Double Duty: A Subsidy Design for Quality Promotion and Price Discipline*

Colleen M. Carey† Ying Fan‡ Yiyi Zhou§

Cornell University University of Michigan Stony Brook University

NBER CEPR, NBER

September 11, 2026

Abstract

This paper studies how the design of quality-linked subsidies shapes firm behavior in both quality choices and pricing decisions. We examine a major reform in Medicare Advantage that links subsidies to quality through two distinct policy parameters. We develop and estimate a structural model of demand and firm behavior to evaluate the reform, decompose its effects, and compare alternative quality-linked subsidy designs. A central finding is that a subsidy design that links product quality to both the subsidy level and the pass-through from firm prices to subsidized consumer prices can simultaneously raise quality and lower prices.

Keywords: subsidy design; quality incentives; imperfect competition; Medicare Advantage

JEL Codes: D4, H5, I1, L1, L5

*We thank Benjamin Vatter, Alex Olssen, Amanda Starc, and conference and seminar participants at Berlin IO Day, KU Leuven, Mannheim, the Cowles Conference on Models and Measurement, and the Triangle Micro Conference for helpful comments. All errors are our own.

†Department of Economics & Brooks School of Public Policy, Cornell University, Email: colleen.carey@cornell.edu.

‡Department of Economics, University of Michigan, Email: yingfan@umich.edu.

§Department of Economics, Stony Brook University, Email: yiyi.zhou@stonybrook.edu.

1 Introduction

Governments commonly use subsidies to encourage quality provision in markets of public concern. When subsidies increase with product quality, firms have stronger incentives to provide quality. At the same time, however, firms may also raise prices in response to the outward shift in demand induced by the subsidy, capturing part of the subsidy through higher markups.

In this paper, we study a subsidy design that simultaneously encourages quality provision and disciplines prices. The key feature of this design is that higher quality makes consumer demand more sensitive to firms' prices. Thus, as quality improves in response to the incentive, demand becomes more sensitive to firm prices, which disciplines firms' pricing incentives. We show that a quality-linked subsidy that adds a quality-price interaction can perform "double duty" by both incentivizing quality provision and disciplining prices.

Our setting is Medicare Advantage (MA), a public-private health insurance market in which firms choose prices and quality and compete for enrollment. MA plans collect part of their prices (known as bids) from the government and the remainder through consumer premiums. Thus, government payments can be viewed as subsidies, and consumer premiums as subsidized prices. In 2012, the MA program enacted the Quality Bonus Program (QBP), which created two quality-linked subsidy incentives. The first is a standard quality incentive: higher quality ratings increase a "bonus rate" subsidy parameter which raises the level of subsidies for any given bid. The second incentive performs the double duty described above: higher quality ratings increase a "rebate rate" subsidy parameter, which not only affects subsidy levels but also governs the pass-through from bids (firm prices) to premiums (subsidized consumer prices). Specifically, a one-dollar increase in a plan's bid raises its premium by the rebate rate. We study these two quality incentives and compare their effects on pricing incentives and welfare.

Our analysis begins with a conceptual discussion of MA subsidies and how each of the two subsidy parameters shapes firm incentives. Since the bonus rate affects only the level of subsidies, whereas the rebate rate also changes the pass-through from bids to premiums, changes in the rebate rate affect firm incentives through additional channels. Consider changes in the levels of the subsidy parameters. An increase in either parameter lowers premiums and raises demand, which tends to increase bids. An increase in the rebate rate, however, also increases the bid sensitivity of demand, pushing bids downward. Next we analyze the effects of linking these two parameters to quality as

in QBP. When either parameter is linked to quality, quality begins to affect demand not only directly through consumer preferences but also indirectly through the subsidy system, strengthening firms' incentives to provide quality. Because higher quality increases demand, bids tend to rise. Linking the rebate rate to quality, however, also creates an association between higher quality and greater bid sensitivity, which disciplines firms' pricing incentives. Thus, linking the rebate rate to quality can simultaneously encourage quality provision and restrain prices.

While this conceptual discussion clarifies how subsidy parameters and their linkage to quality affect firm incentives, it also shows that the ultimate effects of quality-linked subsidies are empirical questions. Thus, the overall effects of QBP, the contributions of its constituent policy changes, and the relative performance of alternative quality-linked subsidy designs are ultimately empirical questions.

To address these questions empirically, we develop and estimate a structural model of MA demand and supply that captures quality incentives, pricing incentives, and their interaction. On the demand side, consumers choose either Traditional Medicare or an MA plan based on plan quality, premiums, supplemental benefits, and other plan characteristics, with heterogeneous preferences across demographic groups. On the supply side, insurers compete oligopolistically. MA plans choose bids, plan quality, and supplemental benefits. We estimate the model using both aggregate plan-level data on plan characteristics and enrollment and individual-level data on consumer demographics and plan choices from 2010 to 2017.

Using the estimated model, we conduct three sets of counterfactual analyses. We compare two alternative subsidy designs in which only one subsidy parameter—either the bonus rate or the rebate rate—is linked to quality. We calibrate these two subsidy policies to keep them budget-neutral. We find that linking quality to the rebate rate is more efficient: the rebate-rate linkage generates higher quality, lower bids and premiums, larger gains in both consumer and producer surplus, and a more favorable distribution of surplus toward consumers. These findings are consistent with the mechanisms highlighted in the paper: linking quality to the rebate rate not only rewards quality but also makes demand more price-sensitive at higher levels of quality. This additional price discipline lowers bids and shifts a larger share of the gains toward consumers.

We also evaluate the overall effects of QBP and decompose them into the effects of its constituent policy changes. The reform consists of several changes to subsidy parameters: it linked both the bonus rate and the rebate rate to quality ratings while also weakly

raising bonus rates and strictly lowering rebate rates overall. Taken together, we find that QBP improves MA plan quality but also induces plans to raise bids and reduce supplemental benefits. Overall, QBP improves both consumer and producer surplus while reducing government spending. Through decomposition exercises, we show that reducing rebate rates contributes to the reduction in government spending, although it also decreases total surplus by more than the associated government savings. In contrast, linking subsidies to quality, on its own, generates gains in total surplus twice the increase in government spending.

Finally, we compare settings with no firm response, partial firm response, and full firm response, highlighting the importance of accounting for firms' endogenous responses in both quality and pricing. We find that ignoring firm responses can yield qualitatively incorrect predictions about the effects of QBP. Moreover, firms respond not only along the quality margin—the direct target of QBP—but also through bids and supplemental benefits. Ignoring these additional margins leads to misleading welfare conclusions: it overstates consumer gains and understates producer gains.

This paper contributes to the literature on subsidies, which emphasizes that the welfare effects of subsidy programs depend not only on consumer responses but also on firms' endogenous reactions to subsidies.[1] As a result, subsidy design plays a central role in determining welfare outcomes. We study a subsidy design that links product quality not only to subsidy levels but also to the pass-through from firm prices to consumer prices. We show that such a design can promote quality while limiting subsidy capture through higher firm prices.

This paper also contributes to the literature on endogenous quality choice. While much of this literature focuses on how market power shapes firms' quality choices,[2] we study how subsidy design shapes quality incentives. The subsidy design we study embeds quality incentives in subsidies, adding a positive relationship between quality and price sensitivity in addition to any quality-price interaction arising from consumer preferences.

We also contribute to a growing literature on the welfare effects of subsidy design in MA.[3] Like these papers, we incorporate firms' endogenous responses to subsidy policy and study how those responses affect both welfare and its distribution between firms and

[1]For example, Weyl and Fabinger (2013); Cabral, Geruso and Mahoney (2018); Polyakova and Ryan (2021); De Groote and Verboven (2019); Decarolis, Polyakova and Ryan (2020); Fan and Zhang (2022).

[2]For example, Fan (2013); Crawford, Shcherbakov and Shum (2019); Fan and Yang (2020).

[3]For example, Curto, Einav, Levin and Bhattacharya (2021); Cabral, Geruso and Mahoney (2018); Miller, Petrin, Town and Chernew (2023); Zahn (2025); Vatter (2025).

consumers. Our paper is distinct in analyzing quality incentives embedded in subsidy design and in highlighting the separate roles of the two subsidy parameters in shaping firms' incentives for both quality provision and pricing.

Finally, we provide policy insights for a program that accounts for roughly six percent of federal spending. Our paper is the first comprehensive evaluation of QBP, the most recent major reform to MA. We highlight two avenues for improvement. First, concentrating quality incentives on the rebate-rate parameter generates substantial gains because it creates an association between higher quality and greater bid sensitivity, causing bids to fall as quality rises. Second, because the rebate-rate parameter governs the pass-through from bids to premiums, the reduction in rebate rates implemented under the 2012 reform nearly offset the reform's welfare gains. Increasing the rebate rate would strengthen price discipline in MA.[4]

More broadly, our paper speaks to how subsidy design affects quality, prices, and welfare. When subsidies encourage quality provision, the resulting increase in demand may induce firms to raise prices. We show that, relative to a flat subsidy, a quality-linked subsidy can instead reduce prices while improving quality if it increases price sensitivity for higher-quality products. We demonstrate this mechanism in a setting where the same program employs two distinct quality incentives: one that operates through the standard demand-expansion channel and another that additionally increases price sensitivity at higher quality levels. Our results show how subsidy design can simultaneously promote quality and discipline prices, improving efficiency without increasing government spending.

2 The MA Market and the Roles of Subsidy Parameters

This section first provides a simplified description of the Medicare Advantage (MA) market, summarizing its basic structure while abstracting from some institutional details. It then clarifies how subsidy parameters and their linkage to quality affect firm incentives.

2.1 A Simplified Description of the MA Market

Although the MA market is institutionally rich, its basic structure can be analogized to a market with consumption subsidies. MA plans are private insurance plans offered to

[4]Curto, Einav, Levin and Bhattacharya (2021) make a similar point for the pre-QBP period.

Medicare beneficiaries under rules administered by the Centers for Medicare & Medicaid Services (CMS). Each MA plan covers the basic benefits of Traditional Medicare (TM), and many plans also provide supplemental benefits. For basic benefits, a plan submits a price, called the bid (denoted by b), to CMS. CMS pays part of this amount, and consumers pay the rest. Consumers pay for supplemental benefits, denoted by x, in full. Thus, a plan receives a payment from CMS and collects $b + x - CMSpayment$ from consumers. Interpreting the CMS payment as a subsidy, the firm receives the unsubsidized price $b + x$, while consumers pay the subsidized price (i.e., premium) equal to $b + x - subsidy$.

Subsidies and Premiums To determine the subsidy amount, CMS sets a benchmark B, adjusts it by a benchmark bonus rate λ, and compares λB with the plan's bid b. If $b < \lambda B$, CMS pays the bid b plus a fraction $\eta < 1$ of the gap $(\lambda B - b)$ as a "rebate." If $b \geq \lambda B$, CMS sets the subsidy equal to λB. Formally,

$$\text{subsidy}(b; \lambda, \eta) = \begin{cases} b + \eta(\lambda B - b) & \text{if } b < \lambda B, \\ \lambda B & \text{if } b \geq \lambda B. \end{cases}$$

The key subsidy policy parameters are therefore the benchmark adjustment factor λ (known as the bonus rate) and the fraction of the benchmark–bid gap η (known as the rebate rate). An increase in either parameter raises the subsidy. In the data, most plans (94%) bid below the benchmark, so we focus on that empirically relevant case in this section.[5]

The premium paid by consumers is the difference between the unsubsidized price and the CMS subsidy: $p(b; \lambda, \eta) = b + x - \text{subsidy}(b; \lambda, \eta)$. For $b < \lambda B$, this simplifies to

$$p(b; \lambda, \eta) = x - \eta(\lambda B - b) = x - \eta\lambda B + \eta b.$$

Thus, the premium consists of a fixed component, $x - \eta\lambda B$, and a bid-dependent component, ηb. The rebate rate η therefore determines the pass-through from bids to premiums.

Quality Bonus Program (QBP) Beginning in 2012, CMS implemented QBP, which changed both subsidy parameters, λ and η. Before QBP, $\lambda = 1$ and $\eta = 0.75$. After 2012, both depend on plan quality, measured by the star rating r.

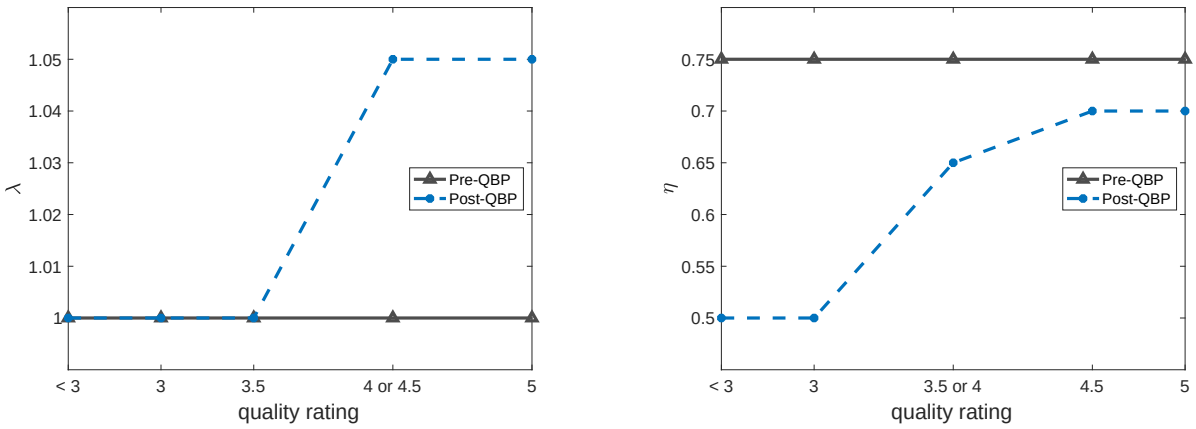
Figure 1 plots the subsidy parameters against r before and after QBP. It shows three changes: both λ and η became increasing functions of quality; λ increased weakly across

[5] In Supplemental Appendix SA, we clarify why a majority of plans bid below the benchmark.

quality levels; and η decreased strictly at all quality levels.

Figure 1: Subsidy Parameters Pre and Post QBP

(a) Bonus Rate λ (b) Rate Rate η



Note: Post-QBP represents 2017 subsidy parameters.

2.2 Roles of Subsidy Parameters

To illustrate how subsidy parameters and their linkage to quality affect firm incentives, we consider changes in the subsidy parameters both in levels and in their dependence on quality. Because subsidy parameters enter profits only through premiums in demand,[6] we focus on how they affect demand.

