

NBER WORKING PAPER SERIES

LET THEM HAVE CHOICE:
FROM SHIFTING AWAY FROM EMPLOYER-SPONSORED HEALTH INSURANCE AND TOWARD AN INDIVIDUAL EX

Leemore Dafny
Katherine Ho
Mauricio Varela

Working Paper 15687
<http://www.nber.org/papers/w15687>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2010

We are grateful for helpful comments by seminar participants at the Yale Economics Department and the Yale School of Public Health. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Leemore Dafny, Katherine Ho, and Mauricio Varela. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Let them Have Choice: Gains from Shifting Away from Employer-Sponsored Health Insurance
and Toward an Individual Exchange

Leemore Dafny, Katherine Ho, and Mauricio Varela

NBER Working Paper No. 15687

January 2010

JEL No. I11,L1

ABSTRACT

Most non-elderly Americans purchase insurance through their employers, which sponsor a limited number of plans. We estimate how much employees would be willing to pay for the right to apply their employer subsidy to the plan of their choosing. We make use of a proprietary dataset containing information on plan offerings and enrollment for 800+ large employers between 1998 and 2006; the dataset represents over 10 million Americans annually. We estimate a model of employee preferences using the set of plans they are offered. Using the estimated parameters from this model, we predict employees' choices in a hypothetical world in which additional plans in a market are available to them on the same terms, i.e. tax-free and subsidized by their employers. Holding employer outlays constant, we estimate that the median welfare gain from expanding choice amounts to roughly 20 percent of premiums. For the vast majority of employee groups and alternative model specifications, the gains from choice are likely to outweigh potential premium increases associated with a transition from large group to individual pricing.

Leemore Dafny
Department of Management and Strategy
Kellogg School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL 60208-2001
and NBER
l-dafny@kellogg.northwestern.edu

Mauricio Varela
Dept of Management and Strategy
Kellogg School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL 60208-2001
m-varela@kellogg.northwestern.edu

Katherine Ho
Columbia University
Department of Economics
1037 International Affairs Building
420 West 118th Street
New York, NY 10027
and NBER
kh2214@columbia.edu

1. Introduction

Over 60 percent of nonelderly Americans purchase employer-sponsored health insurance (ESI).¹ Although there are no legal impediments to offering a broad array of plans, in practice employers offer a very limited set of choices: a 2005 survey by the Kaiser Family Foundation/Health Retirement Education Trust finds 80 percent of employers who offer insurance provide only one option. The restriction of employee choice may prevent individuals and families from selecting the healthplan that best suits their needs, and from trading off added benefits against the associated premium increases.

With the U.S. preparing to embark on the most aggressive healthcare reform since the introduction of Medicare in 1965, the possibility of leveling the playing field between group and individual insurance (through a variety of means) has come to the fore. Although many have expressed concerns about the erosion of employment-based coverage, partly because of higher predicted costs of individual plans, the benefits associated with expanded choice have never been systematically examined, let alone quantified. In this paper, we use a large panel dataset on employer offerings and employee choices to infer the gains consumers would enjoy were they able to select from a broader spectrum of plans in their local market, holding constant employers' spending on employee subsidies and the full tax-deductibility of premiums. By quantifying the gain to individuals from being able to select any plan available in their local market, we back out the amount by which prices would have to increase to fully offset this gain. In so doing, we provide policymakers with guidance regarding the implementation and design of reforms that bolster individual choice. In a companion paper (Dafny, Ho and Varela 2010), we examine the

¹ Using data from the March 2009 Current Population Survey, Fronstin (2009) reports 61.1 percent of the non-elderly population had employment-based health benefits in 2008. The analogous figure for 2000 is 68.4 percent .

distributional consequences of expanded choice and contrast the characteristics of plans selected by employers and those that would be selected by employees if they were available on the same terms.

We use a unique dataset of employer plan offerings and employee plan selections for a national sample of 800+ large U.S. employers during the period 1998-2006, representing over 10 million employed Americans in every year. Our approach consists of three distinct components. First, we estimate a discrete choice model of employee demand for healthplans, conditioning on the set of plans offered by the relevant employer in the relevant geographic market and year. The parameters from this model reflect the values placed by employees on individual plan characteristics. Second, we estimate a hedonic model of premiums that permits us to predict the premiums a given employee would face for each plan offered in her local market ignoring premium increases due to reductions in group size. Third, we use the demand estimates, together with the predicted premiums, to predict employee choices of plans and their expected utility when offered additional plans currently existent in that market and year. The counterfactuals are budget-neutral for employers; that is, their total contributions to health insurance are held constant. Conceptually, the counterfactual is akin to granting employees a voucher equal to their employer's present contribution to health insurance, valid for the purchase of insurance plans on the individual market (which could be a regulated "exchange"). We use the results to estimate the amount by which premiums would need to increase (relative to the levels predicted by our hedonic model, which assumes group-based pricing) to fully offset the net gain in consumer surplus.

We find choice is worth quite a bit for most individuals: in our most conservative hypothetical scenario the median employee would enjoy a surplus gain of roughly 20% of

combined employer and employee premium contributions. In year 2000 dollars, this gain is approximately \$2,025 for a family of four. Combining these figures with data on employer subsidies, we find the median employee would be willing to forego 27 percent of her employer subsidy simply for the right to use what remains toward a plan of her choosing.² (As an analogy, consider the employer who offers her employee a choice of heavily subsidized vehicles: the Ford Focus or the Cadillac Escalade. The employee would trade a non-trivial percent of the employer subsidy in exchange for the freedom to use the subsidy toward her most-preferred vehicle, assuming it is available at the same price as currently paid by employers who buy in bulk. Of course, we do not anticipate the net benefits of choice to be as large as our estimates, since premiums are likely to increase with the devolution of insurance to the individual marketplace, but as we discuss below, available estimates of cost increases associated with shifting this population to an individual marketplace are generally smaller than the estimated gains.

We caution that our results provide a conservative estimate of the value of choice (or, equivalently, a low estimate of the premium increases that would offset the benefits of choice). First, our data enable us to build a very rich logit model of choice for a given set of employees, but we do not incorporate differences in preferences across individuals *within* employee groups except through a random error term. There may be substantial gains from better matching of individuals to plans. Second, due to the well-known limitations of the logit choice model, we do not expand the choice set to include *all* plans we observe in a given market, except to provide an upper-bound estimate of the value of choice.³ In fact the conservative scenario mentioned above

² This estimate is obtained from our preferred specification, described below; estimates from other models are also presented.

³ This upper bound is itself underestimated, as we observe only a sample of plans available in each market and year, and our expanded choice sets deliberately exclude plans observed only a small number of times (to ensure that included plans are truly active in a given marketplace, and that their premiums can be accurately predicted for each employee group).

holds constant the number of plans in the choice set and simply switches the observed options with those that are most preferred by employees in the relevant firm and market.

The paper proceeds in seven sections. Section 2 discusses the recent trends in the degree of choice in employer-sponsored plans and summarizes related research. Section 3 describes the data. Section 4 presents the estimation strategy and results for the demand and hedonic models used as inputs into the simulations presented in Section 5. Section 6 discusses the implications of the results, and Section 7 concludes.

2. Background

2a. Employer-Sponsored Insurance Plans: How Much Choice Is There and How is This Changing?

Most workers who receive insurance through their employers have a choice of plans but this choice can be quite limited. The Kaiser Family Foundation/ Health Research and Educational Trust Employer Health Benefits Annual Survey (hereafter Kaiser/HRET survey) studies the percentage of workers with job-based coverage, the kinds of plans employees offer and the choices made. Approximately 2,000 randomly selected employers are surveyed, covering a range of industries and both public and private firms. The survey indicates that sixty percent of firms, and ninety-eight percent of firms with over 200 workers, offered health benefits in 2005. As mentioned earlier, eighty percent of firms offered a single plan. However, Figure 1 shows that larger firms offered more choice than smaller firms: twenty-seven percent of large firms (those with 1000-4999 workers) and only seventeen percent of firms with over 5,000 workers offered a single plan. Overall, sixty-three percent of covered workers could choose from multiple healthplans.

The most common healthplan offered to workers in 2005 was a PPO plan: 82% of covered workers had access to this type of plan. Figure 2 documents that 28% of covered workers had access to a POS plan, 44% had access to an HMO and only 12% had access to a conventional indemnity plan. Indemnity plans have become less widely available over time while the availability of PPO plans has increased dramatically since 1988. The patterns in the dataset we use are similar to those in the survey, although our sample is skewed towards larger firms so that choice is less limited than is the case for the average employee. For example, in 2005, about half of the employee groups in our sample are offered a single option. The choice sets observed in our data are discussed further in Section 3.

2b. How Do Employees and Employers Choose Among Plans?

Several studies in the health economics and health policy literatures investigate the factors influencing employees' choice of healthplans. A much smaller set of papers examine employer decision-making, with an emphasis on whether healthplan quality affects employer choices. To our knowledge, no study combines empirical analysis of both decisions, preventing any quantitative assessment of the tradeoffs associated with allocating decision rights to one or the other party. In the review that follows, we focus exclusively on studies that pertain to the working population, as our data includes only active employees.

2b.i. Employee Choice of Healthplans

Most studies in this category focus on the sensitivity of employees to variations in plan price and quality, as measured by items included in the Health Plans Employer Data and Information Set (HEDIS) and the Consumer Assessment of Healthcare Providers and Systems (CAHPS) survey.⁴ The range of price elasticities estimated in these papers is quite broad.

⁴This survey primarily addresses consumer satisfaction, and is maintained by the governmental Agency for Healthcare Research and Quality (AHRQ).

Several studies using data on university employees find elasticities of demand exceeding one, including Cutler and Reber (Harvard; 1996), Royalty and Solomon (Stanford; 1999) and Buchmueller and Feldstein (University of California; 1997). However, Carlin and Town (2009) report low semi-elasticities of demand among employees and students of the University of Minnesota, where a \$100 increase in the annual employee contribution translates to a less than 1% decrease in market share.

By contrast, estimates of elasticity based on the behavior of non-university populations are consistently low. Chernew, Frick, and McLaughlin (1997), using a sample of single workers in small businesses in seven metropolitan areas, find that, while participation of low-income workers in employer-sponsored plans is higher when net premiums are lower, even large subsidies will not induce all to participate. Blumberg, Nichols, and Banthin (2001), using the MEPS data set that contains a nationally representative sample of 6,500 workers offered insurance, find a small price elasticity of take-up (below 0.05) for families, and a very small and insignificant elasticity for single persons. Gruber and Washington (2003), using personnel records for all federal employees from 1991 through 2002, find a small elasticity of employer insurance take-up with respect to its after-tax price (a value of approximately 0.02).

The relevant studies that consider the sensitivity of employee decisions to healthplan quality include Wedig and Tai-Seale on federal employees (2002), Beaulieu on Harvard employees (2002), and Scanlon et al (2002) and Chernew et al (2008) on General Motors employees. Generally speaking, these studies find modest reactions to quality information. It is possible these aggregate effects mask larger responses by populations with stronger incentives to respond, however, the evidence to date on this matter is mixed.⁵

⁵ Scanlon et al. find new hires and plan switchers are more responsive to quality measures as well as price. Using the same study population, Chernew et al report “no significant evidence of heterogeneity” in the valuation of plan

2b.i. Employer Choice of Healthplans Research on how employers make decisions regarding which plans to offer, and how many, is limited by comparison. We focus here on empirical analyses of plan offerings, as opposed to analyses of surveys that ask employers to report what factors affect their decisions (e.g. Rosenthal et al. 2007). The most relevant papers for our purposes include Bundorf (2002) and Chernew et al (2004). These papers focus on whether employers' decisions reflect the *assumed* needs of their employees. For example, Chernew et al (2004) uses data on the HMO plans offered by 17 large employers in 2000 to see whether CAHPS scores affect the propensity any given plan is offered; they find that employers are more likely to offer plans with strong absolute and relative CAHPS performance measures. In related work, Bundorf (2002) finds employers' offerings correlate with employee characteristics. For example, firms whose employees have greater variation in healthcare expenditures are more likely to offer a choice of plans.

