per-capita counts of MDs, hopsitals, skilled nursing facilities, and hospice facilities. We also collect the hospital readmission rate for Medicare patients and the "preventable" hospital admission rate.

Table 16 reports the means of 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 the current policy sorts counties along several dimensions, including income, unemployment rate, and MDs and hospitals per capita.[27] The bottom panel shows the optimal benchmarks create dispersion in TM spending across quartiles. The benchmarks themselves track TM spending more closely. The optimal policy also disrupts the sorting of the current policy – the number of beneficiaries, the unemployment rate, and the number of hospitals change non-monotonically across quartiles. The optimal policy breaks the monotonicity in the number of MA insurers and plans in each county. These results provide 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.

This result raises questions of implementability. To understand the degree to which our optimal benchmarks could be reproduced using market-level data, we regressed the logarithm of the optimal benchmarks on the logarithm of TM costs and the other variables in Table 16, except for the number

[27]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 plans and firms. The full results are presented in Appendix Table B.2. 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 1% 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, roughly half of the size of the interquartile range of the difference between the optimal benchmarks and current policy.

Appendix Table B.3 compares the surplus generated by the optimal benchmarks to the regression-fit benchmark. 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.

8.5 Alternative social welfare functions

The increased mean consumer surplus in the optimal policy comes with an increase in the variance, from 47,920 under the 2015 policy to 61,504. We also explore the family of objective functions represented by Equation 14 in Table 17. Column (1) repeats the 2015 policy outcomes. Column (2) places 99% of the weight on 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 similar. Insurers capture, on average, $162.76 in variable profit per Medicare beneficiary. Appendix Figure B.4 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%.

Column (4) of Table 17 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) though variable profits are 5.3% lower.

Table 16: 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 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
Optimal 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

Notes: This table reports county characteristics from CMS, Census, and Area Health Resource File data across benchmark quartiles. The top panel defines benchmark quartiles according to the 2015 policy and sorts counties according to their 2015 benchmark. The bottom panel defines quartiles according to the optimal policy and sorts counties accordingly.

Table 17: Surplus results under 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

Notes: This table reports results of solutions to variations of the optimal subsidy problem presented in Section 7. The first column reports data for the 2015 policy. The second and third columns reports the results of maximizing consumer surplus with a penalty for the variance in consumer surplus per Equation 14. The fourth column maximizes consumer surplus under the condition that no benchmark is decreased per Equation 15. 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 is $9,797 across all specifications.

9 Conclusion

Seeking to reduce the perceived inefficiency of government-provided goods and 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 consumers. 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 good over time. In many cases, the goods provided by firms have differentiated characteristics which are relevant to consumers. Additionally, these goods 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.

Angrist, J. D. and Pischke, J.-S. (2010). The credibility revolution in empirical economics: How better research design is taking the con out of econometrics. *Journal of Economic Perspectives*, 24(2):3–30.

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

Baum-Snow, N. and Marion, J. (2009). The effects of low income housing tax credit developments on neighborhoods. *Journal of Public Economics*, 93(5):654–666.

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.

Bjornerstedt, J. and Verboven, F. (2016). Does merger simulation work? evidence from the swedish analgesics market. *American Economic Journal: Applied Economics*, 8(3):125–64.

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.

Hoxby, C. M. (2000). Does competition among public schools benefit students and taxpayers? *American Economic Review*, 90(5):1209–1238.

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.

Levin, H. M. (2002). A Comprehensive Framework for Evaluating Educational Vouchers. *Educational Evaluation and Policy Analysis*, 24(3):159–174.

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*.

Maruyama, S. (2011). Socially optimal subsidies for entry: The case of medicare payments to hmos. *International Economic Review*, 52(1):105–129.

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.

Nevo, A. and Whinston, M. D. (2010). Taking the dogma out of econometrics: Structural modeling and credible inference. *Journal of Economic Perspectives*, 24(2):69–82.

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.

Peters, C. (2006). Evaluating the performance of merger simulation: Evidence from the u.s. airline industry. *The Journal of Law & Economics*, 49(2):627–649.

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.

Weinberg, M. C. (2011). More evidence on the performance of merger simulations. *American Economic Review*, 101(3):51–55.

Appendices

A 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 do such approximations work 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.

