

NBER WORKING PAPER SERIES

THE OPTIMAL GEOGRAPHIC DISTRIBUTION OF 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 June 2026

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, UCLA, Harvard/MIT, Northwestern University, University of Pennsylvania, University of Maryland, New York University, 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.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Keaton S. Miller, Amil Petrin, Robert Town, and Michael Chernew. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Optimal Geographic Distribution of Managed Competition Subsidies
Keaton S. Miller, Amil Petrin, Robert Town, and Michael Chernew
NBER Working Paper No. 25616
February 2019, Revised June 2026
JEL No. I11, L13, L51

ABSTRACT

Governments often subsidize the purchase of goods or services from competing firms rather than producing them directly. One common mechanism they use is ‘managed competition’, where firms compete for consumers and a subsidy that follows them. We introduce a framework for determining the optimal market-level subsidy schedule that features heterogeneous consumers and firms that set prices and product characteristics in response to 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 product characteristic policy functions empirically and solving for Nash equilibria in prices. The consumer-welfare-maximizing budget-neutral schedule increases aggregate annual consumer welfare by \$7 billion over the current policy.

Keaton S. Miller
University of Oregon
Department of Economics
keatonm@uoregon.edu

Amil Petrin
University of Minnesota
Department of Economics
and NBER
petrin@umn.edu

Robert Town
University of Texas at Austin
Department of Economics
and NBER
robert.town@austin.utexas.edu

Michael Chernew
Harvard University
Harvard Medical School
and NBER
chernew@hcp.med.harvard.edu

1 Introduction

Under “managed competition”, the government does not provide public services directly – rather it subsidizes consumers to purchase them from competing private firms (Laffont and Tirole, 1993). The economic rationale behind this procurement mechanism is that competition among firms disciplines costs and promotes efficient product offerings. Central to the performance of managed competition is the setting of subsidies. In practice, subsidies are often based on coarse measures of the government’s cost of providing the service that ignore heterogeneity in demand and production across suppliers. This approach to setting subsidies can result in substantial welfare loss.

In this paper, we develop an approach for calculating the optimal managed competition subsidy policy across geographically heterogeneous markets. We take the government’s budget constraint as exogenous and model firms as multiproduct oligopolists who choose prices and other product characteristics in equilibrium responses to the subsidy. We tractably endogenize non-price characteristics, which is important because the welfare impact of policy interventions differs when non-price characteristics can respond (Spence, 1975, 1976).¹

We apply this framework to the Medicare Advantage (MA) program in the United States. Under MA, beneficiaries can opt out of Traditional Medicare (TM) fee-for-service and enroll in privately administered plans that compete on premiums, benefit design, and provider networks. Plans receive risk-adjusted, per-enrollee payments tied to a county-specific benchmark rate. Benchmark rates vary substantially across counties, generating meaningful geographic heterogeneity in market conditions. Because these benchmark payments directly affect plan revenues, they shape firms’ incentives over both pricing and product design, making them a central policy lever for enrollee welfare and government spending. More broadly, this institutional structure—where private firms compete for government subsidies in the provision of public services—appears in other settings, including other health insurance programs (Gruber, 2017), education (Hoxby, 2000), and housing (Baum-Snow and Marion, 2009) suggesting that our approach has wider applicability.

Our key methodological contribution is to develop an approach to compute equilibria

¹Appendix A provides a simple example showing the importance of modeling non-price characteristics.

under counterfactual subsidies when firms endogenously choose a high dimensional set of product characteristics. The traditional approach is to solve firms’ best-response functions (Fan, 2013, Wollmann, 2018, Vatter, 2025, Barahona et al., 2023). That approach is impractical here due to the number of markets, plans, and product characteristics. Holding characteristics fixed and solving for price alone is also insufficient since many plans charge a \$0 premium making the other characteristics essential to product differentiation. Instead, we estimate policy functions for product characteristics from the data, predict characteristics under counterfactual benchmarks, and solve firms’ first-order conditions for premiums. We provide Monte Carlo evidence that this approach well-approximates the equilibria.²

Our empirical analysis starts by estimating Medicare beneficiaries’ preferences over plan characteristics using detailed individual panel data with information on demographics, choice sets, and realized choices from 2008–2017. The model allows for endogenous plan characteristics and beneficiary switching costs. The estimates are generally sensible. Our estimates imply a mean premium elasticity of -6.29 . Premium sensitivity decreases with income, MA is more attractive to younger consumers, non-Whites, and those with lower educational attainment. The estimated average switching cost between MA insurers is \$285 (compare to the average annual premium of \$415).

We then invert firms’ first-order conditions to infer marginal costs. The implied marginal costs align with cost estimates from other, independent sources. With preferences and plan cost information in hand we then assess welfare. In 2017, MA generated \$3.18 billion in consumer surplus and \$2.13 billion in variable profit, with \$126 billion in government payments and \$314 billion in TM spending. The last building block needed to estimate the optimal subsidies are the policy functions. Policy function estimates indicate that increasing the benchmark increases benefit generosity and reduces out-of-pocket expenses. The policy function estimates also imply that, on average, MA subsidies are too generous at the margin: an across-the-board \$1 decrease in benchmarks would decrease consumer welfare by \$8 million but save \$9 million.

Our paper concludes by presenting the estimates of the optimal geographic distribution of benchmarks. The optimal benchmark schedule meaningfully differs from the 2017 schedule

²See Appendix D.

by an average of \$537 (about 5%). The optimal policy increases aggregate consumer welfare to \$10.2 billion per year by increasing the total MA share from 29.8% to 47.0%. The mean compensating variation for MA enrollees increases from \$254 to \$515. Changes in product characteristics are responsible for roughly 20% of the total change in consumer welfare. There are winners and losers – relative to the 2017 schedule, the optimal policy reduces welfare for 37% of beneficiaries. We also explore other social welfare functions and find benchmark schedules that increase aggregate welfare while reducing the variance in the geographic distribution of the benchmark.

1.1 Literature review

Our paper builds on a large literature on Medicare Advantage (MA); see McGuire et al. (2011) for a review. It is most closely related to Town and Liu (2003) and Curto et al. (2021), who estimate demand for MA plans and use these estimates to recover consumer surplus and insurer profits. We extend this framework along several dimensions. First, we incorporate richer demographic and plan-level data. Second, we endogenize both premiums and non-premium characteristics, allowing firms to adjust multiple margins of competition. Third, we evaluate counterfactual subsidy policies. These differences are quantitatively important: we estimate larger premium elasticities and correspondingly lower consumer surplus than Curto et al. (2021), and we find that uniform increases in benchmark payments reduce welfare, in contrast to their findings.

A related literature studies the pass-through of benchmark payments to consumers (Cabral et al., 2018, Duggan et al., 2016, Song et al., 2013). We complement this work by modeling firms’ plan design decisions and estimating consumer valuations over both price and non-price attributes, capturing additional channels through which subsidies affect welfare.

Our analysis also contributes to the literature on optimal subsidy design in insurance markets (Tebaldi, 2025, Jaffe and Shepard, 2020, Einav et al., 2026, Decarolis et al., 2020, Polyakova and Ryan, 2023) and on premium regulation (Bundorf et al., 2012, Ericson and Starc, 2015). Relative to this literature, we emphasize the role of non-premium characteristics and study how optimal subsidies shift when firms adjust these attributes in equilibrium.

Finally, our counterfactual analysis builds on approaches that use estimated policy func-

tions to evaluate welfare effects of product changes (Goolsbee and Petrin, 2004, Sweeting, 2013). We extend this approach by integrating policy functions with equilibrium pricing and product design, allowing us to characterize the full equilibrium response to alternative subsidy policies.

2 The Medicare Advantage Program

Medicare Advantage (MA) traces its origins to the risk-based contracting experiments that paid private insurers capitated rates to cover Medicare beneficiaries as an alternative to fee-for-service Medicare. The modern program emerged with the Balanced Budget Act of 1997, which created “Medicare+Choice” to expand private plan participation. The program was rebranded as Medicare Advantage and expanded under the Medicare Modernization Act of 2003, which increased payments and integrated prescription drug coverage. Under MA, private insurers contract with the Centers for Medicare and Medicaid Services (CMS) to cover beneficiaries who forgo Traditional Medicare (TM) fee-for-service for a managed-care-style plan. MA plans receive a per-enrollee payment that varies with observable risk (diagnoses and demographics) but not realized expenditures. In 2023, MA enrollment reached 50% of the Medicare population. In our 2017 data (the last year of our analysis), payments to plans were \$126 billion and TM spending on individuals totaled \$314 billion.

Medicare beneficiaries typically enroll in MA plans during an open enrollment period each fall. Both TM and MA enrollees pay a mandatory ‘Part B’ premium, and MA enrollees may pay an additional premium. Insurers compete via benefits, premiums, provider networks, and marketing (Aizawa and Kim, 2018). MA plans are required to provide at least TM-equivalent benefits, and generally provide a more generous benefit package, such as including dental, vision, hearing, and prescription drug coverage and reduced cost sharing.

The MA market is relatively concentrated. The top five firms nationwide, Aetna/CVS, the Blue Cross Blue Shield Association, Humana, Kaiser Permanente, and UnitedHealth Group, account for 65% of total enrollment. In our 2017 data, the average beneficiary had access to 10 plan options; 64% of beneficiaries had access to 5 or more plans, and 25% of beneficiaries had access to 3 or fewer plans.

2.1 Benchmarks

MA subsidies are based on a “benchmark” rate which varies across counties and over time (Newhouse et al., 2012). CMS updates the benchmarks annually using a nonlinear formula applied to the county-level historical moving average of risk-adjusted per-capita TM spending. The benchmark for counties in the top quartile of average FFS spending is set to 95% of spending. The benchmark for the second quartile is 100% of 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. We denote county markets by m and years by t , and we use B_{mt} to denote the market-year level benchmark.

Figure 1 illustrates the 2017 benchmark policy and the resulting market outcomes with county-level maps of the US. The left map illustrates the ratio of the benchmark to average FFS spending in 2017. The right map illustrates the total MA share in each county. If the ratio of the subsidy to costs were the sole determinant of enrollment, these two maps would be closely correlated. The imperfect correspondence suggests that there may be welfare gains from reallocating subsidies across markets.

2.2 Bids

Each year, after benchmarks are published by CMS, insurers submit their MA plan ‘bids’ detailing their benefit and cost-sharing designs. For a plan j , the ‘bid amount’ b_{jt} is an offer to provide services in chosen counties to an average-risk enrollee in exchange for per-enrollee revenue that must be related to the plan’s projected costs. Insurers that bid at or above the benchmark receive only the benchmark as government payment and must charge premiums to enrollees for the difference between the bid and the benchmark. Firms that bid below the benchmark receive their bid and a portion of the difference as a ‘rebate’ that must be passed on to consumers through decreased cost-sharing or non-TM benefits (e.g. dental). Supplemental benefits may also be added and paid for by an additional premium. In 2017, the average bid was 90% of TM costs (MedPAC, 2017).

The rebate payment varies from 50% to 75% of the benchmark minus the bid based on the CMS ‘star rating’ of insurer quality. The rating summarizes past service performance, including the fraction of enrollees receiving influenza vaccinations and the 30-day hospital

readmittance rate. We use λ_{jt} to denote the plan-year rebate percentage. Under current policy, high quality insurers also receive a 5% bonus to the benchmark rate. We denote this bonus as ϕ_{jt} , and define the effective plan-level benchmark $B_{jt} \equiv B_{mt} \times \phi_{jt}$.³

This bidding system adds complexity to firms' pricing and benefit design problem. Given b_{jt} , the payment from CMS to insurers for enrollee i with risk adjustment factor $Risk_{it}$ is

$$Payment_{imjt} = Risk_{it} \times \begin{cases} B_{jt} & \text{if } b_{jt} \geq B_{jt} \\ (b_{jt} + \lambda_{jt} \times (B_{jt} - b_{jt})) & \text{if } b_{jt} < B_{jt} \end{cases} \quad (1)$$

This setup allows for empirically relevant corner solutions where the rebate and/or annual premium is \$0, which can be difficult to accommodate in standard pricing models. CMS bidding rules, however, can be used to transform the problem into a more tractable form by noting that rebates must fund benefits beyond TM and cost-sharing reductions can *only* be funded through rebates.⁴ Let x_{jt} be the vector of plan design characteristics, and let $x_{c,jt}$ be the subvector restricted to cost-sharing characteristics. Denote $incr_{jt}(x_{c,jt})$ as the 'incremental' cost of providing cost-sharing characteristics $x_{c,jt}$ relative to a TM-equivalent plan. For plans that bid below the benchmark, the CMS rules can now be interpreted as a constraint: $incr_{jt}(x_{c,jt}) \leq \lambda_{jt}(B_{jt} - b_{jt})$.

This constraint binds with equality if plan designers minimize costs given benchmarks and bids. The key friction is that firms capture only a fraction $\lambda_{jt} < 1$ of the gap between the benchmark and the bid. For example, suppose $\lambda = 0.5$ and the plan designer wishes to increase supplemental benefits by \$1. To do so, the firm must lower its bid by \$2, which increases the rebate by \$1. In essence, the implicit tax rate on rebate-funded benefits is $\frac{1-\lambda}{\lambda}$. Benefits funded through premiums face no such tax.

This discussion implies that the bid is determined by the cost-sharing benefits, $x_{c,jt}$, and

³The star rating, and therefore λ and ϕ , are determined at the 'contract' level. A contract generally covers plans of the same type (e.g. HMO or PPO) within a state. For most firms, all plans share a single contract and thus the same λ and ϕ ; larger firms may have multiple contracts with different ratings. Both variables are exogenous from the firms' perspective as CMS sets and frequently redefines the rating criteria, and so we use plan-level notation throughout for simplicity.

⁴In our empirical analysis, "cost-sharing reductions" consist of reductions in the Part B premium, deductibles, and copays.

the incremental cost function, $incr_{jt}$. There is therefore a map from $x_{c,jt}$ onto b_{jt} :

$$b(x_{c,jt}; B_{jt}, \lambda_{jt}, \cdot) = B_{jt} - \frac{1}{\lambda_{jt}} incr_{jt}(x_{c,jt}). \quad (2)$$

The plan-level subsidy is the minimum of the bid plus the rebate and the benchmark: $sub(b_{jt}; B_{jt}, \lambda_{jt}) = \min\{b_{jt} + \lambda_{jt}(B_{jt} - b_{jt}), B_{jt}\}$. Since MA plans' benefits must be at least equivalent to TM, we assume $incr_{jt}(\cdot) \geq 0$. Substituting (2) into the subsidy expression,

$$sub(x_{c,jt}; B_{jt}, \lambda_{jt}, incr_{jt}) = B_{jt} - \frac{1 - \lambda_{jt}}{\lambda_{jt}} incr_{jt}(x_{c,jt}). \quad (3)$$

Note that the subsidy is weakly positive for all plans, including those that charge a positive premium. The maximum subsidy is the benchmark B_{jt} itself.

Insurers receive revenue from two sources, government subsidies and premiums, with bidding and benefit design decisions affecting both revenue sources. Given product characteristics and the insurer's per-enrollee revenue, r_{jt} , the premium faced by consumers is $p_{jt} = r_{jt} - sub(x_{c,jt}; \cdot)$. Thus, under the assumption that plan designers are cost minimizers, the bidding system can be recast as a game in which firms choose per-enrollee revenues (premiums plus government subsidies) and product characteristics. This is analogous to a standard pricing game where firms' per-enrollee revenue is simply the price.⁵ Note that plans for which $r_{jt} = sub(\cdot)$ charge zero premiums; thus in this formulation whether a plan's premium is zero or positive is an endogenous outcome of the firm's revenue and characteristic choices rather than a separate discrete decision.

3 Model

In this section we introduce our model of the MA segment and the approximation method we use to evaluate counterfactual benchmark policies.⁶ We begin by formalizing the social planner's problem, dropping time subscripts where non-essential. The planner chooses county-level benchmarks $\{B_m\}$ to maximize consumer welfare subject to an exogenous bud-

⁵In practice the software used by firms to submit bids to CMS requires firms to choose revenues; bids and premiums are calculated by the software after accounting for the cost data input by firms.

⁶While we develop our framework in the context of Medicare Advantage, the approach applies to any managed competition setting in which a government sets market-level subsidies for differentiated oligopolists. See Section 6 for further discussion.

get constraint G and given the strategic responses of firms. Recall that r_m , p_m , and x_m denote the vectors of per-enrollee revenues (prices plus subsidies), prices, and observable product characteristics in market m . Firms' strategic variables are r and x ; consumer-facing prices are determined residually. Let ξ_m denote unobservable (to the econometrician) product characteristics which are exogenous to the benchmark but potentially correlated with prices and product characteristics. The optimal subsidy problem is

$$\begin{aligned}
& \max_{\{B_m\}} \sum_{m=1}^M CS_m(p_m, x_m; \xi_m) \\
& \text{s.t. } \sum_{m=1}^M GovExp_m(r_m, x_m; B_m, \xi_m) = G; \quad \text{and} \\
& \forall m, (r_m, x_m) \text{ form Nash equilibrium in } \pi_f \text{ given } \xi_m, B_m.
\end{aligned} \tag{4}$$

The solution to this problem maximizes the ‘bang for the buck’ in terms of the welfare returns to government spending. The pass-through rate in each market depends on the demand elasticity, the marginal value of the product characteristics, firms’ costs, and the intensity of competition.

Solving this problem requires finding equilibrium firm policies as functions of B_m , which is generally computationally intractable when firms choose multiple product characteristics. A common approach assumes non-price characteristics are exogenous, but this is restrictive in our context. Over a third of MA products charge zero premiums so firm responses to benchmark changes must account for potential non-price responses. Our solution is to estimate policy functions for product characteristics and solve the Nash equilibrium in revenues taking predicted characteristics as given.

