INSURER PRICING AND CONSUMER WELFARE: EVIDENCE FROM MEDIGAP*

Amanda Starc†

February 22, 2012

Abstract

While adverse selection is often blamed for inefficiently high insurance premiums, imperfect competition is also a pervasive feature of many health insurance markets. In Medicare Supplement Insurance (Medigap), two firms control nearly three-fourths of the market, and premiums exceed claims by 30 percent. I find that, while adverse selection can restrain markups, low price elasticity and consumers' brand preferences create incentives for firms to engage in substantial marketing and price above cost. I conclude that the strategic behavior of insurers facing relatively inelastic demand is critical in explaining poor market performance and loss of consumer surplus, and find that insurers do not capture all of the rents in this market.

JEL Codes: I13

*I am grateful to numerous seminar participants at Harvard, Wharton, University of California-San Diego, University of Chicago-Booth, Duke, Northwestern-Kellogg, Princeton, and Stanford, Susan Athey, Michael Dickstein, Jon Kolstad, Greg Lewis, Julie Mortimer, Chris Nosko, Mark Pauly, Bob Town, Dennis Yao, and especially David Cutler and Ariel Pakes for helpful comments. All errors are my own.

†The Wharton School, University of Pennsylvania, 3641 Locust Walk, Philadelphia, PA 19104, astarc@wharton.upenn.edu.

1 Introduction

The Patient Protection and Affordable Care Act, signed into law in March 2010, represents a substantial health insurance expansion. The legislation relies on the creation of exchanges and regulation of individual insurance markets to accompany an individual mandate to purchase insurance. Underlying these exchanges is the idea that comparison shopping by consumers will lead to more efficient outcomes. However, there is little empirical evidence on the nature of consumer demand for differentiated insurance markets or the operation of individual health insurance markets.[1] Indeed, many such markets are characterized by high concentration, high premiums, and poor outcomes for consumers.

The market for Medicare Supplement Insurance, commonly called Medigap, is one example. In this market, consumers face dramatic price differences for contractually identical policies, and two firms control nearly three-fourths of the market. Further, the markup of premiums over claims is nearly 30 percent, a full 15 percent higher than the guidelines in the heath reform bill. These features are not unique to Medigap insurance. In the past decade, the growth in health insurance premiums has rapidly outpaced inflation, and this combination of rising premiums, profits, and concentration has led to arguments for increased antitrust enforcement and stricter regulation from the left, as well as counterarguments questioning the efficacy and prudence of such policies from the right.

In this paper, I use a structural model to describe and analyze the Medigap market. This insurance insures consumers against the cost-sharing provisions associated with Medicare, the public health insurance program for elderly Americans. The policies in the Medigap market are standardized, and there is a large set of potential suppliers. Rather than experiencing adverse selection, Fang, Keane, and Silverman [2008] argue that Medigap policies are purchased by relatively healthy consumers, implying that the market experiences ad-

[1] As noted by Dafny, Ho, and Varela [2010], most non-elderly Americans purchase policies through an employer, who mediates the exchange and typically offers a small number of choices. They argue that this reduction in the set of choices available to consumers creates substantial deadweight loss.

vantageous selection on the extensive margin.[2] The market is heavily regulated: firms may only offer standardized policies and are required to pay out 65 percent of the premiums they collect. Under many circumstances, they cannot price based on an individual's health status.

Despite a concerted policy effort to ensure a well-functioning market, the market does not seem to lead to efficient outcomes for consumers. I consider a number of explanations for the structure of the Medigap market. First, adverse selection could lead to higher premiums or an underprovision of insurance. Ultimately, I argue that both the scope for potential selection and its empirical importance are limited in this setting. However, in addition to asymmetric information, a low demand elasticity coupled with brand preferences may cause welfare losses for consumers. Insurance policies are complicated financial products, and it is extremely difficult for average Americans, let alone a potentially vulnerable and less technology-savvy group of elderly Americans, to understand or research all of the contractual details. Consumers are willing to pay more for insurance from a "trusted name," potentially driving up premiums as consumers trade-off brand preferences for price. In turn, insurers will compete to be such a trusted name, often incurring large marketing expenditures.

I model demand using a discrete choice framework that allows for the calculation of elasticities and estimation of the relationship between premiums and claims (allowing preferences over plan characteristics to be correlated with costs). I incorporate consumer heterogeneity in preferences to allow for flexible forms of adverse selection, which I define loosely as a correlation between consumer heterogeneity in preferences and variation in claims. In my model, this relationship could be generated by differential search or limited cognition, but ultimately, the goal is to control for the potential that selection is leading to suboptimal market outcomes. The estimates of consumer preferences and the relationship between consumer characteristics and claims are used in a model of profit-maximizing firm behavior. The moments generated by the supply-side assumptions allow for the estimation of any

[2] This study will, by contrast, focus on adverse selection on an intensive margin (price).

marginal administrative costs that firms incur in addition to claims.[3] Finally, assumptions about entry behavior determine fixed costs of advertising and operation for the two largest insurers.

The results show that demand for Medigap insurance is best characterized by a low demand elasticity and very strong brand preferences. This average own-price elasticity, at -1.1, allows insurers to price far above marginal cost, while the strong brand preferences provide big incentives for insurers to invest in brand-promoting activity, such as marketing. Not all of the rents are captured by insurers: many of these rents are realized by other participants in the market, such as the interest group AARP, while other rents are dissipated through the payout of commissions to insurance brokers.

Using these estimates and the equilibrium model, I show that policies that control insurer incentives on the margin, such as perfect risk adjustment or capped commissions, are not effective at improving consumer welfare. To change market outcomes dramatically, public policy needs to find a instrument that shifts consumer preferences. For example, limiting the extent to which insurers can differentiate their products with marketing leads to stiffer price competition associated with a 9 percent fall in premiums. Subsidizing a firm to compete directly with UnitedHealth reduces premiums by 20 percent.

While the potential for welfare losses due to adverse selection is well documented in the theoretical and empirical literature[4], less attention has been paid to the structure of insurance markets.[5] This study is a natural complement to these two strains of the literature; I will

[3] The combination of claims data and assumptions that firms profit maximize profits allows me to separate profits from administrative loads, and, in turn, estimate fixed and sunk costs.

[4] Chiappori and Salanie [2000] propose a positive correlation test: in the presence of either adverse selection or moral hazard, insurance coverage and claims will be positively correlated. A number of studies use this framework, and some fail to find evidence of imperfect information (Cawley and Philipson [1999], Cardon and Hendel [2001]). Cutler, Finkelstein, and McGarry [2006] explain these conflicting results as the result of multidimensional consumer characteristics; a consumer with a low risk ex ante may also be extremely risk averse (the "worried well"). A more recent literature has attempted to quantify the welfare effects of selection (Einav, Finkelstein, and Cullen [2010]).

[5] Cebul et al. [2010] draw attention to the extensive use of marketing and brokers in employer-sponsored health insurance, showing that search costs create frictions in the market. Dafny et al. [2010] show that a portion of the rapid increase in claims can be attributed to changes in market structure. However, these studies have not analyzed the impact of selection and insurer behavior in an equilibrium model.

draw on a fair amount of institutional detail and show how market structure, strategic firm decisions, and consumer welfare are linked when asymmetric information affects demand for insurance. The analysis will highlight the importance of marketing costs and strategic insurer behavior in driving market outcomes. This paper is most similar in methodology to . [2011] and Bundorf, Levin, and Mahoney [2010], but is interested in a different question. I argue that the nature of consumer demand, coupled with the strategic response of insurers, has a major affect on total welfare in this market and determines to whom the rents generated by imperfect competition accrue.

The paper is organized as follows. Section 2 describes the Medicare Supplemental Insurance market in greater detail and presents reduced form evidence of selection. Section 3 outlines the model, while Section 4 describes the estimates, which confirm the reduced form results and additionally imply a small elasticity of demand driven by brand preferences. Section 5 provides an analysis of the interaction between asymmetric information and imperfect competition, and shows that, in the face of imperfect competition, adverse selection can be good for consumer welfare. Section 6 describes the role of marketing in shaping consumer welfare, and Section 7 concludes.

2 The Medigap Market

Medicare, the federally run health insurance for the elderly, has significant cost-sharing provisions. The program was largely designed to resemble private insurance available at the time of its inception and, as a result, the cost-sharing provisions left Medicare beneficiaries exposed to a large degree of risk. Part A, which covers hospitalization and is free to all those who are eligible, has a deductible of approximately one day of inpatient care, or around $1000. Part B, which covers outpatient services, requires 20 percent copayments for services with no stop-loss. As a result, a large proportion of seniors have some form of wrap-around supplemental insurance; despite the prevalence of additional coverage, out-of-pocket costs

are substantial.[6]

Medigap insurance arose to meet the demand for supplemental coverage. Legislation was enacted in 1990 to standardize the plans; additional plans have been added and prescription drug coverage eliminated since that time.[7] During my sample period from 2006 to 2008, consumers could choose only from a federally mandated menu of Medigap policies. These plans are described in Table I, which also shows the distribution of consumers across policies. With the exceptions of policies K and L, which have substantial cost-sharing arrangements, the plans are fairly similar. Plans B through J cover the Part A hospital deductible and plans B through J cover the coinsurance on Part B outpatient services. Plan F is by far the most popular, covering nearly half of the Medigap population, and plans C and J capture significant portions of the market as well.[8] Plans are both fairly tightly packed in characteristic space and horizontally, rather than vertically, differentiated.

[Insert Table I]

The only unobserved dimension on which policies may differ is their pricing (rating) method. Firms may choose to base premiums on the age at which the policy was issued (issue-age) or the current age of the consumer (attained-age) or based on a community rating. The premiums for attained-age policies may initially be cheaper but increase over time. The data do not contain information on the rating method, but firms often only use one form of rating; for example, UnitedHealth only community rates policies. To the extent this is true, brand dummies can capture the rating method; if there is variation within a brand (which is less likely), the unobserved profile of future premiums is part of the unobserved characteristic

[6] This is especially true for the near-poor. See Gross et al. [1999] for a discussion.

[7] Finkelstein [2004] discusses the reasons for and impact of the implementation of minimum standards.

[8] In addition to supplemental coverage through a previous employer or Medigap, seniors can choose a Medicare HMO or a prescription drug plan; some seniors are also eligible for additional benefits from Medicaid or the Veterans' Administration. After turning 65 and enrolling in Medicare Part B, consumers have six months to choose a Medigap policy without being subjected to medical underwriting, in which insurers can ask about and price to pre-existing conditions and current health status. During this period, insurers can only price based on age, gender, and smoking status. Evidence from Robst [2001] indicates that most consumers are not subject to medical underwriting; premiums are a function of state of residence, age, and gender. The individual level sample will be constructed to take account of this fact.

in the demand system. In the modeling section, I discuss the potentially dynamic aspect of the consumer's and firm's problem that this rating issue creates.

2.1 Data

Market level data from the National Association of Insurance Commissioners (NAICS), a regulatory administrative database, contains premiums, quantities of new policies sold, and claims at the state-insurer-policy level.[9] An observation in the market level data is a lettered contract-firm-state-year combination. Because the state-year combination is the relevant market definition, this is essentially market-product level data. For each data point, I observe the number of consumers purchasing a plan (which allows me to construct market shares), the total claims incurred (which allows me to construct average claims), and plan characteristics. Taken together, these data alone would allow me to estimate a demand system and claims function. However, I use individual level data from the Medicare Current Beneficiary Survey (MCBS) to augment the NAIC market level data. The MCBS is a survey of a nationally representative panel of Medicare beneficiaries that is linked to administrative records from Medicare. It provides information on expenditures, coverage, and a wide variety of demographic characteristics. I use several variables from this data, including demographic characteristics such as age and income, Medicare expenditures as a proxy for health, and a dummy for having any Medigap coverage (constructed as in Fang, Keane, and Silverman [2008]).

The level of observation in the demand system is a lettered plan-firm-market combination.

[9] The data are aggregated into three year periods: the 2006 data contain new policies sold from 2004 to 2006, the 2007 data contains new policies sold from 2005 to 2007, and the 2008 data contains new policies sold from 2006 to 2008. I treat each state in each of these three-year sets as a separate market. This also masks any within-state price variation. Maestras, Schroeder, and Goldman [2008] document that within-state price dispersion is dwarfed by cross-state price dispersion. Furthermore, the market shares are adjusted to account for consumers who appear in multiple cross-sections, and the standard errors are corrected using the method described by Bhattacharya [2005]. The data contain information about new policies sold; there are a few advantages to this measure. First, I observe consumers the first time they purchase a given policy, which makes it less likely that they are being subjected to medical underwriting. This limits the ability of the insurer to price to individual risk, making the examination of selection more interesting. Furthermore, because they have no history in this market, consumer inertia or switching costs will not bias my demand estimates, allowing me to abstract away from the potentially dynamic problem facing consumers and firms.

