

How does Insurance Competition Affect Medical Consumption?

Conor Ryan, Pennsylvania State University*

February 14, 2023

Abstract

Competition in insurance markets affects not only the premium but also the cost-sharing terms—e.g. copays and coinsurance rates— which may affect a patient’s medical decisions and health outcomes. Using medical claims data linked to insurance product choices, I estimate a model in which consumers select an insurance plan and make medical consumption decisions given the cost-sharing terms of their insurance. Firms compete on both the premium and the copay for primary care. A \$10 increase in the copay leads to an 8% decrease in medical consumption and a 0.2 percentage point increase in inpatient mortality. Mergers have heterogeneous effects on the primary care copay, leading to between a \$6 reduction and \$24 increase in mean annual medical consumption. At typical estimates of the value of a statistical life, mergers that increase medical consumption improve welfare as the additional resource use is outweighed by a reduction in mortality risk.

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 each of which differently affect consumer behavior. 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

*Email: conor.ryan@psu.edu; Phone: (814) 865-1457

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.

In this paper, I provide a framework to evaluate the effect of changes in market structure on the health and health care use of insurance beneficiaries through the cost-sharing terms of insurance. I estimate a model of insurance competition, medical consumption, and health in Medicare Advantage (MA). There are three main findings. First, a reduction in competition via a merger leads to heterogeneous effects on both the premiums and cost-sharing of merging products, driven by adverse selection between products of the merging parties and the firms' trade-off between premiums and cost-sharing. Second, consumers respond to cost-sharing terms in their demand for medical care, consistent with other findings in the literature (?, ?). Moreover, consumers in the lowest-income zip codes are the most elastic in their medical consumption demand. Finally, 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 to a 0.2 percentage point increase in inpatient mortality. Mortality effects are also concentrated in low-income zip-codes and higher-risk consumers. In the merger with the largest *decrease* in consumer cost-sharing, the average primary care copay decrease by \$0.20, leading to an average increase in annual medical spending of \$24 per person. By combining the estimates of cost from the merger analysis with the effect on inpatient mortality, I find that the average increase in spending per reduction in expected deaths is roughly \$500 thousand dollars. This is well below standard estimates of the value of a statistical life, suggesting that mergers that increase total spending on health care may still be welfare improving to due the benefits of medical consumption.

The model of consumer demand consists of two stages. In the first stage, consumers make a discrete choice over the available MA 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. I show that the effect of competition on cost-sharing terms is ambiguous for two reasons. The first reason is due to selection: not all consumers generate positive expected profit. In a standard model, aggressive pricing steals profitable consumers away from competitors. This intuition can flip in the presence of adverse selection: some firms may benefit their competitors by attracting unprofitable consumers (? , ? , ? , ?). The second reason follows from incentives facing a firm that competes in both price and a non-price quality that consumers value. The balance between cost-sharing (or any product quality) and premium that firms will provide depends on the preferences of the marginal consumers, and changes in competition can alter this balance by directly affecting firm incentive and altering which consumers are marginal (? , ?). The direction and magnitude of the effect of a particular merger can be theoretically characterized through the interaction two incentives but is ultimately an empirical question.

In order to quantitatively evaluate these mechanisms, it is crucial to characterize consumer preferences for insurance, elasticities of medical consumption with respect to cost-sharing, 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 MA market in Massachusetts. Using the medical claims, I can construct detailed information on health status—e.g. specific 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 that incorporates information on medical diagnoses and their interaction with the cost-sharing terms of insurance. The willingness to pay to reduce the primary care copay by \$10 is relatively large and increasing in the total expected medical expenditure. Consumers in the 5th and the 95th percentile of expected cost are willing to pay \$54

and \$96 per month, respectively. 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 ?.

In the second stage, I estimate consumers' elasticity of medical consumption with respect to cost-sharing terms using within product variation over time in the cost-sharing terms of insurance. Since consumers face considerable inertia in plan choices (?, ?, ?), this with-in product variation is a good instrument for with-in consumer variation in cost-sharing terms (?). 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 between a 12% and 39% decline in medical consumption as measured by total medical spending for consumers in the lowest-income zipcodes. In higher-income zipcodes, the semi-elasticity of total medical spending ranges between a 19% decline and a 14% increase. Additionally, I find that, within the lowest-income zipcodes, the least healthy individuals (measured by a medical risk score) are the *most* elastic with respect to primary care copays. This gradient with respect to medical risk diminishes among higher-income consumers. These estimates imply an arc elasticity of between -0.06 and -0.19 among low-income zip codes. These findings reflect the elasticity of *total* spending with respect to only primary care cost-sharing and are smaller magnitudes than service-specific demand elasticities found elsewhere in the literature (?,?).

Next, I investigate the relationship between primary care copays and inpatient mortality. I estimate the effect of the primary care copay on inpatient mortality (patient deaths in hospitals or hospice care facilities) using similar identification intuition as in the medical consumption demand model. I find that a \$10 increase in the primary care copay leads to a 0.2 percentage point increase in inpatient mortality, with these effects concentrated among higher risk consumers and lower-income consumers. The magnitude of the effect is in line with other estimates on the causal

differences in health outcomes among insurance plans in MA (?, ?).

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 MA market in Massachusetts and focus on monthly premium and the primary care copay as the key endogenous features. Among the merging parties, each merger leads to increases in the average premium and the average copay for primary care. But when taking into account consumer switching patterns, some mergers lead to decreases in the average copay for primary care, decreases in the average premium, or decreases in both across consumers in the market. Across each merger and each local market, the mean increase in the average primary care copay is \$0.03. These findings are consistent with empirical evidence using a national, market-level panel of competition and cost-sharing terms (including the primary care copay) which shows that, on average, less competition leads to greater consumer-cost sharing (?). However, the model highlights substantial heterogeneity across different products, markets, and mergers.

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 of \$24.1 per person per year to a decrease of \$5.6 per person per year. These effects are concentrated among consumers in low-income zip codes that are the most elastic in their medical consumption. Among these consumers, the effect on medical consumption is between an increase of \$46.5 per person per year to a decrease of \$10.5 per person per year, magnitudes that are much larger than premium effects in two of the three studied mergers.

The changes in medical consumption lead to an important tradeoff. A reduction in cost-sharing terms as a result of a merger will increase the total spending on medical consumption but also decrease expected mortality. By combining the estimates of cost from the merger analysis with the effect on inpatient mortality, I find that the average increase in spending per reduction in expected deaths is roughly \$500 thousand dollars. This is well below standard estimates of the value of a statistical 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 (?). This suggests that the welfare benefit of mergers that decrease total spending on health care through greater cost-sharing is outweighed by the welfare cost of an increase in mortality. Alternatively, mergers that lead to a reduction in cost-sharing and greater resource cost of health care may still be welfare improving.

Relation to the Literature

This paper makes two main contributions. First, I estimate a model of imperfect competition between insurance firms that includes both adverse selection and moral hazard in consumer behavior. I tractably incorporate a second-stage, medical consumption decision into an imperfectly competitive, differentiated products framework. 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 (?, ?, ?, ?, ?, ?, ?).

The model of consumer insurance choice can depend flexibly on consumer medical conditions and the cost-sharing characteristics of the insurance plan but does not necessarily assume any rational expectations of health expenditure by the consumer. The two-stage aspect of the consumer's problem builds 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 (?, ?, ?). Drawing on insights that consumers make mistakes when selecting health insurance (?, ?, ?, ?, ?), I specify a version of this two-stage model that does not require rational expectations and can be more tractably estimated in a setting with complex cost-sharing terms.

This first stage of the estimations extends ? by adding data on medical diagnoses and service-specific cost-sharing terms to standard discrete choice insurance demand estimation methods (?, ?,?, ?, ?). In the second stage, I estimate the reduced form price-elasticity of medical care in multi-product, non-group insurance markets. Beginning with the RAND Health insurance experiment, a randomized experiment on insurance benefits (?), the literature has studied medical consumption in larger contexts using natural experiments (?, ?), contract non-linearity (?,?),

variation in the choice set (?), and instrumental variables (?). 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 supply side of the model includes imperfectly competitive firms that can choose the cost-sharing terms of insurance, in addition to the premium, in an environment with both adverse selection and moral hazard. This builds on a growing literature that explores the mechanisms through which firms seek a more favorable risk pool (?). There is a related literature on managed care in MA that documents mechanisms and incentives to screen for profitable consumers through the generous (or sparing) provision of certain types of service (?).

The second contribution is estimating the effect of a change in competition on medical consumption and inpatient mortality through cost-sharing terms. This contributes to a literature that studies market structure in health insurance (?). Recent research has begun to move beyond the focus on the insurance premium to study competition in the context of contracting with provider networks (?), the MA bidding rules (?), and the ways that insurance product design feeds back into the market structure of the provider industry (?). This paper builds on this work to study the effect market structure has 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 (?).

2 Setting: Medicare Advantage

The Medicare Advantage (MA) 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 health insurance plan available to U.S. residents over the age of 65 or disabled. MA is

a program through which insurance firms compete to offer insurance plans to the same beneficiaries that cover at least the same services as TM. By allowing firms to compete, the government hoped it could provide greater benefits to consumers at a lower cost (?).

MA prioritized making the market attractive for insurance firms in order to generate competition by offering large subsidies adjusted for risk. The program has been successful in generating substantial participation by both insurance firms and Medicare beneficiaries (?). Despite these successes, the degree of competition still varies substantially widely. 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.¹ Still more mergers have been consummated that have not risen to such high levels of antitrust concern.²

Due to the large subsidies and associated rules, competition between firms is often concentrated on the cost-sharing parameters rather than the monthly premium. Many insurance firms offer products with no monthly premium. And while it is possible to set negative premiums via a rebate, it is rare. Instead, firms offer more generous benefits in order to attract consumers.

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

¹These 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.

²For instance, Aetna-Coventry in 2013 and United-PacifiCare in 2005.

premium and cost-sharing terms. Using the model, I characterize why the effect of a potential merger on premiums and cost-sharing terms is ambiguous.

