Designing Physician Incentives and the Cost-Quality Tradeoff: Evidence from Accountable Care Organizations

Kole Reddig*

October 2, 2024

Abstract

This paper estimates a structural model of Medicare provider expenditure to investigate the cost-quality tradeoff in health care and design contracts for a large physician incentive program. The setting involves Accountable Care Organizations (ACOs), which are groups of health care providers that receive incentive pay for spending below a cost target on shared patients. Estimation of the structural model and counterfactual simulations reveal a cost-quality tradeoff affects less than one third of ACOs. Penalties for overspending have little effect on simulated quality of care scores; however, after accounting for voluntary participation in the incentive program, expected program savings is larger without penalties.

1 Introduction

In the United States health care sector, public and private insurers often implement physician incentive programs and pay-for-performance initiatives to control the cost of care. Designing payment contracts for these programs requires facing a fundamental challenge: physicians may decrease the quality of care they provide in order to reduce cost. In this paper, I estimate a structural model of Medicare expenditure to identify the extent to which quality of care decreases when providers are incentivized to decrease expenditure. I use the structural model to conduct counterfactual analyses that highlight the role of the cost-quality tradeoff in designing physician incentive programs.

The setting of this study is the Medicare Shared Savings Program (MSSP), a large incentive program that involves 10.8 million Medicare beneficiaries and over $2.5 billion in provider incentives each year.[1] The MSSP gives incentive pay to Accountable Care Organizations (ACOs), which are joint ventures of physicians, group practices, and hospitals that form to coordinate care of their shared patients. An ACO earns incentive pay through the MSSP if its members collectively reduce expenditure on health services. Because providers might decrease the quality of care they provide in order to reduce expenditure, both tasks (monetary savings and quality of care) determine ACO payment.

*John Chambers School of Business and Economics, West Virginia University. Email: kole.reddig@mail.wvu.edu. An earlier version of this paper titled "Designing Contracts for Multitasking Groups: A Structural Model of Accountable Care Organizations" was a chapter of my doctoral thesis at Carnegie Mellon University. I would like to thank Martin Gaynor and Maryam Saeedi for support throughout this project. I would also like to thank Chris Cronin, Dennis Epple, Akshaya Jha, Rebecca Lessem, Albert Ma, Hakkı Özdenören, and Amanda Starc, as well as participants at several seminars and conferences for their feedback.

[1]Source: The Centers for Medicare and Medicaid Services. "Shared Savings Program Fast Facts." `https://www.cms.gov/files/document/2024-shared-savings-program-fast-facts.pdf`.

What role does the cost-quality trade off play in determining optimal physician incentives? The answer to this question is central the design of the MSSP, and furthermore will inform incentive and contract design throughout the health care sector. I answer this question by building and estimating a structural model of Medicare providers. In the model, Medicare providers with a shared patient strategically choose expenditures based on their own preferences and the incentive pay offered by the MSSP. The expenditure choices form a Nash equilibrium that describes the expenditure of a beneficiary assigned to an ACO.

In counterfactual analyses, I solve for the contract between ACOs and Medicare that maximizes the quality-weighted cost savings of providers in the MSSP.[2] In this principal agent problem, Medicare is restricted by federal regulation in such a way that it can only choose contracts that pay groups of providers (the ACO), and not individual providers (the ACO participants). Federal regulation also restricts the specific ways that Medicare can define incentive pay for ACOs. First, Medicare can set the generosity of the contract for a given level of quality of care and expenditure, and second, Medicare can require ACOs to make penalty payments to Medicare if their expenditure exceeds its target. Knowing the magnitude of the tradeoff between cost reduction and quality of care is critical for choosing these contract parameters. Depending on the structural relationship between cost reduction and quality of care in health care, enforcing penalties with low-powered contracts has the potential to greatly diminish quality of care. Moreover, because participation in the MSSP is voluntary, contracts must not penalize overspending so much that ACOs exit the program.

My research design exploits observed variation in the well-defined contracts between Medicare and ACOs in the MSSP to identify structural parameters. In the MSSP, an ACO is assigned a benchmark expenditure for the health care services provided to its participants' patients. If a year's actual Medicare expenditure on those beneficiaries is less than the benchmark amount, an ACO earns a portion of the difference, adjusted by a quality score, as incentive pay (hence "sharing savings" with Medicare). These contracts are designed by Medicare and are public information, so I observe cross-ACO variation in the marginal dollar of group incentive pay for a given level of cost savings and quality of care. Under an equilibrium assumption, this identifies parameters describing provider altruism and responsiveness to incentive pay.

I use observed ACO quality scores to infer the structural relationship between quality of care and health expenditure. In short, I assume quality of care is determined by beneficiary health, and beneficiary health is determined by a non-monotone health production function with a global maximum at ideal expenditure. Within-ACO variation in quality scores and expenditure relative to benchmark expenditure are used to identify the ideal expenditures of the representative beneficiaries of ACOs. Hence, I estimate the presence of a cost-quality tradeoff at ACO-level, which is indicated by actual expenditure being less than ideal expenditure, and on the domain of the health production function that is increasing in expenditure.

The conclusions of this work apply to both economics and policy audiences. I estimate a small structural tradeoff between Medicare savings and quality of care. In the context of the MSSP, I provide evidence through counterfactual simulations that the savings-quality tradeoff is neither large enough nor sufficiently

[2]For comparison, I also solve for contracts that are optimal by other definitions, including per-beneficiary objectives and objectives accounting for voluntary participation in the MSSP.

common among ACOs to create significant quality decreases when ACOs are incentivzed to reduce cost. I find imposing penalties on ACOs for having expenditure larger than benchmark expenditure, compared to contracts without penalties, increases average program savings by $144 per beneficiary (0.3 standard deviations). These "two-sided" contracts actually increase quality of care (as measured by lower ACO quality scores) by driving expenditure from above the ideal level to below the ideal level. However, due to the shape of the health production function, the change in quality score is very small, at less than 0.03 standard deviations, or 0.3% of mean quality score.

Under "one-sided" contracts, where ACOs are not penalized for overspending, simulations show per-beneficiary savings is maximized at $47 with the optimal contract. The optimal sharing rate needed to achieve this savings is 74%, which is on the high end of contracts offered by the MSSP. ACO quality scores are maximized at sharing rate of 53%.

Two-sided contracts, where ACOs are penalized for having an expenditure larger than benchmark expenditure, are able to incentivize ACOs to have a lower expenditures at lower sharing rate. Because of the penalties, simulations show that two-sided contracts maximize total program savings at a slightly weaker 64% sharing rate. ACO quality scores are maximized at a much lower sharing rate of 12%. Weighting program savings by quality scores decreases the optimal sharing rates by only a few percentage points.

In order to account for voluntary participation in the MSSP, I estimate a model of ACO exit, and compute objectives weighted by the probability ACOs remain in the MSSP. These objectives imply that very high sharing rates—above 90%—are optimal for both one-sided and two-sided contracts.

Ultimately, the results show that there is little Medicare can do, within the scope of the contract parameters analyzed, to meaningfully reduce costs. This paper empirically verifies the premonitions of Frandsen & Rebitzer (2015) that organizational incentives within ACOs impose inefficiency larger than the efficiency introduced by coordination. This paper goes a step further: the contract that maximizes (non-weighted) per beneficiary program savings achieves a savings of $191 per beneficiary. Medicare Parts A and B spending per enrolled beneficiary was $11,080 in 2021, and increased by an average of $166 every year since 2013.[3] This means applying the *ceiling* level of ACO performance (ignoring incentive program participation and quality of care) to the entirety of FFS Medicare would yield only a 1.7% cost savings. If the rate of increase in spending is not also decreased, spending would return to the same level in less than two years. And, if the population of health care providers in ACOs is representative of providers the in the general health care market, this paper shows both public and private insurers should expect small efficiency gains from shared savings programs.

Related Literature. This paper contributes to economics literature concerning health care provider payment systems and provider behavior in organizations (Gaynor, Rebitzer, & Taylor, 2004; Encinosa, Gaynor, & Rebitzer, 2007; Choné & Ma, 2011; Rebitzer & Votruba, 2011; Ho & Pakes, 2014; Frandsen & Rebitzer, 2015; Grassi & Ma, 2016; Frandsen, Powell, & Rebitzer, 2019). More generally, this paper aligns with

[3]Source: Kaiser Family Foundation. "Medicare Spending per Beneficiary." `https://tinyurl.com/a56w8pbp`.

the literature that studies the supply-side of health care, and examines the incentives faced and decisions made by physicians, hospitals, and insurers (Gaynor, 2006; Chandra, Cutler, & Song, 2011; Gaynor, Ho, & Town, 2015; Gowrisankaran, Joiner, & Léger, 2017; Einav, Finkelstein, & Mahoney, 2018; Eliason, Grieco, McDevitt, & Roberts, 2018; Hackmann, 2019; Gaynor, Mehta, & Richards-Shubik, 2023).

Few studies in economics have discussed ACOs directly. Frandsen & Rebitzer (2015) calibrate a model of ACO performance to examine the size-variance tradeoff in group payment mechanisms like the MSSP, and they argue that ACOs will be unable to self-finance. That is, there is no contract with strong enough incentives to overcome the incentive to free-ride among a group of physicians. The authors conclude with a skeptical look at the MSSP, and mention the untenability of integrated organizations in the now very fractured US health care market. Frech et al. (2015) study county-level entry of private and public ACOs. The authors find small markets generally discourage ACO entry, and that public ACO entry is largely predicted by higher Medicare spending, higher population, and lower physician site concentration. Frandsen et al. (2019) discuss the MSSP's impact on health care in the United States in the context of common agency, where several payers motivate the same agent to improve care delivery and integration. The authors find that unique equilibrium contracts from payers are lower powered in the presence of shared savings payments, and ACO entry can possibly inspire other shared savings contracts in the private sector if they do not already exist.

Two studies evaluate actual or simulated ACO performance with results comparable to those of this paper. Aswani, Shen, & Siddiq (2019) study how to design MSSP ACO contracts. The authors focus on asymmetric information between Medicare and ACOs, and write contracts such that ACO payment is a function of underlying ACO type (benchmark expenditure per beneficiary). Unlike this paper, Aswani et al. (2019) do not consider the cost-quality tradeoff in health care. Their results, in simulations that are comparable, are similar. For example, optimal sharing rates for two-sided contracts in both papers are found to be about 50-60%. Wilson et al. (2020) is short literature review of studies addressing ACO performance. Projections of MSSP ACO cost savings from the reviewed studies are on the same magnitude of this paper, ranging between -\$107 to \$857 per beneficiary.

This paper continues as follows: Section 2 gives a detailed overview of the MSSP and ACOs, including descriptive ACO statistics. I outline my model of Medicare expenditure in Section 3. I describe identification and estimation of model primitives in Section 4, and estimation results and model fit are in Section 5. I present counterfactual analyses, including computation of savings-maximizing contracts between ACOs and Medicare, in Section 6, and Section 7 concludes.

2 Background and Data

The MSSP, a part of the Patient Protection and Affordable Care Act of 2010 (ACA), is a policy response to increasing health care costs in the United States. The premise of the program is that the United States is inefficient at providing health care because care delivery is *fragmented*. That is, unique to the United States, patients tend to see several distinct providers that belong to separate businesses with little incentive to coordinate care. Patients therefore receive haphazard and/or redundant care, implying increased utilization, cost, and risk of adverse health outcomes. A recent empirical analysis by Agha, Frandsen, & Rebitzer (2019) examined the fragmentation of health care in the United States. Using Medicare claims data, the authors show utilization increases in areas with more fragmented care, and patients substitute low cost services for high cost services in areas with high fragmentation.

The MSSP gives providers financial motivation to integrate care delivery. To overcome institutional boundaries to care integration, the program explicitly evaluates and pays Medicare providers based on group performance. By law, a provider must start or join an existing Accountable Care Organization, or ACO, to earn incentive pay from the MSSP. Formally, ACOs are joint ventures of health care providers; nearly any Medicare provider—including individual physicians, group practices, and large hospital systems—can start or participate in an ACO. Applications to join the MSSP are submitted by ACOs to the Centers for Medicare and Medicaid Services (CMS), and CMS approves the applications of ACOs that meet a pre-defined set of criteria. In particular, an ACO must appoint a governing board to oversee the clinical and administrative aspects of operation, and it must show the presence of formal contracts between itself and its member participants (including the distribution of any earned incentive pay). If approved for participation in the MSSP, the ACO then enters into a five year agreement with CMS.[4] Each year of the agreement is referred to as a "performance year."

To explain how ACOs earn incentive pay through the MSSP, I will first describe how Medicare beneficiaries are assigned to ACOs. Then, I will describe how an ACO's spending target—called "benchmark expenditure"—is determined. Finally, I will describe the determination of incentive pay, or "shared savings," given to ACOs.

2.1 Beneficiary Assignment

To be assigned to an ACO in a given performance year, a Medicare beneficiary must be enrolled in both Parts A and B of Medicare (also known as Traditional or fee-for-service (FFS) Medicare). The beneficiary must not be enrolled in any other Medicare group health plan, such as Medicare Advantage (MA) or the Program of All-Inclusive Care for the Elderly (PACE). They also must not be assigned to providers participating in a different Medicare incentive program.

Beneficiaries meeting this criteria are assigned to ACOs based on billed expenditure (specifically, Medicare-allowed charges) on primary care services in the performance year. First, a consider a Medicare beneficiary that has at least one primary care service furnished by a primary care practitioner.[5] The beneficiary is

[4]Before July 2019, agreements lasted three years.

[5]CMS defines primary care services based on CPT code: these include the standard office visit codes 99201-99205 and 99211-99215, as well as codes for nursing facility care, home health care services, and specific outpatient hospital services. Primary care practitioners are defined as those with specialty codes corresponding to general practice, family practice, internal medicine, pediatric medicine, geriatric medicine, nurse practitioner, clinical nurse specialist, and physician assistant. See "Shared Savings and

assigned to the ACO of the provider(s) by which the plurality of primary care services (determined by expenditure) are furnished. If the beneficiary receives more primary care services from primary care practitioners not in an ACO, then the beneficiary is not assigned to an ACO. If the beneficiary does not receive any primary care services from a primary care practitioner, then they are assigned to the ACO of the specialist practitioner(s) by which the plurality of primary care services are furnished. Again, if more primary care services are furnished by specialists not participating in an ACO, then the beneficiary is not assigned to an ACO. In short, beneficiary assignment is determined hierarchically, such that primary care expenditure by primary care practitioners first determines assignment; if primary care is not provided by primary care practitioners, then primary care expenditure by specialists determine assignment.

Because beneficiary assignment is not determined until the end of a performance year, ACOs do not know the exact list of beneficiaries that they are assigned until after the year has ended and expenditure decisions have been made. CMS provides all ACOs with prospective beneficiary assignment lists that are updated quarterly throughout the performance year to keep providers up-to-date regarding beneficiaries that will likely be assigned to their ACO. Lee, Polsky, Fitzsimmons, & Werner (2020) examine whether the patient demographics of a provider changes after joining an ACO, and find no evidence of selection or cream-skimming. CMS has a rule requiring ACOs to have at least 5,000 assigned beneficiaries; however, this rule was seldom enforced, and ACOs were granted exceptions with the expectation they grow in future years. Accordingly, this rule will be removed, and starting 2025, there is no minimum beneficiary requirement.

2.2 Determining Benchmark Expenditure

An ACO's benchmark expenditure is based on the Medicare expenditure on beneficiaries that would have been assigned to the ACO in the three years prior to the agreement period; these are called the "historical benchmark years." First, annualized expenditures are computed for every beneficiary assigned to the ACO in each historical benchmark year. The per beneficiary expenditures of the first two years are trended forward to match the third historical benchmark year based on the national growth rates in FFS Medicare expenditure, where the national growth rate applied depends on the enrollment type of the beneficiary: Aged, Non-Dual (typical Medicare beneficiaries qualifying through being age 65 or older); Aged, Dual (Medicare beneficiaries qualifying through being age 65 or older that are also enrolled in Medicaid); Disabled (Medicare beneficiaries qualifying through Social Security benefits); and ESRD (Medicare beneficiaries qualifying through diagnosis of kidney-failure from the enactment of the Social Security Amendments of 1972).[6] Next, per beneficiary expenditures in each historical benchmark year are risk-adjusted to the third benchmark year using CMS Hierarchical Condition Category (CMS-HCC) risk scores.[7] These per-beneficiary expenditures are then

Losses and Assignment Methodology: Specifications." `https://www.cms.gov/Medicare/Medicare-Fee-for-Service-Payment/sharedsavingsprogram/Downloads/Shared-Savings-Losses-Assignment-Spec-V6.pdf`.

[6]Enrollment types are hierarchical; if a beneficiary classifies as more than one enrollment type, then the assignment progresses ESRD to Disabled to Aged, Dual to Aged, Non-Dual.

[7]CMS-HCC risk scores are an index representing the ratio of expected Medicare expenditure of a beneficiary to the Medicare population. They are computed using demographic information (age and gender) and previously assigned diagnosis codes.

averaged across assigned beneficiaries within enrollment types and historical benchmark years, yielding the average per beneficiary expenditure of assigned beneficiaries by enrollment type and year.

The risk-adjusted and trended historical benchmark year average expenditures are combined into a single historical benchmark expenditure with a weighted average that puts 10% weight on the first benchmark year, 30% weight on the second benchmark year, and 60% weight on the third benchmark year, where the third benchmark year is the year immediately preceding the first performance year of the agreement period. This yields a weighted average annual per beneficiary expenditure for each enrollment type. To aggregate over enrollment types, these average expenditures are weighted by the share of assigned beneficiaries in the third historical benchmark year by enrollment typed and summed across enrollment types. This yields the historical benchmark per beneficiary expenditure of the ACO.

To obtain the ultimate spending target, the ACO's benchmark expenditure, for ACOs in their first agreement period, CMS updates the historical benchmark expenditure for changes in ACO participants and trends it forward to the performance year in question based on national growth in Medicare FFS spending by enrollment type and shares of assigned beneficiaries in the performance year by enrollment type.

If ACOs are in their second agreement period, the historical benchmark per beneficiary expenditure is "rebased" because the historical benchmark years of the second agreement period are the years of the first agreement period; ACOs would otherwise be vulnerable to a "rachet" effect under which low expenditure would disadvantage the ACO in later years with a lower benchmark expenditure. Accordingly, the historical benchmark is rebased by taking a weighted average of the historical benchmark and a historical benchmark computed from spending in the ACO's regional service area; 25% weight is placed on the regional service area's benchmark for ACOs with higher spending than the region, and 35% weight is placed if the ACO has lower spending than the region.[8]

2.3 Shared Savings and Losses

Payment of an ACO depends on the performance year's Medicare expenditure on beneficiaries assigned to the ACO, a quality of care score, and the contract the ACO has with Medicare. Payment is determined by the difference between benchmark expenditure and the realized total Medicare expenditure on beneficiaries assigned to the ACO. Importantly, it is not just the expenditure of providers participating in an ACO that determines its expenditure; rather, it is the total Medicare expenditure on beneficiaries that are assigned to the ACO that defines ACO expenditure, whether or not the spending was billed by a provider in the ACO.

A composite quality score between 0 and 1 also plays a role in determining payment. An ACO's overall quality score is determined by the combination of 30-40 sub-measures of care quality. These sub-measures fall into the domains of "Patient/Caregiver Experience," "Care Coordination/Patient Safety," "Preventative Health," and "At-Risk Population." Some sub-measures are survey responses (e.g., "ACO2: How Well Your Doctors Communicate"), while others are computed from Medicare Claims and aggregated to the ACO-level

[8] The regional service area of an ACO is defined as all the counties of residence of the beneficiaries assigned to an ACO.

