Does Medicare Benefit the Poor? New Answers to an Old Question*

Jay Bhattacharya†
Stanford University and NBER

Darius Lakdawalla‡
RAND and NBER

May 5, 2003

Abstract

Previous research has found that Medicare benefits flow primarily to the most economically advantaged groups and that the financial returns to Medicare are consequently higher for the rich than for the poor. Taking a different approach, we find very different results. According to the Medicare Current Beneficiary Survey, the poorest groups receive the most benefits at any given age. In fact, the advantage of the poor in benefit receipt is so great that it easily overcomes their higher death rates. This leads to the result that the financial returns to Medicare are actually much higher for poorer groups in the population and that Medicare is a highly progressive public program. These new results appear to owe themselves to our measurement of socioeconomic status at the individual level, in contrast to the aggregated measures used by previous research.

*We are grateful to Henry Aaron, Melinda Beeuwkes Buntin, David Cutler, Mark Duggan, Alan Garber, Dana Goldman, Frank Lichtenberg, Michael Schoenbaum, and seminar participants at the Spring 2002 NBER Health Care Program Meeting, and the 2003 American Economic Association Meetings for numerous helpful comments and suggestions. We are especially grateful to Jonathan Skinner for numerous discussions and suggestions. Melinda Beeuwkes Buntin kindly shared her data on Medicare price deflators.

†Stanford University, Center for Primary Care and Outcomes Research, 117 Encina Commons, Stanford, CA 94305-6019. jay@stanford.edu

‡RAND, 1700 Main Street, Santa Monica, CA 90407. darius@rand.org

1 Introduction

Medicare is one of the most significant public entitlement programs in the United States, and certainly the most important program in health. In 1998, Medicare benefit payments alone accounted for 13% of the Federal budget and 2.5% of GDP (US Department of Health and Human Services, 1999). The sheer size of Medicare makes it important for us to understand and quantify its value. In particular, we need to know about the distribution of Medicare benefits and costs, whether Medicare actually benefits anyone, and if so whom.

Though Medicare's primary financing mechanism involves substantial intergenerational transfers of income,[1] whether Medicare has actually benefited particular groups in a cohort depends also on how it redistributes money within generations. Medicare may have improved the lot of the average person in today's elderly cohorts, but not necessarily the lot of the average poor person or the average rich person. It is clear that Medicare taxation transfers resources from the rich to the poor. However, a great deal of previous research has argued that Medicare's benefit structure undoes this by transferring resources back to the rich (Long and Settle, 1984; Gornick et al., 1996; McClellan and Skinner, 1997). This body of research finds that the poor use fewer Medicare resources at any given age, and that their earlier mortality further deprives them of Medicare benefits.[2]

In this paper, we present some new evidence on socioeconomic status and Medicare benefits, and we reach very different conclusions. We find that Medicare spends far more

[1]The two primary mechanisms by which Medicare effects intergenerational transfers are variations in cohort size and growth in medical care expenditures over time. Since the introduction of Medicare, expenditures per beneficiary have grown nearly 4% per annum, as documented in footnote 11.

[2]This research mirrors analogous literature on the progressivity of Social Security. The early literature on this topic (Burkhauser and Warlick, 1981; Hurd and Shoven, 1985; Boskin and Puffert, 1988; Duggan et al., 1993) ignores the lower mortality rates faced by members of disadvantaged groups. Shoven et al. (1987) find that the progressivity of Social Security is considerably flattened when the differential mortality of smokers is taken into account. Similarly, Garrett (1995) finds that differences in mortality between the poor and rich eliminates the "progressive spread in returns" to Social Security across income categories.

on the poor than on the rich at any given age. Reversing the direction of the relationship between socioeconomic status and Medicare benefits has a huge impact on the evaluation of Medicare as a welfare program. In particular, we calculate that the net actuarial value of the Part A benefits received by the 1931-41 birth cohort was much higher for the poor than for the rich. While Medicare is actuarially unfair for college graduates, high school dropouts almost double their money. This is in contrast to previous research, which has found that the actuarial value of Medicare is much higher for the rich, or that Medicare results in a lifetime financial transfer from poor to rich (McClellan and Skinner, 1997). Our finding becomes even more striking in light of the previous research that finds the financial returns to Medicare understate the relative welfare benefit to the poor. Accounting for the non-financial benefits of Medicare would likely increase the extent of progressivity.

Unlike much of the previous literature, we use an individual's own educational attainment as a measure of socioeconomic status. In contrast, other researchers have used average income in an individual's area of residence. We find evidence that aggregation bias may be the reason for the discrepancy between the two methods. In a single data set where less educated individuals consume far more Medicare benefits than the more educated, individuals who live in richer areas still appear to receive more benefits than those who live in poorer areas. This could be because people with high demands for medical care have incentives to move to richer areas, where medical care may be of higher quality. As a result, more benefits may be paid in richer areas, even if richer individuals themselves are not receiving more.

2 A Framework for Measuring Progressivity

To analyze the welfare effects of Medicare for different socioeconomic groups, we first have to identify what we mean by socioeconomically advantaged and disadvantaged. We argue that an individual's own educational attainment is a useful way to identify the disadvantaged. We then present a simple framework for estimating the financial returns to Medicare for each group, and interpret this in the broader context of the total utility value of Medicare.

2.1 Measuring Socioeconomic Status

Conceptually, socioeconomic status ought to be measured using permanent income, which is the closest thing economists have to a definition of socioeconomic status. Ideally, we would like to study how Medicare benefits vary across permanent income levels. Unfortunately though, permanent income is unobserved. Only its correlates are available for analysis.

One often-used correlate is the average income in an individual's zip code or area of residence. This correlate smooths out life-cycle and idiosyncratic fluctuations in income, but it introduces the possibility of aggregation bias (Geronimus et al., 1996). There are at least two issues to consider. First, richer areas may have higher quality medical facilities.[3] People with high demands for medical care thus have incentives to move to such areas, holding their permanent income constant. Even if individual wealth has no effect on health expenditures, health expenditures could be higher in wealthier areas. Second, richer areas will tend to be older, for life-cycle income reasons, even though this may not reflect differences in permanent income. Since older people have higher demands for medical care, this too can create a positive relationship between health expenditures and area wealth. Several previous

[3]For example, Chandra and Skinner (2002) show that areas with a higher percentage of white residents are likely to have higher quality medical facilities.

studies have found a positive relationship between Medicare expenditures and area income (cf, Long and Settle, 1984; McClellan and Skinner, 1997). Is this an artifact of aggregation bias, or do wealthier people actually cost Medicare more?

One way to assess this is to measure SES at the individual level. Unfortunately, an individual's current-period income is a rather poor measure of permanent income, because it is subject to idiosyncratic and life-cycle fluctuation. A more theoretically sound measure is an individual's educational attainment. An individual's permanent income is generated by the returns to human and nonhuman wealth. Human wealth consists of schooling and unobserved ability. A great deal of research in labor economics suggests that schooling is a very good measure of human wealth, and that unobserved ability is not a very significant component of it (cf. Card, 1995). In addition, the vast majority of aggregate wealth in the economy is human wealth (Jorgenson and Fraumeni, 1995). Brown and Weisbenner (2002) find that life-cycle (labor) income is three times more important than bequests and inter vivos transfers, which account for 20 to 25% of aggregate wealth. Moreover, even this amount is concentrated among a relatively small percentage of households. Since schooling is perhaps the best feasible measure of human wealth, and since human wealth makes up the majority of total lifetime wealth, schooling is a reasonable way to measure permanent income and SES at the individual level. As a result, we take an individual's educational attainment as our measure of socioeconomic status, and we calculate the returns to Medicare across different education groups.

2.2 Measuring the Returns to Medicare

In measuring the returns to Medicare, we need to consider both how the financial returns to Medicare should be measured, and how these financial returns relate to the total impact on welfare?

2.2.1 The Financial Returns to Medicare

There are two plausible approaches to calculating the financial returns to Medicare, or more broadly the progressivity of Medicare. Fortunately in our case they yield the same answers. The first calculates the expected net present value of Medicare for each socioeconomic group, and then divides by the expected net present value of lifetime income. This would be analogous to a calculation of income tax incidence, which divide a group's tax bill into its income to arrive at its percent tax incidence. While this would be the clearly preferred method for a static program, such as an income tax system, a lifetime program could also admit an internal rate of return calculation. Such a calculation yields the rate of return one would have to earn on a lifetime annuity in order to obtain the expected net present value of the public program. The theory here is that a lifetime public program functions as an additional financial market instrument and thus ought to be evaluated as such.[4]

To construct the net present value of Medicare, as in the first approach, define B_{it} as the Medicare benefits received by the average individual in group i at age t and define τ_{it} as the Medicare taxes paid by i at t. Finally, define S_{it} as the probability that i survives to age t. If the real risk-free rate of interest is r—commonly estimated at around 3% per annum (Siegel, 1992)—the expected net present value of Medicare transfers to i at age 18 is equal

[4]Further discussion of these issues can be found in the literature on the returns to Social Security (cf, Hurd and Shoven, 1985; Duggan et al., 1993; Garrett, 1995).

to:

$$\sum_{t \geq 18} S_{it} \frac{(B_{it} - \tau_{it})}{(1+r)^{t-18}} \tag{2.1}$$

We find that the expected net present value is actually higher in absolute terms for the poorer groups. There is thus no need to go through the exercise of calculating expected net present lifetime income and computing the "lifetime tax incidence."

The internal rate of return is obtained by rewriting equation (2.1). In particular, it is the scalar ρ that solves the following equation:

$$\sum_{t \geq 18} S_{it} \frac{(B_{it} - \tau_{it})}{(1+\rho)^{t-18}} = 0 \tag{2.2}$$

In words, the internal rate of return on Medicare is the real rate of interest that would have to obtain to set the net present dollar value of Medicare to zero. We show later that the empirical rates of return are higher for poorer groups.

Conceptually, the internal rate of return would tell us everything we would need to know about welfare if a complete private market for old-age medical insurance existed without Medicare. In other words, if people could buy policies when young and receive benefits when old, the private market would price these policies such that their internal rate of return equaled the real rate of interest: this would be the zero profit condition. Therefore, abstracting from the market incompleteness that might generate an insurance value of Medicare, individual i derives a welfare benefit from Medicare if his internal rate of return on it exceeds the real rate of interest.

Estimating the net present value and the internal rate of return of Medicare for different education groups requires that we calculate a survival profile S_{it}, a benefits profile B_{it}, and a tax profile τ_{it} for different socioeconomic groups i. Sections 3, 4, and 5 document our estimates for these three quantities.

We will confine our financial return calculation to the hospital insurance component, or Part A, of Medicare. This represents about two-thirds of Medicare's annual budget. It is easy to identify the taxes that are earmarked for Part A, which is the component of Medicare funded entirely by payroll taxes paid by the young. While we have data on the lifetime payroll tax liabilities of the 1931-41 birth cohort, we do not have similar data on federal income tax liabilities for this or any other elderly cohort. However, we later present some calculations designed to bound the progressivity of Part B. These suggest that it too transfers resources to the poor.

2.2.2 The Welfare Impact of Medicare

Evaluating Medicare's welfare impact is difficult because Medicare pays out not cash, but insurance coverage, which is not easy to value. In the perfect world of complete and efficient markets, the net present value of the benefits paid by an insurance policy measure the policy's full value, at least in the sense of foregone consumption. In this case, the benefit stream is equal to the price of the policy, which is the same as the consumption the individual has given up in exchange for it. However, Medicare is not competitively delivered, nor should we assume that markets for health insurance absent Medicare would be complete. Fortunately, the prior research of McClellan and Skinner (1997) provides us with a means of approaching this problem.

They find that considering only the pure dollar transfers of Medicare *understates* the relative benefit to the poor, who disproportionately benefited from the increased access to insurance provided by Medicare. They argue in particular that prior to the introduction of Medicare, the burden of uninsurance and underinsurance fell most heavily on the poor, and they estimate that Medicare's provision of universal insurance coverage for the elderly disproportionately benefited the poor. Based on this argument, if Medicare's pure financial transfers benefit the poor, its overall transfers of welfare are likely to benefit the poor also. Therefore, our estimates of the financial progressivity of Medicare ought to be taken as a lower bound on the progressivity of Medicare in the sense of overall welfare.

3 Mortality and Socioeconomic Status

Standard life tables tend not to report survival probabilities by education group. To calculate S_{it} for different education groups i, we start with standard Social Security Administration life-tables and adjust them to reflect mortality differences across education groups. Using microdata from the National Mortality Followback Surveys of 1986 and 1993, we calculate the ratio of the group-specific death rate to the total death rate. Applying this ratio to the overall life-table then allows us to compute group-specific survival probabilities.

We construct a 1990 period life table for each sex and education group.[5] The 1990 US Vital Statistics period life table gives us a series $\bar{S}_t$, the average probability of survival to age t. Using data from the 1993 National Mortality Followback Survey (NMFS), we then estimate $\frac{S_{it}}{\bar{S}_t}$ for several age groups. The NMFS contains individual-level data on a sample of decedents from 1992. It is designed to be nationally representative, while oversampling young decedents. Based on interviews with next-of-kin, the NMFS collects demographic information about each decedent, including age, sex, race, education, smoking status, and

[5]Since we are calculating the rate of return for a specific birth cohort, the best thing to do would be to construct a cohort life table, but we know of no source for cohort-specific death rates by education group. Since mortality rates declined more rapidly among the better educated from 1960 onwards (Feldman et al., 1989), it is likely that this strategy will understate the progressivity of Medicare.