3.1 The Environment

Consumers

Consumers, indexed by i , face a two stage decision following ? and ?. 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.

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

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

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 Z_i is a vector of consumer characteristics which may provide 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, $m^*(\omega, X_j, Z_i)$, which I assume does not depend on the premium or non-financial plan characteristics.³

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

$$v_{ijt} = V(\varepsilon_{ij}, \mathcal{E}[U_{ij}^*(\omega; p_j, X_j, W_j, Z_i)|Z_i]) \quad (2)$$

where ε_{ij} is an unobserved idiosyncratic preference of consumer i for product j and

³This requires that the income effect of the premium is small, which is reasonable given the low level of premiums in MA.

\mathcal{E} is the consumers subjective expectation of their second stage utility given their characteristics, Z_i .

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

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

In this paper, I avoid explicitly specifying the beliefs that form \mathcal{E} . Instead, I approximate $\mathcal{E}[U_{ij}^*(\omega)|Z_i]$ with a polynomial in p_j , X_j , W_j , and Z_i . Similar to approaches in the literature, I assume that $E[\varepsilon_{ij}\omega_i|Z_i] = 0$: the idiosyncratic insurance preference (ε) is uncorrelated with the medical demand shock (ω) conditional on patient characteristics (?). This assumption is not required for consistent estimates of consumer preference parameters, but it is necessary for solving the firms' profit maximization problem below.⁴

There are two reasons I make these alternative assumptions. First, this approach allows me to separately and tractably estimate s_{ijt} and m^* . Consumer cost-sharing in this market is primarily determined by service-specific copays, which means out-of-pocket expenses cannot be characterized by piece-wise linear functions of total medical spending as in other applications of this model (?,?). Furthermore, it requires beliefs about demand for particular types of service (e.g. primary care, inpatient stays). These factors present significant but surmountable modeling and computational obstacles to explicitly specifying consumer beliefs about out-of-pocket spending risk.

Second, and more importantly, there is evidence consumers themselves are not very good at predicting their medical expenditures (?, ?, ?). In Section 5.3, I discuss that the findings of this paper, interpreted through the lens of this two-stage model, suggest substantial information frictions or biases. This approach avoids having to make explicit functional form assumptions over the character of consumers' behavioral biases. The primary cost of the chosen approach is inability to include medical

⁴For the insurance demand identification assumptions, see Section 5.2. For the medical consumption demand identification assumptions, see Section 6.2

consumption in a welfare calculation, but the presence of information frictions already suggests that we should be wary using consumer choices to make statements about welfare.

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. 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} .

$$\begin{aligned} \Pi_{jt} &= \int_i s_{ijt}(p_{jt}, X_{jt}, p_{-jt}, X_{-jt}) (p_{jt} + b_{ijt}(p_{jt}, X_{jt}) - mc_{ijt}(X_{jt})) di & (4) \\ p_{jt} &\geq 0 ; x_{jtk} \in X_{jt} \geq 0 \end{aligned}$$

where p_{-jt} and 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. In the case of MA, firms are allowed to send premium rebates to consumers via their social security checks. However, this is rare generally and non-existent in the Massachusetts market, despite a significant portion of plans with a premium equal to zero. In the model, I 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

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

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 a_{jt} is a per-member administrative and drug cost. 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, $\phi_{jt}M$. Details of specifying and estimating the effective coinsurance rate are presented in Appendix Section E, and details on measuring administrative and drug costs can be found in Appendix Section C.6.

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 and depends on the characteristics of the plan. In Appendix Section F, 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 MA 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 ,

$$(p_{jt}^*, X_{jt}^*) = \arg \max_{(p, X)} \sum_{k \in J_{f(j)}} \Pi_{kt}(p, X, p_{-jt}, X_{-jt}) \quad (6)$$

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

3.2 Effect of a Merger

Merger effects are ambiguous for two reasons: adverse selection may flip the firms’ standard incentives to raise prices (premiums and copays, in this case) due to re-captured profit, and the pass-through of these incentives from a merger depends on the willingness to pay for low copays among a firm’s marginal consumers. In this section, I will outline the intuition for each of these sources of ambiguity.

The first source of ambiguity is due to selection: not all consumers generate positive expected profit. Consider a merger between two single product firms, j and

k. Equation 7 shows the post-merger first order condition 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.

$$\underbrace{\frac{1}{\frac{\partial s_j}{\partial x_j}} \int_i s_{ij} \left(\frac{\partial b_{ij}}{\partial x_j} - \frac{\partial mc_{ij}}{\partial x_j} \right) di + p_j}_{\text{Pre-Merger Marginal Revenue}} = \underbrace{-\frac{1}{\frac{\partial s_j}{\partial x_j}} \int_i \frac{\partial s_{ij}}{\partial x_j} (mc_{ij} - b_{ij}) di}_{\text{Pre-Merger Marginal Cost}} - \text{GePP}_{jk}^{\text{copay}} \quad (7)$$

$$\text{GePP}_{jk}^{\text{copay}} \equiv \frac{\int_i \frac{\partial s_{ik}}{\partial x_j} (p_k + b_{ik} - mc_{ik}) di}{\frac{\partial s_j}{\partial x_j}} \quad (8)$$

The expression, $\text{GePP}_{jk}^{\text{copay}}$, refers to the generalized pricing pressure (GePP) of a merger with product k on the copay of product j (?). This term represents the recaptured diverted sales that create unilateral incentives to alter premiums and copays after a merger and reduces to the standard upward pricing pressure (UPP) in price competition without selection (?).

In the presence of adverse selection, this term need not be positive. An increase in the copay (or premium) of product j may *lower* the profit of firm k because the consumers expected to switch to product k generate more cost than revenue. In this case, $\text{GePP}_{jk}^{\text{copay}}$ is negative and a merger between j and k exerts downward pressure on the copay of product j .

The intuition behind a merger that reduces copays is that high copays were rationalized pre-merger in part because they divert costly consumers to other products. However, the merged firm internalizes this negative externality on its newly acquired competitor and no longer has this incentive to keep copays high.

The second source of ambiguity is not directly related to competition. Rather, it follows from the tradeoff facing a firm that sets both a price (premium) and a non-price quality (cost-sharing) that consumers value. The level of cost-sharing that firms will provide depends on the willingness to pay for low cost-sharing among the marginal consumers and the cost of providing it. A merger can affect this trade-off in either direction, both by directly altering the incentives facing the firm and the identities of the marginal consumers (?, ?, ?, ?).

One perspective on this trade-off facing the firm is its pass-through of costs to premiums and copays. Consider a firm that faces some marginal cost increase, such as a per-person tax. The firm will pass that cost through to consumers taking into account the trade-off between higher premiums and higher copays among marginal consumers and potential changes in this trade-off as the margin shifts to new consumers.

The same intuition can be applied to a how a firm will optimally pass through a change in its incentives due to a merger. In an oversimplification, the price effects of the merger are the product of the vector of incentives—the GePP generated by the merger for both prices and copays—and some pass-through matrix, \mathbb{P} . In Appendix Section B, I show that the merger approximation methods of ? can be extended to characterize the appropriate pass-through matrix in this setting and show that this first order approximation closely matches the magnitude and direction of the full simulated merger effects presented in Section 7.

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.

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 sub-

mitted as a part of each medical claim, I construct indications for whether each consumer is diagnosed with a set of medical conditions as well as a summary risk score that aggregates all diagnoses for a consumer. To construct these measures, I use the Center for Medicaid and Medicare Services Hierarchical Conditions Categories (CMS-HCC) algorithm and risk coefficients. Second, the APCD provide a direct measure of medical consumption. The baseline measure is total medical spending, which is common in the literature (?, ?, ?, ?).

A novel aspect of this paper is linking medical consumption to the insurance choices of consumers. While the names of the MA firms are identified in the claims data, the names of the products are 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 MA 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 more detail on sample selection, linking the APCD and PBP data, measuring health status, and other data details, see Appendix Section C.

4.2 Descriptive Results

The first set of results show that 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.⁵ Appendix Table A1 shows descriptive data for all product characteristics by number of firms in the market.

I follow ? in using the potential market size 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

⁵The facts described in this section are also present in descriptive statistics of the estimation sample from the 14 counties of Massachusetts between 2013 and 2017.

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 (?, ?) and that local and national insurance plans are differentiated (?).

The estimation equation is given by (9). The dependent variable, y_{mt}^s is the enrollment weighted cost-sharing characteristic s in county m and year t . The parameter of interest is the number of firms, N_{mt} , and the excluded instrument is the log of market size, measured by population over 65. I include county-level 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 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.

$$y_{mt}^s = \beta \hat{N}_{mt} + \gamma' X_{mt} + \varepsilon_{mt} \quad (9)$$

The estimation shows that competition has significant and negative effects on cost-sharing parameters.⁶ Table 1 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.

The next set of results show that primary care is both commonly used and a large portion of out-of-pocket spending. Table 2 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

⁶These findings are consistent with ?, which finds that a reduction in competition via a large-scale exit of one plan type in MA led to higher expected out-of-pocket spending by the beneficiaries.

	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. 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.

Table 1: Evidence that Competition Reduces Cost-sharing Levels

is specialist visits, which are only used by roughly half of beneficiaries.⁷

	% 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: Office visits 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, mean out-of-pocket spending on each category by all consumers, and the mean out-of-pocket spending conditional on using the service.

Table 2: Primary Care is a Large Component of Out-of-pocket Spending

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.

Importantly, MA 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.

⁷These clinical categories depend on the procedure code billed by the physician, not the physician's specialty.

Following the results of this section, I will focus on primary care copays as the key strategic aspect of cost-sharing. In the following sections, I detail how these data identify the key parameters that allow me to characterize the important heterogeneity in firm and consumer behavior.

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 -2.29. This is lower than the mean semi-elasticity with respect to primary care, -14.4.