The NAIC data contain plan characteristics, such as premiums and letter, and claims. I then simulate potential consumers from the micro-level data. For each observation in the demand system, I draw simulated consumers from the micro-level MCBS data and hold the characteristics of these consumers, as well as whether or not the purchased a Medigap plan, in memory.[10] This provides me with a set of consumers to simulate characteristics without making any assumptions about the joint distribution of consumer characteristics. To construct a sample of potential Medigap enrollees, I follow Fang, Keane, and Silverman [2008], dropping Medicare HMO enrollees. In addition, because I am considering first-time enrollees, I limit the sample to those under age 72.

About 30 percent of the sample purchases Medigap insurance, which is representative of the Medicare eligible population as a whole. The sample, which I describe in Section 4, is restricted to those who are likely to be purchasing Medigap insurance for the first time, those younger than 72 with Part B charges. I observe consumers the first time they purchase a given policy. So long as they are not frequently switching between policies arbitrarily (and they have no incentive to do so), they are in an open enrollment period and not subject to medical underwriting. This limits the ability of the insurer to price to individual risk, making the examination of selection more interesting. Furthermore, because new enrollees have no history in this market, consumer inertia will not bias my demand estimates. Also, I am able to abstract away from the potentially dynamic problem facing consumers and firms.[11]

The data highlights the potential for imperfect competition in this market. First, health insurance is a highly concentrated industry, and Medigap is no exception. In 2008, the American Medical Association reported that nearly all acute care health insurance markets qualified as "highly concentrated" by the standards of the Department of Justice.[12] The

[10] The observation draws are done at the market level so that the shares sum properly.

[11] Because consumers are subject to medical underwriting after the open enrollment period, they may be locked in to the policy they purchase. Therefore, they should consider not only the current premium, but also the expected future premiums. A similar problem would arise if the inertial consumer were self-aware of their lower elasticity in future periods.

[12] See "Competition in Health Insurance: A Comprehensive Study of U.S. Markets"

national four firm concentration ratio for Medigap insurance is a staggering 82%. This stands in direct contrast to other lines, or types, of insurance; for example, commercial auto insurance, purchased by most Americans, has a four firm concentration ratio of only 21 percent, according to the NAIC.

High concentration may lead to higher premiums, but in addition, a low price elasticity driven by brand preferences incentivizes insurers to heavily market their policies and subsequently price them above cost. I measure an average own-price elasticity of -1.1; in the data, this is driven by the fact that consumers simply do not choose the cheapest plans available to them. My model explores the extent to which this can be rationalized by brand preferences (or, alternatively, consumer search) and insurer response to these consumer preferences. However, Table II shows Medigap market share data at both the national and state level. UnitedHealth is the dominant firm in the market, capturing 46 percent of the national market and nearly 50 percent of each state's market on average. Mutual of Omaha sells nearly a fourth of all policies sold nationwide; together, these two firms sell nearly three-fourths of all Medigap policies. Market share declines quickly; beyond the top ten firms, no insurer captures more than 1 percent of the market.[13] In the Medigap market, the firms are highly asymmetrical, with a few large firms that capture the majority of the market, driving the high concentration ratios.[14] This concentration has potential implications for consumer welfare. Table III shows prices, claims, and profit rates, defined as the percent markup over claims. There is both substantial variation in prices across policies and a sizable gap between prices and claims: for Plans F and J, load or margin before administrative cost, average 32 percent and 35 percent, respectively.

[Insert Table II]

[Insert Table III]

[13] Additionally, among the small and mid-sized firms, there are some recognizable companies. For example, State Farm, a firm normally associated with property and casualty, captures approximately 2 percent of the market, while Nationwide, another property and casualty company, is a very small player.

[14] Sutton [1992] suggests that the sunk costs (potentially of marketing) can lead to asymmetric firm sizes similar to those seen in this market. Ultimately, I measure fixed and sunk costs of entry and argue that only an entrant who incurs large costs can have a positive impact on the market.

Given these facts, the focus of this paper is to highlight the forces that drive market power in insurance markets. While the scope for adverse selection is relatively small in this market, I attempt to control for any correlation between consumer preferences and underlying risk in a convincing way. Despite the limited degree to which selection may affect market outcomes, the market is far from efficient due to imperfect competition. Therefore, the key contribution of this paper is to highlight the features of insurance markets that give rise to imperfect competition, where rents accrue, and to consider the scope for public policy remedies.

3 Model

3.1 Consumer Preferences

In this market, the significant cost-sharing provisions of basic Medicare create demand for insurance coverage. In my model, insurers offer a wide range of differentiated products; even within Medigap's standardized policies, consumers choose from a wide range of vertically and horizontally differentiated options.[15] Price dispersion exists even within individual contracts, indicating that consumers either have brand preferences or incur search costs or both. I model consumers as having preferences over individual brands; therefore, a consumer could obtain higher utility from an AARP-sponsored policy A than a Mutual of Omaha branded policy A, even if the latter has lower premiums.[16]

As a result, following the discrete choice literature, contracts are defined as a bundle of characteristics: the lettered policy x_{jm} and brand b_{jm} such that $\phi_j = g(x_{jm}, b_{jm})$. I allow consumers to vary along two dimensions, a proxy for health status (which may be

[15] Additional details on how such a model can arise from underlying primitives is available in Appendix A.

[16] While the plans are identical and claims are largely automated, some firms, such as UnitedHealth, offer services such as a nurse call line. The standardization has limited the ability of insurers to offer differentiated contracts, and this leads companies to compete based on level of service and company reputation.

Alexcih et al. [1997] note that the regulations allow individual states to approve "innovative benefits"; however, in practice, companies choose to compete on brand reputation rather than product differentiation.

continuous or discrete) z_i and income ω_i, but no unobserved characteristics other than the error term. In the specifications I present in the text, I use a measure of uncovered Medicare expenditure (generated from a regression run on the entire population, not just the pool of Medicare beneficiaries) as a proxy for risk. However, the results are robust to various other (potentially less contaminated) measures: in Appendix A, I present results using self-reported health. In both cases, I tie together the individual and market level data by drawing consumers from the MCBS and using micro-moments to help me pin down random coefficients.

Consumer preferences over health and income drive the value they place on different insurance contracts. Income is likely to affect consumers' willingness to pay for insurance through price sensitivity. Therefore, the model allows α to depend on income. Furthermore, if expenditure depends on severity of illness, the trade-off between facing a sure premium and an uncertain expenditure varies by health status. Therefore, the model will allow α to depend on a proxy for health status that captures expected expenditure. Finally, health status also affects the probability of needing medical care; therefore, the utility of holding any insurance policy will depend on health status. The trade-off that governs both interactions is the same: those with higher expected expenditure prefer a product that reduces their own financial exposure or are willing to trade off more certain premiums for reduced exposure to uncertain payments. In addition, choosing these particular interactions is not crucial to the analysis; Appendix A considers numerous robustness checks.[17]

Formally, a set of products J, defined by the contract letters above, are offered by a set of firms F in a set of markets M, where a market is defined as a state-year. The consumer chooses one such policy (or the outside option) to maximize the indirect utility function.

[17] The only specification that isn't deeply explored in the Appendix is one in which consumers vary in an unobserved consumer characteristic. The estimation relies on micro-moments, which allows me to pin down the random coefficients (and which are not particularly helpful in estimating the distribution of an unobserved consumer characteristic). Initial attempts to estimate such a model indicate that allowing for such heterogeneity did not substantively change the results.

Consumer i's valuation from product j in market m is given by:

$$\begin{aligned} v_{ijm}\left(x_{jm}, \xi_{jt}, p_{jm}, \omega_i, z_i; \theta^d\right) &= x_j'\beta_1 + b_f'\beta_2 + \xi_{jm} + \alpha p_{jm} + \mu_{ijm} + \varepsilon_{ijm} \\ \mu_{ijm} &= \alpha^{\omega} p_{jm}\omega_i^k + \alpha^z p_{jm} z_i^k + \beta^k z_i^k. \end{aligned}$$

The observed product characteristics are a vector of dummy variables for both lettered plan x_j and insurer brand b_j; the unobserved product characteristic is denoted by ξ_{jm}. The estimates of α^{ω} and α^z capture the impact of consumer characteristics on price sensitivity. Finally, β^k allows consumers to vary in their preference for holding *any* Medigap policy. The error in the utility function is an independent and identically distributed type-I extreme value ε_{ijm}.

However, price is likely to be correlated with the unobserved characteristic; if brokers differentially "push" certain policies within a company's line of products or insurers vary their rating methods across policies, these characteristics will both be correlated with consumer demand and price. I use two sets of price instruments to correct for the endogeneity problem. First, retaliatory tax regimes, in which states can levy taxes on premiums differently depending on an insurer's state of incorporation, are used. In the demand specifications, in which brand preferences are important, I also include a second set of instruments: the average premiums of policies in other markets.[18] The unobserved characteristic contains anything about the insurance product that is not captured by plan and brand dummies. Therefore, the exclusion restriction could be violated if, for example, firms use different rating methods in different markets based on local demand. In this particular case, large firms in the market, such as UnitedHealth, community rate all of their policies. Therefore, in the presence of the brand dummies (Nevo 2001), a significant amount of potential for price endogeneity has been controlled for in my specifications.

[18] Hausman [1999]. I include these in the demand specifications, and not the reduced-form results because brand fixed effects are critical in the demand system and unimportant to the relationship between premiums and claims. The tax regime instruments do not have enough variation alone.

Finally, it is worth taking note of the options that compose the outside good in this market. Some consumers may have access to supplemental insurance through a former employer, while others may qualify for Medicaid, the health insurance program for the poor, or benefits through the Veteran's Administration. All consumers are eligible to purchase any Medicare Advantage ("Medicare HMO") plan available to them. Ultimately, Medigap plans appeal most to relatively high income consumers without access to insurance through other sources; previous research has indicated that these consumers are healthier on average than consumers who do not purchase Medigap (Fang, Keane, and Silverman 2008). However, for the purposes of the model, the utility of all outside goods is normalized to zero. The results are robust to different market size definitions; this would mainly be problematic in my specification if the value of the outside good varied dramatically across markets. To potentially alleviate some of these concerns, nested logit results (in which all inside goods are in one nest and the outside good is in the other) are presented in the Appendix A. The results are quantitatively similar to the preferred specification.

3.2 Claims Function

In the demand estimation literature, incorporating consumer heterogeneity is key to avoiding undesirable features of logit demand specifications, such as the independence of irrelevant alternatives. In insurance markets, consumer heterogeneity is also a key feature of the asymmetric information problem: firms care about the types of consumers they attract because different types of consumers will incur higher or lower claims. Therefore, both the elasticity of demand and the degree of selection, measured by the shape of the claims function, will affect equilibrium prices. The joint estimation of claims and demand is also the key modification in the model that accounts for the possibility of asymmetric information.

In this market, there is limited scope for selection along a vertical characteristic, such as plan generosity. The policies are horizontally, rather than vertically differentiated, and they are, additionally, tightly packed in product space. However, consumer preferences

may still be correlated with claims in a way that affects insurer decisions. There is a striking relationship between premiums and claims in the data, even for identical policies in identical markets. That is, if Company A offers Policy F at a price of $1500 and Company B offers a *contractually identical* Policy C at a price of $1200 in the *same* market, the latter will experience lower claims. The relationship holds both conditional on policy, state, and brand fixed effects and when using plausibly exogenous variation in premiums.

This relationship between premiums charged could arise from a variety of models of consumer behavior. First, a consumer's willingness to pay could be correlated with health status, giving rise to a form of adverse selection. This form of selection does not suggest that preferences over plan generosity are correlated with underlying risk, as in theoretical models such as Rothschild-Stiglitz, but that willingness to pay is correlated with underlying risk. This particular form of selection is the form measured in Einav, Finkelstein, and Cullen [2010] in their empirical application. However, differential search could account for this pattern as well. Consumers may differentially search in a way that is correlated with their underlying risk. For example, healthy consumers may have lower search costs, perhaps due to increased cognitive function, as pointed out by Fang, Keane, and SIlverman [2008], and search more intensively, choosing the lowest priced plan from their "consideration" set. Finally, healthier consumers with higher cognitive function may have a better understanding of products and may be more apt to realize that all same-lettered plans are contractually identical. This certainty or lack of bias may lead them to choose lower-priced plans or may make them less likely to see price as a signal of underlying quality.

