Competition in Employer Sponsored Health Insurance: Implications for a Public Option

Stuart V. Craig*

November 1, 2024

Click here for most recent version

Abstract

In response to concerns about market power in the private insurance industry, many have proposed a government run insurance plan, or "public option," which would operate within existing insurance markets. I explore the conditions under which a public option attracts employer groups or generates additional premium competition. Using unique data on employer-insurer contracting, I document evidence of large switching costs and variation in expected costs across employers. Insurers respond to switching costs by setting low premiums for new employers and raising them for those who are locked-in. I estimate a model of supply and demand in this market, and use it to evaluate the impact of many potential public plans. Naturally, the public option is most attractive when it has higher quality and lower cost. Uniform premium setting (i.e. community rating) substantially limits the public option's ability to compete with private insurers who set employer-specific premiums. A public option that offers experience rated premiums and discounts to new employers is substantially more attractive and therefore has a larger impact on premium competition. A community rated public option can lower premiums in the market with a large enough cost advantage, but it does so by crowding out existing insurers rather than generating additional price competition.

*Wisconsin School of Business, University of Wisconsin–Madison (email: svcraig@wisc.edu). I am grateful for the support and guidance of many in improving this manuscript, chiefly my dissertation committee — Martin Gaynor, Matthew Grennan, Aviv Nevo, and Mark Pauly — as well as many other colleagues and friends. Zarek Brot, Benjamin Chartock, Zack Cooper, Uli Doraszelski, Hanming Fang, Atul Gupta, João Granja, Phil Haile, Jonathan Holmes, Tomás Larroucau, Claudio Lucarelli, Joseph Martinez, Steve Schwab, Mark Shepard, Amanda Starc, Ashley Swanson, and Kevin Williams have all provided valuable support and feedback. I would also like to thank Lynda Taylor at Wrangle 5500 for her generous guidance on handling the data, and several industry participants who provided helpful insight on institutional context. All errors are, of course, my own.

1 Introduction

Health care reform has become a central issue in US policy debate due to concerns about access and affordability of health insurance. The vast majority of non-elderly Americans — 156 million individuals — obtain health insurance through the employer-sponsored insurance market, where premiums account for the bulk of consumer health expenditure and are substantial. Average premiums for single coverage reached $7,470 in 2020 (Claxton et al. 2020). A growing body of evidence demonstrates that market concentration in insurance markets leads to higher premiums (Dafny 2015, Dafny et al. 2012, Guardado et al. 2013), which are passed on to workers in the form of lower wages (Baicker and Chandra 2006, Arnold and Whaley 2020). Existing reform proposals range from small regulatory changes in existing markets to Medicare for All, in which the government would step in to act as a universal insurer. A public option represents an intermediate reform that would create a government run or financed insurer to operate within existing markets. Such a policy could potentially lower health insurance premiums by regulating the prices of health care services, setting premiums without a markup, reducing administrative overhead, and thus cause insurers to compete more aggressively (Hacker 2008).

This paper examines the impact a public option would have on the existing private insurance market. I study the employer-sponsored insurance market to identify the conditions under which a public option would generate enrollment or provide additional premium competition if offered to employers. The impact of a public option depends on how attractive it is to consumers. Prior work evaluating and forecasting the public option focus on issues of cost and quality (Fiedler 2020, Blumberg et al. 2019), with much attention focused on the degree to which the public option would exercise a strong cost advantage by regulating payments to health care providers. However, the ability of the public option to attract consumers, and therefore generate additional competition, also depends on its ability to deal with common frictions in insurance markets.

I focus, in particular, on the interaction between premium setting and adverse selection.

Whereas large employers obtain employer-specific premium offers, publicly provided health insurance plans tend to favor uniform, or "community rated," premiums. Hoffman et al. (2021) note that adjusting prices across employers groups "cuts squarely against coverage and equity goals" of a public plan. At the same time, a public option that sets uniform premiums has the potential to attract only the most expensive employers, who obtain high risk-rated premium offers in the private market. If the public option has to maintain a break-even constraint, it will tend to unravel (Akerlof 1978). The degree to which this is likely to occur depends critically on markups in the existing market and the differences in costs of providing insurance to different employers. In the absence of an existing policy allowing employers to purchase insurance through a public option, I estimate a structural model of employer-insurer contracting and simulate the market-wide effects under varying assumptions about the cost, quality, and premium setting rules of a government-run insurance plan.

Prior work on insurer competition in this setting is limited due to the lack of quality data on prices and quantities (Dafny et al. 2011, Dafny 2015). Many existing studies focus on competition between insurers within large firms who offer multiple plans to their employees (Cutler and Gruber 1996, Ho and Lee 2017, Tilipman 2018), or aggregate data on premiums and market shares (Dranove et al. 2003). In a study examining choice frictions and resulting price dispersion, Cebul et al. (2011) use multiple sources of survey data. These data provide information on cancellation rates and premiums, but are limited in their ability accommodate markups that result from insurers offering differentiated products (e.g. networks). Finally, using data from a commercial brokerage service, Dafny et al. 2012 and Dafny 2010 demonstrate the existence of market power and price discrimination in the employer market. However, these data capture a select subset of large national employers, preventing them from characterizing demand and competition in any local market.

I leverage a unique source of panel data containing insurer choices and premiums, representing the vast majority of the large group market in Massachusetts from 2001-2018, where three insurers — Blue Cross and Blue Shield of Massachusetts (BCBS), Harvard

Pilgrim Health Care (HPHC), and Tufts — serve over 80 percent of the market. I find that premiums vary considerably across employers. Moreover, switching between health insurers is rare despite potential savings revealed by the switches that do occur — at least in the short run. Analyzing employers' choices, I find evidence of large switching costs, and that insurers price discriminate between old and new contracts by offering low premiums to "invest" in acquiring new employers, and raising premiums to "harvest" rents from those who are locked in.

I then use the premiums and choices in these data to estimate a model that characterizes premiums as a function of (1) employer preferences over health insurers, (2) costs that reflect their underlying health risk or taste for health care consumption, and insurer's relative advantages in negotiating favorable contracts with up-stream health care providers, and (3) insurers' strategic response to switching costs.

I estimate the model in two steps. First, I estimate demand to recover employers' preferences for insurers and their switching costs. In order to capture the variation in preferences that would lead to variation in markups and credibly separate brand preferences from switching costs, I implement a mixed-logit specification that allows for a flexible distribution of preferences (Train 2016). I find that the average employer is willing to pay $489 (per person per year in 2018 dollars) for their most preferred insurer. This result adds to findings from Dranove et al. (2003) — who show that national and local insurers tend to serve distinct segments of the market — by showing that employer groups have strong preferences, even in a market dominated by three local insurers. I also find that switching costs are substantial, averaging $3,020 per person per year, which is consistent with prior work measuring switching costs for individual plan choices (Tilipman 2018). Second, on the supply-side, I specify a dynamic premium setting game, in which forward-looking insurers compete by setting employer-specific premiums. The model both rationalizes the invest-harvest premium setting observed in the data, and demonstrates the potential for dynamic incentives to create downward pressure on premiums despite highly inelastic demand. I leverage this model to

recover a distribution of expected costs across employer-groups, which differ considerably in their taste for health care consumption. I also find that HPHC and Tufts each have a moderate cost advantage relative to BCBS.

I then use the model to simulate the impact of allowing employers access to a public option. To accommodate the policy uncertainty over what a public option in this market would look like, I recompute counterfactual equilibria under a range of assumptions about the quality of a public plan, and the degree to which it exercises a strong cost advantage. I find that enrollment and competitive effects are larger when the public option enjoys a large cost advantage and provides levels of quality commensurate with existing private insurers. I also find that a public option that replicates the cost and quality of existing insurers, but cannot price discriminate, has only a modest impact on the market due to strong brand preference, switching costs, and adverse selection into the public option. In most cases, only the most expensive employers choose to opt in to a public option, resulting in high premiums that further limit demand. This finding is consistent with prior work by Shepard et al. (2020), who show that a uniform public health plan is unlikely to serve the specific needs of individuals with varying tastes for health care consumption. A public option that exercises a cost advantage by regulating provider payments is able to lower the unit cost of health care consumption, increasing the number of employers who find it optimal to abandon the private market. However, employers who remain in the private market do so because they have expected costs well below the premium offer of the public option, and therefore experience limited additional premium competition. Naturally, the effect of any policy is smaller in cases where the public option also has a quality shortfall relative to existing insurers.

I test counterfactual pricing policies to examine the effectiveness of a public option that set premiums more flexibly. A public option that sets experience rated premiums or provides introductory discounts to new employers is more attractive to all employers, increasing enrollment substantially. Both of these premium setting policies provide more

pressure for private insurers to offer lower premiums. In particular, a public option that offers experience rated premiums generates more premium competition than a community rated public option with a 50 percent cost advantage.

While providing insight into the potential impacts of a public option, this paper makes two additional contributions to existing literature. First, I offer new evidence on price discrimination and competition in the employer-sponsored health insurance market. Thus far, prior work has focused on the relationship between market structure and premiums (Dafny et al. 2012, Trish and Herring 2015), estimated demand for insurers using individuals' health plan selections within a single employer(Cutler and Gruber 1996, Gruber and McKnight 2016, Ho and Lee 2017, Tilipman 2018), or inferred measures of market power using aggregate data on premiums and markups (Dranove et al. (2003)). Most recently, Lin and Zhou (2021) make use of the same data I use in this paper to study the relationship between insurer market structure and premiums. Dafny (2010) uses panel data on employers across many markets to demonstrate the existence of first-degree price discrimination, and Cebul et al. (2011) use survey data to study the impact of search frictions. I add to this literature by separately estimating the components of demand that are driven by employers' heterogeneous preferences over insurers and the costs employers face in switching between insurers. My modeling approach incorporates employer-specific pricing that allows for discrimination on dimensions related to employer-specific expected costs, preferences over the differentiated products offered by insurers, and switching costs. To my knowledge, this is the first study to estimate demand for insurers using data from multiple employer-groups. Whereas prior work often focuses on employees' demand for insurance within a choice set defined by the employer, my model allows me to characterize insurer incentives and provide insight into competition and premium setting at the stage where employers select a health insurer. I provide new evidence that employers have strong preferences over insurers as well as large switching costs, and show that switching costs play a complex role in premium setting. Moreover, I develop a procedure to recover measures of consumer preferences and market

power that can be extended to answer a range of policy questions in any state market, and requires only publicly available data.

Second, I contribute to a substantial literature on inertia and competition in health insurance markets.[1] In Medicare Part D, inertia has been attributed to inattention, and multiple studies have found that this leads to higher premiums (Ericson 2014, Ho et al. 2017). In employer-sponsored settings, inertia has been attributed to information and "hassle costs" (Handel and Kolstad 2015), and the cost of switching doctors under a new network (Dahl and Forbes 2014, Tilipman 2018). I show that the magnitude of switching costs are similar for employer-groups selecting health insurers, providing evidence that these inertial forces are present even when the agent is a plausibly more sophisticated purchaser. Moreover, theoretical predictions regarding the equilibrium effects of switching costs are mixed and depend on the balance of selling firms' dynamic incentives.[2] Indeed, empirical work in other settings has identified markets in which switching costs generate lower prices in equilibrium (Dubé et al. 2009). I contribute to this literature by documenting the impact of switching costs on premiums in the employer-sponsored insurance market, while addressing the potential for dynamic incentives to restrain market power.

The remainder of this paper is structured as follows. In Section 2, I describe the relevant institutional features of the employer sponsored market and the public option. I also describe a simple model of price discrimination to provide intuition for the selection issues that govern the extent to which a public option might cause insurers to compete more aggressively. Section 3 introduces the data on employer-insurer contracting and describes the variation in insurance premiums. In Section 4, I present a model of employer-insurer contracting that accounts for employer-specific premium setting, and the dynamic incentives of insurers. I describe the method I use to estimate the model in Section 5, and report the

[1] Employer-sponsored insurance (Handel 2013, Handel and Kolstad 2015, Dahl and Forbes 2014, Tilipman 2018), Medicare Part D (Abaluck and Gruber 2011, Ketcham and Powers 2016, Ericson 2014, Ho et al. 2017, Polyakova 2016), and the ACA exchanges (Drake et al. 2020).

[2] Arie and Grieco 2014, Cabral (2016), Farrell and Klemperer (2007), Farrell and Shapiro (1988), Klemperer (1995)

results of the estimation procedure in Section 6. In Section 7, I use the model to simulate public option policies under various assumptions about cost and quality. Section 8 concludes.

2 Setting and Motivation

The public option is distinct from proposals for more dramatic policies changes — like Medicare for All — in that it would allow for consumers to continue purchasing private coverage if they preferred to do so. For those who are currently uninsured — particularly for those without access to coverage through an employer — the public option is expected to increase voluntary coverage by providing an insurance option at a lower premium than existing public plans. For this reason, the public option would be incorporated into the exchanges established by the ACA under all existing proposals. For those already insured, the public option is expected to generate additional competition among the existing insurers in the market, which many view as charging exorbitant premiums due to consolidation and other sources of market power.

At least eight different bills were introduced in the 116^{th} Congress, all with different design features (Kaiser Family Foundation, 2019). Two of these bills – the Choose America Act (S. 1261/H.R.2463, hereafter CMA) and the Medicare for America Act of 2019 (H.R. 2452, hereafter MAA) – would provide large employers with access to a publicly run insurance plan.

Provider Reimbursements and Network Inclusion In addition to running a public plan that forgoes a markup, the CMA, MAA, and most other proposals envision a public option that generates cost savings for consumers by reducing the reimbursements paid to health care providers. Typically these proposals assume rates will be tied to reimbursements currently set by Medicare or Medicaid, which are substantially lower than those negotiated by private insurers (Cooper et al. 2019, Clemens and Gottlieb 2017). In some cases, policy proposals or bills suggest Medicare or Medicaid rates as a starting point or minimum payment

rate. MAA would set provider payment rates using the larger of current Medicare or Medicaid rates. The CMA specifies that the public option will set reimbursement rates no lower than existing Medicare and no higher than existing private insurance plans.

Obtaining low reimbursement rates may be possible if the public option ties its network to Medicare and Medicaid participation. Both the CMA and MAA would require Medicare and Medicaid participating providers to participate in the public option network. The majority of health care providers currently accept Medicare patients. But there is no guarantee they would continue to do so if participation in a large public option were tied to Medicare participation – particularly if doing so allowed the public option to out-compete private insurers with higher reimbursement rates.

The welfare consequences of regulating lower provider prices are ambiguous, and depend on the ex ante efficiency of prices currently negotiated in the private market. Prior work has highlighted market power and moral hazard as important drivers of health care prices (Cooper et al. 2019, Gowrisankaran et al. 2015). However, recent studies have also shown that these prices create important dynamic incentives for providers. Garthwaite et al. (2020) find that hospitals with access to larger private markets do more to invest in quality improvement and technology. Clemens et al. (2020) show that privately negotiated prices also create incentives for physicians to invest in expanded capacity. Moreover, a public option that aggressively regulates prices may find it difficult to replicate the network breadth and quality of existing private plans.

Both the overall welfare impact of regulating these payments, and the feasibility of maintaining a broad provider network while providing substantially lower reimbursement rates, are beyond the scope of this paper. In order to address uncertainty around this issue, in Section 7, I simulate the impact of a public option with varying degrees of cost savings and quality differentials, while noting that not all of these possibilities may be feasible.

Premium Setting and Adverse Selection Neither the CMA or the MAA would tailor premiums to specific employers. In a proposal aimed at a public option for employers, Hoffman et al. (2021) suggests requiring employers to opt in or out of the public option completely, to avoid a case where only the least healthy employees select into the public plan. This insight builds on a robust literature demonstrating the potential for adverse selection *within* employer groups (see for example, Cutler and Reber 1998, Einav et al. 2010, Handel 2013). However, variation in expected cost *across* employers plays a key role in the model I use to study the public option. I sketch the main intuition here and provide a more formal theoretical model in Appendix A.

Because private insurers can set employer-specific premiums, employers who prefer less health care consumption will always have the option of obtaining an experience rated plan from a private insurer. In this context, a public option setting premiums at average cost will likely unravel until it serves only the most costly employers. The degree to which this occurs will depend on how cost effective the public option is, relative to existing insurers, and the variance in employers' expected costs. An adversely selected public option also has implications for its ability to provide price competition to employers who continue to purchase from private insurers. Employers with sufficiently low expected costs will always prefer an experience rated premium offer to a public option that is more expensive, and only employers that are near the indifference point between their private plan and the public option will receive lower premium offers. The magnitude of this barrier depends on the dispersion in cost across employers relative to the markups charged by private insurers, which are key objects of interest to be recovered by the model presented in Section 4.

3 Data

The primary data used in this paper come from the Department of Labor's (DoL) Form 5500, which is an administrative filing required of nearly all employers with at least 100 employees,

in compliance with the Employee Retirement Income Security Act of 1974 (ERISA).[3] The form was jointly developed by the DoL and the Internal Revenue Service (IRS) as a compliance and disclosure document for a range of benefits plans, including health plans. Filings are audited regularly and employers face financial penalties for incomplete or delinquent filings.

I focus on the market for health insurance in Massachusetts, for which data are available from 2001-2018. For insurance-based health arrangements (i.e. fully-insured health plans), the data include employer identifiers and addresses, as well as detailed information about the duration of the contract, identity of the insurer, the number of employees and enrollees, the type of plan (e.g. HMO or PPO), and total premiums paid for each plan-year. All premiums in this market are adjusted based on individual employers' expected utilization. However, roughly 28 percent of plans in the data are explicitly experience rated and therefore include a measure of total claims in the Form 5500. I focus on average premiums and claims per enrollee, which capture a measure of prices and costs that is invariant to fluctuations in firm size.[4] The data are subjected to a number of restrictions and cleaning rules, which are detailed in Appendix B. The final sample contains an unbalanced panel of 1,061 employers from 2001-2018. In 2018, these employers represent 258,133 covered lives — 24 percent of the total large group enrollment reported in regulatory filings reported to the Centers for Medicare and Medicaid Services (CMS).[5]

[3]ERISA is a federally enacted policy regulating health and benefit plans offered by employers. On many dimensions, fully-insured health plans are typically regulated by the state in which a business operates. The term "ERISA plans" is often used in industry and the literature to refer to employer plans that are self-insured by an employer, in which case ERISA preempts many state regulations. However, both fully and self-insured plans are subject to ERISA regulations in situations where there is no conflicting state law, and are therefore required to file the Form 5500 unless otherwise exempt from ERISA altogether. Examples of exempt employers include 403(b) non-profits, government plans, and plans maintained outside the US for which the majority of members are non-resident aliens.

[4]I observe enrollment counts at the beginning and end of each year. Total enrollees are calculated as the average of these two values.

[5]CMS regulates markups in this industry by instituting Medical Loss Ratios (MLR). Under ACA regulation, insurers in the large group must meet an MLR threshold of 85 percent. That is, insurers must spend at least 85 percent of total premiums collected on health care claims. However, these regulations are rarely binding (Cicala et al. 2019), and apply to the market segment as a whole rather than to any individual employer.

Table 1 summarizes the employers in my sample. The average employer is a medium sized firm with 274 employees and 461 enrollees, including dependents. Employers primarily offer a mix of HMO and PPO products. The average employer is observed for 7 years in the data.

Average premiums and claims in the data match closely to published estimates of aggregate premiums in the CMS filings. In Appendix B.3 I present multiple comparisons to external data sources to benchmark the premiums, claims, and market shares against CMS's Medical Loss Ratio filings and the Kaiser Family Foundation's Employer Health Benefits Survey (KFF-EHBS). For years in which MLR reporting is available (2011-onward), the series closely matches aggregate premiums for Massachusetts in the CMS data. In earlier years, published data specific to Massachusetts is unavailable, but average premiums in the Form 5500 are reasonably close to those in the KFF-EHBS. Although the KFF-EHBS is a national sample of employers, and therefore not a perfect comparison group, I find that the premiums in both data sources are similar in magnitude for years prior to 2011.

This level of detail data on employer-insurer contracting provides an unprecedented look at the functioning of this market. However, the data have several important limitations. First, the Form 5500 provides limited information on self-insured employers, who are therefore excluded from the sample. In Appendix C, I conduct multiple exercises to assess the magnitude of this issue, and its implications for my analysis. For the demand analysis supporting the main policy question, employers may have an option to self-insure. However, self-insured employers typically contract with a health insurer to provide Administrative Services Only (ASO). These services comprise the most likely sources of product differentiation in the fully insured market, including access to a network, claims administration, customer service, and utilization management. Self-insured employers face a similar market for ASO providers even if they choose to take on their own utilization risk. In Massachusetts the insurers providing administrative services are broadly the same as those who are active in the large group market.

Table 1: Characteristics of Employer Panel, 2001-2018

	Mean	Standard Deviation
Employer and Plan Characteristics		
Number of Employees	274	303
Number of Enrollees	461	520
Offers HMO	0.879	0.326
Offers PPO	0.651	0.476
Offers Indemnity	0.025	0.156
Average Premiums (2018)	6,249	1,337
Average Claims (2018)	5,411	1,408
Insurer Market Shares		
BCBS	0.628	
HPHC	0.112	
Tufts	0.132	
Other Insurer	0.128	

Notes: Table contains statistics on 1,051 employers, in the sample representing 7,660 employer-year observations. Premiums and claims are reported for 2018. Appendix figure A.1 plots premiums over time in comparison with other published data sources.

I also restrict my analysis to firms who contract with a single insurer at any given time. In some situations, the structure of the data make it difficult to distinguish the enrollment counts across multiple insurers within an employer. Additionally, focusing on the firms who contract with only one insurer simplifies the modeling in a discrete choice framework. According to data from the Kaiser Family Foundation's Employer Health Benefits Survey (KFF-EHBS), 75-80 percent of employers only offer a single plan to their employees, implying that an even greater share only contract with a single insurer. However, these employers tend to be smaller.