Our project builds on this research by quantifying – in dollars - the loss to consumers associated with restricted choice, and comparing these estimated losses to premium increases likely to occur if employees are free to apply their employer subsidies to other plans offered in their marketplace.

attributes based on observable or unobservable employee characteristics. Evidence from a different population – namely Medicare enrollees – is mixed as well. Using enrollment data surrounding the release of Medicare HMO report cards in 2000 and 2001, Dafny and Dranove (2008) find no differences in responses by demographic characteristics at the county level, but stronger evidence of non-report-card-related learning about quality (“market-based learning”) in counties with greater HMO penetration, more private report card data, and more stable populations.

3. Data

We use a proprietary panel database on healthplans offered by a sample of large, multi-site employers from 1998-2006. The dataset, which we call the “Large Employer Health Insurance Dataset” (LEHID), was provided by a major benefits consulting firm which assists employers with designing or purchasing benefits from health insurers.⁶ The unit of observation is the plan-year. A plan is defined as a unique combination of an employer, geographic market, insurance carrier and plan “type” (HMO, POS, PPO and indemnity), e.g. Company X’s Chicago-area Aetna HMO. The full dataset contains information from 813 employers and 139 geographic markets in the United States. The markets are defined by the data source and typically delineate metropolitan areas and ex-metropolitan areas within the same state, e.g. Arkansas – Little Rock and Arkansas – except Little Rock.⁷ The number of enrollees covered in the data averages 4.7 million per year. Given an average family size above 2, this implies more than 10 million Americans are represented in the sample in a typical year. After excluding observations with missing or problematic data⁸, the sample contains 811 employers, 139 markets and 356 carriers. Most employers are active in a large number of markets (45 for the median employer-year). Descriptive statistics are set out in Table 1. For additional details of the data, see Dafny (2009).

⁶ Using a survey of 21,545 private employers, Marquis and Long (2000) find that external consultants were employed by nearly half of the smallest firms (<25 workers) and nearly two-thirds of the largest firms (>500 workers). This suggests that the results of our study will be generalizable beyond our specific sample.

⁷ Dafny (2009) includes a map of the geographic markets, which occasionally span state lines.

⁸ We drop 347 observations with a missing industry code, 2752 observations associated with employer-market-years in which the employee share of premiums for one or more plans is negative, and 304 observations with missing data. We also consolidate the four plans that appear twice in the data because the employer self-insures some enrollees and fully-insures others.

Premium is the average annual charge, normalized to year 2000 dollars using the CPI, per person-equivalent covered by a plan.⁹ It combines employer and employee contributions. The definition of premium depends on whether a plan is self-insured or fully insured. Many large employers choose to self-insure, outsourcing benefits management and claims administration but paying realized costs of care. Such employers can spread risk across large pools of enrollees, and often purchase stop-loss insurance to limit their exposure. Per ERISA (the Employee Retirement Act of 1974), these plans are also exempt from state regulations and state insurance premium taxes, enabling firms to reduce their insurance costs and/or standardize plan benefits across multiple sites. Reported self-insured plan “premiums” are actually estimates of employers’ projected healthcare expenditures, including any administrative fees and stop-loss premiums.¹⁰

Demographic factor is a measure that captures the family size, age, and gender of enrollees in a given plan-year. It can be construed as the mean number of “person equivalents” per enrollee. *Plan design* captures the generosity of benefits for a particular plan-year, including the level of copayments required of enrollees. Both factors are calculated by the source, and the formulae were not disclosed to us.

Our empirical analyses use the employer-market-year as the unit of observation. If an employer appears in the sample in a given year, all healthplans it offers in any market are included in the data. However, the panel is unbalanced: on average, 240 employers appear in the sample each year. Of the unique employer-market pairs in the data, 46 percent appear only once, and 17 percent appear twice. We do not observe the total number of employees offered

⁹The original data reports the average premium per *enrollee*. Thus, this average premium is larger for employee groups whose enrollees cover more dependents. We follow the practice of our data source and divide this figure by *demographic factor* to obtain the premium “per effective enrollee.”

¹⁰ This definition of premiums for self-insured plans is common to all employer surveys, including the KFF/HRET survey described in Section 2.

insurance, hence our analyses are limited to employees who do take up coverage. As additional employees may take up coverage if more options are available, this likely understates the total gains from expanding choice.¹¹

Before moving to the empirical analysis, we present statistics on the state and evolution of choice within our sample and study period (Figure 3). As expected, choice is more common among employers in our sample than among the universe of employers sampled by the Kaiser/HRET survey, however nearly half of the employer-market-years offer access to only one plan. Over 75% offer at most two plans. Fifty percent of those offering two plans offer an HMO and a PPO; 14% offer a POS plan and a PPO. The figure also shows that the amount of choice offered has fallen over time and (consistent with the survey evidence) that PPO plans have increased in popularity while indemnity plans have become less popular.

4. Empirical strategy

We conduct our analysis in three steps. First we use our data on consumer choices of health plans conditional on the options offered by their employers to estimate a utility equation describing employee preferences for plan characteristics. Second, we estimate a hedonic equation that describes the relationship between the premiums we observe in the data and plan, employer and market characteristics. We use the coefficient estimates from this equation to predict the combined employer and employee premiums that employees in a given firm, market and year would face for every plan offered in their market and year, assuming large-group

¹¹ To estimate the share of employees who do not take up employer-sponsored insurance, we matched our data to total employment figures reported in the Compustat Financial Database. However, Compustat is limited to large, publicly-traded firms, substantially reducing our sample size. In addition, the employment figures are very noisy, particularly as some firms report employment for North America rather than for the United States. The implied mean enrollment rate across employer-years was 46%, much lower than the 67% reported by the Kaiser-HRET survey for large firms (200+ workers) offering health benefits in 2005. We concluded that the analyses using this matched data sample are less informative than those presented in the paper.

pricing prevails. Last, we use the results of both analyses to predict employee choices and expected utility under our counterfactual scenarios in which additional plans are made available on the same terms (i.e. a fixed percentage subsidy for a given set of employees, group rates and full tax-deductibility).

Although we are interested in the effect of expanding consumers' choice sets to encompass all possible options, the structure of our utility equation (which includes a logit error term with unbounded support) implies that adding *all* available plans to the choice set would over-estimate the welfare gains of choice. We therefore investigate three counterfactual scenarios. First, we maintain the same *number* of plans in the choice set for each employer-market-year, but we substitute the most preferred plans for those currently offered (that is, if the employer does not choose optimally for its employees). We call this the "plan swapping" scenario. Second, we assume that employees within each employer-market-year triple gain access to their preferred option within each of three plan types: HMO, POS and PPO plans.¹² (We exclude indemnity plans because they are rarely offered in our data. Employers already offering indemnity plans receive their most-preferred indemnity plan in the counterfactual to ensure a strictly expanded choice set.) We call this the "all plan types" scenario. Third, we make all plans in the market-year available to all employees (the "all plans" scenario). The changes in consumer surplus predicted by the "plan swapping" and "all plans" scenarios provide lower and upper bound estimates, respectively, of the value of greater choice, with the "all plan types" scenario falling in between.

¹² If an employer previously offered more than one option within a given plan type, we retain the same number of options within that plan type in the simulation.

4.a. Demand Model

The first step is to estimate a model of consumer demand for healthplans. We use a logit model, including in the consumer's choice set only the plans that are offered by the relevant employer in the relevant market and year. We denote a "plan" as a unique employer-market-carrier-plantype-year quintuple: indeed, this is the unit of observation for our data. Consumer i 's utility from plan $emcjt$ in year t is modeled as:

$$(1) \quad u_{imcjt} = \delta_{emcjt} + \varepsilon_{imcjt}$$

where $\delta_{emcjt} = x_{emcjt} v_{emt} \beta + p_{emcjt} v_{emt} \alpha + \xi_{emcjt}$: a linear combination of observed characteristics of the plan (denoted x), premium (p), observed characteristics of the employer (v), and an unobserved quality variable (ξ). The term ε_{imcjt} is consumer i 's idiosyncratic preference for carrier c and plantype j in market m at time t .

Before discussing the details of our estimation, we offer remarks on our use of a simple logit model rather than a nested logit or random coefficients model. The most intuitive nested logit model, in which the first nest is the choice of plan type (such as HMO or non-HMO) and the second is the choice of plans within type, requires eliminating most of the data because choice sets typically contain at most one of each plan type. A random coefficients model is computationally infeasible given the number of fixed effects included in the utility equation. As we discuss below, these fixed effects are extremely important for capturing quality differences, for example among insurance carriers within a given market. We attempt to capture heterogeneity among employees by including a large set of interactions between plan characteristics and observable characteristics of the relevant employee population. These permit

the coefficients on the key explanatory variables to differ across observably different groups of consumers.

Berry (1994) shows that the parameters in equation (1) can be estimated using the following linear equation, which explicitly lists all covariates:

$$\begin{aligned}
 (2) \quad \ln(s_{emcjt}) - \ln(s_{em0t}) = & \alpha + \xi_e + \nu_m + \psi_c + \eta_j + \delta_t + \zeta_{em} + \omega_{mc} + \varphi_{mt} + \chi_{mj} + \sum_i \lambda_{ijt} I(\text{industry}_e = i) \\
 & + \alpha_1 p_{emcjt} + \alpha_2 p_{emcjt} * \text{demographic factor}_{emcjt} + \sum \alpha_{3i} I(\text{industry} = i) * p_{emcjt} \\
 & + \sum_i \alpha_{4i} I(\text{industry} = i) * \text{demographic factor}_{mcejt} * p_{emcjt} \\
 & + \psi \text{plan design}_{emcjt} + \sum_i \mu_i I(\text{industry} = i) * \text{plan design}_{emcjt} \\
 & + [\pi \text{self-insured}_{emcjt}] + \xi_{emcjt}
 \end{aligned}$$

In equation (2), s_{emcjt} is the market share of plan $emcjt$ and s_{em0t} is the market share of the outside option in the relevant employer-market-year triple. We define the “outside option” to be the least generous plan in the employer-market-year triple, which implies normalizing its unobserved quality to zero. Other plans’ observed characteristics are measured relative to those of this baseline plan.¹³ For robustness, we also report results obtained when the outside option is the most frequently-offered plan in the relevant market-year.¹⁴

The covariates include several fixed effects, two continuous measures - plan design and the employee’s contribution to premiums (“price”) - and interactions including these two measures. We discuss each in turn.

The fixed effects include all of the “main effects,” that is, dummies for each employer, market, carrier, plan type, and year. However, the dummies for employer, market and year are differenced out when we normalize the characteristics of each plan with respect to those of the

¹³ The least generous plan is defined using plan type and premium. Indemnity plans are the most generous, followed by PPOs, POS plans and HMOs in that order. Within a particular plan type, the cheaper of a pair of plans is defined as the less generous plan.

¹⁴ In the case of a tie the most frequently-offered plan is designated as the plan with the largest number of enrollees in the market-year.

baseline plan, which implies we obtain coefficient estimates only on carrier and plan type dummies. Carriers and plan types with the largest coefficient estimates generate higher utility, *ceteris paribus*, for enrollees.

