ESSAYS ON HEALTH INSURANCE MARKETS

Krisztina Horvath

A dissertation
submitted to the Faculty of
the department of Economics
in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy

Boston College

Morrissey College of Arts and Sciences

Graduate School

March, 2020

ESSAYS ON HEALTH INSURANCE MARKETS

Krisztina Horvath

Advisors:

Michael D. Grubb, PhD, Julie Holland Mortimer, PhD,

Christopher F. Baum, PhD, Samuel Richardson, PhD

Abstract

The first chapter studies behavioral mechanisms to expand health insurance coverage. In health insurance markets where regulators limit insurers' ability to price on the health status of individuals, a traditional regulatory intervention to protect the market from adverse selection and expand coverage among young and healthy people is mandating insurance coverage. In this chapter, I analyze an alternative, behavioral mechanism in the context of the Affordable Care Act Marketplaces: the automatic enrollment of the uninsured with possible opt-out. I build a theoretical model which shows that this nudging policy increases coverage rates, and the size of its benefit depends on the strength of consumer inertia. Using an individual-level panel dataset on health insurance plan choice and claims, I estimate a structural model of health insurance demand and supply in the presence of switching costs. Simulating the effects of the policy, I find that auto-enrollment can increase enrollment rates by over 60% and reduce annual premiums by $300. Moreover, I show that taking into account the heterogeneity of preferences is essential when designing default plans for auto-enrolled consumers. Defaulting everyone into the same contract type leads to more quitting due to inefficient matching and it may also indirectly increase adverse selection on the intensive margin through the price adjustment mechanism. The results of this paper suggest that in order to avoid these problems and maximize the benfits of auto-enrollment in selection markets, it is important to design smart default policies.

The second chapter explores how changes in cost sharing affect consumers' demand for health care. Cost sharing reduction (CSR) subsidies are a less well-known provision of the Affordable Care Act (ACA) that aimed to make private health insurance coverage

more affordable. These subsidies discontinuously increase the share of expenses paid by the insurer as enrollee income crosses the eligibility cutoffs. This specific subsidy design provides a unique setting to identify moral hazard in health care utilization from observational data that is a major empirical challenge in the literature. In this chapter, I combine individual-level post-subsidy premium data from an All Payer Claims Database with information on plan-level base prices to recover the amount of the premium subsidy. Applying the ACA's premium subsidy formula backwards, I am able to estimate family income. Using this imputed income, I exploit a sharp regression discontinuity design to study the impact of changes in actuarial value on consumer behavior. I find significant increases in health care utilization at income levels associated with the CSR subsidy eligibility cutoffs. These results imply that individuals tend to use more health care services only due to the fact that the insurer becomes responsible for a larger share of their expenditures. These results provide insights about the price elasticity of demand for medical care in a new context.

The third chapter evaluates the impact of the ACA on HPV vaccination. Rates of completion of the HPV vaccine series remain suboptimal in the US. The effects of the ACA on HPV vaccine completion are largely unknown. The aim of this study was to examine the associations between the ACA's 2010 provisions and 2014 insurance expansions with HPV vaccine completion by sex and health insurance type. Using 2009-2015 public and private health insurance claims, we conducted a logistic regression model to examine the associations between the ACA policy changes with HPV vaccine completion as well as interactions by sex and health insurance type. Among females and males who initiated the HPV vaccine, 27.6% and 28.0%, respectively, completed the series within 12 months. Among females, the 2010 ACA provision was associated with increases in HPV vaccine completion for the privately-insured and Medicaid enrollees. The 2014 health insurance expansions were associated with increases in vaccine completion for females with private insurance and Medicaid. Among males, the 2014 ACA reforms were associated with increases in HPV vaccine completion for the privately-insured and Medicaid enrollees. Despite low HPV vaccine completion overall, both sets of ACA provisions increased completion among females and males. Our results suggest that expanding Medicaid across the remaining states could increase HPV vaccine completion among publicly-insured youth and prevent HPV-related cancers.

Contents

List of Figures

List of Tables

Chapter 1

Adverse Selection and Switching Costs in Health Insurance Marketplaces:

Using Nudges to Fight the Death Spiral

> *"By improving the effectiveness and efficiency of Government, behavioral science insights can support a range of national priorities ,...; enabling Americans to lead longer, healthier lives"*
>
> (Barack Obama, 2015)

1.1 Introduction

Despite having the highest per capita health care spending in the world, high uninsured rates are a historical problem in the United States (OECD, 2015). To protect consumers from the growing financial burden of medical bills, one of the main goals of the Affordable Care Act (ACA) was to expand access to health insurance. Two key provisions to achieve this objective have been the establishment of the Health Insurance Marketplaces as standardized platforms to purchase health insurance coverage, and the introduction of a community rating system which limited insurers' ability to price on information related to individuals' health status.

However, this new regulation on the premium setting process amplifies information asymmetries between the demand and supply sides of the market, opening the room to an endogenous sorting of consumers based on risk type, i.e. high risk individuals demanding more insurance (Rothschild and Stiglitz, 1976). To protect the market from adverse selection, the ACA implemented a traditional government intervention by mandating individuals to purchase coverage. However, despite the tax penalty imposed on the uninsured, early Marketplace data revealed lower than expected enrollment rates, among several features consistent with the symptoms of adverse selection, causing concerns about the long run sustainability of these new health insurance markets.

Motivated by these alarming facts and recent advances in behavioral economics, the goal of this paper is to study an alternative, behavioral solution for the problem of adverse selection in the ACA Marketplaces: the automatic enrollment of the uninsured with the possibility of opting out. This idea is a recurring proposal in the ongoing health care debate in the US and its popularity stems from two main facts.[1,2] First, this behavioral policy has the potential to increase coverage rates because similar nudges have proved to be very successful in other settings, such as the market for retirement savings (Madrian and Shea, 2001). Second, this soft paternalistic policy tool does not impose any restriction on the freedom of choice, as individuals could still choose to opt out from the default (Thaler and Sunstein, 2009). Moreover, since ACA's mandate tax penalty was eliminated as part of the 2017 tax bill as of January 2019, it became more important than ever to study alternative mechanisms that could expand health insurance coverage among younger and healthier people.[3]

To study the potential effects of this nudging policy, I use a newly available individual-level panel dataset on health insurance plan enrollment and medical claims for 2013-2016 from the Colorado All Payer Claims Database. I begin by documenting reduced form evidence for the presence of adverse selection and inertia in consumer choice in this market. I show that while insurance take-up rates are low, the uninsured population became significantly healthier after the ACA's health insurance reforms in 2014. In addition, I find that following an increase in the premium level, healthier consumers are more likely to drop out from the market, consistent with the predictions

[1] "... states would be responsible for designating several insurance plans as default options to which these individuals would be assigned on a random basis if they failed to sign up for coverage on their own."(Patient CARE Act, 2014)

[2] "States will have the option to auto-enroll individuals. If auto-enrollment is selected, individuals will be allowed to opt-out of coverage. The auto-enroll feature eliminates the need for either an individual or employer mandate." (Patient Freedom Act, 2017)

[3] This paper analyzes ACA Marketplace data from 2014 through 2016, when the mandate tax penalty was in effect.

of adverse selection. I also provide descriptive evidence for a strong choice persistence in this market. I show that while the choices of new enrollees reflect current market conditions, individuals with incumbent health plans – despite the significant changes in the choice set over time – do not tend to update their choices (as in Handel (2013)), suggesting the presence of switching frictions.

I then develop a stylized theoretical model, which combines switching costs and adverse selection on the extensive margin (between uninsurance and enrollment) within the same framework. The model shows that under the current regulation with active enrollment, the switching cost generated by the cognitive effort costs of the enrollment process amplifies adverse selection by weakening the incentives of healthier consumers to sign up. The theoretical model also provides insights about the potential effects of the auto-enrollment policy. Intuitively, changing the design of the default reverses the direction of the switching friction, and the resulting increase in the relative price of uninsurance generates higher enrollment rates. I also show that the marginal enrollees are healthier than the current risk pool and exert a positive externality on the entire market by reducing the premium level, generating further increases in enrollment.

Based on the insights of the theoretical framework, I estimate a structural model of health insurance demand and supply in the presence of consumer inertia using methods from the empirical industrial organization literature (Handel, 2013; Ho et al., 2017). For demand estimation I use a discrete choice framework to model consumers' valuation of insurance products as a function of health insurance plan characteristics. For the identification of the demand parameters, I take advantage of the panel structure of the data and the variation in enrollment status, and I compare the choices of different cohorts over time. This strategy allows me to separately identify two types of switching costs: enrollment costs that represent the effort costs of completing the administrative procedure of enrollment, and decision costs that are associated with cognitive costs of evaluating these complex financial products. As the benefit of the auto-enrollment policy depends on the level of consumer inertia, understanding the micro-foundations of choice persistence allows me to predict the impacts of the counterfactual policies more accurately. To close the model, I employ machine learning techniques to build an empirical model of insurers' contract pricing decisions that fits well to the observed supply side dynamics of the market.

Using the structural parameter estimates of the empirical model, I simulate changes in consumer enrollment, plan choice and premiums under the counterfactual auto-enrollment policy. The results predict that changing the default from opt-in to opt-out can increase enrollment rates by over 60%. The decline in the severity of adverse selection is reflected by the younger and healthier risk pool: the average age of enrollees decreases by 4 years, leading to a 13% drop in mean annual spending. As a result, the average Silver plan premium declines by $300.

In addition to quantifying the overall impacts of this behavioral policy, the simulations also shed light on the importance of choosing the right default plans for auto-enrolled individuals. These exercises are useful to understand which potential policy would be the most effective in terms of reducing adverse selection. Specifically, I compare the benefits of auto-enrolling everyone in the same low-coverage contract type (naive default) to a more sophisticated policy, where the default plan varies across individuals based on observable demographics (smart default). I find that naive defaults generate more opt-out to uninsurance due to the inferior matching quality. The simulation results also reveal an interesting trade-off in case of naive default assignment algorithms: an increase in total enrollment rates and reduced adverse selection on the extensive margin, however a more acute selection on the intensive margin, i.e. sorting of enrolled consumers across plans with different coverage generosity.

The main mechanism driving this result is that the large inflow of healthier consumers into the low coverage plans increases the relative price of higher-coverage contracts by improving the risk pool of the default plan. This change in relative prices generates incremental selection by causing some of the existing consumers to switch to lower-coverage plans.

In contrast, smart default policies, which take into account the heterogeneity of preferences during the default assignment, perform better on both dimensions: they not only lead to higher enrollment rates overall, but also spread out enrollment across different contract types. Therefore, I find that personalized smart defaults can increase coverage rates among healthier people without adversely affecting the allocation of risk inside the market, and lead to better coverage among the enrolled.

In sum, this paper shows that simple behavioral policies have the potential to increase enrollment rates and maintain the stability of the private individual health insurance market in the long run. These results provide important new insights for health care policy design, especially given the recent repeal of the most important stabilizing tool of the ACA, the individual mandate.

This paper contributes to several major areas of the literature. First, it connects to the large literature studying the industrial organization of health insurance exchanges. Most of this literature focuses on the Massachusetts Health Insurance Exchange that served as a model for the ACA Marketplaces (Hackmann et al., 2015; Ericson and Starc, 2015; Jaffe and Shepard, 2016). Recently, as individual-level ACA Marketplace data becomes more accessible for researchers, a growing number of papers study different aspects of these new markets. For instance, Tebaldi (2017) and Orsini and Tebaldi (2016) investigate the supply side incentives created by the premium subsidy mechanism and the age-based pricing rule using data from California's exchange. Related to adverse selection, Panhans (2018) provides evidence for extensive margin sorting by exploiting geographic variation in premiums generated by rating area boundaries. Diamond et al. (2019) analyze the spending behavior of drop-out consumers.

Second, the paper contributes to the wide literature studying consumer switching frictions in health insurance markets. Many papers have documented the negative consequences of suboptimal switching behavior in the context of Medicare Part D, a privately provided publicly subsidized drug insurance program designed for the elderly (Abaluck and Gruber, 2011, 2016; Ho et al., 2017; Polyakova, 2016*a*). For instance, Polyakova (2016*a*) studies the interaction of switching costs and risk sorting across contracts and finds that policies that target the elimination of switching costs reduce the severity of adverse selection. However, in an employer-sponsored health insurance setting, Handel (2013) shows that the reduction of choice persistence might create incremental adverse selection. These mixed empirical findings are due to the fact that theoretically the net effect of the interaction of switching costs and adverse selection is ambiguous, as shown by Handel et al. (2019). The main contribution of this paper to this strand of literature is modeling the interaction of extensive margin adverse selection and switching costs, and showing how changing the default from opt-in to opt-out can manipulate the direction of this relationship. Moreover, most of the empirical literature treats switching frictions as a black box without attempting to understand their micro-foundations. The current paper addresses this gap by estimating of two different sources of consumer inertia in the context of ACA Marketplaces.

There has been an extensive literature analyzing adverse selection in insurance markets since the seminal work of Rothschild and Stiglitz (1976). Einav et al. (2010*a*) provide a review of the most influential papers of this area focusing on health insurance. The most related to this paper is Einav et al. (2010*b*) who propose an empirical method to estimate the loss from adverse selection by using variation in premiums. I contribute to this literature by incorporating switching costs into this classic framework to illustrate how behavioral interventions can affect the direction of the relationship between adverse selection and switching costs.

Finally, this paper connects to the behavioral economics literature, especially to studies analyzing default effects. In one of the most seminal papers of this literature, Madrian and Shea (2001) look at the effects of automatic enrollment to the 401(k) retirement saving program and show that changing the default from opt-in to opt-out almost doubled participation rates. However, as the survey of Chandra et al. (2019) on behavioral health economics points out, very few papers study smart defaults and their potentials for designing nudging policies in health insurance markets. Ericson (2016) and Handel and Kolstad (2015) are two notable exceptions. This paper addresses this gap in the literature by studying auto-enrollment as a potential mechanism to expand health insurance coverage and contrasting the effects of naive versus personalized default designs.

The rest of the paper is structured as follows. Section 2.3 describes the institutional setting and the data. Section 1.3 presents motivating reduced form evidence for the presence of adverse selection and consumer choice frictions. Section 1.4 introduces a theoretical framework to illustrate how adverse selection and switching costs interact in the ACA Marketplaces. Section 1.5 outlines the empirical model of the insurance market and reports the estimation results. Section 1.6 presents the simulation results and discusses the policy implications. Finally, section 1.7 concludes.

1.2 Institutional Background and Data

The Patient Protection and Affordable Care Act, signed into law in 2010 by President Obama, was the most significant healthcare reform in the history of United States since the passage of the public health insurance programs, Medicare and Medicaid in 1965. The law was aimed to help more people get affordable health insurance coverage, receive better medical care, and reduce the growth of health care spending in the U.S.

To increase the historically low health insurance coverage in the US, a key element of the health care reform was the establishment of the Health Insurance Marketplaces.[4] The role of the Marketplaces is to provide a portal for individuals, not covered by employer-sponsored insurance and not eligible for Medicaid, to purchase qualified healthcare plans that are highly standardized in terms of benefits. The ACA Marketplace plans are provided by private insurance companies, with premiums subsidized by the federal government in form of advance tax credits (APTC) for low income people with earnings up to 400 percent of the federal poverty level.

Another relevant key reform of the ACA is that health insurance premiums can no longer be determined based on individuals' pre-existing medical conditions or projected health status. Instead, insurers are required to set prices based on community rating,

[4]The ACA also expanded the Medicaid program, established health insurance exchanges for small businesses and introduced an employer mandate for those with 50 or more full-time employees. This paper focuses on the elements of the law that affected the private individual health insurance market. Note that the mandate for employers with more than 200 employees was implemented through auto-enrollment with possible opt-out (Kaiser Family Foundation, 2017).

relying on no other information than age and geographic location.[5] Moving to a community rating system with guaranteed issue is another tool of the ACA to decrease the uninsured rate, since prior to the health care reform a common reason for remaining uninsured was that many people with pre-existing health conditions were refused coverage or priced out from the individual health insurance market.

Although an insurance market with such restrictions on the premium setting process protects consumers from reclassification risk, these rating regulations also amplify information asymmetries between the demand and supply sides of the market, driving healthier consumers out of the market (Handel et al., 2015; Fleitas et al., 2019). To overcome this classic tradeoff between adverse selection and reclassification risk, the ACA mandated everyone to enroll in a health insurance plan that meets the criteria of the minimum essential coverage or to pay a tax penalty.[6]

Although the benefits of the health insurance plans are standardized, shopping at the Marketplaces is a complex choice problem for consumers since they have to understand and compare several different dimensions of these financial products. First, there are four metal tiers (Platinum, Gold, Silver and Bronze) based on the advertised actuarial value, i.e. the expected share of health care expenses paid by the insurance company.

[5]The ACA also introduced a federally established age curve with maximal age rating ratios of 3:1, however states can use own age curves with less than the standard ratios. In some states, insurers can also charge higher premiums for tobacco users with no more than a 1.5:1 rating ratio (CMS, 2017).

[6]Note that the penalty could not exceed the national average premium of the most basic type plans (Bronze tier) sold on the Marketplaces, however with subsidies, some consumers could purchase a plan for less than the cost of the penalty.

Within these categories, there are different types of plans based on the accessibility and size of the health care provider networks. Furthermore, plans also differ in terms of financial characteristics, such as premiums, deductibles and out-of-pocket costs. Comparing the elements of the choice set based on all these dimensions to find the best match requires consumers to make costly cognitive efforts. Since the Marketplaces are primarily designed to extend health insurance coverage among the previously uninsured population, it is important to note that a large fraction of the potential buyers may have very limited prior experience with similar complex choice problems in the context of private health insurance (Kaiser Family Foundation, 2014).

Finally, another relevant feature of the Marketplaces is the default automatic renewal of the previous years' choice. This means that if individuals are enrolled in a health insurance plan and do not actively search for a new option for the next year, they will be assigned to their previous choice by default. Since it is a well known observation in behavioral economics that defaults can significantly influence consumer choice, in markets where the choice set changes dynamically over time, auto-renewal might significantly contribute to consumer choice inefficiencies.

1.2.1 Data

The main data source of this paper is Colorado's newly available All Payer Claims Database (APCD) (CIVHC, 2017). This large administrative database collects information from all insurance companies operating in the state. The extract I use is an individual-level panel dataset of health insurance plan enrollment and medical claims for 2013-2016.

This dataset consists of two major components. The first part of the data contains de-identified individual-level enrollment data including all Colorado residents covered by health insurance, regardless of whether they submitted any health insurance claim during the sample period. This enrollment file includes detailed information on the plan choice, such as the name of the insurance company, the market category of the plan (individual, employer-sponsored, etc.), the product type of the plan (HMO, PPO, EPO), the Marketplace metal tier, and the enrollment period. I also have access to individual-level demographics, such as age, gender and zip code. The second main part of the dataset is the medical claim file. This file includes essential information for the cost side of my analysis since it contains data on diagnosis codes and health care spending (financial characteristics of each claim).

The Colorado APCD has several features that make it ideal for analyzing the research question of this paper. To quantify the potential benefits of the automatic enrollment policy, I need to know by how much enrollment would increase, and how this change in enrollment would affect insurers' costs. I answer the first question using the enrollment file that allows me to estimate consumer inertia and preferences for health plans. To address the change in the risk profile of the enrollment pool, I use the linked individual-level health care spending data obtained from the claims file. Furthermore, as the Marketplaces started to operate in 2014, the dataset also allows me to track the choices of the Marketplace enrollees over time, starting from their initial decisions. This feature of the data will be essential for the identification of switching costs. Finally, having access to 2013 enrollment and claims data allows me to compare many important aspects of the individual health insurance market before and after the introduction of the ACA's major insurance reforms.

Although the APCD provides detailed data on many aspects of the market, it contains very limited information on the characteristics of the choice set. Therefore, the second dataset I use is a self-collected database on health insurance plan characteristics that was constructed by downloading the summary of benefit forms of each health plan in the choice set from the System for Electronic Rates and Form Filings (SERFF). This dataset contains information on the premium and coverage generosity of each plan, such as deductibles, out-of-pocket limits, co-insurance and co-payment amounts. Note that since information on premiums comes from the plan-level files, individual-level premium subsidies are unobserved. Furthermore, the state of Colorado requires insurance companies to submit network adequacy reports to the Division of Insurance. These files allowed me to measure the breadth of provider and facility networks for each health plan sold on the Marketplace during the sample period. Then I linked these plan-level datasets to the APCD in order to obtain an enrollment dataset with a rich set of plan characteristics for modeling consumers' health plan choices.

I also use individual-level demographic information about the uninsured population from the American Community Survey (Ruggles et al., 2017) to construct a sample of potential buyers for the counterfactual simulations. The last data source I use for my analysis is the Colorado Health Access Survey that provides detailed individual-level information on health insurance coverage and the use of health care services in the state (Colorado Health Institute, 2018). This dataset allows me to study changes in the health status of the uninsured after the introduction of the Marketplaces in 2014 in order to provide descriptive evidence for the presence of extensive margin adverse selection in this market.

1.2.2 Descriptive facts

Enrollment. I begin the descriptive analysis of Colorado's individual health insurance market by exploring enrollment trends and changes in the number of uninsured in the state. Figure 1.1 shows that the number of Marketplace enrollees increased continuously during the period, reaching 150,000 by 2016, the last year I observe. In 2013, the year before the Marketplace opened, about 730,000 Colorado residents had no health insurance coverage, resulting in an uninsured rate of 14%.[7] Figure 1.1 also shows that during the sample years, the number of uninsured declined substantially.[8] By 2016 the number of people with no health insurance coverage dropped to 415,000, reducing the uninsured rate to 7% in the state (Ruggles et al., 2017).

[7] Colorado's uninsured rate was close to the national level (14.5%) in 2013, when Massachusetts had the lowest (3.7%) and Texas highest (22.1%) rates (Ruggles et al., 2017).

[8] Note that Colorado has also expanded the Medicaid program as of January 2014 (Colorado Health Institute, 2018).

Panel A of Table 1.1 presents the state-level enrollment shares of the different metal categories on the Colorado Exchange. The most popular choices are the Silver and Bronze plans with actuarial values defined at 70% and 60%, respectively. A special feature of Colorado's ACA Marketplace is that Bronze plans have a much higher market share (about 40%) compared to the national share of this metal tier (20%). The large market share of Bronze plans in Colorado relates to multiple facts. First, the state has one of the highest premium levels nationally (Rau, 2014). Second, Colorado is a relatively high income state, therefore a lower share of the population qualifies for the ACA premium subsidies than nationally (54% vs. 80%) (C4HCO, 2017).[9] Hence, buying even a lower cost Bronze plan imposes a large financial burden on consumers in this state.

Supply side. Colorado's ACA Marketplace was quite dynamic in terms of the entry and exit of insurers during these early years. Panel B of Table 1.1 shows the participating insurance companies and their market shares. In 2014 ten companies offered plans through the Marketplace. In 2015 two new insurers entered the market and two other companies left in 2016. Since the table shows insurer participation at the state-level, it does not fully reveal the dynamics of the market. In fact, insurance companies decide on entry and exit at the rating-area level, defined as sets of counties. Therefore, many of these local markets experienced even larger variation in the choice set due to changes in the number of insurers and offered plans during these years.

[9] The high income of the state also implies that a lower share of potential enrollees qualify for cost sharing subsidies, which only apply to Silver plans.

The market share data show that the Colorado Marketplace is dominated by a few major insurance companies. Kaiser has the highest market share, over 40% percent in each year of the sample. Furthermore, many of the smaller insurers serve only a limited number of regions. For instance, Denver Health sells plans only in the Denver area. As a result, most of the local markets are even more concentrated than suggested by the state-level shares.

Rating areas. Rating areas are important concepts of the ACA because they define the boundaries of local health insurance markets. Furthermore, the newly introduced community rating system requires insurance companies to offer the same premium to individuals of the same age living in a given rating area. Figure 1.2 displays the current borders of the rating areas in Colorado. Originally, there were 11 rating areas in 2014 when the Marketplace opened, however, due to the very high premium levels in the mountain areas, regulators decided to merge some regions from 2015. This division of the state generates important variations both on the demand and the supply sides of the market.

Figure 1.3 presents two dimensions of the variation in demand across the rating areas. On the one hand, panel 1.3a shows substantial differences in the distribution of the Exchange enrollees across different parts of the state. Enrollment is very low in most regions, while over 50% of the total enrollees are concentrated to the central Denver rating area. On the other hand, panel 1.3b shows that the average age of the enrollment pool in the central area of Colorado is much lower than in the rest of the state. Comparing these two figures suggests that in the eastern and western mountain areas, the risk pool is smaller and sicker than in the central parts of the state.

Figure 1.4 provides a summary of supply side variation across the rating areas. Consistent with the patterns documented on the demand side, panel 1.4a shows that insurance companies prefer to enter in the central area of the state where the composition of the enrollment pool is more balanced in terms of size and age. The regional variation in the average premium of Silver plans, shown in panel 1.4b, also coincides with the previously documented facts: the mountain regions are much more expensive than the central areas of the state.

These patterns provide a first piece of evidence consistent with the presence of adverse selection in this market. In the next section I provide additional reduced form evidence for this inefficiency.

1.3 Reduced Form Evidence

The stability of the ACA Marketplaces is a central question of the health care debate in the US. Much of the political and media attention was due to a series of alarming facts from the early years of the market, such as lower than expected enrollment rates, rising premiums and the exit of large insurers, who blamed the high risk profile of the market segment (Abelson, 2016). These facts are consistent with the presence of adverse selection, suggesting that the mandate tax penalty was not strong enough to incentivize young and healthy people to purchase coverage. In this case, an alternative behavioral policy that can increase insurance take-up rates among the younger and healthier population may improve the stability of the market in the long run.

Furthermore, as I will show later, the benefits of auto-enrolling the uninsured are determined by how much this nudging policy would expand health insurance coverage. These effects of the policy depend on how large consumer inertia is that keeps out uninsured people from the market under the current regulation with active enrollment; and whether this choice persistence is strong enough to keep the auto-enrolled individuals in the market under the counterfactual policy. Hence, measuring consumer inertia is a crucial step to evaluate the potential benefits of this behavioral policy.

Therefore, in this section I document reduced form evidence consistent with the presence of adverse selection and switching frictions in this market in order to motivate the analysis in the rest of the paper.

1.3.1 Adverse Selection

The rating area-level variation in enrollment, age, market concentration and premiums presented in Section 2.3 was the first descriptive pattern consistent with the presence of adverse selection in this market. I now provide further reduced form evidence for adverse selection, focusing on sorting on the extensive margin, the decision between uninsurance and enrollment.

As a first step, Figure 1.5 compares the age distribution of the Exchange enrollees and the uninsured population. The graph shows that while most of the uninsured are younger with a median age of 36 (mean uninsured age is 37.5), the median Exchange enrollee is ten years older (mean enrollee age is 44.4). Note that under the ACA rating regulation, insurers cannot fully price on age, which correlates with health care expenditures.[10] The figure also reveals a large variation in insurance take-up rates across age groups. Consistent with adverse selection, take-up is almost complete in older age groups and much lower among the younger uninsured.

The simplest formal tests for adverse selection are based on a comparison of average costs across contracts with different coverage generosity (Chiappori and Salanie, 2000; Einav et al., 2010*a*). In this case, the ideal positive correlation-type test would compare health care utilization of the uninsured to that of the Marketplace enrollees. However, applying these classic simple tests to detect extensive margin adverse selection in this market is not easy because of data availability problems. The reason is that very limited information is available about the uninsured population, not to mention their health care spending.

I present a new piece of evidence for the presence of adverse selection on the ACA Marketplaces using individual-level information on the health status of the uninsured

[10] Colorado uses the federal age rating curve set by the CMS that limits the premium ratio of a 64 year old individual old to a 21 year old to be 3:1 (CMS, 2017).

from the Colorado Health Access Survey (CHAS). This survey tracks individuals over time and collects information on insurance coverage type and different health care utilization measures. In my analysis I rely on two important features of this dataset: the data not only covers the uninsured population but the time dimension of the panel dataset also allows me to compare changes in the health of the uninsured and the individual market enrollees after the ACA's insurance reforms in 2014.

