Preliminary. Do Not Cite or Circulate Without Permission.
Comments Welcome

Adverse Selection and the Choice of Risk Factors in Insurance Pricing: Evidence from the U.K. Annuity Market

Amy Finkelstein
MIT and NBER

James Poterba
MIT and NBER

July 2005

Abstract: We propose a new test of asymmetric information in insurance markets based on observing characteristics of the individual not priced by the insurance market. Asymmetric information is detected if there is a relationship between these unpriced characteristics and both the claims experience of the insured and the quantity of insurance purchased. This test avoids the potential confounding effects of unobserved preferences that have hampered the standard empirical test for asymmetric information. It can also provide a direct test of adverse selection, as distinct from moral hazard. We implement the test using a new data set of annuities in the United Kingdom which contains information on the annuity buyer's place of residence, a characteristic not used in pricing annuities. We find evidence of asymmetric information: individuals who live in higher socio-economic status neighborhoods self-select into annuity contract that provide greater insurance; they are also higher risk (longer lived). External evidence that socio-economic status and longevity are positively correlated, independent of insurance coverage, allows us to further interpret these findings as direct evidence of ex-ante adverse selection, as distinct from moral hazard. Our findings also raise broader questions about how insurance companies select the set of relevant individual attributes that they use in setting prices.

We thank Richard Disney, Carl Emmerson, Casey Rothschild, Michael Orszag and Michael Wadsworth for helpful discussions, Hui Shan for outstanding research assistance and the National Institute of Aging and the National Science Foundation (Poterba) for financial support. We are particularly grateful to the employees at the company who provided us with the data and patiently answered numerous questions about it.

Theoretical and empirical work has long recognized that asymmetric information can undermine the efficient operation of insurance markets. However, the empirical importance of asymmetric information in specific insurance markets remains a subject of considerable controversy. The standard empirical test for asymmetric information examines the correlation between insurance purchases and risk type, concluding that a positive correlation is evidence of asymmetric information. This test has been widely applied in a number of different insurance markets, including health insurance, life insurance, annuities, and automobile insurance – often with conflicting findings across different markets.

Recent studies of insurance markets, however, including de Meza and Webb (2001), Chiappori et al. (2055) and Finkelstein and McGarry (2003), have shown both theoretically and empirically that a positive correlation between insurance quantity and risk of loss is neither a necessary or sufficient condition for private information about risk type. The logic underlying the standard test works well provided individuals may differ along only a single dimension that is not observed by insurance companies, namely their risk type. However, individuals may have private information about their preferences for insurance, as well as their risk type. Depending on the correlation between unobserved risk type and unobserved preferences, the standard "positive correlation" test can produce either false positives or false negatives.

In this paper, we therefore develop and implement a more refined test for the presence of asymmetric information about risk type that avoids the potential confounding effect of unobserved preferences. We argue that the identification of individual characteristics that are not used in pricing insurance, but that are correlated with both demand for insurance and with subsequent risk experience provides evidence of private information about risk type in the insurance market. An additional step also allows us to use the test as a direct step of adverse selection, as distinct from moral hazard. By also drawing on external information on the relationship between individual characteristics and risk type *not due to an intervening effect of insurance*, we can interpret evidence that individuals with certain unpriced characteristics select more insurance as direct evidence of adverse selection per se.

Distinguishing empirically between moral hazard and adverse selection is both conceptually important and empirically challenging. The standard "positive correlation" test for asymmetric information cannot identify selection distinct from moral hazard (Chiappori, 2000), nor can more recently developed tests of asymmetric information (Chiappori et al. 2005). However, these two very different forms of asymmetric information – one based on ex-ante private information about risk type and the other on ex-post private information about behavior – have very different implications for public policy. Through the power to mandate insurance coverage, the government can potentially ameliorate the inefficiencies produced by adverse selection (see e.g. Akerlof, 1970); however, the government tends not to have a comparative advantage in redressing the inefficiencies caused by moral hazard. Being able to detect adverse selection distinctly is therefore of considerable interest. Our direct test for adverse selection provides a complement to other recent approaches designed to detect moral hazard using dynamic panel data with exogenous price changes (e.g. Abbring et al., forthcoming) or to distinguish between selection and moral hazard using a field experimental approach (Karlan and Zinman 2005).

We apply our test to the U.K. annuity market, using a new data set on the policies issued by one of the large insurance companies operating in this market. The annuitant characteristic that we focus on is residential location, as indicated by the annuitant's postcode. It is observed by the insurance company but not used in pricing. We show that conditional on the annuity buyer's age and gender, the two attributes that are used to price annuities, an annuity buyer's postcode contains information about her socio-economic status that helps to predict future mortality experience (i.e. risk type). We also find that, conditional on age and gender, annuitants in postcodes with higher socio-economic status tend to purchase larger annuities, which establishes the link between this attribute and the quantity of insurance purchased. These two findings indicate the presence of asymmetric information in this insurance market. Moreover, together with external evidence that the relationship between socioeconomic status and mortality exists independent of annuity coverage, they provide direct evidence of adverse selection in particular. Our findings also suggest that at least some of the private information in this market along which selection occurs is related to socio-economic status.

The paper is divided into five sections. The first outlines the test that we develop and places it in the context of previous research. Section two describes the U.K. annuity market and the dataset that underlies our analysis. The third section implements our test with these data and presents our empirical findings. Section four discusses the puzzle raised by our findings of why, in the absence of regulatory restrictions, insurance companies choose not to condition policy prices on factors such as place of residence. A brief conclusion discusses the implications of our finding of socio-economic selection in annuity markets for equilibrium in other insurance markets.

1. Testing for Asymmetric Information in Insurance Markets

1.1 The standard positive correlation test

The standard test for asymmetric information in insurance markets searches for a positive correlation between the amount of insurance an individual buys and his ex-post risk experience, conditional on the observable characteristics that are used in pricing insurance policies. Cawley and Philipson (1999) and Chiappori and Salanie (2000) point out that many models of both adverse selection and moral hazard predict that, in equilibrium, those with more insurance should be more likely to experience the insured risk. With moral hazard, insurance coverage lowers the cost of an adverse outcome and thus increases the expected loss. With adverse selection, the insured is assumed to have *ex-ante* better information than the insurance company about risk type; since the marginal utility of insurance at a given price is increasing in the risk of the insured event, those who know that they are high risk will select contracts with more insurance than those who know that they are low risk.

Empirical evidence on the correlation between risk type and quantity of insurance is mixed. A large number of studies find a positive correlation in health insurance markets, but not in auto insurance, although in both cases exceptions exist.[1] Studies in several countries find no evidence of a positive correlation between insurance coverage and risk occurrence in life insurance, but evidence of this positive

[1] See Cutler and Zeckhauser (2000) for a review of the extensive evidence for the positive correlation in health insurance – and Cardon and Hendel (2001) as an exception. In auto insurance, Chiappori and Salanie (2000), Dionne et al. (2001), and Chiappori et al. (2005) are all unable to reject the null hypothesis that risk type and insurance demand are uncorrelated, although Cohen (2001) finds evidence of a positive correlation.

correlation in annuity markets (Finkelstein and Poterba 2002, 2004, McCarthy and Mitchell 2003, Cawley and Philipson 1999). These findings have been interpreted as evidence of private information in annuity markets, but not in life insurance markets which insure the opposite mortality risks; annuities insure against living "too long" while life insurance insures against the risk of dying "too soon".

However, another possible explanation for the apparent disparity across the two markets is that insurance purchases are driven not only by private information about risk type but also by private information about risk preferences. Such preferences could be affected by a household's risk aversion or wealth level, both of which are likely to be positively correlated with demand for both annuities and life insurance. These variables are likely to be *negatively* correlated with the "life insurance risk" of dying early and *positively* correlated with the "annuity risk" of living a long time.

If individuals have private information about their preferences for insurance, the standard test of asymmetric information can produce misleading results. In this setting, the correlation between unobserved risk type and unobserved preferences that affect insurance demand becomes critical. If individuals with stronger preferences for insurance are also (unobservably) lower risk, preference-based selection may offset risk-based selection to produce an equilibrium in which there is no positive correlation between insurance coverage and risk occurrence; the standard test will fail to detect the presence of private information about risk type even though it exists and is likely to impair market efficiency. Similarly, if risk aversion and risk type are positively correlated, preference-based selection which prompts the more risk averse to purchase more insurance may produce an equilibrium which mimics the standard asymmetric information setting of a positive correlation between insurance coverage and risk occurrence, even if information about risk type is symmetric. Theoretical models illustrating these possibilities have been developed by de Meza and Webb (2001), Jullien et al. (2002) and Chiappori et al (2005). These analyses suggest that the well-documented positive correlation between insurance coverage and risk occurrence in annuity markets but not in life insurance may, contrary to the standard interpretation, in fact be consistent with the presence of private information about risk type in both markets, or with symmetric information in both markets.

