Geographic Variation in Health Care Spending

$10,300 to $13,900
$8,600 to $10,300
$7,800 to $8,600
$6,900 to $7,800
$5,200 to $6,900

Medicare Spending per Beneficiary, 2005

FEBRUARY 2008

A CBO PAPER

Geographic Variation in Health Care Spending

February 2008

The Congress of the United States ■ Congressional Budget Office

Notes

The map on the cover was generated by the Congressional Budget Office from data supplied by the Centers for Medicare and Medicaid Services. The map illustrates Medicare spending per beneficiary in the fee-for-service program on the basis of beneficiaries' residences. (The data were adjusted for age, sex, and race.) The geographic unit is the hospital referral region, as defined by the Dartmouth Atlas of Health Care. Unshaded areas are places without residents, such as national parks, forests, lakes, and islands.

Unless otherwise noted, all years are calendar years. Because of rounding, the sums of numbers in the text and tables may not equal totals.

Preface

Per capita health care spending and patterns of medical practice vary widely across the United States. In this paper, written at the request of the Chairman of the Senate Budget Committee, the Congressional Budget Office (CBO) examines the amount of geographic variation in spending, the reasons for that variation, and its implications for evaluating the efficiency of the health care system. In keeping with CBO's mandate to provide objective, impartial analysis, the paper makes no policy recommendations.

Peter R. Orszag
Director

February 2008

Contents

Table

Figures

Box

Geographic Variation in Health Care Spending

Summary and Introduction

Per capita health care spending varies widely across the United States. In 2004, as an example, per capita spending ranged from roughly $4,000 in Utah to $6,700 in Massachusetts (see Figure 1). The variation is even greater among smaller geographic units and among individual medical providers (see Figure 2). Among large hospitals in California from 1999 to 2003, Medicare spending per patient in the last two years of life ranged more than four-fold, from less than $20,000 to almost $90,000 (Wennberg and others 2005). Researchers affiliated with the Dartmouth Atlas of Health Care estimate that among groups of Medicare beneficiaries who are otherwise similar, individuals who live in high-spending areas receive approximately 60 percent more in services than do those who live in low-spending areas (Fisher and others 2003a).[1] The amount of spending involved is quite large—one report indicated that Medicare spending would fall by 29 percent if spending in medium- and high-spending regions were the same as that in low-spending regions (Wennberg and others 2002).

Large differences across the country in spending for the care of similar patients could indicate a health care system that is not as efficient as it could be, particularly if that higher spending does not produce commensurately better care or improved health outcomes. Given the importance of health care spending in the nation's long-term fiscal outlook, identifying and encouraging patterns of care that are more efficient is clearly important. Because Medicare is paid for in part by federal taxes, high spending in one area is, in effect, funded to some extent by taxpayers in other areas, raising additional concerns.

This paper by the Congressional Budget Office (CBO) examines the amount of and trends in geographic variation in health care spending, and the root causes of that variation. It also examines the relationship between spending and quality of care, and it discusses what those findings imply about how health care is produced in the United States and how it could be made more efficient. The paper focuses primarily on spending in the Medicare program because there are more data available about the cost of providing health care to Medicare beneficiaries than there are for other populations. Also, as the largest federal health care program, Medicare is highly relevant to and directly influenced by federal policy.

The results of current research indicate the following:

■ Two factors—the prices of health care services and severity of illness—are important in explaining geographic variation in health care spending. Different studies support different conclusions about the relative importance of those two factors, but most concur that together they account for less than half (and possibly much less than half) of the geographic variation in spending.

■ Income and the preferences of individuals for specific types of care appear to explain little of the variation in spending.

---

1. According to its Web site (www.dartmouthatlas.org), the Dartmouth Atlas Project "works to accurately describe how medical resources are distributed and used in the United States. The project offers comprehensive information and analysis about national, regional, and local markets, as well as individual hospitals and their affiliated physicians, in order to provide a basis for improving health and health systems."

Figure 1.

Health Care Spending per Capita, 2004

(Dollars)

Source: Congressional Budget Office based on data from the Centers for Medicare and Medicaid Services.

■ A substantial portion of the variation remains unexplained after those factors are considered. Unmeasured differences in demand for care could be important, but some of the variation in medical practice probably is attributable to regional differences in the supply of medical resources (specialist physicians or health care facilities, for example) and the propensity to take advantage of the financial incentives provided by Medicare or other payers in developing and using those resources.

■ Some regions appear more prone to adopt low-cost, highly effective patterns of care whereas others are more prone to adopt high-cost patterns of care and to deliver treatments that provide little benefit or are even harmful.

■ Geographic variation in total health care spending per capita has shown an upward trend in recent years; over the past three decades, in contrast, variation in Medicare spending has narrowed sharply. That reduction could be the result of changes in Medicare's reimbursement policies.

The evidence suggests that efficiency gains in the health care system are possible: Spending in high-spending regions could be reduced without producing worse outcomes, on average, or reductions in the quality of care. But policies that reduce spending in high-spending areas would not necessarily lead to increased efficiency—and could result in worse health outcomes—unless the reductions targeted ineffective or harmful treatments. Reforms that are designed to increase efficiency in the health care sector generally could, as a side effect, reduce geographic variation. Some of those proposals are discussed briefly toward the end of this paper. (CBO is undertaking an expanded effort to examine options for modifying the health care system in the United States. The discussion of those options will include their potential impact on geographic variation.)

Measuring Geographic Variation

Researchers generally use relatively large areas, such as states, as the unit of analysis to examine geographic variation in health care spending. Doing so allows them to identify factors that vary systematically across areas and that affect regional patterns of care. Researchers also have analyzed variation among smaller geographic units, such as hospital referral regions (HRRs), counties, and

Figure 2.

Medicare Spending per Beneficiary, by Hospital Referral Region, 2005

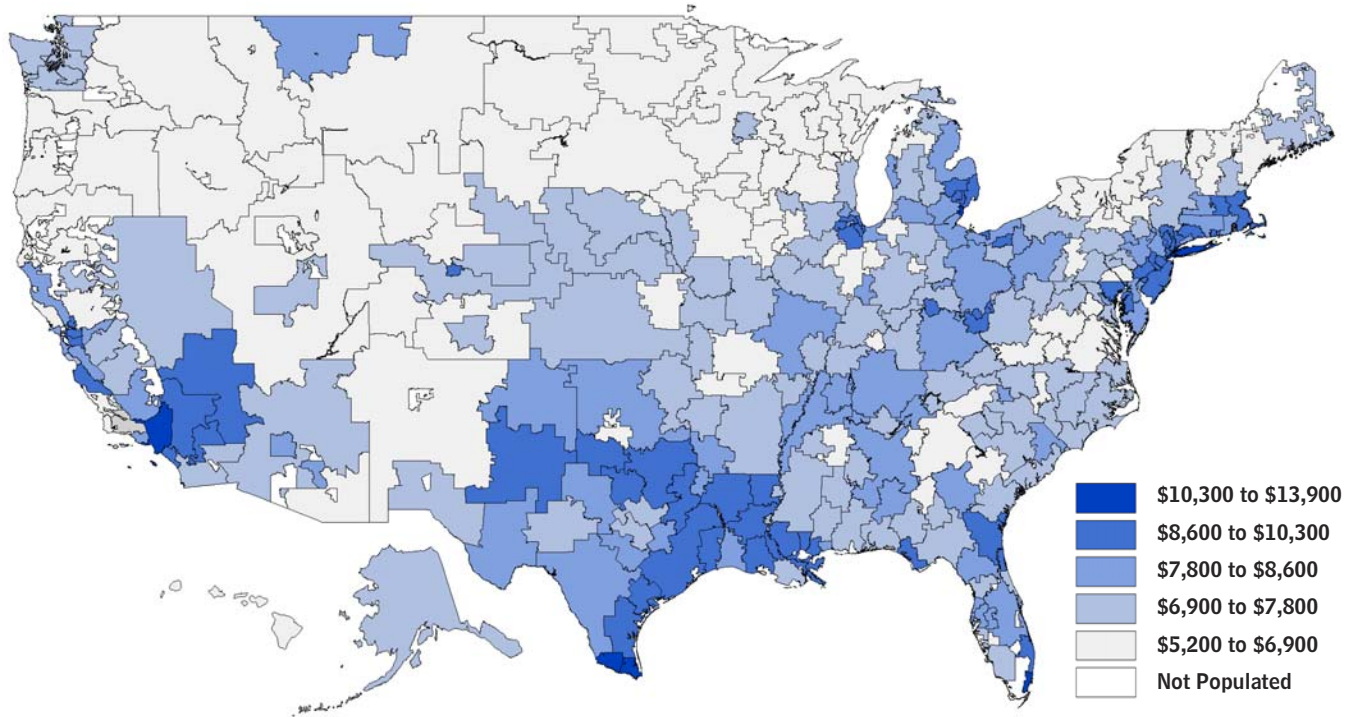|	
	$10,300 to $13,900
	$8,600 to $10,300
	$7,800 to $8,600
	$6,900 to $7,800
	$5,200 to $6,900
	Not Populated

Source: Congressional Budget Office based on data from the Centers for Medicare and Medicaid Services.

Note: The data are for Medicare spending per beneficiary in the fee-for-service program on the basis of beneficiaries' residences and adjusted for age, sex, and race. The geographic unit is the hospital referral region, as defined by the Dartmouth Atlas of Health Care. Areas labeled ''Not Populated'' include places without residents, such as national parks, forests, lakes, and islands.

metropolitan statistical areas (MSAs), and among individual hospitals or medical centers.[2]

Much of the research on geographic variation in health care spending has focused on Medicare spending, at least in part because of the availability of billing records for Medicare's fee-for-service patients. The data provide details about where beneficiaries live (state, county, and zip code) and on the use of and spending for services covered by Medicare. There is no comparable data source for the privately insured population that both covers a large segment of the population and includes detailed information on geography and spending. Variation among states in Medicare spending per beneficiary in 2004 was similar to the variation in total personal health care spending: It

ranged from $5,600 per beneficiary in South Dakota to $8,700 in Louisiana.

The coefficient of variation (COV) is a commonly used statistic for quantifying the degree of variation in a variable. The COV is the standard deviation divided by the mean.[3] (As a simple illustration, the mean height of men in the United States is about 69 inches, and the standard deviation is about 3 inches—a fairly narrow distribution—so the COV is 3 divided by 69, or 0.04.) Geographic variation in Medicare spending per beneficiary is substantially larger: In 2005, the COV in state-level Medicare spending per beneficiary was 0.11.

Another way to measure geographic variation is to calculate the amount by which total spending would be

---

2. The Dartmouth Atlas Project defined HRRs on the basis of referral patterns for inpatient surgical procedures. The United States is divided into 306 HRRs, most of which are larger than counties and some of which cross state boundaries (Center for the Evaluative Clinical Sciences 1999).

3. The standard deviation of a set of values equals the square root of the variance. The variance equals the mean of the square of the difference between each value and the mean of the set of values.

reduced if spending in high-spending areas were reduced to that of low-spending areas. The Dartmouth Atlas researchers undertook such an analysis using Medicare data (Wennberg and others 2002). They calculated that Medicare spending would fall by 29 percent if spending in medium- and high-spending regions were the same as in their benchmark regions, defined as those with spending in the lowest decile.

Geographic Variation in Context

As a first step in explaining the significance of geographic variation in health care spending, it is useful to analyze how that variation has changed since the 1970s, how variation in Medicare spending compares with variation in total health care spending, and how variation in health care spending compares with variation in spending on other goods and services. It also is useful to examine how the United States compares with other countries and how Medicare compares with the Department of Veterans Affairs (VA) health care system.

The following conclusions can be drawn:

■ Geographic variation in total health care spending per capita has been growing in recent years. Geographic variation in Medicare spending, in contrast, has dropped sharply over the past three decades and recently has been slightly lower than the variation in total health care spending.

