Do private Medicare firms have lower costs?

Job Market Paper

Keaton Miller*

Department of Economics
University of Minnesota

December 9, 2014

Abstract

The US government's Medicare Advantage (MA) program offers subsidies to private insurers who then compete to provide Medicare and supplemental benefits to seniors. Today, the subsidies paid to firms exceed Medicare's costs, but almost all plans provide a number of benefits on top of those provided by traditional Medicare. To evaluate the welfare effect of the MA program, I estimate the cost to firms of providing Medicare-equivalent benefits and the consumer welfare gains from supplemental benefits. I introduce a model of supply and demand for MA that focuses on supply-side dynamics. In my model, firms choose prices and generosities for multiple plans and take into account the inter-temporal incentives generated by the existence of switching costs on the demand side. I estimate the model's parameters using panel data on consumers and plans from 2008-2010 and find that, on average, private firms spend \$5,293 to provide Medicare-equivalent benefits to a healthy individual, whereas Medicare spends \$4,390 on similar individuals. The average plan spends \$184 on supplemental benefits, generating \$328 in consumer welfare. In counterfactual simulations, I explore the effect of a reduction in the subsidies offered to firms and policies that make it easier to switch providers. I find both policies increase total welfare, primarily by moving customers to traditional Medicare.

*I am extremely grateful to Amil Petrin for advice and support, as well as Thomas Holmes, Joel Waldfogel, and Robert Town for feedback throughout this project. Thanks are due to Michael Chernew and Harvard Medical School for data support. Finally, I thank Kailin Clarke, Boyoung Seo, Thomas Youle, Scott Petty, and the participants of the applied microeconomics workshop at the University of Minnesota for helpful feedback. This work was supported in part by the University of Minnesota Supercomputing Institute. The most recent version of this paper can be found at `http://www.econ.umn.edu/~mill2171/miller-jmp.pdf`

1 Introduction

Medicare, the U.S. federal health insurance program for seniors, provides a standard level of coverage for hospital and medical expenses with several well-known gaps (Moon et al., 2000). Seeking to improve the welfare of seniors by expanding the available choices to fill these gaps, Congress introduced a partial-privatization program known as Medicare Advantage (MA). Under this policy, the government offers subsidies to private insurers, who then offer Medicare-replacement plans to seniors during an open enrollment period each year. These plans are required to cover the same services as Medicare covers, but may offer some additional benefits, such as coverage for dental or vision services.[1] To help pay for these benefits, firms may charge a premium on top of the subsidy they receive from the government. Today, over 30% of eligible seniors are enrolled in a MA plan.

Proponents of the program point to the prevalence of supplemental benefits as proof that the program is welfare enhancing. Additionally, private insurers may face lower costs than the government for traditional Medicare benefits through superior negotiation with local providers and "managed-care" restrictions, such as referral and network requirements.[2] Detractors have voiced concerns about the high administrative costs of the program and the relative value of the supplemental benefits offered to consumers (Pear and Bogdanich, 2003).

In this work, I study the overall welfare of the MA program taking into account the subsidies paid to firms by the government, the supplemental benefits offered, firm costs, and traditional Medicare expenditures. Currently, the subsidies are set at a level higher than Medicare's average costs. At the same time, almost all plans provide additional benefits beyond those of traditional Medicare. Therefore, to compare the programs, I must put them on an equal footing and estimate both the cost to firms of providing Medicare-equivalent benefits as well as the consumer welfare generated by the supplemental benefits.

I follow the literature[3] and employ a revealed preference approach: I construct a model of supply and demand for Medicare Advantage and estimate the model's parameters using panel data on consumers and plans from 2008-2010. I use detailed information on traditional Medicare expenditures to compare the costs of MA plans to Medicare's cost. In contrast to previous work on Medicare Advantage, I allow for the presence of switching costs on the demand side to impact firm behavior in an imperfectly competitive environment, which requires a dynamic model. To achieve tractibility, I extend the Oblivious Equilibrium concept of Weintraub, Benkard, and Van Roy (2008) to models with switching costs.

I find that, on average, firms spend $5,293 per year to insure a healthy individual and

[1] The Medigap program offers insurance that solely supplements traditional Medicare. In contrast, MA plans replace traditional Medicare benefits and also provide supplemental benefits.

[2] Medicare is organized as a Fee-For-Service system, whereas most Medicare Advantage plans are offered by Preferred Provider Organziations or Health Maintenance Organizations. The Medicare Advantage system is largely based on the idea of managed competition promulgated by Enthoven (1978).

[3] See, for example, Curto, Einav, Levin, and Bhattacharya (2014), Lustig (2011)

$14,609 to insure an unhealthy individual when they provide Medicare-equivalent coverage. In contrast, Medicare spends an average of $4,390 on healthy individuals and $11,453 on unhealthy individuals. While these costs are higher than Medicare's costs, they are below the average subsidy rate ($5,826 for healthy individuals and $16,419 for unhealthy individuals), allowing firms to capture significant profits despite the presence of many competitors. On average, firms spend $184 on supplemental benefits, generating an average of $328 in welfare for individuals enrolled in their plans.

My findings are driven by the interactions between several relevant market features, including large switching costs. Miller, Petrin, Town, and Chernew (2014, hereafter MPTC) use a comprehensive panel survey of Medicare recipients along with detailed information on plan features to estimate the demand for MA plans and find seniors who wish to switch plans incur a cost of up to $1,300, well above the mean annual premium of $825. These switching costs are largely driven by network restrictions: switching insurance providers often involves establishing relationships with an entirely new set of care-givers. The existence of switching costs implies that a firm's current customer base is a large determinant of its future profits (Farrell and Klemperer, 2007). Firms therefore face an intertemporal tradeoff between "locking in" customers with attractive but relatively unprofitable plans and "harvesting" their market share with unattractive but more profitable offerings.[4]

In addition to switching costs, the MA market is characterized by a high degree of heterogeneity in health among seniors, which, when combined with a requirement that firms offer the same products at the same prices to all seniors, leads to adverse selection: firms have an incentive to enroll the most profitable individuals through careful management of plan prices and benefits. To address this incentive, the Medicare Advantage system adjusts the payments offered to firms based upon the characteristics of each of their enrollees. These adjustments, however, are based upon Medicare's expenditures on individuals with similar characteristics. If firm costs differ from Medicare's costs unevenly, firms will still have an incentive to enroll particular groups of individuals.[5]

Finally, unlike the highly regulated Medigap program, MA firms may offer a wide variety of supplemental benefits on top of the minimum "package" required by law. Firms may set a number of parameters for their plans, ranging from the number of days patients are allowed to spend in the hospital per year to the copay for primary care visits. Most firms offer more than one plan with a range of prices and characteristics to appeal to individuals with different preferences. Since these plans generally use the same network, firms must

[4]By default, seniors are automatically re-enrolled in their current plan if they take no action during the open enrollment period. While there are restrictions on the changes firms can make to their set of offered plans each year, they essentially amount to a requirement that the rank-ordering of plans with respect to price stays constant. In other words, a firm can raise prices on its enrollees from year-to-year and those enrollees must take action in order to avoid these price increases.

[5]For example, suppose there are two health conditions in the market. Condition A costs firms $200 to treat and condition B costs firms $500 to treat. The risk-adjusted payments to firms for individuals with conditions A and B are $250 and $700, respectively. Since condition B is more profitable for firms, they have an incentive to use whatever information they have about demand to attract individuals with condition B.

worry about potential self-cannibalization: if a lower priced plan is too generous, consumers may switch.

My model incorporates these features into a dynamic environment with imperfect competition. In the model, consumers are described by their health status and a number of demographic characteristics. They face a discrete choice between a number of different plans defined by a price and a generosity index. Choices that require consumers to change insurance providers incur a switching cost. Firms choose product characteristics for multiple plans and take into account the dynamic incentives generated by the switching costs. Firms face marginal costs and subsidies that vary according to the demographics and health of their enrollees.

I estimate the model in two steps. First, I estimate the demand-side parameters largely following the procedure outlined in MPTC. which in turn is based on the discrete choice literature (Berry, 1994, Berry, Levinsohn, and Pakes, 1995, Goolsbee and Petrin, 2004). In the second stage, I form moments from the model's predictions of optimal firm behavior and compare them to their sample equivalents to estimate the parameters of the firms' cost function (Newey and McFadden, 1994). I take advantage of the fact that the government pays firms different subsidies for similar individuals in different counties to assist in identifying county-level parameters.

My cost estimates are driven by the combination of factors, including the price elasticity of demand and the size of the switching costs. In some sense, the existence of switching costs creates a group of consumers over which the firm faces far less competition. In the preferred specification, these factors combine with the intertemporal dynamics to determine the markup the firm can set. The preferred estimates differ significantly from those obtained with a static version of my model. If switching costs are ignored on the demand side, the estimated costs are up to 20% lower than the preferred specification. This result is largely driven by the low enrollment rate: a majority of seniors in my sample never enroll in a Medicare Advantage plan. The absence of switching costs biases the rest of the demand system and generates a much lower mean price elasticity. This in turn implies the premiums I observe are marked up well above firm costs.

On the other hand, if switching costs are included but firms are assumed to be myopic, the estimated costs are up to 25% higher. This is driven by the connection between market share and firm behavior. When firms are myopic, they will freely charge high prices without regard to what effects those prices will have on their future market power. The inclusion of the dynamic incentive dampens that connection.

In a counterfactual simulation, I explore the impact of a 50% reduction in the switching costs incurred by consumers. This increased "fluidity" in the demand system implies that firms with a low market share don't need to offer plans that are as attractive as the baseline to win customers. In equilbrium, firms with larger shares respond to this change by lowering the attractiveness of their plans as well – as the competitive pressure they face has weakened.

As the plans in the market become less attractive, consumers respond by leaving – the overall Medicare Advantage enrollment rate is almost cut in half. The exodus is biased: the consumers that remain in Medicare Advantage are less healthy than those in the baseline scenario, which leads to higher mean per-enrollee profits for firms. Given the relatively low value consumers place on the supplemental benefits offered by MA plans, the change increases overall welfare by 8%, driven by the movement of consumers from expensive MA plans to cheaper Medicare Fee-for-Service (FFS) benefits.

A second counterfactual explores the effect of a 5% reduction in the subsidies offered to firms, similar to the policy included in the Affordable Care Act (Kaiser Family Foundation, 2014). The change in subsidies implies that zero-premium plans are no longer profitable for firms. Firms stop offering these plans and reduce the attractiveness of their remaining plans to compensate for the reduced subsidies. These effects combine to produce a drastic reduction in MA enrollment, consumer welfare and total firm profits. These reductions, however, are more than offset by a large reduction in total Medicare expenditures.

I simplify firms' information and action sets to achieve computational tractibility. Instead of explicitly tracking the full states and actions of each of their competitors, firms keep track of the average 'competitive environment' they face. This approach is similar to the Oblivious Equilibrium concept of Weintraub, Benkard, and Van Roy (2008); though I allow for a continuous state space and firms in the model keep track of detailed information on the effect of competitors' actions on individuals across the demographic spectrum.[6] Instead of modelling each feature of plans, I follow Lustig (2011) and use a generosity index to measure the relative desirability of different combinations of benefits.