Consumer demand and effective demand To fix ideas, consider a demand function $s(p, r)$ that is decreasing in premium p and increasing in quality r. We refer to this as "consumer demand" because it is a function of plan characteristics relevant for consumer choice. We define "effective demand" as a function of firm choices (bid b and quality r) by plugging the premium mapping into consumer demand: $\tilde{s}(b, r; \lambda, \eta) = s\big(p(b; \lambda, \eta), r\big)$. Effective demand therefore embeds both underlying consumer preferences and the institutional rule that maps the bid into the premium. This is the relevant demand function for analyzing firms' incentives.

Level changes in subsidy parameters An increase in either λ or η lowers the premium and raises demand for a given (b, r). Effective demand therefore shifts outward, pushing bids upward.

[6]The profit margin (the unsubsidized price − marginal cost) is independent of subsidies.

An increase in η has an additional effect: because η governs bid pass-through to premiums, it makes premiums more responsive to bids and thus increases the bid sensitivity of effective demand. In other words, it rotates effective demand in a way that increases bid sensitivity, which pushes bids downward, offsetting the outward-shift effect.

Linking subsidy parameters to quality When λ or η depends on quality, effective demand becomes $s\big(p(b;\lambda(r),\eta),r\big)$ or $s\big(p(b;\lambda,\eta(r)),r\big)$. Quality then affects demand both directly through consumer preferences and indirectly through the subsidy system. Adding this indirect channel, which arises from the linkage, therefore affects effective demand and firm incentives.

Specifically, linking either λ or η to quality increases demand sensitivity to quality and therefore strengthens incentives for quality provision: higher quality raises demand not only directly, but also indirectly through lower premiums.

Linking η to quality has a second effect. Because an increase in η also makes effective demand more bid-sensitive, higher quality indirectly raises bid sensitivity through the induced increase in η. This adds a positive association between higher quality and greater bid sensitivity, which pushes bids downward. Intuitively, linking η to quality creates a relationship between quality and bid–premium pass-through, and thus introduces a positive force linking quality and bid sensitivity.

Difference between λ and η The key distinction is that λ affects only the level of premiums, whereas η affects both premium levels and bid pass-through. As a result, there are additional channels through which changes in η affect effective demand and firm incentives.

An increase in either λ or η raises demand and tends to push bids upward. But an increase in η also increases the bid sensitivity of demand, pushing bids downward. Likewise, linking either parameter to quality strengthens incentives for quality provision. Because demand is higher at the induced higher quality, bids tend to increase. However, linking η to quality additionally adds an association between higher quality and greater bid sensitivity, pushing bids downward. Formally, the bid sensitivity is $\tilde{s}_b = s_p p_b = s_p \eta$. While quality r affects $\tilde{s}_b$ through s_p (consumer preference) under both the λ-linkage and the η-linkage, it affects $\tilde{s}_b$ mechanically through η (the bid–premium pass-through) only under the η-linkage. Therefore, the linkage $\eta(r)$ additionally adds a component $s_p \eta_r$ into $\tilde{s}_{br}$ (the relationship between r and $\tilde{s}_b$). Since $s_p \eta_r < 0$, this term makes the slope of effective demand with respect to the bid more negative at higher quality. As a result, linking

the rebate rate η to quality can both encourage quality provision and discipline bids.

3 Institutional Details and Data

In this section, we begin by providing a more detailed description of the MA environment, which motivates our empirical model in Section 4. We then describe our data and document key data patterns.

Throughout the paper, we denote a plan by j, a year by t, a county by c, and a contract (a set of plans offered by the same insurer in the same state) by k. We refer to the parent insurer of a plan as its firm.

3.1 Institutional Details

Medicare is a federal health insurance program for people age 65 or older and younger people with disabilities. Medicare beneficiaries choose to receive Medicare coverage through either MA or TM based on benefits, premiums, and quality.

Benefits All MA plans provide basic benefits bb_{jt}, with cost-sharing set at the TM-equivalent level cs_{jt}. Some MA plans provide supplemental benefits (denoted by x_{jt}), which are benefits beyond TM-covered services, such as reductions in cost-sharing below cs_{jt} and extra services (e.g., vision, dental, and hearing coverage). In addition, most MA plans have a linked Part D plan for prescription drug coverage, which is not included in TM. Since a consumer enrolled in such an MA plan is automatically enrolled in that linked Part D plan, we consider an MA plan as a bundle that includes both medical and prescription drug coverage. Let db_{jt} denote plan j's Part D benefit. Then, the core benefits are $cb_{jt} = bb_{jt} - cs_{jt} + db_{jt}$, and the supplemental benefits are x_{jt}. In our empirical model, we treat the core benefits as exogenous, while modeling the supplemental benefits as an endogenous choice made by the firm.

Bids and Premiums MA plans submit their bids for basic benefits to CMS by filling out a spreadsheet, the embedded formulas of which codify the CMS rules regarding bids, premiums, and subsidies. According to the formulas, plan j's bid for basic benefits is determined as the sum of four components:

$$b_{jt}(\rho_{jt}) = bb_{jt} - cs_{jt} + c_{jt}^{admin} + \rho_{jt}, \tag{1}$$

where c_{jt}^{admin} is administrative expenses, and ρ_{jt} is a positive or negative amount referred to as the gain-loss amount. In our empirical model, we treat the first three components as exogenous and assume that plans choose the gain-loss amount ρ_{jt}, which is equivalent to choosing the bid b_{jt} for purposes of our model.[7]

The formulas embedded in the CMS spreadsheet dictate that (i) the premium that consumers pay equals the unsubsidized premium minus the subsidy; (ii) the unsubsidized premium is $b_{jt}(\rho_{jt}) + x_{jt} + d_{jt}$, where d_{jt} is the Part D premium; and (iii) the MA subsidy is determined by comparing an adjusted benchmark with the bid so that it is $b_{jt}(\rho_{jt}) + \eta_{kt}(\lambda_{kt}B_{jt} - b_{jt}(\rho_{jt}))$ if $b_{jt}(\rho_{jt}) \leq \lambda_{kt}B_{jt}$ and $\lambda_{kt}B_{jt}$ otherwise. The combination of (i)–(iii) determines that the premium is:

$$p_{jt}(\rho_{jt}, x_{jt}) = \begin{cases} d_{jt} + x_{jt} - \eta_{kt}(\lambda_{kt}B_{jt} - b_{jt}(\rho_{jt})) & \text{if } b_{jt}(\rho_{jt}) \leq \lambda_{kt}B_{jt}, \\ b_{jt}(\rho_{jt}) + d_{jt} + x_{jt} - \lambda_{kt}B_{jt} & \text{if } b_{jt}(\rho_{jt}) > \lambda_{kt}B_{jt}, \end{cases} \tag{2}$$

where the bonus rate λ_{kt} and the rebate rate η_{kt} vary at the contract level, indexed by k. They are constant 1 and 0.75 before QBP and depend on star ratings after QBP.[8] The term B_{jt} denotes the benchmark for plan j in year t.[9]

Both bids and benchmarks—and as a result, subsidies and premiums—are risk adjusted. See Appendix SD for more details on risk adjustment.

Star Ratings Since 2007, CMS has used "star ratings" as a salient measure of contract quality. First, a continuous quality score q_{kt} is calculated for each contract k as a weighted average of measures across five categories: outcomes, intermediate outcomes, access to care, patient experience, and process.

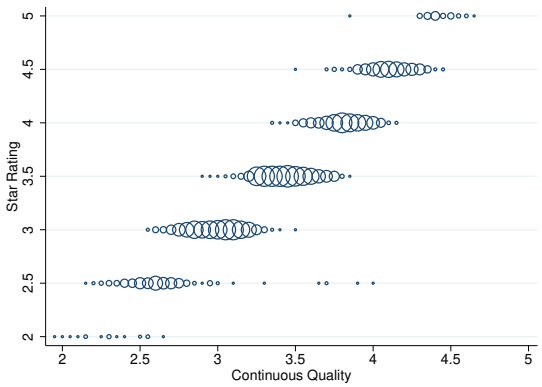
Then, a star rating r_{kt} is determined based on $q_{kt} + \tau_{kt}$, where τ_{kt} is an adjustment factor that captures how the plan's quality compares with that of competing plans as well

[7]Despite the isomorphism, it is cleaner to consider firms choosing ρ_{jt} because b_{jt} actually depends on another endogenous variables, x_{jt}. More precisely, the bid for basic benefits is different from the expression in (1). It is $b_{jt} = bb_{jt} - cs_{jt} + \psi_{jt}(c_{jt}^{admin} + \rho_{jt})$. Correspondingly, the supplemental premium is $x_{jt} + (1 - \psi_{jt})(c_{jt}^{admin} + \rho_{jt})$. Here, $\psi_{jt} = \frac{bb_{jt} - cs_{jt}}{bb_{jt} - cs_{jt} + x_{jt}}$ is the share of the net basic benefits $bb_{jt} - cs_{jt}$ in total benefits (net basic plus supplemental). Therefore, strictly speaking, b_{jt} depends on ρ_{jt} as well as x_{jt} through ψ_{jt}. However, because the ratio ψ_{jt} exceeds 0.9 for most plans, we ignore this detail here for expositional simplicity but fully incorporate it in our empirical analysis.

[8]The full dependence of λ and η on r, which varies in the years after QBP and with a plan's services area, is described in Appendix SB.

[9]This benchmark is calculated by taking the weighted average, across counties, of a proportion of each county's average TM spending, where each county's contribution is weighted by the plan's projected enrollment in that county.

Figure 2: Continuous Quality and Star Rating



Note: This figure reports a scatter plot of star ratings (y-axis) against the continuous quality score q_{kt} binned in 0.05 increments for plan-years used in the analysis.

as with its own quality in the previous year. Formally, the star rating r_{kt} is given by

$$r_{kt} = \text{round}_{0.5}(q_{kt} + \tau_{kt}) \tag{3}$$

where $\text{round}_{0.5}$ indicates rounding to the nearest half. The rating r_{kt} is bounded below by "< 2" and above by 5. We collect data on both q_{kt} and r_{kt} from CMS (CMS, 2021d).

In our empirical model, we model a firm's quality decision as the choice of the quality score q_{kt}, which is then stochastically mapped into a star rating r_{kt}. This is because a firm does not have complete control over its realized star rating in the presence of the adjustment factor τ_{kt}. Figure 2 plots the realized star rating r_{kt} (y-axis) against the binned quality score q_{kt} (x-axis). Each point in the figure corresponds to an observed (q_{kt}, r_{kt}) pair, with the size of the point indicating how many contract-year observations it represents. The figure shows that contracts with similar q_{kt} can receive different realized star ratings, implying that firms may not be in complete control over their realized star ratings.

3.2 Data

Our analysis dataset combines plan-level and individual-level information and includes all plan-counties from 2010 through 2017. Appendix SB provides additional details about the sample and the measurement of key variables.

3.2.1 Plan and County Data

Plan Attributes Our primary source of plan data is the spreadsheets that plans submit to CMS, which are released publicly with a five-year lag (CMS, 2022a). From these spreadsheets, we observe basic benefits bb_{jt}, supplemental benefits x_{jt}, TM-equivalent cost-sharing cs_{jt}, bids b_{jt}, benchmarks B_{jt}, premiums p_{jt}, as well as the bonus rate λ_{kt} and the rebate rate η_{kt}.[10] The bid b_{jt} and benchmark B_{jt} both reflect the plan's risk profile (see Appendix SD for a detailed discussion).

To characterize the prescription drug benefits of MA plans, we collect the Part D premium d_{jt} and invert it using CMS rules to obtain the Part D bid db_{jt}, which represents the expected spending on a risk-neutral beneficiary (CMS, 2021a) and thus measures Part D benefits.[11,12]

Using these data, we derive a risk-neutral core benefits measure to use in demand estimation.

MA Plan Networks A key difference between MA plans and TM is that MA plans typically limit coverage to providers within their contracted networks, whereas TM permits beneficiaries to visit any provider nationwide who accepts Medicare. Following Feyman, Pizer, Shafer, Frakt and Garrido (2024), we infer the scope of an MA plan's network using prescription drug claims. We then classify MA plan-counties as having narrow, medium, or broad networks according to the terciles of this measure. MA plan-counties with insufficient data in the 20% prescription drug claims are assigned "medium".

Characteristics of Traditional Medicare In our empirical analysis, we define TM enrollees as choosing the outside option. Since we treat an MA plan as a bundle of medical and prescription drug benefits, we correspondingly consider the outside option as a bundle that also includes prescription drug benefits and label this outside option as TM. Therefore, the core benefits of TM in county c and year t are $cb_{0ct} = bb_{0ct} - cs_{0ct} + db_{0ct}$. Here, bb_{0ct} is average TM spending in a county (CMS, 2022b), cs_{0ct} is the expected cost-sharing under the TM cost-sharing rules (measurement details in Appendix SB), and db_{0ct} is the Part D benefits of the most popular stand-alone Part D plan in a county (CMS,

[10] For some plans, certain bidding components are censored due to CMS disclosure limitations; we impute these missing entries using the method outlined in Appendix SB.2.

[11] Part D premiums for MA enrollees are commonly offset by rebates as is shown in (2). The Part D premium d_{jt} refers to the Part D premium prior to rebate offsets throughout the paper.

[12] For the small fraction of MA plans that do not offer a linked Part D plan, we assume that enrollees choose the most popular stand-alone Part D plan in the county.

2021c). The TM premium, p_{0ct}, is set equal to the premium of this most popular Part D plan.

County Characteristics We also collect county characteristics at the annual level from various sources. These include county demographics (CMS, 2022b), the number of primary care physicians per capita (AHRF, 2021), the number of specialist physicians per capita (AHRF, 2021), and the Herfindahl–Hirschman Index for hospital systems (AHA, 2016). As a measure of market size in each county-year, we use the total number of Medicare enrollees in that county-year (CMS, 2022b).

Summary Statistics on Plan Characteristics Table 1 reports the mean values of plan-year and plan-county-year-level variables for our analysis sample, separately showing simple averages and averages weighted by plan enrollment. In this sample, 93% of MA plans bid below their benchmark. On average, plans choose a gain-loss amount of $339, which translates into a bid of $9,331 — that is, $1,251 below their bonus-rate-adjusted benchmark. The average unsubsidized premium is $11,982, whereas consumers pay an average subsidized premium of $1,819, which is slightly above the outside option.

On average, plans provide core benefits worth $8,592. This amount reflects basic benefit of $8,991, less $1,420 in cost sharing, plus approximately $1,021 in Part D benefits. The value of their core benefit package is $776 less than the outside option. In addition, they typically provide about $662 in supplemental benefits.

