**RESEARCH STATEMENT**

**Amy Finkelstein**
Assistant Professor of Economics, MIT
www.nber.org/~afinkels
March 2007

The provision of social insurance – particularly medical insurance and pension annuities – is a key function of modern government in developed economies. Social insurance is also a substantial share of public expenditures in these countries, and its share is expected to grow even larger as the baby boomers age and medical expenditures continue to grow.

One of the economic rationales for social insurance stems from a rich theoretical literature on the potential for asymmetric information to impair the efficient operation of private insurance markets. This literature – which dates back to the work of Akerlof (1970) and Rothschild and Stiglitz (1976) and was a major basis for the Nobel Prize awarded in 2001 – also emphasizes that when would-be buyers of insurance have private information about their risk type, there may well be scope for welfare-improving government intervention.

Despite the long-standing theoretical interest in and policy relevance of asymmetric information in insurance markets, empirical work in this area is only in its infancy. My research focuses on the development and implementation of empirical methods that investigate the role of asymmetric information in insurance markets and that make it possible to describe the consequences of government intervention in these markets. In the particular context of health insurance markets, my work has also emphasized that the consequences of government intervention include not only its static effects on insurance coverage, but also important dynamic effects on technological change.

I have studied these economic questions in the context of markets and policies that are important in their own right. As one example, my work on asymmetric information has led to a series of related papers examining the reasons for the very limited size of the private long-term care insurance market in the United States. As another example, my work on the dynamic consequences of health policy has included several papers that shed light on the impact of the introduction of the U.S. Medicare program on the health care sector, and on the program's elderly beneficiaries.

In conducting my research, I do not rely on a particular method or approach. Rather, I endeavor to select the methods and data that are best suited to the question at hand. As a result, my work employs a disparate set of empirical methods, including the use of data to test particular equilibrium theories, structural estimation, and calibrated dynamic stochastic life cycle models, as well as reduced form analysis. In addition, most of my papers identify and assemble new data sets, often from historical archives or from companies' proprietary records.

The remainder of this statement is divided into two sections. The first reviews my work on detecting asymmetric information in insurance markets and determining its effect on market equilibrium. The second section discusses my work on the dynamic effects of government intervention in health insurance markets on technological change in the health care sector.

## I. Asymmetric information in insurance markets

### *A. Detecting asymmetric information in insurance markets*

The first generation of empirical work on asymmetric information in insurance markets focused, quite naturally, on the question of how to design empirical tests for the existence of asymmetric information.

My initial work in this area built on the insight of Chiappori and Salanie (2000) that if either adverse selection or moral hazard exists in an insurance market, it should generate a positive correlation in equilibrium between the amount of insurance individuals have and their ex-post risk experience. For example, in models of adverse selection, would-be buyers of insurance have private information about their likelihood of incurring an insurance claim. As a result, individuals who know they are higher risk than the insurance company expects are more likely to choose to buy more comprehensive insurance, thus driving up the equilibrium price of insurance and producing a market in which those who have more insurance have higher expected probability of the insured event.

In a paper with James Poterba in the *Journal of Political Economy* on "Adverse Selection in Insurance Markets" [1], we extend this insight to note that asymmetric information can affect not just the quantity of insurance but also the form of the insurance contract. Using proprietary data on pension annuity purchases and ex-post annuitant mortality experience from a large insurance company in the United Kingdom, we examine several features of the annuity contract – such as an increasing time profile of annuity payments – which, at a given price, should be more valuable to individuals with private information that they are higher risk than the insurance company expects. We find that these features are in fact disproportionately chosen by higher risk individuals, and that the equilibrium price of these features is correspondingly higher. However, we find no evidence of adverse selection on the amount of payment in the event the insured risk occurs. This measure of the quantity of insurance had been the only feature of insurance contracts modeled in the theoretical literature, and the near-exclusive focus of previous empirical work on detecting asymmetric information. Our work therefore highlights the importance of considering the detailed features of real-world insurance contracts when testing theoretical models of asymmetric information.

