DOES SIZE REALLY MATTER?
AN ANALYSIS OF THE RELATIONSHIP BETWEEN ACO SIZE AND FINANCIAL SUCCESS

By:

Joseph Aoun

Economic Statistics 305

December 5, 2018

Does Size Really Matter?
An Analysis of the Relationship between ACO Size and Financial Success

ABSTRACT

The 2017 financial performance of Medicare Accountable Care Organizations ("ACOs") varied considerably, with some ACOs generating multi-million dollar savings and others producing losses. The conventional view is that ACOs that have a larger number of beneficiaries are more likely to be successful. Drawing from a sample of the 2017 actual ACO financial results, this paper examines the relationship between ACO size and financial success. Following an analysis of correlation, this paper finds that, contrary to the conventional view, there is not a strong positive correlation between these two variables. The advantages associated with economies of scale (in terms of spreading administrative costs across a larger population or gaining leverage in negotiating discounts with providers) do not measurably impact success in the program as currently structured. Although large ACOs may not be more likely to achieve financial success than non-large ACOs, they do enjoy greater predictability in health care costs.

INTRODUCTION

The need to control health care costs is compelling: the most recent Medicare Annual Report projects that fund assets will be depleted by 2026.[1] The Patient Protection and Affordable Care Act, commonly referred to as the ACA, contains various provisions to reign in health care costs. These provisions run the gamut of reduced Medicare payment rates to providers (or more accurately, reduced annual increases in payment rates) to new programs designed to promote

[1] 2018 Annual Report of the Boards of Trustees of the Federal Hospital Insurance and Federal Supplementary Medical Insurance Trust Funds. The Medicare program provides health coverage to individuals who are over the age of 65 or otherwise have qualifying disabilities, such as end stage renal disease. In 2017, approximately 58 million individuals were covered by Medicare.

accountability for the overall cost of care. One such program, the Medicare Shared Savings Program, is a voluntary program where groups of providers - referred to as accountable care organizations or ACOs - try to manage costs below a budget. If the ACO is successful, it can receive a share of the difference between budgeted and actual expenditures.

To participate in the Medicare Shared Savings Program, the ACO must have participating physicians who provide services to at least 5,000 Medicare beneficiaries. In 2017, more than 470 ACOs participated in the Medicare Shared Savings Program. Data from the federal Medicare agency that administers the program (the Centers for Medicare and Medicaid Services or CMS) indicates that ACOs ranged considerably in size, with some ACOs having responsibility for the minimum number of beneficiaries (5,000) and several ACOs having responsibility for more than 100,000 beneficiaries.

The 2017 ACO financial results show that ACOs saved the Medicare program more than $1.0 billion, and other researchers have postulated that the savings from ACOs are larger.[2] Thus, the ACO model may prove to be an effective cost control tool, and the model is capable of being replicated for use in other health insurance programs, such as Medicaid and employer group insurance. Although the ACO model has potential for significant savings, the performance among the ACOs in 2017 was mixed: only 60% generated savings in relation to their benchmarks. Further, to receive a share of the savings from CMS, the ACO's savings rate must surpass a minimum

[2] The savings of $1.0 billion is measured as the difference between budgeted and actual expenditures. Recently, some researchers have proposed that a better measure of success would be to look at what Medicare expenditures would have been in the absence of ACOs (counterfactual). The Medicare Payment Assessment Commission (established by statute to advise Congress on Medicare payment policy) has reviewed the counterfactual research literature and concluded that overall Medicare savings arising from ACOs is in the range of 1-2%. *See* Report to the Congress: Medicare and the Health Care Delivery System, June 2018, Ch. 8. Although a relatively small percentage, it is nonetheless significant given the size of Medicare expenditures (2017: $710 billion).

savings rate. Only 34% of the participating ACOs in 2017 cleared this "hurdle." (Bleser et al. 2018)

Both CMS that administers the program, and the ACOs that participate, desire successful performance. An important question is whether the number of beneficiaries assigned to an ACO is related to success. Stated differently, are larger ACOs more likely to achieve savings than smaller ACOs? The answer to this question has important policy implications. If larger ACOs are more likely to be successful, this could stimulate consolidation activity in order to achieve scale, e.g., hospitals acquiring physician practices (vertical) or physician groups merging with other group practices (horizontal). These larger organizations may also wield considerable clout in health insurance markets outside of Medicare, e.g., when negotiating with commercial insurers such as Blue Cross or United Healthcare. There is ample research concluding that provider consolidation can result in higher health care costs because providers with greater market power tend to command higher payment rates. (Neprash et al. 2017).

