Adverse Selection and Government Intervention in Life and Health Insurance Markets

by

Amy Nadya Finkelstein

A.B., Government
Harvard University, 1995

M.Phil., Economics
Oxford University, 1997

Submitted to the Department of Economics
in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy in Economics
at the
Massachusetts Institute of Technology

June 2001

Signature of Author...
Department of Economics

Certified by...
James M. Poterba
Mitsui Professor of Economics
Thesis Supervisor

Certified by...
Jonathan Gruber
Professor of Economics
Thesis Supervisor

Certified by...
Jerry Hausman
John & Jennie S. MacDonald Professor of Economics
Thesis Supervisor

Accepted by...
Richard S. Eckaus
Ford International Professor of Economics, Emeritus
Chairman, Departmental Committee on Graduate Studies

Adverse Selection and Government Intervention
in Life and Health Insurance Markets

by

Amy Nadya Finkelstein

Submitted to the Department of Economics
in partial fulfillment of the requirements for the Degree of
Doctor of Philosophy

Abstract

This thesis examines the workings of insurance markets. The first two papers examine the effect of government tax and regulatory policy in markets for supplementary health insurance. The third paper presents new evidence of the importance of adverse selection in insurance markets.

The first paper examines the empirical consequences of imposing binding minimum standards on the market for private health insurance for the elderly in the United States. I find robust evidence of a substantial (40 percent) decline in insurance coverage associated with imposing these minimum standards. The standards are also associated with a reduction in coverage of non-mandated benefits among the insured. The minimum standards therefore, while requiring additional insurance coverage among the insured, were also associated with both extensive and intensive declines in insurance coverage. Considering all of these various changes, I estimate that the standards were, on net, welfare reducing.

The second paper presents new evidence of the effect of the tax subsidy to employer-provided health insurance on coverage by such insurance. I study the effects of a 1993 tax change that reduced the tax subsidy to employer-provided supplementary health insurance in Quebec by over half. Using a differences-in-differences methodology in which changes in Quebec are compared with changes in other Canadian provinces not affected by the reform, I estimate an elasticity of employer coverage with respect to the tax price of –0.46 to –0.49. The tax subsidy appears much more critical to the provision of supplementary health insurance in small firms than in larger ones.

The third paper, written jointly with James Poterba, re-examines the importance of adverse selection in insurance markets. We use a unique data set of all annuity policies sold by a large U.K. insurance company since the early 1980s to analyze mortality differences among individuals who purchased different types of policies. We find systematic relationships between ex-post mortality and annuity policy characteristics that are consistent with models of asymmetric information in insurance markets. We confirm that the pricing of features of annuity contracts is consistent with the self-selection patterns we find in mortality rates.

Thesis Supervisor: James M. Poterba
Title: Mitsui Professor of Economics

Table of Contents

Acknowledgements

I have been privileged to receive the advice and encouragement of numerous people throughout my time in graduate school.

My primary debt of gratitude is to my advisers, Jim Poterba, Jon Gruber and Jerry Hausman. I benefited enormously throughout each stage of graduate school from their advice and enthusiasm. They continually encouraged me to try more ambitious projects and to persevere through challenging times. I am extremely grateful for the vast amount of time that they have taken to teach me how to do good economic research and to show me by example how much fun this endeavor can be.

I have also benefited greatly from the accessibility and help of many other faculty members. I particularly wish to thank Daron Acemoglu, David Autor, David Cutler, Peter Diamond, Sendhil Mullainathan and Scott Stern.

The staff at both MIT and NBER have, in numerous ways, made my graduate school experience more manageable and enjoyable. I particularly wish to thank Gary King and John Arditi for all of their good natured help. I am also grateful to both of these institutions for financial support for my research.

My fellow graduate students have been a constant source of help and support. This especially includes: Pierre Azoulay, Emek Basker, Dan Bergstresser, Hoyt Bleakley, Jeff Brown, Courtney Coile, Sue Dynarski, Jeff Furman, Lakshmi Iyer, Ben Jones, Robin McKnight, Paras Mehta, Ben Olken, and Melissa Schettini. Above all, I wish to thank Sarah Reber, who has continually set new standards for both a good economist and a good friend, and without whom I would often be, both figuratively and literally, in the dark.

Last but certainly not least, I want to thank my parents, to whom I dedicate this thesis. They have been an unflagging source of guidance, support, inspiration, love, advice of all kinds, and good humor throughout my life. I look forward to many more years of their skepticism of research endeavors in the social sciences.

Introduction

Economists have clearly defined a set of market failures under which the fundamental welfare theorems no longer hold, and government intervention in the private market can be justified on efficiency grounds. An economic rationale for government intervenion by itself, however, provides very little guidance as to the expected consequences of particular government policies. The consequences of these policies depend both on the equilibrium and structure of the affected market and on the magnitude of possible behavioral responses.

This thesis examines the empirical consequences of tax policy and other forms of government regulation in private health insurance markets. Government intervention in private insurance markets is ubiquitous, and health insurance markets are no exception. Indeed, in the United States, the tax exclusion of employer-provided health insurance from the individual income tax is the largest single Federal tax expendtiure. Increasingly, health insurance markets are also directly regulated, with requirements on what they must cover or whom they must serve.

I study the consequences of both of these types of regulation – tax subsidies and minimum standards – on supplementary private health insurance markets that complement a non-comprehensive public health insurance program. The first market studied is the Medigap market in the United States. Medigap insurance policies cover some portion of the health costs not covered by Medicare, the universal public health insurance for the elderly in the United States. The second market studied is the market for private health insurance in Canada. This private health insurance serves to supplement the national public health insurance system.

Both the U.S. and Canadian public health insurance programs leave substantial medical risks uncovered. The Candian program finances about three-quarters of Canadian health expenditures while the U.S. Medicare system pays just under half of all health care expenses for the elderly in the United States. As a result, private plans to supplement each public program are widespread. Over three-quarters of the relavent population has private insurance to supplement the Canadian public insurance system and about two-thirds have private insurance to supplement the U.S. Medicare program.

In each case, the government is not only a major provider of insurance through the public health insurance program, but also intervenes in the private market for supplementary health insurance. The consequences of this private market regulation is the subject of this thesis. The first chaper examines the effects of state-imposed minimum standards regulations on the non-group Medigap market. The second examines the tax subsidy to employer-provided supplementary health insurance on the supplementary private insurance market in Canada.

The emphasis of the empirical work is two-fold. First, I investigate the empirical consequences of the tax policy or minimum standard on insurance coverage rates. The policies I study diretly affect only one possible source of insurance coverage (either the workplace-based market or the indiviudal market). I therefore examine the effects of the policies both on insurance coverage in the directly affected market and in the markets that are not directly affected by the regulation but provide an alternative source of insurance coverage.

Second, I investigate whether there appears to be any efficiency rationale, or net welfare gain, from the government intervention. In the case of the tax subsidy, I examine whether the tax subsidy plays an important role in creating workplace-based insurance pools and hence minimizing the adverse selection pressures present in the individual market. In the case of the minimum standards, I estimate the net welfare change associated with imposing minimum standards. These minimum standards may increase the amount of insurance for those who remain in the market, but they are also associated with a large exit from the regulated market and hence declines in insurance coverage. The net welfare impact of these changes in risk bearing is therefore ex-ante ambiguous.

The possibility of an efficiency-enhancing role for government policy in insurance markets motivates the empirical analysis of the impact of government tax policy or minimum standards regulation. But such a role must be predicated on a pre-existing market failure that results in non-optimal private insurance purchases. The possibility of asymmetric information and adverse selection in insurance markets provides a potential economic rationale both for subsidizing workplace-based insurance pools through the tax system and for minimum standards regulations The final chapter of this thesis directly examines evidence of the role of asymmetric information in insurance markets.

Outline of Thesis

In the first chapter of the thesis, I examine the effects of imposing binding minimum standards on the voluntary market for private health insurance for the elderly in the U.S. This insurance – commonly known as Medigap – is designed to supplement the public, universal health insurance provided through Medicare. The potential for consumer misnformation or adverse selection in the private supplementary health insurance provide potential economic rationales for minimum standards. Yet the potential private market failure must be balanced against the potential for unintended, negative consequences of the minimum standards. The possibility of negative consequences of minimum standards make them a contentious policy intervention both in health insurance markets and in other settings. Two primary concerns are that minimum standards will cause consumers to exit the market for the regulated insurance product and that they may affect the equilibrium provision of non-regulated aspects of the policy among those who retain coverage.

The empirical evidence presented in the paper sheds light both on the effects of the minimum standards on both these margins and the market forces and market failures operating in this insurance market. I find robust evidence of a large "quality-quantity" tradeoff. The imposition of minimum standards is associated with a short-term decline in coverage in the regulated market of 6 percentage points, and a long term decline of 13 percentage points (40 percent). There is no evidence of substitution to other, unregulated sources of supplementary health insurance for the elderly. I also find evidence that is consistent with the minimum standards reducing the coverage of non-mandated benefits among the insured. I explore what market mechanisms are likely behind this finding and conclude that the evidence is consistent with a model of asymmetric information in which minimum standards hamper the ability of insurance companies to screen individuals into different risk classes. Finally, I estimate the net welfare effects of the minimum standards. The empirical results point to substantial declines in insurance coverage – on both the extensive and the intensive margin – associated with the imposition of the minimum standards. Yet the minimum standards themselves required considerable additional insurance for those who maintained coverage. My estimates of these various induced changes in insurance coverage suggest that, even under conservative assumptions, the imposition of the minimum standards was, on net, welfare reducing.

In the second chapter of the thesis, I examine the effects of the tax subsidy to employer provided health insurance. I again examine the effects of the policy on a supplemental private health insurance market. In this case, I look at the market for private health insurance in Canada that supplements the national public health insurance program. Employer-provided health insurance is subsidized through the tax system in both the United States and Canada. This tax subsidy is potentially efficiency-enhancing if it is necessary for the creation of workplace-based health insurance pools which help to overcome the adverse selection problems that plague the non-group health insurance market.

I study the effects of a 1993 tax change that reduced the tax subsidy to employer-provided health insurance in Quebec by almost 60 percent. I estimate the effects of this large-scale change in the tax subsidy to employer-provided health insurance by comparing changes in Quebec with changes in other provinces not affected by the tax reform. The results indicate that the tax subsidy has a large effect both on employer-provided supplementary health insurance coverage and on total insurance coverage. I estimate that the reduction in the tax subsidy in Quebec coincided with a 13 to 14 percentage point drop in workplace coverage. This corresponds to an elasticity of employer coverage with respect to the tax price of –0.46 to –0.49. Although there is some evidence of a slight offsetting rise in non-group coverage in response to the reduction of the tax subsidy in the group market, the net decline in private health insurance coverage associated with the reduction in the tax subsidy is still substantial.

The evidence also indicates substantial dispersion across firms of different sizes in the sensitivity of workplace coverage to the tax subsidy. In particular, the results suggest that the tax subsidy is critical to the maintenance of the workplace-based health insurance pool in small firms. It appears to play substantially less of a role in larger firms, where gains from pooling and reduced administrative costs are likely to play a larger role in the maintenance of workplace-based insurance pools.

A potential problem in both of the insurance markets studied – and one that might motivate government intervention of the types studied – is the existence of adverse selection in these markets. The final chapter of the thesis, which is joint work with James Poterba, presents new evidence of the importance of adverse selection in insurance markets. We test two simple predictions of asymmetric information models using data from the annuity market in the United Kingdom. The first is that higher risk individuals self-select into insurance contracts that offer features that, at a given price, are more valuable to them than to lower risk individuals. The second is that the equilibrium pricing of insurance policies reflects the fact that the risk pool varies across different policies. Such self-selection across policies by risk type would not occur if the insurer and the insured had symmetric information. Our empirical work finds support for both of these predictions.

A particularly noteworthy finding concerns the dimensions of the insurance product along which we observe self-selection. We find systematic relationships between ex-post mortality and annuity policy characteristics, such as whether the annuity will make payments to the estate in the event of an untimely death and whether the payments from the annuity rise over time. However, we do not find evidence of self-selection on the initial amount of annuity payment. This initial amount of annuity payment is analogous to measures of the accident-induced payment that are used in other empirical papers on asymmetric information in insurance markets. Our research therefore highlights the importance of paying careful attention to the detailed policy features of real-world insurance contracts when testing theoretical models of asymmetric information in insurance markets.

Chapter 1:
Minimum Standards and Health Insurance Regulation: The Case of Medigap

"We may tell the society to jump out of the market frying pan, but we have no basis for predicting whether it will land in the fire or a luxurious bed"

– George Stigler, The Citizen and the State: Essays on Regulation, 1975.

1. Introduction

Government-imposed minimum standards are a near-ubiquitous form of private market regulation. Examples range from safety requirements in housing construction to minimum allowable staff-child ratios in private child care centers to the regulation of employer-provided pension plans. Private health insurance markets are no exception. Minimum benefit standards have been applied or proposed in several different contexts, from state requirements that mental health benefits be included in employer-provided health insurance packages to Federal proposals for a "Patients' Bill of Rights" that would impose minimum standards on HMOs.

In a perfect information, perfect competition model of health insurance markets, there is no economic rationale for minimum standards. Individuals face actuarially fair prices for health insurance and make health insurance purchase decisions that are consistent with expected utility maximization. Relaxing these assumptions, however, can create a role for government intervention in the insurance market. For example, if consumers underestimate the probability of various medical risks, their voluntary insurance purchases will be sub-optimal. Alternatively, asymmetric information can destroy the market for the insurance product or, short of that, result in sub-optimal amounts of insurance for certain risk classes. Minimum standards, by setting a floor on the amount of insurance purchased by the insured, may help to counteract a tendency to purchase sub-optimal amounts of insurance.

As the Stigler quotation above suggests, however, the desire to solve potential private market failures must be balanced against the potential for unintended, negative consequences of minimum standards. A primary concern is that minimum standards will cause people to exit the market for the regulated insurance product. These people may either substitute toward unregulated forms of insurance, or drop coverage all together. An additional concern is that minimum benefit standards may affect the equilibrium provision of non-mandated insurance benefits in the regulated insurance market. Knowledge of the magnitude of the response to minimum standards along both these margins is critical to evaluating the merits of this form of regulation.

Despite this, there is virtually no empirical evidence on the effect of minimum standards on health insurance markets.[1] In this paper, I therefore examine the empirical consequences of large and binding minimum standards on the voluntary, private supplementary health insurance market for the elderly. Such insurance is commonly known as Medigap or Medicare supplement insurance. These insurance policies cover some portion of the health costs not covered by Medicare, the universal public health insurance for the elderly in the United States. There are two main types of coverage gaps in Medicare. First, uncapped cost sharing provisions in services covered by Medicare leave the elderly exposed to considerable out of pocket risk. Second, services that Medicare covers either with stringent restrictions, such as care in a skilled nursing facility, or not at all, such as most outpatient prescription drugs, also leave considerable out of pocket risk exposed.

The market failures that provide potential rationales for minimum standards may well be present in the Medigap market. Consumer misinformation, not only about medical risks but also about insurance needs, may well be a problem for many Medigap consumers. They must decide whether and how to supplement a public health insurance plan that they may not fully understand. Indeed, concerns that the elderly were making uninformed decisions about their Medigap purchases was one of the motivations for the minimum standards (Merritt and Potemken 1982). Problems of asymmetric information and adverse selection – another potential rationale for minimum standards – are also present in the Medigap market (Ettner 1997).

[1] An exception is Gruber (1994) who studies the effect of state mandated benefits for employer-provided health insurance. He finds no evidence of an effect of these mandates on insurance coverage. He notes, however, that this may simply be due to the fact that these mandates were not binding, since most firms who offered insurance were already covering the mandated benefits.

In the late 1970s and early 1980s, almost all states followed a federal "recommendation" to impose minimum standards on the non-group Medigap market. The regulations did not require that individuals purchase Medigap insurance. Nor did they regulate the price of Medigap insurance. What they did was to specify certain gaps in Medicare coverage that any non-group Medigap policy must cover. The coverage of other gaps was left to the market.

The rest of the paper is organized as follows. In Section two, I provide background on the Medigap market and the specific nature of the minimum standards imposed. Section three discusses theoretically the effects of minimum standards in a voluntary, private insurance market. I begin with the benchmark case of perfect competition, constant returns to scale, and perfect information. In this setting, binding minimum standards can produce only declines in the proportion of the population with private insurance coverage. They can also result in changes (of either sign) in the amount of non-mandated insurance coverage among those who retain private insurance coverage. I then consider departures from the benchmark model, such as imperfect competition and asymmetric information, that allow the minimum standards to also affect the supply of insurance. By allowing for the possibility that the minimum standards can affect the price of insurance benefits, these departures provide additional avenues by which the minimum standards can affect both the proportion of the population with private insurance coverage and the amount of coverage of non-mandated benefits among those who retain insurance. The net direction of these effects, however, is theoretically ambiguous. The rest of the paper examines these effects empirically.

Section four examines the consequences of the regulation for insurance coverage rates. I use approximately biannual coverage data from the National Health Interview Survey (NHIS) from 1976 through 1986. Using the variation in the timing of different states' adoption of the minimum standards to identify their effects, I find robust evidence of a large "quality-quantity" tradeoff. The imposition of minimum standards is associated with a short-term decline in coverage in the regulated market – non-group Medigap policies – of 6 percentage points (15 percent), and a long term decline of 13 percentage points (40 percent). I find no evidence of substitution to the two other potential sources of supplementary insurance coverage: Medicaid, which provides free public health insurance that is more comprehensive than Medicare to the indigent elderly, or the group Medigap market. I discuss possible explanations for the large net decline in insurance coverage associated with the minimum standards.

In Section five, I examine the effect of the minimum standards on coverage rates for non-mandated benefits among those with non-group coverage. For this analysis, I use detailed information on benefits provided by individuals' Medigap policies from the 1977 National Medical Care Expenditure Survey (NMCES) and the 1987 National Medical Expenditure Survey (NMES). I compare changes in the provision of non-mandated benefits among affected individuals in the non-group market to changes for individuals in two different control groups. Where they overlap, these two control strategies yield similar findings. Together, this evidence suggests that the mandated minimum benefit standards were associated with a reduction in non-mandated benefits among those with non-group coverage. I explore which of the various mechanisms discussed in Section three appears likely to explain this finding.

These two sets of results point to substantial declines in insurance coverage – on both the extensive and the intensive margin – associated with the imposition of binding minimum standards. Yet the mandated minimum benefits themselves provided considerable additional insurance for most individuals with non-group coverage. In Section six, I consider the overall impact of the regulation on the health expenditure risk faced by the elderly and calculate the net welfare implications of these changes in risk bearing. I find that, even under conservative assumptions, the minimum standards appear to be, on net, welfare reducing. Section seven concludes.

2. Minimum standards in the Medigap market.

2.1. *The Medigap market*

Universal among the elderly in the United States, Medicare provides only partial health insurance coverage. Indeed, in 1977, the period right before the minimum standards regulation, Medicare contributed just under half of all payments for health care for the elderly. As a result, private insurance coverage that supplemented Medicare was widespread. In 1977, almost two-thirds of Medicare beneficiaries had private health insurance coverage to supplement their Medicare benefits. (Cafferata

1984). This Medigap insurance was obtained, in roughly equal proportions, from group and non-group sources.[2] The non-group market, which is the market to which the minimum standards applied, was a highly concentrated market. In 1984, Blue Cross and Blue Shield plans accounted for three-quarters of non-group Medigap premiums nationwide. Three companies accounted for over 50 percent of the remaining non-group premiums earned by commercial companies. (U.S. General Accounting Office, 1986).

Expenditures for private supplemental health insurance were substantial. To give a sense of the magnitude, the total annual amount spent by the elderly on Medigap premiums was approximately one-tenth of the total direct medical expenditures (not including insurance premiums) on the elderly in 1977. (Cafferata 1984). The average annual premium for a non-group policy was $553 in 1999 dollars.[3]

Medigap insurance is designed to cover some of the "gaps" in Medicare. Medicare consists of two different programs. Medicare Part A (Hospital Insurance Program) is mandatory and covers some non-physician inpatient hospital care expenses, and some care in skilled nursing facilities or home health care. Medicare Part B (Supplementary Medical Insurance Program) helps pay for physician fees for covered services and other medical and health services.

Gaps in Medicare coverage fit into three main categories. First, there are cost-sharing provisions for the health services that Medicare covers. These include annual deductibles and co-payments for Part B covered-expenses and for Part A-covered hospital stays. The co-payments for both Part B and Part A are uncapped. In addition, the Part A co-payments increase as a percent of total expenses as the length of the hospital stay increases. As a result, these cost sharing provisions create substantial exposure to medical expenditure risk. Second, there are certain health services that Medicare covers only partially and/or with severe restrictions, such as care in a skilled nursing facility or home health care. Third, there are health

[2] Author's calculation based on NHIS data described in section 4. Throughout this paper, I refer to health insurance purchased through a current or former employer or union as "group" insurance. I use the term "non-group" insurance to refer to non-employment related health insurance. This is purchased either directly through a private company or through a non-employment related association such as the AARP.

[3] Throughout this paper, numbers reported in 1999 dollars are adjusted using the CPI-U.

services that Medicare does not cover at all. These include such benefits as outpatient prescription drugs and hospital stays beyond the maximum number of days covered by Medicare.

2.2. Minimum Standards: The Baucus Amendments

Prior to state adoption of minimum benefit standards in the late 1970s, the non-group Medigap market was essentially unregulated (Van Ellet (1979), McCall, Rice and Hall (1983)). This makes the Medigap market a uniquely useful setting for studying the consequences of minimum standards in insurance markets. Most insurance markets are heavily regulated and have been so for a long time. In such markets, isolating the consequences of one particular type of regulation can be difficult.

Starting in the late 1970s, a small number of states began introducing minimum standards in the non-group Medigap market. In 1978, the National Association of Insurance Commissioners (NAIC) issued a set of model standards for regulation of the non-group Medigap market. These model regulations were incorporated in the 1980 Federal Baucus amendments which provided "encouragement" to the states to adopt them.[4] Presumably as a consequence, 42 states and the District of Columbia shortly thereafter enacted minimum standards for non-group Medigap policies (the remaining states already had regulation in effect that satisfied the Baucus requirements). Table 1 gives the first full year that the regulations were in effect in each state.

Most states simply adopted the minimum standards set forth first in the NAIC model regulations and subsequently in the Federal Baucus amendments. However, I identify 8 states in Table 1 that adopted regulations that placed greater constraints on the allowable policy space than did the Baucus amendments. In the analysis below, I am unable to find any evidence of a differential effect of the more stringent regulations on coverage rates.

The minimum standards applied only to non-group policies. They did not apply to policies purchased from a current or former employer or union ("group policies"). Nor did they apply to policies that were

[4] Since insurance regulation is the prerogative of the states, the federal regulations were technically voluntary. Merritt and Potemken (1982) and McCall et al. (1983) describe the institutional structure established by the Baucus amendments to encourage states to adopt these regulations.

converted from group policies to non-group policies. Finally, the regulations applied only to new, non-group policies, not to renewals of existing policies.[5]

The minimum standards placed a limit on the exclusion of pre-existing conditions for more than 6 months and specified a minimum set of benefits that any policy sold after the enactment of the regulations must cover. These consisted of coverage of the Part A copays for hospital days 61 to 150, and coverage of 90 percent of the cost of hospital stays above 150 days for at least an additional 365 days.[6] Coverage of the Part B copay for 20% of approved physician charges, subject to a maximum deductible of $200 and a maximum benefit of no less than $5,000, was also mandated. Finally, the policy had to cover the annual deductibles in both Part A and Part B for the first three pints of blood used, but not the general Part A and Part B deductibles. Appendix A provides more detail on the gaps in Medicare coverage and on the specific requirements of the minimum standards.

The Baucus amendments arose from two main concerns. First, there was a concern that the elderly were unable to make informed choices about their insurance coverage (see e.g. Merritt and Potemken 1982). The minimum benefit standards described above were designed to address these concerns. In addition, there was also concern about fraud and abuse practiced by a very small segment of the industry that exploited the needs, fears and lack of information of the elderly (McCall et al 1983, U.S. House of Representatives 1978). These concerns motivated several other provisions of the Baucus amendments besides the minimum benefit standards which are the focus of the paper. These included, for example, allowing the purchaser a 30 day "free look" period during which time the policy may be returned for a full refund, standardization of terminology used in the policy, and a requirement that cancellation and termination clauses be prominently displayed. In contrast to the minimum benefit standards, these policies did not affect the vast majority of buyers. (McCall et al., 1983, Merritt and Potemken 1982, U.S. House of Representatives 1978). In the analysis of the empirical results below, I therefore attribute estimated

[5] All non-group Medigap policies are sold on an annual basis.

[6] After the annual deductible, Medicare Part A fully covers all hospital inpatient expenses for 60 days, after which there is a co pay for hospital days 61-90, and another, higher co pay for hospital days 91-150. Thereafter, Medicare Part A coverage ceases.

effects of the reforms to the minimum benefit standards which were the aspects of the reform that affected the vast majority of buyers.

2.3 The "bite" of the minimum standards

The 1977 NMCES data indicate that, prior to the enactment of the regulations, less than 7 percent of individuals with non-group policies had policies that would have met the minimum standards that are measurable in these data. 10 years later, the 1987 NMES data indicate that the Baucus requirements were strongly enforced. Of individuals with non-group policies in 1987 who were young enough to have bought their policy after the enactment of the minimum standards, 94 percent had policies that met the Baucus requirements measurable in the 1977 data.[7] The most binding of the measurable mandated benefits was the requirement of coverage for 365 hospital days beyond the first 150 days, followed by the requirement for coverage of the Part B copayment. Only 11 percent of individuals with non-group policies in 1977 had full coverage for up to 365 hospital days past 150 days, and only 52 percent had full coverage of the Part B copayment. In contrast, 70 percent had coverage for the Part A copayments for hospital days 91-150, and 87 percent had coverage for Part A copayments for hospital days 61 to 90.

The potential out of pocket liability covered by the mandated benefits was substantial. On average, the amount of additional insurance required to upgrade a pre-reform plan to comply with the minimum standards was about one-fifth of the total amount of insurance in the pre-reform plan.[8] The mandated Part A benefits provided insurance against rare but potentially catastrophic financial risks associated with long hospital stays.[9] Part A copays were $85 per day in 1999 dollars for hospital days 61 to 90 and $171 per day in 1999 dollars for hospital days 90-150. An individual without private insurance to cover these Part A copays would therefore have to pay $12,788 in 1999 dollars for a 150-day hospital stay. For someone who stayed another 365 days without private coverage (and we have seen that only 11 percent had such

[7] These estimates do not include any information on compliance with the regulations regarding pre-existing conditions or coverage of the Part A and Part B blood deductibles, but include information on compliance with all other mandated benefits. The data are described in more detail in Section 5.

[8] I compare risk premiums across insurance arrangements to quantify the difference in the amount of insurance. The data and method behind this type of calculation are described in more detail in section 6. The estimate reported here assumes a constant relative risk aversion utility function and a coefficient of relative risk aversion of 3.

coverage), this would cost $170,638 in 1999 dollars.[10] Moreover, the 20% copayment for physician charges (for which the Baucus amendments also mandated coverage) left an uncapped and potentially large exposed risk as well.

3. Minimum standards in a voluntary private insurance market

This section considers theoretically the effects of these minimum standards. Consider a 65-year old, newly eligible for Medicare, who is deciding whether and how much supplementary health insurance to purchase in the private market. Individuals choose more or less supplementary health insurance as shown in Figure 1. Their budget constraint is given by the line AB. Individuals who place a high value on insurance will purchase more comprehensive policies such as E, while those who value insurance less will choose a point such as D.

Now suppose that the government imposes a minimum standard on the private, supplementary insurance market. The government does not require that individuals purchase supplementary insurance. However, it mandates that anyone who does purchase insurance must purchase at least an amount m. The individuals' budget set now becomes the point A and the solid line CB. The dashed line AC (a subset of the original budget constraint AB) is no longer part of the individual's choice set.

Consider first the case of perfect competition, constant returns to scale, and perfect information. Under these assumptions, the relative price of health insurance is unaffected by the regulation. As a result, individuals whose insurance purchases already satisfy the minimum standards (i.e. they either purchase no insurance at point A or they purchase more than m, such as at point E) will experience no change in their consumption decisions as a result of the reform. An individual who purchases less than m (such as the individual at point D), however, must now compare his utility from getting no insurance (point A) to his maximal utility from purchasing a policy that complies with the minimum standards. Assuming strictly convex preferences, this maximal utility will be achieved by purchasing exactly the minimum required amount of insurance (point C). As drawn, the individual now prefers purchasing no insurance

[9] Data from a 20% sample of Medicare beneficiaries in 1984 indicate that only 0.5% of the elderly had hospital stays beyond 60 days and only 0.1% had hospital stays beyond 150 days.

(point A) to complying the minimum requirements (point C). Figure 1 thus illustrates how minimum standards can produce declines in the proportion of individuals with private insurance coverage and reduce welfare.

One important way in which Figure 1 abstracts from the empirical example studied in this paper is that it assumes that supplementary health insurance varies on only one dimension: quantity. In practice, we have seen that Medigap policies are multi-dimensional; they may provide coverage for any number of different benefits. Individuals whose non-group Medigap policies would have met the minimum standards prior to their enactment were more likely to have coverage for all of the six non-mandated benefits measurable in the data than individuals whose policies would not have met the minimum standards (38 percent versus 7 percent). Still, coverage for some of the non-mandated benefits was quite common among individuals who did not have coverage for all of the subsequently-mandated minimum benefits. For example, 98 percent of individuals whose policies did not cover all of the subsequently-mandated benefits had coverage for the (non-mandated) Part A deductible, and over one-fifth had coverage for the out-patient prescription drugs (also not mandated).[11]

We can therefore enrich the analysis to think of a three-good world: mandated minimum insurance benefits, non-mandated insurance benefits, and all other goods. Under the same assumptions as above (perfect competition, constant returns to scale, and perfect information), individuals whose insurance policies already satisfy the minimum standards will not change their purchase of the non-mandated benefits with the imposition of minimum standards. Consider, however, the case of an individual whose policy did not already satisfy the minimum standards and who chooses to upgrade his policy to comply with these standards. A sufficient condition for this individual's purchase of non-mandated benefits to be unaffected by the minimum standards is that preferences are quasi-linear in all other goods and additively separable in the two components of insurance. For more general preferences, the model generates ambiguous changes in the amount of non-mandated benefits purchased by

[10] This assumes a hospital charge of $467.5 per day in 1999 dollars (American Hospital Association, 1978).
[11] Author's calculations based on 1977 NMCES. These data are described in more detail in Section 5.

individuals who upgrade their policies to meet the minimum standards. The sign of these changes depends on both the income effects of the minimum standards and on whether or not the mandated and non-mandated benefits are complements or substitutes.

Additional effects of the minimum standards on the decision of whether to purchase insurance and on the chosen level of non-mandated benefits may be found if we move away from the assumptions of perfect competition, constant returns to scale and perfect information. In particular, a variety of supply side mechanisms can affect the prices of mandated and non-mandated benefits. When the minimum standards produce a change in insurance prices, they may now produce changes in the insurance purchase decisions for individuals who already met the requirements (i.e. those who were not purchasing any insurance or those whose insurance purchases satisfied the minimum standards) as well as for individuals who did not.

Several mechanisms could produce additional declines in the proportion of individuals with insurance and in coverage of non-mandated benefits among the insured. For example, suppose that there are joint costs incurred by insurance companies in producing different insurance products. In this case, minimum standards, by requiring a company to drop its non-compliant policies, can increase the share of costs borne by the remaining policies and hence their prices.

In the insurance market context, asymmetric information provides an additional mechanism by which minimum standards may produce further insurance declines. Consider the standard Rothschild and Stiglitz (1976) framework in which there are two types of individuals, high risk and low risk, whose risk type is private information. In equilibrium, high risk types are able to purchase optimal (full) insurance at their actuarially fair (high) price, but low risk types are constrained to purchase less than the optimal (full) insurance at their actuarially fair (low) price. If the low risk type were not constrained, the high risk type would instead deviate to purchase the insurance product priced for the low risk type. In such a setting, sufficiently large minimum quantity standards can destroy the separating equilibrium by preventing the low risk type from purchasing a low enough quantity of insurance to maintain incentive compatibility for the high risk type. If the separating equilibrium is destroyed, the amount of insurance purchased by the high risk type will decrease, and both types will pool at an amount of insurance above the minimum. For sufficiently large minimum standards, the result can be a destruction of the market for any policy with more than the minimum amount of insurance (Neudeck and Podczeck, 1996).[12]

On the other hand, relaxing the assumption of perfect competition could result in minimum standards exerting downward pressure on prices. Ronnen (1991), for example, develops an oligopoly model with differentiated products in which the introduction of minimum standards results in an increase in both the proportion of the population purchasing the regulated product and the provision of non-mandated benefits. The intuition is that, by reducing the space over which firms can try to differentiate their products, minimum standards enhance price competition. Under the model's assumptions of price competition in the presence of fixed costs, firms that were already offering products that meet or surpass the minimum standards increase the quality of their product further to try to alleviate some of the added price competition.