Our data afford us the opportunity to include second, third, and fourth-order fixed effects as well. Such terms have the advantage of enabling a better fit of the model, but there are four important disadvantages. First, they absorb variation in continuous regressors of interest such as price, leaving little to identify the coefficients on these measures. Second, many of these terms cannot be included in the counterfactual scenarios. For example, employer-carrier fixed effects would capture the mean utility of different carriers to employees of a particular firm. In a counterfactual that expands the choice set to include carriers *not* presently offered by that firm, it would not be possible to estimate the utility of the new options. Third, even if a coefficient could technically be estimated for a third or fourth-order interaction term, the number of observations identifying it would be small and therefore unlikely to yield a representative estimate valid for counterfactual simulations. Last, some of these terms raise endogeneity concerns. For example, employer-year interactions would capture the fixed utility associated with the set of plans offered by an employer in a given year, but this depends on choices currently on offer, and would presumably change when the choice set changes.

In recognition of these issues, we include a parsimonious set of second-order interaction terms that control for the most important unobservable correlates of utility while permitting estimation of our counterfactual scenarios.¹⁵ Two of five interactions we include will be “differenced out” in our specifications: employer-market fixed effects and market-year fixed

¹⁵ Of the 10 possible second-order interaction terms, we include five for reasons detailed in the text that follows. We exclude employer-year interactions due to endogeneity concerns, and employer-carrier and employer-plan type interactions because these are not compatible with our counterfactual simulations. Finally, carrier-plan type and carrier-year interactions are unlikely to be important determinants of unobserved quality. Indeed, we find our price coefficient is unaffected by excluding these terms.

effects. We mention them here to clarify the sources of identification for coefficients in the demand model, and because these terms appear (and are not differenced out) in the hedonic premium model we discuss below. Conceptually, employer-market interactions absorb time-invariant differences across specific sets of employees. For example they capture the fact that employees of a firm in some markets are particularly well-educated, have a particularly high income or are particularly risk averse and therefore place a high value on health insurance. They also absorb any fixed variation in price for a set of enrollees that may be correlated with time-invariant differences in risk profiles and demographics. The market-year fixed effects pick up market-specific shocks to utility such as a reduction in provider quality due to the closure of a hospital. Importantly, their inclusion implies that changes in market-level prices for plans do not identify the price coefficient. Rather, identification relies on changes in the relative prices of plans offered to employees in a given market. Because we also include plantype-year fixed effects interacted with 18 industry dummies, relative price changes that reflect general pricing trends that differ across plan types also do not identify the price coefficient¹⁶ In terms of the utility model, the plantype-year interactions capture broad changes in consumer preferences (such as the decline in popularity of HMOs in the 1990's) and in plan management (such as HMOs' decision to engage in less utilization review) over time, and the interactions with industry dummies increase the flexibility of the model while still providing estimates of all parameters needed for our counterfactual scenarios.

We include market-carrier fixed effects to capture the “fixed utility” associated with each market-carrier combination. For example, Blue Cross Blue Shield of Illinois may be an

¹⁶ The industry categories are: Chemicals, Consumer Products, Energy, Entertainment & Hospitality, Financial Services, Government & Education, Health Care, Insurance, Manufacturing, Pharmaceuticals, Printing & Publishing, Professional Services, Retail, Technology, Telecommunications, Transportation, Utilities, and Unclassified.

especially attractive Blue Cross Blue Shield carrier because it has a very large network of hospitals. This interaction is therefore important to capture unobserved plan quality. Finally, we include market-plan type interactions to capture differences across markets in the utility associated with particular plan types. For example, HMOs are more highly-valued in areas where they have a long history and this may be important for demand.

In addition to these fixed effects, we include two continuous measures: *plan design* and the employee's contribution to the annual premium (hereafter "price"), denoted p_{emcjt} . We interact both with dummies for industry categories to incorporate potentially different valuations of these characteristics by employee populations in different industries. We also include interactions between our measure of family size (demographic factor) and price. Finally we interact both price and the price-demographic factor interaction with industry category dummies. This functional form exploits the richness of the dataset, allowing, for example, Firm X's employees in Industry Y to have less price sensitivity than Firm Z's employees in Industry Y due to their larger hypothetical family size..

Our model takes price to be exogenous to unobserved plan quality, conditional on the many covariates included. The rich set of fixed effects and interaction terms we include mitigates concerns about endogeneity, specifically that price will be positively correlated with unobserved quality, yielding a downward-biased coefficient estimate. For example, unobserved quality of a particular carrier is absorbed in the carrier fixed effects and the market-carrier interactions. Unobserved differences in quality across types of plan are absorbed in the plan type variable, the market-plan type interactions and the plan type-year-industry category interactions. We considered several instruments, including for example the average price of plans offered by the same employer-year in different markets and different plan types, but conditional on all of the

fixed effects in our model there is insufficient variation in these potential instruments to predict the remaining variation in price. As we discuss in the next section, our estimates of price elasticity fall in the middle of the range of estimates from other studies of healthplan choice.

Last, we estimate models with and without an indicator for whether plan j is self-insured. Although self-insurance primarily affects the way in which employers finance their plans – and an indicator for it is therefore included in the hedonic premium model – it is possible that self-insured products are observably different to consumers on dimensions other than price (for example, because they may exclude state-mandated benefits, or because employers may perform some administrative functions for these plans).

4.a.i. Demand Results

The demand estimates are summarized in Table 2. Columns 1 and 2 display results for the model using the least generous plan as the outside option (“LG Model”), excluding and then including an indicator for self-insured plans, respectively. Columns 3 and 4 present the same for the model using the most frequent plan as the outside option (“MF model”). The coefficients for price, the price-demographic factor interaction, and plan design differ across industries because all three are interacted with industry category dummies. We display the estimates for the manufacturing industry, the largest in the data.

We begin by observing that estimates on all coefficients of interest are fairly similar between specifications with and without the self-insurance dummy. The coefficient on this dummy is positive and significant, implying that consumers prefer self-insured plans. As expected, we also find that plan design has a significant positive effect on utility, and its estimated coefficient does not vary much across specifications.

The interaction of price with the demographic factor makes the price coefficients difficult to interpret from the simplest table of results. The mean demographic factor for the manufacturing industry, together with the implied price coefficient for each specification, is provided in the second panel of the table. The price coefficient is statistically significant and nearly identical in magnitude across all specifications.

Table 3 reports the implied price coefficients, together with estimated price elasticities, for the seven largest industries in the data. The price coefficients are negative and significant at $p=0.05$ for all industries and specifications. The elasticities in the LG model vary from -0.07 in the financial industry to -0.47 in the retail industry. The ranking of elasticities by industry is intuitive: in general there are higher elasticities for lower-wage industries. There is some variation in estimates across the LG and MF models, but the ranking of elasticities is similar for most industries. In all cases the elasticities are within the range estimated in the previous literature. Below, we present simulations using demand estimates from both models; these do not reveal meaningful differences in the distribution of estimated utility gains overall.

4.b. Hedonic Equation

We use a hedonic regression model to predict the price at which each plan will be made available to the population in a particular employer-market-year in our simulations. Simply using the average of observed premiums for each plan is undesirable because premiums vary with the composition of the relevant employee population. It is worth noting that we do not expect our estimates to approximate the price that would prevail on an “exchange” for individually-purchased plans; the reduction in group size implied by individual shopping may lead to a substantial price increase, a subject we address in Section 6. Instead we use our predicted prices to estimate the consumer surplus increase from expanding choice, *ceteris paribus* (that is, with

continued price-setting at the employer-market-year level). This model implicitly assumes that all buyers are treated similarly. For example if insurance carrier A’s HMO carries a ten percent premium relative to insurance carrier B’s HMO then all aspiring enrollees will also face a ten percent premium for this plan (they may also face a price increase or reduction due to the characteristics of their employer group and market).

Our model takes the following form:

$$\begin{aligned}
 (3) \quad \ln(\text{premium})_{emcjt} &= \alpha + \xi_e + \mathbf{v}_m + \boldsymbol{\psi}_c + \boldsymbol{\eta}_j + \boldsymbol{\delta}_t + \boldsymbol{\varsigma}_{em} + \omega_{mc} + \varphi_{mt} + \chi_{mj} + \kappa_{jt} \\
 &\quad + \eta \text{plan design}_{emcjt} + \sum_i \varphi_i I(\text{industry} = i) * \text{plan design}_{emcjt} \\
 &\quad + \gamma \text{self-insured}_{emcjt} + \varepsilon_{emcjt}
 \end{aligned}$$

We regress log premium per effective enrollee (combined employer and employee premium contributions) on plan design (interacted with industry dummies), a self-insurance indicator, and the same first and second-order fixed effects included in the utility equation.¹⁷

We anticipate a negative coefficient estimate on the self-insurance dummy. Self-insurance should be cheaper, *ceteris paribus*, because the employer bears some (or all) of the risk of medical expenditures, self-insured plans are exempt from state mandates and premium taxes, and employers often fulfill many of the administrative functions that insurers perform for fully-insured plans (such as explaining benefit coverage to enrollees). We expect the employer-market interactions to be particularly important because they capture unobserved demographic information that is likely to affect health risk and therefore the cost of insurance.

¹⁷ Note that, compared to the utility model, the hedonic model excludes interactions between industry dummies and the plan type* year interactions; this omission is intended to reduce “overfitting” of the data, which could result in misleading predictions of premiums.

4.b.i. Hedonic Results

The results of the hedonic regression are summarized in Table 4. As a measure of the fit of our model, Panel A describes the distribution of the ratios of the regression residuals to the actual premiums. The fit is good: the fifth percentile of this distribution is -0.39 and the 95th percentile is 0.27. That is, the smallest residuals are roughly 40% of premiums and the largest residuals are roughly a quarter of the premiums. The adjusted R-squared of the regression is 0.792.

The discussion thus far pertains to goodness of fit of the regression *within* sample. However, we are interested in predicting premiums out of sample; goodness of fit for this purpose is illustrated by Panels B and C. Column 1 of Panel B gives the distribution of predicted premiums for all (hypothetical and observed) employer-plan combinations.¹⁸ This distribution compares very favorably to the distribution of observed premiums, reported in column 3. For example, the mean predicted premium is \$2460, compared to an observed mean premium of \$2437. However, to ensure that our simulations are not overly sensitive to outliers, we take the extra precaution of censoring predicted premiums at the 5 percent tails before performing our simulations.¹⁹ The distribution of censored premiums is given in column 2.

Our final summary of the predictions implied by the hedonic model is given in Panel C, which presents the distribution of a statistic we term the “span ratio.” The span ratio equals the difference between the largest and smallest predicted premiums, divided by the mean predicted

¹⁸By construction, the number of observations is very large: the average market-year has 15 carrier-plan type combinations offered by at least 3 employers. Given there are 115,440 employer-market-year units the total number of observations exceeds 1.7 million.

¹⁹ Premiums in the low and high tails are replaced by the 5th percentile and 95th percentile of premiums within the relevant market-year, respectively.

premium, for a given set of observations.²⁰ Columns 1 and 2 define this set as the employer-market-year, providing a snapshot of the range of premiums from which an employee *can* choose given their current set of options (column 1, which only includes employer-market-years in which more than one option is available) or *could* choose if all plans were made available (column 2). The median figure in column 2 is 19 percent, as compared to 9 percent in column 1, implying that in a market with full choice, employees would have a wider range of price points from which to choose. Of course, even the 9 percent figure overstates the current range of price points as only 55 percent of employer-market-years offer any choice at all. We note that “span” is defined for combined employer and employee premiums; the span of employee contributions may certainly differ.

