

NBER WORKING PAPER SERIES

DOMINATED OPTIONS IN HEALTH-INSURANCE PLANS

Chenyuan Liu
Justin R. Sydnor

Working Paper 24392
<http://www.nber.org/papers/w24392>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2018

We thank the Kaiser Family Foundation and the Health Research Educational Trust for access to the data from the Survey of Employer Health Benefits used in this analysis. We thank Keith Ericson, Amy Finkelstein, Ben Handel, Nathaniel Hendren, Nicola Lacetera, Devin Pope, John Mullahy and participants at the American Risk and Insurance Association Meetings and University of Wisconsin Center for Financial Security Household Finance Workshop for comments and suggestions. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Chenyuan Liu and Justin R. Sydnor. 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.

Dominated Options in Health-Insurance Plans
Chenyuan Liu and Justin R. Sydnor
NBER Working Paper No. 24392
March 2018
JEL No. D22,G22,I13

ABSTRACT

Recent studies have found that many people select into health plans with higher coverage (e.g., lower deductibles) even when those plans are financially dominated by other options. We explore whether having dominated options is common by analyzing data on plan designs from the Kaiser Family Foundation Employer Health Benefits Survey for firms that offered employees both a high-deductible (HD) health plan and a lower-deductible (LD) option. In 65% of firms the high-deductible option would result in lower maximum possible health spending for the employee for the year. We estimate that the HD plan financially dominates the LD plan at roughly half of firms across a wide range of possible health spending needs employees might anticipate. The expected savings from selecting the HD plan are typically over \$500 per year, often with no increase in financial risk. We present evidence that these patterns may arise naturally from employers passing through large average-cost differences between HD and LD plans to their employees. We discuss the implications of those dynamics for the nature of transfers between employees and the efficiency of health spending.

Chenyuan Liu
University of Wisconsin at Madison
975 University Ave
Madison, WI 53706
chenyuan.liu@wisc.edu

Justin R. Sydnor
Wisconsin School of Business, ASRMI Department
University of Wisconsin at Madison
975 University Avenue, Room 5287
Madison, WI 53726
and NBER
jsydnor@bus.wisc.edu

As part of a broader pattern of expanding insurance choice in the U.S., employers are increasingly giving employees the option to select plans with different coverage levels as part of their benefit packages (Bundorf, 2016; Claxton, 2016). For example, many firms offer both plans with lower deductibles and a qualifying high-deductible health plan that can be matched with a health savings account. It is natural to think about these menus as offering employees tradeoffs between expected spending and risk. However, a few recent studies have analyzed situations where higher-deductible plans financially dominate lower-deductible options for employees (Handel, 2013; Bhargava, Loewenstein and Sydnor, 2017). This raises the question of whether these are unique (and valuable) case studies for analyzing individual decision making or if there are more systematic patterns causing plan options with lower coverage levels to be relatively financially attractive and even dominating.

We explore this question by analyzing data on employer-sponsored health plans from the Kaiser Family Foundation (KFF) and Health Research Educational Trust (HRET) Survey of Employer Health Benefits. The KFF/HRET survey includes detailed information on health plan designs, including employee premium contributions and plan cost-sharing rules. Our analysis focuses on the details for employee-only coverage.¹ To isolate firms offering employees a choice based on different levels of insurance coverage, we identify 331 firms that offered both a high-deductible (HD) health plan paired with a health savings or reimbursement account (HSA or HRA) and a second plan with a lower deductible (LD). Under the assumption that the unobservable quality of the provider network of an HD plan offered by an employer is at least as good as the network in the lower-deductible option, we can use the details about the financial plan designs to compare plans.²

We find that high-deductible plans often offer employees *lower* worst-case spending risk. This occurs because the combination of employee premium savings and the contributions to health savings/reimbursement accounts that employers make for those who opt into the high-deductible plans often exceeds the increase in out-of-pocket maximums employees face with high-deductible plans. On average the difference in deductibles between the high-deductible and low-deductible options for employee-only coverage is around \$1,300 and the differences in maximum-out-of-

¹ The KFF data do not have full details of plan designs for the family-level coverage, but from the data we can observe we suspect that the patterns we document would likely be at least as strong for family-level coverage.

² In Section 2 we discuss this assumption and argue that it is likely to be reasonable in this setting.

pocket limits between these plans is around \$550. The direct savings that employees receive from selecting the HD options, through employee-premium reductions and firm contributions to HSAs and HRAs, averages over \$1,000. Ultimately, at 65% of the firms, the direct savings for those who select HD plans exceeds the increase in maximum-out-of-pocket limits.

We further estimate that at a substantial fraction of firms the HD plan offers employees stochastically dominating, and frequently even strictly dominating, distributions of total yearly health costs. This analysis requires us to account for the complexities of the full schedule of cost-sharing for each plan. In order to make this comparison in a tractable way, we use an approach developed by Ericson, Kircher, Spinnewijn and Starc (2015) that transforms a highly complex health plan with many different procedure-specific cost-sharing rules into a simpler approximation of that plan based on the same deductible and maximum-out-of-pocket limit but a single coinsurance rate. With these simplified plan representations, we can compare the spending consequences for employees of different plan choices for different levels of total medical bills they may incur for the year. We find that with the simplified plan representations, at one third of firms the employee's total costs for the year with the HD plans strictly dominates the total costs with the LD plan for any level of total medical spending. When we further combine the information on plan schedules with possible ex-ante distributions of total medical spending employees might face, we find that the HD plan second-order stochastically dominates at about half of the firms across a wide range of possible ex-ante spending distributions.

These differences in plan value might be insignificant if few employees select into lower-deductible plans at the firms where those plans have especially high relative costs. Prior research on individual decisions in these settings, however, reveals that high shares of employees select lower-deductible plans even at extremely high costs due both to inertia (Handel, 2013) and active choice into those plans (Bhargava et al., 2017). While the KFF survey does not have individual-choice data, firms do report the total share of enrollment in their different plan categories. On average the majority of employees are enrolled in LD options, even at firms where we estimate that the HD plan dominates.

The expected financial savings to employees who select high-deductible plans are non-trivial. We estimate that the average savings to employees of selecting the high-deductible option across all firms offering both types of plans is over \$500 per year for a wide range of different ex-

ante medical-spending distributions the employee might face. At the firms where the simplified representation of the high-deductible plan strictly dominates the alternative, the expected savings from selecting the HD plan are frequently over \$1,000 for the year for a typical employee. The sizeable level of savings for an HD plan is not affected much by instead measuring risk-adjusted savings using various levels of constant absolute risk aversion.

A natural question is why employers would design menus that present employees with dominant and dominated options. The data available in the KFF survey do not allow us to answer this question conclusively and it is generally difficult to disentangle the potential motivations behind employers' benefit designs. However, we present analysis that helps explain how these situations can arise rather naturally in employer-sponsored settings. The mechanism is that a combination of adverse selection and moral hazard may cause a wider disparity in the average costs (from the insurer perspective) for enrollees in the HD plans versus the LD plans than would be expected simply from the reduction in coverage levels between those plans. Firms appear to roughly equalize their total contributions to health care across plans in their menu, which means they pass through these large average cost differences to employees enrolling in the HD plans. This can generate substantial savings for those who enroll in the HD option and when the difference in deductibles and out-of-pocket limits between plans are not large, can cause the HD option to become financially dominant.

We discuss how the implications of these dynamics depend on the extent to which excess disparities in average costs across plans arise due to adverse selection versus moral hazard. If the differences were entirely driven by selection on ex-ante health risk, then there is a large transfer from LD plan enrollees to HD plan enrollees relative to a situation where all employees pool in a single plan option. We quantify the size of that transfer using a simplified counter-factual analysis in the spirit of Handel and Kolstad (2015) and find that it averages around \$200 per year. The other possibility is that HD plans have lower average cost because employees reduce medical spending substantially in response to higher cost-sharing. In this case, the emergence of dominating HD options that do not attract all employees suggests sizeable economic inefficiencies. Frictions people face in choosing plans or anticipated behavioral inefficiencies under higher cost-sharing (ala Baicker et al., 2015) could make it difficult for some employees to take advantage of the available savings with HD options they could have even if they did not want to alter their medical spending under the HD plan. The existing literature has documented strong evidence for

both substantial adverse selection and moral hazard responses in employers-sponsored health insurance. While we cannot disentangle these forces for the firms in the Kaiser Family Foundation data, this analysis highlights that these dynamics are important for better understanding the effects of the proliferation of coverage-level options in employers-sponsored insurance.

This study presents the first systematic analysis of the net financial value of plan options with different levels of coverage in employer-sponsored health insurance. Prior research has exploited situations where plans with lower deductibles are dominated by higher-deductible options to evaluate individual insurance decisions relative to the benchmark of standard insurance-demand models (Handel, 2013; Bhargava et al., 2017). Our results show that these situations are not only interesting case studies, but are common in employer-sponsored health insurance. In fact, we find that HD plans are equally likely to dominate LD plans across firms with very different characteristics, such as size and employee salary levels. These findings are significant, in part, because prior research has documented that employees with chronic health conditions, lower incomes and lower education levels are more likely to select costly lower-deductible plans in these situations (Bhargava, et al., 2017). For example, Handel (2013) documented that families of 3 or more who were earning less than \$41,000 per year were often choosing dominated options, costing them at least an additional \$1,000 per year for sure. The fact that high-deductible options offer substantial expected savings relative to their additional out-of-pocket risk, even to the point of dominating lower-deductible options in many cases, implies that the expansion of plan options may be creating substantial benefit disparities between groups of employees based on health, financial security, and insurance literacy.

Our findings also highlight an interesting disconnect between the popular perception of high-deductible options and the reality of the economic value of these plans in employer-sponsored settings. There has been substantial focus in recent years on the fact that deductibles in employee-sponsored health plans are rising faster than wages and concern that this trend is exposing employees to financial harm (e.g., Altman, 2016; Evans, 2017). Some of this concern may, however, be exaggerated if analysis does not also account for the size of premium reductions and contributions firms make to health savings and reimbursement accounts for those in high-deductible plans. Yet it is also clear that aversion to higher-deductibles is widespread in insurance (Sydnor, 2010; Abaluck and Gruber, 2011; Bharshegyan et al., 2013; Handel and Kolstad, 2015; Bhargava et al., 2017). This highlights the need for more research into decision aids that help

people better understand the consequences of insurance options (Johnson et al., 2013; Loewenstein et al., 2013; Samek and Sydnor, 2017). It also suggests a need for more research into alternative motives driving insurance demand, including liquidity constraints and budgeting problems (Ericson and Sydnor, 2018) as well as concerns people may have about how higher levels of cost sharing might cause them to avoid getting care in suboptimal ways (Brot-Goldberg et al., 2017; Baicker, Mullainathan and Schwartzstein, 2015). These considerations may well imply that the choice of lower-deductible options that appear to be dominated from classical insurance-demand perspectives are rational for some people. What our analysis shows is that at many companies employees who select LD options over HD options, whatever the mix of rationality and confusion behind that choice, are receiving a lower level of net health benefits for the year.

Finally, the results call into question the role of corporate human resource departments and benefit consultants in designing choice options for health insurance. While there may be a number of motives for firms to offer high-deductible plans that are relatively attractive, we suspect that relatively few benefits departments at companies where HD options dominate have analyzed the plan options from that perspective or communicate that fact clearly to their employees in benefit materials. Our findings complement recent research by Abaluck and Gruber (2016), who find that school districts with more plan options tended to have more low-value options in their menu, suggesting that some benefit managers fail to eliminate low-value options as they expand choice. Overall, it may be important to invest more research into helping firms better understand the consequence of their insurance-option designs. Regulation of employer-sponsored insurance may also benefit from focusing on the details of plan designs and how employees' choices across insurance levels may generate disparities in the overall benefit value employees receive.

2. Data

Our primary dataset is the Kaiser Family Foundation Employer Health Benefits Survey (KFF EHBS). The EHBS is an annual survey on employer-sponsored health insurance plans that is widely cited in health care research.³ The survey draws its sample from a list of U.S. private and

³ See for example, Claxton et al. (2015)

public employers with three or more workers, covering ten industry categories and six firm-size categories.⁴ We use data from the 2015 survey.⁵

The key benefit of the survey for our purposes is that firms are asked to provide rich details about the plans they offer to their employees, including information on cost-sharing features and premiums. Each plan reported by a firm is classified within the survey into one of four plan types: preferred provider organization (PPO), health maintenance organization (HMO), point-of-service plans (POS), and high deductible health plans (HDHP). Plans are classified as HDHPs if they have both a deductible for single (i.e., employee-only) coverage larger than \$1000 and either a paired health saving account (HSA) or health reimbursement arrangement (HRA) option.⁶ We refer to these as HDHPs or simply high deductible (HD) plans throughout. All other plans are classified as either preferred provider organization (PPO) plans, health maintenance organization (HMO) plans or point-of-service (POS) plans. We refer to the latter three as either non-HD or equivalently as low deductible (LD) plans throughout the analysis. For each type of plan offered, the survey collected information including premium, deductible, cost-sharing etc. If a firm offers more than one plan in each type, only details of the plan of that type with the largest enrollment are recorded.

Our analysis focuses on firms in the survey that report data for both an HDHP plan and exactly one of the low deductible (LD) plan types. We compare HDHP with an LD alternative in part because the anecdotal cases of health plan menus with dominated options reported in previous studies have been situations where higher deductible plans dominated lower deductible plans. Furthermore, our ideal experiment would be a situation where the other aspects of plans, including the generosity of their network and restrictions on access to care or covered services, are held fixed. The Kaiser Family Foundation data do not have information about the networks covered in the different plans. By their very nature, HMO and PPO plans almost surely differ on some of these dimensions and for this reason we do not compare HMO and PPO offerings to one another. In contrast, HD plans often have similar networks as LD plans within the same firm (e.g., Handel 2013, Brot-Goldberg et al., 2017), or have more generous network access since high deductibles

⁴ <http://kff.org/health-costs/report/2015-employer-health-benefits-survey/>

⁵ Some firms responded with options offered in 2014, some with 2015 options. That means all plans are offered at the same time within a firm, but the plan-year covered in the data may differ across firms.

⁶ For a single-coverage plan to be eligible for a paired Health Savings Account per IRS rules it needs a deductible of \$1250 in 2014 and \$1300 in 2015.

can substitute for other forms of cost control (Ellis and Zhu, 2016)⁷. A key assumption for the analysis to follow is that the non-financial features of the HD and LD plans offered by the same firm are similar enough (or potentially slightly biased in favor of HD plans) so as to make a comparison of the financial features of those plans meaningful.

