

Health Expenditure Risk and Annuitization: Evidence from Medigap Coverage

Daniel W. Sacks

August 2012

BWP2012-01

Boettner Center Working Paper

Boettner Center for Pensions and Retirement Research

The Wharton School, University of Pennsylvania

3620 Locust Walk, 3000 SH-DH

Philadelphia, PA 19104-6302

Tel: 215.898.7620 Fax: 215.573.3418

Email: prc@wharton.upenn.edu

<http://www.pensionresearchcouncil.org/boettner>

The project described was supported by the National Institute on Aging, P30 AG-012836-18, the Boettner Center for Pensions and Retirement Security, and National Institutes of Health – National Institute of Child Health and Development Population Research Infrastructure Program R24 HD-044964-9, all at the University of Pennsylvania. The content is solely the responsibility of the author and does not necessarily represent the official views of the National Institute of Aging or the National Institutes of Health.

Health Expenditure Risk and Annuitization: Evidence from Medigap Coverage

Abstract

Theoretical research suggests that health expenditure risk can have an ambiguous influence on the annuitization decisions of the elderly. I provide empirical evidence on this linkage, by estimating the impact of supplemental Medicare insurance (Medigap) coverage on the annuity demand of older Americans. I use local variation in prices as an instrumental variable to address the possible endogeneity of Medigap coverage, an identification strategy motivated by the fact that Medigap policies are not medically underwritten, and Medigap insurance is required by law to be standardized, so prices reflect neither individual characteristics nor product quality. Medigap coverage has a strong impact on annuitization: the extensive margin elasticity is 0.39, the overall elasticity of private annuity income with respect to Medigap coverage is 0.56. These results are robust to controls for health, wealth, and preferences, as well as other robustness tests. They imply that medical expenditure risk has a large impact on underannuitization.

Daniel W. Sacks

The Wharton School of the University of Pennsylvania

dansacks@wharton.upenn.edu

Health Expenditure Risk and Annuitization: Evidence from Medigap Coverage

Daniel W. Sacks

How does health insurance affect annuitization? An emerging theoretical literature explores the impact of medical expenditure risk on annuitization, but there is little consensus on the quantitative importance of this channel. Using an instrumental variables strategy, this paper measures the impact of supplemental Medicare insurance—Medigap—on the annuity demand of the elderly.

The literature on health expenditure risk and annuitization attempts to provide an explanation for the under-annuitization puzzle. The puzzle is that, while economic theory suggests strongly that elderly individuals should annuitize much of their wealth (Yaari, 1965; Davidoff, Brown, and Diamond, 2005), in fact people annuitize little (e.g. Brown 2007). Health expenditure risk can help resolve this puzzle, since health risk decreases the demand for annuities by driving up the value of liquidity. If health expenditures require large amounts of cash-on-hand, then annuities, with their slow and steady payouts, would be unattractive. The quantitative importance of this channel, however, is unclear. While Turra and Mitchell (2008) and Peijnenburg, Nijman, and Werker (2011a) find that medical expenditure risk may decrease the demand for annuities, Pashchenko (2010) and Yogo (2011) suggest that this channel is quantitatively unimportant, and Peijnenburg, Nijman, and Werker (2011b) argue that full

annuitization remains optimal even in the presence of medical expenditure risk. Theory is therefore ambiguous about the of health expenditure risk for (under-)annuitization.¹

Much of what we know about the impact of health risk on annuitization comes from calibrated or estimated structural models. Solving these models requires making untested assumptions about the functional forms for utility and for the distribution of medical expenditure risk. The impact of health expenditure risk on annuitization, in particular, appears to depend on assumptions about access to credit and about the timing of medical expenditure shocks. Peijnenburg, Nijman, and Werker (2011b) conclude that if the elderly can borrow against future annuity income (or save out of current annuity income), then health expenditure risk does not reduce annuitization. This is not surprising, since with perfect access to credit, annuities are just as liquid as any other investment.

In this paper, I take an alternative approach to measuring how medical expenditure risk affects savings: I directly estimate the impact of health insurance on annuity demand. If indeed medical expenditure risks loom large, then health insurance should complement annuitization. I focus on Medicare Supplemental Insurance, i.e. Medigap insurance, a regulated health insurance product that provides extra coverage beyond Medicare. Estimating the impact of Medigap on savings or annuitization is made difficult by the endogeneity of health insurance. The demand for insurance depends on wealth, risk preferences, and riskiness, and all of these factors also affect savings decisions independently. To address these concern, I use an instrumental variables

¹ There are of course many other explanations for the under-annuitization puzzle. Finkelstein and Poterba (2004) and Mitchell et al. (1999) explore adverse selection and insurance loads. Friedman and Warshawsky (1990) and Lockwood (2012) consider the importance of bequest motives. Chai et al. (2011) consider many sources of risk and complex portfolio choice. Brown (2001) contends that lifecycle models have difficulty explaining all of annuity demand, and Brown (2007) therefore argues that psychological factors may contribute to underannuitization. Brown et al. (2008, 2011) present evidence for a framing explanation of the underannuitization puzzle.

strategy, using the zip-code level price of Medigap as an instrument for coverage, conditioning on extensive controls for health, wealth, and preferences in the Health and Retirement Study.

For this empirical strategy to succeed, I require that Medigap prices, conditional on other controls, be uncorrelated with underlying determinants of the demand for annuities. The institutional features of the Medigap market suggest that this condition may be satisfied. By law, Medigap insurance cannot be medically underwritten for individuals purchasing it at age 65. Prices in a given year thus may depend only on the buyer's location, age, and sex. While insurers might charge prices based on average health in an area, it is plausible to assume that, conditional on an individual's own health and sex, the price of health insurance that she faces is uncorrelated with her underlying annuity demand. The variation in insurance prices underlying my estimates comes from variation in local spending habits unrelated to the individual's demand for medical care. Below I present some evidence verifying this identification assumption, showing that health insurance prices appear uncorrelated with individual preferences or health, conditional on demographic variables and permanent income.

Medigap prices, conditional on demographic and other controls, have a very strong impact on Medigap coverage. Across specifications, I estimate a stable and precise price elasticity of demand for Medigap coverage of about -1.5, with large F-statistics that allay weak instrument concerns. The first stage is not robust, however, to controlling for state fixed effects, because these controls absorb almost all the variation in the instrument. An important weakness of this empirical strategy, therefore, is that I cannot control for unobserved state-level factors affecting Medigap prices and annuity or savings demand.

With this concern in mind, I next examine the reduced form relationship between annuitization, and Medigap prices. Nonparametric estimates of this relationship illustrate the

variation that underlies the instrumental variables regression. The result show a clear and negative relationship between Medigap prices and annuitization, measured as either the probability of having positive pension or annuity income, or as the amount of pension or annuity income. The instrumental variables estimates quantifying the relationship between Medigap coverage and annuitization suggest that health insurance has a strong impact on annuitization: the elasticity of non-Social Security annuity income with respect to Medigap coverage is about 0.56. About two-thirds of the response is on the extensive margin; I estimate an extensive margin elasticity of about 0.39. Social Security income also responds to Medigap coverage, with an elasticity of about 0.11. All of these results are robust across a variety of specifications. While the point estimates are perhaps too large to be believed, they nonetheless indicate an important role for medical expenditure risk in annuity demand.