The impact of preference-based heterogeneity is not just a theoretical possibility. In the U.S. long-term care insurance market, Finkelstein and McGarry (2003) show that individuals have private information about their preferences for insurance that is negatively correlated with risk type, resulting in an equilibrium in which insurance coverage and risk occurrence are not positively correlated, despite the presence of private information about risk type as well. In the Israeli automobile insurance market, Cohen and Einav (2005) find evidence of private information about risk aversion that is positively correlated with risk type and that reinforces the positive correlation created by private information about risk type.

1.2 An alternative test based on unused observables

The foregoing concerns with the potential for the standard test for asymmetric information to produce both false positive and false negative results underlie our development of a new test for asymmetric information. Our test builds on Finkelstein and McGarry (2003) who observe that, with symmetric information, conditional on the risk class in which the insurance company places an individual, the econometrician should not be able to observe *any* individual characteristic that is not used by the insurer but that is correlated in *any* direction with both insurance coverage and risk of loss. Therefore the null hypothesis of symmetric information can be rejected any time there exists an individual characteristic that is unknown to or unused by the insurer, and that is correlated with both insurance coverage and ex-post risk of loss after conditioning on the information that the insurer uses in setting prices. Such informational disparities have negative efficiency consequences, regardless of the sign of the correlation between the preference for insurance and risk type, since the private information of the buyer affects the payoff to the insurance seller (Chiappori et al, 2005).

The test does not require the uniquely rich and specialized data used by Finkelstein and McGarry (2003), who use individuals' own assessment of their risk type as the unused observable. Indeed, the test can have much broader application by exploiting the fact that the econometrician frequently observes characteristics about the insured that are not used in pricing. Examples include geography and gender in long-term care insurance, location and wealth in mortality-related insurance markets such as life insurance and annuities, and geography and occupation in automobile insurance. This allows us to apply the test for a correlation between unused characteristics on the one hand and insurance coverage and risk occurrence on the other hand in many different insurance markets, without requiring special survey data on beliefs. In addition to greater ease of application, utilization of more standard characteristics such as gender or geography instead of individuals' own risk assessments, also sheds light on the nature and source of private information in the market. It also raises the question of why insurance companies forego the use of observable, risk-related characteristics in insurance pricing; we discuss some possible explanations and their likely importance in Section four.[2]

Our course, the test based on unused observables is only one-sided. Failure to find such characteristics might simply reflect a lack of sufficiently rich data, rather than the absence of asymmetric information. However, evidence that such characteristics exist provides more compelling evidence of asymmetric information than the existing positive correlation test. Our test therefore provides a useful complement to Chiappori et al.'s (2005) recent proposals of more robust tests for asymmetric information than the standard positive correlation test. While the one-sided nature of our unused observables test represents an important limitation, the test also offers several advantages over the alternative tests developed by Chiappori et al. (2005). In particular, our test can be used to examine selection across multiple insurance contracts rather than only pairwise comparisons of contracts. In addition, our test does not require us to make assumptions about the cost structure of the insurance company.

Perhaps most importantly, the unused observables test can be used to test directly for adverse selection as distinct from moral hazard by selecting characteristics of the individual which are not priced by the insurance company but for which we have additional external evidence of their correlation with risk type. Such individuals face the same premium for insurance but different *effective prices* since the markup relative to the actuarially fair premium will vary with the unpriced, risk-related characteristics. By

[2] Of course, in some cases insurers may be prohibited by regulation from using certain characteristics; cases in which regulators create asymmetric information by prohibiting the use of pricing on certain factors are relatively uninteresting settings to *test* for asymmetric information. However, many cases, including the above examples in the long-term care insurance and annuity markets, do not involve regulatory restrictions.

looking at whether individuals who face ex-ante different effective prices due to unpriced components of their risk type select different quantities of insurance, we can test directly for selection.

A key aspect of the ability to test separately for adverse selection is the availability of external information on the relationship between these unpriced characteristics and risk type. For example, simply observing in insurance data that a characteristic not priced by annuity companies is positively correlated with insurance quantity purchased and ex post risk occurrence is not sufficient to distinguish selection from moral hazard; these data are consistent both with individuals selecting more insurance because they face an ex ante lower price (selection) and the purchase of more insurance changing behavior so that more risk is incurred (moral hazard). The existence of external information that certain characteristics are correlated with risk occurrence *for reasons other than insurance coverage* is key to allowing us to identify ex ante selection from the relationship between individual characteristics and insurance purchases. Of course, such evidence does not preclude the existence of moral hazard as well. [3]

2. U.K. Annuity Market and Data

2.1 The UK Annuity Market

We implement the test described in the previous section using insurance company data from the I.K. annuity market. Annuities are insurance contracts that pay a pre-specified payment stream to the insured as long as he is alive. They provide a way of spreading an accumulated stock of resources over a retirement period of uncertain length and thus provide insurance against the risk of outliving one's resources. Mitchell, Poterba, Warshawsky, and Brown (1999) show that annuities have the potential to play an important welfare-improving role for retirees, and are an important component of public and private defined benefit pension systems. In light of this potential, the small size of private voluntary annuity markets in many countries has been something of a puzzle. Brown et al. (2001) discuss several possible explanations for the small size of the market, of which asymmetric information is one. Others

[3] In practice, selection can be either adverse or advantageous. For example, higher wealth individuals may demand more insurance but be lower risk (as in the case of life insurance). Either however may lead to negative efficiency consequences as individuals of different wealth – and hence different risk type – face the same premium and therefore cannot all face an actuarially fair price on the margin. Our discussion focuses on adverse selection since it is the more common one studied in the theoretical literature and the one we detect in our empirical work below.

include bequest motives, defined benefit pensions, and the need for buffer stock savings to pay for uninsured medical and long-term care needs.

Our analysis centers on the compulsory annuity market in the United Kingdom. Here, individuals who have accumulated savings in tax-preferred retirement saving accounts, the equivalents of IRA(s) or 401(k)s in the United States, are required to annuitize a large portion of their accumulated balance. In 1998 (the end of our sample period), annual annuity payments to annuitants in the compulsory market totaled £5.4 billion, according to the Association of British Insurers (1999). Although participation in this market is mandatory for the relevant individuals, annuitants in the compulsory annuity market have considerable discretion in the amount that they annuitize and in the timing of their annuitization. They can also choose among a number of annuity options that affect the effective quantity of insurance in the annuity contract (Finkelstein and Poterba 2002).

From the perspective of an insurance company, a high-risk annuitant is a long-lived annuitant. Such an individual will receive payments from the company for a longer time than expected. Annuity companies therefore charge higher prices to individuals who exhibit observable characteristics that are correlated with longer life expectancy. There are currently no restrictions in the U.K. on the characteristics that may be used in pricing annuities. However, during our sample period, the vast majority of annuities, including all of the ones sold by our particular company, were priced solely on the annuitant's age and gender at the time of purchase (Ainslie 2000).

2.2 Insurance Company Data and Descriptive Statistics

There is existing evidence of a positive correlation between the amount of insurance in the contract and the ex-post risk experience of the annuitant (Finkelstein and Poterba 2004). These findings have been interpreted as evidence of asymmetric information in annuity markets but, as discussed above, the potential for additional unobserved heterogeneity in preferences for annuities makes these findings neither necessary nor sufficient for the existence of asymmetric information.

Our proposed, more refined, test searches for characteristics of the annuitant that are not used in pricing but that are correlated with both annuity choice and risk type (i.e. ex-post mortality). To implement a test for adverse selection per se, we desire an unpriced annuitant characteristic for which there exists a documented relationship with mortality operating for reasons other than through the intervening effect of annuity choice. A natural choice is some measure of socio-economic status. Insurance companies do not use measures of socioeconomic status in pricing annuities. Moreover, the Office of National Statistics (1997), Attanasio and Emmerson (2001), and others document a positive correlation between socio-economic status and longevity in the United Kingdom population. Similar evidence of a positive socioeconomic status – longevity gradient exists in the United States as well (see e.g. Attanasio and Hoynes, 2000). More generally, there is a well-documented positive correlation between social economic status and health in many countries (including the United Kingdom); longevity insurance (i.e. annuities) are not considered an important causal factor (see e.g. Smith 1999).

