Reallocating Bonus Payments Through Competition to Improve Medicare Advantage Plan Quality: A Dynamic Game Approach

Rajarshi Bhowal*

March 22, 2026

Abstract

This paper explores how competition among firms can be used to improve the quality of plan offerings in a managed care setting like Medicare Advantage through changes in reimbursement policy. In this managed market, private firms provide government-sponsored services at regulated prices and compete for subsidies. We study how firms offering Medicare Advantage plans compete in terms of quality and evaluate how markets would evolve under a regional comparison-based bonus payment system rewarding plans based on their relative quality in a local market. We introduce a dynamic discrete game model of firm quality investment and use the estimated model to simulate counterfactual equilibria. Our results show that the distribution of county-level average ratings shifts to the right under the counterfactual payment rule relative to their observed outcomes, with underperforming counties improving more and the average rating increasing by 0.44 stars.

JEL Codes: L15, I11, L13, M38,

*Nazarbayev University, Department of Economics, rajarshi.bhowal@nu.edu.kz; I would like to thank my dissertation committee, Jeremy Fox, Isabelle Perrigne, Xun Tang, Maura Coughlin. I also extend my gratitude to Marek Waretka, Mallesh Pai, Naoki Aizawa, and Vivian Ho for their advice and guidance. I acknowledge the helpful comments I received from the seminar participants at SEA Conference 2021, RES-SES Conference 2023, EWMES Conference 2024, and IIT Kanpur 2021. I also thank Maksat Abasov, Amir Kazempour, Hira Farooqi, and Dibya Mishra for their valuable comments and discussions. I am indebted to Ruchira Pratihar for her valuable research assistance. I gratefully acknowledge the financial support of the Social Policy Grant at Nazarbayev University. All other errors are mine.

1 Introduction

Medicare is a national health insurance program in the United States that provides government-sponsored insurance coverage to eligible beneficiaries, mostly comprised of citizens over the age of sixty-five. Traditional Medicare works on a fee-for-service (FFS) basis, where the government pays for a portion of a beneficiary's costs of any covered medical service. A potential drawback of such government-operated healthcare services is the inefficiency that stems from the lack of incentives for the organization to improve and innovate due to the absence of competition. Managed competition in healthcare systems lets private entities deliver these public health insurance services with the aim of reducing the cost of such provisions. Competition among profit-maximizing firms is believed to provide incentives to improve efficiency, which can be exploited in an adequately designed market. Such private delivery of public health insurance has been a growing phenomenon in the United States in recent times (Gruber, 2017) [21] primarily based on this rationale.

The Medicare Advantage (MA) program allows private insurers to provide traditional Medicare services to eligible beneficiaries, and the government pays subsidies to the insurers for every beneficiary they enroll in their plans. Beneficiaries can choose from a set of available plans in a local market and forgo traditional Medicare services. Private insurers compete with each other as in a standard insurance market for beneficiary enrollment, but generate revenue primarily through government reimbursements. However, designing such a managed healthcare market is challenging and comes with a lot of unanswered issues that economists and policymakers are debating. One such pressing concern is monitoring the quality of service provided by private insurers. Competition in health insurance markets may fail to improve health outcomes if consumers are not able to identify high-quality plans (Abaluck et al., 2020) [1] and thus may require interventions.

This paper explores how firms in a managed care setting, like Medicare Advantage, compete in quality and how it can be improved by generating more competition at the local market level. We introduce a dynamic discrete game model where firms choose the quality of their offered insurance plans in each period with a costly investment. Our model captures and quantifies the strategic interaction among firms while choosing plan quality. It also captures the intertemporal nature of the game, where investment in quality in one period affects the outcomes and payoffs for the next period. The primary objective of the paper is to study the supply side effects of such bonus payments and designing the payment rules that can improve access to quality healthcare.

We use the Centers for Medicare and Medicaid Services (CMS) assigned star ratings as an objective measure of insurance quality. The star ratings for the available plans are publicly available and provide the beneficiaries with quality information. CMS also pays bonus subsidies to plans with higher star ratings. We provide the details regarding their calculations in Section 2. Thus, under the condition that the consumers value quality, a firm's payoff is dependent on both its own plan quality and also the quality of the competing plans through the market share. However, a key feature of the current quality bonus payment system is that benchmark bonuses are triggered by a plan's absolute star rating relative to a fixed national threshold, rather than by its performance relative to competing plans within the same market. There exists some empirical evidence casting doubt on the ability of the current financial incentives to improve quality (Layton and Ryan, 2015)[26].

This paper aims to predict if the quality of Medicare Advantage plans can be improved by fostering quality competition among firms through changes in the reimbursement policy. We do this by introducing a new quality bonus payment system where firms are paid based on their relative performance compared to their peers in a local market area. In other words, firms earn quality bonuses if their plan quality is better than other plans in the same local market. We also consider bonus payments as transfers from low-performing to high-performing plans to maintain budget neutrality. We predict if this payment policy change can generate more competition among firms and evaluate its effects on quality outcomes.

Two critical aspects of program design are often debated in the context of pay-for-performance programs that our counterfactual bonus payment rule addresses. On the one hand, it moves from an absolute target-based payment system to a relative performance-based payment system. On the other hand, this transfer-based payment system uses a mixture of penalties and rewards instead of injecting extra bonus payments for satisfactory performance. We find empirical evidence in the context of Medicare Advantage whether these changes can lead to improved outcomes and thus address some key issues regarding the design of pay-for-performance programs. To our knowledge, there are no theoretical results that can compare the effectiveness of these two methods for implementing a performance bonus payment. This paper provides empirical results and insights regarding comparing the two forms of pay-for-performance in the context of Medicare Advantage Quality Bonus Payments, which provides us with the appropriate setting. However, these results can also be potentially helpful in other settings where similar performance bonus payments are implemented

It is essential to understand the nature of the strategic interaction among firms in terms of quality to assess the effects of this proposed payment policy change. We introduce a framework where beneficiaries choose plans based on their characteristics, plan-specific heterogeneity, and quality ratings. Multiproduct firms choose prices and whether to invest in quality improvement in each period. Profit-maximizing firms consider their competitors' strategies when investing in quality. We estimate the cost of such quality improvement initiatives in the context of Medicare Advantage using our dynamic discrete game model. The study of quality investment choice in a managed healthcare market under strategic interactions is an essential contribution of this paper.[1]

Implementing a dynamic game model in managed healthcare markets is challenging. The payoff relevant state space for the firms is large as they take into consideration characteristics of their own, their competitors, and also market characteristics. Some of these variables are continuous, providing computational challenges while solving the value function. Moreover, their relationship with the payoffs is not straightforward. We deal with this problem of dense state space by assuming a linear parametric approximation of the value function like Sweeting (2013) [32], which reduces the computational burden of our estimation procedure. Given the estimates of our structural parameters, we simulate the Medicare Advantage markets forward by calculating the equilibrium policy function of the firms under the new bonus payment rule and predict how each market would evolve in terms of offered plan quality.

[1]Gowrisankaran and Town (1997) [20] studies quality investment choice in hospitals using a dynamic equilibrium framework.

Our counterfactual analysis has important policy implications. The current bonus payment system in Medicare Advantage rewards plans if their quality ratings are above a national threshold. This can generate regional differences in government spending, as shown in Figure 1 with three example markets. Under this payment system, government expenditure on quality bonuses is skewed towards Market 3, where all the plans have a quality rating above the threshold, whereas Market 1 receives none. In other words, government expenditure on quality bonuses is skewed towards markets where more plans exceed the threshold, neglecting regional differences and varying incentives for improvement. Also, the incentives to improve may vary across markets. Plans closer to the national threshold have more incentives to improve as smaller investments take them to bonus status. Thus, markets with more marginal firms may see more quality investments and, consequently, more reimbursements from the government.

We implement a counterfactual bonus payment rule based on the reform proposal discussed in the 2019 report to Congress by the Medicare Payment Advisory Commission (MedPAC). The report highlights distortions created by the current Medicare Advantage quality bonus structure and proposes replacing it with the Medicare Advantage Value Incentive Program (MA-VIP). The MA-VIP framework is designed to improve the alignment between bonus payments and plan performance by rewarding contracts based on their relative quality within a market rather than relying on sharp star-rating thresholds. As illustrated in Figure 2, under the counterfactual payment rule, some plans in every market get access to quality bonus payments. Also, this payment system distorts the firms' incentives to invest in quality. In contrast to the threshold-based payment rule, a small improvement in quality for low-performing plans can earn bonus payments if it makes them better than other plans in the market. This can potentially generate more quality competition in every market regardless of their initial conditions as the firms will try to outperform each other in terms of quality.

In the context of our empirical analysis of Medicare Advantage, this is a relevant problem. The threshold-based quality bonus payment system rewards plans that have a rating greater than or equal to four stars. Thus, marginal plans can receive quality bonuses by an incremental increase in their quality ratings. In contrast, the same level of improvement comes with no additional payment for plans with a lower star rating. In this paper, we will provide evidence from the data that these marginal contracts show more quality improvement initiatives to reach a four-star rating compared to lower-performing contracts, to motivate our question. Also, there exists regional disparity in terms of average plan quality across counties, as shown in Figure 6. We explore if the quality of care in the underperforming Medicare Advantage markets can be improved under the new payment policy we study in the counterfactual.

We show that our model performs well in predicting the observed transition of the market outcomes in terms of the quality rating of the offered plans. The counterfactual results predict that under the new bonus payment system, 75% of the studied counties perform better when compared to the observed data under the status quo payment policy and that more counties will have a quality rating greater than four stars. This improvement is mostly seen in historically underperforming markets that are observed to have lower average ratings in the data. The results indicate that the quality of care can be improved by introducing payment rules that induce more competition among firms. We also observe that most markets having an average quality rating between 4 and 4.25 stars,

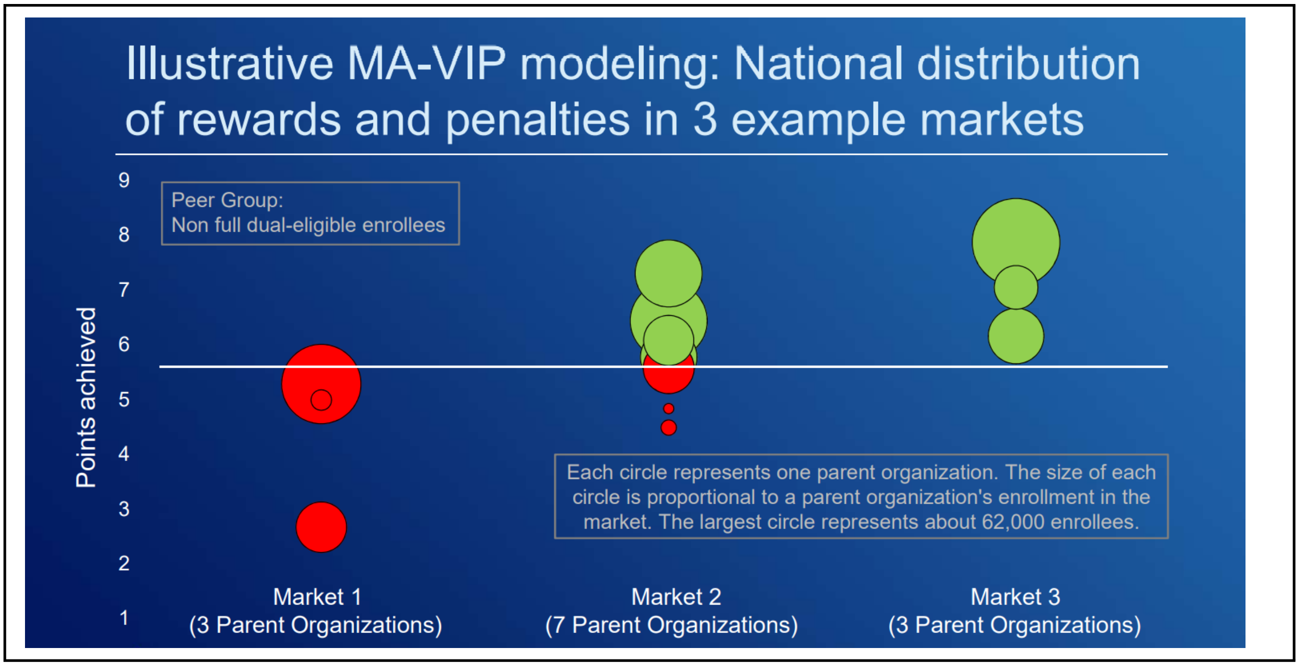

Figure 1: Bonus distribution in example markets for threshold-based bonus payment system. Source: MedPAC

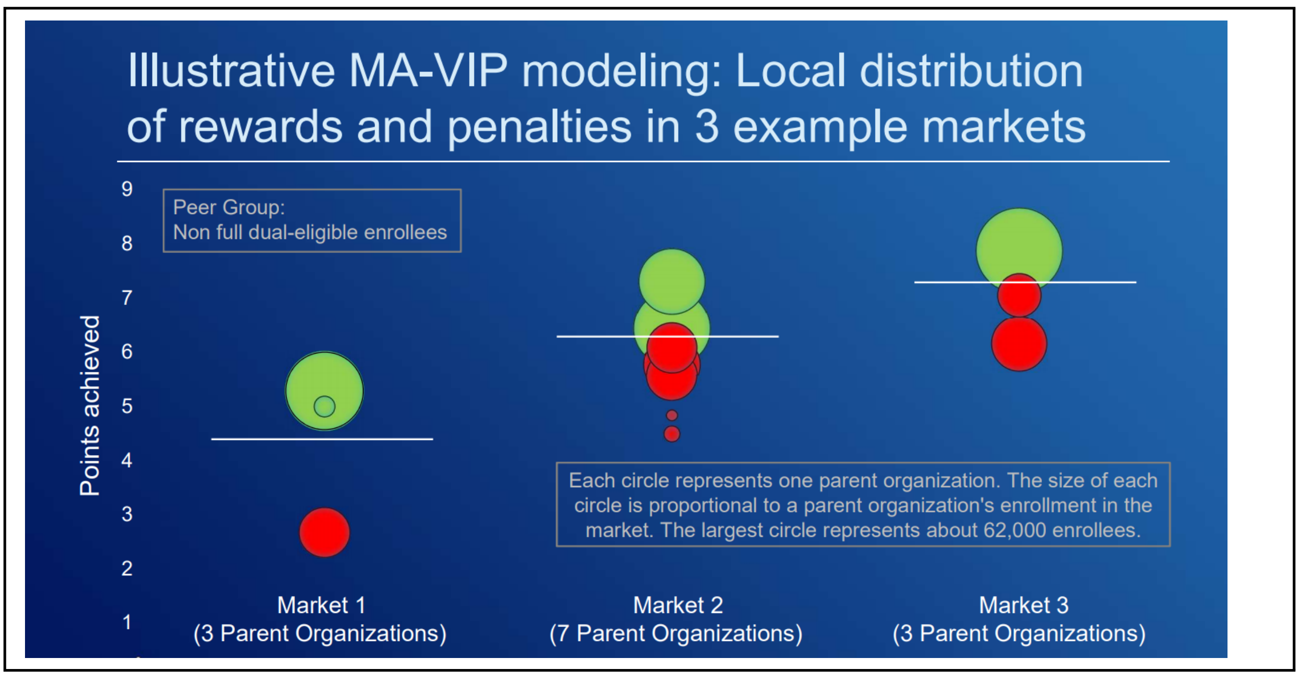

Figure 2: Bonus distribution in example markets for market-level comparison based bonus payment system. Source: MedPAC

reducing the heterogeneity in access to high quality across regions. Among the counties that are observed to improve in the counterfactual, the mean increase in their average increase in the star is 0.44. Though our estimates are specific to the Medicare Advantage markets, these results provide key insights regarding the proper use of financial incentives to improve the quality of care that may apply to other similar markets.