(e.g., "ACO21: Proportion of Adults who had blood pressure screened in past 2 years").[9]

If the ACO's savings rate, defined as $\frac{\text{Benchmark Expenditure}-\text{Expenditure}}{\text{Benchmark Expenditure}}$, exceeds a predetermined minimum called the Minimum Savings Rate (MSR), and if the ACO meets minimum quality of care standards, it earns and distributes to its members the amount

$$\text{Shared Savings} = \text{Sharing Rate} \cdot \text{Quality Score} \cdot (\text{Benchmark Expenditure} - \text{Expenditure}) \ , \tag{1}$$

where "Sharing Rate," a number between 0 and 1, is determined by the type of contract the ACO has with Medicare. The MSR is determined by the number of beneficiaries assigned to an ACO, ranging from 2% for ACOs with 60,000 or more assigned beneficiaries to 12.2% for ACOs with less than 500 assigned beneficiaries.

The overwhelming contract choice of ACOs from 2014 to 2018—the period of time considered by the empirical analysis— was "Track 1," which has a sharing rate of 50%. Under this contract, if a hypothetical ACO with a benchmark expenditure of $186 million and an MSR of 0.02 had an expenditure of $180 million with a quality score of 0.90, it would earn

$$0.5 \cdot 0.9 \cdot (\$186 \text{ million} - \$180 \text{ million}) = \$2.7 \text{ million} \tag{2}$$

in shared savings. Its savings rate is $(186 - 180)/186 = 0.03$, so the MSR is exceeded. Though paying a subsidy, Medicare saves money as well: on net, this ACO contributed a $3.3 million decrease in Medicare expenditure, as it was paid $2.7 million for saving $6 million.

While uncommon in the first few years of the MSSP, some contracts also penalize ACOs for having expenditure *larger* than benchmark expenditure. These are called "two-sided" contracts. "Track 2" and "Track 3" ACOs have two-sided contracts, and pay Medicare "Shared Losses" if their savings rate is less than the Maximum Loss Rate (MLR, equal to the negative of the MSR):

$$\text{Shared Losses} = (1 - \text{Sharing Rate} \cdot \text{Quality Score}) \cdot (\text{Benchmark Expenditure} - \text{Expenditure}) \ . \tag{3}$$

Note that the value of this equation is negative when expenditure exceeds benchmark expenditure, and that the value can be diminished by having a higher sharing rate or quality score. During the early years of the MSSP, two-sided contracts came with higher sharing rates: 60% for Track 2 and 75% for Track 3 ACOs. In 2018, Track 1+ was introduced to the MSSP. This contract offers ACOs up to 50% of savings as incentive pay, but requires ACOs to pay 30% of losses to Medicare if expenditure is much larger than benchmark expenditure and savings is below the minimum loss rate.

For example, if a Track 3 ACO with a benchmark expenditure of $30 million and an MLR of -0.03 had

[9] See `https://go.cms.gov/2xHy7Uo` for a full list of ACO quality scores for every performance year.

an expenditure of $33 million and quality score of 0.80, then they would have shared losses

$$(1 - 0.75 \cdot 0.8) \cdot (\$30 \text{ million} - \$33 \text{ million}) = -\$1.2 \text{ million}\,. \tag{4}$$

This hypothetical ACO would offset the $3 million in expenditure above the benchmark with a $1.2 million payment to Medicare.

The first six years (2013-2018) of the MSSP produced modest decreases in Medicare expenditure (McWilliams et al., 2018). In an attempt to improve ACO performance, CMS made several changes to the MSSP with its final rule named "Pathways to Success" (or "Pathways"). Changes in Pathways pertinent to this paper regard the contracts between ACOs and Medicare. Tracks 1, 1+, 2, and 3 are replaced with two Tracks: "Basic" and "Enhanced." Under the Basic Track, there are five levels, "A" through "E." Under levels A and B, ACOs earn up to 40% of savings and do not pay shared losses if expenditure exceeds benchmark expenditure. Under levels C, D, and E, ACOs earn up to 50% of savings and pay an increasing amount of shared losses if expenditure exceeds benchmark expenditure. An ACO is automatically advanced one level (e.g., from level A to B) after each performance year. The Enhanced Track is equivalent to Track 3. Various other changes were made to the MSSP in Pathways, including beneficiary assignment methodology, benchmark calculation, and assigning new ACO classifications ("low-renvenue" and "experienced") that impact the payment contracts available to an ACO.

2.4 Summary Statistics

In this paper, counterfactual predictions consider two dimensions of contracts: the fraction of savings shared with an ACO and the presence of downside risk. These dimensions broadly account for all previous (Tracks 1, 1+, 2, and 3) and current (Basic and Enhanced Tracks) contract options. Accordingly, in this analysis I use data from all ACOs of all Tracks from 2014 to 2018. ACOs in the first year of the program, 2013, are all assigned a quality score of 1, so I do not use them in this analysis. I stop at 2018 due to the changes in program structure introduced with Pathways in 2019.

The data I use for all empirical analysis is from MSSP ACO Public Use Files, MSSP Participant Lists, and MSSP ACO Performance Year Results Public Use Files. In short, the public data consists of ACO expenditures, benchmark expenditures, quality scores (along with every quality sub-measure), various assigned beneficiary demographics, and various provider statistics. Little public information is available on the characteristics of specific providers. Table 1 displays statistics describing ACOs. There is substantial heterogeneity in the number of providers that join an ACO—some large hospitals are able to form an ACO independently, and others are joint ventures of hundreds of providers.

ACO benchmark expenditures and realized expenditures are large: the means of total expenditure and benchmark expenditures on assigned beneficiaries are approximately $190 million, with several ACOs having expenditure over $1 billion. From 2014 to 2018, ACOs saved money on average. However, less than one

Table 1: **Summary Statistics**

	Mean	S.D.	Min.	Median	Max.
Number of Assigned Ben.	18163	18052	152	11979	149633
Prop. of Assigned Ben: under 65	0.169	0.073	0.038	0.155	0.662
Prop. of Assigned Ben: over 85	0.122	0.032	0.028	0.119	0.309
Prop. of Assigned Ben: male	0.428	0.020	0.356	0.428	0.575
Prop. of Assigned Ben: nonwhite	0.165	0.146	0.015	0.123	0.951
Prop. of Assigned Ben: Aged, Non-Dual	0.772	0.131	0.122	0.806	0.967
Prop. of Assigned Ben: Aged, Dual	0.081	0.090	0.003	0.056	0.808
Prop. of Assigned Ben: Disabled	0.138	0.070	0.023	0.122	0.645
Prop. of Assigned Ben: ESRD	0.009	0.006	0.000	0.008	0.045
Total Providers	737	1073	1	340	11783
Provider per Assigned Ben.	0.041	0.051	0.000	0.027	0.673
$\mathbb{1}\{\text{Exited Next Year}\}$	0.122	0.327	0.000	0.000	1.000
Age	2.711	1.452	1.000	2.000	7.000
Sharing Rate	0.511	0.049	0.500	0.500	0.750
$\mathbb{1}\{\text{Two-Sided}\}$	0.074	0.262	0.000	0.000	1.000
Per Ben. Expenditure ($ thousand)	10.776	2.200	4.692	10.372	27.472
Expenditure ($ million)	186.859	184.506	1.413	122.975	1632.292
Per Ben. Benchmark Expenditure ($ thousand)	10.920	2.288	5.115	10.488	29.223
Benchmark Expenditure ($ million)	188.801	186.677	1.476	124.919	170.496
Quality Score	0.826	0.089	0.107	0.845	0.976
$\mathbb{1}\{\text{Earned Incentive Pay}\}$	0.328	0.470	0.000	0.000	1.000
Earned Shared Savings/Losses ($ million)	1.478	3.601	-5.670	0.000	49.968
Earned Shared Savings/Losses per Provider ($ million)	0.012	0.109	-0.024	0.000	3.458

Notes: $N = 2180$. The sample includes all MSSP ACOs from 2014 to 2018, with 740 ACOs in an unbalanced panel.

third of ACOs had a savings rate at least as large as their minimum savings rate, meaning most ACOs do not actually earn incentive pay during this period. Average earned incentive pay is $1.5 million per ACO. Per provider, average earned incentive pay is $12,000.

3 Model of Health Care Expenditure and Patient Health

This section lays out a model of expenditure decisions made by Medicare providers. Decisions are made in a simultaneous move game among the providers that care for a specific Medicare beneficiary—that is, among the *care team* of a beneficiary. Each provider chooses expenditure on a beneficiary to maximize their individual payoff, which is the sum of profit (revenue from Medicare reimbursement minus the cost of providing care) and the altruistic utility derived from the health of the beneficiary. Health is determined by a health production function that depends on the total Medicare expenditure on a beneficiary across all providers in their care team relative to the ideal level of expenditure that maximizes that beneficiaries health.

After solving for Nash equilibrium Medicare beneficiary expenditure, I extend the model to include incentive pay for cost-reduction and quality improvement from the MSSP when a Medicare beneficiary is assigned to an ACO. For a Medicare beneficiary assigned to an ACO, providers on their care team (which may or may not be members of the ACO) simultaneously choose expenditure on the beneficiary to maximize their individual payoffs, which is the sum of profit, altruistic utility, and incentive pay from the MSSP. This extension of the model ultimately yields a function for Medicare expenditure that depends MSSP contract parameters and ACO performance (ACO savings and quality score) in addition to provider characteristics that otherwise determine the provision of care. Crucially, I assume that the quality score of an ACO is a measure of the average health production of beneficiaries assigned to the ACO.

The model is static, and I take the populations of Medicare providers and Medicare beneficiaries as given. Formally, providers are indexed by $i \in \mathcal{I}$, and beneficiaries are indexed by $a \in \mathcal{A}$. I also take the care teams of beneficiaries as given; these are denoted with the sets $\mathcal{I}_a \subset \mathcal{I}$. Finally, ACOs indexed $j \in \mathcal{J}$ are taken as given, as is ACO beneficiary assignment and provider participation in ACOs. I denote the set of Medicare beneficiaries assigned to ACO j with $\mathcal{A}_j \subset \mathcal{A}$, and $\mathcal{I}_j \subset \mathcal{I}$ is the set of Medicare providers that are members of ACO j.[10] Because beneficiaries and providers are members of at most one ACO, $\mathcal{A}_j \cap \mathcal{A}_{j'} = \varnothing$ and $\mathcal{I}_j \cap \mathcal{I}_{j'} = \varnothing$ for $j \neq j' \in \mathcal{J}$. However, it not uncommon for a provider to be on the care team of a beneficiary that is assigned to an ACO without themselves being a member of the same ACO: formally, there exists $i \in \mathcal{I}_a$ with $a \in \mathcal{A}_j$ and yet $i \notin \mathcal{I}_j$.

3.1 Health Production and the Cost-Quality Tradeoff

The first component of the model to define is the *health production function*, which in this context is a mapping of the total Medicare expenditure on a beneficiary to the beneficiaries health. Health production functions are common in both theoretical and empirical health economics research. For example, the seminal work of Grossman (1972) models the consumption of health services and the "healthy days" of an individual as a function of their depreciating stock of health. More recently, Glied & Hong (2018) assumes a monotone relationship between patient utility (i.e. health) and health care consumption to study demand spillover from insurance expansions. Ho & Lee (2023) (based on the model of Einav, Finkelstein, Ryan, Schrimpf, & Cullen (2013)) use a non-monotone (i.e. quadratic) health production function to study moral hazard in

[10] Participation of Medicare providers in ACOs, and more generally the formation of ACOs, is deliberately not made endogenous in the model. This simplification is necessary, as a formal model of endogenous ACO formation that addresses optimal group size, provider networks, and multiplicity of equilibria is prohibitively complicated and not within the scope of this paper given the available data. I discuss the consequences of this simplification, including the steps necessary for identification of model parameters given endogenous ACO formation, in Section 4. I account for endogenous ACO exit when conducting counterfactual simulations by estimating a simple model of ACO exit; this is discussed in Sections 4 and 6.

the context of insurance plan choice. Gaynor et al. (2004) and Gaynor et al. (2023) likewise apply quadratic health production functions in models of altruistic provider choice of care.

The non-monotonicity of a health production function is fundamental to both the cost-quality trade off in health care and the implementation of physician incentive programs such as the MSSP. For regions of the health production function that are increasing in health expenditure, there is a cost-quality trade off, as increases in health expenditure are required for increases in health and thus quality of care. If expenditure is at a level such that health is decreasing, then cost-savings and quality of care are complementary. Whether an incentive program for cost-reduction in health care carries with it a risk for decreasing quality of care depends on whether expenditure on a beneficiary is at an increasing or decreasing part of the health production function.

To model patient health and the quality of care of an ACO, I follow the aforementioned studies and specify a health production function that is quadratic in expenditure. The health of Medicare beneficiary $a \in \mathcal{A}$ is given by

$$h\left(\tilde{e}_a - e_a\right) = -\frac{1}{2}\left(\tilde{e}_a - e_a\right)^2 \,, \tag{5}$$

where $e_a > 0$ is the total Medicare expenditure on health services for beneficiary a, and $\tilde{e}_a > 0$ is the ideal health expenditure of a; clearly $\tilde{e}_a \equiv \arg\max_{e_a > 0} h\left(\tilde{e}_a - e_a\right)$. Given the above health production function, it is straightforward to derive a definition for a cost-quality tradeoff in health care. Beneficiary a faces a cost-quality tradeoff if

$$h'\left(\tilde{e}_a - e_a\right) > 0 \iff \tilde{e}_a > e_a \,. \tag{6}$$

In other words, if health expenditure on e_a is less than ideal expenditure $\tilde{e}_a$, then there is a cost-quality tradeoff for beneficiary a. Policies like the MSSP can induce cost-savings without decreases in quality of care if there is initially overprovision of care, such that $\tilde{e}_a < e_a$. This is indeed the presumption of shared savings programs, which acknowledge reimbursements from insurers for health services may induce health expenditure that is larger than the ideal level.

3.2 Provider Expenditure Choice without Incentive Pay

The second component of the model is the game played by Medicare providers when choosing expenditure on a beneficiary. This game determines the equilibrium level of Medicare expenditure on a beneficiary, hence determining whether health expenditure is larger or smaller than ideal expenditure and whether the patient faces a cost-quality tradeoff.

Beneficiary a receives services from providers on their care team $i \in \mathcal{I}_a$. Accordingly, the Medicare expenditure on a is equal to the sum of expenditures billed by each member of their care team, where

$$e_a = \sum_{i \in \mathcal{I}_a} e_{ia} \,, \tag{7}$$

and $e_{ia} > 0$ is the Medicare expenditure of provider i on beneficiary a. The profit earned by provider i billing Medicare the amount e_{ia} for services provided to beneficiary a is given by the function

$$\pi_i \left(e_{ia} \right) = e_{ia} - \frac{e_{ia}^2}{2\gamma_i} \,, \tag{8}$$

where $\gamma_i > 0$. This functional form is straightforward to interpret: provider revenue is equal to reimbursements from Medicare for services provided to a, which is simply e_{ia}. The parameter γ_i is the amount of reimbursements on a that maximizes i's profit.[11] The parameter γ_i varies over providers due to factors on both the supply and demand side. For example, a capacity-constrained provider that is on the care team of many individuals has a relatively small γ_i, whereas a provider in a low-demand elasticity market has a large γ_i.

To decide expenditure on patient a, a provider maximizes their profit plus altruistic utility derived from the health of a. Providers $i \in \mathcal{I}_a$ solve

$$\max_{e_{ia} > 0} \pi_i \left(e_{ia} \right) + \alpha_i h(\tilde{e}_a - e_{ia} - e_{-ia}) \,, \tag{9}$$

where $\alpha_i > 0$ parameterizes physician i's level of altruism and $e_{-ia} = \sum_{i' \neq i} e_{i'a}$ is the expenditure on a billed by other providers. Taking the expenditure decisions of other providers as given, the first order condition for optimal expenditure e_{ia}^* is

$$\gamma_i - e_{ia}^* + \alpha_i \gamma_i \left(\tilde{e}_a - e_{ia}^* - e_{-ia} \right) = 0 \,. \tag{10}$$

Following previous literature on team production in health care (Gaynor et al., 2004; Frandsen & Rebitzer, 2015), I find the pure strategy Nash equilibrium expenditure of providers.[12] Accordingly, given Equation 10 holds for all $i \in \mathcal{I}_a$, the Nash equilibrium expenditure of provider i on a is

$$e_{ia}^* = \gamma_i + \frac{\alpha_i \gamma_i}{1 + \sum_{\ell \in \mathcal{I}_a} \alpha_\ell \gamma_\ell} \left(\tilde{e}_a - \sum_{\ell \in \mathcal{I}_a} \gamma_\ell \right) \,. \tag{11}$$

Equilibrium total expenditure on beneficiary a is therefore

$$e_a^* = \sum_{i \in \mathcal{I}_a} e_{ia}^* = \frac{\sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i}{1 + \sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i} \cdot \tilde{e}_a + \frac{\sum_{i \in \mathcal{I}_a} \gamma_i}{1 + \sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i} \,. \tag{12}$$

[11] This is because $\gamma_i \equiv \arg\max_{e_{ia} > 0} \pi_i \left(e_{ia} \right)$.
[12] Proofs regarding the existence and uniqueness of Nash equilibrium are in Appendix A.

Patient health in equilibrium, and likewise the equilibrium cost-quality trade off faced by beneficiary a, is given by the difference between ideal expenditure and equilibrium expenditure:

$$\tilde{e}_a - e_a^* = \frac{\tilde{e}_a - \sum_{i \in \mathcal{I}_a} \gamma_i}{1 + \sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i} \,. \tag{13}$$

This Nash equilibrium is straightforward to interpret. First, note that because provider provider profit is in dollar units, γ_i must also be in dollar units, and α_i is in dollar^{-1} units. Thus, the term $\alpha_i \gamma_i$ is unit-less and describes the tendency of provider i to provide an expenditure at the ideal level. This tendency depends on α_i due to provider altruism, and on γ_i from the curvature of i's profit function. Specifically, a larger value of profit maximizing expenditure γ_i implies the provider's payoff derived from altruism is larger relative to cost, so the effect of α_i on expenditure is amplified by γ_i.

Equation 11 shows that provider i's expenditure on a is equal to the profit maximizing expenditure for i, γ_i, plus an adjustment term that depends on provider i's altruism relative to their peers on the a's care team, $\frac{\alpha_i \gamma_i}{1 + \sum_{\ell \in \mathcal{I}_a} \alpha_\ell \gamma_\ell}$, and the difference between the ideal total expenditure on a and the total profit maximizing expenditure of providers on a's care team, $\tilde{e}_a - \sum_{\ell \in \mathcal{I}_a} \gamma_\ell$. If ideal expenditure is larger than the expenditure that maximizes aggregate profit of $\mathcal{I}_a$, then the adjustment term is positive, and provider i increases their expenditure above their profit maximizing level. If ideal expenditure is smaller than aggregate profit maximizing expenditure, then the adjustment term is negative, and the provider decreases expenditure below their profit maximizing level. In both cases, the adjustment term is increasing in the level of altruism of providers on the care team, and providers with larger altruistic preference adjust to the ideal level expenditure by larger amounts.

Equation 12 shows that equilibrium total expenditure on a is linear in $\tilde{e}_a$. The slope of this line depends on the aggregate tendency to provide ideal care, $\sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i$, such that more altruistic preference on the care team of a implies a closer relationship between $\tilde{e}_a$ and e_a^*. The intercept depends on total profit maximizing expenditure relative to aggregate tendency to provide ideal care.

Finally, Equation 13 shows that Nash equilibrium expenditure on a is equal to the ideal expenditure if and only if $\sum_{i \in \mathcal{I}_a} \gamma_i = \tilde{e}_a$. In other words, equilibrium expenditure on a is equal to ideal expenditure on a only if ideal expenditure also happens to be the aggregate-profit maximizing expenditure of the providers on a's care team. Otherwise, whether a receives less than ideal expenditure, and hence faces a cost-quality tradeoff, depends on whether aggregate profit maximizing expenditure is below the ideal level. This captures the presumption of physician incentive programs mentioned above: if care teams $\mathcal{I}_a$ for $a \in \mathcal{A}$ tend to have $\sum_{i \in \mathcal{I}_a} \gamma_i > \tilde{e}_a$, then reimbursements for services provided are above marginal cost, and there is overprovision of health services and a margin by which expenditure can be reduced and quality of care improved.