The predicted effects of the minimum standards on both the proportion of individuals with insurance coverage and on the insured's purchase of non-mandated benefits are thus theoretically ambiguous. The next two sections examine each of these effects empirically.

4. Effect of Minimum Standards on Insurance Coverage

4.1. Data and Empirical Strategy

I use repeated cross-sections from the National Health Interview Surveys (NHIS) to examine the effect of the minimum standards on the probability of being covered by non-group Medigap. The NHIS is an annual U.S. household survey. Supplementary questions on individuals' source of private health insurance (i.e. group or non-group), if any, were asked in the 1976, 1978, 1980, 1982, 1983, 1984, and 1986 NHIS. The earliest data, therefore, pre-dates the introduction of the minimum standards in all but one of the states. The major drawback to the NHIS is that these data contain no information on health insurance premiums or on the benefits covered by the individual's insurance. In the next section,

[12] Encinosa (forthcoming) shows that this result is not robust to relaxing the equilibrium requirement in Rothschild and Stiglitz (1976) that no contracts offered in equilibrium make negative expected profits.

therefore, I will turn to an alternative data source to examine the effect of the minimum standards on the provision of non-mandated benefits.[13]

For the main analysis, I restrict the sample to those aged 65-68 who are covered by Medicare. This yields a sample of 17,649 individuals.[14] I look only at these "young" old because these are the people who are most likely to be buying Medigap policies after the regulations went into effect. Older individuals are likely to have purchased their policies before the regulations went into effect and are therefore less likely to be affected by them. In additional analyses below, I examine the sensitivity of the estimates to the particular set of ages included in the "young" old, and also investigate whether the reform affected the insurance purchase decisions of older individuals.

The dependent variable is the binary variable COVERAGE. It indicates whether the individual has non-group private health insurance, defined as insurance that was not "obtained through an employer or union." Individuals who do not have private health insurance or whose insurance was "obtained through an employer or union" are coded as not having non-group private health insurance.[15] The main analysis pertains to non-group coverage rates since the regulations applied to this market. Further analysis described below indicates that coverage rates in the group market were not affected by the reform. The estimated decline in non-group coverage therefore represent a net decline in private insurance coverage.

I exploit the variation in the timing of different states' adoption of the minimum standards to identify the effect of these standards on coverage rates in the non-group market. The empirical strategy is to compare non-group coverage rates after the reform has been imposed to non-group coverage rates prior to its imposition, while controlling for other possible confounding changes. The basic estimating equation (referred to as the "basic specification" in subsequent discussions) is:

---

[13] Unfortunately, the absence of high frequency data with both premiums and detailed benefit information prevents any analysis of the effect of the minimum standards on premiums.

[14] I would like to further restrict the sample to those not on Medicaid and those without military health insurance. Unfortunately, data on these types of coverage are not available until the 1982 survey. I therefore re-analyzed the data using only data from 1982 and subsequent years. The estimates from this subsample of the data – which did not differ from the estimates obtained using the whole sample – were not sensitive to whether those with Medicaid and military health insurance were excluded from the data.

$$\text{COVERAGE}_{ijt} = \alpha + \text{STATE}_j + \text{YEAR}_t + X_{ijt}\beta + \lambda^{-2}\text{ADOPT}_{ijt}^{-2} + \lambda^{-1}\text{ADOPT}_{ijt}^{-1} + \lambda^{0}\text{ADOPT}_{ijt}^{0} + \lambda^{1}\text{ADOPT}_{ijt}^{1} + \lambda^{2}\text{ADOPT}_{ijt}^{2} + \lambda^{3}\text{ADOPT}_{ijt}^{3} + \varepsilon_{ijt} \quad (1)$$

STATE and YEAR are fixed effects that control respectively for any fixed differences across states in coverage rates and for any yearly changes in coverage rates that are common across states. X is a vector of covariates. It is included to control for compositional changes in the sample along dimensions that may be related to the propensity to hold non-group coverage. It consists of a series of dummies for gender, race (white or non-white), education (less than high school, high school graduate, some college, college graduate and higher), marital status (married or not), and self-reported health status (excellent, very good or good versus fair or poor).[16] It also includes age, which is controlled for linearly.

ADOPT_{ijt}^{k} is coded 1 if it is k periods after implementation of the minimum standards in state j, and 0 otherwise. Because the data are, for the most part, biannual, I define a "period" as a two year grouping, except for the endpoints where periods are lengthened because of sample size concerns. Period 1 is therefore the first and second years after adoption and period 2 is the third and fourth years after adoption. Similarly, period 0 constitutes the year of adoption and the year prior to adoption. Period –1 is 2 or 3 years prior to adoption. The earliest period, period –2, however, consists of years –4 through –7 prior to adoption and the last period, period 3, consists of years 5 through 12 since adoption.[17]

The key parameters of interest are the pattern on the λ's on the ADOPT_{ijt}^{k} variables. λ^k indicates the estimated change in coverage rate after the reform has been in effect for k periods. Table 1 indicates substantial variation in this timing across states. The identifying assumption is that, absent the introduction of the minimum standards, there would have been no change in coverage rates in the periods that are after the reform relative to the periods that are before the reform, after controlling for state and year fixed effects and observable characteristics of the individuals. For non-positive values of k, the

[15] 4% of the sample has both group and non-group coverage. For these individuals, COVERAGE is coded as 1 since the individual has non-group coverage. In results not reported here, I found that the estimated effect of the reform is not sensitive to how such individuals were coded.
[16] I group health status this way because until 1982 the individual is given a choice of reporting their health status as "excellent", "good", "fair" or "poor." In 1982 and in subsequent years, the individual also has the option of reporting "very good".
[17] The percentage of observations in each category ranges from 21 percent in ADOPT^{1} to 14 percent in ADOPT^{-1}.

$ADOPT_{ijt}^{k}$ variables allow for a partial test of this assumption: I examine whether there are significant changes in coverage rates in the periods prior to the reform. I also examine the sensitivity of my findings to alternative specifications of equation (1), as described below.

4.2. Basic Results

Table 2 presents the results of estimating the basic specification in equation (1) by OLS.[18] The reference category is $ADOPT_{ijt}^{0}$ so that all the λ's are interpreted as the change in coverage relative to the period immediately prior to adoption. The first column presents the results from estimating equation (1) without covariates; the second column presents the results with covariates added. The estimated effects are not sensitive to the inclusion of covariates.

The results in Table 2 indicate a statistically significant 6 percentage point decrease in coverage associated with the regulation being in effect for one period (i.e. one or two years), relative to period 0. The estimated effect of the regulation increases monotonically as the time since the reform increases. When the reform has been in effect for three periods (5 years or more), the estimated effect is a 13 percentage point reduction in coverage relative to period 0. This long term effect is statistically significantly larger at the 1 percent level than the effect after only one period. The results in Table 2 also indicate that coverage rates are essentially flat in the periods prior to the adoption of the regulation.

On average, in period 0, 33 percent of the sample had non-group insurance. Therefore even the smallest estimated effects of the reform, in just the first period of its being in place, suggest that the imposition of minimum standards reduced non-group coverage by over 15 percent. The long-term estimates suggest a reduction in coverage of almost two-fifths.[19]

I perform several specification checks designed to check the robustness of the results and the validity of the identifying assumption. Appendix B reports the results of estimating equation (1) by probit, instead of by OLS. The estimated effects are not sensitive, in either sign, significance, or magnitude, to

[18] All of the standard errors from OLS estimates in Section 4 are adjusted for heteroscedasticity in the linear probability model and for state*year correlation in the error term.

this alternative functional form.[20] In another specification check, I add state-specific linear time trends to the basic specification in equation (1). This allows each state to be on its own trend. λ^k now measures the change in coverage k periods after implementation of the reform, after detrending for the state-specific linear time trend. A final set of specification checks stems from concern about the fact that the composition of states used to estimate the coefficient on any given $ADOPT^k$ varies with k. For example, since the last year that I use data from is 1986, the coefficient on $ADOPT_{ijt}^3$ (i.e. 5 or more years since adoption) is identified only by individuals in states where the regulation's first full year in effect was 1982 or earlier. If the effect of the reform varied across different states, or if the pre-period trends differed across states, my results could be contaminated by the changing composition of states used in identifying the coefficient on various ADOPT variables. To test for this, I re-estimate equation (1) on two different balanced panels of states. In one, I use states in which regulations were first in effect in 1979 or later. In the other, I use states in which regulations were first in effect in 1981 or earlier.

The estimated λ's from these alternative specifications, together with the estimates from the basic specification, are all graphed in Figure 2.[21] Each line represents the λ's estimated from a particular specification. Because state and year fixed effects are included in the regression that produced these coefficients, the graph adjusts for the fact that the reform occurred in different states in different years. It also controls for changes in coverage rates due to compositional changes along the observed X's. Figure 2 indicates that the basic pattern of coverage relative to the time that the reform was enacted is robust to these alternative specifications. Like the basic specification, all three alternative specifications indicate essentially flat coverage rates in the periods prior to the reform, a sharp decrease in coverage in the first period after the reform, and a continued decline in coverage in subsequent periods. The estimated

---

[19] Even these large estimates may understate the effect of the reform: some of the 65 to 68 year old individuals included in the sample may have bought their policies prior to the enactment of the reform (since one may buy a Medigap policy starting at age 65).

[20] In the remainder of Section 4, all reported results are from the OLS specification, but are not sensitive to estimation by probit instead.

[21] I choose as a starting value the mean non-group coverage rate in period –2 (35 percent).

magnitude of the effect of the reform in the basic specification lies in the middle of a tight range of estimates produced by the alternative specifications.

Table 3 presents the p values from various F-tests designed to confirm statistically what Figure 2 indicates graphically. The first row indicates the results from the basic specification. The other three rows report the results from the three alternative specifications. Column 1 indicates that the decrease in coverage between period 1 and period 0 is statistically significant at the 1 percent level in all four specification. Column 2 indicates that the long-term effect of the reform is statistically significantly greater than the short term effect of the reform at at least the 5 percent level in all specifications. The last column of Table 3 reports the results of a falsification test designed to examine the validity of the identifying assumption that absent the introduction of the legislation, there would have been no significant change in coverage in the periods after the introduction of the legislation relative to period 0. The test indicates that – across all four specifications – there is no evidence of a significant change in coverage in period 0 relative to period –1. This suggests that the finding of an effect of the reform is not merely due to spurious noise in coverage rates. It also suggests that the timing of the reform was not a response to pre-existing trends in non-group coverage rates.

I also investigated whether my results are sensitive to the particular age cut of the "young old" chosen. In results not reported here, the results are robust in sign, significance and magnitude to using a sample of 65 to 67 year olds or a sample of 65 to 69 year olds instead. Moreover, I find no evidence of any effect of the reform for older age groups. Specifically, I find no evidence of an effect of the reform on non-group coverage for 70 to 74 year olds or for individuals aged 75 and over. This is consistent with the conventional wisdom that individuals tend to buy their Medigap policies shortly after becoming eligible for Medicare.

In a sample of just 65 year olds, the basic pattern and magnitude holds, although only the long term effect of the reform (and not the short term) is statistically significant at the 5 percent level once I have reduced the sample size to one-quarter of the original sample of 65 to 68 year olds. The fact that the long term effect of the reform continues to be (statistically significantly) larger than the short term effect even among just 65 year olds suggests that the estimate of an increased effect of the regulations over time is not simply due to the fact that – with a sample of 65-68 year olds – more individuals are buying policies under the new regulations as time passes. Unfortunately, lacking high frequency pricing and detailed benefit information, I am unable to distinguish between the hypotheses that the increased effect of the regulations over time is due to a lag in enforcement and that it represents changes as the market readjusts to a new equilibrium.

4.3 Additional Results

The previous sub-section established that the minimum standards in the non-group market were associated with a large, robust, and statistically significant decrease in coverage by non-group private health insurance. In this section, I examine four additional issues.

First, I examine whether the decline in non-group coverage represented a net decline in insurance coverage or whether there was substitution toward other sources of insurance. The alternative sources of supplementary insurance are Medicaid and private group insurance. I therefore re-estimate equation (1) with a dependent variable that indicates whether the individual has Medicaid coverage and with a dependent variable that indicates whether the individual has group coverage.[22] I find no evidence that the reform is associated with a change in Medicaid coverage either in the whole sample or when the sample is restricted to those in the lowest education category (who presumably are mostly likely to be eligible for Medicaid). Similarly, I find no evidence of a change in group coverage associated with the introduction of minimum benefit standards in the non-group market. This lack of substitution from the newly-regulated non-group market to the unregulated group market or to Medicaid is consistent with the fact that group coverage and Medicaid both tend to be more comprehensive and cheaper than non-group policies. The non-group market, presumably, consists of individuals without access to group insurance or Medicaid.

Second, I examine whether the effect of the minimum standards varied across demographic groups or according to the stringency of the state's regulation. There is no evidence of a differential effect of the reform on coverage in the non-group market for the 8 states identified in Table 1 as adopting more stringent regulations than the Baucus minimum standards. There is also no evidence to suggest that the effects of the reform varied across observable characteristics of the individual such as martial status, gender, race, educational attainment or health status. [23]

However, I find a striking difference in the effect of the reform for Blue Cross and Blue Shield (BCBS) plans compared to other private plans.[24] There is no evidence of an effect of the reform on coverage rates for non-group Blue Cross and Blue Shield plans. The effect is entirely on coverage rates for other non-group plans.[25] The 1977 NMCES data indicate that BCBS plans were somewhat more likely than non-BCBS plans to meet the Baucus requirements (11 percent compared to 3 percent). More importantly, BCBS plans were substantially more likely to cover each of the six non-mandated benefits measurable in the data. For example, almost one-third of BCBS plans covered prescription drugs, compared to 10 percent of non-BCBS plans. Additionally, 95 percent of the BCBS plans covered the Part B deductible, compared to 56 percent for the non-BCBS plans. It is therefore not surprising that, since BCBS plans offered substantially more

coverage, individuals considering purchase of a BCBS plan were less likely to have their purchase decisions affected by the minimum standards.[26]

Third, I examine whether spouses appear to make their coverage decisions independently or jointly. In the period prior to the reform, the most common outcome for couples is for either both members of the couple to have non-group insurance (27 percent) or for neither to have it (67 percent); in only 6 percent of

---

[22] Because I only have a good measurement of Medicaid coverage starting in 1982, I restrict the analysis of whether there was substitution to Medicaid to the 1982, 1983, 1984 and 1986 data. The results for non-group coverage reported above are robust to a similar restriction.

[23] For these latter two results, to try to account for the fact that decisions may be made jointly within a married household, I verified that the results were not sensitive to estimation separately by gender.

[24] Less than 1 percent of non-group plans are HMOs. The others are split evenly between BCBS and other private plans.

[25] For the 11 percent of individuals with non-group policies in my sample who have multiple policies, I define policy type based on the main policy held. These results are not sensitive to restricting the sample to only those with single policies.

[26] Unfortunately, the 1987 NMES data do not distinguish between BCBS and other private plans. In the next section, I am therefore unable to examine the effects of the reform on non-mandated benefits separately for these two classes.

couples did one spouse and not the other have non-group health insurance. A multinomial logit of the effect of the reform on these three possible outcomes indicates that the reform was associated with a decline in insurance coverage for either both members of a couple or for neither. This may be because couples make joint decisions about insurance purchases, or because they have similar unobserved preferences over health insurance.

Finally, I consider how the reform affected the holding of multiple Medigap policies. As discussed above, one of the main motivations for the minimum standards was a concern that individuals were making uninformed decisions. The holding of multiple policies was taken as one indication of this (Fox, Rice and Alecxih, 1995). On average, 11 percent of the 65 to 68 year olds who had non-group coverage in period 0 had more than one non-group policy. I find no evidence that the reform was associated with a switch from multiple plans to single plans. Instead, the reform appears to be associated with a switch from either single or multiple plans to no plans, in similar proportions.

4.4 Understanding the magnitude of the decline in non-group insurance coverage

The 40 percent long-term reduction in non-group insurance coverage associated with the regulation suggests that, for many individuals, the mandated minimum benefits were not valued at their cost to the individual. There are several possible explanations for this. If markets function perfectly and individuals make rational decisions about insurance purchases, then the mandated minimum benefits function simply as a tax on the purchase of non-group insurance. The gross tax is equal to the premium increase associated with upgrading policies to comply with the minimum standards. The net tax is lower than the gross tax by the utility gain from the additional insurance.

Using detailed information on medical expenditures and sources of payment from the 1977 NMCES, I estimate that, ignoring moral hazard effects and adverse selection effects, the mandated Baucus benefits would have raised the expected payments (and presumably the premium) from a non-group insurance policy by $168 in 1999 dollars. This represents an average 30 percent increase in non-group premiums. If the marginal utility associated with the additional mandate insurance is zero, then this 30 percent increase in premiums represents a 30 percent net tax on the purchase of non-group insurance.

Most estimates of the price elasticity of demand for health insurance lie in the range of –0.5 to –1 (Cutler 1996). If the insurance is priced actuarially fairly and the marginal utility from complying with the mandate is zero, the predicted premium increase could therefore explain up to three-quarters of the long-term decline in non-group coverage associated with the minimum standards.

Several factors may make the marginal utility from the additional insurance very small. First, the existence of de-facto partial insurance for some of the mandated benefits reduces the value of private coverage. Partial insurance for long hospital stays may be available through states' medically needy programs, which provides Medicaid to elderly individuals who have high medical expenses relative to their income in a given year, and through the provision, by some hospitals, of uncompensated care for those unable to pay. In addition, large medical expenses are partially co-insured through the federal income tax system which allows the deduction of medical expenses in excess of a certain fraction of adjusted gross income (AGI). Second, consumer misinformation – itself a potential motivation for the imposition of minimum standards – can help explain why the mandated benefits are not valued or are valued at less than their actuarial cost. If individuals overestimate the amount of coverage provided through Medicare, they will undervalue the insurance protection offered by the mandated minimum benefits. Similarly, if individuals have difficulty evaluating low probability, high cost, risks, they may underestimate the value of coverage for the financially catastrophic but extremely rare event of a long hospital stay.

Finally, several other departures from a perfect markets, perfect information world may increase the price of insurance above its actuarial cost and thus increase the gross tax associated with complying with the mandated minimum benefits. For example, administrative loads, insurance company profits, and adverse selection may all serve to increase the price of the mandated minimum benefits.[27] As the price increases and hence the gross tax on purchasing insurance increases, larger declines in insurance coverage associated with imposing the minimum standards are to be expected.

5. Effect of Minimum Standards on Coverage of Non-Mandated Benefits

5.1 Data and Empirical Strategy

The only available detailed data on the benefits covered by an individual's Medigap policy are the 1977 NMCES and its companion survey, the 1987 NMES.[28] The long time period between these two cross-sections make the results of this section necessarily more speculative than those in the previous section. Nevertheless, I am able to use two different control groups who have Medigap insurance but are not affected by the minimum standards to try to control for other changes that may have occurred during this 10 year period in the demand or supply of Medigap insurance.

I examine the effect of the reform on the probability of the provision of the six different non-mandated benefits measurable in the data. Two benefits – coverage of the Part A and Part B deductible – cover the remaining cost sharing provisions in Medicare beyond those included in the mandated minimum package. Three benefits cover services for which Medicare provides only limited coverage: home health care coverage, coverage of care in a skilled nursing home facility, and coverage of inpatient psychiatric treatment. The final benefit– prescription drug coverage – covers a service not covered by Medicare.

I limit the sample to individuals who are covered by private health insurance and by Medicare, and are not covered by Medicaid. I limit the analysis to policyholders and do not included dependents since dependents presumably have the same benefits as policyholders and this would involve double counting. I also limit my sample to those who are retired, whom I define as those who are not in the labor force.[29] The data are at the policyholder level. I therefore exclude anyone who has both non-group insurance and group insurance (approximately 7 percent of the sample) because in such cases I cannot tell which market a given benefit comes from.

---

[27] Moral hazard – another potential cause of higher insurance prices – is unlikely to be as much of a source of price increases in the Medigap market compared to other insurance markets. Most of the moral hazard costs of Medigap are born by the public Medicare program rather than private insurers (Ettner 1997).

[28] Benefit information is obtained by contacting each individual's source of insurance for policy details and then coding up these details. This is considerably more reliable than self-reported benefit information.

[29] The control strategy used in this section necessitate the restriction to retirees, as discussed below. This is not a severe restriction as most individuals over 65 are retired. The results in section 3 are not sensitive to this restriction.

The empirical strategy is to compare changes in benefit coverage rates for privately insured individuals affected by the reform (the treatment group) to changes in benefit coverage rates for a control group of privately insured individuals who were not affected by the reform. The basic estimating equation is:

$$BENEFIT = \beta_1 Y1987{*}MINSTAND + \beta_2 Y1987 + \beta_3 MINSTAND + X\beta_4 + \varepsilon \qquad (2)$$

The dependent variable BENEFIT is a binary indicator of whether the individual's health insurance coverage includes a given benefit or not. I estimate equation (2) separately for the six different non-mandated benefits measurable in the data that are described above.

MINSTAND is an indicator variable for whether the individual is in a group ultimately subject to the minimum standards or not. This treatment group is defined as individuals with non-group coverage who were 65 to 68 when a reform was introduced in their state. This age cut was chosen in keeping with the results of the previous section. Because I do not know which state an individual is in in my data, my treatment group becomes any individual with non-group insurance aged 65 and 71.[30] Y1987 is an indicator variable for whether the year is 1987 or not. X is a matrix of covariates, similar to that used in Section four, that controls for observed compositional changes. It consists of dummies for gender, region of the country, whether the individual lives in an SMSA, race (white or non-white), marital status, education (less than high school, high school degree, some college, college degree or higher) and self-reported health status relative to others' their age.[31] In addition, age is controlled for linearly.

The key variable of interest – Y1987*MINSTAND – measures changes in benefit coverage between 1977 and 1987 among the treatment group relative to changes in benefit coverage over the same period for a control group. The identifying assumption is that trends in benefit coverage rates among these two groups would have been similar if the reform had not occurred. Below, I discuss a check of this identifying assumption.

[30] Anyone who is 71 or younger in 1987 was no older than 68 in 1984 (the last year that was the first full year for regulation to be in effect in any state) and therefore is in the age group that I found was affected by the reform in the previous analysis.

I estimate equation (2) using two different definitions of the control group. Since the minimum standards applied to the non-group market and not to the group market, Control Group I consists of individuals aged 65 to 71 who have group health insurance.[32] Control Group II consists of individuals with non-group insurance who are 78 and older. Individuals 78 and older in 1987 were at least 70 in 1979, which is the first year after 1977 that the regulations were in effect. Such individuals are therefore likely to have purchased their Medigap policy prior to the enactment of the Baucus regulations.[33] They could therefore continue to renew these policies without being subject to the minimum standards. Moreover, they could change any benefit that was attached to the policy as a "rider benefit" without becoming subject to the minimum standards regulation. Evidence suggests that coverage for the Part A and Part B deductibles are commonly sold as riders, whereas riders for other benefits are rare.[34] Therefore Control Group II only serves as a control for confounding supply or demand changes in Part A and Part B deductible coverage.

It is worth recalling that the evidence presented in Section four suggested that neither of these putative control groups were themselves directly affected by the minimum standards. Individuals do not appear to have switched from the non-group to the group market in response to the minimum standards, and there was no evidence of an effect of the regulations on coverage rates in the non-group market for individuals aged 70 and on older.

---

[31] State identifiers are not available in the NMCES and NMES data. I therefore use the 9 region classifications and the SMSA designation to try to capture geographic effects.

[32] Federal legislation introduced in 1982 required that employers offer the same health insurance packages to employed workers under and over age 65. Although Glied and Stabile (forthcoming) find very little compliance with (or effects of) the legislation, this change in regulation for employed workers with group coverage cautions against including the employed in the analysis. I therefore limit my sample to retirees. The 1982 legislation did not apply to group coverage of retirees.

[33] Most individuals purchase non-group Medigap policies shortly after becoming eligible for Medicare. Although there are no good data on initial age of purchase, conversations with individuals in state regulatory agencies, individuals active in the debate over and design of the Baucus regulations, and academics who study the Medigap market indicated that all shared the sense that – at the time period studied – most people bought their non-group Medigap policies shortly after becoming eligible for Medicare, if they bought them at all.

[34] Conversations with several people familiar with the market in the late 70s and early 80s each suggested that coverage for the Part A and Part B deductibles are commonly sold as riders. In addition, data from all policies sold by the top 5 insurance companies in 5 states in 1991 also indicate that riders for Part A and Part B deductible coverage are very common. Rice, Graham and Fox (1997) provides information about these data. I am extremely grateful to Tom Rice for providing these data.

A potential problem with both of these control strategies, however, is that in the 1977 data four states have already introduced the minimum standards. Three states implemented minimum standards in 1976 and California implemented them in 1974. To the extent that the effect of these reforms was already partly felt by 1977, I will underestimate the effect of the reform on non-mandated benefit coverage. This potential downward bias is larger with the second (age-based) control strategy because it has the added issue that some of the control group is mis-classified as unaffected in the post-period (i.e. individuals who are older than 78 in 1987 could have been as young as 68 in 1976 and therefore in the four early states, may have bought policies under the Baucus regulation).

5.2 Results

Table 4 presents mean coverage rates and the estimated effect of the reform on coverage of the non-mandated cost sharing provisions of Medicare: the Part A and Part B deductible. The empirical strategy is, a priori, most convincing for these particular benefits for two reasons. First, the deductibles are arguably the benefits for which the 10 year time lag in the data is least troubling. It is harder to think of demand and supply shocks affecting deductible coverage than affecting coverage of other benefits for which the underlying nature of the risk being insured may well be changing between 1977 and 1987. Second, these are the only benefits on which I can estimate the effect of the reform using both control strategies. For the other benefits, I will have to rely on only the group control strategy.

The pre-treatment mean coverage rates given in column (1) look very similar across the treatment and control groups for each of the benefits. A comparison of columns (1) and (2) indicates that coverage for each of the deductibles is decreasing between 1977 and 1987 for all three groups. Columns (3) and (4) report the estimated coefficient on Y1987*MINSTAND from estimation of equation (2) by OLS, estimated without and with the X vector of covariates respectively. The results are not sensitive to the inclusion of covariates. The top portion of Appendix C indicates that these results are also not sensitive in sign or significance to estimation of equation (2) by probit instead of by OLS.

The two control strategies yield similar results. There is no evidence of a decline in coverage of the Part A deductible in the treatment group relative to either control group. Both control strategies, however, indicate a large and statistically significant decline in coverage of the Part B deductible among the treatment group relative to the control group.

In the OLS specification (although not in the probit), the magnitude of the estimated reduction in the Part B deductible associated with the minimum benefit standards is considerably lower when the older age group (Control Group II) is used as a control than when the group market (Control Group I) is used as a control. This is consistent with the potential, discussed above, for a larger downward bias in the estimated effect of the reform using Control Group II rather than Control Group I. Nevertheless, both control strategies indicate a substantial and statistically significant decrease in coverage of the Part B deductible associated with the reform.

Table 5 presents mean coverage rates and the estimated effect of the reform on the provision of four benefits not covered by Medicare. As discussed above, for these four benefits, I can only use those in the group market (Control Group I) as a control group. Column (1) indicates that there are (statistically) significant differences in benefit coverage rates in 1977 between the group and the non-group market for these four benefits. For three of the four benefits, coverage rates are higher in the group market than in the non-group market for each of these benefits. A comparison of columns (1) and (2) indicates that coverage rates for all of the benefits but skilled nursing home care are falling between 1977 and 1987 in the treatment group. In contrast, coverage rates for all four benefits are rising over the time period in the control group. Columns (3) and (4) of Table 5 present the estimated coefficient on Y1987*MINSTAND from estimation of equation (2) by OLS, without and with controlling for covariates respectively. They indicate that the reform is associated with a substantial and statistically significant decline in the coverage of outpatient prescription drug benefits, inpatient psychiatric care, and home health care, but not in the coverage of skilled nursing home care. The estimated effects are statistically significant at the 1 percent level and are robust in significance and magnitude to the inclusion of covariates in the regression. The bottom portion of Appendix C indicates that they are also robust to estimation of equation (2) by probit instead of by OLS

The empirical results are thus consistent with the minimum standards reducing coverage of some of the non-mandated benefits in the regulated insurance market. But a causal interpretation of the estimated coefficients on Y1987*MINSTAND requires the identifying assumption that, absent the minimum benefit standards, benefit coverage rates would have been following the same trends in the treatment and control groups between 1977 and 1987. Ideally, benefit coverage data from a period prior to the introduction of the minimum benefit standards would be used to help assess the validity of this identifying assumption. Although such data are not available, I can point to a weaker piece of evidence in favor of the identifying assumption. The NHIS data allow me to examine whether insurance coverage rates (but not benefit coverage rates conditional on having insurance) are trending similarly in the treatment and control groups in state-years prior to the introduction of the reforms. For each year (1978, 1980, and 1982), I cannot reject the null hypothesis that the year fixed-effects are the same in the treatment group relative to each of the control groups. Nor can I reject the joint hypothesis that all three year fixed effects are the same in the treatment group relative to each of the control groups.

5.3 Understanding the decrease in coverage of non-mandated benefits.

As discussed in Section three, there are several theoretical explanations for the estimated decrease in coverage of non-mandated benefits associated with imposing minimum benefit standards. Although the available data do not present a means of testing definitively among these alternative explanations, it is possible to shed some light on their relative merits. To begin, it is worth noting that there are two different classes of explanations for the estimated declines in non-mandated benefits in the regulated insurance market. These could be the result either of compositional changes in the pool of privately insured or of changes in the demand or supply of non-mandated benefits. However, we saw in Section four that the declines in non-group insurance coverage were confined to those with non-BCBS plans and that these plans were substantially less likely to cover non-mandated benefits than the BCBS plans. Therefore compositional changes in the insured population are unlikely to explain the observed decline in non-mandated benefit coverage among the insured.

The demand-side mechanisms in a perfect information, perfect competition model are unlikely to explain all of the observed declines in non-mandated benefits. Coverage of the Part A and Part B deductibles may well be substitutes for coverage of the mandated Part A and Part B copayments. Therefore the decline in coverage of the Part B deductible could potentially be explained by a substitution effect. However, other benefits such as inpatient psychiatric care and prescription drug coverage are more likely to be complements to the mandated coverage for hospital stays and doctor visits. Therefore substitution effects cannot explain their decline. Furthermore, while the income effect of the price increase that is likely associated with providing the mandated benefits may be part of the explanation for the decrease in coverage of non-mandated benefits, it is unlikely that it can fully explain the magnitude of the estimated decreases. As discussed above, I estimated that, ignoring moral hazard and adverse selection, the mandated Baucus benefits would have raised the expected payments (and presumably the premium) from a non-group insurance policy by $168 in 1999 dollars. This $168 represents 0.8% of median income among those with non-group Medigap policies in the NMCES. Estimates of the income elasticity of demand for medical care range from 0.2 to 1.[35] Given median total health care spending among the non-group privately insured in the NMCES of $490 in 1999 dollars, an income elasticity of 1 would translate into a reduction in demand for health care of less than $3.90 in 1999 dollars. Yet the NMCES data suggest that the loss of the Part B deductible coverage, alone, would result on average in a loss of health benefits of ten times this amount.

This suggests a role for supply-side effects of the minimum standards in helping to explain the decline in non-mandated benefits among those with non-group insurance. The empirical evidence does not, however, support a joint costs explanation for the decrease in non-mandated benefits. For the NHIS data indicated no change in non-group coverage rates in age groups from 70 up. In the presence of substantial joint costs, we would expect that the forced dropping of non-compliant policies for the "young old" would have affected the profitability of offering such policies to the "older old." We should therefore

[35] Newhouse (1992). Of course demand for medical care and demand for health insurance are not the same thing. But presumably the income elasticities are roughly similar.

have seen changes in coverage for all ages and not just those for those ages directly affected by the minimum standards.

Two pieces of evidence lend support to the possibility that, under asymmetric information, the mandated benefits may, by restricting the ability of insurance companies to offer a spectrum of policies designed to screen individuals by risk type, impair the functioning of the market for policies that supply more insurance than the mandated minimum. For one thing, it is interesting to note that there is evidence of a decline in the Part B deductible coverage but not in the Part A deductible coverage. Given that the requirement of coverage of the Part B copay was more binding than the requirement of coverage of the Part A copays, the fact that we observed a decrease in the Part B deductible is consistent with the notion that the Part B copay had previously been being used as a screen.

