NBER WORKING PAPER SERIES

OPTIMAL MANAGED COMPETITION SUBSIDIES

Keaton S. Miller
Amil Petrin
Robert Town
Michael Chernew

Working Paper 25616
http://www.nber.org/papers/w25616

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2019, Revised January 2022

The authors would like to thank Mark Colas, Francesco Decarolis, Clare Evans, David Evans, Gautam Gowrisankaran, Ben Hansen, Kendall Houghton, Van Kolpin, Scott Petty, Jeremy Piger, Christopher Snyder, Glen Waddell, Wes Wilson, anonymous referees, and seminar and conference participants at the University of Wisconsin-Madison, University of Oregon, the University of Minnesota, Harvard/MIT, Northwestern University, University of Pennsylvania, University of Maryland, University of Texas - Austin, the 2019 NBER Winter IO and Insurance Meetings, 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. All errors are our own. We are grateful for support from the National Institute on Aging under grant P01 AG032952. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Optimal Managed Competition Subsidies
Keaton S. Miller, Amil Petrin, Robert Town, and Michael Chernew
NBER Working Paper No. 25616
February 2019, Revised January 2022
JEL No. I11,L13,L51

ABSTRACT

When markets fail to provide socially optimal outcomes, governments often intervene through 'managed competition' where firms compete for per-consumer subsidies. We introduce a framework for determining the optimal subsidy schedule that features heterogeneity in consumer preferences and inertia, and firms with heterogeneous costs that can set prices and product characteristics in response to changes in the subsidy. We apply it to the Medicare Advantage program, which offers Medicare recipients private insurance that replaces Traditional Medicare. We calculate counterfactual equilibria as a function of the subsidies by estimating policy functions for product characteristics from the data and solving for Nash equilibria in prices. The consumer-welfare-maximizing budget-neutral schedule increases aggregate annual consumer welfare by $4.6 billion over the current policy and is well-approximated with a linear rule using market-level observables.

Keaton S. Miller
University of Oregon
Department of Economics
1285 University of Oregon
Eugene, OR 97403-1285
keatonm@uoregon.edu

Amil Petrin
Department of Economics
University of Minnesota
4-101 Hanson Hall
Minneapolis, MN 55455
and NBER
petrin@umn.edu

Robert Town
Department of Economics
The University of Texas at Austin
2225 Speedway
BRB 1.116, C3100
Austin, TX 78712
and NBER
robert.town@austin.utexas.edu

Michael Chernew
Harvard Medical School
Dept. of Health Care Policy
180 Longwood Avenue
Boston, MA 02115
and NBER
chernew@hcp.med.harvard.edu

1 Introduction

When the private market's provision of a good or service deviates from the socially optimal outcome, welfare may be improved through government intervention. Often the obvious intervention is 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 incentives to efficiently design and deliver goods (McKean and Minasian, 1966). This concern has led to the idea that governments can either regulate private firms or procure goods directly from those firms. These approaches 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 idea that profit motives and market pressures will push firms to provide the optimal quantity, variety, and quality at a price nearing marginal cost. This "managed competition" approach is employed by the US government to provide health insurance where market failure is a long-standing concern (Arrow, 1963, Grossman, 1972, Rothschild and Stiglitz, 1976). For example, under the "Medicare Advantage" program (MA) we study, Medicare beneficiaries can forgo Traditional Medicare (TM) fee-for-service benefits and enroll in one of a variety of health plans offered by private insurers. The insurer assumes financial and logistical responsibility for the enrollee's care and receives a risk-adjusted per-capita payment from the government based on a county-specific "benchmark rate" that varies considerably across counties. Similar approaches are used by Medicare Part D and the insurance markets created by the Affordable Care Act (Gruber, 2017). Elements of this approach appear in education, where public, charter, and private primary and secondary schools compete on program offerings, education quality, and productive efficiency (Poterba, 1996, Hoxby, 2000), as well as housing policy, where construction is influenced by differences in tax credits across geographies (Baum-Snow and Marion, 2009).

In this paper, we develop an approach for calculating the optimal subsidy schedule across heterogeneous markets in managed competition settings. We take the government's budget constraint as exogenous—solving for the optimal budget amount is a much more challenging problem although we provide evidence on the efficiency of the current budget level. In our setup, firms choose prices and product characteristics in response to the subsidy set by the government and other competitive conditions. Consumers are heterogeneous and choose plans based on observable plan characteristics and unobservable (to us) plan-specific quality. We take the mechanism that links a market-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.

To notate the problem, consider a government which seeks to maximize consumer welfare by allocating a fixed budget $\bar{B}$ across M markets denoted by m. Under managed competition, the government chooses a schedule of market-level subsidies $\{B_m\}$ (i.e. the benchmark rates). Let $CS_{mi}(B_m)$ be the welfare for consumer i in market m and $GovExp_{mi}(B_m)$ be the government spending on that consumer as a function of the subsidy. The optimal subsidy problem is

$$\max_{\{B_m\}} \sum_{m=1}^{M} \int_i CS_{mi}(B_m)di \quad \text{s.t.} \quad \sum_{m=1}^{M} \int_i GovExp_{mi}(B_m)di = \bar{B}. \tag{1}$$

The solution depends on the derivatives of the CS and $GovExp$ functions, which in turn depend on equilibrium interactions between firms and consumers.[1] The CS generated by a dollar of government spending—i.e. how much of that dollar is passed-through to consumers—in any given market depends on the demand elasticity, firms' cost functions, and competition. These objects likely vary across markets and therefore the optimal subsidy should also vary.[2] In practice, however, the subsidy schedule may be determined by summary measures which do not take supply and demand factors fully into account. For example, MA benchmark rates are set as a function of average risk-adjusted county-level TM costs, which may differ from private firms' costs and may be unrelated to demand conditions. Similarly, many charter schools receive government funding based on the per-pupil cost of public schools in the area (Hoxby, 2000). Political dynamics can also affect subsidy rates (Adrion, 2020).

[1] We do not include the cost of public funds in Equation (1) since as the budget is fixed reallocating subsidies does not change the cost of public funds. We also do not include firm profits as we view the government's normative objective as maximizing the direct well-being of its citizenry. Including firm profits in the problem is straightforward.

[2] Interactions between all of these effects are important. While a planner would seek to move resources to areas where consumers are more elastic *ceterus paribus*, if those areas also feature firms which do not pass-through subsidies to benefits at a high rate, it may be optimal to decrease the subsidies in those areas.

Solving for the optimal subsidy schedule requires calculating outcomes under counterfactual subsidies for each market. The traditional approach to computing counterfactual equilibria is to search for a fixed-point in firms' best-response functions (e.g. Fan, 2013, Wollmann, 2018). This approach is impractical in our setting due to the large number of markets and the complexity of MA products, yet, given that many plans are offered at a supplemental premium of $0, modeling non-premium plan features is essential. We introduce a new counterfactual approach which we believe to be of independent methodological interest. 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 premiums taking those characteristics as given. We provide Monte Carlo evidence that this approach well-approximates the equilibria calculated by explicitly solving the firms' best response functions.

We apply our approach to calculate the optimal MA subsidy schedule. The first step is to estimate Medicare beneficiaries' preferences over MA plans. Using detailed individual panel data on consumer demographics, choice sets, realized choices, and aggregate market-level plan shares for the years 2008-2017, we estimate a flexible demand system. Our demographic variables include a self-reported health status, age, race, educational attainment, and income which allow us to capture plan preferences which vary with these variables. The panel nature of our data allows us to estimate switching costs, which are relevant due to the prevalence of narrow provider networks.

To estimate the demand parameters, we first estimate the individual-specific parameters via maximum likelihood and recover plan-specific mean utility estimates. We then use these mean utilities to estimate individual-invariant plan characteristic preferences. Our model implies that premiums and plan characteristics are endogenous which invalidates many standard instruments. We leverage detailed plan cost data submitted to the government as part of the regulatory process and construct, for each plan, the mean costs of that plan's competitors for required coverage (i.e. the items and services covered by TM). Under a set of assumptions (which we test and verify as likely holding), we derive instruments for plan characteristics from the panel nature of the data.

Our estimates imply beneficiaries are on average price-sensitive with mean implied plan premium elasticities of -5.48. Higher income beneficiaries are less premium sensitive. MA plans are more attractive to younger Medicare recipients, non-Whites, and those with lower educational attainment. The average cost to switch between MA insurers is \$340, which is comparable to the enrollment-weighted average annual plan premium of \$415. The average consumer values an average dental, vision, and hearing coverage package at \$570 per year.

The second step in calculating the optimal benchmarks is to estimate plan marginal cost functions. We invert firms' implied first-order conditions for premiums to infer marginal costs accounting for regulatory requirements. We then estimate the relationship between plan characteristics and the implied marginal cost allowing for carrier-specific differences in the cost of providing a given benefit. The marginal cost function estimates align with independent utilization information. Combined, our estimates imply that in 2017, MA generated a total of \$5.70 billion in consumer surplus as measured by aggregating individual-level compensating variation and \$3.51 billion in variable profit with \$122 billion in total government payments to MA plans in the category we study.[3]

The final step prior to calculating the optimal benchmark is to estimate the plan characteristic policy functions. Our policy function estimates appear sensible – increasing the benchmark increases plan benefit provision and reduces patient out-of-pocket expenses. We combine the demand estimates, the plan marginal cost function, the policy functions and the implied premium setting first-order conditions together and use a multistart search to solve Equation (1).

We find the optimal MA subsidies are meaningfully different from the current policy – the average absolute difference between the optimal benchmark and the 2017 benchmark is \$475 or 4.84%. We find that the optimal benchmarks increase aggregate consumer welfare to \$10.26 billion per year through a combination of increasing the total share of MA from 29.2% to 41.1% and increasing the mean compensating variation for MA enrollees from \$464.68 to \$593.97. Changes in product characteristics are responsible for 36.2% of the total change in consumer welfare. Relative to the 2017 benchmark policy, the optimal policy creates winners and losers—average compensating variation for MA enrollees in markets that receive a higher (lower) benchmark increases (decreases) by \$159 (\$34). We show that the

[3]We exclude plans designed for those who are 'dual-eligible' for Medicare and Medicaid, see Section 3.

derivatives of the CS and $GovExp$ functions are related to market-level observables, and that a linear rule using these observables is able to obtain over 95% of the gains from the optimal policy with a 0.30% increase in spending. We explore other social welfare functions and find benchmark schedules that increase the aggregate consumer welfare and reduce the variance relative to the 2017 policy. We test our counterfactual equilibria by calculating the profitability of deviations from our predicted choices and find that firms' first-order conditions are reasonably satisfied.

To understand the intuition behind our results, it is helpful to note that our estimates imply that MA subsidies are, at the margin, too generous from a social surplus perspective. We find that an across-the-board increase in benchmarks by $1 from the 2017 policy increases MA expenditures by $141 million and increases aggregate consumer welfare by $10.4 million. This, plus the large variation in benchmarks, implies that the marginal utility of a dollar of benchmark (mediated through the competitive interactions of plans) is likely relatively low in current high payment areas whereas in lower payment areas the marginal utility of an extra dollar of benchmark is higher. Reallocating subsidies from high benchmark areas to low can increase welfare because of the differential marginal value. The oversimplification in the above discussion is that we allow the marginal value of an increase in benchmark (to enrollees) and the plans' marginal cost of providing benefits to vary across markets so that our optimal policy may in fact reallocate benchmark away from some low benchmark counties if their marginal utility from an increase in the benchmark is low.

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), Aizawa and Kim (2018) and Curto et al. (2021). 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. (2021) estimate a similar model using more recent data and find that the program generated approximately $600 in per-capita 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. We innovate with respect to these papers by adding rich data on demographics and product characteristics, and considering counterfactual subsidy policies. Aizawa and Kim (2018) estimate a demand model that is similar to ours in order to explore the role of advertising in equilibrium selection. Our demand model innovates upon theirs by adding additional heterogeneity in switching costs.

There is also 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 pass-through rate of 54%, while Duggan et al. (2016) use variation in the benchmark across urban and rural counties and estimate a smaller pass-through. Song et al. (2013) calculate a pass-through from benchmarks to plan bids, which are a measure of premiums and the actuarial value of benefits, of 53%. We expand upon this literature by considering firms' plan design decisions in response to benchmark changes and measuring consumer valuations of those plans.

Our work is also related to research on optimal subsidy structures in 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 premium regulation in an ACA-like insurance exchange. Bundorf et al. (2012) study health-status-linked premiums for employer-sponsored plans. 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 reasonably successfully (Newhouse et al., 2015), and therefore we do not model the role of enrollee selection not captured by risk-adjustment.[4] Decarolis et al. (2020) examine the optimality of using vouchers versus the current subsidy strategy in Medicare Part D and find that the two systems generate similar welfare. We innovate by examining the impact of subsidies on non-premium plan characteristics and calculating counterfactual outcomes.

Finally, our counterfactual 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 premiums and product characteristics to calculate the welfare gains caused by the introduction of satellite TV. Sweeting (2007) uses a

[4]We test for the presence of residual selection after risk-adjustment and, consistent with Newhouse et al. (2015), fail to find meaningful selection in our data.

similar approach to study radio station repositioning. 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 an equilibrium in premiums and calculate the welfare effects of policy changes.

We discuss the institutional details of the Medicare Advantage program in Section 2 and detail our data on Medicare beneficiaries and MA plans in Section 3. We present a model of demand for MA in Section 4 and discuss the supply side in Section 5. We describe our estimation procedure in Section 6 and present estimates in Section 7. We describe our counterfactual approach and present our results in Section 8. We conclude in Section 9.

2 The Medicare Advantage Program

Enacted in 1965, Traditional Medicare (TM) provides health care to seniors (age 65 or older) through its Part A (hospital) and Part B (physician and outpatient) insurance programs. Under TM, Medicare pays service providers according to a pre-set fee-for-service (FFS) reimbursement schedule while beneficiaries pay applicable co-pays and/or coinsurance. Eligibility has since expanded to include those eligible for federal disability benefits and end-stage renal disease (ESRD) patients.

In response to Medicare's increasing costs, in 1982 Congress authorized Medicare administrators to engage in a series of "Part C" trials based on the ideas of Enthoven (1978) 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. To the extent that the rise in cost was driven by principal-agent problems, this mechanism was seen as a way to ensure that providers bore more of the financial risk of medical decisions (Smith et al., 1997). This program was brought to the entire country in 1997 under the name Medicare+Choice.

Medicare+Choice initially struggled to attract plans and nationwide enrollment hovered near 5 million – less than 10% of those eligible. Critics blamed low subsidy rates and the fact that flat payments incentivized firms to cream-skim relatively healthy individuals from the risk pool. The Medicare Prescription Drug, Improvement, and Modernization Act of 2003 aimed to remove this incentive by risk-adjusting payments. 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; a score of 1.0 indicates average risk. 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 program was renamed Medicare Advantage (MA). By 2015, 95% of Medicare beneficiaries had access to MA and enrollment reached 16.3 million.[5]

MA enrollees forgo TM benefits and receive medical benefits from their MA plans exclusively. MA enrollees pay the Medicare Part B premium and may pay a private plan premium as well. Insurers compete along the dimensions of benefit design, premiums, and provider networks, and often heavily market their plans (Aizawa and Kim, 2018). Plans generally offer a set of 'in-network' providers which enrollees may utilize with lower cost-sharing than 'out-of-network' providers. MA plans generally provide a more generous benefit package than TM, such as including dental, vision, and/or hearing coverage (DVH). Many plans include a drug benefit. Plans may also offer a reduction in the Part B premium.

The enrollee-specific subsidy from CMS to insurers is based on a "benchmark" rate for each county, which varies across geographies and over time and is not influenced by MA firms (Newhouse et al., 2012). CMS calculates the benchmark schedule each year using the average risk-adjusted per-capita FFS Medicare spending within the county. Counties are ranked by average spending and placed into quartiles. The benchmark for counties in the top quartile is set to 95% of their FFS spending. The benchmark for the second quartile is 100% of FFS spending, the third quartile benchmark is 107.5% of spending, and the bottom quartile benchmark is 115% of spending. A floor that varies by urban/rural status applies.

Each year, after benchmarks are published by CMS, insurers submit detailed proposals to provide MA plans. These 'bids' include benefit and cost-sharing designs, and detailed information about the insurer's expected revenues and expenses for both TM-covered and

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

non-TM-covered services. Ultimately, the 'bid amount' represents the insurer's offer to provide all services covered by TM to a person of average risk in the plan's coverage area in exchange for a particular level of revenue. The bid amount must be related to the firm's projected costs and may be above or below the benchmark rate. Firms that bid above the benchmark must charge premiums to enrollees. Firms that bid below the benchmark receive a portion of the difference as a 'rebate' that must be passed on to consumers through decreases in cost-sharing (e.g. reductions in copays) or by offerings of services not covered under TM (e.g. dental). Supplemental benefits may also be paid for by an additional premium. MA plans that offer a prescription drug benefit submit a separate bid which maps in a similar way to a Part D premium.

The rebate payment varies across firms and over time based on the CMS 'star rating' measure of insurer quality; payments in our data vary from 50-75% of the difference between the benchmark and the bid. Under current policy, firms with at least four stars (out of five) also receive a 5% bonus to the benchmark rate. The star rating itself is a summary of multiple measures of past service quality which change throughout our study period, such as the fraction of plan members receiving influenza vaccinations, the 30-day hospital readmittance rate, and enrollee assessments of care quality.

Beneficiaries can enroll in plans during an fixed 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 certain life events. These rules are designed to reduce adverse selection.[6] After enrollment, firms collect and transmit risk-adjustment information to CMS.

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

$$Payment_{ijt} = \begin{cases} B_{mt} \times \phi_{jt} \times Risk_{it} & \text{if } bid_{jt} \geq B_{mt}\phi_{jt} \\ (bid_{jt} + \lambda_{jt} \times (B_{mt} \times \phi_{jt} - bid_{jt})) \times Risk_{it} & \text{if } bid_{jt} < B_{mt}\phi_{jt} \end{cases} \tag{2}$$

where ϕ_{jt} captures any bonus to the benchmark rate and λ_{jt} is the rebate percentage. We

[6] Since 2014, enrollees have been allowed to switch to a "5 star" plan at anytime. As only 1% of enrollees switch plans mid-year in our data, we treat any mid-year switchers as choosing the plan in which they spent the most time.

denote the market-level (i.e. county-year level) benchmark with B_m and denote risk-neutral (i.e. $Risk = 1.0$) plan-specific payments with $B_{jt} \equiv B_{mt} \times \phi_{jt}$.[7]

MA is a significant component of the federal budget. In 2017, payments to plans in our data were $122 billion and TM spending on the individuals in our data totaled $317 billion. MA is also relatively concentrated: the top five firms nationwide, Aetna, Blue Cross Blue Shield, Humana, Kaiser Permanente, and UnitedHealth Group, have 65% of total enrollment. The average beneficiary has access to 10 plan options with 64% of beneficiaries having access to 5 or more plans. 25% of beneficiaries in our 2015 data have access to 3 or fewer plans. The average bid is 90% of TM costs (MedPAC, 2017).