It is also important for the relevance of our analysis that employees actually have a choice between these plan options and that the plans we observe in the data are not offered to separate subsets of employees. The Kaiser Family Foundation survey does not explicitly address the issue, but we believe for most firms employees face a choice. First, the survey excludes any plan that is offered and administered exclusively by a union. Second, Bundorf (2012, 2016) reports that large firms generally offer high-deductible options alongside other plans. Recent empirical studies have also relied on data from firms offering employees choices between coverage options (e.g., Einav, Finkelstein and Cullen, 2010; Handel, 2013; Handel and Kolstad, 2015; Bhargava, Loewenstein and Sydnor, 2017). Finally, a recent study by Willis Towers Watson (2015) on U.S. firms with 1,000 or more employees reports that while the number of firms offering only a HDHP is rising, that a large number of firms “still offer alternatives”. That same study reports that around half of those firms are planning to more highly subsidize HDHP plans in an effort to encourage employees to enroll in those options, again implying that employees are being given a choice.

There are a total of 1771 firms that responded to the 2015 version of the survey and reported offering at least one health insurance plan to their employees. Among those firms, we have complete and consistent plan information for 2372 plans offered by 1529 firms. About 30% of all plans are HD plan, and they are offered by about 46% firms. For our analysis sample, we focus on firms reporting details for 1 HD plan and 1 LD plan. After dropping firms with plans with certain complicated cost-sharing structures for which an actuarial value cannot be calculated us the CMS Actuarial Value Calculator, we are left with a match sample of 331 firms offering 662 plans. Detailed cleaning steps are in Appendix I.

We analyze information for the single-coverage (i.e., employee-only) tier at each firm. The KFF data does not contain complete plan information for family-level coverage, which prevents us from doing our full analysis on the family-coverage level. However, we present some limited

⁷ This does not rule out the possibility that firms offering HD plans may tend to have more restrictive networks than firms that do not offer HD plans, but within firm we expect the HD plans to have a weakly more generous network.

comparison between single-coverage and family-coverage on the variables where we have information below.

Table 1 presents summary statistics of key plan financial features and firm demographics both in the broader sample of surveyed firms and for the 331 firms for our analysis sample. We discuss the details of these statistics for our analysis sample in more detail in the next section. At this stage we note primarily that the financial features of HD and non-HD plans look broadly similar when we compare the full sample and the analysis sample of firms offering both types of plans. Key differences are that firms in the analysis sample are somewhat more likely to have self-insured their plans and are somewhat more likely to report having more than 1,000 employees.

In parts of our analysis we supplement the data on plan options with information on distributions of total medical spending. For those analyses we use data from the Center for Medicare and Medicaid Services and the Medical Expenditure Panel Survey. We describe the relevant details of these data sources below when we incorporate them in the analysis.

3. Analysis of Maximum-Spending Risk using Raw Plan Designs

We begin our analysis by examining how plan designs and contributions to health savings/reimbursement accounts map into maximum-spending risk for an employee. This analysis allows us to get a first simple look at the financial value of different insurance options before delving into the complexities of cost-sharing rules, which we address in the next section.

In Table 1, consider first the average lower-deductible option. For these plans, the average annual deductible is \$846 and the maximum-out-of-pocket limit averages \$3,464. The employee is responsible for a total of \$1,450 for employee-paid premiums on average for these plans. So an employee that hits her deductible would expect to pay approximately \$2,300 on average for the year in health costs (premium + deductible). The worst-case total health spending in these plans would average just over \$4,900 (premium + max out-of-pocket).⁸

⁸ The numbers we report here and throughout the analysis are pre-tax dollars. While health-insurance premiums are paid with pre-tax dollars, out-of-pocket medical costs are sometimes paid with post-tax dollars. For all HD plans considered here, however, employees have access to HSA/HRA accounts that allow them to pay for out-of-pocket costs with pre-tax dollars. Many LD plans can also be paired with medical spending accounts that similarly allow them to cover out-of-pocket costs with pre-tax dollars, though these accounts are more limited than HSA accounts. Our focus on pre-tax dollars, if anything, creates a slight bias that makes LD plans look more financially attractive

The HD plans at these firms naturally have higher deductibles (avg. \$2,166) and maximum-out-of-pocket limits (avg. \$4,011), but also have substantially lower premiums (avg. \$841) and additional contributions by the firms to health savings and reimbursement accounts. When we compare the plan types we see that an employee selecting an HD plan would on average be exposed to a \$1,320 increase in the deductible and a \$547 increase in the maximum-out-of-pocket limit. An employee in the average HD plan, however, would pay \$609 less in employee-paid premiums over the course of the year. This alone is enough to cover the difference in maximum out-of-pocket spending for the year and is just under half of the deductible difference.

Firms, however, also tend to contribute to a health savings account or a health reimbursement account for employees who select the HD options. A contribution made by the firm to a health savings account is fully vested immediately for the employee. Money in an HSA can be used to pay for medical spending with pre-tax dollars at any time and can be withdrawn for non-medical spending with rules similar to those in tax-advantaged retirement accounts. An HSA contribution by the firm is very similar to a reduction in the required employee premium. Contributions firms make to a health reimbursement account (HRA) are used to reimburse employees for uninsured health spending, but unlike an HSA these funds are the property of the employer and unused funds do not follow the employee if the employee leaves the firm. In our sample we observe that 86% of the firms offer an HSA and 14% offer an HRA, with most firms making contributions into these accounts for their employees. As Table 1 shows, the average contributions to HSAs (\$473) and HRAs (\$758) are sizable, which substantially reduces an employee's maximum-spending risk with the HD option.

Figure 1 shows the cumulative distribution function for the difference in worst-case spending between HD and LD plans across the 331 firms in our analysis sample. We define worst-case spending here as employee premium plus maximum out-of-pocket limit minus any contribution to HSA or HRA for the plan. We find that at 65% of the firms the worst-case spending is *lower* with the HD plan after accounting for both employee-paid premiums and health savings/reimbursement contributions by the firm. On average, the worst-case spending risk is \$576 lower with the HD

than they really would be because some out-of-pocket costs for LD plans might require post-tax dollars while in principle out-of-pocket costs can be paid with pre-tax dollars in almost all cases with an HD plan.

plans. The HD option reduces worst-case spending risk by over \$1,000 at 38% of firms, while it increases the worst-case spending risk by over \$1,000 at only 14% of firms.

We can also do a limited version of this same analysis for the family-coverage data. The KFF data do not have information about the maximum-out-of-pocket limit for family coverage tiers. However, they do have information about the family deductible level for each plan. That information still may be somewhat incomplete for family coverage, because family coverage often included individual-specific sub-deductibles along with the overall family deductible. However, analyzing the information we do have gives us at least suggestive evidence that the issues we document here for single-coverage tiers likely also apply for family coverage. In particular, we compare the net difference in premiums (including HSA/HRA contributions) between HD and LD plans to the net difference in deductibles between plans for both single and family coverage. We find that for single coverage, 140 of the 331 firms offer a greater net premium reduction (employee premium + HSA/HRA contribution) for choosing the HD plan than the deductible difference between HD and LD. For the family-coverage tier we find that 186 of the 331 firms have a larger net premium difference than deductible difference, which suggests this issue is likely to be at least as prevalent for family coverage. We further find that of the cases where the single coverage tier shows this pattern of higher net premium differences than deductible differences, that also holds for the family coverage tier at those firms in 85% of cases.

The analysis in this section shows that the worst-case spending risk is lower with the HD option at nearly two-thirds of the firms offering an HD and LD option. This pattern is consistent with the possibility that the total health spending distributions employees face with HD plans could dominate what they face with LD alternatives. However, it is also important to note that the worst-case level of spending will not be reached by most people and assessing the net insurance value of different options requires us to do more to account for the full range of spending consequences an employee could face with each plan.

4. Analysis of Dominance using Simplified Plan Representations

We aim in this section to analyze the distribution of total-spending consequences an employee would face under the different plan options. At a basic level, most health insurance plans create a non-linear schedule that maps from total medical bills to the insured's out-of-pocket spending in

three segments: 1) the region under the deductible, where medical bills transfer to out-of-pocket costs on a one-to-one basis 2) a partial cost-sharing region above the deductible where the individual pays a fraction of each marginal medical bill and 3) a maximum-out-of-pocket region where the limit has been hit and the individual's out-of-pocket spending no longer increases with medical bills. The details of most plans are, of course, more complicated than that, but this basic structure guides the approach for our analysis.

Figure 2 provides an illustration of four different cases we might see when comparing these types of non-linear mappings from medical bills to out-of-pocket costs across plans. For each figure on the x-axis we graph total medical bills for the year. The y-axis shows the employee's total health spending for the year, which in these graphs is the employee's premium plus out-of-pocket costs. We discuss how we incorporate HSA and HRA contributions for HD plans below. Case 1 is an example of a classic tradeoff between HD and LD plans and is likely what most economists would have in mind when considering a comparison of an HD and LD option. In this example, the premium is lower for the HD plan, but the higher deductible and maximum out-of-pocket limits for the HD plan cause the schedules to cross for relatively moderate levels of total medical bills (in this case around \$3,000). The other three cases, however, show examples where the worst-case spending risk is lower with the HD option. In Cases 2 and 3 the HD plan would result in higher costs for the employee only for intermediate levels of medical bills. In a situation like Case 3 the HD plan is likely to second-order stochastically dominate the LD plan. Finally, in Case 4 we see an example where the employee's costs strictly dominate under the HD plan. In Section 4.1 we describe how we transform complex schedules of cost-sharing rules observed for actual plan designs into simplified plan representations that allow us to make the type of comparisons presented in Figure 2.

4.1 Creating Simplified Plan Representations

For our analysis we transform the schedule of cost-sharing rules for each plan into these types of simplified schedules. These simplified representations allow us to compare how different plans map to employee spending (premiums + cost-sharing) for different levels of total medical spending. We follow a procedure introduced by Ericson, Kircher, Spinnewijn and Starc (2015). Their procedure involves first calculating the actuarial value (AV), i.e., share of medical spending

covered by insurance, of a plan using the full cost-sharing provisions of the plan. The AV is calculated using the Actuarial Value Calculator established by CMS for estimating the AV of plans for the Affordable Care Act exchanges. The next step is to solve for a single co-insurance percentage – i.e., fraction of bills the individual must pay once the annual deductible has been met – that generates a plan with the same AV, holding fixed the original deductible and maximum-out-of-pocket limit from the original plan. This generates a simplified plan with only the basic cost-sharing features (deductible, coinsurance, and max out-of-pocket) so that the out-of-pocket consequences of each plan are simply a function of total medical spending for the year. The coinsurance rate for this simplified plan will approximate the average fraction of medical bills paid out-of-pocket under the original plan design at different levels of total health spending. The simplified coinsurance rate is based on the average mix of health services used by people at different levels of total spending.

For example, one of the firms in our analysis sample offers a plan with a \$1,500 deductible, \$5,000 maximum-out-of-pocket limit, a 20% employee coinsurance rate on most medical services once the deductible is hit, a \$20 copay per primary-care office visit, and a 25% employee coinsurance rate on most prescription drugs. We calculate an actuarial value for this plan based on the full set of cost-sharing features of 74.1%. We then calculate that a simpler plan design with the same \$1,500 deductible and \$5,000 maximum out-of-pocket limit would need a single coinsurance rate of 21% to have the same actuarial value.

There are a few other complications that arise in trying to map complex service-specific cost-sharing rules into a simpler structure based on total spending. The main issue is that in many plan designs, some specific services are subject to fixed copayment amounts even before the individual has hit the deductible for the year. This issue tends to most often affect cost-sharing for prescription drug coverage and physician office visits and is much more likely to occur in LD plans than HD plans because federal rules for HD plans typically limit their ability to provide partial coverage for services before the deductible is met. We use information contained in the continuance tables that underlie the CMS AV calculator to break down each total spending level into different categories so that we can get a simplified plan representation that accounts for these nuances of how different services apply to the deductible. Appendix B outlines the steps of our full procedure in detail.

The other slight complication in generating the final simplified plan representations is accounting for contributions by firms to HSAs and HRAs. Since HSA contributions are fully vested for employees immediately, we treat them simply as equivalent to a reduction in the premium paid. As such, they represent a constant shift down in the entire schedule of employee costs at every level of total medical spending. For HRA contributions, in contrast, the money can only be used to offset out-of-pocket medical spending. This then creates in initial section of the employee-cost schedule that is flat under the assumption that medical bills are paid from HRA funds before employee cost-sharing kicks in. Again, the details of this structure are outlined in Appendix B.

It is worth clarifying the way in which our simplified plan representations do and do not rely on distributional assumptions. These representations create a mapping between total medical spending levels and an employee's total costs. The mapping works for people who might end up at different total spending levels and does not rely on assumptions about the distribution over total spending levels. Where the mappings are affected by underlying distributions is that they rely on information on the share of total medical bills going to different medical services. For each level of total medical bills, the CMS continuance tables provide a breakdown of spending by services (e.g., office visits, inpatient, etc...) that is based on the average mix of services used by people with that level of medical bills in the broader population. It is possible, though, for two people to reach a given level of total medical bills with different mixes of services, which means the approach is accurate for people with a typical mix of services given a level of expenditure but will not perfectly represent all situations.

There are a number of reasons that relying on the average mix of services conditional on the level of medical spending is not likely to be a substantial concern for our analysis. First, 22% of all plans start out as simplified plans to begin with and are unaffected by the procedure. Furthermore, there is strong correlation between the level of cost-sharing for different services within plan designs. Plans with low coverage on one service tend to have low coverage on other services as well, which mutes the potential effects of different services mixes.⁹ Finally, using a

⁹ For example, we calculate that there is a correlation of 0.72 between a plan's coinsurance rate for outpatient surgery and the effective coinsurance rate on primary-care office visits (transforming copays to coinsurance rates using the average office visit cost in the CMS actuarial value calculator of \$147). There are also positive correlations between spending on different services. Those who have high drug expenditures, for example, are also likely to have high

typical mixture of underlying services for each level of medical spending is likely to be an appropriate approach for analyzing the insurance value of different plan designs for most people. Most people face fairly substantial uncertainty about the level of total medical spending they will incur in a year and it is that uncertainty which drives the first-order differences in insurance value. Most people, in contrast, have little reason to anticipate that they will use unusual mixes of services conditional on a level of spending and the degree of uncertainty generated by variation in the mix of services conditional on the level of medical spending is likely to be more second order.

We use the procedures outlined in this section to map each of the 662 plans in our analysis sample to simplified plan representations. Appendix C (intended for online publication only) has graphs of the two simplified plan representations (LD and HD) for each of the 331 firms in the analysis sample.

4.2 Results

Once we have converted each plan into a simplified plan representation, it is straightforward to classify the firms based on which of their plans exposes the individual to greater worst-case spending risk and whether one plan strictly dominates the other. That is, we can map the firms to the cases we highlighted in Figure 2. This analysis simply compares the different schedules of employee spending (net premiums plus out-of-pocket costs) as a function of total medical spending and does not require additional assumptions about the ex-ante distribution of total medical spending the individual faces.

