College of Population Health Faculty Papers   Jefferson College of Population Health

Analyzing the Health Care Cost Curve: A Case Study

Robert Lieberthal
Jefferson School of Population Health

Recommended Citation

Analyzing the Health Care Cost Curve:
A Case Study

Robert D. Lieberthal, PhD

Abstract

This case study uses data from a self-insured employer plan to perform an analysis into the properties of the health care cost curve. The analysis shows that one statistical property of the health care cost curve is that costs rise continuously, not on an annual or monthly basis. Graphical analysis indicates that managed care techniques used to restrain costs can also smooth utilization, producing the continuously growing cost curve observed. The analysis further illustrates that there is no one "cost curve"—analysis must be segmented by population. Finally, the power of predictive models to fit the cost curve varies by population. To the extent that these results generalize to other health plans, this analysis should be used to inform the implementation of strategies to bend the cost curve. Population health management programs and health policy should be based on continuous analysis and adaption rather than implemented as one-off changes. (*Population Health Management* 2013;16:341–348)

Introduction

THE HEALTH CARE COST CURVE has played a major part in the debate about health reform. In order to reform the health care system properly and bend the cost curve, impactful analyses of the cost curve will be required. Such analyses will need to be statistically rigorous and be matched to clinical reality. For example, although a chronic disease may have costs that are incurred over a long period, an acute episode may have costs that are incurred mostly over a short period. Although the traditional units of analysis for health care costs are per member per month (PMPM) and per member per year (PMPY), these metrics are based on placing financial outcomes first rather than merging financial, clinical, and quality outcomes to assess value. However, the history of financial analysis of the health care cost curve suggests that financial outcomes have been analyzed in their own silo, so that the analysis of the cost curve has been separated from analysis of trends in population health. The need to develop and investigate alternative models for fitting the health care cost curve is the main motivation for this study.

Background

Demographic factors play a large role in the level of spending by health insurers. In addition to cross-sectional differences in spending, health care spending tends to rise longitudinally. Spending for a 65-year-old male today tends to be more expensive than spending 10 years ago for an equivalent 65-year-old male. Models that assess how costs rise longitudinally for all, or most, groups are referred to as the "health care cost curve."[1]

Actuarial models typically focus on PMPY or PMPM costs. This is a result of the managed care payer perspective.[2] These frequencies of measurement are designed to conform to the way that actuaries measure, evaluate, and assess the cost of an insured population. For example, one important model of medical spending, forecast of the health care cost curve, comes from Medicare. Medicare collects aggregate figures for medical spending and breaks down the spending by payer and provider. Medicare also projects spending levels for 10 years based both on current law assumptions and reasonable alternative assumptions. The most recent estimate pegs medical spending at 17.9% of GDP in 2010, rising to 19.6% of GDP by 2021.[3] The major sources of uncertainty in the study include the effect of the Patient Protection and Affordable Care Act and the overall path of economic growth.

From the perspective of an insurance plan member, the health care cost curve is an annual phenomenon because plans assess and set the premiums once per year. From the insurer or employer's point of view, it could be a monthly phenomenon if costs steadily rise month by month. What is important to insurers is the effect of general inflation and spending trends on the growth in spending on insured services over time. However, according to a study by Bundorf et al, inflation alone is not sufficient to generate a cost curve because quantity changes are an important part of medical spending growth.[4] In the same study, the authors also note

that studies of spending growth have generally focused on Medicare, and especially physician services within Medicare. Studies such as one by Sisko et al[5] have focused on the rate of growth in spending in the overall economy. These studies may help to determine trend assumptions for health insurance plans but are not population specific enough for many population health management applications.

Motivation

There is a lacuna in the literature that should be addressed by answering the question: With what frequency should the health care cost curve be measured? One way to demonstrate this gap is to look at how large databases of claims data have been analyzed in the past. Looking at publications resulting from the MarketScan data utilized in this study, no prior analysis has looked at the implications of analyzing health care spending at the daily level over the prior 15 years.[6] The reason that this question is important is that the frequency of measurement should be determined based on when interventions can be impactful. If the cost curve changes once per year, then disease management programs, plan renegotiations, and public policy need only be aimed at affecting the once yearly change in the level of spending. If the cost curve is continuously rising, then continuous process improvement may be the best policy. By determining the answer to the question of what the cost curve really looks like, this case study can inform efforts to bend the cost curve, maximizing the efficacy of such efforts.