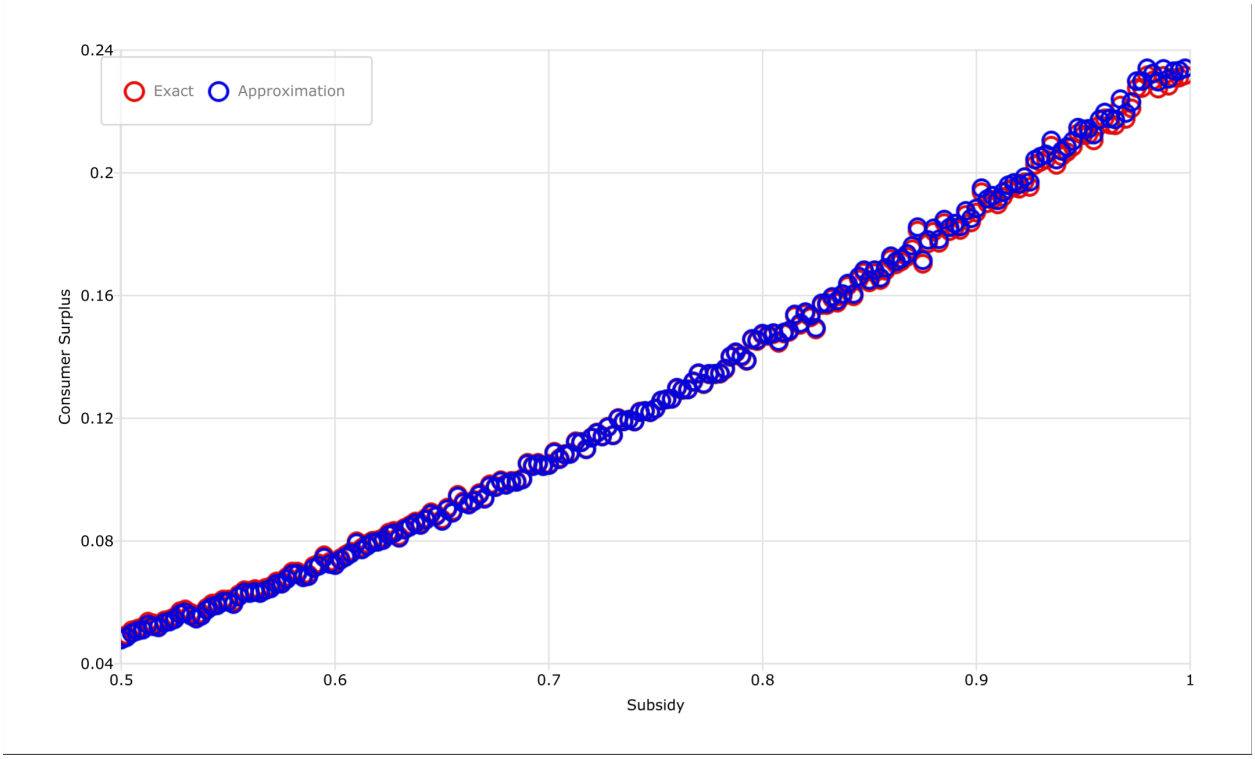
Table A.1: Monte Carlo Evidence

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

Notes: This table presents the results of a Monte Carlo exercise which compares our equilibrium approximation approach to actual equilibrium solutions for a simple model in which firms choose a single price and a single product characteristic.

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.

Figure A.1: Monte Carlo Results for Seven Firms

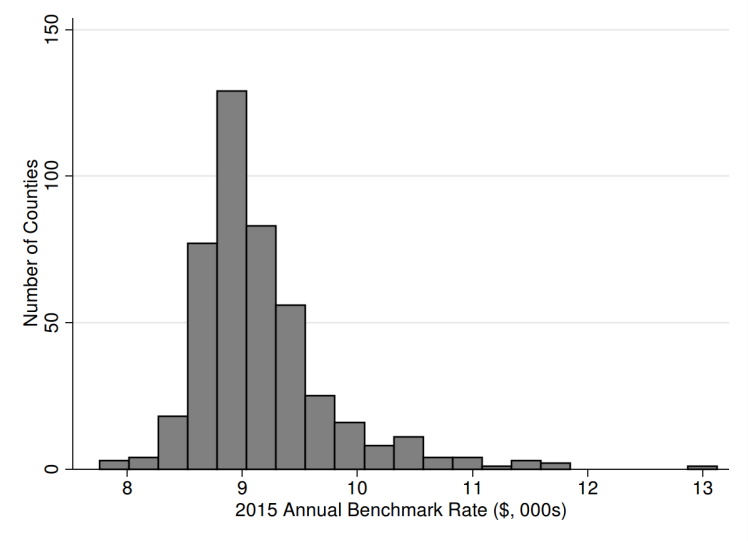


Notes: This figure illustrates the results of our Monte Carlo exercise with seven firms per market and 200 simulated markets. The horizontal axis is the subsidy from the government and the vertical axis is the equilibrium consumer surplus.

