Insurance without Commitment: Evidence from the ACA Marketplaces

Rebecca Diamond[1 *], Michael J. Dickstein[2 *], Timothy McQuade[3], and Petra Persson[4 *]

[1]Stanford Graduate School of Business
[2]Stern School of Business, New York University
[4]Haas School of Business, University of California Berkeley
[4]Department of Economics, Stanford University
[*]NBER

First draft: September 5, 2017
Last compiled: October 4, 2023

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. These dropouts spend more on health care while covered 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 spending per month of part-year vs. full-year enrollees. Using our empirical model, we quantify changes in premiums and welfare after the imposition of penalties targeting drop-out. We find that overall welfare declines with a ban on drop-out: young and healthy consumers—those who can more easily re-time their health spending as well as those who value the option to exit—choose to forego coverage entirely, leading to higher average costs among the insured population and thus higher premiums.

We thank Anais Galdin, Maximilian Grotz, Abhisit Jiranaphawiboon, Ziao Ju, Jiayu Lou, Lucienne Oyer, Ryan Shyu, and Analysia Watley 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@haas.berkeley.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, similar market design issues arise with drop-out. 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. We show the sign of this externality, however, is ambiguous 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 advantageously attract low-cost enrollees to the market. If healthy types more often enter and then drop out early, paying a fraction of the annual premium costs, insurers can effectively price discriminate between low- and high-cost populations. In this scenario, allowing drop-out can generate greater total welfare.

We study the effect of one-sided commitment empirically 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 U.S. 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 structural model to link the attrition behavior to changes in premiums and welfare in the market.

[1] Currie (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. With these data, we can examine alternative explanations for both drop-out and the re-timing of health expenditures, including shocks to liquidity or changes in overall spending on credit and debit cards.

We focus our analysis on households in California, the state with the largest individual insurance market. Using both administrative data from California's state-based marketplace and our bank transactions data, 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. Analyzing spending, we find, on average, a 25% increase in health 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] Finally, while we find drop-out and health spending to be strongly correlated with income fluctuations in the pre-ACA period, we find a substantially weaker correlation after the implementation of the ACA. Thus, at least in our sample period, liquidity shocks do not seem to explain fully the pattern of exit and the change in health care consumption we observe.

To analyze how this drop-out behavior affects consumer welfare, we start with a conceptual 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 this 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

[2]Periods 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.

[3]Since 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.

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 insured 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 level for full-year enrollees. This behavior would drive up average cost and decrease market efficiency.

Our framework thus suggests a connection between the rate of drop-out and the level of consumer welfare. However, the direction of the welfare effect from restricting drop-out is ultimately an empirical question. In the two-type special case, we can show that one need only measure the ratio of monthly spending between dropouts and non-dropouts; if dropouts spend more per month covered, allowing attrition will reduce welfare. To capture richer settings, we extend our model to allow enrollees to differ both (a) in their ability to re-time expenditures and (b) in their underlying health states.

In our dynamic structural model, enrollees face a choice to consume health care in three periods over a year. In the initial period, the household chooses a plan based on the plan's characteristics, including its premium, level of coverage, and factors like plan or brand reputation. The household also faces dynamic considerations; it can endogenously shift spending to earlier months of the year, depending on the ease with which it can re-time its health spending, and it can choose to exit in periods two or three, saving premiums costs but losing coverage in those months.

We estimate this dynamic model by combining household enrollment data from the Covered California marketplace with household health transaction information from our bank data. We allow our key parameters to differ by household demographics, including all combinations of age by income and by household size bins. We find that younger and richer households are both more able to re-time their expenditures and more likely to drop coverage conditional on signing up initially.

Interestingly, conditional on expected spending needs, households with a greater ability to re-time seem to choose more comprehensive plans—with fewer out-of-pocket costs—during the months they remain enrolled.

We use these parameter estimates to consider policies that enforce commitment: if insurers could penalize early drop-out, thereby requiring enrollees to remain insured for a greater portion of the plan year, would consumer welfare decline overall? We quantify changes in premiums and welfare after the imposition of penalties targeting drop-out. We find that overall consumer welfare falls with the new penalties. Some young and healthy households choose to be uninsured in the new environment, given that the penalties eliminate both the gains from re-timing expenditures and also any option value the household places on exiting early. As a result, the pool of insured households becomes more adversely selected, leading to higher average costs to the insurer. In a competitive insurance environment, we predict higher baseline premiums for all plan types, but particularly for silver plans that attract more of the households who choose to forego coverage.

Literature. 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 (Forthcoming) 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 insurance in the United States under the ACA, and outlines our data and the algorithm we use to measure health spending. In Section 3, we provide evidence on the lack of commitment of enrollees and the pattern of health spending around enrollment. In Section 4, we describe a conceptual model showing how the presence of dropouts can affect equilibrium premiums and welfare; we then estimate a structural version of the model using data on plan choices, drop-out, and health care spending. In Section 5, we use our structural estimates to quantify the effect of penalties targeting drop-out. Section 6 concludes.

2 Data and setting

In our empirical analyses, we explore the welfare consequences of one-sided commitment in the setting of the individual insurance market in California. Before moving to our analyses, we first describe our setting and the nature of commitment in the market. We then describe the data we assemble to address our key research questions.

2.1 Setting: Individual insurance

The individual insurance market serves households who do not receive coverage through an employer or through public insurance programs. Prior to the passage of the Affordable Care Act in 2010 and its implementation in 2014, regulation of the market varied by state. In June 2012, for example, California, along with 31 other states, had no rate restrictions on premiums in 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. Since 2014, insurers must issue plans to all consumers who apply during an annual open enrollment period, without regard to a consumer's past illnesses or past attrition.[4] 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.

In the years of our main sample, the ACA also required uninsured households to purchase coverage or face a tax penalty.[5] To ensure plans were affordable to low-income households, the law provided subsidies to cover premium and out-of-pocket costs, with the level of subsidy varying based on a household's income and the specific plan they choose.[6]

[4]Insurers 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 consider enrollment decisions in the open enrollment period separately from those occurring after qualifying life events.

[5]The 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.

[6]In detail, the premium subsidy design sets a cap on how much of a household's income must be spent to enroll in

Finally, in our demand analysis, we benefit from a particular market regulation that requires plan offerings to 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).[7] 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, Forthcoming). 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.

2.2 Data

To study the prevalence of drop-out and then model the consequences of early exit for market outcomes, we need data that capture (a) insurance enrollment, including plan choices and the decision to exit, and (b) household health spending before and after enrollment. We collect these variables from two sources: administrative data from California's insurance marketplace and micro-data on consumers' financial transactions. We describe each source in turn.

Data from the Covered California Marketplace. Our first data source is the official enrollment data from the Covered California Marketplace (CCM), obtained from the marketplace under the California Public Records Act. The records we use for analysis cover households that purchase private health insurance through the CCM from 2014 through 2019, totaling roughly 4.2 million unique households over those years. We observe roughly 1.1 million households enrolling on average each year, with 6.5 million total household-enrollment years.

We observe household-month level information on enrollment, gross and net (of subsidies) premium payments, and household demographics, including household size, the ages of household members, income measures, and geographic market within California. Using the gross and net premium amounts, we classify households as unsubsidized in a given year when their net premium equals their gross premium. We report summary statistics on the key variables from this dataset for the 2015 plan year in Panel A of Table 1.

Given dates of enrollment in the data, we construct variables capturing the month of sign-up, month of drop-out, and the duration of coverage.[8] To be counted as an enrollee, we require that

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.

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

[8] We take the month of enrollment to be the month recorded in the administrative data, with one exception: in cases when enrollment occurs during the open enrollment period for coverage in the subsequent year (typically in

a household makes at least two premium payments.[9] Finally, we use the month of enrollment to classify households into "open" vs. "off-cycle" enrollment periods during the plan year. For some later analyses, we focus only on households who enroll during the open enrollment period, since this population is most likely to be subject to the enrollment penalties we design.

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.

We obtained a random sample of 9.2 million "active user" households in the United States.[10] 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.

To match our enrollment data from Covered California, we narrow our focus in the bank data 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 799,851 "non-mover" California households.[11]

These data have three benefits in our empirical analyses: we observe enrollment in individual health insurance; we observe health spending before, during, and after coverage; and we observe income and other types of spending over the same time period. We discuss each in turn.

First, as in the CCM data, we can identify enrollment in individual health insurance. Unlike households who obtain insurance through an employer, for example, where premium payments would flow through the employer's payroll, enrollees in individual insurance must pay premiums to insurers. We isolate records for premium payments via debit and credit cards, and apply the same definitions of sign-up and drop-out as we do in the CCM data.[12] These data do not include enrollees

November and December), we take the month of enrollment to be January of the subsequent year. We take the month of drop-out to be the month before the administrative recorded exit from coverage. Appendix A provides additional details on the construction of our enrollment variables.

[9] CCM actuaries and staff noted that mistaken enrollments or refunded enrollments might appear as one month premium payments. 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 from our sample.

[10] Our data provider defines an active user as a client with frequent transactions.

[11] We 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.

[12] Helpfully, in the period of our data in California, subsidized enrollees in Covered California were required to pay, at a minimum, a premium of $1 per member per month. As a consequence, even heavily subsidized households will appear in our enrollment counts if they paid this monthly amount via a bank transaction.

who use physical checks or cash to pay their premiums; further, they do not include households that pay premiums through banks that are not in our data. Relative to the CCM data, these limitations mean we do not see the universe of Covered California enrollees. However, unlike data from the marketplace, in the bank data we observe enrollees who buy individual health insurance outside of the marketplace, i.e., directly from insurance carriers or via brokers. Off-marketplace enrollment is substantial: in the years 2014 through 2019, between 30 and 50% of all individual insurance enrollment occurred outside the official marketplace. The two sources of enrollment data are thus complementary, and we will use both in our analyses.

A second key benefit of the bank data is that we can construct measures of out-of-pocket health spending before, during, and after enrollment in individual insurance. Specifically, we use a machine learning algorithm that exploits the text description of each transaction to classify it as "health", "drug", or "non-health".[13] This categorization of out-of-pocket health and drug spending is surely less precise than the patient costs one might observe in individual claims data; however, the bank data offers at least one key advantage: Even if we were to obtain all-payer claims data from a particular state, we would miss out-of-pocket spending completed during periods in which (a) the household is uninsured or (b) when it obtains coverage under self-insured large group plans that are not required to submit claims to all-payer datasets.[14]

A third strength of our bank data is that it allows us to, in a similar fashion, create measures of other types of household spending, as well as household income, to better understand household behavior and potential liquidity issues around periods of insurance participation. We define "other spending" as total spending minus health insurance premium payments, drug spending, and health spending. From this broad spending category, we further break out specific periodic expenses, such as payments for auto insurance, utilities, and video streaming services.[15] 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. We then use our income measure to impute each household's probability of Medicaid eligibility.[16]

Panel B of Table 1 summarizes our key spending variables. The median household in the insured group has monthly income of $4,774 and makes a total of 34 transactions per month. On average, an insured household conducts 0.31 out-of-pocket health transactions per month, resulting in an out-of-pocket spending of $23 on health care.

[13]For this procedure, we first create a training dataset. 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, such as lifestyle activities including fitness, yoga, and spas; insurance premiums that contain health-related words; and dental and vision transactions. We provide more detail on the algorithm in Appendix B.

[14]After the 2016 US Supreme Court ruling in Gobeille v. Liberty Mutual Insurance Co., self-funded or self-insured employers no longer must adhere to a state's requirement to send insurance claims to their state's all-payer dataset.

[15]Appendix Figure A1 illustrates the top 100 words identified in each of the main spending categories.

[16]See Appendix E for details about this imputation.

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.

Validation of Transactions Data. Because of the novelty of our transactions data, we validate our key measures before conducting our analyses. Where available, we collect similar variables observed in public data sources. We briefly summarize these validation exercises here.

First, we calculate the market share of each insurer by year using (a) the plan premium payments we observe in our transactions data and (b) the enrollment reported in both CCM data and federal data. 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 total individual market enrollment for each insurer in California in each year. Insurers must submit these enrollment counts by year to the federal government to comply with rate review 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 measures of household income. Figures 1(b) and 1(c) show that the two income distributions generally match well; for example, the correlation between median 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.[17]

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.[18] Appendix Table A3 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

[17] As 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.

[18] 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.

transactions data also match the quantiles of the MEPS spending distribution. We describe this comparison in detail in Appendix C.

Fifth, we validate our measure of drop-out in several ways, to address the concern that if a consumer switches from paying her premiums on a card we observe in the data to a card we do not, we will falsely interpret the switch as drop-out. We take several steps to lessen this concern. To start, we take the sample of households that is the least likely to suffer from this type of bias. We describe this procedure in Appendix A; in brief, we drop households with inactive accounts, as they may be more likely bank elsewhere. To gauge the likelihood of payment account switches, we collect data from the 2014 Survey of Consumer Payment Choice, housed at the Federal Reserve Bank of Atlanta. The median consumer in the survey has only one checking account, one debit card, and one credit card. In the lowest income groups (less than $20,000), the median consumer carries only a debit card. Next, in Section 3.1, we directly validate the drop-out rate in the bank data by showing that drop-out rate is similar in our administrative CCM data. Finally, in Section 3.2, we explore how sign-up and drop-out correlate with other types of card spending; this allows us to rule out more directly that plan exit is an artifact of consumers' credit card usage.

In sum, the fact that our transactions-derived measures of insurer market share, income, overall household spending, out-of-pocket health spending, and drop-out match statistics in comparable public surveys gives us confidence in the reliability of our transaction classification and our sampling procedure.

3 Evidence on Commitment

This section establishes two key economic facts that emerge from these data: there is substantial drop-out; and patients' out-of-pocket health spending co-moves with drop-out.

3.1 Fact 1: Substantial drop-out

Figure 2(a) plots the distribution of drop-out across months for households enrolling in coverage for the first time during the 2014 open enrollment period. In both CCM and bank data, approximately 35% of the enrollees drop out over the plan year. The rate of exit is relatively flat across the year.

Figure 2(b) shows this drop-out behavior, for the same enrollment period, broken out by income level. We use our bank data in this figure, as we do not observe as granular a measure of income in the CCM data. By January 2015, only 55 - 65% of 2014 open enrollment enrollees maintain their coverage. While poorer enrollees drop out at slightly higher rates, a similar pattern of exit exists in all income groups, including households with annual post-tax income between $100,000 and $200,000.

Figure 2(c) shows the six-month drop-out share for the 2014 open enrollment period, as well as in the pre-ACA period and among 2015 open enrollees. The drop-out rate is roughly similar among 2014 and 2015 open enrollees, while both years show substantially higher drop-out than in the pre-ACA period.

In sum, the raw data presented in Figure 2 suggests widespread drop-out from individual ACA plans. The fact that we observe these patterns in both administrative data and our bank transactions data is significant, given possible concerns that consumers who carry multiple credit and debit cards might switch their card use, which would appear as a change in enrollment in the bank data. As we discuss in Section 2.2, we exclude inactive accounts from our sample to assuage this concern.[19]

3.2 Fact 2: Higher health spending during periods of insurance coverage

Our second key empirical fact is that consumers who drop out appear to conduct more of their annual health spending in the early months of the enrollment year. We illustrate this fact in several ways. First, in Figure 3, we plot the overall relationship between the decision to exit and the fraction of spending that a household conducts in the first four months of the year. In the 2015 plan year, holding income, region, and the identity of the household's insurer fixed, households who drop coverage during these early months also carry out roughly 4% more of their health spending in those months.

