Estimating Costs When Consumers Have Inertia: Are Private Medicare Insurers More Efficient?

Keaton Miller*
University of Oregon

June 30, 2019

Abstract

When consumers have inertia, firms have an intertemporal pricing incentive to capture and then harvest market share. This implies that marginal cost estimators derived from static profit-maximization first-order conditions may be biased depending on the distribution of shares in the data. I examine this empirically with the Medicare Advantage program, through which seniors may obtain health insurance through private firms rather than the traditional public system. I present evidence that firms behave dynamically, and estimate a dynamic multiproduct oligopoly model which incorporates inertia and adverse selection. I find MA firms face costs which may exceed the government's, contrary to static estimates.

JEL Codes: L11, C61, D43, I13

Keywords: Cost Estimation, Switching Costs, Medicare, Industrial Organization, Medicare Advantage, Dynamic Oligopoly.

---

*keatonm@uoregon.edu. I thank Lanier Benkard, Mark Colas, Clare Evans, David Evans, Ben Hansen, Thomas Holmes, Kendall Houghton, Michael Kuhn, Van Kolpin, Grant McDermott, Amil Petrin, Scott Petty, Boyoung Seo, Christopher Snyder, Jeff Thurk, Robert Town, Glen Waddell, Joel Waldfogel, Wes Wilson, seminar participants at the University of Oregon, the University of Minnesota, the University of Connecticut, the University of Indiana, and conference participants too numerous to list for feedback and support. All errors are mine. This work benefited from access to the University of Oregon high performance computer, Talapas, and support from the University of Minnesota Supercomputing Institute.

1 Introduction

The broad strokes of the pricing-under-consumer-inertia story are straightforward. If consumers exhibit inertia (e.g. through explicit or implicit switching costs), firms have an incentive to engage in dynamic pricing (Farrell and Klemperer, 2007): When a firm has a small market share, it reduces prices to "capture" consumers, and once it has obtained enough share, it raises prices to "harvest" those consumers even if costs have remained constant over time.

The thrust of this article is to empirically examine the interplay between that story and the normal story used in the industrial organization literature to estimate marginal costs through revealed preferences. My setting is the Medicare Advantage (MA) program in the United States, where costs have potentially first-order policy relevance. MA, which I describe in-depth in Section 2, allows Medicare beneficiaries to forgo the Traditional Medicare (TM) fee-for-service benefit structure and enroll in one of many differentiated health plans offered by private firms. The firm assumes the financial and logistic responsibilities of the enrollee's care and in return receives premiums from the enrollee and a risk-adjusted per-capita payment from the government.[1] Firms must cover everything that Medicare covers, and may offer supplemental benefits.

These competing systems have been the source of much public policy debate, as many commentators have pointed to the U.S.'s fragmented mechanism for financing health care as a source of inefficiency; a single-payer system, they argue, would deliver health services at a lower total cost to society (Woolhandler et al., 2003, Weisbart, 2012). Others have argued that much of the rise in spending over the past decades has come from other factors such as a lack of competition at key points in the supply chain and the solution is to increase competition throughout the health care system (Cogan et al., 2005, Parente, 2018). A direct cost comparison between TM and MA – that is, a comparison of the dollars spent by the risk bearing agent (either the government or a private firm) to provide promised services – could have implications for the entire U.S. population.

TM expenditures are available through public records. Expenditures by MA firms, however,

[1] I define the price of an MA plan to be "premiums paid by the enrollee which do not vary with utilization" and use "price" and "premium" interchangeably. I do not include out-of-pocket costs such as copays as those vary with the number of services used. I account for these other factors in my demand system separately.

are trade secrets, and so many have used a "static" revealed preference approach to infer costs in MA (Town and Liu, 2003, Dunn, 2010, Curto et al., 2018).[2] However, as Handel (2013) found in other health insurance settings, Miller, Petrin, Town, and Chernew (2019, hereafter MPTC) estimate significant switching costs in MA and thus firms may engage in dynamic pricing and product design strategies.

I analyze MA using microdata on consumer demographics and choices, and administrative data on plan options and market shares, from 2008-2015. I start by describing these data and presenting reduced-form evidence of dynamic behavior in Section 3: firms which earn a higher market share in one period increase their prices and decrease their generosity more in the next period than firms which earn lower shares. I then write a dynamic model of oligopoly in Section 4 focusing on the intertemporal incentives driven by consumer inertia. On the demand side I adopt the discrete choice model of MPTC which in addition to switching costs captures heterogeneity in preferences for plan premiums and benefits – in other words those who are sicker may prefer more generous plans and may have different incomes and thus different price elasticities of demand. On the supply side, forward-looking multiproduct firms choose prices and benefits to maximize profits given their costs – which may vary by consumer thus generating the potential for adverse selection – and information about their competitive environment including the distribution of consumer heterogeneity.

I take the dynamic model to the data in Section 5. Computing a full-information Markov-Perfect Equilibrium is computationally intractable due to the number of competitors in this setting (the median number of firms in the markets I consider is 9), the degree of consumer heterogeneity, and the corresponding dimensionality of the state space. I therefore define a restricted-information equilibrium concept inspired by the Oblivious Equilibrium framework of Weintraub et al. (2008) which may be of independent methodological interest. The restricted information set focuses on the *effect* of competitive pressure on a firm's ability to attract consumers rather than the *sources* of that pressure yet maintains the potential for correlation between consumers' health and price elasticity. I estimate the model's parameters in two steps. First, I estimate the demand-side parameters

[2] The inferred costs referenced here (along with the cost estimates I produce) are "economic" costs, not accounting costs, and can be interpreted as the dollar amount that MA firms must be spending on marginal enrollees to justify the firms' observed behavior. I discuss this nuance further when I present the results in Section 6.

following the maximum likelihood and instrumental variables approach of MPTC. Second, I form moments from the restricted-information model's predictions of optimal firm behavior and compare them to their sample equivalents to estimate the parameters of the firms' marginal cost function (Newey and McFadden, 1994).[3]

In Section 6 I present my primary finding that, averaged across the markets in my data, firms expect to spend \$7,784 per year to insure a healthy individual and \$19,280 to insure an unhealthy individual when they provide coverage with no supplemental premium (i.e. at the same cost to the consumer as TM). In contrast, TM spends an average of \$8,080 on healthy individuals and \$17,250 on unhealthy individuals. This stands in contrast to the previous literature, which generally finds lower costs for MA firms.

This result is driven by the data and the difference between the way in which the static and dynamic approaches interpret that data. The static approach assumes that firms play a series of one-shot Bertrand-Nash games in prices and uses the first-order conditions of the firms' profit-maximization problems to "back out" marginal costs (Berry et al., 1995). *Ceterus paribus*, if this estimator observes a firm with low prices, it infers low costs. The dynamic estimator, however, assumes that period games are intertemporally linked via market shares. If it observes a firm with low prices which also has a low share, it concludes that some portion of that low price stems from the incentive to capture customers, and therefore that the firm's costs are higher than the price alone implies. On the other hand, if it observes a low-price, high-share firm, it concludes the firm is harvesting and thus the firm's cost may be even lower than the static model would estimate.

The data are characterized by a preponderance of firms with low market shares and low prices which drive my estimates of cost. I re-estimate the model with the static approach, which only sees the low prices, and find costs which are 13% lower than the dynamic estimates. I also show that if I observe a number of firms with low prices and high shares, the dynamic estimator concludes that harvesting is taking place and can produce lower cost estimates than the static estimator. Thus, in general, static estimates of dynamic pricing games can be biased in either direction, depending on the extent to which the data is characterized by firms with high or low shares.

[3] MacKay and Miller (2019) discuss the interchangeability of instruments and supply-side model assumptions.

I therefore contribute to the ongoing discussion about ways in which dynamic incentives can affect the pricing behavior of firms and thus bias static estimates of costs. Besanko et al. (2014) find similar pricing dynamics in environments with learning-by-doing. Hendel and Nevo (2006) show demand-side dynamics in the form of stockpiling bias static estimates of elasticity (which drive marginal cost estimates). Goettler and Gordon (2011) estimate a model of dynamic duopoly to investigate innovation incentives in durable goods markets.

My work also contributes to the broader literature on inertia in health insurance. Ho et al. (2017) study consumer behavior in the Medicare Part D (prescription drug) market and use a static counterfactual exercise to conclude that inertia leads to greater expenditures for both consumers and the government. Handel (2013) studies inertia and selection in employer-provided insurance and finds that reducing inertia leads to increased selection. As in that work, I assume that consumers are myopic to simplify my analysis. Indeed, Abaluck and Gruber (2011, 2016) argue that Part D consumers did not improve their ability to select plans over time as the market matured, though Ketcham et al. (2012) provide some evidence that consumers update their beliefs about future drug usage from past experiences.

Finally, my results on the relative profitability of different individuals contribute to the debate around selection and risk adjustment in MA. Brown et al. (2014) use data on Medicare expenditures to study the changes in incentives created by the introduction of the risk adjustment system used by the government to compute MA subsidies. Newhouse et al. (2012) and Newhouse et al. (2015) argue that firms have little incentive to cream-skim under the risk-adjustment system, though my results correspond with other evidence that selection incentives remain (Aizawa and Kim, 2018).

