WORKING
PAPER

Overview of the COMPARE Microsimulation Model

Federico Girosi, Amado Cordova, Christine Eibner, Carole Roan Gresenz, Emmett Keeler, Jeanne Ringel, Jeffrey Sullivan, John Bertko, Melinda Beeuwkes Buntin, Raffaele Vardavas

WR-650

January 2009

RAND HEALTH

COMPARE MODEL OVERVIEW

CHAPTER 1
INTRODUCTION

In 2007, approximately 17 percent of the non-elderly U.S. population was uninsured (DeNavas-Walt, Proctor, and Smith, 2008), and health care spending was expected to increase by nearly 7 percent relative to expenditures in 2006 (Keehan et al., 2008). Concerns about the rates of uninsurance in the United States, coupled with rising health care costs, have made changes in health policy a priority on the U.S. public policy agenda.

In the 2008 presidential race, candidates proposed a variety of policies aimed at expanding health care access and affordability, including individual mandates, requiring all individuals to purchase health insurance; employer mandates, requiring most businesses to offer insurance; changes in the tax treatment of insurance; and safety-net expansions. Yet, understanding the likely effect of these policies on costs, coverage, and population health is challenging in view of the complex effects that such changes may have on individual and employer behaviors, and the limited prior experience with health system changes in the United States.

RAND Health researchers developed the COMPARE microsimulation model as a way of projecting how households and firms would respond to health care policy changes based on economic theory and existing evidence from smaller-scale changes (e.g., changes in Medicaid eligibility). A microsimulation model uses computer software to develop a synthetic U.S. population made up of individuals, families, firms, and the federal and state governments. Individuals, firms, and other *agents* (the general name given to entities that can take actions) in our model make decisions using a customized "rule book," which takes into account such factors as individual and family characteristics, prices, and government regulations. For example, if an offer is available, an individual in our model would make the choice to enroll in employer-sponsored health insurance or not after considering the following:

- Whether he or she was eligible for other options, such as Medicaid
- The cost of employer-sponsored insurance, overall and relative to other options
- Individual characteristics, such as total family income and health
- Whether the government offered an incentive to enroll in insurance, such as a tax credit, or a penalty for non-enrollment.

The individual's decision in the status quo might change after a policy intervention. For example, a person who declines employer-sponsored insurance in the status quo might opt to enroll in an insurance plan if the government introduced an individual mandate with a substantial non-enrollment penalty. Firms in our model also follow a rule book, opting to offer health insurance after considering the value of insurance as a recruitment and retention tool, the expected cost of offering a policy, and any government regulations that might provide an incentive or disincentive to offer insurance.

An advantage of the microsimulation approach is that it allows us to incorporate interactions among agents (firms, households, and the government) in the model. For example, a Medicaid expansion might cause some newly eligible workers to drop employer-sponsored health insurance in favor of public coverage. Employers in our model can respond to this behavior by reassessing the benefit of providing health insurance to workers. If a substantial share of workers becomes newly eligible for Medicaid, then the firm may decide to stop offering insurance. Similarly, an employer mandate that imposes a penalty on non-offering firms may cause some businesses to begin offering health insurance. In response, some workers in these firms might opt to take employer coverage.

The COMPARE microsimulation model is currently designed to address four types of coverage-oriented policy options: individual mandates, employer mandates, expansions of public programs, and tax incentives. The model is flexible and allows us to expand the number and variety of policy options addressed. We plan to add options over time to reflect the current policy discussion.

The first step in using the microsimulation is to compute the status quo--the way things now stand. It is crucial that the status quo configuration provide a realistic picture of the U.S. population at a point in time. For example, insurance premiums predicted by the model must match observed premiums with reasonable accuracy. We replicate health insurance cost and coverage patterns observed in the United States in 2007.

The second step is using the model to simulate a policy option. We simulate policy options by altering the values of appropriate attributes (e.g., health insurance premiums, regulatory requirements, worker preferences) and allowing the agents to respond to these changes and settle into a new equilibrium. We can then compute the outcome of the policy option by comparing the new equilibrium with the status quo. The model not only predicts the effect of various health policy options on spending, coverage, and health outcomes, but it also predicts how specific design features influence the effects of a policy option. For example, depending on the magnitude of the noncompliance penalty and the degree to which small firms are excluded from the mandate, an employer mandate may have a very different effect on health insurance coverage.

Data for developing the model population and predicting household and firm behavior come from nationally representative surveys conducted by government agencies and private foundations. Key data sources used in the model include the Survey of Income and Program

Participation (SIPP), the Medical Expenditure Panel Survey (MEPS), the Kaiser Family Foundation/Health Research and Educational Trust Employer Survey (Kaiser/HRET), and the Survey of U.S. Businesses (SUSB). We also draw from published literature, as well as from documentation published by other modelers--most especially, by Jonathan Gruber of MIT and the Urban Institute.

In the following chapters, we describe the model in greater detail, sometimes focusing on technical aspects of the model that may be of limited interest to a general audience. First, we discuss how we populate the model with firms and individuals and how we compute health insurance premiums in the employer and non-group insurance markets (Chapter 2). Then, we describe how we create the rules that determine the ways in which firms and households will make decisions, both in the status quo and in response to policy changes (Chapter 3). In Chapter 4, we describe the four coverage policy options considered in the initial COMPARE microsimulation modeling, and we discuss assumptions that are specific to these options. Chapter 5 describes how spending changes when previously uninsured individuals gain coverage, and Chapter 6 describes how health changes when individuals acquire insurance.

CHAPTER 2
POPULATING THE MODEL

2.1 Individuals

The individuals in our model come from the 2001 SIPP. The SIPP, a longitudinal study of households conducted by the U.S. Census Bureau, contains data on demographics, household composition, health insurance status, income, assets, and labor force participation. Our data come from a snapshot of the 2001 SIPP taken in spring 2002. The SIPP is only one of the possible choices in terms of population surveys that include information on health insurance status, family composition, and labor force participation. The Congressional Budget Office (CBO)—which has developed a model to analyze policy options related to health insurance coverage--chose the SIPP as well, but alternatives include the Current Population Survey (CPS; used by the Urban Institute and Jonathan Gruber) and the MEPS (used by the Lewin Group). Each of these data sets has its own merits, and no single choice is unequivocally better then the others. We use the 2001 SIPP rather than the more recent 2004 panel because complete data for the 2004 SIPP were released only within the past several months. In updates to this model, we will replace the core sample with data from 2004. For now, we updated the 2001 data to reflect the U.S. population in 2007 by reweighting the data to reflect Census population estimates. Weights are based on race, age, and sex. We have the capability to project population demographics based on age through 2050, although we did not exploit this feature in the current model. Our synthetic population includes individuals of all ages; however, when we present results for the current policy options, we focus on the population under age 65.

We use data from the SIPP to define health insurance eligibility units (HIEUs), which are *groups* of family members that tend to be eligible for health insurance coverage through employer-sponsored insurance (ESI) plans.[1] Specifically, HIEUs in our data consist of adults, their spouses, and dependent children under the age of 18; we also allow for exceptions whereby children older than 18 may be covered through parents.