This paper contributes to the literature on private Medicare plans and competition within health care, an area reviewed by Gaynor and Town (2011). Curto et al. (2014) also estimate costs for MA firms. They estimate a logit demand model and firm costs under the assumption that the subsidy system perfectly captures the relative heterogeneous risk of enrollees. Duggan, Starc, and Vabson (2014) study the effect of changes in the subsidies offered to firms in particular metropolitan areas. Aizawa and Kim (2013) study the effect of advertising on demand for MA plans in a static setting, taking into account the risk-adjustment system.

The concept of switching costs I use is broadly related to the literature on consumer inertia in health. Ho, Hogan, and Morgan (2014) study consumer inertia in the Medicare Part D (prescription drug) market and calculate a static counterfactual environment in which inertia is removed. They find, as I do, the elimination of interia would lead to substantial savings for both consumers and the government. Abaluck and Gruber (2013) also study Medicare Part D and focus on decomposing observed consumer inertia into demand- and supply-side factors. They find little improvement in the ability of consumers to choose plans

[6] Qi (2013) uses an alternative modification of the Oblivious Equilibrium concept to estimate a model of dynamic oligopoly in cigarette advertising.

over time. Handel (2013) studies the interaction between inertia and selection in employer-provided insurance and finds that a reduction in inertia leads to increased selection. Cebul, Rebitzer, Taylor, and Votruba (2011) study search frictions in the commercial insurance market and find that frictions increase premiums and insurance turnover.

My results on the relative profitability of different types of individuals are similar to those of Brown et al. (2011), who use data on Medicare expenditures to understand the changes in incentives brought about by the introduction of the risk adjustment system used by the government to compute Medicare Advantage payments. They find the risk adjustment system significantly increased the profitibility of unhealthy people and led firms to change their selection patterns. Newhouse et al. (2014) futher examines the current behavior of Medicare Advantage plans and finds evidence of selection.

Finally, I contribute to the empirical study of dynamic firm behavior in environments with endogenous product characteristics, an area reviewed by Crawford (2012). My estimation procedure allows firms to simultaneously choose characteristics of a "line-up" of plans each period. Many recent studies of firm behavior[7] use two-step estimators in which policy and transition functions are estimated semi- or non-parametrically and structural parameters are recovered separately;[8] however, given the number of covariates I include, my sample is too small to reliably estimate the transition matrix. Instead, I compute the information sets of firms directly from the data and solve a single firm problem.

The remainder of the paper is structured as follows. Section 2 provides a brief history and description of the Medicare Advantage program. Section 3 introduces my model of dynamic competition. Section 4 describes my data. Section 5 details the empirical implementation of my model including details of estimation, counterfactual, and computation. Results are described in section 6. Section 7 concludes.

2 Medicare Advantage

Medicare was created in 1965 to address the lack of health insurance among senior citizens.[9] While the original law provided basic hospital (Part B) and medical (Part B) insurance to seniors (age 65 or older), reforms have slowly expanded Medicare's role within the U.S. health care system. Eligibility was extended to individuals under 65 with certain disabilities and illnesses. The range of services covered by the program increased. These expansions had a serious impact on the cost of the program. In 1970, Medicare composed about 0.5% of GDP. By 1980, Medicare had more than doubled in size to 1.1% of GDP.[10]

[7]For examples, see Ryan (2012), Collard-Wexler (2013), Youle (2014)

[8]For examples, see Hotz and Miller (1993), Bajari, Benkard, and Levin (2007), Aguirregabiria and Mira (2007), Pakes, Ostrovsky, and Berry (2007a)

[9]Much of the historical information in this section is compiled from `http://www.cms.gov/`

[10]See `http://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/NationalHealthExpendData/NationalHealthAccountsHistorical.html` for information on health care spending in the U.S.

This growth led policy-makers to begin experimenting with different cost-containment and care-delivery strategies in the 1980s. While many efforts focused on broad reforms, such as changing the way Medicare reimbursed care-givers, the Centers for Medicare and Medicaid Services (CMS, then known as the Health Care Financing Administration) began a series of limited trial programs based in part on the ideas of Enthoven (1978) in which the government contracted with Health Maintenance Organizations (HMOs) to manage the care of select groups of enrollees.

HMOs, which had become popular after the passage of the Health Maintenance Organization Act of 1973, provided health care to their customers under a fundamentally different model. Previously, most health insurance was operated under a Fee-for-Service (FFS) model, in which doctors charged patients and insurers for each individual service performed. Given the level of asymmetric information present in the doctor-patient relationship, many feared the system made it too easy for doctors to perform unnecessary procedures (Arrow, 1963, Chernew, 2003). HMOs changed that by signing pre-paid contracts with physicians and hospitals (Markovich, 2003).

Today, HMOs take advantage of a number of other components of the so-called managed care model (Glied, 2000). Patients are generally required to see a primary care physician for a referral before they can visit a more expensive specialist. Preventative care is usually provided at little-to-no charge to enrollees with the idea that regular checkups can detect illnesses before expensive procedures are necessary. HMOs regularly review the performance of their doctors to ensure the number of services they each perform are in line with expectations. Preferred Provider Organizations (PPOs) have arisen as a slightly more flexible but more expensive alternative (Gabel et al., 1988).

The Balanced Budget Act of 1997 expanded and formalized the Medicare managed care program into Part C. The new program closely followed common models of employer-provided health insurance and had several key components. Each spring, firms submitted county-by-county plan proposals to CMS. CMS verified that the plans met the minimum requirements and covered the same conditions as Medicare, though many plans chose to offer additional benefits. In the fall, CMS operated an open-enrollment period, during which Medicare recipients could freely choose between original Medicare or any of the plans available in their county of residence. Coverage began on January 1st, and firms received a flat subsidy from CMS, known as a capitation payment, each month (CMS, 2014).

After several enrollment periods passed, policy-makers grew concerned about the direction of the program (Pear and Bogdanich, 2003). Entry by firms was largely limited to suburban areas and many rural and inner-city residents did not have access to the new plans. Additionally, since the capitation payment was the same for all enrollees, firms had an incentive to tailor their plans to appeal to only the healthiest (and therefore most profitable) consumers. Enrollment hovered around 15% of eligible seniors, or six million people (Kaiser Family Foundation, 2014).

The Medicare Prescription Drug, Improvement, and Modernization Act of 2003 sought to address these concerns by reforming Part C. Plan providers were given new flexibility to manage the care of their enrollees – particularly with respect to the provision of non-emergency care. The subsidies offered to plans were significantly increased to encourage entry in more geographic areas. Finally, the reimbursement system became risk-adjusted. Under the new system, firms submit demographic and diagnostic information about their enrollees each month. CMS "scores" each enrollee's risk according to the cost of similar individuals enrolled in the traditional Medicare system – the average senior has a risk score of 1.0. CMS sets a benchmark rate for each county and multiplies this rate by the individual risk score for each enrollee to determine the subsidies paid to firms (CMS, 2014).

The changes have had the desired effects: now almost all seniors have the option of at least one Medicare Advantage plan and most can choose between two or more (in addition to traditional Medicare). Studies of the risk adjustment system have concluded that it effectively reduced the tendency of firms to prefer healthy enrollees (Brown et al., 2011). Enrollment surged and today over 30% of seniors, almost 16 million people, are enrolled in a Medicare Advantage plan (Kaiser Family Foundation, 2014).

3 Model

Taking these institutional details into account, I build a model of supply and demand for Medicare Advantage. On the demand side, heterogeneous consumers face a discrete choice of plans described by a price and a generosity index. They incur a cost if they switch insurers. Following Handel (2013), I assume consumers are myopic and do not take into account future switching costs.[11] On the supply side, symmetric firms simultaneously choose the price and generosity of multiple plans. They solve a recursive value function that takes into account the dynamic tradeoff generated by the switching costs.

I simplify the problem to ensure that my estimation exercise is tractible while maintaining a degree of flexibility in my specifications of utility and firm costs. In particular, I allow firms to keep track of their market shares of healthy and unhealthy individuals (as opposed to allowing firms to keep track of the full, joint distribution of enrollee characteristics). Firms calculate their per-enrollee profits from a distribution over consumer characteristics conditioned on their health status. This allows firms to keep track of the dimension most correlated with cost while keeping a relatively low-dimension state space. while firms face a recursive problem – which allows firms to alter their markups based upon the switching costs faced by consumers – entry and exit is assumed to be exogenous.[12] While firms

[11]In essence, allowing this would require consumers to forecast future prices – meaning consumers would have to keep track of who will purchase which plans – and their own future health statuses over a potentially lengthy time-frame Handel (2013).

[12]In my modification of the Oblivious Equilibrium concept (explained in detail below), firms keep track of the competition they face in terms of the utility of competing plans. In some sense, firms don't care who those products are offered by (or how many competitors there are) – just the level of competition in the

store summarized market information – which allows firm behavior across markets to differ based on the degree of competition – they do not keep track of the details of each of their competitors.

3.1 Environment

Time is discrete and denoted by t; each period represents a year. The world is divided into a number of discrete markets (representing individual U.S. counties) denoted by m. Each market contains N_m consumers, with individual consumers denoted by i. Each market also has a vector of observables Y_m, including the benchmark subsidy rate B_m.

In each period, for every market, there is an independent set of firms F_m^t, determined exogenously. Individual firms are denoted by f, and these firms offer plans in a set J_f denoted by j, where the size of the set is determined exogenously. Each plan j consists of a premium p_j and a generosity index g_j. Firms enter each period with a market share of healthy people s_{fh} and a market share of unhealthy people s_{fu}. In all markets, the outside good, good 0, is Medicare, and the set of all plans within a market is denoted J_m.[13]

3.2 Timing

In each period t, actions are taken as follows:

1. Firms observe their information set I_f^t, described in detail below.

2. Firms simultaneously choose actions $\sigma_f = \{p_j, g_j\}_j$.

3. Consumers choose plans from J_m.

4. Profits and the next state of the market are realized.

3.3 Consumers

Individual i has a vector of demographic characteristics Z_i which includes age, gender, health status, education, race, and current plan enrollment. Z_i has distribution C_m. Additionally, each individual is assigned to one of several income groups, represented by the dummy variables d_{wi}.

market.

[13]There are two other substantial components of the post-65 insurnace system: Private plans provided by employers as part of a retirement package or pension system, and Medigap, which provides supplementary insurnace on top of traditional Medicare benefits. As individuals with employer-provided plans have very low Medicare Advantage enrollment rates, I treat the employer-based system as a separate entity and remove those individuals from the market. Since Medigap plans are highly regulated and consistent across geographies, I abstract from the variance in those plans and consider them part of the outside good.

Consumer i considering plan j faces the following choice-specific utility function:

$$
\begin{aligned}
U_{ij} =& \alpha_0 p_j + \sum_w \alpha_w p_j d_{wi} + SW_k * 1\{\text{switch}_{ij}\} + \beta_z Z_i + \beta_g g_j \\
&+ \beta_{zg} Z_i g_j + \xi_j + \epsilon_{ij}
\end{aligned} \tag{1}
$$

In this equation, $\alpha = \{\alpha, \alpha_w\}$ represents income-specific price sensitivity. β_z captures tastes for the inside good that vary by demographic characteristics. β_g captures the mean taste for generosity and β_{zg} captures demographic specific tastes for generosity. $SW_k * 1\{\text{switch}_{ij}\}$ represents the cost SW_r that consumer i must pay if j is offered by a different firm than its current plan. SW_k is allowed to vary across switch types: switching between traditional Medicare and Medicare Advantage may incur a different cost than switching between different Medicare Advantage providers. Finally, ξ_j represents the component of plan quality that is unobserved to the econometrician, and ϵ_{ij} represents the individual choice-specific unobservable, which is assumed to be independently drawn according to a Type-I extreme value distribution.