■ There also is geographic variation in per capita spending on non-health care items, such as housing, food, and transportation (see Box 1). The degree of variation in Medicare spending per beneficiary is relatively high compared with those other spending categories, but it is not completely out of line.

■ In recent years, geographic variation in health care spending has been much higher in the United States than in Canada, and somewhat higher than in the United Kingdom. Financing of health care in those countries is more centralized than it is in the United States.

■ In recent years, geographic variation in spending in the VA health care system has been similar to that in Medicare, despite the fact that the VA system uses an explicit allocation formula to distribute funds to regions.

Trends in Geographic Variation

To examine changes in geographic variation over time, CBO analyzed two state-level measures of health care spending: total health care spending per capita and Medicare spending per Medicare beneficiary. Total spending per capita was calculated from data for 1991 through 2004 published by the Centers for Medicare and Medicaid Services, or CMS (2007). Medicare spending per beneficiary was calculated using the Continuous Medicare History Sample, a dataset created by CMS that includes spending data for a sample of 5 percent of Medicare's beneficiaries from 1974 through 2005.

The COV for total health care spending per capita rose gradually during the 1990s and early 2000s, from a low of 0.112 in 1991 to a high of 0.123 in 2004 (see Figure 3). In contrast, with the exception of the early- to mid-1990s, the COV for Medicare spending showed a dramatic downward trend. From a peak of 0.200 in 1976 it fell sharply, to 0.125 by 1991, and then rebounded in the early 1990s before resuming a sharp decline, ending at 0.110 in 2005.

To examine the reasons for the trends in Medicare, CBO decomposed the geographic variation in spending into

Figure 3.

Variation in State-Level Medicare and Overall Health Care Spending per Capita

(Coefficient of variation)

Source: Congressional Budget Office based on data from the Centers for Medicare and Medicaid Services.

Figure 4.

Contributions of Major Service Categories to State-Level Variation in Medicare Spending per Beneficiary

(Coefficient of variation)

Source: Congressional Budget Office based on data from the Centers for Medicare and Medicaid Services.

four components representing major categories: spending on care at "short-stay" hospitals (as opposed to longer-term care in another kind of facility), spending for services provided by physicians and laboratories, spending on post–acute care (provided in skilled nursing or long-term care facilities, as home health care, or as hospice care), and outpatient spending (for "in-and-out surgery," for example; see Figure 4). Each category's contribution to the overall COV for Medicare depends both on its share of total spending and on the degree of geographic variation within the category. That decomposition reveals several trends:

■ In the 1970s, spending on inpatient hospital care accounted for a substantial amount of geographic variation in Medicare spending. But the amount of variation attributable to hospital spending declined throughout the 1970s and 1980s, leveled off in the 1990s, then declined again in the early 2000s.

■ The large spike in geographic variation in Medicare spending that occurred in the early 1990s was attributable to increasing variation in post–acute care spending.

■ In the early 2000s, geographic variation in Medicare spending declined, and each major spending component contributed to that decline.

Why did geographic variation in Medicare spending change so dramatically? One hypothesis is that revisions in Medicare's payment policies contributed substantially.[4] In the 1980s, Medicare gradually began to phase out cost- and charge-based reimbursement and to implement a set of formula-based prospective payment systems. In a cost-based system, Medicare payments to providers are determined by the costs incurred, including labor and capital costs. In a charge-based reimbursement system, payments are determined by the charges or fees submitted. Under cost- and charge-based systems, providers had considerable influence over payment rates, which in turn allowed for substantial geographic variation in payments.

CBO tested the hypothesis concerning reimbursement methods by examining average payment rates for hospital stays and for physicians' and laboratory services.[5] Beginning in 1983, hospitals were switched by Medicare to a formula-based payment system that uses the discharge as the unit for payment. Since 1983, there has been much less state-to-state dispersion in average Medicare payment rates for hospitals (see Figure 5). However, the dispersion in Medicare's hospital payment rates already was declining sharply before then, so other forces (such as the anticipation of the policy change or a stricter review of hospital charges leading up to the policy enactment) could have been at work. In 1992, physicians' reimbursement changed from the cost-based system to Medicare's "resource-based relative value scale" (see Figure 6 on page 9). The dispersion of average payment rates for physicians' and laboratory services was increasing in the

---

4. The link between the implementation of a new rate-setting system in Medicare and a reduction in geographic variation in payment rates was noted by Pope and others (1989).

5. Average payment rates are calculated by dividing total Medicare payments by the units of service provided. The payment rates are not standardized to a uniform basket of services but instead reflect changes in service mix and payment rate for each type of service. Limitations in the underlying data prevented CBO from calculating a standardized payment rate. In addition to medical services provided in a doctor's office, billing for "physicians' and laboratory services" includes billing for services that are provided in a hospital but billed separately from the hospital's charges. The category includes surgical services, laboratory and diagnostic services, and durable medical equipment.

Box 1.

Geographic Variation in Medicare Spending Compared with Spending on Food, Housing, and Transportation

The Congressional Budget Office (CBO) used published results from the Bureau of Labor Statistics' 2004–2005 Consumer Expenditure Survey to compare geographic variation in Medicare spending with variation in spending for food, housing, and transportation.[1] The survey's sampling strategy prevents its use for generating estimates at the state level, but it can be used to measure per capita spending in 24 large metropolitan areas.[2]

CBO measured geographic variation for spending on food, housing, and transportation because, like health care, those categories represent substantial shares of the economy. To make the comparison with Medicare spending valid, Medicare spending per beneficiary was measured separately for each metropolitan area, and a coefficient of variation (COV) for Medicare spending was calculated from data only from those areas.

Geographic variation was fairly similar among the spending categories examined. Per capita spending on

food ranged from $1,880 in Baltimore to $3,010 in Boston, with a metropolitan COV of 0.120 (see the table to the right). There was greater variation in per capita spending on housing: Pittsburgh was lowest, at $5,231, and San Francisco was highest, at $8,802 (COV, 0.143). Per capita spending for transportation ranged from $2,416 in Miami to $5,038 in Anchorage (COV, 0.143). The analogous COV in Medicare spending per beneficiary was 0.148, which is slightly higher than the COVs for housing and transportation.

Differences in income may account for some of this variation. CBO used a regression analysis to examine the extent to which geographic variation in income accounts for variation in spending per person on Medicare and on other goods and services (see the table). A COV was calculated to measure the degree of geographic variation, after controlling for income.

Spending per capita on food and on housing are strongly and positively correlated with income. For Medicare spending (and for transportation), in contrast, there was little or no relationship between per capita income and per capita spending, so the income-adjusted and unadjusted COVs are similar for those spending categories. (Medicare spending could be less tied to income because the program's spending is financed largely by federal taxes, rather than purchased by the individual.) After controlling for income, geographic variation in Medicare spending appears even greater relative to variations in spending for food and housing; it remains similar to geographic variation in spending for transportation.

---

1. Bureau of Labor Statistics, *Current Metropolitan Statistical Areas Tables, 2004–2005,* www.bls.gov/cex/home.htm. Data limitations prevent an analogous analysis of total per capita health care spending by metropolitan area.

2. Anchorage, Alaska; Atlanta, Georgia; Baltimore, Maryland; Boston, Massachusetts; Chicago, Illinois; Cleveland, Ohio; Dallas-Fort Worth, Texas; Denver, Colorado; Detroit, Michigan; Honolulu, Hawaii; Houston, Texas; Los Angeles, California; Miami, Florida; Minneapolis-St. Paul, Minnesota; New York, New York; Philadelphia, Pennsylvania; Phoenix, Arizona; Pittsburgh, Pennsylvania; Portland, Oregon; San Diego and San Francisco, California; Seattle, Washington; St. Louis, Missouri; and Washington, D.C.

Box 1.

Continued

Analysis of Geographic Variation in Spending on Medicare and Other Goods and Services, 2004 to 2005

| | Mean Spending per Person (Dollars) | Unadjusted Coefficient of Variation | Results of Regression of Spending on Income | | Income Elasticity | Income-Adjusted Coefficient of Variation |
			Estimated Coefficient on Income	R^2 of Regression Equation		
Food	2,537	0.120	0.046 [a]	0.331	0.472	0.098
Housing	6,955	0.143	0.226 [a]	0.757	0.854	0.071
Transportation	3,339	0.143	0.011	0.007	0.084	0.143
Medicare	7,814	0.148	-0.104	0.105	-0.349	0.140

Sources: Bureau of Labor Statistics, 2004-2005 Consumer Expenditure Survey, Current Metropolitan Statistical Area Tables; and Congressional Budget Office calculations based on sample data from the Centers for Medicare and Medicaid Services' Continuous Medicare History Sample data.

Notes: R^2 = estimated share of metropolitan-level variation in spending that is explained by income.

Spending per person is measured for 24 selected metropolitan statistical areas (MSAs). Medicare spending is measured per beneficiary in 2005, for fee-for-service beneficiaries only. The unadjusted coefficient of variation (COV) is calculated using unadjusted spending per person and is weighted by population. The estimated coefficient on income is from a weighted ordinary least squares (OLS) regression of MSA-level spending per person on MSA-level pretax income per capita; it represents the estimated change in spending per person (in dollars) for goods in the indicated category in response to a one-dollar increase in pretax income per capita. The MSA-level income elasticity is calculated at the overall mean spending and income; it equals the estimated coefficient on income divided by the ratio of mean spending per person to mean income. The income-adjusted COV is calculated using adjusted spending (overall mean plus the residual from the OLS model) and is weighted by population.

a. $p < 0.01$.

Although the variation in Medicare spending is not completely out of line with that observed in other sectors of the economy, it does warrant closer examination. It is reasonable to assume that the value of the goods and services consumed in most other sectors is readily apparent. For example, if people in a given area spend a relatively large amount on housing, it is reasonable to assume that they are aware of the attendant benefits (for example, the area might provide better job opportunities, have a mild climate, offer access to cultural amenities, or have good public schools) and they choose to spend more as a result. That assumption does not necessarily hold for health care. Health care providers usually have a strong influence on the choice of treatment, and the quality or value of the benefits received from higher spending is much more difficult for patients to discern.

Figure 5.

Dispersion in State-Level Mean Medicare Payments per Hospital Stay

(Ratio)

Source: Congressional Budget Office based on data from the Centers for Medicare and Medicaid Services.

years leading up to 1992, but it declined sharply after the new system's introduction. (Note that the dispersion of *overall* spending for physicians' and laboratory services, shown in Figure 4 on page 5, did not decline as much.)

Geographic Variation in Canada and the United Kingdom

Geographic variation in health care spending in the United States could be related to idiosyncrasies in the nation's system of health care financing and delivery. The United States differs from most other high-income countries in having a relatively decentralized system with a relatively large role for private insurers. The share of the population with health insurance varies from region to region, as do the type and the comprehensiveness of that insurance coverage. It has been hypothesized that, relative to other countries, the United States might therefore exhibit a high degree of geographic variation in health care use and spending.

To test that hypothesis, CBO used publicly available data to compare variation in health care spending per capita among states in the United States, among provinces in Canada, and among regions in the United Kingdom (see Figure 7).[6] Those countries were chosen because they are similar to the United States in many respects (they have

comparable per capita income and systems of governance, for example) and because regional data on health care spending are available for all three countries.

Geographic variation in health care spending has consistently been much higher in the United States than in Canada and somewhat higher than in the United Kingdom in the years for which data are available. From 1991 through 2004, the COV in state-level health care spending per capita in the United States varied between 0.112 and 0.123. Over the same period, the COV in per capita spending by province in Canada (for public and private spending) varied between 0.059 and 0.088, with an increase in recent years. In the United Kingdom, the COV by region has varied in recent years between 0.091 and 0.107.

The greater variation within the United States is not surprising given that the health care systems in Canada and the United Kingdom are explicitly designed to distribute funds from the central governments to the province or region according to "needs-based" formulas. In Canada, health care is financed jointly by the federal, provincial, and territorial governments and other sources. Funds are explicitly allocated from the federal government, partly on a uniform per capita basis through the Canada Health Transfer and partly through the Equalization Program, which is designed to counteract disparities among provinces in the capacity to provide comparable health services.

