

Impact of Retiree Health Plans on Faculty Retirement Decisions[†]

John Rust
University Maryland and NBER

Prepared for: TIAA-CREF Institute Conference on
“Recruitment, Retention and Retirement:
The Three R’s of Higher Education in the 21st Century”
New York City, April 1-2, 2004

[†]Preliminary draft, not for quotation. This research was supported in part by a grant from the University of Michigan’s Economic Research Initiative on the Uninsured. Direct comments to the author at Department of Economics, University of Maryland, College Park, MD 20742 e-mail: jrust@gemini.econ.umd.edu.

Abstract

This paper considers the effect of retiree health insurance coverage on retirement decisions of academics, and attempts to answer two key questions of interest in this conference volume: 1) how do retiree health insurance plans affect retirement decisions? and 2) if retiree health plans are cut back or eliminated, by how much will faculty delay their retirements? I argue that the risk of uninsured health care costs is a major consideration affecting faculty retirement decisions, and via computer simulations, I show that naive attempts to cut costs by eliminating retiree health insurance can end up increasing rather than saving total costs of compensation, since it induces faculty to delay retirement and thus increases the cost of wages, health insurance and other fringe benefits over longer employment durations. These issues are of increasing relevance, since the rapid rise health care costs and health insurance premiums is putting increasing pressure on employers to limit health insurance coverage, or in some cases to eliminate retiree health plans altogether. Thus, colleges and universities face a difficult cost/benefit tradeoff in designing retirement and compensation packages. This paper develops a prototype “life cycle model” of a faculty member’s retirement decision that accounts for the incentives created by their pension plan provisions (including defined benefit and defined contribution plans), the effect of private and public disability insurance plans, and the provisions of their health insurance coverage — both while employed as well as their retiree health insurance coverage, if any, during retirement. I use this model to illustrate the important impact of changes in retiree health insurance coverage on retirement decisions, in experiments that predict by how much academics will delay retirement if retiree health insurance is eliminated, or the level of coverage is reduced. I use the model to perform cost/benefit calculations from the standpoint of an academic institution as to whether it is more cost effective to provide more generous retiree health insurance coverage or limit these benefits in various ways. I show that there are “compromise packages” such as continuing retiree health insurance until Medicare eligibility age (currently age 65), that are significantly less costly than the policy of providing retiree health insurance regardless of age, or eliminating retiree health insurance benefits altogether.

Keywords: retiree health insurance, retirement incentives, life cycle models.

JEL classification: H5

1 Introduction

Tenured faculty members at academic institutions in the U.S. (i.e. colleges and universities) are privileged to hold one of the most secure job contracts available. With essentially no risk of being fired, in a job that offers great working conditions, minimal levels of physical exertion, a high degree of flexibility and personal freedom, in addition to fringe benefit packages that are typically far more generous than those in the “private sector,” it should not be surprising that most academics have little incentive to retire. Indeed, many academics privately regard their jobs as fully paid “virtual retirement packages.”

At the same time, colleges and universities are under increasing financial pressure as a result of recent cutbacks in government funding and student financial aid, the decline in the stock market, increased competition for students and research dollars, and rapidly rising costs — including the costs of providing “fringe” benefits such as health care to their employees. These pressures are especially acute at some of the smaller liberal arts colleges that do not have large endowments and Federal research support, and whose revenues are therefore highly dependent on tuition. Even though tuition at the major public universities have increased rapidly in recent years in order to offset reductions in state and federal funding, the large percentage increases were made from a much smaller base tuition level, so that public school tuition is still generally far cheaper than tuition at private schools. As a result, many small to medium-sized private liberal arts colleges have much less room to raise their already very high tuition rates without risking significant declines in enrollment. In economic terminology, many private colleges may be at a fairly elastic point on their tuition “demand curves” so that further increases in tuition could start to reduce overall revenues. At these schools, the primary focus is on how to reduce costs given that the possibilities for raising additional revenue (apart from soliciting donations from corporate sponsors and alumni to build endowments) has become far less promising in recent years.

Since faculty salaries and fringe benefits are typically a college or university’s biggest single expense, it is not surprising that the focus of most cost-cutting efforts is on either reducing the level or at least constraining the rate of growth in salaries and benefits. In particular, Sylvester Scheiber’s survey in this volume shows that the cost of health insurance benefits is rising at an unsustainable rate of over 10% per year in recent years, and that “if played out over the next

decade or so, could drive total health costs equal to something between 15 and 25 percent of cash wages paid to workers.” (p. 16).

Just as many non-academic businesses have chosen to cut health care benefits and replace full-time workers with lower paid “temporary workers” for whom they do not have to pay health insurance and other types of benefits (or in the most extreme cases, “outsourcing” or “downsizing” by eliminating the jobs entirely), colleges and universities are under similar pressures to reduce costs by cutting benefits and replacing tenured faculty with lower paid untenured assistant professors, instructors, and lecturers, (in many cases saving the costs of the generous fringe benefit packages offered to their “permanent” employees) and in the extreme, to leave the slots created by retirements unfilled for the indefinite future.

The nature of the tenure contract imposes a huge “rigidity” on colleges and universities that most regular firms do not face: tenured faculty cannot generally be “fired” but must be persuaded to leave (i.e. retire) voluntarily. Further, anti age discrimination legislation such as the 1967 Age Discrimination in Employment Act and its 1986 ammendment that banned the use of mandatory retirement, further reduces a college or university’s flexibility in how they can structure financial incentives to encourage older tenured faculty to retire. In general, it is becoming increasingly difficult to use actuarially unfair payout formulas in “defined benefit” pension schemes to create incentives for early retirement, and due to tax advantages, more and more academic institutions are switching to “defined contribution” pension plans that are basically neutral in terms of the incentives they create for early versus late retirement.

Even though a university or college may be bound by its tenure contract, there is typically nothing beyond ordinary competitive pressure and a sense of moral obligation to provide generous fringe benefit packages to its staff. In addition to typically generous health insurance and retiree health insurance packages, many academic institutions also provide private supplemental disability benefits and other types of fringe benefits such as college tuition subsidies for children of academic staff. If the tenure contract does not obligate the academic institution to provide these costly fringe benefits, then it should not be surprising that many of these costly “perks” associated academic job are at risk for cutbacks or even complete elimination.

A 2002 survey of 387 colleges and universities by the College and University Professional Association (CUPA) for Human Resources found that 67 percent provide some type of retiree

health insurance coverage – a rate of provision that is significantly higher than for non-academic employers. However the financial pressures facing colleges combined with the unprecedented double digit rates of increase in health care costs in the U.S. in recent years (with typical health insurance premiums rising at a rate of between 10 to 15% per year), has forced many academic institutions to reconsider their retirement packages, and shift an increasing share of the premiums and costs of health care onto faculty retirees. The CUPA survey found that for two thirds of the colleges surveyed, health care insurance rates increased by more than 10% and one in seven colleges reported rates of increased of more than 20% from the previous year.

In the face of rapidly escalating health care costs, over 80% of U.S. companies have been either had to pass on higher premiums to their employees or reduce the generosity of their coverage of health care costs — generally either by increasing deductibles and co-insurance rates. Indeed, in 2002 33% of U.S. companies reduced or eliminated their health care coverage entirely. Higher educational institutions have dealt with these problems in much the same way: “In an effort to preserve coverage for employees without busting their budgets, colleges and universities in recent years have raised premiums, deductibles, and the co-payments for drugs and doctor visits. Some institutions have questioned how long they can continue to take care of retirees’ health care.” (*Chronicle of Higher Education*, February 27, 2004).

These problems also extend to retiree health coverage, and are not just limited to health insurance coverage for current employees. A 2003 survey by the Henry J. Kaiser foundation found that 10 percent of firms with 1,000 or more employees had cut or eliminated their retiree health benefits. Although I am not aware of comparable data for educational institutions, it is reasonable to predict that increasing number of academic institutions will be forced to reduce or even eliminate their retiree health benefits as part of a series of cost-cutting measures necessary to keep them afloat in increasingly difficult times.

There is no easy solution to the financial pressures facing colleges and universities, and naive efforts to cut costs by reducing or eliminating retiree health insurance benefits could backfire and end up costing more than continuing with the *status quo*. When a college eliminates or substantially reduces the generosity of its retiree health plan, it also significantly reduces the incentive for its existing faculty to retire. It may be tempting for a college or university to cut retiree health insurance benefits under the hypothesis that retirees are a less powerful group that is

“unseen and unheard” compared to more visible political opposition to reducing health insurance benefits that could be generated by existing employees. However if a college attempts to save costs via the “least painful” strategy of cutting benefits for a group that it perceives to be the least politically troublesome — i.e. retirees — existing faculty are sure to take note of the change in policy and will have even less incentive to retire than they currently have.

Thus, colleges and universities have to balance the cost savings from cutting back on their retiree health plans with the potential cost increases due to delayed retirement by its existing tenured faculty. In general if a college wants to encourage a tenured faculty member to retire, they must be keenly aware of the fact that their faculty members already have a very secure and comfortable position. The risk of uninsured health care costs may well be the biggest single risk facing their tenured faculty members. Since a tenured faculty also has the default option of not retiring, a cost-cutting measure that has the effect of increasing the level of risk facing a retiring faculty member during his or her retirement years must be compensated in some way (i.e. by a higher retirement benefit) to maintain incentives for the faculty to retire. However to the extent that faculty are risk averse and do not have good alternative private health insurance options prior to their eligibility for Medicare at age 65, it could end up costing a college more in terms of higher retirement benefits to compensate its faculty for the increased risks of health care costs than it saves from the cost-cutting measures. Alternatively, if the college simply cuts retiree health care benefits without an offsetting increase in retirement benefits, then the likely outcome is that its tenured faculty will delay their retirements, and this alternative could also be more costly than the *status quo* (or even more costly than incurring greater costs to increase the level of generosity of the college’s retiree health insurance package). *In general, when a firm has employees that are sufficiently risk averse, it is typically cheaper to pay them with a compensation package that includes insurance against their most pressing risks, than to forgo the insurance and pay them entirely in cash and force them to “self-insure” their risks.*

The goal of this paper is provide a conceptual framework that could prove useful in the design of faculty compensation, retirement, health insurance and fringe benefit packages. To my knowledge no previous study has provided such a framework. Instead, the design of retirement “schemes” chosen by academic institutions and private firms seems to be largely the result of *ad hoc* experimentation, and recommendations of specialized consulting companies, whose advice seems

guided more by intuition and previous experience than by an extensive body of scientific research. While I would not discount the importance of intuition, common sense, and practical experience in the design of compensation and retirement packages, I believe that economic methods — both theoretical models and econometric studies — could significantly improve our understanding of the factors influencing faculty retirement decisions and as a by-product, enable higher educational institutions to design more effective compensation, fringe benefit, and retirement plans.

In this paper I adapt a version of the classical “life cycle model” in economics to the specifics of the academic job contract, accounting for the job security and employment incentives provided by tenure and the types of pension plans and health care coverage that are typically offered at universities and colleges in the U.S. The life cycle model assumes that decision makers are rational, forward-looking planners who are risk averse, and choose how much to work and how much to consume and when to retire (i.e. when to give up their tenured teaching appointment and start collecting their pension benefits) in order to maximize the expected discounted utility over their remaining lifetime. Unlike many other occupations, academics are modeled as having a much stronger attachment to their work (i.e. a lower disutility of effort). This is a strong part of the explanation, independent of financial incentives, why academics are more reluctant to retire compared to individuals in most other industries and occupations. Even though they have tenure, faculty members do face risks, including risk of health problems leading to disability, the risk of uninsured or partially insured health care costs wiping out accumulated retirement savings, and uncertainty about their future earnings, both while employed at their academic job and subsequent earnings at a possible post-retirement job.

The life cycle model predicts that faculty will adopt an optimal consumption, labor supply, and pension/Social Security acceptance and accumulation/decumulation strategy to maximize their life-time utility, and will undertake an optimal strategy to deal with the risks they face. One risk-reducing strategy is to remain employed on their tenured job. This provides a steady stream of income and access to the university’s employee health care plan. However retiring and giving up their tenure is essentially an irreversible decision. The model is able to reflect differential expectations about post-retirement job opportunities: while the most productive older faculty members may expect to get attractive job offers and possibly even tenured positions at other colleges or universities, many older faculty members may have significant concerns about the

possibility of being subjected to implicit age discrimination that could make it very difficult for them to get another equivalent and well compensated position elsewhere after giving up tenure at their current institution. Furthermore, faculty who expect to work in consulting or in an untenured teaching or research position may expect that their post-retirement earnings opportunities will be significantly less attractive, involving lower salaries or wage rates, more year to year variability in earnings, and potentially inferior working conditions. The model can account for these various beliefs and scenarios, as well as the concerns of faculty who may find it increasingly difficult to carry a full time load as they get older. These individuals would welcome a partial or “phased retirement” option that allows them to gradually reduce their hours of work but without fully giving up the security of their tenured position. In certain cases, the ability to gradually phase out of full time work into full retirement is preferable to making a sudden discontinuous transition from full time employment to full retirement.

The other main strategy that faculty have for dealing with risks is the level of precautionary savings they undertake. Some savings will be done voluntarily through optional tax-sheltered retirement plans such as the 403-B plan offered to educational institutions. Other saving will be done by the employer via an “employer match” to the faculty member’s defined contribution account. The life cycle model also accomodates situations where the university’s main retirement plan is of the defined benefit type, or possibly a combination of DB and DC plans. The accumulation of retirement savings serves to augment the faculty member’s DB and DC pension receipts as well as their Social Security receipts, as well as providing an important “buffer stock” that can be used to pay for health care costs that are not fully covered by their retiree health insurance plan or Medicare insurance plans (if they are over age 65).

A major advantage of the life cycle model is that it generates predictions of how a rational individual’s behavior would change in response to changes in the rules governing the payment of their pensions, and the details of their health insurance coverage. Thus, the model allows me to undertake a series of “computational experiments” to assess how on a person’s retirement decision, consumption and savings, and overall level of welfare would be affected by increasing or decreasing the level of generosity in their health insurance, either before or after retirement. Confirming previous research by Rust and Phelan (1997) (who modeled retirement decisions of blue collar workers), I find that even for moderate levels of risk aversion, the concern about low-

probability, but potentially “catastrophic” out of pocket health care costs has a significant impact on retirement decisions. If a university attempts to save costs by eliminating or substantially increasing the co-insurance rates in its retiree health plan without making an equivalent reductions in the degree of generosity of its health insurance plan for current employees, the increased level of risk this imposes on retirement years relative to working years lead to significant postponements in retirement ages. There are ranges of parameter values for which such a policy is more costly than the *status quo*, i.e. maintaining retiree health insurance benefits even though they may be relatively costly.

Even though the life cycle model developed here is just an illustration that is in many respects an oversimplified caricature of the full complexities of actual retirement decisions, the model is able to capture the main tradeoffs facing faculty members at the approach their end of their academic careers. Even though the model has so far only been crudely calibrated to match several of the “stylized facts” of employment and retirement behavior of academics, with access to better data, it is possible that a more elaborate version of this model could provide a reasonable approximation to actual behavior.

