

How does Insurance Competition Affect Medical Consumption?

Conor Ryan*

University of Minnesota - Economics Department

August 26, 2021

[*For the latest version, click here.*]

Abstract

Competition in insurance markets affects not only the monthly premium but also the cost-sharing terms—e.g. copays and coinsurance rates— of the offered products. These terms determine the out-of-pocket price of medical care, which affect a patient’s medical decisions and thus the patient’s health outcomes. In this paper, I estimate a model of imperfect competition in which firms set both the premium and the cost-sharing terms of their products. Consumers select an insurance plan and make medical consumption decisions given the cost-sharing terms of their insurance. Using medical claims data linked to insurance products, the model incorporates adverse selection in insurance demand through the observed relationship between medical diagnoses, realized cost, and insurance choices. I identify the effect of cost-sharing terms on medical consumption and health using within product variation in cost-sharing terms and consumer inertia in insurance plan choice. First, I show that, on average, less competition leads to higher levels of cost-sharing but multi-product firms respond by increasing the cost-sharing levels of some products and decreasing others. Second, I find that medical consumption and health respond to cost-sharing terms. A \$10 increase in the primary care copay leads to a 5.4% decrease in medical consumption and a 0.1 percentage point increase in inpatient mortality. Putting these results together, I find that a reduction in competition via a merger leads to up to a 4% increase in the primary care copay, an average reduction in medical spending of \$17 per person, and an additional six inpatient deaths per year. At estimates of the statistical value of a life, the reduction in spending is more than outweighed by the cost of additional deaths.

*I would like to express appreciation for the guidance of my committee, Tom Holmes, Amil Petrin, Naoki Aizawa, and Stephen Parente, and for helpful comments from Joel Waldfogel, Roger Feldman, Kevin Williams, Sergio Salgado, Kurt See, and Keaton Miller, as well as all the participants of the Minnesota Applied Microeconomics Workshop.

1 Introduction

Competition in health insurance markets affects not only the monthly premium but also the cost-sharing terms—e.g. copays and coinsurance rates— of the offered products. Firms can raise revenue (or reduce cost) through these different product attributes which affect consumer behavior differently. Cost-sharing terms are particularly important because they determine the out-of-pocket price of medical care and therefore, unlike the monthly premium, may affect a patient’s medical decisions. These decisions may then affect the health outcomes of the patient. While there has been substantial research on the effect of competition on the monthly premiums for insurance, there is relatively little research on how competition affects cost-sharing terms and the subsequent effects on medical consumption and health.

Medicare Advantage (MA)—a subsidized, private insurance alternative to the government-run traditional Medicare program—is an important setting to study the effect of competition on medical consumption via cost-sharing terms. Insurance firms compete to attract beneficiaries by choosing the monthly premium, cost-sharing terms, and other characteristics of their products, which must meet standards required by Medicare. Firms are provided with large, risk-adjusted subsidies allowing many products to be offered at low or \$0 monthly premiums. As a result, other dimensions of the product such as the cost-sharing terms are especially important to understand consumer demand, competition among firms, as well as the effect of a myriad of policies and regulations that govern the market.

In this paper, I provide a framework to evaluate the effect of changes in the market structure or market design of the insurance market on the health and health care use of the beneficiaries through the cost-sharing terms of insurance. I estimate a model of insurance competition, medical consumption, and health in Medicare Advantage. There are three main findings. First, I show that, on average, less competition leads to higher levels of cost-sharing, but a merger between multi-product firms may lead the cost-sharing levels of some products to increase and others to decrease. Second, I find that consumers respond to cost-sharing terms in their demand for medical care. A \$10 increase in the primary care copay leads to a 5.4% decrease in medical consumption. Finally, I find that the cost-sharing terms of insurance have an effect on the health outcomes of patients. In particular, a \$10

increase in the primary care copay leads is associated with a 0.1 percentage point increase in inpatient mortality. Using the fully estimated model, I find that a reduction in competition via a merger leads to up to a 4% increase in the primary care copay, an average reduction in medical spending of \$17 per person, and an additional six inpatient deaths per year.

The model of consumer demand consists of two stages. In the first stage, consumers make a discrete choice over the available health insurance plans. In the second stage, consumers make a sequence of monthly medical consumption decisions given their choice of insurance plan. The model incorporates whether consumers respond to higher cost sharing by decreasing medical consumption (moral hazard), whether insurance preferences are correlated with expected cost (adverse selection), and whether insurance preferences are correlated with medical consumption elasticities (selection on moral hazard).

The model of supply consists of strategic, multi-product firms that simultaneously select both the monthly premium and the cost-sharing terms of differentiated insurance products.¹ I show that the effect of competition on cost-sharing terms is ambiguous for two reasons. The first reason follows from standard incentives facing a firm that competes in both price and a non-price quality that consumers value. The level of cost-sharing (or more generally, any product quality) that firms will provide depends on consumer willingness to pay to reduce cost-sharing and the marginal cost to the firm of doing so (Spence 1975, Schmalensee 1979). This substitution between price and quality is further complicated by the second reason: not all consumers generate expected profit. The standard intuition of competition assumes that when a product lowers its price (or cost-sharing level) it attracts more profitable sales, but this may not be the case in markets with adverse selection (Mahoney and Weyl 2017, Veiga and Weyl 2016, Lester et al. 2019).

In order to quantitatively evaluate these mechanisms, it is crucial to characterize consumer preferences for insurance, elasticities of medical consumption, expected cost, and the relationship between each of these features. I accomplish this by using data on insurance plan choices linked to insurance claims data in the Medicare Advantage market in Massachusetts. Using the medical claims, I can construct detailed information on health status—e.g. specific

¹This model is similar to one in which firms set a price and an aspect of product quality.

diagnoses—and link these characteristics to an individual’s choice of an insurance plan with particular cost-sharing terms. Importantly, I can directly relate this data on choices to the expected cost of insuring this group of consumers.

In the first stage, I estimate discrete choice demand for insurance, extending the methodology of Miller et al. 2019 to incorporate information on medical diagnoses and their interaction with the cost-sharing terms of insurance. I find that consumers are elastic to the monthly premium and the primary care copay of insurance, but elasticities with respect to other cost-sharing terms are relatively small. The median willingness to pay to reduce the primary care copay by \$10 is relatively large at \$43 per month, suggesting that consumers may be willing to pay for lower cost-sharing terms in anticipation of its effect on their medical consumption. Consumers in the 95th percentile of expected cost are willing to pay \$63 per month, about 50% greater than the median consumer. Due to high variance in the costs of the high-risk population, willingness to pay is U-shaped in net-cost when risk-adjusted subsidies are taken into account, which is consistent with the findings of Brown et al. 2014.

In the second stage, I estimate consumers’ elasticity of medical consumption with respect to cost-sharing terms using within product variation in the cost-sharing terms of insurance. This source of variation is common in settings where the sample is enrolled in a single product (Brot-Goldberg et al. 2017), and I extend this intuition to a multi-product setting where consumers face considerable inertia in plan choices (Ho, Hogan, and Scott Morton 2017, Miller et al. 2019, Drake, Ryan, and Dowd 2020). While a price elasticity of medical consumption could be identified using the non-linear features of the insurance—as in Aron-Dine et al. 2015 and Ellis, Martins, and Zhu 2017—this strategy directly estimates the elasticity of interest: the average change in medical consumption that will result for a change in the cost-sharing terms of insurance. I focus this identification on the primary care copay, which has substantial year-to-year variation within products. Roughly 65% of the sample experiences a change in the primary care copay of the product in which they are enrolled at some point during the sample period.

I find that a \$10 increase in the primary care copay leads to a 5.4% decline in medical consumption, as measured by total medical spending. Additionally, I find that the healthiest

and least healthy individuals (measured by a medical risk score) are the *least* elastic with respect to primary care copays, while consumers in the middle of the medical risk distribution have the greatest elasticity. These estimates imply coinsurance elasticities between -0.09 and -0.25, which are consistent with other estimates from the literature (Manning et al. 1987, Ellis, Martins, and Zhu 2017).

Next, I investigate the relationship between primary care copays and inpatient mortality. Using rich controls for plan-level mortality and consumer health status, I estimate the direct effect of the primary care copay on inpatient mortality (patient deaths in hospitals or hospice care facilities). I find that a \$10 increase in the primary care copay is associated with a 0.1 percentage point increase in inpatient mortality. The magnitude of the effect is in line with other estimates on the causal differences in mortality among insurance plans in Medicare Advantage (Abaluck, Hull, and Starc 2020).

With data on the cost of insurance linked to estimates of the demand for insurance and medical consumption, I study the effect of competition on medical consumption by reducing the number of firms in the market through potential mergers. I study three potential bilateral mergers among the three largest firms in the Medicare Advantage market in Massachusetts and focus on monthly premium and the primary care copay as the key endogenous features. Each merger leads to increases in both the average premium and the average primary care copay of the products offered by the merging firms. I also show that these findings are consistent with reduced form estimates from a national, market-level panel of competition and cost-sharing terms (including the primary care copay), following the methodology of Bresnahan and Reiss 1991.

The largest premium increases occur alongside the smallest copay increases. In a merger between the largest and second largest firms, the average premium increases by \$12.3 (10.2%) and the average primary care copay increases by only \$0.1 (0.07%). This highlights that the premium and primary care copay are substitutes and shows the importance of evaluating the effect of a merger on cost-sharing terms. Among the mergers studied, I find that those with the smallest premium effect have the largest effect on the primary care copay and medical consumption.

Changes in the primary care copay affect the total medical spending in the market. The average effect on medical spending in each merger ranges from an increase in \$7.3 per person per year to a decrease of \$17.0 per person per year. These effects average a substantial degree of heterogeneity. In the merger between the largest and second largest firm, 35% of consumers experience decreases in the primary care copays of the products in which they were enrolled pre-merger. Spending by these consumers increases by \$77.4 per person per year. Because of the distribution of cost across consumers, the mean effect of this merger is a spending increase, despite a small average increase in the copay.

This leads to an important tradeoff. An increase in cost-sharing terms as a result of a merger will decrease the total spending on medical consumption but also increase expected mortality. By combining the estimates of cost from the merger analysis with the effect on inpatient mortality, I find that the reduction in spending per expected additional death is \$320 and \$388 thousand dollars. This is well below estimates of the statistical value of a life, which range between \$4 and \$10 million for the general population and exceed \$1 million per life even for individuals over the age of 80 (Aldy and Viscusi 2007). This implies that the welfare benefit of a decrease total spending that results from a merger is outweighed by the welfare cost of an increase in mortality.

Relation to the Literature

There is a substantial body of literature that studies premium competition among insurance firms. However, there is comparatively little research on the effect of competition on the cost-sharing terms of insurance, and still less known about how this dimension of competition affects the medical consumption and health care outcomes of insurance beneficiaries.

This paper makes two main contributions. First, I estimate a model of imperfect competition between insurance firms that incorporates both adverse selection and moral hazard in consumer behavior. This builds on a literature that estimates models of differentiated products to study the effects of adverse selection and market concentration in health insurance markets (Ryan 2020, Miller et al. 2019, Jaffe and Shepard 2017, Shepard 2016, Tebaldi 2020, Ericson and Starc 2015, Starc 2014, Saltzman 2021). More specifically, there is a growing

literature that explores the mechanisms through which firms seek a more favorable risk pool (Cao and McGuire 2003, Brown et al. 2014, Newhouse et al. 2015, Newhouse et al. 2015, Aizawa and Kim 2018, Decarolis and Guglielmo 2017, Geruso, Layton, and Prinz 2019). And in particular, there is a literature on managed care in Medicare Advantage that documents mechanisms and incentives to screen for profitable consumers through the generous (or sparing) provision of certain types of service (Glazer and McGuire 2000, Frank, Glazer, and McGuire 2000, Ellis and McGuire 2007). I build on this work by estimating a model in which firms can choose the cost-sharing terms of insurance, in addition to the premium, in an environment with both adverse selection and moral hazard.

I am building on a literature that estimates the two-stages of consumer decision making in health insurance markets: the purchase of insurance and the consumption of medical care (Marone and Sabety 2020, Einav et al. 2013, Cardon and Hendel 2001). Marone and Sabety 2020 estimates a model in which consumers have beliefs over their out-of-pocket expenditures and preferences over insurance plans that depend directly on their distributions of out-of-pocket spending. Building on insights that consumers make mistakes when selecting health insurance (Handel and Kolstad 2015, Handel, Kolstad, and Spinnewijn 2019, Afendulis, Sinaiko, and Frank 2015, Dalton, Gowrisankaran, and Town 2020, Bhargava, Loewenstein, and Sydnor 2017), I estimate a model where consumer insurance demand can depend flexibly on consumer medical conditions and the cost-sharing characteristics of the insurance plan but does not necessarily assume any un-biased projection of health expenditure by the consumer. This first stage of the estimations adds data on medical diagnoses and service-specific cost-sharing terms to standard discrete choice insurance demand estimation methods (Town and Liu 2003, DeLeire et al. 2017, Miller et al. 2019, Tebaldi 2020, Drake 2019, Geruso 2017).

In the second stage, I extend work on estimating the reduced form price-elasticity of medical care to multi-product, non-group insurance markets. Beginning with the RAND Health insurance experiment, a randomized experiment on insurance benefits (Manning et al. 1987), the literature has studied medical consumption in larger contexts using natural experiments (Duarte 2012, Brot-Goldberg et al. 2017), contract non-linearity (Aron-Dine

et al. 2015, Ellis, Martins, and Zhu 2017), variation in the choice set (Lavetti, DeLeire, and Ziebarth 2019, Marone and Sabety 2020), and instrumental variables (Kowalski 2016). In this paper, I exploit inertia in consumer insurance choices and year-to-year changes in the copays applicable to each type of service in order to estimate the elasticity of consumer spending to key variables set by the insurance firm.

The second contribution is estimating the effect of a change in competition in Medicare Advantage on medical consumption and inpatient mortality through cost-sharing terms. This contributes to a literature that studies market structure in health insurance (Cutler and Reber 1998, Town 2001, Dafny, Duggan, and Ramanarayanan 2012). Previous research moves beyond the focus on the insurance premium to study competition in the context of contracting with provider networks (Shepard 2016, Ho and Lee 2017, Dafny, Ho, and Lee 2018), the Medicare Advantage bidding rules (Cabral, Geruso, and Mahoney 2018, Curto et al. 2021), and the ways that insurance product design feeds back into the market structure of the provider industry (Capps, Dranove, and Satterthwaite 2003, Gowrisankaran, Nevo, and Town 2015). This paper builds on this work to study the effect market structure on medical consumption and patient health through the cost-sharing terms of insurance. My findings also contribute to a more broad literature of how competition and mergers affect product quality (Bloom et al. 2015, Fan 2013).

2 Setting: Medicare Advantage

The Medicare Advantage market is a regulated market in which private insurance firms offer subsidized insurance plans to individuals eligible for the Medicare program. This market is an important setting to study the importance of competition and cost-sharing terms for three reasons: i) the program design is based on the notion that encouraging competition will benefit consumers and save money for the government, ii) the degree of competition varies substantially across local markets and merger activity is common, and iii) equilibrium premiums are low and occasionally zero, which encourages competition on cost-sharing parameters.

The traditional Medicare program (TM) is a government-sponsored, fee-for-service health insurance plan available to U.S. citizens or permanent residents who are over the age of 65 or disabled. In 2003, the government formalized Medicare Advantage (MA): a companion program available to the same individuals in which insurance firms receive substantial government subsidies and compete to offer insurance plans which cover at least the same services as TM. An important motivation for establishing MA was the notion that, by allowing firms to compete and offering consumers more choices of insurance, the government could provide greater benefits to consumers at a lower cost than could be achieved through TM (Bush 2002).²

While MA was not the first program to make private insurance plans available to Medicare beneficiaries, it prioritized making the market attractive for insurance firms in order to generate competition. The MA program provided larger subsidies that were more accurately adjusted for risk, and the result was both a substantial increase in the number of participating insurance firms and a steady increase in the number of Medicare beneficiaries that opt into MA (McGuire, Newhouse, and Sinaiko 2011). Despite these successes, the degree of competition still varies substantially across the nation. Only a single insurance firm offered insurance through MA in roughly one out of seven counties between 2011 and 2019, while many of the largest counties had more than 10 competing insurance firms.

The MA market is also a frequent stage for merger activity. Since 2003, the Antitrust Division of the Department of Justice has sued to prevent or require divestitures in three health insurance mergers because of potential anti-competitive effects in MA.³ Still more mergers have been consummated that have not risen to such high levels of antitrust concern.⁴

Due to the large subsidies and associated rules, competition between firms is often concentrated on the cost-sharing parameters rather than the monthly premium. Insurance

²This reasoning is not unique to Medicare Advantage, and played a role in the creation of the Medicare prescription drug program, Part D, and the health insurance marketplaces created by the Affordable Care Act.

³These mergers include Aetna-Humana, blocked in 2018; Humana-Arcadian Management Services, consummated with divestiture in 2012; and United-Sierra Health, consummated with divestiture in 2008. MA was not necessarily the only antitrust concern in each case.

⁴For instance, Aetna-Coventry in 2013 and United-PacifiCare in 2005.

plans submit a “bid” to the government, which determines the base level of subsidy that the plan will receive for each enrollee. Most plans submit a bid that is below a set benchmark level and are rewarded with an additional subsidy, termed a “rebate”, which the plan is required to return to consumers via either a lower premium or more generous cost-sharing terms. However, if the insurance plan wishes to reduce its premium below zero, it must do so through a rebate to consumers in their social security checks. While many insurance plans have premiums that are low or exactly equal to zero, these rebates are rare—perhaps due to low salience among the consumers. As a result, MA plans compete on the broad set of cost-sharing terms and tend to provide substantially more generous cost-sharing terms than TM.⁵

3 Model

This section presents a model with three components: i) a model of consumer demand for insurance that incorporates adverse selection and moral hazard, ii) a model of medical consumption given the cost-sharing terms of the chosen insurance plan, and iii) a model of competition between insurance firms that set both a monthly premium and cost-sharing terms. The model can then be used to characterize the effect of a change in regulation or market structure on the cost-sharing terms of insurance and the medical consumption of its beneficiaries.

3.1 The Environment

Consumers

Consumers, indexed by i , face a two stage decision following Cardon and Hendel 2001 and Dubin and McFadden 1984. In the first stage, consumers select an insurance plan, j , during an annual period for open enrollment. In the second stage, consumers face a realization of medical needs and consume an amount of medical care each month, m , at the out-of-pocket prices set by the insurance plan in which they are enrolled.

⁵While MA cost-sharing terms are typically more generous than TM, a substantial fraction of TM enrollees purchase additional secondary insurance coverage, or “Medigap” insurance.