Conversely, if the number of beneficiaries is not a determinant of success, then there may be policy initiatives that could increase the attractiveness of the program to encourage smaller ACOs to participate. For example, CMS could consider lowering the minimum savings rate for smaller ACOs. Under current regulations, an ACO with 5,000 - 6,000 beneficiaries has a minimum savings rate in the range of 3.6 - 3.9%, whereas an ACO with more than 60,000 beneficiaries has a minimum savings rate of 2.0%. This differential may disincentivize participation by smaller ACOs. Likewise, under current regulations, ACOs are required to assume risk for losses after so many years in the program, and CMS has recently proposed regulations that would accelerate this progression.[3] If size is not a determinant of success, and CMS desired to improve the attractiveness

[3] *See* 83 Fed. Reg. 41786 (August 17, 2018).

of the program to smaller ACOs, it could provide more time for smaller ACOs to participate in the program without having to assume downside risk.

THEORY

As organizations increase in size, they are able to achieve economies of scale. Economies of scale refers to the dynamic where average total costs decline as quantity increases. Organizations have fixed costs that, by definition, do not vary with production, i.e., whether 500 widgets are produced or 10,000, there are some costs that will not vary, such as the cost of acquiring the production facility. Organizations also incur costs that vary with production, such as most labor and materials. As more widgets are produced, more raw materials are acquired. Total costs refers to the sum of fixed and variable costs, and average total costs is total costs divided by the quantities of production.

Although often discussed in the context of production, the principles relating to fixed and variable costs apply equally in the services context. For example, with hospital services, a hospital has fixed costs relating to facility maintenance. Regardless of the number of patients admitted (occupancy rate), there will be some costs incurred to heat the building, maintain the exterior and provide basic management. Increases in the occupancy rate will result in increased costs (e.g., more nurses to care for more patients). Still, there will be productivity gains as there are incremental changes in occupancy: the number of nurses on a floor may not increase even if the number of patients increases from 12 to 16. Thus, firms in the service sector can also achieve economies of scale as average total costs decline with increasing service quantities. Graphically, this can be presented as follows:

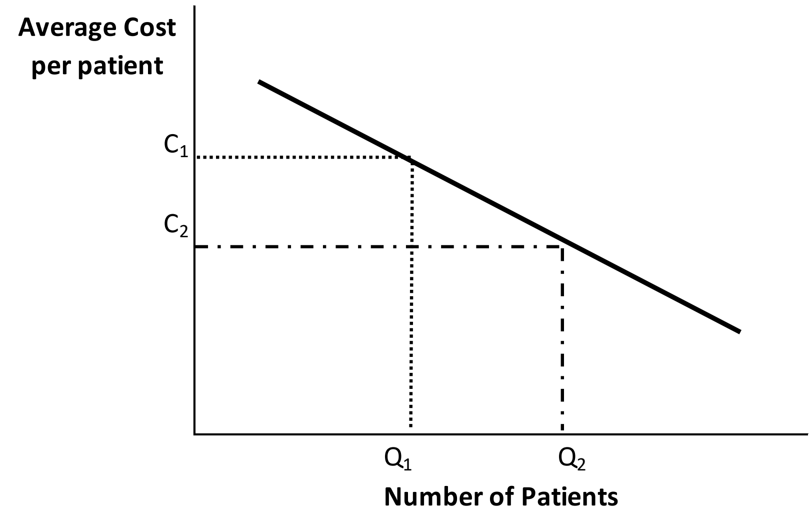

As a provider increases the number of patients from Q_1 to Q_2, the average cost per patient (both fixed and variable) will fall from C_1 to C_2.

ACOs are similar to health plans in the sense that they have financial responsibility for the health care costs of a defined group of individuals (CMS attributes Medicare beneficiaries to an ACO based on the ACO's participating physicians). In the context of health plans, such as health maintenance organizations and insurance companies, economies of scale offer an advantage in terms of administrative costs. As the number of enrollees increase, the administrative cost per enrollee decreases. As one researcher has commented, "administrative costs exhibit economies of scale that create a cost advantage for large HMOs."[4] Since premiums are set based on administrative costs, among other items, the savings from scale could result in the plans gaining market share by establishing lower premiums. Even if premiums are not adjusted, plans with economies of scale may enjoy greater profitability.