Following Berry et al. (1995), I can decompose the utility obtained from good j into a mean:

$$
\delta_j = \alpha_0 p_j + \beta_g g_j + \xi_j
$$

and an individual specific deviation:

$$
\mu_{ij} = \sum_w \alpha_w p_j d_{wi} + SW * 1\{\text{switch}_{ij}\} + \beta_{zg} Z_i g_j + \epsilon_{ij}
$$

Given a set of plans J_f, the probability that consumer i chooses plan j (often known as the share function) can now be written as:

$$
Pr(i \text{ chooses } j) = s_{ij} = \frac{\exp(\delta_j + \mu_{ij})}{1 + \sum_{j' \in J_m} \exp(\delta_{j'} + \mu_{ij'})} \tag{2}
$$

3.4 Firms

As mentioned previously, firms choose prices p_j and generosity g_j for some number of plans J. The firm's problem requires the firm to evaluate the expected profits of different combinations of prices and generosities as a function of the information they have about the market. I develop this problem in stages, starting with the firm's per-enrollee cost function. Computational restrictions prevent me from considering a full-information model. I therefore adapt the Oblivious Equilibrium concept of Weintraub et al. (2008, 2010) to this environment and define the information set of firms based on what they need to calculate demand and cost and, thus, profits. The translation is made more complicated by the inclusion of switching costs on the demand side and therefore the need for firms to track their market shares.

With the information set established, I describe the market timing and write down the firm's problem as a recursive value function.

3.4.1 Cost function

In order to evaluate the outcomes of various actions, the firm must first evaluate the cost of providing insurance. The per-capita expected cost of providing plan j to individual i is a function of the generosity index g_j and is conditioned on the demographic characteristics Z_i and market characteristics Y_m, such as the population density, the doctor/population ratio, and the average per-capita income:

$$c_{ijm}(g_j|Z_i, Y_m) = (\gamma + \gamma_m Y_m) + \gamma_z Z_i + (\gamma_g + \gamma_{gm} Y_m) g_j + \gamma_{gz} Z_i g_j + \gamma_{g2} g^2 + \zeta_f \tag{3}$$

In this equation, γ represents the average marginal cost of providing Medicare-level benefits and γ_m represents deviations due to market-specific factors. γ_z represents the deviations in marginal cost due to consumer demographics (e.g. individuals in poor health cost more to insure). Similarly, γ_g is the mean marginal cost of generosity (with γ_{gm} capturing market-specific factors) and γ_{gz} captures the deviations from that mean due to consumer demographics. I allow for the possibility of a quadratic component to the cost of generosity.[14] Finally, firms receive an i.i.d. cost shock ζ_f each period.

3.4.2 Information set

The contents of the information set of firms is driven by the need (of the firms) to calculate the expected profits for a given action and therefore the need to calculate demand and costs. The demand for a particular plan j can be calculated by integrating out the share function s_{ij} over consumers:

$$s_j(p_j, g_j) = \int_i s_{ij}(p_j, g_j, Z_i) dC(Z_i) di$$

Since costs may be particularly correlated with health, firms can individually calculate the fraction of healthy and unhealthy people who enroll in the plan:

[14] Intuitively, if a more generous plan makes individual claims more attractive to the consumer, the consumer will make additional claims, each of which may cost the firm more money. This quadratic component provides additional curvature in the firm's value function which is particularly helpful during the estimation procedure.

$$s_{jh} = \int_i s_{ij}(p_j, g_j, Z_i) dC(Z_i|h) di$$
$$s_{ju} = \int_i s_{ij}(p_j, g_j, Z_i) dC(Z_i|u) di$$

These plan-level shares can be aggregated into firm-level shares: $s_f = \sum_j s_j; s_{fh} = \sum_j s_{jh}; s_{fu} = \sum_j s_{ju}$. For simplicity, I drop s_{fu} and consider only s_{fh} – the same calcuations are made for s_{fu}.

To compute the numerator of equation 2, the firm must know how its products map into the utility obtained by consumers $\delta_j + \mu_{ij}$. Since consumers who are currently enrolled in other MA firms or original Medicare face switching costs when considering the firm's products, the firm must know its own shares s_{fh} and how many consumers are enrolled in traditional Medicare versus other MA plans. To calculate the denominator of equation 2, the firm must know something about the other products available in the market.

The classic dynamic oliogopoly approach would be to allow firms to observe the full state space and action set of their competitors in the manner of Ericson and Pakes (1995). Equilibrium would be defined as a set of policy functions that obtained the supremum of the recursive value function. As Medicare Advantage markets often have more than 10 incumbents, computational limitations prevent me from calculating optimal firm actions as a function of all of their competitors' states and actions.[15] To reduce the computational burden, I adopt the Oblivious Equilibrium solution concept developed by Weintraub, Benkard, and Van Roy (2008, 2010). However, there are key differences between their model and mine that I must address.

In the Weintraub, Benkard, and Van Roy (2010) version of the Ericson and Pakes (1995) dynamic oligopoly model, consumers are identical and the state space of firms is discrete, representing (depending on interpretation) their level of capital or efficiency. Weintraub et al. simplify the information set of their firms by defining a "long-run average state" vector $\bar{s}$, which is the expected number of firms (which may be fractional) with each state n at any given period in equilibrium. The components of $\bar{s}$ can be calculated as:

$$\bar{s}_n = E_t \left[\sum_f 1\{f \text{ in state } n\} \right]$$

Weintraub et al. write down their firm's problem as a function of the firm's own state, and condition on this long-run average state. In their logit application, they use this state to calculate the share s_j the firm receives as a function of the price the firm charges for its good. Recall that, for the generic logit model:

[15] Even with recent improvements in solution techniques for Markov-Perfect Nash Equilibria, such as the stochastic algorithm introduced by Pakes and McGuire (2001), the number of unique states visited by the market is simply too large.

$$s_j = \frac{\exp(\delta_j)}{1 + \sum_{j'} \exp(\delta_{j'})}$$

Since firms in their model are symmetric, which implies that each firm with the same state n will offer a good with the same $\delta_j = \delta(n)$, they can rewrite their share function using the components $\bar{s}_n$ of $\bar{s}$:[16]

$$s_j = \frac{\exp(\delta_j)}{1 + \sum_n \bar{s}_n \exp(\delta(n))} \tag{4}$$

Unfortunately, consumers in my model are heterogenous, and have a utility composed of $\delta_j + \mu_{ij}$. However, a modification addresses the problem. Instead of taking the expected state of firms, I take the expected impact (where the expectation is over time) of the products they offer on the share equation for i:

$$\bar{q}_i = E_t \left[\sum_j \exp(\delta_j + \mu_{ij}) \right] \tag{5}$$

In other words, given some belief about the future actions of competitors, the firm constructs $\bar{q}_i$ for each individual by summing their exponentiated utility terms for each product in every period in the future and then taking the average over all periods.[17] In the traditional OE setup, these beliefs come from the optimal strategies of firms in different states. In my estimation exercise, I form this expectation by taking the observed actions of firms in my sample period. In my counterfactual, I follow Weintraub et al. (2010) and solve for a self-consistent set of firm strategies and beliefs.

As this expression is essentially the denominator of the traditional logit share function, the share function used by the firm becomes:

$$s_{ij} = \frac{\exp(\delta_j + \mu_{ij})}{1 + \bar{q}_i} \tag{6}$$

In practice, I calculate integrals over the distribution of individuals by taking a number of discrete draws from the conditional dC distributions. I form a vector $\bar{Q}$ by calculating q_i for each draw:

$$\bar{Q} = \{\bar{q}_1, \bar{q}_2, \cdots \bar{q}_n\} \tag{7}$$

The interpretation of $\bar{q}_i$ is slightly different from $\bar{s}$. $\bar{q}_i$ is essentially the denominator of the share function and in some sense measures how attractive of a product firm j must

[16]For ease of exposition, in the following discussion I abstract away from the effect that firm j has on the denominator of the share function. My implementation of the information set described in this section involves a small modification to $\bar{q}_i$ to remove the effect of one firm and an opposite adjustment to s_{ij} to explicitly include the effect of product j.

[17]This expression plays the same role in my model as equation (4.1) of (Weintraub et al., 2008, p. 1386) does in the traditional OE setup.

offer (through δ_j) in order to achieve market share. As $\bar{q}_i$ increases, the firm must somehow increase δ_j in order to achieve the same share. Therefore, I call $\bar{Q}$ the *expected competitive pressure* of the market.[18]

Finally, as mentioned previously, the response of consumers to a product described by p_j and g_j depends upon their current enrollment since they may be subject to a switching cost if they choose plan j. The firm does not need to know where each consumer is – merely whether or not they are enrolled in original Medicare, a competitor's MA plan, or in one of the firm's own plans.[19] The firm keeps track of it's own shares by health status s_{fh}. Additionally, the firm knows the average share of individuals enrolled in any MA plan $\bar{s}_h = E_t[\sum_f s_{fh}]$. With these numbers, the firm can calculate the number of individuals enrolled in a different MA plan and the number of individuals enrolled in original Medicare. If $e \in \{0, 1, 2\}$ represents enrollment in original Medicare, a competitor's plan, and one of firm's own plans respectively, the firm can calculate $S_{e,h}$ as:[20]

$$S_{2,h} = s_{fh}; S_{1,h} = \bar{s}_h - s_{fh}; S_{0,h} = 1 - S_{1,h} - S_{2,h}$$

The firm can therefore calculate it's share function with the modified s_{ij} of equation 6 as

$$s_j = \sum_{e,h} S_{e,h} \int_i s_{ij} dC(Z_i|e,h) di \tag{8}$$

The firm knows $\bar{Q}$, the conditional distributions of consumers $dC(Z_i|e,h)$, which is invariant over time, and the average share of individuals enrolled in MA plans $\bar{s}_h$. Each period, the firm observes an information set that includes:

$$I_f = \left\{ s^t_{fh}, \zeta^t_f \right\} \tag{9}$$

3.4.3 Profit function

The firm forms expected plan-level profits by integrating over the conditional distributions of consumers. Plan-level profits are a function of the plan characteristics under consideration by the firm, and are conditioned on the information set of the firm. Firms receive a risk-adjusted subsidy for each consumer based upon the market-level benchmark and the consumer's individual risk score r_i, where $r_i = f(Z_i)$ is a function of the consumer's demographic characteristics:

[18] Calculating $\bar{Q}$ in this way also addresses the fact that the state space in my model is continuous.

[19] Since there is no switching cost between plans within a firm, the demand will be the same across individuals enrolled in all of a single firms' plans and therefore firms only need to keep track of shares at the firm level.