This is a very specific form of adverse selection, but due to the standardization of policies (and their similarity), there is limited scope for significant adverse selection across horizontally (not vertically) differentiated policies. However, this correlation between consumer preferences for price and underlying risk is likely to affect prices. The claims function relates claims to plan and consumer characteristics. Let γ_{ijm} be the claim of individual i

under policy j in market m and write:

$$\gamma_{ijm} = \theta_0 + x_j'\theta_1 + \omega_i'\theta_2 + z_i'\theta_3 + g_m'\theta_4 + \vartheta_{ijm},$$

where the error ϑ_{ijm} is a function of both an individual level and plan level error $\vartheta_{ijm} = \varepsilon_{jm} + \eta_i$. Therefore, while claims may depend on the observed characteristics in a plan and the plan's characteristics, the unexplained part of an individual's claim does not depend on the plan chosen, which implies that $E_i\left(\eta_i | J = j\right) = 0$; this specification assumes no moral hazard that is plan-individual specific. However, claims are not observed at the individual level, and I must take expectations over individuals in a given plan. To minimize the extent to which outliers drive my results I model claims in logs, rather than levels. The log plan level claims γ_{jm} are given by:

$$\log\left(\gamma_{jm}\right) = \log\left(E_i\left[\gamma_{ijm} | J = j\right]\right) = \theta_0 + x_j'\theta_1 + E_i\left[\omega_i | J = j\right]'\theta_2 + E_i\left[z_i | J = j\right]'\theta_3 + g_m'\theta_2 + \varepsilon_{jm},$$

and this equation can be estimated using the firm level claims data and information from the demand system. The derivative of claims with respect to price, critically important for the insurer pricing decision, are given by a combination of the random coefficients from the demand system and the parameters of this equation.

3.3 Estimation and Discussion

To estimate the random coefficients, I rely on both variation across markets and micro-moments.[19] For the consumers drawn from the MCBS data, I construct the probability of purchasing *any* Medigap plan, as well as the premiums paid *conditional on purchase*. I construct the probability of purchase given consumer characteristics in the model and

[19] Furthermore, Petrin's [2002] example, building on the intuition of Imbens and Lancaster [1994] highlights the value of micro-moments in pinning down the relationship between consumer demographics and utility. This basic framework can be extended to insurance markets; the key difference from more familiar settings is that in addition to estimating the demand system in a sufficiently flexible way, the researcher must estimate a claims (a component, but not the only component, of marginal costs) function as well.

compare it to the expectation of purchase given consumer characteristics in the micro-level MCBS data. Furthermore, I compare premiums paid in the MCBS data and the model. The estimation strategy assumes these are equal in expectation.[20]

In addition, I made a number of modeling decisions. First, the demand system and insurer's problem are both static simplifications of a dynamic problem. Firms must compete for this new pool every year. This both simplifies the problem and reflects that potential that consumers - already confused by a wide array of options - are not completely forward looking. In addition to modeling the problem as a static consumer decision, I choose to model consumers as having brand preferences when choosing policies; I argue that marketing is a central feature of this market. Marketing insurance is a difficult and expensive task; search costs may be high and employing insurance brokers may be costly (Cebul et al. [2010]). Furthermore, this problem is likely to be exacerbated with an elderly population that may not be technologically savvy enough to search effectively on the Internet and may be experiencing cognitive declines.[21] However, search costs are not the only feature of this market that leads to the existing market structure; consumer anxiety and preferences for a trusted name are important even for contractually identical policies.[22] Incorporating brand effects allows me to rationalize a significant amount of price heterogeneity for seemingly homogeneous products. However, a different, but equally reasonable choice would be to model consumers as having search costs: obtaining additional price quotes is costly, and

[20]The model predicts the probability that individual i purchases a Medigap policy is $1 - s_{i0m}$ and that the premium paid by any individual is $\Sigma_{j \in m} p_{ijm} s_{ijm}$. The GMM routine chooses the random coefficients to help match this information from the individual level data along with moments from the demand and claims functions; thus, while the mean price coefficient is identified by the assumption of the exogeneity of the price coefficients and variation in choice sets (and, by extension, prices) across markets, micro-moments help identify the random coefficients, along with geographic variation in demographics across markets. The relationship between consumer characteristics and claims is unbiased if the average consumer characteristics are uncorrelated with the error term; the specification includes both policy and state fixed effects, denoted by x_j and g_m, respectively.

[21]See Fang, Keane, and Silverman [2008] for a discussion of cognitive ability and insurance purchasing behavior.

[22]Alexcih et al. indicate that insurers attempt to differentiate themselves along this dimension. Simple logit models of demand show that elasticities are not statistically different from one another if brand effects are included or not. However, significantly more of the variation in δ_{jm} is explain when brand dummies are included.

they therefore choose from a smaller choice set, leading them to purchase policies that are not the cheapest available in a given geographic region. Search costs have been modelled in Hortascu and Syverson [2004], Cebul et al [2010], and (in this market) Maestras, Goldman, and Schroeder [2008]..

Incorporating both search costs and brand preferences in a model of consumer choice is challenging: both will give rise to similar patterns in aggregate data.[23] While both models give rise to similar incentives for insurers, the two models have different implications in counterfactual exercises; when discussing the impact of eliminating brands, I am much more agnostic about changes in consumer surplus. Ultimately, I see brand preferences and search costs as complementary explanations of the patterns seen in the data, as consumers may want to purchase a plan from a trusted name because the information is limited in a way that prevents the consumer from making a fully-informed decision. This idea is further discussed in Section 5.

A third central modeling choice is the form selection takes in the demand system. Here, selection involves consumer sorting by willingness to pay for insurance. Specifically, I find that sicker consumers sort into higher price policies, leading to adverse selection along this dimension. This is different from selection along the extensive margin into Medigap coverage, which is believed to be advantageous. This modelling choice is both data driven and substantive. The information in the individual level data is best suited to identifying heterogeneity in preferences for price because it contains information on how much each consumer spent on insurance, but not the firm from which it was purchased.[24]

Finally, this approach to the problem of unpriced heterogeneity in consumer preferences (giving rise to selection), is not new to the literature. Bundorf, Levin, and Mahoney [2010]

[23] In order to separately identify search costs from brand preferences, there must be a variable that shifts search costs without changing brand preferences. It is difficult, though not impossible to think of such a variable. Furthermore, it is difficult to model the relationship between search costs and selection in the data. Therefore, in this application, I will model consumers as having brand preferences, rather than incurring search costs. This assumption, while central in determining consumer choice, does not drive the estimates of structural parameters.

[24] By contrast, one could choose to model heterogeneity in brand preferences; robustness checks are provided in the Appendix A.

consider selection and the impact of risk rating using a model of consumer demand, and find that more sophisticated, risk based pricing would do little to improve consumer surplus. Lustig [2009] presents a very similar model of consumer choice in which he uses assumptions about firm pricing to identify (unobserved) costs. Both studies implicitly argue (along with an increasingly voluminous literature, including Einav, Finkelstein, and Cullen [2010]) that selection is unlikely to cause decreases in consumer surplus of any significant magnitude.

I expand the work of Lustig [2009] and Bundorf, Levin, and Mahoney [2010] that attempts to estimate selection by providing a mechanism through which market failure occurs in insurance markets, even in the absence of selection. By examining the Medigap market, we can understand the types of factors that make the market for health insurance imperfectly competitive, and the implications of these factors for consumer welfare. Specifically, I show that the demand system above gives insurers significant incentives to market their prices heavily and price above cost. Because I have data on claims, I can separately identify marketing and administrative costs, such as broker commissions, which I find are a substantial cost to firms (and largely ignored in the previous literature on selection and insurance demand). The strength of my paper, relative to previous work, is that I combine market-level cost data with a model of insurer pricing. This allows me to estimate administrative and marketing costs in the form of broker commissions. The latter are absolutely critical to understanding the market and the importance of imperfect competition as a friction in insurance markets.

3.4 Insurer Profits and Variable Costs

In addition to claims, firms face an administrative cost of marketing and administering an insurance policy. Facing the task of marketing their policies, insurers may choose among a number of potential channels. They can sell their policies directly through independent brokers, who sell a variety of insurance and financial products of a variety of different firms, or proprietary agents. The second largest firm in the Medigap market, Mutual of Omaha,

pays 16 percent of premiums of Medigap policies to agents in the form of commissions.[25] The largest firm, UnitedHealth, is both a large player in other lines of health insurance and has purchased the brand of the AARP, which according to Bloomberg collected nearly five hundred million dollars in royalties and fees in 2007 and $222 million from UnitedHealth alone in 2008. The AARP brand name allows it to capture a large proportion of the Medigap market by selling their policies directly to consumers.[26] This section incorporates marking expenditures into insurer decision making.

Formally, firms set prices p_{jfm} for all J_{fm} policies to maximize variable profits, cognizant of the fact that the vector of market prices $\mathbf{p}_m$ affects not only shares but also claims and potentially administrative costs (if the administrative cost is a percentage commission, for example). Let the econometrician's approximation to firm $f's$ variable profits from policy j in market m with market size $\mathbf{M}$ be equal to

$$\pi_{jfm} = \left(p_{jfm} - \gamma_{jfm}\left(\mathbf{p}_m\right) - a_{jfm}\left(\mathbf{p}_m\right)\right) s_{jfm}\left(\mathbf{p}_m\right)\mathbf{M},$$

where a_{jfm} denotes the administrative cost of policy j offered by firm f in market m. To calculate firm level profits, π_{jfm} is summed over the policies the firm offers, subtracting any fixed or sunk costs of operation, which are discussed in subsection 4.3.

While I assume a Nash in prices equilibrium in this market, two important institutional features play an important role in the estimation. The first, the potential for dynamic commitment, echoes concerns about dynamic consumer behavior on the demand side. Consumers are not simply choosing a policy today, but are locked into rate increases in the future, as they have little incentive to switch policies. Furthermore, firms can rate policies

[25] Brand effects are critical in explaining consumer behavior on the demand side, and are drivers of consumer behavior. The commission rate is taken from a Mutual of Omaha Agent Guide.

[26] The strategies of the firms with a small market presence are to be a niche players, capturing consumers who happen to prefer their brand. Often these firms, such as Aetna, specialize in other lines of insurance, but will engage in direct sales of Medigap on their website or over the phone. As a result, they do not sell many policies, but also do not need to pay commissions on the policies they sell. They can choose to incur low costs, both variable and fixed, knowing that they will sell very few policies. The next section formalizes the game played by insurers to calculate costs that correspond to these marketing strategies.

in a number of ways: community rating, age-attained, or issue-age, which have different profiles of premiums. A firm may choose to set a low price initially and then hold-up consumers ex post in an exploitation phase. Finally, firms' cost structure may vary over time if commission rates decrease from year to year.

This is a real concern; however, for two reasons, I feel the static representation is a fair approximation. First, as pointed out in the Mutual of Omaha literature, commissions are not only for the first year of a policy and do not depend on the age of the enrollee. Insurers want to limit the ability of brokers to churn consumers from one policy to another from year-to-year. Therefore, while commissions may be slightly higher than average during the first year of a policy, they do not fall off dramatically. In addition, the prices I see in the data are certainly not at or below cost in the hopes of exploiting locked-in consumers in the future. For these reasons, I feel that the commission estimates I obtain are a reasonable approximation of the firm's true underlying cost structure.

The second is minimum loss ratio regulation, which limits the amount of information I can infer from firm pricing. If a policy violates the requirement that 65 percent of premiums be paid out in claims, the insurance regulator could ultimately force the firm to rebate the difference to consumers. In practice, the insurer often contests the potential fine and the rates are allowed to stand (a fact which is obvious in the data). Nearly half of all policies violate minimum loss ratio regulation. This information must be incorporated into estimation: ignoring this fact will lead to misleading estimates. By contrast, simply throwing out the observations for which the minimum loss ratio regulation is binding will leave a selected sample and lead to misleading results as well. Therefore, I use all of the data in estimation, but I do not take a stand on about the exact impact of regulation on the profit function; the model simply assumes that firms cannot increase profits by lowering the price of policies which currently violate the minimum loss ratio regulation.

To order to use all of the data in my estimation procedure, I relax the standard assumption that firms set their first-order condition equal to zero. Instead, I assume that they

maximize profits under the shadow of regulation. I implement this using a slightly less stringent assumption. For those firms for which regulation does not bind, the standard profit maximizing assumption is used. However, if, in the data, the regulation should bend, I take a slight more nuanced approach by assuming that the firm could become more profitable by decreasing its price. This allows me to incorporate the regulation in a way that uses all the data without being rejected by the data.