Incentive dilution, as discussed in Gaynor et al. (2004) and Frandsen & Rebitzer (2015), is present as providers do not account for the altruistic utility derived from other team members when choosing expenditure. The provider surplus-maximizing total expenditure on a, e_a^{PS}, is

$$e_a^{PS} = \frac{\left(\sum_{i\in\mathcal{I}_a}\alpha_i\right)\left(\sum_{i\in\mathcal{I}_a}\gamma_i\right)}{1+\left(\sum_{i\in\mathcal{I}_a}\alpha_i\right)\left(\sum_{i\in\mathcal{I}_a}\gamma_i\right)}\cdot\tilde{e}_a + \frac{\sum_{i\in\mathcal{I}_a}\gamma_i}{1+\left(\sum_{i\in\mathcal{I}_a}\alpha_i\right)\left(\sum_{i\in\mathcal{I}_a}\gamma_i\right)}. \tag{14}$$

Because $\alpha_i, \gamma_i > 0$, health production is always larger in the socially-optimal case: $h\left(\tilde{e}_a - e_a^{PS}\right) \geq h\left(\tilde{e}_a - \tilde{e}_a^*\right)$. Note, however, $e_a^{PS} < \tilde{e}_a$ if and only if $e_a^* < \tilde{e}_a$, since provider surplus-maximizing expenditure is likewise equal to ideal expenditure if and only if $\sum_{i\in\mathcal{I}_a}\gamma_i = \tilde{e}_a$.

3.3 Expenditure Choice with Incentive Pay

To examine the cost-quality tradeoff in the MSSP, I extend the model of expenditure choice to include the incentive for additional income (or penalties) from the MSSP. First, consider an ACO $j \in \mathcal{J}$, which has assigned beneficiaries $a \in \mathcal{A}_j$. Incentive pay to ACOs depends on total Medicare expenditure on assigned beneficiaries, which is given by

$$E_j = \sum_{a\in\mathcal{A}_j} e_a. \tag{15}$$

It will be useful to define the total ideal expenditure for beneficiaries assigned to ACO j, given by

$$\tilde{E}_j = \sum_{a\in\mathcal{A}_j} \tilde{e}_a, \tag{16}$$

as well as the per-beneficiary expenditure terms $e_j = \frac{E_j}{n_j}$ and $\tilde{e}_j = \frac{\tilde{E}_j}{n_j}$, where $n_j = |\mathcal{A}_j|$ is the total number of beneficiaries assigned to ACO j. I refer to the quantities e_j and $\tilde{e}_j$ as the expenditure and ideal expenditure of the average or representative beneficiary assigned to ACO j.

In the MSSP, Medicare's incentive contracts pay ACOs a fixed proportion of the cost-savings of an ACO, weighted by the ACO's quality score. ACOs with one-sided contracts earn incentive pay only if expenditure on assigned beneficiaries, E_j, is sufficiently less than the benchmark expenditure, BE_j. Specifically, the savings rate of an ACO, given by $\frac{BE_j - E_j}{BE_j}$, must be greater than the minimum savings rate, $\underline{S}_j$, in order for the ACO to earn a portion of the savings $BE_j - E_j$ as incentive pay. The payment contract for one-sided ACOs therefore has the form

$$P_j^{OS}\left(Q_j, BE_j - E_j\right) = F_j Q_j \left(BE_j - E_j\right) \mathbb{1}\left\{BE_j - E_j \geq \underline{S}_j BE_j\right\} \tag{17}$$

where $Q_j \in (0,1)$ is the quality score of the ACO, $F_j \in [0,1]$ is the sharing rate, and $\mathbb{1}\{\cdot\}$ is the indicator function.[13] Two-sided contracts offer the same incentive pay as one-sided contracts, but also penalize ACOs

[13] Note that I do not include a threshold for Q_j to represent meeting quality reporting standards. This is because the few ACOs in the sample that fail to meet quality standards also *all* fail to meet the savings threshold for earning incentive pay.

for savings rates less than the maximum loss rate:

$$P_j^{TS}(Q_j, BE_j - E_j) = F_j Q_j (BE_j - E_j) \mathbb{1}\{BE_j - E_j \geq \underline{S}_j BE_j\} \\ + (1 - F_j Q_j)(BE_j - E_j) \mathbb{1}\{BE_j - E_j \leq -\underline{S}_j BE_j\}. \quad (18)$$

For simplicity, let $T_j \in \{0, 1\}$ indicate whether the ACO j has a two-sided contract, and define the ACO contract with paramters F_j and T_j as

$$P_j(Q_j, BE_j - E_j; F_j, T_j) = \overbrace{F_j Q_j (BE_j - E_j) \mathbb{1}\{BE_j - E_j \geq \underline{S}_j BE_j\}}^{\text{Shared Savings}} \\ + \underbrace{T_j (1 - F_j Q_j)(BE_j - E_j) \mathbb{1}\{BE_j - E_j \leq -\underline{S}_j BE_j\}}_{\text{Shared Losses}}. \quad (19)$$

I assume an ACO's quality score Q_j measures the average health of beneficiaries assigned to the ACO. Specifically, ACO j's quality score Q_j is given by

$$Q_j = Q\left(n_j^{-1} \sum_{a \in \mathcal{A}_j} h(\tilde{e}_a - e_a) + \xi_j\right) \quad (20)$$

where $Q : \mathbb{R} \to (0, 1)$ is a strictly increasing function, and $\xi_j \in \mathbb{R}$ represents ACO-specific factors determining quality scores through determinants other than expenditure-determined health, $h(\tilde{e}_a - e_a)$. I assume ξ_j is observed by providers, and is hence endogenous. In Section 4, ξ_j is treated as ACO heterogeneity that is unobserved by the econometrician.

The discontinuity of ACO payment contracts at expenditures $E_j = (1 - \underline{S}_j) BE_j$ and $E_j = (1 + \underline{S}_j) BE_j$ presents a challenge for determining provider expenditures. In the extreme case, the payoffs of providers in a two-sided ACO ($T_j = 1$) may have three local-maxima (one at each region $E_j \in (0, (1 - \underline{S}_j) BE_j]$, $E_j \in ((1 - \underline{S}_j) BE_j, (1 + \underline{S}_j) BE_j)$, and $E_j \in [(1 + \underline{S}_j) BE_j, \infty)$), and so Nash equilibrium expenditures may not be unique. To achieve a unique optimal expenditures, I again follow Gaynor et al. (2004) and Frandsen & Rebitzer (2015) and assume that providers are uncertain about the cost savings generated by the ACO. A sufficient level of uncertainty effectively smooths the shared savings contracts, and results in provider objectives that are strictly concave, hence yielding a single optimum of the provider's maximization problem. Uncertainty is also important for model fit—it is ultimately a behavioral assumption regarding the extent to which ACO providers are aware that ACO expenditure is near the threshold to earn shared savings or avoid shared losses.[14] Finally, the MSSP was designed with provider uncertainty in mind: the minimum beneficiary threshold of 5000 was implemented so that providers were ultimately subject to lower per-beneficiary uncertainty from idiosyncratic beneficiary health shocks.

[14]Specifically, a model without uncertainty generates ACO expenditure predictions that are more frequently just below $(1 - \underline{S}_j) BE_j$ and just below $(1 + \underline{S}_j) BE_j$. A model with uncertainty more often predicts expenditures that barely miss earning shared savings or avoiding shared losses, which happens to be the case in the data from the MSSP.

To formalize the uncertainty faced by ACO providers, I assume that providers view the benchmark expenditure of ACO j, BE_j, as a random variable that is correlated with the the ideal expenditures of assigned beneficiaries. Specifically, I assume

$$BE_j = \sum_{a \in \mathcal{A}_j} \tilde{e}_a + \eta_a \tag{21}$$

where $\eta_a \in \mathbb{R}$ are independent and identically distributed across beneficiaries in the same ACO. Specifically, $\mathbb{E}[\eta_a | a \in \mathcal{A}_j] = \beta_j$ and $\text{Var}(\eta_a) = \sigma^2$. The random variable η_a can be thought of as the *bias* or *leniency* of benchmark expenditure for assigned beneficiary a. For example, a beneficiary that has larger health needs than what is accounted for by the benchmark expenditure BE_j would have a negative η_a. Given the large number of beneficiaries assigned to ACOs, I appeal to the central limit theorem, and conclude that $BE_j \sim N\left(\tilde{E}_j + n_j \beta_j, n_j \sigma^2\right)$ from the perspective of ACO providers.

The objective of provider $i \in \mathcal{I}_a$ for $a \in \mathcal{A}_j$ is

$$\max_{e_{ia}} \pi_i(e_{ia}) + \alpha_i h\left(\tilde{e}_a - e_a\right) + \omega_{ij} \mathbb{E}\left[P_j\left(Q_j, BE_j - E_j; F_j, T_j\right)\right]. \tag{22}$$

The provider maximizes their usual payoff, plus the expected value of incentive pay from the performance of ACO j. The expected value is taken over the distribution of BE_j.

The parameter $\omega_{ij} \in [0, 1]$ is i's share of ACO j's earnings. This parameter is particularly important, as it ultimately determines the responsiveness of provider i's expenditure choice on a to the incentives offered by the MSSP. I allow $\omega_{ij} \equiv 0$ for any provider $i \in \mathcal{I}_a$ but $i \notin \mathcal{I}_j$, as in this case the provider is on the care team of a beneficiary assigned to ACO j but does not benefit from the performance of ACO j. This means that $\sum_{i \in \mathcal{I}_a} \omega_{ij} \leq \sum_{i \in \mathcal{I}_j} \omega_{ij}$, where equality holds only if $\mathcal{I}_a \subseteq \mathcal{I}_j$. Furthermore, I allow $\sum_{i \in \mathcal{I}_j} \omega_{ij} \leq 1$, such that not all shared savings is necessarily allocated to providers—this is the case if the ACO keeps a portion of incentive pay for operating expenses.

Parameterization of the allocation of incentive pay among ACO providers accounts for unobserved contracts between ACOs and ACO participants (known as "ACO Participant Agreements"), which are generally not publicly available.[15] This approach is more general than that of Gaynor et al. (2004), who assume that HMO group incentive pay is allocated among the group according to physician patient shares, and Frandsen & Rebitzer (2015), who assume providers are identical and split incentive pay evenly.

[15] See `https://go.cms.gov/2HiHgus` for more detail.

By integrating over the distribution of BE_j, we get

$$\mathbb{E}\left[P_j\left(Q_j, BE_j - E_j; F_j, T_j\right)\right]$$
$$= F_j Q_j \overbrace{\left[\left(\tilde{E}_j + n_j\beta_j - E_j\right)\Phi\left(\frac{\tilde{E}_j + n_j\beta_j - \frac{E_j}{1-\underline{S}_j}}{\sqrt{n_j}\sigma}\right) + \sqrt{n_j}\sigma\phi\left(\frac{\tilde{E}_j + n_j\beta_j - \frac{E_j}{1-\underline{S}_j}}{\sqrt{n_j}\sigma}\right)\right]}^{\equiv \Psi^+\left(E_j;\tilde{E}_j + n_j\beta_j, \underline{S}_j, \sqrt{n_j}\sigma\right)}$$
$$+ T_j\left(1 - F_j Q_j\right)\underbrace{\left[\left(\tilde{E}_j + n_j\beta_j - E_j\right)\Phi\left(\frac{\frac{E_j}{1+\underline{S}_j} - \tilde{E}_j - n_j\beta_j}{\sqrt{n_j}\sigma}\right) + \sqrt{n_j}\sigma\phi\left(\frac{\frac{E_j}{1+\underline{S}_j} - \tilde{E}_j - n_j\beta_j}{\sqrt{n_j}\sigma}\right)\right]}_{\equiv \Psi^-\left(E_j;\tilde{E}_j + n_j\beta_j, \underline{S}_j, \sqrt{n_j}\sigma\right)} \quad (23)$$

where Φ and ϕ are the standard normal cumulative and probability density functions. I use the shorthand $\Psi^+\left(E_j;\tilde{E}_j + n_j\beta_j, \underline{S}_j, \sqrt{n_j}\sigma\right) \equiv \Psi_j^+$, and likewise $\Psi^{+\prime}\left(E_j;\tilde{E}_j + n_j\beta_j, \underline{S}_j, \sqrt{n_j}\sigma\right) \equiv \psi_j^+$, for the first term in square brackets and its derivative. The terms Ψ_j^- and ψ_j^- are similarly defined for the second term in square brackets.[16]

The first order condition for provider $i \in \mathcal{I}_a$ with $a \in \mathcal{A}_j$ is

$$\gamma_i - e_{ia}^* + \alpha_i\gamma_i\left(\tilde{e}_a - e_a\right)$$
$$+ \omega_{ij}\gamma_i\left\{F_j\left[n_j^{-1}\left(\tilde{e}_a - e_a\right)Q_j'\left(\Psi_j^+ - T_j\Psi_j^-\right) + Q_j\left(\psi_j^+ - T_j\psi_j^-\right)\right] + T_j\psi_j^-\right\} = 0\,, \quad (28)$$

where $Q_j' \equiv Q'\left(n_j^{-1}\sum_{a\in\mathcal{A}_j} h\left(\tilde{e}_a - e_a\right) + \xi_j\right)$. Assuming that Equation 28 holds for all $i \in \mathcal{I}_a$ and taking the sum of their first order conditions, I obtain the aggregated condition for equilibrium expenditure

$$\sum_{i\in\mathcal{I}_a}\gamma_i - e_a^* + \left(\sum_{i\in\mathcal{I}_a}\alpha_i\gamma_i\right)\left(\tilde{e}_a - e_a\right)$$
$$+ \left(\sum_{i\in\mathcal{I}_a}\omega_{ij}\gamma_i\right)\left\{F_j\left[n_j^{-1}\left(\tilde{e}_a - e_a\right)Q_j'\left(\Psi_j^+ - T_j\Psi_j^-\right) + Q_j\left(\psi_j^+ - T_j\psi_j^-\right)\right] + T_j\psi_j^-\right\} = 0\,. \quad (29)$$

[16] To be clear, let $Z_j^+ = \left(\tilde{E}_j + n_j\beta_j - \frac{E_j}{1-\underline{S}_j}\right)\Big/\sqrt{n_j}\sigma$ and $Z_j^- = \left(\frac{E_j}{1+\underline{S}_j} - \tilde{E}_j - n_j\beta_j\right)\Big/\sqrt{n_j}\sigma$. Then,

$$\Psi^+\left(E_j;\tilde{E}_j + n_j\beta_j, \underline{S}_j, \sqrt{n_j}\sigma\right) = \left(\tilde{E}_j + n_j\beta_j - E_j\right)\Phi\left(Z_j^+\right) + \sqrt{n_j}\sigma\phi\left(Z_j^+\right) \equiv \Psi_j^+ \quad (24)$$
$$\Psi^{+\prime}\left(E_j;\tilde{E}_j + n_j\beta_j, \underline{S}_j, \sqrt{n_j}\sigma\right) = -\Phi\left(Z_j^+\right) - \frac{E_j\underline{S}_j}{\sqrt{n_j}\sigma\left(1-\underline{S}_j\right)^2}\phi\left(Z_j^+\right) \equiv \psi_j^+ \quad (25)$$
$$\Psi^-\left(E_j;\tilde{E}_j + n_j\beta_j, \underline{S}_j, \sqrt{n_j}\sigma\right) = \left(\tilde{E}_j + n_j\beta_j - E_j\right)\Phi\left(Z_j^-\right) + \sqrt{n_j}\sigma\phi\left(Z_j^-\right) \equiv \Psi_j^- \quad (26)$$
$$\Psi^{-\prime}\left(E_j;\tilde{E}_j + n_j\beta_j, \underline{S}_j, \sqrt{n_j}\sigma\right) = -\Phi\left(Z_j^-\right) - \frac{E_j\left(2+\underline{S}_j\right) - 2\left(\tilde{E}_j + n_j\beta_j\right)\left(1+\underline{S}_j\right)}{\sqrt{n_j}\sigma\left(1+\underline{S}_j\right)^2}\phi\left(Z_j^-\right) \equiv \psi_j^-\,. \quad (27)$$

Finally Equation 29 can be solved for e_a^* to yield the following *implicit* form for equilibrium expenditure:

$$e_a^* = \frac{\sum_{i\in\mathcal{I}_a}\alpha_i\gamma_i + \left(\sum_{i\in\mathcal{I}_a}\omega_{ij}\gamma_i\right) n_j^{-1}F_jQ_j'\left(\Psi_j^+ - T_j\Psi_j^-\right)}{1+\sum_{i\in\mathcal{I}_a}\alpha_i\gamma_i + \left(\sum_{i\in\mathcal{I}_a}\omega_{ij}\gamma_i\right) n_j^{-1}F_jQ_j'\left(\Psi_j^+ - T_j\Psi_j^-\right)}\cdot\tilde{e}_a + \frac{\sum_{i\in\mathcal{I}_a}\gamma_i + \left(\sum_{i\in\mathcal{I}_a}\omega_{ij}\gamma_i\right)\left(F_jQ_j\psi_j^+ + T_j\left(1-F_jQ_j\right)\psi_j^-\right)}{1+\sum_{i\in\mathcal{I}_a}\alpha_i\gamma_i + \left(\sum_{i\in\mathcal{I}_a}\omega_{ij}\gamma_i\right) n_j^{-1}F_jQ_j'\left(\Psi_j^+ - T_j\Psi_j^-\right)}. \quad (30)$$

This is implicit because each of Q_j, Ψ_j^+, and Ψ_i^- and their respective derivatives are functions of e_a^* through $E_j = \sum_{a\in\mathcal{A}_j} e_a^*$.

Equation 30 is the extension of Equation 12 that accounts for the incentives of the MSSP. First, recall from Equation 12,

$$e_a^* = \frac{\sum_{i\in\mathcal{I}_a}\alpha_i\gamma_i}{1+\sum_{i\in\mathcal{I}_a}\alpha_i\gamma_i}\cdot\tilde{e}_a + \frac{\sum_{i\in\mathcal{I}_a}\gamma_i}{1+\sum_{i\in\mathcal{I}_a}\alpha_i\gamma_i}\,, \quad (31)$$

where aggregate altruism term $\sum_{i\in\mathcal{I}_a}\alpha_i\gamma_i$ disciplines provider expenditures relative to ideal expenditure, and $\sum_{i\in\mathcal{I}_a}\gamma_i$ is the total-profit maximizing expenditure of $\mathcal{I}_a$. The MSSP effects expenditure in two ways. First, like the altruism term $\sum_{i\in\mathcal{I}_a}\alpha_i\gamma_i$, the effect of ACO quality scores is represented by the term

$$\left(\sum_{i\in\mathcal{I}_a}\omega_{ij}\gamma_i\right) n_j^{-1}F_jQ_j'\left(\Psi_j^+ - T_j\Psi_j^-\right)\,. \quad (32)$$

This term represents the expected marginal payment to providers on a care team for obtaining expenditure e_a^* closer to $\tilde{e}_a$. This term is strictly positive, since $Q_j' > 0$ by assumption, $\Psi_j^+ > 0$, and $\Psi_j^+ - \Psi_j^- > 0$.

Second, analogous to the term $\sum_{i\in\mathcal{I}_a}\gamma_i$, the effect of decreasing expenditure relative to the benchmark is to increase the available incentive pay. This is represented by the term

$$\left(\sum_{i\in\mathcal{I}_a}\omega_{ij}\gamma_i\right)\left(F_jQ_j\psi_j^+ + T_j\left(1-F_jQ_j\right)\psi_j^-\right)\,, \quad (33)$$

which represents the marginal increase in expected incentive pay to the care team of the representative beneficiary from an increase in e_a^*. This term is negative, because increasing expenditure always decreases the expected amount of incentive pay earned and increases expected payment of shared losses.