Table 2 shows this classification and reveals that a substantial share of firms have menus where the HD plan strictly dominates the LD plan based on their simplified plan representations. We find that 32% of firms have HD and LD options that present a classic tradeoff (Case 1 from Figure 2) where the HD plan has lower spending for the employee if total medical spending is low but the LD plan has lower worst-case spending for the employee. At 37% of firms, however, the HD plan's simplified plan representation strictly dominates the schedule for the LD plan (Case 4). At another 25% of firms, the HD plan shows the Case 3 profile of having lower spending for the employee under both best-case and worst-case medical-spending needs, but higher spending for

levels of inpatient services and/or specialist visits. This type of positive correlation reduces the amount of variation in service mix for a given level of total medical spending.

some intermediate ranges of medical spending. A small fraction of firms (6%) show other patterns, including 2% where the LD option appears to strictly dominate.¹⁰

The final column of Table 2 shows the average percent of enrollees in an HD option for the firms in each of these classification groups. It is important to note that this may not be the fraction enrolled in the specific HD plan for which we have the plan details, as firms report the plan within each category with the highest enrollment. Nonetheless, these reported fractions are revealing in two ways. First, there is a correlation between the classifications and enrollment. On average 25% of employees are enrolled in HD options at the firms where there is a classic tradeoff between HD and LD, while 47% enrolled in HD plans on average at the firms where HD plans strictly dominate the LD plans. So enrollment in HD plans is higher where standard theory says that the HD option is more attractive. Yet it is perhaps more remarkable that even at firms where the HD plan dominates the LD plan only 47% of employees are enrolled in HD plans. This finding is broadly consistent with the results in Handel (2013) and Bhargava et al., (2017).

In Table 3 we classify firms based on whether either of the plan schedules displays first-order or second-order stochastic dominance relative to the other option.¹¹ This analysis requires us to impose a particular ex-ante distribution of total medical spending. For our baseline analysis we use the distribution of total medical spending from the gold-tier combined (medical plus prescription drug) continuance table from the CMS Actuarial Value Calculator. We discuss the robustness of the results to alternative medical-spending distributions in the next section, but note here that this distribution is based on claims data for privately insured individual and is likely a reasonable ex-ante distribution for the average person covered by employer-sponsored insurance.

The table shows that no additional firms have first-order-stochastically-dominant HD plans beyond those already identified as strictly dominant in Table 2. A higher share (53%) of firms,

¹⁰ These few cases where the LD options strictly dominates implies that the higher-coverage LD plan has a lower employee premium than the HD plan, which may suggest a data error in the survey or a case where the LD plan has a much more restrictive network than the HD option. We keep these firms in the analysis, however, so as not to bias the results in favor of finding dominance for HD plans.

¹¹ For two risky prospects represented by the cumulative distribution function F and G , F first-order stochastically dominates G if $F(z) \leq G(z)$, for all z . All expected-utility maximizers whose marginal utility of wealth is positive will prefer prospect F over G if F first-order stochastically dominates G . Prospect F second order stochastically dominates G if $\int_x^k F(x)dx \leq \int_x^k G(x)dx$, for all k . All risk-averse expected-utility maximizer will prefer prospect F over G if F second-order stochastically dominates G .

though, have HD plans that second-order stochastically dominate the LD option. This implies that at these firms any risk-averse expected-utility maximizer who faces the average distribution of total medical spending should prefer the HD option over the LD option, regardless of their degree of risk aversion.

Beyond classifying the firms based on measures of dominance, it is instructive to consider how much money an average employee could expect to save by selecting the HD option over the LD option. For this analysis we again use the CMS gold-tier continuance table for the distribution of ex-ante total medical spending, though in the next section we explore the robustness of these measures to alternative spending distributions. Table 4 shows the results. We calculate on average across all 331 firms that an employee facing this ex-ante distribution of medical spending could expect to save \$569 by selecting the HD option over the LD option. The 25th percentile of expected savings is \$192 and the 75th percentile is \$884. At the 37% of firms where the HD option strictly dominates, the HD option would save the representative employee an average of \$1,121.

Of course, when considering insurance options, focusing on expected savings is not entirely appropriate because foregoing some money in expectation for a reduction in risk is the ultimate purpose of insurance. In order to address that point we also analyze risk-adjusted measures of spending. For this analysis we assume that the individual has constant absolute risk aversion (CARA) utility over total employee spending for the year, which is a common functional form used for utility in studies of health-insurance demand. For different levels of constant absolute risk aversion (r), we can then calculate the expected utility for each plan option, again integrating over the distribution of total medical spending from the CMS gold-tier continuance table. In order to calculate the risk-adjusted savings measure we then simply solve for the minimum sure amount of money that would have to be given to the individual under the LD plan to equate the utility of the LD plan with that of the HD plan.

The second two columns of Table 4 give the risk-adjusted savings results for $r = 0.0005$ and 0.002 respectively. The bottom rows of the table help to benchmark the implications of these levels of risk aversion. One benchmark is the certainty equivalent a person with CARA utility of that degree would have for a 50% chance of winning \$1,000, which is \$438 for $r = 0.0005$ (i.e., a \$62 risk premium) and \$283 for $r = 0.002$ (i.e., a \$217 risk premium). The second benchmark calculates the gain a person with that risk aversion would need to accept a 50/50 gamble with a

downside of losing \$1,000. For $r = 0.0005$ that gain is \$2,092, while there is no level of gain that would make an individual with CARA $r = 0.002$ accept a 50% chance of losing \$1,000. As such, we can reasonably classify these cases as spanning substantial to extreme levels of risk aversion. On average the risk-adjusted savings levels for choosing HD plans are somewhat lower than, but similar to, the expected savings at \$471 for $r = 0.0005$ and \$440 for $r = 0.002$. So even if we assumed that employees had extreme levels of risk aversion ($r = 0.002$), employees with representative distributions of medical spending would on average see a risk-adjusted savings of over \$400 for opting into HD plans over LD plans. At the 37% of firms where the HD plan strictly dominates, increasing levels of risk aversion actually lead to greater levels of risk-adjusted savings than expected savings, which relates to the fact that HD plans have lower worst-case spending.

4.3 Robustness of Main Results

The analysis of second-order-stochastic dominance patterns and (risk adjusted) savings in Section 4.2 used the CMS gold-tier distribution for total medical spending. While this is a reasonable representation of the ex-ante medical spending for an average employee, it is worth considering how our analysis would look for employees facing different distributions of medical spending. The CMS Actuarial Value Calculator actually has four different underlying continuance tables, one for each tier in the ACA exchange market (bronze, silver, gold, and platinum). We find that the results from Section 4.2 are nearly identical using any of these distributions. However, these distributions are all based on large populations of people with certain levels of insurance coverage and as such may not represent the type of idiosyncratic variation in medical-spending risk we might reasonably see within the employee population at a given firm.

In order to generate additional variation in the medical-spending distributions for our analysis, we turn to data on total medical spending from the 2014 wave of the Medical Expenditure Panel Survey (MEPS). We use data for the population with coverage from an employer or group insurance plan in place throughout the year. We then run an ordinary-least-squares regression of total medical spending for the year on a cubic polynomial in age, gender and indicators for self-reported health status (5-categories) and the interactions of all of those variables.¹² We obtain the

¹² MEPS respondents evaluate their health status according to 5 categories: excellent, very good, good fair, and poor.

predicted values from this regression, which is a measure of the expected total medical spending based on these basic demographic indicators. We then group the individual MEPS respondents in our regression sample into deciles of predicted medical spending. This yields 10 distinct groups of people, with substantial variation across groups in expected medical spending. For each of these ten groups we then consider the distribution of realized medical spending. That is, we generate 10 representative individuals with different ex-ante medical spending distributions (which match the distribution of ex-post realizations for the MEPS respondents in that group).¹³

Figure 3 plots the mean, median and 80th percentiles of realized spending for each of these deciles of predicted medical spending. The 10th decile, for example, has average total medical spending of around \$15,000 and median spending of just under \$6,000, the latter of which is comparable to the mean spending from the CMS gold-tier distribution used in our baseline analysis. Overall, we see that these 10 distributions vary substantially in both their expected spending and the spread of spending from the median to 80th percentile.¹⁴

Figure 4 shows how our classification of firms based on whether the HD plan second-order stochastically dominates the LD plan varies across these different distributions. Consistent with our baseline results, we find that approximately 50% of firms have HD plans that second-order stochastically dominate the LD plan across these very different distributions of medical spending risk. The basic trend is that for lower-expected-spending distributions there is a slightly lower share of second-order-stochastically-dominant HD plans and a somewhat higher share for the highest-spending distributions. This is consistent with the fact that many HD plans have lower employee spending for very high levels of medical spending but not for more intermediate medical-spending outcomes. Hence, these results show that the basic conclusion that around half of firms in this sample have dominant HD plans is robust to a wide range of different ex-ante distributions of medical spending needs one could consider. Moreover, if one were concerned that

¹³ For this analysis we use the sample weights provided in the MEPS in the regression predicting total spending. We also use the weights to calculate discrete probabilities of different total spending levels (84 atoms), consistent with the atoms in the CMS continuance tables used in the primary analysis.

¹⁴ The mean spending in the CMS gold-tier distribution is very close to the mean spending for the 8th decile of expected spending in our MEPS sample. This suggests that the MEPS spending levels are lower than those used in the CMS continuance tables. This partly reflects the fact that our MEPS data is from 2014, while the CMS continuance tables are projected for spending levels in 2016. This pattern is also consistent with analysis by Aizcorbe et al., (2012), who compare total spending levels in the MEPS 2005 data and claims data from the MarketScan database and conclude that MEPS estimates are 10% lower on average.

our baseline analysis might overstate the dominance of HD plans for those who could reasonably anticipate that they have high medical spending results, these results show that if anything the share of firms with dominant HD plans looks larger for those who can expect to have very high levels of medical spending.

Figure 5 shows the robustness of the expected-savings and risk-adjusted-expected-savings results across these alternative medical-spending distributions. Here we see more variation based on the distribution used, but again the basic conclusions that HD plans offer large average and risk-adjusted savings compared to LD plans holds. Specifically, we estimate average savings for HD plans (across all 331 firms) of over \$400 for the year with every one of the medical spending distributions, including the very high-risk distribution in the 10th decile. The risk-adjusted savings for CARA $r = 0.0005$ are again substantial, over \$400, for each of the 10 distributions.

Another potential concern for the robustness of our baseline results is that our procedure for calculating the actuarial value of each plan using the CMS Actuarial Value Calculator requires us to make some judgment calls about certain cost-sharing features that do not perfectly align with the CMS calculator. Those judgment calls are typically quite minor and are detailed in Appendix B. Appendix Table B1 shows the fraction of firms at which the HD option is classified as strictly dominant based on the simplified plan representations and the expected savings to an employee selecting the HD plan for various subsamples of the data. If we limit analysis to the 166 firms where we do not need to make any judgment calls, the fraction with strictly dominant HD plans and the average expected savings measures are nearly identical to those in the main analysis. Similarly, the KFF survey data contains a small number of imputed values, but we find that the results are not affected if we limit to the subsample with no imputed values.¹⁵ Finally, we find that the CMS AV calculator can generate errors at times in situations where the plan cost-sharing rules involve per-day copays for inpatient stays. Again, we find that the results are unchanged if we limit the analysis to firms where neither plan has these types of copays.

¹⁵ The KFF survey imputes values for a limited number of cases with missing values, which occurs for about 5% of observations. In our matched sample, none of the main variables like premium, deductible, HSA contribution, coinsurance rate and copayment amount for inpatient stay are imputed. Some other variables are imputed in rare cases by KFF by replacing “the missing information with observed values from a firm similar in size and industry to the firm for which data are missing” (KFF, 2015).

5. Analysis of Sources of Plan-value Disparities

Our analysis has established that many firms offering a high-deductible option provide employees large enough premium differentials and health savings/reimbursement accounts that the HD plan offers substantial expected savings and frequently dominates. A natural question then is why employers would offer menus of health insurance options for employees where some options dominate.

As a first approach to exploring this question, in Table 5 we analyze whether the likelihood of offering a dominated option is related to firm characteristics in a series of one-way correlations. We find that offering a dominant HD plans is not correlated with major characteristics of the firm, including measures of firm size, employee salaries, employee age, unionization status, or whether the firm is self-insured for health benefits. Multivariate regressions controlling for these features simultaneously similarly show little correlation between firm characteristics and offering a dominant HD option. The fact that dominated options appear at similar rates across different firm situations suggests that the reasons that options include dominant HD plans is likely rooted in some process that is common across many firms. While the data available do not allow us to resolve this issue conclusively, our additional analysis that suggests that dominated options may result from some basic forces in how employee health-insurance benefits are established.

First, we note that the total premiums firms report for their plans are consistent with substantial differences in average medical costs for those enrolled in LD plans versus HD plans. Average cost differences could arise due to differential plan selection by employees with different health status or due to the moral hazard effects of the HD plan.¹⁶ Figure 6 shows that the total premium differences between HD and LD plans are substantially larger than the differences in expected coverage between these plans for an average population. For the population distribution of medical costs we use the CMS gold-tier distribution from our primary analysis. The graph divides firms into four groups based on the difference in the deductible between their HD and LD plans. For each of these groups, the graph plots the average differences in expected out-of-pocket costs (HD

¹⁶ Proportional loading built into premiums to recover administrative costs can also increase the cost differentials between plans with different levels of coverage. Loading is likely, however, to account for only a small portion of the differences. Medical loss ratios (claims pay outs over total premiums) for group policies of this type are typically between .8 and .85, which equates to a loading of 1.18 to 1.25. The ratios of premium differences to expected-coverage differences we observe, however, are generally more than 2 and often over 5.

– LD) and the difference in reported total premiums between the plans (LD – HD). These total premiums include the firm’s contribution and employee contributions. The total premiums are significantly larger than the expected coverage differences. For example, for firms where the average deductible difference is around \$1,000, the average difference in expected out-of-pocket costs for a representative employee based on the CMS gold-tier distribution is \$407. The total premium difference between the HD and LD option in these cases, however, is more than double that level at \$952, which in turn is nearly equal to the deductible difference. Similar patterns hold for other deductibles, with the largest absolute and proportional differences between premium differentials and out-of-pocket differentials occurring for plans with deductible differences around \$500.

Whether average-cost differences between HD and LD plans will lead to situations of dominance depends largely on how employers pass through those cost differentials to employees. Figure 6 shows that on average these total premium differences are fully passed through to employees. We calculate the net employee premium difference as the difference in employee premium between the LD and HD plan plus any HSA contribution an employee would receive from selecting the HD plan. This difference is very close to the total premium differentials at each level of deductible difference. Figure 7 shows another way of seeing the pass-through of premium differentials. This figure plots the difference in total firm contributions (employer paid premiums + any HSA contribution) for firms in our matched sample offering HD plans paired with HSA accounts.¹⁷ Here we see that firms typically contribute very similar total amounts toward these two plan types. In fact, just under 15% of firms have total contributions that differ by less than \$50. Firms, on average, contribute a little more to the HD plans (after account for HSA contributions), but the average difference is only \$16.