Moreover, recall that in the stylized Rothschild and Stiglitz model, mandated minimum benefits can produce a switch from a separating equilibrium to a pooling equilibrium. A less extreme version of this stylized prediction (once, for example, preference heterogeneity is allowed) is that, under asymmetric information, the minimum standards should produce a decrease in insurance plan dispersion. When I compare changes in concentration of plan types in the non-group market between 1977 and 1987 with changes in concentration in the group market over this same period, the results are striking.[36] For example, the Herfindahl measure of plan concentration almost doubles in the non-group market (from 0.11 to 0.21) while remaining constant in the group market (at 0.12). Figure 3 shows plan market share by plan rank in the non-group market prior to the introduction of the minimum standards (1977), and after their introduction (1987). The empirical cumulative distribution function of plan shares in the non-group market in 1987 lies everywhere above the 1977 empirical cumulative distribution function; there is no such clear ranking of the two periods in the group market (not shown). Moreover, using McFadden's (McFadden 1989) test for first order stochastic dominance, I am unable to reject the null hypothesis that the 1977 distribution of plan shares in the non-group market first order stochastic dominates the 1987

[36] I define a plan based on which of the 6 non-mandated benefits measurable in the data are covered. In practice, this produces about 20 different plans purchased in a given year and market.

distribution in the non-group market. I can reject this null in the group market, however, at the 10 percent level. The evidence is thus consistent with the prediction from an asymmetric information model that the minimum standards will be associated with a decrease in plan dispersion in the regulated market.

6. Welfare Implications

We have seen that the minimum standards resulted in a large decline in non-group insurance coverage. For individuals covered by non-group plans after the reforms, however, the newly mandated benefits provided substantial additional insurance coverage. However, there is also evidence that the minimum standards are associated with a decline in coverage of non-mandated benefits in the regulated insurance market. To evaluate the net impact of these various changes in insurance coverage, I examine the distribution of out of pocket health expenditure risk for insured individuals before and after the imposition of the minimum standards. This provides an estimate of how the risk distribution faced by individuals changed as the result of the minimum standards. I then evaluate the welfare impact of these changes in risk bearing.

The estimates are based on the sample of 989 Medicare recipients in the 1977 NMCES, who are not on Medicaid and who have private, non-group insurance.[37] The NMCES provides individual-level data on total health expenditures for several different health expenditure categories. It also provides a decomposition of that expenditure into the portions covered by private health insurance, by Medicare, and by Medicaid, and the portion paid out of pocket.

Figure 4 provides a visual summary of the distribution of out of pocket medical expenditures (i.e. medical expenditure risk faced). The figure shows mean out of pocket medical expenditures by health expenditure decile under four different insurance arrangements. The solid black bars show this distribution for those with non-group private health insurance coverage prior to the reforms. The other

[37] An important mandated benefit from the Baucus regulations is coverage for hospitals stays beyond 150 days at which point Medicare ceases to pay anything. However, because such long hospital stays are extremely rare, no one in my 989-person data set has such long stays. So that the welfare benefit from the mandated coverage for long stays is not undervalued, I adjust the risk distribution in my sample to take account of the fact that data from a 20% sample of Medicare beneficiaries in 1984 indicates that 0.1% of the elderly have hospital stays in excess of 150 days. The median length of stay, conditional on it being above 150 days, is 259 days in total.

three bars simulate the medical expenditure distribution under alternative insurance arrangements for these individuals. The simulations are done by adjusting the portion of particular expenses paid out of pocket and those paid by private insurance to reflect the change in insurance coverage. All of these calculations assume that total medical expenditures are unaffected by the change in insurance status. They thus ignore possible moral hazard effects from gaining or losing insurance.[38]

A comparison of the distribution of out of pocket medical expenditures under different insurance arrangements provides an understanding of the tradeoffs associated with imposing minimum benefit standards. The decline in non-group insurance coverage associated with the minimum standards created an increased risk of out of pocket medical expenditures. This can be seen by comparing the distribution of medical expenditures for individuals with pre-reform non-group insurance plans (the solid black bars) to the simulated distribution of expenditures for these individuals if they lose private coverage (the adjacent speckled bars). The pre-reform insurance lowers the medical expenditure risk faced, particularly at the high end of the distribution.

The bars with horizontal lines shows the simulated medical expenditure risk distribution if the pre-reform plans are upgraded as necessary to comply with the minimum standards. A comparison of this distribution with the distribution for individuals with pre-reform plans (sold black bars) gives some sense of the magnitude of the reduction in risk associated with upgrading the plans to comply with the minimum requirements. Figure 4 indicates that out of pocket expenditures are lower at every decile when the pre-reform plans are upgraded to meet the minimum standards. However, this comparison does not take account of the estimated decreases in non-mandated benefit coverage that are also associated with the minimum standards. It thus provides an upper bound of the increase in insurance coverage associated with the minimum standards for those who maintained coverage.

The health expenditure categories in the NMCES are unfortunately not broken down finely enough to allow simulation of the change in risk exposure associated with the declines in coverage for the non-

[38] As discussed previously, much of the moral hazard effects of the private supplementary insurance are born by the universal, public Medicare program (Ettner 1997). This helps to mitigate the impact of ignoring moral hazard effects

mandated benefits estimated in Section five. I can, however, simulate the distribution of medical expenditure risk if all insurance except for the mandated minimum benefits is dropped. This is shown by the gray bars. These provide a lower bound of the increase in insurance coverage associated with the minimum standards for those who maintained coverage. It is interesting to compare the distribution of out of pocket medical expenditures faced by those with pre-reform plans (sold black bars) with the distribution faced by those with insurance plans that cover the mandated benefits only (gray bars). Figure 4 reveals that the medical expenditure risk faced by those whose insurance plans now just meet the minimum requirements is lower than the risk faced by those with pre-reform plans for low levels of expenditure but higher than the risk faced by those with pre-reform plans at higher levels of expenditure.

Evaluation of the welfare impact of the changes in risk bearing simulated in Figure 4 requires assumptions about the individual's utility function.[39] I assume that each individual's utility can be written as a function of his non-health consumption (c) and his out of pocket medical expenditure (m) as $u(c-m)$.[40] m is a random variable with probability density function $f(m)$ and support $[0,\overline{m}]$. $f(m)$ depends both of the distribution of random health shocks, and on the nature of any health insurance held. As the nature of private insurance changes, $f(m)$ changes, thus altering the individual's expected utility which is given by $\int_0^{\overline{m}} u(c-m)f(m)dm$.

A useful way to capture welfare under different distributions of m is through the risk premium (π). The risk premium is the maximum amount that a risk averse individual would be willing to pay to completely insure against the random variable m. It is defined implicitly by:

(3) $u(y-\pi)=\int_0^{\overline{m}} u(c-m)f(m)dm$

---

when examining out of pocket medical expenditures under alternative private insurance arrangements.

[39] This calculation closely follows that in Feldstein and Gruber (1995). More details on the specifics of the calculation can be found there.

[40] I abstract from the possibility of state-dependent utility functions in which the utility function varies according to health status.

For each individual in the data with non-group insurance, I calculate the risk premium implicitly defined by equation (3) under different insurance arrangements. As the risk an individual faces increases, the risk premium increases as well. Thus an increase in risk premium reflects a decrease in welfare. For each insurance arrangement, I use the empirical distribution of out of pocket medical expenditures simulated and summarized in Figure 4 for $f(m)$, and the individual's income for c.[41] I assume a constant relative risk aversion utility function.

Table 6 reports the average change in welfare across individuals associated with moving from the pre-reform insurance status quo to a given insurance status. This welfare change is defined as the risk premium under the pre-reform insurance plans minus the risk premium under the new insurance status. I report results for two different coefficients of relative risk aversion: 1 and 3. Table 6 indicates that the welfare change associated with a given change in insurance status is increasing (in absolute value) as risk aversion increases.

Column 1 indicates an average welfare loss (in 1999 dollars) of \$18 to \$476 associated with losing non-BCBS pre-reform coverage.[42] Column 2 indicates an average welfare gain of \$4 to \$196 associated with upgrading pre-reform plans to comply with the minimum standards while not altering the provision of non-mandated benefits.[43] This therefore represents an upper bound on the welfare gain for individuals who keep non-group insurance associated with the minimum standards. Finally, the estimates in column 3 provides a lower bound of the average welfare change associated with the minimum standards for individuals who retain coverage. They are based on the assumption that individuals drop all non-mandated benefits. This is estimated to result in a net average welfare *loss* in changing from pre-

[41] Following Feldstein and Gruber (1995), I adjust the various distributions of medical expenditure risk to keep their mean constant. Any change in the mean risk should be captured in a change in premium and thus is simply a transfer between the insurer and the insured.
[42] This estimate is based on a loss of coverage of non Blue Cross and Blue Shield (BCBS) plans only since the results in Section 3 indicate a decrease in coverage only among non BCBS plans. The welfare loss is considerably higher if losses of BCBS plans are simulated as well, since these plans tended to be much more comprehensive.
[43] Of course, in a world of perfect information and perfect markets, the increased insurance coverage experienced by those who retained coverage and now have to buy the mandated benefits are welfare losses. Increases in insurance coverage associated with the minimum standards may produce welfare gains, however, if they are designed to address a market failure that results in sub-optimal insurance purchases.

reform insurance plans to insurance plans that cover only the mandated minimum benefits of $29 to $734.[44]

To calculate the average net welfare change associated with the Baucus regulations, I can compare the welfare change experienced by the two-fifths of privately insured who lost private insurance (column 1) to the welfare change experienced by the three-fifths of the privately insured who retained private coverage (column 2 or 3).[45] Even under the conservative assumption that there was no decrease in non-mandated benefits among those who retain insurance coverage, this calculation (based on columns 1 and 2) suggests an average net welfare loss of $5 in 1999 dollars for a coefficient of relative risk aversion of 1, and an average net welfare loss of $73 in 1999 dollars for a coefficient of relative risk aversion of 3. Under the liberal assumption that all of the non-mandated benefits were dropped, the estimated average net welfare losses (based on column 1 and column 3) rise to $24 and $631 respectively.

While these average welfare losses are not large in absolute dollar terms, even the conservative estimates of welfare loss represent 15 to 28 percent of the average welfare gains associated with having the pre-reform coverage (i.e. the negative of the results in column 1). Moreover, since approximately one-third of the 22 million Medicare beneficiaries in 1977 had private, non-group coverage, the conservative estimates suggests an aggregate welfare loss (in 1999 dollars) of $33 million to $530 million associated with imposing minimum standards on the non-group market for private supplementary health insurance for the elderly.

7. Conclusion

This paper has examined the empirical consequences of imposing large, binding minimum benefit standards on a voluntary private health insurance market. Although the economic rationales for minimum

[44] The estimates indicate that maintaining coverage even when it results in a net welfare loss can be consistent with optimizing behavior. For the estimated welfare loss from dropping coverage is even larger when it is not confined to the non-BCBS plans. In this case, the numbers are $36.9 and $943.0 for coefficients of relative risk aversion of 1 and 3 respectively.

[45] This welfare calculation considers only private welfare. Social welfare losses from the decrease in private coverage will be lower than private welfare losses because of the negative moral hazard externality that the private policies exert on the public Medicare program.

standards in insurance markets have been much discussed, little is known about their empirical consequences. I find that the minimum standards are associated with a substantial decline in insurance coverage. The central estimate suggests a long term decline of 40 percent in non-group Medigap coverage associated with imposing minimum standards on this market. There is no evidence of substitution from the regulated market to alternative sources of insurance. Additional evidence suggests that the minimum standards may have reduced the provision of non-mandated benefits. Finally, a conservative welfare calculation that compares the loss in welfare experienced by those who did not purchase insurance coverage as a result of the minimum standards with the increase in welfare from the additional insurance provided by the minimum standards for those who did purchase coverage indicates a net welfare loss associated with imposing these minimum standards.

This paper has concentrated on evaluating the impact of minimum benefit standards. Implicitly, the evaluation has been made relative to a benchmark of no government intervention. A more nuanced approach, and a fruitful direction for further work, would be to consider the impact of minimum benefit standards relative to alternative policy interventions that the government might undertake.

Consumer misinformation and adverse selection – which are potential economic rationales for minimum standards – also serve to reduce the perceived benefit and increase the cost of complying with the minimum standards. They therefore increase the likelihood that the imposition of voluntary minimum standards will result in large declines in insurance coverage. Alternative policies such as public provision of insurance benefits can avoid the consumer misinformation issue and – through a universal mandate – solve the adverse selection problem. However, such policies come at the cost of removing any role for consumer sovereignty. Minimum standards in a voluntary setting respect consumer sovereignty on the decision of whether to purchase any private insurance, but impose constraints conditional on so doing. If differences in preferences such as risk aversion are large between those who do and do not purchase insurance relative to differences among those who purchase insurance, minimum standards in a voluntary setting may offer a considerable advantage over compulsory public insurance.

References

American Hospital Association. *A.H.A. Guide to the Health Care Field.* Chicago, IL, 1978.

Cafferata, Gail Lee. 1984. "Private Health Insurance Coverage of the Medicare Population." National Health Care Expenditures Study, Data Preview 18. DHHS Publication No. (PHS) 84-3362. U.S. Department of Health and Human Services, National Center for Health Services Research.

Cafferata, Gail Lee. 1985. "Private Health Insurance of the Medicare Population and the Baucus Legislation." *Medical Care*, 23(9):1086-96. September.

Cutler, David. 1996. "Comment" to "Fundamental Tax Reform and Employer-Provided Health Insurance," Jonathan Gruber and James Poterba, in Economic Effects of Fundamental Tax Reform, H. Aaron and W. Gale (eds), Washington: Brookings Institution, pp.125-162.

Encinosa, William. Forthcoming. A comment on Neudeck and Podczeck's "Adverse Selection and Regulation in Health Insurance Markets." *Journal of Health Economics.*

Ettner, Susan. 1997. Adverse selection and the purchase of Medigap insurance by the elderly. *Journal of Health Economics* (16): 543-562.

Feldstein, Martin and Jonathan Gruber. 1995. A Major Risk Approach to Health Insurance Reform. NBER Reprint No. 2151.

Fox, Peter D., Thomas Rice and Lisa Alecxih. 1995. "Medigap Regulation: Lessons for Health Care Reform." *Journal of Health Politics, Policy and Law.* Vol. 20, No. 1, Spring. pp 31-48

Glied, Sherry and Mark Stabile. Forthcoming. Avoiding Health Insurance Crowd-Out: Evidence from the Medicare as Secondary Payer Legislation. *Journal of Health Economics.*

Gruber, Jonathan. 1994. State-mandated benefits and employer-provided health insurance. *Journal of Public Economics* 55: 433-464.

Madrian, Brigitte. 1994. The Effect of Health Insurance on Retirement. *Brookings Papers on Economic Activity*, 1, pp.181-232.

McCall, Nelda, Thomas Rice and Arden Hall. 1983. Medigap: Study of Comparative Effectiveness of Various State Regulations: Final Report. SRI International. Menlo Park, CA.

McCormack, Lauren A., Peter D. Fox, Thomas Rice and Marcia Graham. 1996. "Medigap Reform Legislation of 1990: Have the Objectives Been Met?" *Health Care Financing Review*, Fall, Volume 18, Number 1. pp.157-174.

McFadden, Daniel. 1989. Testing for Stochastic Dominance. In Studies in the Economics of Uncertainty, Thomas Fomby and Tae Kun Seo (eds), New York: Springer-Verlag, pp.113-134.

Merritt, Richard E. and Donna B. Potemken (eds). 1982. Medigap: issues and update. The Intergovernmental Health Policy Project, George Washington University, Washington D.C. June.

Neudeck, Werner and Konrad Podczeck. 1996. Adverse selection and regulation in health insurance markets. *Journal of Health Economics*, (15): 387-408.

Newhouse, Joseph P. 1992. "Medical Care Costs: How Much Welfare Loss?" *Journal of Economic Perspectives*, 6(3): 3 –21.

P.L. 96-265. Jun 9, 1980. "Voluntary Certification of Medicare Supplemental Health Insurance Policies." [a.k.a. The Baucus Amendments]

Rice, Thomas, Marcia Graham and Peter Fox. 1997. The Impact of Policy Standardization on the Medigap Market. *Inquiry* 34: 106-116, Summer.

Ronnen, Uri. 1991. "Minimum quality standards, fixed costs, and competition." *RAND Journal of Economics* Vol 22, No. 4, pp. 490 – 504.

Rothschild, Michael and Joseph Stiglitz. 1976. An Essay on the Economics of Imperfect Information. *Quarterly Journal of Economics* 90 (November), pp.629-649.

U.S. General Accounting Office. 1986. Medigap Insurance: Law Has Increased Protection Against Substandard and Overpriced Policies. Pub. No. GAO/ HRD-87-8. Washington D.C.

U.S. General Account Office. 1991. Medigap Insurance: Better Consumer Protection Should Result from 1990 Changes to Baucus Amendment. Pub. No. GAO/HRD-91-49. Washington D.C.

U.S. House of Representatives, Select Committee on Aging. 1978. Abuses in the Sale of Health Insurance to the Elderly in Supplementation of Medicare: A National Scandal. Comm. Pub. No. 95-160. Washington DC

Van Ellet, T. (ed). 1979. Medigap: State Responses to Problems with Health Insurance for the Elderly. The Intergovernmental Health Policy Project, George Washington University, Washington D.C. October 30.

Table 1: Implementation of Minimum Benefits Standards for non-group Medigap policies.

First Full Year of Regulations in Effect	STATES	Percentage of National Health Interview Survey Sample[46]
1974	CA*	9.7
1977	IL, CT, MN*	8.1
1979	RI, PA, WI*	9.1
1980	MA**	2.5
1981	GA, OR, FL, NH, NV, VT, NE, WY*	10.9
1982	IA, SC, AK, AZ, CO, AL, ND, UT, NJ, AR, VA*, WV*, NY*, WA, TN	27.0
1983	ME, HI, IN, KS, OK, OH, ID, MS, DE, KY, TX, MT, MO, SD, NM, LA, NC, MD	28.5
1984	MI, DC	4.3

* Denotes regulation that established classes of policies each with their own minimum benefit standards. In all of these cases, the least comprehensive category had minimum benefit standards as strict or stricter than the Baucus requirements, and with the addition of other benefits came other requirements.

** Denotes standardization. Three policies were specified in detail and these were the only ones that were allowed to be sold. The least comprehensive policy satisfied the Baucus criteria.

Sources: The above table was compiled based on information in Van Ellet (1979), Merritt and Potemken (1982), McCall, Rice and Hall (1983), U.S. General Accounting Office (1986), and conversations with state regulators in Massachusetts and Wisconsin.

[46] This sample is described in detail in section four.

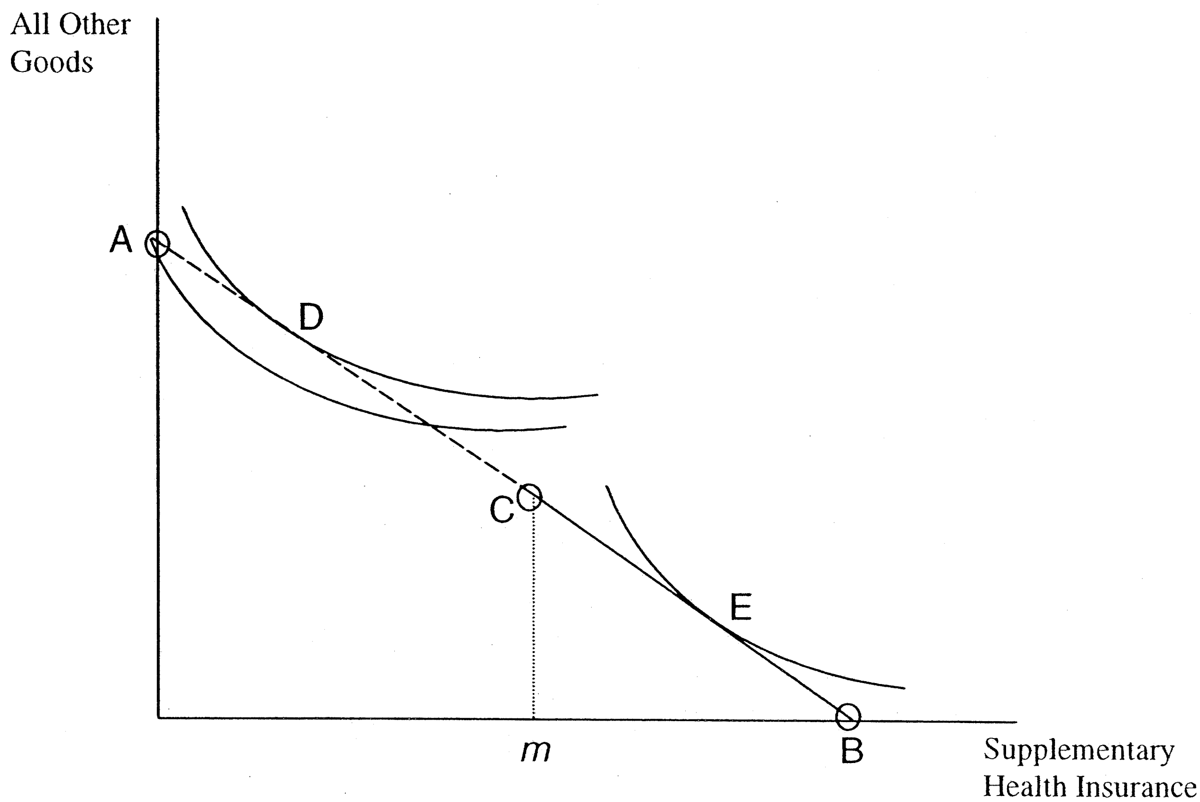

Figure 1: The effect of minimum standards on the decision to purchase insurance

Table 2: Estimated Effect of Minimum Standards on Non-Group Coverage Rate

	Without Covariates	With Covariates
$ADOPT^{-2}$	0.039 (0.022)	0.037 (0.022)
$ADOPT^{-1}$	0.011 (0.015)	0.012 (0.016)
$ADOPT^{1}$	**-0.059**** **(0.013)**	**-0.059**** **(0.014)**
$ADOPT^{2}$	**-0.087**** **(0.018)**	**-0.090**** **(0.018)**
$ADOPT^{3}$	**-0.135**** **(0.027)**	**-0.134**** **(0.027)**
Age		0.016** (0.003)
Male		-0.066** (0.008)
Married		-0.014 (0.010)
High School Degree		0.025** (0.010)
Some College		0.012 (0.014)
College Degree of Higher		-0.026 (0.015)
White		0.150** (0.013)
Health Status Excellent, Very Good or Good		0.025** (0.008)
Constant	0.395** (0.052)	-0.781** (0.212)
R^2	0.031	0.049
N	17,649	17,317

Notes: Coefficients are from OLS estimates of equation (1) using the 1976 through 1986 NHIS data. All regressions include year and state fixed effects. The dependent variable is whether an individual has coverage in the non-group market. Sample is limited to those aged 65 to 68. The reference category for the ADOPT variables is $ADOPT^{0}$. Standard errors are in parentheses. They are adjusted for the heteroscedasticity in the linear probability model and also for state*year correlation in the error term. ** denotes significance at the 1% level. * Denotes significance at the 5% level. The omitted education category is less than high school diploma. The omitted health category is "fair or poor."

Figure 2: Effect of Minimum Standards on Non-Group Coverage (65-68 Year Olds): Specification Checks

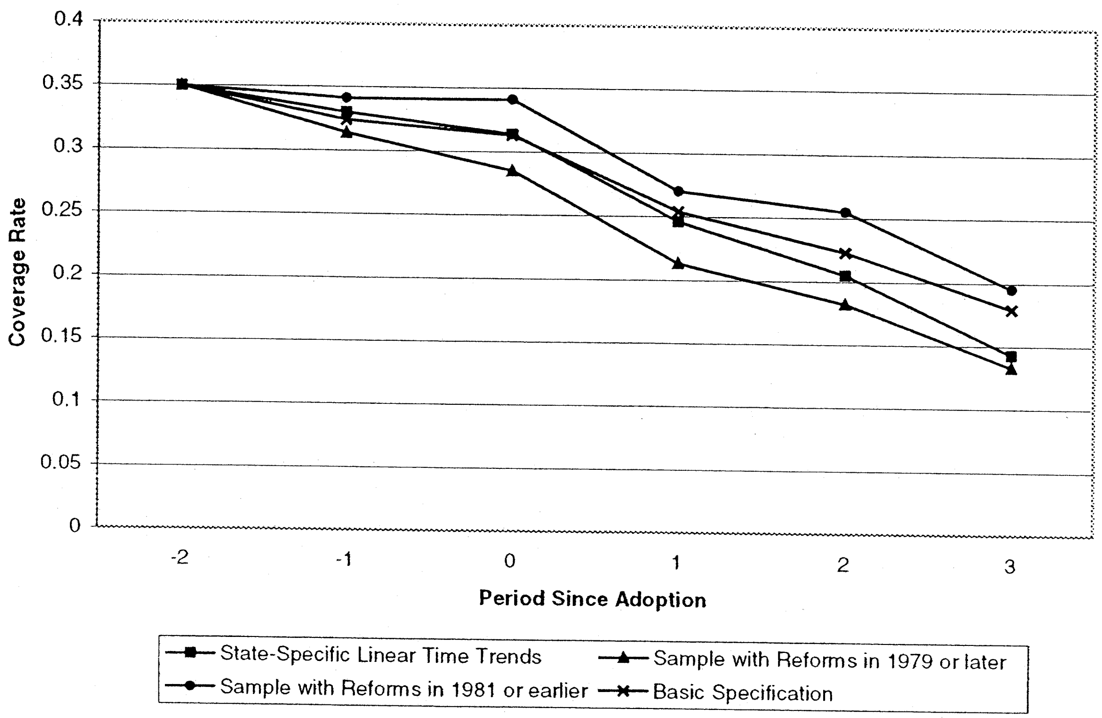

Table 3: Effect of Minimum Standards on Coverage Rates in Non-Group Market: Various Specifications

SPECIFICATION	**Immediate Effect of Reform**: ADOPT^1 versus ADOPT^0 (1)	**Long Term Vs. Immediate Effect**: ADOPT^1 versus ADOPT^3 (2)	**Falsification Exercise**: ADOPT^0 versus ADOPT^{-1} (3)
Basic Specification (BS) (N =17,317)	0.0000	0.0006	0.465
BS Plus State-Specific Linear Time Trends (N=17,317)	0.0001	0.0004	0.383
BS: Sample with Reforms in effect in 1979 or later (N =14,329)	0.0011	0.037	0.232
BS: Sample with Reforms in effect by 1981 (N =6,950)	0.0013	0.0018	0.999

Table reports p values from various F-tests of the λ's from OLS estimates of equation (1) using the 1976 through 1986 NHIS data. The dependent variable is whether an individual has coverage in the non-group market. Sample is limited to those aged 65 to 68. The first row reports results from estimating equation (1). Subsequent rows are based on various modifications of equation (1), as indicated.

Table 4: Estimated Effect Of Minimum Standards on Coverage of Cost Sharing Provisions in Medicare

	1977 (1)	1987 (2)	Difference-In-Differences (3)	Difference-In-Differences With Covariates (4)
Part A Deductible Coverage				
Treatment Group: Those Aged 65-71 in the Non-Group Market	0.99	0.91	----------	----------
Control Group I: Those Aged 65-71 in the Group Market	0.95	0.87	0.0007 (0.034) [N=1,042]	0.004 (0.036) [N = 940]
Control Group II: Those Aged 78+ in the Non-Group Market	0.97	0.91	-0.026 (0.026) [N=1,323]	-0.031 (0.028) [N=1,170]
Part B Deductible Coverage				
Treatment Group: Those Aged 65-71 in the Non-Group Market	0.85	0.39	-----------	--------
Control Group I: Those Aged 65-71 in the Group Market	0.91	0.72	-0.269** (0.054) [N=1,045]	-0.237** (0.057) [N=943]
Control Group II: Those Aged 78+ in the Non-Group Market	0.76	0.46	-0.154** (0.052) [N=1,331]	-0.133* (0.055) [N=1,177]

Note: Data are from the 1977 NMCES and 1987 NMES. Columns (1) and (2) report the weighted means of the dependent variable. Columns (3) and (4) report the estimated coefficient on Y1987*MINSTAND from estimation of equation (2) by OLS. Column (3) reports the estimated coefficient without the X matrix included in the regression. Column (4) reports the estimated coefficient with the X matrix included in the regression. The different rows report the result for different definitions of the control group. Heteroscedasticity-adjusted standard errors are in parentheses. ** indicates significance at the 1% level; * indicates significance at the 5% level.

Table 5: Estimated Effect Of Minimum Standards on Coverage of Services Not Covered By Medicare

	1977 (1)	1987 (2)	Difference-In-Differences (3)	Difference-In-Differences With Covariates (4)
Outpatient Prescription Drugs				
Treatment Group: Those Aged 65-71 in the Non-Group Market	0.27	0.21	----------	----------
Control Group I: Those Aged 65-71 in the Group Market	0.79	0.91	-0.177** (0.051) [N=1,130]	-0.198** (0.054) [N=1,015]
Care In Skilled Nursing Home				
Treatment Group: Those Aged 65-71 in the Non-Group Market	0.62	0.75	-----------	--------
Control Group I: Those Aged 65-71 in the Group Market	0.46	0.68	-0.091 (0.061) [N=1,165]	-0.110 (0.064) [N=1,047]
Inpatient Psychiatric Care				
Treatment Group: Those Aged 65-71 in the Non-Group Market	0.36	0.05	-----------	--------
Control Group I: Those Aged 65-71 in the Group Market	0.67	0.75	-0.389** (0.055) [N=1,165]	-0.389** (0.058) [N=1,047]
Home Health Care				
Treatment Group: Those Aged 65-71 in the Non-Group Market	0.19	0.05	-----------	--------
Control Group I: Those Aged 65-71 in the Group Market	0.28	0.53	-0.396** (0.054) [N=1,165]	-0.404** (0.056) [N=1,047]

Note: See notes from Table 4.

Figure 3: Market Share of Non-Group Plans

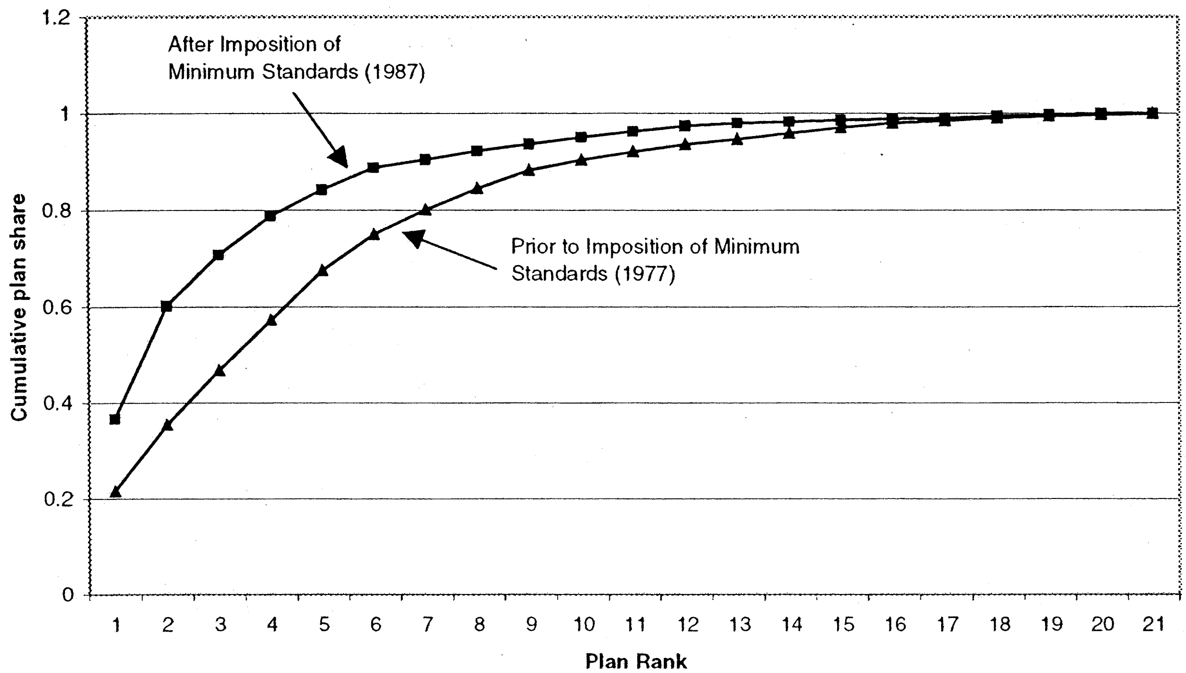

Table 6: Welfare Change Associated with Insurance Plan Changes (1999 $)

Coefficient of Relative Risk Aversion	Losing Pre-Baucus non-BCBS Coverage	Upgrading pre-Baucus plans to comply with Baucus requirements	Upgrading pre-Baucus plans to comply with Baucus requirements but dropping all non-mandated benefits
	(1)	(2)	(3)
1	-17.6	+4.1	-28.6
3	-475.8	+195.5	-733.7

Note: Table entries indicate the negative of the change in risk premium associated with a moving from the pre-reform insurance status quo to a new insurance arrangement. Changes in risk premiums are reported in 1999 dollars.