Related Literature. We contribute to the existing literature regarding managed healthcare markets and Medicare Advantage. Gruber (2017) [21] studies the changing nature of public health insurance in the United States and the extensive use of the "managed competition" approach employed by the government to deliver health insurance. We share similar motivations to explore how these markets work, though they differ methodologically. Previous papers have studied the nature of competition and issues of subsidy design in these markets. Miller et al. (2019) address the issue that the market benchmarks based on which CMS pays subsidies to individual plans are sub-optimally set and find the optimal benchmarks to maximize consumer surplus. Curto et al. (2019) study how market competition works in the managed competition setting of Medicare Advantage and the surplus generated by the competitive bidding process. Decarolis, Plyakova, and Ryan (2020) look at the interaction between strategic insurers and the subsidy mechanism in Medicare Prescription Drug plans. Though the aforementioned papers are similar to our work in exploring issues related to competition and subsidy design in a managed care setting, they study the implications of these factors on consumer surplus without explicitly considering the quality of care. Our paper focuses more on the problems regarding quality offerings in a similar market and explores the scope of improving the same through competition generated by subsidy incentives.

Pay for performance has been increasingly used to provide incentives to health insurance healthcare providers to improve performance.[2]. It is widely applied in both the United States and the United Kingdom and is now being experimented with and implemented in many other countries. Many previous studies have looked into the question of how to structure incentive programs in terms of rewards versus penalties. The rationale for the introduction of penalties for under-performance is that agents may be more responsive to penalties rather than rewards (see Kahneman and Tversky (1979) [25]). However, Kahneman, Knetsch and Thaler (1986) [24] show how penalties might be deemed as unfair by the agents and may fail to elicit appropriate behavioral responses. Damberg et al. (2007) [12] conducts an environmental scan of hospitals under the pay-for-performance programs in the United States.

This paper compares an absolute measure-based payment system with a relative measure-based one. Though absolute targets are transparent and straightforward without any uncertainty involved, they may be inefficient because the bonuses may be paid to insurers or operators already above the target threshold and might fail to incentivize the ones below. A comprehensive study regarding whether financial incentives can help improve healthcare quality has been done by Conrad and Perry (2009) [10]. An absolute measure-based payment system might also generate regional disparity in the incentives. Doran et al. (2008) [14] empirically study whether incentive schemes can increase inequalities in the delivery of care if practices in affluent areas are more able to respond to the incentives than those in deprived areas in the United Kingdom. They study the quality and outcomes framework in the UK, which is a financial incentive scheme to incentivize general practices through bonus payments based on a set of quality indicators. Their results suggest that financial incentive schemes have the potential to make a substantial contribution to the reduction of inequalities in the delivery of clinical care related to area deprivation.

[2] See Eijkenaar (2013) [15] for a detailed review

Previous papers have studied market structure, competition, and quality in other different market structures. Hoxby (2000) [23] studies how competition among public schools affects education quality. Fan (2013) studies the quality of content in newspaper markets and how the quality offerings are affected by mergers under a static framework. We introduce a dynamic model for understanding the relationship between competition and product quality which captures the time lag between investment and outcome.

Our paper contributes to the existing literature regarding quality ratings in Medicare Advantage. Reid et al (2013) [29] study how star ratings affect plan choice in Medicare Advantage by the beneficiaries. Similar work regarding the information and publicly available performance reports and their effect on choice in healthcare and insurance markets have been done by Farley et al. (2002) [18], Scanlon et al. (2002) [30], Wedig et al. (2002) [35], Dafny and Dranove (2008) [11], Darden and McCarthy (2015) [13], and Handel and Kolstad (2015) [22]. Fioretti and Wang (2021) [19] show how the current quality bonus payment system can widen the inequality in accessing social services by private insurers selecting healthier enrollees.However, we look at the supply-side problems in the Medicare Advantage markets associated with plan quality offerings.

Some previous papers explore the supply-side implications of the star ratings and quality bonuses in Medicare Advantage. Vatter (2022) [34] looks at the problem of construction of the star ratings and proposes a welfare-increasing optimal design of the scores, focusing on how the aggregation of multiple quality measures into a single score affects consumer choice and plan incentives. In contrast, we take the star rating system as given and study how the payment rules tied to these ratings influence insurers' quality investment decisions, with particular attention to the distortions created by threshold-based bonus payments. Adrion (2019) [2] finds that contracts operating in more concentrated MA markets receive higher star ratings and relates how market competition affects star ratings through negotiations between private insurers and providers while forming provider networks. Sen et al. (2021) [31] explore in more detail how narrow networks in Medicare Advantage translate into higher star ratings. We differ in our methodology from these papers and provide a structural framework for analyzing the effect of competition on quality offerings by firms. The introduction of a dynamic game to analyze this problem is a novel contribution of our paper.

This paper is methodologically indebted to the extensive existing literature on the estimation of dynamic games. We follow Ericson and Pakes (1995) [16] to assume that firms use stationary Markov Perfect Nash Equilibrium in strategies. Our model and solution procedure for the value function by linear parametric functional approximation is based on Sweeting (2013) [32], and Benitez-Silva et al. (2000) [8]. Our estimation procedure closely follows Pakes, Ostrovsky, and Berry (2007) [28] We implement a nested pseudo-likelihood approach for estimating the dynamic game based on Aguirregabiria and Mira (2007, 2010) [6][7]. Our model of the dynamic game also borrows from Aguirregabiria and Ho (2012) [5], Aguirregabiria (2012) [3], and Blevins (2014) [9]. A complete survey of estimating dynamic games can be found in Aguirregabiria et al. (2021) [4].

Our work is also related to policy discussions on redesigning the Medicare Advantage quality bonus program. For example, the Medicare Payment Advisory Commission has proposed the

MA-VIP system, which evaluates plans relative to competitors within local markets. Unlike that proposal, our paper models firms' dynamic quality investment decisions and evaluates equilibrium responses to such incentives.

We organize the rest of the paper as follows. We discuss details regarding the Medicare Advantage and star ratings in Section 2. In Section 3 we give the details regarding the CMS data we use for our analysis. We introduce our model in Section 4 and proceed to estimation of the parameters in Section 5. In Section 6, we describe how we simulate the counterfactual market outcomes and discuss the results.

2 Empirical Setting and Data

In this section, we provide a detailed discussion regarding Medicare Advantage and our empirical setting.

2.1 Medicare Advantage

The Medicare program is a public health insurance program in the United States whose primary beneficiaries are citizens over the age of sixty-five. The goal of the Medicare program is to provide affordable healthcare services to retired beneficiaries, and it was started in 1965. Though initially, it provided a basic insurance program to cover beneficiaries without any health insurance, over time, it changed to provide more citizens access to quality and affordable healthcare. Medicare consists of four parts: A, B, C, and D. Part A covers hospitals and skilled nursing services, Part B covers outpatient services, and Part D mainly covers self-administered prescription drugs. Part C is an alternative called Managed Medicare or Medicare Advantage, which allows beneficiaries to access at least the services of Part A and Part B through a set of private health insurance plans from which they can choose. Thus, it introduced a choice-based system in a public health insurance program. In this paper, we will concern ourselves with only Part C, the Medicare Advantage. While Part A coverage is provided for free for most people and funded by the Medicare taxes paid by the beneficiaries while working, a monthly premium is charged for Part B

In response to the increasing costs of Medicare, in 1982, Congress authorized Medicare administrators to engage in a series of trials in which the government handed over management of the medical care of selected groups of Medicare enrollees to private insurers in exchange for a fixed payment that did not vary with the realized medical expenditures of each individual. Though it was not a success initially, it laid the foundation for introducing managed competition in Medicare. It went through a series of changes and modernization that led to the formation of the Medicare Advantage (MA) program. Currently, the Medicare Advantage allows each beneficiary to choose from a set of eligible plans that private insurers provide and forgo the traditional Medicare services that are provided directly by the government. In other words, a Medicare-eligible beneficiary can opt for either traditional Medicare or one of the available Medicare-approved private plans in their local area.

Original Medicare provides coverage for the majority of medically necessary services and supplies delivered in hospitals, doctors' offices, and other healthcare facilities. MA plans must cover all these necessary services that traditional Medicare covers. However, they can offer extra benefits that original Medicare does not cover, such as eye exams, most dental care, and routine exams. Any beneficiary choosing a Medicare Advantage plan must also pay the Part B premium they would pay if they chose traditional Medicare. However, some plans may charge an extra premium over the Part B premium for their additional services. When a beneficiary enrolls in a Medicare Advantage Plan, Medicare pays a fixed amount of money for the coverage to the company offering the Medicare Advantage Plan.

Medicare Advantage has been a success, with more and more people choosing an MA plan over traditional Medicare. Enrollment in MA has increased more than double from 12 million or 26% of the total Medicare eligible population in 2011 to 26 million or 42% of the Medicare population in 2021. During this period, enrollment grew by around 10% every year. At $343 billion per year, it comprises around 46% of total Federal Medicare spending and is one of the fastest growing healthcare sectors in the United States.

Enrollment in MA plans takes place during the Open Enrollment Period from October 15 to December 7. Beneficiaries may choose a new plan or switch to a different plan during this period. Each MA beneficiary has to pay a premium called the Part B premium and may need to pay an additional private plan premium as well. This additional premium is determined by a process where each MA plan must report its operation cost to CMS for providing services in a particular county before the enrollment period. Each county is assigned a benchmark that represents a weighted average of fee-for-service (FFS) spending in that county or the expenditure incurred by Medicare. Plans bidding below the benchmark charge a 0 extra premium, while plans that bid above the benchmark charge the difference between the bid and the benchmark as an extra premium.

Medicare Advantage allows beneficiaries to choose from differentiated insurance plans, which is not possible with traditional Medicare. Where traditional Medicare provides uniform benefits to all enrollees, private insurers compete in terms of prices, plan quality, and supplementary benefits and design plans to cater to the needs of the different target populations, thereby increasing the choice set of beneficiaries. Though this might be considered a desirable feature of the market, it is also necessary to make sure that well-informed choices can be made. Thus, in this kind of market, publicly available indicators of plan quality are of utmost importance. To serve this purpose, the CMS uses a 5-star rating system to measure plan quality.

2.2 Star Rating

Though cost-cutting for the government was the main driving factor for the introduction of Medicare Advantage, which allowed private insurers to provide public health insurance services, the need to monitor and improve the quality of care was felt very soon. Since 2009, the CMS has provided comprehensive data regarding MA plan performance through its Star rating program with the goal of encouraging consumers to choose high-quality plans and incentivizing health insurers to improve their service quality. The ratings are assigned at the MA contract level, and all plans under the same contract have the same rating. These ratings serve as the primary publicly reported summary measure of quality for Medicare Advantage plans and play a central role in influencing both consumer choice and plan payments.

The star rating is a summary measure of plan quality constructed by the Centers for Medicare and Medicaid Services (CMS) using a range of observable performance indicators. These indicators fall into several broad categories, including outcome and intermediate outcome measures (such as improving physical health and managing chronic conditions), process measures (such as cancer screenings and flu vaccinations), patient experience measures (such as customer service and access to appointments), and access or administrative performance measures (such as the timeliness of appeals decisions). The underlying measures are compiled from multiple sources, including CMS administrative records, surveys of enrollees, data collected by CMS contractors, and information reported by health and drug plans. CMS aggregates these measures using a weighted scoring system to generate an overall contract-level quality score. Medicare Advantage Star Ratings reported for a year reflect plan performance during a prior measurement period, typically involving data from the prior year and, in some cases, earlier years or multi-year averages. Star ratings are then assigned based on predetermined performance thresholds.

The star ratings not only provide consumers with information about plan quality but also interest private insurers offering Medicare Advantage plans as they represent an opportunity to generate additional revenue. One obvious reason is that if consumers take into account observed ratings while making a plan choice, plans with higher ratings compared to their competitors might capture a larger portion of the market share. Also, a higher star rating in a particular market is usually associated with a higher benchmark and a higher rebate percentage the plan faces, which increases the plan payment for a given bid, also referred to as Quality Bonus Payments (QBP). This process will be explained in detail in the next subsection on plan payment. The bonus payment system was introduced following the Affordable Care Act, which mandates that plan payments should be dependent on quality. The ratings used for payments from 2013 onward were calculated under this revised framework.

Plan quality has improved in MA over the years, especially after the introduction of QBP, as shown in Figure 3. This pattern suggests that private insurers may respond to these bonuses by improving the quality of the plans they offer. However, the increase in average star ratings observed after 2012 may partly reflect administrative changes to the Medicare Advantage rating system introduced under the Affordable Care Act. Beginning in 2012, the Centers for Medicare Medicaid Services revised the construction of the star ratings, expanded the set of quality measures, and linked ratings more directly to quality bonus payments. The changes in the initial years may have affected the distribution of ratings independently of underlying changes in plan quality.

2.3 Plan Payments

Every MA market is characterized by a CMS-assigned benchmark B_{mt} based on which all subsidies are paid. The subscript m denotes a particular market and t a particular period. Each plan reports its cost of providing service b_{jt}. For each individual beneficiary i enrolled in plan j, CMS reimburses

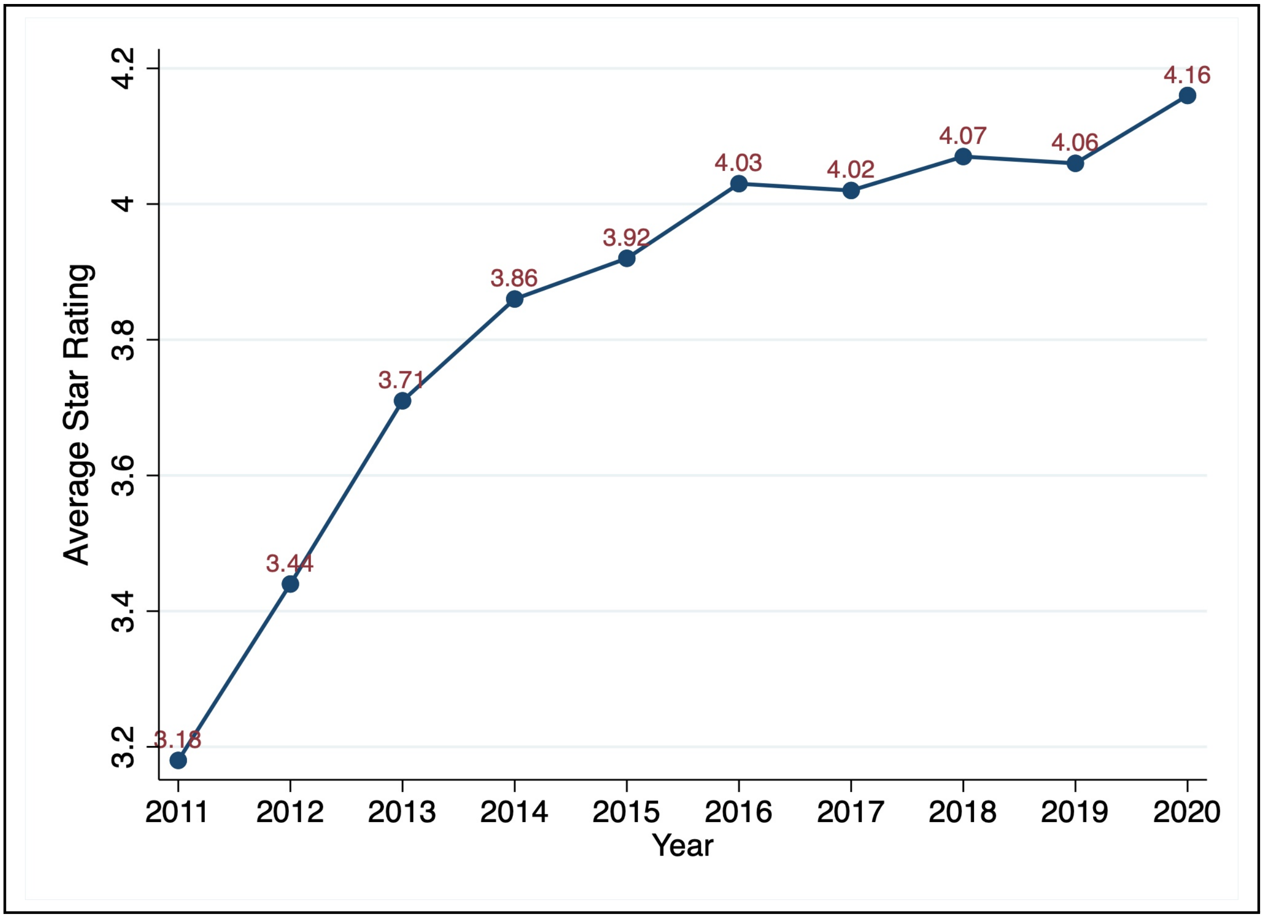

Figure 3: Enrollment weighted average star rating over the years. This figure shows the increasing enrollment trend weighted average star rating after quality bonus payments are introduced.

the plan Reb_{ijt} using the following rule:

$$\text{Reb}_{ijt} = \begin{cases} B_{mt} \times \text{R}_{it} & \text{if } b_{jt} \geq B_{mt} \\ \left(b_{jt} + \lambda^B_{jt} \times \left(B_{mt} - b_{jt} \right) \right) \times \text{R}_{it} & \text{if } b_{jt} < B_{mt} \end{cases} \tag{1}$$

where λ^B_{jt} is the rebate percentage or the portion of the surplus the plans get to keep, and R_{it} is the risk score assigned by the CMS to each individual beneficiary measuring how likely the beneficiary is to incur medical expenses relative to the county benchmark. An average beneficiary in the market receives a risk score of 1. These risk scores are calculated based on individual characteristics and prior medical history. Plans reporting a cost b_{jt} more than the benchmark B_{mt} receive only the benchmark amount, and the difference is charged as an extra premium from the beneficiaries. If the reported cost is less than the benchmark, then the plan receives its reported amount and part of the difference between the benchmark and the reported cost as a reward for cost saving. This rebate percentage is defined by λ^B_{jt}.

