Going for Broker?
Intermediation in Health Insurance Markets*

Anran Li Tong Liu Anthony LoSasso Nicholas Tilipman†

January 29, 2026

Preliminary – Please Do Not Circulate.
Click here for the latest version.

Abstract

This paper studies how insurance brokers affect product choices, premiums, and welfare in the employer-sponsored insurance market. We compile a novel database of contracting relationships among employers, brokers, and insurers in New York State. Exploiting variations in commission schedules, we document two market distortions: First, brokers exhibit traditional agency frictions, steering employers towards more financially lucrative products. Second, commission levels affect ex-ante insurer-broker networks and, in turn, insurers' competitive pressure, leading to anti-competitive distortions. We develop and estimate a structural model of employer insurance demand, insurer pricing, and formation of broker-insurer contracting networks. We use the model to study a commission-cap counterfactual. A one-percentage-point cap reduces broker-induced steering and raises employer surplus by 3%, but the resulting reduction in insurer competition lowers surplus by over 6%, yielding a net decline of about 3%.

JEL Classification Codes: C78, I11, L13

*For helpful comments, we thank seminar participants at AHRQ, UW-Madison, WashU-Olin, Caribbean Health Economics Symposium 2024, ASHEcon 2025, Midwest Health Economics Conference 2025, Southeastern Health Economics Study Group 2025, Chicago Area IO Day 2025, PUC-Chile, Universidad de los Andes. Liu gratefully acknowledges funding support from MIT Sloan Health Systems Initiative (HSI) and MIT Sloan Junior Faculty Research Assistance Program. Tilipman acknowledges funding from the Johns Hopkins Business of Health Initiative Pilot grant and Carey Business School General Research Support Fund.

†Li: Department of Economics, University of Minnesota - Twin Cities, anranli@umn.edu. Liu: Sloan School of Management, MIT, tongl@mit.edu. LoSasso: La Follette School of Public Affairs, University of Wisconsin-Madison, tony.losasso@wisc.edu. Tilipman: Carey School of Business, Johns Hopkins University, tilipman@jhu.edu.

1. Introduction

Intermediaries play a central role in many markets —including finance, housing, used cars, and health insurance —serving as skilled experts who assist in consumer decision-making. While intermediaries can reduce consumer search costs, policymakers have increasingly scrutinized their role due to concerns about conflicts of interest and commission-based steering that may distort market outcomes.[1] In several insurance markets where intermediaries interact directly with consumers, state and federal policymakers considered or enacted caps on broker compensation as a means of curbing these agency problems.[2]

This paper examines the welfare effects of intermediaries in a critical yet understudied setting: health insurance brokers in the employer-sponsored insurance (ESI) market. The ESI market covers over 175 million Americans —nearly half the US population—and accounts for approximately $1.2 trillion in annual health care spending, reflecting approximately 30% of national health expenditures (CMS, 2024). Over the last decade, per-member cost growth for ESI has reached 80%, outpacing that of Medicare and Medicaid (KFF, 2025). A central question is how much of this cost growth reflects frictions in how insurance options are intermediated and chosen by firms, rather than underlying medical costs.

We study how brokers influence employer insurance choices, insurer competition, and ultimately consumer welfare. Unlike intermediaries in Medicare Advantage or Affordable Care Act exchanges, ESI brokers rarely interact directly with employees. Instead, they work with firms' human resources departments, introducing an additional layer of agency between insurers and end consumers. Market outcomes therefore depend not only on consumer preferences, but also on the incentives and contracting behavior of employers, brokers, and insurers. We develop and estimate a model of employer plan choice and endogenous broker–insurer contracting networks to evaluate how these relationships shape prices and welfare, and to assess the consequences of regulating broker compensation.

While existing studies primarily focus on *steering distortion* where commission incentives lead intermediaries to recommend suboptimal products, we identify an additional and novel *anti-competitive distortion*, through which brokers affect welfare by altering the structure of competition among insurers. In ESI, brokers must be licensed to sell carriers, and they contract selectively with only a subset of available insurers in local geographic markets. This selective contracting creates fragmented choice sets that reduce competitive pressure among insurers, leading to higher equilibrium prices even in the absence of steering. This poses a fundamental challenge for regulators: while capping commissions, a common policy approach as described above, might mitigate steering distortions, it would not necessarily do so for anti-competitive distortions.

We begin with a stylized model to illustrate how steering and anti-competitive distortions arise and interact. The key conceptual insight is that these two distortions operate in tension: higher commissions paid by one insurer relative to another exacerbate steering distortions by increasing brokers' incentives to direct employers toward financially lucrative rather than welfare-maximizing products. However, those higher commissions simultaneously encourage brokers to contract with more insurers to capture these commis-

[1] For example, the Medicare Advantage market recently adopted new restrictions on broker compensation (Commonwealth Funds, 2025), while financial markets have implemented fiduciary duty requirements (New York Times, 2020).

[2] The Center for Medicare and Medicaid Services (CMS) recently published a rule that overhauls the regulation of agent and broker compensation in Medicare Advantage and Medicare Part D plans (Sidley Austin LLP, 2024). New York State similarly banned "contingent" commissions in 2005, though the policy was later reversed (Reuters, 2008).

sions, thereby intensifying upstream competition and reducing prices. This tradeoff implies that welfare-optimal regulation cannot focus solely on eliminating steering, as doing so may inadvertently exacerbate the level of competition. The welfare effects of commission regulations depend on how brokers' steering and contracting responses interact with insurers' pricing strategies in equilibrium.

We test these mechanisms using the first comprehensive dataset on the insurer-broker-employer relationship in the New York ESI market. We focus on small-to-medium-sized fully-insured employers, and document their relations with ten major insurers and over seven hundred brokers. The data combine multiple sources: (i) employer-level Form 5500 filings from the Department of Labor with detailed information on insurers, brokers, premiums, and enrollees; (ii) commission schedules obtained through Freedom of Information Act (FOIA) requests and government portals, which identify brokers' financial incentives; and (iii) novel records on broker–insurer contracting networks, including which brokers are licensed to sell for which insurers and the timing of these relationships. We supplement these data with annual enrollment records from a large national insurer and claims data from the Clarivate database.

We proceed to test for the presence of distortions. To identify steering distortions, we exploit sharp discontinuities in broker commission schedules using regression kink designs. Critical to this design is the fact that commission schedules vary *within-insurer* by the number of enrolled employees, thereby mitigating concerns of unobserved insurer quality correlated with commissions. We find no evidence of manipulation of the running variable (number of enrolled employees) around the kink cutoff. However, discrete reductions in marginal commission rates lead to systematic shifts in the choice of insurer identity. For an additional ten enrolled employees, the probability that the employer-broker pair chooses the insurer with commission schedule drop is 3% lower above the commission change cutoff than below the cutoff. Average premiums are lower above the cutoff, and within-employer take-up rates of insurance are higher, suggesting that higher compensation may induce brokers to steer employers towards more expensive and less preferable products.

Turning to anti-competitive distortions, we first investigate how financial incentives affect brokers' contracting networks. We leverage a quasi-experiment in which one insurer raised its per-member per-month (PMPM) commission from \$5 to \$18 and later reduced it to \$12, while other insurers' commissions remained stable. Using an event study design, we document that a 33% decline in commissions results in a 6% reduction in the share of brokers contracting with that insurer, and the increase in commissions has an opposite effect. Using this variation as an instrument for broker–insurer links, we next show that narrower broker contracting networks, i.e., a smaller number of insurers in employers' choice sets, lead to higher per-member premiums. Relative to a control group of brokers that did not contract with the insurer and unexposed to the commission change, a reduction in the contracting network raises per-member annual premiums by about \$1,000 across competing insurers, consistent with softened competition.

To quantify the welfare implications of these competing effects and simulate counterfactual regulations, we next develop and estimate a structural model of demand, pricing, and network formation. The model features four stages. First, given posted commission schedules, brokers select which insurers to include in their product offerings and incur the associated contracting costs. Second, employers quasi-randomly meet with a broker and decide whether to use the brokers' services. This step quantifies how brokers reduce employers' hassle costs of offering insurance plans. Third, insurers simultaneously set product premiums for

employers, depending on whether the employer quotes through a broker, and if so, the broker's contracting set. This captures how selective contracting affects insurer competition. Fourth, employers choose insurers, either from the broker's product offering or from the available set of products on the market, if not working with a broker. Employers weigh their own premium expenses against their employee's surplus, and this relative weights are affected by brokers' commission payments. This step measures how brokers impact match quality between employers and insurers, and the degree that brokers steer consumers towards less-preferred products.

We estimate the model sequentially. On the demand side, we recover employers' preferences using maximum likelihood estimators. By rationalizing differences in employer choice of plan following plausibly exogenous variation in within-insurer commissions, we identify steering, or broker influence on employer product offerings, separately from insurer brand effects. Hassle costs are identidied by comparing firms' propensities to work with brokers across similar geographic markets. On the supply side, we recover insurers' marginal costs using first-order conditions consistent with profit maximization. We estimate brokers' contracting costs using moment inequalities that exploit revealed preferences: adding or subtracting an insurer results in profit losses compared to the status quo contracting network.

Model estimates yield several key findings on the role of brokers. First, employers without brokers value $1 of employee welfare at $1.33 relative to their own premium expenses. In other words, they pay $1.33 to generate $1 in value for an average employee.[3] Second, brokers steer employers towards more generous plans, consistent with our reduced-form results. Brokers increase employers' willingness to pay by $0.71–$0.88 for the same $1 of employee surplus, of which one-fifth is attributable to commission-induced steering. We estimate that a one-percentage-point decrease in commissions reduces the probability of choosing the high-commission insurer by roughly 3 percentage points. On average, brokered employers pay 60% higher premiums than non-broker firms for the same level of employee surplus. Third, brokers provide real value: they reduce employers' hassle costs of offering insurance by $30 PMPM, or 7.8% of the average employer premium share. For a typical 200-employee firm in New York City, this amounts to approximately $71,000 per year, approximately the average salary of hiring an benefits consultant employee. Finally, brokers' selective contracting imposes sizable premium costs. Compared with full-menu contracts, exclusive constracts allows insurers' markup to raise by about 40% for the same client.

Using the estimated model, we evaluate proposed regulations, including caps on commissions and prohibitions on exclusive contracts between brokers and insurers, taking into account both steering and anti-competitive distortions. We first use partial-equilibrium counterfactuals to isolate the welfare effects of broker-induced steering and selective contracting, holding employer-broker matches fixed. Removing commission-driven steering improves plan matching and lowers prices: employee surplus rises by about 6%, employer surplus increases by roughly 2-3%, and average enrollee premiums fall modestly. When insurers are allowed to reoptimize prices in response to the absence of steering, competitive pressure intensifies, leading to a further premium reduction of about 5% and larger gains in both employee and employer surplus. We also examine selective contracting by expanding employers' choice sets to include all insurers active in a market. Broader menus raise employer surplus by more than 10% and modestly increase

[3]This is fairly consistent with the magnitude of the tax exclusion on health insurance, which amounts of 30% on average.

employee surplus, reflecting increased product variety and stronger price competition.

We then study a policy counterfactual that imposes a one-percentage-point cap on broker commissions, allowing different margins of adjustment to operate sequentially. The cap weakens brokers' incentives to steer employers toward high-commission plans, improving allocative efficiency and initially raising employer surplus by about 3% through better choices and lower insurer markups. However, once brokers are allowed to adjust their menus, the cap reduces the profitability of some insurer-broker relationships, leading brokers to narrow their networks. This contraction amplifies anti-competitive forces, raises prices in broker-mediated contracts, and ultimately lowers employer surplus by more than 6%. In equilibrium, the commission cap produces a net decline in employer surplus of roughly 3%, highlighting a trade-off between reducing steering distortions and preserving competitive broker networks.

Our results provide new insight into the mechanisms driving rising health insurance costs in the US and have important policy implications. First, our findings speak directly to the design of intermediary regulation in markets such as Medicare Advantage, where third-party brokers are already subject to extensive federal oversight. Second, they inform antitrust policy by highlighting a vertical friction —selective contracting between insurers and brokers —that can suppress competition and contribute to higher prices. More broadly, our analysis illustrates how intermediaries, while acting as agents for buyers, can also reshape market structure and competition upstream, amplifying or mitigating welfare losses depending on how their incentives are regulated.

Related Literature. This paper contributes to several literatures. First, we contribute to research on information gaps and decision support in health insurance (Gruber et al., 2021; Li et al., 2022; Myerson and Li, 2022; Tong, 2022). While existing research has shown consumers face frictions when choosing insurance products (Abaluck and Gruber, 2015; Handel, 2013; Handel and Kolstad, 2015), we show that employers are no exception. In addition to analyze the effect on choice quality, we present the first equilibrium analysis on how intermediaries may influence downstream choices and upstream competition.

Second, our work speaks to the health economics literature on employer-sponsored insurance and market power (Bundorf et al., 2012; Craig, 2022; Fleitas et al., 2025; Ho and Lee, 2023; Marone and Sabety, 2022; Tilipman, 2022). We provide the first evidence on how intermediaries affect competition in health insurance markets, complementing existing work on insurer-provider bargaining and employer plan design.

Third, we add to the growing body of work on intermediation in financial and insurance markets (Bhattacharya et al., 2019; Boehm, 2023; Grennan et al., 2024; Salz, 2022; Barwick et al., 2017; Robles-Garcia, 2019; Kim, 2024). While previous studies have primarily focused on agency problems between intermediaries and consumers, we highlight how intermediaries affect upstream competition through their product assortment decisions. This mechanism may be relevant in many markets where intermediaries maintain selective relationships with suppliers, for example, in the wholesaler-retailer-consumer context.

Finally, we contribute to the literature on vertical relationships and product variety (Conlon and Mortimer, 2021; Hristakeva, 2022; Lee, 2013). Most work in this area examines retail settings where product assortment is constrained by physical shelf space. We analyze a setting where assortment is endogenously determined by contracting costs and steering incentives, providing new insights into how vertical relationships shape market structure.

Outline. The remainder of the paper proceeds as follows. Section 2 describes the institutional background and our data sources. Section 3 presents a conceptual framework that illustrates how brokers influence product assortment and upstream competition. Section 4 presents reduced-form evidence on steering and anti-competitive effects. Section 5 develops our structural model. Section 6 discusses estimation and identification, and reports estimation results. Section 7 examines policy counterfactuals, and Section 8 concludes.

2. Empirical Setting

2.1. Institutional Background

The ESI Market. Health insurance brokers serve as intermediaries between employers and insurance carriers in the employer-sponsored insurance (ESI) market. ESI remains the primary source of coverage for working-age adults in the US, with 60.4% of individuals under age 65, or roughly 164.7 million people, receiving employment-sponsored health insurance in 2023.[4] To provide this coverage, firms either purchase fully insured group policies or self-fund claims and contract with third-party administrators to manage these claims. Because pricing, benefit design, and provider networks are shaped by both local provider competition and federal and state regulations, health insurance brokerage is inherently localized.

In New York, the setting of this paper, the ESI market operates under close oversight by the Department of Financial Services , and features a mix of national carriers and regionally dominant insurers that operate segmentally across downstate and upstate markets. Large employers are more likely to self-insure, while small and mid-sized firms typically purchase fully insured products. Across firm sizes, however, high health care costs and complex regulatory requirements make skilled intermediation particularly valuable. As a result, brokers are often hired and serve as a critical touchpoint in helping employers navigate plan selection, pricing, and administration.

How Brokers Assist Employers. Brokers assist employers in several ways: they help evaluate and select insurance plans for cost containment, handle administrative tasks related to enrollment and regulatory compliance, and provide ongoing support in benefits management. For larger employers, brokers may also coordinate with pharmacy benefit managers or other specialized vendors. Although some large brokerage firms operate nationally, broker markets are essentially local in scope, and employers generally prefer working with brokers that maintain a physical presence in their communities. Broker-of-record agreements establish the legal authority to represent the employer with carriers, but these agreements are typically structured on a month-to-month basis, even when a fixed annual fee is negotiated.

In addition to selecting insurance plans and handling administrative tasks, brokers can help reduce the insurance costs that employers pay. During our interviews with practitioners, one broker explained: ”The majority of what I do is to use competition to create a marketplace. The initial quote serves as a starting point for negotiations, and I use quotes from other insurers to negotiate reductions... The more you work with, the better your prices, but only up to a point.”

[4] https://www.kff.org/health-costs/health-policy-101-employer-sponsored-health-insurance/?entry=table-of-contents-who-is-covered-by-employer-sponsored-health-insurance

Broker-Insurer Relationships. Broker-insurer contracting further shapes the set of products available to employers. A broker must be licensed in the relevant state and enter a carrier-specific agreement to sell that insurer's plans. Agreements with carriers are typically not exclusive. While brokers may quote plans without a contract, they cannot complete a sale without a formal agreement.

Brokers often choose not to contract with certain carriers when expected client volume is low or when administrative costs are high. This introduces variation in the choice sets available to employers depending on which broker they work with. In practice, brokers tend to favor a limited set of carriers with whom they have long-standing relationships, especially when product features and pricing are broadly comparable. The selective contracting practice is prevalent in the industry, as one broker said: 'There is an optimal number of insurers to contract with—it's not one or two, but it's also not fifteen."

Commissions and Regulations. Brokers are typically compensated by insurers rather than by employers, either through percentage commissions on premiums or through flat fees per enrollee. Carriers generally pass through these payments to employers and have limited incentive to reduce commission levels, as doing so may discourage broker engagement.

In practice, standard commission schedules are usually rigid and rarely vary over time. It applies uniformly across brokers for a given insurer. These schedules are disclosed publicly to all insurance agents and brokers, and in New York, they must also be filed with the Department of Financial Services. Hence, brokers can compare commission structures in advance and know the terms of their compensation even before entering into a contractual agreement with an insurer. Percentage-based commissions are most common for serving small to mid-sized firms, while fee-based contracts are more prevalent among large groups (typically those with over 500 enrollees), reflecting the greater complexity and scale of services provided.[5]

Regulators already intervene to constrain broker compensation schemes. For instance, prior to 2010, New York imposed a cap of 4 percent on commissions for health maintenance organizations (HMOs) and the in-network portion of point-of-service (POS) plans.[6] At the federal level, the Centers for Medicare and Medicaid Services (CMS) sets explicit maximum payments that carriers and Field Marketing Organizations (FMOs) may provide to agents and brokers for enrolling beneficiaries into Medicare Advantage (MA) plans, as specified under 42 C.F.R. §422.2274.[7] These policies illustrate that broker compensation has long been viewed as a lever for influencing market outcomes.

2.2. Data

Our primary data source is the Department of Labor's Form 5500 filings in 2011-2020. Employers that offer health insurance and have at least 100 enrolled employees are required to submit these forms, except for self-insured plans. These filings provide detailed information on employer identifiers, the brokers with whom they contract, the insurance carriers and plan types they select, the number of enrollees, and total

[5]Two other compensation types are feasible but not applicable in our context: Brokers can receive customized commissions for serving large employers, which are negotiable and often depend on group size or premium volume. In addition, brokerage firms can receive contingent commissions tied to their total volume of new business with a carrier. These payments are largely opaque to employers, not disclosed in public filings, and usually apply to large employers outside the scope of our sample.

[6]https://www.dfs.ny.gov/insurance/ogco2010/rg100603.htm

[7]https://ritterim.com/blog/2025-maximum-broker-commissions-for-medicare-advantage-and-medicare-part-d/#hospital-indemnity-guide

premiums paid for each carrier-year, as well as employers' characteristics, such as firm size and industry. Our analysis focuses on fully insured employers in New York, the majority of which have between 100 and 500 employees, with a small number in the 500–1000 range. This data forms the backbone of our analysis by allowing us to link employers, brokers, and insurers over time. Appendix B1 describes our data construction process in detail.

A second key data source consists of commission schedules, which we obtained through a Freedom of Information Act (FOIA) request to New York State (NYS). Under New York Insurance Law, for instance, Section 4235, insurers must file their commission schedules and rate manuals with the NYS Department of Financial Services to ensure compliance with compensation limits. In total, we obtained 76 documents covering all major health insurers in NY over the past two decades. It takes hundreds of hours to extract and hand-code the key variables in the compensation scheme for insurers across years. The records detail each carrier's standard commission rates offered to brokers. Importantly, the schedules exhibit systematic variation along several dimensions, including insurer identity, employer size, and the total volume of premiums sold. They also differ in structure, taking either a flat-fee or a percentage-of-premium form. This heterogeneity generates meaningful cross-sectional and temporal variations in broker incentives, providing a central source of identifying variation for our empirical strategy.