To carry out this test, we obtained data from a large U.K. annuity company consisting of all of the company's compulsory annuities that were in force in 1998 and that were sold between January 1 1988 and December 31 1998. The company is one of the top ten annuity providers in the U.K. We observe the annuitant's date of death if he died over the six-year period between January 1 1998 and February 29, 2004. We also observe detailed information on the type of annuity purchased, and the characteristics of the annuitant used in pricing the annuity. These are the date of purchase, date of birth, and gender of the annuitant. Most importantly, we observe a characteristic of the individual not used in pricing: the individual's post code, or geographic location. We use these data to examine the relationship between an annuitant's place of residence, socio-economic status, subsequent survival experience, and annuity product choice.

For analytical tractability, we restrict our sample in several ways. We limit the sample to the approximately sixty percent of annuities that insure a single life. The mortality experience of the single life annuitant provides a convenient ex-post measure of risk type; measuring the risk type of a joint life policy which insures multiple lives is less straightforward. We also restrict the sample to the approximately eighty percent of annuitants who hold only one annuity policy, since characterizing the features of the total annuity stream for individuals who hold multiple policies is more complicated. We restrict attention to the approximately ninety percent of policies sold in England or Wales; we cannot map postcodes in Scotland into the same type of geographic unit that we can for England and Wales. Finally, we exclude the annuitants who purchased annuities before age 50, and those who purchased annuities with guarantee periods of between one and four years, or between six and nine years; these exclusions affect less than one percent of our sample.

Our final sample consists of 52,826 annuitants. About half of our sample consists of policies sold in the last four years (1995 – 1998) which reflects some combination of greater sales by the company in later years and less opportunity for selection out of the sample due to prior mortality. Selection out of the sample due to death prior to 1998 is non-trivial. For example, we estimate that a little under one-third of the policies sold in the first three years of our sample, 1988-90, terminated before the start of our sample. We discuss below how we account for such left-censoring in our analysis.

Table 1 presents summary information on the sample we analyze. The average age of purchase is 62, and 59 percent of the purchasers are male. When trying to infer market-wide phenomena such as adverse selection from data from a particular firm, there is always a concern with whether the firm being analyzed is representative. Most of the annuitant and annuity characteristics in our sample firm are very similar to those in the sample firm analyzed in Finkelstein and Poterba (2004); the one exception is that our current sample is less disproportionately male (59 percent compared to 77 percent).[4] Characteristics from the current sample also match the UK annuitant population as a whole where market-wide characteristics are available. For example, Murthi et al. (1999) report that the percentage of annuities in the entire U.K. annuity market that are not level nominal annuities is likely to be under 15 percent; in our sample we find that about 10 percent of policies are not level nominal annuities.

2.3 Data on annuitant place of residence

[4] The current data are from a different company than those analyzed in Finkelstein and Poterba (2004). The current data contain all of the information in the previous data set as well as information on postcode. In Table 7 below we document that the positive correlation between insurance quantity and risk type documented in the previous data set obtains in the new data as well.

Our data includes information about the annuitant's place of residence, specifically his postcode. Our analysis uses coarser geographic information than postcode since data on socio-economic status or other characteristics are not publicly available at the postcode level. We therefore map the annuitant postcode to ward-level characteristics from the 1991 U.K. Census. With about 9,000 people per ward compared to only 40 per postcode, wards are considerably larger geographic units than postcodes. As a result, our estimates using ward-level information are a lower bound on the amount of additional information the insurance company could use.

Our sample of annuitants comes from 49,123 unique postcodes and 8,941 unique wards, out of a total possible 1.24 million postcodes, and 9,527 wards in England and Wales. About 200 of these wards have twenty policies or more in them; only 4 wards contain 50 policies or more. We use the May 2004 All Fields Postcode Directory from the Office of National Statistics (ONS) to map the 1988-1998 postcodes in our dataset to the corresponding wards in the 1991 U.K. Census.

The ward-level data in the U.K. census contain two measures of socio-economic status: educational attainment and occupation. Educational attainment is reported in the percent of the ward population aged 18 and over that is "qualified"; to be "qualified", an individual has to have an educational credential above the level of the GCE A-level standard, the equivalent of a good high school degree in the United States. On average, in England and Wales, approximately 15 percent of individuals aged 18 and over in the 1991 Census are "qualified". A positive relationship between education and longevity has been well-documented in England as well as other countries (see e.g. Lleras-Muney 2003 and works cited therein).

Data on occupation consist of the percent of employed people in the ward in different occupational classes, defined as "social classes" in the U.K. data. We compare three groups: professional and managerial (group I and II), skilled manual or non manual (group III), and partly skilled or unskilled (group IV and V). On average in England and Wales, about one-third of the employed individuals are in professional and managerial occupations, two-fifths are in skilled manual or non-manual occupations, and one-fifth are in partly skilled or unskilled occupations. There is existing evidence that for the population in general, social class is positively correlated with population survival probabilities, conditional on age and gender. For example, Figure 1, which is drawn from ONS (1997), shows that the cumulative survival probabilities for males are higher at all ages for individuals in higher social classes.

In addition to these two measures of SES, the ward-level census data also contain a measure of health status: the percent of persons in the ward having a "long-term illness, health problem, or handicap which limits his/her daily activities or the work he/she can do? Include problems which are due to old age." On average in England and Wales, about 12 percent of the population reports having a long-term illness. We investigate this measure as well in our analysis below in order to assess the extent to which annuitants may have private information about factors other than socio-economic status.

Table 2 provides summary statistics for these three measures for England and Wales as a whole and for our sample. We report summary measures based on weighting each ward by its population, and based on weighting each ward by the number of policies our company has in each ward. The table indicates that the proportion of policies sold within each ward is not proportional to ward population. Our sample annuitants are on average from wards of higher SES than the general U.K. population as measured by the proportion of individuals in the ward with an education qualification or from the highest social class (professionals and managers). Our sample annuitants are thus consistent with Banks and Emmerson's (1999) findings from the U.K. Family Resources Survey that annuitants in the country as a whole are of higher SES than non annuitants. Our annuitants are also less likely to be from wards in which a higher proportion of the population has long-term illness.

A key question that our focus on ward-average measures raises is the extent to which the characteristics of the population in an individual's ward are predictive of the individual's own characteristics. We measure the predictive power of a ward-level measure, such as the percentage of the ward that is educationally "qualified," as the ratio of the variance in this measure across wards to the variance in this measure across individuals. This ratio can range from 0, when the average SES or health is the same in all wards, to 1, when there is no within-ward variation in individuals' SES or health. A higher value for the ratio indicates greater power of the ward characteristic in predicting the individual's characteristic. We find that predictive power of ward SES is 0.11 for long-term illness, 0.27 for education qualification, 0.28 for being in social class I or II, 0.16 for being in social class III, and 0.22 for being in Social Class IV or V. Thus the characteristics of the ward population convey some predictive information about the characteristics of a randomly drawn individual within the ward, but substantially less information that knowing the individual's own characteristics directly. Our estimates of asymmetric information using ward-level rather than annuitant-level SES measures therefore represent a lower bound.

3. Results: Survival, Selection, and Socio-Economic Status

This section presents the empirical results from our new test for adverse selection based on the existence of risk-related characteristics not priced by insurance companies. The test proceeds in two simple steps. First, we show that conditional on the characteristics of the annuitant used in pricing the annuity, annuitants from higher SES wards or wards with better health –characteristics not used in pricing annuities – are higher risk (lower mortality). Second, we show that conditional on the characteristics of the annuitant used in pricing the annuity, annuitants from higher SES wards or wards with better health are also purchase more insurance; this second test also provides direct evidence of adverse selection. We conclude, in the last sub-section, with a brief discussion of the implications of our findings for the form of private information in annuity markets.

3.1 Postcodes and Annuitant Survival Rates

As the first step in our two-step test, we examine whether, conditional on the other characteristics of the annuitant already used in pricing the policy, the socio-economic characteristics of the annuitant's ward contains any predictive information about their survival probability. To do so, we estimate a proportional hazard model of the length of time the annuitant lives after purchasing an annuity:

(1) $$\lambda(t, x_i, \beta, \lambda_0) = \exp(x_i'\beta)\lambda_0(t).$$

$\lambda(t, x_i, \beta, \lambda_0)$ denotes the hazard function, in our case, the probability that an annuitant with characteristics x_i dies t periods after 1998, conditional on living until t. The proportional hazard model assumption is that $\lambda(t, x_i, \beta, \lambda_0)$ can be decomposed into a baseline hazard $\lambda_0(t)$ and a "shift factor"

$\exp(x_i'\beta)$ which represents the proportional shift in the hazard caused by the vector of explanatory variables x_i with unknown coefficients β. The main covariates of interest are the socio-economic status measures of the annuitant's ward. We are also careful to control, however, for the annuitant characteristics that are used in pricing.