The main contribution of this paper, therefore, is to provide direct empirical evidence on the relationship between annuity demand and medical expenditure risk. It also contributes to the large literature on precautionary savings among the elderly. Palumbo (1999) and De Nardi, French, and Jones (2010) showed that medical expenditure risk has a large impact on the savings among older Americans. Given the importance of medical expenditure risk for overall savings, it is not surprising that annuitization responds as well.

1. Background on Medigap Insurance

Elderly Americans in the United States face several sources of medical expenditure risk. Medicare does not cover long term care, and this shortcoming represents a serious financial risk to the elderly (cf. Brown and Finkelstein, 2007). Moreover, the health insurance provided by Medicare is incomplete because of its cost-sharing provisions. Medicare Part A, which covers hospitalizations, requires patients to pay for the first day of hospitalization, and Part B, which

covers outpatient services, has a 20 percent coinsurance rate. Medicare did not cover prescription drugs until the introduction of Part D in 2006, and the standard benefit in Part D today includes substantial cost sharing with a deductible and a “donut hole” with 100 percent coinsurance rate for expenditures above a certain amount.

Seniors may buy supplemental insurance to fill in the holes in Medicare’s coverage. These policies are called Medigap plans and they are tightly regulated by the federal government. Insurers may offer policies from a menu of 12 standardized plans (denoted by letters A through N). All plans cover Part A coinsurance and at least part of the Part B coinsurance, but they differ in their coverage otherwise.² A key feature of the Medigap market is that if individuals purchase policies at the time they claim Medicare coverage, then insurers cannot use medical underwriting. Plans instead are priced on sex and location.³ Many plans are also priced based on the buyer’s age of purchase, so that individuals have a strong incentive to purchase a policy at age 65, and the price they face at age 65 is the price they will face for the rest of their life (if they choose to renew their policies). These plans are called “age-issued.” Plans can also be “attained-age,” in which case prices vary with age, or “community rated,” in which case prices do not vary with age, but vary from year to year according to a local risk rating. I focus on attained-age prices, because these prices at age 65 are likely to have a strong and lasting influence on Medigap coverage throughout retirement.

Insurers setting these prices are prohibited from underwriting based on individual health or mortality expectations. The association between local prices at age 65 and long-term Medigap coverage, and the fact that prices cannot depend on individual-specific health/wealth/mortality, motivates my instrumental variables identification strategy. Prices in these markets depend

² For a detailed description of Medigap plans and eligibility/pricing rules, see CMS-NAIC (2012), <http://www.medicare.gov/publications/pubs/pdf/02110.pdf>.

³ Plans may also charge different prices to smokers and non-smokers, but this is very rare in practice.

instead on the underlying area-level demand for Medical care (which I assume derive from from risk preferences, health, and wealth). Even conditional on an individual's demand for medical care, there is substantial variation in prices, since insurance companies cannot perfectly risk adjust (or price discriminate). This variation is plausibly orthogonal to individual savings propensities, so it may be useful as an instrument.

2. Data: Weiss Ratings and the Health and Retirement Study

The data for this study come from two sources: the Health and Retirement Study conducted by the University of Michigan, and a proprietary insurance pricing database created by Weiss Ratings, Inc. The HRS provides information on demographics, preferences, wealth, annuitization, and insurance coverage. I merge the HRS data at the zip code level to the insurance pricing information.

2.1 Weiss Ratings Data

Weiss Ratings is a market research firm that collects Medigap pricing information from firms in most states. Their dataset contains pricing information at the level of year-zip code-firm-plan letter, for 1998-2004. For every zip code it is possible to see every plan offered and the prices charged by each company. There is a fair amount of price dispersion within zip code, as Maestas, Schroeder and Goldman (2007) show, but very little dispersion in mean price across zip codes within a state (given plan characteristics). Although in principle, prices can differ by sex, in practice mean prices for men and women are extremely highly correlated within zip code ($\rho \approx 0.99$). Prices are also very highly autocorrelated over time.

I take the price of Medigap insurance as the mean within-zip code price of an attained-age Plan A policy, allowing prices to differ by sex. There is little lost by using mean Plan A price rather than different plans, or some combination of plan prices, since the prices are highly

correlated. For example, the correlation between Plan A and Plan F prices is 0.95. Since the age 65 price determines the path of future prices, I use the mean price in an individual's zip code when she turns 65.

From the analysis, I exclude Massachusetts, Minnesota, and Wisconsin because these states have their own set of Medigap plans, so the prices and policies are not comparable to those in other states. I exclude Connecticut, Maine, New Jersey, and New York because they lack age-issued policies in the Weiss data. I also lack data on Arizona and Vermont for 2002-2004, and on Washington for 1999 and 2002-2004, but when I do have data for those states, I include them in the analysis.

2.2 Health and Retirement Study

The Health and Retirement Study is a longitudinal survey following older Americans. It began in 1992, collecting detailed information about work, health, and wealth for a representative sample of the cohort born between 1931 and 1941; this is the original "HRS" cohort. In subsequent years, the study added additional cohorts born later and earlier; my analysis here focuses only on heads of household from the original HRS cohort. The HRS offers a wider range of variables than other longitudinal datasets like the PSID or NLS, and these extra variables are essential to the present study. The HRS asks about pension and annuity income, and Social Security income. Pension and annuity income consists of all income from private annuities or employer-provided pensions, including both defined-benefit and defined-contribution pensions. These variables are the main outcomes of interest. I also make use of the HRS's extensive health and preference information as additional controls.

The analysis sample is restricted to heads of household age 65 or older with Medicare coverage not covered by employer-provided health insurance. This sample definition follows

Fang, Keane, and Silverman (2008), who study Medigap demand extensively; like them, I exclude 2,352 observations of 553 people with employer-provided health insurance because I expect Medigap coverage to be superfluous for these people. I also exclude individuals with imputed values for pension/annuity income or for any assets, and people with missing values for any of the outcomes or controls.⁴ The final sample has 3,335 observations of 1,406 people.

Table 1 shows summary statistics on the key variables for the full sample and for the analysis sample. All dollar amounts are in 2010-constant dollars. I present summary statistics separately by whether the respondent has Medigap coverage, defined as having supplemental insurance and receiving Medicare (again following Fang, Keane, and Silverman, 2008). About one-third of the sample have purchased Medigap policies. Relative to the uncovered, people with Medigap coverage generally face lower Medigap prices, as expected. The covered group is also to be female, have fewer children, to be white, and have much higher permanent income, defined as average income over the pre-retirement period. (De Nardi, French, and Jones (2010) and Carroll and Samwick (1997) use a similar permanent income concept.) People with supplemental insurance coverage are healthier than those without: they are more likely to report being in good health, and they suffer less from many chronic diseases. Fang, Keane, and Silverman (2008) explore this advantageous selection in detail. People with Medigap are more likely to have annuity income, have more annuity income, and have higher savings than people without.