For exposition, consider a single, annual medical consumption decision in the second stage.

$$U_{ij}^*(\omega) = U^*(\omega; p_j, X_j, W_j, \eta_i) = \max_m U(m, \omega; p_j, X_j, W_j, \eta_i) \quad (1)$$

where ω is a preference shock for medical demand, p_j is the monthly premium of the insurance plan, X_j is a vector of cost-sharing parameters that govern the out-of-pocket price of medical consumption, W_j is a vector of non-financial insurance plan characteristics, and η_i represents the characteristics of the consumer which may include a signal about ω . The function U represents the indirect utility of an amount of medical consumption, m , given the characteristics of the insurance plan, and the function U^* incorporates the optimal level medical consumption, $m^*(\omega, X_j, \eta_i)$, which I assume does not depend on the premium or non-financial plan characteristics.⁶

In the first stage, a consumer who purchases insurance plan j for the plan year t receives an indirect expected utility given by

$$v_{ijt} = V(\epsilon_{ij}, \mathcal{E}[U_{ij}^*(\omega)|\eta_i]) \quad (2)$$

where ϵ_{ij} is an idiosyncratic preference of consumer i for product j and \mathcal{E} is the consumers subjective expectation of their second stage utility given their characteristics, η .

Consumers select the insurance plan that maximizes the total indirect utility of the insurance plan choice. The probability that a consumer, i , selects an insurance plan j is

$$s_{ijt} = \Pr\{v_{ijt} \geq \max_k v_{ikt}\} \quad (3)$$

The expectations in \mathcal{E} are complex. Because the cost-sharing terms in X_j represent

⁶This requires that the income effect of the premium is small, which is reasonable given the low level of premiums in Medicare Advantage.

detailed categories of service (e.g. primary care visits, diagnostic imaging, urgent care visits, etc.), consumers must have beliefs regarding each specific service.⁷ Not only are these expectations difficult to compute and parameterize, there is evidence that consumers themselves may not be very good at making consistent predictions of their medical use at the time of purchasing insurance (Kling et al. 2012, Handel and Kolstad 2015, Handel, Kolstad, and Spinnewijn 2019).

The empirical estimation focuses on separately estimating s_{ijt} and m^* and avoids explicitly specifying the distribution of consumer beliefs. More information on the specification, identification, and estimation of the two stages of the consumer problem are presented in sections 5 and 6.

Firms

Insurance firms choose monthly premiums, p , and a vector of cost-sharing parameters, X , each year to maximize the static, one-year profit of the firm.⁸ The profit of a single product, j , depends on the probability that each individual will enroll, s_{ijt} , the monthly premium, p_{jt} , an individual-specific subsidy, b_{ijt} , and the expected individual-specific marginal cost, mc_{ijt} .

$$\Pi_{jt} = \int_i s_{ijt}(p_{jt}, X_{jt}, \mathbf{p}_{-jt}, \mathbf{X}_{-jt})(p_{jt} + b_{ijt}(p_{jt}, X_{jt}) - mc_{ijt}(X_{jt}))dF(i) \quad (4)$$

$$p_{jt} \geq 0 ; x_{jtk} \in X_{jt} \geq 0$$

where \mathbf{p}_{-jt} and \mathbf{X}_{-jt} represent the premium and cost-sharing terms for all other products in the market.

Firms cannot set the cost-sharing parameters or the monthly premium to be below zero.

⁷The literature that assumes that consumers evaluate U directly typically has cost-sharing terms governed by only a few general parameters, such as a deductible, coinsurance rate, and out-of-pocket limit (Cardon and Hendel 2001, Marone and Sabety 2020).

⁸The firms do internalize the same switching costs as the consumers. In this way, the dynamic incentives of the firm are restricted to reacting to state-dependence, but not internalizing the effect of current decisions on future payoffs.

In the case of the monthly premium, firms are allowed to send premium rebates to consumers via their social security checks. However, this is rare and non-existent in the Massachusetts market, despite a significant portion of plans with a premium equal to zero. (See Appendix Figure A1). In the model, I abstract from these micro-foundations and treat both constraints as imposed on the firms.

The marginal cost of insuring a particular beneficiary, mc_{ijt} , depends on the health of the beneficiary and the cost-sharing parameters of the product. Marginal costs are given by

$$mc_{ijt}(X_{jt}) = E_{\omega} \left[\sum_{\tau_t} m_{i\tau}(\omega, X_{jt}) - O_{jt}(X_{jt}, \sum_{\tau_t} m_{i\tau}(\omega, X_{jt})) \right] + a_{jt} \quad (5)$$

$$O_{jt}(X_{jt}, M) = \min \left\{ \text{OOP-Limit}_{jt}, \phi_{jt}(X_{jt}, M)M \right\} \quad (6)$$

where τ_t indexes the months of an individual's enrollment during year t , O_{jt} is the function governing the out-of-pocket costs of medical consumption, ϕ_{jt} is an effective coinsurance rate, and *OOP – Limit* is the maximum allowable out-of-pocket spending of the consumer. The total medical spending on an individual is given by the annual sum of monthly medical consumption, $m_{i\tau}^*(\omega, X_{jt})$. The firm covers all of these expenses less the out-of-pocket prices paid by the consumer, O_{jt} . This function is specified as the minimum of the plan-specific maximum out-of-pocket spending limit and an effective coinsurance rate ϕ_{jt} on the total annual spending. The details of specifying and estimating the effective coinsurance rate are presented in Appendix Section B.

In addition to medical consumption, the products face an additional marginal cost, a_{jt} , which includes the average per-member administrative and drug costs for each firm. These costs are assumed to be constant across all enrollees and constant within firms. This is a strong assumption for prescription drug costs, however these costs tend to be small relative to medical costs. Firm-level average drug costs range from \$60 to \$100, without accounting for the subsidies firms receive for providing that coverage.

The per-person subsidy, b_{ijt} , depends on the risk score of the individual and a “bid”

submitted by the plan, which reflects the plan’s risk-adjusted expected costs.⁹ The bid can be modeled as a function, $bid_j(p_j, X_j)$, which depends on the characteristics of the plan set by the firm, holding fixed all other policy variables that affect plan bids. In Appendix Section C, I provide more detail on the formula for the risk adjusted subsidy and show how the bid function is estimated from the national panel of Medicare Advantage product characteristics and payments.

Equilibrium

An equilibrium in this model, for a given year t , is defined as the set of premiums and cost-sharing parameters, $\{(p_{jt}, X_{jt})\}_j$, such that for every product, j ,

$$(p_{jt}^*, X_{jt}^*) = \arg \max_{(p, X)} \sum_{k \in J_{f(j)}} \Pi_{kt}(p, X, \mathbf{p}_{-jt}, \mathbf{X}_{-jt}) \quad (7)$$

where $J_{f(j)}$ indicates the set of products offered by the firm that offers product j , and all other premiums and product characteristics, $(\mathbf{p}_{-jt}, \mathbf{X}_{-jt}) \equiv \{(p_{kt}, X_{kt})\}_{k \neq j}$, are held fixed.

When setting the level of a cost-sharing term, $x_{jt} \in X_{jt}$, the firm faces a trade off between gaining additional sales and increasing the net cost of selling insurance. Consider the first order condition of a single-product firm.¹⁰

$$0 = \underbrace{\int_i \frac{\partial s_{ijt}}{\partial x_{jt}} (p_{jt} + b_{ijt} - mc_{ijt}) dF(i)}_{\text{Profit from Marginal Sales}} + \underbrace{\int_i s_{ijt} \left(\frac{\partial b_{ijt}}{\partial x_{jt}} - \frac{\partial mc_{ijt}}{\partial x_{jt}} \right) dF(i)}_{\text{Infra-marginal Change in Net Cost}} \quad (8)$$

The two terms of equation (8) each represent the two features of consumer behavior that this model must capture. The first term concerns adverse selection: the profitability of additional (or reduced) sales resulting from a reduction (or increase) in the cost-sharing term of product j . This term depends crucially on the covariance between the elasticity of

⁹The submission is not a bid in the economic sense because there is no auction. There is a strategic component to submitting the bid, which has been studied at length (e.g. Miller et al. 2019, Curto et al. 2021)

¹⁰This exposition assumes the non-negativity constraints are non-binding.

demand for product j with respect to this cost-sharing term and the individual-specific net profit of insurance.¹¹ Thus, the incentive to change a cost-sharing term is related to the degree of adverse selection with respect to that particular term. This is also an important feature of the firm’s optimal monthly premium.

The second term concerns moral hazard: the effect of a change of a cost-sharing term on the medical consumption of consumers, which determines the net marginal cost of insurance. For individuals that do not exceed the out-of-pocket spending limit, the change in gross marginal cost is given by the sum of the expected change in medical consumption at the current out-of-pocket-price and the change in the out-of-pocket price of expected medical consumption.

$$\frac{\partial mc_{ijt}}{\partial x_{jt}} = E_{\omega} \left[\sum_{\tau_t} \frac{\partial m_{i\tau}}{\partial x_{jt}} (1 - \phi_{jt}) + \sum_{\tau_t} m_{i\tau} \frac{\partial \phi_{jt}}{\partial x_{jt}} \right] \quad (9)$$

Due to the subsidy rules of MA, the change in gross marginal cost is reinforced by the change in the subsidies. MA requires that subsidies are decreasing in the cost-sharing terms of insurance. Thus, $\frac{\partial b_{ijt}}{\partial x_{jt}}$ is negative.

These two features—adverse selection in the demand for insurance and moral hazard in medical consumption—are at the center of evaluating how any feature of market structure or market design will impact the levels of cost-sharing provided in the market and the consumption of medical care. In the following sections, I outline the data and methods that allow me to identify these features. Finally, in Section 7, I use the fully estimated model to evaluate the effect of a change in the level of competition via a merger.

¹¹The first term can also be written as

$$\int_i \frac{\partial s_{ijt}}{\partial x_{jt}} (p_{jt} + b_{ijt} - mc_{ijt}) dF(i) = E \left[\frac{\partial s_{ijt}}{\partial x_{jt}} \right] E \left[p_{jt} + b_{ijt} - mc_{ijt} \right] - \text{Cov} \left(\frac{\partial s_{ijt}}{\partial x_{jt}}, p_{jt} + b_{ijt} - mc_{ijt} \right)$$

where E evaluates the mean over all consumers.

4 Data and Descriptive Results

The data come from the Massachusetts All Payer Claims Database and the Medicare Advantage Plan Benefits Data provided by the Center for Medicare and Medicaid Services. In this section, I describe the data and two sets of descriptive results. First, I show that the average effect of competition on cost-sharing terms is negative across most cost-sharing terms in the data. This effect is particularly pronounced for the copays for primary care office visits. Second, I show that over 90% of MA beneficiaries in Massachusetts make at least one office visit a year and spend roughly one quarter of all out-of-pocket spending on office visit copays. These results together suggest that cost-sharing terms are an important margin of competition and that primary care copays are a particularly important mechanism.

4.1 Data

The data on consumer behavior come from the 2013 through 2017 Massachusetts All Payer Claims Database (APCD). For each de-identified enrollee, I observe their sex, zip code, age group, a history of plan enrollment from 2013 to 2017, and the contents of their medical insurance claims during that same period. The medical claims data include information on patient diagnoses, the procedures performed by the physician, the total amount paid by the insurance provider, and the value of any copay, coinsurance, or deductible paid by the patient.

These data serve two key functions. First, they provide detailed information on the health status of each consumer. Using the diagnoses codes that are submitted as a part of each medical claim, I can construct indications for whether each consumer is diagnosed with a set of medical conditions as well as a summary risk score that measures the overall health of the consumer. To construct these measures, I use the Center for Medicaid and Medicare Services Hierarchical Conditions Categories (CMS-HCC) algorithm and risk coefficients. This method has the advantage of being designed to measure clinical conditions that are related to high medical consumption.

Second, the APCD provide a direct measure of medical consumption. The baseline

measure is total medical spending, which is common in the literature (Manning et al. 1987, Brot-Goldberg et al. 2017, Aron-Dine et al. 2015, Ellis, Martins, and Zhu 2017). In addition to this measure, I include an adjusted measure of physician service intensity that removes variation in the payment for a particular service across insurance products and years. The measure is created by predicting the payment for a particular service based on the features of the claims data—e.g. principal diagnoses, procedure modifiers, place of service, etc. I use the least absolute shrinkage and selection operator (LASSO) to select the most relevant features, and estimate the prediction coefficients separately for each procedure in the data. The adjusted measure is computed by predicting the price of each claim for a representative firm and year.

A crucial and novel aspect of this paper is linking the medical claims data to the insurance choices of the beneficiaries. The claims data includes identifiers for the firms and the products. While the names of the Medicare Advantage firms are known in the data, the identity of the products is not. I link the product identifiers in the APCD to the publicly available product information using the county-level enrollment panel in each data set.

The key data on product characteristics come from the Plan Benefit Package (PBP) data. The PBP data contain detailed information (over 1,000 features) that describe the cost-sharing terms and covered services of each insurance plan offered in the Medicare Advantage program. The data provide granular cost-sharing terms that govern each type of service (e.g. primary care, medical devices, or diagnostic lab tests). For instance, it indicates whether a plan has a copay, coinsurance rate, or service-specific deductible for a particular type of service and the value of the cost-sharing term if it applies.

The granular data on cost-sharing terms is important for two reasons. First, Medicare Advantage plans typically have no general medical deductible in which beneficiaries have to pay the full cost of care up to some threshold. Instead, the primary form of cost-sharing are type of service specific copays and coinsurance rates. Second, the cost-sharing terms vary widely across different types of service, which can be seen in Table 1. For example, primary care visits generally cost between \$0 and \$30 per visit while emergency room visits range between \$50 to \$150 per visit.

The data are combined to create two analytical data sets. The first is an annual panel of plan enrollment. Plan enrollment decisions are made when the consumers first become eligible Medicare and each subsequent year during the open enrollment period between January and March. These individuals first become eligible three months before turning 65.¹² This data set includes the available demographic information of the enrollees, a summary risk score that corresponds to their expected spending, and indicator variables for whether the consumers are diagnosed with each of a set of clinical disease categories. The data include every insurance product option for each consumer during each year, an indicator for which product was chosen, and interaction terms between the individual-level demographic and clinical variables and the characteristics of the product. The second data set is a monthly panel of medical consumption. For each month that a consumer is in the data, the baseline measure of consumption is the total medical spending by the insurer and the patient.

For more detail on sample selection, linking the APCD and PBP data, measuring health status, and constructing the adjusted medical consumption measure, see Appendix Section A. The appendix also includes details on using Medical Loss Ratio filings to supplement the claims data with data on administrative and prescription drug expenses.

4.2 Descriptive Results

The first set of results show that, on average, cost-sharing terms are lower (i.e. lower out-of-pocket prices for care) in markets with more competition. To show this in detail, I use data on every county (each a local market) in the US from 2011 through 2019, which contain substantial variation in the level of competition. However, the qualitative facts described in this section are also present in the estimation sample from the 14 counties of Massachusetts between 2013 and 2017.

Table 1 presents mean characteristics for the Medicare Advantage, separated by the number of firms that offer plans in each market. The degree of competition varies widely across counties. Only one firm offers insurance products in 13% of counties, while at least 7

¹²Medicare is also available to some individuals under the age of 65 who are disabled and receive social security benefits, but I exclude these individuals from the sample.

firms offer products in 7% of counties. However, even markets with 10 or more participating firms tend to be concentrated, with the top 2 firms accounting for 60% of total Medicare Advantage enrollment.

In general, the average level of cost-sharing is lower across all types of service in markets with more firms. However, the declines appear to be stronger in some categories more than others. As shown in Table 1, the primary care copay falls consistently across each category for a total decline of 73% from monopoly markets to markets with 10 or more firms. The out-of-pocket spending limit, radiology copay, and inpatient copay also falls by substantial amounts. Other cost-sharing characteristics tend to be lower in markets with more firms, but the relationship may not necessarily be monotonic.

To better evaluate the relationship between competition and cost-sharing characteristics, I follow Bresnahan and Reiss 1991 in using the eligible population and other characteristics of the market as an instrument for the number of firms that decide to enter the market. The intuition behind the first stage of this model is that larger markets can support more firms by allowing firms to spread the fixed costs of entry over more sales. This approach has been used in the health insurance literature to show that more competitive health insurance markets have lower average premiums (Abraham et al. 2017, Dickstein et al. 2015) and that local and national insurance plans are differentiated (Dranove, Gron, and Mazzeo 2003).

The first and second stage of the model are given by equations (10) and (11). The dependent variable, y_{mt}^s is the enrollment weighted cost-sharing characteristic s in county m and year t . The first stage predicts the number of firms that will enter a market as a function of the log of market size, M_{mt} , and a vector characteristics about the county, X_{mt} . The second stage then uses the predicted number of firms as an instrument for competition in the county, with the same sets of controls as the first stage. The county-level characteristics are meant to capture aspects other than market size and competition that may affect the supply of insurance. I include measures that may affect demand (average income, race, and senior employment), the use of health care (disability among seniors and population over 85), and bargaining power with health care providers (the number of primary care doctors and

Table 1: More Competitive Markets have Lower Average Cost-sharing Levels

Number of Firms	1	2 - 3	4 - 6	7 - 10	10+
% of Markets	0.13	0.44	0.36	0.06	0.01
Share of Top 2	1.00	0.94	0.78	0.65	0.60
Eligible Population	4,130	8,740	25,900	72,000	261,000
Enrollment Weighted Characteristics					
Premium (monthly)	35.1	27.2	22.1	16.1	2.4
Part B Rebate	0.13	0.08	0.06	2.64	2.15
Deductible	17.6	20.8	17.9	11.5	4.79
OOP Limit	6590	6090	5640	5530	4700
<i>Copays</i>					
Primary Care	15.5	12.6	10.3	8.28	4.20
Specialist	35.5	34.6	33.6	31.9	13.3
Outpatient	121	102	119	108	46
Radiology	80.6	67.5	58.6	45.5	40.3
Lab Tests	4.31	3.78	4.27	4.33	4.36
Emergency	70.0	67.8	68.1	68.7	62.6
Inpatient	295	272	253	250	137
Ambulance	213	195	191	194	167
<i>Coinsurance Rates</i>					
Outpatient	0.102	0.088	0.059	0.051	0.040
Radiology	0.062	0.062	0.69	0.079	0.047
Med Devices	0.190	0.192	0.180	0.171	0.140
Outpt Drugs	0.163	0.162	0.160	0.163	0.141

Note: Cost-sharing terms are lower on average in counties with more participating firms. The data come from Medicare Advantage plans offered in every US county from 2011 to 2019. Each column represents counties in which a certain number of firms offered plans. The top panel displays market characteristics of those counties, and the bottom panel displays the average level of each product characteristic weighted by the number consumers that select each product.