When does expenditure on a equal ideal expenditure? Similar to the case without incentive pay, we have

$$\tilde{e}_a = e_a^* \iff \tilde{e}_a = \sum_{i\in\mathcal{I}_a}\gamma_i + \left(\sum_{i\in\mathcal{I}_a}\omega_{ij}\gamma_i\right)\left(F_jQ_j\psi_j^+ + T_j\left(1-F_jQ_j\right)\psi_j^-\right)\,. \quad (34)$$

Thus, the “optimal” ACO contract that implements ideal expenditure is one that corrects a’s care team’s profit maximizing expenditure (given by the first term $\sum_{i\in\mathcal{I}_a}\gamma_i$) by an amount that exactly offsets its over-expenditure with the contract parameters F_j and T_j. This optimal contract is neither unique nor guaranteed to exist for $F_j \in [0,1]$. For example, given $e_a^* > \tilde{e}_a$ in the absence of incentive pay, the value of F_j that implements $e_a^* = \tilde{e}_a$ is different for $T_j = 0$ and $T_j = 1$. Furthermore, for $e_a^* \gg \tilde{e}_a$, $F_j > 1$ is nessecary to implement ideal expenditure. On the other hand, if $e_a^* < \tilde{e}_a$, then there is no value of F_j and T_j that can implement $e_a^* = \tilde{e}_a$, because the adjustment imposed by the MSSP contract is always negative (decreases e_a^*). For these reasons, counterfactual simulations of this model will focus on find the value of F_j that maximizes (per-beneficiary) program savings for a given $T_j \in \{0,1\}$ across all ACOs.

4 Identification and Estimation

In this section, I walk through the identification and estimation of all model parameters. Above, I showed that the expenditure on $a \in \mathcal{A}_j$ is function of the aggregated parameters of the payoff functions of providers on the care teams of a, $\mathcal{I}_a$. Accordingly, I estimate the parameters α and ω, defined

$$\alpha \equiv \sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i \,, \tag{35}$$

$$\omega \equiv \sum_{a \in \mathcal{A}_j} \sum_{i \in \mathcal{I}_a} \omega_{ij} \gamma_i \,, \tag{36}$$

which are, respectively, the average level of altruism and the average responsiveness to incentive pay of the average ACO-assigned beneficiary's care team. I also estimate γ_j, defined

$$\gamma_j \equiv \sum_{i \in \mathcal{I}_a} \gamma_i \,, \tag{37}$$

which is the aggregate profit maximizing expenditure for the care team of the average beneficiary in ACO j. I also estimate β_j and σ, which are the expected value and standard deviation of per-beneficiary benchmark bias.

As described in Section 3, the parameters α and ω represent the degree of altruism and the responsiveness to incentive pay of Medicare providers. These can be rewritten

$$\alpha \equiv N \cdot (\mu_\alpha \mu_\gamma + \sigma_{\alpha\gamma}) \tag{38}$$

$$\omega \equiv N \cdot (\mu_\omega \mu_\gamma + \sigma_{\omega\gamma}) \,, \tag{39}$$

where N is the team size, or number of providers, on the average beneficiary's care team. The terms μ_α, μ_γ, and μ_ω are the means of α_i, γ_i and ω_{ij} of providers on the care team, and $\sigma_{\alpha\gamma}$ and $\sigma_{\omega\gamma}$ the covariances of γ_i with α_i and ω_{ij}. Thus, larger values of α and ω may correspond to larger team size, higher levels of profit maximizing expenditure, higher levels of altruism, or higher levels of membership of providers on a care team of a beneficiary to the ACO that the beneficiary is assigned to. Interestingly, the efficiency of contracts between ACO participants and ACOs used to distribute shared savings is also captured by ω; this is through the term $\sigma_{\omega\gamma}$. If γ_i and ω_{ij} have a larger covariance, then this means provider i provides more care to a beneficiary and that that that provider also earns a larger proportion of shared savings. This ultimately brings expenditure closer to ideal expenditure and closer to the provider surplus maximizing expenditure.

To estimate model primitives, I use an unbalanced panel of ACO-level data on all ACOs from 2014 to 2018. Track 1 ACOs that have a sharing rate of $F_j = 0.5$ and are one-sided with $T_j = 0$ form the bulk of the sample; Track 1+, Track 2, and Track 3 ACOs are two-sided with $T_j = 1$, and have varying sharing rates of $F_j = 0.5$, $F_j = 0.6$, and $F_j = 0.75$, respectively. Later in this section, I explain how I leverage the variation in ACO contracts to identify model parameters.

I estimate this model in two stages. First, I estimate per-beneficiary expected benchmark bias β_j and standard deviation σ by using the quality score data of ACOs. This yields estimates of these parameters, $\hat{\beta}_j$ and $\hat{\sigma}$, as well as an estimate of ideal per-beneficiary expenditure, $\hat{\tilde{e}}_j$. Second, I estimate the parameters α, ω, and γ_j with the aggregated first order conditions for the representative ACO beneficiary, analogous to Equation 29. Both stages of estimation require accounting for unobserved ACO heterogeneity that is likely correlated with per-beneficiary expenditure; in the first stage, I use ACO fixed-effects for identification, and in the second stage, I instrument for per-beneficiary expenditure. At this point, I append the index $t \in \{2014, 2015, 2016, 2017, 2018\}$ to variables that vary by year.

4.1 Identification and Estimation of Uncertainty Parameters

The first step to estimating this model is to estimate parameters of the distribution of BE_{jt}. Previously, I showed that the benchmark expenditure of ACO j, from the perspective of a provider in ACO j, is a normally distributed random variable: $BE_{jt} \sim N\left(n_{jt}\beta_{jt}, n_{jt}\sigma^2\right)$, where β_{jt} and σ are the mean and standard deviation of the uncertain component of per-beneficiary benchmark expenditure.

First, assuming the representative beneficiary in ACO j has equilibrium expenditure $e^*_{jt} = E_{jt}/n_{jt}$, Equation 20 implies ACO j's quality score in year t is

$$Q_{jt} = Q\left(h\left(\tilde{e}_{jt} - e^*_{jt}\right) + \xi_{jt}\right) \,, \tag{40}$$

where $\tilde{e}_{jt}$ is j's representative beneficiary ideal expenditure in year t, e^*_{jt} is expenditure on j's representative beneficiary, and ξ_{jt} represents other determinants of quality score. I specify the logistic curve $Q(\cdot) = \frac{\exp(\cdot)}{1+\exp(\cdot)}$ for the relationship between quality score and beneficiary health. This sigmoidal form fits observed quality score data, which ranges from 0 to 1, exclusive. $Q(\cdot)$ has the closed-form derivative

$$Q'_{jt} = Q_{jt}\left(1 - Q_{jt}\right) \,, \tag{41}$$

which imposes diminishing returns to quality score for high levels of beneficiary health. This assumption is important for model fit, where quality score varies more for a given ACO at Q_{jt} values closer to 0.5.

Inverting $Q(\cdot)$, I obtain the expression

$$-\log\left(\frac{1-Q_{jt}}{Q_{jt}}\right) = -\frac{1}{2}\left(\tilde{e}_{jt} - e^*_{jt}\right)^2 + \xi_{jt}\,. \tag{42}$$

Note that per-beneficiary benchmark expenditure, $be_{jt} \equiv BE_{jt}/n_{jt}$ can be written

$$be_{jt} = \tilde{e}_{jt} + \beta_{jt} + \eta_{jt} \tag{43}$$

where $\eta_{jt} \sim N\left(0, \frac{\sigma^2}{n_{jt}}\right)$ is the idiosyncratic variation in benchmark expenditure that is uncorrelated with ACO characteristics. This means

$$-\log\left(\frac{1-Q_{jt}}{Q_{jt}}\right) = -\frac{1}{2}\left(be_{jt} - \beta_{jt} - \eta_{jt} - e^*_{jt}\right)^2 + \xi_{jt}\,. \tag{44}$$

Expanding the quadratic term and taking the expected value over η_{jt}, I obtain

$$\mathbb{E}\left[-\log\left(\frac{1-Q_{jt}}{Q_{jt}}\right)\right] = -\frac{1}{2}\left(be_{jt} - \beta_{jt} - e^*_{jt}\right)^2 - \frac{\sigma^2}{2n_{jt}} + \xi_{jt}\,, \tag{45}$$

which follows from $\mathbb{E}\left[\eta^2_{jt}\right] = \mathrm{Var}\left(\eta_{jt}\right)$.

Next, I assume that ξ_{jt} can be decomposed into an endogenous time-invariant component $\bar{\xi}_j$ that is observed by the ACO, and an idiosyncratic component $\dot{\xi}_{jt}$ that is not observed and hence operates like classical measurement error. Then,

$$-\log\left(\frac{1-Q_{jt}}{Q_{jt}}\right) = -\frac{1}{2}\left(be_{jt} - \beta_{jt} - e^*_{jt}\right)^2 - \frac{\sigma^2}{2n_{jt}} + \bar{\xi}_j + \dot{\xi}_{jt}\,, \tag{46}$$

where the expectation operator is removed by absorbing mean-zero uncertainty into $\dot{\xi}_{jt}$.

Finally, I make the following two assumptions to identify β_{jt} and σ^2 and estimate this equation with nonlinear least squares (NLS) with ACO-level fixed-effects.

Assumption 4.1. Conditional on $\bar{\xi}_j$, the time-varying unobserved quality score determinant $\dot{\xi}_{jt}$ is mean-independent of be_{jt}, β_{jt}, e^*_{jt}, and n_{jt}, such that $\mathbb{E}\left[\dot{\xi}_{jt}|\bar{\xi}_j, be_{jt}, \beta_{jt}, e^*_{jt}, n_{jt}\right] = \mathbb{E}\left[\dot{\xi}_{jt}|\bar{\xi}_j\right] = 0$.

Assumption 4.2. The expected value of per-beneficiary benchmark expenditure bias, β_{jt}, is strictly positive.

The first assumption provides moment conditions that are used to identify β_{jt} and σ. Specifically, this assumption states that variation in ideal expenditure and actual expenditure within ACO, as well as variation in the number of assigned beneficiaries within ACO, is uncorrelated with other time-varying factors that influence quality score Q_{jt}. Thus, σ is identified from variation in quality score within the same ACO across years that is correlated with the inverse of the number of beneficiaries assigned to the ACO and not accounted for by variation in benchmark expenditure relative to realized expenditure. In other words, because $\sigma^2/2n_{jt}$

is inversely related to quality score in Equation 46, the variation in quality score that is not explained by expenditure relative to ideal expenditure is larger in smaller ACOs with fewer assigned beneficiaries.

The second assumption is necessary to identify β_{jt} in a quadratic function. For fixed values of Q_{jt}, be_{jt}, e^*_{jt}, n_{jt}, $\overline{\xi}_j$, and $\dot{\xi}_{jt}$, there are *two* possible values of β_{jt} that satisfy Equation 46: one positive, and one negative. Thus, the assumption that $\beta_{jt} > 0$ is required to resolve the indeterminacy and identify the model.

I emphasize that Assumption 4.2 is neither unrealistic nor particularly strong. First, as discussed in Section 2, assignment of the benchmark expenditures of the MSSP follows a procedure that takes several steps to avoid under-assignment. Specifically, an ACO's benchmark expenditure is determined by the history of Medicare expenditure on assigned beneficiaries in the three years prior to the current agreement period; these are referred to as "historical benchmarks." For ACOs in their first agreement period (which are 1592 out of 2180 observations in the estimation sample), historical benchmarks are therefore computed from expenditures determined outside of the MSSP. ACOs in their second agreement period have historical benchmarks based on expenditure during their first agreement period, but these historical benchmarks are adjusted by geographical region to not disproportionately disadvantage an ACO that decreased spending their their first agreement period. Furthermore, historical benchmarks are risk-adjusted and trended to account for national growth in Medicare expenditure.

Second, given $be_{jt} = \tilde{e}_{jt} + \beta_{jt} + \eta_{jt}$, it is possible under Assumption 4.2 that *realized* per-beneficiary benchmark expenditure is less than ideal expenditure; rather, this assumption regards the expected bias. It is unlikely that Medicare providers would join an ACO with an expectation that the benchmark expenditure assignment would be less than the ideal expenditure for their patients. For estimation, I parameterize $\beta_{jt} = \exp\left(\boldsymbol{\beta}'\mathbf{x}^Q_{jt}\right)$, where $\mathbf{x}^Q_{jt}$ is a vector including a constant and variables describing the of characteristics of the beneficiaries assigned to ACO j, and $\boldsymbol{\beta}$ is a parameter vector.

4.2 Identification and Estimation of Provider Payoff Parameters

To estimate provider payoffs, I begin by computing the fitted values of $\hat{\beta}_{jt} = \exp\left(\hat{\boldsymbol{\beta}}'\mathbf{x}_{jt}\right)$, $\hat{\sigma}$, and $\hat{\tilde{e}}_{jt}$ from the first stage of estimation. I use these to compute

$$\hat{\Psi}^+_{jt} = \Psi^+\left(E_{jt}; \hat{\tilde{E}}_{jt} + n_{jt}\hat{\beta}_{jt}, \underline{S}_{jt}, \sqrt{n_{jt}}\hat{\sigma}\right) \tag{47}$$

as well as $\hat{\Psi}^-_{jt}$, $\hat{\psi}^+_{jt}$, and $\hat{\psi}^-_{jt}$. Then, I use these values in the aggregated first order condition for expenditure on the representative beneficiary of ACO j (analogous to Equation 29)

$$\begin{aligned} e^*_{jt} &= \gamma_{jt} + \alpha\left(\hat{\tilde{e}}_{jt} - e^*_{jt}\right) \\ &\quad + \omega\left\{F_{jt}Q_{jt}\left[n_j^{-1}\left(\hat{e}_j - e_j\right)\left(1 - Q_{jt}\right)\left(\hat{\Psi}^+_j - T_j\hat{\Psi}^-_j\right) + \hat{\psi}^+_j - T_j\hat{\psi}^-_j\right] + T_j\hat{\psi}^-_j\right\}, \end{aligned} \tag{48}$$

where $Q'_{jt} = Q_{jt}(1 - Q_{jt})$. Because fitted variables from the previous stage of estimation are used to estimate this stage of the model, I account for error propagation by bootstrapping for standard errors.

The strategy to estimate this model is to treat γ_{jt}, a parameter measuring the profit-maximizing level of expenditure for providers on the care team of the representative beneficiary in an ACO, as endogenous variation in expenditure on the representative beneficiary. Moreover, γ_{jt} represents variation in expenditure that is not explained by a provider's altruistic preferences or by incentive pay through the MSSP, which are accounted for by the second and third additive terms in the regression. Thus, we can interpret the parameter α as the effect of a beneficiaries health state on expenditure, and ω as the effect of incentive pay on expenditure.

Clearly, estimation of this equation with ordinary least squares will lead to biased parameter estimates. The outcome variable e^*_{jt} is present on both sides of the equation, and thus any term on the right hand side that is a function of e^*_{jt} is by definition correlated with unexplained variation in γ_{jt}. Thus, the standard OLS moment conditions will lead to biased estimates. Moreover, because I use the fitted values of several variables, non-classical measurement error threatens to attenuate the coefficients estimates.

I instruments to obtain consistent estimates of α and ω in this linear regression. Valid instruments in this setting are correlated with representative beneficiary expenditure e^*_{jt}, but must be uncorrelated with endogenous variation in γ_{jt} that determines profit-maximizing expenditure of Medicare beneficiaries. Therefore, I can use exogenous variation in incentive pay offered by the MSSP as instruments for the endogenous terms.

First, I specify $\gamma_{jt} = \boldsymbol{\gamma}'\mathbf{x}^e_{jt} + \dot{\gamma}_{jt}$, where $\mathbf{x}^e_{jt}$ is a vector of ACO-specific characteristics describing both provider and beneficiary characteristics. Next, I make the following assumption.

Assumption 4.3. Given $\mathbf{x}^e_{jt}$, $\dot{\gamma}_{jt}$ is mean-independent of $\hat{\beta}_{jt}$, F_{jt}, and T_{jt}, and $\mathbb{E}\left[\dot{\gamma}_{jt}|x^e_{jt}, \hat{\beta}_{jt}, F_{jt}, T_{jt}\right] = \mathbb{E}\left[\dot{\gamma}_{jt}|x^e_{jt}\right] = 0$.

In other words, I assume that conditional on the observed characteristics of ACO j, the expected benchmark bias $\hat{\beta}_{jt}$ and MSSP contract parameters F_{jt} and T_{jt} are uncorrelated with variation in $\dot{\gamma}_{jt}$. While F_{jt} and T_{jt} are chosen by the ACO, I assume conditioning on ACO characteristics *and* incentive pay offered by the ACO (the second additive term in Equation 48) is enough to eliminate selection bias. Moreover, the latent efficiency of ACO operation, which is the most likely avenue for selection of ACOs by F_{jt} and T_{jt}, is parameterized by ω through the covariance $\sigma_{\omega\gamma}$. Thus, this kind of endogeneity of F_{jt} and T_{jt} does not bias ω, but rather is a component of its magnitude.

Thus, I estimate Equation 48 with two-stage least squares (2SLS) using $\hat{\beta}_{jt}$, F_{jt}, and T_{jt} as instruments, and $\mathbf{x}^e_{jt}$ as controls. The first-stage relevance of these instruments is due to the MSSP itself: higher values of each of these instruments increase the marginal payoff to providers for having e^*_{jt} close to $\tilde{e}_{jt}$, and thus identify α and ω through variation that is independent of $\dot{\gamma}_{jt}$.

4.3 Accounting for Voluntary Participation in the MSSP

The final empirical analysis I perform accounts for voluntary participation in the MSSP when simulating counterfactual contracts that may incentivize ACO exit through low or negative incentive pay. I assume a simple non-strategic exit model, where ACO j exits the MSSP after year t depending on the amount of shared savings earned by the ACO and other observed components. Specifically, the probability that ACO j exits the MSSP after performance year t is

$$\Pr\left\{Exit_{jt}|P_{jt}(\cdot),\mathbf{x}_{jt}^{Exit}=1\right\}=\frac{\exp\left(\delta_0 P_{jt}\left(\cdot\right)+\boldsymbol{\delta}_1'\mathbf{x}_{jt}^{Exit}\right)}{1+\exp\left(\delta_0 P_{jt}\left(\cdot\right)+\boldsymbol{\delta}_1'\mathbf{x}_{jt}^{Exit}\right)}, \tag{49}$$

where $Exit_{jt}$ indicates the exit of ACO j after performance year t, $P_{jt}(\cdot)=P_{jt}\left(Q_{jt},BE_{jt}-E_{jt};F_{jt},T_{jt}\right)$ is the observed shared savings (or shared losses) earned by ACO j in year t, and $\mathbf{x}_{jt}^{Exit}$ is a vector of ACO characteristics. I estimate the parameters δ_0 and $\boldsymbol{\delta}_1$ with maximum likelihood.

5 Estimation Results and Model Fit

Table 2 shows the estimates for all parameters of the model. In order to generate the statistics presented, I first estimate Equation 46 with NLS, Equation 48 with 2SLS, and Equation 49 with maximum likelihood with the full estimation sample. Next, the same estimation is performed 1406 times with a re-sampled dataset that has the same number of observations as the original estimation sample.[17] Standard errors and 95% confidence intervals are generated from the empirical standard deviations and quantiles of the recorded estimates.