Taken together these pieces of information suggest a straightforward mechanism for how dominated options emerge. The average medical costs of employees who opt into HD options are substantially lower than the cost for those who opt into LD plans, reflecting the impacts of adverse selection and moral hazard. Those average-cost differences are reflected in the total premiums for plans. Firms equalize their contribution to the two different plan types, which means they pass on

¹⁷ We omit the firms with HD plans paired with HRAs from this analysis because it is unclear how to calculate the expected firm contribution under HRAs.

these large average-cost differentials to the employees through a combination of differential employee-premium contributions across plans and HSA/HRA contributions by the firm. Finally, in cases where the coverage levels are not too different, e.g., deductible differences of \$1,000 or less between HD and LD plan, these large premium and HSA/HRA differentials cause the HD plans to dominate the LD plans.

This basic mechanism is consistent with the differences we see in the patterns of premiums and firm contributions when we compare firms based on our classification from Section 4 of whether or not the HD plan strictly dominates. For example, the difference in deductibles between the HD and LD plans is smaller at the firms where the HD plan strictly dominates, averaging just under \$1,000 at those firms versus just over \$1,500 at the firms without strictly dominant HD plans. The firms with dominant HD plans, though, also report especially high total premium differences between the plans, consistent with strong average cost differences across plans at those firms.¹⁸ The firms with dominant HD plans also come close to equalizing their total contributions between the HD and LD plans, meaning that these large premium differentials are passed through to employees.¹⁹

Of course, not all instances of dominance necessarily result from this exact process. For example, some cases where HD plans dominate arise in situations where the firm has made substantial HSA contributions so that the firm is actually contributing more for the HD plans than the LD plans by a significant margin. In some of these cases, firms may be temporarily offering strong subsidies to induce employees to switch voluntarily to HD options with the intention of lowering firm contributions to that plan over time or eventually limiting to only the HD option. One motivation for such an approach could be avoidance of the possibility of paying the so-called Cadillac Tax on high-actuarial-value health plans established by the Affordable Care Act. We note that Table 5 shows that firms that reported doing some analysis in anticipation of the Cadillac Tax

¹⁸ The average premium difference at firms with strictly dominant HD plans (among firms offering HSAs) is \$1,325, while the expected difference in out-of-pocket costs for a representative employee facing the CMS gold-tier medical spending distribution is only \$343. In contrast, at the firms without strictly dominant HD options, the average reported premium differences are \$963, while the average expected difference in out of pocket costs is \$613.

¹⁹ On average the firms with strictly dominant HD plans contribute \$144 more to the HD plans, while firms without strictly dominant HD plans contribute \$55 less on average to the HD plans. These contribution differentials, however, do not explain much of the dominance patterns. At the firms with the dominant HD plans, the average net difference in costs (including premiums and firm HSA contributions) an employee faces between the plans is \$1,469, which means that the excess contributions to the HD plan account for less than 10% of the differential at these firms.

were slightly more likely to have a strictly dominant HD plan (39% vs 32%). It is also worth noting that our baseline explanation laid out above also leaves open the question of why firms equalize contributions between plans and why reported plan premiums seem to reflect average-cost differentials rather than being set based on plan's actuarial value differences applied to the expected spending for the employee population as a whole. We discuss this question in the concluding section. What the results in this section do suggest, however, is that there is a straightforward mechanism by which dominated menus, and more generally large expected savings from the HD plans, can arise without firms consciously and purposely creating dominated menus.

One important note here is that the mechanism discussed in this section implies that the domination of HD plans is a partial-equilibrium effect. That is, if all employees shifted to the HD plans and did not alter their medical utilization, firms likely would change the contributions they make so that these HD plans are not so attractive. For example, if the large premium differences between HD and LD plans reflect adverse selection into LD plans, employees shifting away from LD plans would cause the premiums for HD plans to rise.²⁰ Similarly, if premium differentials reflect moral hazard effects, employees who currently select LD plans could not shift wholesale to HD plans while keeping their use of medical services the same without that affecting the premium differences that drive the domination in the first place. This means that HD plans are not dominant options in any absolute sense, but instead are dominant within menus that offer multiple coverage levels as options. We elaborate on this point in more detail in the next section.

6. Economic Implications of Health-Plan Menus with Dominated Options

We have documented that the health-plan menus at many firms are setup so that an employee who fails to select the higher-deductible plan is leaving money on the table. In this section, we consider the economic implications of these patterns. Specifically, who benefits from this situation and do these empirical patterns point to any underlying economic inefficiency in employer-sponsored health insurance?

²⁰ The premium levels of HD and LD plans at firms that offer only one type of plan are consistent with this point. The difference in average total premium between HD plans and LD plans reported by these firms is consistent with the difference in average actuarial value between these plan types.

Based on the prior section, if we assume that HD plans dominate because firms pass through large average-cost differences to those who choose the HD plans, then the economic implications of these patterns depends in large part on the extent to which those average cost differences are driven by risk selection versus moral hazard.²¹ In general, plan disparities rooted in risk selection have bigger implications for the nature of transfers between different groups of employees, while disparities rooted in the causal effects of higher cost-sharing on medical spending (i.e., moral hazard) have bigger implications for the efficiency of medical spending.

6.1 Implications When Average Cost Differences are Driven by Selection

Consider first the “selection case” where the discrepancy between the HD and LD plans is solely the result of different average ex-ante medical needs between employees selecting into the different plans. In this case, employees in the HD plan are able to avoid pooling with the LD-plan enrollees who are less healthy on average. Because the firms’ total contribution to benefits, through employer-paid premiums and HSA/HRA contributions, are the same on average across plans, the higher benefits of the HD plan represent a transfer from the LD-plan enrollees to the HD-plan enrollees relative to the situation where they all pool together in a single plan.

We can provide a quantitative estimate of the size of that transfer by considering how the situation for employees enrolling in the HD plans would change in the event that all employees were either forced or effectively nudged to enroll in the HD plan. To get this estimate, we need to have an estimate of how much the total premium for the HD plans would rise if all employees were shifted in the HD plan. In Appendix C we discuss how we create an estimate for these implied premium changes using some assumptions about how total premiums are set and the level of loading factors incorporated into premiums. Table 6 reports our approximations for how the premium levels and the net costs employees face for coverage would change under the counterfactual where everyone is enrolled in the HD plan. We limit this analysis to firms that offer HSAs for their HD plans because HRAs are complicated by the fact that the firm retains control of unused HRA contributions. We assume that the firms would hold their total employer

²¹ Consistent with much of the empirical literature on the effects of insurance coverage on medical spending, we are using the term “moral hazard” here broadly to denote any effect different insurance coverage has on the total amount of medical spending a person incurs for the year. Later in this section we briefly discuss the distinction between reducing spending on “wasteful” versus “valuable” medical services.

contributions fixed and we do not distinguish here between the mix of premium versus HSA contributions firms would make. We calculate that the total premiums for HD plans would rise on average by \$221 if all of the LD enrollees switched to the HD plan. After considering employer contribution adjustments, the net impact to the original HD-plan enrollees would be an increase in their costs of \$208 on average. So relative to the counterfactual where everyone pools in the same HD plan, the current HD plan enrollees on average are receiving an average transfer of about two hundred dollars from the LD enrollees. At the 36% of these firms where our analysis in Section 4 labels the HD plan as strictly dominant, the size of the average transfer in this calculation is \$463.

These counterfactual premium help to highlight Handel's (2013) argument that using techniques like "nudges" to help people spot dominant options may not improve and may sometimes harm overall average welfare. If there is substantial scope for adverse selection, offering plans with different coverage levels should lead to "unraveling" where all employees end up in the HD plan with least coverage (Ericson and Sydnor, 2017). The fact that the HD plans can come to dominate the LD plans without attracting all employees suggests there are frictions in plan choice that prevent that full unraveling. Helping a single employee to overcome those frictions and enroll in the HD plan may have large benefits for the single employee. However, doing so for all of the employees requires the premiums and firm contributions to adjust to the new equilibrium. So the switch to the new equilibrium has a cost for the employees who were originally in the HD plan and also erodes part of the savings for the original LD-plan enrollees. As Handel (2013) illustrates, nudging can potentially lower average welfare and may in some cases even make those who were originally violating dominance worse off in the end. However, the size of the average potential savings we observe for those enrolled in LD plans is large enough that employees enrolled in LD plans would likely be better off at most firms if everyone were enrolled in the HD plan.²² It is also important to note that these discussions of welfare, including the analysis in Handel (2013), do not take a *social* welfare perspective. From a social welfare perspective eliminating transfers from LD-enrollees to HD-enrollees would likely be beneficial, since prior

²² In Section 4 we calculated average savings for those enrolled in LD plans of around \$500 overall and around \$1,000 for those at firms with strictly dominant HD options. Even with premiums rising on average by \$200 to \$500 at these firms in the full-HD-enrollment counterfactual, those originally enrolled in LD plans would likely have average savings of more than \$300 across all firms and over \$500 at firms with strictly dominant HD options. Table 4 suggests that the additional variance in out-of-pocket spending in the HD plan would reduce the value of that option by an additional \$100 on average for reasonable levels of risk aversion. As such, the original LD-plan enrollees would likely benefit from interventions that switched all employees to the HD plan.

research has shown that lower-income employees are more likely to enroll in costly LD options (Bhargava et al., 2017).

Finally, if dominant options are arising in employer-sponsored menus due to the effects of adverse selection, it raises the question of whether offering employees a choice over coverage levels is beneficial to begin with. Ericson and Sydnor (2017) highlight that the primary benefit of offering choice is the possibility to allow people with different levels of risk aversion to sort into different options. They also argue that those benefits are likely to be strongest if there are controls that prevent premiums from reflecting differential selection. Employers do not, in principle, need to pass through the average cost differences between LD and HD plans to employees, yet the widespread presence of plan menus with dominant HD options suggests that in practice employers may be allowing their plan options to be affected by adverse selection. When adverse selection is causing premium differentials between plans to spike, the logical consequence, if people understand their options, should be that the system will unravel toward the lower-coverage option over time. In that case, there is no real value of the “choice” and the real question is whether a menu with only the HD plan is preferable to a menu with only the LD plan.

6.2 Implications When Average Cost Differences are Driven by Moral Hazard

The other possibility is that the differences in average cost between the LD and HD plans are driven not by selection but instead by the causal effects of greater levels of cost sharing on the level of medical care that employees seek. Assume first that the medical spending reductions that result from the increased cost-sharing in HD plans comes entirely from the reduction of wasteful medical services with little value to the individual. In that case, the patterns we document suggest that the frictions that prevent a greater number of people from selecting HD plans, such as inertia and confusion about plan options, result in substantial inefficiencies. Based on our estimates from Section 4, the unexploited savings on wasteful medical spending in this scenario are over \$500 on average and frequently over \$1,000 per employee per year. In fact, if moral hazard is the entire mechanism, our calculations understate the full size of the potential savings at the individual level

because they are based on assuming a fixed distribution of medical spending and the individual would in fact endogenously reduce their spending on low-value services.²³

There is good reason to suspect, however, that some of the reduction in medical spending associated with HD plans involves reductions in *valuable* care (e.g., Brot-Goldberg et al., 2017) and is the result of inefficient *behavioral hazard* (Newhouse, 2006; Baicker et al., 2015). People may find immediate costs for medical care particularly painful or may find it difficult to finance higher out-of-pocket medical spending with HD plans due to budgeting problems and liquidity constraints, even when the premium reductions from an HD plan are enough to cover any increase in out-of-pocket spending. This raises the important possibility that when people enroll in dominated LD plans they are doing it in part because they realize that they would struggle to efficiently access care in a plan with higher deductibles.

The best-case scenario for the efficiency of plan menus at the firms we study may be if the HD enrollees are able to reduce spending efficiently under the HD plan, while the LD enrollees rationally avoid that option because they know it would lead to inefficient medical utilization for them. In that case, the dominated menus we observe represent neither a transfer from LD enrollees to HD enrollees nor inefficiencies due to frictions in plan choice. Instead, the availability of both the LD and HD options allows those who can efficiently reduce spending to do so in the HD plan and those who cannot to remain in a plan with higher insurance coverage.

6.3 Discussion of the Relative Importance of Adverse Selection and Moral Hazard

This section has discussed how the implications of the emergence of dominated health plan menus at firms offering choices between HD and LD plans likely depends on the importance of adverse selection versus moral hazard in driving plan disparities. There is, of course, clear evidence in the literature of both substantial risk selection when employees can choose between plans with different levels of coverage and of causal effects of insurance coverage on the total level of medical spending (see Bundorf, 2016 for a review). The relative importance of these different forces is not something we can disentangle for the firms in the Employer Health Benefits Survey.

²³ In equilibrium the levels of calculated savings would only hold if employees who currently enroll in LD plans would have a similar (or greater) behavioral response to being in the HD plan as those who currently enroll in HD plans. If the LD enrollees were switched into the HD plan but did not change their medical spending, the premiums for the HD plan would rise along the lines of the calculations in Table 6 from the prior subsection for the risk-selection case.

However, in this brief sub-section, we consider what the existing literature on selection and moral hazard in health insurance might suggest for our setting.

Research by Einav et al (2013) suggests that in situations where firms offer plan options, adverse selection may be a larger driver of total spending differentials than moral hazard. They provide evidence from a large employer-sponsored health plan that those who endogenously opt into higher-deductible options on average have lower moral hazard responses than would be observed if higher-deductible options were imposed on the full employee group. In their setting, there is around a \$1,000 total difference in spending between those opting into high and low deductible options and they estimate that around \$130 of that difference would be expected to come from the moral hazard effect.

However, the moral hazard effects needed to explain the premium differences for the firms in our sample are also broadly consistent with estimates from the literature. For example, to explain the entire average differential in total premiums between LD and HD plans we observe, we estimate that the higher cost-sharing in HD plans would need to reduce medical spending by around 10% at the firms that have strictly dominant HD plans.²⁴ We can compare this to Brot-Goldberg et al., (2017), who estimate that a shift from a nearly full-coverage plan to an HD plan with 76% actuarial value at a large employer caused a reduction of 12-14% in total medical spending. While the Brot-Goldberg estimate is consistent with the moral hazard effects needed to explain premium differentials in our sample, it is also likely an upper bound on what we might expect in the Kaiser Family Foundation Sample. The gap in AV between HD and LD plans at the firms with dominant HD plans is much smaller (82% vs 76%) in our sample than in the Brot-Goldberg et al. study. Also, as Einav et al. (2013) highlighted, spending reductions observed for firms where the HD plan is an *option* rather than a requirement would likely be lower. Furthermore, Leive (2018) presents evidence that employees may consider HSA funds as dedicated to medical spending, rather than treating them as more fungible with other wealth, thereby reducing the moral hazard effects for firms that offer substantial HSA contributions along with their HD plans.

