How does Insurance Competition Affect Medical Consumption?

Conor Ryan, Pennsylvania State University*

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 in Medicare Advantage, 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. The optimal premium and primary care copay responses to a merger are inversely related for many products. As a result, mergers lead to higher prices but may reduce the copay for primary care, leading to greater medical consumption and health benefits for consumers.

JEL: I11, L13

Keywords: Insurance, Competition, Market Power, Adverse Selection

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

1 Introduction

Competition affects both product prices and product characteristics. In the context of health insurance, the effect of competition on product characteristics (e.g. copays and coinsurance rates) is particularly important. These cost-sharing terms determine the out-of-pocket price of medical care. Unlike the price (i.e. monthly premium), cost-sharing terms influence a patient's medical decisions which in turn can affect the health outcomes of the patient.

Despite the importance of the industrial organization of the insurance market in determining consumption of medical care, 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 study hypothetical mergers in a model of the Medicare Advantage (MA) market, in which insurance firms compete by selecting both the monthly premium and the copay for primary care, a key source of cost-sharing.

The directional effect of a merger on cost-sharing is theoretically ambiguous for two reasons. First, a firm's cost-sharing and premiums responses to a change in their environment, e.g. a cost shock, may be positively or negatively correlated. The optimal level of cost-sharing balances the willingness to pay among marginal consumers with the infra-marginal effect of cost-sharing on the firm's costs. The optimal trade-off between cost-sharing and premiums depends on the nature of demand and cost heterogeneity. In particular, it depends on how the willingness to pay among marginal consumers changes in response to a change in either characteristic (Spence (1975), Schmalensee (1979), Fan (2013)).

In the context of an insurance market, there is an additional complication. The effect of cost-sharing on the firm's cost is connected to willingness to pay. Consumers that are the most willing to pay for generous cost-sharing may also be those who are most responsive to cost-sharing in their medical consumption. Thus, a reduction in willingness to pay may be associated with a reduction in cost-sharing elasticity of cost. The optimal trade off will incorporate the net effect of this relationship.

Second, in the presence of selection, a merger creates incentives that can put positive or negative pressure on prices (Ryan (2025), Mahoney and Weyl (2017)). In the standard model, a price increase in one product increases the demand for rival products. A newly merged firm internalizes a new portion of this benefit, and the merger creates an incentive to raise price. In the presence of selection, a price increase also affects the cost of rival products through the risk composition of diverted consumers. If this group is adversely selected, rival products are harmed by this diversion through higher cost. A merger might lead a firm to reduce price in order to alleviate selection into the acquired product. Moreover, these incentives may be different for premiums and primary care copays due to different substitution patterns in response to each characteristic.

To capture these incentives, I specify a model in which consumers first select an insurance product and then decide how much medical care to consume based on the terms of their selected product. Insurance firms compete in Nash-Bertrand by selecting both the monthly premium and the primary care copay for a menu of differentiated insurance products. This framework allows for consumers that respond to higher cost sharing by decreasing medical consumption (moral hazard), insurance preferences that are correlated with expected cost (adverse selection), and insurance preferences that are correlated with medical consumption elasticities (selection on moral hazard). I apply the model to the MA market in Massachusetts using data on insurance plan choices linked to data on consumers' medical insurance claims.

I estimate the model in two stages. First, I estimate a discrete choice model of insurance demand (Town and Liu (2003), Tebaldi (2023), Miller et al. (2019), Vatter (2025)) with unobserved and observed heterogeneity. Correlation between consumer choices and health status measures identifies the substitution and selection patterns important in determining merger effects. In the second stage, I estimate a log-linear model of medical care demand. I exploit inertia in consumer insurance choices and year-to-year changes in the copay for primary care to estimate the elasticity of medical consumption (Abaluck et al. (2018)).

Using the estimated parameters, I simulate hypothetical mergers among the largest firms in the data. I find that the optimal premiums and primary care copays are frequently inversely correlated. Facing a marginal increase in per-unit cost, firms would respond for most products through premium increases and offsetting copay decreases. The pricing pressure created by mergers is typically positive, and strongly correlated between premiums and primary care copays. The incentive to raise primary care copays is 4% less, on average, than the incentive to raise premiums.

Due to these features of the market, mergers create stronger effects on raising premiums than on primary care copays. In many markets, the mergers lead to a decrease in the copay for primary care. For the largest mergers, the average premiums for the merging parties increase by 5% - 6%. In these same markets, the change in the average copay for primary care ranges from a 2% increase to a 4% decline.

I evaluate the effect on consumer welfare through four potential channels: total quantity, demand-implied consumer surplus, medical consumption, and inpatient mortality. For small mergers, all of these measures indicate positive effects on consumers. Enrollment and demand-implied surplus are greater after the merger. Medical consumption is also greater and accompanied by reductions in inpatient mortality. For large mergers, the welfare results are mixed. Enrollment and demand-implied surplus decline on average. In the mergers that lead to declines in the copay for primary care, the decreases in inpatient mortality and greater medical consumption provide an avenue for benefits to consumers.

This paper makes two main contributions. First, I contribute to a literature on market structure in health insurance markets by allowing for an endogenous cost-sharing, an important non-premium product characteristic.[1] Existing work studying non-premium dimensions of insurance competition has primarily focused on the provider network (Shepard (2016), Ho and Lee (2017), Dafny et al. (2018),Capps et al. (2003), Gowrisankaran et al. (2015), Tilipman (2022)) or MA payment policy (Cabral et al. (2018), Curto et al. (2021), Vatter (2025)). Work on optimal cost-sharing typically focuses on single agent settings (Marone and Sabety (2020), Ho and Lee (2023)). Horizontal competition with endogenous cost-sharing creates

[1] See Cutler and Reber (1998), Town (2001), and Dafny et al. (2012) for work on market structure in insurance. I also building on a large literature that studies the industrial organization of health insurance markets more generally (Town and Liu (2003), DeLeire et al. (2017), Drake (2019), Geruso (2017), Jaffe and Shepard (2017), Shepard (2016), Tebaldi (2023), Ericson and Starc (2015), Starc (2014), Saltzman (2021), Miller et al. (2019)).

a new channel through which insurance markets affect medical consumption and costs.

This paper also contributes to an ongoing discussion of the potential for positive welfare effects of market power in insurance markets (Mahoney and Weyl (2017), Veiga and Weyl (2016), Lester et al. (2019), Ryan (2025)). In addition to incorporating established selection dynamics in a new empirical setting, I connect competition in the insurance industry to the demand for medical care and health outcomes. With this framework, I am able to evaluate the implications of market power in the insurance industry in the context of the ultimate allocation of care. I find that mergers can lead to lower average cost-sharing among consumers, potentially increasing consumer welfare by encouraging more medical consumption and improving health outcomes.

Second, this paper adds new evidence to 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 (Dickstein et al. (2023), Marone and Sabety (2020), Einav et al. (2013), Cardon and Hendel (2001)). I find that consumers have a high willingness-to-pay for low primary care copays in their insurance decision—\$56 dollars in premium per \$10 reduction in primary care copay for the average consumer. This high willingness to pay is consistent with the elasticity of medical consumption. I find that a \$10 decrease in the copay for primary care leads to an \$80 increase in the average, monthly covered insurance spending. The implied elasticity is -0.1, which is similar to other elasticities in the literature.[2] These findings highlight the importance of primary care copays in determining both demand for insurance and overall access to medical care in Medicare Advantage.

The paper proceeds as follows. In Section 2, I elaborate on the Medicare Advantage setting. In Section 3, I present the model. In Section 4, I describe the data and show some descriptive facts that motivate the empirical exercise. In Sections 5 and 6, I outline the identification strategy and discuss the results for insurance demand estimation and the medical consumption demand estimation, respectively.

[2]See Manning et al. (1987), Duarte (2012), Brot-Goldberg et al. (2017), Aron-Dine et al. (2015), Lavetti et al. (2019), and Kowalski (2016). Chandra et al. (2014) study a policy change that includes a primary care copay increase of \$10 leads to a \$29 decrease in covered spending across prescription drugs, office visits, and hospital visits.

In Section 7, I employ these estimates in counterfactual merger simulations and discuss the key findings. In Section 8, I conclude.

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 (Bush (2002)).

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 (McGuire et al. (2011)). 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.[3] Still more mergers have been consummated that have not risen to

[3]These 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.

such high levels of antitrust concern.[4] The review of these mergers has focused on the potential premium effects of the merger. For instance, in the judicial opinion blocking the merger between Aetna and Humana, “premium” or ”price” are mentioned 202 times while ”cost sharing” or ”copayment” are mentioned only 11 times (D.D.C., 2017, United States et al. vs Aenta, Inc & Humana Inc.).

MA plans are required to offer all of the same benefits included in TM, but the vast majority of plans offer more generous cost-sharing arrangements as well as supplemental benefits, such as vision and dental.[5] In 2015, enrollees in TM were responsible for a $1,260 deductible for hospital care and a separate $147 deductible for physician services. In contrast, more than 99% of MA plans offered in 2015 have a combined hospital and physician deductible less than $147.[6]

Because MA plans rarely require a deductible, the most important source of consumer cost-sharing is service-specific copay and coinsurance terms. These terms are also typically more generous than the 20% coinsurance rate required by TM. Less than 1% of MA plans required coinsurance rates of greater than 20% for medical devices, outpatient services, or outpatient drugs. MA plans are also required to cap consumers' out of pocket expenses.

Medicare Advantage plans competitively offer less cost-sharing through restricted networks and other aspects of managed care that limit use. Descriptive evidence shows that Medicare Advantage enrollees are lower risk and spend less (Agarwal et al. (2021)), but this is challenging to separate from selection (Brown et al. (2014)).

The low cost-sharing is a result of the market design. 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. The Affordable Care Act implemented new rules for MA

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

[5] Many TM enrollees purchase supplemental insurance to cover the high cost-sharing requirements.

[6] The statistics on plan benefit terms in MA are from the author's calculations using 2015 Plan Benefit Package data.

regarding reducing payments to plans and enforcing a medical loss ratio (MLR) minimum, though the minimum is not typically binding.

3 Model

This section presents a model in which consumers first select an insurance product and then make a medical consumption decision that depends on the primary care copay of their selected product. Insurance firms set both a monthly premium and a copay for primary care for a fixed set of differentiated products in Nash-Bertrand competition. Using the model, I characterize why the effect of a potential merger on premiums and copays is ambiguous.

3.1 The Environment

Consumers

Consumers, indexed by i, face a two stage decision following Cardon and Hendel (2001) and Dubin and McFadden (1984). In the first stage, consumers select an insurance plan, j, during an annual period for open enrollment. They observe information about the products in their choice set and their own idiosyncratic preferences, ε_{ij}. In the second stage, consumers observe a realization of their medical need, ω_i and choose how much medical care, m, to consume each month 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.

$$U_{ij}^*(\omega_i) = U^*(\omega_i; p_j, x_j, W_j, Z_i) = \max_m U(m, \omega_i; p_j, x_j, W_j, Z_i) \tag{1}$$

where ω is a preference shock for medical demand, p_j is the monthly premium of the insurance plan, x_j is a copay for primary care—-which affects the out-of-pocket price of medical consumption, W_j is a vector of other insurance plan characteristics, and Z_i is a vector of consumer characteristics which may affect the distribution of

ω. The function U represents the 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, $m^*(\omega_i, x_j, W_j, Z_i)$. I assume medical consumption does not depend on the premium.[7]

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

$$v_{ijt} = V\left(\varepsilon_{ij}, \mathscr{E}_\omega\left[U^*_{ij}(\omega; p_j, x_j, W_j, Z_i)|Z_i\right]\right) \tag{2}$$

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

$$s_{ijt} = \Pr\{v_{ijt} \geq \max_k v_{ikt}\} \tag{3}$$

In this paper, I approximate $\mathscr{E}\left[U^*_{ij}(\omega_i)|Z_i\right]$ with a polynomial in p_j, x_j, W_j, and Z_i, as is common in the literature on insurance demand (Town and Liu (2003), Tebaldi (2023), Miller et al. (2019), Vatter (2025)). There are two key differences from this approach, and other methods that explicitly specify consumer expectations and nest the two stages (e.g. Marone and Sabety (2020), Dickstein et al. (2023), Ho and Lee (2023)).

First, I separately estimate s_{ijt} and m^*, rather than jointly estimating and restricting some parameters to be consistent between the two models. I make this assumption for tractability. Out-of-pocket expenses that depend on service-specific copays complicate recent techniques that jointly estimate both equations. This creates important limitations in the analysis. While the estimation allows for correlation between s_{ijt} and m^* on both observed and unobserved dimensions, I restrict selection on moral hazard to observed consumer heterogeneity.

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

This is similar to some integrated approaches (Dickstein et al. (2023)) but shuts down a possibility of unobserved correlation between moral hazard and risk aversion (Marone and Sabety (2020)).[8] The specified demand for insurance may not sufficiently capture the welfare effects of changing medical consumption. And certain theoretical findings for expected utility consumers, e.g. the relationship between moral hazard and insurance demand, may not be satisfied.

Second, I specify the indirect utility from an insurance choice as linear in x_j. Methods that directly combine the two stages typically depend on higher order terms. For example, Dickstein et al. (2023) show that the expected utility is approximately quadratic in a plan-wide coinsurance rate under constant absolute risk aversion and certain assumptions about moral hazard. A linear specification avoids challenges in using instruments to separately identify the insurance demand elasticities of higher-order cost-sharing terms.[9] Importantly, this simplification is not the same as assuming that consumers are risk neutral. While I do not separately identify risk aversion, it may still play a role in determining consumer preferences over product characteristics.[10] In order to capture similar consumer-level heterogeneity in substitution patterns, I also include a 2^{nd} order polynomial in the plan-level cost-sharing predicted by other product characteristics, W_j.

This method does allow allow for flexibility. There is evidence consumers themselves are not very good at predicting their medical expenditures nor understanding the complexities of insurance products.[11] This approach avoids using explicit functional form assumptions over the character of consumers' information or behavioral biases as a source of identification. In Section 5.3, I discuss the findings of the demand estimation interpreted through the lens of an expected-utility model.

Similar to other potential approaches, I assume that $E[\varepsilon_{ij}\omega_i|Z_i] = 0$: the id-

[8]Ho and Lee (2023) allow for unobserved variation in moral hazard that is uncorrelated with insurance demand.

[9]In a nested approach, the insurance demand elasticity of cost-sharing is identified from the elasticity of medical consumption.

[10]For example, in the two-stage model specified by Dickstein et al. (2023), risk aversion only affects a linear premium sensitivity term, which is also included in the model specified in Section 5.

[11]For work on information frictions in the choice for health insurance, see Kling et al. (2012), Brown and Jeon (2020), Handel and Kolstad (2015), Handel et al. (2019), Afendulis et al. (2015), Dalton et al. (2020), and Bhargava et al. (2017).

iosyncratic insurance preference (ε) is uncorrelated with the medical demand shock (ω) conditional on patient characteristics (Marone and Sabety (2020)). This assumption is not required for consistent estimates of consumer preference parameters, but it is necessary for solving the firms' profit maximization problem below. I allow for unobserved correlation between medical consumption demand and insurance demand through a persistent unobserved component of insurance preferences.[12]

As in Einav et al. (2013), I assume that consumer utility is a combination of health-related benefits, given by H, and non-health related benefits, G.

$$U(m, \omega_i; p_j, x_j, W_j, Z_i) = H(m, \omega_i; Z_i) + G(m; p_j, x_j, W_j, Z_i) \tag{4}$$

Health-related benefits depend only on the medical demand shock, medical consumption, and consumer characteristics. If consumers only value medical consumption through its effect on health, non-related health benefits capture out-of-pocket expenditures and other preferences related to insurance (Einav et al. (2013), Marone and Sabety (2020), Dickstein et al. (2023)). If consumers are not correct about the health benefits of their medical consumption, G may also include this bias (Baicker et al. (2015)).

Through its dependence on optimal medical consumption, the expected health benefits of each consumer is a function of all the characteristics of their chosen insurance

$$H^*(\omega_i; x_j, W_j, Z_i) = H(m^*(x_j, W_j, Z_i), \omega_i; Z_i) \tag{5}$$

A measure of health-related benefits may help provide insight into consumer welfare, but a complete summary measure of health is challenging to identify in the data. Instead, I will estimate the effect of primary care copays, x_j, on an outcome for exceptionally poor health—inpatient mortality. For more details on this specification and estimation, see Appendix Section C.

[12]For the insurance demand identification assumptions, see Section 5.2. For the medical consumption demand identification assumptions, see Section 6.2.

Firms

Insurance firms choose monthly premiums, p, and a copay for primary care, 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}.

$$\Pi_{jt} = \int_i s_{ijt}(p_{jt}, x_{jt}, p_{-jt}, x_{-jt}) \big(p_{jt} + b_{ijt}(p_{jt}, x_{jt}) - mc_{ijt}(x_{jt}) \big) di \tag{6}$$
$$p_{jt} \geq 0 \; ; \; x_{jt} \geq 0$$

where p_{-jt} and x_{-jt} represent the premium and cost-sharing terms for all other products in the market. The functions, $s_{ijt}(\cdot)$, $b_{ijt}(\cdot)$, and $m_{ijt}(\cdot)$, implicitly condition on Z_{it} and W_{jt}, and in the case of the demand function, W_{-jt}.

Firms cannot set the primary care copay 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 is $mc_{ijt}(x_{jt})$. Firms can alter their marginal cost through their choice of the primary care copay, x. The marginal cost function is consumer-product specific, i.e. it can potentially depend on consumer characteristics (Z_i), product characteristics (W_j) and their interaction. In practice, marginal costs will depend on the interaction between a consumer fixed effect and a product fixed effect that converts consumers' demand to realized cost. Details on how marginal costs are constructed are provided in Section 7.1.

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 E, I provide more detail on the formula for the risk adjusted subsidy and show how the bid function is estimated from a 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,

$$(p^*_{jt}, x^*_{jt}) = \arg\max_{(p_k, x_k)} \sum_{k \in J_{f(j)}} \Pi_{kt}(p, x, p_{-jt}, x_{-jt}) \tag{7}$$

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

3.2 Effect of a Merger

Merger effects are ambiguous for two reasons: copays balance the willingness to pay among marginal consumers with the elasticity of medical demand, and selection may flip the firms' standard incentives to raise prices due to recaptured profit. In this section, I will outline the intuition for each of these sources of ambiguity. I focus on a single period and drop period subscripts. I will write S_j to represent the average share across all consumers for product j.

The first source of ambiguity follows from the trade-off facing a firm that sets both a price (premium) and a non-price quality (primary care copay) that consumers value. The level of copays that firms will offer depends on the willingness to pay for low copays among the marginal consumers and the cost of providing it. A merger can affect this trade-off in either direction (Spence (1975), Schmalensee (1979), Hörner (2002), Matsa (2011), Fan (2013)).

Consider a simple case without selection in which the primary care copay affects demand and a constant marginal cost, $mc_j(x_j)$. The monthly premium only affects demand. Let $WTP_j(p, x)$ represent the ratio of the demand slopes with respect to each term. For an explicit setup and derivation of the conditions that follow, see Appendix Section A.1.

$$WTP_j(p,x) \equiv \frac{\frac{\partial S_j}{\partial x_j}}{\frac{\partial S_j}{\partial p_j}} \tag{8}$$

If consumers are homogeneous, this is exactly equal to their willingness to pay for a reduction in primary care copays. If there is demand heterogeneity, the marginal consumers with respect to each term may not be the same, and WTP_j represents an effective average willingness to pay for the demand facing the firm.

The optimal primary care copay is such that the reduction in marginal costs is equal to this willingness to pay.

$$WTP_j(p,x) = -\frac{\partial mc_j}{\partial x_j} \tag{9}$$

The intuition is straight-forward. An increase in the copay for primary care reduces the firms cost, but it also reduces demand. This condition says that, at the optimal copay, the reduction in cost is equal to the compensating reduction in premium required to maintain demand. If this condition did not hold, the firm could profitably raise (or lower) primary care copays with the corresponding premium reduction (or increase) to compensate consumers.

This condition highlights how a firm will respond in its optimal premiums and copays given a change in the environment. If $WTP_j(p,x)$ is unaffected by a change in premium, i.e. there is limited demand heterogeneity, then Equation (9) pins down a unique copay for primary care no matter the premium.

If $WTP_j(p,x)$ varies and is continuous in p_j and x_j, the following relationship can be derived using the implicit function theorem. This represents the relationship between endogenous premiums and primary care copays given, for example, a change in the firms per-unit cost. This condition can also characterize the relationship between optimal premiums and primary care copays, $\frac{\partial x_j^*}{\partial p_j^*}$. Consider a marginal, exogenous change in the product's per-unit cost.

$$\frac{\partial x_j^*}{\partial p_j^*} = -\frac{\frac{\partial WTP_j}{\partial p_j}}{\frac{1}{2}\frac{\partial^2 mc_j}{\partial x_j^2}} \tag{10}$$

Intuitively, if a premium increase leads to a decrease in the willingness to pay for low copays, a premium increase will be associated with a increase in the copay.[13] For example, suppose that a premium increase leads to a disproportionate loss of high-cost consumers with a high willingness to pay. The reduction in the value of copays to the marginal consumer leads the firm to increase the primary care copay as well. This, in turn, affects the degree to which the firm raises its premium.

The presence of selection creates additional complications. Most importantly, a change in the premium may not only affect $WTP_j(p,x)$, but also the elasticity of marginal cost with respect to the copay. For example, a disproportionate loss of high-cost consumers may reduce both the willingess to pay among marginal consumers and the average elasticity of medical consumption. The net effect of these two forces will determine the optimal association between premiums and copays. In Appendix Section A.2, I derive the conditions shown above in the full model.