In the early decades of its National Health Service, the United Kingdom allocated funds to different regions on the basis of historical spending in each region, updated

---

6. To allow valid comparisons, geographic units were chosen for Canada and the United Kingdom that were similar to U.S. states in average population. The weighted-average population per geographic unit is about 8 million for a Canadian province, about 6 million for a region in the United Kingdom, and about 13 million for a U.S. state. COVs are weighted using the population of each province, region, or state as the weight. In Figure 7, data are presented for all years for which data are available. Spending in the United States and Canada is measured by calendar year and includes public and private expenditure. Spending in the United Kingdom is measured by "financial year" (April through March) and includes only public expenditure, which historically has been more than 80 percent of total health care expenditure. Spending data for Canada for 2005 and 2006 are estimated.

Figure 6.

Dispersion in State-Level Mean Medicare Payments per Physician or Laboratory Service

(Ratio)

Source: Congressional Budget Office based on data from the Centers for Medicare and Medicaid Services.

for inflation. Beginning in the 1970s, researchers began to link that approach to financing with unequal regional distributions of funds (Culyer and others 1981, European Observatory on Health Systems and Policies 1999). Those investigations culminated in a plan developed in 1976 by the Resource Allocation Working Group, which laid out a formula for regional health care financing that was based on health care needs and local differences in practice costs. Over the next decade the formula was adjusted to reduce regional disparities.

Variation in Health Care Spending by the Department of Veterans Affairs

The VA health care system is an example of centrally budgeted health care in the United States. It is a large, integrated system that typifies managed care, particularly of the kind practiced by health maintenance organizations, or HMOs. Services are provided primarily through a limited network of staff physicians and hospitals owned and operated directly by the department.

VA's financing structure differs from Medicare's in important ways, allowing for an interesting comparison both with Medicare and with the rest of the U.S. health care system. First, VA operates under a global budget that is

determined by Congressional appropriations. Medicare benefits, in contrast, are paid through an entitlement program that does not require a specific appropriation each year. Second, VA funds are allocated to 21 geographically defined units (called Veterans Integrated Services Networks, or VISNs) on the basis of the number of veterans served and their health care needs. The flow of Medicare funds, in contrast, is determined on the basis of the volume of health care services provided in each region. (Both systems adjust for local input costs.) Because of those differences in financing systems, a reasonable hypothesis is that VA health care would exhibit less variation in spending per capita than Medicare.

CBO used VA data for fiscal years 2001 and 2007 (Department of Veterans Affairs 2001, 2007) and data from the General Accounting Office (2002) to measure geographic variation in VA spending. Variation was measured for two separate years to identify changes over time. The COV by VISN-level allocations per patient in the VA system was 0.085 in 2001 and 0.104 in 2007.[7] To allow a valid comparison with Medicare, CBO grouped Medicare beneficiaries, based on residence, into geographic areas matching those of the VISNs and measured VISN-level Medicare spending per beneficiary. The result was a Medicare COV of 0.141 for 2001 and a COV of 0.116 for 2005 (the most recent year for which data are available).

The calculations show that, in 2001, Medicare exhibited substantially more geographic variation in health care spending per person than the VA system did. Since then, however, the gap appears to have been largely eliminated, as variation in spending in the Medicare program fell while that in the VA program increased. It appears, therefore, that the centrally budgeted VA system does not display much less geographic variation in spending than is exhibited in the unbudgeted Medicare program.

Why did geographic variation in VA spending increase between 2001 and 2007? One possibility is the introduction of a more complex methodology for adjusting allocations based on case mix. (*Case mix* refers to health needs of the population served.) The 2001 VA allocations were

---

7. Patients "in the VA system" are veterans who actually use VA health services, not the broader population of veterans who could be eligible. The number of patients in the VA system is projected separately for each VISN (based on historical trends) as part of the centralized funding process.

Figure 7.

Geographic Variation in Health Care Spending per Capita in Selected Countries

(Coefficient of variation)

Source: Congressional Budget Office based on data from the Centers for Medicare and Medicaid Services, HM Treasury (for United Kingdom data), and the Canadian Institute for Health Information.

based on a case mix system that had only three patient groups and three payment rates. By 2007, the system differentiated among 20 patient groups, each with a separate payment rate. Between 2001 and 2007, the VA methodology also was refined to include an allocation adjustment for treatment of unusually high-cost patients ("outliers"). Both refinements have been described as significant improvements, and both might have contributed to the increase in geographic variation in VA spending.

In addition to exhibiting geographic variation in spending, the VA system shows substantial variation in patterns of clinical practice despite the fact that VA's management tracks providers' compliance with national guidelines for the treatment of many medical conditions. Several studies have documented wide geographic differences within the system in patterns of treatment for several medical conditions: acute myocardial infarction (heart attack), upper respiratory infection, depression, and prostate cancer (Aspinall and others 2005, Fortney and others 1996, Subramanian and others 2002, Wilt and others 1999).

The implication is that local norms can influence practice patterns, even in a relatively centralized system that places a strong institutional emphasis on adherence to clinical guidelines for care.

The evolution of the regional financing system for VA health care has strong parallels with the development of the regional health financing formula in the United Kingdom. In each case, funds initially were allocated to regions, primarily on the basis of historical costs that had been adjusted for inflation. Each system's administrators recognized later that the result was an inequitable distribution of funds; that conclusion in turn led to the implementation of regional allocation formulas based on population, health status, and local practice costs. Iglehart (1996) has reviewed and described changes in the VA financing system.

Explaining Geographic Variation in Health Care Spending

Several researchers have examined explanations for geographic variation in per capita health care spending (see Table 1); most of their studies focus on the Medicare fee-for-service program, largely because better data are available for Medicare than for the private sector. The typical approach has been to measure geographic variation in unadjusted spending per capita and then to measure variation in spending per capita after adjusting for various factors that are believed to affect spending. The contribution of a given factor to geographic variation is measured by the degree to which variation is reduced after adjusting for that factor.

Those factors can be divided into four broad categories, each discussed in detail in the following sections:

■ Prices paid for medical services,

■ Health and illness status of residents of a given region,

■ Regional preferences about the use of health care services (and the determinants of those preferences, such as income), and

■ Residual (unexplained) variation.

Table 1.

Research on Geographic Variation in Health Care Spending

Study	Type of Spending	Explanatory Factors
Welch and others (1993)	Medicare physician spending per beneficiary, 1989	Inpatient hospital admission rate, physicians per capita, proportion of physicians engaged in primary care
Cutler and Sheiner (1999)	Medicare spending per beneficiary, 1995	Health risk behaviors, mortality rates, race, income, education, HMO market share, supply of medical providers
Gage, Moon, and Chi (1999)	Medicare spending per beneficiary, 1995	Share of beneficiary population under age 65, share over age 85
Center for the Evaluative Clinical Sciences (1999)	Medicare spending per beneficiary, various years	Age, sex, race, illness, prices, HMO market share, supply of medical providers
Fuchs, McClellan, and Skinner (2001)	Utilization of Medicare-covered services per beneficiary, 1989–1991	Education, income, cigarette sales, obesity, air pollution, race, region, urbanization
Medicare Payment Advisory Commission (2003)	Medicare spending per beneficiary, 2000	Prices, health status, Medicare Part A and Part B participation rates, special hospital payments
Super (2003)	Medicare spending per beneficiary, various years	Health status, local practice costs, special payments to hospitals, managed care enrollment, intensity of care
Gold (2004)	Medicare spending per beneficiary, various years	Population characteristics, health care needs, prices, intensity of care
Hadley and others (2006)	Medicare spending per beneficiary, 1992–2002	Age, race, urbanization, health status, "end-of-life expenditure index" from the Dartmouth Atlas, tobacco use, educational attainment, income, Medicare payment policy, dual-beneficiary status, percentage of physicians in primary care
Martin and others (2007)	Overall health care spending per capita, Medicare spending per beneficiary, Medicaid spending per beneficiary, 2004	Income, availability of physicians and hospitals, Medicaid eligibility and benefits, age (descriptive analysis only)

Source: Congressional Budget Office.

Notes: HMO = health maintenance organization; Medicare Part A is Hospital Insurance and Part B is Supplementary Medical Insurance.

Complete citations to the studies are given in the references list on page 27.

The studies cited in Table 1 differ in data sources and methods, but several conclusions are possible:

■ Differences among regions in the prices of medical services and in the population's health status explain some of the observed geographic variation in Medicare spending. The amount of variation explained by those factors is most likely less than half of total variation, and possibly much less.

■ Demographic factors (including income, race, and educational attainment) and patients' treatment preferences contribute only a small amount to geographic variation.

■ Much or most of the geographic variation is residual; it cannot be explained by prices, health status, demographics, or treatment preferences.

The existence of residual variation in health care spending implies that populations in different areas—even though they might face similar prices, have similar average severity of illness, and prefer to receive similar types of medical services—nonetheless receive different quantities of services. Because those factors (especially as they relate to treatment preferences) are not measured perfectly, however, the amount of true residual variation might be smaller than implied by empirical analysis.

The explanations for residual variation are difficult to assess quantitatively, but the following factors appear to be significant:

■ Disagreement among medical professionals regarding the appropriateness of some treatments is associated with variations in the use of those treatments.

■ The financial pressures and incentives facing medical providers vary geographically, as does the response of medical providers to those incentives.

■ The supply of physicians and other medical resources varies geographically; is strongly related to use of services and to spending; and appears to be driven, at least in part, by factors unrelated to health status or the demand for health care services. Flexibility in the norms of medical practice might allow such supply variations to persist, particularly in the context of fee-for-service reimbursement.

Variation in Prices and Practice Costs
Health care spending equals the product of the quantity of health services performed and the price paid per service. Here the "price" is the total amount paid to the medical provider in exchange for a specific service, including payment from the insurer and out-of-pocket spending by the patient. The fact that prices for health care services are higher in some regions than in others accounts for some of the geographic variation in health care spending per capita.

There can be many reasons for geographic variation in prices for health care services. The inputs used to produce medical care (such as facilities, supplies, and the services of health professionals) are more costly in some areas than others. In Medicare's fee-for-service program, payments to physicians begin with a uniform national base rate and are adjusted by a measure of local practice costs that includes office rent, malpractice insurance, and the

opportunity cost of the health professional's time (which is estimated from data on the earnings of other local professionals). Payments to hospitals are adjusted by average hospital wages measured at the level of the MSA and other factors.

The private-sector price for a medical service, normally negotiated by the insurer and medical providers, additionally reflects the relative bargaining power of the parties. The Government Accountability Office, or GAO (2005), has reported that, after adjusting for local practice costs, private-sector prices for physicians' services in the highest-priced area were approximately twice those in the lowest-priced area. Hospital prices were distributed even more widely, with rates in the highest-priced area reaching 3.6 times those in the lowest-priced area. GAO attributed the variation in prices paid by insurers (beyond what could be explained by differences in local practice costs) to variation among regions in medical providers' bargaining strength.

Most studies that explain geographic variation control for differences in prices or practice costs but do not report the share of total variation explained by prices. One study, by the Medicare Payment Advisory Commission, or MedPAC (2003), reports that share; it found that variation in prices and practice costs explained about 29 percent of total variation in Medicare spending at the state level.

Variation in Health Status
The populations of some geographic areas are relatively sicker than the populations of others, and so they consume a disproportionately larger amount of health care resources. Policymakers generally consider that source of variation (in addition to variation based on local prices) as less of a policy concern: A basic function of health insurance, after all, is to make health care services available and affordable to people who need them. It is difficult to calculate the degree of geographic variation in health care spending that results from the underlying differences in health status, however, because of limitations in the available data on health status.