²⁴ To get these estimates we start with the total average premium reported for the LD plans. We then estimate how much premiums should fall on average if they simply accounted for the reduction in covered spending when moving to the HD plan, holding fixed the underlying medical utilization. This is equivalent to estimating the ratio of total spending of LD and HD plans, because all total costs difference comes from moral hazard. Using the framework in section 6.1, we calculated the ratio of total premium divided by the ratio of the actuarial value, which reveals how large the moral hazard response would need to be to account for the full premium differential.

Ultimately, it is likely that the dynamics we observe in the plan menus at firms offering different levels of coverage to their employees reflect a mix of the effects of selection and moral hazard. Our reading of the existing evidence suggests to us that adverse selection may be the more important factor overall, but ultimately more research will be needed to better understand the relative importance of adverse selection and moral hazard in determining premium differentials for employer-sponsored health plan options.

7. Conclusion

This study presents an analysis of the insurance value of plan options available to employees at firms that offer both a qualifying high-deductible plan paired with a health savings/reimbursement account and another option with a lower deductible. We find that on average the HD plans offer lower worst-case spending risk and substantial expected savings and risk-adjusted expected savings. Based on simplified representations of plans that transform complex cost-sharing rules into a schedule mapping total medical spending to out-of-pocket costs, we estimate that HD plans strictly dominate in about one third of firms and second-order stochastically dominate LD options at around half of firms. As such, we conclude that at many firms employees do not face a classic tradeoff between expected spending and risk when selecting between health plan options with different coverage levels. Instead, employees who select high-deductible plans at many firms receive a superior net health insurance benefit than employees who select into lower-deductible options.

These findings suggest that those interested in health policy, and in particular the issue of risk pooling across people with different underlying status, may benefit from focusing less exclusively on the private-market exchanges established by the Affordable Care Act (ACA) and more on employer-sponsored insurance. The ACA anticipated that health plans with different levels of coverage (e.g., different deductibles) might attract people with different underlying health status and required that insurers use a “single risk pool” when pricing plans. This implies that premium differences across plans by an insurer reflect the differences in coverage between plans for a common average population. This requirement, and related policies regarding risk adjustments across plans, has received substantial attention, especially during discussions in 2017 around health care reform. Interestingly, however, there is no similar requirement in employer-sponsored health insurance. Employers are prohibited from discriminating against employees on the basis of

health status in their health insurance programs.²⁵ However, there is no regulation governing the premium levels and firm contribution levels across different plans the employer offers to employees. Our results suggest that the principle of a “single risk pool” may often not hold in employer-sponsored health insurance plans that offer multiple plan options.

More research is needed to better understand the extent to which firms actively establish plan options to weaken risk pooling across employees. There is also a need for research to better understand whether offering options with large disparities in net benefits between health plans with different coverage levels has spillovers on labor relations relating to the hiring, retention or wage negotiations with employees of different health status. The results of this paper are also limited to comparisons at firms that offer a very particular mix of high-deductible and lower-deductible options. More research is needed to better understand the average level of benefit disparities experienced by employees across firms throughout the U.S. Given that some firms do not offer choices of plans with different coverage levels, our results may over-state the extent to which health plan options create opportunities for benefit disparities between employees. On the other hand, existing research has documented large differentials (and even dominance) between plan options with varying deductibles that would all be considered low-deductible plans within the Kaiser Family Foundation data. As such, it is possible that situations where plan options varying in coverage level create large net benefit disparities are even more widespread than what we see for the set of firms we analyze here.

Finally, we think the results of this study should encourage those who are investing in research on decision aids that help people to better understand and compare health insurance options. It is quite difficult to compare insurance plan options that differ on the basis of premiums, cost-sharing rules, and even contributions to health savings and reimbursement accounts. Prior to our study one might have reasonably hoped that the ecosystem of human resource departments and benefit consultants and larger employers would have limited the size of financial losses one could incur from not being able to compare plans. Our results show, however, that in employer-sponsored

²⁵ The primary regulation of employer-sponsored health insurance plans comes through the rules imposed by the Employee Retirement and Income Security Act (ERISA). ERISA rules prohibit employers who offer group health insurance plans from discriminating on the basis of employee health status, which generally means that employers charge the same employee premiums to all employees enrolled in the same health plan. The 2010 Affordable Care Act amended ERISA by creating a provision that allows employers to create “health-contingent” wellness programs that can vary the premium and employee pays based on health status. These programs, however, have to follow a number of regulations.

insurance settings being able to make the comparison between plans can often have substantial financial consequences.

References

- Abaluck, Jason, and Jonathan Gruber, "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program," *American Economic Review*, 101(4) (2011), 1180-1210.
- Abaluck, Jason, and Jonathan Gruber, "Improving the Quality of Choices in Health Insurance Markets," National Bureau of Economics Working Paper 22917, (2016).
- Aizcorbe, Ana, Eli Liebman, Sarah Pack, David M. Cutler, Michael E. Chernew, and Allison B. Rosen, "Measuring health care costs of individuals with employer-sponsored health insurance in the U.S.: A comparison of survey and claims data." *Statistical Journal of the IAOS*, 28(1–2) (2012), 43–51.
- Altman, Drew, "The Missing Debate Over Rising Health-Care Deductibles." *The Wall Street Journal*. Dow Jones & Company, 18 September, 2016.
- Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein, "Behavioral Hazard in Health Insurance," *Quarterly Journal of Economics*, 130(4) (2015), 1623-1667.
- Barseghyan, Levon and Molinari, Francesca and O'Donoghue, Ted and Teitelbaum, Joshua C., "The Nature of Risk Preferences: Evidence from Insurance Choices," *American Economic Review*, 103(6) (2013), 2499-2529.
- Bhargava, Saurabh, George Loewenstein, and Justin Sydnor, "Choose to Lose: Health Plan Choices from a Menu with Dominated Options," *The Quarterly Journal of Economics*, 132(3) (2017), 1319–1372.
- Brot-Goldberg, Zarek C., Amaitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad, "What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics," *Quarterly Journal of Economics*, 132(2) (2017), 1261-1318.

Bundorf, M. Kate, “Consumer-directed health plans: do they deliver?” *Policy*, 1(6) (2012).

Bundorf, M. Kate, “Consumer-directed Health Plans: A Review of the Evidence.” *Journal of Risk and Insurance*, 83(1) (2016), 9-41.

Claxton, Gary, Matthew Rae, Nirmita Panchal, Heidi Whitmore, Anthony Damico, Kevin Kenward, and Michelle Long. “Health Benefits In 2015: Stable Trends In The Employer Market.” *Health Affairs*, September 22, 2015, 10.1377/hlthaff.2015.0885.

Claxton, Gary, Matthew Rae, Michelle Long, Anthony Damico, Gregory Foster, Heidi Whitmore, and Lindsey Schapiro. 2016. “Employer Health Benefits 2016 Annual Survey.” *The Kaiser Family Foundation*. Online: <http://files.kff.org/attachment/Report-Employer-Health-Benefits-2016-Annual-Survey>

Einav, Liran, Amy Finkelstein, and Mark R. Cullen. “Estimating Welfare in Insurance Markets Using Variation in Prices.” *The Quarterly Journal of Economics*, 125(3) (2010), 877–921.

Einav, Liran, Amy Finkelstein, Stephen P Ryan, Paul Schrimpf, and Mark R Cullen. “Selection on Moral Hazard in Health Insurance.” *American Economic Review* 103(1), (2013), 178–219.

Ellis, R. P., & Zhu, W.. “Health Plan Type Variations in Spells of Health-Care Treatment,” *American Journal of Health Economics*, 2(4) (2016), 399–430.

Ericson, Keith Marzilli, Philipp Kircher, Johannes Spinnewijn, and Amanda Starc. “Inferring Risk Perceptions and Preferences Using Choice from Insurance Menus: Theory and Evidence.” National Bureau of Economic Research, 2015.

Ericson, Keith Marzilli and Justin Sydnor. “The Questionable Value of Having a Choice of Levels of Health Insurance Coverage,” *Journal of Economic Perspectives*, 31(4) (2017), 51-72.

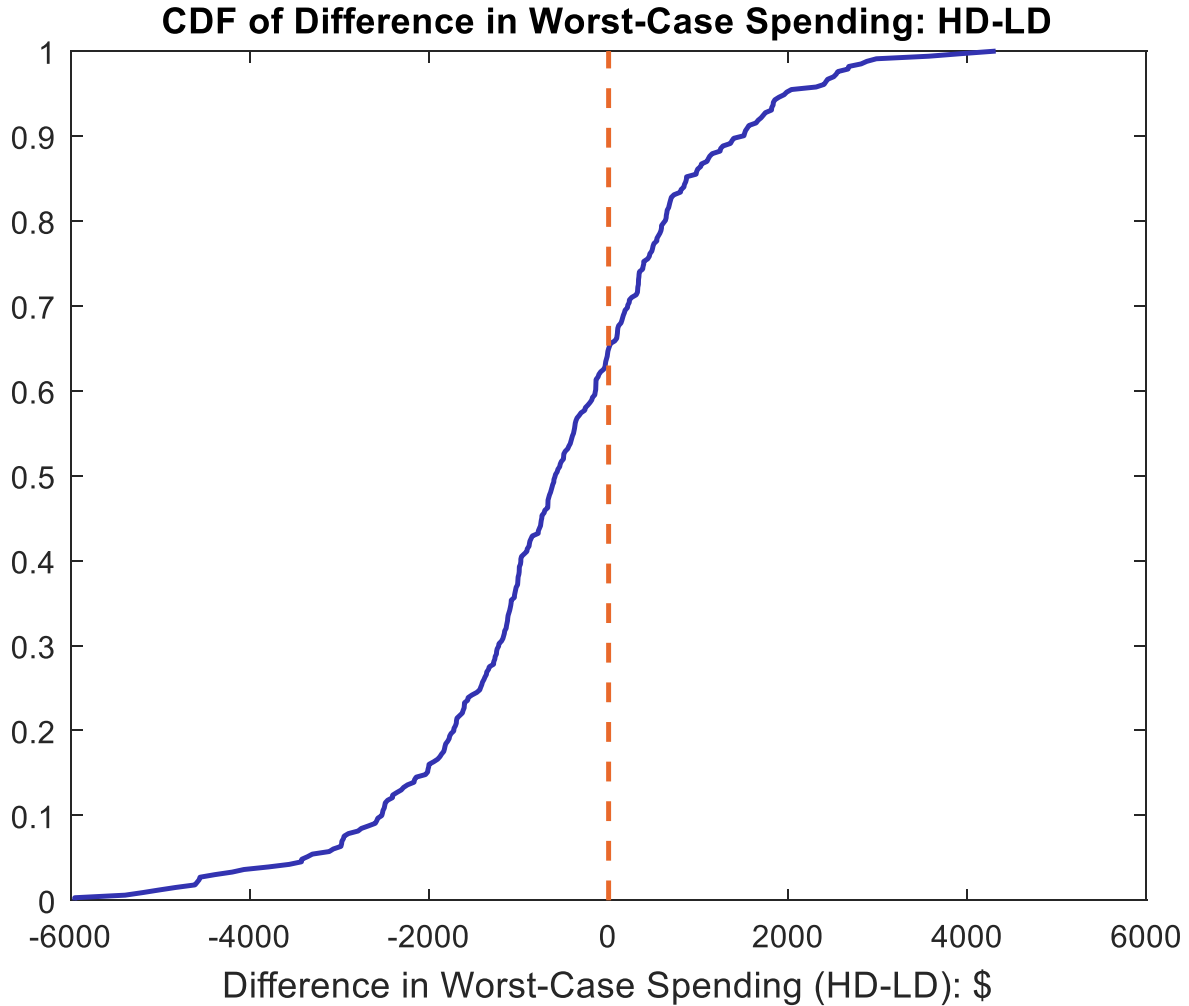
Ericson, Keith Marzilli, and Justin Sydnor. “Liquidity Constraints and the Value of Insurance.” Working Paper, (2018).

Evans, Melanie, Yaryna Serkez and Merrill Sherman. “The The Math Behind Higher Health-Care Deductibles.” *The Wall Street Journal*, 31 Aug. 2017. Accessed September 6, 2017. <http://www.wsj.com/graphics/health-care-coverage/>.

Handel, Benjamin R., “Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts,” *American Economic Review*, 103(7) (2013), 2643-2682.

