Health Insurance Design Meets Tax Incentives for Saving: Consumer Responses to Complex Contracts

Adam Leive*

University of Virginia

June 2018

Abstract

To lower health care costs, health savings accounts (HSAs) offer tax incentives encouraging people to trade off current health care consumption against future income. In this paper, I test whether consumers use HSAs as self-insurance or to reduce their current health insurance deductible (supplemental insurance). I study this question in the context of a large U.S. health insurer that fully replaced its traditional, low-deductible health insurance offerings with a menu of high-deductible health plans and HSAs for its own employees. Using discontinuities in the employer's matching rate for HSA contributions, I estimate a marginal propensity to consume from HSA assets of 0.68. There is no evidence that HSAs crowd out 401(k) saving. The large majority of employees fail to treat HSA money as fungible with retirement saving and make dominated choices costing them $465 annually in foregone tax benefits, on average. In this setting, health spending did not decline as HSA contributions functioned largely as supplemental insurance. To interpret these empirical patterns, I develop a model that incorporates the continuous choice of HSA saving with the discrete choice of deductibles (*JEL* D14, D81, G22, H51, I13).

*Batten School of Leadership & Public Policy, University of Virginia, 235 McCormick Road, Charlottesville, VA 22904 (email: leive@virginia.edu). Funding from the National Institute of Aging through Grant Number T32-AG000186 to the National Bureau of Economic Research, from the Social Security Administration through BC Grant #5002099 to the Center for Retirement Research at Boston College, and from the Ackoff Fellowship from the Wharton Risk Management and Decision Processes Center is gratefully acknowledged. I am grateful to my dissertation advisors Iwan Barankay, Jonathan Kolstad, Mark Pauly, and Robert Town for guidance, and to helpful comments and suggestions from Zarek Brot-Goldberg, Guy David, Hanming Fang, Matt Grennan, Scott Harrington, Jesse Hinde, Bill Johnson, Ambar La Forgia, Olivia S. Mitchell, Elena Prager, Preethi Rao, Jonathan Skinner, Amanda Starc, Ashley Swanson, Richard Thaler, Tess Wilkinson-Ryan, Karen Zhang and seminar participants at CBO, Johns Hopkins, Miami, Minnesota, Northwestern, Virginia, Yale, APPAM, SEA meetings, and the Boulder Summer Conference on Consumer Financial Decision-Making. I thank Andrew Chong and Colin Gray for help in processing the data.

1 Introduction

Health Savings Accounts (HSAs) aim to curb rising health care costs by offering tax incentives to trade off current health care consumption against future income. HSAs function like a 401(k) retirement plan, with the additional feature that withdrawals for health care expenses are always tax-exempt. Such accounts must be paired with a high-deductible health plan (HDHP), with the objective of exposing consumers to the marginal cost of care for moderate expenses. In theory, linking a personal savings account to current benefits brings insurance closer to the first-best contingent-claims contract that smooths consumption without causing moral hazard (Arrow 1963; Zeckhauser 1970). This focus on self-insurance is central to proposed redesigns of other social insurance programs, including unemployment insurance accounts and Social Security privatization (Feldstein 2005; Feldstein and Altman 2007; Setty 2017). Yet HSAs have more complex features than other health insurance products. Consumers may be unaware HSA balances carry forward, heavily discount the future, or fail to recognize the substitutability between health care and other goods from HSA saving. Such behavior could undermine the cost-control incentives of these contracts.

In this paper, I test whether consumers use their HSA as self-insurance or to finance the current insurance deductible (supplemental insurance). Answering this question is challenging due to data requirements and because it requires analyzing a complex life-cycle savings problem with current insurance choices. First, one must observe consumer saving patterns out of an exogenous increase in HSA funds. Second, how consumers treat their HSA must be distinguished from risk preferences and time preferences. Finally, understanding the welfare implications of these accounts requires taking a stand on the optimality of saving decisions, which is difficult because many life-cycle factors are unobserved by the econometrician (Choi et al. 2011).

I examine this question by studying choices of employees at a large U.S. health insurer that fully replaced its traditional, low-deductible health insurance offerings with a menu of

HDHPs and HSAs for its own workers. I combine detailed panel-level data and variation in the employer's insurance offerings and matching rules to overcome these empirical challenges. The data includes insurance deductible choices, contributions by both employees and employer to the HSA and 401(k), medical and pharmacy claims, demographics, and information on salary and job characteristics. I exploit variation in the employer's HSA matching rates by salary level to identify the marginal propensity to consume (MPC) from HSA funds using a regression-discontinuity design. To separate the MPC from other preference parameters, I use variation in deductible levels to pin down risk preferences and observed 401(k) saving decisions to account for preferences over future consumption. For a sub-sample of employees, the similarities between the HSA and 401(k) permit an assumption-free test on whether HSA saving decisions reflect mistakes: employees whose 401(k) contributions exceed the employer match should max out their HSA since the HSA's tax incentives dominate those of the 401(k) past this level.

Consumers in this setting use their HSA largely as supplemental insurance, rather than self-insurance. I estimate an MPC of 0.68 from HSA funds. There is no evidence that HSAs crowd out 401(k) saving. The absence of crowd-out is likely due to consumer misunderstanding: 90 percent of employees make dominated saving decisions by failing to recognize the fungibility between their HSA and 401(k). The magnitude of these optimization errors is sizable. On average, over $1,000 of employee 401(k) contributions are dominated annually, and employees would unambiguously be better off if they re-allocated this saving to their HSA. The average size of these foregone tax benefits amounts to $465 each year.

To interpret these empirical patterns, I develop a model that incorporates the continuous choice of HSA saving with the discrete choice of deductibles. The model includes an MPC parameter measuring the proportion of the marginal HSA dollar saved to finance current out-of-pocket costs. The MPC parameter summarizes the potential effect of various economic fundamentals on insurance plan choices and health care spending, including discounting, mental accounting, and information frictions. The data is consistent with mental accounting

(Thaler 1985, 1990; Prelec and Loewenstein 1998; Shefrin and Thaler 1988). I find less support for information frictions, discounting, or liquidity constraints.

Using the HSA as supplemental insurance blunts the incentives of the deductible, and has implications for both health spending and plan choices. In this setting, moving all employees from traditional coverage to HDHPs produced no reduction in health care costs. Controlling for 401(k) saving, increases in HSA assets strongly correlate with increases in health spending. In addition, almost 30 percent of employees switch deductibles annually. This high rate of plan switching is not consistent with risk aversion and only partly explained by changes in health status. Instead, changes in HSA assets help to account for such behavior.

This paper is the first to empirically study the role of HSA saving on plan choices and spending, and contributes to several strands of literature. It adds to a growing number of studies on health insurance choices, particularly recent work documenting challenges in decision-making with non-linear contracts (Liebman and Zeckhauser 2008; Abaluck et al. 2015; Dalton et al. 2015), complex choice environments (Kling et al. 2012; Bhargava et al. 2017), and price framing on menus (Schmitz and Ziebarth 2015).[1] By incorporating saving decisions into insurance choices, the paper extends the standard insurance choice model beyond claim probabilities and risk aversion, similar to recent work in health (Handel 2013; Einav et al. 2013; Handel and Kolstad 2015) and other contexts (Sydnor 2010; Barseghyan et al. 2013). The spending results stand in contrast with existing research on HDHPs that document moderate cost reductions (Buntin et al. 2011; Borah et al. 2011; Haviland et al. 2015; Brot-Goldberg et al. 2017). The absence of a spending decline in this setting is consistent with evidence from Medigap that supplemental insurance dampens consumer incentives for cost-control (Cabral and Mahoney 2014).

The paper proceeds as follows. Section 2 discusses the empirical setting and data. Section 3 estimates the MPC based on a regression-discontinuity design. Section 4 provides a test of

[1] See e.g. Cardon and Hendel 2001; Carlin and Town 2009; Einav et al. 2010; Abaluck and Gruber 2011; Bundorf et al. 2012; Ketcham et al. 2012; Heiss et al. 2013; Polyakova 2016 for other examples in the plan choice literature.

fungibility between the HSA and 401(k). Implications for health spending and plan choices are discussed in Section 5. Section 6 introduces a model to interpret the findings and Section 7 discusses mechanisms.

2 Setting and Data

The company is one of the five largest health insurers in the U.S. by both market share and revenues, with employees throughout the country. In terms of representativeness, the average salary, age, and tenure among the company's employees are roughly in line with U.S. labor force averages. About two-thirds of employees are female, which is substantially higher than many other settings. Between 2005 and 2010, the company made changes to its health insurance and retirement saving programs for its own employees. It fully replaced its standard insurance contracts with HDHPs in 2008, introducing a menu of HDHP offerings that were differentiated only in the size of the deductible and were otherwise equivalent (e.g. provider networks were identical). Below I describe the specific policies enacted by the employer. Appendix A provides additional background on HSAs more generally.

In 2005, the company began offering employees the choice of HDHPs in addition to traditional health insurance plans. The traditional insurance plans offered first-dollar coverage for physician office visits and prescription drugs. Patients paid $25 or $40 copayments for office visits and received an allowance for each drug fill that varied by tier of the drug. Preventive care—immunizations, physical exams, and certain cancer screening—was free of charge. Patients paid 100 percent out-of-pocket for hospital care and other services until they reached their deductibles. Beyond the deductible, the plan paid for all costs, excluding copayments for office visits which the patient still paid each visit. In the HDHP plan, preventive care was also free and the patient paid the full charge for all other care (physician visits, drugs, hospital care, etc.) until the deductible had been met. After the deductible was reached, the plan paid for all charges.[2] Both types of plans

[2]Starting in 2010, plans also had a 10 percent coinsurance rate until the out-of-pocket maximum, which was $1,500 higher than the deductible for each plan for employee-only coverage and $3,000 higher for family

had identical provider networks.

In 2008, the company dropped its traditional insurance plans and only offered HDHP plans to its employees. Employees had to make an active choice of insurance plan each year during open enrollment. Figure 1 displays the menu of HDHP deductibles between 2007 and 2010 for employee-only (self) coverage. The deductibles for family coverage were double these amounts. The employer contributed a flat amount to each plan premium and then required employees to pay higher costs of additional coverage. Only employees choosing an HDHP plan could open an HSA.

The employer matched employee HSA contributions at different rates based on salary level. For employees earning below $50,000, the employer matched the first $100 of contributions at a rate of 6:1 for those with self coverage, for a maximum employer contribution of $600. The contribution limits were double for family coverage, so that employees could receive $1,200 from employer contributions on the first $200 saved. Between $50,000 and $100,000, the match rate was 4:1, up to maximum amounts of $400 for self coverage and $800 for family coverage. In 2010, these rates were reduced from 6:1 to 5:1 for salaries below $50,000, and from 4:1 to 2:1 for salaries between $50,000 and $100,000.[3] There was no default employee contribution to the HSA. Employees were immediately vested for both their contributions and the employer's contributions.

The HSA operated as both a low-interest rate savings account and an investment account. In this setting, once balances reached $2,000, HSA assets could be invested in a variety of mutual funds provided the amount invested exceeded $1,000 and remained above this level. There were no initial setup fees or monthly fees for the HSA account or the investment account paid by the employee. Low fees and the ability to invest in mutual funds are important because they rule out the argument that consumers should rationally not use their HSA for retirement saving if investment opportunities are poor and transaction costs

coverage.

[3] There was also a 2:1 match rate for employees between $100,000 and $150,000 until 2008. Since only about 5 percent of the sample falls above $100,000, the main analysis focuses on employees below this level.

are high.[4] In addition, the contribution limits established by the IRS are low enough that most people could expect to use all HSA assets to finance health care in retirement: a person could save the maximum each year in the HSA, make no withdrawals until age 65, and accumulate less than the net present value of health care costs not covered by Medicare (Fronstin et al. 2008; Webb and Zhivan 2010; Yamamoto 2013; Fidelity 2014).

The company pursued an extensive communications campaign to inform employees about the HDHP offerings and HSA program. This effort included materials and programs to aid employees in analyzing insurance options and monitoring expenditures. Employees received an annual "Smart Summary" with details on their spending patterns and indicating alternative plans or therapies that may save costs. The employer also provided online budgeting tools, cost calculators, and other resources on their insurance and saving products. In marketing the HSA to its employees, the employer explained the account's tax preferences, and did not explicitly describe it as either a retirement savings vehicle or as a way to "offset" the HDHP. Similar to other settings, the employer provided an HSA debit card to employees.

In terms of retirement benefits, the company offered employees a defined-contribution 401(k) and matched employee contributions up to 6 percent of salary. Prior to 2008, the company matched all employee contributions at 50 percent up to this threshold. Starting in 2008, the company began matching the first percent of employee salary 100 percent and then matched subsequent contributions at 50 percent, up to 6 percent of salary (so that the maximum employer contribution increased from 3 to 3.5 percent of salary). Employee contributions were deducted from each period's paycheck. Unlike with the HSA, the employer provided a higher match limit in absolute dollars to higher salaried workers because the 401(k) match was based on a percentage of salary. If employees did not actively enroll in the 401(k) when they were hired, they were auto-enrolled at a salary contribution of 4 percent, also deducted from each period's paycheck.

[4]Even if HSA assets could not be invested, one might argue that people could rebalance their 401(k) and treat the HSA as the low-yield portion of the portfolio since the money is fungible. This strategy would, nevertheless, entail some hassle costs.

The administrative data includes detailed information on each employee's salary, job characteristics, demographics, medical and pharmacy claims, and choices about retirement saving and health insurance plans. Information on job characteristics and geographic location is measured once in 2011, while information on salary, retirement saving and health insurance choices, and medical and pharmacy claims are measured repeatedly over time. 401(k) contributions (both by the employee and employer) and balances are measured annually. HSA variables—contributions by both the employee and employer, employee withdrawals, balances, and interest—are measured monthly.

The claims data includes information on health expenditure for employees and any dependents covered under the employee's policy. Each claim provides detailed information on diagnoses (ICD-9 codes for medical claims and NDC codes for pharmacy claims), providers, and payment (e.g. patient paid, plan paid), and dates of payment. Each claim also includes an estimate of the employee's health expenditure risk, called the "severity score", that is developed by the employer.

