

Contracting with limited commitment: evidence from employment-based health insurance contracts

Keith J. Crocker*

and

John R. Moran**

Impediments to worker mobility serve to mitigate the attrition of healthy individuals from employer-sponsored insurance pools, thereby creating a de facto commitment mechanism that allows for more complete insurance of health risks than would be possible in the absence of such frictions. Using data on health insurance contracts obtained from the 1987 National Medical Expenditure Survey, we find that the quantity of insurance provided is positively related to the degree of worker commitment. These results illustrate the importance of commitment in the design of long-term contracts, and provide an additional rationale for the bundling of health insurance with employment.

1. Introduction

■ The recent debate over the use of annual re-underwriting in the context of individual health insurance policies highlights the difficulties that insurers face when attempting to craft a stable insurance pool.

Shaneen and Tom Wahl were paying \$417 a month for health insurance when Mrs. Wahl was diagnosed with breast cancer in 1996. Their premiums began rising steadily, and by August 2000, the Wahls were told that their new rate would be \$1,881 a month. Mrs. Wahl, whose cancer was in remission, tried to find out why. Unsatisfied with answers they got on the phone, the . . . couple visited the . . . offices of the insurer. There . . . an executive explained why her premium was soaring: "because of your dread disease" . . . It's called reunderwriting . . . A key challenge in individual health insurance is keeping the healthier people enrolled. Their premiums are needed to subsidize those who get sick. The problem, under the traditional approach, is that the healthiest tend to drop out as rising medical costs gradually drive up premiums . . .

* The Pennsylvania State University; kcrocker@psu.edu.

** Syracuse University; jmoran@maxwell.syr.edu.

The authors gratefully acknowledge support from the National Science Foundation under grant no. SBR-9696260 and the Management Science Group of the Department of Veterans Affairs for their sponsorship of the 13th Annual Health Economics Conference (2002, Carnegie Mellon). This article benefited significantly from the comments of two anonymous referees and the Editor, Martin Gaynor. We also thank Dan Black, Kate Bundorf, Richard Butler, Michael Chernow, Neil Doherty, David Dranove, Bill Encinosa, Gary Engelhardt, William Gentry, Paul Gertler, Jon Gruber, Richard Hirth, Jeff Kubik, Mark Roberts, Lowell Taylor, and participants at the 13th Annual Health Economics Conference for helpful discussions. A portion of this research was conducted while the second author was a Fellow in the Robert Wood Johnson Foundation's Scholars in Health Policy Research Program.

Advocates of reunderwriting argue that because it gives smaller rate increases to the healthy, it keeps more of them enrolled and paying premiums. . .

Chad Terhune, *Wall Street Journal*, April 9, 2002, p. A1

In a setting where an individual's health status is observable but evolving over time, the key to maintaining a successful insurance arrangement is to have healthier members of the group commit to cross subsidize members who experience adverse health outcomes. In the absence of a precommitment mechanism, however, those who turn out to be healthier than average may find it advantageous to leave the risk pool and to purchase alternative coverage at a premium that more accurately reflects their own, observable, health status.¹ This erodes the actuarial integrity of the pool and, consequently, the ability of insurers to offer insurance against both the financial losses associated with a particular illness episode and the classification risk that results when insurance premiums are adjusted to reflect changes in one's health status over time.

This article examines a theoretical model of insurance contract design in which impediments to worker mobility may serve to mitigate the attrition of healthy individuals from employer-sponsored insurance pools, thereby creating a de facto commitment mechanism that allows for more complete insurance of health risks than would be possible in the absence of such frictions. We characterize optimal health insurance contracts in an environment where a worker's evolving health status is publicly observable, and in which employees vary in the transactions costs that they must incur when switching to an alternative employer. Using data on employment-based health insurance plans, we find that the structure of insurance contracts observed in practice is consistent with the predictions of the model, providing empirical documentation for the importance of precommitment in the design of health insurance contracts.

Traditional analyses of contracting in insurance markets have tended to emphasize the role of informational asymmetries, which engender problems of adverse selection (Rothschild and Stiglitz, 1976; Crocker and Snow, 1986) or moral hazard (Shavell, 1979). A recent article by Cardon and Hendel (2001), however, finds no evidence of adverse selection in health insurance markets, suggesting instead that the lack of commitment in long-term contracts examined by Cochrane (1995) may be a more relevant source of market failure in health insurance settings. Accordingly, we consider a full-information model in which insureds are initially identical but anticipate receiving a public signal that will provide information about their expected future health care costs. As a consequence, the risk-averse purchasers of health insurance face two potential sources of uncertainty: the *financial loss* associated with adverse health outcomes and the *classification risk* that results when their future insurance premiums reflect the publicly available information on each individual's likely health-related expenses.

In such a setting, an optimal insurance contract would entail full compensation for all the losses suffered from illness as well as a constant, *ex ante* actuarially fair, premium charged to all individuals independently of the state or their observed health status. Although such a package provides full insurance against both types of uncertainty, the cross subsidy from low- to high-risk individuals in the future, which is inherent in the non-experience-rated premiums required to cover the classification risk, gives those with lower expected health care costs the incentive to exit the insurance pool and to purchase independent coverage at a price that more accurately reflects their own, observable, expected health care costs. As others have noted (Cochrane, 1995; Pauly, Kunreuther, and Hirth, 1995), it is this inability of insureds to precommit to remain in the pool in light of favorable information on health status that provides an impediment to the insurability of classification risk.

Paradoxically, the ability to insure against such risks is enhanced if there exist obstacles to the mobility of insureds that mitigate the erosion of the insurance pool caused by the departure of lower-risk individuals. In an environment where employer-sponsored health insurance is the norm, one of the largest impediments to insured mobility is the cost of switching jobs. To the

¹ This process is empirically documented in a convincing study by Altman, Cutler, and Zeckhauser (1998) that examines the changing enrollments for individuals insured through the Group Insurance Commission of Massachusetts.

extent that health insurance is bundled with employment, any frictions associated with employee mobility across occupations or employers necessarily translate into a de facto commitment to one's employer, which results in a more stable, and hence more insurable, risk pool. Thus, the existence of frictions that impede worker mobility between alternative employments may have heretofore unappreciated beneficial effects through their impact on the insurability of health risks.

The validity of this approach, of course, depends crucially on the requirement that insureds switch jobs in order to obtain alternative insurance coverage. This would appear to be a reasonable proposition in practice, since the vast majority of insured individuals receive their health coverage through employer-sponsored plans. The difficulties associated with the offering of private health insurance are well known, and they include the traditional problem of adverse selection with nongroup insurance (Pauly, 1986) as well as the high fixed costs generally encountered in the administration of individual policies (Diamond, 1992). As a result, the terms associated with independently purchased health insurance, to the extent that it is available, are less than attractive when compared to the employer-sponsored alternative (Gruber and Madrian, 1994).

The role of precommitment in the design of efficient multiperiod contracts has received increasing attention in the literature (Chiappori, 2000). The most germane for this study is a recent article by Hendel and Lizzeri (2003) who, in the context of life insurance, examine the use of "front-loaded" premiums as a mechanism through which insureds can commit credibly to remain in the insurance pool. There is a model of symmetric information in which the purchasers of insurance may face classification risk based on their evolving actuarial status, and in which those who turn out to have a lower risk of mortality may, in the absence of precommitment, exit the life insurance pool. While the prepayment of premiums may be an effective tool to "lock in" customers in life insurance, Cutler and Zeckhauser (2000) note that it is less likely to provide a solution to the commitment problem in health insurance settings due to the greater complexity of that market.² Thus, health insurers must use other commitment devices—such as employer sponsorship of health plans—to craft stable insurance pools.

The article proceeds as follows. In Section 2 we examine a model of employer-sponsored health insurance in which the publicly observed health status of employees changes over time. Under the assumption that the insured workers incur a financial cost when switching jobs to obtain alternative insurance coverage, we characterize the structure of efficient employer-sponsored health insurance contracts. Our model offers several predictions concerning the relationship between these switching costs, which we term "job attachment," and the amount of insurance coverage available through employment-based groups. We find that when employers offer the same contract to all their workers, the optimal contract exhibits a coverage limitation that is inversely related to the amount of job attachment present in the firm. In addition, if employers are able to offer multiple contracts that induce self-selection by insureds, the contracts exhibit more complete coverage of medical expenditures, albeit at premiums that partially reflect the health status of plan participants.

Section 3 presents empirical tests of the model's predictions using data on health insurance contracts obtained from the 1987 National Medical Expenditure Survey matched to proxies for job attachment from the *Dictionary of Occupational Titles* (U.S. Department of Labor, 1977). Consistent with the predictions of our model, we find that the contracts associated with firms that offer a single health insurance policy exhibit coverage limitations that are decreasing in job attachment, whereas firms offering multiple policies have higher levels of coverage that are less sensitive to job attachment. In Section 4 we examine the extent to which our findings might be explained by forces other than job attachment. A final section contains concluding remarks.

2. Theoretical framework

■ We consider an environment in which a continuum of individuals faces the financial risk associated with becoming ill at some future date. All agents are assumed to be identical initially

² Herring and Pauly (2003) argue that appropriately designed front-loaded policies can provide a feasible solution to the commitment problem in the nongroup health insurance market.

and have the same probability, \bar{p} , of sickness. Each agent also anticipates receiving publicly observable information about her own health status prior to the state in which illness may occur. Accordingly, an individual realizes that with probability λ ($1 - \lambda$) she will turn out to be a high (low) risk with probability p^H (p^L) of suffering illness, where $1 > p^H > p^L > 0$ and $\bar{p} \equiv \lambda p^H + (1 - \lambda)p^L$. Individual consumers have an income Y and, in the event of sickness, suffer financial loss S . For the purposes of this model, we assume that workers obtain a health benefits package at the time they become employed but prior to receiving the information on their health status. This package may consist of either a single insurance contract or a menu of potential insurance choices. Before becoming ill, however, workers obtain information on their health status, which may affect their incentives to remain in one of the plans offered by their employer or to select alternative options in cases where multiplan benefit packages are offered. This sequence of events is illustrated with a timeline in Figure 1.

Each individual is assumed to be risk averse and to possess the von Neumann-Morgenstern utility function $U(W_j)$, where W_j is the agent's wealth in the sick ($j = S$) and not-sick ($j = N$) states. An insurance contract $c \equiv \{Z, R\}$ consists of a premium, Z , paid by the insured individual prior to the state in which illness may occur, and a reimbursement, R , received by the insured in the event of illness. The expected utility of an agent who purchases the insurance contract c and has the probability of illness p is written as

$$V(c; p) \equiv pU(W_S) + (1 - p)U(W_N), \tag{1}$$

where $W_S \equiv Y - Z - S + R$, and $W_N \equiv Y - Z$. Since public information about each agent's type will be available prior to the period in which illness may occur, the information structure of the model permits individuals to be treated differently depending on their publicly known types. Letting $c^i = \{Z^i, R^i\}$ denote the contract associated with an individual whose observed health status indicates her to be a type $i \in \{H, L\}$, a health insurance benefits package $C \equiv \{c^H, c^L\}$ results in an expected utility to an agent (prior to obtaining health status information) of

$$V(C) \equiv \lambda v(c^H; p^H) + (1 - \lambda)v(c^L; p^L). \tag{2}$$

Since we assume that insurance firms are risk neutral, we may write the expected profit associated with the offering of package C as

$$\Pi(C) \equiv \lambda \pi(c^H; p^H) + (1 - \lambda)\pi(c^L; p^L), \tag{3}$$

where the expected profit earned on an individual of type i is $\pi(c^i; p^i) \equiv Z^i - p^i R^i$.