To take a clinical example, a readmission prevention program may have interventions within 14 days of discharge with outcomes based on readmissions within a 30-day time horizon.[7] In this case, intervention could bend the cost curve up or down. Although the cost of these programs can be normed to a PMPM basis, this in a way obfuscates the true nature of the program because meaningful progress occurs daily. In this case, it may be necessary to measure costs on a daily basis, too. Such an analysis may offer an advantage over analyzing spending on a PMPY or PMPM basis. Although it is not necessary for clinical outcomes to perfectly match the horizon for financial outcomes, the measurement of outcomes should not be constrained by financial outcomes. In other words, the selection of clinical outcomes should not be dependent on the fact that financial data is available only annually or monthly.

From a statistical point of view, the frequency of measurement should maximize the chance of finding a statistically significant result. A pre/post design for studying any intervention may ignore a large degree of useful variation in outcomes of all kinds. It may be the case that real-time analysis is not currently possible for research, in that population health research cannot utilize hourly or real-time data. However, it is the case that practitioners have such a monitoring system as their ideal in many cases, as computing power and other technological improvements make continuous quality improvement a reality. Given that, the research methods utilized to study such programs also should be advanced toward a real-time analysis. There are several analyses that can potentially address these issues, such as an interrupted time series design. In the current study, a move from monthly to daily costs is investigated, which means the application of traditional (non-interrupted) time series designs to population health.

In general, episodes of care certainly do not follow neat monthly or yearly patterns. As a result, if the consumers of information about the health care cost curve are those in the field of population health management, then it is important to ask whether their needs are served by PMPY or PMPM calculations. If statistical analysis indicates that these models are obscuring important variation, then the result would be that population health management strategies with an aim of bending the cost curve would be underpowered, and therefore biased toward the null of no effect. Thus, it is important to find ways to determine what the statistical properties of the cost curve are in order to best design the protocols of population health management programs with a cost component.

Methods

Conceptual model

In reality, health care costs are incurred in real time. An investigation of health care costs that attempted to model costs over time, matching reality as closely as possible, would not divide the time line of care into discrete periods. Rather than chopping a period of time up into years, months, weeks, or days, such an analysis would model spending continuously over an entire time period as it actually happens. In finance, this is referred to as "continuous time," as opposed to "discrete time," modeling.[8] A closely related concept would be to use finer units of analysis: months as opposed to years, or days as opposed to months.

Consider the figures that show spending on an annual, monthly, and daily basis for a hypothetical plan (Fig. 1A and 1B). This plan has average annual costs per member that rise from $1000 at the beginning of 2000 to $2412 at the end of 2012. This represents a 7% compound annual growth rate in costs. Figure 1A shows what costs look like when viewed on a daily, monthly, and annual basis. Assuming that there is no variance around the trend line, the daily and monthly spending time series look very similar. From an analytic view, they are similar as well. Although the daily time series represents 4750 observations, the monthly time series represents 145 observations. Depending on the power of the study required, monthly data may be sufficient.

Annual data may leave out a large degree of statistically important variation. There are a low number of observations—only 14. Thus, annualized data misrepresent reality to a significant extent. Such data understate costs for a large proportion of the year. This is not a statistical concern for some studies, for example, ones with a pre/post design or a long look back period. However, the annual spending model adversely affects a study with a shorter time horizon or the need for higher frequency monitoring. The performance of a treatment, program, or quality improvement initiative meant to decrease costs, is understated by such a model. The reason is that costs are constantly rising in the background because of reasons unrelated to the particular program. Thus, using annual or even monthly data will tend to bias study designs toward the null hypothesis that they have no effect on costs. For some managers, this type of measurement may be sufficient. For studies that require statistical significance in order to determine whether findings are positive, a bias toward the null is a detriment to the design of the best possible study.