For any given person, health status is a strong predictor of spending on health care. One recent analysis of spending among individual Medicare beneficiaries (as opposed to spending by state or by region) reported that extensive health and disease status measures, including self-reports of health status, explained about 20 percent of total

variation (Hadley and others 2006). However, when people are aggregated into large regional groups, much of that variation is averaged out, necessarily limiting the amount of regional differences in health care spending that can be explained by differences in health status. The MedPAC study (Medicare Payment Advisory Commission 2003) showed that health status explained roughly 16 percent of total variation in Medicare spending by state, but that estimate may be overstated. The health status measure used was the hierarchical condition category (HCC), a measure used for risk adjustment in Medicare. The HCC is constructed in part from diagnoses and conditions as determined by physicians and other providers and, thus, captures regional differences in physician practice and in patient care-seeking behavior in addition to differences in underlying health status. Researchers with the Dartmouth Atlas Project, using data from 1996, concluded that the differences across HRRs in practice costs and in beneficiaries' demographics and health status together explained one-third of the variation in Medicare spending per beneficiary (Center for the Evaluative Clinical Sciences 1999).

By most accounts, differences in the costs of providing health care and in the underlying health conditions and needs of the population thus explain some of the observed geographic variation in Medicare spending. But together they appear to explain less than half the total variation, and possibly much less.

Variation in Demographics and Other Characteristics of Patients

Patients vary in the desire to receive aggressive, expensive care and heroic lifesaving measures. In the aggregate, such treatment preferences might vary from one region to the next because of cultural factors and thus might account for some of the geographic variation in health care spending. Those desires and preferences probably are related to income, which is known to be related to demand for health care (Newhouse 1993), but other demographic factors also could be important.

The quantitative evidence indicates that demographics do not explain much of the geographic variation in Medicare spending. In studies that report separate effects of income, race, and educational attainment on spending, those demographic factors explain less than 5 percent of the total variation (see, for example, Cutler and Sheiner 1999). Other researchers have reported that demographic variables, combined with health status variables, explain

less than one-third of spending overall. Separately, in a review of geographic variation in health care spending, Phelps (2000) noted that the extent to which income varies by region, combined with the extent to which demand for health care spending is related to income (as measured by the RAND Health Insurance Experiment of the 1970s and 1980s, for example), necessarily limits income differences to explaining a small portion of total variation in spending.[8] Any positive effect of income on demand for health care could be negated if people in higher-income regions also tend to be in better health (in ways that are not simultaneously controlled for in the analysis). In fact, the income elasticity reported in the table in Box 1 (page 7) implies that Medicare spending could be inversely related to income in an area.

Treatment preferences also do not appear to explain much of the variation in health care spending. One recent study combined a special survey of a sample of Medicare beneficiaries with data on end-of-life treatment and spending for that set of beneficiaries (Barnato and others 2007). No correlation was reported between patients' preferences for the type and intensity of treatment at the end of life and actual spending at the end of their lives. For example, 21.0 percent of respondents in the lowest-spending areas expressed a preference for mechanical ventilation that might prolong their life by one month, compared with 21.4 percent in the highest-spending regions (that difference was not statistically significant). Even though they are particularly difficult to measure, it seems unlikely that treatment preferences, aggregated by state or even by HRR, vary to a degree that matches variation in spending.

To what extent would any differences in health care spending that are caused by differences in treatment preferences amount to a potential policy concern? The answer is more ambiguous than in the case of differences in health status. If people in one place prefer to receive a set of expensive services and they finance those services locally, it is difficult to make the case for public policy concern. If, however, services are paid for with federal revenue, preference-driven variation in spending might be viewed as inequitable; expensive treatment preferences

---

8.   Higher income usually is associated with greater health care spending, although the magnitude of the relationship varies. For more information about the RAND Health Insurance Experiment, see, for example, Manning and others (1987).

in one area are then to some extent satisfied with money provided by people somewhere else.

Residual Variation

Residual geographic variation in health care spending is variation that cannot be explained readily by observed variation in local practice costs, health status, demographics, or treatment preferences. Researchers have reported that, after controlling for local practice costs, health status, and demographics, between one-half and three-fourths of total variation in spending remains unaccounted for. If one arrayed all regions from highest to lowest spending, after adjusting for the effects of differences in practice costs, health status, demographics, and treatment preferences, the regions in the top quintile (the 20 percent with the highest spending) appear to use between 30 percent and 80 percent more health care, on average, than do regions in the bottom quintile.[9] Differences between the extreme highest- and extreme lowest-spending regions are larger still (see, for example, Center for the Evaluative Clinical Sciences 2007, Fisher and others 2004).

Explanations for residual variation are more difficult to assess quantitatively. A substantial portion might be attributable to more subtle differences in health status, for example, that are difficult to measure. However, researchers generally attribute much of the remaining geographic variation to differences in the way medicine is practiced. The following sections explore those differences and some possible reasons for them.

Medical Uncertainty. There is relatively large geographic variation in the rate at which some surgical procedures are performed. Radical prostatectomy, the total removal of the prostate gland along with portions of surrounding tissue, is one example. Clinically similar patients are much more likely to receive the procedure in some areas of the country than in others. In contrast, rates for other procedures, such as surgical repair of hip fracture, exhibit relatively little geographic variation (Center for the Evaluative Clinical Sciences 1999). Two explanations are evident for the smaller variation: The diagnosis of hip fracture is straightforward, and there is a clear consensus

among clinicians and patients that hospitalization and surgical repair are appropriate. Conversely, there tends to be more geographic variation in the use of surgical procedures if medical professionals disagree about appropriate indications or if alternative nonsurgical treatments are available (Wennberg and others 1982).

Medical uncertainty and disagreement among professionals alone would not necessarily result in persistent variation in spending for large areas. (In particular, random differences among physicians would tend to cancel each other out if no other factors were involved.) The evidence suggests, however, that regional patterns in Medicare spending have been remarkably persistent.[10] Three other factors appear to combine with medical uncertainty to generate geographic variation: financial incentives, the supply of medical resources in an area, and the flexibility of standard patterns or "norms" of practice in medicine.

Financial Incentives. Several of the studies noted above included indicators of the financial environment or financial pressures within regions as factors that might explain variation in spending. One theory is that managed care would be associated with lower spending. "Managed care" encompasses various types of health plans (such as HMOs) and a range of techniques that those plans use to control use of services and to contain spending. Some researchers who have examined geographic variation in health care spending have included measures of the share of patients in HMOs in an area; they have not identified strong effects, perhaps because there is not enough regional variation in the prevalence of managed care to explain a large amount of variation in spending (Cutler and Sheiner 1999, Hadley and others 2006).[11]

Other research has analyzed whether for-profit hospitals charge more than nonprofit hospitals do. Because for-profit hospitals are not distributed uniformly across the country, any ownership-related differences in the way hospitals operate could contribute to geographic variation in spending (Congressional Budget Office 2006). In

---

9. The lower estimate is from work by Hadley and others (2006) and is probably an underestimate because of right-hand-side variables included in the regression that are correlated with the area measures. The higher estimate is for patients suffering heart attacks, reported by Fisher and others (2003a).

10. Cutler and Sheiner (1999) noted a roughly 70 percent correlation between the amount of Medicare spending in an area in 1982 and in that same area 15 years later, in 1997.

11. One explanation for the lack of an association is that HMOs tend to operate in traditionally high-cost areas and to use the savings they achieve to provide additional benefits rather than to reduce costs.

theory, for-profit providers need not be more costly, but some analysts believe that they are more aggressive than are nonprofit providers in maximizing the provision of profitable care (care that generates reimbursement that exceeds costs), which could increase total health care costs if it involves treating patients in a more intensive or costly manner than otherwise. One study reported that higher Medicare spending was associated with for-profit hospitals at a point in time, over time, and in hospitals that converted from nonprofit to for-profit status compared with those that changed in the opposite direction (Wennberg and others 2004). However, the authors were not able to control for all the factors associated with for-profit hospitals that could have led to their higher spending. Another study (Silverman and Skinner 2004) reported evidence that higher Medicare costs at for-profit hospitals were linked to their stronger tendency to categorize patients in higher-reimbursement groups as defined by Medicare (a practice known as upcoding). Areas served by for-profit hospitals also appear to have both higher Medicare spending and faster growth in Medicare spending than do areas served by nonprofit hospitals (Silverman and others 1999).

Supply of Medical Resources. The supply of medical resources, including personnel and physical facilities, varies widely across the United States. Some areas, for example, have roughly three times as many hospital beds or practicing physicians per person that other areas do. The next subsections discuss empirical findings on the effects of supply on variation in health care spending, some possible origins of the supply variation, and some mechanisms by which supply might affect health care spending.

Empirical Association Between Supply and Health Care Spending. A high number of hospital beds per capita and a high ratio of specialists to primary care physicians in an area are two measures of medical resources that several studies have reported are positively associated with higher spending. When researchers on the Dartmouth Atlas Project measured the adjusted rate of hospitalization across HRRs (controlling for age, sex, race, and illness) and compared that rate with the per capita supply of hospital beds, they found the adjusted hospitalization rate to be positively and strongly associated with the supply of hospital beds (Center for the Evaluative Clinical Sciences 1999, 79). Another study showed that (after controlling for demographics and health status) the probability of hospitalization among Medicare beneficiaries was strongly and positively associated with the supply of

hospital beds per capita (Fisher and others 2000). That relationship between supply and admissions was found to be stronger for nonsurgical admissions than for surgical admissions; nonsurgical admissions, they noted, are more influenced by the discretion of admitting physicians.

The percentage of physicians in an area who are specialists also is strongly associated with higher Medicare spending. The effect of a larger concentration of specialists on quality, however, is more ambiguous. One report, while acknowledging that medical specialists can provide better care than nonspecialists for the conditions specific to their specialty, noted that specialists can impose other "costs," such as coordination costs, which tend to hamper overall quality of care (Baicker and Chandra 2004b). According to that study, in areas that had one additional medical specialist and one fewer general practitioner per 100,000 people, providers tended to score lower on patient satisfaction surveys; the cost of care was $120 more per year, on average, per Medicare beneficiary; and aggregate mortality rates in those areas were no lower than elsewhere. Sepulveda, Bodenheimer, and Grundy (2007) review several studies that report similar findings.

Work done through the Dartmouth Atlas Project also indicates that the number of visits to doctors, and the number of doctors involved a single patient's care, vary geographically for otherwise similar patients. One technique the researchers used to assess that type of variation while controlling for the health status of patients in different areas was to derive an "end-of-life index," based on the amount of spending, adjusted for age, sex, and race, of a sample of Medicare patients during the final six months of life. That technique presumes that, because all of the patients died, their health status should have been similar during the six months preceding death. Fisher and others (2003a) found that regions in the highest quintile of Medicare spending, according to the end-of-life index, had 65 percent more medical specialists per capita but 26 percent fewer general and family practitioners. Other work by the same researchers examined care provided in academic medical centers (which generally are viewed as using state-of-the-art medicine and best practices overall) and showed that, in the final two years of life, patients in the highest-spending center received more than twice as many hours of physician visits, but nearly all of that difference was attributable to greater involvement of specialist physicians (3.6 to 1) rather than primary care physicians (1.8 to 1) (Fisher 2007).

Figure 8.

Spending per Beneficiary on Medicare Home Health Care Services in Four States

(Dollars)

Source: Congressional Budget Office based on data from the Centers for Medicare and Medicaid Services.

Possible Explanations for Variation in the Supply of Medical Resources. If geographic variation in the supply of medical resources arises from differences in demand for health care (for example, because of differences in illness severity, income, or treatment preferences), then any effects of supply on the use of services ultimately are driven by demand. In that case, the variation in spending that is related to supply differences is less of a policy concern. If supply variation is related to factors that are separate from the demand for health care (such as preferences of providers), then supply-driven spending variation could have significant policy implications.

In economic studies, the measures of demand that might be used to explain the distribution of medical resources are necessarily imperfect. It is difficult, therefore, to ascertain the extent to which unobserved demand factors might ultimately drive variation in the supply of medical resources. It is not clear, for example, what causes physicians or facilities to be concentrated in urban areas where the cost of living is high. Despite the limitations and difficulties in the analysis, however, there is evidence that at least some of the variation in supply is attributable to factors that are unrelated to demand.