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

$$\delta_{jt} = \alpha_0 p_{jt} + \beta_0' X_{jt} + \gamma_0' W_{jt} + \xi_{jt} \quad (10)$$

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

$$v_{ijt} = \delta_{jt} + \Upsilon' D_{ijt} + \left(\sum_g \alpha_g z_{ig} \right) p_{jt} + \left(\sum_g \beta_g z_{ig} \right)' X_{jt} + \varepsilon_{ijt} \quad (11)$$

where ε_{ijt} 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 .

$$s_{ijt} = \Pr \left(v_{ijt} = \max_k v_{ikt} \right) \quad (12)$$

5.2 Estimation and Identification

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 δ_{jt} , 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 ?. In the first stage, the parameters governing consumer heterogeneity are estimated with maximum likelihood. Following ?, 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 estimate for $\hat{\theta}_z$ and the corresponding values of $\hat{\delta}_{jt}$ maximize the likelihood of the data. In the second stage, the components of δ_{jt} are estimated using linear methods.

The parameters in θ_z govern consumer heterogeneity and are crucial to understanding merger effects, as they will determine correlation between insurance choice preferences and medical consumption demand, i.e. expected cost. These parameters are identified through the richness of the data: the correlation between the detailed, observed health status measures (Z_i) and the shares of consumers selecting particular products. Importantly, this does not require any particular assumption on the information set of a particular consumer. Consumers know their preferences, but do not necessarily need to know each component of Z in the econometric model that identifies those preferences, nor those of their neighbors.

The key identification challenge to estimating the base parameter vector θ_0 from the estimates of $\hat{\delta}$ is the potential correlation between premium, the primary care copay, and the unobserved product characteristic, ξ .⁸ I address this endogeneity concern through several methods—two-way fixed effects and several potential instruments—each with similar quantitative results, displayed in Table 3.

I estimate four specifications that use Hausman instruments with and without county-level fixed effects, demographic instruments, and both Hausman and demographic instruments. The Hausman instruments are the average monthly premium of each product in other counties and the average primary care copay of each product contract in non-contiguous counties.⁹ The demographic instruments are the county-level average risk in both the TM and MA populations, average income, fraction of the population senior population over the age of 75, and firm-level administrative costs.

Across all specifications, the estimate of premium sensitivity is nearly constant. The estimate of consumer sensitivity to the copay for primary care is less

⁸A key unobserved feature of product quality are provider networks, but these are typically constant over time (more easily controlled for). There 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.

⁹Copays do not vary within products across counties, but within contracts—product type groups—there is variation in the benefit package offered. I use non-contiguous counties because neighboring counties often have similar product mixes.

	TWFE	IV			
	(1)	(2)	(3)	(4)	(5)
Monthly Premium	-0.09*** (0.02)	-0.08*** (0.01)	-0.09*** (0.01)	-0.08 (0.05)	-0.08*** (0.01)
Primary Care	-0.37*** (0.08)	-0.35** (0.13)	-0.49*** (0.12)	-2.07*** (0.50)	-0.47*** (0.12)
Out-of-Pocket Limit	-0.06 (0.05)	-0.07 (0.04)	-0.06 (0.04)	0.24* (0.10)	-0.05 (0.04)
Specialist	0.15 (0.13)	0.46*** (0.08)	0.48*** (0.08)	1.00*** (0.19)	0.49*** (0.08)
Outpatient	0.02 (0.02)	-0.02 (0.01)	-0.02* (0.01)	-0.09** (0.03)	-0.02* (0.01)
Inpatient Stay	0.00 (0.00)	-0.00 (0.00)	0.00 (0.00)	0.04** (0.01)	0.00 (0.00)
Emergency Room	-0.02 (0.03)	0.02 (0.03)	0.01 (0.03)	-0.18** (0.07)	0.01 (0.03)
Ambulance	-0.02** (0.01)	-0.04*** (0.01)	-0.04*** (0.01)	-0.05*** (0.01)	-0.04*** (0.01)
Imaging	-0.00* (0.00)	-0.00 (0.00)	0.00 (0.00)	0.00* (0.00)	0.00 (0.00)
Medical Device Coins.	-0.04** (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.00 (0.03)	-0.00 (0.01)
Outpatient Coins.	0.09* (0.03)	0.01 (0.02)	0.00 (0.02)	-0.16* (0.06)	-0.00 (0.02)
Outpt. Drugs Coins.	0.03** (0.01)	-0.03*** (0.01)	-0.03*** (0.01)	0.02 (0.02)	-0.03*** (0.01)
Fixed Effects					
Part D, Rating, Year	✓	✓	✓	✓	✓
County			✓		
Product-County	✓				
Huasman Instr.		✓	✓		✓
Demographic Instr.				✓	✓
F Stat - Premium		742	998	8.65	245
F Stat - Primary Care		277	253	6.16	86.6

Note: The results of different specifications to estimate the base demand parameters are quantitatively similar. Specification (5) is the preferred estimation and used in counterfactual exercises. The monthly premium is denominated in \$10, and all other variables are copays denominated in \$10 with the exception of variables labeled with “coins” (percentage point coinsurance) and the out-of-pocket limit (thousands of dollars). The significance stars ***, **, and * denote statistical significance at the 0.1%, 1%, and 5% level respectively.

Table 3: Estimation Results for Base Parameter Vector

robust, but typically in the range of -0.3 to -0.5. I use the final specification which employs both sets of instruments as the preferred specification throughout the rest

of the paper.

5.3 Results

	Monthly Premium	Primary Care Copay
	<i>All Consumer Semi-Elasticity</i>	
Mean	-2.29 (0.35)	-14.4 (3.86)
10 th Percentile	-8.00	-44.5
90 th Percentile	-0.31	-2.17
	<i>Entering Consumer Semi-Elasticity</i>	
Mean	-8.49 (1.09)	-44.8 (12.2)
10 th Percentile	-9.33	-49.4
90 th Percentile	-7.15	-40.1

Note: The tables shows the mean, 10th percentile, and 90th percentile for the semi-elasticities of demand with respect to the monthly premium and copay for primary care in the demand estimation, both for all consumers and for entering consumers that face no switching cost. The standard errors of the means are displayed in parenthesis. The semi-elasticities represent the percent change in the probability a consumer purchases their chosen plan given a \$10 increase in the characteristic.

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

The implied semi-elasticities of demand are summarized in Table 4. Consumer demand is elastic with respect to the primary care and, to a lesser extent, the monthly premium. The mean semi-elasticity for primary care is -14.4, which implies that the market share of a product will fall by 14.4 percent as the result of a \$10 increase in the primary care copay. This likely reflects that a \$10 increase is a relatively large change (the standard deviation across products is about \$8), more than 90% of the consumers make an office visit during the year, and the average number of office visits is about 8 per year.

The mean semi-elasticity with respect to premium is -2.29. This implies that the average consumer is 2.29 percent less likely to select a plan given a 10\$ increase in the monthly premium. This implies a low elasticity from the consumers perspective, given low and occasionally \$0 premiums. From the perspective of the firm, if the per-person level subsidy is included as the effective premium being paid to the firm, the mean elasticity of demand is -1.6.

The low elasticities are due in part to sizeable switching costs. For entering consumers that face no switching costs, the semi-elasticities for premium and primary care copays are -8.5 and -44.8, respectively. The premium-elasticity of demand for these consumers is -4.5. Consumers face an average switching cost of \$876 per month, which is much greater than the average monthly premium paid by consumers and similar in magnitude to the total per-person monthly revenue received by the firms.

Consumer heterogeneity depends on age, sex, the six most common clinical conditions (listed in order from most to least prevalent), the mean risk score for each consumer across all years, and the mean risk score squared.¹⁰ The estimates for a select number of cost-sharing terms are presented in Table 5, and the remainder are presented in Appendix Table A2. 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 with rheumatoid arthritis and vascular disease (a general category for illnesses related to arteries) have stronger sensitivities to the details of the insurance contract.

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 ex-post distribution of total and net cost. Consistent with adverse selection, the willingness to pay for low cost-sharing increases with the total cost of the consumers (Figure 1a). However, after accounting for risk adjusted subsidies, the forces of selection are more ambiguous. Figure 1b shows that willingness to pay is U shaped in net cost. Variance in spending grows with the mean. Therefore, consumers with high expected spending provide opportunities for both

¹⁰By using the mean risk score across all years, it captures demand based on prior health conditions and anticipated medical conditions.

	Premium (\$10)	Copays (\$10)				
		Primary	Specialist	Outpatient	Inpatient	Imaging
Over 75	0.018*** (0.004)	0.002 (0.022)	0.000 (0.028)	-0.012 *** (0.004)	0.002 (0.001)	-0.010*** (0.003)
Female	-0.004 (0.003)	-0.017 (0.020)	0.016** (0.025)	-0.004 (0.003)	-0.007*** (0.001)	-0.005** (0.003)
Heart Arrythmia	-0.001 (0.005)	0.088*** (0.034)	0.034 (0.040)	0.004 (0.005)	-0.010*** (0.002)	-0.001 (0.004)
Vascular Disease	0.000 (0.005)	0.088** (0.035)	-0.083 ** (0.042)	0.026*** (0.005)	-0.013*** (0.002)	-0.012*** (0.004)
Diabetes w/ Compl.	0.011*** (0.005)	0.016 (0.036)	-0.156*** (0.042)	0.000 (0.005)	0.001 (0.002)	-0.004 (0.004)
Diabetes w/o Compl.	0.003** (0.005)	-0.054 (0.035)	0.081* (0.042)	-0.004 (0.006)	0.002 (0.002)	-0.007 (0.004)
Breast/Prost. Cancer	0.011 (0.006)	0.011 (0.041)	0.028 (0.050)	0.001 (0.006)	-0.013*** (0.003)	-0.000** (0.005)
Rheum. Arthritis	0.013*** (0.007)	-0.212*** (0.049)	-0.034 (0.057)	0.016** (0.007)	-0.018*** (0.003)	-0.007*** (0.006)
Mean Risk Score	0.014** (0.005)	-0.054* (0.031)	0.142** (0.040)	0.000 (0.005)	-0.018*** (0.002)	-0.028*** (0.004)
Mean Risk Score ²	0.000 (0.001)	0.000 (0.005)	-0.012 (0.007)	0.000 (0.001)	0.003*** (0.000)	0.004*** (0.001)

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 0.1%, 1%, and 5% level respectively.