Figure 1 illustrates the 2017 policy and the resulting market outcomes with county-level maps of the US.[8] The left map illustrates the ratio of the 2017 benchmark rate to the average FFS spending in 2017. The right map illustrates the total MA share in each county. As consumer surplus is related to the total MA share, these graphs offer a simple check of the current government policy. If private costs are tightly linked to the government's costs and differences in those costs were the only source of heterogeneity across markets, then we would expect those areas which had larger benchmarks relative to FFS spending to have greater enrollment. Instead, we see significant deviation from this pattern. Some areas with high relative benchmarks, such as much of New Mexico, do not have particularly high enrollment, while other areas with high enrollment, such as Minnesota and southwestern Pennsylvania, do not have particularly high relative benchmarks. This suggests that there may be gains by redistributing government funds across counties.

3 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.

[7] 'Regional PPOs'—certain plans offered in one or more entire states—set premiums and benefits as other plans do, but face a slightly different payment system. For computational tractability, we assume Regional PPO plans operate identically to other plans. As Regional PPO plans have a total market share of 1.0% in our data, our results are not likely to be affected by this assumption.

[8] Appendix Figures E.1 presents another view of the benchmark distribution.

Figure 1: 2017 Medicare Advantage Benchmarks Relative to Traditional Medicare Spending, and Relative Market Penetration, by County

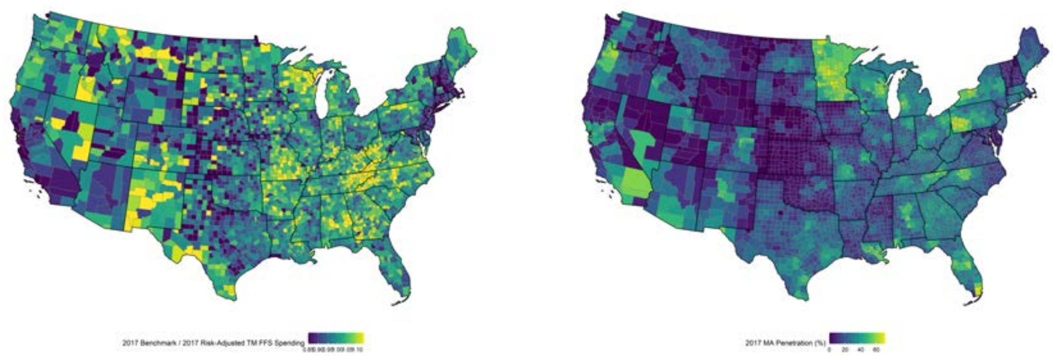

Notes: Data from CMS benchmark and enrollment files. The left map illustrates the ratio of the 2017 benchmark rate to the 2017 risk-adjusted TM spending in each county. To show detail, the data are windsorized at the 5th and 95th percentiles. The right map shows the county-level MA penetration rate, defined as the total number of people enrolled in any MA plan divided by the number of Medicare beneficiaries.

3.1 Medicare Advantage plans

We collect data on all plans offered from 2008 to 2017 from public CMS files. For each plan, we collect county-level enrollment, premiums, the Part B premium reduction, in-network copayment rates for primary care visits and 7-day hospital stays, the star rating, and indicators for basic and expanded drug coverage (as defined by CMS), and dental, vision, and hearing coverage of any type.[9] We also collect benchmark rates. We do not observe the bids directly. Rather, we observe plan-level risk-adjusted payments which, when combined with the above data and Equation (2), allow us to uniquely identify a bid for each plan-county. We combine the enrollment counts with CMS eligibility data to form product shares at the plan-county-year level. Finally, CMS releases detailed costs estimates submitted by firms during the bid process after a five year delay. We obtain these costs for all plans from 2008-2015. While this cost information is extra-ordinarily detailed, we focus on plans' reported risk-adjusted cost of providing TM-equivalent coverage.

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 restrictions, we drop plan-county observations with ten

[9] 2% of plans use coinsurance, which we convert to copayments using the Medicare Physician Fee Schedule and the American Hospital Association Annual Survey.

or fewer enrollees. For consistency, we drop plans outside our micro-data sample area.

Table 1 presents the mean characteristics of our 64,542 plan-county-year observations separated by benchmark quartiles calculated at the market (county-year) level. In the cross section, as benchmarks increase, observable plan benefits generally improve. The fraction of plans offered with zero premium increases from .336 in the first quartile to .486 in the fourth. However, these patterns are not always monotonic: the average deductible increases from \$63.08 in the first quartile to \$99.55 in the third quartile before decreasing to \$87.76 in the fourth quartile. These patterns reflect the fact that benchmarks are set as a function of average TM costs in previous years. While the costs faced by private insurers are surely correlated with average TM costs, there are likely meaningful cost differences which, when combined with heterogeneous demand responses, implies that the benchmark alone is an insufficient statistic for understanding the benefit generosity behavior of firms.[10] In fact, the wedge between the benchmark and mean reported plan TM costs increases in the benchmark suggesting plans' costs differ meaningfully from the government's.

3.2 Medicare beneficiaries

We access data on individual Medicare beneficiaries from the 2008-2017 Medicare Current Beneficiary Survey (MCBS), a rolling-panel survey produced by CMS and Westat. Participants are interviewed repeatedly over three years, and responses are linked to CMS data to ensure accuracy. We observe demographics including income, age, sex, race, education, and county of residence. Respondents self-report their health status, choosing from Excellent, Very Good, Good, Fair, and Poor. We also observe MA plan choices. In some years, the MCBS does not report the plan choice directly and instead reports the insurer choice, along with information about plan premiums and features which we match to plan data.

[10]There are many potential explanations for lack of a monotonic relationship between the benchmark and benefit provision. For example, a change in the benchmark rate could be a signal of a change in the risk distribution of consumers in the market, which could lead plans to try to cream-skim by changing their product characteristics (Decarolis and Guglielmo, 2017). The menu of product features may not be fully salient to consumers (Curto et al., 2021), thus increasing benefits may increase plan costs but yield small increases in enrollment. Our model can account for all of these possibilities (at least to some degree) with the exception of adverse selection. However, as mentioned above, the evidence suggests that the current implementation of the risk-adjustment system effectively reduces incentives for plans to cream-skim (Newhouse et al., 2015).

Table 1: Mean plan characteristics by market-level benchmark quartile

Variable	Benchmark quartile 1st	2st	3rd	4th
Annual premium ($)	636	657	562	501
Enrollment-weighted premium ($)	474	493	444	336
Fraction of plans with zero premium	.336	.361	.419	.486
Annual Part B premium reduction ($)	9.62	14.64	35.42	42.77
Deductible ($)	63.08	91.53	99.55	87.76
Star rating	3.11	3.19	2.45	2.24
Copays				
Primary care ($)	12.90	12.60	14.49	13.99
Hospital stay ($)	1,351	1,288	1,118	1,003
Supplemental coverage indicators				
Basic prescription drug	.791	.691	.771	.766
Enhanced prescription drug	.604	.829	.687	.680
Dental	.686	.698	.586	.639
Vision	.929	.919	.902	.891
Hearing	.620	.684	.718	.744
Annual benchmark ($)	9,469	10,121	10,724	11,939
Plan TM cost ($)	9,198	9,682	9,975	10,582
Enrollment	458	813	746	837
Market-level MA share	.206	.220	.236	.239
Observations	10,750	13,331	17,569	22,892

Notes: An observation is a plan-county-year. Reported figures are unweighted means unless noted. Quartiles are defined at the market (county-year) level. All costs are in 2017 dollars. The annual premium as defined here is the supplemental MA premium – all TM and MA enrollees pay the Part B premium. The star rating ranges from zero to five. Prescription drug coverage indicators are additive. “Plan TM cost” is the plans’ costs of covering required services as disclosed during the bidding process and is limited to 2008-2015.

The MCBS samples Medicare beneficiaries using a multi-level clustered procedure. While we do not observe beneficiaries in every county, within each geography included in the data there is considerable variation in demographics and plan enrollment. The MCBS provides sampling weights which we use to transform our results into a nationally representative form.

The 'standard' set of Medicare beneficiaries studied in the literature includes age 65-plus retirees without outside insurance (Curto et al., 2021). However, the MCBS and CMS data include others who are eligible to purchase MA plans including those with employer-provided insurance, those whose original Medicare eligibility was not age-related, those with ESRD, and those who are not full-year Part A/B enrollees. As these individuals purchase MA plans, we cannot exclude them without violating our assumption that the MCBS draws from CMS enrollment files. We instead create 'administrative' indicator variables. We exclude any individuals who were also eligible for Medicaid during the year and those with missing address information. After applying these exclusions, the sum of the MCBS sample weights differs from the total MA-eligible population in the CMS data by less than 2%.[11]

Medicare beneficiaries have access to non-MA insurance 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 costs 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 plan designs are standardized by CMS and indexed by letters. 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 TM enrollees are responsible for and is the most popular Medigap plan.[12]

Summary statistics on our 78,812 individual-year observations covering 3,851 county-year markets and 42,261 unique individuals are reported in Table 2. 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. Over 75% self-report "Good" or better health. 25% report having college degrees and 16% did not graduate high school. 25% receive employer-

[11] According to our CMS data, in 2017 the total number of Medicare beneficiaries not also eligible for Medicaid was 42.7 million. The total MCBS weight for 2017 is 43.4 million.

[12] 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 individuals in New York, we averaged the Plan C rates offered by all other insurers.

Table 2: Medicare beneficiary micro-data summary statistics

Variable	All observations		By MA enrollment		TM → MA switch	
	Mean	Std. dev.	MA	TM	Yes	No
MA enrollment indicator	.280	.449	1	0	1	0
Income ($)	52,684	76,202	43,117	56,415	43,858	53,172
Age	73.3	9.84	73.5	73.2	72.8	74.8
Medigap price ($)	2,722	674	2,810	2,687	2,654	2,742
Demographic indicators						
Female	.536	.499	.548	.531	.526	.536
Black	.081	.273	.095	.075	.094	.069
Hispanic	.010	.101	.017	.008	.018	.006
Education indicators						
Bachelor's degree or higher	.250	.433	.188	.275	.216	.260
Attended college	.307	.461	.315	.304	.294	.296
Graduated high school	.285	.451	.305	.277	.299	.288
Health status indicators						
Excellent	.177	.381	.175	.177	.174	.165
Very Good	.313	.464	.315	.312	.319	.313
Good	.304	.460	.307	.303	.319	.314
Fair	.151	.358	.155	.149	.149	.153
Poor	.056	.230	.049	.059	.049	.055
Administrative indicators						
Employer-provided insurance	.254	.435	.008	.350	.000	.352
Non-aged eligibility	.144	.351	.147	.143	.193	.141
ESRD	.007	.081	.004	.008	.002	.007
Full-year Part A/B enrollee	.905	.293	.977	.877	.962	.901
Observations	78,812		22,108	56,704	1,345	25,958

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 Medigap price are in 2017 dollars. The Medigap price is the United Healthcare premium for Medigap Plan C (see text for details). Demographic categories are defined by CMS administrative data. Education indicators are mutually exclusive. 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 conditional on observing past-year TM enrollment.

sponsored insurance, and 14% are Medicare-eligible for non age-related reasons. 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.

The third set of columns of Table 2 illustrates the panel nature of our data and focuses on panel observations for which the individual was enrolled in TM in the previous year – 27,297 observations total. We split the data into those who switched from TM to MA, and those who remained. Those who switched are generally similar to the larger group of MA enrollees, though switchers are slightly healthier on average.

Finally, we supplement these data with market-level average demographics from the Area Health Resource File published by the Health Resources and Services Administration. For each market, we collect the median household income, the percent of those 65-and-older in deep poverty, the unemployment rate, the population density, and the number of doctors, Medicare-certified hospitals, skilled nursing facilities, and hospice facilities. Summary statistics for our markets in 2017 by benchmark quartile are reported in Table E.2.

4 Demand

We model the demand for MA plans by extending the discrete choice demand setting of Goolsbee and Petrin (2004) to allow for switching costs between Traditional Medicare to Medicare Advantage and for the possibility of switching costs between MA plans. Agents in the model consist of consumers i and insurers/firms f which exist in markets (counties) m. Consumers are described by a vector of demographic characteristics observable to the econometrician z_i and unobservable characteristics $\nu_i, \{\epsilon_{ij}\}$. Each insurer offers plans $j \in \mathcal{J}_f$. Plans are described by a premium p_j and a vector of characteristics x_j that includes supplemental benefits and cost-sharing rules, and a characteristic ξ_j which is observed by consumers but not by the econometrician. Insurers have time- and market-invariant vertical quality v_f. ξ_j therefore represents plan- and time-specific deviations from that quality.

Let y_i be consumer i's income and h_i be a vector of indicators corresponding to i's health status. Consumers enter the period enrolled in plan k_i. We define three 'switching cost' indicators S_{sij}. Let S_{1ij} equal one if k_i is the outside good – we call this the *Medicare-to-MA* indicator. Let S_{2ij} – the *MA Interfirm* indicator – be one if k_i is offered by a different firm than j. Finally, let S_{3ij} – the *MA Intrafirm* indicator – be one if k_i and j are different plans offered by the same insurer.

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

$$u_{ijt} = \left(\alpha_0 + \alpha_1 y_{it} + \alpha_2 y_{it}^2\right) p_{jt} + \sum_s \beta_s S_{sijt} + \sum_s \sum_h \beta_{sh} S_{sijt} h_{it} \\ + \beta_z z_{it} + \beta_x x_{jt} + \xi_{jt} + v_f + \beta_\nu \nu_{it} + \epsilon_{ijt}. \tag{3}$$

The α parameters capture income-varying premium sensitivity. β_s and β_{sh} capture health-dependent switching costs. β_z captures heterogeneous tastes for MA plans by demographics, and β_x captures mean tastes for plan characteristics x_{jt}. ν_{it} is an unobservable (to the econometrician) preference that consumer i has for MA which is assumed to be drawn independently from a standard normal distribution—β_ν controls the variance of this 'random coefficient.' We have explored specifications in which ν_i is fixed over time and found similar results. ϵ_{ijt} represents the idiosyncratic taste of consumer i for plan j which is assumed to be drawn independently from the Type-I extreme value distribution.

Consumers have access to an outside good, the price of which may vary with demographics $p_{0t}(z_i)$. The utility of the outside good is

$$u_{i0t} = (\beta_{00} + \beta_{01} y_{it} + \beta_{02} y_{it}^2) p_{0t}(z_{it}) + \epsilon_{i0t}. \tag{4}$$

We normalize by subtracting Equation (4) from each u_{ijt}.

We include switching costs due to the consistent finding of inertia in plan enrollment (Nosal, 2012, Aizawa and Kim, 2018).[13] Enrollees in MA face restrictive provider networks that vary across plans. 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

[13] 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, our estimation approach captures the inertia that is salient for our counterfactual analysis.

costless to re-enroll. Similar to Handel (2013), we model these costs directly in utility.

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

$$\delta_{jt} = \alpha_0 p_{jt} + \beta_x x_{jt} + \upsilon_f + \xi_{jt} \tag{5}$$

and an individual-specific deviation from that mean

$$\begin{aligned} \mu'_{ijt} = & \left(\alpha_0 + \alpha_1 y_{it} + \alpha_2 y_{it}^2\right) p_{jt} + \sum_s \beta_s S_{sijt} + \sum_s \sum_h \beta_{sh} S_{sijt} h_{it} + \beta_z z_{it} + \beta_\nu \nu_{it} \\ & - (\beta_{o0} + \beta_{o1} y_{it} + \beta_{o2} y_{it}^2) p_o(z_{it}) + \epsilon_{ijt}. \end{aligned} \tag{6}$$

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

$$s_{ijt} \equiv \Pr(i \text{ chooses } j) = \int_\nu \frac{\exp(\delta_{jt} + \mu_{ijt}(\nu))}{1 + \sum_{k \in \mathcal{J}_m} \exp(\delta_{kt} + \mu_{ikt}(\nu))} d\nu, \tag{7}$$

and the total share of plan j is

$$s_{jt} = \int_{z_i} s_{ijt}(z_i) dz_i. \tag{8}$$

We define consumer welfare in terms of compensating variation: the amount that, if the choice of MA plans was removed, consumer i would have to receive as income in order to achieve the same level of expected utility (Hicks, 1945, Diamond and McFadden, 1974, Nevo, 2000). Let $\alpha_{it} = \alpha_0 + \alpha_1 y_{it} + \alpha_2 y_{it}^2$. The expected consumer welfare for beneficiary i is

$$CS_{it} = E[\max_j u_{ijt}]/\alpha_{it} = \frac{1}{\alpha_{it}} \ln \left(1 + \sum_j \exp(\delta_{jt} + \mu_{ijt}) \right). \tag{9}$$

While the government seeks to maximize the sum of this compensating variation across all markets (as it does not observe ϵ_{ijt}), consumers only accrue welfare from MA if they enroll in an MA plan. We therefore calculate mean compensating variation *for MA enrollees* via

$$\overline{CS}^{cond} = \frac{\sum_i CS_i}{\sum_i s_i}. \tag{10}$$

Following the literature (see, e.g. Petrin, 2002, Town and Liu, 2003), we report the mean compensating variation both per Medicare beneficiary and per MA enrollee, as well as the aggregate consumer welfare $\sum_i CS_i$. This formulation of consumer surplus assumes that our parameterization of demand holds for inter- and infra-marginal consumers (McFadden, 1974). Though we do not observe switches for every consumer, we do observe switches by consumers in each demographic category. As μ_{ij} includes switching costs, our estimates of consumer surplus are net of those costs and in this sense are short run.

5 Supply

Our model of insurers largely follows the multiproduct firm approach in the literature (see e.g. Berry et al., 1995, Petrin, 2002). We incorporate two wrinkles driven by our setting. First, firms choose prices and product characteristics simultaneously. Second, as detailed in Section 2, firms submit a 'bid' b_{jt} to CMS for each plan they offer, which maps into revenue from the government through subsidies and (potentially) from consumers through premiums as a function of the plan's characteristics. In Section A we show that under certain assumptions the CMS rules imply that the mapping is unique and thus we can write the firm's problem in the traditional way in terms prices and product characteristics with the addition of a subsidy that depends on characteristics.

