American Economic Journal: Applied Economics 2022, 14(2): 200–227
https://doi.org/10.1257/app.20200135

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

By Adam Leive*

To lower health care costs, Health Savings Accounts (HSAs) offer tax incentives encouraging people to trade off current consumption against future consumption. This paper tests whether consumers use HSAs as self-insurance over the life cycle. Using administrative data from a large employer and a regression discontinuity design, I estimate the marginal propensity to consume from HSA assets is 0.85 and reject the neoclassical benchmark of 0. Comparisons with 401(k) saving show most employees do not treat HSA money as fungible with retirement savings. In this setting, HSAs did not reduce health spending and instead increased the share that was financed tax-free. (*JEL* D15, D82, G22, G51, I13)

A primary tension in insurance design is between risk spreading and moral hazard (Arrow 1963; Pauly 1968; Zeckhauser 1970). An innovative approach to balance these forces is to link a personal savings account with current insurance benefits. In the context of health insurance, Health Savings Accounts (HSAs) offer tax subsidies encouraging people to trade off current health care consumption against future consumption. HSAs function like a 401(k) retirement plan, with the additional feature that withdrawals for health care expenses are tax-exempt. Such accounts must be paired with a high-deductible health plan (HDHP): the objective is to reduce health care spending by exposing consumers to the marginal cost for moderate expenses while compensating them through subsidies for saving. 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; Kling 2006; Feldstein and Altman 2007; Setty 2017).

* Batten School of Leadership and Public Policy, University of Virginia (email: leive@virginia.edu). Neale Mahoney was coeditor for this article. 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 thank my dissertation committee Iwan Barankay, Jonathan Kolstad, Mark Pauly, and Robert Town for guidance, and am grateful for comments on an earlier draft from Zarek Brot-Goldberg, Guy David, Brent Davis, Hanming Fang, Leora Friedberg, Matt Grennan, Scott Harrington, Jesse Hinde, Bill Johnson, Ambar La Forgia, Lu Li, Olivia S. Mitchell, Elena Prager, Preethi Rao, Jonathan Skinner, Amanda Starc, Ashley Swanson, Justin Sydnor, Richard Thaler, Tess Wilkinson-Ryan, Karen Zhang, and many seminar participants. I thank Andrew Chong and Colin Gray for help in processing the data.
† Go to https://doi.org/10.1257/app.20200135 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

Yet HSAs have more complex features than other health insurance products. Consumers may be unaware that HSA balances carry forward, or they may not understand the more generous tax preferences of HSAs compared to 401(k)s. As a result, consumers may not view the account as a savings vehicle, but rather as a way to offset their current deductible. Such behavior could undermine the incentives of these contracts to reduce spending. Given evidence that many people have limited understanding of the financial dimensions of health insurance plans (Loewenstein et al. 2013; Handel and Kolstad 2015) and make errors in insurance choices (Abaluck and Gruber 2016; Ericson and Sydnor 2017; Bhargava, Loewenstein, and Sydnor 2017; Chandra, Handel, and Schwartzstein 2019), how consumers perceive the features of HSAs will affect how the contracts work in practice. Thirty percent of US workers are now enrolled in these plans (Claxton et al. 2019), but research on HSAs remains limited.

In this paper, I test whether consumers use their HSAs as self-insurance over the life cycle. Answering this question is challenging for several reasons. First, one must observe consumer saving decisions in response to an exogenous increase in HSA funds. Second, the motives that influence how consumers treat HSAs must be distinguished from risk preferences and time preferences, which also influence insurance and saving choices. Finally, understanding the welfare implications of these accounts requires taking a stand on the optimality of saving decisions, which is generally difficult because many life-cycle factors are unobserved by the econometrician.

I examine this question by studying choices of employees at a large US employer that fully replaced its traditional, low-deductible health insurance offerings with a menu of HDHPs and HSAs for its workers. I combine detailed panel-level data and variation in the employer's benefit design to overcome these empirical challenges. The data include employee and employer contributions to the HSA and 401(k), insurance deductible choices, medical and pharmacy claims, demographics, and information on salary and job characteristics. I exploit variation in the firm's HSA matching rates by salary level to identify the marginal propensity to consume (MPC) from HSA funds using a regression discontinuity design. Employees earning less than \$50,000 annually receive a larger employer contribution than employees earning over \$50,000, creating a sharp discontinuity in HSA balances at this salary cutoff. The MPC is calculated as the ratio of the change in HSA withdrawals to the change in total HSA contributions, using the match discontinuity to instrument for total contributions.[1]

One notable aspect of the firm is that it is a large US health insurer. It is reasonable to believe its employees likely possess a relatively high degree of health insurance literacy as a result. The setting may therefore offer a best-case scenario to assess how informed consumers make choices in complex health insurance contracts.

To provide theoretical guidance regarding optimal HSA use, I first develop a life-cycle model that incorporates both HSA and 401(k) saving. The model delivers a neoclassical benchmark for the MPC. Without liquidity constraints, utility maximization entails minimizing the costs to finance the present value of lifetime

[1] More precisely, this ratio is technically a marginal propensity to withdraw. For the purposes of this paper, however, it represents the theory-relevant parameter of interest (see online Appendix B) and I refer to it as the MPC.

health care costs. The MPC should be zero while people are working, because money is fungible. Due to the more generous tax preferences for HSAs, allowing an HSA to grow until age 65 saves money in the long run: people are better off reducing their 401(k) contributions to finance current health care expenses, rather than using their HSA to pay for these expenses. This strategy results in a smaller 401(k) at age 65 than if they withdrew HSA funds immediately, but this loss in the 401(k) is less than the value of the increased HSA assets. The intuition is similar to why people with low mortgage rates are often better off *not* repaying their loan early if they can earn a higher return elsewhere.

I strongly reject the hypothesis that consumers in this setting use their HSA as self-insurance. I estimate an MPC of 0.85 from the HSA, with the lower bound of the 95 percent confidence interval of 0.60. Not only is this magnitude high relative to the neoclassical benchmark of zero, but it also exceeds the MPCs estimated from tax rebates (Sahm, Shapiro, and Slemrod 2010; Shapiro and Slemrod 2009; Parker et al. 2013), out of liquidity (Gross, Notowidigdo, and Wang 2020), regular transfers from the Alaska Permanent Fund (Kueng 2018), or SNAP benefits (Hastings and Shapiro 2018).

The high MPC from HSA assets blunts the deductible's incentive to reduce costs. Total health spending increases in response to additional HSA funds from the more generous employer match. The spending increases are observed across a range of health care services, which corroborates findings from other settings that consumers in HDHPs face difficulty in distinguishing between high- and low-value care (Brot-Goldberg et al. 2017). The largest changes occurred in specialty care and other outpatient care. One of the few services that did not change was preventive care, which is exempt from the deductible.

In fact, health spending did not decline after the firm replaced its low-deductible offerings with HDHPs and HSAs. Instead, a larger share of spending was financed tax-free as HSA contributions exceeded the difference in premiums. The increase in the average tax subsidy was large, exceeding $900 per household. A high MPC from the HSA may therefore fully counteract the cost-reducing incentives of the HDHP, contrary to the objectives of these contracts.

Comparisons with 401(k) saving provide further evidence that most people do not use HSAs to self-insure, and do not view HSA money as fungible with other tax-preferred saving. Employees whose 401(k) contributions exceed the employer 401(k) match should max out their HSA, since the HSA's tax incentives dominate those of the 401(k) past this level. Yet almost 90 percent fail to do so, and the magnitude of these optimization errors is sizable. On average, over $1,300 of employee 401(k) contributions are dominated annually, and employees would unambiguously be better off if they reallocated this saving to their HSA. The average size of these foregone tax benefits amounts to over $550 each year. There is also no evidence that HSAs crowd out 401(k) saving.

Collectively, this behavior is consistent with mental accounting (Thaler 1985, 1990, 1999; Shefrin and Thaler 1988; Prelec and Loewenstein 1998) as well as misperceptions about certain contract features. Mental accounting is a cognitive process whereby people categorize income and expenditure items into separate accounts, even though money is fungible. In supplementary analysis, I provide evidence against two other candidate explanations for the high MPC. HSA withdrawals do not spike at the end of the year, which would be expected if people believed the funds expired, like a Flexible Spending Account (FSA). I also test one version of liquidity constraints in explaining the high MPC: the bulk of the employer's contribution is deposited in a single month, but I do not find evidence of excess sensitivity between health spending and the timing of the employer's HSA contribution.

The paper makes several contributions in an area where the literature is thin. The first contribution is to develop a theory of optimal HSA saving in conjunction with other tax-preferred retirement accounts. While the popular press and personal finance sites have written much about the HSA's tax advantages, the academic literature has paid less attention to formally modeling HSA saving decisions in a life-cycle framework (Baicker, Dow, and Wolfson 2006; Cardon and Showalter 2007; Aaron, Healy, and Khitatrakun 2008; Bundorf 2016). By analyzing the optimal use of these accounts, this paper clarifies the links between health insurance and life-cycle saving that are central to the structure of HSAs. The model's key insight is that consumers without liquidity constraints should limit HSA withdrawals while working, which is not obvious but stems from the fungibility across different financial accounts.

The second contribution is to provide the first causal estimates of the MPC from HSA assets using a transparent regression discontinuity design. Evidence on this important economic parameter complements work that calculates descriptive statistics on HSA saving using administrative data from tax records and large HSA sponsors (Helmchen et al. 2015; Fronstin 2017, 2019).[2] Third, the paper builds upon the growing literature of optimization errors in health insurance choices (Abaluck and Gruber 2011, 2016; Ketcham et al. 2012; Kling et al. 2012; Heiss et al. 2013; Handel 2013; Handel and Kolstad 2015; Bhargava, Loewenstein, and Sydnor 2017; Ericson and Sydnor 2017; Chandra, Handel, and Schwartzstein 2019), and is the first to focus specifically on mistakes related to saving incentives and health insurance. The findings also relate more generally to the large literature on behavioral household finance (Beshears et al. 2019).

The paper proceeds as follows. Section I summarizes the tax preferences of HSAs and discusses optimal HSA saving in a life-cycle framework. Section II describes the empirical setting and data. Section III presents regression discontinuity evidence of the MPC and Section IV presents results for health spending. Analysis of the fungibility between 401(k)s and HSAs is presented in Section V. Section VI briefly concludes.