Columns 3 and 4 of Panel C also report span ratios calculated using the set of observed plans and “all plans,” respectively, but here the set of underlying observations is grouped by market-carrier-plan type-year. Thus, these columns illustrate the variation in premiums for, say, the Aetna POS plan in Chicago in 2003, due to employer-specific characteristics (apart from family size, age, and gender, which are already accounted for as premiums are reported per effective enrollee). The median span ratios are 28 percent (using actual plans on offer, and associated predicted premiums) and 44 percent (all plans, predicted premiums). The sizeable spans are not surprising: the risk profiles of employee populations are very different, and premiums are experience-rated for large groups.. As expected, the span ratio in the “all plans” scenario is larger even within the same market-carrier-plan type and year, as we have expanded the range of employee groups for which each product is available.

²⁰ We report all span ratios using predicted, rather than actual, premiums as our simulation results use predicted premiums to estimate both current and predicted utility under the various scenarios.

In the interest of space we do not report the coefficient estimates from the hedonic model, but we note here that the sign of the coefficient estimate on the self-insured dummy is positive, contrary to expectations. Though statistically significant, the coefficient estimate of 0.005 is economically small: a self-insured plan typically costs 0.5% more than a fully-insured plan, *ceteris paribus*. Together with the estimates from the demand model, in which self-insured plans were found to be more attractive to consumers all else equal, this implies any cost savings associated with self-insurance may be being passed on to employees in the form of higher quality

4.c. Simulations

The next step is to use the estimated coefficients from the demand models to predict employee choices and the resulting consumer surplus if employees are permitted to select among a wider set of healthplans than that offered by their employers, and premiums for these plans are estimated using the hedonic model.

As noted earlier, because our utility equation includes a logit error term which has unbounded support, expanding the choice set to include *all* observed healthplans in each market-year will overestimate the value to consumers of increased choice. Thus we also provide estimates of the “plan swapping” scenario that holds constant the *number* of choices available to each set of employees, but substitutes the most-preferred plans for those currently offered (where the most-preferred are defined as the plans that are estimated to generate the most utility, on average, for employees within the relevant employer-market-year), as well as the “all plan types” scenario that includes access to the most-preferred option within each of the HMO, POS, and PPO plan types. We define a health insurance plan or “option” as an MCPY combination, for example United HealthCare’s Chicago-based PPO in 2005, and we exclude from the counterfactuals plans that are offered by fewer than three employers in the relevant market-

year.²¹ We also ensure that our estimates are conservative by excluding plans whose predicted average utility is below the fifth or above the ninety-fifth percentile of the estimated utility distribution.²² Finally, we drop the small share of employer-market-years for which the estimated price coefficient is positive; the exact percentages are reported in Table 3 and vary depending on the demand specification.

To measure consumer surplus, we use the approach delineated by Nevo (2001) and based on McFadden (1981). Consumer i 's expected gain from a change in the set of healthplans available to him is:

$$(4) \quad \Delta_i = u_i^t - u_i^{t-1}$$

where u_i^t and u_i^{t-1} are define by:

$$(5) \quad u_i^t = E_\varepsilon \max_j (\delta_{jt} + \varepsilon_{ijt}).$$

Note that this is the expected welfare gain from the perspective of the econometrician given the available data. A dollar-valued measure can be obtained using the method suggested by Hicks (1939) to create the equivalent variation (EV). The EV is the change in consumer wealth that would be equivalent to the change in consumer welfare due to the modification in the healthplan choice set. McFadden (1981) shows that:

$$(6) \quad EV_{it} = (u_i^t - u_i^{t-1}) / \alpha$$

²¹ These plans may be offered by carriers who are not truly active in a market, but who rent the networks of active participants in order to provide service in markets where they are not otherwise present. In the counterfactual simulations we allow employer-market-years where these small products are currently offered to keep them; that is, employees never lose MCPY combinations that are currently offered to them.

²² Specifically, plans added in counterfactual scenarios may not fall in the 5 percent tails of the utility distribution for the relevant market and year. To construct this distribution, we calculate the weighted average utility for each plan across all employer-market-year observations. Any plans falling at either extreme of this distribution within the relevant market and year are not included in the choice set for any counterfactual, unless such plans were offered in the original choice set. This reduces the influence of outliers on our estimated surplus gains.

where α is the coefficient on price in the plan utility equation. Integrating analytically over the extreme value distribution of ε implies that the equivalent variation of consumers in employer e , market m and year t is given by:

$$(7) \quad EV_{emt} = (1/\alpha_{emt})[\ln \sum_{c,j \in J_{emt,cft}} \exp(\delta_{emcjt,cft}) - \ln \sum_{c,j \in J_{emt,obs}} \exp(\delta_{emcjt,obs})]$$

where $J_{emt,obs}$ and $J_{emt,cft}$ are the choice sets available to employees of firm e in the observed and counterfactual scenarios respectively and $\delta_{ejmt,obs}$ and $\delta_{ejmt,cft}$ are the values predicted by the demand model.

The inputs to $\delta_{emcjt,obs}$ and $\delta_{emcjt,cft}$ are price, plan design, the fixed effects from the demand model and the unobserved quality ξ_{emcjt} . To calculate $\delta_{emcjt,cft}$ for plans that were originally offered by more than one employer we use the weighted (by number of enrollees) mean plan design and the median unobserved quality (ξ_{emcjt}) for the relevant market-carrier-plan type-year (MCPY). Using the median ξ_{emcjt} reduces the noise caused by particularly high or low estimated unobserved quality for particular employers, although in practice this has very little effect on our estimates.²³

The premiums in both the observed and the counterfactual scenarios are the values predicted by the hedonic regression described above. We incorporate employer subsidies to health insurance by assuming budget neutrality for employers: for every employer-market-year, spending equals actual employer contributions to premiums in the original data. Given most employers in our data appear to pay the same percentage of premium across plans (the median is 79%) rather than the same dollar amount, we retain this feature in our simulations. Thus, in each

²³ There is one exception to these rules. If an employer offers a plan to its employees in the data, in the counterfactual those employees are offered exactly the same plan design, with the same unobserved quality, even if it is offered by fewer than three employers and even if there are multiple plans in the MCPY. This design results in strictly superior choice sets in the counterfactual scenarios.

scenario we solve for the fixed percentage subsidy that yields budget neutrality, and apply this subsidy to plan premiums to determine the price faced by the employee.²⁴

5. Simulation Results

The three panels in Figure 4 summarize the utility gains from the three counterfactuals: plan swapping (4A), all plan types (4B), and all plans (4C). Each panel includes four boxplots that present the distribution of utility changes corresponding to the four distinct demand equations described earlier. The boxes are bounded by the 25th and 75th percentiles of the relevant distribution, and the ends of the vertical lines define the 5th and 95th percentiles.

Every model generates weakly positive changes in consumer surplus. The first plot in each of Figures 4A, B and C uses the demand parameters estimated from the LG model excluding the self-insured indicator. As expected, the plan swapping scenario (Figure 4A) yields lower estimated gains than the all plan types scenario (Figure 4B): \$504 for the median covered person as compared to \$970. The gains from the all plans scenario (Figure 4C) are the highest at \$2,045 for the median covered person. The second plot in each figure is generated using demand estimates from the LG model *with* the self-insured indicator. The simulations represented by these plots assume self-insurance is no longer possible, so any utility benefits associated with self-insurance are eliminated. The plots suggest these benefits are non-trivial. The median gain per covered person declines to \$265, \$728 and \$1,841 for the plan swapping, all plan types and all plans scenarios, respectively.

Although the exact sources of the utility associated with self-insurance cannot be ascertained, models that include an indicator for self-insurance likely overstate its importance

²⁴ This requires us to solve for the fixed point of an equation, since the share of a given plan depends on how much it is subsidized, which in turn depends on how many consumers are choosing that plan.

because self-insurance is correlated with other plan characteristics. For example, Aetna and Blue Cross/Blue Shield plans are far more likely to be self-insured than fully-insured, so a self-insurance indicator will absorb some of the quality associated with plans offered by these carriers. For these reasons, our preferred simulations rely on the demand parameters from the specification excluding the self-insurance indicator.

The second pair of plots in each figure correspond to the first pair, however in this case the demand estimates are derived from the MF model. In every case the estimated benefits are smaller than those simulated using the LG model, but the differences are not great. For example in our preferred LG simulation (using the most conservative scenario - plan swapping – and avoiding overstating the importance of self-insurance by excluding the self-insured dummy from the demand equation) the median benefit declines from \$504 to \$417 per person per year. We take this as evidence that the choice of baseline does not qualitatively affect our results.

6. Discussion

Our findings reveal that, on average, restricting employee choice yields substantial amounts of deadweight loss. This loss is due both to poor matching between employees' preferred plans and employers' offerings, holding the absolute number of choices constant, and to the reduced variety of plans that are offered. In this section we assess whether premium increases that might occur in an expanded-choice scenario are likely to fully offset projected surplus gains.

Our estimates of consumer surplus are calculated under the unrealistic assumption that individuals would enjoy group pricing when choice is expanded. For reasons we detail below, premiums are likely to rise if employer involvement in plan sponsorship is limited to a subsidy.

We begin by presenting data on the amount by which premiums would have to increase to fully offset the gains from choice. We express this figure as a percentage of the average predicted premium for each employer-market year, and present boxplots of the resulting distribution in Figure 5. As in Figure 4, Panel A corresponds to the “plan swapping” scenario, Panel B to the “all plan types” scenario, and Panel C to the “all plans” scenario.

As expected, the numbers reflect the gains reported in Figure 4. The largest figures are associated with the “all plans”/ LG/no self-insured dummy simulation (87 percent of premiums for the median employer-market-year), and the smallest with the “plan swapping”/MF/self-insured dummy simulation (10 percent of premiums for the median employer-market-year). The median increase in premiums needed to offset surplus gains in our preferred (conservative) simulation, defined by the plan-swapping scenario and the demand parameters from the LG model without a self-insurance indicator, is 21 percent.

Of course, to interpret these results we require projections regarding the likely premium increase when the choice set is expanded. We discuss both current differences in loading fees for individual/small group versus large group plans, and projected loading fees for plans to be offered on a hypothetical “individual exchange”, as reported by organizations performing evaluations of recent healthcare reform proposals. As the loading fee represents the difference between dollars paid in as premiums and dollars paid out to providers of healthcare services, we implicitly abstract away from absolute changes in medical spending that may result when the same set of individuals enrolls in different plans.

According to the National Health Expenditure Accounts, which produces estimates of private premiums and insurer outlays on an annual basis, loading fees increased from 10.5

percent of premiums in 1998 to 12.8 percent of premiums in 2006.²⁵ These figures include self and fully-insured plans of all sizes. Other sources report similar aggregate estimates.²⁶ Loading fees can be divided into administrative and non-administrative components, although the categorization of expenses is somewhat arbitrary.²⁷ Non-administrative components include corporate taxes, profits, and additions to capital reserves. Notably, self-insured plans are exempt from premium taxes and solvency requirements, so non-administrative costs for individual and small group plans - which are fully-insured - are clearly higher.

Administrative costs are also higher for small plans. According to a white paper by the American Society of Actuaries (2009), there are four key components to administrative costs of healthplans: marketing (including broker commissions), provider and medical management (e.g. provider network management), account and member administration (includes billing, customer service, and claims processing), and corporate services (including underwriting and associated risk premiums). All but the second of these components will clearly be higher for small plans.