A plan's star rating affects its payments through both the benchmark level and the rebate percentage in every market. Under the current payment rules, contracts with a star rating of 4 stars or higher receive a quality bonus that increases the benchmark by 5 percent (10 percent in qualifying counties). In addition, rebate percentages depend on star ratings: plans with 3 stars or fewer receive a 50 percent rebate, while the rebate increases to 65 percent for plans with 3.5 to 4 stars and 70 percent for plans with 4.5 stars or higher. As a result, holding all else constant, plans with higher star ratings receive more generous payments. Importantly, the benchmark increase at the 4-star threshold creates a discontinuous jump in payments, which may distort plan incentives around this cutoff. Moreover, these bonuses are based on the rating relative to a national threshold rather than on performance relative to competitors within a market. In Section 6, we analyze the effect of a proposed reform to the payment rule in which quality bonus payments would shift from this absolute threshold-based system to a relative performance system, where plans would receive bonuses based on their quality ratings compared to competing plans within the same market, regardless of the absolute star rating level.

Star ratings affect plans through two main channels. Current star ratings influence consumer choice, as beneficiaries observe these ratings when selecting among Medicare Advantage plans. In contrast, plan payments depend on lagged ratings because benchmark adjustments and rebate percentages are determined using ratings from the previous period. To capture these incentives in a tractable way, our model includes both current star ratings, which affect consumer choice, and lagged star ratings, which determine payments in approximating the value function.

2.4 Quality Improvement Initiatives

In an effort to evaluate if the quality bonus payments in practice translated to the increasing trend in the plan quality that we observe in the data and to understand how the QBP may affect organizations' operations, CMS collected information regarding the quality improvement (QI) activities of MA plan sponsors through a contract-level survey and case studies with selected MA sponsors in 2016 [3](Refer to Appendix C for more details). The majority of survey respondents (88 percent) indicated that the budget for the contract's QI activities increased between 2010 and 2013. We enumerate some of the most important characteristics and avenues of the MA contract's QI initiatives as stated in CMS reports in order to motivate our strategic model of firm choice.

i) MA organization's ratings and their competitors' ratings drive organizational star rating strategies: This point illustrates how competition can mitigate or enhance the effects of QBP incentives and the Star Ratings program more broadly to attain higher star ratings. Respondents noted that in a market where all the competitors are 3-star plans, having a 4.5-star rating is good enough. But in markets where there are high-quality contracts on the cusp of very high ratings, they feel acute pressure to achieve five stars. ii) Provider Network formation and provider incentives are an essential pathway of improving ratings: Most experts in this industry agree that provider network formation is an important aspect of improving star ratings. All MA insurers interviewed focused on star ratings when forming narrow networks, though different insurers took different approaches.

This discussion motivates our model, where firms' action choices represent these QI activities in

[3] https://innovation.cms.gov/files/reports/maqbpdemonstration-finalevalrpt.pdf

our empirical setting. In each period, the firm decides whether to invest in these activities. Quality improvement initiatives are costly and improve the quality rating, and earn bonus payments in the future period. We use a dynamic model to capture this intertemporal nature of choice and payoff. Specifically, we use the dynamic game to capture the strategic interaction between firms, as discussed in Section 3. Pay for Performance is increasingly being used to stimulate the quality of services in public healthcare systems. It is widely being used in the United States and the United Kingdom and is gradually but steadily being implemented or experimented with in many other countries. However, evidence of their effectiveness is inconclusive, and their proper implementation is a debated topic in the literature. Proper program design is thus of utmost importance to make such pay-for-performance schemes work efficiently. Some important aspects of such program design are properly defining the performance measures and goals and finding out the best way to implement them, as flaws in doing so will translate into limited success.

One debated aspect of such program design is whether to use rewards or penalties to incentivize the agents to improve their Performance. Another question is whether to use absolute or relative targets to measure the Performance on which their payments will depend. This paper tries to answer both these questions in the context of Medicare Advantage and the effectiveness of the delivery of quality bonus payments to improve plan quality. In particular, we predict the effect of moving from a threshold-based quality bonus payment system to one where bonus payments are made based on the relative Performance of the plans compared to those of their competitors. It should be noted that the status quo bonus payment rule, as discussed in this section, is an absolute target-based one. Also, the counterfactual bonus payment rule, whose effect this paper predicts is a mixture of rewards and penalties, where relatively underperforming plans in a local area are penalized in the form of reduced payment by Medicare, and this reduction is redistributed to the plans that perform better in comparison to them. We would like to see whether this redistribution increases quality competition among the insurers.

Our proposed redistribution mechanism is conceptually related to the Medicare Advantage Value Incentive Program (MA-VIP) proposed by MedPAC, which evaluates plans relative to competitors within local markets and redistributes payments accordingly. While the MedPAC proposal focuses primarily on the fiscal implications and allocation of bonus payments, it does not model how firms adjust their quality investments in response to such incentives. Our dynamic model fills this gap by explicitly incorporating firms' quality investment decisions and the cost of improving ratings. The simulation results suggest that a locally competitive incentive system could substantially increase quality investment, particularly in historically low-performing markets.

2.5 Data

We use publicly available data from various CMS sources for our analysis. We use aggregate MA enrollment for every plan in every market. The data reports monthly enrollment, which we convert to annual enrollment by taking an average. The Plan Benefit Package data provides plan-level information regarding the premium and plan characteristics as provided by the plans during the annual bidding process. The plan payment data provides information regarding the average per month per member Part B payment, the rebate payment to each plan, and the average risk score of the beneficiaries enrolled. The landscape folder provides data regarding service contract area and the performance data provides detailed information regarding the star rating and individual measures.

We have 1,48,392 year-county-plan level observations from 2013 to 2016 comprising HMO and PPO plans. The plans differ in the subsidiary services they offer. Star ratings are calculated only for contracts with sufficient enrollment; for example, new Medicare Advantage contracts do not receive star ratings until they have accumulated at least about 600 enrollees (MedPAC, 2020). Thus, we can't include counties with low MA enrollment due to a lack of star rating data. Among the observed plans, 38% of the plans provide vision coverage, 62% of the plans provide hearing coverage, and 50% of the plans have dental coverage included. Table 1 shows the summary statistics of the MA markets that we use in our analysis. We select these years as most changes in star ratings are observed to occur during this period after the introduction of the quality bonus payments. The average rating of available plans is observed to increase from 2013 to 2016. The number of observations where the rating improves in the next period remains fairly stable across the years. We observe the maximum improvement in the quality ratings from 2013 to 2014. Average MA enrollment went up every year during this period. The average contract revenue in each market, which is the sum of the total premium earnings and government subsidies, shows an upward trend over the years. The increase in average revenue is due to both the expansion of the MA program and also more bonus earnings by star rating improvement.

Table 1: **Summary Statistics of Medicare Advantage Markets by Year**

Variable	2013	2014	2015	2016	Total
Average Star Rating	3.504	3.676	3.810	3.930	3.736
(S.D.)	(0.565)	(0.568)	(0.541)	(0.539)	(0.575)
Number of Observed Rating Improvement	5,120 (30.39%)	4,291 (23.83%)	5,617 (29.24%)		15,028 (20.71%)
Contract Average Enrollment per County (rounded)	718	764	765	842	774
Contract Average Revenue per County	7,232,055	7,707,938	8,780,301	7,252,138	7,688,444
Number of contract-market observations	16,850	18,008	19,208	18,503	72,569

Notes: The reported observations are at the year-market-contract level. The standard deviations for star ratings and percentage of observations observed to improve their ratings are provided in parentheses

In the data, we observe that 15,028 year-market-contract observations (20.71%) improved their star ratings in the data. Among all these observed increases in the star rating over the years, 6,267 (41.7%) of these contracts have a current rating of three and a half. These are marginal plans in our case, as improving the rating by a half-star makes them eligible for quality bonus payments. Figure 4 shows that most of the increase in star ratings is observed for contracts that are rated three and a half stars, followed by four stars (28.85% of the total observed increases).

We explore patterns in the data to reflect on which plans are more likely to improve their ratings in the next period. Table 2 reports the probit estimates of the average marginal effect of being a three-and-a-half star-rated marginal contract, given by the variable Marginal, on the probability of increasing the star rating in the next year to be 0.037 and statistically significant. This might be in-

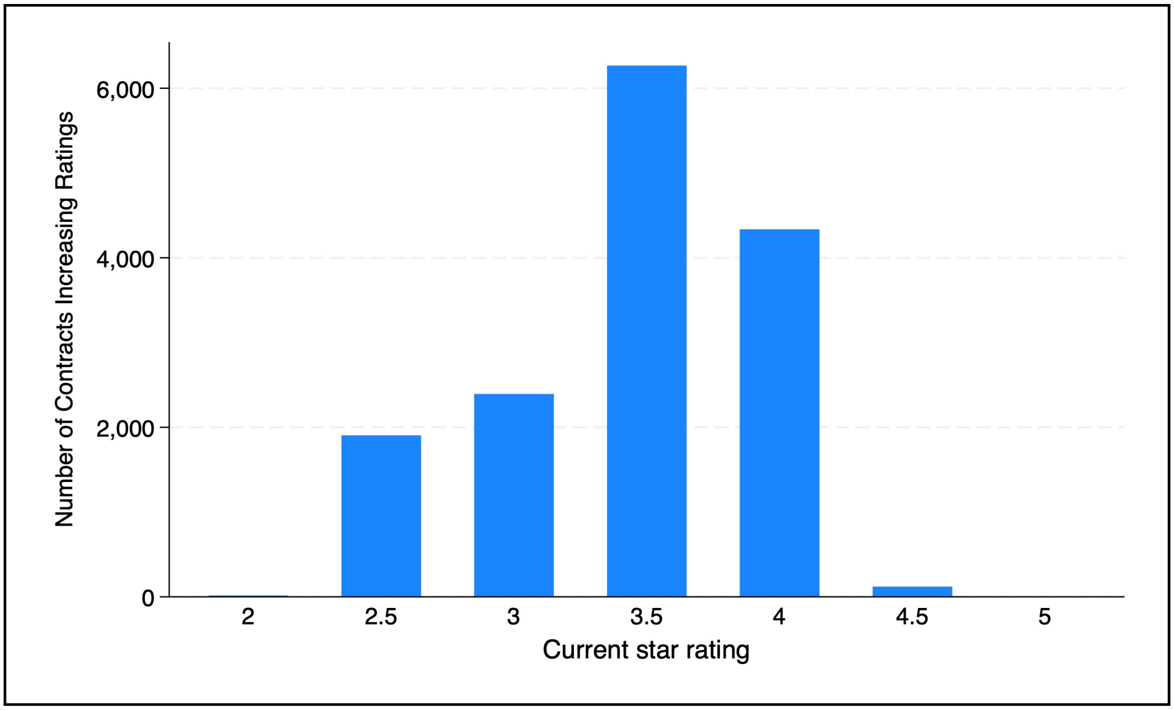

Figure 4: Number of county-contract observations in the data observed to increase their rating in the next period by their current star ratings. Most of the rating increase comes from contracts with star ratings of 3.5 and 4 stars.

dicative of the fact that contracts are trying to cross or be at the national threshold of four stars, and the current bonus payment system can potentially generate a "cliff effect". It also shows a negative effect for the same for contracts equal to or below two stars. The variable Diff_meanstar is the difference between the average star rating of all plans in the market (county) and a contract's current rating. Thus, the positive and statistically significant marginal effect of Diff_meanstar indicates strategic quality responses within markets. Contracts that are relatively disadvantaged compared to their local competitors exhibit a higher probability of improving their ratings, consistent with competitive catch-up behavior. These star ratings are publicly available and are supposed to aid beneficiaries in choosing their MA plans from the set of alternatives. The probability of quality improvement initiatives increases with an increase in the magnitude of the difference between the country's mean quality rating and the plan's quality rating, demonstrating the effort of low-performing plans in a market to catch up.

In the following section, we will introduce a model of firm choice of quality investment initiative where we account for both the above-observed effect of quality bonus payments and quality competition in a market to model a firm's decision to invest in quality rating improvement.

3 Model

We introduce a stylized model of a managed competition setting where firms offer differentiated insurance plans in each market to generate revenue by enrolling beneficiaries in a market. The two main components of our model are the firm's profit function and the beneficiary demand function.

Table 2: **Probit Estimation: Effect of Marginal Contracts on Rating Improvement Probability**

	(1) Increase	(2) Increase
Marginal	0.0375***	0.0376***
	(0.00411)	(0.00411)
Lowstar	-0.198**	-0.199**
	(0.0737)	(0.0737)
Diff_meanstar	0.220***	0.220***
	(0.00397)	(0.00397)
Total_contract		0.000134
		(0.000328)
Year FE	Yes	Yes
State FE	Yes	Yes
N	54,050	54,050

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Marginal effect on the probability of increasing star rating in the next year. The variable Marginal denotes whether the current rating is 3.5, Lowstar is contracted with a rating of 2 stars or less, and Diff_meanstar denotes the difference between the Mean rating of the county and the plan's star rating. Specification (2) controls for the total number of contracts present in the county Total_contract. Observations are at the year-market-contract level.

In each period, a plan's market share is determined by the demand function where the beneficiaries choose from a set of alternatives and their choice depends on the plan characteristics, annual premium[4], observed quality measure provided by the star rating, plan level heterogeneity, and the beneficiary's private taste. For each beneficiary a firm enrolls in its plans, it earns a variable profit, which comprises the annual plan premium, and government subsidies. They incurs a marginal cost of providing service. The government subsidies depend on the given market characteristics and the quality ratings, marginal costs depend on the plan characteristics, both observed and unobserved, and the equilibrium plan premiums are set by the firms. In every period, the firms can decide to invest in quality and incur an investment cost associated with it.

The choice of firms' quality investment determines the quality ratings for the next period. The firms are forward-looking and maximize their expected discounted payoff while investing. This payoff depends on the firms' own quality ratings, as well as those of their competitors, and quality investments are strategic. Our model of the dynamic game closely follows Aguirregabiria and Ho

[4]A positive premium in our model is the excess premium some plans charge over the standard Part A B premiums Medicare beneficiaries already pay

(2012) [5] and Sweeting (2013) [32].