I. Theory: Optimal Saving and Withdrawals from HSAs

This section considers HSA saving and withdrawals in a life-cycle framework of consumption. The purpose is to provide a neoclassical benchmark against which the

[2] Lo Sasso, Helmchen, and Kaestner (2010) use variation across employers in Health Reimbursement Account (HRA) balances, which are not portable and have less generous tax preferences than HSAs. They estimate that a dollar increase in HRA balances raises health spending by the same amount.

estimated marginal propensity to consume (MPC) from HSA assets in Section III can be compared. I begin by explaining the key tax preferences of HSAs relative to taxable savings accounts and 401(k) retirement accounts and then summarize the optimal saving and withdrawal strategies from the HSA in a life-cycle model.

A. *Key Rules and Tax Preferences of HSAs*

HSAs have several features that distinguish them from other savings vehicles. First, contributions are tax-exempt, investments grow tax-deferred, and withdrawals for qualified medical expenses are tax-free.[3] Contributions can only be made when enrolled in an HDHP. Withdrawals without a qualified medical expense are subject to income tax and, if withdrawn before age 65, a penalty tax.[4] Second, employee HSA contributions made through payroll deductions are not subject to FICA (payroll) taxes, unlike employee 401(k) contributions.[5] Third, any past health care expenses while employees are enrolled in an HDHP are eligible for tax-free withdrawals, regardless of how long ago the expense occurred: withdrawals need not correspond to expenses from the same year and consumers can withdraw up to their full balance in any year. Loans are not permitted from HSAs. Finally, HSAs can be inherited by a spouse, who can continue to make tax-exempt withdrawals for qualified expenses.[6] These rules are common across settings. Other provisions such as ERISA protection, fees, and investment opportunities vary across contexts. Online Appendix A provides additional background about HSAs and summarizes relevant literature on descriptive patterns.

Due to the more generous tax preferences, a dollar saved in the HSA is worth more in the future than a dollar saved in other accounts. Consider the choice between saving in taxable accounts, a 401(k), or the HSA, and assume each earns the same rate of return r. After t years, a dollar would be worth the following in each of these accounts:

- Taxable account: $\left(1-\tau_0^y-\tau_0^p\right)\left(1+r\left(1-\tau^i\right)\right)^t$, because contributions are subject to income taxes τ_0^y and payroll taxes τ_0^p and investment returns are subject to taxes τ^i.
- 401(k): $\left(1-\tau_0^p\right)(1+r)^t\left(1-\tau_t^y\right)$, because contributions are subject to payroll taxes, investments grow tax-deferred, and income taxes are paid upon withdrawals.[7]
- HSA: $(1+r)^t$ if used to finance health care and $(1+r)^t\left(1-\tau_t^y\right)$ if withdrawn without a qualifying expense.

These tax advantages of the HSA highlight the rationale for using the account as a savings vehicle.

[3] Qualifying expenses—determined by the IRS—include most out-of-pocket costs incurred while enrolled in an HDHP and dental and vision care, as well as Medicare premiums and out-of-pocket costs, long-term care insurance premiums, and long-term care.

[4] 401(k) withdrawals prior to 59.5 also face a penalty tax unless one of several exemptions are met.

[5] Employer contributions to both accounts are not subject to FICA taxes.

[6] In the case of 401(k)s, a beneficiary can roll the assets into an IRA, but withdrawals still face taxes.

[7] This return is the same as for a traditional IRA. In a Roth account, the dollar would be worth $\left(1-\tau_0^y-\tau_0^p\right)(1+r)^t$, since income tax is paid when contributing rather than when withdrawing.

B. *Optimal HSA Saving over the Life Cycle*

To formally examine HSA saving and calculate the optimal MPC from the HSA, I extend a standard life-cycle model of consumption to incorporate HSAs. Online Appendix B presents the details of the model, including the specification, parametrization of uncertainty, and solution methods. I focus here on describing the setup and summarizing the key results and intuition.

The individual's problem is to choose saving and withdrawals from the HSA and other accounts to maximize the discounted expected utility of consumption over her lifetime. Preferences satisfy constant relative risk aversion (CRRA) and she faces two sources of uncertainty, which are treated as exogenous: survival risk and health care expenditure risk. She is enrolled in an HDHP with an HSA while working and receives Medicare coverage at age 65.[8] I make two key assumptions to obtain sharp predictions and establish a benchmark for the MPC. First, she is not liquidity-constrained and can borrow against future labor income at the risk-free rate. Online Appendix E investigates the importance of this assumption in my empirical setting. Second, lifetime out-of-pocket expenses eligible for HSA reimbursement exceed the maximum that can ever be accumulated in the HSA. This assumption is supported by estimates of lifetime health care spending compared to annual HSA contribution limits. Online Appendix B provides further discussion of these assumptions and other minor ones.

The optimal strategy is to build HSA assets while working—rather than use them for current health spending—and decumulate HSA assets in retirement to pay for lifetime health spending. For a person facing the average mortality and expenditure risks, the MPC from the HSA is zero while working.[9] To finance health care expenses before age 65, reducing either taxable saving or 401(k) contributions while preserving the HSA is cheaper in the long run. Following this strategy results in a smaller 401(k) than if HSA assets were withdrawn before age 65, but the income loss (after taxes) is more than compensated for by the HSA's growth.

This result—that the MPC should be zero while working—is general, and does not depend on risk aversion or discount rates. It follows from the assumption about liquidity: the optimal strategy for HSA contributions and withdrawals will minimize the lifetime costs of financing health care expenses, because doing so maximizes the present value of lifetime consumption. Without liquidity constraints, non-HSA assets can adjust to reach the desired consumption profile over the life cycle. Since money is fungible, the way to minimize the lifetime costs of financing health care is to allow the HSA to grow, given its more favorable tax preferences. All HSA assets can be withdrawn tax-free, while 401(k) withdrawals necessarily incur taxes.

[8] As discussed in online Appendix B, I assume people do not expect to receive Medicaid in retirement or while working. I therefore do not consider the complicated incentives between saving and means-tested programs (Hubbard, Skinner, and Zeldes 1995).

[9] After solving the model, the policy function is used to simulate HSA asset accumulation and decumulation and to calculate the MPC from the HSA across the life cycle. The MPC is calculated as the change in withdrawals for a small change in the state-variable HSA assets.

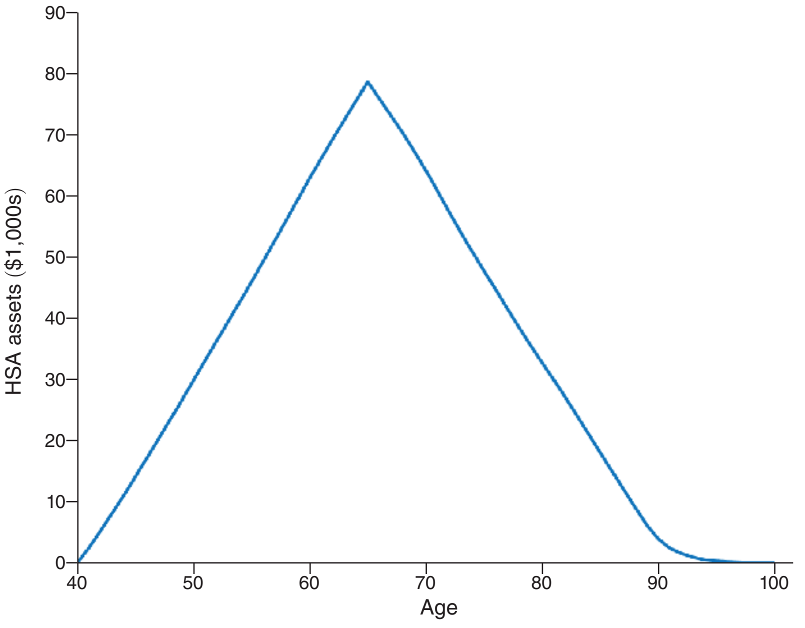

Figure 1. Optimal Path of HSA Assets, 40-year-old

Notes: This figure plots the optimal path of HSA assets for a person who first enrolls in an HDHP with an HSA at age 40 using the life-cycle model in Appendix B. The profile of HSA assets is calculated by averaging 20,000 simulations of the model. The individual faces uncertainty in survival and medical spending. This calibration assumes the employer deposits \$3,000 annually in the individual's HSA, and that the deductible and out-of-pocket maximum under the employer's HDHP is also \$3,000. The person works until age 65 and then receives Medicare coverage. The marginal tax rate is assumed to be 25 percent while working and in retirement, and the employee portion of payroll taxes is 7.65 percent. The interest rate on 401(k) and HSA assets is assumed to be 2 percent. HSA assets rise until age 65 and then are withdrawn increasingly in retirement.

The MPC should be zero even if the employer fully funds the employee's HSA each year while working. As an illustrative example, Figure 1 presents the optimal path of HSA assets for a 40-year-old starting with zero HSA assets. She is assumed to receive \$3,000 each year in employer contributions and her employer insurance has a maximum annual out-of-pocket limit of \$3,000. She faces the average survival and health expenditure risks of a person her age, and the optimal path of HSA assets is calculated by averaging the results of 20,000 simulations of the model. HSA assets rise steadily until age 65 and then are withdrawn increasingly. Withdrawing HSA contributions immediately to finance health care—similar to how Flexible Spending Accounts (FSAs) are used—takes advantage of the tax deductibility of income and payroll taxes, but sacrifices the deductibility of investment returns and compounding. Taking full advantage of the tax incentives requires time for HSA assets to grow. The optimal use of HSAs is therefore as self-insurance over the life cycle, rather than financing the current year's deductible. Incorporating moral hazard only strengthens this result, since more health care spending in the future must be financed through higher 401(k) saving.

II. Setting and Data

This section describes the variation in employer policies used to identify the MPC out of HSA assets and presents descriptive information about the employer and sample statistics. The firm that I study is one of the five largest health insurers in the United States by both market share and revenues, with employees throughout the country. In terms of representativeness, the average salary, age, and tenure among the firm's employees are roughly in line with US labor force averages. The administrative data combine detailed information on employee salary, benefit choices, demographics, and medical and pharmacy claims between 2006 and 2010.