Section 2 reviews existing literature on empirical studies of academic retirements, structural life cycle models of retirement decisions, consumption, savings, and the impact of health insurance, and the literature on “mechanism design” and principal-agent problems which have influenced the modeling approach used in this paper. We refer the reader to the chapter by Sylvester Scheiber (2004) (this volume) for a summary of some of the typical features of employer-provided and retiree health insurance coverage offered by colleges and universities. These features motivate the model I develop in Section 3, which develops a stylized version of the life cycle model that I will use to predict the effect of changes in employee and retiree health plan provisions on the retirement decisions of tenured faculty at colleges and universities in the U.S. Section 4 undertakes several computational experiments designed to illustrate the behavioral and welfare effects of various changes to pension benefits and health insurance coverage, including the impact on retirement ages and the total expected discounted cost of providing these benefits. Section 5 offers some concluding remarks.

2 Previous Work on Health Insurance and Retirement

In many respects, the preferred approach for analyzing the impact of health insurance coverage on retirement would be to conduct a controlled, randomized experiment that compares the retirement behavior of individuals who are randomly assigned to a “treatment group” that will be subject to a new type of health insurance coverage to individuals in a “control group” who retain the type of health insurance coverage they currently have.

To my knowledge, there have been no such controlled experiments, at least in order to evaluate the specific question of the effect of health insurance on retirement behavior.¹ Part of the reason why controlled experiments have not been done is that they are generally very costly and time-consuming. In order to properly evaluate a proposed *permanent* change in health insurance coverage, individuals in the control and treatment group should be observed for at least several years before and after they retire. Further, detailed panel data should be collected in order to measure a range of behavioral outcomes and decisions: in addition to the person’s date of retirement, we are potentially interested in their consumption and savings decisions and implied levels of wealth accumulation, their decisions regarding when to apply for pension and Social Security benefits and engage in post-retirement work, and their utilization of health care services both before and after retirement. The costs, complexities, and long duration of carrying out a well-designed experiment is probably the major reason why no such studies have undertaken.

However even if such a study was done, it is not clear whether it would have more than very limited applicability to other situations. In general, a single experiment can only measure the behavioral effects of a single, specific “treatment” policy. If a firm is interested in evaluating a range of different possible policy changes, some of which depends on parameters (e.g. coinsurance rates) that could be continuously varied, in principle separate experiments would have to be conducted in order to evaluate the effects of each specific proposed policy change. In general, there is no way to extrapolate the results of one particular experiment to predict the outcome of some different policy change. The problem of limited ability to extrapolate and thus limited

¹ The famous RAND Health Insurance Experiment in the 1970s (Newhouse, *et. al.* 1993) is the only controlled, randomized experiment that I am aware of that attempted to measure the effects of changes in health insurance coverage on individual behavior. However the RAND experiment was primarily interested in studying the effect of coinsurance rates and other parameters on levels of utilization and quality of health care and did not focus on the issue of how health insurance coverage affects retirement decisions.

applicability is even greater when one considers whether it would be valid to use the results of an experiment conducted at one firm to predict the impact of implementing the same policy change at another firm. There may be a host of difficult to measure institutional features that have a big effect on worker decisions (e.g. the nature of the job, working conditions, and the overall “work environment”) that it could be hazardous to try to use the results of an experiment conducted at one firm to predict the impact of the same policy change at some other firm unless all of the relevant “background variables” are known to be very similar at both firms.

In the absence of controlled experiments, there are two other approaches that can be used to analyze and predict the impact of changes in health insurance policies on retirement decisions. One approach is to use “reduced form” statistical studies that use data on actual retirement decisions as a basis for extrapolation and prediction. The other approach is to use a “structural model” that is based on an empirically testable and estimable mathematical model of a worker’s behavior that can be used to predict how the worker will respond to a wide variety of different hypothetical policies. Both types of approaches have been used to study and predict the impact of changes of health insurance on retirement for *general* (i.e. blue collar) workers, however I am not aware of any previous reduced-form or structural analyses of retirement decisions that are specifically focused on academics (e.g. where the “worker” is a tenured faculty member).

Reduced-form studies focus on situations where there have been significant changes in health insurance rules or other aspects of compensation policies or work rules. These studies use the policy changes as an “instrument” to assess how workers respond to the changes, basically by comparing the behavior of workers at the firm in question before and after the policy change, or else by comparing the behavior of workers at the firm after the policy change with similar workers at firms whose policies were like the policy used by the firm in question prior to the policy change.

A great deal has been learned from such studies. I would point to work by Ashenfelter and Card (2001) who studied the impact of the elimination of the age 70 mandatory retirement rule in 1994 and Allen, Clark, and Ghent’s (2003) study of the “phased retirement” program at the University of North Carolina as providing very valuable information on how faculty members actually react to changes in retirement rules. Other more recent studies include John Pencavel’s analysis of the impact of faculty buyout plans in the University of California system in this volume. The insights gained from these studies can help university administrators design more effective

retirement schemes. In particular, the Allen, Clark, and Ghent study suggests that the phased retirement program adopted by UNC “seems to be a plus from both employee and employer perspectives” (p. 21) and has not resulted in problems of “adverse selection” that tend to drain the best and brightest faculty while leaving behind the less productive.

However many of the limitations relating to the applicability and generalizability of results of controlled experiments apply to the results of reduced-form studies. There is also a major challenge in finding suitable exogenous “policy instruments” that enable us to infer the impact of a policy change and having adequate controls for the effects of other variables that affect retirement decisions. For example, in the simple case of comparing the behavior of workers at a given firm before and after a change in the firm’s health insurance policy, there is a difficult issue of how to control for possible changes in the economic environment (due to a recession or boom, for example) that could also have an important independent impact on retirement decisions. In a randomized experiment one can control for these other effects by comparing the behavior of the control and treatment groups before and after the policy change, since both groups are operating in the same economic environment before and after the policy change.

This paper adopts a structural approach to predicting the impact of changes in health insurance (and other parameters of academic compensation packages) on retirement decisions as well as on a host of other decisions that are made before and after retirement (i.e. consumption and savings decisions, decisions of whether to apply for disability benefits, Social Security benefits, and whether or not to engage in post-retirement work, and so forth). The main limitation of the current study is that I do not have access to any actual observations on retirement behavior that I could use as a basis for statistical estimation and testing of my model. Instead, the best I can do in this analysis is to formulate a model and to “calibrate” and simulate it in order to argue that the model has the potential to provide detailed and relatively accurate predictions of how real academics would behave, and could thus be a useful tool for universities to use to help them make intelligent decisions about their compensation and fringe benefit packages, with particular focus on health insurance coverage.

As partial compensation for lack of actual empirical results in this paper, I will briefly survey a number of previous papers that did empirically estimate and test structural models of retirement behavior. The first such study of which I am aware is the pioneering work by Gotz and McCall

(1984) who developed a “dynamic retention model” that predicted how various promotion and compensation policies affected retirement decisions of Air Force officers. The dynamic retention model is very similar in spirit to the life cycle model used in this paper. by carefully modeling dynamics and expectations, Gotz and McCall were able to show that their model could provide accurate predictions of retirement decisions under the *status quo*, significantly outperforming several *ad hoc* models known within the Air Force as the PVCOL and ACOL models. They showed that the dynamic retention model generated intuitively plausible predictions of a wide range of hypothetical changes in compensation and promotion policies.

Lumsdaine Stock and Wise (1991) (LSW) used a data set that provides a unique “natural experiment” that allows them to compare the policy forecasts of four competing structural and reduced-form models. The data consist of observations of departure rates of older office employees at a Fortune 500 company. These workers were covered by a “defined benefit” pension plan that provided substantial incentives to remain with the firm until age 55 and then substantial incentives to leave the firm before age 65. In 1982 non-managerial employees who were over 55 and vested in the pension plan were offered a temporary 1 year “window plan”. Under the window plan employees who retired in 1982 were offered a bonus equivalent to 3 to 12 months salary, depending on the years of service with the firm. Needless to say, the firm experienced a substantial increase in departure rates in 1982. Using data prior to 1982, LSW fit four alternative econometric models and used the fitted models to make out-of-sample forecasts of departure rates under the 1982 window plan. One of the models was a reduced-form probit model with various explanatory variables similar in spirit to the PVCOL and ACOL models used by Gotz and McCall.

Using data prior to 1982, LSW fit four alternative econometric models and used the fitted models to make out-of-sample forecasts of departure rates under the 1982 window plan. One of the models was a reduced-form probit model with various explanatory variables similar in spirit to the PVCOL and ACOL models used by Gotz and McCall and Daula and Moffitt. The remaining three models included two version of dynamic programming models and an closely related “option value” model developed by Stock and Wise (1990). All of the three dynamic optimization models fit the pre-1982 data approximately the same, with the dynamic programming models fitting the data slightly better than the option value model. The reduced-form model fit the data slightly better than the dynamic models, although the reduced-form model was a more flexible

specification that have more free parameters than the dynamic programming models. However in “out of sample” predictions of the impact of the window plan, all of the dynamic models generated far more accurate predictions of the impact of the window plan than the reduced-form model, which frequently generated grossly inaccurate predictions of the large increase in departure rates cause by the window plan. Thus, the Gotz and McCall and LSW studies provide a evidence that structural models provide better predictions of the behavior responses to changes in compensation policies than other methods such as reduced-form models.

However neither of these studies attempted to model the impact of health insurance coverage. In the case of the LSW study, there is evidence that the failure to account for health insurance is responsible for the failure of their models to fit the pronounced peak in retirements at age 65, when individuals in the firm become eligible for Medicare coverage. In a subsequent study, Lumsdaine, Stock and Wise (1996) conjectured that this peak in retirements at age 65 was due to a sociological “age 65 retirement effect”. However Phelan and Rust (1997) showed that by estimating a dynamic programming model that carefully models the impact of Social Security and incomplete health insurance markets, one does not need to resort to a sociological hypothesis to explain the large peak in retirements at age 65. The Rust and Phelan model succeeds in explaining the age 65 peak as a result of a substantial fraction of “health insurance constrained” individuals who would prefer to retire prior to age 65 but who do not have access to fairly priced retired health insurance if they were to retire prior to age 65 when they become eligible for Medicare (the Rust and Phelan study was based on the 1969-1979 Retirement History Survey, which was collected before the advent of “COBRA” coverage which gives retiring workers the option to continue health insurance coverage for up to 3 years after leaving the firm at their own cost).

Subsequent work by Blau and Gilleskie (2003) and Khwaja (2003) have also confirmed the importance of modeling health insurance coverage in order to accurately capture individual retirement decisions, although the evidence for the impact of health insurance on retirement decisions is more mixed: “The risk reducing feature of health insurance can fully account for the relatively modest association between retiree health insurance and employment for married men, but can account for only about one third of the large observed association for married women.” (p. 1).

There are differences of opinion in the economics profession about the value of these types of “structural models” for understanding behavior. A number of leading economists believe that

these types of models (including, presumably the model I develop in this paper) are hopelessly oversimplified and misspecified, and thus result in misleading and unreliable predictions of how actual faculty would react to various changes in health insurance and pension benefits. These economists believe that the only reliable way to predict the impact of a change in health insurance or pension benefits is to conduct a controlled, randomized experiment.

In the absence of suitable “historical data” the preferred alternative would be to convince a college or university to conduct a controlled experimental study of the effect of a specified change in their health insurance program. From my perspective, the most promising direction for future research involves combining experimental and structural econometric approaches in a way that enables us to benefit from the strengths of each approach while minimizing each of their weaknesses. Under this approach the first step would be to collect a (panel) dataset on the historical retirement decisions of a sufficiently large sample of faculty members at a number of different universities, operating under a variety of different retirement and health care plans. In addition to reduced-form analyses of the faculty retirement decisions, dynamic structural models similar to the one I propose in this paper could be estimated and tested econometrically. If the behavior predicted by the estimated life cycle model provides a sufficiently good approximation to actual behavior (i.e. if the structural model is not rejected by standard goodness of fit and specification tests), then the next step would be to test how well the structural model would be able to forecast “out of sample.”

The most compelling out-of-sample prediction test would be if one or more university’s or colleges could be persuaded to undertake controlled randomized experiments. The structural model would be able to generate predictions of the full set of observed behavioral responses made by the treatment group in response to the change in the retirement or health coverage provisions. If the actual behavioral response of the treatment group is close to the response predicted by the structural model, then this gives us a greater degree of confidence that the structural model could generate accurate predictions of the behavioral responses to a range of other changes in compensation and retirement package. In such a case, the model could be used as a rapid and cost-effective laboratory for running thousands of computational experiments in order to identify the most promising compensation and pension schemes. This procedure would be the analog of the computerized “crash testing” that automakers use to assess the safety of new model designs

(i.e. they rely on computer simulations of “virtual car crashes” which experiments have shown to be sufficiently accurate approximations to the results of actual car crash tests).

The most ambitious logical extension of this approach could be described as “computational mechanism design” i.e. to systematically optimize over alternative compensation, fringe benefit and retirement packages to find the one that is the most “efficient.” An efficient compensation policy can be viewed as a solution to a “principal-agent” problem between the university (the principal) and the faculty member (the agent). An efficient compensation policy can be defined as one that minimizes the cost to the principal of delivering a given welfare level to the agent. Conceptually we can consider a tenured faculty member as having a given “reservation utility” (i.e. a given level of welfare) under the *status quo*. An efficient compensation policy is one that delivers the this reservation utility level to the faculty member at the least possible cost. Since the welfare of the faculty member is not affected, they would not object if the university adopted an efficient compensation policy, and the university would be strictly better off. The magnitude of the cost savings from an efficient policy relative to the cost of the current system is one way of quantifying the degree of inefficiency in the *status quo*. An example of this approach is the study by Hopenhayn and Nicolini (1997), who estimated that an efficient unemployment insurance system would cost taxpayers 20% less than the *status quo* system currently in effect in the U.S.

While I will not actually attempt to characterize efficient compensation policies in this paper (sufficing instead to evaluating the costs and benefits of a small number of alternatives to the *status quo* using the device of a computational experiment as described above), this does represent the natural logical extension of this line of work. The main impediments to carrying this out in a systematic way to improve the design of retirement and health plans are a) lack of individual data necessary to estimate/calibrate the life cycle model, b) the computational complexity of solving the model. The life cycle model developed in this paper can be solved relatively rapidly: in a few minutes on a standard laptop computer. However to systematically search over *all* compensation/retirement/health insurance policies to find the most efficient one is beyond the computational resources currently at my disposal.

3 A Model of the Academic Retirement Decision

This section develops a specialized version of the life cycle model, which I refer to as the *academic retirement model* (ARM). The model reflects the key features of the tenure contract, and is able to reflect the details of a variety of different academic pension plans (including defined benefit and defined contribution plans such as 403-B plans), and various types of health insurance and retiree health insurance packages. While the focus of this paper is on the effects of variations in health insurance plans on retirement decisions, the ARM model can also be used to evaluate a wide array of other faculty retirement incentives such as phased retirement plans (the subject of Steve Allen’s analysis in chapter XX), and various types of “buy-outs” and “window plans” (the subject of John Pencavel’s analysis in chapter XXX).

The ARM predicts an individual faculty member’s behavior starting at some initial age (I assume age 50 here) through the duration of their employment at their university or college until retirement, including any post-retirement work, until their death. Each period (assumed to be one year in length) the faculty member makes decisions about how much to consume and how much to save, whether or not to retire (if not already retired), whether or not to do either full or part-time post-retirement work (if retired), whether to apply for Social Security benefits (if over the Social Security early retirement age), or wait to apply for Social Security at the normal retirement age, or delay their application further to take advantage of the Social Security *delayed retirement credit*. The ARM includes various sources of risk and uncertainty, including uncertainty about future health status and age of death. If a faculty member becomes disabled and finds it prohibitively difficult to continue working, they also have the option to apply for Social Security disability benefits (these benefits are payable, if awarded, even if the person is below the Social Security retirement age).