Figure 1: Male Survival Curves by Education.

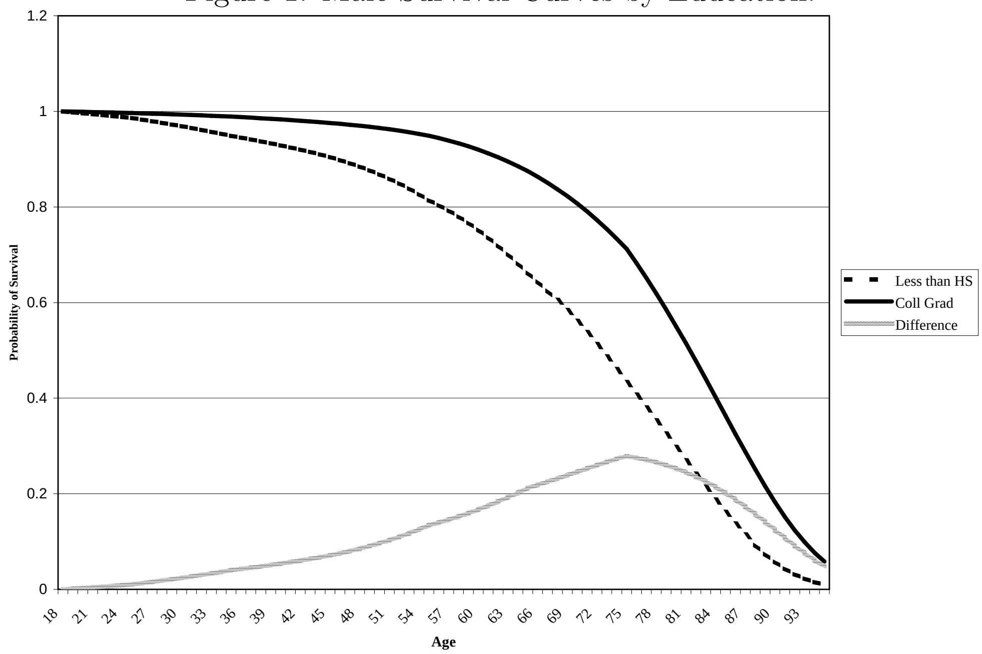

cause of death. Using the weights provided in the NMFS, we are able to estimate the total number of deaths nationwide within each age group, and within each age-education category. To translate the total number of deaths into death rates, we use the National Health Interview Survey to estimate the 1992 population nationwide in each age-education category. The results of these calculations are shown in Table 1. With only a few exceptions, death rates decline uniformly with education group, within an age category. Among very old women, we observe a slight increase in mortality rates between high school dropouts and high school graduates. Among 45-54 year-old men and 55-64 year-old women, we observe mortality rates that are higher for college attendees than high school graduates. Apart from these isolated cases, mortality rates fall with education.

The associated survival curves are graphed in Figures 1 and 2, for men and women. For both men and women, 18 year-old high school dropouts are less likely to reach age 65 than

Table 1: Yearly Deaths per 1000 people, by Age, Sex, and Education, 1985-92.

	Age Group	High Sch Dropouts	High Sch Grads	College Attendees	College Grads	Overall
Males	18-24	2.36	1.98	0.88	0.46	1.60
	25-34	3.83	2.03	1.40	0.67	1.82
	35-44	5.00	3.28	2.61	1.23	2.73
	45-54	9.45	5.50	6.03	2.80	5.62
	55-64	17.47	13.70	12.46	7.43	13.38
	65-74	35.23	33.69	26.13	16.94	30.33
	75+	100.94	95.32	78.02	69.99	92.97
Females	18-24	0.63	0.54	0.36	0.21	0.47
	25-34	1.40	0.67	0.50	0.33	0.65
	35-44	1.98	1.48	0.89	0.88	1.27
	45-54	4.69	3.30	2.17	1.50	3.00
	55-64	10.33	7.58	7.80	5.45	8.13
	65-74	18.65	19.66	15.43	11.82	17.82
	75+	75.88	83.12	73.65	49.05	74.91

Note: Death rates are averages of 1985 and 1992 rates, estimated from the 1986 and 1993 National Mortality Followback Surveys, respectively.

Figure 2: Female Survival Curves by Education.

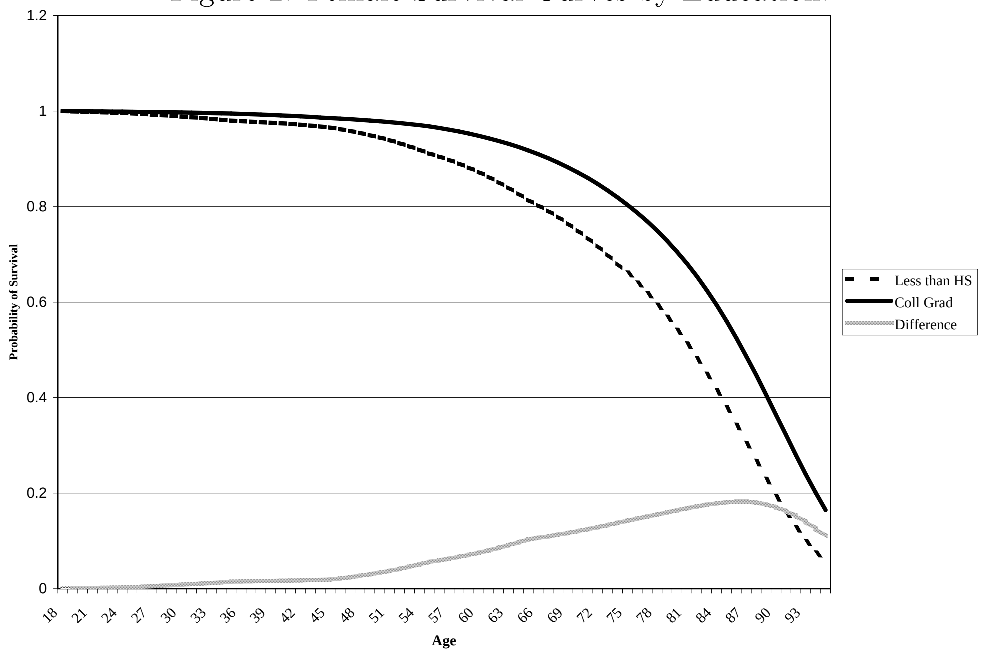

college graduates (or those who will end up as college graduates). However, the difference is twice as large for men as for women. High School dropout males are twenty percentage points less likely to survive to age 65, while females are only about ten percentage points less likely. This is one of the reasons why, from an individual perspective, Medicare is a much better deal for low-skill women than for low-skill men.

4 The Distribution of Medicare Benefits

We calculate age-specific Medicare benefits, B_{it} using the Medicare Current Beneficiary Survey (MCBS). This is a nationally representative, longitudinal, random sample of Medicare beneficiaries, and it includes extensive information about Medicare expenditures, along with demographic information, such as about years of schooling and geographic information.

4.1 Age-specific Medicare Expenditures

The MCBS Cost and Use Files are nationally representative data sets designed to ascertain utilization and expenditures for the Medicare population. They are available every year from 1992 to 1999. The sample frame consists of aged and disabled beneficiaries enrolled in Medicare Part A and/or Part B, although we use only the aged. The oldest-old (85 years of age or over) are oversampled. The MCBS contains demographic data such as age, sex, race, and educational attainment, along with state, county, and zip code of residence. It also contains detailed self-reported information on health, including the prevalence of various conditions, measures of physical limitation in performing daily activities (ADLs) and instrumental activities of daily living (IADLs), and height and weight.

Table 2 presents average real (1997 dollars) per capita Medicare benefits by age group, sex, and educational attainment.[6] The top panel of the table shows the educational gradient in total Medicare benefits, which equals Medicare Parts A and B fee-for-service expenses, plus payments made by Medicare on behalf of its beneficiaries to Medicare HMO's.[7]

In these data, there are consistent negative gradients in education (that is, low SES individuals spent more per capita than high SES individuals). The difference between high school dropouts and college graduates is always at least ten percent (for 75-84 year-old women) and reaches as high as forty-five percent for 65-74 year old men. In addition, there are few instances of increases in per capita benefits across education levels. Per capita benefits rise with education only four out of eighteen times, and two of these times involve

[6] Appendix A describes how expenditure data are collected in the MCBS, and how we identify Part A and B expenditures.

[7] Including the medical expenditures paid by the HMO's themselves would represent double-counting. The actual payment made by Medicare to the HMO represents the public liability. Any difference between these payments and HMO expenditures represent profit or loss for the private firms, not public liability for old-age medical care.

Table 2: Real Per Capita Medicare Benefits by Educational Attainment.

		Females				Males			
		High Sch Dropouts	High Sch Grads	College Attendees	College Grads	High Sch Dropouts	High Sch Grads	College Attendees	College Grads
Total Medicare	65-74	$4,402	$3,422	$3,321	$2,483	$4,821	$3,904	$3,850	$3,480
	75-84	$5,825	$5,294	$5,312	$4,341	$6,126	$5,994	$6,326	$5,805
	85+	$7,321	$6,949	$5,660	$5,860	$7,112	$7,803	$7,228	$7,021
Medicare Part A	65-74	$2,449	$1,731	$1,602	$909	$2,819	$2,004	$1,845	$1,704
	75-84	$3,558	$3,010	$2,969	$2,222	$3,519	$3,460	$3,305	$3,054
	85+	$4,982	$4,623	$3,423	$3,856	$4,643	$5,241	$4,766	$4,285
Medicare Part B	65-74	$1,533	$1,165	$1,096	$1,057	$1,447	$1,312	$1,229	$1,238
	75-84	$1,776	$1,611	$1,599	$1,490	$1,845	$1,775	$2,059	$2,077
	85+	$1,827	$1,709	$1,535	$1,567	$1,744	$1,813	$1,710	$1,956
Medicare HMO	65-74	$420	$526	$623	$517	$556	$587	$776	$538
	75-84	$491	$672	$743	$629	$762	$759	$961	$674
	85+	$513	$617	$701	$438	$725	$749	$752	$781

Source: MCBS, 1992-1999.

Notes: All values are per capita real 1997 dollars.

differences of $120 or less, or about three percent. Most of the negative gradient is driven by variation in Part A, or hospital insurance benefits, but there is a consistent negative gradient in Part B benefits. It is not as large in magnitude, but we will show later that it is likely to be enough to yield progressivity for Part B also. While the patterns are not so clear, there may be a slight positive gradient in Medicare HMO payments, although these are quite small relative to the gradients elsewhere.

Part, though not all, of the negative gradient in Medicare benefits is explained by differences in observed health status. Including self-reported occurrence of diseases and disability in the MCBS erases more than half of the gradient between high school dropouts and college graduates. The rest could be generated by variation in unobserved health, but it could also be related to differences in public and private insurance coverage, or other factors. (Not surprisingly, there is a positive gradient in privately financed medical expenditures, once one controls for health.)

Table 3 demonstrates that high school dropouts enjoyed larger increases in benefits than college graduates during the early 1990s, a result that is consistent with the findings of Lee et al. (1999). However, the middle and bottom panels also show that consistent negative gradients in benefits existed as early as 1992 and widened even further over time. This contrasts sharply with studies that examine gradients in Medicare expenditures by zip code income. These studies, such as Lee et al. (1999), find substantially positive gradients in 1990 that did not turn consistently negative even by 1995.[8]

The negative gradient in benefits across individual education is an important part of our argument, and it accounts for the differences between our results and those of McClellan and

[8] Lee, McClellan and Skinner, for instance, find flat or slightly positive gradients for men aged 65-74 and 75-84.

Table 3: Changes over time in real per capita Medicare benefits by Educational Attainment.

		Females				Males			
		High Sch Dropouts	High Sch Grads	College Attendees	College Grads	High Sch Dropouts	High Sch Grads	College Attendees	College Grads
	1992	$4,743	$3,766	$3,560	$3,093	$4,646	$4,085	$4,142	$4,041
	1993	$4,909	$3,823	$3,822	$3,285	$5,588	$4,166	$4,434	$4,268
	1994	$5,320	$4,104	$4,025	$3,547	$5,694	$4,455	$4,007	$4,104
	1995	$6,142	$4,700	$4,834	$3,267	$5,018	$4,791	$4,704	$3,567
	1996	$5,771	$4,383	$4,531	$3,648	$6,404	$4,531	$5,015	$3,946
	1997	$5,968	$4,253	$4,509	$3,742	$5,586	$5,289	$6,293	$4,645
	1998	$5,329	$4,714	$3,921	$3,676	$6,050	$5,102	$4,697	$4,441
	1999	$5,580	$5,161	$4,440	$3,679	$5,499	$5,318	$4,874	$5,279
1992	65-74	$3,877	$2,950	$2,856	$1,729	$4,002	$3,562	$3,650	$3,645
1992	75-84	$5,013	$4,655	$4,939	$4,626	$5,263	$5,169	$4,668	$5,021
1992	85+	$6,567	$6,699	$3,780	$6,127	$6,131	$5,149	$7,687	$4,547
1999	65-74	$5,049	$4,262	$3,747	$2,990	$4,980	$4,489	$4,093	$4,041
1999	75-84	$5,555	$5,516	$5,003	$4,112	$5,851	$6,094	$5,905	$6,456
1999	85+	$6,713	$7,997	$5,560	$5,558	$6,584	$9,121	$6,021	$10,802

Source: 1992-99 MCBS.