To formalize, suppose firms choose per-enrollee revenues r_j and product characteristics x_j , with premiums and bids determined residually as described in Section 2. Let $\mathcal{J}_f$ be the set of plans offered by firm f which is exogenous. Let r_{-j} indicate the vector of per-enrollee revenue choices for all plans other than j , with an analogous definition for x_{-j} . Denote product shares as a function of prices and product characteristics by $s_j(\cdot)$. Using the

subsidy function $sub(\cdot)$ of (3), plan-level profit is

$$\pi_j(r_j, x_j; r_{-j}, x_{-j}, \xi_m, B_j, \lambda_j) = (r_j - c_j(x_j, \xi_j)) s_j(r_j - sub(x_{c,j}; \cdot), x_j; \cdot) \quad (5)$$

and the profit-maximization problem for firms is⁷

$$\max_{\{r_j, x_j\}_{j \in \mathcal{J}_f}} \pi_f \equiv \sum_{j \in \mathcal{J}_f} \pi_j(r_j, x_j; \cdot). \quad (6)$$

3.1 Policy functions and approximations

The solution to Equation (6) depends on B_j and the determinants of c_j and s_j – the firms’ choices are functions of the subsidy, demand and supply characteristics in the market. In this way, the firms’ product design problem is similar to a traditional production problem (Kim et al., 2016) and our approach is similar to a control function approach. Let $\mathcal{I}_{mt}$ represent these market characteristics. The policy functions are the functions $x_j(B, \mathcal{I})$ and $r_j(B, \mathcal{I})$ such that, when evaluated at the appropriate values, the x_j and r_j form a Nash equilibrium in π_f . We adopt the equilibrium selection assumption of Bajari et al. (2007), ensuring that x_j and r_j are single-valued locally around observed values; see Section 3.4 for discussion.

Our approach is based on estimating an approximation to these local policy functions. Suppose the analyst has data on government policies, market covariates, and firm choices, for several periods denoted by t . We approximate the policy function x_j using a first-order expansion that varies by characteristic, firm, and market. That is, for product j and characteristic l we write

$$\tilde{x}_{l,jt} = \theta_{1,fl} B_{jt} + \theta_{2,lt} B_{jt} M_t + \theta_{3,lt} M_t + F X_l + \eta_{l,jt} \quad (7)$$

Here $\theta_{1,fl}$ captures the firm-specific effect of the benchmark on characteristic l , $\theta_{2,lt}$ captures how market characteristics $M \approx \mathcal{I}$ moderate that effect, $\theta_{3,lt}$ captures level effects of M , and

⁷Selection incentives are a central concern in insurance markets (Geruso and Layton, 2017, Brown et al., 2014). While risk adjustment has mitigated these concerns in MA (Newhouse et al., 2012, 2015), residual selection remains a potential issue. In our data, however, we find little evidence of meaningful selection along these dimensions and so do not incorporate adverse selection into the firm’s profit-maximization problem.

FX_l are firm indicators. The term η_{jt} captures measurement error, approximation error, and characteristic-level cost and demand shocks.⁸

Given parameter estimates $\hat{\theta}$, we compute counterfactual product characteristics by calculating $\hat{\theta}_{lj} \equiv \frac{\partial \hat{x}_{ljt}}{\partial B_{jt}}$ for each product and product characteristic. Suppose the policy we observe in the data is B . Given a counterfactual benchmark B' , we calculate $\hat{x}_{lj}(B') = x_{lj}(B) + \hat{\theta}_{lj}(B' - B)$ where $x_{lj}(B)$ is the observed value of the product characteristic. This implicitly holds the value of M_t fixed and ensures that when we input the policies in the data, our approach recreates the outcomes in the data.

With $\hat{x}_{jt}(B)$ as the vector of predicted non-price characteristics for policy B , and $\hat{\xi}$ as the vector of estimated unobservables, we approximate the firm's problem with

$$\max_{\{r_{jt}\}_{j \in \mathcal{J}_f}} \tilde{\pi}_f \equiv \sum_{j \in \mathcal{J}_f} \pi_{jt}(r_{jt}, \hat{x}_{jt}, \hat{\xi}_{jt}; r_{-jt}, \hat{x}_{-jt}, \hat{\xi}_{-jt}, B). \quad (8)$$

Let $\hat{r}_{mt}$ be the vector of per-enrollee revenues that forms a Nash equilibrium in $\tilde{\pi}_f$ given $\hat{x}, \hat{\xi}$ and B . We are primarily concerned with how the equilibrium moves with the policy B , so, with a slight abuse of notation, we use the tuple $\{\hat{r}_{jt}(B), \hat{x}_{jt}(B)\}$ to evaluate candidate solutions to Equation (4). We discuss limitations and caveats to this approach in Section 3.4.

3.2 Demand and marginal costs

Solving Equation (8) requires defining c_j and s_j , the latter of which requires specifying a demand system. Since insurance beneficiaries are inertial (Aizawa and Kim, 2018), we extend the discrete choice setting of Goolsbee and Petrin (2004) to allow for switching costs.

Let consumer i be described by a vector of potentially time-varying observable characteristics z_{it} . Firms have time- and market-product-invariant quality v_f . Therefore, ξ_{jt} represents product-market- and time-specific factors (e.g. network breadth). Consumers enter period t with product k_{it} , with zero denoting the outside good. We define three switching cost indicators: $W_{1ijt} = \mathbf{1}[k_{it} = 0]$ (TM to MA), $W_{2ijt} = \mathbf{1}[f(j) = f(k_{it}) \wedge k_{it} \neq 0]$ (within-firm), and $W_{3ijt} = \mathbf{1}[f(j) \neq f(k_{it}) \wedge k_{it} \neq 0]$ (across-firm), where $f(j)$ denotes the firm offering j .

Let u_{ijmt} denote the consumer's utility from product j . Let y_{it} be i 's income and h_{it}

⁸The linear form is for simplicity; non-linear and ML alternatives yield similar results.

be demographic indicators that may affect switching costs (e.g. health status). Dropping market subscripts, the choice-specific utility for products is

$$u_{ijt} = (\alpha_0 + \alpha_1 y_{it} + \alpha_2 y_{it}^2) p_{jt} + \sum_{n=1}^3 \beta_{wn} W_{nijt} + \sum_{n=1}^3 \sum_h \beta_{whn} W_{nijt} h_{it} \\ + \beta_z z_{it} + \beta_x x_{jt} + \xi_{jt} + v_f + \beta_\nu \nu_{it} + \epsilon_{ijt}. \quad (9)$$

The α parameters capture income-varying price sensitivity; β_{wn} and β_{whn} capture demographic-dependent switching costs; β_z and β_x capture tastes for demographics and non-price characteristics, respectively.⁹ v_f is a firm fixed effect. $\nu_{it} \sim N(0, 1)$ is an unobservable MA preference and ϵ_{ijt} is an i.i.d. Type-I extreme value taste shock.

The price of the outside good may vary with demographics $p_{0t}(z_i)$. The utility of the outside good (which we normalize by subtracting from each u_{ijt}) is

$$u_{i0t} = (\beta_{o0} + \beta_{o1} y_{it} + \beta_{o2} y_{it}^2) p_{0t}(z_i) + \epsilon_{i0t}. \quad (10)$$

Marginal costs. We assume marginal cost is constant in enrollment after risk adjustment. Let the vectors $x_{o,jt}$ and $x_{c,jt}$ correspond to distinct sets of plan features. The vector $x_{o,jt}$ collects plan-level non-cost-sharing characteristics — basic drug coverage, enhanced drug coverage, dental, vision, and hearing — while $x_{c,jt}$ collects cost-sharing characteristics — the Part B premium reduction, the deductible, the primary-care copay, the specialist copay, and the hospital copay. The logarithm of the marginal cost function is

$$\ln(c_{jt}) = c_{0,jt} + \gamma_{o,jt} x_{o,jt} + \gamma_{c,jt} x_{c,jt} + \omega_{c,jt} \quad (11)$$

where $c_{0,jt}$ is the log cost of providing a TM-equivalent plan, absorbing all plan-level cost heterogeneity not attributable to observable characteristics or unobservable cost-sharing characteristic costs; $\gamma_{o,jt}$ and $\gamma_{c,jt}$ are plan-varying cost coefficient vectors on non-cost-sharing and cost-sharing characteristics, respectively; and $\omega_{c,jt}$ is an unobservable cost associated

⁹In contrast to this specification, Curto et al. (2021) assume that the effect of plan premiums and characteristics on utility is completely captured by the plan's bid relative to the benchmark.

with cost-sharing characteristics.¹⁰

Computing Equation (8) is then straightforward: after computing product characteristics from policies, we calculate firm costs and subsidies. We then solve firms' revenue problems using the fixed-point iteration approach of Morrow and Skerlos (2011).

To connect this cost model to the discussion of Section 2, note that the incremental cost of providing cost-sharing characteristics $x_{c,jt}$ relative to a TM-equivalent plan is

$$incr_{jt}(x_{c,jt}) = \exp[c_{0,jt} + \gamma_{c,jt}x_{c,jt} + \omega_{c,jt}] - \exp[c_{0,jt}]. \quad (12)$$

This is the incremental cost function that enters the subsidy function (3). Note that for reasonable values of $\gamma_{c,jt}$ and β_x , this functional form implies that firms will have an interior solution for the choice of cost-sharing benefits. In what follows, we assume that the data is generated from interior solutions.

3.3 Consumer welfare and government spending

Our consumer welfare measure is compensating variation – the additional income consumer i would need to achieve the same expected utility without MA (Nevo, 2000). Following Berry et al. (1995), we rewrite u_{ijt} into a product-level mean

$$\delta_{jt} = \alpha_0 p_{jt} + \beta_x x_{jt} + v_f + \xi_{jt} \quad (13)$$

and an individual-specific deviation from that mean

$$\begin{aligned} \mu'_{ijt} = & (\alpha_1 y_{it} + \alpha_2 y_{it}^2) p_{jt} + \sum_{n=1}^3 \beta_{wn} W_{nijt} + \sum_{n=1}^3 \sum_h \beta_{whn} W_{nijt} h_{it} + \beta_z z_{it} + \beta_\nu \nu_{it} \\ & - (\beta_{o0} + \beta_{o1} y_{it} + \beta_{o2} y_{it}^2) p_o(z_{it}) + \epsilon_{ijt}. \end{aligned} \quad (14)$$

Let $\mu_{ijt} = \mu'_{ijt} - \epsilon_{ijt}$. Given our distributional assumptions on ϵ_{ijt} and ν_{it} , the probability

¹⁰In contrast, Curto et al. (2021) assume that on the margin plan costs have a one-to-one relationship with plan bids.

that consumer i chooses product j (i.e. the share function) is

$$s_{ijt} \equiv \Pr(i \text{ chooses } j) = \int_{\nu} \frac{\exp(\delta_{jt} + \mu_{ijt}(\nu))}{1 + \sum_{k \in \mathcal{J}_m} \exp(\delta_{kt} + \mu_{ikt}(\nu))} d\nu, \quad (15)$$

and the total share of product j is

$$s_{jt} = \int_{z_{it}} s_{ijt}(z_{it}) dz_{it}. \quad (16)$$

The expected welfare for i is now

$$CS_{it} = \frac{E[\max_j u_{ijt}]}{\alpha_0 + \alpha_1 y_{it} + \alpha_2 y_{it}^2} = \frac{1}{\alpha_0 + \alpha_1 y_{it} + \alpha_2 y_{it}^2} \ln \left(1 + \sum_j \exp(\delta_{jt} + \mu_{ijt}) \right). \quad (17)$$

We report both per-consumer and per-MA enrollee compensating variation, where the latter conditions on choosing an MA plan:

$$\overline{CS}_t^{cond} = \frac{\int CS_{it} di}{\int s_{it} di}, \quad (18)$$

where s_{it} is the probability that consumer i chooses any product in period t . We report the mean compensating variation both per consumer and per purchaser, and the aggregate welfare $CS_t = \int_i CS_{it}$ (Petrin, 2002, Town and Liu, 2003).¹¹ As μ_{ijt} includes switching costs, our surplus estimates are net of those costs.

The total government expenditure on Medicare consists of the sum of the MA payments given by $sub(\cdot)$ and TM spending. 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 $Risk_{mt}$. Thus, dropping time subscripts,

$$GovExp_i(B_m) = \sum_j s_{ij} sub_j(B_j; \cdot) Risk_m + \left(1 - \sum_j s_{ij} \right) TM_m Risk_m. \quad (19)$$

¹¹This formulation of surplus assumes that our demand model holds for inter- and infra-marginal consumers (McFadden, 1974).

3.4 Caveats and limitations

Our approach relies on several notable assumptions that, if violated, could bias the results. One important concern is the potential existence of multiple equilibria. If the data are generated by behavior reflecting different strategy profiles from different equilibria, our policy function estimates will be biased. We adopt the equilibrium selection assumption of Bajari et al. (2007)—that observations come from a single strategy profile—and test it by re-estimating policy functions across sub-national geographies and levels of the current policy. We cannot reject the null of a single strategy profile.¹² We compute counterfactual Nash equilibria using the fixed-point method of Morrow and Skerlos (2011) starting from observed choices.

We make several additional assumptions that could bias our results if violated. First, we assume that the market-level variables, M , capture the information firms condition on when making decisions. If M is incomplete, omitted-variable bias may arise (Sweeting, 2013). Second, we treat ξ as exogenous.¹³ If firms also adjust ξ in response to benchmark changes, our model would attribute too much of the supply-side response to prices and observed characteristics. The implications for aggregate compensating variation are ambiguous: we would overstate the welfare gains from lower prices while understating the gains from improvements in unobserved quality, and these biases may partially offset. Finally, we abstract from product entry and exit. We do not view this as a major limitation because products entering or exiting account for less than 1% of market share, and the optimal policy keeps benchmark payments within the observed range.

Our approach approximates counterfactual outcomes and is not necessarily a strict improvement over methods that solve firms’ full optimization problem. Our goal is a feasible approximation when full solutions are intractable. In Appendix D, we explore the accuracy of the approximation in a Monte Carlo exercise and find that error increases with model complexity and decreases with sample size. In Appendix E, we explore endogenizing one additional characteristic in a reduced set of markets.

¹²These tables are omitted due to space constraints; contact the corresponding author for details.

¹³In our setting, an important component of ξ is the local provider network. This assumption implies that networks are fixed at the time products are designed.

4 Data and Estimation

We use public CMS files for all MA plans offered 2008–2017. For each plan we collect county-level enrollment, premiums, the Part B premium reduction, in-network copayment rates for primary care and specialist visits and 7-day hospital stays, the star rating, and CMS-defined indicators for dental, vision, hearing, and drug (basic and expanded) coverage; coinsurance is converted to copayments using the Medicare Physician Fee Schedule. Risk-adjusted payments and benchmarks yield each plan-county’s bid via the payment formula. Enrollment counts and CMS eligibility data yield plan-county-year shares. CMS releases cost estimates submitted during the bid process with a five-year delay; we obtain 2008–2015 and focus on plans’ reported cost of providing TM-equivalent coverage.

We focus on the individual MA market of Section 2, and drop employer-sponsored plans and plans for “dual-eligible” Medicare/Medicaid beneficiaries, which operate under different payment and benefit structures. Per CMS rules, we drop plan-counties with ten or fewer enrollees. We also drop plans outside our micro-data sample area.¹⁴

Table 1 summarizes these data by county-year benchmark quartiles. As benchmarks increase, benefits generally improve. The zero-premium share increases from 0.336 in the first quartile to 0.486 in the fourth. These patterns are not always monotonic (e.g. deductibles) which reflects the fact that benchmarks are set as a function of past average TM costs and private insurer costs are imperfectly correlated with TM costs. Indeed, the difference between the benchmark and plan TM costs increases in the benchmark.

We also use the 2008–2017 Medicare Current Beneficiary Survey (MCBS), a rolling-panel survey of Medicare beneficiaries linked to CMS administrative data. We observe demographics (income, age, sex, race, education, county), self-reported health status, and MA plan choices. We use MCBS sample weights throughout; after minor cleaning, the weighted sample is within 2% of the CMS-reported eligible population. To price the outside option, we obtain Medigap Plan C rates from Weiss Ratings which we match to each person-year. Plan C covers most TM cost-sharing and is the most popular Medigap plan.¹⁵

¹⁴Firms may sign multiple contracts with CMS to offer plans in different coverage areas; we abstract from contract identifiers and focus on firms.

¹⁵Massachusetts, Minnesota, and Wisconsin have alternative plan definitions; we use the plan closest to Plan C. Additionally, United Healthcare did not offer plans in New York during our study period. For

Table 2 reports summary statistics. MA enrollees have lower income, are less likely to be White, and have lower educational attainment; TM-to-MA switchers are generally similar to the larger group of MA enrollees.

4.1 Demand 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 (13) with an instrumental variables approach.

Let $\Psi = \{\alpha_1, \alpha_2, \beta_w, \beta_{wh}, \beta_z, \beta_\nu, \beta_o\}$ be the set of parameters which determine μ'_{ij} . For a candidate value $\tilde{\Psi}$ we use the Berry (1994) inversion with the Berry et al. (1995) contraction mapping to compute the unique set of product mean utilities $\delta_j(\tilde{\Psi})$ that match predicted shares to the observed shares. Let C_{ijt} be an indicator variable that is equal to one if person i chose product j at time t . The likelihood function is

$$L_{it}(C_{ijt}; \tilde{\Psi}, \delta(\tilde{\Psi})) = \prod_j s_{ijt}^{C_{ijt}}, \quad (20)$$

where s_{ijt} is given by Equation (15). In the first stage of our estimation procedure, we apply the MCBS sample weights w_{it} and maximize the weighted log likelihood function

$$l(C; \tilde{\Psi}) = \sum_t \sum_i \ln(L_{it}) w_{it}. \quad (21)$$

At the estimate $\hat{\Psi}$ we store the unique $\hat{\delta}_j(\hat{\Psi})$ and regress it on observable product characteristics to estimate the parameters of Equation (13).