Table 2: Evidence that Competition Reduces Cost-sharing Levels

First Stage		IV Estimates					
Firms		Prem.	Primary	Spcl.	Emerg.	Radio.	Inpt.
# of Firms		-3.37*** (0.12)	-1.50*** (0.04)	-1.00*** (0.05)	-0.04*** (0.02)	-2.08*** (0.44)	-9.96*** (0.52)
Log Market Size	0.86*** (0.01)						
Income (\$000)	-1.56*** (0.11)	-3.82*** (1.33)	-1.10*** (0.38)	-2.56*** (0.58)	0.95*** (0.18)	-2.55 (4.72)	-4.15 (5.55)
% White	0.00 (0.00)	0.18*** (0.02)	0.00 (0.01)	0.02** (0.01)	0.01 (0.00)	0.48*** (0.08)	-0.25** (0.10)
<i>Among Eligible</i>							
% over 85	1.33*** (0.33)	9.50** (3.90)	15.16*** (1.11)	6.11*** (1.71)	-1.52*** (0.53)	-9.60 (13.81)	66.86*** (16.24)
% Employed	1.27*** (0.29)	1.37 (3.43)	8.83*** (0.98)	5.84*** (1.51)	-0.73 (0.47)	15.81 (12.16)	67.41*** (14.31)
% Cog. Dis.	-0.85*** (0.30)	11.56*** (3.61)	5.13*** (1.03)	1.17 (1.59)	0.42 (0.49)	-61.40*** (12.79)	30.27** (15.04)
<i>Resources (per 1000)</i>							
PC Docs	-0.38*** (0.03)	2.26*** (0.36)	-0.01 (0.10)	-0.36** (0.16)	-0.11** (0.05)	-2.37* (1.28)	1.99 (1.50)
Hosp. Beds	-0.00 (0.00)	0.08*** (0.03)	0.02*** (0.01)	0.05*** (0.01)	0.00 (0.00)	-0.14 (0.10)	-0.15 (0.12)
<i>Fixed Effects</i>							
State & Year	✓	✓	✓	✓	✓	✓	✓
<u>Effect</u>							
Data Mean		-0.18 (0.20)	-0.16 (0.10)	-0.03 (0.01)	0.04 (0.04)	-0.04 (0.03)	-0.04 (0.02)

Note: An additional firm leads to lower cost-sharing levels, and this effect is large for the primary care copay relative to the mean level. As displayed in final row of column three, an additional firm leads to a 16% decline in the primary care copay. The unit of observation is a US county in a given year between 2011 and 2019. The dependent variable is the enrollment weighted average of a product characteristic: prem - monthly premium; primary - primary care copay; spcl - specialist copay; emerg - emergency room copay; radio - radiology copay; and inpt - inpatient copay. Each effect is negative with the exception of the outpatient copay, which may display substitution with an outpatient coinsurance rate.

hospital beds per capita). I also include state fixed effects to control for the local regulatory environment and year fixed effects to control for time trends.

$$N_{mt} = \alpha \log(M_{mt}) + \gamma' X_{mt} + \epsilon_{mt} \quad (10)$$

$$y_{mt}^s = \beta \hat{N}_{mt} + \gamma^{s'} X_{mt} + \epsilon_{mt} \quad (11)$$

The estimation shows that competition has significant and negative effects on many, but not all of the cost-sharing parameters. Table 2 presents the results of this estimation for a selected set of cost-sharing parameters. I find that an additional firm decreases the average primary care copay by \$1.50, 16% of the mean value. Aside from the monthly premium, this is the largest effect relative to the mean. The other cost-sharing parameters generally have significant negative effects of -3 to -4% relative to their mean values. Appendix Tables A1 and A2 display the results for a much broader set of cost-sharing terms.¹³

The next set of results show that primary care is both commonly used and a large portion of out-of-pocket spending. Table 3 displays annual summary statistics on use and out-of-pocket across a number of clinical categories, as identified by Berenson-Eggers Type of Service Codes (BETOS). More than nine out of ten medicare beneficiaries have a office visit, the clinical category for primary care doctor visits, at least once during the year. The next most frequent category of use is specialist visits, which are only used by roughly half of beneficiaries.¹⁴

Despite the copays that typically range from \$0 to \$30 dollars, the mean out-of-pocket spending on office visits is \$116, which suggests that the average beneficiary pays a copay to see the doctor several times throughout the year. As a result, any changes to the primary care copay are felt multiple times over by the beneficiaries. The average out-of-pocket spending on office visits is the largest of any category and constitutes roughly one quarter of all out-of-pocket spending. The category with the next highest amount of out-of-pocket spending

¹³These findings are consistent with Pelech 2018, which finds that a reduction in competition via a large-scale exit of one plan type in Medicare Advantage led to higher expected out-of-pocket spending by the beneficiaries.

¹⁴These clinical categories depend on the procedure code billed by the physician, not the physician's specialty. And while it is likely that office visits are primarily billed as primary care visits and specialist visits primarily billed as such, some office visits may be billed to insurance firms as specialist visits or another cost-sharing category.

Table 3: Primary Care is a Large Component of Out-of-pocket Spending

	% Use	Mean	Out-of-pocket Spending Conditional Mean
Office Visit	0.912	116	128
Specialist Visit	0.516	21.5	41.7
Maj/Min Procedure	0.346	65.2	188
Imaging	0.340	43.9	130
Lab Tests	0.259	12.0	46.5
Emergency Room	0.202	16.3	96.6
Inpatient	0.169	107	695
Ambulance	0.154	25.9	199
Medical Devices	0.130	10.2	80.1
Outpatient Drugs	0.034	6.46	188
Other	0.202	24.4	121.1

Note: The primary care copay is an important aspect of medical consumption for Medicare Advantage beneficiaries. Office visits, the clinical category associated with primary care, are frequently used and make up roughly one quarter of all out-of-pocket spending. The service categories are defined using CPT procedural codes and BETOS service categories. The tables displays the percent of beneficiaries which use that service during the year, total mean out-of-pocket spending on each category by all consumers, and the mean out-of-pocket spending conditional on using the service. The data come from the Massachusetts APCD.

is inpatient hospital stays, which affect only about one in six beneficiaries but feature high prices conditional on use.

Importantly, Medicare Advantage plans typically have no deductible which would require the consumer pays the full cost of care before reaching some threshold. Instead, the primary source of out-of-pocket spending on medical care comes from the copays and coinsurance rates on frequently used services. The summary in Table 3 does not include spending after individuals have reached the maximum out-of-pocket spending limit, but less than 1% of beneficiaries reach that threshold.

These results provide two key facts. First, the average effect of competition is to decrease the level of cost-sharing. As detailed in Section 7.1, this should be expected in an environment where different products do not face widely different costs to insure the same individual.

Second, the primary care copay is an important mechanism in the Medicare Advantage market. The large effect of competition on primary care copays suggest that consumer demand responds elastically to these copays. And the frequency and spending on office visits suggest that primary care copays are also important in determining consumer spending on medical consumption.

These average effects and descriptive facts hide important heterogeneity. In the following sections, I detail how I use these data to identify the key parameters in the model that allow me to characterize product-level best response of cost-sharing terms with respect to a change in the market structure or market design and the subsequent effect on medical consumption and health. Following the results of this section, I will focus on primary care copays as the key strategic aspect of cost-sharing in the counterfactual merger analysis.

5 Estimating Consumer Demand for Insurance

This section outlines the discrete choice model of consumer demand for health insurance. The model follows a logit demand system with switching costs, as is standard in consumer demand for health insurance. Unlike typical demand estimations in this market, I am able to incorporate detailed heterogeneity on consumer health status by linking the diagnosis information in the claims data with insurance choices. The mean estimated semi-elasticity with respect to a \$10 increase in monthly premium is -2.9. This is lower than the mean semi-elasticities with respect to primary care, -12.7, but greater than the elasticities of most other cost-sharing parameters.

5.1 Specification

The model for consumer choices follows a discrete choice logit model with switching costs and rich heterogeneity in consumer health status. Consumers in the model, indexed by i , are characterized by a set of demographic characteristics, $Z_i = \{z_{ig}\}$, where g indexes the consumers' age, sex, an indication of whether the individual is diagnosed with each of a set of clinical conditions, and a summary medical risk score.

Consumers in the local market r and year t choose among a set of J_{rt} products. I assume the products are market-specific: $J_{rt} \cap J_{r't} = \emptyset$, $\forall r \neq r'$. Products are characterized by a monthly premium p_{jt} , a vector of cost-sharing parameters, X_{jt} , a vector of non-financial characteristics, W_{jt} , and an unobserved quality ξ_{jt} . Consumers also face a three-component switching cost, $D_{ijt} = \{d_{ijtk}\}$, where k indicates either a switch to a new product, a switch to TM from MA, or a switch to MA from TM.

The base level of indirect utility from purchasing product j in year t , common across all consumers, is specified as

$$\delta_{jt} = \alpha_0 p_{jt} + \beta'_0 X_{jt} + \gamma'_0 W_{jt} + \xi_{jt} \quad (12)$$

In addition to the base utility, δ , the total indirect utility to a particular consumer depends on their demographics and the switching costs. The total indirect utility, v_{ijt} , that consumer i receives from product j in year t is specified as

$$v_{ijt} = \delta_{jt} + \Upsilon' D_{ijt} + \left(\sum_g \alpha_g z_{ig} \right) p_{jt} + \left(\sum_g \beta_g z_{ig} \right)' X_{jt} + \epsilon_{ijt} \quad (13)$$

where $\epsilon_{ij\tau}$ is an i.i.d. type I extreme value idiosyncratic preference. Consumers have heterogeneous preferences over premium and cost-sharing parameters that depend on the components of their demographics, $z_{ig} \in Z_i$. Importantly, this heterogeneity can capture that consumers with particular medical conditions may seek out plans with specific cost-sharing characteristics that suit their expected medical needs.

Consumers select the plan that maximizes their indirect utility during the year. While there is state-dependence in the choice, via the switching cost terms, consumers are assumed to be myopic and do not consider how state-dependence will affect future decisions. I write s_{ijt} to express the probability that a consumer i selects plan j in year t .

$$s_{ijt} = \Pr \left(v_{ijt} = \max_k v_{ikt} \right) \quad (14)$$

5.2 Estimation

The parameters governing the consumer demand for insurance can be split into those governing consumer heterogeneity, $\theta_z = (\Upsilon, \{\alpha_g, \beta_g\}_g)$ and those governing the base level of product quality $\theta_0 = (\alpha_0, \beta_0, \gamma_0, \{\xi_{jt}\})$. These two sets of parameters are estimated in two stages, following Goolsbee and Petrin 2004.

In the first stage, the parameters governing consumer heterogeneity are estimated with maximum likelihood. Following Berry, Levinsohn, and Pakes 1995, I can compute the values of δ_{jt} for each product given any candidate of the heterogeneity parameter, θ_z , such that the aggregate predicted markets shares of each product precisely match the observed share in the data.

The parameters in θ_z are identified by the correlation between the shares of consumers with a particular value of Z that select products with particular premiums and cost-sharing terms, (p, X) . This requires the assumption, as expressed in equations (12) and (13) that there is no unobserved component of utility that is common to some groups consumers and not to others. This assumption is standard among the demand estimation literature (e.g Goolsbee and Petrin 2004, Tebaldi 2020).

The base parameter vector, θ_0 , is estimated with a linear, two-way fixed effects model using the panel structure of the data. The identifying assumption is that the transient unobserved quality component has the form $\xi_{jt} = \bar{\xi}_j + \nu_{jt}$, and that the first difference of the monthly premium and cost-sharing parameters are uncorrelated with the first difference of the transient unobserved utility component, ν_{jt} .

$$\delta_{jt} - \delta_{jt-1} = \alpha_0(p_{jt} - p_{jt-1}) + \beta'_0(X_{jt} - X_{jt-1}) + \nu_{jt} - \nu_{jt-1} \quad (15)$$

The primary threat to identification is that insurance products change in some unobserved way that is recognized by the consumers and reflected in updates to the premium or a particular cost-sharing characteristic. One potential source of endogeneity are changes in the provider network. However, I find that the networks change very little from year to year. In this sample, 95.6% percent of medical claim reimbursements go to providers that

are also in the network the following year. And among the providers that are no longer in the network, more than two thirds exit the data entirely, which suggests that it may be due to organizational changes rather than true network exits.¹⁵ Given the high dimensionality of insurance contracts, and the difficulty of consumers to assess attributes that I do not observe, these changes are unlikely to be a significant problem (Abaluck and Gruber 2011, Kling et al. 2012).

As a robustness check, I also use Hausman instruments—the average of a particular product’s characteristic in other markets—for the monthly premium and the primary care copay. In addition to this IV strategy, I can also use the insurance firm’s decisions on each characteristic to infer the demand parameters, following the methodology of Petrin and Seo 2019. Because I observe data on marginal cost, the only missing parameters in the firm’s first order condition are demand parameters. Results from this method are forthcoming.

5.3 Results

The implied semi-elasticities of demand are summarized in Table 4. Consumer demand is most elastic with respect to the primary care, followed by the monthly premium. Consumers are largely inelastic with respect to other cost-sharing parameters, with the estimates either close to zero or statistically insignificant. The coefficient estimates for the base-level of indirect utility are presented in Table A4 using both the two-way fixed effect approach and Hausman-style instruments. Because the results are of similar magnitude and the instrumental variable specification has larger standard errors, I use the two-way fixed effects results as the baseline specification.

The mean semi-elasticity with respect to premium is -2.9. This implies that the average consumer is 2.9 percent less likely to select a plan given a 10\$ increase in the monthly premium. This elasticity is low relative to consumer preferences over the copays for primary care, which has a mean semi-elasticities of -12.7. This may be due to the fact that more than 90% of the consumers make an office visit during the year, and typically multiple times

¹⁵There is more variation at the product level, but 96% of member months are in plans where at least 90% of reimbursements are paid to providers that are in network the following year.

Table 4: Insurance Demand Responds Elastically to the Primary Care Copay

	All Consumers		Entering Consumers	
	Semi-elasticity	Std. Error	Semi-elasticity	Std. Error
Premium	-2.91	0.54	-10.9	1.73
Primary Care	-12.74	2.36	-37.6	7.54
OOP Limit	5.80	2.73	18.8	7.59
Specialist	4.94	4.00	9.21	12.8
Outpatient	0.57	0.52	2.51	1.68
Outpatient Coins	2.67	1.02	8.84	3.28
Inpatient Stay	-0.05	0.10	0.73	0.30
Emergency Room	-0.09	0.96	0.03	3.06
Ambulance	-1.01	0.27	-3.20	0.85
Medical Devices	-0.92	0.43	-2.55	1.36
Outpatient Drugs	1.02	0.31	3.15	0.99
Imaging	-1.18	0.39	-2.90	1.28

Note: Consumer demand is most elastic with respect to the primary care, followed by the monthly premium. The tables shows the mean semi-elasticities of demand with respect to each cost-sharing term in the demand estimation, both for all consumers and for entering consumers that face no switching cost. Each characteristic denotes a copay with the exception of premium, the out-of-pocket limit, and the outpatient coinsurance. All semi-elasticities represent the percent change in the probability a consumer purchases their chosen plan given a \$10 increase in the characteristic. For the outpatient coinsurance rate, the elasticity is computed with respect to a single percentage point increase.

throughout the year. Therefore, consumers likely expect to pay the copays for primary care and specialists several times throughout the year.

The low elasticities are also due in part to the sizeable switching costs. For entering consumers that face no switching costs, the semi-elasticities for premium and primary are copays are -10.9 and -37.6, respectively. Consumers face an average switching cost of \$640 per month, which is much greater than the average monthly premium.

Consumer heterogeneity depends on age, sex, the six most common clinical conditions (listed in order from most to least prevalent), the aggregate risk score, and the aggregate risk score squared. The estimates for a select number of cost-sharing terms are presented in Table 5, and the remainder are presented in Appendix Table A3. These estimates show that demand for insurance depends on consumer health in important ways that go beyond aggregate measures of health status. For instance, consumers tend to have low sensitivity to cost-sharing for out-patient drugs, which governs chemotherapy. The exception is consumers that have been diagnosed with breast or prostate cancer.

Table 5: Health Status is an Important Determinant of Insurance Preferences

	Premium (\$10)	Copays (\$10)				
		Primary	Specialist	Outpatient	Inpatient	Imaging
Over 75	0.027*** (0.003)	0.000 (0.021)	0.132*** (0.026)	-0.004 (0.004)	-0.021*** (0.001)	-0.017*** (0.003)
Female	0.008*** (0.003)	-0.069*** (0.019)	0.053** (0.024)	0.000 (0.003)	-0.002 (0.001)	-0.008*** (0.002)
Heart Arrhythmia	-0.018*** (0.004)	0.065** (0.033)	0.008 (0.038)	0.008 (0.005)	-0.017*** (0.002)	-0.010** (0.004)
Vascular Disease	0.009* (0.005)	0.080** (0.033)	-0.037 (0.040)	0.016*** (0.005)	-0.023*** (0.002)	-0.005 (0.004)
Diabetes w/ Compl.	0.024*** (0.005)	-0.044 (0.034)	-0.122*** (0.040)	-0.001 (0.005)	0.005** (0.002)	-0.011*** (0.004)
Diabetes w/o Compl.	0.010** (0.005)	-0.001 (0.033)	0.077* (0.040)	-0.001 (0.006)	0.002 (0.002)	-0.002 (0.004)
Breast/Prost. Cancer	0.006 (0.006)	0.063 (0.039)	0.002 (0.047)	0.010 (0.006)	-0.017*** (0.002)	-0.013** (0.005)
Rheum. Arthritis	0.022*** (0.007)	0.036 (0.047)	-0.082 (0.055)	0.003 (0.007)	-0.008*** (0.003)	-0.018*** (0.006)
Agg. Risk Score	0.003 (0.003)	-0.162*** (0.024)	0.044 (0.030)	-0.034*** (0.004)	0.018*** (0.001)	0.001 (0.003)
Agg. Risk Score ²	0.001 (0.000)	0.013*** (0.003)	-0.002 (0.004)	0.004*** (0.000)	-0.002*** (0.000)	0.001 (0.000)

Note: Demand for insurance is heterogeneous in the observed measures of health status. This table displays the coefficients of the demand estimation that govern the heterogeneity in demand for insurance. Negative values for copays imply that consumers are more willing to pay a high monthly premium in order to have a low level of cost-sharing in that category. The significance stars ***, **, and * denote statistical significance at the 1%, 5%, and 10% level respectively.

The median willingness to pay to reduce the primary care copay is relatively large at

\$43 per month. If medical consumption were inelastic and deterministic, this would only be rationalized if the median used a primary care service more than four times per month, which is substantially higher than observed in the data. The average Medicare Advantage beneficiary has a primary care office visit about 7 to 8 times per year.¹⁶ This suggests that consumers may be willing to pay for lower cost-sharing terms in anticipation of its affect on their medical consumption.

The relationship between health status and demand for insurance is summarized in Figure 1, which plots average willingness to pay to reduce the primary care copay by \$10 across the distribution of total and net cost. As expected, the willingness to pay for low cost-sharing increases with the total cost of the consumers (Figure 1a). This reflects the key feature of adverse selection: those with the highest value of insurance are also the most costly to insure. However, after accounting for risk adjusted subsidies, the forces of selection are not so clear. Figure 1b shows that willingness to pay is U shaped in net cost. This is likely due to the fact that cost variance spending grows with the mean. As a result, consumers with high expected spending provide opportunities for both adverse and advantageous selection relative to the risk adjusted subsidy. This phenomenon is also documented by Brown et al. 2014.

6 Estimating Elasticities of Medical Consumption

This section outlines the model for medical consumption given plan benefits and the heterogeneous health of consumers. The model follows the literature in estimating a log-linear demand equation for medical consumption (Buntin and Zaslavsky 2004, Aron-Dine et al. 2015, Ellis, Martins, and Zhu 2017). The elasticity of consumption with respect to primary care copays can be identified through year-to-year changes in the copay within insurance products and inertia in consumer choice. I find that the semi-elasticity with respect to a \$10 increase in the primary care copay is -5.4%. I also find suggestive evidence that a \$10 increase in the primary care copay is associated with a 0.1 percentage point increase in

¹⁶This the average number of office visit claims during the year, which is likely an upper bound on the actual number of office visits.