There are few sources that compare the loading fees for large group vs. individual/small group plans, and none to our knowledge that separate these estimates by expense category. Before discussing available figures, we note three reasons the difference in current loading fees is likely to overstate premium increases in our hypothetical scenario: (1) the risk premium due to adverse selection in the current individual marketplace exceeds that which would likely prevail if all employer-sponsored enrollees were included in the pool of insureds; (2) as the pool grows,

²⁵ Source: Table 12, <http://www.cms.hhs.gov/NationalHealthExpendData/downloads/tables.pdf>

²⁶ The Congressional Budget office puts the figure at 11 percent. The Sherlock Company, a health care financing firm, reports the median BCBS plan spends 10.4 percent of premiums on administration, and the Lewin Group, a health policy and management consulting firm, estimates the figure is 13.4 percent.

²⁷ For example, risk premiums are viewed as an administrative expense, while additions to capital reserves are not. The NHEA views premium taxes and profits as “non-administrative” expenses.

statutorily-required capital reserves should decline as a percent of premiums; (3) state premium tax rates should also decline as the taxable base increases.

These caveats notwithstanding, the best available estimate of the difference in loading fees between the smallest groups (fewer than 100 employees) and the largest (>1,000 employees) is 10 percent of premiums (Karaca-Mandic, Abraham, Phelps).²⁸ A 2006 study by the Council for Affordable Health Insurance reports a load difference between individual and large group policies of 17.5 percent, although the methodology is not provided. Finally, the Congressional Budget office reports that administrative costs range between 7 percent of premiums for firms with at least 1,000 employees to nearly 30 percent of premiums in the individual insurance market, yielding a maximal loading fee differential of 23 percent of premiums. Again, the source of these figures is not reported.

Evaluations of healthcare reform proposals constitute an alternative source of cost estimates for plans offered through an exchange. The Lewin Group estimated that administrative costs for an exchange with only private plans would be 10.7 percent if all workers are eligible to participate (as opposed to only small groups and individuals, as some reform proposals specify). However, the figures underlying these estimates are also not reported.²⁹

²⁸ These estimates are based on insurance plans selected by 6,115 individuals with employer-sponsored insurance from 2,842 different employers, who appear in the 1997, 1998, 1999, and 2001 linked Medical Expenditure Panel Survey-Insurance Component (MEPS-IC). The authors estimate two regression models – one to predict insurer payments on behalf of each individual, and a second to predict premiums at the individual level. In the latter model, they regress total (employer and employee) premiums on firm size dummies, projected insurer payments from the first model, interactions of firm size dummies and projected insurer payments, demographic and health status measures, employer characteristics, market covariates, and state and year dummies. The indicators for firm size are statistically significant, but the interactions with projected payment are not. Using the parameter estimates from both models, they calculate loading fees for a “typical” employer in each size category. Because healthcare utilization is underreported in the MEPS-IC, they believe their estimates of loading fees may be overstated. Thus, they also report figures that adjust expenditures to levels estimated by NHEA. These are the figures we use to construct loading fees as a percent of premium.

²⁹ “The Impact of House Health Reform Legislation on Coverage and Provider Incomes,” The Lewin Group Testimony before the Energy and Commerce Committee, U.S. House of Representatives, June 25, 2009. Downloaded 9/23/2009 from <http://www.lewin.com/content/publications/June25TestimonyUpdate.pdf>.

In sum, it appears the value of choice is rather likely to exceed its cost, at least for most employees and their dependents. Of course, our analysis overlooks some costs (such as the “shopping cost” to individuals), but omits some benefits as well (such as the reduction in costs to employers who currently do much of the shopping as well as administering of plans).

7. Conclusions

In its current incarnation, employer-sponsored insurance in the U.S. is characterized by a very limited choice, if any, for workers fortunate enough to be eligible to enroll. Our research makes use of a large panel of employer healthplan offerings and employee plan selections to quantify the surplus foregone as a result of restricted choice in the employer-sponsored system. By examining employees' choices among the set of plans they are offered, we obtain estimates of their preferences that enable us to identify their most preferred plans (and corresponding dollar-valued utility) from the entire set available in their marketplace. We find the median employee would be willing to forego roughly 27 percent of her subsidy for the right to apply the remainder to any plan she chooses.

In Dafny, Ho and Varela (2010), we analyze the distributional effects of expanding choice and explore the differences between plans that are offered to workers in our data and those they would select under expanded choice. Importantly, we do not find evidence that employers “overweight” premiums when making healthplan selections (that is, offer cheaper plans than employees would be willing to pay for), as surveys of employers suggest (e.g. Gabel et al. 1998, and Maxwell, Temin and Watts 2001). Our analyses indicate that employees would choose similarly-priced plans, but these plans would differ along other dimensions such as plan type and insurance carrier. We discuss the possible reasons for the (apparent) misalignment of employer and employee preferences.

Our estimates probably represent a lower bound on the value of increased choice. We observe only a subset of plans available in the market, and the U.S. experience with the introduction of Medicare Part D suggests that even more choice would become available in a subsidized individual exchange setting (Abaluck and Gruber 2009). (Of course, some of the plans currently provided by carriers to selected firms might disappear, particularly if adverse selection persists. However, most exchange proposals are accompanied by insurance reforms prohibiting selection tools such as exclusion of pre-existing conditions.) Our estimates also understate the benefits of choice because we do not model consumer heterogeneity within employer groups. This technical shortcoming precludes estimation of surplus gains associated with better matching of plans to individuals within a given employee group.

In addition to surplus gains from choice, reducing the reliance on employer-sponsored insurance would mitigate the labor market frictions arising from the lack of insurance portability between jobs and in or out of the labor force. In a thorough review of the empirical literature on this subject, Gruber and Madrian (2004) conclude there is compelling evidence of “job lock” associated with the current system, as well as other important effects on labor supply decisions.

While there is little political support for explicitly dismantling the employer-based system, reform proposals actively considered during 2009-2010 would all reduce the relative reliance on employer-sponsored insurance.³⁰ We find the value of choice is a nontrivial benefit of such a transition, and may more than offset the higher costs associated with an individual marketplace.

³⁰ Indeed, a bi-partisan proposal by Senators Ron Wyden (D-Oregon) and Robert Bennett (R-Utah) to do precisely that was widely lauded by health care analysts, but did not receive serious consideration.

References

1. Abaluck, Jason and Jonathan Gruber (2009), "Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program," *NBER Working Paper 14759*.
2. American Academy of Actuaries, (2009), "Critical Issues in Health Reform: Administrative Expenses," available at <http://www.actuary.org/health.asp>.
3. Beaulieu, N. D., 2002. "Quality Information and Consumer Health Plan Choices." *Journal of Health Economics*, Vol. 21, pp. 43-63.
4. Berry, S. 1994. "Estimating Discrete Choice Models of Product Differentiation", *The RAND Journal of Economics*, 25(2): 242-262.
5. Blumberg, L.J. and L.M. Nichols. 2002. "Worker Decisions to Purchase Health Insurance". *International Journal of Health Care Finance and Economics*, 1, 305-325.
6. Buchmueller, T.C. and P.J. Feldstein, 1997. "The Effect of Price on Switching Among Health Plans." *Journal of Health Economics*, Vol. 21, pp. 43-63.
7. Bundorf, M. Kate, 2002. "Employee Demand for Health Insurance and Employer Health Plan Choices," *Journal of Health Economics*, Vol. 21, pp. 65-88.
8. Carlin, C. and R. Town. 2009. "Adverse Selection, Welfare and Optimal Pricing of Employer-Sponsored Health Plans". U. Minnesota Working Paper.
9. Chernew, M., Frick, K. and C.G. McLaughlin. 1997. "The Demand for Health Insurance Coverage by Low-Income Workers: Can Reduced Premiums Achieve Full Coverage?". *Health Services Research*, v32(4): 453-470.
10. Chernew, M., Gowrisankaran, G., McLaughlin, C. and T. Gibson, 2004. "Quality and Employers' Choice of Health Plans," *Journal of Health Economics*, Vol. 23, pp. 471-492.

11. Chernew, M., Gowrisankaran, G., and D. Scanlon, 2006. "Learning and the Value of Information: Evidence from Health Plan Report Cards," *NBER Working Paper* 8589.
12. Cutler, D.M. and Reber, S. 1998. "Paying for Health Insurance: The Tradeoff between Competition and Adverse Selection". *Quarterly Journal of Economics*, 108(2): 433-466.
13. Cross, Margaret Ann, 2005. "Momentum Shifts Toward Consumer-Directed Plans," *Managed Care Magazine*, July.
14. Dafny, L. 2009, "Are Health Insurance Markets Competitive?", *American Economic Review*, forthcoming.
15. Dafny, L., Ho, K. and M. Varela. 2010. "Expanding Choice in Employer Sponsored Health Insurance: Who Gains and Why?", *American Economic Review (Papers and Proceedings)* forthcoming.
16. Fronstin, Paul, 2009, "Sources of Health Insurance and Characteristics of the Uninsured: Analysis of the March 2009 Current Population Survey," *Employee Benefit Research Institute Issue Brief No. 334*.
17. Gabel, Jon R., Kelly A. Hunt, and Kimberley Hurst, 1998, "When Employers Choose Health Plans: Do NCQA Accreditation and HEDIS Data Count?" *KPMG Peat Marwick, LLP* working paper.
18. Gruber, J. and B. Madrian, 2004, "Health Insurance, Labor Supply, and Job Mobility: A Critical Review of the Literature," in Catherine McLaughlin, ed. Health Policy and the Uninsured. Urban Institute Press.
19. Gruber, J. and E. Washington. 2005. "Subsidies to Employee Health Insurance Premiums and the Health Insurance Market". *Journal of Health Economics*, v24(2,Mar): 253-276.

20. Health Research and Educational Trust and Kaiser Family Foundation, “Employer Health Benefits 2005 Annual Survey”. Available at <http://www.kff.org/insurance/ehbs-archives.cfm>.
21. Hicks J. 1939. *Value and Capital*. Clarendon Press; Oxford.
22. Jin, Ginger Z. and Alan Sorensen, 2006. “Information and Consumer Choice: The Value of Publicized Health Plan Ratings,” *Journal of Health Economics*, Vol. 25 , pp. 248-275.
23. Marquis, M., M. Buntin, J. Escarce, K. Kapur, T. Louis and J. Yegian, 2006. “Consumer Decision-Making in the Individual Health Insurance Market”, *Health Affairs*, web exclusive, pp. W226-W234, available at <http://content.healthaffairs.org/cgi/content/abstract/hlthaff.25.w226v1?rss>
24. Matthews, Merrill. Medicare’s Hidden Administrative Costs: A Comparison of Medicare in the Private Sector (Alexandria, VA: Council for Affordable Health Insurance, 2006).
25. Maxwell, James, Peter Temin, and Corey Watts, 2001, “Corporate Health Care Purchasing Among Fortune 500 Firms,” *Health Affairs* May/June.
26. McFadden, D. 1981. “Econometric Models of Probabilistic Choice.” In *Structural Analysis of Discrete Data*, Manski C., McFadden D., (eds). MIT Press: Cambridge, MA; 198-272.
27. Nevo, A. 2001. “New products, quality changes and welfare measures computed from estimated demand systems.” NBER Working Paper 8425.
28. Royalty and Solomon. 1999. “Health Plan Choice: Price Elasticities in a Managed Competition Setting”. *Journal of Human Resources*, 34(1): 1-41.
29. Scanlon, D., Chernew, M., McLaughlin, C., and G. Solon, 2002. “The Impact of Health Plan Report Cards on Managed Care Enrollment.” *Journal of Health Economics*, Vol.21, pp. 19-41.