4.1.1 Instruments

We assume ξ_{jt} is observed by market participants and thus is likely correlated with the premium and other characteristics, so identifying α_0 and β_x requires instruments. Some characteristics, such as basic drug coverage, are time-invariant and likely exogenous. The individuals in New York, we averaged the Plan C rates offered by all other insurers.

star rating is set at the insurer level with a two-year lag and is also plausibly exogenous.

Our cost function includes a geographic component so costs are correlated across plans in a market. For each plan, we average competitors' TM-covered service costs, weighted by conditional shares. Exclusion requires $E[c_{-jt}^{TM}\xi_{jt}] = 0$. Since these costs are private (released with a 5-year lag), firms are unlikely to choose ξ_{jt} based on particular competitors' costs.¹⁶

The panel nature of our data provides additional instruments. If characteristic-provision costs are correlated over time and benchmark updates are independent, lagged residuals from regressing plan characteristics on benchmarks and time-invariant state variables are valid instruments – they proxy for characteristic costs with the benchmark's common effect removed. ξ is orthogonal to these residuals under the additional assumption that shocks to ξ are uncorrelated with benefit-provision costs.¹⁷

To examine instrument validity, we first test 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 test residual persistence. We find that the correlation coefficient between the contemporaneous and lagged residuals ranges from 0.67 (enhanced drug coverage) to 0.88 (hospital copay). Taken together, the evidence suggests that the necessary conditions for our lagged residual instrumenting strategy hold. Instrument strength is reported below.

4.2 Demand estimates

We report maximum likelihood estimates of individual-specific parameters in Table 3. Higher income consumers are less price-sensitive. Switching costs are highest from TM to MA. Inter-firm switches are less costly and intra-firm cheaper still, suggesting that the primary component of switching costs is the disutility of changing providers. Younger beneficiaries, non-Whites, and those with less education have stronger MA preferences. The MA-sector random coefficient suggests idiosyncratic MA preferences matter.

¹⁶We weight by conditional shares as many markets feature plans with small enrollment; the cost information submitted by these plans is therefore noisier.

¹⁷More formally, write $x_{l,jt} = g_{f,l}(B_{mt}, a_j) + u_{l,jt}$ and $\xi_{jt} = g_{f,\xi}(B_{mt}, a_j) + u_{\xi,jt}$ for characteristic x_l and time-invariant state a_j . If benchmark updates are independent of u and $E[u_{l,jt-1}u_{\xi,jt}|B_{mt}, a_j] = 0$ — information relevant to the choice of x_l in $t-1$ is irrelevant to ξ in t after conditioning — then $E[u_{l,jt-1}\xi_{jt}] = 0$, and we instrument for $x_{l,jt}$ with $\widehat{u_{l,jt-1}}$.

Table 4 reports estimates of the mean utility parameters. Column (1) presents OLS estimates assuming prices and characteristics are exogenous. Column (2) reports IV estimates when the premium is instrumented with our cost instrument. Column (3) reports the results when prices and product characteristics are both treated as endogenous and instrumented with our full set of instruments. Consistent with attenuation bias on price, the IV premium coefficients are larger in magnitude than OLS. Characteristic coefficients are also generally larger when treated as endogenous and the estimates are somewhat noisier.

Joint underidentification and weak identification statistics are reported at the bottom of Table 4; individual partial F-statistics are reported in Appendix Table B.1. We reject underidentification, and each partial F-statistic exceeds the Stock and Yogo (2005) critical value for 10% maximal IV relative bias.

We focus on specification (3). For positive-premium plans, the average price elasticity is -6.29. The average semi-elasticity of increasing premiums by \$1 is -0.073, in line with the literature; Aizawa and Kim (2018), using an earlier sample, estimate -0.075. Combined with Table 3, the median consumer is willing to pay roughly \$425 annually for prescription drug coverage, \$300 for hearing coverage, and \$190 for a reduction of \$1,000 in the copay for a hospital stay. The cost incurred by a median-income individual switching from TM to MA is approximately \$650, similar to the mean annual premium in our data, while the same individual switching between plans within an MA insurer incurs a cost of only \$140.

4.3 Marginal costs estimates

The solution to Equation (6) is partially characterized by the first-order conditions

$$\{r_{jt}\} : \quad s_{jt} + \sum_{k \in \mathcal{J}_f} (r_{kt} - c_{kt}) \frac{\partial s_{kt}}{\partial p_{jt}} = 0 \quad (22)$$

where the share derivatives are now with respect to prices rather than per-enrollee revenues — we have $\frac{\partial s_{kt}}{\partial p_{jt}}$ from demand, and $\frac{\partial p_{jt}}{\partial r_{jt}} = 1$ leaves the solution unaffected. Following Berry et al. (1995), we define a $J \times J$ matrix Δ_t whose (j, k) elements are

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

The first-order conditions (22) can then be solved in vector form to recover plan costs:

$$\mathbf{c}_t = \mathbf{r}_t - \Delta_t^{-1} * \mathbf{s}_t. \quad (23)$$

Because revenues are observed for all plans—including zero-premium plans—this inversion applies uniformly.¹⁸

Characteristic-level cost parameters can be recovered from the first-order conditions for product characteristics. Because the subsidy depends on cost-sharing characteristics x_c but not on non-cost-sharing characteristics x_o , the conditions take different forms. For non-cost-sharing characteristics, $\frac{\partial \text{sub}_{jt}}{\partial x_{o,jt}} = 0$ and the first-order condition simplifies to

$$\gamma_{o,jt} = \sum_{k \in \mathcal{J}_f} (r_{kt} - c_{kt}) \frac{\partial s_{kt}}{\partial x_{o,jt}} (s_{jt} c_{jt})^{-1}. \quad (24)$$

For cost-sharing characteristics, the subsidy derivative is non-zero. From Equations (3) and (12), $\frac{\partial \text{sub}_{jt}}{\partial x_{c,jt}} = \frac{\lambda_{jt}-1}{\lambda_{jt}} \gamma_{c,jt} \tilde{c}_{jt}$, where $\tilde{c}_{jt} \equiv \exp(c_{0,jt} + \gamma_{c,jt} x_{c,jt} + \omega_{c,jt})$ is the component of marginal cost net of non-cost-sharing characteristics. Substituting and solving yields

$$\gamma_{c,jt} = \sum_{k \in \mathcal{J}_f} (r_{kt} - c_{kt}) \frac{\partial s_{kt}}{\partial x_{c,jt}} \left(s_{jt} \left[c_{jt} - \frac{\lambda_{jt}-1}{\lambda_{jt}} \tilde{c}_{jt} \right] \right)^{-1}. \quad (25)$$

In both cases r and s are data, c is recovered via Equation (23), and $\frac{\partial s_{kt}}{\partial x_{jt}}$ is calculated from the demand parameters. Equation (24) is applied first; $\tilde{c}_{jt}$ in (25) is then computed by stripping the recovered $\gamma_{o,jt}$ from c_{jt} . As we show in Appendix C, the observed rebate $\lambda_{jt}(B_{jt} - b_{jt})$ then identifies the baseline cost $c_{0,jt}$ and the unobservable cost-sharing component $\omega_{c,jt}$.

Table 5 reports the implied marginal-cost estimates and the contribution of plan characteristics to the logarithm of marginal cost, Equation (11). We report means and standard deviations of these characteristic-level costs and overall plan-level cost for the five largest firms grouping smaller firms into a sixth category.

Average MA costs for \$0-premium plans are \$874 per enrollee-month relative to \$906 for TM. Inflation-adjusted, Curto et al. (2021)’s comparable MA cost is \$830 per enrollee-

¹⁸In practice we implement the inversion using the Morrow and Skerlos (2011) form.

month. Our estimated economic profit margin (profit/revenue) is 2.1% in 2017; MedPAC’s 2017 accounting margin was 2.7% (MedPAC, 2019).

Our cost function estimates are broadly consistent with external measures of healthcare utilization. The estimates imply that MA enrollees have 4.4 annual physician visits, compared to an average of 4.98 visits in the Medicare population (Ashman et al., 2019). Similarly, the model implies 0.152 inpatient admissions per enrollee annually, relative to 0.252 admissions in Medicare overall (Kaiser Family Foundation, 2021). We estimate that providing prescription drug coverage costs \$442 for basic coverage and \$288 for enhanced coverage, while dental benefits add an additional \$202.¹⁹

4.4 Policy function estimates

Recall from Section 3.1, the policy function includes the benchmark and the benchmark interacted with insurer indicators. M_{jt} (also interacted with the benchmark) captures market-level cost and demographic shifts: the lagged number of MA contracts, log median income, the fraction of those 65-and-older who are White and Black, and the Medicaid-enrolled share. The regression also controls for the plan’s other characteristics, lagged market-level average characteristics across other firms’ plans, and the lagged number of plans offered by the plan’s parent and by competing contracts.

Benchmark updates are likely orthogonal to ξ but may correlate with insurers’ characteristic-level *costs*, making the benchmark potentially endogenous in Equation (7). We follow Duggan et al. (2016) and use the fact that payment floors differ by county-level CMS urban-rural status. Small population changes shift CMS’s urban-rural determination and thus the benchmark and so it is plausibly orthogonal to characteristic costs. We use the CMS urban-rural code interacted with year fixed effects as an instrument for B_{jt} .²⁰

Our estimates of $\theta_{1,fl}$ and $\theta_{2,l}$ from Equation (7) are reported in Table 6. The policy functions predict the Part B premium reduction, deductible, primary-care/specialist/hospital copays, enhanced drug coverage, and dental, vision, and hearing coverage. Basic drug coverage is held fixed as no plan in our data changed this indicator. We treat the unobserved

¹⁹We also test whether marginal costs vary with realized risk after risk adjustment. A 1% increase in risk-adjusted bids is associated with only a 0.0047% increase in realized risk (t-statistic = 0.77), suggesting that residual selection is quantitatively small.

²⁰We thank anonymous referees for suggesting this strategy.

quality ξ_{jt} as exogenous to the benchmark. We also report a policy function for per-enrollee revenue (column (10)) which we use in an alternative specification.²¹

The top panel reports the coefficients on the benchmark interacted with firm indicators and market-level variables. At the bottom of the table, we report the mean and standard deviation of the implied $\hat{\theta}_l$. In general, higher benchmarks increase plan characteristics. There are two exceptions. Both the deductible and vision characteristics have counter-intuitive signs. This result is likely driven by limited across-plan variation — by the end of our sample most plans had zero deductible and over 90% offered vision coverage. The implied magnitudes of both of these coefficient estimates are small.

For several characteristics, the standard deviations of the implied benchmark responses are large relative to the means (e.g. primary copay). This heterogeneity comes mainly from insurer responses, not market or demographic variation.²² This is consistent with firms maintaining distinct benefit-design strategies driven by insurer-specific cost or marketing differences (Starc, 2014).

5 Optimal Geographic Distribution of MA Subsidies

Figure 2 presents the distribution of the optimal benchmarks across counties for 2017. Panel (a) shows the level of the benchmarks while panel (b) shows the difference in the optimal benchmark from the 2017 benchmark level. Two patterns are of note. First, the optimal benchmark varies markedly across counties. The unweighted median is \$9,948 and standard deviation is \$927. Second, the differences are material but generally modest — the interquartile range of the difference is from -\$157 to \$607. The mean change is \$243 and the median is \$1. In percentage terms, roughly 87% of changes are less than 10% of the 2017 benchmarks; in 381 markets the absolute change is less than \$1,000.

While the benchmark changes are modest, the welfare gains from the optimal benchmark are large. Table 7 presents the counterfactual welfare impact. Consumer surplus more than doubles to \$10.2 billion and MA penetration rises from 29% to 47%. Below, we decompose

²¹See Appendix F.2.

²²A variance decomposition of the implied $\hat{\theta}_l$ into firm and demographic components attributes a majority of the variance to firm identity for seven of nine characteristics, with firm shares ranging from 66% (dental) to 95% (vision); demographic variation accounts for the larger share for the Part B premium reduction and hearing coverage.

the drivers of the welfare change. Aggregate firm variable profits nearly double from \$2.1 billion to \$4.2 billion.²³ That is, both Medicare enrollees and insurers are better off with the optimal benchmarks.

5.1 Prices and plan characteristics

Figure 3 shows box-and-whisker plots of prices and characteristics under the 2017 and optimal policies — for each characteristic, the first box illustrates the 2017 data and the second shows the optimal policy.²⁴

Two notable patterns emerge. First, the optimal policy generally moves product characteristics in a more generous direction – premiums and copays decrease while supplemental benefits improve. Higher benchmarks raise per-enrollee revenues which firms pass through as more generous benefits and lower premiums. Second, the pass-through is uneven across characteristics. Hospital copays respond more in percentage terms than specialist copays likely reflecting differences in the marginal cost of benefit provision. The share of zero-premium plans changes modestly (38.0% vs. 39.1% under the 2017 policy) reflecting the asymmetric effect of the zero-premium constraint. Benchmark decreases quickly push plans to set positive premiums, while benchmark increases induce plans to zero premiums gradually. One concern is that the equilibrium using our policy function approach deviates from the Nash equilibrium. We verify that our policy function approximation is close to the Nash equilibrium by computing profit deviations from small perturbations in product characteristics. The results indicate that our implied equilibrium is very close to a Nash equilibrium.²⁵

5.2 The directional selection of benchmark changes

Our optimal subsidy problem implies markets are selected for benchmark increases based on the marginal effect of the benchmark on consumer welfare and government expenditures. Figure 6 illustrates the distribution of the derivatives of CS , $\overline{CS}^{cond}$, $GovExp$, and ‘total surplus’ (defined as $CS - GovExp$) with respect to the benchmark rate at the 2017 policy.

²³Detailed firm-level results are reported in Appendix Table F.1.

²⁴The ‘enhanced prescription drug’, ‘dental’, ‘vision’, and ‘hearing’ characteristics are indicators. We interpret values between 0 and 1 as less generous than the 2017 average for that coverage type, and values above 1 as more generous. We model them as continuous for consistency and computational simplicity. We have explored alternative specifications with similar results.

²⁵Details are in Appendix Table B.2.

We separate markets by the direction of the benchmark change.

The optimizer’s goal of maximizing the “bang for the buck” can be seen in the split in the derivatives of total surplus. In every market where the optimal benchmark decreases, the derivative is negative at the 2017 policy, and 199 of the markets with benchmark increases have a positive derivative.

Among the components, the CS and $\overline{CS}^{cond}$ derivative distributions overlap substantially across markets with increases and decreases, so compensating variation does not solely determine direction. In contrast, the $GovExp$ derivative shows greater variance. Fiscal-cost variation matters more than welfare variation in selecting markets for changes. In many markets, the derivative of $GovExp$ is negative at the 2017 benchmark. This occurs when benchmarks are below TM spending, and the cost-savings from beneficiaries that move from TM to MA outweigh the cost increases for inframarginal MA enrollees. We explore this in Appendix F.1.

Taken together, these results imply that the current benchmark methodology fails to fully account for the fiscal cost channel. It sets benchmarks too high in markets where marginal subsidy increases are especially costly to the government, and too low in markets where increasing the benchmark imposes relatively little fiscal cost.

5.3 Winners and losers

Figure 4 compares annual compensating variation (CV) across groups of beneficiaries under the 2017 and optimal policies. The first set of bars illustrates the mean CV for all Medicare beneficiaries and conditional on MA enrollment. The second set of bars splits markets by the direction of the benchmark change. The optimal policy increases the benchmark in 262 of 439 markets, where mean MA-enrollee CV rises from \$234 to \$528; in the 177 markets with lower benchmarks, mean CV falls from \$250 to \$245.

Previous work has identified racial and income inequalities in Medicare and MA (Li et al., 2017). The third and fourth sets of bars examine the welfare changes by race and income. While all groups gain on average, CV rises more in percentage terms for White and Hispanic enrollees than for Black enrollees. High-income beneficiaries benefit more in percentage terms than low- or medium-income beneficiaries.

Table 7 decomposes aggregate outcomes by the direction of the benchmark change. The increase in aggregate CV is driven both by moving individuals from TM to MA and by making MA more valuable to those already signed up for an MA plan: though the increase in share in percentage terms is smaller than the percentage increase in the mean compensating variation per MA enrollee. The optimal policy transfers \$89 billion from TM to MA. We find no evidence of economically meaningful TM cost externalities from increased MA enrollment (Lakdawalla and Yin, 2015): regressing log risk-adjusted TM cost on log MA share with county and year fixed effects yields a coefficient of -0.007 ($t = 2.19$).

The reallocation of subsidies is primarily driven by the gap between TM and MA costs. As reported in Table 8, markets in the top tercile of risk-adjusted TM spending receive benchmark increases in 98% of cases, with a mean increase of \$803 per enrollee. In contrast, 87% of markets in the bottom tercile receive benchmark decreases, with a mean decrease of \$318. This pattern reflects the fiscal logic identified in Section 5.2: where TM spending is high relative to MA costs, shifting enrollees from TM to MA generates savings that fund more generous subsidies. Geographically, the optimal policy shifts subsidies toward the Northeast and urban markets. These patterns largely follow from the distribution of TM costs across regions. Market concentration plays a secondary role. Markets with more competing contracts are more likely to receive increases, but the relationship is modest.