Using the severity score and the medical and pharmacy claims, I construct distributions of expenditure risk for each employee and his or her dependents. Appendix B provides the details of this procedure, which follows the approach of other studies of insurance plan choices (Handel 2013; Handel and Kolstad 2015). As an overview, the strategy is to group people with similar health risks in year t together and use the distribution of actual expenditures from year $t+1$ to generate beliefs about expenditure risk at the time the insurance plan was chosen in year t. Each person in the same risk group is assumed to have the same beliefs about his or her health risk. More specifically, I first group each insured individual into 60 different risk groups based on age, sex, and severity score. For each of the 60 risk groups, I record the empirical proportion of individuals with zero expenditures the following year. For those with positive expenditures the following year, I fit a Weibull distribution to the observed expenditure, estimating the shape and scale parameters of this distribution. Expenditures on preventive care are excluded since such costs are covered free of charge by all plans.

Only claims from in-network providers are considered, which comprise nearly all spending. For each risk group, I then construct a "modified" Weibull distribution using the group's estimated shape and scale parameters for positive expenditure and the empirical probability of zero expenditures. For each risk group, I take 100 draws from their corresponding modified Weibull distribution. Within each family, the expenditures for each draw are summed so that each family has 100 draws corresponding to the sum of expenditures for each of its members from that particular draw. This statistical object represents the family's beliefs about its total expenditure risk.

Sample composition

The main analysis sample is constructed by starting with all employees appearing in the employer's Human Resource records in plan years 2007 to 2010 (roughly 26,000) and restricting to those who were (i) enrolled in one of the company's health insurance plans, (ii) did not switch the number of covered dependents during the year on their insurance plan, (iii) had coverage the entire year when insured, (iv) were younger than age 59, and (v) actively enrolled in the 401(k). Restrictions (i) - (iii) are simply to isolate those whose insurance status is not fragmented, and reduces the sample size by 5,522 employees (21 percent of the full sample). I exclude the small number of employees aged 59 years and older (1,210 employees) because 401(k) assets can be withdrawn penalty-free for any reason starting at age $59\frac{1}{2}$. I finally exclude the 18 percent of remaining employees who passively default into saving four percent of salary in their 401(k) because the analysis of HSA saving assumes that the employee's 401(k) saving decisions captures their inter-temporal preferences and retirement saving objectives. This assumption is difficult to justify for people who auto-enroll in the 401(k). After these restrictions, the sample size totals 15,908 employees. Appendix C presents analyses of the employees who passively default into the 401(k) to explore robustness to this restriction.

Table 1 presents summary statistics of the analysis sample, overall and by type of insurance coverage. The average employee age is 40 years, the average tenure with the firm is 7.3 years, and the average salary is \$61,939. Sixty-seven percent of the sample is female. Annual HSA saving, including employer contributions, averages \$1,416 for self-only coverage and \$2,784 for family coverage. Over 97 percent of the sample receives the full employer HSA match and 5.5 percent contribute the maximum to their HSAs. By comparison, just under half receive the full employer 401(k) match and 3 percent contribute the maximum to their 401(k)s. The lower contribution thresholds for the HSA explain the higher share of employees obtaining the full HSA match. For each account, the large majority of employees are not at a corner solution of contributing zero or the maximum amount. Nearly 86 percent of annual HSA contributions are withdrawn that year, on average. Almost 60 percent of the sum of existing balances and contributions are withdrawn. Appendix Table C.1 presents summary statistics stratified by whether the employee adopted the HSA prior to the forced switch in 2008 to provide a sense of the nature of selection into the HSA. On average, early adopters have higher incomes, contribute and withdraw more from their HSA, have higher expected health spending, and have higher 401(k) saving than employees who wait until 2008 to join. These patterns are inconsistent with adverse selection on health risk, in which low-cost employees are more likely to enroll in the HDHP and HSA, but consistent with selection based on expected tax benefits.

3 Regression-Discontinuity Analysis

The employer's matching schedule for HSA contributions provides an exogenous source of variation in HSA assets. To estimate the marginal propensity to consume (MPC) from the HSA, I quantify the change in HSA withdrawals at the match discontinuity and compare it to the change in HSA contributions. To test whether HSAs crowd out 401(k) saving, I estimate the change in 401(k) contributions at the HSA match discontinuity. Figures 2-4 preview the regression-discontinuity results by plotting HSA contributions, HSA withdrawals, and 401(k) contributions by salary level. Figure 2 plots the average level of employer contributions to the HSA within \$500 salary bins for employees earning less than \$100,000 annually. There is a visible drop in employer contributions at the match discontinuity.[5] Total HSA contributions increase with salary, but there is a jump downwards in the regression function at $50,000, indicating that employees do not fully reduce their own HSA contributions in response to the match. Figure 3 presents a similar figure for HSA withdrawals, and also shows a decrease at the match discontinuity. Figure 4 reveals there is no visible change in the level of 401(k) contributions across the HSA match discontinuity, which indicates that HSA saving does not crowd-out 401(k) saving. To establish validity of the RD design, I verify that the density of salary, the running variable, is smooth across the $50,000 cutoff (Appendix Figure C.1) and that controls are balanced across the match discontinuity (Appendix Table C.2).

The RD regression models take the following specification:

$$y_{it} = \alpha + \beta_1 M_{it} + \beta_2 salary_{it} + \beta_3 M_{it} \times salary_{it} + \varepsilon_{it} \tag{1}$$

where y_{it} denotes either the level of HSA contributions, HSA withdrawals, or 401(k) contributions for employee i in year t, *salary* is the employee's annual salary re-centered at $50,000, and M is an indicator variable for employees with salaries between $50,000 and $99,999. The coefficient of interest is β_1, which measures the change in the outcome that occurs at the match discontinuity. The bandwidth for the regression is selected to minimize the mean-squared error (MSE) following the methods of Calonico, Cattaneo and Titiunik (2015). I also estimate quadratic polynomials for robustness. My preferred specification is the local linear model presented in equation (1) because higher-order polynomials can have undesirable properties in terms of inference, sensitivity to polynomial degree, and weighting (Gelman and Imbens 2018). Standard errors are clustered by employee.

Table 2 reports the RD estimates from equation (1) that correspond to Figures 2, 3, and 4. Panel A reports an estimated $323 drop in employer contributions at the cutoff using the local linear model (Column 1). The estimated drop in HSA withdrawals at the match

[5] This figure combines both self and family coverage.

cutoff is \$220 (Panel B, Column 1). Taking the ratio of these estimates yields an MPC of 0.68 (Panel C), with the standard error calculated by a block bootstrap. Using quadratic polynomials yields a higher MPC estimate of 0.81. Finally, Panel D reveals an imprecisely estimated change in 401(k) saving at the match discontinuity, consistent with the absence of crowd-out. Similar results are obtained when including controls (Appendix Table C.3). As an additional check on the validity of the research design, I estimate the RD models using placebo cutoffs of \$25,000 to \$75,000 in \$5,000 increments, where no change in outcomes is expected (Appendix Table C.4). The point estimates are all close to zero and one in twenty is statistically significant at the 5 percent level, as predicted by chance.

The estimated magnitude of the MPC from the HSA is large. An MPC of 1 would imply consumers treat their HSA fully as supplemental insurance, like an FSA. The MPC need not be 0 to treat the HSA as self-insurance, but an estimated MPC of over two-thirds is on the high side of results from studies that measure the change in annual spending in response to tax rebates (Sahm et al. 2010; Shapiro and Slemrod 2009; Parker et al. 2013). Carroll et al. (2017) conclude that most of the literature documents a value of the MPC between 0.2 to 0.6, measured over the course of a year.[6] When interpreted in this context, consumers appear to use the HSA largely as supplemental insurance, rather than as self-insurance.

Appendix Table C.5 presents results for different sub-samples to examine heterogeneity across employees and robustness to sample restrictions. In particular, the MPC is estimated to be higher for self coverage than family coverage and for early adopters of HSAs versus late adopters. Perhaps surprisingly, the MPC is higher for those without a 401(k) loan than for those with a 401(k) loan. If those taking 401(k) loans are more likely to be liquidity constrained, as argued by Lu and Mitchell (2010), then one would expect a higher MPC for such employees. This pattern suggests liquidity constraints are not the mechanism driving the high MPC. Finally, the estimated MPC is 0.71 for the group of employees who auto-enrolled into the 401(k), demonstrating robustness to this sample restriction.

[6] Jappelli and Pistaferri (2010) review earlier studies on the relationship between consumption and income, focusing on the difference between anticipated and unanticipated income.

4 Tests of Fungibility

In general, determining whether saving is sub-optimal is difficult due to unobserved information and often requires making untestable assumptions. The similarity between HSAs and 401(k)s provides a setting to examine the frequency of mistakes, however. By revealed preference, any employee whose 401(k) contribution exceeds the employer match and who is not maxing out the HSA would be unambiguously better off by reallocating some 401(k) savings to the HSA. Without the employer's 401(k) match, the tax incentives on the last dollar contributed are more generous in the HSA than the 401(k) for two reasons. First, withdrawals for medical care are always tax-exempt. Second, HSA contributions are not subject to FICA taxes, but 401(k) contributions are. Most employees in the sample fall below the payroll tax cap, making this difference in FICA tax treatment relevant. The other tax preferences are identical between accounts. As a result, the 401(k) is dominated once employer matching is exhausted. I define a dominated choice as failing to contribute the HSA maximum among the set of employees who save beyond the 401(k) match. Over half of the analytic sample contributes beyond the employer 401(k) match in at least one year, and these employees have a higher average salary than those who are at or below the match.

To describe the characteristics associated with making decision errors, I run linear probability models on the sub-sample of employees saving beyond the match:

$$y_{it} = \alpha + X\delta + \tau_t + c_s + \varepsilon_{it} \tag{2}$$

where y_{it} is an indicator for employee i making a dominated choice in year t, X denotes a vector of employee characteristics that includes salary, tenure, years with the HSA, self coverage, age, sex, race, and job characteristics, τ_t are time effects, and c_s are state effects. To gauge the magnitude of decision errors, I also estimate models in which the dependent variable is equal to the amount of 401(k) contributions that are dominated, conditional on making a dominated choice. This amount is calculated as the minimum of (1) the quantity of 401(k) contributions exceeding the match, and (2) the difference between the HSA contribution limit and the employee's total HSA contribution. On average, this amount is $1,095, ranging from $348 for those in the lowest income quintile and $1,715 for those in the highest income quintile. Multiplying this amount by the employee's marginal tax rate yields an estimate of the foregone tax benefits each year.[7] The size of these mistakes are non-negligible. On average, employees forego an estimated $465 in tax benefits, on average, each year.

As shown in Table 4, income is negatively associated with making a dominated choice (Column 1). Employees in the top income quintile are 18.1 percentage points less likely to make a dominated choice than employees in the lowest income quintile, although over 70 percent of the highest income quintile still do so. The number of years of experience with an HSA also reduces the probability of making a dominated choice. Actuaries and employees who work in finance or accounting are slightly less likely to make mistakes than other employees, but the estimate is not statistically significant. Column 2 presents OLS models of the amount of dominated 401(k) contributions, restricted to those who make a dominated choice. Given their larger saving levels, higher-income employees have larger amounts of dominated 401(k) contributions than lower-income employees. Age is also positively associated with the magnitude of the decision error. Column 3 shows similar patterns in terms of the foregone tax benefits, with employees in the top two salary quintiles leaving over $300 more on the table than those in the lowest quintile. Age and tenure with the employer are again both positively associated with larger mistakes.

[7]The marginal tax rate is estimated using NBER TAXSIM and adds the federal, state, and FICA tax rate (including both employer and employee shares). If an employee were to use the money for non-health consumption in retirement, then the tax benefits would be based on the FICA tax rate and so would be smaller.

5 Implications for Plan Choices and Health Spending

5.1 Health spending

In this setting, spending did not decline after the employer fully replaced first-dollar coverage with HDHPs and HSAs. Figure 5 plots the time series of monthly per person spending between 2006 and 2009 for the employees who did not enroll until 2008. The vertical line represents the final month of the 2007 plan year (corresponding to June 2008), before all employees moved to HDHPs. Two adjustments are made to make spending comparable over time. First, spending is inflated based on the medical care component of the consumer price index. Second, spending is adjusted to account for aging within the sample. This second correction is made by regressing monthly spending on age and other covariates among employees and dependents, and then adjusting each observation's spending for the predicted change in spending from aging one year. A one-year increase in age is associated with an $11 increase in monthly spending. Without adjusting for aging, spending would have a tendency to rise over time simply because employees and dependents grow older. Spending is expressed based on prices and ages in January 2009. In addition to the adjusted mean, Figure 5 presents the 95 percent winsorized mean and the median to remove the effect of outliers. In both cases, there is no visible evidence of a sustained spending decrease after 2008. Spending increases in the final month of each plan year, before falling in the subsequent month, and this is pattern is also observed for the winsorized mean and median. The fact that spending did not decline after the introduction of HDHPs and HSAs contrasts with studies from other settings that document moderate spending reductions in the range of 10 to 15 percent (Buntin et al. 2011; Borah et al. 2011; Haviland et al. 2015; Brot-Goldberg et al. 2017).[8]

To directly analyze how HSAs influence spending, I regress the log of annual health

[8] Whether such reductions are welfare-improving depends on the value of services that are cut. Brot-Goldberg et al. (2017), for example, document consumers reduce consumption indiscriminately across all services, including preventive care that is not subject to the deductible. To the extent consumers misperceive the benefits of care—known as behavioral hazard (Baicker et al. 2015)—HSAs may increase welfare relative to an HDHP alone if they prevent the reduction of valuable services.

spending against quintiles of HSA balances, controlling for quintiles of 401(k) assets. A positive correlation between HSA assets and health spending would be consistent with using the HSA as supplemental insurance. The purpose of conditioning on 401(k) saving is to account for differences across employees in preferences for retirement saving. Table 4 presents the results of these regressions, which also account for demographics and other controls. HSA assets positively correlate with health spending. Total health spending is 14.8 percent higher annually among employees in the second-lowest quintile of HSA assets compared to those in the lowest (Column 1). These percentages increase with each quintile. The highest quintile, for example, spends 37.3 percent more than the lowest quintile. Columns 2 and 3 add employee fixed effects, with Column 2 excluding other controls except for 401(k) saving and year fixed effects and Column 3 including time-variant controls. The estimated magnitudes are larger using within-employee variation: compared to the bottom quintile, the second-lowest quintile spends 16.5 percent more, the middle quintile spending 33.4 percent more, the fourth quintile spends 46.0 percent more, and the highest quintile spends 66.7 percent more (Column 3).