The second source of ambiguity comes from the incentives generated by a merger in the presence of selection, which can lead to positive or negative pricing pressure (Ryan (2025)). For exposition, I assume that the non-negativity constraints are non-binding and drop the t subscripts. Let AC_j represent the average cost of selling insurance to the consumers that purchase product j, net of subsidies.

$$AC_j = \frac{1}{S_j}\int_i s_{ij}\left(mc_{ij}(x_j) - b_{ij}(p_j,x_j)\right) \tag{11}$$

Consider a merger between two single product firms, j and k. Equation 12 shows the post-merger first order condition with respect to the primary care copay, x_j.

[13] Given a convex cost function, the denominator of Equation (10) is positive.

$$\underbrace{\frac{1}{\frac{\partial s_j}{\partial x_j}}\int_i s_{ij}\left(\frac{\partial b_{ij}}{\partial x_j}-\frac{\partial mc_{ij}}{\partial x_j}\right)di+p_j}_{\text{Pre-Merger Marginal Revenue}}=\underbrace{\frac{1}{\frac{\partial s_j}{\partial x_j}}\int_i \frac{\partial s_{ij}}{\partial x_j}\left(mc_{ij}-b_{ij}\right)di}_{\text{Pre-Merger Marginal Cost}}+\text{GePP}_{jk}^{copay} \quad (12)$$

$$\text{GePP}_{jk}^{copay}=\frac{\frac{\partial S_k}{\partial p_j}}{-\frac{\partial S_j}{\partial x_j}}(p_k-AC_k)+\frac{S_k}{\frac{\partial S_j}{\partial x_j}}\frac{\partial AC_k}{\partial x_j} \quad (13)$$

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 (Jaffe and Weyl (2013)). The first term represents the standard recaptured diverted sales. This is a positive externality between firms that is newly internalized after a merger, creating unilateral incentives to raise premiums and primary care copays (Farrell and Shapiro (2010)).

The presence of selection creates the second term: the primary care copay (and premium) of product j affects the average cost of product k through the composition of consumers that choose each product. This term is negative if an increase in the primary care copay of product j leads high cost consumers to switch to product k, increasing its average cost. The newly merged firm internalizes this negative externality, creating downward pressure on primary care copays. If this term is negative and large enough, a lower primary care copay can benefit the acquired product by alleviating this selection, potentially leading to lower primary care copays (or premiums) as a result of a merger. Intuitively, the sign of the merger incentive depends on whether the most elastic consumers are also the most costly. Ryan (2025) discusses some general scenarios under which this selection term and the overall merger incentive is positive or negative.

The generalized pricing pressure for premiums and copays also need not be the same. As shown in Equation 13, the pricing pressure depends on the diverted consumers with respect to each specific term. Consumers that respond to a premium increase may have different substitution patterns than those that respond to a increase in their copay. And even if consumers agree about their substitutes, consumers that respond to one term or the other might have different compositions of risk.

A merger in an insurance market combines these two incentives. The first mechanism captures the firm's optimal trade off between premiums and primary care copays. The second mechanism captures the degree to which softening competition between two firms puts pressure on either premiums, copays, or both. In Section 7, I use the estimated demand and cost curves to evaluate these two mechanisms and show how they influence merger outcomes in the Medicare Advantage market.

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 discuss two sets of descriptive results that motivate using the primary care copay as a key endogenous cost-sharing term.

4.1 Claims 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.

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).

A novel aspect of this paper is linking medical consumption to the insurance choices of Medicare Advantage 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. For more detail on linking the APCD and PBP data, see Appendix Section B.

From these data, I construct an annual panel of insurance enrollment decisions made by Medicare beneficiaries. I include only beneficiaries that are eligible due to their age by excluding all consumers who began medicare coverage before the age of 65. These consumers are likely eligible due to a disability (1.9% of the over 65 sample).

Most consumers are enrolled in a single plan for the entire calendar year. If a consumers switches to a new product, they do so during the open enrollment period that takes place from January to March at the beginning of each year. When assigning a plan choice to a beneficiary for the year, I ignore products selected for only the first 3 months of the year. A small set of consumers (0.9%) switch products outside of this open enrollment period. For these consumers, I treat the plan with the longest enrollment as the plan choice for that particular year. If the enrollment period is equal, I assign the first plan. This abstracts from idiosyncratic special enrollment windows that some consumers may experience during the year.

I drop consumers for which I am unable to identify their enrolled product or are enrolled in plans that are not in the observed choice sets for the county in which they live (5.0%). Next, I drop consumers that are ever enrolled in employer-group plans or a special needs plan, as these consumers have difference choice sets (1.0%).[14] I exclude the two smallest counties from the analysis, which have minimal MA enrollment (0.6%).

I treat individuals over the age of 65 that are not enrolled in any MA plan as eligible to enroll but selecting traditional Medicare (TM). In order to balance the important sources of identification and the computational burden of the large data set, I include a random sample of consumers that choose TM and never switch, sampled by county. I include a 10% sample of consumers that are newly eligible in during the sample and select TM, and a 5% sample of existing incumbents

[14]Some employers continue to offer health plans to their employees or retirees that are subsidized by Medicare but effectively employer sponsored insurance.

that never switch. This dramatically reduces the computational burden without sacrificing much of the identifying variation in the data. These consumers are only important in establishing the relative value of MA and TM, which is identified by the relative size of these populations.

The data contain a limited set of non-health demographic information. I observe the consumers sex, whether they are aged 65-75 or over 75, and the zip code in which they live. Using zip code, I construct a proxy for income using the median, zip code level income. I assign consumers an income quartile based on the median income of their zip code, weighted according to the distribution of the sample across zip codes.

Consumer health status is measured in two ways. The first is through a set of binary variables that indicate whether the consumer is diagnosed with a particular disease, and the second is a summary risk score. Both of these variables are constructed using the Center for Medicaid and Medicare Services Hierarchical Conditions Categories (CMS-HCC) algorithm and risk coefficients, designed to risk-adjust for the Medicare population. The methodology can be reproduced using SAS code made publicly available by CMS. I measure health status using claims from the prior year. In most specifications, I include the 6 most prevalent disease diagnoses in this population: heart arrhythmia, vascular disease, diabetes with complications, diabetes without complications, breast or prostate cancer, and rheumatoid arthritis.

The health status for two populations requires additional assumptions. First, the medical claims of members of United Healthcare cannot be linked properly to the enrollment panel.[15] However, the distribution of health status is known, conditional on variables contained in both data. In the demand estimation and counterfactual exercises, I integrate over the distribution of each of these consumers' possible health states, conditional on the plan year, sex, zip code, and insurance product.

Second, the medical claims of traditional Medicare beneficiaries that do not enroll in a supplemental insurance plan do not appear in the APCD, and as a result, these health measures cannot be constructed.[16] In the demand estimation, I

[15]The linking ID does not function correctly for this firm.

[16]The data contain only commercial claims. Claims for traditional Medicare enrollees only ap-

assume that these enrollees have the same distribution of health status as traditional Medicare enrollees with Medigap.

I use a monthly panel of this same group of patients to estimate the model of medical consumption demand, with three additional restrictions. First, I exclude all TM beneficiaries, who have observed variation in cost-sharing and for which I observe only a limited set of claims. Second, for the reasons stated above, I must exclude United Healthcare beneficiaries, because their medical consumption cannot be directly linked to the insurance choices (14% of member months). Finally, 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 also 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 of mortality.

4.2 Descriptive Results

The first set of results show that primary care copays and premiums are lower in markets with more competition. The main analysis of this paper focuses on Massachusetts, which is covered by the claims data, but these markets exhibit relatively little variation in competition throughout the sample. The data contain no mergers and the entry of one small firm.

For evidence on the relationship between product characteristics and competition, I use national data on every county (each a local market) in the US from 2011 through 2017, which contain substantial variation in the level of competition. Appendix Table A1 shows descriptive data for all product characteristics by number of firms in the market.[17]

I follow Bresnahan and Reiss (1991) in using the market size and other market characteristics as an instrument for the number of firms that decide to enter. The

pear if they are also submitted to a supplementary, private insurer for secondary coverage. These secondary claims are present for 64% of TM beneficiaries in the data.

[17] These descriptive statistics are qualitatively similar for the 14 counties of Massachusetts between 2013 and 2017.

motivation behind the first stage of this model is that larger markets can support more firms by allowing firms to spread the fixed costs of entry over more sales. This approach has been used in the health insurance literature to show that more competitive health insurance markets have lower average premiums (Abraham et al. (2017), Dickstein et al. (2015)) and that local and national insurance plans are differentiated (Dranove et al. (2003)).

Instead of directly using the market size, I use the size of the population under the age of 65 as the excluded instrument. The intuition is similar to market size based instruments: a large non-medicare population encourages insurance firms to overcome the fixed costs of establishing a local network to participate in the employer-sponsored and non-group market segments. Through economies of scope across market segments, this should also lead more firms to participate in the local medicare advantage market. Indeed, the first stage is strong with an F-statistic of 74. This allows me to control for the size of the medicare market itself, which could affect product characteristics through risk pooling or selection in retirement.[18]

The structural equation is given by Equation (14). The dependent variable, y^s_{mt} is an enrollment weighted product characteristic s in county m and year t. The parameter of interest is the number of firms, N_{mt}. 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. I estimate Equation (14) for premiums and a selected set of cost-sharing terms.

$$y^s_{mt} = \beta^s N_{mt} + \gamma^{s\prime} X_{mt} + \varepsilon_{mt} \tag{14}$$

The estimation shows that competition has significant and negative effects on premiums and primary care copays.[19] Table 1 presents the results of this estima-

[18] An estimation that uses the size of the medicare-eligible population as an instrument arrives at quantitative and qualitatively similar results.

[19] These findings are consistent with Pelech (2018), which finds that a reduction in competition

tion for a selected set of cost-sharing parameters. I find that an additional firm decreases the average primary care copay by \$3.30, 35% of the mean value. This effect is smaller but similar in magnitude to the effect on premiums. The other cost-sharing parameters generally have small, insignificant effects. However, for some terms, more competition leads to greater cost-sharing. In Table 1 , I show that urgent care copays increase by 32% relative to their mean with an additional firm. These effects could be in response to lower revenue from other aspects of the insurance contract with more salient demand responses. Table 1 shows only a selected set of cost-sharing terms, but the reduction in the primary copay relative to the mean is large. Aside from the monthly premium, the only cost-sharing terms that exhibit greater magnitudes are either rarely used—e.g. coinsurance rates for emergency room visits—or coinsurance and copays used in conjunction. For example, competition leads to a substitution from coinsurance rates to copays for outpatient services.

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 that includes primary care doctor visits, at least once during the year. The next most frequent category of use is specialist visits, which are only used by roughly half of beneficiaries.[20]

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. The mean out-of-pocket spending on office visits is \$116, the largest of any category and constitutes roughly one quarter of all out-of-pocket spending. While these categories do not necessarily map directly to the firms' cost-sharing policy, it highlights the importance of office-based copays.

via a large-scale exit of one plan type in MA led to higher expected out-of-pocket spending by the beneficiaries.

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

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

5 Estimating Consumer Demand for Insurance

This section outlines the discrete choice model of health insurance demand. The model features a multinomial logit discrete choice with rich observed and unobserved heterogeneity, following other literature on this market (Town and Liu (2003), Miller et al. (2019), Aizawa and Kim (2018), Vatter (2025)). I am able to incorporate detailed heterogeneity on consumer health status by linking the diagnosis information in the claims data with insurance choices, as well as unobserved heterogeneity in demand for Medicare Advantage relative to Traditional Medicare. The mean estimated elasticity with respect to monthly premiums is -1.8, similar in magnitude to the mean elasticity with respect to primary care copays, -1.7. Through the observed and unobserved heterogeneity, the model can capture variation in demand with medical spending, a key feature of selection.[21] In Section 7.2, I connect this model to the model for medical consumption estimated in Section 6.1 and discuss implications for merger effects.

5.1 Specification

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 in the prior year, and an indicator for the quartile of the consumers' average risk score over the sample.[22]

[21]The variation also allows for other dimensions of preference heterogeneity, such as risk aversion.

[22]The clinical conditions included in the model are heart arrhythmia, vascular disease, diabetes with complications, diabetes without complications, breast or prostate cancer, and rheumatoid arthritis.

Consumers in the local market r and year t choose among a set of J_{rt} products. I assume the products are market-specific: $J_{rt} \cap J_{r't} = \emptyset$, $\forall r \neq r'$.[23] Products are characterized by a monthly premium p_{jt}, a primary care copay, x_{jt}, a vector of other observed product characteristics, W_{jt}, and an unobserved quality ξ_{jt}. Consumers also face a switching cost, ΥD_{ijt}, where D_{ijt} is an indicator for whether consumer i purchased product j in year $t-1$.

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

$$\delta_{jt} = \alpha_0 p_{jt} + \beta_0 x_{jt} + \gamma_0' W_{jt} + \xi_{jt} \tag{15}$$

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

$$v_{ijt} = \delta_{jt} + \Upsilon D_{ijt} + \big(\sum_g \alpha_g z_{ig}\big) p_{jt} + \big(\sum_g \beta_g z_{ig}\big) x_{jt} + \big(\sum_g \gamma_g z_{ig}\big)' W_{jt} + \varepsilon_{ijt} \tag{16}$$

where $\varepsilon_{ij\tau}$ is an i.i.d. type I extreme value idiosyncratic preference. Consumers have heterogeneous preferences over product characteristics that depend on the components of their demographics, $z_{ig} \in Z_i$. This heterogeneity captures how differences in consumer risk aversion and preferences over premiums, primary care copays, and other characteristics relate to medical risk.

The vector of product characteristics includes the copay for primary care; the monthly premium; the out-of-pocket maximum; the copays for specialists visits, outpatient procedures, inpatient stays, emergency room visits, urgent care visits, and ambulance rides; and the coinsurance rates for outpatient procedures and outpatient drugs.[24] In the demand estimation, I exclude copays for imaging services and coinsurance rates for outpatient drugs, because these do not vary across prod-

[23]The product options are local HMO/PPO plans that are targeted to single county or small set of counties. The only exception is a single, state-wide product offered by AARP.

[24]With the exception of outpatient procedures, either copays or coinsurance rates are used by all insurance firms and products for a particular service. For instance, there are no insurance plans that require a coinsurance rate for primary care.

ucts for many markets. I also allow for an interaction between risk quartile and a firm fixed effect to capture other persistent heterogeneity across firms, and I allow premium coefficients to additionally vary with respect to the income quartile of the consumers' zip code.

This specification is linear in cost-sharing terms, including the primary care copay. Dickstein et al. (2023) show that the indirect utility for consumers with rational expectations about their insurance costs approximately depends on a quadratic polynomial of the coinsurance rate of the insurance product. To incorporate similar flexibility, the vector of product characteristics also includes a second-order polynomial of an expected product-specific coinsurance rate.

For identification purposes, I exclude the primary care copay from the construction of the expected coinsurance rate. Because the expected coinsurance depends on how medical consumption and out-of-pocket expenses are affected by cost-sharing terms, it is difficult to disentangle endogenous relationships with respect to the underlying cost-sharing terms. To prevent contamination of the elasticities with respect to the primary care copay, it is excluded in the construction of the coinsurance rate. For more details on the construction of the product-specific coinsurance, see Appendix Section D.

Consumers can choose the outside option of Traditional Medicare (TM). The value of the outside option is given by

$$v_{i0t} = \sigma \nu_i + \varepsilon_{i0t} \tag{17}$$

where ν_i is a persistent unobserved preference which includes aspects of the TM choice that I cannot easily observe, such as the consumers' expected out-of-pocket expenditures and preferences over secondary coverage. I assume that ν_i is a standard normal random variable.

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.

$$s_{ijt} = \Pr\left(v_{ijt} = \max_k v_{ikt}\right) \tag{18}$$

5.2 Estimation and Identification

The parameters defining consumer demand for insurance consist of those governing consumer heterogeneity, $\theta_z = (\Upsilon, \{\alpha_g, \beta_g\}_g, \sigma)$, and the base level of product quality, δ_{jt}. The estimation proceeds in two stages, following Goolsbee and Petrin (2004). In the first stage, θ_z, σ, and δ_{jt} are estimated with maximum likelihood (Grieco et al. (2024)). I estimate these parameters separately for each market. To evaluate the likelihood function, I integrate over two distributions. I integrate over the normally distributed unobserved preference for TM using Gaussian quadrature. For consumers enrolled in United Healthcare products, I also integrate over the empirical distribution of their demographic and health characteristics.

The parameters in θ_z determine 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.

This identification treats consumer health status as exogenous to their own idiosyncratic demand and anything unobserved about the product. Recent work suggests that Medicare Advantage firms have an incentive to strategically inflate consumer risk scores in order to receive more favorable payments (Geruso and Layton (2020)). To address this, I use only very prevalent diagnoses categories that are less likely to be marginal, i.e. breast cancer, and risk quartiles, which reflect the ordering rather than the magnitude of consumer risk.

The switching cost is an important feature of demand in this market. Only 3.1% of all consumers change their products every year, though the switch rate excluding Traditional Medicare enrollees is higher—13.4%. This parameter is identified through two important assumptions: new enrollees have the same preferences as recently new enrollees, and the the switching cost is constant with age. These assumptions are required because new enrollees are the youngest beneficiaries in the market due to the features of Medicare eligibility. Because of this institutional feature, this is also the standard approach in the literature for this market (Miller et al. (2019), Aizawa and Kim (2018)).

Identification of σ is similarly identified from the choice data. Because the unobserved component of preferences v_i is persistent, the panel structure of the data identifies σ from the choice behavior of switching consumers. The frequency with which switching consumers choose other MA products rather than the outside option of TM identifies σ, the variance of idiosyncratic preferences for MA relative to TM. Appendix Table A3 shows that about half of all switching consumers choose Traditional Medicare. The switchers that stay within MA tend to choose plans with cheaper premiums and similar primary care copays. I also allow this variance to be different across counties, potentially reflecting different values of Traditional Medicare across markets.

In the second stage, the parameters that define δ_{jt}, $\theta_0 = (\alpha_0, \beta_0, \gamma_0, \{\xi_{jt}\})$, are estimated using linear instrumental variables. Standard errors in the second stage are estimated by bootstrapping draws from the estimated distribution of $\hat{\hat{\delta}}_{jt}$.

The key challenge to recovering 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.[25]

In the main specification, I use two sets of instruments.[26] The first set are "Waldfogel Instruments" (Berry and Haile (2021), Waldfogel (2003)), which consist of average demographics in each county. Conditional on a consumer's own demographics, the demand of their neighbors does not affect their demand except

[25] A key unobserved feature of product quality are provider networks, but these are typically constant over time: 96% of member months are in plans where at least 90% of reimbursements are paid to providers that are in network the following year.

[26] Other product characteristics, e.g. the copay for specialist care, are also plausibly endogenous. The Sargan J-test statistics are reported in Table 3 and small, which suggests that the instruments are not contaminated by the potentially endogenous controls, and the elasticities are consistently estimated.

through the equilibrium price. Specifically, I include the average risk score in the county and the fraction of consumers that are over the age of 75. The second set of instruments are cost-shifters. I include the annual per-member per-month administrative costs of each firm. Administrative expenses range between $11 and $283 per member-month and include expense items like claims adjudication and cost-containment efforts. An exception is Harvard Pilgrim, which entered the MA market in 2014 and spent $2,023 in administrative expenses per person that year. I include a dummy for these firm-year observations, as this likely includes fixed costs. See Appendix Section B.3 for more details on measuring administrative costs. I also report estimates from OLS and a two-way fixed effect estimation that controls for year and product-market fixed effects.

The results of the IV estimation are presented in Table 3. The instruments lead to an increase in the estimated demand elasticity for premium and primary care copay. Many cost-sharing terms are co-moving and potentially endogenous.[27] The Sargan J-test statistic for the IV estimation is small, suggesting the the instruments are not likely contaminated by potentially endogenous controls. To further test the sensitivity of these estimates, Appendix Table A2 contains specifications that include different sets of product characteristics. Specifically, I drop specialty copays—which have strong associations with movements in primary care copays—and the predicted coinsurance rate that summarizes all other characteristics, together and separately.[28] In one specification, I include only premium and copays for primary care. The coefficient on premium is stable across all of these estimations, and the coefficient on copays for primary care varies between -3.07 and -1.68, within the confidence interval of the baseline estimate. The Sargan statistics are also consistently small. The p-value for the final specification without any controls is 0.5.

[27] Appendix Table A4 shows how the levels and year-over-year changes in premiums and primary care copays are related to other product characteristics.

[28] Appendix Table A5 shows more detail on the frequency and magnitudes of year-over-year changes in product characteristics.

5.3 Results

The implied elasticities and semi-elasticities of demand are summarized in Table 4. The average premium-elasiticity of demand is fairly low, -1.8. This low premium elasticity is a product of both large switching costs and subsidized premiums that are very low. Because the premium of some plans is \$0, I also compute the average semi elasticity. A \$1 increase in the monthly premium leads to an average reduction of 2 percent in a plans market share. The median elasticities are smaller at -1 for the elasticity and -1.1 for the semi-elasticity, reflecting a long tail in the elasticities.[29] These estimates are greater but similar in magnitude to other work on this market. Vatter (2025) finds an average premium elasticity of -0.9. Miller et al. (2019), Aizawa and Kim (2018), and Curto et al. (2021) find semi-elasiticites that imply a \$1 premium increase leads to between a -0.7 percent and -1.4 percent decline in market share.