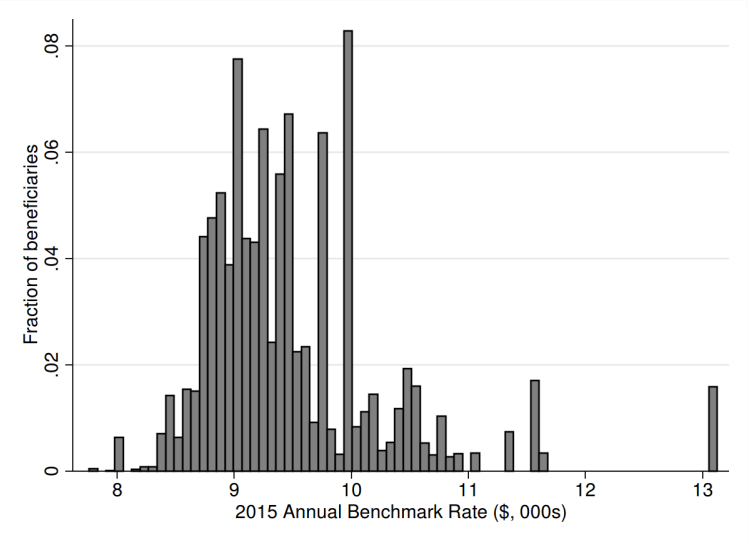
B Additional tables and figures

Figure B.1: Medicare Advantage Benchmark Distribution, 2015

(a) Benchmarks across counties