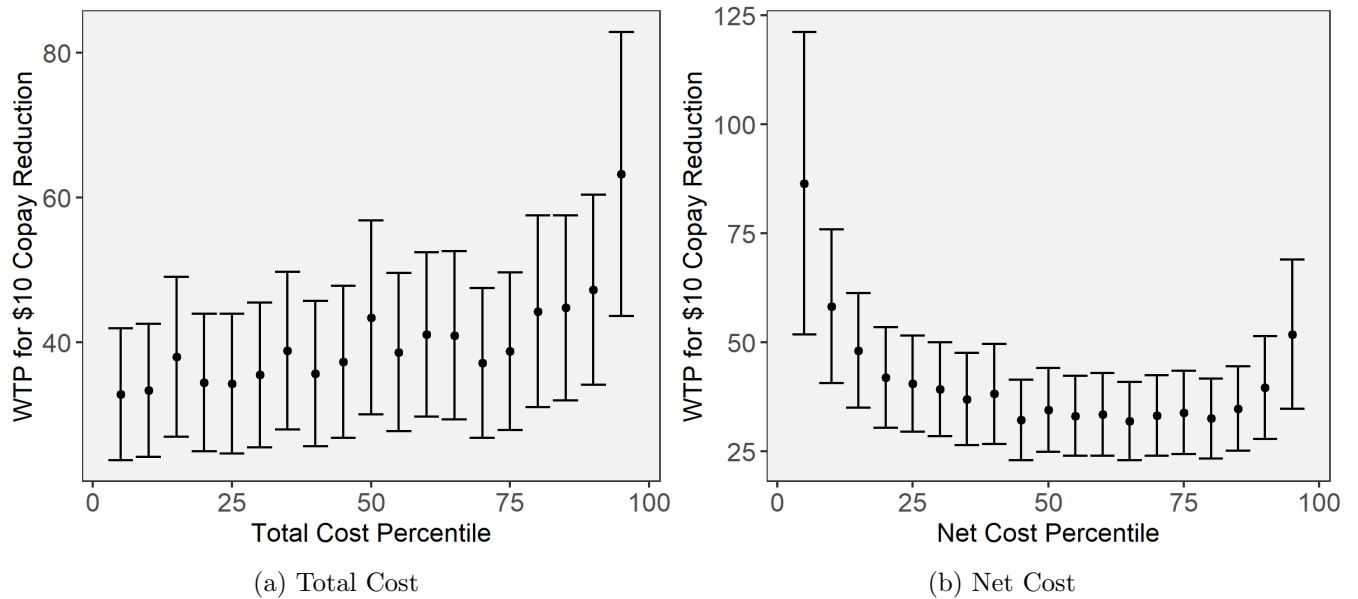


Figure 1: Adverse Selection Plays an Important Role in Insurance Demand

Note: The willingness to pay for low primary care copays is increasing in gross cost but U-shaped in net cost. This figures shows the average willingness to pay for a \$10 reduction in the primary care copay at each percentile of expected cost. The left panel plots willingness to pay across the distribution of gross costs. The right panel plots willingness to pay across the distribution of net cost, after accounting for risk adjusted subsidies.

inpatient mortality.

6.1 Specification

The model of medical consumption is log-linear in plan characteristics, an unobserved individual fixed effect, monthly fixed fixed effects, and an idiosyncratic medical demand error. Let $m_{i\tau}$ be a measure the total medical spending of an individual i in month τ . Let $X_{j(i)\tau}$ be the vector of cost-sharing parameters of product j , in which individual i is enrolled in month τ . Each individual has a constant idiosyncratic health status, η_i and a monthly idiosyncratic medical demand $\omega_{i\tau}$.

Medical consumption is specified as

$$\log\left(m_{i\tau} + \frac{1}{12}\right) = \eta_i + \beta' X_{j(i)\tau} + \lambda_\tau + \gamma' F_{j(i)} + \omega_{i\tau} \quad (16)$$

where λ_t and $F_j(i)$ are month and firm fixed effects. Unless necessary, I will simplify the $j(i)$ notation to j .

The log specification of medical consumption follows a long literature on predicting medical expenditures and estimating elasticities (Manning et al. 1987, Aron-Dine et al. 2015, Ellis, Martins, and Zhu 2017). I follow Ellis, Martins, and Zhu 2017 in using $m_{i\tau} + \frac{1}{12}$ in order to allow elasticities to be comparable to annual elasticity estimates that use $m_{i\tau} + 1$.¹⁷

6.2 Estimation

The central obstacle to consistently estimating the elasticity of patients with respect to the primary care copay is that individuals may select into plans with certain cost-sharing characteristics with knowledge of their future medical needs. The two-way fixed-effect regression specified in equation (16) may produce biased estimates of β because a potential correlation between $X_{j(i)\tau}$ and $\omega_{i\tau}$.

The standard identification approach in the literature, which uses non-linear nature of health insurance contracts during the benefit year to identify the elasticity of consumers to cost-sharing parameters, is not suitable for in this setting (Brot-Goldberg et al. 2017, Aron-Dine et al. 2015). Estimates that exploit non-linearities in the contract recover elasticities that are local to those non-linearities. This is particularly important in the context of Medicare Advantage, where deductibles are typically zero and the out-of-pocket maximum is only reached by 0.5% of beneficiaries. These consumers, who typically have among the highest total medical spending, may not be representative of the whole population. These elasticities are likely not the same as those internalized by firms setting cost-sharing parameters which affect spending throughout the year.

¹⁷While more than 90% of beneficiaries use some medical service during the year, only about 60% of beneficiaries have non-zero spending in any given month.

This paper exploits the plausibly exogenous variation for within-product changes in cost-sharing terms to consistently identify the elasticity. In order for this strategy to be valid, it must be the case that consumers have some degree of inertia in their plan choice and do not optimally select a new plan each year. Fortunately, there is a large literature on consumer inattention and switching costs that documents this to be the case (Heiss et al. 2016, Ho, Hogan, and Scott Morton 2017, Miller et al. 2019, Drake, Ryan, and Dowd 2020). As a result, the change in the primary care copay for the product that an individual was enrolled in during the prior year is a strong predictor of the change in the individual’s actual primary care copay.

This approach has the benefit of using the variation that firms are interested in when making product design decisions: the change in medical consumption caused by a change in a product’s cost-sharing term. However, the estimation abstracts from the relationship between short-run elasticities and dynamic price responses, which may be important for consumers that expect to approach the out-of-pocket spending limit. In a forthcoming robustness exercise, I re-estimate the model 12 times using only observations from one month during the year in each regression.

Identification is focused on the consistent estimation of the primary care copay, which was shown in Section 5 to be an important aspect of insurance demand. Additionally, year to year changes in the primary care copay—the key source of identifying variation—are frequent. In Table 6, I show that 65% of the sample experience are at some point enrolled in a product that changes its primary care copay in the following year. The members that experience an increase in their product’s primary care copay in a particular year are similar to those that do not in their demographics and medical risk.¹⁸ There is only one instance of a product reducing its copay, which affected only a small portion of consumers. Consumers in insurance plans that increase the primary care copay have a lower average switching rate. This disparity is partially mechanical: primary care copay increases only occur in plans that are offered in consecutive years, while a portion of the total switches in a year are due to plans that are no longer offered.

¹⁸While these groups are not required to be as good as randomly assigned, the similarity on observed attributes should provide readers some confidence that they experience parallel trends.

Table 6: Changes in Primary Care Copays are Common

	Product Copay Change		
	Increase	Decrease	No Change
Unique Members	100,244	1,220	57,410
Member - Months (000s)	1,135	13	4,328
Switch Rate	0.031	0.046	0.078
Copay Change			
Product-level	11.92	-5.00	–
Consumer-level	12.13	-5.03	-0.086
Risk Score	1.040	0.79	1.032
Female	0.588	0.506	0.582
Over 75	0.526	0.247	0.473

Note: Roughly 65% of unique Medicare Advantage beneficiaries are enrolled in an insurance plan which changes its primary care copay for the following year, providing the key source of variation in estimating the medical consumption elasticity. This table shows summary statistics for consumers that experience increases, decreases, and no change in their primary care copay. The unit of observation used to compute averages is a member month.

Equations (17) and (18) formally describe the estimated model. To be explicit, I separate the primary care copay from the vector of cost-sharing terms, $x \in X$.

$$\Delta_i \tilde{m}_{i\tau} = \beta_o \widehat{\Delta_i x_{j(i)\tau}} + \beta'_{-o} \Delta_i X_{j(i)\tau} + \Delta_i \lambda_\tau + \gamma' \Delta_i F_j(i) + \Delta_i \omega_{i\tau} \quad (17)$$

$$\Delta_i x_{j(i)\tau} = \rho \Delta_j x_{j(i)\tau} + \mu_{i\tau} \quad (18)$$

The Δ_i operator represents a 12-month, forward difference at the individual level. For example, $\Delta_i x_{j(i)\tau}$ is the difference in the primary copay applicable to consumer i in month τ and month $\tau + 12$. The Δ_j operate is a 12-month, forward difference at the product level. For example, $\Delta_j x_{j(i)\tau}$ is the difference in the copay of product j in time τ (at which time consumer i is enrolled in product j), and the copay of product j in time $\tau + 12$, regardless of whether or not consumer i remains enrolled in that product. The first stage of the estimation uses the latter operator as an instrument for the former.

While the identification strategy addresses the possibility of selection on the primary

care copay, the model relies on the assumption that product level changes in the primary care copay are exogenous with respect to individual medical needs, $E[\Delta_j x_{j(i)\tau} \Delta_i \omega_{i\tau}] = 0$. One way in which this assumption may be violated is if specific products or firms foresee a cost increase via higher negotiated rates with a physician group, for example. This would appear as an increase in the medical expenditures of its beneficiaries, given a particular quantity, and could bias the response to a change in the copay. I address this concern by including a specification that uses predicted physician services spending given the characteristics of the procedures and patients as a measure of medical consumption. Appendix Section A.4 details how this measure is constructed.

6.3 Results

Table 7 displays the results of the medical consumption estimation. I estimate six versions of the medical consumption model. First, I estimate the model separately for four groups of consumer separated into approximate risk quartiles. I divide the sample based on the pooled distribution of aggregate risk scores with one adjustment: the first quartile is expanded to include all individuals with no clinical diagnoses. The only differences in this group are due to age, which is not measured precisely in the data. The risk quartiles are defined by the aggregate risk score of the consumer in the first period of the 12-month difference.¹⁹ Finally, I estimate the model on the full sample for both the baseline and adjusted measure of medical consumption.

The mean semi-elasticity of medical spending with respect to the primary care copay is -5.4%. All variables are copays denominated in \$10, with the exception of outpatient coinsurance (percentage points) and out-of-pocket limit (dollars). The coefficients represent the semi-elasticities, or the percentage decrease in the dependent variable associated with a unit increase in the independent variable.

I find that the semi-elasticity of primary care copays is lowest among the consumers with the lowest and highest medical risk. The right panel of Table 7 contains the results

¹⁹Risk scores are persistent but do vary over time. 51.3% of consumers are in a single risk quartile throughout the entire sample and 89.7% are in two or fewer.

Table 7: Medical Consumption Responds to Primary Care Copays

	Risk Quartiles				Full Sample	
	1 st	2 nd	3 rd	4 th	Baseline	Adjusted
Primary Care	-0.029*** (0.001)	-0.058*** (0.002)	-0.058*** (0.002)	-0.039*** (0.002)	-0.054*** (0.001)	-0.077*** (0.001)
Specialist	-0.004*** (0.001)	0.001 (0.001)	-0.021*** (0.001)	-0.063*** (0.001)	-0.033*** (0.001)	-0.009*** (0.001)
Outpatient	-0.037*** (0.000)	-0.044*** (0.000)	-0.042*** (0.000)	-0.039*** (0.000)	-0.036*** (0.000)	-0.040*** (0.000)
Outpatient Coins	-0.024 (0.049)	-0.241*** (0.070)	-0.093 (0.061)	0.174* (0.082)	-0.019 (0.030)	-0.055 (0.030)
Inpatient Stay	0.000 (0.000)	-0.003*** (0.000)	-0.004*** (0.000)	-0.004*** (0.000)	-0.003*** (0.000)	0.001*** (0.000)
Emergency Room	-0.018 (0.009)	-0.009 (0.013)	-0.089*** (0.010)	0.498*** (0.011)	0.162*** (0.005)	0.019*** (0.005)
Ambulance	-0.011*** (0.000)	-0.010*** (0.000)	-0.008*** (0.000)	-0.000* (0.000)	-0.009*** (0.000)	-0.012*** (0.000)
Medical Devices	0.034*** (0.001)	0.041*** (0.002)	0.048*** (0.002)	0.045*** (0.002)	0.039*** (0.001)	0.038*** (0.001)
Outpatient Drugs	-0.001 (0.001)	-0.002 (0.002)	-0.005*** (0.002)	0.003 (0.002)	0.001 (0.001)	-0.007*** (0.001)
Imaging	0.004*** (0.000)	0.001** (0.000)	-0.001*** (0.000)	-0.003*** (0.000)	0.004*** (0.000)	0.009*** (0.000)
OOP Limit	-0.004*** (0.000)	-0.003*** (0.000)	-0.003*** (0.000)	-0.004*** (0.000)	-0.004*** (0.000)	-0.005*** (0.000)
Monthly Fixed Effects	✓	✓	✓	✓	✓	✓
Firm Fixed Effects	✓	✓	✓	✓	✓	✓
N (000s)	1,875	992	1,404	1,205	5,478	5,478

Note: All consumers are responsive to primary copays in their medical consumption, but those with the lowest and highest risk scores are the least elastic. The tables shows the results of the medical consumption elasticity estimation, where the baseline measure of medical consumption is total monthly spending. All variables are copays denominated in \$10 with the exception of outpatient coinsurance (percentage points) and OOP limit (dollars). In each regression, the primary care copay coefficient is identified using the product-level 12-month change as an instrument. The left panel displays the results of separately estimating each risk quartile of consumers, defined the aggregate risk score in the first period of the 12-month difference. The first quartile is expanded to include all individuals without any HCC diagnoses. The right panel displays the pooled estimation as well as a specification that uses the adjusted measure of physician work intensity as a measure of medical consumption.

for each risk quartile. This U-shaped elasticity-health relationship is likely a result of two countervailing forces. Individuals in the first quartile of risk visit the doctor the least often, and as a result, an increase in the copays for each doctor visit is not likely to have a large effect on their expected annual out-of-pocket costs. The out-of-pocket price increase for consumers in the higher risk quartiles is more substantial because they are charged the copay many times a year. Of course, these consumers are also medically needy and perhaps less likely to forgo necessary medical care. In the highest risk quartile, this second force takes over and consumer elasticity falls.

When using the adjusted measure of medical consumption that controls for provider price differences across products and years, I find that the semi-elasticity is greater but similar: -7.7% relative to -5.4% in the baseline measure. This suggests that the firm's responding to provider price increases may be biasing the baseline estimates downward. In the counterfactual analysis, I use the estimates using the baseline measure instead of the adjusted measure because it is not straightforward to convert the adjusted measure into the actual costs incurred by the insurance firms. I believe this is still reasonable given that the two elasticity estimates are quantitatively similar.

The mean implied coinsurance elasticities range between -0.09 and -0.25. This measure is more comparable to standard estimates in the literature and can be computed by dividing the medical consumption elasticity estimate by the effect of the copay on the effective coinsurance rate (0.033), and multiplying by the mean effective coinsurance rate, 10.5%. Manning et al. 1987 find an overall coinsurance arc-elasticity of -0.2 in the Rand Health Insurance experiment. More recent estimates that target more specific types of service, including primary care office visits find elasticities of similar magnitudes (Aron-Dine et al. 2015, Ellis, Martins, and Zhu 2017). Details on estimating the effect of the primary care copay on the effective coinsurance rate are in Appendix Section B.

6.4 Cost-sharing Parameters and Health

Policy makers are not only concerned about the cost of medical care but also the resulting health of its beneficiaries. It is not a trivial task to identify the effect of changes in

cost-sharing parameters on consumer health. In this section, I present a model that takes advantage of the level of detail available in claims data to provide suggestive evidence on the relationship between primary care copays and inpatient mortality. However, because the identification strategy discussed in Section 6.2 cannot be applied to a dependent variable like mortality, these results are suggestive.

Let $d_{i\tau+s}$ be an indicator variable that represents whether consumer i has died in an inpatient facility, i.e. a hospital or hospice facility, by month $\tau + s$. I specify the following linear probability model.

$$d_{i\tau+s} = \eta_i + \beta^{s'} X_{j\tau} + \lambda_\tau^s + \zeta_j^s + \omega_{i\tau}^s \quad (19)$$

$$\eta_i = \Gamma^{s'} Z_i + \nu_i^s \quad (20)$$

Since individual inpatient mortality is an absorbing state, the estimation equation cannot be differenced to control for the individual fixed effect as in the medical consumption equation. Instead, I specify individual health status as function of Z_i , a vector of demographics and clinical conditions. I include 50 clinical conditions that have higher than a 0.5% prevalence in the Medicare Advantage population.

Figure 2 shows the coefficient estimates and confidence intervals for the relationship between a \$1 increase in the primary care copay and inpatient mortality. Figure 2a shows the absolute effect size and Figure 2b shows the magnitude of the association relative to the mean level of inpatient mortality for a given time in the future. The estimates show a steadily increasing absolute effect which is consistent at nearly every date with a 0.6% percent increase in inpatient mortality. At twelve months, this corresponds to roughly a 0.01 percentage point increase in inpatient mortality.