Let x_{jt} and ξ_{jt} be the product characteristics as defined above, and let p_{-jt}, x_{-jt} and ξ_{-jt} be the set of prices and product characteristics for all plans other than j. Let $c_{jt}(x, \xi)$ be the per-enrollee marginal cost incurred by the firm. Let $sub(x_{jt}; B_{jt}, \lambda_{ft})$ be the function that maps product characteristics and the benchmark B_{jt} $(= B_{mt}\phi_{ft})$ to the subsidy received by the firm where λ_{ft} is the firm's rebate percentage and ϕ_{ft} is the firm's benchmark bonus, taken to be exogenous.[14] Let N_m be the number of Medicare beneficiaries in market m. Plan-level profit is

$$\pi_{jt}(p_{jt}, x_{jt}, \xi_{jt}; \cdot) = (p_{jt} + sub(x_{jt}; B_{jt}, \lambda_{ft}) - c_{jt}(x_{jt}, \xi_{jt}))N_m s(p_{jt}, x_{jt}, \xi_{jt}; p_{-jt}, x_{-jt}, \xi_{-jt}) \quad (11)$$

[14] CMS uses past values of performance measures (two years before the plan year) to calculate the star rating (and thus λ and ϕ), and changes the characteristics used from year to year. Insurers therefore likely find it difficult to manipulate specific characteristics to obtain higher rebates. We thank an anonymous referee for clarifying this point.

and the firm's profit is

$$\max_{\{p_{jt}, x_{jt}, \xi_{jt}\}_{j\in\mathcal{J}_f}} \sum_{j\in\mathcal{J}_f} \pi_{jt}(p_{jt}, x_{jt}, \xi_{jt}; \cdot). \tag{12}$$

Before turning to the cost model we note that we do not model plan entry and exit due to the computational requirements and the fact that there is less than 1% market share change due to entry or exit in any particular year.

5.1 Costs

We assume the marginal cost of offering a plan is constant with respect to the number of enrollees, and further assume that risk-adjustment is effective, so that c_{jt} does not vary with enrollees' health status. We model the (log) marginal cost function as

$$\ln(c_{jt}) = \gamma_f + \gamma_r + \gamma_m + \gamma_t + \gamma_{f,x} x_{jt} + \gamma_{f,\xi}\xi + \omega_{jt}, \tag{13}$$

where γ_f is a firm-specific cost, γ_r is a star-rating-specific cost, γ_m is a market-specific cost, and γ_t is a time-varying cost. These parameters are fixed effects to be estimated. $\{\gamma_{f,x}, \gamma_{f,\xi}\}$ are the firm-varying costs of providing x and ξ, respectively. ω_{jt} is an unobservable (to the econometrician) plan-level cost.

The solution to Equation (12) is partially characterized by the first-order conditions

$$\{p_{jt}\}: \quad 0 = s_{jt} + \sum_{k\in J_f} (p_{kt} + sub(x_{kt}; \cdot) - c_{kt}) \frac{\partial s_{kt}}{\partial p_{jt}}. \tag{14}$$

Following Berry et al. (1995) we define a $J \times J$ matrix Δ_t whose (j,k) elements are given by

$$\Delta_{jkt} = \begin{cases} -\frac{\partial s_{kt}}{\partial p_{jt}}, & \text{if } \{j,k\} \in \mathcal{J}_f \\ 0, & \text{otherwise.} \end{cases}$$

The first-order conditions can then be solved in vector form to obtain costs:

$$\mathbf{c}_t = \mathbf{p}_t + \mathbf{sub}_t - \Delta_t^{-1} * \mathbf{s}_t. \tag{15}$$

5.2 Government spending

Evaluating candidate solutions to Equation (1) requires calculating the total government expenditure on the Medicare program, which consists of the sum of the MA payments given by $sub(\cdot)$ and spending on TM. Let TM_{mt} be the average risk-adjusted TM spending in the market. As we do not observe individual-specific risk scores, we calculate MA and TM spending using the average risk level in the market. Thus, dropping time subscripts,

$$GovExp_i(B_m) = \sum_j s_{ij} sub_j(B_m; \cdot) Risk_m + \left(1 - \sum_j s_{ij}\right) TM_m Risk_m, \tag{16}$$

where TM_m is the average per-enrollee TM spending in the market. In other words, as consumers are attracted to MA or pushed back into TM, $GovExp$ includes their costs across both sectors. As we do not observe individual-specific risk scores, we calculate MA and TM spending using the average risk level in the market—in other words, we set $Risk_{it} = Risk_{mt}$ in Equation (2) for all i. We treat TM_m as exogenous due to the risk adjustment system—i.e. we assume MA enrollment does not change within-county risk-adjusted TM spending, though we test this assumption after estimating the parameters of firms' marginal cost function. However, as benchmarks change and beneficiaries move between TM and MA in a county, the across-county average TM and MA risk scores and average costs change as well.

5.3 Existence of policy functions

To implement our counterfactual approach, described in Section 8, we estimate policy functions. We first establish that policy functions exist in this setting. For ease of exposition, we drop market and time subscripts and let X_f be the (finite) vector of choices made by firm f, and let Z_f be the (finite) vector of payoff-relevant information observed by the firm. The benchmark is B. Let X_{-f} be a vector capturing the actions of firm f's competitors. We rewrite Equation (12) as

$$\max_{X_f} \pi_f(X_f; X_{-f}, Z_f, B), \tag{17}$$

and define equilibrium as a vector $X^* = \{X_f^*\}_f$ that satisfies

$$\frac{\partial \pi_f}{\partial X_{fl}}(X_f^*; X_{-f}^*, Z_i, B) = 0 \text{ for all } f, l \quad \text{and} \quad H_f(X_f^*) \text{ negative definite for all } f \tag{18}$$

where X_{fl} is the l-th element of X_f and $H_f(X_f^*)$ is the Hessian of π_f evaluated at X^*. These conditions imply that $\frac{\partial^2 \pi_f}{\partial B \partial X_f}$ is symmetric and $\left(\frac{\partial^2 \pi_f}{\partial X_f \partial X_f'}\right)^{-1}$ exists in some neighborhood around any equilibrium X^* for all f. By the implicit function theorem,

$$\frac{\partial^2 \pi_f}{\partial X_f \partial X_f'} \cdot \frac{\partial X_f}{\partial B} + \frac{\partial^2 \pi_f}{\partial B \partial X_f} = 0$$
$$\Rightarrow \frac{\partial X_f}{\partial B} = -\left(\frac{\partial^2 \pi_f}{\partial X_f \partial X_f'}\right)^{-1} \left(\frac{\partial^2 \pi_f}{\partial B \partial X_f}\right).$$

Thus, in some neighborhood around X^*, there is a one-to-one relationship from B to X_f for all f—in other words, given a set of Z_f, the policy correspondences $X_f(B, Z_f)$ exist. As multiple equilibria are possible, we adopt the Equilibrium Selection (ES) assumption of Bajari et al. (2007).

Assumption ES—Equilibrium Selection: The data are generated by a single equilibrium strategy profile $\mathbf{X}(B; \cdot)$.

This assumption ensures that $X_f(B, Z_f)$ is a function.

6 Estimation

We estimate the parameters of the demand system following the two-stage approach of Goolsbee and Petrin (2004). First, we estimate parameters which capture individual-level variation in MA preferences—those parameters that define μ'_{ij}— via maximum likelihood and recover mean plan-level utilities δ_j. We then estimate the parameters of Equation (5) with an instrumental variables approach.

Let $\theta_I = \{\alpha_1, \alpha_2, \beta_s, \beta_{sh}, \beta_z, \beta_\nu, \beta_0\}$ be the set of parameters which determine μ'_{ij}. For candidate value $\tilde{\theta}_I$ we use the Berry (1994) inversion with the Berry et al. (1995) contraction mapping to compute the unique set of product mean utilities $\delta_j(\tilde{\theta}_I)$ that match predicted shares to the aggregate county-level market shares observed in the CMS data. Let C_{ij} be an indicator variable that is equal to one if person i chose product j. The likelihood function is

$$L_{it}(C_{ijt}; \tilde{\theta}_I, \delta(\tilde{\theta})_I) = \prod_j s_{ijt}^{C_{ijt}}, \tag{19}$$

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

$$l(C; \tilde{\theta}) = \sum_i ln(L_{it})w_{it}. \tag{20}$$

In practice, we form an empirical analog of the share function by numerically integrating over draws of the ν_{it} distribution. The rolling panel design of the MCBS implies that a fraction of our observations have no past enrollment data with which to form the S_{sijt} variables. We solve this problem by drawing from from the distribution implied by shares of the plans offered in the previous period. Finally, in 2017 CMS updated the weighting methodology to ensure that the MCBS matched average MA enrollment. For consistency across years we reweight the pre-2017 data to match current methodology. At the estimate $\hat{\theta}_I$ we store the unique $\hat{\delta}_j(\hat{\theta}_I)$ and regress it on observable product characteristics to estimate the parameters of Equation (5).

Just as ξ_{jt} is likely to be correlated with p_{jt}, it is also likely to be correlated with firm costs, and thus if we were to try to estimate the parameters of Equation (13) without accounting for ξ_{jt}, our estimates would be biased. After estimating the demand parameters, we calculate $\hat{\xi}_{jt}$ for each plan and assume that $\omega_{c,jt}$ and are uncorrelated with $\hat{\xi}_{jt}$ and our observables.

Finally, we estimate policy functions with a first-order approximation. That is, for product j and characteristic x_l, we write

$$x_{ljt} = \beta_{f,l} \times B_{jt} + \beta_{f,z} \times \tilde{Z}_{ft} + \epsilon_{ljt}, \tag{21}$$

where $\beta_{f,l}$ is the firm-level first-order approximation of the effect of the change in the benchmark, $\tilde{Z}_{ft}$ is a finite approximation of the information vector Z_{ft}, and ϵ_{ljt} captures measurement error, approximation error, and other factors that influence product characteristics such as plan-product-characteristic-level cost and demand expectation shocks.

6.1 Instruments

Since ξ_{jt} is chosen by firms and observed by consumers, but not observed by us, it is likely to be correlated with the plan premium and other product characteristics. To identify the α_0, β_x and β_y coefficients we must therefore find instruments for premiums and plan characteristics. First, we note that not all product characteristics are likely to be endogenous: for each plan, the basic drug coverage indicator remains constant over time and so it is plausibly exogenous. Furthermore, the star rating is set at the insurer level reflecting health outcomes with a two year lag and therefore is also plausibly exogenous.

We construct one instrument from the observation that our cost function includes a geographic component; costs are therefore correlated across plans in a given market. Our data includes detailed information on insurers' cost projections for TM-covered services submitted during the bidding process. These projections must be a) related to the plan's past realized costs, and b) certified by a professional actuary. For each plan, we calculate the average total cost of TM-covered services across competitors weighted by their conditional shares. This instrument is excluded from the demand system if competitors' TM-covered service costs c^{TM}_{-jt} are uncorrelated with ξ_{jt} i.e. $E[c^{TM}_{-jt}\xi_{jt}] = 0$. Since these costs are private information—these data are not released by CMS until five years after the plan year has concluded—it is not likely that firms choose ξ_{jt} based on the costs of particular competitors.

The discussion in Section 5.3 and the panel nature of our data points to additional potential instruments. First, we note that observable and unobservable plan characteristics as well as premiums are likely functions of the benchmark and hence correlated with each other thereby invalidating BLP-type instruments in our setting. However, if the costs of plan characteristic provision are correlated over time and if county benchmark updates are independent, using lagged values of the residuals from a regression of plan characteristics on benchmarks and other time-invariant state variables should be valid instruments. Intuitively, the residuals proxy for plan characteristic costs as the common impact of the benchmark will have been removed from the insurer's choice of plan characteristics. If updates to the benchmark are independent and shocks to ξ are uncorrelated with benefit provision costs, then ξ will be orthogonal to these lagged plan characteristic residuals.[15]

[15]More formally, for characteristic x_l and time-invariant 'state' variables a_j we can write $x_{ljt} =$

We examine the validity of these instruments first by testing the independence of benchmark updates. We estimate that a \$1 increase in $(B_{mt-1} - B_{mt-2})$ is associated with a \$0.03 decrease in $(B_{mt} - B_{mt-1})$ (t-statistic = 0.17 when clustering by year) and conclude benchmark updates are approximately independent. Next, we examine the section assumption of persistence of the plan characteristic residuals. We find that the correlation coefficient between the contemporaneous and lagged residuals ranges from 0.6659 (enhanced drug coverage) to 0.8754 (hospital copay). First-stage F-statistics testing the explanatory strength of instruments in accounting for plan characteristic variation range from 298 (annual premium) to 3,396 (hospital copay). Taken together, the evidence suggests that the necessary conditions for our lagged residual instrumenting strategy seem to hold.[16]

While these results suggest that benchmarks updates are exogenous from ξ, they nonetheless may be correlated with plan-product-characteristic-level *costs* of insurers making them endogenous in the policy functions. Therefore, we need instruments to consistently estimate Equation (21). We take advantage of the difference in the payment floors coming from county-level differences in urban/rural status and leverage the identification strategy of Duggan et al. (2016). These benchmark differences are driven by small population differences across counties that map into CMS's definition of urban and rural that are very likely orthogonal to plan characteristic costs. We obtain the Rural-Urban Continuum Code from the Area Health Resources File and instrument benchmarks with rural-urban category identifiers interacted with year fixed effects.[17]

7 Results

In this section we first describe our estimates of demand for and costs of MA plans. We then turn to a discussion of the firm-varying product characteristic policy functions.

$g_{f,l}(B_{mt}, a_j) + u_{l,jt}$ and $\xi_{jt} = g_{f,\xi}(B_{mt}, a_j) + u_{\xi,jt}$. Here, the u terms are mean-zero random variables. Define $o_{mt} \equiv B_{mt} - B_{mt-1}$. Suppose o_{mt} is a random variable distributed independently across time and with respect to u. Further suppose $E[u_{l,jt-1} u_{\xi,jt} | B_{mt}, a_j] = 0$; time-varying information relevant to the choice of x_l in period $t-1$ is not relevant to the choice of ξ in period t after conditioning on the benchmark and a_j. Under these assumptions, $E[u_{l,jt-1}\xi_{jt}] = 0$. We thus instrument for x_{ljt} with $\widehat{u_{l,jt-1}}$.

[16] This is similar to our policy function regression. The key difference is that the $\tilde{Z}_f$ of Equation (21) may include time-varying components such as demographics and competitors' past choices.

[17] We thank anonymous referees for suggesting this strategy.

7.1 Demand

We report maximum likelihood estimates of individual-specific parameters in Table 3. Higher income consumers are less price-sensitive than lower income consumers. 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. We interact the switching costs with indicators for self-reported health status, with 'Poor' as the excluded group. The point estimates indicate that healthier individuals face lower costs of switching, consistent with the provider-changing hypothesis, though the standard errors prevent us from making clear inferences between adjacent health statuses.

Our demographic estimates imply that younger consumers have a higher preference for MA than older consumers. Non-Whites have a stronger preference for MA plans, as do beneficiaries with lower levels of education. These results align with other findings that MA enrollment of Black and Hispanic beneficiaries has grown faster than enrollment of White beneficiaries (Meyers et al., 2021). Our administrative indicators enter with appropriate signs and reasonable magnitudes: those with employer-provided insurance or ESRD are extremely unlikely to choose an MA plan. Finally, our random coefficient enters significantly with a magnitude roughly equal to the inter-insurer switching cost suggesting that idiosyncratic preferences for MA are important.

Table 4 reports estimates of Equation (5). The first column presents OLS estimates assuming prices and characteristics are exogenous. The second column reports IV estimates when the premium is instrumented with our cost instrument. The third column reports the results when prices and product characteristics are both treated as endogenous and instrumented with our full set of instruments. Consistent with OLS estimates on price being biased towards zero, the IV premium coefficients are larger in magnitude than the OLS coefficient. Furthermore, in general the coefficients on product characteristics are larger in magnitude when they are treated as endogenous, though the estimates are noisier.

We focus our attention on specification (3). The parameter estimates in this specification are quite sensible. For the plans with a positive premium the average plan elasticity is -5.48.

Table 3: Maximum likelihood estimates of individual-specific preferences

Variable	Coeff.	Std. Err.	Variable	Coeff.	Std. Err.
Price × Income	0.1760	0.0343	MA × Demographics		
Price × Income^2	-0.0027	0.0015	Age	2.3135	0.1917
			Age^2	-0.1616	0.0131
TM-to-MA switch ×			Female indicator	-0.1304	0.0385
Constant	-5.0089	0.1093	Black indicator	0.5360	0.0755
Excellent health	0.4724	0.1015	Hispanic indicator	0.4605	0.1782
Very good health	0.3832	0.0968	Graduated high school	-0.1890	0.0624
Good health	0.2295	0.0948	Attended college	-0.3875	0.0624
Fair health	0.1787	0.0987	College degree or higher	-0.9327	0.0673
Inter-Insurer switch ×			Administrative indicators		
Constant	-2.0480	0.1097	Employer-provided insurance	-6.4027	0.1267
Excellent health	0.0619	0.1257	Non-aged eligibility	0.4759	0.071
Very good health	0.0784	0.1186	ESRD diagnosis	-2.0265	0.2456
Good health	-0.0007	0.1183	Full year enrollment	3.1896	0.0824
Fair health	-0.0294	0.1274			
			Outside good (Medigap) price ×		
Intra-Insurer switch ×			Linear	-0.4180	0.0632
Constant	-0.9308	0.1333	Income	0.2165	0.0135
Excellent	-0.0640	0.1533	Income^2	-0.0043	0.0005
Very good	-0.0150	0.1449			
Good	-0.0163	0.1445	Random Coefficient	2.0323	0.0424
Fair	-0.0549	0.1533			
Weighted Log Likelihood			-72,504		
Obs.			78,812		

Notes: An observation is an individual-year. MA and outside good prices are measured in thousands of 2017 dollars. Income is measured in hundreds of thousands of 2017 dollars. The omitted group for the switching cost interactions is 'Poor' health. Standard errors are robust to heteroskedasticity.

The average semi-elasticity of increasing premiums by $1 is -.062, similar to estimates from the literature. For example, using an earlier sample period, Aizawa and Kim (2018) estimate an average MA semi-elasticity of -.075. Combining these estimates with Table 3, the median consumer is willing to pay roughly $450 for prescription drug coverage, $290 for hearing coverage, and $220 for a reduction of $1,000 in the copay for a hospital stay. The cost incurred by an median-income individual switching from TM to MA is $877, roughly twice the mean annual premium in our data, while the same individual switching between plans within an MA insurer incurs a cost of only $178.

