Optimal Risk Adjustment in Markets with Adverse Selection: An Application to Managed Care

By JACOB GLAZER AND THOMAS G. MCGUIRE*

It is well known that adverse selection causes distortions in contracts in markets with asymmetric information. Taxing inefficient contracts and subsidizing the efficient ones can improve market outcomes (Bruce C. Greenwald and Joseph E. Stiglitz, 1986), although regulators rarely seem to implement tax and subsidy schemes with adverse-selection motives in mind. Contracts are often complex and "incomplete," and it is the "inefficient" elements of the contract that are difficult to verify and hence tax or subsidize. This is precisely the reason that in health insurance markets, rather than subsidizing *contracts,* regulators and payers contend with adverse selection by taxing and subsidizing the price paid to insuring health plans on the basis of observable characteristics of the *persons* joining the plan—a practice known as "risk adjustment."

Risk-adjusted premiums are paid to "managed-care" plans—plans that ration care by management, rather than by conventional approaches like coinsurance and deductibles, and offer a bundle of characteristics (quality, access for many services) that is fundamentally outside the scope of direct regulation. Selection-related incentives threaten the efficiency and fairness of this organization of health insurance markets by inducing plans to distort the quality of the services they offer to discourage high-cost persons from joining the plan. As managed care becomes the predominant source of health care for residents of the United States and many other countries, payers attempt to address this incentive by setting a risk-adjusted price that pays more for more-expensive enrollees.[1]

As it is conventionally practiced, risk adjustment sets prices for people proportional to their expected cost based on observable characteristics. The federal Medicare program, for example, has used age, sex, welfare status, and county-of-residence adjusters to set prices to managed-care plans.[2] To convey how what we term "conventional" risk adjustment works, suppose age is the risk adjuster for a Medicare population over 65. If it is determined that the 75- to 84-year-old population costs 20 percent more than the overall average in Medicare, the assumption in conventional risk adjustment is that the premium paid to plans for someone in this group should be 20 percent above the average.

We fundamentally disagree that this is the right way to think about and do risk adjustment.

* Glazer: Faculty of Management, Tel Aviv University, and the Gertner Institute for Health Policy Research, Tel Aviv, Israel (e-mail: glazer@post.tau.ac.il); McGuire: Department of Economics, Boston University, Boston, MA 02215 (e-mail: tmcguire@bu.edu). McGuire's research was supported by Grant K05 MH01263 from the National Institute of Mental Health. We thank Randy Ellis, Richard Frank, Marty Gaynor, Hsien-Ming Ling, Albert Ma, Joseph Newhouse, Michael Riordan, and Wynand P. M. M. van de Ven for helpful discussions. Advice from two referees is gratefully acknowledged. The authors alone are responsible for the analysis and conclusions.

[1] For representative discussions in the U.S. contexts, see Joseph P. Newhouse (1994), David M. Cutler (1995), and Alain C. Enthoven and Sarah J. Singer (1995). Risk adjustment was to be part of President Clinton's proposed health-care reform. See also Netanyahu Commission (1990) for Israel, Rene C. J. A. van Vliet and Wynand P. M. M. van de Ven (1992) for The Netherlands, and Donald W. Light (1998) for Ireland. Risk adjustment is an integral part of many state-based health-reform proposals centering on the poor and uninsured. For a discussion of these reforms, see John Holohan et al. (1995).

[2] Medicare's risk-adjustment system is called the average adjusted per capita cost (AAPCC), and is used to pay HMOs for Medicare beneficiaries that choose to join. Medicare calculates the expected cost in the unmanaged fee-for-service sector for a beneficiary based on the above-mentioned characteristics, and then pays 95 percent of this to the HMO. A substantial amount of favorable risk selection by HMOs not captured by these factors has taken place within the Medicare program. See Harold S. Luft (1995) and Katherine Swartz (1995). Medicare is revising its risk-adjustment policy to moderate the impact of geographic adjusters, and to add indicators of diagnosis from previous hospitalizations. See Medicare Payment Advisory Commission (1998).

We argue that risk adjustment should be viewed as a way to set prices for different individuals, to address adverse-selection problems. Solving for prices in an explicit optimizing framework, rather than simply paying according to average costs, magnifies the power of risk-adjustment policies. Specifically, we show in this paper that conventional risk adjustment is never the optimal policy for a regulator in a market with asymmetric information. Furthermore, we solve for alternative weights on persons' observable characteristics that improve the efficiency of the market for health plans in relation to conventional risk adjustment.

This paper presents a model of the market for managed-care plans in which an individual's true health-cost risk, which we refer to as an individual's "type," is private information. A regulator can use observable characteristics (e.g., age), which are correlated with type and which we refer to as "signals," to pay plans. In these terms, conventional risk adjustment pays, for each enrollee, expected cost given the signal. Conventional risk adjustment attenuates the inefficiency that results from adverse selection, but does not make the best use of the information contained in the signal. Optimal risk adjustment pays higher than conventional risk adjustment for persons with the "bad" signal (the old), and lower than conventional risk adjustment for persons with the "good" signal (the young), a result that we contend is general and of practical importance to health policy in the United States and other countries. "Optimal" risk adjustment, of course, can fully solve the adverse-selection problem only in a simple model. We present such a model to clarify the properties of the idea we are proposing. The main point of the paper, however, is this: in comparison to conventional risk adjustment, the "overpayment" policy we advocate can improve the efficiency properties of health insurance markets in realistic contexts.

The intuition behind our finding is as follows: Consider a plan's incentive to provide a service that might attract the sick "types" in the population. The "service" might be care for cancer, which could be done at a higher or lower quality. The plan evaluates the costs and revenues brought in by providing a higher-quality service. Conventional risk adjustment, because of the weak empirical association of costs with signals at the individual level, cannot do much to raise the premiums paid for the high-cost types who value cancer care. (Medicare's four variables mentioned earlier, for instance, explain only about 1 percent of health cost variance in the elderly.) Under conventional adjustment, the plan is likely to have strong incentives to underprovide care for cancer to discourage high-cost types from joining. Note, however, that the group of high-cost types contains relatively more old people. Optimal risk adjustment can pay more for enrolling high-cost types by giving a heavy weight ("overpaying") for the old, thereby rewarding the plan for spending on cancer care. Our simple insight is that the payment weight on age may be chosen for its incentive properties, and need not—indeed should not—be the same as the coefficient on age from a regression explaining average cost.

I. Risk Adjustment and Managed Care

Recently, managed care has supplanted coverage policy (deductibles, copayments, limits, coverage exclusions) as the main control on moral hazard in employment-based health insurance and many public programs.[3] Managed care is a set of strategies health plans use to direct the quality and quantity of health care provided to their enrollees, including limitation of the hospitals and physicians a patient may see for treatment or drugs the patient may take, specification of clinical protocols to be followed in the case of certain illnesses, application of criteria for access to services, limitations on authorized length of stay or visits, and so on (Institute of Medicine, 1989). A managed-care plan integrates the health insurance and health care functions (Joseph P. Newhouse, 1996) and, as the name suggests, influences the quality of

[3] Virtually all private health insurance and Medicaid plans include some elements of managed care. About 75 percent of the privately insured population are in managed care according to Gail A. Jensen et al. (1997). The Health Care Financing Administration's website reports that as of 1996, 40 percent of the Medicaid population were in managed care, up from 29 percent just the year before. The largest group that has been slow to enter managed care is the elderly in Medicare, just 14 percent at the end of 1997 (Medicare Payment Advisory Commission, 1998) but this number also has been increasing rapidly of late.

care through internal management processes largely immune from direct regulation.