30. Swamy, P., 1970. "Efficient Inference in a Random Coefficient Regression Model",
Econometrica, 38(2): 311-323.
31. Wedig, G. J. and M. Tai-Seale, 2002. "The Effect of Report Cards on Consumer Choice in
the Health Insurance Market." *Journal of Health Economics*, Vol. 21, 1031-1048.

Figure 1: Distribution of Firms Providing a Choice of Health Plans, 2005

Source: Kaiser/HRET Survey of Employer-Sponsored Health Benefits

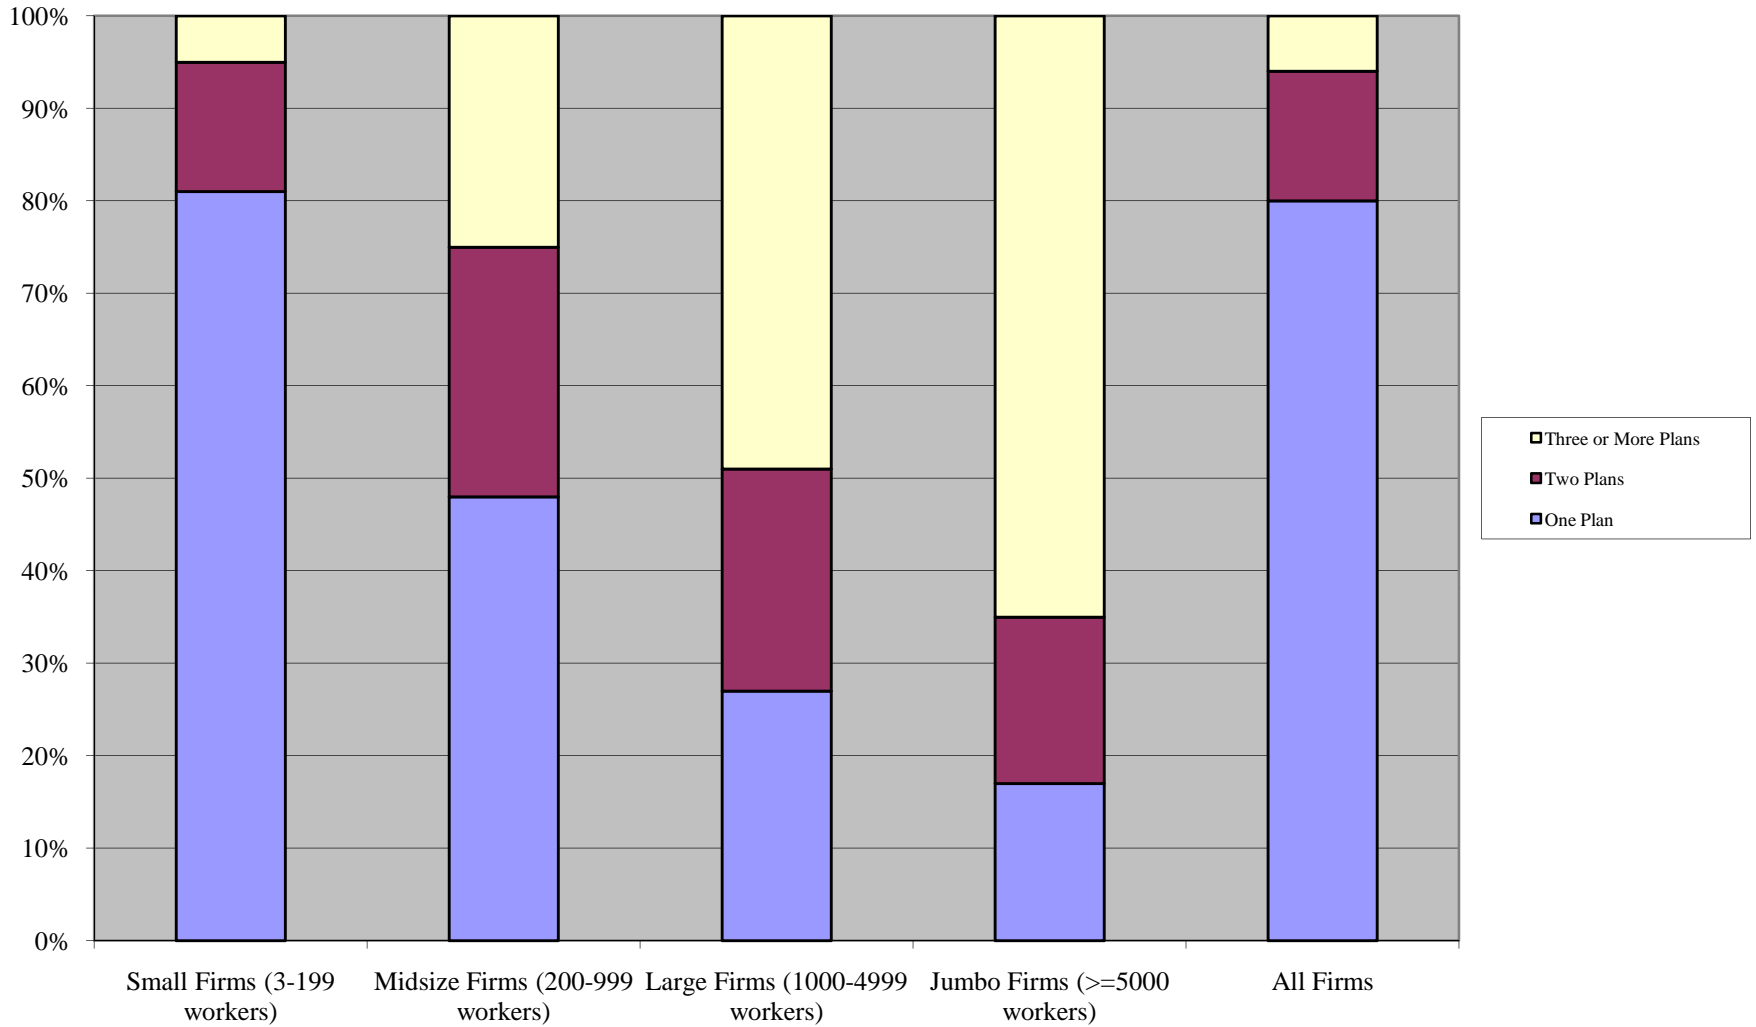


Figure 2: Percentage of Covered Workers with a Choice of Indemnity, HMO, PPO or POS Plans, (1996-2005)

