

Insurance without Commitment: Evidence from the ACA Marketplaces

Rebecca Diamond^{1 *}, Michael J. Dickstein^{2 *}, Timothy McQuade¹, and Petra Persson^{3 *}

¹Stanford Graduate School of Business

²Stern School of Business, New York University

³Department of Economics, Stanford University

*NBER

First draft: September 5, 2017

Last compiled: March 31, 2021

Abstract

We study the dynamics of participation and health care consumption in the Affordable Care Act's health insurance marketplaces. Unlike other health insurance contexts, we find individuals commonly drop coverage midyear—roughly 30% of enrollees exit within nine months of sign-up. While covered, dropouts spend more on health care than in the months before sign-up or after exit. We model the consequences of drop-out on equilibrium premiums and consumer welfare. While dropouts generate a type of adverse selection, the welfare effect from their participation is ambiguous and depends on the relative costs per month of part-year vs. full-year enrollees. In our empirical setting, we find that imposing a penalty that incentivizes participation for at least 3.5 months would lower premium levels and improve overall consumer welfare.

We thank Anais Galdin, Abhisit Jiranaphawiboon, Ziao Ju, Jiayu Lou, Lucienne Oyer, and Ryan Shyu for outstanding research assistance. Diamond appreciates the support of the Laura and John Arnold Foundation. We also thank numerous seminar and conference participants for helpful comments and suggestions. This paper subsumes prior research reported in the working paper, "Take-Up, Drop-Out, and Spending in ACA Marketplaces." Diamond: diamondr@stanford.edu; Dickstein: michael.dickstein@nyu.edu; McQuade: tmcquade@stanford.edu; Persson: perssonp@stanford.edu.

1 Introduction

Governments often subsidize social insurance to protect constituents against adverse risks to their income and health. Despite the risk protection, take-up of these programs is less than universal, due to factors including stigma, a lack of information, and transaction costs (Currie, 2006).¹ For social insurance programs in which enrollees must pay a premium for coverage, price constitutes an additional hurdle to enrollment.

While the academic literature and policymakers often focus on take-up, initial enrollment alone does not guarantee an insurance benefit. Under the one-sided commitment design typical of these insurance programs, enrollees can opt out in the middle of the plan year. During the remaining months, consumers who drop coverage save premium costs but lose risk protection. Consumers who exit also generate an externality: when enrollment is fluid, consumers may enroll based on need and may concentrate their subsidized consumption into a short period of enrollment. The sign of this externality, however, is unclear in theory. Allowing consumers to drop coverage can lead to adverse selection and moral hazard, driving up insurance premiums for all, including those consumers most in need of full-year coverage. Alternatively, the ability to drop out might attract low-cost enrollees to the market. In effect, drop-out can offer insurers a means to price discriminate between low- and high-cost populations, leading to increased welfare.

We study the effect of one-sided commitment in insurance contracts in the context of government-subsidized individual health insurance. The market for individual health plans expanded under the Patient Protection and Affordable Care Act (ACA), passed in March 2010. The ACA established health insurance marketplaces, where individual consumers could shop for health coverage. The law also regulated the types of plans available for purchase and the ability of insurers to reject applicants. In the first years of the marketplaces, the number of US residents covered by individual insurance rose from an average monthly level of 10.6 million in 2013 to 17.4 million in 2015 (The Kaiser Family Foundation, 2016).

In this setting, we have two research aims. First, we measure the extent of attrition and compare the health spending behavior of consumers who exit insurance early with those who remain covered. We find significant drop-out from individual insurance, with roughly 30% exiting coverage within nine months of sign-up. Households who exit early appear to carry out more health transactions in the period of individual coverage than before or after enrollment. Second, because attrition generates an externality on full-year enrollees, we develop a model to link the attrition behavior to changes in premiums and welfare in the market. Our model shows that the presence of dropouts has ambiguous welfare consequences. We then take this model to the data, and examine how the presence of dropouts affect welfare in our empirical context.

¹Currie (2006) reviews take-up of various programs, which is approximately 75% for Earned Income Tax Credit (EITC) and as low as 6-14% for the State Children’s Health Insurance Program (SCHIP). Also see Kleven and Kopczuk (2011) and Bhargava and Manoli (2015).

A central empirical challenge to the study of enrollee behavior before, during, and after enrollment in the individual insurance marketplace is that it requires measures of health spending for a household under all means of coverage, including when the household lacks insurance. Thus, even with access to all-payer claims data, one would still miss household health spending while uninsured.² We overcome this data hurdle by exploiting transaction-level bank account and credit card micro-data. These data allow us to identify new enrollees using premiums paid directly to insurers. We observe enrollment at a granular level, along with measures of out-of-pocket health spending. A shortcoming of these data is that relative to the type of claims data more commonly used to study health insurance markets, our health measures are not as rich. However, we are able to observe a household’s health consumption consistently across periods when claims data would not – including during periods of uninsurance, during job and insurance transitions, and during moves across states. We also directly observe granular measures of income and household expenditures on non-health goods and services. These additional measures of spending enable use to investigate mechanisms potentially driving consumers to drop insurance mid-year.

Using our transaction data for households in California, the state with the largest individual insurance market, we show that attrition from individual health insurance is common across all income groups, including among newly enrolling households who receive government premium subsidies. Mid-year drop-out rates increased over 100% under the ACA expansion, relative to drop-out rates in the individual market prior to expansion. We validate our analysis using administrative data from California’s state-based marketplace and find similar exit rates.

The optimal policy response to drop-out depends on the basis for this behavior. We therefore use the rich information in our transactions data to examine the financial circumstances surrounding a household’s drop-out decision. First, we assess mechanisms we categorize as “non-health-driven,” including drop-out due to job transitions or changes in public insurance eligibility. While we find drop-out to be strongly correlated with income fluctuations in the pre-ACA period, we find a substantially weaker correlation after the implementation of the ACA. Second, we assess whether the data is consistent with “health-driven” mechanisms, where enrollment decisions reflect a household’s health status or expected health care needs. Specifically, we examine whether enrollees increase their health care spending upon enrollment and decrease it after exit. We find, on average, a 25% increase in spending during coverage for the drop-out population. Importantly, for these consumers, we do not see a similar spike and fall for non-health spending or for specific periodic payments, like utility bills or video streaming subscriptions.³ Thus, insurance incentives appear, at least in part, to drive the distinct pattern we observe in health care consumption.

²Periods of uninsurance appear common in the population we study. According to data from the Medical Expenditure Panel Survey, the average household who enrolled in the individual market in 2014 lacked insurance for almost six months in 2013.

³Since other monthly bills do not decrease at the same time, it does not appear that the consumer has changed their means of payment but rather stopped paying for health insurance.

To analyze how this drop-out behavior affects consumer welfare, we develop a simple framework in the spirit of [Einav et al. \(2010\)](#). Our key innovation is to introduce one-sided commitment by adding heterogeneity in the ability of households to re-time their health spending to the early part of the plan year and subsequently exit coverage.

Even with a relatively simple framework, we can make a key insight on welfare. We start by describing a standard reading of the economics of drop-out. In the absence of traditional adverse selection, the presence of dropouts always lowers overall welfare. This finding is intuitive: even if dropouts and full-year enrollees have the same annual health care costs, the drop-out household's ability to concentrate its spending and enrollment into fewer months raises the household's per month costs. In a competitive insurance market where insurers set monthly premiums once per year, the presence of dropouts leads insurers to raise premiums in equilibrium to compensate for dropouts' higher average costs per month enrolled. In this equilibrium, too many dropouts and too few non-dropouts enroll.

However, when we enrich the model to include both drop-out and traditional adverse selection, we show attrition need not always harm welfare in the insurance market. In fact, in this alternate reading of our setting, one-sided commitment contracts can raise welfare *for all* households. We illustrate this result using a two-type special case of our model. We show that if the enrollees who would choose to drop coverage mid-year spend less on health care than do full-year enrollees on an annual basis, allowing the option to drop out can improve market efficiency. Without the option to drop coverage mid-year, healthier dropouts might choose not enroll in insurance at all, as the annual premium may be too high given their low level of health care needs. The option to drop out provides a type of price discrimination that can encourage these low-cost households to enroll, bringing down the average monthly cost in the insurance pool. Conversely, if dropouts instead excessively concentrate their spending into a few months, their cost per month of enrollment may be higher than the full-year enrollees. This behavior would drive up average cost and decrease market efficiency.

Our framework suggests a connection between the rate of drop-out and the level of consumer welfare. In the two-type special case, we show that one needs only measure the ratio of monthly spending between dropouts and non-dropouts; if dropouts spend more per month covered, allowing attrition will reduce welfare. We compute this ratio using our California sample from years 2014-2015 and find dropouts appear to lower welfare: relative to full-year enrollees, dropouts conduct on average 23% more health transactions per month covered.

Of course, while the two-type set-up is helpful to illustrate our model's key intuition, a more realistic analysis allows for greater heterogeneity in drop-out propensity and risk aversion or willingness to pay for insurance. In particular, if we aim to calibrate a penalty for our observed market setting we need to allow for this heterogeneity. We thus move to apply our model in a setting with enrollees who differ in the timing of their plan exit and who differ in their sensitivity to premium levels.

Using this framework, and given our empirical estimates, we discuss alternative penalty designs to minimize possible distortions from drop-out.

In our California market setting, we again find greater health transactions per month for dropouts relative to full-year enrollees, defining enrollment categories of 2-4 months, 5-8 months, and 9-12 months. We estimate a price sensitivity for each enrollment category using a difference-in-differences estimator using differential time variation of prices across carriers within each geographic region. Combining our estimates of price sensitivity with the ratio of costs, we apply our model to design a penalty to maximize consumer welfare. We find the optimal penalty would incentivize households who enroll to maintain some form of insurance for at least 3.6 months.

Finally, we use our model to compare drop-out penalties with mandate penalties. If we narrowly consider policies intended to limit the consequences of drop-out, we find that the optimal penalty targets drop-out rather than non-enrollment. That is, rather than design a penalty around months of uninsurance, penalties could instead take the form of high deductibles or prepayment requirements, for example. Consumers who never sign up would not be subject to a penalty.

Our description of consumers as multi-dimensional – with a specific health spending need and a propensity to drop out – relates to past empirical work testing for sources of informational advantages in health markets, including [Finkelstein and McGarry \(2006\)](#), [Fang et al. \(2008\)](#), and [Shepard \(2016\)](#). In our setting, the unknown proportion of the population that will re-time its health care expenditures and subsequently drop coverage generates adverse selection. This reasoning is similar to [Cabral \(2017\)](#), who shows that re-timing of claims in the dental market can generate adverse selection.

Our paper also contributes to a growing literature on participation, selection, and pricing in individual insurance markets, including [Hackmann et al. \(2015\)](#), [Ericson and Starc \(2015\)](#), [Finkelstein et al. \(2019\)](#), and [Handel et al. \(2015\)](#). Specific to ACA reforms, [Antwi et al. \(2015\)](#) and [Simon et al. \(2017\)](#) analyze changes in health behaviors exploiting variation from the ACA’s Medicaid expansion. [Dafny et al. \(2015\)](#), [Dickstein et al. \(2015\)](#), [Kowalski \(2014\)](#), [Panhans \(2019\)](#), [Saltzman \(2019\)](#), and [Tebaldi \(2017\)](#) provide early evidence on the relationship between entry, pricing and adverse selection. Our contribution is to analyze the role of drop-out behavior on pricing and welfare.

More broadly, our analysis of drop-out behavior contributes to the literature on selection on moral hazard ([Einav et al., 2013](#)) and dynamic contracting with one-sided commitment in insurance, including long-term insurance contract design ([Hendel and Lizzeri \(2003\)](#), [Herring and Pauly \(2006\)](#), [Ghili et al. \(2020\)](#), and [Atal et al. \(2020\)](#)). That a form of adverse selection arises due to enrollees’ ability to drop coverage suggests that the optimal penalty may involve front-loaded payments, as in long-term contracts.

The remainder of the paper proceeds as follows. Section 2 describes the market for individual insur-

ance in the United States under the ACA. Section 3 outlines our data and the algorithm we use to measure health spending. In Section 4, we provide evidence on the lack of commitment of enrollees. In Sections 5 and 6, we model how the presence of dropouts can affect equilibrium premiums and welfare, and how penalties targeting drop-out can affect welfare. Section 7 concludes.

2 Institutional Details

We study the market for individual insurance. In 2013, prior to the implementation of the ACA, just 4% of the U.S. population purchased individual insurance coverage. More commonly, households receive coverage through an employer (50%) or through public insurance programs, including Medicaid, Medicare, and military and veterans health care (33%). The remaining 13% of households lacked insurance coverage (The Kaiser Family Foundation, 2016). After the ACA, in 2014 and 2015, the share covered by individual market insurance rose to 6 and 7%, respectively. The individual market grew in part due to a new insurance mandate and new regulations on the market. We describe each in turn.

First, to increase the rate of coverage, the ACA mandated that most US citizens and legal residents obtain health insurance. For the poorest adults, those under 138% of the federal poverty guidelines, the ACA expanded access to means-tested insurance via Medicaid.⁴ For households not covered by public insurance or insurance through an employer, the mandate required them to purchase individual insurance or face a tax penalty.⁵

To purchase a plan, consumers could contract directly with an insurer, use a broker-intermediary, or shop on newly created marketplaces. Depending on their income and the specific plan chosen, consumers purchasing in the marketplaces could be eligible for premium subsidies and cost-sharing subsidies.⁶ Premium subsidies, which depend on both the individual’s income and the premium of a benchmark plan in the consumer’s geographic region, are “advanceable”—in most cases, the government pays the subsidy directly to the consumer’s chosen insurer. In our data we observe premiums paid by bank or credit accounts; our measure of premiums is thus net of any subsidies the federal government pays the insurer in advance.

Second, the ACA created national regulations for the individual market. Prior to the passage of the ACA, the requirements on insurers serving the individual market differed by state. In June 2012, for example, California, along with 31 other states, had no rate restrictions on premiums in

⁴A subsequent Supreme Court ruling in 2012, however, allowed states to opt out of the expansion; as of May 2020, 37 states including the District of Columbia chose to expand Medicaid. The eligibility thresholds also differ (and are often lower) for parents of dependent children relative to childless adults.

⁵The tax bill signed into law in December 2017 eliminated the tax penalty underlying the individual mandate as of January 2019, but left other features of the law intact.

⁶In detail, the premium subsidy design sets a cap on how much of a household’s income must be spent to enroll in the benchmark plan, for households below 400% of the federal poverty guidelines. In 2017, for example, an individual with income between 100 and 133% of the poverty line would be required to pay no more than 2.04% of her income.

the individual market. In 2012, only six states required insurers to issue individual insurance to any applicant ([The Kaiser Family Foundation, 2012](#)).

In effect, the ACA harmonized the rules across states. Insurers must now issue plans to all consumers who apply during an annual open enrollment period, without regard to a consumer’s past illnesses.⁷ The law also limits the ability of insurers to set prices freely. Premiums can vary only according to four factors: (1) age, with at most a 3:1 ratio of premiums for the oldest to youngest enrollees; (2) geographic rating area; (3) family composition; and (4) tobacco use, limited to a 1.5:1 ratio. Insurers wishing to offer plans in the individual market must declare their interest in entering a particular geographic market and detail specific plan options and monthly premiums before the plan year. Those premiums will then be fixed over the plan year.

The plan offerings themselves also must fit into standardized bins based on actuarial value (AV). Those bins include: Bronze (60% AV), Silver (70% AV), Gold (80% AV), and Platinum (90% AV).⁸ California imposes several additional regulations on insurers beyond those required under the ACA. In particular, when insurers choose to serve a geographic market in California, they must offer at least one plan of each metal type ([Tebaldi, 2017](#)). Furthermore, and in contrast to other state marketplaces, the plan offerings are uniform within a metal tier. That is, California dictates the exact co-payments, deductibles, and out-of-pocket maximums for all plans in a given metal tier. Only the premiums and the network of physicians and hospitals included in the plan’s network may differ by insurer. We report these plan design features in Appendix Table [A1](#) for the 2014 and 2015 enrollment years.

3 Data

We use two principal sources of data: micro-data on consumers’ financial transactions, and household enrollment from the Covered California marketplace. We describe each source in turn and then validate our transactions sample using external data.

3.1 Financial transactions data

Our financial transactions micro-data come from a company that provides services to many banks, including five of the top ten U.S. banks. The company collects transaction-level data from bank clients’ accounts and any bank-issued credit cards. Unlike platforms such as Mint.com, which requires active user opt-in, our data provider collects all bank client data; selection into this dataset thus depends on a consumer’s choice of financial institution.

⁷Insurers must also allow individuals 60-day special enrollment periods for potential enrollees with a qualifying life event. Such life events include losing health coverage, getting married, having a baby, or adopting a child. In our empirical work, we focus on enrollment decisions in the open enrollment period separately from those occurring after qualifying life events.

⁸Catastrophic plans with lower actuarial value could be offered to young adults under 30 ([Patient Protection and Affordable Care Act, 2010](#)).

Sample definition We obtained a random sample of 9.2 million “active user” households in the United States.⁹ For each active user, the data link all bank accounts and credit card accounts within a given financial institution. While a user can be either an individual or a household, we show later that the income measures we observe closely match the household income distribution in Census data. Further, analysis of the 2013 Survey of Consumer Finance shows that 70% of households have all their checking and savings accounts at single bank. Among the remaining 30%, these households do the vast majority of their transactions at a single bank. We therefore describe users as households.

For each household account, we observe line-item transaction data for the years 2012 through 2015. We observe each transaction’s date, amount, merchant name and address, as well as a text description. We also have information about the broad category of the transaction (restaurants, utilities, and so on), as labeled by the data vendor. Appendix Table A2 provides examples of transactions and key variables.

For reasons of data availability, we narrow our focus to households in California. We identify the location of each household using the modal county of the physical merchants with whom the household transacts. After excluding households with ambiguous California residency and other data quality issues, we restrict attention to 798,085 “non-mover” California households.¹⁰

Next, to identify enrollment in individual health insurance among our California households, we compile a list of all health insurers participating in the Covered California marketplace in 2014 or 2015. We then label “enrollees” as those households we observe paying premiums to these insurers. In this way, our classification omits households who obtain insurance through their employer, as premium payments would flow through the employer’s payroll and not a household’s personal account. By the nature of our data, we also miss enrollees who use physical checks or cash to pay their premiums. We verify later through public data sources that this restriction does not lead us to sample systematically from a selected set of insurance carriers.

For duration of enrollment, we define *sign-up* to occur in the year and month of the first premium payment, and *drop-out* to occur in the year and month of the last premium payment. In this way, we allow flexibility in the observed payments—households can skip one or more months of premium payments and make a catch-up payment afterward without falling into our drop-out category.

With these definitions, our baseline analysis sample includes 106,904 households we identify as enrollees in individual insurance, the “insurance group,” as well as a reference group. Our reference group contains a random sample of the households that we never observe enrolling in individual

⁹Our data provider defines an active user as a client with frequent transactions. We further restrict our sample by looking for transaction activity during the months we observe the user’s account. To enter our sample, users must carry out transactions in greater than half of the months we observe.

¹⁰We exclude households that move to or from California during the sample period, but we allow moves within California. In Appendix A, we describe how we assign households to counties and how we construct our sample.

health insurance. We use the 139,000 households in the reference group in some analyses as a comparison set but not a true control from a causal perspective; the collection of reference households are representative of the rest of California and not balanced according to the characteristics in the insurance group.