The supply of hospital beds in the United States ranges from 1.7 to 4.8 per 1,000 population (using HRRs as the geographic unit). One study (Clayton and others 2005) reported that, at most, half of that variation was attributable to health-related variables and that much of the remainder appeared to be a consequence of population density many years earlier and of regulatory restrictions (such as Certificate of Need programs).[12]

Further evidence that supply variation is not driven entirely by demand comes from experience in the 1990s with the Medicare home health care benefit. In 1988, a lawsuit led to a loosening of criteria for eligibility and coverage, which led in turn to rapidly increased spending. Home health spending eventually was reined in during the late 1990s with new reimbursement policies required by the Balanced Budget Act of 1997. The boom-and-bust phenomenon in home health care was a nationwide phenomenon, but it was particularly pronounced in some states, including Texas and Louisiana (see Figure 8). The pattern has been attributed partly to increases in demand that resulted from a decline in the length of time, on average, that people spent in the hospital, but it appears primarily to have been the result of an influx of home health providers that has been ascribed to "low requirements to qualify as a provider, limited scrutiny of the appropriateness of claims, and outright fraud" (Fishman, Penrod, and Vladeck 2003, 111).

There is evidence that some of the uneven geographic distribution of physicians in the United States is unrelated to demand for physicians' services (see Figure 9). The number of physicians per capita is higher in areas with high population density (Rosenthal, Zaslavsky, and Newhouse 2005). Although some of the distribution can be explained by higher demand in urban areas, some is attributable to the preferences of physicians. For example, a specialist physician might seek a practice location with a critical mass of patients nearby and a critical mass of referring physicians, even if demand for that doctor's services is not higher in higher-density areas.

Early literature described the uneven distribution of physicians as evidence of market failure (in that the geographic distribution of physicians did not appear to be responsive to demand), but later work noted that

---

12. Certificate of Need programs are state-based regulatory vehicles used to permit or block new medical facilities; such programs were more active in earlier decades than they are today.

Figure 9.

Physicians per 100,000 Residents, 2004

State	
Idaho	
Mississippi	
Wyoming	
Nevada	
Arkansas	
Oklahoma	
Utah	
Alabama	
Alaska	
South Dakota	
Iowa	
Georgia	
Texas	
Indiana	
Montana	
Arizona	
South Carolina	
Kentucky	
Kansas	
New Mexico	
Nebraska	
North Dakota	
North Carolina	
West Virginia	
Florida	
California	
Louisiana	
Tennessee	
Wisconsin	
Virginia	
Washington	
Missouri	
New Hampshire	
Colorado	
Oregon	
Delaware	
Minnesota	
Illinois	
Michigan	
Ohio	
Hawaii	
Maine	
Pennsylvania	
New Jersey	
Rhode Island	
Vermont	
Connecticut	
Maryland	
New York	
Massachusetts	

Source: Congressional Budget Office based on data from the American Medical Association and the Bureau of the Census.

physicians were simply maximizing their own utility and perhaps trading off higher incomes for other amenities related to their preferred locations (Newhouse and others 1982). Whatever the case, the uneven distribution of physicians has led to federal incentives and programs, which persist today, to attempt to offset the imbalance by providing financial inducements for physicians to work in rural areas.[13]

How Supply Could Affect the Use of Health Care Services

There are several mechanisms through which the supply of health care facilities and personnel (holding population health, prices, and other factors constant) could affect the use of health care services. Some mechanisms could apply to any good or service; others are specific to health care.

Nonmonetary Costs. The nonmonetary costs of using health care services include, for example, time spent traveling to the site of care and time spent in the waiting room. An increase in the supply of health care services will tend to reduce those costs, because of shorter distances traveled and shorter waits for appointments.[14] Escarce (1992) examined the positive relationship between the supply of surgeons per capita and Medicare enrollees' use of surgical services. He reported that a greater supply of surgeons was associated with a higher number of initial consultations with surgeons, which he attributed to beneficiaries' facing lower nonmonetary costs. An international review of physician supply and waiting times for elective surgery reported a significant association between the number of physicians per capita in a country and waiting times (Siciliani and Hurst 2003).[15] Such a link between supply and nonmonetary costs applies to the consumption of any service and is not specific to health care.

Competition Among Providers. An increase in the supply of health care providers, holding all else constant, would

---

13. The Health Resources and Services Administration of the Department of Health and Human Services has programs to help balance the geographic distribution of physicians (http://bhpr.hrsa.gov/shortage/).

14. For discussions of the value of time in the consumption of health care and its effects on the use of services, see Acton (1975) and McLafferty (1988).

15. That study reported that a marginal increase of 1 physician per 10,000 population was associated with about a one-week decrease in the waiting time for elective surgery.

be expected to increase competition, which in turn might lead providers to receive lower fees.[16] The Government Accountability Office (2005) demonstrated a strong inverse relationship between the degree of competition among medical providers and what those providers are paid by private insurers. If increased supply lowers prices, that could in turn lead to reductions in the out-of-pocket payments that individuals face, which could increase the volume of services provided. In addition, lower prices could reduce insurance premiums and thereby increase insurance coverage and the volume of services provided.

Flexible Norms. In many clinical situations there are no hard-and-fast guidelines for appropriate care: "[M]edical texts and journals, for example, are silent on the incremental value of three-month versus six-month intervals between physician visits for patients with such conditions as diabetes or hypertension" (Wennberg and others 2002, 101). Where clear guidelines are lacking, clinicians might be more likely to adjust the recommended course of care to the medical resources available.

That adjustment might be made consciously in response to observed supply constraints, or it might evolve over time through gradual adjustments in local standards of care. Aaron, Schwartz, and Cox (2005), for example, compared kidney transplants in the United Kingdom and the United States. They showed that the attitudes of physicians concerning the medical appropriateness of performing such transplants on older patients reflected differences from one country to another and over time in the availability of financing for the procedure.

The flexibility of medical practice style facilitates the relatively uneven geographic distribution of physicians. One study has shown that where there are more cardiologists per capita, for example, Medicare patients visit cardiologists more frequently, suggesting that physicians' criteria for recommending a visit are adjusted on the basis of the number of cardiologists available (Center for the Evaluative Clinical Sciences 2007).[17] Teachers, by contrast, have little or no ability to influence demand for their services and generally must respond to the changing needs of the school-age population either by relocating of by entering

exiting the profession. It is not a surprise that the supply of teachers by state varies considerably less than does the supply of physicians. The state-level COV in teacher supply per capita is 0.13; for physicians the COV is 0.23.

Over short periods, the availability of hospital beds has been shown to affect the clinical threshold for admissions and discharges (Strauss and others 1986). In that study of the intensive care unit of a single hospital, the researchers found that when more beds were available, the patients admitted were healthier and stayed longer, on average, than when fewer beds were available. The physicians who worked in the unit appeared to reframe their clinical decisionmaking depending on the availability of beds. The authors did not interpret the variability in admission and discharge decisions as reflecting substandard care or demand inducement. Instead, they hypothesized that physicians were acting in the interests of a pool of patients—including those currently in intensive care and those who might be admitted—and to reduce their own workload and justify current staffing.

This short-run link between the availability and the use of services could have longer-run analogies across broad geographic areas. Boston, Massachusetts, and New Haven, Connecticut, have been used as examples of metropolitan areas that are generally similar but that differ both in availability and in use of hospital services. Wennberg and others (1987) showed that, despite the similarity of the two cities' populations, the hospitalization rate in Boston was 44 percent higher than in New Haven. The authors attributed that difference to Boston's relatively greater supply of hospital beds, which was, in turn, attributed to an unusually high concentration of universities and medical schools. In neither Boston nor New Haven were practice patterns alleged to be substandard—the authors noted that both areas are dominated by academic medical centers, which they took to imply that practice patterns in both cities met high standards of care. Furthermore, physicians in New Haven did not perceive a shortage of beds and did not perceive themselves as denying needed care. The fact that such different practice patterns can arise, even between areas dominated by leading academic medical centers, suggests that the

---

16. The assumption that prices for medical care are set through a conventional market-clearing mechanism has been questioned (Cromwell and Mitchell 1986, Feldstein 1970). In the Medicare fee-for-service program, prices are set administratively and are therefore unaffected by competition among providers.

17. Several researchers, including Gruber and Owings (1996), concur with the findings in *Supply-Sensitive Care* (Center for the Evaluative Clinical Sciences 2007) on the ability of physicians to influence demand for their services. The topic is somewhat controversial, however. For a review, see Phelps (2000).

standard of care is flexible and can shift in response to the availability of services.

More instances of the conformance of practice patterns and attitudes to resource availability can be found in recent work that compares physicians' attitudes and practices in areas with different amounts per capita of Medicare spending. Sirovich and others (2005) reported that 82 percent of physicians in high-spending areas said they would recommend an MRI (magnetic resonance imaging) scan for patients with back pain and mildly abnormal nerve function, compared with only 69 percent of physicians in low-spending areas. More recently, Sirovich and others (in press) reported that 49 percent of physicians in high-spending regions would recommend a physician visit within three months for a patient with isolated high blood pressure, compared with 22 percent of physicians in low-spending regions. Other researchers have reported that, in areas with high rates of cesarean section, physicians generally use more lenient criteria for judging when the procedure is indicated (Baicker, Buckles, and Chandra 2006). That is, regional variations in use of the procedure are driven not by differences in patients but by differences in the norms applied by physicians.

The Dartmouth Atlas researchers use the following scenario to generally describe how a greater supply of specialists and hospital beds may interact with financial incentives to increase spending: When confronted with patients for whom hospitalization may or may not be indicated, physicians are more likely to recommend hospitalization if beds are available. Once in the hospital, patients are easier to manage, from the physician's point of view, because "there are no late-night calls. But hospitalization lowers the threshold for further intervention: it is now easier to order tests, perform minor discretionary surgeries, or consult with other specialists, who in turn order their own tests and treatments. … And in the background, the fee-for-service system rewards everyone for doing more" (Mahar 2007, 4).

Although the description is anecdotal, the mechanism is consistent with the findings of other researchers. For example, one analysis of care received by Medicare beneficiaries with the same chronic conditions in high- versus low-spending regions of the country (as classified by the end-of-life index) reported rates of surgery and other major procedures that differed little across regions (Fisher and others 2003a). Yet rates of hospitalization, inpatient

consultations with specialists, and use of many diagnostic tests varied more than twofold, as did the number of days spent in intensive care in the final six months of life—all circumstances over which physicians have considerable influence.

Evidence of how fee-for-service reimbursement could amplify supply differences into spending differences comes from studies of physicians' behavior under different payment scenarios. One study placed physicians randomly into different reimbursement schemes and showed, for example, that fee-for-service reimbursement led physicians to schedule more patient visits than did salary-based reimbursement (Hickson, Altmeier, and Perrin 1987). There is little variation across the United States in how medical specialists are paid (fee-for-service is most common), so that in itself would not explain a great deal of the regional variation in spending. Yet if specialists were salaried, for example, their use of additional available resources (such as diagnostic machines) would be disconnected from their earnings. Under the fee-for-service model, in contrast, there could be incentives for specialists to purchase new equipment or to use available supply (which differs by geographic region) to obtain additional reimbursements—especially where such use might result in higher margins (reimbursements minus marginal costs) (Pham and others 2007).

Another type of care involving considerable discretion is post–acute care; that is, care after hospitalization. Research has established a clear link between the availability of post–acute care facilities and their use. For example, Buntin and others (2004) showed that when patients are discharged from a hospital, they are significantly more likely to be sent to post–acute care facilities if the hospital is near to or owns such a facility (many hospitals own skilled nursing facilities, for example). Further evidence of variability in the use of post–acute care comes from CBO's analysis of Medicare spending. In 2005, there was considerably greater variation across states in Medicare spending per beneficiary on home health care (COV, 0.43), hospital-based long-term care (0.40), and hospice care (0.40) than in outpatient care (0.11) or short-stay hospital care (0.15).