In the ARM, an individual conditions their decisions on their current information which includes their current age, wealth, health, and employment status. In addition to health and mortality, the individual faces uncertainty about future health care costs (which depend on health status), and future earnings. Earnings are assumed to be relatively predictable prior to retirement, reflecting the relative security of tenure (i.e. nominal wage cuts almost never occur, and the main uncertainty is by how much real wages will increase). However after a person retires,

there may be considerably more uncertainty about earnings levels reflecting the lower security level provided by “ordinary” untenured jobs, temporary positions, consulting opportunities, etc. The faculty member can support him or herself in their retirement years via a combination of post-retirement earnings, defined-benefit pension and Social Security payments, and the decumulation of their private savings and defined contribution pension balances. However due to the possibility of uninsured or uncovered health care costs, the individual will want to maintain a significant precautionary or “buffer stock” of assets to cover these costs, or to insure against protracted periods of low earnings in their retirement years.

3.1 Health and Mortality

I assume that health and mortality are *exogenous* in the sense that these variables evolve according to stochastic processes that depend on the individual’s age and health status, but do not depend on their employment status, wealth accumulation, or retirement decisions. For example, for simplicity I assume that an individual’s health status h_t takes one of two possible values, *good health* ($h_t = 0$), and *disabled* ($h_t = 1$). Health status evolves according to an age-dependent Markov chain. Thus, between the initial age, 50, and age 51, health status is governed by a transition probability $\pi_{50}(h_{51}|h_{50})$ which can be represented as a 2×2 Markov transition probability matrix Π_{50} :

$$\Pi_{50} = \begin{bmatrix} \pi_{50}(0|0) & \pi_{50}(1|0) \\ \pi_{50}(0|1) & \pi_{50}(1|1) \end{bmatrix} = \begin{bmatrix} .97 & .03 \\ .15 & .85 \end{bmatrix}, \quad (1)$$

where $\pi_{50}(0|0) = .97$ is the probability that a 50 year old who is in good health remains in good health at age 51, and $\pi_{50}(0|1) = .15$ is the probability that a 50 year old who is disabled returns to good health at age 51, etc. There are separate health transition matrices for each age, and for simplicity, I generate these as a simple age-weighted convex combination of Π_{50} and Π_{100} given by

$$\Pi_{100} = \begin{bmatrix} \pi_{100}(0|0) & \pi_{100}(1|0) \\ \pi_{100}(0|1) & \pi_{100}(1|1) \end{bmatrix} = \begin{bmatrix} .90 & .10 \\ .02 & .98 \end{bmatrix}. \quad (2)$$


Thus, the health transition matrix at age $t \in [50, 100]$ is given by

$$\Pi_t = \frac{(100 - t)}{(100 - 50)}\Pi_{50} + \frac{(t - 50)}{(100 - 50)}\Pi_{100} \quad (3)$$

and the generic element of these transition matrices will be denoted $\pi_t(h_{t+1}|h_t)$.

At each age t the probability that the person survives to age $t + 1$ is denoted by $s_t(h_t)$ which depends on health status h_t . Figure 1 displays the survival probabilities that I assumed, which are significantly higher (but generally parallel to) the survival curves for the U.S. population (from the Social Security or Census life-tables) since faculty are generally healthier than the average American — something that may be partly due to their higher levels of education and partly due to better working conditions.

Figure 1: Assumed Survival Probabilities by Age and Health Status



3.2 Health Care Costs

I model out of pocket health care costs (excluding health insurance premiums) as an exogenous random process, denoted by $\{o_t\}$, that depends on age, health status, and health insurance coverage. Realized health care costs also depend on wealth to the extent that if a person should ever incur out of pocket health care costs that exceeds their available assets, all available assets are used to pay off the costs, but the excess balance is “forgiven.” This implicit payment of residual health care costs in event of “bankruptcy” can be viewed as an implicit backstop form of health insurance. However in the model I add a term to the utility function to represent a severe

disutility to such an event, reflecting the possibility that a person might receive inferior health care or be rationed in some manner in an event that they were unable to pay all of the out of pocket health care costs that they incurred in a given year.

In reality, health care costs are “partially endogenous” in the sense that an individual can affect realized health costs by deciding whether or not to incur or postpone certain medical procedures, doctor visits, and so forth. In an analysis of health care expenditures in the Health and Retirement Survey Boz *et. al.* (2004) find evidence of such endogeneity. In particular out of pocket health care costs for individuals without health insurance are about the same as health care costs of individuals who are covered. We would expect, all other things equal, that out of pocket health care costs (excluding insurance premiums) for individuals without health insurance to be greater than for those with health insurance coverage. The explanation for this finding is partly due to *self selection* (i.e. individuals who are healthier than average may be more likely to forgo health insurance coverage), and partly due to *moral hazard* (i.e. individuals who do not have health insurance are more likely to reduce their utilization of health care since they bear 100% of the costs). While a more complete life cycle model would require expanding our treatment of health care to including health care utilization decisions and decisions about whether or not to pay private health insurance premiums (when health insurance coverage is available), for the purposes of this analysis the assumption of exogenous health care expenditures seems like a good first approximation and that little additional insight would be gained from the additional complexity of modeling the extra decisions about health care. Unless health insurance premiums are prohibitively expensive I would expect that few faculty members would voluntarily forgo health insurance coverage, or would significantly ration their use of health care services due to binding budget constraints.


In the ARM, out of pocket health care costs are modeled as realizations from a flexible yet parsimonious three-parameter family of distributions which can be represented as a probability mixture of a Uniform distribution on the interval $[0, b]$ and a Pareto distribution on the interval (b, ∞) . Thus, with probability p out of pocket expenditures are drawn from the uniform distribution on $[0, b]$ and with probability $1 - p$ out of pocket expenditures are drawn from the Pareto distribution on $(0, \infty)$, which implies the following cumulative distribution function for out of

pocket health care costs:

$$F(o|a, b, p) = Pr\{\tilde{o} \leq o\} = \begin{cases} \frac{po}{b} & \text{if } o \leq b \\ p + (1 - p) \left[1 - \left(\frac{b}{o}\right)^a\right] & \text{otherwise} \end{cases} \quad (4)$$

Figure 2 graphs some examples of this distribution for differing values of the (a, b, p) parameters.

Figure 2: Examples of Possible Distributions of Out of Pocket Health Care Costs



This distribution is motivated by previous empirical work by Rust and Phelan (1997) who found that the upper tail of out of pocket health care expenditures is very well approximated by the Pareto distribution. The Pareto distribution has a density function ab^a/o^{a+1} that decreases relatively slowly as a function of o compared to other “exponential tail” distributions such as the Normal distribution. The relatively “fat tails” of the Pareto distribution enable it to capture the small but greatly feared risk of catastrophic health care costs. Interestingly, Rust and Phelan found that the Pareto distribution provided an extremely good fit to the upper tail of realized out of pocket health care costs not only for people without health insurance, but also for people with various types of health insurance coverage.

The explanation for this finding is that as long as there are coinsurance clauses and uncovered expenditures, having health insurance does not provide an absolute guarantee that one will not

incur arbitrarily large out of pocket health care costs. What health insurance does do is mitigate these risks, and this is reflected by a reduction in the b parameter and an increase in the a parameter of the Pareto distribution. Reducing b reduces both the mean and variance of health costs, and increasing a causes the tail of the Pareto density to decrease to zero at a faster rate as $h \rightarrow \infty$, which also decreases the mean and variance.

To illustrate how this 3-parameter family can capture various types of health insurance contracts, consider first variations in the coinsurance rate. Let $F_u(o|a, b, p)$ denote the distribution of *gross* health care costs (i.e. health costs before any health insurance reimbursement). Then suppose a person purchase a health insurance policy with a coinsurance rate of $c \in [0, 1]$ (e.g. if the person's total health care costs are $\tilde{o}$, then the person's share of these costs which must be paid out of pocket is $c\tilde{o}$). It is easy to see that with a coinsurance rate of c , the distribution of out of pocket health care costs is also a member of this family with parameters (a, cb, p) since

$$F_o(o|a', b', p') = Pr\{c\tilde{o} \leq o\} = Pr\{\tilde{o} \leq o/c\} = F_u(o/c|a, b, p) = F_u(o|a, cb, p). \quad (5)$$

Via appropriate changes in the other parameters a and p a variety of different types of health insurance coverage can be modeled. For example, the empirical work by Rust and Phelan (1997) estimated that the a coefficient for individuals without health insurance is $a = 1.04$, for individuals with Medicare it was $a = 1.76$, for individuals with employer-provided health insurance coverage it was $a = 2.29$, and for individuals with Medicare plus retiree or Medigap coverage it was $a = 2.37$. Although the data set that Rust and Phelan used pertains to the 1970s and thus is out of date, we use these qualitative relationships between the values of the parameters for different types of health insurance coverage to guide our choices of the (a, b, p) parameters for the illustrative calculations and simulations below.

To assist the interpretation of the results, it is useful to note that the mean and variance of out of pocket health care costs implied by this 3 parameter specification is

$$\begin{aligned} E\{\tilde{o}\} &= p\frac{b}{2} + (1-p)\frac{ab}{a-1} \\ \text{var}\{\tilde{o}\} &= \left(p\frac{b^2}{3} + (1-p)\frac{ab^2}{a-2}\right) - \left(p\frac{b}{2} + (1-p)\frac{ab}{a-1}\right)^2 \end{aligned} \quad (6)$$

Thus, as $a \downarrow 2$ the variance of out of pocket health care costs tends to ∞ , and as $a \downarrow 1$ the mean of out of pocket health care costs tends to ∞ . The variance is also proportional to b^2 , so that both the mean and the standard deviation of out of pocket health costs is proportional to b .

This makes b a convenient “scale parameter” for this family of distributions. In our simulation analysis in section 4 we allow the b parameter to grow over time as a simple way of modeling the steady escalation in health care costs. For example, if b grows at $x\%$ per year, then the mean and standard deviation of health care costs will also grow at $x\%$ per year.

With this simplified way of modeling out of pocket health care costs, combined with various assumptions about premium payments associated with different types of health insurance, I can model a wide array of different types of health insurance coverage, both while employed (employer-provided health insurance), and when retired (retiree health insurance). The ARM model also reflects Medicare coverage, and other types of private health insurance such as Medigap, and private health insurance coverage such as Blue Cross-Blue Shield, which is not necessarily associated with an employer or retiree health insurance plan.

3.3 Wage and Employment Dynamics

The ARM is specialized to focus on retirement decisions of academics who have tenure in their main career job and a high amount of “human capital” that gives them many post-retirement job opportunities. Therefore the model excludes the possibility “involuntary unemployment” on the main tenured job. A closely related aspect of the high degree of job security provided by tenure is the fact that in the absence of mandatory retirement rules, an older faculty member has virtually full control over the timing of his/her retirement. The main risk that I model is the possibility that an unexpected disability or severe health problem could occur that would make it very difficult for the faculty member to continue their full-time duties, and which would therefore prompt them to voluntarily retire earlier than they might have desired to retire had the faculty member remained in good health.

However I assume that post retirement work is not governed by a tenure contract and if the person chooses to do post-retirement work, there is more uncertainty about earnings. This uncertainty reflects both uncertainty about finding post-retirement job opportunities, and the possibility of involuntary unemployment or unexpected loss of a post-retirement job opportunity that could result in a drop in earnings until a suitable alternative job is found. The uncertainty in earnings also reflects a higher degree of uncertainty in the rate of pay in various post-retirement jobs. In contrast, nominal wage cuts for tenured faculty members are extremely rare, and usually

there is only relatively small uncertainty about the whether nominal wage increases, if any, will exceed the rate of inflation. Also, to the extent that faculty get up to 1/3 of their salary from research grants or summer teaching assignments that are not absolutely guaranteed, there is could be additional uncertainty in their annual earnings. However overall, I assume there is a relatively small amount of uncertainty about whether real wages will go up or down, compared to a much higher degree of uncertainty associated with earnings from post-retirement jobs.

I can reflect these considerations mathematically in the following simple autoregressive model of earnings dynamics:

$$\log(y_{t+1}) = \begin{cases} a_0 + b_0 \log(y_t) + \sigma_0 \epsilon_t & \text{if } rd_t = 0 \\ [a_1 + b_1 \log(y_r) + \sigma_1 \epsilon_t] \gamma_t(l_t) & \text{if } rd_t = 1 \end{cases} \quad (7)$$


In this equation, the variable y_t denotes earnings, rd_t denotes the retirement decision, and l_t denotes the post-retirement work decision. If the person decides not to retire, $rd_t = 0$, then earnings are given by the top autoregressive earnings equation, with coefficients (a_0, b_0, σ_0) , where ϵ_t is a standard normal random variable representing the uncertainty a faculty member has about how much earnings he/she will make on their tenured job next year if they decide not to retire. I make the standard deviation of these unexpected shocks, σ_0 , relatively small to reflect the relatively small degree of uncertainty about future earnings in the main tenure job.

However if the faculty member retires, with a ending year total earnings of y_r , their subsequent earnings in post-retirement work, given in the second equation of formula (7) above, has a significantly higher standard deviation, σ_1 , reflecting the greater uncertainty discussed above. Further the coefficients a_1 and b_1 are adjusted so that earnings in post-retirement work is not expected to be as high as earnings in the main tenured job. Further, if the individual decides to work part time instead of full time, the factor $\gamma_t(l_t)$ acts to pro-rate the full-time equivalent earnings rate given in the bracketed expression in the second equation of formula (7) into their actual part time earnings. This pro-ration term $\gamma_t(l_t)$ can also reflect the fact that part time jobs do not generally earn as high a rate of pay as full time jobs. Also, if the person decides not to work at all, $l_t = 1$ (denoting decision for full-time leisure, see below), then $\gamma_t(1) = 0$ and realized earnings are zero.

Figure 3 shows how all of these considerations are reflected in the ARM. It plots two example realizations of pre and post-retirement earnings from the base case simulations of ARM that will

be discussed in the next section. We see that prior to retirement, although the growth in real wages is not perfectly predictable, the uncertainty about future real wages while remaining on the main tenure job is relatively insignificant. In both cases we see a trend of modest real wage growth right up until the time of retirement. However we also see that the post retirement earnings are generally lower and much more volatile than the pre-retirement earnings stream. This is due in part to the choice of the (a_1, b_1, σ_1) parameters, which make post-retirement earnings relatively more uncertain, but also with a lower mean value than the comparable pre-retirement earnings. However a significant amount of volatility in earnings is endogenous, in the sense that it results from decisions about whether to work full or part time or not at all. Thus, the post-retirement earnings are in a sense more “opportunistic” and depend on a variety of factors including the person’s health status and wealth levels. Unexpected health care shocks that lead to unexpected decumulations in wealth can prompt an individual to engage in post retirement work to restore “precautionary” wealth holdings in order to maintain consumption levels and as a buffer against further future unexpectedly high out of pocket health care costs.

Figure 3: Examples of Earnings Before and After Retirement



A final thing to note is that once a person retires their subsequent earnings are conditioned on their final earnings in their tenured job, whereas prior to retirement, their earnings in year $t + 1$ are conditioned on their earnings in year t . This means that earnings are serially correlated prior to retirement, but in the post retirement period, earnings are conditionally independent, and thus

less “persistent” and (combined with the higher standard deviation term σ_1) harder to forecast. I can also make the coefficients (a_i, b_i, σ_i) be functions of health status, age and other variables to reflect reductions in earnings with age, which can be especially important in part time work.