[20] It is possible that $s_{fh} > \bar{s}_h$, in which case $S_{1,h}$ will be negative. While this situation does not occur in the data, I must solve a value function across the entire state space, and the continuation values obtained for states in which this occurs will affect the value function in all other states. In the results presented below, I bound $S_{1,h}$ from below by 0 and adjust $S_{0,h}$ accordingly. Other methods of handling this situation, such as allowing negatives, do not substantially change the results.

$$\pi_j(p_j, g_j|I_f) = N_m \sum_{e,h} S_{e,h} \int_i (p_j + Br_i - c_{ij})\, s_{ij}\, dC(Z_i|e,h)di \tag{10}$$

In other words, the firm calculates plan-level profits by considering six different groups of individuals: healthy and unhealthy people who are enrolled in traditional Medicare, a competitor's plan, or one of the firm's own plans. For each group of people, the firm knows the distribution of their demographics, conditional on belonging to that particular group, as well as the size of that group in the market. The firm uses the $\bar{Q}$ vector to calculate s_{ij} for each individual across the conditional distribution. The firm also uses the the conditional distribution to calculate the cost incurred and subsidy obtained by insuring a particular individual.

3.4.4 Firm's problem

With these ingredients in hand, I can formulate a recursive value function for firms. Their value is a function of their current share of consumers across health statuses and is conditioned on their information set:

$$V(I_{fm}) = \max_{\sigma_f} \sum_j \pi_j(p_j, g_j|I_{fm}) + \beta E_{I'}\left[V\left(I'_{fm}(\sigma_f)\right)\right] \tag{11}$$

In this recursive problem, β is the discount factor, which is constant across firms. Dynamics are embedded in the evolution of I_{fm}. In particular, I_{fm} contains the firm's market shares and cost shock as well as the proportion of individuals enrolled in the Medicare Advantage system.

3.5 Equilibrium

I adopt the Oblivious Equilibrium notion of Weintraub, Benkard, and Van Roy (2008) to this case. An equilibrium in my model is an integral number of firms F, a firm strategy σ and competitive pressure vector $\bar{Q}$ such that:

1. Given $\bar{Q}$, σ is the solution to equation 11.
2. When F firms play according to σ, $\bar{Q}$ satisfies equation 7.

4 Data

MPTC construct a comprehensive dataset of the MA market to estimate a detailed demand system. I employ this dataset with a few modifications to align with my model of firm behavior. Broadly speaking, my data falls into three categories: plans, consumers, and geographies. In this section, I briefly describe each of these categories with a focus on the differences between the data used in this paper and the data used in MPTC.

4.1 Sample selection

I restrict the temporal and geographic spread of my sample due to the needs of my equilibrium notion and estimation procedure. I employ a variation of the Oblivious Equilibrium concept of Weintraub, Benkard, and Van Roy (2008) which imposes a notion of stationarity.[21] Additionally, my estimation procedure requires demographic-specific measures of market share. To satisfy these requirements, I select 39 markets throughout the United States where I observe a reasonable sample of individuals and do not observe significant changes in the total number of firms and plans present in the period I study. Additionally, I restrict my attention to plans offered starting January 1, 2008, two years after the implementation of the current Medicare Advantage system.

4.2 Consumers

My data on individual consumerss comes from the Medicare Current Beneficiary Survey (MCBS), an overlapping-panel survey of a nationally representative sample of Medicare recipients sponsored by the Centers for Medicare and Medicaid Services (CMS) and produced by Westat. I use the Cost and Use files from the 2007-2010 data releases to obtain data on individual plan choices[22] and demographic characteristics including age, race, education, and income. I use a self-reported health status variable to separate individuals into healthy and unhealthy categories.[23]

Summary statistics on the individuals used to produce my estimates are in table 1. Despite selecting on the 39 counties with the highest sample available, individuals in my subset are similar to the fuller sample of MPTC, if very slightly richer (the average income in the full sample is \$43,378 as opposed to \$46,198 in my sample). The second column reports the standard deviation of the yearly means for each of the demographic characteristics I include. In particular, the demographic distribution does not shift much between periods. Though the nationwide Medicare Advantage enrollment rates increased significantly throughout the period, enrollment rates are relatively stable in the counties I consider.

The Cost and Use files also contain information on Medicare payments to service providers for individuals enrolled in traditional Medicare. This data is constructed from a combination of patient interviews and administrative records.[24] Table 2 summarizes Medi-

[21] In particular, I must assume that the data I observe is drawn from the ergodic set of market states generated by the OE strategies.

[22] The MCBS does not track the specific plan number chosen by the individual. Instead, it reports the contract number and a number of variables (constructed from survey responses and administrative records) related to the benefits offered by the plan. Following MPTC, I rank all of the plans offered under the recorded contract by their closeness to the plan described by the survey participant and assume the true choice is the closest match.

[23] The question asks responders to rate their own health as "excellent," "very good," "good", "fair," or "poor." I group the first three responses into "healthy" and the last two (as well as any non-responses or refusals) into "unhealthy." Nyman et al. (2007) use a similar question in the Medical Expenditure Panel Survey to develop quality-of-life measures across the US population for cost-utility analyses.

[24] For more on the construction of this file, see Eppig and Chulis (1997).

care expenses for different groups of individuals. I use this information to construct the risk-adjusted subsidies CMS pays to firms for various individuals across the demographic distribution.[25]

4.3 Plan data

The Centers for Medicare & Medicaid Services maintains a public database of the characteristics of all MA plans offered each year. This database includes detailed information on plan costs, benefits, and options, at the contract-plan-segment level. For each plan, I extract the price, coverage area, and a number of plan characteristics, including copays for doctor and hospital visits, as well as flags for drug coverage, dental coverage, and vision coverage.

MPTC estimate the relevative preferences of individuals for these different plan features. I use their estimates to transform the multi-dimension plan characteristics into a single generosity index g_j, for each plan j.[26] Since MA plan providers are required to offer coverage for each service that Medicare covers, I define $g_j = 0$ to be Medicare-level coverage and linearly transform the generosity index to ensure non-negativity. Examples of plans with different generosities can be found in table 5.

Table 3 contains summary information on plans. Notably, 81% of firms offer more than 1 plan, and 73% of those have at least one plan with zero annual premium. Indeed, 43% of all plans offered across markets do not charge a premium. I use these facts to simplify the firm's problem in my estimation routine: I fix the number of plans per firm at 2 and restrict the first plan's premium to zero. For firms with more than two plans, I weight the attributes of the plan by the plan's individual share. Figure 1 shows the results of this weighted average for plans with positive prices.

Table 4 shows the average number of firms and plans for each of the years I consider. While there is some variance in firms from year to year, the yearly means do not differ substantially from the overall mean of about 15 firms. There is a drop in plans between 2009 (21.8 plans on average) and 2010 (16.8 plans). The biggest single change was in Franklin County, Ohio, (in which Columbus, Ohio is located) which went from 28 plans in 2009 to 17 plans in 2010.

4.4 Geographies

I obtain data on individual markets (which are defined as counties) from the Area Health Resources Files maintained by the Health Resources and Services Administration within the US Department of Health and Human Services. The files combine and summarize data

[25] The true risk adjustment formula uses ICD-9 diagnostic codes (in addition to other demographic variables) to determine payments (CMS, 2014). Though the MCBS asks a number of questions about diagnoses and illnesses, it does not contain ICD-9 codes and I therefore use self-reported health status as a proxy.

[26] This essentially assumes that, near the selections offered, consumers and firms view different plan attributes as perfect substitutes.

from multiple sources, including the Census Bureau and other Health and Human Services sources, into a single county-level dataset.

For each county, I extract the population, the median income, the number of practicing medical doctors, the number of hospitals, and the number of nursing homes to use as cost covariates. I additionally extract the "contiguous county" file which allows me to identify neighboring counties for construction of instruments. Finally, I extract the "benchmark" per-capita subisdy rate for MA plans. The benchmark rate is the subsidy paid to a firm for a person of average risk and varies by market according to Medicare's average costs. I use this market-specific benchmark rate along with the the average Medicare expenditures for individuals in different groups to construct individual-specific subsidy rates.

Table 6 summarizes my geographic data. Since I form market shares from the MCBS data, I restrict my analysis to those markets with the greatest MCBS sample size. This in turn means the county markets I consider are significantly larger than the average across the United States.

5 Empirical implementation

I use a multistep approach to estimate the model parameters for the demand and supply sides $\theta = \{\theta^D, \theta^S\}$, and calculate counterfactuals using a modification of the Weintraub et al. (2010) algorithm. For computational simplicity, I limit each firm to offering two plans. Since 40% of plans observed in the market have no premium, I restrict one of the firm's plans to be a zero premium plan. The firm's action therefore consists of three components: $\sigma_f = \{g_0, p_1, g_1\}$.

5.1 Estimation

5.1.1 Preliminaries

Since the policy function generated by the model is dependent on the starting shares of the firm, I must construct these shares for each firm-year observation. I use the MCBS observations to form these shares for 39 markets.[27] For each firm present in the data, I form share-weighted averages of price and generosity to match the three components of the firm's action in my model and adjust the decision of consumers accordingly.

[27] The MCBS has relatively low sample size in most of the counties it observes. Gandhi et al. (2013) show that errors in market shares caused by small sample size can bias demand estimates. MPTC avoid this issue by using monthly enrollment data from CMS. However, this data is only available at the aggregate level. As I must separate out shares by health status, I must read shares from the MCBS directly and therefore must limit myself to markets in which there is sufficient sample.

5.1.2 Demand

I estimate the demand model following the two-stage approach of MPTC, which builds on the discrete choice estimation approach of Berry (1994) and Berry et al. (1995) by adding panel data on consumers' choices. I start by re-writing the utility model as a combination of individual specific terms and product fixed effects δ_j:

$$U_{ij} = \sum_r \alpha_r P_j d_{ri} + F * 1\{\text{switch}_{ij}\} + \beta_z Z_i + \beta_{zg} Z_i g_j + \delta_j + \epsilon_{ij}$$

For a given guess of individual specific demand parameters $\theta_Z^D = \{\alpha_r, F, \beta_z, \beta_{zg}\}$, I use the Berry (1994) contraction to find the unique set of $\delta_j(\theta_Z^D)$ that match predicted shares to observed market shares.[28] Using the individual choice data from the MCBS, and its panel structure to calculate switches, I construct the likelihood function for an individual as follows, where C_i represents the choice of individual i:

$$l_i(j; \theta_Z^D) = s_{ij}^{C_i = j}$$

In the first step of the demand estimation, I maximize the likelihood function over the space of θ_Z^D.[29] At the point estimate, $\hat{\theta}_Z^D$, I store the unique $\hat{\delta}_j$. In the second step of the demand estimation, I regress these $\hat{\delta}_j$ on on observable product characteristic according to the terms in the original demand equation, where ξ_j represents product-specific unobservables:

$$\delta_j = \alpha_0 P_j + \beta + \beta_g g_j + \xi_j$$

I instrument for price using the average benchmark in the surrounding counties as well as the average generosities of competitors in the same market per Berry et al. (1995).

5.1.3 Supply

With demand estimates in hand, I construct the information set I_f for each firm. This consists of calculating the average competitive pressure observed in each market for each draw from that market's demographic distribution. As I take the number of firms in each period as exogenous, I create distinct measures of competitive pressure for each firm-year observed in the data and then average across firm-year observations.[30]

[28] As the choice set is the same for all individuals in the market, I need only construct market-level shares, as opposed to demographic-specific market shares. I therefore use CMS Enrollment files, which cover the entire population of a county, to construct these shares, thus avoiding measurement error problems within the Berry (1994) contraction.

[29] This step utilizes the Broyden-Fletcher-Goldfarb-Shanno (BFGS) maximization routine.