5.2 Plan switching

Between 2008 and 2010, nearly 30 percent of consumers switch between the four deductible levels each year. Thirty-eight percent switch in one of the two years, with roughly equal numbers increasing and lowering their deductibles. This high switching rate is remarkable, considering that no plan was dominated, in contrast to other research documenting lower switching rates in settings with dominated plans (Handel 2013, Bhargava, Loewenstein, Sydnor 2017). Such rates are comparable to enrollees in the Affordable Care Act's exchanges (Sanger-Katz 2015). The two common sources of consumer heterogeneity used to model plan choice are risk preferences and expenditure risk. Using HSAs as supplemental insurance may be a third factor that predicts deductible choices over time. I take advantage of the panel data to examine the correlation between HSA assets and plan choices by estimating fixed effects logits. These regressions specify the probability of choosing the highest deductible as:

$$\Pr(D_{kt}^{HIGH} = 1|\varsigma_k, \beta) = \frac{\exp(\varsigma_k + x_{kt}\beta)}{1 + \exp(\varsigma_k + x_{kt}\beta)} \tag{3}$$

where D_{kt}^{HIGH} is a binary variable equal to 1 if the highest deductible is chosen and 0 otherwise, ς_k is a an employee fixed effect, and and x_{kt} represents the vector of regressors that include HSA and 401(k) assets, expected health risk, and other controls. I also run the same regression with the lowest deductible chosen as the dependent variable. The fixed effect is intended to capture time-invariant preference parameters, such as risk aversion, or other characteristics that may differ between employees and predict plan choices.

The first two columns of Table 5 present odds ratios of choosing the lowest and highest deductibles, respectively. Within-employee increases in HSA assets raise the likelihood of choosing the highest deductible and lower the probability of choosing the lowest deductible, holding expected spending and other factors constant. 401(k) saving is not a statistically significant predictor of plan choice. Since plans differ only in the size of their deductibles, Column 3 reports results from random effects ordered logit models in which the dependent variable is one of the four deductible tiers. Higher HSA balances positively correlate with higher deductible choices across regressions. In terms of heterogeneity, employees who withdraw larger shares of their HSA assets switch plans at higher rates than those who withdraw less. The positive relationship between the share withdrawn and switching holds after controlling for demographics, 401(k) saving, and the level of HSA assets. These patterns are consistent with using the HSA as supplemental insurance. Switching rates are similar between employees with self and family coverage, and between early and late HSA adopters.

6 A Model of Insurance Choice and HSA Saving

To interpret these patterns, this section presents a model that integrates HSA saving into a standard model of deductible choice by adding the MPC as a preference parameter. In the neoclassical model of insurance choice, employees choose the plan that maximizes expected utility given their risk aversion, claim probability, and marginal tax rate. I follow the plan choice literature in assuming preferences satisfy constant absolute risk aversion (CARA), so that for consumption x, $u(x) = -\frac{1}{\gamma}e^{-\gamma x}$, where γ is the coefficient of absolute risk aversion.[9] For family k in year t, a consumption draw under insurance plan j is specified as:

$$x_{kj} = (y_k - \pi_j)(1 - \tau_k) - OOP_{kj} \tag{4}$$

where y_k denotes income, π_j denotes the plan premium, τ_k is the family's marginal tax rate, and OOP_{kj} is an out-of-pocket realization under plan j based on family k's ex ante cost distribution F_k. The equation reflects the tax preference for insurance that allow premiums to be paid with pre-tax dollars, while out-of-pocket payments are paid with after-tax dollars. Employees have a discrete choice of J insurance plans and choose the plan that maximizes their expected utility, defined as:

$$\begin{aligned} U_{kj}\left(\gamma_k, F_k\left(OOP_{kj}\right)\right) &= \int_0^\infty u\left(x_{kj}\right) dF_k\left(OOP_{kj}\right) \\ &= \int_0^\infty -\frac{1}{\gamma}\exp\left(-\gamma x_{kj}\right) dF_k\left(OOP_{kj}\right) \end{aligned} \tag{5}$$

For a given cost distribution F and marginal tax rate τ, the choice of plan identifies a set of risk aversion parameters for each employee consistent with expected utility maximization because there is a finite menu of insurance contracts to choose between. If employees could instead choose from a continuous menu of deductibles and corresponding premiums, then the choice of coverage would point-identify risk aversion for each employee conditional on his cost distribution and marginal tax rate.

Insurance choice with HSAs

First, consider HSA contributions used as supplemental insurance to finance the

[9] See e.g. Einav et al. (2013); Handel (2013); Handel and Kolstad (2015). As noted by Rabin (2000) and Rabin and Thaler (2001), expected utility struggles to explain risk preferences when gambles are modest. In my setting, gains and losses are on the order of $1,000 and I follow existing studies of health insurance choices that assume expected utility and risk aversion in explaining plan choices.

deductible. HSAs provide a continuous choice of insurance contract, similar to an FSA. People can purchase a dollar of additional coverage at a premium equal to 1 minus their marginal tax rate. Figure 6 displays the schedule of out-of-pocket payments versus medical expenses for a \$3,000 deductible and different levels of HSA contributions. With no HSA contribution, the insurance contract is represented by the solid black line where the employee pays for all costs until \$3,000, and the plan pays for any additional costs. If used as supplemental insurance, HSA contributions shift the benefit schedule down towards the horizontal axis. Any level set within the shaded area is now a feasible contract. For example, the dashed line denotes an HSA contribution of \$1,000, so that the first \$1,000 of costs are paid with pre-tax HSA funds, the next \$2,000 of costs are paid from after-tax funds, and costs beyond \$3,000 are paid by the insurance plan.

When used like an FSA, HSA contributions h then represent a second insurance premium by making the consumption draw:

$$x_{kj}(h|\pi,\tau) = (y_k - \pi_j - h_{kj})(1-\tau_k) - OOP_{kj} + \underbrace{\min\{OOP_{kj}, h_{kj}\}}_{\text{HSA withdrawal}} \tag{6}$$

Each OOP realization may now be financed with pre-tax funds withdrawn from the HSA, represented by the last term. The employer matches the first L dollars of employee contributions according to a rate m, where $m = 1$ denotes a dollar-for-dollar match. A higher match rate provides for additional coverage, so that a consumption draw is then represented as:

$$\begin{aligned} x_{kj}(h|\pi,\tau,m,L) = {} & (y_k - \pi_j - h_{kj})(1-\tau_k) - OOP_k \\ & + \min\{OOP_k, \min[h_{jk}(1+m_k), h_{kj} + m_k L_k]\} \end{aligned} \tag{7}$$

I now introduce an MPC parameter η that captures the share of HSA assets that will be withdrawn to finance out-of-pocket costs. Denoting HSA balances at the beginning of the year for family k as H_k, a consumption draw is then represented as

$$x_{kj} = (y_k - \pi_j - h_{kj})(1 - \tau_k) - OOP_{kj} + \min\{OOP_{kj}, \eta(H_k + \min[h_{kj}(1 + m_k), h_{kj} + m_k L_k])\} \tag{8}$$

This formulation nests the standard model represented in equation (1) as a special case when $\eta = 0$. In that case, the deductible choices would be identical whether or not the employee had an HSA. If $\eta = 1$, then HSA contributions function entirely as supplemental insurance, just like an FSA.

This model takes a static perspective to explaining insurance and saving choices. There is option value from HSA contributions that roll over if unused, unlike FSA contributions. Incorporating dynamics is potentially important, although in this context, the option value is likely small because consumers can change contributions monthly, and can make additional HSA contributions over and above their payroll deductions. So a consumer planning to use his HSA to finance this year's deductible could make a small payroll contribution and then contribute more upon incurring medical expenses. There is also little option value from building up HSA assets to hedge against future deductible increases in this setting because the deductible is never larger than the statutory maximum HSA contribution. Consequently, an employee is never at risk for paying more out-of-pocket than he could contribute to his HSA that same year. A richer theoretical model would necessarily include such dynamics, however.

An illustration of the HSA model on plan choices

If used as supplemental insurance, HSAs can induce different deductible choices compared to the standard model. The intuition is that existing HSA balances are viewed as "house money." People planning to finance out-of-pocket costs from the HSA may thus choose a higher deductible than without an HSA, since the deductible is reduced, in effect.[10] Figure

[10] Cardon (2012) theoretically proves that the preferred deductible is weakly higher with FSA contributions

7 displays the deductible chosen in 2009 as predicted from the HSA plan choice model in equations (5) and (8) as a function of the MPC (plotted on the vertical axis) and the size of beginning-year HSA balances (plotted on the horizontal axis). This particular calibration is presented for a female between ages 35 and 44 in the median quintile of health risk, and assuming a coefficient of absolute risk aversion γ of 2.5×10^{-4}, a marginal tax rate of 20 percent, and an employer match of 6:1 on the first \$100 contribution. The standard model—corresponding to $\eta = 0$—predicts a deductible choice of \$1,250. For $\eta > 0$, different combinations of balances and η can induce higher deductible choices, with the preferred deductible increasing in both the MPC and the size of the HSA balance. This model is consistent with both the large MPC and the positive correlation between plan switching and HSA withdrawals observed in the data.

7 Conclusion

This paper studies the saving decisions and plan choices of employees at a large health insurer that fully replaced its traditional health insurance offerings with HDHPs and HSAs. In theory, these contracts offer an innovative design to balance incentives against risk protection. Linking a personal, tax-preferred savings account to insurance coverage encourages consumers to trade off current health care against future income. For the high deductible to increase sensitivity to costs as intended, consumers should view HSA contributions as future cash, rather than as supplemental insurance, like an FSA or Medigap. Using a regression-discontinuity design, I estimate an MPC from HSA funds of 0.68. Employees in this setting use their HSA largely to reduce the deductible, rather than to self-insure. Replacing traditional coverage with these contracts produced no decline in health care spending in this setting. I assess the frequency of consumer mistakes in HSA saving decisions by examining saving for the sub-sample of employees whose 401(k) contributions

under constant absolute risk aversion.

exceeds the employer match. For these employees, the tax preferences of the marginal HSA dollar dominate those of the marginal 401(k) dollar, and so employees should be contributing the maximum to their HSA. Ninety percent fail to do so, however, which violates the fungibility of money.

The MPC summarizes the potential effect of several economic fundamentals on plan choice and consumption, including discounting, mental accounting, information frictions, and financial literacy. Understanding the micro-foundations behind HSA saving decisions may be important for policy; interventions to encourage desirable choices—whether through information, financial incentives, defaults, etc.—are likely to be more effective if they are targeted at the relevant part of the decision process.

While I am unable to study the role of different mechanisms directly, mental accounting is the most plausible mechanism that is consistent with the data. Mental accounting, which is an example of framing, assumes households group income and expenditure items into separate accounts (e.g. current income, future income) and that the marginal propensity to consume differs between accounts (Thaler 1985, 1990; Prelec and Loewenstein 1998; Shefrin and Thaler 1988). Households may also earmark funds for different purposes. The "envelope method" of budgeting—whereby households allocate cash to different physical envelopes for monthly spending on food, gas, etc.—is an example of such behavior. Holding separate accounts violates the fungibility of money. Prior research has documented violations of fungibility for particular expenditure items like gasoline (Hastings and Shapiro 2013), grocery purchases (Milkman and Beshears 2009), restaurant meals (Abeler and Marklein 2013), children's clothing (Kooreman 2000), and Supplemental Nutritional Assistance Program (SNAP) benefits (Hastings and Shapiro 2018). Similarly, other research provides evidence that financial borrowing decisions by some consumers violate the no-arbitrage condition: many take payday loans when lower interest credit is available (Agarwal, Skiba and Tobacman 2009) or simultaneously hold both high-interest credit card debt and low-yield assets (Gross and Souleles 2002).

In the context of HSAs, people may view their HSAs as accounts designated to cover health care expenses, while their 401(k)s are designated for retirement saving, even though the money is fungible. The high rate of dominated saving choices points to this explanation. Such empirical anomalies may ultimately stem from a gap in financial literacy, which includes not only information about financial products but also the mathematical skills and conceptual knowledge required to make sophisticated financial decisions (Hastings, Madrian and Skimmyhorn 2013; Lusardi and Mitchell 2014).

There is less empirical support for other mechanisms, such as information frictions about health insurance, liquidity constraints, or discounting. In particular, employees appear to understand that the money rolls over because HSA withdrawals do not spike at the end of the year. Employees also appear knowledgeable about the coverage provided by HDHPs: there is no decline in the number of preventive visits, total spending on preventive care, or the number of screening mammographies after enrolling in an HDHP. This pattern contrasts with research from other settings documenting that consumers are generally unaware that preventive care was covered in HDHPs and cut back on these services (Reed et al. 2012; Brot-Goldberg et al. 2017). In terms of liquidity constraints, the higher estimated MPC among those without 401(k) loans than those with 401(k) loans is inconsistent with this mechanism.

Discounting is a final candidate to explain the empirical patterns. If people have strong preferences for current consumption, they may choose to contribute and withdraw HSA funds to finance current health care costs even though they recognize their future consumption demands. The debit card arguably makes HSA withdrawals easier than 401(k) withdrawals.[11] Models of present-biased preferences (Laibson 1997; O'Donoghue and Rabin 1999) help to explain the demand for commitment devices like 401(k)s for retirement saving (Laibson et al. 1998). One might argue that the high rate of dominated choices is due to the five

[11] Yet even the 401(k)—clearly intended by policy as a retirement savings vehicle—still gives account holders ample flexibility for penalty-free withdrawals, known as "leakage" (Beshears et al. 2015; Munnell and Webb 2015).

additional years people must wait until HSA assets can finance non-health consumption without penalty, compared to the 401(k). Yet the result that older employees make larger mistakes than younger employees runs counter to discounting; older employees have fewer years until they can use HSA assets as cash, and so should have lower amounts of 401(k) saving past the deductible than younger employees if discounting explained these patterns. While the evidence points to mental accounting, this exploration of potential mechanisms should nonetheless be interpreted as suggestive. Directly testing mechanisms that drive HSA use is an important topic for future research.

References

Abaluck, Jason and Jonathan Gruber, "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program," *American Economic Review*, 2011, *101* (4), 1180–1210.

_, _, and Ashley Swanson, "Prescription Drug Use Under Medicare Part D: A Linear Model of Nonlinear Budget Sets," 2015.

Abeler, Johannes and Felix Marklein, "Fungibility, Labels, and Consumption," 2013.

Agarwal, Sumit, Paige Marta Skiba, and Jeremy Tobacman, "Payday Loans and Credit Cards: New Liquidity and Credit Scoring Puzzles," *American Economic Review Papers & Proceedings*, 2009, *99* (2), 412–17.