5.4 The role of changes in product characteristics

Next, we decompose the contribution to consumer welfare attributable to changes in each product characteristic. We apply the gradient theorem to approximate the change in surplus as the sum of characteristic-specific contributions across products and consumers.²⁶

Figure 5 illustrates this decomposition with box-and-whisker plots of the beneficiary-level percentage contributions of each product characteristic to the change in consumer welfare. In subfigure (a), we compare the impact of changes in all non-premium characteristics to changes in the premiums. These results highlight the importance of modeling product characteristics.

²⁶Let $X(0)$ and $X(1)$ denote the vectors of premiums and product characteristics under the 2017 and optimal policies. By the gradient theorem, $CS(X(1)) - CS(X(0)) = \nabla_X CS(X(0)) \cdot dX(1) + o(\|dX(1)\|)$. Expanding over consumers i , products j , and characteristics k : $\sum_i \sum_j \sum_k \nabla CS_{ijk}(X(0)) \cdot dX(1)_{jk} + o(\|dX(1)\|)$. We approximate each term as $\frac{1}{\alpha_i}(s_{ij}(\delta(0), \mu(0)) + s_{ij}(\delta(1), \mu(1)))\beta_{ik}\Delta X_{jk}$. Summing over products and consumers for each characteristic k yields its total contribution.

Across all markets, non-premium product characteristics contribute an average of roughly 20% of the change in compensating variation.

In subfigure (b), we detail the contribution of each characteristic. The largest mean contributions come from enhanced drug benefits, hearing coverage, and hospital copays. This result is driven by both demand- and supply-side factors. On the demand side, the elasticity with respect to hospital stay copays is larger than for other copays, and the preferences for enhanced prescription drug and hearing are higher than other coverage types (except basic prescription drug, which does not change at the plan level). The policy function estimates on the supply side modulate these effects somewhat. The hospital copay changes (on average) less than primary care or specialist copays, but that difference is not large enough to overcome demand-side preferences. In contrast, the average change in enhanced drug and hearing benefits is larger than for other coverage types.

6 Conclusion

Policy makers have implemented public-private partnerships in which they subsidize private firms to support consumers' choices. We provide a framework for calculating the optimal market-level subsidy schedule that combines policy function estimation with equilibrium pricing to endogenize prices and product characteristics.

We apply our framework to Medicare Advantage using micro-level panel data. We find that the optimal budget-neutral subsidy schedule differs substantially from the current policy: the 2017 policy generates \$76 in consumer welfare per beneficiary per year, while the optimal policy generates \$242. These gains come from both moving beneficiaries from TM to MA and improving benefits for existing MA enrollees. Changes in non-price product characteristics account for roughly 20% of the welfare gains, underscoring the importance of modeling the product characteristic channel.

Our framework applies to differentiated products settings where consistent estimates of firm policy functions are available and the counterfactual policy space is well-represented in the data. Caution is warranted regarding equilibrium multiplicity and the inherent tradeoff between tractability and accuracy.

References

- Aizawa, N. and Kim, Y. S. (2018). Advertising and risk selection in health insurance markets. *American Economic Review*, 108(3):828–67.
- Ashman, J. J., Rui, P., and Okeyode, T. (2019). Characteristics of office-based physician visits, 2016. *NCHS Data Brief*, no 331.
- Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370.
- Barahona, N., Otero, C., and Otero, S. (2023). Equilibrium effects of food labeling policies. *Econometrica*, 91(3):839–868.
- 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.
- 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.
- Brown, J., Duggan, M., Kuziemko, I., and Woolston, W. (2014). How does risk selection respond to risk adjustment? new evidence from the medicare advantage program. *American Economic Review*, 104(10):3335–64.
- Bundorf, M. K., Levin, J., and Mahoney, N. (2012). Pricing and welfare in health plan choice. *American Economic Review*, 102(7):3214–48.
- Cabral, M., Geruso, M., and Mahoney, N. (2018). Does privatized medicare benefit patients or producers? evidence from the benefits improvement and protection act. *American Economic Review*.
- Curto, V., Einav, L., Levin, J., and Bhattacharya, J. (2021). Can health insurance competition work? evidence from medicare advantage. *Journal of Political Economy*, 129(2):570–606.
- 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.
- 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. (2026). Market design in regulated health insurance markets: Risk adjustment vs. subsidies. *Journal of Political Economy Microeconomics*.
- 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. *Econometrica*, 72(2):351–381.
- Gruber, J. (2017). Delivering public health insurance through private plan choice in the united states. *Journal of Economic Perspectives*, 31(4):3–22.
- Hoxby, C. M. (2000). Does competition among public schools benefit students and taxpayers? *American Economic Review*, 90(5):1209–1238.
- Jaffe, S. and Shepard, M. (2020). Price-linked subsidies and imperfect competition in health insurance. *American Economic Journal: Economic Policy*, 12(3):279–311.
- Kaiser Family Foundation (2021). State health facts: Medicare service use: Hospital inpatient services. Technical report, Kaiser Family Foundation.
- Kim, K. i., Petrin, A., and Song, S. (2016). Estimating production functions with control functions when capital is measured with error. *Journal of Econometrics*, 190(2):267–279. Endogeneity Problems in Econometrics.
- 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.
- 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.
- MedPAC (2017). Status report on the medicare advantage program. In *Report to the Congress: Medicare Payment Policy*, chapter 13, pages 345–379. MedPAC.
- MedPAC (2019). *Report to the Congress, Medicare Payment Policy*. MedPAC.
- Morrow, W. R. and Skerlos, S. J. (2011). Fixed-point approaches to computing bertrand-nash equilibrium prices under mixed-logit demand. *Operations Research*, 59(2):328–345.
- Nevo, A. (2000). Mergers with differentiated products: The case of the ready-to-eat cereal industry. *The RAND Journal of Economics*, pages 395–421.
- 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.
- Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, 110:705–729.
- Polyakova, M. and Ryan, S. P. (2023). Market power and redistribution: Evidence from the affordable care act. Technical report, Tech. rep., Stanford and Washington University in St Louis.

- 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.
- Spence, A. M. (1975). Monopoly, quality, and regulation. *The Bell Journal of Economics*, pages 417–429.
- Spence, M. (1976). Product differentiation and welfare. *The American Economic Review*, 66(2):407–414.
- Starc, A. (2014). Insurer pricing and consumer welfare: Evidence from medigap. *The RAND Journal of Economics*, 45(1):198–220.
- Stock, J. and Yogo, M. (2005). Asymptotic distributions of instrumental variables statistics with many instruments. *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg*, 6:109–120.
- Sweeting, A. (2013). Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry. *Econometrica*, 81(5):1763–1803.
- Tebaldi, P. (2025). Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca. *Review of Economic Studies*, 92(1):586–620.
- Town, R. and Liu, S. (2003). The welfare impact of Medicare HMOs. *Rand J Econ*, 34(4):719–736.
- Vatter, B. (2025). Quality disclosure and regulation: Scoring design in medicare advantage. *Econometrica*, 93(3):959–1001.
- Wollmann, T. G. (2018). Trucks without bailouts: Equilibrium product characteristics for commercial vehicles. *American Economic Review*, 108(6):1364–1406.

7 Tables

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

Variable	Benchmark quartile			
	1st	2nd	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
<i>Copays</i>				
Primary care (\$)	12.90	12.60	14.49	13.99
Specialist visit (\$)	37.36	35.36	32.32	28.57
Hospital stay (\$)	1,351	1,288	1,118	1,003
<i>Supplemental coverage indicators</i>				
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: This table summarizes our MA plan data across quartiles of the base benchmark defined at the market (county-year) level. An observation is a plan-county-year. Reported figures are unweighted means unless noted. All dollars are 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.

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
<i>Demographic indicators</i>						
Female	.536	.499	.548	.531	.526	.536
Black	.081	.273	.095	.075	.094	.069
Hispanic	.010	.101	.017	.008	.018	.006
<i>Education indicators</i>						
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
<i>Health status indicators</i>						
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
<i>Administrative indicators</i>						
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. 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 across all observations. The remaining columns report means for the relevant group. Income and Medigap price are in 2017 dollars. All statistics use MCBS sampling weights.

Table 3: Maximum likelihood estimates of individual-specific preferences

Variable	Coeff.	Std. Err.	Variable	Coeff.	Std. Err.
Price $\times$ Income	0.1766	0.0321	MA $\times$ Demographics		
Price $\times$ Income ²	-0.0028	0.0014	Excellent health	0.5421	0.2050
			Very good health	0.5849	0.1923
TM-to-MA switch $\times$			Good health	0.4503	0.1908
Constant	-4.1212	0.2132	Fair health	0.4346	0.2027
Excellent health	-0.0797	0.2184	Age	2.0606	0.1810
Very good health	-0.2134	0.2053	Age ²	-0.1442	0.0125
Good health	-0.2131	0.2052	Female indicator	-0.1123	0.0348
Fair health	-0.2577	0.2191	Black indicator	0.4814	0.0690
			Hispanic indicator	0.4215	0.1595
Inter-Insurer switch $\times$			Graduated high school	-0.1661	0.0562
Constant	-1.9003	0.1061	Attended college	-0.3464	0.0575
Excellent health	-0.0861	0.1215	College degree or higher	-0.8418	0.0669
Very good health	-0.0637	0.1148			
Good health	-0.1138	0.1143	Administrative indicators		
Fair health	-0.1353	0.1234	Employer-provided insurance	-5.8128	0.1792
			Non-aged eligibility	0.4633	0.0663
Intra-Insurer switch $\times$			ESRD diagnosis	-1.7561	0.2158
Constant	-0.8171	0.1334	Full year enrollment	2.8909	0.1090
Excellent health	-0.1770	0.1530			
Very good health	-0.1426	0.1449	Outside good (Medigap) price $\times$		
Good health	-0.1224	0.1446	Linear	0.3834	0.0581
Fair health	-0.1348	0.1527	Income	-0.1986	0.0135
			Income ²	0.0040	0.0005
Random Coefficient	1.6243	0.1266			
Weighted Log Likelihood			-73,967		
Observations			78,812		

Notes: The coefficients reported here are maximum-likelihood estimates of the individual-specific components of utility, i.e. the parameters of Equation (14). 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. Age is measured in decades. Educational indicators are mutually exclusive. The Medigap price is the price of Plan C, see text for details. Standard errors are robust to heteroskedasticity.

Table 4: Estimates of mean preferences for plan characteristics

Variable	(1)	(2)	(3)
Annual Premium (per \$1000)	-0.4005 (0.0417)	-6.9249 (1.8581)	-6.7368 (2.7489)
Part B reduction (per \$1000)	0.1368 (0.0910)	0.6617 (0.1987)	1.8192 (0.7343)
Deductible (per \$1000)	-0.3062 (0.0577)	-1.3622 (0.3593)	-1.8048 (0.7113)
<i>Copays</i>			
Primary care visit	-0.0270 (0.0034)	-0.0074 (0.0083)	-0.0499 (0.0195)
Specialist visit	0.0140 (0.0037)	-0.0404 (0.0195)	-0.0438 (0.0378)
Hospital stay copay (per \$1000)	0.0068 (0.0459)	-0.8844 (0.2727)	-1.2624 (0.6264)
<i>Supplemental coverage indicators</i>			
Prescription drug	0.5928 (0.0621)	3.6512 (0.9012)	2.8598 (1.0738)
Enhanced prescription drug	0.3558 (0.0509)	0.1912 (0.1451)	1.8271 (0.5163)
Dental	0.0684 (0.0419)	0.9969 (0.2750)	1.3355 (0.6329)
Vision	-0.0041 (0.0597)	0.2470 (0.1314)	0.6098 (0.4780)
Hearing	-0.0290 (0.0463)	1.1064 (0.3681)	2.0865 (0.9388)
CMS-assigned star rating	0.1498 (0.0217)	0.2535 (0.0545)	0.3355 (0.1318)
Firm fixed effects?	Yes	Yes	Yes
Endogenous annual premium?	No	Yes	Yes
Endogenous product characteristics?	No	No	Yes
Observations	50,439	50,439	36,077
Mean implied elasticity (if < 0)	-0.3077 (0.2100)	-6.4684 (4.4004)	-6.2908 (4.2793)
Mean ds_j/dp_j	-0.0035 (0.0061)	-0.0745 (0.1257)	-0.0725 (0.1222)
<i>IV diagnostics</i>			
Kleibergen-Paap LM statistic		18.27 [0.0000]	7.84 [0.0051]
Cragg-Donald Wald F statistic		156.28	6.04
Kleibergen-Paap Wald F statistic		21.92	0.86

Notes: The coefficients reported here are estimates of the components of the utility function which are common across individuals i.e. the parameters of Equation (13). An observation is a plan-county-year. Estimates in Column (1) are formed via OLS. 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 consider all product characteristics as endogenous and 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 standard errors clustered at the county level in the coefficient panel and standard deviations in the elasticity panel. P-values for the LM statistic are in brackets. Individual partial F-statistics are reported in Appendix Table B.1.

Table 5: Marginal cost parameter estimates

Variable	γ_{xj}					
	Aetna	BCBS	Humana	Kaiser	UHG	Other
<i>Cost-sharing characteristics</i>						
Part B reduction (per \$1000)	0.0229 (0.0026)	0.0236 (0.0027)	0.0237 (0.0025)	0.0243 (0.0024)	0.0243 (0.0024)	0.0240 (0.0027)
Deductible (per \$1000)	-0.0185 (0.0021)	-0.0191 (0.0022)	-0.0191 (0.0020)	-0.0196 (0.0020)	-0.0196 (0.0019)	-0.0194 (0.0022)
Primary care copay	-0.0004 (0.0000)	-0.0004 (0.0000)	-0.0004 (0.0000)	-0.0004 (0.0000)	-0.0004 (0.0000)	-0.0004 (0.0000)
Specialist copay	-0.0008 (0.0001)	-0.0009 (0.0001)	-0.0009 (0.0001)	-0.0009 (0.0001)	-0.0009 (0.0001)	-0.0009 (0.0001)
Hospital copay (per \$1000)	-0.0135 (0.0015)	-0.0140 (0.0016)	-0.0140 (0.0015)	-0.0144 (0.0014)	-0.0144 (0.0014)	-0.0142 (0.0016)
<i>Supplemental coverage characteristics</i>						
Prescription drug	0.0383 (0.0050)	0.0400 (0.0046)	0.0406 (0.0042)	0.0393 (0.0042)	0.0418 (0.0044)	0.0401 (0.0048)
Enhanced prescription drug	0.0251 (0.0033)	0.0262 (0.0030)	0.0266 (0.0028)	0.0258 (0.0028)	0.0274 (0.0029)	0.0263 (0.0031)
Dental	0.0201 (0.0026)	0.0210 (0.0024)	0.0213 (0.0022)	0.0206 (0.0022)	0.0219 (0.0023)	0.0210 (0.0025)
Vision	0.0006 (0.0001)	0.0006 (0.0001)	0.0006 (0.0001)	0.0006 (0.0001)	0.0007 (0.0001)	0.0006 (0.0001)
Hearing	0.0264 (0.0034)	0.0276 (0.0031)	0.0280 (0.0029)	0.0271 (0.0029)	0.0288 (0.0030)	0.0277 (0.0033)
Mean marginal cost (\$)	11,524	10,966	10,805	11,184	10,498	10,957
Std. dev. marginal cost (\$)	1,520	1,282	1,191	1,302	1,170	1,353
Observations	3,857	9,102	12,800	1,719	4,852	32,506

Notes: We estimate these parameters by inverting the first-order conditions for profit maximization. Observations are county-year-plans. Reported values are means. Standard deviations are in parentheses.