The average star rating across all plan-years is 3.54. Seventy percent of plans provide linked Part D benefits, nearly half offer dental and hearing benefits, and the majority include vision coverage. In addition, 29% of enrollees in the average plan are in a county where the plan has a medium network, while 44% of enrollees in the average plan are in a county where the plan has a broad network. Average enrollment is 1,137 members per plan per county.

A comparison of unweighted and enrollment-weighted averages shows that enrollment is disproportionately concentrated in plans with lower premiums, more generous supplemental benefits, higher star ratings, integrated Part D coverage, and broad provider networks.

Table 1: Summary Statistics on Plan Characteristics

	Unweighted average	Weighted average
Panel A: Bid and Premium ($)		
Bid b_{jt}	9,331	9,358
Gain-loss ρ_{jt}	339	421
$\mathbb{1}$(Bid>Benchmark)	0.07	0.03
Benchmark $\lambda_{kt}B_{jt}$	10,582	10,987
Unsubsidized premium P_{jt}	11,982	12,153
Premium p_{jt}	1,819	1,701
Premium relative to TM_{jct}	68	-82
Panel B: Benefits ($)		
Supplemental benefits x_{jt}	662	811
Core benefits cb_{jt}	8,592	8,582
Basic benefits bb_{jt}	8,991	8,996
Cost-sharing cs_{jt}	1,420	1,420
Part D benefits db_{jt}	1,021	1,006
Core benefits relative to TM_{jct}	-776	-1,142
Panel C: Other Characteristics		
Star rating $r_{j(k)t}$	3.54	3.77
$\mathbb{1}(\text{linked Part D Plan})_{jt}$	0.70	0.82
Share of enrollment in medium network_{jt}	0.29	0.22
Share of enrollment in broad network_{jt}	0.44	0.46
$\mathbb{1}(\text{dental})_{jt}$	0.49	0.50
$\mathbb{1}(\text{vision})_{jt}$	0.84	0.90
$\mathbb{1}(\text{hearing})_{jt}$	0.52	0.63
Panel D: Enrollment and Market Share		
Enrollment_{jct}	1,137	
$\text{Market share}_{jct}$	0.02	
# Plan-year	14,271	
# Plan-county-year	61,461	

Notes: The table presents summary statistics on the characteristics of MA plans from 2010 to 2017. The first column shows unweighted average values, and the second column shows averages weighted by enrollment. Premium relative to TM, core benefits relative to TM, enrollment, and market share are measured at the plan–county–year level, while the remaining variables are measured at the plan-year level.

Table 2: Summary Statistics on Demographic Variables

	All	TM	MA
Age			
< 65	0.157	0.167	0.102
$\in [65, 74)$	0.471	0.467	0.491
≥ 75	0.372	0.366	0.407
High SES	0.542	0.549	0.501
Male	0.448	0.455	0.407
White	0.778	0.793	0.701

Notes: The table presents the summary statistics on the demographic characteristics of Medicare beneficiaries for the years 2010 through 2017, where each observation corresponds to an individual-year. The data includes 12.815 million MA individual-year observations and 68.098 million TM individual-year observations.

3.2.2 Individual-Level Data

We obtain individual-level data on consumer demographics and plan choice from the Medicare Basic Summary File (MBSF), which is a random 20% sample of Medicare beneficiaries.[13] For each Medicare beneficiary in MBSF, we observe her plan choice (including TM), county, age, sex, race, Medicaid enrollment, and zip code. To construct a measure for beneficiary socioeconomic status (SES), we first obtain the Social Deprivation Index for each zip code (Robert Graham Center, 2018; Butler et al., 2013) and then assign a beneficiary high SES if she is not Medicaid-enrolled and does not live in a zip code in the lowest tercile of the Social Deprivation Index.

We report summary statistics on demographics by enrollment in Table 2. The first column reports the summary statistics of the 20% random sample. The comparison of the second and third columns shows that MA enrollment rates differ by demographics; for example, individuals under age 65 (Social Security Disability Insurance recipients) are much less likely to enroll in MA. Moreover, high SES, male, and white beneficiaries are more likely to enroll in TM. These differences are accommodated in our empirical model by demographic-specific valuations of plan characteristics and MA (versus TM) in

[13] We do not use survey data from the Medicare Current Beneficiary Survey (MCBS), as in Miller, Petrin, Town and Chernew (2023) and Vatter (2025). Although the MCBS is attractive for demand estimation because it contains richer demographic information, it covers only 10% of counties, raising concerns for supply-side estimation. For the remaining 90% of counties, we lack the demographic information needed to model plan demand and, in turn, to accurately capture firm incentives for plans whose enrollment in those counties represents a substantial share of a plan's total enrollment.

general.

3.3 QBP and Market Responses to QBP

As previously noted, QBP introduces three main changes: it links bonus and rebate rates to star rating, raises the bonus rate, and lowers the rebate rate. The program began taking effect in 2012, with 2012–2014 serving as a transition phase as the bonus-rate and rebate-rate schedules vary over time. These schedules have been stable since 2015. Appendix Figure SB.1 presents the bonus and rebate rates by star rating for each year from 2010 to 2017.

To see consumer responses to QBP, we plot in Figure 3(a) both the share of contracts with 4 or more stars and the share of total enrollment accounted for by these high-quality contracts. The figure shows that the high-quality contract share increases over time, with a slightly higher slope after QBP. More importantly, the enrollment share of high-quality contracts is larger than the high-quality contract share, and the gap widens following QBP. This pattern is consistent with consumers shifting toward high-quality plans after QBP introduced a negative relationship between quality and premiums.

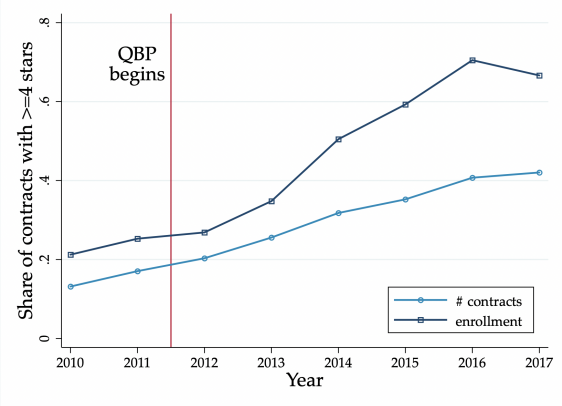
To see firm responses, we plot in Figure 3(b) percentage changes in the average endogenous plan attributes for each year relative to 2010. The figure shows that, on average, by the end of the sample period MA insurers increase their quality by 20%, raise their bids by 10%, and decrease their supplemental benefits by 10%.

However, comparing consumer demand and plan characteristics before and after the implementation of QBP does not identify its causal effect, because market conditions were simultaneously changing along several dimensions. Over this period, the Medicare-eligible population expanded, its demographic composition shifted—for example, the average age declined—and underlying healthcare costs and utilization patterns also changed. These secular trends may affect both consumer demand and firms' plan design decisions, thereby confounding a simple before-and-after comparison.

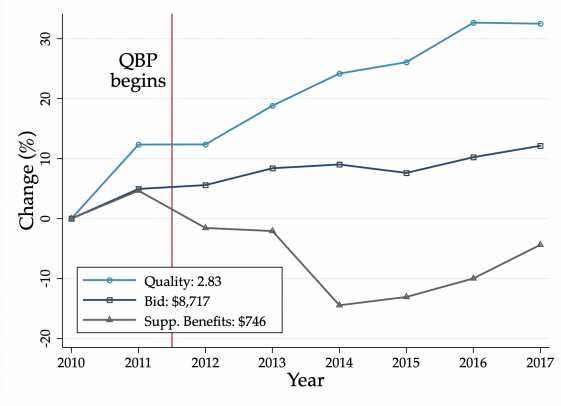
To isolate the effect of the policy from these contemporaneous changes, we next estimate a structural model of demand and supply that explicitly incorporates observed and unobserved heterogeneity across consumers, firms' strategic behavior, and time-varying cost and demographic fundamentals.

Figure 3: Share of High-Quality Plan Enrollment and Plan Attributes Over Years

(a) High-Quality Shares



(b) Plan Attributes



Notes: Figure (a) presents the share of contracts with high star ratings and the share of enrollment in those high-quality contracts from 2010 to 2017. Figure (b) displays the average percentage changes in quality, bid, and supplemental benefits across plans, measured relative to their 2010 baseline levels, when the respective averages are 2.83, $8,717, and $746.

4 An Oligopoly Model of Medicare Advantage

This section develops our full empirical model of the MA market, incorporating the economics of quality-dependent subsidies from Section 2 and the institutional details from Section 3. In our oligopolistic model, consumers (Medicare beneficiaries) choose among MA plans available in the county or the TM outside option. In each market (state-year), a firm (insurer) chooses the quality score q_{kt} for each of its contracts and the gain-loss ρ_{jt} and supplemental benefits x_{jt} for each plan.

4.1 Demand

Consumer i in county c derives utility u_{ijct} from MA plan j in year t:

$$u_{ijct} = \beta_i^p p_{jt} + \beta_i^{cb} cb_{jt} + \beta_i^x x_{jt} + \beta_i^r r_{k(j)t} + \boldsymbol{y}_{jct}\boldsymbol{\beta}^y + \boldsymbol{z}_i\boldsymbol{\theta}^z + \xi_{jct} + \epsilon_{ijct}, \tag{4}$$

where p_{jt} is the premium, cb_{jt} denotes the core benefits, x_{jt} denotes the supplemental benefits, and $r_{k(j)t}$ is the star rating for the contract to which plan j belongs, indexed by $k(j)$. The vector $\boldsymbol{y}_{jct}$ includes dummy variables capturing other plan characteristics, such as network breadth. The vector $\boldsymbol{z}_i$ captures consumer i's demographics. The term ξ_{jct}

captures unobserved demand shocks for plan j in county c in year t, while ϵ_{ijct} captures consumer i's idiosyncratic taste, which we assume to be i.i.d. and to follow a type-I extreme value distribution. We allow consumer tastes for plan premium and characteristics to depend on demographics. Specifically, we assume $\beta_i^p = \bar{\beta}^p + \boldsymbol{z}_i\boldsymbol{\theta}^p$, $\beta_i^{cb} = \bar{\beta}^{cb} + \boldsymbol{z}_i\boldsymbol{\theta}^{cb}$, $\beta_i^x = \bar{\beta}^x + \boldsymbol{z}_i\boldsymbol{\theta}^x$, and $\beta_i^r = \bar{\beta}^r + \boldsymbol{z}_i\boldsymbol{\theta}^r$.

The utility from choosing TM is normalized to be:

$$u_{i0ct} = \beta_i^p p_{0ct} + \beta_i^{cb} cb_{0ct} + \epsilon_{i0ct}, \tag{5}$$

where p_{0ct} and cb_{0ct} are the premium and benefits of TM.

We define the "relative" mean utility, δ_{jct}, as the difference between the mean utility from plan j and that from TM:

$$\delta_{jct} = \bar{\beta}^p(p_{jt} - p_{0ct}) + \bar{\beta}^{cb}(cb_{jt} - cb_{0ct}) + \bar{\beta}^x x_{jt} + \bar{\beta}^r r_{k(j)t} + \boldsymbol{y}_{jct}\boldsymbol{\beta}^y + \xi_{jct}. \tag{6}$$

We collect consumer-specific terms, excluding the idiosyncratic shock, as:

$$\mu_{ijct} = \boldsymbol{z}_i\boldsymbol{\theta}^p \cdot (p_{jt} - p_{0ct}) + \boldsymbol{z}_i\boldsymbol{\theta}^{cb} \cdot (cb_{jt} - cb_{0ct}) + \boldsymbol{z}_i\boldsymbol{\theta}^x \cdot x_{jt} + \boldsymbol{z}_i\boldsymbol{\theta}^r \cdot r_{k(j)t} + \boldsymbol{z}_i\boldsymbol{\theta}^z. \tag{7}$$

Under the type-I extreme value assumption for ϵ_{ijct}, the probability that consumer i in county c chooses plan j in year t is:

$$\frac{\exp(\delta_{jct} + \mu_{ijct})}{1 + \sum_{j' \in \mathcal{J}_{ct}} \exp(\delta_{j'ct} + \mu_{ij'ct})}, \tag{8}$$

where $\mathcal{J}_{ct}$ denotes the set of MA plans in county c in year t.

The market share of plan j in county c in year t is:

$$s_{jct}(\boldsymbol{p}_{ct}, \boldsymbol{x}_{ct}, \boldsymbol{r}_{ct}) = \int \frac{\exp(\delta_{jct} + \mu_{ijct})}{1 + \sum_{j' \in \mathcal{J}_{ct}} \exp(\delta_{j'ct} + \mu_{ij'ct})} dF_z(\boldsymbol{z}_i), \tag{9}$$

where $F_z(\cdot)$ denotes the distribution of demographics $\boldsymbol{z}_i$, the vector $\boldsymbol{p}_{ct}$ collects premiums p_{jt} for all MA plans in county c, and vectors $\boldsymbol{x}_{ct}$ and $\boldsymbol{r}_{ct}$ are defined analogously.

As in Section 2, we define an effective demand function that links consumer demand to firm decisions. First, we substitute the mapping from firm choices (ρ_{jt}, x_{jt}) to premiums $p_{jt}(\rho_{jt}, x_{jt})$ (see (2)). Let $\boldsymbol{p}_{ct}(\boldsymbol{\rho}_{ct}, \boldsymbol{x}_{ct}, \boldsymbol{r}_{ct})$ denote the resulting vector of premiums. Note that the argument $\boldsymbol{r}_{ct}$ reflects that the mapping depends on the subsidy parameters, which in

turn depend on star ratings post QBP. Second, we plug in the stochastic mapping from the firm's quality choice q_{kt} to the star rating r_{kt}, i.e., $r_t(q_{kt} + \tau_{kt})$ in (3), where τ_{kt} follows a normal distribution. Because this mapping rounds $(q_{kt} + \tau_{kt})$ to the nearest 0.5 point, it corresponds to an ordered Probit model. Let $\boldsymbol{r}_{ct}(\boldsymbol{q}_{ct} + \boldsymbol{\tau}_{ct})$ denote the resulting vector of star ratings. The effective market share is therefore:

$$\tilde{s}_{jct}(\boldsymbol{\rho}_{ct}, \boldsymbol{x}_{ct}, \boldsymbol{q}_{ct}) = \mathbb{E}_{\boldsymbol{\tau}_{ct}}\left[s_{jct}\left(\boldsymbol{p}_{ct}(\boldsymbol{\rho}_{ct}, \boldsymbol{x}_{ct}, \boldsymbol{r}_{ct}(\boldsymbol{q}_{ct} + \boldsymbol{\tau}_{ct})), \boldsymbol{x}_{ct}, \boldsymbol{r}_{ct}(\boldsymbol{q}_{ct} + \boldsymbol{\tau}_{ct})\right)\right], \tag{10}$$

and the effective demand equals $M_{ct}\tilde{s}_{jct}(\boldsymbol{\rho}_{ct}, \boldsymbol{x}_{ct}, \boldsymbol{q}_{ct})$, where M_{ct} is the market size.