2 The Medicare Advantage Program

Medicare was created in 1965 to address the lack of health insurance among senior citizens.[4] The original system provided basic hospital (Part B) and medical (Part B) insurance to seniors (age 65 or older). Under its fee-for-service (FFS) structure, private health-care professionals provide

[4]See McGuire et al. (2011) for a comprehensive history of the Medicare Advantage program.

services to beneficiaries and receive payments from Medicare according to a pre-set reimbursement schedule. Beneficiaries pay applicable copays and/or coinsurance that vary according to the service.

Over time, reforms and changing demographics have slowly expanded Medicare's role within the U.S. health care system. In 1972, eligibility was extended to individuals under 65 with certain disabilities and illnesses. These extensions, as well as expansions to the range of services covered by the program, led to increased costs. In 1970, Medicare composed about 0.5% of GDP. By 1980, Medicare had more than doubled in size to 1.1% of GDP.

This growth led policy-makers to experiment with different cost-containment and care-delivery strategies. In 1982 Congress authorized Medicare administrators to engage in a series of "Part C" trials based on the ideas of Enthoven (1978) in which the government handed over management of the medical care of select groups of Medicare enrollees to private insurers in exchange for a payment that did not vary with the realized medical expenditures of each individual. This program was brought to the entire country in 1997 under the name Medicare+Choice.

The Medicare+Choice program struggled to attract plans – entry by firms was largely limited to suburban areas and many rural and inner-city residents did not have access to any plans (Pear and Bogdanich, 2003). Additionally, since the per-person "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 (Newhouse et al., 2012). As a consequence, nationwide enrollment hovered near 5 million – less than 10% of the Medicare population.

The Medicare Prescription Drug, Improvement, and Modernization Act of 2003 aimed to fix these issues with a nearly universal increase in subsidy rates and the introduction of a comprehensive risk adjustment mechanism supported by research using inpatient encounter data (Ellis et al., 1996, Pope et al., 2000). Under this system, firms submit demographic and diagnostic data about enrollees to the Centers for Medicare and Medicaid Services (CMS) at the time of enrollment. CMS assigns each enrollee a score based on its FFS expenditures on similar individuals in Traditional Medicare (TM). Payments to firms are then adjusted according to these risk scores. Proponents argued that this mechanism would compensate firms for taking on risk without reimbursing specific procedures

– thus maintaining the profit motive which would (in theory) lead to cost reductions. The program was renamed Medicare Advantage (MA). By 2015, 95% of Medicare beneficiaries had an MA plan in their county and enrollments had increased to 16.3 million.[5]

Currently, Medicare beneficiaries may enroll in an MA plan during an "Open Enrollment" period in the fall of each year.[6] Those who do forgo TM benefits and receive care exclusively through the MA plan. MA enrollees pay the Medicare Part B premium and any premium charged by the MA plan. Firms compete through pricing, benefit design, and provider networks, and often heavily market their plans (Aizawa and Kim, 2018). As a consequence, MA plans generally provide a more generous benefit package than TM, including coverage of service categories not included in TM such as dental and optometry services. Many plans offer prescription drug benefits.

Firms also differ in their model of care delivery. Most plans are offered by health maintenance organizations (HMOs), though preferred provider organization plans (PPOs) and private fee-for-service plans (PFFS) are available in many geographies. HMOs generally employ providers and operate facilities directly (Markovich, 2003), while PPOs negotiate contracts to form a network of independent providers (Gabel et al., 1988). Like TM, PFFS plans offer a fixed reimbursement rate per procedure and providers choose to accept patients on a case-by-case basis. HMOs generally require recipients to choose a primary care provider and obtain referrals from that provider to receive specialist services. PPOs may or may not require referrals. Both HMOs and PPOs charge additional copays or coinsurance for services obtained from 'out-of-network' providers. These features arguably allow HMOs and PPOs to more tightly control both the degree of utilization by plan enrollees as well as the cost of each service, thus creating the potential for savings over FFS plans.[7]

The enrollee-specific subsidy paid to insurers is based on a "benchmark" rate for each county

[5]The Affordable Care Act contained provisions to reduce payments to MA plans (CMS, 2010, Obama, 2016). In the analysis in Section 6, I focus on MA plans with zero premiums to capture those plans which are closest to TM in benefits. To the extent that these cuts forced firms to reduce benefits and more closely match TM, this policy change only makes my comparison more appropriate. In any case, an earlier draft of this article which only included data from 2008-2010 (i.e. before the majority of the cuts were implemented) found similar qualitative results.

[6]Beneficiaries may also enroll in MA when they become newly Medicare eligible and after certain life events (e.g. relocation, death in the family, etc.).

[7]MA allows firms to register as 'Regional PPOs' if they cover a large area – at minimum an entire state. Regional PPOs face a slightly different subsidy formula but otherwise operate similarly to other plans. As these plans receive roughly 1% market share I group them with other PPOs.

which varies across geographies and over time (Newhouse et al., 2012). CMS calculates the schedule each year by computing county-level averages of risk-adjusted per-capita FFS Medicare spending. Counties are then ranked by this average spending and grouped by quartile. The provisional benchmark rate for each county is set equal to the average spending measure multiplied by a percentage which varies with the quartile: 95% for counties in the top quartile, 100% for the second, 107.5% for the third, and 115% for the bottom quartile. Finally, CMS applies a cap and floor that varies by Census-defined urban/rural county status.

Each year, after benchmarks after released, insurers submit plan-level 'bids.' The bid for each plan represents the insurer's offer to provide a particular set of benefits to a person of average risk residing in a particular county starting in January of the next year in exchange for a stated payment – firms therefore may make decisions at the county-year level and I thus define markets at this level for analysis.[8] The bid amount may be above or below the benchmark rate. Plan bids which are above the benchmark rate must charge premiums to enrollees. Plan bids which are below the benchmark rate earn 75% of the difference as a 'rebate.' MA plans that offer a prescription drug benefit submit a separate bid which maps in a similar way to a Part D premium.

At the beginning of my data, the rebate payment was equal to 75% of the difference between the bid and the benchmark. In 2008, CMS introduced a system to measure insurer quality by assessing performance along multiple dimensions and assigning a summary 'star rating' to firms. In 2012, after multiple iterations of the criteria, CMS began using the rating to determine the rebate for each firm's plans. Under this new approach, firms with at least 4.5 stars (out of five) earn 70% of the difference between the benchmark and the bid. New entrants and those with 3.5 or 4 stars earn 65% of the difference. All others earn 50%. Additionally, firms with at least 4 stars have a 5% bonus applied to the benchmark rate itself (CMS Office of the Actuary, 2017).

Let ϕ_{jt} capture this star bonus (if any) to the benchmark applied to plan j in year t and let λ_{jt} capture the rebate percentage. The payment from CMS to insurers for an enrollee i living in

[8] I do not model the bidding process explicitly, though I use information about bids to determine the payments offered to each firm.

county m enrolled in plan j in year t based on a benchmark B_{mt} can be calculated with

$$Payment_{ijt} = \begin{cases} B_{mt} \times \phi_{jt} \times RiskAdjustment_{it} & \text{if } bid_{jt} \geq B_{mt}\phi_{jt} \\ (bid_{jt} + \lambda jt \times (B_{mt} \times \phi_{jt} - bid_{jt})) \times RiskAdjustment_{it} & \text{if } bid_{jt} < B_{mt}\phi_{jt} \end{cases}. \quad (1)$$

For simplicity, I denote the market-level (i.e. county-year level) benchmark with B_m and denote risk-neutral (i.e. $RiskAdjustment = 1.0$) plan-specific payments with B_j.

3 Data and Reduced Form Evidence

To analyze the pricing behavior of firms and infer costs from that behavior, I combine CMS administrative data on plan design and enrollment with microdata on consumer demographics and choices from the Medicare Current Beneficiary Survey and geographic data from the Census Bureau.

3.1 Medicare beneficiaries

My data on the Medicare population come from the Medicare Current Beneficiary Survey (MCBS), a nationally-representative, rolling-panel survey sponsored by CMS and produced by Westat. Participants are interviewed in-person and over-the-phone over three years and responses are linked to CMS administrative data to ensure accuracy. The Medicare population is sampled via a multi-level clustering procedure. First, several focus geographies consisting of metropolitan statistical areas or groups of contiguous non-metropolitan counties are chosen semi-randomly to reflect the diversity of the United States. Within each focus area, Medicare beneficiaries are chosen according to a weighted randomization procedure. This clustering ensures that, although I do not observe beneficiaries in every county, within a given area there is variation in demographics and plan enrollment. The MCBS provides sampling weights which I use throughout my analysis.

I collect the MCBS responses for 2008-2015.[9] I observe detailed demographic information, including income, age, sex, race, and education. Respondents self-report their health status, choosing between Excellent, Very Good, Good, Fair, and Poor. I collapse these into two categories:

[9]The MCBS did not release data for 2014.