Variable construction Our central outcomes of interest capture how utilization of health care evolves around sign-up and drop-out. We use two main measures: the dollar value of out-of-pocket health care spending and the number of health care transactions. To identify these transactions, we use a machine learning algorithm that exploits the text description for each transaction.

We provide more detail on this algorithm in Appendix B. In brief, to create our health measures, we use a Laplacian Corrected Naive Bayes machine learning algorithm to classify each bank transaction we observe into categories including ‘health’ and ‘drug’. For this procedure, we first create a training dataset in which we manually sort a sample of transactions based on their descriptions. Our data provider created its own classification of spending categories; we use this classification as a baseline and then construct our final training set by manually removing transactions that do not fit our classification needs. For the health spending category, for example, we eliminate transactions that private health insurance is unlikely to cover.¹¹ We calculate the frequency of each word appearing in these transactions in our training dataset, apply a modified version of Laplacian correction, and use the resulting probabilities to classify our full set of transactions.

For our health measures, we exclude drug spending. Drug spending proves harder to measure in our data because we only observe spending at pharmacy locations; we cannot distinguish dollars spent at pharmacies on drugs vs. other products, such as toiletries and food.¹²

Further, in parts of our analysis, we explore categories of non-health and non-drug spending as well as spending specific to households with children. We define “other spending” as total spending minus health insurance premium payments, drug spending, and health spending. In this pool of transactions, we can construct variables capturing specific types of periodic expenses, such as payments for auto insurance, utilities, and video streaming services. To identify child-related spending, we employ our machine learning algorithm in the same way we did to identify health spending. Appendix Figure A1 illustrates the top 100 words identified in each spending category.

In addition to spending, our financial transactions data provide a measure of household income. We calculate monthly post-tax income by aggregating all deposits into the user’s bank accounts and excluding transfers between accounts, withdrawals from brokerage accounts, wire transfers, and loan disbursements. To explore the interaction between individual insurance participation

¹¹For example, for health spending, we remove words classified as drug, as well as veterinary and animal-related services; lifestyle activities such as fitness, yoga, and spas; insurance premiums that contain health-related words; and dental and vision transactions.

¹²We nonetheless create two auxiliary outcomes, drug spending and the number of drug transactions. We repeat our main analyses using these drug variables in Appendix G.

and eligibility for Medicaid, we use our income measure to impute each household’s probability of Medicaid eligibility. We provide details of this imputation in Appendix E.

Finally, we categorize the health transactions observed in our household-level data. We use these categories to define urgent and essential care from more postponable care among a household’s expenditures. To avoid researcher bias in this definition, we design a survey for Amazon Mechanical Turk (MTurk) workers to classify each of our health-related words on a scale of *essentialness* and *urgency*. We describe the classification process in detail in Appendix F.

Panel A of Table 1 presents the mean and standard deviation of household-level variables, separately for the insured group and for our reference group. Over the four years in our sample, an insured household maintains individual insurance for 17 months on average; the reference group (by construction) maintains zero months. Roughly 81-82 percent of all households live in the same California county throughout the 2011-2015 sample period. 58 percent of the insured group and 48 percent of the reference group have at least one credit card linked to their bank accounts. Panel B of Table 1 summarizes our key spending variables. The median household in the insured group has monthly income of \$3,861 and makes a total of 29 transactions per month. On average, an insured group household conducts 0.31 health out-of-pocket transactions per month and spends \$23 on out-of-pocket health care per month.

Appendix Figure A2 illustrates the distribution of households across California in our baseline sample. The distribution, driven by bank prevalence, differs from the population distribution observed in American Community Survey. For example, our data contains an out-sized sample of households in the counties surrounding San Francisco. To adjust for this difference, we create weights that scale up households in under-represented counties in our data and scale down households in over-represented counties. We use these weights in all of our analyses.

3.2 Data from the Covered California Marketplace

Our second data source is the official enrollment data from the Covered California Marketplace (CCM), collected by Tebaldi (2017).¹³ The records cover all households that purchase private health insurance through the CCM from 2014 through 2016, totaling roughly 1.8 million households.

We observe household-month level information on enrollment, gross and net (of subsidies) premium payments, and household demographics. We construct variables capturing the month of sign-up, month of drop-out, and the duration of coverage. Using the gross and net premium amounts, we classify households as unsubsidized in a given year when their net premium equals their gross premium.

There are two key distinctions between our data sources. First, the CCM data includes the universe of enrollees who sign up through the marketplace. In contrast, the financial transactions data only

¹³We thank Pietro Tebaldi for generously providing access to this enrollment data.

include enrollees whose bank obtains services from our third party data provider. To enter our bank sample, households must also pay premiums using a debit or credit account that we observe. This requirement excludes households who pay by other means or who face no out-of-pocket premium for a subsidized plan.

Second, the CCM enrollment data do not include enrollees who buy individual health insurance outside of the marketplace, i.e., directly from insurance carriers or via brokers. In contrast, the financial transactions data allow us to observe sign-ups both on and off the CCM. Off-marketplace enrollment is substantial: between 2014 and 2016, 55-60% of all individual insurance enrollment occurred outside the official marketplace.¹⁴ The two data sources are thus complementary, and we will use both in our analyses.

3.3 Validation of Transactions Data

A key benefit of employing financial transactions data to study health insurance enrollment is the ability to examine health spending before and after enrollment and to examine *non*-health-related household behavior. In particular, we observe the household’s financial position – inflows as well as spending – both before and after health insurance enrollment. Further, we observe these financial variables for individuals who never choose to enroll. We exploit these unique features of our data to examine a range of potential economic mechanisms driving insurance participation.

Because of the novelty of our data, we validate our key income and spending measures before conducting our analyses. Where available, we collect similar variables observed in public data sources. We briefly summarize four validation exercises here.

First, we calculate the market share of each insurer by year using (1) the plan premium payments we observe in our transactions data and (2) the enrollment reported in both CCM data and federal data in 2014. Our goal is to show that our use of debit and credit premium payments does not bias our analysis toward specific insurers in California. However, we cannot simply compare the transactions-based shares to the CCM data because the latter do not reflect off-marketplace enrollment. To make the two sources comparable, we augment CCM data with overall individual market enrollment for each insurer. Insurers must submit these enrollment counts by year to the federal government to comply with Medical Loss Ratio regulation. Appendix D describes the construction of the market shares. Figure 1(a) shows that the transactions data closely match the administrative data.

Second, we compare the household-level income distributions in our transactions data against 2014 American Community Survey measure of household income. Figures 1(b) and 1(c) show that the two income distributions generally match well; for example, the correlation between median

¹⁴Appendix Table A4 compares the enrollment counts in CCM data against carrier-reported enrollment in federal government statistics.

household income by county is 0.87 when comparing ACS and our transactions data.

Third, Figure 1(d) benchmarks average household expenditure in 2014 in the transactions data against two other data sources, the Consumer Expenditure Survey (CEX) and the National Income and Product Accounts (NIPA). Our mean expenditure level in the transactions data matches closely the NIPA average expenditure.¹⁵

Fourth, we compare our measure of out-of-pocket (OOP) health spending to the OOP health spending of individuals in the Medical Expenditure Panel Survey (MEPS). We focus on MEPS survey respondents who report obtaining marketplace private health insurance coverage. While MEPS is helpful for validation, we cannot use it for our main analysis because of small sample concerns: in 2014, for example, only 402 MEPS survey participants enrolled in individual insurance under the ACA nationwide.

Appendix Table A5 compares summary statistics and the distribution of OOP health care costs in MEPS and in our bank data. The mean OOP health spending appears similar in each year (Panel A), at between \$205 and \$210 in 2012, rising to between \$247 and \$276 in 2015. Panel B shows that the transactions data also match the quantiles of the MEPS spending distribution. We describe this comparison in detail in Appendix C.

We also use MEPS to validate our transactions count measure. For many of our later analyses, we use the count of transactions for full vs. part-year enrollees as our proxy for insurer costs. In Appendix Table A8, we show that the percentage change in health transactions after individual insurance enrollment closely matches the percentage change in total health charges insurers incur in MEPS.

That our transactions-derived measures of insurer market share, income, overall household spending, and out-of-pocket health spending match statistics in comparable public surveys gives us more confidence in the reliability our transaction classification and our sampling procedure. There remains, however, the possibility of selection in our bank data when consumers shift transactions from one bank’s checking or credit card to another bank’s cards. In those instances, we would fail to observe the expenditure that occurs outside the initial bank. The truncation that results from such transitions could cause problems in our measurement of insurance duration in particular.

To lessen this concern, we design a sample of households least likely to suffer from this type of bias. As we describe in Appendix A, for example, we require households to maintain an active account with transactions observed in more than 50% of observed months. We can also directly assess whether our measure of drop-out incorrectly picks up households who switch their method of premium payment to a means that we cannot observe in our data. Specifically, we can exclude

¹⁵As noted in [Attanasio and Pistaferri \(2014\)](#), spending is generally under-reported in the CEX, whereas the NIPA, which reports the aggregate amount as measured by GDP, is often more accurate.

households who have credit cards for which we cannot observe transactions records.¹⁶ Appendix Figure A3 shows that the drop-out rates are essentially unchanged when restricting our sample to households who do not have un-linked credit cards. Thus, those households who may change their means of payment for premiums do not appear to drive our measure of plan exit.

4 Evidence on Commitment

Using both our transactions data and public enrollment statistics from the Covered California Marketplace (CCM), we start by examining participation in the individual insurance marketplace. We provide evidence on enrollee commitment to insurance and then explore the mechanisms driving participation decisions.

4.1 Participation rates

In Figure 2(a), we illustrate the share of exiters in our raw data. In the transactions data, we define the exit month as the last month in which we observe an insurance premium payment. Among enrollees who enroll in individual coverage for the first time during the open enrollment period in 2014, approximately 40% exit their plan before year end. The rate of exit by month is relatively flat across the year, with some slight increase in the beginning and end of the year.

We validate our measurement of drop-out by comparing our statistics to administrative enrollment from the Covered California marketplace (CCM). As we describe in Section 3 and Appendix D, the marketplace data and our transactions data cover different enrollee populations. The CCM data only track enrollment via the Covered California insurance platform while our financial data track both on- and off-marketplace enrollment. To make the measures comparable, we re-weight the unsubsidized households in the CCM to approximate off-marketplace enrollment. In Figure 2(a) we plot the distribution of the drop-out rate by months of enrollment for the 2014 open enrollment period. We compare the rates in the re-weighted CCM data and the transactions data. The exit patterns closely match across the two sources.

In Figure 2(b), we again focus on enrollees in the 2014 open enrollment period. We now illustrate the share of enrollees continuing coverage by income group. Here, by January 2015, only between 55 and 65% of enrollees maintain coverage. Crucially, while poorer enrollees drop out at slightly higher rates over the year, we observe a similar pattern of exit in all income groups, including households with annual post-tax income between \$100,000 and \$200,000.

While Figures 2(a) and 2(b) characterize drop-out rates for consumers enrolling in the 2014 open

¹⁶We identify households who have un-linked credit cards as follows: Using the bank account data, we observe the total amount the household spends paying off credit card balances, regardless of whether the credit card is linked in our dataset. We then compare these overall payments to those payments used to pay off balances on linked credit cards. If the payments to the linked credit cards are less than 85% of the total credit card spending out of the bank account, we consider this household to have an un-linked card.

enrollment period, a similar pattern of exit exists among both off-cycle enrollees in 2014 and among open and off-cycle enrollees in 2015. In Figure 2(c) we report the six-month drop-out rate by enrollment period. In both the transactions and re-weighted CCM data, enrollees new to insurance appear to drop out at *higher* rates in later enrollment periods. In our transactions data, we can also measure the rate of part-year exit before the implementation of the ACA. Pre-ACA, the dropout rate was 9%, which is about half of that seen in the years 2014 and 2015.

The fact that we observe drop-out before and after the ACA is not surprising per se – individual insurance has historically operated with one-sided commitment. The sharp increase in drop-out after 2014, however, suggests either a change in the composition of enrollees, as a function of new government subsidies, or a change in the consequences for early exit. Prior to the ACA, California did not require insurers to guarantee issue plans to all individuals. Illness or lack of coverage could serve as a basis for rejecting an insurance application. The one class of applicants in California protected against rejection under federal statute included individuals leaving group coverage for individual coverage. Even these consumers, however, could be rejected if they showed more than a 63-day gap in their insurance coverage ([The Kaiser Family Foundation, 2012](#)).

In sum, we observe widespread plan exit across all income groups, enrollment groups, and years, in both the transaction data and the administrative marketplace data.

Before discussing the mechanisms behind this drop-out, we briefly return to our discussion from Section 3 about consumers’ credit and debit card usage. If a consumer switches from paying her premiums and health consumption on a card we observe in the data to a card we do not, we will falsely interpret the switch as drop-out.

As mentioned earlier, we limit this concern by focusing on active accounts with transactions observed in more than 50% of observed months. In addition, our comparison to CCM data also provides support. If bank card switches drove our measure of drop-out, we would expect to find more dropouts than in CCM data, particularly at higher incomes where banking with multiple institutions is more common. In our data, we observe broadly similar drop-out rates to those in CCM data. We also find relatively lower drop-out rates at higher levels of income in the transactions data.

4.2 Mechanisms behind drop-out

The consequences of this widespread drop-out for insurance markets may depend on the underlying drivers of exit. We therefore use the rich information in our transactions data to examine the financial circumstances surrounding households’ drop-out decision. Our goal is not to identify a single mechanism and rule out others, but instead to leverage the unique content of our data to shed light on several potential drivers.

We distinguish two broad mechanisms. In the first, which we label “non-health-driven” enrollment,

insurance entry and exit decisions are not a reaction to a household member’s observed or perceived health. Instead, enrollment decisions depend on the availability of other coverage or reflect behavioral biases, such as forgetting to pay one’s insurance premium. In contrast, in the second broad mechanism, which we label “health-driven” enrollment, enrollment decisions depend primarily on a household’s health status or expected health care needs. We distinguish these two mechanisms as they have distinct consequences for consumer welfare in the health insurance market, an issue that we return to in Section 5.

In our subsequent analysis, we define the *drop-out population* as the set of households paying more than one but less than nine months of premiums in a calendar year. Further, we make one modification to this rule: we do not classify households exiting in November or December as dropouts, regardless of the duration of premium payments prior to drop-out. This definition of drop-out matches the insurance contract design under the ACA. Specifically, contracts for subsidized households under the ACA have a 90 day grace period in which households can make up for delinquent premium payments. Further, the ACA’s original mandate allows for a two month “short-gap” exemption from penalties, in effect allowing households to drop coverage for the last two months of a calendar year without penalty.¹⁷

4.2.1 Non-health-driven enrollment

To gauge the importance of non-health-driven enrollment, we study how enrollment evolves around (a) job changes, proxied by significant income changes, (b) liquidity issues, proxied by shifts in overall spending, (c) behavioral biases, proxied by changes in spending on periodic expenses, and (d) Medicaid eligibility.

In Figure 3, we illustrate income and expenditure dynamics around the period of enrollment for the drop-out population, as defined above.¹⁸ All months of coverage appear in the period labeled zero in the figure, to emphasize trends before and after coverage. In the plots, we compare the spending dynamics around enrollment in the period prior to the passage of the ACA with spending in the post-ACA period.

Figure 3(a) displays the relationship between changes in income and consumers’ decisions to sign up and drop out. In the pre-ACA period, shown as a dashed line, we see large percentage changes in income before and after the coverage period. Income rises 10-20% in the two months prior to enrollment and drops around 10% in the two months immediately following exit. This income

¹⁷We require more than one month of observed premium payments for a household to be considered an enrollee. Covered California actuaries and staff noted that mistaken enrollments or refunded enrollments might appear as one month premium payments in our transactions data. To avoid this potential measurement issue and to be conservative in our definition of drop-out, we exclude households with only one observed premium payment.

¹⁸In addition to the sample definitions outlined above, in this analysis of income and expenditure dynamics we only include households that appear in the transactions data for at least 10 months prior to sign-up and 10 months following drop-out and that have post-tax income less than \$100,000. We top-code all spending variables at the 99th percentile value within each income group.

pattern points to job changes or affordability as drivers of participation in the individual insurance market prior to the ACA. In contrast, the pattern post-ACA, shown as a solid line, illustrates a weaker correlation between income and participation. ACA enrollees exhibit a substantially smaller (5%) jump in income in the two months prior to sign-up, and no significant decrease in income upon drop-out. This distinct pattern suggests a smaller role for affordability in sign-up and drop-out post-ACA.

To examine the role of liquidity constraints, we examine the evolution of various household expenditure categories around sign-up and drop-out, again differentiating between the period before and after the ACA. We highlight overall non-health spending (Figure 3(b)), as well as spending on video streaming services (Figure 3(c)), utilities (Figure 3(d)), and auto insurance (Figure 3(e)). Similar to the patterns observed for income, we observe sharp increases in non-health expenditure immediately before insurance enrollment and sharp decreases upon drop-out in the pre-ACA period. This pattern suggests that liquidity constraints play an important role in households' insurance participation decisions in this period.

We observe distinct behavior after the implementation of the ACA. While spending does increase by about 10% upon sign-up, there is no drop in spending after drop-out. We see a similar pattern for expenditures on periodic items, such as auto-insurance payments, utilities, and video streaming services. This payment pattern suggests that both liquidity and behavioral biases, such as forgetting to pay bills for periodic payments, play a smaller role in drop-out after the ACA.

Finally, we examine whether the attrition pattern we observe is consistent with households transitioning into Medicaid. While we observe household income, we also need household family size and composition to determine Medicaid eligibility. To predict family size, we match our data to US Census demographic data by household income and county. We then construct a household's probability of being Medicaid-eligible at the monthly level. Appendix E contains our imputation procedure.

While our family size prediction is imperfect, changes in this Medicaid eligibility measure over time are driven by changes in household income, which we observe precisely. Appendix Figure A4 illustrates the dynamics of Medicaid eligibility around a household's observed sign-up and drop-out, both overall and separately by post-tax income bin. We see that a typical drop-out household has a probability of Medicaid eligibility around 20-30%, with an increase in eligibility of around 2-4 percentage points in the month following drop-out. The rate of increase is slightly higher among the lowest income population, but still changes only modestly. Thus, switching to Medicaid does not appear to be a substantial driver of the drop-out behavior we observe.

4.2.2 *Health-driven enrollment*

To examine whether the data is consistent with households choosing to enroll around expected or potential health care needs, we study patterns in health spending around enrollment. If households follow a pattern of health-driven enrollment, we would expect to observe spikes in health consumption around the time of sign-up and declines in health spending immediately following drop-out. While health-driven, such spikes need not reflect strategic motives. For example, while some consumers may sign up for insurance to defray the costs of discretionary health care needs, dropping coverage once they have satisfied these needs, others may also simply discover they are healthy after an initial visit; those healthy consumers may exit coverage after the positive health shock.

We plot health spending around enrollment in Figure 3(f). In the pre-ACA period, we do not observe a sharp decline in health spending upon drop-out. Spending remains roughly stable for six months after exit. This pre-ACA spending pattern could reflect households shifting to other private or public insurance coverage with similar out-of-pocket costs upon drop-out. Alternatively, we may observe stable spending upon drop-out if households, when facing the full cost of care without insurance, choose to consume less care at higher prices.