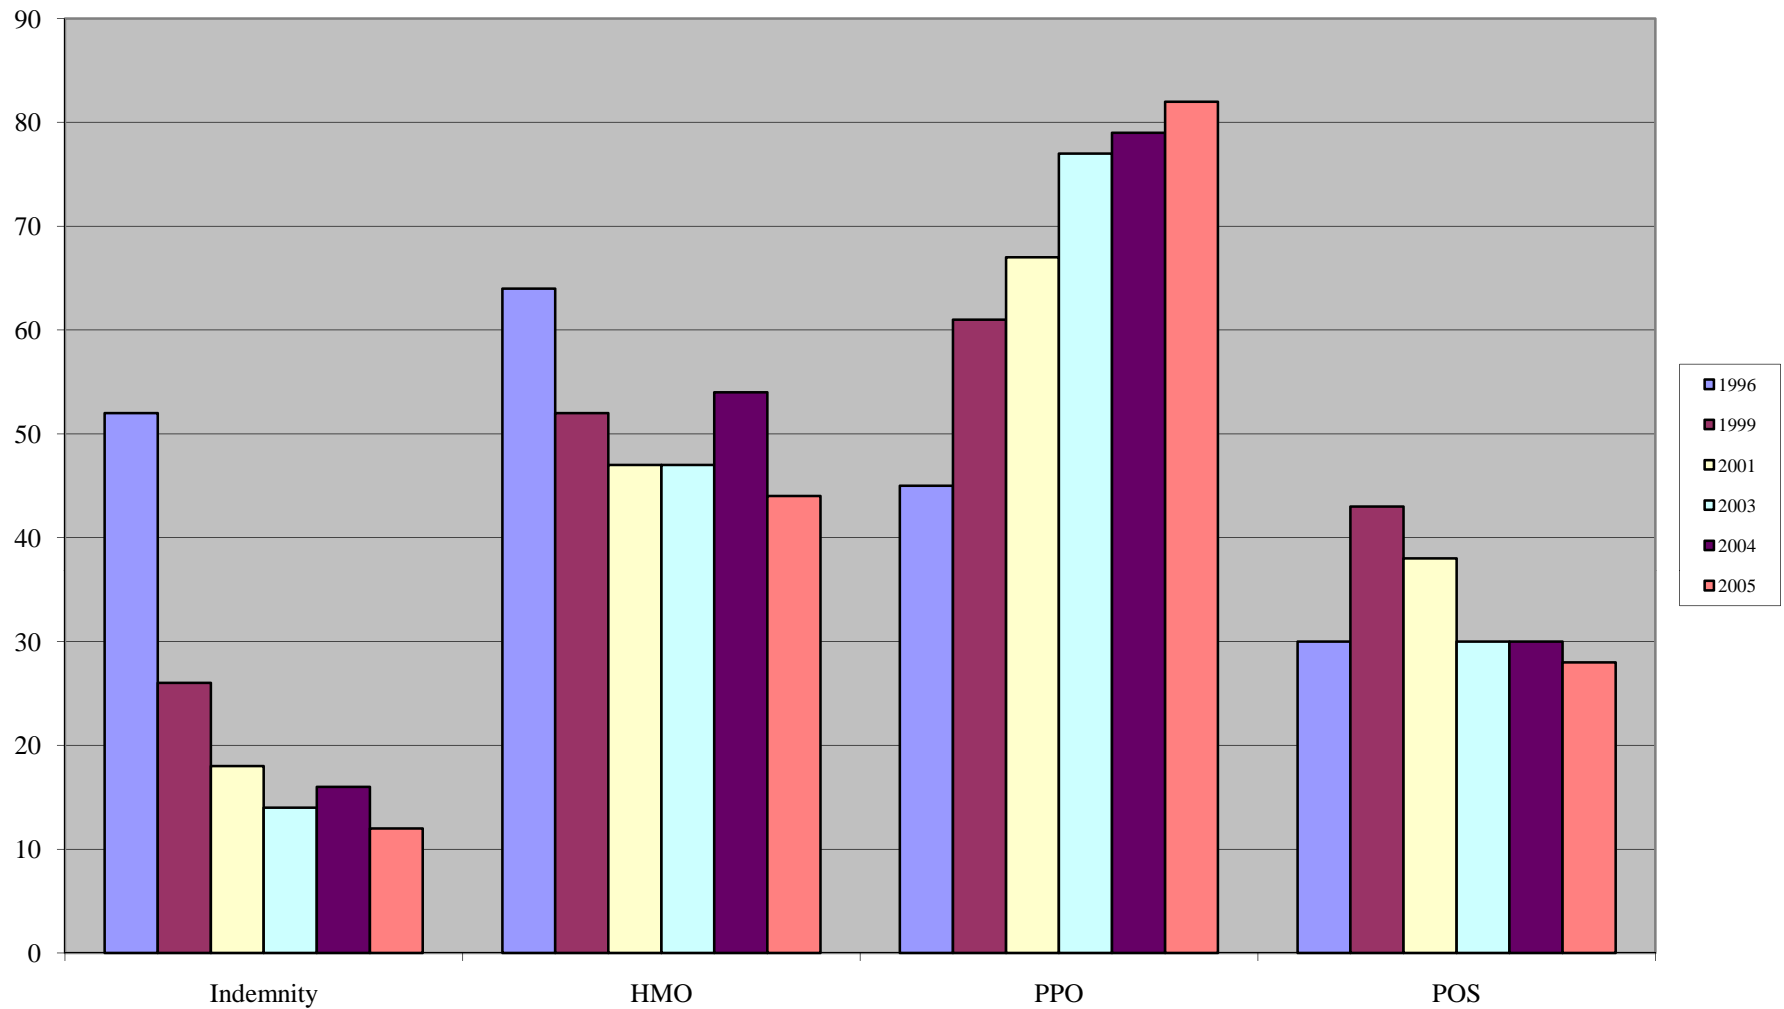


Figure 3: Observed Choice Sets

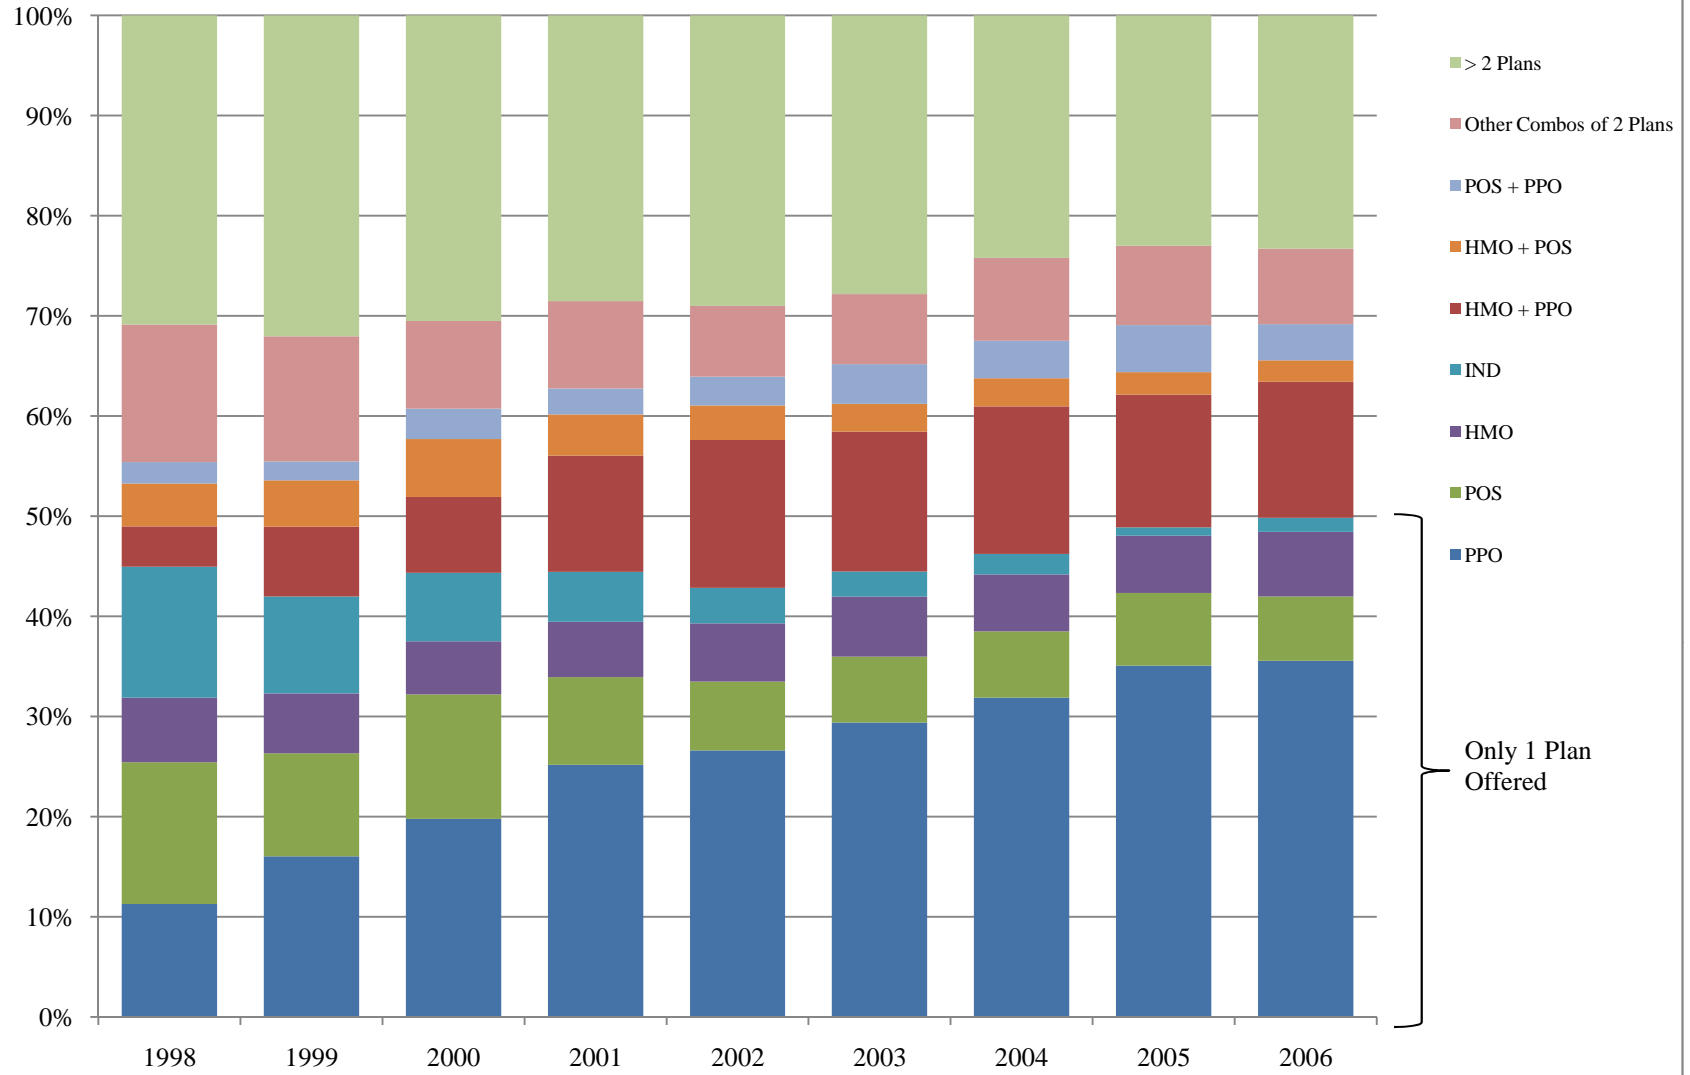


Figure 4A: Utility Gains From Choice Set Including Employees' Preferred Plan(s), Holding The Number Of Options Constant

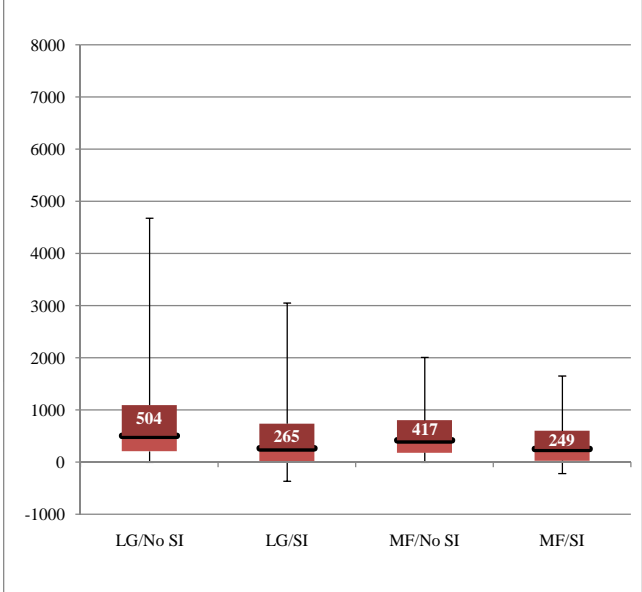


Figure 4B: Utility Gains From Choice Set Including Employees' Preferred Plans, Ensuring at Least 1 PPO, 1 HMO and 1 POS

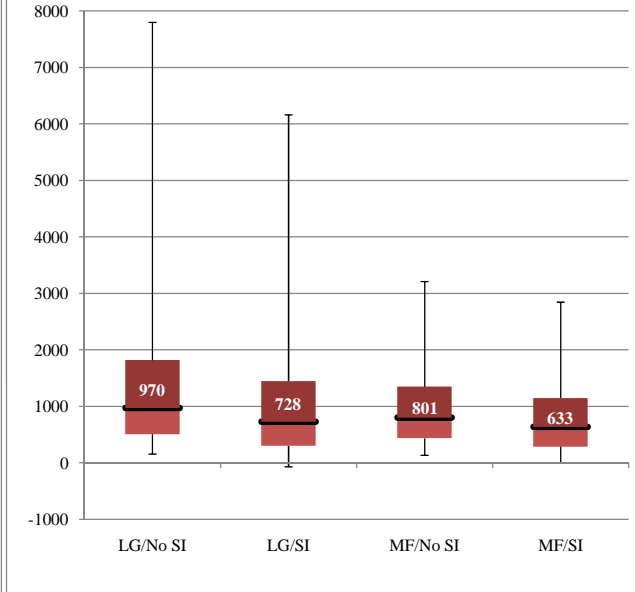


Figure 4C: Utility Gains From Including All Potential Plans in Choice Set

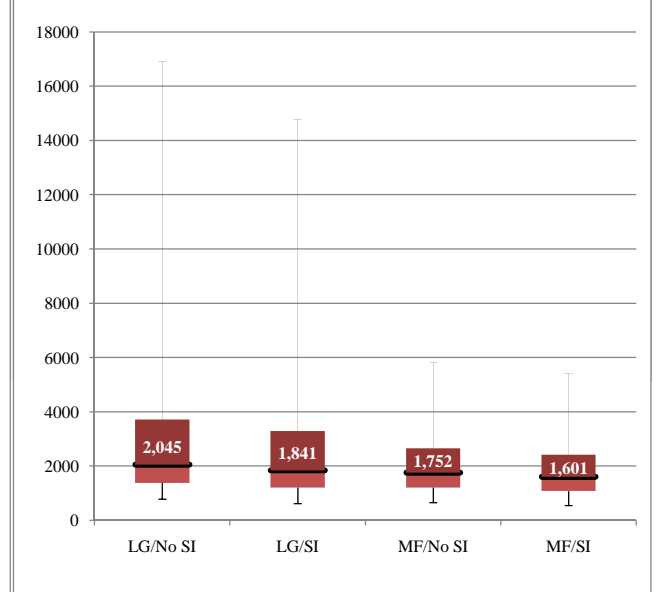


Figure 5A: Percentage Increase in Premiums Required to Offset Gains From Employees' Preferred Plan(s), Holding The Number Of Options Constant

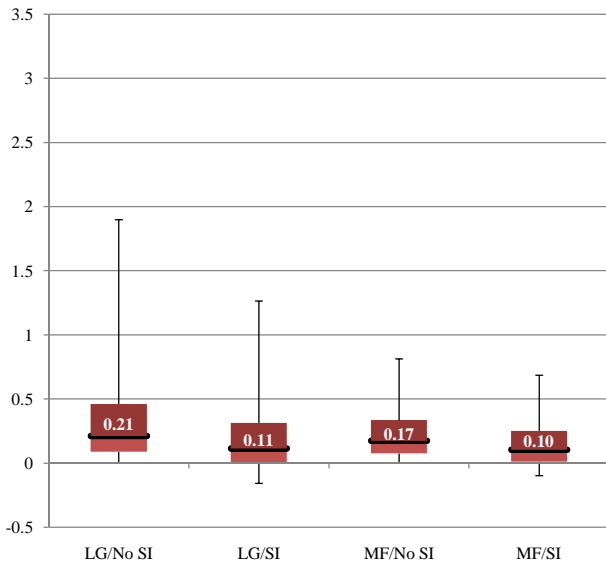


Figure 5B: Percentage Increase in Premiums Required to Offset Gains From Choice Set Including Employees' Preferred Plans, Ensuring at Least 1 PPO, 1 HMO and 1 POS