Next, consider what costs would look like for a single year if costs grow continuously. This is presented in Figure 1B. In
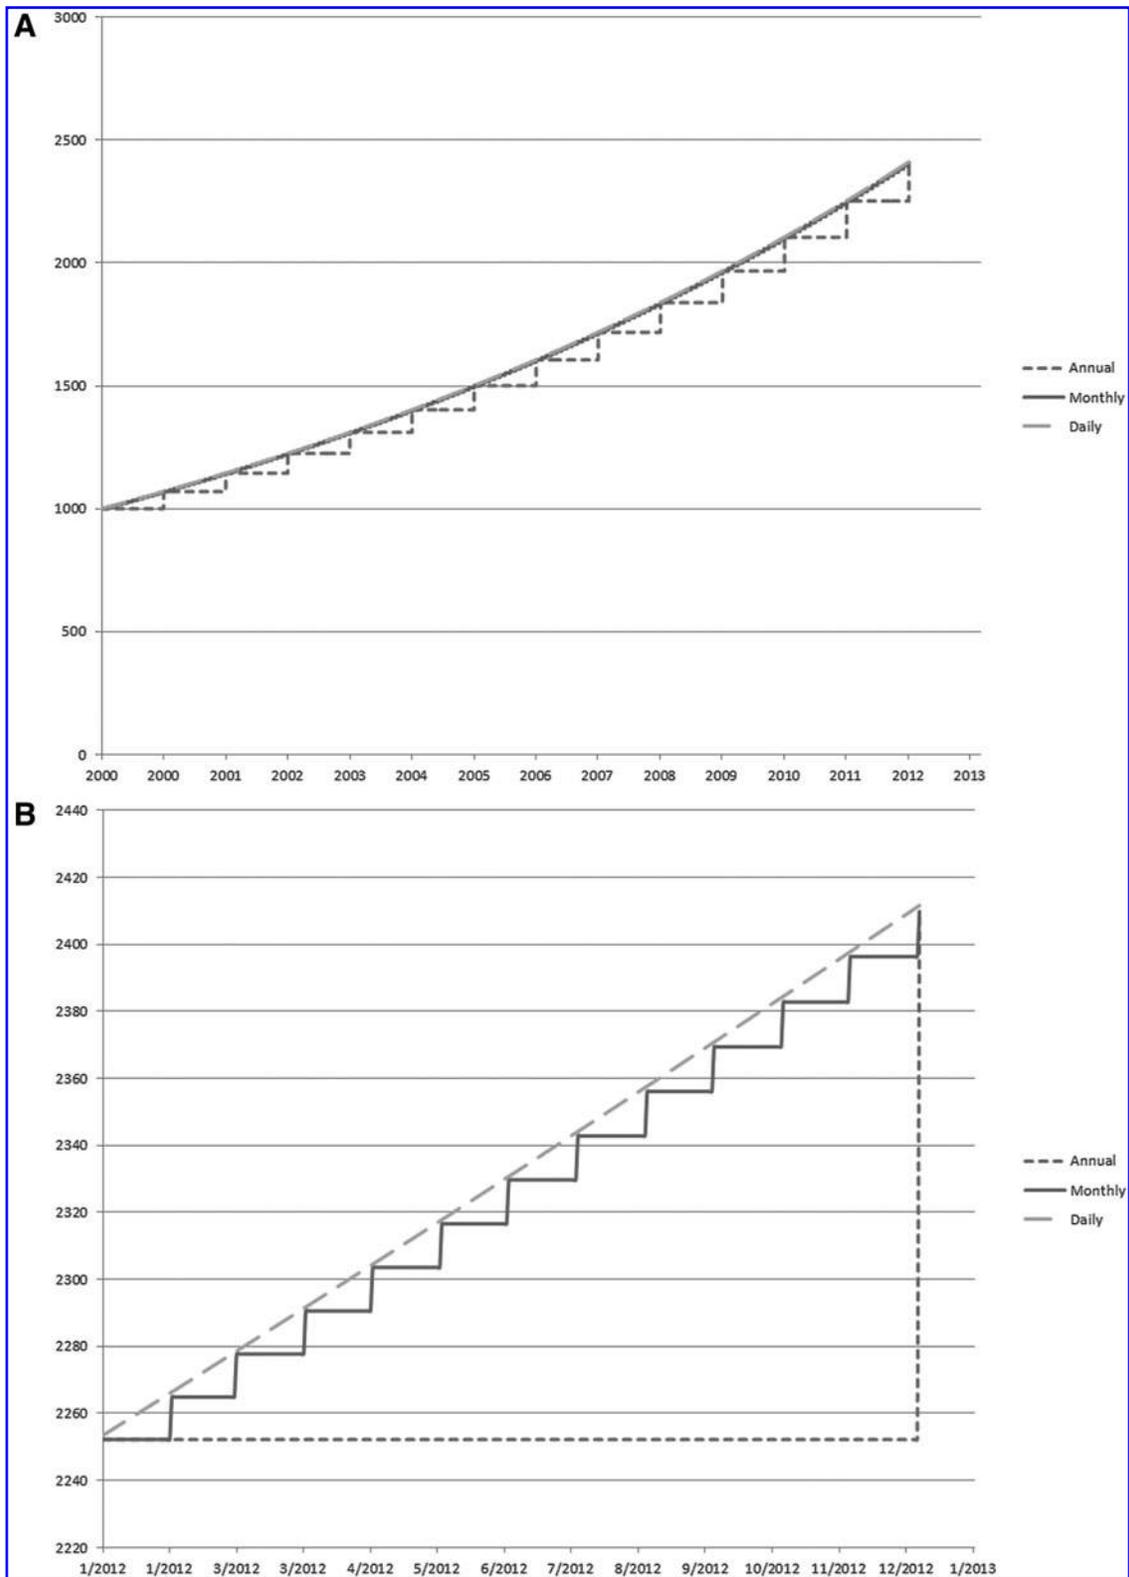
FIG. 1. Conceptual model of annual, monthly, and daily spending.