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

adverse and advantageous selection relative to the risk adjusted subsidy (?).

The willingness to pay to reduce the primary care copay is relatively large and would be challenging to capture in a model where consumers have rational expectations about their medical consumption. For context, the predicted average willingness to pay to reduce the primary care copay would be about \$6.5 per month if consumers were risk neutral. Using constant absolute risk aversion with a pa-

parameter of -0.0018 (? , ? , ?), the predicted average consumer willingness to pay to reduce *all* out of pocket spending risk is about \$72 per month.

The high willingness to pay suggests there could be information frictions, either between stages of consumer demand (e.g. behavioral moral hazard ? or biased beliefs (?)) or about aspects of the products (?). Because the primary downside of modeling each stage separately is the inability to use demand for welfare calculations, these results should encourage the reader that this empirical approach does not forfeit much in its pursuit of tractability and flexibility.

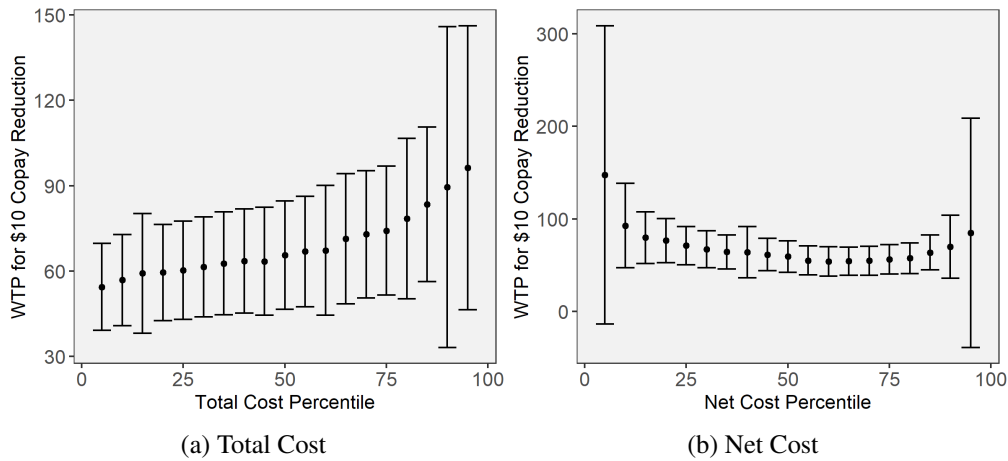


Figure 1: Adverse Selection Plays an Important Role in Insurance Demand

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.

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 (? , ? , ?). 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, on average, the semi-elasticity with respect to a \$10 increase in the primary care copay is -8.2%. I also find evidence that a \$10 increase in the primary care copay leads to a 0.2 percentage point increase in inpatient mortality.

6.1 Specification

The model of medical consumption, $m^*(\cdot)$, is specified as a log-linear in plan characteristics, an unobserved individual fixed effect, monthly fixed effects, and an idiosyncratic medical demand error. Let $m_{i\tau}$ be 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 . Medical consumption depends on patient characteristics via the individual-specific constant term, η_i , and consumers have monthly idiosyncratic medical demand shocks, $\omega_{i\tau}$.

Medical consumption is specified as

$$\log(m_{i\tau} + \frac{1}{12}) = \eta_i + \beta' X_{j(i)\tau} + \lambda_\tau + \gamma' F_{j(i)} + \omega_{i\tau} \quad (13)$$

where λ_τ 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 (Gerdtham, 1991; Gerdtham and Hansson, 1994; Gerdtham and Hansson, 1995). I follow Gerdtham (1991) in using $m_{i\tau} + \frac{1}{12}$ in order to allow elasticities to be comparable to annual elasticity estimates that use $m_{i\tau} + 1$. The results of this section are robust to other adjustments, such as setting the constant at the minimum or 10th percentile of positive monthly consumption.¹¹

¹¹While 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.

6.2 Estimation and Identification

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 (13) may produce biased estimates of β because a potential correlation between $X_{j(i)\tau}$ and $\omega_{i\tau}$.

This paper exploits within-product changes in cost-sharing terms to consistently identify the elasticity (?). Due to consumer switching, the observed consumer-level variation in cost-sharing terms may be correlated with consumer beliefs about their future consumption. To address this, I use the change in the primary care copay for the product that an individual was enrolled in during the prior year as an instrument for the change in the individual's actual primary care copay. Due to strong consumer inertia, the instrument is a strong predictor of consumer-level variation (? , ? , ? , ?). This approach has the benefit of using the variation that firms are interested in when making product design decisions: the change in average medical consumption caused by a change in a product's cost-sharing term.

To be explicit, I separate the primary care copay from the vector of cost-sharing terms, $x \in X$, as shown in equation (14).

$$\Delta_i \tilde{m}_{i\tau} = \beta_o \Delta_i x_{j(i)\tau} + \beta'_{-o} \Delta_i X_{j(i)\tau} + \Delta_i \lambda_\tau + \gamma' \Delta_i F_j(i) + \Delta_i \omega_{i\tau} \quad (14)$$

$$(15)$$

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 instrument for $\Delta_i x_{j(i)\tau}$ is $\Delta_j x_{j(i)\tau}$, the 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.

One potential identification concern is that insurance firms have foresight about changes in negotiated rates that appear as changes in total medical spending at the consumer level. In Appendix Section C.4, I construct a measure of physician

service intensity that controls for this possibility and re-estimate the model using this constructed measure. The results are qualitatively and quantitatively similar.

This identification strategy is also extended to estimate the relationship between the primary care copay and inpatient mortality. For details on this estimation and the results, see Appendix Section D.

6.3 Results

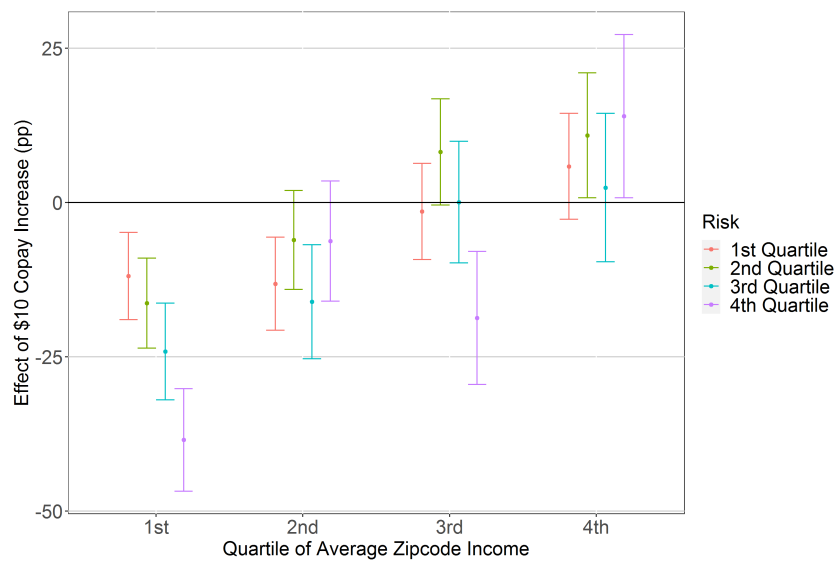


Figure 2: Medical Consumption is Elastic With Respect to the Primary Care Copay

Note: Consumer semi-elasticity with respect to the primary care copay is highest among consumers in the lowest income zipcodes or with the highest medical risk. This figure shows the results of the medical consumption elasticity estimation. The primary care copays are denominated in \$10 and effect sizes are shown in percentage points. The estimation also includes monthly fixed effects, firm-level fixed effects, and all contract characteristics included in Table 3. Consumers are divided by income of the zipcode in which they live and their average risk score throughout the sample. Confidence intervals are shown at the 0.1% level.

Changes in the copay for primary care occur in the data across 17 different products affecting 64% consumers at some point during the sample period. Nearly all observed changes in the copay for primary care are increases from \$0 to \$10 (70% of treated observations), \$15 (19% of treated observations), or \$20 (9% of

	Risk Quartiles			
	1 st	2 nd	3 rd	4 th
<u>1st Income Quartile</u>				
Primary Care	-0.119*** (0.028)	-0.163*** (0.029)	-0.241*** (0.031)	-0.385*** (0.032)
1 st Stage F-statistic	5.5×10 ⁴	7.7×10 ⁴	8.8×10 ⁴	11.7×10 ⁴
N (000s)	286	381	378	362
Individuals (000s)	9.1	10.5	10.3	11.0
<u>2nd Income Quartile</u>				
Primary Care	-0.132*** (0.029)	-0.061* (0.031)	-0.161*** (0.036)	-0.063* (0.038)
1 st Stage F-statistic	5.7×10 ⁴	7.8×10 ⁴	7.9×10 ⁴	10.0×10 ⁴
N (000s)	295	371	345	325
Individuals (000s)	9.5	10.3	9.4	9.8
<u>3rd Income Quartile</u>				
Primary Care	-0.015 (0.030)	0.082** (0.034)	-0.000 (0.039)	-0.187*** (0.042)
1 st Stage F-statistic	7.5×10 ⁴	10.4×10 ⁴	11.6×10 ⁴	13.6×10 ⁴
N (000s)	307	373	351	332
Individuals (000s)	9.7	10.2	9.4	10.0
<u>4th Income Quartile</u>				
Primary Care	0.058* (0.033)	0.109*** (0.040)	0.024 (0.047)	0.140*** (0.052)
1 st Stage F-statistic	10.9×10 ⁴	13.6×10 ⁴	14.2×10 ⁴	15.6×10 ⁴
N (000s)	340	390	337	303
Individuals (000s)	10.1	10.6	9.0	9.0

Note: Medical consumption semi-elasticity is highest among consumers in the lowest income zipcodes and with the highest medical risk. The primary care copays are denominated in \$10. Each panel displays estimation results for a particular risk- and income-quartile combination. The estimation also includes monthly fixed effects, firm-level fixed effects, and all contract characteristics included in Table 3. ***, **, and * denote statistical significance at the 0.1%, 1%, and 5% level respectively.