These sample limitations create potential internal and external validity concerns. Model estimates for smaller, fully insured employers may not generalize to the full distribution of firms. Moreover, because the counterfactuals focus on an equilibrium in which multiple employers may choose to purchase the same insurance product, these unobserved employers

could potentially change the equilibrium premium a public option would be able to offer. As noted in Section 2, the primary issues governing the degree of adverse selection – and therefore the impact a public option would have on premium competition – are the size of markups employers face, and the differences in cost between employers. In Appendix C, I use data from the KFF-EHBS to analyze premium variation across the distribution of firm size and between the fully and self-insured markets. I find that average premiums, as well as the *variation* in premiums across employers, are similar for large and medium sized employers and regardless of whether they self-insure.

Finally, the Form 5500 data are sparse in the sense that not all employers purchase from all insurers, and I only observe choices that were made. In a market with individually determined premiums, modeling employers' choices requires some assumptions about the premium offers an employer would obtain if they were to choose a different health insurer. These premium offers will be a function of the employers' underlying tastes for health care, and the health status of the employees. They will also be a function of the cost at which each insurer can provide those services through their network, and the degree to which employers prefer the differentiation of those services. I discuss these issues in more detail in Sections 4 and 5.

3.1 Premium Variation and Dynamic Insurer Incentives

A unique feature of the large group market, relative to the small group or individual markets, is that premiums are tailored specifically to each employer. Employers may differ in expected cost because of their health status or underlying taste for utilization, and insurers incorporate this by experience rating premiums. These premiums may also reflect differences in the extent to which certain employers prefer the insurer they currently contract with, relative to other available options (Dafny 2010). Moreover, the aggregate distribution of premiums may reflect differences in insurers' underlying cost structure stemming from their ability to negotiate more favorable contracts with upstream health care providers (Craig et al. 2021,

Panhans et al. 2018). Ultimately, the purpose of the model in Section 4 is to carefully identify the contributions of each of these channels. However, I describe the premium variation in this section.

I begin by decomposing premiums into employer and insurer components in the spirit of decomposition methods used to separate worker and firm wage effects (Abowd et al. 1999). I estimate

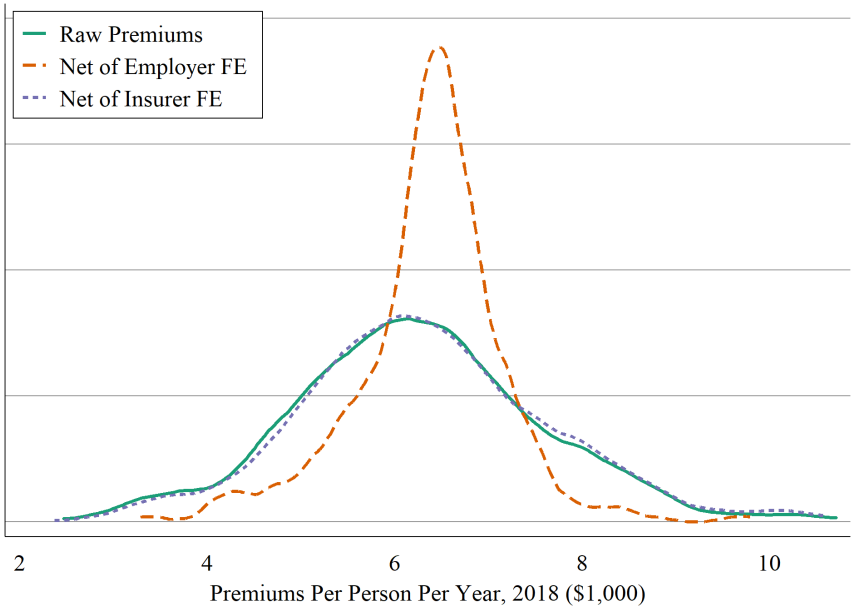
$$p_{ijt} = \gamma_i + \gamma_j + \gamma_t + \varepsilon_{ijt} \tag{1}$$

using OLS. γ_t is a year fixed-effect, which captures aggregate shocks to premiums, which are substantial over 2001-2018. γ_i separates the employer's contribution to premiums, that is common across insurers that it contracts with, while γ_j represents the average premium differences employers' experience when they switch between insurers.

Figure 1 displays the raw distribution of premiums in the data for 2018, as well as the distribution net of γ_i and γ_j respectively. Employer differences explain a large fraction of the variation. Insurer differences appear much smaller, but are economically significant. The table below shows average premiums across insurers, and estimates of γ_j from a version of equation (1) that is estimated in logs. The variation provided by employers who switch insurers reveals that, conditional on employer, Harvard Pilgrim and Tufts offer premiums that are approximately 5 percent lower than BCBS on average. While these differences appear small relative to the overall scale of premium variation, the average employer in the data spent 2.9 million dollars in 2018, meaning that 5 percent lower premiums would result in an average annual savings of 145 thousand dollars. These between-insurer differences could be due to differences in insurer cost structure, or to differences in employers' willingness to pay for Harvard Pilgrim and Tufts relative to BCBS among employers who switch. Nonetheless, these estimates provide evidence that employers remain with a more expensive insurer when potential savings are available.

The results in Figure 1 could also arise if the timing of the switches that identify γ_i in equation (1) are coincident with shocks to employers' underlying demand for health care

Figure 1: Variation in Premiums, 2018



		(1) Overall	(2) BCBS	(3) HPHC	(4) Tufts	(5) Other
Mean p_{ijt}	Mean	6,249	6,231	6,444	6,516	5,882
	(SD)	(1,337)	(1,313)	(1,248)	(1,075)	(1,677)
γ_j	Coefficient			-0.055**	-0.056**	-0.080
	(SE)			(0.021)	(0.018)	(0.032)

Notes: Figure presents distribution of raw premiums, premiums net of fixed-effects in equation (1) as described in the text. The table presents means and estimated premium differences across insurers. Mean p_{ijt} presents the raw average premiums. The second row presents estimates of γ_j from (1) using logged premiums as the dependent variable.

services. To investigate this issue, I estimate an event study regression focusing on the years surrounding the first observed switch for each employer. To ensure common support across all of the estimates, I restrict to the 171 employers who I observe for at least three years before and after the switch. In place of separate insurer and time fixed-effects, I now control for insurer-by-year shocks.[6] As a comparison group, I also include employers who never

[6]I use employer-time fixed effects to be as restrictive as possible about the relative time trends of each insurer. However, the results are qualitatively similar with the fixed effect specification used in equation (1).

switch throughout the panel to contribute to estimating λ_{jt}. Formally, the specification is

$$log(p_{ijt}) = \lambda_i + \lambda_{jt} + \sum_{\tau-3}^{3} \beta_\tau \tau_{it} + \beta_X X_{it} + \varepsilon_{ijt}, \tag{2}$$

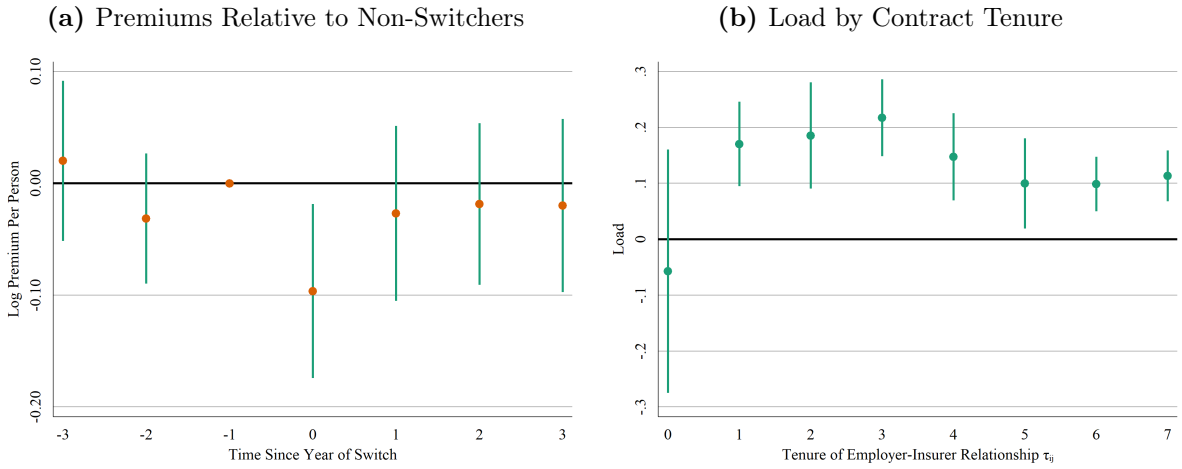
where β_τ estimates the difference between premiums among switchers and non-switchers in year τ, where τ_{it} indexes the year relative to an employers' switch. To ensure that observed trends in premiums do not represent coincident changes in plan quality – other than the identity of the insurer – I also control for a vector of plan characteristics, X_{it}.[7]

In Figure 2, I plot the estimates from the regression in equation 2. Employers who switch display premium trends that are broadly similar to other employers enrolled with each insurer, before and after they switch. However, premiums in the year of the switch are substantially lower than would be predicted by the switching employers' average premium levels and the premium trends of non-switching employers who contract with the same insurer. Indeed, employers appear to receive a one-time discount of 9 percent in the year they switch. This result is consistent with models of dynamic pricing in the presence of switching costs, in which supplier firms set low prices to *invest* in acquiring new market share, and raise prices to *harvest* rents from locked-in consumers. This intuition is built into all models of switching costs (Beggs and Klemperer 1992, Farrell and Shapiro 1988, Klemperer 1995, Farrell and Klemperer 2007). However, these models typically address the balance of invest-harvest incentives for firms setting uniform prices. Individually set premiums allow me to observe the invest-harvest dynamics because insurers can price discriminate between incumbent employers and new customers, as in Cabral (2016).

A competing explanation for this pattern is that employers' health care utilization is lower in the year of a switch. This could occur if, for example, employers face frictions in navigating a new referral system or searching for new physicians. In that case, employers may pay less because they have lower expected costs in the first year of a contract. To

[7]The results of the event study regression in equation (2) are robust to inclusion or exclusion of plan controls.

Figure 2: Dynamic Premium Setting



Notes: Panel (a) shows estimates of β_τ from equation (2). Range bars show 95% confidence intervals, with two-way clustered standard errors at the employer and insurer level. Panel (b) shows the average load by time relative to the switching event (τ_{ijt}). Range bars present 95% confidence intervals.

test whether low first-year premiums are driven by underlying costs, I use average claims reported for the subset of experience rated plans to calculate the load (or margin) of the insurance contract, where $load = premiums/claims$. In Panel (b) of Figure 2, I plot the average load over the tenure of observed employer-insurer relationships (τ). The average load for $\tau = 0$ indicates that insurers tend to set premiums *below* cost for in the first year of a contract, with the expectation of earning positive margins in future years. The data on claims are sparse – not all employers report their claims experience, and those who do report claims often do not report consistently year-after-year. Therefore, the estimates in Panel (b) are not based on a balanced panel. The prevalence of claims reporting in $\tau = 0$ is especially rare, making that estimate particularly noisy. However, taken together, these results are consistent with invest-harvest premium setting.

3.2 The Role of Switching Costs

The results on dynamic pricing provide two important insights for the purpose of modeling employer choices. First, because premiums reflect each employer's costs and preferences,

the data do not contain premiums for employer-insurer contracts that did not occur. The presence of these discounts reveals that the counterfactual premium for switches that do not occur is actually a "sale" price.

Second, the presence of switching costs reveal that employer decisions are the result of factors other than preferences over insurers and their marginal utility of income. Instead, demand may be less elastic simply because employers prefer to remain with the same insurer for multiple periods, irrespective of the identity of that insurer. To understand the magnitude of this issue, consider the simple example of a random utility model for consumer i choosing between insurers j and k. I index whether a choice represents the previously chosen insurer – or *home* insurer – by superscript $\{h, -h\}$. In the absence of switching costs, the probability of $s_{ij}^{h} = s_{ij}^{-h}$, but switching costs imply that $s_{ij}^{h} > s_{ij}^{-h}$.

Assuming employers have time-invariant preferences over insurers, the likelihood of observing two sequences with the same *set* of choices, but different numbers of switches, quantifies the magnitude of switching costs. By examining sequences of three choices, a higher prevalence of sequences containing a "stay" followed by a "switch," relative to a "switch" followed by a "return switch" to the choice observed in the first period, indicates the presence of switching costs. That is, observing employers who choose the sequence BCBS-BCBS-HPHC more frequently than BCBS-HPHC-BCBS provides information about the magnitude of switching costs.

This exercise is based on the intuition from Honoré and Kyriazidou (2000), and is in the spirit of prior work assessing state dependence in binary choice settings (Halliday 2007), and under stronger assumptions over heterogeneity in brand preference (Drake et al. 2020). In Appendix D, I demonstrate that this test is robust to the number of choices in the market, the underlying distribution of brand preferences, and the initial conditions of the sequences. I find that, on average, employers are 18.9 times more likely to choose a sequence that involves one switch instead of two, and strongly reject a null hypothesis of no switching costs. The probability of observing any sequence declines in the number of switches it involves, despite

the fact that switching involves obtaining a short-term discount in premiums.

These patterns indicate that accounting for switching costs in demand estimation is important for separating true preference over insurers from state-dependence. Switching costs also have implications for modeling the supply side. A model of dynamic pricing will predict that insurers will set higher premiums for locked-in employers than they do for those they wish to attract from competing insurers. However, a forward-looking insurer may be restrained from fully extracting the surplus associated with switching costs because they value retaining the employer in future periods.[8] For the purpose of supply estimation, failing to account for dynamic incentives could lead to biased estimates of costs.

4 Model

In this section, I describe a model of employer preferences and insurer premium setting to separate the components of premiums described in Section 3. Employers obtain utility from contracting with each insurer, and face costs associated with switching between insurers. Insurers compete for employers in a dynamic price setting game, by setting premiums in each period over an infinite horizon. For each employer, insurers compete in a Nash-Bertrand sub-game each period, where utilities and premiums depend on the identity of the insurer the employer contracted with in the previous period.

4.1 Demand

Each period, an employer chooses among insurers $j = \{1, 2, 3...\} \in \mathcal{J}$ with utility

$$u_{ijt|h_{-1}} = \theta_{ij} - \alpha_i p_{ijt|h_{-1}} + \eta_i \mathbb{1}\{j = h_{-1}\} + \epsilon_{ijt}. \tag{3}$$

[8] These forces are at the crux of the theoretical work on switching costs and competition, which yields theoretically ambiguous predictions regarding whether switching costs raise or lower prices on net.

θ_{ij} is an employer-specific preference for insurer j, and is equivalent to specifying that θ_{ij} represents preferences over employer i most favored plan attainable through insurer j. Although I assume that these preferences are fixed over time primarily for tractability, this assumption is not overly restrictive. Employers in this market exercise a high level of discretion in specifying the characteristics of the health plan. In Appendix E, I present excerpts from a representative request for proposals (RFP), in which the employer specifies that bids from competing insurers must "be based on exact duplication of the existing plan benefits." Moreover, preferences for the quantity of health care consumption will be captured as expected cost in the supply model. θ_{ij} therefore captures only the degree to which an employer prefers that a plan be provided by one insurer over another.

There is no outside option in the model. Empirically, nearly all employers of the relevant size purchase insurance, and self-insured firms contract with the same set of insurers to provide networks and other administrative services. h indexes the employer's choice in the current period, and $h_{-1} \in \mathcal{H}$ indicates the identity of the insurer chosen in the previous period. η_i therefore captures the additional utility employer i obtains from remaining with the same insurer they chose in the previous period. Because there is no outside option in the model, this is equivalent to specifying a negative cost to switching. α_i is a price (premium) coefficient and ϵ_{ijt} is a type-1 extreme value shock. I denote the employer-specific, and state-dependent choice probabilities by

$$s_{ijt|h_{-1}}(\mathbf{P}_{it}) = \frac{e^{\theta_{ij} - \alpha_i p_{jt|h_{-1}} + \eta_i \mathbb{1}\{j=h_{-1}\}}}{\sum_{k \in \mathcal{J}} e^{\theta_{ik} - \alpha_i p_{kt|h_{-1}} + \eta_i \mathbb{1}\{k=h_{-1}\}}}, \tag{4}$$

where $\mathbf{P}_{it}$ contains the vector of premiums available from all insurers.

4.2 Supply

I now specify a model of insurer premium setting that allows for insurers to incorporate the employer's state-dependent utility into their pricing policy by setting low premiums to

attract new employers, and raising them for incumbent employers. My general approach adheres closely to the details of the model in Cabral (2016), in which switching costs not only result in invest-harvest incentives, but "customer-specific" pricing allows selling firms to distinguish between new and locked-in consumers with varying levels of switching cost.[9] In addition to rationalizing the premium setting observed in the Section 3, this model provides a means to estimate costs and compute counterfactuals in which insurers strategically alter their premium setting in the presence of a public option.

For a given employer i, insurer j sets premiums to maximize the sum of expected flow profit in the current period, as well as the expected value of future periods over an infinite horizon. The value of winning the sale in period t can be written

$$\begin{aligned} V_{ijt}(h|h_{-1}) = \pi(p^*_{ijt|h_{-1}}, c_{ij}) \cdot \mathbb{1}\{h = j\} + 0 \cdot \mathbb{1}\{h \neq j\} \\ + \beta \left[\sum_{h_{+1} \in \mathcal{H}} \mathbb{P}(h_{+1}|h, p_{ijt|h}, p_{-ijt|h}) \cdot V_{ijt+1}(h_{+1}|h) \right]. \end{aligned} \tag{5}$$

If the employer chooses insurer j, then $h = j$ and the insurer earns flow profit $\pi = p_{ijt|h_{-1}} - c_{ij}$, which depends on the premiums it sets and the *expected* cost of insuring the employer. I assume that insurers have full information about the employer's preferences determining insurer choice, as well as their expected costs. The RFP in Appendix E shows that insurers receive substantial financial detail, allowing them to accurately forecast the expected cost of an employer. For simplicity, I also assume that the expected cost varies only by employer and insurer. Employers may have additional time-varying costs due to transitive (i.i.d.) shocks to their health care utilization, which the insurer provides insurance over. In this case, c_{ij} can be thought of as the unbiased expectation of the employers' cost realization, and π is easily interpreted as the profit the insurer expects to earn upon winning the sale but before health care consumption is realized. Assuming that costs vary only by (i, j) greatly simplifies the model and estimation because the possible states are observed directly and

[9] Dubé et al. (2009) present a similar model for illustrative purposes, but empirically consider the market for packaged goods, in which uniform prices reflect the balance of invest-harvest incentives.

fully characterized by the set of possible choices in the previous period ($\mathcal{H}$). Although doing so is beyond the scope of the current paper, it may be desirable to relax this assumption to allow for employers' costs to evolve over time and for insurers to incorporate shifts in cost in premium setting. In particular, employers may experience different expected costs over time if employee turnover generates some drift in health status or preferences of employees, or if shocks to income generated by the employer's output market shift its employees' willingness-to-pay for health care consumption.

In the case where expected costs are at least as coarse as (i, j, h_{-1}), time enters each insurer's value function only through h_{-1}, and the problem can therefore be written as stationary, provided that each insurer only sets one premium for each employer-state combination. Then, equation (5) can be re-written

$$\begin{aligned} V_{ij}(h|h_{-1}) = \mathbb{1}\{h = j\} \cdot (p_{ij|h_{-1}} - c_{ij}) + \mathbb{1}\{h \neq j\} \cdot 0 \\ + \beta \left[\sum_{h_{+1} \in \mathcal{H}} \mathbb{P}(h_{+1}|h, p_{ij|h}, p_{-ij|h}) \cdot V_{ij}(h_{+1}|h) \right]. \end{aligned} \tag{6}$$

I denote the ex ante value of a given state

$$V_{ij}(h) = \sum_{h_{+1} \in \mathcal{H}} \mathbb{P}(h_{+1}|h, p_{ijt|h}, p_{-ijt|h}) \cdot V_{ijt+1}(h_{+1}|h) \tag{7}$$

Then, the ex ante value of h_{-1} is defined recursively as

$$\begin{aligned} V_{ij}(h_{-1}) &= \sum_{h \in \mathcal{H}} \mathbb{P}(h|h_{-1}, p_{ij|h_{-1}}, p_{-ij|h_{-1}}) \cdot V_{ij}(h|h_{-1}) \\ &= \mathbb{P}(j|h_{-1}, p_{ijt|h_{-1}}, p_{-ijt|h_{-1}}) \cdot (p_{ij|h_{-1}} - c_{ij}) \\ &+ \beta \left(\sum_{h_{+1} \in \mathcal{H}} \mathbb{P}(h|h_{-1}, p_{ij|h}, p_{-ij|h}) \cdot V_{ij}(h) \right). \end{aligned} \tag{8}$$

Because the employer's choice fully characterizes the state transition, I can substitute the

choice probabilities into equation (8) to obtain

$$\begin{aligned} V_{ij|h_{-1}} &= V_{ij}(h_{-1}) \\ &= s_{ij|h_{-1}} \cdot (p_{ij|h_{-1}} - c_{ij}) \cdot + \beta \left(\sum_{h \in \mathcal{H}} s_{ih|h_{-1}} \cdot V_{ij|h} \right), \end{aligned} \tag{9}$$

where I have moved indexing of states to the subscript for parsimony.