In the post-ACA period, we observe a different pattern. Spending increases 15-20% at sign-up and decreases roughly 20% at drop-out. The graphical evidence in Figure 3(f) suggests health spending needs may drive enrollment decisions more in the post-ACA period relative to the pre-ACA norm. We probe the relationship with two additional empirical exercises.

First, we use a difference-in-difference specification to examine households' health care consumption. Specifically, we compare households enrolling in the 2014 open enrollment period to households in the reference group who never enroll in individual insurance. For both groups, we compare spending levels before 2014, the start of the ACA's individual market reforms, to spending levels after 2014. We conduct our analysis separately by household income, as measured in 2013, and we include household fixed effects in the specification. Table 2 presents the results.

While the two groups of households we compare differ on some dimensions, the results nonetheless suggest substantial increases in health care consumption among ACA enrollees after they gain coverage. The pattern appears most pronounced at the lower end of the income distribution. For example, for households with less than \$20,000 in post-tax income per year, we see an 80% increase in out-of-pocket health care spending and an 87% increase in the number of health transactions. Given a pre-treatment mean of 0.96 transactions per year, the latter estimate effectively translates into about one extra health care visit per year, on average, among enrolling households.

Second, focusing only on the households that drop out, we use an event study framework to explore how their monthly health spending changes at sign-up and drop-out relative to the period of coverage. As in Figure 3, we collapse all months of coverage into one period. We compare dropouts

in the pre-ACA period, prior to July 2013, with dropouts who enroll at the beginning of 2014. Table 3 reports the results separately for the pre- and post-ACA periods and by income group. In our specifications we include household fixed effects and control for monthly income.

For the pre-ACA period, reported in panel (a) of Table 3, we observe no statistically significant change in health spending upon enrollment or after drop-out. By contrast, after the implementation of the ACA, enrollees who sign up and subsequently drop out have larger health care consumption during the period of coverage, especially at the lower end of the income distribution.¹⁹ For example, as reported in panel (b), health spending increases by 39% at sign-up and falls by 39% upon drop-out among households in the lowest income group. The pattern persists at higher income levels, although the magnitude of the effects in percentage terms falls. In panels (c) and (d), we repeat this exercise using the change in counts of health transactions instead of health spending. The results show similar changes in behavior around sign-up and drop-out.²⁰

Given the nature of our data, we can unpack the behavior behind these observed spikes. Specifically, we can use the wording in each health transaction to examine, for example, whether households that exit insurance use their coverage for discretionary and re-timeable health care spending. To conduct this analysis, we employ Amazon Mechanical Turk (MTurk) workers to classify health-related key words from our financial transactions data on a scale of essentialness and urgency. We describe our survey approach in detail in Appendix F.

Using the MTurk workers’ classification, we run the following difference-in-difference specification:

$$share_{it} = \alpha_i + \delta dropout_i + \lambda enrolled_{it} + \gamma dropout_i \times enrolled_{it} + \varepsilon_{it},$$

where $share_{it}$ refers to the percentage share of health spending household i completes at time t that is classified as either very essential or as very urgent.

For our analysis, we define very essential and very urgent health care as the set of transaction descriptions that include words that score in the 95th percentile of the MTurk score distributions across words (for urgency and essentialness, respectively). The variable $dropout_i$ is a binary variable equal to one if household i satisfies our definition of drop-out. The variable $enrolled_{it}$ is a binary variable equal to one for a household during the 10 months after sign-up and equal to 0 during the 10 months prior to enrollment.

¹⁹One potential interpretation of the results for the post-ACA period is that households transition into Medicaid; because Medicaid often involves zero copayments, we will not observe transactions following transitions to Medicaid. To control for this possibility, in our analysis we condition on monthly income and average lagged monthly income over the prior three months in our regression specifications.

²⁰In Appendix Table A6, we present a parallel analysis for prescription drug usage. While the results are noisier due to measurement error in drug spending, as discussed in Section 3 above, they are consistent with the results on health care spending.

Table 4 presents the results of our difference-in-difference analysis. In each panel, the regression estimates illustrate the difference in the types of spending that drop-out consumers complete relative to non-dropouts, comparing the period of enrollment to the period prior to enrollment. Panel (a) reports results for the pre-ACA period while panel (b) reports results following the implementation of the ACA.

In the post-ACA period, households that drop coverage increase their very essential and very urgent health care during their enrollment by 1.8-1.9 percentage points *less* than the non-dropout group. Relative to the overall health care consumption in our bin of very essential/urgent care, the smaller increase seen among dropouts represents a 35% decline in these categories. This pattern looks quite different in the pre-ACA period, as shown in panel (a). Before the ACA, those who dropped out increased their very urgent and essential care by *more* than those who did not drop out. Taken together, these patterns suggest that in the post-ACA period, we are more likely to find part-year enrollees shifting less urgent and less essential health spending to the period of coverage.

5 Participation and Welfare

Two clear empirical facts emerge from the analysis above. First, there is substantial early drop-out from individual health insurance plans. Second, households that drop out consume more health care during the period of coverage than during the months prior to sign-up or following drop-out. But how does this drop-out behavior affect the functioning of the individual insurance market? How does drop-out affect welfare?

To answer these questions, we develop a framework similar in structure to [Einav et al. \(2010\)](#) and the related literature on traditional adverse selection, where households with higher willingness to pay for insurance have higher expected costs. Our innovation relative to this framework is to introduce “drop-out types”. Drop-out households can re-time their health spending to the early part of the plan year and subsequently exit coverage. To isolate the effect of drop-out types on the insurance market, we first shut down traditional adverse selection. Following this analysis, we examine the effect of drop-out types in the presence of adverse selection.

Our framework delivers two key insights. First, in the absence of traditional adverse selection, the presence of drop-out types always lowers overall welfare. This is intuitive: dropouts pay premiums for only a portion of the plan year and hence are more expensive to insurers on a member-month basis. To recoup these costs, insurers must raise premiums on non-dropouts. Relative to the competitive equilibrium without drop-out, we show that too few non-dropouts enroll and too many dropouts enroll, leading to deadweight loss.

Our second result, however, suggests that in a richer and more realistic version of the model, drop-out types do not always harm insurance markets. In particular, in the presence of traditional adverse selection, permitting drop-out can be beneficial *even to full-year enrollees*, and can raise

overall welfare. Our framework provides the conditions under which we expect welfare to increase; we later match empirical measures to these conditions to determine whether the drop-out behavior that we document improves or harms welfare.

5.1 Model Set-Up

A household faces two possible coverage options. The first is to enroll in a single standardized insurance plan once a year. The premium is fixed for a year but paid in twelve monthly installments.²¹ Once a household enrolls in the plan, it chooses each month whether to continue paying for insurance for that month or to lapse and lose coverage for the remaining months of the year. The second option is to remain uninsured for the entire year, which entails no risk protection but also no premium payments. We label the two options “with insurance” (W) and “uninsured” (U).

We assume there are a finite number of household “drop-out” types $j \in \mathcal{J}$, which indexes the ability of households to re-time their healthcare consumption within the year. In particular, households of type j can accomplish all of their required health spending within the first $\phi_j \leq 12$ months of the year. We let N_j denote the mass of each type.

We normalize the utility from no insurance coverage to zero for all households. The utility from insurance depends on (a) the value placed on coverage, which we denote π_i for household i ; (b) the monthly premium, p ; and (c) the ability to re-time health spending. The utility from coverage for a household i of type j is given by:

$$U_{j,i}^W = \pi_i - \phi_j p, \quad (1)$$

where the idiosyncratic coverage valuations π_i have distribution $G_j(\pi)$. Finally, we denote the total expected health costs over the year for households of type j as c_j .²² With this formulation, our simple framework allows for both traditional adverse selection and for adverse selection from the ability to re-time consumption. Traditional adverse selection would arise, for example, if $c_j > c_{j'}$ and $G_j(\cdot)$ first-order stochastically dominates $G_{j'}(\cdot)$ for some $j, j' \in \mathcal{J}$.

Demand for insurance coverage is given by:

$$D(p) = \sum_{j \in \mathcal{J}} N_j [1 - G_j(\phi_j p)] \quad (2)$$

²¹The assumption that the premium is fixed for a year matches our empirical setting. We abstract in our model from the possibility of mid-year enrollment, which can occur in the insurance marketplaces when enrollees fall into one of several special categories, including, for example, the sudden loss of employer coverage.

²²More generally, we could assume a household i of type j receives quasilinear utility $u_j(\pi_i, q_i) - q_i p$ from $0 \leq q_i \leq 12$ months of coverage and allow households to optimize over months of coverage. Costs to the insurer would be given by $c_j(q_i)$. We recover our current specification by setting $u_j(\pi_i, q_i) = 0$ for $q_i < \phi_j$ and $u_j(\pi_i, q_i) = \pi_i$ for $q_i \geq \phi_j$. Then $c_j \equiv c_j(\phi_j)$. All of our key economic insights carry over to this more general framework. We focus on our preferred specification to provide more straightforward intuition and to link to our empirical exercise.

Average costs per year, on an enrollee member-month basis, are given by:

$$AC(p) = \frac{\sum_{j \in \mathcal{J}} N_j [1 - G_j(\phi_j p)] c_j}{\sum_{j \in \mathcal{J}} N_j [1 - G_j(\phi_j p)] \phi_j} \quad (3)$$

We assume competitive Bertrand pricing. That is, the equilibrium is determined by average cost pricing $p = AC(p)$, so that insurers make zero profits.

Social surplus is given by:

$$W = \sum_{j \in \mathcal{J}} N_j \int_{\pi_j^*}^{\infty} (\pi - c_j) dG_j(\pi), \quad (4)$$

where π_j^* denotes the marginal enrollees of type j . Social surplus is maximized when $\pi_j^* = c_j$; that is, households whose willingness to pay exceeds their marginal cost will purchase insurance while households whose willingness to pay lies below their marginal cost will not.

5.2 Equilibrium and Welfare

We next consider equilibrium and welfare. To make the intuition clear and allow for easier graphical illustration, we first consider a setting with two types of households. There is a mass N_D of households of “drop-out” type who can accomplish all of their required health spending within the first $\phi < 12$ months of the year. There is a mass N_{ND} of households of “non-drop-out” or “full-year” type who cannot re-time their consumption and need the entire year to accomplish all of their required health spending. That is, $\phi_{ND} = 12$. All of the key economic insights will carry over to the more general framework with more than two household types. Later, when we connect the model to data and study optimal drop-out penalties in Section 6, we use the more general model with more than two types to calibrate the penalty. Here, for graphical illustration purposes, we further assume that willingness to pay is uniformly distributed in both the drop-out and full-year populations. In particular, we assume that π_D is distributed uniformly between $[\pi_D^L, \pi_D^H]$ and π_{ND} is uniformly distributed between $[\pi_{ND}^L, \pi_{ND}^H]$.

5.2.1 No Traditional Adverse Selection

To isolate the effect of the novel friction that we introduce – the presence of drop-out types – we first consider a setting without traditional adverse selection. That is, we assume that $c_D = c_{ND}$. When $c = c_D = c_{ND}$ and in the absence of any drop-out, the competitive equilibrium achieves the efficient outcome. Since all households have the same expected costs, the only monthly price at which insurers break even is $p = c/12$. Thus, average cost pricing coincides with marginal cost pricing, maximizing social surplus. Graphically, the average cost and marginal cost curves are flat and coincide with each other.

If we start from this efficient outcome and introduce drop-out types, welfare unambiguously falls. Intuitively, when a subset of the population is able to drop out early, its actions create a distinct

form of adverse selection. We illustrate this outcome in Figure 4. Because drop-out types pay fewer monthly installments but obtain the same amount of health care (concentrated into the subset of months for which they pay), drop-out types are effectively less sensitive to the monthly premium. Thus, at high enough monthly premiums, only drop-out types purchase insurance; for lower premiums, some full-year types choose to purchase insurance as well. This purchase behavior generates a kink in the demand curve, illustrated by the solid line in Figure 4. Further, as drop-out types are costlier for the insurer on a member-month basis, the average cost curve slopes down in the range where both drop-out types and full-year types purchase insurance. We see this average cost curve as a dashed line in the figure.

The intersection of the demand and average cost curves determines the equilibrium premium. Because drop-out types are costlier than non-drop-out types on a member-month basis, the monthly premium in any equilibrium with dropouts always exceeds $c/12$, the monthly premium in the absence of dropouts. Thus, the presence of dropouts unambiguously harms full-year enrollees, who must pay higher monthly (and annual) premiums to allow insurers to break even when dropouts stop paying premiums mid-year. Under the higher equilibrium monthly premium, there is under-insurance among non-dropouts. Because the non-dropouts effectively subsidize the premiums of the dropouts, there is over-insurance among dropouts.²³

5.2.2 Drop-out with Traditional Adverse Selection

Our analysis so far omits traditional adverse selection. However, it seems likely that households who can re-time their consumption may have lower annual expenditures as well as lower willingness to pay for insurance than those who cannot. That is, the ability to drop out may be correlated with traditional adverse selection. We therefore now consider a setting in which $c_D < c_{ND}$, the annual costs for drop-out types are lower than for non-drop-out types, and $G_{ND}(\cdot)$ has first-order stochastic dominance over $G_D(\cdot)$. For our illustration, we continue to assume households' willingness to pay follows a uniform distribution.

When there is both drop-out and traditional adverse selection, the earlier insight – that drop-out always reduces welfare – ceases to hold. Under the right circumstances, allowing drop-out can instead increase social surplus relative to a ban on drop-out.

To understand this outcome, we start from a setting without drop-out. Here, we set $\phi = 12$ so that the drop-out types spread their health care consumption over the full plan year. This assumption brings us back to the standard adverse selection case with two types (non-drop-out types and drop-out types), where dropouts have both lower costs and lower willingness to pay. As usual, traditional adverse selection yields a declining average cost curve, an equilibrium price that exceeds

²³In the model, drop-out types recognize that they will not pay a full year of premiums when purchasing insurance. If instead drop-out types incorrectly anticipate paying the full annual premium, prices would still be higher and welfare lower relative to a world without drop-out. Enrollment of both drop-out and full-year types would fall.

marginal cost, and under-provision of insurance to the drop-out population (the low-cost type) in equilibrium.

Next, we consider a setting in which dropouts can concentrate their annual health care spending into fewer months—i.e. $\phi < 12$. We first examine a special case, $\phi/12 = c_D/c_{ND}$, where the fraction of the year drop-out types enroll ($\phi/12$) is equal to the ratio of the costs between the drop-out and non-drop-out populations (c_D/c_{ND}). Here, a monthly premium of $p = c_{ND}/12$ is an equilibrium outcome. Insurers will break even on the full-year types. On the drop-out types, the insurer’s annual expenditures are c_D . Their annual revenues are $\phi * p = \phi/12 * c_{ND} = c_D$, implying that insurers also break even on the drop-out types. Thus, $p = c_{ND}/12$ is indeed an equilibrium outcome. Moreover, since the implied annual premium for non-dropouts is c_{ND} and the implied annual premium for dropouts is c_D , this choice of p achieves the maximal social surplus. Households purchase insurance if and only if their willingness to pay exceeds marginal cost.

In short, when traditional adverse selection is present along with re-timeability for some households, drop-out can be beneficial: in effect, drop-out allows for a form of price discrimination in which high-cost enrollees pay higher prices than low-cost enrollees, all without insurers having to identify a household’s type at the time of enrollment. This price discrimination allows for greater provision of insurance while still allowing insurers to break even.

In the special case considered above, drop-out is so beneficial that it overcomes the entire traditional adverse selection problem. More generally, drop-out often alleviates – but does not fully solve – the traditional adverse selection problem. Furthermore, drop-out is more beneficial when the difference in annual health care costs between the two types is larger; intuitively, in such a setting, the value of effective price discrimination is higher.

To illustrate this finding, we describe the welfare outcome in two cost scenarios in Figure 5. In the upper panel, the simulated gap in annual health care costs between dropouts and non-dropouts ($c_{ND} - c_D$) is smaller; in the lower panel, the gap is larger. In each panel, we plot (a) the overall welfare, (b) the equilibrium premium, and (c) the consumer surplus, by enrollment type, as we vary the ability of drop-out types to re-time their health care consumption.

Specifically, in each graph, the x-axis is $\phi/12$, the fraction of the year that dropouts need to enroll in order to fulfill their yearly health care spending needs. As we move to the right along the x-axis, dropouts spread their annual health care costs over a gradually larger share of the year. The rightmost endpoint of $\phi/12 = 1$ is equivalent to a drop-out ban – the drop-out types must enroll for the full year to conduct all of their annual health care spending, equivalent to full-year types. Thus, at $\phi/12 = 1$ we are back to the standard adverse selection case with two types (full-year types and drop-out types), but where dropouts have lower costs.

In Panel 1 (a), when $\phi/12 = c_D/c_{ND}$, the market achieves the efficient outcome and maximizes social surplus, as discussed above. Moreover, in a range of $\phi/12$ close to $\phi/12 = c_D/c_{ND}$, welfare

exceeds the level achieved under a drop-out ban. Intuitively, in this range, the advantageous price discrimination effect dominates the adverse selection effect. For sufficiently low values of $\phi/12$, however, the market for full-year types collapses completely and only drop-out types enroll, leading to steep welfare losses.

Panels 1 (b) and (c) show that, in this cost scenario, the welfare gains relative to a drop-out ban accrue only to the dropouts. Here, full-year types always achieve their highest consumer surplus under a ban. The key reason for this outcome is the equilibrium premium that prevails when we vary $\phi/12$. As we move to the left in the figure, dropouts are able to concentrate the same annual health care spending into fewer months. Given their lower willingness-to-pay, the ability to pay fewer months of premiums leads more dropouts to enroll. However, the change in population also raises the insurer's average costs per month, and hence equilibrium premiums increase (panel 1 (b)). The increased premiums, in turn, reduce the consumer surplus of the full-year types (panel 1 (c)).

In contrast, Panel 2 (b) and (c) show that full-year types are not always harmed by the presence of early dropouts in the market. Indeed, in the simulation in Panel 2, drop-out types have substantially lower annual health care spending than full-year types. In this scenario, as we move from a drop-out ban ($\phi/12 = 1$) to allow some drop-out ($\phi/12 < 1$), we again observe that dropouts enroll in greater numbers. However, because dropouts have lower health spending needs in this scenario, they *reduce* the average annual health costs among all enrollees when we permit re-timing and early exit. In Panel 2 (b), for example, we observe a decrease in equilibrium premiums and an increase in the consumer surplus of full-year types when the share of the year required for drop-out enrollment falls from 100% to roughly 78%.

5.2.3 Characterizing welfare for non-drop-out households

Given that in theory the presence of dropouts can either improve or harm the welfare of non-dropouts, we develop a characterization of when we expect full-year enrollees to face unambiguously lower welfare with drop-out. We later connect this prediction to our empirical setting. In short, non-dropouts suffer when dropouts prove more expensive on a per-month basis. Formally:

Prediction 5.1. *Suppose that $c_D/\phi > c_{ND}/12$. Then the consumer surplus of the non-dropouts is (weakly) lower in the presence of dropouts.*^{24,25}

²⁴This result holds both when dropouts (incorrectly) expect to pay the full annual premium and when they expect to pay less than the full year of premiums.

²⁵Consumer surplus of non-dropouts will be strictly lower with drop-out as long as $G_{ND}(12p_{ND}) < 1$, where p_{ND} is the monthly premium under no drop-out. This condition holds when not all full-year types enroll in the absence of dropouts.