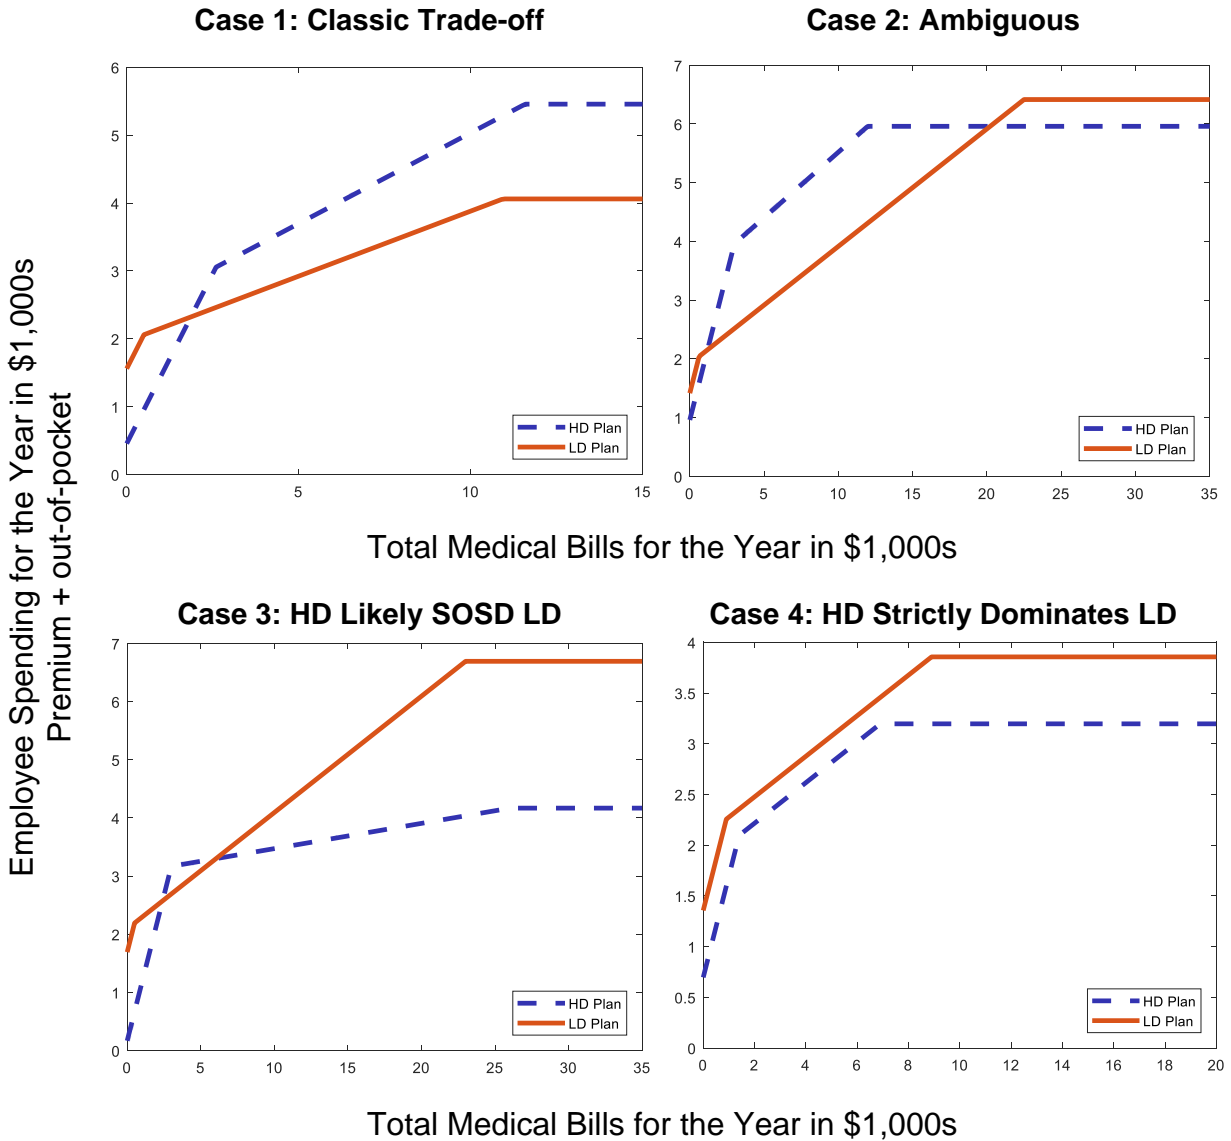
- Handel, Benjamin R., and Jonathan T. Kolstad. "Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare," *American Economic Review*, 105(8) (2015), 2449-2500.
- Haviland, Amelia M., Matthew D. Eisenberg, Ateey Mehrotra, Peter J. Huckfeldt, and Neeraj Sood. "Do Consumer-Directed Health Plans Bend the Cost Curve Over Time?" *Journal of Health Economics*, 46 (2016), 33-51.
- Johnson, Eric J., Ran Hassin, Tom Baker, Allison T. Bajger, and Galen Treuer, "Can Consumers Make Affordable Care Affordable? The Value of Choice Architecture," *PLoS ONE* 8(12) (2013), e81521.
- Leive, Adam. "Health Insurance Design Meets Tax Incentives for Saving: Consume Response to Complex Contracts," Working Paper, (2018).
- Loewenstein, George, Joelle Y. Friedman, Barbara McGill, Sarah Ahmad, Suzanne Linck, Stacey Sinkula, John Beshears, James J. Choi, Jonathan Kolstad, David Laibson, Brigitte C. Madrin, John A. List, and Kevin G. Volpp, "Consumers' Misunderstanding of Health Insurance," *Journal of Health Economics* 32(5) (2013), 850-862.
- Newhouse, Joseph P. "Reconsidering the Moral Hazard-Risk Avoidance Tradeoff." *Journal of Health Economics*, 25(5) (2006), 1005-1014.
- Samek, Anya, and Justin Sydnor, "Simplifying Health Insurance with Consequence Graphs," Working Paper, 2017
- Sydnor, Justin. (Over)insuring Modest Risks. *American Economic Journal: Applied Economics*, 2(4) (2010), 177-199.
- Willis Towers Watson. "2015 Best Practices in Health Care Employer Survey," 2015. Retrieved from <https://www.towerswatson.com/en-US/Insights/IC-Types/Survey-Research-Results/2015/11/full-report-2015-towers-watson-nbgh-best-practices-in-health-care-employer-survey>.

Figure 1. Distribution of Worst-Case Spending Difference: HD – LD



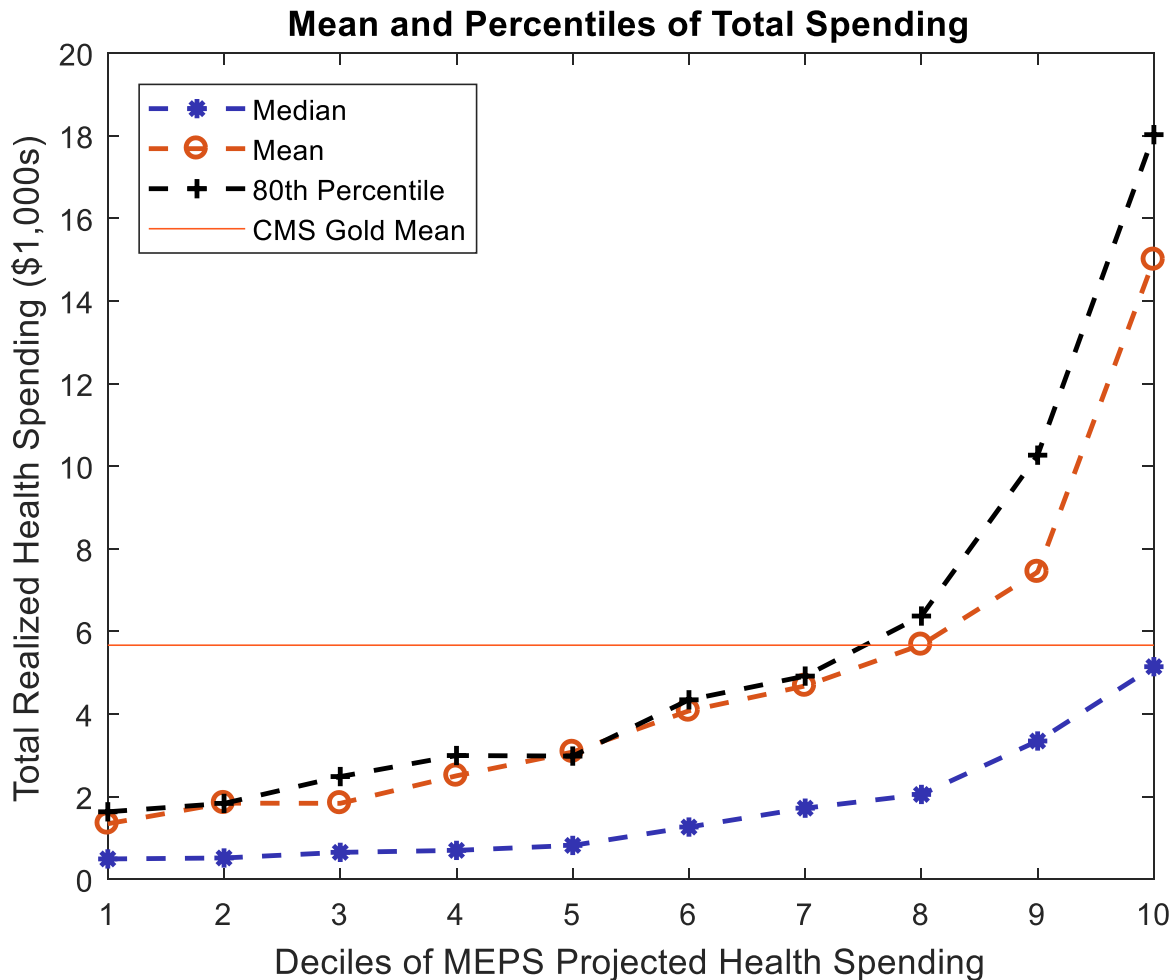
Notes: This graph plots the CDF of worst-case spending differences across firms, such that the y-axis shows the fraction of firms with a worst-case spending difference (HD – LD) smaller than x. For each plan, we calculate the worst-case spending as the annual premium paid by the employee plus the maximum out-of-pocket limit on the plan. For high-deductible (HD) options, we subtract from that amount any contribution the firm makes to a health savings account or health reimbursement account. A negative (positive) number means the HD plan has lower (higher) worst-case spending than the LD plan at that firm.

Figure 2. Illustration of Dominance Classifications



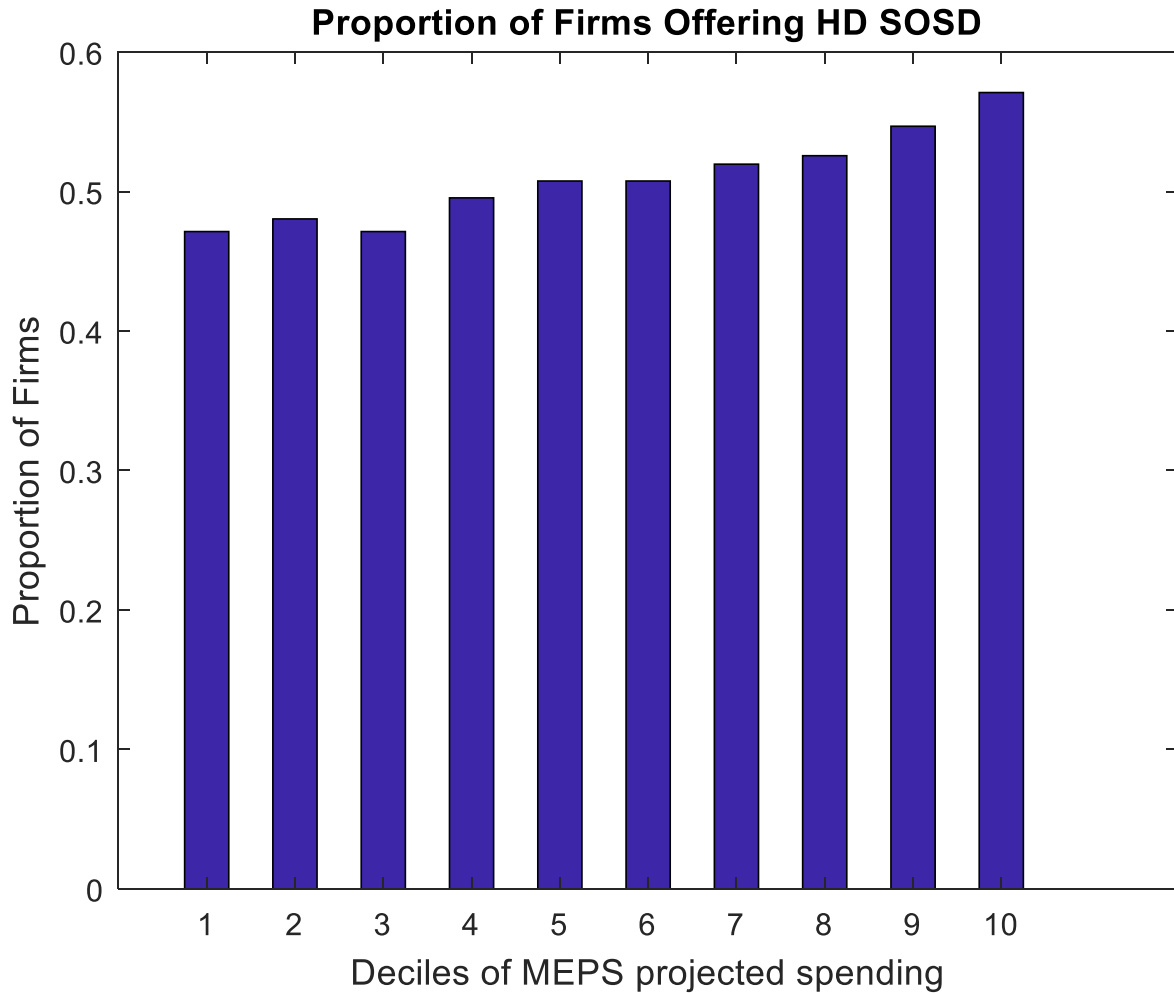
Notes: Illustration of different types of plan comparisons. We transform each plan into a simplified plan based on deductible, coinsurance and maximum out of pocket. We then classify them based on the relative size of net premium (the y-intercept), net premium + maximum out of pocket (the flat range) and the cost sharing.

Figure 3. MEPS Projected Health Spending



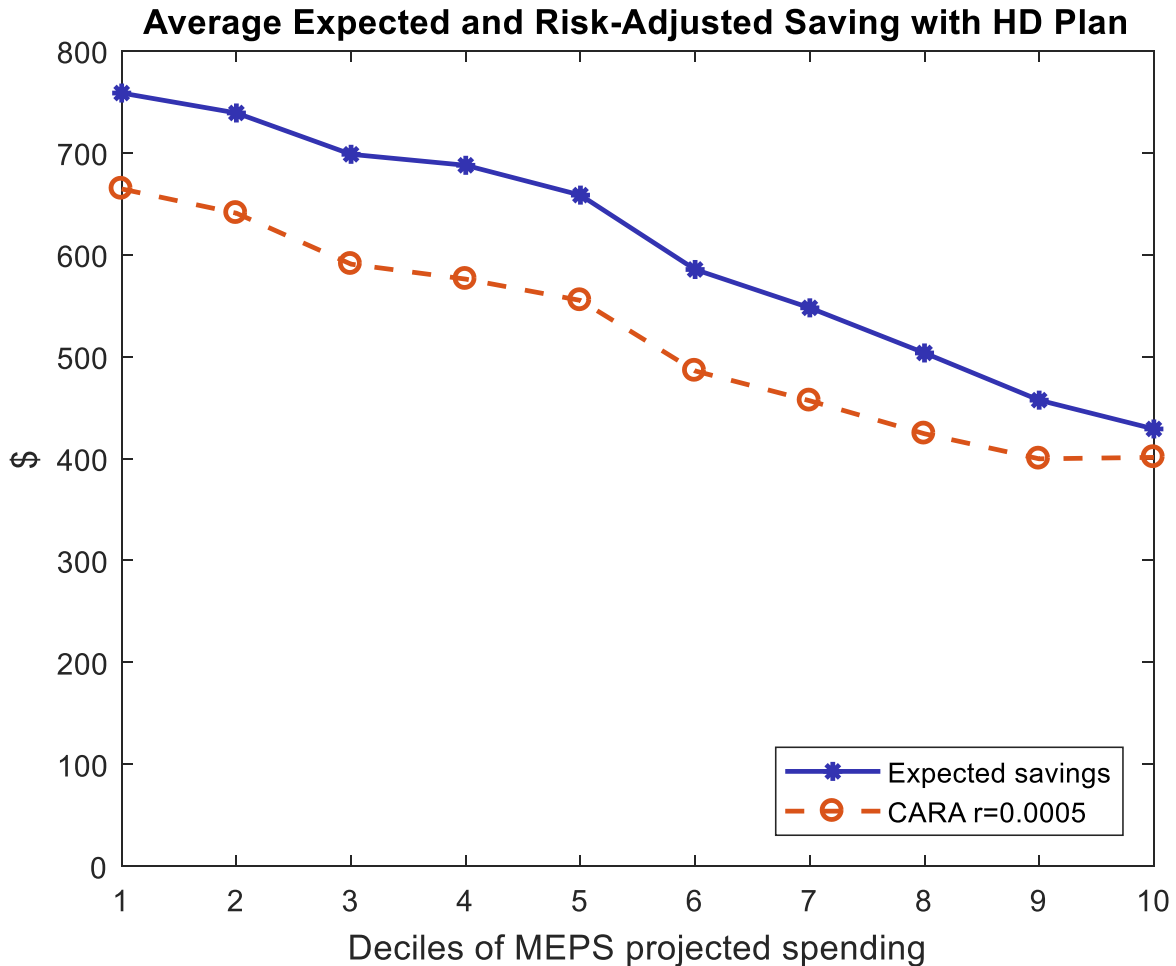
Notes: This graph characterizes the realized health spending for ten groups of people defined by deciles of projected health spending as described in the text. This graph uses data on total spending from the Medical Expenditure Panel Survey (MEPS) in 2014 for the population covered by employer/group insurance throughout the year. We regress total medical spending on a fully interacted model including age, gender, and self-reported health status (as explained in the text) and generate predicted levels of total spending based on those characteristics. We then group MEPS participants into 10 groups (deciles) based on their predicted level of medical spending from this regression, which gives us 10 groups with different ex-ante health risk. The graph provides the mean and percentiles of realized spending for people in each of these 10 MEPS deciles.