3.1 Framework

Medicare Advantage contracts, henceforth firms $f = 1, 2, ..., \mathcal{F}_m$ in a particular market m plays a discrete time infinite horizon game with periods $t = 1, ..., \infty$. Each firm offers a set of plans in each county, which we will consider as a separate market in this model. Markets are assumed to be independent of each other.[5]. In each period, a firm observes its payoff-relevant state variables, denoted by S_{fmt}, which include the firm's own characteristics, such as its star ratings, competitors' characteristics, and market characteristics. These state variables are assumed to be public information observed by all the firms in a market. A market is characterized by the CMS assigned subsidy benchmark B_{mt}, the total number of MA-eligible population M_{mt}, and the FFS spending quartile. County benchmarks determine the plan-specific subsidy amount using the reported plan cost and the CMS payment rule as described in the previous section. We assume that in each period, MA contracts (firms) observe the market characteristics, their own characteristics, and competitors' characteristics, set the premiums of each plan they offer, and make a quality investment choice. Premiums affect the profit of the current period, whereas investment in quality improvement affects their ratings and consequently payoffs for the next period. We drop the market index m below for notational convenience.

In each market period, a firm f offers a set of plans $\mathcal{J}_f$ and generates revenue when an eligible beneficiary chooses their plan. The set of firms and plans is assumed to remain the same over time. For each individual beneficiary who enrolls in one of their plans, a firm receives the extra annual premium and the per-member subsidy assigned by CMS and incurs a marginal cost. The firm also chooses whether to invest in quality improvement. Any decision to improve quality results in an increase in star ratings in the next period and incurs a quality investment cost. We assume that this decision is made at the contract level. The star ratings are provided at the contract level, and the insurers make strategic decisions like exit decisions, network formation, administrative and marketing infrastructure, etc., at the contract level and shared across all plans.

In every period, the firms choose an action $a \in A(S_{ft})$ from a discrete set of possible actions that determines their quality rating in the next period. A firm may decide to improve its rating or not improve at all. The action choices are state-dependent, as, for example, contracts with a current rating of 5 stars cannot improve their rating in the next period. Each action choice of the firm is associated with private information, independent and identically distributed (i.i.d.) payoff shock ν_{aft}. These shocks are assumed to follow a Type 1 extreme value distribution with dispersion parameter σ_ν. The sequence of the firm's decision is as follows:

1. Each firm observes payoff relevant state variable S_{ft} at the beginning of the period t
2. Each firm sets a premium for each plan $j \in \mathcal{J}_f$ offered in the market and reports their cost of providing service to CMS

[5]This assumption, though restrictive, is essential for the identification of the structural parameters of our dynamic model

3. Each firm observes the vector of action specific payoff shocks $\nu_{ft} = \{\nu_{aft} : a = 0, 1, 2\}$ and chooses an action $a \in A_{S_{ft}}$ to maximize the discounted sum of future payoffs

4. Each firm earns variable profit $R_{ft}(S_{ft}; \beta^{dd}, \gamma^{mc})$ from beneficiary enrollment where β^{dd} and γ^{mc} are the demand and marginal cost parameters respectively.

5. Each firm incurs an investment cost for quality improvement $I(a, S_{ft}; \theta)$ based on their action and receives a payoff shock ν_{aft}

6. The state variables evolve to the next period according to firms' choices and transition rule

The firm's flow profit function given payoff relevant state S_{ft} is given by

$$\Pi_{ft}(a, S_{ft}, \nu_{ft}) = R_{ft}(S_{ft}; \beta^{dd}, \gamma^{mc}) - I(a, S_{ft}, \theta) + \nu_{aft} \tag{2}$$

We now describe in detail each component of the profit function and the state transition rules.

The variable profit is earned by firms by beneficiary enrollment in their MA plans and is given by:

$$R_{ft} = \sum_{j \in \mathcal{J}_f} \{(p_{jt} + Reb_{jt} - mc_{jt}) \times s_{jt} \times M\}, \tag{3}$$

where M is the total Medicare eligible population in the market, s_{jt} is the market share of each plan j obtained from the demand model. We will later describe how the market shares are determined in this framework and how it depends on the observed star ratings.

The amount of subsidy each plan receives per beneficiary is given by the payment equation (1) and depends on the market benchmark, the star rating, and the reported costs[6]. The star ratings affect firms' variable profits by increasing government subsidies in the form of quality bonus payments. Plans with a star rating greater than or equal to four stars receive bonus payments in the form of higher rebate percentages and higher benchmark, as detailed in Section 2.4. In our counterfactual analysis, we will change the current national threshold-based bonus payment system to a comparison-based quality bonus payment system. Thus, in our counterfactual, the relationship between a plan's star rating and its annual revenue will be altered by this component of our model. Given the timeline of the game, the equilibrium premium for each plan in a market is determined by maximizing the variable profit R_{ft} with respect to $p_{jt} \forall j \in \mathcal{J}_f$.

In each period, a firm takes a quality improvement decision and incurs an investment cost $I(a, S_{ft}, \theta)$ based on its choice of action a and parameter θ. We assume the following parametric specification of investment cost:

$$I(a_{ft} = 1, S_{ft}) = \theta_1 + \theta_2 \mu_{ft} + \theta_3 PS_{ft} \tag{4}$$

[6]The reported cost is assumed to be a function of the true marginal cost. This assumption is based on the fact that the Affordable Care Act mandates each MA plan to have a minimum medical loss ratio of 85% which is observed by the CMS through its auditing process, and failure to comply results in punitive actions. A similar assumption has been used in the literature previously by Miller et al. (2019). [27]

We let this investment cost depend on the current star rating μ_{ft} and whether quality improvement occurred in the previous period PS_{ft}, which takes the value 1 if the firm improved its rating in the previous period and 0 otherwise. In the data, we observe that most contracts do not improve ratings in consecutive periods. The scope of improvement is usually diminished for a contract that increased its ratings recently, as the firms have to find new avenues through which the measures can be improved. This effect is captured by PS_{ft}. The set of parameters θ is our structural parameters of interest that determine the investment cost of improving quality that we will estimate. We can not identify any fixed cost of maintaining quality for a contract, and this quality investment cost is interpreted as the incremental cost for the contracts that are improving their star rating over the ones that are not. Although we do not observe investment decisions of the firms explicitly in the data, observed changes in the contract level star rating are used to identify the parameters of Equation 4.

As we are not directly observing the investment decisions but interpreting our quality improvement cost as the incremental cost for the contracts observed to improve their star ratings, we investigate whether these contracts improve their raw scores of the quality measured as compared to the ones with unchanged star ratings and the improvement in the star ratings are not purely a construct of how these measures are being aggregated or changes in the definition of the cutoffs. We explore the association between changes in the five groups of raw quality measures reported by the CMS used to construct the star ratings and the changes in the aggregate star rating for the contracts in Table A1. We observe that contracts that improve their star ratings are associated with significantly higher average raw scores for staying healthy, managing chronic conditions, and member experience. The positive coefficients are mildly significant for customer services and insignificant for complaints and performance [7]. This provides reduced-form evidence that contracts improving their star ratings also experience improvements in several underlying quality measures and that rating gains are not purely due to mechanical changes induced by the cutoff-based construction of the star rating system.

The star rating for the next period is assumed to evolve based on the firm's action choice as follows:

$$\mu_{ft+1} = \begin{cases} \mu_{ft} + \epsilon^{a_0}, & \text{if } a_{ft} = 0 \\ \mu_{ft} + 0.5, & \text{if } a_{ft} = 1 \end{cases} \tag{5}$$

where ϵ^{a_0} is a discrete random variable with the following probability mass function

$$P(\epsilon^{a_0}) = \begin{cases} P_0^{\epsilon^a}, & \text{if } \epsilon^{a_0} = 0 \\ P_{0.5}^{\epsilon^a}, & \text{if } \epsilon^{a_0} = -0.5 \\ P_1^{\epsilon^a}, & \text{if } \epsilon^{a_0} = -1 \end{cases}$$

where the values of $P_0^{\nu}, P_{0.5}^{\nu}, P_1^{\nu}$ are exogenously given. During estimation, they are empirically calculated as some of the contracts are observed to see a reduction in rating[8].

[7] The category Complaints and Performance also includes the measure known as "improvement stars," which rewards contracts that improve their ratings relative to the previous year. The absence of a significant effect for this category suggests that increases in overall star ratings are not primarily driven by past improvements.

[8] The star rating for a contract is mostly reduced due to institutional reasons difficult to incorporate in the model.

We next outline the determination of plan-level market shares. Although the primary objective of the paper is to analyze the supply-side implications of alternative quality bonus payment schemes, we develop the following demand model to generate predicted market shares, equilibrium prices, and marginal costs. These objects serve as key inputs for our counterfactual simulations and welfare analysis. In each period, t consumer i in market m chooses a MA plan $j = 0, 1, 2, ..., J_m$ from all available plans offered in a market where $j = 0$ is the outside option of not choosing any MA plan. We let individual utility depend on observed plan characteristics, plan type, and observed and unobserved quality. The utility individual i receives from a plan j in a given period t is given by the following utility function

$$U_{ijt} = \beta^p p_{jt} + \beta^\mu \mu_{jt} + X'_{jt}\beta^x + \xi_{jt} + \epsilon_{ijt}, \tag{6}$$

where p_{jt} is the annual plan premium,[9] μ_{jt} is the CMS assigned star rating of the plan, X_{jt} is the vector of plan types and plan characteristics and includes dummy variables indicating whether certain subsidiary services like vision, hearing, and vision, are covered by the plan. ξ_{jt} is the plan level characteristic that is observed by the consumers and the firms but is unobserved by the econometrician. This captures the plan-market level heterogeneity and also plan qualities that are not captured by the star ratings but can affect the demand. ϵ_{ijt} is the beneficiary-plan specific idiosyncratic taste shock assumed to be Type I extreme value distributed. We consider X_{jt} for each plan to be exogenous. Firms choose p_{jt} and μ_{jt} and are possibly correlated with ξ_{jt}. These two variables are treated as endogenous in our model.

The coefficients β^p, β^μ, and β^x are the structural demand parameters of our model. The market share for each plan is derived from the individual choice probabilities. A consumer chooses a particular plan if the utility of that plan is greater than all other available options in the market. Under the distributional assumption of ϵ_{ijt} the choice probability is given by the following logit form

$$s_{ij} = Pr(\text{i chooses j}) = \frac{exp(\delta_j)}{1 + \sum_{k=1}^{J_m} exp(\delta_k)}$$

where the plan mean utility δ_j of plan j is given by

$$\delta_{jt} = \beta^p p_{jt} + \beta^\mu \mu_{jt} + X'_{jt}\beta^x + \xi_{jt}$$

Most plans in our data are observed to be \$0 plans. It does not however mean that the beneficiaries of these plans pay a \$0 premium. Instead, they pay the standard part B premium for traditional Medicare. A positive premium is thus any extra amount the plan charges in addition to this. Since all individuals choosing a plan in our setting are medicare eligible, traditional fee-for-service Medicare or a PFFS plan are considered to be the outside option.

Failure to comply with CMS rules or report data on time is associated with a reduction to a one star score for the associated measure which often leads to deterioration in rating.

[9] This is the premium over the standard Part B premium that the beneficiaries need to pay. If a plan doesn't charge any extra premium other than the Part B premium they pay to Medicare, it is considered a 0 premium plan.

It should be noted that the star ratings are assigned at the contract level in a particular market and not plan level. This means that plans under the same firm uses the star rating of the contract they are under. Let $\mathcal{J}_f$ be the set of plans offered by the firm f. Thus $\mu_{jt} = \mu_{ft} \forall j \in \mathcal{J}_f$.

We assume that ξ_{jt} evolves according to the following AR(1) process:

$$\xi_{jt+1} = \rho^{\xi} \xi_{jt} + \zeta_{jt}, \tag{7}$$

where ζ_{jt} is distributed with mean 0, and standard deviation σ_{ζ} and can be interpreted as an unanticipated innovation shock that affects demand.

The marginal cost of enrolling an extra beneficiary in a plan is assumed to be constant and depends linearly on plan characteristics as follows:

$$ln(mc_{jt}) = X'_{jt}\gamma^x + \gamma_1^{\mu}\mu_{jt} + \gamma_2^{\mu}\mu_{jt}^2 + \alpha_{s(j)} + \omega_{jt}^{mc}, \tag{8}$$

where $\alpha_{s(j)}$ is the state fixed effect and ω_{jmt}^{mc} is the unobserved marginal cost component assumed to be i.i.d. distributed.

3.2 Value Function and Dynamic Game

We assume that firms use a stationary Markov Perfect Nash Equilibrium in pure strategies in every market. A Markov Perfect strategy Γ_f of a firm is a mapping from (S_{ft}, ν_{ft}) to an action a_{ft}. We assume that the firms are forward looking and they maximize their sum of the discounted expected future payoff. A firm's value function given strategies of all firms Γ is given by the following Bellman equation:

$$\begin{aligned} V^{\Gamma}(S_{ft}, \nu_{ft}) = max_{a \in A(S_{ft})}[R_{ft}(S_{ft}; \beta^{dd}, \gamma^{mc}) - I(a, S_{ft}, \theta) + \nu_{aft} \\ + \beta E\{V^{\Gamma}(S_{ft+1})|a, \Gamma_{-f}, S_{ft})\}], \end{aligned} \tag{9}$$

where $\beta = 0.95$ is the given discount factor assumed and $V^{\Gamma}(S_{ft+1})$ is the firm's value in the particular state before the realization of ν_{ft} takes place. Thus $E\{V^{\Gamma}(S_{ft+1})|a, \Gamma_{-f}, S_{ft})\}$ is the firm's expected future value given the current state, action of the firm and strategies of other firms Γ_{-f}. The expectation is taken over the possible realization of the state variables in the next period where a firm's own star rating evolves depending on firm's action a, the star rating of the competing firms evolve following Γ_{-f}, and all other state variables follow their respective transition rules.

We define the action specific value $v_f^{\Gamma}(a, S_{ft}, \Gamma_{-f})$ as the sum of the flow profit and discounted expected future value at any given state and for a particular action choice as follows:

$$v_f^{\Gamma}(a, S_{ft}, \Gamma_{-f}) = R_{ft}(S_{ft}; \beta^{dd}, \gamma^{mc}) - I(a, S_{ft}, \theta) + \beta E\{V^{\Gamma}(S_{ft+1})|a, \Gamma_{-f}, S_{ft}).$$

Under the Type I extreme value distributional assumption of the action specific payoff shocks, the conditional action choice probability of a firm can be written as a logistic function of the action specific values as:

$$P_f(a|S_{ft},\Gamma) = \frac{exp(v_f^{\Gamma}(a, S_{ft}, \Gamma_{-f})/\sigma_\nu)}{\sum_{a'\in A(S_{ft})} exp(v_f^{\Gamma}(a', S_{ft}, \Gamma_{-f})/\sigma_\nu)}. \tag{10}$$

3.2.1 Approximating the Value Function

The state space in our setting is exceptionally large which poses a problem for solving the value function. A firm's payoff relevant state variables include the firm's and competitors' own characteristics, characteristics of its competitors, and also market characteristics. Some of these state variables are also continuous variables. Solving for the value function for all states is computationally an infeasible task. In order to address this issue we follow Benitez-Silva et al. (2000) [8], Aguirregabiria (2012) [3] and Sweeting (2013) [32] and use a parametric approximation of the value function.

We approximate the value function using a linear representation based on a finite set of basis functions. Let S denote the state vector. We assume that the value function can be represented as a linear combination of K known functions $\{\phi_k(\cdot)\}_{k=1}^K$, such that

$$V(S) = \sum_{k=1}^{K} \lambda_k \phi_k(S). \tag{11}$$

Here, $\{\lambda_k\}_{k=1}^K$ are unknown coefficients to be estimated, and $\phi_k(S)$ are predetermined transformations of the state variables. For expositional clarity, we refer to $\{\phi_k(S)\}$ as the *approximating functions* (or basis functions) for the value function. This linear approximation substantially reduces the dimensionality of the dynamic programming problem. Rather than solving directly for the value function at every possible point in the (potentially large or continuous) state space, we solve instead for the finite-dimensional parameter vector $\lambda = (\lambda_1, \ldots, \lambda_K)$. Consequently, the computational burden shifts from functional iteration over the state space to estimation of a relatively small set of linear parameters.