Table 1.2 presents the change in the health status of the uninsured between 2013 and 2015, before and after the Marketplace opened in 2014. I use two survey questions to proxy for health status: a self-reported health condition and an indicator for visiting any health care provider in the past 12 months. The first row of the table reveals that a smaller share of the uninsured population reported poor health in 2015 than in 2013. However, this improvement in the health of the uninsured might be driven by increased health care utilization due to different ACA reforms. The next row addresses this concern by showing a decline in the fraction of uninsured visiting any health care provider. These patterns are consistent with the self-selection of higher risk types to become insured.[11] The table also reveals a large heterogeneity in the health status changes of different age groups: the overall improvement in the health of the uninsured is mainly driven by a large change in the 35-54 age group. This evidence combined with take-up rates varying by age (as shown in Figure 1.5) suggests an acute selection among the middle-aged uninsured. The more stable health status of the uninsured in

[11] Jacobs et al. (2016) found similar results using data from the National Health Interview Survey.

other age categories might be related to higher take up rates in the 55-64 cohort and the limited use of health care services by young adults.

However, the improved health of the uninsured might reflect a general trend in the population, instead of being specific to this group. Therefore, I use a difference-in-differences regression analysis to compare the health status of the uninsured and the individual market enrollees before and after the ACA insurance reforms in 2014.[12] The regression specification I estimate takes the following form:

$$poor\ health_{it} = \beta_0 + \beta_1 uninsured_{it} + \beta_2 after + \beta_3 uninsured_{it} * after + \varepsilon_{it}, \quad (1.1)$$

where $poor\ health_{it}$ is an indicator for individual i reporting poor or fair health in year t.[13] Column (1) of Table 1.3 shows that the sign of the interaction term's coefficient estimate is negative, suggesting that the uninsured became healthier after the Marketplaces opened in 2014. Column (2) breaks down this interaction term by age and reveals that this improvement in the health of the uninsured is the largest for the 35-54 age group, consistent with the descriptives in Table 1.2. However, it should be noted that comparing the uninsured population to the individual market

[12] The CHAS data does not distinguish between different types of individual market enrollees, therefore this category includes both Marketplace enrollees and consumers with other individual market plans.

[13] The 2009-2013 waves of the CHAS pool fair and poor health categories together, while the 2015-2017 waves report them separately.

enrollees is problematic due to different factors correlated with both insurance status and health risk. To control for these factors, I re-estimate the model on a matched sample based on propensity scores calculated using income, education, household size and employment status. The parameter estimates in column (3) show that the main result is robust to these observable confounders.

As a final exercise to detect adverse selection in this market, I use the time dimension of the APCD data to analyze the relationship between the change in premiums and the costs of insurers. The theory of adverse selection predicts that higher risk individuals are likely to demand better coverage, implying that the highest willingness to pay buyers are the sickest consumers. Therefore, as the premium level increases, healthier people would be the first to drop out from coverage and switch back to uninsurance. I use the increase of the premium level over time in this market to test this prediction empirically. The regression specification I estimate is the following:

$$P(exit_{it} = 1) = \Psi(\beta_0 + \beta_1 cost_{it-1} + \beta_2 premium_{it} + \beta_3 premium_{it} * cost_{it-1}), \quad (1.2)$$

where $exit_{it}$ is an indicator for consumer i dropping out from the Marketplace in year t, $cost_{it-1}$ denotes her yearly health care spending in her last enrollment period, $premium_{it}$ is the premium increase of her default plan, and Ψ denotes the cumulative logistic distribution. The estimation results in Table 1.4 show that higher spenders are less likely to drop out from the market and an increase in the price of the default plan also increases the probability of leaving. The sign of the interaction term suggests that following a premium increase, higher spenders are less likely to leave the market, consistent with the presence of adverse selection. However, there might be different consumer types based on the reason behind the exit decision: while some individuals switch to uninsurance, some of them might obtain employer-sponsored coverage. Since the type of exit can be correlated with health care spending, this might bias the previous estimates. To address this endogeneity concern, I take advantage of the fact that the data that allows me to track individuals switching plans over time and across markets. Therefore, the final specification is column (3) repeats the previous analysis but excludes individuals who switch to employer-sponsored plans.[14] The results do not change qualitatively after excluding these consumers.

These reduced form results are consistent with the presence of adverse selection in this market. For a formal proof of extensive margin adverse selection on the Colorado Marketplace using exogenous variation in prices see Panhans (2018).

[14] My APCD extract includes only commercial plans, therefore I cannot distinguish consumers leaving to uninsurance and Medicaid. However, Medicaid eligibility is likely to be positively correlated with health risk (reduced income due to the deterioration of health, or the reverse), therefore this would bias the cost estimates toward zero.

1.3.2 Inertia in plan choice

In this section, I document reduced form evidence for the presence of consumer inertia in this market. It should be noted that analyzing switching frictions requires variation in the choice set over time since otherwise it would be a natural consequence that individuals choose the same health plan in each year. However, as shown in Section 1.2.2, the choice set was quite dynamic during these early years of the market: insurance companies entered and exited, changed the number of contracts offered and adjusted the premiums. In such a dynamic environment, an active re-optimization of plan choice would be a crucial aspect of consumer's decisions.

However, despite these significant changes in the market environment, the enrollment data in Table 1.5 show that consumer choice was persistent: over 70% of the Marketplace enrollees kept their plan for another year. The table also reveals a large increase in the number of switches from 2015 to 2016. The reason of this increase is that two insurance companies, Colorado Access and Time exited the Colorado Marketplace in 2016, and Rocky Mountain Health Plan reduced its Marketplace participation drastically by exiting from most rating areas. These insurer exits generated many forced switches in 2016, and these these types of decisions will be important for the identification strategy discussed in Section 1.5.1.

Next I investigate whether there are any systematic differences among the contract choices of different enrollee cohorts. This is a commonly used approach in the literature to detect switching frictions (Handel, 2013; Polyakova, 2016*a*). Specifically, I compare the choices of cohorts over time who entered to the Marketplace in different years.[15] Table 1.6 shows that the choice of each cohort in the years following their initial enrollment reflects the market conditions in the first enrollment period. For instance, among individuals who first enrolled in 2014 the Silver tier was the most popular choice in each year, although new enrollees in 2015 and 2016 were more likely to choose Bronze plans due to the substantial increase in premiums over time.[16] Analyzing insurers' shares reveals a similar picture: 59% of consumers who first entered to the market in 2014 chose Kaiser and this payer kept its dominance within this cohort through the following year although newly entering consumers in the subsequent years were more likely to choose Anthem due to changes in relative premiums. Table 1.6 also shows that metal tier shares are more stable within enrollee cohorts than the choice of the insurer. These patterns together with the substantial changes in the choice set over time provide descriptive evidence consistent with the presence of switching frictions.

Finally, I turn to the analysis of factors that influence the probability of switching. I estimate an individual-level logistic regression of switching status on demographics and rating area-level choice set characteristics. The estimation results are presented in

[15]Individuals who were forced to switch their plans any time during their enrollment spell due to the insurer's exit are excluded from this analysis.

[16]An alternative explanation for the same pattern could be that that the preferences of 2014 cohort are different due to their worse health status since they purchased coverage immediately after the ACA insurance reforms banned pricing on pre-existing conditions. However, Table A.3 in the Appendix shows that the enrollee cohorts are similar based observable demographics, and the 2014 new enrollees are not the highest spenders among the three groups.

Table 1.7. The parameter estimates in column (3) show that females, older and higher income enrollees are more likely to switch. Experiencing a negative health shock in the previous year also increases the probability of switching, similarly to an increase in the current plan's premium.[17]

Table 1.7 shows that people who face a smaller choice set are more likely to switch. This result is likely to be the net effect of different forces. First, a larger choice set might increase the frequency of switching since the better matches increase the expected benefit of doing so. However, as the number of options increases, the decision problem becomes more complex, and a wide range of empirical evidence shows that people often avoid such choices. Finally, in this market the size of the choice set is likely to be correlated with the severity of adverse selection in a given rating area, as documented by the descriptive facts in Section 2.3. The dynamic consequences of adverse selection can generate changes in the choice set over time, which in turn affect the probability of switching. Finally, the table also shows that the exit of insurers correlates with more frequent switches, as expected.

Overall, these reduced form results suggest the importance of allowing for consumer heterogeneity in choice models with switching frictions.

[17] This result is consistent with Abaluck and Adams (2017), who show that the probability of being attentive to changes in the choice set depends on the characteristics of the default plan.

1.4 Model

In this section I build a stylized theoretical model to illustrate the intuition of how adverse selection and switching costs interact in this market.

Now consider consumer i who chooses whether to sign up for insurance plan j based on her utility derived from the plan. This utility is given by the quasi-linear function

$$U_{ij} = V(c_i, \Phi_j) - p_j(\bar{c}) - \eta. \tag{1.3}$$

The term $V(c_i, \Phi_j)$ is individual i's valuation of plan j. This value depends on the expected health costs and of agent i (c_i) and the characteristics of plan j, denoted by Φ_j. As agents with higher expected health cost gain more value from insurance, $V(c_i, \cdot)$ is a strictly increasing function. The term p_j denotes the premium charged for plan j which is assumed to be increasing in insurers' average costs ($\bar{c}$) and defined as the net premium (p) faced by consumers after the mandate tax penalty. Parameter η denotes the switching cost capturing the effort cost of understanding the characteristics of plans and completing the administrative procedure of enrollment.

Note that for illustrative purposes I now assume that η is constant across consumers and I do not model different types of switching costs separately. For tractability, I analyze a market with a single product, perfect competition, and linear demand and cost curves as in Einav et al. (2010*b*). These assumptions will be relaxed later in the empirical model.

Additionally, I assume away both moral hazard (c_i does not depend on j) and heterogeneity in risk preferences (V depends only on c_i). Considering moral hazard would not alter the predictions of the theoretical model since if its size was the same across potential enrollees, it would only shift the costs up without altering the distribution of consumers. If moral hazard was positively correlated with health risk, the model would produce the same results qualitatively, and a negative correlation is not likely. Furthermore, moral hazard can also be regarded as part of c_i since people might know their change in spending at the time of the enrollment decision and select contracts based on their expected change in care utilization (for a detailed discussion of selection on moral hazard see Einav et al. (2013)). Heterogeneity in risk preferences would only qualitatively change the predictions of the model if risk aversion was so negatively correlated with health risk that it would lead to advantageous selection, as discussed in Cutler et al. (2008). However, the descriptive evidence suggests that this is not the case in this market.

1.4.1 Active Enrollment

Under the current regulation where the default is uninsurance and active sign-up is required for enrollment, individual i chooses to enroll to plan j if $U_{ij} \geq 0$, that is:

$$V(c_i, \Phi_j) \geq p_j(\bar{c}) + \eta. \tag{1.4}$$

Then, for all plans j there is a threshold value c^* such that $V(c^*, \Phi_j) = p_j(\bar{c}) + \eta$. Each consumer with expected health cost $c < c^*$ will choose not to sign up and each individual with $c \geq c*$ will enroll in a plan. The presence of adverse selection is shown by the fact that only individuals with expected health care costs higher than the cutoff enroll.

Note that adverse selection is even more severe in the presence of switching costs: in a frictionless world where $\eta = 0$, the threshold value c_f^* is given by

$$V(c_f^*, \Phi_j) = p_j(\bar{c}), \tag{1.5}$$

and since $\eta > 0$ and $V(c_i, \cdot)$ is strictly increasing, we must have

$$c^* > c_f^*. \tag{1.6}$$

1.4.2 Automatic Enrollment

Now I consider a counterfactual nudging policy that automatically assigns each eligible uninsured to a set of default insurance plans but allows them to opt-out to preserve the freedom of choice, in the spirit of Thaler and Sunstein (2009).

Under this counterfactual policy, opting out will require individuals to make some costly efforts, such as the cognitive cost of comparing their default plans to their outside options or completing the administrative steps of the cancellation process. These cancellation costs are denoted by the same parameter η. The key intuition of the model is that under this counterfactual policy, switching costs work in the opposite direction and help increase enrollment rates in the insurance market.

Now consider consumer i who is enrolled in plan j. She decides to opt out if the sum of her valuation of the plan and her switching cost is smaller than the premium she is paying. Otherwise she decides to remain enrolled (when the premium she has to pay is smaller than the sum of her valuation of the plan and the switching cost she has to pay). Therefore, individual i opts out of plan j if her utility from opting out is non-negative. Since

$$U'_{ij} = p_j(\bar{c}) - V(c_i, \Phi_j) - \eta, \tag{1.7}$$

$U'_{ij} \geq 0$ holds if

$$V(c_i, \Phi_j) \leq p_j(\bar{c}) - \eta. \tag{1.8}$$

As in the previous case, this condition also gives a cutoff value, c_d^* with $V(c_i, \Phi_j) = p_j(\bar{c}) - \eta$. Therefore, individuals with expected health costs below the cutoff choose to opt-out, while agents with expected health costs above the cutoff remain enrolled. Again, since $\eta > 0$ and $V(c_i, \cdot)$ is strictly increasing we have

$$c^* > c_f^* > c_d^*. \tag{1.9}$$

As the cutoff values represent the enrollee with the lowest expected health cost under each scenario, a higher threshold implies a more acute adverse selection problem on the market. Note that c_f^* represents the cutoff for the case of no switching frictions, where manipulating the defaults does not make any difference. Therefore, the higher cutoff value corresponding to the current regulation with active enrollment implies that the presence of switching costs amplifies adverse selection compared to the frictionless case. In the counterfactual with auto-enrollment and possible opt-out, however, the presence of inertia helps reducing adverse selection: the expected health care spending of the marginal enrollee shifts down and the average cost of the risk pool decreases. As the premium level is determined by insurers' average costs, this change in enrollment generates a lower premium level on the market compared to the current regulation with active enrollment. Therefore we have

$$p_j(\bar{c}) > p_j(\bar{c}_f) > p_j(\bar{c}_d). \tag{1.10}$$

It is important to point out that the reason of this reduction in adverse selection is that some marginal consumers are locked into being covered under the counterfactual policy due to switching costs. For some individuals we have $V(c_i, \Phi_j) \leq p_j + \eta$ but $V(c_i, \Phi_j) > p_j - \eta$, meaning that these individuals would not sign up actively, but they remain insured under the counterfactual policy. First, consumers with $p_j < V(c_i, \Phi_j) < p_j + \eta$ are better off in the counterfactual, since they would be insured in a frictionless environment. However, an interesting feature of my model is that even some of the consumers with valuations $p_j - \eta < V(c_i, \Phi_j) < p_j$ are better off in the counterfactual although they would choose to remain uninsured in a choice frictionless world. The reason is that the increasing enrollment rates reduce the severity of adverse selection and the average cost of the risk pool, resulting in lower premiums compared to choice frictionless world. In other words, the healthier marginal enrollees exert a positive externality on the entire market. As a result of this price adjustment mechanism, it becomes efficient to insure some of the locked-in consumers who would not be insured in a frictionless environment because of the high premium level due to adverse selection.

1.4.3 Graphical Analysis

To illustrate the price adjustment mechanism, I introduce a graphical representation of the insurance market, building on the framework developed by Einav et al. (2010*b*).

Figure 1.6 shows the case of an adversely selected insurance market with no switching frictions. The horizontal axis represents the fraction of people enrolled on the Marketplaces. The graph shows that even in a frictionless environment, there is adverse selection in this market. Adverse selection is represented by the downward sloping marginal cost curve. In this case individuals with the highest willingness to pay for insurance are also those who have the highest marginal costs to insure, therefore as the premium falls, more and more individuals with lower marginal costs will sign up for coverage. Due to asymmetric information or rating restrictions, the competitive equilibrium is determined by the intersection of the demand and average cost curves, resulting in under-insurance (q) compared to the efficient rate (q_{eff}).

I then introduce a switching cost to the baseline frictionless model to illustrate how it interacts with adverse selection. Figure 1.7 represents the current regulation where individuals have to make costly cognitive efforts in order to actively sign up for coverage on the Marketplaces. The graph shows that in this case the demand curve shifts down by the amount of the switching cost and this leads to a more acute selection problem with lower enrollment rates, sicker enrollees and a higher premium level. The intuition is that under active enrollment, the switching cost makes the outside option – uninsurance – more attractive.

Figure 1.8 illustrates the effects of the auto-enrollment policy. In this case the switching cost works in the opposite direction, shifting the demand curve up by making the outside option less attractive. Note that this mechanism is the same as how the mandate tax penalty works (see Hackmann et al. (2015)), however this behavioral policy achieves the same effect without imposing any restriction on the freedom of choice. This change in the direction of the switching cost results in a higher enrollment, even compared to the frictionless environment. The gain in enrollment compared to the current regulation is represented by individuals between q' and q''. The graph also reveals the indirect effect of the price adjustment mechanism in the model: consumers between q and q'' are also insured in the counterfactual although they are not insured in the frictionless case.[18] The reason is that the higher enrollment rate in the counterfactual reduces the average cost of the risk pool, allowing insurance companies to charge lower premiums, and at this new price level more consumers with lower marginal cost decide to enroll.

Note that this stylized model does not consider subsidies provided to low income consumers. Absent subsidies, the model shows that in a market with adverse selection (shown by downwards sloping cost curves in Figure 1.6) the nudging policy furthers regulators' twin goals of expanding coverage and reducing premiums.

These theoretical analyses suggest that automatically enrolling the uninsured may increase enrollment rates and reduce the premium level. However, the potential gains of the policy might be offset by the costs of the implementation. Therefore, from a policy perspective it is essential to quantify the potential benefits of this nudging policy – estimating these impacts is the goal of the remaining sections.

[18] Note that these consumers are efficient to insure since their risk premium is positive, i.e. willingness to pay exceeds the marginal cost.

1.5 Empirical Framework

In this section, I introduce an empirical framework to model consumers' choice of health insurance contracts and insurers' price setting decisions.

1.5.1 Demand for Health Insurance

To estimate the demand for health insurance plans, I use a discrete choice framework. I augment the standard model with a switching cost parameter since the descriptive analysis in Section 1.3 suggests that this friction is an important determinant of consumer choice in this market. Therefore, the utility of consumer i from plan j at time t is given by

$$u_{ijt} = \alpha p_{jt} + \beta_{it} X_{jt} + \eta_{it} \mathbb{1}\left[s_{it-1=j}\right] + \varepsilon_{ijt}, \tag{1.11}$$

where p_{jt} is the premium of plan j in year t and X_{jt} denotes the characteristics of plan j in year t: deductibles, out-of-pocket maximums, the advertised actuarial value of the Marketplace metal tier and an indicator for PPO plans. $\mathbb{1}\left[s_{it-1=j}\right]$ is an indicator function for individual i choosing the same plan as in $t-1$, therefore η_{it} is the switching cost parameter. ε_{ijt} denotes an idiosyncratic preference shock which is assumed to be independently and identically Type 1 Extreme Value distributed.

Consumers' preferences for health insurance products might show a large heterogeneity due to differences in horizontal tastes, health status and attitudes towards risk. To capture these rich individual-specific preferences for health insurance plans, I allow the parameters of plan characteristics to have random coefficients:

$$\beta_{it} = \beta D_{it} + \mu_i^{\beta}, \text{ where } \mu^{\beta} \sim N(\mu^{\beta}, \sigma^{\beta}), \tag{1.12}$$

where D_{it} is a vector of demographics including age and gender.

Furthermore, I assume that the switching cost parameter is individual-specific and can be decomposed into enrollment costs and decision costs, both dependent on demographics, in the following way:

$$\eta_{it} = \underbrace{(\eta^{enr} D_{it} + \mu^{enr}) \mathbb{1}\left[continous\ enrollee_{it}\right]}_{\text{enrollment cost}} + \underbrace{(\eta^{dec} D_{it} + \mu^{dec})}_{\text{decision cost}}. \tag{1.13}$$

The expressions in the two brackets denote enrollment costs and decision costs, respectively. Enrollment costs represent the effort costs of completing the time consuming administrative process of enrollment, while the decision cost parameter denotes the cognitive effort cost of comparing different plans in the choice set, taking into account the individual's projected health care spending and preferences, and finding the individual-specific best match. This method of decomposing switching costs is similar in spirit to the one used by Luco (2017) in the Chilean retirement investment setting.

As a result of the distributional assumption on ε_{ijt}, the choice probability of consumer i choosing plan j in year t has a closed form solution conditional on the realizations of the random coefficients:

$$L_{ijt}(\beta_i) = \frac{\exp(\alpha p_{jt} + \beta_{it} X_{jt} + \eta_{it} \mathbb{1}\left[s_{it-1=j}\right])}{\sum_{k=1}^{K} \exp(\alpha p_{kt} + \beta_{it} X_{kt} + \eta_{it} \mathbb{1}\left[s_{it-1=k}\right])}. \tag{1.14}$$

Therefore, the conditional choice probability of individual i's observed sequence of choices over time can be written as:

$$S_i(\beta_i) = \prod_{t=1}^{T} L_{ijt}(\beta_i). \tag{1.15}$$

However, due to the presence of unobserved heterogeneity in the model, the unconditional choice probability requires integrating over all values of the random coefficients:

$$P_i(\theta) = \int S_i(\beta) f(\beta|\theta) d\beta. \tag{1.16}$$

This probability integral cannot be solved analytically, therefore the model is estimated via Maximum Simulated Likelihood (McFadden, 1973; Train, 2009; Hole, 2007). The MSL estimator is given by:

$$\underset{\theta}{\operatorname{argmax}} \sum_{i=1}^{N} \ ln\left(\frac{1}{R} \sum_{r=1}^{R} S_i(\beta^r)\right), \tag{1.17}$$

where R denotes the number of draws and β^r is draw r from $f(\beta|\theta)$.

Identifying switching costs based on a revealed preference argument is challenging due to the difficulty of distinguishing them from the unobserved preference heterogeneity. The root of this problem is that switching costs and the unobserved preferences have the same consequence observationally: inertia in consumer choice. To overcome this problem, I take advantage of the panel structure of the dataset and the fact that the sample covers the first years of the market. These features of the data are essential elements of my identification strategy since I can track the choices of enrollees over time, starting from the initial choices when they first purchased a Marketplace plan. This structure allows me to compare the observed choices over time and across different enrollee cohorts.

Specifically, under the current regulation, initial choices reflect active choices for which all plans incur some enrollment and decisions costs because there is no default assignment. Therefore, the initial choices of each enrollee – together with the variation in the choice set – allows me to identify the unobserved preference heterogeneity which is captured by the random coefficients in the model. Besides initial choices, I also employ a market-specific phenomenon for the identification of preferences: the frequent exit of insurers and the gradual reduction of the product mix. These supply side dynamics force some consumers to switch their plans in order to remain enrolled after their incumbent plan is no longer available on the market. Therefore, similarly to initial choices, these forced decisions are also active choices that can be used for preference identification since they reflect consumers' tastes given the market conditions rather than a structural state dependence.

As the descriptive analysis suggests, in the years following each consumer's initial enrollment period, the observed choices of existing enrollees become distorted by switching frictions. Therefore, I use the choices following the first enrollment period for the identification of switching costs. However, as a next step I divide the previously enrolled consumers into two groups based on how switching costs affect their decisions.

This approach allows me to estimate two types of switching costs and it follows the intuition of the argument in Luco (2017) who extends the classic switching cost identification strategy of Handel (2013). The first cohort I use is the group of continuous enrollees who were enrolled during all three years of the sample. These consumers can avoid both enrollment costs and decision costs by choosing the same plan as in the previous period. As a result, the choices of continuous enrollees are affected by both types of switching costs.

The next enrollee cohort I use for the identification of different switching costs is a special feature of the institutional setting: the group of returning enrollees who were enrolled initially and return to the Marketplace after a gap in their coverage. This phenomenon of frequent Marketplace drop-outs is well documented in the literature (see for example Diamond et al. (2019) and Gordon et al. (2019)). Marketplace enrollees can leave and return due to different reasons: they might receive and lose employer-sponsored coverage or they can decide to remain uninsured for a period. The only important fact from the perspective of my identification strategy is that returning enrollees have to go through the enrollment process again, regardless of the plan they choose. As a consequence, the chosen alternative is not affected by the enrollment cost. On the other hand, decision costs can be avoided for returning enrollees if they choose the same plan as they have selected at their previous choice point. Therefore, decision costs can be large enough such that this cohort chooses the same plan as before the drop-out. Hence, comparing the observed choices of continuous and returning enrollees can separately identify enrollment costs. A similar comparison between the choices of new enrollees and returning enrollees identifies decision costs.

The endogeneity of prices is a fundamental concern in the empirical IO literature. However, in this particular market, a set of institutional features help the identification of this parameter. First, as described in the earlier sections, the ACA's premium setting regulations generate variation in prices across different demographic cells. Second, insurers make entry decisions and adjust product characteristics at the rating area-level, therefore consumers living in the same region face the same set of plan-level unobservables that can be controlled for using fixed effects. Therefore, the combination of these features of the market environment create exogenous variation in premiums within rating areas, providing a convenient way to estimate demand using region fixed effects (Tebaldi, 2017).

In this model of health plan demand, individuals choose the alternative which results in the highest indirect utility, i.e. where $u_{ijt} \geq u_{ikt}$ for all $k \neq j$. Note that the observed choices reflect consumers' revealed preferences that might be affected by multiple choice frictions such as limited attention to less salient plan characteristics, overoptimism about future health risk or mis-weighting of probabilities. Therefore, the plans that yield the highest indirect utility in this model are not necessarily those that ex-post provide the highest protection against risk at the lowest costs, given the specific health care needs of the consumer (Chandra et al., 2019).

Finally, it should be also noted that this model of plan choice assumes myopic consumers. However, due to the instability of the market and the large political uncertainty around the future of the ACA, it is quite realistic to assume that the best prediction consumers can make about the future is their current information. Due to these specific features of the environment, it is reasonable to assume that a static framework describes better the decision making process of consumers in this market than a dynamic model.

1.5.2 Premium Setting Process

To close the model, I need to specify how insurance companies adjust prices to changes in risk allocation following different policy decisions. Modeling these interactions allows me to take into account the equilibrium effects of supply side responses in the counterfactual simulations.

There is an extensive literature studying supply side responses to consumer choice frictions (see Farrell and Klemperer (2007) and Grubb (2015) for reviews). The standard theoretical prediction is that inertia raises equilibrium prices, however this phenomenon often arises as a net effect of "investing" in high future market shares by charging low prices to new consumers and "harvesting" the benefits of sticky consumer choice combined with large market shares by setting high prices later on (Farrell and Klemperer, 2007). Clearly, in markets where firms cannot discriminate between incoming and existing consumers, applying the "invest and harvest" pricing strategy induces a trade-off in pricing decisions. In such cases the dominating effect depends heavily on the relative weight of the two groups and consumer expectations. However, even under uniform prices, usually the latter effect dominates, resulting in higher equilibrium prices compared to the frictionless case.

In the context of health insurance contract pricing, these theoretical predictions are confirmed by the empirical work of Ericson (2014) who shows that in the Medicare Part D prescription drug market, the demand side switching frictions lead to an "invest and harvest" pricing strategy, where prices are kept low in the early years of the market then raised rapidly later on.

As the main goal of this paper is to analyze how nudging policies affecting enrollment decisions change the severity of adverse selection, when modeling the supply side I focus on how insurers adjust premiums in response to policy changes affecting the risk composition of the enrollment pool. Therefore, I neither model strategic responses in terms of the set of contracts offered, nor endogeneize plan characteristics other than premiums.[19] Although modeling these supply side decisions would provide interesting insights about the equilibrium of effects of different policies affecting consumer choice, these aspects are beyond the scope of this current paper and provide an important area for future research. Moreover, the parameter estimates below reveal that premiums are one of the most salient features determining consumer choice, therefore the determinants of the price setting process are also directly related to changes in the risk composition of the enrollment pool. Hence, focusing on the evolution of prices is not likely to limit the main goals of the paper.