Arrow, Kenneth, "Uncertainty and the Welfare Economics of Medical Care," *American Economic Review*, 1963, *53* (5), 941–973.

Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein, "Behavioral Hazard in Health Insurance," *Quarterly Journal of Economics*, 2015, *forthcoming.*

Barseghyan, Levon, Francesca Molinari, Ted O'Donoghue, and Joshua Teitelbaum, "The Nature of Risk Preferences: Evidence from Insurance Choices," *American Economic Review*, 2013, *103* (6), 2499–2529.

Beshears, John, James Choi, Joshua Hurwitz, David Laibson, and Brigitte Madrian, "Liquidity in Retirement Saving Systems: An International Comparison," *NBER Working Paper 21168*, 2015.

Bhargava, Saurabh, George Loewenstein, and Justin Sydnor, "Choose to Lose: Health Plan Choices from a Menu with Dominated Option," *Quarterly Journal of Economics*, 2017, *3*, 1319–1372.

Borah, Bijan, Marguerite Burns, and Nilay Shah, "Assessing the Impact of High Deductible Plans on Health-Care Utilization and Cost: A Changes-in-Changes Approach"," *Health Economics*, 2011, *20* (9), 1025–1042.

Brot-Goldberg, Zarek, Amitabh Chandra, Benjamin Handel, and Jonathan Kolstad, "What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics," *Quarterly Journal of Economics*, 2017, *123* (3), 1261–1318.

Bundorf, Kate, Jonathan Levin, and Neale Mahoney, "Pricing and Welfare in Health Plan Choice," *American Economic Review*, 2012, *102* (7), 3214–48.

Buntin, Melinda, Amelia Haviland, Roland McDevitt, and Neeraj Sood, "Health Care Spending and Preventive Care in High-Deductible and Consumer-Directed Health Plans," *American Journal of Managed Care*, 2011, *17* (3), 222–230.

Cabral, Marika and Neale Mahoney, "Externalities and Taxation of Supplemental Insurance: A Study of Medicare and Medigap," 2014.

Cardon, James, "Health Insurance Choice with Flexible Spending Accounts," *The Geneva Risk and Insurance Review*, 2012, *37* (2), 208–222.

_ and Igal Hendel, "Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey," *RAND Journal of Economics*, 2001, *32* (3), 408–427.

_ and Mark Showalter, "An Examination of Flexible Spending Accounts," *Journal of Health Economics*, 2001, *20* (6), 935–954.

_ and _, "Insurance Choice and Tax-Preferred Health Savings Accounts," *Journal of Health Economics*, 2007, *26* (2), 373–399.

Carlin, Caroline and Robert Town, "Adverse Selection, Welfare, and Optimal Pricing of Employer Sponsored Health Plans," 2009.

Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew White, "The

Distribution of Wealth and the Marginal Propensity to Consume," *Quantitative Economics*, 2017, *8*, 977–1020.
Chen, Song, Anthony Lo Sasso, and Aneesh Nandam, "Who Funds Their Health Savings Account and Why?," *International Journal of Health Care Finance and Economics*, 2013, *13* (3-4), 219–232.
Choi, James, David Laibson, and Brigitte Madrian, "$100 Bills on the Sidewalk: Violations of No-Arbitrage in 401(k) Accounts," *Review of Economics and Statistics*, 2011, *93* (3), 748–663.
Cubanksi, Juliette, Christina Swoope, Anthony Damico, and Tricia Neuman, "How Much is Enough? Out-of-Pocket Spending Among Medicare Beneficiaries: A Chartbook," Technical Report, Kaiser Family Foundation 2014.
Dalton, Christine, Gautam Gowrisankaran, and Robert Town, "Myopia and Complex Dynamic Incentives: Evidence from Medicare Part D," *NBER Working Paper 21104*, 2015.
DeNardi, Mariacristina, Eric French, and John Jones, "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 2010, *118* (1), 39–75.
Einav, Liran, Amy Finkelstein, and Mark Cullen, "Estimating Welfare in Insurance Markets Using Variation in Prices," *Quarterly Journal of Economics*, 2010, *125* (3), 877–921.
_ , _ , Stephen Ryan, Paul Schrimpf, and Mark Cullen, "Selection on Moral Hazard in Health Insurance," *American Economic Review*, 2013, *103* (1), 178–219.
Engen, Eric, William Gale, and Cori Uccello, "The Adequacy of Household Saving," *Brookings Papers on Economic Activity*, 1999, *2*, 65–187.
Feldstein, Martin, "Rethinking Social Insurance," *American Economic Review*, 2005, *95* (1), 1–24.
_ and Daniel Altman, "Unemployment Insurance Savings Accounts," in "Tax Policy and the Economy," Vol. 21, MIT Press, 2007, pp. 65–82.
Fidelity, "Retiree Health Costs Hold Steady," Technical Report 2014.
Fronstin, Paul, "Health Savings Accounts and Health Reimbursement Arrangements Assets, Account Balances, and Rollovers 2006-2014," Technical Report 409, Employee Benefits Research Institute, Washington, DC 2015.
_ , "Health Savings Account Balances, Contributions, Distributions, and Other Vital Statistics, 2015: Estimates from the EBRI HSA Database," Technical Report, Employee Benefits Research Institute, Washington, DC 2016.
_ , "Trends in Health Savings Account Balances, Contributions, Distributions, and Investments, 2011â2016: Statistics from the EBRI HSA Database," Technical Report 434, Employee Benefits Research Institute, Washington, DC 2017.
_ , Dallas Salisbury, and Jack VanDerhei, "Savings Needed to Fund Health Insurance and Health Care Expenses in Retirement: Findings from a Simulation Model," *EBRI Issue Brief 317*, 2008.
Gelman, Andrew and Guido Imbens, "Why High-Order Polynomials Should Not be Used in Regression Discontinuity Designs," *Journal of Business and Economic Statistics*, 2018.
Gross, David and Nicholas Souleles, "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence From Credit Card Data?," *Quarterly Journal of Economics*, 2002, *117* (1), 149–185.
Handel, Ben and Jonathan Kolstad, "Health Insurance for Humans: Information Frictions, Plan Choice, and Consumer Welfare," *American Economic Review*, 2015, *forthcoming.*
Handel, Benjamin, "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, 2013, *103* (7), 2643–82.
Hastings, Justine and Jesse Shapiro, "Fungibility and Consumer Choice: Evidence from

Commodity Price Shocks," *Quarterly Journal of Economics*, 2013, *128* (4), 1449–1498.
_ **and** _ , "How Are SNAP Benefits Spent? Evidence From a Retail Panel," *American Economic Review*, 2018.
_ **, Brigitte Madrian, and William Skimmyhorn**, "Financial Literacy, Financial Education, and Economic Outcomes," *Annual Review of Economics*, 2013, *5*, 347–373.
Haviland, Amelia, Matthew Eisenberg, Ateev Mehrotra, Peter Huckfeldt, and Neeraj Sood, "Do "Consumer Directed" Health Plans Bend the Cost Curve Over Time?," *NBER Working Paper 21031*, 2015.
Heiss, Florian, Adam Leive, Daniel McFadden, and Joachim Winter, "Plan Selection in Medicare Part D: Evidence from Administrative Data," *Journal of Health Economics*, 2013, *32* (6), 1325–1344.
Helmchen, Lorens, David Brown, Ithai Lurie, and Anthony Lo Sasso, "Health Savings Accounts: Growth Concentrated Among High-Income Households and Large Employers," *Health Affairs*, 2015, *34* (9), 1594–1598.
Hurst, Eric, "The Retirement of a Consumption Puzzle," *NBER Working Paper 13789*, 2008.
Jappelli, Tullio and Luigi Pistaferri, "The Consumption Response to Income Changes," *Annual Review of Economics*, 2010, *2*, 479–506.
Ketcham, Jonathan, Claudio Lucarelli, Eugenio Miravete, and Christopher Roebuck, "Sinking, Swimming, or Learning to Swim in Medicare Part D," *American Economic Review*, 2012, *102* (6), 2639–73.
KFF, "2016 Employer Health Benefits Survey," Technical Report, Kaiser Family Foundation 2016.
Kling, Jeffrey, Sendhil Mullainathan, Eldar Shafir, Lee Vermeulen, and Marian Wrobel, "Comparison Friction: Experimental Evidence from Medicare Drug Plans," *Quarterly Journal of Economics*, 2012, *127* (1), 199–235.
Kooreman, Peter, "The Labeling Effect of a Child Benefit System," *American Economic Review*, 2000, *90* (3), 571–583.
Laibson, David, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 1997, *112* (2), 443–478.
_ **, Andrea Repetto, and Jeremy Tobacman**, "Self-Control and Saving for Retirement," *Brookings Papers on Economic Activity*, 1998, *1*, 91–196.
Liebman, Jeffrey and Richard Zeckhauser, "Simple Humans, Complex Insurance, Subtle Subsidies," *NBER Working Paper 14330*, 2008.
Lu, Timothy and Olivia Mitchell, "Borrowing from Yourself: The Determinants of 401(k) Loan Patterns," 2010.
Lusardi, Annamaria and Olivia Mitchell, "The Economic Importance of Financial Literacy," *Journal of Economic Perspectives*, 2014, *52* (1), 5–44.
Milkman, Katherine and John Beshears, "Mental Accounting and Small Windfalls: Evidence from an Online Grocer," *Journal of Economic Behavior and Organization*, 2009, *71* (2), 384–394.
Munnell, Alicia and Anthony Webb, "The Impact of Leakages on 401(K)/IRA Assets," Technical Report 2015-2, Boston College Center for Retirement Research 2015.
_ **, Matthew Rutledge, and Anthony Webb**, "Are Retirees Falling Short? Reconciling the Conflicting Evidence," Technical Report 2014-16, Boston College Center for Retirement Research 2014.
O'Donoghue, Ted and Matthew Rabin, "Doing It Now or Later," *American Economic Review*, 1999, *89* (1), 103–124.
Parente, Stephen and Roger Feldman, "Do HSA Choices Interact with Retirement Savings Decisions?," in James Poterba, ed., *Tax Policy and the Economy*, Vol. 22 2008, pp. 81–108.

Parker, Jonathan, Nicholas Souleles, David Johnson, and Robert McClelland, "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 2013, *103* (6), 2530–53.

Peter, Richard and Petra Steinorth, "The Optimal Use of Health Savings Accounts: A Consumer's Perspective," 2012.

_ , Sebastian Soika, and Petra Steinorth, "Health Insurance, Health Savings Accounts, and Healthcare Utilization," *Health Economics*, 2015, p. forthcoming.

Polyakova, Maria, "Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D," *American Economic Journal: Applied Economics*, 2016, *8* (3), 165–195.

Poterba, James, Steven Venti, and David Wise, "The Asset Cost of Poor Health," *NBER Working Paper 16389*, 2010.

Prelec, Drazen and George Loewenstein, "The Red and the Black: Mental Accounting of Savings and Debt," *Marketing Science*, 1998, *17* (1), 4–28.

Rabin, Matthew, "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica*, 2000, *68* (5), 1281–92.

_ and Richard Thaler, "Anomalies: Risk Aversion," *Journal of Economic Perspectives*, 2001, *15* (1), 219–232.

Reed, Mary, Ilana Graetz, Vicki Fung, Joseph Newhouse, and John Hsu, "In Consumer-Directed Health Plans, A Majority of Patients were Unaware of Free or Low-Cost Preventive Care," *Health Affairs*, 2012, *31* (2), 2641–2648.

Sahm, Claudia, Matthew Shapiro, and Joel Slemrod, "Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications," *Tax Policy and the Economy*, 2010, *24*, 69–110.

Schmitz, Hendrik and Nicolas Ziebarth, "Does Price Framing Affect the Consumer Price Sensitivity of Health Plan Choice," *Journal of Human Resources*, 2015, *forthcoming.*

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun, "Are Americans Saving "Optimally" for Retirement?," *Journal of Political Economy*, 2006, *114* (4), 607–43.

Setty, Ofer, "Unemployment Insurance and Unemployment Accounts," *Journal of the European Economic Association*, 2017, *15* (6), 1302–1340.

Shapiro, Matthew and Joel Slemrod, "Did the 2008 Tax Rebates Stimulate Spending?," *American Economic Review*, 2009, *99* (2), 374–79.

Shefrin, Hersh and Richard Thaler, "The Behavioral Life-Cycle Hypothesis," *Economic Inquiry*, 1988, *26* (4), 609–643.

Skinner, Jonathan, "Are You Sure You're Saving Enough for Retirement?," *Journal of Economic Perspectives*, 2007, *21* (3), 59–80.

Steinorth, Petra, "Impact of Health Savings Accounts on Precautionary Savings, Demand for Health Insurance, and Prevention Effort," *Journal of Health Economics*, 2011, *30* (2), 458–465.

Sydnor, Justin, "(Over)insuring Modest Risks," *American Economic Journal: Applied Economics*, 2010, *2* (4), 177–99.

Thaler, Richard, "Mental Accounting and Consumer Choice," *Marketing Science*, 1985, *4* (3), 199–214.

_ , "Anomalies: Saving, Fungibility, and Mental Accounts," *Journal of Economic Perspectives*, 1990, *4* (1), 193–205.

Webb, Anthony and Natalia Zhivan, "What is the Distribution of Lifetime Health Care Costs from Age 65?," Technical Report 10-4, Boston College Center for Retirement Research, Chestnut Hill, MA 2010.

Yamamoto, Dale, "Health Care Costs-From Birth to Death," Technical Report 2013-1, Society of Actuaries 2013.

Zeckhauser, Richard, "Medical Insurance: A Case Study of the Tradeoff Between Risk Spreading and Appropriate Incentives," *Journal of Economic Theory*, 1970, *2* (1), 10–26.