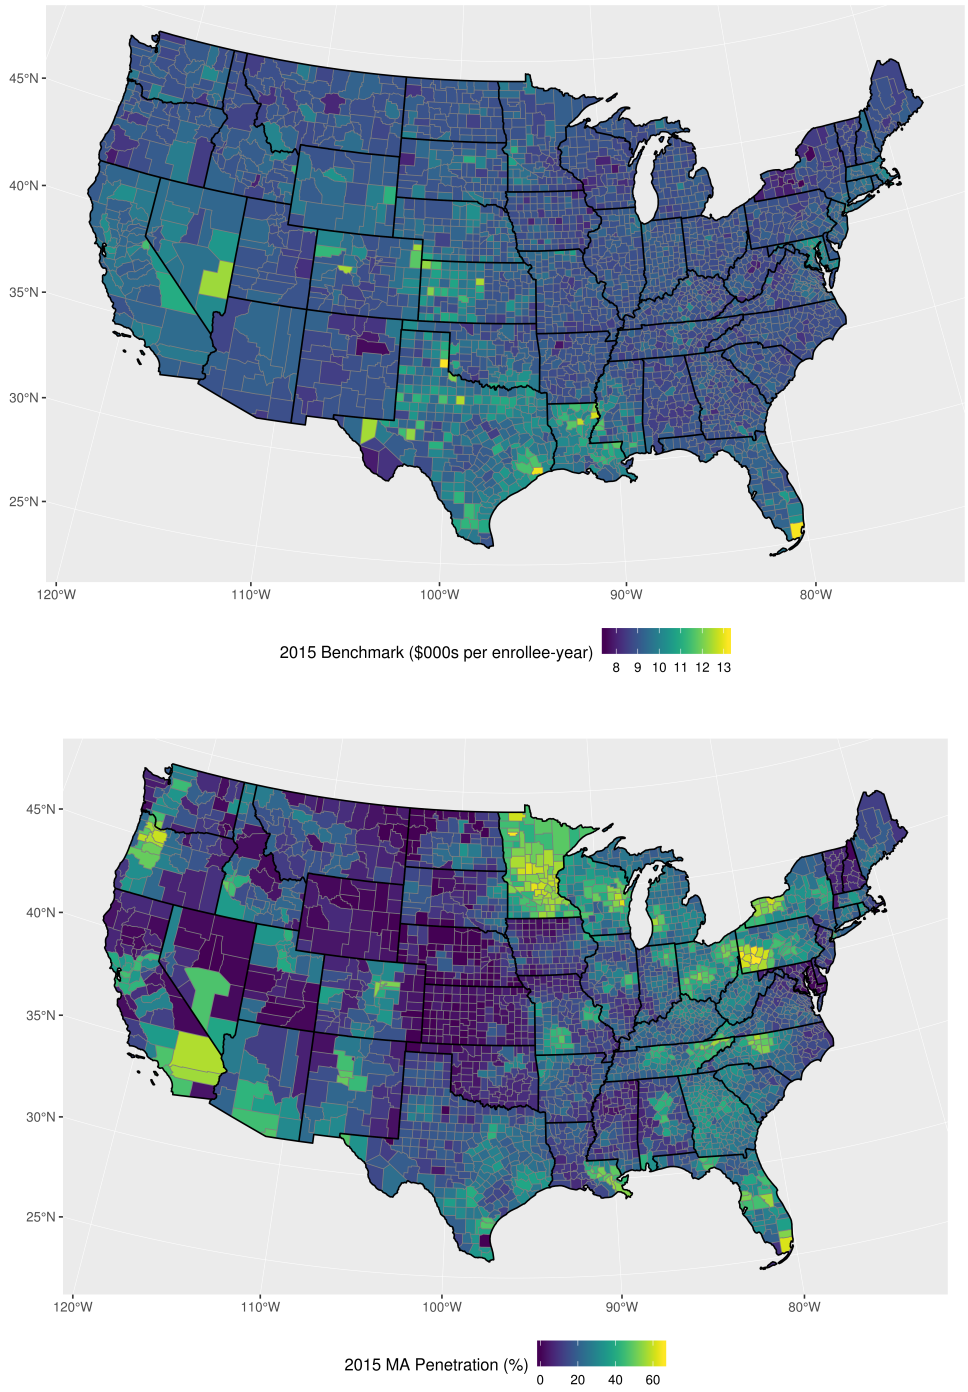


(b) Benchmarks across beneficiaries



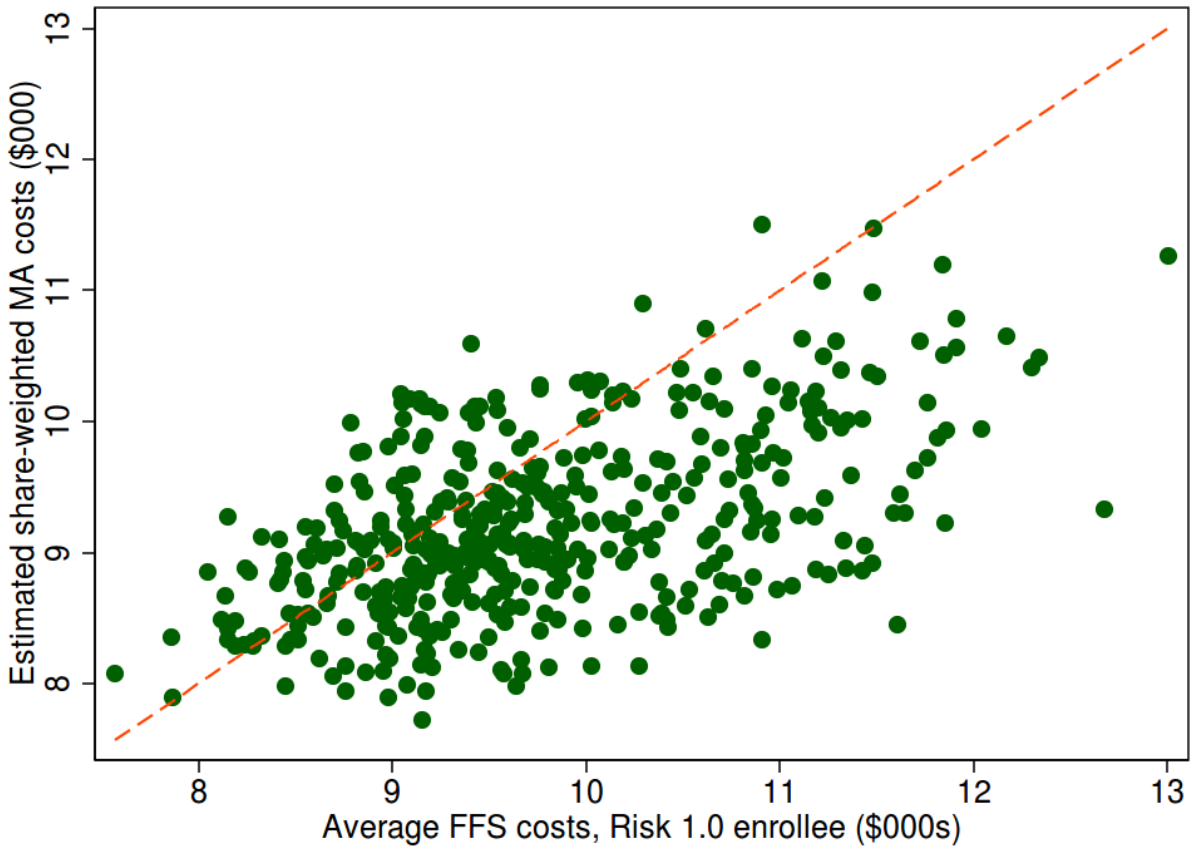
Note: Includes counties observed within the Medicare Current Beneficiary Survey.