Next, we conduct an event study around the period of coverage for households that drop out before the end of the plan year. Figure 4 (a) plots health spending over time, with all months of coverage appearing in the period labeled zero to emphasize trends prior to sign-up and after drop-out.[20] In the post-ACA period, health spending increases at sign-up and decreases at drop-out. This pattern is unique to the post-ACA period—in the pre-ACA period, health spending remained high after exit from individual coverage.[21]

Table 2 presents regression estimates from an event study specification that allows for separate effects from sign-up and drop-out. The post-ACA pattern of increased health spending during the period of coverage appears across income groups. For example, in the population with post-

[19]We provide further supporting evidence in Appendix Figure A3. In the figure, we compare the six-month drop-out rate in our main sample—as shown in Figure 2(c)—as well as in a subsample of our transactions data where we use credit card payment information to identify households unlikely to have an unlinked credit card. We cannot reject that the drop-out rates in the main sample and the sample of households without unlinked cards are the same.

[20]In this analysis we restrict the sample to 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.

[21]The distinct patterns in the two periods suggests that, prior to the ACA, individual coverage might have served as a kind of "weigh station" between employment spells: households would purchase individual insurance after losing access to employer-sponsored coverage. Expecting more generous benefits in later periods, the household may have spent less on health care during the period of individual coverage relative to employer coverage. In contrast, after the ACA, spending appears to increase during the individual coverage time. The generosity of coverage either before sign-up or after drop-out may have been lower.

tax income between $\$20,000$ and $\$40,000$ per year, the number of health transactions increased 24.5% upon sign-up and fell 24.7% upon drop-out. A similar pattern of changes, though sometimes without statistical significance, appears for both lower income and higher income households.

The fact that health spending is higher during the period of insurance coverage motivates us, in the next section, to model the decision to drop out jointly with the decision of when to incur discretionary health spending. Before introducing the model, however, we use the rich information in our transactions data to assess potential alternative and more "mechanical" reasons underlying the empirical fact that health spending is higher during the period of insurance coverage, including (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.

Figure 4(b) displays the evolution of income around sign-up and drop-out. In the pre-ACA period, we see large percentage changes in income before and after the coverage period, which points to job changes or affordability as drivers of participation decisions. After the ACA, in contrast, individual market enrollees exhibit a substantially smaller 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 insurance participation decisions after the implementation of the ACA.

Next, to examine the role of liquidity constraints, we plot the evolution of overall *non*-health spending in Figure 4(c). Similar to the patterns observed for income, in the pre-ACA period we observe sharp increases in non-health expenditure immediately before insurance enrollment and sharp decreases upon drop-out. After the ACA, however, the correlation is weaker: While spending does increase by about 10% upon sign-up, there is there no drop in spending after drop-out. This suggests that liquidity constraints may have been a more important driver of participation decisions in the period before the ACA took effect.

To shed light on the role of behavioral biases, such as forgetting to pay one's credit card or swapping cards, we plot various periodic expenditure categories: video streaming services (Figure 4(d)), utilities (Figure 4(e)), and auto insurance (Figure 4(f)). 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. In addition, the pattern of spending on periodic expenses provides some evidence to support our measure of drop-out. As described earlier, one alternative interpretation of our drop-out measure is that it reflects households switching recurring monthly charges to a bank card that we do not observe in our data. Under this mechanism, we would expect the last premium payment to coincide with a fall in other recurring monthly charges. In Figure 4, we do not see similar changes in other recurring charges around drop-out, suggesting that bank card switching may play a minor role in our measure of drop-out.

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 U.S. 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.

Appendix Figure A5 shows the breakdown of these expenditure and income patterns separately around entry into, and exit from, individual market enrollment. The figure shows a pronounced reduction in health spending upon drop-out and illustrates that this drop is substantial compared to the smoother evolution of income and other spending categories at the time of drop-out.

4 Participation and Welfare

Two clear empirical facts emerge from these initial analyses. First, there is substantial early drop-out from individual health insurance plans. Second, households that purchase insurance and later drop out appear to 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 begin with a simple conceptual framework, in the spirit of Einav et al. (2010), to highlight the possible welfare consequences of drop-out in an insurance market with one-sided commitment. We show that when consumers are allowed to drop out and the market also features traditional adverse selection, drop-out need not always harm welfare. The welfare outcome is ambiguous, and depends on (a) the ease with which households can re-time expenditures and (b) the level of health spending by those households who choose to drop coverage early or choose to be uninsured.

Given the theoretical ambiguity, we then build a structural model to measure the welfare consequences of drop-out in our empirical setting. Our model features both endogenous household exit from insurance coverage and endogenous timing of health care consumption. With estimates of the model parameters, we predict how drop-out rates and consumer welfare would change if regulators imposed penalties for early drop-out or banned early exit.

4.1 Conceptual Framework

We consider a market in which a household chooses, at the beginning of the year, whether to enroll in insurance or remain uninsured. When a household enrolls in coverage, the premium is fixed for a year but paid in twelve monthly installments. Once a household enrolls in the plan, it chooses each period, say a month, whether to continue paying for insurance for that period or to lapse and lose coverage for the remaining months of the year.

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:

$$U_{j,i} = \pi_i - \phi_j p, \tag{1}$$

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. With this formulation, our simple framework allows for both traditional adverse selection and for adverse selection from the 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:

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

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

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

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:

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

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.

To simplify our discussion of equilibrium and welfare outcomes below, we consider a setting with only two types of households. There is a mass N_D of "drop-out"households who can accomplish all of their required health spending within the first $\phi < 12$ months of the year. There is a mass N_F "full-year" type households who cannot re-time their consumption. That is, $\phi_F = 12$.

4.1.1 Equilibrium without traditional adverse selection

To highlight the welfare consequences of drop-out, we first consider a setting without traditional adverse selection. Here, $c = c_D = c_F$. Without drop-out or adverse selection, 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.

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. Because drop-out types pay fewer monthly installments while concentrating their spending c into the subset of months for which they pay, drop-out types are effectively less sensitive to the monthly premium. At high enough monthly premiums, only drop-out types purchase insurance. As premiums fall, some full-year types choose to purchase insurance as well. This purchase behavior generates a kink in the demand curve.[22] 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.

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. 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.

4.1.2 Equilibrium with traditional adverse selection

The outcome becomes less clear when we allow both traditional adverse selection and drop-out. Suppose, for example, that households who can re-time their consumption have lower annual expenditures as well as lower willingness to pay for insurance than those who cannot. That is, $c_D < c_F$ and $G_F(\cdot)$ has first-order stochastic dominance over $G_D(\cdot)$. Then, we can show that allowing drop-out can increase social surplus relative to a drop-out ban.

To understand this outcome, we consider a special case where $\phi/12 = c_D/c_F$—that is, where the fraction of the year that drop-out types enroll is equal to the ratio of the costs between the

[22] We illustrate this outcome in Appendix Figure A8.

drop-out and full-year populations. Here, a monthly premium of $p = c_F/12$ 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 and their annual revenues equal $\phi * p = \phi/12 * c_F = c_D$, implying that insurers also break even on these enrollees. Moreover, since the implied annual premium for non-dropouts is c_F 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.[23] Intuitively, without the option to drop coverage mid-year, drop-out types 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 can encourage these low-cost households to enroll, bringing down the average monthly cost in the insurance pool.

4.2 Structural Model

In our conceptual model, the presence of drop-out type consumers could generate an equilibrium with greater or lower overall welfare, depending on the relative health care needs of drop-out households and the ease with which they could re-time their health expenditures across months of the year. We now develop a structural model that endogenizes the household's plan choice, drop-out decision, and allocation of spending over the year. Combining enrollment data from Covered California and spending patterns from our bank data, we estimate the relevant parameters of households' spending and participation decisions to quantify the welfare effects from policies that discourage drop-out.

4.2.1 Period-specific value functions

To illustrate the household's dynamic considerations, we start by specifying period-specific value functions. The value functions depend on two key parameters. First, we specify H to capture the normalized total annual level of health spending for a household, including those costs the plan covers as well as the household's out-of-pocket costs. H reflects non-discretionary health care services that the household anticipates consuming over the year. We assume this level does not change with insurance coverage and thus we rule out moral hazard. The household can divide its

[23] We describe a range of possible equilibrium outcomes in Appendix G.

total spending H over a finite horizon of T equally sized periods of enrollment. We denote optimal health care consumption in each period, h_t, as a fraction of H, $f_t = h_t/H$. Here, $\sum_t^T f_t = 1$.

Our second key parameter, ν, dictates the degree to which a household can substitute its health spending needs from one period of the year to another. As ν approaches 0, the household faces little penalty from shifting its spending away from a smooth consumption path. To make this penalty explicit in the model, we add a linear term to the utility over plan options that specifies a (dollar-equivalent) penalty as a function of the average squared deviations around the smooth consumption path[24]:

$$Penalty = \nu H \left[\sum_{t=1}^{T} \left(f_t - \frac{1}{T} \right)^2 \right]$$

In our empirical implementation, we allow (ν, H) to vary by demographic group, characterized by household size, income, and the age of the head of household. We also allow these parameters to vary by year and the geographic market in which the household lives. Finally, to simplify the empirical implementation, we set $T = 3$ representing three periods over the plan year, each with four months of enrollment. We describe the value functions below starting with the terminal period, period 3, and then moving backwards to period 1.

Period 3. We write the household's value from choosing plan j in period 3 as a function of several components: choice-specific and time period-specific fixed effects, $(\delta_{1j}, \delta_{2j}, \delta_{3j})$, that affect utility for plan j in each consumption period; the household-specific actuarial value of a plan, α_{ij}, multiplied by the household's spending needs, H_i; the per-period premium, P_{ij}; the penalty for non-smoothed consumption; and, a shock, ε_{itj}, that is specific to the household, time period, and choice, but which the researcher does not observe. We use the parameter β_3 to convert dollars of health care spending and dollars of premiums into utility units. We normalize σ_3, the variance of ε_{i3j}, equal to 1. Dropping i subscripts for ease of notation, the value of choice j, V_{3j}, and the value of the outside option, V_{30}, equal:

$$V_{3j} = \underbrace{\delta_{1j} + \delta_{2j} + \delta_{3j} - \nu H \beta_3 \left[\sum_{t=1}^{3} (f_t - 1/3)^2 \right] - \alpha_j \beta_3 H - 3\beta_3 P_j}_{v_{3j}} + \varepsilon_{3,j}$$

$$V_{30} = \underbrace{\delta_{1j} + \delta_{2j} - \nu H \beta_3 \left[\sum_{t=1}^{3} (f_t - 1/3)^2 \right] - \alpha_j \beta_3 H (f_1 + f_2) - \beta_3 H (1 - f_1 - f_2) - 2\beta_3 P_j}_{v_{30}} + \varepsilon_{3,0}$$

Here, V_{30} illustrates the consequences of dropping out in period 3: the household saves one period's worth of premium payments, but then must pay the full cost of the remaining health spending,

[24]One advantage of this functional form is that when we take the first order conditions of the value functions in each period, treating the drop-out probabilities as data, the resulting expressions are linear.

$H(1-f_1-f_2)$.

Given that $\varepsilon_{j,3}$ follows a Type 1 extreme value distribution, the probability of drop-out takes the following form:

$$\psi_{3j} = \frac{1}{1+\exp\left(v_{3j}-v_{30}\right)} \tag{5}$$

$$= \frac{1}{1+\exp\left(\beta_3\left(1-\alpha_j\right)H(1-f_1-f_2)+\delta_{3j}-\beta_3 P_j\right)} \tag{6}$$

Period 2. Shifting back to period 2, we can define the value of coverage through insurance plan j in period 2, given the decision to drop or maintain coverage in period 3. When the household chooses optimal period 2 health consumption $h_2 = Hf_2$ under plan j, we can write the value from maintaining coverage as below:

$$V_{2j} = \underbrace{\arg\max_{h_2} \log\left[\exp\left(v_{30}\right)+\exp\left(v_{3j}\right)\right]}_{v_{2j}} + \sigma_2\varepsilon_{2,j}. \tag{7}$$

where we define v_{3j} and v_{30} above. The value of not having insurance in period 2 is:

$$V_{20} = \underbrace{\delta_{1j}-\beta_3 b_{20}-\alpha_j\beta_3 Hf_1-\beta_3 H\left(1-f_1\right)-\beta_3 P_j+E\left[\varepsilon\right]}_{v_{20}} + \sigma_2\varepsilon_{i20}$$

where

$$b_{20} = \nu H\left[\left(f_1-1/3\right)^2+2\left(\frac{1-f_1}{2}-1/3\right)^2\right]$$

As in period 3, if the household exits early, it saves the cost of premiums in later periods but loses the insurance contribution. Here, the household pays one period of premiums, P_j, and the insurer covers $(1-\alpha_j)Hf_1$ dollars of health spending. The household must pay $\alpha_j Hf_1 + H(1-f_1)$, its out-of-pocket share of period 1 spending as well as the entire cost of spending in periods 2 and 3. In addition, if the household chooses to concentrate more of its spending into the period with coverage, it will face a penalty, b_{20}. Any deviation from a 1/3 share in period 1 incurs a penalty that scales with the household's cost of substitution, captured by ν. We assume the household splits its uninsured spending in periods 2 and 3 equally, thereby minimizing the penalty in periods in which it has already dropped coverage.

Given the distributional assumption on ε_{tj}, we can write the share of households who drop coverage in period 2 using a logit form:

$$\psi_{2j} = \frac{1}{1+\exp\left(\frac{v_{2j}-v_{20}}{\sigma_2}\right)} \tag{8}$$

Period 1. Finally, in period 1, we can write the value of choosing j or the outside option (0) using the expressions from periods 2 and 3:

$$V_{1j} = \underbrace{\arg\max_{h_1} \log\left[\exp\left(v_{2j}\right) + \exp\left(v_{20}\right)\right]}_{v_{1j}} + \sigma_1 \varepsilon_{1j}. \tag{9}$$

The value of the outside option is:

$$V_{10} = \underbrace{-\beta_3 H + \left(\sigma_2 + \sigma_3\right) E\left[\varepsilon\right]}_{v_{10}} + \sigma_1 \varepsilon_{10}$$

where $\sigma_2 = 1/\beta_2$ and $\sigma_3 = 1$. The probability that a plan j is chosen is given by:

$$\psi_{1j} = \frac{\exp\left(v_{1j}/\sigma_1\right)}{\exp\left(v_{10}/\sigma_1\right) + \sum_{j'} \exp\left(v_{1j'}/\sigma_1\right)}. \tag{10}$$

4.2.2 First-order conditions

As we show later in the estimation section, to recover our parameters of interest, we first need to find the optimal share of annual spending that a household conducts in periods 1 and 2 as a function of plan characteristics and the household's cost of re-timing expenditures. We can do so by finding the first order conditions of the value functions in period 1 with respect to f_1 and in period 2 with respect to f_2. We show the steps to find these first order conditions in Appendix Section H.1. Here, we report the results.

For period 2, we find:

$$-\nu\left[2\left(f_2 - 1/3\right) - 2\left(1 - f_1 - f_2 - 1/3\right)\right] + \psi_3\left(1 - \alpha_j\right) = 0 \tag{11}$$

For period 1, we find:

$$\begin{aligned} \psi_{2j}\left\{-\nu\left[2\left(f_1 - 1/3\right) - 2\left(\frac{1 - f_1}{2} - 1/3\right)\right] + 1 - \alpha_j\right\} \\ + \left(1 - \psi_{2j}\right)\left\{-\nu\left[2\left(f_1 - 1/3\right) - 2\left(1 - f_1 - f_2 - 1/3\right)\right] + \psi_{3j}\left(1 - \alpha_j\right)\right\} = 0 \end{aligned} \tag{12}$$

Thus, we have two first order conditions that are a function only of the parameter ν and observed drop-out shares, plan actuarial values, and the shares of health spending in periods 1 and 2, (f_1, f_2).