□ **Spot markets and efficient benefit packages.** Before proceeding, we consider two benchmark cases. The first is the equilibrium that would occur in the spot market were insurance not bundled with employment and workers were to purchase actuarially fair insurance after their

FIGURE 1

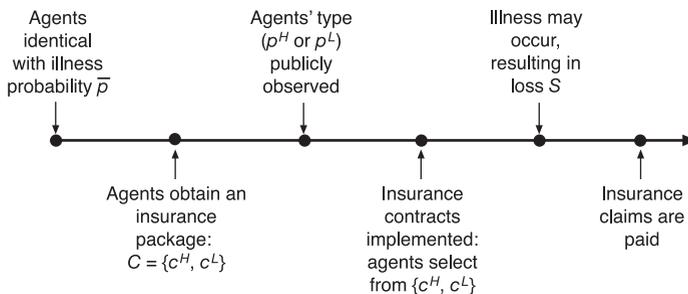
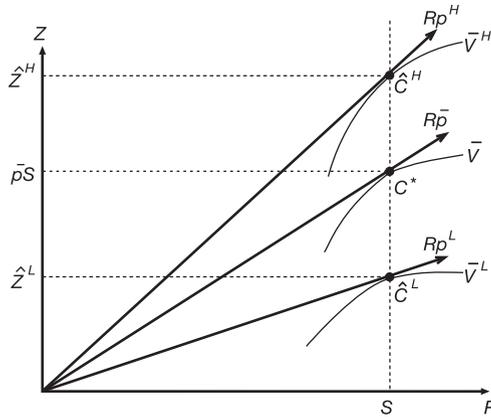


FIGURE 2



health status had been observed. It is straightforward to demonstrate that the equilibrium is characterized by a solution to the problem of selecting C to maximize the expected utility of insureds (2) subject to the zero-profit constraint $Z^i = p^i R^i$ for each $i \in \{H, L\}$. While the resulting contract provides full insurance against the financial loss generated by an illness, so $R^i = S$ for each i , the insured is completely exposed to the classification risk associated with alternative health outcomes through the experience-rated (future) premiums, Z^i . The spot market contracts are depicted in Figure 2, where the insured receives \hat{c}^H (\hat{c}^L) if she is viewed to be high (low) risk.

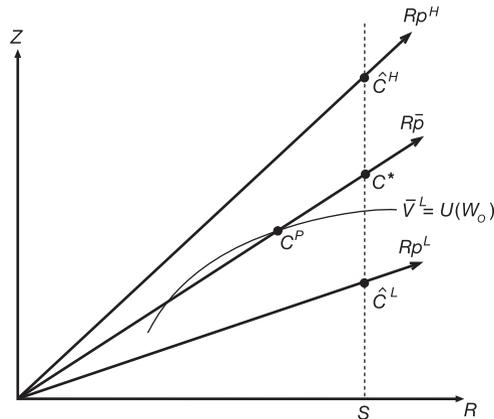
As an alternative to the spot market, insureds could opt for a contract negotiated prior to receiving knowledge of their health status. Such a full-commitment insurance package is characterized by a solution to the problem that maximizes insured expected utility (2) subject to $\Pi(C) \geq 0$, which is depicted as C^* in Figure 2 and results in full coverage of all illness-related expenses ($R^H = R^L \equiv S$) as well as a constant, *ex ante* actuarially fair, premium that does not depend on revealed health status ($Z^H = Z^L \equiv \bar{p}S$). Although the full-commitment package completely insulates insureds from both the financial risk of illness and the classification risk associated with observable changes in health status, the premium structure entails a cross subsidy in which low-risk (high-risk) individuals pay a premium that is above (below) their publicly known actuarially fair rate. Since the courts will not generally enforce long-term contracts against insurance purchasers, the low-risk (p^L -type) insureds may successfully renege on their promise to pay the pooling premium $\bar{p}S$ unless other mechanisms to enforce compliance can be implemented.³

□ **Partial commitment through employment bundling.** When health insurance is provided only as part of an employee benefits package, frictions associated with movement to alternative employers may serve to impede the ability of those individuals who find that they are low risk from leaving the insurance pool. To investigate the extent to which such mobility frictions can serve as a precommitment device, we assume that the low-risk workers can obtain full and actuarially fair health insurance after their health status is revealed but in so doing must incur a switching cost, K , of moving to an alternative employer. Accordingly, the extent to which an insurance package can insure against classification risks is limited by the feasibility constraint

$$v(c^L; p^L) \geq U(W_0), \tag{4}$$

³ See Epstein (1997) for a discussion of the legal problems associated with enforcing long-term insurance contracts. A concise summary of the problem can be found in Cochrane (1995, p. 468), who notes that “courts often reinterpret insurance contracts *ex post*, judge the merits of each clause separately rather than how the clauses fit together to form a reasonable contract, and will not enforce severance payments or bond forfeitures against consumers.”

FIGURE 3



where $W_0 = Y - K - p^L S$, which guarantees that a low-risk individual prefers her current insurance plan to the (actuarially fair, full-insurance) outside option.

The existence of the job attachment conferred by the switching cost K permits some degree of commitment on the part of insureds to remain in the insurance pool even if they turn out to be healthier than average. We will examine the effect of this ability to precommit in two insurance settings, the first of which requires that insurers offer the same contract to all insureds independent of their revealed health status, while the second permits insurers to design benefit packages that induce the members of the pool to voluntarily select insurance contracts that result in experience-rated premiums.⁴

Single-contract benefit packages. One method of mitigating classification risk is by offering an insurance package that does not take into account each insured's evolving health status. Let C^P denote a pooling package in which $c^H = c^L \equiv c^P$. The optimal pooling package consists of the contract $c^P \equiv \{Z^P, R^P\}$ that maximizes (2) subject to the zero-profit (3) and feasibility (4) constraints. The contract associated with the optimal pooling package may be described with reference to Figure 3.⁵

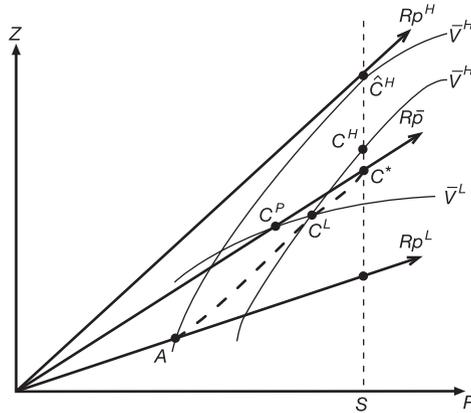
For sufficiently low levels of worker commitment, K , the indifference curve associated with the feasibility constraint, $\bar{V}^L = U(W_0)$, lies everywhere below the pooling zero-profit locus, $\bar{p}R$. In such a setting, there is no pooling contract acceptable to low-risk individuals, so that a pooling package cannot be implemented. As K becomes larger, resulting in higher degrees of employee commitment, eventually a critical value (K_1) is reached at which the indifference curve associated with the feasibility constraint intersects the pooling zero-profit locus, resulting in a feasible pooling insurance package. Further increases in K cause the contract c^P to migrate up the zero-profit pooling locus, which is the situation depicted in Figure 3. As K continues to increase, eventually a critical value K_2 is reached that permits the contract C^* to be achieved. The implication is that firms offering a single health insurance plan should have coverage limitations that are decreasing in the degree of worker commitment, K .

Multicontract benefit packages. In some cases, firms may offer multiple choices to workers, who choose their preferred coverage from among the various insurance options. Although employer-sponsored insurance plans rarely engage in mandatory experience rating of individual participants in an insurance pool, some benefit designs provide the mechanism by which a form of voluntary experience rating can be effected through the offering of insurance contracts that induce self-selection by the insureds. Formally, an insurance package $\{c^H, c^L\}$ can discriminate based upon

⁴ When workers are free to choose among plans, inducing self-selection is necessary even with observable types.

⁵ The formal proof of this result, and of the form of the optimal multicontract benefit package discussed below, is available from the authors upon request.

FIGURE 4



health status as long as

$$v(c^H; p^H) \geq v(c^L; p^H), \tag{5}$$

which requires that high-risk individuals prefer the insurance contract c^H to c^L .⁶ An optimal multicontract benefit package is a solution to the problem that selects $\{c^H, c^L\}$ to maximize (2) subject to the zero-profit (3), feasibility (4), and self-selection (5) constraints.

An optimal insurance package with multiple contracts is illustrated as $\{c^H, c^L\}$ in Figure 4, where the efficient pooling contract, c^P , has been included for comparison purposes. The locus AC^* depicts the set of contracts sold to low-risk individuals which, when coupled with a full-coverage ($R^H = S$) contract that is equally valued by high-risk individuals (so $v(c^H; p^H) = v(c^L; p^H)$), satisfies the aggregate zero-profit constraint (3) with equality. One element of this contractual class is $\{\hat{c}^H, A\}$, which entails no cross subsidization between the two risk types.⁷ Other members of this class of contracts will, however, contain a cross subsidy. As the low-risk contract migrates toward C^* along the AC^* locus and the high-risk contract (on the same \bar{v}^H indifference curve) moves from \hat{c}^H toward C^* , the magnitude of this cross subsidy increases. One such contract is that denoted as $\{c^H, c^L\}$ in Figure 4, where profit earned on low-risk individuals subsidizes losses incurred on the contract sold to higher risks. The limiting case occurs when the contracts for both risk types coincide at the pooling allocation C^* .

In the absence of the feasibility constraint imposed by the potential exit of low-risk individuals, the optimal contract would be the full-commitment package C^* . When that feasibility constraint binds, however, the closest the insurer can get to the full-commitment contract is the multicontract package depicted as $\{c^H, c^L\}$ in Figure 4, which exposes the low-risk insureds to a financial loss in the event of illness, and the high-risk insureds to some classification risk.

There are several aspects of the optimal multicontract package that are worthy of note. First, $\{c^H, c^L\}$ represents a Pareto improvement over the pooling contract c^P , since high risks are strictly better (and low risks no worse) off.⁸ Second, the multicontract package provides some insurance against classification risk, although that protection is not complete owing to the need to keep the low risks in the insurance pool. Third, the optimal pooling contract always provides less coverage than do either of the contracts offered under the multicontract package ($R^P < R^L < S = R^H$).

⁶ Formally, we must have $v(c^L; p^L) \geq v(c^H; p^L)$ as well. We ignore this constraint because it is never binding.

⁷ These contracts correspond to the well-known Rothschild-Stiglitz separating allocation in which each contract makes zero profit, and they are what we would expect to observe were there no mechanism for effecting cross subsidization within the pool. This would be the case if c^H and c^L were issued by different insurers and employee premium contributions were not adjusted by the employer.