this case, daily data provide a superior view of actual cost trends than monthly or annual data. Daily data over the full year include 366 observations, while monthly data include 13 observations and annual data include 2 observations. PMPM analyses will tend to understate the effect of the program on bending the cost curve. These data are inadequate to support a study with multiple interventions, assessments of progress, and revised implementation throughout a single year.

The present study fit all 3 models to the data to determine which most accurately reflected reality. Did health care costs rise mainly on an annual basis, or is this a convenient unit of analysis for policy makers? Did health care costs rise on a monthly basis, or is this a convenient unit of analysis for managed care? Did health care costs rise on a daily basis?

Case study design

This study utilized a case study methodology in order to determine the characteristics of the health care cost curve for a population segmented by demographic factors. Taking advantage of a large database of employer-sponsored health insurance plans, the study prioritized several factors in considering which plans to include. The priorities were to find a large employer that was relatively stable, to find an employer that offered a single plan design, and to find an employer that reported both medical and drug costs. Out of 7 possible employers, a single large employer was chosen. Thus, the case study consisted of a population of continuously insured individuals who were salaried, nonunion workers (and their covered spouses) with a single insurance design. The cost curve was analyzed within this homogeneous sample.

Data

The MarketScan data utilized contained patient-level claims data for private insurers surveyed by the health care business of Thomson Reuters.[9] Thompson used the claims from the employers that submit data to create Health Insurance Portability and Accountability Act-compliant, limited-use data sets. The database was provided under a program of the National Bureau of Economic Research. The data files contained the reimbursements for inpatient, outpatient, and pharmaceutical encounters. The data files included population and enrollment files and benefit design files, as well as basic demographic information (age and sex).

Geographic information, such as county and 3-digit zip code also were available, but were restricted to private use. There also was information on the total claim amount for any encounter and the split between the net amount paid by the plan and the amount paid by the employee.

The data selected from the database came from the health insurance plan of an anonymous employer in the Manufacturing, Nondurable Goods industry. Over the 7-year period of 2000–2006, only 1 plan was offered by this employer: a point-of-service (POS) plan with capitation. The employer's choice to offer only 1 type of plan, and not to change the plan offerings, allowed this study to isolate medical spending growth from other changes or plan switching behavior. This study did not control for changes in the benefit design because the composition of the basket of available medical care should change with changes in medical practice behavior. This study also did not control for the degree of capitation utilized in this POS plan or the variation in the type of capitation arrangement utilized over time.

The sample consisted of adults aged 18–64 years, who were covered as salaried, nonunion employees or the spouses of such employees. The sample only included individuals in a given year who were covered by the plan for the entire year. The total number of covered members varied year to year from 18–30 thousand members. Membership in the plan by sex within age group is shown in Table 1. The average population in the plan was 25,491, with no fewer than 18,000 in any given year. This study used the age groups defined in the MarketScan data (18–34, 35–44, 45–54, and 55–64). The population was virtually equally split between males and females, and larger for ages 18–34 and 35–44 than 45–54 and 55–64. Within age groups, only the oldest group, ages 55–64, ever had fewer than 1,000 members, with only 841 members in 2000. The split by sex was generally even, both overall and within age groups.

Expenditure data were available for encounters—drug fills, inpatient episodes, and outpatient visits. All files contained both the total payment made for each encounter, as well as the plan payment made by the insurer. The drug file includes more extensive information, including coinsurance, co-pays, and deductibles. The outpatient file included co-pays and deductibles in all years 2000–2006, but only included coinsurance beginning in 2005. The inpatient file only contained total and plan payments until 2005, when coinsurance, co-pay, and deductible payments also were recorded. All payments