4.3 Estimation

To estimate our model, we need data on (a) household plan choices (b) household drop-out rates (c) plan premiums and actuarial values, and (d) the share of annual health spending that households consume at each time period over the year. For the first three data elements, we rely on data collected directly from Covered California. We observe the plan choice of each household, the months of coverage they maintain in a year, and demographic details that allow us to compute the premium and actuarial value that a household would face under all plan options in their geographic market. For the share of annual spending a household consumes in each period of the year, we rely on our detailed bank data, which contain the timing of consumers' health transactions over the year.

Our estimation approach proceeds in three steps. In the first step, we recover ν by combining our model with data on the timing of transactions and drop-out by plan and demographic group. Second, given an estimate of ν, we recover (β_2, δ_{2j}) and (β_3, δ_{3j}) for all $j \in J$ using logistic regressions run on the drop-out shares in periods 2 and 3, respectively. We follow the approach of Tebaldi (Forthcoming) to account for endogeneity affecting the estimation of the premium sensitivity parameter. Finally, third, we identify (β_1, δ_{1j}) along with H, the household's annual expected health spending level, by running a conditional logit on observed plan choices. We describe each step in turn.

First stage: solving for ν To recover ν, we exploit variation across demographic groups in their observed propensity to (a) exit coverage prior to the end of the year and (b) re-time their annual health spending. Our model relates the fraction of spending that each demographic group re-times into the first period of coverage with the rates of drop-out for plans of different generosity. We define demographic groups by enrollees' ages, income as a share of the federal poverty level, and family size.

The relationship between household spending in period 1 and drop-out rates can be seen through the first order conditions in equations 11 and 12. Solving this system of equations for f_1, the fraction of spending completed in period 1, we find:

$$f_1 - 1/3 = (1 - \alpha_j) \frac{2\psi_{2j} + (1 - \psi_{2j})\, \psi_{3j}}{6\nu}. \tag{13}$$

The expression provides some intuition about ν. All else equal, we see that as plans increase $(1-\alpha_j)$, covering a greater share of health spending, households will optimally shift more of their spending to period 1. Similarly, if particular demographic groups find it easier to re-time expenditures—for example, younger households with fewer chronic conditions or richer households with greater liquidity—we will see a lower ν and a larger share of spending in period 1.

With data on f_1 by plan and household demographic bin, we could run a regression to recover ν

specific to a demographic bin, using variation in plan generosity, $(1-\alpha_j)$, and different exit rates by plan. However, our bank data only directly identifies a household's insurer, not their specific plan. Thus, we aggregate the model predictions to the carrier level to match to our observed spending data.

To illustrate our aggregation, we introduce additional notation. We denote a geographic region by r and an insurer by ι. We label distinct demographic groups by age, income, and family size using z. Then:

$$f_{1r\iota} - 1/3 = \sum_z \frac{\varsigma_{zr\iota}}{\nu_z} \left[\sum_{k \in \mathcal{K}_{r\iota}} \psi_{1zk|r\iota} \left(1 - \alpha_k\right) \frac{2\psi_{2zk} + \left(1 - \psi_{2zk}\right) \psi_{3zk}}{6} \right],$$

Here, $\varsigma_{zr\iota}$ is the fraction of insurer $\iota's$ enrollment by demographic z in market r. The share $\psi_{1zk|r\iota}$ is the probability of demographic z choosing one of the $\mathcal{K}$ plans that insurer ι offers.

Second stage: identify $(\beta_2, \delta_{2j}, \beta_3, \delta_{3j})$ Given our model and our estimate of ν by demographic group z, we can recover the choice and period-specific fixed effects, $(\delta_{2j}, \delta_{3j})$ for all plans $j \in J$ in periods 2 and 3. We can also recover the price sensitivity, β_3, and period 2 scale parameter, $\beta_2 = \frac{1}{\sigma_2}$. To do so, we carry out two logistic regressions.

For period 2, with data on period 2 drop out decisions, ψ_{2ij}, we run a logistic regression to recover (β_2, δ_{2j}), following Equation 8:

$$\psi_{2ij} = \frac{1}{1 + \exp\left(\beta_2 X_{2ij} + \tilde{\delta}_{2j}\right)}$$

where $\tilde{\delta}_2 = \beta_2 \delta_{2j}$ and we define X_{2ij} as follows:

$$X_{2ij} = \beta_3 \left(b_{2j} - b_{20}\right) - \beta_3 P_{ij} + \left(1 - \alpha_{ij}\right) \beta_3 H f_2 - E\left[\varepsilon\right] + \log\left[\left(1 + \exp\left(v_{3j} - v_{30}\right)\right)\right]$$

with

$$b_{2j} = -\nu H \left[\sum_{t=1}^{3} \left(f_t - 1/3\right)^2 \right],$$
$$b_{20} = -\nu H \left[\left(f_1 - 1/3\right)^2 + 2\left(\frac{1 - f_1}{2} - 1/3\right)^2 \right]$$

In period 3, with data on whether a household drops coverage in period 3, denoted ψ_{3j}, we can run a logistic regression to recover (β_3, δ_{3j}), as in Equation 5:

$$\psi_{i3j} = \frac{1}{1 + \exp\left(\beta_3 X_{3ij} + \delta_{3j}\right)}$$

where:

$$X_{3ij} = (1 - \alpha_{ij})\, H_i(1 - f_{i1} - f_{i2}) - P_{ij}.$$

In our implementation, we allow the household's financial share of health spending, α_{ij}, to account for the household's eligibility for federal cost-sharing subsidies.

Households also face different premiums. For individual market plans regulated by provisions of the Affordable Care Act (ACA), premiums increase with age according to a fixed age curve. However, from actuarial tables, we also know that health spending scales with age, meaning that costs and preferences for plan generosity move mechanically with the variation in premiums. Thus, similar to Tebaldi (Forthcoming), we worry about bias in our estimate of price sensitivity if consumers with higher costs face plans with higher premiums. To address this concern, we adjust our specification to include a variable that acts like a control function—here, we use the household's baseline premium—to proxy for the household's underlying spending needs.

Our estimation then relies on the remaining premium variation that arises due to federal premium subsidies. Under the ACA, the household's income determines the level of the premium subsidy. Once we control for the household's base premium, comparing two households with similar age profiles but different incomes generates sharp differences in the out-of-pocket premiums the households face. In our expression for ψ_{i3j}, we specify δ_{3j} to include both a constant term and our control variable.

Finally, in these logistic regression specifications, we need the share of a household's optimal health spending that it conducts in periods 1 and 2, denoted (f_1, f_2). We can find these implicitly through our model's first order conditions, equations 11 and 12, which depend on the first stage estimate of ν, on the parameter H, and data. Equation 13, for example, illustrates the solution for f_1. We can use equations 11 and 12 to solve for f_2 in a similar manner.

Third stage: identify H Finally, to recover the household's spending level, H, we use the first period choice of insurance plan. Equation 10 shows the probability that the household chooses plan j. We use data on that probability, ψ_{i1j}, to recover (δ_{1j}, β_1) and H.

We estimate the following conditional logit model:

$$\psi_{1ij} = \frac{\exp\left(\beta_1 X_{ij} + \tilde{\delta}_{1j}\right)}{\exp\left(\beta_1 X_{i0}\right) + \sum_{j'} \exp\left(\beta_1 X_{ij} + \tilde{\delta}_{1j}\right)}$$

where:

$$
\begin{aligned}
X_{ij} &= -\beta_3 H \left[\nu \sum_{t=1}^{3} (f_t - 1/3)^2 + \alpha_{ij} \right] - 3\beta_3 P_{ij} - 1/\beta_2 \log (1 - \psi_{2ij}) - \log (1 - \psi_{3ij}) \\
X_{i0} &= -\beta_3 H + (1/\beta_2 + 1)\, E[\varepsilon] \\
\tilde{\delta}_{1j} &= \beta_1(\delta_{1j} + \delta_{2j} + \delta_{3j})
\end{aligned}
$$

We recover H using this logit specification, given estimates of (β_3, β_2, ν) and data on premiums, actuarial value, and the drop-out shares in periods 2 and 3.[25] In Appendix Section H.2, we show the steps needed to derive this conditional logit model in period 1 given our model of behavior in periods 2 and 3.

4.4 Results

We report our parameter estimates from the three stage estimation in Table 3. We report estimates of the re-timing cost parameter, ν, by demographic group z. Here, we find that households with higher incomes have lower ν, meaning they can more easily front-load their health spending in the early part of the plan year. The ν we recover for low income households is at the bounds of our optimization; the implication of the data is that these households cannot easily re-time their spending. Interestingly, among the set of higher income households, those led by household heads below age 40 appear to be able to re-time more easily.

We interpret the finding on poorer households as evidence that liquidity issues might explain at least part of the reluctance to re-time spending. That younger households—both single-membered households and households with spouses and/or dependents—can more easily re-time spending matches actuarial data showing these households suffer fewer chronic conditions that require regular office visits.

In additional columns in Table 3, we report our period and demographic-specific price sensitivity parameters. We also illustrate our estimates of a household's annual spending needs, H, as implied by the household's observed plan choices, drop-out, and spending patterns. We plot both the distribution of model-implied H, and the relationship between H and age in Figure 5.

From the histogram in Figure 5(a), we observe significant dispersion in expected spending levels, where spending includes both costs the insurer covers and household out-of-pocket costs. The level ranges from \$828 to \$18,400, with significant portions of the population expected to spend around \$6000 per year, \$11,000 per year, and between \$12,000 and \$18,000 per year.

Across ages, we observe an increasing relationship between average annual health spending and age.

[25] In particular, we interpret the coefficient on the drop-out penalty element as $\beta_1\beta_3 H$, and use parameters from earlier steps to back out the value of H alone.

We predict the lowest spending among young, single households. Older households, both single and with spouses or families, show the largest annual spending. These results are similar to estimates from Dickstein et al. (Forthcoming), who find spending levels between \$3500 and \$8500 for single households, and as high as \$18,700 for the largest family sizes.

5 Policy Discussion

We can use our model of insurance participation to discuss the design of penalties targeting drop-out. Given that, in theory, the presence of traditional adverse selection can mean drop-out can benefit total welfare by enabling a form of price discrimination, our goal is to measure the welfare change were policymakers to discourage drop-out through penalties.

In our analysis, we ignore "mandate" penalties, which are assessed whenever a household fails to sign up for insurance, say during an open-enrollment period; instead, we analyze 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.

5.1 Equilibrium algorithm

We use our model of plan choice, along with our data and estimated parameters, to simulate an equilibrium in which we add a drop-out penalty, call it m_t. We let m_t vary with the number of time periods the individual enrolls prior to exit.

We start by using the model estimates and data to compute the total premium revenues each carrier receives for plans it offers in the observed setting.[26] We also compute the total implied costs the carrier incurs for the household spending among enrollees. This includes only the actuarial value or fraction of the plan's health spending that the insurer pays; the remaining costs are the household's responsibility.[27] We regress these revenues on costs to recover the parameter φ. Here, $\varphi > 1$ can indicate plan markups. If $\varphi \leq 1$, we assume the plan loses money on premium revenue alone, but breaks even due to additional monetary transfers from the government, including in the form of risk-adjustment, risk-corridor, and reinsurance payments. We hold φ fixed in our counterfactual simulation with drop-out penalties, under the assumption that the transfers or markups scale with health spending costs.

We describe our counterfactual algorithm in more detail in Appendix I. Here, we sketch the equilibrium price search. We start by adding a penalty m_t to the household's period-specific value

[26]The premium payment for a household depends on the underlying premium the carrier sets for the plan, multiplied by the household's rating factor. The rating factor varies with age of each household member, including up to three dependents under age 18.

[27]If a low-income household is eligible for cost-sharing subsidies when it purchases a silver plan, we assume the carrier is nonetheless responsible only for 70% of the costs, with the government paying the carrier back for additional subsidized out-of-pocket costs.

functions, as described in Section 4.2.1, such that the penalty caps the household's probability of drop-out at a specific threshold. We vary this threshold from a probability of 0%—implying a ban on drop-out— to a threshold of 40%, which rarely binds in our sample.

With the penalty in the model, we allow the households to optimally choose (a) the fraction of spending to consume in each time period, $(f_1, f_2, 1 - f_1 - f_2)$ and (b) its preferred plan, including the choice to remain uninsured throughout the plan year. We then compute the total costs the insurer incurs in a plan, which is a function of the expected spending of the enrolled households for the duration of months the household is covered. We then use our estimated φ to compute, for each plan, the predicted premium revenue the plan will collect in total. We divide this revenue by the sum of household rating factors, where we pro-rate the rating factors of each household for the fraction of the year in which the household enrolls and pays premiums. This scaled measure is the new equilibrium "baseline" monthly premium.

We repeat our procedure of predicting choices and spending for each household using this new proposed set of equilibrium premiums. In the household's choice vector, we convert the baseline premium to a household-specific premium by multiplying by each household's rating factor. We also keep the household's premium subsidy fixed; that is, we compute the subsidy amount the household receives for the second cheapest silver plan in the baseline model and provide this subsidy level to the household under the new premiums.

With the predictions from the model, we repeat our procedure to predict baseline premiums. We continue to iterate between computing premiums and determining household choices until we converge to an equilibrium in which baseline premium levels stabilize. During the process of convergence, premiums rise for plans that enroll sicker patients and fall for those that enroll healthier types.

5.2 Welfare under alternative penalties

When we estimate the baseline relationship between costs and premiums, we find φ varies between 0.6 and 0.8, depending on the region and plan. That the estimates lie below one reflects net payments from the government to the insurers. We iterate to convergence under penalties that set caps or thresholds on the probability of drop-out ranging from 0 (banning drop-out) to 0.4, which is rarely binding in our setting.

In Table 4, we report the market outcomes under our alternative thresholds. We report market shares and average monthly premiums by metal tier. We also report the change in consumer surplus for households under each threshold, relative to a baseline threshold that approximates the market outcomes in our sample without penalties.

We find that banning drop-out by requiring a threshold equal to 0 leads a share of healthier consumers, on average, to forego purchasing coverage. In Panel 1 of Table 4, we see the share of "no insurance" under a drop-out ban rises to over 60%, from a baseline of 51%. The drop in share comes mostly from consumers who purchased silver and bronze plans without penalties.

When the healthier consumers exit coverage, the average costs of the insured population increases. Under our assumption of average cost pricing, we see that average monthly premiums also increase. In Panel 2 of Table 4, we report average premiums, scaled to reflect the price for a 40 year-old adult purchasing single coverage. The average is taken across plans and regions, weighted by enrollment. We see that premiums increase across all metal tiers as we increase the penalties for drop-out—i.e. as we decrease the cap on drop-out probabilities toward a total ban at 0% probability. Premiums increase the most for bronze and silver plans, a total of between 12 and 16%. The penalties drive healthy enrollees, who have lower tastes for insurance and who value the option to exit early, to choose to be uninsured; these enrollees disproportionately choose bronze and silver plans in the equilibrium without penalties.