Even though most of the literature emphasizes managed care as a device to control moral hazard [for a recent empirical study, see Dana Goldman et al. (1995)], a health plan can also use managed care to affect the plan's risk selection.[4] Newhouse (1996) writes that, in spite of regulations requiring plans to offer "open enrollment," plans

> Can alter their product to influence [enrollee] choice. ... Their staffing may discourage some types of risks and encourage others; for example, they may stint on oncologists (cancer specialists) but have numerous pediatricians (families with children are better risks). Staffing choices seem especially hard to regulate, because of numerous sensible opportunities for substituting less highly trained personnel for specialists. They could offer incentives to gatekeeper physicians not to refer patients to specialists, thereby discouraging enrollment by the chronically ill.

Another form of adverse selection takes place when a health plan discriminates in favor of or against particular applicants. If a plan knows that the expected costs of an applicant are different from the premium the plan gets if the person joins, the plan may encourage or discourage that individual from joining. An extreme example is simply denying enrollment to an applicant likely to be unprofitable. Seeking favorable risks is often referred to as "cream skimming" (see van de Ven and van Vliet, 1992), whereas seeking to shed bad risks is called "dumping" (Ching-to Albert Ma, 1994; Randall P. Ellis, 1998). Regulation can forbid these risk-selection activities, by, for example, requiring open enrollment, but subtle forms of risk selection of this form are probably hard to eradicate.

[4] Attendees at a conference on risk selection noted that a health plan has an incentive to select professional staff (including type and number) and invest in facilities to encourage favorable selection (Anne K. Gauthier et al., 1995). Luft (1995, p. 27) observes that regulation can deal with some elements of a plan's tactics to attract good risks, like coverage policy, but regulation cannot prevent competition for good risks on the basis of "quality."

The presence of adverse selection in health insurance has been recognized for a long time, and evidence continues to accumulate attesting to its empirical importance. Data from the federal Medicare program indicate that the typical Medicare beneficiary joining a managed-care plan has costs 35 percent lower in the year before joining than his/her nonjoining counterpart; the typical beneficiary leaving a plan has costs 60 percent higher than a nonjoiner (Physician Payment Review Commission, 1997). In the case of a private employer, David M. Cutler and Sarah Reber (1998) contend that the most generous plan fell victim to an adverse-selection-induced "premium death spiral." At the same time, there have been few formal analyses of adverse selection in a context of competing health plans, and of the role of risk adjustment in correcting misallocations. Cutler and Richard Zeckhauser (1998), in an environment without risk adjustment, analyze the employer's problem of setting contributions to health plan choices, when employees differ in their costs and in their tastes for more- or less-generous plans. Plans' characteristics are regarded as given, and enrollment prices are assumed to be set by plans at the average cost of all enrollees in a plan. Employees joining a plan pay the difference between the plan's average cost and the employer's contribution to the plan. The efficiency issue addressed in that paper is how the employer's subsidy policy serves enrollees' taste to sort among more- or less-generous plans. In an observation akin to Greenwald and Stigliz (1986), Cutler and Zeckhauser note the efficiency value of an employer contribution that varies with plan generosity. Subsidizing expensive plans substitutes for risk adjustment, because it ends up paying more for more-expensive employees who tend to choose the generous plan. Tracy Lewis and David Sappington (1996) consider the incentives to cream skim when a health plan may gain private information about a person's health risk by incurring a screening cost. If this screening cost is sufficiently low, Lewis and Sappington show that it will be optimal for the principal to offer a menu of contracts to the plan to induce the plan to reveal the true "type" of its enrollees. Their analysis applies to a much different information structure and regulatory environment than that considered here.

The literature on risk adjustment consists almost exclusively of empirical research on the statistical determinants of health costs.[5] Current research on risk adjusters focuses on clinical information such as diagnosis and health care use in past periods as well as demographics. See, for example, Ellis et al. (1996). Empirical models that use lagged values of components of past health care costs explain less than 10 percent of the variance in health care expenditures at the individual level, leaving many observers pessimistic about how well conventional risk adjustment can deal with selection incentives.[6]

II. The Model

Suppose that there are two types of consumers L and H, who can contract two illnesses a and c. Illness a we call an acute illness and both types of people have the same probability of contracting this illness. To simplify the presentation, we normalize the probability of each type contracting illness a to 1. The two types are distinguished in their probability of contracting the chronic illness, illness c. Let P_i, $i \in \{H, L\}$, denote the probability that a person of type i contracts illness c. Then, $P_H > P_L > 0$. The proportion of H types in the population is λ.

If a person (of either type) has illness j, $j \in \{a, c\}$, his/her utility from treatment will be increased by $v_j(m_j)$, where m_j is the dollar value of resources devoted to treat this illness, $v_j' > 0$ and $v_j'' < 0$. Thus, we make the simplifying assumption that the benefits from treatment are independent of one another and the same for all people. If a person has both illnesses, utility will simply be increased by $v_a(m_a) + v_c(m_c)$.

We assume plans have no copayments to focus on the key aspect of managed care. A contract will thus be of the form (m_a, m_c, r), where r is the premium the consumer pays. Later on, we allow for a regulatory policy in which the consumers do not pay premiums directly to plans. In such a case, regulation can introduce a difference between the premium the *consumer pays* and the premium the *plan receives* for that consumer. In such cases, the "contract" (m_a, m_c, r) will always refer to the premium the consumer pays.

Given some contract (m_a, m_c, r), type i's expected utility, if he/she chooses this contract, will be

$$(1)\quad V_i(m_a, m_c, r) = v_a(m_a) + P_i v_c(m_c) - r, \text{for } i = H, L.$$

We thus assume no income effects (and no risk aversion).[7]

A. *The Socially Desired Contract*

The regulator is concerned with the efficiency and equity of markets for health insurance. Efficiency will be judged in terms of the degree of managed care (which in turn determines treatment). The health plan sets the level of care that will be provided to a patient if he/she becomes ill. With respect to equity, the regulator's goal is to distribute the burden of health care costs equally among the healthy and the sick.[8] (We ignore differences in income in our model.) "Community rating" health insurance is one way to achieve this objective. Equalizing the cost of health insurance to all could be viewed as a form of social insurance, insuring individuals against the risk of being an unhealthy type, and reinterpreted as an efficiency objective. Having noted this interpretation, however, we

[5] For a summary of some of the empirical literature on risk adjustment, see van de Ven and Ellis (2000).

[6] A risk-adjustment system needs to explain only the "predictable" part of the variance in health costs. Consumers or insurers can act only on something they know. Existing empirical risk adjusters, however, are not regarded as coming very close to this standard. See Newhouse (1996) for review and discussion.

[7] We disregard the option to buy no insurance contract at all, and simply assume that premiums never drive consumers out of the pool. Such an assumption readily applies to cases of national health policy where the "premium" to the consumer takes the form of a compulsory tax (and cannot be avoided by not choosing a plan) or to an employer's health benefit plan where the worker's "premium" takes the form of a fringe benefit (and also cannot be avoided).

[8] This involves a redistribution accomplished through taxes or some other collective financing mechanism between the healthy and the sick. Our results about optimal risk adjustment would continue to hold for other feasible redistributions, including no redistribution through regulation.

will continue to refer to equalizing the financial burden of illness as an "equity" objective.

The efficient level of managed care equalizes the marginal benefit of treatment to its marginal cost, 1. Thus, we define m_a^* and m_c^*, the efficient levels of treatment for the two diseases as follows:

$$(2) \qquad v_a'(m_a^*) = 1, \quad v_c'(m_c^*) = 1.$$

High- and low-risk types have different probabilities of becoming ill, but once ill, receive the same utility from treatment. Thus, the efficient level of managed care and treatment for each illness is the same for both types.

Next, define the premium r^*:

$$(3) \quad r^* = m_a^* + [\lambda P_H + (1-\lambda)P_L]m_c^*.$$

r^* is the cost of a contract offering managed care of (m_a^*, m_c^*) averaged across the entire population.