Table 6: Policy function estimation results

	(1) Part B reduction	(2) Deduct- ible	(3) Prim. copay	(4) Spec. copay	(5) Hospital copay	(6) Enhanced drug	(7) Dental	(8) Vision	(9) Hearing	(10) Revenue
<i>Benchmark (\$000s) ×</i>										
Aetna	0.0419 (0.0705)	-0.1239 (0.1813)	1.6338 (4.8952)	-9.5337 (6.3543)	0.3520 (0.2969)	0.5490 (0.2142)	0.6476 (0.3000)	-0.2121 (0.1365)	1.4424 (0.3395)	-3.8310 (1.3289)
BCBS	0.0410 (0.0698)	-0.0881 (0.1814)	1.3330 (4.9320)	-9.9032 (6.3236)	0.3941 (0.2962)	0.6054 (0.2164)	0.6191 (0.2998)	-0.2381 (0.1386)	1.4687 (0.3388)	-3.7671 (1.3303)
Humana	0.0351 (0.0695)	-0.0920 (0.1793)	0.4466 (4.9087)	-8.7643 (6.3233)	0.3345 (0.2951)	0.6561 (0.2186)	0.5735 (0.2983)	-0.3578 (0.1318)	1.4923 (0.3377)	-3.7487 (1.3277)
Kaiser	0.0404 (0.0694)	-0.1831 (0.1799)	-1.7580 (4.9963)	-7.0697 (6.2721)	0.3159 (0.3071)	0.5455 (0.2205)	0.6269 (0.2992)	-0.2322 (0.1350)	1.4885 (0.3400)	-4.0122 (1.3259)
UHG	0.0365 (0.0690)	-0.1373 (0.1805)	1.8214 (4.9313)	-10.6591 (6.2879)	0.4234 (0.2945)	0.7021 (0.2240)	0.7028 (0.3001)	-0.2437 (0.1358)	1.4259 (0.3389)	-3.7616 (1.3297)
Other	0.0493 (0.0706)	-0.1652 (0.1807)	1.6649 (4.9189)	-10.5762 (6.3069)	0.3796 (0.2967)	0.5764 (0.2198)	0.6293 (0.2989)	-0.2209 (0.1362)	1.4513 (0.3386)	-3.7511 (1.3309)
Lagged # contracts	0.0008 (0.0002)	-0.0001 (0.0006)	0.0393 (0.0141)	-0.0184 (0.0216)	-0.0000 (0.0009)	-0.0016 (0.0006)	0.0013 (0.0009)	0.0003 (0.0003)	-0.0026 (0.0008)	0.0233 (0.0065)
Log median income	-0.0009 (0.0058)	0.0198 (0.0147)	-0.1325 (0.4020)	0.9900 (0.5241)	-0.0367 (0.0259)	-0.0259 (0.0188)	-0.0500 (0.0243)	0.0230 (0.0111)	-0.1163 (0.0284)	0.3530 (0.1079)
% 65+ White	-0.0191 (0.0127)	-0.0933 (0.0358)	-0.3214 (0.9527)	-1.0093 (1.2686)	-0.0658 (0.0506)	-0.0890 (0.0336)	-0.0707 (0.0575)	-0.0109 (0.0225)	-0.0697 (0.0499)	-0.1095 (0.2698)
% 65+ Black	-0.0404 (0.0178)	-0.0817 (0.0413)	-4.7558 (1.2378)	2.4925 (1.5836)	-0.1116 (0.0634)	-0.0249 (0.0526)	-0.0080 (0.0679)	-0.0686 (0.0276)	0.0528 (0.0646)	0.5240 (0.3953)
% Medicaid	0.0001 (0.0001)	-0.0005 (0.0004)	0.0046 (0.0114)	-0.0071 (0.0151)	-0.0005 (0.0006)	-0.0013 (0.0004)	0.0000 (0.0007)	0.0001 (0.0003)	-0.0025 (0.0006)	0.0181 (0.0041)
Obs.	37,998	37,998	37,998	37,998	37,998	37,998	37,998	37,998	37,998	37,998
R^2	0.1931	0.2051	0.5442	0.7003	0.5796	0.2779	0.3412	0.2242	0.5015	0.7790
$\hat{\theta}_t$ statistics										
Mean	0.0262	-0.0150	-0.2343	-0.1454	-0.1042	0.1560	0.0324	-0.0099	0.0512	0.7036
Std. dev.	0.0093	0.0353	0.8432	0.9665	0.0273	0.0915	0.0388	0.0548	0.406	0.2689

Notes: This table reports estimates of the effect of changes in the benchmark on product characteristics – i.e. estimates of θ from Equation (7). Each column reports the coefficients of a regression with the specified product characteristic as the dependent variable. The independent variables for each regression include the benchmark interacted with firm indicators and with the market-level moderators shown (lagged number of MA contracts, log median income, and demographic shares). The regression also controls for the plan's other characteristics, lagged market-level average characteristics across other firms' plans, and the lagged number of plans offered by the plan's parent and competing contracts (omitted 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. Observations are county-year-plans. Standard errors clustered at the county level are in parentheses.

Table 7: Aggregate outcomes under 2017 policy and optimal policy

<i>Summary across all markets</i>			
	2017 Policy	Optimal Policy	% Change
Total MA share (%)	29.8	47.0	57.7
Aggregate consumer welfare (\$ billion)	3.18	10.16	219.7
Government spending on MA (\$ billion)	126.3	214.9	70.1
Government spending on TM (\$ billion)	314.2	225.7	-28.2
Total government expenditures (\$ billion)	440.5	440.5	0.0
<i>Markets in which benchmark increases (262 markets, 63.2% population share)</i>			
	2017 Policy	Optimal Policy	% Change
Total MA share (%)	28.5	65.3	129.2
Aggregate consumer welfare (\$ billion)	1.88	9.54	408.8
Government spending on MA (\$ billion)	76.7	191.4	149.4
Government spending on TM (\$ billion)	213.2	100.5	-52.8
Total government expenditures (\$ billion)	289.9	291.9	0.7
<i>Markets in which benchmark decreases (177 markets, 36.8% population share)</i>			
	2017 Policy	Optimal Policy	% Change
Total MA share (%)	32.1	15.6	-51.4
Aggregate consumer welfare (\$ billion)	1.3	.617	-52.6
Government spending on MA (\$ billion)	49.6	23.5	-52.6
Government spending on TM (\$ billion)	101.0	125.1	23.9
Total government expenditures (\$ billion)	150.6	148.6	-1.3

Notes: We calculate all statistics at the individual level and aggregate using the MCBS sample weights. All dollars are 2017 dollars. Totals differ slightly due to rounding.


Table 8: Benchmark changes under the optimal policy by market characteristics

	N markets	% with increase	Mean ΔB_m (\$)	Pop.-weighted mean ΔB_m (\$)
<i>By risk-adjusted TM cost tercile</i>				
Low	147	12.9	-318	-331
Medium	146	68.5	284	145
High	146	97.9	803	719
<i>By Census region</i>				
Northeast	81	74.1	861	1055
Midwest	124	50.0	93	153
South	185	59.5	103	170
West	49	61.2	239	323
<i>By urban/rural status</i>				
Urban (RUC 1–3)	321	63.9	307	426
Rural (RUC 4–9)	118	48.3	115	64
All markets	439	59.7	255	413

Notes: This table reports the distribution of benchmark changes under the optimal policy across market characteristics. “Low,” “Medium,” and “High” TM cost refer to terciles of risk-adjusted per-capita TM spending. Urban markets are those with a Rural-Urban Continuum code of 1–3. ΔB_m is the change from the 2017 benchmark to the optimal benchmark. All dollars are 2017 dollars.


8 Figures

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




Notes: Map (a) illustrates the ratio of the 2017 benchmark rate to the 2017 risk-adjusted TM (FFS) spending in each county. To show detail, the data are windsorized at the 5th and 95th percentiles. Map (b) illustrates the county-level MA penetration rate in 2017, defined as the total number of people enrolled in any MA plan divided by the number of Medicare beneficiaries in the county. All data from CMS.

Figure 2: The distribution of the optimal benchmarks across markets




Notes: Graph (a) illustrates the distribution of benchmarks under the policy that solves Equation (4). Graph (b) illustrates the distribution of differences between the optimal policy and the 2017 policy.

Figure 3: The distribution of product characteristics under the 2017 policy and the optimal policy




Notes: These box-and-whiskers plots illustrate the unweighted plan-level distribution of product characteristics under the 2017 policy (dark) and the optimal policy (light) across all markets. Outliers excluded for clarity.

Figure 4: Compensating variation under the 2017 policy and the optimal policy




Notes: This chart illustrates the mean annual compensating variation for individuals in various groups under the 2017 policy and under the policy that solves Equation (4). We calculate compensating variation for all Medicare beneficiaries via Equation (17) and compensating variation conditional on MA enrollment via Equation (18). The optimal policy increases benchmarks in 262 markets and decreases benchmarks in 177 markets. All dollars are 2017 dollars. White, Black, and Hispanic groups are defined by CMS. We weight all calculations by the MCBS sample weights.

Figure 5: The contribution of product characteristics to changes in compensating variation



Notes: To calculate these percentages, we decompose the total change in compensating variation into the changes stemming from changes in the premium and each of the product characteristics when moving from the 2017 to the optimal policy using the gradient theorem at the individual level.

Figure 6: Derivatives of market-level welfare and spending functions with respect to the benchmark



Notes: These box-and-whiskers plots illustrate the market-level distribution of the derivatives of welfare and spending functions with respect to the benchmark. We calculate “CV per Medicare beneficiary” via Equation (17), “CV per MA enrollee” via Equation (18), and “Government expenditures” via Equation (19). Markets are categorized according to the direction of the benchmark change when moving from the 2017 policy to the policy that solves Equation (4).

Appendices

A The role of non-price characteristics

The analysis in this paper is predicated on the idea that firms may change price and non-price characteristics of their products in response to per-capita subsidies offered by the government. In this Appendix we explore this phenomenon in the form of a simplified model that illustrates the importance of incorporating changes in non-price characteristics into the analysis.

Consider a monopolist offering a single product at price p to a measure of consumers. The product has a single non-price ‘quality’ characteristic x , also chosen by the firm. Demand for the product follows a discrete choice framework. Consumer i has price sensitivity α_i and preference for quality β_i . The utility of purchasing the good is given by $u_i = \alpha_i p + \beta_i x + \epsilon_i$, where ϵ_i is taken to be Type-I Extreme Value. Given a distribution of consumer preferences, demand is given by $s(p, x) = \int \frac{\exp(\alpha_i p + \beta_i x)}{1 + \exp(\alpha_i p + \beta_i x)} di$.

The marginal cost of providing the product is a function of quality: $c(x) = \exp(c_b + c_x x)$ where c_b and c_x are parameters. The firm earns a per-capita subsidy from the government B . The firm’s problem is therefore

$$\max_{p, x} (p - c(x) + B)s(p, x).$$

Suppose that the market consists of two types of consumers with equal frequency. Type 1 consumers have $\alpha_1 = -1.5, \beta_1 = 0.25$. Relative to type 1 consumers, type 2 consumers are more price sensitive and have a stronger taste for quality: $\alpha_2 = -4.0, \beta_2 = 2.0$. Let $c_b = -0.5$ and $c_x = 1.0$. Figure A.1 illustrates outcomes under these assumptions. In each subfigure, the solid lines indicate equilibrium outcomes as the subsidy varies when the firm can freely choose both the price and the quality of the good. The dashed lines indicate outcomes when quality is held fixed at the level the firm chooses when it is unconstrained and it receives zero subsidy.

Subfigure (a) illustrates the firm’s optimal choices. As the subsidy increases, the price decreases and the quality increases. We note that this need not always hold. For some

combinations of consumer preferences and firm costs the firm's quality policy function is decreasing in the subsidy. The price decreases more under the constrained scenario than the unconstrained scenario, though we note the difference in prices remains small relative to the price level even as the subsidy increases.


Subfigure (b) illustrates the choices of the two types of consumers. Type 1 consumers are more likely to purchase the good across the range of subsidies, and the difference between their choices in the constrained and unconstrained scenarios remains small relative to the likelihood of purchase. A significantly larger difference is seen in the behavior of Type 2 consumers. As the subsidy increases, the increase in the probability of purchase under the unconstrained scenario is more than double the increase in the probability of purchase under the constrained scenario.

Finally, subfigure (c) illustrates government spending and firm profits. As the subsidy increases, the firm's profits increase. This increase is larger in the unconstrained scenario. At a subsidy of $B = 1.0$, the firm earns 12% higher profits when it is able to change both the price and quality of the product. The gap in government spending is even larger. At $B = 1.0$, the government spends 32% more when the firm can change its quality.

To summarize, in this model, taking non-price product characteristics as given when they are in fact endogenous would significantly underestimate the impact of subsidies on consumer welfare and government spending. As these are key inputs into the optimal subsidy schedule problem, it is likely that such an analysis would generate an optimal subsidy schedule that is considerably different than the schedule generated by an analysis that incorporated endogenous non-price characteristics.

These differences are apparent even in the absence of the zero-lower-bound on prices or any heterogeneity in costs across consumers. In other words, risk selection incentives are not necessary to generate responses in non-price characteristics. Here, the response is driven solely by the heterogeneity in demand. As the subsidy increases, the firm chooses to attract more type 2 consumers by increasing the quality of their product.

Figure A.1: Choices and outcomes in a simplified model



Notes: This figure plots choices and outcomes for the model described in Appendix A. In the model, a monopolist offers a single product with a single non-price ‘quality’ characteristic to a measure of consumers consisting of two types. See text for details and parameters. In each graph we plot outcomes in two scenarios. In the first scenario (solid lines), we allow the firm to freely choose both price and quality. In the second scenario (dashed lines), we hold the quality fixed at the same level as the optimal choice when the firm is not subsidized.

B Additional diagnostics

Table B.1: Diagnostics for mean preference regressions

	(2)	(3)
<i>Joint underidentification test</i>		
Kleibergen-Paap LM statistic	18.27	7.84
	(0.0000)	(0.0051)
<i>Joint weak identification test</i>		
Cragg-Donald Wald F statistic	156.28	6.04
Kleibergen-Paap Wald F statistic	21.92	0.86
<i>Sanderson-Windmeijer “Partial F” weak identification test</i>		
Annual Premium	21.92	17.67
Part B deduction		62.29
Deductible		28.13
<i>Copays</i>		
Primary care		940.65
Specialist visit		22.14
Hospital visit		22.39
<i>Supplemental coverage</i>		
Enhanced drug		55.54
Dental		29.46
Vision		357.39
Hearing		18.70

Notes: This table reports test statistics for the instrumental variables specifications reported in Table 4. All statistics calculated using the `ivreg2` package in Stata. Critical values for the weak identification tests are available in Stock and Yogo (2005). P-values are reported in parentheses where available.

Table B.2: The profitability of a 1% deviation from predicted product characteristics at the optimal policy

<i>A increase in the characteristic at the optimal policy by 1% of the mean...</i>										
	Enhcd. drug	Dental	Vision	Hearing	Part B reduc.	Deduct.	Prim. copay	Spec. copay	Hosp. copay	
<i>All markets</i>										
25th Percentile	-0.0114	-0.0069	-0.0000	-0.0147	0.0011	-0.0517	-0.1708	-0.8905	-0.4996	
Median	-0.0033	-0.0018	-0.0000	-0.0043	0.0069	-0.0167	-0.0550	-0.2864	-0.1607	
75th Percentile	-0.0005	-0.0002	-0.0000	-0.0007	0.0214	-0.0025	-0.0084	-0.0435	-0.0245	
<i>262 markets with increases</i>										
25th Percentile	-0.0155	-0.0097	-0.0001	-0.0195	0.0004	-0.0536	-0.1770	-0.9211	-0.5170	
Median	-0.0041	-0.0025	-0.0000	-0.0051	0.0059	-0.0143	-0.0471	-0.2454	-0.1375	
75th Percentile	-0.0003	-0.0002	-0.0000	-0.0004	0.0222	-0.0009	-0.0031	-0.0160	-0.0090	
<i>177 markets with decreases</i>										
25th Percentile	-0.0069	-0.0036	-0.0000	-0.0094	0.0028	-0.0495	-0.1633	-0.8542	-0.4781	
Median	-0.0026	-0.0013	0.0000	-0.0036	0.0081	-0.0197	-0.0649	-0.3371	-0.1895	
75th Percentile	-0.0008	-0.0003	0.0001	-0.0012	0.0205	-0.0068	-0.0225	-0.1169	-0.0657	
<i>112 markets with 4 or fewer firms</i>										
25th Percentile	-0.0106	-0.0063	-0.0001	-0.0140	0.0018	-0.0522	-0.1722	-0.9016	-0.5044	
Median	-0.0034	-0.0017	-0.0000	-0.0045	0.0070	-0.0168	-0.0555	-0.2875	-0.1620	
75th Percentile	-0.0007	-0.0003	0.0000	-0.0011	0.0216	-0.0042	-0.0140	-0.0723	-0.0407	
<i>111 markets with 11 or more firms</i>										
25th Percentile	-0.0129	-0.0079	-0.0000	-0.0165	0.0011	-0.0574	-0.1895	-0.9881	-0.5541	
Median	-0.0040	-0.0023	-0.0000	-0.0052	0.0079	-0.0191	-0.0632	-0.3278	-0.1844	
75th Percentile	-0.0006	-0.0003	-0.0000	-0.0008	0.0238	-0.0027	-0.0091	-0.0473	-0.0266	
<i>58 markets with changes greater than \$1,000</i>										
25th Percentile	-0.0209	-0.0142	-0.0003	-0.0250	0.0000	-0.0488	-0.1613	-0.8416	-0.4720	
Median	-0.0004	-0.0003	-0.0000	-0.0005	0.0012	-0.0028	-0.0094	-0.0488	-0.0275	
75th Percentile	-0.0000	-0.0000	-0.0000	-0.0000	0.0202	-0.0000	-0.0000	-0.0001	-0.0001	
<i>381 markets with changes less than \$1,000</i>										
25th Percentile	-0.0109	-0.0065	-0.0000	-0.0141	0.0016	-0.0521	-0.1722	-0.8955	-0.5031	
Median	-0.0036	-0.0020	-0.0000	-0.0048	0.0076	-0.0184	-0.0607	-0.3155	-0.1772	
75th Percentile	-0.0007	-0.0004	-0.0000	-0.0010	0.0216	-0.0039	-0.0129	-0.0670	-0.0378	

...leads to a X% change in variable profit.

Notes: This table investigates the approximation error introduced by our policy function approach. 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 in each of the product characteristics by 1% of the mean starting from the prices and characteristics at the optimal subsidy schedule. We update marginal costs and hold prices and all other product characteristics fixed for all plans in the market.

C Derivation of Characteristic-Level Cost Parameters

This appendix derives the expressions for the characteristic-level cost parameters $\gamma_{o,jt}$ and $\gamma_{c,jt}$ used in the estimation of the marginal cost function (11). The key feature of the derivation is that the share function depends on prices, and prices depend on subsidies through $p_{jt} = r_{jt} - \text{sub}_{jt}$. Because the subsidy (3) depends on cost-sharing characteristics $x_{c,jt}$ through the incremental cost function (12), but does not depend on non-cost-sharing characteristics $x_{o,jt}$, the first-order conditions for the two types of characteristics take different forms.