Following Cox (1972, 1975), we estimate a continuous-time, semi-parametric, partial likelihood proportional hazard model. This allows us to estimate the β coefficients without specifying the form of the continuous baseline hazard function $\lambda_0(t)$; we thus avoid having to impose a parametric assumption on the form of the baseline hazard, which would result in inconsistent estimates of β if the functional form were mis-specified. As described above, our data are both left censored and right censored; both types of censoring are easily handled in the Cox model.

Table 3 presents our findings. The first column shows the results controlling only for the annuitant characteristics used in pricing. The only coefficient shown is for the indicator variable identifying male annuitants; not surprisingly, mortality hazards are higher for males. The other covariates are single year- and age-specific indicator variables; to conserve space they are not reported. Columns (2) and (3) add ward-level SES measures to this basic specification. The results provide the first piece of evidence for our two-step test. They show that, conditional on annuitant characteristics used in pricing, the socioeconomic status of the annuitant's ward are statistically significantly positively correlated with annuitant survival (risk type). Column (2) indicates that annuitants from wards with a greater proportion of the population with educational qualifications have a statistically significantly lower mortality hazard; a one standard deviation (i.e. 8.1 percentage point) increase in the proportion of the annuitant's ward that is educationally qualified is associated with a decrease in the mortality hazard about one-fifth the magnitude of the decrease in the mortality hazard associated with being female instead of male. Column 3 indicates that those from wards with a greater proportion in managerial and professional occupations (group I & II) have a statistically significantly lower mortality hazard than both those in wards with a greater proportion in skilled occupations (group III) and those in our reference category (wards with a greater proportion in partly skilled or unskilled occupations, groups IV and V). Finally, column 4 indicates that annuitants from wards with a greater proportion of the population who suffer from long-term illness have a statistically significant higher mortality hazard; a one standard deviation (i.e. 3.1 percentage point) increase in the percentage of the annuitant's ward that has long-term illness is associated with an increase in the annuitant's mortality hazard that is just over 10 percent of the magnitude of the increase in the annuitant's mortality hazard associated with being male instead of female

A useful way to gauge the magnitude of the estimated coefficients is to use the estimate of the baseline hazard (not shown) to translate these mortality hazards into cumulative survival probabilities. Figure 2 shows the results from the social class measures in column (3); the higher survival probabilities for annuitants from wards that are populated by higher social classes is consistent with the survival probability differences by social class for the population at large shown in Figure 1. Table 4 translates the hazard model coefficients in Table 3 into the implied difference in the probability of dying within 5 years after 1998. We show the results separately for a 65 year old man and for a 65 year old woman who purchased an annuity in 1994 and survived until 1998. They indicate, for example, that a 65 year old male annuitant who purchases a policy in 1994 in a ward with the average proportion of qualified individuals and survives until 1998 has a 10.7 percent chance of dying within the next five years; the same individual from a ward one standard deviation above the average in the proportion educationally qualified has only a 9.7 percent chance of dying. Similarly, the same 65 year old male has a 9.3 percent chance of dying if he is from a ward that is one standard deviation above the average in the proportion from Social Class I or II, compared to a 10.7 percent chance if he is from a ward that is average in this regard.

Survival differences of this magnitude should have an important effect on the expected present discounted value of an annuity payout stream. To investigate this, we examined how much annual annuity payments would change were annuity companies in the United Kingdom to adjust prices in an actuarially fair way to account for the relationship between ward-level socio-economic status and annuitant mortality; this illustrative calculation of course ignores any demand response to such price changes on the part of the annuitant. The actuarially fair annual payment from an annuity depends on the characteristics of the annuity (we consider a nominal annuity with no guarantee), the annuitant mortality table used, and the interest rate. Since we can only estimate five years of annuity mortality in our data, for our illustrative calculation we use the annuitant mortality tables for the compulsory annuity market as a whole; Finkelstein and Poterba (2002) provide a detailed description of these data. We consider a 65 year old male and female who purchases an annuity at the end of our sample period (1998). To discount future annuity payments we again follow Finkelstein and Poterba (2002) and use the zero-coupon yield curve of the nominal U.K. Treasury securities; we select January 1 1998 as our purchase date for our sample calculation. Table 5 shows the results. It indicates for example, that relative to a male 65 year old annuitant in an average (in an annuity-weighted manner) ward in terms of the proportion of the ward that is educationally qualified, a male 65 year old annuitant in a ward one standard deviation above (below) this average would receive a 3% lower (higher) annual annuity payment. For a 65 year old woman the comparable effect is 2.4%; because women are longer lived the actuarially fair proportional adjustment to the annual payment is smaller. Similarly, the proportional change in annual payment increases with the age of purchase since there is less expected time to do the adjustment over (not shown). A one standard deviation increase in the proportion of the ward in social class 1 or 2**,** is associated with a 3.6% (2.8%) decline in the annual annuity payment for a 65 year old male (female).

To put these numbers in perspective, the average load for a 65 year old male in the compulsory annuity market in 1998 was about 5 percent (Finkelstein and Poterba 2002). Thus the changes in pricing associated with using ward-based SES in pricing are of the same order of magnitude as a 50 to 75 percent change in the load. Presumably the changes in pricing associated with using annuitant SES – rather than ward-level SES – would be even higher.

3.2 Postcodes as Predictor of Product Selection

The results in the previous section suggest that the socio-economic characteristics of the annuitant's ward are predictive of his risk type (survival), even after conditioning on the characteristics of the annuitant that are used in pricing. The second component of our test requires that we examine whether these characteristics are also correlated with the amount of insurance purchased, again conditional on the characteristics of the annuitant used in pricing. Since we also have external evidence that SES and survival are positively related even without the annuity intervening effect, evidence of a relationship between the annuitant's (unpriced) ward-level SES and his quantity of insurance purchased also provides direct evidence of adverse selection, since it suggests that people who are *ex ante* different risks self-select into different annuity contracts.

We relate insurance purchases and ward characteristics using the following regression specification:

(2) $INSURANCE_{iw} = \alpha * X_i + \beta * WARD_w + \varepsilon_{iw}$

where $INSURANCE_{iw}$ denotes the quantity of insurance purchased by annuitant i in ward w. The covariates X_i control for the annuitant characteristics that are used in pricing. As in the hazard model analysis in equation (1), these consist of indicator variables for annuitant's age at time of purchase, the year in which the policy was purchased, and the gender of the annuity buyer. The main coefficient of interest is β, which describes the conditional correlation between a ward-level characteristic (WARD) and the amount of insurance coverage purchased.

There is no single, obvious measure of the quantity of insurance in the annuity contract to use for the dependent variable. Finkelstein and Poterba (2004) focus on three features of the annuity that affect the effective quantity of insurance. The amount of insurance is increasing in the initial annual payment and in the "tilt" of the annuity payment profile, and decreasing in the guarantee period. In terms of "tilt', Table 1 indicates that ninety percent of the policies provide a constant nominal "level" payment stream (as opposed to a payment stream that increases in nominal terms over time). The guarantee period specifies the period of time after purchase during which the company will continue to make payments to the annuitant's estate even if the annuitant dies before the guarantee period expires; since annuity payments are not life-contingent during the duration of the guarantee period, the effective amount of insurance in the annuity is diminishing in the length of the guarantee. Compulsory annuitants may choose to purchase a "guarantee" of up to ten years. Table 1 indicates that eighty two percent of annuities in our sample contain a guarantee; about ninety percent of the guaranteed policies have a 5 year guarantee, and the rest have a 10 year guarantee.

We follow two strategies for measuring the quantity of insurance in a particular contract. First, we stratify the sample into sub-samples of contracts that vary on only one contract dimension. We can then look at the relationship between the socio-economic characteristics of the annuitant's ward and the amount of insurance purchased along the dimension of quantity that is free to vary. Specifically, we stratify the sample into level annuities with no guarantee, level annuities with 5-year guarantees and level annuities with 10-year guarantees. Within each of these three sub samples, we examine the relationship between ward-level SES and the quantity of insurance as measured by the log of the initial annual annuity payment. We use the log of the initial payment as the dependent variable because of the skewness in the distribution of initial payment. We were not able to study policy variation related to escalator (i.e. "tilt") clauses, because of heterogeneity in the various escalation provisions and small sample sizes for policies that are not level payment nominal contracts.