Larger enrollment also positions health plans to negotiate better contracts with health care providers. While administrative costs are a component in premium setting, they typically comprise

[4] Wholey et al. 1992 at p. 831.

less than 15% for health plans, and the bulk of the premium is to cover medical expenses (typically around 80%). Health care expenses are a function of price and quantity - the number of services needed (e.g., inpatient hospital admissions) and the amount paid per unit of service (e.g., $5,000 per admission). Health plans with larger enrollment represent more volume to providers and thus enjoy greater leverage in rate negotiations than plans with small enrollment. (Wholey et al. 1992). Other things being equal, a health plan with larger enrollment will achieve better discounts from providers. Consequently, a health plan with larger membership can be expected to achieve lower medical costs, which positions it more favorably in terms of market share gains (lower costs should result in lower premiums) and profitability.

Health plans with larger membership can also be expected to enjoy more predictable health care expenses. There should be less variation in costs when a plan has a large number of enrollees, and the plan should be better positioned to absorb outliers, such as the costs associated with neonate triplets. From an actuarial perspective, this is often times referred to as the law of large numbers. As one leading actuarial firm has noted, "from the standpoint of projecting claims experience, the credibility of historical or projected experience increases as the size of the underlying population increases."[5]

Generally, ACOs with a larger number of primary care physicians (such as internists and family practice physicians) will have more members attributed to them. Based on the attributed members, CMS establishes a cost benchmark, e.g., $12,000 per member per year. If actual expenses are below the benchmark, the ACO may be entitled to share a portion of the savings with CMS.

[5] Milliman Healthcare Reform Briefing Report, *Payer and Provider "Checklist" for Alternative Payment Arrangements* (February 2015) at p. 5

In order to receive savings, an ACO's quality performance must meet certain targets and its savings rate must be greater than the specified minimum savings rate (MSR). CMS establishes the MSR based on the number of beneficiaries that are attributed to the ACO:

Number of Beneficiaries	MSR - low end	MSR - high end
5,000 - 5,999	3.90%	3.60%
6,000 - 6,999	3.60%	3.40%
7,000 - 7,999	3.40%	3.20%
8,000 - 8,999	3.20%	3.10%
9,000 - 9,999	3.10%	3.00%
10,000 - 14,999	3.00%	2.70%
15,000 -19,999	2.70%	2.50%
20,000 - 49,999	2.50%	2.20%
50,000 - 59,999	2.20%	2.00%
60,000+	2.00%	

Source: 42 CFR 425.604(b)

CMS claims that the MSR is set to reduce the likelihood of paying savings due to random variation as opposed to ACO performance:

> "Establishing an MSR on the basis of standard inferential statistics that take into account the size of an ACO's beneficiary population provides confidence that, once the savings achieved by the ACO exceed the MSR, the change in expenditures represents actual performance improvements by the ACO as opposed to normal variations….
>
> The MSR would be established for each ACO based on increasing nominal confidence intervals for larger ACOs so that an ACO with the minimum 5,000 assigned beneficiaries would have an MSR based on a 90 percent CI; an ACO with 20,000 assigned beneficiaries would have a MSR based on a 95 percent CI and an ACO with 50,000 assigned beneficiaries would have an MSR based on a 99 percent CI."
>
> 67 Fed. Reg. at 19611 (April, 7, 2011).

This paper examines the relationship between ACO size and the likelihood of earning savings. In particular, this paper explores whether the benefits of scale that are well-recognized in the health plan context apply to ACOs. Other things being equal, this paper theorizes that larger ACOs are more likely to generate savings than smaller ACOs due to economies of scale. It is anticipated that there will be a significant positive correlation between the variables of ACO size and savings rates. This paper also theorizes that larger ACOs will have more predictable health care expenditures than smaller ACOs given the so-called law of large numbers. It is anticipated that the variation in savings rates among larger ACOs will be smaller when compared to the variation of other ACOs.

METHODOLOGY

CMS posts ACO financial results on the https://data.cms.gov website annually. Among other things, the 2017 performance data identifies each participating ACO, the number of attributed beneficiaries, the benchmark and actual expenditures, the ACO's savings rate, and MSR. The data also identifies by ACO whether it "generated savings" (i.e., ACO expenditures less than benchmark) and whether it "earned savings" (i.e., ACO savings rate exceeds MSR). The ACO data is cross-sectional, quantitative data.

For this paper, the focus is on ACO size and savings rates. The variable size is reported by CMS as the number of beneficiaries attributed to the ACO. The attribution methodology is set forth in the ACO regulations (42 CFR Part 425). The Medicare Shared Savings Program operates in the traditional Medicare fee for service program. Unlike the Medicare managed care program (known as Medicare Advantage), beneficiaries do not affirmatively enroll in an ACO. They remain in Medicare fee for service and may receive services from any Medicare participating provider. Because there is no formal enrollment in an ACO, CMS attributes a beneficiary to an ACO if the beneficiary receives a plurality of primary care services from a physician participating in the ACO.