Figure 4. Proportion of Firms where HD Second Order Stochastically Dominates LD for Different Spending Distributions




Notes: For each of the deciles of MEPS projected spending, we use the empirical distribution of realized spending for people in that decile and compare whether for that distribution the HD plan second order stochastically dominates the LD plan or not. The graph shows the proportion of firms (across the 331-firm sample) offering HD plan that is second order stochastically dominant for the distribution in each MEPS decile.

Figure 5. Expected and Risk-adjusted Average Savings from Choosing the HD Plan over the LD Plan for Different Spending Distributions



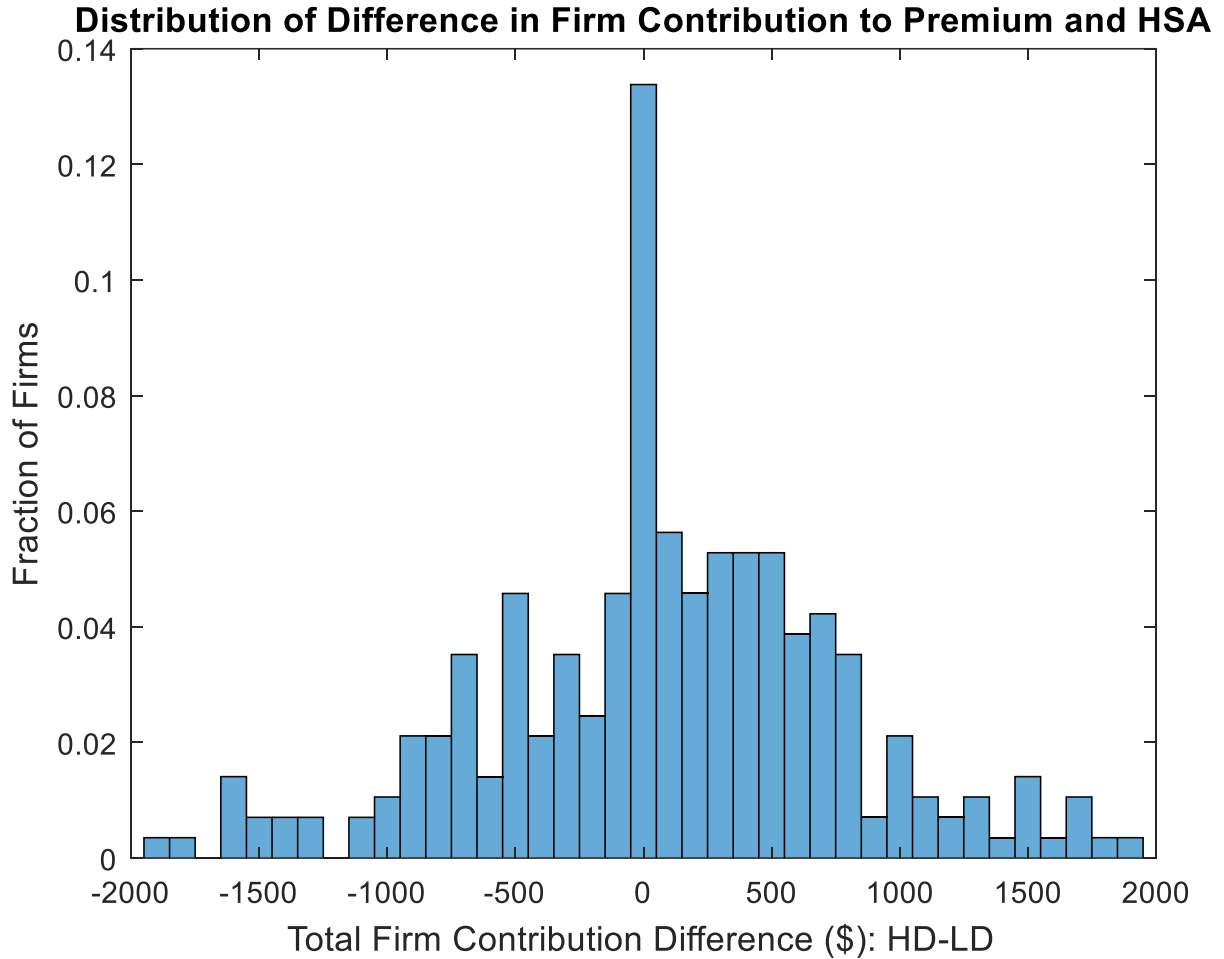
Notes: We calculated the expected savings or risk-adjusted expected savings from choosing the HD plan over the LD plan for each of the 331 firms. We do this separately for each of the 10 different realized spending distributions defined by the deciles of predicted values in the MEPS. The lines then graph the average values of those savings across the 331 firms. The risk-adjusted savings are created by assuming Constant Absolute Risk Aversion Utility with an $r = 0.005$. The risk-adjusted savings is then the amount of money a person facing that particular spending distribution would need to be given with certainty to remain in the LD plan and have the same expected utility as instead switching to the HD plan.

Figure 6. Average Differences in Premiums and Expected Out-of-Pocket Costs by Deductible Difference



Notes: Total premium is the annual premium for single coverage, including contribution from the employer and employee. Net premium is the premium paid by employee, minus any HSA contribution from the firm. Expected out-of-pocket costs is the expected enrollee out-of-pocket spending under the CMS Gold distribution. We subtract any HRA contribution by the firm, as long as the out-of-pocket spending is non-negative. We then calculate the difference of the first two variables as LD value minus HD value, and the difference of expected out-of-pocket costs as the HD value minus LD value. For firms where the HD plan has a strictly higher deductible (309 firms), we put them into 4 groups based on deductible difference: (0,\$750],(\$750-\$1250],(\$1250,\$1750],(\$1750,larger). For each group, we plot the mean and also the 95% confidence interval on the mean (the error bar at each data point).

Figure 7. Distribution of Difference in Total Firm Contributions to Plans



Notes: This graph shows the distribution of the difference in total firm contribution to health plans (HD – LD) for firms in the sample that offer an HD plan with an HSA account (284 firms). For the LD plans this is simply the firm’s contribution toward the premium on the plan. For the HD plan it is the sum of the firm’s contribution to premium and the firm’s contribution to the HSA account. The bin size is set to \$100. So, for example, the bar centered at 0 shows that at just under 15% of these firms the total contribution difference is within \$50.

Table 1. Summary Statistics

	Full Sample of Firms Offering at Least One ...		Matched Sample Offering Both HD and Non-HD	
	LD Plan	HD Plan	LD Plan	HD Plan
Panel A: Plan Level Variables				
Deductible	\$711 (\$900)	\$2157 (\$1004)	\$846 (\$813)	\$2166 (\$986)
Maximum Out-of-Pocket	\$3282 (\$1672)	\$4129 (\$1370)	\$3464 (\$1488)	\$4011 (\$1354)
Premium	\$6700 (\$1904)	\$5668 (\$1639)	\$6482 (\$1748)	\$5453 (\$1463)
Percentage of Worker's Contribution to Premium	0.21	0.17	0.23	0.16
Worker's Contribution to Premium	\$1349 (\$1008)	\$911 (\$728)	\$1450 (\$825)	\$841 (\$616)
Self-insured	0.55	0.64	0.71	0.74
Share of Firm offering HSA		0.78		0.86
HSA Contribution Conditional on Offering		\$505 (\$501)		\$473 (\$404)
Share of Firm Offering HRA		0.22		0.14
HRA Contribution Conditional on Offering		\$1068 (\$909)		\$758 (\$583)
Actuarial Value*	85% (7%)	75% (6%)	83% (6%)	74% (5%)
Number of plans	1663	709	331	331
Panel B: Firm Level Variables				
Firm with More Than 1000 Employee	0.51	0.57	0.64	
More than 35% Employee Age 50 or More	0.47	0.47	0.46	
More than 35% Employee Earn \$58k or More	0.45	0.51	0.50	
Private For-Profit	0.56	0.65	0.63	
Have Union Worker	0.31	0.29	0.31	
Number of firms	1332	709	331	331

Notes: Means with standard deviations in the parentheses. Data from Kaiser Family Foundation Employer Benefits Survey (2015). *Authors' calculation using CMS 2015 Actuarial Value Calculator Gold tier. All plans in the matched sample (Columns 3 & 4) have calculated AV values, while Columns 1 & 2 report AV for a subsample (1,259 firms) with plan features compatible with the AV calculator.

Table 2. Strict Dominance Classification for Simplified Plan Representations

Fig 2. Case	Plan w/ lower net employee premium	Plan w/ lower net premium + max out of pocket	Strictly Dominant Plan	Number of Firms	Fraction of Firms	Avg Enrollment Share in HD Plan
NA	LD	LD	LD	8	2%	46%
NA	LD	LD		3	1%	56%
NA	LD	HD		9	3%	41%
Case 1	HD	LD		106	32%	25%
Case 2&3	HD	HD		84	25%	32%
Case 4	HD	HD	HD	121	37%	47%
Total				331	100%	36%

Notes: The premium is employee paid premium for single coverage of a year, minus firm contribution into HSA account if any. Max out of pocket is the maximum out-of-pocket per enrollee for single coverage of a year, minus firm contribution into HRA account if any. Net premium is the net spending when enrollee uses no medical services (the best-case scenario), and net premium + net max out of pocket is the spending when enrollees uses infinite medical services (the worst-case scenario). HD Strictly Dominant are firms offering an HD plan with lower employee spending under 84 atoms of Gold tier total medical spending distribution from CMS AV Calculator data. The enrollment share is only illustrative: the share is for all HD plans offered by the firm, while we only observe plan information of the one with the largest enrollment. Data from Kaiser Family Foundation Employer Benefits survey (2015).

Table 3. First and Second Order Stochastic Dominance Classification

	FOSD	SOSD
LD dominates	2%	11%
Neither dominates	61%	33%
HD dominates	37%	55%
Total	100%	100%

Notes: The table shows the fraction of the 331 firms in our sample offering menus where either LD dominates, neither dominates or HD dominates. First and second order stochastic dominance calculations are based on comparisons of simplified representations of the plans, using employee spending under 84 atoms of Gold-tier total medical spending distribution from CMS AV Calculator data.

Table 4. Average Savings and Risk Adjusted Savings with the HD Plan

	Expected Employee Savings	Risk-adjusted Net Savings with CARA utility function	
		r=0.0005	r=0.002
All plans n=331	\$569 (\$734)	\$471 (\$832)	\$440 (\$1275)
Plans where HD is strictly dominant n=121	\$1121 (\$723)	\$1127 (\$739)	\$1334 (\$936)
Plans where HD is not strictly dominant n=210	\$251 (\$521)	\$93 (\$623)	\$-75 (\$1156)
CARA in context - Cert. Equiv. for 50% chance at winning \$1k with that r:		\$438	\$283
CARA in context - Gain needed to accept 50/50 Lose \$1k, gain G with that r:		\$2092	∞

Notes: Means with standard deviations in parentheses. Expected utility calculation is based on constant absolute risk aversion (CARA) model. The risk-adjusted savings numbers show the average equivalent money amount the enrollee would be willing to pay to switch from the LD plan to the HD plan. The underlying health spending distribution is the Gold tier distribution from CMS AV Calculator data. Plan data is from Kaiser Family Foundation Employer Health Benefits Survey (2015).