Now, faced with an employer in state h_{-1}, insurer j's optimal premium decision is characterized by

$$V^*_{ij|h_{-1}} = \max_{p_{ij|h_{-1}}} \left\{ s_{ij|h_{-1}}(p_{ij|-h} - c_{ij}) + \beta \left(\sum_{k \in \mathcal{J}} s_{ik|h_{-1}} V^*_{ij|h=k} \right) \right\}. \tag{10}$$

Solving the insurer's first order condition for the optimal premium offer yields

$$p^*_{ij|h_{-1}} = c_{ij} + \overbrace{s^*_{ij|h_{-1}} / (-\partial s^*_{ij|h_{-1}} / \partial p^*_{ij|h_{-1}})}^{\text{Static Markup}} - \beta \underbrace{\left(V^*_{ij|j} - \sum_{k \neq j} \left(\frac{\partial s^*_{ik|h_{-1}} / \partial p^*_{ij|h_{-1}}}{-\partial s^*_{ij|h_{-1}} / \partial p^*_{ij|h_{-1}}} \right) V^*_{ij|k} \right)}_{\text{Marginal Value of Being } h \text{ in next period}}. \tag{11}$$

In setting premiums, the insurer optimizes over the balance of obtaining the sale and earning flow profit, as well as the dynamic incentive to obtain or return to a position as the home insurer in future periods. Equation (11) demonstrates that premiums are a function of the cost of insuring employer i, the standard markup incentives implied by a static trade off between higher prices and higher market share, as well as the discounted marginal value of being the home insurer in the next period.

Each insurer has a policy function containing premium offers for each state ($h \in \mathcal{H}$). In a model with no public option $\mathcal{H} = \mathcal{J}$ because every employer purchases insurance in each period. I use a Markov perfect equilibrium solution concept, in which the premium setting game for employer i is in equilibrium when the vector of each insurers' premium offers is optimal given the expected future values of states generated by the premium offers of all other insurers in all states. The solution in this case is a set of $\mathcal{J} \times \mathcal{H}$ premiums — one for each

insurer-state combination. Even in the case where costs are not state-dependent, each insurer may make $\mathcal{H}$ distinct premium offers if the probability of winning a sale differs depending on the identity of h_{-1}. This $\mathcal{J} \times \mathcal{H}$ matrix of premium offers, along with the demand parameters, determines a set of employer choice probabilities and equilibrium continuation values ($V^*_{ij|h}$), and therefore fully characterizes the equilibrium.

Both my method of estimating costs and the counterfactuals in which I introduce a public option require me to solve the model. I follow prior work in dynamic oligopoly settings (Pakes and McGuire 1994), which use Gaussian methods, to compute the MPE. I solve the system of $\mathcal{J} \times \mathcal{H}$ first-order conditions implied by equation (11) using a Gauss-Seidel updating procedure (as in Benkard 2004 and Doraszelski and Judd 2004). I begin with a guess of the value functions. I then cycle over each state, computing the optimal policy function for each insurer and updating the premium offers and value functions at each step until the system converges.

I do not know whether an equilibrium is guaranteed to exist, nor if any particular solution is a unique equilibrium. In practice, the solutions I obtain are (i) insensitive to starting guesses, (ii) smooth in parameters, and (iii) consistent with the comparative statics demonstrated by Cabral (2016), who derives both existence and uniqueness results for a similar model with two selling firms.

5 Estimation

In order to estimate the complete model, I pursue a two-step estimation strategy. In the first step, I estimate demand using a mixed-logit model to generate a measure of choice probabilities, which form the transition probabilities between states. For unobserved states and choices, I use a Lasso regression to estimate the unobserved premium offers. In the second step, I use the insurers' first-order conditions to recover costs.

The key output I am interested in obtaining is a joint distribution of preferences and

costs, with which I can simulate the market-level impact of introducing a public insurance option. That is, I wish to recompute the equilibrium premiums and choices, as described in Section 4, with a choice set that includes a public option with assumed characteristics θ_{ij}^{PO} and premiums $\overline{p}^{PO}$. The joint distribution of preferences and costs matters primarily because it will impact the break-even premium that a public option could offer based on which employers choose to enroll at a given level of quality and premiums. However, the equilibrium premium offers from private insurers continue to be made on an employer-by-employer basis.

5.1 Data

The data include a panel of employer-insurer contracts from Massachusetts over 2001-2018. I observe total enrollment, premiums, and identifiers for employers and insurers. There are three insurers that account for nearly 90 percent of the market share in any given year and I consolidate the remaining insurers into a composite choice. In order to credibly identify the demand model, I restrict the data to employers who I observe for at least three years, and who switch at least once. This restriction reduces the number of employers from 1,051 to 384. Although this reduction in sample is substantial, it is largely driven by short employer panels. In Appendix C, I conduct additional analyses testing the relationship between sample inclusion, attrition, and characteristics of the employers' choices. I find that the number of firms that meet Form 5500 filing requirements has increased over time, generating a large number of short panels. Moreover, because switching is infrequent, shorter panels provide fewer opportunities to observe a switch between insurers. However, conditional on an employer's first appearance in the data and tenure, switchers are no more likely to attrit from the sample than non-switchers, suggesting that these restrictions do not systematically bias the sample toward employers with different preferences.

As in the preceding analyses, I compute a per-person per-year measure of premiums to standardize premiums across employers of varying sizes, and within employers as they

experience changes to the number of employees. In order to account for aggregate shocks to health care costs, I inflation adjust premiums to 2018 dollars using the managed care producer price index.

5.2 Demand Estimation

I use simulated maximum likelihood to estimate a flexible distribution of employer preferences. This method is based on Train (2016), which builds on the approaches from Bajari et al. (2007) and Fox et al. (2011), and is similar to the approach taken in Nevo et al. (2016). My general strategy is to simulate choice sequences for each employer over a range of potential parameter values. I then use a simulated likelihood procedure to choose probability weights over the distribution of preferences that generate sequences of choices matching most closely to the data. The key insight from Train (2016) is that using a logit to estimate the simulated likelihood function allows for a smooth but flexible distribution of preferences over a fine grid of parameter values with a limited number of parameters, while constraining the weights to be positive and sum to 1.

In the first step, I specify a grid of over a finite parameter space to generate predicted choice sequences for each employer over the varying parameter values. I then use a second logit to estimate probability weights over the grid. Formally, the probability that employer i made a given sequence of choices is the product of choice probabilities over observed choices implied by preferences θ_r:

$$L_i(\theta_r) = \prod_{t \in T_i} s_{ij_{(it)}t|h_{(it)}}(\theta_r).$$

θ_r is a specific draw of parameters from a finite preference space S_i. The simulated log likelihood for the estimator is

$$\mathcal{L} = \sum_i log \left(\sum_{r \in S_i} L_i(\theta_r) w_i(\theta_r|\alpha) \right),$$

where $w_i(\theta_r|\pi)$ is the probability that $(\theta_i = \theta_r)$ as a function of π,

$$w_i(\theta_r|\pi) = \frac{e^{\pi' z(\theta_r)}}{\sum_{s \in S_i} e^{\pi' z(\theta_s)}}.$$

The estimator recovers the vector of parameters, π, for a basis function $z(\theta)$ that maximizes the simulated log likelihood. The basis function $z(\theta)$ can be specified in a number of ways. I estimate using spline functions with 5 knots for each model parameter for a total of 30 $(= (5+1) \times 5)$ estimated (π) parameters.

Finally, the recovered distribution relates back to the data through

$$h_i(\theta_r|\hat{\pi}) = \frac{L_i(\theta_r) w_i(\theta_r|\hat{\pi})}{\sum_{sin S_i} L_i(\theta_s) w_i(\theta_s|\hat{\pi})}. \tag{12}$$

Borrowing intuition provided by Nevo et al. (2016), choosing a grid over θ_r is akin to specifying a uniform prior over the parameter space. Each type r has an optimal set of choice probabilities, and therefore likelihood of sequences, given the premiums it faces. As the number of sequences I observe increases, the estimation procedures recovers the distribution of types. Equation (12) represents the posterior probability, or probability mass at θ_r, given employer i's observed premiums and choices. The estimated values for $h_i(\theta_r|\hat{\pi})$ allow me to generate employer-specific predicted values $\hat{s}_{ijt|h}$, which form the transition probabilities used in the supply estimation.

Unobserved premium offers: One key challenge to estimating demand using a multinomial logit is that I only observe premiums for choices that actually occurred. The average employer appears in the data for approximately 7 years and visits two insurers, meaning that many ij pairs are unobserved throughout the entire sample. However, even ij pairs that are sometimes observed require imputed prices for years in which the employer makes a different choice. For example, an employer who switches from insurer 1 to insurer 2 halfway through the panel needs an estimate of premiums for insurer 2 in the early years and insurer 1 in the

later years. To address this, I impute premiums for unobserved choices in observed states *and* for all choices in unobserved states. These imputed premiums facilitate the estimation of the demand system. They also form the "policy function" used to compute costs in the supply estimation.

In order to use the imputed premiums as data to provide an unbiased estimate of the demand parameters, I must assume that the imputed premiums respect any differences between the premiums for i, j, h_{-1} combinations that were not chosen and those that were. In order to limit potential violations of this assumption, I use a large number of observables to capture employer-specific costs, and adjustments for a large number of characteristics that proxy for any preferences insurers would account for in setting premiums:

- Employer i's average observed premiums for insurer j in states that are observed and employer's average premiums with other insurers
- Summary measures of the employers' choice sequence (shares of observed choices each insurer, percent of choices that represent a switch)
- Indicators for whether the averages for certain states were missing for each employer (not all employers are observed in all states)

In order to provide the most flexible specification, and maximize the inclusion of appropriate observables, I estimate a Lasso regression to predict logged premiums as a high-dimensional function of the above characteristics. The intuition behind this approach is that observed premiums for employer i, in other states and with other insurers, will capture information about employer-specific contributions to cost. Allowing these to vary by insurer will capture insurer specific cost advantages. Incorporating the choice sequence measures will capture relative discounts that are driven by price discrimination on employer preferences (θ_{ij}, α_i and η_i). For example, an employer whose choice sequence reveals lots of persistence should correspond to a larger invest-harvest motive. Additional detail about the estimation approach, and fit of the model is provided in Appendix F.

5.3 Recovering Costs

I recover the expected costs using a nested fixed point approach. Given the demand parameters, and a guess of costs, I compute equilibrium premium offers for each insurer-state combination. For each employer, I select the vector of cost parameters that, together with the demand parameters, minimize the distance between predicted and observed premiums.

Identification The cost parameters can be obtained from the insurers' first-order conditions from equation (11), which imply an estimate of c_{ij}.

$$\begin{aligned} c_{ij} = p^*_{ij|h_{-1}} - s^*_{ij|h_{-1}} / (-\partial s^*_{ij|h_{-1}} / \partial p^*_{ij|h_{-1}}) \\ + \beta \left(V^*_{ij|j} - \sum_{k \neq j} \left(\frac{\partial s^*_{ik|h_{-1}} / \partial p^*_{ij|h_{-1}}}{-\partial s^*_{ij|h_{-1}} / \partial p^*_{ij|h_{-1}}} \right) V^*_{ij|k} \right) \end{aligned} \tag{13}$$

or in logit notation

$$c_{ij} = p^*_{ij|h_{-1}} - \frac{1}{\alpha_i (1 - \hat{s}_{ij|h_{-1}})} + \beta \left(V^*_{ij|j} - \sum_{k \neq j} \frac{\hat{s}_{ik|h_{-1}} V^*_{ij|k}}{1 - \hat{s}_{ij|h_{-1}}} \right), \tag{14}$$

This is analogous to the inversion commonly used to recover costs in static settings (Bresnahan 1981, Berry et al. 1995, Nevo 2001), where costs are inferred as the residual between observed premiums and the markups implied by demand elasticities. Respecting the dynamic incentives only require that these markups be corrected for the marginal value of being the home insurer, itself a function of the premiums of other insurers and the resulting discounted future profits.

In equilibrium, the complete set of premium offers and costs imply a set of equilibrium continuation values, $V^*_{ij|h_{-1}}$, which depend on the costs, demand parameters, and optimal premium offers. Given estimates of the demand parameters, costs are identified from the system of equations implied by equation (14) as long as there is an estimate of $p^*_{ij|h_{-1}}$ (and therefore $\hat{s}_{ij|h_{-1}}$) for each employer-insurer-state. For each employer, there exist $\mathcal{J} \times \mathcal{J}$

first-order conditions, one for each insurer in each state. Since each employer has only $\mathcal{J}$ cost parameters, the system is over-identified, and I use a simplex algorithm to find the vector of cost parameters that minimizes the sum of squared differences between observed premiums and those implied by the model. Because insurers set premiums based on the employer's expected cost, the insurer has already integrated over any i.i.d. cost shocks. I interpret the remaining error resulting from the over-identifying restrictions as measurement error resulting from imperfect measurement of premiums. This specification of the error term is also consistent with fluctuations in premiums that reflect actual pricing decisions. For example, measurement error may also enter the data if insurers' imperfectly predict employers' expected costs. This kind of "measurement error" is modeled in Crawford and Yurukoglu (2012), and is equivalent to noise generated by imperfect measurement of premiums or enrollment as long as it does not systematically reflect changes to an employers' cost or preferences over time. I describe potential sources of measurement error in Appendix H.

Assuming that employers' costs and preferences are constant over time is an assumption I impose for tractability. In practice, employers may experience changes on both of these dimensions due to employer turnover or other exogenous factors. However, this limitation stems from the use of the panel data to identify the model. Inferring preferences from choice sequences, and therefore the costs from premiums, requires an assumption that the choices and premium offers in different periods were made under similar conditions. Such violations would be especially problematic if they are correlated with switching behavior. That is, if employer switching reflects unobserved shifts in preferences causing them to purchase a dramatically different plan from the new insurer, my approach could lead me to attribute these changes to permanent cost differences across insurers. While I cannot rule this out in all cases, the event study presented in Figure 2 provides evidence that employers' premiums are similar before and after switching events, up to an adjustment for the insurer and time shocks. Moreover, in a competitive labor market, where employees sort into employer groups based on the value of health insurance in their total compensation, new employees will have

similar preferences to existing employees.

Since I do not know whether the equilibrium of my model is unique, the nested fixed point is not guaranteed to find the equilibrium observed in the data. In Appendix G, I propose an alternative estimation approach that uses the preferences recovered in demand estimation along with observed premiums to simulate the insurers' value functions. I then recover costs using a simple first order condition, treating the simulated optimal continuation values V^* as data. Estimates from this approach are qualitatively similar to the those obtained using the nested fixed point, providing supportive evidence that those parameters were not obtained from the wrong equilibrium.

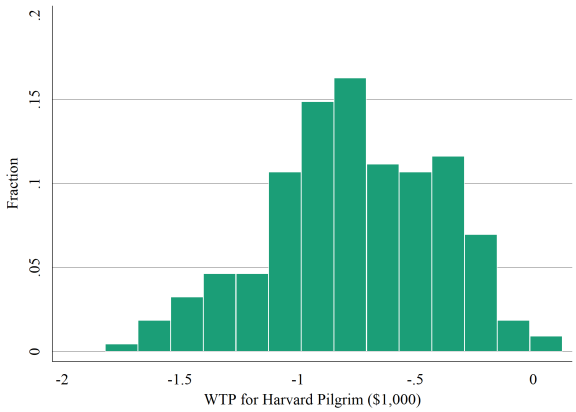
6 Results

Demand Figure 3 presents the results of the demand estimation. As aggregate market shares would imply, the average coefficients on HPHC, Tufts, and Other insurers are negative. They are also widely dispersed. Panels (a) and (b) present distributions of employers' willingness to pay for insurers, calculated by dividing $\widehat{WTP}_{ij} = \hat{\theta}_{ij}/\hat{\alpha}_i$. Employers display strong brand preferences. The average employer is willing to pay \$489 per person, per year for its most preferred insurer relative to its second most preferred insurer.

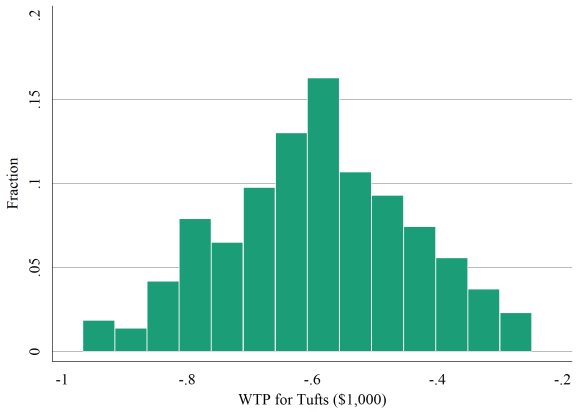
Switching costs are also large. The average employer has switching costs of \$3,020, which is roughly 49 percent of average premiums. This estimate is in line with prior work examining switching costs of households choosing among plans in the Massachusetts Group Insurance Commission (GIC), which provides an insurance exchange for public employees in the state. Tilipman (2018) estimates individual consumers choosing among health plans in the GIC face average switching costs of \$250 per household per month, or \$3,000 per year. These estimates are therefore consistent with employers' switching costs representing aggregations of those faced by their employees, rather than additional agency problems or contracting frictions faced by the employer-group.

Figure 3: Estimated Demand Parameters

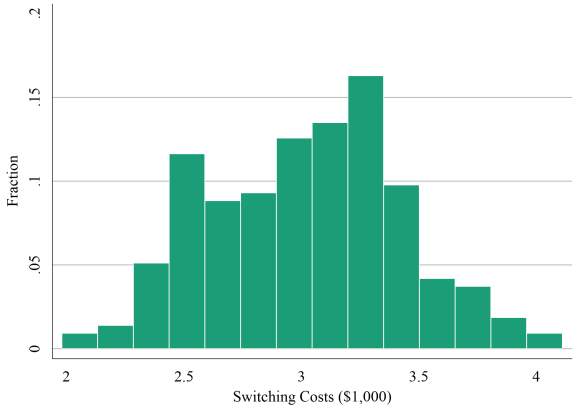
(a) Preferences for Harvard Pilgrim



(b) Preferences for Tufts



(c) Switching Costs



	Mean	SD
θ_{HPHC}	-0.274	0.142
θ_{Tufts}	-0.213	0.065
θ_{Other}	-0.364	0.118
η	1.088	0.236
α	0.357	0.035

Notes: Estimates of parameters from the demand system as described in Section 5. Panels (a) and (b) show the distributions of willingness-to-pay measures for Harvard Pilgrim and Tufts relative to Blue Cross and Blue Shield. Panel (c) displays the distribution of switching costs. The table in the Southeast corner presents the mean and standard deviations of each parameter in equation (**??**).

Table 2: Own-Price Elasticities by Insurer and State

	(1)	(2)
	Home ($j = h_{-1}$)	Non-Home ($j \neq h_{-1}$)
BCBS	-0.402	-6.530
HPHC	-0.462	-5.534
Tufts	-0.390	-5.715
Other	-0.650	-5.854

Notes: Table displays own-price elasticities for each insurer for the home state ($j = h_{-1}$) (column (1)) and all other states (column (2)). All averages weighted by steady-state choice probabilities.

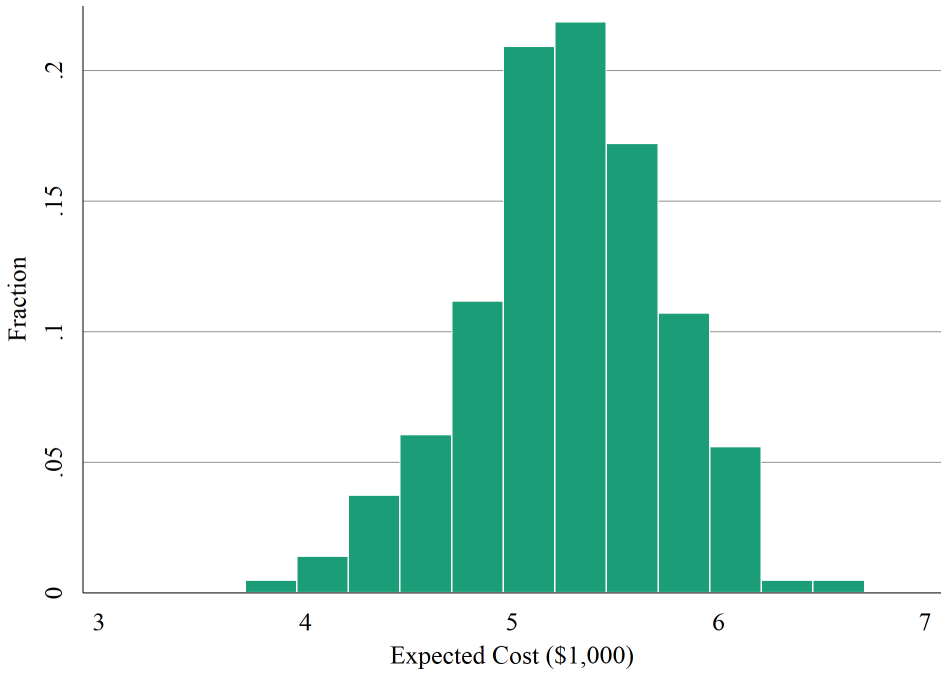
In Table 2, I present the own-premium elasticities for each insurer among switchers and non-switchers. Elasticities are low for employers who remain with the home insurer. The prices that generate these elasticities are inconsistent with a model of premium setting in which insurers set premiums to compete in a static Nash-Bertrand equilibrium. Higher elasticities for switchers indicate that discounted premiums are optimized to a more elastic employer.

Taken together, these results demonstrate that both brand preference and switching costs are large in absolute terms. However, the importance of these magnitudes depend on the degree to which they translate into markups over cost, which I turn to in the next section.

Costs In Figure 4, I show the results of the cost estimation. The accompanying table summarizes the estimates. The average employer has expected costs of $5,267, consistent with aggregate estimates of average claims per covered life published by CMS of $5,400. Columns (1)-(3) of the accompanying table show that HPHC and Tufts enjoy a small cost advantage relative to BCBS. This is consistent with both the premium variation documented in Section 3, and prior work showing that HPHC and Tufts negotiate lower rates with hospitals for common health care services (Craig et al. 2021).