The quality of the approximation depends on the richness of the chosen basis. Under standard regularity conditions, as K increases and the basis becomes sufficiently flexible, the approximation error can be made arbitrarily small. In practice, however, there is a tradeoff between approximation accuracy and computational tractability, since larger values of K increase estimation complexity and may lead to overfitting. In the implementation, we use payoff-relevant state variables that include a firm's current star ratings, lagged star ratings, and the star ratings of market competitors. Lagged star ratings are included to capture the fact that plan payments depend on ratings from the previous period due to the timing of the Medicare Advantage payment rules. Including lagged ratings, therefore, allows the model to account for the payment incentives created by past performance. This representation provides a flexible and computationally feasible approach for solving the dynamic model and conducting counterfactual and welfare analysis.

4 Estimation Results

We proceed by first estimating the demand parameters and the plan level unobservables ξ_{jt}. We then back out the marginal costs from the firm profit maximization condition and estimate the model of marginal cost. With all the observed and estimated state variables, we estimate our dynamic parameters.

4.1 Estimation of Demand Parameters

In our demand estimation, we address two distinct sources of endogeneity. In addition to the standard price endogeneity that arises in differentiated-product demand models, we also confront potential endogeneity in observed star ratings. The plan-level unobservable demand component, ξ_{jt}, follows an AR(1) process:

$$\xi_{jt} = \rho^{\xi}\xi_{jt-1} + \zeta_{jt}, \tag{12}$$

where ζ_{jt} denotes the innovation. At the same time, current star ratings depend on the plans' action choices in the previous period. Because these actions may be correlated with past realizations of ξ_{jt}, and because ξ_{jt} is persistent over time, observed star ratings and current demand unobservables may be correlated. This creates an additional endogeneity concern beyond the usual simultaneity between prices and demand shocks.

Prior work on endogenous quality choice, such as Fan (2013)[17], addresses this issue using demographic characteristics of competitors' markets as instruments for quality. In contrast, we exploit the panel structure of the data together with the AR(1) transition of ξ_{jt} to construct moment conditions. In particular, we assume that the innovation ζ_{jt} is uncorrelated with μ_{jt}, since star ratings in period t are determined by actions taken in period $t-1$, prior to the realization of ζ_{jt}. Under this assumption, the innovation is orthogonal to the contemporaneous star ratings, i.e. $E[\zeta_{jt}|\mu_{jt}] = 0$

Let Z denote the full set of instruments, which includes exogenous product characteristics X_{jt}, star ratings μ_{jt} under the maintained timing assumption, and standard price instruments. For the latter, we employ BLP-style demand instruments following conlon2023incorporating. The key orthogonality condition is

$$\mathbb{E}\left[Z'\zeta\right] = 0, \tag{13}$$

which forms the basis for identification of the demand parameters.

From equation (12), the innovation can be written as

$$\zeta_{jt} = \xi_{jt+1} - \rho^{\xi}\xi_{jt}, \tag{14}$$

where

$$\xi_{jt} = \delta_{jt} - \beta^{p}p_{jt} + \beta^{\mu}\mu_{jt} + X_{jt}\beta^{x}. \tag{15}$$

Substituting this expression into the innovation equation allows ζ_{jt} to be written entirely in terms of observables and the parameter $\beta^{p}, \beta^{\mu}, \beta^{x}, \rho^{\xi}$. We estimate the demand parameters using a quasi-difference moment approach, which leverages the AR(1) structure to difference out the persistent unobservable component.

Marginal costs are not directly observed. We recover them using firms' static profit-maximizing first-order conditions. For each product j belonging to firm f, the first-order condition is given by

$$s_{jt} + \sum_{j' \in \mathcal{J}_f} \frac{\partial s_{j't}}{\partial p_j} \left(p_{j't} + Reb_{j't} + \widehat{mc}_{j't} \right) = 0, \qquad \forall j \in \mathcal{J}_f, \tag{16}$$

where $\mathcal{J}_f$ denotes the set of products offered by firm f.

Solving this system of linear equations yields implied marginal costs $\widehat{mc}_{jt}$. We then parameterize marginal costs as a function of observed cost shifters and estimate the marginal cost parameters γ^{mc} using OLS.

4.2 Estimation of Dynamic Model

We begin our estimation of the dynamic model by defining the expected profit for each state before the realization of action specific shock given a conditional choice probability (CCP) P_f:

$$\widetilde{\Pi}\left(P_f, \mathcal{S}_{ft}\right) = R\left(\mathcal{S}_{ft}; \beta^{dd}, \gamma^{mc}\right) + \sum_{a \in A(\mathcal{S}_{ft})} P_f\left(a \mid \mathcal{S}_{ft}\right)\left(-I(a, \mathcal{S}_{ft}\theta) + \sigma_v\left(\chi - \log\left(P_f\left(a \mid \mathcal{S}_{ft}\right)\right)\right)\right) \tag{17}$$

where χ is the Euler's constant. This expression follows from the assumption that the choice-specific shocks are i.i.d. Type I Extreme Value. Under this distributional assumption, the expected value of the maximum over actions admits a closed-form representation in terms of conditional choice probabilities. Given the conditional choice probability and firm's state variables, following Aguirregabiria and Mira (2007) [6] the Bellman equation can be rewritten in terms of the expected value function. Let $V(\mathcal{S}_{ft})$ denote the ex-antevalue function, i.e. before the realization of the choice specific shocks. The Bellman Equation becomes

$$V(\mathcal{S}_{ft}) = \widetilde{\Pi}\left(P_f, \mathcal{S}_{ft}\right) + \beta E_{P_f}[V(\mathcal{S}_{ft+1}) | \mathcal{S}_{ft}]$$

where the expectation is taken over next-period states induced by the transition probabilities and the CCPs P_f. Substituting the value function using equation [**?**], where we approximate the value function as a linear representation of finite set of basis functions and using linearity of expectation, yields

$$\sum_{k=1}^{K} \lambda_k \phi_k(S_{ft}) = \tilde{\Pi}(P_f, S_{ft}) + \beta \sum_{k=1}^{K} \lambda_k E_{P_f}\left(\phi_k(S_{ft})\right). \tag{18}$$

Given conditional choice probabilities, parameters and observables, $\widetilde{\Pi}\left(P_f, \mathcal{S}_{ft}\right)$ can be calculated. The linear parameters for the value function approximation can be easily estimated from the above equation using a simple OLS.

$$\sum_{k=1}^{K} \lambda_k \left[\phi_k(S_{ft}) - \beta E_{P_f}\left(\phi_k(S_{f,t+1}) \mid S_{ft}\right)\right] = \widetilde{\Pi}(P_f, S_{ft})$$

We employ an Iterative Nested Pseudo-Likelihood estimation procedure as follows:

1. We start with an initial conditional choice probability P^0 by using a reduced form multinomial logit model of observed actions on a set of firm characteristics and an initial guess θ^0 of the structural parameter

2. At every step of the iteration, we calculate the expected profits given the state variable and conditional choice probability P_{iter} and parameter $\hat{\theta}$

3. Estimate the parameters $\hat{\lambda}^{P_{iter}}$ given the choice probabilities

4. Use $\hat{\lambda}^{P_{iter}}$ to calculate the choice specific value function, using equation 10 to form the pseudo-likelihood function and minimize it for the estimated values of $\hat{\theta}'$

5. Use equation 10 to update the conditional choice probabilities to P'_{iter}

6. The procedure stops if absolute differences between P_{iter} and P'_{iter} , and $\hat{\theta}$ and $\hat{\theta}'$ are less than the tolerance level set at 10^{-4}

For the identification of the quality investment cost parameter, we assume that the markets are independent and the same Markov Perfect Bayesian Nash equilibrium is used in every market, which is a standard identification assumption in the literature. The conditional choice probability function can be identified by exploiting the variations in the state variables and action choices across market. Further details regarding the identification and estimation procedure can be found in Sweeting (2013) [32] and Aguirregabiria et al. (2021) [4].

4.3 Estimation Results

This section presents the empirical results of our model using the aforementioned estimation procedure.

Table 3 reports the estimated parameters of our demand model. These estimates, in general, correspond to sensible priors and consistent with the institutional details of the Medicare Advantage (MA) market. The coefficient on the annual premium (measured per $1,000) is -1.62 implying a downward-sloping demand curve. The estimate is economically meaningful: holding other characteristics constant, higher premiums significantly reduce enrollment. The magnitude suggests substantial price sensitivity, which is consistent with the presence of close substitutes within local MA markets.

The coefficient on star ratings is positive and statistically significant. This indicates that beneficiaries value plan quality when making enrollment decisions. A higher star rating increases market share and, through higher enrollment, raises variable profits. Importantly, this finding implies that quality competition is an active margin in the MA market. Plans have incentives to improve their ratings not only because of direct financial bonuses from the government, but also because higher ratings shift demand in their favor relative to competitors. We also find that beneficiaries prefer plans that offer vision coverage but have no significant preference for dental or hearing coverage. Plan type also plays a major role in demand as beneficiaries prefer PPOs compared to HMOs.

Finally, the persistence parameter for the demand unobservable, is estimated to be 0.875 (s.e. 0.007). This high degree of persistence implies that unobserved demand shocks are strongly serially correlated over time. Economically, this suggests that plan-level reputation, brand capital, or other persistent quality components evolve gradually rather than fluctuating idiosyncratically from year to year. This finding validates the importance of modeling the dynamic structure of unobserved demand.

Table 3: **Estimates for Demand Parameter**

Variable	Coefficient	S.E.
Annual Premium (per $1000)	-1.62	0.82
Star	0.19	0.09
Vision	3.33	0.91
Hearing	-0.3	0.38
Dental	-0.03	0.46
Local HMO	-7.09	0.83
Regional PPO	-6.08	0.74
Local PPO	-6.9	0.64
rho (ρ^{ξ})	0.875	0.007

Notes: Estimated demand parameters parameters of equation (6) and (7) based on year-market-plan observations. The standard errors are clustered at the county level.

Table 4 reports the estimated marginal cost parameters. The coefficient on star ratings is positive and statistically significant, indicating that higher quality ratings are associated with higher marginal costs of serving an additional beneficiary. This suggests that improvements in measured quality require greater resource use, such as enhanced provider networks, better care coordination, or higher service intensity. At the same time, the coefficient on the squared star term is negative and significant. This implies a concave relationship between quality and marginal cost: while increases in star ratings raise marginal costs, the incremental cost of further quality improvements diminishes at higher rating levels. In other words, moving from a low to a moderate rating is more costly than moving from a high to an even higher rating. This pattern is consistent with fixed investments in quality infrastructure that yield decreasing marginal cost increases once core systems are in place.

The coefficient on the demand unobservable, ξ_{jt}, is negative and statistically significant. One interpretation is that plans with higher unobserved demand-side quality—captured by positive demand shocks—may possess cost-saving capabilities or operational efficiencies not directly observed in the data. For example, better management practices or more efficient care coordination may both increase consumer valuation and reduce per-enrollee costs. In addition, the cost of providing service is higher for PPO plans compared to HMO plans, consistent with the broader provider networks and greater flexibility typically offered by PPOs.

Table 5 we report the parameters of our dynamic model which are the estimates of improving the quality rating of a firm in a market. These estimates report the cost of improving a firm's star rating by half a star and are reported in millions of dollars.

We estimate that contracts with 2-star ratings undergo a fixed investment cost of $ 0.64 million dollars per county when they improve their rating by 0.5 stars. In our data, this is the lowest star

Table 4: **Estimates of marginal cost function**

	(1)
	ln_mc
eye	0.0233***
	(0.00184)
hearing	-0.0100***
	(0.00126)
dental	0.00509***
	(0.00132)
star	0.0523***
	(0.00845)
star_square	-0.00467***
	(0.00115)
localhmo	8.985***
	(0.0173)
ppo	9.010***
	(0.0167)
unobservables (ξ_{jt})	-0.00260***
	(0.000331)
State FE	Yes
N	146908

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Estimation of equation 8 based on year-market-plan observations. The dependent variable is the log of the estimated marginal cost of each plan, and the dependent variables are the plan characteristics with state-fixed effects. The standard errors are clustered at the county level

rating observed. This cost increases by 0.04 million dollars per one-star rating increase in their current rating. This cost increases by $0.04 million for each one-star increase in the contract's current rating. For example, a three-star-rated contract faces an additional $0.04 million relative to a two-star contract, resulting in a total fixed investment cost of $0.68 million per county for a 0.5-star improvement. These estimates suggest that improving quality becomes progressively more costly as the current star rating increases. In addition, the cost of improving the star rating in a given period rises by $0.304 million if the contract improved its rating in the previous period.

Intuitively, these results are consistent with the idea that opportunities for further quality improvements become more limited as a contract's existing quality level increases. Higher-rated con-

Table 5: **Parameters for Quality Improvement Investment Cost**

Investment Cost for Half Star Improvement	$ million
Baseline cost 2 star contract	0.6448 (0.004)
Incremental cost per additional current star above 2	0.0396 (0.0084)
Incremental cost if rating changed in previous period	0.3039 (0.002)
Scale Parameter (σ_v)	3.6785 (0.0431)

Notes: Reported costs are for a Medicare Advantage Contract per county in dollar millions based on parameter estimation of equation (4). The bootstrapped standard errors clustered at the county level are in parentheses for 1000 bootstrap samples.

tracts often rely on specialized professionals and consultants to identify and implement strategies for incremental improvements in their ratings, and our estimates reflect the increasing cost associated with such efforts. The estimated quality investment cost parameters also help explain how Medicare Advantage markets have evolved under the "cliff effect" created by the national threshold rule. Under this policy, plans receive bonus payments only if their star rating reaches four or above. We observe that markets that initially had higher-quality plans evolved over time to exhibit higher average plan quality, whereas markets that began with relatively low-quality plans appear to have remained trapped at those lower quality levels. Our estimated dynamic investment cost parameters, together with the current quality bonus payment policy and its threshold-based structure, can therefore partly explain the persistent geographic variation in average plan quality observed across markets.

Based on these estimates, the total investment cost for a 3.5-star rated contract to reach bonus status (i.e., move from 3.5 to 4 stars) without any prior quality improvement activity is approximately $0.78 million per county. In contrast, a 2.5-star rated contract attempting to reach the four-star threshold over two periods faces a substantially higher total cost of about $2.575 million. This investment corresponds to roughly one-third of the average annual contract revenue per county, which ranges between $7.2 million and $8.8 million in our sample. In percentage terms, the required investment represents approximately 29–36% of annual contract revenue, indicating that improving ratings to the four-star threshold entails a substantial but economically feasible cost relative to plan revenues. Thus, the cost of reaching bonus status is considerably higher for lower-rated contracts than for those that are already close to the threshold, implying that marginal contracts just below the cutoff face much lower investment costs than poorly rated contracts.

5 Counterfactual

We now return to the problem of implementing the competitive bonus payment rule. After estimating the structural model parameters, we explore how markets might evolve in terms of plan rating if quality bonuses are paid to contracts performing better than their peers in the local market.

5.1 Procedure

We begin by adjusting the variable profit of each plan in a market based on the observed star rating. We follow a budget-balanced redistribution procedure, where plans with a rating lower than the median star rating of the market get a benchmark reduction of 5%, and this deducted amount is redistributed equally to plans above the median star rating. The rebate rate remains at 50% for plans at or below the median rating and 70% for plans above the median. This change in payment rule is intended to induce more competition among peers in the market, where being better quality than the others in a market fetches more subsidies in terms of quality bonus payment, and relative underperformance comes with a penalty. In our counterfactual, the bonus payment is distributed to the contracts based on the current star ratings.

Due to a change in the quality bonus payment structure as described above, we begin our counterfactual by adjusting the per-member subsidy of each plan in a market as stated. Given the plan characteristics, we calculate the new equilibrium prices by iteratively solving for the set of first-order conditions given by the $\frac{\partial s_j}{\partial p_j}$ matrix starting from observed prices with a non-negativity restriction similar to Miller et al. (2019) [27]. To make sure that this procedure works correctly, we precisely match the observed prices in the data when implementing this procedure of finding equilibrium prices with the existing payment rule. We then proceed to estimate the market shares given the plan characteristics and the new equilibrium prices.