3. Empirical strategy and first stage results

3.1 Empirical Strategy

In what follows, I use local price variation in Medigap policies as an instrument for Medigap coverage, conditioning on the individual-level characteristics that might affect savings

⁴ In results not reported, I experimented with including individuals with imputed values. In all cases, I found slightly stronger impacts of Medigap on annuitization.

decisions and correlate with local prices. To see why this strategy is necessary, and why it might be successful, let y_{it} be the outcome of interest (savings or annuitization) for individual i at age t and let $Medigap_{it}$ indicate Medigap coverage at t . I assume these variables are related as follows:

$$y_{it} = \alpha(Medigap_{it}) + X_{it}\beta + \epsilon_{it} \quad (1)$$

$$Medigap_{it} = \gamma \ln price_i + X_{it}\theta + \omega_{it} \quad (2)$$

The object of interest is α , the impact of *Medigap* on annuitization. Annuitization decisions, however, depend not only on health insurance coverage but also on wealth, preferences (i.e. risk aversion and the intertemporal elasticity of substitution), and health and mortality expectations. But since individuals choose whether to purchase insurance, $Medigap_{it}$ depends on these variables as well, creating an endogeneity problem.

Using local Medigap prices as an instrumental variable potentially overcomes this problem, since identification comes from variation in insurance prices that is plausibly orthogonal to individual wealth, preferences, or health. That is, insurance prices are a function of aggregate (zip-code) level variables only. These variables likely include aggregate (determinants of) demand for medical care (and also medical care costs, administrative costs, and market structure in the Medigap industry). Because similar individuals—individuals with the same underlying demand for savings products—can live in zip codes with very different demands for medical care, there is a great deal of variation in the price of Medigap coverage, even conditional on savings propensities. Local prices can therefore isolate variation in Medigap coverage, independent of savings attitudes.

One concern in using cross zip-code variation in Medigap prices, however, is that these prices may be systematically correlated with zip-code level characteristics that directly affect

savings behavior.⁵ For example, it may be that areas in which people have unusually high medical expenses have both high Medigap prices and high savings. In principle, people might decide to live in a given zip code because it has low Medigap prices. Here the control variables are useful: the threat to identification comes from the correlation between prices and individual health, wealth, and preferences. But using the rich HRS data, I can condition on these variables. To the extent that the results are robust to these controls, it is unlikely that additional unobserved determinants of savings or annuity demand are also correlated with local prices. The instrumental variables estimator therefore relies on variation in prices across people with similar characteristics to identify the impact of insurance on savings or annuitization.

In what follows, I show that, conditional on demographics and permanent income, Medigap prices are uncorrelated with health and preferences. These results provide some reassurance of the validity of the instrument, since these prices appear uncorrelated with many individual-level determinants of savings and annuity demand. One remaining concern, though, is that high Medigap prices could reflect high administrative costs for insurance companies in that region, and therefore high prices might also signal high annuity prices (i.e. low returns); if so, this could bias estimates towards finding a positive impact of Medigap coverage on annuitization.

Two considerations cast doubt on this possibility. First, in their analysis of annuity premiums from a single large UK insurer in the 1980s and 1990s, Finkelstein and Poterba (2004) report that firms set geographically uniform prices, despite important variation in mortality across regions. This fact suggests that annuity premiums are not likely to be systematically higher in areas with high Medigap prices. It is not clear, however, whether the experience of that

⁵ Local prices can affect annuity or savings demand through income effects as well as substitution effects, but for my purposes this distinction is irrelevant: all that matters is that the prices alone (and not some omitted variable) is affecting annuity or savings.

firm generalizes to the United States in the 1990s and 2000s.⁶ State-level data on annuity prices would let me address this concern directly, but I am unaware of any such data. As a second piece of evidence, I use premium data for several annuity providers collected in 1995 by A.M. Best (examined by Mitchell et al., 1999). I match this pricing information to the provider's state (as listed by A.M. Best in their ratings) and correlate annuity premiums with state average Medigap premiums (in 1998, the earliest year for which I have data). This exercise reveals that the correlation between the price of a \$10/month annuity and state Medigap price is virtually nil (-0.04). This procedure is of course far from perfect, but it suggests there is very little geographic correlation between Medigap prices and annuity premiums.

3.2 First Stage Results

Table 2 presents the first stage estimates of equation (2). Each column presents the estimate γ obtained from using a different set of controls, as well as the F-statistic on $\ln price$, which provides a test of instrument strength. Column (1) reports estimates of γ obtained with a basic set of controls: year fixed effects, a quartic in age, and indicators for race, sex, marital status, educational attainment (high school, some college, and college or more), and dummy variables for having one, two, or three or more children. These controls serve two functions. First, the year and age controls and the sex dummy are directly correlated with Medigap prices, since prices are set yearly and have increased over time. To the extent that there are differences in annuity demand by age, year, or sex, it is therefore important to control for these variables. Second, the demographic variables are closely related to lifetime earnings potential and savings demand, and this may have a large impact on both savings and demand for medical care (through

⁶ Finkelstein and Poterba (2004, pp. 190-1) write, "...like other firms in the market, [the firm we study] varies the annuity price on the basis of only age and gender, and not on the basis of the individual's geographic location," suggesting that this pricing practice is common at least in the UK.

income effects). Controlling for these variables helps control for the correlation between local demand for medical care (which influences Medigap prices) and own demand for annuities.

The coefficient on $\ln price$ (γ) is -0.48 and, with a standard error of 0.06; it is precisely estimated and highly statistically significant. Here and throughout, all standard errors are clustered at the individual level since I pool multiple years of data for each individual. The F-statistic on log-price is 59.4, alleviating concern about a possible weak instrument. The sample mean of *Medigap* is 0.33, implying an average elasticity of about -1.5, somewhat larger than Starc's (2012) estimate of an own price elasticity of demand for Medigap plans of about -1.1.

Figure 1 provides the graphical analog of the first stage relationship. The figure shows the non-parametric relationship between Medigap coverage and Medigap price, both net of controls. To construct the figure, I regress Medigap coverage and $\log price$ on the age, year, and demographic controls (separately), form the residuals, and then estimate a local linear regression of residual Medigap coverage on residual $\log price$; this is exactly the variation underlying the first stage regression. As the Figure shows, there is a clear negative relationship between Medigap coverage and its price. This relationship seems to level off for a range of prices before declining.

While the basic controls adjust for standard demographic differences across people, they may still control inadequately for aspects of insurance demand that are also correlated with local prices. Column (2) adds controls for permanent income: a cubic in the log of permanent income, plus log permanent income interacted with age. One concern with the identification strategy is that, because health care and savings are normal goods, high-income areas will have high savings and high Medigap prices. Controlling for permanent income addresses this concern, and since the estimate of γ hardly changes, there is little evidence that own permanent income is closely

related to Medigap demand or prices. Permanent income is perhaps an indirect measure of lifetime wealth, but controlling for wealth directly would be inappropriate since wealth is an outcome of interest. Nonetheless, in results not reported here, I control for wealth at age 65 rather than permanent income; the estimates are very similar. In the robustness section below, I also control for pre-retirement wealth, and again the results are confirmed.