Finally, we report measures of welfare that account for both preferences and changes to plan premiums. We report the change in average consumer surplus in dollar units per month and per household, relative to the average surplus at baseline. We proxy for the baseline using the threshold of 40%, since it rarely binds.

We find that overall welfare falls with penalties on drop-out. We see this pattern across demographic groups that vary in their estimated ability to re-time expenditures, as measured by the parameter ν, and their expected annual health spending, measured by the parameter H.

The surplus losses from banning drop-out particularly affect the population of households with low values of ν, who we estimate can more easily re-time their expenditures. They experience an average decline in surplus of $232 from losing the ability to re-time and exit coverage early; for the consumers who cannot re-time, with high values of ν, we estimate surplus losses of only $75.[28] There is also a population of consumers across demographics that, even without re-timing, see welfare losses from losing the option value to exit early. After a few months of use, for example, they may learn that the insurance has poor coverage for the services they consume. In a world with penalties, consumers can no longer exit early as a function of these poor experiences. The utility of a full-commitment plan may even drop below the outside option of no insurance, leading to our prediction that more households choose to be uninsured under larger penalties.

6 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

[28]In our sample, we estimate 40.4% of households to have a ν parameter with a value we define as "low".

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. Given the ambiguous theoretical effect of drop-out, we develop and estimate a structural model of household plan choice that features both endogenous re-timing of health spending and endogenous drop-out. We use these estimates to predict counterfactual market outcomes were policymakers to institute a range of penalties on households who exit early. We find that consumer surplus falls with the introduction of penalties that curb or ban drop-out entirely. In our setting, healthy enrollees disproportionately choose to be uninsured when they cannot exit early; as a result, average costs and premiums increase in equilibrium.

Settings outside of California might involve distinct optimal drop-out penalties, depending on the correlation between health spending and drop-out as well as the degree of adverse selection in the market. We emphasize that the effect on welfare from penalizing drop-out depends on the expected health spending of those households induced to enter or exit the insurance pool when insurers can enforce an enrollment commitment.

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.

Azevedo, E. M. and D. Gottlieb (2017): "Perfect Competition in Markets With Adverse Selection," *Econometrica*, 85, 67–105.

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.

Dickstein, M., K. Ho, and N. Mark (Forthcoming): "Market Segmentation and Competition in Health Insurance," *Journal of Political Economy*.

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. GEISSLER, 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. (Forthcoming): "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA," *The Review of Economic Studies*.

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

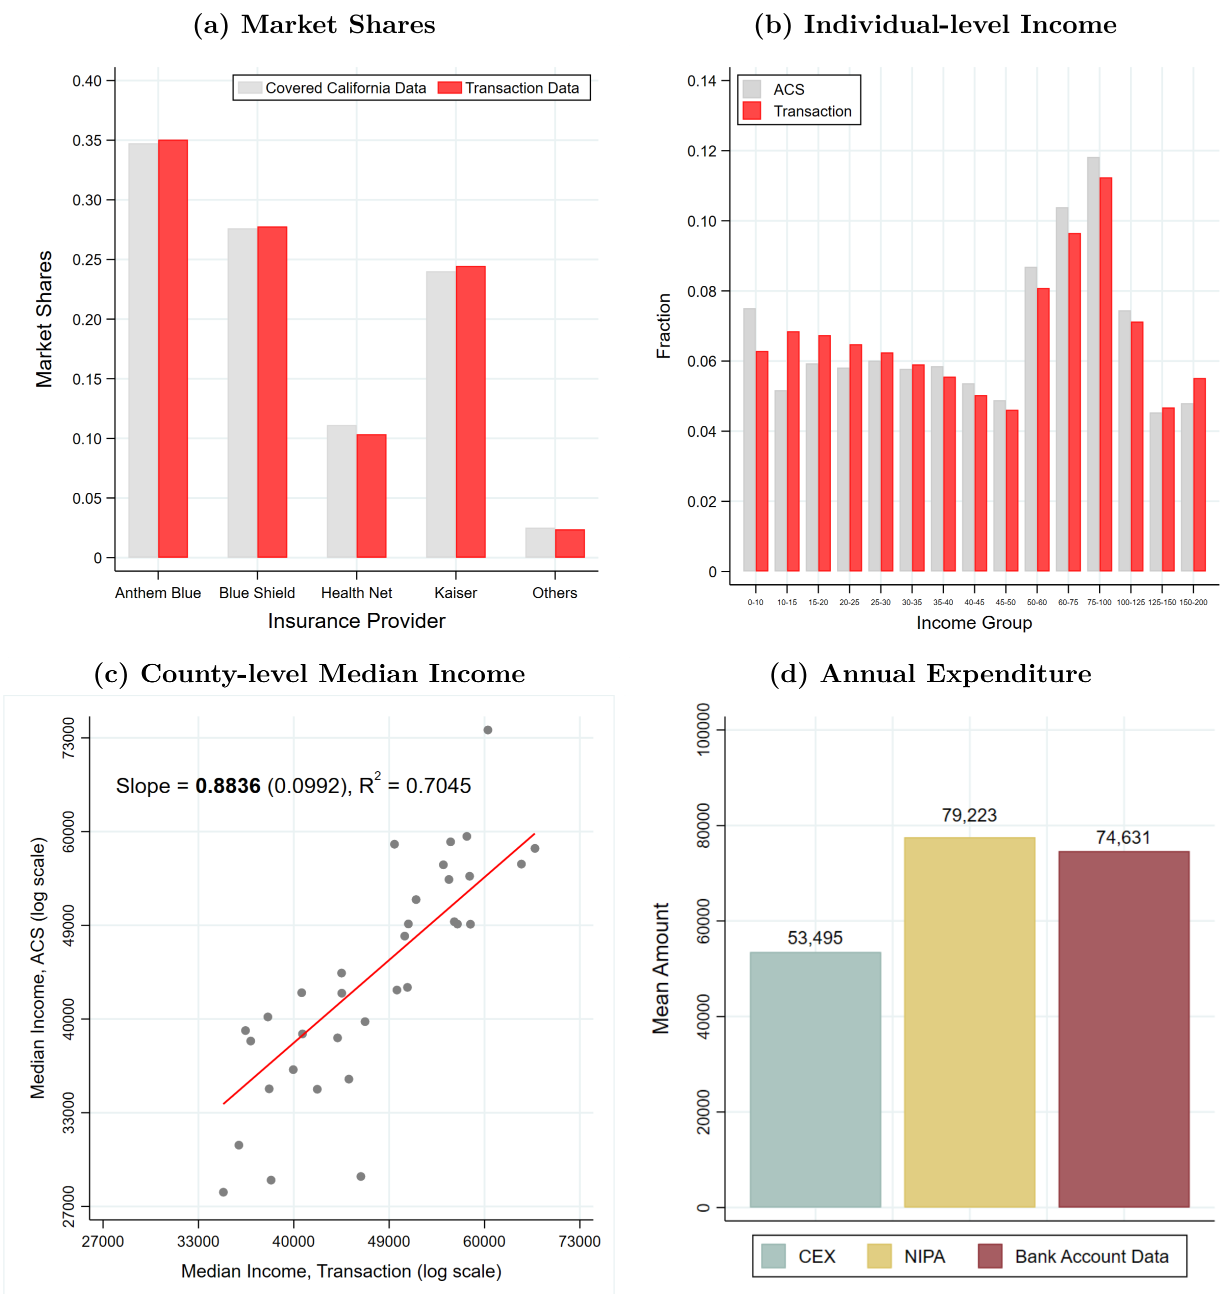

(a) Market Shares

(b) Individual-level Income

(c) County-level Median Income

(d) Annual Expenditure

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**

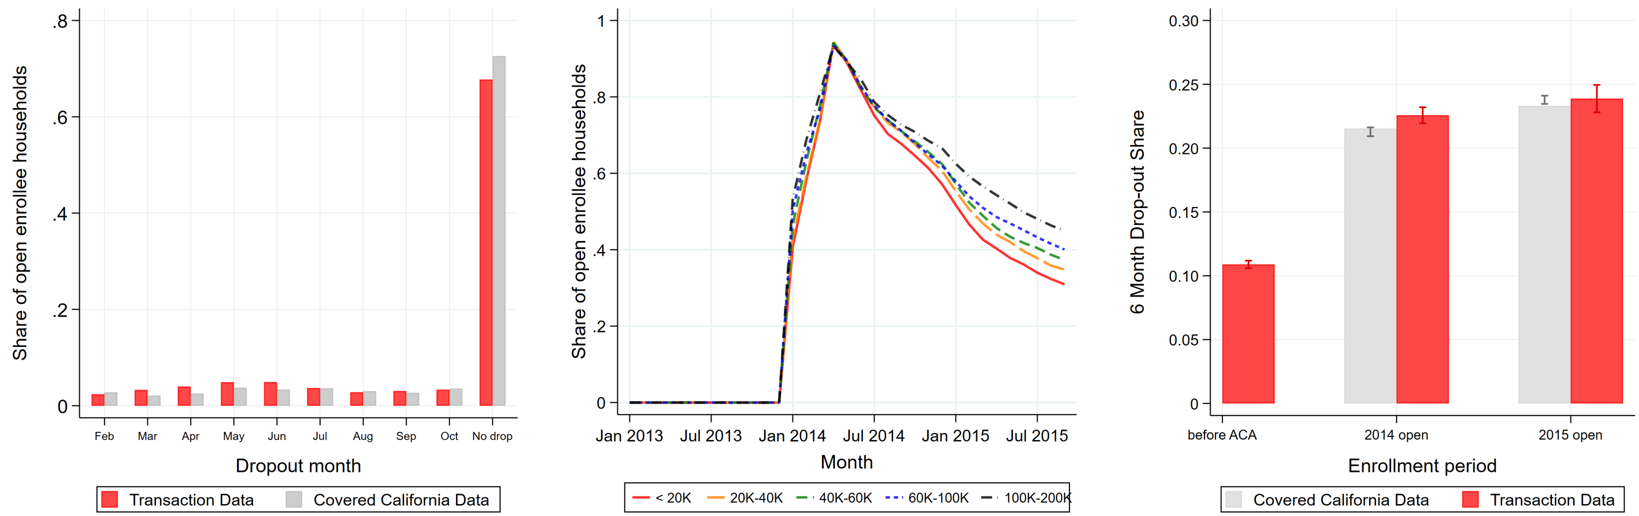

Notes: This figure describes participation in California's individual insurance market in 2014-2015. In **Panel (a)**, we plot the distribution of enrollment months in the financial transactions data and in weighted Covered California marketplace (CCM) data for the 2014 open enrollment period. In **Panel (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: **Relationship between Healthcare Spending and Household Drop-out**

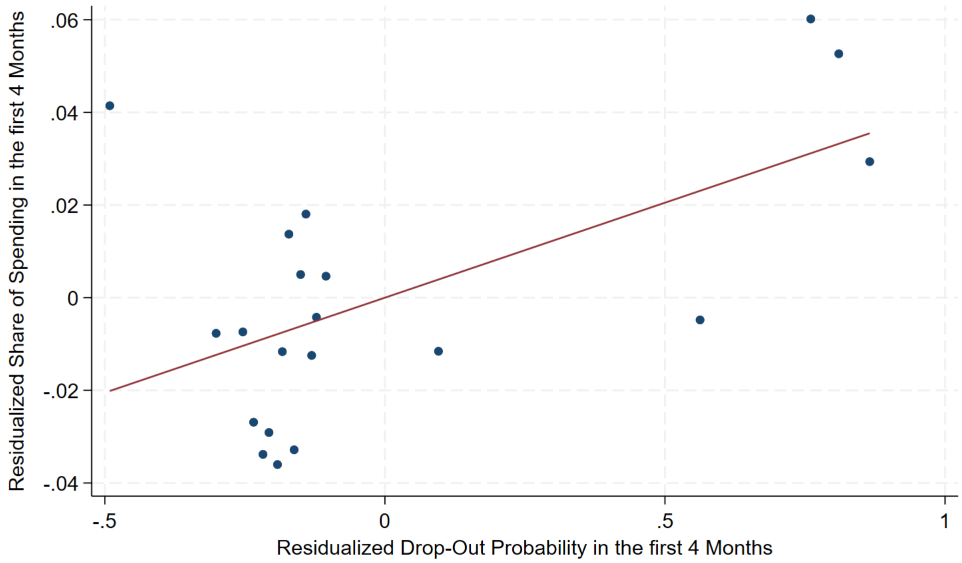

Notes: This figure plots the relationship between (a) the probability that a household drops coverage in the first four months of the plan year and (b) the share of its annual health spending that the household undertakes in the first four months of the year. We plot the residualized first period spending share as a function of the residualized drop-out share, where we form the residualized variables by controlling for the identity of the region, insurer, and for household income. Each dot in the figure represents a 5% bin of the sample, where we order the bins by the residualized drop-out share. The slope of the best fit line shown corresponds to the coefficient of interest when regressing spending share on drop-out probability with the controls mentioned above. The coefficient equals 0.041. Data for the plot comes from the 2015 plan year in California using the bank transactions data

Figure 4: Income and Expenditure Dynamics Around Sign-up and Drop-out

(a) Health Spending **(b) Income** **(c) Other Spending**

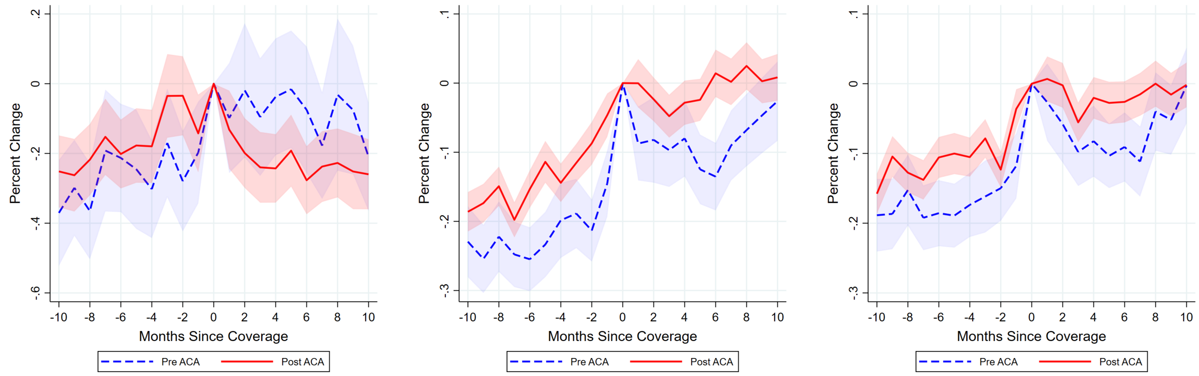

(d) Netflix/Hulu **(e) Utilities** **(f) Auto Insurance**

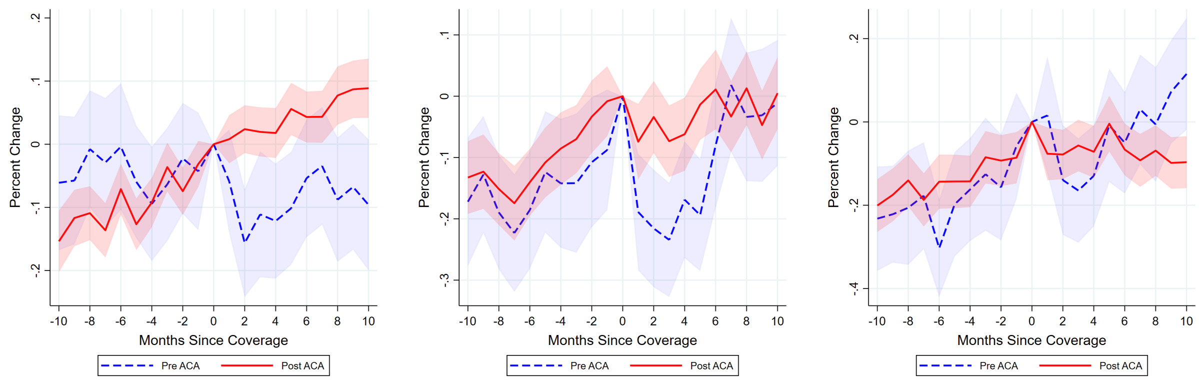

Notes: This figure compares the dynamics of monthly income and various types of monthly dollar spending for households who drop coverage. We examine these dynamics separately for households enrolling in the pre- vs. post-ACA time periods. Other spending, reported in **Panel (b)**, includes all transactions 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 3 (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 5: **Relationship between Age and Model-Implied Household Annual Health Spending**