2.2 Individual Health Care Expenditures

Because the SIPP does not include data on health expenditures, spending estimates in our model are based on data from the Medical Panel Expenditure Survey, Household Component

[1] Unlike the choice to enroll in ESI, the choice to enroll in Medicaid or non-group coverage is made at the individual, rather than the HIEU, level.

(MEPS-HC). Model estimates exclude spending on vision and dental care. To increase sample size, we pooled MEPS data from the years 2002 and 2003. In addition, since MEPS-HC expenditures underestimate national health spending (Sing et al., 2006), we inflated the MEPS-HC spending estimates to match the National Health Expenditure Accounts (NHEAs), using the detailed procedure found in Sing et al. (2006).[2] We also inflated the MEPS-HC spending estimates to 2007 levels, using health care inflation factors based on the NHEA (the factors inflate health spending by approximately 7 percent in each year).

We linked expenditures from the MEPS to individuals in the SIPP using semi-constrained statistical matching. First, we stratified both the MEPS-HC and the SIPP into demographic cells based on age, insurance status, health status, region, and income.[3] We then randomly assigned each person in the SIPP expenditure data using information from a demographically matching individual in the MEPS-HC. We then computed weighted expenditures in each MEPS-HC demographic cell and in each matching SIPP demographic cell. If weighted expenditures differed by more than 0.5 percent, we repeated the procedure until expenditures differed by less than 0.5 percent. We checked that, overall and for both adults and children, this methodology preserved the distribution of health care expenditures in the United States.

2.3 Employers

Employers in our simulation model are based on data from the 2006 Kaiser/Health Research and Educational Trust (Kaiser/HRET) Annual Survey of Employers. The Kaiser/HRET data contain information on firm characteristics and health insurance benefit plans

[2] NHEA estimates are considered to be more accurate than MEPS estimates because NHEA figures are based on administrative reports from health care providers, as opposed to consumer reports of health care spending.
[3] Using additional variables, such as race and education, did not help the matching process significantly, so we did not use them.

for a nationally representative sample of U.S. employers with three or more workers. Benefits data available in the Kaiser/HRET survey include the number and type of plans offered, the cost of these plans, employee contribution and cost-sharing requirements, and covered benefits. We matched firms in the Kaiser/HRET data to workers in the SIPP based on Census region, firm size, industry, and whether or not the firm offers health insurance (in the SIPP data, we know whether or not the worker was offered insurance, regardless of whether the worker accepted it).

Because the estimated size of the labor force differed in the two surveys, and because the sample size of the SIPP was far larger than the Kaiser/HRET data, we assumed that the SIPP workforce count was accurate. We adjusted the weights in the Kaiser/HRET file to match the SIPP. We further adjusted the Kaiser/HRET data to reflect the number of firms enumerated in the Survey of U.S. Businesses (SUSB) (2005), stratified by firm size. Our approach ensures that employer health insurance offer rates among offering firms will reflect the Kaiser/HRET data, that characteristics of workers reflect the worker distribution found in the SIPP, and that the number and type of businesses in the United States reflect the distribution found in the SUSB.

For the current version of the model, this approach does not differ significantly from the approach used by other researchers, such as the Urban Institute, which does not make use of the HRET. Although the HRET offers rich information about plan choices, the details of such plans are not explicitly taken into account in the behavior of the HIEU. Instead, we made the assumption that all HIEUs receiving offers from firms of a given size are offered the same "average" plan and that the actuarial values of the offered plans are a function of firm size, as detailed below. However, by incorporating the Kaiser/HRET data into our model, we have the opportunity to make better use of the more detailed information in the survey in future iterations of COMPARE.

2.4 Employer Health Insurance Premiums

Data for calculating group health insurance premiums were derived from aggregate MEPS expenditures. MEPS data include employer-sponsored health insurance premiums (referred to in the displays of modeling results as *group premiums*) that follow either state rate band limitations (for employers with less than 50 employees) or are based on the employer's own claims experience. Our methodology assigns individuals to a non-specified "average" insurance plan, rather than creating firm- or worker-specific premiums, because the MEPS-HC does not include information about the characteristics of health plans; it reports only the dollar value of premiums. An alternative strategy is to use the restricted-access MEPS-linked Household and Insurance Component (MEPS HC-IC) data, in which one can observe the characteristics of the plans offered to an HIEU, together with the choice the HIEU made. We pursued this strategy initially, but we later abandoned it because the data are only available for restricted-access use on site at the Agency for Healthcare Research and Quality (AHRQ). In order to incorporate the HC-IC data into our model, we would need to have the data in-house in order to link it to the SIPP.

We assume that there are only two types of plans available in the market—single and family plans. *Family plans* combine traditional family plans, two-adult plans, and single-adult-plus-children plans.

2.4.1 Single-Employee Premium Pools

We used 12 pools to compute group premiums for single employees, based on a combination of four regions (North East, South, Mid West, and West) and three firm sizes (<25, 25–99, and 100+). We limited our analysis to no more than 12 pools because sample sizes became thin in certain cells when we attempted to further stratify the data. Since most self-insured firms are large and likely to provide health insurance, they are grouped with firms with 100+ employees for the purposes of the model.

Each of the 12 pool premiums are calculated using the expected value of medical expenditures of people in the pool, multiplied by an actuarial value factor for cost-sharing and an administrative loading factor, defined as $1/_{1-\delta}$ where δ (the administrative loading fee) represents the percentage of the total premium used for administrative expenses and profits. The average actuarial value factor varies by the size of firm (0.78 for firms with <25 employees, 0.85 for firms with 25–99 employees, and 0.88 for firms with 100+ employees). Table 2.1 reports administrative costs. These factors vary by size of the employer and are derived from published research by the Urban Institute.[4]

Table 2.1
Administrative Loading Fees by Firm Size

Size of Firm	Administrative Loading Fees (%)
<25 workers	20.0
25–99 workers	13.0
100+ workers	8.3

NOTE: The administrative loading fee represents the percentage of the total premium used for administrative expenses and profits.

There is large variation in premiums across regions and firms. Although the standard deviation of monthly premiums is approximately $750, the HIEU's elasticity of demand[5] for group insurance with respect to price is generally quite small, around –0.03 (Blumberg et al., 2001). Thus, the model is fairly insensitive to changes in group premiums.

[4] The numbers used by the Urban Institute may appear to be different because they are presented as a percentage of the amount of money paid out by the pool, rather than as the premium.
[5] The *elasticity* represents the percentage change in the probability of taking an insurance policy that results from a 1-percent change in the price of insurance.

2.4.2 Family Premium Pools

The structure of the family premium pools mirrors the 12 pools used for single employee premiums. After much simulation using individual insurance units (e.g., dependent spouses, dependent children), the modeling team found that the procedure tended to underestimate the observed amount of family premiums by about 8 percent, on average. To address this issue, we assumed that family premiums were 2.7 times those of single employee premiums, based on factors observed in the Kaiser/HRET data. Because the elasticity of demand (a measure of the HIUE's sensitivity to price when making the decision to purchase insurance) for group health insurance is so small, any differences in family premiums would make a small difference in the demand for insurance.

2.4.3 Employer/Employee Premium Share

We assumed that the share of premium paid by the employee was 16 percent for single coverage and 26 percent for family coverage. The premium contribution for single coverage comes from the Kaiser/HRET survey, and the premium contribution for family plans is an average based on data from HRET and MEPS. For the reforms that we have simulated to date, it has not been necessary to make assumptions about the degree to which the employer contribution to insurance is effectively paid by workers through reduced wages.