Figure 4: Distribution of Out of Pocket Medical Expenditures Under Different Insurance Arrangements

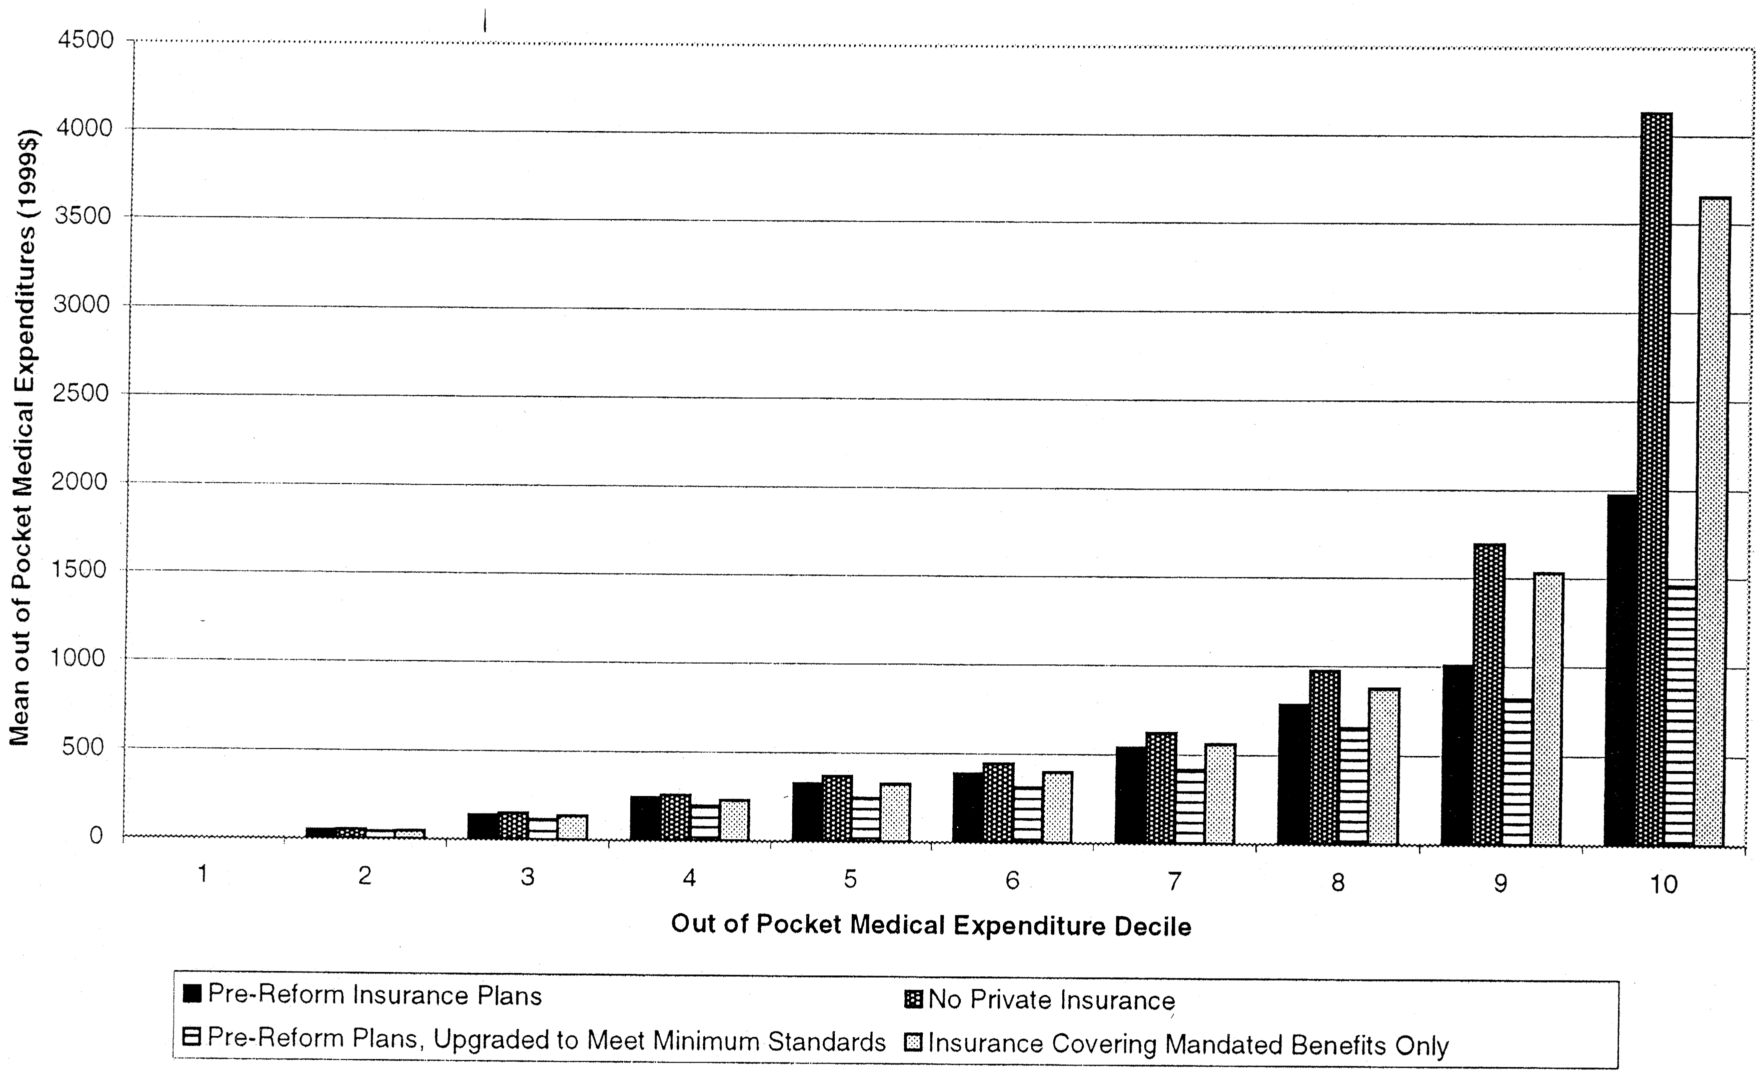

Appendix A: Gaps in Medicare (1977-1987)

Bolded benefits are those that the minimum standards required to be provided in any non-group Medigap policy

I. Cost Sharing Provisions in Medicare

A. Part A

1. Annual Deductible. ($124 in 1977; $520 in 1987)
2. **Copayment for hospital days 61-90 ($31 per day in 1977; $130 per day in 1987)**
3. **Copayment for lifetime reserve hospital days 91-150 ($62 per day in 1977; $260 per day in 1987)**[47]
4. **Annual Deductible for first three pints of blood used in the hospital**

B. Part B

1. Annual Deductible ($60 in 1977; $75 in 1987)
2. **20% Copay for approved physician charges**[48]
3. **Annual Deductible for first three pints of blood used.**

II. Services Covered Only Partially / With Restrictions By Medicare

1. Care in a Skilled Nursing Facility
2. Home Health Care Visits
3. Part A coverage for inpatient psychiatric care

III. Services Not Covered By Medicare

1. Outpatient prescription drugs
2. Dental care
3. Vision Care
4. Hearing Care
5. Preventive care: routine physical examinations, diagnostic tests and some immunizations
6. Care in custodial (not skilled) nursing homes
7. **Hospital stays beyond the lifetime reserve of 150 days**[49]
8. Physician charges above the "reasonable" rate reimbursed by Medicare Part B.[50]

---

[47] Beyond 90 days in the hospital, Medicare Part A provides a "lifetime reserve" of an additional 60 days that will be covered (with copayments) only once in a person's lifetime

[48] The regulation requires that the Medigap policy cover the 20 percent Part B copay for approved physician charges, subject to a maximum deductible of $200 and a maximum benefit of no less than $5,000.

[49] The regulation requires that the Medigap policy pay 90% of coverage of stays above the lifetime reserve maximum for a lifetime maximum of 365 additional days.

[50] The "reasonable charge" is defined as the lowest of the doctor's charge, the customary charge, or the prevailing charge in the area.

Appendix B: Probit Estimates of the Effect of the Minimum Standards on Non-Group Coverage Rates

	Without Covariates	With Covariates
$ADOPT^{-2}$	0.042	0.043
	(0.023)	(0.023)
$ADOPT^{-1}$	0.011	0.012
	(0.016)	(0.016)
$ADOPT^{1}$	**-0.059****	**-0.060****
	(0.013)	**(0.013)**
$ADOPT^{2}$	**-0.085****	**-0.089****
	(0.016)	**(0.017)**
$ADOPT^{3}$	**-0.131****	**-0.131****
	(0.024)	**(0.024)**

Notes: Reported coefficients are the estimated marginal effects from probit estimates of equation (1), evaluated at the means of the independent variables. See notes to Table 2 for more details on the data and specification. ** denotes significance at the 1% level. * Denotes significance at the 5% level.

Appendix C: Probit Estimates of Effect of the Minimum Standards on Non-Mandated Benefit Coverage

	Difference-In-Differences	Difference-In-Differences With Covariates
	Cost Sharing Provisions in Medicare	
Part A Deductible Coverage (Control Group I)	-0.041 (0.035)	-0.044 (0.036)
Part A Deductible Coverage (Control Group II)	-0.026 (0.026)	-0.031 (0.028)
Part B Deductible Coverage (Control Group I)	-0.192** (0.075)	-0.163* (0.079)
Part B Deductible Coverage (Control Group II)	-0.154** (0.052)	-0.133* (0.055)
	Coverage of Services Not Covered by Medicare	
Outpatient Prescription Drug Coverage (Control Group I)	-0.275** (0.071)	-0.302** (0.077)
Care in Skilled Nursing Home (Control Group I)	-0.075 (0.064)	-0.093 (0.067)
Inpatient Psychiatric Care (Control Group I)	-0.448** (0.040)	-0.453** (0.040)
Home Health Care (Control Group I)	-0.297** (0.029)	-0.293** (0.029)

Note: Reported coefficients are the estimated marginal effects of Y1987*MINSTAND from probit estimates of equation (2), evaluated at the means of the independent variables. Control Group I consists of individuals aged 65-71 in the group market. Control Group II consists of individuals aged 78+ in the group market. See notes to Table 4 for more details on the data and specification. ** denotes significance at the 1% level. * Denotes significance at the 5% level.

Chapter 2:
The Effect of Tax Subsidies to Employer-Provided Supplementary Health Insurance: Evidence from Canada

1. Introduction

In both the United States and Canada, private health insurance is primarily obtained as an employee benefit. In the United States, almost 90% of the non-elderly with private health insurance are covered through their employer (Employee Benefits Research Institute, 1995). In Canada, almost all private health insurance – which primarily covers out of hospital prescription drugs since these are not covered by the public health insurance system – is provided through an employer.[1] There may be efficiency reasons for the prevalence of such insurance through the workplace. Economies of scale in the administration and underwriting of policies make it cheaper for firms (particularly large ones) to provide benefits. In addition, by pooling workers of different health risks in a workplace-based insurance pool, employer provision of health insurance can reduce the scope for adverse selection that is present in the market for individual health insurance.

The predominance of employer provision may also be a function of the tax system. Both Canada and the United States subsidize employer provision of health benefits by excluding employer contributions to these benefits from the employee's taxable income. These tax subsidies constitute major expenditures for their respective governments. In the United States, the tax exclusion of employer-provided health insurance is the largest single tax expenditure, costing the federal government $US 72.5 billion in foregone federal income tax revenue in fiscal year 1999 (Office of Management and Budget 1999). In Canada, the exclusion cost the federal government approximately $CA 1.6 billion in lost federal tax revenues in 1998. (Government of Canada 1998).[2]

---

[1] There do not appear to be any comprehensive statistics on the individual health insurance market in Canada. However, calculations by the author, based on the annual surveys of the Canadian Health and Life Insurance Association, suggest that less than 3 percent of individuals with private health insurance in Canada are covered by the non-group market. These data (and their limitations) are discussed in more detail in section 4.4.

[2] Both the U.S. and Canadian figures include lost revenue from the tax subsidy provided to medical expenditures above a certain fraction of income. The total loss in revenue from the tax subsidy to employer-provided health insurance is considerably higher once foregone revenues from state (or provincial) and payroll taxes are considered.

There is little empirical evidence that would provide an estimate of the expected effects of a reduction in the tax subsidy to employer-provided health insurance on the extent of coverage through the workplace. Most of the existing literature is based on comparing health insurance coverage across workers with different marginal tax rates and hence different tax subsidy rates. As is discussed in greater detail below, such analyses are unlikely to estimate consistently the effect of the tax subsidy on coverage by employer-provided health insurance. However, a recent major reform to the tax subsidy in Canada provides an opportunity to consistently estimate the effect of the tax subsidy on coverage by employer-provided supplementary health insurance. In May 1993, the Quebec government removed the exclusion of employer contributions to health and dental benefits from an employee's provincial taxable income. The other provinces and the federal government kept the tax exclusion in place. The reform cut the total tax subsidy to employer-provided supplementary health insurance in Quebec by almost 60 percent. This change in the subsidy is substantially larger than changes studied in previous work. It therefore provides a unique opportunity to consider the effects of large-scale changes in the tax subsidy to employer provided supplementary health insurance.

Although universal public health insurance in Canada makes private health insurance coverage less critical in Canada than in the United States, substantial gaps in the Canadian public system have resulted in widespread use of supplementary private health insurance. In particular, the Canadian public health insurance system, like the U.S. Medicare program for those over 65, does not cover out of hospital prescription drugs. These are a rapidly rising component of health costs in Canada, amounting to 15 percent of total health expenditures in 1993 (World Health Organization 1996). As a result of this and other gaps in the public health insurance system, about 80 percent of the non-elderly, non-indigent adult population in Canada has supplementary health coverage through a private plan (Mercer 1995).

The effects of tax subsidies to employer-provided supplementary health insurance in Canada have been analyzed previously by Stabile (1999). Stabile explores the effects of marginal tax rates on the propensity to hold employer-provided supplementary health insurance by using variation in marginal tax rates across individuals at a point in time. The current paper builds on this work by exploiting a richer data set that allows us to look across time over the period of the Quebec reform as well as across individuals in estimating the effects of a change in the tax subsidy to employer-provided supplementary health insurance.

The results of this paper indicate that the tax subsidy to employer-provided supplementary health insurance has a large effect both on employer-provided supplementary health insurance coverage and on total insurance coverage. The reduction of the tax subsidy in Quebec coincided with a 13 to 14 percentage point drop in workplace coverage. This drop constitutes an 18 to 19 percent decrease in workplace coverage and corresponds to an elasticity of employer coverage with respect to the tax price of –0.46 to –0.49. This estimate is robust to alternative specifications. Changes in coverage reflect the joint outcome of changes in employer offering and changes in employee take-up. For reasons discussed below, I believe that the observed response is more likely to be occurring on the offering margin than the take-up margin, although I cannot distinguish empirically between these two effects.

Evidence from the individual (non-group) market suggests that non-group supplementary health insurance coverage rose slightly in Quebec relative to other provinces in response to the reduction in the tax subsidy to employer-provided (group) coverage. However, the increase in coverage in the non-group market offset only 10 to 15 percent of the decrease in coverage through an employer. The reduction in the tax subsidy to employer-provided supplementary health insurance is therefore associated with a substantial net decline in total private supplementary health insurance coverage in Quebec.

The evidence presented here also indicates a substantial dispersion across firms of different sizes in the sensitivity of workplace coverage to the tax subsidy. The reform is associated with a 19 to 26 percentage point reduction in coverage in firms with less than 20 employees, compared to only a 6 to 7 percentage point reduction in firms with more than 500 employees. These results suggest that the tax subsidy is critical to employer provision of supplementary health insurance in small firms, where gains from pooling and reduced administrative costs are likely to be small if present at all. The tax subsidy appears less critical in larger firms where these other factors may play a larger role in the creation of workplace-based insurance pools.

The rest of the paper proceeds as follows. In Section 2, I provide background on the Canadian public and private health insurance systems and on the tax treatment of employer-provided supplementary health insurance. Section 3 describes the data and estimation strategy used in this paper.

The results of the tax reform are presented and discussed in section 4. I first present estimates of the effect of the Quebec reform on coverage by employer-provided supplementary health insurance. I compare this estimate to previous estimates of the effect of the tax subsidy. I then extend the analysis to consider the effect of the Quebec reform on different types of workers and on coverage by employer-provided dental benefits. I also estimate the effects of a smaller tax reform that occurred at the same time in Ontario. Finally, I analyze the effects of the Quebec reform on the non-group market for supplementary health insurance in Quebec.

Section 5 examines the relative role of the tax subsidy – compared to other advantages of the workplace as a source of health insurance provision – in making the workplace the predominant source of such insurance. By comparing the effect of the Quebec reform on workers in firms of different sizes, I conclude that the tax subsidy plays a much larger role in the decision to provide insurance in small firms than in larger ones.

Section 6 presents several tests of the underlying assumption in the paper that the observed changes in supplementary health insurance coverage in Quebec are due to the Quebec tax reform rather than to other factors. Section 7 concludes.

2. Background–

2.1. The Canadian Health Care System

Canada has a predominantly publicly financed health insurance system that provides universal coverage with no user costs for a wide range of physician and hospital services. Yet significant gaps exist in the national insurance program.[3] As a result, private health insurance that supplements the public insurance system is widespread in Canada. The single largest expense not covered by the public system is expenditures on outpatient prescriptions drugs. Other services not covered by the national insurance include semi-private or private hospital accommodation, eyeglass and hearing aid plans, certain medical equipment such as prostheses and wheelchairs, rehabilitation, private nursing care, cosmetic procedures, out-of-country medical and hospital coverage, and dental benefits. Indeed, in 1993, only 73% of Canadian health expenditures were publicly funded. Private spending on health insurance totaled $CA 19 billion; of this, spending on private insurance amounted to $CA 6.2 billion, or 8.7% of total health expenditures. In addition, some portion of the $CA 7.6 billion dollars of out-of-pocket spending was due to cost sharing in supplementary plans. (World Health Organization, 1996).

As discussed previously, private health insurance in Canada is provided predominantly through the workplace. Survey evidence suggests that participation in employer health plans is not usually optional; over half of employer plans require employee participation, sometimes allowing for opt outs only if the individual is already covered as a dependent under a spouse's plan. Often employees will have a choice between individual or family coverage, but there is no indication of other elements of choice. The evidence also suggests that only about a third of plans require employee contributions to health premia, with average contributions of about 40 percent of premia (Wyatt 1993). Evidence from the Canadian Health and Life Insurance Association – which reports on total premia and number of policyholders – suggests that premia for group supplementary health insurance in 1993 were around $CA 300 to $CA 350.

2.2. The Tax Subsidy to Employer-Provided Supplementary Health Insurance in Canada.

The corporate income tax treats wages and employer contributions to health insurance plans symmetrically: Both are deductible from the employer's corporate income tax base. However, these two forms of compensation are treated differentially by the personal income tax. Unlike wages, employer contributions to health and dental plans in Canada are excluded from an individual's federal income and payroll tax bases. In addition, all provinces except Quebec follow the federal system and exclude employer-provided health and dental benefits from an individual's provincial taxable income. Quebec exempted employer contributions from the Quebec personal income tax until May 1993, when, in an

[3] Some provinces choose to cover some of these gaps through provincial public insurance programs.

effort to raise revenue, the Quebec government eliminated the tax exempt status of employer contributions to private health and dental insurance plans.[4]

The exclusion of employer contributions to group health insurance from the personal income tax base provides a subsidy to such insurance. While compensation paid in the form of wages and salary is subject to personal income and payroll taxation, compensation paid in the form of health insurance is not. If the tax code treated employer contributions to group health insurance and employer contributions to wages symmetrically, the price of health insurance in terms of foregone, after-tax consumption would be one. Since, however, employer contributions to group health insurance are exempt from taxable income, this price is less than one.

Following the existing literature, I define the "tax price" of employer-provided health insurance premia as the cost to the employee of a dollar of health insurance premia in terms of foregone, after-tax consumption. To represent this tax price, let τ_{fed} denote the federal marginal income tax rate on the marginal dollar of earned income, τ_{prov} denote the provincial marginal income tax rate, $\tau_{payroll,\ worker}$ denote the payroll tax levied on the employee for a marginal dollar of earned income, and $\tau_{payroll,\ firm}$ denote the payroll tax levied on the employer for a marginal dollar of earned income. In a competitive labor market, the employer will be indifferent between contributing a dollar to health insurance premia or a dollar to wages. Given the tax-exempt status of health insurance, if the employer contributes a dollar to health insurance premia, the employee receives a dollar in health insurance premia. If the employer contributes a

---

[4] The Quebec government estimated that, in 1992, the tax exclusion of employer contributions to private health insurance premia cost the Quebec government \$CA 149 million in foregone revenue (Government of Quebec, 1996). The Quebec reform was part of a general pattern at both the provincial and federal level to expand the tax base by levying new taxes on employer-provided insurance plans. At the same time that Quebec removed the tax exemption for employer contributions to private health insurance plans, it also removed the previously existing tax exemption to employer contributions for the first \$CA 25,000 of life insurance. The federal government followed suit in 1994 when it eliminated the federal tax exemption for the first \$CA 25,000 of employer contributions to group life insurance. It has debated, but failed to enact, a reform in the federal budget that would have followed Quebec's 1993 health insurance tax reform at the federal level. Nor were the 1993 Quebec reforms the first time the Quebec government had expanded its tax base by levying new taxes on employee insurance plans. In 1985, it levied a 9% retail sales tax on group insurance premia and in 1990 it extended the sales tax to cover self-insured employee benefit plans. Ontario followed this example in 1993 when it levied an 8% retail sales tax on insured and self-insured group benefit plans and extended its 2% premium tax, previously levied only on insured group benefit plans, to self-insured plans. (Koskie 1995; Nielson 1998).

dollar to wages, however, the employee receives only $\frac{1}{(1+\tau_{payroll,firm})}$, since the employer must pay payroll taxes on any wages paid to the employee, but not on compensation paid in the form of health insurance.[5] The employee must, in turn, pay federal and provincial income taxes and payroll taxes on any wages received, but not on any health insurance premia paid by the employer. Finally, group health insurance faces certain consumption taxes (specifically, premium and sales taxes) that are not applied to other services that the employee could buy.[6] As a result of these tax rules, the tax price of a dollar of health insurance premia, or the cost of health insurance premia to the employee in terms of foregone, after-tax consumption, is given by the following:[7]

$$\text{TAXPRICE} = \{ (1 + \tau_{\text{cons}}) * (1 - \tau_{\text{fed}} - \tau_{\text{prov}} - \tau_{payroll,\ worker}) \} / \{ 1 + \tau_{payroll,\ firm} \} \quad (1)$$

Since τ_{cons} is empirically small, this tax price is less than 1. The resulting tax subsidy to employer-provided health insurance therefore reduces the price of this insurance for the employee compared to the price of other consumption. I make use of a change in the tax subsidy to study the relationship between the tax price of employer-provided health insurance and the demand for it.

2.3. The 1993 Quebec Reform

The 1993 Quebec tax reform made employer contributions to health insurance taxable to the employee under the Quebec provincial income tax. For insured plans, the value of the taxable benefit is assessed as the premium paid by the employer. For self-insured plans, the value of the taxable benefit is the employee's pro-rated share of the plan's benefit payments and administrative expenses incurred with third-party operators (Koskie, 1995).

The tax price of employer-provided supplementary health insurance in Quebec after the reform is therefore given by the following:

---

[5] I assume that labor supply is inelastic and therefore the full incidence of income and payroll taxes are borne by the worker.

[6] I assume that the incidence of these consumption taxes is on the consumer.

[7] If we were instead to consider the tax price of a dollar of health insurance benefits (rather than a dollar of health insurance premia), equation (1) would have to be adjusted to take into account the administrative load (or excess of

$$\text{TAXPRICE} = \{ (1 + \tau_{cons}) * (1 - \tau_{fed} - \tau_{payroll,\ worker}) \} / \{ 1 + \tau_{payroll,\ firm} \} \quad (2)$$

The reform increased the tax price of employer-provided supplementary health insurance in Quebec by the following:

$\{ (1 + \tau_{cons}) * (\tau_{Quebec}) \} / \{1 + \tau_{payroll,\ firm}\}$. Basic provincial marginal tax rates in Quebec range from 16 to 24 percent.

A tax subsidy to employer-provided supplementary health insurance still remains in Quebec after the reform because of the exclusion of this benefit from federal income and payroll taxes. Absent the Quebec reform, the average employee in Quebec would have faced a choice between a dollar in supplementary health insurance premia or 57 cents of after tax consumption; because of the reform, the choice was between a dollar in supplementary health insurance or 82 cents of after tax consumption. As a result, the reform increased the tax price of supplementary health insurance on average by 25 cents. The tax subsidy fell from 43 cents to 18 cents, or by 58%. The components of the tax subsidy to employer-provided supplementary health insurance in Quebec before and after the tax reform are given in Table 1. For comparison purposes, Table 1 also reports the components of the tax subsidy in the control provinces that are used in the estimation strategy.[8]

3. Data and Estimation Strategy

3.1. Data Source and Sample

The analysis in this paper is based primarily on a repeated cross-section formed from the 1991 and 1994 Canadian General Social Surveys (GSS). The GSS is an annual, stratified random telephone survey of the non-institutional population over age 15 in the 10 Canadian provinces. Since the Quebec reform was effective starting in May 1993, I use the 1991 survey for the "before" period and the 1994

premia over claims as a percentage of claims) on group health insurance. This would simply represent a scaling of equation (1) and would not affect the substantive analysis.

[8] Table 1 indicates that the average tax price in Quebec in 1991 was 60 cents. The difference between this price and the 57 cents that the tax price would have been in 1994 absent the reform is due simply to slight compositional differences between the 1991 and 1994 sample that are controlled for in the analysis.

survey for the "after" period.[9] I restrict the sample to individuals aged 25-64 who have paid employment, and exclude the self-employed.[10]

3.2. The dependent variable

The dependent variable is a binary variable indicating whether the individual receives supplementary health insurance from his or her employer.[11] The variable therefore combines information on employer offering with information on employee take-up, conditional on offering. Changes in this variable, however, are most likely due to changes in employer offering rather than employee take-up. As discussed above, many employees do not have a choice of take-up if the employer offers health benefits, and employers pay the bulk of their employees' health insurance premia. Both of these features suggest that it is unlikely that changes in coverage reflect changes in take-up.

The decision to offer insurance as a function of the tax subsidy is, at one level, a firm decision. In equilibrium, however, the firm behavior reflects a joint decision, or bargain, between the worker and the firm. Observed changes in firm offering, therefore, likely reflect underlying changes in demand by some group of workers in the firm.

The surveys distinguish between two types of insurance: medical and surgical benefits and dental benefits.[12] As discussed above, the predominant risk insured by medical and surgical benefit plans is expenditure on out of hospital prescription drugs. Other services not covered by the public insurance system (such as semi-private or private hospital accommodation, private nursing care, or chiropractors and other professionals) are often covered as well by these private plans (Wyatt 1993). Private dental

---

[9] The 1991 GSS is the latest before the 1993 reform to ask about employer-provided benefits. The 1994 GSS is the only GSS after the 1993 reform to ask about employer-provided benefits. Both surveys were conducted evenly throughout the 12 months of the year.

[10] Both surveys come with weights that are designed to sum to the population of Canada. So that different years are not given different weight in the analysis (due either to population growth or to the selected subsample being a different fraction of the total sample in the different years), I rescale the weights so that, in each year, in the selected subsample, the weights sum to 1.

[11] The survey questions read: "Does your employer provide you with medical/surgical benefits beyond those provided by your provincial health care system" and "Does your employer provide you with dental care benefits." The interviewer instructions clarify that "by 'provide' we mean that the employer subsidizes or pays for all or part of the items listed."

[12] Dental benefits tend to be insured separately in Canada from other potential medical expenditures.

benefits include basic treatment, major restorative treatment, and orthodontics.[13] I study separately the effect of the reform on medical and surgical benefits (hereafter, " supplementary health insurance") and dental benefits.

3.3. Identification Strategy: Difference-in-Differences

The approach taken in this paper is to compare the change in coverage by employer-provided supplementary health insurance between 1991 and 1994 in Quebec to the change in coverage over the same period in a group of 7 control provinces that are not affected by the Quebec tax reform.[14] Ontario is excluded from the control provinces because in 1993 – at the same time that Quebec removed its provincial income tax subsidy to employer-provided supplementary health insurance – Ontario imposed an 8% provincial retail sales tax on employer-provided supplementary health insurance and extended its existing 2% premium tax for insured employer health plans to self-insured employer plans. I also exclude Saskatchewan because in 1993 it undertook a reform in the generosity of its provincial public drug insurance program which provides some provincial prescription drug insurance to all residents of the province.

The basic estimating equation is:

$$\text{INSURANCE} = \beta_0 + \beta_1(\text{QUEBEC*AFTER}) + \beta_2\text{QUEBEC} + \beta_3\text{AFTER} + X\beta_4 + \varepsilon \qquad (3)$$

The dependent variable INSURANCE is a binary measure of whether the individual has employer-provided supplementary health insurance. QUEBEC is a dummy variable that equals 1 if the individual resides in Quebec, and 0 if he resides in the control provinces. AFTER is a dummy variable equal to 1 if the individual was surveyed in 1994, and 0 if he was surveyed in 1991. QUEBEC*AFTER is an interaction of these two dummies. X is a matrix of covariates.

---

[13] No information is available in the survey about other margins of coverage which may also be affected by the tax subsidy, such as co-pay rates, services covered, or employee contributions to premia. I therefore follow the bulk of the literature in considering just the binary coverage margin.

[14] The 7 control provinces are: Alberta, British Columbia, Manitoba, New Brunswick, Newfoundland, Nova Scotia, and Prince Edward Island. The population of Quebec is about three-quarters of the combined population of the 7 control provinces.

QUEBEC controls for any fixed, regional differences in coverage rates for employer-provided supplementary health insurance. AFTER controls for any nationwide trend in employer-provided supplementary health insurance coverage between 1991 and 1994. With these controls in place, β_1 – the parameter of interest – measures the change in the probability of having employer-provided supplementary health insurance in Quebec between 1991 and 1994 relative to the change in the probability of having employer-provided supplementary health insurance in the control provinces between 1991 and 1994. The identifying assumption is that there was no Quebec-specific time trend in coverage between 1991 and 1994 that would have caused trends in coverage to differ between the two regions if the 1993 Quebec tax reform had not been enacted.

Compositional changes in the treatment group relative to the control group in characteristics that are correlated with coverage by employer-provided supplementary health insurance could drive differences in trends in coverage between the treatment and control group over the sample period. For example, if union membership decreases in Quebec relative to the control provinces, and union members are more likely to be covered by employer-provided supplementary health insurance than non-union members, this change in union membership could drive differences in the trends in insurance coverage in the two regions. I therefore control flexibly for this possibility by including the X matrix of covariates in equation (3). These covariates consist of dummies for union membership, marital status, gender, spousal employment (conditional on marriage), age, educational attainment, number of children under 25, self-reported health status, full time versus part time work, full year versus part year work, personal income, and occupation.[15] I also include the provincial unemployment rate as a covariate to control for any differences in the macroeconomic cycle across provinces.

I would like to control for firm size – since there is a well-known positive correlation between firm size and coverage by employer-provided health insurance – but the 1991 data do not include information on firm size. Since data on firm size are available in the 1989 GSS, I estimated the effect of the reform using the 1989 and 1994 samples with and without controlling for firm size. The inclusion or exclusion of controls for firm size had no effect on the estimated effect of the reform using the 1989 and 1994 samples. This suggests that my results from the 1991 to 1994 comparison are unlikely to affected by the exclusion of firm size controls.

With the covariates included in the regression, the identifying assumption is that, conditional on X, there is no Quebec-specific time trend in coverage by employer-provided supplementary health insurance. I explore several ways of examining the validity of this identifying assumption below. I examine whether coverage by employer-provided supplementary health insurance in Quebec in years prior to the reform is following the same time trend as coverage by this insurance in the control provinces. I also examine whether coverage rates for other employee benefits that are not affected by the Quebec tax reform (such as pensions) are following the same trend in Quebec and in the control provinces between 1991 and 1994.

In the analysis below, I report results from three different specifications of equation (3). In the "difference-in-differences" specification, I do not control for any covariates. This regression therefore includes a Quebec fixed effect, a time fixed effect, and the key variable of interest, the interaction of the Quebec and time fixed effect: QUEBEC*AFTER.

I also report the results from two different specifications with covariates. There is a concern that some of the covariates included in X may be jointly determined with coverage by employer-provided supplementary health insurance. For example, there is a well-known correlation between income and benefits, with "good jobs" offering higher pay and better benefits. If individuals have characteristics unobservable to the econometrician that determine whether they are in good jobs or not, then income and benefits are jointly determined. Similarly, individuals who desire supplementary health insurance may choose to work in occupations that are more likely to offer such insurance, or to choose to work full time rather than part time in order to be eligible for benefits.