4.2 Supply

MA firm f in each market (indexed by state s and year t) offers a set of plans denoted by $\mathcal{J}_{fst}$, which are contained within a set of contracts denoted by $\mathcal{K}_{fst}$. Firms simultaneously choose plan-level gain-loss ρ_{jt} and supplemental benefits x_{jt} and contract-level quality score q_{kt}. Let $\mathcal{J}_{kt}$ denote the set of plans belonging to contract k in year t and $\mathcal{C}_{jt}$ denote the service area of plan j in year t.

Firm f's profit maximization problem is:

$$\max_{\{q_{kt}\}_{k\in\mathcal{K}_{fst}}, \{\rho_{jt},\, x_{jt}\}_{j\in\mathcal{J}_{fst}}} \sum_{j\in\mathcal{J}_{kt}} \left[(\rho_{jt} + x_{jt} + a_{jt}) - mc_{jt}(x_{jt}, q_{k(j)t})\right] \sum_{c\in\mathcal{C}_{jt}} M_{ct}\tilde{s}_{jct}(\boldsymbol{\rho}_{ct}, \boldsymbol{x}_{ct}, \boldsymbol{q}_{ct}), \tag{11}$$

where $a_{jt} = bb_{jt} - cs_{jt} + c_{jt}^{admin} + d_{jt}$ captures the exogenous part of the unsubsidized premium, so that $\rho_{jt} + x_{jt} + a_{jt}$ represents the full unsubsidized premium, and $mc_{jt}(\cdot)$ denotes marginal cost.

We assume the marginal cost of each enrollee is:

$$mc_{jt}(x_{jt}, q_{kt}) = (\gamma^x + \omega_{jt}^x)x_{jt} + (\gamma^q + \omega_{kt}^q)q_{kt} + \zeta_{jt}, \tag{12}$$

where (γ^x, γ^q) are parameters to be estimated,[14] $(\omega_{jt}^x, \omega_{kt}^q)$ are shocks to the slope of marginal cost with respect to x_{jt} and q_{kt}, respectively, and ζ_{jt} captures exogenous components such as the cost of providing core benefits.[15]

[14] In Appendix SF we report estimates of a model that allows γ^q that vary over time to accommodate potential changes in the quality score's underlying composition components and their weights (Vatter, 2025). We find a mild time trend which does not affect our core supply-side estimates.

[15] We do not parameterize this term (as a function of core benefits, for example) because it is independent of firm decisions and is unlikely to change in counterfactual scenarios we study.

Letting $Q_{jt} = \sum_{c \in \mathcal{C}_{jt}} M_{ct} \tilde{s}_{jct}(\boldsymbol{\rho}_{ct}, \boldsymbol{x}_{ct}, \boldsymbol{q}_{ct})$ denote the effective demand for plan j, the first-order conditions with respect to the three choice variables $(\rho_{jt}, x_{jt}, q_{kt})$ are, respectively,

$$Q_{jt} + \sum_{j' \in \mathcal{J}_{fst}} (P_{j't} - mc_{j't}) \frac{\partial Q_{j't}}{\partial \rho_{jt}} = 0, \tag{13a}$$

$$\Big(1 - \frac{\partial mc_{jt}}{\partial x_{jt}}\Big) Q_{jt} + \sum_{j' \in \mathcal{J}_{fst}} (P_{j't} - mc_{j't}) \frac{\partial Q_{j't}}{\partial x_{jt}} = 0, \tag{13b}$$

$$- \sum_{j \in \mathcal{J}_{kt}} \frac{\partial mc_{jt}}{\partial q_{kt}} Q_{jt} + \sum_{j' \in \mathcal{J}_{fst}} (P_{j't} - mc_{j't}) \frac{\partial Q_{j't}}{\partial q_{kt}} = 0, \tag{13c}$$

where $P_{jt} = \rho_{jt} + x_{jt} + a_{jt}$ is the unsubsidized premium and $\mathcal{J}_{kt}$ denotes the set of plans within a contract k. In (13), $\frac{\partial Q_{j't}}{\partial \rho_{jt}} = \frac{\partial Q_{j't}}{\partial x_{jt}} = 0$ when plan j' does not overlap in service area with plan j (i.e., $\mathcal{C}_{jt} \cap \mathcal{C}_{j't} = \emptyset$). Similarly, $\frac{\partial Q_{j't}}{\partial q_{kt}} = 0$ when plan j' does not overlap with any plans in contract k.

4.3 Discussion

The model makes several simplifying assumptions. First, the model is static: consumer and firm decisions do not affect future choices. Capturing such dynamics would require a fully dynamic model. Despite these simplifications, the model captures the key economic forces in the MA market. In particular, quality affects profits through three channels: (i) quality directly increases demand; (ii) under QBP, quality indirectly increases demand by raising bonus and rebate rates, thereby reducing premiums; and (iii) providing quality is costly. Thus, the model captures the mechanisms central to our research questions.

Second, within our static model we abstract from a timing issue: in year t, contract k's bonus and rebate rates actually depend on its star rating in the previous year. In our baseline specification, we assume that firms make quality decisions based on current market conditions. This is a good approximation if preferences, costs, and market structure evolve slowly. In Appendix SC, we clarify why this simplification does not affect demand estimation, explain how it can affect supply estimation, and estimate a specification accommodating such a timing by adding an assumption on firms' beliefs about both the market environment and the endogenous choices for all firms in the following period. We find very similar supply-side estimates.

Third, the model abstracts from selection on enrollee risk. Since the risk composition of a plan's enrollees affects both marginal cost and the risk score that CMS uses to scale its plan payments, firms may engage in risk selection. In Appendix SD, we explain in detail how risk scores enter a firm's profit function and clarify that the risk score that depends on actual enrollment affects only one component of the profit function. We also explain why our model captures a reduced-form relationship between marginal cost and enrollment composition. Finally, in a robustness analysis, we no longer hold plan risk scores fixed at their observed values in counterfactual simulations and instead update each plan's risk score to reflect the demographic composition of its enrollment under each counterfactual scenario, using a demographic projection of risk scores recovered from the data. Our results are very similar when accommodating risk score updating.

5 Estimation

5.1 Demand Estimation

We carry out the demand estimation in two steps, as in Goolsbee and Petrin (2004).[16] In the first step, we leverage individual-level data and use maximum likelihood to estimate the parameters that capture heterogeneity in consumer tastes, $\boldsymbol{\theta} = (\boldsymbol{\theta}^p, \boldsymbol{\theta}^{cb}, \boldsymbol{\theta}^x, \boldsymbol{\theta}^r, \boldsymbol{\theta}^z)$. The parameters $\boldsymbol{\theta}^p, \boldsymbol{\theta}^{cb}, \boldsymbol{\theta}^x$, and $\boldsymbol{\theta}^r$ govern how consumer tastes for plan characteristics vary with demographics, while $\boldsymbol{\theta}^z$ governs how consumer preferences for MA plans relative to TM vary with demographics. We also recover the mean utility δ_{jct} for each plan-county-year in this step. In the second step, we use an instrumental-variables approach to regress mean utility δ_{jct} on plan characteristics to obtain estimates of parameters capturing average tastes, $\boldsymbol{\beta} = (\bar{\beta}^p, \bar{\beta}^{cb}, \bar{\beta}^x, \bar{\beta}^r, \boldsymbol{\beta}^y)$.

To conduct the first step, we construct the likelihood function based on each individual's choice probability. We denote the individual choice probability in (8) by $s_{ijct}(\boldsymbol{\delta}_{ct}, \boldsymbol{\theta})$, where $\boldsymbol{\delta}_{ct} = (\delta_{jct}, j \in \mathcal{J}_{ct})$ collects δ_{jct} for all plans in county c in year t. This function depends on $\boldsymbol{\theta}$, while the parameters $\boldsymbol{\beta}$ are absorbed into $\boldsymbol{\delta}_{ct}$. Let j_i denote the observed choice of individual i. The likelihood of observing individual i choosing plan j_i is then $s_{ij_ict}(\boldsymbol{\delta}_{ct}, \boldsymbol{\theta})$.

Rather than jointly searching over $\boldsymbol{\theta}$ and δ_{jct} for all plan-county-year combinations, we invert out $\delta_{jct}(\boldsymbol{\theta})$ and search only over the parameter space $\boldsymbol{\theta}$. Specifically, following Berry

[16] Other studies that use a similar two-step estimation procedure include Train and Winston (2007), Murry and Zhou (2020), and Bachmann, Ehrlich, Fan, Ruzic and Leard (2023).

(1994) and Berry, Levinsohn and Pakes (1995), we invert out $\delta_{ct}(\boldsymbol{\theta})$ by matching predicted market shares to observed market shares for all plans in county c and year t. That is, $\delta_{ct}(\boldsymbol{\theta})$ solves $s_{jct}(\delta_{ct}, \boldsymbol{\theta}) = s_{jct}$ for all $j \in \mathcal{J}_{ct}$, where the right-hand side is the observed market shares from plan-level enrollment data. Plugging $\delta_{ct}(\boldsymbol{\theta})$ into the likelihood yields the log-likelihood function:

$$L(\boldsymbol{\theta}) = \sum_t \sum_c \sum_{i \in \mathcal{I}_{ct}} \log s_{ij_ict}(\delta_{ct}(\boldsymbol{\theta}), \boldsymbol{\theta}), \tag{14}$$

where $\mathcal{I}_{ct}$ denotes the set of consumers in county c and year t in the MBSF data. This step yields estimates of $\boldsymbol{\theta}$ and the mean utilities δ_{jct}.

In the second step, we regress the estimated mean utilities δ_{jct} on plan premiums and characteristics based on equation (6) to estimate $\boldsymbol{\beta}$, the parameters capturing average tastes. In our setting, both the plan premium p_{jt} and the plan characteristics x_{jt} and r_{kt} are potentially endogenous. Premiums are endogenous because they depend on firms' choices of gain-loss amounts and supplemental benefits, which may respond to unobserved demand shocks. Supplemental benefits are directly chosen by firms, while star ratings depend on firms' endogenous quality choices; both may therefore also be correlated with unobserved demand shocks.

To address the endogeneity of the plan premium, we use a plan's Part D benefits as an instrumental variable. Part D benefits are assumed to be uncorrelated with unobserved demand shocks. Moreover, they enter the premium formula directly and therefore provide variation in p_{jt} that is distinct from the variation in x_{jt} and r_{kt}. To address the endogeneity of plan characteristics, we follow Fan (2013) and exploit a key feature of the MA market: while demand is at the county level, premiums and plan characteristics are constant across counties served by the same plan. When choosing plan characteristics, a firm considers market conditions across the areas in which its plans operate. Given the demographics of a county, the demographics of a plan's broader service area affect that demand in that county only through the plan's endogenous characteristics. We construct instrumental variables as market-size-weighted averages:

$$IV_{jt} = \frac{\sum_{c \in \mathcal{C}_{jt}} M_{ct} D_{ct}}{\sum_{c \in \mathcal{C}_{jt}} M_{ct}},$$

where D_{ct} includes county-level demographics such as average age, racial composition of Medicare beneficiaries, per-capita numbers of primary care physicians and specialists,

and the hospital-system Herfindahl-Hirschman index.

Table 3 reports the demand estimation results. Focusing first on the key endogenous plan characteristics—premiums, supplemental benefits, and star ratings—we find that utility declines with premiums and that high-SES consumers are less sensitive than low-SES consumers. The estimated premium coefficient is -1.172 for low-SES consumers and -0.810 for high-SES consumers. We also find that consumers value supplemental benefits

Table 3: Demand Estimates

	Est	SE
Premium ($1,000)	-1.172	0.027
$\times$ High SES	0.362	0.001
Supplemental benefits ($1,000)	0.901	0.067
$\times$High SES	0.122	0.001
$\times$ Age 75+	0.314	0.001
Star rating	0.900	0.025
$\times$ High SES	0.051	0.001
$\times$ Age 75+	0.108	0.000
Core benefits ($1,000)	0.164	0.007
$\times$ Age 75+	0.316	0.000
Other plan characteristics ($\boldsymbol{\beta}^y$)		
$\mathbb{1}$(linked Part D Plan)	0.717	0.013
$\mathbb{1}$(medium network)	0.210	0.015
$\mathbb{1}$(broad network)	0.641	0.014
Demographics ($\boldsymbol{\theta}^z$)		
Age $\in [65, 74)$	0.665	0.001
Age 75+	-0.007	0.001
Female	0.123	0.001
White	-0.265	0.001
High SES	-1.423	0.002

and star ratings, with stronger preferences among high-SES and older consumers. For example, a one-point increase in star ratings is equivalent to a premium reduction of $768 ($= 0.900/1.172 \times 1000$) for low-SES consumers and $\$1,174$ ($= (0.900 + 0.051)/(1.172 - 0.362) \times 1000$) for high-SES consumers.

Turning to other plan characteristics and demographics, we find that consumers value core benefits, with older consumers valuing them more. Consumers prefer integrated Part D coverage and larger networks. Finally, older consumers, males, whites, and individuals with high socioeconomic status derive greater utility from TM relative to MA.

Table 4: Estimates of Cost Parameters

	Est	SE
Supplemental benefits (γ^x)	1.087	0.001
Quality (γ^q)	1.679	0.036

5.2 Supply Estimation

To estimate supply-side parameters, we first recover mc_{jt}, $\frac{\partial mc_{jt}}{\partial x_{jt}}$, and $\frac{\partial mc_{jt}}{\partial q_{kt}}$ from the first-order conditions in (13) and then estimate the parameters using the Generalized Method of Moments.

Table 4 reports the estimated cost parameters. Providing an additional \$1 of supplemental benefits increases marginal cost per enrollee by approximately \$1.09. A one-point increase in the continuous quality score increases marginal cost by \$1,679 per enrollee. On average, marginal cost is \$8,242 per enrollee, implying an average markup of \$1,609, or 17%. These estimates are comparable to those in previous studies: Vatter (2025) reports an average markup of 11.2%, while Curto, Einav, Levin and Bhattacharya (2021) report an average markup of 19%.

In Appendix SF, we assess the robustness of our results by allowing the marginal cost of quality to vary across years. The estimates suggest that the cost of quality varies somewhat over time and shows a slight upward trend, while the estimated cost of supplemental benefit is essentially unchanged. Our main counterfactual exercises under this more flexible specification yield very similar qualitative conclusions.