Control variables in $\mathbf{x}_{jt}^{Q}$ include a constant term, the proportion of assigned beneficiaries in ACO j that are under 65 years old, the proportion of assigned beneficiaries in ACO j that are over 85 years old, the proportion of assigned beneficiaries in ACO j that are male, and the proportion of assigned beneficiaries in ACO j with a Medicare-reported race other than "white." These estimates indicate the benchmark bias is slightly increasing in the share of young and old beneficiaries assigned the ACO, but these estimates are imprecise. Bias is increasing in the share of nonwhite beneficiaries. The benchmark bias sharply decreases in the share of male beneficiaries. The average value of $\hat{\beta}_{jt}$ across all ACOs is \$417.55, meaning, benchmark expenditure exceeds ideal representative beneficiary expenditure by just over \$400 per beneficiary.

The estimate of σ is quite large at nearly \$50,000 *per beneficiary.* There are several factors that justify this high estimate. First, average expenditure per beneficiary in the estimation sample is \$10,776; however, among beneficiaries enrolled in Medicare through a diagnosis of end-stage renal disease (ESRD), average per beneficiary expenditure is \$80,218. The high estimate of σ is likely caused in large part by the disproportionate uncertainty from ESRD patients. Second, note that σ is the uncertainty in per-beneficiary

[17] I estimate the model 1500 times, and the NLS estimation fails to converge 94 times, yielding 1406 estimations from which to compute standard errors and confidence intervals. The failure of NLS convergence is likely due to the sensitivity of numerical optimization to initial conditions.

Table 2: **Parameter Estimates and Confidence Intervals**

Estimates from variation in Q_{jt}

	Parameter	Estimate (S.E.)	95% CI
	σ (in \$ thousands)	49.616 (10.815)	[12.638,54.872]
$\boldsymbol{\beta}$	Constant	6.360 (2.714)	[1.845, 12.926]
	Proportion of assigned beneficiaries: under 65	0.901 (1.749)	[-3.689, 3.242]
	Proportion of assigned beneficiaries: over 85	2.282 (4.607)	[-11.256, 7.973]
	Proportion of assigned beneficiaries: male	-18.751 (7.401)	[-37.399, -7.402]
	Proportion of assigned beneficiaries: nonwhite	1.316 (0.395)	[0.663, 2.190]

Notes: $N = 2180$. Includes ACO fixed-effects. Standard errors (in parentheses) and 95% confidence intervals are computed by bootstrapping with 1406 replications.

Estimates from variation in e^*_{jt}

	Parameter	Estimate (S.E.)	95% CI
	α	9.751 (1.581)	[7.731, 13.782]
	ω	11.631 (3.331)	[6.143, 19.108]
$\boldsymbol{\gamma}$	log(ACO providers per assigned beneficiary)	0.727 (0.202)	[0.404, 1.216]
	Proportion of assigned beneficiaries: under 65	-30.580 (27.685)	[-87.842, 22.003]
	Proportion of assigned beneficiaries: over 85	48.276 (19.565)	[-2.542, 79.842]
	Proportion of assigned beneficiaries: male	-54.392 (16.716)	[-83.169, -17.762]
	Proportion of assigned beneficiaries: nonwhite	6.690 (2.568)	[1.000, 10.990]
	Proportion of assigned beneficiaries: Aged, Non-Dual	34.086 (8.083)	[16.895, 48.607]
	Proportion of assigned beneficiaries: Aged, Dual	34.513 (8.113)	[17.117, 49.298]
	Proportion of assigned beneficiaries: Disabled	67.433 (28.083)	[11.980, 122.118]
	Proportion of assigned beneficiaries: ESRD	135.072 (54.431)	[30.239, 243.730]

Notes: $N = 2180$. Standard errors (in parentheses) and 95% confidence intervals are computed by bootstrapping with 1406 replications. First-stage F-statistics are α: 63.3; ω: 129.0. Sargan J test of over-identifying restrictions: $p = 0.51$.

Estimates from variation in $Exit_{jt}$

	Parameter	Estimate (S.E)	95% CI
	δ_0	-0.155 (0.048)	[-0.259, -0.079]
$\boldsymbol{\delta}_1$	Intercept	1.816 (1.124)	[-0.393, 4.051]
	$\mathbb{1}\{$ACO age is equal to 3$\}$	1.290 (0.151)	[0.995, 1.587]
	$\mathbb{1}\{$ACO age is equal to 6$\}$	1.168 (0.320)	[0.474, 1.734]
	$\log(n_j)$	-0.530 (0.117)	[-0.763, -0.312]
	log(ACO providers per assigned beneficiary)	-0.223 (0.079)	[-0.364, -0.064]

Notes: $N = 2180$. Standard errors (in parentheses) and 95% confidence intervals are computed by bootstrapping with 1406 replications.

benchmark expenditure *perceived* by providers in ACO j. Thus, σ also captures uncertainty regarding beneficiary assignment to ACOs, as well as (unmodeled) uncertainty regarding the expenditure choices of other providers. Finally, though σ is large, the overall variance of BE_j remains reasonable: the average of $\sqrt{n_j}\sigma$ is \$6.084 million, meaning the total uncertainty in BE_j is, on average, 4.37%. The benefit of a large estimate of σ is that there is no question as to whether provider payoffs are sufficiently smooth to obtain a unique equilibrium.

The control variables in $\mathbf{x}^e_{jt}$ include all but the constant in $\mathbf{x}^Q_{jt}$, plus the share of assigned beneficiaries by enrollment type. The shares of patient enrollment type sum to one, so a constant is not estimated in this model. I also include the log number of ACO providers per beneficiary in $\mathbf{x}^e_{jt}$; this is to control for the effect of team size, which increases γ_{jt}.

The principal parameters of the model, α and ω, are identified with considerable precision. The first-stage F-statistics for the 2SLS estimation far-exceed the threshold to establish relevance (e.g. the well known threshold of Stock, Wright, & Yogo (2002) is $F > 10$).

Finally, control variables used in $\mathbf{x}^{Exit}_{jt}$ included to explain variation in ACO exit that is driven by factors other than incentive pay. Indicators for ACO age being equal to 3 or 6 are included because these are the last years of the three year agreement periods of ACOs in the MSSP. The parameter δ_0 translates (at the mean of $P_{jt}(\cdot)$) to approximately a 0.02 decrease in the probability of exit for each additional \$1 million earned in incentive pay.

5.1 Model Fit

I address model fit by simulating various measures of ACO performance and MSSP program savings in the data.[18] Table 3 shows the means of observed and predicted expenditure per beneficiary (in \$ thousand) and ACO quality score, where simulations are denoted with a $\hat{\cdot}$. The first row displays the means of observed variables next to simulated variables in the full sample of 2180 observations. The second and third rows show observed versus simulated means in the sub-samples with only one-sided ($T_{jt} = 0$) ACOs and only two-sided ($T_{jt} = 1$) ACOs. Finally, the fourth, fifth, and sixth rows shows model fit by the type of two-sided contract, where the sharing rate F_{jt} is 0.5 for Track 1+, 0.6 for Track 2, and 0.75 for Track 3 ACOs.

Clearly, the simulations of ACO performance variables are extremely close to observed values. Simulated per beneficiary expenditure, even among the small sub-samples of two-sided contracts, matches observed per beneficiary expenditure up to the dollar. Quality scores are also nearly equal between observed and simulated means.

Next, Table 4 displays model fit for variables related to ACO payment and exit: $P_{jt}(\cdot)$ is the earned shared savings (or losses) of an ACO (in \$ million), $Qual^+_{jt} = \mathbb{1}\left\{BE_{jt} - E_{jt} \geq \underline{S}_{jt}BE_{jt}\right\}$ indicates whether the ACO achieved savings low enough to qualify for incentive pay, $Qual^-_{jt} = \mathbb{1}\left\{BE_{jt} - E_{jt} \leq -\underline{S}_{jt}BE_{jt}\right\}$ indicates whether the ACO expenditure was high enough to pay a penalty, and $Exit_{jt}$ indicates whether

[18] The process of simulating counterfactual equilibrium ACO performance values is described in detail in Section 6.

Table 3: **Model Fit: ACO Expenditure and Quality Score**

Sample	N	e^*_{jt}	$\hat{e}^*_{jt}$	$\hat{\hat{e}}_{jt} - e^*_{jt}$	$\hat{\hat{e}}_{jt} - \hat{e}^*_{jt}$	Q_{jt}	$\hat{Q}_{jt}$
Full	2180	10.776	10.776	-0.274	-0.274	0.826	0.829
One-sided	2018	10.749	10.749	-0.287	-0.287	0.825	0.827
Two-sided	162	11.114	11.113	-0.113	-0.112	0.849	0.850
Track 1+	52	11.283	11.282	-0.111	-0.111	0.852	0.852
Track 2	25	10.718	10.719	0.088	0.087	0.848	0.845
Track 3	85	11.127	11.126	-0.172	-0.172	0.847	0.851

Notes: This table compares various measures of ACO performance calculated in data and in simulations (denoted with $\hat{\cdot}$). e^*_{jt} is expenditure per beneficiary (in \$ thousand); $\hat{\hat{e}}_{jt} - e^*_{jt}$ is the difference between ideal and actual expenditure per beneficiary (in \$ thousand); Q_{jt} is ACO quality score.

Table 4: **Model Fit: ACO Payment and Exit**

Sample	N	$P_{jt}(\cdot)$	$\hat{P}_{jt}(\cdot)$	$Qual^+_{jt}$	$\hat{Qual}^+_{jt}$	$Qual^-_{jt}$	$\hat{Qual}^-_{jt}$	$Exit_{jt}$	$\hat{Exit}_{jt}$
Full	2180	1.478	1.514	0.328	0.328	0.187	0.187	0.122	0.122
One-sided	2018	1.394	1.432	0.306	0.307	0.189	0.189	0.119	0.120
Two-sided	162	2.527	2.529	0.599	0.593	0.160	0.167	0.154	0.154
Track 1+	52	1.802	1.839	0.635	0.635	0.077	0.077	0.096	0.197
Track 2	25	2.704	2.670	0.720	0.720	0.120	0.120	0.240	0.170
Track 3	85	2.919	2.910	0.541	0.529	0.224	0.235	0.165	0.124

Notes: This table compares various measures of ACO performance calculated in data and in simulations (denoted with $\hat{\cdot}$). $P_{jt}(\cdot)$ is shared savings earned by the ACO (in \$ million); $Qual^+_{jt} = \mathbb{1}\{BE_{jt} - E_{jt} \geq \underline{S}_{jt} BE_{jt}\}$; $Qual^-_{jt} = \mathbb{1}\{BE_{jt} - E_{jt} \leq -\underline{S}_{jt} BE_{jt}\}$; $Exit_{jt}$ indicates ACO exit.

ACO j exited the MSSP after performance year t.[19] Again, simulated means are close to observed means: average ACO earned incentive pay has an error of about 2.4% of mean pay. The proportions of ACOs that qualify for incentive pay and qualify for paying penalties are matched to the hundredths place, with the exception of Track 3 ACOs, which is off by 0.12. Observed exit proportions are also match closely in the full sample and the sub-samples by one-sided and two-sided contracts; however, among two-sided contracts, the model over-predicts exits among Track 1+ ACOs and under predicts for Track 2 and Track 3. This error is likely due to all Track 1+ ACOs in the estimation sample being in their first performance year, as this contract only became available in 2018.

Next, Table 5 shows model fit for variables regarded ACO savings and total program savings. The

[19] Note that $Qual^-_{jt}$ is presented for the one-sided sub-sample as well as the two-sided sub-sample; only the two-sided ACOs actually pay shared losses if expenditure exceeds the MLR.

Table 5: **Model Fit: Program Savings of the MSSP**

Sample	N	SPB_{jt}	$\hat{SPB}_{jt}$	TS_{jt}	$\hat{TS}_{jt}$	PSB_{jt}	$\hat{PSB}_{jt}$	TPS	$\hat{TPS}$
Full	2180	0.144	0.144	1.942	2.008	0.039	0.040	1011	1078
One-sided	2018	0.130	0.130	1.747	1.818	0.031	0.033	712.4	777.4
Two-sided	162	0.312	0.312	4.372	4.384	0.135	0.136	298.9	300.6
Track 1+	52	0.265	0.265	4.268	4.319	0.151	0.151	128.2	128.9
Track 2	25	0.561	0.560	5.482	5.466	0.279	0.281	69.45	69.89
Track 3	85	0.267	0.268	4.110	4.106	0.084	0.084	101.2	101.7

Notes: This table compares various measures of ACO performance calculated in data and in simulations (denoted with $\hat{\cdot}$). $SPB_{jt} = be_{jt} - e^*_{jt}$ is savings per beneficiary (in $ thousand); $TS_{jt} = BE_{jt} - E_{jt}$ is ACO savings (in $ million). PSB_{jt} average program savings per beneficiary (in $ thousands); TPS is total program savings (in $ million) (see Section 6 for definitions).

variable $SPB_{jt} = be_{jt} - e^*_{jt}$ is the savings per beneficiary of ACO j (in $ thousand); $TS_{jt} = BE_{jt} - E_{jt}$ is the total savings of an ACO (in $ million). Savings per beneficiary is matched to the dollar. For ACO total savings, TS_{jt}, simulation error is slightly larger, where in the full sample the average observed ACO savings is $1.942 million, while the simulated average is $2.008 million (a difference of $66,000 per ACO or 3.4%). This error is due to aggregation: simulations are calculated for the variables e^*_{jt}, so error in simulation is magnified when computing E_{jt} by multiplying by the number of beneficiaries n_{jt}.

The term PSB_{jt} is the per-beneficiary program savings of an ACO (in $ thousands), and TPS is the total programs savings across all ACOs in the sample (in $ millions). These variables are computed by subtracting ACO payment from ACO total savings, and represent the monetary savings of the incentive program. Note that subtracting ACO payment from ACO total savings may result in a larger value if the ACO is two-sided and pays a penalty. The total savings variables, along with their quality and exit probability weighted analogs, are the counterfactual simulations of interest presented in Section 6, where they are defined formally and discussed in more detail. For now, they can be interpreted as a representation of Medicare's preferences.

The error of average per beneficiary program savings is less than $10 for all sub-samples. Simulated total program savings is also quite close to the observed value; despite the propagation of error from expenditure per beneficiary to total ACO expenditure, and additional propagation when totalling program savings at the ACO level across ACOs, observed total pogram savings in the full sample is $1.011 billion, while simulated total program savings in the full sample is $1.078 billion: this is a difference of $67 million, or 6% of observed total program savings. Among the sub-samples, error for total program savings is much smaller due to the smaller sample sizes over which program savings is aggregated.

Finally, because the relationship between ACO expenditures and quality scores is at the heart of the theoretical and empirical questions of this paper, Table 6 shows the covariance and correlation matrices of observed and simulated expenditure per beneficiary and quality scores. In both matrices, cells that

Table 6: **Model Fit: Covariation in Expenditure and Quality Score**

Covariance Matrix

	e^*_{jt} Data	e^*_{jt} Simulated	Q^*_{jt} Data	Q^*_{jt} Simulated
e^*_{jt} Data	4.840			
e^*_{jt} Simulated	4.842	4.845		
Q^*_{jt} Data	-0.058	-0.058	0.008	
Q^*_{jt} Simulated	-0.068	-0.069	0.006	0.007

Correlation Matrix

	e^*_{jt} Data	e^*_{jt} Simulated	Q^*_{jt} Data	Q^*_{jt} Simulated
e^*_{jt} Data	1.000			
e^*_{jt} Simulated	**0.999**	1.000		
Q^*_{j} Data	-0.296	-0.294	1.000	
Q^*_{j} Simulated	-0.379	-0.380	**0.819**	1.000

Note: This table compares various relationships of observed and predicted ACO per beneficiary expenditures e^*_{jt} and quality scores Q^*_{jt}. Model fit is indicated by numbers in same-colored cells sharing similar values. Bold numbers are Pearson correlation coefficients of observed and simulated variables.

contain comparable pairs of data and simulated values are shaded the same color. The covariance of e^*_{jt} (in \$ thousands) in the data is 4.840 and in simulations it is 4.854—taking square roots, these values correspond to a standard deviation of expenditure per beneficiary across ACOs of \$2,200 thousand and \$2,201, respectively. Similarly, the values in the covariance matrix indicate the standard deviation of Q^*_j is 0.089 in the estimation sample and 0.084 in the simulation. Comparing the green and yellow boxes in the covariance and correlation matrices shows how well the model simulations can match the relationship of ACO expenditures and quality scores in the data. The bold numbers in the correlation matrix are simply the Pearson correlation coefficients of observed and simulated values.

5.2 Estimates of the Cost-Quality Tradeoff

To address the principal question of this study, I display histograms of ideal per beneficiary expenditure (in \$ thousand) $\hat{\hat{e}}_{jt}$ and actual per beneficiary expenditure e^*_{jt} in Figure 1, and their difference $\hat{\hat{e}}_{jt} - e^*_{jt}$ in Figure 2. The average difference between ideal and actual expenditure is -\$274 per beneficiary, meaning the average ACO, even including the incentives of the MSSP, does not face a cost-quality tradeoff, and can decrease expenditure without decreasing beneficiary health. In the sample, 27.2% of observations have ideal

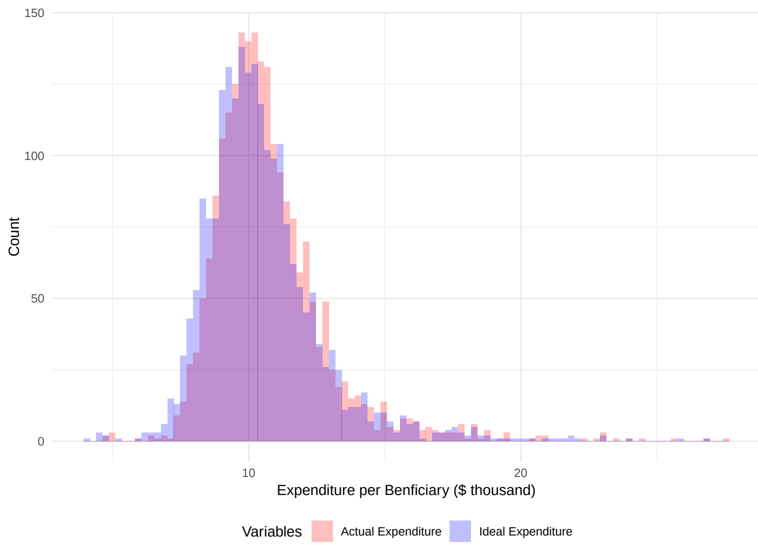

Figure 1: **Histograms of Ideal and Actual per Beneficiary Expenditure**

Note: This figure plots overlapping histograms of ideal expenditure $\hat{\tilde{e}}_{jt}$ and actual expenditure e^*_{jt}. The distribution of ideal expenditure is shifted slightly to the left of actual expenditure.

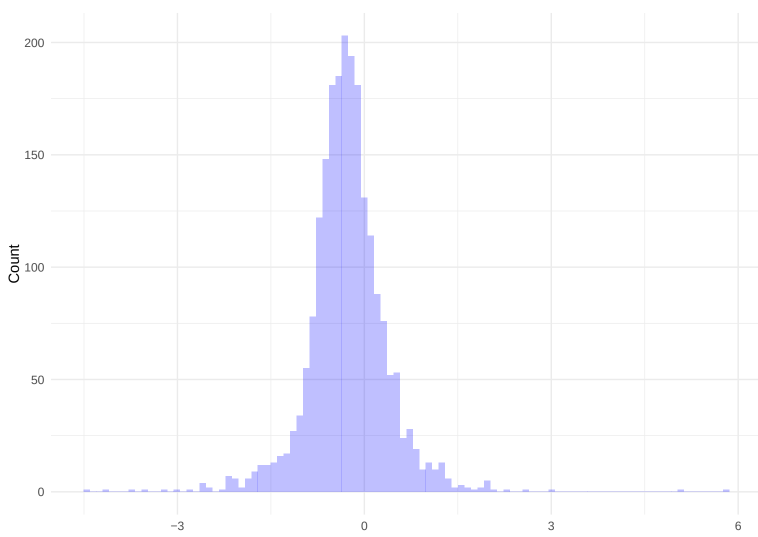

Figure 2: **Histogram of the Cost-Quality Tradeoff**

Note: This figure plots a histogram of of ideal expenditure minus actual expenditure $\hat{\tilde{e}}_{jt} - e^*_{jt}$. Positive values indicate a cost-quality tradeoff. On average, ideal expenditure is $274 less than actual expenditure; health is increasing in expenditure and a cost-quality tradeoff is present for 27.2% of observations in the estimation sample.

expenditures larger than actual expenditures. As indicated in Table 3, two-sided ACOs tend to have closer ideal and actual expenditures, with the difference $\hat{\hat{e}}_{jt} - e^*_{jt}$ about \$170 per beneficiary closer to zero.