⁸ The pooling contract C^P is, by construction, forced to accommodate the participation constraint (4) entirely through a distortion in the financial risk margin, while multicontract packages can (more efficiently) spread that distortion over two margins. Many firms, however, have insufficient enrollment to support multiple plans.

Finally, since c^P and c^L both converge to C^* as K increases (and coincide with C^* at K_2), the effect of increased worker commitment on coverage limitations is more pronounced in the optimal pooling package than in *either* of the contracts offered in the efficient multicontract package. These results will be useful when we implement our empirical tests because they eliminate the need to identify which type of contract (c^H or c^L) is held by policyholders in firms that offer multicontract insurance packages. We therefore have the following testable implications of the theory.

Proposition 1. Firms that offer a single insurance contract to their workforce should have lower levels of coverage that are more sensitive to worker commitment than firms that offer multicontract packages. Thus, $R^P < R^i$ and $dR^P/dK > dR^i/dK$, for $i \in \{H, L\}$.

Proof. See the Appendix.

The importance of job attachment ($K > 0$) is that it permits a limited degree of cross subsidization from the low-risk to the high-risk customers to be implemented, without inducing the former to leave the insurance pool. In settings with no job attachment ($K = 0$), there is no way of effecting such a subsidy, which is exactly the situation encountered in the individual health insurance market and the reason that insured individuals end up exposed to classification risk. We now turn to an examination of the data.

3. Empirical analysis

■ From an empirical perspective, the main testable implication of the model is the finding that the amount of insurance contained in employment-based health insurance contracts should be positively related to agents' ability to precommit to remain in a particular insurance pool. Because workers often must change jobs to obtain more favorable health insurance coverage, measures of the transactions costs associated with moving among employers can be used to proxy for insured commitment, thereby permitting a direct test of the effect of precommitment on the design of contractual agreements.

□ **Data.** Our empirical analysis is conducted using detailed employment and health insurance data from two components of the 1987 National Medical Expenditure Survey (NMES): The Household Survey and the Health Insurance Plans Survey (HIPS), matched to proxies for job attachment from the 1977 *Dictionary of Occupational Titles* (DOT). The DOT provides information on the physical demands, environmental conditions, and educational and vocational preparation associated with each of 12,000 occupations.

The NMES Household Survey is a stratified random sample of the civilian noninstitutionalized population of the United States containing primarily individual-level data on the medical expenditures, demographic characteristics, employment status, and health insurance coverage of some 35,000 individuals in 14,000 households. Household Survey respondents who reported coverage from private insurance were reinterviewed in the HIPS to obtain more detailed information on their type and level of coverage, premiums, deductibles, maximum benefits, and covered illnesses. The HIPS was designed to provide a random sample of all private health insurance policyholders in the civilian population of the United States at the end of 1987. The data available in HIPS further supplements the Household Survey data by providing firm-level information on the characteristics of respondents' employers as well as their employer-sponsored health insurance plans.

The 1987 NMES is well suited for our purposes because it contains data from a period when managed care organizations were relatively minor players in the health insurance market. This is important because we need easily quantifiable measures of the overall amount of insurance contained in each contract, something that is considerably harder to measure when insurers and health care providers are vertically integrated.

□ **Measuring commitment/job attachment.** To test the predictions of our model, we require an observable proxy for the degree of job attachment present in firms. A number of impediments

to job mobility have been identified in the literature, with particular attention being focused on human capital specialization, which refers to the subset of a worker's knowledge or skills that are differentially valued by a particular firm or within a particular occupation or industry. In general, the more specialized one's skills become, the costlier it will be to change employers. In the case of firm-specific training, these costs reflect reductions in productivity and earnings at rival firms.⁹ Alternatively, when skills are occupation- or industry-specific, the transactions costs of changing employers may reflect either reduced productivity, if the switch entails leaving one's occupation or industry, or simply the increased costs of finding a job as employment becomes more specialized.

To proxy for job attachment and, by implication, the degree of commitment to a particular employment-based insurance pool, we use a measure from the DOT of the training specificity required in various occupations. The variable, known as "Specific Vocational Preparation" (SVP) is defined as "the amount of time required to learn the techniques, acquire information, and develop the facility needed for average performance in a specific job-worker situation" (U.S. Department of Labor (1981), p. 473) and is based on the nine categories of vocational preparation shown in Table 1.

A proxy for worker-level job attachment was obtained by imputing an SVP value to each worker in the Household Survey. Because the Household Survey occupation codes are based on the occupation codes used in the 1980 Census, it was possible to impute an SVP value to each worker using the Commerce Department's Standard Occupational Classification (SOC) codes, in conjunction with the Census Bureau's Classified Index of Industries and Occupations, to form a crosswalk between the DOT and the Household Survey. The DOT provides a finer occupational classification than the Census, so there were often multiple SVP values associated with each Census occupation code. To impute a unique SVP value to each Census occupation code, we took a simple average of the SVP values associated with each Census occupation.

The SVP variable, although perhaps not familiar to many readers, offers several advantages over other possible measures of job attachment. First, unlike job tenure or turnover, it is not a function of the quality of health insurance offered by the respondent's employer. Second, we are not using the actual amount of training received by a particular worker, which is also potentially endogenous, but rather the amount of training typically required in the worker's *occupation*. These observations suggest that one can plausibly treat SVP as exogenous when considering the amount of insurance offered by a particular firm.¹⁰

Based on the definition alone, it is not possible to determine definitively whether SVP captures the type of specialized training that tends to tie workers to their employers. On an intuitive level,

TABLE 1 Specific Vocational Preparation (SVP)

Description	Numerical Value
Short demonstration	1
Anything beyond a short demonstration up to and including 30 days	2
Over 30 days up to and including 3 months	3
Over 3 months up to and including 6 months	4
Over 6 months up to and including 1 year	5
Over 1 year up to and including 2 years	6
Over 2 years up to and including 4 years	7
Over 4 years up to and including 10 years	8
Over 10 years	9

Source: The 1977 *Dictionary of Occupational Titles*. See the Appendix for a complete definition of SVP.

⁹ Available evidence suggests that the returns to specialized training are substantial. Topel (1991, p. 147), for example, finds that "10 years of job seniority raises the wage of the typical male worker in the United States by over 25% relative to what he could obtain elsewhere."

¹⁰ Of course, it may be that SVP is correlated with other worker characteristics, e.g., risk aversion, that also influence insurance coverage. We will explore the role of such omitted variables later in the article.

TABLE 2 Worker-Level Job Tenure and Wage-Tenure Regressions
 Dependent Variable: 1–2 = Number of Years Spent with Worker's
 Current Employer; 3–4 = Natural Log of Worker's Hourly Wage

	Job Tenure		Wages	
	(1)	(2)	(3)	(4)
Constant	–6.948 (.593) [.000]	–6.150 (.688) [.000]	.897 (.072) [.000]	.799 (.069) [.000]
Specific Vocational Preparation (SVP)	.360 (.062) [.000]	.384 (.057) [.000]	.081 (.008) [.000]	.076 (.006) [.000]
Years of Education	–.139 (.034) [.000]	–.124 (.033) [.000]	.035 (.005) [.000]	.037 (.004) [.000]
Job Tenure	—	—	.009 (.003) [.003]	.007 (.003) [.010]
Job Tenure × SVP	—	—	–.002 (.0004) [.001]	–.002 (.0004) [.000]
Years of Education × SVP	—	—	.001 (.0003) [.005]	.001 (.0003) [.002]
Industry Dummies	No	Yes	No	Yes
R ²	.377	.383	.377	.403
N	11,426	11,426	11,426	11,426

Notes: Specifications include controls for worker demographics (age, sex, race, marital status, family size) and firm characteristics (region, urban location, union membership, employer size) and, in the job tenure regressions, wages. Specifications in columns 2 and 4 also include 12 industry dummies. Results for these variables are available upon request. Standard errors are in parentheses; *p*-values shown in brackets. Standard errors adjusted for both heteroskedasticity and within-occupation error correlation.

we would expect that if SVP does measure specific training, it should be positively related to the length of time that workers remain with a given employer.¹¹ General education, on the other hand, typically increases worker mobility, which should be reflected in shorter job tenures (Ehrenberg and Smith, 2003). In columns 1 and 2 of Table 2, we regress job tenure, defined as the number of years that a worker has been with his current employer, on the worker's SVP value, his educational attainment, and a number of other variables (listed at the bottom of the table) using a sample of all currently employed, non-self-employed workers from the NMES Household Survey. In column 1, we include controls for worker demographics and firm characteristics; column 2 adds 12 industry dummies. In both specifications, SVP exerts a positive and statistically significant influence on job tenure. This is in contrast to general education, which, as shown in Table 2, reduces job tenure.

A sharper test of whether SVP measures specific, as opposed to general, training is provided by human capital theory. As summarized by Ehrenberg and Smith (2003), standard training models predict that general training will steepen wage-tenure profiles because workers must pay, in the form of lower wages, for any training that is transferable across employers. Since this training is typically received early in one's career, the lower wages that workers initially receive are followed

¹¹ Since job tenure and turnover depend on both the costs and *benefits* of changing employers, they are not good measures of job attachment for our purposes. An industry could experience low turnover because of high mobility costs or, alternatively, because both the costs and benefits of mobility are low; only in the former case would switching based on the desirability of one's health insurance be inhibited.

by higher earnings later when the training begins to increase productivity. Because these skills are transferable, firms must pay workers their full marginal product in the post-training period. Conversely, firm-specific training, or any other type of training that is not fully transferable across employers, flattens wage-tenure profiles because the firm must make an upfront investment in the training (by paying workers in excess of their marginal product during the training period) but recoups its investment later once the workers are locked in. Thus, to the extent that SVP proxies for the presence of skills that are firm-specific or otherwise nontransferable, we would expect it to attenuate the effect of job tenure on wages in our sample. By contrast, general training, which we proxy for with years of education, should lead to a steeper pay-tenure gradient, all else equal.¹²

To test these predictions, we regressed the natural logarithm of wages on job tenure, SVP, years of education, job tenure interacted with SVP, and job tenure interacted with years of education, using the same estimation sample and control variables as were used in our job tenure regressions. Results are presented in columns 3 and 4 of Table 2.¹³ As expected, we find that SVP, general education, and the number of years with one's current employer all raise wages. Focusing on column 4, we find that all else equal, each additional year with the same employer increases wages by .70%. The effects of SVP and general education on the wage-tenure profile are captured by the interaction terms, with a positive (negative) coefficient indicating that the gradient is made steeper (flatter) by an increase in the interacted variable. Using the estimates from column 4, we find that a one-standard-deviation increase in educational attainment raises the return to an additional year of job tenure by about .28%, while a one-standard-deviation increase in SVP *lowers* the return by about .37%. Relative to a baseline return of .70% per year, these changes respectively represent a 40% increase, and a 53% decrease, in the annual returns to job tenure. These findings provide evidence that SVP does capture training that is largely firm-specific in nature, and that the acquisition of this training ties workers to their employers in a manner consistent with human capital theory. The attenuation of returns to job tenure associated with higher levels of SVP also provides a tangible empirical analog to the generic employer-switching costs specified in our theoretical model.

