

ASPIRIN, ANGIOPLASTY, AND PROTON BEAM THERAPY: THE ECONOMICS OF SMARTER HEALTH CARE SPENDING

Katherine Baicker and Amitabh Chandra*
Harvard and NBER

September 9, 2011

Abstract: The growing share of the U.S. economy devoted to health care spending is cause for concern both because of the share that is publicly financed and because of the inefficiency with which those resources are used. Fueled in part by demographic transitions, unchecked increases in entitlement spending will necessitate some combination of tax increases, the elimination of other public spending, or public debt levels that far exceed those currently observed in Greece. Despite these spending levels, highly cost-effective treatments like aspirin and flu shots are underused, while angioplasty is used in both lifesaving and inappropriate cases and exorbitantly unproven procedures such as proton beam therapy are generously reimbursed by public programs. Driven by expenditures on such expensive care, public insurance spending is rising unsustainably at the same time that a fifth of the population goes uninsured. In an efficient system, more spending on health care would not be a cause for concern. Achieving this means improving the incentives and infrastructure for providers to deliver – and patients to consume – high-value care, as well as wrestling with the difficult question of whom to cover versus what to cover in public insurance programs. But the wide prevalence of inefficiencies also offers hope, for there is scope for fundamental reforms to improve the productivity of health care spending and the fiscal health of the U.S.

* Prepared for Jackson Hole Economic Policy Symposium.

Katherine Baicker is a Professor at the Harvard School of Public Health. Amitabh Chandra is a Professor at the Harvard Kennedy School. Both are Research Associates of the NBER. The authors thank Jonathan Skinner for invaluable collaborations and Jacob Robbins for expert research assistance. The work in this paper draws on Baicker and Chandra (2008), Baicker and Chandra (2010), Chandra, Jena and Skinner (2011) and Chandra and Skinner (2011).

I. INTRODUCTION

Health care spending in the United States now represents one sixth of GDP and is projected to comprise more than a quarter within 25 years (Congressional Budget Office 2011). Spending a large and increasing share of GDP on health care in and of itself is not necessarily troubling: as societies grow richer, it would seem natural that an increasing share of their resources should be devoted to “purchasing” health and longevity. After all, living an extra year in good health may be worth more on the margin than an additional home, a longer vacation, or a more sumptuous meal. Some studies suggest that, given the relative productivity of health spending, we might optimally spend as much as a third of GDP on health by mid-century (Fogel 2008; Hall and Jones 2007), and that the improvements in health outcomes driven by health care spending outweigh the costs of care (Murphy and Topel 2006). If the aggregate health benefits of care truly outweigh the total social costs, there is no a priori reason to be concerned about devoting a quarter of GDP – or more – to living longer and healthier lives.

Why, then, is there such distress over rising health expenditures? There are two reasons to be less than sanguine about health care spending in the U.S. – and indeed in most developed countries (see Table 1). First, public expenditures account for almost half of U.S. health care spending (Centers for Medicare and Medicaid Services 2010). Federal spending on public programs (including public insurance programs for the poor and elderly and newly enacted subsidies for the poor to purchase private insurance) is anticipated to grow from 5.5 percent of GDP now to almost 14 percent or more in 2060 (Congressional Budget Office 2011). The economic distortions associated with raising revenues to finance the public share come with a sizable macroeconomic cost (Baicker and Skinner 2011; Congressional Budget Office 2007). Unchecked, increases in entitlement spending would double the federal budget (as a share of

GDP), necessitating some combination of tax increases, the elimination of other public spending, or public debt levels that far exceed those currently observed in Greece (Chernew, Baicker, and Hsu 2010). Second, there is growing consensus that health care resources are not being spent efficiently (and may not even be the primary driver of improved outcomes): we are neither allocating resources efficiently between health and other uses, nor getting as much health as we could for every dollar spent – making it difficult to evaluate how much we “should” be spending on health care.¹ Such inefficiencies are clearly of first-order importance as health care spending encompasses a larger and larger share of total resources.

We begin by describing the landscape of spending in the U.S. in the context of the macroeconomy, how it has evolved over time, and how it compares to spending patterns and trends in other developed countries. The U.S. health care system may on first examination appear to be an intractable morass of regulations and acronyms, and it is tempting to view its failures as the sum of disparate inefficiencies. But such a viewpoint motivates piecemeal modifications that may in fact perpetuate the deeper inefficiencies. To help avoid this trap, we outline a simple framework through which we interpret the principle impediments to a well-functioning health care system. Such a system should exhibit productive efficiency, meaning that health care resources are put to the best use possible and produce as much health as they can, and allocative efficiency, meaning that the right share of resources is being devoted to health care versus other goods in the economy. The peculiar financing of health care in the U.S. fosters both allocative and productive inefficiencies, and different policies are needed to address each. Some have suggested that with system-level improvements, it would be possible to achieve the same level of health with thirty percent less spending (Fisher et al. 2003, 2003).

¹ David Cutler’s interpretation is somewhat different. He writes, “Money matters in health care as it does in few other industries. Where we have spent a lot, we have received a lot in return (Cutler, 2004).”

In an efficient system, more spending on health care would be a sign of prosperity and a harbinger of improved health and longevity, not a cause for concern. Achieving this means improving the incentives and infrastructure for providers to deliver – and patients to consume – high-value care, as well as wrestling with the difficult question of whom to cover versus what to cover in public insurance programs.

II. THE U.S. HEALTH CARE SPENDING LANDSCAPE

In order to gauge the causes, consequences, and potential remedies for rising health care spending, it is important to understand the drivers of spending in the current system (and how they do or do not align with incentives for efficient use) as well as the implications of that spending for overall economic welfare.

A. Health Spending in the Economy

Figure 1a shows the growth of health care spending in the U.S. Real health care spending (measured in 2009 dollars) has increased from \$330 billion in 1966 to an estimated \$2.5 trillion in 2009, and from 6% of GDP to 16%. Hospital spending is the largest category, although spending on prescription drugs has been growing more quickly. It may be tempting to explain this growth as being a consequence of “Baumol’s Cost Disease,” where labor-intensive sectors exhibit less productivity growth than other industries. But this should have happened all over the world, and Table 1 demonstrates that U.S. health care spending growth has been exceptionally high. Nor can these facts be explained by a model of increasing incomes leading to higher spending on health—other countries had similar income growth but did not expand health care spending as rapidly.

Health care is different from other goods in many ways that may interfere with the efficient allocation of resources based on marginal costs and marginal benefits: patients have limited information about the benefits associated with the care that they purchase (relying on the suppliers of that care for advice), and often need to make decisions in difficult circumstances; both insurers and providers often operate with limited competition. The way we finance health care adds an important additional set of wedges between the marginal cost of care and the marginal health benefit that care produces. Figure 1b shows how health spending in the U.S. is financed. Most care in the U.S. is purchased through private insurance (largely obtained through employment) or public insurance programs (largely Medicare and Medicaid, introduced in 1965).² Each mechanism has particular implications for the economy overall and for the efficiency with which health care resources are deployed, introducing inefficiencies in the use of health care resources that are not seen in the markets for most other goods and services.

B. Private Health Spending

The majority of private health insurance in the U.S. today is obtained through employer-sponsored plans, which cover more than 170 million lives (compared with under 20 million non-elderly covered through private non-group insurance plans) (DeNavas-Walt, Proctor, and Smith 2010). The dominance of employer-based health insurance in the private market has historical roots: during a period of war-time price controls it was determined that such plans were not subject to wage controls and were not taxable, making them a favored form of compensation. At

² It is worth noting that much of the debate over health reform confused the concepts of health care and health insurance. Health insurance is of course fundamentally about risk – it is valuable not just because health care is expensive, but because it is expensive and uncertain. Uninsured sick people do not need health insurance (once the illness is known it is no longer insurable), they need health care (Baicker and Chandra 2008). It is also important to acknowledge that many of the inputs into long run health outcomes are independent of the health care system – such as nutrition and exercise, smoking, environmental hazards, etc. The line between such inputs and preventive care is sometimes blurry (and of course not all care labeled preventive is in fact cost effective).

more than \$250 billion annually, this tax subsidization of employment-based private health insurance through its exclusion from payroll and income tax bases is as large as Federal spending on Medicaid (Kleinbard 2008) and larger than the mortgage interest deduction. It disproportionately favors those with higher incomes and jobs that offer generous health insurance benefits, and comes with the deadweight loss associated with the tax expenditures that finance the subsidy (discussed more below). Employer plans, however, form the primary mechanism for risk-pooling among the privately insured: the tax subsidy drives participation in health insurance by healthy as well as sick at pooled premiums. One consequence of this is that it makes changing jobs costly for those in poor health who would be unable to obtain insurance at favorable rates if they lost their employer plan (Madrian 1994).

Figure 2 shows the increase in premiums for family insurance plans purchased through employers. There is much public discussion of how the rise in health insurance costs affects U.S. businesses' international competitiveness, but simple economic theory notes that in the long run the cost of these benefits is borne entirely by workers (as long as they value the benefits) (Summers 1989). Similarly, the share of premiums nominally paid by workers is not indicative of the true burden (and although the dollar cost paid by workers has increased over time, the share has actually remained relatively flat).

In the short run, however, there are constraints that affect the ability of firms to shift costs to workers, and thus rising costs can result in fewer people covered by group policies (Baicker and Chandra 2006; Currie and Madrian 2000; Cutler and Madrian 1998). While the vast majority of large employers offer health benefits, over the last 10 years the share of small firms offering health benefits has dropped from 69 percent to 61 percent and rising costs are cited as one of the main factors influencing firms' decisions to discontinue offering benefits (Kaiser

Family Foundation and Health Research and Education Trust 2008) – meaning that rising health care costs can erode the risk-pooling in the private market and exacerbate adverse selection.

While most of the burden of rising private health care spending is borne in slower or stagnant wage growth (as an increasing share of compensation is devoted to health insurance benefits), the consequences of rising premiums can be particularly severe for low-income populations at greatest risk of losing jobs when the costs of benefits rise (Baicker and Levy 2008).

In addition to the labor market consequences of employer-sponsored health insurance, the tax preference for this insurance relative to out-of-pocket spending promotes first-dollar coverage plans that exacerbate moral hazard. There is always a trade-off between insurance protection and incentives: insurance protects against the bad state of the world in which individuals have bad health and need more medical care, but it lowers the marginal cost of that care thereby promoting consumption of care whose marginal benefit is above the patient's marginal copay but below the true marginal cost of the resources. The tax preference for insurance moves enrollees towards plans with even lower cost-sharing. Furthermore, while the cost of the employer-sponsored insurance is ultimately borne by workers, in the intermediate term that trade-off is not dollar-for-dollar for any individual worker, obscuring the trade-off and limiting the market discipline that would foster high-value insurance design and the optimal division of compensation between health insurance, wages, and other benefits.

There are additional constraints on insurance contracts that limit the availability of lower-cost, higher-value plans. In a world of evolving medical technology and complex care management, is impossible to fully-specify contingent contracts outlining all of the care that individuals can receive in every state of the world, and private insurers have limited ability to deny coverage for procedures with unproven benefits when Medicare covers these technologies

(Ferguson, Dubinsky, and Kirsch 1993). U.S. corporate laws also make it difficult for individual insurers and hospitals to reduce the use of technologies with variable payments: insurers and hospitals are not permitted to interfere with the medical judgment of physicians. State laws also require insurers to pay for any service deemed medically necessary by a physician. Because enrollees eventually leave private insurers (either for other plans or to enroll in Medicare), insurers face a dulled incentive to invest in care with health benefits that accrue only in the future. Together, these undermine the incentive for insurers to act as a residual claimant, leaving insurers to rely more heavily on negotiating lower prices with providers rather than restraining utilization to control spending.