TABLE 1. COVERED LIVES BY SEX WITHIN AGE GROUP

			Age							
			18–34		35–44		45–54		55–64	
	Total	Total								
Year	Male	Female	Male	Female	Male	Female	Male	Female	Male	Female
2000	9643	9667	3400	3596	3868	3873	1905	1827	470	371
2001	11,922	12,140	4242	4596	4639	4688	2454	2363	587	493
2002	14,115	14,674	4934	5533	5427	5611	3041	2927	713	603
2003	14,406	14,997	4694	5338	5585	5739	3317	3231	810	689
2004	14,409	15,038	4502	5169	5461	5675	3550	3400	896	794
2005	13,914	14,931	4182	4960	5094	5514	3671	3581	967	876
2006	8737	9849	2485	3102	3249	3661	2341	2460	662	626
Average	12,449	13,042	4063	4613	4760	4966	2897	2827	729	636

included dates of service and dates of payment, so date of service was used to allocate a claim amount to a claim date. Although the drug and outpatient experiences had a single service date attached to them, an inpatient episode could span multiple days. For that reason, the admission date was used when allocating inpatient expenses for an episode of care to a claim date.

Analysis

Disaggregating the cost curve by demographic group, this study compared and contrasted the spending growth rates of different groups, including which may be larger or more variable, and which were more predictable. A descriptive analysis was the first step to analyze the level of spending over time. This included an analysis of mean spending in total (insurance plan plus member spending) and plan spending alone per member per day. This spending was analyzed by year for plan years 2000–2006, and by month. Spending was compared on weekdays versus weekends, and by demographic groups including sex, age categories, and sex–age category subgroups. Age categories were defined in the MarketScan data (18–34, 35–44, 45–54, and 55–64).

A descriptive analysis of the growth in spending was the next step, in order to calculate and analyze the rate of change in spending over time. Both the rate of change in total spending and plan-only spending were analyzed. The average rate of change in spending PMPY, PMPM, and per member per day were computed. The analysis also was split by sex, age categories, and sex-age category subgroups. The rate of change in spending, or trend, was converted to a logarithmic basis. This also allowed the rate of change in spending to be modeled as a continuous time process. Thus, the rate of change in spending was calculated according to the following equation:

$$\tau_{i,t} = \ln\left(\text{Spending per capita}_{i,t} / \text{Spending per capita}_{i,t-1}\right)$$

For group i for time t $<$ Eq.1 $>$

Equation 1: Log daily trend

Graphical analysis was applied as part of an exploratory modeling process to communicate what the results mean and to explore causes for the findings noted. Spending and encounter counts were explored for all claim types, and separately for drug fills, inpatient episodes, and outpatient visits. The differences in patterns between counts and spending were assessed to determine whether volume or prices could account for any results. Outpatient out-of-pocket payments were analyzed to determine the success of managed care payment techniques in managing and smoothing utilization. Although such analysis for all spending types would be ideal, only outpatient payments are available for the entire 2000–2006 period, which is a limitation of the data. Time series models were applied to remove noise in the data and to test the appropriateness of a continuous time model (additional detail available at http://jdc.jefferson.edu/health-policyfaculty/60/).

It was hypothesized that the health care cost curve would have a daily component as well as annual discontinuities. The daily component of health care spending is the time series modeled in Equation 1. The annual discontinuity would reflect the fact that managed care plans change benefit plans annually and renegotiate annually or every 2 to 3 years. As a result, there would be an annual "reset," meaning a change in patient and provider behavior as they adjusted to the new plan design. This reset should have implications for health care costs, which could rise more or less in a given year depending on the success of managed care in holding down costs. Monthly discontinuities also were tested, with the hypothesis that because managed care finance is commonly managed on a monthly basis, there would be statistically significant differences in spending between months over the course of a year.

Once the final model was identified, it was fit to the data. The final model applied to generate the trend for group i for time t and the prediction error $\hat{e}_{i,t}$ was:

$$\tau_{i,t} = \beta_0 + \beta_1 \cdot \tau_{i,t-1} + \beta_2 \cdot \text{Monday} + \beta_3 \cdot \text{Saturday}$$
$$\hat{e}_{i,t} = \hat{\tau}_{i,t} - \tau_i \qquad\qquad\qquad < \text{Eq.2} >$$

Equation 2: Model for trend and error function

The coefficients of this model and the error terms were assessed to determine model goodness of fit and prediction error.

Results

Level of spending

The per capita counts of encounters showed daily patterns but no clear trend over time. The inpatient counts were discrete, with few claims on any given day leading to "levels" in the graph of counts per capita by claim date and the histogram of counts. The drug and outpatient claims had 2, and possibly 3, different claim count levels corresponding to weekdays and weekends, with more claims on Saturday than Sunday. There was no discernible upward trend in encounters over the years, although there did appear to be a break in the trend in outpatient counts in 2003 (figures available from the author upon request). The overall count levels were mirrored in those of subgroups, such as males versus females. Thus, the count of encounters could not be used to account for the increased spending over time.