A final important consideration is the faculty member’s ability to control hours of work. In the initial simulations of the model that will be presented in the next section, I assume that there is no phased retirement program at the college/university where the person’s main tenured “career job” is at. This means that prior to retirement the person does not have a flexible choice over hours of work, but instead is required to carry a full time load. Thus, a faculty member does not have the option to gradually phase down their work activity, and this can also be a motivation to retire (since a faculty member is allowed to choose to work only part time once retired).

It might seem that this assumption is inconsistent with my comments in the introduction, namely, that given the high degree of personal freedom and control over hours of work makes some faculty members view their full time tenured jobs as “virtual retirement”. However when I describe preferences below, you will see that I can implicitly account for the pleasant, flexible working conditions via the magnitude of the “disutility of work.” For faculty members who are in good health, this disutility of work may be very low. However I assume there is a sense of “moral suasion” within academics that creates a sense of guilt or embarrassment experienced by older faculty members who “check out” and do not carry their “full load.” Further, if a faculty member’s research starts to slow down as they age, their university or department may have rules that require them to do more teaching or committee work, and it may be hard for the faculty member to avoid these duties. Thus, I think the assumption that prior to retirement the faculty member faces a dichotomous choice between working full time versus retiring and does not have the opportunity to gradually phase down their hours of work and effort levels is a reasonable approximation to the situation in most universities that do not have phased retirement programs, especially when you consider the relatively flexible way I model the disutility of work, to be explained below.

3.4 Social Security, Disability and Pension Benefits

The ARM has sufficient flexibility to accommodate both defined benefit (DB) and defined contribution (DC) pension benefits from the person’s main job. Defined benefit pension rules can take a

variety of forms, and the formula providing the benefit payments during retirement can be quite complicated. A commonly used pension formula takes the form

$$\text{pension benefit} = \text{benefit factor} \times \text{years of service} \times \text{final salary} \quad (8)$$

where the “benefit factor” term is a constant that determines the overall level of generosity of the pension plan. For example if factor = .02, then a person with 40 years of would receive a pension that provides an 80% “replacement rate” for their final salary — a generous pension plan. Defined benefit pension plans also frequently have a minimum retirement age such as age 55, and benefits cannot be collected prior to this age. The benefit factors also typically contain actuarial adjustments that increase or decrease the pension benefit based on when the employee retires relative to a target or normal retirement age (often the same as the Social Security normal retirement age). These actuarial adjustments are required by law to be approximately “actuarially fair” – i.e. the expected present value of benefits is independent of the age of retirement – otherwise the plan could violate anti-age discrimination laws.

Defined contribution plans are more common than defined benefit plans among academic employers, and over time, are becoming the dominant type of pension plan among U.S. employers more generally. DC Many institutions have a basic plan where the employer and employee make joint contributions, which are invested with an intermediary such as Vanguard or TIAA-CREF which manages the accounts and provides a number of broad-based portfolios from which the faculty member can choose to invest in. As soon as the person is over the minimum retirement age, they have the option of converting their accumulated DC pension balances into an annuity, or withdrawing their balances as they need them, or some combination of partial annuitization and withdrawal. In addition to the employer’s DC pension plan, a faculty member can choose to invest additional funding in a supplemental voluntary defined contribution plans such as 403-B plans. These are similar to individual retirement accounts (e.g. 401-K plans) in that a faculty member can contribute a portion of their salaries up to a certain annual limit, and these contributions are deducted from their taxable income. Taxes are only paid upon withdrawal, once the individual reaches age 55. There are significant penalties for early withdrawals from these plans, although it is typically possible to borrow at favorable rates, using the plan balances at collateral.

In the simulations presented in the next section I assume the faculty member has a defined contribution pension plan, but no defined benefit plan. For simplicity, I treat the person’s defined

contribution as part of their overall net worth. That is, the model does not have separate variables that record the person's pension balances, 403-B balances, but rather merges these balances into their overall net worth. This means that the model ignores possible early withdrawal penalties to the extent that the person decides to draw on either of these balances prior to the minimum retirement age (assumed to be age 55 in this model). For an individual who is age 50 and who has substantial net worth apart from their DC pension balances, it would be extremely unlikely that they would ever need to make early withdrawals from their DC pension balances. Further, as noted above, the ability to borrow against their pension balances at favorable interest rates would have virtually the same effect as an early withdrawal, but enables them to avoid the tax penalties. For these reasons, ignoring early withdrawal penalties in this version of the model has little effect on the model's predictions.

The ARM also has a detailed and fairly accurate treatment of Social Security Old Age, Disability, and Medicare benefits. Medicare benefits are available at age 65. There are two parts to Medicare, Part A (Hospital Insurance) and Part B (Medical Insurance). There are no premiums for Part A coverage, however Part B coverage is optional and requires payment of a premium, which in 2003 amounted to \$720 per year. I assume that all individuals in the ARM elect the part B coverage. The combined coverage does involve a 20% coinsurance rate for most covered procedures and services. In addition there are a number of uncovered services, such as nursing homes. In the solution and simulation of the model I adjust the parameters of the 3-parameter distribution of out of pocket health care costs discussed above to reflect Medicare coverage once a person reaches age 65, as well as the interaction of Medicare and retiree health insurance and Medigap policies on the distribution of out of pocket health care costs.

In the ARM, a person can apply for Social Security Disability benefits. Any person who is not already receiving SSDI benefits is eligible to apply for SSDI benefits provided they are younger than the normal retirement age. If they are over the normal retirement age, then they only have the option to apply for Social Security Old Age benefits. Prior to the normal retirement age the person has the option of applying for SSDI, or if they are over the early retirement age, they can also apply for early retirement benefits, which are actuarially reduced based on the number of years prior to the normal retirement age at which they first start receiving their OA benefits. While there is a 100% probability of being awarded Old Age benefits if one applies and is age-eligible,

the probability of being awarded SSDI benefits is considerably less than 1. Using data from the HRS and also aggregate program data from the SSA web site, we estimated a probabilistic model of the SSDI award process (see Benítez-Silva, Buchinsky, Chan, Cheidvasser, and Rust 1999).²

Many academic institutions provide private disability insurance benefits that supplement SSDI benefits which are generally slightly above the poverty line. An award of private disability benefits is typically contingent on being awarded SSDI benefits. Once awarded, the combination of SSDI and private disability benefits can often provide relatively high “replacement ratios” for example, raising the replacement ratio of only 20-30% for SSDI only to higher than 70% for the more generous private disability insurance plans. In the simulations in the next section we will assume the faculty member has no private disability insurance plan and their only coverage is SSDI. As a result of the lower replacement ratio, almost no faculty members apply for SSDI benefits even when they become disabled since they find it preferable to try to carry out their teaching duties in the presence of their disability rather than retire and try to live on a much lower disability benefit. However when I allow for more generous private disability insurance, then the simulations do reveal a significant number of faculty members who elect “disability retirement” when they experience a disabling health condition.³

Social Security Old Age benefits are based on the person’s *primary insurance amount* (PIA) which is a piece-wise linear, concave function of their *average indexed monthly earnings* (AIME). The AIME is an average of the 35 highest indexed earnings in the person’s earnings history. We have found that the AIME is well approximated by a simple moving average of indexed earnings (truncated at the Social Security maximum earnings limit), taken over the entire earnings history (i.e. we have not dropped out the lowest 5 years of earnings). This moving average wage, aw_t ,

² The probabilities of being awarded benefits depend on the individual’s health status and labor supply decision in the year of application. The probability of being awarded SSDI benefits is higher the worse one’s health status is, but it is zero if they choose to work at a level such that their earnings exceed a small threshold (currently \$800 per month) known as the *substantial gainful activity* (SGA) level. A person is allowed to appeal a denial, and also to repeatedly apply and/or appeal for SSDI benefits. This “repeated game” aspect of the SSDI program means that the “ultimate award rate” is approximately 70%, much higher than the 50% “initial award rate” that would be inferred by looking only at decisions by the disability determination services (DDS) for individuals who are aged 50 or older.

³ Even if a person is awarded SSDI benefits, they will not necessarily remain on SSDI until they die or reach the normal retirement age (at which point they will be automatically transferred from the SSDI program to the OA program). Individuals who decide to return to work on a full or part time basis will be terminated from the SSDI roles after a 9 month trial work period provided their earnings exceed the SGA level. Also there is a small probability that an individual will be involuntarily terminated as a result from random audits that are known internally in SSA as *continuing disability reviews* (CDRs). We allow the probability of being terminated due to a CDR to be a function of a person’s health status, with persons in good health status having a substantially greater risk of termination than those who are in poor health or who are disabled.

can be written recursively as

$$aw_{t+1} = \frac{t}{t+1}aw_t + \frac{1}{t}y_t. \quad (9)$$

If we regress aw_t on the exact AIME (calculated from the person’s earnings history using SSA’s *ANYPIA* program), the R^2 of the regression is 98%, which confirms that aw_t is an accurate predictor of AIME. The advantage of using aw_t instead of AIME is that aw_t becomes a “sufficient statistic” for the person’s earnings history. Thus we need only keep track of aw_t and update it recursively using the latest earnings according to equation (9) above, rather than having to keep track of the entire earnings history in order to determine the 35 highest earnings years, which the AIME requires.

The DP model also accounts for actuarial reductions in old age benefits claimed prior to the NRA, and for the delayed retirement credit (DRC) for benefits claimed after the NRA. In the ARM we assume the relevant NRA is 66 and the delayed retirement credit rate is 7% per year, i.e. the values applicable to individuals who were born between 1939 and 1940. If an individual retires early, i.e. prior to the NRA, the PIA is permanently reduced by an actuarial reduction factor of $\exp(-g_1k)$, where k is the number of years prior to the NRA but after age 62 that the individual first starts receiving OA benefits. The actuarial reduction rate for the 1931 to 1941 cohort is $g_1 = .0713$, which results in a reduced benefit of 80% of the PIA for an individual who first starts receiving OA benefits at age 62. Note that a person who is accepted into the DI program prior to the NRA receives the full PIA regardless of his/her age. However, the SSA does apply an actuarial reduction to DI benefits that are awarded after the ERA.

As we will see in our simulations in the next section, because of their better than average health, their later than average retirement age from academics, and due to the loss of Social Security benefits due to the *earnings test* if a person applies for Social Security Old Age benefits while continuing to work, very few academics apply for early retirement benefits from Social Security. However the earnings test has been abolished for individuals over the NRA, so we see a rapid rise in applications for Social Security benefits after this age, with a peak at age 70. The reasons applications peak at this age is that the delayed retirement credit only increases benefits up until this age: further delays in retirement result in no further increase in retirement benefits. Thus, it would be irrational to delay applying for Social Security benefits after age 70 since any delays reduce the total amount of benefits one will collect from Social Security.

A final aspect of the Social Security rules that ARM accounts for is taxation of benefits. Individuals whose combined income (including Social Security benefit) exceeds a given threshold must pay Federal Income taxes on a portion of their Social Security benefits. ARM incorporates this taxation well as the 15.75% Social Security payroll tax and the Federal income tax on wage earnings.⁴

3.5 Preferences for Consumption and Leisure

Perhaps the most important assumption driving the predicted labor supply and retirement behavior produced by the ARM is on individual preferences for consumption and leisure. I assume that individuals have a *utility function* $u_t(c_t, l_t, h_t)$ that specifies their preferences for consumption c_t and leisure l_t as a function of their age t and health status h_t . I assume that the individual's utility function is given by

$$u_t(c, l, h) = \frac{c^\gamma - 1}{\gamma} + \phi(t, h) * \log(l) - 2h \quad (10)$$


I measure consumption in (real) dollars, and it is treated as a continuous decision. Leisure l_t is normalized to 1 representing full leisure (i.e. no work) during the course of a year, and it is treated as a discrete decision. That is, in addition to not working, I assume that the individual can either choose to work full time or part time. The amount of leisure associated with full time work is assumed to be $l_t = .543$ under the assumption that full time work requires 2,000 hours per year, leaving the fraction $l = .543 = (12 * 365 - 2000)/(12 * 365)$ of the person's waking hours during the year for leisure (I assume that sleep and other basic bodily functions require 12 hours per day, leaving 12 hours for discretionary activity that can be divided between work and leisure). Similarly, I assume a part time job requires an average of 800 hours per year, which corresponds to a choice of leisure equal to $l = .817 = (12 * 365 - 800)/(12 * 365)$.

The parameter γ in equation (10) indexes the individual's level of risk aversion. As $\gamma \rightarrow 0$ the utility of consumption approaches $\log(c)$. In the simulations reported in the next section I assumed a value of $\gamma = -.37$, which corresponds to a moderate degree of risk aversion. The

⁴ The ARM also incorporates a detailed model of taxation of other income, including the progressive Federal income tax schedule (including the negative tax known as the EITC – Earned Income Tax Credit), and state and local income, sales and property taxes. However given the high income of most academics, almost none of them are eligible for EITC.

left hand panel of Figure 4 plots the utility of consumption (i.e. the first term in equation (10), $(c^\gamma - 1)/\gamma$, as a function of c).

Figure 4: Utility of Consumption and Leisure Functions Assumed for the ARM



The function $\phi(t, h)$ can be interpreted as an age and health-dependent weight on the relative disutility of work. Notice that since leisure $l \in (0, 1]$ we have $\log(l) \leq 0$, so that $\phi(t, h) \log(l)$ can be either viewed as the utility of leisure, or equivalently as the disutility of work. When $l = 1$, the person is not working and engaging in 100% leisure, and we have $\log(l) = 0$, so the disutility of work is zero. Since $\log(.543) < \log(.817) < 0$, there is a higher disutility from working full time, or equivalently, higher utility from leisure associated with part time work than full time work. We assume that ϕ an increasing function of age and is uniformly higher for individuals whose health status is worse. The right hand panel of figure 4 illustrates the ϕ function I have used in this analysis. There are two curves, one for individuals who are in good health, and the other for individuals who are “disabled.” We see that the ϕ function increases with age for both health states, but the disutility of work for a disabled person is uniformly and significantly higher higher at all age levels. Note also that the weight on the disutility of work for a person who is in good health is very low up until the mid 80s, after which it starts to rise more steeply. This is intended to reflect the good working conditions of an academic, and the fact that many academics view their work a virtually equivalent to leisure time. This would be exactly true if the weight $\phi(t, h)$ were exactly equal to 0 since then the faculty member would be indifferent between

full time work and full time leisure. However I make the disutility work to be slightly positive up until age 80 to reflect that the fact that while a faculty member may love many aspects of his/her work (e.g. research, writing, attending seminars), there may be aspects of the job such as teaching large undergraduate classes or doing committee work or attending faculty members that the faculty member really dislikes. As long as there are at least some part of the academic job that the faculty member dislikes, it seems reasonable to provide a slight disutility of work to reflect this. After all, a faculty member also has the option to do research and attend seminars and in general pick out the positive aspects of their job and discard the negative ones by retiring and engaging in the parts of the job they like as a type of unpaid leisure activity.