Table 1: Summary Statistics of Sample

	All employees		Self coverage		Family coverage	
	mean	s.d.	mean	s.d.	mean	s.d.
HSA employee contribution ($)	1548.9	1396.8	950.2	798.8	1949.9	1558.8
HSA employer contribution ($)	686.7	382.0	466.2	176.7	834.5	410.3
HSA balance ($)	1022.7	1978.9	799.1	1460.3	1172.5	2248.5
HSA withdrawal ($)	1942.7	1490.2	1134.4	894.9	2484.4	1562.4
401(k) employee contribution ($)	4137.7	4287.3	3572.8	3810.5	4516.1	4540.0
401(k) employer contribution ($)	1842.3	1598.4	1554.8	1244.3	2034.8	1771.1
401(k) balance ($)	47538.9	97809.6	33313.4	69275.5	57054.3	111938.3
Deductible ($)	3267.0	1542.0	2079.6	748.5	4062.7	1420.5
Expected health spending ($)	8248.3	7823.2	5209.8	5348.1	10583.8	8585.5
Salary ($)	61939.2	40837.6	53320.8	29579.8	67712.8	45997.3
Tenure (years)	7.30	6.08	6.50	5.68	7.83	6.27
Age (years)	40.02	9.59	38.98	10.64	40.71	8.75
Female (share)	0.671	0.470	0.690	0.463	0.658	0.474
Married (share)	0.568	0.495	0.293	0.455	0.751	0.432
White (share)	0.726	0.446	0.713	0.452	0.734	0.442
Number of dependents	1.34	1.42	0.00	0.00	2.24	1.16
Job: Finance or Actuary (share)	0.044	0.205	0.044	0.205	0.044	0.206
Office size	1306.3	1478.3	1365.0	1545.8	1267.0	1430.0
Tenure with employer in years	7.297	6.079	6.496	5.680	7.834	6.275
Years with HSA	2.752	1.137	2.632	1.105	2.832	1.151
Zero HSA contribution (share)	0.010	0.102	0.014	0.118	0.008	0.089
Obtain full employer HSA match (share)	0.973	0.161	0.969	0.173	0.976	0.153
Contribute HSA maximum (share)	0.055	0.229	0.060	0.237	0.052	0.223
Zero 401(k) contribution (share)	0.050	0.218	0.048	0.215	0.051	0.220
Obtain full employer 401(k) match (share)	0.498	0.500	0.502	0.500	0.494	0.500
Contribute 401(k) maximum (share)	0.030	0.171	0.019	0.138	0.037	0.189
N	15,908		7,075		9,880	
NT	40,462		17,595		27,398	

Note: This table presents means and standard deviations of the analysis sample by type of coverage. Family coverage also includes coverage for employee plus spouse and employee plus children. N denotes the number of unique employees in the sample and NT denotes the number of employee-years.

Table 2: RD Estimates: HSA and 401(k) saving at match discontinuity

	Linear	Quadratic
	(1)	(2)
Panel A. HSA contributions (employer)		
Estimate	-323.4***	-328.8***
Standard error	(15.2)	(17.2)
Bandwidth ($1,000s)	10.01	15.67
Observations within bandwidth	10,446	17,493
Panel B. HSA disbursements		
Estimate	-220.1***	-267.3**
Standard error	(66.9)	(81.4)
Bandwidth ($1,000s)	8.88	11.49
Observations within bandwidth	8,971	11,890
Panel C. MPC from HSA contributions		
Estimate	0.681	0.813
Standard error	(0.173)	(0.191)
Panel D. 401(k) contributions (Employer + Employee)		
Estimate	113.7	134.3
Standard error	(150.2)	(159.1)
Bandwidth ($1,000s)	7.70	14.47
Observations within bandwidth	7,746	15,430

Note: This table presents results of RD models using either local linear regression (column 1) or quadratic polynomial regression (column 2) for the matching discontinuity at $50,000. Panels A and B present regressions for HSA contributions and HSA withdrawals, respectively. Robust standard errors clustered by employee are shown in parentheses. The MSE-optimal bandwidth and number of observations within the bandwidth are displayed for each model. The estimate of the MPC in Panel C is calculated by taking the ratio of the point estimates from Panels B and A, and the standard errors for Panel C are calculated by a block bootstrap.
* $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Table 3: OLS regressions of dominated saving choices

	Dep var: Dominated choice (1=yes, 0=no)	Dep var: Dominated 401(k) contributions ($)	Dep var: Foregone tax benefits ($)
	(1)	(2)	(3)
Quintiles of salary (rel. to lowest)			
2nd	-0.007	31.212	-1.239
	(0.005)	(27.263)	(10.393)
3rd	-0.013*	233.392***	150.022***
	(0.006)	(31.178)	(13.437)
4th	-0.075***	549.249***	306.628***
	(0.008)	(33.243)	(14.307)
Highest	-0.181***	796.665***	347.662***
	(0.010)	(39.011)	(16.4)
Age in years	0.000	13.206***	5.593***
	(0.000)	(1.363)	(0.558)
Tenure with employer in years	0.001	6.588**	2.428*
	(0.001)	(2.202)	(0.975)
Years with an HSA	-0.020***	-20.985	-8.185
	(0.004)	(16.403)	(7.166)
Finance or actuary	-0.026	40.006	1.724
	(0.017)	(54.920)	(23.741)
Year fixed effects	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes
State fixed effects	Yes	Yes	Yes
Dependent variable mean	0.902	1,095.294	464.718
N	13,830	12,476	12,476
R-squared	0.092	0.331	0.333

Note: This table presents OLS regressions of factors associated with dominated saving decisions. Column 1 includes the sub-sample of employees whose 401(k) contributions exceed the employer match. Columns 2 and 3 include the sub-set of employees in Column 1 who make a dominated saving choice by not contributing the maximum to their HSA. Column 1 presents linear probability models of dominated choices, Column 2 presents regressions of the annual dollar amount of 401(k) contributions that exceed the employer match, and Column 3 shows regressions of foregone tax benefits, calculated as the quantity in column 2 multiplied by the employee's marginal tax rate. Employees in the top two income quintiles are less likely than employees in the lowest income quintile to make a dominated choice, although even in the top quintile, over 70 percent make a dominated choice. Conditional on making a dominated choice, Columns 2 and 3 show that higher-income employees leave more money on the table than lower-income employees. This pattern is explained by the higher 401(k) contributions of higher-income employees. Demographic controls include indicators for self coverage, female, married, and number of dependents. Robust standard errors clustered by employee in parentheses.
* $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Table 4: OLS Regressions of log total health spending

	(1)	(2)	(3)
Quintiles of HSA assets (relative to lowest)			
2nd	0.148***	0.223***	0.165**
	(0.045)	(0.050)	(0.063)
3rd	0.237***	0.435***	0.334***
	(0.048)	(0.058)	(0.073)
4th	0.284***	0.582***	0.460***
	(0.050)	(0.065)	(0.083)
Highest	0.373***	0.774***	0.667***
	(0.057)	(0.079)	(0.100)
Quintiles of 401(k) assets (relative to lowest)			
2nd	0.049	0.092	0.103
	(0.045)	(0.055)	(0.068)
3rd	0.081	0.208**	0.225*
	(0.051)	(0.074)	(0.093)
4th	0.062	0.180	0.139
	(0.059)	(0.098)	(0.120)
Highest	0.058	0.216	0.111
	(0.078)	(0.137)	(0.162)
Year fixed effects	Yes	Yes	Yes
Controls	Yes	No	Yes
Employee fixed effects	No	Yes	Yes
N	33,842	44,784	33,842
R-squared	0.287	0.009	0.022

Note: This table presents regressions of log total health spending (annually) against quintiles of HSA assets, controlling for 401(k) assets and other covariates. All regressions include year fixed effects and a constant. Controls include salary quintiles, deductible chosen, age, number of dependents, tenure in years, and indicators for self coverage, female, married, and early HSA adoption. Robust standard errors clustered by employee in parentheses.
* p<0.05, ** p<0.01, *** p<0.001.

Table 5: Plan choice regressions

	Fixed effects logit		Random effects ordered logit
	Dep. variable = 1 if given deductible chosen, 0 otherwise		Dep. variable equals deductible chosen
	Lowest deductible	Highest deductible	
Variables in $100s	(1)	(2)	(3)
HSA assets	0.989***	1.014***	1.019***
	(0.003)	(0.003)	(0.001)
401(k) assets	1.000	1.000	1.000***
	(0.000)	(0.000)	(0.000)
Expected health spending	1.007***	0.993***	0.979***
	(0.001)	(0.001)	(0.000)
N	5,685	5,163	27,460
Log-likelihood	-1986.12	-1738.31	-27364.41

Note: This table presents regression results that examine within-employee variation in deductible choices. Coefficients are exponentiated and interpreted as odds ratios. Columns 1 and 2 present the results of fixed effects logits of choosing either the lowest deductible (Column 1) or the highest deductible (Column 2). Column 3 presents random effects ordered logits. Regressions also include age, salary, number of dependents, in age, salary, tenure, coverage type, and year effects. Both sets of regressions indicate that higher HSA balances are positively correlated with higher deductible choices. Robust standard errors clustered by employee in parentheses.
* $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Figure 1: Variation in deductible levels, 2007-2010

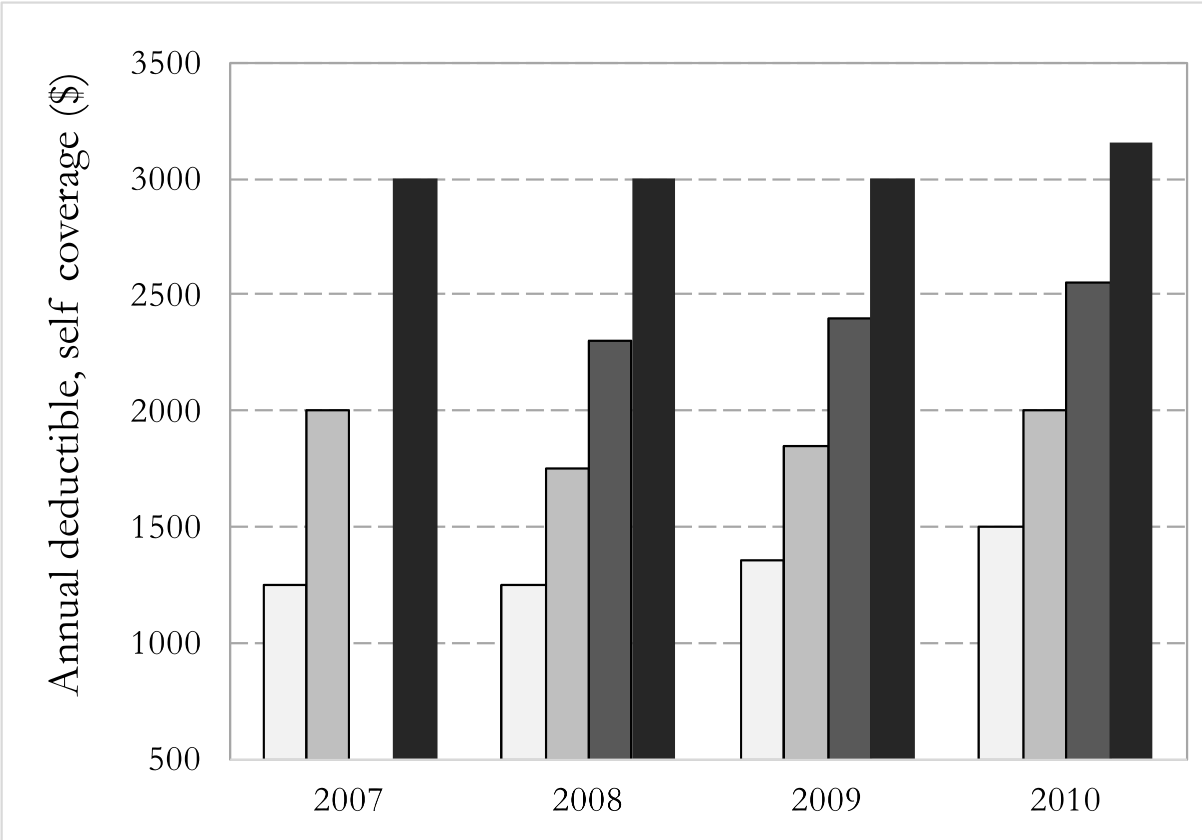

Note: This figure plots the annual deductibles offered to employees for self coverage between 2007 and 2010. Deductibles for family coverage were twice these amounts. In 2008, a fourth plan was added. In 2009, the three lowest deductibles were increased by $100 (and by $200 for family coverage) over 2008 levels. In 2010, each deductible was increased by $150 (and by $300 for family coverage) over 2009 levels.

Figure 2: HSA contributions by salary level

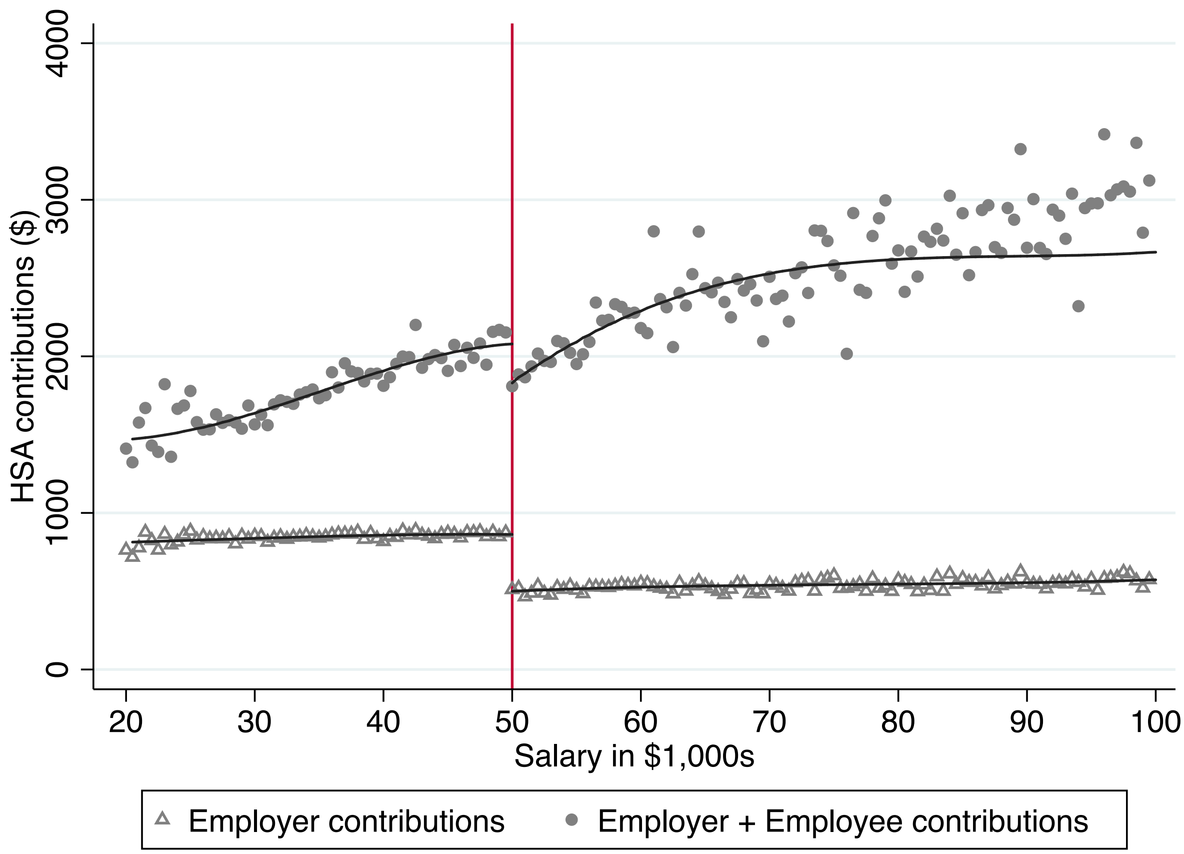

Note: This figure plots means of employer HSA contributions (triangles) and total HSA contributions (circles) within $500 salary bins. For each variable, global third-order polynomials are fit to the data below and above the $50,000 match discontinuity. Data includes both self and family coverage. There is an average drop of about $320 in employer contributions at the discontinuity. Total HSA contributions are increasing in salary. There is a perceptible jump in the regression function for total contributions at the match discontinuity, indicating that employees do not fully reduce their own HSA contributions in response to the match.

