Income and Health Spending: Evidence from Oil Price Shocks*

Daron Acemoglu
MIT and NBER

Amy Finkelstein
MIT and NBER

Matthew J. Notowidigdo
MIT

June 2, 2009

Abstract

Health expenditures as a share of GDP have more than tripled over the last half century. A common conjecture is that this is primarily a consequence of rising real per capita income, which more than doubled over the same period. We investigate this hypothesis empirically by instrumenting for local area income with time-series variation in global oil prices between 1970 and 1990 interacted with cross-sectional variation in the oil reserves across different areas of the Southern United States. This strategy enables us to capture both the partial equilibrium and the local general equilibrium effects of an increase in income on health expenditures. Our central estimate is an income elasticity of 0.7, with an elasticity of 1.1 as the upper end of the 95 percent confidence interval. Point estimates from alternative specifications fall on both sides of our central estimate, but are almost always less than 1. We also present evidence suggesting that there are unlikely to be substantial national or global general equilibrium effects of rising income on health spending, for example through induced innovation. Our overall reading of the evidence is that rising income is unlikely to be a major driver of the rising health share of GDP.

Keywords: health care, income, technology.
JEL Classification: H51, I1.

*We are grateful to Amitabh Chandra, Bob Hall, Guy Michaels, Joe Newhouse, and participants at the NBER Health Care meetings for helpful comments, to James Wang for outstanding research assistance, and to the National Institute of Aging (NIA grant numbers P30-AG012810 and T32-AG000186) for financial support.

1 Introduction

The dramatic rise in health care expenditures is one of the notable economic trends of the postwar era. As seen in Figure 1, health care expenditure as a share of GDP in the United States has more than tripled over the last half century, from 5 percent in 1960 to 16 percent in 2005 (CMS, 2006). A common conjecture is that the rise in the share of income spent on health care expenditures is a direct, or at least a natural, consequence of the secular increase in living standards—because health care is a "luxury good".[1] The *Economist* magazine stated this as a "conventional wisdom" in 1993, writing:

> "As with luxury goods, health spending tends to rise disproportionately as countries become richer." (quoted in Blomqvist and Carter, 1997, p. 27).

This view has recently been forcefully articulated by Hall and Jones (2007). They argue that the optimal share of spending on health increases as incomes rise, since spending money on life extension allows individuals to escape diminishing marginal utility of consumption within a period. The Hall-Jones view also receives indirect support from the very high estimates of the value of life and value of health provided by Nordhaus (2003) and Murphy and Topel (2003, 2006). The fact that most other OECD countries have also experienced substantial growth in their health sector over the last half century (OECD, 2004) also makes the secular rise in incomes a natural candidate to explain the rise in the health share of GDP in the United States.

Understanding the extent to which the rise in the health share of GDP is a direct consequence of the rise in living standards is important for several reasons. First, it enables a proper accounting of the notable growth in the US (and OECD) health care sector over the last half century. Second, it is necessary for forecasting how health care spending is likely to evolve in coming years. Finally, it is a crucial first step towards an assessment of the optimality of the growth of the health care sector. In particular, if health spending is strongly increasing in income, so that rising income can explain most or all of the rising health share, it would be more likely that the increasing share of GDP allocated to health is socially optimal.[2]

[1] Throughout we use the term "luxury good" to designate an empirical income elasticity greater than one (and similarly "necessity" refers to an elasticity less than one). This responsiveness to income may result from preferences, policy or other factors.

[2] Of course a large role for income would only be suggestive, not dispositive. A systematic analysis of social optimality would also have to consider potential externalities in health provision and in health R&D, as well as informational and institutional constraints in the health care market.

The relationship between income and health spending is the subject of a voluminous empirical literature. Remarkably, however, virtually all existing estimates are based on simple correlations of income and health care spending, across individuals, across countries, or over time. These correlations are consistent with income elasticities ranging from close to zero to substantially above one.[3] In light of the paucity of existing evidence, Hall and Jones (2007) conclude their paper by stating that "Our model makes the strong prediction that if one looks hard enough and carefully enough, one ought to be able to see income effects [with elasticities above 1] in the micro data. Future empirical work will be needed to judge this prediction."