[30] Instead of the approach described in this paragraph and the following algorithm, I could follow the approach of Weintraub et al. (2010) and solve for a self-consistent policy function and expected competitive environment, as I do for my counterfactual simulations. Unforunately, their algorithm is not guaranteed to converge and requires significant additional computational complexity. in practice, the application of their approach increases the computational effort by an order of magnitude. The difference in policy functions

5.1.4 Conditional moment restrictions

Using notation from Pakes et al. (2007b), my model creates an approximation $R(\cdot)$ to the true profit function of firms $\pi(\cdot)$ in the following sense:

$$\frac{\partial \pi_{fj}}{\partial X_{fj}} = \frac{\partial R_{fj}}{\partial X_{fj}} + \nu_{1fj}$$

In this expression, X_{fj} represents all of the data I observe about firm f and plan j, including market-level characteristics and the demographic distributions. ν_{1fj} represents both expectational and measurement errors at the firm-plan level. Expectational errors can come from a number of sources, including incomplete information on the environment or asymmetric information on the states of other firms. As $\pi\cdot$ is the result of maximizing the profit function, any error in the profit function will contribute to the error in the policy function as well.

I impose a conditional moment restriction:

$$E[\nu_{1f}(\theta_0)|I_f] = 0\,\forall f$$

This restriction states that, conditional on their information sets, firms choose their actions optimally on average – this is equivalent to a restriction on firm behavior: $E[\sigma^{DATA} - \sigma^{MODEL}|I_f] = 0$

Since I am imposing a conditional moment restriction, there are an infinite number of possible instruments – indeed any function of any component of the information set I_f is a valid instrument. Chamberlain (1987) shows the efficient set of instruments are formed by the derivative of the moment restriction with resepect to each parameter when the derivative is evaluated at θ_0:

$$\boldsymbol{H}_f = E\left[\left.\frac{\partial \boldsymbol{\nu}_{1f}(\theta_0)}{\partial \boldsymbol{\theta}}\right| I_f\right]$$

Following Berry et al. (1999), I approximate the optimal instruments by calculating these derivatives at an initial guess of the parameters and recalculating them when I update the weight matrix during the GMM procedure.

5.1.5 Two-step GMM

I use a a two-step GMM approach. I start by setting my weight matrix to the identity $W_1 = I$ and calculate an approximation to the optimal instruments based upon an initial guess. I minimize the GMM objective function $f(\theta^S) = g'W_1 g$ to obtain an initial estimate

computed by the two approaches at the point estimates described in the subsequent section is approximately 10%.

of the parameters $\hat{\theta}_1^S$.[31] I update the weight matrix according to the sample variance-covariance matrix $W_2^{-1} = \hat{S}(\hat{\theta}_1^S) = \frac{1}{n}\sum_n gg'$. I also update the approximation to the optimal instruments.

I then minimize the modified GMM objective function $f(\theta^S) = g'W_2 g$ to obtain my final parameter estimates $\hat{\theta}$.

5.1.6 Identifcation

I assume that, in expectation, firms choose actions optimally given their information on the markets and the incentives they face from the terms of their value function. Parameters are therefore identified through their impact on the value function of firms and its derivative. The vector of first order conditions for optimality of the firm's value function can be written as follows:

$$\mathbf{0} = \sum_j \frac{\partial \pi_j}{\partial \boldsymbol{\sigma}_f} + \beta E\left[\frac{\partial V}{\partial \boldsymbol{\sigma}_f}\right]$$

This set of first order conditions illustrates the primary incentives faced by firms. Note that σ_f is a vector consisting of the features of all of f's plans. An expansion of $\frac{\partial \pi_j}{\partial \boldsymbol{\sigma}_f}$ would explicitly reveal a cross-product cannibalization issue for firms through the mechanism of s_{ij}: the number of people you attract to an individual plan is a function of the features of all the plans you offer. Since firms must offer plans at the same price and features to every consumer in the market, these inter-plan derivatives are crucial in determining the degree to which firms can price discriminate between consumers with different preferences.

The second term of these first order conditions captures the intertemporal tradeoff between today's profits and tomorrow's share – and the effect of tomorrow's share on tomorrow's profits. This term is the key difference between this model and previous efforts to model the Medicare Advantage market. Its value at various points in strategy and state space determines whether firms are engaging in "collecting" market share or "harvesting" that share (Farrell and Klemperer, 2007, Section 2.6.2). Its scale relative to the first term determines the degree to which that behavior dominates the selection and product cannibalization concerns of the firm.

My estimator identifies the different components of the cost function through intra- and inter-market variance along with the demand distribution. The model generates a policy function for all points in the firm's state space. Differences in observed behavior of firms at different points in that space within the same market identify the constant and health-specific terms in the cost function. Differences in observed behavior of firms at similar points in markets with different attributes identify market-level cost parameters. The other

[31] As analytic derivatives are unavailable and individual function evaluations are extremely expensive (see the computation subsection for more), I use the Nelder-Mead simplex search method.

demographic-specific terms in the cost function are identified by cross-market variation in the distribution of demographic characteristics.

5.2 Counterfactual

In order to evaluate the effects of alternative policies, I must solve for the equilibrium strategies of firms. My estimation approach uses the data to calculate the expected competitive environment faced by firms in each market. In a counterfactual scenario, however, the expected competitive environment is likely to be different from what I observe in the data. For any given strategy σ, I can calculate the expected information set $C(\sigma) = E[I_f|\sigma]$, and for any given information set C, I can calculate the optimal strategy $\sigma(C)$. The challenge is to find a fixed point of these simultaneous equations.[32]

I proceed by fixing the number of firms to the average number of firms observed in the market[33] and adapting the algorithm of Weintraub et al. (2010). The algorithm may be outlined as follows:[34]

1. Start by initializing $\sigma_0 = \mathbf{0} \forall s_{h,f}, \zeta_f$, $t = 0$, and $x = 1$.
2. Loop until $x < tol$:
   (a) Given σ_t, calculate the expected information set C_t by forward simulation.
   (b) Given C_t, solve the single firm problem $\max_{\sigma_f} \sum_j \pi_j(\cdot) + \beta V_t(s'(\sigma_f))$ to obtain σ_{t+1}
   (c) Set $x = ||\sigma_t - \sigma_{t+1}||$ and $t = t + 1$
   (d) Set $\sigma_{t+1} = \lambda\sigma_{t+1} + (1 - \lambda)\sigma_t$

5.3 Computation

The primary computational challenge is calculating $\sigma^{MODEL}(s_{h,f}|\theta^S, I_{fm})$ for a given set of supply-side parameters. The primary advantage of the Oblivious Equilibrium approach is the result that I must merely solve a single firm problem for each market. To do this, I discretize the $s_{h,f}, \zeta_f$ information set space and use a value function iteration algorithm:

[32] In a full-information version of the model, multiple equilibria are possible. Given the restrictions I have placed on firms' information sets, it seems likely (though not inevitable) that the number of possible equilibrium are fewer than in a more complete information environment. In practice, I use a variety of different starting points for each counterfactual and find only a single equilibrium.

[33] Weintraub et al. (2010) include entry and exit, whereas I treat it as exogenous. While my algorithm is easily modified to include entry and exit, doing so would require assumptions about entry- and exit-relevant parameters. Therefore, for consistency with my estimation exercise, I abstract from entry and exit in the results I present below.

[34] As in Weintraub et al.'s case, this algorithm is not guaranteed to converge. Indeed, the algorithm often oscillates around the fixed point without progressing toward it. My implementation detects this condition and employs a local restart strategy with a reduction in the update parameter λ to encourage convergence. In practice, this has worked well.

1. Start by initializing $V_0 = 0 \forall s_{h,f}, \zeta_f$, $t = 0$, and $x = 1$.
2. Loop until $x < tol$:
   (a) Loop over each state $s_{h,f}, \zeta_f$:
      i. Solve the single firm problem $\max_{\sigma_f} \sum_j \pi_j(\cdot) + \beta V_t(s'(\sigma_f))$ to obtain $V_{t+1}(s_{h,f}, \zeta_f)$.
   (b) Set $x = ||V_t - V_{t+1}||$ and $t = t + 1$

To calculate the profit of various actions π_j, I numerically integrate out over the MCBS sample using their sample weights (forming $dC(Z_i)$).[35] To ensure the single firm problem is continuously differentiable, I use bicubic spline interpolation over the value function grid V_t to estimate the firm's continuation value for any future state $s'_{h,f}$. Once the value function converges, I construct sample moments from observed firm behavior by optimizing the strategy at each state observed in the data using the interpolated continuation values.

The estimation procedure spends most of its time in the profit function and its derivative calculating numeric integrals over demographic distributions. The Broyden-Fletcher-Goldfarb-Shanno extrema-finding algorithm used in step 2(a)i can require over 1000 executions of the profit function and 500 executions of the derivative to solve the single firm problem from an arbitrary starting point to the required precision. A complete execution of the value function iteration to the required precision over 1,200 grid points requires roughly 1,000,000 total executions of the profit and derivative functions. These calculations must be repeated for each market, for each guess of the cost parameters. To make matters worse, these functions must be calculated to full 64-bit floating point accuracy to ensure numerical stability of the outer-most GMM minimization routine. To enhance precision, I perform all numerical integration and moment calculations using the summation algorithm provided by Kahan (1965). To achieve sufficient speed, the model is implemented in C++, uses OpenMP technology to parallelize the optimization of individual states within the value function grid and uses Intel MPI technology to coordinate the simultaneous solution of the value functions for different markets across multiple nodes in a high-performance computing cluster.[36]

6 Results

Tables 7 and 8 contain parameter estimates for the demand side. In general, these estimates are in line with other studies of the demand for Medicare Advantage – particularly those from MPTC. The estimates imply that seniors face a cost of roughly \$1,300 when switching from Medicare to Medicare Advantage and \$1,040 when switching between firms within the

[35] It is possible to construct a distribution of preferences such that the firm's problem in step 2(a)i admits multiple solutions. I perform a number of checks to ensure that the distribution created by the MCBS draws and my demand estimates leads to unique solutions of the firm's problem.

[36] My code is available upon request.

Medicare Advantage system. When compared to the average annual premium of \$825, the switching costs have a large impact on the behavior of consumers. However, even when switching costs are incorporated, Medicare Advantage plans are relatively undesirably: the constant disutility of Medicare Advantage is equivalent to \$1,262. On average, a unit of generosity is worth \$284 to consumers.

Table 9 summarizes the taste distribution implied by these estimates. On average, unhealthy people have a lower preference for MA plans – -3.755 for healthy people versus -3.858 for unhealthy people – but have a greater taste for generosity – .745 for healthy people versus .863 for unhealthy people. This result is in line with common models of heterogenous risk and adverse selection: the greater your risk, the greater your demand for insurance against that risk.[37]

Table 10 contain parameter estimates for the firm's cost function. These estimates imply the average base cost of insuring a healthy person is \$5,293; compared to an average subsidy of \$5,783, firms earn a 9.3% margin. The base cost of insuring an unhealthy person is \$14,609; compared to a benchmark of \$16,298, firms earn a slightly higher margin of 11.6%. The average Medicare payments for healthy and unhealthy people are \$4,390 and \$11,452, respectively. These estimates imply that firms spend an average of \$184 on benefits beyond those of Medicare as measured by the generosity index.