Figure 3: HSA withdrawals by salary level

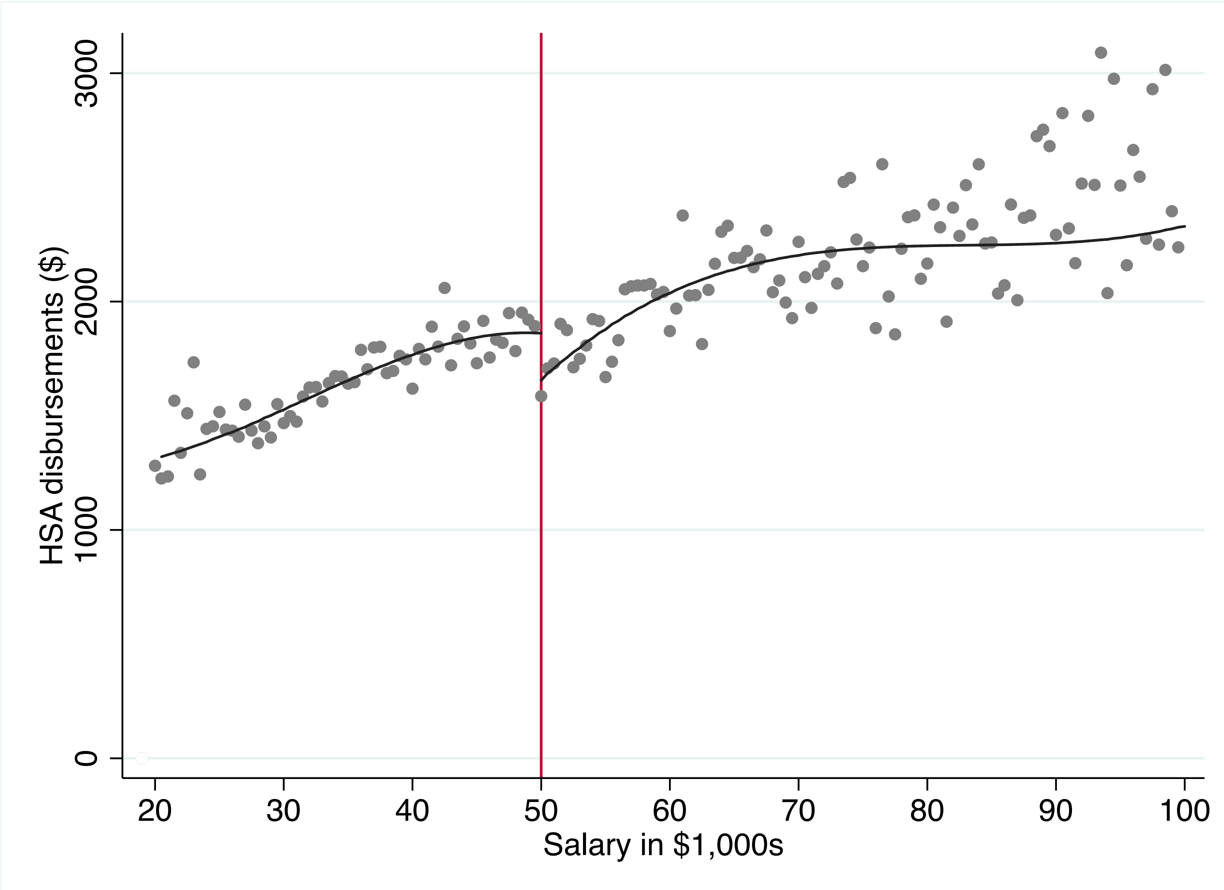

Note: This figure plots means of HSA withdrawals within $500 salary bins, and fits separate global third-order polynomials to the data below and above the $50,000 match discontinuity. Data includes both self and family coverage. Employees with higher salaries make larger HSA withdrawals, but there is a perceptible jump in the regression function at the match discontinuity.

Figure 4: Total 401(k) contributions by salary level

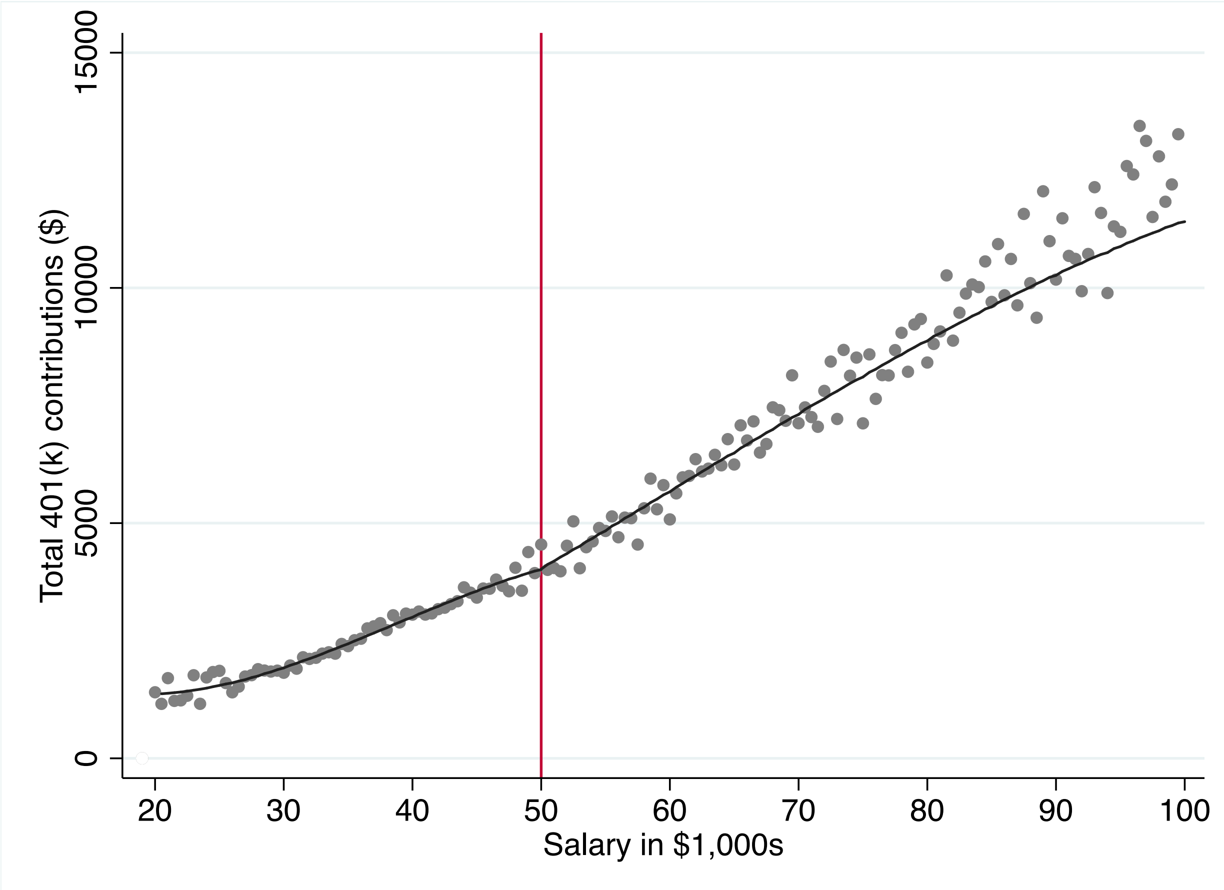

Note: This figure plots means of 401(k) contributions (employer + employee) within \$500 salary bins, and fits separate global third-order polynomials to the data below and above the \$50,000 HSA match discontinuity. Data includes both self and family coverage. There is a strong, positive relationship between salary and 401(k) contributions, and no perceptible jump in the regression function at the HSA match discontinuity. This pattern implies there is no crowd-out of 401(k) saving from an exogenous increase in HSA funds at a \$50,000 salary level.

Figure 5: Time Series of Monthly Spending, CPI and Age-Adjusted

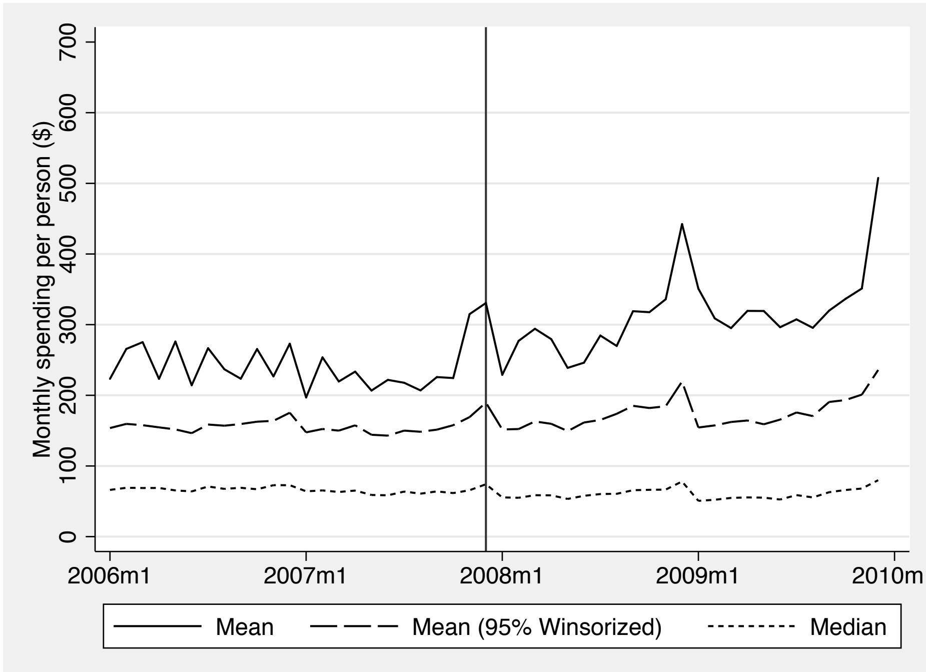

Note: This graph plots monthly spending between 2006 and 2009, adjusted for changes in the medical care component of the consumer price index and for aging as described in the text. The vertical line denotes the final month of the 2007 plan year, the last month before employees enrolled in the HDHP/HSA contract. The graph illustrates that spending did not decline after the introduction of HDHP/HSAs in 2008. Spending rises at the end of each plan year, and then falls again in the first month of the subsequent year. These patterns are also observed when excluding outliers, as shown by trends in the winsorized mean (replacing spending above the 95th percentile to the 95th percentile) and the median.

Figure 6: HSA contributions as supplemental insurance: continuous choice of insurance contract

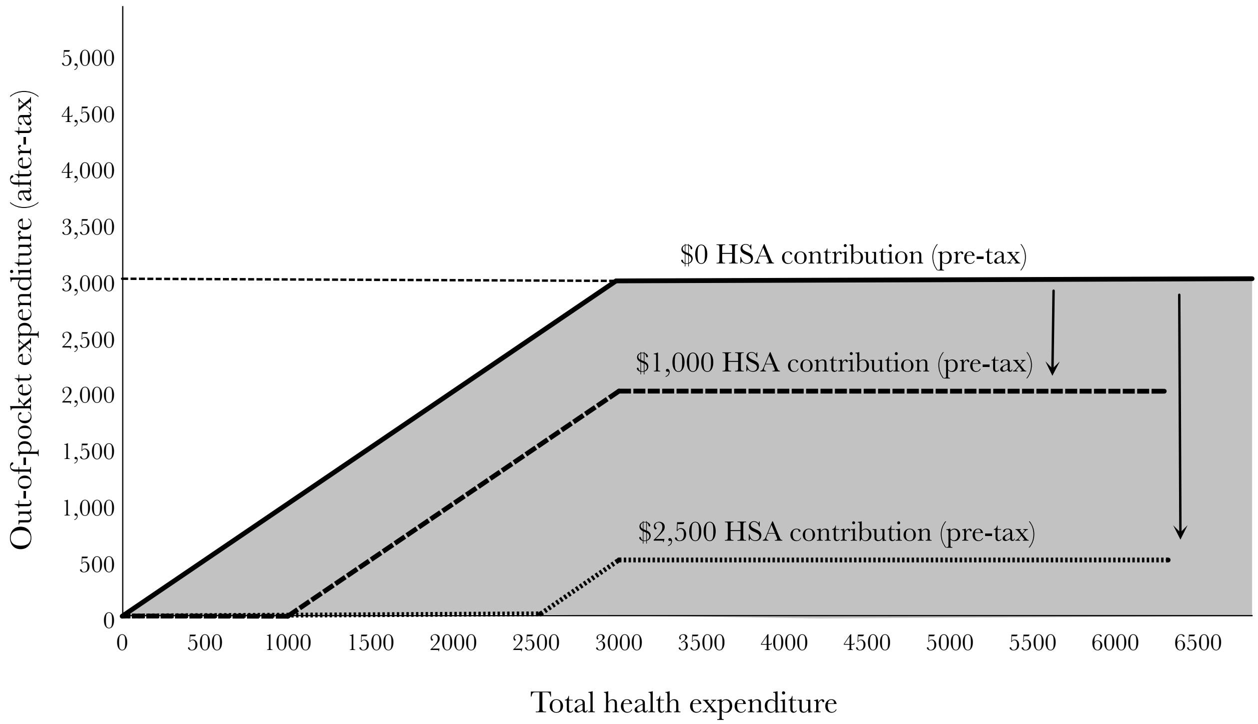

Note: The solid line represents the benefit schedule for a contract with a $3,000 deductible and full coverage past the deductible: the consumer pays for all expenses out-of-pocket until total costs reach $3,000 and the insurance plan pays for all costs beyond this amount. The dashed and dotted lines represent HSA contributions used to provide coverage against out-of-pocket payments associated with the deductible. For example, a $2,500 contribution shifts the benefit schedule towards the horizontal axis, so that the first $2,500 of costs are financed from the HSA contribution with pre-tax funds, the next $500 of costs are financed with after-tax funds, and all costs beyond $3,000 are paid by the insurance plan. With HSA contributions, any contract in the shaded area is feasible so that the choice of contract becomes continuous, in effect.