Therefore, I take the pricing regulation as given and make assumptions that incorporate this institutional feature. If the pricing regulation is not binding for a given policy, I assume that the firm is setting the price of that policy to maximize profit, holding the prices of competitors fixed, and taking into account the fact that the firm may sell a variety of policies within that market. If, by contrast, the pricing regulation is not binding for a given policy, I only assume that the firm cannot profitably deviate by *lowering* its price. The firm may, however, want to raise its price but be unable to do so because of regulation. By making these minimal assumptions, I can at the very least partially identify administrative costs, which play a large role in shaping this market.

The estimation is based on Kuhn-Tucker conditions for profit maximization, where the firm's problem is written as:

$$\max_{p_{jfm}} \Pi_{fm}$$

$$s.t.\ \gamma_{jfm} \geq .65 p_{jfm}.$$

Due to the assumption that regulation need not bind exactly, the complementary slackness condition is relaxed such that $\mu_{jm}\left(.65p_{jfm} - \gamma_{jfm}\right) \leq 0$ while $\mu_{jm} \geq 0$. In this case, if the minimum loss ratio regulation is violated, it must be that $\frac{\partial}{\partial p_{jfm}}\left(\Pi_{fm}\right) \geq 0$, while if the loss ratio regulation is not violated, $\mu_{jm} = 0$ and $\frac{\partial}{\partial p_{jfm}}\left(\Pi_{fm}\right) = 0$.[27]

In words, if minimum loss ratio regulation binds, the estimation routine allows the derivative to be positive, as firms may prefer to raise their price in the absence of regulation. If

[27] The results of a counterfactual in which loss ratio regulation is lifted are in Appendix B. In the absence of regulation, premiums would rise by 14 percent; the regulation does manage to restrain premiums.

all policies were subject to regulation, this would be equivalent to minimizing the negative of the function - estimation based on moment inequalities. If all policies were not subject to regulation, this would be equivalent to the more standard method of moments, the estimation routine would simply pick the variable administrative costs that minimized the expectation of the moments.[28]

The marginal administrative cost is the cost of adding one consumer conditional on the policies offered by an active firm in a market. What composes variable administrative costs? These costs must be truly marginal, in the sense that they are incurred for each additional consumer. Claims processing is automated and done in conjunction with traditional Medicare billing. Therefore, the marginal costs that are not associated with claims must largely be marketing costs, the costs of generating an additional policyholder. The main component of variable administrative cost is sales commissions, which I is a percentage of the price charged; that is, $a_{jfm} = \beta p_{jfm}$. Throughout, I will use the terms marketing and administrative costs interchangeably, as marketing costs are the largest component of administrative costs, and the two cannot be separately identified.

4 Estimates

4.1 Reduced Form Evidence of Selection

Table IV considers the relationship between premiums and claims. Controlling for plan fixed effects, the estimates in the first column indicate that a \$100 increase in premiums is associated with an \$80 increase in claims. This is for policies that are contractually identical; the results cannot be driven by unobserved plan differences, or moral hazard, though reverse causality and geographic and temporal differences may drive the results. Therefore, the

[28] Additionally, I could potentially improve the efficiency of the estimating by jointly estimating administrative costs and the demand and claim moments. However, for computational simplicity and because I will need to subsample to obtain the standard errors of the variable cost estimates, I will do the estimation separately, taking the demand estimates as given. The current standard errors were obtained using a bootstrap.

second column additionally controls for state by year fixed effects. The results are similar when including these fixed effects; these results imply that more expensive policies incur higher claims, even conditional on all contractual variation and the market (state by year) in which the policy is offered. However, in any basic pricing model, prices and costs are likely to be tightly linked, as prices will be equal to claims in a perfectly competitive market, and equal to claims plus some markup in an imperfectly competitive one.

[Insert Table IV]

To estimate a causal relationship between premiums and claims, I use plausibly exogenous variation in premiums due to retaliatory taxes. Retaliatory taxes work as follows: suppose insurance company A is incorporated in Illinois, with a 1 percent premium tax, while company B is incorporated in Indiana, with a 4 percent premium tax. Illinois charges company B a retaliatory tax of 4 percent on the premiums it collects in Illinois. Therefore, different companies in the same state face different tax rates.[29] The instrument is valid if the state of incorporation is uncorrelated with the error term. If firms were incorporated before the Medigap insurance market existed and did not change their state of incorporation in response to their activity in the Medigap market, which seems like a reasonable assumption, the exclusion restriction is satisfied. I instrument for prices using these tax regimes: they imply differences in tax rates on average of two-thirds of a percent on average, ranging up to nearly 3 percent of premiums; the first stage shows that they are passed on directly to the consumer, potentially having a reasonable (roughly \$10-\$45) impact on premiums. These taxes are exempt from rules governing interstate commerce (via a 1981 Western and Southern court decision), and are levied by all states. Their impact varies from a zero percent additional tax burden to a 4.26 percent additional tax burden, and applies to all lines of insurance. These instruments, taken together, have an F-stat of over fifty, indicating that

[29] This has two distinct effects. The first is an entry effect: a company only enters a market in which it will face higher than average taxes if it is a low cost provider who can compete in terms of price. The second effect is a pure price effect, where some of the additional tax is passed on to consumers in the form of higher premiums; the functional form used takes account of both effects.

they have some power in explaining premiums.

In the empirical specification in Column 3, I instrument for price with both a dummy variable that takes on a value of one if the plan faces retaliatory taxes and the difference in taxes between the plan and the (lower) base rate in that state. The coefficients on both variables are statistically significant and in the expected direction. The relationship between premiums and claims falls by about half when instrumenting for price using variation in tax regimes; these results indicate that a $100 increase in premiums attracts a group of consumers that incur approximately $40 more in claims on average. This provides some evidence that consumer preferences may be correlated with underlying risk (perhaps because higher risk individuals are less price sensitive), which could lead to distortions in pricing.

4.2 Structural Parameters

Summary statistics describing consumers are found in Table V and VI.[30] I form an expectation of insured expenditure that conditions on having insurance, as well as income, demographic characteristics, and self-reported health status, and extrapolate to the whole sample. The regression used is detailed in Table VII; the expectation depends heavily on the consumer's own perception of their health as well as demographic characteristics, including income, which may affect utilization of health care services. This measure, which is assumed to proxy for Medigap claims, averages $1538, which is slightly higher than the average claim in the market-level data, which may reflect overlap between Medigap and other private insurance plans.

[30] The income variable indicates an average income of $36,804, with a substantial amount of variation. The mean of private health insurance expenditure is low because it is equal to zero for those consumers without private health insurance, including Medigap. Therefore, I form an expectation of insured expenditure that conditions on having insurance, as well as income, demographic characteristics, and self-reported health status, and extrapolate to the whole sample. The regression used is detailed in Table VIII; the expectation depends heavily on the consumer's own perception of their health as well as demographic characteristics, including income, which may affect utilization of health care services. This measure, which is assumed to proxy for Medigap claims, averages $1537.99, which is slightly higher than the average claim in the market-level data, which may reflect overlap between Medigap and other private insurance plans. Appendix A contains alternative demand specifications and robustness checks.

[Insert Table V]

[Insert Table VI]

[Insert Table VII]

Table VIII shows the demand system estimates. The average elasticity is -1.1714, indicating that a $100 increase in price results in an 7.5 percent reduction in market share. Given that the coverage of lettered contracts is identical and they are likely to be close substitutes, I interpret this as a relatively small price sensitivity. The main effects show that consumers dislike policies K and L, which provide more limited coverage and, unsuprisingly, the popular plans yield high utility for consumers. The elasticity is similar across specifications that allow for different types of consumer heterogeneity: in column 1, consumers are only allowed to vary by income, and in the second specification, the expected individual claim is only allowed to affect preference for insurance, not price sensitivity. The first specification shows that high income consumers are especially price insensitive: a 1 percent increase in income leads to a 0.35 percent decrease in this elasticity.[31]

[Insert Table VIII]

The bottom panel shows the effect of this sorting of consumers on claims. In the first specification, the regression of logged claims on consumer and plan characteristics, which is estimated jointly, shows that a 1 percent increase in income leads to a 0.15 percent increase in expected claim, conditional on health. However, this specification does not capture the sorting of consumers on price by health status. In the preferred specification, the coefficients imply that a $100 increase in expected claim leads to a 14 percent increase in claims. Given

[31] The second specification allows consumers with different expected individual claims to value insurance differently in addition to allowing consumers with different incomes to vary in their price sensitivity. Conditional on income-varying price sensitivity, consumers with higher expected individual claims gain higher utility for holding insurance. In the final, preferred specification, the random coefficient on the interaction between price and income is 0.0002, which is significantly different from zero; again, this indicates that richer consumers are less price sensitive that their poorer counterparts. However, sicker consumers are also significantly less price sensitive, even conditional on income, while the preference for holding insurance does not vary with health status. While allowing consumer preferences to vary with expected individual claim, whether this characteristic affects preference for insurance or price sensitivity is less important.

that the median claim across policies is $944, this estimate implies that, at the median, a $100 increase in the proxy for health status given by average expected private insurance claims (from the individual level data) leads to a $131 increase in average Medigap claim (in the market level data). Furthermore, a $1000 increase in income leads to a .05 percent increase in claims, conditional on health status.[32] Finally, though Table VIII does not include the entire vector of brand effects,[33] Table 5 in Appendix A shows that the value of the UnitedHealth brand in particular is quite high - nearly $1000 per policy. By contrast, the value of Mutual of Omaha's brand is much lower, and their large market share is partially achieved by lower prices, as seen in Table 3.

The demand estimates imply a number of things. First, demand for Medigap insurance is very insensitive to price because consumers are very concerned with the brand name of their insurance. In addition, because consumers have heterogeneous preferences over price, and these preferences are correlated with claims, the price insurers charge is related to their claims. Taken together, these effects imply that a $100 increase in premium leads to a $15 increase in average claims and a 7 percent reduction in market share. The selection effect is due almost exclusively to correlation of elasticities with claims, rather than the fact that claims are correlated with preference for holding insurance. The supply-side estimates will first explore the extent to which medical loss ratios may not reflect insurer profit and the extent to which marketing may shape this market.

Before turning to the estimates of variable costs, it is useful to describe the order of magnitude one might expect. Given that minimum loss ratio regulation theoretically bounds margins net of administrative costs from above at 35 percent of premiums, it would be surprising if estimates were larger than this. Commissions are estimated at 15-20 percent in this industry[34], and this is a useful benchmark for firms who sell Medigap insurance

[32] This captures that the expected private health insurance expenditure in the MCBS tends to be slightly higher than the average claim in the NAIC data; the private health insurance expenditure variable includes other forms of private insurance, including employer-sponsored plans.

[33] There is no constant in the model; therefore, the brand intercepts are for Policy A, and have no immediately clear interpretation.

[34] For an in-depth discussion of the role of brokers, see Hall [2000] or Litow [2006].

using agents or brokers. Table IX contains estimates of marginal administrative costs. The first specification estimates a mean variable administrative cost for all firms, ignoring their different marketing strategies, and finds that the average administrative cost is equal to 8% of premiums. Markups average 29 percent of premiums in this market, and, as a proportion of markups, marketing costs appear substantial, averaging 27 percent of potential profits. However, this estimate masks the heterogeneity of marketing costs.[35] However, UnitedHealth is estimated to have substantially lower administrative costs, which is intuitive, because the firm engages in a substantial amount of direct sales.

[Insert Table IX]

The preferred, third specification replicates the second but allow for more heterogeneity across firms, using a full set of group dummy variables to capture firm-specific administrative costs (although only the top two firms are reported in the table to save space). Again, the estimates for UnitedHealth are small; they pay only 6 percent of premiums collected in administrative costs. By contrast, Mutual of Omaha, who employs a fleet of agents to sell their policies, is estimated to pay nearly 18 percent of premiums out in administrative costs. The stated commissions paid to proprietary Mutual of Omaha agents match the estimate of 16 percent, coupled with a relatively small cost of paying claims.

To understand what mechanically drives the estimates, it is useful to compare United-Health with Mutual of Omaha. These firms have similar prices and claims. Furthermore, they have similar price elasticities and derivatives of claims with respect to price. Yet they have very different estimates of administrative costs. UnitedHealth offers a broader portfolio of plans. Therefore, if it raises its price on any individual plan, some consumers will switch to other plans in its portfolio, on which it still earns a profit. In addition, because Mutual of Omaha faces loss ratio regulation more often (44 percent of policies versus 28 percent of

[35]The second specification allows the estimates to vary across the two dominant firms. The constant represents the mean for omitted firms, and is now estimated to be approximately 12 percent, which is a little less than expected agent or broker commissions of 15-20 percent.

policies), their true markup must be less than the optimal calculated markup. As a result, the estimation forces a higher cost of generating policies for Mutual of Omaha.