Accounting for insurers' dynamic incentives plays an important role in generating these

Figure 4: Cost Estimates



	(1)	(2)	(3)
	Average Cost Estimates		
	BCBS	HPHC	Tufts
Mean	5,579	5,093	5,143
SD	(636)	(473)	(559)

Notes: Estimates reported based on the cost estimation procedure described in 5. Means and standard deviations are weighted according to the steady state market shares from the demand model.

cost estimates. Because switching costs are large, employers typically choose their home insurer and estimated own price elasticities are low in those states. In Appendix G, I re-estimate costs using a myopic insurers' first-order condition. This approach produces average costs of -$8,813. 89% of employers have an expected cost that is negative. The dynamic model used to generate the preferred estimates in Table 4 rationalizes the premiums in the data with more realistic cost estimates. In the dynamic model, forward looking insurers not only offer discounts to induce employers to switch, but also do not fully extract rents from locked-in consumers because they place additional value on remaining the home insurer in

future periods.

7 Implications for a Public Option

I now use the estimated model to recompute the market equilibrium with a public option, under various assumptions about its cost, quality, and premium setting rules. I model the public option as an additional element in each employer's choice set. As noted in Section 2, the details of the public option — both how aggressively it would lower prices for health care services, and the extent to which it would maintain a large network while doing so — are not specified precisely in existing policy proposals. In order to accommodate a range of potential scenarios, I provide predictions for a variety combinations of cost and quality, but I emphasize that some combinations in which a public option provides high quality at a large cost advantage may be infeasible. I also test alternative design features to evaluate the public option's ability to contend with both adverse selection and switching costs.

The results in this section represent stable equilibria in the sense that (1) using the supply model, I allow insurers to set new premiums that account for any additional competitive pressure, and (2) I present employers' enrollment based on forward-simulated choice probabilities in the new equilibrium so that estimated outcomes reflect a "steady state." However, these results reflect a partial equilibrium response in the sense that I do not model upstream price setting between health insurers and health care providers. Fiedler (2020) shows that, under some conditions, private insurers can use the threat of a dominant public option to induce hospitals and physicians to agree to payment rates that are below the status quo but higher than what they would receive from a public option with regulated payments. By holding c_{ij} constant, my results may over-estimate enrollment in the public option and under-estimate the ability of a public option to lower premiums in the private market by lowering the underlying unit cost of health care. I therefore interpret the public option's market share as an indicator for the pressure these providers might face. I discuss

the implications of the distinction for competition and welfare below.

When presenting welfare figures, I focus on changes to consumer surplus. Because the market is complete at baseline, there is no change to total welfare from any of the policies under consideration. Premium adjustments that result from lower markups represent a transfer from insurers to employers. Premium adjustments that result from a public option that lowers cost by regulating provider payments represent a transfer from health care providers to consumers. This is also true for any potential premium adjustments that may result from private insurers' potential ability to leverage a public option with a strong cost advantage to negotiate lower reimbursements (as in Fiedler 2020), but which are outside the model. As discussed in Section 2, lowering rents in the market for health care services may distort incentives for providers. I therefore focus primarily on the savings that result from lower markups in the premiums market, as these are the most likely to provide clear welfare gains for consumers. Nonetheless, policies that increase consumers surplus by shifting substantial market share to the public option, or otherwise lowering the unit cost of health care, may provide substantial gains in consumer welfare.

7.1 Introducing a Public Option

Premium Setting and Equilibrium The public option I consider is non-strategic in the sense that it does not set premiums to maximize profit. However, I assume that the public option must break even so that it charges a premium equal to average cost among employers who purchase it, plus a small administrative load.[10] Hoffman et al. (2021) report that estimates for Medicare's administrative load range from 2-5 percent, and report estimates for private insurers ranging from 5-17 percent. For the main results, I focus on a 5 percent adjustment but replicate the analysis under alternative assumptions in Appendix I. Higher

[10]Whereas premium subsidies are common on the individual exchanges, it is not clear that they are desirable in the employer-sponsored market, where nearly all employers currently purchase plans. Requiring the public option to break even is also consistent with policy literature which highlights the low fiscal footprint as an advantage of the public option relative to more dramatic policies like Medicare For All Hoffman et al. 2021.

administrative loads make the public option more costly and therefore less attractive to employers. However, none of the patterns I discuss in this section are sensitive to this choice.

In most cases, I consider a public option with full community rating, where premiums are equal to the average cost of the employers who enroll. Formally, the public option sets a premium so that

$$\overline{p}_{PO} = \gamma \sum_{i} \sum_{h_{-1} \in \mathcal{H}} f_{ih_{-1}} \times s_{i,PO|h_{-1}} \times c_{i,PO} \tag{15}$$

where $s_{i,PO|h_{-1}}$ is the state-dependent probability that an employer chooses the public option conditional its previous choice, and $f_{ih_{-1}}$ is the steady state probability that the employer is in state h_{-1}. γ captures the administrative load. In order to present premiums offered by private insurers symmetrically, I calculate the average market premiums, weighting by state-dependent choice probabilities and the steady state probability that each employer is in state h_{-1}.

Whereas the model presented in Section 4 is computed independently for each employer, the equilibrium with a community rated public option imposes an interdependence between all employers. To solve for the new equilibrium of each hypothetical policy, I compute counterfactual private premium offers and choice probabilities over a large grid of public option premiums that spans the support of the underlying cost distribution. I then select the lowest feasible premium for each combination of cost and quality, ensuring that the public option maximizes enrollment subject to its break-even constraint. Note that this approach rules out the possibility that private insurers coordinate premium offers across employers with different costs in order to strategically generate negative selection into the public option. If possible, such an outcome would cause me to understate the degree of adverse selection into the public option.

Welfare Welfare follows directly from the demand specification. Denoting the mean utility of each state dependent choice for employer i

$$\delta_{ij|h_{-1}} = \theta_{ij} - \alpha_i p_{ij|h_{-1}} + \eta_i \mathbb{1}\{j = h_{-1}\},$$

expected consumer surplus for a given menu of potential insurers is the probability weighted sum of state-dependent inclusive values:

$$CS_i = \frac{1}{\alpha_i} \sum_{h \in \mathcal{J}} f_{ih} \left[ln \left(\sum_{j=\mathcal{J}} e^{\delta_{ij|h}} \right) \right].$$

Using $\mathcal{J}'$ to denote the set of insurers that includes a public option, the change in consumer surplus from adding a public option can be calculated

$$\Delta CS_i = \frac{1}{\alpha_i} \left[\sum_{h' \in \mathcal{J}'} f'_{ih'} ln \left(\sum_{j'=\mathcal{J}'} e^{\delta_{ij'|h'}} \right) - \sum_{h \in \mathcal{J}} f_{ih} ln \left(\sum_{j=\mathcal{J}} e^{\delta_{ij|h}} \right) \right]. \tag{16}$$

In order to disentangle direct substitution from increased competition, I recompute the change in consumer surplus, applying the premiums offered by private insurers in the presence of a public option, but without a public option in the choice set itself. Because steady-state choice probabilities adjust to changes in the choice set, this exercise does not provide an exact decomposition of the fraction of consumer surplus gained from price changes. Nonetheless, these results provide insight into the degree to which the public option impacts the market by generating downward pressure on insurers' premium offers.

Cost Advantage I evaluate policies where the public option has varying degrees of cost advantage over private insurers. Insurers' commonly negotiate prices with hospitals and physicians, and prior work has shown these prices are approximately twice as high as Medicare payment rates (Chernew et al. 2020, Cooper et al. 2019, Maeda and Nelson 2018). In each of the counterfactuals, I use BCBS as a benchmark and test three scenarios where the

public option has a cost advantage of 0, 25 and 50 percent. The 0 percent cost advantage represents a case in which the public option replicates the cost structure of BCBS. The 50 percent cost advantage therefore approximates a polar case in which the public option regulates provider prices in a manner similar to Medicare. The 25 percent cost advantage represents an intermediate policy, where the public option has some additionally monopsony leverage over providers or otherwise regulates provider prices at a somewhat higher level than Medicare rates. This policy is similar to the proposed Colorado Affordable Healthcare Option or Washington State's recently enacted "Cascade Care," which would fix or set maximum allowable provider payments at 155 and 160 percent of Medicare payment rates (Scott 2021).

Quality I consider a public plan that closely replicates plans offered by BCBS. I then consider policies that offer a public option with a "quality shortfall" relative to BCBS. Placing the public option in employers' preference space in this way requires an assumption that employers will have homogeneous preferences for it, relative to BCBS. In practice, employer groups may have heterogeneous preferences depending on the nature of the quality differences. If BCBS is primarily differentiated from HPHC and Tufts on network structure, and the quality shortfall of a public option represents characteristics other than network, there may be additional heterogeneity in employer demand for the public option.[11] In this case, the impact of vertical differentiation on enrollment may be approximated by counterfactuals considering larger administrative cost adjustment, since this adjustment imposes a disutility for the public option that increases at higher levels of cost. However, such comparisons may understate the ability of such a public option to resolve adverse selection by restraining its own premiums.

Since the model allows for a flexible distribution of preferences over insurers, these preferences do not necessarily have a clear vertical interpretation. However, there are reasons

[11]This may have material implication for the results, particularly regarding selection into the public option. Polyakova (2016) studies selection between public and private insurance in Germany and finds that preference heterogeneity over horizontal quality limits adverse selection in that context.

to expect that BCBS is vertically differentiated, and that the estimated brand-preferences reflect employers' varying tastes for quality. BCBS is the most commonly preferred insurer in the market, and has the weakest cost position and highest premiums. BCBS also maintains the largest coverage in the state, making it a natural starting point for thinking about what a public option with a complete network would attempt to replicate. Notwithstanding, this necessarily means that a public option which replicates BCBS quality will compete most directly for consumers with weaker tastes for HPHC, Tufts, or other insurers. A public option with lower simulated quality will compete more directly with HPHC and Tufts for most employers.

7.2 Performance of a Public Option

Cost and Quality The first row of Table 3 contains premium estimates and market shares for a public option that perfectly replicates BCBS on both cost and quality with a 5 percent administrative load. Such a policy obtains 14 percent market share and charges a premium of $6,600. Allowing the public option to more aggressively exercise a cost advantage lowers the unit cost of health care and therefore public option premiums. A cheaper public option is more attractive: a plan that perfectly replicates BCBS quality while paying 50 percent less than BCBS for all health care services absorbs 98.2 percent of the market. This pattern is stable across policies, holding fixed the quality of the public option.

Holding cost fixed, a lower quality public option is less attractive to employers and therefore experiences lower enrollment. A public option with a quality shortfall of $500 represents a plan with similar characteristics to the average employers' second-best choice in the private market. Such a public option experiences roughly half of the enrollment than one with no quality shortfall. With a large enough cost advantage, the public option is attractive to a large number of employers, despite large quality differences. A public option with an extreme quality shortfall of $2,000 obtains 53 percent market share if it lowers costs to levels resembling Medicare payment rates.

Table 3: Equilibrium Under a Public Option

Cost Advantage	Quality Shortfall	(1) PO Enrollment	(2) PO Premium	(3) Market Premium	(4) ΔCS (with PO)	(5) ΔCS (prices only)
0%	0	0.143	6,600	5,866	166	17
	–500	0.076	6,600	5,961	70	9
	–1,000	0.035	6,700	6,018	23	3
	–2,000	0.011	6,600	6,048	4	1
25%	0	0.715	4,800	5,153	1,559	61
	–500	0.494	4,900	5,415	891	50
	–1,000	0.315	4,900	5,635	481	36
	–2,000	0.086	5,000	5,946	84	10
50%	0	0.982	3,100	5,215	3,526	38
	–500	0.947	3,200	4,998	2,876	49
	–1,000	0.876	3,200	4,989	2,293	58
	–2,000	0.532	3,300	5,369	991	53
Baseline				6,058		

I report the average consumer surplus changes in Column (4) of Table 3. Changes to consumer surplus scale with the attractiveness of the public option. A public option with cost and quality identical to BCBS saves consumers $166 on average. By contrast, a public option that provides BCBS quality at a 50 percent cost advantage generates an additional $3,526 per person per year. In Column (5), I report the change in consumer surplus that would result from the new private market premiums, but without the public option in the choice set. Across all of the policies considered, the changes in consumer surplus reported in column (4) are substantially larger than those in column (5), indicating that welfare gains from increased competition are typically small relative to the direct effect of enrollment in the public plan. Premium competition provides small welfare gains in absolute terms: even a public option that perfectly replicates BCBS with a 50 percent cost advantage provides consumers with an additional $38 per person per year.

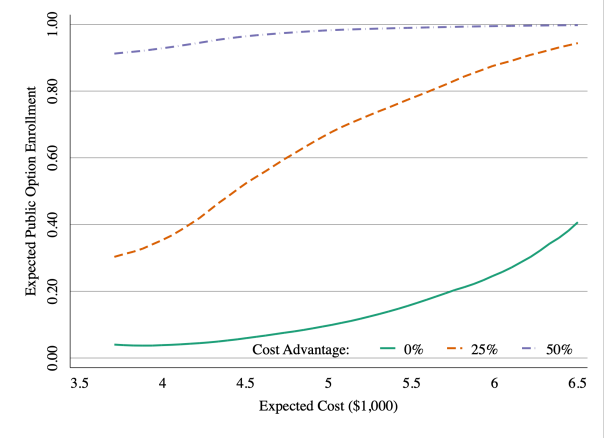
Adverse Selection and Competition The premium patterns in Table 3 suggest that the public option is negatively selected. For the public option with no cost advantage, public option premiums are \$6,600-\$6,700 — higher than the average employers' baseline premium offer. Changes to average premiums reported in Column (3) are large relative to the changes in consumers surplus reported in Column (5), indicating that the premium declines in Column (3) primarily reflect adverse selection into the public option with only modest shifts in insurers' premium offers.

Adverse selection limits both enrollment and premium competition because insurers can continue to compete by setting employer-specific, experience rated premiums. By attracting the most costly employers, the community rated public option enrolls mostly employers that require a high average premium, which is unattractive to employers with lower expected costs. In Panel (a) of Figure 5, I plot average enrollment probabilities across the distribution of employers' expected costs. Employers at the high end of the cost distribution are the most likely to enroll. This result is consistent with theoretical work by Shepard et al. (2020), who show that a uniform health insurance plan would involve substantial dead-weight loss because of wide variation in preferences for health care consumption. In this setting, where employers may opt in or refrain from participating in the public plan, only the highest cost employers tend to enroll, and the market unravels.

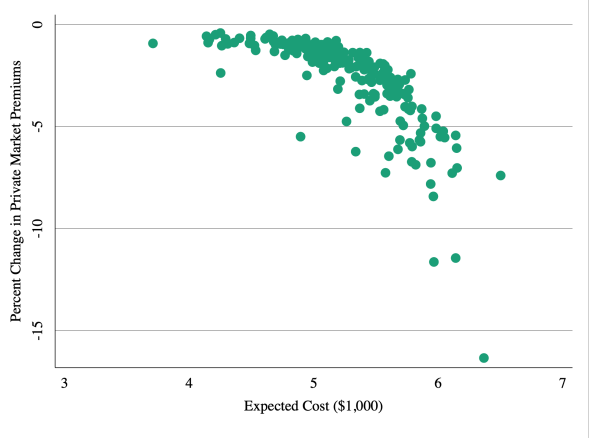
A public option with a cost advantage can dampen this effect by lowering the unit cost of underlying health care spending. However, doing so merely shifts the margin of selection to a lower point in the cost distribution. A public option with a 50 percent cost advantage nearly eliminates the issue, but only by lowering costs to the point that all employers are "covered." That is, because 50 percent of the highest cost employer is approximately equal to the cost of the lowest cost employer, this policy ensures that nearly all employers find it optimal to pool into the public option, as the public option can insure the most expensive employer at a premium that is close to the private market offer available to the lowest cost employer.

Figure 5: Adverse Selection into the Public Option

(a) Enrollment



(b) Private Market Premiums



Notes: Panel (a) shows estimates of steady state enrollment probabilities across employers for a public option that replicates BCBS quality. For parsimony, each line presents a local polynomial of the underlying data. Panel (b) presents estimates of changes to expected premium offers across employers for the case of a public option that perfectly replicates BCBS on both cost and quality. Premium offers and expected costs are weighted using the steady state market shares.

Adverse selection limits the ability of a public option to generate premium competition by restricting the set of employers for whom it represents a credible outside option. Employers with expected costs that place them sufficiently below the indifference point between their preferred private insurers and the public option are priced out, and therefore experience only modest changes in private premium offers. In Panel (b) of Figure 5, I present the percent change to private market premium offers when the market is faced with a public option with BCBS quality and no cost advantage. Premium offers decline substantially for marginal employers — those with high costs — but only modestly for employers who are less likely to choose the public option.

7.3 Alternative Premium Setting Policies

I now present two counterfactual policies intended to mitigate frictions in the market. First, I consider the case where the public option sets experience rated premiums. The first row of Table 4 presents the "baseline" case, reproduced from Table 3, in which the public option replicates BCBS cost and quality. In the second row, I allow for the public option to set

employer-specific, experience rated premiums. Under this policy, enrollment in the public option increases from 14 percent to 36 percent, resulting in a transfer of $592 to consumers. Moreover, allowing the public option to flexibly set premiums generates additional premium competition among private insurers. Lower premiums among private insurers result in a consumer surplus gain of $46 per person per year, nearly triple the value provided by a community rated public option. By contrast, a community rated public option with a 50 percent cost advantage generates only $38 in surplus from premium competition, while enrolling 98 percent of the market. The *overall* change to consumer surplus is larger in the case of a cost advantage, but only because the public option reduces health care costs directly — an effect that is more likely to generate quality distortions in the upstream market for health care services.

Table 4: Equilibrium Under a Alternative Design Features

	(1) **PO Enrollment**	(2) **PO Premium**	(3) **Market Premium**	(4) **ΔCS** (with PO)	(5) **ΔCS** (prices only)
Baseline PO	0.143	6,600	5,866	166	17
Experience Rating	0.364	6,202	5,659	592	46
10% Discounts	0.264	6,600	5,720	402	31

In Table 4, I focus on a public option with no cost advantage, where the averse selection problems are most apparent. I replicate the analysis for the full set of cost and quality combinations in Appendix Table A.5. The results are broadly consistent across potential policies. Except in the case where the public option has a 50 percent cost advantage — where the community rated public option mechanically overcomes the adverse selection problem — public option enrollment and additional savings from premium competition are larger under experience rating.

Next, I allow the public option to offer discounts to deal with inertia from switching

costs. Existing policy proposals have acknowledged the importance this issue, proposing introductory discounts in the first years of the public option to attract new employers (Hoffman et al. 2021). However, the counterfactuals I consider represent the steady state market shares of all plans, where switching costs allow insurers to continually engage in invest-harvest premium setting. Therefore, a public option offering one-time introductory premium offers would not generate this effect in my model, as market shares and premiums would eventually converge to the baseline equilibria in Table 3 over time. Instead, the plan I simulate replicates BCBS cost and quality, but offers a discount to any switching employers in any year. The third row of Table 4 displays outcomes for a public option that offers a 10 percent discount, approximating the pricing behavior of private insurers documented in Section 3. This public option obtains 26 percent market share, and transfers $402 to the average consumer. The premium offers from this policy alone generate $31 of consumer surplus, nearly double the baseline price effect.

Taken together, these results highlight the degree to which cost and quality determine how attractive the public option is to employers. A more attractive public option also provides more price competition in the private market. The reimbursement rates a public option sets are therefore a crucial determinant of the impact the policy will have. The key insight the model provides is that the attractiveness of the public option also depends on how well it deals with frictions in the insurance market. A public option that mimics insurer premium setting by offering discounts and experience rated premiums obtains substantially higher enrollment and larger savings through premium competition. While a public option can overcome adverse selection with a sufficiently large cost advantage, these alternative design choices generate competition across a range of policy parameters.

8 Conclusion

This paper studies the impact of a public option on crowd out and premium competition in the market for employer-sponsored health insurance. I combine panel data on employer-insurer contracting with a structural model of supply and demand to estimate the strength of employers' preferences and associated costs. The key empirical challenges are that premiums in this market are set individually, reflecting differences in employers' preferences and costs. I use a flexible estimation procedure to recover the distribution of employer preferences, and account for insurers' dynamic incentives in recovering the distribution of employers' expected costs. I find that employers have both strong brand preferences and experience large switching costs, and that both of these forces generate market power for insurers. Insurers price discriminate between employers with differing costs, brand preference, as well as between employers who are locked-in and those they wish to compete away from rival insurers.

I find that a public option that does not price discriminate between employers has moderate-to-low take up, even in the case where it perfectly replicates the cost and horizontal quality of the most preferred insurer in the market. Take up is limited by both negative selection into the public option and large switching costs. Allowing the public option to exercise a cost advantage by regulating payments to health care providers allows the public option to attenuate negative selection, but potentially involves distortions to the provider market. A public option only generates premium competition in the private market for employers who are close to the indifference point between the public plan and existing private insurers. Therefore, allowing the public option to offer introductory discounts and experience rated premiums generates additional competition relative to a plan with uniform premium setting.

My results also highlight the importance of insurers' dynamic incentives in setting premiums. I find that insurers set premiums in a forward looking way, which keeps them from fully extracting the rents generated by switching costs. While the model does not explicitly

account for bargaining, my results suggest that dynamic incentives are an important mechanism restraining insurers from fully extracting surplus in setting premiums, despite inelastic demand. Future work should consider the role of switching costs and dynamic incentives as they relate to unexplained bargaining ability in cross sectional studies.

This paper focuses specifically on the large group market in Massachusetts, where employers receive experience rated premium offers. These findings may or may not generalize to other states, where insurers may be more or less differentiated, or employer may have lower switching costs. Self-insured employers contract with the same insurers to provide networks and other administrative services, which represent the most obvious forms of insurer differentiation. However, my findings may not fully generalize to the self-insured market.