(a) Distribution of H

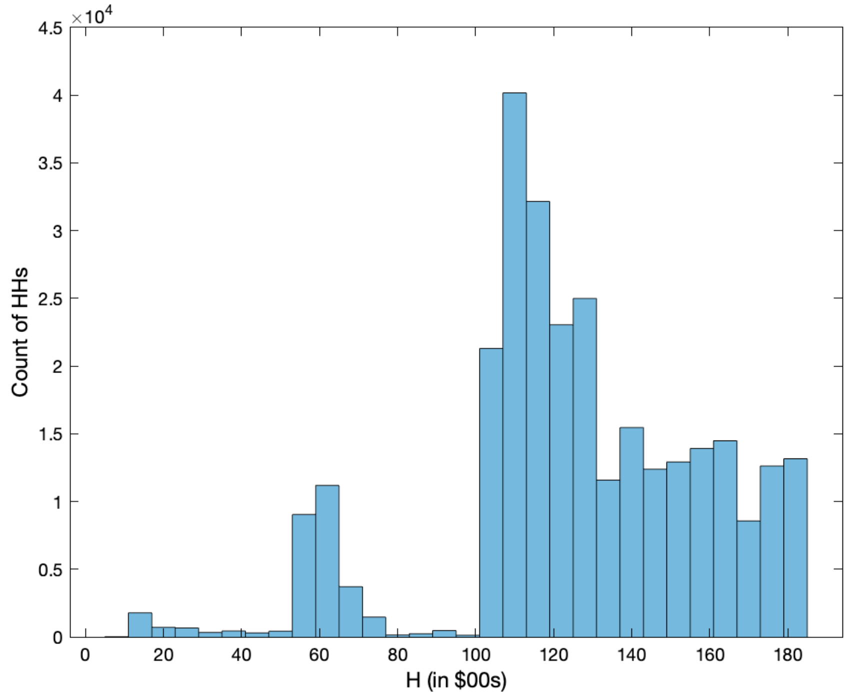

(b) Age vs. H

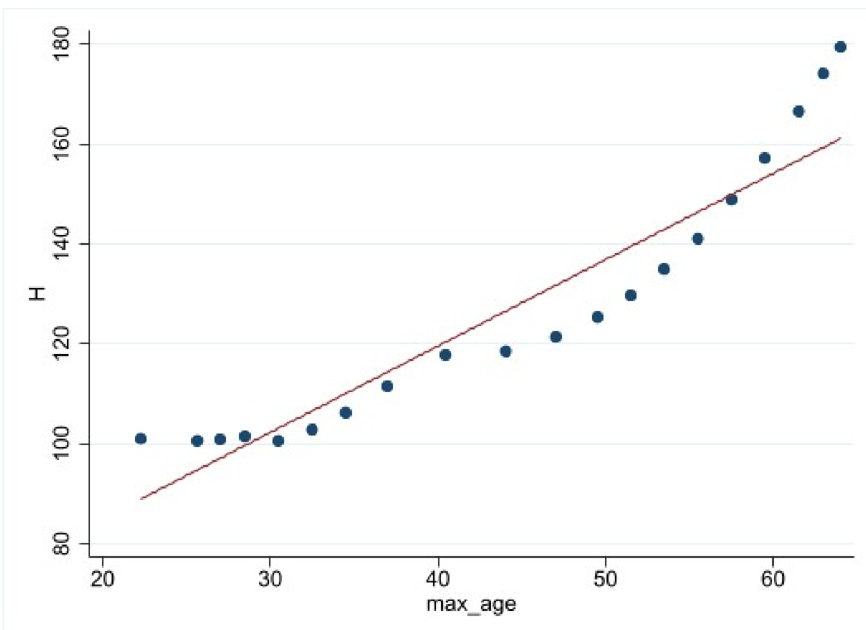

Notes: This figure plots (a) the distribution of H, the model-implied level of the household's annual expected health spending, and (b) the relationship between the age of the household head and H. We estimate H, reported in $100s, via the model described in Section 4. We report the distribution of H and the relationship between age and H for the set of all individual market enrollees, including households of various family sizes and income levels.

Table 1: Sample Summary Statistics

Panel A. Covered California Data Variables

	N=1,309,060			
	Mean	SD	Min	Max
Household size	1.58	0.98	1.00	19.00
Average age	37.69	14.28	0.00	64.00
138% FPL or less	0.05	0.21	0.00	1.00
138% FPL to 150% FPL	0.09	0.28	0.00	1.00
150% FPL to 200% FPL	0.19	0.40	0.00	1.00
200% FPL to 250% FPL	0.12	0.32	0.00	1.00
250% FPL to 400% FPL	0.21	0.40	0.00	1.00
400% FPL or greater / Unsubsidized Application	0.35	0.48	0.00	1.00
Net premium	297.44	279.22	-1,341.92	3,226.37

Panel B. Transaction Data Variables

	N=106,904		
	Mean	Median	SD
Monthly transaction dollar amount			
- Income	6,108.12	4,773.95	5,868.36
- Health out-of-pocket spending	22.45	6.92	57.12
- Premium payments	79.14	26.20	142.16
- Drug out-of-pocket spending	26.90	12.90	45.41
- Netflix and Hulu spending	2.69	0.00	5.66
- Utilities spending	64.41	27.92	281.21
- Auto insurance spending	56.30	12.12	105.88
- Other non premium/health/drug spending	6,260.69	4,696.10	7,178.96
Monthly number of transactions			
- Health out-of-pocket spending	0.31	0.13	0.53
- Other non premium/health/drug spending	39.66	33.72	29.01
Medicaid eligibility probability	0.22	0.12	0.25

Notes: This table reports summary statistics on the households who purchase individual insurance through the Covered California marketplace in Panel A. The statistics include characteristics for the sample of new enrollees in the plan years of 2014 and 2015. In Panel B we report data on the panel of households who purchase individual insurance, as observed in our bank transactions data for the years 2011 through 2016. For these statistics we only keep individuals with an annual income of less than $200k. Health out-of-pocket spending excludes spending on insurance premiums and spending at pharmacies. Other non-health spending excludes spending on health, pharmacy, and insurance premiums.

Table 2: 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)	0.293*	0.126	0.252*	0.320***	0.105
	(0.161)	(0.091)	(0.128)	(0.105)	(0.108)
Drop-out (% change)	-0.256	0.064	-0.093	-0.110	0.090
	(0.179)	(0.128)	(0.132)	(0.115)	(0.122)
Pre Sign-up Mean Health Spending	13.68	15.37	16.22	23.52	31.01
Number of Observations	6,456	9,354	7,945	9,883	8,660

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)	0.408***	0.197***	0.125*	0.041	-0.067
	(0.108)	(0.064)	(0.068)	(0.069)	(0.062)
Drop-out (% change)	-0.419***	-0.214***	-0.199***	-0.158**	-0.109*
	(0.099)	(0.073)	(0.066)	(0.062)	(0.061)
Pre Sign-up Mean Health Spending	9.79	12.08	18.35	31.03	39.69
Number of Observations	15,566	30,391	21,609	24,152	17,586

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)	0.560***	0.103	0.222**	0.188**	0.151*
	(0.172)	(0.086)	(0.104)	(0.094)	(0.082)
Drop-out (% change)	-0.350**	0.068	-0.203**	-0.059	-0.071
	(0.147)	(0.101)	(0.101)	(0.090)	(0.088)
Pre Sign-up Mean Health Transactions	.18	.25	.29	.35	.37
Number of Observations	6,456	9,354	7,945	9,883	8,660

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)	0.545***	0.245***	0.125**	-0.007	0.011
	(0.105)	(0.063)	(0.057)	(0.048)	(0.054)
Drop-out (% change)	-0.412***	-0.247***	-0.264***	-0.138***	-0.118**
	(0.094)	(0.060)	(0.056)	(0.046)	(0.049)
Pre Sign-up Mean Health Transactions	.13	.2	.35	.48	.43
Number of Observations	15,566	30,391	21,609	24,152	17,586

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 3: **Model estimates by Demographic Group**

Demog group	1{Below Age 40 }	1{Single}	1{Low Income}	ν	β_1	β_2	β_3
1	0	0	0	8.609	1.412	1.351	0.009
2	0	0	1	49.819	-0.034	0.381	0.236
3	0	1	0	4.768	6.663	0.568	0.009
4	0	1	1	49.852	0.585	0.714	0.065
5	1	0	0	0.133	0.548	0.449	0.065
6	1	0	1	49.879	0.304	0.714	0.052
7	1	1	0	0.709	2.209	0.357	0.059
8	1	1	1	49.753	0.920	0.824	0.061

Notes: This table reports the estimates from the dynamic choice model in which households can endogenously choose whether to exit and how much health care to consume in each period. We divide the sample of households into demographic groups and recover separate parameter estimates for each group. The eight demographics differ depending upon whether (a) the head of the household is below age 40 or at or above age 40, (b) whether the household insures only one member ("single" membered) or more than 1, and (c) whether the household is below 250% of the federal poverty line.

Table 4: Counterfactual Predictions under Drop-out Penalties

Panel 1: Market shares

	Drop-out threshold				
Plan type	0	.10	.20	.30	.40
No insurance	60.3	54.2	51.5	50.9	50.8
Bronze	6.8	8.0	9.3	9.5	9.5
Silver	24.6	27.4	28.8	29.2	29.3
Gold	3.1	3.8	3.9	3.9	3.9
Platinum	5.3	6.6	6.6	6.5	6.5

Panel 2: Average premiums, in $/month

	Drop-out threshold				
Plan type	0	.10	.20	.30	.40
Bronze	264.06	243.58	229.73	227.15	227.04
Silver	333.12	309.31	298.63	295.60	295.13
Gold	371.22	354.98	351.38	352.47	352.54
Platinum	426.26	404.76	398.75	397.68	397.52

Panel 3: Change in average surplus, in $/month

	Drop-out threshold				
Population	0	.10	.20	.30	.40
Overall	-138.6	-47.2	-8.5	-1.0	0.0
Low ν	-232.3	-72.9	-11.1	-1.1	0.0
High ν	-75.2	-29.8	-6.8	-0.9	0.0
$H \leq$ 20th percentile	-110.2	-51.1	-7.0	-0.3	0.0
H 20-40th percentile	-127.5	-56.7	-9.2	-1.3	0.0
H 40-60th percentile	-141.1	-47.6	-9.5	-1.2	0.0
H 60-80th percentile	-111.2	-31.5	-6.3	-0.8	0.0
$H >$ 80th percentile	-187.7	-53.1	-10.2	-1.2	0.0

Notes: This table reports the market shares, average monthly premiums, and average consumer surplus in dollar terms, as predicted from the dynamic choice model in which households face penalties for drop-out. We design penalties that set a cap or threshold on drop-out probabilities above which households face penalties. For illustration, we report a range of thresholds from 0 (ban drop-out) to 0.4, where 0.4 is rarely binding in our sample. We report market shares and premiums as weighted averages across regions and plans in California, conditional on plan metal tier. We report premiums scaled for a 40 year-old single household. For consumer surplus, we report differences in surplus relative to a threshold of 0.4. In the overall category, we average over all households in all plans and regions. We also report consumer surplus by dividing the sample of households by their ability to re-time spending (ν) and their expected total health spending, H. Low ν households represent 40% of the sample.

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

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

October 4, 2023

Table of Contents

Diamond: Stanford Graduate School of Business and NBER; Dickstein: Stern School of Business, New York University and NBER; McQuade: Stanford Graduate School of Business; Persson: Department of Economics, Stanford University and NBER.

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. 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 model 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 marketplace[1]); 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 medicaid eligibility, as discussed in Appendix Section E. Finally, with these household-monthly panels, we

[1]These 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.

[2]We examine these particular spending categories in Figure 4. 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.

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 A2, along with an accounting for the observations excluded sequentially under each criterion. These sequential drops result in 799,851 households in our final sample, accounting for 70.1% of all California non-mover households. Because the reference group is large and because we only use

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

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.

[4] We 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 $w(t) = w_1 w_2 ... w_{n(t)}$, where w_i denotes the ith word in $w(t)$ for $i = 1, 2, ..., n_t$ and n_t denotes the length or the number of words in $w(t)$. By assuming conditional independence of probabilities, we have

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

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

[5]Kaiser 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. 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 1.25 months in 2013. Overall, because the average household has some form of health insurance coverage for medical expenses for more than half of the year, the 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 A3 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 F.

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 A3 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 A5 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 A6. 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

[6] In 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.

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

[8] In 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 A5 to impute total health charges and changes in charges over time.

A6 shows that in MEPS, transactions increase 21% following marketplace enrollment in 2014, but with large standard errors. In our transactions data, the change is 24%, though now significant given our larger sample. The health out-of-pocket spending in the transactions data also increase by 21%, while MEPS shows an insignificant positive change, with wide standard errors.

Panel B of Appendix Table A6 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 A5, 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. Over this time period, 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

[9] The raw CCM data 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.

[10] For 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.

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

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

[12] Specifically, 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.

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 is 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 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

[13] The 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).

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 Additional drug consumption analyses

In addition to the analyses of consumption using health transactions and health spending in Section 3, we conduct similar analyses using drug transactions and drug spending. We separate out the drug analysis because, as discussed in Appendix Section C, our drug measure contains all pharmacy spending, including non-prescription purchases.

Our results appear consistent with the findings on general health spending. In Appendix Table A4, 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 2 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.

G Conceptual model details

To illustrate this finding, we describe the welfare outcome in two cost scenarios in Figure A9. 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

[14] In 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.

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%.

H Structural model details

We provide additional detail of the derivation of the dynamic empirical model we outline in Section 4.

H.1 First Order Conditions

To recover our parameters of interest, we first need to find the optimal share of annual spending that a household conducts in periods 1 and 2 as a function of plan characteristics and the household's cost of re-timing expenditures. We can do so by finding the first order conditions of the value functions in period 1 with respect to f_1 and in period 2 with respect to f_2.