4.3 Fixed Costs and Sunk Costs

There are a large number of potential entrants in the Medigap market, and variable profits indicate the possibility of a profitable challenge to the current market structure. An additional large firm may be able to both be profitable and improve consumer surplus. In this section, I examine the extent to which national and state level sunk costs may potentially exist in this market. First, the previous section illustrates that, for some firms, broker costs are quite substantial. These costs imply that medical loss ratio does not necessarily translate into high insurer profits. Specifically, in the case of Mutual of Omaha, the marginal cost of marketing and administering policies, largely driven by broker costs, reduce potential margins by nearly two-thirds. In the presence of such high commissions, Mutual of Omaha is not wildly profitable per se.

However, while Mutual of Omaha ultimately does not have extremely high margins, UnitedHealth's margins are largely unaffected by marginal costs of marketing and administration (they remain around 18 percent). Their lack of dependence on brokers is due to the access to branding thought the AARP, which allows them to sell a large number of policies direct to consumers. This branding is not free: United paid AARP's for-profit arm $222 million dollars in 2008. Still, a central question remains: would it be profitable for another firm to enter this market? If not, why not? Do the marketing costs in this market not only differentiate otherwise homogeneous plans, but create barriers to entry as well?

I examine this by simply asking how a new entrant would potentially fare in this market. If two firms commanded the brand power of UnitedHealth, and competed on equal ground, would the firms be profitable enough to cover the sunk cost of a national marketing campaign? While additional details on this exercise are available in Appendix B, I outline them briefly here. I create a firm that carries the same brand fixed effect as UnitedHealth, such that

the only differentiation between the new firm and United comes from the logit error. I then allow the firms to reprice their entire menu of policies, taking into account both the increased competition and the effects of selection, and calculate potential profits of the entrant. This exercise can provide suggestive evidence that marketing creates a barrier to effective entry in this market.

The results indicate that an entrant could not cover the national sunk cost of marketing required to effectively enter this market and challenge United's dominance, or alternatively, that AARP captures the surplus it creates through its trusted name. This is perhaps the most surprising feature of this study: while premiums are somewhat surprisingly high in this market, the rents do not accrue to insurance companies.[36] These numbers illustrate several things. First, marketing costs create and effective barrier to entry; therefore, the market structure is likely to remain extremely concentrated. In addition, the rents in this market are not captured by insurers. Instead, the money flows to parties in the supply chain with the scarce inputs - the ability to sell policies - such as brokers and, especially, AARP.

5 Welfare

Insurance markets are highly regulated, often with the market failure of adverse selection in mind. Many regulations, including the medical loss ratios in this market, constrain insurers while taking demand as given. This section will illustrate the potential folly of this approach. Given the strong consumer brand preferences and low own-price elasticities, policies that attempt to adjust insurer incentives on the margin are ineffective. The ultimate problematic feature of this market still is, in some sense, also an information problem: consumers are either unaware of every policy or potentially incorrectly infer quality, giving firms an incentive to heavily market their policies and subsequently price them above cost. In this section, I consider a number of counterfactuals. The first two speak to changing insurer incentives.

[36] The profits of the counterfactual firm are around one hundred million dollars, less than half of the transfer from United to AARP. Furthermore, if the brand is "degraded" or becomes less valuable by even a small amount, the profits fall dramatically.

However, only the final counterfactual, which alters the nature of demand, has a significant effect on consumer welfare.

5.1 Insurer Incentives

5.1.1 Perfect Risk Adjustment

I first argue that, in contrast to the welfare gains from perfect risk adjustment in perfectly competitive markets presented by Einav, Finkelstein, and Cullen (2010), imperfect competition can lead to an ambiguous relationship between selection (specifically of the kind estimated here) and consumer welfare. While adverse selection is welfare reducing in competitive insurance markets,[37] this need not be true if insurers have market power. Consider a monopolist selling a single policy with demand $D(p)$ and first consider the case in which the average claim γ, here the only component of insurer costs, is temporarily assumed to be independent of the premium charged. The standard analysis shows that the percent markup is equal to the inverse of elasticity. However, if average claim γ is a function of the premium charged, there is an extra term: the markup is equal to the inverse of the elasticity plus an additional term that depends on the amount of selection the policy faces: $\frac{1-\frac{\partial \gamma}{\partial p}}{\varepsilon}$, so long as $\frac{\partial \gamma}{\partial p} < 1$ (implying positive markups). Similarly, in a market in which imperfectly competitive insurers play a Nash-in-prices pricing game, *conditional on the prices of all the other plans in the market*, markups in the case of no selection will be higher than markups in the case of imperfect risk adjustment, because firms have an additional incentive (besides loss of market share) not to raise prices (the possibility of getting sicker consumers).

In this exercise, the derivative of claims with respect to premiums is set to zero for all policies and optimal prices are recomputed. Given this change in firm pricing, I calculate the change in the average premium paid, consumer surplus, and lives covered. If insurers maximize profits without taking selection into account, optimal prices are higher. Table X

[37] If insurers price at average, rather than marginal, costs there will exist a set of consumers for whom the price of insurance exceeds their willingness to pay, but not their underlying risk or claim. For an extensive discussion, see Einav and Finkelstein [2010].

shows the median enrollment-weighted premium across market rises to $1407, a 9 percent increase as compared to the current equilibrium.[38] In turn, this increase in prices shrinks the size of the market by 18 percent. The expansion of coverage under adverse selection is especially striking; under the assumption of perfect competition (zero-profit condition), adverse selection leads to the underprovision of insurance. Furthermore, consumers would give up $258 to face prices *conditional on adverse selection.* This is not just a transfer from producers to consumers: the absence of adverse selection actually reduces total surplus by approximately 7%. This shows that asymmetric information can actually improve welfare conditional on imperfect competition. This effect is generated by the type of selection on the intensive margin with respect to willingness-to-pay that affects insurer incentives. While this certainly affects the ability to generalize the results here, as insurer pricing decisions are potentially driven by different kinds of selection, this direct effect of prices on consumer sorting is likely to be of first-order importance.[39]

[Insert Table X]

[38] As a rule, I take the average across consumers within a market and report the median across markets.

[39] By contrast, imperfect competition has major implications for consumer welfare. Appendix B describes calculations of equilibrium in which prices are set equal to marginal or average cost or claims, and Table XI shows the results of such an exercise. Average claim pricing leads to a 44 percent reduction in prices; unsurprisingly, this dramatic price reduction leads to big gains in consumer welfare. The price reduction comes from two sources. First, insurers no longer have profit margins of nearly 30 percent: this simply represents a transfer from producers to consumers. In addition, lower prices bring more consumers into the market and these consumers are healthier on average, driving down average claims. This leads to better allocation of insurance. While the price reduction leads to a compensating variation of $644, the compensating variation net of the change in producer profits is equal to an average of $229 for the average consumer. This should be thought of as the increase in allocative efficiency, as the reduction in prices is simply a transfer from firms to consumers.

Marginal claim pricing leads to a 45 percent reduction in prices and a compensating variation of only $670 for the average consumer: the size of the gap between average and marginal claim is small. This is not surprising given that the wedge between marginal and average claim gets larger as consumers become more price sensitive; consumers are relatively price insensitive in this market. Therefore, it is reasonable to conclude that the major source of inefficiency in the Medigap market is imperfect competition and relative insensitivity of consumers to premiums, rather than imperfect information. The second panel of Table XI calculates prices under marginal and average *costs*, which include both the marginal or average claim and the administrative costs associated with marketing and administering each policy. The price reductions, while still substantial, are unsuprisingly smaller.

5.1.2 Commissions

Given that perfect risk adjustment does not lead to large gains in consumer surplus, the rents transferred to brokers through commissions are a natural public policy target. If policy interventions can limit broker commissions, this may improve outcomes. This is not an unreasonable strategy: regulation of minimum loss ratios, as discussed in the model, have the potential to limit the amount insurers can pay out to other players in their supply chain.[40] This counterfactual can also be framed as a window into a scenario in which brokers are not capturing substantial rents in this market.

The effect of reduced broker commissions on prices is ambiguous. While costs fall with commissions, reduced broker commissions also give insurers an incentive to increase prices. Because brokers previously captured a large percentage of the margin, insurers had less incentive to increase their prices. If insurers got to keep a larger percentage of their margins, they might raise prices. (Of course, they could also choose to structure payments to brokers as something other than a percent commission.) This counterfactual sets all broker commissions at 6.13% - the estimated cost for UnitedHealth. The results in Table X show that reducing commissions does little to increase consumer welfare; premiums fall only slightly compared to the current equilibrium.

Ultimately in this market, insurers are responding to the incentives generated by the nature of consumer demand. Consumer demand in this market gives insurers incentives to invest in consumer brand preferences (or, alternatively, availability to consumers who engage in costly search) and subsequently capture and distribute rents to other players in the supply chain, such as AARP. Policies aimed at marginal changes to insurer preferences will do little to improve outcomes for consumers. By contrast, policies aimed at improving the information available to consumers, reducing their search costs, and simplifying the problem may have large effects. The next subsection considers potential improvements of

[40] In PPACA, minimum loss ratios are set to 85% of premiums, potentially curtailing the ability of brokers to earn substantial commissions, especially in the small group and individual markets.

this type.

5.2 Information Provision

Broadly speaking, in this market, advertising can solve the problem that consumers may not know about all policies and prices (largely informative), or advertising can solve the fear of the unknown by providing a trusted name for an insurer, which affects the utility function. Regardless, from a policy perspective, regulators are faced with an information provision problem: how to ensure that consumers are aware of all policies available and are aware that all policies of the same letter are contractually equivalent. The policy would increase the amount of information available to consumers and focus competition on contractually relevant features such as price. This could be accomplished by designing health insurance exchanges that limit the role of brokers or providing a seal of approval for all policies. In this exercise, I assign all policies the same brand effect, recompute optimal prices, and compare prices to the current equilibrium. This requires an important choice: the level of the brand effect. This requires taking a stand on how consumers will perceive policies in the absence of marketing; this normalization will affect consumer utility and the consumer surplus calculation, but, generally speaking, not the price level. Therefore, despite the change in the utility function, every consumer is equally likely to purchase insurance in the new scenario.[41] Given this parameterization of consumer utility, I recompute the equilibrium. The median enrollment-weighted claims across markets falls by 9 percent as compared to the current equilibrium. This further confirms that the source of high premiums in the Medigap market is not due to adverse selection. Instead, consumers trade off brand preferences against lower prices, leading to higher premiums in equilibrium. In turn, firms compete to be the preferred brand, rather than the lowest price supplier. Furthermore, this result is not sensitive to the exact normalization of utility.[42] However, a critical assumption is that, despite that fact that consumers view policies of the same letter as close substitutes

[41] Consumers value brands equally, so heterogeneity is not a problem in this exercise.

[42] Robustness checks are in Appendix B.

(up to their draw on the error term), elasticities are unchanged. Given this assumption, the dramatic reduction in premiums paid is due to a reallocation of consumers across policies, rather than substantially lower list prices. If consumers additionally had a higher distaste for price, the effects might be larger; such an exercise is detailed in the Appendix B.[43]

The effects of advertising on consumers, and brand effects in consumer demand, are often difficult to understand and interpret. Despite this, they are pervasive features of many markets, including Medigap insurance. Policies aimed at increasing consumer surplus in this market must focus on imperfect competition arising from product differentiation through branding, barriers to entry, and large marketing expenditures if they are to be effective.

6 Conclusion

Although the literature focuses on asymmetric information in insurance markets as a driver of inefficiency, the estimated welfare effects of adverse selection are small, and imperfect competition is often the driver of inefficiency. Furthermore, adverse selection has been analyzed in the benchmark case of average cost pricing, ignoring the impact of strategic insurer pricing, and, therefore, potentially overstating its negative welfare effects. This paper argues that the intense focus on adverse selection obscures other features of insurance markets, particularly imperfect competition, that have a large impact on consumer surplus. Strategic insurer behavior, both pricing and marketing, plays a large role in shaping equilibrium outcomes. Therefore, in the Medigap market, perfect risk adjustment leads to a 7 percent reduction in total surplus, and, in the absence of marketing, prices fall by 9 percent.

[43] In fact, listed prices fall very little. The premiums of some lower-priced policies actually rise as these policies are now more valuable in eyes of consumers. In Appendix B, counterfactuals in which the mean price effect (and, therefore, elasticity) is increased by 10 percent are also computed. In the first, policies still have different brand effects. In this equilibrium, the posted prices fall dramatically, as does the average premium paid. By contrast, if marketing is also banned, the posted prices actually rise slightly, but the median average premium paid across markets falls by a similar amount (and the list prices rise less than the scenario in which elasticities are as estimated). Therefore, increasing elasticities leads to lower list prices, while eliminating brand effects leads to a reallocation of consumers into lower priced policies.