The disadvantage of this approach is that we can only examine selection on one dimension of the contract at a time, while stratifying on other (potentially endogenous) features of the contract. Our second approach to measuring the quantity of insurance therefore combines the different features of the annuity product into a single measure of insurance quantity. Conditional on the degree of "tilt" in a policy, we measure the expected present discounted value (EPDV) of the insurance component of the policy's payments. An annuity policy with a guarantee has both a bond component and an insurance component. Since we care only about the insurance component, which is the component that reflects the transferring of assets from one state of nature to the other, we subtract the present value of the bond component from the EPDV of the entire payment stream, and define this as the value of insurance. Insurance quantity is therefore defined as:

$$\text{(3)} \qquad \text{Quantity} = \sum_{t=1}^{T} \frac{AS_t}{\Pi_{j=1}^{t}(1+i_j)} - \sum_{G} \frac{A}{\Pi_{j=1}^{t}(1+i_j)}$$

where A denotes the annual annuity payment, S_t denotes the probability that the annuitant survives until period t, G denotes the number of years in the guarantee period, and i_j denotes the expected nominal short-term interest rate at time period j. Our insurance quantity measure is thus increasing in the amount of initial payment and decreasing in the length of the guarantee.

Computation of insurance quantity in equation (3) requires both a table of survival probabilities (S_t) and a term structure for discounting future payments (i_j). For each contract in our data, we use a common survival table, the U.K. population cohort mortality table provided by the Government Actuaries' Department. This mortality table provides current and projected future mortality rates by age and sex, and we find the relevant rates by gender and age for each annuitant based on the table in use in the year of purchase.[5] For the term structure of interest rates used to discount future annuity payouts, as in Finkelstein and Poterba (2002) we use the zero-coupon yield curve of nominal U.K. Treasury securities of the first day of the month and year in which the annuity was purchased; these data are provided by the Bank of England.

Table 6 reports results from estimating equation (2) using both the stratified-sample approach and the uni-dimensional measure of quantity approach. In both cases we exclude non-level annuity policies, which represent ten percent of our sample, for reasons described above. We have verified that the rest of the results in the paper are robust to this sample restriction (results not shown). The different column headings in Table 6 indicate the approach taken to measuring the quantity of insurance. The table shows the results of using both measures of ward SES (education and occupation) and ward level health.

Across all dependent variables and all ward-level measures, the results tend to suggest that individuals in wards of higher socio-economic status or better health are statistically significantly likely to purchase a greater quantity of annuity insurance; the one exception from the 12 regressions reports in Table 6 is the wrong sign on the coefficient of Social Class III when the EPDV measure of quantity is used in column (4). All the findings are statistically significant. A potential concern with these results,

[5] We use population mortality tables rather than annuitant mortality tables since annuitant mortality tables are updated about once every decade, not annually as population mortality tables are.

however, is that our sample of policies is left-censored; the annuitant must survive from purchase until 1998 for his policy to be included in the sample. While such left-censoring is easily handled in the hazard model analysis in Table 3, it may bias the analysis of linear regression models in Table 6. We therefore verified that our results are robust to limiting the sample to the approximately 13 percent of policies sold in 1998, for whom the issue of left censoring does not arise (results not shown).

While statistically significant, the magnitude of the relationship between ward-level characteristics and annuity quantity is modest. For example, the results indicate that compared to an annuitant from a ward in which no one is educationally qualified, an annuitant from a ward in which 100 percent of individuals are educationally qualified purchases an annuity that is 1.6 percent to 2.7 percent larger. A one-standard deviation, or 8.1 percentage point, increase in the proportion of the annuitant's ward that is qualified is thus associated with only a 0.13 to 0.22 percent increase in the quantity of annuity purchased. Results using the occupational measure of socio-economic status or the percentage of the population suffering from long-term illness are similarly small in magnitude.

3.3 Insights on the Source of Asymmetric Information

Together, the results from the previous two sections provide evidence of asymmetric in the U.K. annuity market: socio-economic characteristics of the annuitant's ward that are positively correlated with risk type (survival probability) are also positively correlated with the quantity of insurance purchased. The fact that the external literature tells us that the positive relationship between socio-economic status and survival is not simply due to differences in insurance coverage by SES allows us to interpret to relationship between annuitant's ward SES and quantity of insurance purchased more specifically as evidence of adverse selection: individuals who ex ante face lower effective prices (because they are higher risk – i.e. higher SES – than the insurance company's prices based on only age and gender assume) self-select into annuity contracts with more insurance.

Our findings also provide some information about the *form* of the private information in annuity markets. The correlation between ward-level socio-economic status and annuity demand suggests that some of the selection into annuity markets may reflect selection based on socio-economic status. This may reflect “active” or “adverse” selection by prospective annuity buyers who recognize the implications of their socio-economic status for their risk type, or it could reflect “passive” or “preference-based” selection if socio-economic status directly affects demand for insurance. Either mechanism produces the same negative efficiency consequences. The evidence on long-term illness is consistent with some of the selection occurring along the standard, active dimension of private information about risk type, as long-term illness is less likely to be a marker for preferences independent of risk type.

Ward-level health and socio-economic characteristics are highly correlated, so it is difficult to determine the relative importance of these two types of selection factors. When we include the ward-level health measure together with a socio-economic measure in the hazard model analysis or quantity analysis, the coefficient on the health measure tends to attenuate substantially in size and lose statistical significance, while the coefficient on the socio-economic measure remains largely unchanged in magnitude and retains statistical significance (results not shown).

As discussed, previous research as demonstrated a positive correlation between annuity quantity and ex-post risk type, conditional on the risk classification done by the insurance company (Finkelstein and Poterba, 2004). It is interesting to consider whether basically all of the selection on product choice within the annuity market can be explained by these ward level health and SES characteristics, or whether there are likely to be other unobservable factors along which selection occurs. Table 7 provides some results. The first column shows that the basic “positive correlation” between insurance coverage and risk occurrence appears in our current data as well. The analysis is a Cox proportional hazard model of annuitant mortality with covariates consisting of characteristics of the individual used in pricing and the annuity product characteristics. As with the data in Finkelstein and Poterba (2004), these data show clear evidence of the positive correlation property. Individuals who purchase annuities with a larger initial payment (i.e. more insurance) are statistically significantly lower mortality (higher risk). The amount of insurance is smaller for guaranteed compared to non-guaranteed annuities and for level compared to backloaded annuities; people who purchase these annuities also have higher mortality (lower risk), although the results are only statistically significant for the guarantee.

The remaining columns of Table 7 add controls for the socio-economic characteristics of the annuitant's ward or the health of the ward to the analysis in the first column. These ward-level characteristics remain statistically significant even when controls for product characteristics are included. However, the addition of ward-level characteristics into the analysis does little to attenuate the positive correlation between dimensions of the insurance contract that provide additional coverage and ex-post risk type. This suggests that there are other unobserved characteristics of the annuitant along which selection is occurring that we are not capturing with our ward-level measures of SES and health.

4. The Determinants of the Attributes Used in Pricing Insurance Policies

Our empirical results suggest that U.K. annuitants have private information about mortality risk, conditional on the individual attributes that are currently used in pricing annuities. These findings raise a broad and important question of how insurers determine which variables they will use in setting insurance prices. In particular, why do insurance companies not price on the basis of geographic location if this is an important predictor of demand and of risk type? This issue does not arise only in the U.K. annuity market, but in many other insurance markets as well in which insurance companies appear to forgo pricing on characteristics that are likely to be highly correlated with risk type. In the U.S. long-term care insurance market, for example, Brown and Finkelstein (2004) find that premiums do not vary across locations or by gender, even though there is substantial geographic variation in the cost of nursing homes and substantial gender differences in expected nursing home utilization, and no regulatory restrictions on pricing based on these characteristics.

One reason insurance companies may not use all of the easily available, relevant information in pricing policies is that the asymmetric information that would be eliminated by the use of such information is not quantitatively important. If it were, standard profit maximization arguments would suggest that firms should not forgo seemingly simple ways of reducing the amount of this informational asymmetry. While this may explain why some characteristics are not used in pricing, the results in Section 3.1 suggested that the association between ward-level SES and annuitant mortality implied non-trivial changes in annuity payments – on the order of 5 percent for a one standard deviation change – if these characteristics were used in pricing. Presumably the relationship between annuitant SES and annuitant mortality is even larger. Similarly, Brown and Finkelstein (2004) document that the unisex pricing of long-term care insurance together with risk differences by gender generates a substantial – approximately 50 percentage point – disparity in the loads in this market by gender.