C. Public Health Insurance

The landscape of health insurance coverage in the U.S. is largely shaped by the creation of Medicare, the federal program for the elderly and disabled (covering 49 million and costing over \$500 billion in 2009) and Medicaid, the joint federal-state program for the poor and other medically needy groups, but with demographic eligibility requirements that vary (covering more than 60 million - some of whom are also on Medicare – and costing about \$375 billion annually, \$250 billion from federal revenues and \$125 billion from states and localities). SCHIP expanded coverage for low-income children above the poverty level beginning in 1997 (covering 8 million and costing about \$11 billion). There are numerous other smaller programs covering, for example, military personnel and veterans, as well as the direct provision of care through community health centers. All told, public dollars account for about half of national health expenditures, even before accounting for the tax subsidy for private insurance described above. The Affordable Care Act expands Medicaid coverage and introduces additional subsidies for

low-income adults to purchase private insurance – adding substantially to public spending on insurance (offset at least in part by additional revenue sources).

While Medicaid poses an enormous fiscal challenge for state governments, Medicare poses the greater challenge for the Federal government. Medicare expenditures account for 15% of federal spending and 3.6% of GDP (Boards of Trustees of the Federal Hospital Insurance and Federal Supplemental Medical Insurance Trust Funds 2011). Medicare spending grew 2.5 percentage points faster than GDP from 1975 through 2008 (Congressional Budget Office 2011). The health care used by Medicare beneficiaries is financed by a combination of dedicated taxes and general revenues (as well as supplemental plans and the 25% of care that beneficiaries pay through premiums or out-of-pocket). The payroll taxes dedicated to financing hospital spending are paid into a trust fund, but since 2009 spending has grown more quickly than the tax stream, with the trust fund forecast to be exhausted in 2024. Program costs for physicians and other outpatient care and drugs are financed mostly through general revenues and beneficiaries' premiums. Rapid spending growth on these components means that general revenues are an increasingly important component of Medicare financing. While the aging of the population is clearly one driver of higher spending, the majority of rising expenditures can be attributed to higher spending per beneficiary (Congressional Budget Office 2011). Public health insurance programs thus comprise a rising share of both GDP and federal revenues (see Figure 3a).

Public insurance programs provide benefits that are extremely valuable to beneficiaries. Medicaid redistributes resources towards low-income and sick populations, and Medicare pools risk in a retiree population without access to employer pools. Finkelstein estimates that in 1963 only 25% of seniors had comprehensive insurance, but shortly after the advent of Medicare

virtually all did (Finkelstein 2007). This redistribution – from healthy to sick and from wealthy to poor – is a key feature of social insurance that is not achievable through private insurance.

These social benefits must be balanced, however, against the costs associated with the programs' financing. Raising taxes to pay for public insurance exerts a structural drag on the economy (deadweight loss). The health sector jobs created by low-value health care spending divert resources from higher-value uses. Deficit spending on health care also carries an economic cost: taxes are required to pay interest and principal on any borrowed money, and rising debt-to-GDP ratios eventually have serious adverse effects on the country's future ability to borrow. Projections of growing federal debt largely reflect anticipated increases in health care spending (see Figure 3b). Even if policy could halve the gap between the growth in health care spending and the growth in GDP, some estimates suggest that our debt-to-GDP ratio would drop only from 280% to 200% by 2050 (Kogan, Cox, and Horney 2008).

There is a substantial body of research exploring the magnitude of the deadweight loss created on the margin in current the U.S. tax structure – with many estimates around 0.3, meaning that every dollar of public spending comes with an additional cost of lowering economic activity by 30 cents (Gahvari 2006; Feldstein 2006, 1999, 1973; Fullerton and Henderson 1989). These tax distortions are not just static in nature (lowering economic activity at a point in time), but can also affect economic growth by distorting investment decisions (Engen and Skinner 1996).

Financing current trends in public program costs would necessitate dramatic increases in tax rates. The Congressional Budget Office (CBO) estimated (before the Affordable Care Act) that income tax rates would have to increase by more than 70% to finance health care spending that grew 1 percentage point faster than GDP, and would have to increase by more than 160% by

2050 to finance growth at the historical rate of 2.5 percentage points faster than GDP growth. Even with just 1 percentage point excess growth in health care spending, the CBO estimates that the tax increase would reduce GDP by 3 to 14% (Congressional Budget Office 2007). This additional cost of public health insurance must be included in any cost-benefit analysis of the program.

Most forecasts of health care spending growth assume that there must be a break in current trends – health care spending simply cannot continue to grow at historical rates – but it is not clear what will cause that slow-down. Getzen postulates that health care spending at the country level is driven in large part by the ability to finance health care, and this hypothesis is consistent with evidence presented in Baicker and Skinner that countries with a higher tax-to-GDP ratio in 1979 experienced significantly slower growth in health care spending between 1980 and 2008 (Baicker and Skinner 2011; Getzen 1992). Rising costs of taxation (and debt) may thus ultimately serve as a brake on health care spending growth – creating the impetus to overcome stakeholder resistance to policies that improve the efficiency of spending.

The deadweight loss associated with the public financing of programs like Medicare is not the only inefficiency they generate. Medicare performs no evaluation of the benefits associated with new medical technologies, and in its fee-for-service incarnation does not ask if care could be better managed. Both features discourage cost-saving innovation and efficient insurance offerings. As an example, consider that traditional Medicare does not coordinate hospital, outpatient, ambulatory and prescription drug care. This means that care is often “siloeed,” with little attempt made to choose modes or sites of delivery that work best and with provider reimbursements not informed by their broader effects on care. It also means that if a private insurer providing prescription drug coverage changes program features in a way that

reduces prescription drug use but results in patients requiring more hospital care, the Medicare program bears that cost. Chandra, Gruber and McKnight (2010) find evidence of exactly these offsets.

Other inefficiencies are not restricted to the Medicare program, affecting physician practice style in the non-Medicare population. When Medicare covers Provenge for prostate cancer (at a cost of over \$90,000 for a few weeks of survival) or the latest cyclotron-based proton-beam therapy (an unproven treatment with a 120 million price tag requiring a football field sized accelerator), private insurers are likely to follow the coverage decision to avoid litigation where their patients claim that insurers are withholding valuable care. Medicare regulatory boards evaluating new technology focus on whether drugs or procedures provide positive benefits, rather than costs.³

There is also a substantial literature at the provider level showing that practice pattern norms drive similar care for all of the patients that a provider sees, regardless of individual insurance status – so that changes in the incentives applying to a large share of patients (say, Medicare beneficiaries) can drive changes in the care received by all patients (Glied and Zivin 2002; Baker and Corts 1996; Baker 1999; Frank and Zeckhauser 2007). Medicare’s administered pricing scheme is frequently out-of-step with real resource costs and subject to political manipulation (as witnessed by the regular override over the “sustainable growth rate” formula intended to cut physician fees when total Medicare spending rises too quickly).

³ In other countries, boards such as England’s National Institute for Health and Clinical Excellence (NICE) explicitly use costs to determine coverage decisions. Typically these rulings are limited to specific and discrete choices, for example whether specific drugs can be used for a specific disease. In practice, treatments with heterogeneous benefits are more difficult to classify under “cover/not cover” decision rules.

On the patient side, in an attempt to control the overuse of care, the basic Medicare benefit requires beneficiaries to pay substantial cost-sharing – almost 20 percent coinsurance for outpatient care. But almost all beneficiaries buy additional Medigap wrap-around plans, eliminating most copayments and leading to greater moral hazard (which imposes additional costs on the Medicare program itself). The Medicare Advantage program was intended to promote innovation in insurance coverage and provision of care, but there has been limited evidence of success to date; the proposals for “premium support” (in essence an insurance voucher) have similar goals.

Thus neither the patient- nor provider-side of the system fosters decision-making that weighs costs against benefits. Nor is there an evaluation of the benefits of additional spending through one program against the benefits of additional spending through another. The spending on care of questionable marginal benefit fostered by the Medicare program, for example, might come at the expense of covering more people in the Medicaid program who would consume infra-marginal care of high value (Baicker and Chandra 2010). Recent research suggests that Medicaid coverage substantially expands the use of preventive care in low-income adult populations, but such populations are not covered by Medicaid in most states (Finkelstein et al. 2011).

III. ASSESSING THE EFFICIENCY OF HEALTH SPENDING

The above discussion has focused on specific inefficiencies in the design and financing of public and private health insurance plans. The tax-preferred nature of financing private plans, the limited use of cost-effectiveness analysis, and the overly generous public reimbursement of care with questionable health benefit might suggest that the U.S. spends too much on health care relative to other goods. At the same time, we have seen dramatic increases in life expectancy

and reductions in disease burden over the last century (Cutler, Deaton, and Lleras-Muney 2006). Some new interventions are cheap (antibiotics, statins, beta-blockers, vaccines – many of which are even underused) and others are more expensive but still remain cost-effective (HAART therapy, primary PCI after heart-attacks), but medical technology has high health returns on average (Cutler, Rosen, and Vijan 2006). These facts might suggest that the U.S. spending on health care has been worth it, and may even be too low in some areas.

How can we reconcile these facts? Drawing on Chandra and Skinner, we develop a framework for understanding the sometimes conflicting evidence and the drivers of differential productivity of health spending (Chandra and Skinner 2011; Weinstein and Skinner 2010). We make the key distinction between *productive efficiency*, meaning that health care resources are put to the best use possible and produce as much health as they can, and *allocative efficiency*, meaning that the right amount of resources are being devoted to health care, versus other goods in the economy. When the evidence on the value of health care spending is evaluated in light of this framework, we see that what might at first appear to suggest that we are spending too little on health care may actually be evidence only that we have suboptimally deployed what we are already spending. It is in fact impossible to evaluate allocative inefficiency in the presence of productive inefficiency.

A. Two Kinds of Efficiency

We begin with the world of productive efficiency, in which health care resources are put to the best use possible. Figure 4 illustrates the association between spending on factor inputs (health resources like physicians, scans, hospitals) on the horizontal axis and survival/quality of life on the vertical axis. A concave production possibility frontier (PPF_1) illustrates the aggregate health that is achievable for a given level of inputs for the United States, or any

particular delivery system which could be the UK, or the Mayo Clinic. This is the relationship between health and spending in a world where we have solved all of the problems that interfere with efficient production, described below. The PPF assumes diminishing returns to spending: higher levels of spending are associated with inputs that generate smaller and smaller health benefits; CT scanners, proton beam accelerators, and chemotherapy will have lower benefit per dollar spent relative to aspirin, beta blockers, and hospital care for pneumonia. Under diminishing returns, physicians are first giving treatments to the patients who benefit most from the treatment, and then moving into patients with lower benefit. Technological advancements can push the PPF up and out.