6 Counterfactual Simulations

In this section, we conduct three sets of counterfactual simulations. In the first set, we compare the effects of incentivizing quality provision through the bonus rate parameter versus through the rebate rate parameter. In the second set, we quantify the overall effects of QBP and decompose these effects into those attributable to each of the program's three changes. In the third set, we highlight the importance of considering firm responses in both pricing and quality when studying the effects of QBP. We conduct all counterfactual simulations for 2017, the last year of our sample.

Throughout this section, we use a no-QBP scenario as the benchmark scenario for comparison. In this counterfactual scenario ("No-QBP"), we restore the subsidy parameters

to their pre-QBP levels, $\lambda^{pre} = 1$ and $\eta^{pre} = 0.75$.

For each counterfactual scenario, we hold the set of MA plans fixed and recompute equilibrium bid b_{jt}, supplemental benefits x_{jt}, and contract-level quality q_{kt}. We simulate star ratings r_{kt} at the equilibrium q_{kt} using the empirical model described in Section 3.1, calculate outcomes such as the MA share at each simulation draw, and then take the average across simulation draws.

For welfare, we report consumer surplus, producer surplus, and government spending. Total government expenditure on the Medicare program equals the sum of CMS payments to MA plans and CMS spending on TM enrollees. We use the average CMS spending per TM enrollee at the county-year level for computing the CMS TM expenditure.[17]

6.1 Incentivizing Quality Provision Through $\lambda(r)$ or $\eta(r)$

As highlighted in Section 2, linking λ to r and linking η to r have different economic implications. Motivated by this insight, we now compare two subsidy designs in which only one subsidy parameter—either λ or η—is linked to quality. Our goal is to assess which approach is more effective at promoting quality and improving welfare without increasing government spending.

Specifically, we conduct two counterfactual simulations. In CF1, we link the bonus rate to star ratings via $\kappa_\lambda \lambda^{QBP}(r)$ while holding the rebate rate at its pre-QBP level, $\eta^{pre} = 0.75$. In CF2, we link the rebate rate to star ratings via $\kappa_\eta \eta^{QBP}(r)$ while holding the bonus rate at its pre-QBP level, $\lambda^{pre} = 1$. We calibrate the scale factors to $\kappa_\lambda = 0.971$ and $\kappa_\eta = 1.163$ so that total government spending in both CF1 and CF2 matches the level in the no-QBP scenario, making each subsidy policy budget-neutral.

Table 5 reports the simulation results, with the upper panel reporting plan-level averages and the lower panel reporting aggregate market outcomes. These result show that linking η to r leads to higher quality and lower prices than linking λ to r. Specifically, column (1) shows that a budget-neutral bonus-rate linkage induces small adjustments in firm decisions: On average, star ratings increase by 0.12, supplemental benefits decrease

[17] Following many papers in the literature, for example, Miller, Petrin, Town and Chernew (2023), our counterfactuals treat the spending of a TM enrollee as fixed at its county-year data value. In other words, we assume that as consumers switch into or out of the MA market, the average TM spending in that county does not change, ruling out changes in the risk profile or realized utilization. This assumption is relaxed in the robustness exercise in Appendix SD, showing only small differences in spending when accounting for risk re-sorting.

Table 5: Bonus-Rate Linkage vs. Rebate-Rate Linkage

	(1) Bonus-Rate Linkage $(\kappa_\lambda \lambda^{QBP}(r),\ 0.75)$ (CF1 − No-QBP)	(2) Rebate-Rate Linkage $(1,\ \kappa_\eta \eta^{QBP}(r))$ (CF2 − No-QBP)
Plan-Level Average Outcomes		
(1) Star rating	0.12	0.34
(2) Supplemental benefits ($)	-5	11
(3) Gain-loss ($)	21	-47
(4) Bid ($)	19	-44
(5) Rebates ($)	-43	28
(6) Subsidy ($)	-28	-8
(7) Premium ($)	44	-28
Aggregate Outcomes		
(8) MA share (%)	1.12	7.28
(9) Government spending ($b)	-0.01	-0.04
(9a) Payment to MA ($b)	4.14	29.56
(9b) Payment to TM ($b)	-4.14	-29.60
(10) Total CS ($b)	0.79	4.56
(11) Total PS ($b)	1.01	2.93

by $5, and bids increase by $19. In contrast, column (2) shows that a budget-neutral rebate-rate linkage leads to a larger increase in star ratings, with an average increase of 0.34, or about 10% of the average star rating in the data. It increases average supplemental benefits rather than reducing them. It also leads to opposite price changes from a budget-neutral bonus-rate linkage: average gain-loss amounts and bids decrease under the rebate-rate linkage, whereas bids increase slightly under the bonus-rate linkage. As a result, the average premium increases under the bonus-rate linkage but decreases under the rebate-rate linkage (row (7)). Combined with the stronger increase in star ratings, the rebate-rate linkage leads to a larger increase in MA enrollment (7.28% versus 1.12%, row (8)).

Turning to welfare, we find that the rebate-rate linkage generates substantially larger welfare gains at essentially the same government spending. Specifically, consumer surplus increases by $4.56 billion under the rebate-rate linkage, compared with $0.79 billion under the bonus-rate linkage, an increase about four times as large. Similarly, producer surplus increases by $2.93 billion under the rebate-rate linkage, compared with a $1.01 billion increase under the bonus-rate linkage. Therefore, the rebate-rate linkage leads to both a larger total efficiency gain and a more favorable division of the gains between

consumers and producers.

Overall, these results indicate that the rebate-rate linkage is more cost-effective: holding government spending fixed, it generates higher quality, lower premiums, larger gains in both consumer and producer surplus, and a more favorable division of total surplus between consumers and producers.

These findings are consistent with the mechanisms described in Section 2. Linking either the bonus rate λ or the rebate rate η to quality strengthens incentives for quality provision. Because higher quality shifts demand outward and thereby increases the marginal profit of raising bids, such a linkage can raise bids. However, when η is quality-dependent, it also creates a positive relationship between quality and the pass-through from bids to premiums. As a result, higher quality makes effective demand more sensitive to bids, which pushes bids downward and shifts a larger share of the gains to consumers.

6.2 Quality Bonus Program

In this section, we quantify QBP's overall effects and to decompose them into the contributions of QBP's three constituent changes: (I) making both λ and η dependent on star rating r; (II) increasing the bonus rate λ on average; and (III) reducing the rebate rate η on average.

Figure 4 summarizes the counterfactual design. The scenario labeled "Data" at the right end of Figure 4 corresponds to the post-QBP environment in 2017, in which each plan's subsidy depends on its star rating r according to $\lambda^{\text{QBP}}(r)$ and $\eta^{\text{QBP}}(r)$. Comparing observed outcomes in "Data" and the simulated outcome in "No-QBP" therefore yields the overall effects of QBP.

Figure 4: CF Design: Effect of QBP and Decomposition



Note: $(\Delta_\lambda, \Delta_\eta)$ solves $avg_r(\lambda^{\text{QBP}}(r) - \Delta_\lambda) = \lambda^{\text{pre}}$ and $avg_r(\eta^{\text{QBP}}(r) + \Delta_\eta) = \eta^{\text{pre}}$.

To decompose these overall effects, we consider two additional counterfactual scenarios, CF3 and CF4. Specifically, we introduce the three changes sequentially as we move from "No-QBP" to "Data": we add Change I when moving from No-QBP to CF3, Change II from CF3 to CF4, and Change III from CF4 to "Data". In CF3, the subsidy

parameters are $\lambda^{\text{QBP}}(r) - \Delta_\lambda$ and $\eta^{\text{QBP}}(r) + \Delta_\eta$, where $\Delta_\lambda = \text{avg}_r(\lambda^{\text{QBP}}(r)) - 1$ and $\Delta_\eta = 0.75 - \text{avg}_r(\eta^{\text{QBP}}(r))$, where the average is taken over the star rating categories. Thus, moving from No-QBP to CF3 incorporates Change I (the positive dependence of subsidies on star ratings) while holding the average values of the subsidy parameters at their pre-QBP levels. In CF4, we further incorporate Change II by increasing the bonus rate by Δ_λ. Finally, moving from CF4 to "Data" incorporates Change III, which reduces the rebate rate by Δ_η.

The nonlinearity of the model implies that, in principle, the order in which each change is implemented could affect the decomposition results. In Appendix SE, we demonstrate robustness to all possible alternative orders.

Table 6: Effects of QBP

	(1) QBP Overall Effect (Data−No-QBP)	(2) Change I Star Rating Dependent (CF3−No-QBP)	(3) Change II Higher Bonus Rate λ (CF4−CF3)	(4) Change III Lower Rebate Rate η (Data−CF4)
Plan-Level Average Outcomes				
(1) Star rating	0.25	0.35	0.06	-0.16
(2) Supplemental benefits ($)	-55	7	21	-83
(3) Gain-loss ($)	92	-15	24	83
(4) Bid ($)	93	-15	18	91
(5) Rebates ($)	-87	174	203	-463
(6) Subsidy ($)	42	182	246	-386
(7) Premium ($)	-5	-190	-201	386
Aggregate Outcomes				
(8) MA share (%)	2.46	11.40	5.86	-14.80
(9) Government spending ($b)	-2.39	6.86	7.02	-16.28
(9a) Payment to MA ($b)	6.28	53.14	30.82	-77.69
(9b) Payment to TM ($b)	-8.67	-46.28	-23.80	61.41
(10) Total CS ($b)	1.71	7.71	4.51	-10.51
(11) Total PS ($b)	2.59	6.03	3.89	-7.33

We find that the QBP, which was designed to incentivize plans to provide higher quality, indeed increases quality. The average star rating rises by 0.25 (row (1), column (1)), which is about 7% of the average star rating and is equivalent to a premium decrease of about $192 for low-SES consumers and $294 for high-SES consumers. This increase in star ratings is driven almost entirely by making subsidies dependent on star ratings (row (1), column (2)). When subsidies depend on star ratings, higher plan quality increases demand not only directly (because consumers prefer plans with higher star ratings) but

also indirectly through higher subsidies and lower premiums (because consumers prefer plans with lower subsidized premiums).

At the same time, QBP reduces average supplemental benefits by $55 (row (2), column (1)), which is about 7% of the average level of supplemental benefits. This reduction is driven by QBP lowering the rebate rate η (row (2), column (4)). Because a lower rebate rate shifts effective demand inwards, it reduces a plan's marginal profit from increasing supplemental benefits.

QBP also leads to higher gain-loss and therefore higher bids (rows (3) and (4), column (1)), again mainly driven by the decrease in the rebate rate η (column (4)). Although a reduction in the rebate rate shifts effective demand inward, it also rotates effective demand so that demand becomes less sensitive to bids. In our simulations, the latter effect dominates, resulting in bids that are $91 higher on average (row (4), column (4)). This rotation effect is strong because a reduction in the rebate rate mechanically reduces the pass-through from bids to premiums, making demand less sensitive to bids.

Average rebates decrease (row (5), column (1)). Specifically, the effects of the three constituent changes in QBP on rebates largely offset one another, yielding an overall change of -$87 per plan on average (about 4% of the average rebate).

Overall, the average subsidy (i.e., CMS payment) increases by $42 per enrollee (row (6), column (1)), while the average premium decreases by $5 (row (7), column (1)). The increase in subsidies is a result of higher bids and higher bonus rates despite lower rebate rates. Premiums barely change because lower supplemental benefits and higher subsidies are offset by higher bids.

Turning to the aggregate outcomes, we find that these changes in plan quality, premiums, and other characteristics increase the share of Medicare beneficiaries enrolled in MA and reduce government spending (rows (8)–(9), column (1)). Specifically, the MA share increases by 2.46 percentage points. The increase in MA enrollment, together with higher subsidies per enrollee, implies that total CMS payment to MA plans rises by $6.28 billion. At the same time, CMS saves $8.67 billion in TM expenditures. The net effect on government spending is therefore a decrease of $2.39 billion.

Overall, QBP increases both consumer and producer surplus while reducing government spending. Consumer surplus increases by $1.71 billion, corresponding to $45 per beneficiary (row (10), column (1)). Producer surplus increases by $2.59 billion (row (11), column (1)).

The decomposition results show that while Change III contributes to the reduction

in government spending, Changes I and II contribute to the increases in consumer surplus and producer surplus. Specifically, Change III (lowering the rebate rate) reduces quality and supplemental benefits while raising bids and premiums, resulting in a $10.51 billion decline in consumer surplus (row (10), column (4)). It also reduces subsidies per enrollee and MA enrollment, leading to a $7.33 billion reduction in producer surplus (row (11), column (4)). While total surplus falls by $17.84 billion, government spending falls by a smaller amount ($16.28 billion in row (9), column (4)). In contrast, Changes I and II increase total surplus more than government spending, making up the lost surplus without undoing savings in government spending. These two changes together increase consumer surplus by $12.22 billion (row (10), columns (2)–(3)), producer surplus by $9.92 billion (row (11), columns (2)–(3)), and government spending by $13.88 billion (row (9), columns (2)–(3)).

In sum, QBP increases MA plan quality, but plans also bid higher and reduce supplemental benefits. Overall, consumer surplus and producer surplus increase while government spending decreases. The decomposition shows that Changes I and II increase total surplus by more than the associated increase in government spending. Moreover, the increase in consumer surplus is about two thirds of the increase in producer surplus. Change III appears necessary to curb government spending growth; however, it has negative welfare effects, reducing total surplus by more than the associated savings in government spending.

6.3 Firm Responses

To highlight the importance of accounting for firm responses in both quality and bids when evaluating the QBP, we now analyze two counterfactual scenarios where firms do not respond to QBP at all in one scenario (CF5) and respond only by adjusting quality—the direct target of QBP—in the other scenario (CF6).

Specifically, in CF5, we restore all endogenous firm variables—quality, supplemental benefits, and gain-loss amounts—to their no-QBP levels and then compute how QBP affects subsidies, premiums, consumer behavior, and welfare outcomes. In CF6, we allow firms to adjust only quality by restoring supplemental benefits and gain-loss amounts to their no-QBP levels and then compute equilibrium quality and the resulting market outcomes under QBP. The comparison between CF5 and No-QBP therefore captures the effects of QBP absent any firm responses. The comparison between CF6 and No-QBP captures the effects of QBP when firms respond only in quality, while holding their other

Table 7: QBP Effects under No, Partial, or Full Firm Responses

	(1) No Firm Responses (CF5−No-QBP)	(2) Partial Firm Responses (Quality) (CF6−No-QBP)	(3) Full Firm Responses (Data−No-QBP)
Plan-Level Average Outcomes			
(1) Star rating	–	0.32	0.25
(2) Supplemental benefits ($)	–	–	-55
(3) Gain-loss ($)	–	–	92
(4) Bid ($)	–	–	93
(5) Rebates ($)	-151	12	-87
(6) Subsidy ($)	-119	61	42
(7) Premium ($)	119	-60	-5
Aggregate Outcomes			
(8) MA share (%)	-2.16	6.55	2.46
(9) Government spending ($b)	-2.70	-1.29	-2.39
(9a) Payment to MA ($b)	-12.73	24.18	6.28
(9b) Payment to TM ($b)	10.03	-25.48	-8.67
(10) Total CS ($b)	-1.08	4.11	1.71
(11) Total PS ($b)	-1.29	1.64	2.59

decisions fixed. We report these results in Table 7, with CF5−No-QBP shown in column (1) and CF6−No-QBP shown in column (2). For comparison, we also reproduce column (1) of Table 6, which reports the effects of QBP with full firm responses, as column (3) of this table.