"healthy" which includes Excellent, Very Good, and Good, and "unhealthy" which includes Fair and Poor. I also observe the respondent's county of residence and their MA plan, if any.[10]

To model the MCBS micro-data as draws from the process that generates the plan-level enrollment data described below, I first construct a number of administrative indicators to capture individuals in the MCBS data who do not fall into the standard Medicare-eligible set often studied in the literature (i.e. age 65-plus retirees without outside insurance). As these individuals can and do purchase MA plans I cannot exclude them entirely. These variables include indicators for individuals with employer-provided insurance plans, those whose original Medicare eligibility was not age-related, those with end-stage renal disease (ESRD), and those who are not full-year enrollees in Part A and Part B. I then exclude Medicaid-eligible individuals and those with missing address information. After applying these exclusions, the sum of the MCBS sample weights differs from the CMS-reported number of eligible people by less than 1%.

Those who do not enroll in MA have access to other insurance options, and variation in the price of those options may change the demand for MA plans. I focus on Medicare supplemental insurance (a.k.a. Medigap) which pays for care not covered by TM. The design of Medigap plans is standardized by CMS: all "Plan A" policies offer the same benefits regardless of the insurer. For each person, I obtain the rate for Medigap Plan C offered by United Healthcare that year from Weiss Ratings. Plan C covers most of the coinsurance and deductibles that beneficiaries are responsible for under TM and is the most popular Medigap plan.[11]

Summary statistics on my 58,444 individual-year observations covering 2,947 county-year markets and 32,993 unique individuals are reported in Table 1. The first two columns report means and standard deviations across all observations. The mean age of individuals in my data is 73. Slightly more than half of my observations are of females. Over 90% of individuals are coded by CMS as White, with 7.9% Black and 1.0% Hispanic. Over 75% self-report "Good" or better health. 24%

[10] In some years, the MCBS does not report the plan choice directly and instead reports which firm the individual has chosen, along with information about plan features and any premiums paid by the individual. I match this data to the plan data to identify each individual's choice.

[11] Massachusetts, Minnesota, and Wisconsin have alternative plan definitions; in those states I use the rate for the plan closest to Plan C. Additionally, United Healthcare did not offer plans in New York during my study period. For individuals in New York, I averaged the Plan C rates offered by all other insurers.

report having college degrees and 18% did not graduate high school. 33% receive some insurance from a current or previous employer, and 16% are Medicare-eligible for reasons other than turning 65. The second set of columns splits the data by MA enrollment. On average, MA enrollees have lower income, are less likely to be White, and have lower educational attainment. Individuals with employer-provided insurance or with ESRD are less likely to be enrolled in MA.

The third set of columns illustrates the panel nature of the data and focuses on panel observations for which the individual was enrolled in TM in the previous year – 23,278 observations total. I split the data into those who switched from TM to MA (4.3% of observations), and those who remained on TM. Those who switched are generally similar to the larger group of MA enrollees. Some differences are seen in health status – switchers from TM to MA are slightly healthier on average than the broader MA population.

3.2 Medicare Advantage plans

CMS maintains a public website which allows consumers to compare the benefits of plans offered in their location during each Open Enrollment period. I obtain the database driving this website from CMS for the years 2008 to 2015. I extract plan premiums, the Part B premium reduction (if any), an indicator for prescription drug coverage of any kind, and indicators for HMO and PPO plan types. Separately, I collect firm-level "Star" ratings, plan-level bids,[12] and county-level benchmarks from other CMS sources. I also collect average risk-adjusted TM spending at the county level.

Each year, CMS calculates and publicizes out-of-pocket cost (OOPC) estimates for each plan. CMS creates these estimates by forming a representative bundle of services used by TM enrollees in different demographic groups. CMS then calculates the out-of-pocket costs for that bundle under each plan's benefit structure, which implicitly assumes that service-level consumption patterns do not change between TM and MA. I collect the OOPC estimate for each plan for an enrollee of self-determined "Good" health between the age of 70 and 74.

I collect county-level plan coverage and enrollment counts to transform the data into the plan-

[12] I do not observe the bids directly and instead observe plan-level risk-adjusted payments, including the rebate payment, county-level enrollments and benchmark rates. Combined with the CMS payment rules described in Section 2, I can uniquely identify a bid for each plan-county-year observation.

county-year level and calculate market shares. I focus on the market for individual insurance described in Section 2 and therefore drop plans sponsored by employers and plans designed for individuals who are "dual-eligible" for Medicare and Medicaid as these plans operate under a different administrative structure. Due to CMS data restrictions, I drop plan-county-year observations with fewer than ten enrollees. For consistent presentation, I focus on plans which fall within my microdata sample area.

In what follows, it is useful to collapse the various non-price plan characteristics into a single measure reflecting the relative generosity of each plan (Lustig, 2011). Which the out-of-pocket cost estimate provides a potential index (and one used by CMS), MPTC and others have shown that other plan characteristics such as drug coverage or the restrictions required by HMO plans influence both consumer behavior and costs. As these characteristics are captured in the form of indicator variables, and likely have a significant subjective component (e.g. the disutility of a primary care visit to obtain a referral), the various characteristics do not immediately map to a single continuous measure. I use the demand system estimated by MPTC: I take the estimated coefficient on each characteristic to measure the tradeoff between alternative plan characteristics and construct a "generosity index" for each plan.[13] Appendix Table A.1 reports these parameters and the distributions of the relevant characteristics.

Table 2 presents the mean characteristics of my 50,593 plan-county-year observations separated by price quartiles. As 43% of my observations have a price of zero, I separate those observations into their own column. In general, more expensive plans are associated with increased benefits measured both by changes in the various characteristic measures and through the generosity index. More expensive plans also tend to be offered by firms with a higher Star rating, though firms across the distribution of Star ratings tend to offer plans at a variety of price points.[14]

[13]MPTC estimates coefficients for a number of additional plan features such as the plan deductible and specific copays. They interpret these coefficients as the relative value of specific financial terms of the plan, holding everything else (including the overall out-of-pocket cost) constant. As in reality these features co-move with the out-of-pocket cost estimate, I exclude them from the generosity index to avoid double-counting.

[14]It is for this reason, that I do not incorporate the Star rating directly into the generosity index as well as the difficulties caused by the changing definition of the Star rating over time. Appendix Table A.2 reports summary statistics by generosity quartile and shows that in general more generous plans are offered by firms with a higher Star rating.

3.3 Reduced form evidence of dynamic behavior

The existence of consumer inertia in the form of switching costs suggests that firms may have an incentive to engage in dynamic pricing and product design behavior: offer a low price in one period and then raise the price (or lower product quality) once some market share has been obtained (Cabral, 2016).

Figure 1 illustrates an example of this behavior by a particular plan in the data by plotting the share obtained by the plan in the previous year (i.e. the number of enrollees "locked in" to the plan's network) and the price in the current year. When the plan's share increases, its price increases and vice-versa. The correlation coefficient between the "prior-year share" and "current-year price" for this particular example is 0.731.

While this example is illustrative, it is anecdotal at best. To test for this behavior more comprehensively, let p_{jt} and g_{jt} be the price and generosity of plan j in period t, respectively. Define $\Delta p_{jt} = p_{jt} - p_{jt-1}$ and $\Delta g_{jt} = g_{jt} - g_{jt-1}$. If firms are engaging in dynamic behavior, Δp and Δg should be correlated with s_{jt-1}, the share obtained by plan j in the last period. I test for this by taking advantage of the panel nature of my plan data and estimating the coefficients of

$$\Delta p_{jt} = \beta s_{jt-1} + \text{FX} + \epsilon_{jt} \tag{2}$$

where FX stands for fixed effects and ϵ_{jt} is the unobservable term.

To estimate this equation (and the analogous expression for Δg), I restrict my attention to plans which are offered across multiple years – i.e. plans which are offered under the same CMS contract and with the same unique identifier. As firms often re-organize their plan offerings, this limits me to roughly half of my plan-county-year observations.

I report OLS estimates in Table 3. The first set of columns uses Δp as the left-hand-side variable and the second set of columns uses Δg. For each dependent variable, I report results with no fixed effects, and with combinations of firm and county fixed effects. Across specifications, there is a positive association between s_{jt-1} and Δp_{jt}; in other words, those plans that obtained a higher market share last year increased their prices more this year. The magnitude of the association is even stronger when firm and geographic fixed effects are included, suggesting that cost shocks that occur at the firm or county level are not to blame for this pattern. The results for Δg are weaker: there is a negative correlation between last year's market share and the change in generosity across years in the raw data, but the connection becomes weaker when firm fixed effects are considered.[15]

I conclude that the data are consistent with a dynamic pricing story – that firms are likely to pursue capture and harvest opportunities mainly through the price channel and possibly through the benefit design channel. Identifying the extent to which this behavior biases estimates of firm costs, requires a structural model of supply and demand which incorporates the dynamic incentives generated by inertia.