Based on the results presented above, SVP would appear to be a good proxy for job attachment at the level of an individual worker. But because decisions about health insurance are made at the firm level, the best conceptual measure for our purposes is one that measures job attachment at the firm level, that is, one measuring the average amount of job attachment in a particular firm. Given the need to obtain a firm-level measure, there were two ways in which we might have proceeded. One possibility was to simply use the SVP value assigned to each worker as a proxy for the average degree of skill specialization in that worker's firm. This approach would have been desirable if we believed that within-firm human capital heterogeneity was not very large. An alternative approach, and the one adopted in this article, was to construct a measure of the average degree of skill specialization in various industries and use this as a proxy for the typical amount of specialization arising in particular firms within those industries. This latter approach will be preferable if, as appears likely, there is less skill heterogeneity across firms (in narrowly defined industries) than within firms.

To construct a measure of job attachment at the firm level, we first computed the average SVP value in each worker's industry, labelled *Industry SVP 1980*, by averaging the SVP values by industry for all employed persons in the 1980 Census Public Use Microdata Sample (PUMS) (5% sample). To control for income differences at the industry level, which may confound *Industry SVP 1980* if more specialized training is associated with higher point-in-time earnings, we also computed average income in each industry using the PUMS data. Because the industry codes in the NMES Household Survey are at the three-digit level, representing some 230 distinct industries,

¹² Although we focus on general education rather than general training, the two should have the same effect on the wage-tenure gradient: although firms typically pay for neither general training nor general education, both lead to the acquisition of skills that are readily transferable across employers.

¹³ All the estimated coefficients reported in Table 2 are statistically significant at the .01 level or better.

we believe that *Industry SVP 1980* is likely to provide a good measure of the average amount of job attachment present in individual firms.

One problem with this measure is that it is based on the distribution of occupations existing in 1980. If between 1980 and 1987 (the year of the NMES survey) there were important changes in the distribution of occupations within industries, then *Industry SVP 1980* will be subject to measurement error. To avoid this problem, we calculated a second job-attachment variable, *Industry SVP 1987*, using the 18,000 persons in the NMES Household Survey who reported a Census occupation code.¹⁴ This variable captures the within-industry distribution of occupations that existed at the time of the survey. However, unlike *Industry SVP 1980*, which is based on millions of observations, and therefore thousands of observations per industry, *Industry SVP 1987* is based on only 18,000 observations in total, so the number of observations per industry is quite small in some cases. As a result, *Industry SVP 1987* may also be subject to measurement error. Since these two potential sources of measurement error should be unrelated, we will use both measures of job attachment in our analysis and examine the robustness of our findings across the two variables. In addition, we will reestimate each model using *Industry SVP 1980* as an instrument for *Industry SVP 1987*. It is well known that the instrumental variables estimator remains consistent in the presence of a mismeasured explanatory variable, even in cases where the instrument itself is measured with error, provided there is not a common component to the measurement error across the two variables.

□ **Measures of insurance.** We use two standard contract provisions as measures of the amount of insurance available through each policy. The first is the lifetime limit on benefits, which is positively related to the amount of insurance coverage specified in a plan, and probably the best overall measure of the extent to which long-term, costly medical conditions are insured. The second is the annual stop-loss, defined as the threshold level of medical expenditures above which the policyholder is no longer required to make co-payments. In contrast to the lifetime benefit caps, the annual stop-losses are less closely tied to catastrophic medical expenditures but are more likely to be reached, and therefore they could be viewed as a more immediately relevant measure of the amount of insurance contained in a plan. The stop-losses are also of interest because they apply on an annual, as opposed to a lifetime, basis. This allows for an additional check on the robustness of our findings.

We make use of these two contract provisions, in lieu of others such as the annual deductible or the plan actuarial value, because we believe they map most directly to our theoretical model, which is fundamentally a model of *insurance* provision.¹⁵ We emphasize the word “insurance” because many health insurance plans combine prepaid medical care (i.e., care for common, low-cost afflictions) with true “insurance,” by which we mean protection against less common, but significantly more costly, medical conditions. In particular, because our article deals with the role of precommitment in maintaining viable risk pools, we would expect our measure of job attachment to matter most for insuring losses that require large, foreseeable cross subsidies from other members of the insurance pool. This is most likely to occur in cases where a coworker develops a serious medical condition that is both chronic in nature and costly to treat, such as diabetes, cancer, or heart disease. For such conditions, we would expect that the overall amount of financial protection conferred by the plan would be determined disproportionately by contract provisions that apply to the right tail of the loss distribution.¹⁶

□ **Estimation samples.** To create our samples, we applied a common set of restrictions to the 6,549 HIPS respondents reporting employment-related health insurance coverage. Specifically,

¹⁴ To preserve comparability, we also calculated a measure of average industry income using this group of workers.

¹⁵ Actuarial values measure the fraction of medical expenditures reimbursed by the plan for the average person in the population. Thus, for a representative individual, a generous plan is one that provides a lot of “first-dollar” coverage, even if it omits coverage for (low-probability) catastrophic illnesses.

¹⁶ From an insurance perspective, it is exactly these low-probability, high-expenditure conditions that risk-averse individuals will most want to insure. As a result, the actuarial value is not a particularly good measure of insurance coverage, even though it may accurately reflect plan generosity in a different sense.

we dropped any person who was classified as unemployed, or for whom a link to a “current main job” was unavailable. We also dropped any person who did not hold an employer- (or union-) sponsored plan, and anyone whose employer did not offer group coverage or who was employed by a subchapter S corporation.¹⁷ We also deleted a small number of persons who held policies from more than one employer. These restrictions, coupled with observations lost from missing or incomplete data, resulted in final samples of 2,130 policyholders for the lifetime benefit sample and 1,984 policyholders for the annual stop-loss sample.

To control for factors other than job attachment that might affect the level of coverage provided under employment-based health insurance contracts, we include controls for various employer and health plan characteristics, as well as industry income.¹⁸ We include a measure of average industry income in all of our regressions to ensure that *Industry SVP* is not simply picking up income differences across industries that are potentially correlated with insurance purchases. In the stop-loss samples, we also include controls for whether the stop-loss applies to all covered expenses or only the policyholder’s out-of-pocket expenses, and whether the annual deductible is counted toward the stop-loss.

The mean lifetime benefit in 1987 was approximately \$850,000, and there is substantial variation in these benefits; in our sample the range runs from a low of \$10,000 to a high of \$25,000,000.¹⁹ Although amounts below \$50,000 may seem unusual, we have only included lifetime maximums that apply to coverage categories likely to yield high expenses, such as hospital room and board charges, inpatient surgical benefits, and inpatient physician fees. We imposed this restriction to ensure that these lower limits were the limits for the policy as a whole and not for specialized types of care that sometimes have their own maximum benefits, such as mental health or substance abuse treatment. Nonetheless, our results do not change if we eliminate observations with maximum benefits below \$50,000 or \$100,000 from the sample.

The mean annual stop-loss is approximately \$2,400, and there is considerable variation in this variable as well. We have less reason to worry that the reported stop-losses do not apply to the policy as a whole because there is generally only one stop-loss reported for each policy. Again, our results do not change if we drop potential outliers, such as observations with stop-loss values below \$500 or above \$25,000.

There are a few other characteristics of the samples that are worth mentioning. For both the lifetime benefit and stop-loss samples, we see that roughly 60% of employers offer only one health insurance policy. About 45% of employers self-insure to some degree, and only about 1% of health plans are classified as HMOs. Approximately 70% of the establishments are for-profit, roughly a third have more than 500 employees, and about a third are either fully or partially unionized.

□ **Baseline empirical results.** The effect of job attachment on the amount of insurance coverage available from employment-based health insurance contracts can be estimated from a simple OLS regression of the type shown below:

$$\text{Insurance}_i = \beta_0 + \beta_1 \text{Industry SVP}_i + \gamma' \mathbf{X}_i + \varepsilon_i, \quad (6)$$

where \mathbf{X}_i is a vector of employer and insurance plan characteristics for worker i , Industry SVP_i is a proxy for the average degree of job attachment in the firm that employs worker i (either *Industry SVP 1980* or *Industry SVP 1987*), and ε_i is a random error term. The dependent variable in our analysis, Insurance_i , is a measure of the total amount of insurance coverage provided by the

¹⁷ The coefficient estimates on the variables of interest are little changed when we include S corporations and account for their influence with a dummy variable. However, the use of this variable reduces the precision of our estimates substantially because of missing observations, which shrink our estimation samples by roughly one-third.

¹⁸ These controls are described in Tables 3 and 4.

¹⁹ Summary statistics for the lifetime benefit and annual stop-loss samples are available upon request from the authors.

policyholder's insurance contract (either the maximum lifetime benefit or the annual stop-loss).²⁰ In terms of (6), the estimated value of β_1 provides a direct measure of the effect of commitment on the amount of insurance provided by a given employment-based health insurance policy. To address the possibility that either *Industry SVP 1980* or *Industry SVP 1987* could be measured with error, we estimate equation (6) using each variable separately and also using an instrumental variables (IV) procedure in which *Industry SVP 1980* is used as an instrument for *Industry SVP 1987*.

To examine the main predictions from our theoretical model, which concern the differential effects of job attachment in firms offering a single health insurance plan ("single-plan firms") relative to firms offering multiple plans ("multiple-plan firms"), we will run (6) separately for both firm types. To test for significant differences between the two, we will run a difference-in-differences model on the pooled sample of all firms, as shown below:

$$\begin{aligned} \text{Insurance}_i = & \beta_0 + \beta_1 \text{Industry SVP}_i + \beta_2 \text{Single Plan}_i \\ & + \beta_3(\text{Single Plan}_i \times \text{Industry SVP}_i) + \gamma' \mathbf{X}_i + \varepsilon_i, \end{aligned} \quad (7)$$

where *Single Plan_i* is a dummy variable indicating that worker *i*'s employer offers a single health insurance policy.^{21,22} Observe that β_2 captures the mean difference in insurance coverage between single- and multiple-plan firms, while β_3 measures the differential effect of job attachment in firms offering a single policy relative to firms offering multiple policies. When the lifetime maximum benefit is the dependent variable, our theoretical model implies that $\beta_2 < 0$ because, all else equal, single-plan firms should offer less insurance than multiple-plan firms. More important, we would expect $\beta_3 > 0$, since the effect of job attachment on the quantity of insurance should be larger in firms offering a single policy than in firms offering multiple policies.²³ Notice that these predictions are reversed when the annual stop-loss is the dependent variable because a lower stop-loss translates into more insurance coverage.

Lifetime benefits. Results for lifetime benefits are presented in Table 3. Columns 1–6 display estimates of equation (6); columns 1–3 are based on a sample of firms offering a single health insurance plan, while columns 4–6 pertain to firms offering multiple plans. In each case, separate estimates are presented for *Industry SVP 1980* and *Industry SVP 1987* to provide a check on the robustness of our results with respect to the two potential sources of measurement error discussed earlier. In addition, in columns 3 and 6, an IV model is estimated in which *Industry SVP 1980* is used to instrument for *Industry SVP 1987*.