To construct broker-insurer contracting networks, we draw on two complementary datasets. The first is Discovery Data, a proprietary database that records the set of carriers with which each broker has contracts and is licensed to sell. While this provides a comprehensive snapshot of potential broker affiliations, it does not capture the timing of active contractual relationships. To address this limitation, we supplement Discovery Data with hand-collected records from the NYS Department of Financial Services portal. These records provide detailed information about brokers, including their names, business types (individual or business entity), license numbers, statuses, emails, home states, addresses, and lines of business (e.g., life, accident, health, or property and casualty). Most importantly, the portal also reports the NAIC number and name of insurers with which brokers have ever entered into a contract, along with the start and end dates of those contracts. By combining these sources, we can track the evolution of broker-insurer networks over time, capturing both the extensive margin (i.e., the insurers a broker can represent) and the intensive margin (i.e., when these contractual relationships are active).

Lastly, to obtain detailed enrollment, plan characteristics, and spending information, we use two novel datasets that allow us to observe granular enrollee-level records. The first one is from a large insurer operating in New York State. We observe all employers that choose this insurer, providing information on the number of plans offered, the number of eligible and enrolled employees, as well as plan characteristics such as plan type (EPO, PPO, POS), coinsurance, deductible, out-of-pocket maximum, and total premium. We also observe enrollee-year level demographics and health plan choice from the employer's offerings, as well as linked enrollee claims data. Although this data only represents the plans offered by a single large insurer, it complements the Form 5500 data, which do not contain enrollment records across plans of different characteristics within insurers. This facilitates our analysis of employee preferences for insurance products. Appendix B2 reports further details on the dataset and summary statistics.

We also draw on the proprietary Clarivate DRG Real World Data (Liu, 2022; Arnold et al., 2025). The

dataset comprises billions of patient claims, providing detailed information on providers, patients, payers, diagnoses, procedures, as well as a granular breakdown of payer and patient payments, including deductibles, coinsurance, and copayments. The ability to track patients and payers over time allows us to construct insurer-market-year-level plan characteristics, such as deductibles and coinsurance, with meaningful variation across insurers and markets. Since the data is sourced from a claims-processing intermediary between insurers and providers, they are broadly representative of plan characteristics across a wide range of commercial insurers. The Clarivate data complements the Form 5500 data, which do not report product characteristics, and allows us to accurately characterize preferences for products from distinct insurers. Appendix B3 provides details on the dataset, sample construction, and the imputation of product characteristics.

3. Conceptual Framework

To illustrate how broker incentives affect menu choices and steering decisions—and how these incentives interact—we present a stylized model.

3.1. Setup

Consider a market with two insurers, A and B, each offering a single product, and a broker who receives flat commissions, γ_A and γ_B.[8]

The game proceeds in three stages. In the first stage, the broker observes commissions[9,10] and chooses one of three possible menus: contracting with both insurers, contracting only with A, or contracting only with B. In the second stage, employers meet with the broker and receive price quotes from insurers. Insurers set prices based on the contracting arrangement: competitive prices p_A^c, p_B^c, and monopoly prices p_A^m, p_B^m when the broker contracts with either insurer exclusively. We assume $p_j^m > p_j^c,\ j \in \{A, B\}$. In the final stage, employers choose among insurers A, B, or remaining uninsured (U).

The flow utility of employers from option $j \in \{A, B, U\}$ under contracting arrangement $l \in \{c, m\}$ is given by:

$$V_j^l = \delta_j - \alpha p_j^l - \rho\gamma_{-j} + \epsilon_j, \tag{1}$$

where δ_j is the employer's brand preferences for product j, p_j^l is the price under contracting arrangement l, and ϵ_j is an independent and identically distributed draw from a type-I extreme-value distribution.

Steering enters employer utility through the term $\rho\gamma_{-j}$, which reduces the perceived utility of the non-steered option when the rival insurer offers a higher commission. The parameter ρ captures brokers' steering ability: a larger ρ implies that commissions exert greater sway over employers' choices through broker recommendations.

[8] The intuition and qualitative predictions still hold if the broker instead receives commissions as a percentage of premiums.

[9] Because the goal of this section is to analyze how regulatory commission caps affect broker menu and employer welfare, we treat commissions as exogenously given and do not model insurers strategically choosing commissions.

[10] If insurers were allowed to strategically set commissions before brokers chose menus, they would compete by raising commissions until the point where one insurer's net profit is driven to zero. This logic is in a similar spirit to manufacturers paying lump-sum transfers to retailers to carry their products (Shaffer, 2005; Hristakeva, 2022). However, such settings raise concerns about equilibrium existence when upstream firms (insurers or manufacturers) are not homogeneous or single-product because best responses may be discontinous (Rey and Vergé, 2010, 2016).

The broker's profit equals the expected commission income minus contracting costs. Depending on the contracting strategy, the broker earns:

$$\Pi_{AB} = s_A^c \gamma_A + s_B^c \gamma_B - 2c, \quad \Pi_A = s_A^m \gamma_A - c, \quad \Pi_B = s_B^m \gamma_B - c,$$

where c is the fixed cost of contracting with a single insurer, and s_j^l is the market share of insurer j under arrangement l.

We assume that employers have inherent preferences for insurer A over B, formalized as

$$s_B^m > s_B^c, \quad s_A^c \gamma_A - c > 0. \tag{2}$$

The former inequality states that when both insurers are available, the probability of choosing B is lower relative to the probability that B would be chosen if offered exclusively (relative to the outside option). The latter inequality states that A is sufficiently attractive that under combined contracting, it remains profitable for the broker even if B offers no commission and incurs no contracting costs.

3.2. Broker Menu Without Steering

We start by considering the scenario of the broker menu without steering. In this baseline case ($\rho = 0$), broker profits respond to commission changes in a straightforward manner: higher commissions increase the expected returns from offering the corresponding insurer. Holding γ_A fixed, we analyze how changes in insurer B's commission, γ_B in the range of $[\gamma_A, \infty]$, affect the broker's menu choice, as illustrated in Figure 1a.

The profit from offering only A is unaffected by B's commission, resulting in a horizontal line. By contrast, the profits from offering A+B or only B increase with B's commission at rates determined by the share of employers choosing B under combined or exclusive contracting, respectively:

$$\frac{\partial \Pi_A}{\partial \gamma_B} = 0, \quad \frac{\partial \Pi_{AB}}{\partial \gamma_B} = s_B^c, \quad \frac{\partial \Pi_B}{\partial \gamma_B} = s_B^m.$$

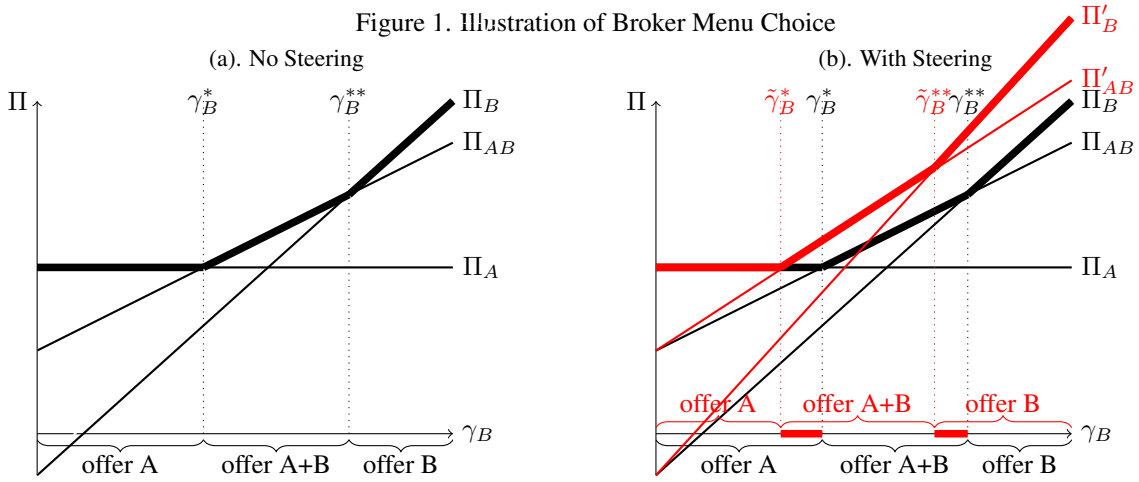
Our assumptions on employer preferences (equation (2)) imply that the slope of the profit line for offering only B is steeper than that for offering both insurance products, while the intercept of offering both products exceeds that of offering only B: $\frac{\partial \Pi_B}{\partial \gamma_B} > \frac{\partial \Pi_{AB}}{\partial \gamma_B}$, $\Pi_{AB}|_{\gamma_B=0} > \Pi_B|_{\gamma_B=0}$.

The broker chooses the menu that maximizes its profit, corresponding to the upper envelope of the three profit lines in Figure 1a. When γ_B is low, the broker offers only A, reflecting both employers' inherent preference for A and the limited expected revenue from B relative to contracting costs. As γ_B rises, the expected revenue from including B rises, and the broker expands the menu to offer both insurance plans once a threshold γ_B^* is reached. At sufficiently high γ_B, expected revenues from offering B dominate, and the broker optimally contracts only with B, forgoing the contracting costs of A and further increasing B's monopoly market share.

This analysis highlights the ambiguous effect of commission capping on brokers' menu choices. Small changes in γ_B may leave contracting decisions unchanged unless they cross thresholds γ_B^* or γ_B^{**}. Starting from a high baseline commission, an exogenous reduction from above to below γ_B^{**} would expand the

broker's menu from B alone to A+B, as catering to employer preferences increases expected profits. Conversely, when the baseline commission is low, an exogenous reduction from above to below γ_B^* would shrink the menu from A+B to only B, as the expected commission income from B is no longer sufficient to justify contracting costs.

Figure 1. Illustration of Broker Menu Choice



Notes: This plot shows how brokers' profits from different product menus change, along with Insurer B's commission payment, γ_B. In panel (b) of the case with steering, Π_B and Π_{AB} can be convex functions of γ_B. For simplicity, we use linear plots to illustrate the idea. We plot in the range where $\tilde{s}_B^m < (1+\rho\gamma_B)/(2\rho\gamma_B)$ and γ_A is sufficiently small to ensure that $\partial\Pi_B/\partial\gamma_B > \partial\Pi_{AB}/\partial\gamma_B$ in the with steering case. Otherwise, at larger values of γ_B, it is possible that $\partial\Pi_B/\partial\gamma_B < \partial\Pi_{AB}/\partial\gamma_B$, such that at some large γ_B, the broker changes from offer B to offer $A+B$.

3.3. Broker Menu With Steering

We next analyze the case with steering, $\rho > 0$. Without loss of generosity, we assume that steering occurs only towards insurer B, relative to A and the outside option U.[11] Because steering reduces the perceived utility of non-steered options, the employer's flow utility when B is offered is

$$V_A^c = \delta_A - \alpha p_A^c - \rho\gamma_B + \epsilon_A; \quad V_U^l = \delta_0 - \rho\gamma_B + \epsilon_0,\ l \in \{m, c\}.$$

When B is absent from the menu, the perceived utility of A or U remains the same as in the baseline, where $\rho\gamma_B$ is not subtracted from the flow utility. Since we hold γ_A fixed, γ_B can equivalently be interpreted as the commission differential $\gamma_B - \gamma_A$, which is greater than or equal to zero.

With steering, commissions affect broker profits in two ways: directly through higher per-employer payments and indirectly by raising the likelihood that employers follow the broker's recommendation to choose the non-preferred products. Denoting variables under steering with tildes, the derivatives of broker profits with respect to B's commission become:

$$\frac{\partial\tilde{\Pi}_{AB}}{\partial\gamma_B} = \tilde{s}_B^c + \frac{\partial\tilde{s}_B^c}{\partial\gamma_B}\gamma_B + \frac{\partial\tilde{s}_A^c}{\partial\gamma_B}\gamma_A = \tilde{s}_B^c + \rho\tilde{s}_B^c(1-\tilde{s}_B^c)\gamma_B - \rho\tilde{s}_A^c\tilde{s}_B^c\gamma_A > \tilde{s}_B^c > s_B^c = \frac{\partial\Pi_{AB}}{\partial\gamma_B}.$$

$$\frac{\partial\tilde{\Pi}_B}{\partial\gamma_B} = \tilde{s}_B^m + \frac{\partial\tilde{s}_B^m}{\partial\gamma_B}\gamma_B = \tilde{s}_B^m + \rho\tilde{s}_B^m(1-\tilde{s}_B^m)\gamma_B > \tilde{s}_B^m > s_B^m = \frac{\partial\Pi_B}{\partial\gamma_B}.$$

[11] The intuition and qualitative predictions still hold if steering occurs towards both directions. In that case, one would need to rescale the brand fixed effects to get similar derivations as reported in the main text.

Thus, steering increases the responsiveness of broker profits to commissions compared to the no-steering case, effectively steepening the slopes of $\frac{\partial \Pi_B}{\partial \gamma_B}$ and $\frac{\partial \Pi_{AB}}{\partial \gamma_B}$.[12]

Figure 1b illustrates how steering affects menu choices. Greater steering ability expands the range of commission over which B is offered, either exclusively or alongside A. This occurs because the broker can more effectively capture higher commissions from B by shifting employers away from their natural preference for A.

This analysis reveals that the effect of steering on brokers' menu choices is ambiguous. If the baseline commission level lies between γ_B^* and $\tilde{\gamma}_B^*$, the additional revenue from shifting employers to B outweighs contracting costs, expanding the menu from A to A+B. If the baseline commission level lies between γ_B^* and $\tilde{\gamma}_B^*$, revenue from B is high enough that the broker forgos contracting costs with A, shrinking the menu from A+B to B. A higher steering ability ρ increases the range of commissions where menu changes may occur, amplifying the sensitivity of broker contracts to commissions.

3.4. Welfare and Regulation

We now consider employer surplus, taking into account the broker's menu choice. Expected surplus under each broker menu choice follows a standard log-sum format (ignoring the constant term):

$$\begin{aligned}
\text{CS}_{A+B} &= \ln\left(\exp(\delta_A - \alpha p_A^c - \rho\gamma_B) + \exp(\delta_B - \alpha p_B^c) + \exp(\delta_0 - \rho\gamma_B)\right), \\
\text{CS}_A &= \ln\left(\exp(\delta_A - \alpha p_A^m) + \exp(\delta_0)\right), \\
\text{CS}_B &= \ln\left(\exp(\delta_B - \alpha p_B^m) + \exp(\delta_0 - \rho\gamma_B)\right).
\end{aligned}$$

Comparing the A-only and B-only menu is straightforward: employer surplus is higher with the A-only menu, both because employers prefer A to B ($\delta_A > \delta_B$) and because the B-only menu imposes a steering penalty ($-\rho\gamma_B$) on non-steered options.

Assessing the welfare effect of moving from a single-product to a multi-product menu is more nuanced. We compare employer surplus at the commission level $\tilde{\gamma}_B^*$, where the broker is indifferent between offering A+B versus only A. The difference in surplus can be decomposed into two components:

$$\begin{aligned}
\text{CS}_{A+B} - \text{CS}_A = &\underbrace{\ln\left[\frac{\exp(\delta_A - \alpha p_A^c) + \exp(\delta_B - \alpha p_B^c) + \exp(\delta_0)}{\exp(\delta_A - \alpha p_A^m) + \exp(\delta_0)}\right]}_{\text{Competitive Effects (+)}} \\
&+ \underbrace{\ln\left[\frac{\exp(\delta_A - \alpha p_A^c - \rho\tilde{\gamma}_B^*) + \exp(\delta_B - \alpha p_B^c) + \exp(\delta_0 - \rho\tilde{\gamma}_B^*)}{\exp(\delta_A - \alpha p_A^c) + \exp(\delta_B - \alpha p_B^c) + \exp(\delta_0)}\right]}_{\text{Steering Effects (-)}}. \qquad (3)
\end{aligned}$$

The **competitive effect** captures the welfare gain from introducing B to compete with A. It can be

[12] The first inequality in the $\partial\tilde{\Pi}_{AB}/\partial\gamma_B$ expression holds, because $\rho\tilde{s}_B^c(1-\tilde{s}_B^c)\gamma_B - \rho\tilde{s}_A^c\tilde{s}_B^c\gamma_A = \rho\tilde{s}_B^c(1-\tilde{s}_B^c)\gamma_B - \rho(1-\tilde{s}_B^c-\tilde{s}_U^c)\tilde{s}_B^c\gamma_A = \rho\tilde{s}_B^c(1-\tilde{s}_B^c)(\gamma_B-\gamma_A) + \rho\tilde{s}_U^c\gamma_A > 0$.

further decomposed into:

$$\underbrace{\ln\left[\frac{\exp(\delta_A-\alpha p_A^c)+\exp(\delta_0)}{\exp(\delta_A-\alpha p_A^m)+\exp(\delta_0)}\right]}_{\text{Price Competition (+)}}+\underbrace{\ln\left[\frac{\exp(\delta_A-\alpha p_A^c)+\exp(\delta_B-\alpha p_B^c)+\exp(\delta_0)}{\exp(\delta_A-\alpha p_A^c)+\exp(\delta_0)}\right]}_{\text{Product Variety (+)}}>0.$$

The first term represents the price competition effect: the welfare gain from A's price falling from monopoly (p_A^m) to competitive levels (p_A^c). The second term represents the product variety effect: the welfare gain from adding insurer B as a new option to the choice set. Both components are always positive.

The **steering effect** represents the welfare loss from directing employers away from their true preferences toward the higher-commission product. We can rewrite it as:

$$\ln\left(\frac{\exp(-\rho\tilde{\gamma}_B^*)(\exp(\delta_A-\alpha p_A^c)+\exp(\delta_0))+\exp(\delta_B-\alpha p_B^c)}{\exp(\delta_A-\alpha p_A^c)+\exp(\delta_0)+\exp(\delta_B-\alpha p_B^c)}\right)=\ln\left(1-(s_A^c+s_U^c)(1-\exp(-\rho\gamma_B^*))\right)<0,$$

where s_A^c, s_U^c are the share of employers choosing A or U without steering. This expression shows that the welfare loss from steering is proportional to both brokers' steering incentive $\tilde{\gamma}_B^*$, steering ability ρ, as well as the baseline share of employers preferring a non-steered option, $s_A^c + s_U^c$. The steering effect is always negative, representing a pure deadweight loss from distorted choices.

Menu expansion improves welfare when the competitive effect dominates the steering effect. Figure 2a illustrates this case. In the leftmost region, when the broker chooses an A-only menu, employer surplus is flat as B's commissions have no effect. At γ_B^*, there is a discrete jump in surplus driven by the introduction of B and the reduction in prices. As γ_B increases further, surplus slopes downward within the A+B region because higher commissions increase the steering distortion. When γ_B becomes very high, such that the broker contracts only with B, surplus drops below the A-only level and continues to decline as employers incur a steering penalty of the outside option.

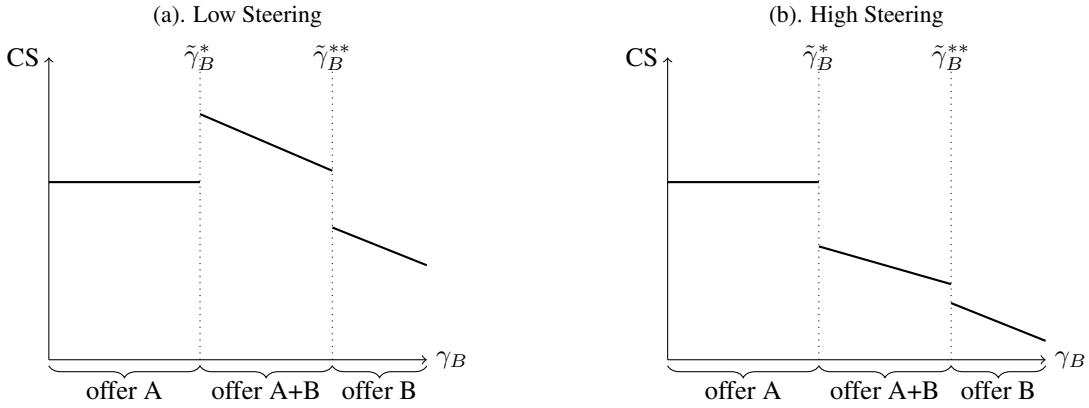
From a regulatory perspective, the B-only menu is clearly dominated. Consumers perceive no benefits from competition and suffer a substantial steering penalty. Regulating B's commissions down improves welfare significantly, and the optimum occurs at γ_B^*, when the broker offers A+B. Further caps on commissions would cause the broker to drop B from the menu, leaving only the most preferred option, but would raise monopoly prices and reduce employer surplus.