Proof. The consumer surplus of non-drop-out types is given by:

$$CS_{ND} = N_{ND} \int_{\pi_{ND}^*}^{\infty} (\pi_{ND} - 12p) dG_{ND}(\pi_{ND}), \quad (5)$$

where p is the monthly premium. This surplus is (weakly) decreasing as the monthly premium rises. Thus it suffices to show that the monthly premium allowing drop-out, p_D , is greater than the monthly premium under a drop-out ban, p_{ND} . If there is no drop-out, then the equilibrium premium p_{ND} must satisfy $c_{ND}/12 \geq p_{ND} \geq c_D/12$, since otherwise insurers would not make zero profits. On the other hand, with drop-out, it must be the case that the equilibrium monthly premium $p_D > c_{ND}/12$. Suppose otherwise. Then, the profits from non-dropouts would be negative. Moreover, since dropouts are more costly on a per month basis than non-dropouts, the profits from dropouts would be:

$$\phi p_D - c_D = \frac{\phi}{12} 12 p_D - c_D \leq \frac{\phi}{12} c_{ND} - c_D < c_D - c_D = 0.$$

Insurers would therefore make negative profits on both types of consumers, so $p_D \leq c_{ND}/12$ cannot be an equilibrium. Thus, $p_D > p_{ND}$ as desired. $\square$

5.3 Empirical evidence on model predictions

Our model provides a prediction of the welfare consequences from drop-out: the presence of dropouts lowers the welfare of non-dropouts when dropouts spend more per month while covered than full-year enrollees do. Thus, to comment on welfare, we can compute the ratio of the monthly costs of a drop-out household (during the months of coverage) to the monthly costs of full-year enrollees.

There are two caveats to taking this prediction from our stylized model to data. First, we note again that we do not observe insurers' costs directly in our transactions data. Instead, we employ our best available proxy, the count of transactions for households of different enrollment tenures. Second, in the individual market, insurers receive risk adjustment transfers (Geruso et al. (2019), Einav et al. (2019)), which may change their incentives to enroll patients of different disease characteristics and spending levels outside our simple model.

In particular, Ericson et al. (2018) note that at the start of the ACA, the risk adjustment formula used the same algorithm to predict partial-year and full-year enrollee's annual spending, but adjusted for the fraction of the year the enrollee participated. These authors illustrate how a fractional adjustment method systematically under-predicts spending. Even in 2017, when the federal government added the duration of coverage as a factor in its risk adjustment formulas, Dorn et al. (2018) show that these factors fail to fully account for the differences in spending.²⁶

²⁶Using data from two carriers in 2015, Dorn et al. (2018) found that the risk adjustment formula underpaid insurers for part-year members, both for enrollees in the special-enrollment periods and enrollees who signed up during open-enrollment periods and subsequently dropped coverage.

While the current risk adjustment policy redistributes the costs of early drop-out across individual market insurers, albeit imperfectly, our stylized model points to a more fundamental friction: even with perfect re-distribution, the presence of drop-outs raises average costs and premiums. In our model, for example, insurers are homogeneous and set premiums equal to average costs; there is no scope for redistribution. The added costs from drop-out, which we find can influence equilibrium premiums and thus consumer welfare, would persist with risk adjustment. Thus, we find it valuable to explore how the level of spending varies by drop-out type in our empirical setting.

In the model discussion above, we illustrate our model and its implications using two duration types. For our empirical examination, we expand to multiple durations. In Figure 6, we distinguish enrollees by their duration of coverage and income, and plot the monthly number of health care transactions they conduct while covered. We show patterns of spending among enrollees in the pre-ACA period vs. the post-ACA period in separate lines in the plot. We normalize the monthly health transactions to one for enrollees who stay enrolled for a full year or more (“13+ months”). We then compare the percentage difference in the average number of monthly health transactions across enrollees of different durations relative to the full-year benchmark.

In the post-ACA period in 2014, we find enrollees who remain covered for 2-4 months consume on average 60% more health care transactions per month when covered than those enrolled for the full year or more. This number falls with the duration of coverage, where those enrolled for 9-12 months may consume *less* health care per month than the long-term enrollees. In the pre-ACA period, the pattern is similar but smaller: 2-4 month duration enrollees spend only 20% more than their full-year counterparts.²⁷ Thus, in light of the model’s predictions and the observed transaction counts, we find evidence that drop-out likely lowers the welfare of full-year enrollees in our empirical setting.²⁸

6 Policy Discussion

We can use our model of insurance participation to discuss optimal policy design. Here, we focus on the design of penalties used to eliminate the distinct form of adverse selection that stems from drop-out.²⁹ We differentiate between two types of penalties: (1) a “mandate” penalty assessed whenever a household fails to sign up for insurance, say during an open-enrollment period, and

²⁷We calculate full-year enrollees’ monthly transaction counts using data from the corresponding months for dropouts. For example, we compare Jan-April consumption for both 2-4 month enrollees and full-year enrollees. In this way, we confirm that the added monthly spending we find for dropouts is not simply an artifact of seasonality. We also repeat our analysis using the full-year enrollees’ annual transaction count divided by 12, assuming an equal spacing of transactions throughout the year. With this approach, we find very similar results.

²⁸As an additional validation, we repeat our analysis using monthly health care out-of-pocket spending in place of counts of health care transactions. The findings are qualitatively similar in magnitude and sign.

²⁹When we discuss penalties, we do so narrowly, taking the set of consumers in our setting as given. For example, when we consider a “mandate” penalty, we ignore the role such a mandate plays in changing the distribution of costs among the full-year types in the population.

(2) a “drop-out” penalty assessed when a household signs up for insurance but becomes uninsured before the full plan term. Here, our model treats insurance exit as the decision to become uninsured. If a consumer switches from individual insurance to a form of group insurance or public insurance, they would not be subject to the drop-out penalty.

Our goal in this section is twofold. First, we calibrate the optimal drop-out penalty given data from our California setting. Second, we discuss the effects of combining a mandate and drop-out penalty.

6.1 Calibrating the penalty amount

Given our data on costs, plan choices, and premiums, we can quantitatively calibrate the penalty design in the California market that would maximize consumer welfare. We approximate demand using linear demand functions and assume insurer price at average cost. We assume a mandate penalty of \$0; later we consider the effect of pairing a mandate penalty with a drop-out penalty. Our policy goal is to use our penalty to discourage the type of drop-out that harms total welfare.

In our penalty analysis, we use the more general version of our model in Section 5. That is, rather than assume there are only two types of consumers, we now specify multiple drop-out types. In doing so, we can capture more flexibly the responses of different enrollee types to the imposition of a penalty.

We optimize over the following linear penalty design, similar to the actual implementation of the ACA. The penalty an enrollee pays for dropping out after m months is:

$$F(m) = p \max(m^* - m, 0), \quad (6)$$

where p is the monthly insurance premium and m^* is the required enrollment period chosen by the government. Given our model framework, it is clear any agent with $\phi_j < m^*$ will find it optimal to continue paying premiums for m^* months. Any enrollee with $\phi_j > m^*$ will pay ϕ_j months.³⁰ The government thus optimizes over m^* , the required enrollment period.

Of course, outside of the simplest two-type case, we are not able to design policy to achieve the first-best outcome. That is, with only one drop-out type, we showed the optimal penalty fixes the number of required months of participation equal to the ratio of costs between dropouts and non-dropouts. With multiple drop-out types and only a single penalty instrument m^* , we are left with second-best solutions.

Our approach is to search for the penalty design that will maximize overall consumer welfare across

³⁰It is without loss of generality to consider the linear penalty design above instead of the more general design $F(m) = \xi \max(m^* - m, 0)$ for some $\xi > 0$. If $\xi < p$, this is strictly worse than no penalty (i.e. $m^* = 0$), because no agent will delay dropping out and the system will simply incur deadweight losses. If $\xi > p$, the design achieves exactly the same outcome as $\xi = p$.

three broad enrollee types: 2-4 months tenure, 5-8 months tenure, and 9+ months tenure, which we defined earlier as our full-year population. To do so, we compute the ratio of costs between dropouts and non-dropouts as well as the market shares and the shape of demand by type.

In Table 5, we report empirical estimates needed for our penalty analysis. First, we estimate each group’s price sensitivity using differences-in-differences. We regress carrier-specific market share within each drop-out group, where we define market by year and geographic region, on a weighted average of the carrier’s premiums in that market.³¹ We control for region×year fixed effects and region×carrier fixed effects. With this fully balanced panel two-way fixed effect design, we identify consumers’ price sensitivities using the variation in price changes over time across carriers within a geographic region. This approach differences out any time-invariant characteristics of carriers’ plans within each region. It also controls for any region-specific trends in overall health insurance prices. We run this regression separately for sub-populations of our California households defined by their drop-out timing. We see from Table 5 that price sensitivity increases with tenure: short-term enrollees are less elastic in demand than long-term enrollees.

The second empirical measure we use to calibrate the optimal drop-out penalty is the ratio of costs between dropouts and non-dropouts. In panel (b) of Table 5 we show the ratio of costs for each drop-out tenure relative to the costs of non-dropouts. We compute these ratios using counts of health transactions in years 2014-2015. As in Figure 6, we see more health transactions per month for shorter tenures.

Combining these two empirical measures with the shares of each enrollee type, we simulate our model of demand and calculate welfare in a setting with perfect competition. Our calibration provides the number of months m^* the optimal penalty should target for enrollment. In Figure 7, we plot consumer surplus under these alternative targets. We distinguish the changes in surplus for 2-4 month drop-out types, 5-8 month drop-out types, and full-year types.

From the figure, we see that the welfare effects of a drop-out penalty are heterogeneous by drop-out type. For 2-4 month drop-out types, any penalty that requires more than three months of participation reduces their surplus. These consumers exit insurance as the required enrollment period increases; at four months of required prepayment, all 2-4 month drop-out types choose to be uninsured. A similar phenomenon occurs with the 5-8 month drop-out types. These households begin to exit at penalty levels that require more than six months of insurance.

For the full-year enrollment types, surplus is highest at a penalty design requiring approximately seven months of enrollment. In Figure 7, we illustrate this optimal number of required months with a solid vertical line. Relative to a market without drop-out penalties, full-year types achieve a surplus that is 8.8% higher under this optimal penalty level.

³¹The weighted average premium uses plan-specific gross premiums for a carrier, weighted by plan enrollment. We use the Covered California data for this analysis, adjusting for the size of off-marketplace enrollment.

Interestingly, while full-year enrollees prefer any strictly positive penalty to a policy with no penalty, the optimal policy from their perspective is not one that requires a full year of coverage, consistent with the theoretical intuitions laid out in the previous section. We find there is a cost to pushing the required window too high: too large a required enrollment period m^* discourages participation from all drop-out types, leaving only full-year types enrolled. Given the higher annual spending needs of the full-year types, premiums rise higher than when the market contains some healthier drop-out types who pay the required months of premiums. Thus, consumer surplus among full-year enrollees falls when the penalty design’s required enrollment period is too long.

When we consider the optimal penalty in the market overall, maximizing surplus across both drop-out and full-year types, the required enrollment period m^* is significantly shorter than the period that is optimal for the full-year enrollees alone. In Figure 7, we plot a dashed vertical line indicating the required months of enrollment that maximizes overall surplus. In our calibration, we find an optimal penalty requires enrollment of 3.6 months per year.

Relative to a scenario with no penalty on drop-out, overall surplus increases 1%. However, this change reflects a redistribution in the population. Those consumers who would typically drop out within 2-4 months fare the worst, receiving only 18% of the surplus they would achieve with no drop-out penalty. Full-year enrollees gain 4% in surplus because the presence of some healthier drop-out types under this policy lowers premiums in the market. The 5-8 month drop-out types gain the most: their surplus rises 14% relative to a market without drop-out penalties. Under the optimal penalty, households of this intermediate drop-out type benefit because they prefer to remain in the market but now face lower premiums.

Overall, we find welfare gains from relatively short required enrollment periods when there is a modest ratio of per-month health spending between dropouts and non-dropouts. If that ratio were larger, say in another time period or geographic market, penalties incentivizing longer enrollment would improve overall welfare.

6.2 Mandate vs. Drop-out Penalties

Finally, we use our model and the simulated setting from Section 5 to assess the welfare effect of a policy combining mandate and drop-out penalties. The original penalty design of the ACA, before the elimination of the mandate penalty in January 2019, took this form: Households without insurance faced a penalty that depended on household income and scaled with the number of months of non-enrollment in a year. We ask whether, when targeting penalties against drop-out, mandates can lower welfare in the presence of a drop-out penalty.

Our discussion in Section 5 in the two-type case provides our main insight. There, we found that a penalty that encourages drop-out types to maintain coverage for $\phi/12 = c_D/c_{ND}$ portion of the year— i.e. for the fraction of the year equal to the ratio of costs between dropouts (c_D) and full-year

enrollees (c_{ND}) – would achieve the first best level of welfare. Thus, under such an optimal drop-out penalty, any mandate penalty would necessarily reduce welfare by inducing households to sign up who would not in the first-best outcome.

Our welfare result depends specifically on drop-out households’ participation decisions. Under the optimal drop-out penalty, we observe that full-year consumers whose valuations exceed their costs sign up. For drop-out consumers, the penalty also induces only those households whose valuation exceeds their costs to sign up. With a mandate penalty, some drop-outs with valuations below their costs would face incentives to enroll, leading to lower welfare under our model.

In the absence of a drop-out penalty, however, a mandate penalty can improve welfare. As described in Section 5, when there are no penalties, the presence of dropouts forces insurers to charge higher premiums to break even. With higher premiums, some full-year types choose not to purchase even though their willingness-to-pay exceeds their own marginal costs. At the same time, with the ability to exit early, drop-out types who would not optimally purchase at the annual premium may purchase at the price of a few months of premiums. The result is a welfare loss. Under the parameters of our simulation, as show in Appendix Figure A5, we can find a mandate penalty that improves welfare by encouraging greater enrollment of full-year types.

7 Conclusion

A crucial component of social insurance design is to ensure the programs reach the targeted beneficiaries. While both the academic literature and policy responses have focused on take-up, we contribute a new analysis of drop-out from social insurance. We show that attrition is an important feature of such insurance programs that require beneficiaries to pay a recurrent premium, which differentiates them from employer-sponsored insurance and fully subsidized public insurance.

In particular, in the context of the ACA-established health insurance marketplaces for individual insurance, we document that attrition is widespread among new enrollees, even among the poorest newly enrolling households that receive government premium subsidies. In the 2014 and 2015 open enrollment years, roughly a third of all new enrollees in California drop out within nine months of sign-up. The households who drop out appear to spend more on health during the months of coverage relative to the period before enrollment and after drop-out.

Such attrition can have fundamental effects on market outcomes for both households who exit early and households who remain enrolled. Our model illustrates that one-sided commitment contracts can improve welfare if dropouts are relatively healthy. We provide (a) an empirical measure to determine when we expect attrition to harm overall welfare and (b) a method to calibrate drop-out penalties to maximize welfare. In our California setting, we find that, far short of an annual commitment, a drop-out penalty that requires as little as 3.6 months of enrollment upon sign-up can maximize welfare in the marketplace.


References

- ANTWI, Y., A. MORIYA, AND K. SIMON (2015): “Access to Health Insurance and the Use of Inpatient Medical Care: Evidence from the Affordable Care Act Young Adult Mandate,” *Journal of Health Economics*, 39, 171–187.
- ATAL, J. P., H. FANG, M. KARLSSON, AND N. R. ZIEBARTH (2020): “Long-term health insurance: Theory meets evidence,” National Bureau of Economic Research.
- ATTANASIO, O. AND L. PISTAFERRI (2014): “Consumption inequality over the last half century: some evidence using the new PSID consumption measure,” *American Economic Review*, 104, 122–26.
- BHARGAVA, S. AND D. MANOLI (2015): “Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment,” *American Economic Review*, 105, 3489–3529.
- CABRAL, M. (2017): “Claim Timing and Ex Post Adverse Selection,” *The Review of Economic Studies*, 84, 1–44.
- CURRIE, J. (2006): “The Take-up of Social Benefits,” in *Poverty, the Distribution of Income, and Public Policy*, ed. by Alan Auerbach, David Card, and John Quigly, Russel Sage, chap. 3.
- DAFNY, L., J. GRUBER, AND C. ODY (2015): “More Insurers Lower Premiums: Evidence from Initial Pricing in the Health Insurance Marketplaces,” *American Journal of Health Economics*, 1, 53–81.
- DICKSTEIN, M., M. DUGGAN, J. ORSINI, AND P. TEBALDI (2015): “The Impact of Market Size and Composition on Health Insurance Premiums: Evidence from the First Year of the Affordable Care Act,” *American Economic Review*, 105, 120–125.
- DORN, S., B. GARRETT, AND M. EPSTEIN (2018): “New Risk-Adjustment Policies Reduce But Do Not Eliminate Special Enrollment Period Underpayment,” *Health Affairs*, 37, 308–315.
- EINAV, L., A. FINKELSTEIN, AND M. CULLEN (2010): “Estimating Welfare in Insurance Markets using Variation in Prices,” *The Quarterly Journal of Economics*, 125, 877–921.
- EINAV, L., A. FINKELSTEIN, S. P. RYAN, P. SCHRIMPF, AND M. R. CULLEN (2013): “Selection on Moral Hazard in Health Insurance,” *American Economic Review*, 103, 178–219.
- EINAV, L., A. FINKELSTEIN, AND P. TEBALDI (2019): “Market Design in Regulated Health Insurance Markets: Risk Adjustment vs. Subsidies,” Stanford University Working Paper.
- ERICSON, K. M. M., K. H. GEISLER, AND B. LUBIN (2018): “The impact of partial-year enrollment on the accuracy of risk-adjustment systems: A framework and evidence,” *American Journal of Health Economics*, 4, 454–478.

- ERICSON, K. M. M. AND A. STARC (2015): “Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange,” *The Review of Economics and Statistics*, 97, 667–682.
- FANG, H., M. KEANE, AND D. SILVERMAN (2008): “Sources of Advantageous Selection: Evidence from the Medigap Insurance Market,” *Journal of Political Economy*, 116, 303–350.
- FINKELSTEIN, A., N. HENDREN, AND M. SHEPARD (2019): “Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts,” *American Economic Review*, 109, 1530–1567.
- FINKELSTEIN, A. AND K. MCGARRY (2006): “Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market,” *American Economic Review*, 96, 938–958.
- GERUSO, M., T. LAYTON, AND D. PRINZ (2019): “Screening in Contract Design: Evidence from the ACA Health Insurance Exchanges,” *American Economic Journal: Economic Policy*, 11, 64–107.
- GHILI, S., B. HANDEL, I. HENDEL, AND M. D. WHINSTON (2020): “Optimal Long-Term Health Insurance Contracts: Characterization, Computation, and Welfare Effects,” Cowles Foundation Discussion Paper.
- HACKMANN, M. B., J. T. KOLSTAD, AND A. E. KOWALSKI (2015): “Adverse Selection and an Individual Mandate: When Theory Meets Practice,” *American Economic Review*, 105, 1030–66.
- HANDEL, B., I. HENDEL, AND M. D. WHINSTON (2015): “Equilibria in health exchanges: Adverse selection versus reclassification risk,” *Econometrica*, 83, 1261–1313.
- HENDEL, I. AND A. LIZZERI (2003): “The Role of Commitment in Dynamic Contracts: Evidence from Life Insurance,” *The Quarterly Journal of Economics*, 118, 299–328.
- HERRING, B. AND M. V. PAULY (2006): “Incentive-compatible guaranteed renewable health insurance premiums,” *Journal of Health Economics*, 25, 395–417.
- KLEVEN, H. AND W. KOPCZUK (2011): “Transfer Program Complexity and the Take Up of Social Benefit,” *American Economic Journal: Economic Policy*, 3, 55–65.
- KOWALSKI, A. E. (2014): “The Early Impact of the Affordable Care Act, State by State,” *Brookings Papers on Economic Activity*, 277–333.
- PANHANS, M. (2019): “Adverse Selection in ACA Exchange Markets: Evidence from Colorado,” *American Economic Journal: Applied Economics*, 11, 1–36.
- PATIENT PROTECTION AND AFFORDABLE CARE ACT (2010): Pub. L. No. 111-148, 124 *Stat.* 119.
- SALTZMAN, E. (2019): “Demand for health insurance: Evidence from the California and Washington ACA exchanges,” *Journal of Health Economics*, 63, 197 – 222.