Looking at single- and multiple-plan firms separately, we see that the data are consistent with the predictions of our theoretical model. Focusing on the results for *Industry SVP 1980*, we see that a one-unit increase in job attachment increases the lifetime maximum benefit by approximately \$114,000 in single-plan firms, but by only \$10,000 in multiple-plan firms (in 1987 dollars). For single-plan firms, this effect is statistically significant at the 10% level, while in multiple-plan firms the hypothesis that the effect of job attachment is zero cannot be rejected with any reasonable level of confidence. As can be seen from column 2, results are similar, but

²⁰ Because *Industry SVP* is measured at the industry level while the dependent variable and other explanatory variables are measured at the firm level, it is possible that our regression residuals could be correlated within industries. To allow for this possibility, we implement a variant of the standard error correction proposed by Moulton (1986) using the CLUSTER command in STATA. This command also adjusts the standard errors for heteroskedasticity using White's procedure (White, 1980).

²¹ We assume that the number of health plans offered by the firm is exogenous. This is consistent with recent pension and health insurance studies, where the assumed exogeneity of the benefits choice set is the norm (Poterba, Venti, and Wise, 1995; Gale, 1998; Cardon and Hendel, 2001). We return to this issue later in the article.

²² IV estimation of (7), analogous to that used for (6), is not possible because of the need to employ multiple instruments once the effect of *Industry SVP* is allowed to vary across single- and multiple-plan firms. As we will see, however, the estimates obtained from (7) will generally be consistent with those from (6).

²³ Recall from Proposition 1 that these predictions hold regardless of which contract (c^H or c^L) is held by respondents in multiple-plan firms.

TABLE 3 Effect of Job Attachment on Lifetime Maximum Benefits
 Dependent Variable: Policyholder's Maximum Lifetime Dollar Benefit

	Single-Plan Firms			Multiple-Plan Firms			Difference-in-Differences	
	Industry SVP 1980 (1)	Industry SVP 1987 (2)	IV Model (3)	Industry SVP 1980 (4)	Industry SVP 1987 (5)	IV Model (6)	Industry SVP 1980 (7)	Industry SVP 1987 (8)
Intercept	356,236 (284,562) [.212]	402,292 (204,754) [.051]	340,836 (292,563) [.245]	1,243,714 (350,632) [.001]	1,112,866 (317,903) [.001]	1,177,936 (365,327) [.002]	1,102,305 (248,954) [.000]	1,138,929 (262,181) [.000]
Industry SVP	114,312 (65,908) [.084]	84,749 (40,154) [.036]	101,094 (69,506) [.147]	9660 (33,041) [.770]	22,737 (32,326) [.483]	5926 (45,128) [.896]	-14,793 (35,250) [.675]	-35,708 (41,291) [.388]
Single-Plan Firm	—	—	—	—	—	—	-657,074 (403,782) [.105]	-753,196 (382,166) [.050]
Single-Plan Firm × Industry SVP	—	—	—	—	—	—	133,032 (82,033) [.106]	147,833 (74,871) [.050]
R ²	.035	.038	.037	.048	.043	.045	.032	.033
N	1,254	1,277	1,254	839	853	839	2,093	2,130
Partial F-statistic for instrument	—	—	336.95	—	—	100.22	—	—

Notes: Standard errors are in parentheses; *p*-values shown in brackets. Standard errors adjusted for both heteroskedasticity and within-industry error correlation. IV models use *Industry SVP 1980* as an instrument for *Industry SVP 1987*. All regressions include controls for type of coverage (single, two-party, family, or other), self-insurance status of employer (fully self-insured, partially self-insured, or commercial insurance), type of insurance (HMO or indemnity), type of employer (for profit, nonprofit, government, or other), region (northeast, midwest, south, or west), urban location (SMSA or non-SMSA), establishment size (indicator variables for: <10 employees, 10–25 employees, 26–100 employees, 101–500 employees, >500 employees), employer unionization (fully unionized, partially unionized, or nonunion), and industry income. Results for these variables are available upon request.

more precisely estimated, when *Industry SVP 1987* is used to proxy for job attachment. A similar, albeit less precisely estimated, coefficient is obtained using the IV model shown in column 3. The consistency of the point estimates across these alternative specifications, in both the single- and multiple-plan samples, suggests that measurement error is not a large problem in this case. To gain some perspective on the magnitude of these effects, observe that in single-plan firms, a one-standard-deviation increase in *Industry SVP 1980* raises the maximum lifetime benefit by roughly \$83,000. Relative to the mean lifetime benefit in our sample, this represents about a 10% increase in available health insurance benefits over one's lifetime.

To further examine the predictions of our theoretical model, we next turn to the difference-in-differences models shown in columns 7 and 8. Again focusing on *Industry SVP 1980*, we can check for consistency between the results from this specification, which pools single- and multiple-plan firms, and the more general (but less efficient) specifications that examine each type of firm separately. First, note that given the way the difference-in-differences model is specified, the intercept in column 7 corresponds to the mean lifetime benefit in multiple-plan firms and is very similar to the corresponding estimate from column 4. The mean lifetime benefit for single-plan firms can be calculated from the difference-in-differences model by adding the intercept term to the coefficient on the *Single-Plan Firm* indicator; doing so yields a mean benefit of \$445,231, which is fairly close to the mean benefit of \$356,236 shown in column 1. Similarly, the effect of job attachment in multiple-plan firms is given by the coefficient on *Industry SVP*, which is small and not statistically different from zero at any reasonable level of significance. The effect of job attachment in single-plan firms is given by the sum of the coefficient on *Industry SVP*

and the coefficient on the interaction variable, *Single-Plan Firm* \times *Industry SVP*. Adding these coefficients yields a value of \$118,239, which is virtually identical to the corresponding estimate of \$114,312 from column 1. Thus, the difference-in-differences model appears to yield similar estimates to those obtained when single- and multiple-plan firms are considered separately.

The difference-in-differences specification provides a natural framework for testing the predictions of our theoretical model. As discussed previously, we would expect the coefficient on *Single-Plan Firm* to be negative because our model predicts that, *ceteris paribus*, single-plan firms should offer less insurance than multiple-plan firms. In addition, the coefficient on the interaction variable, *Single-Plan Firm* \times *Industry SVP*, should be positive because the effect of job attachment should be more pronounced in firms offering a single insurance policy.

Focusing again on *Industry SVP 1980*, we find that both of these predictions are upheld. All else equal, lifetime benefits are approximately \$657,000 lower in single-plan firms. Furthermore, a one-unit increase in *Industry SVP 1980* raises lifetime benefits by approximately \$133,000 more in single-plan firms than in multiple-plan firms, where the effect of job attachment is found to be negligible. These differences are statistically significant at approximately the 10% level and, as shown in column 8, become significant at the 5% level when *Industry SVP 1987* is used as our measure of job attachment

Annual stop-losses. Results for the annual stop-losses are shown in Table 4. We proceed in exactly the same fashion as in the previous section, with one difference: Given that a lower stop-loss is associated with more insurance coverage, the predicted signs of the coefficients from equations (6) and (7) are reversed. Specifically, we would expect that the coefficient on *Industry SVP* should be negative and larger in absolute value for single-plan firms than for multiple-plan firms. In the difference-in-differences models, we would expect the *Single-Plan Firm* dummy to carry a positive coefficient, consistent with single-plan firms offering less insurance, *ceteris paribus*,

TABLE 4 Effect of Job Attachment on Annual Stop-Losses
Dependent Variable: Policyholder's Annual Stop-Loss^a

	Single-Plan Firms			Multiple-Plan Firms			Difference-in-Differences	
	Industry SVP 1980 (1)	Industry SVP 1987 (2)	IV Model (3)	Industry SVP 1980 (4)	Industry SVP 1987 (5)	IV Model (6)	Industry SVP 1980 (7)	Industry SVP 1987 (8)
Intercept	1733.08 (676.11) [.011]	753.80 (688.32) [.275]	2166.06 (833.65) [.010]	447.90 (944.52) [.636]	-74.26 (788.37) [.925]	243.25 (1185.12) [.838]	89.07 (767.13) [.908]	-474.75 (649.50) [.466]
Industry SVP	-336.95 (164.81) [.042]	-196.63 (159.47) [.219]	-551.80 (250.49) [.029]	152.97 (154.31) [.323]	336.81 (138.71) [.016]	255.76 (265.35) [.337]	85.04 (145.24) [.559]	176.88 (133.49) [.187]
Single-Plan Firm	—	—	—	—	—	—	1802.66 (930.64) [.054]	1608.06 (831.39) [.054]
Single-Plan Firm \times Industry SVP	—	—	—	—	—	—	-372.43 (168.40) [.028]	-307.18 (151.95) [.045]
R ²	.173	.145	.167	.224	.226	.227	.178	.155
N	1,160	1,181	1,160	788	803	788	1,948	1,984
Partial F-statistic for instrument	—	—	351.87	—	—	124.06	—	—

Notes: See notes to Table 3. The nature of the stop-loss provisions necessitated the inclusion of two additional controls: an indicator for whether the stop-loss applies to covered or out-of-pocket expenses and an indicator for whether the annual deductible is counted toward the stop-loss. Results for these variables are available upon request.

^a Defined as the expenditure threshold above which the policyholder is no longer required to make co-payments.

while the *Single-Plan Firm* \times *Industry SVP* interaction variable should be negative, indicating that job attachment lowers stop-losses by more in single-plan firms.

Both of these implications are supported by the data. Focusing on *Industry SVP 1980*, a comparison of columns 1 and 4 reveals a negative effect of job attachment on stop-losses in firms offering a single health insurance policy, but not in firms that offer multiple policies.²⁴ The coefficient on *Industry SVP* in column 1 implies that a one-standard-deviation increase in *Industry SVP* lowers the annual stop-loss by approximately \$250. Relative to a mean stop-loss of \$2,400, this represents approximately a 10% reduction in the threshold set for limiting a policyholder's annual loss exposure. The difference-in-differences estimates reported in column 7 are consistent with single-plan firms having significantly higher stop-losses that are more negatively related to job attachment, as implied by our model. These differences are statistically significant at approximately the 5% level.

4. Additional evidence

■ In the previous section we documented that workers with more specific vocational preparation (SVP) have both longer job tenures and flatter wage-tenure profiles, all other things equal. These findings support the view that workers with more specialized training, as measured by SVP, face higher costs of changing employers than workers with lower levels of SVP. The fact that our proxy variable, SVP, conforms to the predictions of human capital theory is encouraging but does not rule out the possibility that it may simultaneously proxy for other worker characteristics that independently influence the amount of health insurance provided by employers. In particular, one could conjecture that workers who have more specialized skills may also have different underlying preferences, or might simply be "better" workers along some unobservable dimensions. If the former were true, then the possibility arises that our earlier results could be biased by a correlation between SVP and unobservable worker preferences for health insurance. Under the latter scenario, our results would be affected if some of the additional compensation that high-SVP workers receive were to come in the form of a higher lifetime limit on their health insurance policy or a lower annual stop-loss. Another potential concern is that firms offering a single health insurance plan might attract different types of workers than firms offering multiple plans, which could cause the coefficient on our *Single-Plan Firm* indicator to be estimated with bias.