The age group categories defined in the data (18–34, 35–44, 45–54, and 55–64) and sex were used to break up the spending by demographic groups. Overall, each age group had significantly different spending even when compared to the closest (ie, the adjacent) age groups, with the older groups more expensive than the younger ones. The same was true within each year for total and plan spending. Separating weekdays from weekends, the difference remained only for weekday spending, suggesting that weekend spending, which is largely inpatient driven, is probably generated by emergencies that have nothing to do with age-related medical care. Spending was higher for females than males, and the difference between weekdays and weekends was more pronounced. These results in the levels confirmed the need to analyze medical trend by demographic group.

Medical trend

Total payments were rising both between and within years. Table 2 shows rising average nominal payments per member per day in each year, both overall and for the plan

TABLE 2. NOMINAL SPENDING PER MEMBER
PER DAY, BY YEAR

Year	Daily	Annualized	Change	SD	N
Total spending					
2000	5.28	1926		2.62	366
2001	5.75	2098	9%	2.87	365
2002	6.77	2470	18%	3.35	365
2003	6.97	2544	3%	3.37	365
2004	8.06	2941	16%	3.91	366
2005	9.08	3314	13%	4.58	365
2006	10.25	3740	13%	5.21	365
Plan spending					
2000	4.06	1483		2.15	366
2001	4.70	1716	16%	2.38	365
2002	5.68	2074	21%	2.91	365
2003	5.50	2006	−3%	2.77	365
2004	6.27	2289	14%	3.23	366
2005	7.07	2580	13%	3.62	365
2006	7.84	2860	11%	4.23	365

FIG. 3. Daily plan spending per capita.

(net of member payments). Mean daily spending was sta-
tistically significantly different in each year. The average
total spending (by plan and individuals) rose from $1926 in
2000 to $3740 in 2006, a 94% increase spread over 6 years.
The corresponding compound annual growth rate was
11.7%, but the annual rates of change in mean spending
ranged from 3% in 2003 to 16% in 2004. Plan spending in-
crease was a nearly identical 93%, but this change masked
even larger variation, including a 21% increase in 2002 and a
3% decrease in 2003. Figures 2 and 3 show the increase in
spending over time for total spending and plan spending.
The summarized results for the daily change in spending are
in Table 3. The absolute value of average daily log change in
spending by group was smaller than 0.001 in all cases. In
some cases, the average was negative, but all figures are
close to zero. However, the medians were all negative,
whereas the skews were all positive, which is indicative of
the fact that most spending growth was generated on a small
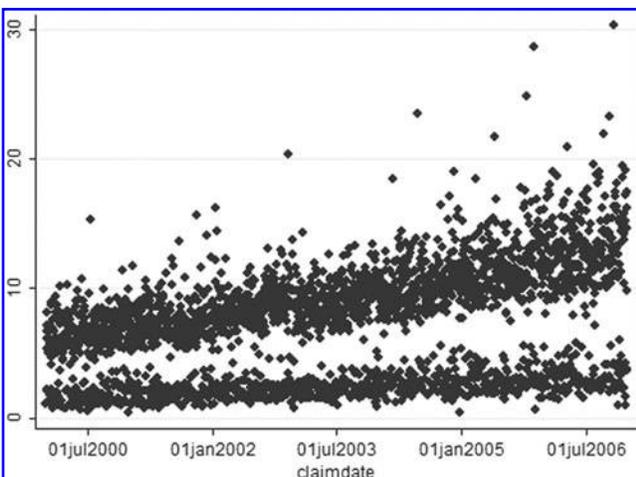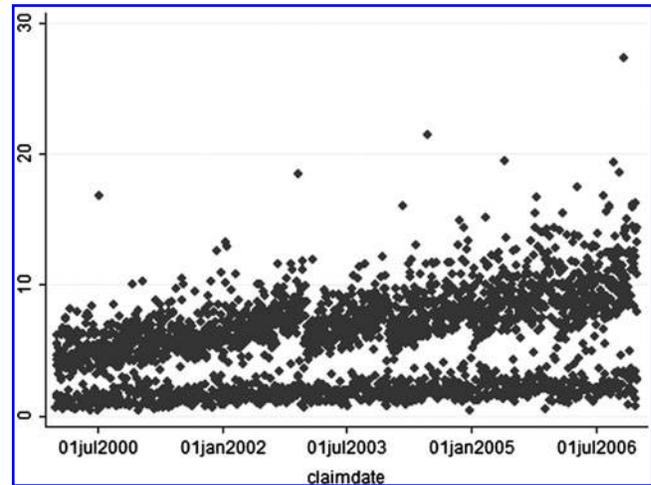number of days. The standard deviations were large enough