- SHEPARD, M. (2016): “Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange,” *NBER Working Paper No. 22600*.
- SIMON, K., A. SONI, AND J. CAWLEY (2017): “The Impact of Health Insurance on Preventive Care and Health Behaviors: Evidence from the First Two Years of the ACA Medicaid Expansions,” *Journal of Policy Analysis and Management*, 36, 390–417.
- TEBALDI, P. (2017): “Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA,” Becker Friedman Institute for Research in Economics Working Paper No. 2017-05.
- THE KAISER FAMILY FOUNDATION (2012): “Health Insurance Market Reforms: Rate Restrictions,” .
- (2016): “State Health Facts: Health Insurance Coverage of the Total Population,” .


Figure 1: Transactions Data vs. Public Data




Notes: This figure benchmarks bank transaction data against public data sources. **Panel (a)** compares market shares of individuals by insurance carrier between the Covered California marketplace (CCM) data and the transactions data in 2014. To make the CCM data comparable, we construct weights to scale up all unsubsidized individuals in the CCM data (See Appendix D). **Panel (b)** compares the household-level income distribution in our data to ACS data for the same time period. **Panel (c)** regresses the log of 2013 median post-tax household income from our data against the log of 2011-2015 median post-tax household income from ACS data. Observations are at the county level. **Panel (d)** compares average household annual expenditure in 2014 among three data sources: CEX, NIPA, and our data.

Figure 2: Participation in California's Individual Insurance Market, 2014-2015


(a) Distribution of Dropouts Across Months



(b) Continued Enrollment by Income Group




(c) Six Month Drop-out Rates by Enrollment Period




Notes: This figure describes participation in California's individual insurance market in 2014-2015. In **Panel (a)**, we plot the distribution of enrollment months in the financial transactions data and in weighted Covered California marketplace (CCM) data for the 2014 open enrollment period. In **Panel (b)** we use the transactions data to plot the net share of households who continue to pay insurance premiums in each month, conditional on participating in the 2014 open enrollment period. We plot this share by bin of 2013 annual post-tax income. In **Panel (c)** we compare the six month drop-out rate by enrollment period using both the transactions data and CCM data. We define "six-month drop-out" as the share of enrollees exiting within 6 months after sign-up, excluding dropouts in November/December of the enrollment year. This procedure ensures off-cycle enrollees who sign up later in the year are not labeled dropouts. In all panels, we use a county weight to match the California-wide population distribution. We weight the CCM data in panels (a) and (c) to make it comparable to the transactions data, which includes both on and off-marketplace enrollees (See Appendix D). We exclude enrollees who have coverage for one month or less.

Figure 3: Income and Expenditure Dynamics Around Sign-up and Drop-out




Notes: This figure compares the dynamics of monthly income and various types of monthly dollar spending for households who drop coverage. We examine these dynamics separately for households enrolling in the pre- vs. post-ACA time periods. Other spending, reported in **Panel (b)**, includes all transactions less health premium payments, drug spending, and health spending. In all panels, the x-axis measures months around coverage, where negative values indicate months before enrollment and positive values indicate months since drop-out. All months of coverage are collapsed into the event time period 0. The pre-ACA period includes households who both sign up and drop out prior to July 2013; the post-ACA period includes households whose sign-up begins as of the 2014 ACA open enrollment period. For both periods, we further restrict our sample to households who: (i) satisfy our drop-out definition in Section 4 (ii) have at least 10 months of data prior to sign-up and at least 10 months after drop-out; (iii) have pre-period annual income <100,000 dollars; and (iv) have insurance coverage for more than one month. All outcome variables have been top-coded at the 99th percentile value within each income group before being included in the regressions. Units are measured as a percentage change relative to the average monthly value during coverage. All regressions use a county weight to account for sampling differences across counties and absorb household fixed effects. 95% Confidence intervals based on robust standard errors are reported as shaded areas.

Figure 4: Demand and Average Cost in a Market with Dropouts




Notes: This figure illustrates the demand curve and member-month average cost curve in a market with drop-out types present but no traditional adverse selection (i.e. annual household health costs are equal across drop-out and full-year types). The equilibrium monthly premium occurs where the demand curve intersects the average cost curve. The plots reflect a numerical example in which $\phi/12 = 1/2$, $N_D = N_{ND} = 100$, $c_D = c_{ND} = 40$, and $G_D(\cdot) = G_{ND}(\cdot) = Unif([0, 100])$.

Figure 5: Consumer Surplus by Drop-out Type, with Differential Health Costs



Notes: This figure illustrates the welfare consequences of drop-out in the presence of traditional adverse selection. In panel (1) we assume a small difference in annual health care costs by type: dropouts have cost $C_D = 20$ while non-dropouts have a cost $C_{ND} = 40$. In panel (2), we assume a larger difference in annual costs by type: $C_D = 20$ while non-dropouts have a cost $C_{ND} = 49$. In both panels, welfare in subfigures (a) shows that modest levels of drop-out increase social welfare, with the social planner's optimum achieved when $\phi/12 = c_D/c_{ND}$. Subfigures (b) show the monthly premium as a function of the fraction of the year drop-out types pay, $\phi/12$. Subfigures (c) illustrate the consequences of drop-out on consumer surplus, separately by household type. In this particular numerical example we set $N_D = N_{ND} = 100$, $G_D(\cdot) = \text{Unif}([0, 50])$, and $G_{ND}(\cdot) = \text{Unif}([50, 100])$.


Figure 6: Health Consumption by Length of Coverage




Notes: This figure compares a count of health transactions during the period of health insurance coverage across households by length of coverage. We provide an analysis for two periods: the pre-ACA period includes households who sign up and drop out before December 2013 and the post-ACA period includes households who sign up during the 2014 open enrollment period. Regressions use cost data from the period of health insurance coverage. We restrict to households who (i) have data at least 10 months before sign-up and at least 10 months after drop-out; and (ii) have a coverage period of at least two months. We recover these estimates from a regression of the count of monthly health transactions on indicators for coverage length (2-4 months, 5-8 months, 9-12 months, and 13+ months). We include monthly income and lagged monthly income as controls the regressions. Quantities of monthly healthcare usage are reported relative to the corresponding months of the 13+ months coverage group—for example, we divide the count of monthly health transactions of the 5-8 month group by the count of transactions during the first 8 months of the 13+ month group. We use county weights to account for sampling differences across counties in our data. The 95% confidence intervals shown use robust standard errors clustered by household.

Figure 7: Consumer Surplus under a Drop-out Penalty, by Enrollment Type


(a) 2-4 Months Drop-out Type



(b) 5-8 Months Drop-out Type



(c) Full-year Type



Notes: In this figure, we report the normalized consumer surplus under alternative drop-out penalties. The surplus level is normalized to one for the setting without a drop-out penalty—that is, for the setting in which zero months of enrollment are required. Each subplot represents the normalized surplus within each enrollment type. The dashed vertical line represents the welfare maximizing number of months of required enrollment when considering the joint surplus of all enrollee types. The solid vertical line represents the number of required enrollment months that would maximize the full-year enrollee type’s consumer surplus alone. We conduct the calibration using the model in Section 5 and the parameters of cost and demand reported in Table 5 for the years 2014-2015.

Table 1: Sample Summary Statistics

Panel A. Household-Level Variables						
Variable	Insurance Group ($N = 106,904$)		Reference Group ($N = 139,000$)			
	Mean	S.D.	Mean	S.D.		
Lives in the same California county every year	0.81	0.39	0.82	0.39		
Has a credit card account	0.58	0.49	0.48	0.50		
Number of months covered by health insurance	17.34	17.02	-	-		
Share of “drop-out” households	0.31	0.46	-	-		

Panel B. Household-Monthly-Level Variables						
Variable	Insurance Group ($N = 4,686,107$)			Reference Group ($N = 5,866,040$)		
	Mean	Median	S.D.	Mean	Median	S.D.
Monthly number of transactions						
- Health out-of-pocket	0.31	0.00	0.97	0.21	0.00	0.82
- Other non-health	39.00	29.00	34.55	37.02	27.00	34.14
Monthly transaction dollar amount						
- Income	6,143.93	3,861.40	17,519.22	5,034.01	3,165.88	13,278.70
- Premium payments	78.01	0.00	267.29	0.00	0.00	0.00
- Health out-of-pocket spending	22.69	0.00	144.78	11.65	0.00	101.46
- Other non-health spending	6,268.94	3,875.45	19,637.03	4,962.89	3,077.81	12,577.19
- Netflix and Hulu spending	2.67	0.00	15.36	2.29	0.00	24.37
- Utilities spending	65.93	0.00	460.54	49.72	0.00	143.09
- Auto insurance spending	57.18	0.00	243.30	46.24	0.00	210.13
Medicaid eligibility probability	0.21	0.00	0.35	0.31	0.01	0.40

Notes: This table reports summary statistics in our sample of households who purchase individual insurance between January 2011 and December 2015, along with a reference group that represents a random sample of households who we never observe paying premiums. For the monthly transaction counts and dollar amounts, we report statistics for consumers with incomes less than \$200,000. Households below this income cap form our sample for the main analyses on spending. We define “drop-out” households in Section 4.2: households who pay more than one but less than nine months of premiums in a calendar year, excluding households exiting in November or December. Health out-of-pocket spending excludes spending on insurance premiums and at pharmacies. Other non-health spending excludes spending on health, pharmacy, and insurance premiums.

Table 2: Health Consumption for 2014 Open Enrollees

Panel A. Change in Health Spending					
Annual Income:	(1) ≤20K	(2) 20K-40K	(3) 40K-60K	(4) 60K-100K	(5) 100K-200K
Treat × Post ACA	42.196*** (4.694)	35.301*** (5.092)	35.961*** (8.619)	48.769*** (11.352)	60.616*** (20.006)
Post ACA	8.796*** (0.722)	12.583*** (1.230)	14.177*** (1.996)	18.018*** (2.511)	26.148*** (3.911)
Pre Period Treatment Mean	52.94	90.79	137.67	191.72	266.13
% Change	0.797*** (0.089)	0.389*** (0.056)	0.261*** (0.063)	0.254*** (0.059)	0.228*** (0.075)
Number of individuals in treatment	2,108	2,965	2,172	2,229	1,735
Number of individuals in reference	37,712	28,821	19,944	23,212	19,144
Panel B. Change in Health Transactions					
Annual Income:	(1) ≤20K	(2) 20K-40K	(3) 40K-60K	(4) 60K-100K	(5) 100K-200K
Treat × Post ACA	0.834*** (0.076)	0.773*** (0.083)	0.640*** (0.121)	0.598*** (0.134)	0.637*** (0.163)
Post ACA	0.180*** (0.013)	0.158*** (0.022)	0.107*** (0.032)	0.121*** (0.034)	0.162*** (0.040)
Pre Period Treatment Mean	0.96	1.95	2.73	3.46	3.93
% Change	0.869*** (0.079)	0.396*** (0.043)	0.235*** (0.044)	0.173*** (0.039)	0.162*** (0.041)
Number of individuals in treatment	2,108	2,965	2,172	2,229	1,735
Number of individuals in reference	37,712	28,821	19,944	23,212	19,144

Notes: This table examines the change in health transactions and spending among 2014 open enrollees after sign-up. We compare these transaction and spending levels to the same levels in our reference group in 2014. Observations are at the household-year level. “Treat” is an indicator set equal to 1 if a household enrolls during the 2014 open enrollment period and 0 if the household belongs to the reference group. “Post ACA” equals 1 in 2014-2015 and 0 in 2011-2013. We run a diff-in-diff regression separately by income group, where income is defined as the household’s 2013 annual post-tax income. We top-code health transactions and health spending at the 99th percentile value within each income group. In all regressions, we include household fixed effects and apply a county weight to account for sampling difference across counties in our data. We restrict our individual market sample to households who have coverage for at least two months. Robust standard errors are clustered by household. * p<0.1, ** p<0.05, *** p<0.01.

Table 3: Changes in Health Consumption around Sign-up and Drop-out

Panel A: Change in Health Spending: Pre-ACA Sign-up/Drop-out					
Annual Income:	(1) ≤20K	(2) 20K-40K	(3) 40K-60K	(4) 60K-100K	(5) 100K-200K
Sign-up (% change)	0.307 (0.190)	0.115 (0.102)	0.175 (0.133)	0.164 (0.124)	0.159 (0.144)
Drop-out (% change)	-0.292 (0.210)	0.072 (0.142)	-0.158 (0.142)	-0.029 (0.138)	-0.071 (0.119)
Pre Sign-up Mean Health Spending	14.59	16.84	14.11	22.14	26.69
Number of Observations	5,006	7,458	5,981	7,452	6,599
Panel B: Change in Health Spending: Post-ACA Sign-up/Drop-out					
Annual Income:	(1) ≤20K	(2) 20K-40K	(3) 40K-60K	(4) 60K-100K	(5) 100K-200K
Sign-up (% change)	0.390*** (0.104)	0.230*** (0.075)	0.142 (0.099)	0.091 (0.080)	-0.014 (0.070)
Drop-out (% change)	-0.386*** (0.097)	-0.220*** (0.081)	-0.181* (0.094)	-0.199*** (0.065)	-0.116* (0.069)
Pre Sign-up Mean Health Spending	9.11	9.94	15.32	27.61	35.94
Number of Observations	14,079	27,352	18,991	20,241	15,137
Panel C: Change in Health Transactions: Pre-ACA Sign-up/Drop-out					
Annual Income:	(1) ≤20K	(2) 20K-40K	(3) 40K-60K	(4) 60K-100K	(5) 100K-200K
Sign-up (% change)	0.587*** (0.204)	0.099 (0.099)	0.181 (0.114)	0.087 (0.104)	0.224** (0.093)
Drop-out (% change)	-0.338* (0.172)	0.075 (0.118)	-0.245** (0.109)	0.061 (0.099)	-0.140 (0.094)
Pre Sign-up Mean Health Transactions	0.17	0.25	0.27	0.35	0.33
Number of Observations	5,006	7,458	5,981	7,452	6,599
Panel D: Change in Health Transactions: Post-ACA Sign-up/Drop-out					
Annual Income:	(1) ≤20K	(2) 20K-40K	(3) 40K-60K	(4) 60K-100K	(5) 100K-200K
Sign-up (% change)	0.604*** (0.122)	0.232*** (0.067)	0.162** (0.065)	0.047 (0.053)	0.061 (0.060)
Drop-out (% change)	-0.402*** (0.109)	-0.257*** (0.067)	-0.276*** (0.064)	-0.157*** (0.050)	-0.156*** (0.058)
Pre Sign-up Mean Health Transactions	0.11	0.19	0.35	0.43	0.41
Number of Observations	14,079	27,352	18,991	20,241	15,137

Notes: This table examines changes in healthcare consumption of drop-out households after sign-up and after drop-out during the pre-ACA period vs. post-ACA period. Observations are at the household-month level. “Sign-up” indicates the period after sign-up. “Drop-out” indicates the period after drop-out. The pre-ACA period includes households who both sign up and drop out prior to July 2013, while the post-ACA period includes households who sign up in the 2014 ACA open enrollment period and drop out after January 2014. We run regressions separately by income group, where income is defined as the household’s 2013 annual post-tax income. For both periods, we restrict our sample to “dropouts”: households who pay more than one but less than nine months of premiums in a calendar year, excluding households exiting in November or December. We also restrict our sample to include only households who appear in the transactions data at least 10 months before sign-up and at least 10 months after drop-out. We top-code health transactions and health spending at the 99th percentile value within each income group. All regressions control for household fixed effects, monthly income, average lagged monthly income from the past three months, and use county weight to account for sampling differences across counties in our data. Units are measured as a percentage change, relative to the average healthcare consumption amount in the 10 months leading up to sign-up. Robust standard errors are clustered by household. * p<0.1, ** p<0.05, *** p<0.01.

Table 4: Urgent/Essential Health Spending for Continuing Enrollees vs. Dropouts

Dependent Variable:	% Very Essential (1)	% Very Urgent (2)
Panel A. Pre ACA Period		
Drop-out consumers	-2.042* (1.112)	-1.173 (1.343)
Enrolled in insurance	-0.913* (0.512)	-0.746 (0.527)
Drop-out consumers $\times$ Enrolled in insurance	4.085** (2.026)	2.158 (1.547)
Enrolled in insurance + Drop-out consumers $\times$ Enrolled in insurance	3.172 (1.960)	1.413 (1.454)
Mean	6.4	6.4
Std. dev.	20.1	20.5
Observations	5,682	5,682
Panel B. Post ACA Period		
Drop-out consumers	0.065 (0.888)	-1.097 (0.830)
Enrolled in insurance	-0.307 (0.443)	-0.082 (0.453)
Drop-out consumers $\times$ Enrolled in insurance	-1.562* (0.919)	-1.676* (0.895)
Enrolled in insurance + Drop-out consumers $\times$ Enrolled in insurance	-1.869** (0.803)	-1.758** (0.771)
Mean	5.4	5.2
Std. dev.	18.2	18.2
Observations	7,298	7,298

Notes: This table reports essential and urgent healthcare purchasing patterns of drop-out consumers relative to non-dropouts. Panel A includes pre-ACA households who sign up prior to December 2013; Panel B includes post-ACA households who sign up from December 2013 onward. In both panels, we restrict to households who: (i) appear in the transactions data at least 10 months prior to sign-up and at least 10 months following drop-out; (ii) have coverage for more than 1 month; (iii) have positive income; and, (iv) has non-missing health transactions containing any of 134 words rated as health relevant (See Appendix F). We collapse all months of coverage into an “enrolled in insurance” period labeled month 0. “Drop-out consumers” equals 1 for households who pay more than one but less than nine months of premiums in a calendar year, excluding households exiting in November or December, and equals 0 otherwise. The two outcome variables, “% Very Essential” and “% Very Urgent”, measure the percentage of total health transactions that are considered very essential and very urgent, respectively, where we define “very” as a word rating at least in the 95th percentile among the set of health-relevant words. All regressions control for log income and household enrollment status interacted with income group fixed effects. We use the county weight to account for sampling difference across counties in our data. Robust standard errors are clustered by household. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Parameter Estimates for Model Calibration

Dependent variable:	Market shares by carrier		
Coverage months	2-4	5-8	9+
Share of total enrollees	0.182	0.192	0.626
Cost ratios (base = 9+)	1.221	1.077	1.000
	(1)	(2)	(3)
Annual premium (\$1,000)	-0.145*** (0.046)	-0.294** (0.129)	-0.284*** (0.091)
Constant	0.830*** (0.166)	1.374*** (0.469)	1.362*** (0.332)
Full set of controls	Yes	Yes	Yes
R^2	0.934	0.883	0.896
N (year-region-carrier)	114	114	114

Notes: This table presents estimation of parameters we use in our model calibration. To calculate market shares, we first restrict our transactions data to 2014-2015 open enrollees and then classify them into three drop-out groups based on length of coverage: 2-4, 5-8, and 9+ months. Within each year $\times$ region, we then calculate a share of enrollees that is made up of each carrier. To calculate annual premium, we use data on gross premiums for 40-year-old individuals. We calculate average premium at the year $\times$ region $\times$ carrier level by averaging premiums across all metal tiers, using weights based on the Covered California marketplace (CCM) enrollment. We adjust the CCM weights to account for both on- and off-marketplace enrollment. Annual premiums are measured in thousands of dollars. All regressions employ a level-level specification and control for region, year, carrier, region $\times$ year, and region $\times$ carrier fixed effects. Observations are weighted by the number of enrollees within each year $\times$ region $\times$ carrier. The cost ratio amounts are based on estimates from a regression of the monthly count of health transactions among 2014-2015 open enrollees (normalized relative to the 9+ group mean) on drop-out group indicators, evaluated at mean monthly income and lagged monthly income. We also report the share of total enrollees that compose each drop-out group in the table. Robust standard errors are reported in parenthesis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Online Appendix for “Insurance without Commitment: Evidence from the ACA Marketplaces”

Rebecca Diamond, Michael J. Dickstein, Timothy McQuade, Petra Persson

March 31, 2021

Table of Contents

A	Sample construction	1
B	Classifying health and drug spending	3
C	Validating health care consumption via MEPS	4
D	Validating participation behavior via Covered California enrollment	6
E	Determining Medi-Cal eligibility	7
F	Classifying transactions by urgency, essentialness	9
G	Additional health and drug consumption analyses	9
H	Appendix figures and tables	11

A Sample construction

We form the main analysis sample using account-linked transactions from our data vendor. In the vendor’s data, for each unique user, we observe an identification variable that links the user to her bank account and any credit cards from the same bank. Within the accounts, we observe all account transactions, with a post date, transaction date, brief transaction description, dollar amount, credit/debit indicator, merchant name, transaction category, merchant street address, city, county, state, and zip code. See Appendix Table A2 for some example transactions. A user must have a bank account to enter our sample but need not have a card account.