Finally, while I do not consider the details of any particular policy proposal, the model provides general insight into the way that employers and insurers would respond to the entry of a public insurance program. The results suggest that a public option that exercises a larger cost advantage will do more to crowd out private insurers. Therefore, the efficiency implications of this crowd out depend on the ex-ante efficiency of the prices insurers' currently negotiate with providers. Cicala et al. (2019) study the introduction of regulation mandating minimum medical loss ratios (MLR) and find that private insurers who operate below the MLR raise claims to comply with regulation – a finding that suggests employers are more elastic to quality than they are to premiums. Additional work should explore the implications of employer preferences and switching costs on insurer investments in quality and cost containment.

References

Abaluck, J. and Gruber, J. (2011). Choice inconsistencies among the elderly: evidence from plan choice in the medicare part d program. *American Economic Review*, 101(4):1180–1210.

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Akerlof, G. A. (1978). The market for "lemons": Quality uncertainty and the market mechanism. In *Uncertainty in economics*, pages 235–251. Elsevier.

Arie, G. and Grieco, P. L. (2014). Who pays for switching costs? *Quantitative Marketing and Economics*, 12(4).

Arnold, D. and Whaley, C. (2020). Who pays for health care costs? the effects of health care prices on wages. *The Effects of Health Care Prices on Wages (July 21, 2020)*.

Baicker, K. and Chandra, A. (2006). The labor market effects of rising health insurance premiums. *Journal of Labor Economics*, 24(3):609–634.

Bajari, P., Fox, J. T., and Ryan, S. P. (2007). Linear regression estimation of discrete choice models with nonparametric distributions of random coefficients. *American Economic Review*, 97(2):459–463.

Beggs, A. and Klemperer, P. (1992). Multi-period competition with switching costs. *Econometrica: Journal of the Econometric Society*, pages 651–666.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pages 841–890.

Blumberg, L. J., Holahan, J., Wengle, E., and Elmendorf, C. (2019). Is there potential for a public option to reduce premiums of competing insurers? *The Urban Institute, Washington DC*.

Bresnahan, T. F. (1981). Departures from marginal-cost pricing in the american automobile industry: Estimates for 1977–1978. *Journal of Econometrics*, 17(2):201–227.

Cabral, L. (2016). Dynamic pricing in customer markets with switching costs. *Review of Economic Dynamics*, 20:43–62.

Cebul, R. D., Rebitzer, J. B., Taylor, L. J., and Votruba, M. E. (2011). Unhealthy insurance markets: Search frictions and the cost and quality of health insurance. *American Economic Review*, 101(5):1842–71.

Chernew, M. E., Hicks, A. L., and Shah, S. A. (2020). Wide state-level variation in commercial health care prices suggests uneven impact of price regulation: An examination of state-level price variation in the commercial market, relative to medicare, for a broader set of states and a wider set of services than had been previously examined. *Health Affairs*, 39(5):791–799.

Cicala, S., Lieber, E. M., and Marone, V. (2019). Regulating markups in us health insurance. *American Economic Journal: Applied Economics*, 11(4):71–104.

Claxton, G., Damico, A., Rae, M., Young, G., McDermott, D., and Whitmore, H. (2020). Health benefits in 2020: Premiums in employer-sponsored plans grow 4 percent; employers consider responses to pandemic: The annual kaiser family foundation employer health benefits survey of the cost and coverage of us employer-sponsored health benefits. *Health Affairs*, 39(11):2018–2028.

Clemens, J. and Gottlieb, J. D. (2017). In the shadow of a giant: Medicare's influence on private physician payments. *Journal of Political Economy*, 125(1):1–39.

Clemens, J., Gottlieb, J. D., and Hicks, J. (2020). How would medicare for all affect health system capacity? evidence from medicare for some.

Cooper, Z., Craig, S. V., Gaynor, M., and Van Reenen, J. (2019). The price ain't right? hospital prices and health spending on the privately insured. *The Quarterly Journal of Economics*, 134(1):51–107.

Craig, S. V., Ericson, K. M., and Starc, A. (2021). How important is price variation between health insurers? *Journal of Health Economics*.

Crawford, G. S. and Yurukoglu, A. (2012). The weflare effects of bundling in multichannel television markets. *American Economic Review*, 102(2):643–685.

Cutler, D. M. and Gruber, J. (1996). Does public insurance crowd out private insurance? *The Quarterly Journal of Economics*, 111(2):391–430.

Cutler, D. M. and Reber, S. J. (1998). Paying for health insurance: the trade-off between competition and adverse selection. *The Quarterly Journal of Economics*, 113(2):433–466.

Dafny, L., Dranove, D., Limbrock, F., and Morton, F. S. (2011). Data impediments to empirical work on health insurance markets. *The BE Journal of Economic Analysis & Policy*, 11(2).

Dafny, L., Duggan, M., and Ramanarayanan, S. (2012). Paying a premium on your premium? consolidation in the us health insurance industry. *American Economic Review*, 102(2):1161–85.

Dafny, L. S. (2010). Are health insurance markets competitive? *American Economic Review*, 100(4):1399–1431.

Dafny, L. S. (2015). *Evaluating the impact of health insurance industry consolidation: Learning from experience*. Commonwealth Fund New York (NY).

Dahl, G. B. and Forbes, S. J. (2014). Doctor switching costs in health insurance. *Unpublished. Frank, Matthew B., John Hsu, Mary Beth Landrum, and Michael E. Chernew*.

Drake, C., Ryan, C., and Dowd, B. E. (2020). Sources of inertia in health plan choice in the individual health insurance market. *Available at SSRN 3556857*.

Dranove, D., Gron, A., and Mazzeo, M. J. (2003). Differentiation and competition in hmo markets. *The Journal of Industrial Economics*, 51(4):433–454.

Dubé, J.-P., Hitsch, G. J., and Rossi, P. E. (2009). Do switching costs make markets less competitive? *Journal of Marketing research*, 46(4):435–445.

Einav, L., Finkelstein, A., and Cullen, M. R. (2010). Estimating welfare in insurance markets using variation in prices. *Quarterly Journal of Economics*.

Ericson, K. (2014). Consumer inertia and firm pricing in the medicare part d prescription drug insurance exchange. *American Economic Journal: Economic Policy*, 6(1):38–64.

Farrell, J. and Klemperer, P. (2007). Coordination and lock-in: Competition with switching costs and network effects. *Handbook of industrial organization*, 3:1967–2072.

Farrell, J. and Shapiro, C. (1988). Dynamic competition with switching costs. *The RAND Journal of Economics*, pages 123–137.

Fiedler, M. (2020). Capping prices or creating a public option: How would they change what we pay for health care? Technical report, USC-Brookings Schaeffer Initiative for Health Policy.

Fox, J. T., Kim, K. I., Ryan, S. P., and Bajari, P. (2011). A simple estimator for the distribution of random coefficients. *Quantitative Economics*, 2(3):381–418.

Garthwaite, C., Ody, C., and Starc, A. (2020). Endogenous quality investments in the us hospital market. Technical report, National Bureau of Economic Research.

Gowrisankaran, G., Nevo, A., and Town, R. (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review*, 105(1):172–203.

Gruber, J. and McKnight, R. (2016). Controlling health care costs through limited network insurance plans: Evidence from massachusetts state employees. *American Economic Journal: Economic Policy*, 8(2):219–50.

Guardado, J. R., Emmons, D. W., and Kane, C. K. (2013). The price effects of a large merger of health insurers: a case study of unitedhealth-sierra. *Health Management, Policy and Innovation*, 1(3):16–35.

Hacker, J. S. (2008). The case for public plan choice in national health reform: Key to cost control and quality coverage. *Institute for America's Future*.

Halliday, T. J. (2007). Testing for state dependence with time-variant transition probabilities. *Econometric Reviews*, 26(6):685–703.

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 103(7):2643–82.

Handel, B. R. and Kolstad, J. T. (2015). Health insurance for" humans": Information frictions, plan choice, and consumer welfare. *American Economic Review*, 105(8):2449–2500.

Ho, K., Hogan, J., and Scott Morton, F. (2017). The impact of consumer inattention on insurer pricing in the medicare part d program. *The RAND Journal of Economics*, 48(4).

Ho, K. and Lee, R. S. (2017). Insurer competition in health care markets. *Econometrica*, 85(2):379–417.

Hoffman, A. K., Jackson, H. E., and Monahan, A. (2021). A public option for employer health plans. *Public Law and Legal Theory Research Paper Series*, 21-12.

Honoré, B. E. and Kyriazidou, E. (2000). Panel data discrete choice models with lagged dependent variables. *Econometrica*, 68(4):839–874.

Ketcham, Jonathan D., N. V. K. and Powers, C. A. (2016). Choice inconsistencies among the elderly: Evidence from plan choice in the medicare part d program: Comment. *American Economic Review*, 106(12):3932–61.

Klemperer, P. (1995). Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade. *The review of economic studies*, 62(4):515–539.

Lin, H. and Zhou, Y. (2021). Distributional effect of insurer concentration on premiums: Evidence from the employer-sponsored health insurance market. *Working Paper*.

Maeda, J. L. K. and Nelson, L. (2018). How do the hospital prices paid by medicare advantage plans and commercial plans compare with medicare fee-for-service prices? *INQUIRY: The Journal of Health Care Organization, Provision, and Financing*, 55:0046958018779654.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342.

Nevo, A., Turner, J. L., and Williams, J. W. (2016). Usage-based pricing and demand for residential broadband. *Econometrica*, 84(2):411–443.

Panhans, M., Rosenbaum, T., and Wilson, N. E. (2018). Prices for medical services vary within hospitals, but vary more across them. *Medical Care Research and Review*.

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

Scott, D. (2021). The public option is now a reality in 3 states. *Vox.com*.

Shepard, M., Baicker, K., and Skinner, J. (2020). Does one medicare fit all? the economics of uniform health insurance benefits. *Tax Policy and the Economy*, 34(1):1–41.

Thisse, J.-F. and Vives, X. (1988). On the strategic choice of spatial price policy. *The American Economic Review*, pages 122–137.

Tilipman, N. (2018). Narrow physician networks, switching costs, and product variety in employer markets.".

Train, K. (2016). Mixed logit with a flexible mixing distribution. *Journal of choice modelling*, 19:40–53.

Trish, E. E. and Herring, B. J. (2015). How do health insurer market concentration and bargaining power with hospitals affect health insurance premiums? *Journal of health economics*, 42:104–114.

APPENDICES

A Price Discrimination and Adverse Selection

Here, I provide a simple model of first-degree price discrimination to provide intuition for the conditions under which an employer would choose a public option, and under which a public option generates additional price competition in a market where private insurers can flexibly set premiums. Under a Thisse and Vives (1988) style model of price discrimination, insurers would set premiums so that an employers' most preferred insurer extracts the full social surplus relative to the employer's next best option. A given employer i, with willingness to pay θ_j for each of the insurers $j \in \mathcal{J}$, experiences utility according to $u_j = \theta_j - p_j$, and chooses an insurer according to

$$d_j = \mathbb{1}\{u_j \geq u_k, \forall j, k \in \mathcal{J}\}. \tag{17}$$

In the event of a tie, the employer chooses the insurer that maximizes the joint surplus $(\theta_j - c_j)$. The insurer with the highest joint surplus wins the sale by setting premiums

$$\begin{aligned} p_1^* &= c_1 + [(\theta_1 - c_1) - (\theta_2 - c_2)] \\ &= \theta_1 - (\theta_2 - c_2), \end{aligned} \tag{18}$$

and all other insurers, $j \geq 2$, set premiums according to $p_j^* = c_j$. The employer always chooses the insurer with the highest joint surplus, and that insurer fully extracts its cost-adjusted quality advantage over the next best alternative.

To understand the impact a public option would have in this setting, I assume that a the new plan does not charge a markup and instead sets premiums to cover average costs, but that it cannot price discriminate. Therefore, if each employer in the market costs $c_{PO(i)}$ to insure under the public option, the premium for the public plan is

$$\overline{p}_{PO} = \frac{\sum_{i \in PO} N_i \cdot c_{PO(i)}}{\sum_{i \in PO} N_i}, \tag{19}$$

where each employer has its own utility over the public option $\theta_{PO(i)}$, so that the set of employers who choose the public option ($i \in PO$) is determined in equilibrium by whether or not $\theta_{PO(i)} - \overline{p}_{PO}$ is higher or lower than the utility each employer can obtain from the private insurers, who continue to price strategically.

For a given employer, introducing a non-strategic public option in this setup can result in three distinct outcomes depending on the cost and quality of the public plan. In the following results, I add (i) to the subscripts of the primitives above to emphasize that insurers compete in a separate game for each employer.

First, if the public option provides lower utility than the joint surplus of the second-best insurer $(\theta_2 - c_2)$, then equilibrium is completely unchanged. That is, if the public option has quality that is too low, or the premium for the public option exceeds the cost of insuring employer i for all relevant competitors in the market, then it will not present a desirable outside option to the employer.

Second, if the employer obtains a higher net utility from the public option than insurer 1 could possibly offer with positive profits $(\theta_{PO(i)} - \overline{p}_{PO}) > \theta_{1(i)} - c_{1(i)})$, then the employer will abandon the private market in favor of the public option. A public option with sufficiently

high quality or low costs will crowd-out the private market for consumer i.

In the third case, the public option presents the consumer with net utility that is between $(\theta_{1(i)} - c_{1(i)})$ and $(\theta_{2(i)} - c_{2(i)})$, in which case the second-best insurer no longer competes for the employer and insurer 1 prices as a monopolist relative to the employer's new outside option

$$p^*_{1M(i)} = \theta_{1(i)} - (\theta_{PO(i)} - \overline{p}_{PO}). \tag{20}$$

In this third case, the public option generates additional price competition by reducing the employer's excess demand relative to the outside option. The employer continues to purchase from the insurer with the highest joint surplus, but it obtains a lower premium.

This simplified model helps clarify the conditions under which a public option will either supplant private insurers or cause them to offer lower premiums. The results also highlight the scope for adverse selection to limit the role of a public option. If employers who choose the public option are more costly to insure, then $\overline{p}_{PO}$ will be higher, and therefore less likely to be a desirable first- or second-best choice for other employers. Because of the public option sets premiums according to average cost, the impact of such a policy will depend on the joint distribution of θ and c across employers. Holding preferences constant, the magnitude of cost dispersion across employers relative to the size of existing markups will determine the set of employers for whom it is optimal to chose the public option, and the set of employers who will experience additional price competition. On the other hand, selection issues could be limited if lower cost employers have a larger gap between $\theta_{1(i)}$ and $\theta_{2(i)}$.

B The Form 5500

B.1 Identifying Insurer Groups

While the Form 5500 contains information about insurers, that information contains measurement error. Primarily this comes from two sources. First, insurer IDs are recorded using the National Association of Insurance Commissioners' company codes. However, these codes have a many-to-one mapping to insurance carriers. Second, forms submitted prior to 2009 were not submitted electronically. Available data from this time period was digitized, resulting in transcription errors. The data also include a text field containing the name of the carrier. This field is also subject to transcription errors, but provides an additional source of information to verify insurer identifiers. In order to properly group codes into insurance carriers, I implement the following procedure:

1. Match insurers to NAIC group codes using their NAIC company code and directories published by NAIC for 2004-2018
2. Based on insurer names of successful matches, perform string matching
3. Reconcile conflicts from steps (1) and (2)
4. I perform a small number of manual assignments based on obvious mappings missed by string matching (e.g. "Connecticut General" and "Cigna" do not pass string matching, but are in fact the same company).

B.2 Measuring Employer Size

In order to construct a standardized measure of premiums, I must construct a measure of employer size that will allow me to express premiums in terms of dollars per covered life year. Firms can fluctuate substantially in size from 2001-2018. While the cleanest approach would be to focus on firms that remain a consistent size over the study period, such a restriction would severely limit the size and generalizability of the sample.

The combined 5500 base form and Schedule A contain three separate variables that serve to measure the number of plan participants. These variables are the most consistently reported, and most closely capture the relevant employer size for the purpose of standardizing premiums. I denote them here as:

- N^{boy}: Line 6a(1) of the Base Form 5500 reports the total number of active participants at the beginning of the plan year. This number reflects the number of enrolled employees across all schedule A filings attached to a given base form.
- N^{eoy}: Line 6a(2) of the Base Form 5500 reports the total number of active participants at the end of the plan year. This number reflects the number of enrolled employees across all schedule A filings attached to a given base form.
- N^{cov}: Line 1e of the Schedule A reports the "approximate number of persons covered at end of policy or contract year."

N^{cov} is the measure that closely captures the relevant employer size for computing average per-member premiums. However, it is clear from the data and documentation that reporting is inconsistent. In a preparer's manual for the 5500, Fisher and Anderson note that

> *"The DOL says dependents should be included in the count reported on line 1e* (name of *Ncov* variable)*, although whether dependents are include or excluded in the data provided by an insurance company varies depending on the carrier's own internal procedures. Generally, preparers simply use the information provided by the insurance company without further analysis. Dependents are not counted for any other purposes on the Form 5500 or its schedules"* (text in parentheses mine).

Patterns in the data indicate that not only does reporting vary across insurers, but also within insurer-employer dyads over time.

A straightforward solution would be to simply rely on the N^{boy} and N^{eoy} variables. There are two issues with this approach. First, employers often report several policies as part of a "wrap plan," including separate contracts for life insurance, stand-alone dental plans, and other benefits. In these cases, the correspondence of N^{boy} and N^{eoy} will depend greatly on the existence and relative size of other plans offered by the employer. Second, even restricting to health plans where N^{cov} and N^{eoy} align closely will generally bias premium measures upward across all plans.

To solve this issue, I implement the following approach.

1. Identify observations of N^{cov} that definitely represent enrollment because $N^{cov} >> N^{act}$. In this step, I also manually reviewed cases representing outliers or marginal cases. I exclude 86 employers for whom enrollment was not sufficiently clear in all years of their panel.

2. For observations where $N^{cov} >> N^{act}$, I calculate the ratio of N^{cov}/N^{act} and apply the nearest ratio (in time) to adjust each non-enrollment observation.

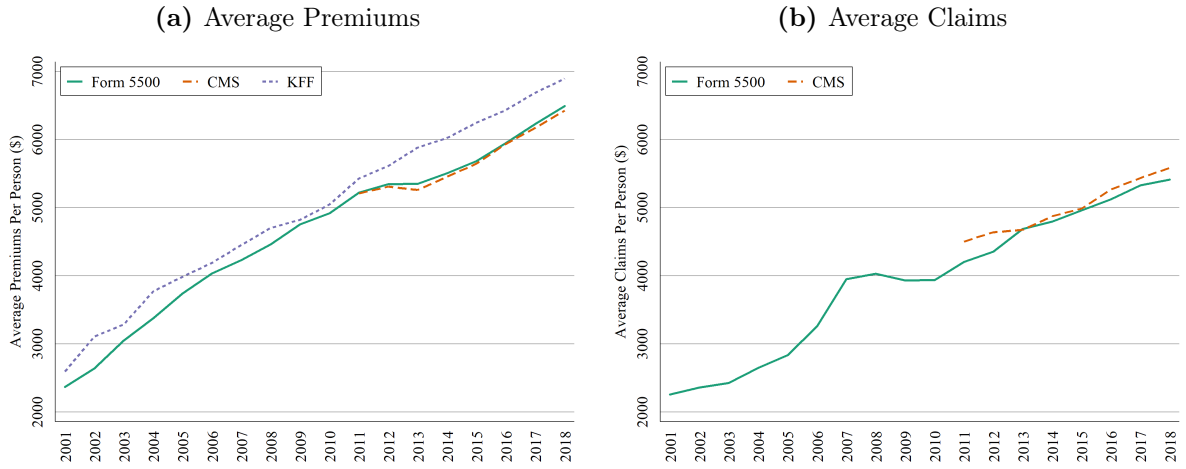
   Additional detail regarding premium measurement are included in the following section.

B.3 Benchmarking the Form 5500

In order to assess the validity of the measures from the Form 5500, I compare premiums, claims, and market shares to other published sources. I rely heavily on published data from CMS's Medical Loss Ratio (MLR) Filings[12]. These data are more representative in the sense that they are based on each insurers' entire book of business in the large group market. However, they are aggregated to the insurer-year level and would not be sufficient to perform the analysis in the main text. The CMS data were collected to enforce a provision of the Affordable Care Act, and are therefore only available from 2011-onward. For early years of the data, during which the CMS data are unavailable, I use the the Kaiser Family Foundation's Employer Health Benefits Survey (KFF-EHBS). The KFF-EHBS is a nationally representative survey of employers and is available for 2001-2018. However, it does not contain geographic identifiers and may be limited in its ability to capture Massachusetts specific trends in premiums

[12] https://www.cms.gov/CCIIO/Resources/Data-Resources/mlr

Figure A.1: Premiums and Claims in the Form 5500 and Other Data Sources



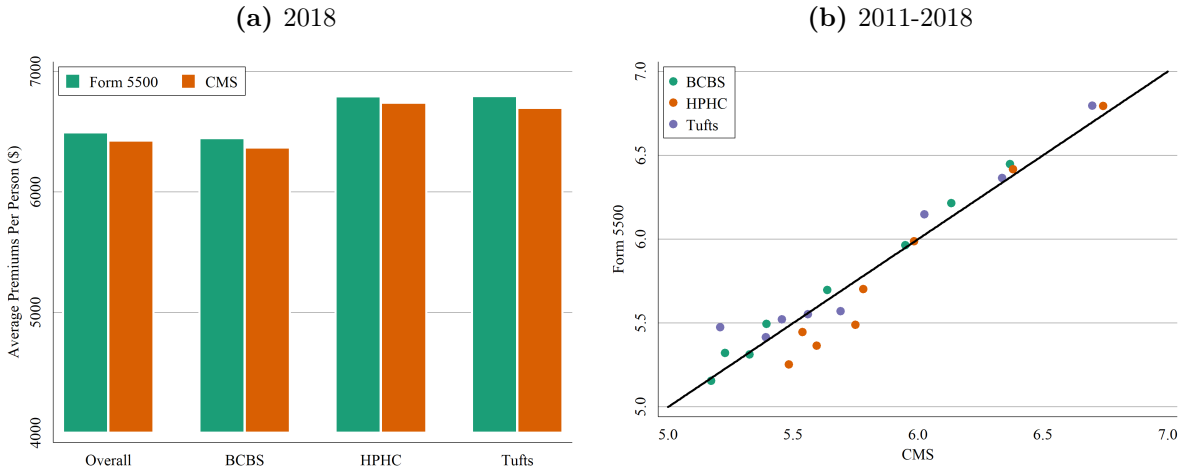
Notes: Panel (a) presents adjusted premiums per person from the final series, along with data from the CMS MLR filings for Massachusetts, and national average premiums for the KFF-EHBS. Data from the Form 5500 of weighted by the adjusted number covered lives as described in B.2. Data from the KFF-EHBS are weighted to reflect each employers' number of covered workers. Panel (b) presents average claims per person for the Form 5500 and the CMS MLR filings.