In this section, we attempt to address these issues in two ways. We begin by examining the sensitivity of our baseline results to the addition of an unusually detailed set of explanatory variables that proxy for many of the typically unobservable person-specific factors that could influence the demand for health insurance. Of course, it is unlikely that any set of control variables could completely capture all the factors that determine the demand for insurance, or the relative desirability of workers. We therefore augment the use of these additional controls with two "falsification" tests. The first examines the insurance purchases of the self-employed. If *Industry SVP* were acting as a proxy for insurance demand, then one would expect to observe a positive relationship between this variable and the propensity of self-employed individuals to purchase a health insurance policy or to obtain coverage through other means. The second examines the hypothesis that high-SVP workers are superior to other workers in ways not accounted for by our control variables and that, as a result, firms compensate them more, not just in wages but in more generous health insurance coverage as well. If this were the case, one would expect to observe a positive relationship between *Industry SVP* and the propensity of firms to offer similar benefits. We focus here on life insurance because it presumably appeals to the same types of workers as health insurance. However, in contrast to health insurance, a viable nongroup market exists because life insurers have been able to solve the commitment problem through the use of front-loaded premiums (Hendel and Lizzeri, 2003). This allows nongroup policies to be sold on a guaranteed renewable basis at competitive rates, which makes individual policies a close

²⁴ These conclusions continue to hold when *Industry SVP 1987* is used in lieu of *Industry SVP 1980*, and when *Industry SVP 1987* is instrumented with *Industry SVP 1980*.

TABLE 5 Effect of Job Attachment on Lifetime Maximum Benefits
Difference-in-Differences Models with Additional Controls
Dependent Variable: Policyholder's Maximum Lifetime Dollar Benefit

	Industry SVP 1980				Industry SVP 1987			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Intercept	1,102,305 (248,954) [.000]	1,088,220 (278,275) [.000]	1,052,797 (476,759) [.028]	735,796 (514,332) [.154]	1,138,929 (262,181) [.000]	1,104,600 (294,723) [.000]	1,125,982 (473,718) [.018]	802,315 (520,037) [.125]
Industry SVP	-14,793 (35,250) [.675]	-43,035 (39,398) [.276]	-31,672 (47,257) [.504]	-28,626 (47,779) [.550]	-35,708 (41,291) [.388]	-65,518 (47,108) [.166]	-58,967 (59,781) [.325]	-64,461 (64,898) [.322]
Single-Plan Firm	-657,074 (403,782) [.105]	-701,376 (448,537) [.119]	-767,087 (525,616) [.146]	-809,671 (570,852) [.158]	-753,196 (382,166) [.050]	-800,648 (422,346) [.059]	-881,519 (506,170) [.083]	-949,677 (562,065) [.093]
Single-Plan Firm × Industry SVP	133,032 (82,033) [.106]	144,445 (92,007) [.118]	160,958 (109,957) [.145]	168,767 (117,259) [.152]	147,833 (74,871) [.050]	159,342 (83,463) [.058]	177,886 (101,569) [.081]	190,239 (111,238) [.089]
Demographics	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Insurance Demand	No	No	Yes	Yes	No	No	Yes	Yes
Personal Characteristics	No	No	No	Yes	No	No	No	Yes
R ²	.032	.041	.060	.072	.033	.042	.060	.071
N	2,093	1,970	1,752	1,626	2,130	2,005	1,781	1,654

Notes: Standard errors are in parentheses; *p*-values are shown in brackets. Standard errors adjusted for both heteroskedasticity and within-industry error correlation. All regressions include controls for employer and health plan characteristics and industry income described in the previous tables. "Insurance Demand" includes answers to the following questions: "I am too healthy to need health insurance," "health insurance is not worth the cost," "can get well without a doctor," "medical care easily available without cash," "more than average risk taker," and "how often respondent wears a seat belt." Also included in this category is a variable for self-reported health status, and indicator variables for whether the respondent has been told by a doctor that they have any of eleven medical conditions (stroke, cancer, heart attack, gall bladder disorder, high blood pressure, arteriosclerosis, rheumatism, emphysema, arthritis, diabetes, and heart disease). "Personal Characteristics" includes answers to "health limits kind of work," "weight in pounds," "regular exercise," "can share feelings with someone," "social outings past month," "time respondent was 'very down' last month," "had blood pressure checked last year," "how often respondent brushes teeth," and "respondent's native language is English." Results for variables not shown are available upon request.

substitute for policies sold through employer groups and allows good risks to utilize the private market without having to change employers. As a result, it seems unlikely that any quantitatively important cross subsidies could exist within employer groups in the case of life insurance.²⁵

□ **Robustness checks.** As a supplement to the NMES Household Survey, the Agency for Health Care Policy and Research released an additional data tape (NMES Public Use Tape 9) titled "National Medical Expenditure Survey, 1987: Household Survey, Health Status Questionnaire and Access to Care Supplement." This supplement contains information on the health status, insurance and medical care preferences, and personal attributes of respondents in the Household Survey. Of particular interest for our purposes are the answers to questions grouped in the categories labelled "Insurance Demand" and "Personal Characteristics" in Tables 5–8. Since many of these questions permit multiple responses (often on a five- or six-point scale), we enter all variables with a separate dummy for each possible response.²⁶

In Tables 5 and 6, we examine the robustness of our baseline difference-in-differences models to the inclusion of these additional controls. Columns 1 and 5 present, for ease of comparison, the

²⁵ Moreover, employer-sponsored life insurance policies typically require applicants to submit a medical questionnaire or take a physical exam before increasing their level of coverage. In addition, life insurance premiums are almost universally age rated and are often adjusted based on whether the insured is a smoker, further limiting the scope for cross subsidization within the group.

²⁶ A complete list of these variables is provided at the bottom of Table 5.

TABLE 6 Effect of Job Attachment on Annual Stop-Losses
 Difference-in-Differences Models with Additional Controls
 Dependent Variable: Policyholder's Annual Stop-Loss

	Industry SVP 1980				Industry SVP 1987			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Intercept	89.07 (767.13) [.908]	-17.06 (842.72) [.984]	-379.46 (1829.01) [.836]	-245.16 (2228.58) [.913]	-474.75 (649.50) [.466]	-507.02 (765.27) [.508]	-770.69 (1810.11) [.671]	-1012.54 (2301.41) [.660]
Industry SVP	85.04 (145.24) [.559]	89.46 (142.08) [.530]	90.29 (153.71) [.558]	74.07 (157.52) [.639]	176.88 (133.49) [.187]	205.08 (124.94) [.102]	189.82 (156.74) [.227]	234.91 (152.17) [.124]
Single-Plan Firm	1802.66 (930.64) [.054]	1832.03 (919.72) [.048]	1933.54 (1010.36) [.057]	2212.65 (1116.96) [.049]	1608.06 (831.39) [.054]	1587.70 (833.63) [.058]	1613.49 (943.40) [.089]	2148.21 (954.27) [.026]
Single-Plan Firm × Industry SVP	-372.43 (168.40) [.028]	-381.16 (166.57) [.023]	-399.44 (183.54) [.031]	-459.81 (202.46) [.024]	-307.18 (151.95) [.045]	-304.54 (152.63) [.047]	-308.39 (173.85) [.078]	-410.51 (173.39) [.019]
Demographics	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Insurance Demand	No	No	Yes	Yes	No	No	Yes	Yes
Personal Characteristics	No	No	No	Yes	No	No	No	Yes
R ²	.178	.179	.186	.196	.155	.157	.167	.172
N	1,948	1,848	1,624	1,516	1,984	1,881	1,649	1,539

See notes to Table 5.

baseline estimates from Table 3 (Lifetime Maximum Benefits) and Table 4 (Annual Stop-Losses). In columns 2 and 6 we add the following set of worker-level demographic variables: age, gender, race (black, white, or Hispanic), marital status, family size, union membership, hourly wages, and years of education. In columns 3 and 7 we add the “Insurance Demand” variables and in columns 4 and 8 we add the “Personal Characteristics” variables.

Focusing first on the Maximum Benefit results in Table 5, we see that for both *Industry SVP 1980* and *Industry SVP 1987*, the coefficient estimates for *Single-Plan Firm* and *Single-Plan Firm* × *Industry SVP* become monotonically larger (in absolute value) as each successive set of covariates is added. This suggests that the omission of these additional control variables from our baseline specification did not bias the coefficient estimates on either of our variables of interest, or perhaps biased them downward in absolute value. Adding this set of additional covariates does reduce the precision of our estimates, both because of the large number of variables added (67 in total for columns 4 and 8) and because of a reduction in sample size, but all the estimates based on *Industry SVP 1987* (displayed in columns 5–8) remain statistically significant at the 10% level or better.

Turning to the Stop-Loss results in Table 6, we find again that our main coefficient estimates become larger in absolute value as successive sets of control variables are added. This pattern is more evident for the results based on *Industry SVP 1980* than for the results based on *Industry SVP 1987*, where the coefficient estimates remain roughly constant as the demographic and “Insurance Demand” variables are added but increase in absolute value following the addition of the “Personal Characteristics” variables. Note also that the addition of these variables has little impact on the statistical significance of our Stop-Loss results, with all of the *Single-Plan Firm* and *Single-Plan Firm* × *Industry SVP* coefficients remaining statistically significant at the .10 level and many significant at the .05 level.

Although they are obviously not conclusive, we interpret these findings as providing some evidence against the view that our SVP variables are picking up other worker characteristics that independently influence lifetime benefit limits or annual stop-loss amounts. Were this the case, one would expect that adding a fairly extensive set of controls for worker demographics, insurance

preferences, health status, and personal attributes would begin to attenuate our coefficient estimates. In a similar vein, we believe that the robustness of the estimated coefficient on *Single-Plan Firm* suggests that worker sorting across firms based on the number of offered plans, to the extent that it occurs, is perhaps not systematically related to the specific contract provisions we study.

□ **Falsification tests.** Because the previous regressions may still omit relevant explanatory variables, it is useful to supplement those results with findings from a pair of falsification tests. We begin by looking at the insurance purchases of the self-employed. This is a group for whom job attachment is unlikely to have much of an effect on their ability to obtain insurance²⁷ but, based on their being employed, should be at least somewhat similar to other workers. In our sample of self-employed persons, 62% hold a private health insurance policy and 75% have coverage from a private policy (either their own or someone else's). These numbers are not dramatically different from the figures for the non-self-employed population: 78% and 82%, respectively. In addition, self-employed workers in our sample exhibit approximately the same mean, variance, and range of *Industry SVP* as the workers in our earlier samples. Thus, if *Industry SVP* were a marker for the demand for health insurance, we would expect it to significantly influence the likelihood that self-employed individuals would either hold a private policy or obtain private health insurance from another source.²⁸

To test this proposition, we examine all self-employed individuals in the NMES Household Survey who are currently employed, who are not working in a self-described "retirement job," and who do not have health insurance coverage from a government program. Results for both probit and linear probability (LPM) models are displayed in Table 7. All specifications include a set of controls for worker demographics (listed at the bottom of Table 7) and industry income. In columns 2 and 5 we add the "Insurance Demand" variables and in columns 3 and 6 we add the "Personal Characteristics" variables. For the probit models, marginal effects are shown in braces directly beneath the coefficient estimates.