Having access to linked individual-level plan choice and medical claims data makes modeling the supply side convenient because in this case the marginal cost of insuring each enrollee is observed. In such circumstances, a commonly used approach in the empirical literature is estimating a policy function that provides an empirically stable relationship between between premiums and the risk composition of the enrollment pool, instead of explicitly modeling insurers' strategic pricing interactions (Handel, 2013; Polyakova, 2016*a*; Hackmann et al., 2015). Since in health insurance markets, pricing decisions are determined based on the past quality of risk pool, estimating the relationship between average plan costs and premiums is a standard procedure in the literature to close the model in the presence of claims data.[20]

[19]Note that the ACA imposed strong restrictions on the characteristics of the insurance plans sold on the Marketplaces. Therefore, this regulatory product standardization limits insurers' ability to endogeneize most product characteristics in their decisions.

[20]In adversely selected insurance markets, insurance companies cannot fully price on the risk characteristics of individuals either because of the presence of private information on risk status or due to rating restrictions, such as the community-rating provision of the ACA. Therefore, the equilibrium prices are determined by average costs instead of marginal costs (Einav et al., 2010*b*).

In selection markets, changes in enrollment directly affect the cost of insurers. Therefore, the price effects of the counterfactual policies can be pinned down by the new choice pattern that affects the input variables of the price setting model. In this paper, I follow this tractable approach to model how changes in the risk profile of the enrollment pool affect premiums.

The literature often augments the average cost based pricing rules with different plan or market characteristics. The usual approach is that the researcher selects a set of variables that are assumed to play a role in the premium setting process based on theory or different institutional features of the given market. In this paper, I contribute to this literature by employing machine learning techniques to select the most important features determining insurers' pricing decisions in order to improve the predictive power of the premium setting model. This approach allows me to estimate the counterfactual changes in premiums more accurately.

In particular, I use LASSO (Least Absolute Shrinkage and Selection Operator (Tibshirani, 1996)) to select the most relevant variables and to estimate the parameters of the contract pricing model. The estimator is given by

$$\hat{\beta}_{LASSO}(\lambda) = \underset{\beta}{\operatorname{argmin}} \left(\frac{1}{n} \sum_{j=1}^{J} (p_{jt} - X_{jt}\beta)^2 + \lambda \sum_{k=1}^{K} |\beta_{kt}| \right), \tag{1.18}$$

where p_{jt} denotes the premium of plan j and X is a vector including different plan characteristics, market shares, moments of the past claim expenditures and demographics of the risk pool. The term λ penalizes the sum of absolute values of the coefficients, with $\lambda = 0$ being equivalent to a standard linear regression. Because of the presence of this term, one of the most important properties of the LASSO estimator is that it shrinks the least squares estimators towards zero, with $\hat{\beta}_{LASSO} = 0$ for some j's. Due to this property LASSO is able to perform variable selection in regression models. I take advantage of this property of the estimator by allowing the algorithm to select the most important features that determine insurers' pricing decisions. In addition, it can be shown that LASSO performs better than least squares in terms of minimizing prediction error, especially as the dimensionality of the data increases (Bühlmann and van de Geer, 2011; Hastie et al., 2009). These properties of LASSO fit well with my goal of building an empirical pricing model to predict premiums accurately, given the policy changes affecting the demand side of the market.

1.5.3 Results

Table 1.8 shows the parameter estimates of the choice model, allowing for both observable and unobservable preference heterogeneity, as well as switching frictions. The parameter estimates have the expected signs: enrollees dislike premiums, deductibles and out-of-pocket payments, and they are more likely to choose plans with higher advertised actuarial values conditional on other characteristics. The estimates of the demographic interactions and random coefficients suggest a large heterogeneity in consumers' preferences for health insurance contracts. For instance, while older enrollees are less sensitive to deductibles and out-of-pocket maximums, they are more likely to choose metal tiers with more comprehensive coverage and PPO plans.

The coefficient estimates of the switching cost parameters are also significant, confirming the reduced form evidence for the presence of these choice frictions. As Table 1.9 shows, enrollment costs are higher than decision costs on average: \$1,473 and \$1,120, respectively. These results are consistent with the economically significant cost switching costs found by the health insurance literature in various settings.[21] The reduced form way of modeling switching costs implies that these estimates can be interpreted as consumers' willingness to pay to remain enrolled in their default plans.

The parameter estimates also show that enrollment costs show a larger variation across the population and they are decreasing in age. The latter result is not surprising as older consumers might care more about having insurance coverage due to their higher health risks. On the other hand, decision costs slightly increase with age, consistent with the literature on Medicare Part D that documented a wide range of empirical evidence suggesting that the choice persistence might be related to older people's higher cognitive costs of interpreting financial information related to health plans. The results also imply that the ratio of enrollment costs to decision costs is significantly higher among younger consumers. This pattern may be related to the fact that older people have more incentive to enroll due to their worse risk status, and

[21]For instance, in an employer-sponsored health insurance market, Handel (2013) estimated switching costs to be about \$2,000. Polyakova (2016*a*) found switching costs to be on the order of \$1,000 in the context of the Medicare Part D market. In Medicare Advantage, Nosal (2012) estimated switching costs to be about \$4,000.

the opportunity cost of the time-consuming administrative process of enrollment may also decrease with age.

Table 1.10 reports the estimation results of the contract pricing model. The input variables fed into the LASSO algorithm included different lags of own and competitors' market shares, lags of the plan enrollees' average claims, demographics and a rich set of plan characteristics. The penalty term λ was chosen by cross-validation to minimize the mean-squared prediction error. The parameter estimates reveal that premiums increase in the enrollment pool's past average spending, consistent with the literature using average cost-based pricing models (Einav et al., 2010*b*; Handel, 2013; Polyakova, 2016*a*; Hackmann et al., 2015). The table also shows that plans with older enrollees charge higher base prices. This result can be explained by the fact that the ACA does not allow insurers to fully underwrite age since the price setting rules maximize the ratio of premiums charged to the oldest and youngest enrollees. Indeed, Orsini and Tebaldi (2016) show that base premiums follow the age profile of the enrollment pool. Therefore, in rating areas with a higher share of older population, younger enrollees also face higher base prices. Premiums are also increasing in the size of the plan's provider network and the advertised actuarial value of the plan. Insurers with larger market shares charge lower premiums, possibly due to a combination of different factors, such as economies of scale, less severe adverse selection and larger insurers better ability to charge lower prices in order to gain market share as in Farrell and Klemperer (2007) or to negotiate lower prices with providers as in Ho and Lee

(2017). Finally, premiums are decreasing in out-of-pocket maximums, as expected.

Table 1.11 shows the observed market shares and premiums along with those predicted by the model. The table also reports some key moments of the observed and predicted age distribution that will be important to analyze changes in the risk profile of the enrollment pool in the counterfactual analysis. Overall, the choice and contract pricing models perform well in terms of replicating the observed outcomes. However, in order to account for simulation errors, I use the predicted outcomes as a comparison base for the counterfactual simulations presented in the next section.

1.6 Counterfactual Analysis

The theoretical analysis of the interaction of adverse selection and switching costs presented in Section 1.4 suggests that a policy that automatically enrolls the uninsured into ACA Marketplace plans would be beneficial if enrollment rates were low due to adverse selection. Reduced form evidence suggests that this condition holds in this market. However, these gains might be partially offset by the costs of implementing the policy. Therefore, in order to understand the potential benefits of automatic enrollment for health insurance market design, it is essential to predict and quantify the expected impacts of this nudging policy on outcomes the main outcomes targeted by policy-makers. Hence, in this section I simulate the auto-enrollment of the uninsured, relying on the structural parameter estimates of the empirical model.

I start the simulation algorithm by assigning a default plan to the uninsured. In order to understand how the choice of the default plan affects market outcomes, I create four variants of the auto-enrollment policy based on the default assignment algorithm: a random assignment of Bronze tier plans, a uniform assignment of a high deductible health plan (HDHP), a smart default policy with a personalized a plan assignment based on demographics, and the assignment of an individual-specific "best match" plan. The first two default plan designs are motivated by existing policy proposals and practices implemented in other health insurance markets. The personalized default counterfactuals are suggested by the demand estimates that revealed the importance of preference heterogeneity in consumer choice.

As a next step, I allow consumers to re-optimize their plan choice by comparing the utilities derived from each choice alternative. The resulting changes in enrollment generate a reallocation of market shares across insurers and contract types, and also affect the risk profile of the enrollment pool.

To predict the adjustment of premiums, I need to analyze how the new allocation of risk induced by the changes in enrollment affects insurers' costs, which is a crucial determinant of the pricing decision as shown above. To compute the change in the cost of the risk pool, I use two methods to measure the annual health care spending for each individual in the sample. First, for the currently enrolled population I use the individual-level claims data to calculate the actual realized annual spending. Second, while a fraction of the currently uninsured might not opt out from the assigned default plan in the counterfactuals, which is a main driver of the indirect effects, the yearly health care spending of this population is unobserved. Therefore, to predict the effect of the changes in the risk profile of enrollment pool on insurers' costs, I have to measure the expected yearly health care spending of the uninsured. To do so, first I project the observed spending of the currently insured on demographics, then I use these estimates to calculate the expected spending of the uninsured.

It should be noted that this strategy assumes that the health care spending of the insured and uninsured are similar conditional on observable demographics. This assumption results in conservative estimates for the change in average costs, and premiums. The reason is that this approach does not take into account the selection between insurance and uninsurance based on unobserved health risk. However, as the reduced form analysis above suggest, the uninsured might choose to forgo coverage due to their better risk profile. In this case, my strategy is likely to underestimate the decline in insurers' average costs, resulting in underestimates of the simulated premium reduction associated with the coverage expansion among healthier consumers. Therefore, these estimates should be interpreted as conservative lower bounds. Note that by assuming the similarity of the uninsured and insured population conditional on observable demographics, I can use data on the insured to choose a smart default for the uninsured. Finally, I assume away moral hazard by keeping the annual spending of individuals constant across different default plan assignments in the simulations.

1.6.1 Simulation Results

I simulate four variants of the auto-enrollment policy based on the default plan assigned to the uninsured. First, I use a randomly chosen Bronze plan as default, which is the most basic contract type on the market. Second, I assign a new catastrophic health plan to the uninsured and I determine the characteristics of this high deductible plan following the recommendations of existing policy proposals.[22] After simulating these naive policies that assign the same type of low-coverage plan to everyone, I turn to the analysis of a more sophisticated smart default policy. In this third counterfactual, I use the choice model to determine the individual-specific best match based on observables and I assign the uninsured to these plans by default. In the final default plan assignment, I repeat the smart default policy but now I find the "best-match" plan for each consumer based on both observables and the realized values of unobservables. Although the last simulation is not feasible to implement as it relies on private information unobservable for policy makers, it provides a useful exercise to learn about the upper bounds of the benefit of auto-enrolling the uninsured.

Table 1.12 reports the predicted outcomes for a set of key policy-relevant moments for each different default plan design. The simulation results predict that total enrollment increases for each potential policy, however there is substantial variation across the

[22]To construct this plan, I follow the recommendations of policy proposals that suggest the assignment of a catastrophic health plan with deductibles in the \$15,000-\$20,000 range and very low premiums. To check the robustness of the results to plan characteristics, I simulate multiple versions of this counterfactual where I change deductibles and premiums. The main results are not affected qualitatively by the choice of these plan characteristics as the actuarial value is kept constant across the simulations.

different types of default assignment algorithms. Column (2) shows that the random assignment generates the lowest change with a 32% growth in enrollment compared to the baseline case. The new HDHP default presented in column (3) performs better than the random assignment, with a 43% increase in the number of insured. The results in column (4) predict that the smart default policy raises enrollment by 69%. Column (5) shows that, as expected, the highest enrollment growth (93%) is generated by assigning the individual-specific "best-match" plan to auto-enrollees. The results imply that the last simulation represents the first best outcome of automatically enrolling the uninsured to ACA Marketplace plans, while the best feasible equilibrium in terms of the number of enrollees is the smart default policy.

The differences in predicted enrollment growth across the simulations suggest the importance considering the heterogeneity of preferences when designing default options in order to maximize the enrollment rates. In particular, the simulations predict that when the same type of low-coverage contracts are assigned to everyone, a higher fraction of auto-enrollees opt out to uninsurance than in the case of the two other sophisticated policies that take into account the rich individual-specific tastes for health insurance plans. The intuition underlying this result is the following: when an uninsured person is assigned to a plan which is not a good match to her, the expected benefit from switching back to uninsurance is more likely to exceed her switching cost. However, when the default plans are more aligned with individual-specific preferences, the utility difference between uninsurance and the given plan is less likely to exceed the amount of the switching cost, leading to lower opt-out rates.

Table 1.12 also reveals that the large inflow of the previously uninsured people affects the risk profile of enrollment pool. The decline in the average age and spending of enrollees indicates an improved extensive margin selection in each case. The table also shows that the change in the risk profile of enrollees in turn leads to a lower premium level on average.

However, looking at changes in the average premium level hides important variation in relative prices across the different counterfactual policies, as presented in Figure 1.9. The graph shows that although the two naive policies, which assign the same type of low-coverage plan to everyone, lead to a decline in the average premium level, this overall effect on the price level is driven by a large decrease in the Bronze tier premiums. However, the relative premiums of higher metal tier plans increase at the same time. Therefore, the changes in relative premiums suggest that while naive policies make low-coverage contracts more affordable, they increase the incremental premium of more comprehensive coverage as the price schedule becomes steeper.

To shed light on the mechanisms leading to these differential impacts on relative premiums, Figure 1.10 presents how the design of the default plan affects the distribution of enrollees across metal tiers. The simulations predict that naive policies raise enrollment overall, however this increase comes only from the Bronze tier. Moreover, the naive assignment of the lowest-coverage plans also generates a reallocation of consumers across metal tiers, leading to a decline in the market shares of more generous plans.

The key mechanism behind this result is the price adjustment process. In particular, the large number of new healthier enrollees in the Bronze tier, which is the metal tier of the default, decrease insurers' average costs for these plans. The contract pricing model implies that this change in the risk profile of enrollment pool leads to a decline in the premiums of Bronze plans. However, the resulting changes in relative prices make Bronze plans more attractive for consumers, even for those who would choose higher-coverage contracts under the baseline price vector. Therefore, while the naive default assignment algorithms expand coverage overall and reduce adverse selection on the extensive margin, they also indirectly create incremental sorting on the intensive margin by "stealing" healthier existing consumers from the higher metal tiers.[23]

Figure 1.10 also reveals, however, that smart default policies, which take the heterogeneity of preferences into account when assigning a default plan to the uninsured, spread out enrollment across contract types with different levels of coverage generosity. Therefore, smart default policies expand coverage among the healthier marginal consumers without indirectly generating adverse selection on the intensive margin. In addition, smart default policies lead to an equilibrium with better coverage for the enrolled since a fraction of new enrollees stays in their higher-coverage default plans. Moreover, it should be noted that analyzing a longer time horizon might reveal a larger divergence across the impacts of naive and personalized defaults due to the dynamic nature of adverse selection.

[23]These results are consistent with the theoretical predictions of Azevedo and Gottlieb (2017) about the implementation of a mandate.

These simulation results show that personalizing the default plan assignment mechanism using only simple demographic information, such as age and gender, can increase enrollment rates by about 27 percentage points compared to naive algorithms. Since coverage rate is an objective targeted by policy makers, the results of this paper suggest that it is important to invest in designing personalized smart defaults in order to maximize the benefits of automatic enrollment policies in health insurance markets. As these types of default plan designs rely heavily on understanding individual-specific preferences, the findings of this paper also highlight the importance of building and maintaining high-quality data infrastructures to support health care policy making and health insurance market design.

1.7 Conclusion

Early data on the ACA Marketplaces suggested that the enforcement of the individual mandate was not strict enough to achieve high health insurance coverage in the younger population. However, enrolling low-risk individuals is essential to protect these new health insurance markets from the harmful consequences of adverse selection in the long run.

In this paper, I analyze an alternative, behavioral approach to expand health insurance coverage in the ACA Marketplaces that uses the exact same mechanism that currently prevents consumers from enrolling under the market design with active sign-up: the persistence of consumer choice. I show that a nudging policy that automatically enrolls the uninsured into ACA Marketplace plans would reverse the direction of the cognitive costs associated with the enrollment process into the right direction, to increase enrollment rates. Moreover, due to the tight link between demand and costs in selection markets, enrolling the healthier marginal consumers generates a positive externality on the premium level.

Using a theoretical model, I show that this policy might be beneficial if coverage rates are inefficiently low, and its benefit depends on the strength of consumer inertia. I also document a set of new reduced form evidence consistent with the presence of adverse selection and switching frictions in this market. To quantify the potential impacts of the counterfactual auto-enrollment policy, I run simulations using the structural parameter estimates of an empirical model of the insurance market. I find that this nudging policy would lead to a substantial expansion in coverage among younger people and reduce the premium level.

The results of the paper also highlight the importance of the optimal default plan design. Naive policies, which assign the same type of low-coverage plan by default to the uninsured, can have dangerous indirect effects on adverse selection conditional on enrollment. However, I also show that this trade-off between extensive and intensive margin selection can be avoided by personalizing the default option. Therefore, the results suggest that designing personalized smart defaults is essential for maximizing the benefits of auto-enrollment policies in health insurance markets.

In general, I show that simple behavioral policies have the potential to increase enrollment rates and maintain the stability of the private individual health insurance market in the long run. However, I also find that implementing auto-enrollment policies in selection markets can result in unexpected indirect effects due to the price adjustment mechanism. These results provide important new insights for healthcare policy design, especially given the recent repeal of the most important stabilizing tool of the ACA, the individual mandate.

Figures and Tables

Figure 1.1: Enrollment and Uninsurance by Year

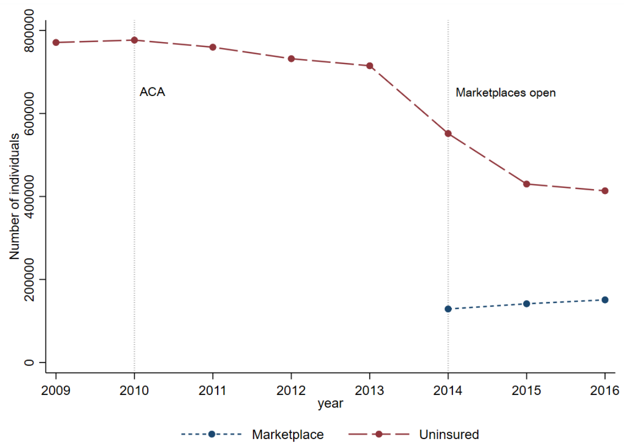

Notes: The figure displays the number of uninsured and Marketplace enrollees by year. Data on the number of uninsured was obtained from the American Community Survey. The number of ACA Marketplace enrollees was released by Connect for Health Colorado. Colorado also expanded the Medicaid program as of January 2014.

Figure 1.2: ACA Geographic Rating Areas in Colorado

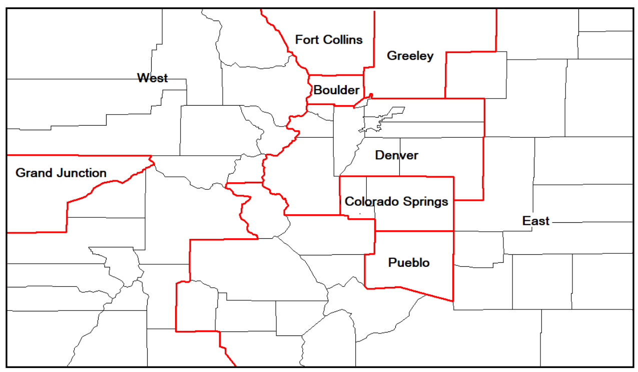

Notes: The map displays the current borders of the ACA geographic rating areas in the state. In Colorado, the ACA rating areas are defined as sets of counties. Originally, in 2014, there were 11 rating areas, however due to very high premiums in some regions, regulators decided to merge some of them from 2015. The current East rating area was formed by merging Southeast and Northeast regions. The current West region was also divided into two areas initially: West and Resort.

Figure 1.3: Demand Side Variation Across Rating Areas

(a) Number of Enrollees

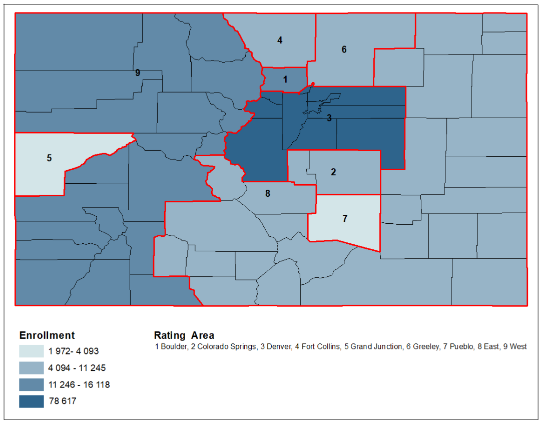

(b) Average Age of the Enrolled

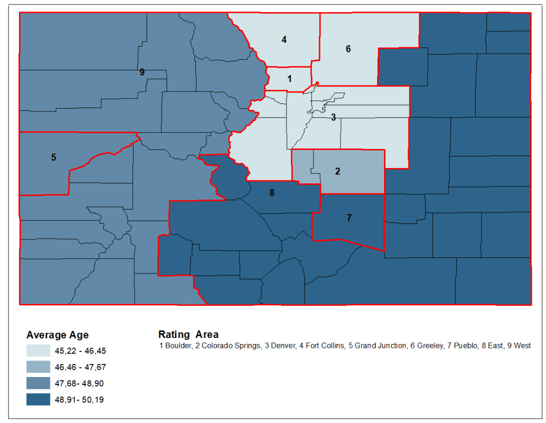

Notes: Panel 1.3a shows the distribution of Marketplace enrollment across rating areas in Colorado. Panel 1.3b shows the average age of the Marketplace enrollees by rating area. Sample: Colorado APCD, 2015 enrollment data.

Figure 1.4: Supply Side Variation Across Rating Areas

(a) Number of Insurers

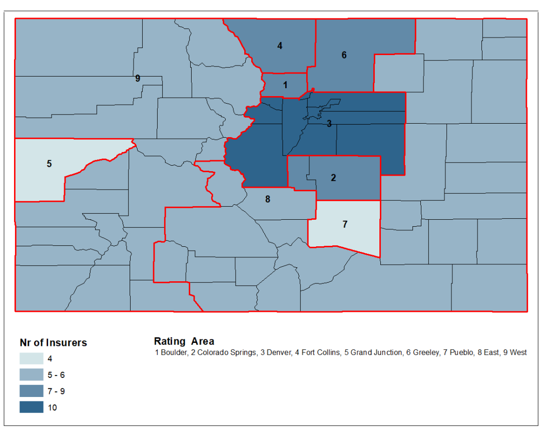

(b) Average Premium

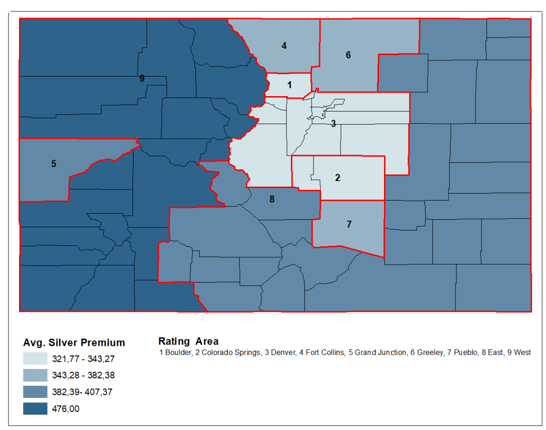

Notes: Panel 1.4a shows the number of insurance companies offering at least one plan in any metal tier in the rating area. Panel 1.4b displays the variation in the average monthly premium of Silver plans across rating areas. The figures were generated based on 2015 data.

Figure 1.5: Enrollment and Uninsurance by Age

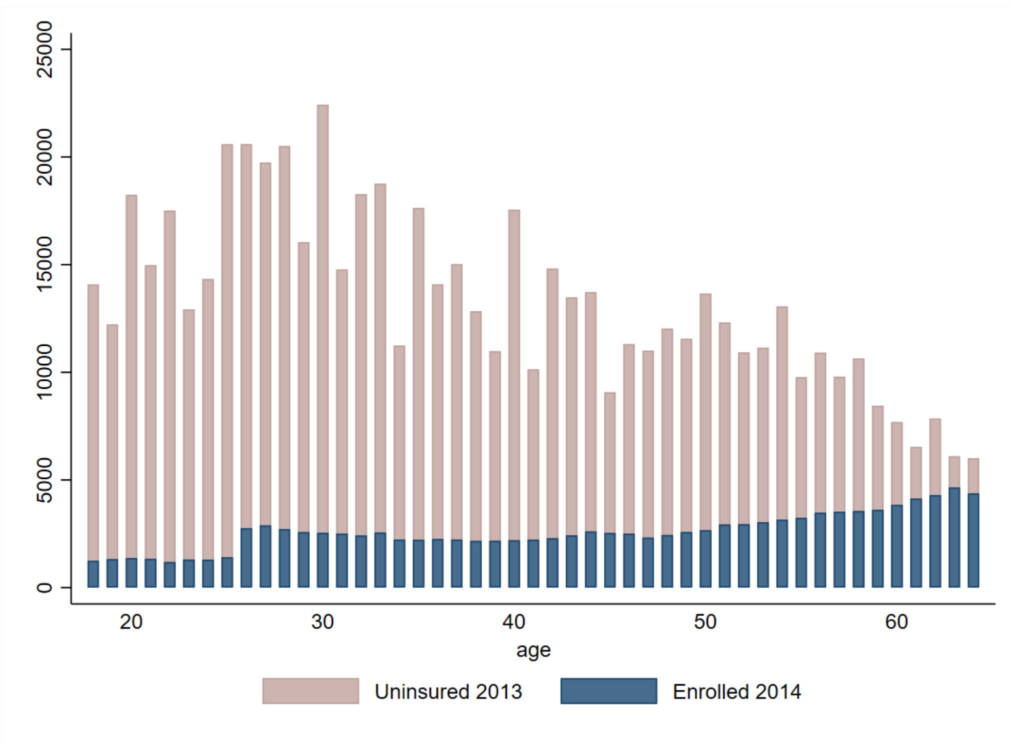

Notes: The figure displays the number of ACA Marketplace enrollees and uninsured by age in Colorado. Data on the number of uninsured was obtained from the American Community Survey. The enrollment data reflects information in the eligibility file of the Colorado APCD.

Figure 1.6: No Switching Frictions

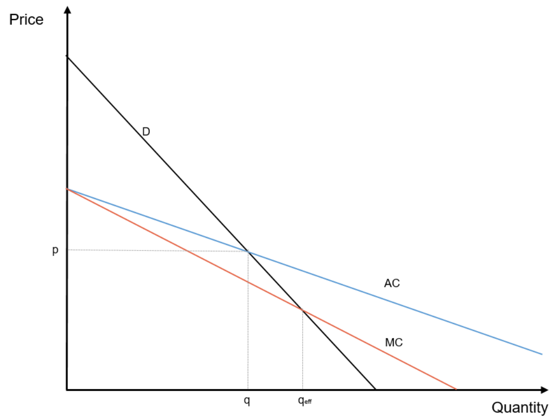

Notes: The figure shows an adversely selected health insurance market in the absence of switching costs, based on Einav et al. (2010*b*). The downward sloping MC curves implies that consumers with higher willingness to pay are also riskier. Due to asymmetric information or rating restrictions, there is under-insurance (q) compared to the efficient rate (q_{eff}).