Our data cover the period from January 2011 through December 2015. In this time frame, there are 9,223,300 unique users/households. To form our sample, we first identify the locations of all physical (non-online) merchants and assign a county to each transaction. We identify the location of each household using the modal county of the set of merchants with which the household transacts most frequently within a given year. Because we aim to collect all “California non-movers”, we look for households who have a modal county in California for every year in which we observe their data. In other words, we allow moves between California counties, but do not allow moves to or from other states. We identify 1,141,592 California households, accounting for 12.4% of all users in our data.

For each household, we use its 2013 modal county as the modal county across all years. Whenever 2013 data are missing, we take the mode of the household’s modal counties from other years as the modal county across all years. Since the distribution of households across counties in our data is driven by bank prevalence and differs from the population distribution observed in the National Historical Geographic Information System (NHGIS), we create a weight to scale up households in under-represented counties and to scale down households in over-represented counties in our data. For each county, we compute the ratio of the number of households in NHGIS to the corresponding number of households in our transaction data, yielding the county weight. We use this weight for all of our analyses.

With our California population, we construct a household-monthly panel, where each household’s month ranges from the month we first observe a transaction in its bank or card account to the month of the last transaction we observe. From the raw transactions, we count the number of transactions and sum the dollar amount for various variables at the monthly level. These variables include credits or debits from the bank and card account; bank credits net of transfers; debit taxes from both bank and card accounts; bank credits related to income; health insurance premiums (determined by our premium-related keywords for carriers participating in California’s ACA marketplace¹); and specific types of periodic transactions such as payments for auto insurance, utilities, and streaming services (Netflix and Hulu).² In addition, we classify all debit transactions in our California sample into drug, health, child-related, and other transactions using machine learning algorithms discussed in Appendix Section B. We also impute a household’s probability of

¹These carriers include: Anthem Blue Cross of California, Blue Shield of California, Chinese Community Health Plan, Contra Costa Health Plan, Health Net, Kaiser, L.A. Care Health Plan, Molina, Sharp Health Plan, and Western Health Advantage.

²We examine these particular spending categories in Figure 3. Netflix/Hulu spending consists of all debit transactions that contain the words “netflix” or “hulu”. Utilities spending includes all debit transactions that belong to the “Utilities” transaction category classified by our data vendor. We identify auto insurance spending by picking ten popular auto insurance providers and filtering through all debit transactions that contain any of these providers’ names: Allstate, American Family, Farmers, GEICO, Liberty Mutual, Nationwide, Progressive, State Farm, Travelers, and USAA.

medicaid eligibility, as discussed in Appendix Section E. Finally, with these household-monthly panels, we assign 0 to missing or negative values and top-code monthly health and drug spending at 100 transactions and 10,000 dollars per month.³

With this household panel, we define our variables of interest. First, for households who pay individual health insurance premiums in the sample, we define “sign-up” as the month we first observe the household’s premium payment and “drop-out” as the month we last observe a premium payment. We use this sign-up variable to sort households into six enrollment statuses: a reference group; pre-ACA enrollees; 2014 open enrollees; 2014 off-cycle enrollees; 2015 open enrollees; and 2015 off-cycle enrollees. The reference group includes households we never observe paying individual health insurance premiums. Pre-ACA enrollees are households who sign up for individual private health insurance prior to December 2013. Open enrollment includes new sign-ups during California’s open enrollment period in a given year. In 2014, for example, we label open enrollees as those consumers who newly sign up from December 2013 through March 2014. Off-cycle enrollment includes sign-ups outside the open enrollment period in a year. We also flag and exclude individuals who sign up for a coverage period of one month or less. In conversations with staff and actuaries at Covered California, we learned such short enrollment periods may reflect incomplete or mistaken enrollments that were subsequently refunded rather than true enrollment. To avoid this measurement error, we exclude these households from our analyses.

We define “active months” as the number of months between sign-up and drop-out (inclusive), or equivalently, the length of the insurance coverage period. We create a drop-out type indicator that takes a value of 1 if a household enrolls for less than nine months, excluding drop-outs in November and December. We apply the November/December exclusion because it remains ambiguous in those months whether a new enrollee dropped out or signed up late under off-cycle eligibility. As some of our analyses involve changes in spending around sign-up and drop-out, we create post-signup and post-dropout indicators for months at or after sign-up and for months after drop-out, respectively. Other analyses deal with changes in spending relative to the enrollment period, so we create event time indicators, where all enrollment months are collapsed into the event time period of 0; negative integers indicate months before sign-up and positive integers indicate months after drop-out. We also flag households that satisfy a ‘balanced panel’ requirement—i.e. those who have data at least 10 months before sign-up and at least 10 months after drop-out.

We build post-tax income by spreading annual debit taxes proportionally to monthly pre-tax post-transfer bank credits across all months within a given year. We take the difference between the two measures, yielding monthly post-tax post-transfer income. Then, we classify households into six income categories based on their 2013 annual post-tax income: 0-20K, 20K-40K, 40K-60K, 60K-100K, 100K-200K, and 200K and above. We create another version of the income group classification that is based on average monthly income from 10 months leading up to the initial sign-up month; this indicator has the same six income brackets. Next, we build an indicator for having children based on our classified child-related transactions. For each county, we label the 10% of households with the highest average monthly number of child-related transactions as “having children” and the remaining 90% as “having no children”.

Finally, we form our analysis sample by excluding observations that do not meet our inclusion criteria. We list these criteria in Appendix Table A3, along with an accounting for the observations excluded sequentially under each criterion. These sequential drops result in 798,085 households in our final sample, accounting for

³These thresholds lie above the 99th percentile for transactions and spending.

69.9% of all California non-mover households. Because the reference group is large and because we only use this group in a handful of analyses, we take a 20% random sample of this group to minimize computational cost. We use this sample in almost all analyses throughout this paper.

B Classifying health and drug spending

We describe in detail our classification of transactions into the categories: drug, health, child-related, and other. First, we collect the universe of transactions of the 1,141,592 California non-mover households we have identified in Section A spanning January 2011 to December 2015. Over all users, we save transaction details, including the post date, transaction date, brief transaction description, merchant name, transaction category, and dollar amount.⁴ Because the transaction description string can contain typographical inconsistencies, we clean it by (i) converting to lowercase letters; (ii) removing numbers and symbols (e.g. operations, punctuation, parentheses, brackets, braces, quotations) so that there are only letters left; (iii) removing repeated de-identifier x’s, stand-alone letters (a-z), and repeated white spaces. We clean merchant names in a similar manner. Because Kaiser-related transactions are complex, we remove them from the pool of transactions here and analyze them separately. We then split the cleaned description string into individual words by parsing white spaces.

We use a random sample of 10 million cleaned transactions as our training data and classify them into four categories: drug, health, child, and other, defined as follows. For the drug category, we take all transactions that belong to the Healthcare/Medical category in the vendor data and contain drugstore-related words such as “walgreens”, “rite aid”, “drug”, “pharmacy”, and “prescription”. For the health category, we begin with all transactions assigned to the Healthcare/Medical category in the vendor data and then remove transactions that are unlikely to be covered by private health insurance. More precisely, for health spending, we remove words classified as drug, as well as veterinary and animal-related services; lifestyle activities such as fitness, yoga, and spas; insurance premiums that contain health-related words; and dental and vision transactions. For the child-related category, we take transactions associated with popular merchants in the “Child/Dependent Expenses”, “Education”, and “Entertainment” categories in our data. We assign the remaining transactions to the other category.

Next, we construct a frequency table of words that appear in transaction descriptions. This table shows how frequently each word appears in each category. We also save the total number of transactions within each category. We remove words that are not useful for classifying transactions. The list includes common English words, 26 English alphabet letters, bank and card transaction-related words, state names and abbreviations, date and time, location-related words, all California cities and counties, couplets of stand-alone English letters (except md and dr), and different variations of association/center/foundation/group/institution. To ease computation, we also remove words that appear less than five times in all four categories.

With this list, we apply a modified version of Laplacian correction: for each category, we compute the mean count across all words and then add 0.01% of that mean to the original count within that category. This step adds a small positive value so that the count is never zero. In this way, the posterior probability of a transaction string being in a category will not equal zero simply because one word within the transaction has a zero mean count.

⁴We use post date for most transactions because it is highly populated. Whenever post date is missing, we use transaction date.

We construct a probability look-up table that contains three probabilities: $Prob(category)$ or fraction of transactions that are in *category*; $Prob(w)$ or probability that word w appears; and $Prob(w|category)$ or probability that word w appears within *category*. Next, we log-transform these probabilities and merge them back into our cleaned transactions in the entire dataset. We proceed iteratively on each individual word split out from the transaction description. This process allows us to calculate the probability that transaction t is classified into a particular *category* given its description string $w(t) = w_1w_2...w_{n_t}$, where w_i denotes the i th word in $w(t)$ for $i = 1, 2, \dots, n_t$ and n_t denotes the length or the number of words in $w(t)$. By assuming conditional independence of probabilities, we have

$$\begin{aligned} Prob(category|w(t)) &= Prob(category|w_1, w_2, \dots, w_{n_t}) \\ &= \frac{Prob(w_1w_2...w_{n_t}|category) \times Prob(category)}{Prob(w_1w_2...w_{n_t})} \\ &= \frac{\prod_{i=1}^{n_t} Prob(w_i|category) \times Prob(category)}{\prod_{i=1}^{n_t} Prob(w_i)} \\ \log Prob(category|w(t)) &= \sum_{i=1}^{n_t} \log Prob(w_i|category) - \sum_{i=1}^{n_t} \log Prob(w_i) + \log Prob(category) \end{aligned}$$

Lastly, we classify transactions into drug, health, child-related, and other transactions. We classify Kaiser transactions into drugs, health, and premiums, using specific keywords. For non-Kaiser transactions, we assign each transaction a category for which the log likelihood is the highest among the four categories.⁵

C Validating health care consumption via MEPS

We use our unique credit and bank account data to measure health spending among enrollees in individual insurance under the ACA. We observe a measure of out-of-pocket health care spending for each household, independent of insurance enrollment. We do not, however, observe a household’s total health care spending, including costs covered by the insurer. Our focus on out-of-pocket spending raises two concerns. First, new enrollees under the ACA may have consumed charity care or had full coverage under Medicaid in prior years; for such households, our data would find no health transactions or out-of-pocket spending. Second, if households are uninsured prior to the ACA and pay the full cost of medical bills during that period, out-of-pocket spending in periods prior to ACA enrollment may appear higher than during enrollment because observed spending upon enrollment does not include insurer costs. The enrollee’s actual consumption may have in fact increased when accounting for the total cost of care. Both of these concerns could create problems when mapping out-of-pocket spending to health care consumption pre- versus post-enrollment.

We examine both issues using the Medical Expenditure Panel Survey (MEPS). This survey collects information on health care and health insurance coverage for households over a two-year period. Importantly, the survey data includes a variable indicating households who purchased individual health insurance coverage through an ACA marketplace. MEPS also contains data on health care consumption, regardless of who paid

⁵Kaiser drug transactions contain either: (i) the word “kaiser” and one of the following words: “phar”, “cpp”, “mail”, “rx”, “downey”, and “livermore” or (ii) the word “kp” and one of these words: “mailorder”, “rx”, and “drug”. Kaiser health premium transactions contain the word “kaiser” and one of these words: “due”, “health pl”, “hps”, “direct pay”, “bill pay”, and “online pay”. Kaiser health transactions contain either (i) the word “kaiser” and remaining keywords not in the premium or drug categories (e.g. “perm”, “prmnt”, and “prnte”) or (ii) the words “kp” and “medical”.

for the health care, including charity care. We can thus compare the spending of those individuals reporting marketplace private coverage in MEPS to spending in our transactions data. A key downside of the MEPS data, however, is power: in 2014, only 201 individuals in the MEPS data, nationwide, report purchasing ACA marketplace coverage.

First, in MEPS, we examine the types of health insurance 2014 ACA marketplace enrollees had in 2013. Appendix Figure A6(d) shows that the average 2014 ACA enrollee had no health insurance for 5.5 months, employer-sponsored health insurance for 3.75 months, individual insurance for 1 month, and public insurance coverage (Medicaid, etc) for 1.75 months in 2013. Overall, the average household has some form of health insurance coverage for medical expenses for more than half of the year, suggesting that pervasive consumption of charity care would likely be low.

Next, we directly compare the levels of annual out-of-pocket (OOP) spending on health and drug care from 2012 through 2015 in MEPS data and our transactions data. In Table A5 we see the average spending on OOP health costs in MEPS and our transactions data are quite similar each year, at around \$205 in 2012, rising to between \$250 and \$275 in 2015.⁶ Our drug spending levels are roughly twice as large as MEPS reported levels. Due to limitations in our transaction data, we define drug spending to include all spending at drug stores, including non-prescription purchases. The reported drug spending in MEPS includes only prescription drug purchases. For this reason, in our main analysis we focus only on health spending; we leave our drug spending analysis for Appendix Section G.

In addition to the mean OOP drug and health spending, we can also compare the distribution of health spending between our transactions data and MEPS data. Panel B of Appendix Table A5 compares the two datasets by year at various quantiles of the spending distribution. The quantiles track very closely between the two datasets, including for the highest quantiles of spending.

Finally, we use MEPS data to validate our reported changes in health care consumption among new ACA enrollees. We analyze the within-household change in health consumption between 2013 and 2014 for households enrolling in coverage during the 2014 open enrollment period. In addition to OOP spending and transactions counts, MEPS data also reports total health charges, which reflect the total "retail cost" of health care consumed.⁷ We do not observe a similar measure in our transactions data, and so impute charges in our data using the relationship between OOP costs and total charges we observe in MEPS data. Appendix Table A7 reports the output of the regression we run to impute charges from OOP costs.⁸

The results of our comparison of within-individual changes in health care consumption appear in Appendix Table A8. We first consider transaction counts, as these counts are reported in MEPS in a format comparable to the count of health care transactions we observe in our transactions data. Panel B of Appendix Table

⁶In detail, MEPS data come in two consecutive years. For each given year, we use two panels to calculate summary statistics. For example, the 2012 statistics above are based on 2012 data from both the 2011-2012 and 2012-2013 panels. As MEPS data do not report detailed geography, we restrict our sample to include only respondents living in the Western region in the US. These respondents must also complete all five surveys within the year. We use the longitudinal weight in MEPS to calculate the summary statistics. Our transaction data reflect spending in California only.

⁷The retail cost variable contained in MEPS data does not reflect negotiated prices with insurers.

⁸In brief, we regress health charges on out-of-pocket spending in the MEPS data among marketplace enrollees. We run separate regressions in 2013 and 2014 to allow the relationship between health charges and health out-of-pocket to differ between these periods. We then take the observed out-of-pocket spending in our transactions data and use the regression estimates in Appendix Table A7 to impute total health charges and changes in charges over time.

A8 shows that in MEPS, transactions increase 21% following marketplace enrollment in 2014, but with large standard errors. In our transactions data, the change is 25%, though now significant given our larger sample. The health out-of-pocket spending in the transactions data also increase by 23%, while MEPS shows an insignificant positive change, with wide standard errors.

Panel B of Appendix Table A8 also reports the change in health charges in MEPS data following enrollment. In MEPS, charges increase by 28% among these households who sign up for individual insurance. Although the MEPS estimates are noisy, we find the increase in health care consumption looks similar regardless of whether we use health charges or a count of health care transactions. Because our estimates of health transactions line up closely with those in MEPS, we interpret the change in health care transactions we observe in bank and credit card data as a reasonable proxy of the overall quantity of health care consumed. Using the regression estimates in Appendix Table A7, we impute total health charges and find that the changes in charges over time equal 37%, within the 95% confidence interval of the MEPS estimate.

D Validating participation behavior via Covered California enrollment

We validate our statistics on enrollment using complementary enrollment data from the Covered California marketplace, henceforth CCM. We use the universe of CCM household enrollment data for the years 2014-2016 from Tebaldi (2017). The records cover roughly 1.8 million households who purchase private health insurance through the Covered California marketplace.⁹ We observe monthly enrollment details by household along with household demographics and the level of the household’s gross and net premium payments (net of subsidies). With these records, we construct variables including sign-up month, drop-out month, number of coverage/active months, enrollment status (open enrollment vs. off-cycle enrollment), and an allowed dropout indicator, using the same definitions we applied to premium payments in our transactions data. With the gross premium and net premium reported by household, we construct an unsubsidized indicator, which takes value 1 for a household-year in which the net premium is equal to the gross premium.

We first use the CCM data to validate statistics in our transaction data, including the 2014 market shares by insurance carrier, six/eight-month dropout rates by enrollment status, and shares of 2014 open enrollees by length of coverage. However, before comparing the statistics, we adjust for a key distinction between the datasets: our transactions data covers households who purchase private health insurance both via the CCM and off the marketplace (i.e. directly from insurance carriers or via brokers) while the CCM data omits off-marketplace insurance purchases. The latter constitutes a non-trivial share of individual insurance enrollment in California.¹⁰ We do not have enough information to distinguish these two groups of purchasers in our data.