We begin with period 2, using the value function for continuing coverage under plan j, denoted V_{2j} in Equation 7. We take the first order condition with respect to f_2, yielding:

$$\psi_{3j}\frac{\partial v_{30}}{\partial f_2} + (1-\psi_{3j})\frac{\partial v_{3j}}{\partial f_2} = 0$$

where:

$$\begin{aligned}
\frac{\partial v_{30}}{\partial f_2} &= -\nu H\beta_3\left[2\left(f_2 - 1/3\right) - 2\left(1 - f_1 - f_2 - 1/3\right)\right] + \beta_3 H\left(1-\alpha_j\right)\\
\frac{\partial v_{3j}}{\partial f_2} &= -\nu H\beta_3\left[2\left(f_2 - 1/3\right) - 2\left(1 - f_1 - f_2 - 1/3\right)\right]
\end{aligned}$$

Simplifying, we find:

$$-\nu\left[2\left(f_2 - 1/3\right) - 2\left(1 - f_1 - f_2 - 1/3\right)\right] + \psi_3\left(1-\alpha_j\right) = 0 \tag{14}$$

For period 1, we equivalently take the first order condition of the value function for period 1, in Equation 9:

$$\psi_{2j}\frac{\partial v_{20}}{\partial f_1} + (1-\psi_{2j})\frac{\partial v_{2j}}{\partial f_1} = 0$$

where:

$$\frac{\partial v_{20}}{\partial f_1} = -\nu H \beta_3 \left[2\left(f_1 - 1/3\right) - 2\left(\frac{1-f_1}{2} - 1/3\right)\right] + \beta_3 H \left(1 - \alpha_j\right) = 0$$

and:

$$\begin{aligned}\frac{\partial v_{2j}}{\partial f_1} &= \frac{\partial}{\partial f_1} \log\left[\exp\left(v_{30}\left(f_1\right)\right) + \exp\left(v_{3j}\left(f_1\right)\right)\right] \\ &= \psi_3 \frac{\partial v_{30}}{\partial f_1} + \left(1 - \psi_3\right) \frac{\partial v_{3j}}{\partial f_1}\end{aligned}$$

with:

$$\begin{aligned}\frac{\partial v_{30}}{\partial f_1} &= -\nu H \beta_3 \left[2\left(f_1 - 1/3\right) - 2\left(1 - f_1 - f_2 - 1/3\right)\right] + \beta_3 H \left(1 - \alpha_j\right) \\ \frac{\partial v_{3j}}{\partial f_1} &= -\nu H \beta_3 \left[2\left(f_1 - 1/3\right) - 2\left(1 - f_1 - f_2 - 1/3\right)\right]\end{aligned}$$

Simplifying, we find:

$$\begin{aligned}&\psi_{2j} \left\{-\nu \left[2\left(f_1 - 1/3\right) - 2\left(\frac{1-f_1}{2} - 1/3\right)\right] + 1 - \alpha_j\right\} \\ &+ \left(1 - \psi_{2j}\right)\left\{-\nu\left[2\left(f_1 - 1/3\right) - 2\left(1 - f_1 - f_2 - 1/3\right)\right] + \psi_{3j}\left(1 - \alpha_j\right)\right\} = 0\end{aligned} \tag{15}$$

Thus, from the first and second periods, we have two first order conditions that are a function only of the parameter ν and observed drop-out shares, plan actuarial values, and the shares of health spending in periods 1 and 2, (f_1, f_2).

H.2 Steps to identify H

In this section, we provide more details to show how to use our model to estimate H by generating a prediction of the probability of choosing each option j in period 1.

To see this, we can use our model to write the value of choosing option j in period 1:

$$v_{1ij} = \sigma_2 \log\left(\exp\left(\frac{v_{2ij}}{\sigma_2}\right) + \exp\left(\frac{v_{2i0}}{\sigma_2}\right)\right).$$

Using our result from period 2, we also know that the probability of remaining enrolled in period 2 takes the form:

$$1 - \psi_{2ij} = \frac{\exp\left(\frac{v_{2ij}}{\sigma_2}\right)}{\exp\left(\frac{v_{2i0}}{\sigma_2}\right) + \exp\left(\frac{v_{2ij}}{\sigma_2}\right)}$$

Taking the log of both sides, we find:

$$\log\left(1 - \psi_{2ij}\right) = v_{2ij}/\sigma_2 - \log\left(\exp\left(\frac{v_{2ij}}{\sigma_2}\right) + \exp\left(\frac{v_{2i0}}{\sigma_2}\right)\right)$$

Multiplying both sides by σ_2 and recognizing the last term on the right as v_{1ij}, we find:

$$v_{1ij} = v_{2ij} - \sigma_2 \log\left(1 - \psi_{2ij}\right).$$

We can use the same approach to translate the log-sum expression for v_{2ij}, the period 2 value of choice j, into a form that looks similar but involves period 3 variables:

$$\begin{aligned}v_{2ij} &= \log\left(\exp\left(v_{3ij}\right) + \exp\left(v_{3i0}\right)\right) \\ &= v_{3ij} - \log\left(1 - \psi_{3ij}\right).\end{aligned}$$

Substituting this expression for v_{2ij} into v_{1ij} and letting $\beta_2 = 1/\sigma_2$, we find:

$$\begin{aligned}v_{1ij} &= v_{3ij} - 1/\beta_2 \log\left(1 - \psi_{2ij}\right) - \log\left(1 - \psi_{3ij}\right) \\ &= \delta_{1j} + \delta_{2j} + \delta_{3j} - \beta_3 H \left[\nu \sum_{t=1}^{3} \left(f_t - 1/3\right)^2 + \alpha_{ij}\right] - 3\beta_3 P_{ij} - \sigma_2 \log\left(1 - \psi_{2ij}\right) - \log\left(1 - \psi_{3ij}\right).\end{aligned}$$

When a household chooses to be uninsured, its value equals:

$$v_{10} = -\beta_3 H + \left(1/\beta_2 + 1\right) E\left[\varepsilon\right]$$

After re-writing, we can estimate the following conditional logit:

$$\psi_{1ij} = \frac{\exp\left(\beta_1 X_{ij} + \tilde{\delta}_{1j}\right)}{\exp\left(\beta_1 X_{i0}\right) + \sum_{j'} \exp\left(\beta_1 X_{ij} + \tilde{\delta}_{1j}\right)}$$

where:

$$\begin{aligned}X_{ij} &= -\beta_3 H \left[\nu \sum_{t=1}^{3} \left(f_t - 1/3\right)^2 + \alpha_{ij}\right] - 3\beta_3 P_{ij} - 1/\beta_2 \log\left(1 - \psi_{2ij}\right) - \log\left(1 - \psi_{3ij}\right) \\ X_{i0} &= -\beta_3 H + \left(1/\beta_2 + 1\right) E\left[\varepsilon\right] \\ \tilde{\delta}_{1j} &= \beta_1(\delta_{1j} + \delta_{2j} + \delta_{3j})\end{aligned}$$

We recover H using this logit specification, given estimates of (β_3, β_2, ν) and data on premiums, actuarial value, and the drop-out shares in periods 2 and 3. In particular, we interpret the coefficient on the drop-out penalty element as $\beta_1\beta_3 H$, and use parameters from earlier steps to back out the value of H alone.

I Counterfactual algorithm details

In this appendix section, we provide more detail on the algorithm we use to compute our counterfactual equilibrium. The approach follows the equilibrium concept from Azevedo and Gottlieb (2017) for a competitive insurance market, as applied previously in the individual insurance market (Dickstein et al., Forthcoming). We adjust the algorithm for the specific regulatory environment in California.

We begin by collecting the parameters from our model estimation. From our three stage routine, which relies on household plan choices, drop-out shares by plan and household demographic, and the observed pattern of household health care spending over distinct periods of the year, we collect the following parameters in

θ:

$$\theta = \begin{pmatrix} \nu \\ \beta_1 \\ \beta_2 \\ \beta_3 \\ \delta_{1j} \forall j \in J \\ \delta_{2j} \forall j \in J \\ \delta_{3j} \forall j \in J \\ H \end{pmatrix}$$

In the equilibrium search, our goal is to find p_{jr}, the plan-region baseline premium the insurer sets in 2015 to equate revenue and costs in equilibrium for all plans and regions, accounting for possible government transfers. We can translate the normalized baseline premium to (a) gross premiums for household i, p_{ijr}, and (b) net (subsidized) premiums for household i, p^s_{ijr}.

First, we use regulated age-rating factors in effect in California in our sample period to convert baseline premiums to household-specific premiums. We label these factors $\gamma_{k,i}$ for household member k in household i:

$$p_{ijr} = p_{jr} \sum_{k \in i}^{K_i} \gamma_{k,i}$$

Second, We simplify the premium subsidy formula for the purpose of searching for counterfactual premiums. In the true formula, premium subsidies for individual market households depend both on household income relative to regulated thresholds and the premium for the second cheapest silver plan offered in a market and year. To simplify our equilibrium search while still capturing the key features of premium subsidies, we fix subsidy levels before the equilibrium search algorithm. We also account for the minimum premium charged in California, which equals $1 per household member:[15]

$$p^s_{ijr} = max\{p_{ijr} - subs_{ir}, K_i\}$$

I.1 Baseline mark-up or government transfer share

In the counterfactual algorithm, we keep fixed the relationship between premium revenues and the total covered health costs that enrollees incur in the plan. We start by using the model estimates and data to compute the total premium revenues each carrier receives for plans it offers in the observed data. The premium payment for a household depends on the underlying premium the carrier sets for the plan, multiplied by the household's rating factor. The rating factor varies with the age of each household member, including up to three dependents under age 18:

$$p_{ijr} = p_{jr} \sum_{k \in i}^{K_i} \gamma_{k,i}$$

where K_i includes all household members over 18 and only the three oldest household members under 18. To compute the premium revenues to the carrier, we omit any subsidies, since the carrier collects the full gross premium including the household's net payment and the government's subsidy payment.

[15] Outside of our sample period, on January 1, 2022, the minimum premium fell to $0. California legislation AB 133 created a "California Premium Credit," allowing the state to subsidize the remaining $1 per member per month premium for the most subsidized households.

We also compute the total implied costs the carrier incurs for the household spending amongst enrollees. This includes only the actuarial value or fraction of the plan's health spending, α_j, that the insurer pays; the remaining costs are the household's responsibility. If a low-income household is eligible for cost-sharing subsidies when it purchases a silver plan, we assume the carrier is nonetheless responsible only for 70% of the costs, with the government paying the carrier back for additional subsidized out-of-pocket costs. We label α_j^{base} the baseline actuarial value that does not include cost-sharing subsidies for low-income households.

We run the following OLS regression using the observed premium revenues for a plan on the left-hand side and the observed costs based on implied H_i for household i on the right-hand side:

$$\sum_{i \in N_{jr}}^{N_{jr}} p_{ijr} g_i = \varphi \sum_{i \in N_{jr}}^{N_{jr}} \alpha_j^{base} H_i \frac{g_i}{12}$$

where p_{ijr} is the gross premium for household i choosing plan j in region r for each month. We scale both the premium and the health costs by the number of months of coverage observed in the data, g_i for each household i.

We recover the parameter φ. Here, $\varphi > 1$ can indicate plan markups. If $\varphi \leq 1$, we assume the plan loses money on premium revenue alone, but breaks even due to additional monetary transfers from the government, including transfers in the form of risk-adjustment, risk-corridor, and reinsurance payments. We hold φ fixed in our counterfactual simulation with drop-out penalties, under the assumption that the transfers or markups scale with health spending costs.

I.2 Steps in the equilibrium search

With these components, our algorithm proceeds as follows:

1. Compute the lth guess at the household's (potentially) subsidized premium, $p_{ijr}^{s,l}$. Using the baseline premium for plan j, p_{jr}^l, we know $p_{ijr}^l = p_{jr}^l \sum_{k \in i}^{K_i} \gamma_{k,i}$ and $p_{ijr}^{s,l} = max\{p_{ijr}^l - subs_{ir}, K_i\}$. We compute this premium for every plan j and every household i, since in the counterfactual households will choose all plans with some probability, ψ_{i1j}.

2. Compute each household's value functions for plan j for periods 1, 2, and 3 given the drop-out penalty (m_1, m_2) and $p_{ijr}^{s,l}$. We add m_t to the household's period-specific value functions, as described in Section 4.2.1. Specifically, we add m_2 to V_{30} and m_1 to V_{20}. Unlike a mandate penalty, if the household chooses never to enroll in insurance, it never faces a penalty. That is, there is no penalty added to V_{10}.

3. With the penalty in the model, we allow the households to choose optimally three elements: (a) the fraction of spending to consume in each time period, $(f_1, f_2, 1 - f_1 - f_2)$; (b) enrollment in each time period; and (c) its preferred plan, including the choice to remain uninsured throughout the plan year. For each household i, We save the plan choice probabilities, ψ_{i1j}, and the probabilities of exit in periods 2 and 3, ψ_{i2j} and ψ_{i3j}, respectively.

4. We then compute the total costs the insurer incurs in a plan, which is a function of the expected spending of the enrolled households for the duration of months the household is covered. For household

i, the expected spending under plan j is:

$$E[spend_{ij}] = \alpha_j^{base} H_i[\psi_{i2j} f_1 + (1-\psi_{i2j})\psi_{i3j}(f_1+f_2) + (1-\psi_{i2j})(1-\psi_{i3j})(1)]$$

We sum this over all households, multiplied by the probability the household chooses plan j:

$$\sum_i^N \psi_{i1j} E[spend_{ij}]$$

5. We then use our estimated φ to compute, for each plan, the predicted premium revenue the plan will collect in total. Predicted total premium revenue for iteration l equals:

$$R_j^l = \hat{\varphi} \sum_i^N \psi_{i1j} E[spend_{ij}]$$

We divide this revenue by the sum of household rating factors, where we pro-rate the rating factors of each household by the fraction of the year in which the household enrolls and pays premiums. For household i, the expected number of months of the year covered equals:

$$g_i = \psi_{i2j}(4) + (1-\psi_{i2j})\psi_{i3j}(8) + (1-\psi_{i2j})(1-\psi_{i3j})(12)$$

The household's pro-rated rating factor is then:

$$\frac{g_i}{12} \sum_{k \in i}^{K_i} \gamma_{k,i}$$

The following scaled measure is the new equilibrium "baseline" monthly premium for plan j for this iteration of algorithm:

$$p_{jr}^{l+1} = \frac{R_j^l}{\sum_i^N \psi_{i1j} \frac{g_i}{12} \sum_{k \in i}^{K_i} \gamma_{k,i}}$$

6. We test for convergence by comparing the plan-specific price vector p_{jr}^{l+1} from the $(l+1)$st iteration against the lth iteration across all plans. We compute the mean across plans $j = 1, ...J_r$ in all regions r, which total J possible plans:

$$\frac{1}{J} \sum_{\forall (j,r)}^{J} \left(\frac{p_{jr}^{l+1} - p_{jr}^{l}}{p_{jr}^{l}} \right)$$

If the mean percentage difference in prices (excluding the outside good) is less than .001, we stop the search algorithm and define the equilibrium price vector, p^{eql}, of length J. An entry in p^{eql} is equal to $p_{jr}^{eql} = p_{jr}^{l+1}$. If our condition on the mean percentage difference in premiums is not satisfied, we return to Step 1, using p_{jr}^{l+1} in place of p_{jr}^{l}.

I.3 Counterfactual sample

For each counterfactual scenario we examine, we define the set of households able to choose a plan in the individual insurance market in 2015. We include households in the data who purchase an individual market plan plus the set of uninsured households in that year.