Table 6: Medical Consumption Responds to Primary Care Copays

treated observations). A small fraction of consumers (1%) experience a decrease from \$25 to \$20. Each of product changes affect large consumers across the entire medical risk and zip code-level income distribution.

Figure 2 displays sixteen versions of the medical consumption model, divided by both local average income and individual medical risk. The coefficients represent the semi-elasticities, or the percentage decrease in the dependent variable caused by a \$10 increase in the copay for primary care, and the effect sizes are displayed in percentage points.

The sample is divided into four quartiles of risk based on the pooled distribution of mean risk scores with one adjustment: the first quartile is slightly reduced to include only individuals with no clinical diagnoses. The risk quartiles are defined by the mean risk score of the consumer throughout the sample.¹² Next the sample is further divided into four quartiles of income. Since the data does not contain consumer-level income, I divide the sample according to the zip code-level average income, weighted according to the distribution of the sample across zip codes. The model is estimated separately for each risk-income quartile combination, each of which contains roughly 300 thousand consumer-month observations and 10 thousand individual consumers. Table 6 contains more information on the estimation and sample sizes.

The average semi-elasticity of primary care copay across all groups is -0.08, and the elasticity is greatest among consumers in low-income zip codes.¹³ Within consumers in low-income zip codes, the elasticity with respect to the copay for primary care is increasing in the level of medical risk. This result may be driven by two features of medical consumption demand. First, consumers that expect to have more doctor visits throughout the year face a higher out-of-pocket price increase as the result of greater primary care copays. Second, procedure prices are often

¹²Risk scores are persistent but do vary over time. If I were to construct this measure based on annual risk scores, 51.3% of consumers would be in a single risk quartile throughout the entire sample and 89.7% in two or fewer.

¹³For consumers in the higher income zip codes, medical consumption elasticities are positive, though not significant at the 0.1% level. Notably, positive elasticities are not necessarily counter to classic economic incentives. Since there are many channels through which consumers can seek medical care, it is possible that greater copays for primary care lead consumers to seek out more costly alternatives to care.

increasing in the patient's risk. As results, similar decisions, e.g. not pursuing a minor surgery, may appear as larger spending elasticities for higher risk patients.

Because most of the changes are increases from \$0, the implied arc-elasticity is roughly half the semi-elasticity. The arc-elasticities are between -0.06 to -0.19 among the low-income population. This is consistent with, though a bit lower than, magnitudes found elsewhere in the literature (?, ?, ?), including studies of low-income populations specifically (?).

Two aspects of these findings are new to the literature. First, these estimates reflect the effect of primary copays on total spending, not just office visit spending, which highlights a gateway channel of medical consumption. Second, the elasticities of higher-income consumers are very small, suggesting that the only low-income consumers respond significantly to copays in the this setting. This has potential implications for other policy interventions, such as cost-sharing subsidies or restrictions on copay levels.

Appendix Figure A1 shows the results of the identical estimation procedure using whether or not a consumer receives a particular medical service in a month as a dependent variable. Results are displayed for primary care visits, minor outpatient procedures, major procedures (typically require general anesthesia), and hospital visits. Elasticities are greater for primary care visits and minor outpatient procedures but negligible for major procedures and hospital visits. This is consistent with the idea that while the primary care copay primarily affects the probability of primary care office visits, this has follow on effects for other types of elective procedures and less of an effect more urgent care. There are no positive elasticities with respect to primary care visits themselves but also no evidence of diversion to other services.

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 MA market in Massachusetts: Tufts Health Plan (Tufts), Blue Cross Blue Shield of Massachusetts (BCBS), and United Healthcare (United).

A merger in this model is characterized as maximizing the joint profit of the two sets of products owned by each of the merging parties. In order to mitigate potential problems of multiple equilibria, the mergers are simulated via a homotopy. I check for robustness of the solved pre-merger equilibrium by re-solving assuming that each merging party has a 1% interest in the other, and then solving for the pre-merger equilibrium again from this 1% merger starting point. Then, the full post-merger equilibrium is solved incrementally in which the merging firms gain a 5% interest in each other during each step. The summary statistics for all six firms that operate in the state are displayed in Table A3.

Table 7 displays the change in HHI from each merger (using pre-merger market shares), the consumer-weighted mean values of the premium and primary care copay in the simulated pre-merger equilibrium, the average effect of the merger on each product weighted by post-merger market share, and the consumer-weighted mean values of the premium and primary care copay in the post-merger equilibrium.

The re-solved equilibrium generates substantially greater premiums than the observed premiums shown in Table A3. In Appendix Section G, I show results are qualitatively similar from a merger simulation under an alternative assumption that better matches the observed pre-merger equilibrium.

The average merger effects hide a substantial amount of heterogeneity in product-level effects. In each of the three mergers, at least 25% of consumers will experience a *decrease* in their primary care copay as a result of the merger. Fewer consumers will experience decreases in the premium, but even in the largest merger, the premiums fall for 5% of consumers. In Appendix Figure A2, I show that this heterogeneity in the direction of merger effects is closely tied to the merger incentives outlined in Section 3.2.

The effects of these changes on consumers are displayed in Table 8. To characterize the heterogeneity in the effects, I separate the consumers into two roughly equal groups: consumers that live in the highest income zip codes and consumers that live in the lowest income zip codes.¹⁴ The consumer level effects on medical

¹⁴For consistency with previous sections and across mergers, this division is equal for all of Massachusetts, not only the markets affected by the merger.

	Mean Δ HHI	Pre Baseline		Merger Effect		Post Mean	
		Premium	Copay	Premium	Copay	Premium	Copay
Tufts - BCBS	2,240						
<i>Merging Parties</i>		374	11.8	22.1	0.11	406	12.3
<i>All Firms</i>		332	13.2	15.7	0.07	351	13.5
Tufts - United	727						
<i>Merging Parties</i>		377	9.9	13.5	0.14	379	10.3
<i>All Firms</i>		322	13.8	2.7	0.02	310	14.0
BCBS - United	488						
<i>Merging Parties</i>		271	13.6	9.8	-0.08	275	15.1
<i>All Firms</i>		321	13.8	4.0	0.01	326	14.1

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 MA market. The average Δ HHI is computed as the predicted change in HHI using pre-merger market shares, weighted by market size. The pre-merger and post-merger average values are weighted by pre-merger and post-merger enrollment, respectively. The middle group of merger effects weights the product-level merger effect by post-merger enrollment, which controls for changes in market composition due to switching.

Table 7: Mergers Lead to Higher Average Premiums and Primary Care Copays

consumption are limited to those that result from changes in the copay for primary care and do not include, for example, the effects of switching to a low-cost firm for which I do not have causal identification.

Mergers lead to a change in market composition as some consumers leave MA for Traditional Medicare to escape increasing premiums and copays. In each merger, the Massachusetts MA market loses between 4% and 7% of the pre-merger consumers in affected markets. Due to limited data about Traditional Medicare, I limit the description of consumer-level merger effects to those that remain in the MA market post-merger.

Focusing only on consumers that remain in the market post merger, only the merger between the two largest firms, Tufts and BCBS, leads to an increase in the primary care copay and a corresponding decrease in medical consumption. In the other two mergers, consumers pay face lower primary care copays on average, and in the Tufts-United merger, consumers also pay lower premiums.

Changes in medical consumption are concentrated among consumers in low-income zipcodes. Despite facing similar changes in the primary care copay, the changes in medical consumption are larger by an order of magnitude. In the Tufts-BCBS and the BCBS-United mergers, consumers in low-income zipcodes reduce their medical consumption by \$10.5 and \$8.8 per year on average respectively, similar magnitudes to the premium effect of the merger. And in the Tufts-United merger, consumers in low-income zipcodes increase medical consumption by \$46.5 per year.

	Primary Care Copay Effect	Premium Effect	Medical Cons. Effect (\$/year)	Mortality Effect (bp)	Savings per Life (\$000s)
Tufts - BCBS					
Mean Effect	0.07	18.5	-5.6	0.112	502
High Income	0.10	24.2	-1.0	0.097	103
Low Income	0.05	12.6	-10.5	0.128	821
Tufts - United					
Mean Effect	-0.21	-7.2	24.1	-0.446	518
High Income	-0.14	-6.1	3.2	-0.146	216
Low Income	-0.28	-8.3	46.5	-0.807	576
BCBS - United					
Mean Effect	-0.01	3.6	-3.6	-0.006	-
High Income	-0.03	4.6	0.9	-0.034	255
Low Income	0.00	2.6	-8.8	0.023	3,819

Note: While the average effects of a merger are small, the medical consumption of consumers in low income zip codes responds elastically to increases or decreases in their primary care copay. This table displays the consumer level effects of each merger, averaged across all consumers affected by the merger, and separately for consumers in high and low income zip codes. The mortality effect, displayed in basis points or hundredths of a percent, is the product of the primary copay effect and the IV estimates in Table A4. The final column is the result of dividing the predicted change in medical consumption by the predicted change in twelve month inpatient mortality. Averages are weighted by post-merger enrollment in MA.

Table 8: Consumer-level Effects of a Merger are Heterogeneous

Elastic medical consumption increases both the potential for increased moral hazard, i.e. wasteful medical consumption, and more access to valuable care. In order to put these results in context, I combine them with estimates on the effect

of the primary care copay and inpatient mortality based on their risk and zip code-level income group. For details on these estimates, see Appendix Section D. These estimates imply that changes in the average primary care copay due to each merger will lead to an average change in 12-month inpatient mortality of between 0.112 and -0.446 basis points.

In Table 8, *savings per life* refers to the savings (or spending) associated with these increases (or decreases) in mortality risk. The average effects corresponds to \$502 and \$518 thousand in medical spending per expected life, excluding the Tufts-United merger which leads to a decrease in both mortality and medical spending. The savings per life is again concentrated among consumers in low-income zip-codes.

These figures are well below most estimates for the value of a statistical life (VSL). The federal government uses a value of \$10 million for all individuals when assessing policy cost-benefit analysis. Even when taking into consideration a reduced life expectancy for senior citizens, VSL estimates for individuals in the Medicare eligibility age range exceed \$1 million (?). However, recent work using the choices made by the MA population estimates that the VSL is \$402 thousand for individuals aged 67 and declining with age (?).