Table 4: Estimates of mean preferences for plan characteristics

Variable	(1)	(2)	(3)
Annual Premium (per $1000)	-0.4472	-5.4345	-5.8986
	(0.0214)	(0.4293)	(0.6277)
Part B reduction (per $1000)	0.1982	0.5414	1.4465
	(0.0729)	(0.0866)	(0.1892)
Deductible (per $1000)	-0.2115	-1.0752	-1.6567
	(0.0419)	(0.1011)	(0.2023)
Copays			
Primary care	-0.0223	-0.0316	-0.0808
	(0.0017)	(0.0028)	(0.0085)
Hospital stay (per $1000)	0.1429	-0.8961	-1.3120
	(0.0265)	(0.0980)	(0.2095)
Supplemental coverage indicators			
Basic prescription drug	0.6131	2.9099	2.6799
	(0.0448)	(0.2074)	(0.2834)
Enhanced prescription drug	0.3741	0.2801	1.5109
	(0.0407)	(0.0622)	(0.1705)
Dental	0.0840	0.8152	1.1568
	(0.0332)	(0.0777)	(0.1834)
Vision	-0.1181	0.1402	0.4814
	(0.0483)	(0.0718)	(0.2334)
Hearing	0.1510	0.9988	1.7161
	(0.0416)	(0.0952)	(0.1976)
Fixed effects			
Star rating	Yes	Yes	Yes
Firm	Yes	Yes	Yes
Endogenous variables			
Annual premium	No	Yes	Yes
Product characteristics	No	No	Yes
Mean implied elasticity (if < 0)	-0.3500	-5.0405	-5.4773
	(0.2393)	(3.4350)	(3.7325)
Mean ds_j/dp_j	-0.0041	-0.0575	-0.0624
	(0.0070)	(0.0961)	(0.1045)
Observations	50,439	50,439	36,073

Notes: An observation is a plan-county-year. In Column (2) we instrument for the annual premium using the share-weighted average of competitors' plans projected costs of TM-covered services. In Column (3) we add lagged residuals from regressions of the characteristics on the benchmark as instruments. See text for details. All dollar values are in 2017 dollars. Parentheses indicate robust standard errors in the top panel and standard deviations in the bottom panel.

7.2 Supply

Table 5 reports parameter estimates of Equation (13) – the marginal cost equation. Throughout, we focus insurer-varying coefficients on the five largest firms nationally and group smaller firms into a sixth category. Supplemental benefits and cost-sharing characteristics generally enter with the correct sign (with the exception of vision), though some estimates are noisy. The coefficient on the demand unobservable is positive and significant at the 1% level for all firms, suggesting that specifications which did not take the correlation between demand and cost unobservables into account would be biased.

To help interpret these parameters, it is useful to calculate the impact of changes in product characteristics. Given our estimates of marginal costs, a $1 increase in the primary care copay would decrease marginal costs by an average of $8.50, implying that the MA population visits doctors an average of 8.50 times per year. This is roughly in line with the Centers for Disease Control and Prevention estimate of 4.98 visits per year per individual age 65-or-older in 2016 (Ashman et al., 2019). Similarly, a $1 increase in the hospital visit copay would decrease marginal costs by $0.194, which is comparable to the Kaiser Family Foundation estimate of 0.252 hospital visits per year per TM enrollee in 2015. This finding, when combined with our estimate of the consumer valuation of hospital copays, is consistent with previous findings of behavioral hazard in the use of care (Loewenstein et al., 2013, Baicker et al., 2015). Adding basic drug coverage to plans without drug coverage costs an average of $407, whereas adding enhanced drug coverage costs an additional $482. Dental coverage costs $192.

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. 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. (2021) and compare the share-weighted estimated MA cost for zero-premium plans to TM costs at the county level. On average, we estimate MA costs to be $823 per enrollee-month, 3.2% less than TM costs of $851 per enrollee-month. The two cost measures are positively correlated with a coefficient of .561. After adjusting for inflation, the comparable average MA cost reported by Curto et al. is $830 per enrollee-month.

Finally, our approach assumes that the MA risk-adjustment system is effective and that marginal costs do not vary by the realized risks of the enrollees. If this assumption was violated, and firms faced higher (lower) marginal costs for higher-risk enrollees *after risk adjustment*, our estimated marginal costs would be biased upward (downward), which could influence our counterfactual calculations in Section 8. We test for this in two ways. First, we estimate the relationship between the bid and the realized risk of the plan. After accounting for firm fixed effects and aggregating to plan-year observations as realized risk is reported at that level, we estimate that a 1% increase in bids increases risk by 0.0047% (t-statistic = 0.77). Second, we re-estimate the regression of Table 5 with the addition of the plan's realized risk as a covariate. We estimate that a 1% increase in average enrollee risk increases marginal costs by 0.014% (t-statistic = 3.91). While these results suggest that the risk-adjustment system may slightly under-compensate plans with higher-risk enrollees, selection with respect to the benchmark is likely second-order.

7.3 Policy functions

To estimate the parameters of Equation (21) for each product characteristic, we must first define $\tilde{Z}_{ft}$, the approximation of the information vector Z_{ft}. The profit function (11) suggests that demographics and cost shocks likely influence product characteristics. We therefore include market-level average demographics, including the fraction of those 65-and-older who are White, Black, and Hispanic, the fraction of those 65-and-older in deep poverty, the median household income, the unemployment rate, and the population density. We also include lagged market-level averages of all product characteristics in each regression.[18]

The results are reported in Table 6. As no plan in our panel changed their basic drug coverage indicator, we do not consider changes to that indicator in our counterfactual. In

[18] We treat enhanced drug and DVH coverage indicators as continuous variables throughout this exercise for simplicity and consistency with our other characteristics. See Section B for details. Our results are robust to logit and probit specifications for these variables.

Table 5: Marginal cost parameter estimates

	$\ln(c_j)$					
Variable	Aetna	BCBS	Humana	Kaiser	UHG	Other
Cost-sharing characteristics						
Part B reduction (per \$1000)	-0.2740	-0.0302	0.0062	0.0000	0.0000	0.0002
	(0.0195)	(0.0169)	(0.0040)	(0.0000)	(0.0000)	(0.0015)
Deductible (per \$1000)	-0.0168	-0.0223	-0.0182	0.0000	-0.0037	-0.0283
	(0.0029)	(0.0021)	(0.0016)	(0.0000)	(0.0035)	(0.0011)
Primary care copay	-0.0012	-0.0011	-0.0011	0.0011	-0.0001	-0.0007
	(0.0001)	(0.0001)	(0.0001)	(0.0002)	(0.0001)	(0.0000)
Hospital copay (per \$1000)	-0.0245	-0.0146	-0.0060	-0.0297	-0.0169	-0.0225
	(0.0015)	(0.0014)	(0.0011)	(0.0030)	(0.0014)	(0.0006)
Supplemental coverage characteristics						
Basic prescription drug	0.0457	0.0288	0.0716	N/A	0.0005	0.0304
	(0.0033)	(0.0020)	(0.0016)		(0.0024)	(0.0011)
Enhanced prescription drug	0.0543	0.0367	0.0643	0.0337	0.0185	0.0411
	(0.0026)	(0.0016)	(0.0011)	(0.0011)	(0.0017)	(0.0007)
Dental	0.0146	0.0197	0.0032	0.0434	0.0159	0.0207
	(0.0019)	(0.0013)	(0.0012)	(0.0058)	(0.0014)	(0.0009)
Vision	-0.0035	0.0282	0.0191	0.0244	-0.0285	-0.0050
	(0.0184)	(0.0022)	(0.0013)	(0.0076)	(0.0101)	(0.0016)
Hearing	0.0376	0.0357	0.0177	0.0183	0.0018	0.0281
	(0.0073)	(0.0019)	(0.0011)	(0.0025)	(0.0047)	(0.0011)
Demand unobservable (ξ_j)	0.0160	0.0142	0.0110	0.0109	0.0086	0.0120
	(0.0002)	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0001)
Fixed effects	Star rating, firm, county, year					
Observations	64,542					
R^2	0.8675					

Notes: Observations are county-year-plans. Estimates are formed via OLS. All Kaiser plans in our sample had basic prescription drug coverage. Robust standard errors are in parentheses.

general, the signs of the coefficients line up with the prior that plans should improve benefits when the benchmark increases; the net effect of an increase in the benchmark for each plan is an increase in mean utility.[19]

In our context, Assumption ES may be restrictive as multiple equilibria are possible. In Appendix Table E.1, we explore the possibility of multiple equilibria by re-estimating the policy functions for each Census Region. While region-within-firm point estimates differ, the confidence intervals generally overlap. We conclude that the existence of isolated markets with widely disparate equilibrium behavior is unlikely.[20]

[19]Though we present a linear specification, we have explored non-linear functions of B_j, first-differenced specifications, and machine learning techniques and did not obtain meaningful improvements over this specification.

[20]We repeated this exercise splitting markets by the level of the benchmark and found similar results.

Table 6: Policy function estimation results

	(1) Part B reduction	(2) Deduct-ible	(3) Prim. care copay	(4) Hospital copay	(5) Enhanced drug	(6) Dental	(7) Vision	(8) Hearing	(9) ξ
Benchmark ($000s) ×									
Aetna	0.0047	0.0364	-1.3122	-0.1846	0.1607	0.0688	-0.0174	0.0817	0.6632
	(0.0028)	(0.0068)	(0.1831)	(0.0124)	(0.0109)	(0.0102)	(0.0049)	(0.0095)	(0.0969)
BCBS	0.0052	0.0383	-1.2529	-0.2022	0.1514	0.0622	-0.0286	0.0564	0.6657
	(0.0028)	(0.0069)	(0.1837)	(0.0125)	(0.0109)	(0.0102)	(0.0049)	(0.0095)	(0.0973)
Humana	0.0050	0.0405	-1.5049	-0.1817	0.1438	0.0753	-0.0327	0.0175	0.6570
	(0.0029)	(0.0069)	(0.1839)	(0.0125)	(0.0110)	(0.0103)	(0.0050)	(0.0096)	(0.0973)
Kaiser	0.0043	0.0304	-0.6334	-0.1492	0.1756	0.0885	-0.0254	0.0317	0.6026
	(0.0028)	(0.0068)	(0.1819)	(0.0123)	(0.0108)	(0.0101)	(0.0049)	(0.0095)	(0.0962)
UHG	0.0033	0.0292	-1.4517	-0.1706	0.1445	0.0316	-0.0187	0.0809	0.6553
	(0.0028)	(0.0069)	(0.1836)	(0.0125)	(0.0110)	(0.0103)	(0.0049)	(0.0096)	(0.0973)
Other	0.0087	0.0336	-1.6246	-0.2119	0.1494	0.0591	-0.0214	0.0645	0.6665
	(0.0028)	(0.0069)	(0.1832)	(0.0125)	(0.0109)	(0.0102)	(0.0050)	(0.0095)	(0.0974)
Obs.	37,974	37,974	37,974	37,974	37,974	37,974	37,974	37,974	37,974
R^2	0.1371	0.0758	0.2429	0.2826	0.0638	0.1457	0.1228	0.3140	0.0920

Notes: Observations are county-year-plans. The dependent variable in each regression is the product characteristic. The independent variables include the benchmark interacted with firm indicators, lagged values of market average product characteristics and market average demographics. We omit other covariates for space. We instrument for the benchmark with the Census Rural/Urban Continuum category interacted with year fixed effects. All dollars are 2017 dollars. No plan in our sample changed basic drug coverage. Robust standard errors are in parentheses.

8 Optimal Geographic Variation in MA Subsidies

We now turn to the problem of Equation (1): setting benchmark rates to maximize aggregate consumer welfare keeping government expenditures constant, where consumer welfare is given by Equation (9) and government expenditures are given by Equation (16). Evaluating candidate solutions to Equation (1) requires finding equilibrium vectors (p, x, ξ) for counterfactual benchmarks. A traditional approach involves searching for equilibria through analyzing agents' first-order conditions. The products in this setting are numerous and detailed; solving for a single equilibrium in one market this way requires a fixed point search over potentially hundreds of first order conditions, and solving Equation (1) requires searching over many equilibria in each of many markets.

We solve this challenge by using our estimated policy functions to reduce the dimension of the problem. In particular, for any B_m, we 1) use the estimated policy functions to update x and ξ; 2) use the estimated γ parameters to update c; and 3) solve for equilibrium prices. For convenience, we focus on the results of this process here and leave the technical details to Section B. We test the performance of our approach with Monte Carlo exercises in Section C and explore examples of the counterfactual outcomes in Section D.

We begin in Section 8.1 by describing the consumer-welfare maximizing policy for 2017.

In Section 8.2 we detail the contribution of each product characteristic to changes in welfare, and in Section 8.3 we explore the selection of markets for increased or decreased benchmarks based on the derivatives of the CS and $GovExp$ functions. In Section 8.4, we consider the performance of a linear policy rule. We explore alternative social welfare functions in Section 8.5 and summarize other results and robustness checks in Section 8.6.

8.1 The consumer welfare maximizing policy

Table 7 summarizes the results of solving Equation (1). The first column reports consumer welfare under the 2017 policy across a number of dimensions. The mean compensating variation for MA enrollees is \$464.68. The MA program increases aggregate consumer welfare by \$5.70 billion per year. The second column reports outcomes under the optimal policy and the third column calculates percentage changes between the two policies. The optimal policy increases mean compensating variation for enrollees to \$593.97 and aggregate consumer welfare to \$10.26 billion. The aggregate increase is driven more by an increase in the total share of MA (from 29.2% to 41.1%) than by the increase in compensating variation for enrollees; the optimal policy thus obtains the majority of its gains by moving individuals from TM to MA, as opposed to making MA more valuable for current enrollees.

The second panel of Table 7 splits markets by the direction of the benchmark change relative to the 2017 policy. The optimal policy increases the benchmark in 300 out of the 439 markets in our data, and in those markets the mean compensating variation for MA enrollees increases from \$429.37 to \$588.03, while in the 155 markets where the benchmark is lower, the mean compensating variation decreases from \$465.06 to \$431.44.

The third panel examines the changes by race and income as previous work have identified inequalities in the Medicare and MA systems (Ayanian et al., 2014, Li et al., 2017). While all groups gain on average, compensating variation increases more in percentage terms for White and Hispanic enrollees than Black enrollees. High income enrollees benefit more in percentage terms than low or medium income enrollees. We note that this result does not imply that *uniform* changes in MA benchmarks would disproportionately affect any particular group.

The benchmark changes are detailed in Figure 2. The left-hand histogram illustrates

Table 7: **Comparing the 2017 policy and the optimal policy**

	2017 Policy	Optimal Policy	Pct. Change
Mean compensating variation per Medicare beneficiary ($)	135.71	244.20	79.9
Total MA share (%)	29.2	41.1	40.8
Mean compensating variation per MA enrollee ($)	464.68	593.97	27.8
Aggregate consumer welfare generated by MA ($ billion)	5.70	10.26	79.9
Mean compensating variation for MA enrollees in...			
300 markets with benchmark increases ($)	429.37	588.03	37.0
139 markets with benchmark decreases ($)	465.06	431.44	-7.23
Mean compensating variation for MA enrollees who are...			
White ($)	436.53	563.17	29.0
Black ($)	462.73	576.58	24.6
Hispanic ($)	542.71	705.01	29.9
Low income ($)	458.28	585.35	27.7
Medium income ($)	440.48	559.42	27.0
High income ($)	407.72	553.77	35.8

Notes: This table reports the results of solving Equation (1). Compensating variation is calculated via Equation (9) and is reported in 2017 dollars per year. Compensating variation conditional on enrollment is calculated via Equation (10). White, Black, and Hispanic groups are defined by CMS. "Low income" refers to the lowest quartile of income, "high income" refers to the highest quartile, and "middle income" refers to the second and third quartiles. All calculations are weighted by MCBS sample weights.

the distribution of benchmarks under the optimal policy, and can be compared to Appendix Figure E.1. The right-hand histogram shows the distribution of changes in the benchmark. The changes are generally modest – the interquartile range of the difference between the optimal and 2017 benchmarks is from -$6 to $659. The mean change is $267 and the median is $151. In percentage terms, roughly 90% of the changes in the benchmark are of less than 10% of the 2017 benchmarks; in 379 markets the absolute change is less than $1,000.

The magnitude of the changes in the mean surplus across the benchmark increase/decrease split suggests that the aggregate changes in these markets are also substantial. Table 8 reports the aggregate consumer welfare and government spending under the 2017 policy and the optimal policy, split by the direction of the benchmark change. The aggregate consumer welfare generated by MA increases in markets with benchmark increases from $3.82 billion per year to $9.17 billion per year. This change comes with a decrease in spending on TM of

Figure 2: The distribution of the optimal benchmarks across markets

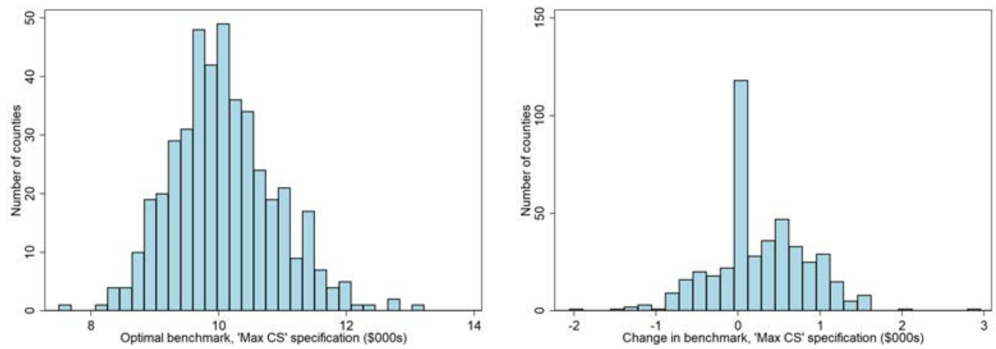

Notes: These graphs illustrate the solution of Equation (1). The left-hand graph shows the distribution of benchmarks under the optimal policy and the right-hand graph illustrates the distribution of the change in benchmarks from the 2017 policy to the optimal policy.

$74.4 billion, and an increase in MA spending of $75.8 billion.[21] Aggregate consumer welfare decreases in the other markets from $1.88 billion to $1.09 billion, while spending transfers from MA to TM to balance the government's budget constraint.

8.2 Changes in product characteristics

A key element of our contribution is the ability to incorporate changes in product characteristics into our evaluation of different benchmark schedules. A natural question is the extent to which changes in product characteristics contribute to consumer surplus. For any set of benchmarks $\{B_m\}$, our approach calculates a set of premiums and product characteristics for each market, X_m. Dropping the market subscripts, let $X(t)$ be a line through product space with $X(0)$ reflecting outcomes in the data and $X(1)$ reflecting outcomes at the optimal policy. The total gains to consumers by moving from current policy to the optimal policy can be written as $CS(X(1)) - CS(X(0))$. To understand how the gains are realized across different product characteristics, we decompose the overall gain using the gradient theorem.