A. *Variation in Employer Benefits*

The firm began offering employees the choice of HDHPs with HSAs instead of its traditional health insurance plans in 2005. Starting in 2008, the firm dropped its traditional plans and only offered HDHPs with HSAs. Employees chose between four deductibles ranging from $1,250 to $3,150, with double these amounts for family coverage. As in other settings, the employer contributed a flat amount to each plan premium and then required employees to pay the higher costs of additional coverage. Preventive care was free and the patient paid the full charge for all other care until the deductible had been met.[10]

The firm matched employee HSA contributions at different rates based on the employee's annual salary. Employees with annual salaries below $50,000 received a more generous HSA match than those earning more than this amount. This discontinuity in matching rates provides the exogenous source of variation I use to estimate the MPC from HSA assets and to measure the responsiveness of health spending to HSA funds. In particular, employees who face the higher match rate for that year will receive larger employer contributions than employees who face the lower match rate. This matching schedule therefore creates an arbitrary difference in HSA balances for employees just below and just above this salary cutoff.

The HSA matching rates were as follows. For employees earning below $50,000, the employer matched the first $100 of contributions at a rate of 6:1 for those with employee-only coverage, up to a maximum employer contribution of $600. The contribution limits were double for family coverage, so employees could receive $1,200 in employer contributions on the first $200 they saved. For employees earning between $50,000 and $100,000, the match rate was 4:1 up to maximum amounts of $400 for employee-only coverage and $800 for family coverage. These matching rates also varied over time. 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. There was no default employee contribution to the HSA. Employees were immediately vested for both their contributions and the employer's contributions.

[10] Prior to 2010, the out-of-pocket maximum in each HDHP equalled the deductible. 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 coverage.

There were no other employer policies that varied discontinuously at the matching threshold.

In this setting, HSA assets could be invested in a variety of mutual funds once balances reached $2,000. There were no initial setup fees or monthly fees for the HSA account or the investment account paid by the employee.[11] 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 to save for future consumption if investment opportunities are poor and transaction costs are high.[12]

The firm pursued an extensive communications campaign to inform employees about its HDHP offerings and HSA benefits. 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. 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 deductible. The employer provided an HSA debit card to employees.

As for other retirement benefits, the firm offered employees a defined-contribution 401(k) and matched employee contributions up to 6 percent of salary. Prior to 2008, the firm matched all employee contributions at 50 percent up to this threshold. Starting in 2008, the firm began matching the first percent of employee salary at 100 percent and then matched subsequent contributions at 50 percent, up to 6 percent of salary. Employee contributions were deducted from each period's paycheck. 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.

B. *Data Description*, *Sample Composition*, *and Descriptive Statistics*

The administrative data include detailed information on each employee's salary, job characteristics, demographics, medical and pharmacy claims, and choices about retirement saving and health insurance plans. 401(k) contributions and balances are measured annually. HSA variables—contributions by both the employee and employer, employee withdrawals, balances, and interest—are measured monthly. Job characteristics and geographic information is recorded once, at the end of the sample period.

I use detailed claims data to test for balance in chronic conditions and expected spending at the match discontinuity and to examine health spending responses to more generous HSA matching. The claims data include information on health expenditures for employees and any dependents covered under the employee's policy. Each claim provides detailed information on diagnoses (ICD-9 and CPT codes for medical claims), providers, and payment (e.g., patient paid, plan paid), and month

[11] Investing HSA assets required an active choice, unlike the 401(k). The default for HSA balances was cash unless the employee opened an investment account.

[12] Moreover, people can transfer their HSA balances once a year to a different plan administrator (e.g., one with lower or zero fees) if they choose. See online Appendix B for more discussion of this option.

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 expected spending risk for each employee and their dependents (see online Appendix C for details).

The main analysis sample is constructed by starting with all employees appearing in the employer's payroll records between 2006 and 2010 (roughly 26,000) and restricting to those who were (i) enrolled in one of the firm'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), as recorded by the auto-enrollment indicator equalling zero. Restrictions (i)–(iii) are 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 59.5. I exclude the 18 percent of remaining employees who default into saving 4 percent of salary in their 401(k) because some analyses compare HSA saving to 401(k) saving. The assumption in this comparison is that the employee's 401(k) saving decision captures their intertemporal preferences and retirement saving objectives, which is more difficult to justify for people who auto-enroll in the 401(k). In robustness checks, I include employees who default into the 401(k) to examine sensitivity to this restriction. After these restrictions, the sample includes 15,908 employees.

Table 1 presents summary statistics of the analysis sample—overall and by type of insurance coverage—for years 2008 to 2010, when all employees were enrolled in HDHPs.[13] The average employee age is 40 years, the average tenure with the firm is 7.3 years, and the average salary is $61,933. Sixty-seven percent of the sample is female. Annual HSA saving, including employer contributions, averages $1,403 for employee-only coverage and $2,756 for family coverage. Over 88 percent of annual HSA contributions are withdrawn the same year, on average. Almost 60 percent of the sum of existing balances and contributions are withdrawn each year.[14]

III. Regression Discontinuity Estimates of the MPC

This section tests whether the MPC from the HSA is zero, as predicted by the life-cycle model, using a regression discontinuity design. To do so, I compare the change in HSA withdrawals at the match discontinuity to the change in HSA contributions. Figure 2 previews the regression discontinuity results by plotting the means of HSA contributions and withdrawals by salary level (within $1,000 bins) for employees earning less than $80,000 annually. There is a visible drop in employer contributions at the match discontinuity, driven by the matching rules. Total HSA contributions increase with salary, but there is a jump downward in the regression

[13] The main RD analysis uses data from 2008–2010 to weight employees who adopted HSAs early and those not adopting HSAs until 2008 equally.

[14] Online Appendix Table D.1 provides statistics on other measures of HSA and 401(k) saving and other variables. Online Appendix Table D.2 presents summary statistics stratified by whether the employee adopted the HSA prior to the forced switch in 2008.

Table 1—Summary Statistics of Sample

	All employees		Employee-only coverage		Family coverage	
	Mean	SD	Mean	SD	Mean	SD
HSA employee contribution ($)	1,558	1,424	954	813	1,975	1,595
HSA employer contribution ($)	646	372	449	179	781	408
HSA balance ($)	1,061	2,067	837	1,522	1,215	2,359
HSA withdrawal ($)	1,943	1,507	1,155	914	2,487	1,593
401(k) employee contribution ($)	4,104	4,319	3,546	3,835	4,488	4,584
401(k) employer contribution ($)	1,872	1,626	1,580	1,263	2,073	1,807
401(k) balance ($)	45,717	94,715	32,230	67,003	54,992	108,823
Total health spending ($)	8,924	22,319	4,850	12,488	11,731	26,731
Out-of-pocket spending ($)	2,198	1,940	1,126	1,049	2,936	2,064
Salary ($)	61,933	40,765	53,369	29,556	67,833	46,038
Tenure with employer (years)	7.31	6.05	6.53	5.64	7.85	6.26
Age (years)	40.01	9.60	38.95	10.61	40.74	8.76
Female (share)	0.67	0.47	0.69	0.46	0.66	0.47

Notes: This table presents means and standard deviations of the analysis sample between 2008 and 2010 by type of coverage. $N = 34{,}628$ employee years for full sample (columns 1–2), $N = 14{,}124$ employee years for employee-only coverage (columns 3–4), and $N = 20{,}504$ employee years for family coverage (columns 5–6). Family coverage also includes coverage for employee plus spouse and employee plus children. The HSA balance denotes the balance at the beginning of the year, prior to that year's contributions and withdrawals.

function at $50,000, indicating that employees do not fully reduce their own HSA contributions in response to the match. HSA withdrawals also rise with salary, and there is a clear drop in withdrawals at the match discontinuity.

The baseline RD specification is a local linear model, with separate regressions for annual contributions and withdrawals:

$$(1) \qquad contributions_{it} = \beta_0 + \beta_1 M_{it} + \beta_2 salary_{it} + \beta_3 M_{it} \times salary_{it} + \varepsilon_{it},$$

$$(2) \qquad withdrawals_{it} = \pi_0 + \pi_1 M_{it} + \pi_2 salary_{it} + \pi_3 M_{it} \times salary_{it} + e_{it},$$

where $salary_{it}$ is the annual salary of employee i in year t that has been recentered to $50,000 and M_{it} is an indicator for employees with salaries above $50,000. The MPC equals the ratio π_1/β_1, which measures the change in withdrawals relative to the change in contributions at the match discontinuity. This is analytically equivalent to an instrumental variables (IV) regression using the match discontinuity to instrument for total contributions: equation (1) is the first stage and equation (2) is the reduced form. I calculate the MSE-optimal bandwidth based on the methods in Calonico, Cattaneo, and Titiunik (2014) and use the bandwidth from the reduced-form regression in estimating π_1/β_1 via IV. I subsequently show that the results are not sensitive to bandwidth choice. The main specification excludes covariates, uses a uniform kernel, and clusters standard errors by employee.

Before presenting the RD results, I first test that controls are balanced and the density of salary is smooth at the match discontinuity. Table 2 presents balance tests using a variety of covariates as the dependent variable in equation (1). Nearly all covariates are smooth at the cutoff. One exception is a statistically significant jump in age of 1.8 years, which is equal to 4.5 percent of the control mean. The p-values from omnibus balance tests at the bottom of Table 2 reveal that all

Figure 2. HSA Contributions by Salary Level

Notes: This figure plots means of employer HSA contributions (triangles), total HSA contributions (circles), and HSA withdrawals (diamonds) within $500 salary bins. For each variable, linear regressions are separately fit to data on both sides of the match discontinuity at $50,000. Data include both employee-only and family coverage. There is an average drop of about $400 in employer contributions at the discontinuity. Total HSA contributions are increasing with salary. There are perceptible jumps in the regression functions for both contributions and withdrawals at the match discontinuity.

covariates excluding age are jointly balanced. I demonstrate robustness of the MPC estimates to including controls, and the main RD regressions for health spending in Section IV include controls to account for the influence of any small imbalance in age. Importantly, the absence of jumps in expected spending or chronic conditions suggests there are not substantial differences in health status at the discontinuity. Online Appendix Figure D.1 presents the McCrary (2008) test that the density of the running variable is smooth at the cutoff. There is no evidence of manipulation of salary on either side of the discontinuity, confirming that salary is held constant when the employer's HSA contribution changes.[15]

Figure 3 graphically displays the RD estimate from estimating equations (1) and (2) by IV.[16] The estimated MPC equals 0.855 and is not statistically distinguishable from 1. The lower bound of the 95 percent confidence interval is 0.60. Online Appendix Table D.4 presents the results from seven other specifications to examine robustness to different modeling assumptions: including or excluding covariates, using a triangular or uniform kernel, and a linear or quadratic polynomial in salary. The estimated MPCs from these alternative models are all above 0.77. To explore sensitivity to the choice of bandwidth, online Appendix Figure D.2 presents estimates that vary the bandwidth from $500 to $20,000. The estimated MPC is again above 0.8 in nearly all cases.