6 Counterfactual Simulations

In this section, I use the estimated model of ACO performance to evaluate counterfactual contracts between ACOs and Medicare in the MSSP. The main goal of this section is to simulate ACO performance under stronger savings incentives (a higher sharing rate and two-sided incentives) to see whether the cost-quality tradeoff causes decreases in the quality of care. In other words, I address whether increasing the incentive pay per dollar saved or imposing two-sided contracts that include penalties for over-expenditure drive expenditure per beneficiary low enough to decrease the quality of care measured by ACO quality scores.

This section begins with a simulation of ACO performance under all possible sharing rates $F \in [0,1]$, and with and without two-sided risk $T \in \{0,1\}$. I simulate expenditure per beneficiary, $\hat{e}^*_{jt}(F,T)$, ACO quality score $\hat{Q}_{jt}\left(\hat{e}^*_{jt}; F, T\right)$, as well as exit probability $\hat{\varphi}_{jt}\left(\hat{e}^*_{jt}; F, T\right) \equiv \hat{\Pr}\left\{Exit_{jt} = 1 | \hat{P}_{jt}(\cdot; F, T), \mathbf{x}^{Exit}_{jt}\right\}$. The first results I display are the mean of simulated values of each of these variables, as well as the mean of $\hat{\hat{e}}_{jt} - \hat{e}^*_{jt}(F,T)$, where all ACOs have the same contract (F,T).

After showing simulated expenditures, quality scores, and exit probabilities, I define various objectives as potential representations of Medicare's preferences. The objectives include total program savings, a calculation of monetary savings; quality-weighted program savings, which weights monetary savings by the quality score of an ACO; exit-weighed savings, which weights program savings by the probability the ACO remains in the MSSP; and finally quality-and-exit-weighted savings.

The results regarding two-sided incentives and quality of care are surprising: I find that penalties for overspending actually *increase* quality of care; however, the difference is extremely small. This is because the average ACO spends more than the ideal expenditure, so two-sided incentives uniformly bring ACOs closer to ideal spending. The change is small because of the non-monotone health production function. Going from one-sided to two-sided incentives makes ACOs go from overspending to underspending, but the health production function and accordingly the ACO's quality score yields the same value for symmetric amounts of over and underspending.

Unsurprisingly, penalties for over spending increase program savings per beneficiary by about \$200. This in turn yields an increase in total program savings across all ACOs increases by nearly \$5 billion; however, most of this increase is due to penalty payment or the threat of penalties. By weighing program savings by ACO exit probabilities, one-sided contracts without penalties offer higher expected program savings than two-sided contracts.

The optimal contract parameters of this section are not very different from those found by Aswani et al. (2019). In particular, the optimal sharing rate F for one-sided contracts is in the range of 0.5 to 0.6 in Aswani et al. (2019), and here it is around 0.6 to 0.8, depending on quality weighting. When weighing by the probability of ACO exit, however, optimal values of F I find are much higher, near 0.9.

The counterfactual simulations follow the these steps.

1. Define the functions

$$\hat{Q}_{jt}\left(e_{jt}\right)=\frac{\exp\left(-\frac{1}{2}\left(\hat{\tilde{e}}_{jt}-e_{jt}\right)^{2}-\frac{\hat{\sigma}^{2}}{2n_{jt}}+\hat{\bar{\xi}}_{j}\right)}{1+\exp\left(-\frac{1}{2}\left(\hat{\tilde{e}}_{jt}-e_{jt}\right)^{2}-\frac{\hat{\sigma}^{2}}{2n_{jt}}+\hat{\bar{\xi}}_{j}\right)} \tag{50}$$

$$\hat{\Psi}^{+}_{jt}\left(e_{jt}\right)=\Psi^{+}\left(n_{jt}e_{jt};\hat{\tilde{E}}_{jt}+n_{jt}\hat{\beta}_{jt},\underline{S}_{jt},\sqrt{n_{jt}}\hat{\sigma}\right) \tag{51}$$

$$\text{(and likewise for } \hat{\Psi}^{-}_{jt}\left(e_{jt}\right),\ \hat{\psi}^{+}_{jt}\left(e_{jt}\right),\text{ and } \hat{\psi}^{-}_{jt}\left(e_{jt}\right)) \tag{52}$$

$$\hat{\varphi}_{jt}\left(e_{jt};F,T\right)=\frac{\exp\left(\hat{\delta}_{0}P_{jt}\left(\hat{Q}_{jt}\left(e_{jt}\right),BE_{j}-n_{jt}e_{jt};F,T\right)+\hat{\boldsymbol{\delta}}_{1}'\mathbf{x}^{Exit}_{jt}\right)}{1+\exp\left(\hat{\delta}_{0}P_{jt}\left(\hat{Q}_{jt}\left(e_{jt}\right),BE_{j}-n_{jt}e_{jt};F,T\right)+\hat{\boldsymbol{\delta}}_{1}'\mathbf{x}^{Exit}_{jt}\right)} \tag{53}$$

2. Specify $F\in[0,1]$ and $T\in\{0,1\}$ for the simulation.

3. Set tolerance for fixed point iteration: I set $tol=1\times10^{-6}$.

4. Guess an initial value of per beneficiary expenditure, e^{0}_{jt}.

5. Compute

$$e^{1}_{jt}=\frac{\hat{\alpha}+\hat{\omega}n^{-1}_{jt}F\hat{Q}_{jt}\left(e^{0}_{jt}\right)\left[1-\hat{Q}_{jt}\left(e^{0}_{jt}\right)\right]\left(\hat{\Psi}^{+}_{jt}\left(e^{0}_{jt}\right)-T\hat{\Psi}^{-}_{jt}\left(e^{0}_{jt}\right)\right)}{1+\hat{\alpha}+\hat{\omega}n^{-1}_{jt}F\hat{Q}_{jt}\left(e^{0}_{jt}\right)\left[1-\hat{Q}_{jt}\left(e^{0}_{jt}\right)\right]\left(\hat{\Psi}^{+}_{jt}\left(e^{0}_{jt}\right)-T\hat{\Psi}^{-}_{jt}\left(e^{0}_{jt}\right)\right)}\cdot\tilde{e}_{j} \tag{54}$$

$$+\frac{\hat{\gamma}_{jt}+\hat{\omega}\left(FQ_{jt}\left(e^{0}_{jt}\right)\hat{\psi}^{-}_{jt}\left(e^{0}_{jt}\right)+T\left[1-FQ_{jt}\left(e^{0}_{jt}\right)\right]\hat{\psi}^{-}_{jt}\left(e^{0}_{jt}\right)\right)}{1+\hat{\alpha}+\hat{\omega}n^{-1}_{jt}F\hat{Q}_{jt}\left(e^{0}_{jt}\right)\left[1-\hat{Q}_{jt}\left(e^{0}_{jt}\right)\right]\left(\hat{\Psi}^{+}_{jt}\left(e^{0}_{jt}\right)-T\hat{\Psi}^{-}_{jt}\left(e^{0}_{jt}\right)\right)}. \tag{55}$$

6. Compute $\Delta_{jt}=|e^{1}_{jt}-e^{0}_{jt}|$.

   (a) If $\Delta_{jt}>tol$, replace $e^{0}_{jt}=e^{1}_{jt}$ and go back to step 5.

   (b) If $\Delta_{jt}<=tol$, set $\hat{e}^{*}_{jt}=e^{1}_{jt}$ and go to step 7.

7. Simulated quality score and exit probabilities are

$$\hat{Q}_{jt}\left(\hat{e}^{*}_{jt};F,T\right)\ , \tag{56}$$

$$\hat{\varphi}_{jt}\left(\hat{e}^{*}_{jt};F,T\right)\ . \tag{57}$$

8. Compute counterfactual ACO performance measures: program savings TPS, quality-weighted program savings PSQ, exit-weighted program savings PSX, and quality-and-exit-weighted program savings $PSQX$ (defined below). Theses values are taken to represent different specifications of Medicare's preferences for cost reduction and quality of care as the principal in the optimal contracting problem,

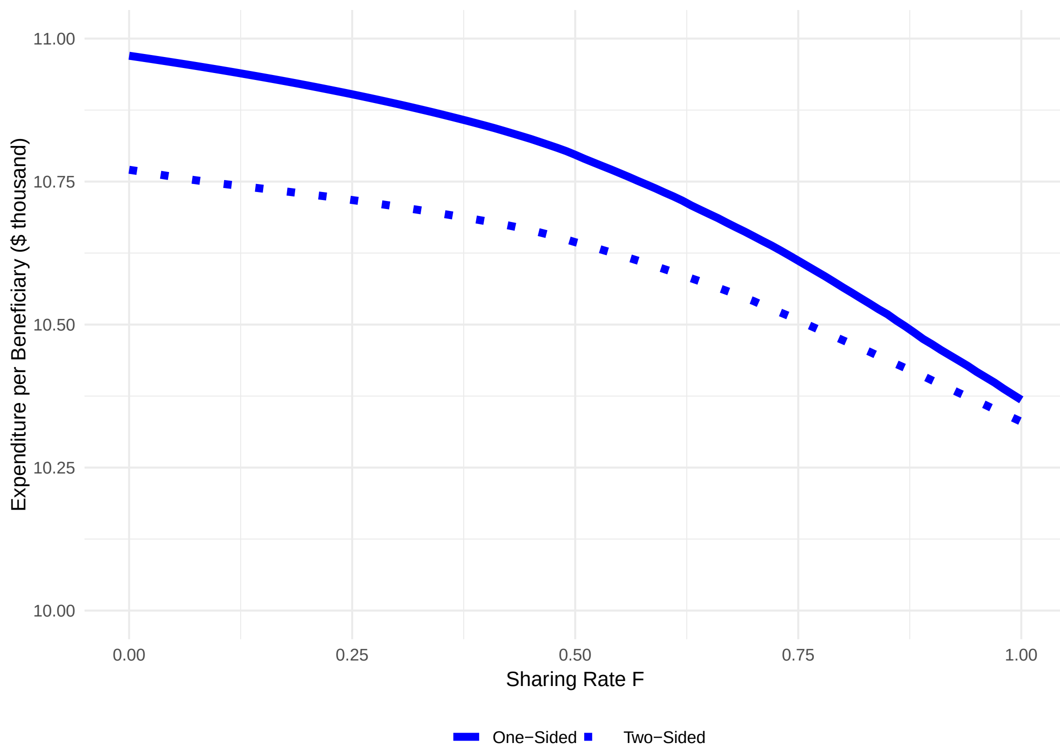

Figure 3: **Simulated Expenditure per Beneficiary**

Note: This figure plots simulated average per beneficiary expenditure across ACOs, e^*_{jt}.

and each is defined formally as a function of counterfactual performance values $\hat{e}^*_{jt}(F,T)$, $\hat{Q}_{jt}(\hat{e}^*_{jt}F,T)$, and $\hat{\varphi}_{jt}\left(\hat{e}^*_{jt};F,T\right)$ in the following subsection.

6.1 Simulations of Expenditure, Quality Score, and ACO Exit

Figures 3, 4, 5, and 6 show the simulated values of per beneficiary expenditure, ideal minus actual expenditure, quality score, and exit probability for values of $F \in [0,1]$ (horizontal axis) and $T \in \{0,1\}$ (solid and dashed lines, respectively). Figure 3 shows that the effect of two-sided contracts, should they be imposed on all ACOs, is to decrease per-beneficiary expenditure by \$199 ($F = 0$) to \$38 ($F = 1$) relative to one-sided contracts. Two-sided contracts achieve this additional savings across all ACOs through the threat of penalization for over-expenditure.

Figure 4 shows that $\hat{\tilde{e}}_{jt} - \hat{e}^*_{jt}(F,T)$ steadily grows larger for larger values of F. Furthermore, because of penalties imposed by two-sided contracts, a lower rate of shared savings is required to to achieve average beneficiary expenditure closer to ideal expenditure. Figure 4 also indicates values of F and T that are plausibly optimal from the perspective of Medicare or an average health-maximizing social planner. Specifically, $(F,T) = (0.88,0)$ and $(F,T) = (0.77,1)$ both implement $\hat{e}^*_{jt} = \tilde{e}_{jt}$ on average across ACOs.

Figure 5 displays simulated ACO quality scores. Two-sided ACOs have slightly higher quality scores on

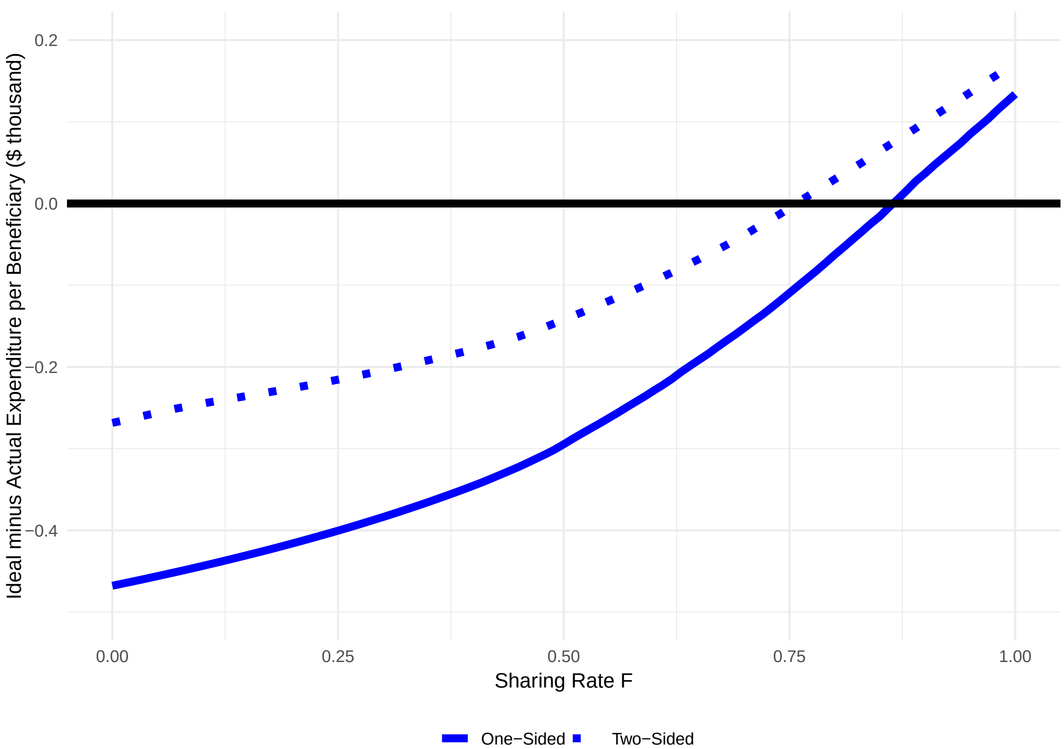

Figure 4: **Simulated Ideal minus Actual Expenditure per Beneficiary**

Note: This figure plots the average ideal expenditure per beneficiary minus simulated expenditure per beneficiary across ACOs, $\hat{\hat{e}}_{jt} - e^*_{jt}$. The intersection at $\hat{\hat{e}}_{jt} = e^*_{jt}$ occurs at $F = 0.88$ for one-sided contracts and $F = 0.77$ for two-sided contracts.

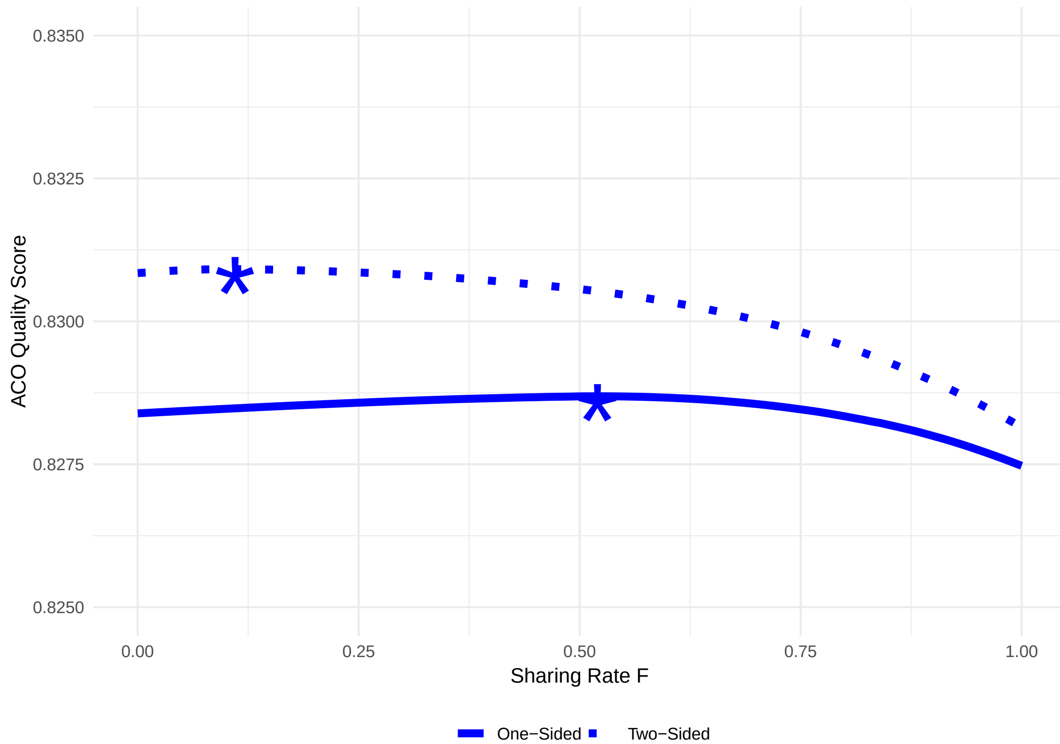

Figure 5: **Simulated ACO Quality Scores**

Note: This figure plots the average simulated quality score across ACOs, $\hat{Q}_{jt}$. The maximum occurs at $F = 0.53$ for one-sided contracts and $F = 0.12$ for two-sided contracts.

average. This is ultimately because the penalties on ACOs for overspending push expenditures closer to the ideal level. Despite expenditure for two-sided ACOs at times decreasing below ideal expenditure, the effect, on average, is still expenditures that are closer to the ideal level than one-sided contracts. Because health is non-monotone in the difference $\tilde{e}_{jt} - e_{jt}$, being slightly above or slightly below ideal expenditure has the same effect on quality score. This ultimately leads to very small changes in simulated quality score for different values of F and T.

Note that average quality score for two-sided contracts decreases above a low sharing rate of $F = 0.12$; likewise, one-sided contracts are high-powered enough to decrease quality score above $F = 0.53$. The change in quality score is so small, however, that this will not effect total program savings after weighting by quality score. These values that maximize quality score are different than those that impose $\hat{\tilde{e}}_{jt} = e^*_{jt}$. This is because average quality score is maximized where the *average* of $\left(\hat{\tilde{e}}_{jt} - e^*_{jt}\right) \cdot \hat{Q}_{jt} \left(1 - \hat{Q}_{jt}\right)$ is equal to zero.

Finally, Figure 6 displays simulated exit probabilities for one-sided and two-sided contracts. On average, the probability of exit increases by 0.06 to 0.13 when imposing two-sided contracts. Approximately one-in-five ACOs are predicted to exit the MSSP should two-sided contracts be imposed; otherwise, the rate of exit is about one-in-ten ACOs for one-sided contracts.