3.6 Solving the ARM

Solutions to the ARM are computed numerically, using dynamic programming. The solution proceeds by backward induction, starting from the terminal age, which is assumed to be 100. Let $s_t(h)$ be the probability that a person of age t survives to age $t + 1$. Starting the backward induction at age 100 amounts to the implicit assumption that $s_{100}(h) = 0$, i.e. nobody survives beyond age 100. As we can see from figure 1, the survival probabilities at age 100 are actually well above .5. It turns out that due to the combined effects of discounting of future utility and uncertain survival prior to the terminal age, the predicted behavior in the relevant age (age 50 to 80 in this analysis) is insensitive to the assumed terminal age: increasing the terminal age to 110 or 120 has little effect on the model’s prediction in the relevant age range. Thus, starting the backward induction at age 100 saves on computation without having a significant effect on the results.

In order to state the Bellman equation, which formalizes the calculations that are made in the backward induction process to solve the ARM, tables 1, 2 and 3 summarize the relevant decision and state variables, and the notation for the “laws of motion” of the state variables. As we see from table 1 there are 4 decision variables in the ARM: 1) a consumption decision c_t , 2) a work/leisure decision l_t (constrained to the three options, don’t work, work full time or work part time), 3) the decision to retire from their tenured “career job” $rd_t = 1$, and 4) the decision to apply for Social Security benefits, ssd_t . Table 1 does not fully reflect the constraints on these choices, which can depend on the person’s age and state.

Table 1: Summary of Decision Variables in the ARM

Symbol	Description	Type
c_t	consumption expenditures	continuous, c_t
l_t	leisure choice	discrete, $l \in \{1, .817, .543\}$
rd_t	retirement decision	discrete, $rd \in \{0, 1\}$
ssd_t	Social Security decision	discrete, $ssd \in \{0, 1, 2\}$

Table 2: Summary of State Variables in the ARM

Symbol	Description	Type
t	age	discrete, $t \in \{50, 51, \dots, 100\}$
w_t	net worth	continuous, $w_t \in [0, \infty)$
y_t	annual earnings	continuous, $y \in [0, \infty)$
o_t	out of pocket health care costs	continuous, $o \in [0, \infty)$
aw_t	average indexed earnings	continuous, $aw \in [0, \text{ssmax}]$
h_t	health status	discrete, $h \in \{0, 1\}$
r_t	retirement status	discrete, $r \in \{0, 1\}$
ss_t	Social Security status	discrete, $ss \in \{0, 62, 63, \dots, 70\}$

For example, the budget constraint (to be presented below), combined with the assumption that a person cannot borrow against future earnings or Social Security benefits implies that c_t cannot be larger than the person's accumulated net worth, w_t . I also assume that the retirement decision is irreversible, so that once a person retires from their tenured job, $rd_t = 1$, they remain retired for the remainder of their life. Finally, the Social Security decision can take 3 possible values, where $ssd_t = 0$ denotes the decision not to apply for Social Security benefits, $ssd_t = 1$ denotes the decision to apply for Old Age Benefits, and $ssd_t = 2$ denotes the decision to apply for Social Security Disability Benefits. A person is not eligible to apply for Old Age benefits until they are 62, and once they start receiving Old Age benefits, they do not have the option to stop receiving them (except if they die). Individuals can apply for disability benefits at any age below the normal retirement age: after this age there is no distinction between Old Age and Disability benefits. Thus, prior to age 62 the Social Security disability choice set consists of $\{0, 2\}$ (i.e. whether or not to apply for disability benefits, or leave the SSDI roles if they are already receiving SSDI), and between age 62 and the normal retirement age, they have the full choice set $\{0, 1, 2\}$ unless they are already receiving old age benefits (in which case they have the two

options $\{0, 2\}$), and after the normal retirement age they only have the choice of whether or not to apply for old age benefits, $\{0, 1\}$, unless they are already receiving Social Security benefits, in which case they have no further decisions. Finally, since there is no actuarial increase in benefits past age 70, it is never optimal to delay applying for old age benefits after this age, so that in fact, there are no further Social Security decisions in the ARM past age 70.

The Social Security state variable ss_t reflects the Social Security decision. It takes on 10 possible values, $ss \in \{0, 62, 63, \dots, 70\}$ where $ss = 0$ denotes not receiving Social Security, $ss = 62$ denotes an individual who first started receiving Social Security benefits at age 62, and so forth. Individuals who are receiving Social Security disability benefits are denoted by state $ss = NRA$, where NRA is the normal retirement age. We can distinguish a person who is receiving SSDI from someone who first started receiving old age benefits at the normal retirement age by looking at the persons age: if they are younger than NRA then they are an SSDI recipient, otherwise they are an Old Age insurance beneficiary. It is necessary for the ARM to “remember” the age that a person first started receiving Old Age benefits due to the permanent actuarial adjustment in benefits (either a reduction in benefits for retiring before NRA or an increase in benefits for retiring after the NRA) as discussed above.

Table 3: Laws of Motion for State Variables in the ARM

Symbol	Description	Type
$s_t(h)$	survival probability	discrete
$\pi_t(h' h)$	health transition probability	discrete
$\rho_t(ss' ss, ssd, h)$	Social Security state	discrete
$g_t(y' y, l, r)$	distribution of earnings	continuous
$f_t(o h, r, ss)$	distribution of OOP health care costs	continuous
$prem_t(h, ss, r, y)$	health insurance premiums	continuous
$wp_t(w, y, aw, o, c, r, ss, rd, ssd)$	budget constraint	continuous
$ssb_t(aw, y, r, ss)$	Social Security benefits	continuous
$pb_t(y, r)$	pension benefits	continuous
$pc_t(y, r)$	pension contribution (from employer)	continuous
$\tau_t(y, w, o, r, ss)$	Federal, state and local taxes	continuous

Table 3 summarizes the various laws of motion for the state variables in the ARM. The probability functions $s_t(h)$ and $\pi_t(h'|h)$ for mortality and health status have already been discussed

above. Similarly we have also discussed the distributions governing out of pocket health care costs and earnings, but a few comments on notation are in order. Although we originally wrote the distribution of out of pocket health care costs depends on three parameters, (a, b, p) (see section 3.2), we allow these parameters to depend on the person's age and state. Thus, we write $f_t(o|h, r, ss)$ to reflect the conditional distribution of health care costs given the person's health, retirement, and Social Security status. As noted in section 3.2, it is necessary to account for the state variables r and ss since these determine whether the person is covered by their employer health insurance, retiree health insurance, and Medicare, and thus affects the distribution of out of pocket health care costs that they expect to incur.

The notation $\rho_t(ss|ss, ssd, h)$ is required to reflect stochastic elements in the law of motion for Social Security status due to the probabilistic nature of the disability award process. Thus, if a person who is not currently receiving SSDI decides to apply, $ss = 0$, $ssd = 2$, the probability $\rho_t(ss_{t+1} = NRA|0, 2, h)$ denote the probability that this applicant will actually be awarded benefits next period. Also, if a person is currently receiving SSDI, the ρ_t function also reflects the possibility that the person could be randomly audited and terminated from the roles. The probability function ρ_t depends on the person's health, reflecting the fact that those who are disabled are more likely to be awarded benefits and less likely to be removed from the SSDI roles due to an audit than are those who are in good health. The laws of motion for Social Security status for those receiving old age benefits are deterministic and quite simple, e.g. $ss_{t+1} = ss_t$ with probability 1 (i.e. once someone starts receiving old age benefits, they continue receiving them until they die). The ρ_t function encompasses both situations.

The functions τ_t , pc_t , pb_t , and ssb_t encode the rules governing taxes, the employer's contributions to the employee's defined contribution pension plan, defined benefit pension payments (if any), and Social Security old age or disability benefits. As noted above, Social Security benefits depend on the person's average indexed monthly earnings, aw_t which is used to calculate the primary insurance amount (PIA), and the PIA is then actuarially adjusted to reflect early or late retirement (as encoded by the Social Security state variable ss).

Perhaps the most important of the law of motion formulas is the budget constraint, which can be written as

$$w_{t+1} = wp_t(w_t, y_t, aw_t, o_t, c_t, r_t, ss_t, rd_t, ssd_t), \quad (11)$$

where the function wp_t is defined by

$$\begin{aligned} wp_t(w, y, aw, o, c, r, ss, rd, ssd) = & R(w - c) + y - o - prem_t(h, ss, r, y) + ssb_t(aw, y, ss, r) + \\ & pc_t(y, r) + pb_t(y, r) - \tau_t(yh, w, o, r, ss). \end{aligned} \quad (12)$$

This equation simply says that end of period wealth equals the interest plus principal on the amount $(w - c)$ that was invested at the start of the current period, plus earnings y , pension benefits, employer pension contributions and Social Security benefits, less out of pocket health care costs, o , health insurance premiums, and taxes. I assume that all payments for taxes, health care costs, etc. are due at the end of the period, and that wage and pension payments and Social Security benefits are also paid at the end of the period. Thus, interest is only earned on the portion of beginning of period wealth $w - c$ that was not used to finance consumption expenditures, where the latter are assumed to be incurred at the beginning of each period.

Now that all of the notation is defined, the Bellman equation governing the backward induction solution procedure for the ARM can be written as

$$V_t(w, y, aw, ss, r, h) = \max_{\substack{c \leq w, r \in \{0, 1\} \\ l \in \{1, .817, .543\} \\ ssd \in \{0, 1, 2\}}} \left[\begin{aligned} & u_t(c, l, h) + \beta s_t(h) EV_{t+1}(w, y, aw, ss, r, h, c, l, rd, ssd) \\ & + \beta [1 - s_t(h)] EB_{t+1}(w, y, aw, ss, r, h, c, l, rd, ssd) \end{aligned} \right] \quad (13)$$

where V_t denote the person's expected discounted utility at age t and EV_{t+1} denotes the conditional expectation of discounted utility from age $t + 1$ onwards, where all of its arguments are the conditioning variables in the calculation of the expectation. Similarly, EB_{t+1} denotes the conditional expectation of bequests in the event the person dies. We can write the EV_{t+1} function in terms of the notation given above as

$$\begin{aligned} EV_{t+1}(w, y, aw, ss, r, h, c, l, rd, ssd) = & \sum_{h'} \sum_{ss'} \int_{o'} \int_{y'} V_{t+1}(wp_t(w, y', aw, o', c, r, ss, rd, ssd), y', aw, ss', rd, h') \pi_t(h'|h) \rho_t(ss'|ss, ssd, h) \\ & \times f_t(o'|h, r, ss) g_t(y'|y, l, r) do' dy' \end{aligned} \quad (14)$$

The formula for EB_{t+1} is similar except that instead of taking the expectation of the next period value function V_{t+1} , we take the expectation of the *bequest function* $b_t(w)$, where w is the wealth

that is (accidentally or intentionally) bequeathed at death. I assume that the bequest function is given by

$$b_t(w) = \delta \left[\frac{t}{1+t} \right] \left[\frac{w^\gamma - 1}{\gamma} \right] \quad (15)$$

where δ is a parameter determining the strength of the bequest motive, and γ is the same coefficient of risk aversion that I used in the person's utility function (10).

Note that the equation (14) embodies a simplifying assumption that $aw_{t+1} = aw_t$ with probability one. As I discussed in section 3.4, since aw_t is an average of the 35 highest years of earnings, by age 50 (the initial age in this analysis) additional years of earnings change the person's average wage by a relatively insignificant amount. For this reason I decided to treat aw as a fixed, time invariant parameter in the analysis that follows. This results in considerable simplification of the solution of the dynamic programming problem since aw no longer has to be carried around as an explicit (continuous) state variable.

3.7 Results

Using numerical methods (see Rust, 1994), I solved the dynamic programming recursions given in equation (13) for a set of parameters intended to capture the beliefs of a "typical academic" at a university with a defined contribution pension plan and a generous employee and retiree health insurance plan. I assumed that at age 50, the person had an average wage of $aw = 40$ (amounts in thousands of dollars). The solution is based on the utility function illustrated in figure 4, the health and mortality probabilities described in section 3.1, the earnings dynamics described in section 3, and a discount factor of $\beta = .92$. I assume that the faculty member can invest his/her savings in a fund that earns a real rate of return of 4% per year,

One of the most important assumptions driving the results in the next section is the assumption about health insurance coverage and health care costs. Under the "base case" scenario, I assume that the faculty member is covered under a generous health insurance plan provided by his/her university, who pays the premiums both before and after retirement, until the death of the faculty member. I assume the plan is a "family plan" that covers all members of the household, and the plan has cost and coverage characteristics similar to the coverage offered by Duke University to its faculty and staff. The annual premium for this plan is \$9,000. Via some rough estimates, I assume that given the deductibles and coinsurance rates for this plan, the distribution of out

of pocket health care costs under this plan can be approximated by the 3 parameter family of distributions discussed in section 3.2, with parameters $(a, b, p) = (2.3, 2, .9)$ if the person is in good health and $(a, b, p) = (2.1, 2, .8)$ if the person is disabled. To provide some context, note that using the formulas for the mean and variance of health care costs in equation (6) the implied mean and standard deviation of out of pocket health care costs is \$1,253 and \$1,641, respectively, for individuals who are in good health, and \$1,563 and \$3,927, respectively, for individuals who are disabled.

The current rate of growth in health insurance premiums — both private health insurance and Medicare — is in excess of 10% per year. As Schieber’s analysis (this volume) points out, this rate of increase is not sustainable in the long run since it exceeds the real growth rate of the economy. In the base case analysis, I assume that premiums and overall health care costs will grow at 14% (real) for the next 15 years, after which a major national health care reform occurs that causes health premiums and costs to stop growing entirely, flattening out at the values it reaches 15 years from now and staying fixed at that level thereafter. Specifically under the base case solutions of the ARM, I assume that both health insurance premiums and the b (scale) parameter of the distribution of out of pocket health care costs increase at 14% per year for the next 15 years (from age 50 to age 65) and then remain constant at their age 65 values thereafter.

I assume that as long as the faculty member has not retired, their real earnings grow stochastically, with the growth rate being lognormally distributed with location parameter μ_t and scale parameter σ_t . I assume that $\mu_{50} = .015$ and $\sigma_{50} = .013$, which implies that at age 50 the expected growth rate in real earnings is 1.5%. As the person ages, both μ_t and σ_t decline linearly towards $\mu_{100} = .005$ and $\sigma_{100} = .007$, so that the real rate of growth in earnings is steadily decreasing with age. If the faculty member retires, I assume that earnings in post retirement work involves a lower effective full time wage rate, and higher variability compared to their tenured career job. The lognormal μ_t parameter for a faculty member who has retired declines linearly from $\mu_{50} = -.3$ to $\mu_{100} = -.6$ and σ_t declines linearly from $\sigma_{50} = .02$ to $\sigma_{100} = .01$. This implies that at age 50, the expected full time earnings from post retirement work is only 75% as high as the expected earnings on their tenured career job.

The solution to the ARM consists of 4 *decision rules* which are functions of the state variables

that specify the optimal choices for the 4 decision variables in the problem:

$$\begin{aligned}
c_t &= c_t(w, y, ss, r, h) \\
l_t &= l_t(w, y, ss, r, h) \\
rd &= rd_t(w, y, ss, r, h) \\
ssd &= ssd_t(w, y, ss, r, h)
\end{aligned} \tag{16}$$

Figure 5 illustrates the function rd_t as a function of (w, y) for an individual who is age 70 ($t = 70$), in good health ($h = 0$), who has not retired yet ($r = 0$) and who is not receiving social security benefits ($ss = 0$). The figure shows that retirement decision can be represented as a *threshold rule* in wealth:

$$rd = \begin{cases} 1 & \text{if } w > \lambda_t(y, ss, h) \\ 0 & \text{otherwise} \end{cases} \tag{17}$$

From figure 5 we see that the retirement threshold is an increasing function of pre-retirement earnings, y .