With the counterfactual variable profits for each observed state of the firm and the parameters for quality investment initiative costs, we estimate the new parameters of the value function λ^{count} using an iterative process similar to the one used in the estimation procedure. We begin with a guess of counterfactual conditional choice probabilities P_0^c for each state, calculate the expected profit before realization of payoff shocks in equation 17, estimate λ^{count,P_0^c} using equation 18. These λ^{count,P_0^c} and the estimated investment cost parameter $\hat{\theta}$ are used to update the conditional choice probabilities given by equation 10 to P_1^c. These updated conditional choice probabilities P_1^c are compared to the initial guess P_0^c. The iterative procedure continues till the maximum absolute value is less than the tolerance (10^{-6}) and convergence is achieved. The final iteration provides us with the counterfactual parameters λ^{count}.

We then simulate the market forward based on λ^{count} to see how markets evolve under the new payment rule. We start from 2013, and the previous step provides the equilibrium choice probabilities for observed states for this year. These probabilities and the estimated transition process of ξ are used to move the model forward one period. As this process takes the market to a state not observed in the data, from this step, we proceed by solving for equilibrium choice probabilities for every period we simulate the model forward. We use initial choice probabilities approximated by a multinomial logit model. [10]. We solve for the equilibrium choice probabilities in each market by iterating over the best response function.

5.2 Simulation Results

We start from the observed states in 2013 and simulate the markets forward. We do this in order to compare the simulated markets with observed data. In our simulation, we exclude states such as Alaska, which have very low MA market penetration during our study period, and non-continental states like Puerto Rico, dominated by a few large local insurers. We also exclude counties with insufficient star-rating data in 2013, the starting period of our simulation. Figure 5 shows the county average ratings in the United States in 2013. This was the initial period when the rating system and the national threshold bonus payment system were introduced. We observe that most regions in the United States perform poorly in terms of offered plan quality, except regions on the West Coast, the Midwest, and the Northeast, which perform better than other regions. Figure 6 shows the county average rating of 2016 observed in data, where we find that certain markets have evolved to have higher average plan quality compared to others like counties in the south.

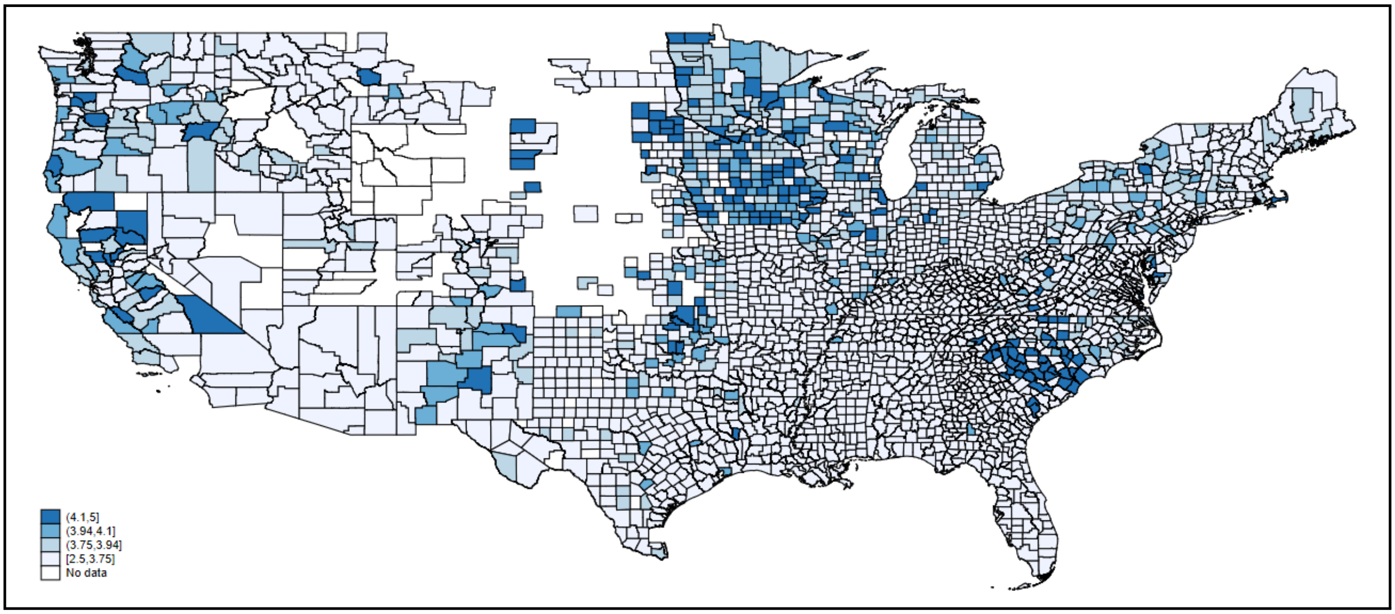

Figure 5: Geographical distribution of average MA plan star rating of counties in 2013

It is important to note that counties with higher average ratings in 2016 are mostly concentrated in geographic areas with better initial conditions in 2013, as measured by average rating. This might be due to two facts that our model captures:

[10]We initialize this iteration with a reduced form estimate of choice probabilities using a simple multinomial logit model of observed action on potential subsidy increase for each possible action if no other firms change and use this model to predict the choice probabilities under the new scheme by changing the covariates. See [33] Sweeting, 2013 Supplement

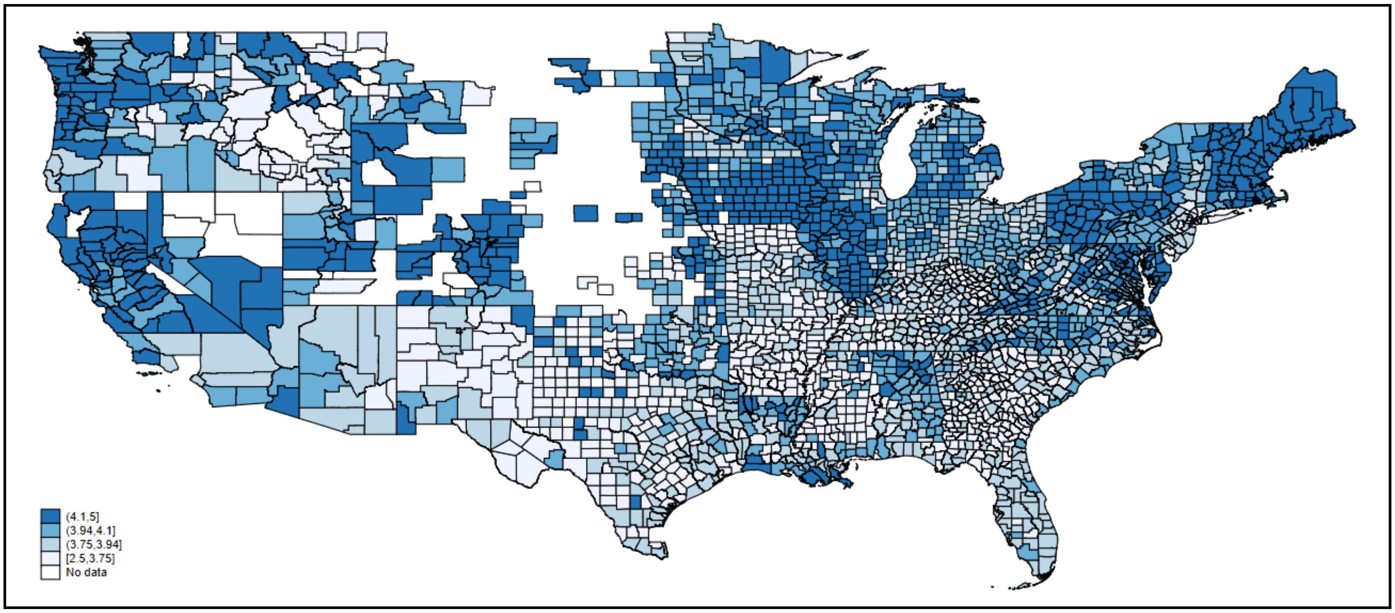

Figure 6: Geographical distribution of average MA plan star rating of counties in 2016

- There are more marginal firms in these regions that have incentives to improve their rating, as a relatively smaller investment makes them quality bonus eligible.
- Non-marginal firms are incentivized to improve their ratings because marginal firms are improving theirs, increasing quality competition in these markets through their effect on the demand side.

Our model of quality improvement initiatives captures both the overall and the regional distribution of mean county ratings if we simulate the model without implementing the new payment policy, as shown in Figure 7, starting from the initial market conditions, and predicts reasonably well the regions that evolved to a higher average quality in the data. The ability of our model to predict the market outcomes indicates that the model assumptions regarding the strategic interactions in terms of quality improvement initiatives are plausible. As the initial payment system pays quality bonuses only to plans with a rating greater than four, it is less costly for plans with better initial ratings to move to bonus status. For example, a plan with a 3.5 rating can move to bonus status by increasing its rating by only half a star, and this increase is associated with higher government subsidies. The cost of moving to a four-star rating is much higher for plans with poorer initial ratings. This, coupled with the strategic interaction among firms, drives the results for our model simulation. This is consistent with our estimates of investment cost parameter in Table 5.

We then introduce the competitive bonus payment system in our counterfactual, simulate our model forward, and compare it with observed data. Figure 8 shows the regions that would perform better under the competitive payment system. Low-performing regions like West Texas, New Mexico, and Florida. are predicted to improve under the new payment system. Among the counties studied in our counterfactual, 75% of the counties performed better under the counterfactual in terms of their average ratings than their observed average rating in 2016. Among the counties that are observed to improve in the counterfactual, the mean increase in their average increase in the

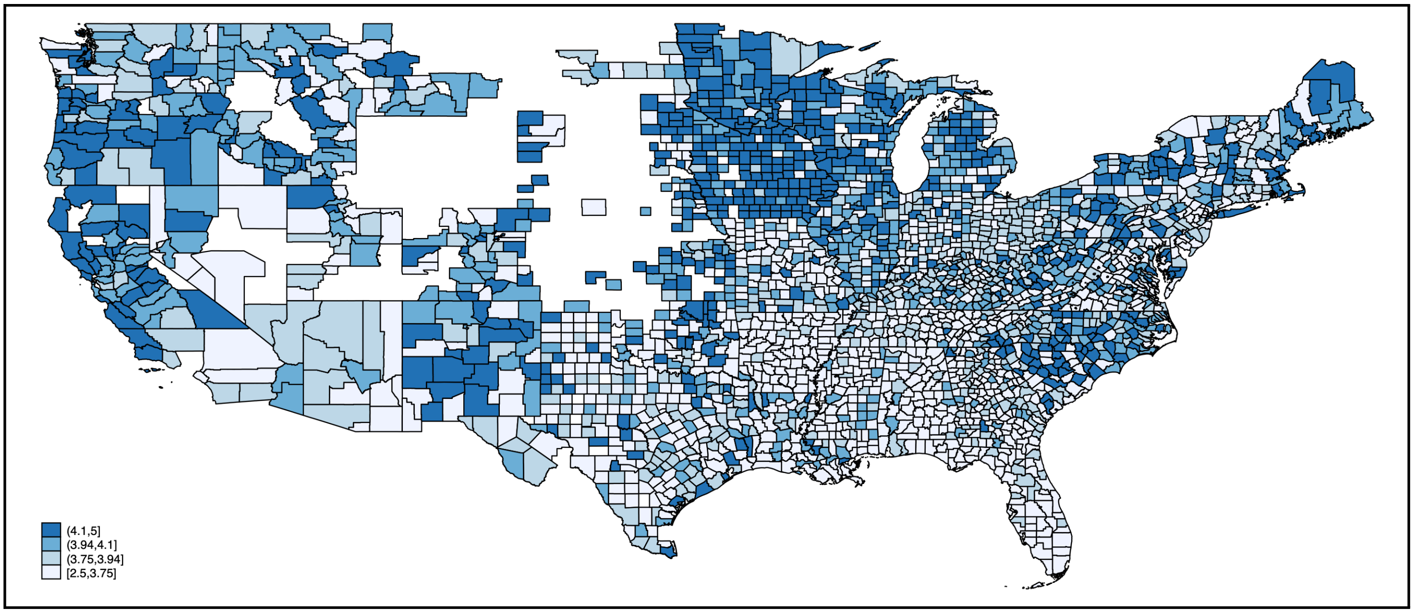

Figure 7: Geographical distribution of model-simulated average MA plan star rating of counties in 2016

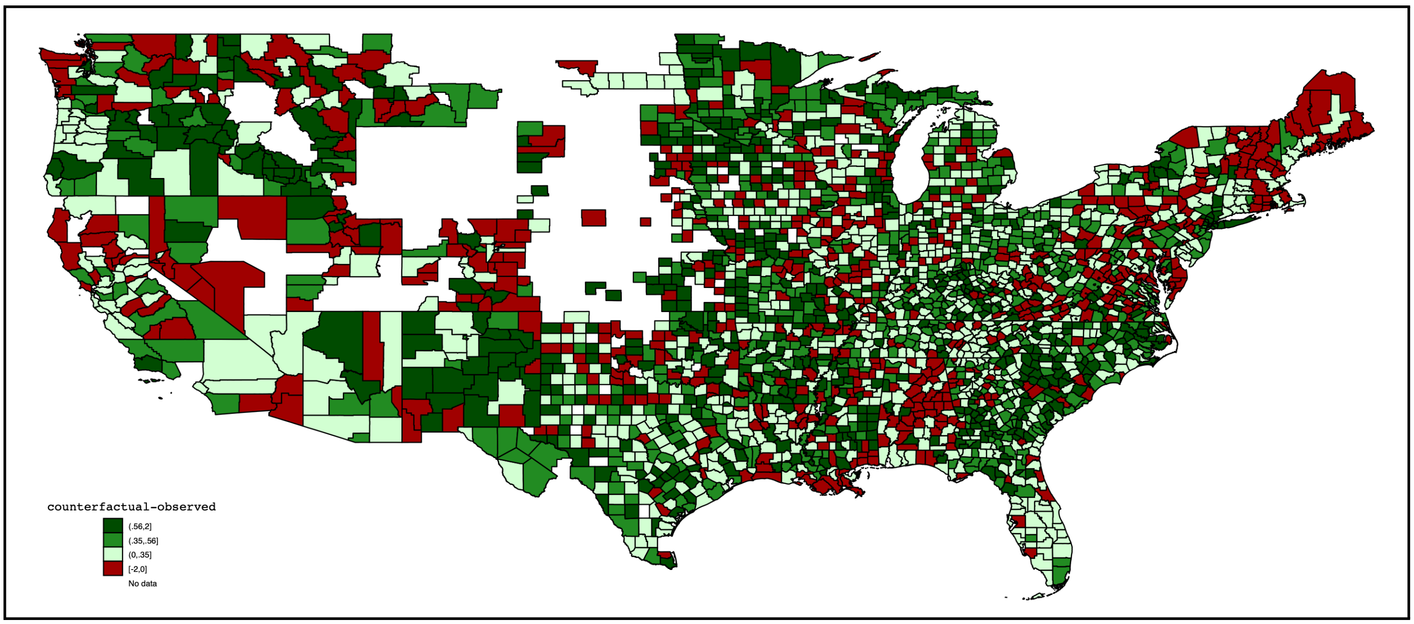

Figure 8: Counties improving mean star rating under the counterfactual payment rule

star is 0.44, and the median increase is 0.4.