The efficient and fair allocation will be called the "socially desired" allocation. The *socially desired contract* is (m_a^*, m_c^*, r^*). This contract provides the efficient level of managed care. Both types pay the same premium r^*, and thus the equity goal is achieved as well. Low-risk types subsidize the high-risk types. It is easy to confirm that a plan attracting a random distribution of the population will break even with this contract.[9]

B. *Risk Adjusters*

Suppose that the regulator and the plans get some signal about each consumer's type. Some proposed risk adjusters, such as evidence of health care utilization in prior periods, could be controlled by an individual or a plan and, therefore, subject to moral hazard. We disregard this feature of some signals, and assume the signal is entirely exogenous. The signal s can take a value of 0 or 1. The signal contains information in the sense that a type H person is more likely than a type L person to get a signal of value 1. From time to time we refer to 0 as the "good" signal, and 1 as the "bad" signal. In general, the signal is imperfect. Some L types get a signal 1, and some H types get a signal 0. Let q_i, $i = H, L$, be the probability that consumer of type i gets a signal with a value of 1. We assume $1 \geq q_H > q_L \geq 0$. With a perfect signal, $q_H = 1$ and $q_L = 0$. The probability of getting signal of value 0 is just $1 - q_i$ for each type.

Let λ_s be the posterior probability the consumer is of type H given the signal s, $s = 0, 1$. Since the signal is informative, using Bayes' rule one can show that: $1 \geq \lambda_1 > \lambda > \lambda_0 \geq 0$. Thus, if a consumer got the signal 1, that person is more likely be of type H than if he/she got the signal 0. When the signal is fully informative, we have $\lambda_1 = 1$ and $\lambda_0 = 0$. Let

$$(4) \quad \mathbf{P_s} = P_H\lambda_s + P_L(1-\lambda_s), \quad \text{for } s = 0, 1$$

and

$$(5) \quad \mathbf{r_s} = m_a^* + \mathbf{P_s} m_c^*, \qquad \text{for } s = 0, 1.$$

$\mathbf{P_s}$ is the probability that a person with signal s will contract illness c and $\mathbf{r_s}$ is the expected health care costs of such a person at the efficient levels of care. Clearly, $\mathbf{P_1} > \mathbf{P_0}$ and $\mathbf{r_1} > \mathbf{r_0}$. Hereafter, variables representing probabilities and premia will be bold typeface when they refer to values conditioned on signals, and will be regular typeface when they refer to true types. $\mathbf{r_s}$ is what we later refer to as the "conventional" risk adjuster. One can readily confirm that if plans are paid in this way, and consumers are randomly distributed across plans, plans break even providing the optimal levels of care. This "pooling," however, does not constitute an equilibrium.

C. *Equilibrium*

We assume that each plan can offer only one contract, and that each consumer chooses only one contract. Our definition of competitive equilibrium is similar to that of Michael Rothschild and Stiglitz (1976). A *competitive*

[9] As we later show, all of our results regarding optimal risk adjustment will hold if the regulator only wished to implement the efficient levels of care (m_a^* and m_c^*), ignoring redistribution.

equilibrium is a set of contracts such that, when consumers choose contracts to maximize expected utility, (i) no contract in the equilibrium set makes negative expected profits; and (ii) there is no contract outside the equilibrium set that, if offered, will make a positive profit.[10]

Consumers know their type, whereas the regulator and the plans observe only the signal defined previously about consumers' type. The assumption that consumers "know their type" is not at all restrictive. One need not interpret this assumption literally. A refined interpretation is that consumers have some additional signal about their expected health costs that is not available to the regulator. Consumers may know their family medical history, for example, and a regulator may be unable to use this in premium regulation.

Furthermore, it is not really important what the regulator knows or does not know—what is important is what is used in the risk adjustment system. Many informative signals about health care costs may be "known," but may not be suitable for use in premium regulation for various reasons. One important reason that a signal may not be used is that it is subject to moral hazard. Some empirical research on risk adjusters studies how health care use in past periods predicts health care use in the future. In general, past health care use is a better predictor than sociodemographic factors. The obvious problem with past health care use as a risk adjuster is that a plan or a consumer has incentives to manipulate use to affect the plan's payment. Although past use may not be suitable for a risk adjuster, it is known to the consumer (and plan), and even may be known to the regulator. It is irrelevant to the impact of premium regulation whether the regulator "knows" past use, so long as it is not part of the risk-adjustment system. Thus, our result in this section applies to a wide range of important and relevant situations in which the consumer knows more information about his/her expected health care costs than the regulator uses in a risk-adjustment system.

We start the analysis by examining the market equilibrium when no regulation takes place. We then consider the impact of the "conventional" form of risk adjustment, and finally, solve for the optimal risk-adjustment policy.

Since signals are observed by plans, in the unregulated case, plans can condition their contract on the consumer's signal. In principle, a plan can offer a contract only to those who got the signal 0 but not 1, or vice versa. One should, therefore, analyze separately the equilibrium contracts in each of the two "markets": the one that serves those who got the signal 0 and the one that serves those who got the signal 1. We show, however, that this separation is artificial since signals are completely ignored by plans in equilibrium and the only thing that affects the contract a consumer obtains in equilibrium is his/her type.[11]

Let us assume that there are two markets and consider first the market for those who got the signal 0. Equilibrium in the unregulated case will be similar to the standard Rothschild-Stiglitz one. Consumers of type H will purchase their best (full information) contract, namely the contract (m_a^*, m_c^*, r_H^*), where m_a^* and m_c^* have been defined previously and are the efficient levels of quality for the acute and chronic illness, respectively, and $r_H^* = m_a^* + P_H m_c^*$. The contract for the L types is their best separating contract—the zero-profit contract that maximizes the utility of the L types subject to the contract not being preferred by the H types. Solving this maximization problem (a complete proof is given in the Appendix), one can show that type L consumers purchase the contract (m_a^*, m_c', r_L'), where $m_c' < m_c^*$ and $r_L' = m_a^* + P_L m_c'$.

[10] It is well known that in this type of model, if the proportion of the L types is large enough, a competitive equilibrium may not exist. We shall study the competitive equilibrium under several informational assumptions. In some cases an equilibrium may not exist, whereas in some others an equilibrium does exist but it is not the socially desired one. However, under the risk-adjustment policies we suggest, an equilibrium always exists, and it is always the socially desired one.

[11] It should be mentioned, however, that the conditions for a separating equilibrium to exist are different in the presence of signals. Without signals, a separating equilibrium will exist if the proportion of the type L consumers, in the entire population, is not too large, since, otherwise, a plan may enter offering a contract that attracts both types and makes a positive profit. In the presence of signals, the condition for a separating equilibrium to exist is that *within each signal group* the proportion of the type L consumers is not too large.

Note that the L types get a contract with the efficient level of acute care, but less than the efficient level of chronic care, and pay a premium lower than that paid by the high-cost types.[12] Even though the premium is lower, the H types do not purchase the contract that L types purchase, since the reduction in premium is not enough to compensate them for the reduction in the coverage for illness c.

One can see that the fact that consumers got signal 0 is completely ignored in equilibrium. In a similar way it can be shown that in the market for those who got signal 1, equilibrium contracts depend only on the consumer's type and are independent of the signal. Furthermore, the contracts are precisely the same contracts as those obtained in the 0 signal market. The reason that signals are ignored in equilibrium is that what really affects plans' profit and consumers' utility are consumers' types. The signal is nothing but a piece of information about the type. Since consumers know their type and since plans know that consumers choose contracts according to their type, signals are ignored.

This equilibrium is inefficient because the low-cost types do not get the efficient level of care, and unfair as we have defined it because the high-cost types are not subsidized by the low-cost types.