In table 11, I compare these results to those obtained using a static model. Ignoring switching costs completely results in much lower cost estimates, shown in column I. This change is driven by the low overall enrollment rate: roughly two-thirds of seniors in my sample never enroll in a Medicare Advantage plan. The demand estimator must rationalize this behavior and does so in part by significantly lowering the price elasticity.[38] This in turn implies that firm margins are considerably higher and therefore costs are much lower. Though consumers are much less price sensitive, the estimated value of Medicare advantage is roughly the same as the specification with switching costs. Column II of table 11 illustrates the results of incorporating switching costs on the consumer side but assuming firms are myopic on the firm side. This specification results in significantly higher costs for some enrollees. Both of these alternative specifications infer a negative marginal cost of aging.

The results of these alternative specifications are driven by the differences in the policy functions. Figure 2 illustrates a "slice" of the three components of the policy function for a firm in Brown County, Wisconsin, using the preferred estimates of demand and cost under both models of firm behavior. When switching costs are ignored, firms behave uniformly

[37] Lustig (2011) estimates the demand and supply sides simultaneously. This allows for increased flexibility on the demand side with respect to individual product features while still using a generosity index on the supply side – in essence forcing the average taste for generosity to be equal to 1. Data use restrictions prevented the implementation of this approach.

[38] In some sense, the inclusion of the panel data allows the estimator to push some of the observed distaste for MA plans into a switching cost. The switching cost itself is then identified through the implied utility that would have been obtained if the consumer had switched.

across their possible market shares. Per Farrell and Klemperer (2007), when switching costs are included, myopic firms set higher prices and lower generosities than firms who take into account the future value of market share.

Taken together, these results reveal that MA firms do not face lower costs than the federal government. These increased costs may reflect increased administrative costs relative to Medicare, or difficulties negotiating favorable contracts with service providers.

The results of a counterfactual simulation for Brown County, Wisconsin in which the switching costs are lowered by 50% are in column 2 of table 12. The reduced switching cost results in fewer consumers enrolled in Medicare Advantage overall – 11.6% in the counterfactual compared to 23.7% in the baseline scenario. The increase in consumer "liquidity" causes firms with low shares to reduce their generosities – they don't have to offer as generous of a plan to obtain the same share. In equilibrium, firm with larger shares respond by lowering their generosities as well, as they don't have to compete with as generous plans. This results in an overall reduction of the quality of plans offered in the market, causing many consumers to leave. Those that remain are slightly less healthy than enrollees in the baseline scenario. The difference in margins combined with lowered generosities leads to roughly a 10% increase in per-enrollee firm profits – though the total profits are much smaller, as fewer individuals are participating in Medicare Advantage.

Total welfare, including consumer surplus from the MA program, firm profits, and government spending on traditional Medicare benefits and MA subsidies, increases from negative \$224.64 million to negative \$205.82 million. This is largely driven by changes in government spending: in the baseline scenario, the government spends \$79 million on MA subisides and \$158 million on FFS benefits.[39] With the reduction in switching costs, and accompanying reduction in MA enrollment, the government spends \$40 million on subsidies and \$176 on FFS benefits, a total savings of \$21 million.

Column 3 of table 12 reports the results of a second counterfactual scenario in which subsidies are cut by 5%. This change leads to firms no longer offering the zero-premium plan and significantly increasing the premium for their second plan: the positive premium plan in the baseline scenario is offered for \$302 and the counterfactual plan is offered for \$772. Generosity is also reduced – the welfare from generosity for the second plan in the baseline is \$397 and drops to \$311 in the counterfactual plan. These effects combine to allow firms to maintain a degree of profitibility on a per-enrollee basis, but they have a drastic impact on enrollment: only 5.8% of consumers stay in Medicare Advantage.

Aggregating up to the market level, the change results in an improvement in total welfare of \$23.06 million. As in the lower switching cost scenario, the change is driven by government spending and is partly offset by a reduction in total firm profits and consumer

[39] These baseline numbers are computed using model-predicted outcomes. For comparison, CMS reported \$163 million in Medicare FFS spending for Brown County in 2009, the middle of my sample period. (Source: `http://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/Dashboard/Geo-Var-County/GeoVar_County.html`)

welfare.

7 Conclusion

The cost of Medicare, both as a percentage of the federal budget and as a percentage of GDP, has risen steadily since its introduction despite the presence of many 'gaps' in coverage relative to the insurance plans sponsored by employers. Seeking in part to eliminate these gaps cost-effectively, policy makers have implemented the Medicare Advantage system, which offers subsidies to private firms who offer plans to consumers. For historical reasons, the subsidies are currently set well above Medicare's cost – in other words a senior who switches from Medicare to Medicare Advantage increases their burden on taxpayers. However, almost all firms provide supplemental benefits on top of the mandated set of Medicare-equivalent services, meaning that a direct comparison is difficult.

To understand the welfare impact of the Medicare Advantage program, I estimate the cost structure of insurers using a revealed-preference approach. In contrast to previous work on the subject, I introduce a dynamic model of the Medicare Advantage market. The dynamic incentives in my model are driven by the existence of switching costs on the consumer side. Firms know that if they lower their prices today, they can attract a greater share to harvest tomorrow.

I find that the costs of private firms are higher than Medicare's by a significant margin. My estimates also suggest that the risk adjustment formula used by Medicare overcompensates firms for unhealthy enrollees, relative to healthy enrollees. My results highlight the importance of switching costs in this environment – alternative estimation approaches that ignore these costs produce significantly different results.

In a counterfactual simulation, I find a reduction in the switching cost leads to a reduction in the number of Medicare Advantage enrollees. Those that remain are slightly sicker than in the baseline scenario, and firms make greater per-enrollee profits. These changes are driven by the increased "fluidity" of demand: low-share firms don't need to offer particularly generous plans in order to attract share, and high-share firms reduce their generosities as well to save costs. In equilibrium, the average level of generosity offered in the market decreases which pushes many consumers out of the system. Since the subsidies are higher than Medicare's costs, the change leads to an increase in total surplus of $18.9 million per year for a mid-sized county.

My second counterfactual simulation shows that reducing the subsidies offered to firms would also have a substantial positive effect on total welfare. However, this net effect is driven by changes in government spending and comes with a significant loss in welfare for individuals.

There are a number of avenues for future work. While the median number of plans offered by each firm is 2, many firms offer additional plans. Endogenizing the precise number of plans could lead to insights about the administrative costs of marginal plans. Endogenizing entry and exit could provide information about the overall fixed costs of the program.

These results support a view of Medicare as a relatively tax-efficient way to provide medical services to seniors in the United States. While private firms may offer attractive supplemental benefit packages, these currently come at a high cost to taxpayers.

Table 1: Summary statistics: individuals

Variable	Mean	Std. dev. across years
Income	$46,198	3391
Age	75.2	.171
Pct. healthy	56%	1.9
Pct. female	54.5%	0.96
Pct. black	6.1%	0.43
Pct. hispanic	1.5%	0.11
Pct. w/ bachelor's degree	25.5%	0.93
Pct. enrolled in MA	26.0%	0.62
Obs.	9,346	

Note: Observations are at the year-individual level. All calculations use MCBS sample weights.

Table 2: Average Medicare payments across groups of individuals

Category	Mean	Std. dev.
All	$6,743	16,177
Health status		
Healthy	4,390	10,761
Unhealthy	11,452	21,816
Gender		
Male	8,199	18,877
Female	7,117	15,406
Age		
65-69	4,366	14,008
70-74	6,395	18,742
75-79	8,031	15,132
80-84	8,976	15,561
85+	10,731	20,431
Obs.	6,792	

Note: Observations are at the year-individual level. All calculations use MCBS sample weights.

Table 3: Summary statistics: Plans

Variable	**Mean**	**Std. dev.**
Pct. w/ 0 premium	43%	
Annual Premium (if $>$ 0)	$825	629
Generosity	1.736	.809
Med. plans per firm	2	
Avg. plans per firm	2.85	1.4
Pct. firms with $>$ 1 plan	81%	
Pct. of those with 0 prem plan	73%	
Obs.	2306	

Table 4: Competition across years

Year	**Firms**	**Plans**
2008	15.6 (4.69)	20.3 (6.06)
2009	17.1 (5.23)	21.8 (6.14)
2010	13.4 (4.85)	16.8 (6.00)

Note: Figures averaged across 39 counties. Standard deviations in parentheses.

Table 5: Example plan generosities

Variable	**Generosity** ≈ 1	**Generosity** ≈ 2
Drug coverage	Yes	Yes
Vision coverage	No	Yes
Dental coverage	No	No
Primary care copay	$15	$10
Specialist copay	$30	$25
Out-of-pocket limit	$4,000	$2,500
Generosity index	.998	2.03

Table 6: Summary statistics: Markets

Variable	**Mean**	**Std. dev.**
Num. MA firms	15.21	3.633
Total population	1,435	1,835
Medicare population	180.3	207.9
Per-capita income	$40,825	8,330
Num. doctors	3,782	4,831
Num. hospitals	19.1	22.6
Num. nursing facilities	47.9	63.4
Benchmark rate	$10,267	865
Obs.	39	

Note: Markets are defined as counties. Population measured in thousands of people. Per-capita income and benchmark rate are annual figures.

Table 7: Estimates: First stage demand parameters

Variable	**Coefficient**	**Std. Err.**
Income-level price effects (per $1000)		
Price for medium income group	0.0069	0.072
Price for high income group	0.048	0.070
Switching costs		
Medicare-to-Medicare Advantage	-3.971	.056
Inter-contract Medicare Advantage	-3.18	.098
Demographic level effects		
Age	.0049	.012
Female	.0035	.166
Black	-.107	.4289
Hispanic	-1.857	2.255
Graduated high school	-.373	.269
Some college	-.554	.289
Bachelor's degree	-.372	.298
Healthy	.132	.187
Demographic-generosity interactions		
Age	-.001	.005
Female	-.012	.073
Black	.063	.182
Hispanic	.718	.915
Graduated high school	.118	.118
Some college	.129	.126
Bachelor's degree	-.122	.131
Healthy	-.010	.081
Weighted log likelihood	-11,241	
Sample size	12,806	

Note: All dollar amounts are in thousands per year.

Table 8: Estimates: Second stage demand parameters

Variable	**IV**
Premium	-3.05 (.184)
Generosity	0.865 (.105)
Constant	-3.85 (1.03)
Sample size	22,717

Note: Standard errors are in parentheses. All dollar amounts are in thousands per year.

Table 9: Implied taste distribution by health status

Variable	**Mean utils**	**Std. dev.**
Healthy		
α	-3.04	.020
Taste for MA plans (β_i)	-3.755	.249
Taste for generosity (β_{ig})	.745	.130
Unhealthy		
α	-3.03	.021
Taste for MA plans (β_i)	-3.858	.309
Taste for generosity (β_{ig})	.863	.138
Obs.	4122	

Note: Observations are at the year-individual level. All calculations use MCBS sample weights. α is per $1,000.

Table 10: Estimates: Cost parameters

Variable	**Mean**	**Std. err.**
Base cost		
Constant	4.416	0.095
Market population	0.0589	0.0238
Unhealthy	9.142	0.116
Age	0.0850	0.0100
Female	0.0272	0.0591
Generosity cost		
Constant	.0494	.0093
Market population	0.003	0.032
Unhealthy	.0309	.0269
Age	.0095	.0065
Female	-0.0118	.0245
Generosity quadratic cost		
Constant	0.0450	.0240
Unhealthy	0.0004	0.0392
Obs.	1,779	

Note: Observations are at the year-firm level. Costs are measured in thousands of dollars per enrollee per year.