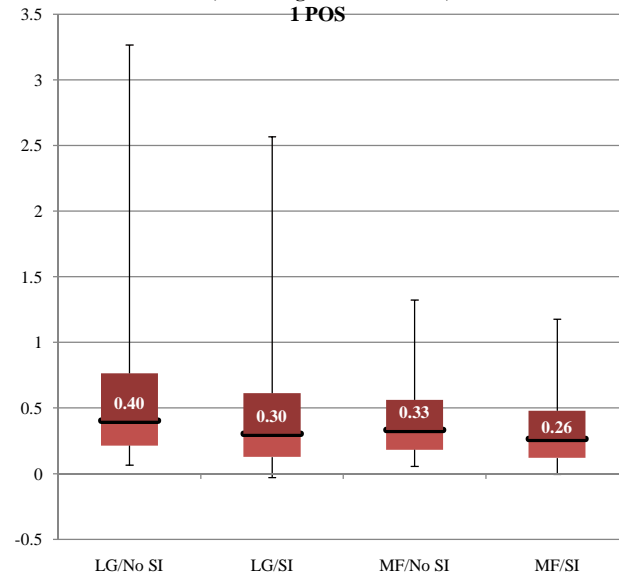


Figure 5C: Percentage Change In Premiums Required to Offset Gains When Including All Potential Plans in Choice Set

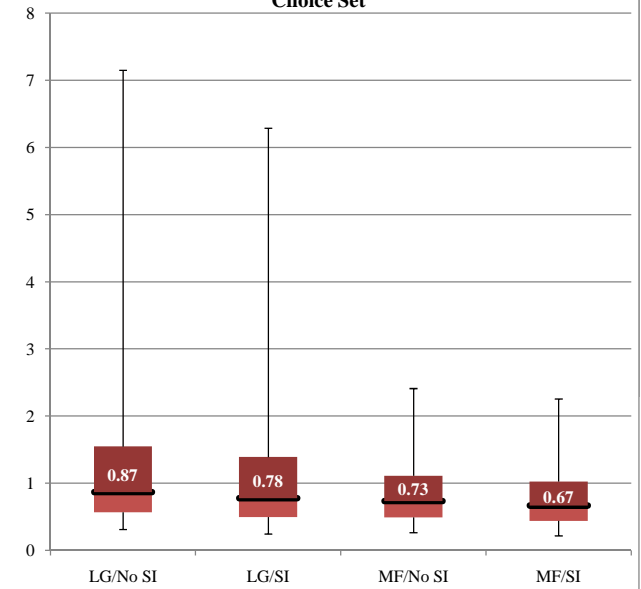


Table 1. Descriptive Statistics

Premium (\$)	2436 (704)	Plan Type (%)	
Employee Contribution (%)	0.212 (0.122)	HMO	37.8
Enrollment	175 (601)	POS	35.3
Demographic Factor	2.243 (0.449)	PPO	16.5
Plan Design	1.038 (0.087)	Indemnity	10.3
Self Insured (%)	0.682	Industry (%)	
Employer-Markets in the Panel		Manufacturing	13.4
Appear in every year	3%	Retail	12.0
Appear once	46%	Financial	11.4
Appear twice	17%	Technology	9.4
Appear thrice	13%	Consumer Products	7.3
Other	22%	Pharmaceuticals	6.4
Total # of EM's	39,633	Insurance	6.3
Number of employers	811	Transportation	5.6
Number of markets	139	Telecommunications	5.1
Number of carriers	356	Entertainment & Hospitality	5.1
Number of EMY	115,440	Professional Services	3.5
Number of Observations (EMCPY)	237,253	Health Care	3.4
		Energy Production/Transmission	2.9
		Printing & Publishing	2.8
		Utilities (Gas & Electric)	2.4
		Chemicals	1.7
		Government/Education	0.8
		Unclassified	0.3

Note: Industry breakdown percentages obtained using the employer-market-year as the unit of observation, as this is the unit of interest for the choice models.

Table 2. Demand Estimates
(Reported for Manufacturing Industry)

	<u>LG Model</u>		<u>MF Model</u>	
	<u>w/o SI</u>	<u>with SI</u>	<u>w/o SI</u>	<u>with SI</u>
	(I)	(II)	(III)	(IV)
Price (\$ 000s)	0.509 (0.394)	0.532 (0.393)	1.345 (0.395)	1.350 (0.394)
Price x DemoFactor	-0.769 (0.164)	-0.764 (0.163)	-1.139 (0.164)	-1.121 (0.164)
PlanDesign	0.945 (0.193)	0.951 (0.192)	1.349 (0.210)	1.385 (0.209)
Self Insured	NA	0.330 (0.011)	NA	0.279 (0.010)
Mean Demofactor for Manufacturing	2.44	2.44	2.44	2.44
Implied price coefficient for Manufacturing	-1.37	-1.33	-1.43	-1.38
Adj R-Sq	0.362	0.367	0.316	0.320

Notes: "LG Model" uses the least-generous plan within each employer-market-year as the outside option, while "MF Model" uses the most-frequently offered plan within a market-year as the outside option. Estimates in this table are for the manufacturing sector, as price, price x demofactor, and plan design are interacted with industry dummies. All implied price coefficients are statistically different from zero at $p \leq 0.05$. $N = 237,253$

Table 3. Average Implied Price Coefficients and Elasticities for Selected Industries

	<u>LG Model</u>		<u>MF Model</u>	
	<u>w/o SI</u>	<u>with SI</u>	<u>w/o SI</u>	<u>with SI</u>
	(I)	(II)	(III)	(IV)
Price Coefficients				
Manufacturing	-1.36	-1.33	-1.43	-1.38
Financial	-0.13	-0.15	-0.65	-0.66
Retail	-1.28	-1.23	-1.59	-1.55
Technology	-0.74	-0.72	-1.08	-1.06
Consumer Products	-1.22	-1.18	-1.42	-1.37
Telecommunications	-0.05	-0.44	-1.14	-1.11
Pharmaceuticals	-0.61	-0.58	-0.94	-0.90
Elasticities				
Manufacturing	-0.23	-0.22	-0.24	-0.23
Financial	-0.07	-0.08	-0.26	-0.27
Retail	-0.47	-0.44	-0.59	-0.57
Technology	-0.17	-0.16	-0.26	-0.26
Consumer Products	-0.19	-0.19	-0.23	-0.22
Telecommunications	-0.06	-0.05	-0.14	-0.14
Pharmaceuticals	-0.07	-0.07	-0.12	-0.12
% of E-M-Y with positive implied price coefficients	8.9%	8.7%	1.5%	1.5%
# of Employer-Market-Years included	105,214	105,435	113,709	113,700

Notes: "LG Model" uses the least-generous plan within each employer-market-year as the outside option, while "MF Model" uses the most-frequently offered plan within a market-year as the outside option. Estimates in this table are for the manufacturing sector, as price, price x demofactor, and plan design are interacted with industry dummies. All implied price coefficients are statistically different from zero at $p < 0.05$. Average elasticities across employer-market-year-plan combinations within each industry are reported. Elasticity is defined as price coefficient * (1 - share) * price.

Table 4. Hedonic Results

Panel A: Residual Ratio		Panel B: Predicted Premiums			Panel C: Span Ratio			
					<u>Grouped by E-M-Y</u>		<u>Grouped by M-C-P-Y</u>	
	(1)	(1)	(2)	(3)	(1)	(2)	(3)	(4)
		<i>Uncensored</i>	<i>Censored</i>	<i>Observed Data</i>	<i>Actual Choice Set</i>	<i>All Plans</i>	<i>Actual Choice Set</i>	<i>All Plans</i>
1%	-1.019	1,407	1,486	1,387	0.00	0.00	0.04	0.33
5%	-0.394	1,634	1,651	1,605	0.01	0.07	0.09	0.37
25%	-0.108	2,003	2,012	1,973	0.04	0.15	0.20	0.41
50%	0.000	2,373	2,370	2,341	0.08	0.19	0.28	0.44
75%	0.099	2,799	2,790	2,781	0.14	0.24	0.35	0.48
95%	0.274	3,655	3,585	3,650	0.28	0.35	0.45	0.56
99%	0.437	4,343	4,313	4,349	0.37	0.45	0.51	0.66
Smallest	-657	238	999	183	0.00	0.00	0.00	0.21
Largest	1394	13,785	5,814	12,710	1.28	1.60	0.76	1.04
Mean	-0.022	2,460	2,456	2,437	0.10	0.20	0.28	0.45
Std. Dev.	4.377	626	600	635	0.09	0.09	0.11	0.06
# Obs	237,253	1,957,013	1,957,013	237,253	63,125	115,376	18,797	18,797

Notes: Premiums are per effective enrollee.

Panel A: Residual Ratio = Residuals / Actual Premiums.

Panel B: Censored predicted premiums are censored at the 5% and 95% values within each market-year.

Panel C: Span Ratio = (Largest Premium in group - Smallest premium in group) / Average Premium in Group. All figures in Panel C use censored predicted premiums. E-M-Ys with only 1 plan are excluded from Panel C, column 1. M-C-P-Ys offered by fewer than 3 E-M-Ys are only included in Column 2 of Panel C.

Appendix Table 1A. Distributions of Utility Gains

	Plan Swapping				All Plan Types				Including All Plans in Choice Set			
	Base Plan: Least Generous		Base Plan: Most Frequent		Base Plan: Least Generous		Base Plan: Most Frequent		Base Plan: Least Generous		Base Plan: Most Frequent	
	w/o SI	with SI	w/o SI	with SI	w/o SI	with SI	w/o SI	with SI	w/o SI	with SI	w/o SI	with SI
1%	0	-2,079	0	-591	0	-481	0	-193	374	197	304	175
5%	0	-368	0	-218	157	-71	134	-11	776	618	656	543
10%	16	-228	16	-148	262	61	224	87	981	819	843	725
25%	208	0	180	26	507	300	436	285	1,381	1,207	1,210	1,081
50%	504	265	417	249	970	728	801	633	2,045	1,841	1,752	1,601
75%	1,094	738	803	599	1,817	1,447	1,348	1,146	3,711	3,290	2,652	2,415
90%	2,560	1,701	1,405	1,120	4,243	3,278	2,266	1,973	9,103	8,038	4,261	3,925
95%	4,678	3,051	2,007	1,652	7,802	6,162	3,209	2,845	16,903	14,785	5,821	5,432
99%	19,620	12,849	4,656	3,874	34,040	26,456	7,693	6,674	71,580	65,806	13,596	12,402
Mean	2,393	1,198	808	652	4,054	2,932	1,511	1,340	10,754	7,632	3,061	2,854
Std. Dev	67,194	29,851	10,477	30,888	98,608	54,140	31,267	54,536	662,118	153,031	58,814	83,815
Smallest	-60	-1,379,029	-68	-100,902	-177	-352,426	-70	-13,444	-227	-28,819	-148	-7,224
Largest	1.79E+07	6.46E+06	2.11E+06	9.10E+06	2.60E+07	1.08E+07	8.44E+06	1.71E+07	2.00E+08	2.76E+07	1.63E+07	2.60E+07
Obs	105,212	105,430	113,703	113,694	105,212	105,430	113,703	113,694	105,210	105,431	113,703	113,694

Notes: This table corresponds to Figure 4.