D. *Regulation*

Our next purpose is to see whether regulation can improve on the unregulated equilibrium and implement the socially desired outcome. As already noted, we assume that managed care cannot be regulated directly.[13]

Suppose, first, that the regulator does not use risk adjustment and simply sets the premium at r^*. The regulator stipulates that plans must accept all applicants. Thus, each plan chooses a combination (m_a, m_c), consumers choose plans, and each consumer pays r^* to the plan. In this case, if λ is sufficiently small, an equilibrium will not exist; however, if λ is sufficiently large, equilibrium will exist, and it will be characterized by two contracts: H types purchase the contract

$$(6)\quad (m_{aH}, m_{cH}, r^*) = \text{argmax } v_a(m_a) + P_H v_c(m_c)$$

subject to $m_a + P_H m_c - r^* = 0$,

where r^* is given by (3),

and L types purchase the contract

$$(7)\quad (m_{aL}, m_{cL}, r^*) = \text{argmax } v_a(m_a) + P_L v_c(m_c)$$

subject to $m_a + P_L m_c - r^* = 0$,

where r^* is given by (3), and

$$V_H(m_a, m_c, r^*) = V_H(m_{aH}, m_{cH}, r^*).$$

The equilibrium is described in Figure 1. The lines r^{*i}, $i = H, L$, represent all combinations of (m_a, m_c) that will break even if a plan attracts only consumers of type i, each paying premium r^*. The points i, $i = H, L$, depict the equi librium levels of care of a person of type i. The point D depicts the socially desired levels of care. We can see, therefore, that setting the premium at r^* does not implement the socially desired outcome.

Since setting the premium at r^* does not implement the socially desired outcome, the question arises as to whether risk adjustment can help. Risk adjusters will depend on the

[12] All the "distortion" in the contract for the low-cost types is in the level of managed care for the chronic illness here because of our assumption of separability of the benefits of the two illnesses and the absence of risk aversion. In general, the level of managed care would be distorted for both illnesses for the low-cost types. Some further analysis and discussion regarding these results are provided in the Appendix.

[13] Our approach in this respect is the same as James Baumgardner's (1991) model of a managed care organization, and similar to the many papers concerned with the regulation of health care that assume a hospital or a health plan offers a "quality" of care, which is prohibitively costly to regulate directly [e.g., Ma (1994) or William P. Rogerson (1994)]. This is obviously an extreme assumption; some elements of quality can be regulated. However, as long as some dimensions of quality cannot be regulated, our results will be of interest.

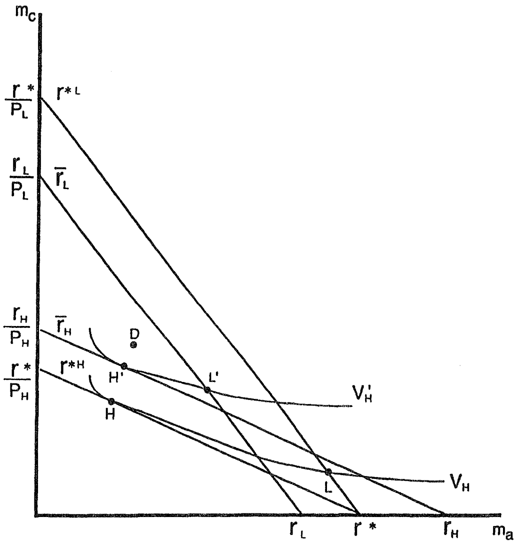

FIGURE 1. EQUILIBRIUM WITH CONVENTIONAL RISK ADJUSTMENT

signal since this is the only additional information that the regulator can use about consumers' type. The natural risk-adjustment policy to consider first is conventional risk adjustment. We consider conventional risk adjustment in the case in which everyone pays r^* and the regulator enforces an open enrollment policy. To be explicit, the model of conventional risk adjustment we analyze can be described in the following four stages:

Stage 1. All consumers pay the regulator r^*.
Stage 2. Plans choose levels of care (m_a, m_c).
Stage 3. Consumers choose plans; plans must accept every applicant.
Stage 4. The regulator pays the plans; for each applicant who received signal s, $s = 0, 1$, the plan receives the conventionally risk-adjusted premium $\mathbf{r_s}$, where $\mathbf{r_s}$ is given by (5).

Under the preceding policy, equilibrium consists of two contracts. H-type consumers, like all consumers pay r^*, and get m'_{aH} and m'_{cH} defined as follows:

(8) $$(m'_{aH}, m'_{cH}) = \text{argmax } v_a(m_a) + P_H v_c(m_c),$$

subject to $m_a + P_H m_c - r_H = 0$,

where r_H is given by[14]

$$q_H \mathbf{r_1} + (1 - q_H)\mathbf{r_0} = r_H.$$

L-type consumers get m'_{aL} and m'_{cL}, where

(9) $$(m'_{aL}, m'_{cL}) = \text{argmax } v_a(m_a) + P_L v_c(m_c),$$

subject to $m_a + P_L m_c - r_L = 0$,

where r_L is given by

$$q_L \mathbf{r_1} + (1 - q_L)\mathbf{r_0} = r_L, \qquad \text{and}$$

$$v_a(m_a) + P_H v_c(m_c) = v_a(m'_{aH}) + P_H v_c(m'_{cH}).$$

In Figure 1, the zero-profit line for a plan attracting only the high-cost types will shift outward relative to the no-risk adjustment case, and to compensate, the zero-profit line for a plan attracting only the low-cost types will shift inward by the risk adjustment. These lines are labeled $\bar{r}_H$ and $\bar{r}_L$, respectively, in Figure 1, with the "bars" over the premiums for the high- and low-cost types, indicating that these are the average premiums plans would receive in the presence of conventional risk adjustment. The new equilibrium will be a separating equilibrium, given the zero-profit lines altered by risk adjustment.

Because all consumers pay r^*, they will simply choose the plan that maximizes their expected benefits. All H-type consumers will choose the same plan in equilibrium, the plan

[14] Recall that q_i is the probability that a person of type i, $i = L, H$ gets a signal of value 1. A health plan getting consumers of type H only would have a share q_H of its enrollees with a signal of 1 and $(1 - q_H)$ with signal zero. With conventional risk adjustment paying premiums $\mathbf{r_0}$ and $\mathbf{r_1}$ for consumers getting the signals 0 and 1, respectively, the plan attracting H types would receive only an average payment of r_H, where this average premium is defined in the text. A similar formula applies to r_L. Since $q_H > q_L$, $r_H > r_L$.

that maximizes their expected utility given the average per-person payment to the plan r_H, which is greater than r^* because of risk adjustment and because H types are more likely to get the signal 1. A plan that attracts the H types will offer a combination of managed care levels that maximizes their expected utility given that it breaks even with an expected payment of r_H. A plan that attracts the L types will offer the best combination to those consumers, given that it breaks even at a premium of r_L, and given it does not attract the H types. Solving (8) one can see that $m'_{aH} < m^*_a$ and $m'_{cH} < m^*_c$. Thus, under conventional risk adjustment, H types get *less than* the efficient levels of care. Solving (9), one can see that $m'_{aL} > m^*_a$ and $m'_{cL} < m^*_c$. Thus, the L types also do not get the efficient levels of care.

We can see, therefore, that the forces that break the efficient pooling equilibrium when premiums are not risk adjusted will also break the efficient pooling equilibrium when premiums are conventionally risk adjusted. Since conventional risk adjustment does not pay expected cost given type, plans will try to attract the low-cost types within each signal group.

In Figure 1, the points H' and L' depict the levels of care of types H and L, respectively, in equilibrium under the conventional risk adjustment, and V'_H is the indifference curve of a type H consumer in equilibrium. Although conventional risk adjustment can improve the efficiency of the equilibrium in insurance markets with managed care, it is not the best the regulator can do with the information available.