The regression results reported in Table 7 provide clear and consistent evidence that neither *Industry SVP 1980* nor *Industry SVP 1987* influences the probability that a given individual holds a private health insurance policy, either in statistical or in economic terms. The results are virtually identical if the dependent variable is set equal to one if the respondent simply has coverage from a private health insurance policy. Thus, at least among the self-employed, *Industry SVP* does not appear to proxy for unobservable factors that determine an individual's demand for health insurance.

As discussed earlier, another potential confounder of our baseline results is the possibility that workers with high SVP values are more valuable to firms than workers with low SVP values, for reasons unrelated to their training. To the extent that this is true, one would expect firms that rely heavily on such workers to provide benefits designed to attract or retain workers that possess the desired traits. One might argue that health insurance is such a benefit, and that offering more generous coverage, in the form of higher lifetime benefit limits or lower annual stop-losses, serves as an inducement to these workers. We test this hypothesis by examining whether *Industry SVP* influences the decision to offer life insurance among the same set of firms that made up our base sample for the lifetime limit and stop-loss regressions.

We chose life insurance because it is the employee benefit that most closely resembles health insurance and, as such, should appeal to the same types of workers as health insurance. Unfortunately, we do not have the same kind of detailed information on the provision of life insurance policies as we had for health insurance policies. However, in contrast to health insurance,

²⁷ With individual policies, there is no pooling of health risks. Zellers, McLaughlin, and Frick (1992) find that at the time of our survey, the "small group" market for health insurance functioned in much the same way as the market for individually purchased policies, with similar restrictions on enrollment and coverage.

²⁸ One potential concern is that the self-employed may have lower demand for insurance than other workers. This need not concern us as long as there is not a good reason to believe that the relationship between *Industry SVP* and insurance demand is different for the two groups. Our results for the self-employed are also not sensitive to the inclusion of the "Insurance Demand" and "Personal Characteristics" variables described earlier.

TABLE 7 Insurance Purchases Among the Self-Employed
 Dependent Variable = 1 if Respondent Holds
 Private Health Insurance Policy, 0 If Uninsured

	Industry SVP 1980			Industry SVP 1987		
	(1)	(2)	(3)	(4)	(5)	(6)
Industry SVP (Probit)	-.006 {-.002} (.089) [.948]	.003 {.001} (.087) [.972]	.035 {.012} (.083) [.676]	.012 {.005} (.089) [.890]	.019 {.007} (.087) [.831]	-.005 {-.002} (.092) [.957]
Industry SVP (LPM)	-.007 (.029) [.812]	-.002 (.026) [.928]	.006 (.023) [.799]	.004 (.030) [.905]	.007 (.028) [.810]	.002 (.028) [.945]
Insurance Demand	No	Yes	Yes	No	Yes	Yes
Personal Characteristics	No	No	Yes	No	No	Yes
R ² (Probit)	.158	.275	.328	.160	.268	.320
R ² (LPM)	.191	.312	.357	.193	.304	.347
N	923	764	705	944	780	720

Notes: For the probit models, marginal effects are shown in braces below the coefficient estimates. Standard errors are in parentheses; *p*-values shown in brackets. Standard errors adjusted for both heteroskedasticity and within-industry error correlation. All regressions include controls for gender, age, race (white, black, or Hispanic), region (northeast, midwest, south, or west), urban location (SMSA or non-SMSA), marital status, family size, education, and industry income. Results for variables not shown are available upon request. Data on establishment size, union membership, and hourly wages were not available for self-employed individuals.

the NMES Survey does contain a question that specifically asks whether the employer *offers* life insurance. Using this question, we create a dummy variable equal to one if the firm offers life insurance and zero otherwise. Approximately 92% of the roughly 3,700 firms in our baseline estimation sample offer life insurance.

We examine the effect of *Industry SVP* on life insurance provision using exactly the same set of firm-level control variables as were used in our lifetime limit and stop-loss regressions. Results for both probit and linear probability models are presented in Table 8. Columns 2 and 6 add worker demographics, columns 3 and 7 add our “Insurance Demand” variables, and columns 4 and 8 add the “Personal Characteristics” variables. For the probit models, marginal effects are displayed in braces directly below the coefficient estimates. Looking first at columns 1–4, we see that *Industry SVP 1980* has essentially no effect on the likelihood that a firm offers life insurance. One may worry that this is due to the fact that a large majority of firms in our sample offer life insurance, leaving little variation in the dependent variable. However, many of the control variables (not shown, but available upon request) do exert a statistically significant effect on the probability that a firm offers life insurance.²⁹ By contrast, the marginal effect of *Industry SVP 1980* is always below 1%, and in the specifications with more control variables, columns 3 and 4, it is substantially smaller. Turning next to *Industry SVP 1987*, we find that the marginal effects are a bit larger, on the order of 1%, but are *negatively* signed in all specifications. Interestingly, this negative effect becomes statistically significant in columns 7 and 8 as additional covariates are added, which provides further evidence that there is enough variation in the dependent variable to detect significant effects.

Taken together, the results from this section provide evidence against two alternative hypotheses that could potentially confound our results. The fact that, among the self-employed, *Industry SVP* does not predict either health insurance purchases or the propensity to obtain

²⁹ In every specification in Table 8, the coefficients on the employer-size dummies are monotonically increasing and are statistically significant at the .01 level or better. Other variables that are statistically significant in most or all of the specifications are unionization, type of employer, and region.

TABLE 8 **Probability Employer Offers Life Insurance**
Dependent Variable = 1 if Respondent's
Employer Offers Life Insurance, 0 Otherwise

	Industry SVP 1980				Industry SVP 1987			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Industry SVP (Probit)	.059 {.006} (.058) [.307]	-.016 {-.001} (.060) [.796]	.002 {.0002} (.063) [.970]	.013 {.0009} (.059) [.833]	-.079 {-.008} (.071) [.268]	-.128 {-.011} (.076) [.094]	-.169 {-.013} (.075) [.024]	-.156 {-.011} (.079) [.049]
Industry SVP (LPM)	.009 (.007) [.205]	-.000 (.007) [.995]	.001 (.007) [.852]	.002 (.007) [.754]	-.009 (.009) [.329]	-.015 (.009) [.115]	-.018 (.009) [.047]	-.017 (.009) [.071]
Demographics	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Insurance Demand	No	No	Yes	Yes	No	No	Yes	Yes
Personal Characteristics	No	No	No	Yes	No	No	No	Yes
R ² (Probit)	.157	.163	.190	.214	.165	.166	.197	.219
R ² (LPM)	.091	.085	.097	.106	.096	.086	.099	.108
N	3,708	3,489	3,076	2,847	3,778	3,557	3,133	2,901

Notes: For probit models, marginal effects are in braces below the coefficient estimates. Standard errors are in parentheses; *p*-values shown in brackets. Standard errors adjusted for both heteroskedasticity and within-industry error correlation. All regressions include controls for type of employer (for profit, nonprofit, government, or other), region (northeast, midwest, south, or west), urban location (SMSA or non-SMSA), establishment size (indicator variables for: <10 employees, 10–25 employees, 26–100 employees, 101–500 employees, >500 employees), employer unionization (fully unionized, partially unionized, or non-union), and industry income. Results for variables not shown are available upon request.

coverage from another source indicates that, at least among this group, SVP is not serving as a marker for insurance demand. Similarly, the fact that firms with high levels of *Industry SVP* are no more prone to offer life insurance than other firms suggests that it is unlikely that any additional compensation paid to high-SVP workers comes in the form of higher lifetime health insurance benefits or lower annual stop-loss amounts.

5. Conclusions

■ In many insurance settings it may be argued that private information, and the attendant problem of adverse selection, is the primary cause of market failure. The evidence increasingly suggests, however, that the market for health insurance is more accurately characterized as one in which participants possess symmetric information about the evolving health status of insurance purchasers. The key ingredient for a successful health insurance package is to have those who turn out to be healthier than average agree to cross subsidize the members of the pool who are less fortunate. But, as illustrated by the increasing use of re-underwriting in the individual health insurance market, the inability of low risks to precommit to remain in the insurance pool may eliminate the ability to insure against the classification risk that results when premiums are adjusted to reflect observed health status.

The evidence presented in this article demonstrates the importance of precommitment in the provision of long-term health insurance. When insurance is bundled with employment, we find that frictions in worker mobility impart a *de facto* commitment to employment-based insurance pools that can be exploited to provide more complete insurance of health risks than would be available in a competitive market. In the absence of such a commitment device, insurers anticipate that the healthier individuals would, over time, exit the insurance pool to obtain coverage elsewhere at more favorable terms. Thus, the ability of insurers to offer benefit packages that insure individuals against the classification risk associated with changes in their health status is enhanced when workers experience an attachment to their job that impedes their exit from the pool.

Consistent with this view, we find that the amount of coverage provided under employment-based health insurance contracts is increasing in the degree of worker commitment. The measures of coverage that we consider in the analysis are the lifetime coverage cap and the stop-loss amounts, since these are the contract provisions most likely to be encountered by individuals who have chronic, long-term illnesses, and whose costs of treatment are likely to result in higher premiums for the healthier members of the group. As predicted by the theoretical model, we find that the distortions in coverage are most pronounced in cases where firms offer a single insurance contract to all of their workers. In contrast, the coverages associated with multiple-plan benefit packages, which permit the voluntary experience-rating of premiums through the selection of contracts by insureds, are much less sensitive to job attachment.

Our results indicate that in addition to the advantages conferred by preferential tax treatment, employer-sponsored health insurance is also attractive because the bundling of insurance with employment results in a more stable, and therefore insurable, risk pool. If employment were to become less stable or the incentive to invest in job-specific skills were to decline, insurers would need to utilize institutions other than employer sponsorship to craft insurable risk pools. From a policy perspective, our results suggest caution when considering proposals that might weaken the link between health insurance and employment.

At a broader level, our findings support the view that limitations on insured commitment are a potentially important impediment to the implementation of the long-term contracts required to fully insure chronic illnesses. Thus, our results also provide a glimpse into the often hidden costs of unreliable contract enforcement.

Appendix

■ The proof of Proposition 1 and the complete definition of SVP follow.

Proof of Proposition 1. Since $R^H = S$, it follows immediately that $dR^P/dK > dR^H/dK = 0$. The rest of the proof will demonstrate that $dR^P/dK > dR^L/dK$.

The contract C^P is characterized by the conditions

$$Z^P - \bar{p}R^P = 0 \tag{A1}$$

and

$$p^L U(W_S^P) + (1 - p^L)U(W_N^P) - U(W_0) = 0, \tag{A2}$$

where $W_S^P \equiv Y - Z^P - S + R^P$, $W_N^P \equiv Y - Z^P$, and $W_0 \equiv Y - K - p^L S$. Total differentiation of (A1) and (A2) yields the result

$$\frac{dR^P}{dK} = \frac{U'(W_0)}{\bar{p}(1 - p^L)U'(W_N^P) - p^L(1 - \bar{p})U'(W_S^P)}. \tag{A3}$$

It is straightforward to demonstrate that (A3) is positive at C^P , since \bar{V}^L is flatter than the pooling zero profit locus, (A1).