The next two columns control for health and then preferences more directly. Column (3) adds indicators for health status (excellent, very good, good, fair, or poor), as well as indicators for ever having one of several chronic diseases or medical conditions (high blood pressure, diabetes, cancer, lung disease, heart disease, stroke, or arthritis). Column (4) controls explicitly for the HRS's measures of risk preferences and mortality expectancy. The risk preference measures are the average value of risk tolerance implied by a person's answer, over several rounds, to questions about large hypothetical gambles (Barksy et al., 1997.) Controlling for risk preferences can be important if area-level variation in risk preferences induces low Medigap prices while simultaneously inducing high annuitization. Fang, Keane, and Silverman (2008) demonstrate that advantageous selection is an important feature of the Medigap market. As columns (3) and (4) show, the first stage results are virtually unchanged by introducing these controls.

The first stage results are therefore highly robust to controlling for individual-level factors that might be correlated with insurance prices while directly affecting savings. The main results are similarly robust, as I show below. However, to the extent that permanent area-level differences in preferences or health may be driving differences in prices across areas, the most direct control is simply to include state (or zip code) fixed effects. As Maestas, Schroeder and Goldman (2007) show, however, most of the geographic variation in Medigap prices is between

(rather than within) states; controlling for state fixed effects eliminates almost all the variation in prices. Controlling for state fixed effects shows this clearly: the coefficient on log price falls and its standard error increases. These state fixed effects eliminate the first stage relationship because there is no variation remaining to (precisely) identify γ . Indeed, if I regress the instrument on the other controls, I obtain an R^2 of 0.91, but when I include state fixed effects, the R^2 rises to 0.97, implying that the state fixed effects absorb about two-thirds the remaining variation in the instrument. Since the first stage F-statistics are so low, state fixed effects cannot be used.

3.3 Assessing instrument validity

The first stage results are essentially unchanged when additional controls are included for permanent income, health status, and preferences. This robustness provides a first hint that the instrument may be valid: since these controls do not affect the first stage estimate of γ , either they are uncorrelated with Medigap coverage, or they are uncorrelated with $\ln price$. To the extent that observable determinants of savings and annuitization are uncorrelated with the instrument, it is likely that the unobservables are also uncorrelated.

Table 3 presents direct evidence that the instrument is uncorrelated with these additional controls. Each column of the table reports estimates from a regression of $\ln price$ on the indicated controls; F-statistics test that the indicated set of controls is jointly equal to zero. Column (1) again has only the basic controls; here we see, as expected, that the year and age controls are highly correlated with Medigap prices. The race indicators and education year fixed effects are also significantly correlated. Sex is, somewhat surprisingly, uncorrelated with price, while the children dummies are at most weakly correlated. Column (2) adds controls for permanent income, which is correlated with the instrument. In regressions below, controlling for permanent income does not affect results because permanent income turns out not to be closely

correlated with annuitization, conditional on the demographic variables. Column (3) adds controls for self-reported health status and for chronic diseases. Health appears uncorrelated with prices. Finally, Column (4) controls for self-assessed mortality and risk tolerance. These preference variables are also uncorrelated with the instrument.

In sum, the instrument is largely uncorrelated with health, or preferences, conditional on basic demographics and permanent income, although there is some correlation between local prices and indicators for having chronic disease. These results may ease concern that the instrument might be systematically correlated with unobserved determinants of savings or annuitization.

4. Medigap coverage increases annuitization

I now turn to estimating the impact of Medigap coverage on annuitization. I begin with income from private (i.e. non-Social Security) pension and annuity income, looking first at the probability of having any pension or annuity income, and then the amount of yearly pension and annuity income. When I look at the amount of pension income, I work in levels despite the skewness of the data because many people have zero income. The results in logs are highly imprecise. I exclude Social Security income because over 90% of the sample has positive Social Security income, making it difficult to detect any effect on the extensive margin. Nonetheless, people can increase their Social Security income by claiming later (as well as other mechanisms), so I examine the impact of Medigap coverage on Social Security income receipt.

Table 4 shows the impact of Medigap coverage on the probability of having positive pension or annuity income, estimated with a linear probability model.⁹ Panel A provides the reduced form result, where the coefficient on $\ln price$ is obtained by regressing an indicator for

⁹ The marginal effects of $\ln price$ implied by probit estimates (not reported) are very similar to the reduced-form coefficients, however.

positive pension or annuity income on $\ln price$ and the indicated controls. The coefficient on $\ln price$ is a statistically significant -0.20 , indicating that increasing the price of Medigap coverage by 10 percent would reduce the annuitization rate by about two percentage points. Figure 2, constructed analogously to Figure 1, illustrates the reduced form relationship between annuitization coverage and prices: it shows the nonparametric fit between (the residual) annuitization probability and (residual) prices (net of controls). There is a clear negative relationship, although it is not as precise as the first stage relationship.

Panel B of Table 4 presents the instrumental variables estimate of α , the impact of Medigap coverage on annuitization, estimated via two stage least squares with $\log price$ as the excluded instrument. The coefficient is 0.41 . To interpret this number, it is helpful to calculate the elasticity of (aggregate) annuitization it implies. With mean Medigap coverage of 0.33 and mean annuitization probability of 0.35 , the elasticity of annuitization with respect to Medigap coverage is about 0.39 . A policy change that increases Medigap coverage by 10 percent would increase the probability of having private pension or annuity income by about 4 percent.

Columns (2), (3), and (4) include successively more controls, and results change very little. The point estimates fall by two percentage points when controlling for permanent income, and a further two when controlling for the health variables. Controlling for preferences has no further impact on the estimate. This suggests that Medigap coverage appears to have a large and robust positive impact on annuitization.

These results, of course, reflect only the extensive margin, that is, whether people have any private annuity income at all. Equally important is the intensive margin—the amount of income derived from private annuities and pensions. Table 5 shows the impact of Medigap coverage on average annuity income. Increases in annuity income reflect both intensive and

extensive margin responses and therefore summarize the overall impact of Medigap on the demand for annuity income. Figure 3 illustrates this relationship graphically, showing the relationship (net of age, year, and demographic controls) between private pension or annuity income and $\ln price$.

The reduced form results in Panel A indicate a strong inverse relationship between Medigap prices and annuity income. Increasing Medigap prices by 10 percent reduces annuity income by about \$550, relative to an annual mean of about \$6,300. This number is fairly consistent across specifications, although including health or preference controls reduces it slightly. The instrumental variables results in Panel B yield a point estimate of \$9,500-11,600, with a fairly wide confidence interval. Since mean annuity income is about \$6,200, these point estimates imply an overall elasticity of annuity income with respect to Medigap coverage of about 0.56. Comparing this to the extensive margin elasticity implies that about perhaps seventy percent of the impact of Medigap coverage on annuitization comes from the extensive margin, and the remainder from the intensive margin.

Figure 4 turns to the impact of Medigap coverage on Social Security income, plotting the reduced form relationship between Social Security income and Medigap prices. The figure shows a negative relationship between Medigap prices and Social Security income, especially at high prices. Table 6 quantifies this relationship. The reduced form results in Panel A of table 6 indicate an average slope of -1,900, with additional controls attenuating this relationship somewhat. The instrumental variables results in Panel B indicate that Medigap coverage increases Social Security income by \$3,600-\$4,000. The implied elasticity of Social Security income is about 0.11.