To make the two datasets comparable, we therefore re-scale the CCM data to be representative of enrollees in the entire market in a given year. Assuming that we know the number of individuals who purchase private insurance in a given year from a given carrier, we can subtract out the number of individuals who participate in the CCM by carrier (i.e. all individuals we observe in the CCM data for a carrier), yielding the number

⁹The raw CCM data from Tebaldi (2017) includes 1,874,970 households. After dropping households with missing variables (region, carrier, or enrollment status) and dropping those households that enrolled in catastrophic insurance plans, we observe 1,791,389 households from 2014-2016.

¹⁰For the carriers in our California sample, enrollment counts in Medical Loss Ratio data suggest roughly 60% of individual insurance enrollees purchased off the marketplace in 2014.

of off-marketplace households by carrier. We also know that all consumers who purchase off-marketplace do so without subsidies.

These facts combined suggest the following procedure to construct weights. First, we collect data from The Centers for Medicare and Medicaid Services’ Medical Loss Ratio (MLR) report in California by carrier and year.¹¹ The MLR data provides a total count of member months and covered lives by carrier in the individual insurance market. We use this count as a measure of the number of households in California who buy individual insurance from a CCM carrier in given year, both via the CCM, via direct purchase, or via an insurance broker.¹² We then subtract out the number of subsidized CCM households we observe in the CCM data, yielding the carrier-specific number of unsubsidized individuals both *on and off* the CCM in that year. Then, we take a ratio of this number to the number of unsubsidized individuals in the CCM and assign this ratio as a weight to all unsubsidized individuals in the CCM. Finally, we normalize the weights such that, once summed up at the carrier level, they match the market shares by carrier in the MLR data. Appendix Table A4 illustrates the construction of this set of weights.

E Determining Medi-Cal eligibility

To determine the probability that a household might be eligible for Medi-Cal in a given month, we develop a three part procedure. First, we convert our observed post-tax income into a pre-tax measure. This conversion depends on the household’s tax rate, which in turn depends on the size and composition of the household. We do not observe family characteristics directly in our data and therefore make the conversion under a range of family size assumptions. Second, we collect monthly income eligibility thresholds based on the federal poverty level set for each family size and for each year. We compare our pre-tax income by family size against the thresholds to determine Medi-Cal eligibility in each family size bin for each household in our data. Finally, third, we collect data on the empirical distribution of family sizes by income within the population of California residents who purchase individual insurance. We then multiply the likelihood that a household falls into each family size bin by the eligibility determination. We sum these values to find a single weighted measure of the probability of Medi-Cal eligibility in each month for each household in our bank data. We describe this procedure below in more detail.

Calculating pre-tax income

To convert our post-tax income measures into pre-tax measures, we need to assign average tax rates for federal, California, and federal (FICA) payroll taxes to each household in our dataset. To do this, we use the National Bureau of Economic Research’s TAXSIM software package. Specifically, we use the following procedure:

1. Using our observed post-tax annual income measure, create a grid of pre-tax income around the observed post-tax amounts. We set the end points of the grid at 0 and the max of our post-tax income divided by $(1 - .4)$, to approximate a 40% average tax rate. We fill in the grid with approximately

¹¹This data is available publicly at <https://www.cms.gov/CCIIO/Resources/Data-Resources/mlr>.

¹²Specifically, we divide each carrier’s count of member-months by 12 to approximate the number of households enrolling. We verify that our numbers for the three largest carriers match those reported in public data by the Kaiser Family Foundation from 2014 to 2016.

70,000 points per year. Our grid has larger spacing between grid points as the grid approaches the maximum income value.

2. Create a matrix of incomes and family sizes to feed into NBER’s TAXSIM (v27) software. We assign the state to California and consider 12 alternative hypothetical household compositions.¹³
3. Feed the income and family size-specific combinations into the TAXSIM software. The program produces output that includes federal, state, and FICA payroll taxes in dollars. We convert these dollars of tax to an average tax rate by dividing the tax dollars by the total pre-tax income. We then sum the state, federal and fica tax rates into a single average rate, τ_i .
4. Convert the income in this simulated dataset to be post-tax by multiplying the pre-tax income by $1 - \tau_i$.
5. To be able to match to the observed post-tax income in the bank dataset, create bins of the hypothetical post-tax income. Within each bin, we find the average of τ_i across all i s in the bin and label it $\bar{\tau}_i$.
6. For each observed post-tax income in our bank data, match the income to one of the bins created above. We assign to that observation the $\bar{\tau}_i$ for the bin. We then calculate pre-tax income by dividing the observed post-tax income by $(1 - \bar{\tau}_i)$.
7. Smooth over fluctuations around the borders between months by calculating a rolling average of pre-tax income by averaging periods

$$(t - 1, t, t + 1)$$

income together. In doing so, we lose the first and last period’s observed income for each individual.

Eligibility determinations

Given a level of pre-tax income for the household under each hypothetical family size, we compare these pre-tax income measures to the monthly Federal Poverty Line (FPL) thresholds appropriate for adults aged 18 to 64 under the ACA. We use 138% of the FPL as our threshold. We use this threshold in 2013 as well, even though eligibility under the ACA guidelines didn’t come into effect until 2014. Using the year 2013 data allows us to examine how hypothetical eligibility changed for open enrollment participants in 2014 in the period before sign-up compared to the period after drop-out.

Family size weighting

We construct weights for each of the 12 household compositions. To do so, we combine data from two sources. We use pooled 2011-2015 NHGIS data to build weights that vary by county, income group, and household composition. Specifically, for each of our six income groups (0-20K, 20K-40K, 40K-60K, 60K-100K, 100K-200K, and $\geq 200K$) and for each county, we calculate fractions of households in each of our hypothetical household bins. The combined weights sum to one within each year, county, and income group.

Computing the eligibility probability by household

Finally, to calculate a single Medicaid eligibility probability per household, we begin with the household-monthly panel described in Section A and merge in the 12 household-composition-specific monthly pre-tax

¹³The 12 household compositions comprises two marital statuses (single and married) and six household sizes (number of children under thirteen equal to 0, 1, 2, 3, 4, and 5).

incomes. Next, we compare the annual pre-tax income levels for each household size against the 138% of the federal poverty level threshold to determine Medicaid eligibility (0 or 1) by household-month. Then, we merge in the 12 household-composition-specific weights from the above step by year-county-income group and calculate a weighted average of our 12 Medicaid eligibility indicators. This process yields a single measure of Medicaid eligibility probability by household-month. We plot eligibility in Figure A4.¹⁴

F Classifying transactions by urgency, essentialness

While our design assessed differences in spending and enrollment patterns between “drop-out” types and “long-term” enrollees, we did not directly answer the question of whether “drop-out” types have fundamentally different health consumption needs. To do so, we study our sample of 1,522 unique root words classified under the “health” spending category,

We seek to classify these words along two dimensions: essentialness and urgency. Such a classification task is inherently difficult for a computer to perform and thus requires human intelligence. To perform this classification task without introducing researcher bias, we relied on the service of Amazon Mechanical Turk workers (MTurk). MTurk is a marketplace which allows access to a on-demand scalable workforce, and thus to perform thousands of classification tasks quickly.

We designed a classification survey so that 10 different workers classify each health-related word on the scales of essentialness and urgency. We define ‘1’ as the lowest level of urgency/essentialness, while 10 represents an extremely urgent/essential health care need. We chose to work exclusively with MTurk workers that the platform had awarded “Masters” qualification— that is, workers that MTurk identified as high-performing workers across a range of tasks.

To help workers in our classification task, we provided standard examples of activities considered urgent and essential. We defined urgent health care spending as a service a patient cannot postpone. A patient with a heart attack, for example, must go to the emergency room immediately or risk death. Similarly, some visits to the doctor are essential: a patient with a chronic condition like diabetes needs to receive regular treatment. We also asked workers to tag whether a word was irrelevant for health care. We omit from our word analysis any root word for which at least two of the ten MTurk workers flagged the word as irrelevant.

The final surveys had a total of 15,220 assignments (1,522 words to classify, aka “HITS” in MTurk, times 10 workers per word). In total, 124 distinct workers completed the survey, with a mean number of assignment per worker of 857 words. Workers took in average 42 seconds to classify a word. Of the initial list of 1,522 words, 134 were rated as “health-relevant”, defined as words for which at least 9 out of 10 workers rate them as relevant.

G Additional health and drug consumption analyses

In addition to the analyses of consumption using health transactions and health spending in Section 4, we conduct similar difference-in-difference analyses using drug transactions and drug spending. We separate

¹⁴In Figure A4(b), we observe a large decrease in Medicaid eligibility around insurance sign-up for the population with income below \$20,000 in 2013. This is an artifact of the sample definition. Households with post-tax income below \$20,000 in 2013 are very likely to be eligible for Medicaid in 2013. That a household signed up for private coverage in 2014 suggests its income increased and its likelihood of eligibility fell.

out the drug analysis because, as discussed in Appendix Section C, our drug measure contains all pharmacy spending, including non-prescription purchases. We also repeat our health and drug analysis focused on households with young children.


Looking first at drug spending, our results appear consistent with our findings from general health spending. In Appendix Table A6, we replicate our event study analysis to explore how monthly drug spending changes at sign-up and drop-out relative to the period of coverage. Our results are noisier due to measurement error in our drug spending variable. However, the changes in drug spending match the pattern in the health category shown in Table 3 in the main text. In particular, we observe an increase in spending upon sign-up in the post-ACA period, followed by a decrease upon drop-out. The magnitude and significance of the changes are greatest for the lowest income groups.

Second, we examine the effect of ACA enrollment on the health spending of households with children. We label families based on the number of child-related transactions we identified in the household’s transactions using our machine learning algorithm, as discussed in Appendix B; those households with the highest child-related transactions are classified as having children, up to the share defined in Appendix A.

We repeat our spending analysis for households with children only. We report our results in Appendix Table A9. Again, we find that take-up of insurance leads to an increased number of health transactions and more dollars spent for lower income households. Interestingly, the effect of ACA enrollment on total dollars spent on health care is even more pronounced for lower income households with children than for the average lower income household. ACA plan enrollment increases health transactions by almost 56% relative to the reference group for households with income less than \$20K and by over 34% for households with income between \$20K and \$40K. The dollar level of spending increased by 106% for the lowest income group, suggesting that the spending burden of ACA plans appears higher for families than for individuals. Even though families increased their number of health care transactions by the same amount as all enrollees, they had to spend substantially more out of pocket. This is consistent with higher deductibles for family insurance plans relative to plans covering individuals. For drug spending, our estimated effects for families with children are noisier, but are consistent with the patterns seen in health spending.


H Appendix figures and tables

Figure A1: Word Clouds for Machine Learning Predictions, Health and Drug Spending




Notes: This figure reports the top 100 words classified by our machine learning algorithm as health spending, drug spending, child-related, or other types of spending. The font size indicates the frequency with which the word appears in the transactions data.

Figure A2: Geographic Distribution of Financial Account Members in California




Notes: This map illustrates the geographic distribution of households in our transaction data relative to the number of households in the 2011-2015 NHGIS data. Specifically, for each county, we count the number of unique households in our data, divide by the corresponding number in NHGIS, and then multiply by 100 to obtain percentages. To be included in our sample, each household must have a “modal county” in California in every year that we observe it (See Appendix A for details on sample construction). Given our restrictions, we have 798,085 households in California from 2011-2015, which accounts for about 7.0% percent of the actual number in ACS. The bank customers we observe are not randomly distributed in the state; as shown in the figure, counties in the Bay Area such as Alameda, Santa Clara, and San Francisco counties, are over-represented in our sample. We observe roughly 10% of the population in the Bay Area. By contrast, in many rural counties, our sample represents less than 3% of the total population. To account for this sampling in most of our analyses, we use the inverse of these geographic shares as a weight.

Figure A3: Six Month Drop-out Rate Validation: Full Sample vs Households with an Unlinked Credit Card




Notes: This chart compares our measure of drop-out among our entire transaction sample vs. among households who do not have an unlinked credit card. We define households as having an unlinked credit card as follows: Using the bank account data, we observe the total amount the household spends paying off credit card balances, regardless of whether the credit card is linked in our dataset. We then compare these overall payments to those payments used to pay off balances on linked credit cards. If the payments to the linked credit cards are less than 85% of the total credit card spending out of the bank account, we consider this household to have an un-linked card.

Figure A4: Medicaid Eligibility Probability by Pre-ACA Income Category




Notes: This figure shows the probability of medicaid eligibility among households by income group. We assign income groups using pre-ACA annual post-tax income from 2013. The x-axis contains the months surrounding the coverage period under individual health insurance; all months of coverage are collapsed in the event month 0. The solid line represents the average probability among individual households in the income bin. The shaded areas represent the 95% confidence interval around this average at each month in the event study. The figures contain all households that (i) sign up for individual insurance coverage under the ACA, including both on and off-cycle enrollees from 2014-2015; (ii) have at least 10 months of data before sign-up and after drop-out; and (iii) subsequently drop coverage less than 9 months after initial sign-ups excluding drop-outs in November/December; and (iv) make premium payments for at least two months. All regressions use county weight to account for sampling differences across counties and absorb individual fixed effects.

Figure A5: Social Surplus under Alternative Mandate Policies



Notes: This figure illustrates the welfare consequences of a mandate penalty for non-enrollment. We plot welfare when the mandate penalty is paired with either (a) the optimal drop-out penalty or (b) no drop-out penalty. In each case, we vary the mandate penalty from \$0 to \$40 in our simulated setting. The optimal penalty is set such that drop-outs enroll for $\phi/12 = c_D/c_{ND}$ fraction of the year. In this numerical illustration we set $\phi/12 = 1/2$, $N_D = N_{ND} = 100$, $c_D = 20$, $c_{ND} = 40$, $G_D(\cdot) = Unif([0, 50])$, and $G_{ND}(\cdot) = Unif([40, 100])$.

Figure A6: Individual Insurance Enrollment, 2012-2015



Notes: This figure provides summary statistics of enrollees in individual insurance from 2012 to 2015 (inclusive). **Panel (a)** shows new enrollee sign-ups and drop-outs in each month from January 2012 to October 2015. A new sign-up or drop-out is defined based on the observed first or last monthly premium payment, respectively. **Panel (b)** shows the absolute count of members in each enrollment period. The “open enrollment” group includes households who pay their first premium in the open enrollment period; the “off-cycle” group includes households who pay their first premium between April and November of the enrollment year. **Panel (c)** plots the median 2013 household income of ACA enrollees by ACA enrollment cycle. We use county weight when calculating median income. The reference group consists of households who we do not observe paying individual insurance premiums. We restrict our sample to households who pay premiums for at least two months in panels (a), (b), and (c). **Panel (d)** uses MEPS data to plot the mean number of months of insurance coverage in 2013 for the population of marketplace enrollees during the 2014 open enrollment period.

Table A1: Standard Benefit Design by Metal Tier

Panel A: 2014				
Coverage Category	Bronze	Silver	Gold	Platinum
Percent of cost coverage	Covers 60%	Covers 70%	Covers 80%	Covers 90%
Preventive Care Copay	No cost	No cost	No cost	No cost
Primary Care Visit Copay	\$60 for 3 visits	\$45	\$30	\$20
Specialty Care Visit Copay	\$70	\$65	\$50	\$40
Urgent Care Visit Copay	\$120	\$90	\$60	\$40
Emergency Room Copay	\$300	\$250	\$250	\$150
Lab Testing Copay	30%	\$45	\$30	\$20
X-Ray Copay	30%	\$65	\$50	\$40
Generic Medicine Copay	\$19 or less	\$19 or less	\$19 or less	\$5 or less
Annual Out-of-Pocket, Maximum Individual and Family	\$6,350 individual and \$12,700 family	\$6,350 individual and \$12,700 family	\$6,350 individual and \$12,700 family	\$4,000 individual and \$8,000 family
Panel B: 2015				
Coverage Category	Bronze	Silver	Gold	Platinum
Percent of cost coverage	Covers 60%	Covers 70%	Covers 80%	Covers 90%
Annual Wellness Exam	\$0	\$0	\$0	\$0
Primary Care Visit	\$60	\$45	\$30	\$20
Specialist Visit	\$70	\$65	\$50	\$40
Emergency Room	\$300	\$250	\$250	\$150
Laboratory Tests	30%	\$45	\$30	\$20
X-Ray	30%	\$65	\$50	\$40
Imaging	30%	20%	20%	10%
Preferred Drugs	50%	\$50	\$50	\$15
Generic Drugs	\$15 or less	\$15 or less	\$15 or less	\$5 or less
Deductible	\$5,000	\$2,000 medical \$250 brand drugs	\$0	\$0
Annual Out-of-Pocket, Maximum Individual and Family	\$6,250 individual \$12,500 family	\$6,250 individual \$12,500 family	\$6,250 individual \$12,500 family	\$4,000 individual and \$8,000 family

Notes: This table reports the standard benefit design required for plans offered in the individual insurance marketplace in California in 2014 and 2015 by metal tier. The ‘percent of cost coverage’ represents a share of average annual cost.

Table A2: Examples of Financial Transaction Records

Description	Date	Amount	Merchant	Street	City	State	Zip
A. Health Transactions							
CHANG ACUPUNTURE	2011-11-11	61.68	Acupunture	4840 Irvine Blvd	Irvine	CA	92620
EPIONE MEDICAL CORP BEVERLY HILLS CA	2012-08-10	4890.0	Epione Medical	444 N Camden Dr	Beverly Hills	CA	90210
KAISER PERM XXXXXXXXX SAN FRANCISCO CA	2013-02-18	20.00	Kaiser	2425 Geary Blvd	San Francisco	CA	94115
GLOBAL ONCOLOGY MONTEREY PARK CA	2014-09-23	20.00	Global Oncology	600 N Garfield Ave	Monterey Park	CA	91754
DIABETES & ENDOCRINOLOGY XXX-XXX-XXXX CA	2015-03-28	102.09	Center For Diabetes	400 Camarillo Ranch Rd	Camarillo	CA	93012
B. Drug Transactions							
RITE AID STORE 5456 LOS ANGELES CA	2011-04-05	73.53	Rite Aid	11321 National Blvd	Los Angeles	CA	90064
KAISER LIVERMORE CA - DRUG STORES PHARMACIES	2012-06-09	5.00	Long's Drug Stores	1500 1st St	Livermore	CA	94550
WALGREENS #0063 BELMONT CA	2013-01-02	185.00	Walgreens	900 Ralston Ave	Belmont	CA	94002
CVS PHARMACY #3951	2014-03-10	99.50	CVS Pharmacy	305 S Highway 101	Solana Beach	CA	92075
MILART PRESCRIPTION PH BEVERLY HILLS CA	2015-10-12	24.96	Milart Pharmacy	300 S Beverly Dr	Beverly Hills	CA	90212
C. Child-Related Transactions							
TOYSRUS SHIN-KAMAGAYA TIBAKEN JPN	2011-10-22	100.00	ToysRUs	1737 Post St	San Francisco	CA	94115
GYMBOREE XXXXXXXXXXXXX GLENDALE CA	2012-08-18	27.71	Gymboree	1314 Glendale Galleria	Glendale	CA	91210
DISNEY EVENTS XXXXXXXXXXXXX FL	2013-10-25	304.5	Disney	2522 Capital Cir NE	Tallahassee	FL	32308
KUMON MATH & READING C CHINO CA	2014-11-22	1460.00	Kumon	5420 Philadelphia St	Chino	CA	91710
BUYBUYBABY#3084 XXXXX MISSION VIEJO CA	2015-03-28	4.31	Buy Buy Baby	25322 El Paseo	Mission Viejo	CA	92691
D. Other Remaining Transactions							
CINEMARK MOVIES #12QPS DANVILLE CA	2011-02-19	12.50	Cinemark USA, Inc.	4175 Blackhawk Plaza Cir	Danville	CA	94506
WALMART 5435 SE2 XXXXX0047 SAN JOSE CA	2012-06-01	57.80	Walmart	777 Story Rd	San Jose	CA	95122
MCDONALD'S F855 RIVERSIDE CA	2013-03-23	6.69	McDonald's	2242 University Ave	Riverside	CA	92507
COSTCO GAS #0420 XXXXXXXXXOXNARD CA	2014-03-08	82.43	Costco Gas	2001 E Ventura Blvd	Oxnard	CA	93036
FEDEX OFFICE XXXXXXXXX SAN DIEGO CA	2015-02-11	7.84	FedEx	111 W Harbor Dr	San Diego	CA	92101

Notes: The rows above provide examples of financial transactions we observe in our data. For each transaction, we observe a brief description that we use in our machine learning algorithms to classify the transaction as either health spending, drug spending, child-related spending, or other types of spending (See Appendix B). We use the physical address of transaction locations to help identify each household's home geographic region.