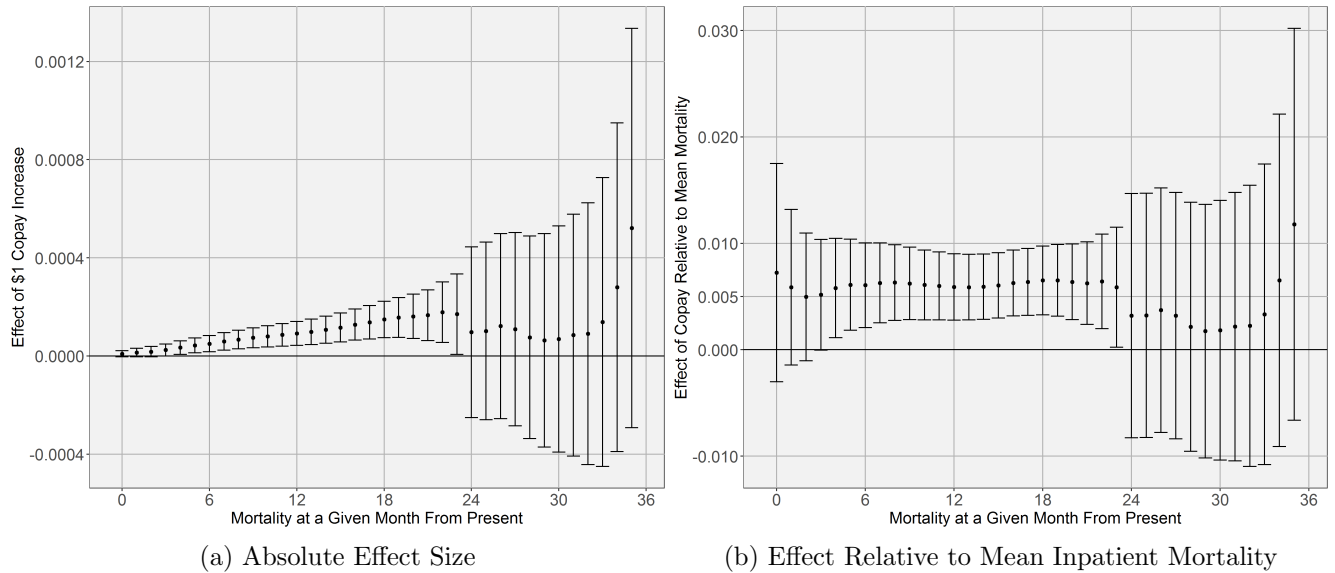


Figure 2: Higher Primary Care Copays are Associated with Higher Inpatient Mortality

Note: A \$1 increase in the primary care copay is associated with a 0.6% increase in the 12-month inpatient mortality rate. The left panel plots the coefficients and 95% confidence interval from estimation of equation (19), where the x-axis indexes the dependent variable of the regression from concurrent inpatient mortality to 36-month inpatient mortality. The right panel presents these results normalized by the mean inpatient mortality rate in the sample.

7 Merger Analysis

To assess the effects of competition, I study three counterfactual mergers between two of each of the three largest firms in the Medicare Advantage market in Massachusetts: Tufts Health Plan (Tufts), Blue Cross Blue Shield of Massachusetts (BCBS), and United Healthcare (United).

7.1 Effect of a Merger

The effect of a merger on the cost-sharing terms and the premiums of the merging products is ambiguous for two reasons. The first follows from the incentives facing a firm that competes in both price and a non-price quality that consumers value. The level of cost-sharing (or more generally, any product quality) that firms will provide depends on the willingness to pay for low cost-sharing among the marginal consumers and the cost of providing low cost-sharing.

A reduction in competition through firm exit or a merger may alter this trade-off in either direction. For example, suppose the consumers of a particular firm have a below average willingness to pay for low cost-sharing and instead prefer low premiums. If this firm exits, the remaining firms may lower premiums, offset by cost-sharing increases, in order to attract these consumers. Of course, this incentive would be flipped if this group of consumers instead had a greater-than-average willingness to pay for low cost-sharing. This observation has been theoretically and empirically investigated in many different settings (Spence 1975, Schmalensee 1979, Hörner 2002, Matsa 2011).

The second reason concerns the characteristics of consumers. To illustrate this, consider the first order condition of a single product firm j with respect to the primary care copay, $x_j \in X_j$. For exposition, I assume that the non-negativity constraints are non-binding and drop the t subscripts.

$$0 = \underbrace{\int_i s_{ij} \left(\frac{\partial b_{ij}}{\partial x_j} - \frac{\partial mc_{ij}}{\partial x_j} \right) dF(i)}_{\text{Infra-marginal Profit Increase}} - \underbrace{\int_i \frac{\partial s_{ij}}{\partial x_j} (mc_{ij} - b_{ij} - p_j) dF(i)}_{\text{Loss of Diverted Consumers}} \quad (21)$$

If firm j increases the primary care copay, it will earn additional profit on the consumers enrolled in its product through a reduction in net cost—first term in equation (21)—and lose the profit of marginal consumers which decide not to purchase its product—the second term. The first order condition is met when the firm optimally balances these two considerations.

The primary care copay of firm j also affects the other firms in the market. Consider another single-product firm k . An increase in the primary care copay of firm j will cause some individuals to decide to purchase from firm k instead. In Farrell and Shapiro 2010, this effect is called upward pricing pressure in the context of mergers in standard product market, because in those standard markets, this effect will always lead to an increase in the price after a merger. The concept was further generalized to generalized pricing pressure by Jaffe and Weyl 2013 to incorporate non-Bertrand competition.

Because the effect concerns the cost-sharing terms as well as the premium and the pres-

sure will not always be upward, this paper will refer to the effect on the profit of firm k from an increase in characteristic l of firm j is referred to as the *diverted average profit* (DAP_{jk}^l). For exposition, I will write the expression for the primary care copay, $\text{DAP}^{\text{copay}}$.

$$\text{DAP}_{jk}^{\text{copay}} = \int_i \frac{\partial s_{ik}}{\partial x_j} (p_k + b_{ik} - mc_{ik}) dF(i) \quad (22)$$

In a standard market, DAP_{jk}^l is positive because firms operate at a price that exceeds marginal costs and the costs do not depend on the characteristics of a consumer. In this case, if firms j and k were to merge, the new merged firm would internalize the additional *benefit* to firm k from an increase in the copay. The tradeoff expressed in equation (21) would no longer be balanced at the pre-merger equilibrium and the merged firm would have an incentive to raise price.

However, in the market for health insurance many consumers do not generate expected profit. And importantly, these consumers may have systematically different preferences for insurance, as is the case in the presence of adverse selection. As a result, there may be pairs of products jk and characteristics l for which DAP_{jk}^l is negative.

If DAP_{jk}^l is negative, an increase in the copay of product j *lowers* the profit of firm k because the consumers that switch to product k generate more cost than revenue. If firms j and k were to merge, the new firm would internalize this cost. The trade-off in equation (21) would again be unbalanced but now in the opposite direction, and the merged firm would have an incentive to raise total profit by lowering the copay. The intuition is that high copays can be rationalized in part because they divert costly consumers to other products. This reduction in average cost helps to offset the reduced sales. However, the merged firm no-longer realizes this reduction in average cost, as those consumers still select one of the firm's products. Thus, the optimal primary care copay is lower.

To see this formally, the post-merger first order condition for product j given a merger with product k can be written as the combination of the original first order condition

and $\text{DAP}_{jk}^{\text{copay}}$. The equation below is written in terms of marginal revenue and marginal cost.

$$\underbrace{\frac{1}{\frac{\partial s_j}{\partial x_j}} \int_i s_{ij} \left(\frac{\partial b_{ij}}{\partial x_j} - \frac{\partial mc_{ij}}{\partial x_j} \right) dF(i) + p_j}_{\text{Marginal Revenue}} = \underbrace{-\frac{1}{\frac{\partial s_j}{\partial x_j}} \int_i \frac{\partial s_{ij}}{\partial x_j} (mc_{ij} - b_{ij}) dF(i)}_{\text{Marginal Cost}} - \frac{\text{DAP}_{jk}^{\text{copay}}}{\frac{\partial s_j}{\partial x_j}} \quad (23)$$

The left-hand side of the equation represents the marginal revenue, the first term on the right represents marginal cost, and the final term is the DAP divided by the total number of diverted consumers. This final term is analogous to the upward pricing pressure in the antitrust literature (Farrell and Shapiro 2010).

Finally, while a merger between two particular products is unclear, the average effect of competition may be less ambiguous. At the pre-merger equilibrium, the average profitability to firm j of all the diverted consumers (to product k as well as other products and the outside good) must be positive. This term is represented by the loss of diverted consumers in equation (21). As long as higher primary care copays generate additional revenue per consumer—the infra-marginal profit increase is positive—then a pre-merger equilibrium must have positive loss from the diverted consumers. If all products were identical in their per-consumer expected profit, this would imply that the mean DAP is also positive. However, this may not be the case if different firms and products face different expected costs to insure the same individual.

Importantly, this mechanism exists for the premium as well as the cost-sharing terms. In Ryan 2020, I investigate this aspect of competition in the Affordable Care Act exchanges where product characteristics are fixed and firms compete primarily on the monthly premium. However, it is more likely that DAP_{jk}^l will be negative for characteristics on which there is more selection in consumer demand. For instance, if preferences for the primary care copay are tied more closely to health status than preferences for premium, we would expect to see more variance in $\text{DAP}_{jk}^{\text{copay}}$ than in $\text{DAP}_{jk}^{\text{prem}}$. And because the firms set both features in equilibrium, the DAP for the primary care copay will affect the firms optimal premium, and vice versa.

7.2 Approximating a Merger

This paper follows the local approximation approach of Jaffe and Weyl 2013 to predict the effect of a merger. This approach is desirable for two reasons.

First, it uses only the local properties of the pre-merger equilibrium: first and second derivatives of the demand and cost of each insurance products. This requires less reliance on the functional form assumptions of demand for insurance and medical consumption. Rather, the approach focuses on making predictions from the estimated elasticities that make up the crucial mechanisms, as outlined in Section 7.1.²⁰ Second, it avoids the computational difficulty of locating an equilibrium in a market where firms with many products choose two strategic variables per product that each affect the costs of the firm and its competitors.

To reiterate the firm's problem expressed in equations (21) and (23), consider a single product firm j . I will write the pre-merger first order condition of the firm with respect to the primary care copay, $x_j \in X_j$, as

$$0 \leq f_j^x(\mathbf{P}, \mathbf{X}) \equiv \frac{1}{\frac{\partial s_j}{\partial x_j}} \int_i s_{ij} \left(\frac{\partial b_{ij}}{\partial x_j} + \frac{\partial mc_{ij}}{\partial x_j} \right) dF(i) + p_j + \frac{1}{\frac{\partial s_j}{\partial x_j}} \int_i \frac{\partial s_{ij}}{\partial x_j} (mc_{ij} - b_{ij}) dF(i) \quad (24)$$

where \mathbf{P} and \mathbf{X} are the vectors of the premiums and primary care copays of all the products in a market. There also exists an analogous pre-merger first order condition with respect to price given by $0 \leq f_j^p(\mathbf{P}, \mathbf{X})$.

Now consider a merger between j and another single product firm k , the post-merger first order condition can be expressed as $h_j^p(\mathbf{P}, \mathbf{X}) \equiv f_j^p(\mathbf{P}, \mathbf{X}) + g_{jk}^x(\mathbf{P}, \mathbf{X})$ where g represents the diverted average profit normalized by the total amount of diverted consumers.

$$g_{jk}^x(\mathbf{P}, \mathbf{X}) = - \frac{\text{DAP}_{jk}^{\text{copay}}}{\frac{\partial s_j}{\partial x_j}}$$

Analogous functions, h_j^p and g_{jk}^p , exist for the merger incentive and post-merger first

²⁰The functional form assumptions still play a role in determining the second order properties of demand and cost.

order conditions with respect to the premium.

The insight of Jaffe and Weyl 2013 is that the changes to a firm's first order condition due to a merger, g , affects the firm's decision in the same way as a change in their marginal cost. The incentive will be passed-through to consumers via a pass-through matrix. The methodology involves characterizing the merger incentives, which depend on the diverted average profit for premiums and cost-sharing parameters, and the pass-through matrix of the merged firm.

This methodology can be easily extended to a setting where firms choose both a premium and a component of the cost-sharing parameter vector. I focus the counterfactual on the monthly premium and the primary care copay, and assume that firms hold all other aspects of their insurance plans fixed.

Jaffe and Weyl 2013 show that pre-merger premium pass-through can be characterized as the derivative of the first order conditions, $\frac{\partial f^p}{\partial p}$, and that post-merger premium pass-through is the derivative of the post-merger first order conditions, $\frac{\partial h^p}{\partial p}$. In Appendix Section D, I show that the logic easily extends to the combined first order conditions of premium and copay.

In the case that all product first-order conditions are binding with equality in the pre-merger equilibrium, a first-order approximation of the merger can be expressed as

$$\begin{bmatrix} \Delta p_j \\ \Delta x_j \end{bmatrix} = \begin{bmatrix} \frac{\partial h_j^p}{\partial \mathbf{P}} & \frac{\partial h_j^x}{\partial \mathbf{P}} \\ \frac{\partial h_j^p}{\partial \mathbf{X}} & \frac{\partial h_j^x}{\partial \mathbf{X}} \end{bmatrix}^{-1} \bigg|_{(\mathbf{P}_0, \mathbf{X}_0)} g_{jk}(\mathbf{P}_0, \mathbf{X}_0) \quad (25)$$

where $(\mathbf{P}_0, \mathbf{X}_0)$ are the pre-merger equilibrium vectors of premium and primary care copays and g_{jk} is a stacked vector of g_{jk}^p and g_{jk}^x .

Because some products have premiums or primary care copays that are equal to 0, it is unlikely that the first-order constraints for those products are just-binding in the pre-merger equilibrium. In the event that the conditions are not binding, there are two cases. In the first case, the post-merger first order is non-negative at the pre-merger equilibrium, e.g. $h_j^p(\mathbf{P}, \mathbf{X}) \geq 0$ and $p_j = 0$. In this case, the additional pressure from the merger is not

large enough to justify any incentive to raise the premium and the local approximation of the effect of the merger on this particular product is 0.

In the second case, the pre-merger first order condition is not binding, but the incentive from the merger is still large enough to justify a premium increase, i.e. $h_j^p(\mathbf{P}, \mathbf{X}) < 0$. In this case, the effect of the merger depends on the degree of slack in the pre-merger first order condition. Measuring the slack in the first order condition requires imposing the model predicted marginal revenues and marginal costs and making assumptions about which deviations in the data are due to structural errors and which are due to the optimal response to constraints.

Instead of imposing these assumptions on equilibrium, I assume in the baseline predictions that all slack first order conditions in the pre-merger equilibrium remain slack in the post-merger equilibrium and the effect of the merger on those product characteristics is zero—a lower bound on the effects of a merger. In a forthcoming appendix section, I show results with the alternative assumption that all first order conditions are just-binding, which represents the upper bounds on the effects of the merger.

7.3 Results

I predict the effects of a merger between two of each of the three largest firms in the Medicare Advantage (MA) market in Massachusetts: Tufts Health Plan (Tufts), Blue Cross Blue Shield of Massachusetts (BCBS), and United Healthcare (United). The summary statistics for all six firms that operate in the state are displayed in Table 8.

Tufts is the largest firm in the state and insures nearly half of all MA beneficiaries. Tufts also attracts the highest average risk of all the firms and charges the lowest average primary care copay. BCBS is the next largest and covers 26% of the MA market. BCBS products have a similar average premium to Tufts, but charge a higher primary care copay and enroll a healthier base of consumers. United is the third largest firm, with 14% of the market. United also charges the lowest average premium of any firm in the market and enrolls the lowest risk pool of consumers.²¹

²¹Table 8 shows state wide averages, but there is substantial variation in county-level competition. For

Table 8: Firms are Differentiated in their Premiums, Copays, and Risk Distributions

	MA Share	Average	Average	Avg. Risk		Risk Adj. Cost	
		Premium	Copay	Data	Model	Data	Model
Tufts	0.47	112	11.6	1.24	1.28	556	579
BCBS	0.26	113	19.0	0.98	0.97	628	595
United	0.14	26.5	16.0	0.82	0.85	606	687
Fallon	0.06	92.6	20.3	1.20	1.13	629	530
Health New Engl.	0.05	130	20.1	1.19	1.11	579	554
Harvard Pilgrim	0.03	117	13.4	1.14	1.13	592	591

Note: This table shows summary statistics for the six firms that offer Medicare Advantage plans in Massachusetts. The market share, average premium, and average copay are matched precisely to the data. The average risk and cost comparisons show that the model can capture the risk heterogeneity among the firms.

Table 9 displays the mean pre-merger values of the premium and primary care copay, the average effect of the merger, and the average DAP for each characteristic. Each merger results in an increase in monthly premium and the primary care copay, on average. The largest average price increase is also associated with a low average primary care copay increase.

Table 9: Mergers Lead to Higher Average Premiums and Primary Care Copays

	Pre-merger Mean		Merger Effect		DAP	
	Premium	Copay	Premium	Copay	Premium	Copay
Tufts - BCBS	120	13.5	12.3	0.11	24.3	24.2
Tufts - United	93.9	11.6	4.07	1.08	6.51	4.42
BCBS - United	86.4	16.9	8.42	1.33	3.93	0.33

Note: In all three mergers analyzed, the mean premium and mean primary care copay increase as a result of the merger, with the largest effect on the premium occurring with the smallest effect on the primary care copay. This table shows the mean effects of the merger analysis of three hypothetical mergers among the three largest firms in the Massachusetts Medicare Advantage market. The mean pre-merger values and the mean diverted average profit are presented for context.

Figure 3a explores these dynamics at the product-level in the merger between the two largest firms, Tufts and BCBS. Each dot represents the merger effect on a single product in a single county, and the size of the dot is relative to the pre-merger enrollment of that product. The figures make it clear that the average conceals a great deal of heterogeneity in the effect

example, United is popular in western Massachusetts, but not in the Boston area, possibly due to the composition of its network.




Figure 3: High Premium Effects are Associated with Low Copay Effects

Note: Large premium increases as the result of a merger are typically compensated with primary care copay reductions, and the products with large opposite effects tend to be those with negative diverted average profitability values. This figures shows the product-level effect of the Tufts-BCBS merger. The size of each dot represents the relative size of pre-merger product enrollment. The left panel shows the relationship between the effects on premium and primary care copay. The right panel shows the relationship between the primary care copay effect and the diverted average profit with respect to the copay.

of the merger. Figure 3b shows that while the majority of merger effects are concentrated in modest price and copay increases, there are a number of products with sizeable decreases in the primary care copay and some with decreases in the premium. Moreover, Figure 3b demonstrates the substitutability of premium and the primary care copay. Large increases in the premium are typically paired with decreases in the primary care copay and vice versa.

From the averages presented in 9, the average DAP does not have a clear relationship with the average merger effect. Figure 3 shows that the relationship is more clear at the product-level. Nearly every product with a negative DAP value experiences a decrease in the primary care copay as a result of the merger. The interaction of the two strategic variables through the substitution effect weakens this correlation. There are still many negative effects of the merger even for products with positive values of DAP.

The effects of these changes on consumers are displayed in Table 10. To characterize the heterogeneity in the effects, I separate the consumers into two groups: those that experience a reduction and those that experience an increase in their share-weighted primary care copay. Intuitively, this is the effect of their merger on the expected primary care copay of an individual if their product choice probabilities are held fixed. In each merger, there are large groups of consumers in each category.

For example, in the merger between the largest two firms, Tufts and BCBS, 98 thousand individuals face a primary care copay increase as a result of the merger, and the merger leads their copays to increase by an average of \$1.08 dollars. These consumers respond by decreasing their medical consumption by \$31 per person per year. In this same merger, 54 thousand individuals face a primary care copay decrease with an average effect of -\$1.92. These consumers increase their medical consumption by an average of \$77.4 dollars per person per year. Averaging across both of these groups, the primary care copay effect is negligible and total spending increases slightly by \$7.33 dollars per person per year.