[15] Employees just below the threshold do have slightly higher compensation than those above due to the employer's HSA contribution, but this difference is far too small to explain the high MPC from the HSA.

[16] Online Appendix Table D.3 reports the estimates from the first stage and reduced-form regressions corresponding to both Figure 2 and Figure 3.

Table 2—Covariate Balance

Dependent variable	Estimate	SE	p-value	Control mean	Percent difference
Expected spending	−367	(421)	0.383	8,438	−4.3
Diabetes Dx	−0.015	(0.013)	0.261	0.091	−16.1
Hypertension Dx	−0.008	(0.021)	0.695	0.254	−3.3
Atrial Fibrillation Dx	−0.020	(0.020)	0.335	0.239	−8.3
COPD Dx	−0.008	(0.011)	0.491	0.076	−10.1
Stroke Dx	−0.002	(0.004)	0.597	0.010	−21.1
Ischemic heart disease Dx	−0.008	(0.008)	0.272	0.031	−27.4
Cancer Dx	0.007	(0.006)	0.298	0.014	48.6
Asthma Dx	−0.018	(0.015)	0.236	0.139	−12.9
Liver Dx	−0.009	(0.007)	0.177	0.033	−28.8
Hyperlipidemia Dx	−0.039	(0.020)	0.055	0.259	−15.0
Osteoporosis Dx	−0.005	(0.004)	0.236	0.011	−44.8
Rheumatoid arthritis Dx	−0.005	(0.014)	0.712	0.103	−5.0
Age in years	−1.807	(0.471)	0.000	40.488	−4.5
Tenure in years	0.358	(0.341)	0.294	7.982	4.5
Female	0.000	(0.023)	0.995	0.691	0.0
Married	0.015	(0.026)	0.557	0.516	2.9
White	0.060	(0.023)	0.008	0.684	8.8
Years in HSA	0.048	(0.052)	0.357	2.913	1.7
Employee-only coverage	0.026	(0.025)	0.308	0.439	5.9
Number of dependents	−0.040	(0.070)	0.566	1.233	−3.2
Early HSA adopter	−0.005	(0.025)	0.840	0.638	−0.8
Omnibus test, all controls			0.002		
Omnibus test, excluding age			0.252		

Notes: This table presents RD estimation results of covariates using the $50,000 salary cutoff. Each row presents the results from a different RD model using local linear regression, uniform kernel, and the MSE-optimal bandwidth. Diagnoses of chronic conditions from the previous year are coded as indicators and denoted by the series of "Dx" variables. For most covariates, the point estimates are not statistically significant from zero. The third column lists the p-value from the test in which the discontinuity equals zero. The fourth column lists the control mean, calculated as the predicted value of the dependent variable immediately to the left of the discontinuity. The fifth column expresses the point estimate as a percentage of the control mean. Standard errors clustered by employee in parentheses.

In terms of magnitudes, the estimated MPC of 0.855 is very large. In the neoclassical benchmark discussed earlier, the optimal MPC should be zero for these employees. I soundly reject the null of zero, which would be consistent with self-insurance. By contrast, I cannot reject the null that the MPC equals 1. The estimated MPC is also higher than other contexts measuring the change in annual spending in response to tax rebates (Sahm, Shapiro, and Slemrod 2010; Shapiro and Slemrod 2009; Parker et al. 2013), out of liquidity (Gross, Notowidigdo, and Wang 2020), dividends from the Alaska Permanent Fund (Kueng 2018), or SNAP benefits (Hastings and Shapiro 2018).[17]

To explore heterogeneity in the MPC, online Appendix Table D.5 presents the RD estimates for different subsamples. For most employee characteristics examined, the MPC exceeds 0.8. The MPC is higher for workers who are older, who have longer tenures, who have family coverage, or who have lower beginning-year HSA balances.

[17] As another parameter of interest, online Appendix Figure D.3 shows that the estimated ratio of the change in out-of-pocket payments to HSA contributions at the discontinuity is equal to 0.399.

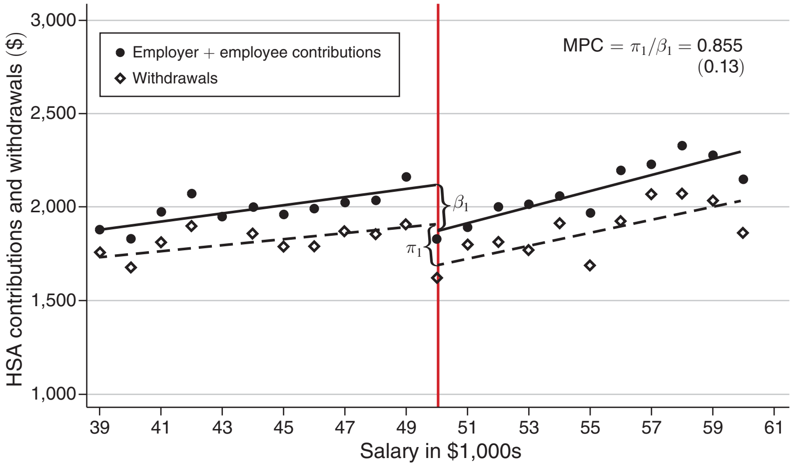

FIGURE 3. RD ESTIMATE OF THE MPC

Notes: This figure graphically displays the RD results (local linear models, uniform kernels, and no controls) used to calculate the MPC from HSA assets. The figure plots means of HSA contributions and withdrawals within $1,000 salary bins and fits separate linear regressions to the data below and above the $50,000 match discontinuity within the MSE-optimal bandwidth of $10,752 for withdrawals. Data includes both employee-only and family coverage. The MPC is estimated using IV as the ratio of the jump in withdrawals (π_1) to the jump in contributions (β_1) at the match discontinuity, following equations (1) and (2) in the text. The estimated MPC is 0.855 with a standard error (clustered by employee) equal to 0.130.

IV. HSAs and Health Spending

Having documented a high MPC from HSA assets, I now estimate the effect of HSAs on health spending. The estimated MPC close to 1 strongly rejects the neoclassical model, but does not directly imply results for total health spending. For example, consumers may withdraw HSA assets to finance more of their deductible tax-free, but not change how much care they consume. I use two sources of variation to analyze the influence of HDHPs with HSAs on total health spending. First, I continue to use the discontinuity in matching rates to estimate how more generous HSA funding influences health spending within the HDHP. This comparison provides insight into the alternative of offering an HDHP without an HSA. Next, I compare the contracts to a traditional insurance plan by using employees who switched to the HDHP/HSA early as a control group for those who waited until the firm discontinued its traditional plans. This comparison provides insight into whether HSAs, when not used as self-insurance, change health spending compared to a low-deductible plan without a savings vehicle.

A. *RD Results for Total Health Spending*

There is evidence more generous HSA funding increases total health spending in HDHPs. Using the discontinuity in matching rates, I estimate the effect of a dollar increase on spending via IV by replacing withdrawals in equation (2) with total health spending. Increasing HSA funds by $1 is estimated to raise total spending by

Table 3—RD Results: Spending by Category and Service Type

	Estimate	SE
Panel A. Total spending		
Total spending	−1,942.2	(887.0)
Total spending, 95 percent Winsorized	−685.7	(393.7)
ΔTotal spending/ΔHSA	6.53	(3.11)
ΔTotal spending, 95 percent Winsorized/ΔHSA	2.50	(1.42)
Panel B. Spending by category		
Primary care	−56.8	(38.4)
Specialty care	−332.2	(158.5)
Other outpatient	−664.5	(344.7)
Inpatient	−888.3	(630.8)
Prescription drugs	209.7	(221.1)
Emergency room	−16.0	(74.2)
Other	−41.9	(33.1)
Panel C. Spending by service type		
Radiology	−115.7	(121.9)
Mental health	−19.8	(51.1)
Behavioral health	−38.3	(41.0)
Preventive care	13.1	(26.1)

Notes: This table presents results of RD models using local linear regression with a uniform kernel for the matching discontinuity at $50,000 corresponding to different categories or services of health care. Each regression includes controls for number of dependents, indicators for coverage type, state of residence, deciles in age, female, married, White, and lagged diagnoses (from the previous year) of chronic conditions. The first column reports the estimated jump in the outcome at the cutoff and the second column reports the standard error (clustered by employee). MSE-optimal bandwidth calculated separately for each outcome. Categories in panel B are mutually exclusive and combine both office visits and outpatient hospital care into the outpatient category.

$6.53 (Table 3). The reduced form is shown graphically in Figure 4 and estimates a $1,942 change at the match discontinuity. Both estimates are statistically significant.[18] The magnitude of the point estimate is large, but the confidence interval is quite wide. The lower bound of the 95 percent confidence interval is a $0.43 increase in total spending for a $1 increase in HSA funds. As a robustness check, there are also statistically significant increases in Winsorized spending in which values greater than the ninety-fifth percentile of spending are top-coded at this percentile, as shown in online Appendix Figure D.4 and Table 3.

One explanation for the large point estimate is that HSA funding increases the likelihood of hitting the deductible, and most spending occurs beyond the deductible. Specifically, the additional HSA funds increase the probability of hitting the deductible by 2.2 percentage points (roughly a 6 percent increase) and raises the likelihood that spending exceeds amounts ranging from $12,000 to $20,000 (online Appendix Table D.6). My interpretation of this collective set of results is that HSAs increase total health spending, but any multiplier effect cannot be precisely estimated.[19]

[18] Regressions control for deciles in age, the number of dependents, indicators for coverage type, state of residence, female, married, White, and diagnoses from the previous year of chronic conditions.