Figure 7: Example: Deductible choices by MPC and HSA balances

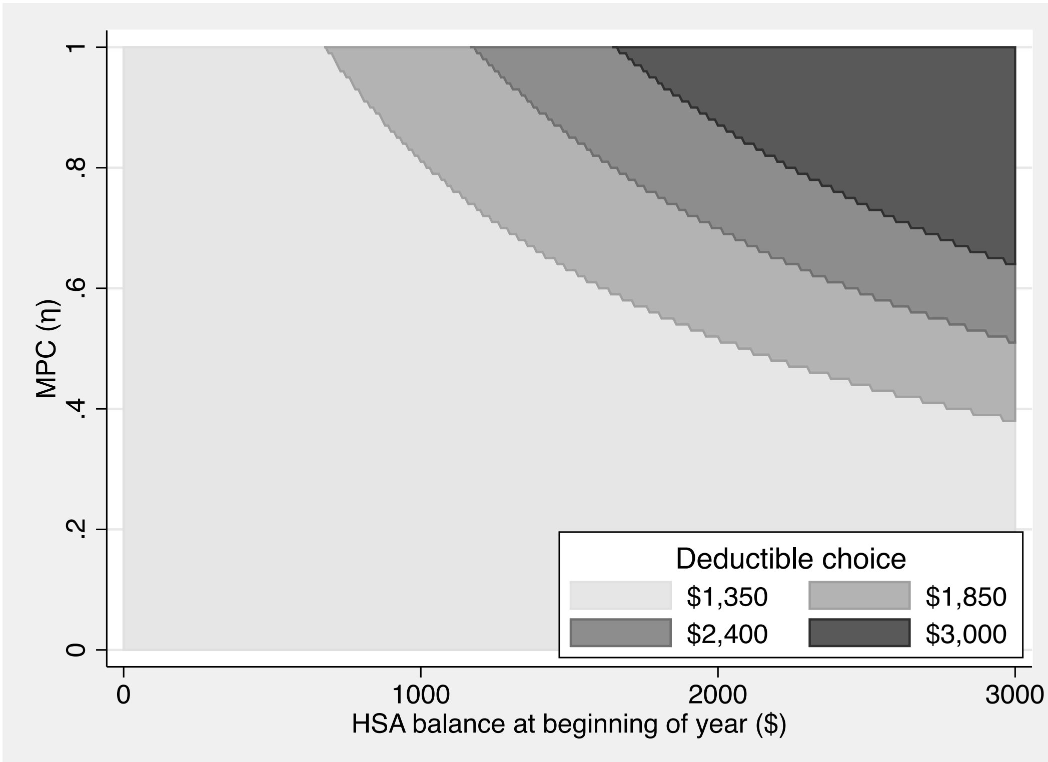

Note: This graphs plot the chosen deductible assuming a coefficient of absolute risk aversion of 2.5x10^{-4}, the distribution of spending for a 35-44 year old female in the median quintile of health risk, a marginal tax rate of 20 percent, and a 6:1 employer match on eligible HSA contributions (income below $50,000). The construction of quintiles of health risk are described in Appendix B. Both a higher MPC and higher HSA balances can induce consumers to choose higher deductibles.

Appendix A: Background on Health Savings Accounts [For Online Publication]

Created with the 2003 Medicare Modernization Act, an HSA is a portable financial account that must be paired with a HDHP. In 2016, the statutory minimum deductible—the amount paid by the consumer before coverage begins—for a HDHP was $1,300 for self coverage and $2,600 for family coverage. HSA account holders cannot use their HSA if they switch to other types of health insurance while working. HSA contributions roll over each year, unlike Flexible Spending Accounts (FSAs) where the enrollee loses unused balance at year's end.[12,13] Medicare beneficiaries can make HSA withdrawals, but not contributions.

HSAs, which are owned by individuals, closely resemble 401(k)s as shown in Table 1: contributions to HSAs are deductible from taxable income ("above the line"), interest grows tax-deferred, and withdrawals for non-medical consumption are subject to income taxation and a penalty if before age 65.[14] HSAs also offer survivor benefits, similar to 401(k)s.[15] However, HSAs provide superior tax advantages to 401(k)s since withdrawals at any age to finance qualified medical care are not counted as taxable income. Qualified expenses, which must be incurred after the HSA has been established, includes out-of-pocket payments associated with a HDHP while working, as well as premiums for Medicare, COBRA, or long-term care insurance.[16] HSA balances cannot be used tax-free to finance premiums for employer health insurance or Medigap. Another tax advantage relative to 401(k)s is that employee contributions made as payroll deductions are not subject to FICA taxes (Social Security and Medicare) and employer contributions are not subject to FICA or FUTA (unemployment insurance) taxes.[17]

[12] In October 2013, the Treasury Department announced FSA balances up to $500 could be rolled over from one year to the next. This paper studies pre-2010 choices, before this rule went into effect.

[13] FSAs are compatible with traditional insurance and do not need to be paired with a HDHP as HSAs do. HSAs are also distinguished from Health Reimbursement Arrangements (HRAs) and Archer Medical Savings Accounts (MSAs). HRAs are owned by the employer, not the employee, and cannot be invested. The employer funds the account for qualified expenses and may decide whether HRA funds roll over from year to year or are forfeited at year's end. HRAs do not have to be paired with a HDHP. Archer MSAs, created as a pilot program in the mid-1990s and ended in 2007, were more similar to HSAs in that funds roll over, accounts are portable, and HDHP insurance is required. Eligibility for Archer MSAs was restricted to self-employed or employers with fewer than 50 employees.

[14] The penalty for early withdrawals was 10 percent before 2011, equal to that of 401(k) withdrawals, and increased to 20 percent beginning in 2011.

[15] HSAs obtained through employers may or may not qualify as an ERISA plan, depending on the employer's involvement in the plan. HSAs not opened through an employer are considered personal savings vehicles by the Department of Labor and not protected under ERISA.

[16] Until 2011, over-the-counter drugs without a prescription were included as qualified medical expenses. Starting January 1, 2011, a prescription was needed for over-the-counter drugs to be financed tax-free with HSA funds.

[17] A policyholder can also pay out-of-pocket (with taxable savings or current income) for health expenses while enrolled in an HSA, save the receipts, and reimburse herself from the HSA at any time in the future, even decades later. In effect, this strategy transforms part of the HSA balance into a Roth IRA because

In 2016, the annual HSA contribution limit including both employer and employee contributions was \$3,350 for self-only coverage and \$6,750 for family coverage. These limits, which the IRS increases over time for cost of living adjustments, rose from \$2,850 and \$5,650, respectively, in 2007. Individuals over age 55 could also make annual "catch-up" contributions of an extra \$1,000 in 2016 to their HSAs, up from \$800 per year in 2007.[18] By way of comparison, the IRS's annual limit on employee 401(k) elective deferrals was \$18,000 in 2016 and employers could provide an additional \$53,000 in defined contributions to an employee's 401(k) in the same year.

Table A.1: Comparison between 401(k) and HSA features

	401(k)	HSA
Contributions exempt from income tax	Yes	Yes
Contributions exempt from FICA taxes	No	Yes
Interest grows tax-deferred	Yes	Yes
Withdrawal for medical care always tax free	No	Yes
Penalty for early withdrawal before 2011	10%	10%
Age when can withdraw penalty-free	$59\frac{1}{2}$	65
Annual contribution limit (incl. employer), 2016	\$53,000 self	\$3,350 self /\$6,750 family

Note: This table presents the major similarities and differences between HSAs and 401(k)s. In 2011, the penalty for early withdrawn from the HSA was raised from 10% to 20%. The IRS rules for catch-up contributions also differ between accounts. People can contribute an extra \$6,000 starting at age 50 to their 401(k) versus \$1,000 starting at age 55 to their HSA.

While HSA saving limits are substantially lower than 401(k) limits, the costs of health care in retirement is large enough for the average consumer to consider HSAs as an important vehicle for long-term saving. On average, Medicare beneficiaries spend close to \$5,000 annually out-of-pocket on premiums, long-term care facilities, and other services (Cubanksi et al. 2014). The net present value of out-of-pocket expenses not covered by Medicare at age 65 has been estimated at between \$220,000 and \$376,000 dependent on the time period and whether long-term care is included (Fronstin et al. 2008; Webb and Zhivan 2010; Yamamoto 2013; Fidelity 2014). Since there is a distribution around these costs, risk-averse workers will want to save more than these averages to guard against high cost realizations. For example, Webb et al. (2010) estimate that out-of-pocket costs exceed \$570,000 with a five percent probability in 2009—the last year of my sample period. In order to finance these costs with HSA funds, a consumer would need to fully fund his HSA without any withdrawals while working. This comparison is important because it indicates the average consumer can meaningfully trade off current versus future health care consumption in making HSA saving

taxes are paid on the out-of-pocket payments while working, interest earned on HSA balances is tax-exempt, and the HSA withdrawal in the future is tax-exempt.

[18]Between 2004 and 2006, annual HSA contributions were limited to the lesser of the statutory maximum or the chosen deductible. This restriction biases contribution and deductible choices to be positively correlated. This rule was repealed on December 20, 2006 by the Tax Relief and Health Care Act of 2006. As described in Section 4, I exclude employees who opened an HSA prior to 2008 from my analysis to avoid introducing possible correlation between deductible and saving choices induced by legislation.

decisions.[19]

Since their creation in 2004, HSAs have grown to include 20 million accounts and assets of $37 billion, with the majority of accounts opening since 2011 (Fronstin 2017). The take-up of HDHPs and HSAs has increased dramatically over the last decade, covering almost 30 percent of people who obtain insurance through their employer (Kaiser Family Foundation 2016). Over one-quarter of firms now offer an HDHP/HSA option, and nearly half of firms with over 1,000 workers do. Despite their increased popularity, contributions, have been modest. Based on the Employee Benefits Research Institute (EBRI) database comprising one fifth of accounts and assets nationwide, the average contribution (including employer contributions) in 2015 was about $2,000, with about half receiving employer contributions (Fronstin 2016). Roughly 10 percent of account holders contributed the maximum amount (Fronstin 2015). On average, end-of-year account balances amounted to $1,844 in 2015, up from $1,332 the previous year (Fronstin 2016). Assets in the large majority of accounts are not invested in financial markets.[20] In a first study using tax records, Helmchen et al. (2015) find older and higher-income workers opened and fully funded their HSAs more often than did younger and lower-income workers.[21]

Evidence of how HDHPs affect health care costs is still emerging, but it points to moderate declines in spending compared to traditional insurance of around 15 percent (Buntin et al. 2011; Borah et al. 2011; Haviland et al. 2015; Brot-Goldberg et al. 2017). In one setting, the move from no out-of-pocket payments to an HDHP produced spending cuts across all services (Brot-Goldberg et al. 2017). Many consumers do not appear to be forward-looking even over the course of a year: those with low shadow prices based on their expected spending relative to the deductible reduced spending while below the deductible, thus reacting to spot prices instead. If such demand responses reflect behavioral hazard as well as moral hazard (Baicker et al. 2015), then such documented consumption drops cannot quantitatively estimate a reduction in moral hazard.

Limited research exists on how employees fund HSAs in relation to 401(k)s or how

[19]Although there is some debate about whether Americans are prepared for retirement (see e.g. Engen et al. (1999), Scholz et al. (2006), Hurst (2008), Munnell et al. (2014)), there is consensus that risks to health represent an important component in retirement saving decisions. Both the direct cost of poor health in terms of medical care and the indirect costs, such as the inability to substitute home production for purchased goods, reduce assets of retirees (Skinner 2007; Poterba et al. 2010; DeNardi et al. 2010).

[20]Keeping HSA funds in low-yield saving accounts may reflect account rules on minimum balances or a lack of consumer information. However, since money is fungible, such behavior may actually be a financially savvy investment strategy if consumers can rebalance their 401(k) and IRA by selling bonds and cash, keeping their total asset allocation unchanged, because HSAs provide option value before retirement in terms of tax-free withdrawals for medical care.

[21]Many states exclude HSA assets in determining Medicaid eligibility, so low contributions among lower-income workers are likely not due to consideration about Medicaid enrollment.

HSAs interact with deductible choices. Parente and Feldman (2008) find a weak positive correlation between contributions to HSAs and other tax-deferred retirement saving vehicles among one set of University employees. Yet, their sample included just 63 HSA accounts from a sample of 16,000—a take-up rate of just 0.4 percent—and their results were not robust to alternative specifications. Analyzing over 160,000 accounts held at United Health Group and OptumHealth Bank, Chen et al. (2013) also take a descriptive approach and find HSA contributions were negatively correlated with employer contributions and positively correlated with age, income, education, and health care spending. They estimate a positive but insignificant correlation with the deductible, and they lacked data on 401(k)s or other saving accounts. Peter and Steinorth (2012) simulate a life-cycle model where health care spending and lifespan were stochastic (with certain labor income), but made the assumption that individuals max out HSA contributions, which is at odds with observed contribution levels. Other studies on HSAs focus on the choice of insurance plan, comparing traditional insurance to a high-deductible health plan (Cardon and Showalter, 2007; Steinorth, 2011; Handel, 2013; Handel and Kolstad, 2015), but do not study saving decisions. Cardon and Showalter (2001) and Cardon (2012) discuss the role of FSA contributions as supplemental insurance coverage, which also applies to HSAs. Peter et al. (2015) theoretically analyze how HSA saving may influence insurance decisions, but do not consider HSAs in the context of other retirement saving.