In this section, we therefore briefly discuss several other potential factors that may reduce the amount of information insurance companies use, even in situations in which asymmetric information has a quantitatively important effect on market equilibrium or insurer profits. While our discussion of these factors focuses on the U.K. annuity market, where we have conducted a number of in-depth interviews with companies involved in designing annuity pricing, many of these factors are general in nature and likely apply to other insurance markets as well.

One possibility is that the predictive content of place of residence is limited by the extent to which this characteristic is mutable, and the extent to which introducing pricing differentials on the basis of this characteristics provides incentives to change the characteristic. For a sufficiently large difference in expected annuity payments based on geographic location, would-be annuitants might try to arrange for a different mailing address, or perhaps even to move. While this explanation may be relevant for the use of postcodes in pricing annuities, it seems unlikely to be a general answer since there are other less easily mutable characteristics – such as gender in the long-term care insurance market or educational attainment in the annuity market – that are also not used in pricing.

The incentive for an insurance company to begin using a previously-unexploited buyer characteristic for pricing may be tempered by the fact that considerable up-front investment is required to determine the appropriate pricing structure. Indeed, in studying the development of impaired life annuities in the U.K., which offer substantial discounts to smokers or other individuals likely to be in poor health, we learned that the initial pricing of these products involved both considerable fixed investments and considerably risk and uncertainty. For example, one of the new firms engaged a large reinsurer and used their information on the medical records from life insurance sales around the world to try to predict the relationship between various medical conditions and mortality in the annuity pool. Another one of the new firms contracted with one of the U.K. health authorities for their data on the mortality of individuals in nursing homes and hospitals and then spent several months analyzing these data to try to derive the relationship between mortality and health conditions. Even with these efforts, there was considerable uncertainty surrounding initial estimates of the prices at which impaired life annuities would break even.

Interestingly, the firms that introduced new impaired annuity products do not appear to have been concerned about other firms free riding on their pricing decisions, without paying the costs of determining the appropriate pricing structure. This seemingly natural potential explanation for the limited use of potentially relevant characteristics in insurance pricing may therefore not, in practice, be an important factor. We learned in our discussions with individuals who started the "impaired life" annuity business that one of the incentives to enter this market early is that by selling these products, the company builds up a statistical experience base which it can then use to refine its pricing. This offers an informational advantage relative to later-entering competitors. More importantly, simply observing the new policy's pricing structure will not enable a competitor to mimic the innovator's risk pool. Potential imitators will not know the innovator's underwriting rule, and in particular which applicants they deny coverage. Firms that seek to emulate the innovator by introducing policies with similar pricing may therefore suffer from a form of the winner's curse in which individuals who were denied policies by the innovator obtain insurance from the emulator. This could reduce the profitability of emulation.

A comparison of pricing practices in annuity markets and in life insurance markets – which insure the opposite risk but use much more extensive financial and medical information in pricing – suggests that insurance policy pricing reflects the relative costs and benefits of using additional information. It is likely to be considerably more costly to verify socio-economic information in the annuity market than in the life insurance market. Annuities make many payments to their beneficiaries, while life insurance policies typically make only one. Therefore any verification in an annuity contract must be done up front and continuously, while such verification in a life insurance contract would only need to be done once. It is also likely to be more difficult to verify the veracity of reported socio-economic information in the annuity market than in the life insurance market. For annuities, the consumer has an incentive to try to hide some wealth to reduce the insurance company's belief about his survival prospects, whereas for life insurance, the consumer would want to fabricate higher wealth holdings to increase the insurance company's belief about his survival prospects. It may be easier to hide assets than the fabricate them.

The benefits to pricing based on additional characteristics may also be lower in annuity markets than in life insurance markets. Second-degree price discrimination, screening buyers based on contract design, can serve as a partial substitute for third-degree price discrimination, in which prices are set based on observable characteristics. The availability of the former reduces the net profits that accrue from adopting the latter. There is evidence of such second-degree priced discrimination in annuity markets. For example, Table 2 indicates that individuals in equilibrium are sorted across insurance contracts based on their underlying risk type. However, Cawley and Philipson (1999) report that there appears to be no such sorting in life insurance markets.

Finally, political economy issues may also play an important role in affecting the choice of characteristics used in pricing. Introducing additional price distinctions can have large public relations costs. Surprisingly, this appears to be the case even in a market in which it is the wealthy who will face the higher prices. Several U.K. insurers recently considered using annuitant postcodes to condition payouts. There was a sharply negative public reaction to such proposals, illustrated by newspaper stories on "Postcode Prejudice" (*Sunday Times*, July 13 2003), and "Postcode Peril" (*Manchester Evening News* July 7, 2003). Insurance firms may be concerned about the direct costs of negative publicity, as well as by the prospect of triggering new regulatory initiatives in the largely unregulated annuity market.

Political economy concerns may have less impact on small firms or new entrants who do not internalize the costs of increased regulation or lost public good will to the same extent that large existing firms do. The recent introduction in the U.K. of "impaired life annuities" is consistent with this hypothesis. Ainslie (2000) reports that these products were originally introduced by new companies formed expressly for the purpose of offering the impaired annuity products to individuals in observably poor health (i.e. good risks from the annuity company perspective). Incumbent firms did not follow suit, until, about 5 years after the introduction of these products, the impaired life market had grown to the point where the cream skimming of good risks by the impaired life companies created pressure on the existing companies to expand their pricing system.[6] The eventual response of other companies to the new rating factors illustrates another likely deterrent to introducing additional rating factors. If other firms feel compelled to follow suit with the additional rating factor to avoid being selected against, the result may simply be an equilibrium in which all firms incur more up-front costs in pricing insurance products, but no one has gained a competitive edge over another.

5. Conclusion

This paper introduces a new test of asymmetric information in insurance markets based on potential identification of individual characteristics that are not used in pricing insurance but that are correlated with both insurance coverage and risk occurrence. This test offers a more refined approach for testing for asymmetric information than the standard positive correlation test used in much of the existing literature. It can also provide a direct test of adverse selection, as distinct from moral hazard.

We implement this new test using information on the geographic location of U.K. annuitants. We show that geographic identifiers predict a potential buyer's socio-economic status, survival probability, and amount of insurance purchased. Information on such identifiers is not used, however, in pricing annuities. External evidence of a positive socio-economic status – longevity gradient – independent of annuity contracts – further indicates that our findings provide direct evidence of adverse selection – as distinct from moral hazard – in U.K. annuity markets. It also suggests that part of the private information consists of information about socio-economic status. Other forms of private information not related to socio-economic status also appear to exist.

Our findings have implications beyond the operation of annuity markets. Evidence of socio-economic selection in annuity markets provide a potential unifying explanation for the observed disparities across insurance markets in the risk type of the insured relative to the population. There is no a priori reason to expect socio-economic selection to operate in the same direction in different insurance markets. Indeed,

[6] In 1998 – the end of the sample period of the data analyzed in this paper – however, sales of discounted or "impaired" life annuities represented only 3 percent of total sales in the compulsory market (Ainslie 2000).

the available empirical evidence suggests that it does not. In the annuity market, we have found that socio-economic selection draws higher risk individuals into the annuity market and therefore reinforces any selection based directly on private information about risk type. In the life insurance market, however, Banks and Tanner (1999) find that selection based on socio-economic status appears to draw lower risk individuals into the market. Finkelstein and McGarry (2003) find a similar pattern for the long-term care insurance market. These patterns of selection may explain why insured individuals do not appear to be higher risk in either the life insurance or long-term care insurance market, but do in annuity markets..

The ability of the test developed here to test directly for adverse selection distinctly from moral hazard is likely to be of greater interest in insurance markets other than annuities. Despite the theoretical potential for moral hazard in annuities (Philipson and Becker 1998), the empirical relevance is arguably weaker than in other insurance markets. However, in other insurance markets such as health or automobile insurance in which moral hazard is likely to be empirically more important and many public policies are predicated on the assumption of at least some adverse selection, the ability to apply the test developed here to test explicitly for the existence of adverse selection offers an exciting and important avenue for further research.

References

Abbring, Jaap, Pierre-Andre Chiappori, and Jean Pinquet. (forthcoming). "Moral hazard and dynamic insurance data." *Journal of the European Economic Association.*

Ainslie, Ross. 2000. "Annuity and Insurance Products for Impaired Lives." Presented to the Staple Inn Actuarial Society, May 9.

Akerlof, George. 1970. "The Market For Lemons: Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*

Attanasio, Orazio and Carl Emmerson. 2001. *Differential Mortality in the UK*. Institute for Fiscal Studies working paper W01/16. http://www.ifs.org.uk/wps/wp0116.pdf

Attanasio, Orazio and Hilary Hoynes. 2000. "Differential mortality and wealth accumulation," *Journal of Human Resources* 35(1): 1-29.