[19] The spending response is also higher than the responses to an increase in other retirement wealth within the sample. The employer made annual contributions equal to 4 percent of employee salary to a separate firm retirement account after employees reached two years of tenure. Online Appendix Table D.11 shows there is no evidence health spending rises in response to this increase in other retirement assets.

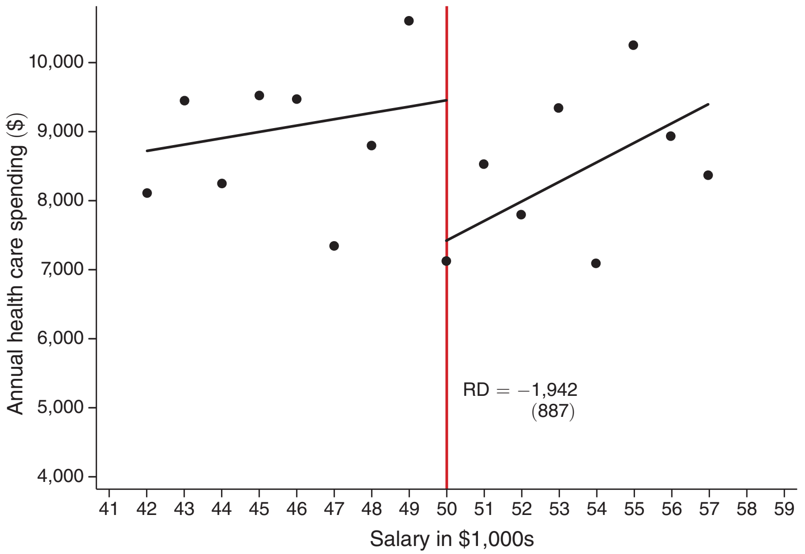

Figure 4. RD Estimate of Total Health Spending

Notes: This figure presents the results of RD models (local linear regression) of total health spending. There is an estimated $1,942 decrease in total health spending at the threshold, with a standard error (clustered by employee) equal to $887. Regressions include controls for number of dependents, indicators for coverage type, state of residence, deciles in age, female, married, White, and lagged diagnoses (from the previous year) of chronic conditions. Points plot the mean of spending within $1,000 salary bins within the MSE-optimal bandwidth after residualizing this set of controls.

The welfare implications of these spending results depend on what health care is reduced in response to lower HSA contributions. The demand curve is no longer sufficient to quantify moral hazard if consumers misjudge the benefits from care or make other optimization errors in their health care consumption decisions—which Baicker, Mullainathan, and Schwartzstein (2015) define as “behavioral hazard.” If the high deductible leads people to forego care worth more than its cost, then the reduced spending from lower HSA contributions may worsen health and decrease welfare relative to an HDHP alone. If, instead, the services displaced are low-value, then reducing HSA funds may improve welfare. Recent studies have found that higher cost-sharing reduces both high-value and low-value care by similar amounts (Brot-Goldberg et al. 2017; Lavetti, DeLeire, and Ziebarth 2019).

To decompose the nature of the spending reductions, panels B and C of Table 3 report the reduced-form estimates for different types of spending. Negative entries indicate reductions without the HSA match. There are estimated reductions for most categories of spending (panel B), with large and statistically significant declines in outpatient specialty care and other outpatient care. There is an imprecisely estimated increase in prescription drugs. Panel C breaks spending into several types of services that begin to point to high- and low-value care. There is no evidence of declines in preventive care, which is not subject to the deductible. This result is to be expected if consumers understand that preventive care is covered free of charge.

Table 4 examines changes in the quantities of services and specific types of low-value care or high-value care to further assess the welfare implications of the

Table 4—RD Results: Quantities of Health Care Services

	Estimate	SE	Control mean
Panel A. Visits, tests, and prescriptions			
PCP visits, all	−0.367	(0.337)	6.720
PCP visits, non-preventive care	−0.385	(0.300)	5.346
PCP visits, preventive care	0.064	(0.176)	3.392
Mammograms	−0.038	(0.034)	0.412
Specialist visits	−1.327	(0.614)	9.310
Mental health visit	−0.534	(0.441)	3.879
Behavioral health visits	−0.334	(0.231)	1.274
ER visits	−0.003	(0.022)	0.276
CT scans	−0.072	(0.051)	0.462
MRIs	−0.015	(0.038)	0.290
Prescription fills	−1.737	(1.524)	27.318
Panel B. Low-value care			
Head imaging for uncomplicated headache	−0.135	(0.111)	0.637
Back imaging for nonspecific low back pain	−0.407	(0.274)	1.488
Antibiotics for acute respiratory infection	−0.006	(0.028)	0.239
Concurrent Rx for opioids and benzodiazepines	−0.029	(0.013)	0.077
Long-term use of opioids, non-cancer patients	−0.055	(0.022)	0.247
Standardized index of low-value care	−0.077	(0.025)	
Panel C. High-value care			
Preventive visits	0.064	(0.176)	3.392
Physical therapy visits	−0.212	(0.248)	1.628
Diabetes drugs	0.073	(0.171)	0.873
Antidepressants	−0.059	(0.251)	2.097
Hypertension drugs	0.076	(0.175)	1.555
Lipid-lowering drugs	−0.042	(0.177)	1.224
Standardized index of high-value care	0.012	(0.020)	

Note: This table presents results of RD models using local linear regression for the matching discontinuity at $50,000 corresponding to changes in the amount of services consumed. Each regression includes controls for number of dependents, indicators for coverage type, state of residence, deciles in age, female, married, White, and lagged diagnoses (from the previous year) of chronic conditions. The control mean is calculated as the predicted value of the dependent variable immediately to the left of the discontinuity. Classifications for low-value care and high-value care use procedure and diagnosis codes following the definitions from Schwartz et al. (2014) and the National Quality Forum (2019). The standardized indexes summarize the effects for high- and low-value care, accounting for covariance between the estimates, using the seemingly unrelated regression approaches from Kling, Liebman, and Katz (2007); Clingsmith, Khwaja, and Kremer (2009). The standardized indexes are interpreted in terms of standard deviations of health care consumption within each domain (high- or low-value care). MSE-optimal bandwidth calculated separately for each outcome.

spending reductions. In panel A, the point estimates on all utilization measures are negative, with the exception of preventive primary care visits. Specialist visits are the single measure that is statistically significant. This finding suggests one potential mechanism for the large response in total spending: after a patient initiates a visit, subsequent treatment decisions are often driven by the provider.

Finally, panels B and C of Table 4 present several measures of low-value care and high-value care using detailed procedure and diagnosis codes, following the definitions from Schwartz et al. (2014) and the National Quality Forum (2019). There are estimated reductions in low-value care, with the largest results for two measures of opioid consumption: long-term use of opioids for non-cancer patients, and concurrent prescriptions for opioids and benzodiazepines. To provide a summary metric for the consumption of low-value care, I construct a standardized index that weights each measure equally and accounts for covariance between the estimates, using the seemingly unrelated regression approaches from Kling, Liebman, and Katz (2007) and Clingsmith, Khwaja, and Kremer (2009). The estimate for the standardized index is interpreted in terms of standard deviations of health care consumption within the low-value care domain. There is a reduction in the index, driven by the opioid measures. The results for high-value care are mixed in sign and generally noisier.[20]

Collectively, this set of analyses point to reductions in total spending in HDHPs due to lower HSA funds. There are particularly large reductions in the quantity and spending for specialty care and outpatient care, but no evidence of reductions in preventive care.

B. *HDHP/HSA versus Traditional Insurance*

To examine how spending changed after the introduction of HDHPs and HSAs, I compare employees who switched to HDHPs early to employees who did not. As a visual preview, Figure 5 plots the time series of monthly per-person spending between 2006 and 2010 separately for employees who selected the HDHP early (dashed line) and employees who did not enroll in the HDHP until traditional plans were dropped (solid line). The vertical line represents the final month of the 2007 plan year (corresponding to June 2008), before all employees moved to HDHPs.[21] Spending is expressed based on prices and ages in January 2009. In addition to the adjusted mean, Figure 5 presents the median, Winsorized mean, and enrollment counts.[22]

There is no visual evidence of a spending decrease after HDHPs replaced traditional plans in 2008. Among employees enrolled in HDHPs, spending rises in the final month of each plan year, before falling in the subsequent month. This pattern is observed for those who chose the HDHP early as well as those who only enrolled after 2008. There does not appear to be a large break in spending in 2008 for either group. The absence of a trend for the employees who chose the HDHP early reveals there are no relevant secular trends coincident with the firm's 2008 policy change.

To formally test whether spending changed, I estimate difference-in-difference regressions that use the early switchers as a control group for late switchers. Since early switchers are not an exogenous control group and they differ on observables, I perform a reweighting procedure to match early and late switchers based on lagged

[20] Online Appendix Table D.7 replicates Table 3 and online Appendix Table D.8 replicates Table 4 without controls and show qualitatively similar patterns.

[21] 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.

[22] Since the sample is restricted to those with coverage over the entire plan year, new employees with a partial year of coverage are not included until the next year. The increase in enrollment among "late switchers" in 2009, rather than 2008, reflects new employees who joined the firm partway through 2008 when the HDHP/HSA was the only option.