We denote by β_{x_o} and β_{x_c} the subvectors of the demand characteristic coefficient β_x from (9) corresponding to $x_{o,jt}$ and $x_{c,jt}$, respectively. The firm chooses revenues r_{jt} and characteristics x_{jt} for each product $j \in \mathcal{J}_f$ to maximize

$$\pi_f = \sum_{j \in \mathcal{J}_f} (r_{jt} - c_{jt}) s_{jt}.$$

The first-order condition with respect to revenue for plan j yields (22), which we use below as a substitution identity:

$$\sum_{k \in \mathcal{J}_f} (r_{kt} - c_{kt}) \frac{\partial s_{kt}}{\partial p_{jt}} = -s_{jt}. \quad (26)$$

Non-cost-sharing characteristics. Because the subsidy does not depend on $x_{o,jt}$, the price $p_{jt} = r_{jt} - \text{sub}_{jt}$ is unaffected when the firm adjusts $x_{o,jt}$ (holding revenue fixed). Shares therefore respond only through the direct demand channel, $\frac{\partial s_{kt}}{\partial x_{o,jt}} = \frac{\partial s_{kt}}{\partial \delta_{jt}} \beta_{x_o}$, and the cost derivative from the log-linear specification is $\frac{\partial c_{jt}}{\partial x_{o,jt}} = \gamma_{o,jt} c_{jt}$. The first-order condition is

$$\sum_{k \in \mathcal{J}_f} (r_{kt} - c_{kt}) \frac{\partial s_{kt}}{\partial \delta_{jt}} \beta_{x_o} = \gamma_{o,jt} c_{jt} s_{jt},$$

which yields (24) directly.

Cost-sharing characteristics. When the firm adjusts $x_{c,jt}$, the subsidy responds because the incremental cost changes. From (3) and (12),

$$\frac{\partial \text{sub}_{jt}}{\partial x_{c,jt}} = -\frac{1 - \lambda_{jt}}{\lambda_{jt}} \gamma_{c,jt} \tilde{c}_{jt},$$

where $\tilde{c}_{jt} \equiv \exp(c_{0,jt} + \gamma_{c,jt} x_{c,jt} + \omega_{c,jt})$. The price therefore moves: $\frac{\partial p_{jt}}{\partial x_{c,jt}} = \frac{1-\lambda_{jt}}{\lambda_{jt}} \gamma_{c,jt} \tilde{c}_{jt}$. Shares now respond through two channels—the direct characteristic effect and the induced price change:

$$\frac{\partial s_{kt}}{\partial x_{c,jt}} = \frac{\partial s_{kt}}{\partial \delta_{jt}} \beta_{x_c} + \frac{1-\lambda_{jt}}{\lambda_{jt}} \gamma_{c,jt} \tilde{c}_{jt} \cdot \frac{\partial s_{kt}}{\partial p_{jt}}.$$

The first-order condition with respect to $x_{c,jt}$ is

$$\sum_{k \in \mathcal{J}_f} (r_{kt} - c_{kt}) \frac{\partial s_{kt}}{\partial x_{c,jt}} = \gamma_{c,jt} c_{jt} s_{jt}.$$

Expanding the share derivative and grouping terms:

$$\beta_{x_c} \sum_{k \in \mathcal{J}_f} (r_{kt} - c_{kt}) \frac{\partial s_{kt}}{\partial \delta_{jt}} + \frac{1-\lambda_{jt}}{\lambda_{jt}} \gamma_{c,jt} \tilde{c}_{jt} \underbrace{\sum_{k \in \mathcal{J}_f} (r_{kt} - c_{kt}) \frac{\partial s_{kt}}{\partial p_{jt}}}_{=-s_{jt}} = \gamma_{c,jt} c_{jt} s_{jt}.$$

Substituting (26) and collecting $\gamma_{c,jt}$:

$$\beta_{x_c} \sum_{k \in \mathcal{J}_f} (r_{kt} - c_{kt}) \frac{\partial s_{kt}}{\partial \delta_{jt}} = \gamma_{c,jt} s_{jt} \left[c_{jt} + \frac{1-\lambda_{jt}}{\lambda_{jt}} \tilde{c}_{jt} \right],$$

which yields

$$\gamma_{c,jt} = \frac{\beta_{x_c} \sum_{k \in \mathcal{J}_f} (r_{kt} - c_{kt}) \frac{\partial s_{kt}}{\partial \delta_{jt}}}{s_{jt} \left[c_{jt} + \frac{1-\lambda_{jt}}{\lambda_{jt}} \tilde{c}_{jt} \right]}. \quad (27)$$

This is an explicit expression for $\gamma_{c,jt}$: the revenue first-order condition eliminates the indirect price channel, so that the numerator depends only on the demand parameter β_{x_c} and the share derivatives with respect to mean utility δ_{jt} , not on the price derivatives. In both (24) and (27), revenues and shares are data, costs are recovered via (23), and $\frac{\partial s_{kt}}{\partial \delta_{jt}}$ is computed from the estimated demand parameters. Equation (24) is applied first; $\tilde{c}_{jt}$ is then computed by stripping the recovered $\gamma_{o,jt}$ from c_{jt} .

Decomposing costs using the observed rebate. The first-order conditions recover the cost coefficients $\gamma_{o,jt}$ and $\gamma_{c,jt}$, and the revenue inversion (23) recovers the marginal cost level

c_{jt} . To complete the decomposition of (11) into its constituent parts, we must identify the baseline cost $c_{0,jt}$ and the characteristic-cost terms. Define the following shorthand:

$$\mathbf{x}_{jt} \equiv \gamma_{o,jt} x_{o,jt},$$

$$\mathbf{y}_{jt} \equiv \gamma_{c,jt} x_{c,jt} + \omega_{c,jt}.$$

The marginal cost identity (11) and the incremental cost identity (12) then give a system of two equations in two unknowns ($c_{0,jt}$ and $\mathbf{y}_{jt}$):

$$c_{jt} = \exp(c_{0,jt} + \mathbf{x}_{jt} + \mathbf{y}_{jt}), \quad (28)$$

$$\text{rebate}_{jt} = \exp(c_{0,jt} + \mathbf{y}_{jt}) - \exp(c_{0,jt}). \quad (29)$$

The knowns are: c_{jt} from the revenue first-order condition (23); $\text{rebate}_{jt} = \lambda_{jt}(B_{jt} - b_{jt})$, which is observed in the data; and $\mathbf{x}_{jt}$, which is computed from the recovered $\gamma_{o,jt}$ and observed $x_{o,jt}$. Solving (28) for $\exp(c_{0,jt} + \mathbf{y}_{jt}) = c_{jt}/\exp(\mathbf{x}_{jt})$ and substituting into (29) yields

$$c_{0,jt} = \ln\left(\frac{c_{jt}}{\exp(\mathbf{x}_{jt})} - \text{rebate}_{jt}\right), \quad (30)$$

and $\mathbf{y}_{jt}$ follows as the residual: $\mathbf{y}_{jt} = \ln(c_{jt}) - \mathbf{x}_{jt} - c_{0,jt}$. In the counterfactual, $c_{0,jt}$ and $\omega_{c,jt}$ are held fixed while the observable terms $\gamma_{o,jt} x_{o,jt}$ and $\gamma_{c,jt} x_{c,jt}$ adjust according to the recovered coefficients.

D Monte Carlo analysis

How well does our approximate counterfactual approach work? In this Appendix, we write a simplified model and simulate data by solving for equilibria explicitly. We use this data to estimate policy functions and approximate outcomes as subsidies change. We then compare the exact and approximated outcomes in terms of both firms' actions and consumer welfare. We also compare the error when prices and product characteristics are *all* approximated by policy functions and the error when *only* product characteristics are approximated by policy functions and prices are solved through first-order conditions.

Markets are denoted by m . Each market has consumers denoted by i and F firms denoted

by f . Each firm is present in each market, and offers J products denoted by j . Each product has a price p_{jm} and an $X \times 1$ vector of non-price characteristics δ_{jm} . Individual-choice-specific utility is $u_{ijm} = \alpha_i p_{jm}^2 + \beta'_i \delta_{jm} + \epsilon_{ijm}$ where α_i is the price sensitivity of consumer i , β_i is i 's vector of characteristic preferences, and ϵ_{ijm} is consumer i 's idiosyncratic preference for product j , assumed to be i.i.d. Type-I Extreme Value.

Marginal costs are constant at the product level and given by $c_{jm} = \exp(-0.3 + \gamma' \delta_{jm}^2 + \nu_f + \omega_m)$ where δ_{jm}^2 is element-wise squaring of product characteristics, γ is a 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 drawn from i.i.d. $N(0, 0.1)$.

Firms receive a subsidy payment b_m . The firm's problem is

$$\max_{p_{fm}, \delta_{fm}} \pi_{fm} = \sum_{j=1}^J \int_i (p_{jm} - c_{jm}(\delta_{jm}) + b_m) s_{ijm}(p_m, \delta_m) di \quad (31)$$

where p_{fm} and δ_{fm} represent the vectors of prices and product characteristics for the firm in that market and s_{ijm} is the (logit) probability that consumer i purchases product j .

Equilibrium in m is a vector (p, δ) for all firms such that each firm's choices solve Equation (31) when taking the competitors' choices and the subsidy as given.

This model could admit both multiple equilibria and trivial equilibria in which firms offer identical products. We address this through the distribution of consumer preferences. If X is the number of product characteristics, we define X consumer types with equal proportions. Let β_{nx} be the x th element of the preference vector for type n consumers. We assume $\beta_{nx} \sim N(2.5, 0.1)$ if $n = x$ and $\beta_{nx} \sim N(0.1, 0.01)$ if $n \neq x$. We set $\alpha_i = \alpha = -2.5$. This ensures that when $J \leq X$, there exists an equilibrium in which each product features a high value of a single characteristic and low values of other characteristics.²⁷

We solve for equilibria by iterating over best responses. We generate data by giving each market a subsidy ranging from 0.25 to 0.75 with equal spacing. We then consider a counterfactual in which the order of subsidies is reversed and implement our approach. After

²⁷This preference configuration also lends itself to computational expediency, as we can start the equilibrium solver with a set of products and characteristics with this structure.

estimating demand and marginal costs per the method outlined in Section 4, we estimate policy functions for each characteristic. Let δ_{fjx} be the x th characteristic of product j for firm f . We fit

$$\delta_{fjx} = \theta_{jx}b_m + FE_f + \epsilon_{fjx} \quad (32)$$

where θ_{jx} is the parameter of interest, FE_f is firm fixed effects, and ϵ_{fjx} is an error term.

We approximate non-price characteristics by using the estimated $\hat{\theta}_{jx}$ with


$$\hat{\delta}_{fjx}^{ctf} = (b_m^{ctf} - b_m)\hat{\theta}_{jx} + \delta_{fjx} \quad (33)$$

where the ctf superscript means ‘counterfactual’. We calculate marginal costs at these characteristics using our cost parameter estimates.

Figure ?? illustrates output for a 50-market run with two consumer types, four firms, two products per firm, and two non-price characteristics. Graph (a) depicts the product design for firm 1’s first product in each market. A higher subsidy lowers prices and increases characteristic 1; variation comes from market-specific cost shocks. The other three graphs illustrate approximation errors when computing counterfactual equilibria. Graph (b) illustrates the error coming from policy function approximation: the reported values are the difference between the approximated characteristics in each market and the true counterfactual characteristics. The noise observed in graph (a) attenuates the estimate of θ_{11} , which means that the policy function overestimates (underestimates) the value of the characteristic when the counterfactual subsidy is low (high). Graph (c) illustrates the difference between approximated prices and true prices under two approaches. First, we use policy functions for prices. These underestimate (overestimate) prices when the counterfactual subsidy is low (high). Second, we solve for prices taking approximated characteristics as given. This approach generates an opposite pattern.

The impact of these patterns is seen in graph (d), which illustrates the error in the consumer welfare in the market as a percentage of the true value. When using the pure policy function approach, the errors in characteristics and prices are ‘additive’: when the counterfactual subsidy is low (high), the product is both too generous (not generous enough) and too cheap (too expensive) relative to the true outcome, implying that welfare is substantially

Figure D.1: Example Monte Carlo simulation output



Notes: 50 markets were simulated, each with four firms and two consumer types. Each firm offered two products with two characteristics. The top-left graph illustrates the policy function for firm 1 product 1 across markets. The top-right graph compares the exact and approximate characteristics under counterfactual subsidies. The bottom-left graph compares prices under an estimation approach and a solution approach. The bottom-right graph compares the consumer welfare in each market under each price approach.

higher (lower) than the true outcome. In contrast, the second approach generates smaller errors, as the errors in prices and characteristics work in the opposite direction.

We define two performance statistics. Let $D(x, y)$ be the Euclidean distance between x and y , and let a_m be the vector of prices and product characteristics for all products in market m . We define $Err^a = \frac{1}{M} \sum_m D(a_m^{Exact}, a_m^{Approx}) / \|a_m^{Exact}\|$. Second, we calculate consumer welfare per Equation (17) and define the mean absolute logarithm error as $Err^{CW} = \frac{1}{M} \sum_m |\log(CW_m^{Exact}) - \log(CW_m^{Approx})|$. Both of these measures are strictly positive. We calculate these both when policy functions are used for prices and when prices are solved.

Results for several runs are reported in Table D.1. Increasing the complexity of markets increases the mean simulation error when the number of markets is low. Increasing the number of markets decreases the error. Even in the worst case, when using the price solver

Table D.1: Monte Carlo simulation results

(a) Approximation error in action space

Firms per mkt.	Prods. per firm	Chars. per prod.	50 Markets		100 Markets		200 Markets	
			Err_{pol}^a	Err_{aug}^a	Err_{pol}^a	Err_{aug}^a	Err_{pol}^a	Err_{aug}^a
2	2	2	0.0026	0.0023	0.0012	0.0009	0.0129	0.0133
2	4	4	0.0026	0.0023	0.0013	0.0009	0.0126	0.0131
2	6	6	0.0025	0.0022	0.0013	0.0008	0.0122	0.0128
2	8	8	0.0025	0.0022	0.0014	0.0009	0.0118	0.0126
4	2	2	0.0029	0.0024	0.0016	0.0009	0.0132	0.0137
4	4	4	0.0027	0.0024	0.0015	0.0009	0.0127	0.0134
4	6	6	0.0026	0.0023	0.0014	0.0009	0.0122	0.0131
4	8	8	0.0025	0.0023	0.0014	0.0009	0.0118	0.0128

(b) Absolute logarithm error in consumer welfare

Firms per mkt.	Prods. per firm	Chars. per prod.	50 Markets		100 Markets		200 Markets	
			Err_{pol}^{CW}	Err_{aug}^{CW}	Err_{pol}^{CW}	Err_{aug}^{CW}	Err_{pol}^{CW}	Err_{aug}^{CW}
2	2	2	0.0052	0.0008	0.0019	0.0003	0.0275	0.0046
2	4	4	0.0047	0.0007	0.0023	0.0003	0.0233	0.0040
2	6	6	0.0043	0.0006	0.0024	0.0002	0.0212	0.0035
2	8	8	0.0040	0.0006	0.0023	0.0002	0.0196	0.0033
4	2	2	0.0048	0.0006	0.0030	0.0002	0.0220	0.0032
4	4	4	0.0041	0.0005	0.0025	0.0002	0.0189	0.0027
4	6	6	0.0037	0.0004	0.0023	0.0002	0.0175	0.0024
4	8	8	0.0033	0.0004	0.0021	0.0002	0.0163	0.0022

Notes: In all simulations, the number of consumer types is equal to the number of non-price product characteristics. Err^a 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.


the errors in the action vector and consumer welfare are less than 2% and 1% respectively.

E Endogenizing additional characteristics: A case study

Our counterfactual approximation approach focuses on explicitly endogenizing a single characteristic, per-enrollee revenue, and using policy functions to calculate other product characteristics as a function of counterfactual benchmarks. In the previous appendix we investigate the performance of this approach in a Monte Carlo environment. Here, we investigate the performance in our application with a ‘case study.’ In particular, we endogenize two characteristics, price and hospital copays, and use policy functions for the remaining characteristics. We calculate the optimal county-level benchmark policy for this specification, our preferred specification, and a specification in which we hold product characteristics constant and solve for counterfactual benchmarks alone. Due to the computational complexity involved in this approach, we limit our sample to 10 markets chosen at random from the subset of markets in our data with fewer than 20 plans.

The optimal policies are reported in Table E.1. In general, the policies generated by the “two endogenous characteristics” specification are similar to those generated by our preferred specification (with one endogenous characteristic). Outcomes are reported in Figure E.1. For this figure, outcomes are calculated using the relevant counterfactual model (i.e. we evaluate outcomes for the “Nash equilibrium in revenue” specification by using the “Nash equilibrium in revenue” model itself). The three models find optimal policies that result in similar increases to consumer surplus, though the “two endogenous characteristics” specification performs slightly better.

Figure E.1: Compensating variation under various counterfactual approaches



Notes: This chart illustrates the mean compensating variation for individuals in various groups under the policy that solves Equation (4) for different counterfactual approaches. Markets are limited to the sample reported in Table E.1. All dollars are 2017 dollars. White, Black, and Hispanic groups are defined by CMS.

Table E.1: Policies for 10 markets

FIPS	2017 policy	Optimal policy when using		
		Nash Eq. in revenue	Nash Eq. in rev. plus policy funcs.	Nash Eq. in rev. and hosp. copay plus policy funcs.
06005	\$10591	\$11266	\$11392	\$11919
06113	\$10061	\$10057	\$10058	\$10053
13075	\$9714	\$8943	\$8935	\$8966
21145	\$9820	\$9152	\$9147	\$9163
22077	\$9885	\$9217	\$9212	\$9228
29201	\$9579	\$9156	\$9150	\$9160
37155	\$9294	\$9420	\$9421	\$9427
48193	\$10309	\$10502	\$10507	\$10475
53045	\$9836	\$10156	\$10162	\$10137
55061	\$9555	\$7500	\$7500	\$7500

F Additional Counterfactual Specifications

F.1 Alternative social welfare functions

The optimal policy detailed in Section 5 generates substantial negative changes for some consumers. It is therefore natural to explore social welfare functions that incorporate measures of equity, which may more closely model the preferences of policy-makers. Consider the family of social welfare functions

$$\max_{\{B_m\}} \left(\zeta \sum_m \int_i CS_{im}(B_m) di \right) - (1 - \zeta) Var(CS) \quad \text{s.t.} \quad \sum_m GovExp_m(B_m) = G. \quad (34)$$

Equation (34) is similar to Equation (4) with the addition of a penalty for variance in surplus parameterized by ζ . We report aggregate outcomes in Table F.1 and illustrate mean surplus by group in Figure F.1. This variance penalty has large effects: a weight of merely 0.01 on $Var(CS)$ results in benchmark decreases nearly everywhere in order to fund increases in a few markets. The aggregate consumer welfare increases to \$3.79 billion. Reducing the weight to 0.001 results in a policy that raises aggregate welfare to \$9.30 billion by increasing the benchmark in 326 counties and raising the total MA share to 54.9%.