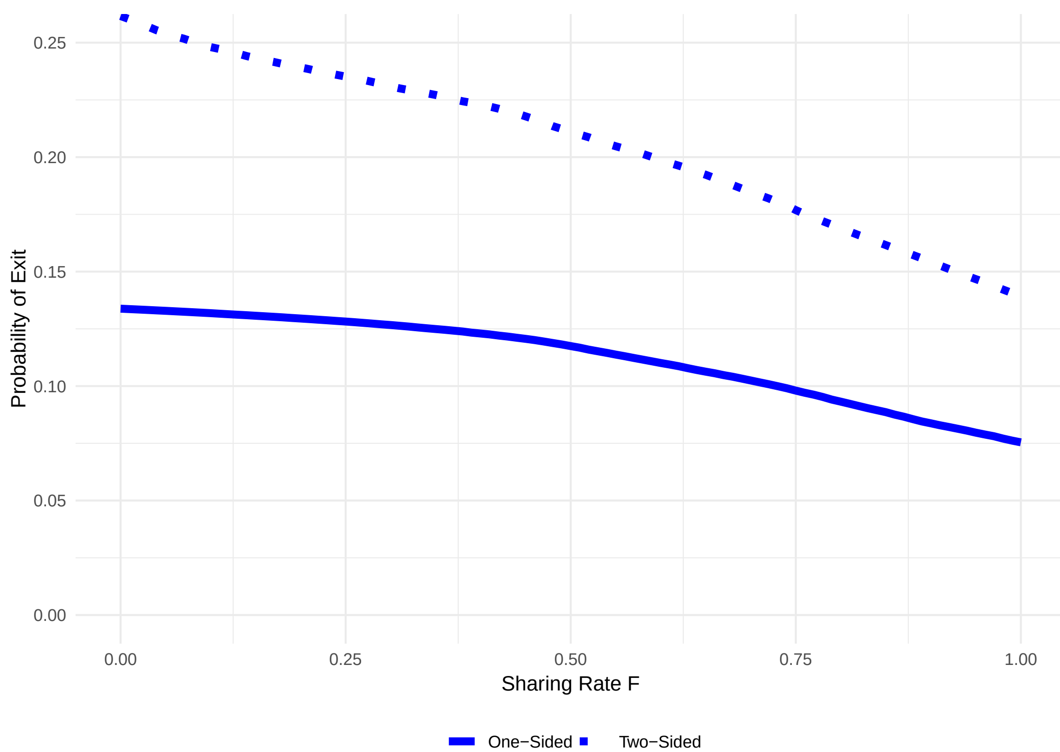

Figure 6: **Simulated ACO Exit Probabilities**

Note: This figure plots the average simulated exit probability of ACOs, $\hat{\varphi}_{jt}$. On average, two-sided contracts increase the probability of exit by 13 percentage points (at $F = 0$) to 6 percentage points (at $F = 1$).

6.2 MSSP Contract Comparisons

Here, I model incentive design in the MSSP as an optimal contracting problem and solve for contracts that maximize a wide range of objectives. This model departs from typical models of incomplete contracting in several ways. First, there are no actions or information that are necessarily hidden— rather, contracts are incomplete and do not implement first best outcomes due to regulatory constraints on the manner in which ACOs can be paid through the MSSP. Accordingly, I limit counterfactual simulations to varying the sharing rate $F \in [0,1]$ and presence of penalties $T \in \{0,1\}$.

Second, individual ACO contracts are not offered by the MSSP. Rather, the MSSP offers a limited menu of contracts that ACOs can choose from, and does not define a specific contract for a specific ACO. Accordingly, I solve for optimal contracts across the sample of ACOs, finding the contract that maximizes an objective which is computed from the performance of all ACOs in the sample. This also means individual rationality constraints are not included in simulations, as requiring that ACO pay not be negative for optimal two-sided contracts *for all ACOs in the sample* results in a solution that is determined by the lowest value of F_{jt} for which the worst-performing simulated two-sided ACO does not earn negative pay.[20]

The objective functions I define should be interpreted as representations of the preferences of Medicare, who is the principal in the optimal contracting problem that is designing ACO payment models in the MSSP. The first definition is total program savings, TPS, defined as the sum of all ACO savings (according to the defined benchmark expenditure BE_j) minus the sum of all ACO incentive pay. Formally, let $\left\{\hat{e}^*_{jt}\left(F,T\right), \hat{Q}_{jt}\left(\hat{e}^*_{jt}; F, T\right)\right\}_{j\in\mathcal{J}}$ be the simulated expenditure per beneficiary and quality scores of the sample of ACOs under counterfactual sharing rate $F \in [0,1]$ and risk model $T \in \{0,1\}$. Total program savings is therefore

$$
\begin{aligned}
&TPS\left(\left\{\hat{e}^*_{jt}\left(F,T\right), \hat{Q}_{jt}\left(\hat{e}^*_{jt}; F, T\right)\right\}_{j\in\mathcal{J}}\right) \\
&\qquad = \sum_{j\in\mathcal{J}}\left[BE_{jt} - n_{jt}\hat{e}^*_{jt}(F,T) - P_{jt}\left(\hat{Q}_{jt}\left(e^*_{jt}; F, T\right), BE_j - n_{jt}e^*_{jt}; F, T\right)\right]. \qquad (58)
\end{aligned}
$$

The first term in the summation, $BE_{jt} - n_{jt}\hat{e}^*_{jt}(F,T)$, is the predicted dollar amount of savings (or losses if negative) of ACO j. The second term, $P_{jt}\left(\hat{Q}_{jt}\left(e^*_{jt}; F, T\right), BE_j - n_{jt}e^*_{jt}; F, T\right)$, is the payment to the ACO. According to this definition, Medicare's objective is to maximize the total dollar savings to Medicare for the ACOs in the simulation sample $\mathcal{J}$.

Medicare's preferences, as defined by TPS, has two shortcomings. First, TPS depends on quality score $\hat{Q}_{jt}\left(e^*_{jt}; F, T\right)$ only through ACO payment $\hat{P}_{jt}(\cdot)$, and so it is *decreasing* in quality score. Second, particularly for two-sided contracts, ACOs may exit the MSSP if they earn sufficiently low incentive pay. This means TPS would predict high program savings for two-sided ACOs, where it would in fact be unlikely that those

[20] There are in fact several ACOs that, even with $F = 1$, earn negative incentive pay. This is because, despite the large incentive to decrease expenditure, responsiveness to incentive pay (measured by ω) is not large enough to reduce expenditure to the point penalties are not paid.

ACOs would remain the the MSSP (or have joined in the first place) if such contracts were imposed.

To address these problems, I define quality-weighted program savings PSQ, exit-weighted program savings PSX, and quality-and-exit-weighted program savings $PSQX$:

$$PSQ\left(\left\{\hat{e}^*_{jt}\left(F,T\right),\hat{Q}_{jt}\left(\hat{e}^*_{jt};F,T\right)\right\}_{j\in\mathcal{J}}\right) = \sum_{j\in\mathcal{J}}\left\{\left[BE_{jt}-n_{jt}\hat{e}^*_{jt}(F,T)-P_{jt}\left(\hat{Q}_{jt}\left(e^*_{jt};F,T\right),BE_j-n_{jt}e^*_{jt};F,T\right)-\underline{\Upsilon}\right]\hat{Q}_{jt}\left(\hat{e}^*_{jt};F,T\right)+\underline{\Upsilon}\cdot\overline{\hat{Q}}\right\} \tag{59}$$

$$\begin{aligned}PSX&\left(\left\{\hat{e}^*_{jt}\left(F,T\right),\hat{Q}_{jt}\left(\hat{e}^*_{jt};F,T\right),\hat{\varphi}_{jt}\left(\hat{e}^*_{jt};F,T\right)\right\}_{j\in\mathcal{J}}\right)\\ &=\sum_{j\in\mathcal{J}}\left\{\left[BE_{jt}-n_{jt}\hat{e}^*_{jt}(F,T)-P_{jt}\left(\hat{Q}_{jt}\left(e^*_{jt};F,T\right),BE_j-n_{jt}e^*_{jt};F,T\right)-\underline{\Upsilon}\right]\right.\\ &\qquad\left.\times\left[1-\hat{\varphi}_{jt}\left(\hat{e}^*_{jt};F,T\right)\right]+\underline{\Upsilon}\cdot\overline{\hat{\varphi}}\right\}\end{aligned} \tag{60}$$

$$\begin{aligned}PSQX&\left(\left\{\hat{e}^*_{jt}\left(F,T\right),\hat{Q}_{jt}\left(\hat{e}^*_{jt};F,T\right),\hat{\varphi}_{jt}\left(\hat{e}^*_{jt};F,T\right)\right\}_{j\in\mathcal{J}}\right)\\ &=\sum_{j\in\mathcal{J}}\left\{\left[BE_{jt}-n_{jt}\hat{e}^*_{jt}(F,T)-P_{jt}\left(\hat{Q}_{jt}\left(e^*_{jt};F,T\right),BE_j-n_{jt}e^*_{jt};F,T\right)-\underline{\Upsilon}\right]\right.\\ &\qquad\left.\times\hat{Q}_{jt}\left(\hat{e}^*_{jt};F,T\right)\left[1-\hat{\varphi}_{jt}\left(\hat{e}^*_{jt};F,T\right)\right]+\underline{\Upsilon}\cdot\overline{\hat{Q}\hat{\varphi}}\right\}.\end{aligned} \tag{61}$$

The quality-weighted measures multiply the savings of each ACO by quality score before summing over the observations in the sample. Similarly the exit probability weighted measures multiply the savings of each ACO by the probability they do not exit the MSSP before summing over observations. The term $\underline{\Upsilon}=\min_{j\in\mathcal{J},F\in[0,1],T\in\{0,1\}}BE_{jt}-n_{jt}\hat{e}^*_{jt}(\cdot)-\hat{P}_{jt}(\cdot)$ is a normalizing value chosen so that only positive values are multiplied by $\hat{Q}_{jt}(\cdot)$ and $\hat{\varphi}_{jt}(\cdot)$. Normalization is necessary to make the quality-weighted objectives increasing in ACO quality score for all values of expenditure—without normalization, PSQ and $PSQX$ would be decreasing in $\hat{Q}_{jt}(\cdot)$ if $\hat{e}^*_{jt}$ is sufficiently high. Normalization is also performed for the exit-weighted objective.[21]

The quality-weighted objectives can be interpreted as the savings of ACOs such that every one dollar of cost reduction of ACO j is worth $\hat{Q}_{jt}\in(0,1)$ dollars to Medicare. Likewise, the exit-weighted objectives are (approximately) expected program savings, where with probability $1-\hat{\varphi}_{jt}(\cdot)$ the ACO did not join the MSSP because they anticipated low (or negative) incentive pay.

The contract parameters (F,T) that maximize total program savings are *not* the same as the parameters

[21] Note that Υ is multiplied by $\overline{\hat{Q}}$, $\overline{\hat{\varphi}}$, or $\overline{\hat{Q}\hat{\varphi}}$ and added back onto weighted-savings. $\overline{\hat{Q}}$ is the mean of simulated quality score across all simulations of F and T, and is similarly defined for $\overline{\hat{\varphi}}$, or $\overline{\hat{Q}\hat{\varphi}}$. This adjustment is performed so that the levels of PSQ, PSX, and $PSQX$ are roughly the same as TPS.

that maximize *per-beneficiary* savings. This is because ACOs that have the largest cost reductions have disproportionately large number of assigned beneficiaries. In other words, the optimal F and T from the above defined objectives place more weight on ACOs with a larger number of beneficiaries. Therefore, contract parameters chosen to maximize total savings favor ACOs that have large numbers of beneficiaries, while contract parameters chosen to maximize per-beneficiary savings (at the ACO level) may be more representative of optimal contracts for the average Medicare provider. For this reason, I present average per-beneficiary savings for an objective:

$$PSB\left(\left\{\hat{e}^*_{jt}\left(F,T\right),\hat{Q}_{jt}\left(\hat{e}^*_{jt};F,T\right)\right\}_{j\in\mathcal{J}}\right) = |\mathcal{J}|^{-1}\sum_{j\in\mathcal{J}}\left[\frac{BE_{jt}-n_{jt}\hat{e}^*_{jt}(F,T)-P_{jt}\left(\hat{Q}_{jt}\left(e^*_{jt};F,T\right),BE_j-n_{jt}e^*_{jt};F,T\right)}{n_{jt}}\right]. \quad (62)$$

Quality-weighted and exit-weighted per-beneficiary objectives are presented as well.

Figure 7 plots the objective functions for all four of of the different preference specifications for Medicare. Each line in the figure is a graph of F vs. TPS, PSQ, PSX, or $PSQX$ for $F\in[0,1]$, $T=0$ and $T=1$. Figure 8 shows the per-beneficiary average program savings.

6.2.1 Optimal Sharing Rates

Optimal sharing rates for the various objectives are marked in Figure 7 by stars. These values range from a sharing rate of 59% (two-sided quality-weighted) to a sharing rate of 91% (two-sided exit-weighted), which are higher-powered than most existing MSSP contracts that offer a sharing rate of only 50%, 60% , or 75%. When not weighting by ACO exit probability, optimal contracts are weaker. However, due to the increase in exit probability for lower earned incentive pay, exit-weighted objectives push the optimal sharing rate to 88% and above. Weighting by quality score has the overall effect of shifting the levels of these objective functions, but the shape of the objective functions are mostly unchanged. Conditional on exit-weighting, optimal sharing rates are hardly changed between quality weighted and non-quality weighted objectives.

The per beneficiary objectives in Figure 8 yield similar solutions for all objectives, except for two-sided contracts *without* exit-weighting. Specifically, the per-beneficiary average savings is maximized for two-sided contracts at very low values of 8% (two-sided) and 14% (two-sided quality weighted). This is due to both decreased expenditures due to the threat of penalty and actual penalties paid for overspending. When ACO exit is accounted for, the optimal sharing rate changes drastically, and becomes as high as 99% (two-sided exit-weighted).

6.2.2 One-Sided vs. Two-Sided Contracts

Implementing two-sided contracts, where health care providers are penalized for overspending, is a primary goal of policy-makers involved with the MSSP. These simulations indicate why: ACOs are predicted to save

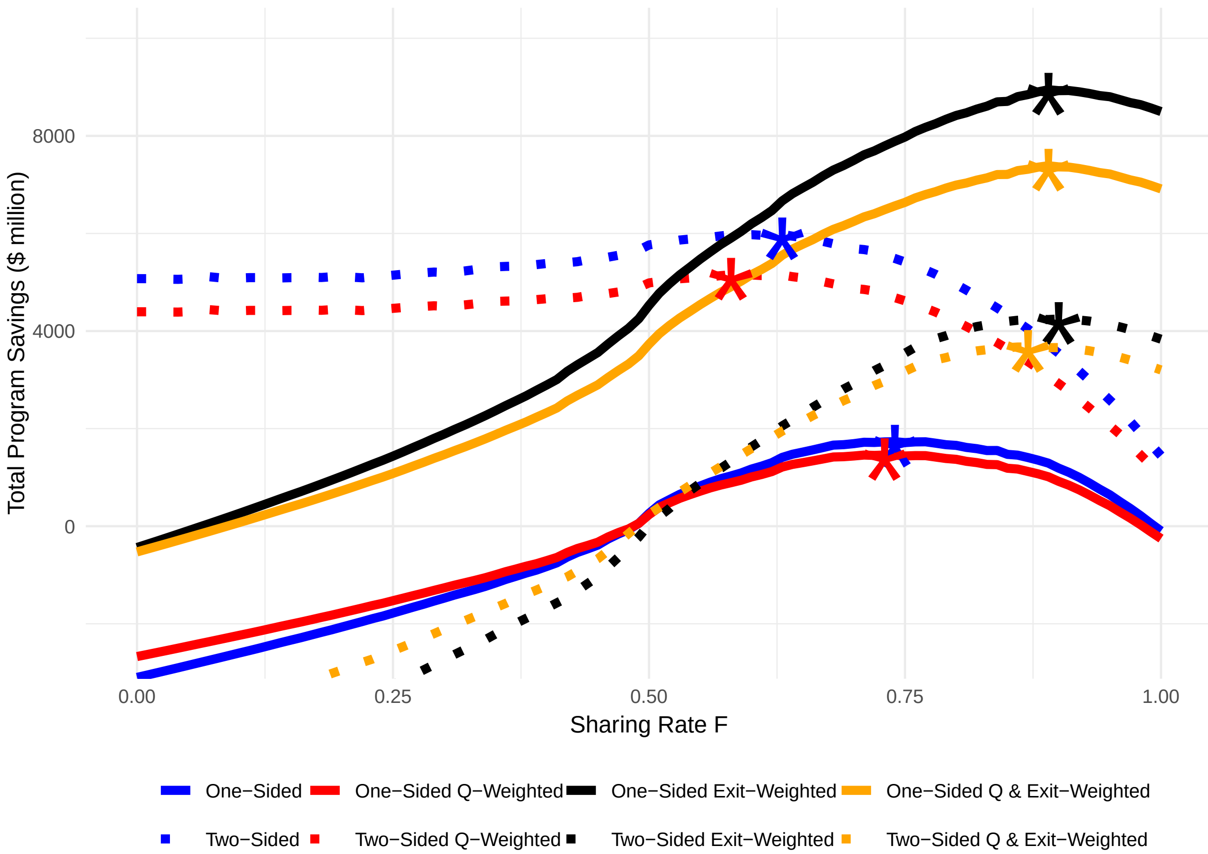

Figure 7: **Simulated Total Program Savings**

Note: This figure plots simulated total program savings. Simulations vary MSSP contracts by sharing rate (horizontal axis) and presence of penalties (solid line is without penalties, dotted line is with penalties). Objectives vary by color depending on their quality score or exit probability weight: blue is no weighting, red is quality score weighting, black is exit probability weighting, and gold is quality score and exit probability weighting. Stars are placed at the maximum of their respective line.