Skinner (1997). The cross-sectional gradient is sufficiently large that lifetime differences in the expected value of Medicare benefits are quite small. However, since richer groups will pay substantially more in expected taxes, the net result is progressivity for Medicare.

4.2 Lifetime Medicare Benefits

We use the data in Table 2 to construct the expected present value of lifetime Part A benefits,[9] but we cannot simply apply the survival profiles estimated earlier to the cross-sectional Part A benefits profiles. The benefits that will be received by the 1931-41 birth cohort—our cohort of interest—will grow over time, past the levels that are currently observed in the MCBS. Therefore, we calculate the expected value of Medicare benefits under various

[9]This includes payments made in both fee-for-service and HMO arrangements. Appendix A explains how we decompose both types of payments into Part A and B components.

Table 4: Expected Net Present Value of Real Part A Medicare Benefits, by Sex and Education.

	Real Interest Rate	Male				Female			
		Less than HS	HS Grad	Coll Attendee	Coll Grad	Less than HS	HS Grad	Coll Attendee	Coll Grad
No Growth	0%	$44,753	$47,066	$49,251	$49,803	$41,100	$42,744	$43,742	$40,306
	1%	$25,337	$26,161	$27,344	$27,499	$22,790	$23,458	$23,973	$21,759
	2%	$14,497	$14,702	$15,353	$15,356	$12,786	$13,027	$13,293	$11,885
	3%	$8,379	$8,350	$8,714	$8,669	$7,256	$7,318	$7,455	$6,566
	4%	$4,890	$4,792	$4,997	$4,946	$4,162	$4,157	$4,226	$3,669
	5%	$2,880	$2,777	$2,895	$2,850	$2,413	$2,387	$2,421	$2,072
4% Ann. Growth	0%	$70,283	$79,572	$83,837	$86,548	$70,642	$76,407	$78,503	$76,569
	1%	$38,944	$43,275	$45,497	$46,734	$38,218	$40,942	$42,028	$40,429
	2%	$21,827	$23,807	$24,984	$25,532	$20,931	$22,205	$22,772	$21,598
	3%	$12,368	$13,245	$13,878	$14,108	$11,601	$12,186	$12,485	$11,672
	4%	$7,084	$7,450	$7,796	$7,883	$6,505	$6,766	$6,924	$6,380
	5%	$4,098	$4,234	$4,426	$4,452	$3,689	$3,799	$3,882	$3,526

assumptions about future real growth in Medicare benefits.

Suppose Medicare benefits are assumed to grow at some rate X. We construct the lifetime path of benefits by first assuming that the real Part A benefit data from 1992-99 approximately represent the benefits the 1931-41 cohort will be receiving exactly at age 65. We then suppose that benefits will be $X\%$ higher at age 66, an additional $X\%$ higher at age 67, and so forth. If B_{it} represents the average observed benefit of group i at age t, we construct the age t benefit as $B_{it} * (1 + \frac{X}{100})^{t-65}$.[10] We explore the impact of real benefit growth that ranges from zero to four percent annually, since the latter figure has been the benefit growth rate that Medicare has experienced since its introduction.[11]

Table 4 documents the results of the lifetime benefit calculation, for various real interest rates and two real benefit growth rates. Lifetime benefits are once again shown in terms of

[10] Specifically, B_{it} is estimated within the following age intervals: 65-69, 70-74, 75-79, 80-84, and 85+. Within each interval, real benefits are assumed to be constant. We group the data within intervals to smooth out estimated benefits, because the data are too sparse to estimate benefits for every single age group reliably.

[11] Data from the Health Care Financing Administration (http://www.hcfa.gov/stats/hstats98/blustat4.htm, downloaded on March 8, 2002) on total Medicare outlays and total Medicare enrollees, shows that per capita benefits grew four percent annually from 1966-2000.

constant 1997 dollars. Adjusting for survival and accounting for benefit growth favors the more educated groups, because of their greater longevity. However, even after accounting for longevity differences, male high school dropouts are only at a slight disadvantage, receiving 14% fewer lifetime benefits than college graduates at a 3% real rate of interest and 4% real rate of benefit growth. Female high school dropouts are just about level with the other groups. In contrast, high school dropouts earn about half as much as college graduates, so (as we will confirm) the gradient in lifetime taxes paid will be substantially steeper. Table 4 turns out to be the reason why our conclusion that Medicare transfers resources to the poor is robust to a variety of different estimation assumptions. The expected net present value of Medicare benefits varies little with education. However, no matter how one assembles the data, tax liabilities will be strongly positively correlated with education.

5 The Lifetime Incidence of Medicare Taxation

The last step in estimating the expected net present value of Medicare is the construction of τ_{it}, expected Medicare taxes paid by group i at time t. We use data on actual Medicare tax rates, and earnings data from the Health and Retirement Study (HRS) to construct the expected lifetime tax liability of the 1931-1941 birth cohort.[12] The HRS is a nationally-representative longitudinal household dataset with detailed demographic and financial data on respondents. The great advantage of the HRS is that it can be linked with the restricted-use[13] Social Security earnings file. Based on Social Security Administration records, the restricted file contains, for every quarter from 1951-1990, respondents' earnings that were

[12]Since we are evaluating the Part A hospital insurance component of Medicare, we restrict ourselves to Medicare payroll taxes, which are the sole source of funds for Part A coverage.

[13]Clearance was received from the Institute for Social Research, at the University of Michigan, to use these data for this project.

subject to payroll taxation. These are entirely from administrative records, rather than retrospective self-reports. In particular, for 9537 of the 13,478 people present in Wave 1 of the HRS, we have the quarterly earnings subjected to Social Security taxation from 1951 to 1991. Over this period of time, the earnings subjected to Medicare taxation was identical. From 1991 to 1999, we have detailed self-reported data on wage and self-employment earnings from the HRS itself. From these two data sources, we construct the payroll tax payments of each individual, from the inception of Medicare until 1999.[14]

There are two important problems to solve in these data. First, even though the males in this cohort have much higher labor force attachment and much higher payroll tax outlays, it would be misleading to allocate all of this to men. If market and home work are shared within a family, so too are market wages and market taxes. Therefore, taxes should be calculated for married couples rather than individual workers. Moreover, the true rate of return on Medicare ought to be calculated at the family level, rather than the individual level, since a husband derives benefit from his wife's Medicare claims, and vice-versa. For the purposes of this calculation, we take the pure life-cycle view that couples share their *lifetime* wealth with each other. Therefore, we think of couples as a "family unit" even before their actual date of marriage, and after their date of divorce. The latter assumption relies on the notion that there is an implicit or explicit (i.e., alimony) sharing rule even after divorce. Since this "pure life-cycle approach" has some important limitations,[15] we consider a very different approach later with data that allow more flexibility on this question. The HRS eases the task of computing family taxes, because it contains earnings histories for most of the couples that were still together at the HRS baseline, but we are left with the task of

[14]Details on the construction of earnings and taxes are presented in Appendix B.

[15]We are forced to ignore remarriage of any kind, and we must impute some data for unobserved spouses, as discussed below and in Appendix B.

Table 5: Expected Net Present Value of Medicare Payroll Tax Liability faced by Families of HRS Cohort Members.

Real	Male				Female			
Interest	High Sch	High Sch	College	College	High Sch	High Sch	College	College
Rate	Dropouts	Grads	Attendees	Grads	Dropouts	Grads	Attendees	Grads
0%	$20,298	$28,676	$32,385	$45,565	$15,406	$24,007	$29,059	$35,777
1%	$14,971	$20,983	$23,485	$32,393	$11,458	$17,743	$21,235	$25,659
2%	$11,165	$15,534	$17,233	$23,298	$8,612	$13,256	$15,695	$18,621
3%	$8,417	$11,633	$12,793	$16,953	$6,540	$10,009	$11,731	$13,672
4%	$6,413	$8,810	$9,607	$12,481	$5,016	$7,635	$8,864	$10,155
5%	$4,937	$6,745	$7,296	$9,296	$3,884	$5,883	$6,770	$7,629

Note: All figures are real 1997 dollars, from the perspective of an 18 year-old in the HRS cohort.

estimating the family's tax burden for some families that dissolved prior to the HRS baseline due to divorce or death.

Second, the HRS does not separate taxable wage income from taxable self-employment income, even though the two income sources were taxed at different rates for much of Medicare's history. We thus need some way of estimating the share of income from self-employment earnings. Both these issues force us to impute portions of data for a piece (around 15-20%) of the HRS sample. In Appendix B, we discuss in detail the imputation procedures used to address these two problems. In Section 6.2, we describe sensitivity analysis designed to assess the impact of the imputation procedure and of the life-cycle sharing assumption.

Based on historical tax rates, we can estimate taxes paid using earnings data from the HRS. All these calculations result in an age-profile of real Medicare income (i.e., income subject to Medicare taxes) for couples, as well as age-profiles of real Medicare taxes paid, by education group. Using our estimated survival curves, we calculated the expected net present value of a family's Medicare tax liabilities across education groups and sex. When we report the tax liabilities of a man (or a woman), we are reporting the liability faced by the family of the average man (or woman) in that category. On average, the families of college graduates can expect to pay about twice as much in Medicare payroll taxes as the families of high school dropouts. This result was quite insensitive to various manipulations of our assumptions about self-employment income or the imputation of spousal income. Not surprisingly, the vast majority of variation across education groups is generated simply by earnings differences.

Notice also that the families of women are expected to have fewer liabilities than the families of men. There are two reasons for this. The first is the timing of Medicare for the HRS cohort. Since women tend to have older spouses, their families' income profiles peak earlier. Therefore, they earned a larger portion of their lifetime income before the introduction of Medicare taxes in 1966. When we calculated what expected tax liabilities would have been if Medicare had been introduced in 1950, about three-quarters of this gap disappeared, although our conclusions for progressivity were entirely unchanged even under this counterfactual. The second reason is the higher earnings of male workers. For unmarried or widowed individuals, men have higher earnings than women. This is compounded by the fact that women are more likely to be widowed; in other words, the size of the "average family" over the life course is somewhat smaller for women than for men.

6 The Value of Medicare

We can now take the expected net present value of tax liabilities and net present value of Medicare benefits to arrive at estimates for the returns to Medicare. After presenting these, we check the robustness of our results by using a different data set, and we present an extension of our analysis to Medicare Part B.

6.1 Estimated Returns

To arrive at final dollar returns from Medicare, we have to reconcile the tax liabilities of families, in Table 5 with the expected medical benefits of individuals, in Table 4. Our strategy is to convert the data on individual medical benefits to family benefits by matching men and women. We keep with the earlier assumption from the HRS calculations that families are formed through marriage and dissolved only at the death of one spouse. Therefore, our task is to impute the Medicare benefits received by an individual as well as his current spouse, or his living ex-spouse.[16]

To impute the average family Medicare benefit for, say, X year-old college-educated males, we use the proportion (in the MCBS) of this population that has a living spouse or ex-spouse, along with the distribution of spousal education for 65 year-old college-educated males in the HRS. The average Medicare family benefit is then equal to the individual's benefit plus the average spousal benefit. The latter term is taken to be the probability of having a living spouse (or ex-spouse) within the age-sex-education cell, multiplied by the weighted average of Medicare benefits for X year-old females, where the weights are given by the distribution of spousal education observed for 65 year-old college-educated males in the HRS. When we alter our life-cycle sharing rule in Section 6.2, we calculate rates of return based only on individual Medicare benefits. The analysis there serves as a check on the importance of the family benefits imputation also.

After converting Table 4 to a family basis, we can compute the expected dollar value (benefits minus costs) from Medicare for the families of people of a specific sex and educational attainment. The net flows of Medicare resources are depicted in Table 6. Given

[16]Since virtually no elderly people get married for the first time, we do not consider the problem of benefits for "potential spouses."

Table 6: Expected Net Present Dollar Flows from Medicare Part A for Families of HRS Cohort Members, by sex and education of the cohort member.

	Real Interest Rate	Male				Female			
		Less than HS	HS Grad	Coll Attendee	Coll Grad	Less than HS	HS Grad	Coll Attendee	Coll Grad
No Growth	0%	$54,514	$50,636	$50,064	$37,290	$43,173	$39,607	$37,250	$27,565
	1%	$28,018	$24,028	$23,262	$14,311	$22,166	$18,509	$16,511	$10,066
	2%	$13,786	$10,271	$9,543	$3,298	$10,885	$7,617	$6,009	$1,736
	3%	$6,205	$3,308	$2,695	-$1,656	$4,877	$2,130	$872	-$1,955
	4%	$2,237	-$76	-$561	-$3,598	$1,734	-$507	-$1,477	-$3,345
	5%	$226	-$1,591	-$1,964	-$4,088	$143	-$1,657	-$2,400	-$3,634
4% Ann Growth	0%	$100,402	$106,015	$108,604	$99,466	$79,961	$82,760	$82,724	$75,333
	1%	$53,002	$53,976	$54,885	$47,796	$42,144	$41,851	$41,144	$35,767
	2%	$27,516	$26,613	$26,782	$21,493	$21,837	$20,360	$19,475	$15,690
	3%	$13,821	$12,307	$12,178	$8,319	$10,938	$9,150	$8,300	$5,689
	4%	$6,500	$4,923	$4,702	$1,919	$5,118	$3,396	$2,657	$880
	5%	$2,634	$1,210	$983	-$1,010	$2,050	$532	-$78	-$1,278

Note: All figures are in real 1997 dollars, from the perspective of an 18 year-old in the HRS cohort.

a real rate of interest at 2% or higher, the net dollar flows uniformly fall with education. Therefore, it will certainly be true that the dollar flows are progressive, in the sense that they replace a greater percentage of income for the poorest groups. Notice also that, on a family basis, Medicare is more valuable to men than for women. This is primarily because men spend a greater proportion of their lives as married; this causes their average family Medicare benefits to be significantly higher, even though they live fewer years than women on average. Of course, this effect largely washes out in the internal rate of return calculation, since men invest more family tax payments into Medicare. Adjusting for this investment, therefore, rates of return are about equal for men and women.