If we use high values for the VSL (exceeding \$1 million), the results from the merger analysis in Table 8 would imply that increases in the primary care copay and decreases in medical consumption are welfare reducing (and opposite effects welfare improving), because the reduction in health status outweighs the benefit of lower resource use on health care. If we use the lower values of the VSL from ?, the welfare result is more ambiguous. It is also important to note that the effects on inpatient mortality are only one piece of the total health effect of reduced medical consumption, most of which is challenging to measure or assign welfare values. This is an important area for future research.

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. Moreover, the magnitude of the effect of a merger on medical consumption and consumer health is similar 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. Consumers respond to lower (higher) levels of cost-sharing by increasing (decreasing) their medical consumption. And lower levels of cost-sharing are associated with 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. While reductions in competition via mergers on average lead to higher levels of cost-sharing, the effects are heterogeneous across products, markets, and mergers.

In the merger with the largest effect on the primary care copay, average medical spending increases by \$24 per person per year and the likelihood of an inpatient death in a twelve month period decreases by 0.004 percentage points. This implies a mortality reduction at the cost of about \$518 thousand per expected death. At typical estimates of the value of a statistical life, the benefit of reduced deaths far outweigh the resource cost of additional medical spending.

Other ways in which insurance firms compete includes the design of the hospital and physician network (? , ? , ?), 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. The extension of this model to incorporate these other mechanisms, especially non-financial mechanisms, is an important agenda for future research.

Acknowledgments

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, Keaton Miller, and Paul Grieco, as well as all the participants of the workshops and seminars where the paper has been presented.

References

ONLINE APPENDIX

A Supplemental Figures and Tables

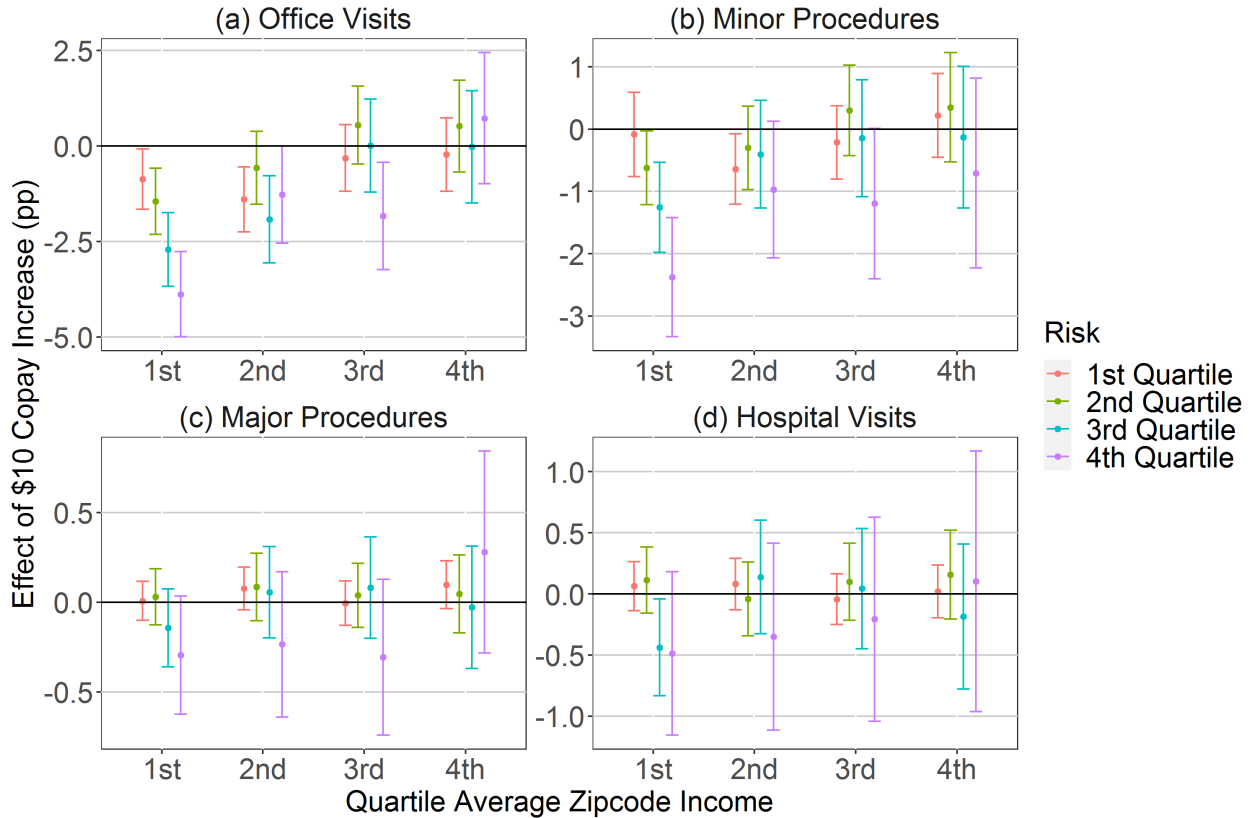


Figure A1: Primary Care Copays have Larger Effects on Office Visits and Minor Procedure

Note: Higher primary care copays affect medical consumption primarily through office visits and minor procedures. This figure shows the results of following the identical estimation procedure outlined in Section 6.2 using as the dependant variable whether a consumer receives a particular kind of service in each month: primary care visits (a), minor procedures (b), major procedures (c), and hospital visits (d). The primary care copays are denominated in \$10 and effect sizes are shown in percentage points. The estimation also includes monthly fixed effects, firm-level fixed effects, and all contract characteristics included in Table 3. Consumers are divided by income of the zipcode in which they live and their average risk score throughout the sample. Confidence intervals are shown at the 0.1% level.

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 MA 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 A1: More Competitive Markets have Lower Average Cost-sharing Levels

	OOP Limit	Copays (\$10)		Coinsurance Rates (pp)		
	(\$1000)	Emergency	Ambulance	Outpatient	Med Device	Drug
Over 75	-0.041*** (0.010)	0.006 (0.009)	-0.008*** (0.003)	-0.013** (0.006)	0.000 (0.005)	0.000 (0.002)
Female	0.025*** (0.009)	-0.003 (0.008)	0.002 (0.002)	-0.014*** (0.005)	0.011*** (0.004)	-0.001 (0.002)
Heart Arrythmia	0.009 (0.016)	0.006 (0.014)	-0.007* (0.004)	0.004 (0.008)	0.002 (0.006)	-0.006* (0.004)
Vascular Disease	0.031* (0.017)	-0.012 (0.014)	-0.008* (0.004)	0.011 (0.009)	0.001 (0.006)	-0.001 (0.004)
Diabetes w/ Compl.	-0.066*** (0.017)	-0.010 (0.014)	0.006 (0.004)	0.019** (0.009)	0.022*** (0.007)	0.023*** (0.004)
Diabetes w/o Compl.	0.036*** (0.015)	-0.009 (0.014)	-0.003 (0.004)	0.000 (0.009)	-0.002 (0.007)	-0.004 (0.004)
Breast/Prost. Cancer	0.021 (0.019)	-0.037** (0.017)	0.004 (0.005)	0.001 (0.010)	0.005 (0.008)	-0.006 (0.004)
Rheum. Arthritis	-0.010 (0.023)	-0.055*** (0.019)	-0.066*** (0.006)	0.003 (0.012)	0.000 (0.009)	0.015*** (0.005)
Agg. Risk Score	-0.008 (0.014)	-0.018 (0.013)	0.009** (0.004)	0.002 (0.009)	-0.002 (0.006)	-0.002 (0.004)
Agg. Risk Score ²	0.006** (0.003)	0.001 (0.002)	-0.002** (0.001)	0.000 (0.001)	-0.002** (0.001)	-0.001** (0.001)

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 0.1%, 1%, and 5% level respectively.

Table A2: Estimates of Demand Heterogeneity - Continued

	MA Share	Avg Prem	Avg Copay	Avg. Risk		Risk Adj. Cost	
				Data	Model	Data	Model
Tufts	0.47	112	11.6	1.04	1.07	672	672
BCBS	0.26	113	19.0	0.82	0.82	765	765
United	0.14	26.5	16.0	0.80	0.82	629	846
Fallon	0.06	92.6	20.3	0.93	0.93	777	704
Health New Engl.	0.05	130	20.1	1.00	1.01	696	664
Harvard Pilgrim	0.03	117	13.4	0.89	0.96	762	693

Note: This table shows summary statistics for the six firms that offer MA plans in Massachusetts. The average risk and cost comparisons are computed at observed premiums and copays to show how well the model can capture the risk heterogeneity among the firms and consumers.

Table A3: Firms are Differentiated in their Premiums, Copays, and Risk Distributions

B First Order Approximation of Merger Effects

In this section, I restate the proof of Theorem 1 in ?, with minor extensions to accommodate an environment with both copays and premium. I show how this first order approximation (FOA) incorporates the incentives of the merger outlined in Section 3.2 and is closely tied to the simulated merger effects.

To reiterate the firm's problem expressed in equation (7), consider a single product firm j . I will define $f_j^l(P, X)$, the first order condition with respect to $l \in \{\text{Copay, Premium}\}$ as a function of the vectors of premiums and primary care copays in the market. For example, the pre-merger first order condition of the firm with respect to the primary care copay is

$$0 \leq f_j^{\text{copay}}(P, X) \equiv \frac{1}{\frac{\partial s_j}{\partial x_j}} \int_i s_{ij} \left(\frac{\partial b_{ij}}{\partial x_j} + \frac{\partial mc_{ij}}{\partial x_j} \right) di + p_j + \frac{1}{\frac{\partial s_j}{\partial x_j}} \int_i \frac{\partial s_{ij}}{\partial x_j} (mc_{ij} - b_{ij}) di \quad (16)$$

There also exists an analogous pre-merger first order condition with respect to price given by $0 \leq f_j^{\text{prem}}(P, X)$.