Social Security benefits are a function of lifetime earnings, claiming age, and work experience post-claiming.¹⁰ Since I condition on permanent income, Medigap coverage can affect Social Security income primarily by encouraging people to claim benefits at older ages. In results not detailed here, I examine the impact of Medigap coverage on the age of claiming, defined as the age in which a person first reports positive Social Security results. The results are imprecise but indicate that Medigap coverage raises the claiming age by about two-thirds of a year, relative to a mean of 64.3. During the sample period, delaying claiming raises future benefits by about 7 percent, so either the impact of Medigap on Social Security claiming is concentrated among high-income earners, or Medigap affects claiming through channels other than claiming age. In particular, the Earnings Test can affect benefits in complex ways, reducing them in years with high earnings, but increasing them after the Normal Retirement Age.

5. Robustness: Controlling for pre-determined endogenous variables

I have shown that the impact of Medigap coverage on annuity demand is robust to extensive controls for income, health, and preferences. Nonetheless, it is possible that the instrument remains correlated with unobserved determinants of savings or annuity demand. In this section, I explore the robustness of my main results to additional controls.

One threat to identification comes from a possible correlation between own medical expenditures and area medical expenditures. If own medical expenditures and annuity or savings demand are also negatively correlated, then the instrumental variables results will be biased upward. Such a correlation could arise if people with low discount factors, present-biased preferences, or high mortality expectations underinvest in health when young; when old, they would have high health expenditures and little demand for savings. If people sorted across areas

¹⁰ Friedberg (1998) provides detailed explanation of the relationship between claiming age and benefits. Because of the Social Security Earnings Test, an individual with sufficiently high earnings after claiming benefits loses her benefits in the year with high earnings, but this is offset by higher benefits after she reaches Normal Retirement Age.

on the basis of demand for medical care, then there would be a positive correlation between own health expenditures and aggregate expenditures, and hence also Medigap prices.

A natural way to address this issue is to simply control directly for individual-level medical expenditures. A problem with this approach is that medical expenditures depend on insurance coverage, so controlling for medical expenditures may absorb part of the effect of Medigap coverage. As a partial solution, I control for mean out-of-pocket expenditures prior to age 65, since these expenditures are likely unaffected by Medigap prices. It is possible that in areas with high prices, people use employer-provided health insurance more intensively immediately before retiring. This seems unlikely, however, since Card, Dobkin and Maestas (2009) show that medical care utilization *increases* discontinuously at age 65, when people become eligible for Medicare, suggesting that people typically delay treatment until after retiring. If pre-retirement demand for healthcare when young (conditional on observables) reflects the unobservable determinants of demand for healthcare while old, then these health expenditures proxy for unobservables.

A second way to address this problem is to control for pre-retirement savings, again using pre-determined endogenous variables as a proxy for the unobserved determinants of savings and annuity demand. Here assets are defined as total net assets including savings, real estate, and businesses, but not the asset value of pensions or annuity holdings. This approach suffers from the same problem as using pre-retirement health expenditures: people may save before retirement in anticipation of wanting more precautionary wealth after retirement. Because of these endogeneity problems, I view these additional control variables as robustness checks rather than preferred specifications.

Table 7 shows the results of controlling for these pre-retirement endogenous variables in addition to the basic set of controls. Column (1) presents results for the probability of having private pension or annuity income; Column (2) shows results for the level of pension/annuity income, and Column (3) reports results for Social Security income. For comparison, I also report baseline results. The point estimates are essentially unchanged for private annuities when controlling for medical expenditures; if anything, this strengthens the estimated impact of Medigap coverage on annuity income. Controlling for assets has a similarly small impact on the point estimates, and controlling for both assets and expenditures simultaneously also does not affect estimates. The point estimates for Social Security income do respond to the additional controls, falling (in absolute value) with controls for medical spending but rising with controls for pre-retirement wealth, with little net impact. Even taking the smallest point estimates (in absolute value), there still remains an important role for Medigap coverage in affecting annuity decisions. Overall, then, my findings are robust to the inclusion of these pre-retirement variables, and to the extent that these variables proxy for unobserved determinants of savings or annuitization demand, they suggest that correlation between the instrument and unobservables do not drive the results.

6. Conclusions

This paper examined the impact of Medigap coverage on annuitization and savings decisions, using local prices of Medigap coverage as an instrument for coverage. I found that that Medigap coverage has a large impact on annuitization, with an aggregate elasticity of private annuity income with respect to Medigap coverage of about 0.56. About 70 percent of this effect operates through the extensive margin. Medigap coverage also appears to induce higher Social Security income, with an elasticity of about 0.11.

While the results are robust to extensive individual-level controls for income, health, and preferences, I cannot control for state-level factors affecting insurance prices and savings demand. One such factor is that medical expenditure risk may vary at the state level because of differences in physician behavior (e.g. Gottlieb et al. 2010). To the extent that this risk variation drives annuitization or savings decisions, then our results here may be biased upward -- although they would still support the view that medical expenditure risk has a large impact on annuitization decisions. An alternative threat to identification is that state-level differences in wealth or prudence might boost the demand for medical care, driving up Medigap prices, and hence have a direct effect on annuity demand. The variables used here can only imperfectly control for these differences in lifetime wealth and preferences, and so this concern remains. Nonetheless, my results suggest that medical expenditure risk may be an important part of the underannuitization puzzle, as people hold their wealth in a liquid form in order to finance unexpected health care needs.