(Each ACO submits a physician participation list, and CMS reviews beneficiary claims information by physician as part of its attribution methodology.)

The second variable of focus in this paper is savings rate. This rate is also reported by CMS and is computed as the difference between actual and budgeted expenditures divided by the budgeted expenditures. If an ACO's actual expenditures exceed budget, CMS will report a negative savings rate for that ACO. The calculation of savings rates only focuses on health care expenditures. Administrative costs are not considered in the calculation.

It is also important to note that the benchmark and actual expenditure data reflect the volume of covered services paid at prevailing Medicare payment rates. Unlike health plans, ACOs do not negotiate rates of payment with participating providers. Instead, providers are paid at traditional Medicare fee for service rates, and these rates (with some adjustments) are used by CMS when determining the benchmark and actual expenditures. Given the dynamic of administrative prices, the ACO's ability to generate savings is primarily through reductions in utilization (e.g., lowering the number of inpatient admissions) or redirection of services to providers that have lower Medicare rates (e.g., moving outpatient surgery from hospitals that have higher Medicare prices to free-standing surgery centers that have lower Medicare prices).

The 2017 ACO financial results were obtained from the above-referenced website. The population of participating ACOs in 2017 was 472. Using the random number generator function in Excel, a random sample of 100 ACOs was selected. This sample was used for all of the descriptive and inferential statistical analysis discussed in this paper, except as noted below.

This paper also explores the differences in the variance in savings rates between "large ACOs" and non-large ACOs. "Large ACOs" are defined as those having more than 40,000 beneficiaries which is approximately equal to the population mean number of beneficiaries

(19,053) plus one standard deviation (19,122). The data was separated between large and non-large ACOs, and then a random sample of 40 large ACOs and 40 non-large ACOs was obtained. These random samples were used solely for purposes of the variance analysis.[6]

Descriptive Analysis

Tables 1 and 2 summarize the descriptive statistics for size and savings rates:

Table 1

Sample Number of Beneficiaries	
Mean	19,281
Standard Error	1,815
Median	12,091
Mode	#N/A
Standard Deviation	18,154
Sample Variance	329,555,314
Skewness	2.431282274
Range	112,098
Minimum	4,509
Maximum	116,607
Count	100
Confidence Level(95.0%)	3602

Table 2

Sample Savings Rate	
Mean	1.68%
Standard Error	0.41%
Median	1.84%
Mode	#N/A
Standard Deviation	4.09%
Sample Variance	0.17%
Skewness	-0.09071
Range	19.36%
Minimum	-8.22%
Maximum	11.14%
Count	100
Confidence Level(95.0%)	0.81%

The mean number of beneficiaries attributed to ACOs in the sample was 19,281. Among the 100 selected ACOs, the mean represents the average size, and it will serve as the point estimate for later inferential analysis. Interestingly, the median size (12,091) is smaller than the average size (19,281). The median represents the point where one-half of the ACOs in the sample have a larger number of beneficiaries and one-half have a smaller number. When the median is smaller than the mean, it suggests that the mean is being pushed up due to the impact of some very large ACOs. A few very large ACOs could be accounting for this result, and the range indicates a

[6] The random sample of 100 ACOs which was used for the descriptive and inferential analysis (other than the variance analysis) only had 15 "large ACOs." To improve the integrity of the variance analysis, a second random sample of 40 large ACOs and 40 non-large ACOs was utilized.

sizeable difference between the smallest size ACO (4,509)[7] and the largest (112,098). The measure of skewness (+2.43) also shows that the distribution is skewed to the right, consistent with the observation that the median is smaller than the mean. Both the standard deviation and variance illustrate that there is a fair amount of dispersion in ACO size around the mean.

In the case of the savings rate, the mean (1.68%) is actually smaller than the median (1.84%), illustrating that the distribution is skewed slightly to the left (skewness: -0.09). Opposite of what is happening with the size variable, in the case of the savings rate, it appears that poor performing ACOs are bringing down the average to a level below the median. There is also a very large range of savings rates, with one ACO in the sample achieving a rate greater than 11% and another having a negative savings rate (loss rate) greater than 8%. The standard deviation of 4.09% confirms a fairly wide dispersion in ACO savings rates around the mean.