Now consider a merger between j and another single product firm k , the post-merger first order condition can be expressed as $h_j^{\text{copay}}(P, X) \equiv f_j^{\text{copay}}(P, X) + g_{jk}^{\text{copay}}(P, X)$ where g represents GePP as defined in equation for copays in equation 8. Analogous functions, h_j^{prem} and g_{jk}^{prem} , exist for the GePP and post-merger first order conditions with respect to the premium.

Let $Q = [P, X]$ be the stacked vector of premiums and copays selected for all products in a market. Let $f(Q)$ be the stacked vector of pre-merger first order conditions and $g(Q)$ be the stacked vector of GePP, such that $f + g = h$, the post-merger first order conditions for all products for both premiums and primary care copays.

Assumption B.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. Now I can extend the main result of ?.

Theorem B.1. *Given assumption B.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

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

□

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 the product of the merger \mathbb{P} and the vector of GePP, evaluated at the pre-merger equilibrium.

$$\begin{bmatrix} \Delta p_j \\ \Delta x_j \end{bmatrix} = \mathbb{P}(P_0, X_0) g_{jk}(P_0, X_0) \quad (17)$$

$$\mathbb{P}(P_0, X_0) = \begin{bmatrix} \frac{\partial h_j^p}{\partial P} & \frac{\partial h_j^x}{\partial P} \\ \frac{\partial h_j^p}{\partial X} & \frac{\partial h_j^x}{\partial X} \end{bmatrix}^{-1} \Bigg|_{(P_0, X_0)} \quad (18)$$

where (P_0, 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 .

Merger effects are ambiguous for two reasons: adverse selection may flip the firms' standard incentives to raise prices (premiums and copays, in this case) due to recaptured profit, and the pass-through of these incentives from a merger depends on the willingness to pay for low copays among a firm's marginal consumers. The first incentive is captured by GePP and the second is captured by the merger pass-through matrix, \mathbb{P} .

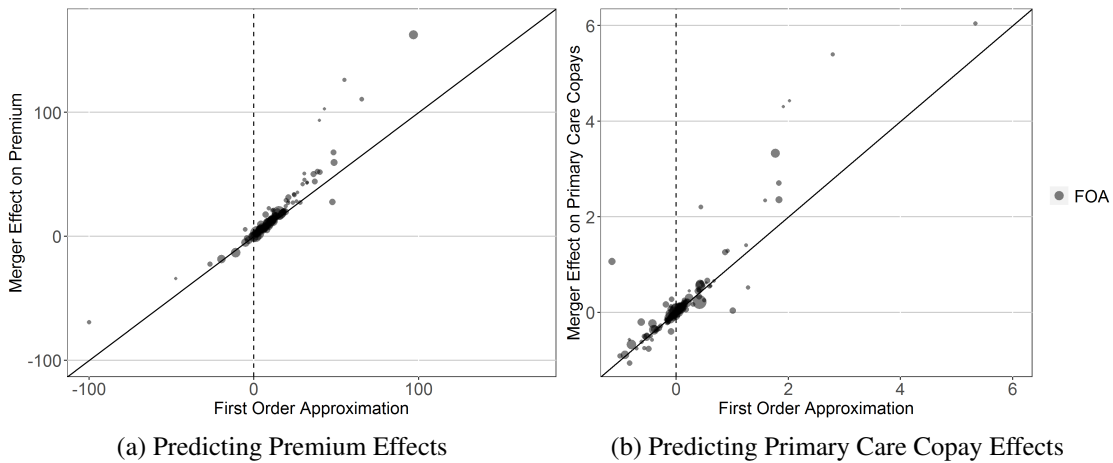


Figure A2: Merger Effects are Associated with FOA and GePP

Note: The first order approximation (FOA) outlined in Section 3.2 is a prediction of the direction and magnitude of merger effects on both the primary care copay and the premium, and the GePP is a slightly noisier prediction of direction. Each dot represents a merger-market-product, and the size of the dot represents post-merger enrollment. Each observation has both an FOA prediction and GePP prediction, connected via a light red line. The black line is the 45 degree line through the origin. The plot does not display some prediction outliers, and only products that are not constrained either pre-merger or post-merger are displayed.

Figure A2 plots the simulated premium and copay merger effects for each product in each merger and its relationship to the FOA measures. The plot shows two important results: first, negative merger effects are present for both premiums and primary care copays, and second, the magnitude and direction of the merger effects are closely related to the FOA. there are 49 simulated negative copay merger

effects.¹⁵ All but two negative simulated effect have negative FOA values. This demonstrates that the merger incentives outlined in Section 3.2 are tightly connected to simulation results.¹⁶

C Data Processing

C.1 Linking Medical Claims to Products

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 second 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 MA. However,

¹⁵This excludes products that are constrained both pre- and post-merger.

¹⁶It has an additional value of demonstrating that the simulation effect can be predicted from pre-merger information on consumer substitution patterns, consumer-level costs, and firm pass-through rates.

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 APCD 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 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.

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.

C.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. In counterfactual exercises, I use a 5% sample of consumers that ever select an MA plan, and the other samples scale accordingly.

C.3 Sample Selection for Medical Consumption Estimation

Conditional on being over the age of 65, this data exclude two populations. First, it excludes any member-months of traditional Medicare enrollment. Second, the linking IDs between the insurance enrollment panel and the medical claims data for members of United Healthcare are often incorrect (i.e. do not correspond to valid IDs in both sets of data). Because this breaks the primary source of identification in the estimation, I 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 its 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.

C.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

$$m_{ipjt} = \Gamma'_p L_{ip} + \iota_{pjt} + \varepsilon_{ip} \quad (19)$$

where ι_{pjt} is a procedure-product-time fixed effect that accounts for differences in billing practices across insurance plans. L_{ip} is a vector of features that appear on the claim bill: including 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.

The goal is to estimate $\hat{\Gamma}_p$ and use the predicted value of $\hat{\Gamma}'_p L_{ip}$ 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.

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 LASSO tuning parameter is selected for each procedure to minimize the mean squared error of prediction on a sample withheld for cross-validation. 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.

In Figure A3, I compare the results for the baseline measure of medical consumption (Figure A3(a)) and the adjusted measure of medical consumption (Figure 2(b)). Using the adjusted measure leads to quantitatively and qualitatively similar results. 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.

C.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. The binary diagnosis indicators reflect current-year diagnoses (i.e. the year of the insurance plan selection) and the risk score reflects the average risk across the whole sample period. Both of these variables are constructed using the risk score methodology that CMS uses administer the risk adjusted subsidies associated with the MA program. The methodology can be reproduced using SAS code made publicly available by CMS.

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. 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 ob-

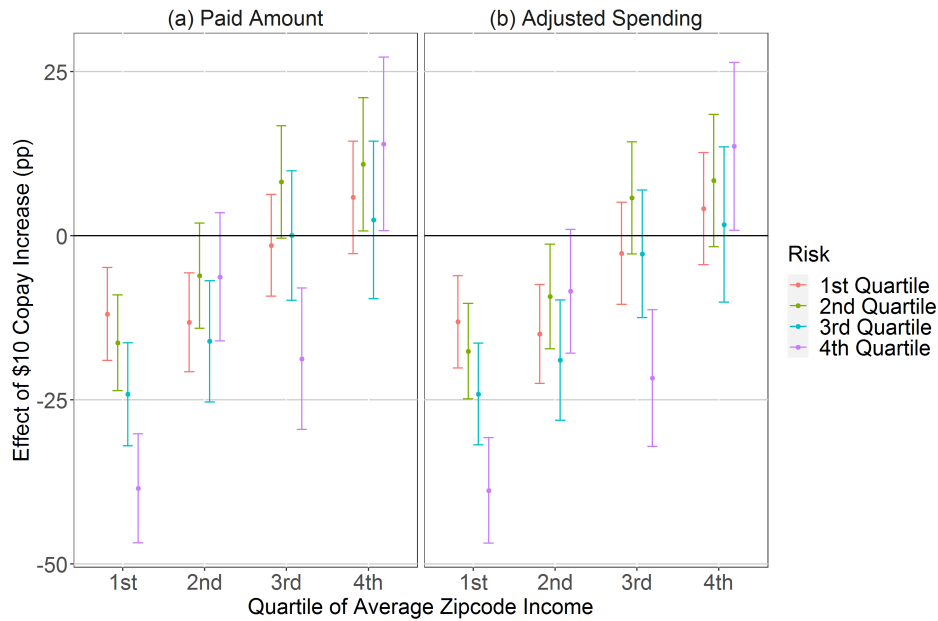


Figure A3: Robustness on Measures of Medical Consumption

Note: Consumer semi-elasticity with respect to the primary care copay is highest among consumers in the lowest income zipcodes or with the highest medical risk. This figure shows the results of the medical consumption elasticity estimation using total monthly spending as the baseline measure of medical consumption (Panel a) and the adjusted measure of medical consumption (Panel b). The primary care copays are denominated in \$10 and effect sizes are shown in percentage points. The estimation also includes monthly fixed effects, firm-level fixed effects, and all contract characteristics included in Table 3. Consumers are divided by income of the zipcode in which they live and their average risk score throughout the sample. Confidence intervals are shown at the 0.1% level.

served 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.

C.6 Measuring Additional Sources of Marginal Cost

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. I assume 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.

D 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 evidence on the relationship between primary care copays and inpatient mortality.

Let $d_{i\tau}$ be an indicator variable that represents whether consumer i has died in an inpatient facility, i.e. a hospital or hospice facility, within 12 months of month τ . I specify the following linear probability model.

$$d_{i\tau+s} = \beta^{s'} X_{j\tau} + \lambda_{\tau}^s + \zeta_j^s + \omega_{i\tau}^s \quad (20)$$

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 in Section 6.2. For this reason, the exogeneity assumption of identification is stronger than that employed in the medical consumption estimation. The within-product change in the copay for primary care must be exogenous with respect to the *level* of mortality risk of the product's enrolled population, conditional on controls. Controlling directly for health status is potentially undesirable because health status likely mediates effects on mortality. I believe this assumption is plausible (?), and specifications that include health status controls produce similar qualitative and quantitative results.