References

- Barsky, Robert B, Thomas F. Juster, Miles S. Kimball, and Matthew D. Shapiro. 1997. Preference parameters and behavioral heterogeneity: An experimental approach in the Health and Retirement Study. *Quarterly Journal of Economics*, 112(2): 537-579.
- Brown, Jeffrey R. 2001. Private pensions, mortality risk, and the decision to annuitize. *Journal of Public Economics*, 82, 29-62.
- Brown, Jeffrey R. 2007. Rational and behavioral perspectives on the role of annuities in retirement planning. NBER Working Paper 13537.
- Brown Jeffrey R., Jeffrey R. Kling, Sendhil Mullainathan and Marian V. Wrobel. 2008. Why don't people insure late-life consumption? A framing explanation of the underannuitization puzzle. *American Economic Review Papers and Proceedings*, 98(2), 304-309.
- Brown, Jeffrey R. and Amy Finkelstein. 2007. Why is the market for long-term care insurance so small? *Journal of Public Economics*, 91, 1967-1991.
- Brown, Jeffrey R. Arie Kapteyn, and Olivia S. Mitchell. 2011. Framing Effects and Expected Social Security Claiming Behavior. NBER Working Paper No. 17018
- Card, David, Carlos Dobkin and Nicole Maestas. 2009. Does Medicare Save Lives? *Quarterly Journal of Economics*, 124(2), 597-636.
- Carroll, Christopher D. and Andrew A. Samwick. 1997. The nature of precautionary wealth. *Journal of Monetary Economics*, 40, 41-71.
- Chai, Jingjing, Wolfram Horneff, Raimond Maurer and Olivia S. Mitchell. 2011. Optimal portfolio choice over the life cycle with flexible work, endogenous retirement, and lifetime payouts. *Review of Finance*, 15(4), 875-907.
- Davidoff, Thomas, Jeffrey R. Brown and Peter A. Diamond. 2005. Annuities and individual welfare. *American Economic Review*, 95(5): 1573-1590.
- De Nardi, Mariacristina, Eric French, and John B. Jones. 2010. Why do the elderly save? The role of medical expenses. *Journal of Political Economy*, 118(1), 39-75.
- Fang, Hanming, Michael P. Keane and Dan Silverman. 2008. Sources of advantageous selection: evidence from the Medigap insurance market. *Journal of Political Economy*, 116(2), 303-350.
- Feenberg, Daniel and Jonathan Skinner. 1994. The risk and duration of catastrophic health care expenditures. *Review of Economics and Statistics*, 76(4), 633-647.
- Finkelstein, Amy and James Poterba. 2004. Adverse selection in insurance markets: Policyholder evidence from the U.K. annuity market. *Journal of Political Economy*, 112(1), 183-208.
- French, Eric and John Bailey Jones. 2004. On the distribution and dynamics of health care costs. *Journal of Applied Econometrics*, 19(6), 705-721.
- Friedman, Benjamin and Mark J. Warshawsky. 1990. The cost of annuities: Implications for saving behavior and bequests. *Quarterly Journal of Economics*, 105(1), 135-154.
- Gottlieb, Daniel J., Weiping Zhou, Yunjie Song, Kathryn Gilman Andrews, Jonathan S. Skinner and Jason M. Sutherland. 2010. Prices don't drive regional Medicare spending variations. *Health Affairs*, 29(3), 1-7.
- Lockwood, Lee M. 2012. Bequest motives and the annuity puzzle. *Review of Economic Dynamics*, 15, 226-243.
- Maestas, Nicole, Mathis Schroeder and Dana Goldman. 2009. Price variation in markets with homogenous goods: The case of Medigap. NBER Working Paper 14679.

- Mitchell, Olivia S., James Poterba, Mark J. Warshawsky and Jeffrey R. Brown. 1999. New evidence on the money's worth of individual annuities. *American Economic Review*, 89(5), 1299-1318.
- Palumbo, Michael G. 1999. Uncertain medical expenditures and precautionary saving near the end of the life cycle. *Review of Economic Studies*, 66(2), 395-421.
- Paschenko, Svetlana. 2010. Accounting for non-annuitization. Federal Reserve Bank of Chicago Working Paper 2010-03.
- Peijnenburg, Kim, Theo Nijman and Bas J.M. Werker. 2011a. Health cost risk: A potential solution to the annuity puzzle. Working paper. Available at <http://www.kimpeijnenburg.com/research.html>.
- Peijnenburg, Kim, Theo Nijman and Bas J.M. Werker. 2011b. The annuity puzzle remains a puzzle. Working paper. Available at <http://www.kimpeijnenburg.com/research.html>
- Starc, Amanda. 2012. Insurer pricing and consumer welfare: Evidence from Medigap. Unpublished.
- Turra, Cassio M. and Olivia S. Mitchell. 2008. The impact of health status and out-of-pocket medical expenditures on annuity valuation. In John Ameriks and Olivia S. Mitchell, eds., Recalibrating Retirement Spending and Saving. Oxford University Press, 227-250.
- Yaari, Menahem E. 1965. Uncertain life time, life insurance, and the theory of the consumer. *Review of Economic Studies*, 32(2), 137-150.
- Yogo, Motohiro. 2011. Portfolio choice in retirement: Health risk and the demand for annuities, housing, and risky assets. Working paper. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1085306.

Table 1: Summary Statistics

	All	Medigap	No Medigap
Has Medigap (%)	0.33 (0.47)	--	--
Price of Medigap (dollars/year)	1030 (275)	980 (259)	1055 (278)
Female (%)	0.51 (0.50)	0.57 (0.50)	0.48 (0.50)
Age (%)	69.2 (2.7)	69.4 (2.6)	69.0 (2.8)
Black (%)	0.10 (0.31)	0.03 (0.17)	0.15 (0.35)
Married (%)	0.50 (0.50)	0.54 (0.50)	0.48 (0.50)
Number of children	3.32 (2.11)	3.06 (1.96)	3.45 (2.17)
Excellent/very good health (%)	0.36 (0.48)	0.42 (0.49)	0.33 (0.47)
Ever high blood pressure (%)	0.60 (0.49)	0.61 (0.49)	0.60 (0.49)
Ever diabetes (%)	0.24 (0.43)	0.20 (0.40)	0.25 (0.44)
Ever cancer (%)	0.16 (0.37)	0.16 (0.36)	0.17 (0.37)
Ever lung disease (%)	0.14 (0.35)	0.11 (0.31)	0.16 (0.37)
Ever heart disease (%)	0.27 (0.44)	0.28 (0.45)	0.27 (0.44)
Ever stroke (%)	0.08 (0.27)	0.07 (0.25)	0.09 (0.28)
Ever arthritis (%)	0.68 (0.47)	0.67 (0.47)	0.68 (0.47)
Permanent income (thousands \$/year)	73.0 (93.6)	87.1 (95.9)	66.1 (91.8)
Has pension or annuity income (%)	0.35 (0.48)	0.39 (0.49)	0.33 (0.47)
Private pension/annuity inc (thousands \$/year)	6.34 (26.6)	6.76 (24.67)	6.14 (27.6)
Has Social Security Income (%)	0.95 (0.22)	0.98 (0.13)	0.93 (0.25)
Social Security Income (thousands \$/year)	11.3 (5.40)	12.5 (5.06)	10.7 (5.46)
N People	1406	348	1058
N Observations	3335	994	2341

Notes: Table shows mean and, in parentheses, the standard deviation of selected variables, for the analysis sample and by Medigap coverage status.

Table 2: First Stage Results

<i>Dependent Var.:</i> <i>Has Medigap</i>	(1) Basic	(2) Permanent income	(3) Health	(4) Preferences	(5) State FEs
Coefficient on log(Medigap price)	-0.48*** (0.06)	-0.47*** (0.06)	-0.47*** (0.06)	-0.47*** (0.06)	-0.09 (0.09)
R^2	0.16	0.18	0.19	0.19	0.26
F-Statistic	67.7	65.3	68.4	66.7	1.0
Controls for:					
Age quartic	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes	Yes
Permanent inc.		Yes	Yes	Yes	Yes
Health			Yes	Yes	Yes
Preferences				Yes	Yes
State FEs					Yes

Note: Table shows the results of regressing Medigap coverage on the log of price, plus the indicated controls. The price is the zipcode-average price of plan A coverage in effect the year of turning 65. See the text for a definition of the sample. The sample consists of 3,335 observations of 1,406 individuals. Demographic controls are a quartic in age, plus dummies for female, black, married, educational attainment (high school, some college, college or more), married, and number of children (one, two, or three plus). The permanent income control is a cubic in the log of permanent income (average income before retirement), plus an interaction between log permanent income and age. The health controls add a set of dummy variables for several chronic conditions, plus indicators for self-reported health status. The preference controls are controls for subjective mortality expectations and subjective risk tolerance. Heteroskedasticity-robust standard errors, clustered on individuals, are in parentheses. Asterisks indicate statistical significance, with *: $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table 3: Tests of instrument validity