Both Table 1 and Table 2 present the margin of error for the 95% confidence level (last line in each table). In the case of size, an interval estimate of 3,602 means that there is a 95% probability that the population mean will fall within the range of 15,679 (mean less 3,602) and 22,883 (mean plus 3,602). Likewise, in the case of savings rates, an interval estimate of .81% means that there is a 95% probability that the population mean savings rate will fall within the range of .87% to 2.49%. This matter is reviewed in more detail in the discussion concerning the inferential analysis.

Based on the sample data, I examined whether there is a linear relationship between size and savings rates through a covariance analysis. As illustrated below, the covariance is negative, indicating an inverse linear relationship; that is, as the variable size increases, there is a decrease

[7] Although an ACO must have at least 5,000 attributed beneficiaries to join the program, changes during the year in physician composition and beneficiary utilization patterns can result in an ACO having responsibility for a group of beneficiaries that is smaller than 5,000.

in the savings rate variable. This contradicts the theory that size and savings rates have a positive relationship.

Covariance

	Size	*Savings Rate*
Size		
Savings Rate	-52.7541	

To measure the strength of this linear relationship, a correlation analysis was performed on the sample. The correlation analysis shows a weak negative relationship (-7.18%). This is close to 0 which represents no correlation. Consequently, the bivariate analysis does not support the theory that larger ACOs are more likely to achieve savings.

Correlation

	Size	*Savings Rate*
Size		
Savings Rate	-0.0718	

I also conducted a probability analysis on the sample. Among the 100 ACOs, there were 15 that had more than 40,000 beneficiaries (large ACOs). The likelihood that the large ACOs in the sample would have a savings rate greater than their MSR was around 33%, whereas the probability that an ACO in the sample would beat the MSR (whether large or not) was 42%. Consequently, based on the sample data, the large ACOs were less likely to achieve savings than other ACOs:

	Beats MSR	**Does Not Beat MSR**	*Total*
Large ACO	5.00	10.00	15.00
Other ACO	37.00	48.00	85.00
Total	42.00	58.00	100.00

Probability of beating MSR: 42/100 = 42.00%
Probability of beating MSR|Large: 5/15 = 33.33%

INFERENTIAL ANALYSIS

For convenience, the inferential analysis is arranged by first examining the means between the dependent and independent variables, followed by an analysis of their correlation. The inferential analysis concludes with a review of the variances of the two variables.

A. MEANS

In this analysis, the dependent variable - savings rate - has a point estimate or mean of 1.68% (Table 2). The point estimate (x-bar) is a consistent and unbiased estimator of the population mean. The central limit theorem provides that the distribution of x-bar will approach the normal distribution as the size of the sample increases, regardless of the actual population distribution. A general requirement is that the sample size have at least 30 observations; in the sample selected for this paper, there are 100 observations.

To estimate the confidence interval for the population mean, the first calculation is the margin of error ("MOE") at the given confidence level. The higher the confidence level, the greater the margin of error. If the population standard deviation is not known, the MOE is determined by multiplying the t-critical value by the standard error (sample standard deviation/square root of the sample size). For the 95% confidence level, where there are 99 degrees of freedom (sample size minus 1), the t-critical value is 1.982. Per Table 2, the standard error is 0.41% (4.09%/√100). Consequently, the MOE is equal to 1.9842 * 0.41% or 0.81%. This amount is also reflected in the last line of Table 2.

The formulas described above can be expressed as follows:

$\bar{x} \pm t_{n-1, \alpha/2} * (s/\sqrt{n})$

$t_{n-1, \alpha/2} = 1.9842$ (confidence level = 95%; α = 5%)

1.68% ± 1.9842 * (4.09%/√100)

1.68% ± 0.81%

Upper Confidence Limit: 1.68% + 0.81% = 2.49%

Lower Confidence Limit: 1.68% - 0.81% = 0.87%

Given the above, there is a 95% probability that the population mean savings rate will fall within the range of .87% to 2.49%.

I also completed a hypothesis test of the population mean. The null hypothesis is that the population mean savings rate is equal to or greater than 2%. The alternative hypothesis is that population mean savings rate is less than 2%. Given that the population standard deviation is not known, this test will be done using the t-statistic. Because the alternative hypothesis is that the savings rate is less than 2%, this is a one-tail test where the focus is the left tail. The formula for calculating the t-statistic is expressed as follows:

$$t = \frac{\bar{x} - \mu}{s/\sqrt{n}}$$

$$t = \frac{1.68\% - 2.0\%}{4.09\%/\sqrt{100}} = -0.772$$

Using Excel, it is possible to calculate the critical value and the p-value given the t-statistic of -0.772. Assuming a 1% significance level, the calculations are summarized below:

Hypothesis Test - Savings Rate

H_0 : The population mean savings rate is =>2.0%

H_a : The population mean savings rate is <2.0%

Hypothesized Mean (μ_0)	2.00%
Sample Mean (**Xbar**)	1.68%
Population Standard Deviation (σ)	NA
Sample Standard Deviation (**s**)	4.09%
Standard Error (**SE**)	0.41%
Type of Test	t-test; left tail
Sample Size (**n**)	100
Degrees of Freedom (**df**)	99
Significance Level (α)	0.01
Test Statistic	-0.772
p-value	0.221
critical value	-2.365

The critical value is -2.365, meaning that the null hypothesis should be rejected if the t-statistic is smaller than -2.365. In this case, -0.772 is larger than the critical value. Consequently, there is not enough evidence to reject the null at the 1% level of significance. The same result arises under the p-value approach. If the p-value is less than the significance level (1% in this case), then the null hypothesis should be rejected. Here, the p-value is greater than the significance level (0.221>0.01), so the null hypothesis cannot be rejected. Stated differently, there is not enough evidence at the 1% significance level to accept the alternative hypothesis that the population mean is less than 2%.

With respect to the independent variable - the number of beneficiaries in the ACO - the point estimate is the sample mean (x-bar), and the central limit theorem applies given the sample size is greater than 30. For the 95% confidence interval, the same formula as described above was employed:

$\bar{x} \pm t_{n-1, \alpha/2} * (s / \sqrt{n})$

$t_{n-1, \alpha/2}$ = 1.9842 (confidence level = 95%; α = 5%)

19,281 ± 1.9842 * (18,154/√100)

19,281 ± 3,602

Upper Confidence Limit: 19,281 + 3,602 = 22,883

Lower Confidence Limit: 19,281 - 3,602= 15,679

Given the above, there is a 95% probability that the population mean number of beneficiaries will fall within the range of 15,679 and 22,883.

The hypothesis test for this variable is that the population mean number of beneficiaries is equal to 20,000, and the alternative hypothesis is that the population mean is not equal to 20,000. Following the same framework that is described above with respect to the hypothesis test of the savings rate, I determined a t-statistic of -0.396. This is a two tail test. The right tail critical value at a 1% significance level and 99 degrees of freedom is 2.6264 (per t-table, where $\alpha/2$ = .005). Correspondingly, the left tail critical value is -2.6264. The t-statistic of -0.396 is less than +2.6264 and greater than -2.6264, thus falling in the "do not reject" region. The calculations are summarized below:

Hypothesis Test - Number of Beneficiaries

H_0 : The population mean number of beneficiaries is = 20,000

H_a : The population mean number of beneficiaries is <> (not equal to) 20,000

Hypothesized Mean (μ_0)	20,000
Sample Mean (**Xbar**)	19,281
Population Standard Deviation (σ)	NA
Sample Standard Deviation (**s**)	18,154
Standard Error (**SE**)	1,815
Type of Test	t-test; two tail
Sample Size (**n**)	100
Degrees of Freedom (**df**)	99
Significance Level (α)	0.01
Test Statistic	-0.396
critical value - right tail	2.6264
critical value - left tail	-2.6264
p-value - left tail	0.3465
p-value combined	0.6930

Because the test statistic falls in the range of the two critical values (-2.6264, +2.6264), there is not enough evidence at the 1% significance level to reject the null hypothesis. This conclusion is also confirmed by using the p-value approach. The p-value calculation per Excel for the left tail is 0.3465. Because this is a two-tail test, this amount is doubled to determine the overall p-value (2*0.3465 = 0.6930). Here, the p-value (0.6930) is significantly higher than the significance level (0.01). As a result, there is not enough evidence at the 1% significance level to reject the null hypothesis that the population mean is equal to 20,000.

B. CORRELATION

As noted earlier, the sample correlation between ACO size and savings rates was -7.18%. The sample correlation coefficient (r) is a consistent and unbiased estimator of the population correlation coefficient (ρ). Consequently, the point estimate for inferential analysis of correlation will be r. To determine the interval estimate, the margin of error is added to (or subtracted from) the point estimate, as illustrated by the formula below. Consistent with earlier interval estimates, the confidence interval is set at 95%:

$r \pm t_{n-2,\, \alpha/2} * \sqrt{(1-r^2/ n-2)}$

$t_{n-2,\, \alpha/2}$ = 1.9845 (confidence level = 95%; α = 5%)

$-7.18\% \pm 1.9845 * \sqrt{(1-7.18\%^2/98)}$

-7.18% ± 19.99%

Upper Confidence Limit: -7.18% + 19.99% = 12.81%

Lower Confidence Limit: -7.18% - 19.99% = -27.18%

There is a 95% probability that the population correlation coefficient (ρ) will fall within the range of 12.81% to -27.18%. The confidence interval results support the conclusion that ACO size and savings rates are not strongly correlated (either positively or negatively). Neither value at the upper and lower confidence limits (12.81%, -27.18%, respectively) evidences a strong linear relationship. The results of this interval analysis cast considerable doubt on the proposition advanced in this paper: namely, that ACO size and savings rates have a strong positive correlation.