Association of British Insurers. 1999. Insurance Statistics Year Book. London: ABI.

Banks, James, and Carl Emmerson, 1999. "U.K. Annuitants," *IFS Briefing Note,* May 1999. London: Institute for Fiscal Studies.

Banks, James and Sarah Tanner. "Household Savings in the UK," October 1999. *IFS Reports*, R62.

Brown, Jeffrey and Amy Finkelstein. 2004. "Supply or Demand: Why is the Market for Long-Term Care Insurance So Small." NBER Working Paper 10782.

Brown, Jeffrey, Olivia Mitchell, James Poterba and Mark Warshawsky. 2001. The Role of Annuity Markets in Financing Retirement. MIT Press: Cambridge, MA.

Cardon, James and Igal Hendel. 2001. "Asymmetric information in Health Insurance: Evidence from the National Medical Expenditure Survey." *Rand Journal of Economics* 32: 408-27.

Cawley, John and Tomas Philipson. 1999. "An empirical examination of information barriers to trade in insurance." *American Economic Review* Vol 89, no 4 pp.827-846.

Chiappori, Pierre-Andre. 2000. Econometric Models of Insurance under Asymmetric Information. In Georges Dionne, ed. Handbook of Insurance Economics. London: Kluwer.

Chiappori, Pierre-Andre and Bernard Salanie. 2000. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy*, 108, 56-78.

Chiappori, Pierre-Andre, Bruno Jullien, Bernard Salanie, and Francois Salanie. 2005. "Asymmetric Information in Insurance: General Testable Implications." *Rand Journal of Economics.*

Cohen, Alma. 2001. "Asymmetric Information and Learning in the Automobile Insurance Market." Mimeo, Harvard University, Department of Economics.

Cohen, Alma and Liran Einav. 2005. "Estimating Preferences from Deductible Choice." Unpublished working paper.

Cox, David. 1972. "Regression models and life-tables (with discussion)." *Journal of the Royal Statistical Society*, Series B 34: 1987-220.

Cox, David. 1975. "Partial Likelihood." *Biometrika* 62: 269-276.

Cutler, David and Richard Zeckhauser. 2000. The Anatomy of Health Insurance, in The Handbook of Health Economics, Culyer AJ and Newhouse JP (eds)

de Meza, David and Webb, David. 2001. "Advantageous selection in insurance markets." *Rand Journal of Economics*. Summer.

Dionne, Georges, Christian Gourieroux and Charles Vanasse. 2001. "Testing for evidence of adverse selection in the automobile insurance market: a comment." *Journal of Political Economy* 109(2): 444-451.

Finkelstein, Amy and Kathleen McGarry. 2003. "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market." Unpublished mimeo.

Finkelstein, Amy and James Poterba. 2002. "Selection Effects in the Market for Individual Annuities: New Evidence from the United Kingdom." *Economic Journal* 112, 28-50.

Finkelstein, Amy and James Poterba. 2004. "Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market." *Journal of Political Economy.*

Jullien, Bruno, Bernard Salanie, and Francois Salanie. 2002. "Screening risk-averse agents under moral hazard." Unpublished mimeo.

Karlan, Dean and Jonathan Zinman. 2005. "Observing Unobservables: Identifying Information Asymmetries with a Consumer Credit Field Experiment." Memo.

McCarthy, David and Olivia Mitchell. 2003. "International Adverse Selection in Life Insurance and Annuities." NBER Working Paper 9975.

Mitchell, Olivia S., James M. Poterba, Mark Warshawsky, and Jeffrey R. Brown. 1999. "New Evidence on the Money's Worth of Individual Annuities." American Economic Review, 89, 1299-1318.

Murthi, Mamta, J. Michael Orszag, and Peter Orszag. 1999. The Charge Ratio on Individual Accounts: Lessons for the U.K. Experience. Birkbeck College, University of London Working Paper 99-2.

Office for National Statistics. 1997. *Health Inequalities.* London: The Stationary Office Limited.

Rothschild, Michael and Joseph Stiglitz. 1976. "An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics* 90: 629-649.

Smith, James. 1999. "Healthy Bodies and Thick Wallets: The Dual Relation between Health and Economic Status," Journal of Economic Perspectives, volume 13, number 2, Spring 1999, pp 145-166

Table 1: Summary Statistics on Annuitant Population at Sample Firm

Number of policies	52,826
Number (%) of annuitants who die within sample period	593 (10.6%)
Number (%) of annuitants who are male	31,330 (59.3%)
Average age at purchase	62.2
Number (%) of policies that are level nominal payout	47,372 (89.7%)
Number (%) of policies that have guarantees	43,261 (81.9%)
Mean initial payment	1,820
Median initial payment	901
Standard deviation of initial payment	3,683
Average premium	19,554

Note: The sample consists of single life compulsory annuities sold between 1988 and 1998 and still in force in 1998 (see text for further sample restrictions). Mortality experience covers the period January 1 1998 through February 29, 2004. Policies that are do not have level nominal payouts have payouts that increase over time in nominal terms. Policies with guarantees continue to make payments to annuitant estate if the annuitant dies during the guarantee period. Premium and initial payment are converted to 1998£ using the UK annual Retail Prices Index (RPI).

Table 2: Summary Statistics on Ward Characteristics: Socio-Economic Status and Health Status

	Population-weighted		Annuitant-weighted	
	Average	Std. Dev	Average	Std. Dev.
Qualified	13.4%	8.00	15.9%	8.15
Social Class: Professional and Managerial (I & II)	31.6	12.13	36.1	12.13
Social Class: Skilled (III)	43.6	6.95	41.7	7.48
Social Class: Partly Skilled or Unskilled (IV & V)	21.6	8.03	19.4	2.47
Long-term illness	12.1	3.44	11.4	3.12

Note: Based on ward-level statistics from 1991 UK census. Population-weighted estimates are constructed weighting each ward by its population; annuitant-weighted estimates are constructed weighting each ward by the number of policies the sample firm has in that ward. 3 percent (2.8 percent) of the population-weighted (annuitant-weighted) sample are in the omitted social class: armed forces, government schemes, and unknown (not shown).

Table 3: Hazard Models Relating Annuitant Mortality Experience to Ward Characteristics

	(1)	Education (2)	Occupation (3)	Illness (4)
Male	0.638***	0.629***	0.628***	0.636***
	(0.0349)	(0.0347)	(0.0348)	(0.0347)
Percentage of Ward that is Educationally Qualified		-0.0150***		
		(0.0017)		
Percentage of Ward in Professional or Managerial Occupations (Group I & II)			-0.0118***	
			(0.0017)	
Percentage of Ward in Skilled Occupations (Group III)			-0.0029	
			(0.0027)	
Percentage of Ward with Long Term Illness				0.0248***
				(0.0043)

Note: Coefficients are from Cox Proportional Hazard Model of time lived since 1998. N = 52,824. All regressions contain dummies for age at purchase and year of purchase. Heteroscedasticity-robust standard errors clustered at the ward level are in parentheses. In column 3, omitted category is percentage of ward in partly or unskilled occupations (Group IV or V). ***, **, * denotes statistical significance at the 1 percent, 5 percent and 10 percent level respectively.

TABLE 4: WARD CHARACTERISTICS AND IMPLIED 5-YEAR MORTALITY PROBABILITY

	Fraction of Ward Qualified		Fraction of ward in Social Class I or II		Fraction of ward with Long-Term Illness	
	Average	One Std Dev Above Average	Average	One Std Dev Above Average	Average	One Std Dev BelowAverage
Male	10.7	9.7	10.7	9.3	10.9	10.2
Female	4.3	3.7	4.3	3.7	4.2	3.8

Notes: Table reports the subsequent to 1998 5-year cumulative mortality probability of an individual who purchased an annuity at age 65 in 1994, conditional on having survived until 1998. Cumulative mortality probabilities are derived from the coefficient estimates in Table 3 and the associated estimated baseline hazard (not reported). For the change in the proportion of the ward in Social Class I or II, the individuals are moved to Social Class IV or V.

Table 5: Implied change in annuity payment for 65-yr old annuitant in ward one std dev from average

	Proportion Qualified	Proportion in Social Class I or II	Proportion with Long Term Illness
Male	3.0%	3.6%	1.9%
Female	2.4%	2.8%	1.5%

Note: Table reports the proportional increase (decrease) in level, non-guaranteed, annual annuity payments for a 65-year old if annuity pricing were to incorporate ward-level characteristics in an actuarially fair manner and the individual lives in a ward one standard deviation from the mean in terms of the ward characteristic. For the change in the proportion of the ward in Social Class I or II, the individuals are moved to Social Class IV or V.