It is unclear how to address this endogeneity problem given the available data. Lacking instruments for all of the potentially endogenous variables, I try estimating two different specifications of

---

[15] I have not controlled for industry because of differences in industry categories in the 1991 and 1994 surveys.

equation (3). In the "limited covariates" specification, I include only the most plausibly exogenous covariates: union membership, age, marital status, gender, spousal employment, number of children under 25, educational attainment, and provincial unemployment rates. In the "full covariates" specification, I add to these covariates the potentially endogenous covariates: occupation, income, full time versus part time, full year versus part year, and health status.

I estimate equation (3) using ordinary least squares. Non-linear estimation techniques such as probit yield similar results. In all regressions, standard errors are adjusted for heteroscedasticity and, when provincial unemployment is included as a covariate, for province*year correlation in the error term.

3.4. Comparison to other empirical approaches

Much of the previous literature on the relationship between tax subsidies to employer-provided health insurance and coverage by such insurance has been plagued by problems of identification.[16] Time series studies have compared coverage by health insurance across points in time in which the price of health insurance differs.[17] Such studies suffer from the concern that other things – such as the cost of health care – may have been changing over the time period studied and these things would have an independent effect on coverage rates. Cross-sectional studies have compared health insurance coverage at a point in time across individuals who face different tax subsidies for health insurance.[18] After controlling for income and family structure, the variation in the tax subsidy across individuals is driven predominantly by non-linearities in the treatment of income and family structure by the tax system. For the estimated effect of the tax subsidy to have a causal interpretation requires that the effects of income and family structure on health insurance demand have been fully captured by the controls; otherwise, their effect "loads on" to the estimated effect of the tax subsidy.

[16] For a more detailed review of the literature on the price responsiveness of employer-provided health insurance, see Gruber (1999).

[17] Examples of such studies include Long and Scott (1982) and Turner (1987).

[18] Examples of such studies include Phelps (1973), Goldstein and Pauly (1976), Taylor and Wilensky (1983), Woodberry (1983), Holmer (1984), Sloan and Adamache (1986), Leibowitz and Chernew (1992), and Gentry and Peress (1994). These studies all use U.S. data. Stabile (1999) performs a cross-sectional analysis on Canadian data.

The difference-in-differences approach employed here overcomes many of the identification problems of the time series and cross-sectional literature. By comparing changes in Quebec to changes in the control provinces, it controls for any time trend that is common to Quebec and the control provinces in the demand or supply of employer-provided supplementary health insurance. It also controls for any fixed regional differences across provinces in coverage by employer-provided supplementary health insurance.

This difference-in-differences approach has been used previously by Gruber and Poterba (1994) to study the effect of the tax subsidy on health insurance coverage. They study the effect of the creation of a tax subsidy in the United States for health insurance purchased by the self-employed on the probability that a self-employed individual has health insurance. Their estimate of the effect of the tax subsidy is based on a comparison of the change in health insurance coverage for the self-employed before and after the creation of the tax subsidy to the change in coverage for the employed, who were not affected by the new tax subsidy. The primary drawback to this study is that it studies only the self-employed. The self-employed tend to differ from the employed in identifiable ways and may therefore have a different elasticity of demand for health insurance.[19] Moreover, the self-employed are able to tailor their health insurance package to their personal preferences. In contrast, employer-provided health insurance has aspects of a local public good in that the employer cannot offer separate health insurance packages to each individual (Goldstein and Pauly 1976). As a result, the effect of a tax subsidy may be different in an employment context, in which some collective decision rule is used to decide on the benefits package(s) offered, than for an individual who can choose his own package.

4. Results

4.1. Basic Results: Changes in Coverage by Employer-Provided Health Insurance in Quebec

Table 2 presents sample means for characteristics of the treatment and control groups before and after the 1993 reform. There are some pre-treatment differences between Quebec and the control

[19] For example, they tend to be older, more educated and more likely to be male and non-black (Gruber and Poterba (1994)). They may also be less risk averse than those in paid employment.

provinces. In particular, the unionization rate and the unemployment rate are statistically significantly higher in Quebec, and patterns of educational attainment differ, with more people in Quebec having less than a high school degree and fewer having received some college education. The last column of Table 2 looks at whether differences between Quebec and the control provinces are stable over time. For the most part they are, which allays concerns that the two regions might have been moving on different underlying trends over the sample period. One exception is that there appears to be an increase in the number of people who report their health status compared to others in their age group as excellent in Quebec relative to the control provinces. There is also some evidence of an increase in the percentage of people who are married in Quebec relative to the control provinces. I can control for marital and health status in the analysis below, to control for any effect these changes might have on the demand for employer-provided supplementary health insurance.

Table 3 presents the basic difference-in-differences result. Between 1991 and 1994, coverage by employer-provided supplementary health insurance fell by 8.0 percentage points in Quebec. During the same period, coverage rose by 5.2 percentage points in the control provinces. The difference in these differences indicates a 13.2 percentage point decrease in coverage by employer-provided supplementary health insurance in Quebec relative to the control provinces between 1991 and 1994.

Table 4 presents the results in a regression context. The first column reports the same analysis as in Table 3. Columns 2 and 3 report the effect of adding various covariates to the analysis in the first column. The estimated magnitude of the effect of the reform is not sensitive to the addition of covariates: Depending on the specification, the Quebec reform is associated with a 13.1 to 13.6 percentage point drop in coverage by employer-provided supplementary health insurance in Quebec relative to the control provinces.[20] This effect is statistically significant at the 1% level in all specifications. Given the baseline of 72.3 percent coverage in Quebec in 1991, these percentage point drops represent an 18 to 19 percent decrease in coverage by employer-provided supplementary health insurance in Quebec.

With the richest set of controls, there is no significant time trend in coverage by employer-provided supplementary health insurance in the control provinces. Furthermore, while the probability of coverage by employer-provided supplementary health insurance in Quebec is about 4 percentage points higher than in the control provinces prior to the reform, this difference is significant only at the 10% level. Part of this difference may be due to variations across provinces in their provincial public insurance programs. For example, Manitoba and British Columbia have public drug programs that provide all residents with some insurance for prescription drugs, albeit with considerable co-payments and deductibles.[21] The lower tax price of supplementary health insurance in 1991 in Quebec compared to the control provinces (see Table 1), may also help account for the higher rate of pre-reform coverage in Quebec. As long as the non-tax factors contributing to the higher coverage rate in Quebec relative to the control provinces in 1991 were not changing between 1991 and 1994, they do not present a problem for this analysis.

To check the sensitivity of these results to the use of the linear probability model, I re-estimated the effect of the Quebec reform using the non-linear probit model. In results not reported here, I find that, with the values of the covariates set to their 1991 (pre-reform) Quebec means, the probit model estimates that the effect of the reform was to reduce the probability of coverage by employer-provided supplementary health insurance in Quebec by 13.4 to 17.6 percentage points. These estimates are of slightly larger magnitudes than the estimated 13.1 to 13.6 percentage point decrease from the linear probability model. As in the linear model, the estimated effects of the reform from the probit model are all statistically significant at the 1 percent level.

4.2. Estimated Elasticity and Comparison to Previous Estimates

As discussed above, many papers have examined the relationship between tax subsidies and coverage by employer health insurance. The ones that are most directly comparable to the estimate here

[20] In results not reported here, the estimated effect of the reform is also robust to including interactions of the covariates (X) with AFTER in the regression. This specification allows the effect of covariates (such as union membership) on coverage by employer-provided supplementary health insurance to vary over time.

are those that look at the relationship between the tax subsidy and the probability that the employer offers health insurance. Of course, my coverage measure potentially combines changes in offering with changes in take-up conditional on offering. However, as discussed above, I believe that the observed changes are most plausibly occurring on the offering margin.

Gentry and Peress (1994) use variation in U.S. state tax rates to identify the effect of the tax subsidy on the percentage of blue collar workers in a city that are offered health insurance. They estimate an elasticity of employer offering with respect to the tax rate of –1.8 for blue collar workers; they find little evidence of an effect for white collar workers. Royalty (1999) also uses cross-state variation in U.S. state income tax rates to estimate the effect of the tax subsidy on employer-provided health insurance. She controls for state fixed effects by comparing the effect of variation in state tax rates on employer offering of health insurance with the effect of this variation on employer offering of sick leave, which does not enjoy the tax advantages of health insurance. She estimates an elasticity of employer offering with respect to the marginal tax rate of –0.68. Finally, Gruber and Lettau (2000) use variation in the tax price of health insurance that stems from variation across U.S. states at a point in time as well as from changes in U.S. state and federal tax rates over time to estimate an elasticity of employer offering with respect to the tax price of –0.32.

To compare my estimate to these, I translate the percentage point drop in the probability of coverage by employer-provided supplementary health insurance associated with the Quebec reform into an elasticity of employer coverage of supplementary health insurance with respect to the tax price. There is some degree of arbitrariness in reporting elasticities of coverage with respect to the tax price, as opposed to elasticities with respect to the overall price, or load, to which the tax subsidy contributes. I use the tax price rather than the overall price since the former is the metric used by the literature and therefore provides a better way of comparing my results to those of the previous literature.

---

[21] The estimated effect of the Quebec reform is not sensitive – in either magnitude of significance – to the exclusion of these two provinces from the group of control provinces.

I use the Quebec reform – and the associated change in the tax price in Quebec – to estimate the magnitude of the decrease in coverage associated with a given increase in the tax price. I therefore regress coverage by employer-provided supplementary health insurance on the individual's tax price, using the reform (QUEBEC*AFTER) as an instrument for the tax price.

The basic estimating equation is now:[22]

$$INSURANCE = \beta_0 + \beta_1 TAXPRICE + \beta_2 PROVINCE + \beta_3 AFTER + X\beta_4 + \varepsilon \qquad (4)$$

TAXPRICE is then instrumented for using QUEBEC*AFTER. The variation in the individual's tax price used to identify the relationship between the tax price and coverage by workplace-based insurance therefore comes from the Quebec reform. This approach allows me to translate the estimated effect of the Quebec reform into a parameterized estimate of the relationship between the tax price of employer-provided supplementary health insurance and coverage by such insurance. This parameterized estimate can then be compared to estimates that have used other sources of variation to estimate the relationship between the tax price of insurance and coverage by such insurance.

To calculate the tax price, I impute federal, provincial, and payroll marginal tax rates for individuals in my data set.[23] The imputation takes into account federal and provincial taxes and surtaxes, various tax credits, and federal payroll taxes for unemployment insurance and for the Canadian Pension Plan / Quebec Pension Plan. It also takes account of the differential consumption tax treatment of supplementary health insurance compared to other services due to provincial sales taxes on group supplementary health insurance and provincial premia on insured group supplementary health insurance plans.[24]

Table 5 reports the results from estimating equation (4) by OLS and by IV. The OLS estimate of the effect of the tax price of employer-provided supplementary health insurance on coverage by such

---

[22] I have controlled for each province separately rather than simply controlling for whether the individual lives in Quebec or not. I do this to control for differences in TAXPRICE that are driven by differences in provincial tax systems. Controlling for each province separately does not alter the estimated effect of the reform in Table 4.
[23] Information on these rates comes from Canadian Tax Foundation (1991) and Canadian Tax Foundation (1994). I am grateful to Kevin Milligan for his help in navigating the intricacies of the Canadian tax system.

insurance is –0.322. This is almost half of the magnitude of the IV estimates which are –0.560 or –0.595 depending on the specification. In the only other estimate I know of of the effect of the tax price of employer-provided supplementary health insurance in Canada, Stabile (1999) estimates a coefficient on tax price of –0.417. This estimate is close to my OLS estimate. This is not surprising: Although the estimates use different data sets and different years (as well as slightly different specifications), Stabile's estimate (like my OLS estimate) is identified off of non-linearities in the tax schedule, after controlling for province of residence, income and other personal characteristics. My results suggest that such an identification strategy produces a considerable downward bias in the estimated effect of the tax price of employer-provided supplementary health insurance on coverage by such insurance.

The IV estimate of equation (4), which uses variation in the tax price from the Quebec reform as an instrument for TAXPRICE, is identified solely by changes in the tax price due to the 1993 Quebec reform. The IV results indicate that a ten cent increase in the tax price of employer-provided supplementary health insurance reduces the probability of coverage by this insurance by about 6 percentage points. The implied elasticities are reported in the second row of Table 5. These elasticities are calculated with reference to the pre-reform means in Quebec of the tax price (57 cents) and the probability of the employer providing supplementary health insurance (73 percent). The IV estimates imply an elasticity of coverage by employer-provided supplementary health insurance with respect to the tax price of –0.46 to –0.49. This estimate lies between the estimate of –0.32 reported by Gruber and Lettau (2000) ~~and~~ that of –0.68 reported by Royalty (1999). It is considerably lower than the –1.8 elasticity estimated for blue collar workers by Gentry and Peress (1994), although they find virtually no effect for white collar workers so the average elasticity is presumably lower than the reported one. In surveying the U.S. literature more broadly, Cutler (1996) concludes that estimates of the elasticity of demand for employer-provided health insurance tend to lie between 0 and –2.0, with the "consensus" on elasticities in the range of –0.5 to –1.0.

[24] A more detailed description of the way in which tax price is calculated is provided in the earlier working paper version of this paper.

The elasticity produced from the Quebec reform thus appears to lie well within the range of estimates reported in the literature. Yet there are several reasons we might have expected *ex ante* that the Quebec reform could have produced substantially different estimates. First, as discussed above, many of the previous papers suffer from problems of identification that could bias the estimated effects. Second, the magnitude of the Quebec reform, which reduced the subsidy to employer-provided health insurance by almost 60 percent – is considerably larger than the variation in tax price studied in previous work. For example, the cross-state variation in U.S. state income tax rates used by studies such as Royalty (1999), is only 10 percentage points. There could be non-linearities in the effect of the tax subsidy on workplace-based insurance due, for instance, to the fixed costs of setting up an employee benefit plan or to an adverse selection spiral that may result once healthy workers are no longer subsidized enough to make them willing to pool with less healthy workers.[25] The potential for such non-linearities makes estimates based on large variation in the tax price particularly useful for assessing the likely impact of a complete removal of the tax subsidy to employer-provided health insurance.

Third, almost all of the previous estimates are from the United States. The responsiveness of coverage by employer-provided health insurance in Canada may differ from that in the United States. Private health insurance in the United States is considerably more comprehensive than in Canada. We might expect the demand for supplementary coverage to be more price elastic than the demand for more comprehensive coverage. On the other hand, unionization rates are much higher in Canada than in the United States, and, as discussed below, I find a substantially larger response to the tax change among non-union members than union members. This might lead us to expect a lower elasticity in Canada than in the United States.

4.3. Additional Results for the Group Market

The preceding sections have presented the central estimate of the 13.1 to 13.6 percentage point decrease in employer provided supplementary health insurance in Quebec associated with the tax reform and the resulting elasticity estimate of coverage by employer-provided supplementary health insurance

[25] These issues are considered in more detail by Pauly (1986) and Gruber and Poterba (1996a).

with respect to the tax price of –0.46 to –0.49 In this section I consider three other additional effects of interest.

First, I examine whether certain types of workers were more likely than other types to lose employer-provided supplementary health insurance coverage in response to a change in the tax subsidy in Quebec. To do this, I estimate equation (3) separately for different groups of workers. Some of the results are reported in Table 6. I find that the effect of the Quebec tax reform on coverage by employer-provided supplementary health insurance varies substantially by gender. Indeed, while the reform is associated with a decline in coverage for men from 17.6 to 19.3 percentage points, the decline in coverage for women ranges only from 5.2 to 6.5 percentage points. The estimated effect of the reform is statistically significantly different for men and for women at at least the 5% level in all specifications. This result is puzzling as there is no a priori reason to expect coverage to be less sensitive to the tax price for women than for men. Nor is the estimated effect for men and women similar in percentage terms. The reform is associated with a 22 to 24 percent decline in coverage for men but only an 8 to 10 percent decline for women.

A more intuitive result is that the drop in coverage associated with the reform is substantially larger for non-unionized workers than for unionized workers. The estimated effect of the reform is a 7.6 to 9.7 percentage point drop in coverage for union members compared to a 14.9 to 17.9 percentage point drop for non-union members. The difference between the estimated effect of the reform for non-union workers compared to the estimated effect for union workers is statistically significant in all specifications. Finally, I also find evidence that the effect of the reform was largest for the oldest group of workers (aged 55 to 64). The reform is associated with a 21.4 to 29.6 percentage point drop in coverage by employer-provided supplementary health insurance for workers aged 55 to 64. In both specifications with covariates, this estimated effect for the oldest age group of workers is statistically significantly larger than the estimated effects for workers in any of the other three age groups (ages 25-34, 35-44, and 45-54). There are no statistically significant differences in the estimated effect of the reform among the other three age groups. The estimated effect of the reform for these other groups ranges from 10.0 percentage points to 14.7 percentage points.

I find no evidence of a differential effect of the reform for workers with different potential experience (measured by age minus education minus 6), or for workers with different labor force attachments (part-time versus full-time, or part-year versus full-year). I also examined whether I may be underestimating the effect of the reform by looking at its effect only in the year immediately following the reform. I estimate the effect of the reform separately for those individuals surveyed in the first half of 1994 and those surveyed in the second half of 1994. The estimated effect of the reform is an 18 percentage point drop for individuals in the second half of the year, compared to a 7 percentage point drop for individuals in the first half of the year; these differences are statistically significant. However, when I compare the effect in the first quarter of 1994 to that in the last quarter of 1994, the difference is smaller (10 percentage points compared to 14 percentage points) and is not statistically significant. These results suggest that I may be underestimating somewhat the effect of the reform.

Second, since the Quebec reform affected the tax treatment of dental benefits as well as health insurance, I replicate the basic analysis for dental benefits. I re-estimate equation 3 using coverage by employer-provided dental benefits as the dependent, binary variable. In results not reported here, I find that there is a statistically significant drop in dental benefits in Quebec compared to the control provinces. As with supplementary health insurance benefits, the magnitude and significance of the result is robust to the inclusion of other covariates and to the inclusion of time-varying covariates. The magnitude of the drop ranges between 6.9 and 7.9 percentage points depending on the specification, and is statistically significant at at least the 5% level in all specifications. However these estimates should be interpreted with some caution. In 1992, the Quebec government instituted minor cutbacks in its public dental benefits program. This public insurance program covers oral surgery for all residents and dental services for children. In 1992, the government reduced the maximum age of eligibility for children, from 15 to 12 years or from 13 to 10 years depending on the specific procedure. Although these reforms were considered relatively minor,[26] they nonetheless may have increased the demand for private dental insurance in Quebec during the period in which the tax reform increased the price of private dental insurance. The reported results are therefore probably an underestimate of the effect of the tax reform on coverage by employer-provided dental benefits.

And finally, I consider the effect of a smaller tax reform that occurred in Ontario at the same time as the Quebec reform. In 1993, the Ontario government imposed an 8 percent sales tax on group health and dental benefits and extended the 2 percent premium tax previously in place for insured group plans to uninsured group plans. This 8 to 10 percentage point increase in taxes on group health and dental benefits in Ontario is roughly two-fifths of the average 25 percentage point increase in the tax price from the Quebec reform. In results not reported here, I repeated the above analysis using Ontario as the treatment province instead of Quebec, and using the same 7 unaffected provinces as the control provinces. I find that coverage by employer-provided dental benefits experienced a statistically significant drop in Ontario relative to the control provinces of 4.6 to 7.5 percentage points depending on the specification; this drop is significant at the 1 percent level in all specifications. Coverage by employer-provided supplementary health insurance dropped in Ontario relative to the control provinces by 0.01 to 2.0 percentage points depending on the specification. However, this effect is never statistically significant.

4.4. Effect of the Quebec Reform on Coverage by Non-Group Supplementary Health Insurance

We have seen that the reduction in the tax subsidy coincided with a substantial decrease in coverage by employer-provided supplementary health insurance in Quebec. It remains to examine the effect of the reform on total insurance coverage. I therefore examine the effect of the decrease in the subsidy to employer-provided supplementary health insurance in Quebec on coverage by individual (non-group) supplementary health insurance, which is not subsidized through the tax code.[27] Marquis and Long (1995) have examined the effect of regional differences in price in the non-group market on the effect of

[26] Conversation with Irene Klatt, Director of Health Policy, Canadian Life and Health Association, October 1999.
[27] In the United States, a tax subsidy to health insurance for the self-employed who buy in the individual market has been in effect since 1988. A similar subsidy was introduced for the self-employed in Canada in 1998. In both

non-group coverage in the U.S. Gruber and Poterba (1994) have examined how the self-employed responded to the introduction of a subsidy to their purchase of non-group insurance in the U.S. However, I know of no work that has looked at the effect of the tax subsidy to employer-provided health insurance on both group and non-group coverage.

Data on the individual health insurance market in Canada are limited. There are no available micro data. The data on the individual market analyzed here come from the Canadian Life and Health Insurance Association's (CLHIA) annual survey of its member companies.[28] CLHIA data are available on an annual basis at the provincial level. Data consist of the number of individuals covered, the total dollar value of premia paid, and the total dollar value of claims made.[29]

CLHIA's membership consists of over 95% of the premia written for group health insurance by for-profit insurance companies, and just under 90% of the premia written for individual health insurance by for-profit insurance companies. The high, stable market share of CLHIA's member companies suggest that their data will provide an accurate reflection of changes that occurred in coverage sold by for-profit companies. The major limitation of the data, however, is that they do not include sales by non-profit insurance companies such as provincial Blue Cross organizations. CLHIA member companies therefore account for only about 70 percent of the total number of private insurance policies in Canada.[30] Since the CLHIA data tell us nothing about changes in coverage sold by non-profit insurers, the data will present a misleading picture if trends in non-profit insurance sales by province differed from trends in for-profit insurance sales by province.

Figure 1 shows trends in individual (non-group) coverage of supplementary health insurance in Quebec relative to the control provinces. Annual data are from December 31 of the year shown. The

---

countries, employed individuals who are not offered – or choose not to purchase – health insurance through their employers are not eligible for any tax subsidy in the individual market.

[28] The CLHIA is an industry-wide organization of life and health insurance companies in Canada. Their data are widely used; for example, the Canadian Department of Finance uses CLHIA data in estimating the federal tax expenditure from the tax subsidy (Government of Canada, 1998).

[29] I am grateful to Alice Freeburn, Director of Statistics at the CLHIA, for providing me with these data.

[30] In 1994, CLHIA companies covered 14.2 million individuals with supplementary health insurance. Mercer (1995) reports that in 1994, there were approximately 20 million Canadians with some supplementary health coverage under a private plan.

figure indicates that between 1989 and 1992, trends in the number of policyholders in the individual market are roughly similar in Quebec and the control provinces, although there is a slight increase in Quebec relative to the control provinces (about 12,000 policyholders between 1989 and 1992). Between the end of 1992 and the end of 1993 – the period of the Quebec reform – the number of individual policyholders fell sharply in the control provinces relative to Quebec. There is also some indication of a larger increase in the number of policyholder during 1994 and 1995 in Quebec relative to the control provinces. Taken together, the evidence is suggestive of an increase in coverage in the individual market in Quebec relative to the control provinces.[31]

But the magnitude of the increase is small relative to the decrease observed in the group market. The increase in coverage in the individual market in Quebec relative to the control provinces between 1991 and 1994 is about 38,500 policies. Allowing for the fact that the CLHIA data encompass about 70 percent of the total market for health insurance in Canada, this suggests an increase in coverage in the individual market in Quebec relative to the control provinces of roughly 55,000 policies. This increase in the individual market represents only 15% of the estimated 360,000 policy decrease in group coverage in Quebec relative to the control provinces between 1991 and 1994.[32] Of course, this may be an overestimate of the effect of the reform on coverage in the individual market, given the slight upward trend in coverage in the individual market in Quebec relative to the control provinces before the reform. Adjusting for this trend – about 4,000 policies per year – the increase in coverage in the individual market between 1991 and 1994 attributable to the reform is then only about 38,000 policies, or just over 10% of the estimated decrease in group coverage.

The reduction in the tax subsidy to employer-provided supplementary health insurance is therefore associated with a substantial net decline in insurance coverage. With any good, we expect that the reduction of a subsidy – i.e. an increase in price – results in a net decrease in demand. Without

[31] The analysis looks the same when coverage rates in the adult population are analyzed instead of number of policies.

[32] This estimate for the group market is based on GSS data of a sample of all ages, to make it comparable to the estimate for the individual market. It is not restricted to 25-64 year olds, as is the analysis in the rest of the paper.

external evidence about the price elasticity of demand for supplementary health insurance, it is not possible to say whether the lack of a more substantial offset in the individual market in Quebec is attributable solely to the fact that the individual market is not subsidized or whether it is also due to the higher administrative costs and greater scope for adverse selection in the individual market relative to the group market.[33]

5. The Role of the Tax Subsidy in Creating Workplace Based Insurance Pools

We have seen that the reduction in the tax subsidy to employer-provided supplementary health insurance in Quebec coincided with a substantial decrease in coverage by employer-provided supplementary which was not offset by an increase in coverage by the non-group market. Yet the tax subsidy is not the only advantage offered by workplace-based insurance over the individual market. A central question is the degree to which the predominance of health insurance through the workplace is a function of the tax subsidy rather than of the economies of scale and reduced adverse selection through pooling that workplace-based insurance also offers.

This section tries to shed some light on the relative role of tax subsidies in creating workplace-based insurance pools by examining how the effect of the reduction of the tax subsidy in Quebec varied across firms of different size. As noted by Stabile (2000), the tax subsidy to employer-provided supplementary health insurance applies to firms regardless of their size, but the other potential gains associated with workplace-based health insurance – reduced administrative costs and adverse selection problems – are much more of a factor in large firms rather than small firms. The role of the tax subsidy may therefore be quite different in large and small firms, and this differential role can shed light on the importance of other advantages of workplace-based insurance.

Data on firm size are available in the 1989 GSS and the 1994 GSS but not in the 1991 GSS. Results by firm size are therefore reported based on estimation of equation (3) using a pooled cross-section of the 1989 and 1994 GSS. All prior estimation of equation (3) used a pooled cross-section of the

[33] The CLHIA data indicate that administrative loads – measured as the excess of premia over claims as a percentage of claims paid – are an order of magnitude lower in the group market compared to the individual market (21 percent

1991 and 1994 GSS. As is discussed in more detail in the falsification exercises below, coverage by employer-provided supplementary health insurance is trending similarly in Quebec and the control provinces between 1989 and 1991. Therefore any difference between Quebec and the control provinces in the trend in coverage between 1989 and 1994 is presumably driven by the same factors that drives such differences between 1991 and 1994.

Table 7 reports the estimated effect of the reform for workers in firms of different sizes. The effect of the reform decreases mostly monotonically with the size of the firm.[34] The largest difference in the effect of the reform by firm size is observed between firms with less than twenty employees and all other sized firms. Indeed, in all of the specifications with covariates, the difference between the effect of the reform in firms with less than 20 employees and the effect of the reform in any other size category of firm is statistically significant at at least the 5% level. Differences in the effect of the reform between any of the other three size groupings are never statistically significant. The effect of the tax subsidy therefore appears to be substantially greater in firms with less than 20 employees than in firms of larger size.

The estimated effect of the reform in firms with less than 20 employees ranges from a 19 percentage point decrease to a 26 percentage point decrease. In firms with 500 or more employees, by contrast, the reform is associated with only a 6 to 7 percentage point decrease in coverage by employer-provided supplementary health insurance. Before the reform, the probability of receiving employer-provided supplementary health insurance in Quebec was 52 percent in firms with less than 20 employees, compared to 95 percent in firms with more than 500 employees. Therefore, the Quebec reform was associated with a 36 to 50 percent decrease in group coverage in firms with less than 20 employees but only a 6 to 7 percent decrease in firms with more than 500 employees.

A linear extrapolation from these effects to the effect of a complete removal of the tax subsidy, which would have raised the tax price from 57 cents to $1.11 (because of the excess consumption taxes), suggests that complete removal of the tax subsidy would reduce coverage by employer-provided

---

compared to 213 percent).

supplementary health insurance by 13 to 15 percentage points in firms with more than 500 employees and by 41 to 56 percentage points in firms with less than 20 employees. Results based on out of sample predictions need to be interpreted with caution and are sensitive to the linearity assumption. With these caveats in mind, however, this extrapolation, together with the initial coverage levels in small and large firms, is supportive of the idea that complete removal of the tax subsidy to employer-provided supplementary health insurance in Canada could completely eliminate employer-provided supplementary health insurance in small firms, while reducing coverage in large firms by only about 15 percent.

These results therefore suggest that the tax subsidy plays a very important role in getting small firms to offer supplementary health insurance. Since problems of adverse selection and high administrative loads are likely to be almost as severe in firms with less than 20 employees as in the non-group market, it is not surprising that the tax subsidy is critical to getting such firms to offer insurance.[35] In large firms, where gains from pooling and economies of scale are substantially larger, the tax subsidy appears to have a small but not solitary role in getting such firms to provide insurance. This suggests that while the prevalence of insurance in large firms is not primarily a function of a the tax subsidy, the high rates of insurance provision in very small firms compared to the individual market is likely to arise almost entirely from the subsidy.

6. Falsification Exercises

The interpretation of the main results in the preceding sections is based on the identifying assumption ~~that,~~ absent the tax reform in Quebec, coverage by employer-provided supplementary health insurance in Quebec and in the control provinces would have followed similar trends between 1991 and 1994. In this section, I take two different approaches to testing the validity of this identifying assumption. First, I look at whether coverage by employer-provided supplementary health insurance was trending

---

[34] This finding of a larger effect of the tax subsidy in smaller firms is consistent with similar comparisons by firm size done by Gruber and Lettau (2000) and Stabile (2000).

[35] For example, in the United States, administrative loads decline monotonically from 40 percent in firms with 1-4 employees to 5.5 percent in firms with more than 10,000 employees (Council of Economic Advisers 1994).

similarly in Quebec and the control provinces before the reform was enacted.[36] To do this, I compare changes in coverage by employer-provided supplementary health insurance in Quebec between 1989 and 1991 to changes in the control provinces between 1989 and 1991.[37] In other words, I re-estimate equation (3) using 1989 as the "before" period and 1991 as the "after" period.

Panel A of Table 8 reports the coefficient on the "treatment effect" (QUEBEC*AFTER) from these regressions. The small and insignificant estimated effect of QUEBEC*AFTER indicates that coverage by employer-provided supplementary health insurance is following the same trend in Quebec relative to the control provinces between 1989 and 1991. Figure 2 illustrates this result graphically. This supports the identifying assumption for health insurance.[38]

The second type of falsification exercise looks at changes in coverage rates for two employee other benefits between 1991 and 1994: retirement pensions and paid parental leave. We might be concerned that there was a Quebec-specific shock between 1991 and 1994, other than the tax reform, which caused the relative drop in coverage by employer-provided supplementary health insurance in Quebec. One way to look for the presence of such a shock – such as a worsening business climate or increased administrative costs of running employee benefit plans – is to see whether other employee benefits that were not affected by the tax reform also experienced a relative decline in Quebec.[39]

---

[36] Although in principle I could carry out a similar exercise after the reform was enacted, in practice the GSS does not provide any information on employer-provided supplementary health insurance after the 1994 Cycle.

[37] 1989 is the first year before 1991 in which the GSS asks about employer-provided supplementary health and dental benefits. The 1989 GSS follows the same survey design as the 1991 and 1994 GSSs. There is a minor difference in the phrasing of the question used for the dependent variable for medical benefits. While the 1991 and 1994 surveys ask about employer provision of medical benefits "beyond those provided by your provincial health care system," the 1989 question does not include this caveat. However, people whom I contacted felt that individuals would not confuse this question with publicly provided insurance (Conversations with Alice Freeburn, Director of Statistics Canadian Life and Health Association, and Nancy Turner, Housing, Family, and Social Statistics Division of Statistics Canada. October 1999). Moreover, even were there to be some confusion, this would only contaminate the analysis if the probability of misinterpreting the question differed systematically between residents of Quebec and residents of the control provinces. Again, there is no evidence to suggest that this would be the case.

[38] In results not reported here, I also perform this exercise for dental benefits. The results for dental benefits are not supportive. Coverage by employer-provided dental benefits is rising in Quebec relative to the control provinces between 1989 and 1991. In fact, the relative rise is of slightly larger magnitude than the relative drop between 1991 and 1994. This again points to caution in interpreting the estimated effect of the 1993 reform on dental benefits.

[39] Neither retirement pensions nor paid parental leave were affected by the 1993 Quebec reform. Employer contributions to private retirement pensions are excluded from an employee's taxable income. Paid parental leave is a taxable benefit (Koskie 1995; Nielson 1998).

I therefore re-estimate equation (3) using coverage by employer-provided retirement pensions and coverage by employer-provided paid parental leave as the dependent binary variables, instead of supplementary health insurance which was used in all of the other analyses.[40] Panel B of Table 8 reports the results of this set of falsification exercises. Figure 3 shows the results graphically for employer-provided pensions. Across all specifications, there is no significant change in coverage rates of either pensions or paid parental leave in Quebec relative to the control provinces between 1991 and 1994. The estimated changes in these employee benefits in Quebec relative to the control provinces are not only statistically insignificant, but are also extremely small in magnitude compared to the effects observed for supplementary health insurance. For example, for pensions, the largest estimated effect is an insignificant drop of 0.5 percentage points in coverage in Quebec relative to the control provinces.