Similar to the earlier analysis relative to sample means, a hypothesis test of the population correlation was conducted. The null hypothesis is that there is a positive correlation between ACO size and savings rates (ρ = > 0). The alternative hypothesis is that there is a negative correlation between these two variables (ρ < 0). This is a one-tail test, with the focus on the left tail. The t-statistic is utilized for this test and is determined by the following formula:

$$t = \frac{r - \rho}{\sqrt{(1 - r^2)/n-2}}$$

$$t = \frac{-7.18\% - 0}{\sqrt{(1 - -7.18\%^2)/98}}$$

t = -.7128

At the 5% level of significance, the critical value is -1.661. The t-statistic is larger than the critical value, thus falling in the “do not reject” region. Similarly, the p-value (0.239) is substantially higher than the significance level 0.05, also confirming that there is not enough evidence to reject the null. A summary of the Excel calculations of the critical value and p-value is below:

Hypothesis Test - Correlation

H_0 : There is a positive correlation between ACO size and savings rate in the population, ρ = > 0

H_a : There is a negative correlation between ACO and savings rate in the population, ρ < 0

Hypothesized Correlation (ρ_0)	0
Sample Correlation (**r**)	-7.18%
Sample Size (**n**)	100
Degrees of Freedom (**df**)	98
Significance Level (α)	0.05
Test Statistic	-0.7128
p-value	0.239
critical value - left tail	-1.661

The hypothesis test was repeated where the null hypothesis was that the population correlation was equal to or greater than 0.5. This level was selected on the basis that, if the population correlation coefficient was greater than 0.5, one would conclude that there is a strong positive correlation between ACO size and savings rates. The second hypothesis test produced a t-statistic of -5.6763 which would fall in the left tail rejection region. Consequently, the null hypothesis can be rejected, and there is strong evidence supporting the alternative hypothesis ($\rho<0.5$) at the 5% significance level. The results of this test are contrary to this paper's initial theory that ACO size and savings rates have a strong positive correlation.

C. VARIANCE

The inferential analysis concludes with an examination of the variances of the two variables. Even if ACO size and savings rates are not positively correlated, there may still be some advantages to size in the sense that larger ACOs may have more predictable health care expenditures. As noted earlier, for this analysis a different random sample was used than the sample in the preceding analyses. For this analysis, a random sample of 40 large ACOs (ACOs with attributed beneficiaries greater than 40,000) and a random sample of 40 non-large ACOs was taken. The variance in savings rates among the large ACOs was 0.08%, and the variance in savings rates among the non-large ACOs was 0.22%. The testable hypothesis was whether large ACOs in the population have a smaller variance than non-large ACOs in the population at the 5% significance level.

This is one-tail test, with a focus on the left tail. The f-test was utilized to compare the two variances. Given the testable hypothesis, the f-statistic was determined by dividing the large ACO variance (0.08%) by the non-large ACO variance (0.22%), yielding 0.35126. Using Excel, the critical value and p-value were determined for this statistic at a 5% significance level:

Hypothesis Test - Two Sample Variances

H_0: The variance of Large ACOs is equal to Non-Large ACOs - $\sigma_1^2 = \sigma_2^2$

H_a: The variance of Large ACOs is less than Non-Large ACOs - $\sigma_1^2 < \sigma_2^2$

Significance Level:		**5%**
	Large ACOs	Non-Large ACOs
Mean	0.09%	0.06%
Variance	0.08%	0.22%
Observations	40	40
df	39	39
F	0.35126	
P(F<=f) one-tail	0.00075	
F Critical one-tail 5%	0.58669	

As noted above, the f-statistic (0.35126) is smaller than the critical value (0.58669). Consequently, there is enough evidence at the 5% level of significance to reject the null hypothesis and accept the alternative hypothesis. This also confirmed because the p-value (.00075) is substantially less than the significance level (0.05). Consequently, there is enough evidence at the 5% level of significance that the variance in savings rates among large ACOs in the population is smaller than the variance in savings rates among non-large ACOs in the population.