Column (1) of Table 7 corresponds to the scenario with no firm responses. Rows (5)–(7) confirm that QBP mechanically affects rebates, subsidies, and premiums. In particular, the average per-enrollee rebates and subsidy both decrease. This is primarily driven by the decrease in rebate rates. In the scenario with no firm responses, the averages of firm decisions in rows (1)–(4), of course, do not change. However, their sales-weighted averages change under the QBP. Specifically, the sales-weighted average star rating increases by 0.18, indicating that consumers reallocate from low-quality plans to high-quality plans in response to QBP.

Comparing the no-response scenario in column (1) with the full-response scenario in column (3) shows that ignoring firm responses leads to wrong predictions of the effects of QBP. Column (3) shows that QBP leads firms to increase plan quality, reduce supplemental benefits, and raise bids. These firm responses alter the effects driven by consumer

responses alone: both consumer surplus and producer surplus increase instead of decreasing (rows (10)–(11)), and government savings are smaller (row (9)).

Comparing the partial-response scenario in column (2), in which firms adjust only quality, with the full-response scenario in column (3) shows that ignoring firm responses in bids and supplemental benefits leads to an overestimation of the increase in consumer surplus and an underestimation of the increase in producer surplus. Consumer gains are overstated because this scenario ignores increases in bids and reductions in supplemental benefits, both of which reduce consumer welfare. Producer surplus is understated because this scenario omits key margins through which firms adjust to the new market environment.

Overall, these results demonstrate that the welfare effects of QBP depend on both consumer and firm responses and that evaluating such programs requires accounting jointly for firms' responses in quality, pricing, and other product attributes.

7 Conclusion

This paper studies how subsidy design affects prices, quality, enrollment, and welfare in markets where private firms provide publicly subsidized services. We focus on Medicare Advantage program and its Quality Bonus Program, which links key subsidy parameters to plan quality. Using a structural model of consumer choice and insurers' pricing, quality, and benefit decisions, we quantify the effects of QBP and compare alternative approaches to linking subsidies to quality.

We find that QBP increased plan quality and MA enrollment, while inducing firms to raise bids and reduce supplemental benefits. Overall, consumer and producer surplus increased while government spending declined. Our decomposition shows that linking subsidy parameters to quality accounts for most of the welfare gains, whereas the reduction in the rebate rate restrained government spending but generated welfare losses.

A central finding of the paper is that the design of the linkage matters. Linking the rebate rate—the parameter governing pass-through from bids to premiums—to quality is particularly effective. This design strengthens incentives for quality provision while simultaneously disciplining prices by increasing demand sensitivity to prices at higher quality. As a result, it delivers larger welfare gains and a more favorable distribution of surplus between consumers and firms without an additional fiscal cost.

More broadly, our results highlight the importance of accounting for firms' endoge-

nous responses in both pricing and quality when evaluating subsidy programs. They also illustrate how creating a linkage between quality and both the subsidy amount and the pass-through from firm to consumer prices can promote quality and discipline prices simultaneously. These insights are relevant beyond Medicare Advantage. They provide a principle for subsidy design: linking a desirable product attribute not only to the subsidy level but also to price pass-through can strengthen incentives for that attribute while limiting firms' ability to capture the subsidy through higher prices.

References

American Hospital Association (AHA) (2016) "American Hospital Association Annual Survey: 2012-2015."

Bachmann, Ruediger, Gabriel Ehrlich, Ying Fan, Dimitrije Ruzic, and Benjamin Leard (2023) "Firms and Collective Reputation: a Study of the Volkswagen Emissions Scandal," *Journal of the European Economic Association*, 21 (2), 484–525.

Berry, Steven (1994) "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 25 (2), 242–262.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995) "Automobile Prices in Market Equilibrium," *Econometrica*, 63 (4), 841–890.

Butler, Danielle C., Stephen Petterson, Robert L. Phillips, and Andrew W. Bazemore (2013) "Measures of Social Deprivation that Predict Health Care Access and Need Within a Rational Area of Primary Care Service Delivery," *Health Services Research*, 48 (2), 539–559, Part 1.

Cabral, Marika, Michael Geruso, and Neale Mahoney (2018) "Do Larger Health Insurance Subsidies Benefit Patients or Producers? Evidence from Medicare Advantage," *American Economic Review*, 108 (8), 2048–2087.

Centers for Medicare & Medicaid Services (2015) "Medicare 2015 Part C & D Star Ratings Technical Notes,"Technical report, Centers for Medicare & Medicaid Services, `https://www.cms.gov/files/document/2025-star-ratings-technical-notes.pdf`, Annual series; index of all years at https://www.cms.gov/medicare/health-drug-plans/part-c-d-performance-data. Accessed July 2026.

——— (n.d.) "Medicare Managed Care Manual, Chapter 8: Payments to Medicare Advantage Organizations," Internet-Only Manual Pub. 100-16, Centers for Medicare & Medicaid Services, https://www.cms.gov/regulations-and-guidance/guidance/manuals/downloads/mc86c08.pdf, Accessed July 2026.

Centers for Medicare and Medicaid Services (2021a) "Landscape Files," https://www.cms.gov/medicare/coverage/prescription-drug-coverage.

——— (2021b) "Monthly Enrollment by Contract/Plan/State/County," https://www.cms.gov/data-research/statistics-trends-and-reports/medicare-advantagepart-d-contract-and-enrollment-data/monthly-enrollment-contract/plan/state/county.

——— (2021c) "Part C and D Performance Data," https://www.cms.gov/medicare/health-drug-plans/part-c-d-performance-data.

——— (2021d) "Plan Payment Data," https://www.cms.gov/medicare/health-drug-plans/plan-payment-data.

——— (2022a) "Bid Pricing Tools and Instructions," https://www.cms.gov/medicare/payment/medicare-advantage-rates-statistics/bid-forms-instructions.

——— (2022b) "Geographic Variation Public Use File," https://data.cms.gov/summary-statistics-on-use-and-payments/medicare-geographic-comparisons/medicare-geographic-variation-by-national-state-county.

——— (2026) "CMS Program Statistics - Medicare Part A & Part B - All Types of Service," https://data.cms.gov/summary-statistics-on-use-and-payments/medicare-medicaid-service-type-reports/cms-program-statistics-medicare-part-a-part-b-all-types-of-service.

Crawford, Gregory S., Oleksandr Shcherbakov, and Matthew Shum (2019) "Quality Overprovision in Cable Television Market," *American Economic Review*, 109 (3), 956–995.

Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya (2021) "Can Health Insurance Competition Work? Evidence from Medicare Advantage," *Journal of Political Economy*, 129 (2), 570–606.

De Groote, Olivier and Frank Verboven (2019) "Subsidies and Time Discounting in New Technology Adoption: Evidence from Solar Photovoltaic Systems," *American Economic Review*, 109 (6), 2137–72.

Decarolis, Francesco, Maria Polyakova, and Stephen R. Ryan (2020) "Subsidy Design in Privately Provided Social Insurance: Lessons from Medicare Part D," *Journal of Political Economy*, 128 (5), 1712–1752.

Fan, Ying (2013) "Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market," *American Economic Review*, 103 (5), 1598–1628.

Fan, Ying and Chenyu Yang (2020) "Competition, Product Proliferation, and Welfare: A Study of the US Smartphone Market," *American Economic Journal: Microeconomics*, 12 (2), 99–134, 10.1257/mic.20180182.

Fan, Ying and Ge Zhang (2022) "The Welfare Effect of a Consumer Subsidy with Price Ceilings: the Case of Chinese Cell Phones," *The RAND Journal of Economics*, 53 (2), 429–449.

Feyman, Yevgeniy, Steven D. Pizer, Paul R. Shafer, Austin B. Frakt, and Melissa M. Garrido (2024) "Measuring Restrictiveness of Medicare Advantage Networks: A claims-Based Approach," *Health Services Research*, 59 (1), e14255.

Goolsbee, Austan and Amil Petrin (2004) "The Consumer Gains from Direct Broadcast Satellites and the Competition with Cable TV," *Econometrica*, 72 (2), 351–381.

Health Resources and Services Administration (HRSA) (2021) "Area Health Resources File (AHRF) 2020-2021 Release," `https://data.hrsa.gov/data/download`.

Medicare Payment Advisory Commission (2021) "Rebalancing Medicare Advantage Benchmark Policy," report to the congress: Medicare and the health care delivery system, ch. 1, Medicare Payment Advisory Commission, `https://www.medpac.gov/wp-content/uploads/import_data/scrape_files/docs/default-source/reports/jun21_ch1_medpac_report_to_congress_sec.pdf`, Accessed July 2026.

——— (2024) "Medicare Advantage Program Payment System," payment basics, Medicare Payment Advisory Commission, `https://www.medpac.gov/wp-content/uploads/2024/10/MedPAC_Payment_Basics_25_MA_FINAL_SEC.pdf`, Accessed July 2026.

Miller, Keaton S, Amil Petrin, Robert Town, and Michael Chernew (2023) "Optimal Managed Competition Subsidies," *NBER Working Paper 25616*.

Murry, Charles and Yiyi Zhou (2020) "Consumer Search and Automobile Dealer Colocation," *Management Science*, 66 (5), 1909–1934.

Polyakova, Maria and Stephen P Ryan (2021) "Subsidy Targeting with Market Power," *NBER Working Paper w26367*.

Robert Graham Center (2018) "Social Deprivation Index (SDI)," https://www.graham-center.org/rgc/maps-data-tools/sdi/social-deprivation-index.html.

Scheinker, David, Kevin Schulman, Barak Richman, Arnold Milstein, and Richard Kronick (2026) "Medicare Advantage Plan Bid Risk Underestimates and Corresponding Insurer Revenues," *Health Services Research*, 61 (5), e70158, https://doi.org/10.1111/1475-6773.70158, e70158 HSR-26-0038.

Train, Kenneth E and Clifford Winston (2007) "Vehicle Choice Behavior and the Declining Market Share of US Automakers," *International economic review*, 48 (4), 1469–1496.

Vatter, Benjamin (2025) "Quality Disclosure and Regulation: Scoring Design in Medicare Advantage," *Econometrica*, 93 (3), 959–1001.

Weyl, E Glen and Michal Fabinger (2013) "Pass-through as an economic tool: Principles of incidence under imperfect competition," *Journal of Political Economy*, 121 (3), 528–583.

Zahn, Matthew V. (2025) "Entry and Competition in Insurance Markets: Evidence from Medicare Advantage," *Working Paper*.

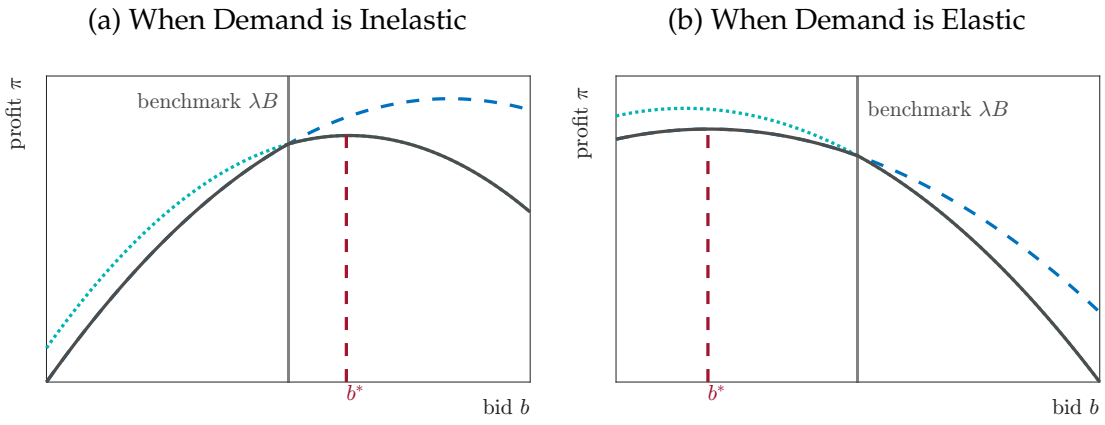
SA Bid Relative to Benchmark

Although this issue is not central to our analysis, we clarify why most plans bid below the benchmark.

At first glance, one might think that plans should always bid above the benchmark. When a plan bids above the benchmark, a one-dollar increase in its bid b leads to a one-dollar increase in the subsidy it receives. By contrast, when a plan bids below the benchmark, a one-dollar increase in b leads to an increase in the subsidy of only $1 - \eta < 1$ dollars. However, this intuition is misleading because, as clarified by the stylized model, a plan's profit margin is independent of the subsidy payment. Whether a plan bids above or below the benchmark, a one-dollar increase in the bid b always increases the profit margin $(b + x - mc)$ by one dollar.

The key consideration instead is the discrete change in the relationship between bid and premiums at the benchmark. This incentivizes plans to bid below the benchmark, because the subsidy affects the premium and thus consumer demand. When the plan bids below the benchmark, a one-dollar increase in b leads to an increase in p of $\eta < 1$ dollars, whereas when the plan bids above the benchmark, the same one-dollar increase in b raises p by one dollar. As a result, the marginal profit of increasing the bid exhibits a discrete drop once b crosses the benchmark from below, which tends to push optimal bids below the benchmark.

Figure SA.1: Profit As Bid Varies: $\pi(b, \cdot)$



That said, it is possible for an optimal bid to be above the benchmark, especially when a plan faces an inelastic demand. Figure SA.1 plots the profit of a plan as a function of b when it faces inelastic demand in the left panel and elastic demand in the right panel.

Specifically, the thick black line in both panels represents the profit function $\pi(b, \cdot)$. To illustrate the kink in the profit function (i.e., the discrete decrease in the slope $\frac{\partial \pi}{\partial b}$), we also plot two auxiliary profit functions, namely $(b + x - mc)s(p, r)$, where p is always $b + x - \lambda B$ (the above-benchmark premium) in the dotted line and p is always $x - \eta(\lambda B - b)$ (the below-benchmark premium) in the dashed line. This figure shows that it is indeed possible for some plans to bid below the benchmark while others bid above it.