In such a world, different points on the frontier (such as A_1 or B_1) represent different choices about allocating spending between health and non-health consumption (cars, vacations, education, defense). If the system is allocatively efficient, that choice would be made such that the last dollar spent on health generates the same marginal utility as spending an additional dollar on non-health consumption; otherwise, society would be made better off by moving resources from the lower marginal utility use to the higher. Because the marginal utility of spending on health care is falling, point B_1 produces more health than point A_1 , but it could certainly be the case that the incremental improvement in health outcomes gained by moving from A_1 to B_1 is not worth the opportunity cost of non-health spending that was given up (which may be even higher for publicly financed care that comes with deadweight loss). If so, we would say that choosing point B_1 is allocatively inefficient—we are on the efficient frontier (producing as much health as possible with given health care resources), but have spent too much on health care. Why might a system allocate at B_1 rather than A_1 ? The previous section describes reasons that health care markets do not resemble efficient markets for other goods and services. Several of these factors,

from the tax-preference for employer-sponsored insurance and the moral hazard it generates to fee-for-service nature of Medicare coverage, push towards a greater allotment of resources to health care.

Evidence suggests, however, that we are likely in a world in which we are *not* on the health care frontier, but in its interior (see line labeled PPF₂ in Figure 4). This could happen for a variety of reasons. Fee-for-service reimbursement combined with low levels of competition (most health care is consumed locally) encourages providers to deliver care that they can bill for and underinvest in keeping patients healthy. They fail to do very simple things like hand-washing, prescribing prophylactic antibiotics before surgery, or using beta-blockers for heart-attack patients (Skinner and Staiger 2009). Much (but by no means all) investment in prevention would yield high returns in future health. Similarly, inadequate health IT reduces the productivity of all resources used by increasing fragmentation, errors, and duplication. These sins of omission coexist with sins of commission; estimates suggest that the overuse of CTs and MRIs cause 1.5-2.0 percent of total cancers (Brenner and Hall 2007).

A particularly insidious form of productive inefficiency arises if medical technologies are adopted in the wrong sequence or deployed on patients in the wrong order – potentially even creating a convex production function (see Figure 5). Adopting angioplasty, CT scanners, proton beam therapy, and Provenge for prostate cancer patients before patients with heart-attacks receive aspirin would result in this type of productive inefficiency. This could result from insurers paying for these intensive technologies more generously than other more therapeutic interventions (care management, reduced errors, lower readmissions, aspirin). With these jumbled adoption decisions, it will not be the case that the marginal dollar spent on health care is purchasing the marginal (most intensive) technology—if low-value technologies are adopted

before higher values ones, then an empirical analysis of the relationship between spending and health outcomes would find that increases in spending lead to improvements in health, but this would be entirely a consequence of productive inefficiency in how health care technologies are utilized. In this case, higher spending could even exhibit increasing marginal returns (compare the slopes at points A and B). At the patient level, some may have generous insurance that covers every new technology regardless of effectiveness (Medicare), leaving less money to cover others who would benefit from prevention or low-cost treatments for chronic disease (Medicaid).⁴ In the presence of these productive inefficiencies, analysts might conclude that marginal spending was incredibly valuable and that health spending should be higher. But such an interpretation would miss the tremendous opportunities to improve productivity by reallocating existing spending.

B. Interpreting the Evidence

What is the evidence on the extent of these inefficiencies? The simplest way to evaluate the presence of allocative inefficiencies would be to measure the spending required to save an additional (quality-adjusted) year of life, and to evaluate whether the gain was “worth” the cost. The key idea is that spending on health care may be expensive, but may produce gains that are more than commensurate with the additional spending. Garber and Skinner report that the average cost per life year gained has risen substantially over the decades, from \$64,000 during the 1970s to \$247,000 in the 1990s (assuming that half of the gains in life expectancy are attributable to health care spending) (Cutler, Rosen, and Vijan 2006; Garber and Skinner 2008; Cutler and Madrian 1998). This cost is substantially higher than the threshold of \$100,000 per

⁴Finkelstein et al. (2011) use a novel experimental design to evaluate the effect of extending Medicaid coverage to previously uninsured adults. They find that expanding coverage increases the use of outpatient (including preventive) care, hospitalizations, and prescription drug use; reduces financial strain; and improves self-reported physical and mental health and well-being.

quality-adjusted life year often used to evaluate policy, suggesting that we may be spending too much on health relative to other goods.⁵ And even this number likely underestimates what it costs to produce a year of life, as it is based on averages, and does not measure the benefits from marginal spending where allocative efficiency should be judged.

Chandra and Skinner (2011) provide another cautionary tale about how allocative efficiency can be measured, arguing that improvements in health are often generated because of relatively low-tech technological progress. One recent study provides direct evidence of this theory to explain the decline in mortality from coronary disease during 1980-2000 due to specific factors (Ford et al. 2007). Over 40 percent of decline can be attributed to changes in smoking, physical activity, blood pressure, and cholesterol (and offsets from rising obesity and diabetes). 35 percent of the decline is attributed to inexpensive but highly effective treatments: aspirin, beta blockers, blood-thinning drugs, anti-hypertensives, diuretics, and pharmaceuticals such as ACE inhibitors, statins, and thrombolytics (“clot-busters”). Many of these treatments represent movements from A_2 to A_1 in Figure 4. Expensive treatments, such as angioplasty (stents), bypass surgery, and cardio-pulmonary resuscitation (such as automated defibrillators) explain 11 percent of the improvement, but almost all of the spending increase (movement from A_1 to B_1). These facts caution against attributing the improvement in health to concomitant spending increases—the bulk of the health improvement may not have been caused by spending.

An alternative way to measuring allocative efficiency would be to compare spending and

⁵ Readers may be interested in the evolution of this number; our discussion is adapted from Jena, Chandra and Skinner (2011). The willingness to pay for a life year was initially based on a 1984 Canadian study of patients with end stage kidney disease on dialysis. That study computed these costs at \$50,000 per QALY (Winkelmayer, Weinstein, Mittleman, Glynn, and Pliskin, 2002). Almost thirty years of inflation have increased it to \$100,000 (Lee, Chertow, and Zenios, 2009). Garber and Phelps (2007) offer an alternative rationalization: an annual salary of \$30,000 for a 40-hour work week would lead to a value of a life-year of \$100,000 if leisure time was priced at average market wages.

outcomes in different geographic areas to see the size of the gains in outcomes that accrue to higher spending, and then to compare those to productivity in other sectors. These studies have had mixed findings: some characterize productivity as high (such as (Ong et al. 2009; Romley, Jena, and Goldman 2011; Silber et al. 2010; Bach 2010)). But similar study designs have also pointed to low productivity (Skinner, Staiger, and Fisher 2006; Yasaitis et al. 2009). With these conflicting facts, it is unsurprising that some analysts believe that spending more would improve health by enough to warrant the cost, while others believe that it would not.

We believe, however, that both of these methods are flawed. In other work we have demonstrated that much of the regional variations literature implicitly assumes that all regions are on the same production function, while in reality they are on different production functions (Chandra and Staiger 2007). This would happen if delivery systems specialize in how they deliver care. Some may specialize in technologically intensive treatments while others may specialize in less expensive treatments. Because of this specialization, even though two systems or countries may have similar health outcomes with one achieved at substantially lower cost, simply cutting spending in the high-cost system will not allow it to achieve the results of a low-cost system, but will almost certainly harm patient health in the process. Finding that two systems have similar outcomes at dissimilar levels spending levels is therefore not informative about whether we should spend more or less on health care: this question cannot be answered in the presence of productive inefficiency.

What then, do we know about productive efficiency – could the same health gains be achieved for less? Chandra and Skinner argue that the peculiar financing of the health care system results in many high-tech treatments being adopted before lower tech (but cheaper) ones are. This suggests productive inefficiencies of the type illustrated in Figure 5. Many expensive

technologies are used more extensively in the U.S. than in other countries, such as cardiovascular procedures, at 587 per 100,000 in the U.S. compared with 207 in Denmark (Peterson and Burton 2007). Simultaneously, there are many cases of underuse of effective treatments in the U.S. compared with other nations (Cutler and Ly 2011).

We interpret the geographic variations literature as providing further evidence of productive inefficiency. There is evidence that each of these regions is on its own production function, where places that spend less have comparable health outcomes to higher-spending areas because they have adopted all the right low-cost technologies (comparing point A_1 to B_2 in Figure 4). Figure 6 shows the variability in spending on Medicare fee-for-service beneficiaries across different states, as well as the fact that higher-spending areas are not those in which beneficiaries receive the highest quality care. Figure 7 shows that there is more variability across states in physician and outpatient care than in inpatient care. That there is a negative correlation between the use of relatively low-tech treatments and expensive high-tech treatments of dubious value further suggests inefficient adoption (Yasaitis et al. 2009; Baicker and Chandra 2004). It also suggests the tremendous scope for increasing productive efficiency through policies that foster higher-value use without encouraging low-value use. Once we move closer to productive efficiency, issues of allocative efficiency can then be addressed more effectively.

IV. POLICY LEVERS FOR IMPROVING THE EFFICIENCY OF HEALTH CARE SPENDING

The distinction between productive and allocative efficiency can help guide analysis of policy levers. Different policy levers operate on different sources of inefficiency, and allocative efficiency cannot be evaluated when productive inefficiency persists. We separate our discussion of policy levers into those that operate on the patient side, provider side, and the environment in which care and insurance are purchased.

A. Provider-Side Incentives

Reforming the way public insurance programs pay providers could promote both productive and allocative efficiency. Medicare administered prices play a particularly large role in creating fragmented patterns of care, and moving away from fee-for-service payments could provide much better incentives for efficient allocation of resources. Current payment structures provide little incentive for cost-saving innovation (Cutler 2010) or for providing lower-cost care. For example, a back pain clinic loses revenues by steering patients to low-cost rehabilitation rather than equally-effective but costlier diagnostic tests and back-surgery (Fuhrmans 2007). A change in payment incentives away from traditional fee-for-service thus has the potential to improve health and moderate spending growth. Similarly, system efficiency would be improved by creating incentives not only to adopt cost-effective treatments, but to adopt the most cost-effective treatments first and to apply them first to the patients for whom they produce the greatest value. The imperfect incentives created by current payments reduce productive efficiency by making relatively high payments for is not necessarily of the highest value, and interfere with allocative efficiency by committing (uncapped) resources to care that may have lower value than alternative uses of funds.