The inelastic demand in this market is due, in part, to generous subsidies, which make it optimal for firms to set premiums at a fairly inelastic section of the demand curve. These demand estimates (in conjunction with the cost estimates in the following section) imply reasonable markups that are inline with the literature. At observed premiums, I find an average markup of 11.1%. Vatter (2025) finds an average markup of 10.5%.

Consumer demand is similarly elastic with respect to the copy for primary care with an average elasticity is -1.72. The average semi-elasticity implies that a \$1 increase in the copay for primary care leads to a 10.1 percent decline in market share. A \$1 increase is large—the standard deviation across products is \$8.8—comparable to a premium increase of \$7. But \$5-10 increases in primary care copays are common in the data.

The elasticity of primary care copays are conditional on the other cost-sharing features of the plan. It is possible that consumers are using the primary care copay to infer something about the over-all generosity of the plan, but it is not a good indicator of overall generosity. Appendix Table A4 shows that projecting the primary

[29] All summary statistics in this section are for the 2017 plan-year and computed across consumers, after integrating over unobserved heterogeneity and weighting by the model-predicted choice probability.

care copay on all other plan characteristics explains about 80% of the variation in the copay for primary care. However, the associations are not always the expected direction. The primary care copay is negatively correlated with four of the ten features of cost-sharing included.

The heterogeneity in elasticities for both premiums and primary care copays comes from the demographic information used to estimate consumer preferences. In Appendix Figures A1 and A2, I show how premium and primary care copay elasticity variation depends on income, risk quartiles, disease indications, and geography. Each panel shows the median consumer-level elasticity with error bars representing the inter-quartile range. Low risk consumers and high income consumers tend to be the least elastic with respect to premiums. High income consumers are also less elastic with respect to the copay for primary care, but low risk consumers are more elastic.

The within group heterogeneity is larger than heterogeneity across groups. Within group heterogeneity comes from level-shifts in demand generated by the interactions between consumer demographics and the other set of product characteristics. Approximately a third of the parameters governing consumer-level heterogeneity are statistically significant at the 95% confidence level. For example, I find that the polynomial in predicted, product-level coinsurance is important in determining demand heterogeneity. Additionally, there is a significant amount of unobserved heterogeneity in the preference for TM relative to MA. The coefficients on the random effect range between 0.6 and 2.0 across markets, and are all statistically significant at the 0.1% level.

The switching costs are significant. The estimates imply an average monthly switching cost of $118 per month, which is about twenty percent larger than the average monthly premium in the data of $99. This is close to estimates in the literature. Aizawa and Kim (2018) estimate $342 per month and Miller et al. (2019) report switching costs are similar in magnitude to average premiums. The entering cohort faces no switching cost and makes up only 7 percent of consumers each year.

The estimates imply a high willingness to pay for low primary care copays. The average consumer would pay $56 extra dollars in monthly premiums in order to reduce the copay for primary care by $10. In an expected utility model, the willingness to pay for reduced cost-sharing can be separated into a price effect, improved risk protection, and the benefit of increased medical consumption (Marone and Sabety (2020)). The price effect should be small. While use of primary care is common, the average number of office visits is only about 8 per year. A $10 reduction in the copay saves the average consumer $6.5 per month.

The value of improved risk protection is also not likely to explain such a high willingness to pay. Using constant absolute risk aversion with a parameter of -0.0018 (Handel (2013), Einav et al. (2013), Marone and Sabety (2020)), the predicted average consumer willingness to pay to reduce *all* out of pocket spending risk is about $72 per month.

This leaves the potential benefit from increased medical consumption as the remaining rational explanation. In the following section, I estimate that a $10 increase in primary care copays leads to an 11% reduction in total insurance-covered medical spending due a $10 increase in the copay for primary care. For the average consumer, this amounts to a reduction in $80 per month of insured consumption. To the extent that consumers foresee and value this reduction, it can contribute to their high willingness to pay.

This finding should also be viewed in the context of other work that shows consumers over-emphasize salient features of cost-sharing. Abaluck and Gruber (2011) find that that seniors are willing to pay about an order of magnitude more for reductions in cost-sharing terms than can be rationalized with the expected-utility model, despite under-paying for actual reductions in expected out-of-pocket costs. Bhargava et al. (2017) find that employees at a large US firm pay an average of $31 per month to enroll in a low-deductible plan despite being a financially dominated option.

6 Estimating Elasticities of Medical Consumption

This section outlines the model for medical consumption given plan benefits and the heterogeneous health of consumers. The model follows the literature in estimating a log-linear demand equation for medical consumption (Buntin and Zaslavsky (2004), Aron-Dine et al. (2015), Ellis et al. (2017)). The elasticity of consumption with respect to primary care copays can be identified through year-to-year changes in the copay within insurance products and inertia in consumer choice. I find that, on average, the elasticity of medical consumption with respect to the primary care copay is -0.1.

6.1 Specification

The model of medical consumption, $m^*(\cdot)$, is specified as log-linear in plan characteristics, an individual fixed effect, monthly fixed effects, and an idiosyncratic medical demand error. Let $m_{i\tau}$ be the medical consumption of an individual i in month τ. Let $x_{j(i)\tau}$ be the primary care copay, and $W_{j(i)\tau}$ be a vector of other product characteristics—just as in Equations (15) and (16)—for the 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 demand is specified as

$$\log(\tilde{m}_{i\tau}) = \eta_i + \zeta_i x_{j(i)\tau} + \iota' W_{j(i)\tau} + \lambda_{g\tau} + F_{gj(i)} + \omega_{i\tau} \tag{19}$$

where λ_{gt} and $F_{gj(i)}$ are month and product fixed effects. Unless necessary, I will simplify the $j(i)$ notation to j. I allow month and firm fixed effects to vary by group, g, which I define as combination of the quartile of the consumer's average risk score over the sample and the quartile of the median income of consumer's zip code.

The log specification of medical consumption follows a literature on predicting medical expenditures and estimating elasticities (Manning et al. (1987), Aron-Dine et al. (2015), Ellis et al. (2017)).[30] I follow Ellis et al. (2017) in using $\tilde{m}_{i\tau} \equiv m_{i\tau} + \frac{1}{12}$ 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,

[30]This constitutes a parametric assumption about the shape of the demand function for medical consumption, and the estimation strategy relies on this model. Other work specifies linear demand (Marone and Sabety (2020), Einav et al. (2013)), or linear conditional on a demand shock (Dickstein et al. (2023)).

such as setting the constant at the minimum or 10^{th} percentile of positive monthly consumption.[31]

I estimate Equation (19) with four different definitions for m: total medical spending, total medical spending less out-of-pocket expenses, and two analogous outcomes for a estimated measure of medical consumption intensity. Total medical spending is a standard measure of medical service demand (Manning et al. (1987), Brot-Goldberg et al. (2017), Aron-Dine et al. (2015), Ellis et al. (2017)). But a potential concern is that the prices the firms negotiate with providers may be related to the cost-sharing features of the insurance contract. To address this, I construct a medical consumption intensity measure that values specific services equally across products and years. See Appendix Section B.2 for more details on the construction of the medical consumption intensity measure.

The elasticity important to the firm is the change in medical spending covered by insurance. I estimate Equation (19) using medical spending less out-of-pocket expenses in order to construct estimates for medical demand covered by the insurance firm. Finally, I construct an analogous medical consumption intensity net of out-of-pocket expenses by multiplying the medical consumption intensity measure by the realized fraction of expenses paid by the insurance firm for that consumer in that 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 (19) may produce biased estimates of ζ because a potential correlation between $x_{j(i)\tau}$ and $\omega_{i\tau}$.

This paper exploits within-product changes in the primary care copay to consistently identify the elasticity (Abaluck et al. (2018)). The change in the primary care copay for the product that an individual was enrolled in during the prior year is an instrument for the change in the individual's actual copay. Due to strong

[31]More than 90% of beneficiaries use some medical service during the year, and about 60% of beneficiaries have non-zero spending in any given month.

consumer inertia, the instrument is a strong predictor of consumer-level variation (Heiss et al. (2016), Ho et al. (2017), Miller et al. (2019), Drake et al. (2020)). 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.

The identifying assumption is that a change in a product's copay for primary care is orthogonal to year-to-year variation in idiosyncratic, consumer-level medical consumption variation, separate from month and consumer fixed-effects. This assumption might be violated, for instance, if a products negotiated rates change in the upcoming year. This could affect both medical spending of that product's consumers and the optimal cost-sharing terms of the product. I use the medical consumption intensity measure described above to address this issue.

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 treated observations). Each copay change affects large sets of consumers across the entire medical risk and zip code-level income distribution.

There is some co-movement in contract characteristics. While the primary care copay is not a precise prediction of the overall cost-sharing characteristics of the plans, changes in primary care copays are correlated with changes in other characteristics. Appendix Table A4 shows the relationship between changes in other cost-sharing terms and the copay for primary care. The strongest association is with the copay for specialists.

The potential endogeneity of all contract characteristics threatens the validity of the instrument, in particular if changes in the primary care copay are related to changes in other characteristics. To address this issue, I include this inertia-based instrument for all the contract characteristics that vary sufficiently within products in the data. I include out-of-pocket maximum and the coinsurance rates for outpatient drugs and medical devices as controls but do not attempt to estimate causal elasticities due to very few products exhibiting any year-to-year variation. Outpatient coinsurance is excluded entirely as there is too little year-to-year variation in the sample. Appendix Table A5 shows more detail on the frequency and magnitudes of changes in cost-sharing terms.

$$\Delta_i \log(\tilde{m}_{i\tau}) = \zeta_i \Delta_i x_{j(i)\tau} + \iota' \Delta_i W_{j(i)\tau} + \Delta_i \lambda_{g\tau} + \Delta_i F_g j(i) + \Delta_i \omega_{i\tau} \tag{20}$$

The Δ_i operator represents a 12-month, forward difference at the individual level. For example, $\Delta_i X_{copay,j(i)\tau}$ is the difference in the primary copay applicable to consumer i in month τ and month $\tau + 12$. The coefficient of interest is ζ, and the instruments for the primary copay difference is $(X_{copay,j(i,\tau)\tau} - X_{copay,j(i,\tau)\tau+12})$, the 12-month, forward difference at the product level. In words, the product-level forward difference is defined as 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. I compute clustered standard errors by product-risk-income group.

6.3 Results

The estimates for the effect of a change in cost-sharing parameters are displayed in Table 5. I find that a \$10 increase in the primary care copay leads to an 10.9% decline in total medical spending. I find a slightly larger elasticity in the measure of medical intensity, but the quantitative difference is negligible. The effect of a \$10 increase in the primary care copay on covered expenses by the insurance firm is only an 8% decline. The smaller elasticity in covered expenses seems to imply that the fraction of out-of-pocket expenses paid by the insurance firm increases as a result of greater copays for primary care. Because many cost-sharing terms are fixed copays rather than defined fractions of total spending, this is a plausible outcome of reduced medical demand. It could also reflect a changing composition of medical consumption. However, strong conclusions may not be warranted as I cannot statistically reject that the two estimates are equal.

These estimates reflect relatively small elasticities. Using the average primary care copay, \$9.4, the estimated elasticity is -0.1. This is on the inelastic end of most estimates of medical consumption elasticities in the literature, which range from

-0.1 to as high as -2 for some elective services.[32] The discrepancy likely reflects that the elasticity captures the change in all medical consumption with respect to a service-specific cost-sharing term, albeit an important one.

While it is not straightforward to separate medical consumption into primary care and non-primary care related expenses, I estimate the effect on spending on office visits and all other spending. I also re-estimate the model using an indicator whether or not a consumer receives a particular medical service in a month as a dependent variable. Results are displayed for office visits, minor outpatient procedures, major procedures (typically require general anesthesia), and hospital visits (involving a facility fee billed by a hospital).[33] These procedures do not fit cleanly into cost-sharing categories. But the likelihood that a primary care copay is involved is much greater for office visits than the other categories. Primary care physicians also occasionally bill for minor procedures, but these are also likely billed as outpatient procedures with the according cost-sharing requirement.

I find that spending falls for both office visits and non-office spending in response to an increase in the copay for primary care, but the effect is larger for office visits. Elasticities are significant for primary care visits and minor outpatient procedures but negligible for major procedures and hospital visits. This is consistent with the intuition that the primary care copay primarily affects the probability of primary care office visits, which has follow on effects for other types of elective procedures and less of an effect on more necessary care, such as major procedures. Ellis and McGuire (2007) estimate service-specific elasticities with-respect to service-specific cost-sharing terms to be between two and four times larger in outpatient and non-specialty care. This difference could reflect that primary care copays are particularly low in this market, and the older population may be less elastic in their consumption of medical care. These results are displayed in Appendix Tables A6 and A7.

The estimates imply that a $10 primary care copay decreases office visits by

[32]For elasticities with respect to general cost-sharing terms, see Manning et al. (1987), Brot-Goldberg et al. (2017), Aron-Dine et al. (2015), Lavetti et al. (2019), and Kowalski (2016). For work on service-specific elasticities, see Keeler and Rolph (1988) Duarte (2012), Ellis et al. (2017) and Chandra et al. (2010).

[33]These categories are defined using BETOS codes, as in Section 4.2.

one percentage point, about a 3.5% decline relative to the mean. Similarly, all medical use falls by about 1 percentage point, about 2% of mean medical use. The magnitudes suggest that small changes on the extensive margin are amplified in total medical spending, either through the likelihood of pursuing followup procedures or within month reductions in medical care without forgoing all care.

To test for selection on the copay changes, I also estimate the effect of copay changes on the risk score trends by using a lag first-difference in risk scores as the outcome. The effect is small and not statistically significant.

Among the other cost-sharing terms for which I have estimates, the copay for specialty care, outpatient procedures, ambulance rides, and urgent care have significant negative effects on medical consumption. I find that emergency room visit copays lead to more medical consumption. While theoretically possible, the emergency room copay only changes for one firm in the data and fails the test for selection on risk trends. Thus, I hesitate to draw any strong conclusions.

Specialty care copays behave similar to primary care copays with similar magnitudes. All of the copay terms reduce both office visit spending and non-office visit spending. However, only the copays for specialty care has a significant, albeit small, effect on major procedures. And only the copays for outpatient procedures and ambulance rides have any significant effect on hospital visits, which includes hospital billing for both inpatient and outpatient procedures, and these effects are all very small in magnitude. These results support the intuition that the first order effect of these cost-sharing terms are on the services for which they specifically apply, but these effects can spillover to other kinds of medical care.

The identification strategy relies on estimating a structural parameter from the compliers with respect to the instrument, i.e. the consumers that do not switch their product. There is a possibility that these consumer who "accept" the treatment by not switching products are different in their medical consumption behavior than those that switch. Following the intuition of the marginal treatment effects literature (Heckman et al. (1999), Cornelissen et al. (2016)), I check for heterogeneity in consumer responses using variation in the propensity of each consumer to adopt the treatment, i.e. keep their product from the previous year. Conveniently, the demand estimation in Section 5 provides an estimate of each consumers probability of switching in each period. I interact deciles of the predicted switching probability with both the change in the primary care copay and the inertia-based instrument.[34]

There is some suggestive evidence that consumers more likely to switch are more elastic in their medical consumption responses. Appendix Figure A3 shows that consumers with the median predicted switching probability (about 20%) are roughly twice as elastic as those with a negligible switching probability. However, the estimates for each decile are generally not statistically significant from one another. As a result, I do not include this type of unobserved heterogeneity in the counterfactual simulations.

Instead, I capture potential heterogeneity in medical consumption demand by allowing the elasticity with respect to the primary care copay to depend, separately, on the quartile of the consumer's average risk score over the sample and the quartile of the median income of the consumer's zip code. I estimate this heterogeneity by interacting the inertia-based instrument with the consumer-level information on risk and zip code. This is motivated by existing work that finds consumption elasticities and potential health effects are concentrated among higher risk and lower income seniors (Chandra et al. (2010), Chandra et al. (2014)).

Elasticities are increasing in the consumers average risk and decreasing in the income of the consumers zip code. I plot the estimates for insurance-covered medical consumption in Figure 1. Each point represents the implied point estimates for each consumer group, and the error bars represent a 95% confidence interval. For consumers in higher income zip codes, the effect of a change in the primary care copay is not statistically different from zero. The strongest effects are concentrated among low income zip codes and high risk consumers. These comparisons are robust to the particular measure of medical consumption. The coefficients of this estimation and estimations of heterogeneity in the other outcome measures are presented in Appendix Table A8.[35]

In the counterfactual simulations, the objective is to capture the elasticity of covered expenses by the insurance firm without the contamination of insurance firm

[34]Because of the panel difference estimation method, I use each consumer's average switching propensity in the structural equation, and the period-specific switching propensity in the first stage.

[35]The estimates for observed heterogeneity are robust to allowing for unobserved heterogeneity through the switching probability decile.

specific negotiated prices. To do so, I will use this final specification with consumer-level heterogeneity and the covered medical consumption outcome. The first order conditions of the firms convert the predicted covered medical consumption from this model into expected expenditure. I provide details on this procedure in Section 7.

6.4 Unobserved Correlation Between Medical and Insurance Demand

The selection dynamics in this market depend on the relationship between the demand for insurance and the demand for medical consumption. The primary connection between these two models is observed heterogeneity. The elasticity of medical consumption is correlated with demand elasticities through medical risk and zip code level income. The persistent level of medical consumption demand, determined by the fixed effect η_i, may be freely correlated with insurance demand elasticities through the relationship between the estimated fixed effects and the dimensions of observable heterogeneity included in the insurance demand model.

In addition to this observable heterogeneity, I also allow η_i to be correlated with the unobserved preference for Traditional Medicare, ν_i. This allows me to capture selection into Medicare Advantage based on health status. I specify $\eta_i = \tilde{\eta}_i + \varsigma \nu_i$. Using the estimated insurance demand model, I compute the expected value of ν_i for every consumer given their choices during the sample period. I estimate ς by regressing the recovered medical consumption demand fixed effects on the expected value of ν_i given the choice panel of the consumer. Intuitively, consumers that enroll in MA despite demographic and health characteristics that suggest they should prefer TM have higher expected values of ν_i, and this is especially the case if they stay enrolled in MA despite switching plans.

The estimates imply advantageous selection into Medicare Advantage, with $\hat{\varsigma} = 0.119$, with a standard error of 0.028. This positive correlation between medical consumption demand and the value of traditional medicare is consistent with findings in the literature that Medicare Advantage tends to attract slightly healthier beneficiaries, conditional on observed health status (Brown et al. (2014)). A consumer with a standard deviation greater demand for MA is 12% less costly to insure.

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). The summary statistics for all seven firms that operate in the state are displayed in Table A9.[36]

7.1 Solving Equilibrium

I assume that firms set two strategic variables for each product, the copay for primary care and the monthly premium. These two first-order conditions allow me to identify two dimensions of marginal cost. Following the specification in Section 6, I specify the marginal costs for every potential consumer-product pair as follows:

$$mc_{ij}(x_j) = \kappa_{1j} e^{\varsigma \nu_i + \hat{\zeta}_i x_j + \hat{\eta}_i} + \kappa_{2j}$$

Persistent consumer heterogeneity is summarized by unobserved heterogeneity is captured through the relationship between cost and idiosyncratic preference for TM ($\varsigma \nu_i$), the effect of primary copays (x_j) on costs is given by ζ_i, and an individual fixed effect (η_i). For these parameters, I use the unobserved demand correlation estimated in Section 6.4, the estimated elasticity terms from the final column of Table 5, and the estimated fixed effects. These results represent the estimation of an index of medical consumption with firm-specific prices stripped out and scaled by consumer out-of-pocket contributions. The measure captures the quantity of medical consumption that the insurance firm expects to cover.

[36] AARP offers a Medicare Advantage plan in a partnership with United, which I treat as a separate firm.

In order to convert the estimates of medical consumption demand to cost, I must recover κ_{1j} and κ_{2j}. The first term, κ_{1j}, re-scales the quantity index to firm expenditures, and captures the effects of all other product characteristics, the cost of non-financial characteristics like the provider network and care-management, and the expectation of idiosyncratic medical demand shocks, $E[\exp(\omega_{ij})]$. Importantly, this requires that the demand shocks are independent of the idiosyncratic insurance preferences, $E[\varepsilon_{ij}\omega_{ij}] = 0$. The second term, κ_{2j}, allows for non-risk related marginal costs, including administrative expenses and the expected cost of providing a prescription drug benefit.

I fit these two unknowns about each products marginal costs using the two first order conditions for the primary care copay and the monthly premium. I illustrate how these two terms are identified using a single product firm. The first order conditions imply that marginal revenue is a linear equation of a term summarizing the persistent cost components of the marginal consumers.

$$p_j + \frac{\int_i \left(s_{ij} \left(\frac{\partial b_{ij}}{\partial x_j} - \frac{\partial mc_{ij}}{\partial x_j} \right) + \frac{\partial s_{ij}}{\partial x_j} b_{ij} \right) di}{\frac{\partial s_j}{\partial x_j}} = \kappa_{1j} \frac{\int_i \frac{\partial s_{ij}}{\partial x_j} \left(e^{\hat{\eta}_i + \hat{\zeta}_i x_j} \right) di}{\frac{\partial s_j}{\partial x_j}} + \kappa_{2j} \tag{21}$$

$$\underbrace{p_j + \frac{\int_i \left(s_{ij} \left(\frac{\partial b_{ij}}{\partial p_j} \right) + \frac{\partial s_{ij}}{\partial p_j} b_{ij} \right) di}{\frac{\partial s_j}{\partial p_j}}}_{\text{Estimated Marginal Revenue}} = \kappa_{1j} \frac{\int_i \frac{\partial s_{ij}}{\partial p_j} \left(e^{\hat{\eta}_i + \hat{\zeta}_i x_j} \right) di}{\frac{\partial s_j}{\partial p_j}} + \kappa_{2j} \tag{22}$$

The differences of Equations (21) and (22) identify κ_{1j}, and the remaining residual identifies κ_{2j}.