Table 10: Consumer-level Effects of a Merger are Heterogeneous

	Population Affected (000s)	Primary Care Copay Effect	Medical Cons. Effect (\$/year)	Mortality Effect (pp)	Savings per Life (\$000s)
Tufts - BCBS					
Mean Effect	152	0.02	7.33	0.000	-
Increase	98	1.08	-31.2	0.010	315
Reduction	54	-1.92	77.4	-0.018	441
BCBS - United					
Mean Effect	127	0.30	-10.6	0.002	388
Increase	55	1.61	-57.1	0.015	387
Reduction	72	-0.71	25.2	-0.006	386
Tufts - United					
Mean Effect	127	0.58	-17.0	0.005	320
Increase	39	3.00	-91.6	0.028	332
Reduction	88	-0.49	16.0	-0.005	353

Note: While the average effects of a merger are small, groups of consumers face concentrated increases or decreases in their primary care copay depending on their county of residence and the plans in which they are enrolled. For each group of consumers, the reduction in spending per increase in expected life lost is between \$320 and \$441 thousand. This table displays the consumer level effects of each merger, averaged across all consumers, those that experience an increase in their share-weighted reduction in the primary care copay, and those that experience an increase. The final column is the result of dividing the predicted change in medical consumption by the predicted change in twelve month inpatient mortality.

The results can quantify the resource cost or benefit of increasing or decreasing the medical consumption of Medicare Advantage beneficiaries via changes in the primary care copay. However, in order to discern if additional medical consumption at a higher total cost should be viewed as welfare improving or not, it is important to put these results in the context of their effect on the health of the consumers.

The descriptive evidence on the relationship between inpatient mortality and primary care copays presented in Section 6.4 suggest that a \$1 increase in the primary care copay is associated with a 0.01 percentage point increase in the 12-month inpatient mortality rate. This implies that the two mergers which result in a meaningful increase in the average primary care copay will lead to an average increase in 12-month inpatient mortality of 0.002 and 0.005 percentage points. Relative to the amount of additional spending or savings, this corresponds to between \$320 and \$388 thousand in medical spending per expected life. These figures are well below estimates for the value of a statistical life (VSL), which range between \$4 million and \$10 million. Even when taking into consideration a reduced life expectancy, VSL estimates for individuals in the Medicare eligibility age range exceed \$1 million (Aldy and Viscusi 2007). This is consistent with findings in the literature that patients cut back on all types of care in the face of higher out-of-pocket prices, rather than the most unnecessary or wasteful care (Chandra, Gruber, and McKnight 2010, Baicker, Mullainathan, and Schwartzstein 2015, Brot-Goldberg et al. 2017).

Taken together, these results suggest that mergers in the insurance can have a meaningful impact on medical consumption and health via the cost-sharing terms of insurance. The decline in medical spending that results from an increase in the level of cost-sharing is more than offset by the negative effect on the health of the consumers. Moreover, the magnitude of the effect of a merger on medical consumption and consumer health is not proportional to the effect on premiums, the current focus of competition policy. This shows that a framework that can assess the impact of a merger on the ultimate medical consumption of the insurance beneficiaries should be an important aspect of competition policy.

8 Conclusion

This paper follows from the observation that, by setting the cost-sharing terms of insurance, competition in the insurance industry has an effect on medical consumption and patient health. I estimate a model using detailed data that links insurance product choices to medical claims in order to incorporate adverse selection, moral hazard, and the effect of cost-sharing terms on patient health. I find that this channel is indeed important. Competition between insurance firms reduces the level of cost-sharing, on average. Consumers respond to lower levels of cost-sharing by increasing their medical consumption. And lower levels of cost-sharing lead to lower rates of inpatient mortality.

I combine these estimates with the observed costs of insurance in the claims data to characterize the effect on insurance competition on the cost-sharing terms. In a counterfactual exercise, I vary the degree of competition in the Massachusetts Medicare Advantage market through potential bilateral mergers between the three largest firms and focus on the primary care copay as the endogenous cost-sharing term. Each merger leads to both higher premiums and higher primary care copays, on average. However, the effects on particular products are heterogeneous, with the largest premium increase occurring alongside reductions in the primary care copay and vice versa.

These changes in the primary care copay have implications for both medical consumption and patient health. In the merger with the largest effect on the primary care copay, average medical spending declines by \$17 per person per year and the likelihood of an inpatient death in a twelve month period increases by 0.005 percentage points. This corresponds to an expected 6.3 additional deaths as a result of the decrease in competition, with a savings reduction of about \$320 thousand per death. At estimates of the value of a statistical life, the cost of the additional deaths far outweigh the savings from the reduction in medical spending.

This framework explores the ways in which competition in health insurance affects not only the monthly premium of insurance but also the cost-sharing. The cost-sharing terms are of course not the only important non-price feature of insurance. Other ways in which

insurance firms compete includes the design of the hospital and physician network (Capps, Dranove, and Satterthwaite 2003, Shepard 2016, Ho and Lee 2017), the design of drug formularies, the use of “gate-keepers”, and the use of non-financial ways to allocate medical care such as prior authorization requirements. Each of these may also be a mechanism through which insurance competition affects the amount and type of medical care received by insurance beneficiaries. The extension of this model to incorporate these other mechanisms is an important agenda for future research.

Works Cited

- Abaluck, Jason and Jonathan Gruber. “Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program”. *American Economic Review* 101.4, ISSN: 08966273. DOI: 10.1016/j.jacc.2007.01.076. White (2011): 1180–1210. Print.
- Abaluck, Jason, Peter Hull, and Amanda Starc. “Mortality Effects and Choice Across Private Health Insurance Plans”. *National Bureau of Economic Research*. ISSN: 0898-2937. DOI: 10.3386/w27578 (2020). Web. <<https://www.nber.org/papers/w27578>>.
- Abraham, Jean, Coleman Drake, Jeffrey S Mccullough, · Kosali Simon, B Jean, and Marie Abraham. “What drives insurer participation and premiums in the Federally-Facilitated Marketplace?” *International Journal of Health Economics and Management* 17. DOI: 10.1007/s10754-017-9215-y (2017): 395–412. Web. <<https://link.springer.com/content/pdf/10.1007/s10754-017-9215-y.pdf>>.
- Afendulis, Christopher C., Anna D. Sinaiko, and Richard G. Frank. “Dominated choices and Medicare Advantage enrollment”. *Journal of Economic Behavior and Organization* 119. ISSN: 01672681. DOI: 10.1016/j.jebo.2015.07.009 (2015): 72–83. Print.
- Aizawa, Naoki and You Suk Kim. “Advertising and Risk Selection in Health Insurance Markets”. *American Economic Review* 108.3. ISSN: 00028282. DOI: 10.1257/aer.20151485 (2018): 828–867. Print.
- Aldy, Joseph E and W Kip Viscusi. “Age Differences in the Value of Statistical Life: Revealed Preference Evidence”. *Review of Environmental Economics and Policy* 1.2. DOI: 10.1093/reep/rem014 (2007): 241–260. Web. <<http://reep.oxfordjournals.org/>>.
- Aron-Dine, Aviva, Liran Einav, Amy Finkelstein, and Mark Cullen. “Moral Hazard in Health Insurance: Do Dynamic Incentives Matter?” *The Review of Economics and Statistics* 97.4. DOI: 10.1162/REST_a_00518 (2015): 725–741. Web. <<http://www.mitpress>>.
- Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein. “BEHAVIORAL HAZARD IN HEALTH INSURANCE”. *Quarterly Journal of Economics* 130.4. DOI: 10.1093/qje/qjv029 (2015): 1623–1667. Web. <<https://academic.oup.com/qje/article-abstract/130/4/1623/1916415>>.

- Berry, Steven, James Levinsohn, and Ariel Pakes. “Automobile Prices in Market Equilibrium”. *Econometrica* 63.4. ISSN: 00129682. DOI: 10.2307/2171802 (1995): 841. Web. <<http://www.jstor.org/stable/2171802?origin=crossref>>.
- Bhargava, Saurabh, George Loewenstein, and Justin Sydnor. “CHOOSE TO LOSE: HEALTH PLAN CHOICES FROM A MENU WITH DOMINATED OPTIONS”. *The Quarterly Journal of Economics* 132.3. DOI: 10.1093/qje/qjx011 (2017): 1319–1372. Web. <<https://academic.oup.com/qje/article-abstract/132/3/1319/3769420>>.
- Bloom, Nicholas, Carol Propper, Stephan Seiler, and John Van Reenen. “The Impact of Competition on Management Quality: Evidence from Public Hospitals”. *Review of Economic Studies* 82. DOI: 10.1093/restud/rdu045 (2015): 457–489. Web. <<https://academic.oup.com/restud/article-abstract/82/2/457/1585860>>.
- Bresnahan, Timothy F and Peter C Reiss. “Entry and Competition in Concentrated Markets”. *Journal of Political Economy* 99.5 (1991): 977–1009. Web. <<https://ideas.repec.org/a/ucp/jpolec/v99y1991i5p977-1009.html>>.
- Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad. “What does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics”. *The Quarterly Journal of Economics* 132.3 (2017): 1261–1318. Print.
- Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston. “How does risk selection respond to risk adjustment? New evidence from the Medicare Advantage Program”. *American Economic Review* 104.10, ISSN: 00028282. DOI: 10.1257/aer.104.10.3335 (2014): 3335–3364. Web. <<http://dx.doi.org/10.1257/aer.104.10.3335>>.
- Buntin, Melinda Beeuwkes and Alan M. Zaslavsky. “Too much ado about two-part models and transformation? Comparing methods of modeling Medicare expenditures”. *Journal of Health Economics* 23 (2004): 525–542. Print.
- Cabral, Marika, Michael Geruso, and Neale Mahoney. “Do Larger Health Insurance Subsidies Benefit Patients or Producers? Evidence from Medicare Advantage”. *American Economic Review* 108.8. DOI: 10.1257/aer.20151362 (2018): 2048–2087. Web. <<https://doi.org/10.1257/aer.20151362>>.

- Cao, Zhun and Thomas G. McGuire. "Service-level selection by HMOs in Medicare". *Journal of Health Economics* 22.6. ISSN: 01676296. DOI: 10.1016/j.jhealeco.2003.06.005 (2003): 915–931. Print.
- Capps, Cory, David Dranove, and Mark Satterthwaite. "Competition and market power in option demand markets". *RAND Journal of Economics* 34.4 (2003): 737–763. Web. <<https://pdfs.semanticscholar.org/868f/28b7b1f33c170f2f5e3fbef6d3a4b1ecd7c.pdf>>.
- Cardon, James H. and Igal Hendel. "Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey". *The RAND Journal of Economics* 32.3. ISSN: 07416261. DOI: 10.2307/2696362 (2001): 408. Web. <<http://doi.wiley.com/10.2307/2696362>>.
- Chandra, Amitabh, Jonathan Gruber, and Robin McKnight. "Patient cost-sharing and hospitalization offsets in the elderly". *American Economic Review* 100.1. ISSN: 00028282. DOI: 10.1257/aer.100.1.193 (2010): 193–213. Print.
- Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya. "Can Health Insurance Competition Work? Evidence from Medicare Advantage". *Journal of Political Economy* NBER Working Paper forthcoming (2021). Web. <<http://www.nber.org/data-appendix/w20818>>.
- Cutler, D. M. and S. J. Reber. "Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection". *The Quarterly Journal of Economics* 113.2. ISSN: 0033-5533. DOI: 10.1162/003355398555649 (1998): 433–466. Web. <<https://academic.oup.com/qje/article-lookup/doi/10.1162/003355398555649>>.
- Dafny, Leemore, Mark Duggan, and Subramaniam Ramanarayanan. "Paying a Premium on Your Premium? Consolidation in the US Health Insurance Industry". *American Economic Review* 102.2. DOI: 10.1257/aer.102.2.1161 (2012): 1161–1185. Web. <<http://dx.doi.org/10.1257/aer.102.2.1161>>.
- Dafny, Leemore, Kate Ho, and Robin S Lee. "The Price Effects of Cross-Market Mergers: Theory and Evidence from the Hospital Industry". *RAND Journal of Economics* (2018). Web. <<https://www.cms.gov/research-statistics-data-and->>.

- Dalton, Christina M., Gautam Gowrisankaran, and Robert J. Town. "Salience, myopia, and complex dynamic incentives: Evidence from Medicare Part D". *Review of Economic Studies* 87.2. ISSN: 1467937X. DOI: 10.1093/restud/rdz023 (2020): 822–869. Print.
- Decarolis, Francesco and Andrea Guglielmo. "Insurers' response to selection risk: Evidence from Medicare enrollment reforms". *Journal of Health Economics* 56. ISSN: 01676296. DOI: 10.1016/j.jhealeco.2017.02.007 (2017): 383–396. Web. <<http://linkinghub.elsevier.com/retrieve/pii/S0167629617302060>>.
- DeLeire, Thomas, Andre Chappel, Kenneth Finegold, and Emily Gee. "Do individuals respond to cost-sharing subsidies in their selections of marketplace health insurance plans?" *Journal of Health Economics* 56. ISSN: 0167-6296. DOI: 10.1016/J.JHEALECO.2017.09.008 (2017): 71–86. Web. <<https://www.sciencedirect.com/science/article/pii/S0167629617301005?via=ihub>>.
- Dickstein, Michael J., Mark Duggan, Joe Orsini, and Pietro Tebaldi. "The impact of market size and composition on health insurance premiums: Evidence from the first year of the affordable care act". *American Economic Review*. American Economic Association, 2015. 120–125. Web. <<http://dx.doi.org/10.1257/aer.p20151083>>.
- Drake, Coleman. "What Are Consumers Willing to Pay for a Broad Network Health Plan?: Evidence from Covered California". *Journal of Health Economics* 65 (2019): 63–77. Print.
- Drake, Coleman, Conor Ryan, and Bryan Dowd. "Sources of Inertia in Health Plan Choice in the Individual Health Insurance Market". *Working Paper* (2020). Print.
- Dranove, David, Anne Gron, and Michael J. Mazzeo. "Differentiation and Competition in HMO Markets". *Journal of Industrial Economics* 51.4. ISSN: 0022-1821. DOI: 10.1111/j.0022-1821.2003.00209.x (2003): 433–454. Web. <<http://doi.wiley.com/10.1111/j.0022-1821.2003.00209.x>>.
- Duarte, Fabian. "Price elasticity of expenditure across health care services". *Journal of Health Economics* 31. DOI: 10.1016/j.jhealeco.2012.07.002 (2012): 824–841. Web. <<http://dx.doi.org/10.1016/j.jhealeco.2012.07.002>>.
- Dubin, Jeffrey A. and Daniel L. McFadden. "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption". *Econometrica* 52.2. ISSN: 00129682. DOI: 10.2307/1911493 (1984): 345. Print.

- Einav, Liran, Amy Finkelstein, Stephen P Ryan, Paul Schrimpf, and Mark R Cullen. “Selection on Moral Hazard in Health Insurance”. *American Economic Review* 103.1. ISSN: 0002-8282. DOI: 10.1257/aer.103.1.178 (2013): 178–219. Web. <<http://pubs.aeaweb.org/doi/10.1257/aer.103.1.178>>.
- Ellis, Randall P., Bruno Martins, and Wenjia Zhu. “Health care demand elasticities by type of service”. *Journal of Health Economics* 55. ISSN: 18791646. DOI: 10.1016/j.jhealeco.2017.07.007 (2017): 232–243. Web. <<https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5600717>>.
- Ellis, Randall P. and Thomas G. McGuire. “Predictability and predictiveness in health care spending”. *Journal of Health Economics* 26.1, ISSN: 01676296. DOI: 10.1016/j.jhealeco.2006.06.004 (2007): 25–48. Web. <<https://www.sciencedirect.com/science/article/pii/S0167629606000725>>.
- Ericson, Keith M Marzilli and Amanda Starc. “Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange”. *Review of Economics and Statistics* 97.3 (2015): 667–682. Web. <<http://www.nber.org/papers/w18089>>.
- Fan, Ying. “Ownership consolidation and product characteristics: A study of the US daily newspaper market”. *American Economic Review* 103.5. ISSN: 00028282. DOI: 10.1257/aer.103.5.1598 (2013): 1598–1628. Print.
- Farrell, Joseph and Carl Shapiro. “Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition”. *The B.E. Journal of Theoretical Economics* 10 (2010). Web. <<http://www.bepress.com/bejte/vol10/iss1/art9>>.
- Frank, Richard G., Jacob Glazer, and Thomas G. McGuire. “Measuring Adverse Selection in Managed Health Care”. *Journal of Health Economics* 19 (2000): 829–854. Print.
- Geruso, Michael. “Demand heterogeneity in insurance markets: Implications for equity and efficiency”. *Quantitative Economics* 8. DOI: 10.3982/QE794 (2017): 929–975. Web. <<http://www.qeconomics.org>>.
- Geruso, Michael, Timothy Layton, and Daniel Prinz. “Screening in contract design: Evidence from the ACA health insurance exchanges”. *American Economic Journal: Microeconomics* 11.2. ISSN: 19457685. DOI: 10.1257/po1.20170014 (2019): 64–107. Print.

- Glazer, Jacob and Thomas G. McGuire. “Optimal risk adjustment in markets with adverse selection: An application to managed care”. *American Economic Review* 90.4, ISSN: 00028282. DOI: 10.1257/aer.90.4.1055 (2000): 1055–1071. Web. <<http://pubs.aeaweb.org/doi/10.1257/aer.90.4.1055>>.
- Goolsbee, Austan and Amil Petrin. “The Consumer Gains from Direct Broadcast Satellites and the Competition with Cable TV”. *Econometrica* 72.2. ISSN: 0012-9682. DOI: 10.1111/j.1468-0262.2004.00494.x (2004): 351–381. Web. <<http://doi.wiley.com/10.1111/j.1468-0262.2004.00494.x>>.
- Gowrisankaran, Gautam, Aviv Nevo, and Robert Town. “Mergers when prices are negotiated: Evidence from the hospital industry”. *American Economic Review* 105.1. ISSN: 00028282. DOI: 10.1257/aer.20130223 (2015): 172–203. Print.
- Handel, Benjamin R and Jonathan T Kolstad. “Health Insurance for Humans: Information Frictions, Plan Choice, and Consumer Welfare”. *American Economic Review* 105.8. DOI: 10.1257/aer.20131126 (2015): 2449–2500. Web. <<http://dx.doi.org/10.1257/aer.20131126>>.
- Handel, Benjamin R., Jonathan T. Kolstad, and Johannes Spinnewijn. “Information frictions and adverse selection: Policy interventions in health insurance markets”. *Review of Economics and Statistics* 101.2. ISSN: 15309142. DOI: 10.1162/rest_a_00773 (2019): 326–340. Web. <https://www.mitpressjournals.org/doi/abs/10.1162/rest_a_00773>.
- Heiss, Florian, Daniel Mcfadden, Joachim Winter, Amelie Wuppermann, and Bo Zhou. “Inattention and Switching Costs as Sources of Inertia in Medicare Part D”. *NBER Working Paper* 22765. DOI: 10.3386/w22765 (2016). Print.
- Ho, Kate, Joseph Hogan, and Fiona Scott Morton. “The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program”. *RAND Journal of Economics* 48.4, ISSN: 17562171. DOI: 10.1111/1756-2171.12207 (2017): 877–905. Print.
- Ho, Kate and Robin S Lee. “INSURER COMPETITION IN HEALTH CARE MARKETS”. *Econometrica* 85.2. DOI: 10.3982/ECTA13570 (2017): 379–417. Web. <http://www.sutterhealth.org/about/healthcare_costs.html>.