[21] This transfer of resources from TM to MA may generate concerns about externalities with respect to the government's bargaining power (see e.g. Lakdawalla and Yin, 2015). However, as the government largely sets Medicare reimbursement rates nationally with local cost-of-living adjustments, these externalities are likely to be small over the range of TM and MA enrollment shares estimated in our counterfactuals. To check this, we regressed the log of the risk-adjusted TM cost on the log of the total MA share and county and year fixed effects. The coefficient on the MA share was -0.007 (t-value 2.19) indicating that this is likely a second-order concern.

Table 8: Aggregate market share, consumer welfare, and government spending under 2017 policy and optimal policy by direction of optimal policy change

Markets in which benchmark increases (300 markets, 71.1% population share)

	2017 Policy	Optimal Policy	% Change
Total MA share (%)	28.0	49.6	77.2
Aggregate consumer welfare ($ billion)	3.82	9.17	139.9
Spending on Medicare Advantage ($ billion)	83.1	158.9	91.1
Spending on Traditional Medicare ($ billion)	238.6	164.2	-31.2
Total Medicare Spending ($ billion)	321.7	323.2	0.5

Markets in which benchmark decreases (139 markets, 28.9% population share)

	2017 Policy	Optimal Policy	% Change
Total MA share (%)	32.2	20.2	-37.2
Aggregate consumer welfare ($ billion)	1.88	1.09	-42.1
Spending on Medicare Advantage ($ billion)	38.9	23.9	-38.7
Spending on Traditional Medicare ($ billion)	78.7	92.3	17.3
Total Medicare Spending ($ billion)	117.7	116.2	-1.2

Government spending across all markets

Spending on Traditional Medicare ($ billion)	317.3	256.6	-19.1
Spending on Medicare Advantage ($ billion)	122.1	182.8	49.7
Total Medicare Spending ($ billion)	439.4	439.4	0.0

Notes: Units are billions of 2017 dollars per year. All figures are calculated from individual-level data, aggregated using the MCBS sample weights. Totals may differ due to rounding.

For any line $X(t)$,

$$
\begin{aligned}
CS(X(1)) - CS(X(0)) &= \int_0^1 \nabla CS(X(t)) dX(t) dt \\
&= \nabla_X CS(X(0)) \cdot dX(1) + o(||dX(1)||),
\end{aligned}
$$

where $\nabla_X CS(X(0))$ is a $1 \times \#X$ vector of derivatives of CS with respect to premium and each product characteristic, and $dX(1)$ is a $\#X \times 1$ vector of changes in premiums and product characteristics. This can be rewritten as

$$
\sum_i \sum_j \sum_k \nabla CS_{ijk}(X(0)) \cdot dX(1)_{jk} + o(||dX(1)||), \tag{22}
$$

where i denotes consumers, j denotes products, and k denotes product characteristics (including premiums). In short, the change in surplus for individual i is approximately equal to the sum of the effects of each product characteristic for each good. By re-arranging this expression and summing over goods for each product characteristic, we can compute the effects of changes in each product characteristic as the subsidy schedule changes from the actual policy to our calculated optimal policy—we approximate Equation (22) with

$$
\frac{1}{\alpha}\left(\sum_j \sum_k s_{ij}(\delta(0), \mu(0)) + s_{ij}(\delta(1), \mu(1))\beta_{ik}\Delta X_{jk}\right).
$$

We report the results of this exercise in Table 9. The top and bottom panels reports the effects of changes in markets where in the benchmark increases and decreases, respectively. We report the distribution of percentage contributions of changes in each product characteristic to the overall change in compensating variation to focus on the effect of differences in product characteristics rather than differences in consumer demographics which change the level of welfare. For example, in markets where the benchmark increases, changes in the hospital copay generate an average of 7.45% of the overall change in compensating variation.

The results highlight the importance of modeling product characteristics, and particularly the difference between markets based on the direction of the benchmark change. In markets in which the benchmark increases (decreases), non-premium product characteristics contribute a net average of 40.3% (26.0%) of the change in compensating variation. This difference is driven in part by the zero-lower bound on premium: as the benchmark increases, the annual premium is driven to zero for more plans and so product characteristics take on a greater role. As the benchmark decreases, all plans may increase their premiums.

Despite these results, it is possible that our optimal policy is not a function of product characteristics. For example, if characteristics were held fixed, firms may change the pass-through from the benchmark to the premium. We therefore measure the contribution of product characteristics in an alternative way by recomputing the consumer surplus maximizing policy holding all non-premium characteristics fixed at 2017 levels; changes in the benchmark then represent pure premium subsidies. The mean absolute difference between the benchmarks under this "fixed characteristic" policy and our optimal policy is $114.96, 43.0% of the mean change reported above.

8.3 The selection of markets for increase and decreases

The construction of our optimal subsidy problem implies that markets should be selected for increases based upon the marginal impact of an increase in the benchmark rate on consumer welfare and government expenditures. Table 10 reports the distribution of the derivatives of CS, $\overline{CS}^{cond}$, $GovExp$, and 'total surplus' (defined as $CS - GovExp$) with respect to a $1 increase in the benchmark rate. The first set of rows focuses on markets in which the benchmark increases, and the second set of rows focuses on benchmark decreases.

The distributions of the derivatives of both CS and $\overline{CS}^{cond}$ overlap across the two sets of markets, suggesting that changes in compensating variation alone do not drive the results. Indeed, the mean change in CS is higher in the markets where the optimal policy decreases the benchmark. In contrast, both the distributions of the $GovExp$ and 'total surplus' derivatives are more separated across the two sets of markets. 214 of the markets with benchmark increases have positive total surplus derivatives at the 2017 policy level, and every market which receives a benchmark decrease has a negative total surplus derivative at the 2017 policy level. The derivative of $GovExp$ depends on price elasticities, MA costs relative to TM costs, and the extent to which competition leads firms to pass-through increases in the benchmark to benefits. While it is difficult to disentangle these interrelated factors, to the

Table 9: The contribution of product characteristics to changes in compensating variation

	Mean	25th %-ile	Median	75th %-ile
Markets in which benchmark increases (300 markets, 71.1% pop. share)				
Enhanced prescription drug	6.71	4.58	5.37	7.49
Dental	2.06	1.43	1.64	2.33
Vision	-.351	-.396	-.279	-.236
Hearing	2.81	1.85	2.48	3.27
ξ	19.6	13.5	15.5	21.5
Part B reduction	.252	.168	.219	.308
Deductible	-1.73	-1.94	-1.35	-1.21
Primary care copay	3.47	2.45	2.80	3.81
Hospital copay	7.45	5.22	5.94	8.33
Total non-premium characteristics	40.3	27.7	32.5	44.6
Annual premium	59.7	55.4	67.5	72.3
Markets in which benchmark decreases (139 markets, 28.9% pop. share)				
Enhanced prescription drug	4.28	3.65	3.76	3.96
Dental	1.33	1.05	1.24	1.33
Vision	-.235	-.230	-.212	-.185
Hearing	1.71	1.33	1.68	1.86
ξ	12.7	10.8	11.2	11.7
Part B reduction	.167	.125	.143	.182
Deductible	-1.14	-1.08	-1.02	-.939
Primary care copay	2.31	1.92	2.03	2.18
Hospital copay	4.88	4.10	4.31	4.60
Total non-premium characteristics	26.0	22.2	22.9	24.2
Annual premium	74.0	75.8	77.1	77.8

Notes: All entries are percentages; the numerator is the change in compensating variation attributed to changes in the specified product characteristic, and the denominator is the total change in compensating variation. See text for details. All statistics are across individuals and weighted by the MCBS sample weights.

extent that the optimizer seeks to maximize the “bang for the buck”, these results suggest that the ‘buck’ (how much more the government spends when the benchmark increases by a dollar) matters more than the ‘bang’ (how much additional surplus consumers receive when the benchmark increases by a dollar).

Finally, we note that in many markets, the derivative of GS is negative at the 2017 benchmark. This can occur when the benchmarks are set lower than TM spending, and the cost-savings from beneficiaries that move from TM to MA outweighs the cost increases for inframarginal MA enrollees. We explore this in more detail in Section 8.5.

Table 10: Derivatives of surplus and spending functions with respect to a $1 benchmark increase at 2017 policy by direction of optimal policy change

	Mean	25th %-ile	Median	75th %-ile
Markets in which benchmark increases				
(300 markets, 71.1% pop. share)				
Compensating variation per Medicare beneficiary	.165	.108	.143	.199
Compensating variation per MA enrollee	.150	.094	.133	.180
Government expenditures per Medicare beneficiary	.020	-.154	.011	.227
$(CS - GovExp)$	.145	-.059	.118	.283
Markets in which benchmark increases				
(139 markets, 28.9% pop. share)				
Compensating variation per Medicare beneficiary	.209	.121	.164	.283
Compensating variation per MA enrollee	.193	.127	.185	.230
Government expenditures per Medicare beneficiary	.532	.331	.487	.661
$(CS - GovExp)$	-.323	-.398	-.278	-.219

Notes: Derivatives are calculated at the market level and are weighted by the MCBS sample weights.

Given the importance of these derivatives to determining the direction of benchmark changes, a natural question is the extent to which market-level observables can explain the variance we see in these derivatives. We model the derivatives of the consumer surplus and government expenditure functions as a linear function of our county-level observables collected from the Area Health Resources File detailed in Section 3. Appendix Table E.2 reports the means of these variables by benchmark quartile.

Columns (1) and (2) of Table 11 present standardized regression coefficients when the dependent variable is the derivative of consumer surplus and government expenditures, respectively. Across the two regressions, TM costs, measures of competition, income, and risk enter significantly.

8.4 A linear policy rule

A related question is the extent to which these variables can be used to generate linear policy rules. Analyzing such rules may be particularly policy-relevant, as the difficulties involved in implementing optimal tax and subsidy systems are well-known (see e.g. Scott Morton, 1997, Decarolis, 2015, Jaffe and Shepard, 2017). Column (3) of Table 11 models the optimal benchmark. We include the number of firms offering plans but not the number of plans due to endogeneity concerns. We include racial data due to their importance in the demand system, though we note that we do not intend this as a normative exercise. The optimal benchmark is most strongly associated with measures of competition, cost, and income.

Table 12 compares the 2017 policy to the policy generated by the fitted values of this regression. The linear rule increases aggregate consumer welfare to \$10.14 billion; thus it captures over 95% of the aggregate gains generated by the optimal policy. However, the linear rule increases government spending by 0.296%. Under this rule, 288 markets receive benchmark increases (as opposed to 300). The second panel shows that, on average, the linear rule results in changes which are "too large" relative to the optimal policy; consumers in markets which receive benchmark decreases lose more surplus in percentage terms than under the optimal policy.

8.5 Alternative social welfare functions

The results to this point show the optimal policy creates winners and losers relative to the current policy. We therefore explore other social welfare functions given by

$$\max_{\{B_m\}} \zeta \sum_m \int_i CS_{im}(B_m)di - (1-\zeta)Var(CS) \quad \text{s.t.} \quad \sum_m GovExp_m(B_m) = \bar{B}, \quad \text{and} \tag{23}$$

$$\min_{\{B_m\}} \sum_m GovExp_m(B_m) \text{ s.t. } B_m \geq \bar{B}_m \forall m. \tag{24}$$

Equation (23) is similar to Equation (1) with the addition of a penalty for variance in compensating variation parameterized by ζ. Equation (24) seeks to minimize government expenditures, potentially subject to a floor $\bar{B}_m$. The results are reported in Table 13. We compare the policies generated by these functions to the optimal policy. Columns (2) and

Table 11: Modeling the optimal policy and the derivatives of surplus and spending at the 2017 policy as a function of market-level observables

	(1) Consumer Surplus	(2) Government Expenditures	(3) Log Optimal Benchmark
Log of Risk-adj. per-cap. TM costs	-0.2763	-0.7727	0.0822
	(0.0583)	(0.0382)	(0.0018)
Number of MA firms	0.1594	0.2038	-0.0013
	(0.0634)	(0.0415)	(0.0020)
Log of Medicare beneficiaries	0.1358	0.0188	-0.0079
	(0.0845)	(0.0553)	(0.0027)
Share of 65+ population who is White	-0.5323	0.1095	0.0180
	(0.1870)	(0.1225)	(0.0059)
Share of 65+ population who is Black	-0.5042	0.0976	0.0136
	(0.1473)	(0.0965)	(0.0047)
Share of 65+ population who is Hispanic	-0.3741	0.0899	0.0091
	(0.1198)	(0.0785)	(0.0038)
Average risk score	0.1445	0.1972	-0.0031
	(0.0761)	(0.0499)	(0.0024)
Log of median household income	-0.0507	0.0222	-0.0051
	(0.0835)	(0.0547)	(0.0026)
Share of 65+ population in deep poverty	-0.0482	0.0066	-0.0006
	(0.0554)	(0.0363)	(0.0018)
Unemployment rate	-0.0060	-0.0836	-0.0002
	(0.0655)	(0.0429)	(0.0021)
Population density	0.0319	-0.1389	-0.0081
	(0.0543)	(0.0356)	(0.0017)
MDs per capita	-0.0239	-0.0209	0.0019
	(0.0641)	(0.0420)	(0.0020)
Medicare-qualified hospitals per capita	0.0008	0.0144	0.0011
	(0.0472)	(0.0309)	(0.0015)
Nursing facilities per capita	-0.0399	-0.0364	-0.0044
	(0.0602)	(0.0394)	(0.0019)
Hospice facilities per capita	-0.0270	-0.0334	0.0003
	(0.0478)	(0.0313)	(0.0015)
Medicare hospital readmission rate	-0.0714	-0.0476	-0.0002
	(0.0641)	(0.0420)	(0.0020)
Observations	439	439	439
R-squared	0.1601	0.6397	0.8693

Notes: The independent variables have been normalized to have mean zero and unit variance. The dependent variables for Columns (1) and (2) are the derivatives with respect to the benchmark of GS and $GovExp$ functions as defined in Equations (9) and (16), respectively, and are also normalized. The dependent variable in Column (3) is the optimal subsidy schedule of Table 7, in units of thousands of dollars per year. Estimates obtained by OLS. Robust standard errors are in parentheses.

Table 12: Comparing the 2017 policy to the linear policy rule

	2017 Policy	Linear Rule	Percentage Change
Mean compensating variation per Medicare beneficiary ($)	135.71	241.42	77.9
Total MA share (%)	29.2	39.7	36.1
Mean compensating variation per MA enrollee ($)	464.68	607.53	30.7
Aggregate consumer welfare generated by MA ($ billion)	5.70	10.14	77.9
Total Medicare Spending ($ billion)	439.4	440.7	.296
Mean compensating variation for MA enrollees in...			
288 markets with benchmark increases ($)	439.51	611.08	39
151 markets with benchmark decreases ($)	442.44	351.45	-20.6
Mean compensating variation for MA enrollees who are...			
White ($)	436.53	574.89	31.7
Black ($)	462.73	587.29	26.9
Hispanic ($)	542.71	728.64	34.3
Low income ($)	458.28	594.36	29.7
Medium income ($)	440.48	573.25	30.1
High income ($)	407.72	565.02	38.6

Notes: The linear rule is the policy fit by the regression in Column (3) of Table 11. Compensating variation is calculated via Equation (9) and is reported in dollars per year. Compensating variation conditional on enrollment is calculated via Equation (10). White, Black, and Hispanic groups are defined by CMS. "Low income" refers to the lowest quartile of income, "high income" refers to the highest quartile, and "middle income" refers to the second and third quartiles. All statistics are weighted by MCBS sample weights.

(3) report the results of solving Equation (23) with $\zeta = 0.99$ and $\zeta = 0.999$ respectively. In Column (2), the penalty on variance causes a reduction in benchmarks nearly everywhere in order to fund increases in a few markets; the aggregate consumer welfare generated by MA drops to $2.58 billion. Column (3) reduces the penalty on the variance, which results in a policy very similar to the optimal policy with $8.77 billion in aggregate consumer welfare.

Column (4) seeks to minimize government expenditures. Where the 2017 MA payments are larger than the cost of TM, this is done by reducing the benchmark. However, there are markets where the 2017 policy results in MA payments that are on average lower than TM costs. Increasing the benchmark results in both intensive and extensive margin changes to MA payments: the government must pay more for consumers who were already enrolled in an MA plan, and must transfer payments from the TM system to the MA system for consumers who switch. In 182 markets, the extensive margin impact is larger than the intensive margin impact, and therefore an increase in the benchmark rate results in a decrease in total government expenditures.[22] As a result, the government can reduce total spending on TM and MA by \$3.5 billion, though at the cost of \$4.98 billion in aggregate consumer welfare relative to the optimal policy.

Column (5) seeks a "consumer-Pareto improvement" by minimizing government expenditures subject to the constraint that no benchmark is lowered below its 2017 level. As the benchmark is raised in 182 markets, total consumer welfare increases from the 2017 policy to \$7.23 billion. At the same time, government spending is reduced by \$1.4 billion.

8.6 Other results and robustness

In this subsection, we briefly describe a number of other results. We begin with investigating the robustness of our counterfactual equilibria, as while our approach finds equilibria in premiums taking product characteristics as given, this may not be an equilibrium in the full game described by Equation (17). At the same time, the possibility of adjustment costs implies that 'standard' first-order conditions with respect to product characteristics may not hold with equality. To investigate our counterfactual equilibrium, we calculate the change in variable profits firms would obtain by deviating from the predicted values for each product characteristic at the optimal policy by 1% of the sample mean, and subtract it from the change in profits for deviations from the observed 2017 characteristics. The difference represents error introduced by our counterfactual approach.

The results are reported in Appendix Table E.5. Across all markets, deviations in hospital copays pose the largest potential profit: a 1% increase in hospital copays leads to a mean increase in profits of 0.0073%. For all other characteristics the mean profitability change is less than 0.006% in absolute value, indicating that our approximation error is second order. The Monte Carlo results in Appendix C suggest that approximation error increases as the magnitude of the benchmark change increases. The bottom rows of Appendix Table E.5 split markets into those with changes in the benchmark of more and less than \$1,000. As expected, the 60 markets with changes greater than \$1,000 have more profitable deviations

[22]In Appendix D we illustrate government expenditures as a function of the benchmark for three markets, including a market which features this behavior.