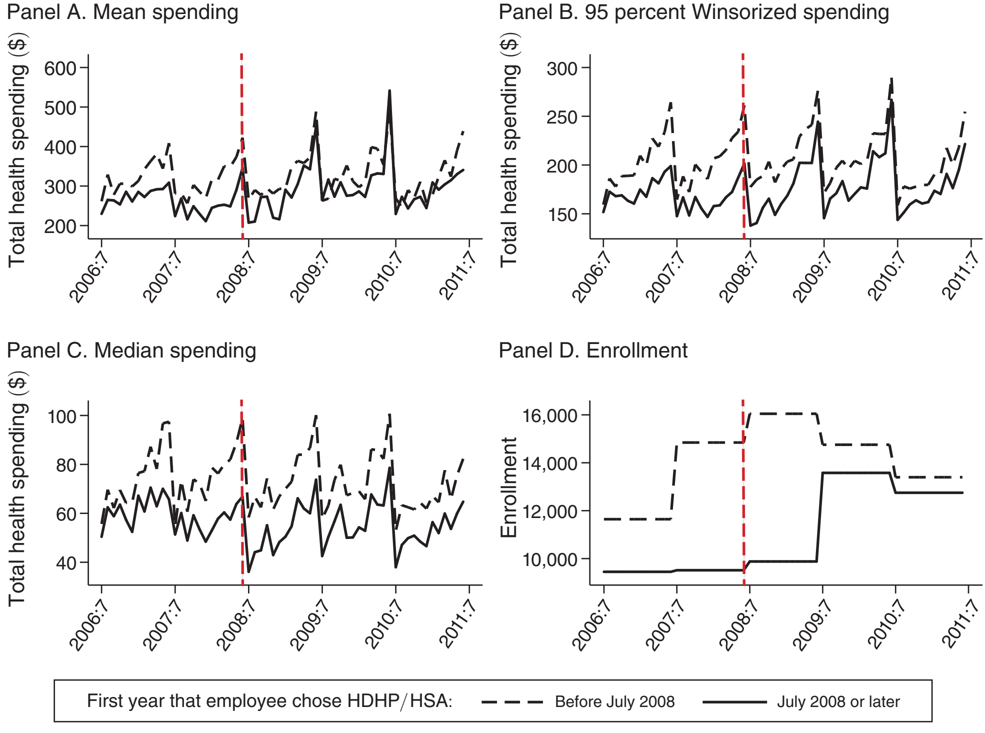

Figure 5. Time Series Patterns of Health Spending before and after HDHP/HSA

Notes: This figure plots monthly health spending between July 2006 and June 2011 separately for employees who chose the HDHP/HSA prior to July 2008 (dashed line), when traditional plans were still available, and those who chose after July 2008 (solid line), when only HDHPs/HSAs were offered. The vertical line denotes June 2008, the final month before the forced switch to HDHPs/HSAs. The sample is restricted to employees with full-year coverage, so new employees joining partway through the year are included the following year. Spending is adjusted for changes in the medical care component of the consumer price index and for aging, as described in the text. The panels plot mean spending (panel A), Winsorized spending top-coded at the ninety-fifth percentile (panel B), median spending (panel C), and enrollment (panel D). Spending spikes in the final month of each year for employees enrolled in the HDHP/HSA but does not exhibit a break in July 2008 once all employees enroll in the HDHP/HSA.

values of health spending. Specifically, I construct vingtiles of lagged health spending for early switchers in 2007 and reweight the late switchers based on their 2007 spending to match the distribution of early switchers. Including covariates, there is a 3.5 percent increase in the first year and a 4.7 percent increase in the first 2 years, both of which are statistically indistinguishable from 0 (online Appendix Table D.9). The lower bounds of the 95 percent confidence intervals rule out decreases of −5.6 percent and −3.5 percent, respectively.

The absence of a spending decline after the HDHP's introduction contrasts with research from other settings documenting reductions between 5 and 15 percent (Buntin et al. 2011; Bundorf 2016; Haviland et al.2016; Brot-Goldberg et al. 2017). Part of the difference may be explained by the extent that coverage generosity decreased across various settings. In my setting, the actuarial value—defined as the fraction of total spending paid by insurance—declined from 77 percent to 69 percent. The decline in Brot-Goldberg et al. (2017) was much larger: from 100 percent (free care) to 78 percent. Differences in health insurance literacy may also partly explain the absence of large declines in this setting. Qualitative research shows most consumers are unaware that preventive care is free under HDHPs (Reed et al. 2012), and Brot-Goldberg et al. (2017) document reductions in preventive care. By contrast, employees in this setting appear aware that preventive care is exempt from the deductible by not reducing their consumption of it while they cut back on other services. Although other studies have not estimated the MPC from the HSA, Haviland et al. (2016) find that spending decreases are not as large in firms that offer more generous employer HSA contributions, consistent with my findings.[23]

A high MPC from the HSA may also lead to the unintended consequence (from society's perspective) of increasing tax expenditures. Prior to 2008, annual premiums for the HDHP were about $140 less, on average, than premiums for traditional insurance. This slight premium reduction for the HDHP was swamped by an average HSA contribution of over $2,300. Based on imputed marginal tax rates for these employees, the net change in tax subsidies for health care increased by over $900 per employee, on average.[24] In this context, replacing traditional plans with HDHPs did not lead to spending reductions and instead increased the share of health expenditure that was financed tax-free.[25]

V. HSAs, 401(k)s, and Fungibility

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. This section tests whether consumers treat HSAs as fungible with other tax-preferred saving through two comparisons. First, I examine HSA saving among employees who make 401(k) contributions in excess of the employer match. These employees should max out their HSA since the HSA's tax incentives dominate those of the 401(k) past this level. Second, I use the same discontinuity in HSA matching rates from earlier to test for crowd-out in 401(k) saving.[26]

A. *Test of Fungibility between HSA and 401(k) Saving*

Any employee whose 401(k) contribution exceeds the employer match and who is not maxing out the HSA would unambiguously be 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). As a result, the 401(k) is dominated once employer matching is exhausted. Over half of the sample contributes beyond the employer 401(k) match in at least one year, and

[23] The average HSA/HRA contribution of the firms in Haviland et al. (2016) is lower than in this setting, while the firm's HSA contribution in Brot-Goldberg et al. (2017) is greater than here.

[24] The marginal tax rate is estimated using NBER TAXSIM and adds the federal, state, and FICA rates (including both employer and employee shares).

[25] This finding relates to evidence from Medigap that supplemental insurance dampens consumer incentives to reduce spending and creates a fiscal externality (Cabral and Mahoney 2019).

[26] There is a long debate on retirement saving crowd-out (Engen, Gale, and Scholz 1996; Poterba, Venti, and Wise 1996; Benjamin 2003; Gelber 2011; Chetty et al. 2014).

these employees have a higher average salary than those who are at or below the 401(k) match.

Yet only 12 percent of employees beyond the 401(k) match max out their HSA. The other 88 percent do not treat money as fungible between accounts. The size of these mistakes is large. On average, employees contribute $1,300 in unmatched 401(k) contributions that should instead be contributed to the HSA. The average size of these foregone tax benefits amounts to over $550 each year.

B. *Analysis of 401(k) Crowd-Out*

There is little evidence that HSAs crowd out 401(k) saving. Figure 6 plots annual 401(k) contributions (employee plus employer) against salary for employees earning less than $80,000 annually. There is no clear jump in 401(k) contributions at the HSA match discontinuity. To formally examine the extent of crowd-out, I estimate the ratio of the jump in 401(k) contributions to the jump in HSA contributions using the same IV approach as before. On average, there is an estimated 0.30 cent decrease in 401(k) contributions for a 1 dollar increase in HSA contributions at the match discontinuity (online Appendix Table D.10). The 95 percent confidence interval includes 0 (no crowd-out) and excludes −1 (full crowd-out) at the 10 percent level.

There is also no statistically significant change in the total amount of tax-preferred saving. Panel B of online Appendix Table D.10 shows RD results for the sum of HSA and 401(k) contributions, less HSA withdrawals. There is a statistically insignificant $40 dollar increase in total tax-preferred saving (last row), which is less than 1.5 percent of the control mean. Relative to the average employee salary, this estimate represents an increase in the saving rate of less than 0.1 percentage points. HSAs do not raise total saving rates, on average.

C. *Interpretation and Potential Mechanisms*

Collectively, these results provide strong evidence that consumers violate the fungibility between HSAs and 401(k)s. One mechanism consistent with this set of patterns is mental accounting. Mental accounting 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. 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 2017), children's clothing (Kooreman 2000), and food stamps (Hastings and Shapiro 2018). Mental accounting offers another explanation for borrowing decisions that violate the no-arbitrage condition, such as taking payday loans when lower interest credit is available (Agarwal, Skiba, and Tobacman 2009) or simultaneously holding both high-interest credit card debt and low-yield assets (Gross and Souleles 2002). These empirical anomalies may ultimately stem from psychological factors like salience

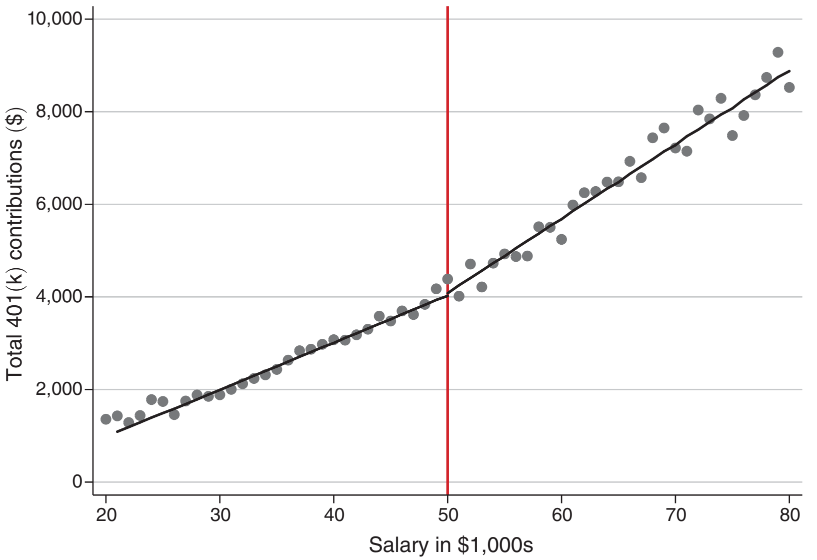

Figure 6. Total 401(k) Contributions by Salary Level

Notes: This figure plots means of 401(k) contributions (employer + employee) within $1,000 salary bins and fits separate linear regressions to the data below and above the $50,000 HSA match discontinuity. Data includes both employee-only 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 suggests there is no crowd-out of 401(k) saving from the exogenous increase in HSA funds at a $50,000 salary level.

(Bordalo, Gennaioli, and Shleifer 2013), rational inattention (Kőszegi and Matějka 2020), or gaps in financial literacy (Hastings, Madrian, and Skimmyhorn 2013; Lusardi and Mitchell 2014).

Mistaken beliefs about the rules of HSAs are another possible explanation. One might suspect that many consumers believe that HSA money does not roll over each year and is lost if not spent, like an FSA. In that case, one would expect a spike in withdrawals in the last months of the year. Online Appendix E shows this pattern is not observed, however. In fact, withdrawals are lowest in the final months of each year. The lower amount of withdrawals in the last quarter also holds when controlling for the level of existing HSA balances. The estimates are precise and provide strong evidence against the hypothesis that people believe that HSA funds expire.