<i>F</i> -statistic for test of joint significance:	(1) Basic	(2) Permanent income	(3) Health	(4) Preferences
Year Fixed Effects	454 (0.00)	452 (0.00)	426 (0.00)	424 (0.00)
Age trend	639 (0.00)	164 (0.00)	164 (0.00)	166 (0.00)
Children dummies	1.05 (0.37)	1.07 (0.36)	0.90 (0.44)	0.95 (0.41)
Female Indicator	5.11 (0.02)	5.74 (0.02)	4.08 (0.04)	5.10 (0.02)
Black	4.47 (0.02)	3.34 (0.07)	3.57 (0.06)	2.20 (0.14)
Educational attainment	2.14 (0.09)	1.88 (0.13)	1.73 (0.16)	1.51 (0.21)
Permanent income		4.88 (0.00)	4.70 (0.00)	4.85 (0.00)
Health Status indicators			0.67 (0.61)	0.70 (0.59)
Chronic disease indicators			1.29 (0.22)	1.33 (0.20)
Self-assessed mortality				1.82 (0.16)
Risk Tolerance				0.17 (0.68)
Controls for:				
Age quartic	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Permanent income		Yes	Yes	Yes
Health			Yes	Yes
Preferences				Yes

Note: The table is constructed by regressing log Medigap price on the indicated controls. Demographic controls are dummies for female, black, married, educational attainment (high school, some college, college or more), and number of children (one, two, or three plus). The permanent income control is a cubic in the log of permanent income (average income before retirement), plus an interaction between log permanent income and age. The health controls add a set of dummy variables for several chronic conditions, plus indicators for self-reported health status. The preference terms control for subjective mortality expectations and subjective risk tolerance. The price is the zipcode-average price of Medigap plan A coverage in effect the year of turning 65. The table reports F-statistics of the test that the coefficients on the indicated controls are jointly equal to zero, with p-values in parentheses. See the text for a definition of the sample. The sample consists of 3,335 observations of 1,406 individuals.

Table 4: Impact of Medigap on probability of having private pension or annuity income

<i>Dependent Var.:</i> <i>Has private pension</i> <i>or annuity income</i>	(1) Basic	(2) Permanent income	(3) Health	(4) Preferences
<i>Panel A: Reduced Form</i>				
Coefficient on log(Medigap price)	-0.20*** (0.07)	-0.18** (0.07)	-0.17** (0.07)	-0.17** (0.07)
<i>Panel B: IV Results</i>				
Coefficient on Medigap	0.41** (0.17)	0.39** (0.17)	0.37** (0.16)	0.37** (0.16)
Controls for:				
Age quartic	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Permanent inc.		Yes	Yes	Yes
Health			Yes	Yes
Preferences				Yes

Note: Table shows the linear-probability model instrumental variables estimate of the impact of Medigap coverage on the probability of having positive income from pension or annuities. The mean of the dependent variable is 0.35, and its standard deviation is 0.48. The excluded instrument is the log of Medigap price; its coefficient is reported in Panel B, and the reduced form coefficient on log price is reported in panel A. Additional controls are included as indicated. See the text for a definition of the sample. The sample consists of 3,335 observations of 1,406 individuals. Demographic controls are dummies for female, black, married, educational attainment (high school, some college, college or more), and number of children (one, two, or three plus). The permanent income control is a cubic in the log of permanent income (average income before retirement), plus an interaction between log permanent income and age. The health controls add a set of dummy variables for several chronic conditions, plus indicators for self-reported health status. The preference controls are controls for subjective mortality expectations and subjective risk tolerance. Heteroskedasticity-robust standard errors, clustered on individuals, are in parentheses. Asterisks indicate statistical significance, with *: $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table 5: Impact of Medigap on pension / annuity income

<i>Dependent Var.: private pension or annuity income (\$ thousands)</i>	(1) Basic	(2) Permanent income	(3) Health	(4) Preferences
<i>Panel A: Reduced Form</i>				
Coefficient on log(Medigap price)	-5.57** (2.74)	-5.22** (2.63)	-4.45* (2.38)	-4.68* (2.39)
<i>Panel B: IV Results</i>				
Coefficient on Medigap	11.6** (5.85)	11.2* (5.83)	9.48* (5.19)	10.35* (5.29)
Controls for:				
Age quartic	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Permanent inc.		Yes	Yes	Yes
Health			Yes	Yes
Preferences				Yes

Note: Table shows the model instrumental variables estimate of the impact of Medigap coverage on income from pension or annuities. The mean of the dependent variable is 6.34, and its standard deviation is 5.85. The excluded instrument is the log of Medigap price; its coefficient is reported in Panel B, and the reduced form coefficient on log price is reported in Panel A. Additional controls are included as indicated. See the text for a definition of the sample. The sample consists of 3,335 observations of 1,406 individuals. Demographic controls are dummies for female, black, married, educational attainment (high school, some college, college or more), and number of children (one, two, or three plus). The permanent income control is a cubic in the log of permanent income (average income before retirement), plus an interaction between log permanent income and age. The health controls add a set of dummy variables for several chronic conditions, plus indicators for self-reported health status. The preference controls are controls for subjective mortality expectations and subjective risk tolerance. Heteroskedasticity-robust standard errors, clustered on individuals, are in parentheses. Asterisks indicate statistical significance, with *: $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table 6: Impact of Medigap coverage on Social Security Income

<i>Dependent Var.:</i> <i>Social Security income</i> <i>(\$ thousands)</i>	(1) Basic	(2) Permanent income	(3) Health	(4) Preferences
<i>Panel A: Reduced Form</i>				
Coefficient on log(Medigap price)	-1.92*** (0.73)	-1.73** (0.68)	-1.70** (0.68)	-1.68** (0.68)
<i>Panel B: IV Results</i>				
Coefficient on Medigap	4.01*** (1.54)	3.72** (1.51)	3.61** (1.47)	3.59** (1.48)
Controls for:				
Age quartic	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Permanent inc.		Yes	Yes	Yes
Health			Yes	Yes
Preferences				Yes


Note: Table shows the instrumental variables estimate of the impact of Medigap coverage on the indicated outcomes. The mean of the dependent variable is 11.3, and its standard deviation is 5.40. The excluded instrument is the log of Medigap price; its coefficient is reported in Panel B, and the reduced form coefficient on log price is reported in panel A. Additional controls are included as indicated. See the text for a definition of the sample. The sample consists of 3,335 observations of 1406 individuals. Demographic controls are dummies for female, black, married, educational attainment (high school, some college, college or more), and number of children (one, two, or three plus). The permanent income control is a cubic in the log of permanent income (average income before retirement), plus an interaction between log permanent income and age. The health controls add a set of dummy variables for several chronic conditions, plus indicators for self-reported health status. The preference controls are controls for subjective mortality expectations and subjective risk tolerance. Heteroskedasticity-robust standard errors, clustered on individuals, are in parentheses. Asterisks indicate statistical significance, with *: $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table 7: Robustness to Pre-determined endogenous variables