My more recent work on detecting asymmetric information has uncovered another important gap between the seminal theoretical work on market equilibrium with asymmetric information – in which there is at most only one source of private information – and the reality of actual insurance markets – in which consumers may have private information about risk preferences as well as risk type. In a paper in the *American Economic Review* with Kathleen McGarry on "Multiple Dimensions of Private Information" [2] we show that if individuals have private information about their risk preferences as well as their risk type, the equilibrium need not involve the positive correlation between the amount of insurance coverage and risk type that the standard uni-dimensional models of asymmetric information predict, even when individuals act based on private information about risk type and thereby distort the market equilibrium.

Our empirical evidence is from the U.S. market for long-term care insurance. There, we find that two types of individuals select into the market: individuals with private information that they are higher risk than the insurance company would expect, and individuals with private information that they have strong preferences for insurance (for example, they are highly risk averse). While the former are higher risk than the insurance company would expect, the latter are lower risk. As a result, these two forms of private information operate in offsetting directions to produce an equilibrium in which the insured are not higher risk than the general population, despite the presence of asymmetric information about risk type that may impair market efficiency.

This work raises important theoretical questions about how to model insurance markets with multiple dimensions of private information. It also presents an empirical challenge as it invalidates the standard "positive correlation" empirical test for asymmetric information. In a follow up paper with James Poterba on "Testing for Adverse Selection with 'Unused Observables'" [3], we develop an alternative approach to testing for asymmetric information that is robust to the potential existence of preference heterogeneity in insurance demand.

*B. The efficiency consequences of asymmetric information in insurance markets and the scope for welfare improving government intervention*

The importance of asymmetric information in insurance markets – at least from a public finance perspective – stems from the potential for government intervention to improve upon the private market allocation. My most recent research in this area therefore moves beyond the question of how to detect whether private information exists, to develop an approach to quantifying the efficiency costs of asymmetric information once it has been detected, and to estimate the likely welfare consequences of possible government interventions.

In a paper with Liran Einav and Paul Schrimpf on "The Welfare Cost of Asymmetric Information" [4] we show that the potential presence of heterogeneous preferences as well as heterogeneous risk types – as documented in my earlier work with McGarry [2] – substantially complicates any such efficiency analysis. We build a simple theoretical model that demonstrates a key insight: even when asymmetric information is known to exist, the reduced form equilibrium relationship between insurance coverage and risk occurrence (i.e. how "adversely selected" the market appears) does not permit even qualitative inferences about the relative efficiency cost of asymmetric information in different insurance markets. In other words, in a market with private information about risk type that looks very adversely selected, we show that there is an underlying joint distribution of unobserved preferences and risk type that can generate this equilibrium with very high efficiency costs, and another underlying distribution that can generate it with very low efficiency costs. Likewise, in a market with private information about risk type that does not look adversely selected at all (as for example, I found in my analysis of the long tem care insurance market [2]), different underlying joint distributions of unobserved preferences and risk type can produce this equilibrium with either very high or very low efficiency costs.

This theoretical insight motivates our development of a structural method that uses insurance company data to recover the joint distribution of (unobserved) risk type and risk preferences of consumers. We show that we can identify these dimensions of unobserved heterogeneity and the relationship between them from variation in the menu of options different individuals face, combined with data of individuals' choices from their menu and their ex-post claims experience. Once we recover these primitives, we can compare individual welfare under the observed equilibrium allocations to their welfare under counterfactual allocations.

We apply our method to the U.K. pension annuity market in which, in a series of previous papers with James Poterba ([5], [1], [3]) I had documented the existence of adverse selection. Einav, Schrimpf and I estimate the efficiency costs of this adverse selection to be about 2.5% of the premiums in this market. We also compare welfare in the asymmetric information equilibrium to what could be achieved under the canonical solution to the problem of adverse selection: mandatory social insurance. We emphasize that once it is recognized that individuals may differ both in their preferences and in their (privately known) risk type, mandates in fact involve a trade-off between the allocative inefficiency produced by adverse selection and the allocative inefficiency produced by a uniform insurance policy. Whether a particular mandate will be welfare improving therefore becomes an empirical question. Indeed, we estimate that some mandatory annuity policies would have substantial welfare improvements over the asymmetric information equilibrium, while others would entail considerable welfare losses.