Figure 5: Optimal Retirement Decision for 70 year olds





Figure 6: Plots of Optimal Consumption, Labor Supply, and Value Functions



The top 2 panels of Figure 6 plot the optimal consumption functions at ages 50 and 70 for an individual who is in good health and who has not retired or applied for Social Security benefits. We see that optimal consumption is an increasing function of both y and w . Note that at very low levels of wealth w the individual is *liquidity constrained* and the optimal consumption rule is to consume 100% of the person's wealth: $c_t = w_t$ when w_t is sufficiently small. In our simulations in the next section we will be focusing on faculty members whose combined pension accumulations and net worth w_t is sufficiently large that they will never be liquidity constrained. When there is no liquidity constraint it is easy to see from figure 6 that the marginal propensity to consume out of wealth, i.e. the derivative of c_t with respect to w , is significantly below 1. This low marginal propensity to consume reflects a *precautionary savings motive*, i.e. the individual wishes to save a significant share of their accumulated wealth as a “buffer stock” against unexpected health care costs o_t , and unexpectedly low post-retirement earnings (which could be partly due to unexpected health problems that make post-retirement work very difficult), and a desire to have enough assets to supplement Social Security benefits in the event that the individual lives longer than expected.

The middle two panels of Figure 6 show the decision rule for labor supply (or equivalently, for optimal choice of leisure). Similar to the retirement decision, the work decision takes the form of a threshold rule, where for any given income level (where if the person is retired, the y variable records the last earnings level in their tenured job), there is a threshold value of wealth such that if the person's wealth is greater than this level, they decide not to work at all. The left hand panel shows the labor supply decision for an individual who has not yet retired from their main tenured career job. In this case, they have only two alternatives: work full time or retire. The right hand panel of Figure 6 shows the labor supply decision rule for a person who has already retired from their main tenured career job. In this case the person has a choice of working part-time at a post-retirement job in addition to working full time or not working. We see that if wealth levels are sufficiently low, the person works full-time, and if wealth levels are sufficiently high the person does not work. The bottom two panels of Figure 6 plot the value functions V_{70} and V_{80} as a function of (w, y) holding the other arguments fixed ($r = 0$, $ss = 0$, $h = 0$). We see that both functions are concave in (w, y) and $V_{80} < V_{70}$. The concavity property implies diminishing marginal utility of wealth and income. The value functions decline with age t since the higher t is, the fewer the remaining number of periods in the person's life.

4 Simulating the Effects of Changes in Health Insurance

This section presents stochastic simulations of the ARM under a “base case” scenario and four alternative scenarios involving several commonly considered changes retiree health plans offered by colleges and universities. For each alternative scenario, I re-solve the dynamic programming problem and generate four corresponding simulations using optimal decision rules that represent the faculty member’s “best response” to each of the assumed changes in the retiree health insurance plan.

It is typically easier to understand the behavior predicted by the ARM by analyzing stochastic simulations rather than by studying its implied decision rules, as was done in the previous section. For example section 3.7 showed that the solution to the ARM consists of 4 optimal decision rules for consumption, labor supply, the retirement decision and the Social Security application decision. Each of these functions depends on 6 variables: age, wealth, income, a retirement indicator, a health indicator, and the person’s Social Security status. It is difficult to visualize these functions and see what they imply about the timing of retirement and how changes in health insurance coverage affects them. This section begins with an analysis of stochastic simulations under the base case in order to obtain further insights into the behavior implied by the ARM. Then I compute simulations under the 4 alternative scenarios and summarize how each change affects behavior, faculty welfare, and the university’s expected discounted cost of compensation relative to the base case.


A *stochastic simulation* of the ARM starts from an arbitrarily assigned *initial condition* specifying the state of the simulated individual at age 50 and then uses the laws of motion for the state variables and the decision rules calculated from the dynamic programming problem to generate a realized sequence of decisions and states over the individual’s lifetime until their death. The simulations depend on a set of *random shocks* which can be generated in advance and stored. If I re-use previously stored values of these random shocks (which are formally a $50 \times k$ matrix of uniformly distributed random numbers, where 50 is the maximum remaining lifespan of a 50 year old simulated individual and k is the number of randomly evolving state variables in the ARM), then I can re-generate the stochastic simulation and obtain exactly the same set of states and decisions that were previously generated. This provides a key method of ‘experimental control’

that is a major advantage of the ARM from the standpoint of evaluating the behavioral impacts of changes in retiree health insurance. By using the same random shocks for each of the four alternative scenarios that were used to generate the base case simulation, I am able to isolate the behavioral impacts of changes in health insurance coverage in each of the alternative scenarios from the effects of health, earnings, and mortality shocks. That is, if I use the same random shocks to generate the alternative scenarios, the individual will experience exactly the same realized values of the “exogenous variables” in the ARM. In particular, the sequence of health states and the age of death will be exactly the same in the base case simulation and in each of simulations of the alternative scenarios. Realized values of the endogenous state variables and decision variables such as earnings and the labor supply and retirement decisions will change, but the changes are a result purely of behavioral changes (e.g. changes in the decision rules) and not due to “luck” as reflected in different realized values of the underlying random shocks.

Clearly, it is not possible to attain this level of experimental control in an experiment with human subjects: the same person can only be given one “treatment” and it is not possible to observe how the same person would have behaved had they been given a different health insurance plan or if some other aspect of their fringe benefit package had changed. Thus, human experiments depend on randomly selected treatment and control groups with sufficiently large numbers that the effect of idiosyncratic shocks can be averaged out. I use this same type of experimental control in the stochastic simulations below. I generate large numbers of independent simulations for a synthetic population of individuals, starting either from either the same or a randomly generated set of initial conditions. Appealing to the Law of Large Numbers, I can average out the effects of the stochastic shocks to provide a clear view of the typical behavior implied under the base case and each of the alternative scenarios. In this analysis I use 500 independent simulations starting from the same set of initial conditions, yielding artificial “panel data sets” that follows 500 *ex ante* identical individuals from age 50 until their deaths.

Figure 7 illustrates the results of the base case simulations, 500 simulated individuals all of whom start from the same initial condition at age 50: good health, not retired or receiving Social Security with an income of \$64,857 and net worth of \$581,508.

Figure 7: Results of Simulations of the Model in the ‘Base Case’




The top two panels of figure 7 show the fraction of survivors and the fraction of survivors who are in good health status. Although the one year survival probabilities are high (at least 60% even at age 100), the cumulative risk of death between age 50 and 100 is substantial, and only 3% of the initial cohort survives to age 100. The fraction of survivors who are in good health generally declines until about age 87, and then it starts to rise. This partially reflects “survivorship bias” but also the averages beyond age 90 are noisier and should be interpreted with caution since they are based on relatively few observations given the small number of survivors at these ages.

The middle two panels of figure 7 show average consumption and earnings, respectively. Consumption declines steadily from \$80,000 at age 50 to just over \$20,000 at age 90. Interestingly, we observe a slight upturn in consumption between ages 65 and 70 and between age 95 and 99. Earnings grow at a near linear rate between age 50 and 55, and start to decline thereafter. As we will see shortly this is due to the fact that nobody in the simulation sample retires prior to age 55, but after that age the retirements are the main factor driving the decline in earnings. However due to post-retirement work, mean earnings at age 80 exceeds \$30,000 per year.

The bottom panels of figure 7 show the trajectories for the average of net worth and expected discounted utility for the sample of survivors. Net worth declines from \$580K at age 50 to just under \$100K at age 100. Wealth declines relatively rapidly between ages 50 and 70, then declines at a relatively slow rate between ages 70 and 85, and then resumes a more rapid rate of decline after age 85. Expected discounted utility declines between ages 50 and 100, reflecting the fewer remaining years of life and thus lower expected future utility as one ages.

The top two panels of figure 8 show the cumulative fraction of survivors who have retired and the fraction who are receiving Social Security benefits, respectively. As I noted previously, nobody in the simulation sample retires prior to age 55, but after this age retirements rise rapidly. However fewer than 50% of the survivors have retired by age 70, and it is not until age 85 when all surviving members of the simulation sample have retired. The middle two panels of figure 8 show the distribution of retirement ages and the distribution of ages of first eligibility for Social Security benefits, respectively. We see that over 60% of the simulation sample coordinate the timing of their retirement and application for Social Security benefits, with both occurring at age 70. The relatively late retirement ages of the faculty in this simulation sample illustrates a pressing concern of real academic institutions.

Figure 8: Results of Simulations in the ‘Base Case’




The top two panels of figure 9 show the fraction of survivors who choose to work full time and part time, respectively. Prior to age 55 100% of the survivors are working full time in their tenured job. There is a sharp drop in the fraction of individuals working full time right at age 55, when over 10% of the survivors choose to retire. We see a corresponding upward jump in the fraction of survivors working part time at age 55. Thus, it appears that there are a significant number of individuals who would like to work part time but do not have this option in their tenured job, so this motivates them to retire so that they can find part time employment. A significant number of these individuals have experienced a disability, which significantly increases their disutility of work. Part time employment rises rapidly between age 70 (when most individuals in the simulation sample choose to retire) and age 85, when part time employment peaks at a rate of over 60%. Part time employment rates decline rapidly after age 85, and by the early 90s only a small fraction of survivors continues to work on a full or part time basis.

These trends are mirrored in their complement, the “non-employment rate” shown in middle left hand panel of figure 9. Starting at age 55 the fraction not working relatively steady increases until age 80, when it reaches 40%. After age 85, as all remaining tenured faculty retire from the university and we pass the peak in part time work, the fraction of survivors who work increases dramatically, reaching virtually 100% by age 94.

The middle right hand panel of figure 9 shows mean and median out of pocket health care costs as a function of age. Median out of pocket health care costs rise from slightly over \$1000 per year at age 50 to about \$9000 per year by age 65 and remain relatively constant thereafter. Of course, this pattern is entirely a consequence of my assumption that health insurance costs will grow at 14% per year for the next 15 years and then flatten out thereafter.

The bottom two panels of figure 9 show health insurance premiums paid by the faculty member and the university, respectively. Under the base case scenario, the university picks up the cost of all health insurance both while employed and during retirement, so the only health insurance premium the faculty member has to pay is the Medicare Part B premium at age 65. These premiums have been growing at about 14% per year in recent years. Thus, I also assume that Medicare premiums will also grow at 14% per year for 15 years until flattening out. This means that the Medicare Part B premiums, which are currently about \$720 per year, are expected to grow to about \$7600 per year by the time the faculty member reaches age 65.

Figure 9: Results of Simulations in the ‘Base Case’



The bottom right hand panel of figure 9 shows the costs of health insurance from the university’s perspective. The blue line shows premiums for employees and the green line shows premiums for retirees. The blue line starts at \$9,000 per year and increases to nearly \$50,000 per year by age 65, reflecting my assumption that health premiums will increase at 14% per year for the next 15 years. However retirements cause the cost of health insurance premiums to faculty to fall steadily after age 65, reaching 0 at age 85 when all remaining faculty members have retired. The cost of retiree health insurance increases steadily after age 55, when faculty members start to retire. The average cost reaches \$72,000 per year at age 85 (when all faculty have retired) and remains constant thereafter. The costs of retiree health insurance rise steadily over a 30 year period, despite the assumption that premiums rise at 14% for 15 years and then flatten out. This is due to the fact that faculty members are retiring over a 30 year interval and university will have to pay the full \$72,000 annual premium for the latest of these retirees. However since this figure is an average over all survivors, it presents a slightly misleading forecast of the university’s costs of retiree health insurance, since it does not account for the cost savings due to cancellations of retiree health insurance benefits to faculty members who die. The discounted cost calculations that I will present shortly will explicitly account for the university’s cost-savings due to mortality.

Although the rapidly increasing health care premiums may somewhat exaggerate the expectations that many employers have about future health care costs, it does serve as a graphic illustration of the problems they confront. I now present an analysis of four scenarios, representing several common strategies that academic institutions and other employers have considered as a way to deal with rapidly increasing health care costs. I use the ARM to predict how faculty will rationally respond, in terms of adjusting their retirement behavior, in response to each scenario, and then simulate each scenario using the same random shocks and initial conditions that I used to generate the base case scenario. Table 4 summarizes the four scenarios that I consider:

Table 4: Summary of Policy Scenarios to be Simulated using ARM

- **Case 0:** The University pays only 50% of retiree health insurance premiums
- **Case 1:** The University cancels retiree health insurance completely.
- **Case 2:** The University continues retiree health insurance (and pays all premiums) up until age 65, when the faculty member is eligible for Medicare.
- **Case 3:** Same as age 2, except that I also assume that fairly priced high quality Medigap insurance is available for purchase in the private health insurance market.

**Figure 10: Comparisons of Base Case and Case 0:
Effects of Shifting 50% of Retiree Health Insurance Premiums to Retirees**





Figure 10 shows that when the University pays for only half of retiree health insurance benefits, there are major offsetting behavioral changes on the part of faculty members in response to what effectively amounts to a large cut in their lifetime compensation. The top left panel of figure 10 shows that the University appears to benefit from this cut since its cost of retiree health insurance benefits falls dramatically. The right hand panel shows the direct effect of this cut on the faculty member, resulting in a large increase in total health insurance premiums in Case 0 scenario compared to the base case scenario. Recall that under the base case, the faculty member only had to pay for their part B Medicare premiums, which were forecasted to be \$7,600 per year by the time they reach 65. However under Case 0 they will also have to pay for 50% of their retiree health insurance premiums, or approximately \$36,000 per year, so that total health insurance premiums will amount to nearly \$45,000 per year under this scenario.

The middle two panels of figure 10 illustrate the faculty member's key behavioral response to this large effective cut in compensation: they significantly delay their date of retirement in order to accumulate enough additional net worth to help cover these large additional health care costs. We see that whereas faculty started to retire after age 55 under the base case, they delay starting to retire by more than 10 years under the case 0 scenario, and the last retirees remain at the University until age 94 compared to 85 under the base case. The right hand middle panel of figure 10 shows that the policy change significantly increases the earnings the University has to pay to its faculty members as a result of the induced delay in retirements.

The bottom left hand panel of figure 10 shows that in addition to higher wage payments, the University will also have to incur significantly larger costs of its most costly fringe benefit: health insurance for its current employees. Since the University has not cut its contribution to health insurance for its existing employees, the delay in retirements entails a significant increase in health insurance premiums. The bottom right hand panel of figure 10 shows that mean out of pocket health care costs do not change significantly in the case 0 scenario. This is due to the fact that even though the faculty member must pay 50% of their retiree health insurance premiums, their overall level of coverage is assumed to remain the same. Out of pocket payments increase slightly due to a slower rate of decumulation of wealth under the Case 0 scenario, reducing the incidence of "bankruptcies" due to health costs that exceed the person's available net worth (recall that I assume that in this event, the person's liability is limited to his/her available net worth).

Figure 11: Comparisons of Base Case and Case 0
Effects of Shifting 50% of Retiree Health Insurance Premiums to Retirees




The top left hand panel of figure 11 shows effect of the policy change on average consumption. We see that there is a significant, approximately uniform drop in annual consumption of about \$15,000 per year. The top right hand panel of figure 11 shows that there is a significantly lower rate of decumulation of wealth under Case 0. Thus, via a combination of reduced consumption and a delay in retirement, the faculty member is able to partially offset the effect of the cut in retiree health benefits. The increase in net worth reflects the increased reliance on “self insurance” due having to incur significant costs to maintain retiree health benefits.