Spending and Quality of Care

Arriving at an understanding of the relationship between health care spending and the quality of care is a critical part of interpreting geographic variation. If high

spending appears to produce high quality of care and good health outcomes, that high spending would not necessarily seem inefficient. (Even in that case, high spending in some areas could still be cause for concern if it is inequitably financed or if it does not produce an improvement in quality that is commensurate with the resources used.) If high spending does not appear to produce high-quality care or improved outcomes, spending should be able to be reduced without compromising medical outcomes.

The evidence on the relationship between spending and quality or outcomes of care is not straightforward to interpret, but the following conclusions can be drawn:

■ Areas with higher-than-expected Medicare spending per beneficiary tend to score no better and, in some cases, score worse than other areas do on process-based measures of quality and on some measures of health outcomes.

■ Patterns of treatment in high-spending areas tend to be more intensive than in low-spending areas. That is, in high-spending areas a broader array of patients will receive costly treatments. Those treatment patterns appear to improve health outcomes for some types of patients, but worsen outcomes for others.

■ Decreasing spending in high-spending areas would reduce geographic variation in health care spending and might or might not harm the quality of care in those areas, depending on how the reduction took place.

Evidence on the Relationship Between Spending and Quality

The relationship between spending and quality is better understood for Medicare spending than for overall or non-Medicare health care spending, and it could be different for Medicare and the rest of the health care sector.[18] Several studies have examined the relationship between average spending and the quality of care provided to the Medicare population in different areas. The quality measures used are limited: Studies generally focus

on relatively easy-to-measure standards of "good medical practice," some of which are noted below.

The evidence does not indicate that higher Medicare spending is associated with better care for Medicare beneficiaries. In fact, it suggests the opposite: After adjusting for other factors, areas with higher Medicare spending tend to score substantially worse on a composite indicator of the quality of care provided to Medicare beneficiaries (see Figure 10).[19] That finding is echoed in the work of the Dartmouth Atlas Project: Areas with higher end-of-life expenditures in Medicare tend to perform worse in several dimensions of quality—particularly those that involve low-cost interventions (Fisher and others 2003a).

Even stronger evidence of the lack of an association between spending and quality of care comes from a state-level study of the way patient care changed as spending changed over time (Baicker and Chandra 2004a). The researchers found that if spending per Medicare beneficiary increased by $1,000 in a state, there was an associated *decrease* in most measures of "good" medical practice, including, for example, the share of heart attack patients who were given aspirin (a 3.6 percentage point decrease) or offered advice about smoking cessation (6.8 percentage points) at discharge, the share of pneumonia patients who received antibiotics within 8 hours of arrival at the hospital (2.0 percentage points), and the share of diabetes patients whose blood sugar concentrations were evaluated (3.2 percentage points).

Other studies have focused on other dimensions of quality, including patient satisfaction, functional status, and mortality rates. Fisher and others (2003a) identified regions (in this case HRRs) as high or low spending on the basis of Medicare expenditures at the end of life and showed that high-spending areas had aggregate mortality rates and mortality rates from several chronic diseases that either were slightly higher than or the same as rates

---

18. Martin and others (2007) noted that states with high Medicare spending do not necessarily have high total spending for health care. Preliminary analyses by CBO suggest a positive relationship at the state level between overall health care spending per capita and some measures of quality of care.

19. The Medicare Payment Advisory Commission (2003) calculated an adjusted state-level measure of Medicare expenditures per beneficiary in the fee-for-service sector that controlled for input prices, beneficiaries' health status, and other factors. States were then ranked according to adjusted spending per beneficiary, and that placement was compared with states' rankings on a composite measure of quality of care. The composite measure was calculated on the basis of the percentage of patients in specific clinical situations who received appropriate treatment as defined by MedPAC. (See also Baicker and Chandra [2004a] and Jencks and others [2003].)

Figure 10.

Relationship Between Quality of Care and Medicare Spending, by State, 2004

(Composite measure of quality of care, 100 = maximum)

State-Level Medicare Spending per Beneficiary (Thousands of dollars)

Source: Congressional Budget Office based on data from the Centers for Medicare and Medicaid Services and the Agency for Healthcare Research and Quality.

Note: The composite measure of quality reflects the provision of recommended care to patients hospitalized with acute myocardial infarction, heart failure, or pneumonia.

in low-spending regions.[20] Higher spending also was found not to be associated either with increased patient satisfaction or with improvements in patient function and health status (Fisher and others 2003b). And according to one study, physicians in high-spending areas noted more difficulty in coordinating care, providing for continuity of care, and communicating with other physicians (Sirovich and others 2006).

Two other recent studies that have analyzed the relationship among spending, medical treatment patterns, and health outcomes provide possible explanations for the lack of an association between higher spending and better health outcomes. Landrum and others (2008) showed

that patients who had colorectal cancer and lived in high-spending regions were more likely to receive chemotherapy than were similar patients in low-spending regions. And that treatment was given, it is critical to note, both to patients for whom it generally is recommended (those with stage III colon cancer) and to some for whom it is not and for whom it might in fact be harmful (those with stage I colon cancer, older patients, and those with multiple accompanying illnesses). The implication is that patterns of higher-cost, higher-intensity treatment could benefit some patients but harm others.

The second study examined treatment for heart attack (Chandra and Staiger 2007).[21] Among heart attack patients, high-cost surgical intervention will be more appropriate in some cases, and low-cost medical management will be more appropriate in others. The researchers report that patients for whom the high-cost surgical treatment was more appropriate fared better if they lived in areas that practiced surgical procedures on more patients, which tend to be high-spending areas. But patients for whom low-cost medical management was more appropriate fared worse in high-intensity, high-spending areas. Both studies suggest that the relationship between high-cost, intensive treatment and health outcomes is complex and depends on the patient population and the disease being treated.

Conceptual Models of Health Care Productivity

The evidence presented so far suggests that, at the state level, high spending on Medicare tends to be associated with care that is poorer in quality and does not necessarily produce improvement in aggregate outcomes. That observation implies that, in many states, Medicare funds are being spent inefficiently. It does not necessarily suggest, however, that high-cost areas could become low-cost areas easily and without detriment to health outcomes. It is difficult to judge how much could be saved or how much efficiency would improve if geographic variation were reduced. To do so requires an understanding of how health care is produced in low- and high-cost areas.

Flat-of-the-Curve Medicine. In one conceptual model, the "flat-of-the-curve" (FOTC) model, the health benefit produced (the output) is a simple function of the amount

---

20. Several groups of patients were analyzed for differences in outcomes: a representative sample of all Medicare patients and subgroups hospitalized for colorectal cancer, hip fracture, and acute myocardial infarction.

21. Another, preliminary, study (Doyle 2007) identified a positive association in higher-spending regions in Florida with survival among heart emergency patients.

that is spent on care (the input). The FOTC model underlies much of the literature on geographic variation, but it appears to be inconsistent with some of the evidence on the differences in practice patterns between high- and low-spending regions, and it could be an over-simplification.

The FOTC model assumes that the mix of effective, ineffective, and harmful care will change with the amount spent. With very low spending, effective care is provided; at higher levels, effective and ineffective care are provided; and at higher levels still, all three types of care are provided. The implication of this model is that, when health care spending is "too high"—that is, when it is beyond the point of producing maximum health benefits—the quality of health care actually can be improved by reducing spending. (Such a reduction moves spending back to the "flat of the curve," where benefits are, if not rising, at least not falling with increasing spending.) The fact that higher Medicare spending is associated with lower quality of care has been interpreted as implying just that.[22]

The evidence is mixed on whether inappropriate care is more likely to be delivered in high-cost areas. Several studies have compared the different mixes of inappropriate and appropriate procedures in high- and low-utilization regions (Baicker, Buckles, and Chandra 2006; Chassin and others 1987; Dyck and others 1977; Guadagnoli and others 2001; Leape and others 1990). Each study examined either a single type of procedure or a small set of procedures. Case histories for individual procedures were examined, and procedures were categorized on the basis of clinical appropriateness. For each geographic region, the share of procedures that was determined to be appropriate was measured, and areas with high rates of service use were compared on this measure with areas where use was lower.

Two of the studies (Chassin and others 1987, Leape and others 1990) used similar approaches and examined several procedures used for diagnosis and treatment of cardiovascular ailments or gastrointestinal conditions. The studies provided some evidence that the share of procedures deemed appropriate was higher in areas with lower rates of service use and the share of procedures

deemed inappropriate was higher in areas with higher rates of use. The authors point out, however, that the overall variation in the use of procedures was attributable mainly to variations in the provision of procedures that were medically appropriate.

Similarly, a study of patients with colorectal cancer gives evidence that high-spending areas provide more services to all types of patients, regardless of appropriateness (Landrum and others 2008). Those findings imply that some services are not provided in a strictly prioritized way and that the FOTC model may oversimplify the relationship between spending and outcomes.

A study of variation in delivery by cesarean section, however, reported evidence that is more consistent with FOTC medicine (Baicker, Buckles, and Chandra 2006). The authors contend that physicians generally rank patients according to the clinical appropriateness of cesarean section delivery for the patient and determine where the dividing line should be between cesarean and vaginal births. That dividing line was found to be strongly affected by nonmedical factors, such as the supply of physicians, exposure to malpractice liability, and local physicians' opinions. The implication of the model is that in areas where cesarean sections are more common, doctors perform the "additional" operations on patients for whom the procedure would be deemed unnecessary in the more conservative areas.

Multiple Production Functions. A more sophisticated conceptual model, one that fits better with some of the evidence on geographic variation, is the multiple-production-functions model (Skinner, Fisher, and Staiger 2006). In this model, different geographic regions adopt different production technologies. At the outset (Time 1 in Figure 11), all regions operate at point A on the initial curve, which represents the production of health outcomes as a function of increasing expenditure. Over time, treatment technology improves, shifting all regions to higher production curves, which represent better outcomes for every level of expenditure than in Time 1.

But some regions travel to a point on a new production function (point C) characterized by small and diminishing returns: Those regions adopt practices that cost more but achieve limited improvement in outcomes. The other regions move to a point on a different production function (point B) representing a much larger improvement in outcomes for the same level of expenditure, and they

---

22. For a discussion of the potential harm from excessive provision of health care services, see Fisher and Welch (1999). A simple discussion of flat-of-the-curve medicine is provided by Fuchs (2004).

Figure 11.

Conceptual Model of Health Care Delivery in Different Regions

Source: Adapted by the Congressional Budget Office with permission from Skinner, Fisher, and Staiger (2006).

do not travel as far along this curve toward diminishing returns. Those regions more readily adopt low-cost, effective treatments (such as those mentioned above for heart attack). Between Time 1 and Time 2, then, both types of regions have increased expenditures and improved outcomes, but the low-spending regions have increased expenditures less and improved outcomes more, resulting in the negative correlation in Time 2 between spending and quality or outcomes (points B and C).

Chandra and Staiger (2007) used Medicare data from 1996 to study the care of heart attack patients for whom two general courses of treatment were available: surgery (a high-intensity treatment) or medical management (a low-intensity treatment). For reasons that are not entirely clear, some regions of the country were found to become better at (to "specialize in") the high-intensity treatment; that is, a patient who receives the aggressive treatment will have a better outcome, and for a lower cost, in areas that perform the treatment more often. At the same time, those high-intensity areas are worse at performing the low-intensity treatment: Outcomes are poorer and costs are higher than in low-intensity areas. "Spillovers" reinforce this tendency; as more and more physicians and supporting institutions become skilled at the high-intensity treatment, medical practice becomes oriented

toward it, and the low-intensity treatment is used less frequently and produces worse outcomes.

This specialization might not result in any aggregate differences in outcomes for patients from one area to the next. Gains in high-intensity areas among patients given the high-intensity intervention could be offset by losses for patients in those same areas who are given the low-intensity treatment. At the same time, one additional consequence of this model tends to lead to higher average costs in the high-intensity areas, despite the fact that the high-intensity treatment is performed for a lower cost per case. Because of their greater efficiency with the high-intensity treatment (and lower efficiency with the low-intensity treatment), providers in high-intensity areas will have an incentive to apply the intensive treatment to more patients; that is, they set the bar lower for deciding who should receive the high-intensity treatment. That treatment is far more costly than the low-intensity treatment is, and thus total spending is higher.