J Appendix figures and tables

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

(a) Health Transactions

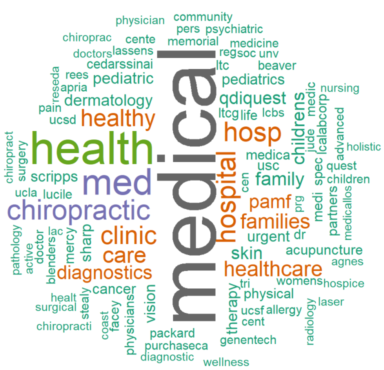

(b) Drugs Transactions

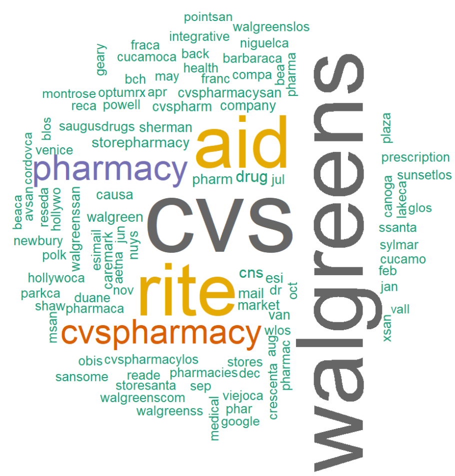

(d) Other Transactions

Notes: This figure reports the top 100 words classified by our machine learning algorithm as health spending, drug spending, 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

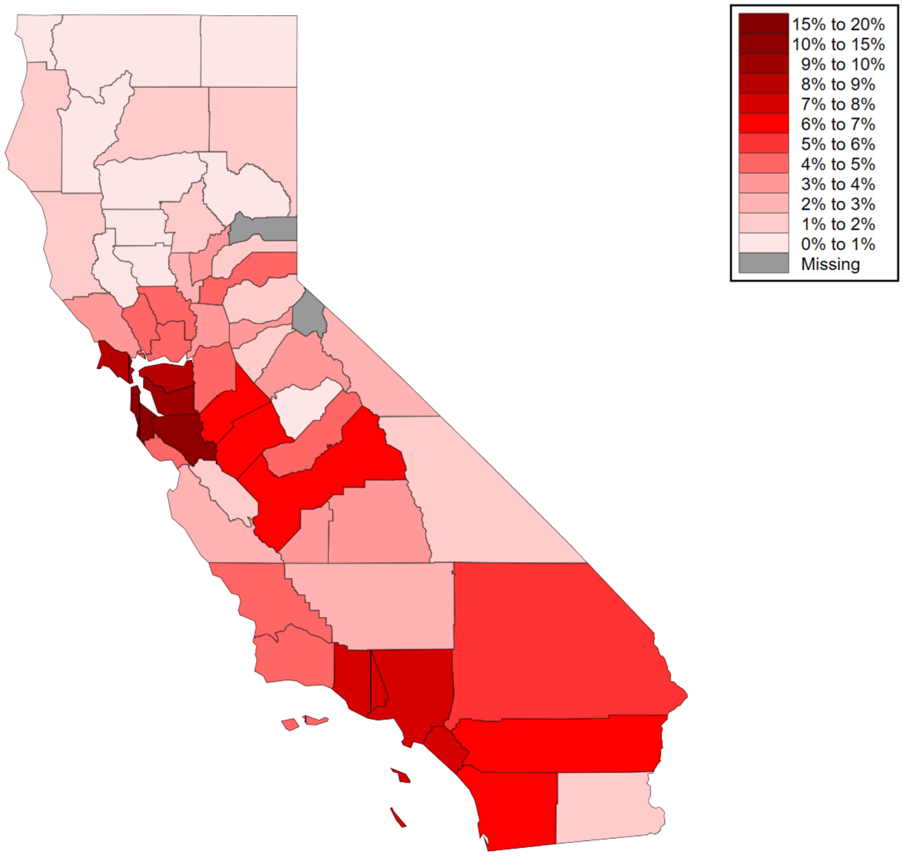

Notes: This map illustrates the geographic distribution of households in our transaction data relative to the number of households in the 2011-2015 NHGIS data. Specifically, for each county, we count the number of unique households in our data, divide by the corresponding number in NHGIS, and then multiply 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 799,851 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**

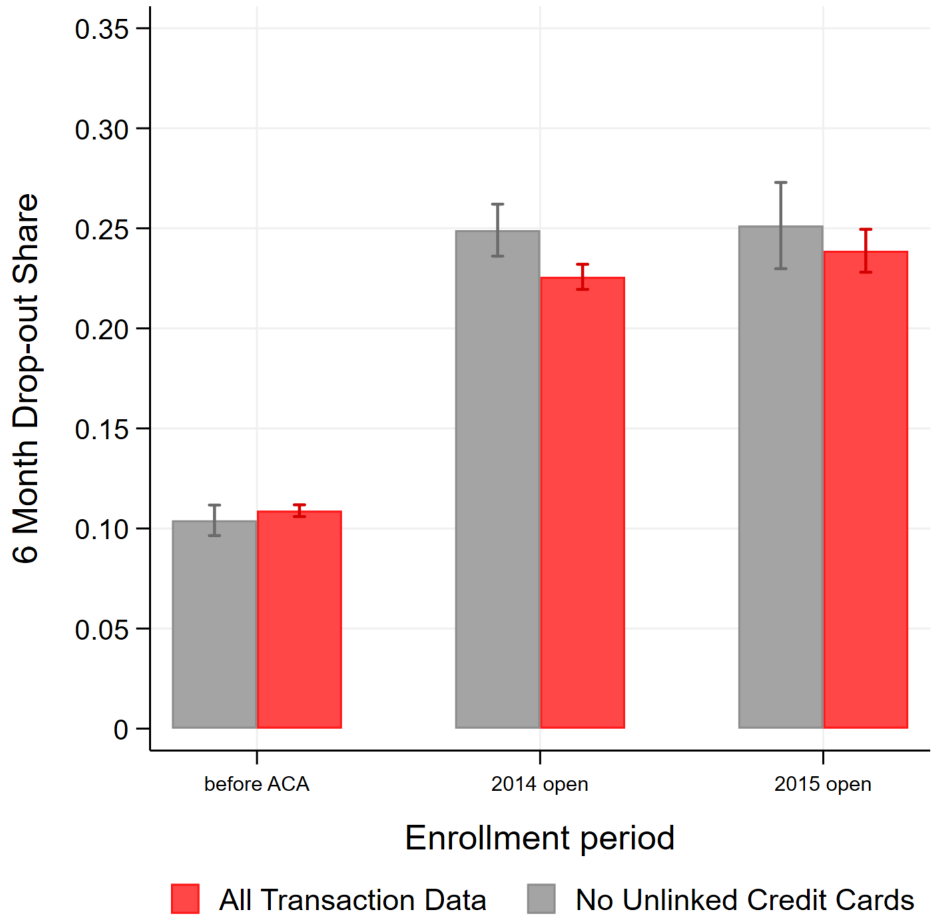

Notes: This chart compares our measure of drop-out among our entire transaction sample vs. among households who do not have an unlinked credit card. We define households as having an unlinked credit card as follows: Using the bank account data, we observe the total amount the household spends paying 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

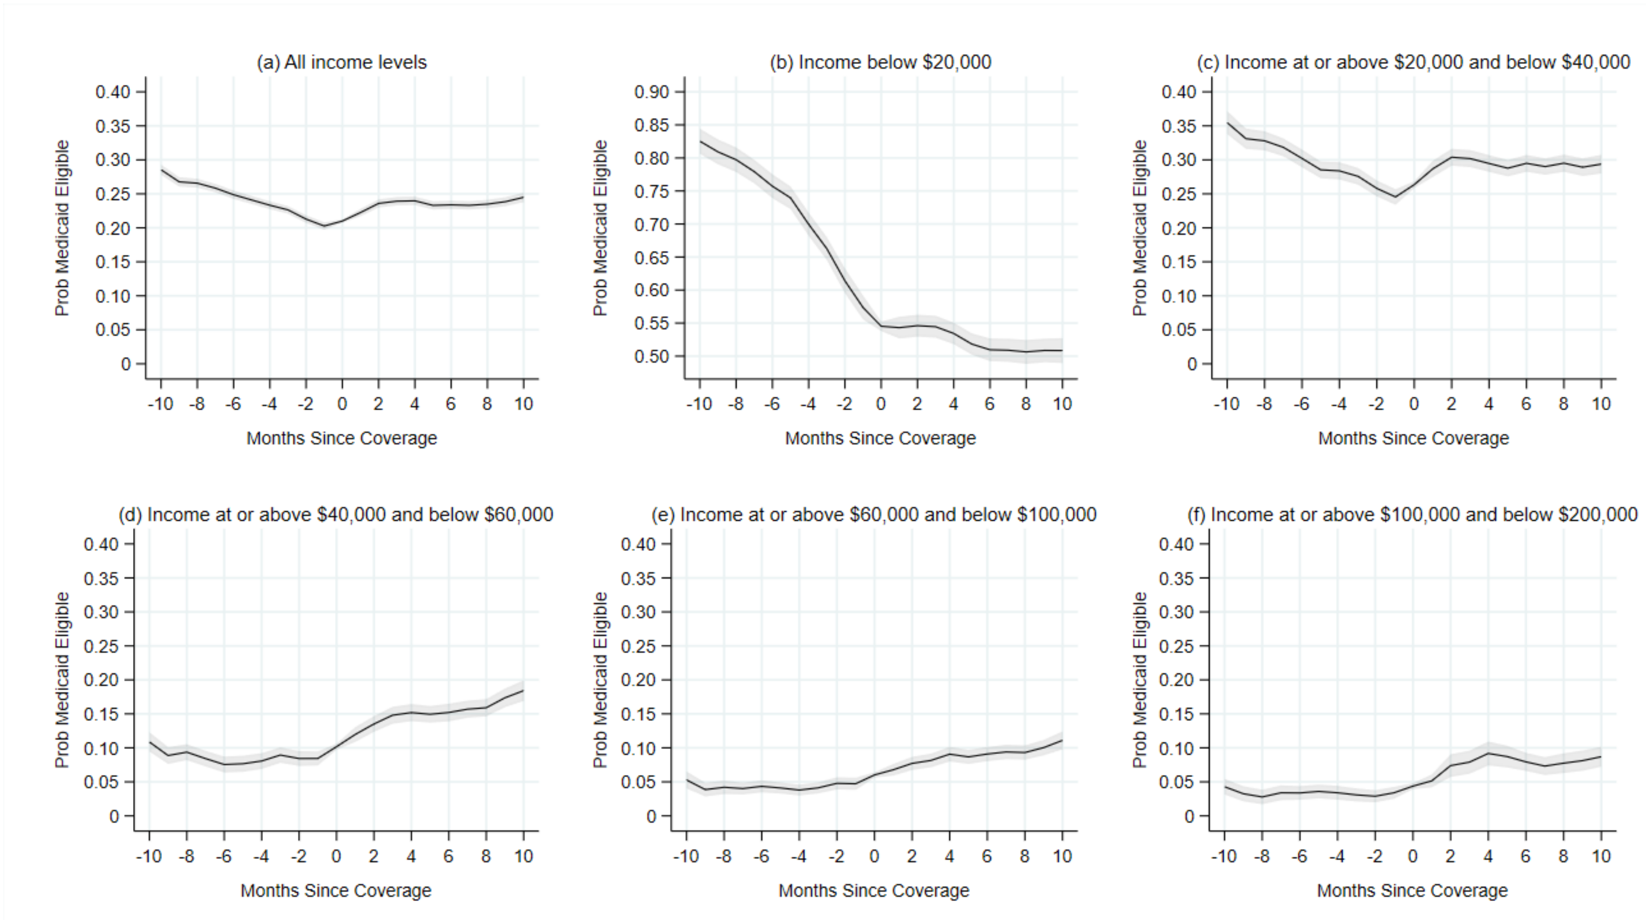

Notes: This figure shows the probability of medicaid eligibility among households by income group. We assign income groups using pre-ACA annual post-tax income from 2013. The x-axis contains the months surrounding the coverage period under individual health insurance; all months of coverage are collapsed 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: Detailed Income and Expenditure Dynamics around Sign-up and Drop-out by Dropout Status