4 A Model of Dynamic Competition

In this section I introduce a model of dynamic competition in Medicare Advantage. The demand side is similar to MPTC: 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) and Aizawa and Kim (2018), I assume consumers are myopic.[16] On the supply side, multi-product firms simultaneously choose the price and generosity of their plans and solve a recursive problem that incorporates the intertemporal tradeoff generated by the switching costs.

4.1 Environment and timing

Let m denote markets (taken to be counties) and t denote periods. For each market and period, there is a set of firms $\mathcal{F}_{mt}$. Each firm f offers a set of plans $\mathcal{J}_f$; these sets are taken to be exogenous.

[15] This pattern persists when I restrict my attention to those plans with zero premium.

[16] In other words, consumers do not consider how their choice of plan today may affect the relative desirability of plans (and thus their realized utility) tomorrow. Ketcham et al. (2012) argue that some patients in Medicare Part D are able to learn from their realized drug expenditures when making future plan choices though they do not consider provider-driven inertia or learning. Relaxing this assumption would likely impose neoclassical intertemporal preferences with a discount factor close to one, though Dalton et al. (2018) argue that that model does not explain Medicare beneficiary behavior well. Nosal (2012) estimates such a demand model and finds extremely high (perhaps implausibly so) switching costs. Incorporating such a demand model into this analysis would be computationally intractable given the dynamic oliogopolistic competition on the supply side. While Goettler and Gordon (2011) estimate a model with dynamic firms and consumers, their environment is a duopoly which affords them several simplifications to firm strategies and the relevant state space.

Each plan j consists of a price p_j, generosity g_j, and unobservable quality ξ_j. Each market also has a set of consumers $\mathcal{I}_{mt}$. Each individual i has demographic characteristics drawn from $f_{mt}^{\mathcal{I}}$ which include the current enrollment in an MA plan, if any.

Each period begins with the government exogenously setting a schedule of benchmarks at the market level $\{B_m\}_t$ and Medigap insurers exogenously determining a price schedule p_{0mt}. B_m maps to plan-specific subsidies via Equation 1 – the parameters ϕ and λ are taken to be exogenous. Firms publicly observe the subsidies and the current market share distribution – i.e. the distribution $f_{mt}^{\mathcal{I}}$ – and privately observe cost shocks ω_j. They simultaneously choose their prices and product characteristics for each of their plans. Consumers then choose a single plan or the outside option of TM. Profits and the next state of the market – i.e. the updated distribution $f_m^{\mathcal{I}}t$ – are then realized.

4.2 Demand

Consumers have characteristics z_i which include an income bracket $w \in W$ and a health type h (healthy or unhealthy). Let w_i be a set of indicators which are equal to one if consumer i is a member of group w, and let h_i be one if the consumer considers themselves 'healthy.'

Consumers enter the period enrolled in plan k. I define two indicators S_{sij} to capture the effect of this previous enrollment on utility. Let S_{1ij} be equal to one if k is the outside good – I call this the *Medicare-to-MA* indicator. Let S_{2ij} – the *MA Interfirm* indicator – be one if k is offered by a different firm than j.

Let u_{ijmt} denote the consumer's utility from enrolling in a particular plan j. Dropping the market and time subscripts, the choice specific utility is given by

$$u_{ij} = \left(\alpha + \sum_w \alpha_w w_i\right) p_j + \beta_g g_j + \beta_{gh} h_i g_j + \sum_s \left(\beta_s + \beta_{sh} h_i\right) S_{sij} + \beta_z z_i + \xi_j + \epsilon_{ij}. \qquad (3)$$

The utility of the outside good varies with the price schedule of supplemental insurance $\{p_0\}$.

Let p_{0i} represent the price faced by individual i. u_{i0}, the utility of the outside good, is given by

$$u_{i0} = \beta_0 p_{i0} + \epsilon_{i0}. \tag{4}$$

I normalize the outside good by subtracting $\beta_0 p_{i0}$ from each u_{ij}.

In these equations, $\alpha + \alpha_w$ captures price sensitivity for income group w. β_g captures the impact of generosity, which may vary by health via β_{gh}. $\beta_s + \beta_{sh}$ captures the (potentially health-status dependent) switching costs through a direct utility term, similar to Handel (2013). β_z captures heterogeneous tastes for MA over TM across demographics. Following standard practice, I decompose the unobservable portion of utility into two components. ξ_j represents the portion of unobserved plan utility that is common across individuals. ϵ_{ij} represents the idiosyncratic taste of consumer i for plan j.

Following Berry et al. (1995), it is useful to rewrite the choice-specific utility into a product-level mean

$$\delta_j = \alpha p_j + \beta_g g_j + \xi_j \tag{5}$$

and an individual-specific deviation from that mean

$$\mu'_{ij} = \sum_w \alpha_w w_i p_j + \beta_{gh} h_i g_j + \sum_s (\beta_s + \beta_{sh} h_i) S_{sij} + \beta_z z_i + \epsilon_{ij}. \tag{6}$$

Let $\mu_{ij} = \mu'_{ij} - \epsilon_{ij}$. Let ϵ_{ij} be drawn independently from the Type-I extreme value distribution. The probability that consumer i chooses plan j then takes a logit form

$$q_{ij} \equiv \Pr(i \text{ chooses } j) = \frac{\exp(\delta_j + \mu_{ij})}{1 + \sum_{k \in \mathcal{J}_m} \exp(\delta_k + \mu_{ik})}. \tag{7}$$

4.3 Supply

Firms start each period with common knowledge of the distribution of consumer demographics, including enrollment status, and the per-enrollee subsidy they receive B_j.[17] Firms choose the

[17] I do not explicitly model bidding, though I account for it in my calculation of B_j.

prices p_j and product characteristics g_j, ξ_j for a fixed set of plans simultaneously. I abstract away from entry and exit decisions and from the possibility of economies of scope (i.e. the firm's profit function is additively separable across markets). I model marginal costs mc_{ij} as a function of the plan generosity g_j, individual health status h_i, market-level characteristics X_m, and unobservable component ω_j with

$$mc_{ij} = \gamma_h h_i + \gamma_g g_j + \gamma_{gh} g_j h_i + \beta_X X_m + \omega_j. \tag{8}$$

The cost "shock" ω_j consists of any component of costs which are constant across individuals and uncorrelated with the observables g_J and X_m. These costs could include the cost of advertising (Aizawa and Kim, 2018), the cost of contracting with specific hospitals or other providers (Lewis and Pflum, 2015), or changes in the cost of capital or reinsurance (Hall, 2010). As many of these factors may influence demand as well, I allow ω_j and ξ_j to be correlated. I do not explicitly model the variance in spending driven by individual-specific health shocks. As a consequence if firms are risk-neutral this marginal cost can be viewed as the expected marginal cost of enrolling person i before their individual specific health conditions are realized. Insofar that firms employ actuaries to calculate expected costs as a normal part of their product design process, this simplification is unlikely to have first-order implications.

Let Q_m be the number of people in the market. Given these costs and the demand model above, and using variables without subscripts to indicate the vector of choices available in the market, the firm's period profit is

$$\pi_f = \sum_{j \in \mathcal{J}_f} \left[\int_{i \in \mathcal{I}_m} (p_j + B_j - mc_j) q_{ij}(p, g, \xi) Q_m di \right]. \tag{9}$$

The switching costs in the demand model imply that q_{ij} depends on the current distribution of consumers across firms and TM. Let s_{ft} be an object that contains all payoff-relevant information about the distribution of consumers enrolled in any plan offered by firm f at time t.[18] Let s_{-ft} be

[18] s_{ft} can be represented with two components: the number of consumers enrolled in one of f's plans, and the distribution of consumer demographics conditional on enrollment in f. Note that since there is no switching cost incurred when switching between plans within the same firm, conditioning on the firm captures all intertemporallly-relevant factors.

the set of consumer information for all firms other than f at time t, s_{0t} be the consumer information for TM at time t, and $s_t = \{s_{0t}, s_{1t}, \cdots, s_{Ft}\}$ be the full set of consumer information at time t. I drop the time subscripts and write the firm's problem using a recursive value function formulation. Given $(p_{-f}, g_{-f}, \xi_{-f})$, the set of products offered by firms other than f, and cost shocks ω_f, firm f maximizes

$$V(s_f, \omega_f; s) = \max_{p_f, g_f, \xi_f} \pi_f(\cdot) + \beta E_\omega[V(s'_f, \omega'_f; s')]. \tag{10}$$

In this equation, the firm chooses prices and characteristics for all of its products simultaneously. s'_f and s' indicate the next period's "share" objects as a function of the choices of all firms. The expectation is over the cost shocks. A Markov Perfect Nash Equilibrium is defined in this context as a set of policy functions $\{p_f, g_f, \xi_f\}_{f \in \mathcal{F}}$ for each firm as a function of their cost shock and the industry state such that each policy function solves the maximization problem in Equation 10 when all other firms play according to their policy function and the evolution of s' is given by Equation 7.