Table 13: Surplus, share, and spending under the optimal policy and alternative social welfare functions

	(1) Optimal Policy	(2) Max 0.99 CS − 0.01 Var	(3) Max 0.999 CS − 0.001 Var	(4) Minimize *GovExp*	(5) Min. Exp. With Floor
Mean compensating variation per Medicare beneficiary ($)	244.20	62.50	229.50	125.78	172.07
Total MA share (%)	41.1	17.8	43.6	26.8	35.1
Mean compensating variation for MA enrollees ($)	593.97	350.35	526.84	468.87	490.16
Aggregate consumer welfare generated by MA ($ billion)	10.26	2.62	9.64	5.28	7.23
Total government expenditures ($ billion)	439.4	439.4	439.4	435.9	438.0
Number of markets with benchmark increases	300	154	346	182	182
Mean compensating variation for MA enrollees in...					
Markets with benchmark increases ($)	588.03	324.84	510.33	470.45	470.47
Markets with benchmark decreases ($)	431.44	333.20	446.83	394.68	N/A
Mean compensating variation for MA enrollees who are...					
White ($)	563.17	326.89	496.19	437.44	459.59
Black ($)	576.58	357.81	540.15	485.51	499.13
Hispanic ($)	705.01	395.82	642.14	595.40	616.99
Low income ($)	585.35	344.67	528.06	467.32	487.28
Medium income ($)	559.42	328.01	497.29	438.82	463.05
High income ($)	553.77	313.24	472.34	422.74	435.41

Notes: Columns (2) and (3) maximize consumer surplus with a penalty on the variance of consumer surplus across individuals. Column (4) minimizes the sum of government expenditures across MA and TM. Column (5) also minimizes the sum of government expenditures, with the constraint that no benchmark can be reduced relative to the 2017 policy. Compensating variation is calculated via Equation (9) and is reported in dollars per year. Compensating variation conditional on enrollment is calculated via Equation (10). White, Black, and Hispanic groups are defined by CMS. "Low income" refers to the lowest quartile of income, "high income" refers to the highest quartile, and "middle income" refers to the second and third quartiles. All statistics are weighted by MCBS sample weights.

– the mean change in variable profit from a 1% change in hospital copays is 0.0204%.

The optimal policy may have political economy implications if the changes result in a large-scale redistribution of government expenditure dollars and consumer welfare across states. Indeed, past changes to the MA payment formula have likely resulted from political considerations during the legislative process (Berenson, 2008). To explore these issues, we summarize the total consumer surplus and government expenditures by state in Appendix Table E.3. Of the 41 states (plus Washington D.C.) included in the 2017 MCBS, 29 receive increases to aggregate consumer welfare. The results suggest the optimal policy does not split cleanly along political divisions.

Appendix Table E.4 reports firm-level market shares and total variable profits under the 2017 policy, the optimal policy, and other policies explored in Table 13. The optimal policy increases aggregate insurer variable profits from $3.51 billion to $5.63 billion. The absolute shares of each of the five largest insurers increase, though Humana and UnitedHealth Group lose share relative to others. The linear rule results in $5.42 billion in aggregate variable profits. The 'Pareto improvement' specification results in $4.45 billion in aggregate variable profits; again, the absolute shares of the largest insurers increases.

9 Conclusion

Seeking to reduce the perceived inefficiency of government-provided goods and services, policy makers have implemented public-private partnerships in which the government provides subsidies to private firms that are tied to the consumers' choices. The firms compete with market forces working to bring down the total cost and increase the benefits of providing the good over time. In many cases, the goods have differentiated characteristics which are relevant to consumers. Additionally, these goods may be offered in geographies with consumers who have substantially different preferences. While the subsidy rates are generally set according to measures of government costs, the optimal subsidies conditional on a fixed budget depend on equilibrium interactions between heterogeneous firms and consumers.

We provide a framework for calculating the optimal subsidies that takes into account both supply and demand responses to alternative subsidy rates. We model demand with a discrete-choice system and avoid the curse of dimensionality in computing counterfactual product characteristics with an approach that combines policy function estimation with a first-order condition solver.

We apply our framework to MA in the United States, through which approximately one-third of U.S. seniors obtain Medicare benefits, and estimate our model using micro-level panel data. We derive demand instruments from detailed data on costs, which are likely to be available in other managed competition environments, and from the panel nature of our plan data. We derive policy function instruments from county-level differences in urban/rural status following Duggan et al. (2016).

We find that the optimal (budget-neutral) subsidies differ substantially from those currently employed by the government. The current policy generates an average of \$135.71 in consumer surplus as measured by compensating variation per Medicare beneficiary per year. By maximizing the mean consumer surplus while holding government spending constant, we find a policy that results in an average of \$244.20 in benefits per beneficiary per year. These gains come more from finding places where it is easy (in the sense of needing fewer government dollars) to move people from TM to MA than from improving benefits for people already utilizing MA. Accounting for endogenous product characteristics is important; changes in product characteristics (as opposed to prices) account for over 35% of the changes in total surplus and a policy computed under the assumption that non-price characteristics remain constant is significantly different. We show that freely-available market-level observables can be used to approximate the optimal policy with a linear rule that captures over 95% of the consumer welfare gains at the cost of an increase in government epxenditures of 0.30% relative to the optimal policy. Apart from finding a particular consumer-surplus-maximizing policy, which may not be implementable for political or other reasons, our framework can be used to evaluate the outcomes of any proposed subsidy schedule.

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

Adrion, E. R. (2020). Politics and policymaking in medicare part c. *Medical care*, 58(3):285–292.

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.

Arrow, K. J. (1963). Uncertainty and the welfare economics of medical care. *American Economic Review*, LIII(5):941–973.

Ashman, J. J., Rui, P., and Okeyode, T. (2019). Characteristics of office-based physician visits, 2016. *NCHS Data Brief, no 331*.

Ayanian, J. Z., Landon, B. E., Newhouse, J. P., and Zaslavsky, A. M. (2014). Racial and ethnic disparities among enrollees in medicare advantage plans. *New England Journal of Medicine*, 371(24):2288–2297.

Baicker, K., Mullainathan, S., and Schwartzstein, J. (2015). Behavioral hazard in health insurance. *The Quarterly Journal of Economics*, 130(4):1623–1667.

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*.

Berenson, R. A. (2008). From politics to policy: A new payment approach in medicare advantage. *Health Affairs*, 27(Suppl1):w156–w164.

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

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

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

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

Curto, V., Einav, L., Levin, J., and Bhattacharya, J. (2021). Can health insurance competition work? evidence from medicare advantage. Technical Report 2.

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.

Decarolis, F. (2015). Medicare part d: Are insurers gaming the low income subsidy design? *American Economic Review*, 105(4):1547–80.

Decarolis, F. and Guglielmo, A. (2017). Insurers' response to selection risk: Evidence from Medicare enrollment reforms. *Journal of Health Economics*, 56:383–396.

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

Diamond, P. A. and McFadden, D. L. (1974). Some uses of the expenditure function in public finance. *Journal of Public Economics*, 3(1):3–21.

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.

Grossman, M. (1972). On the concept of health capital and the demand for health. *Journal of Political economy*, 80(2):223–255.

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.

Hicks, J. R. (1945). The generalised theory of consumer's surplus. *The Review of Economic Studies*, 13(2):68–74.

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.

Kaiser Family Foundation (2021). State health facts: Medicare service use: Hospital inpatient services. Technical report, Kaiser Family Foundation.

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

Lakdawalla, D. and Yin, W. (2015). Insurers' negotiating leverage and the external effects of medicare part d. *Review of Economics and Statistics*, 97(2):314–331.

Li, Y., Cen, X., Cai, X., Thirukumaran, C. P., Zhou, J., and Glance, L. G. (2017). Medicare advantage associated with more racial disparity than traditional medicare for hospital readmissions. *Health Affairs*, 36(7):1328–1335.

Loewenstein, G., Friedman, J. Y., McGill, B., Ahmad, S., Linck, S., Sinkula, S., Beshears, J., Choi, J. J., Kolstad, J., Laibson, D., et al. (2013). Consumers' misunderstanding of health insurance. *Journal of Health Economics*, 32(5):850–862.

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

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. *Frontiers in Econometrics*, pages 105–142.

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

McKean, R. N. and Minasian, J. R. (1966). On achieving pareto optimality-regardless of cost! *Economic Inquiry*, 5(1):14.

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

Meyers, D. J., Mor, V., Rahman, M., and Trivedi, A. N. (2021). Growth in medicare advantage greatest among black and hispanic enrollees. *Health Affairs*, 40(6):945–950.

Nevo, A. (2000). Mergers with differentiated products: The case of the ready-to-eat cereal industry. *The RAND Journal of Economics*, pages 395–421.

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.

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

Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, 110:705–729.

Poterba, J. M. (1996). Government intervention in the markets for education and health care: how and why? In *Individual and social responsibility: Child care, education, medical care, and long-term care in America*, pages 277–308. University of Chicago Press.

Rothschild, M. and Stiglitz, J. (1976). Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *The Quarterly Journal of Economics*, 90(4):629–649.

Scott Morton, F. M. (1997). The interaction between a most-favored-customer clause and price dispersion: an empirical examination of the medicaid rebate rules of 1990. *Journal of Economics & Management Strategy*, 6(1):151–174.

Smith, P. C., Stepan, A., Valdmanis, V., and Verheyen, P. (1997). Principal-agent problems in health care systems: an international perspective. *Health policy*, 41(1):37–60.

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.

Sweeting, A. (2007). Dynamic product repositioning in differentiated product markets: The case of format switching in the commercial radio industry.

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

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

Wollmann, T. G. (2018). Trucks without bailouts: Equilibrium product characteristics for commercial vehicles. *American Economic Review*, 108(6):1364–1406.

Appendices

A Accounting for the CMS Bidding System

In Section 5, we model firms as choosing prices and product characteristics while receiving a subsidy as a function of those product characteristics. In reality in our setting, firms choose a bid b_{jt} and a premium p_{jt} in addition to product characteristics. In this Appendix, we show how the CMS bidding rules can be used to transform the bidding problem into the price-setting problem.

Per Equation (2), the government pays the firm the amount of their bid plus a rebate payment if their bid is less than the benchmark. After taking into account risk adjustment, we can write the rebate payment as a function of the bid b_{jt} and the plan-level benchmark

$B_{jt} = B_{mt} \times \phi_{ft}$ with

$$reb(b_{jt}; B_{jt}, \lambda_{ft}) = \begin{cases} \lambda_{ft}(B_{jt} - b_{jt}) & \text{if } b_{jt} < B_{jt} \\ 0 & \text{if } b_{jt} \geq B_{jt} \end{cases}, \tag{25}$$

where λ_{ft} is the rebate percentage.

CMS requires that rebate payments be used to fund benefits beyond those offered by TM. Both supplemental benefits and cost-sharing reductions may be paid for with rebate funds, but cost-sharing reductions *must* be paid for with rebate funds. We assume this constraint is binding in the following sense: the rebate the firm receives is exactly equal to the incremental cost of providing the plan's cost-sharing benefits over the cost of providing a 'base' TM-equivalent plan. To notate this, let $x_{c,jt}$ be the subvector of product characteristics capturing cost-sharing reductions. We write the incremental cost function as $incr_{jt}(x_{c,jt})$ which is greater than zero as firms must provide at least TM-equivalent coverage, though it may vary by firm. We assume

$$reb_{jt} = incr_{jt}(x_{c,jt}). \tag{26}$$

Combined, Equations (25) and (26) imply that the bid itself is determined by the choice of cost-sharing benefits $x_{c,jt}$ and the incremental cost function $incr_{jt}$. In other words, our assumption can be reinterpreted to mean that insurers bid in such a way that the rebate they receive exactly pays for the incremental cost. There is therefore a continuous function mapping the choice of $x_{c,jt}$ onto the bid b_{jt}:

$$b(x_{c,jt}; B_{jt}, \lambda_{ft}, \cdot) = B_{jt} - \frac{1}{\lambda_{ft}} incr_{jt}(x_{c,jt}). \tag{27}$$

As a consequence, we can model the firm as simply choosing a price p_j and product characteristics. Finally, the government pays the minimum of the bid plus the rebate and the benchmark: $sub(b_{jt}; B_{jt}, \lambda_{ft}) = \min\{b_{jt} + \lambda_{ft}(B_{jt} - b_{jt}), B_{jt}\}$. Thus, since $incr_{jt}(\cdot) \geq 0$,

$$sub(x_{c,jt}; B_{jt}, \lambda_{ft}, incr_{jt}) = B_{jt} - \frac{1 - \lambda_{ft}}{\lambda_{ft}} incr_{jt}(x_{c,jt}). \tag{28}$$

B Computational details

In this Appendix we provide computational details for our counterfactual approach. Given the demand, supply, and policy function estimates of Section 7, we begin by updating product characteristics. For plan j in market m with plan-level benchmark B_j and counterfactual benchmark B', we calculate the counterfactual product characteristic $\widehat{x_{lj}(B')} = x_{lj}(B_m) + \hat{\beta}_{f,l}(B' - B_m)$ where $x_{lj}(B_m)$ is the value observed in the data. This ensures that when we input the benchmarks in the data, the counterfactual recreates the product characteristics in the data. Given $\widehat{X(B')}$, we counterfactual estimate plan costs $\widehat{c_j(B')} = c_j(\widehat{X(B')})$ using our estimates of the marginal cost shocks.

B.1 Incremental costs and bids

While the model introduced in Section 5 can be estimated without defining the incremental cost function, as we observe payments to firms, computing *GovExp* for counterfactual benchmarks requires computing the incremental cost for each plan as plans change product characteristics. Let $c_{b,ft} = \gamma_f + \gamma_s + \gamma_r + \gamma_m + \gamma_t$. Let $\gamma_{xc,jt}$ be the subvector of cost parameters associated with $x_{c,jt}$ and let $\omega_{c,jt}$ be the unobservable component of costs associated with cost-sharing benefits. The log-cost of cost-sharing benefits is then $c_{c,jt} = \gamma'_{xc,jt} x_{c,jt} + \omega_{c,jt}$.[23] We define the incremental cost function as

$$incr(x_{xc,jt}; x_{b,ft}, \gamma_{xc,jt}, \omega_{c,jt}) = \exp[c_{b,ft} + c_{c,jt}(x_{c,jt})] - \exp(c_{b,ft}). \tag{29}$$

As we observe rebates for each plan, after estimating the marginal cost function, we can estimate ω_{jt} and use Equations (26) and (29) to estimate $\hat{\omega}_{c,jt}$ for each plan. Then, after updating x, we calculate $incr(\cdot)$ using this estimated $\hat{\omega}_{c,jt}$. With the incremental cost in hand, we use Equation (27) to calculate bids and Equation (28) to calculate government payments to firms.

B.2 Solving for equilibrium prices

With product characteristics, costs, and subsidies in hand, we turn to the firms' price-setting problem. First, we note that in principle onx could also estimate policy functions for prices;

[23] We have explored specifications with linear cost functions and found similar results.

the assumptions that ensure the consistency of $\widehat{x_{lj}(B')}$ may also apply to p_j. However, in finite samples the error in functions of the equilibrium components may be sizable as, for example, it may be a function of the sum of the magnitudes of the ϵ_{lj}. To increase the accuracy of our estimated equilibrium, and therefore any functions of equilibrium objects, we combine the policy function estimation approach with a fixed-point problem. That is, using the notation of Equation (17), we estimate policy functions for some action subvector X_{1f} and then solve the reduced game with remaining action subvector X_{2f} given by (with a slight abuse of notation)

$$\max_{X_{2f}} \pi_f(X_{2f}; X_{-2f}, \widehat{X_{1f}}, \widehat{X_{-1f}}, Z_f, B) \tag{30}$$

by searching for a fixed point in the first-order conditions.

This reduced game may also admit multiple equilibria. Assumption ES provides a way forward: we use observed actions as a basis for calculating counterfactual equilibria by discarding equilibria inconsistent with our data in the sense that a small change in the benchmark generates a larger change in equilibrium outcomes than is observed in the data. In practice, this does not bind in our setting as we do not find evidence of multiple equilibria.

Given an equilibrium X^*_{2f} in the reduced game, the first-order conditions of Equation (18) with respect to the elements of X_{2f} will be satisfied. In practice, the remaining first-order conditions may not be satisfied exactly due to the errors encapsulated in ϵ_{lj}. We therefore check these conditions after computing the solution to Equation (1). Note, however, that this hybrid procedure will result in estimates of $X^*(B)$ (and therefore estimates of CS and $GovExp$) which are almost certainly 'better' than estimates obtained through policy function estimation alone as measured by the magnitude of deviation from the equilibrium conditions.

B.3 Solving Equation (1)

With equilibrium vectors in hand, we calculate CS and $GovExp$ using Equations (9) and (16), respectively. These functions are non-linear, and so we implement the government spending constraint with a penalty function and address the possibility of multiple local maxima with a multi-start procedure. We restrict the set of counterfactual benchmarks we consider to the range of benchmarks we observe in the data, though in practice this constraint does not bind.

Finally, we obtain substantial computational efficiency by noting that our problem is separable; outcomes in one market do not directly affect outcomes in other markets. We proceed by solving equilibria for each market on a grid of benchmarks in a first stage and then evaluating candidate benchmark schedules using grid interpolation. Both of these steps benefit from parallel computing.

C A Monte Carlo Analysis of our Counterfactual Approach

In Section B, we detail an approach for calculating counterfactual equilibria. In our MA application, we predict changes in the plan characteristics with estimated policy functions and then solve for the equilibrium in prices. An obvious question is how well does this approach work in practice. In this Appendix, we explore the performance of our approach. We write a simplified model and simulate market-level data for two periods by solving for equilibria explicitly. We use the results in the first period as the basis for making predictions of the second period outcome using our approach. We then compare the exact solution to our estimated solution both in terms of the firms' actions (the objects being predicted) and in terms of consumer welfare (the object of interest).

To focus our attention on the uncertainty introduced by our technique and to ensure tractability we adopt the following model. Time is denoted by t, and markets by m. Each market has a unit measure of consumers, denoted by i, and F firms denoted by f. Each firm is present in each market, and offers a constant number of products J, denoted by j. Each product consists of a price p_{jmt} (which varies by market) and a $D \times 1$ vector of non-price characteristics δ_{jmt}. The choice-specific utility obtained by consumer i when purchasing product j is

$$u_{ijmt} = \alpha_i p_{jmt}^2 + \beta_i' \delta_{jmt} + \epsilon_{ijmt} \tag{31}$$

where α_i is the price sensitivity of consumer i, β_i is i's $D \times 1$ vector of preferences for non-price product characteristics, and ϵ_{ijmt} is consumer i's idiosyncratic unobservable preference for product j, assumed to be i.i.d. Type-I Extreme Value.