SB Details on Sample Construction and Measurement

SB.1 Sample Construction

Our plan-county-year level data include all plans except for regional PPO plans, employer group waiver plans, and special needs plans.[18] We exclude counties with market sizes below five thousand Medicare beneficiaries, because such small markets are unlikely to play a major role in firms' plan-design decisions, and county characteristics described in Section 3.2 are commonly missing for these markets. These excluded counties represent 6% of the total market size. Within each county, we aggregate all plans whose enrollment shares are below 1% of total MA enrollment in that county into a single fringe plan and define its characteristics as the average of such plans. We include the fringe plan in each county when estimating and computing demand, but treat it as non-strategic on the supply side.

For our individual-level data, we exclude individuals who enroll in only Medicare Part A (inpatient/hospital care) but not Part B (outpatient/medical care). Most such individuals are enrolled in employer-sponsored insurance and enroll in Part A (which has no premium) as a secondary payer. Essentially all Part-A-only enrollees are in TM as all MA plans cover both Part A and Part B.

SB.2 Imputing Missing Bidding Variables

The CMS spreadsheets recording bidding data are subject to a censoring rule prior to public release that affects certain bidding variables in certain years. If a plan-year experiences censoring, the variables basic benefits bb_{jt}, cost-sharing cs_{jt} and gain-loss ρ_{jt}

[18] Regional PPO plans serve beneficiaries across states, are subject to different CMS payment rules, and have a total MA share of less than 1%. Employer group waiver plans are offered by employers or unions to their Medicare-eligible retirees with standard Medicare rules waived. Special needs plans require consumers to have specific conditions to enroll.

are set to missing. The variables benchmark B_{jt}, bid b_{jt}, supplemental benefits x_{jt}, stars r_{kt} and non-medical costs c_{jt}^{admin} are still observed in those plan-years. Two-thirds of sample plans experience no missingness in any year. Missingness is not strongly related to plan size or whether a plan is newly introduced.

The majority of missingness represents a single year of missingness for a given plan; all bidding variables are observed for that plan in adjoining years. Our strategy leverages this pattern. Specifically, we predict net basic benefits $bb_{jt} - cs_{jt}$ using $\overline{bb}_j - \overline{cs}_j$, b_{jt}, and year fixed effects, where $\overline{bb}_j - \overline{cs}_j$ represents the average value of $bb_{jt} - cs_{jt}$ for this plan in all years observed, b_{jt} is the bid for this plan-year (always observed), and year fixed effects account for time-specific factors. We impute missing values of $bb_{jt} - cs_{jt}$ using the coefficients from this regression; we find that the predicted values have a correlation of 0.98 with the observed values for non-missing observations. We then calculate the gain-loss ρ_{jt} from the accounting identity for bid (equation (1)) using imputed $bb_{jt} - cs_{jt}$ and observed non-medical costs c_{jt}^{admin}.

SB.3 Characteristics of Traditional Medicare as the Outside Option

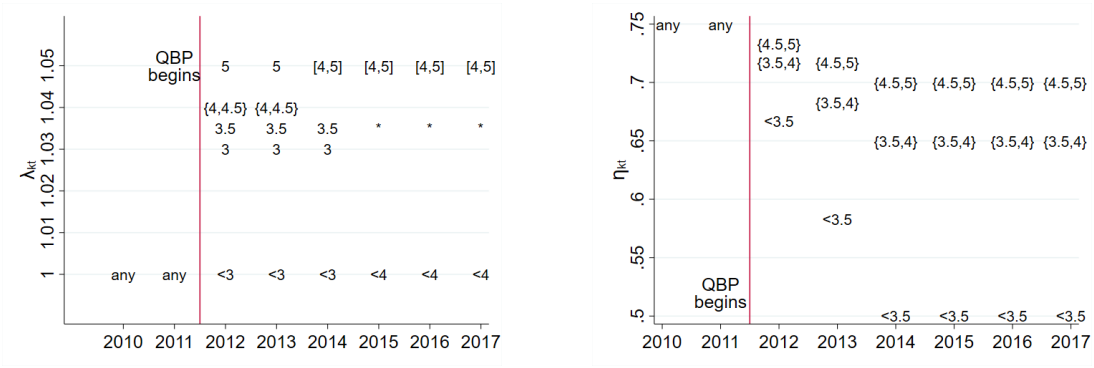
We characterize the outside option TM in a way that is comparable to MA plan benefits. To obtain cost-sharing in TM, we leverage a public dataset where we observe county-year costs in TM across all services and by service category (Centers for Medicare and Medicaid Services, 2022b). Within each service category, we apply the cost-sharing rules. For example, for Part B services, we apply the Part B deductible (since CMS reports nearly all beneficiaries use some Part B services) and infer cost-sharing of 20% for amounts above the deductible. Computing the cost-sharing for Part A services (largely, hospitalizations and skilled nursing facilities) is more challenging. We compute the cost-sharing paid under the inpatient deductible using the dataset's record of the number of individuals per county-year with a covered stay. We assume 2% cost-sharing on Part A services beyond the deductible. The 2% is calibrated to match CMS estimates that cost-sharing represents about 14% of total Medicare costs nationwide (Centers for Medicare and Medicaid Services, 2026).

SB.4 Mapping from Star Rating to Bonus and Rebate Rates

Mappings by Years. While Figure 1 reported the 2017 bonus and rebate rate mappings, the mapping varied slightly over the years 2012-2017. Figure SB.1 presents the bonus and

rebate rates by star rating for each year from 2010 to 2017.

Figure SB.1: Bonus Rates and Rebate Rates from 2010 to 2017



(a) Bonus Rates (b) Rebate Rates

Notes: Panels (a) and (b) represent (standard) bonus rates λ_{kt} and rebate rates η_{kt} over the years 2010–2017. The y-axis represents the bonus or rebate rate in that year, while the numbers denote the star rating that corresponds to that bonus or rebate rate in that year.

Double Bonus and Benchmark Caps In addition to the core QBP policy that maps star ratings to bonus and rebate rates as shown above, there are two county-specific policies: "double bonuses" and benchmark caps.

"Double bonuses" refers to a policy that doubles quality-related benchmark adjustment for around 200 counties designated in the Affordable Care Act (ACA). For example, for a plan with a star rating of at least 4, the standard quality bonus raises the county benchmark by 5 percentage points (see Figure SB.1). In a qualifying "double-bonus county," the benchmark increase is instead 10 percentage points. These double bonuses accrued to plans in proportion to enrollment in double bonus counties.

At the same time, starting in 2015, a statutory benchmark cap can also prevent a plan from receiving the full increase. Under the ACA, a county's final benchmark (including quality bonus adjustments) cannot exceed what the benchmark would have been under the pre-ACA rules for that county. Because historical spending and pre-ACA caps vary across counties, a high-quality plan's 4- or 5-star quality bonus may be fully realized in one county yet capped in another.[19]

These county-specific policies imply that the mapping from the star rating to the bonus rate is actually plan-specific. We ignore this detail in our main text for expositional sim-

[19]These caps were enacted in the ACA but were suspended for 2012-2014 as part of the Quality Bonus Program Demonstration that superseded the ACA subsidy parameters for those years.

plicity but fully incorporate the plan-specific mapping in our empirical analysis.[20]

SC Robustness: Lagged Quality Ratings

Since CMS needs to set subsidy parameters (i.e., the rebate rate and the bonus rate) before firm decisions and consumer choices, CMS uses lagged star ratings to determine the subsidy parameters and thus premiums. However, contemporaneous star ratings are advertised to consumers at the time of enrollment decisions.

Our baseline specification ignores this timing issue, which has implications for the supply-side estimation. This simplification does not affect our demand estimation because we use the true consumer-facing premium (computed using lagged star rating) and the contemporaneous star rating that is advertised to consumers. For the supply side, a firm's quality choice in year t affects its profit through three channels: (1) current marginal cost, (2) current effective demand via consumer tastes for quality, and (3) next year's effective demand via its effect on subsidy parameters and thus premiums. In our baseline estimation, we consider all three channels to be contemporaneous, which is a good approximation if preferences, costs, and market structure evolve slowly. In a robustness analysis, we allow the first two channels to be contemporaneous and the last channel to depend on market conditions in the following period.

Specifically, in the robustness analysis, the effective demand is:

$$\tilde{s}_{jct}(\boldsymbol{\rho}_{ct}, \boldsymbol{x}_{ct}, \boldsymbol{r}_{ct-1}, \boldsymbol{q}_{ct}) = \mathbb{E}_{\boldsymbol{\tau}_{ct}}\left[s_{jct}\left(\boldsymbol{p}_{ct}(\boldsymbol{\rho}_{ct}, \boldsymbol{x}_{ct}, \boldsymbol{r}_{ct-1}), \boldsymbol{x}_{ct}, \boldsymbol{r}_{ct}(\boldsymbol{q}_{ct} + \boldsymbol{\tau}_{ct})\right)\right], \tag{SC.1}$$

and when firm f chooses q_{kt} for $k \in \mathcal{K}_{fst}$, its objective function is the sum of the profit in

[20] The CMS spreadsheets provide county-specific mappings. We construct the plan-specific mapping based on the directly-observed county-specific mapping following the CMS procedure, and confirm that our plan-specific mapping evaluated at each plan's observed star rating is consistent with the observed star rating of the plan.

two periods:

$$\sum_{j\in\mathcal{J}_{kt}}\left[(\rho_{jt}+x_{jt}+a_{jt})-mc_{jt}(x_{jt},\underbrace{q_{k(j)t}}_{\text{channel 1}})\right]\times\sum_{c\in\mathcal{C}_{jt}}M_{ct}\tilde{s}_{jct}(\boldsymbol{p}_{ct}(\boldsymbol{\rho}_{ct},\boldsymbol{x}_{ct},\boldsymbol{r}_{ct-1}),\boldsymbol{x}_{ct},\underbrace{\boldsymbol{q}_{ct}}_{\text{channel 2}})$$
$$+\sum_{j\in\mathcal{J}_{kt+1}}\Big\{\left[(\rho_{jt+1}+x_{jt+1}+a_{jt+1})-mc_{jt+1}(x_{jt+1},q_{k(j)t+1})\right]\times$$
$$\sum_{c\in\mathcal{C}_{jt+1}}M_{ct+1}\mathbb{E}_{\boldsymbol{\tau}_{ct}}\left[\tilde{s}_{jct+1}(\boldsymbol{p}_{ct+1}(\boldsymbol{\rho}_{ct+1},\boldsymbol{x}_{ct+1},\boldsymbol{r}_{ct}(\underbrace{\boldsymbol{q}_{ct}}_{\text{channel 3}}+\boldsymbol{\tau}_{ct})),\boldsymbol{x}_{ct+1},\boldsymbol{q}_{ct+1})\right]\Big\}. \tag{SC.2}$$

In this model, the quality choice in year t, q_{kt}, affects current marginal cost mc_{jt} in channel (1), shifts current demand via consumer taste for quality in channel (2) , and in channel (3) affects next year's demand through affecting future subsidy parameters and thus future premiums. While this specification more accurately captures the fact that a firm's current quality choice affects both current and future profits, it requires the strong assumption that firms have perfect foresight about both the market environment and the endogenous choices for all firms in the following period: $\mathcal{J}_{kt+1}, \mathcal{C}_{jt+1}, \rho_{jt+1}, x_{jt+1}, a_{jt+1}$. In addition, we must drop the data from the last year from estimation.

Table SC.1 reports the supply-side estimates under this model. The estimates of the supplemental-benefit cost parameter is virtually unchanged, while the quality cost parameter is somewhat larger. As we show in Table SC.2, the resulting counterfactual conclusions remain similar. The first two columns compare quality incentives that load fully on the bonus rate (column (1)) or rebate rate (column (2)). As in Table 5, we find that the rebate rate incentive induces higher quality and lower bids and premiums, generating a larger increase in consumer and producer surplus than the bonus rate quality linkage. When considering the effect of QBP, we first regenerate our baseline contemporaneous-quality counterfactual since the effect of QBP are potentially slightly different in 2016. Then we account for the dynamic effects of quality by applying the two-period objective function SC.2. We find the same general effects of QBP: quality improvements from QBP's quality incentive, but higher bids related to the rebate rate reductions. Government spending falls slightly, although the reforms generate consumer surplus nonetheless.

Table SC.1: Estimates of Marginal Cost Parameters: Lagged Quality

	Est	SE
Supplemental benefits (γ^x)	1.088	0.002
Quality (γ^q)	2.017	0.072

Table SC.2: Counterfactual Results: Lagged Quality

	Quality Linkage		QBP Effect	
Objective Function	(1) SC.2 Bonus Rate	(2) SC.2 Rebate Rate	(3) 11 QBP – NoQBP	(4) SC.2 QBP – NoQBP
Plan-Level Average Outcomes				
(1) Star rating	0.09	0.31	0.26	0.18
(2) Supplemental benefits ($)	-12	78	-64	-68
(3) Gain loss ($)	19	-59	91	90
(4) Bid ($)	18	-66	94	93
(5) Rebates ($)	20	91	-86	-85
(6) Subsidy ($)	79	38	61	61
(7) Premium ($)	-42	-19	-34	-39
Aggregate Outcomes				
(8) MA share (%)	1.61	9.01	3.53	2.36
(9) Government spending ($b)	0.13	0.01	-4.48	-1.51
(9a) Payment to MA ($b)	5.96	35.44	8.02	5.73
(9b) Payment to TM ($b)	-5.83	-35.42	-12.50	-7.23
(10) Total CS ($b)	0.97	5.94	2.27	1.00
(11) Total PS ($b)	1.73	5.93	3.08	-1.77

Notes: $\kappa_\lambda = 0.986$ and $\kappa_\eta = 1.342$. The simulations were conducted for the year 2016.

SD Robustness: Risk Selection

Risk Adjustment and Risk Selection. CMS uses risk adjustment to scale its plan payments to the predicted health spending of enrollees. Specifically, it uses a projected risk score rs_{jt}^{proj} (in practice, the plan's lagged risk score) to determine the subsidy (i.e., the CMS payment) for the purpose of determining premiums at the time of consumer enrollment, and uses an updated risk score rs_{jt} based on the actual enrollment to determine the actual CMS payment to plan j. Table SD.1 summarizes key variables at the time of enrollment and whether they are updated for payment based on the actual enrollment, where (B_{jt}^1, b_{jt}^1) are the risk-neutral benchmark and bid and (B_{jt}, b_{jt}) represents their risk-

inclusive values. The table shows that the only post-enrollment update is to one component of the subsidy: the non-rebate part.