The implementation of this procedure is complicated by products that set premiums equal to zero. Because Equation (22) may not hold with equality, the two parameters are not identified. To accommodate this possibility, I fit κ using a two step procedure. In the first step, I assume that all first order conditions hold with equality. For a product j' with a premium of zero, I assign a κ_{1j} value equal to the minimum of the implied $\kappa_{1j'}$ and the average of all other products in the market with $\kappa_{1j} \leq \kappa_{1j'}$. Intuitively, I assume this product is similar to the average among the other low-cost products in the market. In the second step, I re-compute the κ_{1j} values for the unconstrained products holding fixed the constrained products. The κ_{2j} values are then implied in a straight-forward manner. I apply the same methodology for primary care copays set to be zero. There are no products with neither a premium nor a primary care copay.

While the first order conditions hold (or meet an inequality) at the observed data, the data do not necessarily constitute an equilibrium if the profit functions are not concave at these premiums and primary care copays. In order to locate a stable equilibrium, I add a small amount of noise into the premiums and primary care copays and re-solve the equilibrium using iterated best responses. The resulting equilibria are close to the observed data. In Appendix Figure A4, I plot the observed premiums and primary care copays relative to the baseline of the model. The solved equilibria fit premium variation very well. The equilibria slightly under-predict the prevalence of higher copays and over-predict \$0 copays.

Further, because of adverse selection, the model may have multiple equilibria. In order to ensure that I isolate the effects of a merger outlined in Section 3.2 rather than equilibrium selection, I solve for the post-merger equilibrium using a homotopy method (Eaves and Schmedders (1999)). A homotopy method is one in which a hard problem can be solved (locating the post-merger equilibrium) by starting with an easy problem (locating the pre-merger equilibrium) and following a continuous transformation between the easy problem and the hard problem. In this case, the continuous transformation is a continuously consummated merger in which each merging firm continuously gains share of the other. In practice, I incrementally solve small mergers in which each merging firm gains a 5% interest in the profit of the other party. While this doesn't rule out the possibility of multiple equilibria, it holds fixed an equilibrium selection mechanism across the mergers and markets studied and allows me to separate merger effects from comparing possible equilibria.

I solve each merger market-by-market. This assumes that firms can freely adjust their premiums and primary care copays across each market. Premiums for most plans are market-specific in the pre-merger equilibrium. All the merging firms offer different cost-sharing terms to different sets of markets, but a single set of cost-sharing terms is typically offered to more than one market. The extent of geographic segmentation suggests that the cost to do so is not prohibitively high, but the simulated equilibria may overestimate cross-market heterogeneity if firms are restricted to zone strategies for their primary care copays.

To the extent that firms have less than perfect information about consumer health and preference distributions, the estimated marginal costs will reflect the firm strategies and information rather than true costs. The marginal cost estimates, as well as the equilibrium solution relies on the assumption that firms are static, which is common in the literature (Aizawa and Kim (2018), Vatter (2025)). If firms are behaving dynamically, the fixed marginal cost estimates may include invest and harvest incentives for future revenue that is internalized by the firm (Anran Li (2025)). The counterfactual exercise that follows will assume these parameters are fixed, and be unable to incorporate any way in which a merger alters these incentives.

In the merger simulation, firms endogenously respond in their premiums and their primary care copays, but I assume that firms hold fixed all other product features, including other cost-sharing parameters, such as specialist copays, that also have significant effects on medical consumption. I also implicitly hold fixed any non-financial aspects of the plan, such as their physician network or prior-authorization, both of which are potential avenues through which insurance firms can balance their demand and cost. This is a limitation due to both identifying insurance demand elasticities and solving the equilibrium. The endogenous primary care copays give firms an avenue to adjust both their demand and their cost through the equilibrium medical consumption of their beneficiaries, but this trade off may be different when firms have access to their full-suite of tools.

7.2 Characterizing Selection

In Figure 2, I combine the implied marginal cost estimates with demand estimates to demonstrate key features of selection in this market.

Consumer demand features adverse selection on primary care copays. Consumers with higher expected costs have a higher willingness to pay for reducing their copays for primary care. However, this relationship is complicated by the risk-adjusted subsidies. Because higher cost consumers also have higher variance in their spending, there are opportunities for both adverse and advantageous selection relative to the risk-adjusted payments, leading to a U-shape in willingness to pay and net-cost (Brown et al. (2014)).

Despite adverse selection on primary care copays, the most elastic consumers are the most costly. This is primarily due to advantageous selection into the MA program: healthier and younger consumers prefer Medicare Advantage, while Traditional Medicare is a popular option for consumers with higher expected medical spending. This elasticity relationship is especially pronounced for premium, and because the advantageous selection is not fully risk adjusted, the relationship is robust both with and without the risk-adjusted payments.

These features are important for the effects of competition and mergers in this market. The presence of advantageous selection creates downward pressure on the premium effect of a merger, as reductions in quantity are also associated with reductions in marginal cost (Mahoney and Weyl (2017)). Moreover, the interaction between these two kinds of selection highlight the importance of the particular substitution patterns between any two products when evaluating a merger (Ryan (2025)).

There are two key mechanisms that drive the results of a merger in this market, discussed in detail in Section 3.2. The first mechanism concerns the relationship between the willingness to pay of marginal consumers and the elasticity of medical consumption to primary care copays. Figure 2 suggests that the most elastic consumers also have the highest willingness to pay. If a merger raises premiums, this may shift the margin to consumers with less willingness to pay, putting upward pressure on primary care copays as well. However, these high risk consumers are also the most elastic in their medical consumption. If the composition of enrollment shifts towards healthier consumers, the reduced in cost reduction effect of copays may lead to lower optimal copays. The net effect of these two forces will determine the trade-off for firms.

In Panel (a) of Figure 3, I plot relative pass-through rate of primary copays and premiums. This reflects the firms optimal trade-off between primary care copays and premiums in response to a change in its cost. Jaffe and Weyl (2013) show that the relevant pass-through rate for the incentives generated by a merger is derived using the profit of the merged firm, evaluated at the pre-merger equilibrium.[37] I compute this relative pass-through rate for each of the three, hypothetical merged firms.

Optimal premiums and primary copays are often negatively associated. For example, a value of -0.1 indicates that a \$1 increase in the premium is a associated with a \$0.1 decrease in the primary care copay in firms' best responses. Even in the face of upward pricing pressure on premiums, this mechanism will lead many firms to reduce primary care copays as a result of a merger. In each panel, there is a mass of products on the right with an infinitely high pass-through ratio. This reflects the group of products for which the premium is bounded at zero. In this case, any upward pricing pressure will be fully passed through via primary care copays up to the point that the lower bound constraint on premiums no longer binds.

The second mechanism is pricing pressure due to recaptured sales, which may be positive or negative. Pricing pressure might also be unequal with respect to premiums and primary care copays if two products are closer substitutes in one characteristic than another. I find that this is not the case.

Panel (b) of Figure 3, plots the GePP values across all products, markets, and mergers for premiums and primary care copays. Most of the points are close to the 45 degree line, indicating that substitution patterns between products are similar when changing either premiums or copays. On average, the GePP for primary care copays is 4% less than the GePP for premiums, reinforcing a focus on increasing premiums rather than primary care copays as a result of a merger.

Pricing pressure is occasionally negative. For 6% percent of affected products, the merger creates an incentive to reduce price for either the premium or the primary care copay due to a negative externality on the costs of merging products.

[37] I show how their method applies in the setting of this model in Appendix Section A.3.

7.3 Results

Tables 6 and 7 display the average effects of each merger on premiums and copays. I separate the effects into 11 merger-markets with a change in the Herfindahl–Hirschman index of concentration (HHI) of less than 750 (Table 6) and 13 merger-markets with a change in HHI of greater than 750 (Table 7).[38] The merger effects are decomposed into a "within product" effect shows the average product-level premium and copay changes weighted by pre-merger enrollment and a "sorting within market" effect due to new consumer choices across products within the market. I weight effects across markets according to the total market size.

Mergers across all categories lead to increases in the product-level premiums. But mergers in all categories but one lead to decreases in the average copay for primary care.[39] This leads to potentially ambiguous welfare effects of the mergers.

The markets in which the mergers lead to large changes in concentration have correspondingly larger effects. The effect on the average premiums of the merging firms exceeds 5% in each case. But importantly, the mergers in which the merging firms increase their premiums by the most are also those in which the primary care copay declines.

In Figure 4, I explore heterogeneity in merger effects through the lens of the mechanisms discussed in Section 3.2. Each point represents the average effect of the merging firms' products for a particular merger in a particular market. The size of each point is proportional to the enrollment in the merging firms' products pre-merger.

Panels (a) and (b) of Figure 4 plot the premium and copay effects with respect to the size of the merger. Mergers with larger changes in HHI lead to much greater pricing pressure, and accordingly lead to greater increases in premiums. However, the change in HHI is negative correlated with changes in the average primary care copay.

Panels (c) and (d) of Figure 4 plot the premium and copay effects with respect

[38]The mergers are classified according to the change in HHI implied by pre-merger market shares.

[39]The comparison of average effects are indicative of the market-level effects as well. The primary care copay increases in all three large Tufts-United mergers. Small primary care copay increases are present in one market for each of the small merger categories.

to the average copay/premium trade off of the merging firms. The decreases in primary care copays come predominantly from markets and firms for which this trade off is near zero or negative. This relationship also generates the negative relationship between HHI and copays for primary care. For many large mergers, large premium increases and combined with large copay decreases.

For each merger, I compute four potential measures related to consumer surplus. The first is enrollment in any Medicare Advantage plan. The second is the consumer surplus implied by the insurance demand model estimated in Section 5. The demand-implied surplus will only be reliable to the extent that consumers are well-informed about their health, the insurance products, and their future medical consumption when making their insurance choice. Third, I compute the expected annual change in medical consumption spending. I convert percent changes in medical demand to dollars using average medical spending in each market. Finally, I compute the expected change in inpatient mortality using the model estimated in Appendix Section C and the average, market-level change in the primary care copay.

Demand-based welfare measures suggest that consumers are better off in many cases. In markets where the mergers have less of an impact, enrollment in Medicare Advantage increases on average as a result of the merger, and demand-implied consumer surplus is also greater—though small in magnitude. This is primarily a result of how the changing composition of consumers enrolled in non-merging firms leads to a more appealing set of premiums and primary care copays. The total demand for the merging products falls as they extract more profit, but the responses of the rival firms more than compensate consumers. These measures also suggest a consumer benefit in a Tufts-BCBS merger, even in the markets where these firms have large market share. In the other two cases, the markets more greatly affected by the merger show consumer surplus loses.

In mergers where the copay falls, medical consumption increases and inpatient mortality falls. This provides another potential avenue for consumer benefits. Merger effects on medical consumption are of a similar order of magnitude to the average premium for a single month. This is the case even in some markets where the mergers and premium effects are small.

The average effects disguise important heterogeneity among consumers. In Figure 5, I plot the medical consumption effect for every merger, divided by the risk quartile of consumers and the income quartile of their zip code. Each point represents a merger, and its size represents the number of affected consumers within the risk quartile. The gray lines connect merger effects across consumer groups.

The changes in medical consumption are strongly concentrated among the highest risk consumers, who either increase or decrease their consumption substantially as a result of the merger. These consumers reach as high as \$900 in additional spending or as low as a \$300 reduction in spending. Medical consumption effects are similarly concentrated among lower income consumers, although to a lesser extent. Consumers in high income zip codes still have significant responses, driven by the higher risk consumers within those zip codes.

The welfare effect of changes in medical consumption is not straight-forward. A merger that lowers primary care copays and increases total spending on medical care creates a potential inefficiency by increasing the resource cost of medical care. If the additional spending comes from wasteful, unnecessary care in which the benefit to the patient is less than the cost (i.e. moral hazard consumption), then the spending constitutes a net welfare loss. However, if the reduced primary care copays encourage effective medical consumption, the additional cost may be outweighed by the additional well-being of the patient.

To address this question, I use the estimated average effect of primary care copays on inpatient mortality. I find that a \$10 increase in the copay for primary care leads to a 0.3 percentage point increase in 12-month inpatient mortality risk. This finding suggests that additional medical consumption through a reduction in the primary care copay confers a health benefit to patients. For details on these estimates, see Appendix Section C.

The largest effect on medical consumption comes from the large-market Tufts-BCBS mergers, in which copays fall by 5.5% and medical consumption increases by \$106 per year. The mortality estimates imply that this decrease in copays for primary care would decrease the expected inpatient mortality rate by 2.5 basis points. This amounts to an average cost per averted inpatient death of \$424,000. The other results imply similar costs per averted inpatient death. These estimates are much less than estimates of the value of a statistical life, which tend to range between $1 and $10 million for consumers in this age range (Aldy and Viscusi (2007)). This would suggest that increases in medical spending are more than offset by their corresponding welfare benefit through improved health.

There are several important caveats to these findings. First, I find that changes in medical consumption are concentrated among consumers with the highest medical risk, but the estimate of the effect of copays on inpatient mortality are from the whole sample. To the extent that the marginal mortality are consumers suffering from serious medical conditions, a quality-adjusted value of a statistical life may be quite a bit lower.

Second, this analysis focuses on the Medicare Advantage market, but the merger also shifts consumers to (or from) Traditional Medicare. Duggan et al. (2018) estimate that shifting a consumer from Medicare Advantage to Traditional Medicare leads to 35% increase in spending without any associated benefit in health or quality.[40] Mergers that reduce enrollment in Medicare Advantage will carry this additional welfare cost of more costly care. Their estimates imply that a 3% decline in MA enrollment leads to roughly a $80 increase in the annual, per-person medical spending of the pre-merger set of MA enrollees. This mechanism would offset the small gains from the large markets affected by the BCBS - United merger and reinforce the losses of the large markets affected by the Tufts-United merger. For markets in which the mergers result in increases in the Medicare Advantage enrollment, this mechanism represents another potential benefit through a spillover reduction in costs for the TM enrollees drawn in by the merger.

Third, the magnitude of the merger effect is small relative to other potential policy changes. For instance, Chandra et al. (2014) study an increase in cost-sharing among a group of Medicare-age retirees in California that increased office visit copays and some prescription drug copays by $10. This policy change led to a 15% reduction in office visit spending, and a 32% decrease in prescription drug spending. Even in the mergers with the largest effects on medical spending, it

[40] Specifically, Duggan et al. (2018) find that consumers shifted in to TM by plan exits in New York lead to a 52% increase in inpatient charges. They assume that these apply equally across all spending categories and apply a two-thirds ratio of costs-to-charges to arrive at 35%.

amounts to roughly a 1% change in average spending.

And finally, the results of this model may not apply directly to the other large insurance markets in the US. The price effects of mergers in Medicare Advantage are mitigated by the presence of Traditional Medicare, a popular alternative that attracts many high cost consumers. Moreover, the primary care copay is not as prominent a feature of insurance contracts in the employer sponsored and non-group markets. In these markets, deductibles and plan-wide coinsurance rates are common (Ho and Lee (2023), Dickstein et al. (2023)). The substitution and medical consumption behavior with respect to these other cost-sharing terms may be different.

Taken together, these results suggest that mergers in Medicare Advantage 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.

8 Conclusion

Insurance firms compete on both the premiums and cost-sharing characteristics of their products. Competing on both price and quality is a common feature of competitive markets, but cost-sharing has an additional importance in determining the allocation of medical care. By setting the cost-sharing terms of insurance, competition in the insurance industry has an effect on medical consumption and patient health. To begin to shed light on this channel, I study the relationship between competition and the optimal primary care copays of firms in Medicare Advantage.

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 primary care copays on patient health. These estimates, together with a model of the firm, imply that the optimal premiums and primary care copays are often inversely correlated. As a result, mergers that lead to premium increases are often associated with decreases in the copay for primary care.

In the mergers with the largest effects on the primary care copay, average medical spending increases by \$106 per person per year and the likelihood of an inpatient death in a twelve month period decreases by 0.025 percentage points. This provides an important channel for a potential benefit to consumers as a result of a merger.

Other ways in which insurance firms compete includes the design of the hospital and physician network (Capps et al. (2003), Shepard (2016), Ho and Lee (2017)), the design of drug formularies, the use of "gate-keepers", and the use of non-financial ways to allocate medical care such as prior authorization requirements. Work that incorporates the relationship between competition and 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, and Paul Grieco, as well as all the participants of the workshops and seminars where the paper has been presented.

Declaration of Interest

I declare that I have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this article.

Data Availability Statement

Raw data were generated by the Massachusetts Center for Health Information and Analysis and are available for a fee through an application process. The codes used to generate the results of this paper from the raw data are available from the author upon request.

References

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

_ , _ , and Ashley Swanson, "Prescription drug use under Medicare Part D: A linear model of nonlinear budget sets," *Journal of Public Economics*, aug 2018, *164*, 106–138.

_ , Peter Hull, and Amanda Starc, "Mortality Effects and Choice Across Private Health Insurance Plans," *National Bureau of Economic Research*, jul 2020.

Abraham, Jean, Coleman Drake, Jeffrey S Mccullough, · Kosali Simon, B Jean, and Marie Abraham, "What drives insurer participation and premiums in the Federally-Facilitated Marketplace?," *International Journal of Health Economics and Management*, 2017, *17*, 395–412.

Afendulis, Christopher C., Anna D. Sinaiko, and Richard G. Frank, "Dominated choices and Medicare Advantage enrollment," *Journal of Economic Behavior and Organization*, 2015, *119*, 72–83.

Agarwal, Rajender, John Connolly, Shweta Gupta, and Amol S. Navathe, "Comparing Medicare Advantage And Traditional Medicare: A Systematic Review," *https://doi.org/10.1377/hlthaff.2020.02149*, jun 2021, *40* (6), 937–944.

Aizawa, Naoki and You Suk Kim, "Advertising and Risk Selection in Health Insurance Markets," *American Economic Review*, mar 2018, *108* (3), 828–867.

Aldy, Joseph E and W Kip Viscusi, "Age Differences in the Value of Statistical Life: Revealed Preference Evidence," *Review of Environmental Economics and Policy*, 2007, *1* (2), 241–260.

Anran Li, "Commitment, Competition, and Preventive Care Provision," 2025.

Aron-Dine, Aviva, Liran Einav, Amy Finkelstein, and Mark Cullen, "Moral Hazard in Health Insurance: Do Dynamic Incentives Matter?," *The Review of Economics and Statistics*, 2015, *97* (4), 725–741.

Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein, "BEHAVIORAL HAZARD IN HEALTH INSURANCE," *Quarterly Journal of Economics*, 2015, *130* (4), 1623–1667.

Berry, Steven T. and Philip A. Haile, “Foundations of demand estimation,” *Handbook of Industrial Organization*, jan 2021, *4* (1), 1–62.

Bhargava, Saurabh, George Loewenstein, and Justin Sydnor, “CHOOSE TO LOSE: HEALTH PLAN CHOICES FROM A MENU WITH DOMINATED OPTIONS,” *The Quarterly Journal of Economics*, 2017, *132* (3), 1319–1372.

Bresnahan, Timothy F and Peter C Reiss, “Entry and Competition in Concentrated Markets,” *Journal of Political Economy*, oct 1991, *99* (5), 977–1009.

Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad, “What does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics,” *The Quarterly Journal of Economics*, 2017, *132* (3), 1261–1318.

Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston, “How does risk selection respond to risk adjustment? New evidence from the Medicare Advantage Program,” *American Economic Review*, 2014, *104* (10), 3335–3364.

Brown, Zach Y and Jihye Jeon, “Endogenous Information and Simplifying Insurance Choice *,” 2020.

Buntin, Melinda Beeuwkes and Alan M. Zaslavsky, “Too much ado about two-part models and transformation? Comparing methods of modeling Medicare expenditures,” *Journal of Health Economics*, 2004, *23*, 525–542.

Bush, George W., “President’s Radio Address,” 2002.

Cabral, Marika, Michael Geruso, and Neale Mahoney, “Do Larger Health Insurance Subsidies Benefit Patients or Producers? Evidence from Medicare Advantage,” *American Economic Review*, 2018, *108* (8), 2048–2087.

Capps, Cory, David Dranove, and Mark Satterthwaite, “Competition and market power in option demand markets,” *RAND Journal of Economics*, 2003, *34* (4), 737–763.

Cardon, James H. and Igal Hendel, “Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey,” *The RAND Journal of Economics*, 2001, *32* (3), 408.

Chandra, Amitabh, Jonathan Gruber, and Robin McKnight, “Patient cost-sharing and hospitalization offsets in the elderly,” *American Economic Review*, mar 2010, *100* (1), 193–213.

— , — , and — , "The impact of patient cost-sharing on low-income populations: Evidence from Massachusetts," *Journal of Health Economics*, jan 2014, *33* (1), 57–66.

Cornelissen, Thomas, Christian Dustmann, Anna Raute, and Uta Schönberg, "From LATE to MTE: Alternative methods for the evaluation of policy interventions," *Labour Economics*, aug 2016, *41*, 47–60.

Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya, "Can Health Insurance Competition Work? Evidence from Medicare Advantage," *Journal of Political Economy*, 2021, *forthcomin.*

Cutler, D. M. and S. J. Reber, "Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection," *The Quarterly Journal of Economics*, may 1998, *113* (2), 433–466.

Dafny, Leemore, Kate Ho, and Robin S Lee, "The Price Effects of Cross-Market Mergers: Theory and Evidence from the Hospital Industry," *RAND Journal of Economics*, 2018.

— , Mark Duggan, and Subramaniam Ramanarayanan, "Paying a Premium on Your Premium? Consolidation in the US Health Insurance Industry," *American Economic Review*, 2012, *102* (2), 1161–1185.

Dalton, Christina M., Gautam Gowrisankaran, and Robert J. Town, "Salience, myopia, and complex dynamic incentives: Evidence from Medicare Part D," *Review of Economic Studies*, mar 2020, *87* (2), 822–869.

D.D.C., "United States of America, et al., v. Aetna Inc., et al.," 2017.

DeLeire, Thomas, Andre Chappel, Kenneth Finegold, and Emily Gee, "Do individuals respond to cost-sharing subsidies in their selections of marketplace health insurance plans?," *Journal of Health Economics*, dec 2017, *56*, 71–86.

Dickstein, Michael J., Mark Duggan, Joe Orsini, and Pietro Tebaldi, "The impact of market size and composition on health insurance premiums: Evidence from the first year of the affordable care act," in "American Economic Review," Vol. 105 American Economic Association may 2015, pp. 120–125.

Dickstein, Michael, Kate Ho, and Nathaniel Mark, "Market Segmentation and Competition in Health Insurance," *Journal of Political Economy*, 2023.

Drake, Coleman, “What Are Consumers Willing to Pay for a Broad Network Health Plan?: Evidence from Covered California,” *Journal of Health Economics*, 2019, *65*, 63–77.

__ , **Conor Ryan, and Bryan Dowd**, “Sources of Inertia in Health Plan Choice in the Individual Health Insurance Market,” *Working Paper*, 2020.

Dranove, David, Anne Gron, and Michael J. Mazzeo, “Differentiation and Competition in HMO Markets,” *Journal of Industrial Economics*, dec 2003, *51* (4), 433–454.

Duarte, Fabian, “Price elasticity of expenditure across health care services,” *Journal of Health Economics*, 2012, *31*, 824–841.

Dubin, Jeffrey A. and Daniel L. McFadden, “An Econometric Analysis of Residential Electric Appliance Holdings and Consumption,” *Econometrica*, mar 1984, *52* (2), 345.

Duggan, Mark, Jonathan Gruber, and Boris Vabson, “The Consequences of Health Care Privatization: Evidence from Medicare Advantage Exits,” *American Economic Journal: Economic Policy*, feb 2018, *10* (1), 153–86.

Eaves, B. Curtis and Karl Schmedders, “General equilibrium models and homotopy methods,” *Journal of Economic Dynamics and Control*, sep 1999, *23* (9-10), 1249–1279.

Einav, Liran, Amy Finkelstein, Stephen P Ryan, Paul Schrimpf, and Mark R Cullen, “Selection on Moral Hazard in Health Insurance,” *American Economic Review*, feb 2013, *103* (1), 178–219.

Ellis, Randall P. and Thomas G. McGuire, “Predictability and predictiveness in health care spending,” *Journal of Health Economics*, jan 2007, *26* (1), 25–48.

__ , **Bruno Martins, and Wenjia Zhu**, “Health care demand elasticities by type of service,” *Journal of Health Economics*, sep 2017, *55*, 232–243.

Ericson, Keith M Marzilli and Amanda Starc, “Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange,” *Review of Economics and Statistics*, 2015, *97* (3), 667–682.

Fan, Ying, “Ownership consolidation and product characteristics: A study of the US daily newspaper market,” *American Economic Review*, 2013, *103* (5), 1598–1628.

Farrell, Joseph and Carl Shapiro, "Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition," *The B.E. Journal of Theoretical Economics*, 2010, *10.*

Geruso, Michael, "Demand heterogeneity in insurance markets: Implications for equity and efficiency," *Quantitative Economics*, 2017, *8*, 929–975.

— **and Timothy Layton**, "Upcoding: Evidence from Medicare on Squishy Risk Adjustment," *https://doi.org/10.1086/704756*, jan 2020, *128* (3), 984–1026.

Goolsbee, Austan and Amil Petrin, "The Consumer Gains from Direct Broadcast Satellites and the Competition with Cable TV," *Econometrica*, mar 2004, *72* (2), 351–381.

Gowrisankaran, Gautam, Aviv Nevo, and Robert Town, "Mergers when prices are negotiated: Evidence from the hospital industry," *American Economic Review*, jan 2015, *105* (1), 172–203.

Grieco, Paul L.E., Charles Murry, and Ali Yurukoglu, "The Evolution of Market Power in the U.S. Automobile Industry," *The Quarterly Journal of Economics*, mar 2024, *139* (2), 1201–1253.

Handel, Benjamin R, "Adverse Selection and Inertia in Health Insurance Markets :," *American Economic Review*, 2013, *103* (7), 2643–2682.

— **and Jonathan T Kolstad**, "Health Insurance for Humans: Information Frictions, Plan Choice, and Consumer Welfare," *American Economic Review*, 2015, *105* (8), 2449–2500.

Handel, Benjamin R., Jonathan T. Kolstad, and Johannes Spinnewijn, "Information frictions and adverse selection: Policy interventions in health insurance markets," *Review of Economics and Statistics*, may 2019, *101* (2), 326–340.

Heckman, James J., Robert J. Lalonde, and Jeffrey A. Smith, "The Economics and Econometrics of Active Labor Market Programs," *Handbook of Labor Economics*, jan 1999, *3* (1), 1865–2097.

Heiss, Florian, Daniel Mcfadden, Joachim Winter, Amelie Wuppermann, and Bo Zhou, "Inattention and Switching Costs as Sources of Inertia in Medicare Part D," *NBER Working Paper*, 2016, *22765.*

Ho, Kate and Robin S Lee, "INSURER COMPETITION IN HEALTH CARE MARKETS," *Econometrica*, 2017, *85* (2), 379–417.

_ **and Robin S. Lee**, "Health insurance menu design for large employers," *RAND Journal of Economics*, dec 2023, *54* (4), 598–637.

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

Hörner, Johannes, "Reputation and competition," *American Economic Review*, jun 2002, *92* (3), 644–663.

Jaffe, Sonia and E. Glen Weyl, "The First-Order approach to merger analysis," *American Economic Journal: Microeconomics*, nov 2013, *5* (4), 188–218.

_ **and Mark Shepard**, "Price-Linked Subsidies and Health Insurance Markups," *National Bureau of Economic Research Working Paper Series*, 2017, *No. 23104.*

Keeler, Emmett B. and John E. Rolph, "The demand for episodes of treatment in the health insurance experiment," *Journal of Health Economics*, 1988, *7* (4), 337–367.

Kling, Jeffrey R., Sendhil Mullainathan, Eldar Shafir, Lee C. Vermeulen, and Marian V. Wrobel, "Comparison friction: Experimental evidence from medicare drug plans," *Quarterly Journal of Economics*, feb 2012, *127* (1), 199–235.

Kowalski, Amanda, "Censored Quantile Instrumental Variable Estimates of the Price Elasticity of Expenditure on Medical Care," *Journal of Business & Economic Statistics*, 2016, *34* (1), 107–117.

Lavetti, Kurt, Thomas DeLeire, and Nicolas Ziebarth, "How Do Low-Income Enrollees in the Affordable Care Act Marketplaces Respond to Cost-Sharing?," *National Bureau of Economic Research*, 2019.

Lester, Benjamin, Ali Shourideh, Venky Venkateswaran, and Ariel Zetlin-Jones, "Screening and Adverse Selection in Frictional Markets," *https://doi.org/10.1086/700730*, jan 2019, *127* (1), 338–377.

Mahoney, Neale and E. Glen Weyl, "Imperfect Competition in Selection Markets," *The Review of Economics and Statistics*, oct 2017, *99* (4), 637–651.

Manning, Willard G, Joseph P Newhouse, Naihua Duan, Emmett B Keeler, Arleen Leibowitz, and M Susan Marquis, "American Economic Association

Health Insurance and the Demand for Medical Care: Evidence from a Randomized Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment," *Source: The American Economic Review*, 1987, *77* (3), 251–277.

Marone, Victoria R and Adrienne Sabety, "Should There Be Vertical Choice in Health Insurance Markets? ," 2020.

Matsa, David A., "Competition and product quality in the supermarket industry," *Quarterly Journal of Economics*, aug 2011, *126* (3), 1539–1591.

McGuire, Thomas G., Joseph P. Newhouse, and Anna D. Sinaiko, "An Economic History of Medicare Part C," *Milbank Quarterly*, jun 2011, *89* (2), 289–332.

Miller, Keaton, Amil Petrin, Robert Town, and Michael Chernew, "Optimal Managed Competition Subsidies," 2019.

Pelech, Daria, "Paying more for less? Insurer competition and health plan generosity in the Medicare Advantage program," *Journal of Health Economics*, sep 2018, *61*, 77–92.

Ryan, Conor, "Mergers in the Presence of Adverse Selection," *Rand Journal of Economics (forthcoming)*, 2025, pp. 1–41.

Saltzman, Evan, "Managing adverse selection: underinsurance versus underenrollment," *The RAND Journal of Economics*, jun 2021, *52* (2), 359–381.

Schmalensee, Richard, "MARKET STRUCTURE, DURABILITY, AND QUALITY: A SELECTIVE SURVEY," *Economic Inquiry*, apr 1979, *17* (2), 177–196.

Shepard, Mark, "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange," *NBER Working Paper*, 2016.

Spence, A. Michael, "Monopoly, Quality, and Regulation," *The Bell Journal of Economics*, 1975, *6* (2), 417.

Starc, Amanda, "Insurer Pricing and Consumer Welfare: Evidence from Medigap," 2014.

Tebaldi, Pietro, "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA," *Review of Economic Studies*, 2023, *Forthcomin.*

Tilipman, Nicholas, "Employer Incentives and Distortions in Health Insurance Design: Implications for Welfare and Costs," *American Economic Review*, mar 2022, *112* (3), 998–1037.

Town, Robert, "The welfare impact of HMO mergers," *Journal of Health Economics*, nov 2001, *20* (6), 967–990.

_ **and Su Liu**, "The Welfare Impact of Medicare HMOs," *The RAND Journal of Economics*, 2003, *34* (4), 719.

Vatter, Benjamin, "Quality Disclosure and Regulation: Scoring Design in Medicare Advantage," *Econometrica*, 2025, *93* (3), 959–1001.

Veiga, André and E. Glen Weyl, "Product design in selection markets," *Quarterly Journal of Economics*, may 2016, *131* (2), 1007–1056.

Waldfogel, Joel, "Preference Externalities: An Empirical Study of Who Benefits Whom in Differentiated-Product Markets," *The RAND Journal of Economics*, 2003, *34* (3), 557.

APPENDIX

A Derivations of Merger Effects

In this section, I elaborate on the theory presented in Section 3.2, deriving both the conditions presented in the section and the same conditions applied to the full model.

A.1 A Simple Model

Consider a single product firm, facing a demand curve $S_j(p_j, x_j)$ that depends on the premium (p_j) and primary care copay (x_j) set by the firm. The firm has a constant marginal cost, $mc(x_j)$, that is decreasing and convex in primary care copays.

The profit earned by the firm is given by

$$\Pi_j(p_j,x_j) = S_j(p_j,x_j)(p_j - mc(x_j)) \tag{23}$$

The firm has two first order conditions with respect to each strategic variable. Ignoring non-negativity constraints, each equation relates to the firms markup.

$$p_j - mc(x_j) = -\frac{S_j}{\frac{\partial S_j}{\partial p_j}} \tag{24}$$

$$p_j - mc(x_j) = \frac{S_j}{\frac{\partial S_j}{\partial x_j}} \frac{\partial mc_j}{\partial x_j} \tag{25}$$

Combining these two equations leads to the condition that the ratio of demand slopes must be equal to the change in marginal cost.

$$\frac{\frac{\partial S_j}{\partial x_j}}{\frac{\partial S_j}{\partial p_j}} = -\frac{\partial mc_j}{\partial x_j} \tag{26}$$

When combined with the definition of $WTP_j(p,x)$, this leads to the condition in Equation (9). If we assume demand is homogeneous, then the left hand side

($WTP_j(p,x)$) is equal to the willingness to pay for reducing primary care copays among all consumers. Let $u_j(p_j, x_j)$ be the indirect utility of product j.

$$\frac{\frac{\partial S_j}{\partial x_j}}{\frac{\partial S_j}{\partial p_j}} = \frac{\frac{\partial u_j}{\partial x_j}\frac{\partial S_j}{\partial u_j}}{\frac{\partial u_j}{\partial p_j}\frac{\partial S_j}{\partial u_j}}$$
$$\frac{\frac{\partial S_j}{\partial x_j}}{\frac{\partial S_j}{\partial p_j}} = \frac{\frac{\partial u_j}{\partial x_j}}{\frac{\partial u_j}{\partial p_j}}$$

The ratio on the right-hand side reflects ratio of premium change to primary care copay change that will perfectly offset each other in the indirect utility of consumers.

A.2 Full Model

The first order conditions of a single product firm in the full model, ignoring nonnegativity constraints, are given by

$$\frac{1}{\frac{\partial s_j}{\partial p_j}} \int_i s_{ij}\left(\frac{\partial b_{ij}}{\partial p_j}+1\right)di + p_j = \frac{1}{\frac{\partial s_j}{\partial p_j}} \int_i \frac{\partial s_{ij}}{\partial p_j}\left(mc_{ij}-b_{ij}\right)di \tag{27}$$
$$\frac{1}{\frac{\partial s_j}{\partial x_j}} \int_i s_{ij}\left(\frac{\partial b_{ij}}{\partial x_j}-\frac{\partial mc_{ij}}{\partial x_j}\right)di + p_j = \frac{1}{\frac{\partial s_j}{\partial x_j}} \int_i \frac{\partial s_{ij}}{\partial x_j}\left(mc_{ij}-b_{ij}\right)di \tag{28}$$

The right-hand side of each of the above equations represents the net marginal cost with respect to premium and the primary care copay. I will refer to these as MC_p and MC_x, respectively. Each equation has an equality with respect to premium. Combining these conditions leads to the following condition.

$$-\frac{\int_i s_{ij}\frac{\partial mc_{ij}}{\partial x_j}di}{S_j} = \frac{\frac{\partial S_j}{\partial x_j}}{\frac{\partial S_j}{\partial p_j}} + \frac{\frac{\partial S_j}{\partial x_j}}{\frac{\partial S_j}{\partial p_j}}\left(\frac{\int_i s_{ij}\left(\frac{\partial b_{ij}}{\partial p_j}-\frac{\partial b_{ij}}{\partial x_j}\right)di}{S_j}\right) + \frac{\frac{\partial S_j}{\partial x_j}}{S_j}\left(MC_x - MC_p\right)$$

The left-hand side is the average infra-marginal effect on marginal costs due to a change in the copay for primary care, analogous to the left hand side in the simple model. The first term on the right is the same average willingness to pay for the demand facing the firm, WTP_j, as defined in the simple model.

This model has two additional terms on the right. The first term reflects the infra-marginal effect changes in per-person subsidies that result from a change in either the premium or the primary care copay. In the simple model, the firm is trading off a change in marginal cost with a $1 increase in revenue from premiums. In the full model, the change in revenue is $1 plus the effect on subsidies due to the changes in premiums and copays.

The final term reflects an adjustment for the difference in the marginal cost composition of consumers that are marginal with respect to the primary care copay or the premium. If consumers marginal to the primary care copay are more costly than those with respect to the premium, the final term is negative. If marginal consumers with respect to the primary care copay are more costly than those with respect to the premium, the optimal primary care copays will be greater.

A.3 Merger Pass-through

In this section, I follow Jaffe and Weyl (2013) to characterize the pass-through rate of a merging firm.

To reiterate the firm's problem expressed in equation (12), consider a single product firm j. I will define $f_j^l(P, X)$, the first order condition with respect to $l \in \{\text{Copay}, \text{Premium}\}$ 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

$$0 \leq f_j^{\text{copay}}(P,X) \equiv \frac{1}{\frac{\partial s_j}{\partial x_j}} \int_i s_{ij} \left(\frac{\partial b_{ij}}{\partial x_j} + \frac{\partial mc_{ij}}{\partial x_j} \right) di + p_j + \frac{1}{\frac{\partial s_j}{\partial x_j}} \int_i \frac{\partial s_{ij}}{\partial x_j} \left(mc_{ij} - b_{ij} \right) di \tag{29}$$

There also exists an analogous pre-merger first order condition with respect to premium given by $0 \leq f_j^{\text{prem}}(P,X)$. Let $F^l(P,X)$ be all of the stacked first order conditions in the market.

This equation is linear in cost that applies equally to all consumers. If all products were to face an equal increase in per-unit marginal costs, a constant vector t, then the resulting equilibrium could be characterized by:

$$\begin{bmatrix} F^{\text{prem}}(P,X) \\ F^{\text{copay}}(P,X) \end{bmatrix} + t \geq 0 \tag{30}$$

If all first-order conditions are binding, invertible, and the demand and cost-functions are smooth, C^2 functions, then the implicit function theorem characterizes the pass-through of the marginal cost increase into premiums and copays.

$$\begin{bmatrix} \frac{\partial P}{\partial t} \\ \frac{\partial X}{\partial t} \end{bmatrix} = - \begin{bmatrix} \frac{\partial F^{\text{prem}}}{\partial P} & \frac{\partial F^{\text{copay}}}{\partial P} \\ \frac{\partial F^{\text{prem}}}{\partial X} & \frac{\partial F^{\text{copay}}}{\partial X} \end{bmatrix}^{-1} \cdot 1 \tag{31}$$

If any of the first order conditions for premiums or copays are not binding, then the linear system is indeterminate. For these variables, the marginal pass-through rate is 0, and the above system can be solved for the remaining strategic variables.

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 (13). Analogous functions, h_j^{prem} and g_{jk}^{prem}, exist for the post-merger first order conditions and GePP with respect to the premium.

Let $H^l(P,X)$ similarly represent the stacked, post-merger first order conditions. If the post-merger first order conditions are also binding and invertible, the pass-through rate for the merged firm can be derived analogously to the pre-merger equilibrium. Strategic variables that are constrained can be dealt with in the same manner as with the pre-merger pass-through rate.

$$\begin{bmatrix} \frac{\partial P}{\partial t} \\ \frac{\partial X}{\partial t} \end{bmatrix} = - \begin{bmatrix} \frac{\partial H^{\text{prem}}}{\partial P} & \frac{\partial H^{\text{copay}}}{\partial P} \\ \frac{\partial H^{\text{prem}}}{\partial X} & \frac{\partial H^{\text{copay}}}{\partial X} \end{bmatrix}^{-1} \cdot 1 \tag{32}$$

The insight of Jaffe and Weyl (2013) is that the appropriate pass-through rate for incentives created by a merger is the post-merger pass-through rate defined in Equation (32) evaluated at the pre-merger equilibrium. Following this intuition, I compute these pass-through rates for each merger, and then define a product-level trade-off between primary care copays and premiums by computing the ratio of the product-specific pass-through rates. I plot the distribution of these pass-through rates in Figure 3.

$$\frac{\partial x_j}{\partial p_j} = \frac{\frac{\partial x_j}{\partial t}}{\frac{\partial p_j}{\partial t}} \tag{33}$$

B Data Processing

B.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 member, 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, the APCD product IDs are not linked to the public names of the products. The data are matched using aggregate information on the market shares of each plan in each county. In the APCD, MA products are identified in the product file using the line of business and insurance plan market fields. Members in the ACPD are linked to counties through their 5-digit zip code. Where the zip code does not fully identify the county, the observation is given a weight in all counties that intersect that zip code proportional to the distribution of population in the zip code. In Massachusetts, this affects a small number of observations. From this data, I can compute the MA market share of each APCD product ID in each county and month.

This data set can be compared to the county-month level market shares computed to the enrollment data made publicly available by CMS. Market shares from this data are computed among the medical MA plans that are not Senior Care Options plans, which are identified separately in the APCD. Then for each possible pair of a CMS plan ID and APCD plan ID, I compute the percent of 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.

B.2 Measuring Medical Consumption

The typical 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.

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

$$m_{ipjt} = \Gamma_p' L_{ip} + \iota_{pjt} + \varepsilon_{ip} \tag{34}$$

where ι_{pjt} is a procedure-product-time fixed effect that accounts for differences in billing practices across insurance plans and years. 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. Because all of the features are categorical and procedure-specific, I use a linear specification rather than a log.[41]

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

[41] A log specification predicts a similar amount of the variation in the data.

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.

B.3 Measuring Administrative Expenses

The data on administrative 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." I also use this later category for the firm, Health New England, which continues to only report in this category following 2015.

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 are found in Sections 4 and 5 of Part 1 of the MLR filing.

C Cost-sharing Parameters and Health

Policy makers are not only concerned about the cost of medical care but also the health benefits of Medicare Advantage beneficiaries. To provide this context for the merger results, I measure the effect of changes in cost-sharing parameters directly on inpatient mortality risk.