In contrast, when steering ability is very high, steering distortions can outweigh competitive benefits of menu expansion. Figure 2b illustrates this case, where the surplus function has a much lower intercept when offering $A + B$. Here, the welfare-maximizing commission is very low, so the broker offers only A. This highlights that the desirability of broker intermediation depends critically on the relative magnitudes of competitive and steering effects: with high steering ability, policies that expand menus and increase competition may actually reduce welfare due to severe distortions in employer choice.

3.5. Summary

Our analysis identifies three key forces shaping equilibrium contracts and welfare. The first force is brokers' contracting response to commissions. Brokers select which insurers to include in their networks, and

Figure 2. Employer Surplus Conditional on Broker Menu



their decisions reflect both employer preferences and contracting costs. The second force comes from insurers' pricing decisions in response to contracting networks. Insurers set premiums strategically, taking into account both the price sensitivity of employers and their own marginal costs. The third force is brokers' steering response to commissions, namely, how much brokers may direct employers toward particular plans due to variations in commissions.

Notably, the steering effect (the third force) operates in tension with the competition-enhancing impacts (the first and second forces). This tradeoff represents a fundamental tension in the market design, as policies that reduce steering may simultaneously weaken competitive pressures, and vice versa. In what follows, we first present reduced-form evidence of each force at work, and then quantify their relative magnitudes in a structural model.

4. Evidence of Steering and Anti-Competitive Distortions

As analyzed in Section 3, the equilibrium broker-insurer contracts and employer welfare depend on three key forces: (1) the degree to which brokers steer employers toward financially lucrative products, (2) the responsiveness of broker-insurer contracts to changes in commissions, and (3) the impact of these networks on upstream competition among insurers. In this section, we present reduced-form evidence to show the empirical relevance of each of these forces.

4.1. Descriptive Patterns

Employers. Table 1 reports sample statistics for NY employers by broker status. Three main patterns emerge. First, fully insured employers, the focus of this paper, account for 74% of employer-year pairs and 81% of total premiums, but only 34% of covered lives. This pattern reflects the fact that small employers are more likely to offer fully insured plans: among such firms, 84% have fewer than 500 enrollees, while only 5% exceed 1,000. Second, the vast majority of fully insured employers (79%, or 13,368 out of 16,396) use brokers. Brokered employers are smaller on average, with 378 enrolled employees and a 78% take-up rate, compared to 605 employees and a 75% take-up rate for non-brokered firms. Most of these employers (87%) offer plans from only one insurer. Third, premiums differ systematically by broker status. The average annual premium per enrollee is $7,495 for brokered firms versus $6,563 for non-brokered firms.

The $932 difference in premium exceeds the average commission payments ($201), suggesting that brokered employers systematically select different insurers and products than employers who do not use brokers.

Table 1. Employer Sample Statistics

	All employers	Fully-insured employers		
	All	All	With broker	No broker
Sample share, by empl-yr pairs	100%	74%	63%	16%
Sample share, by enrollee	100%	34%	24%	10%
Sample share, by premium	100%	81%	58%	22%
Employer-year pairs	22,150	16,396	13,368	3,621
Average number of enrollees	963	442	378	605
Average takeup rate	77%	77%	78%	75%
Average number of carriers offered	1.2	1.2	1.2	1.4
share only one carriers offered	86%	87%	86%	73%
Average premium per enrollee	3,043	7,213	7,495	6,563
Average broker payments per enrollee	57	140	201	-

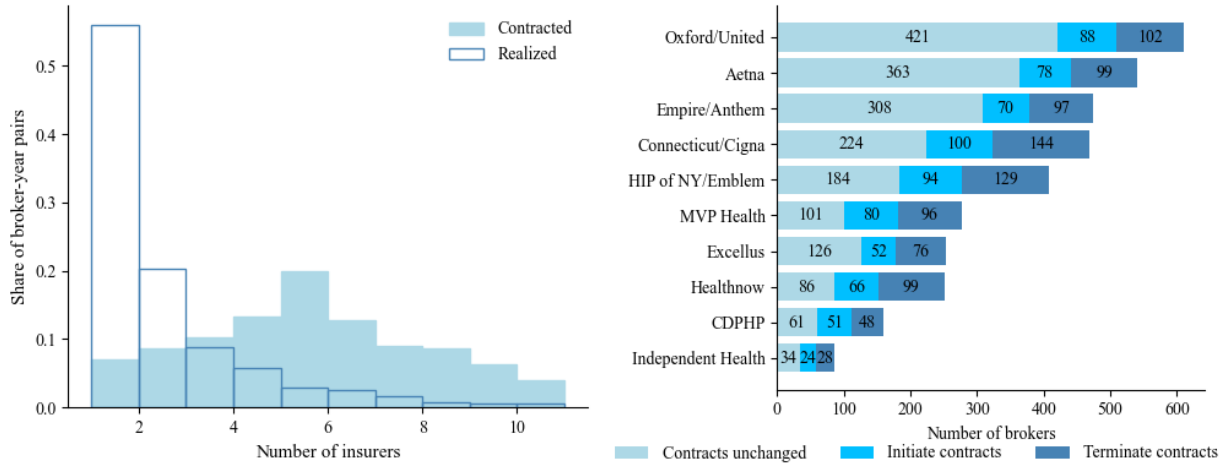
Insurers. Table 2 reports descriptive statistics for the major insurers in our sample and highlights four main patterns. First, the insurer market is relatively concentrated: the top four carriers account for 73.3% market share. As a result, our structural analysis focuses on the ten largest insurers. By contrast, the broker market is far more decentralized, with more than 700 brokers serving in the NY ESI markets. Second, insurer choices differ systematically between employers that work with brokers and those that do not. While the overall distribution across insurers is broadly similar, the top insurer captures a larger market share among brokered firms, whereas lower-cost insurers account for a greater share among non-brokered firms. This may reflect brokers' ability to match employers with higher-quality insurers, but it may also indicate that brokers steer employers to higher-priced insurers. Third, insurers differ markedly in both pricing and contracting strategies. Average annual per-member premiums range from $5,268 to $9,412, and broker networks differ widely: some insurers contract with more than 500 brokers, while others engage only a handful. Notably, the largest insurers, such as Oxford, tend to maintain broader broker networks and charge higher premiums. Finally, not all contracted brokers ultimately recommend an insurer's products. For example, of Oxford's 551 contracted brokers, only 319 ultimately placed clients into an Oxford plan.

Table 2. Insurer Sample Statistics

	Market Share		Premium	Number of brokers	
	No broker	With broker	With broker	Contracted	Recommending
Oxford/United	23.1%	25.3%	9,341	551	319
Aetna	13.3%	13.1%	7,242	474	184
Excellus	11.0%	19.7%	6,562	234	129
Connecticut/Cigna	13.4%	15.1%	6,464	401	177
Empire/Anthem	12.9%	11.5%	9,412	417	177
Healthnow	8.5%	5.9%	6,255	229	76
HIP of NY/Emblem	7.7%	5.9%	7,232	341	127
MVP Health	7.4%	6.1%	5,937	251	82
Independent Health	5.8%	2.3%	5,268	71	36
CDPHP	7.6%	3.0%	6,528	132	51
Other	3.8%	4.9%	5,734	-	205

Broker-Insurer Contracts. Figure 3 illustrates the distribution of broker-insurer contracts and highlights three key patterns. First, selective contracting is widespread. Only 3% of brokers contract with all ten major insurers, while about 8% contract exclusively with a single insurer. The median broker contracts with five insurers. Second, the set of insurers that brokers ultimately recommend (white bars) is only a subset of those with whom they contract (blue bars). Brokers recommend far fewer insurers than they contract with, underscoring the value of our data with precise contract initiation and termination dates: treating either all realized or all potential insurers as the relevant choice set would generate substantial measurement error. Finally, the broker-insurer networks evolve considerably over time, as contracts are formed and dissolved. These dynamics provide useful variation for identifying brokers' primitives in our network formation model.

Figure 3. Broker-Insurer Contracts Among Fully-Insured Employers



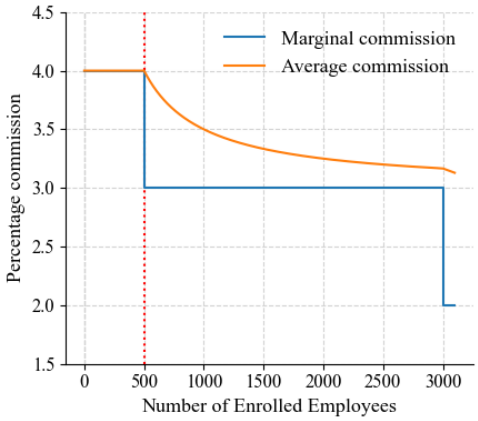
Comission Schedules. Brokers in New York State face heterogeneous commission structures across insurers. Some insurers compensate brokers with a percentage of total premiums, while others pay a flat fee per enrollee. Among those using percentage-based commissions, rates often vary discontinuously with employer size or total premiums. For example, Insurer 1 pays a 4% commission for policies sold with fewer than 500 enrolled employees, but only 3% for additional enrollees beyond that threshold. Similarly, Insurer 2 pays 4% on policies with total premiums below $1 million and 3.5% on premiums exceeding that threshold. Figure 4 presents a detailed example of Insurer 1's commission schedule, and Figure A1 presents the snapshot of commission schedules for all major insurers in our sample in 2017.

4.2. Effects of Commissions on Broker Steering

To examine whether brokers steer employers toward insurance products that offer higher commissions, we exploit the *within-insurer* variation in commission rates at thresholds defined by the number of enrolled employees. As depicted in Figure 4, Insurer 1 reduces its marginal broker commissions from 4% to 3% once enrollment exceeds 500 employees, creating a discrete drop in marginal commissions and a corresponding kink in average commissions. This discontinuity creates two potential channels for broker responses.

The first is that brokers could shift *plan quality* recommendations within Insurer 1, in a similar spirit to Lee and Lemieux (2008): for firms just above the cutoff, brokers may recommend plans with less generous

Figure 4. Commission Schedules, Insurer 1



cost-sharing, higher premiums, or other attributes that employees value less, resulting in fewer employees participating. To test this hypothesis, we implement the density tests proposed by Cattaneo et al. (2018) and McCrary (2008). If brokers manipulate the number of participants on the intensive margin, we would expect a break in the density function of enrollment size at the 500-employee threshold, in other words, "bunching" at the cutoff.

Figure 5a reports the density test results. The distribution of firm size is smooth around the cutoff, suggesting that intensive-margin manipulation is not present. We therefore treat the number of enrolled employees as a valid running variable for a regression discontinuity design.

The second is that brokers may shift the *identity of the insurer* recommended to employers just above or below the commission cutoff. In this case of extensive-margin manipulation, we would expect a kink in the *slope* of the probability of choosing Insurer 1. To test this hypothesis, we focus on the set of employers who have Insurer 1 in their choice sets and employ a regression kink design (RKD):

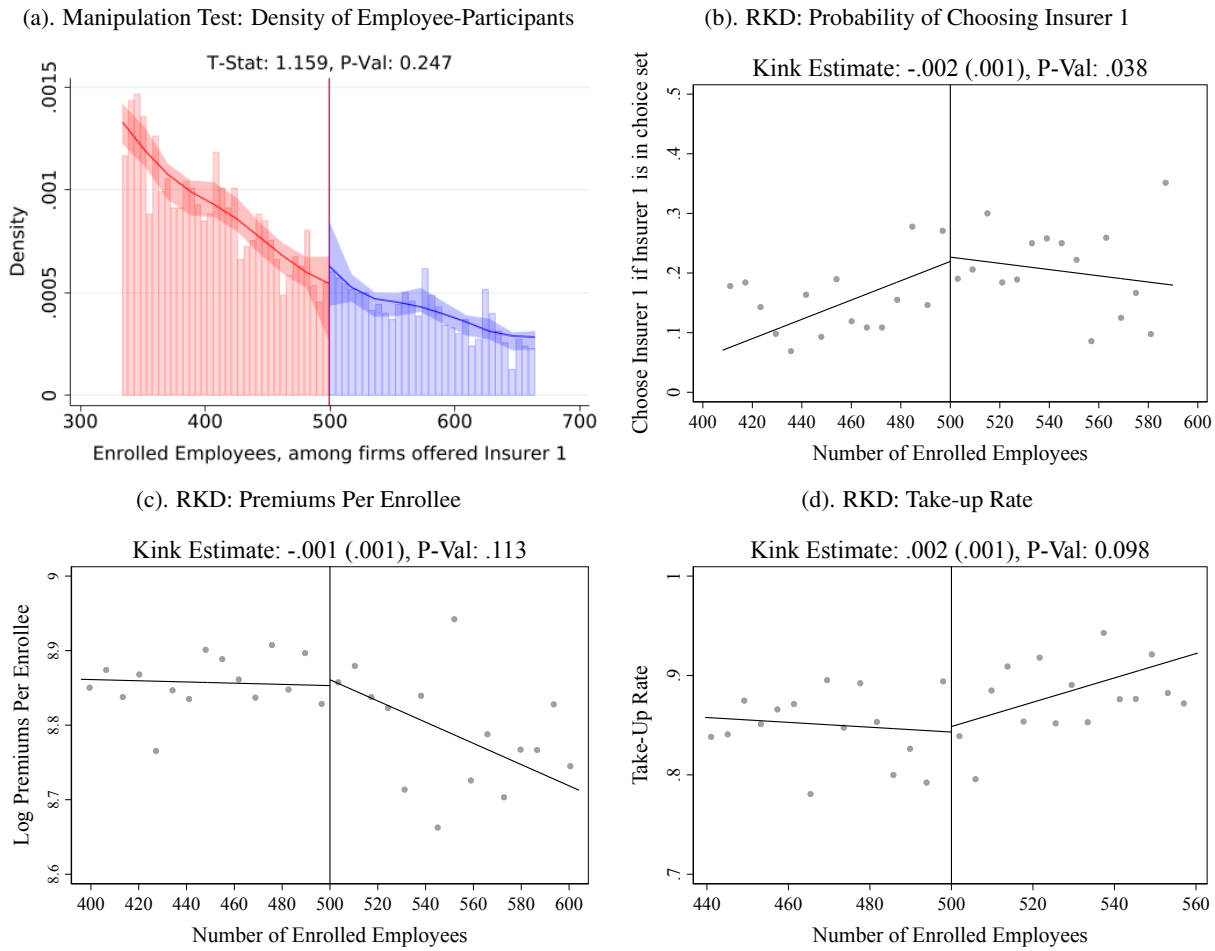
$$Y_f = \alpha + \tau \cdot \mathbb{1}[X_f \geq K] \cdot (X_f - K) + f(X_f - K) + \varepsilon_f, \tag{4}$$

where Y_f denotes whether employer f chooses Insurer 1; K denotes the commission change cutoff (i.e., 500 enrolled employees), and X_f denotes our running variable: the number of enrolled employees in the firm. The coefficient τ is our parameter of interest, which captures the slope change in the probability of choosing Insurer 1 at the threshold. The identifying assumption is that all observables and unobservables are smooth around K, so that any slope break in choice probability can be attributed to the sharp change in commissions. As shown in Figure A2, we fail to reject imbalance of covariates around the threshold at conventional statistical significance levels.

Figure 5b plots the RKD results. The estimated slope change is abrupt and statistically significant: for each additional 10 enrollees beyond the threshold, the probability of offering Insurer 1 declines by about 2% more than it would have at the higher commission level. This provides direct evidence that brokers respond to financial incentives by steering employers towards insurers that compensate them more generously.

We further examine how this extensive-margin steering affects match quality. Steering could, in principle, benefit employers if it directs them to lower-cost insurers or to insurers better aligned with their preferences. However, if insurers with the highest commissions also charge higher markups, steering may instead

Figure 5. Effects of kinked commission on broker recommendations



Notes: This figure depicts brokers' responses to commission incentives, where Insurer 1's commissions drop from 4% to 3% at the cutoff of 500 enrolled employees. The analysis sample is restricted to fully insured employers who have Insurer 1 in their choice set and more than 100 enrolled employees in the carrier of choice. Panel (a) plots the density test of the distribution of the number of employee-participants at the 500 cutoff. Panels (b)-(d) plot the results of a regression kink design, with the number of enrolled employees as the running variable. We use the CER-optimal bandwidth (Calonico et al., 2019).

impose considerable welfare losses. To investigate this, we re-estimate equation (4) using per-enrollee premiums and plan take-up rates as outcome variables. We use low take-up rates as a proxy for low plan quality (i.e., more expensive without necessarily being more generous), because employees may instead opt for spousal insurance at another firm or forego employer insurance entirely if they dislike the plan. Despite lacking statistical power, Figures 5c and 5d exhibit patterns consistent with poorer match quality after broker steering. In Figure 5c, log premiums per enrollee display a negative change in slope with respect to the number of enrolled employees at the threshold, while in Figure 5d, the take-up rate shows a positive change in slope beyond the threshold. Together, these patterns suggest that brokers tend to steer employers toward more expensive plans with lower employee take-up. This indicates that steering does not necessarily improve match quality and may, in fact, reduce welfare by pushing employers into higher-priced, less-preferred insurance options.

4.3. *Effects of Commissions on Insurer-Broker Contracts*

Next, we examine whether commissions influence brokers' contracting decisions with insurers. We exploit plausibly exogenous, time-varying changes in broker commissions. Specifically, Insurer 2, which is the only insurer that reimbursed brokers on a per-member-per-month (PMPM) flat fee basis, as opposed to a percentage of premiums, changed the level of its flat fee twice throughout our sample period. Prior to 2014, Insurer 2 paid brokers $5 PMPM. In 2014, it raised this reimbursement to $18 PMPM, before reducing it to $12 PMPM in 2016. No other insurers changed their commission schedules during this time frame.

We leverage these changes as a natural experiment and estimate the following event-study specification:

$$Y_{ict} = \sum_{k \neq -1} \beta_k \cdot D_i \cdot \mathbb{1}[t - t_i^* = k] + \alpha_i + \alpha_{ct} + \varepsilon_{ict}, \tag{5}$$

where Y_{ict} is the share of brokers contracting with insurer i in county c at time t, D_i is an indicator for Insurer 2, t_i^* is the event year (either 2014 or 2016), and α_i, α_{ct} are insurer and county-time fixed effects. The coefficient β_k is the parameter of interest, which captures the differential effect of commission changes relative to the year prior to the event.

Figure 6 displays the results for the two events separately. Note that we end the post-period of the "2014" event in 2016 and the pre-period of the "2016" event in 2014, so that the two events do not contaminate each other. When Insurer 2 increased commissions in 2014 (blue dots), the share of brokers it contracted with increased compared to other insurers that experienced no change. Conversely, when Insurer 2 decreased its commissions in 2016 (red dots), its broker contracting share declined. Analogous stacked differences-in-differences estimates indicate that a 33% reduction in commissions results in a 5% decrease in contracting brokers. These results suggest that brokers respond to commission incentives in their contracting choices with insurers.

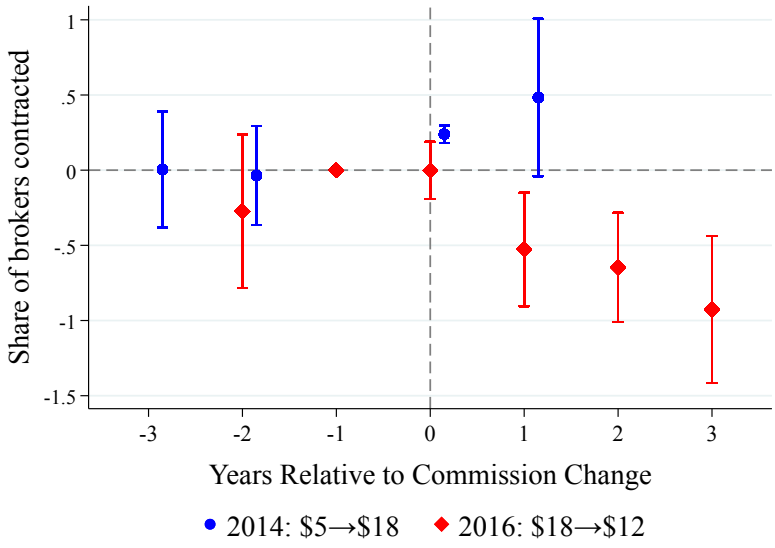
As a placebo test, we re-estimate the specification using insurers' financial characteristics as the outcome variable. This addresses the concern that changes in contracting shares could be driven by confounders, such as shifts in plan quality, other than commission changes. Figure A3 displays the results. While we find some suggestive evidence of increases in cost-sharing following the commission reduction, these effects are not statistically significant and, if anything, work in the opposite direction: more generous plans would be expected to attract more brokers, not fewer.