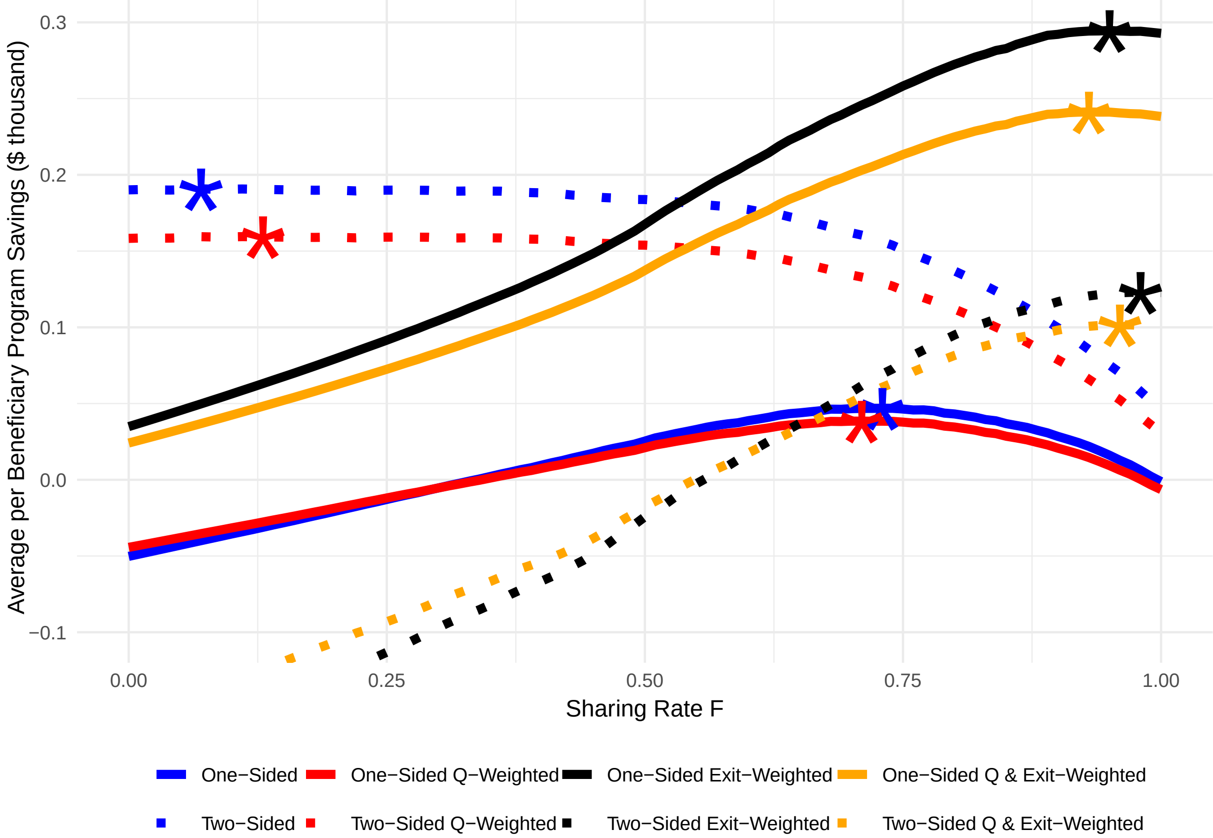

Figure 8: **Simulated Average Savings per Beneficiary**

Note: This figure plots simulated average program savings per beneficiary. Simulations vary MSSP contracts by sharing rate (horizontal axis) and presence of penalties (solid line is without penalties, dotted line is with penalties). Objectives vary by color depending on their quality score or exit probability weight: blue is no weighting, red is quality score weighting, black is exit probability weighting, and gold is quality score and exit probability weighting. Stars are placed at the maximum of their respective line.

substantially more when facing downside risk. Specifically, total program savings would increase by over \$4 billion in the estimation sample if penalties were imposed on all providers. Per-beneficiary spending decreases by \$240 ($F = 0$) to \$47 ($F = 1$). At $F = 0.5$, a 50% sharing rate, savings per beneficiary increases by \$158, which is a 0.25 standard deviation decrease.

When exit-weighting is not used, both quality-weighted and non-quality-weighted objectives are larger for two-sided contracts. However, exit-weighting the objectives reverses this result: even with high-powered (larger F) contracts, the PSX and $PSQX$ are much higher for one-sided ACOs due to their lower probability of exit from not facing downside risk.

7 Conclusion

Incentive design is used by firms, governments, households, educators, and many others to achieve a variety of ends. Accordingly, designing effective incentives is a popular topic of study in all fields of economics. It is particularly important in the United States health care sector, where physician incentive programs and pay-for-performance initiatives impact the quality of life and spending of individuals in 3.5 trillion dollar industry. In this paper, I investigate the empirical role of the cost-quality tradeoff the design of physician incentive programs, specifically in the context of the Medicare Shared Savings Program and Accountable Care Organizations. I identify the cost-savings tradeoff experienced by ACO providers and accordingly design MSSP contracts that maximize the quality-weighted incentive program savings, while accounting for free-riding of health care providers induced by ACO payment.

I estimate a model of Medicare expenditure and simulate the program savings of the MSSP under counterfactual contracts. I find that most ACOs do not face a cost-quality tradeoff, and that two-sided incentives likely will not decrease the quality of care provided to Medicare beneficiaries. Counterfactual policy analysis shows that if ACOs are required to pay penalties to Medicare for spending too much, savings increases modestly, though participation in the program drops significantly. Without accounting for voluntary participation, optimal contracts are two-sided and moderately powered with a sharing rate around 60-70%. However, adjusting by the probability of ACO exit, optimal contracts are one-sided and much higher powered with a sharing rate around 90%.

Appendix

A On Existence and Uniqueness of Equilibrium

This appendix derives conditions under which game in Section 3 has a unique equilibrium. The proof provided shows that this is a concave game, as in Rosen (1965), and thus equilibrium exists and is unique.

A.1 Provider Expenditure Choice without Incentive Pay

Recall from Section 3 that providers $i \in \mathcal{I}_a$ simultaneously solve

$$\max_{e_{ia}>0} \pi_i\left(e_{ia}\right) + \alpha_i h(\tilde{e}_a - e_{ia} - e_{-ia}), \tag{63}$$

where $\pi_i(e_{ia}) = e_{ia} - \frac{e_{ia}^2}{2\gamma_i}$ and $h\left(\tilde{e}_a - e_a\right) = -\frac{1}{2}\left(\tilde{e}_a - e_a\right)^2$, and $\gamma_i, \alpha_i > 0$. The first order condition for provider expenditure is

$$1 - \frac{e_{ia}^*}{\gamma_i} + \alpha_i\left(\tilde{e}_a - e_{ia}^* - e_{-ia}\right) = 0 \tag{64}$$

The second order condition that implies strict concavity is,

$$-\gamma_i^{-1} - \alpha_i < 0\,; \tag{65}$$

this is always true from the assumption that $\alpha_i, \gamma_i > 0$.

To solve this for closed forms of e_{ia}^* and e_a^* (Equations 11 and 12), first assume all providers $i \in \mathcal{I}_a$ choose expenditure e_{ia}^* that solves Equation 64, and let $e_a^* = \sum_{i\in\mathcal{I}_a} e_{ia}^*$. Then

$$\gamma_i - e_{ia}^* + \alpha_i\gamma_i\left(\tilde{e}_a - e_a^*\right) = 0 \tag{66}$$

$$\implies \frac{\gamma_i - e_{ia}}{\alpha_i\gamma_i} = -\left(\tilde{e}_a - e_a^*\right) \tag{67}$$

Then, for providers $i, \ell \in \mathcal{I}_a$,

$$\frac{\gamma_i - e_{ia}^*}{\alpha_i\gamma_i} = \frac{\gamma_\ell - e_{\ell a}^8}{\alpha_\ell\gamma_\ell} \tag{68}$$

This means

$$e_{\ell a}^* = \gamma_\ell - \alpha_\ell\gamma_\ell \cdot \frac{\gamma_i - e_{ia}^*}{\alpha_i\gamma_i} \tag{69}$$

Summing over $\ell \in \mathcal{I}_a$:

$$e_a^* = \sum_{\ell \in \mathcal{I}_a} \gamma_\ell - \left(\sum_{\ell \in \mathcal{I}_a} \alpha_\ell \gamma_\ell \right) \frac{\gamma_i - e_{ia}^*}{\alpha_i \gamma_i} \tag{70}$$

Plugging this in to the FOC:

$$\gamma_i - e_{ia}^* + \alpha_i \gamma_i \left[\tilde{e}_a - \sum_{\ell \in \mathcal{I}_a} \gamma_\ell + \left(\sum_{\ell \in \mathcal{I}_a} \alpha_\ell \gamma_\ell \right) \frac{\gamma_i - e_{ia}^*}{\alpha_i \gamma_i} \right] = 0 \tag{71}$$

$$\gamma_i - e_{ia}^* + \alpha_i \gamma_i \tilde{e}_a - \alpha_i \gamma_i \sum_{\ell \in \mathcal{I}_a} \gamma_\ell + \left(\sum_{\ell \in \mathcal{I}_a} \alpha_\ell \gamma_\ell \right) (\gamma_i - e_{ia}^*) = 0 \tag{72}$$

$$\gamma_i \left(1 + \sum_{\ell \in \mathcal{I}_a} \alpha_\ell \gamma_\ell \right) - e_{ia}^* \left(1 + \sum_{\ell \in \mathcal{I}_a} \alpha_\ell \gamma_\ell \right) + \alpha_i \gamma_i \tilde{e}_a - \alpha_i \gamma_i \sum_{\ell \in \mathcal{I}_a} \gamma_\ell = 0 \tag{73}$$

$$\tag{74}$$

Finally, isolating e_{ia}^*:

$$e_{ia}^* \left(1 + \sum_{\ell \in \mathcal{I}_a} \alpha_\ell \gamma_\ell \right) = \gamma_i \left(1 + \sum_{\ell \in \mathcal{I}_a} \alpha_\ell \gamma_\ell \right) + \alpha_i \gamma_i \left[\tilde{e}_a - \sum_{\ell \in \mathcal{I}_a} \gamma_\ell \right] \tag{75}$$

$$e_{ia}^* = \gamma_i + \frac{\alpha_i \gamma_i}{1 + \sum_{\ell \in \mathcal{I}_a} \alpha_\ell \gamma_\ell} \left(\tilde{e}_a - \sum_{\ell \in \mathcal{I}_a} \gamma_\ell \right) . \tag{76}$$

For e_a^*, simply sum over $i \in \mathcal{I}_a$:

$$e_a^* = \sum_{i \in \mathcal{I}_a} \gamma_i + \frac{\sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i}{1 + \sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i} \left(\tilde{e}_a - \sum_{i \in \mathcal{I}_a} \gamma_i \right) \tag{77}$$

$$e_a^* = \frac{\sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i}{1 + \sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i} \tilde{e}_a + \left(1 - \frac{\sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i}{1 + \sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i} \right) \sum_{i \in \mathcal{I}_a} \gamma_i \tag{78}$$

$$e_a^* = \frac{\sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i}{1 + \sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i} \tilde{e}_a + \frac{\sum_{i \in \mathcal{I}_a} \gamma_i}{1 + \sum_{i \in \mathcal{I}_a} \alpha_i \gamma_i} . \tag{79}$$

Because optimal provider expenditure choices are guaranteed to exist, be unique, and be strictly positive when $\alpha_i, \gamma_i > 0$, then $(e_{ia}^*)_{i \in \mathcal{I}_a}$ is the pure strategy equilibrium of this game.

A.2 Provider Expenditure with Incentive Pay

Characterizing Nash equilibrium pay of beneficiaries assigned to ACOs isn't as straightforward as the game without incentive pay. The game without incentive pay is played among the care team of a, $i \in \mathcal{I}_a$, while the game for a beneficiary assigned to an ACO is among all $i \in \cup_{a \in \mathcal{A}_j} \mathcal{I}_a$. This is because ACO expenditure is the sum of expenditure across all beneficiaries assigned to ACOs, and beneficiary expenditure is the sum of expenditure across their care team: $E_j = \sum_{a \in \mathcal{A}_j} e_a^* = \sum_{a \in \mathcal{A}_j} \sum_{i \in \mathcal{I}_a} e_{ia}^*$.

Nonetheless, for the purposes of this study it will suffice to show that an optimal e^*_{ia} exists and is unique conditional on the expenditures of other providers $(e_{\ell a})_{\ell \in \left\{\cup_{a \in \mathcal{A}_j} \mathcal{I}_a\right\} \setminus \{i\}}$. This is of course done by showing the objective

$$\max_{e_{ia}} \pi_i(e_{ia}) + \alpha_i h\left(\tilde{e}_a - e_a\right) + \omega_{ij} \mathbb{E}\left[P_j\left(Q_j, BE_j - E_j; F_j, T_j\right)\right] \tag{80}$$

is strictly concave, where again

$$\begin{aligned} &\mathbb{E}\left[P_j\left(Q_j, BE_j - E_j; F_j, T_j\right)\right] \\ &\qquad = F_j Q_j \Psi^+\left(E_j; \tilde{E}_j + n_j\beta_j, \underline{S}_j, \sqrt{n_j}\sigma\right) + T_j\left(1 - F_j Q_j\right)\Psi^-\left(E_j; \tilde{E}_j + n_j\beta_j, \underline{S}_j, \sqrt{n_j}\sigma\right). \end{aligned} \tag{81}$$

and

$$\Psi^+\left(E_j; \tilde{E}_j + n_j\beta_j, \underline{S}_j, \sqrt{n_j}\sigma\right) = \left(\tilde{E}_j + n_j\beta_j - E_j\right)\Phi\left(Z_j^+\right) + \sqrt{n_j}\sigma\phi\left(Z_j^+\right) \equiv \Psi_j^+ \tag{82}$$

$$\Psi^{+\prime}\left(E_j; \tilde{E}_j + n_j\beta_j, \underline{S}_j, \sqrt{n_j}\sigma\right) = -\Phi\left(Z_j^+\right) - \frac{E_j\underline{S}_j}{\sqrt{n_j}\sigma\left(1 - \underline{S}_j\right)^2}\phi\left(Z_j^+\right) \equiv \psi_j^+ \tag{83}$$

$$\Psi^-\left(E_j; \tilde{E}_j + n_j\beta_j, \underline{S}_j, \sqrt{n_j}\sigma\right) = \left(\tilde{E}_j + n_j\beta_j - E_j\right)\Phi\left(Z_j^-\right) + \sqrt{n_j}\sigma\phi\left(Z_j^-\right) \equiv \Psi_j^- \tag{84}$$

$$\Psi^{-\prime}\left(E_j; \tilde{E}_j + n_j\beta_j, \underline{S}_j, \sqrt{n_j}\sigma\right) = -\Phi\left(Z_j^-\right) - \frac{E_j\left(2 + \underline{S}_j\right) - 2\left(\tilde{E}_j + n_j\beta_j\right)\left(1 + \underline{S}_j\right)}{\sqrt{n_j}\sigma\left(1 + \underline{S}_j\right)^2}\phi\left(Z_j^-\right) \equiv \psi_j^- \tag{85}$$

for $Z_j^+ = \left(\tilde{E}_j + n_j\beta_j - \frac{E_j}{1-\underline{S}_j}\right) \Big/ \sqrt{n_j}\sigma$ and $Z_j^- = \left(\frac{E_j}{1+\underline{S}_j} - \tilde{E}_j - n_j\beta_j\right) \Big/ \sqrt{n_j}\sigma$. Thus, I must establish

$$-\gamma_i^{-1} - \alpha_i + \omega_{ij}\frac{\partial^2}{\partial e_{ia}^2}\mathbb{E}\left[P_j\left(Q_j, BE_j - E_j; F_j, T_j\right)\right] < 0\,. \tag{86}$$

This is obviously true for any $i \notin \mathcal{I}_j \cap \mathcal{I}_a$, as in this case $\omega_{ij} = 0$ or $e^*_{ia} \equiv 0$ by definition. For the remaining providers that are both on a's care team and a member of the ACO to which a is assigned, this inequality holds if $\omega_{ij}\frac{\partial^2}{\partial e_{ia}^2}\mathbb{E}\left[P_j(\cdot)\right]$ is sufficiently small in comparison to $\gamma_i^{-1} + \alpha_i$. Furthermore, a larger value of $\sqrt{n_j}\sigma$ means $\frac{\partial^2}{\partial e_{ia}^2}\mathbb{E}\left[P_j(\cdot)\right]$ is smaller. Both conditions almost certainly hold given model estimates: in particular, $\hat{\sigma}$ is large.

References

Agha, L., Frandsen, B., & Rebitzer, J. B. (2019). Fragmented division of labor and healthcare costs: Evidence from moves across regions. *Journal of Public Economics*, *169*, 144-159.

Aswani, A., Shen, Z.-J. M., & Siddiq, A. (2019). Data-driven incentive design in the medicare shared savings program. *Operations Research*, *67*(4), 1002-1026.

Chandra, A., Cutler, D., & Song, Z. (2011). Chapter six - who ordered that? the economics of treatment choices in medical care. In M. V. Pauly, T. G. Mcguire, & P. P. Barros (Eds.), *Handbook of health economics* (Vol. 2, p. 397 - 432). Elsevier.

Choné, P., & Ma, A. (2011). Optimal health care contract under physician agency. *Annals of Economics and Statistics*(101/102), 229–256.

Einav, L., Finkelstein, A., & Mahoney, N. (2018). Provider incentives and healthcare costs: Evidence from long-term care hospitals. *Econometrica*, *86*(6), 2161-2219.

Einav, L., Finkelstein, A., Ryan, S. P., Schrimpf, P., & Cullen, M. R. (2013). Selection on moral hazard in health insurance. *American Economic Review*, *103*(1), 178–219.

Eliason, P. J., Grieco, P. L. E., McDevitt, R. C., & Roberts, J. W. (2018, November). Strategic patient discharge: The case of long-term care hospitals. *American Economic Review*, *108*(11), 3232-65.

Encinosa, W. E., Gaynor, M., & Rebitzer, J. B. (2007). The sociology of groups and the economics of incentives: Theory and evidence on compensation systems. *Journal of Economic Behavior & Organization*, *62*(2), 187 - 214.

Frandsen, B., Powell, M., & Rebitzer, J. B. (2019). Sticking points: common-agency problems and contracting in the us healthcare system. *The RAND Journal of Economics*, *50*(2), 251-285.

Frandsen, B., & Rebitzer, J. B. (2015). Structuring incentives within accountable care organizations. *Journal of Law, Economics, and Organization*, *31*(suppl. 1), 77-103.

Frech, H. E., Whaley, C., Handel, B. R., Bowers, L., Simon, C. J., & Scheffler, R. M. (2015, Sep 01). Market power, transactions costs, and the entry of accountable care organizations in health care. *Review of Industrial Organization*, *47*(2), 167–193.

Gaynor, M. (2006, June). *What do we know about competition and quality in health care markets?* (Working Paper No. 12301). National Bureau of Economic Research.

Gaynor, M., Ho, K., & Town, R. J. (2015, June). The Industrial Organization of Health-Care Markets. *Journal of Economic Literature*, *53*(2), 235-284.

Gaynor, M., Mehta, N., & Richards-Shubik, S. (2023). Optimal contracting with altruistic agents: Medicare payments for dialysis drugs. *American Economic Review*, *113*(6), 1530–1571.

Gaynor, M., Rebitzer, J., & Taylor, L. (2004). Physician incentives in health maintenance organizations. *Journal of Political Economy*, *112*(4), 915–931.

Glied, S., & Hong, K. (2018). Health care in a multi-payer system: Spillovers of health care service demand among adults under 65 on utilization and outcomes in medicare. *Journal of Health Economics*, *60*, 165 - 176.

Gowrisankaran, G., Joiner, K. A., & Léger, P.-T. (2017, December). *Physician practice style and healthcare costs: Evidence from emergency departments* (Working Paper No. 24155). National Bureau of Economic Research.

Grassi, S., & Ma, A. (2016). Information acquisition, referral, and organization. *The RAND Journal of Economics*, *47*(4), 935-960.

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

Hackmann, M. B. (2019, May). Incentivizing better quality of care: The role of medicaid and competition in the nursing home industry. *American Economic Review*, *109*(5), 1684-1716.

Ho, K., & Lee, R. S. (2023). Health insurance menu design for large employers. *The RAND Journal of Economics*, *54*(4), 598–637.

Ho, K., & Pakes, A. (2014). Hospital choices, hospital prices, and financial incentives to physicians. *The American Economic Review*, *104*(12), 3841–3884.

Lee, J. T., Polsky, D., Fitzsimmons, R., & Werner, R. M. (2020). Proportion of racial minority patients and patients with low socioeconomic status cared for by physician groups after joining accountable care organizations. *JAMA Network Open*, *3*(5), e204439–e204439.

McWilliams, J. M., Hatfield, L. A., Landon, B. E., Hamed, P., & Chernew, M. E. (2018). Medicare spending after 3 years of the medicare shared savings program. *New England Journal of Medicine*, *379*(12), 1139-1149.

Rebitzer, J. B., & Votruba, M. E. (2011, October). *Organizational economics and physician practices* (Working Paper No. 17535). National Bureau of Economic Research.

Rosen, J. B. (1965). Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, *33*(3), 520–534.

Stock, J. H., Wright, J. H., & Yogo, M. (2002). A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business & Economic Statistics*, *20*(4), 518–529.

Wilson, M., Guta, A., Waddell, K., Lavis, J., Reid, R., & Evans, C. (2020). The impacts of accountable care organizations on patient experience, health outcomes and costs: a rapid review. *Journal of Health Services Research & Policy*, *25*(2), 130-138.