Table 11: Estimation summary

Specification	**(I)**	**(II)**	**(III)**
Switching costs	No	Yes	Yes
Supply model	Static	Static	Dynamic
Demand			
Medicare-to-MA cost	N/A	1,302	1,302
Mean elasticity	.131	1.56	1.56
Mean inclusive value	$223	252	252
Supply			
Base cost	$3,450	5,509	4,416
Additional cost of unhealthy enrollees	$8,808	9,071	9,142
Cost per year of enrollee's age above 65	-$80	-23	85
Mean generosity expenditure	$432	213	184

Note: All dollar amounts annualized. Base cost is for a 65-year-old male in good health enrolled in a plan with 0 generosity. Inferred values (elasticity, inclusive value, generosity expenditures) calculated using observed plan characteristics and MCBS sample weights.

Table 12: Counterfactual: 50% decrease in switching costs

	Baseline	**Lower switching costs**	**Lower subisides**
MA enrollment			
Healthy	24.3%	11.6%	5.7%
Unhealthy	23.2%	11.7%	5.9%
Per-enrollee:			
Welfare from generosity	$328	278	311
Average subsidy	$10,162	10,568	9,728
Mean firm costs	$9,951	10,249	10,005
Mean firm profits	$542	606	495
Total: (in millions)			
Consumer welfare	$8.30	7.97	3.49
Firm profits	$4.23	2.31	0.95
Subsidies to MA firms	$79.29	40.36	18.65
Medicare FFS spending	$157.87	175.73	187.37
Total welfare (millions)	$(224.64)	(205.82)	(201.58)

Note: Figures for Brown County, Wisconsin. Welfare, subsidies, and profits calculated using model-predicted enrollment and MCBS sample weights. Total consumer welfare is calculated using the inclusive value metric across all Medicare recipients and does not include base welfare generated by Medicare FFS benefits. "Lower subsidies" counterfactual calculated with firms only offering a single plan each with unrestricted premium.
For comparison, CMS reported 2010 Medicare FFS spending of $163 million for the county.

Figure 1: Firm shares and actions

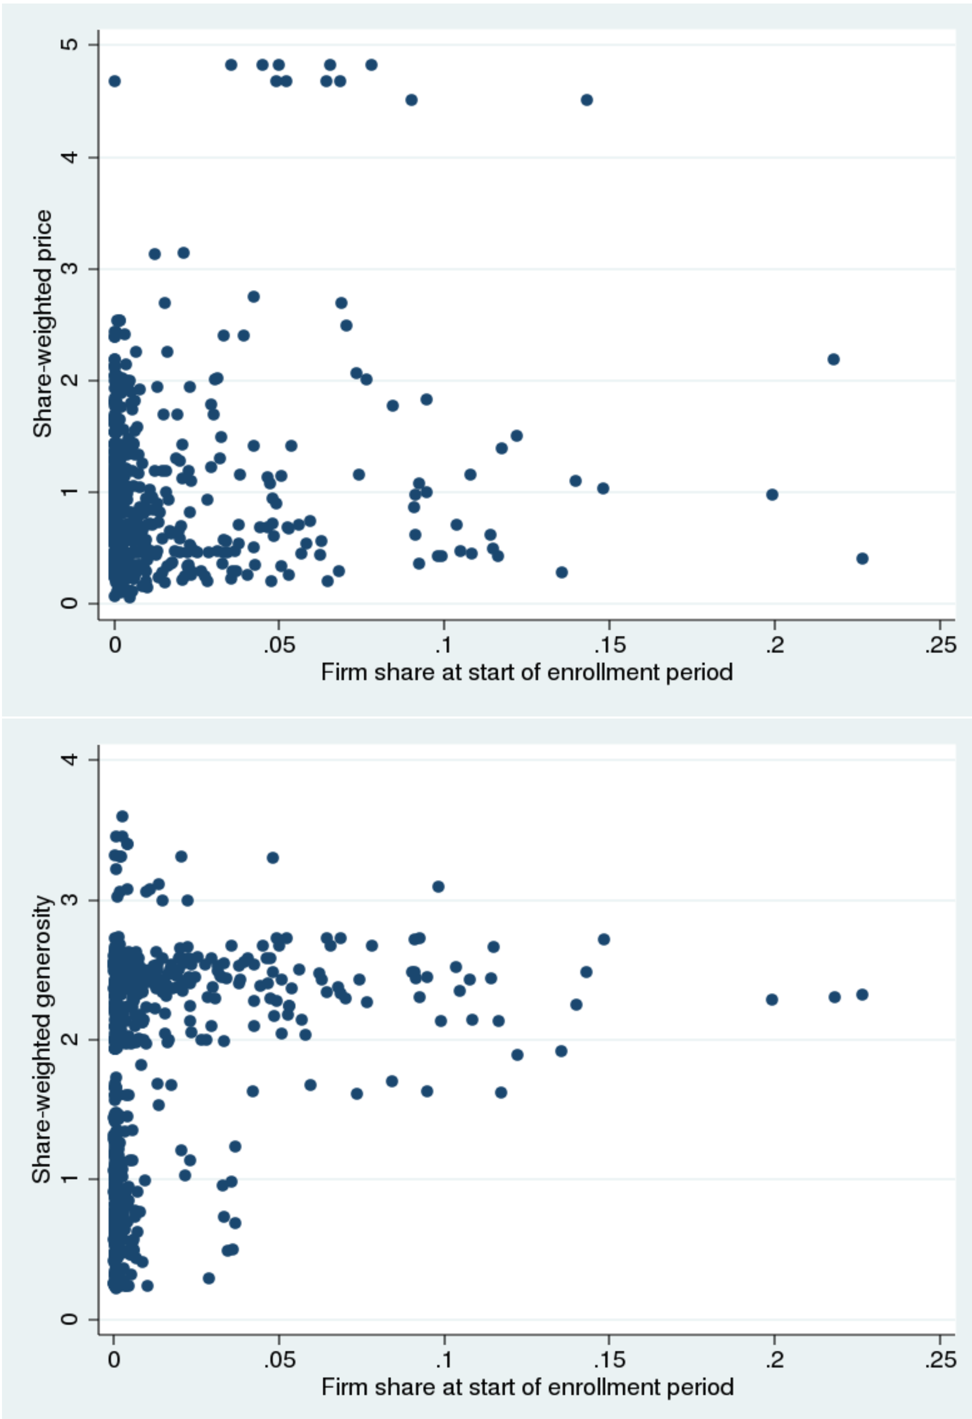

Note: Price measured in thousands of dollars per year. Each dot represents one firm-year. Price and generosity weighted by previous period plan-level shares. Zero price plans excluded.

Figure 2: Policy function comparison

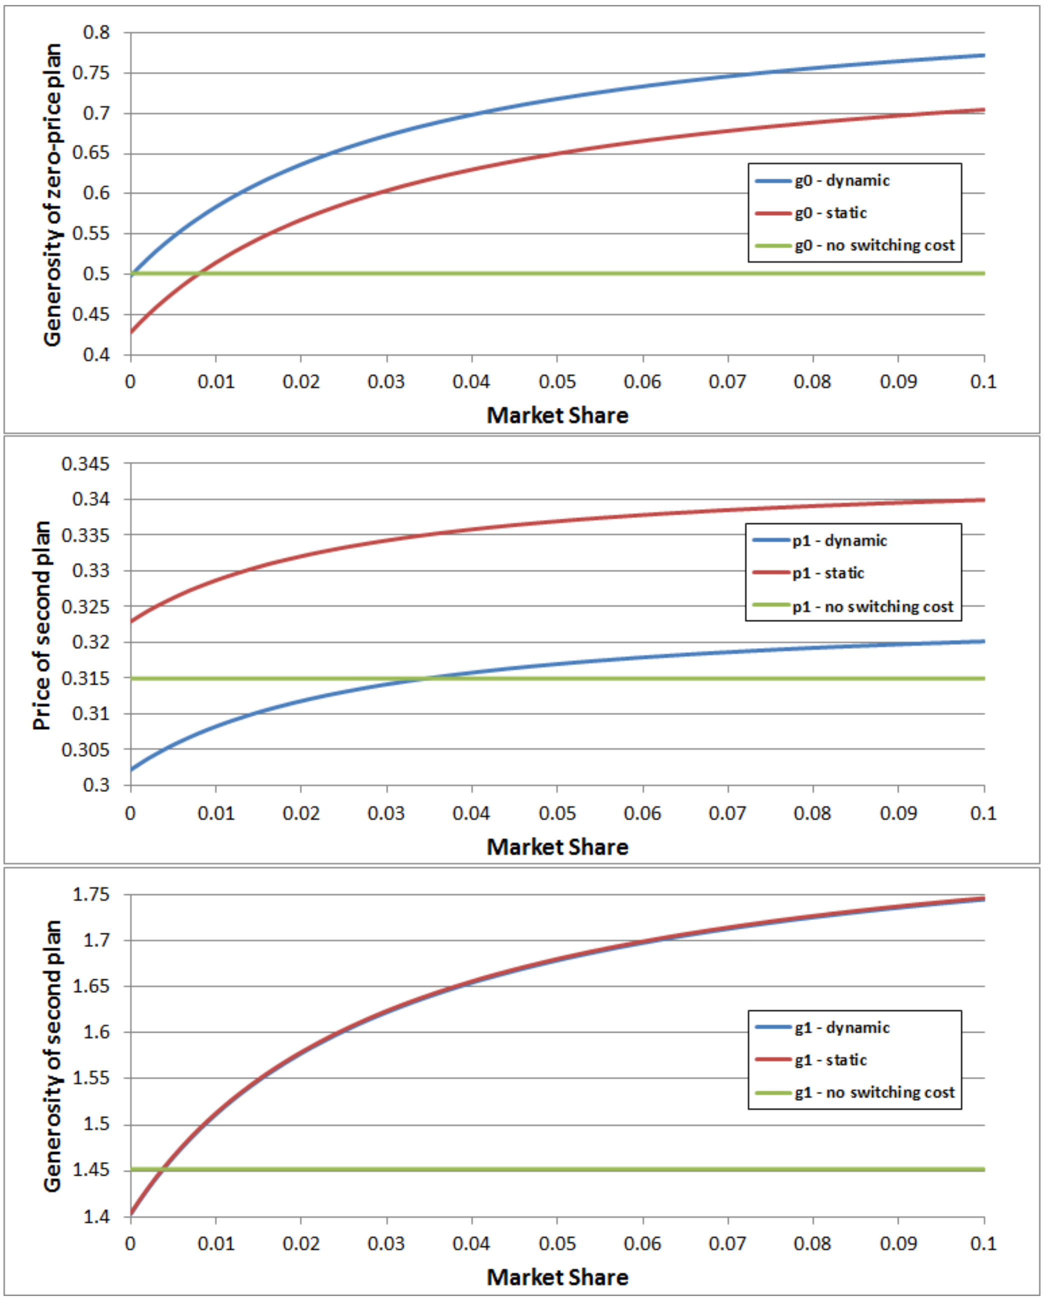

Note: Policy function shown for a firm in Brown County, Wisconsin. Policy functions calculated using preferred estimates for demand and cost. Market share is of both healthy and unhealthy individuals.