In reality, it is possible that the multiple-production-functions model applies better to some fields of medical practice, and the FOTC model applies better to others. The implications of the FOTC model are clear: Cutting spending in high-cost areas does not produce worse

outcomes because those areas are performing care of no or negative value at the margin. Yet to the extent that the multiple-production-functions model applies to some portion of health care delivery, its implications are less clear. In that model, efficiency improvements are still possible, but they are more complicated to achieve. In particular, simply forcing a reduction in spending in the high-spending regions could worsen outcomes (that is, move backward along the middle curve in Figure 11) as long as the area remains on the same production function curve. The key to improving efficiency would be to create an upward shift in the curve on which the high-spending area lies, toward that of the low-spending area, thus causing a convergence in medical practice and an improvement in outcomes for a given amount of money spent.

Reducing Geographic Variation in Spending for Health Care

Some changes have occurred already in the Medicare program to reduce spending variation and improve efficiency. Other policies or interventions can be posited that might further improve the efficiency of health care spending while reducing its geographic variation. (Note that not all interventions that would reduce geographic variation in health care spending would necessarily improve the overall efficiency of medical practice. For example, reducing payments to high-spending areas while increasing payments to low-spending areas would reduce spending variation but could result in worse outcomes if quality of care declines in the high-spending areas more than it improves in places where spending is lower.) CBO is undertaking an expanded effort to analyze options for improving the efficiency of the health care system.

The options include the following:

■ Increase the "bundling" of services in payments to providers (such as those that have been implemented in the Medicare program for payments to hospitals, for example), which could help to curb current incentives to provide increasingly intensive services that produce only modest or no improvement in health.

■ Enhance incentives to provide care consistent with accepted guidelines for low-cost, highly effective care, thus helping to change patterns of medical practice in places that now are characterized by lower-quality, higher-cost care.

■ Generate more information about variations in practice patterns and the relative cost-effectiveness of different procedures for different populations. Such information could help reorient inefficient practice patterns toward greater efficiency, especially if greater oversight or changed financial incentives led to increased pressure to use this sort of information.

Prospective payment or increased bundling of payments could reduce both the financial incentives and the opportunities for providers to distort patterns of care so as to receive greater reimbursements. Providers could be paid a fixed amount for all treatments for a certain patient with a certain condition, or they could be paid for episodes of care that include inpatient care, physician services, and post–acute care, thus removing the prospect of additional payments for additional procedures (Medicare Payment Advisory Commission 2006, Mutti and Lisk 2007). The trends in Medicare's payment rates are generally consistent with the hypothesis that the introduction of formula-based rate-setting systems reduces geographic variation. (A potential drawback of bundling, however, is the incentive to withhold necessary care.)

Identifying and targeting providers whose practices are unusually costly also could help to increase efficiency and reduce geographic variation in health care spending. Welch and Miller (1994) proposed an approach in which the practice patterns of medical staffs of individual hospitals would be measured, and payments would be reduced for hospitals with patterns of unusually high spending. Fisher and his colleagues (2007) set the unit of accountability at the individual hospital and the associated set of providers who are primarily users of that medical facility. Practice patterns vary greatly among the roughly 5,000 such units in the United States. Financial incentives or performance measures targeted that way could affect the decisions that lead to divergent practices, such as unusually aggressive investments in facilities, certain types of personnel, or other resources.

One variation on this idea targets the efficiency of Medicare spending generally but also could serve to reduce geographic variation in spending. That approach is a current initiative in Medicare that provides bonus payments to hospitals that participate in programs to meet certain clinical goals for provision of highly efficient care (such as giving aspirin to and prescribing beta blockers for heart attack patients) (Kahn and others 2006). Such incentives

could reduce variation and improve outcomes by encouraging more national uniformity in patterns of practice.

More intensive oversight, which helped end the rapid increase in spending for Medicare's home health care program, would probably serve to increase the program's efficiency and reduce geographic variation in spending. Oversight of the Medicare program is conducted by two Congressional agencies (MedPAC and GAO) and by CMS. As a first step, those entities could focus on identifying service categories or procedures with high rates of growth over time, or high degrees of geographic variation, or both. If distorted financial incentives or regulatory lapses were found to promote those services, MedPAC could propose legislative remedies and CMS could implement appropriate regulatory changes.

Finally, significant opportunities exist outside of Medicare to improve efficiency while reducing geographic variation in spending. Experts have noted a link between the degree of uncertainty regarding the clinical effectiveness of a procedure and the degree of geographic variation in the use of that procedure. Expansions in knowledge regarding clinical effectiveness appear, therefore, to have some potential to reduce geographic variation in health care spending. For expanded clinical knowledge to be translated into actual changes in practice patterns, however, several steps are required: Newly developed guidelines must be widely disseminated, current practice patterns must be measured and compared against best-practices benchmarks, and financial incentives must be provided to encourage adoption of the new practices.

Even without any expansions in clinical knowledge, the measurement and dissemination of information on variations in practice patterns could reduce geographic variation. Measuring and reporting patterns at small units of analysis (the individual hospital, nursing home, or physician) can be particularly effective in identifying outliers and spurring change. The researchers from the Dartmouth Atlas Project have analyzed variation in spending in individual hospitals for the care of chronically ill Medicare beneficiaries (Wennberg and others 2005). That analysis was a notable departure from previous work, which measured variation only on larger geographic scales. The newer work was applauded by the chairman of a large insurer that is using similar analytical methods to improve efficiency and reduce practice variation (Schaeffer and McMurtry 2005).

References

Aaron, Henry, William Schwartz, and Melissa Cox. 2005. *Can We Say No? The Challenge of Rationing Health Care* (Washington, D.C.: Brookings Institution Press).

Acton, Jan Paul. 1975. "Nonmonetary Factors in the Demand for Medical Services: Some Empirical Evidence," *Journal of Political Economy,* vol. 83, no. 3 (June), pp. 595–614.

Aspinall, S.L., and others. 2005. "Facility-Level Variation in Antibiotic Prescriptions for Veterans With Upper Respiratory Infections," *Clinical Therapy,* vol. 27, no. 2 (February), pp. 258–262.

Baicker, Katherine, and Amitabh Chandra. 2004a. "Medicare Spending, the Physician Workforce, and Beneficiaries' Quality of Care," *Health Affairs,* Web exclusive (April 7), pp. W4-184–W4-197, http://content.healthaffairs.org/cgi/content/full/hlthaff.w4.184v1/DC1.

Baicker, Katherine, and Amitabh Chandra. 2004b. "The Productivity of Physician Specialization: Evidence from the Medicare Program," *American Economic Review*, vol. 94, no. 2, pp. 357–361.

Baicker, Katherine, Kasey Buckles, and Amitabh Chandra. 2006. "Geographic Variation in the Appropriate Use of Cesarean Delivery," *Health Affairs,* Web exclusive (August 8), pp. W355–W367, http://content.healthaffairs.org/cgi/content/full/25/5/w355.

Barnato, Amber E., and others. 2007. "Are Regional Variations in End-of-Life Care Intensity Explained by Patient Preferences? A Study of the US Medicare Population," *Medical Care*, vol. 45, no. 5, pp. 386–393.

Buntin, Melinda B., and others. 2004. *How Much Is Post–Acute Care Use Affected by Its Availability?* NBER Working Paper 10424 (Cambridge, Mass.: National Bureau of Economic Research).

Center for the Evaluative Clinical Sciences. 1999. *The Quality of Medical Care in the United States: A Report on the Medicare Program,* www.dartmouthatlas.org/atlases/99Atlas.pdf.

Center for the Evaluative Clinical Sciences. 2007. *Supply-Sensitive Care,* Dartmouth Atlas Project Brief, www.dartmouthatlas.org/topics/supply_sensitive.pdf.

Centers for Medicare and Medicaid Services. 2007. *Health Expenditures by State of Residence, 1991–2004,* www.cms.hhs.gov/NationalHealthExpendData/downloads/resident-state-estimates.zip.

Chandra, Amitabh, and Douglas Staiger. 2007. "Productivity Spillovers in Health Care: Evidence from the Treatment of Heart Attacks," *Journal of Political Economy*, vol. 115, no. 1, pp. 103–140.

Chassin, Mark, and others. 1987. "Does Inappropriate Use Explain Geographic Variations in the Use of Health Care Services? A Study of Three Procedures," *JAMA: The Journal of the American Medical Association,* vol. 258, no. 18 (November 13), pp. 2533–2537.

Clayton, Lindsay, and others. 2005. *Why Is There Regional Variation in per Capita Hospital Bed Capacity?* (New Haven: Yale University).

Congressional Budget Office. 2006. *Nonprofit Hospitals and the Provision of Community Benefits.*

Cromwell, Jerry, and Janet B. Mitchell. 1986. "Physician-Induced Demand for Surgery," *Journal of Health Economics,* vol. 5, no. 4 (1986), pp. 293–313.

Culyer, A.J., and others. 1981. "Alternative Systems of Health Care Provision: An Essay on Motes and Beams," in Mancur Olson, ed., *A New Approach to the Economics of Health Care* (Washington, D.C.: American Enterprise Institute for Public Policy Research), pp. 131–150.

Cutler, David M., and Louise Sheiner. 1999. "The Geography of Medicare," *American Economic Review*: *Papers and Proceedings,* vol. 89, no. 2, pp. 228–233.

Department of Veterans Affairs. 2001. *Veterans Equitable Resource Allocation: Equity of Funding and Access to Care Across Networks.*

Department of Veterans Affairs. 2007. *Veterans Equitable Resource Allocation: Equity of Funding and Access to Care Across Networks.*

Doyle, Joseph. 2007. *Returns to Local-Area Health Care Spending: Using Health Shocks to Patients Far From Home,* NBER Working Paper 13301 (Cambridge, Mass.: National Bureau of Economic Research).

Dyck, Frank J., and others. 1977. "Effect of Surveillance on the Number of Hysterectomies in the Province of Saskatchewan," *New England Journal of Medicine,* vol. 296, pp. 1326–1328.

Escarce, José J. 1992. "Explaining the Association Between Surgeon Supply and Utilization," *Inquiry,* vol. 29 (Winter), pp. 403–415.

European Observatory on Health Systems and Policies. 1999. *Health Systems in Transition: United Kingdom,* www.euro.who.int/document/e68283.pdf, pp. 85–86.

Feldstein, Martin S. 1970. "The Rising Price of Physicians' Services," *Review of Economics and Statistics,* vol. 52, no. 2 (May), pp. 121–133.

Fisher, Elliott. 2007. "Can We Use the Health Care Workforce More Efficiently?" (presentation, "Health Care Workforce Issues of the 21st Century: The Fourteenth Princeton Conference," May 23–24, Princeton, N.J.), http://council.brandeis.edu/pubs/Prince14/Elliott%20Fisher%20final.pdf.

Fisher, Elliott S., and H. Gilbert Welch. 1999. "Avoiding the Unintended Consequences of Growth in Medical Care: How Might More Be Worse?" *JAMA: The Journal of the American Medical Association,* vol. 281, no. 5 (February 3), pp. 446–453.

Fisher, Elliott S., and others. 2000. "Associations Among Hospital Capacity, Utilization, and Mortality of U.S. Medicare Beneficiaries, Controlling for Sociodemographic Factors," *Health Services Research,* vol. 34, no. 6, pp. 1351–1362.

Fisher, Elliott S., and others. 2003a. "The Implications of Regional Variations in Medicare Spending. Part 1: The Content, Quality, and Accessibility of Care," *Annals of Internal Medicine,* vol. 138, no. 4 (February 18), pp. 273–287.