4.4. *Effects of Contracting Networks on Insurer Competition*

Finally, we examine how broker contracting networks affect competition among insurers. Broker contracting decisions play a central role in shaping the competitive environment because selective contracting means employers are not always exposed to the full set of insurers active in the market. As a result, not all insurers compete directly against one another within a given employer's choice set.

To study this mechanism, we leverage the same natural experiment and results in Figure 6 as a first-stage in determining the effect of plausibly exogenous network changes on premiums. Specifically, we run an event-study of Insurer 2's commission changes on premiums-per-enrollee of *other* insurers. Rather than comparing Insurer 2 against other insurer's, we compare brokers already contracting with Insurer 2 in the

Figure 6. Effect of commission changes on broker networks



Notes: This figure plots the point estimates and 95% confidence interval from the estimation of equation (5). The blue plot represents the estimation results for the event where Insurer 2 increased flat commissions in 2014. The red plot represents the estimation results for the event where Insurer 2 decreased flat commissions in 2016. The dependent variable is the share of brokers contracted at the county-year level. Both plots are normalized to the year prior to the event. Standard errors are clustered at the county level.

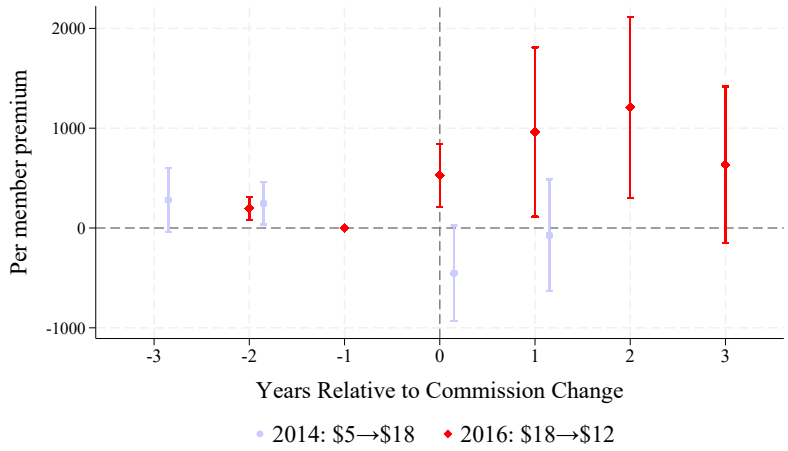
pre-period against brokers who did not. As an example, for the 2014 event (Insurer 2 raising its commission from $5 to $18), we consider the "treatment" group the set of brokers who did not contract with Insurer 2 prior to 2014 and the "control" group the set of brokers who did have contracts with Insurer 2. The logic is that former group would be theoretically exposed to the commission change (and potentially react by including Insurer 2 in its network, as in Figure 6), whereas the latter group would be relatively unexposed. Because Insurer 2's commission changes ought not to affect the premiums charged by other insurers—except through changes to whether brokers contract with Insurer 2 —this specification would specifically isolate how competitive changes affect prices. We apply a similar logic to the 2016 event, but flip the treatment and control groups.

Estimates in Figure 7 show that narrower contracting networks soften competition among insurers. For the 2014 event, we show a small decrease in premiums-per-enrollee as Insurer 2 increased its commissions (and thus networks expanded). The results are not statistically significant. However, for the 2016 event, premiums per enrollee are systematically higher when the choice set includes fewer insurers. On average, the 6% narrowing of networks seen in Figure 6 leads to an approximate $1,000 per enrollee increase in premiums charged.This finding indicates that broker networks effectively determine the scope of rivalry by restricting or expanding the set of insurers available to employers.

4.5. Summary and the Need for a Model

So far, our reduced-form exercise has shown that commissions influence brokers' recommendations and also shape brokers' contracting incentives and the competitive pressure facing insurers. These findings point to an important tradeoff: while lowering commissions weakens brokers' incentives to steer, it also shrinks the size

Figure 7. Effect of contracting networks on insurer pricing



Notes: This figure plots the point estimates and 95% confidence interval from the estimation of an alternate version of equation (5). The dependent variable is premiums-per-enrollee of non-Insurer-2-insurers. The blue plot represents the estimation results for the event where Insurer 2 increased flat commissions in 2014. For this plot, the treatment group is the set of brokers not contracting with Insurer 2 in the pre-period and the control group is the set of brokers that have pre-existing contracts with Insurer 2. The red plot represents the estimation results for the event where Insurer 2 decreased flat commissions in 2016. Here, the treatment group is brokers who did have a contract with Insurer 2 in the pre-period, and the control group is those who did not. Both plots are normalized to the year prior to the event. Standard errors are clustered at the county level.

of broker–insurer networks, thereby softening insurer competition and raising premiums. These competing forces motivate our subsequent structural analysis, which aims to quantify the welfare implications of each channel.

5. Model

Our empirical model features brokers as intermediaries between employers and oligopolistic insurers. The game proceeds in four stages. First, given commission schedules, brokers select which insurers to include in their menus and incur the associated costs of contracting. Second, employers randomly meet with a broker and decide whether to use the broker's services. Third, insurers simultaneously set premiums conditional on the contracting structure, taking employers and brokers as price takers. Fourth, employers select one insurer for their workforce, either from the broker's menu or from the market at large if they do not employ a broker. In what follows, we present the details of each step in reverse order.

5.1. Employers Choose Insurers

Employers choose one insurer to maximize a weighted average of employee welfare and costs.[13] We assume that employers in different geographic markets, industries, and firm-size categories are indexed as distinct

[13] As Form 5500 data only report insurer identity but not insurance plans or their detailed characteristics, we make the simplifying assumption that insurers are single-product oligopolies in a given geographic market, where financial characteristics are held fixed but product pricing could vary by market, employer type, and the competitive environment determined by contracting arrangements.

types e. The flow utility of employer e from choosing product j is:

$$u_{ej} = \overbrace{-\alpha_e \underbrace{(1-\iota)p_j s^w_{ej}}_{\text{employer paid premium}} (1 + \underbrace{B_e}_{\substack{\text{whether} \\ \text{use broker}}} \underbrace{\gamma_j}_{\substack{\text{percentage} \\ \text{commission}}})}^{\text{weight on own expenses}} + \overbrace{\rho_{ej}\mathcal{W}_{ej}}^{\text{weight on employee welfare}} + \delta^e_j + \xi_{ej} + \epsilon_{ej}, \quad (6)$$

where B_e is an indicator of whether the employer works with a broker to offer insurance; γ_j is the percentage commission that the broker receives from selling product j; δ^e_j is insurer-specific brand effects perceived by the employer; ξ_{ej} is unobserved preference shocks; ϵ_{ej} is independent and identically distributed extreme-value type-I error.

The first term in equation (6) measures the employer's weight on its own expenses. α_e is the employers' price sensitivity. $(1-\iota_e)p_j s^w_{ej}$ is the employer's average premium (net of commissions) contribution per employee. If the employer uses brokers, its payment to insurers is inflated by percentage commissions $B_e\gamma_j$.

The second term, W_{ej}, measures the average ex-ante surplus of employees of employer e from having access to insurer j. The more that the plan of insurer j caters to employees' preferences, the higher the share of employees take up the plan, and the larger the ex-ante surplus is. Appendix C1 describes how we derive W_{ej} in detail.

We assume brokers influence employers' choices by adjusting how much employers value their employees' welfare, captured by the second term in equation (6):

$$\rho_{ej} = \rho_0 + \underbrace{B_e\rho_1}_{\text{effect of choice quality}} + \underbrace{B_e\rho_2\gamma_j}_{\text{steering}}. \quad (7)$$

ρ_1 captures brokers' effect on choice qualities. For instance, brokers could reduce informational frictions, enabling employers to better understand their employees' preferences. ρ_2 captures brokers' steering ability, or the degree to which commissions distort brokers' recommendations. We allow ρ_1, ρ_2 to differ by brokers' characteristics, for example, whether the broker is an individual agent or a brokerage firm.[14] Equation (7) thus captures the key force of brokers' steering responses to commissions.

5.2. Insurer Choose Prices

Insurers compete à la Nash-Bertrand. Conditional on the broker menu m that decides its competitors, insurer j sets premiums p_{ejm} for employer type e to maximize profits:

$$p_{ejm} = \text{argmax}\, \pi^I_{ejm} = \text{argmax}\left(s^e_{ejm}\, s^w_{ej}\, (p_{ejm} - mc_{ej})\right),$$

where s^e_{ejm} is the share of employers in type e choosing plan j given menu m, according to the decision rule outlined in equation (6). mc_{je} denotes marginal costs of insurer j for insuring employer type e.

[14] We assume that the intensity of steering is homogeneous across brokers of the same type. This is because we do not observe claims mapped to the employer-broker-insurer level, and thus cannot compute the extent to which brokers differentially leave money on the table. While this precludes heterogeneity in steering intensity, it offers a tractable approximation of average steering effects.

We further project marginal cost mc_{ej} into

$$mc_{ej} = c_e + c_j + \omega_{ej},$$

where c_e captures the riskiness of employer type e; c_j captures variations across insurers' price levels for the same procedure-client pair due to its negotiation power; ω_{ej} are the insurers' variable cost shocks. We assume insurers observe the realizations of demand and marginal cost shocks (ξ, ω) when setting prices, whereas realizations of these shocks are unobservable to econometricians.

Insurers' first-order condition on p_{ejm} is

$$\frac{\partial s^e_{ejm}}{\partial p_{ejm}} s^w_{ej}(p_{ejm} - mc_{je}) + \frac{\partial s^w_{ej}}{\partial p_{ejm}}(p_{ejm} - mc_{je}) + s^e_{ejm} s^w_{ej} = 0. \tag{8}$$

The first two terms measure the decrease in marginal profit due to less share of employers choosing the insurer, and less share of employees taking up the plan. The last term denotes the increase in marginal profit due to the per-enrollee revenue increase. Note that insurers' pricing p_{ejm} responds to the contracting arrangement m, which shapes the competitive environment facing insurer.

5.3. Employer Choose Whether to Use Brokers

Given the set of brokers operating in their market, employers randomly meet a broker and decide whether to use the broker's services, which causes differences on two margins. First, if the employer uses a broker, its choice set will only consist of the menu of insurers with which the broker has contracted $j \in M$. Instead, if it does not use a broker, its choice set will include all available insurers on the market, plus the option of not offering insurance $j \in \mathcal{J} \cup U$. Second, if the employer forgoes broker services, it must pay the hassle costs of completing administrative paperwork or searching for a suitable insurer.

Let V^{wb} and V^{nb} denote utilities with and without a broker. We assume that brokers change employers' choice rules but not their experienced utility, so only ρ_0 and ρ_1 in equation (7) participate in employers' flow utility: $\rho^{\text{true}} = \rho_0 + \rho_1$. In other words, the employers' true preferences for an insurance product are measured by the case where it works with a broker and the broker does not receive any commissions. Following Allcott (2013), we explictly account for misjudged preferences by allowing a wedge between the employers' choice rules and experienced utility:

$$V^b = E[\underbrace{\max_{j \in M}}_{\text{broker menu M}} u^b_{ej}] + \frac{1}{\alpha_e} \sum_j \left(s^e_{ejm}(\underbrace{-\rho_2 \gamma_j}_{=\rho^{\text{nb}} - \rho^{wb}}) \text{WTP}_{ejm} \right), \tag{9}$$

$$V^{nb} = E[\underbrace{\max_{j \in J \cup U}}_{\text{all options on market}} u^{nb}_{ej}] \tag{10}$$

where u_{ej} is detailed in equation (7).

Employer e uses a broker if its expected utility gain exceeds a random hassle cost F:

$$B_e = 1 \quad \text{if} \quad V^{wb} > V^{nb} - \underbrace{F}_{\text{hassle cost}}, \quad F \sim N(\mu_F, \sigma^2_F) \tag{11}$$

We further parameterize F as a function of employers' types, i.e., by industry, firm sizes, geographic markets, and an idiosyncratic normally-distributed error term, with standard deviation σ_F.

Our framework does not explicitly model brokers competing directly against one another. Instead, we assume that brokers face discipline from the outside option of employers forgoing broker services and quoting for insurers directly themselves. This modeling choice generates forces that operate in a manner similar to broker competition. Specifically, brokers are limited in their ability to steer employers, as excessive steering increases the likelihood that employers are matched with a non-preferred insurer, thereby reducing expected utility and discouraging broker use. Likewise, brokers refrain from offering overly narrow networks, as limited product variety also lowers employers' expected utility and, in turn, the likelihood that employers will rely on broker services.

This modeling choice is equilvalent to assuming that employers do not shop for brokers. This can be reasonable because of high search costs: Small employers are most likely to work with a broker located in close geographic proximity, as shown in Appendix C2. Employers do not shop for brokers, just as they don't search for all insurance products.

Our framework may understate the disciplining forces on steering and network design if, in practice, brokers compete as intensely with one another as they do with the outside option. As a result, our estimates may overstate the distortions arising from these frictions. To assess the robustness of our findings, we consider alternative specifications that strengthen the outside option. For instance, we allow the attractiveness of offering insurance without a broker to increase exogenously or to rise with the number of brokers available in the market.

5.4. Broker Choose Menus of Insurers

Given the commission schedule, broker b decides on the menu M of insurers to contract with. Brokers form expectations over the distribution of employers' demand shocks ξ, insurers' variable cost shocks ω, and employers' hassle cost shocks F. All brokers simultaneously decide on the set of insurers they contract with, and these joint decisions collectively shape the network structure of the industry.

The expected commission income from menu m is

$$\Pi^b(m) = \mathbb{E}_{\xi,\omega,F}\Big[\sum_{e\in\mathcal{E}_b} \underbrace{s_{eb}}_{\substack{\text{share employer}\\ \text{use brokers}}} \sum_{j\in m} \underbrace{s^e_{jem}s^w_{jem}}_{\substack{\text{demand from}\\ \text{empl. e, worker j}}} \underbrace{p_{jem}\gamma_j}_{\substack{\text{commission}\\ \text{income}}} \Big], \tag{12}$$

where $\mathcal{E}_b$ is the set of employers in broker b's market.

Brokers choose menu by maximizing its expected profit, which equals its commission income minus contracting cost.

$$m = \text{argmax}\Big(\underbrace{\Pi^b(m)}_{\text{commission income}} - \underbrace{K(m)}_{\text{contracting costs}} \Big). \tag{13}$$

Equation (12) and (13) reveal several factors that affect the brokers' contracting decisions. First, higher commissions from an insurer, γ_b, increase the benefit of including that insurer. Second, a broader menu may attract more employers to use the broker, thereby increasing s_{eb} and expected income. However, a

broader menu also incurs higher contracting costs, namely, a larger $\sum_j K_j$. Third, brokers take into account that their menu decision affects insurer pricing p_{ejm} and consumer choices s^e_{ejm}, s^w_{ejm}. Equation (13) thus yields broker-specific menus m, which are functions of insurers' commission rates, pricing power, product characteristics, and the preferences of employees and employers on the market.

6. Estimation and Identification

6.1. Employers Choose Insurers

Key primitives include employers' weights for their own expenses, α_f, weights for employee surplus, ρ_0, ρ_1, ρ_2 in the no-broker, with-broker, and with-broker-commission cases, and brand perception, δ^e_j.

Estimation. We use the maximum likelihood estimator to estimate the conditional logit model outlined in equation (6). This is equilvalent to adding employer fixed effects in the equation of employers' flow utility, which implicitly controls for time-invariant preference shocks by market, industry, and firm size.

We draw on two additional data inputs to facilitate the estimation. First, since the Form 5500 data only reports premiums for the realized insurance choice, we assume that employers in the same market, type, facing the same insurer and broker menu will get the same price quote. Appendix B1 describes how we compute premiums for each alternative insurer. Second, we compute the average welfare of employees for each alternative insurer $\hat{W}_{ej}$ in the product menu offered by their observed broker, using the estimates of employee preferences, along with the imputed premium and product. See Appendix C1 for details.

Identification. The correlation of choice patterns and employee valuations among firms that do not work with brokers identifies ρ_0. The differences in the choice patterns among firms that work and do not work with brokers identify ρ_1. The identification of the steering coefficient, ρ_2, exploits discontinuities in commission schedules *within* insurers. Conditional on the same insurer, the differential choice patterns of employers who work with brokers and are just below and above the commission change cutoff, along with their correlations with the magnitude of commission changes, identify how commission payments affect choice rules.

The correlation of choice patterns and insurance premiums identifies α_f. Insurers' knowledge of employers' unobserved preferences when making pricing decisions creates a correlation between residuals and prices. To address this endogeneity, we use the Hausman instrument following Nevo (2001). The identifying assumption is that variations in prices in other markets can signal an insurer's cost changes in all markets, which also shifts the equilibrium pricing choices in the focal market. It is plausible that residual demand shocks in other markets are mean independent of residual demand shocks in the focal market, conditional on employer fixed effects that control for time-invariant shocks by market, industry, and firm size, and insurer fixed effects.

Results. Table 3 reports estimates of employers' preferences.

Looking first at the premium coefficients, the results indicate that employers incur disutility from their portion of the total premium paid, as expected. The coefficient indicates that, for firms not using a broker (and hence given the option not to offer insurance), a 10% increase in premiums results in a 4.8 percentage point (pp) decrease in the probability of offering insurance. This is well within established estimates of offer

Table 3. Estimates of Employers' Preferences

Realized Choice	Coefficient	Standard Error
WTP per employee	0.0435	0.0133
has broker × WTP	0.0230	0.0044
has broker × WTP × effective commission	0.1247	0.0266
has broker × WTP × unknown commission	−0.0009	0.0014
Premium per employee	−0.0326	0.0049
Control function for premium	0.0156	0.0049

elasticities and, in particular, is quite similar to estimates from Gruber and Lettau (2004). On the intensive margin, if the price of a major insurer increases by 10%, holding all other prices constant, the probability of a no-broker firm offering that insurer decreases by approximately 13%.

We also compare employers' weights on employee surplus, ρ_e, and their premium expenditure, α_e. A rational employer should give equal weight to these terms, i.e., $\rho_e = \alpha_e$. For instance, if the employer offers more generous plans than the average employee values, they could reduce the coverage level and compensate employees with higher wages, improving employees' total surplus. Therefore, any deviation from equal weights could be seen as a distortion from the optimum.

Three key patterns are outlined. First, employers who do not work with brokers value \$1 of employees' welfare relative to their own premium expenses at \$1.33. In other words, we find that, for each dollar in value of employee surplus, employers pay, on average, \$1.33 in insurance premiums. While suggestive that employers are therefore "overweighting" health insurance relative to average employee valuations, this discrepancy could be rationalized by the tax exclusion on employer-sponsored insurance (Gruber, 2010).

Second, however, brokers tend to increase employers' valuation of employee WTP by \$0.71 to \$0.88, implying that brokers worsen the matches between employers and insurers, and encourage employers to over-insure their workforce. Our estimates imply that, at the most extreme, firms with brokers pay about 67% more in premiums to generate the same dollar of employee value as non-broker firms. This "overweighting" is consistent with the estimates from Tilipman (2022). We note that two possible reasons could explain this difference. Employers with brokers may differ in unobserved ways that render them more likely to offer generous insurance than employers without brokers. Alternatively, it is possible that even in the absence of commissions, brokers may recommend different types of insurers than employers would have chosen on their own.

This leads to our final point: commissions lead to a higher employer weight on employee surplus. Indeed, we estimate that about 20% of the difference in generosity of offers between no-broker and broker firms reported above can be attributed to broker commissions. For a prominent insurer in our data, decreasing commissions by 1pp reduces employers' likelihood of choosing the plan by about 3pp. This is of a similar magnitude to our reduced form estimates of brokers' steering effects in Section 4. It also implies that, given our intensive margin elasticities of employees reported in C1, employers with brokers appear to be more sensitive to the same percentage change in commissions than in raw premium rates.