E. *Optimal Risk Adjustment*

As just shown, conventional risk adjustment redistributes some, but not enough, resources from the low-cost to the high-cost types. In Figure 1, this redistribution appears as a shift in the zero-profit lines relative to the zero-profit lines in the no-risk adjustment case. As we now go on to show, the regulator may shift the zero-profit lines even further than is implied by conventional risk adjustment, by "overpaying" for a consumer who got the signal 1, compensated by "underpaying" for consumers who got the signal 0, and by so doing, bring the market closer to the socially desired outcome. "Overpaying" and "underpaying" are in comparison to the conventional risk-adjustment premiums. In fact, an optimal risk adjustment can be constructed so as to implement precisely the socially desired bundle of services.

Let

$$r^*_H = m^*_a + P_H m^*_c \tag{10}$$

$$r^*_L = m^*_a + P_L m^*_c\,, \tag{11}$$

where r^*_H and r^*_L are the expected costs of the efficient levels of managed care provided to consumers of types H and L, respectively. Let $\mathbf{r}^*_0$ and $\mathbf{r}^*_1$ be the solution of the following two equations:

$$q_H\mathbf{r}^*_1 + (1 - q_H)\mathbf{r}^*_0 = r^*_H, \tag{12}$$

$$q_L\mathbf{r}^*_1 + (1 - q_L)\mathbf{r}^*_0 = r^*_L. \tag{13}$$

These are two linear equations with two unknowns, $\mathbf{r}^*_1$ and $\mathbf{r}^*_0$. So long as $q_H > q_L$, which will be the case so long as the signal is at all informative, there will be a solution for $\mathbf{r}^*_1$ and $\mathbf{r}^*_0$. The premiums $\mathbf{r}^*_s$, $s = 0, 1$, are the optimal risk-adjustment premiums, as is proven in the following proposition.

PROPOSITION 1: *The regulator can achieve the socially desired outcome by the following policy.*[15]

Stage 1. All consumers pay the regulator r^*.
Stage 2. Plans choose levels of care (m_a, m_c).
Stage 3. Consumers choose plans; plans must accept every applicant.
Stage 4. The regulator pays the plans; for each applicant who received signal s, $s = 0, 1$, the plan receives the risk-adjusted premium $\mathbf{r}^*_s$ defined in (12) and (13).

PROOF:

We show that under the preceding policy, the market equilibrium is such that all plans offer the socially desired bundle of services.

[15] If the regulator just wishes to implement the efficient levels of care, and disregards equity objectives, stage 1 of Proposition 1 could be eliminated, and consumers would pay r^*_s directly to the plans.

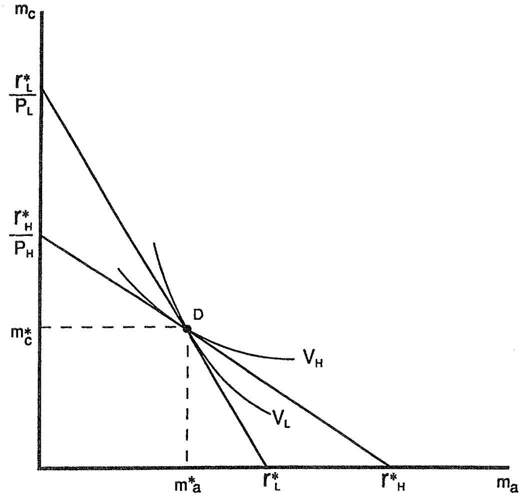

FIGURE 2. EQUILIBRIUM WITH OPTIMAL RISK ADJUSTMENT

Suppose all plans offer the socially desired bundle of services (m^*_a, m^*_c) and consumers randomly choose a plan. Each plan breaks even since its expected cost per person equals the expected premium it receives per person, r^*. (A formal proof of this claim is contained in the next section.) Is there another bundle of services a plan can offer and make a positive profit? There are three cases to consider:

1. A plan can introduce another bundle that will attract both types and will make a positive profit. This, however, is impossible by the efficiency of (m^*_a, m^*_c, r^*).
2. A plan can introduce another bundle of services that will attract only the H-type consumers and will make a positive profit. By (12), if a plan attracts only the H-type consumers, the premium the plan expects to receive for each consumer is r^*_H, since with probability q_H each consumer gets the signal 1 and with probability $(1 - q_H)$ he/she gets the signal 0. However, since the bundle (m^*_a, m^*_c) is the solution to the following problem:

$$(14) \qquad \text{Max } V_H(m_a, m_c)$$

$$\text{subject to } m_a + P_H m_c - r^*_H = 0,$$

a profitable deviation does not exist. [In Figure 2, the line connecting r^*_H on the horizontal axis and r^*_H/P_H on the vertical axis represents the zero-profit line of a plan that attracts only the H-type consumers and the point D is the solution to the maximization problem (14).]
3. A plan can offer a contract that will attract only the L-type consumers and will make a positive profit. If a plan attracts only the L types, the premium it expects to receive for each consumer is r^*_L, since with probability q_L each consumer gets the signal 1 and with probability $(1 - q_L)$ he/she gets the signal 0. By (13), the expected premium for each type-L consumer is r^*_L. However, since (m^*_a, m^*_c) is the solution to the following problem:

$$(15) \qquad \text{Max } V_L(m_a, m_c)$$

$$\text{subject to } m_a + P_L m_c - r^*_L = 0, \qquad \text{and}$$

$$V_H(m_a, m_c, r^*) = V_H(m^*_a, m^*_c, r^*),$$

a profitable deviation that attracts only the L-type consumers also does not exist. [In Figure 2 the line connecting r^*_L and r^*_L/P_L represents the zero-profit line of a plan that attracts only the L types and the solution to the maximization problem (15) is at the point D.]

The essence of Proposition 1 is that it is possible to compensate for a weak signal by a tax/subsidy scheme. If signal perfectly captured type, $q_H = 1$ and $q_L = 0$, and equations (12) and (13) reduce to $\mathbf{r}^*_1 = r^*_H$ and $\mathbf{r}^*_0 = r^*_L$. In general, however, we can solve (12) and (13) for the optimal risk adjustment as

$$(16) \qquad \mathbf{r}^*_0 = [q_H r^*_L - q_L r^*_H]/[q_H - q_L]$$

and

$$(17) \qquad \mathbf{r}^*_1 = [(1 - q_L) r^*_H - (1 - q_H) r^*_L] \div [q_H - q_L].$$

As the signal weakens (q_H gets closer to q_L), $\mathbf{r}^*_1$ gets larger and $\mathbf{r}^*_0$ gets smaller.[16]

[16] Indeed, if q_H is too close to q_L, (16) can call for $\mathbf{r}^*_0$ to be negative. If plans must accept all applicants, there is nothing to stop the regulator within the context of our model

Intuitively, this result can be understood as follows: If the signal is not very precise, the difference in premiums conventional risk adjustment pays for a consumer who got the signal 1 and a consumer who got the signal 0 will be small. Furthermore, the proportion of consumers who got the good signal 0 among the type-L consumers is not much larger than the proportion of consumers who got this signal in the entire population. Thus, by offering a contract that attracts only the L-type consumers, a plan can reduce its costs by a significant amount—as the L-type consumers are much cheaper—without sacrificing much on the premium it is paid per enrollee. If, on the other hand, the premium for a consumer who got the good signal is much smaller than the premium for a consumer who got the bad signal (as our optimal risk-adjustment policy dictates), the plan is severely punished for attracting only the L types, enough so as to deter the plan from taking such an action.