Figure 9 compares the distribution of observed and simulated average county star ratings in the United States in 2016. It can be observed that the distribution shifts rightwards as 68.72% counties are predicted to have an average rating greater than four stars as compared to 32.75% in the data. This is a desirable outcome as the existing payment rule rewards contracts having a rating of four or higher. The competitive payment rule is predicted to achieve this goal more effectively under the

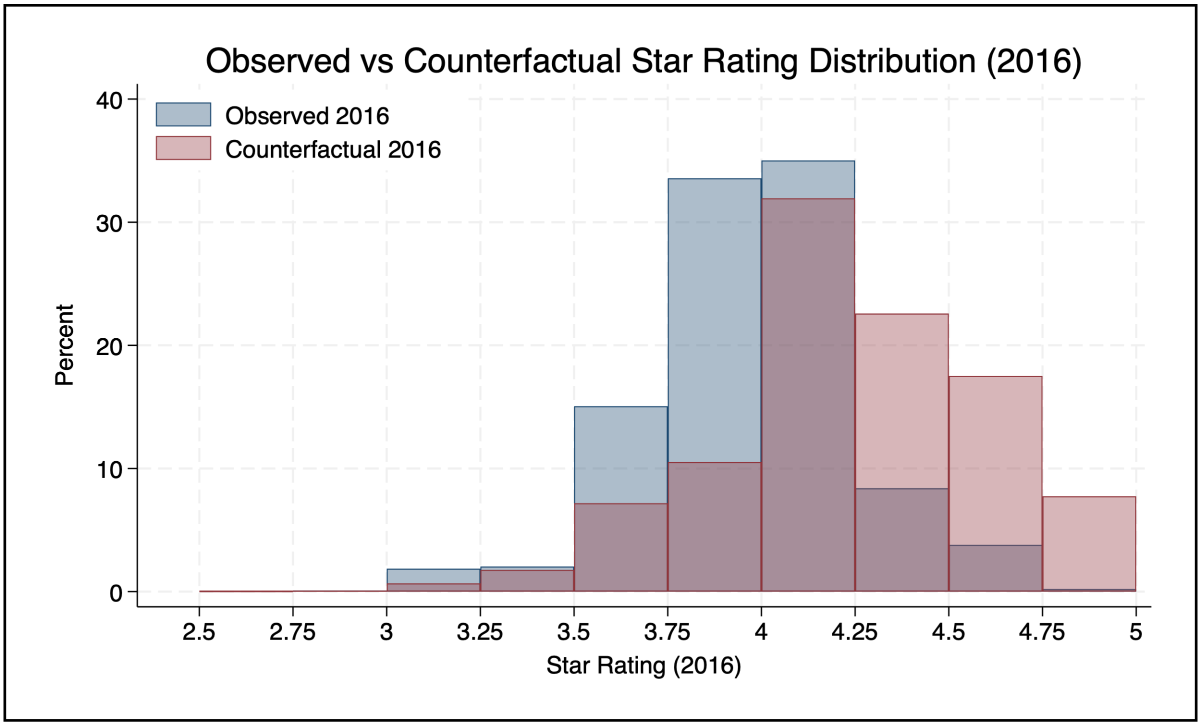

Figure 9: Distribution of mean county star rating in observed data and counterfactual. This figure shows that under the counterfactual payment rule, the distribution shifts to the right.

counterfactual bonus payment rule. Our findings suggest that most markets evolve to have an average rating between 4 and 4.5 stars. This is a desirable outcome both in terms of an overall increase in plan quality and a decrease in regional heterogeneity in the availability of higher-quality plans.

We then analyze which counties perform better under the new payment rule. We calculate the difference between the counterfactual county average rating and the observed mean of 2016 for each county. A positive value of this variable signifies that the county improves in terms of plan quality offering under the new payment rule. In Figure 10, we plot for each county the value of their observed average rating in 2016 against the calculated difference as defined above. We observe that most counties that improve have a lower average rating in the data. Only a small number of counties below an observed average rating of 3.5 worsen in the counterfactual. Most of the counties predicted to worsen under the new payment rule have an observed average rating of 4 stars.

This pattern can be indicative of how the firms might allocate resources for quality improvement initiatives under the new payment rule. As explained before, under the existing rule, it was more profitable for firms to invest in quality improvement of contracts that were already performing better but just below the national benchmark. The cliff effect of the current payment policy generates the incentive to do so. However, when the bonus payments are based on the relative performance of the contract in a market where the contract has to be better than its peers, improving the rating of a low-performing contract in a low-performing market can become profitable. Our counterfactual predictions capture this effect of the competitive bonus payment rule, as illustrated.

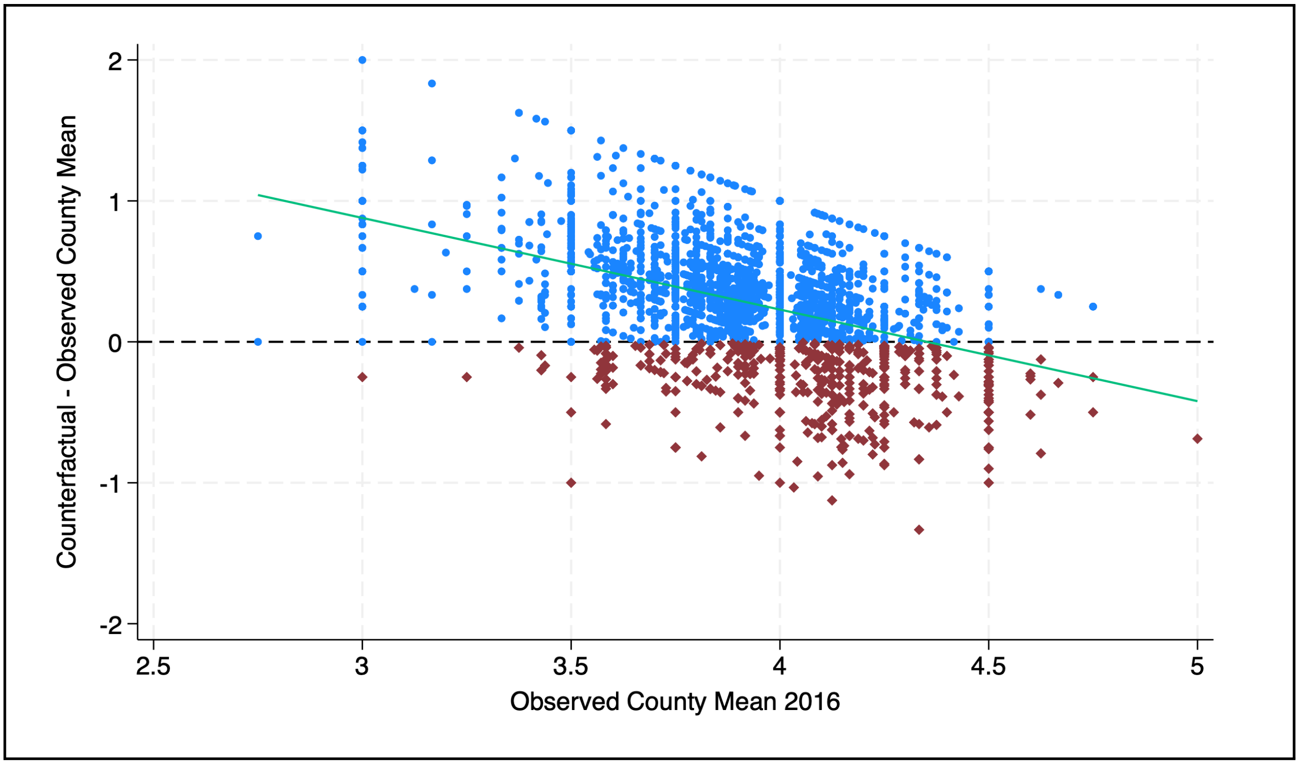

Figure 10: Scatter plot of the observed average rating of a county and its difference from the counterfactual prediction. A value greater than zero for the Y axis indicates that the market performs better in the counterfactual.

Finally, we analyze how the annual premiums will change under the counterfactual payment rule. The new bonus payment system is budget-balanced, as no extra dollar amount per beneficiary is injected into the system and is paid as transfers across plans. It is necessary to see how the firms might react in terms of premiums under such a payment system. In other words, we would like to see if this predicted increase in quality comes at an increased annual premium.

As shown in Figure 11, the distribution of annual premium shifts right as compared to the distribution in the observed data. The average annual premium increases to $318.12 from $237.76 in 2016. The number of plan-county observations that charge $0 extra premium falls to 2,793 from 14,002. In Table 6, we show that for 7.85% of plan-counties we study, the premium remains the same under counterfactual. All these correspond to plans charging no extra premium. Although 92.12% of plan-counties increase their premiums, the mean premium increase is $80.36 (median $88.33). The dispersion in premium adjustments is relatively limited, with a standard deviation of $11.01, suggesting that most plans raise premiums by a similar magnitude. This indicates that the counterfactual policy change leads to a broadly uniform upward shift in premiums rather than highly heterogeneous pricing responses across plan-counties. The reduction in the mass of zero-premium plans is particularly notable, as these plans play an important role in attracting price-sensitive beneficiaries in the Medicare Advantage market. The disappearance of a large fraction of such plans implies that the policy change could reduce the availability of zero-premium options, potentially affecting enrollment decisions for low-income beneficiaries.

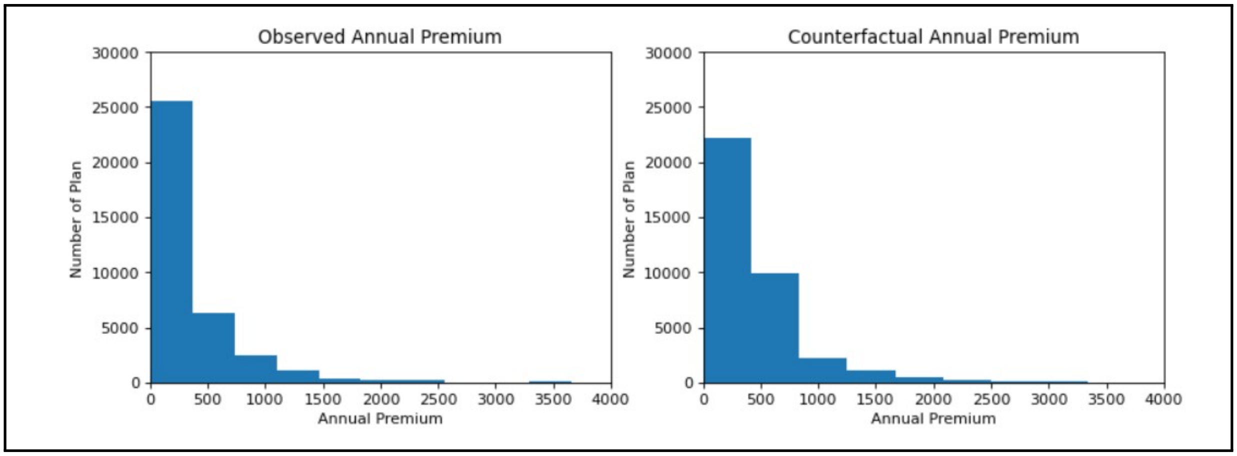

Figure 11: Distribution of Annual Premium in Observed Data and Counterfactual in 2016. This figure shows how the distribution of the annual premium of MA plans changes in the counterfactual.

Table 6: **Difference between Observed and Counterfactual Annual Premium**

Variable	Number of County-Plans (% of total observations)	Mean Difference in $ (S.D)
Observed Price= Counterfactual Price	2,784 (7.85 %)	0
Observed Price<Counterfactual Price	32,687 (92.12 %)	87.22 (11.01)

Notes: The level of observation is county-plan in 2016. This table reports the comparison of observed plan premiums and calculated counterfactual plan premiums in 2016.

Finally, because our counterfactual results show an increase in average star ratings across plans and across counties and also a mean overall increase in the excess premium that Medicare Advantage charges plans, we proceed to calculate the change in the consumer surplus under the new bonus payment system. The consumer surplus of individual i given product characteristics in a market can be written as follows:

$$CS_i = E[max_j u_{ij}]/\beta^p = \frac{1}{\beta^p} ln(1 + \sum_j exp(\delta_j))$$

where

$$\delta_j = \beta^p p_{jt} + \beta^\mu \mu_{jt} + X'_{jt}\beta^x + \xi_{jt}$$

This represents the mean utility gain converted to dollar amounts when any individual beneficiary chooses a Medicare Advantage plan over the outside option. We calculate the mean consumer surplus for each market under the existing payment system and our simulation. Our results predict a mean increase in this consumer surplus by 8.13% under the counterfactual payment system. However, this estimated increase in the consumer surplus should be taken with caution, given that we do not explicitly model the bidding behavior of the MA firms and the parsimonious nature of our demand model and may not fully capture strategic pricing responses or other adjustments by plans.

6 Conclusion

We study how competition among firms can be leveraged in a managed care setting such as Medicare Advantage. While much of the existing literature on managed competition focuses on improving financial efficiency, we instead examine how competitive forces can be used to create incentives for firms to improve the quality of service they provide. To this end, we propose a redistribution mechanism for quality bonus payments. Under this mechanism, plans within each local market are compared based on their measured quality, and bonus payments are financed through transfers from lower-performing plans to higher-performing plans within the same market. Such a design not only increases competitive pressure on firms to improve quality, but also makes quality improvement financially attractive, particularly in markets where overall performance is relatively low.

We provide a framework to analyze how firms would behave under such a competitive bonus payment system and predict how the markets evolve in terms of offered plan quality. We use a dynamic game model where forward-looking profit-maximizing firms strategically choose whether to invest in quality improvement. The firms also take into account their competitor's behavior as the payoffs for improving quality measures under the new proposed rule depend on their relative performance. We use estimated model parameters to calculate the equilibrium quality investment decision of the firms under the new quality bonus payment rule and simulate the model forward to see how they evolve. Our counterfactual predicts that the average quality of plan offerings improve under this new payment rule.

By comparing the simulated market outcomes under the redistribution bonus payment system with the observed 2016 data, we find that approximately 75% of the counties in our sample experience an improvement in average plan ratings. This suggests that the proposed policy substantially strengthens incentives for quality improvement relative to the existing system. Importantly, the gains are not uniform across markets. We observe that counties that historically had lower-performing plans exhibit the largest improvements under the new policy. In these markets, the redistribution mechanism creates stronger incentives for firms to invest in quality by allowing high-performing plans to benefit directly from transfers within the local market. As a result, the policy not only raises overall quality but also helps reduce persistent disparities in plan quality across geographic markets. We estimate an average increase in consumer surplus of 7.58% in our counterfactual.

In our model, we focus on firms' quality investment decisions and do not explicitly incorporate entry and exit behavior. In practice, however, market dynamics may also involve low-performing firms choosing to exit the market, reduce their presence, or merge with better-performing plans. Allowing for endogenous entry and exit decisions would provide a richer characterization of firms' strategic responses to quality-based incentives and is a promising direction for future research.

In addition, our framework does not explicitly model the mechanisms through which quality improvements occur, such as changes in provider network formation or contracting with healthcare providers. In reality, the ability of plans to improve quality may depend on local provider availability, existing healthcare infrastructure, and the structure of provider networks within a market. Incorporating these factors could shed further light on the constraints and opportunities plans face when attempting to improve their ratings. We therefore leave for future work the task of examining how provider network formation influences plan quality and how these dynamics interact with competition in Medicare Advantage markets.

References

[1] ABALUCK, J., CACERES BRAVO, M., HULL, P., AND STARC, A. Mortality effects and choice across private health insurance plans. *The quarterly journal of economics 136*, 3 (2021), 1557–1610.

[2] ADRION, E. R. Competition and health plan quality in the medicare advantage market. *Health services research 54*, 5 (2019), 1126–1136.

[3] AGUIRREGABIRIA, V. A method for implementing counterfactual experiments in models with multiple equilibria. *Economics Letters 114*, 2 (2012), 190–194.

[4] AGUIRREGABIRIA, V., COLLARD-WEXLER, A., AND RYAN, S. P. Dynamic games in empirical industrial organization. In *Handbook of industrial organization*, vol. 4. Elsevier, 2021, pp. 225–343.

[5] AGUIRREGABIRIA, V., AND HO, C.-Y. A dynamic oligopoly game of the us airline industry: Estimation and policy experiments. *Journal of Econometrics 168*, 1 (2012), 156–173.

[6] AGUIRREGABIRIA, V., AND MIRA, P. Sequential estimation of dynamic discrete games. *Econometrica 75*, 1 (2007), 1–53.