Figure B.2: Medicare Advantage Benchmarks and Penetration by County, 2015



Notes: Data from CMS benchmark and enrollment files. Penetration is defined as the total number of people enrolled in any MA plan divided by the number of people eligible for Medicare benefits.

Figure B.3: County-level estimated MA and FFS costs



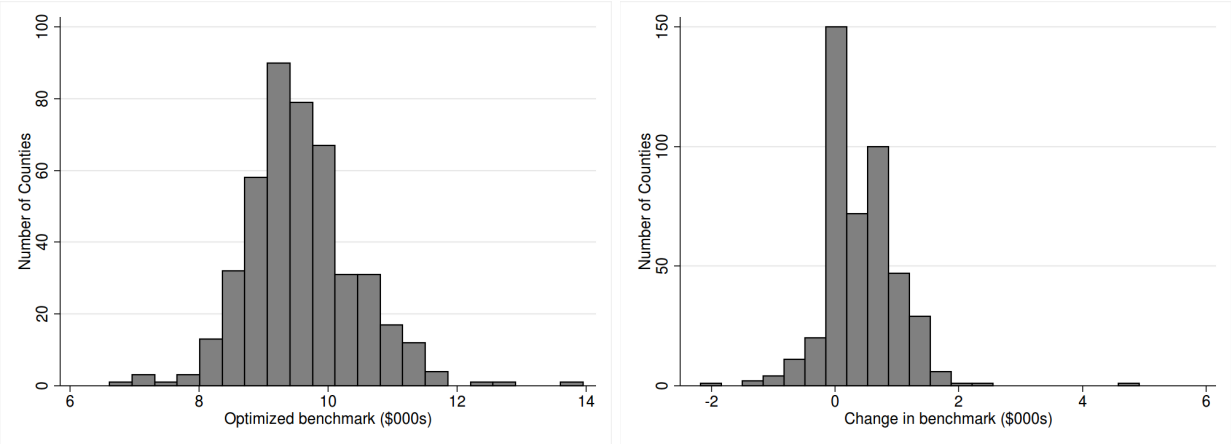
Notes: Each dot in this figure represents a county in 2015. MA costs are implied by the firms' first order conditions for optimal pricing and weighted by market share. FFS costs come from CMS data on realized expenditures and county-level average risk; both MA and FFS costs are normalized to a risk 1.0 enrollee.

Table B.1: Estimates of the out-of-pocket cost policy functions

Demographic Cat.	Benchmark	Plan Rank	Leading Insurer	# Plans	R^2
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

Notes: Each row of this table reports the OLS coefficients for a different regression. The dependent variable for each regression is the out-of-pocket cost estimate for a particular age-health demographic group. The columns are the independent variables as defined in Table 9. Each regression has 50,593 plan-market observations. Robust standard errors are in parentheses.

Figure B.4: Optimized benchmarks by county under variance-penalized objective function



Notes: These graphs illustrate the solution of Equation 14 when the weight on consumer surplus is 0.999 and the weight on variance is 0.001. The left-hand graph shows the distribution of benchmarks under the variance-penalized policy and the right-hand graph illustrates the distribution of the change in benchmarks from the 2015 policy to the variance-penalized policy.

Table B.2: New benchmark regression exercise

VARIABLES	(1) Optimal Benchmark
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

Notes: The dependent variable is the optimal benchmark as determined by solving Equation 13. Estimates obtained by OLS. Robust standard errors are in parentheses.

Table B.3: 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

Notes: This table compares surplus under the 2015 policy (Column 1), our optimal policy (Column 2), and the policy predicted by the regression results reported in Table B.2 (Column 3). 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.