TABLE 6: WARD CHARACTERISTICS AND INSURANCE QUANTITY

	Dependent Variable: Log Initial Payment			Dependent Variable: EPDV of annuity component of policy
	Policies with No Guarantee [N=7,964]	Policies with 5-yr Guarantee [N=35,042]	Policies with 10-yr Guarantee [N=4,366]	[N=47,370]
Panel A: Education measure of SES				
Percentage of Ward that is Educationally Qualified	0.0223*** (0.0017)	0.0271*** (0.0011)	0.0160*** (0.0022)	223.4*** (12.6)
Panel B: Occupational Measure of SES				
Percentage of Ward in Professional or Managerial Occupation (Group I & II)	0.0154*** (0.0018)	0.0201*** (0.0013)	0.0103*** (0.0022)	136*** (9.6)
Percentage of Ward in Skilled Occupations (Group III)	-0.0012 (0.0029)	-0.0010 (0.0020)	-0.0054 (0.0035)	-49.2*** (16.4)
Panel C: Health Measure				
Percentage of Ward with Long-Term Illness	-0.0373*** (0.0046)	-0.0438*** (0.0029)	-0.0284*** (0.0052)	-330.1*** (25.3)
Mean Dependent Variable (1998 £'s)	6.63	6.30	7.23	8,842

Note: Table reports coefficients from estimating equation (2) by OLS. Sample limited to 90 percent of policies that have level (nominal) payouts (i.e. not backloaded). Columns (1) through (3) report results separately for with different guarantee choices; the last column includes all three guarantee levels in the sample. All regressions include indicator variables for age and year of purchase and for gender of annuitant. Each panel x column reports the results from a different regression, with the SES measure described in the panel heading and the sample definition given in the column headings. In panel B, omitted category is partly or unskilled social class (groups IV or V). Standard errors are in parentheses. They are heteroscedasticity-robust standard errors and are clustered at the ward level to allow for within-ward correlation in the error term. ***, **, * denotes statistical significance at the 1 percent, 5 percent, and 10 percent levels respectively.

TABLE 7 – HAZARD MODEL LINKING TYPE OF POLICY, INDIVIDUAL ATTRIBUTES, AND WARD CHARACTERISTICS TO SUBSEQUENT MORTALITY

	(1)	(2)	(3)	(4)
Male	0.630***	0.621***	0.628***	0.620***
	(0.0355)	(0.0354)	(0.0354)	(0.0355)
Level Indicator	0.047	0.048	0.045	0.050
	(0.049)	(0.049)	(0.049)	(0.049)
Guarantee Indicator	0.083**	0.076*	0.079**	0.076*
	(0.0391)	(0.0400)	(0.0400)	(0.0400)
Initial Payment (£1,000)	-0.013***	-0.009	-0.012**	-0.009
	(0.0040)	(0.0058)	(0.0059)	(0.0058)
Percentage of Ward that is Educationally Qualified		-0.014***		
		(0.0018)		
Percentage of Ward with Long Term Illness			0.024***	
			(0.0043)	
Social Class (Omitted Category = % of Ward in Partly or Unskilled Occupations (Group IV or V))				
Percentage of Ward in Professional or Managerial Occupation (Group I & II)				-0.011***
				(0.0017)
Percentage of Ward in Skilled Occupations (Group III)				-0.003
				(0.0027)

Note: Coefficients are from Cox Proportional Hazard Model of time lived since 1998. N = 52,824. All regressions contain individual dummies for age at purchase and year of purchase (1988-1998) and frequency of annuity payments. Standard errors are in parentheses. They are heteroscedasticity-robust standard errors and are clustered at the ward level to allow for within-ward correlation in the error term. ***, **, * denotes statistical significance at the 1 percent, 5 percent and 10 percent level respectively. In column 4, the omitted social class is Group IV and V (partially skilled or unskilled occupation).

Figure 1: Cumulative Survival Probabilities By Social Class, Males (1992-1996)

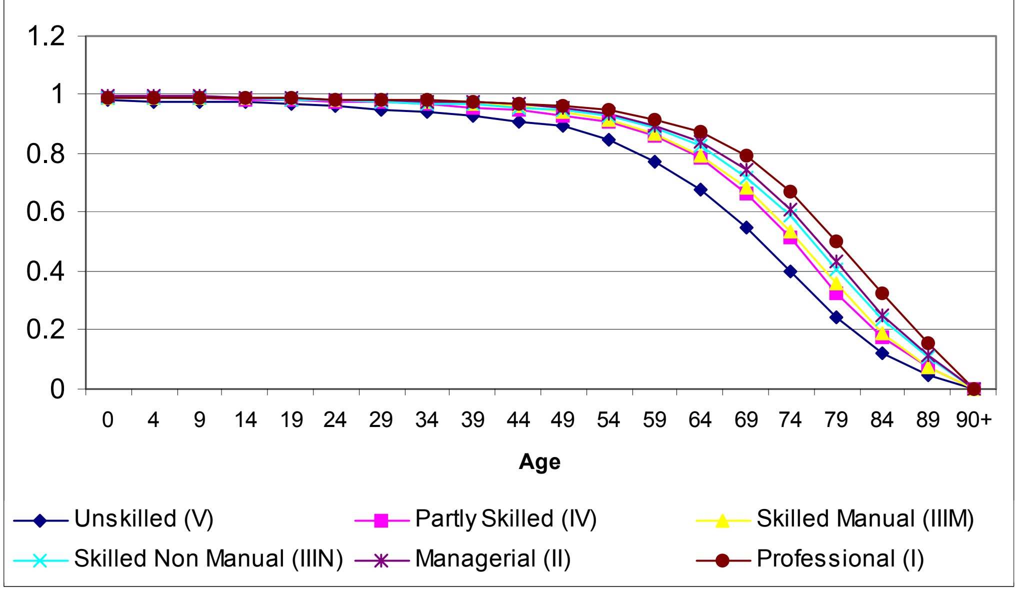

Source: ONS (1997)

Figure2: Survival Curve for Male Annuitants - Social Class

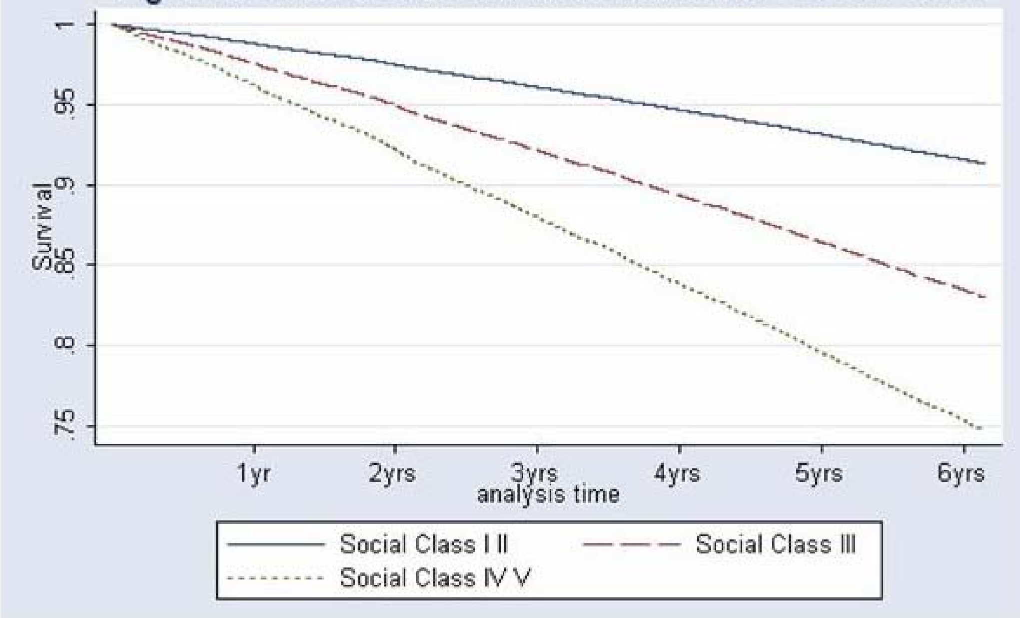

Notes: Figure plots the implied cumulative survival probability curve by social class for male annuitants aged 65 in 1994 from the hazard model estimation in Table 3. This hazard model includes as covariates only age and year of purchase and social class group, and the sample is limited to male annuitants only.