Income

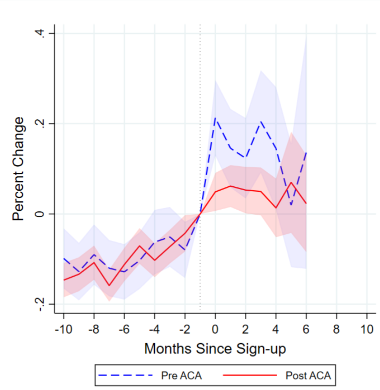

(a) Dropout: Pre Sign-up vs. Enrollment

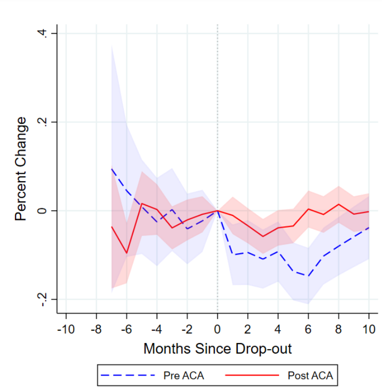

(b) Dropout: Enrollment vs. Post Drop-out

Other Spending

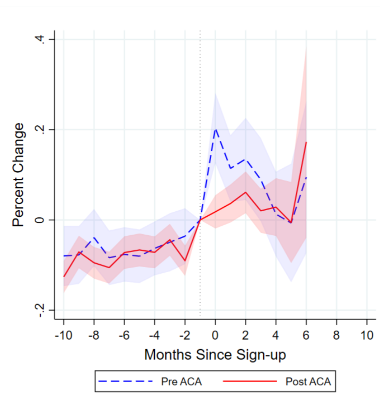

(c) Dropout: Pre Sign-up vs. Enrollment

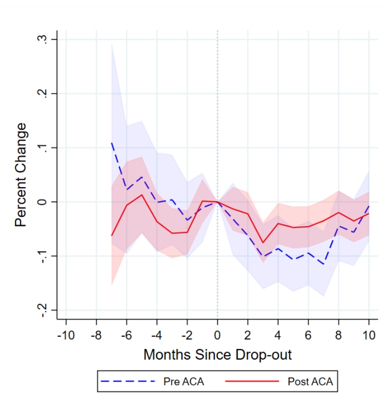

(d) Dropout: Enrollment vs. Post Drop-out

Netflix/Hulu

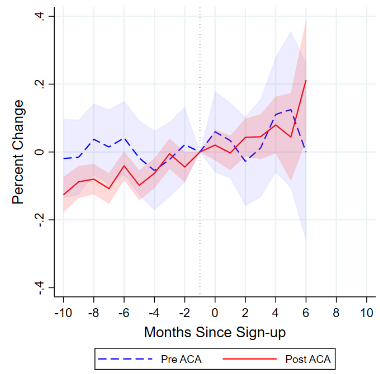

(a) Dropout: Pre Sign-up vs. Enrollment

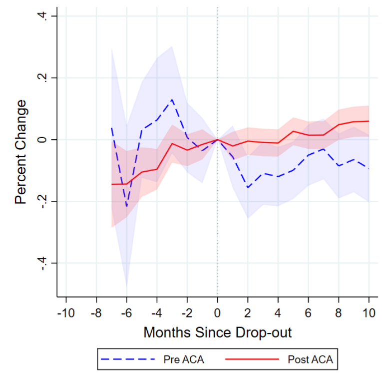

(b) Dropout: Enrollment vs. Post Drop-out

Utilities

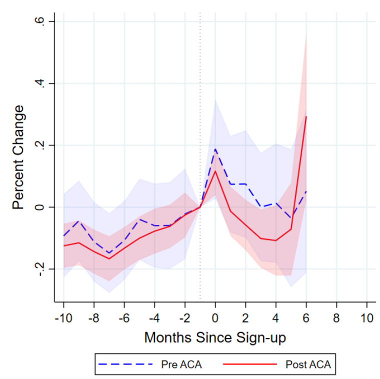

(c) Dropout: Pre Sign-up vs. Enrollment

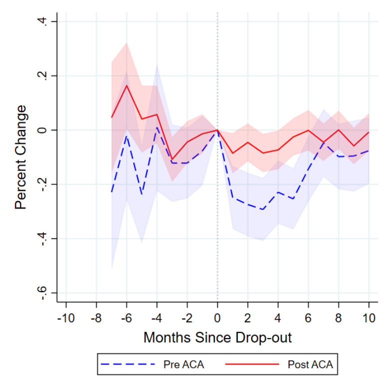

(d) Dropout: Enrollment vs. Post Drop-out

Auto Insurance

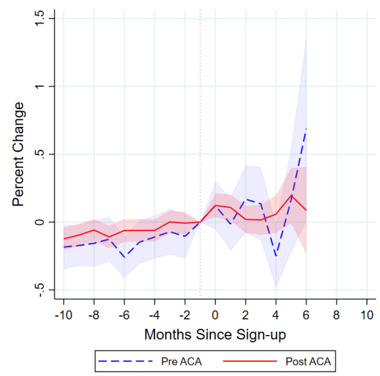

(a) Dropout: Pre Sign-up vs. Enrollment

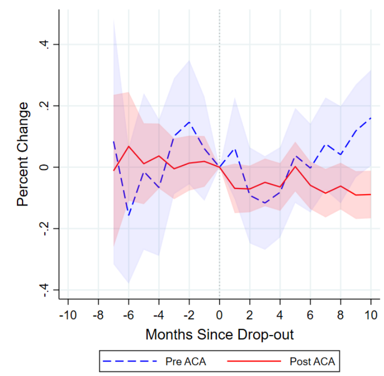

(b) Dropout: Enrollment vs. Post Drop-out

Health Spending

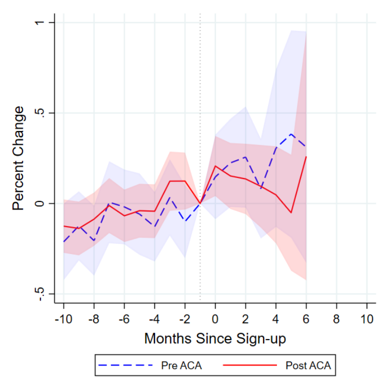

(c) Dropout: Pre Sign-up vs. Enrollment

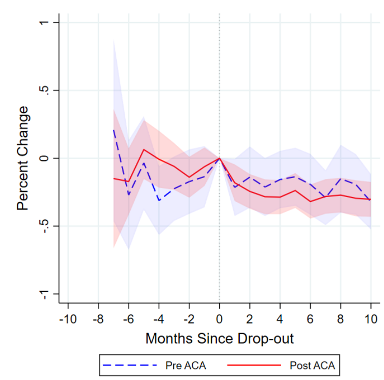

(d) Dropout: Enrollment vs. Post Drop-out

Notes: This figure shows the change in spending before and after enrollment in individually purchased private insurance by dropout type. A pre-ACA dropout is defined as a person insured for less than nine months; not dropping out in November or December; and dropping out before July 2013. A post-ACA dropout is defined as a person insured for less than nine months; not dropping out in November or December; signing up after November 2013; and dropping out after December 2013. Figure 4 shows these same graphs with spending during coverage collapsed to to time 0 and combining dropouts with non-dropouts.

Figure A8: Demand and Average Cost in a Market with Dropouts

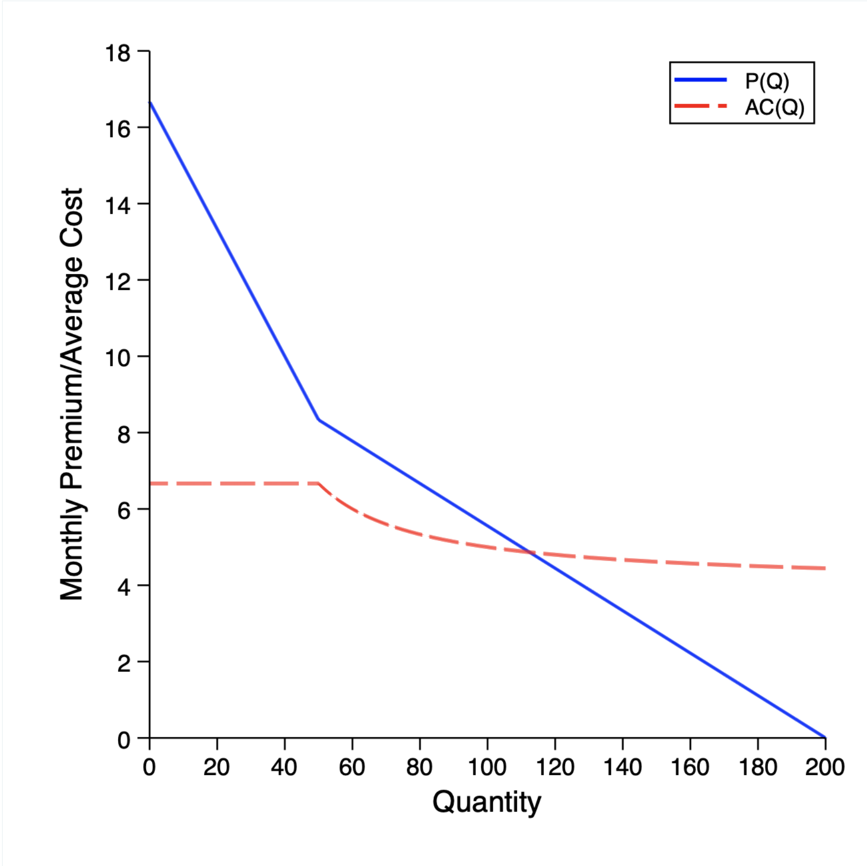

Notes: This figure illustrates the demand curve and member-month average cost curve in a market with drop-out types present but no traditional adverse selection (i.e. annual household health costs are equal across drop-out and full-year types). The equilibrium monthly premium occurs where the demand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

Figure A9: Consumer Surplus by Drop-out Type, with Differential Health Costs

Panel 1: Smaller Difference in Annual Health Costs by Type

(a) Welfare **(b) Premium** **(c) Consumer Surplus**

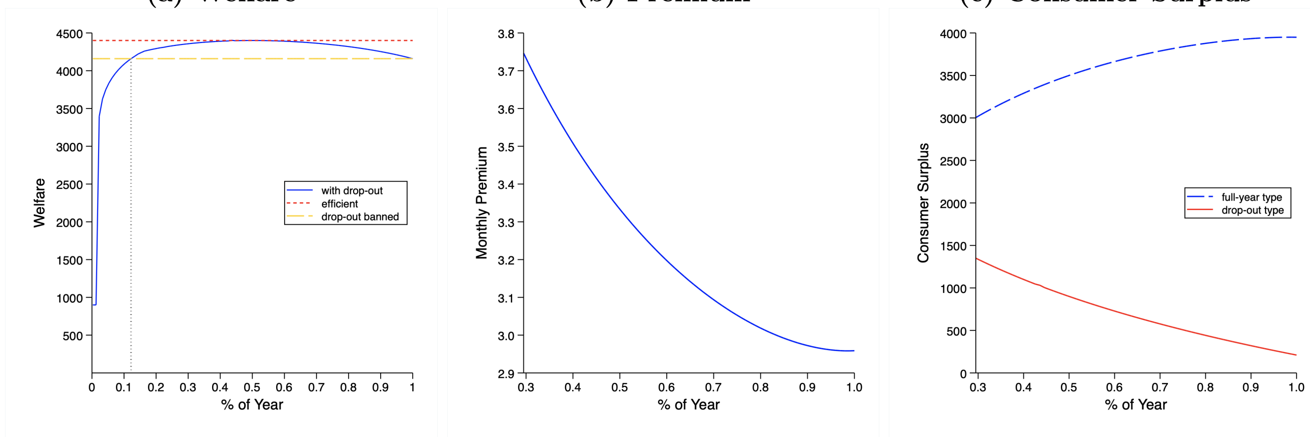

Panel 2: Larger Difference in Annual Health Costs by Type

(a) Welfare **(b) Premium** **(c) Consumer Surplus**

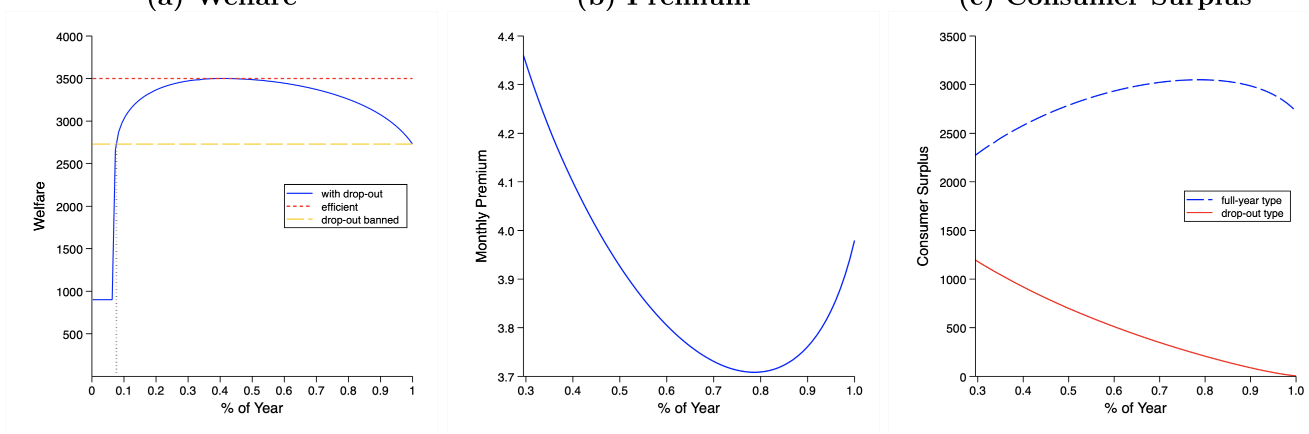

Notes: This figure illustrates the welfare consequences of drop-out in the presence of traditional adverse selection. In panel (1) we assume a small difference in annual health care costs by type: $C_D = 20$ while $C_F = 40$. In panel (2), we assume a larger difference in annual costs by type: $C_D = 20$ while $C_F = 49$. Subfigures (a), (b), and (c) show the effect of drop-out on social welfare, monthly premiums, and consumer surplus, respectively. We break out consumer surplus separately by household type. In this numerical example we set $N_D = N_F = 100$, $G_D(\cdot) = Unif([0, 50])$, and $G_F(\cdot) = Unif([50, 100])$.

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: 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	1,047,873	91.8%
Drop households who enter after Dec 2013 or exit before Jan 2013	1,017,549	89.1%
Drop households with inactive accounts:	846,301	74.1%
- 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	812,654	71.2%
Drop households that are likely to be small businesses	799,853	70.1%
- Average annual credit transaction count $\geq$ 500 (around 99th percentile)		
- Average annual debit transaction count $\geq$ 2,000 (around 99th percentile)		
Drop counties with less than 5 households	799,851	70.1%
California non-mover households (cleaned sample)	799,851	70.1%
Keep individually enrolled from 2011 to 2015	**106,904**	**8.5%**

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 846,301 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. 799,851 households remain after these steps. To form our final analysis sample, we save all individual market enrollees. The final sample contains 106,904 households.

Table A3: **Comparison of Annual Spending with MEPS Data**

Panel A. Summary Statistics of Annual Drug/Health Spending

	MEPS			Transactions		
Year	N	Drug	Health	N	Drug	Health
2012	9,327	141.7	204.8	196,650	258.0	212.7
2013	9,263	122.4	258.0	247,664	324.5	222.3
2014	8,761	119.8	262.5	235,029	323.2	239.2
2015	8,588	114.2	276.1	211,373	303.8	251.3

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	126.5	480.0	928.4	3,094.7
	2013	0.0	0.0	0.0	0.0	0.0	131.4	500.7	983.6	3,247.2
	2014	0.0	0.0	0.0	0.0	3.7	150.0	545.8	1052.7	3,298.8
	2015	0.0	0.0	0.0	0.0	5.0	152.6	560.0	1082.5	3,483.3

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 A4: 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)	0.699***	0.389***	0.335***	0.274***	0.292***
	(0.112)	(0.099)	(0.095)	(0.077)	(0.081)
Drop-out (% change)	-0.196**	0.118	0.001	-0.038	0.046
	(0.097)	(0.087)	(0.111)	(0.077)	(0.092)
Pre Sign-up Mean Drug Spending	11.96	14.79	21.88	26.05	27.76
Number of Observations	6,456	9,354	7,945	9,883	8,660

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)	0.171***	-0.004	0.020	-0.113***	-0.092*
	(0.064)	(0.030)	(0.037)	(0.035)	(0.050)
Drop-out (% change)	-0.133**	-0.104***	-0.150***	-0.083***	-0.115***
	(0.057)	(0.032)	(0.036)	(0.027)	(0.041)
Pre Sign-up Mean Drug Spending	16.39	20.32	28.18	40.68	48.3
Number of Observations	15,566	30,391	21,609	24,152	17,586

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)	0.844***	0.512***	0.396***	0.335***	0.328***
	(0.127)	(0.080)	(0.075)	(0.075)	(0.079)
Drop-out (% change)	-0.185*	0.088	0.008	-0.049	0.086
	(0.108)	(0.080)	(0.079)	(0.070)	(0.084)
Pre Sign-up Mean Drug Transactions	.38	.54	.75	.72	.82
Number of Observations	6,456	9,354	7,945	9,883	8,660

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)	0.199***	0.012	0.064**	-0.031	-0.002
	(0.055)	(0.026)	(0.031)	(0.025)	(0.026)
Drop-out (% change)	-0.187***	-0.122***	-0.150***	-0.089***	-0.136***
	(0.051)	(0.025)	(0.033)	(0.024)	(0.030)
Pre Sign-up Mean Drug Transactions	.7	.92	1.07	1.32	1.49
Number of Observations	15,566	30,391	21,609	24,152	17,586

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 A5: Imputation of Healthcare Charges from Health Out-of-pocket in MEPS

	2013 Health Charges	2014 Health Charges
Health OOP	5.640**	4.580*
	(2.668)	(2.391)
Constant	2658.086**	4159.501**
	(1090.061)	(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 A6: **Healthcare Consumption for 2014 Open Enrollees: Transactions vs. MEPS Data**

	Panel A: Transactions		
	(1)	(2)	(3)
	Health Charges	Health OOP	Health Transactions
Post ACA	1474.134***	47.868***	0.763***
	(37.773)	(7.393)	(0.063)
Pre Period	3970.899***	232.783***	3.156***
	(35.253)	(6.249)	(0.065)
% Change	0.371***	0.206***	0.242***
	(0.012)	(0.035)	(0.022)
Observations	30,927	30,927	30,927

	Panel B: MEPS		
	(1)	(2)	(3)
	Health Charges	Health OOP	Health Transactions
Post ACA	1212.748	3.977	1.187
	(1760.709)	(58.754)	(1.011)
Pre Period	4292.207***	289.758***	5.614***
	(972.820)	(39.366)	(0.789)
% Change	0.283	0.014	0.211
	(0.450)	(0.204)	(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 A5). 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.