Figure 1.7: Opt-in

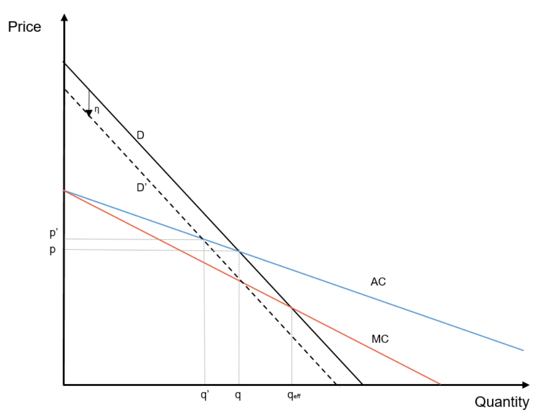

Figure 1.8: Opt-out

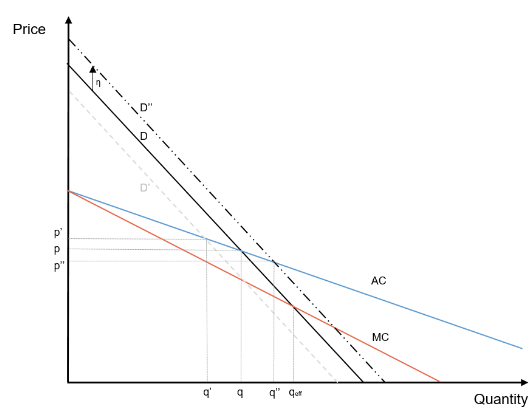

Notes: Figure 1.7 shows that under active enrollment, adding switching costs to the frictionless Einav et al. (2010*b*) framework makes the outside option more attractive. Figure 1.8 shows that the auto-enrollment policy increases enrollment from q' to q'', by making the outside option less attractive.

Figure 1.9: Simulated Premiums under Different Default Plan Designs

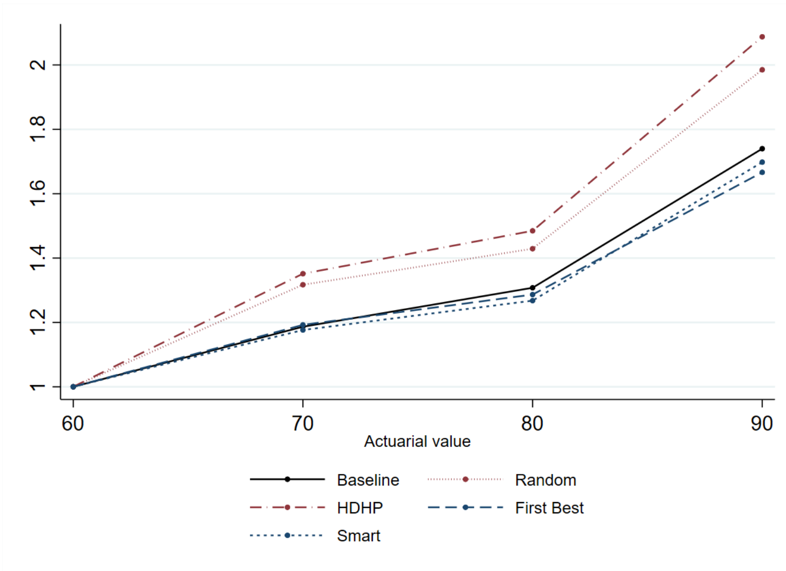

Notes: The figure displays the simulated relative prices by metal tier in the baseline case and the different counterfactuals. The horizontal axis represents the advertised actuarial value of the metal tiers. Bronze, Silver, Gold and Platinum plans correspond to 60%, 70%, 80% and 90%, respectively. Premiums are expressed relative to Bronze tier prices. The steeper price schedule reflects the incremental intensive margin adverse selection in case of counterfactual policies using naive default assignment rules (Random assignment, HDHP).

Figure 1.10: Simulated Enrollment under Different Default Plan Designs

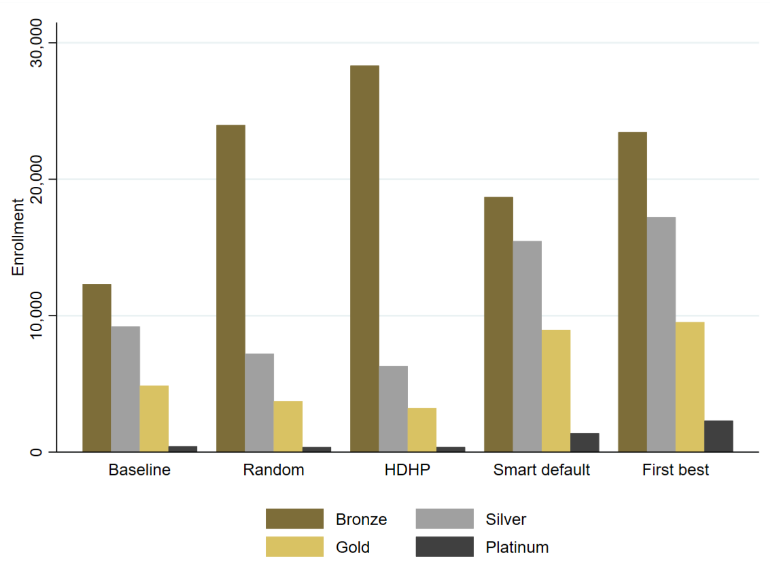

Notes: The figure displays the simulated enrollment by metal tier in the baseline case and the different counterfactuals. The design of the default plan has a large effect on the distribution of enrollment across metal tiers. Policies that assign the same type of plan to everyone (random assignment, HDHP) increase enrollment rates overall but reduce enrollment in the upper metal categories, generating incremental adverse selection conditional on enrollment. However, personalized default policies (smart, first best) expand coverage by more overall and spread out enrollment across metal tiers.

Table 1.1: Insurer Participation and Market Shares

	2014	**2015**	**2016**
Panel A - Metal tier			
Bronze	35.21%	38.72%	42.29%
Silver	45.54 %	41.69%	41.42%
Gold	15.47%	15.28%	13.01%
Platinum	3.78%	4.31%	3.28%
Panel B - Insurer			
Kaiser Permanente	57.60%	45.38%	40.08%
HMO Colorado /Anthem	12.61 %	24.60%	39.10%
Humana Health Plan	7.74 %	8.02%	4.82%
CIGNA	7.31 %	6.63%	6.26%
Rocky Mountain Health Plan	7.11%	5.57%	1.76%
Time	6.38%	4.47%	–
Humana Insurance Company	1.12%	1.15%	0.33%
Colorado Access	0.07%	0.19%	–
Colorado Choice	0.05%	2.84%	5.08%
Denver Health	0.01%	0.21%	0.12%
Freedom Life	–	0.04%	0.06%
United Healthcare	–	0.92 %	3.43%

Notes: The table displays state-level market shares by metal tier and insurer. Market shares are calculated from the CO APCD eligibility file. ACA Marketplace plans are categorized into metal tiers based on actuarial value. Bronze, Silver, Gold and Platinum plans pay about 60%, 70%, 80% and 90% of the insured's health care expenditures, respectively. Insurers' entry decisions take place at the rating area-level and most local markets are even more concentrated. In 2016, Time and Colorado Access exited the ACA Marketplace in the entire state and Rocky Mountain Health Plan limited its participation to only two rating areas, generating substantial changes in the choice set.

Table 1.2: **Change in the Health Status of the Uninsured, 2013-2015**

	overall		**by age**					
			19-34		**35-54**		**55-64**	
	2013	**2015**	**2013**	**2015**	**2013**	**2015**	**2013**	**2015**
Fair/ poor health	24%	19%	18%	20%	33%	16%	35%	33%
Provider visit	51%	45%	42%	37%	53%	47%	60%	50%

Notes: The table displays changes in the health of the uninsured in Colorado over time, before and after the Marketplace opened in 2014. Fair/ poor health is an indicator for a self reported health status. Provider visit measures the % of the uninsured who visited a health care professional in the past 12 months. Sample: Colorado Health Access Survey.

Table 1.3: **Health Status of the Uninsured, Pre- and Post-ACA**

	(1)	(2)	(3)
	poor health	**poor health**	**poor health**
uninsured	0.273***	0.187***	0.187***
	(0.0003)	(0.0004)	(0.0005)
after	0.104***	0.048***	0.160**
	(0.0005)	(0.0009)	(0.0856)
uninsured x after	-0.182***	-0.083***	-0.195**
	(0.0007)	(0.0012)	(0.0857)
uninsured x after x age 35-54		-0.059***	0.099
		(0.0019)	(0.1122)
uninsured x after x age 55-64		-0.002	0.147
		(0.0020)	(0.1087)
N	4,153,573	4,153,573	2,508,848

Notes: The tables displays the parameter estimates of the difference in differences model. The estimates show how the self reported health status of the uninsured population changed compared to the individual market enrollees after the ACA's health insurance reforms in 2014. The dependent variable is an indicator for a self reported poor health. Column (1) reports the results for the baseline specification. The interaction terms in column (2) break down the change in health status by age group. The last specification estimates the second model on a propensity score matched sample based on employment category, education, family size and income. Sample: Colorado Health Access Survey. Standard errors in parenthesis; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 1.4: **Adverse Selection Dynamics**

	(1)	(2)	(3)
	$\mathbb{1}$(**Exiting consumer**)		
Cost	-0.009***	-0.015***	-0.012***
	(0.0018)	(0.0037)	(0.0038)
Premium increase		0.130***	0.139***
		(0.0153)	(0.0158)
Premium increase x Cost		-0.004***	-0.006***
		(0.0014)	(0.0023)
N	87,627	87,627	84,884

Notes: The table displays the estimates from a logit model of consumer exit decision on cost, premium change and a set of region and time fixed effects. The dependent variable is an indicator for a previously enrolled person dropping out of the Marketplace in the current year. Cost measures the annual health care spending in the previous year. Premium increase measures the change in the price of the default plan. The model also controls for rating area and year fixed effects In column (3) individuals switching to employer-sponsored insurance plans are excluded from this analysis. Sample: Colorado APCD. Standard errors in parenthesis; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 1.5: Switching Patterns

	2014		2015		2016	
	Enrollment	**Share**	**Enrollment**	**Share**	**Enrollment**	**Share**
Total enrollment	125227	100%	145,709	100%	150,691	100%
New enrollees	125,227	100%	48,321	33.2%	73,034	48.5%
Re-enrollees			97,388	66.8%	77,657	51.5%
Same plan			86,651	89.0%	60,844	78.3%
Switch			10,737	11.0%	16,813	21.7%

Notes: The table displays the number of enrollees by switching status. The ACA Marketplaces opened in 2014, therefore in the first year every individual is a new enrollee by definition. The large increase in the number of switchers in 2016 is generated by the exit of insurers.

Table 1.6: **Market Shares by Enrollment Cohort and Year**

	Initial enrollment 2014			*Initial enrollment 2015*			*Initial enrollment 2016*		
	2014	2015	2016	2014	2015	2016	2014	2015	2016
Panel A - Metal tier									
Bronze	35%	34%	35%	-	45%	48%	-	-	44%
Silver	46%	46%	46%	-	36%	35%	-	-	41%
Gold	16%	15%	16%	-	15%	15%	-	-	11%
Platinum	3%	4%	4%	-	4%	3%	-	-	3%
Panel B - Insurer									
Kaiser	58%	59%	63%	-	28%	29%	-	-	34%
Rocky M.	7%	6%	4%	-	3%	2%	-	-	1%
Anthem	13%	12%	18%	-	41%	48%	-	-	42%
Humana HP	8%	8%	5%	-	8%	4%	-	-	3%
Cigna	7%	8%	2%	-	6%	1%	-	-	9%
Other	6%	7%	8%	-	13%	16%	-	-	11%

Notes: The top panel shows the market shares of the different metal tiers among different enrollment cohorts based on the year of the first enrollment. The bottom panel displays the choices of the same cohorts by insurer. Enrollees who were forced to switch their plans due to their insurer's exit decision were excluded from the analysis. The table reveals that the choice of each cohort reflects the conditions of the initial enrollment period. Due to the substantial changes in the choice set over time, this pattern suggests a strong inertia in consumer choice.

Table 1.7: The Determinants of Switching Status

	(1)	(2)	(3)
	$\mathbb{1}$[Switch]		
Female	0.098***	0.104***	0.083***
	(0.0079)	(0.0080)	(0.0141)
Age	0.003***	0.002***	0.001*
	(0.0003)	(0.0003)	(0.0005)
Income	0.004**	0.014***	0.009***
	(0.0018)	(0.0019)	(0.0033)
Number of options		-0.018***	-0.029***
		(0.0009)	(0.0017)
Number of exits		0.023***	0.045***
		(0.0018)	(0.0033)
Health shock			0.129***
			(0.0161)
Default premium increase			0.337***
			(0.0155)
Constant	-1.228***	-0.931***	-1.158***
	(0.0189)	(0.0369)	(0.0672)
N	150 467	150 467	48 565
Pseudo R-squared	0.0021	0.0246	0.0313

Notes: The table presents the parameter estimates of a logistic regression of individual's switching status on demographics and rating area-level choice set characteristics. Income is measured at the zip code-level. Health shock is an indicator for at least $1,000 change in yearly health care spending. The number of options measures the number of plans offered in the individual's rating area. Standard errors in parenthesis; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 1.8: Demand Estimates

	Estimate	**S.E.**
Plan characteristics (β)		
Premium	-1.4371***	(0.0149)
Deductible, μ	-0.3332***	(0.0180)
x Age	0.0048***	(0.0004)
x Female	0.0194**	(0.0102)
σ	0.0016	(0.0094)
OOP max, μ	-1.9280***	(0.0756)
x Age	0.0238***	(0.0011)
x Female	0.2007***	(0.0209)
σ	0.0898***	(0.0033)
Advertised AV	0.0261***	(0.0022)
x Age	0.0003***	(0.0000)
x Female	0.0064***	(0.0011)
PPO	0.1038***	(0.0604)
x Age	0.0036***	(0.0012)
x Female	0.0361	(0.0331)
Switching costs (η)		
η^{enr}	3.5915***	(0.5545)
x Age	-0.0292***	(0.0102)
x Female	-0.3149	(0.2448)
η^{dec}	1.6100***	(0.5329)
x Age	0.0068	(0.0078)
x Female	-0.5303***	(0.1321)

Notes: The table reports the estimated parameters of the demand model. The dependent variable takes value 1 for the observed choice and 0 for all other alternatives. The model is estimated via maximum simulated likelihood. The top panel shows the parameter estimates of different plan characteristics, the bottom panel displays the switching cost parameters. Column (1) reports the parameter estimates, column (2) shows the standard errors. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 1.9: Switching Cost Estimates

	Enrollment Cost ($)	Decision Cost ($)
Mean	1 473	1 120
Standard deviation	300	195
Median	1 448	1041
p25	1 225	951
p75	1 712	1 320

Notes: The table presents the estimated enrollment and decision cost implied by the demand parameters. These costs represent the willingness to pay to remain enrolled in the default plan.

Table 1.10: Pricing Equation

Selected feature	Estimate
Lagged spending	0.018
Lagged age	7.922
Lagged market share	-1.203
OOP max	-0.028
Network size	0.002
Gold	-52.398
Silver	-89.510
Bronze	-136.078

Notes: The table shows the selected variables and their coefficient estimates. Variable selection is performed via LASSO. The dependent variable is the monthly plan-level base premium. Lagged spending refers to the first lag of plan enrollees' average claims. Network size is a sum of the number of providers and facilities in the plan's network. The pricing equation is estimated separately for each rating area. The optimal value of the LASSO penalty term is selected by cross-validation.

Table 1.11: Fit of the Model

	Observed	**Predicted**
Market shares		
Uninsured	77.86%	77.66%
Bronze	9.17%	10.71%
Silver	9.73%	7.26%
Gold	3.05%	4.21%
Platinum	0.18%	0.17%
Premiums		
Mean monthly premium	393	409
Risk profile		
Mean age - enrolled	45.41	46.29
Mean age - uninsured	36.17	35.39

Notes: The table displays the observed and simulated market shares, monthly premiums and average age by insurance status. Outcomes are reported for 2015. Simulations are based on the estimated demand and supply parameters, assuming no change in the policy environment.

Table 1.12: **Counterfactual Simulations with Different Default Designs**

	(1) Baseline	(2) Random Bronze	(3) New HDHP	(4) Smart default	(5) First best
Number of enrollees	26,799	35,298	38,234	45,422	51,584
Mean age	46.3	43.2	43.1	42.1	40.7
Mean spending ($)	3,416	3,120	3,073	2,974	2,753
Mean monthly premium ($)	409	384	379	374	362

Notes: The table shows a set of key moments under the current regulation and the different variants of the auto-enrollment policy. Column (2) refers to a policy that assigns a random existing Bronze plan to auto-enrolled individuals. Column (3) shows the predictions in case of a high deductible health plan set as the default. Column (4) is a smart default policy that assigns the best choice to everyone based on observable demographics. Column (5) corresponds to the assignment of a best match plan based on both observables and unobservables.

Chapter 2

Moral Hazard in Health Care Utilization and the ACA's Cost Sharing Reduction Subsidies

2.1 Introduction

Health care has been a central question of policy debates in the United States during the last decade. Much of this large political and media attention was generated by two key features of the US health care system that make it distinctive among developed countries. First, spending on health care has been growing rapidly from 5% of the GDP in the 1960s to 17% in 2017 (OECD, 2017). As a result, today the US has the highest level of health care spending in the world, expressed both as a share of the GDP and in per capita terms.[1] Second, despite having the most expensive health

[1]The OECD average of health expenditure as a share of GDP was 9% in 2017. Per capita health spending was $10,585 in the US, while the OCD average was $3,992 in 2017 (OECD, 2017).

care system in the world, the number of Americans with no health insurance coverage exceeded 46 Million by 2010 (Kaiser Family Foundation, 2018*a*).

As a consequence of these striking trends, another major health policy concern has developed in recent years: the rapid spread of high deductible health plans (HDHPs). According to recent data from the National Health Interview Survey, the share of privately insured consumers enrolled in health insurance plans with annual deductibles exceeding $1,000 reached 47 percent in 2018 (NCHS, 2018).[2] Although a common argument in favor of HDHPs is that these products help consumers get access to health insurance coverage by allowing insurers to provide low premium plans, the value of this catastrophic coverage is a non-trivial question since individuals have to pay the full marginal cost of care out-of-pocket until their expenditures meet the deductible level.

Therefore, it has become a significant health policy concern how the rising enrollment in these types of health insurance products affect health outcomes in the long run. However, to answer this question, the first step is to understand how consumer demand for health care responds to the price of medical care.

The 2010 Affordable Care Act (ACA), also known as Obamacare, aimed both to expand health insurance coverage and reduce the growth of health care spending in the US. Therefore, one of the law's major provisions was the introduction of the

[2] In the 2018 National Health Interview Survey, HDHP was defined as a private health plan with an annual deductible of at least $1,350 for individual coverage or $2,700 for families (NCHS, 2018).

health insurance marketplaces as standardized platforms for individuals to purchase subsidized health insurance coverage. Although insurance coverage reduces the price of medical care, with the rising prevalence of high deductible health plans, using health care services can still result in a large financial burden for the insured population. Therefore, the ACA introduced cost sharing reduction (CSR) subsidies for low income marketplace enrollees to decrease the share of health care expenditures paid by consumers.

In particular, the CSR subsidies induce discontinuous increases in the generosity of health insurance plans in the silver tier when the enrollee's income falls below the subsidy eligibility threshold.[3]

This specific design of the CSR subsidies provides a unique setting to study the impact of increasing insurance coverage generosity on the demand for medical care. Insurance induces an efficiency tradeoff between providing risk protection and changing utilization incentives due to the decline in the marginal cost of medical care – the latter is also known as moral hazard in the context of health insurance (Arrow, 1963; Pauly, 1968).

Identifying moral hazard from observational data is a major challenge in the literature due to the endogeneity of health plan choice. In other words, a simple comparison of

[3] In the ACA marketplaces, health insurance plans are categorized into metallic tiers based on actuarial value (i.e. the share of expenses paid by the insurance company on expectation). The advertised actuarial value of silver plans is defined at 70%; with CSR subsidies this can increase to 73%, 87% or 94%, depending on income.

average spending across plans with different levels of coverage might be misleading because people who choose low-coverage plans are systematically different from those who purchase more generous contracts. Moreover, these differences might be correlated with expected medical spending, causing higher risk types to demand more insurance (Rothschild and Stiglitz, 1976).

Therefore, in order to measure the behavioral response of consumers to changes in cost sharing – conditional on health plan choice – it is important to separate the effect of moral hazard from adverse selection. Clearly, the ideal exercise for this purpose is a random assignment of consumers to health insurance plans with different cost sharing features. However, implementing such large scale randomized control trials is extremely rare in this context. As a result, the RAND health insurance experiment conducted in the 1970s still provides a gold standard for the literature on moral hazard in health insurance (Manning et al., 1987).

In this paper, I take advantage of the discontinuities in the generosity of silver plans – generated by the ACA's cost sharing subsidies – to estimate behavioral responses to changes in the price of medical care. This subsidy design generates variation in the cost sharing features faced by the enrollees of otherwise identical plans, determined only by income. Therefore, the CSR subsidies provide a unique setting to separate moral hazard from adverse selection.

These specific properties of the institutional setting allow me to exploit a sharp regression discontinuity design to estimate the behavioral response to changes in out-of-pocket prices. However, data availability poses a significant barrier to implement this clean identification strategy. In particular, linked individual-level data on income, CSR enrollment and health care spending are typically not available for researchers. In this paper, I overcome this empirical challenge by combining data from different sources with industry regulations. Specifically, I use data on post-subsidy premiums from the Massachusetts All Payer Claims Database and information on plan-level base prices from insurers' rate filings to calculate the amount of the ACA's premium subsidy. Then I apply the ACA's premium subsidy formula backwards, which determines the amount of the premium subsidy as a function of income, in order to obtain an estimate for income.

The regression discontinuity estimates relying on this imputed income reveal significant increases in the health care utilization of silver plan enrollees as income falls below the CSR eligibility cutoffs. For instance, individuals who earn just *below* the CSR eligibility threshold at 250% of the FPL, demand $650 more health care services annually than enrollees of the same plan but whose income falls just *above* the subsidy eligibility cutoff point. Intuitively, the design of the CSR subsidies imply that consumers on the two sides of the cutoffs face different out-of-pocket costs for the same set of medical services, affecting their demand for health care. These results confirm the previous findings of the literature that the demand for health care is elastic, which is typically interpreted as an evidemce of moral hazard in health insurance.

The results of this paper also provide new estimates for moral hazard in health care utilization among the under-studied previously uninsured, low income population – many of whom gained access to insurance coverage for the first time due to the ACA's coverage expansion.

The rest of the paper proceeds as follows. Section 2.2 summarizes the related literature on the concept of moral hazard in health insurance and estimating the price elasticity of demand for health care. Section 2.3 describes the most relevant features of the institutional environment, focusing on the health insurance marketplaces and the ACA's subsidy mechanisms. Section 2.4 presents the data sources and some key descriptive statistics of the sample. Section 2.5 describes the empirical strategy and Section 2.6 presents the results and section 2.7 provides several robustness and placebo checks. Section 2.8 concludes.

2.2 Literature Review

This paper contributes to two major areas of the literature. First, it is related to the large literature studying moral hazard in health insurance and the price elasticity of health care. The concept of moral hazard in health insurance has been recognized by the theoretical literature for a long time (Arrow, 1963; Pauly, 1968).

However, showing whether the demand for health care is indeed price sensitive, and quantifying its extent still provides a major challenge for the empirical literature. These difficulties in identifying moral hazard stem from the fact that it is hard to separate it from adverse selection using observational data since both concepts produce a positive correlation between insurance coverage and health care utilization.

Therefore, the famous RAND health insurance experiment implemented an ideal research design to address this endogeneity concern by randomly assigning individuals to insurance plans with different levels of cost sharing. Studies analyzing these experimental data found significant reductions in health care utilization associated with higher levels of enrollee contributions (Manning et al., 1987; Aron-Dine and Finkelstein, 2013). These results were confirmed later by the Oregon health insurance experiment that studied a randomized expansion of the Medicaid public insurance program in the state of Oregon due to budgetary constraints (Finkelstein et al., 2012).

However, this ideal randomized research design is extremely rare, therefore several papers have attempted to use quasi-experimental variation in different settings to study how insurance coverage affects the demand for medical care (Einav et al., 2017; Duggan et al., 2008). One strand of this literature uses a regression discontinuity design, similar to the empirical model applied in this paper, to identify moral hazard (Polyakova, 2016*b*; Almond and Doyle, 2011). This paper contributes to this literature by taking advantage of the unique institutional setting provided by the ACA's cost sharing subsidies to analyze this classic question in a new context.

This paper also contributes to the growing literature studying different aspects of health insurance exchanges. Many of these papers focus on the Massachusetts health insurance exchange that served as a model for the design of the ACA Marketplaces

(Hackmann et al., 2015; Ericson and Starc, 2015; Jaffe and Shepard, 2016). As individual-level data becomes available of the ACA marketplaces, more and more papers analyze different aspects of this new market (Tebaldi, 2017; Orsini and Tebaldi, 2016; Panhans, 2018; Diamond et al., 2019). Most of the literature that studies the ACA subsidy mechanisms focus on the premium subsidies, while the impacts of the cost sharing reductions are much less understood. The most related studies also focus on the CSR subsidies (DeLeire et al., 2017; Ericson and Sydnor, 2018; DeLeire et al., 2018), however the current paper contributes to this area of research by exploiting an identification strategy that allows to address the potential endogeneity concerns of previous research.

2.3 Institutional Setting

The 2010 Affordable Care Act had three major aims in order to improve the health care system of the United States: expanding access to health insurance coverage, helping people receive better quality care and reducing the growth of health care spending. To accomplish the first goal and make private health insurance markets more accessible for individuals, the ACA established health insurance marketplaces as standardized platforms to purchase publicly subsidized private health insurance plans. Moreover, the ACA's rating regulations do not allow insurance companies to fully underwrite health risk anymore; instead they have to offer the same premiums to sick and healthy individuals within a given demographic cell (by age, rating area and smoking status).

2.3.1 Health Insurance Marketplaces

The marketplaces are online portals offering a wide range of health insurance plans that are provided by private health insurance companies, however the benefits are heavily regulated by the government. States have the option to use either the federally operated platform (available at HealthCare.gov), or to run their own state-based Marketplaces. Since 2014, twelve states (CA, CO, CT, DC, ID, MA, MD, MN, NY, RI, VT, WA) established their own state-based marketplaces (Figure 2.1).[4]

The major regulatory requirements regarding the marketplaces are standardized across states. Enrollees can purchase health insurance coverage during the open enrollment period each year and the health insurance plans sold on these platforms must satisfy the minimum essential benefits determined by the ACA.[5] The ACA also requires insurers to use a set of the pre-determined geographic rating areas as local markets in their contract pricing and entry decisions. These regions are determined by the Centers for Medicare and Medicaid Services, and are usually defined either as a set of counties or Metropolitan Statistical Areas (CMS, 2018). In addition, the health

[4] Five other states (AR, KY, NM, NV, OR) run their state-based exchanges through the federal platform.

[5] Under some special circumstances (certain life events, such as marriage, giving birth or losing other source of coverage), individuals might qualify to buy coverage through the marketplaces during special enrollment periods.

insurance plans offered on the marketplaces must be categorized into one of four metal tiers based on the advertised actuarial value. Bronze, silver, gold and platinum plans cover 60%, 70%, 80% and 90% of the insured's health care expenditures on expectation, respectively.