5 Taking the Model to the Data

The model of the previous section captures the tradeoffs faced by consumers and MA firms. The key challenge in estimating its parameters lies in the intertemporal incentives embedded in Equation 10 and the nature of the payoff-relevant state for firms. Broadly speaking, previous efforts to estimate dynamic oligopoly models have largely imposed separability between play in the period game and the intertemporally-relevant decision. For example, the "quality ladder" model of Ericson and Pakes (1995) can be seen as a static game between firms with differentiated quality levels with an intertemporal linkage driven by entry, exit, and the ability of firms to spend some of their period profits to probabilistically raise their quality level for tomorrow's static game. Since the model lacks credit constraints, the prices set by firms each period are the same as those in a completely static differentiated products environment. As a consequence, the parameters of this game can be estimated from data in stages: the demand and marginal cost parameters from the period game can be inferred from market outcomes without any dynamic concerns. The intertemporally-relevant parameters (such as the cost of investment and the scrap value of firms) can then be identified with patterns of entry, exit, and the evolution of firms' states over time. Alternatively, one may estimate reduced-form policy functions and laws-of-motion for state variables and then use the optimality of behavior in equilibrium to back out structural parameters (Hotz and Miller, 1993, Bajari et al., 2007, Aguirregabiria and Mira, 2007, Pakes et al., 2007).

These approaches run up against two issues in this context. First, there is no separability between the actions in the period game and the evolution of the state space: the firms' actions determine the distribution of consumers that choose to purchase various products, which is precisely the payoff-relevant state variable for the next period. This implies that marginal cost parameters cannot be estimated without considering the dynamic incentives. Second, not only is the payoff-relevant state space continuous, it is also not separable across firms: if one firm manages to obtain a higher share of (for instance) healthy individuals, that must come at the cost of some other firm's share. This adds considerable complication to any attempt to model policy functions and the evolution of the state space in a reduced-form way.[19]

A third estimation approach uses moment inequalities derived from the assumption of optimal play to achieve partial identification (Pakes et al., 2015). However, using this approach in a dynamic setting requires finding sequences of deviations from observed play which only temporarily perturb the state space from its observed behavior so that a finite number of terms appear in the moment condition. For example, Holmes (2011), studying the growth of Walmart, considers single-position swaps in the order in which stores were opened. In this case it is not certain that a sufficient number of such sequences of deviations exist to achieve identification.[20]

I pursue a two-step approach for estimating the parameters of the model. In the first step, I take advantage of the assumption of consumer myopia and estimate demand parameters using

[19]It is possible to re-write the problem posed in Section 4 in terms of the frequency distribution of consumers instead of the probability distribution. This reduces, but does not eliminate, the interconnected nature of the state space, and also introduces the possibility of errors-in-shares discussed by Gandhi et al. (2013).

[20]To see this, note that the dimension of the payoff-relevant demographic characteristics can easily exceed the number of control variables, even when considering all controls across all firms simultaneously. If consumers are homogenous so that the relevant state has dimension equal to the number of firms (i.e. each firm's share), the Berry (1994) inversion implies that any unilateral deviation in one period may be "repaired" in the next period through appropriately chosen deviations. However, those repairs may involve deviations in the play of multiple agents, violating the assumptions of Pakes et al. (2015).

a relatively "off-the-shelf" approach. In the second step, I estimate marginal cost parameters by repeatedly solving a reduced-information version of the model and applying the method of moments.

Previous work has demonstrated that solving models of dynamic oligopoly in the sense of finding a Markov Perfect Nash Equilibrium is computationally difficult (Pakes and McGuire, 1994, 2001), even in cases with a single-dimension discrete state variable. To achieve computational tractability, I follow recent trends in the literature towards relaxing the Markov Perfect solution concept (Fershtman and Pakes, 2012). In particular, I adopt the Oblivious Equilibrium (OE) approach of Weintraub et al. (2008) and extend it to this continuous-state-space "market share" game. Qualitatively, the OE approach models firms as playing against the "average" competitive state of the world, rather than tracking the exact period-to-period state of each of its competitors. In doing so, the model reduces to a series of single-agent problems which can be solved through value function iteration. Weintraub et al. (2008) apply their solution concept to the quality ladder game and use the average number of firms on each step of the ladder as their information set. The key insight of my extension is that the information set of firms in OE need only contain what is necessary to compute a consistent estimate of profits, and need not contain any direct information about the state of its competitors. In this way, my work runs in parallel to other solution concepts which limit the information set of firms in some way such of that of Ifrach and Weintraub (2016).

5.1 Estimating demand parameters

I estimate the parameters of the demand system following the approach of MPTC. First, I estimate the parameters that define μ'_{ij} via maximum likelihood estimation. I then estimate the parameters that define δ_j with an instrumental variables approach.

Let $\theta_I = \{\alpha_w, \beta_{gh}, \beta_s, \beta_{sh}, \beta_z, \beta_0\}$ be the set of parameters which form μ'_{ij} and u_{i0}. For a given candidate value $\tilde{\theta}_I$, I use the Berry (1994) inversion and find the unique set of product-level fixed effects $\delta_j(\tilde{\theta}_I)$ that match the product shares predicted by the model to the observed shares in the data. Let C_{ij} be an indicator which is equal to one if person i chose product j in the data. The likelihood function for i is then

$$L_i(C_{ij}; \tilde{\theta}_I, \delta(\tilde{\theta}_I)) = \Pi_j q_{ij}^{C_{ij}}, \tag{11}$$

I recover parameters by maximizing the sum of the log-likelihood functions across individuals, weighted by the MCBS sample weights.[21] At the point estimate $\hat{\theta}_I$ I store the unique $\hat{\delta}_j$ and regress it on observable product characteristics according to Equation 5.

Since ξ_j is observed by firms and consumers, it may be correlated with the price of the plan. To identify the coefficients α_0 and β_g, I assume that ξ_j is uncorrelated with the observed generosity g_j. This implies both that β_g is identified and that instruments formed from g_j are valid for price. I adopt the summation instruments of Berry (1994). Furthermore, following the logic of Hausman et al. (1994) I use average prices in other counties as additional price instruments under the assumption that some component of marginal cost is constant across firms in different geographic markets (such as different counties within the same state).[22] Since plans are often offered at the same price in geographically contiguous counties, I calculate this instrument using "non-contiguous counties" – that is, for a plan in a county, I calculate the average price of plans in counties of the same state which do not share borders with the county under consideration.

Table 4 reports the results of my demand estimation procedure. The first panel reports parameter estimates and standard errors across both the maximum likelihood and regression stages. Generally speaking, the parameters are similar to the estimates reported by Miller et al. (2019). The coefficient on generosity, 0.770, is slightly less than one, which is to be expected as the index is formed under the assumption that the coefficients are perfect substitutes close to their mean values which is unlikely to hold precisely throughout the distribution of product characteristics. The coefficient on price, -2.02, is slightly lower than the -3.11 reported by MPTC, though the 95% confidence intervals overlap. I interpret these two results as evidence of the tradeoff between flexibility and simplicity – my two observables do a good job capturing the variety of plans offered, but of course a more flexible demand system can tease out more dimensions of differentiation. As a

[21]In practice, data limitations in certain years of the MCBS prevent me from recovering a unique choice for each person. In these cases, I eliminate as many options as I can given what I observe about the individual's selection and take draws from the conditional distribution of the remaining plans.

[22]This is consistent with some components of X_m in Equation 8 being constant across markets.

consequence of this difference in coefficients, the mean elasticity is also slightly smaller in absolute value than those reported by MPTC though still well above 1.

5.2 An Oblivious Equilibrium approach to estimating supply parameters

With the demand parameters in hand, I turn to the problem of estimating marginal costs. To do so, I rewrite the model in a simplified form consistent with the OE approach. I begin by considering the information that firms need in order to calculate profits. First, note that the denominator of share function can be written as the sum of components from plans offered by a given firm and the sum of components from plans offered by other firms.

$$q_{ij} = \frac{\exp(u_{ij})}{1 + \sum_{k \notin \mathcal{J}_f} \exp u_{ik} + \sum_{k \in \mathcal{J}_f} \exp u_{ik}}.$$

The individual $\exp u_{ik}$ terms will vary across time and plans as different firms experience different cost shocks and obtain different market shares. However, if there are many firms the sum $\sum_{k \notin \mathcal{J}_f} \exp u_{ik}$ will have lower variance from year to year. I therefore define the "expected competitive pressure" for consumer i, r_i, as

$$r_i = E_{f,t}\left[\sum_{k \neq J_f} \exp(u_{ik})\right]. \tag{12}$$

In this equation, the expectation is taken across time, over all of the firms in a market, and is specific to individual i. The sum is taken over all products offered by firms other than the one under consideration. In other words, r_i represents the expected extent to which the plans offered by competitors pressure consumer i *not* to choose a plan from firm f. In particular, a firm can use r_i to calculate an "oblivious" approximation to the share function via

$$q_{ij}^O = \frac{\exp(u_{ij})}{1 + r_i + \sum_{k\ inJ_f} \exp u_{ik}}. \tag{13}$$

Define the information set I_f as the set of r_i for each consumer, the set of individual-specific components $\beta_z z_i$ to the $\exp(u_{ij}$ term, the firm's subsidies B_j, and any other relevant market-level variables. Let s_{fi}^{OE} be the current market share of firm f of each consumer i; define s_f^{OE} as the vector of shares for firm f. Then the oblivious approximation to the value function becomes

$$V^O(s_f^O, \omega_f; I_f) = \max_{p_f, g_f, \xi_f} \sum_{j \in \mathcal{J}_f} \left[\int_i (p_j + B_j - mc_j) q_{ij}^O Q_m di \right] + \beta E_\omega [V^O(s_f^{O'}, \omega_f'; I_f)]. \tag{14}$$

An Oblivious Equilibrium in this context is a set of policy functions $\sigma = \{p_f^O, g_f^O, \xi_f^O\}$ and an information set I_f such that the policy functions solve the maximization problem when the firm faces the information set and the share $s_f^{O'}$ evolves according to Equation 13, and the information set is generated by firms who play according to the policy functions in the sense that the r_n are generated by the choices of firms.