*C. Understanding the equilibrium in the private long-term care insurance market in the United States*

My work on estimating the efficiency costs of asymmetric information is motivated by the simple observation that detecting asymmetric information is only a first step. Evidence that market failures exist is necessary, but not sufficient, for them to be an important source of market inefficiency and a rationale for government intervention. In the private long-term care insurance market in the United States, this

insight has prompted a series of research papers designed to investigate whether market failures might account for the market's limited size.

Long-term care expenditures represent one of the largest uninsured financial risks facing the elderly in the United States, and the magnitude of this risk is projected to grow as the baby boomers age and medical expenditures continue to rise. In a series of papers with different co-authors I have documented the existence of several different market failures in the private long-term care insurance market, including asymmetric information [2], and dynamic contracting problems [6]. As a result of these and other market imperfections, Jeffrey Brown and I estimate that prices in this market are quite high relative to the expected benefits from private insurance policies [7].

In a separate paper on "The Interaction of Public and Private Insurance" [8], Brown and I then investigate whether these high prices – and by extension the underlying market failures that produce them – can explain the small size of the private long-term care insurance market. To do so, we develop and calibrate a dynamic, stochastic, life cycle model of a risk averse individual's demand for private insurance. This model suggests that even if we were to eliminate all potential market imperfections so that, contrary to fact, actuarially fairly priced comprehensive private insurance policies were available, about two-thirds of the elderly wealth distribution would still not want to buy this insurance. Once again, this research emphasizes the importance of moving beyond the detection of market imperfections to analysis of their contribution to market equilibrium.

The explanation for our findings of limited private insurance demand even in the (counterfactual) absence of market failures lies in the role of Medicaid, the public insurance program that serves as a payer-of-last-resort for long-term care expenditures. We estimate that Medicaid has a substantial crowd-out effect on demand for private insurance because it imposes a large implicit tax on private policies. Yet because of its means-tested nature, it serves as a very incomplete form of coverage for all but the poorest of individuals, by essentially imposing a deductible of virtually all of one's financial assets. Our analysis therefore suggests that a public insurance system can substantially crowd out private insurance, even when the public insurance itself provides for only very limited reductions in risk exposure. It also suggests that fundamental Medicaid reform is necessary for the private insurance market to expand considerably.

## II. The general equilibrium consequences of government intervention in health insurance markets

There is a long tradition in public finance of estimating the impact of public policies on insurance markets in a static environment, in which the distribution of risk is assumed constant. I have written several papers in this tradition (including [8], [9], [10], [11], and [12]). However, the other major strand of my research has emphasized that in health insurance markets, public policy can also have important dynamic effects on incentives to develop and adopt new medical technologies, and thus the underlying distribution of medical expenditure risk that individuals wish to insure.

### *A. The determinants of technological change in medicine*

There is widespread consensus among health economists that technological change in medicine is the driving factor behind the dramatic increase in real per capita medical expenditures over the last half century. From a theoretical perspective, it has long been recognized that economic factors including relative factor prices and expected market size can influence technological change in general (see e.g. Hicks, 1932 and Schmookler, 1966) and in the health care sector in particular (see e.g. Weisbrod, 1991). Yet there has been relatively little empirical work exploring the determinants of technological change in medicine.

Several of my research papers suggest an important role for health policies – and the economic incentives they create – in affecting the development and diffusion of new medical technologies. In a paper published in the *Quarterly Journal of Economics* on "Static and Dynamic Effects of Health Policy" [13], I examine the dynamic consequences of policies designed to increase vaccination rates against particular diseases. Standard analysis of such policies would have considered only the static consequences of the policies' explicit goal of increasing vaccination rates with the existing vaccine technologies. Using a new data source I uncovered on pharmaceutical clinical trials, I show that such policies, which increased the expected market size for particular vaccines, also encouraged the development of improved versions of these vaccines. Moreover, I estimate that the dynamic welfare consequences of the induced innovation may be larger than the static welfare benefits from increasing the vaccination rate using the existing technology. These findings underscore the importance of including induced technological change in any economic analysis of health policy.