The middle two panels of figure 11 show the impact of the policy change on full and part time work, both pre and post retirement. In addition to delaying retirement, we see a significantly increased incidence of full time work after retirement, as well as a shift to doing part time work even at very late ages in life. To some extent this increase in labor supply is an artifact of the assumption that the retiree has no choice but to pay the large annual retiree health insurance premiums. If I were to extend the ARM to allow a decision over whether or not to maintain retiree health insurance coverage, I would expect to see older individuals dropping their coverage and reducing their labor supply late in life (when there is a significantly higher disutility of work), relying only on their Medicare coverage.

The bottom two panels of figure 11 show the impact of the policy change on total earnings (earnings on the main career job plus earnings in post-retirement jobs) and the impact on expected discounted utility. We see that this change has a significant, very negative effect on the lifetime welfare of faculty members, and for this reason we could expect that the option of cutting the University’s contribution to health insurance premiums would result in active opposition by both current faculty and retired faculty. I will defer a discussion of whether this could actually save costs for the University until I have presented the analysis of the remaining 3 scenarios.

Figures 12 and 13 present a comparison of the base case and an even more extreme alternative that the University might consider: cancelling its retiree health insurance altogether. The top left panel of figure 12 shows the major benefit to the University from this policy: it completely eliminates the cost of retiree health insurance. The right hand panel of figure 12 shows that this also reduces the premiums the retiree will have to pay since I assume that fairly priced private retiree health care coverage is not available unless provided through a University-sponsored plan. The retiree only has to pay the \$7,600 annual Medicare part B premiums, same as the base case.

**Figure 12: Comparisons of Base Case and Case 1:
Effects of Cancelling Retiree Health Insurance**




The remaining panels of figure 12 show that a cancellation of retiree health benefits will have much the same effect on retirements as was observed in case 0: there is a significant delay in the retirement age and this delay causes the university to have to pay out more in earnings and in health insurance premiums for its current employees. The bottom right panel of figure 12 shows that mean out of pocket health care costs increase significantly in the case 1 scenario, and of course this is due to the cancellation of the retiree health benefits,

The top two panels of figure 13 show the effect of the cancellation of retiree health insurance on consumption and wealth. We see that there is a modest reduction in consumption between ages 55 and 65, but afterwards mean consumption is actually slightly higher under the Case 1 scenario. This is due in part to the higher earnings and wealth accumulation. Net worth is significantly higher at nearly all ages, and this is clearly due to the need to maintain higher precautionary savings to cover medical expenses that are not covered under Medicare parts A and B.


The remaining panels of figure 13 show that there is significantly higher full time work under the case 1 scenario, but most of this appears to be a result of the delay in retirement from the University. There is relatively little increase in the amount of full time work after retirement, and the amount of part time work after retirement is generally lower as well. These effects reinforce the conclusion that the behavioral effects predicted under case 0 may be artificially high due to the failure to give the retiree the option to discontinue their retiree health insurance coverage — in effect forcing them to pay a premium that is much higher than they would be willing to pay voluntarily. The unrealistically high levels of labor supply offered at very old ages in the case 1 scenario seems to be an artifact of the need to pay for the incredibly high cost of retiree health insurance. The bottom right panel of figure 13 shows that cancellation of retiree health benefits also uniformly reduces the faculty member’s welfare relative to the base case. In principal it is possible to compute the faculty member’s “willingness to pay” to avoid the cut in retiree health benefits, either via a lump sum payment from their DC pension accumulations or via a reduction in their salary over the course of their employment at the University.

Figures 14 and 15 show the simulated effects of the Case 2 “compromise scenario” in which the University continues retiree health insurance and pays 100% of the premiums, but only until the person reaches age 65 and becomes eligible for Medicare benefits. This option may be increasingly attractive in light of a recent decision by the Equal Employment Opportunity Commission.

Figure 13: Comparisons of Base Case and Case 1
Effects of Cancelling Retiree Health Insurance



**Figure 14: Comparisons of Base Case and Case 2:
Effects of Cancelling Retiree Health Insurance after Age 65**



The EEOC ruled that an employer's decision to discontinue retiree health insurance after age 65 does not violate civil rights laws banning age discrimination. This ruling appears to be a pragmatic compromise to confronting employers with an all or nothing choice of "Give all of your retirees the exact same benefits, which is incredibly difficult, or eliminate your retiree health benefits altogether." (Leo Silverman, member of the EEOC, quoted in *New York Times* front page article, April 23, 2004). For this reason, the New York Times article reported that many labor organizations, including the American Federation of Teachers and the National Education Association supported the EEOC ruling, supported the EEOC ruling.

The top left hand panel of figure 14 shows that compared to the base case, the compromise policy of offering full, employer-paid retiree health insurance only up to age 65 results in a major reduction in retiree health insurance costs to the University. The right hand panel shows that the retiree's health insurance premiums remain the same as in the base case, i.e. the retiree only has to pay for their Medicare Part B premiums, since I assume that the University covers the retiree health insurance premiums prior to age 65, and after age 65 the retiree has no affordable "Medigap option" to supplement their Medicare coverage.

The middle two panels of figure 14 show that the cut in retiree health benefits after age 65 does result in some delay in retirements, but the effect is not nearly as large as in Case 0 or Case 1. We see this also in the right hand panel of figure 14, where the University also pays higher earnings due to delayed retirement, but the shift in expected earnings is not nearly as large as it was in cases 0 or 1.

The bottom two panels of figure 14 show that under the employer can save on benefits paid not only to its retirees, but it can also save on health insurance benefits to its current employees but discontinuing employee health insurance coverage at 65 in addition to discontinuing retiree health insurance. This policy has an added benefit of not appearing to discriminate against retirees in favor of current employees, and has the desirable effect of mitigating incentives to delay retirement in order to maintain the deluxe health care coverage that the University provides to its employees under the base case. The bottom right hand panel of figure 14 shows that the retiree's mean out of pocket health care expenses is significant higher after age 65 under Case 2 compared to the base case, and of course these higher costs are due to the fact that the retiree no longer has the more deluxe retiree health care coverage to supplement their Medicare Part A and B coverage.

Figure 15: Comparisons of Base Case and Case 2
Effects of Cancelling Retiree Health Insurance after Age 65




Figure 15 shows the effect of the Case 2 health coverage changes on consumption, wealth accumulation, labor supply and welfare. We see that prior to age 65 the faculty member increases his/her precautionary savings balances by a combination of delaying retirement and reducing consumption. Indeed, the incidence of full-time work is uniformly higher both before and after age 65 in case 2 relative to the base case. The effect on part time work is more complicated: it is lower up to age 60, higher between ages 62 and 70, significantly lower between ages 70 and 85, and significantly higher thereafter. Overall, we see that earnings are significantly higher at all ages under Case 2, and this causes the faculty member's net worth to be uniformly higher at all ages and decumulate more slowly compared to the base case. The bottom right panel of figure 15 shows that the faculty member's welfare is uniformly lower under Case 2, but the change in welfare between ages 50 and 65 is very small so that at these ages the faculty member is approximately indifferent (or at least would not vigorously oppose) the change in the health insurance coverage. This seems consistent with the National Education Association's endorsement of the EEOC's decision to allow employer's to terminate retiree health care coverage at age 65: relative to the alternative of cutting benefits completely or shifting a large share of the premiums onto retirees, the case 2 policy clearly seems to be the "lesser of evils" from a faculty member's standpoint.

Figures 16 and 17 present the simulation results for the final case, Case 3. This case is the same as Case 2 except that I assume that there is fairly priced Medigap insurance available for individuals over 65. I assume this Medigap coverage is twice as expensive as the Medicare part B premium, or \$15,200 per year and is paid completely by the retiree. The top right hand panel of figure 16 shows that the University's retiree health insurance benefits hardly change under Case 3, but the right hand panel shows that the average premiums paid by retirees do increase significantly as a result of this Medigap insurance.

The middle two panels of figure 16 show that similar to case 2, faculty members delay retirements relative to the base case, resulting in an increase in average earnings paid at each age. The bottom right panel of figure 16 shows that out of pocket health care costs are higher under case 3 than the base case, but as will be more evident shortly, mean out of pocket health care costs are significantly lower in case 3 than in case 2 due to the availability of the Medigap insurance.

Figure 17 presents comparisons of consumption, wealth, labor supply, earnings and welfare for Case 3 relative to the base case. The results are quite close to the results for Case 2.

**Figure 16: Comparisons of Base Case and Case 3:
Effects of Cancelling RHI after Age 65 if fairly priced Medigap coverage exists**





Figure 17 Comparisons of Base Case and Case 2
Effects of Cancelling RHI after Age 65 if fairly priced Medigap coverage exists



I conclude the analysis of the simulation results with figures 18 and 19 which compare the outcomes in Cases 2 and 3, making it easier to see how the existence of fairly priced Medigap insurance complements the compromise policy of eliminating retiree health insurance coverage after age 65. The availability of fairly priced Medigap insurance has a reassuring effect, and causes faculty members to be willing to retire slightly earlier than they would without Medigap. The reason for this can be seen in the bottom right panel of figure 18, which shows that with Medigap, mean out of pocket health care expenditures are significantly lower. Although the average drop in out of pocket expenditures is generally lower than the annual Medigap premium, \$15,200, these risk averse faculty members prefer to purchase Medigap due to the value of insurance it provides that reduces the risk of catastrophic health care costs in old age.

Figure 19 compares mean consumption, wealth, labor supply, earnings and welfare trajectories under Cases 2 and 3. We see that even though mean consumption is slightly lower after age 75 under Case 3 than under Case 2 (due to the payment of Medigap premiums), and wealth accumulation is significantly lower at all ages over 60 (due to a reduced need to accumulated precautionary savings balances due to the Medigap insurance), the lower right hand panel of figure 19 shows that average welfare is slightly higher after age 65 under Case 3, a result that is again entirely attributable to the existence of fairly priced Medigap insurance.

Thus, although the differences between Cases 2 and 3 are relatively subtle, the ARM is able to provide intuitively plausible forecasts of how relatively subtle changes in the environment can affect individual behavior and welfare. In general, the more complete and more fairly priced options that exist for purchasing health insurance in the private market, or from an expansion of government insurance programs, the better it is from the standpoint of both the employee and the employer. As we will see shortly, the employer can take advantage of the insurance opportunities available in the market (or from the government) to avoid the costs involved in providing these opportunities to its employees and retirees on its own.

I conclude this section with Table 5, which compares the expected discounted costs of compensation (both in terms of direct wages and fringe benefits) under the Base case and the four alternative scenarios. Panel A presents the expected discounted values of the various components of compensation, discounted at a real interest rate of 2%. These calculations take account of mortality, and in particular, all benefit payments terminate at death.

Figure 18: Comparisons of Case 2 and Case 3:
Effects of Cancelling RHI after Age 65 with/without fairly priced Medigap coverage




Figure 19: Comparisons of Case 2 and Case 3

Effects of Cancelling RHI after Age 65 with/without fairly priced Medigap coverage




Table 5: Expected Discounted Compensation Under Alternative Scenarios

Panel A: Expected Discounted Values (at 2%) (amounts in thousands)					
Compensation Item	Base Case	Case 0	Case 1	Case 2	Case 3
Wages	1047.17	1497.40	1370.9	1222.48	1183.06
Pension Contributions	52.36	74.87	68.6	61.12	59.15
Employee Health Insurance	537.55	863.41	764.3	303.43	294.82
Retiree Health Insurance	553.13	113.64	0.0	33.79	42.40
Total Compensation	2190.21	2549.31	2203.81	1620.83	1579.43

Panel B: Breakdown as a Percent of Total Expected Discounted Compensation					
Compensation Item	Base Case	Case 0	Case 1	Case 2	Case 3
Wages	47.8	58.7	62.2	75.4	74.9
Pension Contributions	2.4	2.9	3.1	3.8	3.8
Employee Health Insurance	24.5	33.9	34.7	18.7	18.7
Retiree Health Insurance	25.3	4.5	0.0	2.1	2.7
Total Compensation	100.0	100.0	100.0	100.0	100.0

According to Table 5, of the various scenarios considered, the Case 3 scenario is the cheapest from the standpoint of the University, whereas the most expensive is Case 0, which involves shifting 50% of the cost of retiree health insurance onto the faculty retiree. As I noted above, due to the rapid rise in health insurance premiums to perhaps implausibly high levels, and my implicit assumption that retirees do not have the option to cancel this expensive retiree health insurance, the forecasts under the Case 0 scenario may be somewhat distorted relative to the outcomes that would occur if retirees could cancel their retiree health insurance. If this latter option were allowed, most older retirees would cancel their retiree health insurance rather than pay the high premiums, and the outcomes of Case 0 would be very similar to the outcomes in Case 1, where the University cancels retiree health insurance completely.

However the interesting point to note is that the ARM predicts that it is cheaper for the University to maintain its deluxe employee and retiree health insurance package than to adopt a naive policy of trying to cut costs by cutting these benefits completely. This is due to the offsetting increase in wage and employee health insurance benefits under either Case 0 or Case 1 that is due to the faculty member's optimal behavioral response to the University's action. The failure to account for these behavioral responses can lead to ill-advised policies that are not only greatly unpopular with existing and retired faculty, but they can actually end up costing the University

more. These results support a general point that I made in the introduction: *if workers are risk averse, it is cheaper to offer them a compensation package that includes insurance for their most pressing risks rather than to force the employee to “self insure” against these risks.*

However the other major conclusion from this analysis is that when the cost of providing insurance grows too fast and gets too high, there is a point at which some “compromise” cutbacks in health insurance benefits become inevitable. I have illustrated a case where the University can reduce the total costs of its compensation by judiciously cutting its retiree health benefits, taking advantage of the existence of a government program — Medicare — that already provides a good deal of coverage for individuals who are over 65. The policy of cutting retiree health benefits after age 65 was the cheapest among the options considered. Furthermore, the results indicated that when faculty expect to have access to fairly priced Medigap insurance to supplement their Medicare coverage, the savings of the policy of eliminating retiree health at age 65 are even higher. This is due to the fact that the increased security level provided by the option to purchase fairly priced Medigap insurance induces faculty members to retire earlier than they would in the absence of fairly priced Medigap, and this saves the University money.

These results suggest that the recent passage of Medicare prescription drug benefit will have a similar effect: it will induce employers (including Universities) to drop prescription drug coverage from their retiree health insurance benefits (at least for those over age 65), and effectively shift their employees onto the Medicare drug plan. While these actions are good from the standpoint of employers, if the financing of Medicare is not sustainable in the long run (and the most recent Report of the Trustees of the Social Security System forecasts that the prescription drug benefit will add \$8 trillion dollars in unfunded obligations to the \$27 trillion in unfunded obligations of the pre-existing Medicare program over a 75 year horizon), the “fix” provided by such policies may only be temporary and illusory. What is required is an intelligent, sustainable, long term solution to the problem of runaway health care costs in the United States, and neither employees nor employers can really rest easy until such a solution can be found.

5 Conclusions

This paper has developed a model of faculty retirement decisions by extending and specializing the classical “life cycle model” in economics to account for the presence of uncertain out of pocket health care expenditures and the specifics of the academic job contract and typical fringe benefit packages. I refer to this model by its acronym — ARM, for “Academic Retirement Model”. Via solutions and simulations of an informally calibrated version of ARM I have tried to show how such a model can generate “realistic” predictions, and can serve as an important tool to help a University make intelligent decisions about the compensation packages it offers to its faculty.