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.

$$d_{i\tau+s} = \zeta^{mort,s} x_{j\tau} + \iota^{mort,s\prime} W_{j\tau} + \lambda_{\tau}^{mort,s} + \gamma Z_{i\tau} + \lambda_{j}^{mort,s} + \omega_{i\tau}^{mort,s} \tag{35}$$

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. Instead, I estimate a modified model in which I control for the primary care copay in the prior period, and instrument for the change in primary care copay using the with-product change. To control for individual heterogeneity, I include the summary risk score and the 53 hierarchical condition categories that affect more than 5% of the sample.

$$d_{i\tau+s} = \zeta^{mort,s\prime}(\Delta x_{j\tau} + x_{j(\tau-1)}) + \iota^{mort,s\prime} W_{j\tau} + \gamma Z_{i\tau} + \lambda_{\tau}^{mort,s} + \lambda_{j}^{mort,s} + \omega_{i\tau}^{mort,s} \tag{36}$$

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 for product-level characteristics and health status.

Table A10 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. I show the result for 6-month, 12-month, and 18-month mortality. Standard errors are clustered at the same level as the medical consumption regression: product, risk quartile, and income quartile.

A \$10 increase in the copay for primary care leads to 0.3 percentage point increase in 12-month inpatient mortality, where the total population mean is 1.6 percent. This effect, 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 associated with a 4.3 percentage point increase in 12-month inpatient mortality. The increasing effect over a longer time horizon represents a similar fraction of baseline mortality risk.

The magnitude of the effect is in line with other estimates on the causal differences in health outcomes among insurance plans in MA. For instance, Abaluck et al. (2020) find that a standard deviation decrease in plan quality (defined in terms of its effect on mortality) leads to a 0.9 percentage point increase in mortality. My estimates imply that a standard deviation increase in the copay for primary care leads to about a quarter of this overall effect. These results are also consistent with findings in the literature that patients cut back on all types of care in the face of higher out-of-pocket prices, rather than the most unnecessary or wasteful care (Chandra et al. (2010), Baicker et al. (2015), Brot-Goldberg et al. (2017)).

These findings require some caveats. First, I can only observe inpatient mortality, and may be underestimating the effect if it affects the fraction of patients that die in a facility. Moreover, mortality is an extreme health outcome. It does not say anything about the health outcomes or quality of life of the surviving beneficiaries.

Second, the data do not have very much power to investigate such a rare outcome. If I include the risk and income heterogeneity in the estimation, as in Section 6.1, none of the causal coefficients are statistically different from zero. The estimates are also statistically insignificant when simultaneously estimating the causal effects of all the cost-sharing parameters, as in Section 6.1.

To address the potential endogeneity of other cost-sharing terms, I instead estimate the effect of each cost-sharing parameter one-by-one to investigate whether other, concurrent changes in cost-sharing might be affecting the result. I don't find any significant effect of any of the other cost-sharing parameters on 12-month mortality. I also estimate a specification that instruments for both the copay for primary care and the copay for specialty care, which have significant co-movement in the data. I find quantitatively similar effects of the copay for primary care and no effect for the copay for specialty care. The results are presented in Appendix Table A11.

Because of these caveats, I use the full-population estimate of 12-month mortality to provide suggestive context for the extent to which increases in medical expenditures are associated with better health outcomes.

D The Predicted Coinsurance Rate

I estimate the relationship between cost-sharing terms and an expected, plan-level coinsurance rate in order to summarize plan generosity in a way that is interpretable as expected out-of-pocket expenses.

I specify ex-post coinsurance as a consumer-specific, flexible polynomial of cost-sharing terms and health information of the consumers. The outcome fit in the data is the total out-of-pocket expenses of each consumer divided by their total expenses in a given month. The function includes second-order terms of all the cost-sharing characteristics included in the demand model, as well as third-order terms between the cost-sharing characteristics and the out-of-pocket maximum. The function includes firm-specific interaction terms between consumer health information and cost-sharing terms, and year-specific consumer health effects. To compute the plan-level average coinsurance rate, I average the consumer-predicted coinsurance across all consumers in the market.

A notable exception is that I exclude the copay for primary care from this function. The realized coinsurance rate depends on how consumers choose their medical care. Most out-of-pocket expenses are determined by fixed copays, and thus the fraction of covered expenses closely depends on the total amount of medical spending given any visit to the doctor. Because of this endogeneity, I exclude the primary care copay in order to preserve identification of an elasticity in the insurance demand model.

I estimate a probit function to fit the consumer specific, ex-post fraction of payments covered by the insurance firm using LASSO. First, I calibrate the LASSO penalty parameter on a 10 percent sample of all consumers, divided into three parts. I select the penalty value that minimizes out-of-sample prediction error, and re-estimate the LASSO procedure on the full sample. The estimated function contains 294 non-zero parameters, which include some variables from each set of covariates described above: linear terms, demographic-cost-sharing interactions, quadratic cost-sharing terms, and interactions between all of these variables and firm fixed effects. Predicted coinsurance rates explain about 25% of the variation in the out-of-pocket expense fraction among consumers.

E Estimating the Bid Function

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

$$b_{ij} = rs_i\left(\min\{Bench_j, bid_j(p_j,x_j)\} + \Lambda_j \underbrace{\max\{Bench_j - bid_j(p_j,x_j), 0\}}_{\text{Rebate}}\right) \quad (37)$$

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. The distance between a bid and a benchmark is termed the "rebate", and is given back to the plans at a share Λ_j, which depends on the plan's quality rating.

If we define a raw rebate value, $rb_j = \max\{Bench_j - bid_j(p_j,x_j), 0\}$, we can redefine the risk adjusted subsidy as

$$b_{ij} = rs_i\left(Bench_j + (\Lambda_j - 1)rb_j\right) \quad (38)$$

The goal of this section is to approximate rb_j as a function of the premium and primary care copay. I specify this function as additively separable components of premium and cost-sharing: $rb_j(p_j,x_j) = h(p_j) + g_j(x_j)$ and estimate it using national data on MA plans offered between 2011 and 2019.

The data contain information on rebates, but not plan-specific benchmarks. If the plan-specific benchmark level was directly observable, the bid itself could be inferred from equation (37). I follow Curto et al. (2021) in using an approximated plan-specific benchmark from the enrollment weighted average of county-level benchmarks.

In the simplest case, plan bids and premiums are related mechanically. If a plan bid is below the benchmark, and its premium is zero. If a plan bid is above the benchmark, the raw rebate is 0, and the premium is the difference between the bid and the benchmark: $p_j = bid_j(p_j,X_j) - Bench_j$. However, plans can still charge premiums for prescription drug coverage and other supplemental benefits. In the data, it is common for plans with positive premiums to still receive significant rebates. To capture this mechanism but allow for flexibility, I model the relationship between premiums and rebates, $h(p_j)$, as an exponential decay.

The relationship between the plan bids and the copay for primary care is less clear. A plan that receives rebates is supposed to return that rebate to consumers in the form of either lower premiums or more generous benefits. This suggests a negative relationship between primary care copays and rebates. However, the plan bid reflects the expected cost of coverage for the plan's beneficiaries. Since greater primary care copays lead to lower costs, this would suggest a positive relationship between primary care copays and rebates. Because of this more ambiguous relationship between rebates and the primary care copays (and other cost-sharing terms), I model the bid policy as linear in these terms. Figure A5 plots how the rebates vary with plan premiums and cost-sharing terms in the data.

I specify the rebate function as

$$rb_{jt} = \frac{a_1}{1+e^{a_2 p_j}} + \beta x_{jt} + \gamma W_{jt} + o_{jt} \tag{39}$$

The first term models the mechanical relationship between premiums and rebates, parameterized by a_1 and a_2. The rebates also depend linearly on x_{jt} and X_{jt}, which includes the primary care copay, other cost-sharing terms, product-fixed effects, and year fixed-effects. I also include the plan-specific benchmark in W_{jt} to allow for imperfect pass-through of the benchmark (Curto et al. (2021)).

I estimate the parameters using non-linear least squares. The identifying assumption is that there is no idiosyncratic and transient shock, observable to the firm, that affects both the bid and the premium or primary care copay. This assumption requires, for example, that firms do not alter the mapping between their product characteristics and their bid in response to a demand shock. This fits the setup of the model, in which the firms are only setting the observable features of the contract, and is also consistent with other work that assumes bids, premiums, and cost-sharing are three outcomes with only two degrees of freedom (Curto et al. (2021),Vatter (2025)).

The results of the bid estimation are presented in Table A12. The estimates imply that an increase in the premium of \$46 cuts the rebate in half. Most cost-sharing terms are negatively associated with the rebates. There is a small but significant negative relationship between the primary care copay and the rebate. The function explains the variation in rebates reasonably well, with an r-squared value of 0.93.

Main Text Tables and Figures

	First Stage	IV Estimates					
	Firms	Prem.	Primary	Spcl.	Urgent	Radio.	Inpt.
# of Firms		-9.2***	-3.3***	-0.74	8.2***	2.6	-2.3
		(3.0)	(0.97)	(1.2)	(2.3)	(7.2)	(10.6)
Log <65 Pop	0.33***						
	(0.04)						
Log >65 Pop	0.50***	5.0**	1.5*	-0.23	-6.2***	-4.0	-6.6
	(0.04)	(2.5)	(0.83)	(1.1)	(2.0)	(6.1)	(9.2)
Income ($000)	-0.002	0.19***	0.006	0.02	0.009	0.48***	-0.25
	(0.002)	(0.05)	(0.02)	(0.02)	(0.05)	(0.11)	(0.18)
% White	-1.4***	-12.9**	-3.9**	-2.1	11.8**	4.8	7.9
	(0.11)	(5.8)	(1.9)	(2.2)	(5.2)	(12.2)	(17.9)
Among Eligible							
% over 85	1.4***	17.2**	17.5***	5.8**	-15.3***	-15.9	56.7**
	(0.33)	(7.2)	(2.3)	(2.7)	(5.7)	(19.9)	(24.9)
% Employed	0.76***	8.8	11.1***	5.5**	-6.6	9.8	57.7***
	(0.29)	(6.0)	(1.9)	(2.4)	(4.8)	(16.7)	(21.5)
% Cog. Dis.	-0.97***	6.6	3.6**	1.4	2.3	-57.4***	36.8*
	(0.30)	(5.6)	(1.8)	(2.4)	(4.4)	(13.9)	(20.5)
Resources (per 1k)							
PC Docs	-0.36***	0.06	-0.68	-0.26	2.6**	-0.59	4.9
	(0.03)	(1.5)	(0.43)	(0.60)	(1.1)	(3.4)	(5.5)
Hosp. Beds	-0.002	0.06	0.02	0.05*	-0.02	-0.12	-0.12
	(0.002)	(0.05)	(0.02)	(0.03)	(0.04)	(0.16)	(0.21)
Fixed Effects							
State & Year	✓	✓	✓	✓	✓	✓	✓
$\frac{\text{Effect}}{\text{Data Mean}}$		-0.49	-0.35	-0.02	0.32	0.05	-0.01
		(0.55)	(0.21)	(0.02)	(0.14)	(0.10)	(0.02)

Note: An additional firm leads to lower premiums and primary care copays. The unit of observation is a US county in a given year between 2011 and 2017. The dependent variable is the enrollment weighted average of a product characteristic: prem - monthly premium; primary - primary care copay; spcl - specialist copay; urgent - urgent care copay; radio - radiology copay; and inpt - inpatient copay. The final row displays the effect size relative to the mean, with the variance of the mean below in parentheses. The significance stars ***, **, and * denote statistical significance at the 1%, 5%, and 10% level respectively. The first stage F statistic is 74.

Table 1: Evidence that Competition Reduces Cost-sharing Levels

	% 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

	OLS	TWFE	IV
Monthly Premium	-0.073***	-0.162***	-0.496**
	(0.015)	(0.018)	(0.201)
Primary Care	-0.376**	-0.181*	-2.39***
	(0.148)	(0.094)	(0.856)
Out-of-Pocket Limit	-0.114*	-0.137***	-0.523*
	(0.067)	(0.048)	(0.274)
Specialist	0.165	0.121	0.946
	(0.162)	(0.181)	(0.696)
Outpatient	-0.044***	0.006	-0.075
	(0.016)	(0.020)	(0.082)
Inpatient Stay	-0.006	0.020	-0.043
	(0.034)	(0.030)	(0.105)
Emergency Room	0.005	0.001	0.011
	(0.005)	(0.004)	(0.012)
Urgent Care	-0.003	0.010	0.134
	(0.036)	(0.028)	(0.240)
Ambulance	0.061	-0.007	-0.313
	(0.092)	(0.064)	(0.201)
Medical Device Coins.	-0.0009	-0.003	0.021
	(0.014)	(0.013)	(0.063)
Outpatient Coins.	0.003	0.010	0.038
	(0.027)	(0.022)	(0.043)
Predicted Coinsurance	6.26	-2.28	-5.21
	(4.93)	(2.52)	(7.52)
Predicted Coinsurance^2	7.76	-2.91	54.7
	(12.7)	(8.18)	(35.8)
Fixed Effects			
Part D, Rating, Year	✓	✓	✓
Product			✓
Product-County		✓	
F Stat - Premium			12.47
F Stat - Primary Care			22.51
Sargan J-Stat			0.82

Note: 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 1%, 5%, and 10% level respectively. Standard errors are clustered at the product level and bootstrapped using draws from the estimated distribution of δ_{jt}.

Table 3: Estimation Results for Base Insurance Demand Parameter Vector

	Monthly Premium	Primary Care Copay
	Elasticity	
Mean	-1.75	-1.72
25^{th} Percentile	-2.32	-2.35
Median	-1.01	-0.99
75^{th} Percentile	-0.36	-0.56
	Semi-Elasticity	
Mean	-2.00	-10.9
25^{th} Percentile	-3.76	-21.8
Median	-1.12	-6.7
75^{th} Percentile	-0.77	-4.30

Note: The tables shows the mean, median, 25^{th} percentile, and 75^{th} percentile for the semi-elasticities of demand with respect to the monthly premium and copay for primary care in the demand estimation. In the means and quantiles, elasticities represent own-premium elasticities for a consumer-product pair, are weighted by the purchase probability and integrated over unobserved heterogeneity. The semi-elasticities represent the percent change in the probability a consumer purchases their chosen plan given a $1 increase in each characteristic.

Table 4: Insurance Demand Elasticities for Premium and Primary Care Copays

	Monthly Spending	Consumption	Covered Spending	Covered Consumption
Primary Care	-0.109***	-0.114***	-0.085**	-0.083**
	(0.042)	(0.043)	(0.039)	(0.040)
Specialty Care	-0.143***	-0.128***	-0.138***	-0.130***
	(0.040)	(0.041)	(0.039)	(0.040)
Outpatient	-0.046***	-0.044***	-0.047***	-0.046***
	(0.010)	(0.010)	(0.010)	(0.010)
Inpatient	0.000	0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Ambulance	-0.015***	-0.014***	-0.012***	-0.010***
	(0.004)	(0.004)	(0.003)	(0.003)
Emergency Room	0.037***	0.042***	0.029***	0.030***
	(0.009)	(0.009)	(0.008)	(0.008)
Urgent Care	-0.091***	-0.086***	-0.089***	-0.083***
	(0.015)	(0.016)	(0.014)	(0.015)
Imaging	0.001	0.001	0.001	0.001
	(0.000)	(0.001)	(0.001)	(0.001)
Fixed-effects				
Risk-Income-Firm	✓	✓	✓	✓
Risk-Income-Month	✓	✓	✓	✓
Product	✓	✓	✓	✓
Observations	4,881,912	4,881,912	4,881,912	4,881,912

Note: The results of different specifications to estimate the semi-elasticity of medical consumption with respect to each cost-sharing term. All cost-sharing terms are copays denominated in \$10. The estimates for each term displayed in the table are estimated using the inertia-based instrument described in Section 6.2. Additional terms for out-of-pocket maximum, outpatient coinsurance, and outpatient drugs coinsurance are included as controls. The significance stars ***, **, and * denote statistical significance at the 1%, 5%, and 10% level respectively. Standard errors are clustered by product, risk, and income group. The F-statistics are all greater than 1×10^5.

Table 5: Medical Consumption Responds to Primary Care Copays

	Tufts - BCBS		Tufts - United		BCBS - United	
	Parties	All	Parties	All	Parties	All
Mean Δ HHI	187		286		267	
# Markets	2		4		5	
Mean Market Size	91		87		123	
Baseline Premium	88	68	71	78	86	88
Change Within Product	0.2%	-0.1%	1.3%	0.3%	4.6%	1.5%
Sorting Across Products	0.7%	0.2%	-1.6%	-0.8%	0.1%	-0.6%
Average Merger Effect	0.8%	0.1%	-0.3%	-0.4%	4.7%	0.9%
Baseline Copay	6.4	2.3	8.3	8.7	15.7	10.4
Change Within Product	-0.3%	-0.4%	-1.0%	-1.5%	-1.5%	-1.9%
Sorting Across Products	0.7%	0.3%	0.6%	-2.0%	0.0%	-1.3%
Average Merger Effect	0.3%	0.0%	-0.5%	-3.5%	-1.5%	-3.1%
Welfare Measures						
Enrollment	0.1%		2.9%		0.9%	
Annual Consumer Surplus	\$0.4		\$2.2		\$0.4	
Annual Medical Consumption	\$0.5		\$66		\$42	
Expected Inpatient Mortality	0.0 bp		-0.7 bp		-1.0 bp	

Note: This table shows the mean effects of the merger analysis of three hypothetical mergers in markets for which the change in HHI is less than 750. Premium and copay effects are decomposed into product-level changes, holding pre-merger enrollment fixed, and consumer sorting across products. These two effects combine for the effect on the average premium or primary care copay paid by consumers. Averages are computed within market according to enrollment, and across markets according to market size. “Parties” refers to the merging firms and “All” refers to all products in the market.

Table 6: Merger Effects for Small Concentration Changes

	Tufts - BCBS		Tufts - United		BCBS - United	
	Parties	All	Parties	All	Parties	All
Mean Δ HHI	2,310		1,634		1,205	
# Markets	8		3		2	
Mean Market Size	99		142		79	
Baseline Premium	107	95	91	93	55	76
Change Within Product	6.2%	4.9%	2.5%	1.4%	9.6%	4.0 %
Sorting Across Products	0.0%	-5.5%	3.2%	1.8%	-3.6%	-1.1%
Average Merger Effect	6.3%	-0.6%	5.7%	3.2%	6.0%	2.8%
Baseline Copay	14.6	14.7	11.9	13.8	17.1	15.6
Change Within Product	-3.7%	-4.7%	3.7%	2.3%	-0.4%	-1.0 %
Sorting Across Products	-0.6%	-0.9%	-1.3%	0.5%	0.6%	0.1%
Average Merger Effect	-4.3%	-5.5%	2.3%	2.8%	0.2%	-0.9%
Welfare Measures						
Enrollment	1.6%		-2.9%		-2.9%	
Annual Consumer Surplus	$0.9		-4.0		-2.4	
Annual Medical Consumption	$106		-$55		-$3.4	
Expected Inpatient Mortality	-2.5 bp		1.2 bp		-0.4 bp	

Note: This table shows the mean effects of the merger analysis of three hypothetical mergers in markets for which the change in HHI is greater than 750. Premium and copay effects are decomposed into product-level changes, holding pre-merger enrollment fixed, and consumer sorting across products. These two effects combine for the effect on the average premium or primary care copay paid by consumers. Averages are computed within market according to enrollment, and across markets according to market size. “Parties” refers to the merging firms and “All” refers to all products in the market.