Table A4 shows the estimates and confidence intervals for the relationship between a \$10 increase in the primary care copay and inpatient mortality, measured in percentage points. The column labeled OLS displays the full sample results of the regression as specified in equation 20. I estimate four IV specifications of the model for four different groups of consumers divided by low and high income zip codes (above and below the median) and low and high income risk (above and below the median). The IV specifications use the copay of the prior year and the change in the copay of the product in which the consumer was enrolled during the prior year as an instrument for the current level of the primary care copay. Because the copay of the prior year is likely correlated with unobserved health status, I include it in both the first stage and the structural equation.

A \$10 increase in the copay for primary care leads to between a 0.07 and a 0.44 percentage point increase in 12-month inpatient mortality, where the total population mean is 1.6 percent. The strongest relationship is among high risk consumers in low-income zip codes, where the effect is 0.44 percentage points relative to a mean inpatient mortality rate of 5.1 percent. These effects, while economically significant, is still small relative to the baseline heterogeneity in mortality risk. For example, a one standard deviation increase in a patient's medical risk score is asso-

	OLS Full Sample	IV			
		Low Income		High Income	
		Low Risk	High Risk	Low Risk	High Risk
Primary Care	0.04 (0.03)	0.15*** (0.03)	0.44** (0.18)	0.12*** (0.03)	0.07 (0.30)
Primary Care, Prior Year		-0.00*** (0.00)	-0.01 (0.00)	-0.00 (0.00)	0.01* (0.01)
Specialist	0.09*** (0.02)	0.07 (0.05)	1.83*** (0.26)	0.14** (0.05)	-0.83** (0.34)
Outpatient	-0.01 (0.00)	0.02*** (0.00)	-0.03 (0.02)	-0.02*** (0.00)	-0.04 (0.03)
Inpatient Stay	0.00 (0.00)	-0.00*** (0.00)	0.01 (0.01)	0.00 (0.00)	0.01 (0.01)
Emergency Room	0.03*** (0.01)	-0.00 (0.01)	0.05 (0.04)	0.02** (0.01)	0.06 (0.05)
Ambulance	0.01* (0.00)	-0.00 (0.00)	-0.02 (0.02)	-0.01** (0.00)	0.00 (0.02)
Outpatient Coins	-0.02*** (0.01)	-0.01 (0.01)	0.08* (0.05)	0.02* (0.01)	-0.24*** (0.06)
Medical Devices	-0.00 (0.00)	0.01 (0.01)	0.01 (0.03)	0.02*** (0.01)	-0.03 (0.03)
Outpatient Drugs	-0.01*** (0.00)	0.01*** (0.00)	0.05*** (0.01)	0.01*** (0.00)	0.02 (0.02)
Diagnostic Imaging	-0.00 (0.00)	-0.01** (0.00)	-0.04** (0.02)	-0.00 (0.00)	0.01 (0.02)
Out-of-Pocket Limit	0.03***	0.02**	-0.03	-0.01	-0.10
Fixed Effects					
Product	✓	✓	✓	✓	✓
Month	✓	✓	✓	✓	✓
1 st Stage F-Statistic		2.23×10 ⁶	1.40×10 ⁶	4.61×10 ⁵	5.78×10 ⁵

Note: A \$10 increase in the primary care copay is associated with an increase in inpatient mortality, and up to a 0.44 percentage point increase for high risk patients in low-income zip codes. High and Low income are determined by consumers in zip codes above and below the median of zip code-level mean income. High and low risk and determined by consumers above and below the median risk score. All independent variables are copays, denominated in \$10, unless otherwise specified. Coinsurance rates and the OOP limit are denominated in percentage points and thousands of dollars, respectively. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level respectively.

Table A4: Evidence of Relationship between Morality and the Primary Care Copay

ciated with a 4.3 percentage point increase in 12-month inpatient mortality.

These results are 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 (?, ?, ?).

E 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.

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

$$\phi_{ijt} = \beta' X_{jt} + \lambda_t + \gamma_{j1} M_{it} + \gamma_{j2} M_{it}^2 + \omega_{ijt}^{coins} \quad (21)$$

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.

	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. ***, **, and * denote statistical significance at the 0.1%, 1%, and 5% level respectively.

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

F Estimating the Bid Function

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

$$b_{ijt} = r_{si} \left(\min\{Bench_j, bid_j(p_j, X_j)\} + \Lambda_j \max\{Bench_j - bid_j(p_j, X_j), 0\} \right) \quad (22)$$

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 MA 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 can be inferred from equation (22). I follow ? 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.

$$bid_{jt} = \alpha p_{jt} + \beta x_{jt} + \Gamma' X_{jt} + \gamma_j + \lambda_t + \zeta_{jt} \quad (23)$$

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.

G Alternative Merger Simulation Assumption

In the estimation of the model, I do not impose any equilibrium conditions. As a result, when solving the model for the main results of the paper, the model does not precisely match the observed equilibrium conditions. This section outlines an alternative method for solving the pre-merger and post-merger equilibrium that better matches observed premiums and primary care copays.

First, I allow for a wedge between the predicted profit maximizing premium (or copay) and the observed premium (or copay). The estimation procedure does not

	(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: This table shows the estimates of a bid policy function using a national panel of MA plans between 2011 and 2019. The final specification (3) contains the estimates used in the main results of the paper. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level respectively.

Table A6: Estimated Plan Bid Function

assume or enforce that first order conditions for profit maximization in the model are met. Instead, the observed premiums and copays are rationalized by allowing firms to have a direct preference in their objective function over the level of each characteristic, as shown in equation (24).¹⁷ These wedges are held constant

¹⁷Standard practice is to match observed prices through a residual in a product's marginal cost.

throughout the counterfactual exercise.

	Mean Δ HHI	Pre Baseline		Merger Effect		Post Mean	
		Premium	Copay	Premium	Copay	Premium	Copay
Tufts - BCBS	2,262						
<i>Merging Parties</i>		116	13.8	3.9	1.45	120	15.3
<i>All Firms</i>		104	14.4	2.9	1.04	106	15.5
Tufts - United	867						
<i>Merging Parties</i>		101	10.2	4.9	-0.21	103	10.3
<i>All Firms</i>		102	14.2	8.4	-0.33	109	14.4
BCBS - United	396						
<i>Merging Parties</i>		87.6	20.0	0.9	-0.13	81.1	20.0
<i>All Firms</i>		103	14.5	4.0	-0.16	105	14.0

Note: This table produces the identical results as shown in Table 7 with the alternative equilibrium assumption described in Appendix Section G. This table shows the mean effects of the merger analysis of three hypothetical mergers among the three largest firms in the Massachusetts MA market. The average Δ HHI is computed as the predicted change in HHI using pre-merger market shares, weighted by market size. The pre-merger and post-merger average values are weighted by pre-merger and post-merger enrollment, respectively. The middle group of merger effects weights the product-level merger effect by post-merger enrollment, which controls for changes in market composition due to switching.

Table A7: Merger Results Under Alternative Equilibrium Assumption

$$\tilde{\Pi}_j = \Pi_j + \psi^{prem} p_j + \psi^{copay} x_j \quad (24)$$

First, I re-solve for the equilibrium using the estimated model. There is a possibility for multiple equilibria, and it is important to be able to study merger effects separately from equilibrium selection. To address this, the equilibrium is resolved again assuming that each merging party has a 1% interest in the products of the other. Then, using this new set of premiums and copays as the starting vector, I re-solve the pre-merger baseline. Intuitively, the equilibrium solution method is to search slowly along the gradients of each strategic variable for each product until all first

However, when there are two strategic variables per product, one marginal cost residual can no longer match both observed variables. When firms are at the \$0 constraint for either variable, I assume that the first order condition is met exactly at \$0.

order conditions are met. Finally, just as in the main results of the paper, the merger is solved via homotopy in which the model is resolved for each 5% increment of interest that each merging party gains in the products of the other.

	Primary Care Copay Effect	Premium Effect	Medical Cons. Effect (\$/year)	Mortality Effect (pp)	Savings per Life (\$000s)
Tufts - BCBS					
Mean Effect	1.05	2.78	18.4	0.035	-
High Income	1.69	2.33	60.0	0.056	-
Low Income	0.37	3.26	-25.2	0.012	203
Tufts - United					
Mean Effect	-0.28	7.35	21.7	-0.010	228
High Income	-0.28	7.65	-4.6	-0.009	-
Low Income	-0.29	7.04	48.8	-0.010	501
BCBS - United					
Mean Effect	-0.43	2.75	27.4	-0.014	191
High Income	-0.57	2.76	-12.2	-0.019	-
Low Income	-0.30	2.75	68.1	-0.010	694

Note: This table produces the identical results as shown in Table 8 with the alternative equilibrium assumption described in Appendix Section G. This table displays the consumer level effects of each merger, averaged across all consumers affected by the merger, and separately for consumers in high and low income zip codes. The mortality effect is the average of the product of the primary copay effect and the full sample IV estimate in Table A4. The final column is the result of dividing the predicted change in medical consumption by the predicted change in twelve month inpatient mortality. Averages are weighted by post-merger enrollment in MA.

Table A8: Consumer-level Effects under Alternative Equilibrium Assumption

The main results are reproduced for the alternative assumption in Tables A7 and A8. In general, the qualitative results are similar. The effects on premium and primary care copays are heterogeneous across mergers, and the changes in medical consumption are focused among the low-income consumers. In this model, the effects on the primary care copay are much larger relative to premium effects, and among consumers in high income zip codes, the positive copay elasticities dominate the effects: leading to positive associations between the effect of the merger on copays and medical consumption.

While this model is better able to match observed equilibrium, the wedges

interact with merger incentives in a way that is challenging to understand theoretically. For example, a company that prefers low premiums for reasons outside the model might acquire a firm that prefers high premiums for reasons outside the model. As such, this is not the preferred model for understanding merger effects. A similar assumption is to interpret the average of the two wedges as a true adjustment to marginal cost. The results for that assumption are similar to those presented in this section.