2.5 Non-Group Premiums

Non-group premiums are estimated based on predicted expenditures, plus a loading factor that accounts for administrative expenses and profits. The expenditure data that we used to estimate non-group premiums comes from MEPS-HC respondents who are covered in the individual health insurance market. We used six age and health pools to compute premiums: children (0–17), ages 18–34 years old in good health, ages 18–34 years old in poor health, ages 35–64 years old in good health, ages 35–64 years old in fair health, and ages 35–64 years old in poor health. For purposes of modeling, these age-health pools are likely to capture the major variations in premiums that are charged.

Regulations that affect how insurers set prices in the non-group market vary dramatically by state, but the small size of the pools' population prohibits any state-specific modeling. The only exception is for the six states in which community rating regulations, which restrict insurers' ability to charge differential prices to older and sicker individuals, are in effect (Massachusetts, Maine, New Hampshire, New York, New Jersey, and Vermont)[6]. For these states, we added two special pools--18–34 years old and 35–64 years old--independently of health status. However, there are limitations in the MEPS data, since state of residence is not available on the data set and must be inferred from the statistical match to the SIPP. As a result, we have only a good approximation of the individuals living in the six states that have community rated premiums. Children in community rating states were pooled together with children of other states.[7]

From the raw 2002 SIPP data, we computed initial premiums in each pool, using expected total medical expenditures, an actuarial value of 0.65[8], and an administrative loading fee of 0.35 as a share of the premium.[9] Because simulated non-group premiums are based on very limited

---

[6] Although New York does not permit premiums to vary with age, it was not possible to model New York separately due to sample size.

[7] We did so because, in the other states, we could not distinguish between good, fair, and poor health in children, since so few of them are sick.

[8] Buntin et al. (2003) found that, in 2002, the average actuarial value for individual insurance in California was 0.74, while Gabel et al. (2007) found an actuarial value of 0.56. We split the difference between these figures to get a value of 0.65.

[9] We derive the administrative loading fee based on a minimum loss ratio of 0.65, a typical minimum among states that regulate loss ratios in the non-group market (American Academy of Actuaries, 2004).

data for non-elderly people with individual insurance in the MEPS and SIPP, we performed a revenue neutral adjustment of the premiums (*premium smoothing*, a standard actuarial practice) to obtain results that varied more proportionately to insurance industry expectations of risk. Specifically, we recalibrated premiums to ensure that they did not differ by more than 50 percent for people in different health states within the same age category, with the constraint that total revenue that insurers collect from premiums remained the same as originally estimated. We did not apply these adjustments to the community rated market. Table 2.2 shows the premiums in the simulated equilibrium.

Table 2.2
Non-Group Premiums in the Simulated Status Quo for Year 2007

Pool	Premium (2007 $)
Children	1,915
Ages 18–34, community-rated	2,870
Ages 35-64, community-rated	4,305
Ages 18-34, excellent, very good, or good health	2,305
Ages 18-34, fair or poor health	3,460
Ages 35-64, excellent, very good, or good health	4,285
Ages 35-64, fair health	6,430
Ages 35–64, poor health	7,075

CHAPTER 3
DETERMINING BEHAVIORAL RESPONSES IN THE STATUS QUO

3.1 Household Behavior

We used a sequential approach for estimating the insurance choices made by households. First, we predicted whether the household opted for group insurance. If it did not, then we predicted whether eligible individuals enrolled in Medicaid. If the household enrolled in neither group nor Medicaid, we modeled the selection of non-group coverage. Below, we describe the process for predicting each of these choices.

3.1.1 Selects Group Insurance

To estimate household insurance choice, we would ideally want to control for the availability and price of group insurance, Medicaid eligibility, demographic characteristics, family composition, and health. Unfortunately, few data sets are available that contain geographic and income information sufficient to determine Medicaid eligibility and information on employer premiums. To address these limitations, we used a hybrid approach that combines employee premium contribution information from the MEPS-HC with Medicaid-eligibility information that we impute in the SIPP based on the most recent state-specific Medicaid eligibility rules. Specifically, to get the elasticity of enrollment with respect to worker premium contributions, we analyzed a set of regressions using the MEPS-HC Person Round Plan Level (PRPL) file, which contains self-reported data on employee health insurance premium contributions.

We imputed premium contributions for non-enrolled workers based on data from enrolled workers by running a logit regression to predict whether the employer required a premium contribution or not. Then, for workers with nonzero premium contribution requirements, we predicted the premium contribution amount using an ordinary least squares (OLS) regression. Independent variables in the premium imputation regressions include a variable indicating family or single plan, region, urban versus rural, firm characteristics (size, industry category, multiple-location firm, federal or state employer versus private), and worker characteristics (union or not, full-time versus part-time). We also applied an adjustment factor based on Blumberg et al. (2001) to account for the likelihood that non-enrolled workers would be expected to have higher premium contributions (if they opted to take insurance) than demographically similar enrolled workers. Specifically, among non-enrolled workers, we increased predicted premium contributions by a factor of 1.85 for single plans, and by a factor of 1.32 for family plans. Our elasticity of enrollment with respect to worker premium contributions is -0.03, comparable to the elasticities estimated in Blumberg et al. (2001).[10] After calculating selection elasticities, we reestimated the regressions on the SIPP, controlling for Medicaid eligibility and other characteristics, and adjusting for the MEPS-based selection elasticities.

The final regression correctly reproduces the propensity of different groups—including individuals who are eligible for Medicaid—to enroll in employer insurance, while preserving the elasticity estimates derived from the MEPS-HC.

Because some families have the option to enroll in more than one health insurance plan, we estimated three regressions to determine health plan choice:

- Regression 1: For families eligible for only one employer offer, did they enroll or not?
- Regression 2: For families eligible for one employer offer and enrolled, did they take single or family coverage?
- Regression 3: For families with two employer offers, did they take one or two plans?

The vast majority (97 percent) of families with access to two employer policies take at least one family plan. As a result, we did not explicitly model the behavior of the 3 percent of families who have access to two employer plans but take neither of them, either because they take only single coverage or decline group insurance altogether. Instead, we "froze" the behavior of these HIEUs with regard to group insurance decisions. Specifically, rather than using a regression to predict whether or not these families will take group insurance, we assigned group insurance status according to patterns observed in the SIPP. We still used a regression to predict enrollment in Medicaid and non-group insurance for these families, but we did not allow their

[10] We have a high degree of confidence in the elasticity estimates from Blumberg et al. (2001) because these figures were estimated using the MEPS HC-IC, a unique data set where premiums are reported by firms rather than workers, and where premium contribution requirements for non-enrolled workers can be directly observed.

group insurance enrollment decisions to change under any modeled policy change scenarios. This decision had little effect on the simulation, both because only a small fraction of people are frozen (just 3 percent of HIEUs with two group offers, or less than 1 percent of all HIEUs), and because families with access to two offers usually have higher incomes and are, therefore, minimally affected by the policy options we consider.