6.2. Insurer Choose Prices

The key primitive in this stage is marginal costs, mc_{je}, by insurer, employer types, and geographic markets.

Estimation and Identification. We use a generalized method of moments estimator, where the moments are insurers' first-order conditions (equation (8)). The premium levels identify marginal cost. This is in a similar spirit to backing out the marginal costs that rationalize prices using pricing FOCs (Berry, 1994; Berry et al., 1995): given demand elasticities and derived markups, we can back out a one-to-one mapping between observed premiums and total marginal costs.

Results. Table 4 reports our estimates. The magnitude of the monthly per-member marginal cost estimates is sensible. First, the marginal cost-to-price ratio is approximately 85%, consistent with MLR requirements. Second, big insurers have higher markups than small insurers. Third, New York City and urban areas incur higher claims expenses than rural areas, such as the Finger Lakes. Fourth, larger employers have lower claims costs than smaller employers.

Table 4. Per-Member Monthly Marginal Cost Estimates

	Marginal Costs	Premiums	Cost Ratio (%)
By Insurer			
Aetna	677	772	87.74
MVP Health	463	515	89.87
Empire/Anthem	674	766	87.90
Oxford/United	639	773	82.71
Healthnow	618	688	89.90
Connecticut/Cigna	650	730	89.01
HIP of NY/Emblem	597	701	85.35
Excellus	539	643	83.84
CDPHP	586	668	87.71
Independent Health	541	617	87.75
By Region			
Capital Region	551	626	88.09
Central NY	546	630	86.76
Finger Lakes	531	602	88.32
Mid Hudson	624	717	87.16
Long Island	643	744	86.45
Mohawk Valley	522	594	88.05
New York City	658	758	86.90
North County	488	564	86.80
Southern Tier	580	672	86.54
Western NY	549	637	86.22
By Industry			
Natural resources	607	666	91.09
Mining	470	532	88.30
Construction	665	759	87.76
Manufacturing	601	685	87.84
Wholesale and retail trade	610	695	87.77
Transportation and utilities	653	758	86.35
Information	632	732	86.38
Financial activities	652	754	86.61
Professional and business services	648	748	86.76
Education, health, and social services	589	679	86.94
Leisure and hospitality	592	679	87.17
Other services	649	751	86.54
By Employer Size			
100-499 employees	636	730	87.12
500+ employees	574	662	86.82

6.3. Employer Choose Whether to Use Brokers

Primitives to be estimated in this stage are the mean μ_F and the standard deviation σ_F of employers' hassle costs associated with offering insurance without a broker.

Estimation. We parameterize the hassle costs parameters F by industry, firm sizes and regions, and use maximum likelihood estimation based on employers' decisions of whether to use a broker. We match the observed share of employers using brokers in each geographic market.

Identification. The hassle cost parameter F is identified by variation in broker usage rates across markets with different characteristics. All else equal, levels of broker usage rates identify the mean of hassle costs μ_F, while dispersions of broker usage rates identify the standard deviation of hassle cost draws σ_F.

Intuitively, if brokers facilitate the insurance offering process, higher hassle costs should lead to higher fractions of employers working with brokers instead of handling the process on their own. Consider two markets with identical distributions of premiums, commissions, and employer characteristics, but different broker usage rates. Under our model, these different usage rates can only be rationalized by different levels of hassle costs F that employers face when operating without a broker. This is because we have already estimated all other parameters that affect the relative value of using a broker, using variation in plan choices conditional on broker status.

This identification relies on two assumptions. First, previously estimated preference parameters are consistent. Second, unobserved employer characteristics affecting hassle costs and broker choice are not systematically correlated with valuations in broker choice sets relative to the full choice set (from not working with a broker). The second assumption could be violated if, for example, employers with higher unobserved hassle costs disproportionately select into markets with more brokers, larger broker networks of insurers, or where those networks are easier to perceive.

To address this concern, we rely on an instrument for the firm's expected valuation of the broker menu. This instrument exploits hospital mergers within the same geographic market, as well as hospitals being acquired by systems headquartered in other geographic markets. The idea is that when the hospital market gets more concentrated, hospital systems gain higher bargaining power and negotiate higher charges for providing the same medical services (Ho and Lee, 2017). This works as a cost shifter for insurers, as long as the hospitals that experience either merger or acquisition events are in the insurer's network. Thus, this shifts the valuation of a broker's insurer network. However, these mergers are unlikely to affect broker usage rates apart from this channel.

Results. Table 5 reports coefficients of hassle costs. The average hassle cost in the data is estimated to be about \$30 per-member-per-month (7.8% of the average employer premium share), and ranges from a low of about \$14 PMPM (3.6%) to a high of \$52 PMPM (13.5%).

The heterogeneity across industries, regions, and firm sizes follows expected patterns. First, the industries that see the lowest hassle costs are information and leisure and hospitality, likely due to the sophistication of those types of firms (in the case of the former) and simple benefit structures (in the case of the latter). Notably, the industries with the highest hassle costs are construction, agriculture, and mining, which are characterized by highly variable risk profiles among employees and less sophistication than comparable

Table 5. Hassle Cost Estimates

Variable	Hassle Cost ($ PMPM)	Std. err.
Constant	45.04	2.07
Industry (NAICS):		
Construction	−2.15	2.42
Manufacturing	−9.45	1.87
Wholesale and retail trade	−13.73	1.69
Transportation and utilities	−12.08	1.85
Information	−16.66	1.80
Financial activities	−8.10	1.74
Professional and business services	−12.31	1.78
Education, health, and social services	−10.53	1.74
Leisure and hospitality	−17.48	1.74
Other services	−11.89	1.74
Firm Size:		
200-499 employees	−5.64	0.90
500+ employees	−11.39	0.55
Region:		
Central NY	+2.08	0.84
Finger Lakes	−3.23	0.65
Mid Hudson	+2.81	0.69
Long Island	+0.34	0.78
Mohawk Valley	−1.32	0.96
New York City	+2.61	0.56
North County	−5.72	1.21
Southern Tier	+6.54	0.92
Western NY	−0.76	1.15

Notes: Estimates represent hassle costs per member per month. The omitted industry category is mining. The omitted firm size category is 100-199 employees. The omitted region is the Capital Region.

service-industry firms. Second, compared to New York City, the Mid-Hudson, and Long Island, the Finger Lakes, Mohawk Valley, North County, and Western NY all have lower hassle costs. This is likely driven by the fact that the latter are more rural, with fewer competing insurers and less complex benefit offerings. Finally, the hassle costs monotonically decrease by firm size, with firms of 200-499 employees paying about $6 PMPM less than smaller firms, and large firms paying about $11 PMPM less. This is consistent with both the observation that larger firms tend to use brokers at lower rates and that larger firms have more established human resources departments and benefits administration.

The magnitude of the hassle costs is certainly reasonable. As an example, the costs for a 200-employee professional and business services firm in New York City are around 71k a year. This is almost equivalent to the average salary of either an HR specialist or Benefits Administrator in New York City, which is listed at $74k.[15] This makes sense as employers could perceive hassle costs as about the equivalent of hiring a dedicated human resources expert to search for suitable insurers and complete relevant paperwork.

6.4. Broker Choose Menus of Insurers

The key primitive in this stage is the brokers' contracting costs K.

[15] https://www.indeed.com/career/benefits-administrator/salaries/New-York–NY

Estimation. Conducting empirical analysis on broker network formation presents a series of computational and econometric challenges. First, the number of possible network configurations increases exponentially with the number of insurers and thus suffers from the curse of dimensionality. Second, insurers are interdependent. That is, the revenue from the inclusion of an extra insurer depends on the identity of insurers in the existing portfolio. Instead of directly solving the industry equilibrium, we exploit the inequality restrictions implied by brokers' revealed preferences.

We assume that the total contracting costs of the broker menu m are linearly additive across the contracting costs of each included insurer. If a broker contracts with an insurer, its contracting cost must be weakly lower than the commission income from including that insurer in the menu. This delivers an upper bound on the contracting cost. Let m^* denote the observed broker menu, $m^- = m^* \setminus j$,

$$\Pi^b(m^*) - \Pi^b(m^-) = \Delta\Pi^b_{jm}(m^*, m^-) + \nu^1_{bjm} + \nu^2_{bjm} \geq k_j. \tag{14}$$

Conversely, if a broker does not contract with an insurer, the contracting cost must be weakly higher than the counterfactual commission income from including that insurer, $m^+ = m^* \cup j$. This generates a lower bound for the contracting cost.

$$\Pi^b(m^+) - \Pi^b(m^*) = \Delta\Pi^b_{jm}(m^+, m^*) + \nu^1_{bjm} + \nu^2_{bjm} \leq k_j. \tag{15}$$

ν^1_{bjm} is the expectation error of brokers' expected revenue from a contract (Dickstein and Morales, 2013; Prager and Tilipman, 2025). This expectational error term is unknown to brokers when they make their network decisions. Including this term allows the model to rationalize cases in which it would otherwise predict that a broker should contract with an insurer, but the broker does not form the relationship.

ν^2_{bjm} is an error term that is known to brokers at the time they decide on contracts, but is not known to the econometrician. ν^2_{bjm} allows contracting costs to vary across markets and brokers in ways not otherwise captured by this parameterization. However, the presence of this term introduces a selection problem. Brokers that receive high draws of ν^2_{bjm} will be more likely to contract with insurer j.

Identification. To identify contracting costs, we require additional assumptions.

First, we assume that the residual profit difference ν^1_{bjm} satisfies $\mathbb{E}[\nu^1_{bjm}|\mathcal{I}] = 0$, where $\mathcal{I}$ is the brokers' information set when they decide on contracts. In other words, conditional on the information set, when the sample of inequalities grows large, the errors tend to zero in the limit. We thus stack the inequalities across brokers and estimate k_j using sample moment inequalities (Pakes et al., 2015). Similar assumptions have been made in Dickstein and Morales (2013) and Prager and Tilipman (2025).

Second, we assume that the contracting cost with a given insurer is identical across brokers and menus, $\nu^2_{bjm} = \nu^2_j, \forall b, m$. In other words, we rule out the possibility that network formation depends on private information of heterogeneity in contracting costs across brokers. This implicitly assumes away any structural errors that might bias the estimate of k (Eizenberg, 2014; Pakes, 2010), and we derive contracting costs as $k'_j = k_j + \nu^2_j$. Similar assumptions have been made in Ho (2009) and Nosko (2010).

The moment inequalities literature has proposed several alternative approaches to address the structural error ν^2_{bjm}. Ciliberto and Tamer (2009) estimate a distribution for the structural error term. This approach is computationally challenging in our setting, as the network formation problem is combinatorial rather than a

binary-entry problem. Fan and Yang (2025) develop inequalities that allow the estimation of the distribution of sunk costs in settings with product positioning; however, their inequalities require the identification of the best and worst case situations in terms of rivals' actions. Identifying these in our setting is non-trivial, since the worst-case scenario may not be a full insurer menu, but a selective contract by rivals. Wollmann (2018) restricts how the structural error term may vary across observations. This approach is unattractive in our setting because it limits the ability to identify parameters that are common across insurers or years.

Third, k_j reflects the annual fixed costs of contracting incurred irrespective of whether an insurer had a contract with a broker in prior years. That is, we assume that the broker training and licensing to sell process is costly, even for existing contracts. This is motivated by the fact that brokers undertake the costly task of learning about changes in insurers' policies every year, then deciding whether to carry that insurer. This also facilitates our computation by abstracting from the dynamic decisions of long-term contracts.

Under these assumptions, we incorporate the following inequality moments into the GMM estimation by penalizing violations of these conditions proportional to their magnitude.

$$\mathbb{E}_{b,m}[\Delta\Pi^b_{jm}(m^+, m^*) - k'_j)] \leq 0, \quad \mathbb{E}_{b,m}[k'_j - \Delta\Pi^b_{jm}(m^*, m^-))] \leq 0. \tag{16}$$

k_j are identified by differences in expected revenue between different menus.

Results. Table 6 reports the estimates.

Table 6. Contracting costs (rescaled to a monthly basis)

Insurer	Lower Bound	Upper Bound	Midpoint
Aetna	666	1496	1081
MVP Health	635	1213	924
Empire/Anthem	1403	1539	1471
Oxford/United	688	2024	1356
Healthnow	644	1038	841
Connecticut/Cigna	1367	2034	1700
HIP of NY/Emblem	417	585	501
Excellus	2219	2689	2454
CDPHP	1174	1282	1228
Independent Health	1437	1629	1533

7. Counterfactuals

7.1. Quantify Distortions

We conduct two partial equilibrium simulations to quantify the welfare effects of relevant distortions. We hold the employer-broker matches fixed as in the status quo, and focus on the employers who work with brokers. First, we remove commission-induced steering in employers' choices rules to measure the impact of broker steering. Second, we augment employers' choices sets to include the non-contracted insurers to measure the impacts of selection contracting. We report average per-enrollee prices of the resulting choice patterns, employee surplus (W_{ej} in equation (6)), and employer surplus (defined in equation (9)). For tractability, the results reported are in year 2018, across all geographic market in NYS.

Before proceeding to the simulation results, Table A1 reports the model fit. It is reassuring that the simulated equilibrium characterizes the observed pricing strategies and market shares well.

Steering. Table 7 reports the simulated market outcomes under alternative broker incentive regimes. Column (1) shows the baseline case with steering. Columns (2) and (3) remove steering, holding prices fixed and allowing for price re-optimization, respectively. Three key patterns are outlined.

Table 7. Welfare Effects of Steering Distortions

	(1)	(2)	(3)
	Employers with brokers		
	With steering	No steering No repricing	No steering With repricing
Average menu sizes	4.99	4.99	4.99
Average employee takeup rates	64.99	67.00	68.57
Average employee surplus	105.66	112.13	114.64
Average employer surplus (no hassle costs)	237.49	243.30	260.52
Average premium per enrollee	8890	8721	8411
Average premium per employee	5731	5799	5732
Average markup	1.20	1.18	1.14
Average commission per enrollee	216	155	170
Average percentage commission	2.43	1.78	2.02

First, comparing Columns (1) and (2), we find that restoring the optimal choice rule, so that employees choose plans aligned with their preferences rather than broker steering, raises welfare substantially. Average employee surplus increases from 105.7 to 112.13, while employer surplus (net of hassle costs) rises from 237.5 to 243.30. These welfare gains reflect both improved plan matching and lower priced products being choosen. Average per-enrollee premiums decline slightly from $8,890 to $8,721, consistent with reduced form results.

Second, comparing Columns (2) and (3), we find that allowing insurers to reprice in the no steering scenario induces further decreases in average enrollee premiums from $8,721 to $8,411, a 5% reduction from the baseline. Absent steering, employers now place a relatively smaller weight on employee surplus, or in other words, for the same unit of employee welfare, the employers' willingness to pay drops. This disciplines insurers' pricing power and equilibrium price reduction. As a result, average employee surplus further increases to 114.64, while employer surplus (net of hassle costs) rises to 260.52.

Finally, the effects are heterogeneous across insurers, as shown in Table 8. Average per-enrollee premium drops by 7.8% for large insurers like United, while only 3% for regional insurers like MVP Health. Price declines are most pronounced among insurers with higher baseline commission rates and more agents contracted, that are more likely to benefit from broker steering. Eliminating steering make regional lower-priced insurers stronger competitors for national higher-priced insurers. This equilibrium impacts further enhance consumer welfare.

Anti-competition. Table 9 quantifies the impact to anti-competitive effects by expanding the menu of available plans. Column (1) reports the menus in the baseline, while Column (2) reports the scenario where we manually augment all insurers that appear on the geographic market in that specific year to the choice set. Three patterns are outlined.

Table 8. Equilibrium Strategy under With or Without Steering Case

	(1)	(2)	(3)	(4)	(5)
	Market Share			Price	
	With steering	No steering No repricing	No steering	With Steering	No steering
Aetna	13.7	11.5	11.9	9577	9222
MVP Health	3.9	5.5	4.6	6438	6242
Empire/Anthem	10.3	8.0	8.3	9331	9083
Oxford/United	21.9	9.9	14.7	9884	9166
Healthnow	5.8	7.4	6.4	8322	8137
Connecticut/Cigna	12.8	29.7	23.2	8814	8644
HIP of NY/Emblem	2.5	0.7	0.7	8793	8334
Excellus	24.3	24.0	26.7	8155	7759
CDPHP	2.5	1.2	1.4	8459	8129
Independent Health	2.3	2.1	2.1	7744	7448

Table 9. Welfare Effects of Anti-competitive Distortions

	(1)	(2)
	Employers with brokers	
	Baseline	Full
Average menu sizes	4.99	6.65
Average employee takeup rates	63.67	64.98
Average employee surplus	102.90	105.50
Average employer surplus (no hassle costs)	235.70	259.43
Average premium per enrollee	8903	8789
Average premium per employee	5622	5667
Average markup	1.20	1.18
Average commission per enrollee	215	223
Average percentage commission	2.42	2.54

First, expanding the choice set leads to meaningful improvements in welfare and competitive discipline. Employee surplus rises modestly, from 102.9 to 105.5 with steering, reflecting the gains from better plan matching. Employer surplus increases even more sharply from 235.7 to 259.4 under steering, reflecting the gains from both increased product variety and lower prices.

Secoond, prices respond in the expected direction. Average enrollee premiums decline by roughly 1–2 percent following menu expansion. These declines suggest that broader choice intensifies price competition among insurers.

Finally, Table 10 presents insurer-level responses to menu expansion. When new or popular insurers are introduced, market shares shift considerably toward those entrants. For instance, Oxford/United's share rises from 22.5 to 32.6 percent, while Excellus and Healthnow—insurers less frequently chosen in the baseline menus—lose market share. These reallocations highlight the importance of consumer preferences and name recognition in shaping equilibrium enrollment patterns once employees face a broader set of options. The entry of well-known insurers intensifies competition, forcing incumbent plans to reprice aggressively to maintain enrollment.

Table 10. Equilibrium Strategy under Baseline Menu and Counterfactual Full Menu

	(1)	(2)	(3)	(4)
	Market Share		Price	
	Baseline	Full	Baseline	Full
Aetna	14.2	13.4	9573	9412
MVP Health	3.7	1.7	6445	6356
Empire/Anthem	10.0	9.5	9327	9341
Oxford/United	22.5	32.6	9882	9170
Healthnow	6.0	4.1	8329	8602
Connecticut/Cigna	12.4	13.0	8810	8456
HIP of NY/Emblem	2.4	3.0	8809	8631
Excellus	23.7	19.6	8155	8060
CDPHP	2.6	2.0	8463	8486
Independent Health	2.5	1.2	7749	7630

7.2. Commission Caps

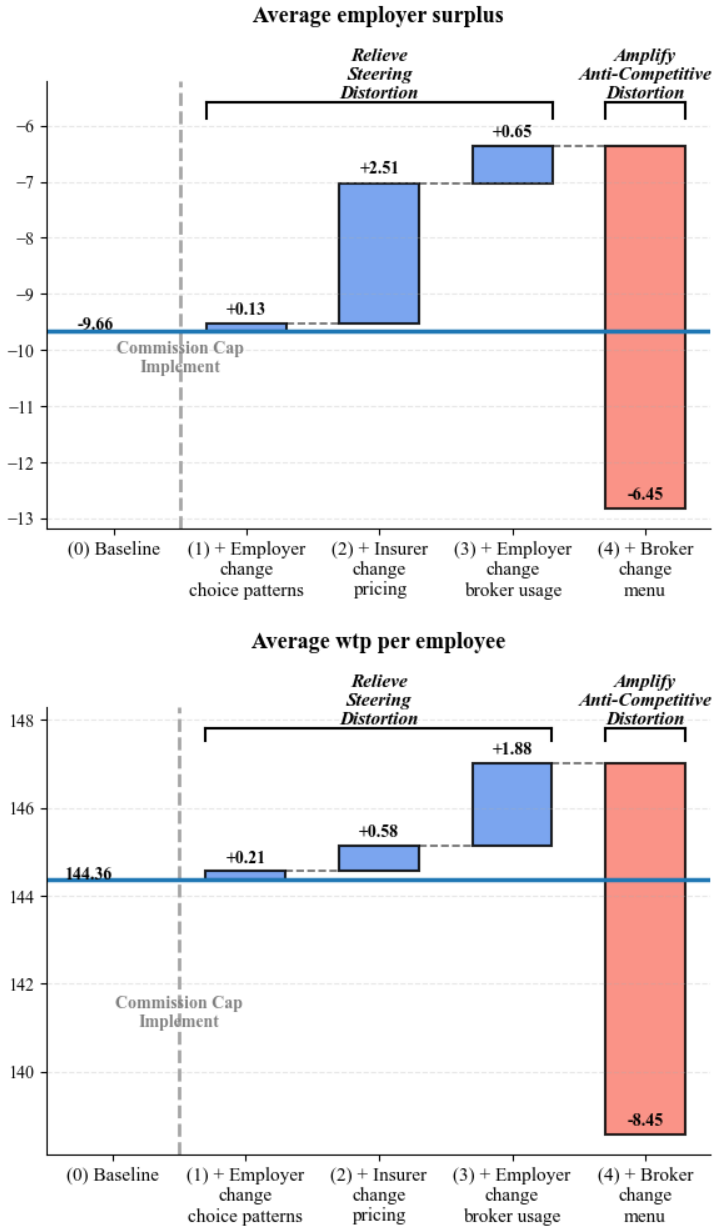
Our counterfactual analysis is motivated by existing policies that constrain intermediaries' compensation. For instance, prior to 2010, New York imposed a cap of 4 percent on commissions for health maintenance organizations (HMOs) and the in-network portion of point-of-service (POS) plans. At the federal level, the Centers for Medicare and Medicaid Services (CMS) sets explicit maximum payments that carriers and Field Marketing Organizations (FMOs) may provide to agents and brokers for enrolling beneficiaries into Medicare Advantage (MA) plans, as specified under 42 C.F.R. §422.2274. These examples illustrate that regulators already intervene in broker compensation schemes in order to limit distortions, yet the welfare consequences of such restrictions remain unclear.