A key contribution of this paper is the ability to explore the sources and consequences of firm market power. Specifically, I show that the nature of consumer demand gives insurers an incentive to heavily market their policies and subsequently price policies above cost. I separate the difference between average claims and average premiums (one minus the minimum loss ratio) into marketing costs and variable profits. This analysis shows both that these marketing costs are potentially large and also vary across firms. Estimates of variable, fixed, and sunk costs indicate that the two largest firms engage in different strategies. UnitedHealth engages in direct sales, but pays a large fee to use the AARP brand, while Mutual of Omaha pays large commissions to brokers, dissipating some of the rents associated with the exertion of market power.

Future work should consider contract design; the absence of policies or policy features is potentially a large source of welfare loss from adverse selection. In addition, because I consider variation in preferences over premiums, my estimates reflect the largest possible price reduction due to selection. Furthermore, the design and implementation of exchanges creates a wealth of interesting and important research questions.

This paper has taken an important first step to illustrate that the exertion of market power by insurers can be a larger threat to consumer surplus than adverse selection, and that adverse selection may actually serve as a check on insurer pricing power. As health reform is enacted, consumers will face differentiated product markets facilitated by exchanges. Policy makers designing and regulating exchanges should be aware of the features of health insurance that shape insurer incentives. Providing better information to consumers, rather than focusing on the potential for adverse selection, can lead to more effective public policy and better outcomes for consumers.

(The Wharton School, University of Pennsylvania)

Appendix A: Demand Robustness Checks

In the most basic model of insurance, consumers value both their health status and consumption c_i. The two are linked because medical expenditures depend on health status and may, in turn, improve health. To simplify, consider the choice between supplement coverage and no supplemental coverage in a model in which consumers have utility over income ω_i and health which is a function of underlying status z_i and medical expenditure $m(z_i)$. If the consumer chooses no supplemental coverage, they face the cost of their medical expenditure, while if the consumer chooses supplemental coverage, they simply pay the policy's premium regardless of the state of the world. Let the states of the world be indexed by s with probability $\pi(s, z_i)$, which vary with individual health status, and expenditure by $m(s, z_i)$ where $j = 1$ if the consumer purchases insurance. Consumer utility is given by:

$$V(c_i, z_i) = \int \left(U(z, \omega_i - p_j)\, 1(j = 1) + U(z_i, \omega_i - m(s, z_i))\, 1(j = 0)\right) d\pi(s, z_i)$$

and with sufficiently restrictive functional forms on the utility over consumption and health and knowledge of the potential states of the world s, the econometrician can estimate coefficients of interest, such as a coefficient of absolute risk aversion and the consumer's (potentially subjective) probabilities $\pi(s)$.

However, this model is already a simplification (full insurance, a single policy) and still cannot be estimated in this context. Except for binary cases or extremely detailed data, it is nearly impossible to even innumerate the states of the world due to data limitations[44]. Therefore, one can consider a simplification of this model where the consumer has a valuation over insurance contracts ϕ_j. In such a model, the valuation, which may be contingent on health status or income ω_i, is denoted by $v(\phi_j, z_i, \omega_i; \theta^d) = f\left(E\left[m\left(\phi_j, z_i\right)\right], z_i; \theta^d\right) - \alpha(\omega_\iota, z_i)\, p_j$, where consumers value their expected medical expenditure under the contract, and utility over insurance depends on health status and price sensitivity is also individual-specific.

In the model detailed in the paper, consumers have two defining characteristics: income and expected claim. The expected claims is constructed using a regression of private insurance claims on demographics (including income) and self-reported health, using only the Medigap population. This prediction may be troublesome if changes in the environment lead to dramatic changes in the population purchasing Medigap (the prediction is out-of-sample). However, simpler specifications produce similar results. Other specifications also include the variance of the expected claim or squared claims to capture consumer desire for "stop-loss" coverage of Part B services; these changes do not affect the qualitative results. Presented here are results from a demand system in which consumers are defined by income and self-reported health. I de-mean income by self-reported health status so that the two are not correlated and repeat the analysis in Section 3. The results are very similar. Wealthier consumers and sicker consumers are less price sensitive, and a $100 increase in premiums

[44] This is the approach taken, for example, in Cohen and Einav (2007) and Handel (2009). In the first case, the authors model deductible choice (high or low), while Handel considers an employer-sponsored context and has extensive data on claims and out-of-pocket costs under various policies.

is associated with a $9 increase in claims (slightly smaller than the previous estimates). The elasticities are nearly identical. Therefore, while it critical to incorporate consumer heterogeneity in the demand estimates, the exact specification is not critical to the results.

Table A1: Nested Logit Demand Estimates

	(1)	(2)	(3)	(4)
Within-Group Share		0.680***		0.681***
		(0.0248)		(0.0251)
Premium	-0.0283***	-0.0285***	-0.0545***	-0.0633***
	(0.00180)	(0.00175)	(0.00621)	(0.00609)
B	0.370***	0.403***	0.484***	0.555***
	(0.0689)	(0.0672)	(0.0741)	(0.0727)
C	0.799***	0.815***	0.947***	1.012***
	(0.0610)	(0.0595)	(0.0701)	(0.0687)
D	0.907***	0.913***	0.986***	1.018***
	(0.0612)	(0.0597)	(0.0642)	(0.0630)
E	0.882***	0.896***	0.939***	0.972***
	(0.0817)	(0.0796)	(0.0833)	(0.0817)
F	1.972***	1.980***	2.057***	2.094***
	(0.0538)	(0.0525)	(0.0576)	(0.0565)
G	1.272***	1.283***	1.342***	1.377***
	(0.0606)	(0.0590)	(0.0631)	(0.0618)
H	-0.0511	-0.0988	-0.0178	-0.0546
	(0.115)	(0.112)	(0.116)	(0.114)
I	0.424***	0.458***	0.507***	0.568***
	(0.106)	(0.103)	(0.108)	(0.106)
J	2.310***	2.304***	2.329***	2.328***
	(0.0799)	(0.0778)	(0.0806)	(0.0790)
K	-0.732***	-0.783***	-0.884***	-0.986***
	(0.109)	(0.107)	(0.115)	(0.113)
L	-0.444***	-0.495***	-0.522***	-0.600***
	(0.105)	(0.102)	(0.107)	(0.105)
Constant	-9.896***	-8.693***	-9.815***	-8.582***
	(0.198)	(0.197)	(0.200)	(0.201)
Fixed Effect	brand	brand	brand	brand
Observations	14,146	14,146	14,146	14,146
R-squared	0.392	0.423	0.382	0.406

Table A1a: Demand Estimates

	(1)	(2)	(3)
Price	−0.0767	−0.1053	−0.1049
	(0.0070)	(0.0070)	(0.0070)
Random Coefficients			
Self-Reported Health		0.0106	0.0104
		(0.0004)	(0.0055)
Income*Price	0.0002	0.0002	0.0002
	(0.000003)	(0.000004)	(0.000010)
Price*Self-Reported Health			0.0023
			(0.1398)
Log Claim Regressed on Consumer Characteristic			
Self-Reported Health	−0.0795	0.6696	0.6581
	(0.0338)	(0.0345)	(0.0278)
Income	0.0075	0.0012	0.0013
	(0.0005)	(0.0008)	(0.0284)
N	14146	14146	14146
Mean Price Semi-Elasticity	-1.1301	-1.1227	-1.1230
Mean $\frac{\partial \gamma}{\partial p}$	0.0337	0.0935	0.0906

Additionally, the demand results are robust to including random coefficients on the level of coverage (defined as the percentage of the Part A deductible or Part B copayment paid by the policy) or including random coefficients on the AARP brand effect. This is meant to capture the fact that consumers who are sicker may have a higher preference for a "secure" policy.

Table A2: Demand Estimates II

	(1)	(2)	(3)
Price	−0.0967	−0.0973	−0.0964
	(0.0075)	(0.0075)	(0.0075)
Random Coefficients			
E(Claim)*Coverage		−0.0006	
		(0.0057)	
E(Claim)*AARP Brand Effect	0.0047		0.0129
	(0.0048)		(0.0004)
Income*Price	0.0005	0.0005	0.0005
	(0.00002)	(0.1233)	(0.00003)
Price*E(Claim)	0.0003	0.0005	
	(0.0001)	(0.0001)	
Log Claim Regressed on Consumer Characteristic			
Income	0.0185	0.0189	0.0175
	(0.0196)	(0.0254)	(0.0006)
E(Claim)	0.0536	0.0634	0.0348
	(0.2087)	(0.1838)	(0.0048)
N	14146	14146	14146
Mean Price Semi-Elasticity	-1.1173	-1.1227	-0.9324
Mean $\frac{\partial \gamma}{\partial p}$	0.1228	0.1433	0.2903

Appendix B: Additional Counterfactuals

Entry

I will consider the potential for a firm to duplicate the strategy of either UnitedHealth or Mutual of Omaha by offering the same mix of products with the same brand preferences induced by similar marketing expenditures. I will estimate the fixed or sunk cost of such a strategy. However, this analysis will require a number of additional assumptions. I will not be able to separately identify all of the fixed and sunk costs. However, institutional detail can focus my analysis and allow for reasonable assumptions that capture the costs most relevant to insurers. First, I assume that the fixed cost of offering an additional policy conditional on being active in a market, ϕ_{jm} is zero or small[45]. Furthermore, firms may incur a fixed cost of entry or marketing in an additional market, denoted by Θ_{jm}. This may encapsulate the cost of setting up a network of agents in a given geographic area to sell policies or the cost of being certified by the state insurance commissioner to sell Medigap insurance. The degree to which insurers choose to enter markets and offer policies is likely to impact pricing and consumer surplus.

Because of Mutual of Omaha's strategy of employing brokers to sell their policies, this is likely to be the most relevant cost they face. UnitedHealth, however, faces one large sunk cost Φ_J - the cost of the AARP endorsement - that is likely to dwarf any market level fixed costs. Therefore, while I cannot separately identify fixed cost of entry into a specific geographic market and sunk costs of national operation, I will focus on the fixed costs of Mutual of Omaha's strategy and the sunk cost of UnitedHealth's strategy. Conditional on this set of potential entrants, the assumption identifying fixed or sunk costs follows Mazzeo's model of entry[46], modified to incorporate an inequality approach. This will require the calculations of the profits of a second firm duplicating the strategy of one of the two largest firms. This estimate is formed by adding a firm with the same brand effect and policies, and recomputing optimal prices. Given optimal prices, I can calculate costs and shares, which are functions of prices, and sum across policies to obtain an estimate of variable profits.

First consider Mutual of Omaha. Because I observe one firm employing Mutual of Omaha's strategy, I assume that two inequalities hold. First, Mutual of Omaha is currently profitable; in addition, a new entrant mimicking the strategy of Mutual of Omaha would be unprofitable, because such a firm does not exist. Taken together, these inequalities will allow me to bound fixed costs of entry into an addition market. In order to simplify notation, let $\Pi_{Mm}(2,1)$ be the profits of a firm employing Mutual of Omaha's strategy in a market in which two firms employ Mutual of Omaha's strategy and one firm employes UnitedHealth's strategy. Given the estimates from the demand system and the estimates of claims and administrative costs, the assumption above, along with the assumption that the firms expectations are, on average, correct, implies:

$$\Pi_{Mm}(1,1) = \sum_{j \in J_f} \pi_M(1,1) \geq \Theta_{Mm} \geq \sum_{j \in J_f} \pi_M(1,2) = \Pi_{Mm}(2,1)$$

[45] See Iyengar et al (2000) for a discussion of the negative impacts of additional choice.

[46] Mazzeo analyzes a game in which hotels choose a differentiated type (low, middle, or high quality) and whether or not to enter a specific market. I will consider two types based on presumed cost structure rather than quality.

The mean of the distribution of the lower bound of fixed costs for Mutual of Omaha can be written as:

$$\underline{\beta}_M = E_m \left(\Pi_{Mm} \left(2, 1 \right) \right)$$

Likewise, the mean of the distribution of the upper bound of fixed costs for Mutual of Omaha can be written as:

$$\overline{\beta}_M = E_m \left(\Pi_{Mm} \left(1, 1 \right) \right)$$

and likewise for high marginal cost firms. These equations imply that the lower bound of Mutual of Omaha's fixed costs is equal to the expectation of variable profits of a firm duplicating their strategy; likewise, the upper bound of firm $f's$ fixed costs is equal to the expectation of variable profits. Estimates of $\underline{\beta}_M$ and $\overline{\beta}_M$ are easily obtained by taking sample averages, and standard errors can be obtained using the bootstrap procedure outlined in Pakes, Porter, Ho, and Ishii (2006). The estimates are based on a single firm across markets; therefore, the standard errors are large.