Table SD.1: Key Bidding and Payment Variables

variable	at enrollment	update for payment
risk score	rs_{jt}^{proj}	rs_{jt}
rebate	$rs_{jt}^{proj} \cdot \eta_{kt}(\lambda_{kt} B_{jt}^1 - b_{jt}^1)$	Not updated
subsidy	$\begin{cases} rs_{jt}^{proj} \cdot b_{jt}^1 + \text{rebate}_{jt} & \text{if } b_{jt}^1 \leq \lambda_{kt} B_{jt}^1 \\ rs_{jt}^{proj} \cdot \lambda_{kt} B_{jt}^1 & \text{if } b_{jt}^1 > \lambda_{kt} B_{jt}^1 \end{cases}$	$\begin{cases} rs_{jt} \cdot b_{jt}^1 + \text{rebate}_{jt} \\ rs_{jt} \cdot \lambda_{kt} B_{jt}^1 \end{cases}$
premium	Equation (2), where $B_{jt} = rs_{jt}^{proj} \cdot B_{jt}^1$ and $b_{jt} = rs_{jt}^{proj} \cdot b_{jt}^1$	Not updated

Sources: Scheinker et al. (2026); Centers for Medicare & Medicaid Services (n.d., 2015); Medicare Payment Advisory Commission (2021, 2024)

Since the composition of enrollees and their utilization of a plan directly affect both the updated risk scores rs_{jt} and marginal costs mc_{jt} for the plan, a firm may take into account how its decisions affect consumer enrollment and utilization, which in turn affect rs_{jt} and mc_{jt} and thus its profit. In other words, firms may engage in risk selection.

How Our Model Treats Selection. We treat the projected risk scores rs_{jt}^{proj} as exogenous to firm decisions. This exogeneity assumption is plausible because at the time of consumer enrollment, the projected risk score is in practice the plan's lagged risk score.

We also treat the updated risk score rs_{jt} in the subsidy as exogenous to firm decisions. Since the updated risk score rs_{jt} reflects actual enrollment, to the extent that firms take into account how their decisions affect rs_{jt}, our model omits such strategic considerations.

We allow marginal cost to depend on both plan quality and supplemental benefits. Therefore, this specification captures both their mechanical effects on costs and their effects on costs through selection—that is, how these plan choices affect the risk composition of enrollment and the utilization of enrollees—in a reduced-form way.

A full treatment of selection would require linking each Medicare beneficiary's demographics and risk score to her plan choice and spending. While the MBSF dataset reports demographic characteristics at the individual level, we do not observe risk scores, diagnoses, or utilization at the individual level.[21]

[21] Vatter (2025) uses data from the Medicare Current Beneficiary Survey (MCBS) to construct proxies for the predictors that CMS uses to compute risk scores. However, as explained in footnote 13, this dataset is not suitable for our analysis due to its limited county coverage.

Robustness: Demographic Projection of Risk Scores Because we treat the risk score as exogenous, we fix it in our counterfactual simulations. This means that we could misspecify CMS payments if plans' counterfactual choices generate substantial changes in the risk composition of enrollment, either across MA plans or between MA and TM.

To address this potential concern, we recompute CMS payments in our counterfactual simulations based on demographic-specific risk scores in this section as a robustness check. Specifically, we carry out this exercise in two steps.

First, we construct a risk score for each demographic group based on realized plan-level risk scores (Centers for Medicare and Medicaid Services, 2021d) and the demographic composition of each plan's enrollees. We do so by regressing rs_{jt} for plan j on its enrollment shares $\text{EnrollmentShare}_{jgt}$ where $g = 1, ..., 24$ indexes demographic groups formed as the combination of sex, race (white / non-white), age category ($<$ 65, 65–74, 75+), and SES (low / high). The estimated coefficients therefore capture the average risk score associated with each demographic group. We run this regression for 2017—the year in which all counterfactual simulations are conducted. The predicted values $\widehat{rs}_{jt}$ from this regression have a correlation of 0.73 with the actual risk scores, indicating that demographic composition explains a large portion of the variation in risk scores across plans. We repeat the exercise for TM: using the TM risk score in each county and the demographic composition of TM enrollment, we generate a second set of average risk scores specific to TM enrollees.

Second, in each counterfactual, we compute an updated risk score for each plan based on the demographic-specific risk score and the equilibrium demand for each demographic group. We then use the updated risk scores to recompute CMS payments to the plans. Next, we multiply the unit cost by the ratio of realized and predicted risk scores. Finally, we recompute plan profits and producer surplus. In addition, we recalculate government spending, reflecting the updates to government subsidies to MA plans and rescaling TM spending in each county by the ratio of realized and predicted risk scores.

Table SD.2 reports these outcome variables for the overall effects of QBP, effects of the bonus-rate linkage, and effects of the rebate-rate linkage based on the updated risk scores. For comparison, we also copy the outcomes in our baseline results in this table. The comparison shows that our results are robust.

Table SD.2: Robustness Analysis: Updating Risk Scores Based on Counterfactual Enrollment

	(1) Bonus-Rate Linkage		(2) Rebate-Rate Linkage		(3) QBP Overall Effect	
	Baseline	Updated	Baseline	Updated	Baseline	Updated
Plan-Level Average Outcomes						
(6) Subsidy ($)	-28	-41	-8	-25	42	39
Aggregate Outcomes						
(9) Government spending ($b)	-0.01	-0.04	-0.04	0.26	-2.39	-1.90
(9a) Payment to MA ($b)	4.14	3.77	29.56	29.26	6.28	5.90
(9b) Payment to TM ($b)	-4.14	-3.81	-29.60	-29.00	-8.67	-7.80
(11) Total PS ($b)	1.01	1.25	2.93	3.06	2.59	2.94

SE Robustness: Alternative Decompositions of the Overall QBP Effects

QBP involves three changes. There are six possible orders to introduce these three changes sequentially for decomposing the overall QBP effects into those attributable to each change. The baseline decomposition follows a specific order: Changes I, II, and III. Figure SE.1 depicts the counterfactual design for the other five alternative orders. Table SE.1 presents the simulation results.

Our decomposition results are robust to ordering. Specifically, the following results are common across all six alternative decompositions. First, Change I increases both consumer surplus and producer surplus, with the increase in total surplus larger than the associated increase in government spending when it increases. Second, while Change II always increases government spending, its effect on consumer surplus or producer surplus can be positive or negative. Third, Change III reduces both consumer surplus and producer surplus, with an associated saving in government spending smaller than the reduction in total surplus.

SF Robustness: Year-Specific Cost of Quality

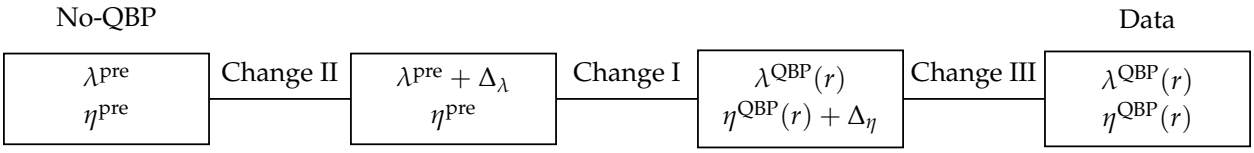
In our baseline specification, we assume that the marginal cost of quality, γ^q, is constant over time. However, the star rating is constructed from a set of underlying quality measures whose composition and relative weights have changed over time (Vatter,

Figure SE.1: Alternative Decompositions of the Overall QBP Effects

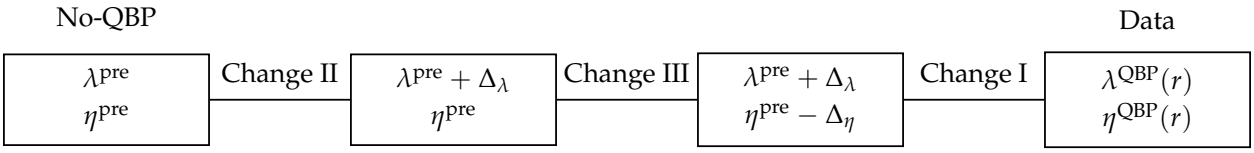
(a) Alternative 1



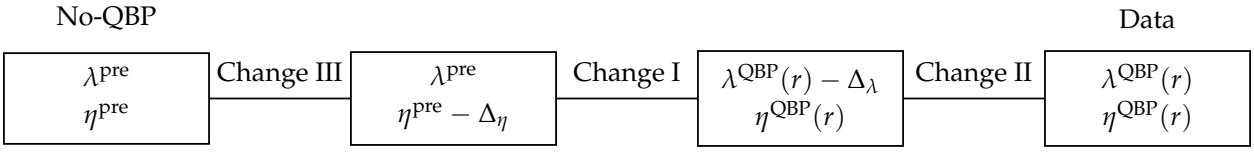
(b) Alternative 2



(c) Alternative 3



(d) Alternative 4



(e) Alternative 5

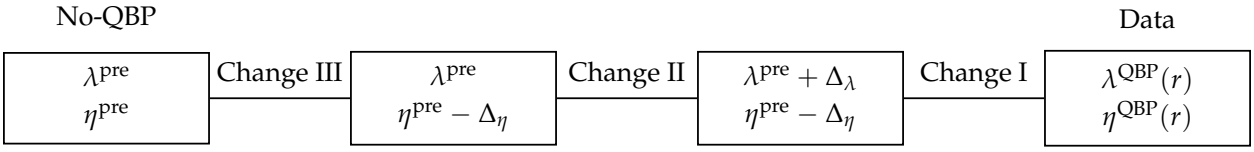


Table SE.1: Effects of QBP: Alternative Decompositions

	Change I Star Rating Dependence	Change II Higher Bonus Rate λ	Change III Lower Rebate Rate η
Alternative 1			
MA share (%)	11.40	2.68	-11.62
Government spending ($b)	6.86	0.99	-10.25
Payment to MA ($b)	53.14	11.65	-58.51
Payment to TM ($b)	-46.28	-10.65	48.26
Total CS ($b)	7.71	1.76	-7.76
Total PS ($b)	6.03	1.83	-5.27
Alternative 2			
MA share (%)	18.10	-0.84	-14.80
Government spending ($b)	11.04	2.85	-16.28
Payment to MA ($b)	84.36	-0.40	-77.69
Payment to TM ($b)	-73.33	3.25	61.41
Total CS ($b)	12.76	-0.53	-10.51
Total PS ($b)	10.40	-0.48	-7.33
Alternative 3			
MA share (%)	12.82	-0.84	-9.52
Government spending ($b)	0.10	2.85	-5.34
Payment to MA ($b)	51.77	-0.40	-45.10
Payment to TM ($b)	-51.67	3.25	39.76
Total CS ($b)	7.28	-0.53	-5.03
Total PS ($b)	7.66	-0.48	-4.59
Alternative 4			
MA share (%)	12.64	2.68	-12.86
Government spending ($b)	-0.33	0.99	-3.06
Payment to MA ($b)	50.93	11.65	-56.31
Payment to TM ($b)	-51.27	-10.65	53.25
Total CS ($b)	6.70	1.76	-6.75
Total PS ($b)	7.64	1.83	-6.88
Alternative 5			
MA share (%)	12.82	2.50	-12.86
Government spending ($b)	0.10	0.56	-3.06
Payment to MA ($b)	51.77	10.81	-56.31
Payment to TM ($b)	-51.67	-10.25	53.25
Total CS ($b)	7.28	1.19	-6.75
Total PS ($b)	7.66	1.81	-6.88

2025). As a result, a one-point increase in our quality score may represent different improvements in the underlying dimensions of care in different years, and thus may entail different costs. In addition, advances in medical technology and insurers' accumulated experience with quality improvement could also cause the cost of providing a given level of measured quality to change over time. To assess whether imposing a constant marginal cost of quality affects our estimates and counterfactual results, we also consider a more flexible specification that allows γ^q to vary by year.

Table SF.1 reports cost estimates under this alternative marginal-cost specification. Allowing a year-specific cost of quality does not change our estimated cost of supplemental benefits, which is identical to the baseline estimate of 1.087. In contrast, the estimated marginal cost of quality exhibits meaningful variation over time. It rises from 1.081 in 2010 to 1.982 in 2013, remains between roughly 1.742 and 1.802 from 2014 to 2016, and then increases to 2.281 in 2017. Hence, while the pattern is not strictly monotonic, the estimates suggest an overall upward trend in the marginal cost of measured quality over the sample period. This pattern is consistent with the possibility that changes in the composition and weights of the underlying measures used to construct the quality score altered the cost of achieving a given increase in measured quality. Importantly, the baseline estimate of $\gamma^q = 1.679$ falls within the range of the year-specific estimates, suggesting that the baseline specification captures a reasonable average cost of quality over the sample period.

Table SF.2 repeats our main counterfactual exercises using the 2017 year-specific cost estimate of \$2,281. The key takeaways remain the same. Thus, allowing the marginal cost of quality to vary flexibly over time affects certain quantitative magnitudes but does not alter the paper's central qualitative conclusions–particularly, the relative advantage of linking the rebate rate rather than the bonus rate to quality.

Table SF.1: Estimates of Cost Parameters: Year-Specific Cost of Quality

	Est	SE
Supplemental benefits (γ^x)	1.087	0.001
Quality (γ^q)		
2010	1.081	0.094
2011	1.302	0.100
2012	1.519	0.096
2013	1.982	0.097
2014	1.802	0.097
2015	1.815	0.106
2016	1.742	0.097
2017	2.281	0.100

Table SF.2: Counterfactual Results: Year-Specific Cost of Quality

	(1) Bonus-Rate Linkage	(2) Rebate-Rate Linkage	(3) QBP Effect
Plan-Level Average Outcomes			
(1) Star rating	0.09	0.25	0.18
(2) Supplemental benefits ($)	-47	6	-51
(3) Gain loss ($)	36	-52	78
(4) Bid ($)	41	-48	81
(5) Rebates ($)	-87	21	-79
(6) Subsidy ($)	-57	-18	39
(7) Premium ($)	47	-28	-12
Aggregate Outcomes			
(8) MA share (%)	-0.12	5.38	1.38
(9) Government spending ($b)	0.05	0.03	-2.23
(9a) Payment to MA ($b)	-0.87	22.09	2.26
(9b) Payment to TM ($b)	0.92	-22.06	-4.49
(10) Total CS ($b)	0.05	3.40	1.07
(11) Total PS ($b)	0.61	1.85	1.76

Notes: $\kappa_\lambda = 0.9618$ and $\kappa_\eta = 1.160$. The simulations were conducted for the year 2017.