We use the model to evaluate a one-percentage-point cap on broker commissions, decomposing the equilibrium effects into distinct mechanisms. Commission caps affect market outcomes through two competing forces identified in our conceptual framework. First, they weaken brokers' incentives to steer employers toward high-commission plans, improving allocative efficiency. Second, they reduce the profitability of broker-insurer relationships, potentially leading brokers to narrow their contracting networks and thereby softening insurer competition. To quantify the magnitude of each mechanism, we conduct a series of counterfactuals that sequentially allow different margins of adjustment.

Figure 8 displays the simulated effects on employer surplus under each counterfactual. Table 11 reports the complete set of equilibrium statistics for all market participants. The first counterfactual, denoted (0), represents the baseline scenario without any commission regulation. The subsequent counterfactuals apply the commission cap and sequentially allow employers, insurers, and brokers to adjust their behavior.

Our counterfactual (1) isolates the effect of the commission cap on employer choice patterns. We compare the baseline with a situation in which the cap mechanically truncates commissions above one percentage point, but insurers' prices, brokers' menu offerings, and employer-broker matches remain fixed at their baseline values. We allow only employers to adjust their plan choices in response to the new commission environment. Under this counterfactual, the cap weakens brokers' incentives to steer employers toward high-commission plans. As a result, employer choices become better aligned with their true preferences, improving allocative efficiency. Moving from counterfactual (0) to (1), employer surplus increases modestly,

Figure 8. Decompose the welfare effects of commission cap at 1pp



reflecting improved plan matching.

We then compute counterfactual (2), quantifying how insurer price responses amplify or attenuate the effects of reduced steering. We simulate a scenario where the commission cap affects both employer choice patterns and insurer pricing, but brokers' menu offerings and employer-broker matches remain fixed. We allow insurers to optimally reset prices given the new employer choice rules. Reduced steering lowers employers' effective willingness to pay for plan generosity, as the coefficient ρ_2 in equation (**??**) diminishes. This increased price sensitivity disciplines insurers' pricing power, leading to lower equilibrium premiums. Moving from counterfactual (1) to (2), average enrollee premiums decline by approximately 5% from baseline, generating a substantial increase in employer surplus—roughly 3% above baseline levels.

Table 11. The welfare effects of commission cap at 1pp

	CF0	CF1	CF2	CF3	CF4
Overall Market					
Broker usage rate	0.598	0.598	0.598	0.622	0.500
Average choice set size	4.700	4.700	4.700	4.647	4.766
Average worker takeup rate	0.652	0.652	0.655	0.663	0.626
Average price per employee	340.04	341.72	339.45	344.03	328.87
Average WTP per employee	144.36	144.57	145.15	147.03	138.57
Average non-offering rate	0.155	0.155	0.155	0.145	0.189
Average price per enrollee	351.43	352.77	350.15	350.22	354.00
Average cost per enrollee	308.52	309.94	309.15	309.05	312.51
Average markup	1.141	1.140	1.134	1.135	1.134
Average commission percentage	0.013	0.006	0.006	0.006	0.005
Average employer choice value	20.47	20.60	23.11	21.99	24.76
Average hassle costs	30.14	30.14	30.14	28.37	37.58
Average employer surplus	−9.66	−9.54	−7.03	−6.38	−12.82
Among firms that use brokers					
Average choice set size	3.827	3.827	3.827	3.824	3.529
Average worker takeup rate	0.771	0.771	0.775	0.776	0.770
Average price per employee	401.14	403.95	400.16	401.79	407.10
Average WTP per employee	170.48	170.83	171.79	172.26	170.68
Average non-offering rate	0	0	0	0	0
Average price per enrollee	348.07	350.32	345.93	346.87	353.43
Average cost per enrollee	302.08	304.45	303.13	304.04	308.70
Average markup	1.155	1.153	1.143	1.143	1.147
Average commission percentage	0.021	0.010	0.010	0.010	0.010
Average employer choice value	5.16	5.38	9.57	8.88	5.41
Average hassle costs	0	0	0	0	0
Average employer surplus	5.16	5.38	9.57	8.88	5.41
Among firms that do not use brokers					
Average choice set size	6	6	6	6	6
Average worker takeup rate	0.476	0.476	0.476	0.477	0.481
Average price per employee	249.09	249.09	249.09	249.10	250.79
Average WTP per employee	105.49	105.49	105.49	105.55	106.53
Average non-offering rate	0.386	0.386	0.386	0.384	0.378
Average price per enrollee	349.67	349.67	349.67	348.92	348.13
Average cost per enrollee	310.97	310.97	310.97	310.11	309.52
Average markup	1.126	1.126	1.126	1.126	1.126
Average commission percentage	0	0	0	0	0
Average employer choice value	46.21	46.21	46.21	46.51	46.85
Average hassle costs	75.19	75.19	75.19	75.17	75.21
Average employer surplus	−28.97	−28.97	−28.97	−28.66	−28.36

Notes: Simulations reported for year 2019, region 0. CF0: baseline. CF1: allow change choice patterns. CF2: allow change insurer pricing. CF3: allow change broker usage decision. CF4: alloq change broker eqlm menu offering

Next, we compute counterfactual (3), measuring how employers' decisions to use brokers respond to the policy change. We simulate a scenario where the commission cap affects employer choices, insurer pricing, and employer-broker matching decisions, but brokers' contracting networks remain fixed. We allow employers to adjust their probability of using a broker given the new market conditions. Because brokers now distort choices less through steering, broker usage increases modestly. This raises welfare through two channels: it reduces employer hassle costs, and fewer employers select the non-offering option, expanding employee access to insurance. Moving from counterfactual (2) to (3), both employee welfare and employer surplus rise further.

Taken together, counterfactuals (1) through (3) capture the relief of steering distortions induced by the commission cap. These adjustments generate welfare gains: employer surplus rises by approximately 3% relative to baseline, reflecting improved plan matching, reduced insurer markups, and expanded broker usage.

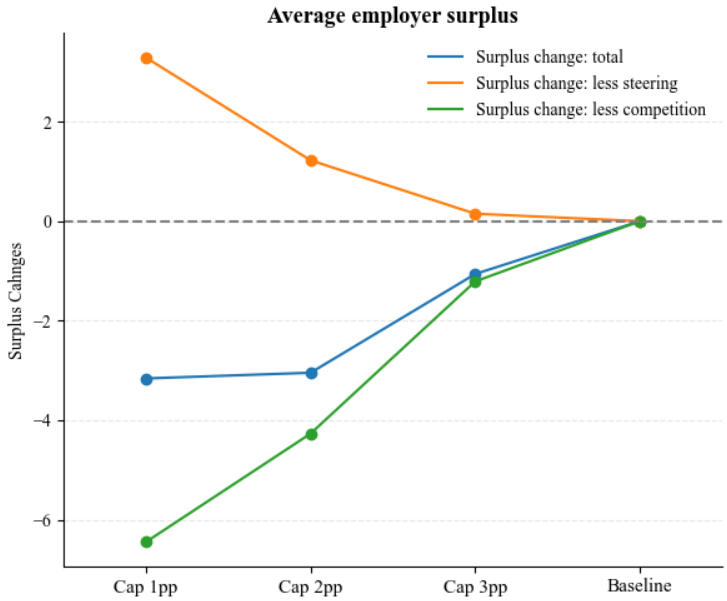
Finally, we compute counterfactual (4) to illustrate the full equilibrium effects, including brokers' contracting responses. This scenario corresponds to the complete equilibrium model, where the commission cap affects all margins: employer choices, insurer pricing, employer-broker matching, and broker-insurer contracting networks. The difference between counterfactuals (3) and (4) isolates the anti-competitive effects of network contraction. The commission cap reduces the profitability of some insurer-broker relationships, leading brokers to drop insurers from their menus. This shrinkage in network breadth amplifies anti-competitive distortions. With narrower choice sets, insurers face less competitive pressure, and prices rise in broker-mediated contracts. Moving from counterfactual (3) to (4), broker usage declines as the value of broker services diminishes, and more employers revert to not offering insurance. Employee welfare falls sharply, and employer surplus declines by more than 6% relative to the gains achieved in counterfactual (3).

The decomposition reveals a fundamental trade-off. Comparing counterfactuals (0) and (3) shows that commission caps improve allocative efficiency by mitigating broker-induced steering and reducing insurer market power, yielding a 3% gain in employer surplus. However, comparing counterfactuals (3) and (4) shows that commission caps also weaken broker incentives to maintain broad networks, producing a 6% reduction in employer surplus through anti-competitive network contraction. In net equilibrium, comparing counterfactuals (0) and (4), the commission cap produces a decline in employer surplus of roughly 3%, as the anti-competitive effects dominate the steering-reduction benefits.

Figure 9 illustrates how employer surplus varies with the stringency of commission caps, decomposing the total effect into its constituent mechanisms. The horizontal axis measures the cap level in percentage points, with lower values corresponding to stricter regulation. The figure plots three series: total surplus change (solid line), surplus change from reduced steering (dashed line), and surplus change from reduced competition due to network contraction (dotted line).

Several patterns emerge from the figure. First, the steering-reduction benefits (dashed line) increase monotonically as caps become stricter—lower commission caps more effectively curtail brokers' incentives to steer employers toward high-commission plans. Second, the competition-reduction costs (dotted line) also increase in magnitude as caps tighten, reflecting greater network contraction when broker-insurer relationships become less profitable. Third, and most importantly, the total effect (solid line) is non-monotonic: at a cap of approximately 3 percentage points, employer surplus roughly equals the unregulated baseline. Stricter caps (below 2 percentage points) produce net welfare losses as network contraction effects dominate the steering-reduction benefits. Conversely, very loose caps (above 3 percentage points) provide only modest welfare gains, as they do little to curtail steering while also preserving existing network structures.

Figure 9. The welfare effects of varying commission caps



8. Conclusion

Our findings highlight the central role of intermediaries in employer-sponsored health insurance markets. Our key insights are twofold. First, brokers influence employers' plan choices and, through their contracting decisions, alter the degree of competition among insurers. Second, regulations on commissions entail an inherent tradeoff: while commission-driven steering distorts employers' choices, the expansion of broker-insurer networks can also intensify competition. This tension highlights the importance of carefully designed policies that strike a balance between the benefits of limiting distortions and the potential costs of weakening competition.

These insights have broader policy relevance. First, similar dynamics may arise in other vertical markets where intermediaries shape consumers' effective choice sets, such as the grocery retailing setting, where store shelves determine which upstream brands are visible to consumers. Second, our results also inform ongoing policy discussions regarding the regulation of payments to intermediaries. By quantifying the welfare effects of steering and network breadth, the paper sheds light on whether commission caps improve efficiency. More generally, the evidence suggests that intermediary frictions may constitute a significant barrier to competition, with implications for explaining the persistently rising healthcare costs.

References

Abaluck, Jason and Jonathan Gruber, "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program," *American Economic Review*, 2015.

Allcott, Hunt, "The welfare effects of misperceived product costs: Data and calibrations from the automobile market," *American Economic Journal: Economic Policy*, 2013, *5* (3), 30–66.

Arnold, Jonathan, Atul Gupta, and Alexander Olssen, "Targeting Efficiencies? Reducing Costs Through Hospital Consolidation," *Unpublished Working Paper*, 2025.

Barwick, Panle Jia, Parag A Pathak, and Maisy Wong, "Conflicts of interest and steering in residential brokerage," *American Economic Journal: Applied Economics*, 2017, *9* (3), 191–222.

Berry, Steven, "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 1994, pp. 242–262.

__ **, James Levinsohn, and Ariel Pakes**, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *63* (4), 841–890.

Bhattacharya, Vivek, Gaston Illanes, and Manisha Padi, "Fiduciary duty and the market for financial advice," Technical Report, National Bureau of Economic Research 2019.

Boehm, Eduard, "Intermediation, Choice Frictions, and Adverse Selection: Evidence from the Chilean Pension Market," Technical Report 2023.

Bundorf, M Kate, Jonathan Levin, and Neale Mahoney, "Pricing and Welfare in Health Plan Choice," *American Economic Review*, 2012, *102* (7), 3214–3248.

Calonico, Sebastian, Matias D Cattaneo, Max H Farrell, and Rocio Titiunik, "Regression discontinuity designs using covariates," *Review of Economics and Statistics*, 2019, *101* (3), 442–451.

Cattaneo, Matis D., Michael Jansson, and Xinwei Ma, "Manipulation testing based on density discontinuity," *The Stata Journal*, 2018, *18* (1), 234–261.

Ciliberto, Federico and Elie Tamer, "Market Structure and Multiple Equilibria in Airline Markets," *Econometrica*, 2009, *77* (6), 1791–1828.

CMS, "The National Health Expenditure Accounts," 2024. link.

Commonwealth Funds, "A New Rule to Protect Medicare Beneficiaries Against Inappropriate Sales Tactics Is Stuck in the Courts," 2025. link.

Conlon, Christopher T and Julie Holland Mortimer, "Efficiency and Foreclosure Effects of Vertical Rebates: Empirical Evidence," *Journal of Political Economy*, 2021, *129* (12), 3357–3404.

Craig, Stuart, "Competition in Employer-Sponsored Health Insurance: Implications for a Public Option," Technical Report, Working Paper 2022.

Cutler, David and Sarah Reber, "Paying for Health Insurance: The Trade-off Between Competition and Adverse Selection," *Quarterly Journal of Economics*, 1998, *113* (2), 433–466.

Dafny, Leemore, Kate Ho, and Maurcio Varela, "Let Them Have Choice: Gains from Shifting Away from Employer-Sponsored Health Insurance and Towards an Individual Exchange," *American Economic Journal: Economic Policy*, 2013, *5* (1), 32–58.

Dickstein, Michael J and Eduardo Morales, "Accounting for expectational and structural error in binary choice problems: A moment inequality approach," Technical Report, National Bureau of Economic Research 2013.

Eizenberg, Alon, "Upstream Innovation and Product Variety in the US Home PC Market," *Review of Economic Studies*, 2014, *81* (3), 1003–1045.

Fan, Ying and Chenyu Yang, "Estimating Discrete Games with Many Firms and Many Decisions: An Application to Merger and Product Variety," *Journal of Political Economy*, 2025, *133* (6), 1886–1931.

Fleitas, Sebastián, Caitlyn Fleming, Gautam Gowrisankaran, and Anthony T Lo Sasso, "Group Health Insurance Versus ACA Marketplaces: Selection, Subsidies, and Welfare," Technical Report, National Bureau of Economic Research 2025.

Grennan, Matthew, Kyle R Myers, Ashley Swanson, and Aaron Chatterji, "No free lunch? welfare analysis of firms selling through expert intermediaries," *Review of Economic Studies*, 2024, p. rdae090.

Gruber, Jonathan, "The tax exclusion for employer-sponsored heatlh insurance," *NBER Working Paper*, 2010.

__ **and Michael Lettau**, "How elastic is the firm's demand for health insurance?," *Journal of Public Economics*, 2004, *88* (7-8), 1273–1293.

__ , **Benjamin R Handel, Samuel H Kina, and Jonathan T Kolstad**, "Managing intelligence: Skilled experts and decision support in markets for complex products," *NBER Working Paper*, 2021, (w27038).

Handel, Benjamin R, "Adverse selection and inertia in health insurance markets: When nudging hurts," *American Economic Review*, 2013, *103* (7), 2643–2682.

__ **and Jonathan Kolstad**, "Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare," *American Economic Review*, 2015.

Ho, Kate and Robin Lee, "Health Insurance Menu Design for Large Employers," *RAND Journal of Economics*, 2023, *54* (4), 598–637.

__ **and Robin S Lee**, "Insurer competition in health care markets," *Econometrica*, 2017, *85* (2), 379–417.

Ho, Katherine, "Insurer-Provider Networks in the Medical Care Market," *American Economic Review*, 2009, *99* (1), 393–430.

Hristakeva, Sylvia, "Vertical contracts with endogenous product selection: An empirical analysis of vendor allowance contracts," *Journal of Political Economy*, 2022, *130* (12), 3202–3252.

KFF, "Health Care Costs and Affordability," 2025. link.

Kim, Gi Heung, "The equilibrium impacts of broker incentives in the real estate market," 2024.

Lee, David S. and Thomas Lemieux, "Regression Discontinuity Designs in Economics," *Journal of Economic Literature*, 2008, *48*, 281–355.

Lee, Robin S, "Vertical Integration and Exclusivity in Platform and Two-Sided Markets," *American Economic Review*, 2013, *103* (7), 2960–3000.

Li, Anran, Becky Staiger, Diane Alexander, and Molly Schnell, "Enrollment Brokers Did Not Increase Medicaid Enrollment, 2008–18: Study examines the effectiveness of Medicaid enrollment brokers.," *Health Affairs*, 2022, *41* (9), 1333–1341.

Liu, Tong, "Bargaining with private equity: Implications for hospital prices and patient welfare," *Unpublished Working Paper*, 2022.

Marone, Victoria and Adrienne Sabety, "When Should There Be Vertical Choice in Health Insurance Markets?," *American Economic Review*, 2022, *112* (1), 304–342.

McCrary, Justin, "Manipulation of the running variable in the regression discontinuity design: A desntiy test," *Journal of Econometrics*, 2008, *142* (2), 698–714.

Myerson, Rebecca and Honglin Li, "Information gaps and health insurance enrollment: evidence from the Affordable Care Act navigator programs," *American journal of health economics*, 2022, *8* (4), 477–505.

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

New York Times, "Financial Brokers Must Now Act in Your Best Interest. What Does That Mean?," 2020. link.

Nosko, Chris, "Competition and Quality Choice in the CPU Market," *Manuscript, Harvard University*, 2010.

Pakes, Ariel, "Alternative Models for Moment Inequalities," *Econometrica*, 2010, *78* (6), 1783–1822.

__ , **Jack Porter, Kate Ho, and Joy Ishii**, "Moment inequalities and their application," *Econometrica*, 2015, *83* (1), 315–334.

Prager, Elena and Nicholas Tilipman, "Regulating Out-of-Network Hospital Payments: Disagreement Payoffs, Negotiated Prices, and Access," Technical Report, National Bureau of Economic Research 2025.

Reuters, "Under the microscope: Spitzer-era ban on insurer fees," 2008. link.

Rey, Patrick and Thibaud Vergé, "Resale Price Maintenance and Interlocking Relationships," *The Journal of Industrial Economics*, 2010, *58* (4), 928–961.

__ **and** __ , "Secret Contracting in Multilateral Relations," 2016.

Robles-Garcia, Claudia, "Competition and incentives in mortgage markets: The role of brokers," *Unpublished working paper*, 2019.

Salz, Tobias, "Intermediation and competition in search markets: An empirical case study," *Journal of Political Economy*, 2022, *130* (2), 310–345.