By contrast, the biggest cost facing UnitedHealth is the sunk, national level cost of acquiring the AARP seal of approval. Therefore, I will assume that Θ_{Um} is equal to zero and focus on this sunk cost Φ_J plus national level advertising. The assumption that identifies this sunk cost is the same: while UnitedHealth is profitable, it would not be profitable for another firm to replicate UnitedHealth's strategy. This can be written as:

$$\sum_{m \in J_{mt}} \sum_{j \in J_f} \pi \left(1, 1 \right) \geq \Phi_U \geq \sum_{m \in J_{mt}} \sum_{j \in J_f} \pi \left(1, 2 \right)$$

The mean of the distribution of the lower bound of sunk costs for UnitedHealth can be written as:

$$\underline{\beta}_U = E_t \left(\sum_{m \in J_{mt}} \Pi_{Um} \left(2, 1 \right) \right)$$

where t denotes a year. Because I only have three years of data, my estimates will be imprecise. Likewise, the mean of the distribution of the upper bound of sunk costs for UnitedHealth can be written as:

$$\overline{\beta}_U = E_t \left(\sum_{m \in J_{mt}} \sum_{j \in J_f} \Pi_{Um} \left(1, 1 \right) \right)$$

The confidence interval for UnitedHealth's sunk cost will be based on only three years of data. Therefore, the standard errors are quite large. Despite this, the estimates of fixed cost of entering another market and the sunk cost of national market provides a number of insights into the Medigap market by illustrating how the strategies of different firms can shape outcomes.

In addition to marginal costs incurred beyond claims, fixed and sunk costs shape this market. Table 8 describes these costs. The fixed cost of a firm duplicating Mutual of Omaha's strategy entering an additional market (a state) is between $450,000 and $796,000; across states, this implies fixed costs between $22 and $98 million. By contrast, the sunk cost of a low marginal cost firm entering all markets is between $100 million and $488 million.

The latter number in each pair represents an aggregation of the variable profits observed in the data. The former represents the mean of variable profits of the firm duplicating UnitedHealth or Mutual of Omaha across years or markets in a counterfactual calculation. The estimates are extremely noisy; this is unsuprising, as the estimates of UnitedHealth's sunk costs are based on three observations: one from each year of the data and the estimates of Mutual of Omaha are based on estimates from a single firm.

However, these estimates contain some insight into the operation of the Medigap market. AARP collects nearly five hundred million - over half of its revenues - from endorsements, royalties, and fees. Nearly half of this is collected from UnitedHealth alone. Though United sells a number of lines with the AARP endorsement, and can spread costs across these lines, they must also maintain some website and phone bank infrastructure as well as paying for advertising on television and on the internet. The sunk costs are in line with these costs; the lower bound is slightly less than the royalties and fees paid to AARP and the upper bound allows for additional costs of advertising. Despite the large standard errors, these estimates reflect the right order of magnitude for sunk costs and explain why an additional firm might not choose to replicate UnitedHealth's strategy: the cost of purchasing an endorsement is likely to be higher than the variable profits of a second firm, but lower than UnitedHealth's profits. Furthermore, the sunk costs estimates illustrate why AARP provides its brand exclusively to a single firm: UnitedHealth's current profits are more than the combined counterfactual profits of UnitedHealth and a second firm with the AARP brand.

Table A4 describes the total costs of marketing to consumers. As expected, I estimate that UnitedHealth spends more on sunk costs of marketing than Mutual of Omaha does on fixed cost of marketing. On a per consumer basis, UnitedHealth spends between \$24 and \$73 marketing its policies, while Mutual of Omaha spends between \$8 and \$15 per consumer. By contrast, Mutual of Omaha spends substantially more on broker costs. As a result, Mutual of Omaha spends almost twice as much - between \$247 and \$253 - per consumer on marketing and administration; however, this is only because UnitedHealth spreads its sunk costs over such a high volume of consumers.

Table A3: Fixed and Sunk Cost Estimates

	Lower Bound	Upper Bound
Sunk Cost,	\$99,261,645.01	\$487,935,210.41
UnitedHealth	(\$1,530,902,861,706.31)	(\$23,031,614,127.02)
Fixed Cost,	\$445,010.32	\$796,342.56
Mutual of Omaha	(\$225,593.04)	(\$3,578,033.82)

Profits are given in dollars. Sunk costs are calculated at a national level, while fixed costs are calculated at the state (market) level. Standard errors are computed using the bootstrap procedure outlined in Pakes, Porter, Ho, and Ishii (2006).

Table A4: Marketing Expenditure and Advertising Value

	UnitedHealth	Mutual of Omaha
L.B. of Sunk (Fixed) Cost/Consumer	$23.65	$8.37
U.B. of Sunk (Fixed) Cost/Consumer	$73.09	$14.81
Average Marginal Cost/Consumer	$98.27	$238.67
L.B. of Total Marketing Cost/Consumer	$121.92	$247.05
U.B. of Total Marketing Cost/Consumer	$171.36	$253.48
Value of Brand to Consumers	$959.10	$104.07

Compensating variation is calculates as the average across consumers within a market using the standard log-sum formula; the number reported is the median across markets.

Marginal and Average Cost and Claim Pricing

Section 5 describes the results of marginal and average cost and claim pricing. Because the average claim depends directly on the price charged, an equilibrium is a fixed point in which $\mathbf{p} = \gamma_j(\mathbf{p})$. In this equilibrium, marginal administrative costs, including broker fees, are set to zero. Likewise, the marginal claim depends directly on the price charged, an equilibrium is a fixed point in which $\mathbf{p} = m\gamma_j(\mathbf{p})$[47]. Marginal claims are calculated using the method proposed by Einav, Finkelstein, and Cullen (2010) as the derivative of total claims with respect to price divided by the derivative of share with respect to price:

$$m\gamma_j(\mathbf{p}) = \frac{\frac{\partial \gamma_{jm}}{\partial p_{jm}} s_{jm} + \frac{\partial s_{jm}}{\partial p_{jm}} \gamma_{jm}}{\frac{\partial s_{jm}}{\partial p_{jm}}}$$

Equilibriums in which prices are set to marginal and average *costs*, equal to claims plus administrative costs, including broker fees, are also estimated. Average cost is equal to the average claim plus estimated administrative costs, and marginal cost is defined as the marginal claim plus the estimated administrative cost.

Subsidies (and, in the case of current reform, targeted tax credits) are often suggested as a remedy for adverse selection in insurance markets; the problem is inefficient under-provision of insurance and a price reduction induced by a subsidy is a natural solution. Furthermore, a subsidy has the added bonus of keeping healthy consumers, who may otherwise choose not to buy insurance if their expect claim is below the premium. By contrast, Medigap insurance lowers the effective price for Medicare-covered services. That is, there is not only a potential moral hazard problem with respect to Medigap covered (insured) expenditure, but expenditure in the primary insurance (Medicare) as well. As a result, the Congressional Budget Office has suggested a 5-10% tax on Medigap policies. However, taxes could potentially exacerbate adverse selection.

This counterfactual simply applies a tax or a subsidy on the premiums of Medigap policies to assess the impact on consumer welfare. The results are in Table A6. Subsidies do little to improve consumer welfare; prices fall, but by substantially less than the amount of the

[47] This is done at the policy level, without incorporating any individual-specific knowledge of claims. In a market with a single policy, if premiums are set at the marginal claim, all consumers for whom the value of insurance is greater than their claim will purchase insurance. Given more information on individual claims and preferences, insurers could effectively price discriminate.

subsidy. Premiums are reduced by 3% under a 10% subsidy and less than 1% under a 5% subsidy. Because consumer demand is relatively unresponsive to price, the "incidence" of the subsidy falls on insurers - they simply increase their prices. Because the subsidy simply represents a transfer, the change in total surplus is effectively zero in both cases. Similarly, taxes raise prices by more that the amount of tax. The welfare effect is unmitigated by better sorting. Furthermore, consumers are hurt by higher prices without an increase in total surplus. However, there is an important caveat to this analysis: it abstracts from the role of moral hazard and the potential for less utilization of Medicare services, and the potential for more efficient funding of Medicare expenditures.

Table A5: Additional Counterfactual Results

	10% Subsidy	5% Subsidy
Median Average Premium Paid	$1258.60	$13.0937
% Change in Premium	-0.0341	-0.0067
% Change in Lives Covered	0.0154	0.0038
Median Compensating Variation	$33.97	$8.38
% Change in Total Surplus	-0.0003	-0.0054
Taxes		
	5% Tax	10% Tax
Median Average Premium Paid	$1392.62	$1442.11
% Change in Premium	0.0495	0.0974
% Change in Lives Covered	-0.0503	-0.0989
Median Compensating Variation	-$72.43	-$144.78
% Change in Total Surplus	-0.0201	-0.0354

Loss Ratio Regulation

In the insurance industry, tight regulation is often used in place of strict antitrust enforcement. In the Medigap market, a key regulation is the mandatory loss ratio, which in theory requires sixty-five percent of premiums to be paid out in claims. This exercise asks how those policies would be priced if the regulation was lifted. However, this regulation is not strictly enforced in the data: nearly two-fifths of policies violate the regulation. In the current equilibrium, the first order condition with respect to the price of policies in violation of regulation is allowed to be positive: firms may like to raise their prices, but are constrained by the regulation. In this exercise, the first order condition with respect to all prices is set equal to zero.

The table below shows the results of this exercise. Without mandatory loss ratio regulation, firms would raise prices; this seems intuitive, as the mere presence of regulation implies that it may be profitable to do so. As a result, the median price paid by consumers increases by 14%. This both reduces the size of the market and reduces consumer surplus. Nearly twenty percent fewer consumers purchase Medigap insurance and the compensating variation indicates that mandatory loss ratio regulation is worth $340 to the median consumer, and mandatory loss ratios increase total surplus by 6%. As a result, stricter enforcement may be an efficient mechanism for increasing consumer surplus.

Table A6: Additional Counterfactual Results

Loss Ratio Regulation		
	$\frac{\partial \pi_{Jm}}{\partial p_{jm}} = 0$	Regulation Binding
Median Average Premium Paid	\$14.9882	\$7.7612
% Change in Premium	0.1368	-0.3187
% Change in Lives Covered	-0.1970	0.3032
Median Compensating Variation	-\$339.77	\$362.83
% Change in Total Surplus	-0.0647	0.0724
% of Policies with Positive profits	0.9167	0.3371

Additionally, there are a number of results that are referenced in footnotes. They are largely robustness checks and are reproduced here:

Table A7: Counterfactual Results Referenced in Footnotes

	Brand "Degradation"
Median Average Premium Paid	\$1131.14
% Change in Premium	-20%
% Change in Lives Covered	5%
Mean Compensating Variation	\$71.54
Mean Increase in Consumer Surplus	\$250,019,535.29
Marketing	
	UnitedHealth Utility
Median Average Premium Paid	\$1284.43
% Change in Premium	-0.1143
% Change in List Prices	5.76%
UnitedHealth Entrant	
Median Average Premium Paid	\$1114.31
% Change in Premium Paid	-21.26%
% Change in Lives Covered	50.02%
Mean Compensating Variation	\$576.29
Mean Increase in Consumer Surplus	\$1,470,141,130.19
Mean Increase in Consumer Surplus due to Lower Premiums	\$553,490,083.63
Upper Bound of Sunk Costs	\$487,935,210.41

References

Alecxih, L., S. Lutzky, P. Sevak, and G. Claxton (1997). Key issues affecting accessibility to Medigap insurance. *New York: Commonwealth Fund*.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica 63*(4), 841–890.

Bhattacharya, D. (2005). Asymptotic inference from multi-stage samples. *Journal of econometrics 126*(1), 145–171.

Bundorf, K., J. Levin, and N. Mahoney (2008). Pricing, matching and efficiency in health plan choice. *mimeograph, Stanford University*.

Cardon, J. and I. Hendel (2001). Asymmetric information in health insurance: evidence from the National Medical Expenditure Survey. *The RAND Journal of Economics 32*(3), 408–427.

Cebul, R., J. Rebitzer, L. Taylor, and M. Votruba (2008). Unhealthy insurance markets: search frictions and the cost and quality of health insurance.

Chiappori, P. and B. Salanié (2000). Testing for asymmetric information in insurance markets. *The Journal of Political Economy 108*(1), 56–78.

Cutler, D., A. Finkelstein, and K. McGarry (2008). Preference heterogeneity and insurance markets: Explaining a puzzle of insurance.

Dafny, L., M. Duggan, and S. Ramanarayanan (2009). Paying a premium on your premium? Consolidation in the US health insurance industry. *National Bureau of Economic Research*.