- Hörner, Johannes. “Reputation and competition”. *American Economic Review* 92.3. ISSN: 00028282. DOI: 10.1257/00028280260136444 (2002): 644–663. Print.
- Jaffe, Sonia and Mark Shepard. “Price-Linked Subsidies and Health Insurance Markups”. *National Bureau of Economic Research Working Paper Series* No. 23104. DOI: 10.3386/w23104 (2017). Web. <<http://www.nber.org/papers/w23104><http://www.nber.org/papers/w23104.pdf>>.
- Jaffe, Sonia and E. Glen Weyl. “The First-Order approach to merger analysis”. *American Economic Journal: Microeconomics* 5.4. ISSN: 19457669. DOI: 10.1257/mic.5.4.188 (2013): 188–218. Web. <<http://dx.doi.org/10.1257/mic.5.4.188>>.
- Kling, Jeffrey R., Sendhil Mullainathan, Eldar Shafir, Lee C. Vermeulen, and Marian V. Wrobel. “Comparison friction: Experimental evidence from medicare drug plans”. *Quarterly Journal of Economics* 127.1. ISSN: 00335533. DOI: 10.1093/qje/qjr055 (2012): 199–235. Web. <www.planprescriber.com>.
- Kowalski, Amanda. “Censored Quantile Instrumental Variable Estimates of the Price Elasticity of Expenditure on Medical Care”. *Journal of Business & Economic Statistics* 34.1. ISSN: 1537-2707. DOI: 10.1080/07350015.2015.1004072 (2016): 107–117. Web. <<https://www.tandfonline.com/action/journalInformation?journalCode=ubes20><https://tandfonline.com/r/JBES>>.
- Lavetti, Kurt, Thomas DeLeire, and Nicolas Ziebarth. “How Do Low-Income Enrollees in the Affordable Care Act Marketplaces Respond to Cost-Sharing?” *National Bureau of Economic Research*. ISSN: 0898-2937. DOI: 10.3386/w26430 (2019). Print.
- Lester, Benjamin, Ali Shourideh, Venky Venkateswaran, and Ariel Zetlin-Jones. “Screening and Adverse Selection in Frictional Markets”. <https://doi.org/10.1086/700730> 127.1. DOI: 10.1086/700730 (2019): 338–377. Web. <<https://www.journals.uchicago.edu/doi/abs/10.1086/700730>>.
- Mahoney, Neale and E. Glen Weyl. “Imperfect Competition in Selection Markets”. *The Review of Economics and Statistics* 99.4. ISSN: 0034-6535. DOI: 10.1162/REST_A_00661 (2017): 637–651. Web. <http://direct.mit.edu/rest/article-pdf/99/4/637/1918443/rest{_}a{_}00661.pdf>.

- Manning, Willard G, Joseph P Newhouse, Naihua Duan, Emmett B Keeler, Arleen Leibowitz, and M Susan Marquis. “American Economic Association Health Insurance and the Demand for Medical Care: Evidence from a Randomized Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment”. *Source: The American Economic Review* 77.3 (1987): 251–277. Web. <<http://www.jstor.org/stable/1804094>>.
- Marone, Victoria R and Adrienne Sabety. “Should There Be Vertical Choice in Health Insurance Markets? ” (2020). Print.
- Matsa, David A. “Competition and product quality in the supermarket industry”. *Quarterly Journal of Economics* 126.3. ISSN: 00335533. DOI: 10.1093/qje/qjr031 (2011): 1539–1591. Web. <<https://academic.oup.com/qje/article/126/3/1539/1854944>>.
- McGuire, Thomas G., Joseph P. Newhouse, and Anna D. Sinaiko. “An Economic History of Medicare Part C”. *Milbank Quarterly* 89.2. ISSN: 0887378X. DOI: 10.1111/j.1468-0009.2011.00629.x (2011): 289–332. Web. <<http://www.ncbi.nlm.nih.gov/pubmed/21676024><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3117270><http://doi.wiley.com/10.1111/j.1468-0009.2011.00629.x>>.
- Miller, Keaton, Amil Petrin, Robert Town, and Michael Chernew. “Optimal Managed Competition Subsidies”. Web. <https://editorialexpress.com/cgi-bin/conference/download.cgi?db{_}name=II0C2018{\&}paper{_}id=424>. NBER Working Paper Series.
- Newhouse, Joseph P., Mary Price, J. Michael McWilliams, John Hsu, and Thomas G. McGuire. “How much favorable selection is left in medicare advantage?” *American Journal of Health Economics* 1.1. ISSN: 23323507. DOI: 10.1162/ajhe_a_00001 (2015). Web. <https://www.journals.uchicago.edu/doi/abs/10.1162/AJHE{_}a{_}00001>.
- Pelech, Daria. “Paying more for less? Insurer competition and health plan generosity in the Medicare Advantage program”. *Journal of Health Economics* 61. ISSN: 18791646. DOI: 10.1016/j.jhealeco.2018.07.002 (2018): 77–92. Print.
- Petrin, Amil and Boyoung Seo. “Identification and Estimation of Discrete Choice Demand Models when Observed and Unobserved Characteristics are Correlated”. *Working Paper* (2019). Web. <<https://docs.google.com/a/umn.edu/viewer?a=v{\&}pid=sites{\&}srcid=dW1uLmVkdXxhbWlsLXBldHJpbXxneDoyNzUwNDIzZThlNW5N2U3>>.

- Ryan, Conor. “Adverse Selection and Market Structure”. *Working Paper* (2020): 1–41. Print.
- Saltzman, Evan. “Managing adverse selection: underinsurance versus underenrollment”. *The RAND Journal of Economics* 52.2. ISSN: 1756-2171. DOI: 10.1111/1756-2171.12372 (2021): 359–381. Web. <<https://onlinelibrary.wiley.com/doi/full/10.1111/1756-2171.12372>><<https://onlinelibrary.wiley.com/doi/abs/10.1111/1756-2171.12372>><<https://onlinelibrary.wiley.com/doi/10.1111/1756-2171.12372>>.
- Schmalensee, Richard. “MARKET STRUCTURE, DURABILITY, AND QUALITY: A SELECTIVE SURVEY”. *Economic Inquiry* 17.2. ISSN: 00952583. DOI: 10.1111/j.1465-7295.1979.tb00307.x (1979): 177–196. Web. <<http://doi.wiley.com/10.1111/j.1465-7295.1979.tb00307.x>>.
- Shepard, Mark. “Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange”. *NBER Working Paper* (2016). Print.
- Spence, A. Michael. “Monopoly, Quality, and Regulation”. *The Bell Journal of Economics* 6.2. ISSN: 0361915X. DOI: 10.2307/3003237 (1975): 417. Print.
- Tebaldi, Pietro. “Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA”. *Becker Friedman Institute for Research in Economics Working Paper*. Chicago. Web. <http://www.pietrotebaldi.com/uploads/8/1/3/6/81364040/eehix{_}aug14{_}17.pdf>. Becker Friedman Institute for Research in Economics Working Paper.
- Town, Robert. “The welfare impact of HMO mergers”. *Journal of Health Economics* 20.6. ISSN: 01676296. DOI: 10.1016/S0167-6296(01)00104-7 (2001): 967–990. Print.
- Town, Robert and Su Liu. “The Welfare Impact of Medicare HMOs”. *The RAND Journal of Economics* 34.4. ISSN: 07416261. DOI: 10.2307/1593785 (2003): 719. Print.
- Veiga, André and E. Glen Weyl. “Product design in selection markets”. *Quarterly Journal of Economics* 131.2, ISSN: 15314650. DOI: 10.1093/qje/qjw007 (2016): 1007–1056. Web. <<https://academic.oup.com/qje/article-lookup/doi/10.1093/qje/qjw007>>.




Figure A1: Histogram of Premiums (National, 2011-2019)

Table A1: Effect of Competition on Insurance Plan Characteristics - Full Results

	First Stage Firms	Prem.	Pt B Reb.	Deduct.	OOP	Primary	Copays Spcl.	Otpt.
Log Market Size	0.86*** (0.01)							
# of Firms		-3.37*** (0.12)	0.25*** (0.01)	-3.24*** (0.81)	-160.06*** (13.18)	-1.50*** (0.04)	-1.00*** (0.05)	4.60*** (0.60)
Income (\$000)	-1.56*** (0.11)	-3.82*** (1.33)	0.90*** (0.16)	-4.51 (8.72)	-415.38*** (141.37)	-1.10*** (0.38)	-2.56*** (0.58)	-2.14 (6.41)
% White	0.00 (0.00)	0.18*** (0.02)	-0.01*** (0.00)	0.37** (0.15)	19.79*** (2.47)	0.00 (0.01)	0.02** (0.01)	-0.36*** (0.11)
<i>Among Eligible</i>								
% Seniors over 85	1.33*** (0.33)	9.50** (3.90)	1.07** (0.46)	41.11 (25.53)	-106.67 (413.83)	15.16*** (1.11)	6.11*** (1.71)	-82.76*** (18.76)
% Seniors Employed	1.27*** (0.29)	1.37 (3.43)	-1.52*** (0.40)	-46.75** (22.49)	-842.45** (364.62)	8.83*** (0.98)	5.84*** (1.51)	8.65 (16.53)
% w/ Cog. Disability	-0.85*** (0.30)	11.56*** (3.61)	-1.45*** (0.42)	-87.46*** (23.65)	-125.67 (383.40)	5.13*** (1.03)	1.17 (1.59)	-27.64 (17.38)
<i>Resources (per 1000)</i>								
PC Docs (per 1000)	-0.38*** (0.03)	2.26*** (0.36)	-0.06 (0.04)	-3.40 (2.36)	-52.87 (38.30)	-0.01 (0.10)	-0.36** (0.16)	2.10 (1.74)
Hosp. Beds (per 1000)	-0.00 (0.00)	0.08*** (0.03)	0.00 (0.00)	0.15 (0.18)	-22.57*** (2.98)	0.02*** (0.01)	0.05*** (0.01)	-0.36*** (0.14)
<i>Fixed Effects</i>								
State & Year	✓	✓	✓	✓	✓	✓	✓	✓
<u>Effect</u>								
Data Mean		-0.18 (0.20)	0.20 (0.68)	-0.22 (0.79)	-0.03 (0.01)	-0.16 (0.10)	-0.03 (0.01)	0.04 (0.04)

Note: This table shows the results of the IV analysis of the effect of competition on insurance cost-sharing characteristics detailed in Section 4.2. The unit of observation is a US county in a given year between 2011 and 2019. The dependent variable is the enrollment weighted average of a product characteristic: prem - monthly premium; pt B reb - reduction in the part B premium; deduct - deductible; oop - out-of-pocket spending limit; primary - primary care copay; spcl - specialist copay; otpt - outpatient services.

Table A2: Effect of Competition on Insurance Plan Characteristics - Full Results (Continued)

	Copays					Coinsurance Rates			
	Radio.	Lab	Emerg.	Inpt.	Amb.	Outpt.	Radio.	Device	Drug
# of Firms	-2.08*** (0.44)	0.05 (0.03)	-0.05** (0.02)	-9.96*** (0.52)	0.98*** (0.38)	-1.10*** (0.05)	-0.43*** (0.05)	-0.09*** (0.02)	0.10*** (0.03)
Income (\$000)	-2.55 (4.72)	1.19*** (0.36)	1.14*** (0.25)	-4.15 (5.55)	8.35** (4.09)	0.58 (0.50)	1.00** (0.49)	0.36* (0.20)	1.06*** (0.37)
% White	0.48*** (0.08)	-0.00 (0.01)	0.01*** (0.00)	-0.25** (0.10)	0.19*** (0.07)	-0.02** (0.01)	-0.04*** (0.01)	-0.00 (0.00)	-0.02*** (0.01)
<i>Among Eligible</i>									
% Seniors over 85	-9.60 (13.81)	3.43*** (1.06)	-2.14*** (0.74)	66.86*** (16.24)	-10.43 (11.96)	10.81*** (1.48)	5.49*** (1.44)	-1.08* (0.58)	3.11*** (1.08)
% Seniors Employed	15.81 (12.16)	6.64*** (0.94)	-1.87*** (0.65)	67.41*** (14.31)	15.83 (10.54)	5.53*** (1.30)	5.72*** (1.27)	0.64 (0.51)	-0.45 (0.95)
% w/ Cog. Disability	-61.40*** (12.79)	4.56*** (0.98)	0.11 (0.69)	30.27** (15.04)	6.53 (11.08)	8.24*** (1.37)	12.84*** (1.33)	-0.72 (0.54)	-0.16 (1.00)
<i>Resources (per 1000)</i>									
PC Docs (per 1000)	-2.37* (1.28)	0.08 (0.10)	-0.21*** (0.07)	1.99 (1.50)	-1.98* (1.11)	-0.33** (0.14)	-0.19 (0.13)	0.22*** (0.05)	-0.04 (0.10)
Hosp. Beds (per 1000)	-0.14 (0.10)	0.01 (0.01)	0.01* (0.01)	-0.15 (0.12)	-0.08 (0.09)	0.05*** (0.01)	0.10*** (0.01)	0.00 (0.00)	0.02*** (0.01)
<i>Fixed Effects</i>									
State & Year	✓	✓	✓	✓	✓	✓	✓	✓	✓
<hr/>									
Effect	-0.04	0.01	0.00	-0.04	0.01	-0.19	-0.06	-0.01	0.01
Data Mean	(0.03)	(0.02)	(0.00)	(0.02)	(0.00)	(0.18)	(0.05)	(0.00)	(0.00)

Note: This table shows the results of the IV analysis of the effect of competition on insurance cost-sharing characteristics detailed in Section 4.2. The unit of observation is a US county in a given year between 2011 and 2019. The dependent variable is the enrollment weighted average of a product characteristic: radio - diagnostic radiology; lab - lab tests; emerg - emergency room visits; inpt - inpatient stays; amb - ambulance; outpt - outpatient services; device - medical devices; drug - outpatient drugs.

Table A3: Estimates of Demand Heterogeneity - Continued

	OOP Limit (\$1000)	Copays (\$10)		Coinsurance Rates (pp)		
		Emergency	Ambulance	Outpatient	Med Device	Drug
Over 75	0.145 (0.010)	0.003 (0.009)	-0.001 (0.003)	-0.001 (0.005)	-0.009** (0.004)	-0.001 (0.002)
Female	0.118 (0.009)	-0.004 (0.008)	0.001 (0.002)	-0.010** (0.005)	0.005 (0.004)	-0.002 (0.002)
Heart Arrythmia	0.754*** (0.015)	0.045*** (0.013)	-0.023*** (0.004)	0.012 (0.008)	0.001 (0.006)	-0.010*** (0.004)
Vascular Disease	-0.202 (0.016)	-0.027** (0.013)	0.025*** (0.004)	0.000 (0.008)	0.014 (0.006)	-0.001 (0.004)
Diabetes w/ Compl.	-0.543*** (0.016)	-0.026* (0.013)	0.036*** (0.004)	-0.004 (0.008)	0.010 (0.006)	0.013 (0.004)
Diabetes w/o Compl.	-0.661*** (0.015)	-0.012 (0.013)	0.007* (0.004)	0.014* (0.008)	0.004 (0.007)	-0.001 (0.004)
Breast/Prost. Cancer	0.605*** (0.018)	-0.016 (0.016)	-0.001 (0.005)	0.003 (0.010)	0.008 (0.008)	-0.020*** (0.004)
Rheum. Arthritis	-0.054 (0.023)	-0.056*** (0.018)	0.051*** (0.006)	-0.018 (0.012)	-0.004 (0.009)	0.004 (0.005)
Agg. Risk Score	-0.436*** (0.012)	-0.010 (0.010)	-0.011*** (0.003)	-0.028*** (0.007)	0.004 (0.005)	0.005 * (0.003)
Agg. Risk Score ²	0.056*** (0.002)	0.000 (0.001)	0.001*** (0.000)	0.006*** (0.001)	-0.002*** (0.001)	0.000 (0.000)

A Data Processing

A.1 Linking Medical Claims to Products

The task of linking publicly available data on insurance products to the patients in the MA APCD requires two tasks. The first is to correctly identify the APCD product identifier in which each patient is enrolled in each month. The member file of the APCD lists the products in which each patient is enrolled and the start and end months for their enrollment, but these records are in general not unique. The membership file is first subset to include only medical insurance for patients in Massachusetts, and only insurance products which are indicated to be the primary source of coverage.

The membership records are de-duplicated for each patient in the following way. First, only records with the highest membership eligibility ID for a particular product and activity month are kept. Next, only records with the most recent activity date for a particular product and start month are kept. Then, for each month between 2013 and 2017, I collect all remaining records with a start date prior to that month and an end date that is either missing or later than that month. The remaining records are prioritized first by coverage type and then by activity month. Highest priority is given to fully insured plans and the most recent record activity. Any remaining duplicate records are randomly assigned. This ambiguity affects the product ID in 0.1% of member-months and the firm ID in less than 0.01% of member months.

The next task is to link APCD product identifier to publicly available information. The MA APCD makes publicly available the identity of some insurance firms in the data, including all of the firms offering plans in Medicare Advantage. However, the APCD product IDs are not linked to the public names of the products. The data are matched using aggregate information on the market shares of each plan in each county. In the APCD, MA products are identified in the product file using the line of business and insurance plan market fields. Members in the ACPD are linked to counties through their 5-digit zip code. Where the zip code does not fully identify the county, the observation is given a weight in all counties that intersect that zip code proportional to the distribution of population in the zip code. In

Massachusetts, this affects a small number of observations. From this data, I can compute the MA market share of each APCD product ID in each county and month.

This data set can be compared to the county-month level market shares computed to the enrollment data made publicly available by CMS. Market shares from this data are computed among the medical MA plans that are not Senior Care Options plans, which are identified separately in the APCD. Then for each possible pair of a CMS plan ID and APCD plan ID, I compute the percent of variation in the vector of county-month market shares in the CMS data that is present in the APCD data, similar to the R^2 of a regression. A pair is considered to be a match if they are close (explained variation exceeds 90%) and have no close match to any other plans in their respective data sets. This match is performed separately for every calendar year, as some APCD product IDs change from year to year. Some plans have ambiguous matches and are manually assigned based on the identity of the firm and the share of enrollees that are enrolled in an identified plan the following year.

Through this methodology, I am able to identify the insurance plan for 93% of all medicare advantage beneficiaries and 97% of those enrolled in one of the three largest firms. I drop all plans that have fewer than 11 individuals from both the APCD and CMS data.

A.2 Sample Selection for Insurance Demand Estimation

The demand for insurance relies on an annual panel of insurance enrollment decisions made by Medicare beneficiaries. I exclude from this sample all enrollees in employer-sponsored MA plans or Special Needs Plans (SNP), and all persons under the age of 65 who may be eligible because of a disability.