Marginal costs are assumed to be constant at the product level and given by

$$c_{jmt} = \exp(-0.3 + \gamma' \delta^2_{jmt} + \nu_f + \omega_m) \tag{32}$$

where δ^2_{jmt} is element-wise squaring of product characteristics, γ is a $D \times 1$ vector of per-characteristic costs, where each component is drawn from i.i.d. $N(0.1, 0.05)$, ν_f is a firm-specific cost shock that is constant across markets, drawn from i.i.d. $N(0, 0.01)$, and ω_m is a market-specific cost shock that is constant across firms, drawn from i.i.d. $N(0, 0.1)$.

Firms receive a subsidy payment b_{mt} from the government and make decisions at the market level. As decisions made in each market are independent, the firm's problem is

$$\max_{p_{fm}, \delta_{fm}} \pi_{fm} = \sum_{j=1}^{J} \int_i (p_{jmt} - c_{jmt}(\delta_{jmt}) + b_{mt})\, s_{ijmt}(p_m, \delta_m) di \tag{33}$$

where p_{fm} and δ_{fm} represent the vectors of prices and product characteristics for the firm in that market and s_{ijmt} is the probability that consumer i purchases product j as a function of all of the prices and product characteristics in the market. As ϵ_{ijmt} is Type-I Extreme Value, s_{ijmt} takes a logit form.

Equilibrium in m is a vector (p, δ) for all firms such that each firm's choices solve Equation (33) when taking the competitors' choices and the subsidy level as given. While this model abstracts from common empirical issues, it is a framework which allows us to explore the performance of our approximation approach not merely as a function of the size of the dataset (here represented by the number of markets simulated), but also as a function of the number of firms in each market, the number of products offered by each firm, and the number of non-price product characteristics. Each of these factors can potentially affect the structure of the equilibrium, and in particular may affect the shape of the response of firms to changes in the subsidy.[24]

The existence of multiple products and product characteristics raises the potential for both multiple equilibria and for "trivial" equilibria, in which firms offer identical products. We address this issue through the selection of the distribution of consumer preferences. In particular, if D is the number of product characteristics, we define D consumer types with equal proportions among the consumer population. Each consumer type n has a strong

[24]We thank an anonymous referee for this point.

preference for it's 'own' characteristic and a weak preference for 'other' characteristics per

$$\beta_{nd} \sim \begin{cases} N(2.5, 0.1) & \text{if } n = d \\ N(0.1, 0.01) & \text{if } n \neq d \end{cases} \tag{34}$$

where β_{nd} is the dth element of the preference vector for consumers of type n. We set $\alpha_n = 2.5$ for all consumers. These choices ensure that when $J \leq D$, there exists an equilibrium in which firms' strategies involve products that are differentiated and specialized: each product features a high value of a single product characteristic and low values of other product characteristics.

We solve for equilibria by iterating over best response functions. In the first period, we give each market a subsidy ranging from 0 to 1.0, with equal spacing. In the second period, we reverse the order of the subsidies so that the market which received the highest subsidy in the first period now receives the lowest and vice versa.

We implement our approach using data from the first period. Specifically, we use the equilibrium outcomes to estimate demand and invert the first-order conditions for price-setting to recover an estimate of marginal costs $\hat{c}_{jmt}$. We regress this estimate on δ_{jmt} and firm and market fixed effects to recover the cost parameters.[25] We then estimate policy functions for each product characteristic. Let δ_{fjd1} be the dth product characteristic of product j for firm f in period 1. We fit

$$\delta_{fjd1} = \theta_{jd} b_{m1} + FE_f + \epsilon_{fjd1} \tag{35}$$

where θ_{jd} is the parameter of interest, FE_f is fixed effects, and ϵ_{fjd1} is an error term.

We approximate non-price characteristics in period 2 using the estimated $\hat{\theta}_{jd}$ with

$$\hat{\delta}_{fjd2} = (b_{m2} - b_{m1})\hat{\theta}_{jd} + \delta_{fjd1}. \tag{36}$$

We calculate marginal costs at these estimated characteristics using our estimates of the cost parameters. We then solve for an equilibrium with marginal costs and product characteristics as given and prices as the only choice variables for firms.

Figure C.1 illustrates the output of this procedure for a 50-market run. Each market

[25]Note that since there are no product-level unobservables, this procedure recovers exact parameters.

Figure C.1: Example Monte Carlo simulation output

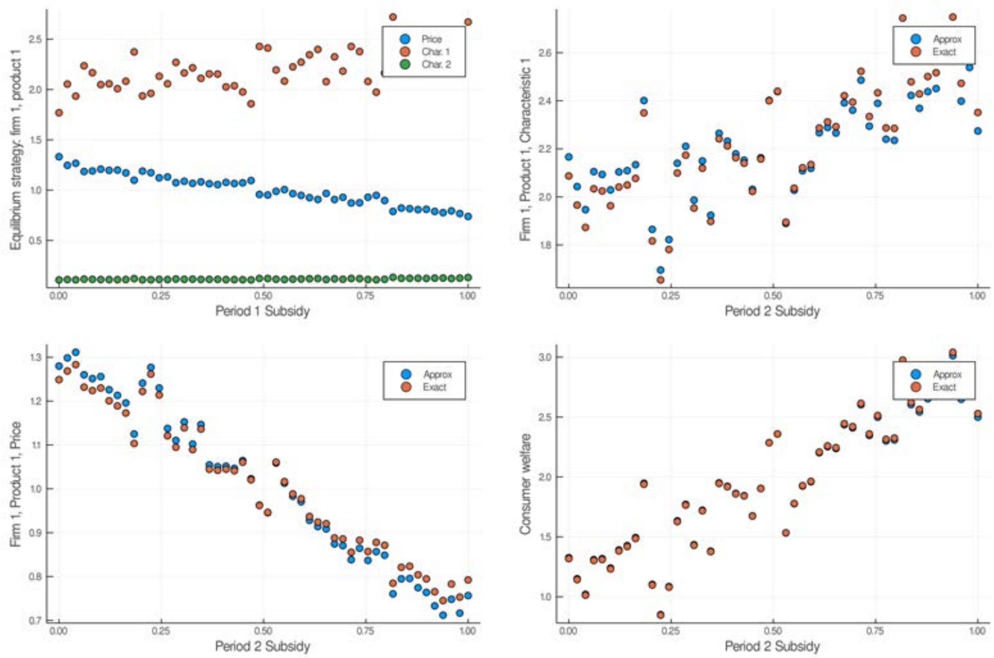

Notes: These graphs illustrate outcomes from a sample Monte Carlo simulation run. For this run, 50 markets were simulated, each with four firms (identical across markets). Each firm offered two products, each of which had two non-price product characteristics. There were two consumer types. The top-left graph illustrates the policy function for firm 1 product 1 across markets, including the price decisions and both product characteristics. The top-right graph compares the exact policy function to the approximated policy function for firm 1, product 1, characteristic 1 in period 2. The bottom-left graph compares the optimal prices when product characteristics are exact and approximated. The bottom-right graph compares the consumer welfare in each market under the exact and approximated solutions.

has two consumer types and four firms. Each firm offers two products with two non-price characteristics. The top-left graph illustrates the product design for firm 1's first product in each market; this product is 'designed' to attract consumers of type 1. A higher subsidy leads to lower prices and 'more' characteristic 1, though there is some variation from the trend due to market-specific cost shocks. The other three graphs compare exact and approximated outcomes in period 2. The top-right graph focuses on the first product characteristic and the bottom-left graph focuses on the price of the first product. The bottom-right graph illustrates the consumer welfare in each market and shows that these differences "net-out": the welfare difference is smaller than the differences in characteristics and prices.

To evaluate the performance of the approximation approach in a systematic way, we compute two statistics. First, we calculate the approximation error in product characteristics by calculating the Euclidean distance between the approximate and exact characteristic vectors and dividing it by the magnitude of the exact characteristic vector: if $D(x, y)$ is the

Table C.1: Monte Carlo simulation results

			50 Markets		100 Markets		200 Markets	
Firms / mkt.	Prods. / firm	Chars. / prod	Err^{δ}	Err^{CW}	Err^{δ}	Err^{CW}	Err^{δ}	Err^{CW}
1	1	1	0.0071	0.0025	0.0098	0.0035	0.0044	0.0015
1	2	2	0.0070	0.0023	0.0098	0.0033	0.0043	0.0014
1	4	4	0.0089	0.0026	0.0116	0.0037	0.0057	0.0017
2	1	1	0.0089	0.0036	0.0111	0.0045	0.0028	0.0011
2	2	2	0.0095	0.0036	0.0123	0.0050	0.0028	0.0009
2	4	4	0.0104	0.0040	0.0227	0.0084	0.0217	0.0086
4	1	1	0.0095	0.0027	0.0119	0.0036	0.0023	0.0006
4	2	2	0.0095	0.0026	0.0121	0.0035	0.0022	0.0005
4	4	4	0.0217	0.0015	0.0130	0.0036	0.0061	0.0006

Notes: In all simulations, the number of consumer types is equal to the number of non-price product characteristics. Err^{δ} and Err^{CW} are defined in the text; reported metrics are means across all markets in the simulation. For ease of comparison, preference and firm-level cost draws are identical across rows, while market-level cost draws are constant across columns.

distance between x and y, we define $Err^{\delta} = D(\delta_m^{Exact}, \delta_m^{Approx})/||\delta_m^{Exact}||$. Second, we calculate consumer welfare per Equation (9) under both the exact solution and the approximated solution and define the absolute logarithm error as $Err^{CW} = |\log(CW_{Exact}) - \log(CW_{Approx})|$. Both of these measures are strictly positive and are designed to capture the amount of error relative to the size of the approximated object.

The results of our simulations are presented in Table C.1. We explore the performance of our approach by varying the complexity of the market both in terms of the number of firms present, as well as the complexity of the product offerings of those firms. We also vary the number of markets simulated to evaluate the performance with datasets of differing lengths. Across scenarios, the performance corresponds with sensible priors. Increasing the complexity of markets tends to increase the mean simulation error when the number of markets is low. Increasing the number of markets tends to decrease the error. Even in the worst case, however, the error in δ remains under 2.5%, and the error in the consumer welfare is less than 1%. This is consistent with Figure C.1 – simulation error in the policy function is offset by firm prices.

D The Non-Local Behavior of Surplus and Expenditures

The results in Section 8.3 show that the derivatives of the CS and $GovExp$ functions at the 2017 policy point to the direction of the optimal policy. However, the derivatives do not

Figure D.1: Prices and product characteristics under counterfactual benchmarks, selected counties

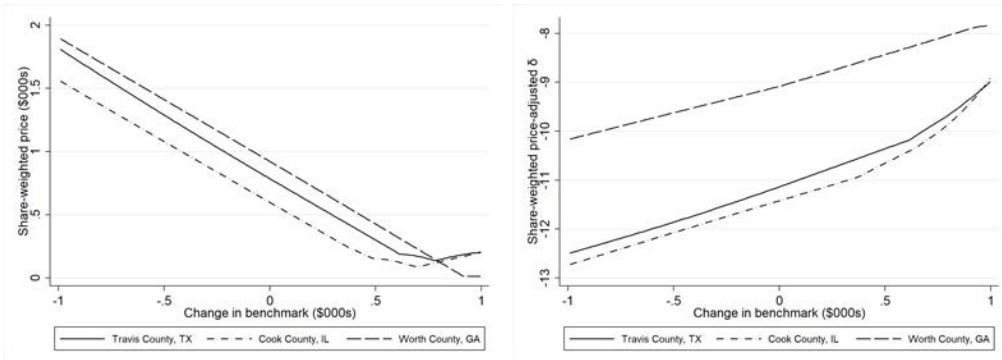

Notes: The horizontal axis is the change in the benchmark relative to 2017. The lines illustrate share-weighted averages across plans.

provide sufficient information to calculate the optimal policy. Though the policy function approximations we use are linear, the pricing behavior and share functions are not and so we should expect the CS and $GovExp$ functions to be non-linear as well. In this Appendix, we explore the non-local behavior of CS and $GovExp$ through illustrated examples.

Figure D.1 illustrates components of the CS function for Travis County, TX (containing Austin), Cook County, IL (containing Chicago), and Worth County, GA (a rural county near Albany), in 2017. We chose these counties due to their different sizes and the typical nature of their counterfactual equilibria. The left-hand graph depicts the share-weighted premium and shows non-monotonicity as the benchmark increases. As the benchmark increases, prices near the zero lower bound. The right-hand graph illustrates the share-weighted $\delta'_j \equiv \delta_j - \alpha_0 p_j$ (i.e. the net utility impact of product characteristics) as a function of the benchmark. All three counties show increases as the benchmark increases. The slope, however, is higher in Cook County and Travis County than in Worth County due in part to differences in the firms present and their associated policy functions, and in part to the amount of competition.

Figure D.2 illustrates components of $GovExp_m$ for the same counties. The left-hand graph depicts the share-weighted plan cost. Cook County sees the highest increases. The middle 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 bids in Travis and Worth Counties increase with a shallower slope. The right-hand graph illustrates the total share of MA (relative to TM). Cook County's share increases the fastest.

Figure D.3 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

Figure D.2: Average plan costs and bids and MA share under counterfactual benchmarks, selected counties

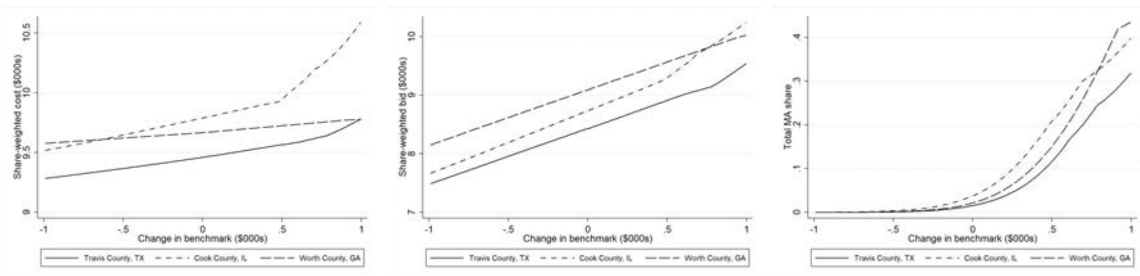

Notes: The horizontal axis is the change in the benchmark relative to the 2017 level.

Figure D.3: Per-capita consumer surplus and government expenditures under counterfactual benchmarks, selected counties

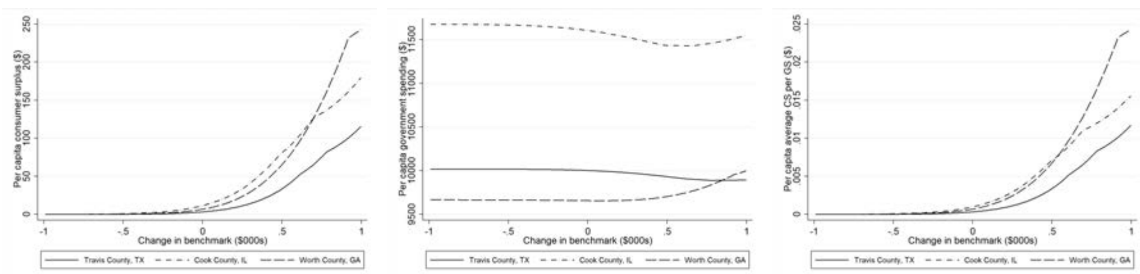

Notes: The horizontal axis is the change in the benchmark relative to the 2017 level.

receive similar surplus. As the benchmark in each county is increased, the average surplus in Worth County grows faster than the others, eventually overtaking Cook County after an increase of $600. 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 final graph of Figure D.3 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.