We also consider minimizing government expenditures subject to a benchmark floor:

$$\min_{\{B_m\}} \sum_m GovExp_m(B_m) \text{ s.t. } B_m \geq B_m^{2017} \forall m. \quad (35)$$

Where the 2017 MA payments are larger than the TM cost, this is done by reducing the benchmark. However, there are markets where the 2017 policy results in MA payments that are lower than TM costs. Increasing the benchmark results in both intensive and extensive margin payment changes. The government must pay more for consumers who were already enrolled in an MA plan, and it must transfer payments from TM to MA for consumers who switch. In 201 markets, the extensive margin impact dominates, and therefore an increase in the benchmark rate *decreases* total government expenditures. As a result, an unconstrained government can reduce total spending on TM and MA by \$5.2 billion, though at the cost of most of the consumer surplus. If the government is prohibited from reducing benchmarks below 2017 levels, spending is reduced by only \$3.0 billion.

We exclude firm profits from Equation (4) as we view the government’s objective as maximizing citizen well-being. That said, our alternative policies affect firm outcomes. Table F.1 reports firm-level market shares and total variable profits under all policies explored here. The optimal policy increases the absolute shares of each of the five largest insurers and most policies result in increases to aggregate variable profits.

F.2 Alternative policy function specifications

Our counterfactual specification can be thought of as a hybrid approach in the sense that we use policy functions to approximate counterfactual product characteristics, and then solve for Nash Equilibria in revenue taking those product characteristics as given. In this section, we explore two other specifications.

First, we explore the use of policy functions for the revenue choices as well. We start by estimating a policy function for revenue. We use the same specification as for the rest of the product characteristics. The last column of Table 6 reports the results. The mean coefficient is sensible: when benchmarks increase by \$1, plan sponsors increase their per-enrollee revenue by a bit less than a dollar (meaning prices decrease and product characteristics improve).


Figure F.2 compares compensating variation outcomes across the two specifications. The

Table F.1: Aggregate outcomes and firm profits under alternative social welfare functions and consequent benchmark schedules

	(1) 2017 policy	(2) Optimal policy	(3) Eq. 34 $\zeta = 0.99$	(4) Eq. 34 $\zeta = 0.999$	(5) Eq. 35 w/o floor	(6) Eq. 35 w/ floor
MA share (%)	29.8	47.0	37.8	54.9	31.4	40.9
$\sum CS$ (\$ billion)	3.18	10.16	3.79	9.30	3.91	5.04
$\sum GovExp$ (\$ billion)	440.5	440.5	440.6	440.6	435.3	437.5
Total var. profits (\$ bill.)	2.13	4.2	2.65	4.4	2.66	3.13
Mkts. w/ increases	N/A	262	326	347	201	201
Frac. plans w/ \$0 premium	.391	.380	.247	.436	.235	.235
<i>Aetna</i>						
Shr. of Medicare benes. (%)	1.16	3.55	1.83	3.8	2.07	2.3
Shr. of MA enrollees (%)	3.90	7.55	4.83	6.92	6.58	5.63
Total var. profits (\$ bill.)	.077	.309	.167	.351	.190	.202
<i>Blue Cross Blue Shield</i>						
Shr. of Medicare benes. (%)	4.22	9.09	6.33	10.3	5.14	6.59
Shr. of MA enrollees (%)	14.2	19.4	16.7	18.7	16.3	16.1
Total var. profits (\$ bill.)	.234	.809	.380	.743	.388	.436
<i>Humana</i>						
Shr. of Medicare benes. (%)	5.97	8.69	9.12	11.8	5.33	7.59
Shr. of MA enrollees (%)	20.0	18.5	24.1	21.5	17	18.6
Total var. profits (\$ bill.)	.314	.528	.477	.713	.324	.400
<i>Kaiser Permanente</i>						
Shr. of Medicare benes. (%)	1.83	2.49	1.56	2.39	2.33	2.36
Shr. of MA enrollees (%)	6.15	5.3	4.13	4.35	7.43	5.78
Total var. profits (\$ bill.)	.235	.392	.196	.342	.325	.329
<i>UnitedHealth Group</i>						
Shr. of Medicare benes. (%)	6.12	6.72	6.63	8.39	5.94	7.98
Shr. of MA enrollees (%)	20.5	14.3	17.5	15.3	18.9	19.5
Total var. profits (\$ bill.)	.523	.608	.575	.726	.555	.686
<i>All Others</i>						
Shr. of Medicare benes. (%)	10.5	16.4	12.4	18.3	10.6	14.1
Shr. of MA enrollees (%)	35.2	35	32.7	33.2	33.8	34.4
Total var. profits (\$ bill.)	.744	1.56	.853	1.52	.873	1.08

Notes: The subsidy schedule in Column (1) is the 2017 benchmarks. The schedule in Column (2) maximizes consumer welfare. The schedules in Columns (3) and (4) maximize consumer welfare with different penalties on the variance of welfare. The schedules in Columns (5) and (6) minimize government expenditures with a non-binding and binding floor set to the 2017 benchmarks, respectively. Dollars are 2017 dollars. Variable profits are computed via Equation (5) 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.

Figure F.1: Compensating variation under various policies




Notes: This chart illustrates the mean compensating variation for individuals in various groups under the policy that solves Equation (4), policies that solve Equation (34) (consumer welfare maximization with a penalty on the variance of welfare) with different choices of ζ , and policies that solve Equation (35) (government expenditure minimization) with a non-binding floor and with a binding floor set equal to the 2017 policy. The markets with benchmark increases and decreases are not identical between the various policies. All dollars are 2017 dollars. White, Black, and Hispanic groups are defined by CMS.

aggregate results are similar: the preferred specification generates an optimal policy with \$241 in compensating variation per Medicare beneficiary, while the optimal policy under the alternative specification generates \$216. Each demographic group receives increases relative to the policy in the data, though for each group these increases are muted under the alternative specification as compared to the preferred specification.

These results are driven by differences in model predictions about firm behavior at counterfactual benchmarks across specifications and differences in the choice of benchmarks for the optimal policy as a result of those different behavioral predictions. Those different benchmarks are illustrated in Figure F.3, panel (a). While benchmarks in most markets are similar (the correlation coefficient between the optimal policies under both specifications is 0.58), the “revenue policy function” specification results in several markets being pushed to extremes. While we imposed bounds of [7500, 13500] on the benchmarks for both specifications (coming from the range of benchmarks observed in the data), that restriction binds for two markets in the primary specification (out of 439) and 46 for the alternative specification. We conclude that a specification that uses only policy functions is likely to perform worse (or generate implausible results) than a specification which solves for at least part of a Nash Equilibrium. Indeed, if we solve for the Nash Equilibria at the benchmarks proposed by the alternative specification, we find that government spending exceeds the budget constraint by roughly \$10 billion (2.3%).


We also explore a more traditional specification in which product characteristics are held fixed throughout the counterfactual analysis. As with the “revenue policy function” specification, this “fixed product characteristics” specification obtains similar, though less favorable, aggregate results: the optimal policy under this specification generates \$229 in compensating variation per Medicare beneficiary. Turning to Figure F.3, panel (b), the benchmarks in most markets are quite similar as well; the correlation coefficient is 0.93. This is perhaps unsurprising given the results discussed above: since non-premium product characteristics provide roughly 15-20% of the changes in compensating variation in the preferred specification, this alternative specification can “get close” through the price channel alone.

Figure F.2: Compensating variation under the optimal policy for three specifications



Notes: This chart illustrates the mean annual compensating variation for individuals in various groups under the policy that solves Equation (4) for three specifications: our preferred specification where we use policy functions for product characteristics and solve for Nash Equilibria in revenue, a specification where we estimate revenue policy functions from the data as well, and a specification where product characteristics are held fixed and only revenues vary. We calculate compensating variation for all Medicare beneficiaries via Equation (17) and compensating variation conditional on MA enrollment via Equation (18). All dollars are 2017 dollars. White, Black, and Hispanic groups are defined by CMS. We weight all calculations by the MCBS sample weights.

Figure F.3: Comparing optimal benchmarks under the preferred specification and two other specifications



Notes: This scatter plot compares the optimal benchmarks (that solve Equation (4)) under our preferred specification (i.e. with policy functions used for product characteristics only) and alternative specifications. In (a) we use policy functions for the revenue choices. In (b) we hold product characteristics fixed. Each marker represents a market. For all specifications, potential benchmarks were restricted to [\$7500, \$13500] to match the range of benchmarks observed in the data.

Response to Editor (Round 2)

Thank you for your continued engagement with our paper and for the detailed, constructive feedback. We appreciate the time you have invested and believe the paper is substantially improved as a result.

Summary of changes

Before turning to the individual comments, we highlight the main changes made in this round:

1. **Restructured the paper** into a linear progression, per your recommendation. Section 2 now describes the MA program in full (institutional setting, benchmarks, the bidding system, and the subsidy function); Section 3 develops the model (the government’s problem, consumer demand, firm profit maximization, and the policy function approximation); Section 4 covers data and estimation; and Section 5 presents the optimal policy results. The earlier interleaving of model and institutions has been eliminated.
2. **Recast the firm’s problem in terms of per-enrollee revenue**, consistent with how CMS structures bidding in practice and the fact that insurer per-enrollee revenue comes from two endogenous sources: subsidies and premiums. Firms choose per-enrollee revenue r_{jt} and characteristics x_{jt} ; consumer-facing premiums $p_{jt} = r_{jt} - \text{sub}(x_{c,jt}; \cdot)$ and bids are determined residually (page 7). This change lets us write profits, first-order conditions, and the cost inversion uniformly for all plans (including those charging zero premiums).
3. **Overhauled the notation.** Parameters now use dedicated Greek letters (α for price, β for demand, γ for cost, θ for policy functions, Ψ for the demand parameter vector). We have eliminated the double-index structure ($\beta_{pf,1fl} \rightarrow \theta_{1,fl}$), the confusing c_c object, and the time-varying $\gamma_{\xi,jt}$ term; ξ is now absorbed into the plan-specific baseline cost $c_{0,jt}$ consistent with its treatment as pre-determined.
4. **Rewrote the bidding and cost section** for a non-MA-expert reader. We derive the incremental cost function, its relationship to the marginal cost identity, and characteristic-level first-order conditions separately for non-cost-sharing and cost-sharing characteristics (Equations 24 and 25, page 19). The full derivation is in Appendix C.
5. **Sharpened Section 5.** Alternative welfare functions and alternative policy-function specifications have moved to Appendix F. Figures that split results by the direction

of the benchmark change have been combined where the split did not add economic insight, and kept where it does (the directional selection of benchmark changes). We have added new analyses of what the current benchmark formula gets wrong (Section 5.2) and of cross-subsidization across markets (Section 5.3, with new Table 8), replacing number-summarizing prose with economic interpretation.

6. **Clarified the policy function specification.** The premium policy function previously reported in Table 6 has been replaced with a revenue policy function, which is the object used in the alternative specification now reported in Appendix F.2. The endogenous vs. fixed characteristics are now explicitly listed (page 20), and the primary specification uses policy functions only for non-price characteristics, solving for equilibrium revenues.

A detailed response to each of your comments follows.

Comment 1: Linear structure

One the reason the paper is too long is your decision to discuss the optimal policy problem in Section 2 before introducing the Medicare Advantage program (Section 3) and the MA bidding system (Section 4.2). The description of the model in Section 2 makes sense only in the context of the MA program, and this organization forces you to change the notation across sections. For instance, you change the notation for the benchmark B several times in the paper. Equation 1 defines the choice of B_m , but later on you use B_j and B_{jt} . You also use $\bar{B}$ to denote the government budget, but I am not sure there is a direct connection between the benchmark rate and the budget. The definition of the “subsidy” function is also very cryptic. The reader should not have to skip 10 pages ahead to understand equation 2.

My first recommendation is to adopt a more linear structure. First, tell us everything we need to know about the MA program. This includes the benchmark and the subsidy function, the bidding system and target revenue. You should use this section also to introduce the relevant notation (and stick with it throughout the paper!). The second section should introduce each component of the model: (a) government objective function, (b) consumer problem, (c) insurers’ profit maximization problem and policy functions (including a description of the variables under the control). The third section should describe the data and the estimation method/results (more or less what you have in Section 4 minus the bid discussion). And finally the last section describes the optimal policy (more or less as you have it).

Response: We deeply appreciate your clear guidance on the structure of the paper. Earlier versions of the paper attempted to preserve a more general managed-competition framework before introducing all of the institutional details of Medicare Advantage. Your comments helped us recognize that this choice made the paper harder to evaluate: the economic content of the model depends critically on the institutional mapping from benchmarks, bids, rebates, premiums, and plan characteristics. We have therefore reorganized the paper around the MA institutional setting from the outset, allowing the formal framework to follow the economics of the application more directly.

The paper now proceeds as follows:

- **Section 2: The Medicare Advantage Program.** This section describes the institutional environment, introduces the benchmark and subsidy function with all relevant notation, and derives the revenue framing.
- **Section 3: Model.** The government’s problem, consumer demand, firm profit maximization, and the policy function approximation.
- **Section 4: Data and Estimation.** Data description, demand estimation, cost recovery, and policy function estimates.
- **Section 5: Optimal Policy Results.**

We have also resolved the specific notation inconsistencies you identify in this comment. As seen on page 7, the government budget constraint is now G (eliminating the confusion with the benchmark B). We now clearly define the county-level benchmark B_{mt} on page 5 and define the plan-level benchmark in relation to the county-level benchmark B_{jt} on page 6. See our response to Comment 4 for further details on the notation overhaul.

Comment 2: Subsidy, revenue function and zero premium

Response: Before addressing the specific points below, we note that the restructuring of the paper has led to a consistent formulation of the firm’s problem in terms of revenue rather than price throughout the paper. In the previous draft, the firm’s problem was written in terms of p_j for consistency with standard IO notation, with the bidding system sort of tucked away late in the paper. In the revision, since we talk about the bidding system up front, we define price as a function of revenue and

the subsidy (which is a function of product characteristics) on page 7. This then allows us to write the firm’s problem on page 9 in terms of revenue. This formulation is more consistent with the institutional reality: in practice, the software used by firms to submit bids to CMS requires firms to choose revenues; bids and premiums are calculated by the software after accounting for cost data.

The resolutions of the issues discussed below are more easily understood under this formulation of the problem (and this is consistent with how the problem is coded-up in our estimation and counterfactual analyses). Two overall points:

- Per-enrollee revenues (subsidies plus premiums) are observed and positive for all plans regardless of whether the consumer-facing premium is zero, so the standard BLP cost inversion technique applies. (Equation (23), page 19)
- CMS’s rules imply that price is a function of per-enrollee revenue and product characteristics, so the “choice” to offer a zero-premium plan is in fact fully endogenous and does not need to be modeled separately.

Again, throughout the recent history of this project, we have conducted the analysis using per-enrollee revenue as the choice variable – previous editorial direction pushed us towards keeping the paper as general as possible as long as possible. We apologize for the confusion and appreciate the opportunity to reframe and simplify.

Section 4.2 is still a mystery to me. I vaguely understand what you mean by a “cost minimizing” assumption, but I am still having a lot of difficulty understanding how the choice of premium and plan characteristics determine the subsidy that enter the firms’ profit function. I hope you can clarify this in the next round by simplifying the notation and by using language that non-MA experts can understand. In particular, I do not understand how you go from equation 18 to 20, and I do not understand the concept of incremental cost (hopefully you can clarify both). Does equation 20 hold only for plans that bid below the benchmark or for all plans (i.e. is the subsidy zero for plans with positive premium)?

Response:

As part of our re-ordering of the paper, we have moved most of the bidding discussion (minus the specific functional form of the incremental cost function) into Section 2.2. The concept of incremental cost comes from the requirement that MA plans offer cost-sharing terms that are at least “equivalent” to TM’s cost-sharing. The idea is that there exists some baseline cost of offering a plan with identical cost

sharing, and then if plans choose to offer more attractive cost-sharing (represented by the subset of cost-sharing characteristics $x_{c,jt}$), the plan will incur incremental per-enrollee costs on top of the baseline cost. We discuss this in the paper on page 6.

The subsidy formula (now Equation 3) holds for all plans. We have added a clarifying sentence on page 7: the subsidy is weakly positive for all plans, including those that charge a positive premium. The maximum subsidy is the benchmark B_{jt} itself.

In the previous round I asked you to explain two points: (i) how do you identify costs from plans with zero premiums?, and (ii) how does your supply model predict the choice of zero vs positive premium?

On the first point, your cost minimization assumption implies a subsidy function equal to:

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

where $incr_{jt} = \exp(c_{b,ft} + c_{c,jf}(x_{c,jt})) - \exp(c_{b,ft})$. I can see that this representation of the subsidy function allows you to infer something about the marginal cost of plans with zero premium. The optimal choice of “cost-sharing characteristics” changes the subsidy level that the firm gets (which allows you to infer the slope parameters), and you can infer a portion of the firm’s marginal cost by connecting the subsidy level with the marginal cost function. However this is not enough for me to understand how you can infer the cost for all plans (including those with zero premium). Equation 8 suggests that the following variables are unknown and need to be inferred from choices and subsidy level: $\gamma_{x,jt}$ and ω_{jt} . Can you provide a clear mapping from choices and subsidy levels to residuals (i.e. from data to model residuals)? Based on your description of the problem, I am worried that there is not a one-to-one mapping between the data and the model residuals.