Rewarding more efficient delivery is a logical strategy for increasing the productivity of health care spending.⁶ Medicare has in fact experimented with alternative payment systems. Hospitalizations are reimbursed by bundled payments based on DRGs (diagnosis-related groups). These prospective payments are imperfect, but can limit the incentives to provide

⁶ Historically, the U.S. has not relied on quantity restrictions in health care planning, except for sometimes ineffective certificate of need programs (Ho, Ku-Goto, and Jollis 2009). In theory, regional restrictions to discourage duplication and overbuilding of surgical units, MRIs, and hospital or ICU beds could control unrestricted growth in procedures of questionable value. But this approach to market-level quantity restrictions would require a tectonic shift in the U.S. regulatory and policy environment.

excess hospital care. Perhaps as a result, since 1992 there has been limited per capita real growth in Medicare spending on these types of DRG-related categories (Weisbrod 1991). Weisbrod captured the dynamics between reimbursement systems and R&D and optimistically described the new prospective payment system:

With a hospital's revenue being exogenous for a given patient once admitted ... , the organization's financial health depends on its ability to control costs of treatment. Thus, under a prospective payment finance mechanism, the health care delivery system sends a vastly different signal to R & D sector, with priorities the reverse of those under retrospective payment. The new signal is as follows: Develop new technologies that reduce costs, provided that quality does not suffer "too much." (p. 537)

Weisbrod was too sanguine about this system: it created DRGs for new procedures (as opposed to illnesses), which unravels the "prospective" part of the payment and leaves the incentives for innovation largely unchanged. There are incentives to "unbundle" outpatient and post-acute care, and exempting certain capital costs from the expenditure cap continues to encourage capital spending. Weisbrod's fundamental insight is still valid: only when health care is reimbursed based on value will innovations be focused on cost-effective treatments. The greatest saving could arise not so much from new cost-saving devices, but instead from reducing the organizational fragmentation inherent in U.S. health care (Cebul et al. 2007). The more than two-fold difference in risk- and price-adjusted costs across top-ranked U.S. academic medical centers (Fisher et al. 2004) is not due to "what" is provided, since nearly all academic medical centers have access to the latest technology. Rather, the difference is related to "how" care is provided or the organization of the health care system – the frequency of follow-up visits, referrals to subspecialists, hospital days, and the intensity of diagnostic testing and imaging procedures.

To confront the challenge of inefficient delivery systems, some have embraced the idea of an integrated delivery system with capitated payments. Because such systems retain the

savings from better prevention, lower readmission rates, and better medication adherence, they have better incentives to avoid therapies of questionable benefit. The 2010 U.S. Affordable Care Act (ACA) focused on “accountable care organizations” (ACOs), where shared-saving “bonuses” are provided to health care organizations that are able to provide high quality care at lower costs. Examples include integrated systems such as Intermountain in Utah or the Geisenger Clinic in Pennsylvania, but also traditional hospital-physician networks (Fisher et al., 2009). These may come closer to the ideal expressed by Weisbrod by providing incentives for cost-saving innovations, and gut the incentives that physicians currently have to engage in “financial entrepreneurship” of the type seen in McAllen, TX, Elyria OH and Redding, California (Gawande 2009). However, we do not know how well ACOs will sidestep cost-ineffective technologies, particularly if the latest shiny innovation increases market share. The viability of ACOs will depend on the receptiveness of physicians to capitated payments—some specialists will see their incomes fall and are unlikely to take these cuts quietly. While their concerns may not resonate with patients, they might if providers claim that valuable care is being withheld. Designers of ACOs are therefore keenly interested in measuring ACO performance and patient satisfaction, but current quality measures only capture truly negligent care.

If payment-system reform is particularly promising on the public side, what about the private side? As noted above, private insurance coverage is heavily influenced by the norms driven by Medicare coverage, so these reforms are likely to have system-wide spillover effects. It is natural to ask why private providers have not adopted ACOs or more bundled payments on their own. This remains a puzzle. One explanation is that it is a coordination problem—all insurers may want to adopt larger bundled payments, but no single insurer can make the transition. This is certainly consistent with the historical record on the adoption of prospective

payment for hospital care. Once it was introduced in Medicare, private plans were quick to adopt it. Similarly, private hospitals were quick to use the federal Government's efforts to measure quality of care even though nothing stopped them from forming consortiums to measure quality before these federal efforts.

What might be the savings from these reforms? Some estimates would put the number at 20-30 percent of spending, but that assumes that low-cost regions are at the frontier. But even low cost regions may not be producing efficiently. Medicare is expressly prohibited from selectively contracting with more efficient physicians. Growth in home-health care, outpatient visits, office visits and diagnostic testing – areas of care where reimbursements are based on volume and the “right rate” is not known – have exhibited very rapid growth in all areas. At least for home health care, the additional spending appeared to have no impact on health outcomes (McKnight 2004). All of this suggests that we may save more than 30 percent if the right incentives are in place; Buntin and Cutler place the figure at more like 50 percent (Buntin and Cutler 2009). It is worth reiterating, however, that some of the savings from lower quantities may be offset with higher prices as ACOs exert market power.

System reform such as ACOs may thus improve productive efficiency by reorganizing the delivery system. Capitated reimbursement can affect not just the level but the growth of spending through slowing the adoption (and innovation) of new medical technologies. But achieving allocative efficiency also requires carefully-calibrated ACO payments and updates over time. Additional tools on the patient side can help promote allocative efficiency.

B. Patient-Side Incentives and Insurance Design

While insurance provides highly valuable protection against financial risk, it comes with traditional moral hazard: patients consume more health care when the cost is lower (Newhouse

and The Insurance Experiment Group 1993). First-dollar private insurance is promoted by the tax code, as described above, and Medicare's inconsistent cost-sharing is undermined by the pervasiveness of wrap-around supplemental insurance. This promotes greater health care consumption than allocatively efficient, and also interacts with provider-side incentives to undermine productive inefficiency by dampening the market discipline that price-sensitive consumers would provide.

Several policy proposals aim to increase the incentives for patients to consume only care that is sufficiently valuable to them, a key condition for allocative efficiency: it is necessary (but far from sufficient) that patients face the right price for the health care that they consume, balancing the increased health care use from moral hazard with the financial protection that insurance provides. The tax exclusion for employer-sponsored health insurance pushes enrollees towards policies with less cost-sharing than optimal. There is widespread support among economists for limiting this tax exclusion, even beyond the limited reform in the ACA (Gruber, 2009). Other proposals include expanding the tax-preference for high-deductible policies.

How might a demand-side approach affect the adoption or diffusion of new technology? There is some anecdotal evidence that higher out-of-pockets payments would dissuade high-income patients from consuming cost-ineffective treatments, suggesting that cost-sharing could discourage some use of expensive and only marginally-effective technologies.⁷ More sophisticated approaches aim to impose different cost-sharing for different patients and different procedures, based on the health value of the care. In most conventional models, cost-sharing should be based solely on the price elasticity of demand for health care (which governs the

⁷ One former corporate finance officer whose insurance covered the \$93,000 prostate cancer drug Provenge was explicit about his willingness-to-pay: "I would not spend that money," because the benefit doesn't seem worth it." (Marchione 2010). Others have pointed out the difficulty of judging "true" demand for health care, particularly among those near the end of life (Becker, Murphy, and Philipson 2007).

degree of moral hazard). For well-informed patients who are aware of marginal costs and benefits, this produces efficient use of care (Pauly and Blavin 2008). But mounting evidence from the behavioral economics literature suggests that patients may not always make optimal decisions about the use of services when faced with higher copayments and coinsurance.⁸ “Value-based” insurance plans aim to increase efficiency by imposing higher patient prices for lower-value care. Prices (or the degree of patient-cost sharing) could be set based on evidence from clinical trials, and could also vary based on patient characteristics. While this approach has been applied to lowering prices of high-value care with some success (e.g. Chernew et al., 2008), insurance plans that impose significant copayments for lower-value care are rare, perhaps because of the difficulty in determining the idiosyncratic benefit that a patient may attach to these services, and because of regulatory or legal barriers. Patient-side approaches thus have the potential to limit low-value spending, but require a more nuanced approach than seen thus far (Baicker and Goldman 2011).

Designing coordinated insurance plans such as ACOs may help patients choose higher-value coverage and move towards allocative efficiency. ACOs could, for example, aim to attract patients with greater coverage of marginal medical technologies, higher quality, and broader provider networks, with patients responsible for the marginal cost of more generous policies (where the threshold is based on the value of alternative uses of funds). This is much broader than using cost-sharing to affect the choice of generic drugs versus formulary drugs, or a CT-scan in a suburban shopping mall versus one in a hospital outpatient setting. Insights from the behavioral economics literature suggest that patients may be stymied by the cognitive demands

⁸ Most evidence suggests that health is not adversely affected by greater patient cost-sharing on average (Chandra, Gruber, and McKnight 2010; Newhouse et al. 1993).

of making many price comparisons, but choice may be facilitated by creating larger bundles for patients to choose among.

C. Health Care and Health Insurance Markets

There are several features of the way health care and health insurance are purchased that affect both providers and patients – and indeed the types of care and insurance products that are available – with implications for the efficiency of use of health resources. One popular explanation for inefficiency offered by physicians and provider groups is that the medical malpractice environment provides massive disincentive for efficient use by fostering “defensive medicine,” with providers delivering more care than is socially optimal to avoid potential lawsuits. Such behaviors seem more likely in uncompetitive markets – otherwise if a given doctor provided more care without associated benefits, others could offer higher-value care that was more attractive to patients and insurers. The evidence on the magnitude of defensive medicine, however, suggests that it is not a primary driver of higher health care spending – accounting for less than 3% of health care spending in the U.S. (Mello et al. 2010). Nor is there a lot of evidence that areas with the greatest malpractice pressure provide the most care (Baicker, Fisher, and Chandra 2007). While malpractice reform is needed to ensure that the victims of negligence are compensated and that compensation is not paid to those who are not, more stringent damage caps or other reforms are not likely to result in first- or even second-order savings.

A more systemic issue is that provider and insurer market power can result in higher-than-competitive health insurance prices and premiums (Dafny, Duggan, and Ramanarayanan 2009; Dafny 2008). Market power does not affect productive efficiency per se, but could certainly affect allocative efficiency by increasing prices and reducing quantities relative to a

competitive market. Enthusiasm for the potential of ACOs to promote coordinated care must be moderated by their potential to increase market power, making antitrust enforcement even more important.

Perhaps the most important contribution that public policy could make to system-wide efficiency would be to generate more information – for both patients and providers – about what care is in fact high value (Weinstein and Skinner 2010). There is as yet little evidence on comparative effectiveness for a vast array of treatments: for example, we don't know whether proton-beam therapy offers *any* advantage over conventional approaches. Moreover, most drug studies compare new drugs to placebos, rather than “head-to-head” with other drugs on the market, failing to generate evidence on which drug works best. One area where cost effectiveness analysis may prove to be particularly powerful is in evaluating the relative efficiency of different delivery systems and institutional organization (rather than the relative effectiveness of particular drugs or procedures) – such as drop-in clinics rather than emergency room care. Tiered networks, with greater cost-sharing for wider or less efficient networks is one way to blend demand side influences with supply-side reforms.

Because such information is a public good, it is underprovided by the market, and that suggests a role for government subsidization of trials. The trials required to ascertain effectiveness will be expensive. At present, the NIH budget is just over \$30 billion a year (\$100 per person); we are devoting less than 1 percent of health care resources to learning what works. Improving our knowledge of what works is a prerequisite for increasing productive efficiency and eliminating medical practices that are unsafe at any price or dominated by other treatments.

The challenge with effectiveness studies is not just how to undertake them, but also how to use the results to inform reimbursement or patient cost-sharing. Medicare is forbidden from

using cost-effectiveness analysis and, as long as this continues, simply providing more information on what works is unlikely to affect coverage and reimbursement decisions by private insurers. Effectiveness studies could still be used to design patient cost-sharing, but even with better demand-side incentives, government agencies must decide whether (and at what price) Medicaid and Medicare will cover new treatments, or how to adapt coverage and cost-sharing rules as the implicit value of additional health benefits evolves over time. It is not clear how decisions about the size and growth of these large tax-financed government programs will be made. Tying their global budgets to GDP growth would restrain expenditures, but might not result in the optimal share of resources being devoted to public insurance versus other responsibilities of government such as education and defense versus private goods.