It's important to note that this definition of equilibrium does not eliminate the dynamic capture/harvest cycle. While the firms are "playing against" some average level of competition which is constant over time, their own strategies can freely vary with their current enrollment. One can interpret this behaviorally as follows: consider an individual firm. This firm knows that in any given period, some of its competitors will be capturing, and others will be harvesting. On average, the proportion of firms engaged in each behavior is constant, and so the individual firm does not need to worry about exactly which firm is engaged in which behavior each period. Against this backdrop, however, the firm knows its own current market share and can decide whether or not to capture or harvest based on that information. I return to this issue below after discussing the empirical implementation of the model.

The benefit of defining equilibrium in this way is that given some information set I_f, the firm problem can be solved by single-agent value function iteration. One can then easily simulate cost shocks and actions to generate r_i and a new I_f. Iterating between these two steps can lead to an equilibrium for some parameter vector θ^S (Weintraub et al., 2010).[23]

Suppose that I can, for any guess of parameters $\tilde{\theta}^S$, find equilibrium information sets and strategies $\{\sigma(\tilde{\theta}^S), I_f(\tilde{\theta}^S)\}$. I use these equilibrium strategies to create a method of moments estimator in the following way. Using notation from Pakes et al. (2015), my model creates an approximation

[23]This process is not a contraction mapping and so therefore some care must be taken during implementation.

π^O to the true period profit function π in the following sense:

$$\frac{\partial \pi_{fj}}{\partial X_{fj}} = \frac{\partial \pi^O_{fj}}{\partial X_{fj}} + \nu_{1fj} \tag{15}$$

In this expression, X_{fj} represents all of the data I observe about firm f and plan j, including market-level characteristics. 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. In particular, in addition to the difference between π^O and π generated by differences between the actual competitive pressure and the expected competitive pressure, firms in my model do not observe each other's cost shocks. As the policy functions $\{p, g, \xi\}$ are the result of maximizing the profit function, any error in the oblivious profit function $\pi^O(\cdot)$ will contribute to the error in the oblivious policy functions $\{p^O, g^O, \xi^O\}$ as well.

I form an estimator by imposing the conditional moment restriction

$$E[\nu_{1f}(\theta_0)|I_f] = 0 \, \forall f. \tag{16}$$

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: $E[\sigma^{DATA} - \sigma^{MODEL}|I_f] = 0$. Let $g_i(\theta) = h_i \cdot (\sigma_i^{DATA} - \sigma_i^{MODEL}(\theta))$ and $g_n(\theta) = \frac{1}{n}\sum_i g_i(\theta)$, where h_i is a vector of instruments for observation i. Then given some weight matrix W, define the objective function $J(\theta, W) = n \cdot g_n(\theta)' W g_n(\theta)$. I form an estimate $\hat{\theta}^S$ by minimizing J.

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 respect to each parameter when the derivative is evaluated at θ_0 with

$$\boldsymbol{H}_f = E\left[\frac{\partial \boldsymbol{\nu}_{1f}(\theta_0)}{\partial \boldsymbol{\theta}} \middle| I_f \right]. \tag{17}$$

I calculate approximations to these optimal instruments following Berry et al. (1999).

5.2.1 Identification

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:

$$\mathbf{0} = \sum_j \frac{\partial \pi_j}{\partial \boldsymbol{\sigma}_f} + \beta E\left[\frac{\partial V}{\partial \boldsymbol{\sigma}_f}\right] \tag{18}$$

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

5.3 Further details

To implement the approach described above, I make a number of further simplifications. First, I condition all demographic characteristics on healthy/unhealthy status, so that each market has a distribution of healthy individuals and a distribution of unhealthy individuals, represented by the MCBS draws. Due to the size of the switching costs, I form competitive pressures for each of these draws under three enrollment scenarios relative to the perspective of any particular firm: "on TM," "on a competitor's plan," and "on my plan." As a consequence, firms in the implemented estimation routine have an information set that contains six vectors: competitive pressures for MCBS draws that are healthy or unhealthy under each of the three enrollment scenarios. This maintains the pattern of correlation between health, income, and inertia. The state space becomes two-dimensional: the market share that each firm has of healthy and unhealthy people. Furthermore, under the assumption that firms are playing equilibrium strategies on average, I calculate the expected competitive pressure vectors using the MCBS draws for each market and the demand parameter estimates. This speeds up computation by several orders of magnitude, as I need only solve the value function once for each market for each guess of parameters.

The innermost loop in the estimator involves maximizing the profit function repeatedly in order to update the current guess of continuation values.[24] For timely execution, it is necessary to restrict the choice set available to firms as much as possible. Table 2 shows that 43% of plans are offered with no annual premium. Furthermore, the modal firm offers two plans. As a consequence I restrict firms to offering two plans, one of which has zero premium. If firms offer more than two plans, I share-weight the price and use the Berry inversion to weight the generosity index in such a way that the predicted share of the combined plan at the demand parameters estimated in Table 4

[24] This maximization is performed on the order of 100 billion times in the course of estimation and is implemented with a high degree of parallelism. Contact me for details.

equals the sum of the shares of the individual plans; thus this simplification does not distort the raw data.[25] I assume cost shocks are distributed normally and discretize them for estimation.

As I am interested in the degree to which insurers might be differentially incentivized to attract consumers of different health given the potential inadequacies of the risk-adjustment system (Brown et al., 2014, Newhouse et al., 2015), I must account for differences in both the potential costs and benefits of serving healthy and unhealthy customers. As discussed in Section 2, CMS uses a number of demographic and diagnostic variables collected by insurers post-enrollment to assign a risk score to each enrollee. While I do not observe all of the variables used by CMS in the MCBS, I use the CMS methodology to obtain an average risk score for healthy and unhealthy individuals in each county. I use these scores to adjust both the expected TM costs and the MA subsidies for each plan across the unhealthy/healthy dimension.

To summarize, though I make simplifications for tractability, my estimator still captures crucial first-order features of the market: consumers have different price sensitivities – which affect their demand – and different health statuses which affect their demand and their costs to firms. The correlation between these consumer demographics (and thus the ability of firms to drive self-selection into particular plans) comes directly from the MCBS draws for each market – thus my estimator captures the spectrum from poor, sick geographies to wealthy, healthy areas. Furthermore, the reduced-information model still generates capture and harvest behavior. Figure 2 illustrates example firm policy functions. The left graph shows the generosity of the zero-premium plan decreasing as the prior-year share increases, while the right graph illustrates the price of the positive-premium plan increasing as the prior-year share increases.

The final form of the estimator follows a two-step GMM approach. I start by setting the weight matrix equal to the identity matrix. I calculate approximations to the optimal instruments and minimize the GMM objective function to find interim parameters. I then update the weight matrix using the sample variance-covariance matrix at these interim parameters and similarly recalculate

[25] This effectively combines the generosity index and the ξ_j term of the demand system into a single product quality index that is observed by firms and consumers. In a previous draft of this article, I used the share-weighted generosity index only and obtained qualitatively similar results. I prefer the approach used here due to the consistency with the observed market shares.

the approximate optimal instruments. I minimize the updated GMM objective function to form my final estimates. As this objective function is non-linear, I start the procedure from several starting values to increase confidence that I have found the global minimum. I calculate standard errors following the asymptotic theory of Hayashi (2000, Ch. 7).