The assumption underlying this falsification exercise is that the determinants of coverage for employer-provided supplementary health insurance are similar to the determinants of coverage for these other employee benefits. Several other results lend support to this hypothesis. First, the coverage rates for these other benefits is on the same order of magnitude as the coverage rate for supplementary health insurance.[41] Second, the effects of various covariates on the probability of coverage by a given employee benefit are extremely similar across the different benefits. And third, I find that in the pre-reform period (1989-1991), pensions and paid parental leave coverage, like coverage by supplementary health insurance, are on a downward trajectory.

7. Conclusion

This paper has presented new evidence of the effect of the tax subsidy to employer-provided supplementary health insurance on coverage by such insurance. The 1993 removal of the Quebec

[40] The General Social Survey provides information on employer provision of one other benefit: counseling services for a variety of personal problems. However, the 1993 Quebec reform that made employer contributions to health and dental benefits taxable also reduced the list of problems for which employer-financed counseling services are tax exempt. In results not reported here, I find some evidence of a decrease in counseling services in Quebec relative to the rest of Canada. The estimated decrease ranges from 4.4 to 5.6 percentage points and is significant at the 5% level in both the limited and full covariates specification.

[41] The probability of coverage for the full sample is 0.67 for employer-provided supplementary health insurance, 0.59 for pensions, and 0.43 for paid parental leave.

provincial tax subsidy to employer-provided supplementary health insurance, while leaving in place the federal tax subsidy and the payroll tax subsidy, reduced the total tax subsidy to employer-provided supplementary health insurance in Quebec by almost 60 percent. I estimate that this reduction in the tax subsidy resulted in a 13 to 14 percentage point decrease in supplementary health insurance coverage through the workplace. This represents a decline in workplace coverage of about one fifth, and corresponds to an elasticity of coverage by employer-provided supplementary health insurance with respect to the tax price of –0.46 to –0.49. This elasticity lies well within the range of previously estimated elasticities in the United States of employer offering of health insurance with respect to the tax subsidy. I also find that decreases in coverage through the workplace associated with the tax reform are only slightly offset by increases in coverage in the non-group market. I estimate that the increase in coverage in the non-group market in Quebec relative to the control provinces was only about 10 to 15 percent of the decrease in workplace coverage.

The evidence also suggests that a tax subsidy to employer-provided supplementary health insurance plays a large role in explaining the provision of such insurance in small firms, but a considerably smaller role in large firms. This result is consistent with the notion that large firms offer other advantages as a source of insurance – such as risk pooling and economies of scale in administration – that would make them attractive venues for insurance even absent the subsidy. The appeal to offering supplementary health insurance in small firms, however, where gains from risk pooling and economies of scale are considerably lower, appears to be much more sensitive to the tax subsidy.

Several interesting questions relating to the effects of the tax subsidy to employer-provided supplementary health insurance were not explored in this paper due to limitations of the data. First, it would be interesting to study the effect of the tax subsidy on other margins of coverage besides the binary measure of whether the individual's employer provides coverage. The response to a change in the tax subsidy may occur partly through changes in intensive margins of coverage, such as decreases in maximum claims, reductions in services covered, and increases in employee co-payments and deductibles. To the extent that such changes occurred, the examination solely of the extensive coverage margin undertaken here underestimates the decrease in insurance coverage due to the tax reform. Second, the reduction of the tax subsidy to employer-provided health insurance reduces the rationale for employer (as opposed to employee) contributions to health insurance premia. Increased employee contributions in turn should affect margins such as take-up conditional on offering and therefore, in turn, the scope for adverse selection in a workplace pool. Finally, it would also be interesting to explore the consequences of the reduction in supplementary health insurance for utilization of privately and publicly insurance health services, and for health outcomes.

References

Canadian Life and Health Association. Various Years. Health Insurance Benefits In Canada. Toronto, Ontario.

Canadian Tax Foundation. 1991. The National Finances.

Canadian Tax Foundation. 1994. The National Finances.

Congressional Research Service. 1990. The Canadian Health Care System.

Council of Economic Advisers. 1994. Economic Report of the President. Washington: United States Government Printing Office.

Cutler, D. 1996. "Comment" to "Fundamental Tax Reform and Employer-Provided Health Insurance," Jonathan Gruber and James Poterba, in Economic Effects of Fundamental Tax Reform, H. Aaron and W. Gale (eds), Washington: Brookings Institution, pp.125-162.

Employee Benefits Research Institute. 1995. EBRI Databook on Employee Benefits. Third Edition: Washington.

GAO. 1992. Prescription Drug Charges in US and Canada. GAO/HRD-92-110.

Gentry, W. and E. Peress. 1994. "Taxes and Fringe Benefits Offered By Employers." NBER Working Paper 4764.

Goldstein, G. and M. Pauly. 1976. "Group Health Insurance as a Local Public Good," in The role of health insurance in the health services sector," ed. R. Rosett. NY: NBER, pp.73-110.

Government of Canada. 1998. Tax Expenditures 1998.

Government of Quebec. 1996. Tax Expenditures.

Gruber, J.. 1999. "Tax Subsidies for Health Insurance: What Do We Need to Know, What Do We Know, and How Could We Learn More?" MIT Mimeo.

Gruber, J. and M. Lettau. 2000. "How Elastic is the Firm's Demand for Health Insurance?" MIT Mimeo.

Gruber, J. and J. Poterba. 1994. "Tax Incentives and the Decision to Purchase Health Insurance: Evidence from the Self-Employed." *Quarterly Journal of Economics,* August, pp.701-733.

Gruber, J. and J. Poterba. 1996a. "Fundamental Tax Reform and Employer-Provided Health Insurance," in Economic Effects of Fundamental Tax Reform, H. Aaron and W. Gale (eds), Washington: Brookings Institution, pp.125-162.

Gruber, J. and J. Poterba. 1996b. "Tax Subsidies to Employer-Provided Health Insurance," in Empirical Foundations of Household Taxation, M. Feldstein and J. Poterba, eds. Chicago: University of Chicago Press, 135-164.

Holmer, M. 1984. "Tax Policy and the Demand for Health Insurance." *Journal of Health Economics,* III, pp.203-221.

Koskie, R. et al. 1995. Employee Benefits in Canada. International Foundation of Employee Benefit Plans, Inc.

Leibowitz, A. and M. Chernew. 1992. "The Firms' Demand for Health Insurance," in Health Benefits and the Workforce, Washington, DC: U.S. Department of Labor.

Long, J. and F. Scott. 1982. "The Income Tax and Nonwage Compensation," *Review of Economics and Statistics,* LXIV, pp. 211-219.

Marquis, S. and C. Phelps. 1987. "Price Elasticity and Adverse Selection in the Demand for Supplementary Health Insurance," *Economic Inquiry,* XXV, pp.229-313.

Marquis, S. and S. Long. 1995. "Worker demand for health insurance in the non-group market," *Journal of Health Economics*, 14: 47-63.

Mercer, William H Limited. 1995. Supplementary Health and Dental Programs for Canadians: Assessment of Coverage and Fairness of Tax Treatment.

Nielson, N. 1998. Taxes and Employee Benefits in Canada. International Foundation of Employee Benefit Plans, Inc.

Office of Management and Budget. 1999. Budget of the United States Government (2000). Washington, D.C.: Government Printing Office.

Pauly, M. 1986. "Taxation, Health Insurance, and Market Failure in the Medical Economy." *Journal of Economic Literature*, XXIV: pp.629-675.

Phelps, C. 1973. The Demand for Health Insurance: A Theoretical and Empirical Investigation. Santa Monica, CA: Rand.

Royalty, A. 1999. "Tax Preferences for Fringe Benefits and Workers' Eligibility for Employer Health Insurance," *Journal of Public Economics*, forthcoming.

Sloan, F. and K. Adamche. 1986. "Taxation and the Growth of Nonwage Benefits." *Public Finance Quarterly,* XIV, pp.115-139.

Stabile, M. 1999. "Private Insurance Subsidies and Public Health Care Markets: Evidence from Canada." Mimeo, University of Toronto.

Stabile, M. 2000. "The Role of Tax Subsidies in the Market for Health Insurance." Mimeo, University of Toronto.

Statistics Canada. 1997. Historical Labour Force Statistics. Ottawa.

Taylor, A. and G. Wilensky. 1983. "The Effect of Tax Policies on Expenditures for Private Health Insurance," in J. Meyer, ed. Market Reforms in Health Care: Current Issues, New Directions, Strategic Decisions. Washington, DC: American Enterprise Institute.

Turner, R. 1987. "Are Taxes Responsible for the Growth in Fringe Benefits?" *National Tax Journal,* XL pp.205-220

Woodberry, S. and D. Hammermesh. 1992. "Taxes, Fringe Benefits, and Faculty." *Review of Economics and Statistics,* LXXIV pp.287-96.

Woodberry, S. 1983. "Substitution between Wage and Nonwage Benefits," *American Economic Review,* LXIII, 166-182.

World Health Organization. 1996. Health Care Systems in Transition: Canada. Copenhagen.

Wyatt. 1993. Survey of Group Benefit Plans Covering Salaried Employees of Canadian Employers. The Wyatt Company, Toronto Ontario.

Table 1: Tax Rates and Tax Prices

	Quebec		Control Provinces	
	1991	1994	1991	1994
Federal Marginal Tax Rate	0.197	0.196	0.227	0.232
Provincial Marginal Tax Rate	0.214	0.235	0.110	0.117
Payroll Marginal Tax Rate (Combined Employer plus Employee Rate)	0.068	0.076	0.068	0.073
Excess Consumption Tax on Group Supplementary Health Insurance Relative to other Services	0.110	0.1115	0.022	0.022
Tax Price of Supplementary Health Insurance	0.597	0.822	0.623	0.607

Table reports average tax rate in each cell. All means are weighted. Tax Price is defined by equation (1) for all cells except Quebec in 1994 where it is defined by equation (2).

Table 2: Weighted Means of Covariates
Quebec vs. Control Provinces

	1991 Control Provinces	1991 Quebec	1994 Control Provinces	1994 Quebec	Difference in Differences
Union Member	0.365 (0.482)	0.439 (0.497)	0.353 (0.478)	0.447 (0.498)	0.020 (0.032)
Married	0.761 (0.427)	0.743 (0.437)	0.745 (0.436)	0.774 (0.419)	0.046* (0.026)
Male	0.538 (0.499)	0.577 (0.494)	0.548 (0.498)	0.558 (0.497)	-0.029 (0.032)
Spouse Employed (among married)	0.731 (0.444)	0.750 (0.434)	0.726 (0.446)	0.720 (0.450)	-0.024 (0.036)
EDUCATION					
Less than High School	0.199 (0.399)	0.277 (0.448)	0.159 (0.366)	0.222 (0.416)	-0.016 (0.028)
High School Diploma	0.151 (0.359)	0.166 (0.373)	0.191 (0.393)	0.194 (0.396)	-0.012 (0.025)
Some College or Postsecondary Schooling	0.465 (0.499)	0.349 (0.477)	0.418 (0.493)	0.329 (0.470)	0.027 (0.031)
(4 yr) College Degree of Higher	0.185 (0.389)	0.207 (0.405)	0.232 (0.422)	0.255 (0.436)	0.001 (0.027)
HEALTH					
Excellent Health	0.270 (0.444)	0.286 (0.452)	0.292 (0.455)	0.382 (0.486)	0.074*** (0.030)
Very Good Health	0.376 (0.485)	0.338 (0.473)	0.375 (0.484)	0.305 (0.461)	-0.032 (0.031)
Good Health	0.285 (0.452)	0.280 (0.449)	0.266 (0.442)	0.223 (0.416)	-0.038 (0.029)
Fair Health	0.057 (0.233)	0.078 (0.269)	0.058 (0.234)	0.067 (0.250)	-0.013 (0.016)
Poor Health	0.012 (0.107)	0.018 (0.134)	0.009 (0.095)	0.023 (0.151)	0.008 (0.009)
Employed Full Year	0.811 (0.392)	0.821 (0.384)	0.808 (0.394)	0.856 (0.351)	0.038 (0.024)
Employed Full Time	0.891 (0.311)	0.911 (0.285)	0.891 (0.312)	0.893 (0.309)	-0.017 (0.019)
Number of Observations (unweighted count)	2,066	827	2,259	834	

Notes: All means are weighted. Standard deviations are in parentheses in the first 4 columns. Standard errors are in parentheses in the difference-in-differences column, which compares means in Quebec before and after the reform to means in the control provinces before and after the reform. Standard errors come from regressing the characteristics on year and region trends and an interaction term, and are adjusted for heteroscedasticity. *** Denotes significant difference-in-differences at the 1% level. ** Denotes significant difference-in-differences at the 5% level. * Denotes significant difference-in-differences at the 10% level.

Table 3: Difference in Differences: Probability of Coverage by Employer-Provided Supplementary Health Insurance, Quebec vs. Control Provinces

	Quebec	Control Provinces	Difference
1991	0.723	0.641	0.082 (0.021)
1994	0.644	0.693	-0.050 (0.022)
Difference	-0.080 (0.025)	0.052 (0.017)	-0.132 (0.030)

Note: All means are weighted. Heteroscedasticity-adjusted standard errors are in parentheses.

Table 4: Difference-in-Differences
Probability of Coverage by Employer-Provided Supplementary Health Insurance, Quebec vs. Control Provinces

	(1) Difference-in-Differences	(2) Limited covariates specification	(3) Full covariates specification
Quebec*After	-0.132***	-0.131***	-0.136***
	(0.030)	(0.026)	(0.020)
Quebec	0.082***	0.076***	0.038*
	(0.021)	(0.022)	(0.019)
After	0.052***	0.048*	0.020
	(0.017)	(0.025)	(0.020)
Union Member		0.296***	0.220***
		(0.019)	(0.014)
Prov. Unemployment Rate		-0.008***	0.001
		(0.002)	(0.003)
Age 35-44		0.039***	-0.018
		(0.009)	(0.011)
Age 45-54		0.021	-0.048***
		(0.013)	(0.014)
Age 55-64		-0.029***	-0.075***
		(0.010)	(0.017)
Male		0.095***	-0.053***
		(0.018)	(0.017)
Married		0.073**	0.022
		(0.030)	(0.028)
High School Diploma		0.111***	0.005
		(0.020)	(0.019)
Some College or Post-Secondary Schooling		0.134***	-0.003
		(0.012)	(0.016)
(4 year) College Degree or Higher		0.218***	-0.014
		(0.012)	(0.011)
Spouse Employed		-0.022	0.001
		(0.019)	(0.016)
1 Child under 25		-0.031	-0.023
		(0.026)	(0.021)
2 Children under 25		-0.016	-0.008
		(0.026)	(0.021)
3+ Children under 25		-0.032	-0.017
		(0.034)	(0.027)
Excellent Health			-0.007
			(0.016)
Very Good Health			-0.004
			(0.015)
Fair Health			0.015
			(0.035)
Poor Health			0.059
			(0.046)
Full Time			0.139***
			(0.037)
Full Year			0.151***
			(0.010)
Occupation Dummies			Yes
Pers. Inc. Bracket Dummies			Yes
Constant	0.641***	0.407***	0.109
	(0.012)	(0.035)	(0.071)
R^2	0.005	0.146	0.343
N	5,986	5,850	5,377
Mean of Dependent Variable	0.667	0.668	0.675

Notes: Coefficients are from estimates of equation (3). Standard errors are in parentheses. They are adjusted for heteroscedasticity and in columns (2) and (3) are also adjusted for province*year correlation in the error term. *** denotes significance at the 1% level. ** denotes significance at 5% level. * denotes significance at 10% level. Regressions are weighted. Omitted age category is age 25-34. Omitted education category is less than high school. Omitted children category is no children under 25. Omitted health category is "good health." There are 34 occupation dummies and 9 personal income bracket dummies.

Table 5: Effect of Tax Price on Probability Employer Provides Supplementary Health Insurance

	Ordinary Least Squares	Instrumental Variables	
	(1) Full Covariates Specification	(2) Limited Covariates Specification	(3) Full Covariates Specification
Coefficient on TAXPRICE	-0.322*** (0.080)	-0.595*** (0.070)	-0.560*** (0.079)
Implied Elasticity of Coverage with respect to the Tax Price.	-0.264*** (0.066)	-0.488*** (0.057)	-0.459*** (0.065)
N	5,377	5,449	5,377

Notes: The first row reports the results from estimating equation (4). The second row reports the elasticity implied by the estimated coefficient on TAXPRICE. Elasticities are reported relative to pre-reform Quebec means. Standard errors are in parentheses. They are adjusted for heteroscedasticity and for province*year correlation in the error term. *** denotes significance at the 1% level. ** denotes significance at the 5% level. * denotes significance at the 10% level. Regressions are weighted.

Table 6: Effect of Quebec Tax Reform on Coverage by Employer-Provided Supplementary Health Insurance for Different Types of Workers

	(1) Difference-in-Differences	(2) Limited Covariates Specification	(3) Full Covariates Specification
Gender			
Men	-0.179***	-0.176***	-0.193***
	(0.040)	(0.037)	(0.033)
Women	-0.062	-0.065***	-0.052***
	(0.045)	(0.019)	(0.017)
Union Status			
Non-Union	-0.179***	-0.173***	-0.149***
	(0.043)	(0.027)	(0.022)
Union	-0.083**	-0.076**	-0.097***
	(0.034)	(0.030)	(0.025)
Age			
25 to 34	-0.117**	-0.115***	-0.115***
	(0.050)	(0.020)	(0.025)
35 to 44	-0.147***	-0.147***	-0.128***
	(0.049)	(0.027)	(0.022)
45 to 54	-0.100	-0.103*	-0.135**
	(0.071)	(0.051)	(0.049)
55 to 64	-0.214**	-0.210***	-0.296***
	(0.096)	(0.029)	(0.042)

Notes: Table reports the estimated coefficient on QUEBEC*AFTER in equation (3). In each row, equation (3) is estimated separately for the group described in the left-hand column. Standard errors are in parentheses. They are adjusted for heteroscedasticity and in columns (2) and (3) are also adjusted for province*year correlation in the error term. *** denotes significance at the 1% level. ** denotes significance at 5% level. * denotes significance at 10% level. Significance refers to whether effect is significantly different from zero. Regressions are weighted.

Table 7: Effect of Quebec Tax Reform on Coverage by Employer-Provided Supplementary Health Insurance for Different Sized Firms

	(1) Difference-in-Differences	(2) Limited Covariates Specification	(3) Full Covariates Specification
500+ employees	-0.067**	-0.058***	-0.068***
	(0.031)	(0.016)	(0.014)
100-499 employees	-0.071	-0.071**	-0.058
	(0.059)	(0.033)	(0.050)
20-99 employees	-0.152***	-0.073	-0.092
	(0.075)	(0.065)	(0.058)
Less than 20 employees	-0.202***	-0.191***	-0.257***
	(0.072)	(0.033)	(0.023)

Notes: Table reports the estimated coefficient on QUEBEC*AFTER in equation (3). Each row reports the estimated coefficients for individuals in a given sized firm. Regressions are based on a pooled cross section of 1989 and 1994 data. Therefore the full covariate specification could not control for health status, since this variable is not available in 1989. Standard errors are in parentheses. They are adjusted for heteroscedasticity and in columns (2) and (3) are also adjusted for province*year correlation in the error term. *** denotes significance at the 1% level. ** denotes significance at 5% level. * denotes significance at 10% level. Significance refers to whether effect is significantly different from zero. Regressions are weighted.

Table 8: Falsification Exercises

Dependent Variable	(1) Difference-in-Differences	(2) Limited Covariates Specification	(3) Full Covariates Specification
Panel A: Probability Employer Provides Benefits 1989 v. 1991 Quebec vs. Control Provinces			
Supplementary Health Insurance	0.033	0.042	0.017
	(0.029)	(0.034)	(0.030)
Panel B: Probability Employer Provides Other Employee Benefits 1991 v. 1994 Quebec vs. Control Provinces			
Pension	0.002	-0.003	-0.005
	(0.032)	(0.035)	(0.045)
Paid Parental Leave	-0.010	-0.021	-0.022
	(0.034)	(0.063)	(0.063)

Notes: Standard errors are in parentheses. They are adjusted for heteroscedasticity and in columns (2) and (3) are also adjusted for province*year correlation in the error term. *** denotes significance at the 1% level. ** denotes significance at 5% level. * denotes significance at 10% level. Regressions are weighted. Panel A: Numbers in the table are the estimated coefficient on QUEBEC*AFTER in equation (3). But here, equation (3) is estimated using 1991 and 1989 data. AFTER is therefore a dummy variable equal to 1 if the observation is from the 1991 survey and 0 if it is from the 1989 survey. Columns have the same controls as the corresponding columns in previous tables (except that there are no health controls in column 3 since health status information is not available in 1989). Panel B: Numbers in the table are the estimated coefficient on QUEBEC*AFTER in equation (3). But here, equation (3) is estimated on 1991 and 1994 data with a binary dependent variable that indicates coverage by the employee benefit shown in the left hand column. Columns have the same controls as corresponding columns in previous tables.

Figure 1: Trends in Individual (Non-Group) Health Insurance

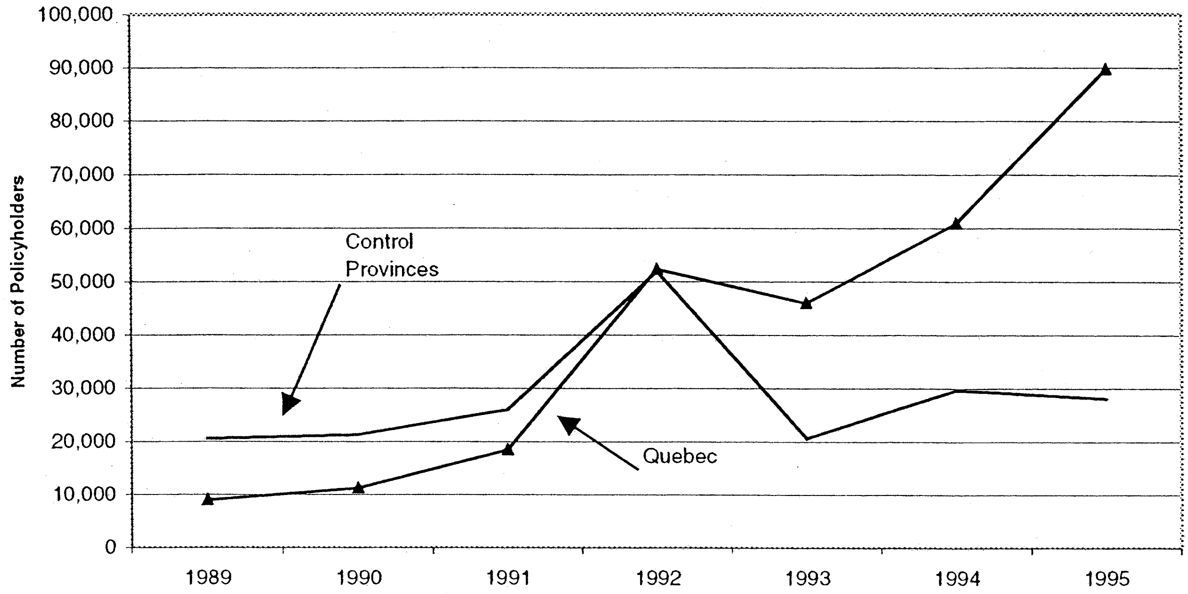

Figure 2: Trends in Employer-Provided Supplementary Health Insurance

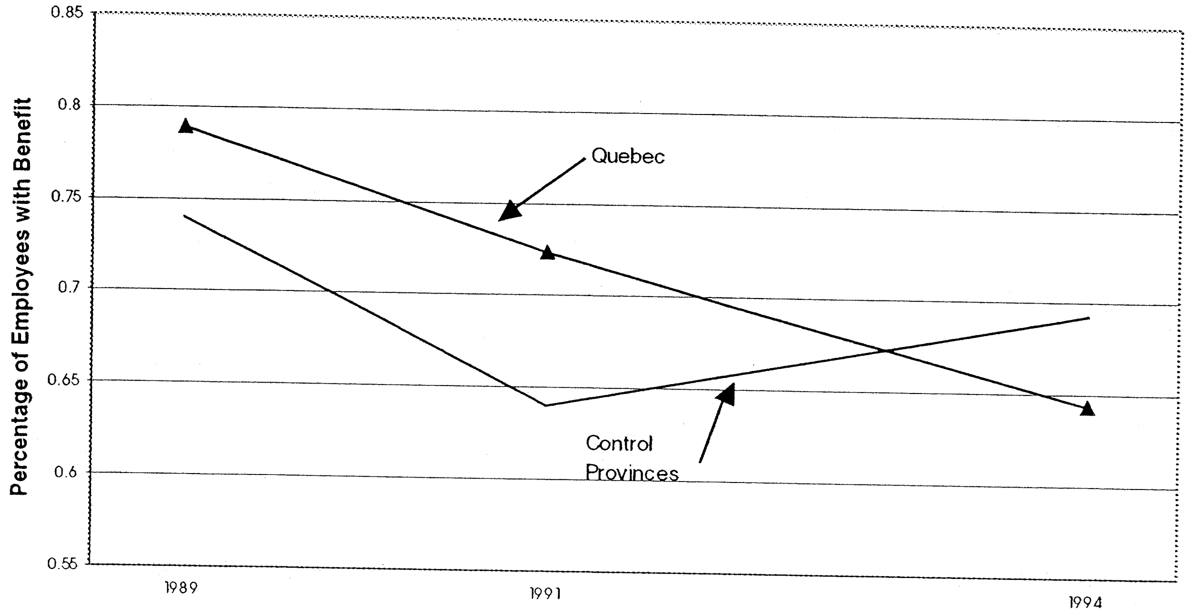

Figure 3: Trends in Employer-Provided Pensions

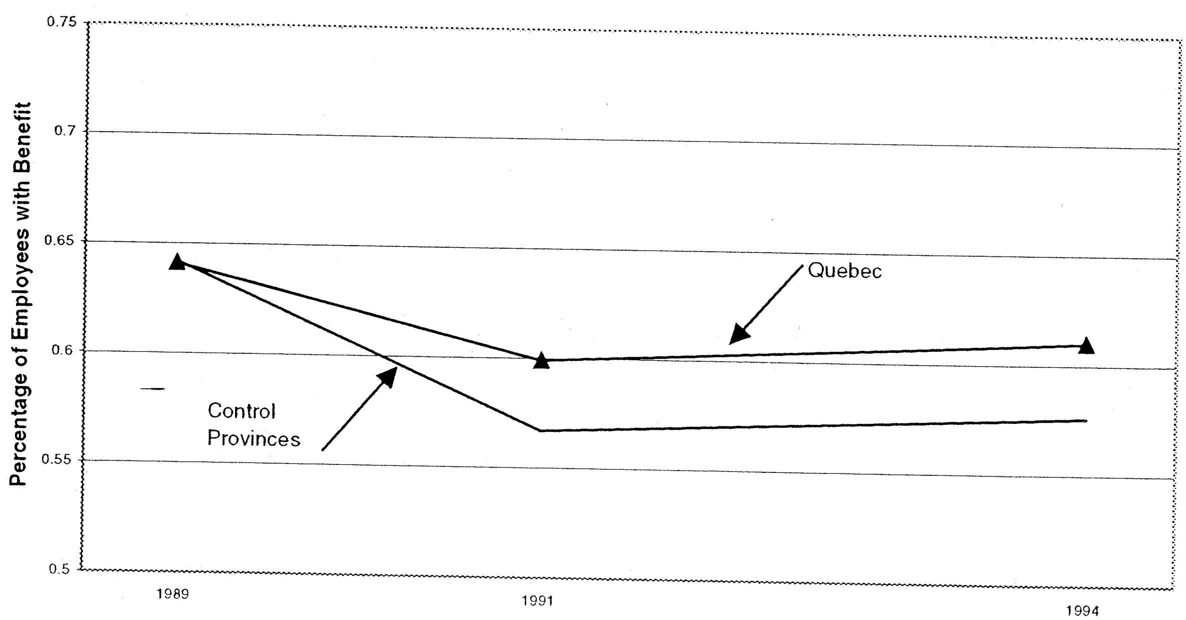

Chapter 3:
Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market

Joint with James Poterba

Theoretical research on insurance markets has long emphasized the potential importance of asymmetric information. The welfare implications of adverse selection, one of the consequences of asymmetric information, are well documented. Yet several recent empirical papers have failed to find evidence of asymmetric information in property-casualty, life, and health insurance markets. Cawley and Philipson (1999), who study the U.S. life insurance market, Cordon and Hendel (1999), who study the U.S. health insurance market, and Chiappori and Salanie (2000), who study the French automobile insurance market, all find it difficult to reject the null hypothesis of symmetric information. These findings raise new questions about whether asymmetric information is, in practice, an important feature of insurance markets.

This paper tests two simple predictions of asymmetric information models using data from the annuity market in the United Kingdom. The first is that higher risk individuals self-select into insurance contracts that offer features that, at a given price, are more valuable to them than to lower risk individuals. The second is that the equilibrium pricing of insurance policies reflects the fact that the risk pool varies across different policies. Such self-selection across policies by risk type would not occur if the insurer and the insured had symmetric information. Our empirical work finds support for both of these predictions.

Annuity markets are an interesting setting in which to study asymmetric information problems. These markets are of substantial interest in their own right. Mitchell, Poterba, Warshawsky, and Brown (1999) emphasize that annuities, which provide insurance against outliving one's resources, play a potentially important welfare-improving role for retirees. But in spite of the potential value of annuity products for households that face uncertain mortality and have fixed resources, voluntary annuity markets in both the United States and the United Kingdom are small. Asymmetric information, in particular adverse selection, has often been suggested as a potential explanation for the limited size of these markets.

Besides their intrinsic interest, there is another reason for studying asymmetric information issues in the context of annuity markets. Most tests for asymmetric information cannot distinguish between adverse selection and moral hazard, even though the welfare implications of the two, and their potential policy implications, are often quite different. Moral hazard seems likely to play a smaller role in annuity markets, however, than in many other insurance markets. While receipt of an annuity may lead some individuals to devote additional resources to life-extension, we suspect that this is likely to be a quantitatively small effect. If the behavioral effects of annuities are small, so that the associated moral hazard problems are limited, testing for asymmetric information in the annuity market provides a particularly valuable test for adverse selection.

This paper is divided into six sections. The first outlines a simple theoretical model of screening and adverse selection. It presents a generalized version of the classic Rothschild-Stiglitz (1976) model to illustrate that when insurance contracts are multi-dimensional, many features of the insurance contract can serve as screening mechanisms. Two empirical predictions of this model form the basis for our empirical tests. This section also summarizes the previous literature that has tested for asymmetric information in insurance markets.

Section two describes the general operation of annuity markets. It notes several specific features of the annuity market in the United Kingdom that are relevant for our analysis. It also refines the model from the first section to apply specifically to the annuity market setting. The third section describes the data set that we have obtained from a large U.K. insurance company, and compares the annuitants at this firm with those in the market at large.

Section four explores the relationship between mortality patterns and annuity product choice. In this section, we develop hazard models that relate the probability of annuitants dying after purchasing annuities to detailed characteristics of the annuity policy, as well as other information about the annuitant that is known to the insurance company. The empirical findings indicate that ex-post annuitant mortality varies systematically across annuity policies with different characteristics. This variation is consistent with theoretical models of asymmetric information.

The fifth section investigates the pricing of different annuity products. We report hedonic pricing models that include different features of annuity policies. We also present some evidence on the pricing formula used by the insurance company whose data we analyze. We find that annuity pricing is consistent with the presence of adverse selection in the annuity market. Features of the annuity that our hazard models indicate are selected by low mortality individuals, the "high risk" annuity buyers from the perspective of the insurance company, are priced higher than features that are selected by high-mortality individuals. We do not find evidence that mortality rates differ, or that marginal prices differ, across annuities with different initial annual annuity payment. The initial annual annuity payment is a crude measure of the size of the annuity policy, and it is our best approximation to the "amount of insurance" that features in most theoretical discussions of adverse selection. This finding underscores the importance of considering adverse selection along a range of different insurance policy attributes.

The final section summarizes our findings and considers explanations for our results that might not involve adverse selection.

1. Adverse Selection and Screening in Insurance Markets: Theory and Evidence

This section describes the nature of equilibrium in insurance markets in the presence of adverse selection, with particular emphasis on the structure of insurance market prices. The discussion is general; the next section applies the analysis to the annuity market. This section also includes a brief summary of existing empirical research on asymmetric information in insurance markets.

1.1 Theoretical Background

The standard model of insurance markets, developed for example in Rothschild and Stiglitz (1976), assumes that there are two types of individuals who differ only in their risk type p. The parameter p denotes the probability the individual will incur a claim against the insurance company. We denote these two types of individuals by p_H and p_L, with $p_H > p_L$. The individual knows his or her risk type, but the insurance company does not. Competition in the insurance market implies a zero profit condition.