	(1)	(2)	(3)	(4)	(5)	(6)
Controls:	Basic controls			All controls		
Outcome:	Has private pension income (%)	Private pension income (thousands \$)	Social Security Income (thousands \$)	Has private pension income (%)	Private pension income (thousands \$)	Social Security Income (thousands \$)
Panel A: Reduced Form						
Original	-0.20*** (0.07)	-5.22** (2.63)	-1.92*** (0.73)	-0.17** (0.07)	-4.68* (2.39)	-1.68** (0.68)
Add OOP expenditures	-0.20*** (0.08)	-5.94** (2.75)	-1.69** (0.74)	-0.17** (0.07)	-4.94** (2.42)	-1.55** (0.69)
Add ln(savings)	-0.20*** (0.07)	-5.58** (2.74)	-1.91*** (0.72)	-0.17** (0.07)	-4.73** (2.39)	-1.67** (0.68)
Both both	0.20*** (0.07)	-5.94** (2.75)	-1.69** (0.73)	-0.17** (0.07)	-4.97** (2.41)	-1.55** (0.69)
Panel B: Instrumental Variables						
Original	0.41*** (0.17)	11.2* (5.83)	4.01*** (1.54)	0.37** (0.16)	9.48* (5.19)	3.59** (1.48)
Add OOP expenditures	0.41*** (0.17)	12.48** (5.95)	3.56** (1.56)	0.36** (0.17)	10.6** (5.38)	3.31** (1.51)
Add ln(savings)	0.41*** (0.17)	11.67** (5.87)	3.99*** (1.52)	0.35** (0.16)	10.2* (5.44)	3.81** (1.48)
Both both	0.41*** (0.17)	12.47** (5.95)	3.56** (1.54)	0.35** (0.17)	10.8* (5.53)	3.52** (1.50)
Controls For:						
Age quartic	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes	Yes	Yes
Permanent inc.				Yes	Yes	Yes
Health				Yes	Yes	Yes


Note: Table shows the instrumental variables estimate of the impact of Medigap coverage on the indicated outcome. The excluded instrument is the log of Medigap price; its coefficient is reported in Panel B, and the reduced form coefficient on log price is reported in panel A. Additional controls are included as indicated. See the text for a definition of the sample. Demographic controls are dummies for female, black, married, educational attainment (high school, some college, college or more), and number of children (one, two, or three plus). The permanent income control is a cubic in the log of permanent income (average income before retirement), plus an interaction between log permanent income and age. The health controls add a set of dummy variables for several chronic conditions, plus indicators for self-reported health status. The preference controls are controls for subjective mortality expectations and subjective risk tolerance. Heteroskedasticity-robust standard errors, clustered on individuals, are in parentheses. Asterisks indicate statistical significance, with *: $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Figure 1: First-stage relationship between Medigap coverage and its price, in logs, net of basic controls



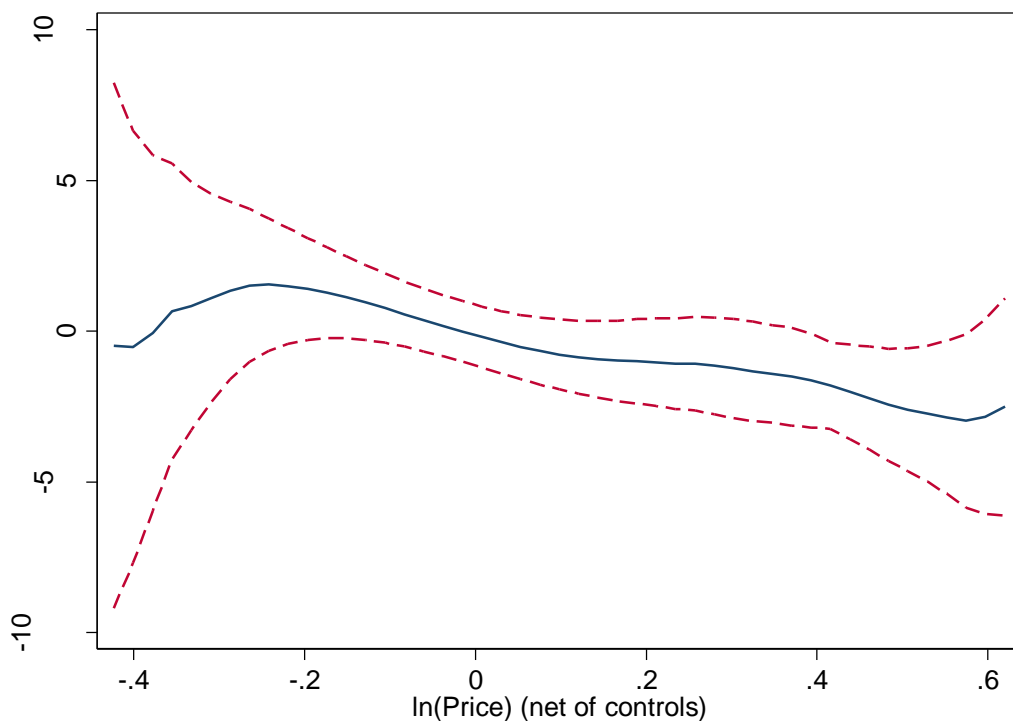
Notes: This figure plots the local polynomial (degree 1) fit between log Medigap price and Medigap coverage, net of controls for age, year, and demographics. The dashed lines show the 95% confidence interval. The figure is constructed in three steps. First, I regress Medigap coverage and log Medigap price (separately) on controls for age, year, and demographics. Second, I form the residuals from this regression. Third, I estimate a local linear regression of residual Medigap on residual log price. I plot the fitted values from this regression.

Figure 2: Reduced-form relationship between Pr(pension/annuity income) and log(Medigap price), net of basic controls




Notes: This figure plots the local polynomial (degree 1) fit between log Medigap price and an indicator for having private pension or annuity income, net of controls for age, year, and demographics. The dashed lines show the 95% confidence interval. The figure is constructed in three steps. First, I regress the indicator and log Medigap price (separately) on controls for age, year, and demographics. Second, I form the residuals from this regression. Third, I estimate a local linear regression of the residual indicator on residual log price. I plot the fitted values from this regression.

Figure 3: Reduced-form relationship between pension/annuity income and log(Medigap price), net of basic controls



Notes: This figure plots the local polynomial (degree 1) fit between log Medigap price and private pension/annuity income, net of controls for age, year, and demographics. The dashed lines show the 95% confidence interval. The figure is constructed in three steps. First, I regress private pension/annuity income and log Medigap price (separately) on controls for age, year, and demographics. Second, I form the residuals from this regression. Third, I estimate a local linear regression of residual private pension/annuity income on residual log price. I plot the fitted values from this regression.

Figure 4: Reduced-form relationship between Social Security income and $\log(\text{Medigap price})$, net of basic controls



Notes: This figure plots the local polynomial (degree 1) fit between \log Medigap price and Social Security income, net of controls for age, year, and demographics. The dashed lines show the 95% confidence interval. The figure is constructed in three steps. First, I regress Social Security income and \log Medigap price (separately) on controls for age, year, and demographics. Second, I form the residuals from this regression. Third, I estimate a local linear regression of residual Social Security Income on residual \log price. I plot the fitted values from this regression.