that none of the means were indistinguishable from zero—on
average, daily spending growth was zero. In addition, the
standard deviation was lowest for the total population,
smaller for younger than older ages, similar for males and
females overall, and increasing by age groups for females,
but not males. Thus, for purposes of analysis of average
spending by demographic group, the total population had
the most robust sample; younger groups had lower sampling
error than older groups. The error rates may have been re-
lated to the size of the population sampled (Table 1). Addi-
tional results on model development are available at http://
jdc.jefferson.edu/healthpolicyfaculty/60/.

Graphical analysis confirmed the continuous increase in
health care costs. Figure 2 shows this for total costs and
Figure 3 shows a similar pattern for costs borne by the in-
surance plan alone. Both also exhibit a separation in spend-
ing, which was later confirmed as the weekday/weekend
effect in the descriptive analysis. Finally, there was a great

FIG. 2. Daily total spending per capita.

TABLE 3. DAILY LOG CHANGE IN TOTAL NOMINAL SPENDING
BY DEMOGRAPHIC GROUP

	Frequency	Mean	Median	SD
All	2556	0	−0.07	0.86
Age 18–34	2556	0	−0.04	0.91
Age 35–44	2556	0	−0.06	0.94
Age 45–54	2556	0	−0.09	1.01
Age 55–64	2556	0	−0.11	1.19
All males	2556	0	−0.07	0.91
All females	2556	0	−0.07	0.91
Males				
Age 18–34	2556	0	−0.07	1.11
Age 35–44	2556	0	−0.08	0.99
Age 45–54	2556	0	−0.1	1.07
Age 55–64	2554	0	−0.09	1.28
Females				
Age 18–34	2556	0	−0.04	0.97
Age 35–44	2556	0	−0.07	1.04
Age 45–54	2556	0	−0.11	1.11
Age 55–64	2554	0	−0.12	1.32

deal of noise in the data. It was possible to remove the seasonality or to fit 2 regression lines to the data. Ultimately, the choice was made to remove the seasonality and fit a single regression line (details available at http://jdc.jefferson.edu/healthpolicyfaculty/60/).

The final model chosen, presented in Equation 2 in the analysis section, contains only daily effects. The hypothesis that there is an annual component to the health care cost curve was not demonstrated in this study. There also were no monthly effects. Once the effects in daily spending were accounted for, month and year dummies were insignificant. As a result, managed care changes around benefit design and plan renegotiation did not show up in the cost curve of this case study as an annual "reset"; the growth in health care costs was continuous over the 7 years of the study.

An analysis of outpatient deductibles was performed as part of the analysis to determine why the health care cost curve was continuous in this case. The evidence strongly suggested that the annual reset in deductibles accounted for at least part of the effect. At the beginning of each year, patients pay high deductibles for care. As the year goes on, the graph suggests that patients start to reach their out-of-pocket limits, and deductibles actually paid fell toward zero before resetting (Fig. 4). This explanation could not be confirmed because of the anonymity involved in the MarketScan data. However, it is suggestive of the power of payment methods for smoothing utilization, and provided hypotheses to be explored further in future work.

Goodness of fit

The power of the model to explain the variation in data varied across groups. The model accounted for approximately 70% of the variation in spending for all groups aggregated based on the adjusted R^2. The ability of the model to fit the cost curve varied from below 50% to above 60% in various age/sex categories. The lowest adjusted R^2 was for the youngest groups, despite the fact that the count of individuals was highest. This may have reflected the fact that, for the youngest group, spending growth is hardest to predict. For these groups, the predicted trend was farthest from the

FIG. 4. Average outpatient deductible by day.

experienced trend, so they should receive particular scrutiny when quality improvement studies are performed.

Discussion

This case study generates 2 findings that are not well explored in the literature on the health care cost curve. The first is that the cost curve is continuous within and across years. Although it is true that each year has higher spending, and each month has higher spending, each day is expected to have higher spending than the one before. Thus, the cost curve is not subject to annual resets or "jumps" when contracts are renegotiated but rather rises throughout the year. The second finding is that there is no single cost curve. The growth rate and predictability of health care spending growth differ by population group, which is distinct from the observation that the level of spending differs by population group.

This also means that models that fit the cost curve on a PMPY or PMPM basis are potentially losing a large amount of explainable variation in rising medical costs. Possible reasons for this smooth rise in medical spending include the tendency for annual limits on deductibles, co-pays, and co-insurance to tend to kick in toward the end of the year. However, this phenomenon needs further research in order to be better understood.