Our objective is to provide "causal" estimates of the effect of income on aggregate health spending. There are (at least) two important challenges in this exercise. The first is that income and health covary at the individual or regional level for a variety of reasons. Therefore, simple correlations are unlikely to reveal the causal effect of income on health spending.

A second challenge is that an investigation of the role that rising income plays in the growth of the health care sector requires incorporating the general equilibrium effects of income on health spending. Partial and general equilibrium income elasticities may differ for a variety of reasons. For example, the general equilibrium effect of rising income may be larger than the partial equilibrium effect if an increase in the demand for health care from a community (a "general equilibrium change") prompts changes in medical practices, including the adoption (and possibly development) of new technologies.[4] Alternatively, if the supply of health care is less than perfectly elastic and the price elasticity of demand for health care is greater than one, the responsiveness of health care expenditures to an increase in income may be lower in general equilibrium than in partial equilibrium. In addition, changes in income may also affect health care policy through a variety of political economy channels, either magnifying or curtailing the direct effect of income on health expenditures. Many of the potential general equilibrium effects are "local" in the sense that they result from changes in incomes in a particular region or local economy. These effects can be detected by looking at the response of health spending to income in the local economy. In addition, there may also exist national or even global general equilibrium effects, which will be harder to detect empirically.

[3] OECD (2006) provides a recent survey of the large empirical literature on the correlation between income and health spending (see particularly Annex 2B). The cross-sectional relationship across individuals between income and health spending tends to be small or negative (e.g., Newhouse and Phelps 1976). In contrast, cross-country analysis tends to suggest income elasticities greater than 1 (e.g., Newhouse 1977, Gerdtham and Jonsson 2000), as do time-series analyses of the relationship between income growth and growth in health spending for individual countries (e.g., Fogel 1999).

[4] Finkelstein (2007), for example, argues that, for such reasons, the general equilibrium effect of health insurance coverage on health spending is larger than the partial equilibrium effect.

We confront both of these challenges. By exploiting potentially exogenous variation in local area incomes, we attempt to estimate causal elasticities that incorporate local general equilibrium effects. On the basis of our estimates and additional evidence, we also argue below that national or global general equilibrium income effects are unlikely to be significant in this instance.

Our strategy is to exploit the time-series variation in global oil prices between 1970 and 1990, which impacted incomes differentially across different parts of the (Southern) United States that vary in the oil intensity of the local economy. In our baseline specification we approximate local economies by economics sub-regions (ESRs), which consist of groups of counties within a state that have strong economic ties. We focus on the South of the United States to increase the comparability of the ESRs, in particular to minimize the likelihood of differential trends in health care expenditure driven by other factors. Our empirical strategy exploits the interaction between global oil prices and ESR-level importance of oil in the economy as an instrument for income. Our main proxy for the importance of oil is the size of pre-existing oil reserves in an ESR. The identifying assumption is that the interaction between global oil price changes and local oil reserves should have no effect on changes in the demand for health care, except through income. We provide several pieces of evidence that are supportive of the validity of this identifying assumption. Using this instrumental-variable strategy we estimate an elasticity of ESR-level hospital spending with respect to ESR-level income of 0.72 (standard error = 0.21). Point estimates of the income elasticity from a wide range of alternative specifications fall on both sides of our baseline estimate, but are almost always less than 1.

Because our instrument impacts incomes at the ESR level (rather than individual income), our estimates correspond to local general equilibrium effects of income changes, but will not capture any global or national general equilibrium effects.[5] Of particular concern is that if the growth of the health care market resulting from the rise in global incomes induced more innovation, our estimates may not incorporate the implications of these induced innovations on health expenditures. Our analysis suggests that significantly larger elasticities resulting from these induced innovation general equilibrium effects are unlikely for two reasons. First, the same induced innovation effects working at the national or global level should manifest themselves as increased technology adoption or entry of new hospitals at the local (ESR) level. However, we find no statistically or substantively significant effects of local income on hospital entry or on various measures of technology adoption at the ESR level. In this light, a significant

[5] We also present results at the state rather than ESR level. This reduces our cross-sectional variation in oil intensity but alllows us to capture general equilibrium effects at a higher level of geographic aggregation than the ESR. The results are similar.