E Additional Tables and Figures

Table E.1: Policy functions for product characteristics by census region

	(1) Part B reduction	(2) Deduct-ible	(3) Prim. care copay	(4) Hospital copay	(5) Enhanced drug	(6) Dental	(7) Vision	(8) Hearing	(9) ξ
Aetna									
Northeast	0.0141	0.0390	-0.9575	-0.1569	0.1514	0.0641	-0.0136	0.0693	0.6849
	(0.0029)	(0.0069)	(0.1851)	(0.0126)	(0.0111)	(0.0102)	(0.0050)	(0.0095)	(0.0987)
Midwest	0.0139	0.0373	-1.1700	-0.1760	0.1596	0.0439	-0.0134	0.0690	0.5873
	(0.0029)	(0.0069)	(0.1875)	(0.0127)	(0.0110)	(0.0104)	(0.0049)	(0.0094)	(0.0994)
South	0.0121	0.0421	-1.0256	-0.1577	0.1585	0.0602	-0.0148	0.0677	0.5443
	(0.0029)	(0.0069)	(0.1862)	(0.0126)	(0.0110)	(0.0102)	(0.0050)	(0.0095)	(0.0989)
West	0.0131	0.0373	-1.2843	-0.1773	0.1558	0.0639	-0.0129	0.0704	0.3595
	(0.0029)	(0.0070)	(0.1925)	(0.0128)	(0.0111)	(0.0104)	(0.0050)	(0.0095)	(0.1010)
BCBS									
Northeast	0.0145	0.0417	-0.9800	-0.1862	0.1446	0.0617	-0.0167	0.0585	0.6407
	(0.0029)	(0.0069)	(0.1845)	(0.0126)	(0.0110)	(0.0102)	(0.0050)	(0.0096)	(0.0985)
Midwest	0.0138	0.0334	-0.9288	-0.1799	0.1453	0.0754	-0.0330	0.0560	0.5706
	(0.0029)	(0.0069)	(0.1849)	(0.0125)	(0.0110)	(0.0101)	(0.0050)	(0.0095)	(0.0982)
South	0.0117	0.0502	-0.9722	-0.1736	0.1502	0.0251	-0.0298	0.0188	0.5542
	(0.0029)	(0.0071)	(0.1865)	(0.0126)	(0.0111)	(0.0103)	(0.0050)	(0.0096)	(0.0990)
West	0.0129	0.0411	-1.2165	-0.1756	0.1415	0.0527	-0.0210	0.0334	0.5842
	(0.0028)	(0.0069)	(0.1868)	(0.0126)	(0.0111)	(0.0102)	(0.0051)	(0.0097)	(0.0988)
Humana									
Northeast	0.0146	0.0384	-1.0963	-0.1323	0.1345	0.0784	-0.0254	0.0119	0.5822
	(0.0029)	(0.0070)	(0.1870)	(0.0128)	(0.0113)	(0.0104)	(0.0053)	(0.0098)	(0.0994)
Midwest	0.0146	0.0493	-1.3379	-0.1497	0.1332	0.0751	-0.0290	0.0109	0.5844
	(0.0029)	(0.0069)	(0.1847)	(0.0126)	(0.0111)	(0.0102)	(0.0051)	(0.0095)	(0.0981)
South	0.0136	0.0430	-1.2088	-0.1655	0.1404	0.0616	-0.0310	-0.0011	0.5524
	(0.0030)	(0.0069)	(0.1852)	(0.0126)	(0.0110)	(0.0102)	(0.0051)	(0.0096)	(0.0985)
West	0.0133	0.0391	-1.2410	-0.1571	0.1477	0.0626	-0.0291	0.0055	0.6271
	(0.0028)	(0.0069)	(0.1857)	(0.0126)	(0.0111)	(0.0102)	(0.0051)	(0.0096)	(0.0986)
Kaiser									
Northeast	*Kaiser had no presence in the Northeast region*								
Midwest	0.0138	0.0364	-1.0883	-0.1550	0.1620	0.0643	-0.0144	0.0641	-0.3968
	(0.0028)	(0.0067)	(0.1790)	(0.0121)	(0.0107)	(0.0098)	(0.0049)	(0.0093)	(0.0953)
South	0.0145	0.0314	-0.8989	-0.1589	0.1722	0.0681	-0.0368	0.0426	0.4715
	(0.0029)	(0.0069)	(0.1992)	(0.0125)	(0.0109)	(0.0104)	(0.0062)	(0.0101)	(0.1006)
West	0.0112	0.0321	-0.3409	-0.1245	0.1706	0.0812	-0.0187	0.0160	0.5463
	(0.0028)	(0.0068)	(0.1822)	(0.0124)	(0.0108)	(0.0100)	(0.0049)	(0.0095)	(0.0968)
UHG									
Northeast	0.0140	0.0345	-1.0138	-0.1491	0.1361	0.0237	-0.0149	0.0717	0.5633
	(0.0029)	(0.0071)	(0.1865)	(0.0127)	(0.0112)	(0.0103)	(0.0051)	(0.0096)	(0.0995)
Midwest	0.0139	0.0294	-1.0690	-0.1390	0.1399	0.0229	-0.0148	0.0715	0.6249
	(0.0029)	(0.0068)	(0.1852)	(0.0125)	(0.0111)	(0.0103)	(0.0050)	(0.0095)	(0.0976)
South	0.0095	0.0313	-1.2485	-0.1406	0.1357	0.0280	-0.0135	0.0676	0.5959
	(0.0029)	(0.0069)	(0.1854)	(0.0127)	(0.0111)	(0.0103)	(0.0050)	(0.0096)	(0.0985)
West	0.0109	0.0318	-1.3601	-0.1595	0.1475	0.0161	-0.0171	0.0611	0.5539
	(0.0028)	(0.0068)	(0.1838)	(0.0126)	(0.0111)	(0.0101)	(0.0050)	(0.0095)	(0.0982)
Other firms									
Northeast	0.0141	0.0401	-1.3511	-0.1969	0.1419	0.0524	-0.0174	0.0547	0.6398
	(0.0029)	(0.0068)	(0.1841)	(0.0125)	(0.0110)	(0.0101)	(0.0050)	(0.0095)	(0.0986)
Midwest	0.0144	0.0405	-1.4959	-0.1984	0.1497	0.0454	-0.0252	0.0522	0.5342
	(0.0029)	(0.0070)	(0.1843)	(0.0126)	(0.0110)	(0.0101)	(0.0051)	(0.0096)	(0.0979)
South	0.0226	0.0340	-1.3084	-0.1809	0.1448	0.0527	-0.0164	0.0486	0.5545
	(0.0029)	(0.0068)	(0.1850)	(0.0126)	(0.0110)	(0.0102)	(0.0050)	(0.0096)	(0.0986)
West	0.0137	0.0310	-1.3560	-0.1809	0.1413	0.0506	-0.0151	0.0521	0.6238
	(0.0029)	(0.0069)	(0.1846)	(0.0126)	(0.0110)	(0.0102)	(0.0050)	(0.0095)	(0.0983)
Observations	37,974	37,974	37,974	37,974	37,974	37,974	37,974	37,974	37,974
R-squared	0.1696	0.0920	0.2518	0.3016	0.0692	0.1799	0.1429	0.3379	0.1039

Notes: This table reports the results of multiple policy function regressions. The dependent variable in each column is the product characteristic. The independent variables include the market-level benchmark (measured in thousands of 2017 dollars) interacted with firm-by-Census-Region indicators, the relevant cost shock, the last period demand unobservable, and market average demographics. Estimates are formed via OLS. No plans in our sample changed basic drug coverage. Robust standard errors are in parentheses.

Table E.2: Mean county characteristics by benchmark quartile, 2017 policy and optimal policy

2017 policy	**0-25th**	**26-50th**	**51-75th**	**76-100th**
Risk-adj. TM costs per capita ($)	9,472	9,830	10,225	11,047
Average risk score	.991	.978	1.00	1.01
Beneficiaries	43,833	43,964	72,933	98,264
Median household income ($)	53,857	52,610	58,314	61,417
Percent in deep poverty, 65+	2.53	2.75	2.50	2.76
Unemployment rate	5.65	5.77	5.33	5.19
Population density (per mi^2)	755	1,280	1,098	1,409
Urban/Rural continuum code	2.90	3.52	2.25	2.25
Resources per 10,000 people				
MDs	21.3	20.0	22.7	24.5
Medicare hospitals	.028	.036	.045	.035
Skilled nursing facilities	.573	.679	.571	.590
Hospice facilities	.132	.190	.113	.101
Medicare hospital readmission rate	17.5	17.3	17.6	18.1
Preventable hospital admission rate	52.0	52.7	51.6	55.0
2017 benchmark ($)	9,318	9,660	9,868	10,421
Number of MA plans	14.8	13.1	17.0	17.5
Number of MA firms	7.2	7.0	9.0	9.3
Observations	110	110	110	109
Optimal policy	**0-25th**	**26-50th**	**51-75th**	**76-100th**
Risk-adj. TM costs per capita ($)	9,037	9,776	10,295	11,458
Average risk score	1.000	.987	.990	1.000
Beneficiaries	35,186	55,237	66,915	101,370
Median household income ($)	49,127	53,853	57,011	66,438
Percent in deep poverty, 65+	2.78	2.53	2.57	2.66
Unemployment rate	6.04	5.51	5.28	5.11
Population density (per mi^2)	354	775	1,081	2,328
Urban/Rural continuum code	3.45	2.64	2.55	2.28
Resources per 10,000 people				
MDs	16.8	23.4	20.8	27.5
Medicare hospitals	.030	.032	.046	.037
Skilled nursing facilities	.647	.522	.641	.604
Hospice facilities	.153	.134	.137	.112
Medicare hospital readmission rate	17.4	17.3	17.6	18.2
Preventable hospital admission rate	52.6	52.0	53.3	53.3
Optimal benchmark ($)	9,117	9,786	10,261	11,168
Number of MA plans	13.3	15.9	15.6	17.5
Number of MA firms	6.7	8.8	8.6	8.4
Observations	110	110	109	110

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 2017 policy. The bottom panel defines quartiles according to the optimal policy.

Table E.3: State-level surplus and expenditures changes from 2017 policy to optimal benchmark schedule

State	# counties in sample	Sum(weight) (000,000s)	Consumer surplus ($ M) 2017	Consumer surplus ($ M) Optimal	Consumer surplus ($ M) %Δ	Government expenditures ($ M) 2017	Government expenditures ($ M) Optimal	Government expenditures ($ M) %Δ
Alabama	14	11.67	98.8	85.3	-13.6	11,444	11,385	-0.51
Arizona	5	11.76	179.4	273.3	52.4	10,894	10,841	-0.49
Arkansas	3	5.25	70.7	24.0	-66.1	4,788	4,697	-1.91
California	17	33.87	776.0	1,641.3	111.5	39,505	39,788	0.72
Colorado	8	3.85	85.3	180.1	111.3	3,711	3,761	1.35
Connecticut	4	6.19	81.0	341.8	322.0	6,914	7,046	1.91
District of Columbia	1	1.43	2.1	9.4	356.6	1,769	1,762	-0.42
Florida	26	36.69	737.5	734.7	-0.4	40,803	40,802	0.00
Georgia	18	11.52	89.2	200.5	124.7	11,198	11,252	0.49
Illinois	15	13.55	88.1	260.9	196.1	14,830	14,941	0.75
Indiana	3	0.66	2.9	3.7	25.0	647	645	-0.25
Iowa	4	2.71	4.7	7.1	50.6	2,347	2,349	0.09
Kansas	2	3.27	17.0	34.7	104.6	3,099	3,116	0.54
Kentucky	12	7.56	57.7	123.3	113.6	7,656	7,687	0.41
Louisiana	7	3.97	112.8	97.4	-13.7	4,328	4,303	-0.56
Maryland	8	4.89	3.7	90.7	2,338.5	5,672	5,654	-0.32
Massachusetts	9	8.52	59.7	541.8	807.2	10,226	10,290	0.62
Michigan	29	22.92	227.2	516.8	127.4	24,483	24,520	0.15
Minnesota	14	8.29	44.5	61.5	38.2	8,219	8,209	-0.12
Missouri	15	8.48	104.8	93.5	-10.8	8,444	8,404	-0.47
Nebraska	6	2.56	11.8	70.9	498.4	2,669	2,658	-0.40
Nevada	2	6.16	145.3	145.9	0.4	7,042	7,043	0.01
New Hampshire	16	16.74	97.2	556.6	472.7	20,358	20,418	0.29
New Jersey	6	9.49	158.7	159.5	0.5	7,090	7,091	0.02
New Mexico	26	26.91	489.5	1,068.9	118.4	30,842	30,439	-1.31
New York	18	18.61	289.9	277.9	-4.1	19,189	18,934	-1.33
North Carolina	24	17.45	219.2	586.0	167.3	17,809	18,107	1.68
Ohio	6	3.85	8.4	8.4	-0.3	3,692	3,691	-0.04
Oklahoma	1	0.02	0.2	0.3	55.2	14	14	0.84
Pennsylvania	25	19.45	385.1	864.4	124.5	20,434	20,722	1.41
Rhode Island	7	3.09	12.2	18.3	50.2	2,935	2,934	-0.04
South Carolina	1	0.02	0.1	0.1	165.3	19	20	0.35
South Dakota	12	10.38	112.9	102.7	-9.0	9,964	9,925	-0.40
Texas	37	29.23	312.1	695.5	122.8	31,612	31,728	0.36
Utah	1	0.01	0.3	0.3	0.0	13	13	0.00
Vermont	1	0.02	0.1	0.3	432.1	16	17	0.79
Virginia	11	7.38	45.5	59.0	29.6	7,196	7,206	0.14
Washington	8	21.23	228.1	156.0	-31.6	18,715	18,538	-0.95
West Virginia	5	3.71	29.5	65.7	123.1	3,527	3,539	0.34
Wisconsin	11	15.81	308.6	90.4	-70.7	14,444	14,083	-2.50
Wyoming	1	0.84	1.7	6.9	317.2	696	699	0.45
Total	439	420	5,700	10,256	79.9	439,254	439,272	0.00

Notes: The MCBS uses a sample of counties and weights observations to be nationally representative; the first column reports the number of counties included in the MCBS in each state, and the second column reports the total MCBS sample weight in the state. Consumer surplus is calculated via Equation (9). Government expenditures include expenditures on TM and MA and are calculated via Equation (16). Surplus and expenditure statistics are calculated using MCBS sample weights.

Table E.4: Market shares and total variable profits for selected firms under 2017 policy and alternative policies

	(1) 2017 policy	(2) Optimal policy	(3) Linear rule	(4) 0.99 CS - 0.01 Var	(5) 0.999 CS - 0.001 Var	(6) Min. *GovExp*	(7) Min exp. w/ floor
Aetna							
Shr. of Medicare benes. (%)	1.16	2.03	2.19	.928	2.15	1.39	1.58
Shr. of MA enrollees (%)	3.96	4.94	5.50	5.20	4.93	5.18	4.50
Total var. profits ($ bill.)	.136	.286	.302	.133	.305	.196	.211
Blue Cross Blue Shield							
Shr. of Medicare benes. (%)	4.17	7.01	6.77	2.73	7.6	4.13	5.42
Shr. of MA enrollees (%)	14.3	17.0	17.0	15.3	17.5	15.4	15.5
Total var. profits ($ bill.)	.388	.836	.812	.285	.816	.497	.581
Humana							
Shr. of Medicare benes. (%)	5.94	7.87	7.40	3.75	9.14	4.75	6.73
Shr. of MA enrollees (%)	20.3	19.1	18.6	21.0	21.0	17.7	19.2
Total var. profits ($ bill.)	.529	.735	.675	.351	.825	.483	.600
Kaiser Permanente							
Shr. of Medicare benes. (%)	1.66	2.41	2.41	.917	2.11	1.83	1.86
Shr. of MA enrollees (%)	5.67	5.86	6.08	5.14	4.84	6.81	5.29
Total var. profits ($ bill.)	.354	.627	.647	.181	.496	.407	.415
UnitedHealth Group							
Shr. of Medicare benes. (%)	6.02	7.53	7.03	3.29	7.98	5.35	7.15
Shr. of MA enrollees (%)	20.6	18.3	17.7	18.4	18.3	20.0	20.4
Total var. profits ($ bill.)	.877	1.18	1.09	.509	1.21	.851	1.06
All Others							
Shr. of Medicare benes. (%)	10.3	14.3	13.9	6.23	14.6	9.38	12.4
Shr. of MA enrollees (%)	35.1	34.7	35.1	34.9	33.5	35.0	35.2
Total var. profits ($ bill.)	1.23	1.97	1.89	.763	1.89	1.26	1.59
Total							
Shr. of Medicare benes. (%)	29.2	41.1	39.7	17.8	43.6	26.8	35.1
Total var. profits ($ bill.)	3.51	5.63	5.42	2.22	5.54	3.7	4.45

Notes: Variable profits are computed via Equation (11) using equilibrium prices and estimated marginal costs adjusted for changes in product characteristics under alternative policies. The entries for Blue Cross Blue Shield include all members of the Blue Cross Blue Shield Association. All statistics are weighted by the MCBS sample weights.

Table E.5: The profitability of a 1% deviation from predicted product characteristics at the optimal policy

...leads to a X% change in variable profit relative to the same deviation at the 2017 policy.	*An 1% increase in the characteristic at the optimal policy...*								
	Enhcd. drug	Dental	Vision	Hearing	ξ	Part B reduc.	Deduct.	Prim. copay	Hosp. copay
All markets									
Mean	-0.0055	-0.0035	0.0007	-0.0053	-0.0015	0.0001	0.0006	0.0043	0.0073
25th Percentile	-0.0077	-0.0083	-0.0044	-0.0062	-0.0016	-0.0003	-0.0001	-0.0017	-0.0007
Median	-0.0000	0.0000	0.0000	-0.0001	-0.0000	0.0000	0.0000	0.0000	0.0001
75th Percentile	0.0045	0.0003	0.0061	0.0008	0.0003	0.0004	0.0007	0.0051	0.0087
284 markets with increases									
Mean	-0.0058	-0.0060	-0.0003	-0.0093	-0.0027	-0.0001	0.0011	0.0082	0.0126
25th Percentile	-0.0140	-0.0076	-0.0095	-0.0134	-0.0040	-0.0010	-0.0000	-0.0007	-0.0000
Median	-0.0001	-0.0003	0.0000	-0.0003	-0.0001	-0.0000	0.0001	0.0003	0.0007
75th Percentile	0.0099	0.0000	0.0098	0.0001	0.0001	0.0003	0.0016	0.0121	0.0173
155 markets with decreases									
Mean	-0.0045	0.0033	0.0035	0.0055	0.0017	0.0006	-0.0007	-0.0063	-0.0073
25th Percentile	-0.0015	0.0000	-0.0003	0.0000	0.0000	-0.0000	-0.0005	-0.0051	-0.0053
Median	-0.0000	0.0001	0.0000	0.0001	0.0001	0.0000	-0.0000	-0.0002	-0.0002
75th Percentile	0.0009	0.0018	0.0025	0.0033	0.0015	0.0004	-0.0000	-0.0000	-0.0000
60 markets with changes greater than $1,000									
Mean	-0.0134	-0.0094	0.0025	-0.0133	-0.0043	-0.0009	0.0019	0.0131	0.0204
25th Percentile	-0.0533	-0.0174	-0.0560	-0.0291	-0.0111	-0.0051	-0.0000	-0.0085	-0.0008
Median	0.0001	-0.0018	0.0002	-0.0052	-0.0015	-0.0000	0.0007	0.0045	0.0076
75th Percentile	0.0529	0.0017	0.0588	0.0051	0.0010	0.0026	0.0036	0.0363	0.0378
379 markets with changes less than $1,000									
Mean	-0.0042	0.0007	0.0004	-0.0041	-0.0010	0.0003	0.0004	0.0029	0.0052
25th Percentile	-0.0048	-0.0044	-0.0025	-0.0036	-0.0009	-0.0002	-0.0001	-0.0014	-0.0006
Median	-0.0000	0.0000	0.0000	-0.0000	-0.0000	0.0000	0.0000	-0.0000	0.0000
75th Percentile	0.0026	0.0061	0.0041	0.0006	0.0003	0.0003	0.0004	0.0029	0.0050

Notes: For each plan, we calculate the percentage change in the variable profit earned by the firm that offers the plan in response to an increase of 1% in each of the product characteristics. We update marginal costs and hold prices and all other product characteristics fixed for all plans in the market. To isolate the effect of our approximation approach, we calculate the percentage difference both at the optimal policy and at the 2017 policy, and subtract the two.

Figure E.1: Medicare Advantage Benchmark Distribution, 2017

(a) Benchmarks across counties

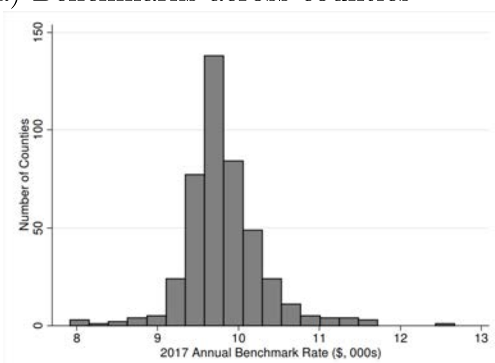

(b) Benchmarks across beneficiaries

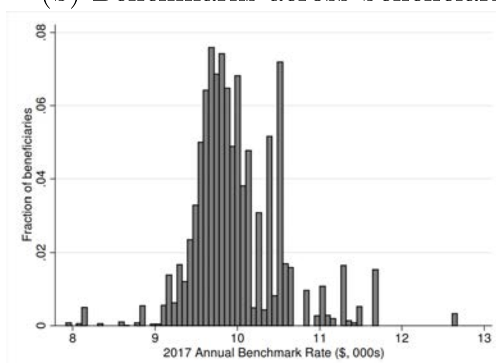

Note: Includes only those counties included in the 2017 Medicare Current Beneficiary Survey.