Response: Again, before diving in to the details here, we note that the re-ordering you suggested to us has led us to significantly update and simplify the notation used for the cost model. As we are no longer performing a cost regression (as had been the case in previous editorial rounds), but instead are only presenting a specification in which we estimate plan-characteristic level cost parameters, we were able to tighten up this section.

To your point about the mapping: the key observation is that because our model frames firms as choosing per-enrollee revenue (subsidy plus premiums)—not prices—and revenues are observed and positive for all plans regardless of whether the pre-

mium is zero, the standard cost inversion applies uniformly. Specifically:

1. Plan-level costs are recovered from the revenue first-order condition (Equation (23), page 19): $\mathbf{c}_t = \mathbf{r}_t - \Delta_t^{-1} \mathbf{s}_t$. The inputs—revenues, shares, and demand derivatives—are observed or estimated for every plan. Zero-premium plans have $r_{jt} = \text{sub}(\cdot)$, but r_{jt} is still observed.
2. Characteristic-level cost parameters are recovered from first-order conditions on product characteristics. Because the subsidy depends on cost-sharing characteristics but not on non-cost-sharing characteristics, these take different forms: Equation (24) recovers non-cost-sharing cost slopes $\gamma_{o,jt}$, and Equation (25) recovers cost-sharing slopes $\gamma_{c,jt}$. Both use revenues, costs from step (a), shares, and demand derivatives—all available for every plan.
3. Cost residuals then follow from the cost function (Equation (11), page 11) after subtracting the estimated components.

The zero-premium constraint does not affect the observability of revenues or the validity of the first-order conditions. We have added a sentence after Equation (23) making this point explicit. Conditional on demand estimates, there is a one-to-one mapping from the data to the model residuals for all plans.

I am also confused about the the exact relationship between the overall cost function defined in equation 8 and the incremental cost function. I guess that the TM-equivalent plan cost is what you later refer as “some baseline” in your response. This seems to correspond to the intercept of the cost function: $c_{b,ft} = \gamma_f + \gamma_s + \gamma_r + \gamma_m + \gamma_t$. I don’t know what all those gamma’s represent. I presume that f , m and t refer to firm, market and time fixed-effects. What about r and s ?

Response: We have clarified this in the revision. The baseline cost is now written as $c_{0,jt}$ (Equation 11), which absorbs all plan-level cost heterogeneity not attributable to observable characteristics, including firm, star-rating, market, and time effects. The “ r ” and “ s ” were fixed effects related to the star rating (we used two, as there was a separate indicator for the plan being too new to merit a star rating). The incremental cost function (Equation 12) measures the additional cost of providing cost-sharing characteristics $x_{c,jt}$ relative to a plan with baseline cost $c_{0,jt}$.

In your response to my comment (point 4), you wrote: “The way we handle the bidding system (transforming the firm’s problem into a choice of revenue and characteristics that then

maps into a premium) both rationalizes the observed data and makes an interior assumption easier to swallow”. Perhaps I am mis-understanding your description, but equation 21 clearly state that insurers choose premium to maximize profits (not r_{jt}). Moreover, I did not see any mention of an “target revenue” policy function being estimated anywhere in the paper. See for instance Table 7.

Response: In the revision the profit function (Equation 5, page 9) and first-order conditions (Equation (22), page 18) are written in terms of per-enrollee revenue (subsidy plus premium). The bid mechanically determines both the subsidy and the premium, so this formulation is isomorphic to plans choosing per-enrollee revenue directly.

In previous drafts, we reported a premium policy function as part of (now) Table 6; we now report a revenue policy function which is the object used in the alternative specification that has moved to the Appendix. We have clarified the text on page 20.

I also (still) don’t understand how your approach deals with the choice of zero vs positive premium. Are you taking as given which plan offers a zero premium? Or is it a choice variable that depends on the policy instruments (i.e. benchmark)? If the latter is true, how does the model predict the decision of plan to offer a zero vs positive premium plan?

Response: We apologize for the continued confusion on this point. As discussed in Section 2.2 (page 7), firms choose per-enrollee revenue (premium plus subsidy), r_{jt} , and product characteristics x_{jt} ; the consumer-facing premium is determined residually as $p_{jt} = r_{jt} - \text{sub}(x_{c,jt}; \cdot)$. Whether a plan’s premium is zero or positive is therefore an endogenous outcome of the firm’s per-enrollee revenue and characteristic choices, not a separate decision. We have added a clarifying sentence to this effect on page 7.

Comment 3: Policy function

Your description of the policy function estimation is now a lot more transparent. However, I could not find any description of the “policy” variables. What characteristics of plans are predicted by the model? What characteristics are assumed fixed and exogenous?

Response: In the revised draft we state this on page 20. The plan characteristics predicted by policy functions in our counterfactual are: the Part B premium reduction, deductible, primary care copay, specialist copay, hospital copay, enhanced drug coverage, dental coverage, vision coverage, and hearing coverage. Basic drug coverage is held fixed as no plan in our data changed this indicator. Per your previous comments, we treat the unobserved quality ξ_{jt} as exogenous to the benchmark.

Table 7 presents the results of the policy functions. The footnote defines the explanatory variables interacted with the benchmark rate. In the text (page 26) you write that one of these variables is the number of firms in the market (lag). In the table you describe the variable as the lagged values of market average characteristics. Which one is it? If the footnote is correct, which characteristics are used in the regression? Is it simply the lagged dependent variable? If so, how do you account for this in the simulations of the optimal policy counter-factuals?

Response: We have corrected this inconsistency. The variable interacted with the benchmark in (now) Table 6 is the lagged number of MA contracts in the market, along with log median income, demographic shares, and Medicaid enrollment. The regression also controls for the plan’s other characteristics, lagged market averages of each characteristic across other firms’ plans, the lagged number of plans offered by the plan’s parent contract and by competing contracts, and firm and market demographic variables in levels. This is now stated clearly on page 20.

The Table also reports results for the “annual premium” policy function. This regression seems to be estimated on a smaller sub-sample (presumably only plans with positive premium). Do you use this policy function in the simulations, or are you repeatedly solving the optimal premium that solves the Bertrand-Nash FOCs?

Response: You are correct about that policy function. As we did not use that policy function for any additional analysis, we have replaced it with the revenue policy function that we use in the alternative specification (now moved to the Appendix). We have clarified the text on page 20.

For completeness, in our counterfactual, we are solving for the optimal choice of revenue that solves the FOCs.

You describe the results of Table 7 by stating that on average an increase in the benchmark improves plan characteristics. The dispersion numbers suggest that for many plans the sign

goes negative. The standard-deviation is in many cases much larger than the mean. I understand that the policy functions are meant to be reduced-form representations, but it seems hard to reconcile this with the model. What does this imply for the estimated cost parameters? Are the two consistent?

Response:

The large standard deviations in the policy function estimates do not necessarily imply inconsistency with the cost estimates. The cost function captures the technology of providing plan characteristics; the policy functions capture equilibrium responses to benchmark changes, which depend not only on own costs but also on the competitive environment, market position, and the strategies of rivals. Firms with similar cost structures may respond very differently to the same benchmark change if they face different competitive conditions or occupy different strategic positions. A variance decomposition shows that firm identity accounts for the majority of the variation in policy function coefficients across most characteristics, consistent with firms maintaining distinct competitive strategies for benefit design (Starc, 2014). The cost estimates and policy function estimates are not derived from the same set of conditions — the former come from first-order conditions of the firm’s problem, while the latter are reduced-form estimates of equilibrium behavior — but they are economically compatible. We have added some brief discussion of this on page 21.

Comment 4: Additional imprecisions

The policy function notation is terribly difficult to understand. You need to change the indexing (for instance $\beta_{pf,1fl}$). You should try to avoid using more than two or three index, and the fact that f shows up twice in the same subscript is very confusing. You should commit to a specific Greek letter to refer to policy-function parameters, a specific greek letter for cost parameters and demand parameters, etc. This comment applies to pretty much all the functions that you are using. If you use β in the demand model, use a different symbol elsewhere rather than adding subscripts.

Response: We agree, and thank you for pushing us on this. We have overhauled the notation to assign a dedicated Greek letter to each parameter family. The following table summarizes the changes:

Parameter type	Before	After
Price sensitivity	$\alpha_0, \alpha_1, \alpha_2$	(unchanged)
Demand	$\beta_w, \beta_x, \dots$	(unchanged)
Cost	$\gamma_f, \gamma_{x,jt}, \dots$	(unchanged)
Policy function	$\beta_{pf,1fl}$	$\theta_{1,fl}$
	$\beta_{pf,2l}$	$\theta_{2,l}$
	$\beta_{pf,3l}$	$\theta_{3,l}$
	$FX_{pf,l}$	FX_l
Estimation param. vector	θ_I	Ψ

Equation 8: The cost function includes $\gamma_{\xi,jt}\xi$, which suggests that the demand residual enters the marginal cost with a time-varying parameter. Is this a typo? I thought the new version of the paper assumed that ξ is pre-determined?

Response: Yes—this was a typo in the previous draft. In the revised cost specification (Equation 11, page 11), the $\gamma_{\xi,jt}\xi$ term no longer appears; because our FOC-based recovery identifies a plan-specific baseline $c_{0,jt}$, any cost component attributable to ξ is absorbed into that term rather than identified as a separate coefficient. We hold $c_{0,jt}$ (and therefore the ξ -component) fixed in the counterfactual, consistent with our treatment of ξ as pre-determined.

Page 16 refers to a “rebate” payment expressed as a percentage of the market-level benchmark. Do you mean φ ? Or λ ? Both? Make sure to introduce the notation when you describe the concept of “plan-level benchmark”, and define the difference between the bonus % and the rebate %. Are both variables exogenous in the model?

Response: We have clarified the notation in Section 2.2. The rebate percentage λ_{jt} (which ranges from 50% to 75% based on the CMS star rating) is exogenous in our model and is now defined where the concept is first introduced. The bonus percentage ϕ_{jt} (which determines the effective plan-level benchmark) is also exogenous and is defined in the same section.

As I explained above, I am still not sure to understand the concept of incremental cost $incr_{jt}$. The text around equation 20 introduces two different cost concepts that appear to be different from the marginal cost c_{jt} defined in equation 8 and 21: $c_{b,ft}$ and $c_{c,jt}$. What are those variables? Once again your choice of index is very confusing (e.g. c_c ?!?).

- What do $x_{c,jt}$ and $\gamma_{xc,jt}$ refer to? Can you clarify in the paper what is the difference between cost-sharing characteristics and x_{jt} ?
- What is the difference between $\omega_{c,jt}$ and ω_{jt} ? Why do you have two separate cost residuals? If they are the same, why do you need to use a subscript c ?
- I didn't find any discussion of how the parameters in the function $c_{b,ft}$ were estimated. This goes back to the point in Comment 2 regarding the estimation the cost parameters.

Response: We agree that c_c was confusing – the thinking was that this represented the cost of cost-sharing characteristics. Obviously that did not come across cleanly.

We have eliminated c_c and rewritten the cost decomposition in the model section (Equation 11, page 11) to clarify these objects. The cost function is now $\ln(c_{jt}) = c_{0,jt} + \gamma_{o,jt}x_{o,jt} + \gamma_{c,jt}x_{c,jt} + \omega_{c,jt}$, where $x_{o,jt}$ denotes non-cost-sharing characteristics (drug, dental, vision, hearing coverage) and $x_{c,jt}$ denotes cost-sharing characteristics (copays, deductibles, Part B premium reductions). The baseline cost $c_{0,jt}$ absorbs all plan-level cost heterogeneity not attributable to observable characteristics, including firm, star-rating, market, and time effects. The residual $\omega_{c,jt}$ captures unobserved costs specific to cost-sharing provision.

This decomposition also simplifies the gamma subscripts per your comment: $\gamma_{o,jt}$ for non-cost-sharing and $\gamma_{c,jt}$ for cost-sharing. The estimation section now presents separate first-order conditions for the two types (Equations 24 and 25, page 19), reflecting the fact that the subsidy depends on cost-sharing characteristics but not on non-cost-sharing characteristics. The non-cost-sharing slopes $\gamma_{o,jt}$ are recovered first; the cost-sharing slopes $\gamma_{c,jt}$ are then recovered using an adjusted denominator that accounts for the subsidy derivative through the incremental cost function (Equation 12, page 12).

The baseline cost $c_{0,jt}$ is then backed out residually: given the overall cost c_{jt} from the revenue FOC and the recovered $\gamma_{o,jt}$ and $\gamma_{c,jt}$, the remaining variation (along with the observed rebates) identifies $c_{0,jt}$ and $\omega_{c,jt}$.

The FOCs for plan characteristics suggest a way to recover the time and firm varying slope parameters. You wrote below equation 23: $\frac{\partial c_{jt}}{\partial x_{jt}} = \gamma_{x,jt}c_{jt}$. Yet the cost function in equation 8 is assumed to be log-linear. Shouldn't $\frac{\partial c_{jt}}{\partial x_{jt}}$ be equal to $\frac{\gamma_{x,jt}}{c_{jt}}$?

Response: No—this is correct as written. The cost function is $\ln(c_{jt}) = \dots + \gamma_{x,jt}x_{jt} + \dots$, so $\frac{\partial \ln c}{\partial x} = \gamma_{x,jt}$. Since $\frac{\partial \ln c}{\partial x} = \frac{1}{c} \frac{\partial c}{\partial x}$, we have $\frac{\partial c_{jt}}{\partial x_{jt}} = \gamma_{x,jt} \cdot c_{jt}$ (not $\gamma_{x,jt}/c_{jt}$).

In the same paragraph you wrote that $\partial \text{sub}/\partial x_{jt}$ can be calculated from the demand parameters. As I wrote above, I don't understand how you construct the subsidy level, but the definition of the incremental cost and the subsidy in equation 20 does not involve the demand parameters. Instead, the slope seems to depend on cost parameter, and on the the variable $\lambda_{f,t}$.

Response: Thank you for catching this imprecision. The characteristics FOCs (Equations (24) and (25), page 19) now make the role of the subsidy explicit: for non-cost-sharing characteristics the subsidy derivative is zero, while for cost-sharing characteristics it enters through the incremental cost function (Equation (12), page 12).

Again, this imprecision was generated in part by our previous attempt to keep the paper as general as possible for as long as possible. The reordering of sections has (we hope!) helped make the exposition easier to follow.

Comment 5: Optimal MA subsidies

In general I find Section 5 well written and clear. However, it is too long, and does not offer much “takeaways.” Most of the discussion in Section 5.1 simply summarizes the numbers, but does not provide particularly useful/interesting economic interpretation. Without expanding the length of the section, I would like you to add a bit more “substance” to the section. Here are a few possible questions that the simulations might shed light on:

What are the implications of your result for the design of the MA program? What does the regulator gets “wrong” in the current environment? Does the optimal policy leads to more or less cross-subsidization between region and/or risk groups? Does the optimal policy lead to better “risk-based” pricing and less misallocation? How do differences in local market concentration (and presumably market-power) affect the optimal geographic dispersion of subsidies?

Response: Thank you for these suggestions—they have significantly improved the section. We have added economic interpretation addressing each of your questions:

- **What the regulator gets wrong.** We have added a paragraph at the end of Section 5.2 (Page 22) explaining that the current benchmark formula misprices market-level subsidies primarily through the fiscal cost channel: it sets

benchmarks too high in markets where raising the subsidy is expensive to the government, and too low where the fiscal cost is small or negative. The current formula does not account for the variance in MA firm costs and instead links benchmarks nearly one-for-one to TM spending.

- **Cross-subsidization.** We have added a paragraph and a new table (Table 8) to Section 5.3 documenting how the optimal policy reallocates subsidies across markets. Markets in the top tercile of TM spending receive benchmark increases in 98% of cases; markets in the bottom tercile receive decreases in 87% of cases. Geographically, the optimal policy shifts subsidies toward the Northeast and toward urban markets. This reflects the fiscal logic of the derivatives analysis: where TM costs are high relative to MA costs, shifting enrollees from TM to MA generates savings.
- **Market concentration.** We find that concentration plays a secondary role. While markets with more competing contracts are somewhat more likely to receive benchmark increases, this relationship is modest after accounting for TM cost levels. We discuss this in the same paragraph.

In order to cut down on pages, I think you need to present fewer results and interpret them better. For instance, I am not sure what is the point of presenting results separately for positive and negative change in the benchmark. This makes the figures unnecessarily long. The analysis of alternative social welfare functions (5.4) is potentially interesting, but the discussion does not lead to any takeaways. Why do we care about this? The same criticism applies to Section 5.5. I didn't see the point of this section.

Response: We have made the following changes to streamline Section 5:

- **Combined figures.** Figures 3 and 5 no longer split results by the direction of the benchmark change. This reduces the complexity of the figures and focuses the reader on the overall patterns.
- **Figure 6 kept split.** We retained the increase/decrease split for the derivatives figure because the split shows the insight: the separation of total surplus derivatives across markets selected for increases vs. decreases illustrates the optimizer's selection rule and helps show what the current policy gets wrong.
- **Sections 5.4 and 5.5 moved to appendix.** The analysis of alternative

social welfare functions and alternative policy function specifications are now in Appendix F. The body of the paper focuses on the primary specification.

- **Rewritten prose.** We have replaced number-summarizing prose with economic interpretation throughout, focusing on pass-through rates, the cost-of-provision channel, the fiscal logic of the optimal reallocation, and the cross-subsidization patterns described above.