Panel (a) of Figure A.1 shows the measure of average premiums from the Form 5500 from 2011-2018. Comparable measures from CMS are available through MLR filings for 2011-2017. Data from the KFF-EHBS represent the national average and show that in early years, the Form 5500 produces premiums that trend consistently with the national average for each year. Panel (b) presents the comparable exercise for average claims per person. Claims are unavailable in the KFF-EHBS, and the figure therefore only contains a comparison for the CMS MLR filings. The claims measure is not available for all employers in all years. As a result, it is substantially noisier than the premium measure, particularly in early years of the data. When available, the CMS data report average claims that are similar to those in the 5500, though I interpret them with more caution than the premium data.

In Figure A.2, I present average premiums across the three main insurers in Massachusetts. The bars in Panel (a) show the average premiums for each payer in 2018. Premiums in the Form 5500 are slightly higher on average than those in the CMS MLR filings. However, average premiums are similar in both data sources, and the relative differences across payers are common across data sources. Panel (b) shows the relationship between average premiums from CMS and those derived from the Form 5500 across across all years for which both data sources are available. Premium measures from both data sources are highly correlated, suggesting that the Form 5500 is accurately respecting the within between insurer and over-time differences in premiums observed in the CMS data. The extent to which the data do differ likely reflects the fact that the Form 5500 only captures the subset of employers who qualify based on the 100 employee filing threshold, comply with the filing requirement, and meet the sample restrictions discussed in Section 3.

In Figure A.3, I compare market shares of the top 3 insurers. Panel (a) contains market shares for 2018. While not exact, the Form 5500 contains employers who represent

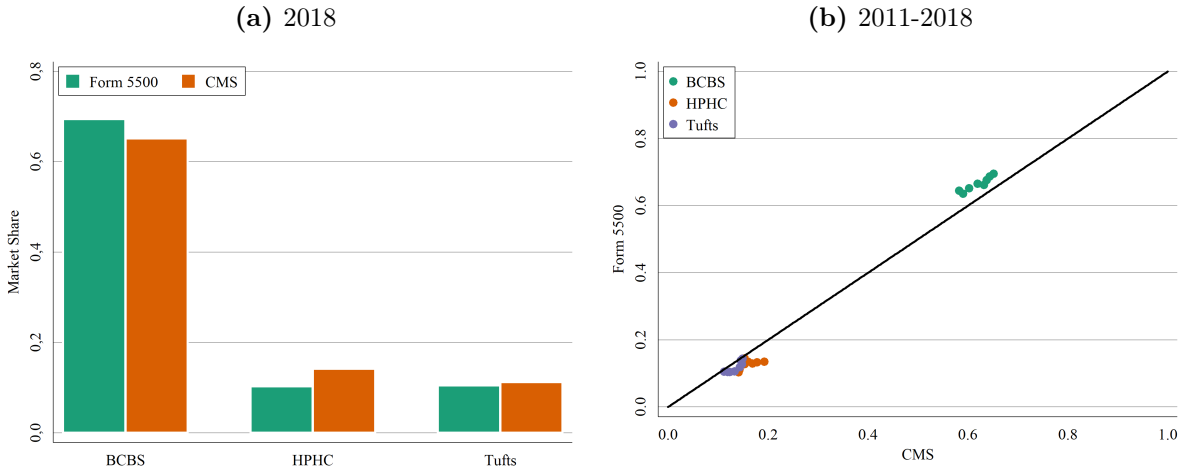
Figure A.2: Average Premiums by Insurer



Notes: Panel (a) presents average adjusted premiums per person for 2018 from both the Form 5500 and the CMS MLR filings. Panel (b) presents the correlation of average premiums from both data sources, within and across payer, for 2011-2018. The diagonal line in Panel (b) is a 45-degree line, indicating the points at which measures from both data sources would match exactly. Data from the Form 5500 of weighted by the adjusted number covered lives as described in B.2.

the distribution of choices for the overall market. Employers in the Form 5500 tend to slightly prefer BCBS to HPHC relative to the market average, potentially reflecting differences in employer size or the composition of employers who purchase multiple plans and are therefore excluded from the Form 5500 sample. Panel (b) displays the same measure plotted for 2011-2018, which indicates that these patterns are consistent across all years for which both data sources are available.

Figure A.3: Market Shares by Insurer



Notes: Panel (a) presents market shares for 2018 from both the Form 5500 and the CMS MLR filings. Panel (b) presents the correlation of market shares from both data sources, within and across payer, for 2011-2018. The diagonal line in Panel (b) is a 45-degree line, indicating the points at which measures from both data sources would match exactly. Data from the Form 5500 of weighted by the adjusted number covered lives as described in B.2.

C Generalizability

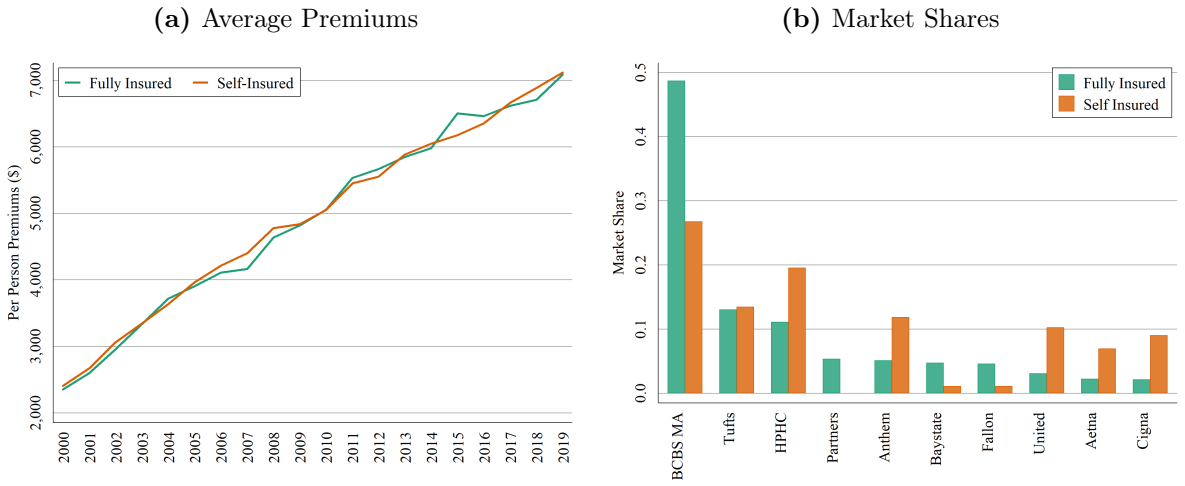
In this section, I present a number of exercises designed to assess the generalizability of the sample used for analysis.

C.1 Self-Insurance and Firm Size

As discussed in Section 3, the employer panel I use does not contain self-insured employers. Self-insured employers are required to submit a Form 5500 filing, though they do so via financial information about the trust through which benefits or third-party administrators are paid. This typically makes it difficult to identify the insurer who provides their administrative services, as the administrative fees are sometimes paid out of the general assets of the employer rather than the trust. Self-insured employers also tend to be disproportionately large firms with many locations, all of whom may file the 5500 under a common tax identifier. These filings make it difficult or impossible to determine which of the insurance contracts correspond with different local markets. Although this feature of the data potentially limits my ability to speak to the full landscape of employer-sponsored insurance, there are several reasons to expect that my findings will generalize to larger self-insured firms.

Both the unavailability of data on self-insured employers and the restrictions on multi-state employers have the potential to bias the results of the counterfactuals. If large and/or self-insured employers have drastically different underlying costs and preferences, the counterfactuals will less accurately reflect the true equilibrium premium the public option would offer.

Figure A.4: Premiums and Market Shares in the Self-Insured Market



Notes: Panel A shows average premiums for self-insured and fully-insured plans. Data come from the KFF Employer Health Benefits Survey. All estimates calculated using the provided employer weights. Panel B shows the market shares of insurers in the fully insured market and the market for self-insured (ASO), using data from HealthLeaders Interstudy.

Prior work has found that levels and trends in premiums match closely between the fully and self-insured markets, and self-insured employers typically contract with the same

Table A.1: Between Employer Premium Variation by Employer Size

	Number of Employees			
	3-49	50-199	200-999	1,000-4,999
Panel A: All Employers				
Mean Premiums	7,390	7,085	7,346	7,392
SD	(2,560)	(2,241)	(2,111)	(1,788)
SD/Mean	0.346	0.316	0.287	0.242
N	270	265	442	688
Panel B: Fully Insured				
Mean Premiums	7,560	7,008	6,982	7,404
SD	(2,552)	(2,257)	(1,987)	(1,876)
SD/Mean	0.337	0.322	0.285	0.253
N	231	201	207	104
Panel C: Self Insured				
Mean Premiums	5,901	7,450	7,712	7,389
SD	(2,140)	(2,146)	(2,173)	(1,768)
SD/Mean	0.363	0.288	0.282	0.239
N	39	64	235	584

Notes: Calculations based on the 2019 KFF-EHBS. All calculations made using the provided survey weights.

insurers to provide administrative services only (ASO) contracts (Altman 2020). In Panel (a) of figure A.4, I reproduce analysis from Altman (2020), using data from the KFF-EHBS to plot average premiums between self and fully insured employers. Levels and trends of premiums in the self-insured market are broadly similar and have followed similar trends over the past two decades.

Self-insured employers bear the risk of underlying claims, however, they contract with health insurers to provide networks, utilization management, and claims processing. Panel (b) of figure A.4 shows the market shares of health insurers in the Massachusetts fully insured and ASO markets using data from HealthLeaders Interstudy. National insurers tend to play a larger role in the ASO market in Massachusetts. However, in at least some cases national insurers who wish to operate in the state do so by renting their network from local players (Wojcik 2004). Since the network is a, – if not the – primary source of product differentiation across insurers, I interpret this as evidence that employers in the self-insured market face the same market for differentiated services as those in the large group market.

The primary concern for the counterfactuals is whether between-insurer variation in costs is substantially different across these groups. While cost is not readily observable, the KFF-EHBS contains measures of premiums, firm size, and insurance type for a nationally representative sample of firms. Table A.1 displays the mean, standard deviation, and coefficient of variation in premiums for fully and self-insured firms across the distribution of firm size. Nearly all subgroups appear to have similar patterns in terms of premium levels of variation across firms. Small (3-49 employee) self-insured firms appear to be an outlier. Because these firms operate in the small group market, where premiums are community rated,

this could be the result of selection into self-insurance by firms with below average costs. However, I also note that the number of firms in this cell is small and has the potential to be less representative.

C.2 Restriction on Estimation Sample

In addition to the cleaning rules discussed in Section 3 and Appendix B, I apply two restrictions to the estimation sample, which are required for identification of the demand model. First, I restricted to employers who have a minimum panel length of three years. Second, I condition on employers who switch at least once. These restrictions reduce the sample substantially. Employer panels may be incomplete due to entry, attrition, or a variety of data quality issues. In this section, I perform a number of exercises aimed at assessing the generalizability concerns associated with these restrictions.

First, Table A.2 shows the count of employers remaining after each of the restrictions. The data contain limited direct information on characteristics of the employers. However, these restrictions do not substantially alter the distribution of employer size. The final sample is more likely to choose HPHC and less likely to choose BCBS than the baseline sample. Since BCBS is the most preferred insurer, there will be more choice sequences of three consecutive years with BCBS than the other insurers. Therefore, conditioning on a switch fundamentally lowers the probability the expected market share for BCBS.

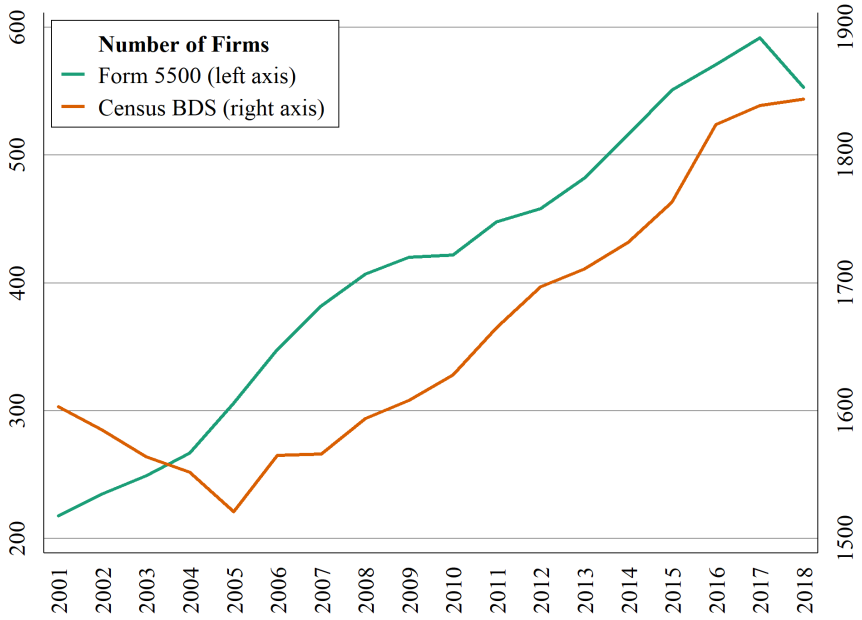
Table A.2: Estimation Restrictions and Sample Size

	N	Average Employer Group Size	Share BCBS	Share HPHC	Share Tufts
Baseline Sample	1,051	468	0.680	0.112	0.119
Has 3+ Years	744	479	0.583	0.156	0.154
Ever Switches	605	474	0.590	0.152	0.150
Both	384	475	0.488	0.198	0.195

Second, to analyze the role of entry in determining panel length, I compare trends in the number of firms who file a 5500 to the subset of firms who are likely to fully insure through the large group market *and* meet the filing threshold of 100 employees. I benchmark these numbers using the Census Bureau's Business Dynamic Statistics (BDS) data focusing on the number of firms in Massachusetts with 100-999 employees – the category most of firms that are are likely to participate in the large group market. In both datasets, the number of firms the number of firms who are likely to purchase fully insured products, and who meet the filing thresholds for the 5500 follow similar trends. The 5500 consistently contains approximately one third the number of firms in this category.

Third, because switches occur infrequently, there is a strong relationship between the probability that a firm switches and the length of time they are observed in the data. Figure

Figure A.5: Count of firms in the 5500 data



Notes Figures show number of firms in the 5500 and the number of firms in Massachusetts with 100-1,000 employees based on the Census Bureau's Business Dynamic Statistics (BDS).

A.6 shows the relationship between the probability of observing a switch for an employer and the length of their panel in the data. Longer panels provide more opportunity to observe a switch. Therefore, employers who are not ever observed making a switch, may simple be in the data for shorter amounts of time, rather than this condition reflecting differences in the employers' underling preferences. In order to evaluate this hypothesis, I estimate a regression to predict attrition as a function of whether or not a switch is ever observed in the data. Because observing a switch is more likely the longer an employer remains in the data, I include fixed effects for the year in which the employer entered the data and the year of the data. Column (1) presents the baseline relationship between whether an employer has previously switched insurers by time t, and an indicator for whether the has exited the data by year t. Column (2) includes fixed effects for the year and cohort-year of the employer. Column (3), which includes an interaction between year and cohort-year most flexibly controls for the increased likelihood of observing a switch in longer panels. In both Columns (2) and (3), I find no evidence that switchers are more likely to attrit from the data.

Figure A.6: Probability of observing a switch and panel length

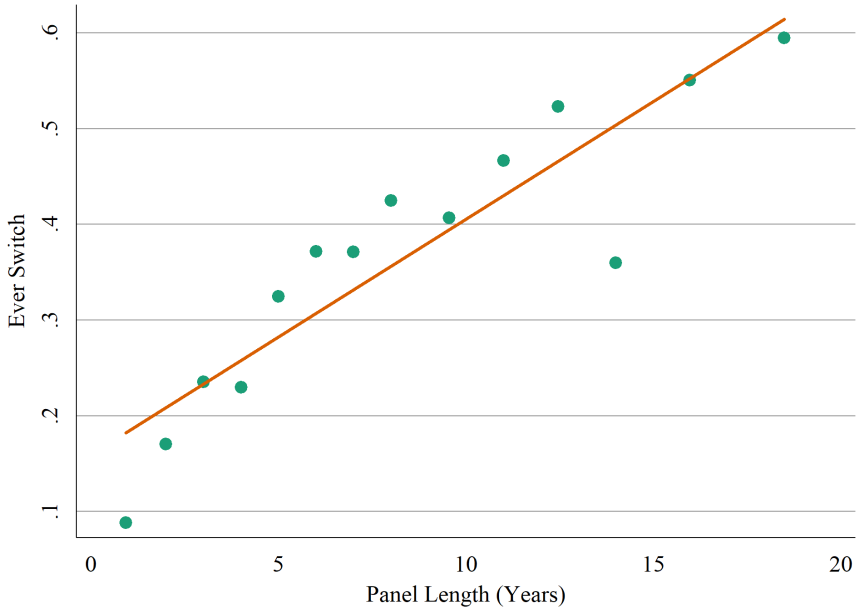


Table A.3: Switching Behavior and Attrition

Dependent Variable:	Out of the Data in Year t		
Has Switched Before t	0.167***	0.008	-0.004
	(0.025)	(0.029)	(0.029)
Fixed Effects	None	Cohort + Year	Cohort X Year
N	21,220	21,220	21,220
N_i	1,051	1,051	1,051

D Testing for State-Dependence

D.1 Properties of Choice Sequences

I explore the panel structure of the data to provide evidence of switching costs using a semi-parametric approach. Consider an employer i with random utility over $j \in \mathcal{J}$ insurers. Any sequence of three choices can be partitioned into five mutually exclusive, collectively exhaustive categories:

- $A1$: Stay, Stay
- $A2$: Stay, Switch
- $A3$: Switch, Stay
- $A4$: Switch, Return
- $A5$: Switch, Switch

These categories provide a way to assess the degree of state-dependence in the data, by comparing the prevalence of sequences containing the exact same *set* of choices, but different numbers of switches.

To illustrate the logic, consider an employer who enters the market and chooses between two insurers. Using superscript $\hat{h} = \{h, -h\}$ to denote whether insurer j is the "home" insurer (the one i chose in the prior period).

$$\begin{aligned} s_{ij}^{h} &> s_{ij}^{-h}, \text{ when there are switching costs and} \\ s_{ij}^{h} &= s_{ij}^{-h}, \text{ when there are no switching costs.} \end{aligned}$$

Each probability $\mathbb{P}_i(A2)$ and $\mathbb{P}_i(A4)$ can be written as the sum of sequence probabilities that would result in each type of sequence.

$$\begin{aligned} \rho_i &= \mathbb{P}_i(A4)/\mathbb{P}_i(A2) \\ &= \left(\frac{s_{i1}^{h} s_{i2}^{-h} s_{i1}^{-h} + s_{i2}^{-h} s_{i1}^{-h} s_{i2}^{-h}}{s_{i1}^{h} s_{i1}^{h} s_{i2}^{-h} + s_{i2}^{-h} s_{i2}^{h} s_{i1}^{-h}} \right), \end{aligned} \tag{21}$$

As long as an employer's preferences for insurers are constant over time, ρ_i will be exactly one in the absence of switching costs, and strictly less than one when there are switching costs. This result is not sensitive to initial conditions, and can be performed using the full set of consecutive choices in the data, regardless of choices preceding each sequence. The sequence probabilities for a consumer who chose insurer j prior to the sequence are:

$$\begin{aligned} \mathbb{P}_i(A2) &= s_{ij}^{h} s_{ij}^{h} s_{ik}^{-h} + s_{ik}^{-h} s_{ik}^{h} s_{ij}^{-h} \\ \mathbb{P}_i(A4) &= s_{ij}^{h} s_{ik}^{-h} s_{ij}^{-h} + s_{ik}^{-h} s_{ij}^{-h} s_{ik}^{-h}. \end{aligned}$$

The sequence probabilities for a consumer who chose insurer 2 prior to the sequence can be written symmetrically. The key insight is that each of these sequences contains the same

set of products, but the sequences that contribute to $A4$ always involve an additional switch relative to their brand-matched counterparts in $A2$. In the absence of switching costs, the sequence probabilities will be the same, and switching costs will always tend to drive ρ below 1. This property also holds for any pair of two insurers, and therefore accommodates a choice set of any size because a ρ_i with a larger choice set simply involves summing all of the equivalent terms for each pairwise combination of jk. In the case of three insurers:

$$\rho_i = \frac{s_{ij}^{-h} s_{ik}^{-h}(s_{ij}^{-h} + s_{ik}^{-h}) + s_{ij}^{-h} s_{il}^{-h}(s_{ij}^{-h} + s_{il}^{-h}) + s_{ik}^{-h} s_{il}^{-h}(s_{ik}^{-h} + s_{il}^{-h})}{s_{ij}^{-h} s_{ik}^{-h}(s_{ij}^{h} + s_{ik}^{h}) + s_{ij}^{-h} s_{il}^{-h}(s_{ij}^{h} + s_{il}^{h}) + s_{ik}^{-h} s_{il}^{-h}(s_{ik}^{h} + s_{il}^{h})}$$

In the population, the estimated ratio will be the average of ρ_i over the distribution of (θ, η). With no state-dependence, $\rho_i = 1$ for all individuals, and therefore the population average of ρ will be 1, regardless of the underlying distribution of time invariant brand preferences. On the other hand, the presence of positive switching costs in at least some part of the population implies that $\rho_i \leq 1$, for all consumers. Since the switching costs drive ρ_i below 1 for any individual, the population average of ρ, forms a test that is sufficient to demonstrate the presence of switching costs. The intuition is the same as adding more products. The population average ρ is simply involves adding more terms of equation (21) to the numerator and denominator.

$$\begin{aligned}\rho &= \mathbb{P}(A4)/\mathbb{P}(A2) \\ &= \frac{\sum_i \sum_{j\in\mathcal{J}} \sum_{k\in\mathcal{J}} s_{ij} s_{ij} s_{ik}}{\sum_i \sum_{j\in\mathcal{J}} \sum_{k\in\mathcal{J}} s_{ij} s_{ik} s_{ij}},\end{aligned}$$

which makes it easy to discern that integrating over the population does not reduce the strength of the test. Without switching costs, these terms all contribute the same quantity to the numerator and denominator. The presence of switching costs will add smaller terms to the numerator than the denominator.