An important assumption in our equilibrium is that when all plans offer the same contract, consumers randomly choose a plan. This is a standard assumption under a "pooling" equilibrium but it deserves some discussion. Obviously, if consumers' preferences over plans do not depend only on the contracts the plans offer but also on some other plans' characteristics (e.g., distance), the equilibrium may not survive, as some plans may not break even. An analysis of this case is discussed later in Section III. However, it is important to observe that consumers having different preferences over plans alone is not sufficient to break our equilibrium; plans must know these preferences, since, otherwise, all plans break even *ex ante* and our equilibrium survives.

One should also notice that the assumption that consumers are indifferent among plans (with the same contract) is common to most adverse-selection models that aim to abstract from issues other than asymmetric information problems. In fact, we view our assumption here as much weaker than the assumption needed for a separating equilibrium to exist under the conventional risk adjustment policy. There, it is absolutely necessary that H types and L types do not mix, even though, from the H type's point of view, all plans offer them contracts with the same expected utility.

Finally, one should observe that there exists another equilibrium under our optimal risk-adjustment policy, one that also implements the socially desired bundle of services. Under this equilibrium, all plans offer the efficient bundle (m_a^*, m_c^*), and consumers fully separate according to type: all H-type consumers go to one (or several) plans and all L types go to another (or several other). We do not view this equilibrium as interesting by itself, but rather as a demonstration of one of the main contributions of our paper. Note that this separating equilibrium is, in fact, the limit of the separating equilibria that emerge under the risk-adjustment overpayment/underpayment policy. Starting with no risk adjustment at all, a separating equilibrium exists, which is far away from the socially desired one. Conventional risk adjustment results in another separating equilibrium, which is closer to the socially desired one. (Refer to Figure 1.) By increasing the payment for the bad signal, 1, and decreasing payment for the good signal, 0, even further, the regulator can induce a separating equilibrium that is even closer to the socially desired one, and the limit of all these equilibria is the separating equilibrium with the optimal contracts.

In practical terms, this is the main point of our paper. Conventional risk adjustment moves us in the right direction, but not far enough. The empirical research that underlies conventional risk adjustment tells the regulator the good and bad signals, and gives the regulator a benchmark, the *minimum* adjustments warranted by the data. The economic reasoning laid out in this paper adds this: magnifying the adjustments implied by conventional risk adjustment further improves matters, by sharpening plans' incentives to provide the efficient level of services.

A problem with our optimal risk-adjustment mechanism is that if risk adjustment does not pay expected cost on the basis of a signal, plans have an incentive to select consumers on the basis of the signal itself. This was not possible in our analysis so far as we assumed that the regulator could enforce an open enrollment policy and a plan could compete for persons only on the basis of their type, not on the basis of their signal. However, in general, if overpay-

from requiring a negative payment for consumers with the good signal, though it would obviously cause protest in a real-world context!

ment on some signal were in place, one could assume that plans would try to provide services attractive to the "old," say, to attract these people. We do not address this problem here, because our main objective was not so much to solve for the optimal risk adjustment, but rather to show that some overpayment is desirable. We have carried out some preliminary analysis of the case where plans can also attract consumers on the basis of signals, and not only on the basis of types. One can show that this boils down to a situation where there are more types than signals, where a risk-adjustment policy that implements the first best may not exist. Conventional risk adjustment, however, will not generally be second best. Competition on the basis of the signal can be viewed as imposing a cost on the overpayment policy, but this cost is very small as "over" payment begins to be made, implying that at least some overpayment part of the second-best policy.

III. Heterogeneous Consumers

In this section we dispense with the assumption that, if all plans offer the same contract, consumers are indifferent among plans. As Cutler and Zeckhauser (1998) point out, one function of a market for health plans is to accommodate consumers' tastes for different plans. In general, we can expect that consumers' choices are a function not only of the contracts plans offer, but also of some other (plans' or consumers') characteristics (e.g., distance).

In what follows we do not provide a general theory of equilibrium under consumer heterogeneity. Instead, we use results regarding optimal risk adjustment from Section II to address a much narrower question: Suppose an equilibrium exists, and in this equilibrium all plans offer the same bundle of services, but consumers do not necessarily randomize across plans. In other words, the proportion of consumers with the signal 0 (1) is not the same in all plans. How should premiums be risk adjusted so that all plans break even?[17] We show that the risk adjustment required for the plans to break even depends on the distribution of consumers (defined by types and signals) across plans. We show that except for some special cases, conventional risk adjustment will not be optimal (i.e., plans will not break even) and overpayment on one signal, compensated by underpayment on another, is generally necessary.

Suppose there is some equilibrium with two plans A and B.[18] Let a denote the number of consumers that choose plan A and let b denote the number of consumers in plan B in equilibrium. Let a_s denote the number of consumers that choose plan A and got signal s, a_i the number of consumers of type i that choose plan A, and a_{is} denote the number of type i consumers who got signal s choosing plan A in equilibrium, where $s = 0, 1$, and $i = H, L$. Define similar notation for plan B.

Let r_i denote the expected cost of type i consumers in equilibrium. We assume that all consumers in all plans receive the same levels of care in equilibrium, and hence the expected cost of a type i consumer is independent of which plan he/she purchases. Given this equilibrium, our purpose is to find the values of r_0 and r_1 so that the plans break even.

For plan A to break even, the following condition must hold:

$$a_0 r_0 + a_1 r_1 = a_H r_H + a_L r_L. \quad (18)$$

The left-hand side of (18) is the total premiums the plan receives, and the right-hand side is total expected cost. Since

$$a_s = a_{Hs} + a_{Ls} \qquad \text{for } s = 0, 1, \quad (19)$$

and

$$a_i = a_{i0} + a_{i1} \qquad \text{for } i = H, L, \quad (20)$$

we can plug (19) and (20) into (18) and get

$$a_{H0}(r_0 - r_H) + a_{L0}(r_0 - r_L) + a_{H1}(r_1 - r_H) + a_{L1}(r_1 - r_L) = 0, \quad (21)$$

[17] For an equilibrium to be sustained, two conditions must hold: (1) all plans break even and (2) there is no profitable deviation. In this section, we study the risk adjustment implied by the first necessary condition alone.

[18] The basic insight obtained by the analysis below will be true also in the more general case when there are more than two plans.

and similarly for plan B we get

$$(22)\quad b_{H0}(r_0 - r_H) + b_{L0}(r_0 - r_L) + b_{H1}(r_1 - r_H) + b_{L1}(r_1 - r_L) = 0.$$

Equations (21) and (22) are the generalizations of equations (16) and (17) from the previous section for any distribution of types and signals across plans. The two unknowns are the risk-adjusted premiums r_0 and r_1 and the given parameters are the distribution of types across plans, $\{a_{is}, b_{is}\}$, $i = H, L$, $s \in 0, 1$, and the expected cost of the contracts for each type, r_H and r_L.

A special case to mention is when consumers are homogeneous. In this case, consumers randomize between plans, and each plan gets the same share of consumers of each type–signal combination. Thus, $a_{ij} = kb_{ij}$, where $i = H, L$ and $j = 0, 1$, and $k > 0$. Hence, equations (21) and (22) are the same equation, and become just one equation with two unknowns. This one equation has many solutions, including, as is easy to show, the conventional and optimal risk adjustment discussed earlier in Proposition 1.[19]

A simple manipulation of (21) will give us the following condition:

$$(21')\quad a_0\left(\frac{a_{H0}}{a_0} r_H + \frac{a_{L0}}{a_0} r_L - r_0\right) + a_1\left(\frac{a_{H1}}{a_1} r_H + \frac{a_{L1}}{a_1} r_L - r_1\right) = 0.$$

From (21′) we can report the following:

PROPOSITION 2: *If* $(a_{H0}/a_0) = \lambda_0$, *and* $(a_{H1}/a_1) = \lambda_1$, *then* r_0 *and* r_1 *that satisfy (21′) are the conventional risk-adjusted premiums.*

The proof of Proposition 2 (and of Proposition 3 below) follows directly from equation (21′). Proposition 2 says that if in a plan, the distribution of types given a signal is the same as the distribution of types given the signal in the entire population, then conventional risk adjustment will cover the plan's costs.