6 The Costs of Medicare Advantage

Table 5 reports GMM estimates of Equation 8. Private firm costs are highly correlated with fee-for-service costs. This is unsurprising given that prior work has found that TM reimbursement rates are often used as a baseline for payment rates from private insurers (Berenson et al., 2015), though private insurers may pay slightly less for some services (Baker et al., 2016). While the point estimates on the generosity-related terms are as expected, only the quadratic term is precisely estimated. The enrollee's age does not enter significantly, suggesting that the risk adjustment system provides appropriate incentives along that dimension. The log of the median household income enters significantly with a negative sign, though the magnitude is small. Since FFS spending is correlated with household income,[26] I interpret this coefficient to suggest that private firms may have some advantage over the government at controlling costs in places where incomes are high. Indeed, the current TM physician payment system includes an adjustment for the local cost-of-living at the service level (Zuckerman et al., 1990, Bai and Anderson, 2017) whereas private insurers may provide cost-of-living adjustments that do not vary with the number of services provided i.e. at the episode level or through hiring salaried providers.[27]

Table 6 combines these estimates to construct the marginal cost of each plan for healthy and unhealthy enrollees, and compares those costs to average TM costs. I focus on zero-premium plans as they effectively compete with TM on price. On average, the marginal cost of a zero-premium MA plan for a healthy enrollee is $7,784, compared to an average TM cost of $8,080. The marginal cost of an unhealthy enrollee for the same plan is $19,280, compared to an average TM cost of

[26] The correlation coefficient between log income and log FFS spending for the markets in my sample is 0.371.

[27] As part of continuing efforts to control costs, in recent years TM has implemented a strategy that bundles payments for some diagnoses at the episode level (Cutler and Ghosh, 2012, Press et al., 2016).

\$17,250. The most straightforward conclusion is thus that MA firms may be slightly more efficient than TM at providing care to healthy enrollees, but face higher costs for unhealthy enrollees.

This interpretation, however, requires some nuance. First, since I am using a revealed-preference approach, a more accurate statement might be "under the assumption of differentiated product competition, MA firms on average *behave as though* their costs of providing care to unhealthy individuals are greater than TM's costs." Second, while the average TM costs I report are accounting costs calculated from county-level CMS data, the marginal costs I estimate are economic costs. If insurers face substantial opportunity costs from providing care to an MA enrollee rather than a member of an employer-sponsored plan, the accounting costs of MA plans may be more in line with TM costs. Indeed, Dafny (2010) finds that insurers are able to exert some market power in their negotiations with employers (whereas MA subsidy rates are set independently of the actions of any individual insurer), suggesting that employer-sponsored-enrollees may be more profitable and thus that significant opportunity costs may exist. Finally, zero-premium plans do occasionally advertise benefits which TM does not offer, though they also come with restrictions which TM enrollees do not face. These differences in benefit design may be non-neutral in the sense that a substantial portion of the difference between the estimated TM and MA costs may be related to supplemental benefits – something which can not be explicitly tested with available data.

Though my method for splitting costs and payments across the healthy/unhealthy categorization matches the CMS method, any flaws in the CMS methodology will flow through to these estimates (Chen et al., 2015). In particular, Geruso and Layton (2015) and Newhouse et al. (2019) argue that MA firms are able to "upcode" the diagnoses of their enrollees and receive a risk score that is between 6%-16% higher than the enrollee "should" receive. Since the risk scores directly affect MA payments, this result implies that firms may receive subsidies that are roughly 10% higher than those I calculate. Given the structure of the profit function in my model (Equation 9), unobserved increased payments would be interpreted by my estimator as decreased costs. As my estimated MA costs for healthy individuals are only 3.7% lower than TM costs, these results cast doubt on any conclusion that firms are more efficient for healthy individuals.[28]

[28] I have also explored specifications which do not explicitly include TM costs as a measure of MA costs and obtained

Table 6 also reports average marginal costs for positive premium plans. The average premium is \$513, and the average plan provides a generosity of 3.69 (relative to an average generosity of 2.55 for zero-premium plans). The cost of these plans (relative to the zero-premium plans) is on average \$436 higher for healthy individuals, and \$500 higher for unhealthy individuals, suggesting that the passthrough rate from premia to benefits is between 85% and 98%. Given CMS's actuarial rules – i.e. that premia must be 'justified' by expenditures on services (CMS Office of the Actuary, 2017) – this result is sensible.

The estimator I develop differs from the static first-order condition approach through the inclusion of the dynamic incentives driven by the switching costs. Table 7 compares my results to the standard approach introduced by Berry et al. (1995). For the purposes of comparison, I weight costs by the ratio of healthy to unhealthy people in the market to obtain a "risk 1.0" cost. On average, the statically estimated costs are 13% lower for zero-premium plans (\$9,335 versus \$10,728) and 9.3% lower for positive-premium plans. The estimated costs are higher under the dynamic approach than under the static approach for 84% of zero-premium plans and 67% of positive premium plans.

After adjusting for observed risk, the average TM cost is \$10,111, or 7.7% higher than MA costs under the static estimator or 6.1% lower than MA costs under the dynamic estimator. In other words, the dynamic estimator flips the sign of the qualitative comparison; even when considering the potential upcoding behavior discussed above, it is reasonable for those using static estimators to conclude that TM costs are largely lower than MA costs, as others have in the literature (Town and Liu, 2003, Dunn, 2010, Curto et al., 2018, Miller et al., 2019).

The joint distribution of prior-year shares and current-year prices is illustrated in Figure 3 – the histogram on the left focuses on zero-premium plans while the scatterplot on the right illustrates the positive-premium plans in the data. The difference between the static and dynamic estimates is driven by the the different way in which the two estimators perceive the data. The static estimator uses elasticities, the subsidy received by the firm from the government, and the assumption of profit maximization to infer marginal costs. When the static estimator sees a low (high) price, it infers low (high) costs. In addition to those features of the data, the dynamic estimator also incorporates

qualitatively similar results.

information about the firm's prior-period market share. When it sees low prices together with a low share, it concludes that the price is driven by capture incentives, and so it infers that costs are higher than the static estimator finds. When it sees low prices together with a high share, it concludes that the firm is harvesting and so therefore the marginal cost may be lower than the static estimator finds. As the preponderance of the data consists of firms with low prices and low shares (75% of the firm-county-year observations in my estimation sample enter the period with less than 2% market share), the dynamic estimator infers higher costs.

If, instead, the data consisted of firms with larger shares, the dynamic estimator could find lower costs. To see this, I alter the dataset by dividing each firm's share by the sum of the shares across firms. In other words, I scale up each firm's share so that the sum of the shares across firms within each market is one, while maintaining the relative share across firms and the prices and generosities of the various plans as seen in the original dataset. I re-estimate the model with this altered dataset and report the results in Column (2) of Table 8. While the estimate of the marginal cost of healthy people is about $300, or 4%, higher than the original estimate, the estimate of the marginal cost of unhealthy people is about $3,000, or 15%, lower. Aggregating across health status, the marginal cost of the zero-premium plan is $9,920, closer to the static estimates reported in Table 7. Column (3) takes this further by keeping only the top two firms by share in each market (if more than two are present) and redistributing the shares accordingly. The estimated costs for healthy people rise slightly once more, but the estimated costs for unhealthy people fall by roughly 20% relative to Column (2). Aggregating across health once more, the marginal cost of the zero-premium plan is $9,191, below the static estimate.[29]

7 Conclusion

For some years the industrial organization literature has been concerned with identifying sources of intertemporal incentives and the extent to which firms engage in dynamic pricing (Chen and

[29]Parameter estimates for these specifications are in Appendix Table A.3. The standard errors are larger, which I interpret to mean that the estimator "rejects" the notion that the pattern of prices and generosities seen in the data is a function of relative shares alone.

Pearcy, 2010). Some dynamic pricing behaviors based on inertia are obvious: advertisements for mobile telephone or high-speed internet service with a temporary "promotional" rate or offering of hardware for new subscribers who sign a long-term contract are ubiquitous (Grzybowski and Pereira, 2011).[30] Yet firms who do not impose explicit switching costs through contract termination fees can engage in capture and harvest behavior as well. MA plans must accept all applicants during the open enrollment period each year and may not price discriminate. Still, reduced-form evidence suggests that these firms engage in dynamic behavior.[31] I show that this behavior biases estimates of marginal cost based on the first-order conditions for profit-maximization by estimating a dynamic model of differentiated products oligopoly.

This analysis is particularly relevant as MA was sold to policy-makers in part based on the idea that private firms may be able to provide the same level of coverage at a lower cost than TM could achieve. With the two systems in place, one public, one partially-private, much political rhetoric in the United States has focused on future path of the Medicare system, particularly as the U.S. spends more on health care per-capita than any other country, and yet many outcomes consistently rank below other industrialized countries (Anderson and Poullier, 1999, Anderson et al., 2005, Papanicolas et al., 2018). While some have proposed a "Medicare-for-all" approach in which non-seniors could buy-in to the Medicare system, others have argued for eliminating TM and delivering legally required Medicare benefits through MA alone (Morone, 2017). Overshadowing this debate is the concern that, under current policy, Medicare expenditures are expected to rise to 6% of GDP by 2043 and revenues are not projected to keep pace (Boards of Trustees, 2019).

I find the costs of zero-premium MA plans are on average 12% higher than TM's costs for unhealthy individuals. MA costs for healthy individuals are 3.7% lower than TM's costs, though if firms are able to engage in upcoding, this apparent advantage may be spurious. Furthermore, I show that this result is driven by the distribution of shares in the data – if I redistribute shares my model produces lower cost estimates.

[30] For examples see `https://www.ispot.tv/ad/IOUV/t-mobile-unlimited-new-year-new-phone` or `https://www.ispot.tv/ad/oYr4/comcast-xfinity-keeping-up-2999-a-month`.