D.2 Results

In this section, I provide estimates of ρ using the Form 5500 employer data. I focus on sequences of three consecutive choices in the data. With this restriction, assuming stable preferences at the individual level implies that the underlying distribution of preferences in the population is also fixed.

The main result is an estimated ρ using the full set of three-year sequences, and is reported in the first two rows of Table A.4. The ratio of sequence probabilities imply that employers are over 18 times more likely to choose a sequence containing one fewer switch, holding the set of chosen products fixed. Standard errors are generated using bootstrapping over $1,000$ re-samples.

I repeat this exercise on subsets of the data, based on the maturity of the employers' health plan. "New Plans" are three-period sequences where the first period is the first observation for a given employer, and the "plan effective date" reveals that it was the first year that the employer offered a health plan. The results using "Mature Plans" relies only on the complement of these observations. The estimated ρ for new plans is 0.2, which is still

Table A.4: Semi-parametric Tests of State Dependence in the Employer Panel

	$\left(\frac{\mathbb{P}(A4)}{\mathbb{P}(A2)}\right)$ ρ	$\mathbb{P}(A1)$ Stay, Stay	$\mathbb{P}(A2)$ Stay, Switch	$\mathbb{P}(A3)$ Switch, Stay	$\mathbb{P}(A4)$ Switch, Return	$\mathbb{P}(A5)$ Switch, Switch
Full Sample	0.053 (0.034, 0.075)	0.837 (0.828, 0.846)	0.075 (0.068, 0.082)	0.079 (0.073, 0.086)	0.004 (0.002, 0.006)	0.004 (0.003, 0.006)
New Plans	0.200 (0.048, 0.404)	0.803 (0.766, 0.845)	0.086 (0.059, 0.114)	0.083 (0.059, 0.110)	0.017 (0.003, 0.031)	0.010 (0.003, 0.021)
Mature Plans	0.042 (0.024, 0.063)	0.840 (0.831, 0.850)	0.074 (0.067, 0.081)	0.079 (0.071, 0.086)	0.003 (0.002, 0.005)	0.004 (0.002, 0.005)
Early Years: 2001-2003	0.097 (0.000, 0.211)	0.826 (0.795, 0.857)	0.079 (0.059, 0.102)	0.082 (0.059, 0.105)	0.008 (0.000, 0.015)	0.005 (0.000, 0.010)
Late Years: 2014-2016	0.051 (0.014, 0.102)	0.851 (0.833, 0.867)	0.068 (0.055, 0.081)	0.072 (0.059, 0.084)	0.003 (0.001, 0.007)	0.006 (0.003, 0.009)
First Observed Sequence	0.108 (0.046, 0.193)	0.808 (0.786, 0.831)	0.080 (0.066, 0.098)	0.096 (0.080, 0.112)	0.009 (0.004, 0.015)	0.006 (0.002, 0.011)
Randomly Chosen Sequence	0.048 (0.000, 0.102)	0.843 (0.822, 0.865)	0.077 (0.062, 0.093)	0.073 (0.058, 0.088)	0.004 (0.000, 0.007)	0.004 (0.000, 0.007)

below 1 but substantially higher than estimates generated from other samples.[13] While it is tempting to conclude that new plans are more likely to reverse a switch (choose $A4$) than mature plans, none of these estimates are significantly different from each other.

I also produce estimates over different time periods to determine whether the results are driven by something that is time varying. Because the distribution of preferences is fixed in all exercises, this result is mainly informative regarding the degree to which the evidence of switching costs is disproportionately driven by something that is changing over time. I estimate ρ using a sample of sequences from 2001-2003 and 2014-2016, which provide estimates that are similar to both each other and the main estimates.

Finally, using the full sample of possible sequences produces observations relying on overlapping observations. A reasonable concern is the degree to which this impacts the

[13]Switching costs could be lower for new employers if, for example, new plans receive a value from searching. Such a pattern would suggest that $\rho > 1$ in the absence of switching costs.

expectation of observed sequences, since these are not independent observations. While I have not fully explored the analytic properties of this issue, I implement two tests specifically designed to address it. I measure ρ over the first three observations for each employer. For employers who offer a health plan for the first time during the sample period, this observation will be the same one underlying the estimate for new plans. For other employers, this will be the sequence beginning in the first three-year sequence I observe for each employer. I also produce an estimate using one sequence per employer, chosen at random. Both of these exercises produce measures that are qualitatively similar and statistically indistinguishable from the main results.

Finally, this exercise relies only on the observed choice probabilities and ignores the contribution of prices. In general, both the empirical exercises in Section **??** and the model in Section 4 indicate that prices are weakly lower for choices that result in a switch because of the dynamics of the invest-harvest motives. Lower prices from switching imply that there may be situations where $s_{ij}^{-h} > s_{ij}^{h}$. In this case, the results provide a lower bound on the degree to which state dependence drives the observed choices.

D.3 Numerical Illustrations

This section provides numerical simulations to illustrate the behavior of ρ, and the interaction between switching costs and stable brand-preference. I simulate choice scenarios with varying degrees of heterogeneity in preferences over products and switching costs. Utilities are given by

$$u_{ij}^{h} = \theta_{ij} + \eta_i + \epsilon_{ij} \tag{22}$$

$$u_{ij}^{-h} = \theta_{ij} + \epsilon_{ij}, \tag{23}$$

θ_{ij} is a persistent taste shock for consumer i and product j. η_i is a persistent switching cost for consumer i, which experienced only if the consumer stays with a previously chosen product, and ϵ_{ij} is an iid logit shock.

I first, consider a set of homogeneous set of consumers with identical brand preferences:

$$\theta^{HOM} = (2, 2, 4)$$

Next, I consider a distribution of brand preferences, which are distributed normally around the same means in θ_1, this time each having a standard deviation of 1.

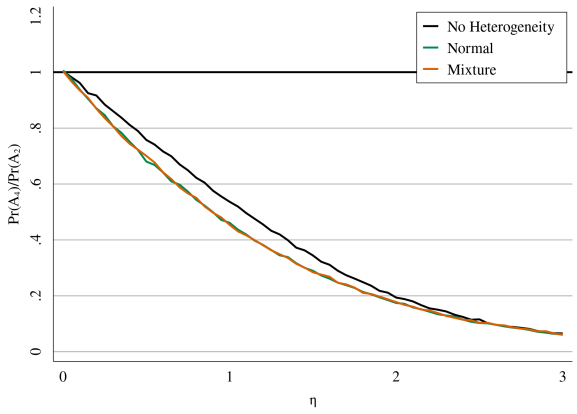
$$\theta^{NORM1} \sim \mathcal{N}\left(\begin{bmatrix} 2 & 2 & 4 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \right)$$

Finally, I allow for a *mixture* of normals, where 50 percent of the sample is drawn from θ^{NORM1}, and 50 percent is drawn from

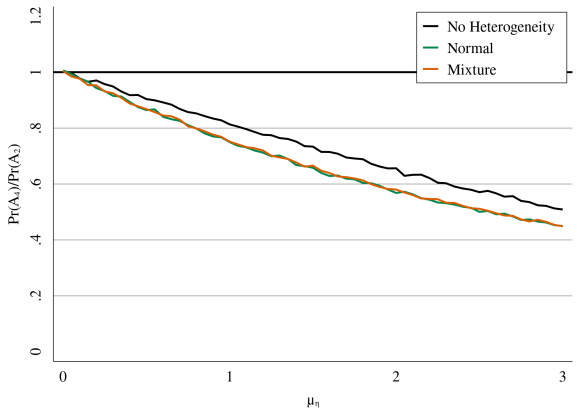
$$\theta^{NORM2} \sim \mathcal{N}\left(\begin{bmatrix} 1 & 4 & 2 \end{bmatrix}, \begin{bmatrix} 3 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix} \right).$$

Figure A.7: Simulated $\rho = \mathbb{P}(A4)/\mathbb{P}(A2)$

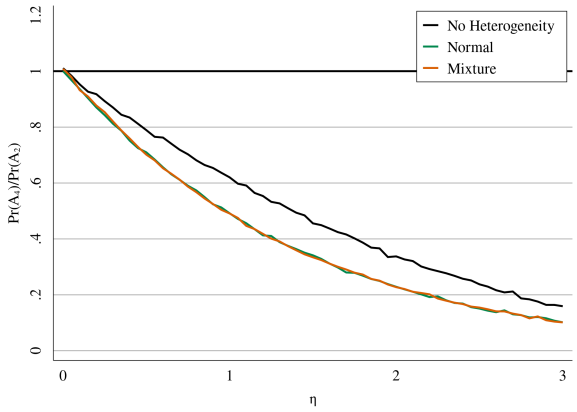
(a) $\overline{\eta}$, Naive Plans



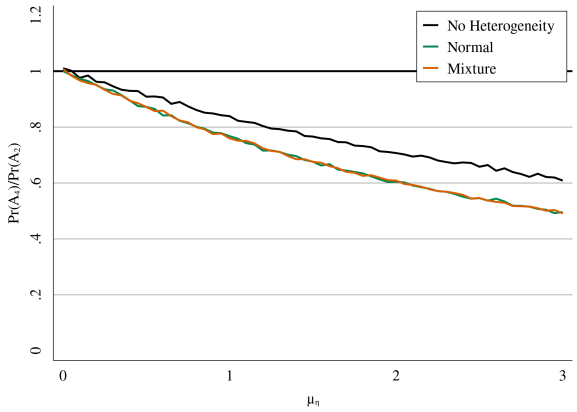
(b) η_i, Naive Plans



(c) $\overline{\eta}$, Mature Plans



(d) η_i, Mature Plans



Notes:

The baseline analysis demonstrates the properties of ρ, across each distribution of consumer preferences, allowing for various values of a fixed switching cost. For each set of simulations, I estimate the expectation of ρ by averaging the choices of $N = 1,000$ randomly drawn consumers over $R = 500$ identically simulated markets.

In a second set of analysis, I allow for heterogeneity in switching costs by drawing consumers from a log-normal distribution with a standard deviation, $\sigma_\eta = 0.5$. I choose a log-normal for this parameter because I am specifically interested in testing the behavior or ρ where switching costs are weakly positive, as assumed in the set-up of the problem. In unreported analysis, I find that the results are qualitatively similar with the unrestricted range of a normal distribution, provided that the standard deviation is not too large. However, for sufficiently small means of switching costs, large negative draws from the left tail can produce some estimates of $\rho > 1$.

Figure A.7 shows the results from these exercises. The left column – panels **(a)** and

(c) – of the table presents results for simulations using homogeneous switching costs. The right column – panels **(b)** and **(d)** – presents results from simulations using the log-normal distribution of switching costs. The top two panels – **(a)** and **(b)** – show these results for "naive plans," where the first choice cannot involve a switching cost by construction. In the bottom two panels, I simulate the results using steady state sequences calculated after a 20-period burn in.

All of these exercises provide intuition that confirms the results above: ρ is close to 1 (up to a small amount of simulation error) when there are no switching costs, and deviates from one as switching costs become larger. Adding heterogeneous brand preferences generates lower ρ for a given value of switching costs. However, the incremental effect from the normally distributed preferences to the mixture model is undetectable in every exercise. Mature plans look qualitatively similar to naive plans. Adding a distribution of switching costs mutes the effect of η on ρ. This is likely driven by the assumed log-normal distribution of η_i, which has a large density below the mean.

E Example "Request for Proposals" (RFP)

This section contains excerpts from an example RFP from a large group employer. Figure A.8 shows the list of plan requirements. Notably, requirement 5 requires the bidding insurer to replicate the benefits of existing plans exactly. This supports the evidence in Section **??** showing that employer switching and premium adjustments are not driven by changes to plan generosity or other benefits.

Figure A.9 shows the level of claims data insurers receive when submitting bids to contract with employers. In this case, the insurer received data for six years of claims history. For the most recent years, they also received a listing of the most expensive conditions and procedures. Finally, Figure A.10 further shows the level of detail that insurers receive on enrollment and characteristics of employees for setting premiums.

Figure A.8: Plan Requirements

1. Proposals are **required** to provide a minimum 12-month rate guarantee, with a contract period of January 1, 2007 through December 31, 2007.

2. Since there are important considerations involved in selecting a carrier in addition to rates, [redacted] will not be required to accept the lowest bid. In addition to gross premium and retention charges, services rendered will also serve as a basis for award of the contract.

3. The Carrier must submit evidence of ability to service the group without undue requirements of [redacted] employees. Each Carrier should list three (3) references that have terminated within the last year as well as three (3) references that are active groups and are approximately our size. (Forms Provided).

4. [redacted] reserves the right to reject any and all proposals and to accept any bid deemed advantageous to [redacted] Any variance from these specifications must be stated in detail with complete reference to the bid specification provision from which the deviation is being made.

5. All proposals must be based on exact duplication of the existing plan benefits unless alternate benefits are requested. Any variance of benefits must be explained in writing and attached to the proposal for consideration. (Plan of current benefits attached.)

6. Proposals must be submitted for coverage on all eligible full-time regular employees and their dependents.

7. Actively at work and dependent confined requirements must be waived.

8. Please complete the appropriate enclosed bid forms that include:
   a. Proposal submission form including a Declaration of Compliance.
   b. Questionnaire
   c. References

   All proposals, including the current carrier, shall complete the proposal forms provided. **All proposal forms submitted must be signed by an authorized official that has the authority to bind the bidder.**

Figure A.9: Claims History

X. MEDICAL & LIFE HISTORY

Medical History and Claims Experience

Carrier History:

1999 -2001	
Carrier	Aetna
Plan	HMO

	January 2001 - December 2004	
Carrier	Pacific Care	
Plan	HMO	
	Premium	Claims
Jan 2002 - July 2002	$ 230,851	131,955
Aug 2002-July 2003	$ 397,802	$364,780
Plan	PPO	
Aug 2003-Dec 2003		
Jan 2004-June 2004	$ 235,320	$ 97,168
Aug 2004 - Dec 2004	Not Available	Not Available

1/1/04 - 12/31/04	Rates	
	E	$292.12
	ES	$702.98
	EC	$538.13
	EF	$867.08

	January 2005 - December 2005			
Carrier	Aetna			
Plan	HMO & PPO			
	Premium	Claims	Subscribers	Members
Jan 2005	$ 41,814	$ 28,907	85	170
Feb 2005	$ 41,272	$ 31,275	87	173
Mar 2005	$ 43,766	$ 22,793	85	171
Apr 2005	$ 40,930	$ 19,548	84	168
May 2005	$ 39,130	$ 19,690	83	167
Jun 2005	$ 39,907	$ 33,032	83	168
Jul 2005	$ 42028	$ 94,884	85	174
Aug 2005	$ 42028	$ 70875	85	174
Sep 2005	$ 42877	$ 42550	84	174
Oct 2005	$ 40608	$ 39406	81	169
Nov 2005	$ 41,559	$ 18,876	83	172
Dec 2005	$ 38,855	$ 31,853	81	167

1/1/05 - 12/31/05	Rates	PPO	HMO
	E	$ 276.86	$ 313.46
	ES	$ 643.09	$ 728.12
	EC	$ 536.34	$ 607.25
	EF	$ 822.38	$ 931.10

	January 2006 - December 2006			
Carrier	Aetna			
Plan	HMO & PPO			
	Premium	Claims	Subscribers	Members
Jan 2006	$ 47732	$ 75,563	86	176
Feb 2006	$ 50,778	$ 22,497	88	182
Mar 2006	$ 47,731	$ 22,961	91	189
April 2006	$ 50,210	$ 53,369	90	185
May 2006	$ 46,126	$ 61289	90	185
Jun 2006	$ 50,193	$ 253,722	94	196
Jul 2006	$			
Aug 2006				
Sep 2006				
Oct 2006				
Nov 2006				
Dec 2006				

1/1/06 - 12/31/06	Rates	PPO	HMO	HSA
	E	$ 314.24	$ 355.78	$ 276.20
	ES	$ 729.91	$ 826.42	$ 641.55
	EC	$ 608.75	$ 689.23	$ 535.06
	EF	$ 933.40	$ 1056.80	$ 820.41

Large Claims			
From	To	Amount	Condition
01/01/2005	07/31/2005	$ 43,115.63	Breast Cancer*
05/01/2005	04/30/2006	$ 115,453	Breast Cancer*
		$ 40,187	Early onset of delivery**
		$ 26,795	Sinoatrial Node Dysfunction***
07/01/2005	06/30/2006	$ 92,656	Breast Cancer*
		$ 42,340	Early Onset of Delivery**
		$ 35,985	Sinoatrial Node Dysfunction***
07/01/2005	06/30/2006	$ 97,956	Vomiting
		$ 42,340	Squamous Blepharitis
		$ 35,555	Intestinal Infection
		$ 52,540	Coronary Atherosclerosis

Figure A.10: Enrollment Census

HMO Census
September 25, 2006

D/O/B	SEX	HMO	STATUS OF COVERAGE	TOTAL
01/07/65	F	+	Employee + Children	1
04/09/62	F	+	Employee + Children	1
09/11/66	F	+	Employee + Children	1
07/20/80	F	+	Employee + Children	1
06/17/45	F	+	Employee + Children	1
07/11/70	F	+	Employee + Children	1
			=	6
04/24/60	M	+	Employee + Spouse	1
11/13/47	M	+	Employee + Spouse	1
			=	2
05/01/56	M	+	Employee + Family	1
07/21/57	M	+	Employee + Family	1
12/27/51	M	+	Employee + Family	1
10/08/69	M	+	Employee + Family	1
			=	4
11/03/67	F	+	Employee + Family	1
09/15/72	F	+	Employee + Family	1
04/29/66	F	+	Employee + Family	1
12/20/70	F	+	Employee + Family	1
01/15/53	F	+	Employee + Family	1
			=	5

PPO Census
September 25, 2006

D/O/B	SEX	PPO	STATUS OF COVERAGE	TOTAL
09/05/49	M	**	Employee + Spouse	1.00
06/10/46	M	+	Employee + Spouse	1.00
08/28/61	M	**	Employee + Spouse	1.00
			=	3.00
05/31/72	F	**	Employee + Spouse	1.00
			=	1.00
03/16/63	M	+	Employee + Family	1.00
06/10/66	M	+	Employee + Family	1.00
05/11/68	M	+	Employee + Family	1.00
04/27/60	M	+	Employee + Family	1.00
07/29/59	M	+	Employee + Family	1.00
02/26/68	M	+	Employee + Family	1.00
03/10/56	M	**	Employee + Family	1.00
			=	7.00
01/18/68	F	+	Employee + Family	1.00

F Estimating Unobserved Premiums

F.1 Estimation

The goal of this estimation exercise is to impute counterfactual premiums for demand and cost estimation. In general, the premiums will reflect the cost at which insurers can provide a given bundle of health care consumption to consumers. Because premiums are set individually, they will reflect the underlying costs associated with insuring each employer, as well as the switching costs and preferences of the employer. Moreover, there may be important match value between an employer's preference for certain providers and the prices a particular insurer has negotiated with those providers.