The contract C^L is characterized by the conditions

$$\lambda(Z^H - p^H S) + (1 - \lambda)(Z^L - p^L R^L) = 0, \tag{A4}$$

$$U(W^H) - p^H U(W_S^L) - (1 - p^H)U(W_N^L) = 0, \tag{A5}$$

and

$$p^L U(W_S^L) - (1 - p^L)U(W_N^L) - U(W_0) = 0, \tag{A6}$$

where $W^H \equiv Y - Z^H$, $W_S^L \equiv Y - Z^L - S + R^L$, and $W_N^L \equiv Y - Z^L$. Solving (A4) for Z^H , substituting the result into (A5), and totally differentiating yields

$$\frac{dZ^L}{dR^L} = \frac{p^L(1 - \lambda)U'(W^H) + \lambda p^H U'(W_S^L)}{(1 - \lambda)U'(W^H) + \lambda p^H U'(W_S^L) + \lambda(1 - p^H)U'(W_N^L)}, \tag{A7}$$

which is the slope of the AC^* locus depicted in Figure 4. Total differentiation of (A6) gives the result that

$$\frac{dR^L}{dK} = \frac{U'(W_0)}{\frac{dZ^L}{dR^L} \left[p^L U'(W_S^L) + (1 - p^L) U'(W_N^L) \right] - p^L U'(W_S^L)}. \quad (A8)$$

To demonstrate the desired result, we must show that the denominator of (A8) is greater than the denominator of (A3). This difference may be written as

$$E_1 \left[\frac{dZ^L}{dR^L} - \frac{p^L U'(W_S^L)}{E_1} - \frac{E_2}{E_1} \left[\bar{p} - \frac{p^L U'(W_S^P)}{E_2} \right] \right], \quad (A9)$$

where $E_1 \equiv p^L U'(W_S^L) + (1 - p^L) U'(W_N^L)$ and $E_2 = p^L U'(W_S^P) + (1 - p^L) U'(W_N^P)$.

By (A2) and (A6), it follows that $E_2/E_1 = 1$, so that (A9) reduces to

$$E_1 \left[\frac{dZ^L}{dR^L} - \bar{p} + \left[\frac{p^L U'(W_S^P)}{E_2} - \frac{p^L U'(W_S^L)}{E_1} \right] \right]. \quad (A10)$$

Since the L -type's indifference curve is steeper at C^P than at C^L , the term in the inner set of brackets is positive. Thus, to complete the proof, it is sufficient to demonstrate that $dZ^L/dR^L > \bar{p}$.

$$\frac{dZ^L}{dR^L} - \bar{p} = \frac{p^L(1 - \lambda)U'(W^H) + \lambda p^H U'(W_S^L) - \bar{p}E_3}{E_3}, \quad (A11)$$

where $E_3 \equiv (1 - \lambda)U'(W^H) + p^H U'(W_S^L) + \lambda(1 - p^H)U'(W_N^L)$. The numerator in (A11) may be written as

$$U'(W^H)\lambda(1 - \lambda)(p^L - p^H) + \lambda p^H U'(W_S^L)(1 - \bar{p}) - \lambda(1 - p^H)U'(W_N^L)\bar{p},$$

which is greater than

$$U'(W^H)\lambda(1 - \lambda)(p^L - p^H) + \lambda p^H U'(W_S^L)(1 - \bar{p}) - \lambda(1 - p^H)U'(W_S^L)\bar{p}, \quad (A12)$$

since $U'(W_S^L) > U'(W_N^L)$. We may then write (A12) as

$$\lambda(1 - \lambda)(p^L - p^H) \left[U'(W_S^L) - U'(W^H) \right] > 0,$$

since $p^H > p^L$, and the concavity of U and (A5) imply that $U'(W_S^L) > U'(W^H)$. Thus, $E_3 > 0$ implies that (A11) is positive, and the proof of Proposition 1 is complete. *Q.E.D.*

□ **Definition of SVP.** (From the 1977 *Dictionary of Occupational Titles*): This [SVP] represents the amount of time required to learn the techniques, acquire information, and develop the facility needed for average performance in a specific job-worker situation. The training may be acquired in a school, work, military, institutional, or a vocational environment. It does not include orientation training required of even every fully qualified worker to become accustomed to the special conditions of any new job. Specific vocational training includes training given in any of the following circumstances: (a) Vocational education (such as high school commercial or shop training, technical school, art school, and that part of college training which is organized around a specific vocational objective); (b) Apprentice training (for apprenticeable jobs only); (c) In-plant training (given by an employer in the form of organized classroom study); (d) On-the-job training (serving as learner or trainee on the job under the instruction of a qualified worker); (e) Essential experience in other jobs (serving in less responsible jobs which lead to the higher grade job or serving in other jobs that qualify).

References

- ALTMAN, D., CUTLER, D.M., AND ZECKHAUSER, R.J. "Adverse Selection and Adverse Retention." *American Economic Review*, Vol. 88 (1998), pp. 122–126.
- CARDON, J.H. AND HENDEL, I. "Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey." *RAND Journal of Economics*, Vol. 32 (2001), pp. 408–427.
- CHIAPPORI, P.-A. "Econometric Models of Insurance under Asymmetric Information." In G. Dionne, ed., *Handbook of Insurance*. London: Kluwer Academic, 2000.

- COCHRANE, J.H. "Time-Consistent Health Insurance." *Journal of Political Economy*, Vol. 103 (1995), pp. 445–473.
- CROCKER, K.J. AND SNOW, A. "The Efficiency Effects of Categorical Discrimination in the Insurance Industry." *Journal of Political Economy*, Vol. 94 (1986), pp. 321–344.
- CUTLER, D.M. AND ZECKHAUSER, R.J. "The Anatomy of Health Insurance." In A.J. Culyer and J.P. Newhouse, eds., *Handbook of Health Economics*, Vol. 1. New York: Elsevier, 2000.
- DIAMOND, P. "Organizing the Health Insurance Market." *Econometrica*, Vol. 60 (1992), pp. 1233–1254.
- EHRENBERG, R.G. AND SMITH, R.S. *Modern Labor Economics: Theory and Public Policy*. Reading, Mass.: Addison-Wesley, 2003.
- EPSTEIN, R.A. *Mortal Peril: Our Inalienable Right to Health Care?* Reading, Mass.: Addison-Wesley, 1997.
- GALE, W.G. "The Effects of Pensions on Household Wealth: A Reevaluation of Theory and Evidence." *Journal of Political Economy*, Vol. 106 (1998), pp. 706–723.
- GRUBER, J. AND MADRIAN, B.C. "Health Insurance and Job Mobility: The Effects of Public Policy on Job-Lock." *Industrial and Labor Relations Review*, Vol. 48 (1994), pp. 86–102.
- HENDEL, I. AND LIZZERI, A. "The Role of Commitment in Dynamic Contracts: Evidence from Life Insurance." *Quarterly Journal of Economics*, Vol. 118 (2003), pp. 299–327.
- HERRING, B.J. AND PAULY, M.V. "Incentive-Compatible Guaranteed Renewable Health Insurance Premiums." Mimeo, Wharton School of Business, University of Pennsylvania, 2003.
- MOULTON, B.R. "Random Group Effects and the Precision of Regression Estimates." *Journal of Econometrics*, Vol. 32 (1986), pp. 385–397.
- PAULY, M.V. "Taxation, Health Insurance, and Market Failure in the Medical Economy." *Journal of Economic Literature*, Vol. 24 (1986), pp. 629–675.
- , KUNREUTHER, H., AND HIRTH, R. "Guaranteed Renewability in Insurance." *Journal of Risk and Uncertainty*, Vol. 10 (1995), pp. 143–156.
- POTERBA, J.M., VENTI, S.F., AND WISE, D.A. "Do 401(k) Contributions Crowd Out Other Personal Saving?" *Journal of Public Economics*, Vol. 58 (1995), pp. 1–32.
- ROTHSCHILD, M. AND STIGLITZ, J.E. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics*, Vol. 90 (1976), pp. 630–649.
- SHAVELL, S. "Risk Sharing and Incentives in the Principal and Agent Relationship." *Bell Journal of Economics*, Vol. 10 (1979), pp. 55–73.
- TERHUNE, C. "Health Insurer's Premium Practices Add to Profits, Roil Customers." *Wall Street Journal*, April 9, 2002, p. A1.
- TOPEL, R.H. "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority." *Journal of Political Economy*, Vol. 99 (1991), pp. 145–176.
- U.S. DEPARTMENT OF COMMERCE, BUREAU OF THE CENSUS. *Classified Index of Industries and Occupations—Census of Population*. Washington, D.C.: U.S. Government Printing Office, 1980.
- . "Census of Population and Housing, 1980 [United States]: Public Use Microdata Sample (A Sample): 5 Percent Sample [Puerto Rico]" [Computer file]. Ann Arbor, Mich.: Inter-University Consortium for Political and Social Research, 1985.
- U.S. DEPARTMENT OF COMMERCE, OFFICE OF FEDERAL STATISTICAL POLICY AND STANDARDS. *Standard Occupational Classification Manual*. Washington, D.C.: U.S. Government Printing Office, 1980.
- U.S. DEPARTMENT OF HEALTH AND HUMAN SERVICES, AGENCY FOR HEALTH CARE POLICY AND RESEARCH. "National Medical Expenditure Survey, 1987: Household Survey, Health Status Questionnaire and Access to Care Supplement [Public Use Tape 9]" [Computer File]. Ann Arbor, Mich.: Inter-University Consortium for Political and Social Research, 1991.
- . "National Medical Expenditure Survey, 1987: Household Survey, Population Characteristics and Person-Level Utilization, Rounds 1–4 [Public Use Tape 13]" [Computer File]. Ann Arbor, Mich.: Inter-University Consortium for Political and Social Research, 1992.
- . "National Medical Expenditure Survey, 1987: Policyholders of Private Insurance: Premiums, Payment Sources, and Types and Source of Coverage [Public Use Tape 15]" [Computer File]. Ann Arbor, Mich.: Inter-University Consortium for Political and Social Research, 1993a.
- . "National Medical Expenditure Survey, 1987: Health Insurance Plans Survey Data, Private Insurance Benefit Database and Linkages to Household Survey Policyholders [Public Use Tape 16]" [Computer File]. Ann Arbor, Mich.: Inter-University Consortium for Political and Social Research, 1993b.
- U.S. DEPARTMENT OF LABOR, EMPLOYMENT AND TRAINING ADMINISTRATION, U.S. EMPLOYMENT SERVICE. *Dictionary of Occupational Titles*, 4th ed. Washington, D.C.: U.S. Government Printing Office, 1977.
- . *Selected Characteristics of Occupations Defined in the Dictionary of Occupational Titles*. Washington, D.C.: U.S. Government Printing Office, 1981.
- WHITE, H. "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." *Econometrica*, Vol. 48 (1980), pp. 817–838.
- ZELLER, W.K., McLAUGHLIN, C.G., AND FRICK, K.D. "Small-Business Health Insurance: Only the Healthy Need Apply." *Health Affairs*, Vol. 11 (1992), pp. 174–180.