Table 7 displays the internal rates of return associated with these expected net present values. At historical rates of Medicare benefit growth, around 4%, the overall rate of return is between 5.1% and 5.2%. We can do a quick "reality check" by comparing these to the rate of return calculated from a simple overlapping generations model. Suppose cohorts live for two periods: during the first period, they work and pay Medicare taxes, and during the

Table 7: Internal rates of return on Medicare Part A by sex, education group, and rates of growth in Medicare benefits.

	Males					Females				
Ben. Gwth.	HS Dropout	HS Grad	Coll Attendee	Coll Grad	Overall	HS Dropout	HS Grad	Coll Attendee	Coll Grad	Overall
0%	4.8%	3.6%	3.4%	2.3%	3.7%	4.8%	3.4%	3.0%	2.2%	3.6%
1%	5.1%	4.0%	3.8%	2.8%	4.1%	5.1%	3.8%	3.4%	2.6%	4.0%
2%	5.4%	4.4%	4.2%	3.3%	4.4%	5.4%	4.1%	3.7%	3.1%	4.4%
3%	5.8%	4.8%	4.6%	3.7%	4.8%	5.7%	4.5%	4.1%	3.5%	4.7%
4%	6.1%	5.2%	5.0%	4.2%	5.2%	6.1%	4.9%	4.5%	4.0%	5.1%
5%	6.4%	5.5%	5.4%	4.6%	5.6%	6.4%	5.3%	4.9%	4.4%	5.5%

second they receive Medicare benefits. Define n_t as the size of the cohort that is working at time t, define B_t as the Medicare benefits paid at time t and τ_t as the taxes paid at time t. Assuming that Medicare is a strictly pay-as-you-go system, we would have the balanced budget constraint:

$$n_t \tau_t = B_t n_{t-1} \tag{6.1}$$

The rate of return earned on Medicare by the time t cohort is:

$$1 + r_{t-1} = \frac{B_t}{\tau_{t-1}} = (1 + \beta_{t-1})(1 + \pi_{t-1}), \tag{6.2}$$

where β_{t-1} represents the rate of growth in benefits from time $t-1$ to time t, and π_{t-1} represents the rate of growth in population over the same period. Taking logarithms yields the approximation:

$$r_{t-1} \approx \beta_{t-1} + \pi_{t-1} \tag{6.3}$$

Thus, the return on Medicare is approximately equal to growth in per capita benefits plus population growth. From 1966 to 2000, the rate of growth in the 18-65 year-old population was approximately 1.4% annually, while the rate of growth in real per capita Medicare benefits was about 4 percent annually. This roughly corresponds to a 5.4 percent annual return on Medicare. This back of the envelope calculation is reasonably close to our estimated internal rates of return.

6.2 Sensitivity Analysis

Partly as a result of data limitations, the earlier analysis took a life-cycle view of perfect sharing of wealth and taxes between spouses. Using different data, we now take the opposite view, that couples share wealth and taxes as little as possible. In particular, we now assume that unmarried people do *not* share with potential spouses. Divorced people do not share with ex-spouses. And married couples share exactly half of their wealth and taxes. This strategy also obviates the need to impute tax data for an unobserved ex-spouse or future spouse, and the need to impute family-level Medicare benefits from the MCBS.

By using data from the 1966-2000 Current Population Surveys (CPS), we are able to implement this analysis. Unlike the HRS, the CPS contains data on marital status at every point in time. Since we know with certainty whether or not individuals are married, we can apportion total family taxes into one half borne by the husband and another half borne by the wife. The two important drawbacks of the CPS are its use of self-reported rather than administrative wage data,[17] and the need to construct a synthetic 1931-41 birth cohort,[18] rather than the actual cohort present in the HRS.

[17]This is not to imply that administrative data are perfect, but they are likely to be better than self-reported data, because individuals and the government have incentives to correct mistakes in administrative data.

[18]In one respect, the use of the synthetic cohort is valuable, because it helps assess the impact of survivorship bias in the HRS cohort.

Table 8: Expected Net Present Value of Medicare Payroll Tax Liability faced per Person, 1931-41 birth cohort.

Real Interest Rate	Male: Less than HS	Male: HS Grad	Male: Coll Attendee	Male: Coll Grad	Female: Less than HS	Female: HS Grad	Female: Coll Attendee	Female: Coll Grad
0%	$9,609	$14,876	$17,895	$25,217	$7,968	$13,592	$16,955	$23,093
1%	$7,152	$10,991	$13,168	$18,151	$5,962	$10,066	$12,456	$16,802
2%	$5,383	$8,212	$9,800	$13,221	$4,510	$7,536	$9,254	$12,365
3%	$4,097	$6,205	$7,376	$9,745	$3,448	$5,703	$6,951	$9,204
4%	$3,151	$4,739	$5,612	$7,267	$2,664	$4,361	$5,278	$6,927
5%	$2,448	$3,657	$4,315	$5,481	$2,078	$3,368	$4,050	$5,271

Note: All figures are in real 1997 dollars. Calculations are based on data from the 1966-2000 CPS.

From the 1966 through 2000 CPS data,[19] we select every household in which at least one person belongs to the HRS cohort. Within each household, we match each individual to his/her spouse if present and calculate the total Medicare payroll taxes paid by the couple. The tax burden of each couple is then split in half and assigned to each partner. If a spouse is absent but married (i.e., not divorced, deceased, or separated), we impute spousal wage income within single-year age, sex, and education cells.[20] In the CPS, this affects just three to four percent of the total observations on spousal income; in comparison, the rate of imputation is about twenty percent in the HRS.

Table 8 displays the resulting estimates of the per-person lifetime Medicare tax liability. It is not possible to compare these numbers directly with the family taxes paid in Table 5, but it is possible to make some rough comparisons. Approximately, average family income should be about twice as high as individual income. Even if the data were perfect, this would not hold exactly, because not everyone is married, and because spouses are not always identically educated. Nonetheless, this simple rule of thumb seems to work fairly well. Departures from

[19]These surveys yield a profile of Medicare payroll taxes exactly as long as the profile we obtained from the HRS.

[20]For example, if there is a 40 year-old, white high school graduate male (married, divorced, or separated) with an unobserved spouse, we assign to him the spousal income observed for other 40 year-old, white high school graduate males.

Table 9: Expected Net Present Dollar Flows from Medicare Part A for Individuals in the HRS Cohort, by sex and education.

	Real Interest Rate	Male				Female			
		Less than HS	HS Grad	Coll Attendee	Coll Grad	Less than HS	HS Grad	Coll Attendee	Coll Grad
No Growth	0%	$35,144	$32,190	$31,356	$24,586	$33,132	$29,152	$26,787	$17,212
	1%	$18,185	$15,170	$14,176	$9,348	$16,827	$13,392	$11,518	$4,957
	2%	$9,114	$6,489	$5,552	$2,135	$8,276	$5,491	$4,039	-$480
	3%	$4,282	$2,146	$1,338	-$1,076	$3,807	$1,615	$503	-$2,637
	4%	$1,739	$53	-$614	-$2,321	$1,499	-$204	-$1,052	-$3,259
	5%	$431	-$880	-$1,421	-$2,631	$335	-$982	-$1,629	-$3,199
4% Ann Growth	0%	$60,674	$64,696	$65,942	$61,331	$62,674	$62,815	$61,548	$53,475
	1%	$31,792	$32,284	$32,329	$28,582	$32,256	$30,876	$29,572	$23,627
	2%	$16,443	$15,595	$15,184	$12,310	$16,421	$14,669	$13,519	$9,233
	3%	$8,272	$7,041	$6,503	$4,363	$8,153	$6,483	$5,534	$2,468
	4%	$3,933	$2,711	$2,184	$616	$3,842	$2,405	$1,645	-$548
	5%	$1,650	$577	$111	-$1,029	$1,611	$430	-$168	-$1,745

Note: All figures are in real 1997 dollars, from the perspective of an 18 year-old in the HRS cohort.

the rule invariably seem to suggest underreporting of taxes in the CPS, relative to the HRS. Indeed, earnings profiles in the CPS are slightly lower than in the HRS administrative data. This could be the result of self-reporting bias, or of mistakes in the administrative data. Fortunately, this difference does not appear to affect progressivity.

This phenomenon also appears in the expected net present value of Medicare and the associated internal rates of return. Table 9 displays the net present flows of Medicare to individuals, assuming that Medicare taxes are split evenly between spouses. According to the CPS tax data, the value of Medicare is higher for most groups, except perhaps for the least educated males. The values in this table tend to be a bit higher than those in Table 6. Of course, we are measuring slightly different concepts in each table: this table reflects the value per person, while Table 6 measured the value for the average family. The scale in both tables should be roughly similar, however, because Table 9 is based on approximately half the tax payments and approximately half the benefits. The slightly higher values of Medicare are consistent with the lower earnings and tax payments in the CPS data. Notice,

Table 10: CPS-Based Internal rates of return on Medicare Part A by sex, education group, and rates of growth in Medicare benefits.

Medicare Benefit Growth	Males					Females				
	High Sch Dropouts	High Sch Grads	College Attendees	College Grads	Overall	High Sch Dropouts	High Sch Grads	College Attendees	College Grads	Overall
0%	5.6%	4.0%	3.6%	2.6%	3.9%	5.5%	3.8%	3.2%	1.9%	3.9%
1%	5.9%	4.4%	4.0%	3.0%	4.3%	5.8%	4.2%	3.6%	2.3%	4.3%
2%	6.2%	4.8%	4.3%	3.4%	4.6%	6.2%	4.6%	4.1%	2.8%	4.7%
3%	6.5%	5.1%	4.7%	3.9%	5.0%	6.5%	5.0%	4.5%	3.3%	5.1%
4%	6.7%	5.5%	5.1%	4.3%	5.3%	6.8%	5.4%	4.9%	3.7%	5.5%
5%	7.0%	5.8%	5.5%	4.7%	5.7%	7.1%	5.8%	5.3%	4.2%	5.8%

however, that the gradient across education groups remains unchanged. Medicare continues to be more valuable in absolute terms to less educated groups.

Table 10 displays the estimated internal rates of return based on the CPS data. These tend to be about two-tenths of a percentage point higher overall than the HRS-based numbers in Table 7. This seems to be because earnings profiles in the CPS are slightly lower than in the HRS administrative data. This could be the result of self-reporting bias, or of mistakes in the administrative data. These data also tend to generate slightly larger spreads in the internal rates of return across education groups. Regardless, this difference does not affect the qualitative results of interest. We continue to find that the rate of return on Medicare falls for the most educated groups. Finally, it is interesting to note that the CPS figures are closer to our "ballpark" estimate of a 5.4% internal rate of return on Medicare. It is of course hard to draw conclusions from this, since our ballpark estimate is nothing more than its name suggests.

6.3 Accounting for Part B

The previous analysis accounts for the lifetime value of Part A. Unfortunately, we cannot produce similar estimates for Part B, because we do not have lifetime data on federal income tax payments for the HRS cohort. Since Part B of Medicare is funded out of general federal revenues, this limitation makes it impossible for us to estimate the exact rates of return on Medicare Part B. However, we do have enough data to show that, given the negative gradients in its benefit structure, Part B is likely to be progressive as well.

We cannot estimate internal rates of return, but we can produce bounds on the standard tax incidence calculation. The net present value of Medicare for group i is given by Equation 2.1. If this net present value represents a larger share of income for poorer groups, the program is progressive according to the standard tax incidence view of progressivity. We can take this approach to show that, at a minimum, Medicare is progressive, although we cannot quantify the exact extent of the progressivity.

First, we use the MCBS data to calculate directly the expected present value of Medicare Part B benefits, net of Part B premia paid by elderly beneficiaries. We net out premia, because the portion of Part B financed by premia does not represent a return on taxes paid. Our estiamtes suggest that the expected present value of Part B benefits represents a larger share of lifetime income for high school dropouts than college graduates. Moreover, the progressivity of the federal tax system implies (and we will show empirically) that expected present income tax liabilities represent a smaller share of lifetime income for high school dropouts. These two results taken together imply that Part B is progressive under the tax incidence view of progressivity.

Data on Part B benefits are taken directly from the MCBS. Appendix A describes how we identify Part B expenditures. The MCBS also allows us to calculate the actual Part B premia paid by respondents, because it reports the number of months each respondent paid for Part B. We combined these data with *Federal Register* information on the monthly Part B premia charged, from 1991 to 1998 (the years covered by the 1992-99 MCBS surveys).