Shaffer, Greg, "Slotting Allowances and Optimal Product Variety," *The BE Journal of Economic Analysis & Policy*, 2005, *5* (1), 000010151515380637 1083.

Sidley Austin LLP, "CMS Finalizes Proposal to Overhaul Medicare Advantage Broker Payments," 2024. link.

Tilipman, Nicholas, "Employer incentives and distortions in health insurance design: Implications for welfare and costs," *American Economic Review*, 2022, *112* (3), 998–1037.

Tong, Joy Tianjiao, "Advising sophisticated customers: Evidence from health insurance brokers," *SSRN Electronic Journal, October. https://doi. org/10.2139/ssrn*, 2022, *4597241*.

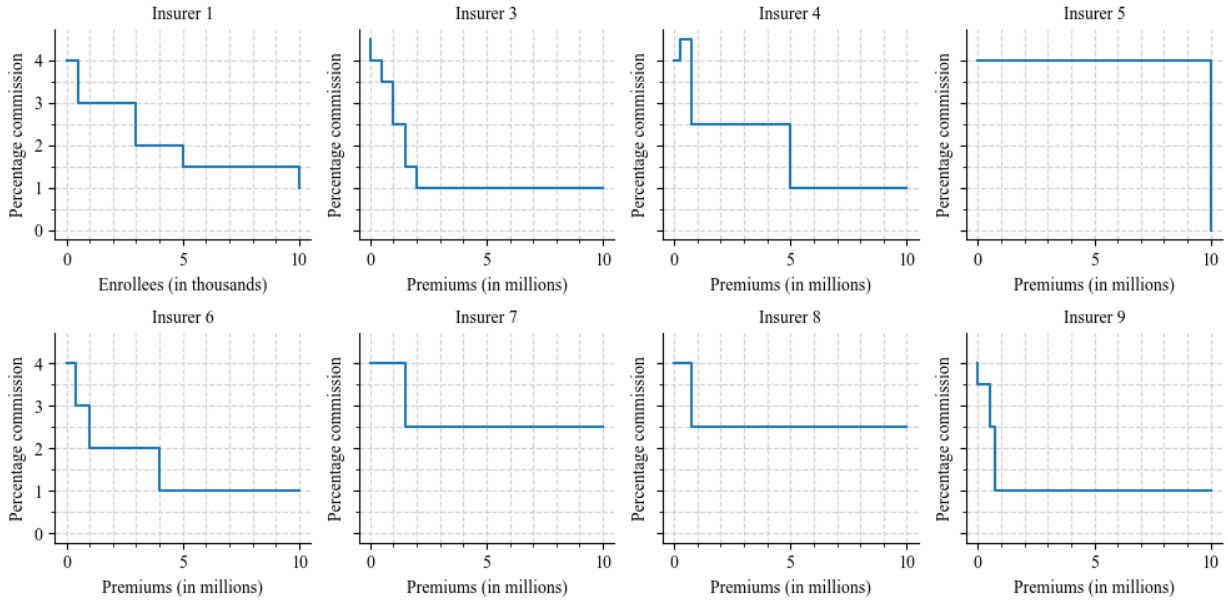
Wollmann, Thomas G, "Trucks Without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles," *American Economic Review*, 2018, *108* (6), 1364–1406.

Wooldridge, Jeff, "What's New in Econometrics? Lecture 6: Control functions and related methods," *NBER. http://www. nber. org/WNE/lect*, 2007, *6*, 420–445.

Online Appendix

A. Additional Tables and Figures

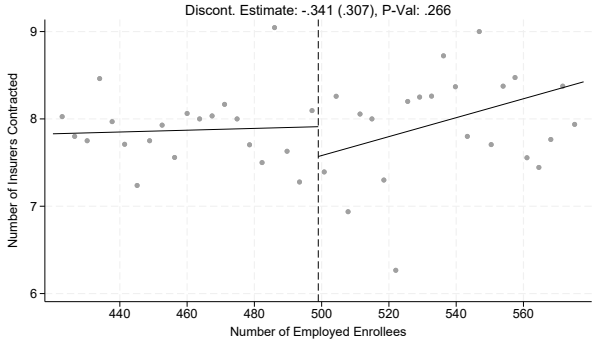
Figure A1. Commission schedules by insurer



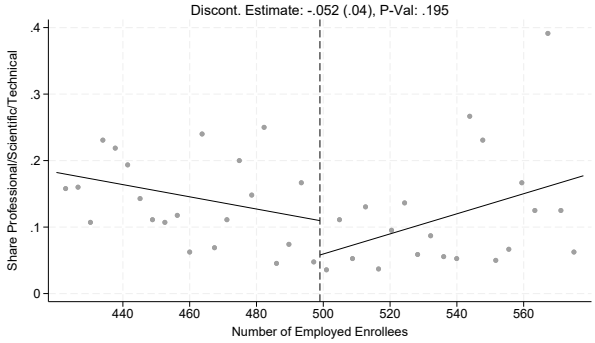
Notes: This figure plots the snapshot of comission schedules in 2017 for major insurers in NY ESI. Insurer 2 has per member per month commissions. Insurer 10's commission schedule is unknown.

Figure A2. RD Covariates Balance

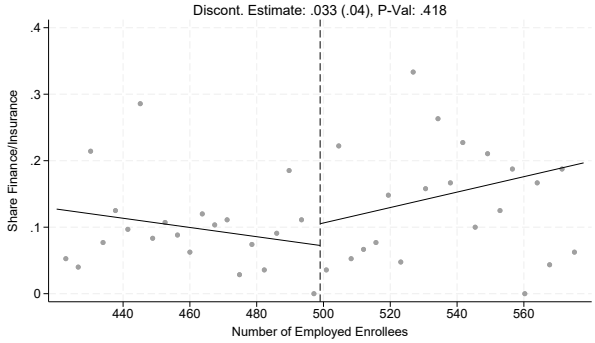
(a). Number of Insurers Contracted



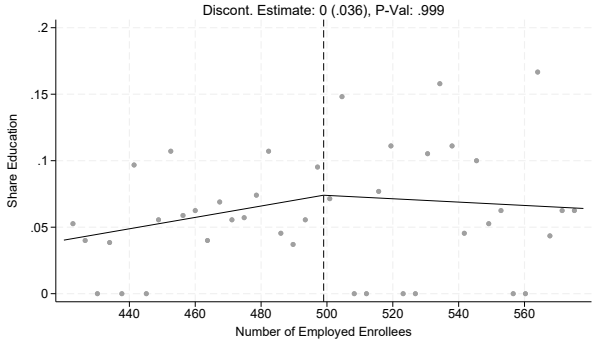
(b). Share Finance/Technical/Scientific



(c). Share Finance/Insurance

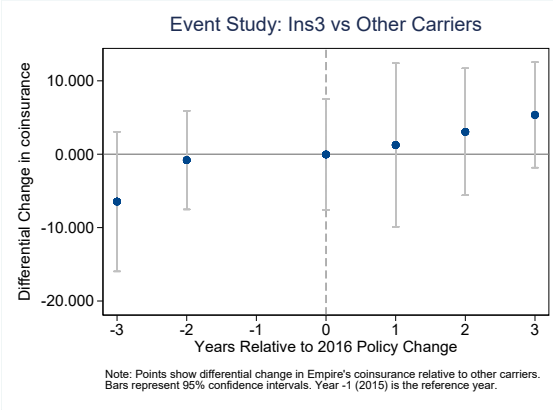


(d). Share Education



Notes: This figure plots the results of a regression discountinuity design that tests balances of covariates around the cutoff, with the number of enrolled employees as the running variable. Insurer 1's commissions drop from 4% to 3% at the cutoff of 500 enrolled employees. The analysis sample is restricted to fully insured employers who have Insurer 1 in their choice set and more than 100 enrolled employees in the carrier of choice. We use the CER-optimal bandwidth (Calonico et al., 2019).

Figure A3. Event Study: Plan Characteristics Response to Commission Change



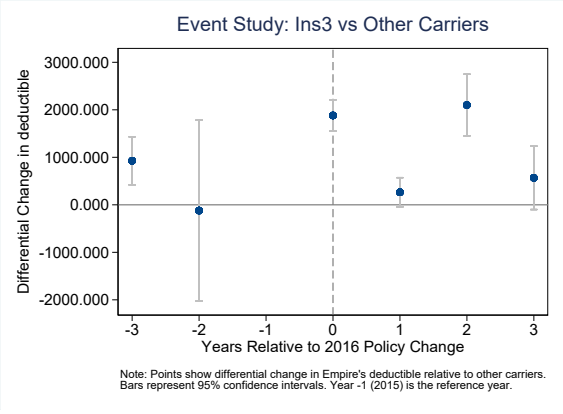


Figure A4. Effect of contracting networks on plan types

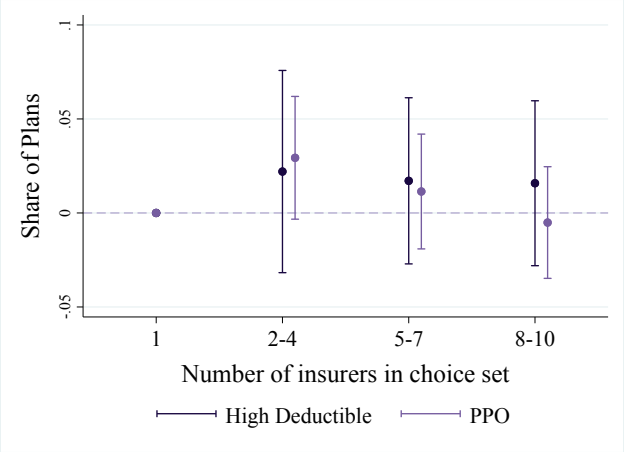


Table A1. Observed versus Simulated Equilibrium

	(1)	(2)	(3)	(4)	(5)
	Market Share			Price	
	Observed	Observed price Simulated choice	Simulated price Simulated choice	Observed	Simulated
Employers with brokers					
Aetna	13.4	14.2	14.2	9621	9573
MVP Health	4.8	3.7	3.7	6974	6445
Empire/Anthem	8.5	10.0	10.0	9651	9327
Oxford/United	24.3	22.5	22.5	9988	9882
Healthnow	6.3	5.9	6.0	8331	8329
Connecticut/Cigna	12.4	12.4	12.4	8955	8810
HIP of NY/Emblem	2.8	2.4	2.4	8809	8809
Excellus	22.5	23.7	23.7	8363	8155
CDPHP	2.6	2.6	2.6	8463	8463
Independent Health	2.3	2.5	2.5	7753	7749
Employers without brokers					
Uninsured	44.0	17.6	22.1	0	0
Aetna	8.2	8.9	11.2	9522	9500
MVP Health	1.6	0.4	0.4	6512	6566
Empire/Anthem	4.9	17.1	10.0	16713	9628
Oxford/United	15.2	22.4	29.4	9302	9221
Healthnow	5.4	2.3	3.0	8815	8852
Connecticut/Cigna	7.6	6.0	7.1	8849	8939
HIP of NY/Emblem	1.6	15.7	3.6	18681	9148
Excellus	2.7	8.1	10.9	8476	8421
CDPHP	2.2	1.1	1.3	8489	8569
Independent Health	6.5	0.5	0.8	7513	7553

B. Additional Details on Data

B1. Panel Construction Using Form 5500 Filings

Panel Construction. Below, we describe the data cleaning steps used to construct the final analytic dataset from Form 5500 filings and related sources.

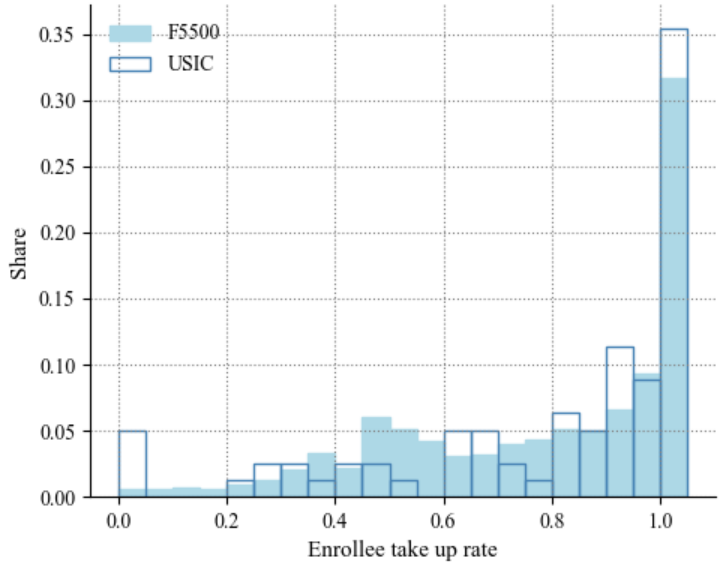
First, we begin with Schedule A filings of Form 5500, which records employers' benefit contracts. Each filing reports the carrier, premiums paid, and the number of covered enrollees (employees and dependents) for each plan. Second, we restrict attention to health insurance filings. Other benefit types, such as life, disability, and unemployment insurance, are excluded. Third, we clean and standardize the carrier identifiers. Schedule A provides insurers' NAIC codes, which we use to identify health insurers. We harmonize these codes and consolidate carriers under common parent companies. Fourth, we limit the sample to New York employers. Fifth, we resolve multiple forms within a reporting-carrier-employer-year. If the form is an amendment, only the final submission is retained. If several forms correspond to different products under the same carrier, we collapse them into a single record. Sixth, we incorporate information on brokers. Schedule A Part I reports broker names, locations, and commission payments. We standardize broker names and merge them to the main Schedule A record at the reporting-carrier-employer-year level. Seventh, we merge the sample with the main Form 5500 database, which adds the employer-year level information, including the total number of enrolled employees that aggregates across carriers, industry, and location. We aggregate to one observation per employer–year–carrier, applying the same collapsing rules as above. If multiple brokers are associated with a filing, we retain the broker accounting for the largest share of premiums. Finally, we augment the dataset with Form 5500 group health plan information from the Bureau of Labor Statistics, which provides a self-insured indicator. We also bring in the BLS-recovered total number of enrolled employees at the employer-year level, aggregating across carriers when multiple carriers are offered.

The resulting dataset is structured at the employer–year–realized insurer level. Each record contains the insurance carrier contracted by the employer, the broker through which the contract was obtained, the number of enrollees (employees plus dependents), premiums paid, the number of enrolled employees (unlike enrollees or premiums, this is an aggregate across carriers), and employer characteristics.

Takeup Rate. We next describe how we measure firm size and insurance take-up. Form 5500 filings contain records on both welfare benefits (including medical, drug, dental, vision, disability, and life insurance, and others) and pension benefits (including defined benefits plans, defined contribution plans, and pension plans). We use these filings to proxy for firm size by taking the maximum number of enrolled employees across all reported welfare and pension plans for each employer, assuming that every employee takes up at least one benefit the employer offers. Because not all employees necessarily take up the benefits offered, this measure may underestimate firm sizes.

We then compute the take-up rate of employer-sponsored insurance as the ratio of enrolled employees in health insurance benefits to our imputed firm size. We benchmark our imputed take-up rates against two data sources. First, the average health insurance take-up rate in our sample is 79%, which is close to 71% for firms of similar sizes in the Medical Expenditure Panel Survey. Second, we compare our imputed take-up rates for the large national insurer against those reported in its records, which provide exact firm size and

Figure A5. Benchmark Takeup-Rate Across Data Sources



Notes: This figure plots imputed takeup-rate for tha large national insurer in Form 5500 and its accurate enrollment records.

enrollment information. Figure A5 shows our imputed takeup rate distribution for that large representative insurer closely resembles the true takeup measure. This similarity provides support for the validity of our firm size and take-up rate measures.

Household Characteristics A limitation of the Form 5500 data is that, while we observe the number of employees who enroll in coverage through the firm, we do not observe the demographics of those employees. We thus constructed a simulated sample of employee demographics relying on distributions from an external data source: the Medical Expenditures Panel Survey (MEPS). The MEPS Household Component (HC) is a nationally representative sample of households and contains information on demographic characteristics and health insurance coverage information. Using MEPS, we identify individuals with employer-sponsored insurance through their own employer and construct demographic distributions of age, gender, and household size by firm size and industry, which we then sample into the Form 5500 dataset.

Premiums The Form 5500 contains the total premiums earned for each employer, which we use as our primarily variable for constructing the prices faced by employers (per employee per month) for each insurance plan. As the Form 5500 data only contains *realized* premiums (i.e. premiums conditional on firms having chosen the insurer), we do not directly observe premiums of the full choice set facing the employer. Thus, we rely in imputation methods to simulate the likely premiums that a firm would face for an unobserved plan.

To do so, we implement a hedonic premium regression in the spirit of Dafny et al. (2013). Specifically, for our New York sample between 2011-2020, we regress the log of premiums per enrollee on a rich set of observables, including insurer fixed effects, geographic region, firm size, year, and industry. We also leverage broker network information in the regression by including indicators for the number of insurers in a broker's network for the specific employer. Given that a key factor we are exploring is the competitive effects of broker-insurer networks, we expect premiums for the same insurer and same employer to vary

depending on the broker.

B2. Data from the Large National Representative Insurer

We supplement our analysis with data from a large health insurance provider. We focus on the data for New York State. The data provides information on health insurance coverage for employers in the small group market. The data cover the years 2013 through 2017. We observe data at the employer-year and enrollee-year levels as well as linked enrollee claims data. The employer-year level data include the number of employees eligible for health insurance coverage, the number of employees enrolled, the number of plans offered by the employer, characteristics of each plan, including plan type (EPO, PPO, POS), coinsurance, physician copay, in-network deductible, in-network out-of-pocket maximum, out-of-network deductible, and out-of-network out-of-pocket maximum, and the total premium paid to the insurer for each plan. The enrollee-year data include age, gender, and health plan choice from the employer's offerings. We calculate a household's premium by multiplying the enrollee's premium amount by the number of enrollees in each household.

One limitation of the data is that while we know the *number* of eligible employees who do not take up offered coverage, we do not know any of their characteristics. Thus, we supplement the data again with demographics from the MEPS HC, as with the Form 5500 sample above. Here, we identify individuals who worked for a small group employer and were offered health insurance from that employer.

Table A2 reports the sample statistics at the enrollee level. The first two columns represent all firms, while the last two columns represent firms with multiple plans. We include the distribution of employees who take up coverage at the firm (i.e., those covered in the large national insurers' records) and those who opt for the outside option (i.e., calibrated from the Medical Expenditure Panel Survey). Employees who do not participate in the employer-sponsored plan tend to be older on average, are less likely to be men, and are more likely to have a household size of one. There is no clear pattern of firm size differences, except that they are less likely to be 100-499 firms, and more likely to be smaller. Comparing the single-plan and multi-plan firms, the age distribution and family size distribution appear similar, but multi-plan firms are clearly larger. Enrollment-weighted plan characteristics, such as premiums and deductibles, are also similar.

Table A3 reports sample statistics at the firm level. Multi-plan firms are bigger and have slightly higher take-up rates. In general, approximately 60% of all firms have 1-49 employees, 30% have 50-99 employees, and 10% have more than 100 employees. However, for multi-plan firms, the rates are 10%, 60%, and 30%, respectively.

B3. Clarivate Data and Imputed Product Characteristics

We derive insurance product characteristics from the Clarivate Real-World Data (RWD) Product, previously known as Decision Resources Group Real-World Data. This dataset covers over 300 million longitudinal patient records in the United States, constructed from medical claims submitted through billing software used by a wide array of healthcare providers, including hospitals, physician offices, long-term care facilities, etc. The raw dataset contains billions of patient claims with detailed information on providers, patients, and payers, and a granular breakdown of patient payments such as deductibles, coinsurance, and copayments. In addition, the data report full diagnosis information in terms of ICD codes, medical procedures via HCPCS

Table A2. Summary Statistics by Coverage Decision and Firm Type

	All Firms				Multi-Plan Firms			
	(1)		(2)		(3)		(4)	
	Inside		Outside		Inside		Outside	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
A. Demographics:								
Age 18-34	0.623	0.485	0.487	0.500	0.619	0.486	0.479	0.500
Age 35-54	0.283	0.451	0.339	0.473	0.293	0.455	0.342	0.474
Age 55-64	0.074	0.262	0.111	0.314	0.070	0.255	0.114	0.317
B. Firm Size:								
Firm Size: 1-49 employees	0.063	0.243	0.082	0.274	0.019	0.136	0.032	0.176
Firm Size: 50-99 employees	0.497	0.500	0.439	0.496	0.507	0.500	0.523	0.499
Firm Size: 100-499 employees	0.437	0.496	0.303	0.460	0.474	0.499	0.445	0.497
Firm Size: 500+ employees	0.003	0.057	0.176	0.381	0.000	0.000	0.000	0.000
Percent Male	0.526	0.409	0.436	0.496	0.535	0.415	0.440	0.496
Household Size	1.874	1.316	1.000	0.000	1.818	1.293	1.000	0.000
C. Plan Characteristics:								
Monthly Premium	103.377	77.757			107.285	75.152		
Deductible	566.032	920.844			634.305	974.874		
Out-of-Pocket Maximum	2407.116	2019.124			2637.779	2024.493		
Coinsurance Rate	0.682	0.432			0.719	0.410		
EPO	0.329	0.470			0.375	0.484		
POS	0.661	0.473			0.618	0.486		
PPO	0.009	0.096			0.007	0.083		

Notes: Summary statistics at the household level. Inside option refers to households enrolled in employer coverage, outside option to those declining coverage. Take-up rate is 59.9% for all firms and 76.2% for firms offering multiple plans.