It is possible that consumers are unaware of some of the HSA's more obscure rules, though, such as the ability to withdraw HSA money to finance past health care expenses or to finance consumption besides health care at age 65 without penalty. Such mistaken beliefs likely reflect mental gaps rather than standard search frictions, in the terminology of Handel and Schwartzstein (2018). Information about the rules of HSAs is not particularly costly to acquire given the many resources freely available online.

Finally, liquidity constraints represent a potential mechanism for the high MPC. If people are unable to borrow at the risk-free rate and instead only have access to credit at much higher rates (e.g., credit card debt), then immediately withdrawing HSA funds may be optimal. Survey data from the Federal Reserve and detailed interviews have documented many middle-income households struggle to finance emergency expenses of $400, well below the minimum deductible in an HDHP (Board of Governors 2014, Morduch and Schneider 2017). Ericson and Sydnor (2018) highlight the importance of borrowing constraints in rationalizing choices for lower deductible insurance plans that would otherwise appear to be dominated by a higher deductible plan.

To explore this mechanism, I exploit the concentrated timing of the employer's HSA contributions to test whether health spending jumps in the month that the employer's contribution is deposited. Under one version of liquidity constraints, employees may delay care if they are unable to finance it without the employer's HSA contribution. Perhaps surprisingly, there is no evidence of excess sensitivity of health spending to the employer's HSA contributions within the year. Online Appendix E provides additional details of this analysis.

VI. Discussion

This paper studies whether employees at a large firm use their HSA as a savings vehicle and how HSAs affect total health spending. In theory, these contracts offer an innovative design to balance incentives against risk protection. Linking a personal, tax-preferred savings account with insurance coverage encourages consumers to trade off current health care against future consumption. HSAs have more complex features than other insurance products, though. For the high deductible to increase sensitivity to costs as intended, consumers should view HSA contributions as savings. I estimate an MPC from HSA funds of 0.85 using a discontinuity in employer matching rates and strongly reject the neoclassical benchmark of zero. Employees in this setting do not use their HSA to self-insure over the life cycle.

The high MPC from the HSA counteracts the incentives of the high deductible to reduce spending, contrary to the contract's objectives. The responsiveness of total spending to an exogenous increase in HSA funds is large in magnitude. In fact, moving to HDHPs and HSAs did not reduce spending in this setting. Instead, the increase in HSA spending had the unintended consequence of increasing the share of health care financed through tax subsidies.

It is important to acknowledge several limitations of the study. First, the estimates are local to $50,000, given the discontinuity design. More importantly, the analysis is limited to a single firm during the initial years after HSA adoption. Learning over time might lead people to use their HSA as a savings vehicle. However, even those employees who enrolled in the HDHP prior to the forced switch and had multiple years of experience with HSAs did not use the accounts as self-insurance and made errors in retirement saving. One must also exercise caution in generalizing from a single context, particularly since the ways that employers implement and market HSAs can differ. This setting is arguably important, though, because the firm is a health insurer. Its employees likely have a high degree of health insurance literacy compared to other settings. Employees do appear to understand several key features about HDHPs and HSAs, including that preventive care is exempt from the deductible and that HSA funds do not expire at year's end, yet mistakes in HSA saving and withdrawals are still prevalent.

Directly testing mechanisms that drive consumer HSA decisions is an important topic for future research. The results in this setting are consistent with mental accounting, as well as misperceptions about certain features of the HSA. For policies that strongly steer consumers toward one decision, such as a default contribution rate or restricting the choice set of insurance plans, pinpointing which behavioral mechanism operates may not matter for welfare evaluation (Handel and Schwartzstein 2018). Yet distinguishing the role of liquidity from behavioral factors is important to inform policy design in this area. The use of HSAs as self-insurance requires sufficient liquidity to meaningfully trade-off current versus future consumption. I find suggestive evidence against some forms of liquidity constraints in this context, at least as the primary explanation for the high MPC. Future work might consider a direct examination of borrowing constraints in the context of HSAs, as well as characterizing which mechanisms and constraints matter most for different types of consumers.

REFERENCES

Aaron, Henry, Patrick Healy, and Surachai Khitatrakun. 2008. "What's in a Name? Are Health Savings Accounts Really Health Savings Accounts?" In *Using Taxes to Reform Health Insurance,* edited by Henry J. Aaron and Leonard E. Burman, 92–118. Washington, DC: Brookings Institution.

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

Abaluck, Jason, and Jonathan Gruber. 2016. "Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance." *American Economic Review* 106 (8): 2145–84.

Abeler, Johannes, and Felix Marklein. 2017. "Fungibility, Labels, and Consumption." *Journal of the European Economic Association* 15 (1): 99–127.

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

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg. 2001. "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation." *Journal of Economic Perspectives* 15 (3): 47–68.

Arias, Elizabeth, and Jiaquan Xu. 2019. *United States Life Tables, 2017.* Hyattsville, MD: National Center for Health Statistics.

Arias, Elizabeth, and Jiaquan Xu. 2020. "United States Life Tables, 2018." National Center for Health Statistics 69 (12), Hyattsville, MD. https://ftp.cdc.gov/pub/Health_Statistics/NCHS/Publications/NVSR/69-12/Table01.xlsx (accessed April 26, 2021).

Arrow, Kenneth J. 1963. "Uncertainty and the Welfare Economics of Medical Care." *American Economic Review* 53 (5): 941–73.

Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein. 2015. "Behavioral Hazard in Health Insurance." *Quarterly Journal of Economics* 130 (4): 1623–67.

Baicker, Katherine, William H. Dow, and Jonathan Wolfson. 2006. "Health Savings Accounts: Implications for Health Spending." *National Tax Journal* 59 (3): 463–75.

Bailey Jones, John, Mariacristina DeNardi, Eric French, Rory McGee, and Justin Kirschner. 2018. "The Lifetime Medical Spending of Retirees." NBER Working Paper 24599.

Benjamin, Daniel. 2003. "Does 401(k) Eligibility Increase Saving? Evidence from Propensity Score Subclassification." *Journal of Public Economics* 87 (5–6): 1259–90.

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian. 2019. "Behavioral Household Finance." In *Handbook of Behavioral Economics,* Vol. 1, edited by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, 177–276. Amsterdam: North Holland.

Bhargava, Saurabh, George Loewenstein, and Justin Sydnor. 2017. "Choose to Lose: Health Plan Choices from a Menu with Dominated Option." *Quarterly Journal of Economics* 132 (3): 1319–72.

Blewett, Lynn A., Julia A. Rivera Drew, Risa Griffin, and Kari C.W. Williams. 2019. "IPUMS Health Surveys: Medical Expenditure Panel Survey, Version 1.1 [dataset]." Minneapolis, MN: IPUMS. https://meps.ipums.org/meps/ (accessed April 26, 2021).

Board of Governors of the Federal Reserve System. 2014. *Report on the Economic Well-Being of U.S. Households in 2013*. Washington, DC: Federal Reserve Board.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2013. "Salience and Consumer Choice." *Journal of Political Economy* 121 (5): 803–43.

Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad. 2017. "What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics." *Quarterly Journal of Economics* 132 (3): 1261–1318.

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

Buntin, Melinda Beeuwkes, Amelia M. Haviland, Roland McDevitt, and Neeraj Sood. 2011. "Healthcare Spending and Preventive Care in High-Deductible and Consumer-Directed Health Plans." *American Journal of Managed Care* 17 (3): 222–30.

Cabral, Marika, and Neale Mahoney. 2019. "Externalities and Taxation of Supplemental Insurance: A Study of Medicare and Medigap." *American Economic Journal: Applied Economics* 11 (2): 37–73.

Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik. 2014. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica* 82 (6): 2295–2326.

Cardon, James, and Mark Showalter. 2007. "Insurance Choice and Tax-Preferred Health Savings Accounts." *Journal of Health Economics* 26 (2): 373–99.

Centers for Medicare and Medicaid Services (CMS). 2021. "National Health Expenditure Data: Health Expenditures by Age and Gender." https://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/NationalHealthExpendData/Age-and-Gender (accessed April 26, 2021).

Chandra, Amitabh, Benjamin Handel, and Joshua Schwartzstein. 2019. "Behavioral Economics and Health-Care Markets." In *Handbook of Behavioral Economics*, Vol. 2, edited by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, 459–502. Amsterdam: North Holland.

Chen, Song, Anthony T. Lo Sasso, and Aneesh Nandam. 2013. "Who Funds Their Health Savings Account and Why?" *International Journal of Health Care Finance and Economics* 13 (3–4): 219–32.

Chetty, Raj, John N. Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen. 2014. "Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark." *Quarterly Journal of Economics* 129 (3): 1141–1219.

Claxton, Gary, Matthew Rae, Anthony Damico, Gregory Young, Daniel McDermott, and Heidi Whitmore. 2019. *2019 Employer Health Benefits Survey.* Washington, DC: Kaiser Family Foundation.

Clingsmith, David, Asim Khwaja, and Michael Kremer. 2009. "Estimating the Impact of the Hajj: Religion and Tolerance in Islam's Global Gathering." *Quarterly Journal of Economics* 124 (3): 1133–70.

Cubanksi, Juliette, Wyatt Koma, Anthony Damico, and Tricia Neuman. 2019. *How Much Do Medicare Beneficiaries Spend Out of Pocket on Health Care?* Washington, DC: Kaiser Family Foundation.

Devenir. 2019. *2018 Year-End Devenir HSA Research Report.* Minneapolis, MN: Devenir Research.

Engen, Eric, William G. Gale, and John Karl Scholz. 1996. "The Illusory Effects of Saving Incentives on Saving." *Journal of Economic Perspectives* 10 (4): 113–38.

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

Ericson, Keith Marzilli, and Justin R. Sydnor. 2018. "Liquidity Constraints and the Value of Insurance." NBER Working Paper 24993.

Fang, Hanming, and Dan Silverman. 2009. "Time-Inconsistency and Welfare Program Participation: Evidence from the NLSY." *International Economic Review* 50 (4): 1043–77.

Fang, Hanming, and Yang Wang. 2015. "Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting, with an Application to Mammography Decisions." *International Economic Review* 56 (2): 565–96.

Feldstein, Martin. 2005. "Rethinking Social Insurance." *American Economic Review* 95 (1): 1–24.