For the remaining families with access to two health insurance policies, we used regression 3 to predict whether they took either one or two family plans. Consistent with patterns observed in the SIPP, we did not allow these families to take single coverage. This assumption has no effect on the number of uninsured, but it does have some effect on the consumer's financial risk. It is relatively common for married couples with no children to take two single policies, and two single policies would typically be less expensive than one family plan. However, HIEUs with two group policies tend not to be affected by the types of policy options we model, so we do not think this issue will have a major effect on our results. In future versions of the model, we plan to update the decision making process to allow for the selection of two single plans, as well as for other coverage permutations.

We ran each of these regressions first using the MEPS-HC and then using the SIPP, as described above. The SIPP specification originally included the same covariates as the MEPS regression, to ensure that we obtained the same results with each approach. We then added new variables to the SIPP specification that we found to be predictive and removed variables that did not have predictive power (even if they were statistically significant). Therefore, the final specifications for the SIPP and MEPS differ somewhat. Table 3.1 lists the covariates used for each data source.

Table 3.1
Covariates Included in ESI-Acceptance Regressions

Variable	MEPS Regression	SIPP Regression
Employee premium contribution[a]	X	X (based on MEPS data)
Family income	X	X
Number of children in the family	X	X
Any children less than 5 years old?	X	X
Any children less than 12 years old?	X	X
Two-adult household (versus one-adult)	X	
Any child in poor health?	X	
Any adult in poor health?	X	X
Age of the oldest adult less than 35	X	X
Age of the oldest adult 35–45	X	X
Any adult in the HIEU a minority?	X	X
Any adult in the HIEU college-educated?	X	X
Two-worker household		X
Single household		X
Single parent		X
Traditional household		X
Any child eligible for Medicaid/SCHIP		X
Any adult eligible for Medicaid		X
State dummy variables		X

[a] Premium data included in the models varies with the regression. In regression 1, we used the contribution for the least expensive available plan. In regression 2, we used the difference in price between the least expensive option to cover the entire family and the least expensive single-coverage option. In regression 3, we used the difference in price between the least expensive 1-family plan versus 2-family plan options.

After estimating group health insurance choice, we used the SIPP data to determine enrollment in public insurance and selection of non-group insurance. Below, we describe the methodology used for these analyses.

3.1.2 Selection of Public Coverage

For individuals who are eligible for public coverage, we predicted enrollment using two regressions—one for children and one for adults. Both regressions control for group enrollment, which is estimated as described above. Other covariates included in the adult regressions are income group, age, sex, health, employed versus unemployed, parent versus nonparent, race/ethnicity, education, and Census region. The child regressions exclude parent status, employment status, and education, but include a control for parent's choice of insurance.

3.1.3 Selection of Non-Group Coverage

We estimated selection of non-group coverage in the status quo using two regressions: one for adults and one for children. Covariates included in these regressions are income category, self-employed or not (adult regression only), sex, participation in the Medicaid program, enrollment in employer coverage, race, health, age, and the non-group premium. In the regression for children we control for whether any of the parents has non-group insurance, a covariate that turned out to have high predictive power. We estimated these regressions in a way that allowed us to set an elasticity of demand for health insurance in a fashion similar to the method used by the Urban Institute. The default value for this elasticity is –0.4 (Marquis and Long, 1995). This elasticity is lower than the –0.65 used by Jonathan Gruber. Our elasticity of take-up for non-group coverage does not vary with income.

3.2 Firm Behavior

The initial distribution of employer offers in our data set is based on the distribution found in the Kaiser/HRET survey. When policy options are implemented, firm behavior changes accordingly. A discussion of how firms respond to each type of policy option is found in Chapter 4.

CHAPTER 4
MODELING PARTICULAR POLICY CHANGES

In modeling policy changes, we did not seek to estimate the effects of policies proposed by particular individuals or groups; rather, we created hypothetical policy options that allowed us to demonstrate how the model works and the likely sensitivity of the results to particular design parameters. We also elected, in the initial round of modeling, to estimate "pure" forms of these policy options; that is, we did not provide estimates of the results from combining two or more of these options. That will be done in subsequent releases.

4.1 Employer Mandate

We modeled an employer mandate that requires firms above a certain size to offer health insurance coverage to their workers. We considered three minimum firm sizes—5 workers, 10 workers, and 25 workers. In addition, we modeled three potential penalties for non-participation, set at 5, 10, and 20 percent of payroll. Since health insurance costs are approximately 11 percent of payroll for firms that currently offer insurance, the scenario with a 20 percent penalty is politically unrealistic and used only as a sensitivity analysis. To determine firms' response to employer mandates, we assumed that firms currently provide benefits if the costs of offering insurance to employees were lower than the benefit to workers of having health insurance. The pay-or-play tax imposed under the employer mandate can be thought of as an increase in the benefit associated with providing health insurance. Under this policy option, firms will offer insurance if the benefit to workers plus the penalty that the firm would face under the option is greater than the cost of providing insurance. The value of the tax is specified as part of the policy change and is typically a share of total payroll at the firm.

We were able to calculate the costs of insurance for all firms (both offering and non-offering) using the synthetic data in our model, containing information on worker expenses and worker decision rules about health insurance acceptance, if offered. Although we could not directly estimate the benefit of insurance to workers, we assumed that this benefit was a function of firm and worker characteristics, including firm size, industry, and the age and wage composition of the labor force. We could then fit a logistic regression that allowed us to empirically estimate the benefit of health insurance for workers, both at offering and non-offering firms. We assumed that the firms that newly offer insurance following the policy change offer group insurance plans that have the same actuarial values as plans currently offered by firms and similar eligibility rules (about 80 percent of employees are eligible for group coverage).

Using this methodology, we estimated that the mean value of insurance among offering firms is 11.7 percent of payroll, and the mean value of insurance among non-offering firms is 8.6 percent of payroll. These results, at least for offering employers, can be validated against published literature, which has found that the typical offering firm spends about 11 percent of total payroll on health insurance.[11] Once we estimated benefits, calculated costs, and assigned the pay-or-play tax for each firm, it was straightforward to determine how a business would respond to an employer mandate.

After the firm decides to offer coverage, we allowed its employees to decide to take group insurance according to the same regressions used in the status quo calculation. We also allowed people on Medicaid to choose group insurance and people on non-group coverage to switch to group coverage after gaining access. We did not make any assumptions about what the federal government does with revenue generated by the non-participation tax levied on employers.

---

[11] See Kaiser Family Foundation and Health Research and Educational Trust (HRET), Snapshots: Health Care Costs, Employer Health Insurance Costs and Worker Compensation page, March 2008 (as of November 19, 2008: http://www.kff.org/insurance/snapshot/chcm030808oth.cfm), and Eibner and Marquis (2008).

Many proposals would use extra revenue to facilitate pooling mechanisms and/or coverage subsidies in the non-group market. In a future iteration of the COMPARE microsimulation, we will model the formation of new pools for those working in the firms that do not comply with the mandate.

4.2 Individual Mandate

In our individual-mandate policy option, we assumed that all individuals are required to obtain health insurance coverage, with no exceptions. To provide access to health insurance to people without group insurance, the government creates a national health insurance purchasing pool, whereby coverage is offered by private insurers who agree to comply with federal rating regulations. Plans offered through the purchasing pool must sell policies to all who apply (i.e., have guaranteed issue), and rates can vary only with age. Age bands allowed in the purchasing pool are 0–17 years, 18–29 years, 30-49 years, and 50–64 years.