Massachusetts, the state studied in this paper, runs a state-based marketplace called the Massachusetts Health Connector. The Connector was the first health insurance exchange of the country, established as part of the 2006 Massachusetts health care reform, prior to the ACA. Due to its stability and success in expanding health insurance coverage, the Connector was considered as model for the ACA marketplaces (Gasteier et al., 2018; Ruggles et al., 2017). In Massachusetts, the premium setting process deviates from the rules used in most states in two ways. First, although the federal age rating ratio allows insurers to charge at most three times higher prices for older enrollees than to younger individuals, Massachusetts uses its own age curve with a rating ratio maximized at 2:1. Second, the state does not allow insurance companies to price on smoking status (CMS, 2017). Moreover, the state's individual mandate remains in effect although the federal government repealed the mandate tax penalty starting from 2019 (Massachusetts Health Connector, 2019). Massachusetts is divided into seven ACA geographic rating areas, defined as sets of 3-digit zip codes (Figure 2.2).

Massachusetts follows the active purchaser health insurance exchange model, meaning that the Health Connector directly negotiates a wide range of conditions (premiums, benefits, networks, number of plans, etc.) with insurers (Gasteier et al., 2018; Norris, 2018*a*).[6] Due to its active role in the negotiating process and longer history, the Health Connector is one of the most stable exchanges of the country. The strict standardization of the plan benefits provide an ideal setting for my analysis since it allows me to study the impacts of changes in plan generosity on consumer behavior within a set of otherwise homogenous contracts.

2.3.2 The ACA's Subsidy Mechanisms

Although the ACA's main goal was to help people get access to health insurance coverage by eliminating previous rating practices that used information on pre-existing medical conditions, keeping health insurance affordable was also an essential part of the reform. Therefore, to reduce the financial burden of health insurance for low income people, the ACA introduced two different subsidy mechanisms available for individuals purchasing plans from the marketplaces. This paper focuses on the impacts of the ACA's cost sharing reduction subsidies, however the empirical strategy relies heavily on the design of the premium subsidies. Therefore, this section describes both mechanisms.

[6] The other health insurance exchange model is where the Marketplaces serve as clearinghouses and accept all plans that meet their criteria(Krinn et al., 2015).

2.3.2.1 Premium Subsidies

Premium subsidies – or Advanced Premium Tax Credits (APTCs) – provide financial assistance for purchasing health insurance coverage for low income individuals up to 400% of the federal poverty level (FPL).[7][8] The minimum eligibility threshold is 100% of the FPL. Premium subsidy eligibility also requires purchasing insurance through the marketplaces. Individuals who have access to affordable employer-sponsored coverage or are eligible to different public insurance programs (Medicaid, Medicare) do not qualify for this financial assistance.

Premium subsidies are designed to reflect both enrollee income and variations in the price level across different geographic areas. The key idea behind the design of this subsidy is that it determines a maximum amount the enrollee is expected to contribute out-of-pocket to the premium of the benchmark plan in the region. Importantly for the empirical strategy described in Section 2.5, these expected premium contributions are determined based on consumers' income. The ACA defines the benchmark plan as the second cheapest silver plan in the rating area that is meant to capture the price level in the region. Since the second cheapest silver plan can change over time as insurers adjust premiums each year before the open enrollment periods, the goal of linking the premium subsidy mechanism to the price of a benchmark plan is also

[7] The amount of the FPL in 2014, when the ACA marketplaces were introduced, was $11,670 for individuals and $27,910 for a family of five.

[8] APTCS are paid directly by the government to the insurance company on the behalf of the insured.

to protect low income enrollees from increases in the price level. The amount of the premium subsidy is calculated as the difference between the price of the rating area's benchmark plan and the individual's expected contribution cap. Therefore, the premium subsidy of individual i in rating area r decreases in income and increases in the cost of the benchmark plan of the geographic area according to the following function:

$$s_{ir}(b, I) = max\left\{b_{ir} - \overline{p}_i(I), 0\right\}, \tag{2.1}$$

where I denotes family income, b_{ir} is the age-rated premium of the benchmark silver plan of the rating area, and $\overline{p}_i$ is the expected premium contribution calculated based on Table A.4. Intuitively, the subsidy covers the part of enrollees' premium expenditures on the benchmark plan that exceeds their expected contributions.

Note that the subsidy amounts are calculated based on a projected income reported during the open-enrollment period, however they are adjusted later to reflect the actual income reported in the tax filing. Therefore – importantly for the empirical strategy of this paper – this feature of the subsidy mechanism prevents incentives to manipulate the reported income levels.

It is also important to point out that although premium subsidies are linked to the price of a silver plan, enrollees can apply these flat amounts towards the purchase of health insurance plans classified into any metal tiers on the marketplaces. This implies that individuals with sufficiently low income levels might be able to purchase Bronze plans for free, or buy a gold or platinum plan at a lower cost than unsubsidized buyers.

2.3.2.2 Cost Sharing Subsidies

Although premium subsidies help low income people get access to health insurance coverage at affordable prices, using health care services can still impose a large financial burden for this population. The trend of increasing cost sharing has been one of the most dominant features of US health insurance markets during the past decade (Kirzinger et al., 2019). As the quick spread of high deductible health plans substantially increased costs faced by the insured, out-of-pocket health care expenditures reached the catastrophic threshold for low income households (Schoen et al., 2005), and became a common reason for delaying or forgoing care (Brot-Goldberg et al., 2017; Wharam et al., 2018). Therefore, it was an important element of the ACA to introduce a mechanism that helps low income individuals overcome the large financial burden imposed by the trend of rising cost sharing. Cost sharing reduction (CSR) subsidies are designed to help people use medical services at lower costs.

Low income enrollees can be eligible for three different CSR variant silver plans based on their income, as shown on Figure 2.3. For enrollees with incomes between 100% and 150% of the FPL, cost sharing subsidies increase the actuarial value – the share of expenses paid by the insurance company on expectation – of silver plans from the baseline 70% to 94%, increasing the generosity of the plan to the Platinum level. For consumers with incomes from 151% to 200%, CSRs increase the actuarial value to 87% (the level of the Gold tier), while for individuals earning between 201 to 250 % of the FPL, the actuarial value of their silver plan increases to 73%. Above 250% of the FPL consumers are not eligible for CSRs, however they can still qualify for the premium subsidies with incomes up to 400% of the poverty line.

There are two key differences between the two subsidies available at the ACA marketplaces. First, while premium subsidies reduce the cost of purchasing coverage, CSR subsidies lower the cost sharing associated with using a plan, i.e. deductibles, co-insurances, co-payments and out-of-pocket limits. Second, unlike the premium subsidies, cost sharing subsidies are only available for consumers who purchase silver plans.

In the empirical strategy of this paper, I take advantage of the discontinuity in plan generosity induced by the CSR eligibility cut-offs, together with the fact that premium subsidies are an increasing function of income. Combining these features of the two subsidy mechanisms allows me to use data on the premium subsidies, with the rule that determines the amount of the subsidy to recover an imputed income level, and use it to identify the behavioral responses to changes in cost sharing induced by the CSR subsidies.

2.4 Data and Descriptives

The main data source of this paper is the Massachusetts All Payer Claims Database (APCD) distributed by the Center for Health Information and Analysis (CHIA). APCDs are large centralized state-based administrative datasets that collect individual-level information from all private and public insurers operating in a given state with the aims of reducing cost, improving quality and population health (CHIA, 2018). As data submission requirements are standardized across payers and states, today APCDs represent one of the most comprehensive and highest quality data sources for analyzing health care markets. Information collected by the APCDs include individual-level health insurance enrollment records along with corresponding medical, dental, and pharmaceutical claims.

Due to the very low uninsured rate in Massachusetts, the states' APCD offers the widest coverage of the population among similar datasets. The sample used in this paper covers years 2014 and 2015, the first two years of the ACA marketplaces. The main analysis of this paper relies on information contained in the eligibility and medical claims files of the MA APCD.

The eligibility file contains individual-level data on health insurance enrollment records, including an identifier that enables tracking people over time and switching across insurers, demographics (age, gender, zip code), and information on the health insurance plan (insurance company, product type and some financial characteristics). For the empirical strategy of this paper, I take advantage of a special feature of the MA APCD: the availability of rich plan-level information, including marketplace metal tiers and premiums. Moreover, three large insurance companies offering plans through the Connector report premiums at the individual-level. Therefore, for these insurers I observe the premium paid by the enrollee net of the ACA's premium subsidy. Since this information is crucial for my identification strategy, I restrict my analysis for the silver plan enrollees of these insurance companies.

The main outcome of interest in this paper is enrollees' spending on health care. I obtained this information from the medical claims file by calculating the total yearly medical expenditures at the individual-level.

The second data source of the paper is the rating tables provided by the Massachusetts Division of Insurance. These files allowed me to obtain the plan-level base premiums by rating area for the insurance plans included in the sample. As the ACA allows insurance companies to underwrite age based on a specific rule, I used the state-specific age curve of Massachusetts to calculate the age-rated pre-subsidy premium amounts. Although the MA APCD does not include plans identifiers, I was able to link the age-adjusted base premiums with the enrollment data based on plan-level observables.

The combination of these two data sources will be the fundamental elements of the empirical strategy described in Section 2.5. Having access to both enrollees' pre- and post-subsidy premiums allows me to calculate the amount of the premium subsidy, which I can use to impute income by applying the ACA's premium subsidy formula backwards.

2.4.1 Descriptives

In Massachusetts, the first open enrollment period was not smooth since the marketplace needed several upgrades to be compliant with the ACA's stronger regulations. These problems caused substantial delays in enrollment, in 2014 only 31,700 individuals enrolled in ACA marketplace plans, however after overcoming the technical difficulties in the first year, the number of enrollees grew to 165,922 in 2015 (Norris, 2018*b*).

Table 2.1 provides market share data by metal tier and insurance company for the Massachusetts marketplace during the sample period.[9] The table shows that in this state, the marketplace is dominated by the silver tier, which is the focus of this paper. Multiple facts might be related to this extreme high market share of the silver tier. First, the share of consumers receiving premium subsidies is very high (80%) in Massachusetts (Kaiser Family Foundation, 2018*b*), therefore for most buyers, silver plans provide the highest risk protection at the lowest cost compared to the benchmark plan which is designed to follow the price level of the silver tier. Second, only consumers enrolling in silver plans are eligible for the ACA's CSR subsidies that increase the generosity of these plans for the same premium.

The analysis of insurance companies' market shares reveals increasing concentration

[9] The market share data are obtained from the APCD sample and might differ from the official enrollment data released by the Health Connector due to data submission issues.

over time. The top 3 insurers in 2014 were Neighborhood, Harvard Pilgrim and Blue Cross Blue Shield that had about 68% of the enrollees. By 2015, Harvard Pilgrim and Blue Cross Blue Shield lost a substantial part of their market shares, and Boston Medical Center Health Net and Network Health took their dominant positions in the market.

Some key characteristics of the contract space are summarized by Table 2.2. The table shows that premiums increase with the generosity of the metal tier: Bronze plans with the lowest actuarial value are the most inexpensive contract types in the market, and prices are the highest in the most generous Platinum tier. These differences in actuarial value are reflected in the higher deductibles and out-of-pocket limits in the Bronze tier, and these characteristics change proportionally as the metal level increases from Bronze to Platinum. The premium data also reveal a significant increase in the average monthly premiums of each metal tier during the sample period. The number of plans available in the market shows a substantial variation by metal tier: the Gold tier offered the most options for enrollees, while Bronze tier had the lowest number of plans. Unlike on the ACA marketplaces of many other states, the number of plans was relatively stable over time in Massachusetts. This stability of the market is related to the longer history of the state's health insurance exchange and the low uninsured rate of the state that alleviates adverse selection problems.

To be able to implement the empirical strategy described in Section 2.5, I had to impose some restrictions on the baseline sample consisting of the raw APCD data.

The goal of these constraints is to obtain a final sample that includes only individuals enrolled in silver plans who purchased individual coverage for themselves. The first restriction allows me to focus on the only metal tier that provides CSR subsidies for eligible enrollees, and to address the concern of possible risk sorting among the different metal levels. The second sample restriction is needed to be able to invert the ACA's premium subsidy formula. As the amount of the premium subsidy also depends on unobserved family size, combining the coverage level code and the member-subscriber variables of the data, I am able to recover single member households. The final sample includes data from 3 major insurers since individual-level premium data were only available for these payers.

Table 2.3 reports the summary statistics of the full sample and the baseline sample. Enrollees in the raw data are 37.5 years old on average and 51% of them are female. Most individuals live in the Boston Rating Area, while the enrollment is the lowest in the Cape Cod region. Comparing the baseline sample to the raw data shows that enrollees in the final analytic sample are slightly older on average (40.5 years), however the two samples are similar in terms of gender and geographic composition.

2.5 Empirical Strategy

This section describes the empirical framework for estimating the impact of the discontinuities in patient cost sharing induced by the CSR subsidies on the demand for medical care, and the different steps taken to impute enrollees' incomes – a key element of the identification strategy – from a combination of different data sources.

2.5.1 Methodology

The main empirical challenge in identifying demand responses to the CSR subsidies is the endogenous self-selection of individuals into plans with different levels of coverage. The empirical literature aiming to estimate moral hazard in health insurance markets has followed various approaches to address this issue. For instance, in the context of the CSR subsidies, DeLeire et al. (2018) addresses this identification problem by combining data from different sources on ex-ante risk.

In this paper, I use a different strategy by exploiting a regression discontinuity (RD) design that relies on the exogenous variation generated by the subsidy eligibility rule – the discontinuities in actuarial value as income crosses the eligibility thresholds.

The main barrier of using this identification strategy in this context is its complex data requirement. In order to implement a regression discontinuity design to study impacts of the CSR subsidies on the demand for health care, one needs linked individual-level data on enrollment, income and medical claims.

Obtaining information on individuals enrolled in CSR plans has two possible sources. The first one is the enrollment data collected by the health insurance marketplaces. These data files typically contain information on the amount of subsidy or income, however it is not possible to link them to information on health care spending at the individual-level. The second possible data source on CSR enrollees are the APCDs available in several states that collect both individual-level enrollment and health insurance claim information. However, due to the strict data protection standards, APCDs do not contain data on income.

Moreover, most APCDs typically do not include premium data either, nevertheless net premiums paid after the ACA premium subsidies. The Massachusetts APCD is a notable exception in this sense, and in this paper I take advantage of this unique feature to overcome the data limitations of using an RD framework.

A key feature of the Massachusetts APCD that makes it ideal for analyzing the research question of this paper is the availability of premium data. Moreover, for three major insurance companies, the net premium paid by the enrollee after the ACA premium subsidy is also available in the data. This information allows me to impute enrollees' incomes in multiple steps.

First, I combined the amount of the individual-level premium net of the subsidy reported in the APCD with data on plan-level age-adjusted (pre-subsidy) base premiums to calculate the amount of the premium subsidy.

The second main step is to invert the ACA's premium subsidy formula to impute income from the amount of the subsidy. For this purpose, first I used plan-level rating tables and the age curve used in Massachusetts to determine the benchmark plan's (the second cheapest silver plan of the rating area) premium in each year and rating area, which is key part of the premium subsidy formula.

Finally, I used the data on the individual-level premium subsidy amount and the plan-level information on the premium of the second cheapest silver plan to invert the premium subsidy formula and calculate imputed income based on equation 2.1. The ACA's formula linking income as percentage of federal poverty level and the premium cap ($\overline{p}_i(I)$) is shown in Table A.4.

A major difficulty of applying this strategy is that the poverty line depends on family size as shown in Table 2.4. Although the limited demographic data available in the APCD does not contain information on household size, by combining data on coverage-level codes and member-subscriber relationship codes, I was able to observe individuals enrolled in the same plans and identify families in data. Since the actual household size might differ from the number of people enrolled in the same plan, I excluded families from the sample. As next step, I assumed that people who appear to be single based on the family coverage information available in the APCD also report family size of one during the enrollment process.

2.5.2 Estimation

The main identification problem is that comparing the average health care spending of individuals enrolled in different CSR-variant plans might be biased due to unobservables likely to be correlated with both income and health status. To overcome this issue, I implement the regression discontinuity design around the CSR cutoff points (Angrist and Pischke, 2009; Calonico et al., 2014*a*). The idea of applying the RD framework in this context is to exploit the discontinuous selection into treatment (enrollment in different CSR plans) based on income in order to determine the treatment effect. Therefore, by comparing the health care spending of individuals with income levels just *below* and just *above* the CSR eligibility cutoffs, I am able to control the selection based on observables, assuming that close enough to the cutoff, the only difference between the treatment and control groups is their CSR eligibility.

Formally, consider the regression model

$$E[C_i|x_i] = E[C_{0i}|x_i] + E[C_{1i} - C_{0i}|x_i]D_i, \tag{2.2}$$

where

$$D_i = \mathbb{1}\Big(x_i \leq CSR\Big). \tag{2.3}$$

The dependent variable of the model is the annual health care spending of individual C_i. D_i denotes a deterministic treatment status (enrollment in a CSR plan) and x_i is income imputed based on the strategy described in Section 2.5.

Hence, the conditional treatment effect – the impact of the CSR enrollment on health care spending – can be obtained by first defining

$$\lim_{x_i \to CSR^+} E[C_i | D_i = 1, x_i] = \lim_{x_i \to CSR^+} E[C_{1i}(x_i) | x_i] \tag{2.4}$$

and

$$\lim_{x_i \to CSR^-} E[C_i | D_i = 0, x_i] = \lim_{x_i \to CSR^-} E[C_{0i}(x_i) | x_i]. \tag{2.5}$$

Therefore, the conditional average treatment effect (CATE) can be written as

$$\begin{aligned} CATE(CSR) &= E[C_1(CSR) - C_0(CSR) | x = CSR] = \\ &= \lim_{x_i \to CSR^+} E[C | D = 1, x] - \lim_{x_i \to CSR^-} E[C | D = 0, x]. \end{aligned} \tag{2.6}$$

This expression implies that the impact of the CSR subsidies on consumers' demand for medical care can be estimated by comparing the mean spending of individuals who are just below and just above the subsidy eligibility income cutoffs.

Not that an important identification assumption of the model is that individuals cannot manipulate which side of the CSR cutoffs their income level falls. Although the costless decrease in cost sharing could incentivize enrollees to misreport their incomes, the regulatory framework of the market alleviates this concern. As described above, the ACA subsidies are subject to reconciliation in the next year's income tax return, therefore if someone submits an incorrectly projected family income during the annual open enrollment period, the final subsidy amount will be adjusted according to the actual earning. Finally, the data does not provide evidence for bunching around the subsidy thresholds.

I estimate the model nonparametrically by fitting local linear regressions around the CSR eligibility cutoffs, where the running variable is imputed income. A triangular kernel function was used to construct the local polynomial estimator and the optimal bandwidth was chosen by a data-driven selecting algorithm based on Calonico et al. (2014*b*).

2.6 Results

This section presents the main empirical findings of the paper. As a first step, I plot the mean annual health care spending of silver plan enrollees by imputed income expressed as percentage of the poverty level. Figure 2.4 shows the yearly health care utilization for each 5 percentage point bin of the poverty level. The difference in mean spending around 250% of the FPL suggest that there is a substantial increase in the demand for health care around the threshold where enrollees become eligible for the CSR subsidies. Figure 2.5 plots the fitted values of a piece-wise linear regression model of mean annual health care spending on income expressed as FPL % with three cut-offs that are associated with the income levels where the cost sharing subsidy changes the actuarial value of silver plans. Consistent with Figure 2.4, the piece-wise linear regression shows a large and statistically significant jump in health care utilization as enrollees' income falls below 250% of the FPL.

Figures 2.6, 2.7 and 2.8 show the discontinuities in detail by estimating a more flexible nonparametric regressions around the three different income cutoffs, where changes in CSR eligibility occur, at 250, 200 and 150 percents of the FPL. The figures present the estimated sample means for evenly spaced bins, and the number of bins was determined by spacing estimators (Calonico et al., 2014*a*). The fitted local polynomial regressions clearly show that health care utilization is higher in the left neighborhoods of the thresholds, where the actuarial value of the same plan is higher than on the right hand side of the eligibility cutoff. Intuitively, these discontinuous changes in actuarial value imply that consumers on the two sides of the CSR cutoffs face different out-of-pocket costs for the same set of medical services, affecting their demand for medical care.

Table 2.5 confirms the graphical results by presenting the formal regression discontinuity estimates. The table shows that the largest jump in health care utilization occurs at 250% of the FPL: on average, CSR eligible enrollees demand more medical services by $618 per year than individuals enrolled in the same silver plans but ineligible for the CSR subsidy. The coefficient estimate at 200% of the FPL is significant only at 10%, however the estimates reveal another substantial and statistically significant discontinuous change at 150%, where health care utilization increases by $570 per year.

In order to compare the elasticities implied by the regression discontinuity results to estimates found by the literature in different contexts, a simple back of the envelope calculation shows that while the responsiveness of demand to cost sharing at 150% and 200% of the FPL (-0.34 and -0.22, respectively) is similar in magnitude to earlier findings, the elasticity at 250% percent (about -2) is substantially larger than the benchmark estimate of the RAND health insurance experiment (-0.2).

There are multiple potential explanations for these results. First, the elasticities found in this paper were calculated based on the percentage change in the advertised actuarial value at the CSR cutoffs. However, when consumers choose health insurance plans through the ACA marketplaces, some specific cost sharing features of the plans (deductibles, out-of-pocket-maximums, some co-payments) are also displayed to consumers. Table A.1 shows that the percentage change in the deductible is much larger at the CSR cutoffs than the change in the advertised actuarial value (about 20% vs. 10% at the 250% FPL cutoff). Therefore, using the change in deductible for the elasticity calculations yields substantially lower estimates (-0.27, -0.16, -1, respectively). The difference between these two sets of estimates suggest that determining the elasticity of demand in this context is complicated by the fact that the price perceived by consumers might differ from average cost sharing, therefore it is a non-trivial problem to decide what price should be used in the calculations (Einav and Finkelstein, 2017). This argument explains why consumers might be more responsive to changes in deductibles than other plan characteristics used to determine the actuarial value – possibly because this financial characteristic is more salient to consumers or because deductibles also affect demand on the extensive margin, i.e. the probability of seeking care. In addition, due to the nature of deductibles and out-of-pocket maximums, health insurance contracts are non-linear, and the spot price of medical care typically differs from the average price or the marginal cost of health care at the end of the year (Finkelstein et al., 2015). Indeed, empirical evidence shows that consumers behave myopically when facing deductibles because they tend to put larger weights on current prices as opposed to future prices (Ellis et al., 2017). This argument provides an additional explanation why consumers are more sensitive to changes in deductibles.

Another possible behavioral explanation for elasticity estimates is related to the issue that the previously uninsured target population of the ACA is likely to have lower health insurance literacy. Therefore, consumers might not fully understand the concept of actuarial value, however they know whether their out-of-pocket spending is subsidized or not. This argument could also explain why the demand elasticity is higher at the 250% FPL cutoff, where the extensive margin of the subsidy changes (subsidized vs. unsubsidized cost sharing) than at the other CSR cutoffs where only the level of the subsidy changes. Note that the RAND health insurance experiment analyzed changes in cost sharing on the intensive margin, explaining why the elasticities found at 150% and 200% of the FPL are closer to these benchmark estimates.

Finally, this paper provides the first elasticity estimates among the low income ACA marketplace enrollees, many of whom gained access to health insurance coverage for the first time due to the ACA's coverage expansion. Therefore, it is reasonable to believe that this low income population is more responsive to changes in cost sharing than the higher income insured populations studied by the existing literature. A possible explanation for the highest elasticity found at 250% of the FPL could be that this is an income range where consumers react substantially to changes in cost sharing by demanding more non-essential care that they could not afford at the unsubsidized prices, however at even lower levels of income, individuals either care less about their health or use only essential treatments that are less responsive to out-of-pocket costs.

Assuming that enrollees in small neighborhoods on the two sides of the cutoffs differ only in their income, these discontinuous jumps in health care spending suggest that individuals tend to use more medical services only due to the fact that the insurer becomes responsible for a larger share of the incurred expenses due to the ACA's cost sharing subsidy.

Clearly, the interpretation of the results depends on whether these changes in health care spending are associated with over-utilization or this low income population consumes sub-optimal amount of health care without the cost sharing subsidy. Understanding the underlying mechanisms behind this phenomenon is important for designing targeted policies and understanding the welfare effects of this subsidy. Therefore, exploring these micro-foundations behind the results of this paper provides a fruitful area for future research.

2.7 Robustness and Placebo Checks

To provide justification for the changes in health care utilization found in this paper in response to variation in cost sharing at CSR subsidy cutoffs, several robustness checks were conducted.

First, I re-estimated the regression discontinuity models on three different sub-samples that contain a symmetric neighborhood of the three CSR cutoffs in order to address the concern that the estimated jumps in health care utilization are driven by parts of the income distribution that are farther away from the relevant cutoffs. Figures A.3 - A.5 show the mean annual health care spending of silver plan enrollees by imputed income expressed as % of the federal poverty level, along with local polynomial regressions fitted around the CSR subsidy cutoffs. For the CSR cutoff at 150% of the FPL, the model was estimated on a sample restricted to individuals with incomes between 100% and 200% of FPL. For the CSR cutoffs at 200% and 250% of the FPL, the restricted sample included income levels between 150% and 250%, and 200% and 400% of the poverty level, respectively. The results show that the change in healthcare utilization is still the largest at the 250% cutoff, and the jump at 150% of the FPL is also substantial. The discontinuity at 200% of the FPL decreases compared to the main estimates.

As a next step, I conducted placebo checks at different cutoffs that are not associated with discontinuous changes actuarial value due to the ACA's cost sharing subsidies.

Figures A.6 - A.8 show three different placebo checks, where the cutoffs were placed in the mid-points between the CSR cutoffs. The results show that no discontinuities in health care utilization occur at these income levels, confirming the validity of the main results.

Another possible placebo check is to investigate whether there are any discontinuities in health care spending at the income levels associated with the CSR cutoffs for individuals who are enrolled in non-silver plans. The idea behind this test is that only consumers enrolled in silver plans are eligible for the ACA's cost sharing subsidies, therefore there are no discontinuities in the actuarial values of plans in the other metal tiers at the same income levels.

However, due to the limitations associated with using imputed income and some special features of the market, I am not able to run similar regression discontinuity models for the other metal tiers. The reason is that my strategy to impute income relies on the ACA'a premium subsidy, therefore I cannot infer the income of non-subsidized enrollees. In addition, the enrollment pattern on the Massachusetts exchange displays a series of special features: the market is highly dominated by the silver tier (87%), 80% of the enrollees receive premium subsidies and 77% cost sharing subsidies (Henry J Kaiser Family Foundation, 2018*b*). Furthermore, subsidized individuals almost exclusively choose the silver tier and I do not observe people receiving subsidies in the gold or platinum tiers. Note that a few subsidized consumers are enrolled in bronze plans, however these are individuals who pay zero premium after the subsidy.

Equation 2.1 implies that the method used in this paper to impute income leads to a censored income distribution since the income of people receiving no premium subsidies or paying zero premium can be determined only up to a bound.

Finally, a concern about a potential compositional change in the enrollment pools of the different CSR plans provides a possible alternative explanation for the results of this paper. The problem is that the probability of choosing a silver plan might also increase discontinuously at the CSR cutoffs, implying that the risk pool on the more generous side of the cutoff might be sicker. Such discontinuous change in enrollment would be problematic for the argument of this paper that the observed changes in health care utilization at CSR cutoffs measure the behavioral response to more generous coverage as opposed to a selection effect.

However, as discussed above, the distribution of income and the plan choice pattern in this market is very special in a sense that all subsidized consumers choose the silver tier, except of a small set of consumers who pay zero premiums for a bronze plan. These features of the market suggest that the main results of the paper are not driven by a compositional change in the enrollment pool of the silver plans at the CSR cutoffs. Note that this concern might be more problematic in other states where the marketplace is not dominated by a single metal tier and the choice pattern of the subsidized enrollees is more diverse.