Table 11: Expected Present Value of Part B Benefits Net of Premia for a Family, 4% Annual Growth.

Real Interest Rate	Male				Female			
	Less than HS	HS Grad	Coll Attendee	Coll Grad	Less than HS	HS Grad	Coll Attendee	Coll Grad
0%	$36,608	$38,146	$43,020	$49,837	$31,239	$30,156	$32,455	$38,770
1%	$20,841	$21,457	$24,035	$27,899	$17,769	$17,001	$18,260	$21,768
2%	$11,982	$12,191	$13,558	$15,773	$10,210	$9,682	$10,377	$12,343
3%	$6,955	$6,994	$7,721	$9,004	$5,925	$5,569	$5,955	$7,067
4%	$4,075	$4,051	$4,438	$5,189	$3,471	$3,234	$3,450	$4,084
5%	$2,409	$2,368	$2,574	$3,019	$2,053	$1,896	$2,017	$2,382

Note: All figures are in real 1997 dollars, from the point of view of an 18 year-old.

We checked these calculations against actual premia paid, which are reported in the 1995-99 MCBS. For every observation, our estimates were within rounding error (to the nearest penny) of the data reported in the MCBS.

Table 11 depicts what the family of the average individual in the given sex-education category can expect to receive from Medicare Part B benefits alone, net of actual premia. These figures assume a 4% rate of benefit growth (premia are assumed to grow at the same rate as benefits), and are discounted to the point of view of an 18 year-old. At a real interest rate of 3%, male college graduates can expect to receive approximately 30% more from Part B than high school dropouts, while female college graduates can expect to receive about 19% more. While these gradients are significant, they are not as large as the gradient in expected lifetime income across these groups.

To gain an appreciation for the gradient in expected lifetime income, we will make some conservative assumptions. First, suppose that mortality rates do not differ across education groups. In reality, differential mortality works to lower the relative lifetime income of the less educated. Ignoring this, we can look at annual gradients in lifetime income without adjusting for survival. Second, consider the gradient in family income for 1975, which exhibits the

Figure 3: Age-Specific Family Income Ratio of College Graduates to High School Dropouts, 1975

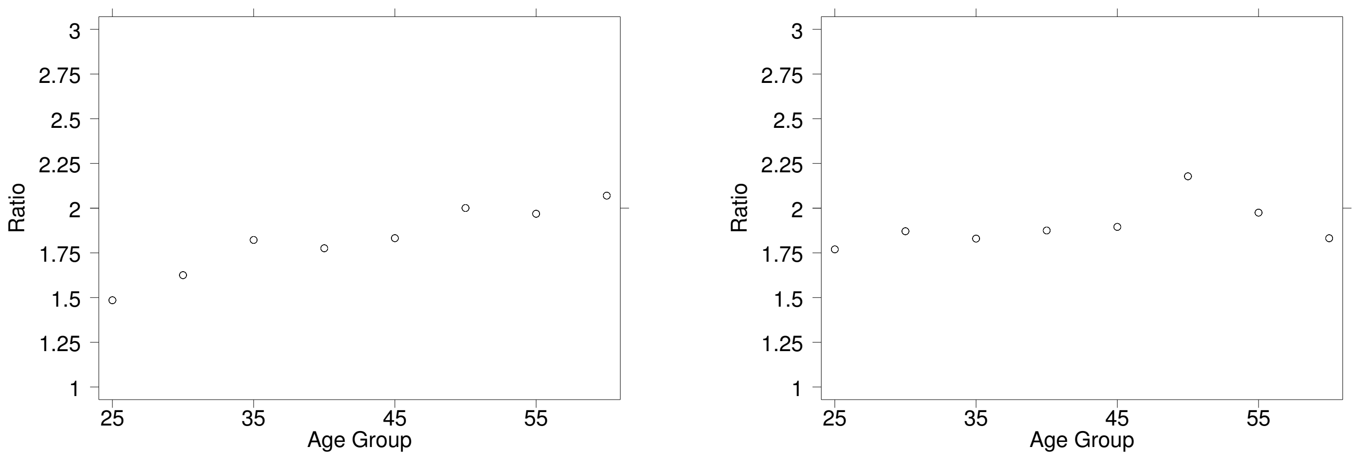

(a) Males

(b) Females

most compressed family income gradients over the entire period 1965-2001.[21] The family income gradients observed in 1975 provide a lower bound on the expected lifetime income gradients. Using the Current Population Survey (CPS), we calculated average family income by education and sex, for 5-year age intervals. Figure 3 depicts the results for the 1975 CPS. The figure shows that the family income ratio is uniformly above 1.5, even for this most compressed of years. This conservative 50% income gradient exceeds the difference in expected Medicare benefits, which is under 30%. Therefore, the expected net present value of benefits represents a larger share of income for poorer groups.

The calculations above show that Part B is progressive on the benefit side. As a result, it must be progressive as a whole, since it is funded by the progressive federal income tax system. The progressivity of the federal income tax system is evident in Table 12. The Table depicts estimated effective federal income tax rates, produced by the Congressional Budget

[21]This statement is based on the authors' calculations using the CPS.

Table 12: Effective Federal Income Tax Rates, by Quintile of the Income Distribution.

	1979	1981	1983	1985	1987	1989	1991	1993	1995	1997
Lowest Quintile	-0.8	-0.2	-0.6	-0.2	-1.3	-1.9	-2.9	-3.4	-6.8	-7.1
Second Quintile	3.9	4.6	3.5	3.9	3.2	3.3	2.7	1.7	1.5	1.6
Middle Quintile	7.5	8.2	6.8	6.8	6.1	6.5	6.3	5.9	6.1	6.2
Fourth Quintile	10.4	11.3	9.5	9.3	8.7	8.9	8.7	8.5	8.7	8.8
Fifth Quintile	16.3	17.1	14.5	14.3	15.1	15	14.7	15.4	16.3	16.8
All Quintiles	11.6	12.6	10.7	10.7	10.8	10.9	10.5	10.8	11.3	11.8

Source: Congressional Budget Office (2001). *Effective Federal Tax Rates, 1979-1997*.

Office (CBO). The CBO starts with detailed income data from the CPS, supplements it with imputations based on IRS and other federal agency data, and finally produces a set of effective tax rates. These data are available from 1979-1997. The lowest quintile of taxpayers pay less than zero tax, while the highest quintile pay approximately 15% of their income.

While the data are not as high-quality as the CBO data, the CPS (from 1980 through the present) itself contains imputations of federal tax liabilities, based on self-reported income data and other federal data sources from the IRS.[22] Using these data, we can estimate the gradient in taxes across education groups, rather than quintiles of the income distribution. The top two panels depict the gradients for the year 2000, while the bottom two depict 1980. The left-hand panels display the data for males, while the right-hand ones show them for females. In every instance, the gradient in taxes paid is steeper than the gradient in income. Interestingly, there has been some expansion in the progressivity of the income tax system over the past 20 years, particularly for males. The figure illustrates empirically the legislated progressivity of the federal income tax system.

[22]The CPS data have three important limitations, relative to the CBO estimates. First, the CPS does not collect data on itemized deductions and capital gains, both of which the CBO imputes. Second, the CPS data are topcoded. Finally, CPS incomes differ significantly from incomes reported on tax returns, which are believed to be more reliable. CBO adjusts the CPS data to bring it in line with tax return data.

Figure 4: Age-Specific Federal Tax and Income Ratios of College Graduates to High School Dropouts, 1980 and 2000

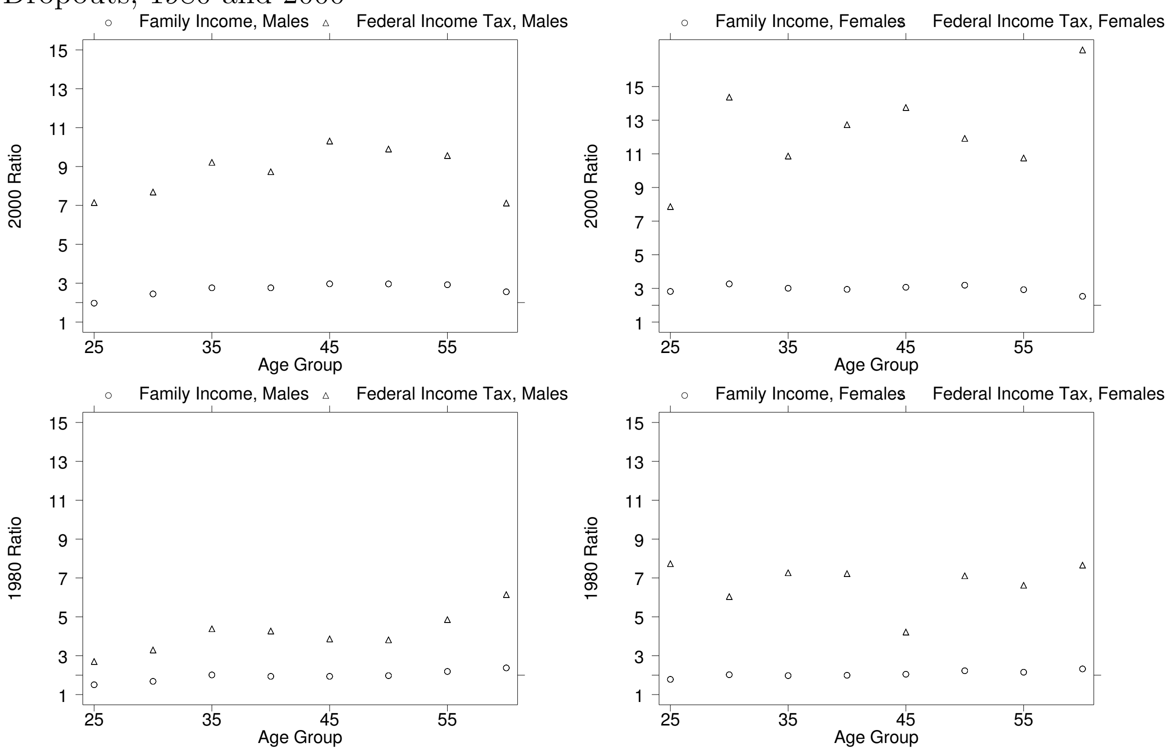

7 Comparison with Previous Research

Unlike McClellan and Skinner (1997), we find that the financial returns to Medicare are much higher for disadvantaged groups, both in absolute terms and *a fortiori* as a percentage of lifetime income. The key source of difference in our results is the finding that Medicare benefits are significantly higher for less educated groups. Other research using aggregate measures of SES find a flat or positive SES gradient in benefits. Using aggregate measures of SES, McClellan and Skinner (1997) find that Medicare transferred dollars from the poor to the rich, but as we have shown, the use of individual-level measures of SES leads to very different conclusions.

In this section, we first show that the difference in our results seems to turn on our use of individual-level data, rather than aggregated data. We then go on to show that aggregate measures of SES suffer from more misclassification error than individual-level measures. In particular, there are more poor people living in rich zip codes than poor people with high educational attainment. As a result, geographic measures of income would overstate health expenditures for rich zip codes and tend to flatten out an otherwise negative gradient in medical expenditures. Finally, we provide evidence from the HRS that medical expenditures are negatively correlated with permanent income when the latter is measured as well as possible.

7.1 The Effects of Aggregation

There are several differences between our approach and that of McClellan and Skinner (1997), but none seems to matter as much as their use of aggregate measures of SES. There are also two significant, but incidental differences: we use a different data source for Medicare benefits, and we use education-based rather than income-based measures of permanent income. These differences are probably incidental, because aggregating within the MCBS data itself also produces the effect of flattening these gradients, and because aggregate measures of education—not just income—also yield flatter SES gradients in Medicare benefits than individual-level measures. These results suggest that the validity of our method is best assessed by examining the validity of our individual-level measure of socioeconomic status.

7.1.1 Aggregating Within the MCBS

To compute gradients across aggregated measures of income, we first compute geographic measures of permanent income in the MCBS. Since the MCBS reports the county of residence for each respondent, we link the MCBS to BEA data on per capita income (described in Appendix A) at the county level for each year of the survey. We then split up the MCBS sample into county income quintiles, using the MCBS sample weights. In essence, we are ranking each year's MCBS respondents by county income, and then dividing up each yearly ranked sample into five quintiles of equal population weight.

Figure 5 reports the results. The left-hand panels depict the benefit gradient across county income quintiles, while the right-hand panels depict the gradient across education groups. The data points in the right-hand panel correspond to the figures reported in Table 2. The gradient across county income quintiles is either flat or somewhat positive, even though in the same data, the least educated individuals receive the most per capita Medicare benefits.

Aggregation seems to affect the size of the gradient, but not *changes* in the size of the gradient, at least qualitatively. For instance, even across county income quintiles, the gradient for females is more negative than for males. This pattern is replicated across

Figure 5: Per Capita Medicare Benefits Across Education Groups and County Income Quintiles.