The ARM has a great deal of flexibility and can be used to predict the impact of a wide range of policy changes that were not considered here, such as the use of “phased retirement programs” and lump sum “buyouts” that are the subject of empirical analyses by Steven Allen and John Pencavel in this volume. More ambitiously, it is possible to use the ARM to systematically optimize over entire classes of compensation packages, in order to identify an “efficient” (i.e. cost-minimizing) compensation package. However as I will discuss below, it would be advisable to extend the ARM to include indicators of faculty “outputs”, rather than relying on the current model which has a near exclusive focus on the faculty’s decision problem.

However before attempting to extend the model in these more ambitious directions, I would place the highest priority on collecting data on retirement decisions at specific universities that could enable me (or others) to estimate the unknown preference parameters of the ARM and statistically test its predictions. The credibility of the model rests on its ability to accurately predict retirement decisions by faculty in the past, and its ability to accurately predict how faculty will change their behavior in response to the changes in health insurance coverage and other aspects of compensation packages. An ideal data set would come from a large university (such as one of the large state systems such as University of North Carolina or California) for which we have large numbers of “observations” on faculty retirements and where there have been interesting changes to compensation plans that provide a testing ground that can enable us to compare actual response to changes in the compensation plan to those predicted by ARM.

I will conclude by commenting on several larger and more general issues about health care coverage that cannot be easily addressed by this simple initial version of ARM.

An important caveat to the principle that it is cheaper to compensate risk averse employees via provision of insurance is the possibility of *moral hazard* which can result in “excessive use” of medical care services. When moral hazard exists it may not necessarily be cheaper to compensate an individual with a significant component of the package being in the form of the usual textbook type of “full insurance” (i.e. a policy that covers all risks in exchange for a premium that equals the expected value of the cost of benefit payments). By exposing the employee to some risk (i.e. forcing the employee to “co-insure”) the employer can reduce the level of moral hazard and thus reduce the costs of providing excessive levels of medical services.

Moral hazard is typically thought to arise at the individual level (i.e. from an individual insured’s “overuse” of medical care), but in my view, a vastly more important source of moral hazard is at the aggregate level, i.e. from doctors, hospitals and pharmaceutical companies that “overcharge” from their services, taking advantage of poorly designed health insurance contracts. In the U.S., it is my belief that most of the problem of runaway inflation in health care costs is due to moral hazard “in the aggregate,” and this problem results from a lack of adequate cost-control measures in overly generous health insurance plans (both private and government-provided such as Medicare). The lack of adequate cost control procedures has, in my opinion, greatly exacerbated the high degree of inflation in health care costs. For example, a number of leading experts predict that the recent Medicare prescription drug bill will turn out another example of “aggregate moral hazard” resulting from poorly designed health care policies. These health care costs must be covered in higher premiums, and have contributed to the current situation where the cost of health insurance is spiraling out of control.

My view that “overpricing” of health care the main source of the “health care problem” facing colleges and universities (and other firms) in the U.S. is consistent with the view in Joseph Newhouse’s (2002) book *Pricing the Priceless*. Newhouse notes “Why the United States saw little or no price competition over medical services for several decades is puzzling. The proximate explanation is that a passive, self-insured employer generally paid the premium for the insurance policy on the behalf of the employee, or paid a large percentage of it, but that explanation, of course, begs the question of why the employer was passive. Employers could potentially have bargained directly with providers for lower fees, but they faced large transactions costs to do so and until relatively recently, they did not do so. Moreover, as long as employers sought to provide

their employees with an insurance policy that covered all or almost all providers, they had little bargaining power.” (Newhouse, p. 10).

However many larger universities and colleges that have medical schools or large on-campus health clinics do have the ability to self-insure by creating their own on-campus HMOs and health care plans in order to take more direct control over their health care costs and utilization. In fact, the CUPA survey showed that in 2002 over 36% of colleges and universities opted for the self-insurance strategy. However self-insurance is an option that is really only feasible for the larger public and private universities that have a medical school and an associated university hospital which receives other sources of government subsidies and has sufficient “scale economies” that cost-saving via self-insurance can be realized after accounting for the substantial fixed costs of creating and administering their own health care plans. By in large, most small and medium scale liberal arts colleges do not have the resources to self-insure and are dependent on contracts with private health insurers. Even for institutions that are big enough to self-insure, it is not always clear that doing so will necessarily result in dramatic reductions in health care costs. A college or university may not have a “comparative advantage” in controlling costs and administering benefits compared to health insurance companies that have substantial experience in the area. “The hefty start-up costs associated with the formation of the Florida colleges’ self-insured plan were just one of the reasons that the University of Denver decided about six months ago that switching to self-insurance would be a bad move.” (Chronicle of Higher Education, May 2, 2003).

In summary, while self-insurance may be one option that can enable the largest colleges and universities to manage the moral hazard problems associated with the provision of medical care to their faculty and staff, self-insurance involves substantial fixed costs and probably is not feasible option for the majority of smaller colleges and universities that do not have major on-campus medical schools and medical centers.

The second general problem associated with health insurance and employer retention and retirement policies is *adverse selection*. In insurance contexts, adverse selection occurs when there is unobservable (or uncontractable) heterogeneity in characteristics of insureds which cannot be accounted via “price discrimination” in insurance premiums. In the case of health insurance, people who are less healthy or who have genetic susceptibilities to certain diseases such as Alzheimers, etc. are the “bad apples” which health insurers would like to exclude from coverage or charge

higher premiums. Instead, insurers can typically only charge a single average premium, which in order to be profitable, must exceed the average cost of benefits provided to the healthier insureds in order to cover the losses incurred on the unhealthy insureds. If the healthier insureds leave the plan and join an alternative plan with lower premiums, then the insurance company must drive up premiums to avoid taking losses. In extreme cases the adverse selection problem can lead to health care costs to spiral out of control and the insurance policy to collapse.

There is no doubt that adverse selection is a significant problem that is contributing to the rapid rise in health care costs, and is largely responsible for the relative lack of fairly priced private health insurance coverage in the U.S. However since faculty members, due in part to their higher levels of education, tend to be healthier than the U.S. population at large, a University can use adverse selection to its advantage by negotiating *group insurance plans* that offer much lower premiums to faculty members than are available to the general population. A university or college would be foolhardy not to attempt to find a group health insurance plan for its faculty at premiums that are far more advantageous than what individuals could obtain on their own.

The problem of adverse selection can also occurs in an employment context, especially for firm that are limited in their ability to fire workers or adjust their pay to reflect differential abilities and productivities. Even when the firm can “observe” which workers are productive and which aren’t, the firm may not be able to fire workers who are less productive, or pay them a lower amount to reflect their lower productivity. This is especially acute for colleges and universities, since the tenure contract generally prevents it from firing its least productive tenured faculty members. Further, compensation policies at many state schools and private universities place limits on their ability to penalize unproductive faculty via wage cuts. If nominal wages cannot be decreased, then these institutions are restricted to a much slower process of real wage reductions in periods of high inflation. Given these constraints on compensation policies, higher educational institutions have tried to design retirement packages that create incentives for the “bad apples” to retire early, while retaining the “good apples.” However there are examples where the effect is exactly the opposite of what was intended: many early retirement incentive plans have enabled the “good apples” to leave the university, collect generous bonuses in addition to their basic pensions, and then “double dip” by being hired at a competing college or university. The “bad apples” are typically unable to get attractive job offers at other schools, and unless the retirement

incentives are sufficiently generous, these individuals prefer to keep their tenure rather than retire for temporary bonuses and then face the subsequent risks of health care costs and wage volatility during retirement.

This problem is not unique to academics: similar problems have occurred in the military, which has also used pension incentives to try to retain its best personnel. However here too there are examples of poorly designed early retirement schemes that induce the best people to take the early retirement option and take a more lucrative job in the private sector. This problem is especially acute with Air Force pilots, the best of whom significantly increase their earnings by accepting a job at a private airline (at least this was true up until the problems faced by the airline industry in the aftermath of the 9/11 terrorist incident).

If it is the case that faculty who are less productive are also more likely to be less healthy, the problems resulting from mistakes in the design of faculty retirement packages can be especially severe, resulting from a compound form of adverse selection: the more healthy and more productive older faculty are induced to leave the university leaving behind less productive and less healthy faculty members with compound consequences from lower educational “output” and higher costs of health care benefits.

In order to adequately address the full range of complexities involved in designing compensation packages, the ARM needs to be extended to include measures of a faculty member’s level of “effort” as well as the “output” produced by faculty member and the values of these outputs to the University. Neither effort nor output are likely to be well measured by simple uni-dimensional measures, and this is part of the complexities involved in extending the economic analysis provided in this paper to a deeper level.

In general, it is not necessarily the case that the University should try to encourage all older faculty to retire in order to make room for younger faculty. Instead the university should try to design compensation packages that are *incentive compatible* and which creates incentives for the least productive faculty to *self-select* and retire, while enabling the university to retain its most distinguished and productive older faculty.

Furthermore, there are a whole range of issues related to the job duties of faculty, and how these duties might change over time to best adapt to their abilities and interests, and whether there is scope for policies such as sabbaticals designed to help the best older faculty “retool” and

keep up with the state of the art. By manipulating these “non-pecuniary” aspects of the job, the university can have a major impact on the “disutility of work” that is a critically important determinant of retirement behavior in the ARM.

Hopefully these general comments offer many directions for readers who might be interested in furthering research on this topic. But most importantly, I hope this paper convinces readers that the design of academic retirement packages is an important problem with significant implications for higher education in the U.S. I hope it also clear why decisions about health insurance coverage — both while employed and during retirement — is an integral part of this problem.

References

- Allen, S.G. (2004) “The Value of Phased Retirement” forthcoming in R. C. Clark (ed.) *Recruitment, Retention and Retirement: The Three R’s of Higher Education in the 21st Century* Edward Elgar Publishers.
- Allen, S. G. R. L. Clark and L.S. Ghent (2003) “Phasing Into Retirement” NBER working paper 9779.
- Ashenfelter, O. and D. Card (2001) “Did the Elimination of Mandatory Retirement Affect Faculty Retirement Flows?” NBER working paper 8378.
- Bazzoli, G.J. (1985): “The Early Retirement Decision: New Empirical Evidence on the Influence of Health,” *Journal of Human Resources*, **20** 214–234.
- Benítez-Silva, H. (1999): “Micro Determinants of Labor Force Status Among Older Americans,” manuscript, Yale University.
- Benítez-Silva, H., M. Buchinsky, and, J. Rust (2001): “Dynamic Structural Models of Retirement and Disability,” manuscript, Yale University.
- Benítez-Silva, H., M. Buchinsky, H-M. Chan, J. Rust, and S. Sheidvasser (1999): “An Empirical Analysis of the Social Security Disability Application, Appeal and Award Process,” *Labour Economics* **6** 147–178.
- Benítez-Silva, H., M. Buchinsky, and J. Rust (2003): “How Large are the Classification Errors in the Social Security Disability Award Process?” manuscript, University of Maryland.
- Blau, D.M. and D.B. Gilleskie (2003) “Health Insurance and the Retirement Behavior of Married Couples” manuscript, University of North Carolina.
- Blau, D.M., D.B. Gilleskie, and C. Slusher (1997): “The Effect of Health on Employment Transitions of Older Men,” manuscript, Department of Economics, University of North Carolina at Chapel Hill.
- Boz, E. Benítez-Silva, H. Buchinsky, M. Nichols, J. Roy, S. Rust, J. and I. Tristao (2004) “The Relationship Between Health Insurance, Health Care Expenditures and Health Status: An Analysis Using the Health and Retirement Survey” manuscript, University of Maryland.
- Chatterjee, K. and R.C. Marshall (2001) “A Model of Academic Tenure” working paper, Pennsylvania State University.
- Ehrenberg, R.G. P. Pieper, and R. Willis (1998) “Do Economics Departments with Lower Tenure Probabilities Pay Higher Faculty Salaries?” *Review of Economics and Statistics* **80-4** 503–512.
- Ehrenberg, R.G. (2002) “Studying Ourselves: The Academic Labor Market” NBER working paper 8965.
- French, E. and J.B. Jones (2004) “The Effects of Health Insurance and Self-Insurance on Retirement Behavior” manuscript, Federal Reserve Bank of Chicago.

- Gotz, Glenn A. and John J. McCall (1984) "A Dynamic Retention Model for Air Force Officers" RAND project Air Force report R-3028-AF, RAND Corporation, Santa Monica, California.
- Hopenhayn, H. and J.P. Nicolini (1997) "Optimal Unemployment Insurance" *Journal of Political Economy* **105-2** 412-438.
- June, A.W. (2004) "Medicare Law May Provide a Lift to College Health-Care Budgets" *Chronicle of Higher Education* February 27, 2004.
- Khwaja, A. (2003) "How would Men Fare without Medicare? A Simulation-Based Evaluation" manuscript, Duke University.
- Li, S. and H. Ou-Yang (2003) "Incentives, Performance and Academic Tenure" manuscript, Duke University.
- Lozier, G. Gregory and Michael J. Dooris (1991) "Projecting Faculty Retirement: Factors Influencing Individual Decisions" *American Economic Review* **81-2** 101-105.
- Lumsdaine, R., James H. Stock and David A. Wise (1992) "Three Models of Retirement: Computational Complexity versus Predictive Validity" in D.A. Wise (ed.) *Topics in the Economics of Aging* University of Chicago Press.
- Retirement Rates So High at Age 65?" in D.A. Wise (ed.) *Advances in the Economics of Aging* University of Chicago Press.
- Lumsdaine, R., James H. Stock and David A. Wise (1996) "Why are Retirement Rates So High at Age 65?" in D.A. Wise (ed.) *Advances in the Economics of Aging* University of Chicago Press.
- Newhouse, Joseph and the Insurance Experiment Group (1993) *Free for All? Lessons from the RAND Health Insurance Experiment* Cambridge, Harvard University Press.
- Newhouse, Joseph (2002) *Pricing the Priceless A Health Care Conundrum* MIT Press.
- Oster, S.M. and D.S. Hamermesh (1998) "Aging and Productivity Among Economists" *Review of Economics and Statistics* 154-156.
- Pencavel, J. (2004) "Faculty Retirement Incentives by Colleges and Universities" forthcoming in R. C. Clark (ed.) *Recruitment, Retention and Retirement: The Three R's of Higher Education in the 21st Century* Edward Elgar Publishers.
- Royalty, Ann Beeson and J. Hagens (2003) "The Effect of Premiums on the Decision to Participate in Health Insurance and Other Fringe Benefits Offered by the Employer: Evidence from a Real-World Experiment" ERIU Working paper 23.
- Rust, J. (1996) "Numerical Dynamic Programming in Economics" in H. Amman, D. Kendrick and J. Rust (eds.) *Handbook of Computational Economics* volume 1, Elsevier, North Holland,
- Rust, J. and C. Phelan (1997): "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets," *Econometrica*, **65-4** 781-831.

- Schieber, Sylvester S. (2004) "The Future of Retiree Health Benefits in Higher Education in the United States" forthcoming in R. C. Clark (ed.) *Recruitment, Retention and Retirement: The Three R's of Higher Education in the 21st Century* Edward Elgar Publishers.
- Stock, J. Wise, D. (1990) "Pensions, The Option Value of Work and Retirement" *Econometrica* **58-5**, 1151-1180.