All plans offered in the purchasing pool have an actuarial value of 70 percent. Administrative costs for plans in the purchasing pool were modeled at 15 percent, although this value can be altered in sensitivity analyses. The federal government fully subsidizes coverage in the purchasing pool for low-income individuals and families with incomes under 100 percent of the federal poverty level; individuals with incomes between 100 and 300 percent of the federal poverty level are eligible for partial subsidies, which decrease linearly as income rises.

We loosely follow the design of the individual mandate enacted in Massachusetts, and restrict access to the purchasing pool: Individuals who are Medicaid eligible and individuals with access to ESI (either through their own firm or through their spouses) cannot buy into the purchasing pool. In future versions of the model we will relax the eligibility constraint related to

ESI, allowing small firms or their employees to access the purchasing pool. We have already modeled a scenario in which the Medicaid eligibility constraint has been relaxed, although we are not showing those results in the current version.

Penalties for nonparticipation in the mandate are expressed as a percentage of the premium that an individual would expect to pay in the purchasing pool, the maximum being 100 percent. We assume that there is a mechanism (for example the tax system) that can be used to enforce the penalty. In future versions of the model we plan to relax this assumption and allow for the penalty not to be enforced in certain population subgroups (such as the very poor who do not pay taxes).

We made a number of assumptions about how individuals and HIEUs respond to the mandate. HIEUs in which all members are fully insured through ESI are unaffected by the mandate. For uninsured singles or HIEUs in which one or more member is uninsured we decided to use a utility maximization framework to determine their choices, since we felt that we were operating at the edge of the domain of validity of the regressions we used to model the status quo. We follow a standard approach in health economics (Goldman, Buchanan and Keeler, 2000; Pauly, Herring and Song, 2002) and model the utility as separable in consumption and health. The consumption terms consists of a quadratic approximation: it includes the negative of the sum of expected out-of-pocket (OOP) expenditures, the OOP (subsidized) premium and the penalty as well as a risk term, modeled as $-0.5 \times r \times Var[\text{OOP}]$, where r is the Pratt risk aversion coefficient. The expected value and variance of the OOP expenditures that one would experience under a certain type of insurance are computed for each individual using the distribution of OOP expenditures of "similar" individuals with that type of insurance, where the definition of similarity depends on personal characteristics such as age, income and health status. The utility of health is measured in willingness-to-pay units, and is approximated by a fraction ε of the expected total medical expenditures. While Pauly et al (2002) used a value of 0.5 for ε we found that a value of 0.3 better calibrates, that is it better reproduces the choices made by individuals and HIEUs in the status quo. Pauly used a value of 0.00095 for r, in 2001 dollars, which corresponds to 0.00082 in 2007 dollars. However, Manning and Marquis (1996) found a value of r of 0.00021 in 1995 dollars, which corresponds to 0.00015 in 2007 dollars. We decided to use the average of these two numbers as our base case, which amounts to 0.00049 in 2007 dollars. Some additional assumptions had to be made. The utility maximization by itself does not explain the presence of a large number of individuals who are Medicaid eligible but uninsured. We model the unobserved factors that make these people uninsured by assigning to each individual a disutility associated to Medicaid. The assignment is made in such a way that it satisfies the following constraints: 1) when the penalty is p%, the take-up rate of Medicaid in this group is also p%; 2) individuals who are predicted by the Medicaid take-up regressions to have a higher take-up probability are assigned lower disutility.

Similarly, the utility maximization framework does not explain the large number of people with access to ESI who remain uninsured. We assume that these individuals do not take up ESI because they face a higher employee premium share. Therefore, we randomly assign them an employee premium share between the average (16 percent for singles) and 80 percent. In order to check the consistency of these assumptions, we test them on a "non-reform", that is we run the model with no penalty, no subsidy and no purchasing pool and verify that the model predicts, within few percent, the observed pattern of insurance selection.

Finally, while the utility associated with being on an ESI plan, or non-group, or Medicaid, or uninsured can be computed using data from the MEPS, which contains information about all of those insurance categories, the utility associated with being on the purchasing pool requires some adjustment because there is not such a plan in the MEPS. However, the purchasing pool is modeled as slightly more generous than the average non-group plan, Therefore we use expenditures of people on non-group insurance to estimate expenditures for people on the purchasing pool, except that we make a small adjustment to the OOP expenditures to reflect the difference in actuarial values.

Using the above assumptions we are able to assign a utility value to each possible combination of insurance choices, for individuals and for HIEUs[12]. There are three main groups of people whose choices we must consider:

1) Uninsured living in HIEUs with no ESI offers. Within this group, those who are not Medicaid eligible have the following choices: non-group, purchasing pool, uninsured. For the Medicaid eligible, Medicaid substitutes for the purchasing pool option.
2) People currently on individual insurance and living in HIEUs with no ESI offer: these people may switch to the purchasing pool, which is slightly more generous but has lower administrative costs and therefore it is likely to have favorable premiums, especially for those in poor health. If enough people switch, the non-group market that exists today may shrink and disappear, or survive by offering inexpensive plans to the very healthy.
3) Uninsured living in HIEUs with an ESI offer: the HIEU may decide to cover its members with any combination of individual insurance, group insurance or Medicaid (for those who are eligible) or to remain uninsured. Notice that the non-group market may disappear, leaving this large group with the option of buying ESI (which is presumably quite expensive for them, since they did not buy it in the status quo) or remaining

[12] The utility of an HIEU is the sum of the utilities of its members.

uninsured. The uptake of insurance in this group is therefore strongly dependent on the size of the penalty.

Firms may play a role in the individual mandate too. The penalty for being uninsured makes health insurance more valuable to employees, which in turn should increase the probability that a firm offers ESI. In the current version we are not assigning any behavior to firms as a result of the individual mandate.

4.3 Tax Credits

We model a refundable tax credit that cannot exceed the value of the health insurance premium. We also assume that eligibility for the tax credit depends on household income. The tax credit is defined by six parameters:

- The size of the credit for singles, which ranges from $1,000 to $5,000 in our modeled scenarios.
- The size of the credit for families, which we assume is 2.5 times the credit for singles.
- The lower income threshold for singles, which we set at $15,000. All singles with incomes below this level received the full credit.
- The upper income threshold for singles, which we set at $30,000. All singles with incomes above this level were ineligible for the credit. Singles with total income between the lower and the upper income thresholds receive a credit on a sliding scale, decreasing from full (at the lower threshold) to 0 (at the upper threshold).
- The lower income threshold for families, which we set at $30,000. All families with incomes below this level receive the full credit.

- The upper income threshold for families, which we set at $60,000. Families with incomes above this level are not eligible for the credit. Families with total income between the lower and the upper income thresholds receive a credit on a sliding scale, decreasing from full (at the lower threshold) to 0 (at the upper threshold).

We assume that tax credits have no effect on the behavior of households that are not eligible for the credits. In addition, we assume that some eligible individuals who are currently enrolled in the non-group market will *not* claim the credit, either because of the transaction cost, or because they are unaware of their eligibility. The probability of claiming the credit is proportional to the size of the subsidy: Those with full subsidy have a probability equal to 90 percent of claiming, and this probability decreases to 40 percent as the level of subsidy goes to 0. We further assume that no one who is currently enrolled in Medicaid will switch to non-group coverage in response to the tax credit.