Most consumers are enrolled in either a single plan for the entire calendar year or they switch into a new plan during the open enrollment period that takes place from January to March at the beginning of each year. For consumers which have two plans during the year, I treat the plan with the longest enrollment as the plan choice for that particular year. This affects only 0.09% of member-years and abstracts from idiosyncratic special enrollment windows that some consumers may experience during the year.

I treat individuals over the age of 65 that are not enrolled in any MA plan as eligible to enroll but selecting traditional Medicare. I normalize the total relative size of the MA and TM population using the MA county-level penetration rate documented in the Area Health Resource File.

In order to balance the important sources of identification and the computational burden of the large data set, I over sample among individuals that ever select a MA plan and individuals that become eligible for MA during the sample period. I draw a random sample of 30% of consumers that ever select an MA plan, and a 60% sample of consumers that become eligible for MA during the sample period. For the remaining population that always select TM, I draw a 1.5% sample. The estimation procedure uses the corresponding probability weights.

A.3 Sample Selection for Medical Consumption Estimation

The estimation of the elasticity of medical consumption with respect to cost-sharing terms relies on a monthly panel of medical consumption for Medicare Advantage beneficiaries. Conditional on being over the age of 65, this data exclude two populations. First, it excludes any member-months of traditional Medicare enrollment. The medical consumption of traditional medicare members is only observed for traditional Medicare enrollees that are also enrolled in a Medigap plan. However, I am unable to link the precise Medigap plan, and therefore do not have full information on the cost-sharing terms of the members' insurance. As a result, I exclude all of these observations from estimation.

Second, there is a problem in the link between the insurance enrollment panel and the medical claims data for members of United Healthcare. As a result, I am unable to directly link the medical consumption of a patient identified in the claims data to the enrollment of a particular member in the insurance enrollment panel. Because this breaks the primary source of identification in the estimation, I also exclude all member-months of United Healthcare enrollment from the estimation data, which account for roughly 14% of all member months.

Additionally, I drop any member-months where there is disagreement in the product

in which a consumer is enrolled between the membership and medical claims data (3% of member months). I drop any member-months after a month in which it's been indicated that a patient died in an inpatient facility. If the patient has non-zero spending, I allow for up to two additional months after the indicated month.

A.4 Measuring Medical Consumption

The baseline measure of medical consumption is the total medical spending—both out-of-pocket and covered expenses—of a patient during a particular month. This measure is convenient because it incorporates a notion of intensity (some medical services are higher value or represent more in-depth care) and it has a direct relationship to the costs of the insurance firms. However, the measure may be contaminated by differences in the negotiated prices paid by each insurance product for a particular medical service in each year.

Ideally, a measure of medical consumption would result in equal quantities if two individuals receive the same care but are enrolled in different insurance products at different times. I construct such a measure to serve as a robustness check for the medical consumption elasticity estimates presented in Section 6.

Consider a patient i , enrolled in product j , that receives a procedure p in year t . The total spending on that procedure is given by

$$m_{ipjt} = q_{ip} + \iota_p K_{jt} \tag{26}$$

where q_{ipjt} represents the medical intensity of the service, K_{jt} is a vector of indicator variables for each product-year and ι_p is the procedure-product-year conversion factor that adjusts the medical intensity of the service into a price paid by the insurance firm to the provider, minus the out-of-pocket expenses of the patient.

The medical intensity itself is modeled as

$$q_{ip} = \Gamma_{p0} + \Gamma'_p L_{ip} + \epsilon_{ip} \tag{27}$$

where g is a procedure specific function that maps a vector of characteristics, L_{ip} , into a quantity of medical intensity. The vector of characteristics contains indications of the hospital revenue code, the principal diagnosis code, the first procedure modifier, the site of service, and the provider specialty that apply to the procedure, each of which is coded as a binary variable on the values that appear in the data for a given procedure. For example, for many procedures, the site of service variables contain indication of whether the procedure was performed in a physician's office, a hospital outpatient center, or an ambulatory surgery center (among other possibilities).

The goal is to estimate $\hat{\Gamma}_p$ and use the predicted value of q_{ip} as an alternative measure of quantity. To estimate the large number of parameters, I use the least absolute shrinkage and selection operator (LASSO) on the data for in-network procedures among all MA patients that receive each procedure. Because this method focuses on procedures themselves (i.e. physician services), I ignore all spending related to medical facilities. The estimator solves

$$\min_{\Gamma_{p0}, \Gamma_p, \iota_p} \frac{1}{N} \sum_{i=1}^N \frac{1}{2} (y_i - \Gamma_{p0} - \Gamma'_p L_{ip} - \iota_p K_{jt}) + \lambda_p \left(\frac{1}{2} \left\| [\Gamma_p; \iota_p] \right\|_2^2 + \left\| [\Gamma_p; \iota_p] \right\|_1 \right) \quad (28)$$

I estimate this model for every procedure in the data where the total number of claims for that particular procedure is at least 25. The vector K_{jt} excludes a large plan-year which has coverage of most procedures in the data in order to provide a consistent interpretation for Γ_{p0} across procedures. The parameter λ_p determines the degree of regularization in the regression and is selected for each procedure to minimize the mean squared error of prediction on a sample withheld for cross-validation.

With the model estimated for each procedure, I predict \hat{q}_{ip} . The adjusted measure of medical consumption is equal to the sum of all predicted medical consumption quantities for all procedures that an individual receives during a given month.

A.5 Measuring Consumer Health Status

Consumer health status is summarized in two ways. The first is through a set of binary variables that indicate whether the consumer is diagnosed with a particular disease, and the second is a summary risk score. Both of these variables are constructed using the risk score methodology that CMS uses to administer the risk adjusted subsidies associated with the MA program. The methodology can be reproduced using SAS code made publicly available by CMS.

For each plan year, I compute the risk score of each consumer enrolled in a Medicare Advantage or Medigap plan. The methodology also assigned hierarchical condition categories (HCCs) as building blocks of the risk score. I use the most prevalent of these HCCs as clinical disease categories. In each case, these measures of health status are concurrent to the plan year. For example, if an individual is indicated as having diabetes, it implies that individual had some procedure during the current year in which diabetes was listed as a relevant diagnosis.

The health status for two populations must be imputed. First, the medical claims of members of United Healthcare cannot be linked properly to the enrollment panel. However, the distribution of health status is known, conditional on the plan year, sex, and insurance product. Therefore, I assign each consumer a random draw from this distribution. I first assign a draw from the empirical distribution of HCC indications. I then assign a random risk score drawn from a parametric log-normal distribution conditional on the plan year, sex, insurance product, and the HCC indications. I truncate the parametric distribution at the observed conditional maximum and minimum risk scores in the data in order to avoid unreasonable outliers.

Second, the medical claims of traditional Medicare beneficiaries that do not enroll in a Medigap plan do not appear in the APCD, and as a result, these health measures can not be constructed. To impute the health status of these enrollees, I follow the same methodology as the previous case and assume that the enrollees in traditional Medicare without Medigap come from the same distribution of health status as traditional Medicare enrollees with

Medigap.

A.6 Expected Medical Spending

The firm's problem requires the model to make predictions about the expected medical consumption of consumers that may or may not have enrolled in a particular product in the data. This presents two problems: computing expected medical spending and computing counterfactual medical spending in other products.

Because the estimation of medical consumption is log-linear, computing the expectations is not straightforward. If there is heteroskedasticity in the error, the typical formula for the mean of a log-normal distribution no longer applies. More worrying are the non-linearities in the out-of-pocket expenses. Rather than compute expectations with more costly computational methods, I take advantage of the feature of the firm model that treats all consumers that are identical in observable characteristics as identical. As a result, I compute the expectation as the empirical mean across all consumers within an observable type.

In Section 6.1, the medical consumption of an individual can be decomposed into a component common to all members in a particular product, $\beta'X_{j\tau} + \gamma'F_j + \lambda_\tau$, and an idiosyncratic component, $\eta_i + \omega_{i\tau}$. I compute the counterfactual medical consumption all the plans in a consumer's choice set by applying the component common to all consumers in each product. The differences in spending across products is not assumed to be causal, and therefore these differences are not included in the estimates of spending changes in the policy analysis.

This computation is further complicated for the two populations that were dropped from the medical consumption estimation: members of United Healthcare and members of traditional Medicare. Because I do not observe the medical spending of these consumers, I assign each a draw from the empirical distribution of the idiosyncratic component of medical consumption, $\eta_i + \omega_{i\tau}$, conditional on the sex, year, the consumer's risk score, the consumer's HCC diagnoses. This maintains the identified relationship between the preferences for insurance (which depend on the measurements of health status) and the expected medical consumption of the consumer. For United Healthcare plans, I also calibrate a plan-level

fixed effect in order to match the total predicted spending of the members in the data with the actual total spending on those members.

A.7 Measuring Additional Sources of Marginal Cost

In addition to the cost of medical claims, insurance firms also incur administrative costs and costs associated to prescription drug claims. I assume that these costs fixed, per-beneficiary expenses and are identical across all products offered by a firm in a particular year.

The data on both administrative and prescription drug expenses come from the Medical Loss Ratio filings (MLR). In years 2015 through 2017, the MLR data separately provide information on each firm's Medicare business in a particular state. Prior to 2015, I use the category designated as "government program plans."

Administrative expenses consist of the sum of expenses related to quality (health outcome) improvement, preventing hospital re-admissions, improving patient safety and reducing medical errors, wellness and health promotion, health IT improvement, cost containment, direct sales salaries and benefits, agent and broker fees, taxes and assessments, fines and penalties, claim adjustment expenses, and other general administrative costs. These make up sections 4 and 5 of part 1 of the MLR filing, with the exception of costs related to the implementation of the ICD-10 standard.

Prescription drug expenses are computed as the total spending on prescription drugs less pharmaceutical rebates. The assumption that prescription drug expenses are constant across products and consumers is quite strong. However, the per-consumer cost of prescription drug coverage net of the subsidies associate with Medicare Advantage Part D is small relative to the medical claims cost of insurance.

B The Effective Coinsurance Rate

I estimate the effect of cost-sharing terms on a plan-level average coinsurance rate for two reasons. First, it allows me to translate elasticity estimates on primary care copays to a coinsurance elasticity that can be more easily compared to estimates in the literature.

Second, it is required to predict the expected change in out-of-pocket expenses charged to each consumer given a change in the primary care copay but holding fixed their medical consumption. This is a component of computing the firms’ expected marginal cost.

The coinsurance rate is modeled as linear in cost-sharing parameters and also depends on a second-order, product-specific polynomial in individual medical spending. This captures the fixed nature of many of the out-of-pocket expenses. The average coinsurance rate is decreasing in total medical spending up to the out-of-pocket spending limit.

The effective coinsurance rate, computed over the year t , is specified as

$$\phi_{ijt} = \beta' X_{jt} + \lambda_t + \gamma_{j1} M_{it} + \gamma_{j2} M_{it}^2 + \omega_{ijt}^{coins} \quad (29)$$

where M_{it} is the total annual spending of consumer i in year t . I restrict the sample to individuals that have non-zero medical spending during the year but do not reach the out-of-pocket spending limit. The results are displayed in Table A5.

C Estimating the Bid Function

The per-person subsidy, risk-adjusted subsidy is given by

$$b_{ijt} = rs_i \left(\min\{Bench_j, bid_j(p_j, X_j)\} + \Lambda_j \max\{Bench_j - bid_j(p_j, X_j), 0\} \right) \quad (30)$$

where rs_i is the individual’s summary risk score, bid_j is the bid submitted by the insurance plan, $Bench_j$ is a plan specific benchmark subsidy level that depends on the counties where the plan is offered, and Λ_j is a “rebate” share that depends on the plan’s quality rating.

The bid function is estimated from a national panel on Medicare Advantage plan characteristics and payment information. While the plan bids are not directly observable, the data do contain the rebate payment, mean risk score, and mean payment level. If the plan-specific benchmark level was directly observable, the bid itself could be inferred from equation (30).

I follow Curto et al. 2021 in using an approximated plan-specific benchmark from the enrollment weighted average of county-level benchmarks. This provides an approximated bid that can be used to estimate the function, bid_{jt} .

The plan bid function is specified as linear in the monthly premium, the primary care copay, and a vector of product characteristics which include other cost-sharing parameters and the plan specific benchmark.

$$bid_{jt} = \alpha p_{jt} + \beta x_{jt} + \Gamma' X_{jt} + \gamma_j + \lambda_t + \zeta_{jt} \quad (31)$$

The parameters, α and β are identified through a two-way fixed effects model. The identifying assumption is that all plans experience parallel trends. In this context, it requires that there is no idiosyncratic and transient shock, observable to the firm, that affects both the bid and the premium or primary care copay.

The results of the bid estimation are presented in Table A6. The monthly premium and primary care copay are each replaced by bid dollars in a slightly less than 1-to-1 ratio, which is consistent with the intended rules of Medicare Advantage and firms that have market power. For firms that bid below the benchmark (nearly all firms), an increase in the bid must be accompanied by an increase in the plan premium, an increase in the plans cost-sharing, or both. If the plan held cost-sharing fixed, the Medicare Advantage rules intend to require a premium increase by between 0.5 and 0.7, depending on the plan's rebate share. Since the plans have market power and the government cannot perfectly observe the plans' costs, these rules may not be followed exactly. A finding that the plan premium would increase by 0.822 on average is consistent with this model of firm behavior. A similar rationale applies to changes in the copay, though it is less clear what the intended substitution between bid dollars and the primary care copay should be.

D First Order Approximation of Merger Effects

In this section, I restate the proof of Theorem 1 in Jaffe and Weyl 2013, with minor extensions to accommodate an environment with both copays and premium.

Let $Q = [P; O]$ be the stacked vector of premiums and copays selected for all products in a market. Let $f(Q)$ be the vector of pre-merger first order conditions and $g(Q)$ be the vector of upward pricing pressure, such that $f + g = h$, the post-merger first order conditions. More detail on these functions is presented in Section ??.

Assumption D.1. *The vector of post-merger first order conditions, h , is locally invertible in a neighborhood \mathcal{B} around Q_0 , the pre-merger equilibrium, such that there is a vector $Q^M \in \mathcal{B}$ with $h(Q^M) = 0$.*

This assumption requires that there is a locally unique equilibrium in the neighborhood of the pre-merger equilibrium and demonstrates one key strength of this approach. If there is a locally unique equilibrium, this first order approximation will point in that direction. This is a conceptual strength, as it likely corresponds to how the firm's themselves internalize their change in incentive. And it is a computational strength, as more general solution methods may have trouble locating the neighborhood of uniqueness. This assumption would fail if there is no post-merger equilibrium in a neighborhood sufficiently small enough that h is invertible throughout. However, if this is the case, it is likely infeasible to evaluate a post-merger counterfactual with any method.

Theorem D.1. *Given assumption D.1, then a first-order approximation of the change in Q induced by the merger is*

$$\Delta Q = - \left(\frac{\partial h}{\partial Q}(Q_0) \right)^{-1} \cdot g(Q_0)$$

Proof. Since $f(Q_0) = 0$, $h(Q_0) = g(Q_0) = r$. The goal is to locate Q^M such that $h(Q^M) = 0$. If h is invertible in a neighborhood that encompasses both the pre- and post-merger

equilibrium, then

$$\begin{aligned}\Delta Q = Q_m - Q_0 &= h^{-1}(0) - h^{-1}(r) = \left(\frac{\partial h^{-1}}{\partial h}(r) \right) (0 - r) + \mathcal{O}(\|r\|^2) \\ &\simeq - \left(\frac{\partial h}{\partial Q}(Q_0) \right)^{-1} \cdot g(Q_0)\end{aligned}$$

□

Table A4: Coefficients of Base-level Indirect Utility

	Two-way	IV
Monthly Premium	-0.119*** (0.018)	-0.110 (0.063)
Primary Care	-0.413*** (0.075)	-0.505* (0.319)
Out-of-Pocket Limit	-0.092* (0.049)	-0.000 (0.000)
Specialist	0.082 (0.129)	-0.000 (0.507)
Outpatient	0.012 (0.017)	-0.004 (0.049)
Outpatient Coins	0.085** (0.033)	0.076 (0.104)
Inpatient Stay	0.018*** (0.003)	0.014** (0.007)
Emergency Room	0.003 (0.031)	-0.016 (0.097)
Ambulance	-0.042*** (0.008)	-0.039* (0.0204)
Medical Devices Coins	-0.030** (0.013)	-0.020 (0.030)
Outpatient Drugs Coins	0.037*** (0.010)	0.039 (0.030)
Diagnostic Imaging	-0.002 (0.001)	-0.015 (0.029)
Fixed Effects		
Year & Product	✓	✓
Offers Part D	✓	✓
Star Rating	✓	✓
Switching Cost Estimates		
Product-Level	5.54*** (0.01)	
MA to TM	2.12*** (0.02)	
TM to MA	1.60*** (0.01)	

Note: The results of the two-way fixed effect and the IV specifications are quantitatively similar. The two-way specification will be used as the baseline specification. All variables are copays denominated in \$10 with the exception of variables labeled with coinsurance (percentage points) and OOP limit (thousands of dollars). The switching cost estimates come from the maximum likelihood estimation. The significance stars ***, **, and * denote statistical significance at the 1%, 5%, and 10% level respectively.

Table A5: Cost-sharing Terms and the Effective Coinsurance Rate

	Effective Coinsurance Rate
Primary Care	0.033*** (0.000)
Specialist	0.010*** (0.000)
Outpatient	0.000** (0.000)
Outpatient Coins	0.0005*** (0.000)
Inpatient Stay	-0.000*** (0.000)
Emergency Room	-0.007*** (0.000)
Ambulance	0.000*** (0.000)
Medical Devices	0.029*** (0.000)
Outpatient Drugs	0.001*** (0.000)
Year	✓
Product-specific Spending Polynomial	✓
Observations	897,030

Note: The tables average estimated effective coinsurance rate as predicted by the cost-sharing terms of the insurance plan. The unit of observation is a person-year. The estimation controls for year fixed effects and a product-specific polynomial in the annual spending of each consumer. Coefficients marked with *** are statistically significant at the 0.1% level.

Table A6

	(1)	(2)	(3)
Benchmark	0.611*** (0.009)	0.561*** (0.010)	0.893*** (0.012)
Premium	1.080*** (0.014)	1.089*** (0.014)	0.822*** (0.027)
Primary Care Copay	2.420*** (0.083)	2.232*** (0.083)	0.484*** (0.095)
Specialist Copay	0.624*** (0.053)	0.629*** (0.052)	0.602*** (0.065)
Outpatient Copay	0.022*** (0.006)	0.022*** (0.006)	0.017*** (0.005)
Outpatient Coinsurance	1.251*** (0.075)	1.151*** (0.074)	0.669*** (0.072)
Inpatient Copay	0.104*** (0.006)	0.105*** (0.006)	0.022*** (0.005)
Emergency Copay	-0.347*** (0.069)	-0.226** (0.074)	0.350*** (0.069)
Ambulance Copay	0.072*** (0.009)	0.097*** (0.009)	0.099*** (0.010)
Med Device Coins	2.421*** (0.168)	2.292*** (0.166)	0.002 (0.183)
Outpatient Drug Coins	0.080	0.199	0.011
Fixed Effects			
Year		✓	✓
Product			✓

Note: *p<0.05; **p<0.01; ***p<0.001