[7] AGUIRREGABIRIA, V., AND MIRA, P. Dynamic discrete choice structural models: A survey. *Journal of Econometrics 156*, 1 (2010), 38–67.

[8] BENITEZ-SILVA, H., HALL, G., HITSCH, G. J., PAULETTO, G., AND RUST, J. A comparison of discrete and parametric approximation methods for continuous-state dynamic programming problems. *manuscript, Yale University* (2000).

[9] BLEVINS, J. R. Nonparametric identification of dynamic decision processes with discrete and continuous choices. *Quantitative Economics 5*, 3 (2014), 531–554.

[10] CONRAD, D. A., AND PERRY, L. Quality-based financial incentives in health care: can we improve quality by paying for it? *Annual review of public health 30* (2009), 357–371.

[11] DAFNY, L., AND DRANOVE, D. Do report cards tell consumers anything they don't already know? the case of medicare hmos. *The Rand journal of economics 39*, 3 (2008), 790–821.

[12] DAMBERG, C. L., SORBERO, M., MEHROTRA, A., TELEKI, S., LOVEJOY, S., AND BRADLEY, L. An environmental scan of pay for performance in the hospital setting. Tech. rep., Rand Health Working Paper WR-474-ASPE/CMS. Santa Monica, CA, 2007.

[13] DARDEN, M., AND MCCARTHY, I. M. The star treatment estimating the impact of star ratings on medicare advantage enrollments. *Journal of Human Resources 50*, 4 (2015), 980–1008.

[14] DORAN, T., FULLWOOD, C., KONTOPANTELIS, E., AND REEVES, D. Effect of financial incentives on inequalities in the delivery of primary clinical care in england: analysis of clinical activity indicators for the quality and outcomes framework. *The Lancet 372*, 9640 (2008), 728–736.

[15] EIJKENAAR, F. Key issues in the design of pay for performance programs. *The European Journal of Health Economics 14*, 1 (2013), 117–131.

[16] ERICSON, R., AND PAKES, A. Markov-perfect industry dynamics: A framework for empirical work. *The Review of economic studies 62*, 1 (1995), 53–82.

[17] FAN, Y. Ownership consolidation and product characteristics: A study of the us daily newspaper market. *American Economic Review 103*, 5 (2013), 1598–1628.

[18] FARLEY, D. O., ELLIOTT, M. N., SHORT, P. F., DAMIANO, P., KANOUSE, D. E., AND HAYS, R. D. Effect of cahps performance information on health plan choices by iowa medicaid beneficiaries. *Medical care research and review 59*, 3 (2002), 319–336.

[19] FIORETTI, M., AND WANG, H. Performance pay in insurance markets: Evidence from medicare. *The Review of Economics and Statistics* (2020), 1–45.

[20] GOWRISANKARAN, G., AND TOWN, R. J. Dynamic equilibrium in the hospital industry. *Journal of Economics & Management Strategy 6*, 1 (1997), 45–74.

[21] GRUBER, J. Delivering public health insurance through private plan choice in the united states. *Journal of Economic Perspectives 31*, 4 (2017), 3–22.

[22] HANDEL, B. R., AND KOLSTAD, J. T. Health insurance for" humans": Information frictions, plan choice, and consumer welfare. *American Economic Review 105*, 8 (2015), 2449–2500.

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

[24] KAHNEMAN, D., KNETSCH, J. L., AND THALER, R. Fairness as a constraint on profit seeking: Entitlements in the market. *The American economic review* (1986), 728–741.

[25] KAHNEMAN, D., AND TVERSKY, A. Prospect theory: An analysis of decision under risk. *Econometrica 47*, 2 (1979), 263.

[26] LAYTON, T. J., AND RYAN, A. M. Higher incentive payments in medicare advantage's pay-for-performance program did not improve quality but did increase plan offerings. *Health services research 50*, 6 (2015), 1810–1828.

[27] MILLER, K. S., PETRIN, A., TOWN, R., AND CHERNEW, M. Optimal managed competition subsidies. Tech. rep., National Bureau of Economic Research, 2019.

[28] PAKES, A., OSTROVSKY, M., AND BERRY, S. Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *the RAND Journal of Economics 38*, 2 (2007), 373–399.

[29] REID, R. O., DEB, P., HOWELL, B. L., AND SHRANK, W. H. Association between medicare advantage plan star ratings and enrollment. *Jama 309*, 3 (2013), 267–274.

[30] Scanlon, D. P., Chernew, M., McLaughlin, C., and Solon, G. The impact of health plan report cards on managed care enrollment. *Journal of health economics 21*, 1 (2002), 19–41.

[31] Sen, A. P., Meiselbach, M. K., Anderson, K. E., Miller, B. J., and Polsky, D. Physician network breadth and plan quality ratings in medicare advantage. In *JAMA Health Forum* (2021), vol. 2, American Medical Association, pp. e211816–e211816.

[32] Sweeting, A. Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry. *Econometrica 81*, 5 (2013), 1763–1803.

[33] SWEETING, A. Supplement to "dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry". *Econometrica 81*, 5 (2013), 1763–1803.

[34] Vatter, B. Quality disclosure and regulation: Scoring design in medicare advantage. *Available at SSRN 4250361* (2022).

[35] Wedig, G. J., and Tai-Seale, M. The effect of report cards on consumer choice in the health insurance market. *Journal of health economics 21*, 6 (2002), 1031–1048.

Appendix

A Additional Tables

Table A1: Regression of Rating Improvement on Next-Year Quality Changes

	(1) Staying Healthy	(2) Managing Chronic	(3) Member Experience	(4) Complaints & Performance	(5) Customer Service
Intercept	4.053***	0.788	-0.472*	-6.374**	4.729**
	(0.507)	(0.636)	(0.249)	(2.040)	(1.669)
Improved_Rating	2.320**	3.244***	1.209***	0.749	5.028*
	(0.744)	(0.938)	(0.364)	(2.993)	(2.449)
Year 2014	3.626***	2.617**	-0.475	7.468*	-1.699
	(0.732)	(0.919)	(0.358)	(2.947)	(2.483)
Year 2015	-13.315***	4.096***	-0.797*	23.073***	-2.016
	(0.820)	(1.029)	(0.401)	(3.300)	(2.837)
Star Rating	-0.857**	-0.268	0.056	2.864*	-1.444
	(0.287)	(0.361)	(0.141)	(1.156)	(0.946)
Improved_Rating × 2014	-1.938	-1.883	0.253	-9.544*	-9.506*
	(1.318)	(1.656)	(0.647)	(5.304)	(4.490)
Improved_Rating × 2015	1.627	-0.729	0.146	5.316	0.669
	(1.217)	(1.530)	(0.598)	(4.897)	(4.206)
Improved_Rating × Star	-0.916*	-0.853	-0.326	1.909	-2.869*
	(0.519)	(0.654)	(0.254)	(2.088)	(1.708)
2014 × Star	-0.765*	-0.689	0.226	-13.629***	0.239
	(0.406)	(0.510)	(0.198)	(1.635)	(1.382)
2015 × Star	2.938***	0.172	0.169	1.211	1.204
	(0.439)	(0.551)	(0.215)	(1.768)	(1.516)
Improved_Rating × 2014 × Star	1.231	1.545	0.075	2.254	5.440*
	(0.911)	(1.144)	(0.449)	(3.664)	(3.077)
Improved_Rating × 2015 × Star	-0.416	0.891	-0.042	-4.074	-0.245
	(0.802)	(1.008)	(0.396)	(3.229)	(2.759)

Notes: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The intercept corresponds to 2-star rated contracts in the omitted baseline year with no improvement action. Improved_Rating is a dummy variable categorizing contracts that improved their MA star ratings, and the dependent variables are the corresponding changes in average raw quality scores for different measure categories. All specifications include full interactions between Improved_Rating, Year indicators, and baseline Star Rating. Contracts that improve their MA star ratings are associated with corresponding increases in raw scores across all categories except Complain & Performance.

B Estimation Details

We use a multinomial logit model to estimate the initial choice probabilities of the firms. The results of the first stage MNLogit is provided in Table A2. The first stage multinomial logit model has an accuracy of 77.9% in predicting the firms' action choice.

Table A2: Determinants of investment incentives in the dynamic model

Variable	Definition	Coefficient
Star	CMS Star Rating (quality measure)	0.1726
Dental_quality	Dental coverage quality (average per contract)	-0.0643
Vision_quality	Vision coverage quality (average per contract)	-0.0710
Hearing_quality	Hearing coverage quality (average per contract)	-0.0368
Unobservables	Unobservable plan/market heterogeneity (residual or fixed effect)	-0.0330
Plan_count	Number of plans under a contract	0.0684
Profit	Plan profit	0.00000083
State_count	Number of states the contract operates in	0.0432
Org_type	Organization type (HMO, PPO)	0.10197
Mean_unobs	Mean unobserved heterogeneity in county	-0.0523
Star_diff	Star rating difference (own star market average)	0.0566
Share	Estimated market share	0.0043
Prev_action	Action taken in previous round	-0.0336

Optimization of the pseudo log likelihood function starts with the predicted choice probabilities using this first-stage model and an initial parameter guess for the θ parameters. Two optimization algorithms are used for minimizing the likelihood function:

- Nelder–Mead Simplex: A derivative-free method, robust for noisy or irregular likelihood surfaces
- BFGS: A quasi-Newton method that uses gradient information to converge faster on smooth surfaces.

Let $J^{(t)}$ denote the value at iteration t. The algorithm is said to converge when

$$|J^{(t)} - J^{(t-1)}| < \varepsilon,$$

where ε is a small tolerance level.

Convergence is monitored by tracking changes in the objective function. If the objective stops improving or stagnates, the algorithm switches from Nelder–Mead to BFGS to ensure efficient convergence. After each iteration, parameters are updated, and probabilities recomputed. The process continues until convergence, yielding the final estimates that maximize the likelihood.

The state variables used for the estimation of the value function are a contract's own star rating, average plan level unobserved heterogeneity from the demand model, star rating of the previous

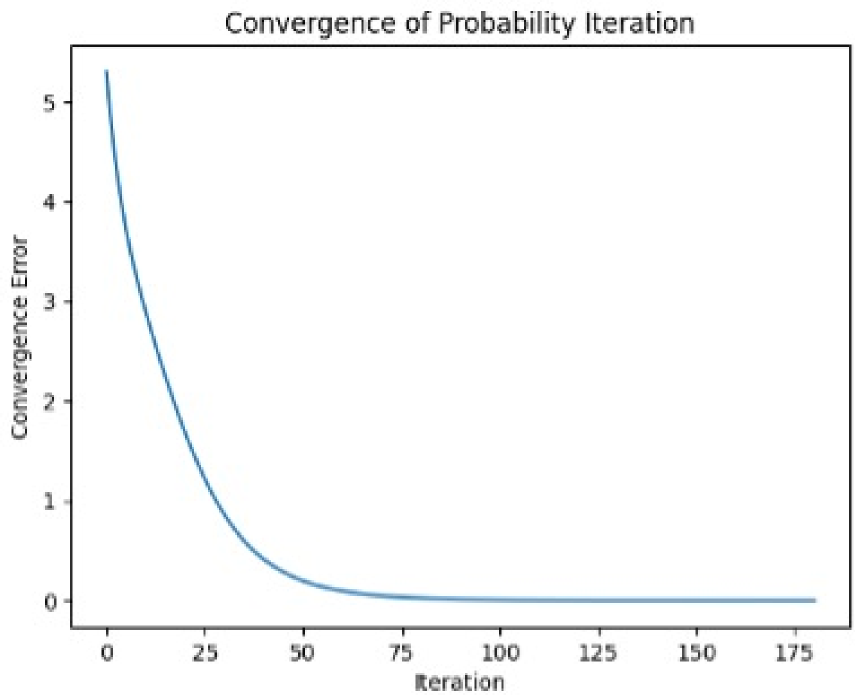

Figure A1: Figure showing monotonic decay of the error in predicted choice probabilities across iterations, indicating convergence.

period to account for lags in MA payment, market level characteristics such as the star rating of the competitors in terms of the mean ratings, and the unobserved mean quality of the competitors. A quadratic basis function is used for the estimation of the value function. The estimation routine converges in 196 iterations. Figure A1 shows the convergence of the errors in predicted choice probabilities over the number of iterations.

C Qualitative Research

Several case study organizations described the demonstration as beneficial in focusing on leadership on Star Ratings and prioritize improving their ratings because it tied funding to QI. MA organizations' Star Ratings influenced their perception of the MA QBP Demonstration: they generally saw the demonstration as an opportunity to improve quality and improve their competitive positions, but case study organizations with lower rated contracts were also concerned about the low performer icon and remaining in the MA program. The MA organizations also described how their competitors' Star Ratings and the alternatives available to beneficiaries in their markets factored into their organization's Star Rating strategies, thus illustrating how competition can mitigate or enhance the effects of QBP incentives and the stars program more broadly. Although most of the participating MA organizations reported that their QI activities and focus on Star Ratings preceded the announcement of the demonstration, several indicated that the demonstration provided added incentives for improving Star Ratings and in some cases reorganizing internally to better focus organizational resources on their Star Ratings.

The QBP demonstration helped define the internal business case for QI

Several organizations indicated that the MA QBP was beneficial in helping focus leadership on Star Ratings and their improvement because it tied funding to QI in the wake of MA payment reductions in the ACA. QI had historically been a cost center, but the opportunity to make up some of the ACA's MA payment reductions through QBPs is changing that. Organizations interviewed described a number of investments in analytics, provider and beneficiary outreach, and changes to organizational structure to focus on Star Ratings

MA organization's ratings and their competitors' ratings drive organizational Star Rating strategies

This illustrates how competition can mitigate or enhance the effects of QBP incentives and the Star Ratings program more broadly to attain higher Star Ratings. Representative from a national level organization said that in a market where all the competitors are 3-star plans, having a 4.5-star rating is good enough. But in markets where there are highquality contracts on the cusp of very high ratings, they feel acute pressure to achieve 5 stars so they can put that information on their products

Provider Network formation and provider incentives are an important pathway of improving ratings

Most experts in this industry agree that provider network formation is an important aspect of improving star ratings. Over half of MA plan's star ratings are based on physicians delivering appropriate services including providing screening tests, vaccines and managing chronic conditions. Such clinical measures cannot easily be improved by plans without significant cooperation from primary care physicians and other clinicians. Therefore, significant provider cooperation and buy-in are necessary for improving star ratings.

All MA insurers interviewed focused on star ratings when forming narrow networks, though different insurers took different approaches. One MA insurer that focuses on HMO products said that they narrow primary care networks because they believe primary care drives star ratings. The large national MA insurers said they generally form narrow networks around already high-performing physician groups or hospital systems that have proven track records on quality and utilization. All the health systems echoed the MA insurers' emphasis on star ratings, and the two health systems that partially or fully own an MA plan said their high star ratings were integral to their success in MA. All but one MA expert also agreed that star ratings are a crucial consideration for MA insurers when forming networks.

Increasingly using data and predictive analytics to support QI programs, with the QBP providing additional impetus for investment

CMS's Star Ratings initiatives collectively provided an impetus for several to invest in these analyses. One MA organization said its information technology (IT) investments were a direct result of the bonus payments. Others contracted with an analytics vendor (either Inovalon or Optum), although in several cases this was seen as the short-term solution until the MA organization had the resources to perform the analyses internally. Some organization's QI staff worked with actuarial staff to develop a graphic illustrating the impact of MA payment cuts and the revenues the organization might recoup through the QBPs. This information was presented to leadership to lobby for additional funding for QI efforts.

Summarizing: The MA organizations invest in quality improvement measures driven by the incentive of the quality bonus payments as they look at it as an extra source of revenue for the plans. The measures that MA organizations undergoes usually is a structural and organizational change which requires a fixed investment and might also increase the marginal cost of providing service(For example: if an MA organization is improving quality by choosing a new provider network, then such an organizational change will need a fixed investment and switching to a better provider network might also increase the marginal cost as indicated by my model of marginal cost) They do not however use it for marketing individual plans and decide on how much to invest based on the competition they face in terms of ratings. The more highly rated contracts appear to have more highly rated competitors, on average.