Feldstein, Martin, and Daniel Altman. 2007. "Unemployment Insurance Savings Accounts." In *Tax Policy and the Economy,* Vol. 21, edited by James M. Poterba, 35–63. Chicago: University of Chicago Press.

Fidelity. 2014. "Retiree Health Costs Hold Steady." http://web.archive.org/web/20140612134019/https://www.fidelity.com/viewpoints/retirement/retirees-medical-expenses.

Fronstin, Paul. 2017. *Trends in Health Savings Account Balances, Contributions, Distributions, and Investments, 2011–2016: Statistics from the EBRI HSA Database.* Washington, DC: Employee Benefits Research Institute.

Fronstin, Paul. 2017. *Trends in Health Savings Account Balances, Contributions, Distributions, and Investments, 2011–2017: Estimates from the EBRI HSA Database.* Washington, DC: Employee Benefits Research Institute.

Fronstin, Paul. 2019. *Enrollment in HSA-Eligible Health Plans: Slow and Steady Growth Continued into 2018.* Washington, DC: Employee Benefits Research Institute.

Gelber, Alexander. 2011. "How Do 401(k)s Affect Saving? Evidence from Changes in 401(k) Eligibility." *American Economic Journal: Economic Policy* 3 (4): 103–22.

Gross, David B., and Nicholas S. Souleles. 2002. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence From Credit Card Data?" *Quarterly Journal of Economics* 117 (1): 149–85.

Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang. 2020. "The Marginal Propensity to Consume Over the Business Cycle." *American Economic Journal: Macroeconomics* 12 (2): 351–84.

Handel, Benjamin R. 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review* 103 (7): 2643–82.

Handel, Benjamin R., and Jonathan T. Kolstad. 2015. "Health Insurance for 'Humans': Information Frictions, Plan Choice, and Consumer Welfare." *American Economic Review* 105 (8): 2449–2500.

Handel, Benjamin, and Joshua Schwartzstein. 2018. "Frictions or Mental Gaps: What's Behind the Information We (Don't) Use and When Do We Care?" *Journal of Economic Perspectives* 32 (1): 155–78.

Hastings, Justine S., and Jesse M. Shapiro. 2013. "Fungibility and Consumer Choice: Evidence from Commodity Price Shocks." *Quarterly Journal of Economics* 128 (4): 1449–98.

Hastings, Justine, and Jesse M. Shapiro. 2018. "How Are SNAP Benefits Spent? Evidence From a Retail Panel." *American Economic Review* 108 (12): 3493–3540.

Hastings, Justine S., Brigitte C. Madrian, and William L. Skimmyhorn. 2013. "Financial Literacy, Financial Education, and Economic Outcomes." *Annual Review of Economics* 5: 347–73.

Haviland, Amelia M., Matthew D. Eisenberg, Ateev Mehrotra, Peter J. Huckfeldt, and Neeraj Sood. 2016. "Do 'Consumer Directed' Health Plans Bend the Cost Curve over Time?" *Journal of Health Economics* 46: 33–51.

Heiss, Florian, Adam Leive, Daniel McFadden, and Joachim Winter. 2013. "Plan Selection in Medicare Part D: Evidence from Administrative Data." *Journal of Health Economics* 32 (6): 1325–44.

Helmchen, Lorens A., David W. Brown, Ithai Z. Lurie, and Anthony T. Lo Sasso. 2015. "Health Savings Accounts: Growth Concentrated Among High-Income Households and Large Employers." *Health Affairs* 34 (9): 1594–98.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy* 103 (2): 360–99.

Jones, Damon, and Aprajit Mahajan. 2015. "Time Inconsistency and Saving: Experimental Evidence from Low-Income Tax Filers." NBER Working Paper 21272.

Ketcham, Jonathan D., Claudio Lucarelli, Eugenio J. Miravete, and M. Christopher Roebuck. 2012. "Sinking, Swimming, or Learning to Swim in Medicare Part D." *American Economic Review* 102 (6): 2639–73.

Kling, Jeffrey. 2006. *Fundamental Restructuring of Unemployment Insurance: Wage-Loss Insurance and Temporary Earnings Replacement Accounts.* Washington, DC: Brookings Institution.

Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz. 2007. "Experimental Analysis of Neighborhood Effects." *Econometrica* 75 (1): 83–119.

Kling, Jeffrey R., Sendhil Mullainathan, Eldar Shafir, Lee Vermeulen, and Marian V. Wrobel. 2012. "Comparison Friction: Experimental Evidence from Medicare Drug Plans." *Quarterly Journal of Economics* 127 (1): 199–235.

Kooreman, Peter. 2000. "The Labeling Effect of a Child Benefit System." *American Economic Review* 90 (3): 571–83.

Kőszegi, Botond, and Filip Matějka. 2020. "Choice Simplification: A Theory of Mental Budgeting and Naive Diversification." *Quarterly Journal of Economics* 135 (2): 1153–1207.

Kueng, Lorenz. 2018. "Excess Sensitivity of High-Income Consumers." *Quarterly Journal of Economics* 133 (4): 1693–1751.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112 (2): 443–77.Laibson, David I., Andrea Repetto, and Jeremy Tobacman. 1998. "Self-Control and Saving for Retirement." *Brookings Papers on Economic Activity* 29 (1): 91–196.

Laibson, David, Andrea Repetto, and Jeremy Tobacman. 2003. "A Debt Puzzle." In *Knowledge, Information, and Expectations in Modern Economics: In Honor of Edmund S. Phelps,* edited by Phillipe Aghion, Roman Frydman, Joseph Stiglitz, and Michael Woodford, 228–66. Princeton, NJ: Princeton University Press.

Laibson, David, Peter Maxted, Andrea Repetto, and Jeremy Tobacman. 2015. "Estimating Discount Functions with Consumption Choices over the Lifecycle." Unpublished.

Lassman, David, Micah Hartman, Benjamin Washington, Kimberly Andrews, and Aaron Catlin. 2014. "Health Spending Trends by Age and Gender: Selected Years 2002–10." *Health Affairs* 33 (5): 815–22.

Lavetti, Kurt J., Thomas DeLeire, and Nicolas R. Ziebarth. 2019. "How Do Low-Income Enrollees in the Affordable Care Act Marketplaces Respond to Cost-Sharing?" NBER Working Paper 26430.

Leive, Adam. 2022. "Replication Data for: Health Insurance Design Meets Saving Incentives: Consumer Responses to Complex Contracts." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E129302V1.

Loewenstein, George, Joelle Friedman, Barbara McGill, Sarah Ahmad, Suzanne Linck, Stacey Sinkula, John Beshears, et al. 2013. "Consumers' Misunderstanding of Health Insurance." *Journal of Health Economics* 32 (5): 850–62.

Lo Sasso, Anthony T., Lorens A. Helmchen, and Robert Kaestner. 2010. "The Effects of Consumer-Directed Health Plans on Health Care Spending." *Journal of Risk and Insurance* 77 (1): 85–103.

Lusardi, Annamaria, and Olivia S. Mitchell. 2014. "The Economic Importance of Financial Literacy." *Journal of Economic Literature* 52 (1): 5–44.

Lu, Timothy, and Olivia Mitchell. 2010. "Borrowing from Yourself: The Determinants of 401(k) Loan Patterns." Unpublished.

McCrary, Justin. 2008. "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test." *Journal of Econometrics* 142 (2): 698–714.

Milkman, Katherine L., and John Beshears. 2009. "Mental Accounting and Small Windfalls: Evidence from an Online Grocer." *Journal of Economic Behavior and Organization* 71 (2): 384–94.

Morduch, Jonathan, and Rachel Schneider. 2017. *The Financial Diaries: How American Families Cope in a World of Uncertainty*. Princeton, NJ: Princeton University Press.

NQF. 2019. "Behavioral Health and Substance Use Final Report: Spring 2018 Cycle." Washington, DC: National Quality Forum.

O'Donoghue, Ted, and Matthew Rabin. 1999. "Doing It Now or Later." *American Economic Review* 89 (1): 103–24.

Parente, Stephen T., and Roger Feldman. 2008. "Do HSA Choices Interact with Retirement Savings Decisions?" In *Tax Policy and the Economy,* Vol. 22, edited by James M. Poterba, 81–108. Chicago: University of Chicago Press.

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

Paserman, M. Daniele. 2008. "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation." *Economic Journal* 118 (531): 1418–52.

Pauly, Mark V. 1968. "The Economics of Moral Hazard: Comment." *American Economic Review* 58 (3): 531–37.

Peter, Richard, Sebastian Soika, and Petra Steinorth. 2016. "Health Insurance, Health Savings Accounts, and Healthcare Utilization." *Health Economics* 25 (3): 357–71.

Poterba, James M., Steven F. Venti, and David A. Wise. 1996. "How Retirement Saving Programs Increase Saving." *Journal of Economic Perspectives* 10 (4): 91–112.

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

Reed, Mary, Ilana Graetz, Vicki Fung, Joseph Newhouse, and John Hsu. 2012. "In Consumer-Directed Health Plans, A Majority of Patients were Unaware of Free or Low-Cost Preventive Care." *Health Affairs* 31 (2): 2641–48.

Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod. 2010. "Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications." *Tax Policy and the Economy* 24 (1): 69–110.

Schwartz, Aaron L., Bruce E. Landon, Adam G. Elshaug, Michael E. Chernew, and Michael McWilliams. 2014. "Measuring Low-Value Care in Medicare." *JAMA Internal Medicine* 174 (7): 1067–76.

Setty, Ofer. 2017. "Unemployment Insurance and Unemployment Accounts." *Journal of the European Economic Association* 15 (6): 1302–40.

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

Shefrin, Hersh M., and Richard H. Thaler. 1988. "The Behavioral Life-Cycle Hypothesis." *Economic Inquiry* 26 (4): 609–43.

Social Security Administration. 2021. "Actuarial Life Table." https://www.ssa.gov/oact/STATS/table4c6.html (accessed April 26, 2021).

Steinorth, Petra. 2011. "Impact of Health Savings Accounts on Precautionary Savings, Demand for Health Insurance, and Prevention Effort." *Journal of Health Economics* 30 (2): 458–65.

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

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

Thaler, Richard H. 1999. "Mental Accounting Matters." *Journal of Behavioral Decision Making* 12 (3): 183–206.

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

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