2.8 Conclusion

This paper studies the impact of the ACA's cost sharing subsidies on the demand for medical care. Since the CSR subsidies do not directly affect insurance premiums but increase the actuarial value of silver plans sold to low income enrollees on the new health insurance marketplaces, the aim of this paper is to analyze whether this subsidy design leads to any behavioral response in health care utilization.

The special design of the CSR subsidies provides a unique setting to study the price elasticity of health care because the share of expenses paid by the enrollees of otherwise identical health plans falls discontinuously when the income crosses the eligibility cutoffs. This feature of the subsidy design provides an ideal setting for exploiting a regression discontinuity design in order to eliminate the main identification challenge in estimating moral hazard in health care: the endogenous risk sorting of individuals across plans. Since I study only individuals enrolled in silver plans, the change consumers' health care utilization at the eligibility cutoff provides an estimate for moral hazard separated from adverse selection – concepts that are typically hard to separately identify in observational data.

The main barrier of implementing this identification strategy is its specific data requirement because it needs both individual level enrollment, claim and income data. Most data sources currently available for researchers provide access either to enrollment and income data (FPL ranges) from the ACA marketplace enrollment records, or health insurance claims data with no information on income. Linking these two types of datasets at the individual-level is typically not allowed due the strict privacy regulations regarding individual-level health care data.

To overcome this problem, in this paper I take advantage of the special feature of the Massachusetts APCD that it not only provides data on enrollment and medical claims, but it also has individual-specific information on insurance premiums (ie. net of the ACA premium subsidies). To complement this data source, I also obtained pre-subsidy, area-age specific premiums from the Division of Insurance. Then, combining the two different premium variables, I was able to calculate the premium subsidy. As a final step, I used the ACA premium subsidy formula backwards to impute the enrollees' incomes from the amount of subsidies.

The regression discontinuity estimates revealed significant increases in the health care utilization of silver plan enrollees whose income falls just below the CSR cutoffs. For instance, individuals who earn just *below* the CSR eligibility threshold at 250% of the FPL, use more medical services by about $650 each year than the enrollees in the same plan whose income falls just *above* the cutoff point.

This paper shows that moral hazard in health care utilization exists and it also quantifies its magnitude by estimating the price elasticity of medical care. These results provide new insights into how consumers in the under-studied previously uninsured population – many of whom first gained access to insurance coverage due to the ACA's different provisions – respond to changes in out-of-pocket costs, which is an important policy-related question due recent trends in patient cost sharing.

Figures and Tables

Figure 2.1: ACA Marketplaces by state and type

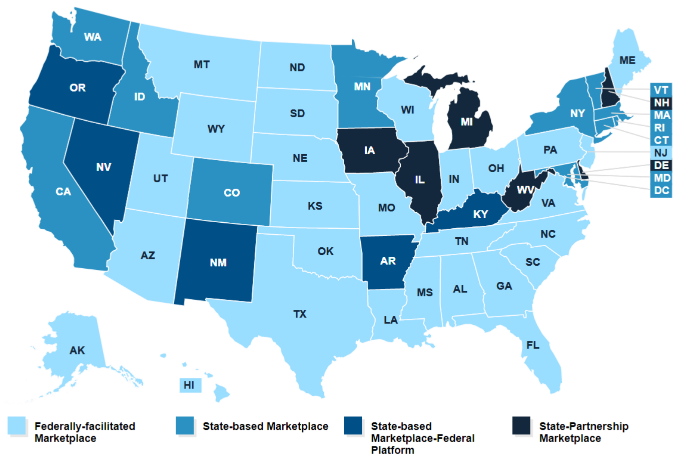

Notes: The map displays the ACA marketplace type of each state in 2019. In state-based marketplaces, enrollees in the small group and individual markets use the websites maintained and established by the state to purchase insurance coverage. In state-based marketplaces that use the federal platform, the state is responsible for operating the marketplace but the underlying IT platform is the provided by the healthcare.gov site. Federally-facilitated marketplaces are operated by HHS and consumers enroll through healthcare.gov.
Source: Kaiser Family Foundation: State Health Insurance Marketplace Types, 2019.

Figure 2.2: ACA Geographic Rating Areas in Massachusetts

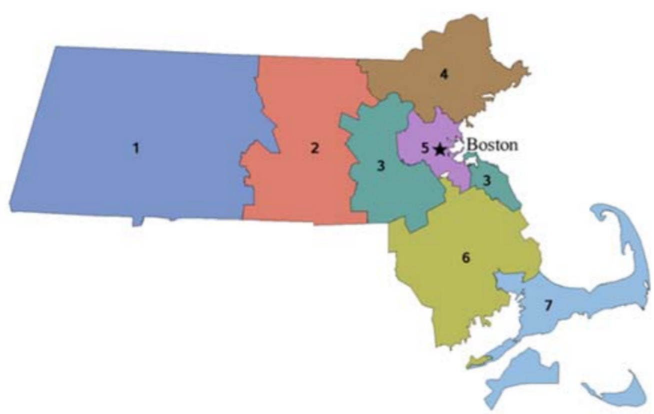

Notes: The map displays the current borders of the ACA geographic rating areas in Massachusetts. In this state, the ACA rating areas are defined as sets of 3 digit zip codes. Source: Andrews, R. and Allison, R. (accessed at https://www.lexjansen.com/sesug/2015/RV-204.pdf on July 2, 2019)

Figure 2.3: The ACA's Subsidy Mechanisms

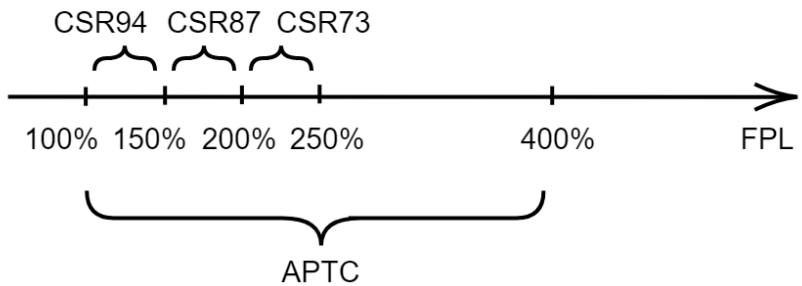

Notes: The figure shows the eligibility thresholds of the ACA's two subsidy mechanisms. Consumers with income between 100% and 400% of the federal poverty level are eligible for the premium subsidies (advance premium tax credits). Individuals enrolling in silver plans and earning less than 250% of the FPL are also eligible for cost sharing reduction (CSR) subsidies. For enrollees with income between 250% and 200% of the FPL, the standard actuarial value (70%) of silver plans increases to 73% due to the cost sharing subsidies. With income between 200% FPL and 150%, the subsidized actuarial value becomes 87%, and 94% below 150% of the FPL. Low income people earning less than 100% of the FPL are not eligible for subsidized marketplace plans.

Figure 2.4: Mean Annual Spending by Income

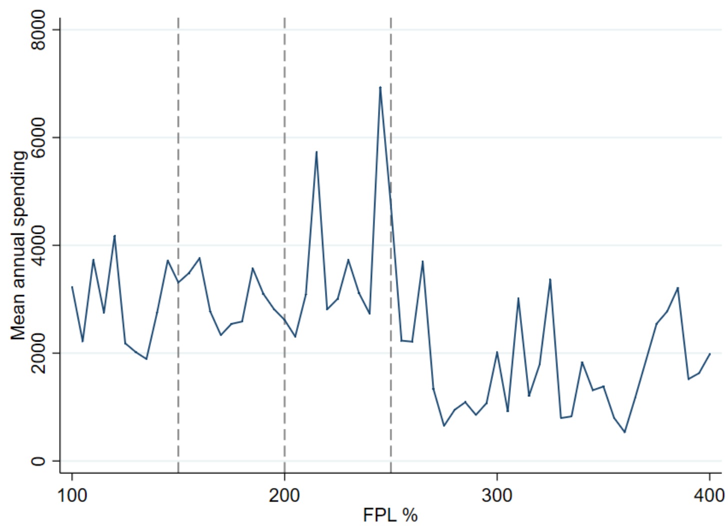

Notes: The figure shows the mean annual health care spending of silver plan enrollees by imputed income expressed as % of the federal poverty level. Mean annual spending was calculated using the allowed amount (i.e. the negotiated payment between the insurer and the provider) at each 5 percentage point bin of the federal poverty level. Calculations are based on individual-level enrollment and medical claims data.

Figure 2.5: Piece-wise Linear Regression with the CSR Cutoffs

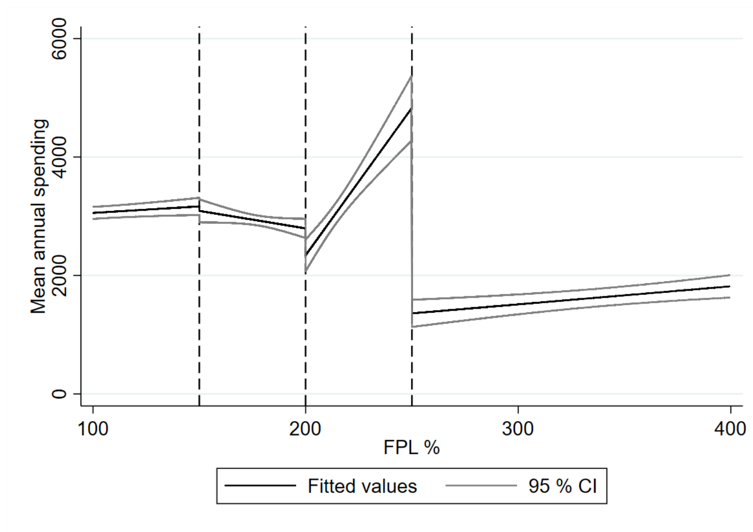

Notes: The figure presents a piece-wise linear regression model of yearly health care spending on income as FPL % with three cut-offs at the income levels where the CSR subsidies change the actuarial values of silver plans. The regression was estimated using individual-level data on health insurance plan enrollment and medical claims data. The figure also displays the 95% confidence intervals.

Figure 2.6: Discontinuity in Annual Spending at 250% FPL

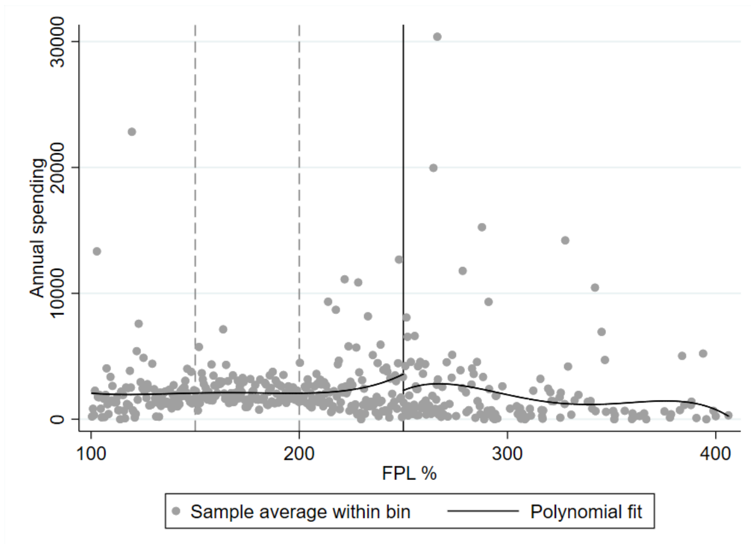

Notes: The figure displays the mean annual health care spending of silver plan enrollees by imputed income expressed as % of the federal poverty level and a fitted local polynomial regression at the discontinuity at 250%.

Figure 2.7: Discontinuity in Annual Spending at 200% FPL

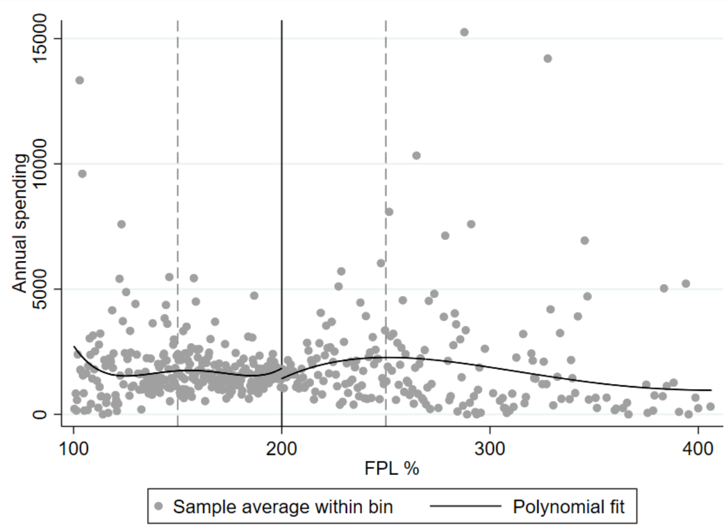

Notes: The figure shows the mean annual health care spending of silver plan enrollees by imputed income expressed as % of the federal poverty level and a fitted local polynomial regression at the discontinuity at 200%.

Figure 2.8: Discontinuity in Annual Spending at 150% FPL

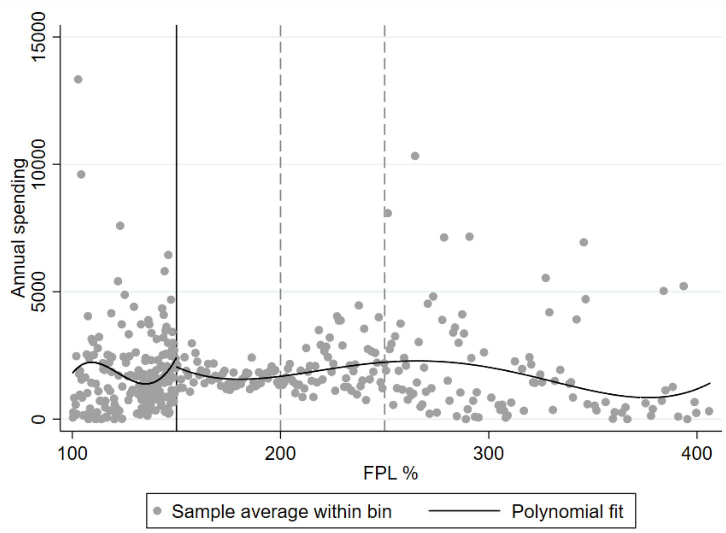

Notes: The figure shows the mean annual health care spending of silver plan enrollees by imputed income expressed as % of the federal poverty level and a fitted local polynomial regression at the discontinuity at 150%.

Table 2.1: Market Shares by Metal Tier and Insurer (%)

	2014	2015
Panel A - Metal tier		
Bronze	13.74	5.36
Silver	71.41	86.64
Gold	9.12	4.38
Platinum	5.73	3.62
Panel B - Insurer		
Neighborhood	37.33	24.18
Harvard Pilgrim	17.95	8.21
BCBS	13.28	1.24
Tufts	10.98	0.77
BMCHN	9.98	21.44
Network Health	3.36	38.07
Other	7.12	6.09

Notes: The table shows market shares by metal tier and insurer on the Massachusetts ACA marketplace. Source: Massachusetts APCD.

Table 2.2: Product Characteristics by Metal Tier

	Bronze		Silver		Gold		Platinum	
	2014	2015	2014	2015	2014	2015	2014	2015
Monthly premium (mean)	196	263	231	309	277	391	318	472
Nr of plans	16	16	19	21	50	49	28	26
Nr of insurers	5	11	9	12	7	10	10	11

Notes: The table displays the characteristics of the supply side by metal tier. Source: Massachusetts APCD, Division of Insurance.

Table 2.3: Summary Statistics

	Full sample		Baseline sample	
	2014	2015	2014	2015
Age (mean)	37.75	41.80	40.50	43.02
Female (%)	51	56	52	62
Rating area shares (%)				
1	8.77	12.82	5.26	7.11
2	9.06	8.89	5.67	9.36
3	15.65	10.51	11.88	7.24
4	18.09	18.09	19.81	25.26
5	31.32	26.90	38.10	23.79
6	9.84	15.10	10.31	17.39
7	7.28	7.23	8.97	9.86

Notes: The table shows the summary statistics of the raw APCD data and the final working sample obtained after applying a set of restrictions (silver plan enrolees of three large insurers with individual-coverage). Source: Massachusetts APCD

Table 2.4: Federal Poverty Level as a Function of Household Size

Persons in family/household	Poverty guideline	
	2014	2015
1	$11,490	$11,670
2	15,510	15,730
3	19,530	19,790
4	23,550	23,850
5	27,570	27,910
6	31,590	31,970
7	35,610	36,030
8	39,630	40,090

Notes: The table presents poverty lines at the time of the 2014 and 2014 Marketplace open enrollment periods. Source: U.S. Department of Health & Human Services, Office of The Assistant Secretary for Planning and Evaluation. https://aspe.hhs.gov/2014-poverty-guidelines

Table 2.5: RD Estimates Around the CSR Subsidy Eligibility Cutoffs

	FPL %		
	150	200	250
Estimate	-569.55***	-376.92*	-617.96***
Standard error	208.03	195.71	172.55
N	32021	32021	32021

Notes: The table shows the regression discontinuity estimates of CSR subsidy eligibility on annual health care spending at the three eligibility cutoffs. The model was estimated nonparametrically by fitting local linear regressions around the CSR eligibility cutoffs, where the running variable is imputed income. A triangular kernel function was used to construct the local polynomial estimator.

Chapter 3

Associations between insurance-related Affordable Care Act policy changes with HPV vaccine completion[1]

3.1 Background

In the US, an estimated 34,800 cancers are attributable to the human papillomavirus (HPV) annually, with cervical and oropharyngeal cancers being the most common (Senkomago et al., 2019). Despite the HPV vaccine being one of the most effective measures to prevent the majority of cervical and other HPV-related cancers (Markowitz

[1]This Chapter is co-authored with Summer Sherburne Hawkins (Boston College, School of Social Work), Jessica Cohen (Harvard T.H. Chan School of Public Health, Department of Global Health and Population), Lydia E. Pace (Brigham and Women's Hospital) and Christopher F. Baum (Boston College, School of Social Work,Department of Economics; German Institute for Economic Research (DIW Berlin), Department of Macroeconomics)

et al., 2014; Senkomago et al., 2019), completion of the multi-dose series remains suboptimal in the US. The Advisory Committee on Immunization Practices (ACIP) has recommended that females receive the HPV vaccine since 2006 (Markowitz et al., 2007) and males since 2011 (for Disease Control et al., 2011). ACIP initially recommended that all 11- or 12-year-olds receive a 3-dose series within a 12-month period until age 26 years if not vaccinated previously for females and until age 21 for males (Markowitz et al., 2014; Petrosky et al., 2015). Revised guidelines in December 2016 permit a 2-dose series for girls and boys who receive their first HPV vaccine at ages 9-14 years, while still requiring a 3-dose series for older or immunocompromised adolescents (Meites et al., 2016). In 2017, the National Immunization Survey–Teen (NIS-Teen) showed that while 68.6% of 13-17-year-old females initiated the HPV vaccine series, only 53.1% completed the recommended sequence; however, corresponding figures for males were 62.6% and 44.3%, respectively (Walker et al., 2018). Increasing the proportion of female and male adolescents who complete the HPV vaccine series is a national priority (Healthy People 2020 IID:11.4 and IID-11.5) (US Department of Health and Human Services, 2020).

The Patient Protection and Affordable Care Act (ACA) rolled out a series of provisions that removed barriers to accessing the HPV vaccine. In September 2010, the ACA required non-grandfathered private plans to cover the HPV vaccine with no patient cost-sharing (US Government, 2010), which removed a significant financial barrier to vaccine uptake (Holman et al., 2014; Tricco et al., 2013; Newman et al., 2018). At that time, the dependent care provision also came into effect allowing young adults up to age 26 years to remain on their parents' private health insurance plans. From January 2014, state and federal Marketplaces offering publicly-subsidized private health insurance plans and insurance companies were banned from denying coverage or charging higher premiums for pre-existing conditions. Also in 2014, newly eligible Medicaid enrollees acquired coverage for the HPV vaccine with ACA's Medicaid expansion (Henry J Kaiser Family Foundation, 2018*a*). As a result of these ACA provisions, over 19 million more people have acquired health insurance coverage (Terlizzi et al., 2018), reducing coverage barriers to HPV vaccine uptake.

Evaluations of the ACA have only examined the effects of the 2010 provisions on HPV vaccination among females. Using self-reported data from national surveys, both studies found that the ACA increased HPV vaccine initiation and completion (Lipton and Decker, 2015; Corriero et al., 2018). However, neither assessed differential effects by health insurance status. Furthermore, we are not aware of any studies that have evaluated the ACA's 2014 insurance-related provisions on HPV vaccine completion among females or males across insurance types. The aim of this study was to examine the associations between the ACA's 2010 provisions and 2014 insurance expansions with HPV vaccine completion by sex and health insurance type.

3.2 Methods

We used All Payer Claims Databases (APCD) from Massachusetts (MA), Maine (ME) and New Hampshire (NH), which collect health insurance claims from insurance companies operating in a given state, therefore covering both the privately- and publicly- (Medicare, Medicaid) insured populations. While the standardized data submission requirements allowed us to pool data from the three states, the length of the study sample differed: MA provided data from January 2011 through December 2015, and ME and NH provided data from January 2009 through December 2015. For NH there were no public claims available from April 2013 through November 2013; personal communication with Rose Hess on April 17, 2018.

Our analytic sample included children and young adults aged 9 to 26 years who had at least one HPV vaccine dose during the study period based on claims associated with Current Procedural Terminology codes 90649 (Gardasil), 90650 (Cervarix) and 90651 (Gardasil 9). We restricted the study period to September 2009 through June 2014 (from September 2011 in MA) in order to ensure the full time period allocated for series completion at either end of data collection. Due to low HPV vaccine initiation prior to the ACIP recommendation for males, we excluded males before October 2011 (Spencer et al., 2018; Hawkins et al., 2019).

Among the 385,998 youth who received at least one HPV vaccine over the study period, we excluded 2,701 individuals who received more than 3 doses because we were unable to observationally distinguish data submission errors from intentional repetitions. This resulted in a total sample size of 383,297 9-26-year-olds, with 194,407 females and 188,890 males. The Boston College Institutional Review Board reviewed this study and considered it exempt.

As the HPV vaccination recommendation over the study period consisted of a series of three shots (Markowitz et al., 2014; Petrosky et al., 2015) we constructed an individual-specific completion measure by aggregating all HPV claims received for a given person. Our primary outcome variable was a binary indicator of HPV completion, defined as receiving a 3-dose series within 12 months from the date of initiation (Spencer et al., 2018).

We examined the associations between two sets of ACA policy changes with HPV vaccine completion during the study period. First, the introduction of the ACA in September 2010 resulted in facilitated dependent care coverage and HPV vaccination without cost-sharing (US Government, 2010). The elimination of cost-sharing allows young people enrolled in non-grandfathered private insurance plans and Medicaid expansion plans (coverage varies by state for traditional Medicaid plans) to receive the HPV vaccine without any copayment, coinsurance or deductible payments (Centers for Disease Control and Prevention, n.d.*a*). However, the three major provisions of the ACA that substantially extended access to health insurance coverage came into effect only in 2014. The reforms that could influence HPV completion included Medicaid expansion in MA (January 1, 2014) and NH (August 15, 2014), but not

ME; the introduction of the health insurance Marketplaces as standardized platforms to purchase publicly-subsidized private health insurance coverage; and the ban of insurers' rating practices based on pre-existing conditions. We created binary indicator variables for each of these policies.

We used demographic and insurer information available in the APCD medical claims files to generate covariates for our analyses: sex (male, female), age groups (9-13, 14-18, 19-26 years) at the initiation of the HPV vaccine, state (MA, ME, NH), insurance type (private, Medicaid), and the year of initiation. Additional participant socio-demographic information, including race/ethnicity, is not recorded in the APCD.

3.3 Statistical analysis

We calculated descriptive statistics for the analytic sample stratified by the different covariates of interest. We calculated the prevalence of HPV completion by the ratio of those who completed the series based on each definition and those who initiated the vaccine during the study period. We first examined the prevalence of HPV vaccine completion within 12 months across the demographic and insurance-related characteristics available in the medical claims data stratified by sex. We then estimated stratified adjusted logistic regression models to examine the predictors of HPV vaccine completion.

Next, we conducted an individual-level regression analysis to examine the associations between the insurance-related ACA policy changes with HPV vaccine completion within 12 months. Since our analytic sample consisted of multiple HPV claims for most individuals, we collapsed our dataset to a single observation per person that contained information of the characteristics of the entire series of HPV claims (total number of doses, date of initiation, number of months between each dose, months between the first and last doses). We estimated logistic regression models where we modeled the probability of series completion as a function of the ACA policy breaks and demographic characteristics. For the dependent variable, we used a binary indictor which takes the value one if our completion definition is satisfied, zero otherwise. In our first, simplest specification, we ran stratified models by sex to identify any heterogeneous policy effects across males and females. Then, we constructed a combined model to directly compare our estimates across these groups. In the final specification we included three-way interaction terms for each policy break by sex and insurance type (males were excluded from the sample at the time of the 2010 ACA break). The model also controlled for the participant's age group, state and year fixed effects.

We then conducted two sensitivity analyses in which we repeated this series of analyses by first, extending the definition of our vaccine completion measure and second, reducing the dose schedule. The former outcome was defined as completion of the 3-dose series within 18 months from the date of initiation (Spencer et al., 2018). In response to the revised 2016 HPV guidelines (Meites et al., 2016), the latter outcome was defined as completion of the 2-dose series within 12 months from the date of initiation (Spencer et al., 2018).

We report the average marginal effects to determine the change in the probability of HPV vaccine completion following each policy change while holding other covariates constant. We calculated differential responses by sex and insurance type as our stratified analysis revealed significant heterogeneity among these groups. Finally, we calculated predictive margins and used F tests to test for the statistical significance of differences by sex and insurance type. We conducted analyses using Stata statistical software version 15.1 (StataCorp, College Station, TX).

3.4 Results

Overall, among females and males who initiated the HPV vaccine over the study period, 27.6% and 28.0%, respectively, completed the series within 12 months. The prevalence of HPV completion was higher among females and males who initiated the series between ages 9 and 14 compared with those who initiated at older ages and privately-insured youth were more likely to complete the series than Medicaid enrollees (Table 3.1). In addition, a higher proportion of youth in MA completed the HPV vaccine series than in ME or NH. Associations persisted in adjusted models and were consistent for females and males. The prevalence of vaccine completion also fluctuated significantly across the study period. In adjusted models, vaccine completion in 2013-2014 was significantly lower than the baseline year for females and males.

Among females, the introduction of the ACA in 2010 was associated with increases in HPV vaccine completion for both the privately-insured (0.043; 95% confidence interval [CI]: 0.036-0.061) and Medicaid enrollees (0.057; 95% CI: 0.032-0.081) (Table 3.2). Similarly, the 2014 health insurance expansions were associated with increases in vaccine completion for females with private insurance (0.094; 95% CI: 0.082-0.107) and Medicaid (0.085; 95% CI: 0.068-0.102). Across both insurance-related ACA policy changes, there were no differences in vaccine completion for females by insurance type ($p=0.3$ and $p=0.3$, respectively).

Among males, the 2014 ACA reforms were associated with increases in HPV vaccine completion for both the privately-insured (0.051; 95% CI: 0.039-0.063) and Medicaid enrollees (0.034; 95% CI: 0.017-0.050) (Table 3.2). Although the differences in vaccine completion among males by insurance type were minimal ($p=0.05$), the 2014 insurance reform had a larger effect among both privately- and publicly-insured females than males ($p<0.01$ and $p<0.01$, respectively).

In the first sensitivity analysis with the extended HPV vaccine completion window to 18 months, 32.5% of females and 34.3% of males completed the 3-dose series. The pattern of associations between participant characteristics and HPV completion within 18 months were consistent with completion by 12 months (Supplemental Table A.5). Similar to the main specification, the prevalence of vaccine completion fluctuated across the study period and adjusted trends showed significantly lower completion in 2013-2014 for females and males. Models using the extended completion window delivered similar results, with coefficients of a comparable magnitude to those from the 12-month model, confirming the robustness of our findings (Table 3.2).