References

Jason Abaluck and Jonathan Gruber. Evolving choice inconsistencies in choice of prescription drug insurance. Working Paper 19163, National Bureau of Economic Research, June 2013. URL http://www.nber.org/papers/w19163.

Victor Aguirregabiria and Pedro Mira. Sequential estimation of dynamic discrete games. *Econometrica*, 75(1):1–53, 2007. ISSN 1468-0262. doi: 10.1111/j.1468-0262.2007.00731.x. URL http://dx.doi.org/10.1111/j.1468-0262.2007.00731.x.

Naoki Aizawa and You Suk Kim. Advertising Competition and Risk Selection in Health Insurance Markets: Evidence from Medicare Advantage. Job market paper, November 2013. URL http://economics.sas.upenn.edu/graduate-program/candidates/you-suk-kim.

Kenneth J. Arrow. Uncertainty and the welfare economics of medical care. *American Economic Review*, LIII(5):941–973, December 1963.

Patrick Bajari, C. Lanier Benkard, and Jonathan Levin. Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370, 2007. ISSN 1468-0262. doi: 10.1111/j.1468-0262.2007.00796.x. URL http://dx.doi.org/10.1111/j.1468-0262.2007.00796.x.

Steven Berry. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262, 1994.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):pp. 841–890, 1995.

Steven Berry, James Levinsohn, and Ariel Pakes. Voluntary export restraints on automobiles: Evaluating a trade policy. *American Economic Review*, 89(3):400–430, 1999. doi: 10.1257/aer.89.3.400. URL http://www.aeaweb.org/articles.php?doi=10.1257/aer.89.3.400.

Jason Brown, Mark Duggan, Ilyana Kuziemko, and William Woolston. How does Risk Selection Respond to Risk Adjustment? Evidence from the Medicare Advantage Program. Working Paper 16977, National Bureau of Economic Research, April 2011. URL http://www.nber.org/papers/w16977.

Randall D. Cebul, James B. Rebitzer, Lowell J. Taylor, and Mark E. Votruba. Unhealthy insurance markets: Search frictions and the cost and quality of health insurance. *American Economic Review*, 101(5):1842–71, 2011. doi: 10.1257/aer.101.5.1842. URL http://www.aeaweb.org/articles.php?doi=10.1257/aer.101.5.1842.

Gary Chamberlain. Asymptotic efficiency in estimation with conditional moment restrictions. *Journal of Econometrics*, 34(3):305–334, 1987.

Michael Chernew. General equilibrium and marketability in the health care industry. In Peter J. Hammer, Deborah Haas-Wilson, Mark A. Peterson, and William M. Sage, editors, *Uncertain times: Kenneth Arrow and the changing economics of health care*. Duke University Press, 2003.

CMS. Medicare Managed Care Manual. Technical report, 2014. URL `http://www.cms.gov/Regulations-and-Guidance/Guidance/Manuals/Internet-Only-Manuals-IOMs-Items/CMS019326.html`.

Allan Collard-Wexler. Demand fluctuations in the ready-mix concrete industry. *Econometrica*, 81(3):1003–1037, 2013. ISSN 1468-0262. doi: 10.3982/ECTA6877. URL `http://dx.doi.org/10.3982/ECTA6877`.

Gregory Crawford. Accommodating endogenous product choices: A progress report. *International Journal of Industrial Organization*, 30:315–320, 2012.

Vilso Curto, Liran Einav, Jonathan Levin, and Jay Bhattacharya. Can Health Insurance Competition Work? The U.S. Medicare Advantage Program. Working paper, 2014.

Leonardo Dagum and Ramesh Menon. OpenMP: an industry standard API for shared-memory programming. *Computational Science & Engineering, IEEE*, 5(1):46–55, 1998.

Mark Duggan, Amanda Starc, and Boris Vabson. Who benefits when the government pays more? pass-through in the medicare advantage program. Working Paper 19989, National Bureau of Economic Research, March 2014. URL `http://www.nber.org/papers/w19989`.

A. C. Enthoven. Consumer-Choice Health plan (second of two parts). A national-health-insurance proposal based on regulated competition in the private sector. *N. Engl. J. Med.*, 298(13):709–720, Mar 1978. [DOI:10.1056/NEJM197803302981304] [PubMed:415241].

F. Eppig and G. S. Chulis. Matching MCBS and Medicare Data: The Best of Both Worlds. *Health Care Financing Review*, 18(3):211–229, 1997.

Richard Ericson and Ariel Pakes. Markov-perfect industry dynamics: A framework for empirical work. *The Review of Economic Studies*, 62(1):53–82, 1995.

Joseph Farrell and Paul Klemperer. Chapter 31 coordination and lock-in: Competition with switching costs and network effects. volume 3 of *Handbook of Industrial Organization*, pages 1967 – 2072. Elsevier, 2007. doi: http://dx.doi.org/10.1016/S1573-448X(06)03031-7. URL `http://www.sciencedirect.com/science/article/pii/S1573448X06030317`.

J Gabel, C Jajich-Toth, G de Lissovoy, T Rice, and H Cohen. The changing world of group health insurance. *Health Affairs*, 7(3):48–65, 1988. doi: 10.1377/hlthaff.7.3.48. URL http://content.healthaffairs.org/content/7/3/48.short.

Amit Gandhi, Zhentong Lu, and Xiaoxiz Shi. Estimating Demand for Differentiated Products with Error in Market Shares. Working paper, 2013.

Martin Gaynor and Robert J. Town. Chapter nine - competition in health care markets. In Thomas G. Mcguire Mark V. Pauly and Pedro P. Barros, editors, *Handbook of Health Economics*, volume 2 of *Handbook of Health Economics*, pages 499 – 637. Elsevier, 2011. doi: http://dx.doi.org/10.1016/B978-0-444-53592-4.00009-8. URL http://www.sciencedirect.com/science/article/pii/B9780444535924000098.

Sherry Glied. Chapter 13 managed care. volume 1, Part A of *Handbook of Health Economics*, pages 707 – 753. Elsevier, 2000. doi: http://dx.doi.org/10.1016/S1574-0064(00)80172-9. URL http://www.sciencedirect.com/science/article/pii/S1574006400801729.

Austan Goolsbee and Amil Petrin. The consumer gains from direct broadcast satellites and the competition with cable tv. *Econ*, 72(2):351–381, March 2004.

Benjamin R. Handel. Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 103(7):2643–82, 2013. doi: 10.1257/aer.103.7.2643. URL http://www.aeaweb.org/articles.php?doi=10.1257/aer.103.7.2643.

Kate Ho, Joseph Hogan, and Fiona Scott Morgan. The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program. Working paper, 2014.

V. Joseph Hotz and Robert A. Miller. Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, 60(3):497–529, 1993. doi: 10.2307/2298122. URL http://restud.oxfordjournals.org/content/60/3/497.abstract.

William Kahan. Pracniques: Further remarks on reducing truncation errors. *Communications of the ACM*, 8(1):40, January 1965. ISSN 0001-0782. doi: 10.1145/363707.363723. URL http://doi.acm.org/10.1145/363707.363723.

Kaiser Family Foundation. Medicare advantage fact sheet. Technical report, 2014. URL http://kff.org/medicare/fact-sheet/medicare-advantage-fact-sheet/.

Josh Lustig. Measuring Welfare Losses from Adverse Selection and Imperfect Competition in Privatized Medicare. Working paper, 2011.

Martin Markovich. *The Rise of HMOs*. PhD thesis, RAND Graduate School, May 2003.

Keaton Miller, Amil Petrin, Robert Town, and Michael Chernew. Health Plan Changes in Response to Medicare Advantage Payment Rates. Working paper, 2014.

Marilyn Moon, Sheila Burke, Roderick DeArment, Judith Feder, Robert Helms, William Hoffman, Walter Maher, Martha Phillps, Anna Rappaport, William Scanlon, James Tallon, and Robert Valdez. Financing Medicare's Future. Technical report, National Academy of Social Insurance, 2000.

Whitney K. Newey and Daniel McFadden. Chapter 36 large sample estimation and hypothesis testing. volume 4 of *Handbook of Econometrics*, pages 2111 – 2245. Elsevier, 1994. doi: http://dx.doi.org/10.1016/S1573-4412(05)80005-4. URL http://www.sciencedirect.com/science/article/pii/S1573441205800054.

Joseph P. Newhouse, Mary Price, J. Michael McWilliams, John Hsu, and Thomas McGuire. How much favorable selection is left in medicare advantage? Working Paper 20021, National Bureau of Economic Research, March 2014. URL http://www.nber.org/papers/w20021.

J. A. Nyman, N. A. Barleen, B. E. Dowd, D. W. Russell, S. J. Coons, and P. W. Sullivan. Quality-of-life weights for the US population: self-reported health status and priority health conditions, by demographic characteristics. *Med Care*, 45(7):618–628, Jul 2007.

Ariel Pakes and Paul McGuire. Stochastic algorithms, symmetric markov perfect equilibrium, and the 'curse' of dimensionality. *Econometrica*, 69(5):1261–1281, 2001. ISSN 1468-0262. doi: 10.1111/1468-0262.00241. URL http://dx.doi.org/10.1111/1468-0262.00241.

Ariel Pakes, Michael Ostrovsky, and Steven Berry. Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *The RAND Journal of Economics*, 38(2):373–399, 2007a. ISSN 1756-2171. doi: 10.1111/j.1756-2171.2007.tb00073.x. URL http://dx.doi.org/10.1111/j.1756-2171.2007.tb00073.x.

Ariel Pakes, Jack Porter, Kate Ho, and Joy Ishii. Moment inequalities and their application. Technical report, cemmap working paper, Centre for Microdata Methods and Practice, 2007b.

Robert Pear and Walt Bogdanich. Some Successful Models Ignored As Congress Works on Drug Bill. *The New York Times*, 2003.

Shi Qi. The impact of advertising regulation on industry: the cigarette advertising ban of 1971. *The RAND Journal of Economics*, 44(2):215–248, 2013. ISSN 1756-2171. doi: 10.1111/1756-2171.12018. URL http://dx.doi.org/10.1111/1756-2171.12018.

John Rust. Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica: Journal of the Econometric Society*, pages 999–1033, 1987.

Stephen P. Ryan. The costs of environmental regulation in a concentrated industry. *Econometrica*, 80(3):1019–1061, 2012. ISSN 1468-0262. doi: 10.3982/ECTA6750. URL `http://dx.doi.org/10.3982/ECTA6750`.

Conrad Sanderson. Armadillo: An Open Source C++ Linear Algebra Library for Fast Prototyping and Computationally Intensive Experiments. Technical report, NICTA, 2010.

US Department of Health and Human Services, Health Sources and Services Administration, Bureau of Health Workforce, Rockville, MD. Area Health Resources Files. Available at `htpp://ahrf.hrsa.gov`, 2013-2014.

Gabriel Y Weintraub, C Lanier Benkard, and Benjamin Van Roy. Markov perfect industry dynamics with many firms. *Econometrica*, 76(6):1375–1411, 2008.

Gabriel Y Weintraub, C Lanier Benkard, and Benjamin Van Roy. Computational methods for oblivious equilibrium. *Operations research*, 58(4-part-2):1247–1265, 2010.

Thomas Youle. How much did manipulation distort the libor? Working paper, 2014.