Fisher, Elliott S., and others. 2003b. "The Implications of Regional Variations in Medicare Spending. Part 2: Health Outcomes and Satisfaction with Care," *Annals of Internal Medicine,* vol. 138, no. 4 (February 18), pp. 288–298.

Fisher, Elliott S., and others. 2004. "Variations in the Longitudinal Efficiency of Academic Medical Centers," *Health Affairs,* Web exclusive, (October 7), http://content.healthaffairs.org/cgi/content/full/hlthaff.var.19/DC3.

Fisher, Elliott, and others. 2007. "Creating Accountable Care Organizations: The Extended Hospital Medical Staff," *Health Affairs,* Web exclusive, vol. 26, no. 1, pp. W44–W57, http://content.healthaffairs.org/cgi/content/full/26/1/w44.

Fishman, Eliot Z., Joan D. Penrod, and Bruce C. Vladeck. 2003. "Medicare Home Health Utilization in Context," *Health Services Research*, vol. 38, no. 1, Part I, pp. 107–112.

Fortney, John C., and others. 1996. "Variation Among VA Hospitals in Length of Stay for Treatment of Depression," *Psychiatric Services,* vol. 47, no. 6 (June), pp. 608–613.

Fuchs, Victor. 2004. "More Variation in Use of Care, More Flat-of-the-Curve Medicine," *Health Affairs,* Web exclusives (October 7), pp. 104–105, http://content.healthaffairs.org/cgi/content/full/hlthaff.var.104/DC2.

Fuchs, Victor, Mark McClellan, and Jonathan Skinner. 2001. *Area Differences in Utilization of Medical Care and Mortality Among U.S. Elderly,* NBER Working Paper 8628 (Cambridge, Mass.: National Bureau of Economic Research).

Gage, Barbara, Marilyn Moon, and Sang Chi. 1999. "State-Level Variation in Medicare Spending," *Health Care Financing Review,* vol. 21, no. 2, pp. 85–98.

General Accounting Office. 2002. *VA Health Care: Allocation Changes Would Better Align Resources With Workload,* GAO-02-338, www.gao.gov/new.items/d02338.pdf, Table 5.

Gold, Marcia. 2004. *Geographic Variation in Medicare per Capita Spending: Should Policymakers Be Concerned?* Synthesis Project Report 6 (Princeton, N.J.: Robert Wood Johnson Foundation).

Government Accountability Office. 2005. *Federal Employees Health Benefits Program: Competition and Other Factors Linked to Wide Variation in Health Care Prices,* GAO-05-856, www.gao.gov/new.items/d05856.pdf.

Gruber, Jonathan, and Maria Owings. 1996. "Physician Financial Incentives and Caesarian Section Delivery," *RAND Journal of Economics*, vol. 27, no. 1, pp. 99–123.

Guadagnoli, Edward, and others. 2001. "Impact of Underuse, Overuse and Discretionary Use on Geographic Variation in the Use of Coronary Angiography After Acute Myocardial Infarction," *Medical Care,* vol. 39, pp. 446–458.

Hadley, Jack, and others. 2006. *Variations in Medical Care Spending per Medicare Beneficiary: The First Stage of an Instrumental Variable Analysis, Final Report,* submitted to Changes in Health Care Financing and Organization Program, Project ID 53269.

Hickson, Gerald, William Altmeier, and James Perrin. 1987. "Physician Reimbursement by Salary or Fee-for-Service: Effect on Physician Practice Behavior in a Randomized Prospective Study," *Pediatrics*, vol. 80, no. 3, pp. 344–350.

Iglehart, John K. 1996. "Reform of the Veterans Affairs Health Care System," *New England Journal of Medicine,* vol. 335, no. 18 (October 31), pp. 1407–1412.

Jencks, Stephen F., and others. 2003. "Change in the Quality of Care Delivered to Medicare Beneficiaries, 1998–1999 to 2000–2001," *JAMA: The Journal of the American Medical Association,* vol. 289, no. 3 (January 15), pp. 305–312.

Kahn, Charles III, and others. 2006. "Snapshot of Hospital Quality Reporting and Pay-for-Performance Under Medicare," *Health Affairs,* vol. 25, no. 1, pp. 148–162, http://content.healthaffairs.org/cgi/content/full/25/1/148.

Landrum, Mary Beth, and others. 2008. "Is Spending More Always Wasteful: The Appropriateness of Care and Outcomes Among Colorectal Cancer Patients," *Health Affairs,* vol. 27, no. 1, pp. 159–168, http://content.healthaffairs.org/cgi/content/full/27/1/159.

Leape, Lucian L., and others. 1990. "Does Inappropriate Use Explain Small-Area Variation in the Use of Health Care Services?" *JAMA: The Journal of the American Medical Association,* vol. 263, no. 5 (February 2), pp. 669–672.

Mahar, Maggie. 2007. "The State of the Nation's Health," *Dartmouth Medicine* (Spring), http://dartmed.dartmouth.edu/spring07/html/atlas.php.

Manning, Willard G., and others. 1987. "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment," *American Economic Review*, vol. 77, no. 3, pp. 251–277.

Martin, Anne B., and others. 2007. "Health Spending by State of Residence, 1991–2004," *Health Affairs,* Web exclusives (September 18), vol. 26, no. 6, pp. W651–W663, http://content.healthaffairs.org/cgi/content/full/26/6/w651.

McLafferty, S. 1988. "Predicting the Effect of Hospital Closure on Hospital Utilization Patterns," *Social Science and Medicine,* vol. 27, no. 3, pp. 255–262.

Medicare Payment Advisory Commission. 2003. *Report to the Congress: Variation and Innovation in Medicare,* www.medpac.gov/documents/June03_Entire _Report.pdf, Chapter 1.

Medicare Payment Advisory Commission. 2006. "Using Episode Groupers to Assess Physician Resource Use," in *Increasing the Value of Medicare,* www.medpac.gov/documents/Jun06_EntireReport.pdf, Chapter 1.

Mutti, Anne, and Craig Lisk. 2007. "Moving Toward Bundled Payments Around Hospitalizations" (Washington, D.C.: Medicare Payment Advisory Commission), www.medpac.gov/transcripts/1107_bundling _hosp_AM_pres.pdf.

Newhouse, Joseph E. 1993. *Free for All? Lessons From the RAND Health Insurance Experiment* (Cambridge, Mass.: Harvard University Press).

Newhouse, Joseph, and others. 1982. "Does the Geographical Distribution of Physicians Reflect Market Failure?" *Bell Journal of Economics*, vol. 13, no. 2, pp. 493–505.

Pham, Hoangmai, and others. 2007. "Redesigning Care Delivery in Response to a High Performance Network: The Virginia Mason Medical Center," *Health Affairs,* Web exclusives, vol. 26, no. 4 (July 10), pp. W532–W544, http://content.healthaffairs.org/cgi/content/full/26/4/w532.

Phelps, Charles E. 2000. "Information Diffusion and Best Practice Adoption," in *Handbook of Health Economics,* A.J. Culyer and J.P. Newhouse, eds., vol. 1, no. 1 (Amsterdam: Elsevier).

Pope, Gregory C., and others. 1989. "Cost of Practice and Geographic Variation in Medicare Fees," *Health Affairs,* vol. 8, no. 3 (August), pp. 117–128.

Rosenthal, Meredith, Alan Zaslavsky, and Joseph Newhouse. 2005. "The Geographic Distribution of Physicians Revisited," *Health Services Research*, vol. 40, no. 6, pp. 1931–1952.

Schaeffer, Leonard D., and Dana E. McMurtry. 2005. "Variation in Medical Care: Time for Action," *Health Affairs,* Web exclusives (November 16), pp. W5-552–W5-554, http://content.healthaffairs.org/cgi/content/full/hlthaff.w5.552/DC1.

Sepulveda, Martin-J., Thomas Bodenheimer, and Paul Grundy. 2007. "Primary Care: Can It Solve Employers' Health Care Dilemma?" *Health Affairs,* vol. 27, no. 1, pp. 151–158, http://content.healthaffairs.org/cgi/content/full/27/1/151.

Siciliani, Luigi, and Jeremy Hurst. 2003. *Explaining Waiting Times Variations for Elective Surgery Across OECD Countries,* OECD Health Working Papers 7 (Paris: Organisation for Economic Co-operation and Development).

Silverman, Elaine, and Jonathan Skinner. 2004. "Medicare Upcoding and Hospital Ownership," *Journal of Health Economics*, vol. 23, no. 2, pp. 369–389.

Silverman, Elaine M., and others. 1999. "The Association Between For-Profit Hospital Ownership and Increased Medicare Spending," *New England Journal of Medicine,* vol. 341, no. 6 (August 5), pp. 420–426.

Sirovich, Brenda E., and others. 2005. "Variation in the Tendency of Primary Care Physicians to Intervene," *Archives of Internal Medicine*, vol. 165, no. 19, pp. 2252–2256.

Sirovich, Brenda, and others. 2006. "Regional Variations in Health Care Intensity and Physician Perceptions of Quality of Care," *Annals of Internal Medicine*, vol. 144, pp. 641–649.

Sirovich, Brenda E., and others. In press. "Discretionary Decision-Making by Primary Care Physicians and the Cost of Health Care in the United States," *Health Affairs*.

Skinner, Jonathan, Elliott Fisher, and Douglas Staiger. 2006. "Is Technological Change in Medicine Always Worth It? The Case of Acute Myocardial Infarction," *Health Affairs,* Web exclusives, vol. 25, W34–W47, http://content.healthaffairs.org/cgi/content/full/25/2/w34.

Strauss, Michael J., and others. 1986. "Rationing of Intensive Care Unit Services: An Everyday Occurrence," *JAMA: The Journal of the American Medical Association,* vol. 255, pp. 1143–1146.

Subramanian, Usha, and others. 2002. "Geographic Variation in Health Care Utilization and Outcomes in Veterans With Acute Myocardial Infarction," *Journal of General Internal Medicine,* vol. 17, no. 8 (August), pp. 604–611.

Super, Norah. 2003. *The Geography of Medicare: Explaining Differences in Payment and Costs,* National Health Policy Forum Issue Brief 792 (Washington, D.C.: National Health Policy Forum).

Welch, W. Pete, and Mark E. Miller. 1994. "Proposals to Control High-Cost Hospital Medical Staffs," *Health Affairs,* vol. 13, no. 4 (Fall), pp. 42–57, http://content.healthaffairs.org/cgi/reprint/13/4/42.

Welch, W. Pete, and others. 1993. "Geographic Variation in Expenditures for Physicians' Services in the United States," *New England Journal of Medicine,* vol. 328, no. 9, pp. 621–627.

Wennberg, John E., and others. 1982. "Professional Uncertainty and the Problem of Supplier-Induced Demand," *Social Science & Medicine,* vol. 16, pp. 811–824.

Wennberg, John E., and others 1987. "Are Hospital Services Rationed in New Haven or Over-Utilised in Boston?" *Lancet* (May 23), pp. 1185–1188.

Wennberg, John E., and others. 2002. "Geography and the Debate Over Medicare Reform," *Health Affairs,* Web exclusives (February 13), pp. W96–W114, http://content.healthaffairs.org/cgi/content/full/hlthaff.w2.96v1/DC1.

Wennberg, John E., and others. 2004. "Use of Hospitals, Physician Visits, and Hospice Care During the Last Six Months of Life Among Cohorts Loyal to Highly Respected Hospitals in the United States," *British Medical Journal,* vol. 328 (March 13), p. 607.

Wennberg, John E., and others. 2005. "Evaluating the Efficiency of California Providers in Caring for Patients With Chronic Illnesses," *Health Affairs,* Web exclusives (November 16),W5-526–W5-543, http://content.healthaffairs.org/cgi/content/full/hlthaff.w5.526/DC1.

Wilt, Timothy J., and others. 1999. "An Evaluation of Radical Prostatectomy at Veterans Affairs Medical Centers: Time Trends and Geographic Variation in Utilization and Outcomes," *Medical Care,* vol. 37, no. 10 (October), pp. 1046–1056.