In related work with Daron Acemoglu on "Input and Technology Choices in Regulated Industries" [14], I have also investigated the role of relative factor prices in affecting hospitals' adoption of new medical technologies. We find that a reform in Medicare reimbursement to hospitals which increased the relative price of labor encouraged hospitals to adopt capital-embodied medical technologies as a substitute for labor. The work with Acemoglu emphasizes, more generally, the role of regulation as an important instrument of government policy, a theme I have explored in other papers ([10] and [15]) as well. This work on regulation complements my work on the more traditional public finance instruments of taxation (see [9] and [16]) and spending (see [8], [12], [17] [18]).

*B. The general equilibrium effects of health insurance: evidence from the introduction of Medicare*

Standard economic analysis of the impact of health insurance on health spending has focused on how health insurance increases demand for existing health technologies. My work on induced technological change suggests that large scale changes in health insurance may, be increasing the expected market size for new medical technologies, also affect the supply side of the health care sector and therefore have disproportionately larger effects on total spending.

I explore this in a paper in the *Quarterly Journal of Economics* on "The Aggregate Effects of Health Insurance" [17], where I investigate the impact of the introduction of Medicare in 1965 on the health care sector. The introduction of Medicare, which provides virtually universal public health insurance coverage for the elderly in the United States, was the single largest change in health insurance coverage in U.S. history. It therefore offers a rare opportunity to study the effect of a market-wide change in health insurance.

Yet prior to my work, there had been virtually no empirical analysis of the impact of Medicare's introduction. This hole in the literature stemmed from two obstacles: a lack of detailed data from the time period, and the empirical difficulty in analyzing the impact of a uniform national program. To surmount the first of these challenges, I tracked down and coded up previously unknown data in the historical archives that contain detailed, annual, hospital-level measures of expenditures, utilization, and technology adoption. To separate the impact of Medicare from any underlying secular changes in the health care sector that would have occurred anyway during this time period, I drew on the insight that the, prior to Medicare, the rate of private insurance among the elderly differed greatly across areas of the country; as a result, the increase in insurance coverage associated with Medicare also varied substantially across different areas of the country.

Using this approach, I estimate that the general equilibrium effects of market-wide changes in health insurance on health spending are about six times larger than what partial equilibrium evidence from individual-level changes in health insurance would have predicted. Estimates of the impact of health

insurance on health spending from the Rand Health Insurance Experiment, one of the largest randomized, individual-level social experiments ever conducted in the United States, had previously suggested that the spread of health insurance was not an important cause of the rise in health spending. In contrast to the partial equilibrium estimates from the Rand analysis, my general equilibrium estimates from the impact of Medicare imply that the overall spread of health insurance, and the public policies that encouraged it, may be able to explain half of the rise in real per capita health spending over the last half century.

One important explanation for these disproportionately larger general equilibrium estimates is that, unlike an individual's change in health insurance, market-wide changes in health insurance, and hence in aggregate demand, can create incentives for health care providers to incur the fixed costs of adopting new practice styles. For example, in the case of Medicare, I find that the introduction of Medicare was associated with substantial new hospital entry and with hospitals' increasing the rate at which they adopted new medical technologies. The evidence of Medicare's impact on technology adoption also suggests that while my findings challenge the view that health insurance has played only a limited role in the dramatic growth in medical expenditures, they are at the same time consistent with the conventional wisdom that technological change in medicine has played a primary role in the growth of health spending. A key insight from this work – as well as my other work on the dynamic effects of health policy – is that it is important to get inside the "black box" of technological change to understand its underlying determinants, and particularly the impact of health policy.

Given the evidence that Medicare dramatically increased health care spending, Robin McKnight and I investigate the benefits from the introduction of Medicare and compare them to the costs of the program. In "What Did Medicare Do (And Was It Worth It)?" [18], we find no evidence of improvements in elderly mortality from Medicare's introduction. However, we do find large increases in financial security among the elderly due to Medicare's substantial reduction in the risk of large out-of-pocket medical expenses. We estimate that, even in the apparent absence of any health benefits (and there may well be improvements in morbidity or long-run mortality benefits that we cannot estimate), the reductions in risk exposure alone may be sufficient to cover about one-third of the costs of Medicare. These findings underscore the importance of considering direct insurance benefits from public health insurance, in addition to any indirect benefits from improvements in health. This straightforward theoretical point is often overlooked in empirical work on the impact of health insurance, which tends to focus almost exclusively on the potential health benefits from health insurance.