In order to impute a set of flexibly estimated premiums for unobserved insurers and states, I implement a Lasso regression predicting logged premiums using as much information as possible about each employer: their premiums in observed states, and characteristics of their choice sequences that would imply the strength of their preferences or the magnitude of their switching costs. Maintaining the employer-specific components of costs is important because the dispersion of costs across employers is a key driver of selection in the counterfactuals. Using employer i's observed premiums as a basis for anchoring their premium levels, the Lasso specification is designed to impute premium offers using other employers with similar choice patterns who do purchase from insurer j at some point.

$$log(p_{ijt}) = f(X_{ijh}, j, h, \tilde{\beta}) + \tilde{\varepsilon}_{ijt} \tag{24}$$

where X_{ij} is a vector of measures summarizing observed prices and choices for employer i. $\tilde{\beta}$ is the set of coefficients on the relevant features. Specifically, X_{ij} contains employer-specific measures of

- Indicators for each insurer j and state h
- Average observed premiums across all insurers (j) and states (h)
- Average observed premiums for employer i at each insurer separately
- An indicator variable for whether employer i has an observed premium for insurer j
- Raw choice probabilities for each insurer
- Fraction of years in which an employer switches between insurers
- An indicator for whether or not a choice represents a switch in a given year for that employer

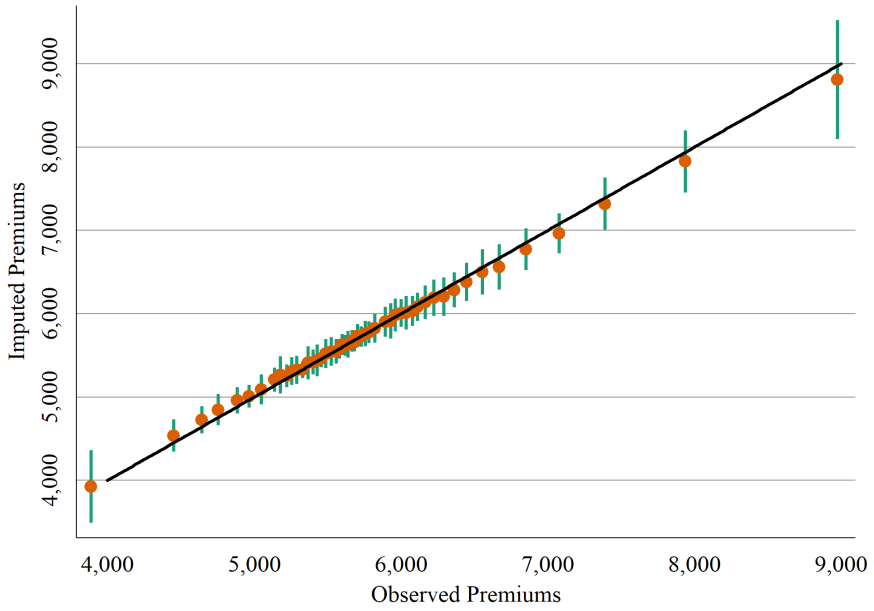
Average observed premiums for employer i for other insurers and states are included to capture the employer-specific components of cost. The indicators for employer j are intended to capture insurer-specific costs. State indicators and the switching rate are included to capture information about each employers' switching costs, and the invest-harvest pricing they face. I also include a third-order polynomial on each pairwise combination of the items listed above. The intuition behind this is to allow for relative shifts in premiums, which

may be informed by aspects of the choice sequence, using observed premiums for employer i as a base. These features may have important, and possibly non-linear interactions. For example, one would expect that a lower switching rate provides some information about the switching cost. This may impact prices, but should likely be interacted with whether or not that insurer represents a switch for the employer in that year.

F.2 Fit

The root mean squared error of the model is \$196.44. To put that number in context, the root mean squared error for a naive model with only a constant term is \$853.60. This results in an R^2 of 0.947 ($= 1 - (196.44^2/853.60^2)$), indicating that the model explains nearly 95 percent of the premium variation in the data. Estimates are slightly less dispersed, which is consistent with premium observations in the tails of the observed premium distribution containing more measurement error.

Figure A.11: Fit of Estimated Premiums



Notes: Figure displays a binscatter showing estimated premiums over the distribution of observed premiums. Data are split into 50 percentile bins of observed premiums. Each point displays the mean predicted value for the bin. Horizontal bars display the standard deviation of predicted values within a bin.

G Additional Detail on Cost Estimation

G.1 Multiplicity

Since I do not know whether the equilibrium of my model is unique, the nested fixed point is not guaranteed to find the equilibrium observed in the data. In order to address this concern, I estimate the cost parameters using the observed premiums in the data, and the corresponding transition probabilities implied by the demand model. Given a guess of costs, I estimate the payoffs for each state and use the transition probabilities to simulate forward to obtain the value function. I then use the insurers' first-order conditions to obtain an optimal premium offer, which I compare to the observed premiums.

Specifically the estimation algorithm proceeds by

1. Guessing a value of c_{ij} for each (i, j, h_{-1})
2. Simulating forward to find the value function given the full set of premium offers and transition policies
3. Solving the first order condition using the value functions in (2) to obtain $\hat{p}_{ij|h_{-1}}$
4. Minimizing the sum of squared residuals between observed and implied premiums:

$$\hat{c}_{ij} = argmin \sum_{ijt|h_{-1}} \left(p^*_{ijt|h_{-1}} - \hat{p}_{ij|h_{-1}}(c_{ij})\right)^2$$

where $p^*_{ijt|h_{-1}}$ is given by the data and $\hat{p}_{ij|h_{-1}}(c_{ij})$ is the implied optimal premium calculated in step (3)

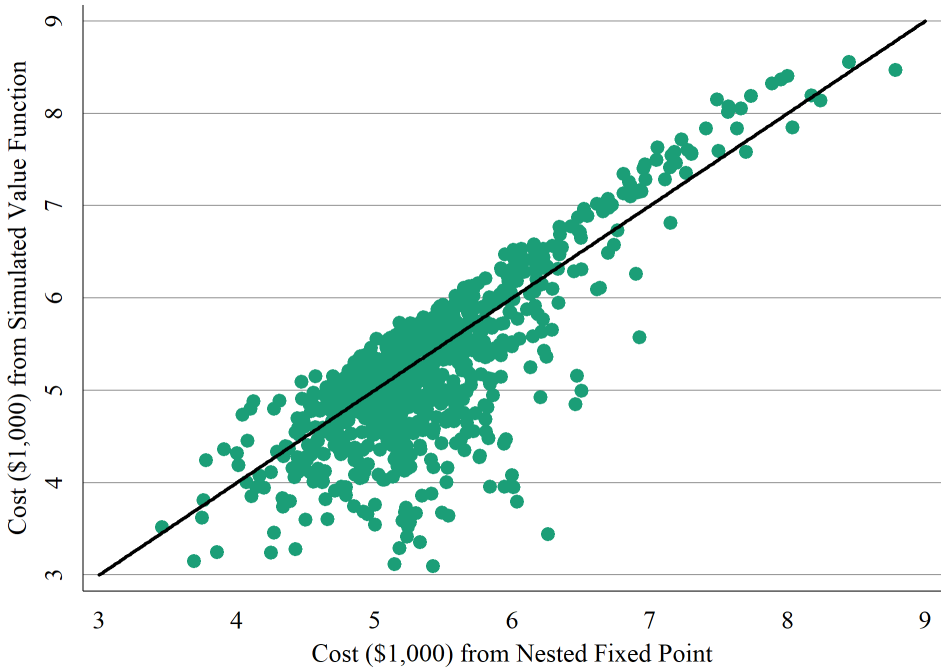
This method of estimating the costs is computationally fast and does not require that I solve for the equilibrium. However, one limitation of this method is that there is no guarantee that the costs I obtain will reproduce prices close to the observed policy function when re-calculating the equilibrium. In this method, the policy function is taken as given when simulating the value functions. Measurement error in observed premiums may propagate the value functions, moving the estimated costs away from the true equilibrium – or indeed any equilibrium that could be implied by the model.[14] Therefore, I interpret these estimates with caution, but view them as a complement to the approach described above. Estimates from this approach that are qualitatively similar to the those obtained using the nested fixed point providing supportive evidence that those parameters were not obtained from the wrong equilibrium.

In Figure A.12, I show the final results of the cost estimation for both methods of estimating costs. Each point shows a unique c_{ij} estimate under each approach. The accompanying table summarizes the estimates. Column (1) shows the estimates based on the nested fixed point approach. Column (2) shows the cost estimates based on the method that uses the equilibrium policy functions to simulate the value functions before solving the first-order condition. Both sets of estimates are similar, revealing that HPHC and Tufts have sizable cost advantages relative to BCBS. The standard deviations reveal that both sets of

[14] Additional detail on sources and magnitudes of measurement error is provided in Appendix H.

estimates also produce a sizable amount of variation across employers. The simulated value functions produce estimates that are on average $300-$500 lower than the nested fixed point and more dispersed across employers, which I attribute to the role of measurement error in premiums.

Figure A.12: Cost Estimates



		(1) Nested Fixed Point	(2) Simulated Values
BCBS	Mean	5,579	5,057
	SD	(636)	(943)
HPHC	Mean	5.093	4,757
	SD	(473)	(618)
Tufts	Mean	5,143	4,817
	SD	(559)	(704)

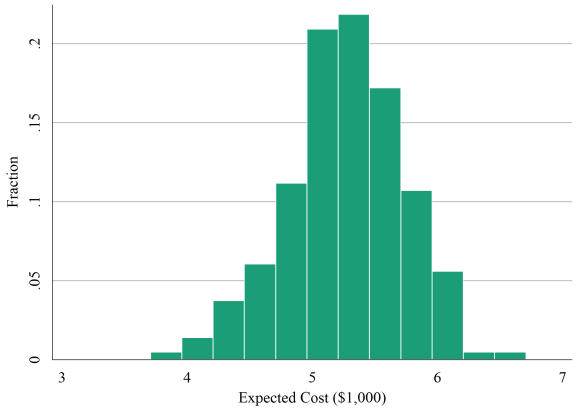
For the counterfactuals in Section 7, I privilege the estimates obtained using the nested fixed point, but interpret this as suggestive evidence that the nested fixed point likely does not suffer from a multiplicity problem. These estimates are also closer to aggregate estimates of average claims per covered life published by CMS of $5,400. Additionally, more dispersion in costs across employers will tend to exacerbate adverse selection in the counterfactuals. I therefore view the results of the nested fixed point as a more conservative estimate for the purpose of this paper.

G.2 Importance of Dynamic Premium Setting

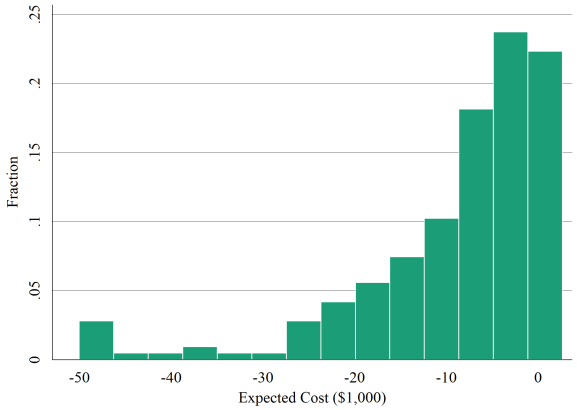
In Table A.13, I recompute cost estimates implied by a static Nash-Bertrand model. Panel (a) shows the main estimates recovered using the dynamic model for comparison. Panel (b) shows the estimates recovered using a myopic first order condition, which results in average expected costs of -$8,813. 89 percent of employers have negative expected costs using this method. This comparison demonstrates the value of accounting for dynamic incentives, which restrain insurer premiums even when the employer is locked in.

Figure A.13: Estimated Cost Parameters

(a) Baseline Cost Estimates



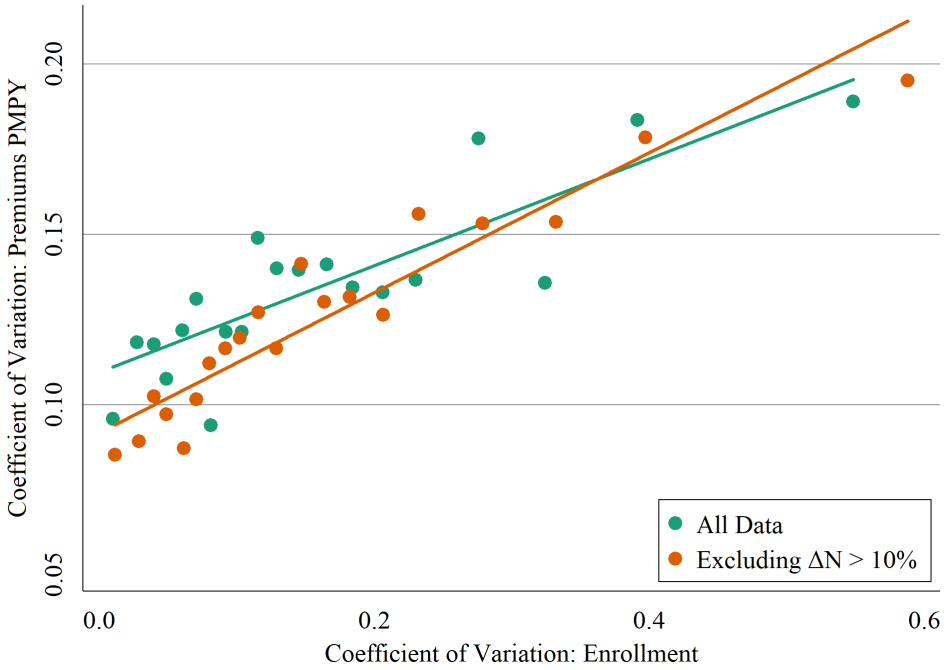
(b) Cost Estimates Using Static Supply Model



H Testing for Measurement Error in Premiums

This Appendix provides additional detail regarding the scale and source of measurement error in premiums. Because premiums are standardized as a per-person-per-year measure, fluctuations in firm size can influence the measure of premiums. Premiums are typically paid on a per-member-per-month basis, whereas I only observe the number of covered lives on a yearly basis. To illustrate the relationship between measured enrollment and premiums, I calculate the coefficient of variation of each variable within employer-insurer pairs. Figure A.14 presents a binned scatterplot showing the strong, positive relationship between the two. The orange series recalculates the same measures, conditioning on years in which the number of covered lives changed by less than 10 percent. This restriction results in lower measured variation in premiums. The average coefficient of variation in premiums is 0.12, indicating that the within ij standard deviation of premiums is just over $700 on average.

Figure A.14: Within-Employer Variation in Premiums and Enrollment



Although employer size is not an important determinant of premium setting in my model, Figure A.14 could also be consistent with a model of employer-insurer bargaining in which larger employers receive lower premiums. Under such a model, fluctuations in firm size would result in changes to premiums due to real economic forces rather than measurement error. To directly test this, I estimate two regressions that capture the relationship between firm size and premiums within an employer-insurer dyad.

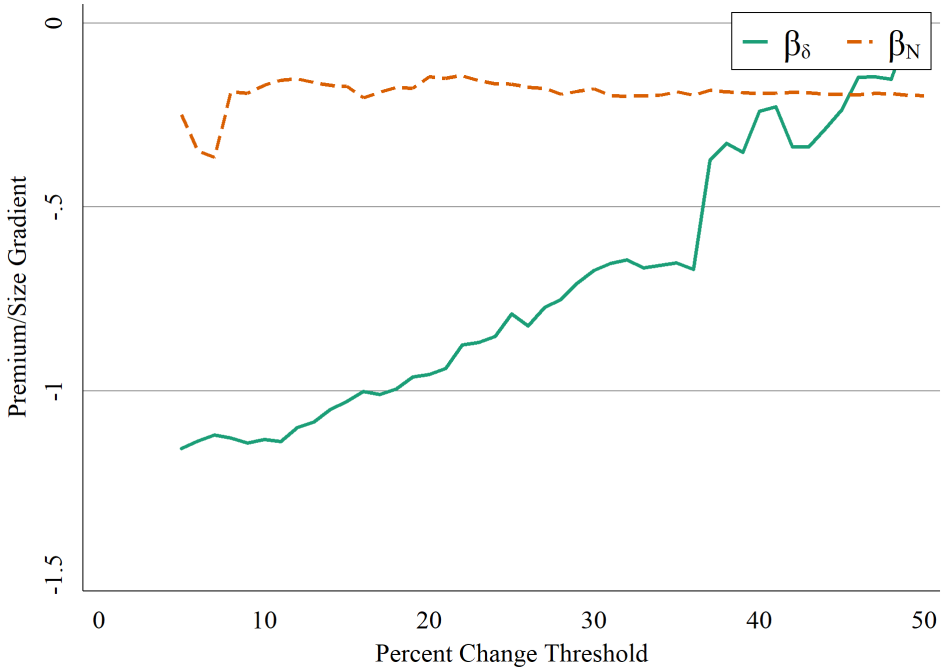
$$p_{ijt} = \beta_N N_{it} + \gamma_{ij} + \epsilon^1_{it} \tag{25}$$
$$p_{ijt} - p_{ijt-1} = \beta_\delta (N_{it} - N_{it-1}) + \epsilon^2_{it} \tag{26}$$
$$\tag{27}$$

I use the first equation to estimate β_N, restricting to years in which the year-to-year change in premiums is small. This specification is intended to capture any "true" price effects, while omitting years in which large transitions take place. Intuitively, β_N captures the premium difference for the same employer-insurer pair between years in which the employer is large or small, but not transitioning in size. By contrast β_δ in the second equation is designed to capture the fluctuations in premiums that occur purely due to the year-to-year change in observed firm size.

I compare estimates of β_N when the percent change in N is small, to estimates of β_δ in years where the percent change in N is large. The difference between these two estimates estimates is informative about the degree to which the correlation between premiums and covered lives is driven by measurement error rather than an employer size discount. If $|\beta_\delta - \beta_N|$ is small then premium fluctuations are likely driven entirely by economic forces, whereas a large difference implies observed "size" discounts that are driven by measurement error in enrollment.

Figure A.15 presents the estimates of this exercise, where the x-axis varies the threshold by which I consider a percent change in N to be large or small. Focusing on the non-differences fixed-effect specification, β_N, the gradient between firm size and premiums is small, and consistent across thresholds. On the other hand, estimates of β_δ, focusing on large year-to-year changes in premiums, show a much stronger relationship between the two measures. As the treshold approaches 50 percent, the estimating samples become more similar and the estimates for both coefficients converge. I view this as strong evidence that a large amount of within employer-insurer premium variation is driven by measurement error in firm size fluctuations.

Figure A.15: Testing for Impact of Fluctuations in Coverage Count



I Additional Counterfactual Results

Table A.5: Equilibrium Under a Public Option with Experience Rating

Cost Advantage	Quality Shortfall	(1) PO Enrollment	(2) PO Premium	(3) Market Premium	(4) ΔCS (with PO)	(5) ΔCS (prices only)
Baseline				6,058		
0%	0	0.364	6,202	5,659	592	46
	–500	0.211	6,202	5,843	281	28
	–1,000	0.115	6,202	5,957	122	14
	–2,000	0.032	6,202	6,041	19	2
25%	0	0.882	4,681	5,259	2,281	65
	–500	0.739	4,681	5,241	1,615	69
	–1,000	0.530	4,681	5,417	973	59
	–2,000	0.185	4,681	5,854	229	24
50%	0	0.984	3,247	5,218	3,720	34
	–500	0.966	3,247	5,120	3,194	43
	–1,000	0.923	3,247	5,037	2,634	55
	–2,000	0.666	3,247	5,216	1,378	62

Table A.6: Equilibrium Under a Public Option with 10% Introductory Discount

Cost Advantage	Quality Shortfall	(1) PO Enrollment	(2) PO Premium	(3) Market Premium	(4) ΔCS (with PO)	(5) ΔCS (prices only)
Baseline				6,058		
0%	0	0.264	6,600	5,720	402	31
	-500	0.135	6,700	5,892	162	16
	-1,000	0.072	6,700	5,976	68	8
	-2,000	0.020	6,700	6,041	11	1
25%	0	0.819	4,800	5,004	2,090	64
	-500	0.667	4,800	5,203	1,465	62
	-1,000	0.449	4,900	5,481	820	48
	-2,000	0.145	5,000	5,875	176	17
50%	0	0.989	3,100	5,282	3,844	35
	-500	0.966	3,200	5,000	3,217	45
	-1,000	0.916	3,200	4,934	2,648	56
	-2,000	0.661	3,200	5,215	1,420	60

Table A.7: Equilibrium Under a Public Option with 2% Administrative Load

Cost Advantage	Quality Shortfall	(1) PO Enrollment	(2) PO Premium	(3) Market Premium	(4) ΔCS (with PO)	(5) ΔCS (prices only)
Baseline				6,058		
0%	0	0.181	6,400	5,812	229	22
	–500	0.086	6,500	5,946	84	10
	–1,000	0.045	6,500	6,004	34	5
	–2,000	0.014	6,400	6,044	6	1
25%	0	0.778	4,600	5,082	1,804	61
	–500	0.570	4,700	5,324	1,097	55
	–1,000	0.348	4,800	5,593	550	39
	–2,000	0.098	4,900	5,930	100	12
50%	0	0.982	3,100	5,215	3,526	38
	–500	0.956	3,100	5,019	2,987	47
	–1,000	0.894	3,100	4,978	2,413	56
	–2,000	0.570	3,200	5,324	1,097	55

Table A.8: Equilibrium Under a Public Option with 8% Administrative Load

Cost Advantage	Quality Shortfall	(1) PO Enrollment	(2) PO Premium	(3) Market Premium	(4) ΔCS (with PO)	(5) ΔCS (prices only)
Baseline				6,058		
0%	0	0.086	7,000	5,946	84	10
	–500	0.045	7,000	6,004	34	5
	–1,000	0.024	7,000	6,033	13	2
	–2,000	0.008	6,900	6,051	2	0
25%	0	0.608	5,100	5,279	1,208	57
	–500	0.418	5,100	5,506	707	45
	–1,000	0.228	5,200	5,748	312	27
	–2,000	0.067	5,200	5,974	58	8
50%	0	0.974	3,300	5,123	3,314	41
	–500	0.936	3,300	4,983	2,763	51
	–1,000	0.855	3,300	5,005	2,172	59
	–2,000	0.494	3,400	5,415	891	50

Table A.9: Equilibrium Under a Public Option with 10% Administrative Load

Cost Advantage	Quality Shortfall	(1) PO Enrollment	(2) PO Premium	(3) Market Premium	(4) ΔCS (with PO)	(5) ΔCS (prices only)
Baseline				6,058		
0%	0	0.143	6,600	5,866	166	17
	–500	0.076	6,600	5,961	70	9
	–1,000	0.035	6,700	6,018	23	3
	–2,000	0.011	6,600	6,048	4	1
25%	0	0.715	4,800	5,153	1,559	61
	–500	0.494	4,900	5,415	891	50
	–1,000	0.315	4,900	5,635	481	36
	–2,000	0.086	5,000	5,946	84	10
50%	0	0.982	3,100	5,215	3,526	38
	–500	0.947	3,200	4,998	2,876	49
	–1,000	0.876	3,200	4,989	2,293	58
	–2,000	0.532	3,300	5,369	991	53