V. CONCLUSION

Technological advancement in medicine will continue to produce innovations that offer great benefit to some patients, but can easily be overused in others. It is imperative that health care systems be designed to foster that innovation and promote its use in patients to whom high health benefits will accrue without incurring massive government debt to cover its use in patients with little or no benefit. “Solving” the fiscal imbalances in current public programs such as Medicare, by raising taxes or crudely cutting reimbursements, will not in and of itself improve the efficiency of the care that is delivered. The need for fundamental reforms is heightened by the demographic pressures of the aging baby-boom generation, but arises more fundamentally from the increasing cost of health benefits per person.

There is no single strategy that is likely to achieve efficient use of health resources. On the provider side, payments through public insurance programs can be bundled to encourage coordination, and providers can share in the financial gains of improving the efficiency with

which they deliver care. On the patient side, more nuanced cost-sharing and leveling the playing field for higher-cost-sharing insurance plans can encourage patient involvement in decision-making and the balancing of resources costs against health benefits, as well as fostering competition. System-wide, better information is needed about which delivery systems – not just which drugs or procedures – are most effective.

The U.S. has yet to wrestle with the question of public policy priorities in a world of scarce resources: even with perfect productive efficiency, we cannot cover all services for all people. When public resources come at a cost of lower economic growth, there must be some explicit consideration of the value of redistribution, and the public priority placed on covering different levels of service for different parts of the population. By first ensuring that health care resources are used more productively, we will be in a much better position to move towards spending the “right” amount on health.

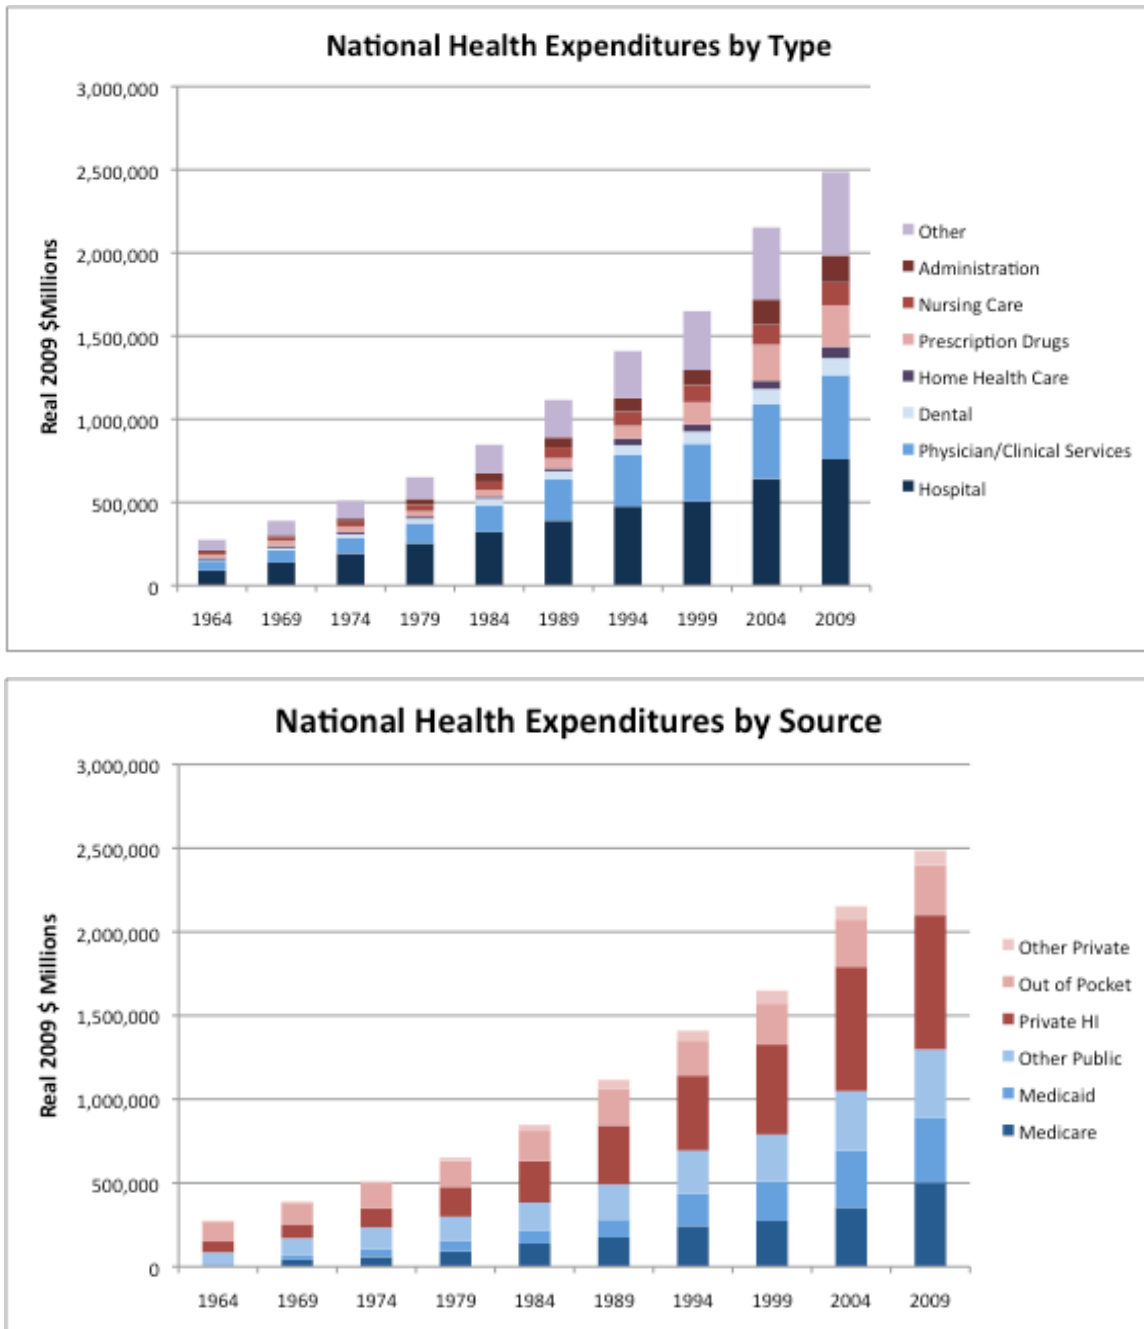
**Table 1: Health Care Share of GDP for OECD Countries
1980 and 2008**

Country	Health Care Share of GDP		Annual Growth Rate (1980-2008)		Public Share of Health Spending
	1980	2008	Health Care Spending	GDP Per Capita	2008
Spain	5.3	9.0	4.5	2.5	72.6
Portugal	5.3	9.9*	4.4	2.4	65.1
UK	5.6	8.7	3.8	2.2	82.4
Greece	5.9	9.7	3.1	1.5	
New Zealand	5.9	9.8	3.0	1.1	80.3
Australia	6.1	8.5	2.9	1.8	68.0
Belgium	6.3	10.2	3.3	1.4	75.0
Finland	6.3	8.4	3.0	1.9	74.4
Iceland	6.3	9.1	2.4	1.0	82.6
Japan	6.5	8.1	2.5	1.8	80.8
Canada	7.0	10.4	2.9	1.5	70.5
France	7.0	11.2	3.1	1.4	77.7
Norway	7.0	8.5	4.3	3.5	84.3
Switzerland	7.3	10.7	2.4	1.0	59.5
Austria	7.4	10.5	2.8	1.6	77.2
Netherlands	7.4	9.9	3.1	2.1	
Ireland	8.2	8.7	4.2	3.8	76.7
Germany	8.4	10.5	1.8	1.6	76.6
Denmark	8.9	9.7	1.9	1.7	84.7
Sweden	8.9	9.4	1.6	1.4	81.5
<i>United States</i>	<i>9.0</i>	<i>16.0</i>	<i>4.0</i>	<i>1.8</i>	<i>46.0</i>

Source: Chandra and Skinner, 2008, drawing on OECD, 2010. Data for either 2008, 2007, or 2006(*). For Germany, the 1980 values are for West Germany.

Alternatively, one can calculate growth in health care spending relative to GDP for 1980-1990 and for 1992-2008 (thereby avoiding the transition). This yields an increase of 0.8 percentage points of GDP in health care expenditures. Share of health care spending comprised of public dollars does not include tax exclusion of employer sponsored insurance.

Figure 1: U.S. National Health Expenditures

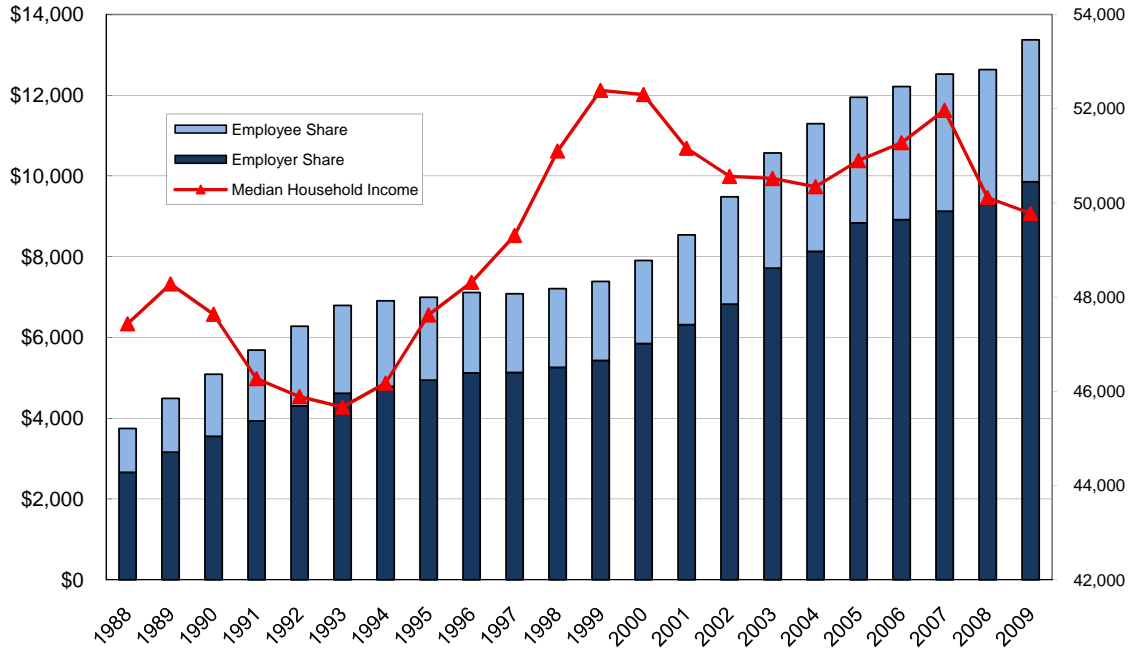


Source: (Centers for Medicare and Medicaid Services, 2011).