Appendix B: Construction of Out-of-Pocket Cost Distributions [For Online Publication]

This section describes in detail the procedure for constructing distributions of out-of-pocket costs for each insured family (employee only or the employee and dependents). It follows similar methods as Handel (2013) and Handel and Kolstad (2015). This cost model assumes that there is no moral hazard and that each person in the same risk group holds the same beliefs about his or her ex ante health expenditure risk. There are four steps to construct the distributions from the inputs of expenditure claims and the employer's severity score.

1. Group each insured individual i into risk group z based on age, sex, and health status
2. For each risk group, construct a Weibull distribution, G_z, that is modified to allow for the possibility of zero expenditure using observed total health expenditure m from the following year
3. For each person in risk group z, simulate expenditure draws from G_z and add up the draws within each family k to create an ex ante distribution of total health expenditure risk G_k for family k
4. For each family k, map the distribution of expenditure risk G_k to out-of-pocket costs under deductible j to create a family-specific ex ante distribution of out-of-pocket costs F_{jk} of choosing deductible j

Risk groups: Each individual i is first categorized into risk group z based on their age, sex, and quintile of the severity score. The age bins used are 0-14, 15-24, 25-34, 45-64, 65 and older. The severity score, which is recorded on each insurance claim, measures the expected health spending for that enrollee using a proprietary formula constructed by the company. The score is not used for payment purposes, but rather represents the employer's actuarial forecast about that person's expenditure risk. I use the severity score captured on the last claim before the start of the plan year as a measure of health status during the open enrollment period. For each age bin and each sex, I classify individuals into quintiles of the severity score within that age-sex cell. I pool years 2007 through 2009 together to ensure adequate sample sizes. This process results in 60 risk groups based on six age bins, sex, and quintiles of severity score.

Expenditure distributions by risk group: After the risk groups are defined, the observed expenditures for each person in the group the following year are used to estimate an ex ante expenditure distribution for that group. Denote the empirical distribution of claims the following year by $\widehat{G}_{I_z}$. In constructing this distribution, expenditures on preventive care are excluded since such services are covered free of charge by all plans. Only claims from in-network providers are considered, which comprise over 95 percent of all spending. I continuously fit this empirical distribution using a Weibull distribution with a mass of claims at zero to generate an ex ante distribution of expenditure risk, consistent with prior work (Handel 2013; Kolstad and Handel 2015).

The creation of this ex ante distribution of expenditure by risk group involves two steps to deal with the mass of expenditure at zero. First, for each risk group k, the empirical probability of zero expenditure is used to construct the mass of expenditure realizations at zero, denoted $\widehat{G}_{I_z}(0)$. Second, a Weibull distribution is fitted to the observed expenditures that are positive in that risk group by maximizing the following likelihood with respect to the scale parameter α and shape parameter β:

$$\prod_{i \in I_z} \frac{\beta_z}{\alpha_z} \left(\frac{m_i}{\alpha_z} \right)^{\beta_z - 1} e^{-\left(\frac{m_i}{\alpha_z}\right)^{\beta_z}}$$

Denote $\widehat{\alpha_z}$ and $\widehat{\beta}_z$ as the estimated parameters and $W\left(\widehat{\alpha_z}, \widehat{\beta}_z\right)$ as the distribution of positive expenditure in risk group z. The (ex ante) distribution for expenditure in risk group z is then:

$$G_z = \begin{cases} \widehat{G}_{I_z}(0) & \text{if } m = 0 \\ \widehat{G}_{I_z}(0) + \frac{W\left(\widehat{\alpha_z}, \widehat{\beta_z}\right)}{1 - \widehat{G}_{I_z}(0)} & \text{if } m > 0 \end{cases}$$

Simulated expenditures: For each insured individual within each risk group, 100 draws are simulated from the corresponding expenditure distribution G_z. Then within each family k, the expenditures for each draw from each member are summed, so that each family has 100 draws corresponding to the family's total expenditure. This statistical object, denoted G_k, represents the beliefs of family k about its total health expenditure risk. Since families differ in their compositions by age, sex, severity score, and size, this classification by risk group results in over 2,500 different combinations of expected spending in the sample.

Appendix C: Additional analyses [For Online Publication]

Table C.1: Characteristics of early and late HSA adopters

	Early HSA adopter	Late HSA adopter
HSA employee contribution ($)	1,528.4	1,087.2
HSA employer contribution ($)	727.6	763.6
HSA withdrawal ($)	1,893.3	1,260.6
401(k) employee contribution ($)	4,357.0	3,469.9
401(k) employer contribution ($)	1,942.3	1,621.0
401(k) balance ($)	52,151	45,913
Deductible ($)	3,172.5	3,013.8
Expected health spending ($)	8,178.8	6,214.7
Salary ($)	62,797	54,312
Tenure (years)	6.26	7.50
Age (years)	39.24	40.04
Female (share)	0.65	0.72
Married (share)	0.58	0.54
White (share)	0.74	0.68
Number of dependents	1.36	1.16
Self coverage (share)	0.39	0.44

Note: This table presents means for employees who adopted the HSA prior to 2008 ("early adopters") and employees who adopted the HSA in 2008 or later after the full replacement of traditional plans. On average, early adopters have higher incomes, contribute and withdraw more from their HSA, have higher expected health spending, and have higher 401(k) saving. The differences are statistically significant for each variable between early and late adopters at conventional levels.

Table C.2: Covariate Balance Tests

Dependent variable	Point est.	s.e.	Control mean	Percent difference
Age in years	-2.80***	(0.520)	40.91	-6.8
Female	0.005	(0.026)	0.597	0.8
Married	0.017	(0.030)	0.659	2.6
White	0.042	(0.025)	0.794	5.3
Black	-0.035	(0.019)	0.078	-44.9
Self coverage	0.026	(0.025)	0.351	7.4
Tenure in years	-0.463	(0.382)	6.647	-7.0
Years in HSA	-0.012	(0.031)	1.305	-0.9
Early HSA adopter	-0.001	(0.027)	0.756	-0.1
Number of dependents	-0.052	(0.072)	1.489	-3.5
Job group: practitioner	0.081*	(0.033)	0.470	17.2
Job group: manager	-0.000	(0.021)	0.227	0.0
Job group: sales	-0.004	(0.007)	0.125	-3.2
Job group: staff	0.022	(0.025)	0.097	22.7
Job group: directors	-0.000	(0.002)	0.028	0.0

Note: This table presents RD regression results of covariates using the $50,000 salary cutoff. Each row presents the results from a different RD model using local linear regression and the MSE-optimal bandwidth. For most covariates, the point estimates are not statistically significant from zero. The third column lists the mean of the variable for those between $50,000 and $99,999 (the "control mean"). The fourth column lists the point estimate as a percentage of the control mean. Robust standard errors clustered by employee in parentheses.
* $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Figure C.1: Manipulation test

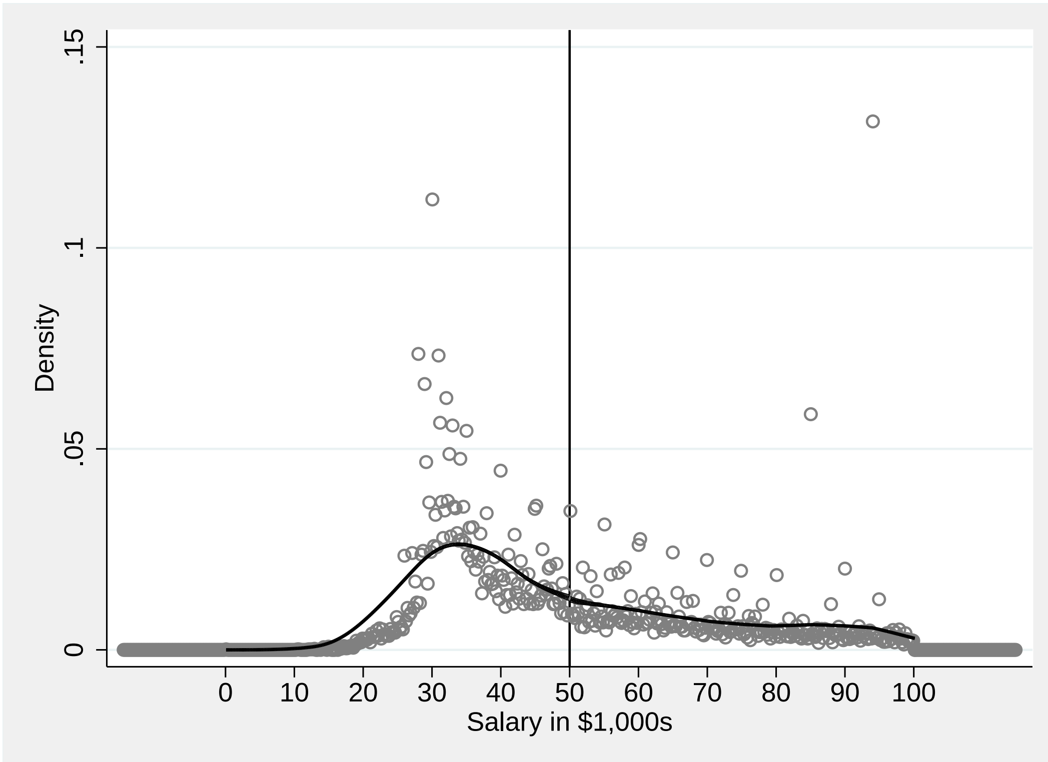

Note: This figure presents the McCrary (2008) test that the density of salary is smooth across the match discontinuity. There is no evidence of manipulation of the running variable on either side of the threshold, providing support to the validity of the RD design.

Table C.3: RD Regressions with Controls

	Linear	Quadratic
	(1)	(2)
Panel A. HSA contributions (employer)		
Estimate	-291.9***	-291.3***
Standard error	(12.7)	(13.4)
Bandwidth ($1,000s)	3.81	7.08
Observations within bandwidth	3,665	7,115
Panel B. HSA disbursements		
Estimate	-191.1***	-197.0**
Standard error	(50.7)	(60.1)
Bandwidth ($1,000s)	11.56	16.39
Observations within bandwidth	11,879	18,640
Panel C. MPC from HSA contributions		
Estimate	0.655	0.676
Panel D. 401(k) contributions (Employer + Employee)		
Estimate	140.0	148.0
Standard error	(146.7)	(161.2)
Bandwidth ($1,000s)	7.69	13.10
Observations within bandwidth	7,631	13,654

Note: This table presents results of RD models using either local linear regression (column 1) or quadratic polynomial regression (column 2) for the matching discontinuity at $50,000. Each models includes as controls. Panels A and B present regressions for HSA contributions and HSA withdrawals, respectively. Robust standard errors clustered by employee are shown in parentheses. The MSE-optimal bandwidth and number of observations within the bandwidth are displayed for each model. The estimate of the MPC in Panel C is calculated by taking the ratio of the point estimates from Panels B and A.
* $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Table C.4: Placebo RD Estimates

	Placebo salary cutoff:									
	$25k	$30k	$35k	$40k	$45k	$55k	$60k	$65k	$70k	$75k
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A. HSA contributions	-4.4	2.0	-9.9	-23.0	39.7*	-26.0	7.9	-12.2	33.4	-11.9
(employer)	(38.1)	(16.9)	(13.7)	(16.7)	(16.9)	(22.3)	(19.7)	(18.9)	(23.2)	(25.3)
Bandwidth ($1,000s)	5.0	5.1	5.1	5.7	7.4	4.4	5.6	8.3	7.7	5.4
Observations in bandwidth	2,889	10,296	12,353	9,247	9,093	3,445	4,024	5,050	3,900	2,519
Panel B. HSA disbursements	-6.9	-38.4	-38.1	-49.4	-108.9	-27.6	-127.7	-43.5	85.7	-229.9
	(108.7)	(54.4)	(42.6)	(68.1)	(59.7)	(67.9)	(84.1)	(101.6)	(106.5)	(128.7)
Bandwidth ($1,000s)	4.4	3.5	5.9	5.1	8.3	13.6	9.9	9.7	9.7	7.0
Observations in bandwidth	2,483	7,391	13,493	7,997	10,279	11,881	6,763	5,783	4,919	3,191

Note: This table presents results of RD models using local linear regression at placebo cutoffs for which there is not a change in matching rates. The magnitudes of the point estimates are generally small compared to the results in Table 2. One of the 20 estimates is statistically different from zero at the 5 percent level, as expected by chance. These patterns reinforce the causal interpretation given to the estimates in Table 4. * p<0.05, ** p<0.01, *** p<0.001.

Table C.5: RD Analysis for Sub-Samples

	Auto-enroll in 401(k)	Self coverage	Family coverage	Early HSA adopter	Late HSA adopter	No 401(k) loan	With 401(k) loan
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. HSA contributions (employer)							
Estimate	-303.5***	-181.3***	-399.4***	-311.2***	-349.3***	-289.8***	-386.1***
Standard error	(42.3)	(11.2)	(15.1)	(19.0)	(30.8)	(23.4)	(29.1)
Bandwidth ($1,000s)	10.6	3.9	5.4	9.8	5.8	5.3	7.0
Observations within bandwidth	1,714	1,636	3,106	6,584	1,908	3,716	2,258
Panel B. HSA disbursements							
Estimate	-214.2	-138.8*	-211.5**	-270.9**	-184.2	-204.1**	-183.5
Standard error	(135.7)	(56.6)	(78.6)	(90.9)	(97.4)	(69.9)	(108.8)
Bandwidth ($1,000s)	15.0	12.2	11.9	7.2	10.2	12.8	10.3
Observations within bandwidth	2,679	5,525	6,959	4,907	3,498	9,064	3,419
Panel C. MPC from HSA contributions							
Estimate	0.706	0.766	0.530	0.871	0.527	0.704	0.475

Note: This table presents results of RD models using local linear regression for the matching discontinuity at $50,000 corresponding to different sub-samples. Panels A and B present regressions for HSA contributions and HSA withdrawals, respectively. Robust standard errors clustered by employee are shown in parentheses. The MSE-optimal bandwidth and number of observations within the bandwidth are displayed for each model. The estimate of the MPC in Panel C is calculated by taking the ratio of the point estimates from Panels B and A. The estimated MPC among those who auto-enroll in the 401(k) is 0.706, which is very similar to the main estimate of 0.681 in Table 2. Columns 2 to 4 reveal heterogeneity in the estimated MPC by type of coverage and early adoption of the HSA, although the MPC exceeds 0.5 in all sub-samples. Columns 6 and 7 show those without a 401(k) loan have higher MPC than those with a 401(k) loan. * $p<0.05$, ** $p<0.01$, *** $p<0.001$.