Table 7: Merger Effects for Large Concentration Changes

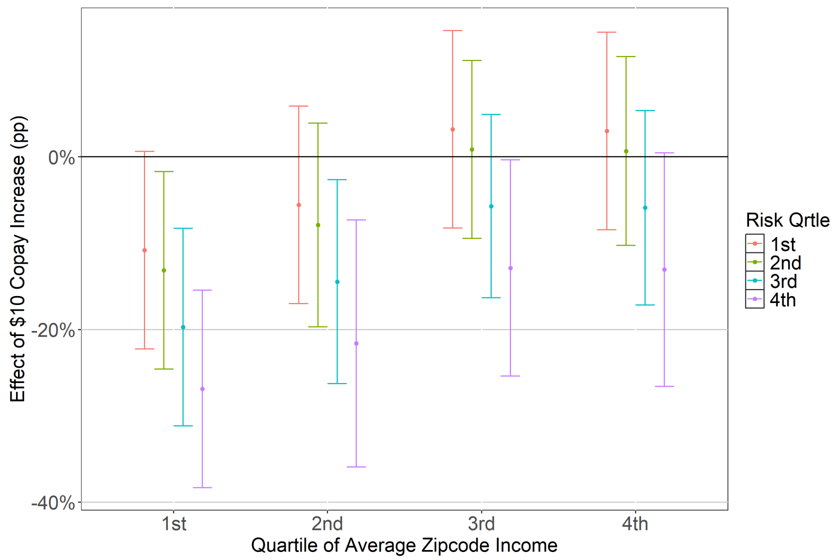

Figure 1: Medical Consumption Elasticity Heterogeneity

Note: Consumer semi-elasticity with respect to the primary care copay is highest among consumers in the lowest income zipcodes and with the highest medical risk. This figure shows the results of the medical consumption demand estimation. The outcome is insurance-covered consumption, as defined in Section 6.1. The primary care copays are denominated in $10 and effect sizes are semi-elasticities shown in percentage points. Consumers are divided by income of the zip code in which they live and their average risk score throughout the sample. Confidence intervals are shown at the 5% level. For more details, see the notes of Table 5

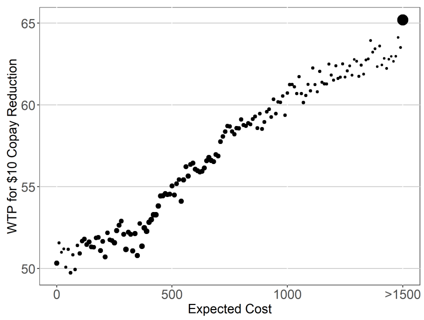

(a) Willingness to Pay and Cost

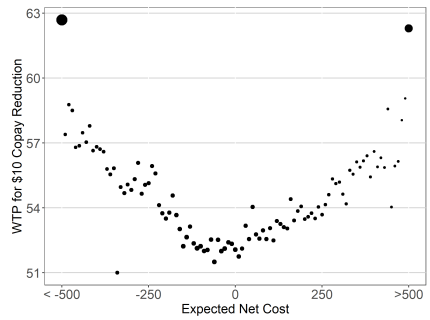

(b) Willingness to Pay and Net Cost

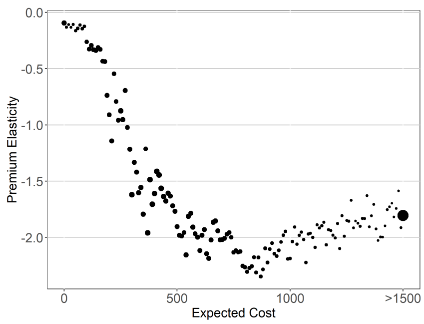

(c) Premium Elasticity and Cost

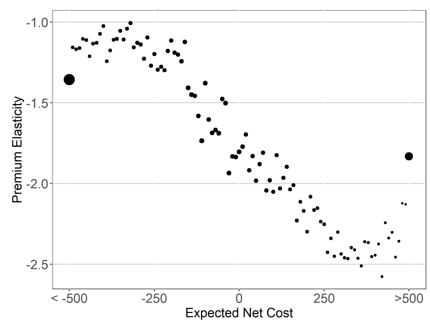

(d) Premium Elasticity and Net Cost

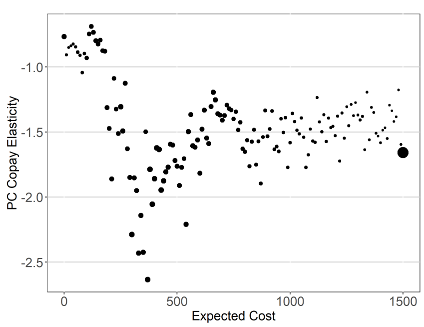

(e) Primary Care Copay Elasticity and Cost

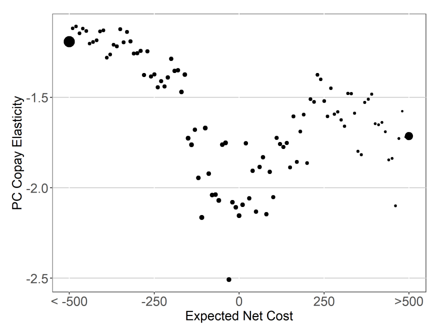

(f) Primary Care Copay Elasticity and Net Cost

Figure 2: Selection on Premium and Primary Care Copays

Note: This figure is a binned scatter plot of selection features of demand. Each row plots, from top to bottom, the willingness to pay for a $10 reduction in the primary care copay, the premium elasticity of demand, and the primary care copay elasticity, with respect to total expected cost (left) and net expected cost (right). The size of each dot corresponds to the size of the bin.

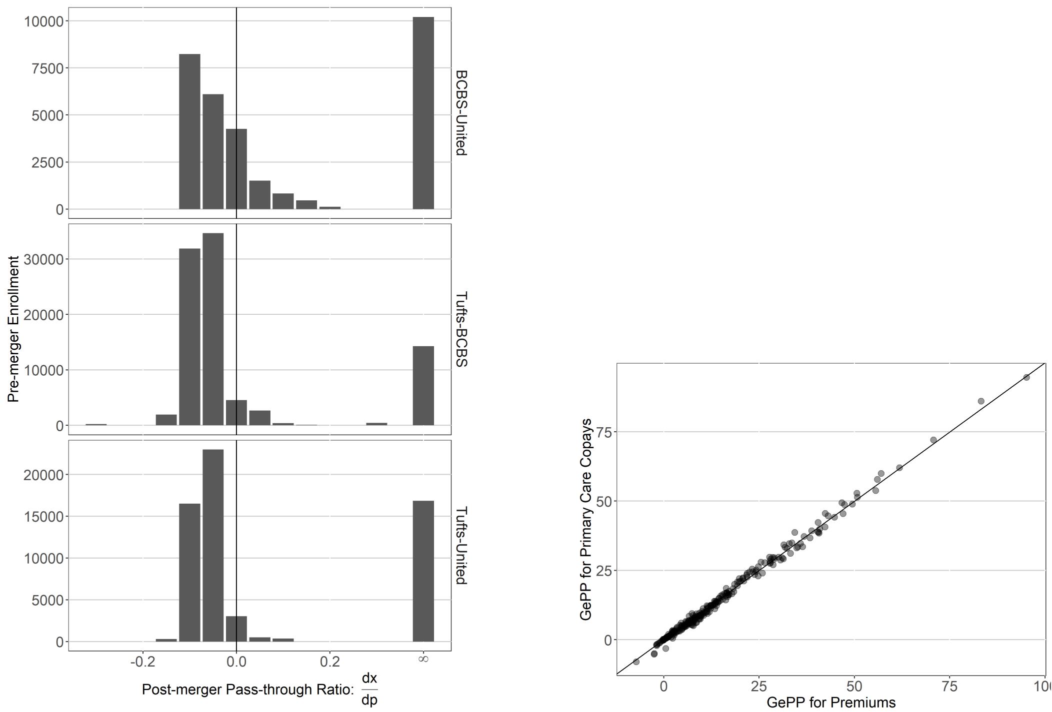

(a) Pass-through Ratio of Primary Care Copay to Premium

(b) Generalized Pricing Pressure for Primary Care Copays and Premiums

Figure 3: Merger Mechanisms

Note: This figure illustrates mechanisms through which a merger affects primary care copays and premiums. Panel (a) plots a histogram, weighted by pre-merger enrollment, of the pass-through ratio between primary care copays and premiums. Panel (b) plots the product-merger specific generalized pricing pressure created by each merger for primary care copays and premiums.

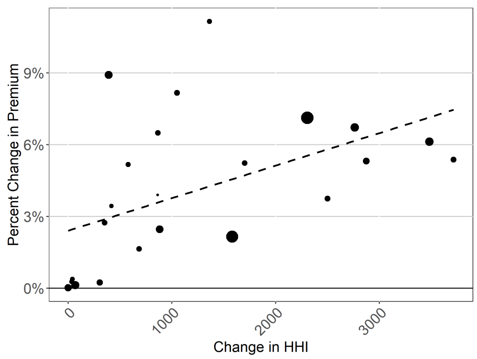

(a) Premium Effect and HHI

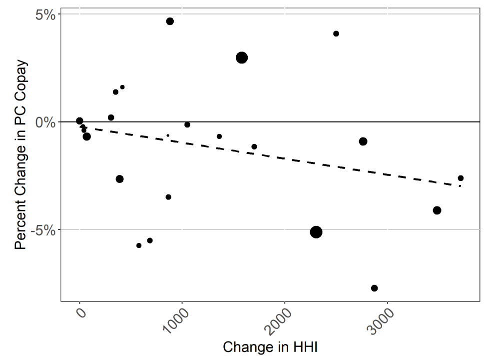

(b) Copay Effect and HHI

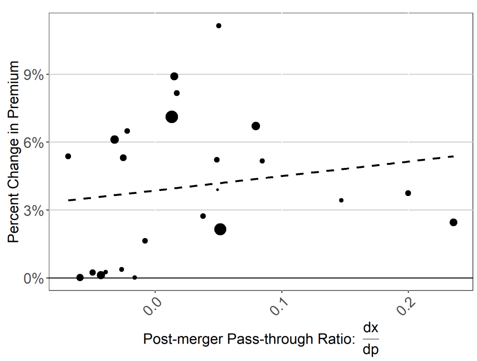

(c) Premium Effect and Pass-through

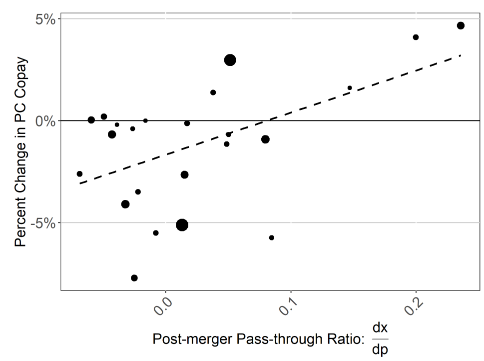

(d) Copay Effect and Pass-through

Figure 4: Merger Effects, HHI, and the Pass-through Rate

Note: Each point represents the average merger effect for merging firms' products in a particular merger for a particular market. The size of each point is proportional to amount of pre-merger enrollment in the merging firm's products. The top panel illustrates merger effects with respect to the change in HHI, computed using pre-merger market shares. The bottom panel illustrates merger effects with respect to the relative pass-through rates of primary care copays to premiums. The dotted line represents a linear fit.

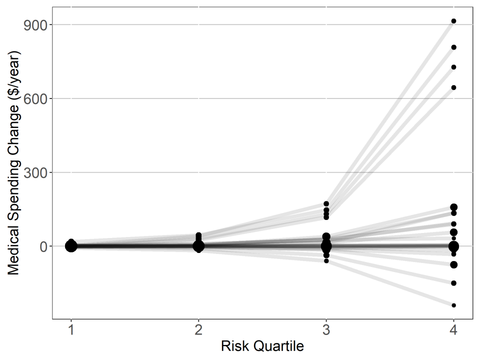

(a) Heterogeneity by Risk Quartile

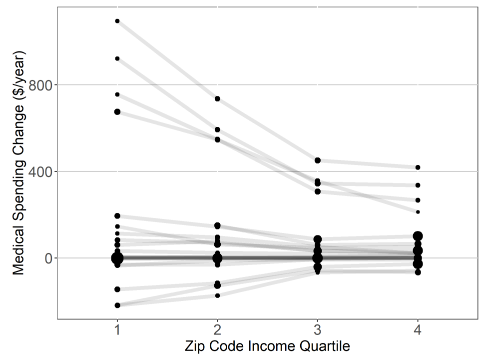

(b) Heterogeneity by Zip Code Income Quartile

Figure 5: Consumer Heterogeneity in Merger Effects on Medical Consumption

Note: This figure plots consumer-level heterogeneity in the effect of each merger on medical consumption. Each point represents the average effect on medical consumption for consumers within each demographic group for a particular merger and particular market. The size of the point represents the number of consumers represented by the average. The gray lines connect points across consumer groups for the same merger-market.

Appendix Tables and Figures

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 Pop (000s)	4.2	8.9	29.0	92.7	343.0
FFS Risk Score	0.89	0.95	0.98	1.01	1.14
Over 75	0.45	0.43	0.42	0.43	0.44
	Enrollment Weighted Characteristics				
Premium (monthly)	34.3	27.7	23.6	17.5	2.4
Part B Rebate	0.1	0.1	0.6	2.6	2.5
Deductible	20.5	22.8	19.4	13.7	5.4
OOP Limit	6,452	5,961	5,581	5,457	4,714
Copays					
Primary Care	15.4	12.9	11.1	8.8	3.8
Specialist	34.1	33.8	34.0	31.5	12.8
Outpatient	103.8	91.2	109.8	96.7	43.1
Imaging	83.4	68.3	59.1	45.5	40.1
Lab Tests	3.9	3.8	4.4	4.5	4.0
Emergency	67.8	65.8	65.8	65.4	60.0
Urgent Care	19.1	24.7	27.7	27.2	18.1
Inpatient	288.9	267.0	255.5	247.7	130.2
Ambulance	205.1	188.7	191.0	187.8	164.0
Coinsurance Rates					
Outpatient	0.112	0.094	0.068	0.058	0.040
Imaging	0.065	0.062	0.073	0.081	0.048
Med Devices	0.190	0.192	0.183	0.169	0.141
Outpt Drugs	0.162	0.158	0.160	0.159	0.140

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 2017. 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. The FFS risk score refers to the risk of traditional medicare enrollees, and Over 75 refers to the fraction of Medicare eligible consumers over the age of 75.

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

	Baseline	(1)	(2)	(3)	(4)
Monthly Premium	-0.496**	-0.500**	-0.495**	-0.499**	-0.485***
	(0.201)	(0.203)	(0.198)	(0.201)	(0.168)
Primary Care	-2.39***	-3.07**	-1.70**	-2.12***	-1.68***
	(0.856)	(1.42)	(0.704)	(0.819)	(0.454)
Out-of-Pocket Limit	-0.523*	-0.574*	-0.511*	-0.565*	
	(0.274)	(0.301)	(0.284)	(0.319)	
Specialist	0.946		0.945*		
	(0.696)		(0.505)		
Outpatient	-0.075	-0.079	-0.029	-0.015	
	(0.082)	(0.097)	(0.060)	(0.064)	
Inpatient Stay	-0.043	-0.061	0.020	0.027	
	(0.105)	(0.129)	(0.078)	(0.086)	
Emergency Room	0.011	0.014	0.009	0.011	
	(0.012)	(0.016)	(0.010)	(0.012)	
Urgent Care	0.134	0.072	0.179	0.148	
	(0.240)	(0.245)	(0.278)	(0.283)	
Ambulance	-0.313	-0.333	-0.210	-0.187	
	(0.201)	(0.265)	(0.170)	(0.189)	
Medical Device Coins.	0.021	0.062	-0.047	-0.037	
	(0.063)	(0.100)	(0.077)	(0.084)	
Outpatient Coins.	0.038	0.045	0.036	0.044	
	(0.043)	(0.051)	(0.032)	(0.035)	
Predicted Coinsurance	-5.21	-12.4			
	(7.52)	(8.19)			
Predicted Coinsurance^2	54.7	87.0			
	(35.8)	(57.9)			
F Stat - Premium	12.474	12.311	13.103	12.293	12.658
F Stat - Primary Care	22.510	15.076	32.203	23.936	37.676
Sargan J-Stat	0.81679	0.45515	1.2397	0.89083	1.3265

Note: This table contains robustness specifications for the final column of Table 3. Each specification drops the missing product characteristics. See Table 3 for other details. The significance stars ***, **, and * denote statistical significance at the 1%, 5%, and 10% level respectively. Standard errors are clustered at the product level and bootstrapped using draws from the estimated distribution of δ_{jt}.

Table A2: Robustness Specifications for Base Insurance Demand Parameter Vector

Panel A: Average Switch Rate	
Full Sample	3.1%
MA Consumers	13.4%
Given Switch from an MA product	
Choose Same Firm	22.3%
Choose New Firm	30.2%
Choose TM	47.5%

Panel B: Distribution of Differences Between Chosen Product and Prior Product Characteristics

	10%	25%	Median	75%	90%
Premium	-129.4	-83.0	-33.8	0.0	34.0
Primary Care Copay	-15	-5	0	5	10

Note: This table shows descriptive statistics for switching rates in the estimation sample. Panel A plots the frequency of switching and a description of choice categories conditional on switching. Panel B plots the differences between the chosen product and the consumers prior product (in the current year) among switching consumers.

Table A3: Switching Statistics

	Level		Annual Differences	
	Premium	PC Copay	Premium	PC Copay
Monthly Premium		-0.01^{**}		-0.06^{***}
		(0.00)		(0.01)
Primary Care	-0.77^{**}		-0.70^{***}	
	(0.26)		(0.12)	
Out-of-Pocket Limit	-0.00	0.00^{***}	-0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)
Specialist	0.42	0.47^{***}	-0.00	0.51^{***}
	(0.22)	(0.02)	(0.11)	(0.03)
Outpatient	-0.28^{***}	-0.03^{***}	-0.00	-0.03^{***}
	(0.03)	(0.00)	(0.02)	(0.01)
Outpatient Coins	-4.00^{***}	-0.97^{***}	0.35	-0.31^{*}
	(0.58)	(0.06)	(0.50)	(0.15)
Inpatient Stay	-0.06^{***}	0.03^{***}	0.00	0.00
	(0.01)	(0.00)	(0.01)	(0.00)
Emergency Room	-0.14^{*}	-0.11^{***}	0.14^{***}	0.03^{***}
	(0.07)	(0.01)	(0.03)	(0.01)
Ambulance	-0.22^{***}	-0.01	-0.08^{***}	-0.01^{*}
	(0.03)	(0.00)	(0.02)	(0.01)
Medical Devices	-5.76^{***}	-0.30^{***}	0.15	-0.00
	(0.33)	(0.04)	(0.18)	(0.05)
Outpatient Drugs	0.99^{***}	0.28^{***}	-0.21	-0.21^{***}
	(0.27)	(0.03)	(0.15)	(0.04)
Diagnostic Imaging	-0.03	0.01^{*}	-0.03	0.01^{*}
	(0.03)	(0.00)	(0.02)	(0.01)
R^2	0.73	0.79	0.18	0.57

Note: This table presents the relationship between product characteristics in the estimation sample. The level regressions show the projection of premium and primary care copay levels on other product characteristics included in the estimation. The difference regressions show the projection of year-over-year changes in premiums and primary care copays on year-over-year changes in the other product characteristics. The significance stars ***, **, and * denote statistical significance at the 1%, 5%, and 10% level respectively.

Table A4: Correlations between Premiums, Primary Care Copays, and other product characteristics

	Fraction of Products with Any Contract Change		Distribution of Contract Changes		
	Unweighted	Enrollment-Weighted	25%	Median	75%
Monthly Premium	0.71	0.73	2	8	11
Primary Care	0.19	0.23	5	10	15
Out-of-Pocket Limit	0.06	0.09	500	1500	1500
Specialist	0.20	0.17	5	5	30
Outpatient	0.40	0.52	-14	25	50
Outpatient Coins	0.002	0.002	-20	-20	-20
Inpatient Stay	0.36	0.40	-5	35	50
Emergency Room	0.26	0.25	10	10	75
Ambulance	0.29	0.36	25	50	50
Medical Devices	0.07	0.12	10	10	10
Outpatient Drugs	0.10	0.08	10	10	15
Diagnostic Imaging	0.29	0.28	25	50	100

Note: This table presents descriptive statistics for year over year changes in product characteristics in the estimation sample. The left panel shows the frequency of any change in the characteristic. The unweighted average is computed across product-years, and the enrollment-weighted average weights by consumers enrolled in the product in the initial year of the difference. The right panel shows the distribution of the value of year-over-year changes, computed across product-years.

Table A5: Year-to-year changes in Product Characteristics

	Probability of Any Use	Office Spending	Non-Office Spending	Prior Year Risk Score
Primary Care	-0.012**	-0.100***	-0.083**	0.002
	(0.005)	(0.030)	(0.036)	((0.007)
Specialty Care	-0.013***	-0.160***	-0.104***	-0.003
	(0.005)	(0.028)	(0.034)	((0.006)
Outpatient	-0.006***	-0.030***	-0.037***	0.003***
	(0.001)	(0.007)	(0.008)	((0.001)
Inpatient	0.000	0.000	-0.001	0.000
	(0.000)	(0.001)	(0.001)	((0.000)
Ambulance	-0.001***	-0.012***	-0.009***	0.003**
	(0.000)	(0.003)	(0.003)	((0.001)
Emergency Room	0.003***	0.031***	0.025***	-0.001
	(0.001)	(0.007)	(0.007)	((0.003)
Urgent Care	-0.009***	-0.076***	-0.087***	0.014**
	(0.002)	(0.013)	(0.014)	((0.006)
Imaging	0.002*	0.006	0.008	-0.002
	(0.000)	(0.005)	(0.006)	((0.001)
Fixed-effects				
Risk-Income-Firm	✓	✓	✓	✓
Risk-Income-Month	✓	✓	✓	✓
Product	✓	✓	✓	✓
Observations	4,881,912	4,881,912	4,881,912	4,881,912

Note: This table displays the results of applying the identical identification strategy in Table 5 to other outcomes. I estimate the effect of cost-sharing terms on any medical use, office spending, and non-office spending. As described in Section 6.1, the outcomes are first-differenced, and the spending outcomes have the same log transformation. I also include a lagged first-difference of consumer risk scores to test selection on future copay changes. The significance stars ***, **, and * denote statistical significance at the 1%, 5%, and 10% level respectively. Standard errors are clustered by product, risk, and income group. See Table 5 notes for additional details.

Table A6: Other Medical Consumption Outcomes with respect to Cost Sharing Terms

	Probability of Using Particular Service			
	Office Visits	Minor Procedures	Major Procedures	Hospital Visits
Primary Care	-0.012***	-0.007***	-0.000	-0.000
	(0.004)	(0.002)	(0.000)	(0.000)
Specialty Care	-0.019***	-0.005***	-0.0004**	0.000
	(0.004)	(0.002)	(0.0002)	(0.000)
Outpatient	-0.004***	-0.002***	-0.000	-0.0003 ***
	(0.000)	(0.000)	(0.000)	(0.0001)
Inpatient	-0.000	-0.000	-0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Ambulance	-0.000	-0.000	0.000	-0.0003 ***
	(0.000)	(0.000)	(0.000)	(0.0001)
Emergency Room	0.004***	0.001***	-0.000	0.0006 ***
	(0.000)	(0.000)	(0.000)	(0.0000)
Urgent Care	-0.009***	-0.005***	0.000	-0.000
	(0.002)	(0.001)	(0.000)	(0.000)
Imaging	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Fixed-effects				
Risk-Income-Firm	✓	✓	✓	✓
Risk-Income-Month	✓	✓	✓	✓
Product	✓	✓	✓	✓
Observations	4,881,912	4,881,912	4,881,912	4,881,912

Note: This table displays the results of applying the identical identification strategy in Table 5 to other outcomes. I estimate the effect of cost-sharing terms on binary indicators of any office, minor procedure, major procedure, or hospital spending in a given month. The service categories are defined using BETOS codes. As described in Section 6.1, the outcomes are first-differenced. The significance stars ***, **, and * denote statistical significance at the 1%, 5%, and 10% level respectively. Standard errors are clustered by product, risk, and income group. See Table 5 notes for additional details.