Note: Figures are in real 2009 millions of dollars, calculated from national health expenditures from CMS deflated by CPI

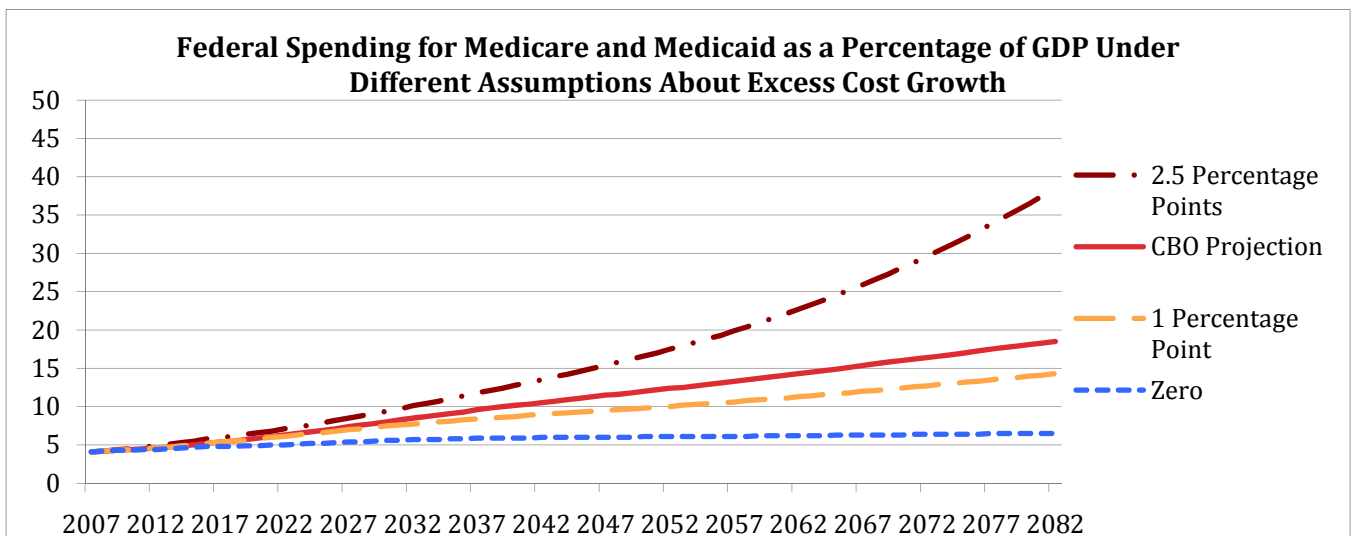
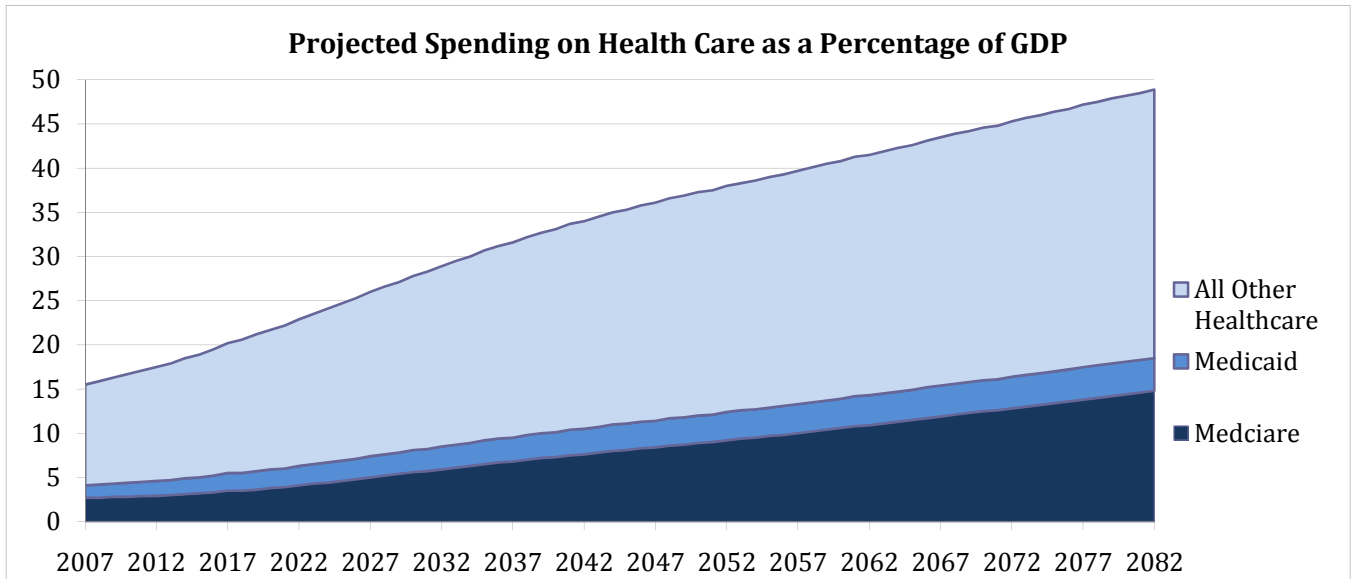
Figure 2: Health Insurance Premiums and Wages

Employer-Sponsored Health Insurance Premiums for Family Plans
(1988-2009, adjusted for inflation)



Sources: Kaiser Family Foundation/Health Research and Education Trust; U.S. Census Bureau
Notes: Premiums on left axis, household income on right axis, both in real 2009 dollars. Health insurance premiums for the following years were interpolated: 1989-1992; 1994-1995; 1997-1998.