Table A3. Firm Characteristics Summary Statistics

	All Firms		Multi-Plan Firms	
	Mean	SD	Mean	SD
A. Firm Size:				
Number of Enrollees	34.5	44.4	69.7	48.8
Number Eligible	45.5	65.6	81.9	38.8
B. Size Distribution:				
1-49 employees	0.6	0.5	0.1	0.3
50-99 employees	0.3	0.5	0.6	0.5
100-499 employees	0.1	0.3	0.3	0.4
500+ employees	0.0	0.1	0.0	0.0

Notes: Summary statistics at the firm-year level. N = 590 total firms, of which 194 offer multiple plans.

codes, patient demographics (e.g., 3-digit ZIP code, gender, age), and insurer type (commercial, Medicare, or Medicaid). A key feature of the dataset is its open-network structure, which enables us to track patients over time and across providers, even when patients switch insurers within the observation period.

We limit our analysis sample to New York and include all commercial claims, which encompass both professional and institutional claims, and cover patient visits to hospitals (inpatient and outpatient), ambulatory surgical centers (ASCs), clinics, and physician offices. Starting from the raw claims data, we restrict the sample to patients aged between 18 and 64 and exclude individuals enrolled in privately administered Medicaid or Medicare plans. We further require all claims to have non-missing information on patient gender, year of visit, total allowed amounts, and detailed out-of-pocket components, including deductibles and coinsurance. Because most health plans do not count preventative care toward deductibles or maximum

out-of-pocket limits, we exclude claims associated with preventative care procedures to avoid bias. Our final analysis sample focuses on the four largest commercial payers in New York State—Aetna, Connecticut/Cigna, Empire/Anthem, and Oxford/United—while pooling all other relatively smaller insurers into a single category.[1] labeled as "Other Insurers." In total, the sample comprises more than 12.5 million commercial claims spanning from 2013 to 2022.

To infer plan characteristics across carriers and regions, we aggregate claims for each patient within a calendar year, summing allowed amounts, deductibles, coinsurance, and other out-of-pocket costs. Although the dataset does not provide the start and end dates of annual health plan cycles, we use the calendar year as a proxy. Deductibles accrue until a patient reaches the plan's annual maximum deductible, at which point the insurer begins covering subsequent medical expenses (with patients still responsible for coinsurance or copayments). To ensure that we are observing patients who likely reached their deductible limits, we focus on those with annual medical expenditures exceeding $6,000. For carrier-region-year cells with fewer than 20 such patients, we relax this threshold to $2,000 to preserve sample size. We define a cell as the interaction of a carrier, region, and year. Regions are based on the ten economic development regions designated by the New York State Department of Economic Development: Western NY, Finger Lakes, Southern Tier, Central NY, Mohawk Valley, Capital Region, Mid-Hudson, New York City, Long Island, and North Country (https://esd.ny.gov/regions). Each patient is mapped to a region according to their ZIP code.

Within each cell, the inferred deductible is calculated as the average deductible paid by patients whose annual medical spending exceeds $6,000. The inferred coinsurance rate is defined as one minus the mean ratio of patients' total coinsurance payments to their total medical expenses net of deductibles, thereby capturing the average share of costs borne by insurers. Summary statistics of the inferred deductibles and coinsurance rates are presented in Table A4.

Table A4. Plan Characteristics

		Sample Mean					
	N	Patient Age	Female	Patient Paid	Insurer Paid	Deductible	Coinsurance (%)
Oxford/United	2,784,846	45.0	0.60	266.1	574.4	702.2	84.8
Aetna	730,811	42.6	0.58	285.9	416.9	996.8	88.4
Connecticut/Cigna	527,818	41.2	0.60	218.0	322.0	1088.2	86.0
Empire/Anthem	472,644	45.3	0.58	399.7	546.8	929.8	70.3
Other Insurers	1,696,603	44.7	0.61	123.3	432.0	817.9	84.4

[1]The cleaned sample provides relatively limited coverage of Excellus. As a result, we combine Excellus with the group of other insurers and instead single out the next largest insurer, Empire/Anthem, as a separate category.

C. Additional Model and Estimation Details

C1. Employee Surplus From Having Access to Employer-Sponsored Insurance

Below, we derive employees' decisions on whether to participate in the insurance plan that their employer sponsors and the surplus they derive from it.

Employees' Decision. The employee i decides whether to participate in the plan offered by insurer j through their employer e. For notational simplicity, we omit the subscript for geographic markets. Employee i's flow utility from choosing product j is:

$$u_{ij} = -(\alpha_0 + \underbrace{\alpha_1 z_i}_{\substack{\text{employee} \\ \text{demo.}}}) \underbrace{\iota p_j}_{\substack{\text{employee paid} \\ \text{premium}}} + (\beta_0 + \beta_1 z_i) \underbrace{x_j}_{\substack{\text{financial} \\ \text{char.}}} + \delta_j^w + \varepsilon_{ij}; \quad u_{i0} = \varepsilon_{i0}. \tag{A1}$$

where p_j is the insurance premium of plan j; ι is the employee's share of premium contribution; x_j is the plan's financial attributes, including deductibles, coinsurance rates; δ_j^w is insurer-specific brand effects perceived by workers. ε_{ij} is independent and identically distributed from the extreme-value type-I distribution. We allow each employee's preferences for insurance attributes to vary by their characteristics z_i, including age, gender, and household size. The outside option $j = 0$ is opting out of employer-sponsored plans, where the employee remains uninsured or enrolls in spousal coverage or other insurance programs.

Let s_{ej}^w denote the share of employees who take up plan j sponsored by employer e. The average welfare of employee i of employer e from having access to the plan j is

$$\mathcal{W}_{ej} = \frac{1}{n_e} \sum_{i \in e} \ln(\exp(u_{ij}) + \exp(u_{i0})) \frac{1}{\alpha_i}, \tag{A2}$$

where n_e is the employer size. We set $\mathcal{W}_{e,j=U} = 0$ if the employer does not offer insurance coverage.

Estimation. Employees' primitives include their premium sensitivity α_i, preferences for financial attributes β_i, and brand perception δ_j^w.

Since Form 5500 lacks information on the exact attributes of insurance plans, we supplement it with enrollment records from a large, nationally representative insurer and conduct the estimation in two steps. In the first step, we exploit enrollment records of fully insured plans of the national insurer to estimate the premium sensitivity α_i and preferences for financial attributes β_i by worker demographics. We use maximum simulated likelihood estimators that derive choice patterns from equation (A1), substituting brand effects with an inside-option fixed effects.[2][3] We assume preference estimates α_i, β_i from this single large insurer accurately depict valuations for plan attributes of the NY employee population and are consistent across insurers (except insurer brand effects).

[2] Since both Form 5500 and the large insurer only include aggregate premium payments, we calibrate the premium contribution of employees to be 25% from the Kaiser/HRET annual employer survey. https://www.kff.org/series/employer-health-benefits-survey/

[3] We observe enrollment records, plan attributes, and demographics for employees who opt into employer-sponsored coverage in USIC data, and simulate the demographics for employees who are eligible for employer coverage but do not take it up using the empirical distributions of demographics from the Medical Expenditure Panel Survey (MEPS) for employees in a similar enrollment situation, based on age, gender, and household sizes.

In the second step, we utilize Form 5500 to estimate brand effects for all insurers, δ_j^w. To do so, we use three inputs from outside the Form 5500 data: insurer-specific distributions of product attributes from RWD Clarivate, market-by-firm-type-specific distributions of employee demographics from MEPS, and employees' preferences for each product attribute estimated from the large national insurer. We then use maximum likelihood to recover δ_j^w that rationalize within-firm take-up rates for employers offering that particular insurer.[4] Appendix B1 describes how we calculate take-up rates in Form 5500.

Identification. In the first step, variations in choice patterns and their correlation with product characteristics identify employees' preferences for prices and financial attributes, α_i and β_i. Since employees at the same firm face the same set of insurance plan choices, this within-employer method controls for unobserved employer-level factors that might affect selection into plans. In particular, preferences for product attribute parameters are primarily identified through enrollment decisions made by employees in firms that offer multiple plan options with different attributes from the same insurer. The overall valuation of the inside option (as opposed to taking up spousal coverage or remaining uninsured) is primarily identified by variation in take-up rates across firms that choose the large national insurer, conditional on attributes and demographics.

Insurers' knowledge of employees' unobserved preferences across plans when making pricing decisions creates a correlation between residuals and prices. We construct a Nevo (2001)-style instrument to address this endogeneity concern. We instrument for prices facing employees using the average prices of plans with the same characteristics (e.g., high versus low deductible) offered by other employers in the same firm size category (e.g., 100-499 or over 500 employees) in different markets. The instrument's relevance is satisfied because plans of similar characteristics offered by employers with similar bargaining leverage with insurers have similar cost structures. We assume that the unobserved preferences are idiosyncratic across employers after controlling for plan types and attributes. Namely, average prices of other employers do not reflect unobserved differences in preferences that affect choice probabilities for one particular employer.[5] We implement the instrument using a control-function approach following Wooldridge (2007).

In the second step, the brand effects, δ_j^w, are identified by variations in take-up rates across employers that offer different insurers in Form 5500. Holding fixed plan attributes and employees' preferences for these attributes constant, variations in the take-up rates reflect the perceived quality of each insurer relative to the same outside option of being uninsured or having non-employer-sponsored insurance. In other words, variations in take-up rates of the employer-sponsored options identify insurer-specific brand effects.

Results. Table A5 reports estimates of employee preferences from the conditional logit model. The first two columns present results with no instruments, and the second two columns present instrumented results. Across all specifications, employees prefer plans with lower premiums or a higher degree of financial generosity, such as low deductibles or coinsurance rates. However, the parameters on other attributes tend to diverge from there. Having a full set of interactions of demographics and attributes (columns 2 and 4) tends to yield sensible results on premium preferences by age. Namely, consumers become less price sensitive as

[4] As Form 5500 does not break the number of enrolled employees by insurer, we conduct second-step estimation only for employers that offer a single insurer, which accounts for 87% of employer-year pairs in our sample.

[5] This exclusion restriction may be violated if unobserved demand shocks correlate across firms in different markets, for example, when all employers with 100-499 employees have similar unobserved preferences for certain plan attributes not controlled for in our model. Given that preferences for plan attributes are quite local, this seems likely in this setting.

the average age of the household increases. Moreover, older consumers receive less utility from a higher out-of-pocket maximum compared to their younger peers, who may be less sick and more liquidity-constrained. Larger households are significantly less sensitive to price, as expected.

Table A5. Estimates of Employees' Preferences

	No Instruments		**With Instruments**	
	(1)	(2)	(3)	(4)
Monthly Premium	-13.296***	-14.386***	-67.093***	-65.231***
	(4.376)	(4.694)	(2.392)	(2.234)
Deductible	-0.230	-0.280	-0.797***	-1.018***
	(0.141)	(0.190)	(0.033)	(0.043)
Out-of-Pocket Maximum	0.339***	0.380***	0.230***	0.485***
	(0.103)	(0.134)	(0.014)	(0.034)
Coinsurance Rate	1.022	0.995	2.934***	2.727***
	(1.246)	(1.260)	(0.170)	(0.173)
Inside Option	-2.136	-2.090	4.748***	3.788***
	(1.401)	(1.421)	(0.377)	(0.396)
Premium Interactions:				
35-54 × Monthly Premium	-0.763	0.990	1.221***	3.156***
	(0.568)	(0.894)	(0.281)	(0.655)
55-64 × Monthly Premium	-2.091**	2.459*	-0.133	2.665***
	(0.999)	(1.441)	(0.474)	(0.959)
x Male	0.850	0.807	-0.407	-0.171
	(0.557)	(0.557)	(0.281)	(0.294)
x HH Size	11.772***	11.859***	13.112***	16.482***
	(1.767)	(2.104)	(0.351)	(0.587)
Deductible Interactions:				
35-54 × Deductible		0.036		0.053
		(0.085)		(0.044)
55-64 × Deductible		0.151		0.088
		(0.149)		(0.061)
x HH Size		0.018		0.199***
		(0.074)		(0.025)
Out-of-Pocket Maximum Interactions:				
35-54 × Out-of-Pocket Maximum		-0.097*		-0.104***
		(0.051)		(0.032)
55-64 × Out-of-Pocket Maximum		-0.259***		-0.152***
		(0.089)		(0.045)
x HH Size		0.010		-0.146***
		(0.047)		(0.020)
Observations	81,412	81,412	76,522	76,522

Notes: Bootstrapped standard errors are reported in parentheses. Monthly premiums, deductibles, and out-of-pocket maximums are measured in thousands of dollars. The estimation sample is fully-insured firms with more than 100 employees. The base age group is 18-34 years old. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The biggest differences, however, arise from the inclusion of instruments (columns 3 and 4). In our preferred specification (column 4), the magnitude of the premium coefficient quadruples in size, suggesting a significant downward bias in the absence of instruments. Moreover, other coefficients flip signs towards a more expected direction. Most notably, the value of the inside option (being insured through the firm) was negative and insignificant without instruments, suggesting consumers do not value insurance. This is clearly the result of the model being unable to rationalize why consumers forego insurance at low levels of price

sensitivity. However, with instruments, consumers can be price sensitive and value insurance, which aligns more intuitively.

We convert the point estimates to derived statistics. First, Table A6 derives own-premium elasticities by demographic groups. On average, a 1% increase in monthly enrollee-share premiums results in a decrease of about 3.5% in market share. This varies significantly by household size, with single households being the most price sensitive, while households with 3 members are significantly less price sensitive (with the latter having an elasticity of -1.04 on average). Within those groups, there is also significant variation: three member households, for instance, vary between elasticities of -0.69 and -1.47.

Table A6. Own-Price Elasticities by Consumer Type

	Mean	Std. Dev.	25th pctl	75th pctl	N
Total	3.59	2.45	1.70	5.26	15198
A. Household Size					
1	4.78	1.98	3.68	5.85	10008
2	2.40	1.20	1.60	3.23	2221
3	1.04	0.63	0.69	1.47	1347
B. Firm Size					
1-49 employees	3.08	2.37	0.93	4.57	1010
50-99 employees	3.65	2.45	1.71	5.21	8051
100-499 employees	3.58	2.44	1.70	5.37	6137

Notes: Elasticity estimates from the conditional logit model with premium interactions and control function.

While higher than existing elasticity estimates, there are a few notes that distinguish our data from prior literature. First, most of the employers in our data are small firms, the enrollees of which are thought to have higher price elasticities. Second, the single households make up a majority of the data, driving the estimates higher. Third, most firms in our data only offer one plan, meaning in a majority of cases, this can be thought of as an extensive margin elasticity, which tends to be higher as individuals could simply take up spousal coverage.[6] Finally, our estimates are broadly consistent with recent literature estimating demand for large firms, including Ho and Lee (2017), which finds elasticities ranging from -1.23 to -2.95 for *total* (enrollee + employer) shares.

We next report the estimated value of the plan, along with its characteristics other than price. A 10pp increase in the coinsurance rate is valued at approximately $112 annually for a household of two, while a $1,000 decrease in deductible is valued at approximately $233 by two-member households and approximately $800-$1000 for larger families. The average surplus from having access to insurers on NY ESI is between $103 - $142 per employee per month, though with considerable variation across firm type, region, and industry, as is reported in Table A7.

C2. Distance is the Primary Factor in Broker-Employer Matches

In this section, we provide empirical evidence that geographic proximity is the most important determinant of employers' broker choices: employers overwhelmingly prefer to work with the closest broker in their

[6]While the avearge elasticities in works such as Cutler and Reber (1998) are between -0.3 and -0.6, the *extensive margin* elasticities are around -2. The remaining discrepancy could be explained by our relatively smaller firm size and higher monthly premiums.

Table A7. Willingness to Pay by Insurer

Insurer	Mean	Min	Max
Aetna	124.22	102.10	147.05
MVP Health	142.72	120.74	164.87
Empire (Anthem)	115.90	67.83	147.63
Oxford (United)	115.05	95.55	142.36
Healthnow	115.27	99.20	145.41
Connecticut (Cigna)	142.86	124.59	182.43
Emblem/HIP of NY	103.63	88.82	118.62
Excellus	128.41	107.01	164.75
CDPHP	113.92	96.86	139.48
Independent Health	114.73	102.05	129.22

Notes: Willingness to pay estimates are in dollars per employee per month.

community.

We estimate a series of logit models of employer–broker matches. For each employer in the Form 5500 data, we define the potential broker choice set as all active brokers having at least one branch in New York State in a year. In the baseline model (Model 1), the only covariate is the log of the distance between the employer's location and the broker's nearest office. Model 2 replaces distance with the number of insurers in the broker's contracting network. In Model 3, we include both the log distance and the number of insurers in the contracting network, as well as an indicator for whether the broker is an individual agent. Finally, Model 4, labeled as the full model, adds interaction terms between all three broker-level variables and the employer's industry code, along with additional employer-level controls: the number of brokers hired, the number of carriers ultimately selected, the number of plan participants, and an indicator for whether the plan is fully insured. We then compare model performance using the Pseudo R-square across specifications, as reported in Column (1) of Table A8.

Table A8. Employer-Broker Matching

	Evaluation Scores	
	(1)	(2)
Models	Pseudo R2	AUC Distance
Logit with Dist (Model 1)	0.1833	0.8340
Logit with #Carriers in Network (Model 2)	0.0056	0.5836
Logit with Dist+#Carriers in Network+Inidividual Agent (Model 3)	0.1921	0.8582
Logit with Full Model (Model 4)	0.1970	0.8595

Notes: This table reports the predictive performance of alternative models in explaining employer–broker matching outcomes.

As shown, the Pseudo R-square is substantially larger in Model 1 (0.1833) than in Model 2 (0.0056), indicating that geographic distance between employers and brokers is far more predictive of matching outcomes than the number of carriers in a broker's network. Adding further covariates provides only modest improvements in explanatory power: Model 3 increases the Pseudo R-square by 0.0088 relative to the baseline, while the full specification in Model 4 raises it by just 0.0137.

Instead of relying on Pseudo R-square, we evaluate model performance using an alternative measure—

the Area Under the Receiver Operating Characteristic Curve (AUC) distance. Specifically, within each choice set, we compare the predictive probability assigned by the logit models to the realized (chosen) option versus each alternative. If an alternative receives a higher predicted probability than the chosen option, it is assigned a score of 1; if lower, a score of 0; and if equal, a score of 0.5. The AUC distance is then computed as the average score across all comparisons in the whole choice panel. By construction, a random choice model, where each option has equal probability $1/N$, yields an AUC distance of 0.5. A higher AUC distance therefore indicates stronger predictive performance.

Column 2 of Table A8 reports the AUC values for our four models. Model 1 achieves an AUC distance of 0.8340, substantially above 0.5, confirming that the baseline model significantly outperforms a random choice model. Consistent with the Pseudo R-square results, Model 1 also exhibits much higher predictive power than Model 2, underscoring that geographic distance is the dominant predictor of employer-broker matches. Adding additional covariates in Models 3 and 4 modestly improves predictive performance, but the gains are small relative to the explanatory power of distance alone.

Taken together, these results highlight distance as the single most important determinant of employer-broker matching. This finding supports our simplifying assumption that, conditional on using a broker, employers are most likely to select the broker located closest to them.