In the second sensitivity analysis with the 2-dose schedule, 46.9% of females and 46.5% of males completed the series within 12 months. The pattern of associations between participant characteristics and HPV completion within 12 months were consistent (Supplemental Table A.6). Similar to the main specification, both ACA provisions were associated with increases in HPV vaccine completion among females and privately-insured males (Table 3.2). In contrast, among publicly-insured males that received their first HPV vaccine, there was some evidence that the 2014 ACA reforms were associated with decreases in completion of the 2-dose series.

3.5 Discussion

We have shown that despite low HPV vaccine completion overall, the 2010 ACA provision modestly increased vaccine completion among females and the 2014 ACA-related health insurance reforms further increased completion among females and males. While both privately- and publicly-insured youth benefitted from the ACA, the 2014 insurance reforms had a larger effect among females than males. Furthermore, extending the HPV vaccine completion window from 12 to 18 months increased completion from approximately 28% to 33-34% in females and males. Despite minimal sex differences in HPV vaccine completion in these three New England states, publicly-insured youth continued to have lower completion than their privately-insured counterparts. Our findings suggest that Medicaid expansion in the 14 remaining states that have not yet expanded Medicaid (Henry J Kaiser Family Foundation, 2017) could increase HPV vaccine completion among publicly-insured youth and reduce their longer-term risk of HPV-related cancers (Markowitz et al., 2014).

State and national estimates of HPV vaccine initiation and completion are based on self-report of the number of HPV doses received without specifying a time frame for receipt (Walker et al., 2018; Lipton and Decker, 2015; Corriero et al., 2018). Although the ACIP recommendation for completion of the 3-dose series at the time of this study required the third dose given at least 6 months after the first dose, this information is not recorded in surveys such as the NIS-Teen. A strength of population-based APCDs and claims-based studies (Spencer et al., 2018; Liu et al., 2016) are the use of objectively recorded medical claims with associated dates of service. Consistent with studies using private claims data (Spencer et al., 2018; Liu et al., 2016), we used a 12-month time frame for the 3-dose series completion from the month the first dose was received. We found an average prevalence of the HPV vaccine 3-dose series completion to be 30.5% among females and 27.5% among males in 2014, among those who initiated the HPV vaccine, with corresponding estimates of 69.3% and 57.8% from the NIS-Teen that same year (Reagan-Steiner et al., 2015). Corroborating results by Spencer and colleagues (Spencer et al., 2018), we found a higher prevalence of

HPV vaccine completion by extending the vaccine window to 18 months—specifically, completion increased to 36.8% among females and 34.2% among males in 2014. This suggests that youth are often completing the required number of HPV doses well-beyond the recommended window. As the new 2016 ACIP guidelines requires only two doses for adolescents starting the series at ages 9-14 years (Meites et al., 2016), it will be important to monitor whether compliance increases due to fewer return visits. In a sensitivity analysis, we found overall consistent effects of both sets of ACA provisions on 2-dose HPV vaccine completion within 12 months except for publicly-insured males. However, as noted by Spencer and colleagues, it is challenging to use historical data to estimate future compliance as both the number of doses and timing of the second dose were modified (Spencer et al., 2018).

This study also contributes to the literature in its inclusion of males. Prior evaluations of the 2010 ACA provision on HPV completion have been limited to females. Lipton and Decker found that the ACA increased HPV vaccine completion by 5.8 percentage points among women aged 19-25 years compared to 18- or 26-year-olds (Lipton and Decker, 2015). Corriero and colleagues also found that after the ACA was implemented, females aged 9-33 years were 5.8 times more likely to complete the 3-dose HPV series (Corriero et al., 2018). Despite the 2010 ACA provision only benefitting the privately-insured, neither study examined differential effects by insurance type (Lipton and Decker, 2015; Corriero et al., 2018). Consistent with these studies, we found that the ACA increased HPV vaccine completion among females; however, our results show that both privately- and publicly-insured females benefited from the policy—increasing vaccine completion by 4.3 and 5.7 percentage points, respectively. Although Lipton and Decker found no effect of the ACA on self-reported awareness of the HPV vaccine (Lipton and Decker, 2015), increased insurance coverage and access to providers as a result of the ACA (Wisk and Sharma, 2019; Adams et al., 2018; Sommers et al., 2013) may have resulted in an increased awareness of vaccination among providers for all young women.

This is one of the first studies to evaluate the 2014 ACA-related health insurance reforms on HPV vaccine completion. We found that both privately- and publicly-insured females and males increased vaccine completion after the reforms came into effect. For the privately-insured, state and federal Marketplaces offered publicly-subsidized private health insurance plans as well as the protections for those with pre-existing conditions. For the publicly-insured, Medicaid expansion due to the ACA provided preventive services with no cost sharing for those who qualified. Our findings demonstrate that increasing health insurance coverage, access to preventive services and reducing costs through the ACA directly benefited youth by increasing HPV vaccine completion rates.

In addition to the lack of insurance coverage and cost (Holman et al., 2014; Tricco et al., 2013; Newman et al., 2018), having a regular medical provider, lack of provider recommendation, and being unaware or forgetting about additional doses have been identified as barriers specific to HPV vaccine completion (Holman et al., 2014). Across these three New England states, youth who initiated the first vaccine at a younger age were more likely to complete the series, consistent with other studies (Spencer et al., 2018; Liu et al., 2016; Agawu et al., 2019). We also found that privately-insured youth had higher HPV vaccine completion than Medicaid recipients. In contrast, the NIS-Teen has reported that publicly-insured adolescents were more likely to initiate and complete the HPV series than their privately-insured counterparts (Walker et al., 2018). Based on the NIS-Teen, New England has the highest prevalence of HPV vaccine series completion (63.3%) across all regions and the US overall (48.6%) (Walker et al., 2018). This may be due to differences in state programs available for Medicaid recipients or social norms related to HPV vaccination that differ regionally. Since we found strong policy effects across states with high levels of HPV vaccine completion, this suggests that other states may have experienced the same or larger gains in HPV vaccine completion in response to both ACA provisions.

3.6 Limitations

There are a number of limitations to note. APCDs only record insured enrollees, so uninsured youth, recognized to have the lowest uptake of HPV vaccine initiation and completion (Walker et al., 2018), were not included. Due to increases in health insurance coverage as a result of the ACA (Terlizzi et al., 2018), the composition of the insured population in the APCDs also likely changed over the study period to include those with continuous coverage as well as the newly-insured. Despite known racial/ethnic disparities in HPV vaccine initiation and completion (Hirth, 2019; Agénor et al., 2018), APCDs also do not consistently collect information on race/ethnicity. There are other factors associated with HPV vaccine completion that we were not able to examine, including provider type or type of private insurance plan (Spencer et al., 2018; Liu et al., 2016; Agawu et al., 2019); however, as these factors are more likely to be a result of ACA implementation, i.e. on the causal pathway, rather than moderate the relationship, they are not likely to confound the associations between implementation of the ACA provisions and HPV vaccine completion. Youth may have received the HPV vaccines through the Vaccines for Children Program, a program providing free vaccines for 18-year-olds and younger who are uninsured, underinsured, eligible for Medicaid, or American Indian or Alaskan Native (Centers for Disease Control and Prevention, n.d.*b*). Vaccines received through the Program or youth who paid for the vaccines out-of-pocket were not recorded in the APCD and may under-estimate the true prevalence of vaccine completion.

3.7 Conclusions

Among cancers attributable to HPV in the US, 92% are associated with HPV types targeted by the 9-valent HPV vaccine (Senkomago et al., 2019). Using APCDs from these three states, we previously found that the 2010 and 2014 ACA provisions increased HPV vaccine initiation rates among males and Medicaid recipients, but females and youth with private insurance did not exhibit these same increases in HPV vaccine uptake (Hawkins et al., 2019). Among youth that initiated the HPV vaccine, we have shown that implementation of both sets of ACA provisions increased completion of the series independent of sex and with similar gains among privately- and publicly-insured youth. Thus, expanding Medicaid across the remaining states (Henry J Kaiser Family Foundation, 2017) could further increase HPV vaccine initiation and completion as well as prevent HPV-related cancers (Markowitz et al., 2014) among this population.

Figures and Tables

Table 3.1: Descriptive statistics for HPV vaccine completion defined as completion of the 3-dose series within 12 months

	Females				Males			
	N	%	% Comp.	Adjusted OR	N	%	% Comp.	Adjusted OR
Age of first dose								
9-14	96702	49.7	30.8	1	72071	38.2	29.7	1
15-18	56407	29.0	26.8	0.79 (0.77-0.81)	87336	46.2	28.6	0.86 (0.85-0.88)
19-26	41298	21.2	21.1	0.58 (0.56-0.59)	29483	15.6	21.9	0.60 (0.58-0.62)
Insurance type								
Medicaid	28827	14.8	17.6	1	22588	12.0	17.8	1
Private	165580	85.2	29.3	2.60 (2.08-3.24)	166302	88.0	29.4	1.93 (1.86-2.01)
Year of first dose								
2009	3565	1.8	21.7	1	-	-	-	-
2010	11397	5.9	28.2	1.42 (1.29-1.56)	-	-	-	-
2011	28463	14.6	26.8	0.94 (0.83-1.07)	6484	3.4	31.8	1
2012	59371	30.5	32.6	1.14 (1.01-1.30)	74231	39.3	34.8	1.09 (1.04-1.16)
2013	64134	33.0	22.2	0.62 (0.55-0.70)	78172	41.4	21.4	0.52 (0.49-0.55)
2014	27477	14.1	30.5	0.75 (0.65-0.86)	30003	15.9	27.5	0.60 (0.55-0.65)
State								
Massachusetts	136351	70.1	29.2	1	151877	80.4	29.5	1
Maine	24761	12.7	24.4	0.76 (0.73-0.78)	15160	8.0	19.3	0.59 (0.56-0.61)
New Hampshire	33295	17.1	23.2	0.67 (0.65-0.69)	21853	11.6	23.6	0.72 (0.70-0.75)

Notes: Abbreviations: CI, confidence interval; HPV, human papillomavirus

Table 3.2: Marginal effects of the associations between insurance-related ACA policy changes and HPV vaccine completion by sex and insurance type

	ACA 2010		ACA 2014	
	Marginal effect (95% CI)	P	Marginal effect (95% CI)	P
3-dose HPV series completion within 12 months[ab]				
Male				
Private insurance	-		0.051 (0.039-0.063)	<0.01
Medicaid	-		0.034 (0.017-0.050)	<0.01
Female				
Private insurance	0.043 (0.036-0.061)	<0.01	0.094 (0.082-0.107)	<0.01
Medicaid	0.057 (0.032-0.081)	<0.01	0.085 (0.068-0.102)	<0.01
3-dose HPV series completion within 18 months[ab]				
Male				
Private insurance	-		0.060 (0.048-0.071)	<0.01
Medicaid	-		0.040 (0.022-0.058)	<0.01
Female				
Private insurance	0.033 (0.014-0.052)	<0.01	0.109 (0.097-0.121)	<0.01
Medicaid	0.069 (0.040-0.098)	<0.01	0.100 (0.082-0.118)	<0.01
2-dose HPV series completion within 12 months[ad]				
Male				
Private insurance	-		0.018 (0.006-0.029)	<0.01
Medicaid	-		-0.026 (-0.046- -0.006)	0.01
Female				
Private insurance	0.094 (0.055-0.134)	<0.01	0.044 (0.032-0.056)	<0.01
Medicaid	0.025 (0.005-0.046)	0.02	0.035 (0.015-0.054)	<0.01

Notes: *a*: Models includes interaction between policy breaks, sex, and insurance type; adjusted for age group, state and year fixed effects. *b*: Recommendation (as of 2015) for females and males aged 9-26 years to receive 3-dose HPV series within 12 months after first dose. *c*: Sensitivity analysis based on recommendation (as of 2015) to extend window of series completion to 18 months after first dose. *d*: Sensitivity analysis based on 2016 recommendation to receive 2-dose HPV series within 12 months after first dose.

Bibliography

Abaluck, J., and J. Adams. 2017. "What Do Consumers Consider Before They Choose? Identification from Asymmetric Demand Responses." NBER Working Paper No. 23566.

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

Abaluck, J., and J. Gruber. 2016. "Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance." *American Economic Review*, 106(8): 2145–2184.

Abelson, R. 2016. "UnitedHealth to Pull Back From Insurance Exchanges, Citing Losses." *The New York Times.* Accessed at https://www.nytimes.com/2016/04/20/business/citing-losses-unitedhealth-to-pull-back-from-obamacare.html on Apr 14, 2018.

Adams, S. H., M.J. Park, L. Twietmeyer, C.D. Brindis, and C.E. Irwin. 2018. "Association between adolescent preventive care and the role of the Affordable

Care Act." *JAMA pediatrics*, 172(1): 43–48.

Agawu, A., A. L. Hanlon, A. M. Buttenheim, L. Song, A. G. Fiks, and K. A. Feemster. 2019. "Disparities in Human Papillomavirus Vaccine Series Completion by Adolescent Males: A Retrospective Cohort Study." *Academic pediatrics.*

Agénor, M., A. E. Pérez, S. M. Peitzmeier, and S. Borrero. 2018. "Racial/ethnic disparities in human papillomavirus vaccination initiation and completion among US women in the post-Affordable Care Act era." *Ethnicity & health*, 1–15.

Almond, D., and J. J. Doyle. 2011. "After Midnight: A Regression Discontinuity Design in Length of Postpartum Hospital Stays." *American Economic Journal: Economic Policy*, 3(3): 1–34.

Angrist, J. D., and J.-S. Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton University Press.

Aron-Dine, A.and Einav, L., and A. Finkelstein. 2013. "The RAND Health Insurance Experiment, Three Decades Later." *Journal of Economic Perspectives*, 27(1).

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

Askin, E., and N. Moore. 2016. *The Health Care Handbook.* Washington University of St.Louis.

Azevedo, E., and D. Gottlieb. 2017. "Perfect Competition in Markets with Adverse Selection." *Econometrica*, 85 (1): 67–105.

Brot-Goldberg, Z. C., A. Chandra, B. R. Handel, and J.n T. Kolstad. 2017. "What does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics*." *The Quarterly Journal of Economics*, 132(3): 1261–1318.

Bühlmann, P., and S. van de Geer. 2011. *Statistics for High-Dimensional Data: Methods, Theory and Applications.* Springer.

C4HCO. 2017. "Connect for Health Colorado Open Enrollment Report." http://connectforhealthco.com/about-us/metrics-reports/.

Calonico, S., M. D. Cattaneo, and R. Titiunik. 2014*a*. "Robust Data-Driven Inference in the Regression-Discontinuity Design." *The Stata Journal*, 14(4): 909–946.

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

Centers for Disease Control and Prevention. "STD Preventive Services." Available at: https://www.cdc.gov/nchhstp/preventionthroughhealthcare/preventiveservices/std.htm. Accessed December 7, 2018.

Centers for Disease Control and Prevention. "Vaccines for Children Program (VFC)." Available at: https://www.cdc.gov/vaccines/programs/vfc/index.html. Accessed Novermber 29, 2018.

Chade, H., and J. Swinkels. 2016. "Disentangling Moral Hazard and Adverse Selection."

Chandra, A., and D. O. Staiger. 2007. "Productivity Spillovers in Health Care: Evidence from the Treatment of Heart Attacks." *Journal of Political Economy*, 115(1): 103–140.

Chandra, A., B Handel, and J. Schwartzstein. 2019. "Behavioral Economics and Health-Care Markets." In *Handbook of Behavioral Economics: Foundations and Applications 2.* , ed. Stefano DellaVigna B. Douglas Bernheim and David Laibson. Amsterdam, Elsevier.

CHIA. 2018. "The Massachusetts All Payer Claims Database." http://www.chiamass.gov/chia-data/.

Chiappori, P.-A., and B. Salanie. 2000. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy*, 108: 56–78.

CIVHC. 2017. "The Colorado All Payer Claims Database." https://www.civhc.org/.

CMS. 2017. "Market Rating Reforms: State Specific Rating Variations." Accessed at https://www.cms.gov/CCIIO/Programs-and-Initiatives/Health-Insurance-Market-Reforms/state-rating.html on Sep 20, 2018.

CMS. 2018. "Market Rating Reforms: State Specific Geographic Rating Areas." Accessed at https://www.cms.gov/CCIIO/Programs-and-Initiatives/Health-Insurance-Market-Reforms/state-gra.html on July 2, 2019.

Colorado Health Institute. 2018. "Colorado Health Access Survey." https://www.coloradohealthinstitute.org/research/colorado-health-access-survey.

Corriero, R., J. L. Gay, S. W. Robb, and E. W. Stowe. 2018. "Human papillomavirus vaccination uptake before and after the Affordable Care Act: variation according to insurance status, race, and education (NHANES 2006-2014)." *Journal of pediatric and adolescent gynecology*, 31(1): 23–27.

Cutler, D. M., A. Finkelstein, and K. McGarry. 2008. "Preference Heterogeneity and Insurance Markets: Explaining a Puzzle of Insurance." *American Economic Review*, 98(2): 157–62.

DeLeire, T., A. Chappel, K. Finegold, and E. Gee. 2017. "Do individuals respond to cost-sharing subsidies in their selections of marketplace health insurance plans?" *Journal of Health Economics*, 56: 71 – 86.

DeLeire, T., K. Lavetti, and N. R. Ziebarth. 2018. "How Do Low-Income Enrollees in the Affordable Care Act Exchanges Respond to Cost-Sharing?"

Diamond, R., M. J. Dickstein, T. McQuade, and P. Persson. 2019. "Take-Up, Drop-Out, and Spending in ACA Marketplaces."

Duggan, M., P. Healy, and F. Scott Morton. 2008. "Providing Prescription Drug Coverage to the Elderly: America's Experiment with Medicare Part D." *Journal of Economic Perspectives*, 22(4): 69–92.

Einav, L., A. Finkelstein, and J. Levin. 2010*a*. "Beyond Testing: Empirical Models of Insurance Markets." *Annual Review of Economics*, 2: 311–336.

Einav, L., A. Finkelstein, and M.R. Cullen. 2010*b*. "Estimating Welfare in Insurance Markets Using Variation in Prices." *Quarterly Journal of Economics*, 125: 877–921.

Einav, L., A. Finkelstein, and P. Schrimpf. 2017. "Bunching at the kink: Implications for spending responses to health insurance contracts." *Journal of Public Economics*, 146: 27 – 40.

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(1): 178–219.

Einav, Liran, and Amy Finkelstein. 2017. "Moral Hazard in Health Insurance: What We Know and How We Know It." NBER Working Paper 24055.

Ellis, R. P., B. Martins, and W. Zhu. 2017. "Health care demand elasticities by type of service." *Journal of Health Economics*, 55: 232 – 243.

Ericson, K.M. 2014. "Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange." *American Economic Journal: Economic Policy*, 6: 38–64.

Ericson, K. M. 2016. "When Consumers Do Not Make an Active Decision: Dynamic Default Rules and their Equilibrium Effects."

Ericson, K.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.

Ericson, K. M., and J. R. Sydnor. 2018. "Liquidity Constraints and the Value of Insurance." National Bureau of Economic Research Working Paper 24993.

Ericson, K. M., J. Kingsdale, T. Layton, and A. Sacarny. 2017. "Nudging Leads Consumers In Colorado To Shop But Not Switch ACA Marketplace Plans." *Health Affairs*, 36(2): 311–319.

Farrell, J., and P. Klemperer. 2007. "Coordination and Lock-In: Competition with Switching Costs and Network Effect." *Handbook of Industrial Organization*, 3: 1967–2072.

Finkelstein, A., K. J. Arrow, J. Gruber, J. P. Newhouse, and J. E. Stiglitz. 2015. *Moral Hazard in Health Insurance.* Columbia University Press.

Finkelstein, A., S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. P. Newhouse, H. Allen, K. Baicker, and Oregon Health Study Group. 2012.

"The Oregon Health Insurance Experiment: Evidence from the First Year*." *The Quarterly Journal of Economics*, 127(3): 1057–1106.

Fleitas, S., G. Gowrisankaran, and A. Lo Sasso. 2019. "Reclassification Risk in the Small Group Health Insurance Market." NBER Working Paper 24663.

for Disease Control, Centers, Prevention (CDC, et al. 2011. "Recommendations on the use of quadrivalent human papillomavirus vaccine in males–Advisory Committee on Immunization Practices (ACIP), 2011." *MMWR. Morbidity and mortality weekly report*, 60(50): 1705.

Gabel, J., H. Whitmore, M. Green, A. Call, S. Stromberg, and R. Oran. 2016. "The ACA's Cost-Sharing Reduction Plans: A Key to Affordable Health Coverage for Millions of U.S. Workers." *The Commonwealth Fund: Realizing Health Reform's Potential.*

Gasteier, A. M., E. Brice, and M. Woltmann. 2018. "Why Massachusetts Stands Out In Marketplace Premium Affordability." Health Affairs Blog, Accessed at "https://www.healthaffairs.org/do/10.1377/hblog20180903.191590/full/" on July 2, 2019.

Gordon, S. H., B. D. Sommers, I. B. Wilson, O. Galarraga, and A. N. Trivedi. 2019. "Risk Factors for Early Disenrollment From Colorado's Affordable Care Act Marketplace." *Medical Care*, 57(1): 49 – 53.

Grubb, M. D. 2015. "Failing to Choose the Best Price: Theory, Evidence, and Policy." *Review of Industrial Organization*, 47(3): 303–340.

Hackmann, M.B., J. Kolstad, and A. Kowalski. 2015. "Adverse Selection and an Individual Mandate: When Theory Meets Practice." *American Economic Review*, 105(3): 1030–1066.

Handel, B., and J. Kolstad. 2015. "Getting the Most from Marketplaces: Smart Policies on Health Insurance Choice." Brookings Hamilton Project Discussion Paper.

Handel, B., I. Hendel, and M. D. Whinston. 2015. "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk." *Econometrica*, 83(4): 1261–1313.

Handel, B., J. Kolstad, and J. Spinnewijn. 2019. "Information Frictions and Adverse Selection: Policy Interventions in Health Insurance Markets." *The Review of Economics and Statistics*, forthcoming.

Handel, B.R. 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review*, 103: 2643–2682.

Hastie, T., R. Tibshirani, and J. Friedman. 2009. *The Elements of Statistical Learning Data Mining, Inference, and Prediction, Second Edition.* Springer.

Hawkins, S.S., K. Horvath, J. Cohen, L.E. Pace, and C.F. Baum. 2019. "The impact of policy changes on HPV vaccine completion: A natural experiment."

Heidhues, P., and B Kőszegi. 2018. "Behavioral Industrial Organization." In *Handbook of Behavioral Economics - Foundations and Applications 1.* Vol. 1 of *Handbook of Behavioral Economics: Applications and Foundations 1*, , ed. B. Douglas Bernheim, Stefano DellaVigna and David Laibson, 517 – 612. North-Holland.

Henry J Kaiser Family Foundation. 2017. "Medicaid expansion enrollment." Available at: https://www.kff.org/health-reform/state-indicator/medicaid-expansion-enrollment/?currentTimeframe=0&sortModel=Accessed March 7, 2019.

Henry J Kaiser Family Foundation. 2018*a*. "The HPV vaccine: Access and use in the U.S." Available at: https://www.kff.org/womens-health-policy/fact-sheet/the-hpv-vaccine-access-and-use-in-the-u-s/. Accessed November 29, 2018.

Henry J Kaiser Family Foundation. 2018*b*. "Marketplace Effectuated Enrollment and Financial Assistance." Available at: https://www.kff.org/other/state-indicator/effectuated-marketplace-enrollment-and-financial-assistance/?currentTimeframe=0sortModel=Accessed March 12, 2020.

Hirth, J. 2019. "Disparities in HPV vaccination rates and HPV prevalence in the United States: a review of the literature." *Human vaccines & immunotherapeutics*, 15(1): 146–155.

Ho, K., and R. S. Lee. 2017. "Insurer Competition in Health Care Markets." *Econometrica*, 85(2): 379–417.

Ho, K, J. Hogan, and F. Scott Morton. 2017. "The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program." *The RAND Journal of Economics*, 48(4): 877–905.

Hole, A. R. 2007. "Fitting mixed logit models by using maximum simulated likelihood." *The Stata Journal*, 7 (3): 388–401.

Holman, D. M., V. Benard, K. B. Roland, M. Watson, N. Liddon, and S. Stokley. 2014. "Barriers to human papillomavirus vaccination among US adolescents: a systematic review of the literature." *JAMA pediatrics*, 168(1): 76–82.

Jacobs, P. D., N. Duchovny, and B. J. Lipton. 2016. "Changes In Health Status And Care Use After ACA Expansions Among The Insured And Uninsured." *Health Affairs*, 35(7): 1184–1188.

Jaffe, S., and M. Shepard. 2016. "Price-Linked Subsidies and Health Insurance Markups." Working Paper.

Kaiser Family Foundation. 2014. "Assessing Americans' Familiarity With Health Insurance Terms and Concepts." Accessed at https://www.kff.org/health-reform/poll-finding/assessing-americans-familiarity-with-health-insurance-terms-and-concepts/ on May 20, 2019.

Kaiser Family Foundation. 2017. "Summary of the Affordable Care Act." Accessed at http://files.kff.org/attachment/Summary-of-the-Affordable-Care-Act on June 16, 2019.

Kaiser Family Foundation. 2018*a*. "Key Facts about the Uninsured Population."

Kaiser Family Foundation. 2018*b*. "Marketplace Effectuated Enrollment and Financial Assistance."

Kirzinger, A., C. Munana, B. Wu, and M. Brodie. 2019. "Data Note: Americans' Challenges with Health Care Costs." Accessed at "https://www.kff.org/health-costs/issue-brief/data-note-americans-challenges-health-care-costs/" on July 11, 2019.

Krinn, K., P. Karaca-Mandic, and L.A. Blewett. 2015. "State-Based Marketplaces Using 'Clearinghouse' Plan Management Models Are Associated With Lower Premiums." *Health Affairs*, 34(1): 161–169.

Lipton, B.J., and S.L. Decker. 2015. "ACA provisions associated with increase in percentage of young adult women initiating and completing the HPV vaccine." *Health affairs*, 34(5): 757–764.

Liu, G., L. Kong, and P. Du. 2016. "HPV vaccine completion and dose adherence among commercially insured females aged 9 through 26 years in the US." *Papillomavirus Research*, 2: 1–8.

Luco, F. 2017. "Switching Costs and Competition in Retirement Investment." forthcoming: AEJ: Microeconomics.

Madrian, B., and D. Shea. 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *The Quarterly Journal of Economics*, 116 (4): 1149–1187.

Manning, W. G., J. P. Newhouse, N. Duan, E. B. Keeler, and A. Leibowitz. 1987. "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment." *The American Economic Review*, 77(3): 251–277.

Markowitz, L. E., E. F. Dunne, M. Saraiya, H. W. Chesson, C. R. Curtis, J. Gee, J. A. Bocchini Jr, and E. R. Unger. 2007. "Human papillomavirus vaccination: recommendations of the Advisory Committee on Immunization Practices (ACIP)." *Morbidity and Mortality Weekly Report: Recommendations and Reports*, 56(RR–2):1–24.

Markowitz, L. E., E. F. Dunne, M. Saraiya, H. W. Chesson, C. R. Curtis, J. Gee, J. A. Bocchini Jr, and E. R. Unger. 2014. "Human papillomavirus vaccination: recommendations of the Advisory Committee on Immunization Practices (ACIP)." *Morbidity and Mortality Weekly Report: Recommendations and Reports*, , (5): 63(RR–05):1–30.

Massachusetts Health Connector. 2019. "Massachusetts Individual Mandate." Accessed at https://www.mahealthconnector.org/stay-covered on July 2, 2019.