Figure 3: Health Care Spending and the Economy



Source: Congressional Budget Office 2011

Figure 4

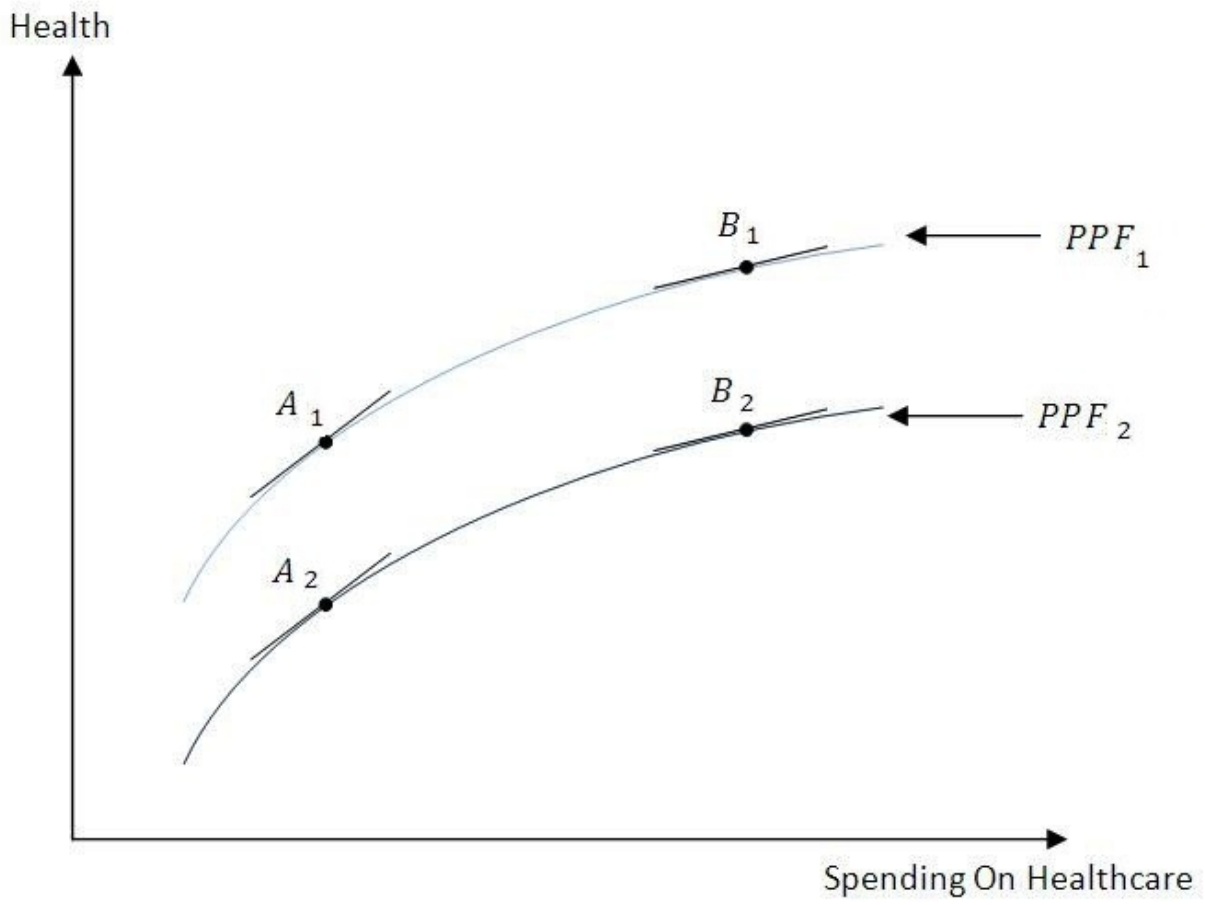


Figure 5

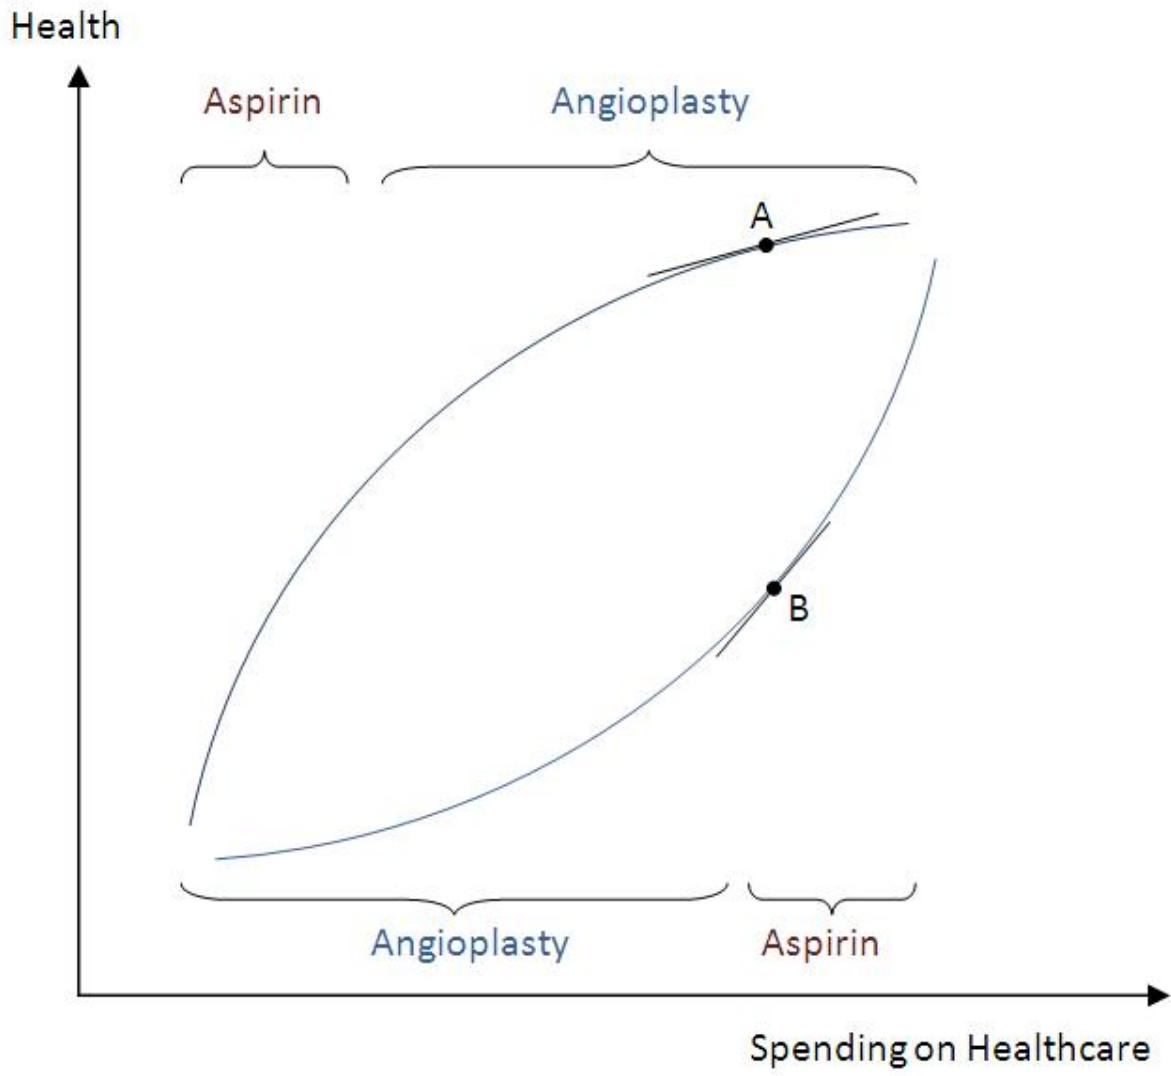
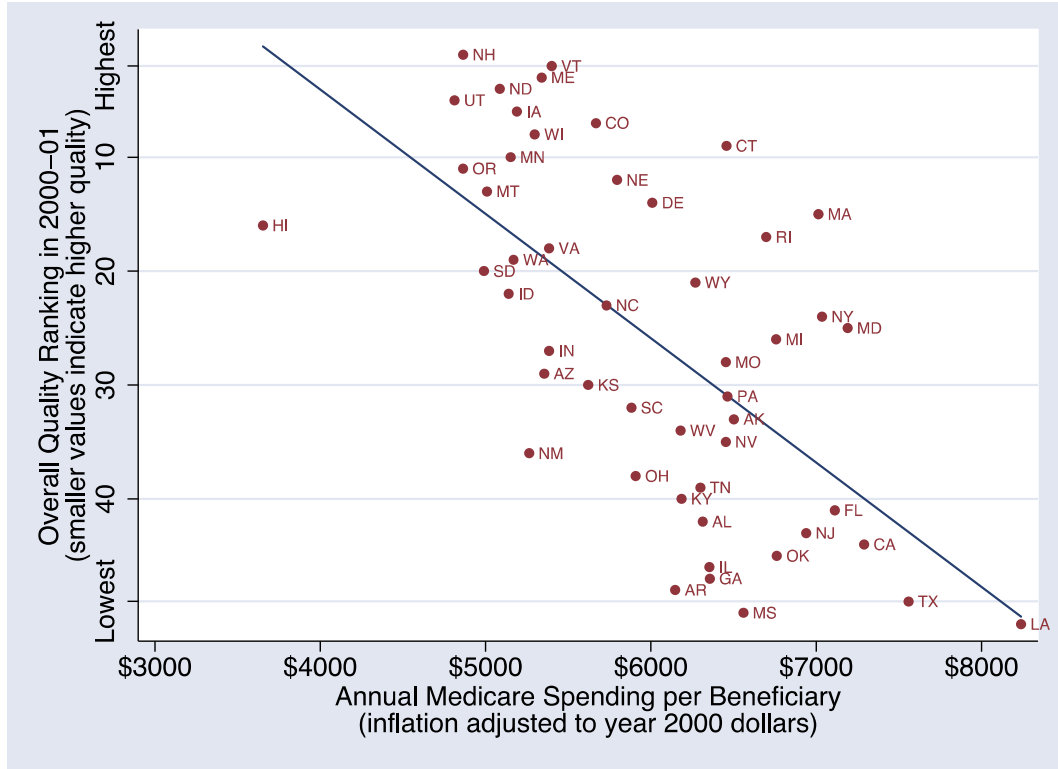
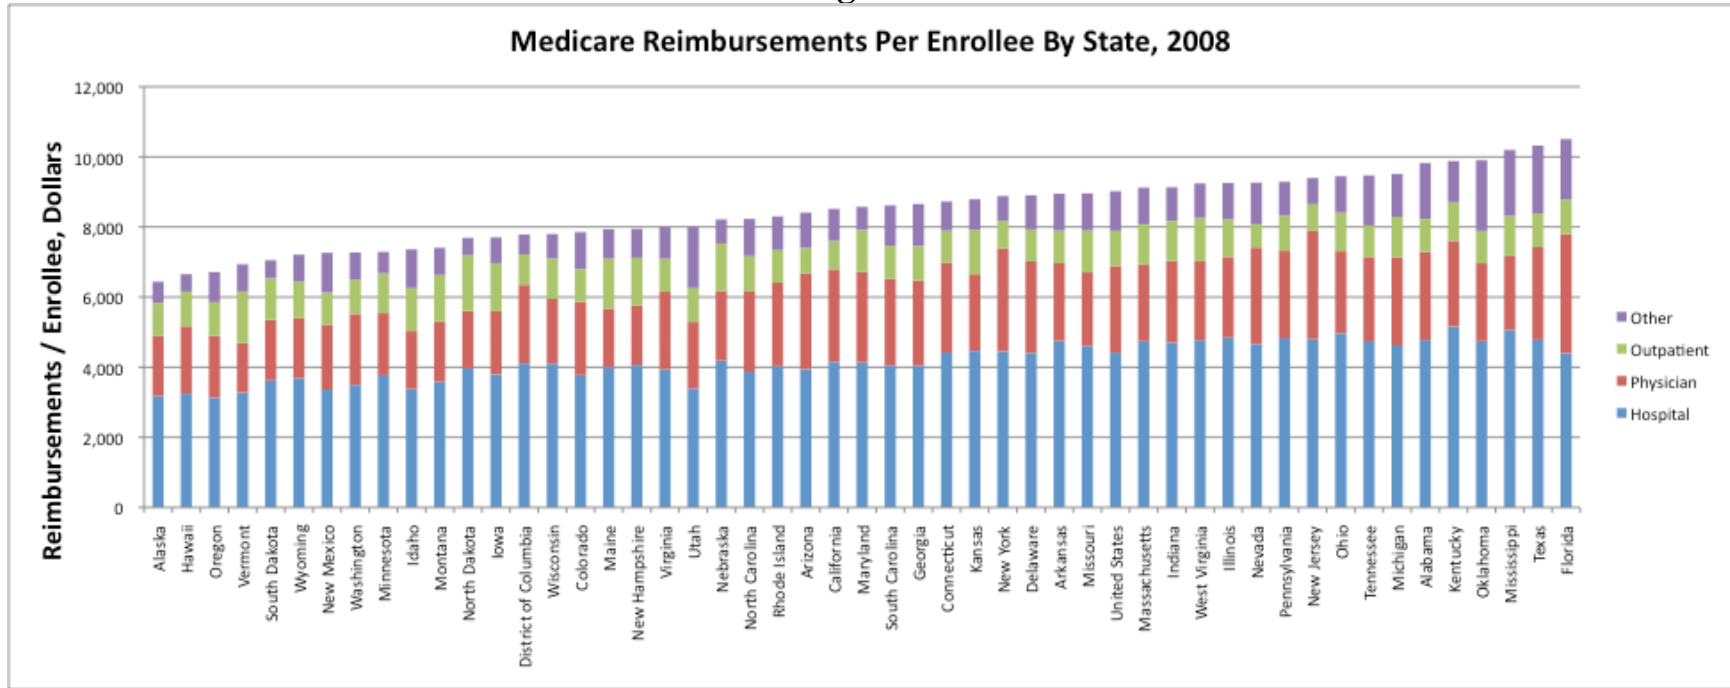


Figure 6: Medicare Spending and the Quality of Care



Source: Baicker and Chandra (Health Affairs, 2004), based on state-level use of care such as aspirin post-heart-attack, blood level monitoring for diabetics, flu vaccination, etc. for Medicare fee-for-service beneficiaries.

Figure 7



Source: Dartmouth Atlas of Health Care (Wennberg, Cooper, and Dartmouth Atlas of Health Care Working Group 2011)

References

- Bach, Peter B. 2010. A Map to Bad Policy - Hospital Efficiency Measures in the Dartmouth Atlas. *New England Journal of Medicine* 362:569-574.
- Baicker, K., and A. Chandra. 2004. Medicare spending, the physician workforce, and beneficiaries' quality of care. *Health Affairs* 23 (3):W4184-W4197.
- . 2008. Myths And Misconceptions About US Health Insurance. *Health Affairs* 27 (6):W533-W543.
- . 2010. Uncomfortable Arithmetic - Whom to Cover versus What to Cover. *New England Journal of Medicine* 362 (2):95-97.
- Baicker, K., E. S. Fisher, and A. Chandra. 2007. Malpractice liability costs and the practice of medicine in the Medicare program. *Health Aff (Millwood)* 26 (3):841-52.
- Baicker, Katherine, and Amitabh Chandra. 2006. The Labor Market Effects of Rising Health Insurance Premiums. *Journal of Labor Economics* 24 (3):609-34.
- Baicker, Katherine, and Dana Goldman. 2011. Patient Cost-Sharing and Healthcare Spending Growth. *Journal of Economic Perspectives* 25 (2):47-68.
- Baicker, Katherine, and Helen Levy. 2008. Employer Health Insurance Mandates and the Risk of Unemployment. *Risk Management and Insurance Review* 11 (1):109-132.
- Baicker, Katherine, and Jonathan S. Skinner. 2011. Health Care Spending Growth and the Future of U.S. Tax Rates. *Tax Policy and the Economy*.
- Baker, L, and K Cortis. 1996. HMO Penetration and the Cost of Health Care: Market Discipline or Market Segmentation? *American Economic Review* 86 (2):389-394.
- Baker, Laurence C. 1999. Association of Managed Care Market Share and Health Expenditures for Fee-for-Service Medicare Patients. *Journal of the American Medical Association* 281 (5):432-437.
- Becker, Gary, Kevin Murphy, and Tomas Philipson. 2007. The Value of Life Near its End and Terminal Care. *NBER Working Paper*.
- Boards of Trustees of the Federal Hospital Insurance and Federal Supplemental Medical Insurance Trust Funds, The. 2011. *Annual Report Of The Boards Of Trustees Of The Federal Hospital Insurance And Federal Supplementary Medical Insurance Trust Funds*. Washington, DC: Government Printing Office.
- Brenner, David J., and Eric J. Hall. 2007. Computed Tomography - An Increasing Source of Radiation Exposure. *New England Journal of Medicine* 357:2277-2284.
- Buntin, Melinda, and David Cutler. 2009. The Two Trillion Dollar Solution: Saving Money by Modernizing the Health Care System. Center for American Progress.
- Cebul, Randal, James Rebitzer, Lowell Taylor, and Mark Votruba. 2007. Organizational Fragmentation and Care Quality in the US Health Care System. *Journal of Economic Perspectives* 22 (4):93-113.
- Centers for Medicare and Medicaid Services. 2010. National Health Expenditures Accounts: Definitions, Sources, and Methods, 2010.
- Chandra, A., J. Gruber, and R. McKnight. 2010. Patient Cost-Sharing and Hospitalization Offsets in the Elderly. *American Economic Review* 100 (1):193-213.
- Chandra, Amitabh, and Jonathan S Skinner. 2011. Technology Growth and Expenditure Growth in Health Care. *National Bureau of Economic Research Working Paper* 16953.