Table A3: Number of Households after Data Creation Steps

Sample Criterion	# Households	Percent
Households in national transactions data	9,223,300	
California non-mover households (starting sample)	1,141,592	100.0%
Drop households with unidentified modal county	1,130,304	99.0%
Drop households with no bank account	1,047,634	91.8%
Drop households with missing enrollment status i.e. sign-up in Dec 2015	1,044,924	91.5%
Drop households who enter after Dec 2013 or exit before Jan 2013	1,014,591	88.9%
Drop households with inactive accounts:	843,766	73.9%
- Fraction of months with observed positive bank credits/debits ≤ 0.5		
- Fraction of months within each year with any observed transactions ≤ 0.5		
Drop households with bad income measure i.e. taxes $> 0.5 \times$ pre-tax income	810,651	71.0%
Drop households that are likely to be small businesses	798,087	69.9%
- Average annual credit transaction count ≥ 500 (around 99th percentile)		
- Average annual debit transaction count $\geq 2,000$ (around 99th percentile)		
Drop counties with less than 5 households	798,085	69.9%
California non-mover households (cleaned sample)	798,085	69.9%
California non-mover households (final subsample)	245,904	

Notes: This table describes the creation of the analysis sample. Nationally, our raw data follows 9.22 million households with accounts at banks serviced by our data provider. We assign each household a “modal county” by taking the county in which it transacts most frequently each year. We include all households with a modal county in CA for every year in which they appear in our data from 2011 to 2015; in effect, we exclude households who leave California in our sample period. We also drop households who: are recorded in counties near the CA border, as we cannot precisely identify their locations; have no bank account; sign up for private health insurance for the first time in December 2015; and drop households who enter our data in 2014-2015 or exit in 2011-2013. These sequential drops leave 1.01 million households.

To further remove households with inactive accounts, we restrict to those for which we observe positive monthly bank credits and positive bank debits for more than half the number of months they appear in the data. We also require transactions for more than half the number of months each year. These two restrictions leave 843,766 households. Next, we drop households with bad income estimates or those with imputed income taxes greater than 50% of their pre-tax incomes. In addition, we drop households with extremely high transaction frequencies because they are likely to be small businesses. Finally, we drop counties that have fewer than five households. 798,085 million households remain after these steps. To form our final analysis sample, we save all individual market enrollees and a 20% random sample of those households who never purchase individual insurance during 2011-2015. The final sample contains 245,904 households.

Table A4: Weight Calculation Using Medical Loss Ratio Data

Year	Carrier	Total Enrollment	Covered California Enrollment			Scale
		Medical Loss Ratio	Total	Subsidized	Unsubsidized	
2014	Anthem	768,080 (36.82%)	247,128	224,759	22,369	24.29
2014	Blue Shield	548,581 (26.30%)	228,400	207,054	21,346	16.00
2014	Kaiser	494,633 (23.71%)	154,226	140,204	14,022	25.28
2014	Health Net	222,756 (10.68%)	172,136	159,643	12,493	5.05
2014	L.A. Care	20,211 (0.96%)	22,383	20,001	2,172	0.10
2014	CCHP	11,751 (0.56%)	9,680	9,028	652	4.18
2014	Sharp	10,189 (0.49%)	8,907	7,773	1,134	2.13
2014	Molina	6,076 (0.29%)	6,408	5,989	332	0.26
2014	Western	3,413 (0.16%)	2,678	2,313	365	3.01
2014	Contra Costa	816 (0.04%)	802	763	39	1.35
2015	Anthem	758,581 (32.20%)	278,422	250,397	28,025	18.13
2015	Blue Shield	631,594 (26.81%)	252,068	222,536	29,532	13.85
2015	Kaiser	629,325 (26.71%)	254,552	226,697	27,855	14.45
2015	Health Net	265,078 (11.25%)	174,563	164,843	9,720	10.31
2015	Molina	19,609 (0.71%)	20,760	16,632	1,151	2.59
2015	Sharp	19,203 (0.82%)	13,407	11,604	1,803	4.21
2015	L.A. Care	15,587 (0.66%)	14,051	12,704	1,347	2.14
2015	CCHP	12,919 (0.55%)	8,987	8,221	766	6.13
2015	Western	6,879 (0.29%)	4,632	3,983	649	4.46
2015	Contra Costa	. (0.00%)	.	.	.	.
2016	Anthem	741,248 (30.67%)	265,920	231,966	33,954	15.00
2016	Blue Shield	733,032 (30.33%)	289,559	256,006	33,553	14.22
2016	Kaiser	615,523 (25.47%)	265,368	229,441	35,927	10.75
2016	Health Net	201,010 (8.32%)	130,703	121,637	9,066	8.76
2016	Molina	66,079 (2.73%)	67,352	63,249	4,103	0.69
2016	Sharp	27,072 (1.12%)	18,378	15,688	2,690	4.23
2016	CCHP	12,913 (0.53%)	10,097	9,128	969	3.91
2016	L.A. Care	10,644 (0.44%)	9,725	8,727	998	1.92
2016	Western	9,521 (0.39%)	6,403	5,479	924	4.37
2016	Contra Costa	. (0.00%)	.	.	.	.

Notes: This table describes the construction of household weights that we use to adjust the Covered California Marketplace (CCM) data to match the transactions sample. The transactions data include households who bought individual insurance both on and off the California marketplace; the CCM data include only on-marketplace enrollment. Thus, to make the CCM data comparable, we scale up households in the CCM data to make them representative of the entire market. We use Medical Loss Ratio (MLR) data, which contain a measure of the total number of enrollees by year-carrier (See Column (3)). For each year-carrier in the CCM data, we classify households into the subsidized and unsubsidized groups shown in columns (5) and (6), respectively. Because all enrollees off the CCM purchase unsubsidized plans, we weight up the CCM unsubsidized population to account for the difference between the CCM and MLR data. The scale reported in the final column is year-specific and carrier-specific. We use this set of weights when employing CCM data in our analyses.

Table A5: Comparison of Annual Spending with MEPS Data

Panel A. Summary Statistics of Annual Drug/Health Spending										
Year	MEPS			Transactions						
	N	Drug	Health	N	Drug	Health				
2012	9,327	141.7	204.8	194,987	257.8	209.7				
2013	9,263	122.4	258.0	245,621	324.5	219.0				
2014	8,761	119.8	262.5	232,966	322.9	235.5				
2015	8,588	114.2	276.1	209,260	303.3	246.6				

Panel B. Distribution of Annual Health Spending										
	year	p1	p5	p10	p25	p50	p75	p90	p95	p99
MEPS	2012	0.0	0.0	0.0	0.0	5.0	110.0	460.0	1,010.0	3,204.0
	2013	0.0	0.0	0.0	0.0	15.0	150.0	589.0	1,185.0	3,676.0
	2014	0.0	0.0	0.0	0.0	10.0	150.0	580.0	1,142.0	3,601.0
	2015	0.0	0.0	0.0	0.0	4.0	145.0	583.0	1,253.0	4,725.0
Transactions	2012	0.0	0.0	0.0	0.0	0.0	125.0	473.6	911.0	3,037.0
	2013	0.0	0.0	0.0	0.0	0.0	130.0	495.0	968.1	3,200.0
	2014	0.0	0.0	0.0	0.0	2.8	148.4	537.4	1,034.8	3,245.6
	2015	0.0	0.0	0.0	0.0	4.2	150.0	548.0	1,059.6	3,399.3

Notes: The Medical Expenditure Panel Survey (MEPS) is conducted in two consecutive years. For each year in our sample, we use two adjacent MEPS panels as our comparison sample. For example, the 2012 MEPS statistics above are based on data from the 2011-2012 and 2012-2013 panels. We restrict our MEPS sample to include only respondents living in the Western region in the US and we adjust the sample statistics using MEPS' weights. Our transactions data reflect spending in California only. We classify drug and health spending using the machine learning algorithm discussed in Appendix B. Because we do not have data on all 12 months for some respondents (e.g. those who enter after January or exit before December), we calculate monthly average health/drug spending each year and multiply that by 12 to obtain the annual figures. When computing the annual statistics above, we use a county weight to account for sampling differences across counties in our data. Here, drug spending in MEPS covers true prescription drug purchases. In our transactions data, we observe drug spending as all spending at drug stores, including non-prescription purchases.

Table A6: Changes in Drug Consumption Around Sign-up and Drop-out

Panel A: Change in Drug Spending: Pre ACA Sign-up/Drop-up					
Annual Income:	(1) ≤20K	(2) 20K-40K	(3) 40K-60K	(4) 60K-100K	(5) 100K-200K
Sign-up (% change)	0.866*** (0.137)	0.349*** (0.115)	0.341*** (0.112)	0.370*** (0.094)	0.449*** (0.104)
Drop-out (% change)	-0.260** (0.112)	0.183* (0.097)	-0.137 (0.114)	-0.019 (0.095)	-0.006 (0.119)
Pre-Signup Mean Drug Spending	11.43	14.88	23.14	19.51	23.91
Number of Observations	5,006	7,458	5,981	7,452	6,599
Panel B: Change in Drug Spending: Post ACA Sign-up/Drop-up					
Annual Income:	(1) ≤20K	(2) 20K-40K	(3) 40K-60K	(4) 60K-100K	(5) 100K-200K
Sign-up (% change)	0.233*** (0.065)	-0.003 (0.030)	0.022 (0.041)	-0.099** (0.040)	-0.036 (0.048)
Drop-out (% change)	-0.172*** (0.057)	-0.105*** (0.034)	-0.146*** (0.040)	-0.118*** (0.031)	-0.136*** (0.043)
Pre-Signup Mean Drug Spending	16.59	19.86	27.82	38.21	44.91
Number of Observations	14,079	27,352	18,991	20,241	15,137
Panel C: Change in Drug Transactions: Pre ACA Sign-up/Drop-up					
Annual Income:	(1) ≤20K	(2) 20K-40K	(3) 40K-60K	(4) 60K-100K	(5) 100K-200K
Sign-up (% change)	1.009*** (0.151)	0.496*** (0.092)	0.410*** (0.089)	0.407*** (0.086)	0.439*** (0.096)
Drop-out (% change)	-0.240* (0.124)	0.148* (0.089)	-0.066 (0.091)	-0.013 (0.089)	0.044 (0.099)
Pre-Signup Mean Drug Transactions	0.38	0.53	0.77	0.68	0.81
Number of Observations	5,006	7,458	5,981	7,452	6,599
Panel D: Change in Drug Transactions: Post ACA Sign-up/Drop-up					
Annual Income:	(1) ≤20K	(2) 20K-40K	(3) 40K-60K	(4) 60K-100K	(5) 100K-200K
Sign-up (% change)	0.236*** (0.056)	-0.006 (0.027)	0.070** (0.033)	-0.036 (0.028)	0.000 (0.030)
Drop-out (% change)	-0.216*** (0.051)	-0.102*** (0.027)	-0.157*** (0.034)	-0.105*** (0.026)	-0.125*** (0.033)
Pre-Signup Mean Drug Transactions	0.72	0.88	1.08	1.31	1.41
Number of Observations	14,079	27,352	18,991	20,241	15,137

Notes: This table examines changes in drug consumption of drop-out consumers during coverage vs. after drop-out during the pre-ACA period vs. post-ACA period. Observations are at the household-month level. “Sign-up” is a indicator equal to 1 in the month of sign-up and in subsequent enrollment months. “Drop-out” is an indicator equal to 1 in the month of drop-out and thereafter. The pre-ACA period includes households who both sign up and drop out prior to July 2013, while the post-ACA period includes households who sign up in the 2014 open enrollment period or later. Regressions are run separately by income group, where income is defined as 2013 annual post-tax income. For both periods, we restrict our sample to households who drop out less than 9 months after initial sign-up excluding November/December dropouts. We also restrict our sample to include only households who have data at least 10 months before sign-up and at least 10 months after drop-out. We top-code drug transactions and drug spending at the 99th percentile value within each income group. All regressions control for household fixed effects, monthly income, and average lagged monthly income from the past three months. We use county weight to account for sampling differences across counties in our data. Units are measured as a percentage change, relative to average drug consumption amount in the 10 months leading up to sign-up. In all regressions, we further restrict our sample to households who have coverage for at least two months. Robust standard errors are clustered by household.* p<0.1, ** p<0.05, *** p<0.01.

Table A7: Imputation of Healthcare Charges from Health Out-of-pocket in MEPS

	2013 Health Charges	2014 Health Charges
Health OOP	5.640** (2.668)	4.580* (2.391)
Constant	2658.086** (1090.061)	4159.501** (1810.254)
Observations	201	201

Notes: This table presents regressions of health charges on health out-of-pocket spending in MEPS. Sample includes MEPS respondents in the 2013-2014 longitudinal file who: (a) report purchasing marketplace coverage at some point from January 2014 through May 2014; (b) participate in all five surveys within the year; and (c) have 2013 annual income less than or equal to \$200,000. Robust standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A8: Healthcare Consumption for 2014 Open Enrollees: Transactions vs. MEPS Data

Panel A: Transactions			
	(1) Health Charges	(2) Health OOP	(3) Health Transactions
Post ACA	1495.301*** (46.712)	53.046*** (9.135)	0.806*** (0.072)
Pre ACA	3984.470*** (42.644)	235.191*** (7.559)	3.177*** (0.075)
% Change	0.375*** (0.015)	0.226*** (0.043)	0.254*** (0.026)
Observations	22,418	22,418	22,418
Panel B: MEPS			
	(1) Health Charges	(2) Health OOP	(3) Health Transactions
Post ACA	1212.748 (1760.709)	3.977 (58.754)	1.187 (1.011)
Pre ACA	4292.207*** (972.298)	289.758*** (39.300)	5.614*** (0.789)
% Change	0.283 (0.450)	0.014 (0.204)	0.211 (0.194)
Observations	402	402	402

Notes: This table compares health consumption in 2013-2014 among enrollees in the 2014 open enrollment period. We compare consumption in our transactions data (Panel A) vs. MEPS (Panel B). Pre ACA and post ACA correspond to 2013 and 2014, respectively. In Panel A, we restrict our sample to households who have 2013 annual pre-tax income less than or equal to \$200,000. In Panel B, we include households who report purchasing marketplace coverage for at least one month from January 2014 through May 2014. In the MEPS data, we define health charges and health out-of-pocket (OOP) costs as charges from provider visits at offices, emergency departments, and inpatient and outpatient facilities. In our data, we only observe OOP spending and transactions. We impute overall health charges in column (1) using the relationship between health charges and health OOP spending in MEPS (See Table A7). We weight regressions in Panel A by a county weight to account for sampling differences across counties in our data; we weight regressions in Panel B by MEPS's longitudinal weight. All regressions in both panels control for income group fixed effects. Robust standard errors are clustered by household. * p<0.1, ** p<0.05, *** p<0.01.

Table A9: Health/Drug Consumption for 2014 Open Enrollees with Children

Annual Income:	(1) ≤20K	(2) 20K-40K	(3) 40K-60K	(4) 60K-100K	(5) 100K-200K
Panel A: Change in Health Transactions					
Treat × Post ACA	1.064** (0.477)	1.056*** (0.334)	0.076 (0.508)	0.860* (0.490)	0.402 (0.527)
Post ACA	0.515*** (0.081)	0.235** (0.094)	0.126 (0.133)	0.281** (0.129)	0.181 (0.124)
Pre Period Treatment Mean	1.91	3.11	4.43	5.79	6.73
% Change	0.557** (0.250)	0.340*** (0.107)	0.017 (0.115)	0.149* (0.085)	0.060 (0.078)
Panel B: Change in Health Spending					
Treat × Post ACA	93.123*** (26.978)	53.167** (22.189)	44.152 (30.417)	69.577** (32.984)	130.551* (78.482)
Post ACA	23.869*** (4.986)	23.697*** (5.440)	15.558** (7.290)	30.368*** (9.024)	30.997*** (10.462)
Pre Period Treatment Mean	88.11	119.92	185.97	264.42	383.85
% Change	1.057*** (0.306)	0.443** (0.185)	0.237 (0.164)	0.263** (0.125)	0.340* (0.204)
Panel C: Change in Drug Transactions					
Treat × Post ACA	1.401* (0.832)	0.749 (0.811)	1.102 (0.869)	0.545 (0.924)	0.067 (0.860)
Post ACA	2.746*** (0.240)	2.911*** (0.237)	3.234*** (0.277)	2.906*** (0.238)	3.149*** (0.236)
Pre Period Treatment Mean	7.24	10.25	12.73	15.35	16.86
% Change	0.193* (0.115)	0.073 (0.079)	0.087 (0.068)	0.036 (0.060)	0.004 (0.051)
Panel D: Change in Drug Spending					
Treat × Post ACA	32.309 (21.173)	4.988 (19.743)	1.126 (20.017)	-25.999 (28.966)	-28.619 (42.796)
Post ACA	60.671*** (6.026)	62.819*** (6.229)	73.773*** (7.263)	86.848*** (6.724)	103.012*** (8.019)
Pre Period Treatment Mean	145.66	210.65	278.60	371.48	479.09
% Change	0.222 (0.145)	0.024 (0.094)	0.004 (0.072)	-0.070 (0.078)	-0.060 (0.089)
Number of individuals in treatment	146	243	253	306	246
Number of individuals in reference	2,014	2,490	2,253	3,427	3,410

Notes: This table examines healthcare and drug purchasing patterns of enrollees in the 2014 open enrollment period. We compare spending and transactions after vs. before sign-up, and for the enrollment population relative to the reference group who never enroll in individual coverage. We restrict our sample to households we classify as having children, as discussed in Appendix B. “Treat” or treatment group indicator equals 1 if an household signs up during 2014 open enrollment period between December 2013 and March 2014 (inclusive) and 0 if a household belongs to the reference group. “Post ACA” equals 1 in 2014-2015 and 0 in 2011-2013. We run a difference-in-difference regression separately by income group, where income is defined as 2013 annual post-tax income computed from household bank account records. We top-code drug transactions and drug spending at the 99th percentile value within each income group before including them in the regressions. We also control for household fixed effects and use county weight to account for sampling difference across counties in our data. In all regressions, we restrict our sample to households who have coverage for at least two months. Robust standard errors are clustered by household. * p<0.1, ** p<0.05, *** p<0.01.