Although the regressions used to model the status quo are appropriate for modeling participation in the current non-group market, they may not be adequate for modeling a tax credit in the non-group market. In the status quo, non-group health insurance is typically an unattractive option for consumers because of its higher prices, disadvantaged tax treatment relative to employer coverage, and restrictive enrollment policies in some states. As a result, acceptance elasticities observed in the status quo are underestimated. A major change in the tax treatment or price of non-group policies could lead to a much larger change in non-group enrollment than the status quo models predict.

To address this issue, we followed the Urban Institute and applied some correcting factors to the probabilities estimated by the logits. A detailed explanation of how these correcting factors are derived is found in the Urban Institute (2008) Health Insurance Reform Simulation Model

(HIRSM) documentation. In short, the idea is to understand how people react to large changes in prices by comparing group insurance participation for people whose employer does not pay any portion of the premium with that of people whose employer pays significant portions of the premium. The method involves reducing the probability of declining non-group insurance by the ratio of insurance take-up in the group market for workers required to pay the full premium relative to take up in the group market for workers whose employer contribution is similar to the subsidy available in the non-group market. Specifically, the adjusted take-up in the non-group market is equal to 1-[(1-*predicted take-up in the non-group market*)*(*group take-up, no employer contribution*)/(*group take-up, employer contribution similar to non-group subsidy*)]. We think that adopting the Urban Institute approach strengthens our ability to model behavior changes stemming from a change in the tax treatment of insurance; however, both individuals and insurers could respond to tax credits in unforeseen ways. For example, insurers might begin to offer new policies once the tax code is changed, which could alter participation decisions by individuals. A limitation of our model is that we cannot fully anticipate how insurers and individuals will respond to major policy changes that fundamentally alter the market for insurance.

In addition, to account for new participation in non-group coverage, we also need to include the possibility that some individuals will switch from group to non-group insurance in response to the tax credits. Our current model of the status quo does not allows us to simulate this effect directly, because the decision to take group health insurance or not is made without incorporating data on the price and availability of non-group policies. Instead, we accounted for switching from group to non-group coverage using Gruber's (2000) approach. Specifically, the probability of switching from group to non-group insurance is a function of the difference in price between out-of-pocket payments for the group policy and the post-subsidy costs of the non-group policy, scaled by the total cost of the group policy and multiplied by an adjustment factor of 0.625. See Gruber (2000, p. 45) for a more detailed discussion.

We further assume that firms will drop coverage if a substantial fraction of their labor force is eligible for the credit. If 100 percent of workers are eligible for a tax credit, the maximum share of firms will drop coverage, based on the predetermined upper bounds shown in Table 4.1. The upper-bound dropping probabilities are based on figures reported in Gruber (2000).

Table 4.1
Upper Bound Percentage of Firms Dropping Coverage,
Under a Tax Credit

Firm Size, based on number of workers	Upper Bound*
<10 workers	45%
10 to 24 workers	35%
25 to 99 workers	20%
100 to 499 workers	5%
500 to 999 workers	3%
1,000 or more workers	0%

*This is the proportion of firms that would drop coverage if all of its workers were eligible to take up the tax credit.

4.4 Medicaid Expansion

The Medicaid expansion considered in our model extends eligibility for Medicaid or SCHIP to all individuals with family incomes varying between 100 and 300 percent of the poverty level. We assume that all individuals below a given income level qualify for the new program and that there is no cost sharing on the part of individuals. If a state's current minimum eligibility requirements exceed the new minimum standard, we assume that the state *does not* change its eligibility threshold.

When eligibility for public coverage expands, we assume that there is some crowd out from group insurance. Crowd out is modeled with the same household level regressions used to simulate the status quo, since in the specification we controlled for the presence of an adult and/or a child eligible for Medicaid. Therefore, the increased eligibility causes some households that previously chose group coverage to switch to Medicaid. However, the regressions are not designed to estimate crowd out, and we found levels of crowd out to be on the low end.

To address this issue, we introduced an adjustment factor that allowed us to increase crowd out up to a maximum level at which all Medicaid-eligible individuals drop group coverage in favor of Medicaid. Many of our model runs assume a Medicaid crowd out rate of 0.50, which is based on a study by Marquis and Long (2003) that used variations from state Medicaid expansions in the 1990s to identify the effect of changes in public coverage eligibility on worker take-up in the group market. However, other studies have found crowd-out rates ranging from less than 10 percent (Shone et al., 2008) to as much as 60 percent (Gruber and Simon, 2008). In sensitivity analyses, we checked how the model responded to different assumptions about the magnitude of this adjustment factor.

We also allowed firms to drop insurance in response to a Medicaid expansion, since this policy change would give workers an alternative source of coverage (and thus lower the value of health insurance benefits for some workers). To model employers dropping insurance coverage, we used a methodology similar to that described in Gruber (2000), and we allowed the share of firms that drop coverage to vary by the generosity of income eligibility. The probability of dropping coverage varies linearly between zero and a predetermined upper bound, depending on the fraction of workers that would be eligible for Medicaid under the policy change. We assumed that, if 100 percent of workers were eligible for Medicaid, the predetermined upper-bound fraction of firms would drop coverage. Table 4.2 lists the predetermined upper-bound dropping probabilities under a Medicaid expansion.

Table 4.2
Upper Bound Percentage of Firms Dropping Coverage, Under a Medicaid Expansion

Firm Size, based on number of workers	Upper Bound*
<10 workers	22.5%
10 to 24 workers	17.5%
25 to 99 workers	10%
100 to 499 workers	2.5%
500 to 999 workers	1.5%
1,000 or more workers	0%

*This is the share of firms that would drop coverage if 100% of workers were eligible for subsidy.

Probabilities of dropping ESI in the face of a Medicaid expansion are one-half the value of the probabilities that firms drop ESI with a tax credit (shown in Table 4.1), to reflect factors, such as stigma and hassle, which may preclude some eligible workers from seeking coverage through Medicaid.

Uninsured individuals who are eligible in the status quo for Medicaid but do not take it are assumed to remain uninsured. However, we modeled a small (5-percent) increase in the overall participation rate when eligibility for coverage is expanded, to take into account a reduction in Medicaid stigma. This increase was held constant at 5 percent regardless of whether the policy change expanded Medicaid to households with incomes below 100, 200, or 300 percent of the federal poverty level.

CHAPTER 5
UTILIZATION CHANGES IN RESPONSE TO BECOMING INSURED

When individuals in our model become insured, we would expect their health care utilization to increase. Yet, it is not straightforward to estimate the change in utilization from obtaining health insurance coverage. Uninsured individuals may tend to be healthier or have a lower propensity to use health care than the currently insured population. As a result, utilization levels of the currently insured population probably overestimate the likely utilization of the currently uninsured population, were they to gain coverage.

Hadley et al. (2008) used data from the MEPS to simulate the cost of care for the uninsured if they were to become covered; the study found that an adult who moves from being full-year uninsured to full-year insured would increase spending by 124 percent. However, because this approach estimates spending for the currently uninsured using observational data from similar individuals who have insurance, it may overstate spending changes.