Summary and Conclusions

The analysis of 2017 ACO financial results does not provide support for the proposition that ACO size and savings rates have a strong positive correlation. The 95% confidence level estimate of correlation concludes that the upper confidence limit is only 12.81%, and the lower confidence limit is -27.18%. Neither of these values imply a strong linear relationship between the two variables. Given that the lower confidence limit (27.18%) is greater than the upper confidence limit (12.81%) in absolute terms, there is some evidence that there is an inverse linear relationship between these two variables in the population. This is not surprising given that the covariance and correlation analyses of the sample also revealed an inverse linear relationship.

The results of the hypothesis testing on correlation also demonstrate that there is not a strong positive correlation between ACO size and savings rates, contrary to this paper's initial theory. Under the first hypothesis test, there was not enough evidence at the 5% significance level to reject the hypothesis that the population correlation coefficient was equal to or greater than 0. The results of this test do not prove that there is a positive correlation; only that there is not enough evidence to conclude that the correlation is negative (alternative hypothesis). The test was repeated where the hypothesized population correlation coefficient was equal to or greater than 0.50. The second test results provide strong evidence to reject this hypothesis at the 5% significance level and accept the alternative (population correlation coefficient is below 0.5). Of course, accepting the alternative hypothesis is, in essence, concluding that this paper's initial theory about the relationship between ACO size and savings rates was incorrect. Stated differently, one can conclude from the inferential analysis that there is not a strong positive correlation between these two variables in the population.

At one level, this seems counter-intuitive: one would expect that, due to economies of scale, larger ACOs would be more successful. They have the ability to implement medical management programs and other cost-cutting initiatives at a lower average cost than smaller ACOs. This advantage, however, is not measured in the ACO financial results. CMS focuses exclusively on health care expenditures and does not factor in administrative costs. Likewise, while health plans having larger enrollment have the ability to obtain better discounts from providers, this is not the case in the ACO program: ACOs do not negotiate payment rates - these are set by Medicare. Consequently, a large ACO does not enjoy any more negotiating leverage than a small ACO.

Smaller ACOs may also have more ability to implement care management programs effectively over a smaller group of participating physicians. There is some research supporting the proposition that, with respect to physician groups, substantial scale is not needed in order to achieve efficiency gains in care management. (Neprash et al. 2017). With larger organizations, the ability to "police" adherence to medical management programs may be administratively infeasible given the number of participating physicians. Moreover, any one physician's share of any CMS savings will be diluted in larger organizations, and this may influence the level of effort exerted to achieve cost savings. In smaller organizations, any payments from CMS will be shared with fewer participants, and this may motivate high performance.

Although larger ACOs may not necessarily have a higher likelihood of success in the ACO program, it does appear that they benefit from the law of large numbers, and there is less variation in their savings rates as compared to non-large ACOs. Large ACOs, therefore, appear to benefit from greater predictability in health care costs.

BIBLIOGRAPHY

1. Medicare Payment Assessment Commission, Report to the Congress: Medicare and the Health Care Delivery System (June 2018), Ch. 8 available at http://www.medpac.gov/docs/default-source/reports/jun18_ch8_medpacreport_sec.pdf?sfvrsn=0.

2. 2018 Annual Report of the Boards of Trustees of the Federal Hospital Insurance and Federal Supplementary Medical Insurance Trust Funds available at https://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/ReportsTrustFunds/Downloads/TR2018.pdf.

3. Bleser, William et al., *Half a Decade In, Medicare Accountable Care Organizations Are Generating Net Savings*, Health Affairs Blog, September 20-21, 2018 available at https://www.healthaffairs.org/do/10.1377/hblog20180918.957502/full/ (Part I) https://www.healthaffairs.org/do/10.1377/hblog20180920.604462/full/ (Part II).

4. Neprash, Hannah et al., *Little Evidence Exists to Support the Expectation that Providers Would Consolidate to Enter New Payment Models*, Health Affairs, Vol 36, Iss. 2 (February 2017) available at https://www.healthaffairs.org/doi/10.1377/hlthaff.2016.0840.

5. Milliman Healthcare Reform Briefing Report, *Payer and Provider "Checklist" for Alternative Payment Arrangements* (February 2015) available at http://www.milliman.com/uploadedFiles/insight/2015/payer-provider-payment-checklist.pdf.

6. Wholey, Douglas et al., *Organization Size and Failure Among Health Maintenance Organizations*, American Sociological Review, Vol. 57, No. 6 (December 1992) available at www.jstor.org/stable/2096126.