The insurance company offers a menu of insurance contracts with different features. It designs this menu to screen customers according to their risk type. The key to this screening mechanism is that if the value of a contract feature differs across individuals of different risk types, the insurance company can, through appropriate pricing, induce individuals to self-select into different contracts on the basis of their (privately known) risk type. Under symmetric information, in contrast, each individual would purchase the same, efficient, insurance contract and individual-specific prices would reflect individual-specific risk.

In the Rothschild-Stiglitz (1976) framework, insurance contracts differ only in the quantity of insurance sold and the marginal price charged. In this case, the screening equilibrium involves high-risk individuals buying full insurance, as in the symmetric information case. Low-risk individuals buy the maximum amount of insurance that they can purchase while leaving the high-risk individuals indifferent between paying the marginal price p_H for full insurance and paying the marginal price p_L to buy the amount of insurance purchased by the low-risk individuals. The key to this equilibrium is that the marginal utility that the high risk type gets from an extra unit of insurance at a given price is higher than the marginal utility that the low risk type gets from an extra unit of insurance at this same price. This Spence-Mirlees (or single crossing) property ensures that there exists some quantity of insurance that the low risk type can purchase at marginal price p_L that leaves the high risk type indifferent between purchasing this quantity of insurance at price p_L and buying full insurance at the higher price p_H.

It is straightforward to generalize this model to allow for a multi-dimensional vector of insurance product characteristics. Smart (2000) discusses extensions of this type. Any features that satisfy the Spence-Mirrlees condition can potentially be used to separate different risk types into different insurance contracts. The basic result that markets characterized by asymmetric information are inefficient hold regardless of the feature, or features, of insurance contracts that are used to screen customers.

Our empirical tests for asymmetric information in the annuity market are based on two predictions that emerge from models like the one sketched above. Chiappori and Salanie (2000) note that these predictions are robust with respect to many modifications of the basic framework. First, there should be a positive correlation between (privately known) risk type and the demand for features of insurance contracts whose marginal value is greater for high-risk types. Second, insurance companies should charge more for features of insurance contracts that are purchased by higher risk types than for features that are purchased by lower risk types.

In practice, the insurance company may observe some risk features of the individual. The stylized model described above therefore applies to individuals who are observationally equivalent to the insurance company, and the foregoing predictions apply <u>conditional</u> on the characteristics of the insured that are observable by the insurance company.

1.2. Empirical Tests of Asymmetric Information in Insurance Markets

Several recent papers, linked by a common empirical strategy, have tested for adverse selection in insurance markets. These papers examine the relationship between ex-post risk type and insurance coverage, conditional on observable characteristics of individuals. Models of asymmetric information predict a positive relationship between risk type and the amount of insurance coverage, if the amount of insurance coverage is used as a screening device. Pueltz and Snow (1994) test this relationship in the U.S. automobile insurance market. They find that higher risk individuals, measured by ex-post accident probabilities, choose policies with lower deductibles, and hence greater insurance coverage. This finding is consistent with the presence of asymmetric information.

Several other recent studies reach a different conclusion, using data from different insurance markets. Chiappori and Salanie (2000) also examine automobile insurance markets, using data from France. They relate both quantity of insurance and ex-post risk type to exogenous variables, and then test for a positive correlation between the errors of the two estimating equations. They cannot reject the null hypothesis that the two sets of residuals are uncorrelated. Using a similar test, Cordon and Hendel (1999) fail to reject the null hypothesis of symmetric information in health insurance markets. In a related study, Cawley and Philipson (1999) examine the relationship between insurance coverage and risk in the U.S. life insurance market. Their results also fail to reject the null hypothesis that risk type and amount of insurance coverage are uncorrelated.

The preponderance of recent research fails to reject the null hypothesis of symmetric information in insurance markets. One common limitation of these studies is their focus on a one-dimensional screening variable: the amount of payment in the event of an accident. Chiappori and Salanie (2000) acknowledge that there are "many different comprehensive coverage contracts on offer," but they focus on a binary measure of whether the individual has anything more than the legal minimum level of coverage. Cawley and Philipson (1999) examine the relationship between risk type and the amount the term insurance policy would pay out in the event of death, while ignoring other potential screening variables such as the renewability of a life insurance policy. In practice, many aspects of insurance contracts may satisfy the Spence-Mirrlees condition and therefore have the potential to serve as screening mechanisms. We emphasize this, and try to fashion appropriate empirical tests, in our subsequent analysis. While we do not find any evidence of screening on the amount of payment that the insurance policy prescribes, we do find evidence of asymmetric information using other screening devices.

Several of these studies also test our second empirical prediction that features of insurance that are selected by higher risk individuals should be priced correspondingly higher. This involves, at a general level, testing the prediction that insurance policies with more comprehensive coverage sell at higher ~~marginal~~ prices. Pueltz and Snow (1994) find a concave premium-deductible schedule, which suggests that low risk individuals, who choose large deductibles, receive lower marginal prices for their choice. Cawley and Philipson (1999), however, find no evidence that the marginal price of insurance rises with the amount of payment of the event of death. This is consistent with their finding of a lack of correlation between risk type and amount of insurance coverage.

2. Background on Annuities and Annuity Markets

This section presents an overview of the annuity market in the United Kingdom. It then applies the framework developed in the last section to this market, and formalizes the model-based predictions that underpin our empirical work.

2.1 Annuity Market Overview

An annuity is a contract that pays a pre-specified amount to a beneficiary, the annuitant, for as long as the annuitant is alive. It thus insures the annuitant against the risk of outliving his accumulated resources. In some ways annuities function as reverse life insurance, insuring against the risk of living too long rather than against the risk of dying too soon. From the standpoint of an annuity insurance company, a high-risk individual is an individual who is likely to live longer than his observable attributes, such as age, would otherwise suggest.

Yaari (1965) documented the welfare-improving role of annuity insurance for individuals facing uncertain mortality. In light of this, the small size of the voluntary annuity market in the United States and the United Kingdom has puzzled many researchers. One possible explanation for the small market size is adverse selection. A number of previous studies suggest that annuity markets are characterized by some adverse selection. The pricing of voluntary annuities in both the U.S. and the U.K. implies that, for a typical individual from the population, the expected present discounted value of the payments from an annuity are only 80 to 85 percent of the annuity's initial premium. Part of the divergence between this expected payout and the cost of the annuity is due to administrative loads, but roughly half appears to be due to adverse selection. If the typical annuitant is longer lived than the typical individual drawn at random from the population, then annuities priced to reflect the longevity of the annuitant population will not be actuarially fair for the typical individual in the population. Indeed, market-wide mortality tables published in the U.S. and the U.K. based on the mortality experience of voluntary annuitants suggest that the typical voluntary annuitant 65-year-old male is expected to live twenty percent longer than the typical 65-year-old male in the population.[1]

While these mortality patterns are consistent with adverse selection into the annuity market, they do not provide evidence on a relationship between mortality rates and the type of annuity purchased by annuitants. We are not aware of any published mortality tables that distinguish annuitants by the type of annuity policy that they purchase. Yet this relationship is the central prediction of asymmetric information models of insurance markets. The data in this paper allow a direct investigation of this relationship.

The U.K. annuity market provides a particularly rich setting for studying adverse selection in annuity markets. There are effectively two annuity markets in the United Kingdom. There is a compulsory annuity market in which individuals who have accumulated savings in tax-preferred retirement vehicles similar to 401(k)'s or IRA's in the U.S. are required to annuitize a large portion of their accumulated balance. There is also a voluntary annuity market in which individuals with accumulated savings may use these accumulated assets to purchase an annuity. Finkelstein and Poterba (1999) suggest that adverse selection, as measured by the average price of annuity contracts, is roughly half as great in the compulsory market as in the voluntary market.

A wide variety of annuity products are sold in both the compulsory and voluntary annuity markets. Annuitants in the compulsory market face virtually no restrictions on the types of annuities that they can purchase. This suggests that there is scope for selection among annuity products in both markets. We focus our analysis on annuities that pay a pre-determined payment stream. We also limit our analysis to annuities that insure a single life, as opposed to joint life annuities that continue to pay out as long as one of several annuitants remains alive.

We pay particular attention to three features of annuity policies that may serve as screening devices. One is the initial annual annuity payment. The payment amount is the analog of the payment in

[1] Finkelstein and Poterba (1999) and Murthi, Orszag, and Orszag (1999) present summary information and mortality comparisons for the U.K. annuity market. Brown, Mitchell, and Poterba (2000), and Poterba and Warshawsky (2000), present related information for the U.S. market.

the event of accident in the auto insurance market. It is a rough measure of the amount of insurance purchased by the annuitant.

A second annuity feature that we consider is the tilt, or degree of backloading. A more backloaded annuity is one with a payment profile that provides a greater share of payments in later years. The most common form of annuity is a nominal annuity, which pays out a constant nominal amount each period. In a world with positive expected inflation, the expected real payment stream from such an annuity is declining over time. An escalating annuity, in contrast, provides a payment stream that rises at a pre-specified nominal rate in each year. Annuities escalate at a nominal rate of anywhere from 3% to 8.5% in our data. These escalating annuities may be rising in real terms depending on the expected rate of inflation. There are also real (i.e. inflation-indexed) annuities, which pay out a constant real amount in each year. The payments from real annuities and from escalating annuities are both backloaded relative to those from nominal annuities.

A third feature that we focus on is whether the annuity may make payments to the annuitant's estate. Some annuities offer guarantee periods. The insurance company continues to make payments to the annuitant's estate for the duration of the guarantee period even if the annuitant dies before the guarantee period expires. Annuities with guarantee periods of one to fifteen years are present in our data sample, although in the compulsory market, regulations forbid the sale of policies with guarantee periods of more than ten years. "Capital protection" is another form of potential payments to the estate. If at the date of the annuitant's death the sum of nominal annuity payments to date is less than the premium paid for the annuity, a capital-protected annuity pays the difference to the estate as a lump sum.

All three of these features -- initial payment amount, backloading, and the presence of payments to the estate -- are potential screening devices for insurance companies with less information than their customers. Our empirical work explores whether the annuitants who choose these different features differ in their mortality experience and whether the pricing of these different features reflects any such mortality differences across features.

2.2 Market Equilibrium in the Annuity Market

To apply the equilibrium framework discussed in the last section to the annuity market, we represent an annuity contract by a vector $x = (a, \tau, g)$. The initial annuity payment is a, τ denotes the tilt or backloading of the annuity payout stream, and g denotes the guaranteed payout amount. We consider a two-period setting, in which an individual of risk type i has a probability of surviving in each period of p_i. Individuals purchase insurance contracts at the beginning of the first period. They then learn if they will survive through the first period, in which case they will receive their annuity payout. If they survive through the second period as well, they receive annuity payouts in both periods. For simplicity, we prohibit saving between the first and second period, and we assume a zero discount rate. We further assume that guaranteed payouts are paid to the annuitant's estate, and that they are only paid, if at all, in the period when the annuitant dies.

The expected utility of individual of risk type i who purchases policy x is therefore:

(1) $V_i(p_i, g_t, \tau, a) = (1 - p_i)U(g_1) + p_i U(a) + p_i^2 U(a(1+\tau)) + (1 - p_i)p_i U(g_2)$.

With probability $(1 - p_i)$, individual i dies in the first period and his estate receives the amount guaranteed if he dies in the first period, g_1. With probability p_i individual i survives the first period and receives initial annual annuity payment a. He then survives to the second period with probability p_i^2, in which case he then receives an annuity payment of $a(1+\tau)$, where τ denotes the tilt factor. Finally, with probability $(1 - \overline{p_i})p_i$ the individual survives the first period but dies in the second period. In this case, the annuitant's estate receives the second period guaranteed payment of g_2.

Screening on a given annuity characteristic requires that the Spence-Mirrlees property be satisfied for that characteristic. Equation (1) indicates that the marginal value of the initial annual annuity payment (a) is increasing in risk type p as is the marginal value of an increase in tilt (τ). Finally, the marginal value of a guaranteed payout (g) is decreasing in p. Thus all three features satisfy the Spence-Mirlees property and are therefore potential screening devices.

These results have an intuitive interpretation. The initial annual annuity payment is the analog of the payment in the event of a claim, or "quantity" in the Rothschild and Stiglitz (1976) model. The analysis of screening on this variable therefore follows the earlier discussion of screening on quantity choice in that model. An annuity with an upward tilt has more of its payments in later periods than an annuity with a flat payment profile. An annuitant with a longer life expectancy is more likely to be alive in later periods when the tilted annuity pays out more than the flat annuity. Such an annuitant therefore expects to gain more, at a given price, from a tilted annuity than does an annuitant with a shorter life expectancy. Similarly, an annuity that pays out more in the event of an early death, either with a guarantee period or with capital protection, is of greater value to an individual who is shorter lived than to one who is longer lived.

This model makes several predictions about the relative mortality patterns of different annuitants, if the annuity market is characterized by asymmetric information. For example, individuals who buy more backloaded annuities should be longer lived, conditional on observables, than those who buy less backloaded policies. In a similar fashion, those who buy annuities that make payments to the estate should be shorter-lived, and that those who buy annuities with larger initial annual payments should be longer lived, conditional on observables.

The foregoing model also makes predictions with respect to the pricing of annuity contracts. More backloaded annuities and annuities with larger initial annuity payments should be priced higher to reflect the fact that in equilibrium they are purchased by longer lived individuals than the buyers of other annuities. Similarly, annuities that make payments to the annuitant's estate should be priced lower to reflect the fact that in equilibrium they are purchased by shorter-lived individuals.

Table 1 summarizes our predictions with respect to pricing and mortality patterns in the annuity market when there is adverse selection. These predictions would not obtain in a setting with symmetric information. Consider the case of the degree of backloading of the annuity. With symmetric information, a longer-lived annuitant has no incentive to buy an annuity with backloaded payments. Whatever annuity he buys, the insurance company will adjust the price charged to reflect his individual mortality prospects.

Since the price adjusts, any preference for an annuity of a given tilt will be influenced only by discount rates, not by mortality prospects. When the annuitant has private information that he is likely to be long-lived, however, when he chooses to buy a particular annuity the price is not fully adjusted to take account of his mortality prospects. The key difference between the symmetric and asymmetric information scenarios is that under asymmetric information, the annuitant acts as a price taker, while under symmetric information, he does not. With asymmetric information, the price of a backloaded annuity will reflect the average mortality prospects of the annuitants who, in equilibrium, purchase backloaded annuities. This price will not reflect an <u>individual</u> annuitant's mortality prospects.

We test for asymmetric information in the annuity market by considering both mortality patterns and annuity prices. The fact that in practice, unlike in the stylized model, individuals differ on dimensions other than simply their risk type does not pose a problem for the interpretation of our empirical analysis. If product choice is driven not by private information about mortality but rather by (privately known) preference parameters such as discount rates, risk aversion, or bequest motives which are correlated with mortality, the effect of such private information is similar to the effect of private mortality information. Anything that is correlated with mortality and is known by the individual but not to the insurance company, even if the individual does not recognize its effect on his mortality, operates just like traditional asymmetric mortality information. A potential annuitant's wealth may operate in just this way: higher wealth annuitants are likely to face lower mortality risks, but insurance companies are not aware of the total wealth of their annuity buyers.

3. Data and Descriptive Statistics

Our data set consists of information on the complete set of both compulsory and voluntary immediate annuities sold by a large U.K. annuity company over a seventeen-year period ending on December 31, 1998. The first year in our sample, 1981, was the first year when the company sold both voluntary and compulsory annuities. At the end of the sample period, the firm was among the ten largest

U.K. sellers of new compulsory annuities.[2] The company's sales mix between compulsory and voluntary annuities shifts within our sample, with sales in the compulsory market growing more rapidly than sales in the voluntary market. This in part reflects the expansion in the late 1980s of the set of retirement savings plans that face compulsory annuitization requirements. In recent years, sales of voluntary annuities have been a small share of the company's total annuity sales.

Our data set includes almost everything that the insurance company knows about its annuitants. The only pieces of information that we did not receive from the company were the annuitant's address and his or her birthday. This information was suppressed to protect the confidentiality of the annuitants' identities. We did, however, obtain information on the annuitants' birth month. We have information on the type of policy purchased by the annuitant, and on the annuitant's day of death, if the annuitant died within the sample period. Death records are current through December 31, 1998. The insurance company collects very little information – only age and gender – about the personal characteristics of annuitants. In particular, it does not collect any information on the annuitants' wealth, income, education, occupation, or another other indicators of socioeconomic status. The information collection practices at the firm we study are typical for insurance companies selling annuities in the U.K.

We restrict our attention to single life annuities, annuities whose payments are based on the mortality of one individual, because the mortality patterns of the single insured lives on each policy provide a straightforward measure of ex-post risk type. Brown and Poterba (2000) note that it is considerably ~~more~~ complicated to analyze the joint mortality patterns of couples with joint life annuities than to analyze the mortality pattern for a single annuitant. Our final sample size is 42,054 annuity policies.

Table 2 provides an overview of the characteristics of the compulsory and voluntary annuities sold by the company that we analyze. The voluntary market accounts for about one tenth of the policies in our sample, and a somewhat higher fraction of premiums. The relative magnitudes of the voluntary

[2] Information on the market share of various U.K. insurance companies in the annuity market may be found at http://insider.econ.bbk.ac.uk/pensions/annuities/experiences/uk/q97s.htm

and compulsory market sales for our sample company are similar to those reported by the Association of British Insurers (various years) for the aggregate U.K. market. Differences between voluntary and compulsory annuitants in our data sample also appear typical of the U.K. market as a whole. Table 2 shows that voluntary annuitants are substantially older than compulsory annuitants. They are also more likely to be female. This is consistent with tabulations from the Family Resources Survey reported in Banks and Emmerson (1999). The product mix of annuities sold by our sample firm also matches the limited aggregate data that exist on the product mix for the U.K. market as a whole. Index-linked and escalating products together make up less than 10% of the voluntary or the compulsory market, with index-linked policies less than 5% in each market. Murthi, Orzsag, Orszag (1999) report a similar preponderance of nominal annuities in a data set that includes all insurance companies selling annuities in the United Kingdom.

The modal policy in the compulsory annuity market for 1997 and 1998, when annuity sales peaked for our sample firm, was a nominal annuity with a five-year guarantee period sold to a 65-year-old man. In the voluntary market, the modal policy in 1984 and 1985, which together account for more than one quarter of the voluntary annuity sales in the sample, was a nominal annuity with no guarantee, sold to a 74-year-old woman. These differences underscore the distinct character of the compulsory and voluntary annuity markets.

4. Annuitant Mortality and Annuity Product Choice

To test the predictions of asymmetric information models regarding the relationship between insurance product choice and risk type, we estimate hazard models for the mortality experience of annuitants who purchase different types of annuities. This section presents our empirical findings on this issue. In the next section, we investigate whether the variation in prices across different types of annuities accords with the estimated mortality differences.

4.1 A Hazard Model Framework for Studying Annuitant Mortality

We estimate mortality differences among different groups of annuitants using the discrete-time, semi-parametric, proportional hazard model used by Meyer (1990) and Han and Hausman (1990). Our duration measure is the length of time the annuitant lives after purchasing an annuity. We let $\lambda(t, x_i, \beta, \lambda_0)$ denote the hazard function, the probability that an annuitant with characteristics x_i dies t periods after purchasing the annuity, conditional on living until t.

The proportional hazard model assumes that $\lambda(t, x_i, \beta, \lambda_0)$ can be decomposed into a baseline hazard $\lambda_0(t)$ and a "shift factor" $\phi(x_i, \beta)$ as follows:

(2) $$\lambda(t, x_i, \beta, \lambda_0) = \phi(x_i, \beta)\lambda_0(t).$$

The baseline hazard, $\lambda_0(t)$, is the hazard when $\phi(\cdot) = 1$. $\phi(\cdot)$ represents the proportional shift in the hazard caused by the vector of explanatory variables x_i with unknown coefficients β. The proportional hazard model in (2) restricts the effects of the explanatory variables (x_i) to be duration-independent.

We adopt one of the common functional forms for $\phi(\cdot)$, $\phi(x_i; \beta) = \exp(x_i'\beta)$. The proportional hazard model is then written as:

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

Specification of the baseline hazard $\lambda_0(t)$ is a key issue in estimating models like (3). We model the baseline hazard non-parametrically as a step function. This allows us to avoid imposing any restrictive functional form assumptions on the baseline hazard. We have seventeen years of data and we allow for seventeen, annual, discrete, time periods. If we let $\delta_t = \int_0^t \lambda_0(s)ds$ denote the integrated baseline hazard, the proportional hazard model in (3) becomes:

(4) $$\lambda(t_i; x_i, \beta, \delta) = 1 - \exp\{\exp(x_i'\beta)(\delta_t - \delta_{t+1})\}.$$

Models in which the hazard function is given by (4) can be estimated by maximum likelihood with the log likelihood function given by

$$\ln(L)=\sum_{i=1}^{n}(1-c_i)\ln(\lambda(t_i;x_i,\beta,\alpha))-\int_0^t \lambda(s_i;x_i,\beta,\alpha)ds \tag{5}$$

where c is an indicator variable that equals 1 if the individual survives until the end of our sample period and 0 otherwise. Eighty-four percent of the voluntary annuitants in our sample, and 47 percent of the compulsory annuitants, survive until the end of our sample.

We also estimate hazard models that account for unobserved heterogeneity across individuals. Such heterogeneity can be modeling by assuming that there is an unobserved regressor that is a proportional shifter of the baseline hazard, just like the observed x_i regressors. In this case, we rewrite the proportional hazard function in (3) as

$$\lambda(t;x_i,\beta,\alpha,v_i)=v_i\exp(x_i'\beta)\lambda_0(t;\alpha) \tag{6}$$

where v_i summarizes the effect of omitted regressors for individual i. We assume a functional form for the probability density function for v_i and we then integrate it out of the conditional density function above to obtain a marginal density function. We assume that v_i follows a gamma distribution with unit mean and variance σ^2, which is independent of x_i and t. This assumption implies that the density of v_i is proportional to $v^{\sigma^{-2}-1}\exp(-v\sigma^{-2})$. This density is chosen mainly for its convenient mathematical form, for it admits an analytical expression for the marginal hazard function:

$$\lambda(t;x_i,\beta,\alpha)=\frac{\exp(x_i'\beta)\lambda_0(t;\alpha)}{1+\sigma^2\exp(x_i'\beta)\int_0^t \lambda_0(s)ds}. \tag{7}$$

When we allow for unobserved heterogeneity in the population, our results are based on maximizing (5) using the definition in (7). Allowing for heterogeneity requires that we estimate one more parameter than in the no-heterogeneity case; this parameter is σ^2.

We estimate hazard models for annuitant deaths as a function of all of the known characteristics of the annuitants and their annuity policies. Our hazard models do not include any measure of "marginal price" because all of the known characteristics of the annuity product and the annuitant, which are in the models, should completely determine both the policy premium and the marginal price. We estimate separate models for annuitants in the voluntary and compulsory markets. In all of the hazard models we include indicator variables for the age at purchase of the annuity (in five-year intervals), the year of purchase of the annuity, and the gender of the annuitant. We also include indicator variables for the frequency of annuity payments (monthly, termly, quarterly, semi-annually or annually).

We include two indicator variables to capture the degree of backloading of the annuity payment. One is an indicator for whether the annuity payments are indexed to inflation, and the other is an indicator for whether the annuity payments are escalating in nominal terms. Level nominal annuities are the omitted category. The theory described above suggests that individuals who buy index-linked or escalating annuities should be longer lived than those who buy nominal annuities. They should therefore have a lower hazard and so the predicted coefficients on the indicator variables for "index-linked" and for "escalating" in the hazard model are negative. We cannot predict the relative magnitude of these two coefficients, since the choice that a long-lived individual would make between an escalating and an indexed annuity will depend on the rate of expected inflation and the individual's discount rate.

We also include two indicator variables to capture payments to the estate. One is for whether the annuity is guaranteed, and the other is for whether the annuity is capital protected. An annuity cannot be both guaranteed and capital protected. The omitted category is an annuity that does not make any payments to the estate. The theory described above predicts that individuals who buy annuities with more payments to the estate will be shorter-lived (i.e. have a higher hazard rate) than individuals who buy annuities that do not make any payments to the estate. The predicted coefficients on the indicator variables for "guaranteed" and "capital protected" in the hazard model are therefore positive. There is no prior prediction concerning the relative longevity of guaranteed and capital protected annuitants, as there is no clear measure of which is relatively more attractive to someone with mortality that diverges from the population average. Thus there is no prediction about the relative size of their coefficients.

Finally, we include one other annuity product characteristic that satisfies the single crossing property and is therefore a potential screening device: the initial annual annuity payment. This "payment" variable corresponds to the amount that would be paid out in life insurance in the event of death or the amount that would be paid out from an automobile insurance policy in the event of an accident. The theory described above predicts that individuals who face greater payments in the event of a claim will be longer lived than individuals who f ace smaller payments in the event of a claim. The predicted coefficient on the "payment" variable in the hazard model is therefore negative. Table 1 summarizes these predictions in columns one and two.

4.2 Basic Results: Annuity Choice and Mortality Patterns

Table 3 presents estimates of the hazard model in equation (4). The first column presents estimates from the compulsory annuity market, while the second column presents estimates from the voluntary market. The results closely match our theoretical predictions of self-selection under asymmetric information. In both the compulsory and the voluntary market, there is strong evidence that individuals who buy more backloaded annuities are longer-lived. The coefficients on indicator variables for index-linked and escalating annuities are negative and statistically significant at the 1% level in both markets, indicating that all else equal, individuals who buy these annuities have a lower mortality hazard rate than individuals who buy nominal annuities.

There is also evidence strong evidence that voluntary annuitants who buy annuities that make payments to the estate in the event of an early death are shorter lived than individuals whose annuities do not make such payments. The coefficient on the indicator variable for guaranteed payouts is positive, indicating that individuals who buy guaranteed annuities have higher hazards (and hence are shorter-lived) than observationally similar individuals who buy non-guaranteed, non-capital protected annuities. This coefficient is statistically significantly different from zero at the 1% confidence level in the voluntary market. Additionally, the coefficient on the indicator variable for a capital-protected annuity is positive in the voluntary market, indicating that individuals who buy these annuities are shorter lived than observationally similar individuals who buy non-guaranteed, non-capital protected annuities. The coefficient on the indicator variable for capital protected annuities is not significantly different from zero in our basic specifications, but once we allow for unobserved heterogeneity, it is.

In the compulsory market, the results in Table 3 indicate that although individuals who buy guaranteed annuities are shorter lived than individuals who buy non-guaranteed annuities, the difference in hazard rates is not statistically significant. However, in results not reported here we grouped annuities into four sets by length of guarantee period: 1-4 years, 5 years, 10 years, and, only in the voluntary market, 15 years. We also grouped annuities by degree of escalation: 3% per annum, 5% per annum, and 8-8.5% per annum. In the compulsory market, we found that although guaranteed annuitants as a group are not significantly shorter lived than non-guaranteed annuitants, annuitants with 10 year guarantee periods, the longest in the compulsory market, are significantly shorter lived than annuitants with non-guaranteed annuities. We also found that the hazard increases monotonically with the length of the guarantee period and decreases monotonically with the degree of escalation. The differences between the hazards for 5 and 10 year guaranteed annuities, and for 3% and 5% escalating annuities, was significant at conventional confidence levels.

The findings in Table 3 suggest that, in both the voluntary and the compulsory market, longer-lived individuals select escalating annuities while shorter-lived individuals select guaranteed annuities. This choice is apparently based on private information about mortality prospects. In addition, in the compulsory annuity market individuals appear to base their choice of the degree of guarantee period or degree of escalation on private mortality information. We do not find evidence of this finer selection in the voluntary market.

Table 3 also provides some evidence that in the compulsory market, annuitants with a higher initial annual payment are longer lived than annuitants with a lower initial annual payment. This is the relationship predicted by the standard models of insurance market equilibrium with adverse selection, provided insurance companies use the amount of the initial payment as a screening variable. In the voluntary market, however, we find a positive coefficient on the initial payment variable. This is not consistent with stylized models of self-selection behavior. Our estimates are broadly consistent with market-wide data on the relationship between the size of the annuity and the mortality of the annuitant. For example, the Institute of Actuaries (1999a, 1999b) reports that while those who buy larger compulsory annuity policies tend to live longer than those who buy smaller policies, there is no discernable relationship between annuity size and mortality in the voluntary market.

It is important to note that in both the voluntary and the compulsory market, the estimated effect of the amount of initial payment is small compared to the magnitude of the other screening devices. To illustrate this, we translate the hazard model coefficients into survival probabilities for individuals with different types of annuities. Table 4 presents information on the probability that a 58-year-old annuity buyer will survive until age 75, conditional on different types of annuity purchases. The "baseline" survival probability, which is shown in the first row, is that for a 58 year old female annuitant who purchases a nominal, non-guaranteed, annual payment, single life annuity in 1981. This annuity is assumed to pay the company's average initial annuity payment, with different averages used in the voluntary and compulsory markets. Subsequent rows show the corresponding survival probabilities that emerge when we change single dimensions of the annuity policy.

The first column of Table 4 shows the results for the compulsory market. The second column of Table 4 shows the results for the voluntary market. Looking across these columns, our estimates indicate substantially lighter mortality in the voluntary market relative to the compulsory market. This is consistent with market-wide mortality tables for the compulsory and voluntary annuity markets (Institute of Actuaries, 1999a and 199b).

The rows in Table 4 suggest large mortality differences among certain classes of annuitants. For example, in the compulsory market, the baseline annuitant has a 64% chance of surviving to age 75. The switch to an index-linked but otherwise-identical annuity is associated with a survival probability of 83%, or a 30% increase in the survival probability compared to the baseline (nominal) annuity. In contrast, a one standard deviation increase in the initial annual annuity payment, from the mean, is only associated with a conditional survival probability of 66%, or a 3% increase in the 64% baseline survival probability.

In the voluntary market, we note that the mortality difference associated with guaranteed payments is somewhat smaller than that associated with backloading. For example, the probability that the baseline annuitant survives to 75 is 92% in the voluntary market. It is 97% for those in the index-linked annuity market and 98% for those in the escalating annuity market, but only 90% for those in the guaranteed annuity market. These results suggest that annuitants self-select to a greater degree based on the amount of backloading of payments than on the presence or absence of a payout guarantee.

To facilitate interpretation of these mortality differences, we can compare the mortality differences in Table 4 to the mortality differences between male and female annuitants. The mortality table for compulsory annuitants suggests that a 58-year-old male annuitant has a 42% chance of surviving until 75. A 58-year-old female annuitant, otherwise identical, has a 64% chance. A 58-year-old female annuitant who purchases an escalating annuity has an 86% survival probability to age 75. Therefore the increase in survival probability associated with purchasing an escalating annuity (34%) in the compulsory market is about the same as the increase associated with switching the gender of the annuitant from male to female. In the voluntary annuity market, the increase in survival probability associated with purchasing an escalating annuity (7%) is even larger than the increase associated with being female (2%).

The results in Tables 3 and 4 suggest substantial mortality selection based on the backloading of the annuity, some mortality selection based on payments to the estate, and very little – if any – screening of annuitants based on the initial annual payment of the annuity. The lack of selection on initial annual annuity payment supports results obtained by Cawley and Philipson (1999) for life insurance and by Chiappori and Salonie (2000) for auto insurance. These papers examine only one potential screening device. In one study it is the amount paid in the event of death, and in the other it is whether the individual has more than the legally-required minimum level of insurance. Both of these screening variables resemble the initial annual annuity payment in our screening model. While our results, like those in the other studies, suggest little screening on this variable, they suggest substantial screening on other margins of insurance policy choice.

There is good reason why insurance companies would not find it effective to screen annuity buyers on the amount of their initial annual payout. Recall that a screening device must be priced to reflect self-selection. Offering marginal prices that increase with the amount of the insurance payment in the event of an accident requires that insurers be able to monitor and to verify the total amount of insurance that each buyer has purchased. This can be difficult, since firms typically know only the amount of insurance that the buyer has purchased from their firm. If firms do not monitor total insurance purchases, then if the marginal price rises with the payment in the event of an accident, individuals who desire a large payment will be better off buying several small policies from different insurance companies rather than one large policy from a single firm.

Monitoring total insurance purchases may be particularly difficult in annuity markets. Unlike life insurance and auto insurance, which pay out only infrequently, if ever, annuities are almost certain to make some payments and may make many payments. For most other insurance products, the insurance company can stipulate that the contract is valid only if the insured has not purchased other insurance, and investigate compliance upon submission of a claim. Doing this with an annuity would require continuous monitoring, and could therefore be quite costly. This has led Abel ((1986) and Brugiavini (1993) to argue that it would be difficult to offer annuities with marginal prices that rose as a function of the annuity payment.[3]

Insurance companies do not face such problems in charging more for backloaded annuities, to reflect the higher average risk pool, or making the price of annuities with payments to the estate lower to reflect lower average risk (higher mortality). Individuals who desire a backloaded annuity cannot replicate such an annuity buy buying several (cheaper) immediate nominal annuities, and someone who wants a guaranteed annuity cannot create such annuity by purchasing multiple non-guaranteed policies.