McFadden, D. 1973. *Conditional Logit Analysis of Qualitative Choice Behavior. In Frontiers in Econometrics, ed. P. Zarembka.* Academic Press: New York.

Meites, E., A. Kempe, and L. E. Markowitz. 2016. “Use of a 2-dose schedule for human papillomavirus vaccination—updated recommendations of the Advisory Committee on Immunization Practices.” *Morbidity and Mortality Weekly Report*, 65(49): 1405–1408.

NCHS. 2018. “Health Insurance Coverage: Early Release of Estimates From the National Health Interview Survey, January–March 2018.”

Newman, P. A., C. H. Logie, A. Lacombe-Duncan, P. Baiden, S. Tepjan, C. Rubincam, N. Doukas, and F. Asey. 2018. “Parents’ uptake of human papillomavirus vaccines for their children: a systematic review and meta-analysis of observational studies.” *BMJ open*, 8(4): e019206.

Norris, L. 2018*a*. “Massachusetts health insurance marketplace: history and news of the state’s exchange.” Accessed at "https://www.healthinsurance.org/massachusetts-state-health-insurance-exchange/" on July 2, 2019.

Norris, L. 2018*b*. “Massachusetts health insurance marketplace: history and news of the state’s exchange.” Accessed at "https://www.healthinsurance.org/massachusetts-state-health-insurance-exchange/" on July 25, 2019.

Nosal, K. 2012. “Estimating Switching Costs for Medicare Advantage Plans.”

Obama, B. 2015. “Using Behavioral Science Insights To Better Serve the American People.” Executive Order 13707.

OECD. 2015. “Health at a Glance 2015: OECD Indicators.” OECD Publishing, Paris.

OECD. 2017. "Health at a Glance 2015: OECD Indicators." OECD Publishing, Paris.

Orsini, J., and P. Tebaldi. 2016. "Regulated Age-Based Pricing in Health Insurance Exchanges."

Panhans, M. 2018. "Adverse Selection in ACA Exchange Markets: Evidence from Colorado." Forthcoming American Economic Journal: Applied.

Patient CARE Act. 2014. "The Patient Choice, Affordability, Responsibility, and Empowerment Act: A Legislative Proposal." Accessed at https://www.ok.gov/health2/documents/The%20Patient%20CARE%20Act_Coburn.pdf on Apr. 18, 2018.

Patient Freedom Act. 2017. Accessed at https://www.cassidy.senate.gov/imo/media/doc/2.14%20One%20Pager.pdf on Oct. 4, 2018.

Pauly, M. V. 1968. "The Economics of Moral Hazard: Comment." *The American Economic Review*, 58(3): 531–537.

Petrosky, E., J.A. Bocchini Jr, S. Hariri, H. Chesson, C. R. Curtis, M. Saraiya, E. R. Unger, and L.E. Markowitz. 2015. "Use of 9-valent human papillomavirus (HPV) vaccine: updated HPV vaccination recommendations of the advisory committee on immunization practices." *MMWR. Morbidity and mortality weekly report*, 64(11): 300.

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

Polyakova, M. 2016*b*. "Risk selection and heterogeneous preferences in health insurance markets with a public option." *Journal of Health Economics*, 49: 153 – 168.

Rau, J. 2014. "The 10 Most Expensive Insurance Markets In The U.S." *Kaiser Health News*. Accessed at https://khn.org/news/most-expensive-insurance-markets-obamacare/ on May 5, 2018.

Reagan-Steiner, S., D. Yankey, J. Jeyarajah, L. D. Elam-Evans, J. A. Singleton, C. R. Curtis, L. E. MacNeil, J.and Markowitz, and S. Stokley. 2015. "National, regional, state, and selected local area vaccination coverage among adolescents aged 13–17 years—United States, 2014." *MMWR. Morbidity and mortality weekly report*, 64(29): 784.

Rothschild, M., and J. Stiglitz. 1976. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *The Quarterly Journal of Economics*, 90(4): 629 – 649.

Ruggles, S., K. Genadek, R. Goeken, Grover, and Sobek. 2017. "Integrated Public Use Microdata Series: Version 7.0 [dataset]. Minneapolis: University of Minnesota." https://doi.org/10.18128/D010.V7.0.

Schoen, C., M. M. Doty, S. R. Collins, and A. L. Holmgren. 2005. "Insured But Not Protected: How Many Adults Are Underinsured?" *Health Affairs*, 24(Suppl1): 289–302.

Senkomago, V., S. J. Henley, C. C. Thomas, J. M. Mix, L. E. Markowitz, and M. Saraiya. 2019. "Human papillomavirus–attributable cancers—United States, 2012–2016." *Morbidity and Mortality Weekly Report*, 68(33): 724.

Shepard, M. 2016. "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange." NBER 22600.

Sommers, B. D., T. Buchmueller, S. L. Decker, C. Carey, and R. Kronick. 2013. "The Affordable Care Act has led to significant gains in health insurance and access to care for young adults." *Health affairs*, 32(1): 165–174.

Spencer, J. C., N. T. Brewer, J. G. Trogdon, S. B. Wheeler, and S. B. Dusetzina. 2018. "Predictors of human papillomavirus vaccine follow-through among privately insured US patients." *American journal of public health*, 108(7): 946–950.

Tebaldi, P. 2017. "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA."

Terlizzi, E. P., R. A. Cohen, and M. E. Martinez. 2018. "Health insurance coverage: early release of estimates from the National Health Interview Survey, January–September 2018." *National Center for Health Statistics National Health Interview Survey Early Release Program.*

Thaler, R. H., and C. R Sunstein. 2009. *Nudge: Improving Decisions About Health, Wealth, and Happiness.* Penguin Books.

Tibshirani, R. 1996. "Regression Shrinkage and Selection via the LASSO." *Journal of the Royal Statistical Society*, 58 (1): 267–288.

Train, K. 2009. *Discrete Choice Methods with Simulation.* Cambridge University Press.

Tricco, A. C., L. Rambout, L. C. Hopkins, and M. Tashkandi. 2013. "Self-reported barriers and facilitators to preventive human papillomavirus vaccination among adolescent girls and young women: a systematic review." *Preventive Medicine.*

US Department of Health and Human Services. 2020. "Topics & objectives: Immunization and infectious diseases." Available at: https://www.healthypeople.gov/2020/topics-objectives/topic/immunization-and-infectious-diseases/objectives. Accessed July 17, 2019.

US Government. 2010. "Patient Protection and Affordable Care Act, Pub. Law 111-148, as amended by the Health Care and Education Reconciliation Act (HCERA), Pub. Law 111–152."

Walker, T.Y., L.D. Elam-Evans, J.A. Singleton, D. Yankey, L.E. Markowitz, B. Fredua, C.L. Williams, S.A. Meyer, and S. Stokley. 2018.

"National, regional, state, and selected local area vaccination coverage among adolescents aged 13–17 years—United States, 2017." *MMWR Morb Mortal Wkly Rep*, 67: 909–917.

Wharam, J. F.and Zhang, F., C. Lu, A. K. Wagner, L. Nekhlyudov, C. Earle, S. Stephen, and D. Ross-Degnan. 2018. "Impact of high-deductible insurance on breast cancer care among lower-income women." *Journal of Clinical Oncology*, 36: 6560–6560.

Wisk, L. E., and N. Sharma. 2019. "Inequalities in young adult health insurance coverage post-federal health reform." *Journal of general internal medicine*, 34(1): 65–74.

Appendix A

Appendix

Figure A.1: Average Annual Deductible, by Metal Tier, CSR, and Employer Plans

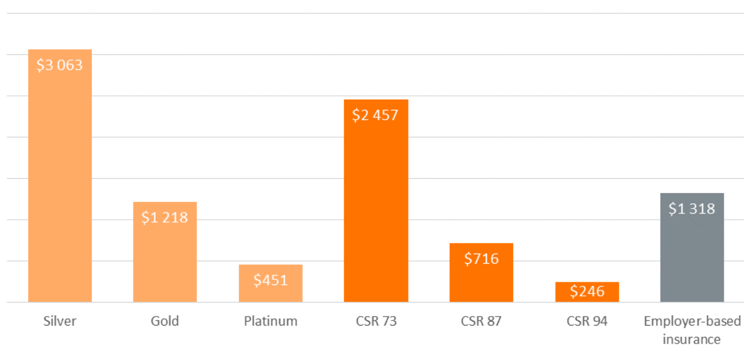

Notes: Base silver plans have an actuarial value of about 0.7, meaning an average of 70 percent of costs are covered; CSR 73, CSR 87, and CSR 94 silver plans have actuarial values of 0.73, 0.87, and 0.94, respectively. The most recent employer-based insurance survey data are from 2015. Sources: The Commonwealth Fund. 2016. The ACA's Cost-Sharing Reduction Plans: A Key to Affordable Health Coverage for Millions of U.S. Workers. Qualified Health Plan Landscape Files for federally facilitated marketplaces; state insurance websites and state marketplace websites for state-based marketplaces.

Figure A.2: Expected Premium Contributions By Percentage of FPL

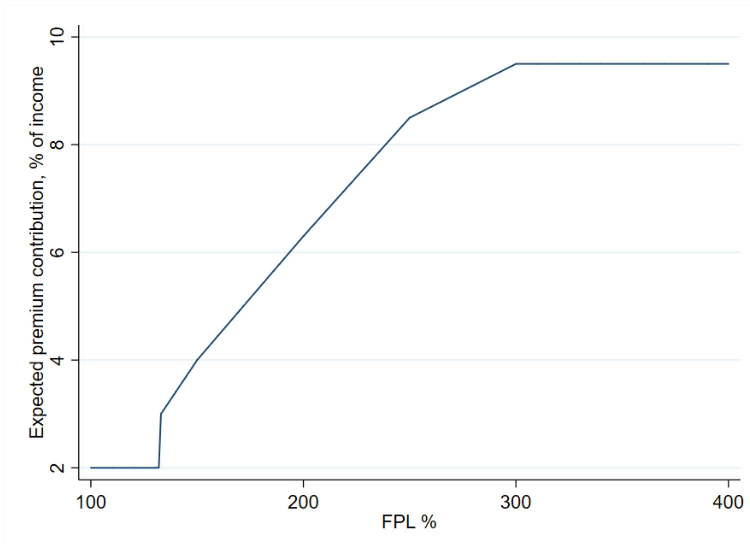

Notes: The graph shows the amounts enrollees are expected to contribute to the premium of the ACA marketplace plans as a function of their income. To determine the amount of premium subsidy, the ACA imposes a maximum amount that enrollees are expected to contribute out-of-pocket to the premium of the benchmark plan of the rating area. Source: IRS: Internal Revenue Code, Section 36B Refundable credit for coverage under a qualified health plan.

Figure A.3: Discontinuity in Annual Spending at 150% FPL, Restricted Sample

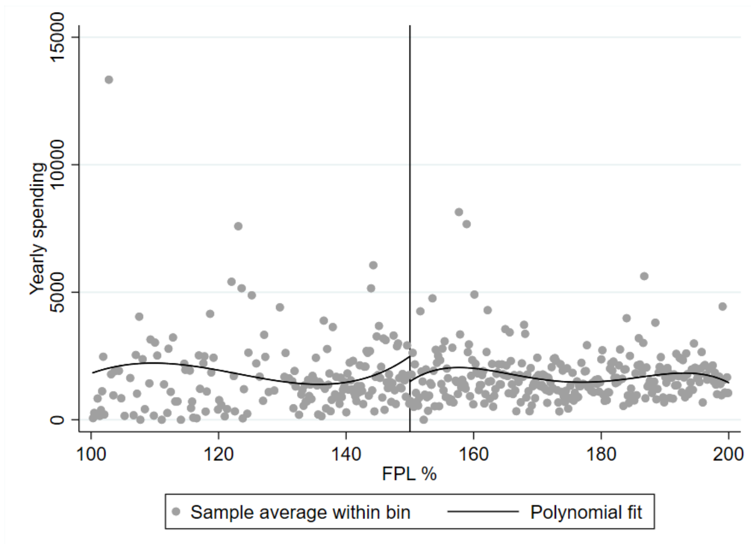

Notes: The figure displays the mean annual health care spending of silver plan enrollees by imputed income expressed as % of the federal poverty level and a fitted local polynomial regression around the discontinuity at 150%. The model is estimated on a restricted sample with income between 100% and 200% of the FPL.

Figure A.4: Discontinuity in Annual Spending at 200% FPL, Restricted Sample

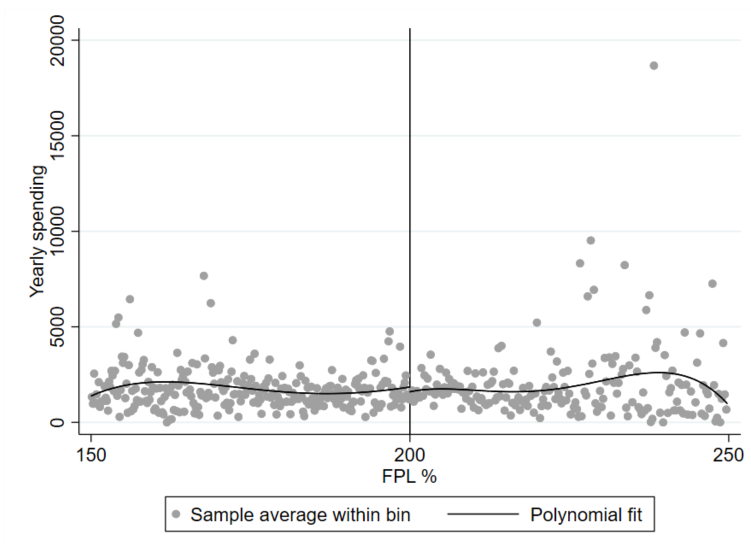

Notes: The figure displays the mean annual health care spending of silver plan enrollees by imputed income expressed as % of the federal poverty level and a fitted local polynomial regression around the discontinuity at 200%. The model is estimated on a restricted sample with income between 150% and 250% of the FPL.

Figure A.5: Discontinuity in Annual Spending at 250% FPL, Restricted Sample

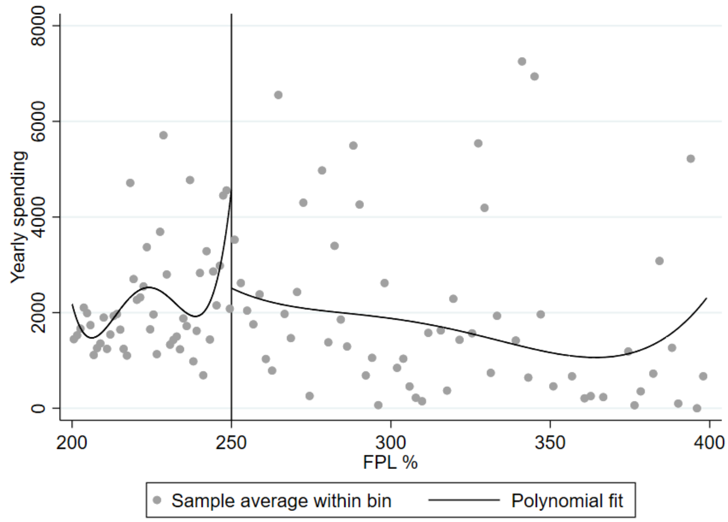

Notes: The figure displays the mean annual health care spending of silver plan enrollees by imputed income expressed as % of the federal poverty level and a fitted local polynomial regression around the discontinuity at 250%. The model is estimated on a restricted sample with income between 200% and 400% of the FPL.

Figure A.6: Placebo Checks at 125% FPL

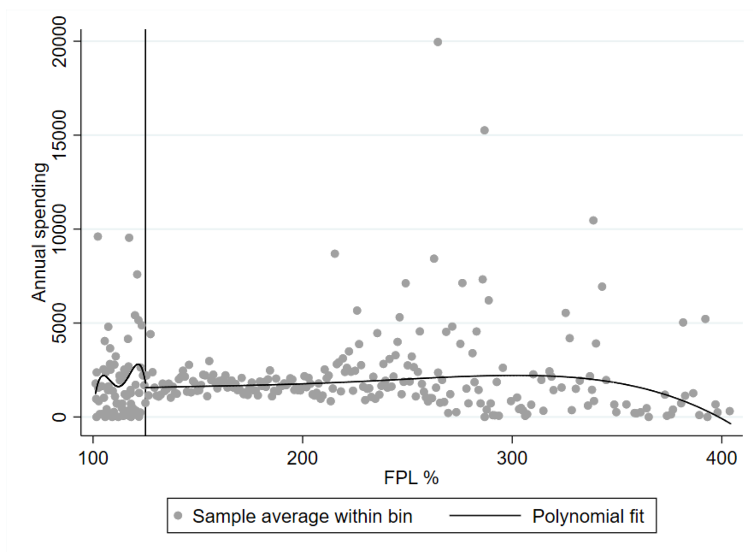

Notes: The figure displays the mean annual health care spending of silver plan enrollees by imputed income expressed as % of the federal poverty level. The model was estimated nonparametrically by fitting local linear regressions around 125% of the FPL, and the running variable was imputed income. A triangular kernel function was used to construct the local polynomial estimator and the optimal bandwidth was chosen by a data-driven selecting algorithm. The ACA's CSR subsidies induce discontinuous increases in the actuarial values of silver plans at 150%, 200% and 250% of the federal poverty level.

Figure A.7: Placebo Checks at 175% FPL

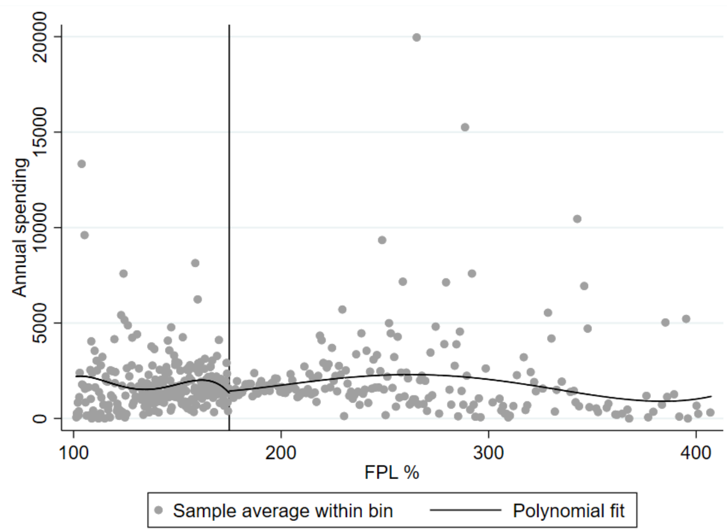

Notes: The figure displays the mean annual health care spending of silver plan enrollees by imputed income expressed as % of the federal poverty level. The model was estimated nonparametrically by fitting local linear regressions around 175% of the FPL, and the running variable was imputed income. A triangular kernel function was used to construct the local polynomial estimator and the optimal bandwidth was chosen by a data-driven selecting algorithm. The ACA's CSR subsidies induce discontinuous increases in the actuarial values of silver plans at 150%, 200% and 250% of the federal poverty level.

Figure A.8: Placebo Checks at 225% FPL

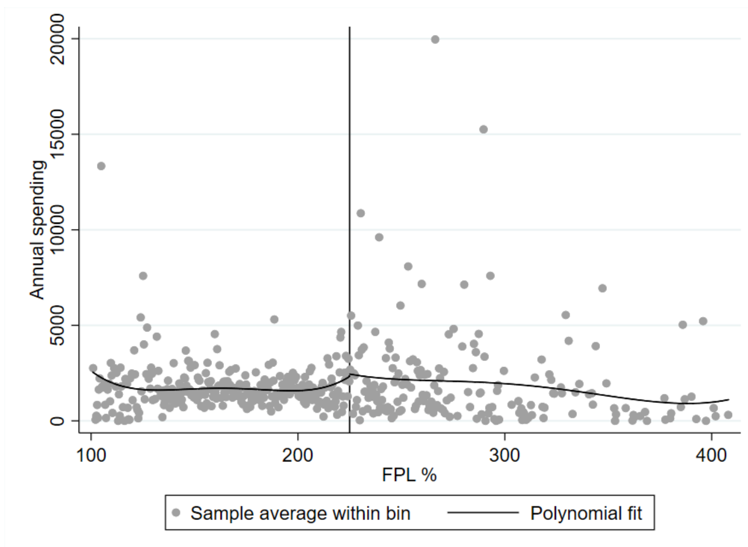

Notes: The figure displays the mean annual health care spending of silver plan enrollees by imputed income expressed as % of the federal poverty level. The model was estimated nonparametrically by fitting local linear regressions around 225% of the FPL, and the running variable was imputed income. A triangular kernel function was used to construct the local polynomial estimator and the optimal bandwidth was chosen by a data-driven selecting algorithm. The ACA's CSR subsidies induce discontinuous increases in the actuarial values of silver plans at 150%, 200% and 250% of the federal poverty level.

Table A.1: Health Care Spending by Market Category

	Spending			
	2013	**2014**	**2015**	**2016**
Employer-sponsored	3 418	3 259	3 163	3 586
Individual	2 910	3 395	3 423	3 842
Individual - pre ACA		3 202	3 420	3 467
Individual - Exchange		4 112	3 585	4 970

Notes: The table displays the average yearly health care spending of enrollees in the individual market and the employer-sponsored market by year. "Individual - pre ACA" denotes enrollees who had individual market health insurance coverage in 2013, prior to the ACA's rating reforms. "Individual - Exchange" shows the spending of enrollees who entered to the Marketplaces starting from 2014 with no prior individual coverage.

Table A.2: Yearly Health Care Utilization by Market and Prior Insurance Status

	(1)	(2)	(3)
	Annual spending	**$\mathbb{1}$[Claims>0]**	**No. of inpatient claims**
Group	Reference		
Enrolled_2013	-31.453**	-0.085***	-0.007***
	(12.708)	(0.002)	(0.001)
Enrolled_2014	854.794***	0.393***	0.004***
	(44.958)	(0.004)	(0.001)
N	3,110,651	3,110,651	3,110,651

Notes: The table shows different health care utilization measures regressed on health insurance status and a set of time fixed effects. The reference category includes individuals with employer-sponsored health plans. The dependent variable in column (1) is the average yearly health care spending of enrollees in the given market. The dependent variable in column (2) is an indicator for the enrollee submitting any claim during the year. Enrolled_2013 is an indicator whether the individual had any individual market health insurance coverage in 2013 prior to the ACA's rating reforms. Enrolled_2014 is an indicator for enrollees who entered to the Marketplaces starting from 2014 with no prior health individual insurance coverage. Standard errors in parenthesis; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.3: Demographics by Enrollment Cohort

	New enrollees			*Returning enrollees*
	2014	2015	2016	2016
Mean age	44.47	43.55	44.98	44.82
Female %	0.58	0.56	0.58	0.57
Mean income	68 353	68 097	67 802	68 369
Mean yearly spending	3 878	3 727	4 065	3 844

Notes: The first three columns compare the demographic characteristics and health care spending of those who first entered to the Marketplaces in 2014, 2015 and 2016, respectively. Column (3) excludes individuals who enrolled in 2014, had no Marketplace coverage in 2015 and returned in 2016. These returning enrollees are shown in column (4). Income is measured at the zip code level.

Table A.4: Expected Premium Contributions By Percentage of FPL

Household income percentage of FPL	Initial premium %		Final premium %	
	2014	2015	2014	2015
Up to 133%	2.00%	2.01%	2.00%	2.01%
133% up to 150%	3.00%	3.02%	4.00%	4.02%
150% up to 200%	4.00%	4.02%	6.30%	6.34%
200% up to 250%	6.30%	6.34%	8.05%	8.10%
250% up to 300%	8.05%	8.10%	9.50%	9.56%
300% up to 400%	9.50%	9.56%	9.5%.	9.56%

Notes: The table shows the amounts enrollees are expected to contribute to the premium of the ACA marketplace plans as a function of their income. To determine the amount of premium subsidy, the ACA imposes a maximum amount that enrollees are expected to contribute out-of-pocket to the premium of the benchmark plan of the rating area. Between the values shown in table, the expected contribution increases linearly. Source: IRS: Internal Revenue Code, Section 36B Refundable credit for coverage under a qualified health plan (2014-2015).

Table A.5: Descriptive statistics for HPV vaccine completion defined as completion of the 3-dose series within 18 months

	Females				**Males**			
	N	%	% Comp.	Adjusted OR	N	%	% Comp.	Adjusted OR
Age of first dose								
9-14	96702	49.7	37.2	1	72071	38.2	37.3	1
15-18	56407	29.0	31.6	0.74 (0.73-0.76)	87336	46.2	34.9	0.82 (0.80-0.84)
19-26	41298	21.2	22.9	0.48 (0.46-0.49)	29483	15.6	25.4	0.51 (0.49-0.52)
Insurance type								
Medicaid	28827	14.8	22.1	1	22588	12.0	23.4	1
Private	165580	85.2	34.3	2.58 (2.09-3.18)	166302	88.0	35.8	1.79 (1.73-1.86)
Year of first dose								
2009	3565	1.8	23.8	1	-	-	-	-
2010	11397	5.9	31.5	1.49 (1.36-1.63)	-	-	-	-
2011	28463	14.6	30.7	0.99 (0.88-1.13)	6484	3.4	38.3	1
2012	59371	30.5	37.8	1.24 (1.10-1.40)	74231	39.3	41.8	1.09 (1.03-1.15)
2013	64134	33.0	27.2	0.69 (0.61-0.78)	78172	41.4	27.0	0.52 (0.49-0.55)
2014	27477	14.1	36.8	0.84 (0.74-0.96)	30003	15.9	34.2	0.59 (0.55-0.64)
State								
Massachusetts	136351	70.1	34.9	1	151877	80.4	36.5	1
Maine	24761	12.7	27.3	0.66 (0.64-0.69)	15160	8.0	22.6	0.50 (0.48-0.52)
New Hampshire	33295	17.1	26.4	0.61 (0.59-0.63)	21853	11.6	27.5	0.63 (0.61-0.65)

Notes: Abbreviations: CI, confidence interval; HPV, human papillomavirus

Table A.6: Descriptive statistics for HPV vaccine completion defined as completion of the 2-dose series within 12 months

	Females				Males			
	N	%	% Comp.	Adjusted OR	N	%	% Comp.	Adjusted OR
Age of first dose								
9-14	96702	49.7	49.8	1	72071	38.2	48.2	1
15-18	56407	29.0	45.5	0.82 (0.80-0.83)	87336	46.2	47.0	0.91 (0.89-0.93)
19-26	41298	21.2	42.3	0.74 (0.72-0.75)	29483	15.6	40.9	0.71 (0.69-0.73)
Insurance type								
Medicaid	28827	14.8	42.3	1	22588	12.0	38.0	1
Private	165580	85.2	48.6	2.17 (1.84-2.56)	166302	88.0	47.7	1.79 (1.73-1.86)
Year of first dose								
2009	3565	1.8	40.1	1	-	-		-
2010	11397	5.9	48.4	1.42 (1.31-1.54)	-	-		-
2011	28463	14.6	47.6	1.20 (1.07-1.34)	6484	3.4	47.5	1
2012	59371	30.5	50.3	1.30 (1.17-1.46)	74231	39.3	52.1	1.18 (1.12-1.24)
2013	64134	33.0	40.8	0.84 (0.76-0.94)	78172	41.4	39.6	0.69 (0.65-0.72)
2014	27477	14.1	53.6	1.51 (1.34-1.71)	30003	15.9	50.5	1.11 (1.03-1.19)
State								
Massachusetts	136351	70.1	47.3	1	151877	80.4	46.9	1
Maine	24761	12.7	48.3	1.08 (1.05-1.11)	15160	8.0	42.9	0.89 (0.86-0.92)
New Hampshire	33295	17.1	44.7	0.88 (0.86-0.91)	21853	11.6	46.6	0.97 (0.95-1.00)

Notes: Abbreviations: CI, confidence interval; HPV, human papillomavirus