- Chandra, Amitabh, and Douglas O. Staiger. 2007. Productivity Spillovers in Health Care: Evidence from Treatments of Heart Attacks. *Journal of Political Economy* 115 (1):103-140.
- Chernew, M. E., K. Baicker, and J. Hsu. 2010. The Specter of Financial Armageddon -- Health Care and Federal Debt in the United States. *New England Journal of Medicine* 362 (13):1166-1168.
- Congressional Budget Office. 2007. Financing Projected Spending in the Long Run: Letter to Senator Gregg from Peter Orszag.
- . 2011. The Long-Term Budget Outlook. Washington, DC.
- Currie, Janet, and Brigitte Madrian. 2000. Health, Health Insurance and the Labor Market. In *Handbook of Labor Economics*, edited by O. Ashenfelter and D. Card. Amsterdam: Elsevier Science.
- Cutler, D. M., A. B. Rosen, and S. Vijan. 2006. The value of medical spending in the United States, 1960-2000. *N Engl J Med* 355 (9):920-7.
- Cutler, David. 2010. How Health Care Reform Must Bend the Cost Curve. *Health Affairs* 29 (6):1131-1135.
- Cutler, David, Angus Deaton, and Adriana Lleras-Muney. 2006. The Determinants Of Mortality. *Journal of Economic Perspectives* 20 (3):97-120.
- Cutler, David, and Dan Ly. 2011. The (Paper)Work of Medicine: Understanding International Medical Costs. *Journal of Economic Perspectives* 25 (2):3-25.
- Cutler, David, and Brigitte Madrian. 1998. Labor Market Responses to Rising Health Insurance Costs: Evidence on Hours Worked. *RAND Journal of Economics*, 29 (3):509-530.
- Dafny, Leemore. 2008. Are Health Insurance Markets Competitive? *National Bureau of Economic Research Working Paper* 14572.
- Dafny, Leemore, Mark Duggan, and Subramaniam Ramanarayanan. 2009. Paying a Premium on Your Premium? Consolidation in the U.S. Health Insurance Industry. *National Bureau of Economic Research Working Paper* 15434.
- DeNavas-Walt, Carmen, Bernadette D. Proctor, and Jessica C. Smith. 2010. Income, Poverty, and Health Insurance Coverage in the United States: 2009. edited by U. S. C. Bureau. Washington, DC: U.S. Government Printing Office.
- Engen, and Jonathan S. Skinner. 1996. Taxation and Economic Growth. *National Tax Journal*:617-42.
- Feldstein, Martin. 1973. The Welfare Loss of Excess Health Insurance. *Journal of Political Economy* 81 (2):251-80.
- . 1999. Tax Avoidance and the Deadweight Loss of the Income Tax. *Review of Economics and Statistics* 81 (4):674-680.
- . 2006. The Effect of Taxes on Efficiency and Growth. Dissertation/Thesis, National Bureau of Economic Research, Inc, NBER Working Papers.
- Ferguson, John H., Michael Dubinsky, and Peter Kirsch. 1993. Court-Ordered Reimbursement for Unproven Medical Technology. *Jama* 269 (16):2116-2121.
- Finkelstein, A., S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. Newhouse, H. Allen, and K. Baicker. 2011. The Oregon Health Insurance Experiment: Evidence from the First Year. *National Bureau of Economic Research Working Paper* 17190.
- Finkelstein, Amy. 2007. The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare. *Quarterly Journal of Economics* 122 (1):1-37.

- Fisher, E. S., D. E. Wennberg, T. A. Stukel, and D. J. Gottlieb. 2004. Variations in the longitudinal efficiency of academic medical centers. *Health Affairs (Millwood)* Suppl Web Exclusive:VAR19-32.
- Fisher, Elliott S., David E. Wennberg, Therese A. Stukel, Daniel J. Gottlieb, F. Lee Lucas, and E.L. Pinder. 2003. The Implications of Regional Variation in Medicare Spending. Part 1: the Content, Quality and Accessibility of Care. *Annals of Internal Medicine* 138 (4):273-87.
- . 2003. The Implications of Regional Variation in Medicare Spending. Part 2: Health Outcomes and Satisfaction with Care. *Annals of Internal Medicine* 138 (4):288-98.
- Fogel, Robert. 2008. Forecasting the Cost of U.S. Health Care in 2040. *NBER Working Paper* 14361.
- Ford, Earl S. , Umed A. Ajani, Janet B. Croft, Julia A. Critchley, Darwin R. Labarthe, Thomas E. Kottke, Wayne H. Giles, and Simon Capewell. 2007. Explaining the Decrease in U.S. Deaths from Coronary Disease, 1980–2000. *N Eng J Med* 356:2388-2398.
- Frank, Richard G., and R.P. Zeckhauser. 2007. Custom Made Versus Ready to Wear Treatments: Behavioral Propensities in Physicians Choices. *Journal of Health Economics* 26:1101-1127.
- Fuhrmans, Vanessa. 2007. Withdrawal Treatment: A Novel Plan Helps Hospital Wean itself Off Pricey Tests. *The Wall Street Journal*.
- Fullerton, Don, and Yolanda Kodrzycki Henderson. 1989. The Marginal Excess Burden of Different Capital Tax Instruments. *Review of Economics and Statistics* 71 (3):435-442.
- Gahvari, Firouz. 2006. On the Marginal Cost of Public Funds and the Optimal Provision of Public Goods. *Journal of Public Economics* 90 (6-7):1251-1262.
- Garber, Alan M., and Jonathan S. Skinner. 2008. Is American Health Care Uniquely Inefficient? *Journal of Economic Perspectives* 22 (4):27-50.
- Gawande, Atul. 2009. The Cost Conundrum. *The New Yorker*.
- Getzen, T.E. . 1992. Population Aging and the Growth of health Expenditures. *Journal of Gerontology* 17:S98-104.
- Glied, S., and J. Zivin. 2002. How Do Doctors Behave When Some (But Not All) of Their Patients Are In Managed Care? *Journal of Health Economics* 21:337-353.
- Hall, Robert E., and Charles I. Jones. 2007. The Value of Life and the Rise in Health Spending. *Quarterly Journal of Economics* 122 (1):39-72.
- Ho, Vivian, Meei-Hsiang Ku-Goto, and James G Jollis. 2009. Certificate of Need (CON) for Cardiac Care: Controversy over the Contributions of CON. *Health Serv Res* 44 (2):483-500.
- Kaiser Family Foundation, and Health Research and Education Trust. 2008. Employer Health Benefits 2008 Survey. Menlo Park, CA.
- U.S. Senate. 2008. Senate Finance Committee. *Chief of Staff, Joint Committee on Taxation*. July 30, 2008.
- Kogan, R , K Cox, and J. Horney. 2008. The long-term fiscal outlook is bleak: restoring fiscal sustainability will require major changes to programs, revenues, and the nation’s health care system. *Center on Budget and Policy Priorities*.
- Madrian, Brigitte. 1994. Employment-Based Health Insurance and Job Mobility: Is There Evidence of Job-Lock? *Quarterly Journal of Economics* 109 (1):27-54.
- Marchione, Marilyn. *\$93,000 Cancer Drug: How Much is a Life Worth?* 2010 [cited 2010]. Available from <http://www.cbsnews.com/stories/2010/09/27/health/main6905221.shtml>.

- McKnight, Robin Lynn. 2004. Home Care Reimbursement, Long-Term Care Utilization, and Health Outcomes. *Journal of Public Economics* 90 (1-2):293-323.
- Mello, Michelle, Amitabh Chandra, Atul Gawande, and David Studdert. 2010. National Costs of the Medical Liability System. *Health Affairs* 29 (9):1569-1577.
- Murphy, Kevin M., and Robert H. Topel. 2006. The Value of Health and Longevity. *Journal of Political Economy* 114 (5):871-904.
- Newhouse, Joseph P., Elizabeth M. Sloss, Willard Manning, and Emmett B. Keeler. 1993. Risk Adjustment for a Children's Capitation Rate. *Health Care Financing Review* 15 (1):39-54.
- Newhouse, Joseph P., and The Insurance Experiment Group. 1993. *Free for all? Lessons from the RAND Health Insurance Experiment*. A RAND Study. Cambridge and London: Harvard University Press.
- Ong, M.K., C. M. Mangione, P. S. Romano, Q. Zhou, A. Auerbach, A. Chun, B. N. Davidson, T. Ganiats, S Greenfield, M. A. Gropper, S. Malik, J. Rosenthal, and J. J. Escarce. 2009. Looking Forward, Looking Back: Assessing Variations in Hospital Resource Use and Outcomes for Elderly Patients with Heart Failure. *Circulation Cardiovascular Quality and Outcomes* 2:548-557.
- Pauly, M. V., and F. E. Blavin. 2008. Moral hazard in insurance, value-based cost sharing, and the benefits of blissful ignorance. *J Health Econ* 27 (6):1407-17.
- Peterson, Chris L., and Rachel Burton. 2007. U.S. Health Care Spending: Comparison with Other OECD Countries. edited by C. R. Service. Washington DC.
- Romley, John A., Anupam B. Jena, and Dana P. Goldman. 2011. Hospital Spending and Inpatient Mortality: Evidence From California. *Annals of Internal Medicine* 154:160-167.
- Silber, Jeffrey H., Robert Kaestner, Orit Even-Shoshan, Yanli Wang, and Laura J. Bressler. 2010. Aggressive Treatment Style and Surgical Outcomes. *Health Serv Res* 25 (6):1872-1892.
- Skinner, J. S., D. O. Staiger, and E. S. Fisher. 2006. Is technological change in medicine always worth it? The case of acute myocardial infarction. *Health Aff (Millwood)* 25 (2):w34-47.
- Skinner, Jonathan, and Douglas Staiger. 2009. Technology Diffusion and Productivity Growth in Health Care. *NBER Working Paper 14865*.
- Summers, Lawrence. 1989. Some Simple Economics of Mandated Benefits. *American Economic Review* 79 (2):177-183.
- Weinstein, Milton C., and Jonathan S. Skinner. 2010. Comparative Effectiveness and Health Care Spending — Implications for Reform. *N Engl J Med* 362:460-465.
- Weisbrod, Burton A. 1991. The Health Care Quadrilemma: An Essay on Technological Change, Insurance, Quality of Care, and Cost Containment. *Journal of Economic Literature* 29 (2):523-552.
- Wennberg, J.E., M.M. Cooper, and Dartmouth Atlas of Health Care Working Group. 2011. *The Dartmouth Atlas of Health Care 2011*. Chicago: American Hospital Publishing.
- Yasaitis, Laura, Elliott S. Fisher, Jonathan S. Skinner, and Amitabh Chandra. 2009. Hospital Quality and Intensity of Spending: Is There an Association? *Health Affairs* 28 (4):566-572.