4.3 Sensitivity Analysis

While our initial findings provide substantial support for the role of selection effects in the annuity market, we also developed several tests designed to explore these results – and their robustness – in greater detail. We examined whether the magnitude of selection effects among annuity products varied across groups of annuitants. For example, we tested whether mortality differences across annuity products were of the same magnitude for men and for women. We included a full set of interactions between all of the covariates and the gender of the annuitant in our basic hazard model, while constraining men and women to exhibit the same baseline hazard. We were unable to reject the null hypothesis that selection effects among annuity products were the same for men and for women.

We also subjected the results to a variety of robustness checks. We estimated the hazard models separately for men and women to allow the baseline hazard to be different for these two groups. We estimated the hazard models on several-year subsamples of our complete data set to check whether our results were contaminated by time trends in the annuity market that could be correlated with product characteristics and mortality. And we re-estimated the hazard model using as our duration measure total age rather than years since purchase of the annuity. Virtually all of the results discussed above were robust in sign and significance to these various checks. One notable finding, however, concerns the variable measuring the initial annuity payment. When we estimated separate hazard models for male and female annuitants, this variable was only statistically significant for men in the compulsory market and for women in the voluntary market, with the signs shown in Table 3. The gender-related difference in this coefficient bears further investigation.

4.4 Allowing for Unobserved Heterogeneity

In addition to the robustness checks described above, we also explored the effect of allowing for unobserved heterogeneity of the form described above. Table 5 presents our coefficient estimates for the covariates in the hazard function models, while Table 6 presents the coefficients that describe the baseline

---

[3]Chiappori (2000) specifically mentions annuities and life insurance as insurance markets where the non-exclusivity of the contract makes a rising marginal price impossible. Cawley and Philipson (1999) find no evidence of a rising

hazard function. The first and third columns of Table 5 reproduce the results in Table 3, while the second and fourth columns show estimates that allow for unobserved heterogeneity in the compulsory and voluntary market respectively. In both markets, the sign and significance of almost all of our substantive results are robust to allowing for unobserved heterogeneity. The one change of note is that when we allow for unobserved heterogeneity, the coefficient on the indicator variable for capital-protected annuities in the voluntary market becomes statistically significantly different from zero.

The parameter σ^2 is the estimate of the variance of the gamma distribution. Values of this parameter that are statistically significantly different from zero suggest the presence of unobserved heterogeneity. Our findings support the presence of unobserved heterogeneity in the voluntary market but not in the compulsory market. This may indicate that, because of fixed costs associated with offering each product class or other factors, insurance companies do not find it optimal in the voluntary annuity market to screen individuals into perfectly homogeneous product-risk categories. In the substantially larger compulsory market, our results are consistent with the view that individual annuitants have been screened based on their characteristics, and the characteristics of the policies that they buy, into homogeneous risk classes.

Allowing for unobserved heterogeneity affects our estimates of the selection effects in the voluntary market. Simply comparing the coefficients in columns (3) and (4) of Table 5 can be misleading, however, because these coefficients represent proportional shifts in the baseline hazard, and the estimated baseline hazard changes when we allow for unobserved heterogeneity, as can be seen in Table 6. To calibrate the difference in the estimated selection effects in the voluntary market with and without accounting for unobserved heterogeneity, Table 7 reports further calculations of how different covariates affect the probability of the baseline, 58 year old, female annuitant surviving until age 75. The first column of Table 7 replicates the results in the second column of Table 4 for the voluntary annuity market, modeled under the assumption that there is no unobserved heterogeneity. The second column

marginal price in the life insurance market.

reports similar estimates for the voluntary market based on the model that allows for unobserved heterogeneity.

Table 7 shows that accounting for unobserved heterogeneity increases our estimates of the selection effects in the voluntary market. For example, without accounting for unobserved heterogeneity, we estimate that buying an index-linked annuity is associated with a 5 percent increase, from 92% to 97%, in the probability that an annuitant survives until age 75. Once unobserved heterogeneity is accounted for, purchasing an index-linked annuity is associated with an increase in survival probability from a baseline of 91% to 98%, an 8 percent increase. Similarly, the reduction in the baseline survival probability associated with a guaranteed annuity is 2 percent when we do not account for unobserved heterogeneity, and 5 percent when we do. Both with and without unobserved heterogeneity, the mortality differences associated with backloading are larger than those associated with gender. The mortality differentials associated with payments to the estate are of the same or smaller size than those associated with gender. These findings underscore the importance of allowing for unobserved heterogeneity, and they suggest that studies of market equilibrium in insurance markets need to recognize the potential presence of such heterogeneity.

5. Pricing Differences Across Annuity Products

Our hazard model results suggest clear relationships between annuity product characteristics and annuitant mortality. We now consider the relationship between annuity product characteristics and annuity prices. Recall that if annuitants self-select among insurance products on the basis of private information about their mortality, then in equilibrium, prices on different annuity features should adjust to reflect feature-specific average mortality.

5.1 Calculating Annuity Prices and the "Money's Worth" Concept

The price of an annuity is not the premium that the annuitant pays when he or she purchases an annuity, but the differential between this premium and the expected present discounted value of annuity payouts. The price is therefore related to the annuity's "money's worth," which earlier studies such as

Mitchell, et al. (1999) have defined as the expected discounted present value of annuity payouts divided by the premium. The expected present discounted value is calculated using the mortality curve of a typical individual in the population, so an annuity that is actuarially fair for such an individual will have a money's worth of unity. Money's worth values may be less than one when there are administrative costs associated with the annuity policy, or if the mortality rate of the individuals buying an annuity is lower than that in the population at large. One minus the money's worth value equals the price of the annuity.

If the individuals who buy annuity product j are, on average, longer-lived than individuals who buy annuity product k, then insurance companies will reflect these mortality differences in their pricing. Assuming that an insurance company's costs and markup are the same across products, a given premium should therefore purchase a lower payment stream for product j than for product k since the buyers of product j will, on average, outlive the buyers of product k. From the standpoint of an individual facing a given mortality table, such as the population mortality table, product j should have a lower expected present discounted value of payments, and hence a lower money's worth, than product k. This implies that annuity product j should have a higher price than annuity product k.

The formula for the money's worth of a nominal, non-guaranteed annuity is:

$$\text{(8)} \qquad MW_{NOM} = \frac{\sum_{t=1}^{T} \frac{A * S_t}{\prod_{j=1}^{t}(1+i_j)}}{P}.$$

In this expression, A denotes the per-period payment from the nominal annuity, P denotes the initial premium payment, S_t denotes the probability that the annuitant survives until payment period t, and i_j denotes the expected nominal interest rate at time period j. The above formula is easily adjusted, as in Finkelstein and Poterba (1999), for the case of index-linked or escalating annuities, and for annuities that make payments to the annuitant's estate.

We calculate the money's worth value for each annuity in our data set using a common mortality 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 use the relevant rates for each individual who buys an annuity in our data set. In each case, we use mortality tables from the year in which the annuity is sold. For example, for a 65-year-old male who purchased an annuity in 1988, we use the 1988 population projections of his survival probability in each successive year. We use the zero-coupon yield curve of nominal U.K. Treasury securities on the day of annuity purchase to measure the term structure of nominal interest rates in equation (8). In calculating the money's worth of inflation-indexed annuities, we use data on the expected rate of inflation, as reported by the Bank of England, on the day of annuity purchase.

To explore how annuity prices are related to product characteristics, we regress the price of the annuity, defined as one minus the money's worth of the annuity, on characteristics of the annuity and the annuitant. The hedonic pricing equation, which we estimate by ordinary least squares, is:

$$\begin{aligned} \text{PRICE}_i &= \alpha + \beta_1 \text{INDEX}_i + \beta_2 \text{ESCALATING}_i + \beta_3 \text{GUARANTEED}_i + \beta_4 \text{CAPITAL}_i \\ &+ \beta_5 \text{PAYMENT}_i + \beta_6 \text{PAYMENT}^2{}_i + \beta_7 X_i + \varepsilon_i \end{aligned} \tag{9}$$

X consists of all known features of the annuitant and the annuity product that are not labeled separately. These include the age of the annuitant at time of purchase (in five year groupings), the gender of the annuitant, the year of purchase, and a series of indicator variables for the frequency of the annuity payments. All of these variables were used in our hazard model analysis above.

Equilibrium requires that the prices of various annuity product features are affected by the selection behavior of different mortality types. We therefore expect a positive coefficient on an indicator variable for whether the annuity is index-linked and on an indicator for whether the annuity is escalating. Because we found the mortality of the annuitants who buy these products to be lower than that of nominal annuitants, we expect that the annual payments offered on these products will be lower than those for nominal annuities. This reflects the fact that escalating and indexed annuitants are likely to live longer than the annuitants who buy plain nominal annuities. As a result of the lower annual payments, the money's worth calculated using a common mortality table will be lower, and the price of the annuity will be higher, for escalating or indexed products than for nominal ones. Similarly, we expect a negative coefficient on indicator variables for whether the annuity is guaranteed and whether it is capital protected. Table 1 summarizes these predictions in columns three and four.

The existence of pricing differentials such as those described in the previous paragraph are critical to the screening equilibrium described above. If insurance features that are purchased by higher risk individuals are not priced higher than those purchased by lower-risk individuals, then the incentive compatibility constraint of the low risk type would be violated. Since low risk individuals get less than their desired insurance, they would want to switch to the full-insurance, high-risk package if prices did not vary across packages.

5.2 Empirical Findings

Table 8 reports the hedonic regression results, which are generally supportive of the pricing patterns predicted by the screening equilibrium model. The estimated coefficients on the indicator variables for guaranteed and capital protected annuities indicate that annuities that make payments to the estate have significantly lower prices in both the compulsory and voluntary market than do annuities that do not make payments to the estate. This is consistent with the foregoing hazard model results that suggest that annuitants who purchase annuities that make payments to the estate are shorter lived (lower risk) than annuitants who purchase annuities that do not make payments to the estate. The results for backloading are similarly supportive. Index-linked annuities are priced significantly higher than nominal annuities in both the compulsory and voluntary market, reflecting the fact that the typical annuitant who purchases an index-linked annuity is longer lived (higher risk) than the typical annuitant who purchases a nominal annuity. In the compulsory market, but not the voluntary market, there is also evidence that, as predicted, escalating annuities have higher prices than nominal annuities.

In results not reported here, we also examined the relative prices of annuities with different degrees of escalation and different guarantee periods. In the compulsory market, consistent with the selection results reported above, the annuity price rises monotonically in the amount of escalation and falls monotonically with the length of guarantee period. The pricing differences among annuities of different guarantee lengths are statistically significantly different from zero, while those among annuities with different escalation rates are not. There is no clear pricing pattern for different degrees of escalation in the voluntary market, but the annuity price does fall monotonically with the length of guarantee period. The price difference between a five- and ten-year guaranteed annuity is statistically significantly different from zero, although the price differences between other guarantees are not.

The hedonic pricing equations also include the amount of the initial annuity payment. The negative coefficients on this variable in both markets are indicative of bulk discounts similar to those found by Cawley and Philipson (1999) in the U.S. life insurance market. The positive coefficient on the square of the initial annual payment is consistent with Rothschild and Stiglitz' (1976) prediction of a higher marginal price for larger quantities of insurance. However the magnitude of the coefficient is tiny, which suggests relatively little mortality screening or pricing response to this variable.

Chiappori and Salanie (2000) and Chiappori (2000) note that estimation of a firm's pricing policy is notoriously difficult. Fixed costs and economies of scale and scope can introduce non-linearities in a firm's pricing schedule. Such nonlinearities can be difficult to distinguish from the predictions of models of asymmetric information. Fortunately, we have direct information on the sample firm's firm's pricing policies in addition to the hedonic pricing relationship estimated above. We were told that for the insurance firm that we are studying, the pricing formula is as follows. Within a given class (guarantee period, tilt, frequency of payment, gender and age), if a £10,000 purchase price buys an annual payment of A, then a purchase price of P buys an annual annuity payment of (P*A)/10,000+([P-10,000]*f)/10,000) where f is the fixed policy fee. This formula applies both in the voluntary and the compulsory market.

To illustrate this pricing formula, consider an example. With £10,000 in 1999, a 65-year-old man could purchase an index-linked, single life, monthly payment, non-guaranteed annuity with an annual payment of £539.64. The fixed policy fee (f) was £18. This formula indicates the presence of bulk discounts for policies of less than £10,000 and it indicates a constant marginal price. Our estimates suggesting a very small bulk discount and an even smaller rising marginal price of increased annual payments are consistent with the pricing policy. More importantly, this linear pricing scheme is consistent with our findings that the insurance company does not use the initial annual payment to screen individuals into risk categories.

6. Conclusion and Discussion

This paper uses a unique data set of the annuitants at a large U.K. insurance company to provide new evidence of the presence of asymmetric information in insurance markets. Our results are consistent with the view that individuals self-select across annuity products based on private information about their mortality. We also find that, as would be expected in equilibrium, the pricing of different annuity products reflects the product-specific average mortality of annuity buyers. This evidence is robust to alternative specifications of the model and to various changes in the sample on which we estimate hazard models for mortality. It is also robust to the extending our basic hazard model to allow for unobserved heterogeneity across individuals.

We find the clearest evidence of self-selection on two attributes of annuity policies that are potential screening devices: the time profile (backloading) of the annuity, and whether it promises any payments to the estate of the annuitant. All else equal, annuitants who are longer lived select annuities with payment streams that are backloaded and that therefore pay out relatively more of their value in later years. Similarly, annuitants who are shorter lived select annuities that make payments to the annuitant's estate in the event of an early death. These selection effects are large. For example, in the compulsory annuity market, the mortality difference between otherwise identical individuals who purchase backloaded and non-backloaded annuities is similar in magnitude to the mortality difference between men and women. In the voluntary market, the mortality difference between backloaded and non-backloaded annuitants is several times larger than that between men and women. We also find evidence that relatively backloaded annuities are priced higher than less backloaded annuities. Annuities with payments to the estate are priced lower than annuities that do not make payments to the estate, which is what we expect if insurance companies take account of selection effects in setting their prices.

Our results broadly confirm the prediction of models of asymmetric information in insurance markets that emphasize a positive correlation between risk type and characteristics of the insurance policy that are of greater value to high risk individuals. Our findings on the differential mortality experience of annuitants who purchase different types of policies are consistent with individual having private information about their mortality experience, and acting on this information in their insurance purchases. These findings are complementary to survey-based studies, such as Hamermesh (1985) and Hurd and McGarry (1997) that have explored the extent to which individuals have informed, and plausible, views about their potential life expectancy. Such studies present evidence that individuals have private mortality information. Our results are not only consistent with individuals having private mortality information but also suggest that they use this information in making financial decisions about their later years.

While our results are consistent with the presence of adverse selection in the annuity market, two other potential explanations of our findings that warrant some discussion. The first is the possibility that our results are an artifact of the particular firm whose annuity sales we have analyzed. This small sample concern is difficult to address without detailed data from other insurance firms, and we do not have such data. However, our comparison of characteristics of the sample of annuities sold by the insurance company whose policies we study with characteristics of the aggregate U.K. market did not suggest obvious differences in the pool of policyholders. Furthermore, although we only have detailed pricing and mortality data from one insurance company, Finkelstein and Poterba (1999) report on the pricing of annuities with different guarantee periods and different degrees of backloading in a sample of the major U.K. annuity companies. The aggregate patterns match those found for the single firm considered in here.

A second concern is that our findings are not the result of adverse selection, but rather moral hazard. If individuals who purchase more insurance change their behavior in a way that results in higher claims against their insurance company, then we would observe the patterns we have documented in the relative mortality of different groups of annuitants. It is notoriously difficult to distinguish empirically between adverse selection and moral hazard. Even without such a distinction, our results provide evidence that is strongly suggestive of some type of asymmetric information in the annuity market. Our findings therefore stand in contrast to the recent claims of symmetric information in insurance markets, such as those by Cawley and Philipson (1999) and Chiappori and Salanie (2000). Moreover, the case for interpreting findings like ours as due to moral hazard is arguably weaker for annuity markets than for most other insurance markets.

For the moral hazard analysis to apply to annuities, the conversion of income to an annuity stream must affect the individual's mortality. It is difficult to think of a convincing mechanism for this. For example, since individuals with an insured income stream can consume more each period because they do not have to save for the possibility of outliving their resources, it could be argued that this increased consumption possibilities frontier improves their mortality. Philipson and Becker (1998) note that in principle the presence of annuity income may have effects on individual efforts to extend length of life, although they suspect that such effects are more likely to be important in developing than in currently developed nations. Even recognizing this potential effect, the importance of moral hazard in annuity markets is likely to be much smaller than that in many other insurance markets. Banks and Emmerson (1999) report that among both voluntary and compulsory annuitants in the U.K., annuity income represents less than one fifth of annual income. This reduces the likelihood that the form of the income stream is affecting mortality.

One of our most important results is that we do not find evidence of self-selection on the initial amount of annuity payment. This initial amount of annuity payment is analogous to measures of the accident-induced payment that are used in other empirical papers on asymmetric information in insurance markets. We find no evidence of a positive correlation between risk type and payment in the event of an accident in annuity markets. This is consistent with the evidence found in other insurance markets such as those studied by Cawley and Philipson (1999) and Chiappori and Salanie (2000). Consistent with the lack of selection on initial annuity payments, we find that the price of additional increments of initial annuity payment is linear. Cawley and Philipson (1999) report similar findings for the life insurance market. We do, however, find evidence of self-selection and pricing responses on other features of the annuity insurance contract. Our research therefore highlights the importance of paying careful attention to the detailed policy features of real-world insurance contracts when testing theoretical models of asymmetric information in insurance markets.

BIBLIOGRAPHY

Abel, Andrew B. 1986. Capital Accumulation and Uncertain Lifetimes with Adverse Selection. Econometrica 54 (September), pp.1079-1097.

Association of British Insurers. 1997. Insurance Statistics Year Book: 1986-1996. London: ABI.

Banks, James and Carl Emmerson. 1999. UK Annuitants. London: Institute for Fiscal Studies. http://www.ifs.org.uk/publications/briefnotes.shtml

Brown, Jeffrey, Olivia Mitchell, and James Poterba. 2000. Mortality Risk, Inflation Risk, and Annuity Products. NBER Working Paper 7812.

Brown, Jeffrey and James Poterba. 2000. Joint Life Annuities and Annuity Demand by Married Couples. Journal of Risk and Insurance (forthcoming).

Brugiavini, A. 1993. Uncertainty Resolution and the Timing of Annuity Purchases. Journal of Public Economics, 50, pp.31-62.

Cawley, John and Tomas Philipson. 1999. An Empirical Examination of Information Barriers to Trade in Insurance. American Economic Review 89(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 (forthcoming).

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

Cordon, James and Igal Hendel. 1999. Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey. Unpublished Mimeo, Princeton University Department of Economics..

Finkelstein, Amy and James Poterba. 1999. Selection Effects in the Market for Individual Annuities: New Evidence from the United Kingdom. NBER Working Paper 7168.

Friedman, Benjamin, and Mark Warshawsky. 1988. Annuity prices and saving behavior in the United States. In Z. Bodie, J. Shoven, and D. Wise, eds. Pensions in the US Economy. Chicago: University of Chicago Press.

Friedman, Benjamin, and Mark Warshawsky. 1990. The cost of annuities: Implications for saving behavior and bequests. Quarterly Journal of Economics 105 (February), pp.135-154.

Hamermesh, Daniel. 1985. Expectations, Life Expectancy, and Economic Behavior. Quarterly Journal of Economics 100(May), 389-408.

Han, Aaron and Jerry Hausman. 1990. Flexible Parametric Estimation of Duration and Competing Risk Models. Journal of Applied Econometrics. 5:1-28.

Hurd, Michael and Kathleen McGarry. 1997. The Predictive Validity of Subjective Probabilities of Survival. NBER Working Paper 6193.

Institute of Actuaries and Faculty of Actuaries, Continuous Mortality Investigation Committee. 1999a. Continuous Mortality Investigation Reports, Number 16.

Institute of Actuaries and Faculty of Actuaries, Continuous Mortality Investigation Committee. 1999b. Continuous Mortality Investigation Reports, Number 17.

Meyer, Bruce. 1990. Unemployment Insurance and Unemployment Spells. Econometrica, 58(4): 757-782.

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(December), pp.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.

Philipson, Tomas and Gary Becker. 1998. Old-Age Longevity and Mortality-Contingent Claims. Journal of Political Economy, 106: 551-573.

Poterba, James and Mark Warshawsky. 2000. The Costs of Annuitizing Retirement Payments from Individual Accounts. In John Shoven, ed. Administrative Costs and Social Security Privatization. Chicago: University of Chicago Press, 173-200.

Pueltz, Robert and Arthur Snow. 1994. Evidence on Adverse Selection: Equilibrium Signaling and Cross-Subsidization in the Insurance Market. Journal of Political Economy, 102: 236-257.

Rothschild, Michael and Joseph Stiglitz. 1976. An Essay on the Economics of Imperfect Information. Quarterly Journal of Economics 90 (November), pp.629-649.

Smart, Michael. 2000. Competitive Insurance Markets with Two Unobservables. International Economic Review 41 (February), 153-169.

Yaari, Menahem. 1965. Uncertain Lifetimes, Life Insurance, and the Theory of the Consumer. Review of Economic Studies 32: 137-150.

TABLE 1:
PREDICTIONS OF ASYMMETRIC INFORMATION MODELS FOR ANNUITY MARKETS

ANNUITY FEATURE / POTENTIAL SCREENING DEVICE	MORTALITY PATTERN (1) Mortality Difference	(2) Coefficient in Hazard Model	PRICING RESPONSE (3) Pricing Difference	(4) Coefficient in Hedonic Price Equation
BACKLOADED ANNUITIES	Backloaded annuity buyers will be longer lived than nominal annuity buyers	Negative	Backloaded annuities will have higher relative prices	Positive
PAYMENTS TO THE ESTATE	Annuitants with annuities that make payments to the estate will be shorter lived	Positive	Annuities with guarantees will have lower relative prices	Negative
INITIAL ANNUAL ANNUITY PAYMENT	Annuitants with larger initial annual payments will be longer lived	Negative	Rising Marginal Price	Positive coefficient on square of variable

TABLE 2: OVERVIEW OF THE COMPULSORY AND VOLUNTARY MARKET

	COMPULSORY MARKET	VOLUNTARY MARKET
Number of policies	38,362	3,692
Number (%) of annuitants who are deceased.	6,311 (16.5%)	1,944 (52.7%)
Number (%) of annuitants who are male	29,681 (77.4%)	1,272 (34.5)
Average age at purchase	63.2	76.4
	Backloaded Annuities	
Number (%) of policies that are index-linked	428 (1.3%)	66 (3.5%)
Number (%) of policies that are escalating in nominal terms	1,492 (3.9%)	175 (4.7%)
	Payments to Estate	
Number (%) of policies that are guaranteed	28,424 (74.1%)	872 (23.6%)
Number (%) of policies that are capital protected	0	843 (22.8%)
	Initial Annual Annuity Payments	
Average Initial Payment (£)	1,151	4,773
Median Initial Payment (£)	627	3,136
Standard Deviation of Initial Payment (£)	1,929	5,229
Average Premium (£)	10,523	25,603

Notes: All monetary figures in the paper are in December 1998 pounds. The first index-linked policy was sold in February 1985; therefore percentage of policies index-linked refers to percentage of policies sold since that date.

TABLE 3: SELECTION EFFECTS AND ANNUITY PRODUCT CHARACTERISTICS

	COMPULSORY MARKET	VOLUNTARY MARKET
Backloaded Annuities		
INDEX-LINKED	-0.839***	-0.894**
	(0.217)	(0.358)
ESCALATING	-1.085***	-1.497***
	(0.113)	(0.253)
Payments to Estate		
GUARANTEED	0.019	0.216***
	(0.029)	(0.060)
CAPITAL-PROTECTED	------------	0.056
		(0.051)
Initial Annuity Payment		
PAYMENT (in £100s)	-0.003***	0.001**
	(0.0006)	(0.0004)
N	38,362	3,692
Number of deaths in sample	6,311	1,944

Note: All estimates are from Han-Hausman discrete-time, semi-parametric proportional hazard models. These are estimated using 17 annual discrete time intervals. Baseline hazard parameters are not reported. Indicator variables for five-year intervals for age at purchase, and for year of purchase, are included in all regressions, along with indicator variables for gender and for the frequency of payments. Standard errors are in parentheses. The omitted category for the "tilt" dummies (index-linked and escalating) is nominal annuities. The omitted category for the guarantee feature dummies (guaranteed and capital protected) is not guaranteed and not capital protected. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE 4: INTERPRETING SELECTION EFFECTS: PROBABILITY OF SURVIVING UNTIL AGE 75 FOR DIFFERENT TYPES OF ANNUITY BUYERS

	COMPULSORY MARKET	VOLUNTARY MARKET
Baseline	0.641	0.916
Index-linked annuity	0.827	0.965
Escalating Annuity	0.862	0.981
Guaranteed Annuity	0.635	0.897
Capital Protected Annuity	---------	0.912
Initial Annual Annuity Payment One Standard Deviation Above the Mean	0.657	0.912
Male Annuitant	0.423	0.893

Note: Baseline survival probability is for a 58 year old female annuitant with a nominal, non-guaranteed, annual payment, single life annuity purchased in 1981 and paying the mean initial annual annuity payment for the compulsory or voluntary market. Other rows reflect the survival probability of an annuitant who has the baseline characteristics and the baseline annuity except for the change noted in the first column.

TABLE 5:
SELECTION EFFECTS AND ANNUITY PRODUCT CHARACTERISTICS: ESTIMATES WITH AND WITHOUT GAMMA HETEROGENEITY

	COMPULSORY MARKET		VOLUNTARY MARKET	
	No Heterogeneity Factor (1)	Allowing For Heterogeneity (2)	No Heterogeneity Factor (3)	Allowing For Heterogeneity (4)
IINDEX-LINKED	-0.839*** (0.217)	-0.874*** (0.222)	-0.894** (0.358)	-1.671*** (0.582)
ESCALATING	-1.085*** (0.113)	-1.114*** (0.122)	-1.497*** (0.253)	-2.442*** (0.352)
MALE ANNUITANT	0.640*** (0.039)	0.663*** (0.045)	0.252*** (0.051)	0.463*** (0.094)
GUARANTEED	0.019 (0.029)	0.024 (0.031)	0.216*** (0.060)	0.423*** (0.105)
CAPITAL PROTECTED	----------	------------	0.056 (0.051)	0.389*** (.105)
SEMI-ANNUAL PAYMENTS	-0.020 (0.048)	-0.025 (0.051)	0.121 (0.111)	0.296 (0.184)
TERMLY PAYMENTS	-0.424 (0.428)	-0.257 (0.436)	--------------	--------------
QUARTERLY PAYMENTS	0.004 (0.038)	0.0009 (0.040)	0.165 (0.108)	0.406** (0.180)
MONTHLY PAYMENTS	-0.006 (0.034)	-0.008 (0.036)	0.256*** (0.099)	0.544*** (0.162)
PAYMENT (in £100s)	-0.003*** (0.0006)	-0.003*** (0.0006)	0.001** (0.0004)	0.001 (0.0009)
σ^2	-------------	0.253 (0.241)	-----------	2.021*** (0.216)

Notes: All estimates are from Han-Hausman discrete-time, semi-parametric proportional hazard models. These are estimated using 17 annual discrete time intervals. Table 6 reports the baseline hazard parameters. Indicator variables for age at purchase (in five-year intervals) and year of purchase are included in all regressions. Standard errors are shown in parentheses. The omitted category for the "tilt" dummies (index-linked and escalating) is nominal annuities, for the guarantee feature dummies (guaranteed and capital protected) is not guaranteed and not capital protected, for the frequency of payment dummies is annual payments. σ^2 is the estimated variance of the gamma heterogeneity. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. Sample size for first two (last two) columns is 38,362 (3,692), with 6,311 (1,944) annuitant deaths within sample.

TABLE 6: BASELINE HAZARD FUNCTION WITH AND WITHOUT GAMMA HETEROGENEITY

	COMPULSORY MARKET		VOLUNTARY MARKET	
Years	No Heterogeneity Factor (1)	Allowing for Heterogeneity (2)	No Heterogeneity Factor (3)	Allowing for Heterogeneity (4)
1	0.0091	0.0090	0.0011	0.0001
	(0.0008)	(0.0008)	(0.0003)	(0.0001)
2	0.0104	0.0103	0.0018	0.0002
	(0.0009)	(0.0009)	(0.0005)	(0.0001)
3	0.0110	0.0110	0.0022	0.0003
	(0.0010)	(0.0010)	(0.0006)	(0.0002)
4	0.0144	0.0144	0.0024	0.0004
	(0.0012)	(0.0013)	(0.0006)	(0.0002)
5	0.0146	0.0148	0.0030	0.0007
	(0.0012)	(0.0013)	(0.0008)	(0.0004)
6	0.0166	0.0169	0.0027	0.0008
	(0.0014)	(0.0015)	(0.0007)	(0.0004)
7	0.0182	0.0187	0.0032	0.0011
	(0.0015)	(0.0016)	(0.0008)	(0.0006)
8	0.0195	0.0202	0.0036	0.0015
	(0.0016)	(0.0018)	(0.0009)	(0.0008)
9	0.0213	0.0224	0.0032	0.0016
	(0.0018)	(0.0021)	(0.0008)	(0.0008)
10	0.0224	0.0237	0.0038	0.0022
	(0.0019)	(0.0023)	(0.0009)	(0.0011)
11	0.0256	0.0275	0.0042	0.0030
	(0.0022)	(0.0028)	(0.0010)	(0.0014)
12	0.0305	0.0331	0.0039	0.0032
	(0.0027)	(0.0037)	(0.0009)	(0.0015)
13	0.0334	0.0368	0.0039	0.0037
	(0.0029)	(0.0044)	(0.0009)	(0.0017)
14	0.0345	0.0387	0.0063	0.0068
	(0.0033)	(0.0051)	(0.0016)	(0.0032)
15	0.0391	0.0446	0.0056	0.0069
	(0.0040)	(0.0069)	(0.0014)	(0.0033)
16	0.0519	0.0603	0.0060	0.0086
	(0.0061)	(0.0105)	(0.0016)	(0.0041)
17	0.0800	0.0956	0.0260	0.0452
	(0.0125)	(0.0215)	(0.0089)	(0.0249)

Note: The estimates are for the baseline hazard parameters from the Han-Hausman hazard models fit to data on the voluntary and compulsory market. Coefficients for the covariates in these models are shown in Table 5. Standard errors (computed using the delta method) are in parentheses. The baseline hazard represents the hazard of a 55-60 year old female annuitant who purchased a nominal, non-guaranteed, annual payment, single life annuity in 1981.

TABLE 7:
THE IMPACT OF MODELLING HETEROGENEITY ON THE ESTIMATED PROBABILITY OF VOLUNTARY ANNUITANT SURVIVING UNTIL AGE 75

	NO HETEROGENEITY FACTOR	ALOWING FOR HETEROGENEITY
Baseline	0.916	0.912
Index-linked annuity	0.965	0.983
Escalating Annuity	0.981	0.992
Guaranteed Annuity	0.897	0.868
Capital Protected	0.912	0.873
Initial annual annuity payment one standard deviation above the mean	0.912	0.908
Male Annuitant	0.893	0.863

Note: Baseline survival probability is for a 58 year old female annuitant with a nominal, non-guaranteed, annual payment, single life annuity purchased in 1981 and paying the mean initial annual annuity payment for the voluntary market. Other rows reflect the probability of survival until age 75 for an annuitant who has the baseline characteristics and the baseline annuity, except for the change noted in the first column.

TABLE 8: HEDONIC MODEL OF ANNUITY PRICING: THE EFFECT OF PRODUCT CHARACTERISTICS

	COMPULSORY MARKET	VOLUNTARY MARKET
INDEX-LINKED	0.096***	0.046***
	(0.004)	(0.007)
ESCALATING	0.004*	-0.032***
	(0.002)	(0.005)
GUARANTEED	-0.014***	-0.037***
	(0.0009)	(0.002)
CAPITAL PROTECTED		-0.081***
	----------------	(0.002)
PAYMENT (in £100s)	-0.002***	-0.0003***
	(0.00003)	(0.00003)
PAYMENT Squared (in £100s)	2.56e-06***	5.31e-07***
	(9.24e-08)	(8.78e-08)
Sample Size	38,362	3,692

Note: Coefficient estimates are from linear regressions of price, defined as 1-Money's Worth, on product characteristics. Regressions include indicator variables for five-year intervals of age at time of annuity purchase, year of annuity purchase, gender of annuitant, and frequency of annuity payments. Standard errors are in parentheses. The omitted category for the "tilt" dummies (index-linked and escalating) is nominal annuities. The omitted category for the guarantee feature dummies (guaranteed and capital protected) is not guaranteed and not capital protected. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.