Dafny, L., K. Ho, and M. Varela (2010). An individual health plan exchange: which employees would benefit and why. *American Economic Review 100*(2), 485.

Einav, L., A. Finkelstein, and M. Cullen (2010). Estimating Welfare in Insurance Markets Using Variation in Prices. *Quarterly Journal of Economics 125*(3), 877–921.

Fang, H., M. P. Keane, and D. Silverman (2008). Sources of advantageous selection: Evidence from the medigap insurance market. *Journal of Political Economy 116*(2), pp. 303–350.

Finkelstein, A. (2004). Minimum standards, insurance regulation and adverse selection: evidence from the Medigap market. *Journal of Public Economics 88*(12), 2515–2547.

Gross, D., L. Alecxih, M. Gibson, J. Corea, C. Caplan, and N. Brangan (1999). Out-of-pocket health spending by poor and near-poor elderly Medicare beneficiaries. *Health Services Research 34*(1 Pt 2), 241.

Hall, M. (2000). The role of independent agents in the success of health insurance market reforms. *Milbank Quarterly 78*(1), 23–45.

Handel, B. (2010). Adverse selection and switching costs in health insurance markets: When nudging hurts. *manuscript, Northwestern University*.

Imbens, G. W. and T. Lancaster (1994). Combining micro and macro data in microeconometric models. *The Review of Economic Studies 61*(4), pp. 655–680.

Iyengar, S. and M. Lepper (2000). When choice is demotivating: Can one desire too much of a good thing?. *Journal of personality and social psychology 79*(6), 995.

Litow, M. and F. Actuary (2006). Medicare versus Private Health Insurance: The Cost of Administration. *Milliman, January 6*.

Maestas, N., M. Schroeder, and D. Goldman (2009). Price Variation in Markets with Homogeneous Goods: The Case of Medigap.

Robst, J. (2006). Estimation of a Hedonic Pricing Model for Medigap Insurance. *Health services research 41*(6), 2097–2113.

Sutton, J. (1992). *Sunk costs and market structure*. MIT press.

Table I. Medicare Supplement Plans

	A	B	C	D	E	F	G	H	I	J	K	L
Part A Coinsurance	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Part B Coinsurance	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	$\frac{1}{2}$	$\frac{3}{4}$
Blood	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	$\frac{1}{2}$	$\frac{3}{4}$
Hospice											$\frac{1}{2}$	$\frac{3}{4}$
Skilled Nursing			✓	✓	✓	✓	✓	✓	✓	✓	$\frac{1}{2}$	$\frac{3}{4}$
Part A Deductible		✓	✓	✓	✓	✓	✓	✓	✓	✓	$\frac{1}{2}$	$\frac{3}{4}$
Part B Deductible			✓			✓				✓		
Part B Excess Charges						✓	$\frac{8}{10}$		✓	✓		
Foreign Travel Emergency			✓	✓	✓	✓	✓	✓	✓	✓		
At-Home Recovery							✓		✓	✓		
Preventative Care Coinsurance	✓	✓	✓	✓								
Preventative Care						✓				✓		
Market Share	.04	.03	.12	.04	.02	.49	.08	.01	.01	.15	.01	.01

Notes: Percentages do not add to one because of rounding.

Table II. Medigap Market Shares

	National Market Share	Share of Markets Active	Avg. Premium ($)
Unitedhealth	.4566	.9796	1534.82
Mutual of Omaha	.2352	.9524	1398.38
Conseco	.0868	.8980	1615.26
American Financial	.0382	.7755	1630.09
HCHSC	.0313	.0544	1815.55
Genworth Financial	.0233	.8776	1517.81
State Farm	.0229	.5918	2159.99
American Republic Mutual	.0195	.5306	1323.05
Universal American Financial	.0120	.7891	1771.63
Guarantee Trust	.0105	.4966	1756.02
Physicians Mutual	.0097	.6803	1596.92
USAA	.0094	.8980	1677.31
American National Financial	.0080	.6735	1247.75
Atlantic American	.0080	.6259	1531.27
Thrivent Financial for Lutherans	.0071	.3810	1629.46
State Mutual Company	.0067	.1633	703.04
Humana	.0066	.6734	1247.23
Liberty National	.0064	.8778	1736.36

Source: NAIC Market Level Data.

Table III. Plan Level Premiums and Claims

Plan	Profit Rate (%)	# of Policies	Weighted Premium ($)	Weighted Claim ($)
A	.2006	1403	1457.25	1223.80
	(.8152)		(743.26)	(1161.64)
B	.2192	1079	1562.31	1218.03
	(.9803)		(493.55)	(554.31)
C	.2387	1764	1729.26	1398.40
	(.5814)		(389.37)	(460.22)
D	.3182	1822	1546.58	1150.81
	(.5145)		(459.44)	(451.93)
E	.3055	668	1691.22	1235.19
	(.3863)		(511.84)	(459.03)
F	.3213	3518	1518.81	1170.77
	(.4834)		(663.37)	(524.24)
G	.3228	1936	1500.26	1094.19
	(.4301)		(446.44)	(380.99)
H	.2414	266	1379.37	1033.05
	(.4582)		(1379.37)	(493.11)
I	.3778	327	1675.13	1252.45
	(.3777)		(352.85)	(310.48)
J	.3539	716	1503.10	1130.17
	(.4335)		(380.23)	(341.00)
K	.4543	308	712.59	477.71
	(.4739)		(196.40)	(183.35)
L	.3600	339	183.35	784.10
	(.5218)		(263.68)	(784.30)

Source: NAIC Market Level Data. Standard deviations in parentheses.

Table IV. Regressions of Claims on Price

Dependent Variable: Average Claim			
	(1)	(2)	(3)
Price	.8025***	.7930***	.4318***
	(.0109)	(.0111)	(.1336)
C	−1.113**	−.9150**	1.328
	(.3976)	(.3974)	(.9237)
F	−2.584***	−2.393***	−1.025*
	(.3480)	(.3492)	(.6206)
J	−1.791***	−1.683***	−1.331**
	(.5031)	(.5034)	(.5379)
First Stage			
Retaliatory Tax Dummy			−1.111***
			(.2166)
% Increase in Taxes Above Local Base			1.482***
			(.1452)
Observations	14146	14146	14146
F-Stat on Excluded Instruments			52.45***
Adjusted R^2	0.2941	0.3013	0.2486
Controls	policy	policy, state*year	policy, state*year

Notes: OLS and IV regressions of claims on premiums and covariates. "% Increase in Taxes Above the Local Base" is the difference between the retaliatory tax rate (if applied) and the local tax rate. "Retaliatory Tax Dummy" takes on a value of one if the policy is subjected to a retaliatory tax on premiums. Significance levels are denoted by asteriks (* p<0.1, ** p<0.05, *** p<0.01).

Table V. Individual Level Summary Statistics

	Medigap Enrollees	Non-Enrollees	Entire Sample
Total Health	$12,605.12	$15,505.71	$14,737.4
Expenditure	(20729.97)	(25401.41)	(24284.55)
Medicare	7552.43	6504.89	6782.365
Expenditure	(16015.92)	(14857.11)	(15178.9)
% Female	.3958	.4309	.4216
	(.4891)	(.4952)	(.4938)
Age (years)	78.88	78.11	78.31
	(7.815)	(8.140)	(8.062)
% College	.1597	.1656	.1641
	(.3664)	(.3719)	(.3703)
% Medigap	1	0	.2649
			(.4412)

Source: Medicare Current Beneficiary Survey. Notes: Medigap coverage is defined as in Fang, Keane, and Silverman (2008). Standard deviations in parentheses.

Table VI.. Subsample Characteristics for New Medigap Enrollees

	Subsample
Income	$\$36,803.60$
	$(57,278.53)$
Private Health Insurance Expenditure	\$389.92
	(1988.80)
E(Private Health Insurance Expenditure)	\$1537.99
	(77.47)
% Medigap	.2324
	(.4224)

Source: Medicare Current Beneficiary Survey. Notes: The subsample was created to capture those consumers who are likely to be purchasing Medigap coverage for the first time; the plan level data contain exclusively new enrollees.

Table VII. Construction of Expectation of Expenditure

Dependent Variable: Claim	
Income	.00001
	(.00002)
Female	2.7254
	(1.2055)
Age-65	.6100
	(.6443)
(Age-65)^2	−.0346
	(.0438)
(Age-65)^3	.0006
	(.0009)
Self Reported Health (Omitted Category=Excellent)	
Very Good	1.5366
	(1.5366)
Good	6.3365
	(1.7777)
Average	15.089
	(2.128)
Poor	32.530
	(3.178)
Adjusted R^2	0.0584
N	2591

Source: Medicare Current Beneficiary Survey. The regression is run on the non-age-restricted sample, conditional on Medigap purchase. Also included are controls for education, race, and region. This construction allows for the computation of expected expenditure for individuals choosing not to purchase insurance.

Table VIII. Key Parameters from the Demand System

	(1)	(2)	(3)
Price	−0.0824	−0.0817	−0.0826
	(0.0075)	(0.0075)	(0.0076)
C	0.5382	0.5341	0.5308
	(0.0883)	(0.0884)	(0.0883)
F	1.5856	1.5836	1.5834
	(0.0769)	(0.0770)	(0.0770)
J	1.8896	1.8881	1.8867
	(0.0959)	(0.0961)	(0.0959)
Random Coefficients			
E(Claim)		0.0174	−0.0003
		(0.0014)	(0.0049)
Income*Price	0.0002	0.0002	0.0002
	(0.00002)	(0.00006)	(0.00006)
Price*E(Claim)			0.0012
			(0.00009)
Log Claim Regressed on Consumer Characteristic			
Income	0.0172	0.0169	0.0047
	(0.0008)	(0.0002)	(0.0047)
E(Claim)	0.0364	0.0301	0.1375
	(0.0084)	(0.0069)	(0.0055)
N	14146	14146	14146
Mean Price Elasticity	-1.1141	-1.1147	-1.1191
Mean $\frac{\partial \gamma}{\partial p}$	0.1034	0.1244	0.1504

Notes: Source: MCBS data, NAIC data, and author calculations describes in the text in detail. Astriks denoting significance are omitted. Throughout, a complete vector of all plan dummies and brand dummies are included in the demand moments and plan and market dummies are included in the claims moments. For simplicity, I report only the coefficients on the three most popular plans in consumer preferences. Claims differences among plans, except for K and L, are qualitatively and quantitatively similar in IV specifications, as seen in Table IV. The main driver of differences in claims is largely captured by price in the regressions in Table IV and by consumer characteristics in these specifications.

Table IX. Variable Administrative Cost Estimates

	(1)	(2)	(3)
Constant	0.0833	0.1234	
	(0.0053)	(0.0243)	
Unitedhealth Group		−0.0633	0.0613
		(0.0249)	(0.0047)
Mutual of Omaha Group		0.0477	0.1844
		(0.0506)	(0.0348)

Notes: Source: MCBS data, NAIC data, and author calculations describes in the text in detail. Standard errors, obtained from bootstrapped estimates, are in parentheses.

Table X. Counterfactuals

	$\frac{\partial \gamma_{jm}}{\partial p_{jm}} = 0$	No Commission	Information Provision
Median Premium Paid	\$1406.87	\$1384.99	\$1296.43
% Change in Premium Paid	9.21%	-1.25%	-9.17%
% Change in Total Surplus	-6.52%	-	-

Notes: median premium paid is calculated as the median average premium paid across all state-year markets. The percentage change in average premium paid is calculated similarly. When noted, the change in total surplus includes both compensating variation and insurer profits. Compensating variation is calculates as the average across consumers within a market using the standard log-sum formula; the number reported is the median across markets.

Table XI. Average and Marginal Claim and Cost Pricing

Pricing	Average Claim	Marginal Claim
Median Premium Paid	\$739.34	\$732.06
% Change in Premium Paid	-44.29%	-45.52%
% Change in Lives Covered	55.31%	56.00%
Median Compensating Variation	\$643.51	\$670.42
Median Compensating Variation Net of Change in Producer Profits	\$229.50	\$229.54
With Administrative Costs		
Median Premium Paid	\$808.99	\$799.41
% Change in Premium Paid	-39.37%	-40.23%
% Change in Lives Covered	46.08%	47.19%
Median Compensating Variation	\$564.50	\$568.44
Median Compensating Variation Net of Change in Producer Profits	\$204.45	\$211.38

Notes: Median premium paid is calculated as the median average premium paid across all state-year markets. The percentage change in average premium paid is calculated similarly. The change in lives covered is also the median across markets. Compensating variation is calculated as the average across consumers within a market using the standard log-sum formula; the number reported is the median across markets.

Figure I: Relationship between Premiums and Claims