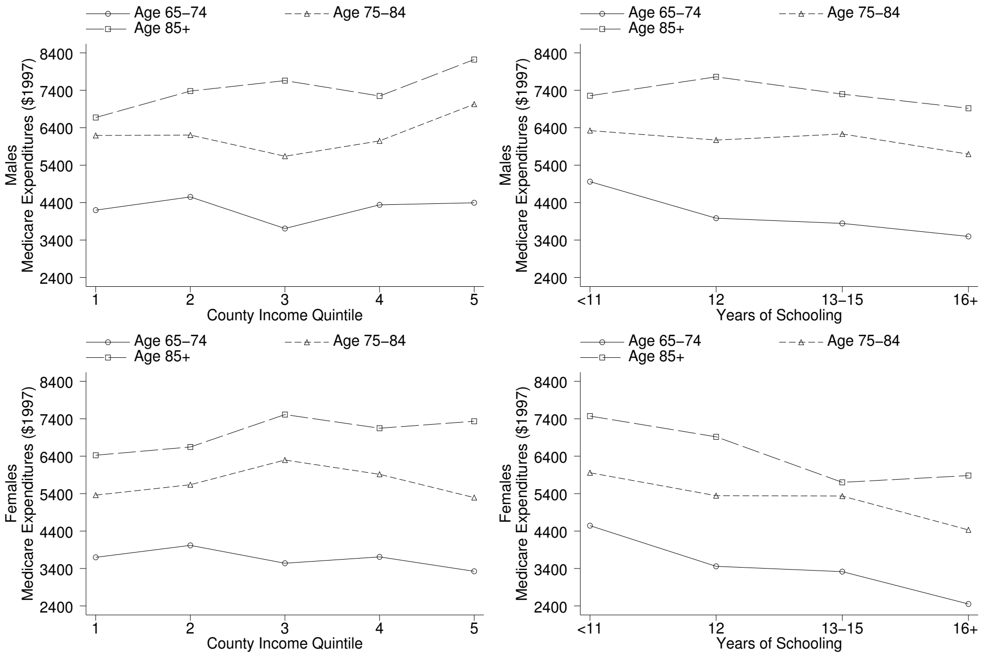

individual education groups as well. Nonetheless, there is a fairly consistent positive trend in benefits across county income quintiles for males 75 and above, and for females over 85. Trends for men aged 65-74 and women aged 65-84 are flat, from the bottom to top quintiles. On the basis of the county income quintile data, we might conclude that residents of richer counties spend more or about the same amount of Medicare's resources, but the individual-level data suggests that the most educated people use by far the least amount of resources.

To compare our methods most directly with those of McClellan and Skinner (1997), we also used zip code income, just as they did. Using the MCBS data on zip code of residence, we link the MCBS to measures of per capita income in each Zip Code from the 1990 Census, which reports 1989 income.[23] Figure 6 reports the results of computing Medicare benefits

[23]The Census data are described in Appendix A.

Figure 6: Per Capita Medicare Benefits Across Education Groups and Zip Code Income Quintiles.

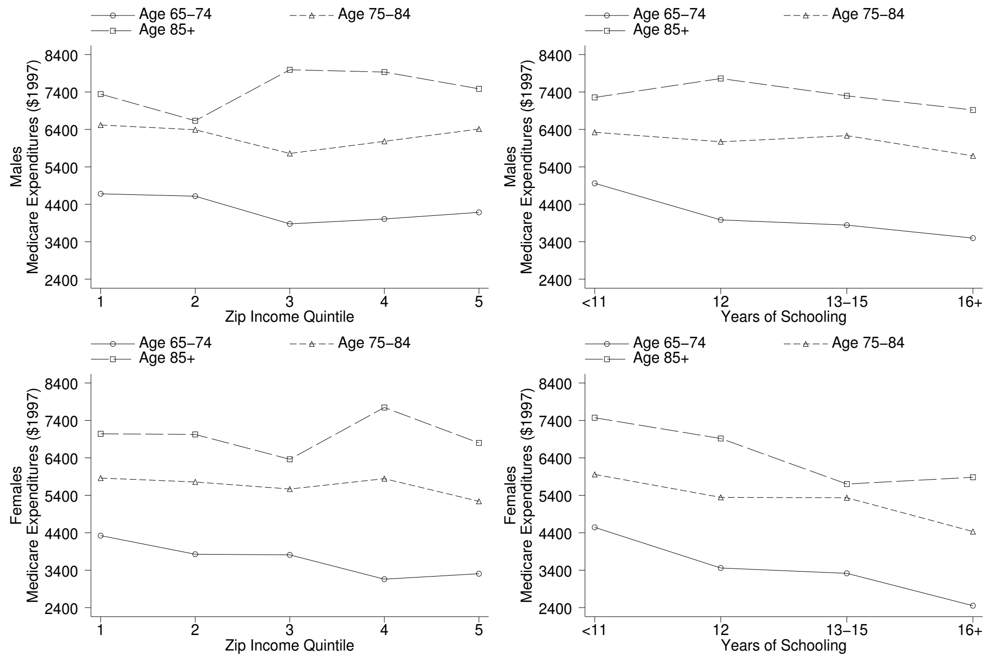

across Zip Code income quintiles.

The curves across Zip Code quintiles are flat or slightly increasing for men and women over age 75, but they do decrease for people aged 65-74. In fact, they show significant negative gradients for this age group. However, the magnitudes are still not close to the magnitudes for the individual-level data. From peak-to-trough, the zip income quintile declines by $1000 for females and $700 for males. In contrast, the gap between high school dropouts and college graduates is $2000 for females and $1500 for males, about twice as large. Assuming we can trust the measurement of zip code income data, it would appear that using this lower level of aggregation does lessen the discrepancy with the individual-level measures, although only partially.

A regression context can provide a more precise sense of these results. To facilitate the

Table 13: Age- and Sex-Specific Distribution of Education in the MCBS, 1992-99.

	Age	No HS	HS Attend	HS Grad	Some Coll	Coll Grad
		Educational Attainment				
Men	65-74	0.20	0.18	0.29	0.16	0.21
Men	75-84	0.26	0.21	0.26	0.14	0.17
Men	85+	0.37	0.23	0.17	0.13	0.13
Females	65-74	0.17	0.19	0.38	0.18	0.11
Females	75-84	0.24	0.23	0.32	0.15	0.10
Females	85+	0.33	0.23	0.24	0.13	0.10

Source: MCBS, 1992-99.

comparison between education groups and zip code income quintiles, we divided the MCBS population into five groups: no high school, some high school, high school graduates, some college, and college graduates. Unfortunately, these are not equally weighted quintiles, but for some age groups, they are close. Table 13 depicts the distribution of education across age- and sex-specific groups in the MCBS. For all except the oldest age groups, the bottom two education groups are roughly equivalent to quintiles. The group of high school graduates, however, tends to be larger than a quintile, while the two college groups tend to be smaller than quintiles. In making comparisons across the distributions of education and zip code income quintile, therefore, it is important to bear in mind that the size of the high school graduates group is often larger than the size of the third quintile, while the sizes of the two college groups are often smaller than the top two quintiles.

Using the zip code income quintiles and the five education categories, we estimate the following regressions separately for age and sex-specific categories:

$$McareTotal_{it} = \beta_0 + \beta_1 ZipQuint_{it} + \lambda_t + \epsilon_{it} \tag{7.1}$$

Table 14: Comparing Medicare gradients across Zip Code Income Quintiles and Education Groups.

	Age	Zip Income Quintile 2nd	3rd	4th	5th	Educational Attainment HS Attend	HS Grad	Some Coll	Coll Grad
Men	65-74	-59	-732	-421	-272	-88	**-974**	**-1008**	**-1380**
		(341)	(327)	(340)	(345)	(342)	(286)	(318)	(303)
	75-84	-223	**-859**	-566	-216	-342	-381	-205	**-817**
		(439)	(383)	(407)	(426)	(339)	(328)	(441)	(353)
	85+	-635	681	406	182	722	785	455	-192
		(615)	(668)	(642)	(612)	(518)	(600)	(625)	(537)
Women	65-74	-381	-520	**-1138**	**-1110**	**654**	**-775**	**-1078**	**-1839**
		(286)	(273)	(245)	(254)	(301)	(231)	(257)	(242)
	75-84	-245	-287	-81	**-637**	-66	**-576**	**-637**	**-1515**
		(297)	(301)	(292)	(296)	(269)	(226)	(274)	(277)
	85+	-91	-682	665	-264	23	**-655**	**-1858**	**-1638**
		(394)	(375)	(412)	(389)	(365)	(332)	(356)	(441)

Source: MCBS, 1992-99.

Note: Data are for total (Parts A+B) Medicare expenditures, adjusted for regional price variation using the GPCI and the hospital wage-price indices. Numbers in the table are based on age- and sex-specific regressions of price-adjusted total Medicare expenditures on dummies for year, as well as dummies for either Zip Income Quintile or Educational Category.

$$McareTotal_{it} = \gamma_0 + \gamma_1 Educ_i + \mu_t + \phi_{it} \tag{7.2}$$

$McareTotal_{it}$ represents individual $i's$ total Medicare expenditures (i.e., Parts A and B) at time t.[24] $ZipQuint_{it}$ represents individual $i's$ quintile in the zip code income distribution at time t, where zip code income is always based on the income of zip codes in the 1990 Census. The variables λ_t and μ_t represent time-specific fixed-effects.

Table 14 reports the results of these regressions, which are identical to the curves in Figure 6, except the regression estimates also remove a year-specific fixed-effect. The numbers in

[24]The differences between gradients are quite similar for the Parts A and B components as well.

bold are significantly different from zero at the 5% level. As in the figure, the gradient across education groups is more steeply negative for women than for men, and more steeply negative for younger men than older men. However, considering that women account for 60% of the elderly Medicare population, and men between the ages of 65 and 74 account for another 20%, the bulk of the elderly Medicare population exhibits a significantly negative gradient across education groups. In contrast to the negative gradient across education groups, the gradient across zip codes is essentially flat among 65-74 year-old men and women over age 85.

To compare accurately the gradients for 65-84 year-old women, we estimated our regression across zip code income deciles. In the education distribution, the difference between the top decile (college graduates) and the bottom quintile (those with no high school) is about $1839 for 65-74 year-old women. In contrast, the difference between the top zip code decile and the bottom zip code quintile is $1162, less than two-thirds of the value across education groups. For 75-84 year-old women, the difference between the top education decile (college graduates) and the bottom quintile (those with no high school) is about $1515. The analogous difference across zip codes is just $360.

7.1.2 Aggregate Measures of Education

We might also ask whether our results owe themselves to the distinction between education and income, or the distinction between aggregate and individual-level measures. To assess this, we looked at gradients across aggregate measures of education. Data from the 1990 Census (described in Appendix A) allow us to compute the fraction of people within each

Figure 7: Per Capita Medicare Benefits Across Education Groups and County Education Quintiles.

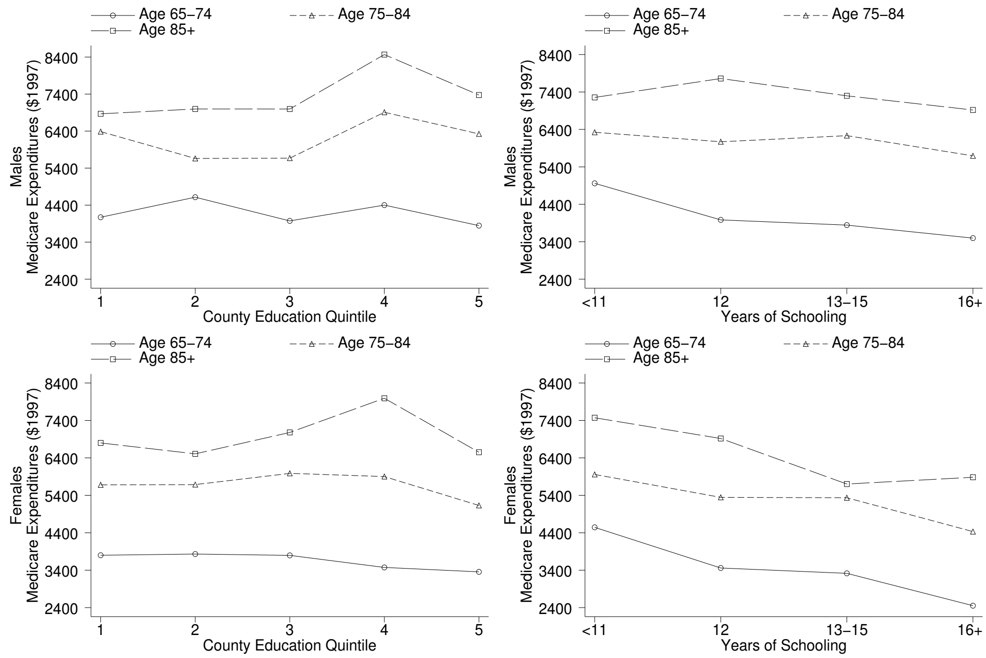

county who had at least a college degree (the average proportion is about 20%) in 1989.[25] Since the MCBS contains data on county of residence, we can link these data to MCBS respondents. Based on the fraction of College Graduates in 1990, we construct county education quintiles. One caveat to note is that we link the 1990 Census Data to the 1992-99 MCBS data and thus could be measuring county-wide education with error. However, our results for the 1992 MCBS are quite similar to those for the 1999 MCBS, suggesting that the expansion of measurement error over time is not affecting the estimated gradients in benefits.