Using data from the RAND Health Insurance Experiment (HIE), which randomly assigned individuals to health care plans, Keeler et al. (1988) estimated that individuals with plans that provide free care would spend 80 percent more than those with a 95 percent coinsurance plan with no limits. Currently uninsured people get some subsidized care, but they face more obstacles than people who pay for care directly, so results from the HIE cannot tell us exactly how much spending will increase when individuals move from being uninsured to being insured. Still, the HIE results suggest that spending increases estimated using observational data may overstate actual spending changes.

To address this issue, we assumed that spending for the full-year uninsured increases by *one-half* the spending change reported in Hadley et al. (2008) when they become insured either through group coverage or Medicaid. When individuals move from being part-year uninsured to insured, we used spending increases reported in Hadley et al. (2008) without making adjustments. Full-year and part-year insurance status are reported in the SIPP. We made the assumption that everyone who becomes insured stays insured for at least a year. However, we also assumed that people who lose insurance coverage when their employer drops insurance remain uninsured for only part of the year. The latter assumption reflects the possibility that individuals who previously had insurance might be more likely than chronically uninsured individuals to find alternative coverage sources if their employer stops offering insurance.

Our approach for estimating spending for the newly insured implicitly assumes that spending levels in group coverage are equivalent to spending levels in Medicaid. On paper, Medicaid is more generous than group coverage, suggesting that the total increase in spending generated when uninsured people switch to Medicaid might be higher than spending levels generated when uninsured people switch to group coverage. However, we assume that access problems for those on Medicaid counteract any such differentials (Berk and Schur, 1998; Cunningham, 2005). Because non-group plans are typically less generous than group coverage, we assume that switching from group to non-group coverage decreases spending by 5 percent.

Table 5.1 shows the complete schedule that we used to estimate spending changes that occur when individuals change insurance status.

Table 5.1
Change in Health Care Spending After Becoming Insured

Type of coverage change	Increase in health care spending
Full-year uninsured to group or Medicaid coverage, adults	62%
Part-year uninsured to group or Medicaid coverage, adults	42%
Full-year uninsured to group or Medicaid coverage, children	37%
Part-year uninsured to group or Medicaid coverage, children	20%
Full-year uninsured to non-group coverage, adults	54%
Part-year uninsured to non-group coverage, adults	35%
Full-year uninsured to non-group coverage, children	30%
Part-year uninsured to non-group coverage, children	14%

We further assume that current health care spending by the uninsured is partially billed as uncompensated care. Using figures reported by Hadley et al. (2008), the American Hospital Association 2007 Annual Survey, and Gruber and Rodriguez (2008), we calculated that uncompensated care spending totaled $52.1 billion in 2007, or $1,100 per uninsured individual. Of this amount, 75 percent was paid for by the government and 25 percent was paid for by private insurers. Using these figures, we could account for savings due to uncompensated care in our model estimates. Savings recouped by private insurers are passed back to consumers as lower premiums. We do not make any assumptions about how the government uses savings generated through reduced, uncompensated care spending.

CHAPTER 6
DETERMINING THE POLICY OPTION'S EFFECT ON HEALTH OUTCOMES

The COMPARE microsimulation does not include a module to analyze the health benefits associated with changes in health care policy options. Rather, we relied on results from the RAND Future Elderly Model (FEM) (Goldman et al., 2004) to determine years of life gained by becoming insured. The FEM is a dynamic microsimulation model based on the Health and

Retirement Study (HRS; Institute for Social Research), which tracks older Americans (aged 51 and above) over time to project their health and various economic outcomes. The current model is an extension of a previous version that focused exclusively on Medicare beneficiaries (Goldman et al., 2004).

We used the distribution of insurance status in the SIPP database, combined with U.S. life tables published by the U.S. Census Bureau, to generate separate life tables for insured and uninsured individuals under age 65 stratified by gender and race (black and non black). Given the assumption that treatment and adherence differences between insured and uninsured people would apply for major chronic conditions, calculations from the FEM showed the difference in mortality rates between insured and uninsured individuals from age 51 to 65. Results from the FEM implied that, on average, a 51 year old insured individual who remains insured until becoming Medicare eligible lives an average of 0.46 years more before age 65 than a similar uninsured individual of age 51 who remains uninsured during the same age period.

To reproduce 0.46 expected life years gained using the U.S. Census Bureau mortality tables, we assumed that for each age *a* between 51 and 65 years, the population averaged mortality rate multiplied by the total population of age *a* is equal to the averaged mortality rate of the insured at age *a* multiplied by the insured population at age *a,* plus the averaged mortality rate of the uninsured at age *a* multiplied by the uninsured population at age *a*. The size of insured and uninsured populations as well as the size of the total population was taken from the SIPP database. The mortality rate of the uninsured was assumed to be larger than that for the insured by a constant percentage difference over the entire age range 51 to 65.

Because evidence suggests substantial catch-up in health after people uninsured in their 50s become eligible for Medicare at age 65, we assumed that the entire mortality effect would occur between ages 51 and 64. To reproduce the 0.46 expected life years gained, the population average mortality difference between insured and uninsured would have to be 9.06% between ages 51 to 65 and 0% for individuals over age 65. If instead we assumed that people gaining insurance have lower mortality after age 65 for the rest of their lives, the mortality difference would only be 4.87%.

The mortality table for those insured and uninsured of age 50 and below was generated on the assumption that this mortality difference of 9.06% applies to all ages below 65 and not just to those in age group 51 to 65. Since mortality rates up to age 50 are comparatively small, this constant value in mortality difference has very little effect between ages 0 to 50 on the overall expected life years gained by an insured individual compared with an uninsured individual.

On the assumption that the SIPP database represents the insured and uninsured population at equilibrium, we calculated the transition probabilities from insured to uninsured and vice versa as a function of age and poverty category. This calculation assumes that the transitions are Markov. From these probabilities we calculated for each poverty category the probabilities of being insured and of being uninsured at age 50, given that one is insured (uninsured) at a previous age. In so doing, we considered all possible paths between being insured and uninsured that an insured (or uninsured) individual of a certain age and poverty category takes before reaching the age of 50. We sought to use this information to give a better estimate of the difference in mortality rates between an insured and uninsured individual of age 50 and below. However, as mentioned previously, mortality rates for those age 50 and below are very small in comparison to those above the age of 50. Moreover, changes from insured to uninsured are probably not Markov, as some individuals may be uninsured for long periods, while others move back and forth between being uninsured and insured. Therefore, we decided not to use this additional information in calculating the extra life years gained of an insured individual, compared with an uninsured individual, between birth and age 50. Instead we assumed that the mortality difference of 9.06% applies equally to those of all ages below 50 as it does to those between 51 and 65.

REFERENCES

Agency for Healthcare Research and Quality, Medical Expenditure Panel Survey–Household Component, Rockville, Md., n.d. As of November 19, 2008: http://www.meps.ahrq.gov/mepsweb/

American Academy of Actuaries. "Report to the NAIC's A&H Working Group of the Life and Health Actuarial Task Force." American Academy of Actuaries Rate Filing Task Force, May 12, 2004.

American Hospital Association, 2007 Annual Survey, Chicago, Ill.: American Hospital Association.