The choice to model trend continuously has both positive and negative effects. The main positive is that medical encounters occur in continuous time, so the growth in spending also could be a continuous time phenomenon (or at least one that is best modeled on a daily basis). Aggregating at the quarterly or annual level could obscure the true time series properties of spending growth if the process is continuous. The main downside is the difficulty in interpreting the results. If the level of spending is the same on 1/1/2000 and 12/31/2006, then the average log trend will be zero even if the spending was generally increasing over time.

Conclusion

In summary, this case study produced findings that are consistent with the literature on the health care cost curve, while also generating new findings. The cost curve is difficult to predict, and any model is likely to leave a significant amount of unexplained variation in medical spending growth. A large part of this is because the health care cost curve is not smooth but rather contains a great deal of noise around the trend line.

In general, studies utilizing the health care cost curve will need to include analyses that are more selective with respect to timing. Researchers should choose to model costs for a population and time horizon that fit their intervention rather than one of convenience. The use of PMPY and PMPM calculations can help to standardize calculations of affect size, but the unit of analysis should fit the unit of intervention (ie, researchers should not use a PMPM measure for programs that do not have a monthly effect).

In order to bend the cost curve, we first must understand it, and developing better models is a logical starting point. Given that the results show that costs do rise on a continuous basis, the daily model presented here should be part of the modeling of the health care cost curve. A "horse race" that compares different methods in different situations would be
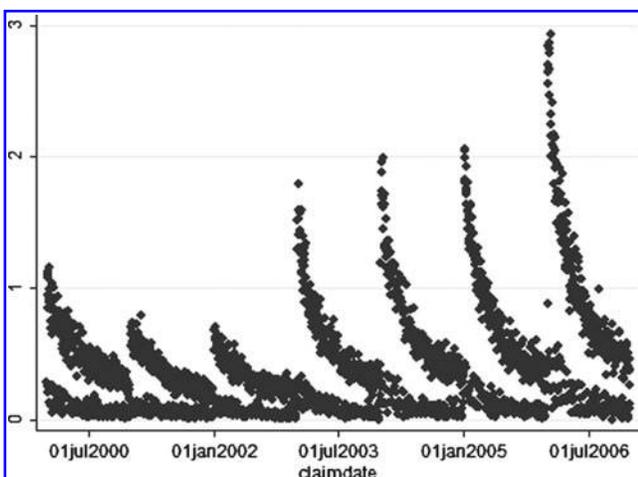
ideal, and is part of the scope for follow-up to the current study. Then, researchers and clinical decision makers can better justify their choice of unit of analysis for the health care cost curve in order to achieve the most accurate, impactful findings.

Acknowledgments

Author Disclosure Statement

References

1. McClellan M. Reforming payments to healthcare providers: The key to slowing healthcare cost growth while improving quality? J Econ Persp 2011;25:69–92.

2. Kongstvedt PR. *Essentials of Managed Health Care.* 6th ed. Burlington, MA: Jones and Bartlett Learning; 2013.

3. Keehan SP, Cuckler GA, Sisko AM, et al. National health expenditure projections: Modest annual growth until coverage expands and economic growth accelerates. Health Aff (Millwood) 2012;31:1600–1612.

4. Bundorf MK, Royalty A, Baker LC. Health care cost growth among the privately insured. Health Aff (Millwood) 2009;28:1294–1304.

5. Sisko AM, Truffer CJ, Keehan SP, Poisal JA, Clemens MK, Madison AJ. National health spending projections: The estimated impact of reform through 2019. *Health Aff (Millwood)* 2010;29:1–9.

6. Thompson Reuters. *MarketScan Studies: Abbreviated Bibliography*. Ann Arbor, MI: Thomson Reuters; 2011.

7. Harrison PL, Hara PA, Pope JE, Young MC, Rula EY. The impact of post-discharge telephonic follow-up on hospital readmissions. Popul Health Manag 2011;14:27–32.

8. Duffie D, Protter P. From discrete to continuous time finance: weak convergence of the financial gain process. Math Fin 1992;2:1–15.

9. Adamson DM, Chang S, Hansen LG, eds. *Health Research Data for the Real World: The MarketScan Databases.* Ann Arbor, MI: Thomson Reuters; 2008.

Address correspondence to:
Robert D. Lieberthal, PhD
Jefferson School of Population Health
Thomas Jefferson University
901 Walnut St.
10th Floor
Philadelphia, PA 19107

E-mail: robert.lieberthal@jefferson.edu