The gradients across individual education groups and county education quintiles are shown in Figure 7. From the top to the bottom education quintile, there is often no change,

[25]Similar results were obtained for the proportion of high school graduates and the proportion of high school dropouts.

or a decline of less than a few hundred dollars. In contrast, the gradient across individual-level education often exceeds $1000 from the bottom group to the top. Moreover, while the gradient across individual education groups is almost always flat or negative, there are several instances of benefits rising across county education quintile.

Qualitatively, Figure 7 looks somewhat similar to Figure 5, which illustrates the gradient across county income quintiles. In both figures, the individual-level education measures generate a more negative and more consistently negative slope than the aggregated measures.

7.2 Aggregation Bias

The previous section documented the differences between aggregated and disaggregated measures of SES, but it is not yet clear that the aggregated measures are less accurate than the individual-level measures. Both, after all, are merely proxies of the underlying variable of interest, which is permanent income. It remains to evaluate the accuracy of individual-level education, compared to zip code income.

To compare the two proxies, we need some third measure of permanent income that we are willing to treat as a "gold standard." Our strategy is to use data on total family income from the HRS, over all 5 waves of the panel. Since the HRS was designed in large part to measure income accurately, its total family income measure is of reasonably high quality, even though it is self-reported. Moreover, its panel aspect provides us a profile of total family income over a nine-year period during which the HRS cohort members were in their 50s. This is a particularly advantageous time in the life-cycle to pick, because earnings profiles are relatively flat over this age range, as evidenced by the relatively flat age-earnings profiles in the HRS. Therefore, it is reasonable to assume that human capital is roughly constant over this age range, and that mean total family income over this period represents a reasonable measure of mean permanent income. We take information on zip code of residence from the HRS Geocode file, and on educational attainment from the HRS main file. Zip code income is the per capita income of the zip code in the 1990 Census.

Using this measure as our gold standard of permanent income, we explore how well education and zip code income predict it. In particular, for various percentile points p in the distribution of permanent income, we assess whether an individual is above or below p in the permanent income distribution, and whether she is above or below p in the proxy variable distribution. If she is above p in the permanent income distribution but *below* it in the proxy distribution, we say that the proxy generates a "false negative." Conversely, if she is below p in the permanent income distribution but above it in the proxy distribution, we say that it generates a "false positive."

The percentiles we chose were driven entirely by the lumpy education distribution. In the HRS, 28% of the sample are high school dropouts, 64% have a high school degree or less, and 83% do not have a college degree. We use these three cut-points in our assessment of false positives and negatives. The results are shown in Table 15. Rates of false negative are almost exactly identical for the two proxies, but the key difference lies in the rates of false positive. No matter which percentile we choose, zip code income produces more false positives than education. This suggests that education serves as a more accurate proxy for permanent income. Moreover, the nature of the bias is consistent with our finding that zip code produces a more positively sloped gradient than individual-level education. There are more poor people in rich zip codes than there are poor people with high educational attainment. If it is true that the poor spend more on medical care, this would tend to bias upward the SES gradient estimated using zip code.

Table 15: Classification Error Using Zip Code Income and Education as Proxies for Permanent Income.

		False Positive Rate		False Negative Rate	
Educational Category	Percentile	Education	Zip Code	Education	Zip Code
HS Dropout	0.28	0.13	0.15	0.12	0.12
HS Grad	0.64	0.15	0.20	0.16	0.15
Coll Attendee	0.83	0.10	0.15	0.10	0.11

Sources: HRS Geocode File and Main File, Waves 1 through 5.
Notes: A false positive is an individual who is above the given percentile in the proxy distribution, but below it in the permanent income distribution. A false negative is one who lies below the given percentile in the proxy distribution, but above it in the permanent income distribution. Permanent income is the mean of real family income across all years of the panel, from 1991-99.

Indeed, in the HRS, it seems that poor people in rich zip codes are in fact sicker and spend more on medical care than others in their zip codes, while richer people in poorer zip codes are healthier and spend less. The HRS collects data on medical expenditures by asking questions about total medical expenditures over the previous two years.[26] In particular, they are asked separate questions about the total cost of: hospital stays, nursing home stays, doctor visits, outpatient surgeries, dental visits, prescriptions drugs, and home health care. In addition, they are asked about costs covered by: Medicare, Medicaid, out-of-pocket, or other insurance. Missing values are imputed based on the HRS bracketing questions. Details of the imputation procedure can be found in StClair et al. (2001). Unfortunately, the HRS does not separately report total Medicare expenditures. Therefore, we concentrate on total expenditures for the HRS population that is below age 65 and not eligible for Medicare. Since this population does not have access to Medicare, it seems that it would be less likely to exhibit a negative gradient across socioeconomic status than the Medicare population.

[26]We divide these numbers by two in order to annualize the estimates.

Table 16: Health Characteristics of Correctly Classified and Misclassified HRS Respondents.

	Family Medical Expenditures	Any ADL's?	Number ADL's	In Hosp. This Year?	Number Hosp. Stays
True Negative	$6,643	0.08	0.14	0.13	0.23
False Negative	$4,665	0.01	0.01	0.08	0.10
True Positive	$5,395	0.01	0.01	0.07	0.09
False Positive	$6,551	0.05	0.07	0.11	0.17

Notes: Family Medical Expenditures are in 1997 dollars. The HRS Activities of Daily Living (ADLs) consist of: walking across a room, dressing, bathing, eating, getting into and out of bed, and using the toilet. Misclassification is assessed according to whether the individual is above or below the 64th percentile in the distributions of permanent income and zip code income.

Table 16 depicts real family medical expenditures, along with various measures of health and health care utilization, for respondents in different classification categories. False positives are uniformly sicker than true positives, while false negatives are healthier than true negatives. All these differences are statistically significant at the 5% level. Therefore, within zip codes, higher socioeconomic status is associated with better health and lower medical expenditures. Moreover, this is true for the near-elderly under 65 who do not have access to Medicare coverage. It appears that zip code income is alone among SES measures in its positive relationship with medical expenditures.

Further evidence for this claim is provided by examining the relationship between family medical expenditures and various measures of SES in the HRS. The relationships between real total family medical expenditures and various measures of income, for all five waves of HRS data, are shown in Table 17. All standard errors are clustered by individual and robust to heteroskedasticity. The first two columns use current period income as a measure of SES, while the next two use our "permanent income" measure, which is an average of real family income over all five waves of the HRS. The latter measure smooths out some of the temporary fluctuations in current period income. We use both the simple value of in-

Table 17: Correlation between Total Real Medical Expenditures and an Index of Permanent Income, for the HRS Population below age 65.

		Current Income		Permanent Income		Instrumented Income
Males	Income Level	-0.003 **		-0.005 ***		-0.007 *
		[0.001]		[0.002]		[0.004]
	Income Quintile 2		-1,128 *		-2,218 ***	
			[592]		[648]	
	Income Quintile 3		-1,822 ***		-1,954 ***	
			[547]		[687]	
	Income Quintile 4		-1,487 ***		-2,584 ***	
			[556]		[644]	
	Income Quintile 5		-2,677 ***		-3,194 ***	
			[492]		[597]	
	Observations	15434	15434	15598	15598	14832
Females	Income Index	-0.005 **		-0.008 *		-0.015 **
		[0.003]		[0.005]		[0.006]
	Income Quintile 2		-1,714 ***		-1,448 ***	
			[396]		[441]	
	Income Quintile 3		-1,533 ***		-1,401 ***	
			[456]		[446]	
	Income Quintile 4		-1,960 ***		-2,121 ***	
			[413]		[460]	
	Income Quintile 5		-2,467 ***		-2,410 ***	
			[452]		[435]	
	Observations	18020	18020	18169	18169	17353

Source: HRS, Waves 1-5, Respondents under age 65.

Notes: Results based on regression of real total family medical expenditures on income measures. Instrumented Income column based on IV regression of real total family expenditures on permanent income, instrumented using Zip Code Income Quintile and Education Group. All regressions include single-year age dummies and single year dummies.

*Significant at 10% level

**Significant at 5% level

***Significant at 1% level

come as a regressor, as well as yearly income quintiles. Regardless of the income measure or the specification, there is a uniformly negative and significant relationship between medical expenditures and income, and this holds true for a population without access to the universal coverage afforded by Medicare. The regressions suggest that a $1000 increase in total family income is associated with a $3 reduction in medical expenditures for males, and a $5 reduction for females. For permanent income, the gradient gets a bit more steeply negative, suggesting that longer-run growth in income is associated with even more of a decline in medical spending. a $1000 increase in average income over the panel is associated with a $5 reduction for males and an $8 reduction for females. In terms of quintiles, the top income quintiles (measured either in terms of current or permanent income) spend between $2500 and $3000 less on medical care.

The last column of the table investigates the correlation between medical expenditures and the portion of permanent income measured by our two proxies—zip code income quintile and education category. It represents the results of an instrumental variables regression of total family expenditures on permanent income (along with age and year dummies), which is instrumented using zip code income quintiles and education categories (high school dropout, high school graduate, college attendee, and college graduate). Here too, the relationship between medical expenditures and socioeconomic status is negative. Using the information content of both proxies at once also yields a negative relationship between medical expenditures and socioeconomic status.

8 Conclusions

At any given age, the poor seem to receive more real Medicare benefits than the rich, at least when poverty is measured using individual-level educational attainment. These gradients are significant enough that they almost exactly offset the effects of early mortality for the poor. As a result of these gradients, it appears that Medicare transfers considerably more resources to the poor than previously thought. In addition, it also appears that measuring SES using aggregated income measures may understate the true benefits of Medicare to the poorest groups.

Compared to previous work, our results are much stronger in favor of the conclusion that Medicare benefits the poor. They would only be strengthened if we also accounted for the observation of McClellan and Skinner (1997) that the nonpecuniary value of Medicare is higher for the poor than for the rich. Not only is Medicare progressive in the sense that its value represents a larger percentage of lifetime income for the poor, but the net present value of Medicare is actually higher in absolute terms for high school dropouts than for college graduates.

While we have investigated the financial returns to Medicare, further research is needed to determine the true welfare consequences of Medicare, which include more than just the pure dollar transfers. A dollar of Medicare benefits may have different values for people in different educational groups, primarily because the alternatives to Medicare may involve very different welfare levels for each educational group. For example, Medicare is much more valuable to a group that would have had no insurance in its absence, and much less valuable to a group that would have preferred to consume less insurance than Medicare mandates were it given the choice. Another possibility here is that the poor might value health insurance less (or perhaps even more, under certain circumstances) than the rich. A more structural empirical approach might be called for to address these questions.

APPENDIX

A MCBS data

The MCBS contains detailed data on health expenditures and especially on Medicare expenditures. MCBS respondents are linked to Medicare administrative data on claims.[27] From the claims data, the MCBS constructs total annual Medicare fee-for-service expenditure for each respondent, as well as the total annual payment made to a Medicare HMO on behalf of each respondent.[28] The sum of the two represents Medicare's total outlay on each individual.

Medicare fee-for-service payments can be further broken down into Part A and B expenditures, by using data on the type of service rendered. MCBS breaks expenditures down into the following service categories: inpatient hospital visits, outpatient hospital visits, institutional utilization stays, facility stays, home health utilization, hospice stays, medical provider visits, prescribed medicine, and dental visits. We take Part A expenditures to be Medicare fee-for-service expenditures for: facility visits, home health utilization, hospice visits, inpatient hospital visits, and institutional utilization. Part B expenditures are Medicare fee-for-service expenditures for: dental visits, medical provider fees, and outpatient hospital visits.

We also need to decompose Medicare HMO payments into Part A and B components. Medicare pays a flat fee to private HMO's, who in turn provide hospital insurance as well as insurance for items that would normally be covered by Part B. To decompose these payments, we use Medicare's payment schedules to calculate these quantities explicitly. Medicare's Part

[27]For details of the linking procedure, see Eppig and Chulis (1997).

[28]About ten to fifteen percent of elderly Medicare beneficiaries are enrolled in Medicare HMOs. These are private HMOs that contract with Medicare to provide medical care in exchange for a flat, per capita fee.

A and B payments to HMO's are determined each year as a function of the individual's age, sex, county of residence, coverage by employer-based insurance, coverage by Medicaid, and ESRD (end-stage renal disease) status. Since we have data from Medicare on the specific payment schedule from 1992 to 2002, and since the MCBS reports all the relevant characteristics for each individual, as well as monthly[29] data on whether an individual is enrolled in an HMO, we explicitly calculate the monthly Part A and Part B payments made by Medicare to HMO's for all respondents in HMO's.

The geographic identifiers in the MCBS allow us to link it to several important databases discussed in the text. The first is a data set containing the GPCI used to deflate Medicare physician payments, and the hospital wage-price index used to deflate hospital payments. The physician GPCI's are used by Medicare to adjust expenditures for differences in area labor costs, practice expenses, and malpractice expenses. The wage-price indices are used to adjust hospital expenditures for differences in labor cost. Due to the exclusion of capital costs (which account for an average of 30% of hospital expenditures), the standard practice is to use (0.7)*(Wage Index)+(0.3) as an index of total hospital expenditures. We adopt this convention. Both indices were matched to beneficiaries in the MCBS sample by county and year. We deflate Part A expenditures using the hospital wage-price index and deflate Part B expenditures and HMO capitation payments by the GPCI deflators.