Berk C, Schur, "Access to Care: How Much Difference Does Medicaid Make?" Health Affairs, Vol 17. No. 3, 1998, 169-180.

Blumberg LJ, Shen Y-C, Nichols LM, et al., *The Health Insurance Reform Simulation Model (HRSIM): Methodological Detail and Prototypical Simulation Results*, Washington, D.C.: The Urban Institute, 2003. As of November 19, 2008: http://www.urban.org/UploadedPDF/410867_HIRSM_Report.pdf

Blumberg LJ, Nichols LM, Banthin JS, "Worker Decisions to Purchase Health Insurance" International Journal of Health Care Finance and Economics, Vol. 1 No. ¾, 2001, pp. 305-325.

Buntin, MB, Escarce JJ, Kapur K, Yegian J, and Marquis MS, "Trends and Availability of Individual Insurance Products in California." Health Affairs, web exclusive, September 24, 2003, w3-449-w3-459.

Congressional Budget Office (CBO), *CBO's Health Insurance Simulation Model: A Technical Description*, Washington, D.C.: The Congress of the United States, Background Paper, 2007. As of November 19, 2008: http://www.cbo.gov/ftpdocs/87xx/doc8712/10-31-HealthInsurModel.pdf

Cunningham, PJ. "Medicaid Cost Containment and Access to Prescription Drugs," Health Affairs, Vol. 24 No.3, 2005, 780-789.

DeNavas-Walt C, Proctor BD, Smith JC, *Income, Poverty, and Health Insurance Coverage in the United States: 2007*, Washington, D.C.: U.S. Census Bureau, Current Population Reports: Consumer Income, P60-235, 2008.

Eibner C, Marquis MS, "Employers' Health Insurance Cost Burden: 1996–2005," *Monthly Labor Review*, Vol. 131, No. 6, 2008, pp. 28–44.

Gabel J, Pickering J, McDevitt R, Whitmore H, Gandolfo L, Lore R, Wilson K, "Trends In The Golden State: Small-Group Premiums Rise Sharply While Actuarial Values For Individual Coverage Plummet" *Health Affairs*, Vol. 26, No. 4, 2007, w488-w499.

Goldman DP, Buchanan JL, Keeler EB, "Simulating the Impact of Medical Savings Accounts on Small Business" *Health Services Research*, Vol. 35, No.1 pt. 1, pp. 53-75.

Goldman DP, Shekelle PG, Bhattachanja J, Hurd M, Joyce GF, Lakdawalla DN, Matsui DH, Newberry, SJ, Panis CWA, Shang B, *Health Status and Medical Treatment of the Future Elderly: Final Report*, Santa Monica, Calif.: RAND Corporation, TR-169-CMS, RAND, Santa Monica, 2004. As of November 19, 2008: http://www.rand.org/pubs/technical_reports/2005/RAND_TR169.pdf

Gruber J, *Tax Subsidies for Health Insurance: Evaluating the Costs and Benefits*, Cambridge, Mass.: National Bureau of Economic Research, Working Paper 7553, 2000. As of November 19, 2008:
http://www.nber.org/papers/w7553.pdf

Gruber J, Simon, KI, "Crowd-out 10 Years Later: Have Recent Public Insurance Expansions Crowded Out Private Health Insurance?" *Journal of Health Economics*, Vol. 27, 2008, 201-217.

Gruber J, Rodriguez D, *How Much Uncompensated Care Do Doctors Provide?* Cambridge, Mass.: National Bureau of Economic Research, Working Paper 13585, 2008.

Hadley J, Holahan J, Coughlin T, Miller D,. "Covering the Uninsured in 2008: Current Costs, Sources of Payment, and Incremental Costs," *Health Affairs*, Vol. 27, No. 5, 2008, pp. w399–w415.

Institute for Social Research, Health and Retirement Study, Ann Arbor, Mich.: The University of Michigan. As of November 20, 2008: http://hrsonline.isr.umich.edu/

Kaiser Family Foundation and Health Research and Educational Trust (HRET) Web site, Employer Health Benefits 2006 Annual Survey page, Menlo Park, Calif., and Chicago, Ill. As of November 19, 2008: http://www.kff.org/insurance/7527/index.cfm

___, Snapshots: Health Care Costs, Employer Health Insurance Costs and Worker Compensation page, March 2008. As of November 19, 2008: http://www.kff.org/insurance/snapshot/chcm030808oth.cfm

Keehan S, Sisko A, Truffer C, et al., "Health Spending Projections through 2017: The Baby Boom Generation Is Coming to Medicare," *Health Affairs*, Vol. 27, No. 2, 2008. w145-w155.

Keeler E, Buchanan J, Rolph J, Hanley JM, Reboussin D, *The Demand for Episodes of Medical Treatment in the Health Insurance Experiment*, Santa Monica, Calif.: RAND

Corporation, R-3454-HHS,1988. As of November 19, 2008: http://www.rand.org/pubs/reports/R3454/

The Lewin Group, *Summary Description of the Health Benefits Simulation Model (HBSM)*, Falls Church, Va.: The Lewin Group, 2007. As of November 19, 2008: http://www.lewin.com/content/Files/HBSMSummary.pdf

Manning WG, Marquis MA, "Health insurance: the tradeoff between risk pooling and moral hazard", *Journal of Health Economics* Vol. 15 No. 5, 1996, pp. 609-639.

Marquis MS, Long, SH, "Worker Demand for Health Insurance in the Non-Group Market," Journal of Health Economics, Vol. 14, No.1., pp. 47-63.

Marquis MS, Long SH, "To Offer or Not to Offer? The Role of Price in Employers' Health Insurance Decisions," *Health Services Research*, Vol. 36, No. 5, 2001, pp. 935–958.

Marquis MS, Long SH, "Public Insurance Expansions and Crowd Out of Private Coverage," *Medical Care* Vol. 41, No. 3, 2003, pp. 344-356.

Pauly MV, Herring B, Song D, "Tax Credits, the Distribution of Subsidized Health Insurance Premiums, and the Uninsured," Forum for Health Economics & Policy, Vol. 5: (Frontiers in Health Policy Research), Article 5, 2002, http://www.bepress.com/fhep/5/5

Shone LP, Lantz PM, Dick AW, Chernew ME, Szilagyi PG. Crowd-Our in the State Children's Health Insurance Program (SCHIP): Incidence, Enrollee Characteristics and Experiences, and Potential Impact on New York's SCHIP," Health Services Research Vol. 43 No. 1, part II, 2008, 419-434.

Sing M, Banthin JS, Selden TM, et al., "Reconciling Medical Expenditure Estimates from the MEPS and NHEA, 2002," *Health Care Financing Review*, Vol. 28, No. 1, 2006, pp. 25–40.

U.S. Census Bureau, Survey of Income and Program Participation, Washington D.C.: Demographics Survey Division, Survey of Income and Program Participation Branch, 2008. As of November 19, 2008: http://www.census.gov/sipp/

U.S. Census Bureau, Statistics of U.S. Businesses, Washington, D.C.: Company Statistics Division. As of November 19, 2008: http://www.census.gov/csd/susb/index.html

Urban Institute Web site, The Health Insurance Reform Simulation Model (HIRSM) page, 2008, section 11.2. As of November 19, 2008: http://www.urban.org/publications/410867.html