Table A7: Service Type Use and Cost Sharing Terms

	Monthly Spending	Consumption	Covered Spending	Covered Consumption
Primary Care	-0.126**	-0.137**	-0.101*	-0.108*
	(0.061)	(0.063)	(0.057)	(0.058)
Risk Quartile				
2nd	-0.028	-0.017	-0.033	-0.023
	(0.048)	(0.048)	(0.044)	(0.044)
3rd	-0.101**	-0.080*	-0.106**	-0.089**
	(0.046)	(0.047)	(0.043)	(0.044)
4th	-0.164***	-0.143**	-0.178***	-0.161***
	(0.059)	(0.060)	(0.057)	(0.057)
Income Quartile				
2nd	0.051	0.048	0.0055	0.053
	(0.066)	(0.067)	(0.0063)	(0.064)
3rd	0.133**	0.133**	0.139**	0.140**
	(0.060)	(0.061)	(0.057)	(0.058)
4th	0.131**	0.130**	0.137**	0.138**
	(0.0061)	(0.0062)	(0.0058)	(0.0059)
Specialty Care	-0.135***	-0.127***	-0.130***	-0.123***
	(0.040)	(0.040)	(0.038)	(0.038)
Outpatient	-0.046***	-0.045***	-0.046***	-0.046***
	(0.010)	(0.010)	(0.009)	(0.010)
Inpatient	-0.000	-0.000	-0.001	-0.000
	(0.001)	(0.001)	(0.001)	(0.001)
Ambulance	-0.014***	-0.013***	-0.012***	-0.010***
	(0.004)	(0.004)	(0.003)	(0.003)
Emergency Room	0.037***	0.041***	0.029***	0.029***
	(0.009)	(0.010)	(0.008)	(0.009)
Urgent Care	-0.097***	-0.090***	-0.096***	-0.090***
	(0.016)	(0.016)	(0.015)	(0.016)
Imaging	0.012	0.013*	0.009	0.010
	(0.007)	(0.008)	(0.007)	(0.008)
Fixed-effects				
Risk-Income-Firm	✓	✓	✓	✓
Risk-Income-Month	✓	✓	✓	✓
Product	✓	✓	✓	✓
Observations	4,881,912	4,881,912	4,881,912	4,881,912

Note: This table presents the coefficients for the heterogeneity used in the counterfactual estimation. The total coefficients for all consumers in the final specification are plotted in Figure 1. For detailed notes, see Table 5. The significance stars ***, **, and * denote statistical significance at the 1%, 5%, and 10% level respectively. Standard errors are clustered by product, risk, and income group. The F-statistics are all greater than 1×10^5.

Table A8: Medical Consumption Elasticity Heterogeneity

	MA Share	Avg Prem	Avg Copay	Avg. Risk Data	Model	Risk Adj. Cost Data	Model
Tufts	0.47	111	11.7	1.07	1.12	812	812
BCBS	0.27	113	19.0	0.84	0.86	897	939
United	0.09	13.1	13.2	0.89	0.97	802	781
Fallon	0.05	93.9	20.2	0.98	1.13	922	936
AARP	0.05	47.0	20.0	0.99	1.03	879	868
Health New Engl.	0.04	133	19.5	0.95	1.14	873	773
Harvard Pilgrim	0.03	116	13.6	0.96	1.07	934	941

Note: This table shows summary statistics for the seven 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 A9: Firms are Differentiated in their Premiums, Copays, and Risk Distributions

	Inpatient Mortality		
	6 - Month	12 - Month	18 - Month
Primary Care	0.106	0.305**	0.393**
	(0.077)	(0.133)	(0.177)
Prior Year PC Copay	0.005**	0.003	0.001
	(0.002)	(0.004)	(0.005)
Specialty Care	-0.049	0.323	0.259
	(0.119)	(0.237)	(0.263)
Outpatient	-0.010	-0.006	0.017
	(0.010)	(0.021)	(0.025)
Inpatient	0.0009	-0.0003	0.0005
	(0.002)	(0.006)	(0.008)
Emergency Room	-0.009	0.019	0.041
	(0.018)	(0.032)	(0.039)
Urgent Care	0.007*	0.006	0.004
	(0.004)	(0.007)	(0.009)
Ambulance	-0.001	-0.004	0.011
	(0.005)	(0.012)	(0.015)
Imaging	0.008	-0.015	-0.017
	(0.008)	(0.015)	(0.018)
Medical Device Coins	-0.021	-0.012	-0.006
	(0.023)	(0.043)	(0.044)
Outpatient Drugs Coins	0.013*	0.020*	0.008
	(0.007)	(0.012)	(0.012)
Out-of-pocket Maximum	0.009	0.041	0.046
	(0.022)	(0.042)	(0.053)

Note: This table shows the effect of primary care copays on inpatient mortality, measured in percentage points. All cost-sharing terms are copays denominated in \$10, except for coinsurance rates (pp) and the OOP max (\$1000). The specification includes additional controls for risk score, 53 HCC diagnoses groups, and fixed effects for firm-risk-income, month-risk-income, and product. The significance stars ***, **, and * denote statistical significance at the 1%, 5%, and 10% level respectively. Standard errors are clustered by product, risk, and income group. The F-statistics are all greater than 1×10^6.

Table A10: Higher Primary Care Copays lead to Higher Inpatient Mortality

	Inpatient Mortality			
	(1)	(2)	(3)	(4)
Primary Care	0.199**	0.220**	0.156	0.207**
	(0.096)	(0.088)	(0.103)	(0.088)
Specialty Care	**1.32**	0.202	0.386*	0.365
	(2.67)	(0.248)	(0.225)	(0.223)
Outpatient	-0.020	**0.025**	-0.008	-0.004
	(0.051)	**(0.031)**	(0.022)	(0.022)
Inpatient	0.003	-0.002	**0.005**	0.003
	(0.006)	(0.006)	**(0.010)**	(0.006)
Emergency Room	0.025	-0.016	0.006	**0.005**
	(0.047)	(0.034)	(0.031)	**(0.036)**
Urgent Care	-0.002	0.003	0.002	0.004
	(0.020)	(0.007)	(0.007)	(0.007)
Ambulance	-0.015	0.007	-0.008	-0.004
	(0.025)	(0.014)	(0.012)	(0.013)
Imaging	-0.018	-0.026	-0.011	-0.014
	(0.021)	(0.019)	(0.016)	(0.016)
	(5)	(6)	(7)	(8)
Primary Care	0.201**	0.225**	0.202**	**0.321***
	(0.089)	(0.088)	(0.086)	**(0.174)**
Specialty Care	0.379*	0.375	0.226	**2.71**
	(0.228)	(0.233)	(0.251)	**(3.54)**
Outpatient	-0.007	-0.008	-0.028	-0.047
	(0.021)	(0.029)	(0.030)	(0.067)
Inpatient	0.003	0.002	0.006	-0.0003
	(0.005)	(0.006)	(0.007)	(0.008)
Emergency Room	0.010	0.016	0.015	0.061
	(0.032)	(0.045)	(0.033)	(0.075)
Urgent Care	**0.005**	0.004	0.010	-0.009
	(0.008)	(0.007)	(0.008)	(0.022)
Ambulance	-0.006	**-0.009**	-0.012	-0.024
	(0.012)	**(0.026)**	(0.014)	(0.032)
Imaging	-0.014	-0.011	**0.018**	0.034
	(0.016)	(0.017)	**(0.038)**	(0.044)

Note: Each specification applies the IV strategy outlined in Appendix Section C to the bolded coefficient(s). The stars ***, **, and * denote statistical significance at the 1%, 5%, and 10% level. The F-statistics are all greater than 1×10^5. See notes of Table A10 for more details.

Table A11: Effects of Other Cost-sharing Terms on Inpatient Mortality

	(1)	(2)	(3)
a_1	299.5***	309.1***	307.3***
	(5.70)	(6.05)	(18.00)
a_2	0.0253***	0.0237***	0.0149***
	(0.0009)	(0.0008)	(0.0010)
Primary Care Copay	-1.818***	-1.602***	-0.4433***
	(0.0777)	(0.0771)	(0.0940)
Benchmark	0.3466***	0.3958***	0.1146***
	(0.0086)	(0.0090)	(0.0120)
Specialist	-0.4254***	-0.4155***	-0.5326***
	(0.0496)	(0.0486)	(0.0651)
Outpatient	-0.0024	-0.0053	-0.0116**
	(0.0050)	(0.0050)	(0.0050)
Inpatient Stay	-0.0807***	-0.0818***	-0.0197***
	(0.0053)	(0.0052)	(0.0054)
Emergency Room	0.4394***	0.3037***	-0.1501***
	(0.0545)	(0.0573)	(0.0505)
Urgent Care	-0.2518***	-0.2621***	-0.1163***
	(0.0339)	(0.0333)	(0.0420)
Ambulance	-0.0325***	-0.0581***	-0.0735***
	(0.0081)	(0.0081)	(0.0095)
Imaging	-0.0281***	-0.0294***	0.0107
	(0.0065)	(0.0064)	(0.0075)
Medical Device Coins.	-2.195***	-2.156***	0.0428
	(0.0995)	(0.0976)	(0.1151)
Outpatient Coins.	-1.061***	-0.9652***	-0.6168***
	(0.0707)	(0.0696)	(0.0697)
Outpt. Drugs Coins.	-0.1582*	-0.1021	-0.1053
	(0.0896)	(0.0878)	(0.0916)
Out-of-Pocket Max	-0.0057***	-0.0063***	-0.0030***
	(0.0003)	(0.0003)	(0.0004)
Fixed-effects			
Year		✓	✓
Product			✓
Observations	9,789	9,789	9,789

Note: This table displays estimates for the rebate policy function. All variables are measured in dollars, with the exception of coinsurance rates, which are measured in percentage points. Specification (3) is used in the counterfactual exercises in Section 7. The significance stars ***, **, and * denote statistical significance at the 1%, 5%, and 10% level respectively.

Table A12: Estimated Plan Bid Function

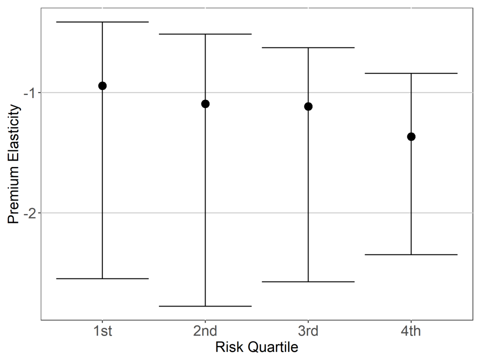

(a) Premium Elasticity and Risk Quartile

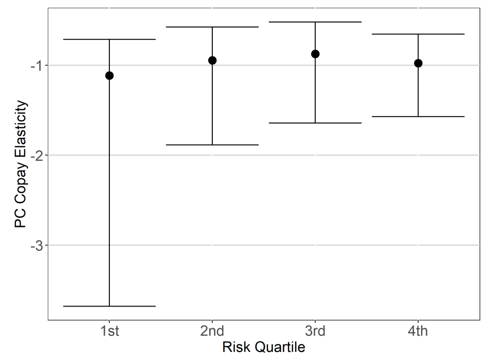

(b) Primary Care Copay Elasticity and Risk Quartile

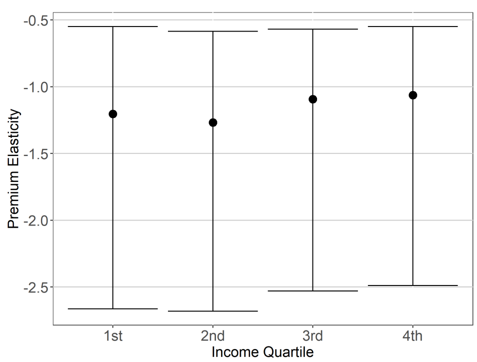

(c) Premium Elasticity and Zip Code Income

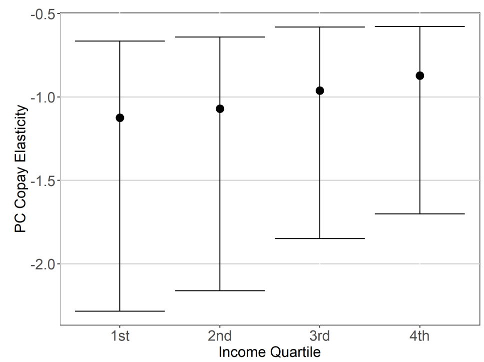

(d) Primary Care Copay Elasticity and Zip Code Income

Figure A1: Elasticity and Observable Heterogeneity

Note: This plots within and across group heterogeneity in the elasticities for premiums and primary care copays. The points represent the median elasticity within each group, and the error bars represent the inter-quartile range. The left column plots premium elasticities and the right panel plots primary care copay elasticities. The order statistics are computed at the consumer level, after integrating over unobserved heterogeneity and averaging across predicted product choices.

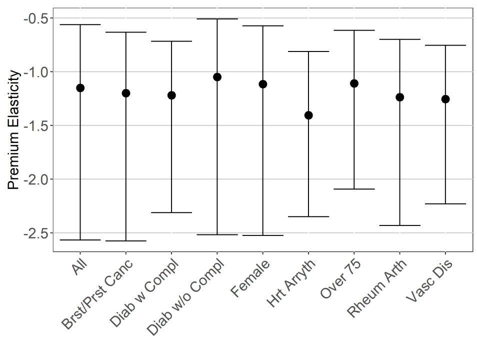

(a) Premium Elasticity and Disease Indication

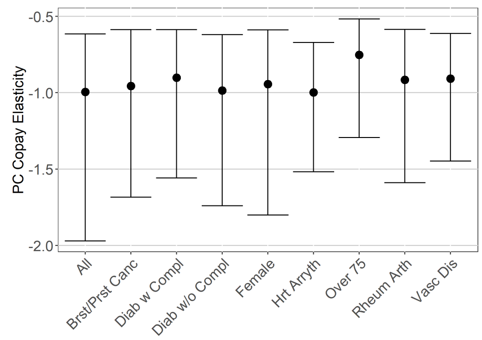

(b) Primary Care Copay Elasticity and Disease Indication

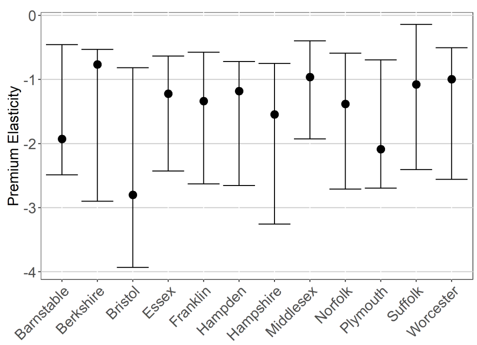

(c) Premium Elasticity and County

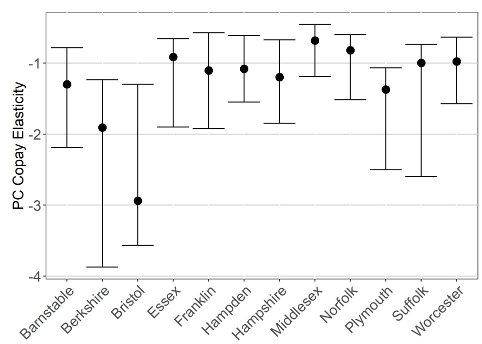

(d) Primary Care Copay Elasticity and County

Figure A2: Elasticity and Observable Heterogeneity (continued)

Note: This plots within and across group heterogeneity in the elasticities for premiums and primary care copays. The points represent the median elasticity within each group, and the error bars represent the inter-quartile range. The left column plots premium elasticities and the right panel plots primary care copay elasticities. The order statistics are computed at the consumer level, after integrating over unobserved heterogeneity and averaging across predicted product choices.

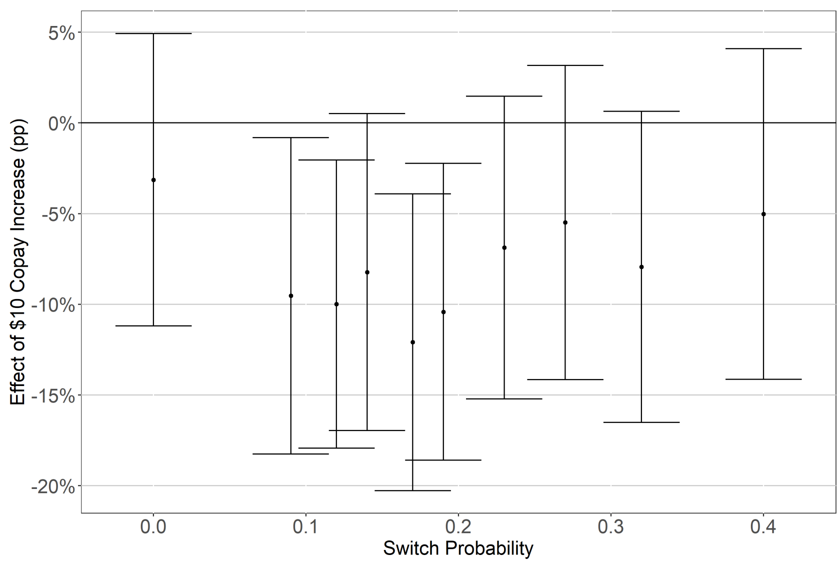

Figure A3: Medical Consumption is Elasticities and Switching Probability

Note: This figure shows the results of the medical consumption demand estimation including an interaction effect with the decile of switching probability. The outcome is insurance-covered consumption, as defined in Section 6.1. The primary care copays are denominated in \$10 and effect sizes are semi-elasticities shown in percentage points. Each point represents the effect of a change in primary care copay for consumers in a particular decile of switching probability, plotted by the minimum probability within that decile. Confidence intervals are shown at the 5% level for a hypothesis test of each decile relative to a null hypothesis of zero. For more details, see the notes of Table 5.

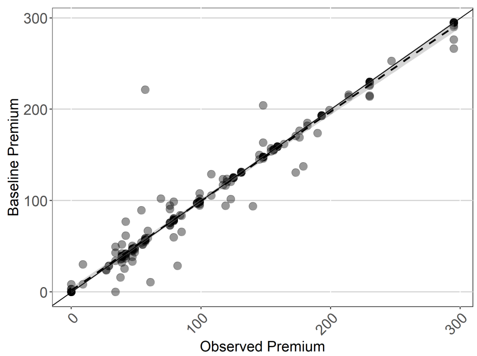

(a) Model Fit for Premiums

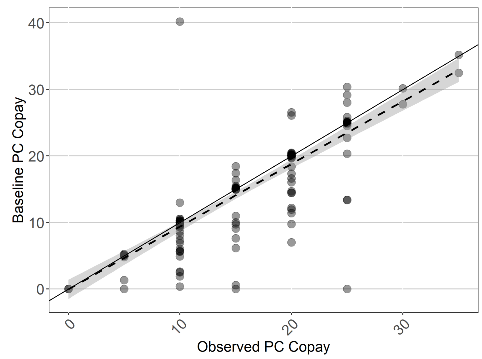

(b) Model Fit for Primary Care Copays

Figure A4: Model Fit

Note: This figure plots the relationship between the observed strategic variables and the baseline equilibrium. Panel (a) plots the baseline and observed premiums. Panel (b) plots the baseline and observed primary care copays. The dotted line represents the estimated linear relationship. The solid line traces the 45 degree line.

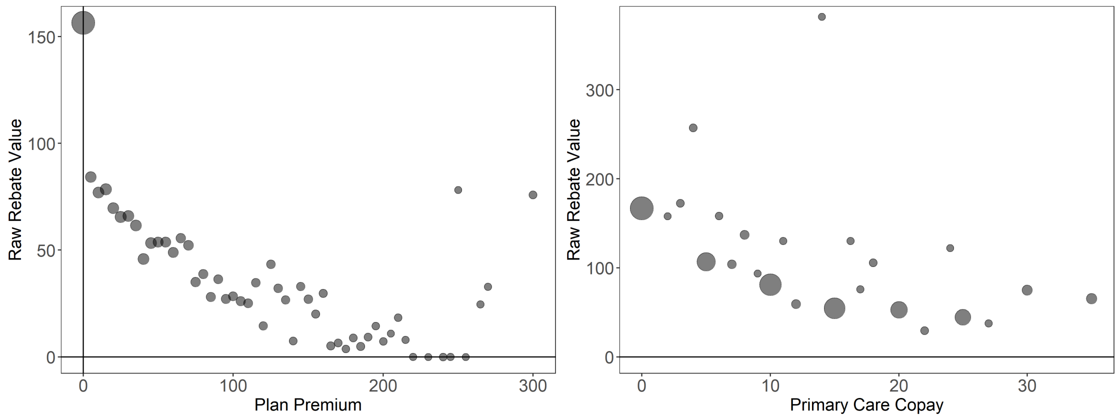

(a) (b)

Figure A5: Rebates, Premiums, and Primary Care Copays

Note: This figure contains a binned scatter plot of the average raw rebate values. The premiums in panel (a) are grouped at $5 increments and capped at $300. The primary care copays are grouped at their integer values. The size of the dots represent the number of products in the computed average.