The second data set contains Bureau of Economic Analysis data on per capita personal income at the county and state level, available by year. Of course, since the BEA classifies counties according to the FIPS scheme, and the MCBS classifies them according to the SSA scheme, a crosswalk is used. The last is data from the 1990 Census on 1989 per capita

[29]In the MCBS, the data on ESRD status, institutionalization status, and employer-based health insurance coverage are also available monthly from Medicare administrative data.

income and area-wide educational attainment at the zip code, county, and state levels.[30] The data contain the number of residents in each zip code, county, and state with a certain educational attainment, along with the total residents. They also contain per capita income data at each level of geographic aggregation. Using these three databases, we are able to augment the MCBS data so that it contains for each respondent: county-level price deflators for all components of Medicare; per capita personal income in state of residence for each year; per capita personal income in county of residence for each year; per capita income in zip code of residence during 1990; and educational attainment in the state, county, and zip code of residence during 1990.

B Health and Retirement Study

The HRS is a longitudinal study of individuals born between 1931 and 1941, who have survived until 1992. The first wave of the HRS was conducted in 1991. The fifth wave collected data for 1999. It can be linked to quarterly Social Security Administration (SSA) earnings records that go back to 1951. This linked file contains earnings records for 9537 HRS respondents present in Wave 1. Between the linked file and the HRS main files, we have quarterly earnings histories from 1951 through 1999. The linked Social Security file contains data on Social Security covered earnings, or the amount of earnings subjected to Social Security payroll taxes. However, from 1966 to 1992, the Medicare earnings maximum was the same as the Social Security earnings maximum.

[30]These data are taken from GeoLytics (1996).

B.1 Interpolation Across Time in the HRS Main File

We use five waves of the HRS data. Waves 1 and 2 record income data from 1991 and 1993, respectively. Wave 3 records it in 1995 or 1996, depending on when the interview was conducted. Wave 4 records it in 1997 or 1998, and Wave 5 records 1999 data. From these data, we exponentially interpolate missing years, but only if we have data on years prior to *and* following the missing year. In other words, we do not extrapolate any data.

B.2 Family Tax Liability

Since Medicare is financed by a payroll tax, the total expected tax liability ought to be calculated at the level of the family. Men tend to work more and pay more taxes than women, but these are taxes borne by the entire family, rather than just the individual man. The HRS data simplifies the task of computing annual taxes paid by couples, since a reasonable number of married couples in the HRS cohort are both present in the HRS data and the linked Social Security earnings data. For these people, we have complete data on the couple's income. The remaining respondents include the never married, widow(er)s, divorce(e)s, and married people whose spouse is simply not present in the linked earnings file. For these people, we must impute spousal earnings, according to an algorithm we describe below.

Table A-1 provides a useful description of the data. There are 13,478 respondents in Wave 1 of the HRS. 3941 of these are not present in the linked Earnings History file. We drop these observations. As long as selection into the Earnings History file is random, this introduces no bias.[31] Another 6668 people (or 3334 couples) are present with their spouses or partners

[31] Haider and Solon (2000) show that, conditional on having a Social Security Number, selection into the SSA file is indeed random.

Table A-1: Availability of data in the HRS Earnings History File.

	Presence in Earnings History File			
Marital Status in 1991	Respondent and Spouse Present	Respondent Only Present	Respondent Not Present	Total
Married, Spouse Present	6427	975	2435	9837
Married, Spouse Absent	12	22	23	57
Partnered	229	54	102	385
Separated	0	222	88	310
Divorced	0	807	270	1077
Widowed	0	457	163	620
Never Married	0	264	98	362
Unknown	0	68	762	830
Total	6668	2869	3941	13478

in the Earnings History file. For each of these people, we are able to calculate earnings for the couple. Of the remaining 2869 people, 264 were never married; as such, individual income is equal to family income, and we drop the 68 respondents for whom marital status is unknown. This leaves 2537 people for whom family income must be imputed. Consider first the 1051 married or partnered respondents in this group. We impute spousal earnings by looking at similar respondents and calculating the earnings of their spouses. Specifically, we compute the real average spousal earnings profile of all similarly aged and educated HRS respondents (of the same sex). The average earnings profile is then assigned to each respondent whose spouse is not present in the data. As discussed above, the 1029 divorced or separated respondents are treated as if they were married; average spousal earnings are imputed for them according to the same procedure. Even if the individual has been divorced more than once, our strategy will not be affected, as long as his spouses have been similarly educated.

This leaves only the 457 widowed respondents. The difficulty with these respondents is estimating the year of their spouse's death, which is not reported in the data. The best we can do is to make use of the HRS variable for ”length of longest marriage.” For those who are currently married in wave one of the HRS, we compute the year they would have been married, assuming that their current marriage is their longest marriage. This yields the most recent year in which they could have been first married. We then compute the average year within the four education groups we are considering, racial category (white, black, or other), and age in 1991. This yields our estimate of year of marriage for widow(er)s. Using the variable for length of longest marriage, we then compute the year in which each widow's spouse would have died. This date is used to truncate the average real spousal earnings profile estimated above, and this finally yields the earnings that the deceased spouse would have contributed to the partnership. As a result of the data limitations we face, this is a highly imperfect strategy, but it is important to stress that it affects less than 5% of our sample. Even if we were to mismeasure income by 50% for these respondents, it would have less than a 3% impact on our estimates of average income.

B.3 Self-Employment Income

The HRS SSA file does not break apart taxable income into self-employment income and wage income, even though Medicare taxed these two types of income at different rates from 1966 to 1983. Today, the worker and firm each pays half the tax on wage earnings. However, through 1983, self-employed people paid at the tax rate faced by the worker alone, which amounts to half the total Medicare tax paid. Prior to this year, therefore, self-employed individuals faced a lower total tax rate than workers. During the years with a Medicare earnings cap, if a worker had earnings both from wage work and self-employment, her wage taxes were calculated first, and then her self-employment tax. For example, suppose a worker

Table A-2: Proportion of Self-Employment income in the HRS Cohort.

		Males				Females			
	Age Group	Less than HS	HS Graduate	Coll Attendee	Coll Grad	Less than HS	HS Graduate	Coll Attendee	Coll Grad
1966	25-29	1.3%	2.6%	3.0%	2.1%	1.4%	1.8%	8.5%	2.1%
	30-34	2.2%	4.8%	2.2%	2.9%	1.4%	1.9%	7.8%	0.8%
	35-39	3.7%	7.0%	7.0%	6.6%	3.6%	3.9%	2.8%	2.1%
1982	40-44	4.8%	6.0%	5.3%	8.8%	5.1%	4.6%	6.3%	2.1%
	45-49	6.9%	12.6%	11.5%	9.1%	8.7%	4.3%	3.8%	9.9%
	50-54	8.4%	9.1%	8.4%	16.1%	4.0%	3.2%	3.7%	5.7%

Source: Current Population Surveys, 1966-82.

in 1967 had \$6000 in wage income, and \$4000 in self-employment income. Taxes would have been collected on all her wage income, but only the first \$600 of her self-employment income. Her total tax would have been: (1.0%)*\$6000+(0.5%)*\$600=\$63.

To decompose the HRS income measures into self-employment income and wage income, we use data from the 1966-83 Current Population Surveys (CPS). The CPS asks respondents about wage income, self-employment income, age, sex, educational attainment, and race. From the CPS, we estimate-for every survey year, 5-year age group, education group, sex, and race-the average proportion of total income subject to Medicare tax that was derived from self-employment. We restrict these calculations to CPS respondents that reported some income during the year. These proportions are then used to impute self-employment income and wage income for the 1966-83 period. In practice, these imputations had very little effect on our estimated rates of return from Medicare. Even ignoring this issue—and treating all 1966-83 income as wage income—yields virtually the same rates of return. Nonetheless, for the sake of consistency, we estimate self-employment income. Table A-2 displays these estimated proportions for the age ranges occupied by the HRS cohort in 1966 and 1982. Self-employment income is relatively for women and young men throughout these age ranges. It is, however, somewhat important for men over the age of 40, and particularly for high school graduates.

C CPS Data

For the most part, the processing of the CPS data was quite straightforward, but there were two issues to address.

C.1 Matching Spouses

The most difficult part of using the CPS data is matching spouses. To aid in this task, the CPS provides a "household relationship" variable, which indicates whether the individual is a head of the family, or the spouse of the head. The CPS also identifies family units. We identify spouses as the individuals (of opposite sex) reporting that they are married, with spouse present, within a single family (or subfamily). Since the CPS defines a family around a single married couple, this procedure works in the vast majority of cases. However, there are a few problematic years: 1966, 1967, and 1972. Apart from these years, there are a total of 276 observations that appear to be in families with more than 2 spouses (the entire CPS data set has more than 2 million observations). However, there are 9277 problematic observations in 1966, 4478 in 1967, and 43,474 in 1972. The problem in 1972 appears to be duplicate family or subfamily identifiers: there are many instances of 4 and 6 spouses within a single family. For the other two years, the CPS identified households, but not subfamilies within households; the problematic observations thus could arise in extended families. Excluding these years from our calculations has little effect. All the results in the paper include them, but randomly match spouses within "families" as identified by the CPS.

C.2 Pre-1964 data

The HRS Earnings History file goes back to 1951, but the CPS only goes back to 1964. To make the average tax calculations comparable, we must account for the fact that the 1931-41 birth cohort did not have to pay Medicare taxes for some portion of its working life, prior to 1964. We simply construct data files for 1951-63 in which covered income is set to zero, but the one complication is how to weight observations in the pre-1964 years for the CPS. Essentially, we assume that the composition of the cohort remains fixed at its 1964 level. This ignores some differential mortality, but since the cohort is relatively young at this point of time, this bias ought not to be severe.

References

Boskin, M., and D. Puffert (1988). “The Financial Impact of Social Security by Cohort Under Alternative Financing Assumptions.” *Issues in Contemporary Retirement*. R. Ricardo-Campbell and E. Lazear, eds. Stanford, CA: Hoover University Press.

Brown, J. R., and S. J. Weisbenner (2002). “Is A Bird in the Hand Worth More than a Bird in the Bush? Intergenerational Transfers and Savings Behavior.” *Working paper 8753*. Cambridge, MA: National Bureau of Economic Research.

Burkhauser, R., and J. Warlick (1981). “Disentangling the Annuity from the Redistributive Aspects of Social Security.” *The Review of Income and Wealth* 27(4):402–421.

Card, D. (1995). “Earnings, Schooling, and Ability Revisited.” *Research in Labor Economics* 14:23–48.

Chandra, A., and J. Skinner (2002). “Geography and Racial Health Disparities.” *Paper prepared for the Workshop on Ethnic Disparities in Aging Health*. National Acadamy of Sciences.

Duggan, J., R. Gillingham, and J. Greenlees (1993). “Returns Paid to Early Social Security Cohorts.” *Contemporary Policy Issues* 11(4):1–13.

Eppig, F. J., and G. S. Chulis (1997). “Matching MCBS and Medicare Data: The Best of Both Worlds.” *Health Care Financing Review* 18(3):211–29.

Feldman, J. J., et al. (1989). “National Trends in Educational Differentials in Mortality.” *American Journal of Epidemiology* 129(5):919–933.

Garrett, D. M. (1995). “The Effects of Differential Mortality Rates on the Progressivity of Social Security.” *Economic Inquiry* 33(3):457–75.

GeoLytics (1996). *CensusCD (Version 1.1) [CD-ROM]*. East Brunswick, NJ: GeoLytics.

Geronimus, A. T., J. Bound, and L. J. Neidert (1996). “On the Validity of Using Census Geocode Characteristics to Proxy Individual Socioeconomic Characteristics.” *Journal of the American Statistical Association* 91(434):529–537.

Gornick, M. E., et al. (1996). “Effects of Race and Income on Mortality and Use of Services Among Medicare Beneficiaries.” *New England Journal of Medicine* 335(11):791–799.

Haider, S. J., and G. Solon (2000). “Non-Response Bias in the HRS Social Security Files.” *Working Paper DRU-2254-NIA*. Santa Monica, CA: RAND.

Hurd, M. D., and J. B. Shoven (1985). “The Distributional Impact of Social Security.” *Pensions, Labor, and Individual Choice*. D. A. Wise, ed. Chicago: University of Chicago Press. 193–215.

Jorgenson, D. W., and B. M. Fraumeni (1995). "The Accumulation of Human and Nonhuman Capital, 1948-1984." *Productivity. Volume 1: Postwar U.S. Economic Growth*. D. W. Jorgenson, ed. Cambridge, MA: MIT Press. 273–331.

Lee, J., M. McClellan, and J. Skinner (1999). "The Distributional Effects of Medicare." *Working Paper 6910*. Cambridge, MA: National Bureau of Economic Research.

Long, S. H., and R. F. Settle (1984). "Medicare and the Disadvantaged Elderly: Objectives and Outcomes." *Milbank Quarterly* 62(4):609–656.

McClellan, M., and J. Skinner (1997). "The Incidence of Medicare." *Working Paper 6013*. Cambridge, MA: National Bureau of Economic Research.

Shoven, J., J. Sundberg, and J. Bunker (1987). "The Social Security Cost of Smoking." *Working Paper 2234*. Cambridge, MA: National Bureau of Economic Research.

Siegel, J. J. (1992). "The Real Rate of Interest from 1800-1990." *Journal of Monetary Economics* 29:227–252.

StClair, P., et al. (2001). "SSA-HRS Data Documentation." *Project Memorandum PM-1172-SSA*. RAND.