Table 5. Dominance Pattern and Firm Demographics

Demographic Split	N	Share with HD Plan Strictly Dominant	Difference
Firm with More Than 1000 Employee	119	0.32	-0.07
Firm with Less Than 1000 Employee	212	0.39	(0.06)
More than 35% Employee Earn \$58k or More	167	0.36	-0.01
Less than 35% Employee Earn \$58k or More	164	0.37	(0.05)
More than 35% Employee Age 50 or More	179	0.38	0.03
Less than 35% Employee Age 50 or More	152	0.35	(0.05)
Have Union Worker	228	0.37	0.02
No Union Worker	103	0.35	(0.06)
HD Self-Insured	87	0.35	-0.03
HD Underwritten by Other Firm	244	0.37	(0.06)
Private For-Profit	121	0.37	0.01
Private Non-Profit/Public	210	0.36	(0.06)
Analysis in Anticipation of Cadillac Tax	222	0.39	0.07
No Analysis in Anticipation of Cadillac Tax	109	0.32	(0.06)

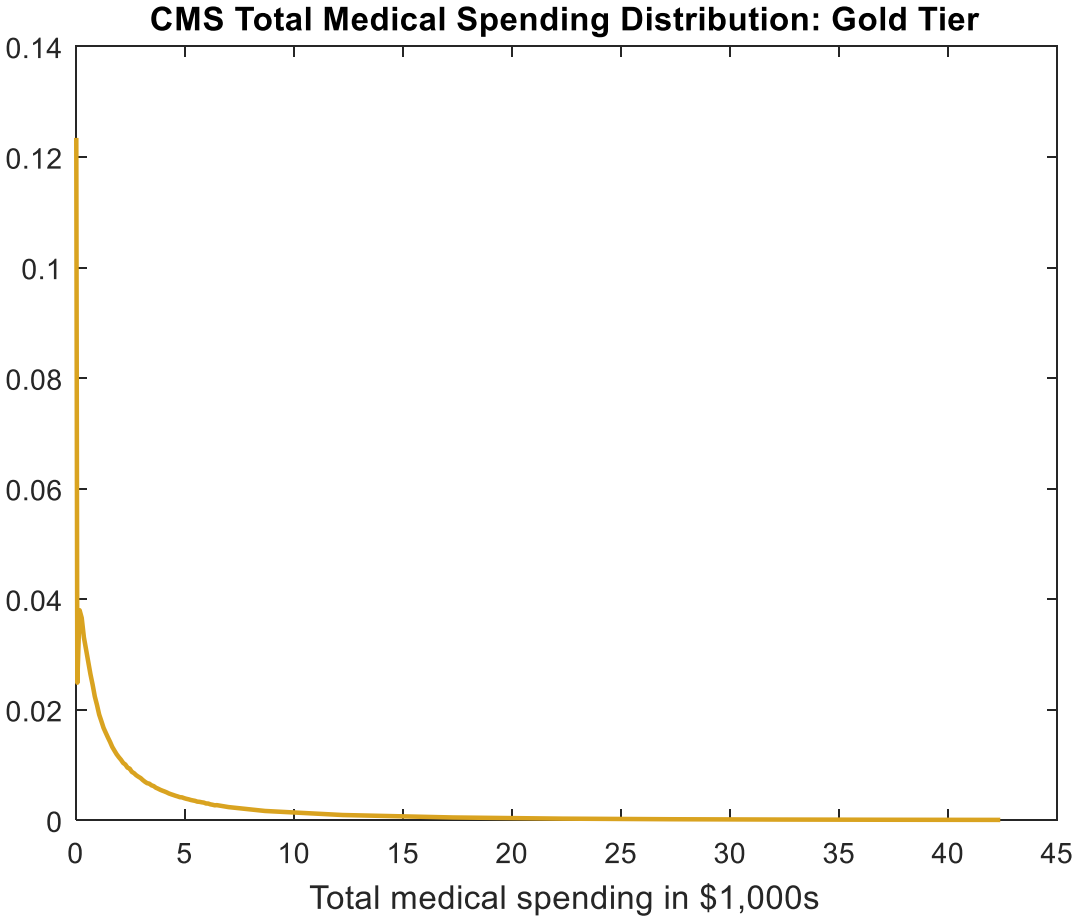
Notes: Standard error of difference in parentheses.

Table 6. Plan Premiums in Counterfactual with Only the HD Option

	Original LD	Original HD	New HD-only
Total Premium	\$6,483 (\$1690)	\$5,391 (\$1476)	\$5,612 (\$1386)
Firm Contribution to Premium/HSA	\$5,047 (\$1755)	\$5,063 (\$1650)	\$5,077 (\$1596)
Employee Net Premium	\$1,436 (\$816)	\$328 (\$756)	\$536 (\$748)

Notes: This table uses the subsample of firms offering HSA accounts for their HD plans (n = 284). We assume a proportional loading factor of 1/0.85. The first two columns report the average for the observed plans at these firms. The third column shows our calculation of the implied total plan premium, firm contributions, and employee net premium in the counterfactual where only the HD plan is offered. This counterfactual assumes that original differences in plan premiums in excess of those predicted by the actuarial-value difference between plans are generated by differences in the ex-ante spending needs of the enrollees in the two plans (i.e., adverse selection) as described in Section 6.1.

Appendix Figure B1. CMS Total Medical Spending Distributions



Notes: Data from continuance tables provided by CMS for estimating the actuarial value of plans in the Affordable Care Act exchanges. The distribution corresponds to Gold metal tier. The plot shows the fraction of individuals with spending within a \$100-bin of spending. The intercept shows the fraction of individuals with zero spending.

Table A1. Sample Construction

	Number of firms	Number of plans
1 Firms reported offering at least one health plan	1771	2714*
2 Firms with complete and consistent plan information (full sample)	1529	2372
3 Firms offering 1 HD plan and 1 LD plan	373	746
4 Firms with plan information consistent with CMS AV Calculator (matched sample)	331	662

Notes: Each row represents a subsample of the above row. *In Row 1 we count the number of plans with any plan information. In some cases, the information is incomplete, which may imply that these are not actually available plans.

Table B1. Robustness Checks

		Number of Firms	Share of HD Strictly Dominant	Average Expected Net Savings with HD plan
B1a	Basic Analysis	331	37%	\$569 (\$734)
B1b	Sample without Judgement Calls	166	35%	\$580 (\$730)
B1c	Sample without Imputation	280	38%	\$572 (\$687)
B1d	Sample without Inpatient Copay	275	36%	\$594 (\$751)

Notes: Standard deviation is in parentheses. The average expected savings is calculated using the CMS Gold-tier continuance table distribution of total medical spending. We made reasonable judgment calls when calculating the actuarial value, as described in Appendix B. B1b takes the subsample of firms without any such judgement calls. KFF data contain ~5% imputation. B1c removes firms with any form of imputation. B1d includes firms that have no copayment as cost-sharing for hospital stays for both the LD and HD plans. (They may have coinsurance rate for hospital stay though.) Data is from Kaiser Family Foundation Employer Health Benefits Survey (2015) and CMS AV Calculator Gold continuance table.

Appendix A Sample Construction

Appendix Table A1 gives a breakdown of our sample construction. There are a total of 1,771 firms that responded to the 2015 version of the Kaiser Family Foundation Employer Health Benefits Survey and reported offering at least one health insurance plan to their employees. Among these firms, we dropped firms with missing or contradictory plan information for any of the plans they reported on, leaving us with 2372 plans offered by 1529 firms. This is our full sample.

In the cleaning step going from firms in category 1 (firms reported offering at least one health plan) and to the subsample of firms in category 2 (firms with complete and consistent plan information), we dropped the firms with any health plan with the following features:

1. Missing plan information
 - a. No cost-sharing information
 - b. No maximum out-of-pocket amount
2. Contradictory plan information
 - a. Plan information not consistent with plan type variable
 - b. HDHP with neither health savings account (HSA) nor health reimbursement arrangement (HRA) information. (By definition of the survey, plans must have either HSA or HRA to be classified as HDHP.)
 - c. Deductible amount is larger than maximum out-of-pocket amount
 - d. Deductible amount is the same as maximum out-of-pocket amount, but there is cost-sharing after the deductible range
 - e. Deductible amount is smaller than maximum out-of-pocket amount, but there is no cost-sharing after deductible for any service
3. Potential data error
 - a. Deductible amount or maximum out-of-pocket amount is not divisible by 5

We also confirmed that the remaining HD plans satisfy the IRS regulation on HDHP.¹

¹ For HDHP, there is regulation on a) caps of maximum out-of-pocket value for HDHP (for individual coverage, \$6450 in 2014 and \$6600 in 2015) and minimum deductible value (for individual coverage, \$1250 in 2014, \$1300 in 2015).

From the main sample we then restrict our sample to firms reporting details for 1 HD plan and 1 LD plan, which is a sample of 373 firms with 746 plans. To be consistent with later analysis, we further dropped 42 firms with at least one plan that has cost-sharing features that are not compatible with the Centers for Medicare and Medicaid Services (CMS) Actuarial Value Calculator. This leaves us with 331 firms offering 662 plans. This is our analysis sample (matched sample).

The cleaning step going from firms in category 3 (Firms offering both 1 HD and 1 LD) to the subsample of firms in category 4 (Firms offering both 1 HD and 1 LD and features consistent with AV calculator), we dropped firms with any health plan with the following features:

1. Maximum out-of-pocket amount larger than \$6850 (this is the upper bound imposed by CMS Actuarial Value calculator) [1 Firm; 1 LD plan]
2. There is a copayment for outpatient surgery (CMS AV calculator only supports coinsurance rate for this type of services). [41 Firms; 40 LD plans; 6 HD plans]

Appendix B Creating Simplified Plan Representations

In our main analysis, we create a simplified plan representation for each plan (deductible, coinsurance and maximum out-of-pocket limit). In our sample, there are 146 plans that are originally simple plans. For the rest, we follow the approach in Ericson et al. (2015) to create a simplified plan representation using the CMS AV Calculator.

When calculating actuarial value and the equivalent coinsurance rate, we ignore the following features reported in the Kaiser Family Foundation Survey data that are not supported by AV Calculator:

1. maximum and/or minimum limit for coinsurance payment
2. separate deductible for hospital and/or outpatient surgery
3. If the cost-sharing rule is paying both coinsurance and copayment, we treat it as paying whichever is higher. The former is not supported by CMS AV calculator.

For drug cost sharing features, the KFF survey uses different category labels than the CMS AV calculator. We make the following assumptions in mapping the KFF data to the CMS AV calculator:

1. We map tier 1-4 cost-sharing rule for drugs in the KFF data to generics, preferred brand drugs, non-preferred brand drugs, and specialty drugs in the CMS AV calculator respectively.
2. If a plan has a 2-tier drug cost-sharing rule, we use the cost sharing information for the first tier for tier 1 drug, and second tier sharing rule for 2-4 tiers. We do the same thing for 1-tier and 3-tier plans.
3. The only exception to 2) is when the plan does not cover specialty drugs. If that is the case, we indicate tier-4 cost sharing rule as not covered by the plan.

As mentioned in the main text, there are a few complications to the basic approach. Specifically, about 40% of plan designs have office visits copays before the deductible and/or separate deductibles for drug coverage. To get the simplified plan representation for these plans, we follow the procedures below.

We split medical spending into three categories: office visits, prescription drugs, and all other services. Following the approach in Ericson et al. (2015), we calculate out-of-pocket spending given total spending \mathbf{x} in the following steps:

Step 1: determine which services are subject to the general deductible. Index them with g . Split \mathbf{x} into spending on each service using information from the Gold continuance table from CMS AV Calculator. Denote the amount of spending on services subject to the general deductible as \mathbf{x}_g . For example, if drugs are not subject to general deductible, and represents 20% of total spending at \mathbf{x} , then g represents all services other than drug, and $\mathbf{x}_g = (1 - 0.2)\mathbf{x} = 0.8\mathbf{x}$.

Step 2: input all plan information into AV calculator and get the actuarial value for the plan. We show the details of how we map KFF data into the AV Calculator below. In this and the next step, we set the underlying metal tier as Gold. (We also tried other metal tiers, the numbers are very close.)

Step 3: keep cost-sharing feature of services not subject to the deductible as given, remove all other copay or coinsurance (i.e., on services subject to the general deductible), and input the deductible and MOOP of the plan. Find a single coinsurance rate paid by enrollee that gives the same AV as calculated in step 2. We label it as co_g . The idea is to convert complicated cost sharing feature into an equivalent single coinsurance rate. If all services subject to general deductible have the same coinsurance rate, then co_g will be that number. Intuitively, co_g is the average amount an enrollee needs to pay for each dollar of medical services after the deductible is met for services subject to the general deductible.

Step 4: calculate the out-of-pocket for services subject to the general deductible at each level of total spending \mathbf{x} as:

$$oop_g = \mathbb{1}(\mathbf{x}_g > ddct_g) \cdot co_g \cdot (\mathbf{x}_g - ddct_g) + \min(\mathbf{x}_g, ddct_g).$$

Step 5: If office visits are not subject to the general deductible, calculate the out-of-pocket spending on office visits oop_{ov} . Let \mathbf{x}_{ov} denote the portion of total spending on office visits. For example, if at total spending level \mathbf{x} , 10% is spent on office visits, then $\mathbf{x}_{ov} = 0.1\mathbf{x}$. Let co_{ov} denote the equivalent coinsurance rate for office visits. If the plan has coinsurance rate

for office visits, then co_{ov} is that coinsurance rate. If the plan has copayment for office visits, then

$$co_{ov} = \frac{copay * \# \text{ of visits per year}}{x_{ov}}.^2$$

Then the out-of-pocket spending on office visits oop_{ov} is

$$oop_{ov} = x_{ov} \cdot co_{ov},$$

Step 6: In the KFF EHBS and AV Calculator, drugs are split into four tiers. If drugs have a separate deductible, calculate the drug cost sharing co_d for each tier in the same way as for office visits (Step 5). Then the drug out-of-pocket spending is:

$$oop_d = \left(\sum_{i=1}^4 \mathbb{1}(x_d > ddct_d) \cdot co_{di} \cdot x_{di} \right) + \min(x_d, ddct_d),$$

where x_d is the total drug spending at total spending level x , x_{di} is the amount spent on each tier beyond the deductible range. We determine x_{di} based on the proportion of spending on each tier. For example, if there is a drug deductible of \$100, and at total spending level x , the spending on drugs tier 1-4 is \$30, \$40, \$30, and \$20 respectively, and exceed the deductible by \$20. Then we calculate $x_{d1} = 20 * \frac{30}{40+30+30+20} = \5 . The calculations for other tiers are similar.

Step 7: Then the total spending over the year is:

$$p + oop = p + \max\{\min(oop_g + oop_{ov} + oop_d, m) - HRA, 0\}.$$

p is the net premium (employee premium minus any firm contribution to HSA), m is the MOOP of the plan, HRA is the amount that the firm contributes into health reimbursement arraignment. According to the definition of HRA contribution, we treat this value as a reduction of out-of-pocket cost: if an enrollee has no out-of-pocket cost, he cannot use the money. He also cannot withdraw the money or take the money if he leaves the firm.

² In a few rare cases, this rate is higher than 1, implying that the copayment amount is higher than x_{ov} . When this happens, we replace co_{ov} by 1. This is consistent with the AV Calculator methodology.

Appendix C Calculating Counterfactual Premium of HD-only Menu

To get the estimate of the counterfactual total premium for the HD plans if all employees were shifted in the HD plan, we need an estimate of the underlying expected medical costs differences between the LD-plan enrollees and the HD-plan enrollees. The KFF survey data does not have direct information about medical claims costs, but we can get an approximation under the assumption that total premiums P_j reported in the survey are based on the following basic formula:

$$P_j = \gamma AV_j \mu_j,$$

where γ is a proportional loading factor, AV_j is the actuarial value of plan j , and μ_j is the average expected level of medical spending for the employees who enroll in plan j . We use that formula to solve for an estimate of the average medical costs of the employees enrolled in each plan type by dividing the total premiums reported for each plan by the plan's actuarial value and a loading factor. For the actuarial value we use the CMS Actuarial Value Calculator estimates as in the main text. For the loading factor, we use a common factor of 1.18 for all plans. That loading is consistent with a medical loss ratio of 85%, which is the regulated minimum ratio allowed for large group plans under the Affordable Care Act.

After getting the average expected level of medical spending μ_H and μ_L , we can calculate the new HD premium as

$$P_{NewHD} = \gamma AV_H (\mu_H * s_H + \mu_L * s_L),$$

where s_H and s_L represent the shares of employees enrolling in each plan category. The weighted average $(\mu_H * s_H + \mu_L * s_L)$ represents the expected medical costs for the full population of employees at each firm. This estimate is, of course, only an approximation and relies on a number of assumptions. Nonetheless, we believe this exercise provides a useful estimate for evaluating the size of the transfer from LD to HD enrollees.