Table B.4: State-level changes in surplus and expenditures from 2015 policy to optimal benchmark schedule

		Sum(weight)	Consumer surplus ($ M)			Government expenditures ($ M)		
State	# counties	(000,000s)	2015	Optimal	%Δ	2015	Optimal	%Δ
Alabama	14	14.63	130.5	166.6	27.6	13,063	13,088	0.20
Arizona	5	8.36	147.1	148.1	0.7	7,111	7,112	0.01
Arkansas	4	4.36	62.0	61.2	-1.3	3,546	3,544	-0.04
California	18	31.48	699.8	916.7	31.0	34,684	34,832	0.43
Colorado	8	3.95	118.1	150.6	27.5	3,508	3,531	0.64
Connecticut	6	6.34	115.0	199.9	73.7	6,663	6,700	0.56
Dist. of Columbia	1	0.92	1.0	2.7	163.8	1,077	1,078	0.05
Florida	22	29.73	675.5	564.8	-16.4	31,790	31,627	-0.51
Georgia	21	11.12	89.6	97.3	8.5	10,087	10,071	-0.16
Illinois	18	12.46	52.1	125.0	139.7	12,691	12,688	-0.03
Indiana	3	0.65	5.0	9.7	93.7	636	640	0.63
Iowa	4	2.74	10.8	17.4	61.3	2,166	2,164	-0.09
Kansas	3	3.59	21.1	37.1	76.2	3,156	3,161	0.19
Kentucky	8	7.34	60.3	84.9	40.7	6,799	6,808	0.14
Louisiana	5	4.44	159.4	61.0	-61.7	4,758	4,582	-3.71
Maryland	8	6.01	4.6	28.7	526.6	6,571	6,553	-0.27
Massachusetts	8	5.71	43.8	153.9	251.1	6,439	6,462	0.36
Michigan	32	21.63	233.6	465.6	99.3	22,325	22,315	-0.04
Minnesota	13	9.80	66.9	131.0	95.9	8,961	8,991	0.33
Missouri	15	8.27	94.2	106.6	13.2	7,696	7,707	0.15
Nebraska	3	2.45	13.7	41.1	200.8	2,235	2,210	-1.12
Nevada	3	4.70	115.3	115.4	0.1	4,822	4,823	0.01
New Hampshire	1	0.03	0.1	0.3	200.1	27	27	0.16
New Jersey	16	15.29	90.1	312.2	246.5	17,495	17,487	-0.05
New Mexico	5	10.60	252.2	259.5	2.9	7,495	7,495	0.00
New York	28	27.57	779.5	931.3	19.5	29,257	29,248	-0.03
North Carolina	25	18.68	290.1	301.2	3.8	17,739	17,690	-0.28
Ohio	30	18.31	228.4	473.0	107.1	17,764	17,928	0.92
Oklahoma	2	4.09	8.6	3.0	-65.2	3,592	3,581	-0.31
Pennsylvania	25	18.43	370.3	623.8	68.5	18,287	18,444	0.86
Rhode Island	1	0.03	0.5	0.5	1.4	30	30	0.02
South Carolina	7	3.67	17.2	28.6	66.4	3,422	3,426	0.12
South Dakota	14	8.26	91.8	91.3	-0.5	7,247	7,246	-0.01
Texas	33	30.77	312.6	399.3	27.7	31,891	31,896	0.02
Utah	1	0.05	0.6	0.6	0.9	46	46	0.01
Vermont	1	0.01	0.0	0.1	241.1	6	6	-0.18
Virginia	9	7.03	42.3	55.7	31.7	6,390	6,400	0.15
Washington	8	16.55	207.5	307.8	48.3	13,470	13,513	0.32
West Virginia	4	3.74	33.6	47.3	40.6	3,048	3,032	-0.52
Wisconsin	11	14.33	290.3	172.5	-40.6	12,360	12,170	-1.54
Wyoming	2	1.47	5.6	5.8	3.5	1,167	1,167	-0.01
Total	445	400	5,940	7,699	29.6	391,517	391,519	0

Notes: This table presents state-level summaries of the total consumer surplus and government expenditures at the 2015 policy and our calculated optimal policy. The first column reports the number of counties included in the MCBS in each state, and the second column reports the total sample weight in the state; the MCBS uses a sample of counties and weights observations to be nationally representative. Consumer surplus is calculated via Equation 8. Government expenditures include expenditures on TM and MA and are calculated via Equation 16. All surplus and expenditure numbers reported here are calculated using MCBS sample weights.