**References**


[1]. "Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market," 2004, *Journal of Political Economy* 112(1) Part 1: 183-208 (with James Poterba).

[2]. "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market," 2006, *American Economic Review* September, 96(4): 938 – 958 (with Kathleen McGarry).

[3]. "Testing for Adverse Selection with 'Unused Observables'." NBER Working Paper 12112. March 2006 (with James Poterba).

[4]. "The Welfare Cost of Asymmetric Information: Evidence from the U.K. Annuity Market" (with Liran Einav and Paul Schrimpf).

[5]. "Selection Effects in the Market for Individual Annuities: New Evidence from the United Kingdom," 2002, *The Economic Journal*, 112(476): 28-50 (with James Poterba).

[6]. "Dynamic Inefficiencies in Insurance Markets: Evidence from long-term care insurance," 2005, *American Economic Review Papers and Proceedings*, May, 95: 224-228 (with Kathleen McGarry and Amir Sufi.)

[7]. "Why is the Market for Long-Term Care Insurance So Small?" forthcoming *Journal of Public Economics* (with Jeffrey Brown).

[8]. "The Interaction of Public and Private Insurance: Medicaid and the Long-Term Care Insurance Market." NBER Working Paper 10989. December 2004. Revised and resubmitted to the *American Economic Review* (with Jeffrey Brown).

[9]. "The Effect of Tax Subsidies to Employer-Provided Supplementary Health Insurance: Evidence from Canada," 2002, *Journal of Public Economics*, 84(3): 305-340.

[10]. "Minimum Standards, Insurance Regulation and Adverse Selection: Evidence from the Medigap Market," 2004, *Journal of Public Economics* 88(12): 2515-2547.

[11]. "The Interaction of Partial Public Insurance Programs and Residual Private Insurance Markets: Evidence from the U.S. Medicare Program," 2004, *Journal of Health Economics*, 23(1): 1-24.

[12]. "Medicaid Crowd-Out of Private Long-Term Care Insurance Demand: Evidence from the Health and Retirement Survey." Forthcoming. *Tax Policy and the Economy* (with Jeffrey Brown and Norma Coe).

[13]. "Static and Dynamic Effects of Health Policy: Evidence from the Vaccine Industry," 2004, *Quarterly Journal of Economics* 119(2): 527-564.

[14]. "Input and Technology Choices in Regulated Industries: Evidence from the Health Care Sector."NBER Working Paper 12254. May 2006 (with Daron Acemoglu).

[15]. "Redistribution by Insurance Market Regulation: Analyzing a Ban on Gender-Based Retirement Annuities." NBER Working Paper 12205. May 2006 (with James Poterba and Casey Rothschild).

[16]. "*EZ-Tax*: Tax Salience and Tax Rates." February 2007. NBER Working Paper 12924.

[17]. "The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare" 2007, *Quarterly Journal of Economics* 122 (1): 1-37.

[18]. "What Did Medicare Do (And Was It Worth It)?" NBER Working Paper 11609. September 2005 (with Robin McKnight).

*Other cited works*

Akerlof, George. 1970. "The Market for 'Lemons'; Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 84(3), 488-500.

Chiappori, Pierre-Andre and Bernard Salanie. 2000. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy*, 108 (1), pp. 56-78.

Hicks, John. 1932. *The Theory of Wages*, Macmillan, London.

Rothschild, Michael and Joseph Stiglitz. 1976. "An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics*, 90 (4) pp. 629-649.

Schmookler, Jacob. 1966. *Invention and Economic Growth.* Cambridge, Harvard University Press.

Weisbrod, Burton. 1991. "The health care quadrilemma: An Essay on Technological Change, Insurance, Quality of Care, and Cost Containment" *Journal of Economic Literature* Vol XXIX (June): pp. 523-552