Two cases that satisfy Proposition 2 are:

(i) complete pooling; and
(ii) one plan gets all consumers that have the signal 0, and the other plan gets all consumers that have the signal 1.

One can see, however, that if the condition in Proposition 2 about the distribution of types within plans does not hold, then r_0 and r_1 will not generally be the conventional risk-adjusted premiums. We now discuss this case.

One special case where the preceding condition does not hold is when $a_L = 0$ and $b_H = 0$. This is the separating equilibrium, and equations (21) and (22) reduce to the equations (16) and (17) from the previous section. In equations (16) and (17), optimal risk adjustment "overpaid" for the "bad" signal, 1, and "underpaid" for the "good" signal, 0. The question we address now is whether this result is special to the case studied in Section II, or whether it is more general. In what follows, we demonstrate that overpayment for the bad signal and underpayment for the good signal is a general property of optimal risk adjustment.

Denote $(a_H/a) \equiv \lambda_a$ and $(b_H/b) \equiv \lambda_b$.

PROPOSITION 3: *If* $\lambda_i > \lambda_j$, $i, j, = a, b$, $i \neq j$, *and*, $a_{H1} = q_H a_H$, $a_{L1} = q_L a_L$, $b_{H1} = q_H b_H$, *and* $b_{L1} = q_L b_L$; *then in equilibrium* r_0 *and* r_1 *must satisfy:*

$$(23)\quad r_1 - r_0 = \frac{r_H - r_L}{q_H - q_L}.$$

Proposition 3 says that if the fraction of type *H* people in one plan is larger than the fraction of type *H* people in the other plan ($\lambda_i > \lambda_j$) and, given a type, the distribution of signals in the plan is the same as the distribution of signals given a type in the entire population, then condition (23) holds. A special case of the previous results is when $\lambda_i = 1$ and $\lambda_j = 0$, $i, j = a, b$, i.e., equilibrium is fully separating. In this case, r_0 and r_1 will satisfy (16) and (17).

Since r_H and r_L are given, condition (23) says that the smaller the difference $q_H - q_L$ (i.e., the less informative the signal), the larger the difference $r_1 - r_0$ (i.e., the greater the difference

[19] We are grateful to a referee for insight on this point.

in payment between the bad signal and the good signal). We can therefore conclude that "overpayment" for the bad signal is general to the form of equilibrium.

IV. Discussion

Risk adjustment of health insurance premiums is part of virtually all market-based proposals to reform health care markets, in the United States and internationally. Weak correlation between observable characteristics, such as gender and age, and health costs at the individual level limit how effectively conventional risk adjustment—paying a plan in proportion to a person's expected costs conditional on the observable signals—can remedy adverse-selection-created incentives. With this paper we hope to suggest that risk adjustment is a more powerful tool than previously thought. A regulator need not confine itself to paying in proportion to expected costs. Indeed, as we have shown, conventional risk adjustment proportional to expected costs is not the best policy of risk adjustment. By viewing risk adjustment as a tax/subsidy scheme based on signals, risk adjusters can improve the allocation over that which can be achieved with conventional risk-adjustment policy. The weakness of empirical signals about health costs can be compensated for by overpaying on the basis of a bad signal, and underpaying for a good signal.

Our paper contains a model of health insurance with managed care within which we formally solve for the optimal risk-adjustment policy. Although we regard this theoretical characterization of the optimal policy to be a contribution to the analysis of risk adjustment, we want to stress here the practical importance of the idea of overpayment in relation to conventional adjusters. Starting with conventional risk adjustment, helpful but hardly dispositive, public and private regulators can take the information contained in conventional risk adjusters and improve the efficiency of health insurance markets by magnifying the weights. We want to be clear that we are not claiming that our proposed method for risk adjusting is a substitute for conventional risk adjustment. Indeed, to be implemented, our ideas rely on the existence of reliable empirical associations between signals and cost. Better conventional risk adjustment improves matters on its own (see Section II) and, furthermore, increases the confidence we would have in the overweighting of the signals correlated with higher costs implied by the theory of optimal risk adjustment developed here.

The most direct application of the ideas developed here would be to some state Medicaid programs for the poor in the United States, to countries such as Israel, or many employers, which require employees to choose a managed care plan paid by risk-adjusted capitation payments. For practical reasons, the present and likely future risk adjusters available to regulators in these systems are simple and few (Israel, for example, uses only "age"). We contend that increasing the weight on the bad signal from that derived from conventional risk adjustment would improve the efficiency of the health care system by redistributing health care spending away from services the healthy are likely to use and toward services likely to be used by the less-healthy members of the pool.

Medicare, the federal program for the elderly, has been beset by selection-related problems since the government opened the choice of managed care plans to the elderly about ten years ago. Medicare continues to experiment with the formula with which to pay plans, currently beginning a downweighting of geographical adjusters, and introducing a new more powerful signal, categories of prior health care use. (See Medicare Payment Advisory Commission, 1998.) Medicare's main complaint about its risk adjustment/managed care policy is that since the healthy elderly tend to join the managed care plans, and conventional adjusters pay only the mildly adjusted average, Medicare pays too much to managed care plans. Though not designed to address Medicare's expenditure minimization perspective, our proposed policy of over- and underpaying would help Medicare, too: the healthy will come disproportionately from those with the good signal (the young elderly), and underpaying for this group will at least partially alleviate Medicare's budget woes.

Stepping back to the more general literature on adverse selection, a tax/subsidy based on a signal is new for this literature, and may have promise for applications outside of health care. The "privatization" of public education and other social services shares key features with the market for managed health care. Fairness and efficiency of schooling would seem to matter in much the same manner as these criteria do in health care, motivating governments to subsidize consumption and to equalize ability to pay for education through voucherlike mechanisms. Students' true differences in "cost" (viewed in terms of either management costs or costs to educate to a certain standard) will be imperfectly correlated with signals like ability scores. In a recent paper, Dennis Epple and Richard Romano (1998) characterize the equilibrium of private and public schools when government sets a (non-risk-adjusted) voucher paid on behalf of heterogeneous students. Lower ability/income students get lower-quality education. A student-based tax/subsidy scheme, amounting to "risk-adjusted vouchers," seems, on the basis of our analysis, to be a promising approach to improving the efficiency and fairness of the market outcome in that case, too.

Appendix

In this Appendix we solve for the equilibrium contracts in the unregulated case and show that only the c service is distorted. We also provide an intuition for this result. In the unregulated case, the equilibrium contract for the L types solves the following problem:

$$\text{(A1)} \qquad \text{Max } v_a(m_a) + P_L v_c(m_c) - r$$

$$\text{subject to } m_a + P_L m_c - r = 0$$

$$\text{(A2)} \qquad v_a(m_a) + P_H v_c(m_c) - r = v_a(m_a^*) + P_H v_c(m_c^*) - r_H^*.$$

The first-order conditions are:

$$\text{(A3)} \qquad v_a'(m_a) - \alpha - \delta v_a'(m_a) = 0$$

$$\text{(A4)} \qquad P_L v_c'(m_c) - \alpha P_L - \delta P_H v_c'(m_c) = 0$$

$$\text{(A5)} \qquad -1 + \alpha + \delta = 0 .$$

Substituting (A5) into (A3) we get:

$$\text{(A6)} \qquad \alpha(v_a'(m_a) - 1) = 0 .$$

If $\alpha = 0$, (A4) and (A5) imply $P_L = P_H$, which is a contradiction. Thus, by (A6) it must be that $v_a'(m_a) = 1$ and, hence, $m_a = m_a^*$. Substituting $m_a = m_a^*$, $r^* = m_a^* + P_H m_c^*$ and (A1) into (A2) we obtain

$$\text{(A7)} \qquad P_H v_c(m_c) - P_L m_c = P_H v_c(m_c^*) - P_H m_c^* .$$

Since $P_L < P_H$, (A7) implies

$$\text{(A8)} \qquad m_c^* - m_c < v_c(m_c^*) - v_c(m_c) .$$

Since $v_c'' < 0$ and $v_c'(m_c^*) = 1$ we obtain $m_c < m_c^*$.

The intuition for this result is as follows: when consumers are risk neutral, the contract (m_a, m_c, r) can, in fact, be viewed as an integration of two contracts, (m_a, r_a) and (m_c, r_c), where the first one covers acute care and the second one covers chronic care, and $r_a + r_c = r$. Since there is no asymmetry of information with respect to acute care, its contract must be the efficient one for both types, namely $m_a = m_a^*$. As for chronic care, since there is asymmetric information regarding the consumer's type, equilibrium here will take the standard Rothschild-Stiglitz form: type H consumers get the efficient contract and pay premium accordingly and the type L consumers are separated by a contract that provides less chronic care and requires a lower premium. One should note, however, that risk neutrality is what enables us to treat the contract (m_a, m_c, r) as an integration of two contracts for the two services. In the presence of risk aversion, the marginal utility depends on the premium and, hence, the premium paid for one treatment, r_c, say, may affect the consumer's utility from treatment of the other illness, m_a.

REFERENCES

Baumgardner, James. "The Interaction Between Forms of Insurance Contract and Types of Technical Change in Medical Care." *Rand Journal of Economics*, Spring 1991, *22*(1), pp. 36–53.

Cutler, David M. "Cutting Costs and Improving Health: Making Reform Work." *Health Affairs*, Spring 1995, *14*(1), pp. 161–72.

Cutler, David M. and Reber, Sarah. "Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection." *Quarterly Journal of Economics*, May 1998, *113*(2), pp. 433–66.

Cutler, David M. and Zeckhauser, Richard. "Adverse Selection in Health Insurance," in Alan M. Garber, ed., *Frontiers in health policy research*, Vol. 1, Cambridge, MA: MIT Press, 1998, pp. 1–31.

Ellis, Randall P. "Creaming, Skimping and Dumping: Provider Competition on the Intensive and Extensive Margins." *Journal of Health Economics*, October 1998, *17*(5), pp. 537–56.

Ellis, Randall P.; Pope, Gregory C.; Iezzoni, Lisa I; Ayanian, John Z.; Bates, David W.; Burstein, Helen and Ash, Arlene S. "Diagnosis-Based Risk Adjustment for Medicare Capitation Payments." *Health Care Financing Review*, Spring 1996, *17*(3), pp. 101–28.

Encinosa, William E. and Sappington, David E. M. "Competition Among Health Maintenance Organizations." *Journal of Economics and Management Strategy*, Spring 1997, *6*(1), pp. 129–50.

Enthoven, Alain C. and Singer, Sara J. "Market-Based Reform: What to Regulate and by Whom." *Health Affairs*, Spring 1995, *14*(1), pp. 105–19.

Epple, Dennis and Romano, Richard E. "Competition Between Private and Public Schools, Vouchers, and Peer-Group Effects." *American Economic Review*, March 1998, *88*(1), pp. 33–62.

Gauthier, Anne K.; Lamphere, Jo and Barrand, Nancy L. "Risk Selection in the Health Care Market: A Workshop Overview." *Inquiry*, Spring 1995, *32*(1), pp. 14–22.

Goldman, Dana; Hosek, Susan D.; Dixon, Lloyd S. and Sloss, Elizabeth M. "The Effects of Benefit Design and Managed Care on Health Care Costs." *Journal of Health Economics*, October 1995, *14*(4), pp. 401–18.

Greenwald, Bruce C. and Stiglitz, Joseph E. "Externalities in Economies with Imperfect Information and Incomplete Markets." *Quarterly Journal of Economics*, May 1986, *101*(2), pp. 229–64.

Holohan, John; Coughlin, Teresa; Ku, Leighton; Lipson, Debra J. and Rajan, Shruti. "Insuring the Poor Through Section 1115 Medicaid Waivers." *Health Affairs*, Spring 1995, *14*(1), pp. 199–216.

Institute of Medicine. *Controlling costs and changing patient care? The role of utilization management.* Washington, DC: National Academy Press, 1989.

Jensen, Gail A.; Morrisey, Michael A.; Gaffney, Shannon and Liston, Derek K. "The New Dominance of Managed Care: Insurance Trends in the 1990's." *Health Affairs*, January/February 1997, *16*(1), pp. 125–36.

Lewis, Tracy and Sappington, David. "Cream Skimming, Cost Sharing, and Subjective Risk Adjustment in the Health Care Industry." Unpublished manuscript, University of Florida, June 1996.

Light, Donald W. "Keeping Competition Fair for Health Insurance: How the Irish Beat Back Risk-Rated Policies." *American Journal of Public Health*, May 1998, *88*(5), pp. 745–48.

Luft, Harold S. "Potential Methods to Reduce Risk Selection and Its Effects," *Inquiry*, Spring 1995, *32*(1), pp. 23–32.

Ma, Ching-to Albert. "Health Care Payment Systems: Cost and Quality Incentives." *Journal of Economics and Management Strategy*, Spring 1994, *3*(1), pp. 93–112.

Medicare Payment Advisory Commission. *Report to Congress: Medicare payment policy.* Washington, DC, 1998.

Netanyahu Commission. *Report of the State Commission of Inquiry into the functioning and efficiency of the health care system.* Jerusalem: The State of Israel, August 1990.

Newhouse, Joseph P. "Patients at Risk: Health Reform and Risk Adjustment." *Health Affairs*, Spring I 1994, *13*(1), pp. 132–46.

———. "Reimbursing Health Plans and Health Providers: Efficiency in Production Versus Selection." *Journal of Economic Literature*, September 1996, *34*(3), pp. 1236–63.

Pauly, Mark V. "Is Cream-skimming a Problem for the Competitive Medical Market?" *Journal of Health Economics*, April 1984, *3*(3), pp. 87–95.

Physician Payment Review Commission. *Risk selection remains problem in Medicare.* PPRC Update No. 21, Washington, DC, July 1997.

Rogerson, William P. "Choice of Treatment Intensities by a Nonprofit Hospital under Prospective Pricing." *Journal of Economics and Management Strategy*, Spring 1994, *3*(1), pp. 7–51.

Rothschild, Michael and Stiglitz, Joseph E. "Equilibrium in Competitive Insurance Markets: An Essay in the Economics of Imperfect Information." *Quarterly Journal of Economics*, November 1976, *90*(4), pp. 629–49.

Swartz, Katherine. "Reducing Risk Selection Requires More Than Risk Adjustments." *Inquiry*, Spring 1995, *32*(1), pp. 6–10.

van de Ven, Wynand P. M. M. and Ellis, Randall P. "Risk Adjustment in Competitive Health Plan Markets," in A. J. Culyer and J. P. Newhouse, eds., *Handbook in health economics*. Amsterdam: North-Holland, 2000 (forthcoming).

van de Ven, Wynand P. M. M. and van Vliet, Rene C. J. A. "How Can We Prevent Cream Skimming in a Competitive Health Insurance Market? The Great Challenge for the 90's." in Peter Zweifel and H. E. Frech, III, eds., *Health economics worldwide*. Dordrecht: Kluwer Academic, 1992, pp. 23–46.

van Vliet, Rene C. J. A. and van de Ven, Wynand P. M. M. "Towards a Capitation Formula for Competing Health Insurers: An Empirical Analysis." *Social Science and Medicine*, 1992, *34*(9), pp. 1035–48.