[31] This point is similar to the theoretical work of Doraszelski and Markovich (2007), who show that advertising alone can create persistent competitive advantages under either "awareness" or "goodwill" paradigms.

These results support a view of Medicare as a relatively efficient way to provide medical services to seniors in the United States, particularly those in poor health. However, this alone does not indicate that a “Medicare-for-all” system would be similarly effective for the working-age population. For one, as private insurers have been working to find efficiencies for that population for much longer, it is possible that their relative costs are lower for that population. Additionally, while the Medicare system does not engage in large-scale rationing, it is not clear that access would be ensured under a Medicare expansion. Finally, given that MA enrollees are not a random sample of all Medicare-eligible individuals, comparisons of health outcomes are difficult at best; it is possible that while private plans may spend more than TM, private plans also deliver some increased quality. A full evaluation of TM and MA, therefore, must consist not only of comparisons of cost, but also measures of value delivered to consumers either through subjective perception (i.e. consumer surplus) or through health outcomes.

References

Jason Abaluck and Jonathan Gruber. Choice inconsistencies among the elderly: Evidence from plan choice in the medicare part d program. *American Economic Review*, 101(4):1180–1210, June 2011. doi: 10.1257/aer.101.4.1180. URL `http://www.aeaweb.org/articles?id=10.1257/aer.101.4.1180`.

Jason Abaluck and Jonathan Gruber. Evolving choice inconsistencies in choice of prescription drug insurance. *American Economic Review*, 106(8):2145–84, August 2016. doi: 10.1257/aer.20130778. URL `http://www.aeaweb.org/articles?id=10.1257/aer.20130778`.

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 and risk selection in health insurance markets. *American Economic Review*, 108(3):828–67, March 2018. doi: 10.1257/aer.20151485. URL `http://www.aeaweb.org/articles?id=10.1257/aer.20151485`.

Gerard F Anderson and Jean-Pierre Poullier. Health spending, access, and outcomes: Trends in industrialized countries: Even with managed care and government efforts to control growth, us health spending per capita grew more rapidly in the 1990s than did spending in the average industrialized country. *Health Affairs*, 18(3):178–192, 1999.

Gerard F Anderson, Peter S Hussey, Bianca K Frogner, and Hugh R Waters. Health Spending In The United States And The Rest Of The Industrialized World. *Health Affairs*, 24(4):903–914, 2005. doi: 10.1377/hlthaff.24.4.903. URL `https://doi.org/10.1377/hlthaff.24.4.903`.

Ge Bai and Gerard F Anderson. Variation in the ratio of physician charges to medicare payments by specialty and region. *Jama*, 317(3):315–318, 2017.

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.

Laurence C Baker, M Kate Bundorf, Aileen M Devlin, and Daniel P Kessler. Medicare Advantage Plans Pay Hospitals Less Than Traditional Medicare Pays. *Health affairs (Project Hope)*, 35(8): 1444–1451, aug 2016. ISSN 1544-5208 (Electronic). doi: 10.1377/hlthaff.2015.1553.

Robert A Berenson, Jonathan H Sunshine, David Helms, and Emily Lawton. Why medicare advantage plans pay hospitals traditional medicare prices. *Health Affairs*, 34(8):1289–1295, 2015.

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.

David Besanko, Ulrich Doraszelski, and Yaroslav Kryukov. The economics of predation: What drives pricing when there is learning-by-doing? *American Economic Review*, 104(3):868–97, March 2014. doi: 10.1257/aer.104.3.868. URL http://www.aeaweb.org/articles?id=10.1257/aer.104.3.868.

Boards of Trustees. The 2019 annual report of the boards of trustees of the federal hospital insurance and federal supplementary medical insurance trust funds, 2019. URL https://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/ReportsTrustFunds/Downloads/TR2019.pdf.

Jason Brown, Mark Duggan, Ilyana Kuziemko, and William Woolston. How does risk selection respond to risk adjustment? new evidence from the medicare advantage program. *American Economic Review*, 104(10):3335–64, October 2014. doi: 10.1257/aer.104.10.3335. URL http://www.aeaweb.org/articles?id=10.1257/aer.104.10.3335.

Luís Cabral. Dynamic pricing in customer markets with switching costs. *Review of Economic Dynamics*, 20:43–62, 2016. ISSN 1094-2025. doi: https://doi.org/10.1016/j.red.2015.09.002. URL http://www.sciencedirect.com/science/article/pii/S1094202515000563.

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

Jing Chen, Randall P Ellis, Katherine H Toro, and Arlene S Ash. Mispricing in the medicare advantage risk adjustment model. *Inquiry : a journal of medical care organization, provision and financing*, 52, 2015. ISSN 1945-7243 (Electronic). doi: 10.1177/0046958015583089.

Yongmin Chen and Jason Pearcy. Dynamic pricing: when to entice brand switching and when to reward consumer loyalty. *The RAND Journal of Economics*, 41(4):674–685, 2010.

CMS. Affordable care act update: Implementing medicare cost savings. Technical report, August 2010. URL https://web.archive.org/web/20130418015959/http://www.cms.gov/apps/docs/aca-update-implementing-medicare-costs-savings.pdf.

CMS Office of the Actuary. Actuarial bid training–2017, 2017. URL https://www.cms.gov/Medicare/Health-Plans/MedicareAdvtgSpecRateStats/BidTraining2017.html.

John F Cogan, R Glenn Hubbard, and Daniel P Kessler. Making Markets Work: Five Steps To A Better Health Care System. *Health Affairs*, 24(6):1447–1457, 2005. doi: 10.1377/hlthaff.24.6.1447. URL https://doi.org/10.1377/hlthaff.24.6.1447.

Vilsa Curto, Liran Einav, Jonathan Levin, and Jay Bhattacharya. Can health insurance competition work? evidence from medicare advantage. Working Paper 20818, National Bureau of Economic Research, November 2018. URL http://www.nber.org/papers/w20818.

David M Cutler and Kaushik Ghosh. The potential for cost savings through bundled episode payments. *New England Journal of Medicine*, 366(12):1075–1077, 2012.

Leemore S. Dafny. Are health insurance markets competitive? *American Economic Review*, 100 (4):1399–1431, September 2010. doi: 10.1257/aer.100.4.1399. URL http://www.aeaweb.org/articles?id=10.1257/aer.100.4.1399.

Christina M Dalton, Gautam Gowrisankaran, and Robert Town. Salience, myopia, and complex dynamic incentives: Evidence from Medicare Part D. Technical report, National Bureau of Economic Research, 2018.

Ulrich Doraszelski and Sarit Markovich. Advertising dynamics and competitive advantage. *The RAND Journal of Economics*, 38(3):557–592, 2007.

Abe Dunn. The value of coverage in the medicare advantage insurance market. *Journal of Health Economics*, 29(6):839–855, 2010. ISSN 0167-6296. doi: https://doi.org/10.1016/j.jhealeco.2010.08.005. URL http://www.sciencedirect.com/science/article/pii/S0167629610001141.

R P Ellis, G C Pope, L Iezzoni, J Z Ayanian, D W Bates, H Burstin, and A S Ash. Diagnosis-based risk adjustment for Medicare capitation payments. *Health care financing review*, 17(3):101–128, 1996. ISSN 0195-8631 (Print).

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].

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.

Chaim Fershtman and Ariel Pakes. Dynamic Games with Asymmetric Information: A Framework for Empirical Work. *Quarterly Journal of Economics*, 127(4):1611–1661, 2012.

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.

Michael Geruso and Timothy Layton. Upcoding: Evidence from medicare on squishy risk adjustment. Working Paper 21222, National Bureau of Economic Research, May 2015. URL http://www.nber.org/papers/w21222.

Ronald L. Goettler and Brett R. Gordon. Does amd spur intel to innovate more? *Journal of Political Economy*, 119(6):1141–1200, 2011. doi: 10.1086/664615. URL https://doi.org/10.1086/664615.

Lukasz Grzybowski and Pedro Pereira. Subscription choices and switching costs in mobile telephony. *Review of Industrial Organization*, 38(1):23–42, 2011.

Mark A Hall. The three types of reinsurance created by federal health reform. *Health Affairs*, 29 (6):1168–1172, 2010.

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.

Jerry Hausman, Gregory Leonard, and J. Douglas Zona. Competitive analysis with differenciated products. *Annales d'conomie et de Statistique*, 34:159–180, 1994. ISSN 0769489X, 22726497. URL http://www.jstor.org/stable/20075951.

Fumio Hayashi. Econometrics. 2000. *Princeton University Press. Section*, 1:60–69, 2000.

Igal Hendel and Aviv Nevo. Measuring the implications of sales and consumer inventory behavior. *Econometrica*, 74(6):1637–1673, 2006. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/4123086.

Kate Ho, Joseph Hogan, and Fiona Scott Morton. The impact of consumer inattention on insurer pricing in the Medicare Part D program. *The RAND Journal of Economics*, 48(4):877–905, 2017. doi: 10.1111/1756-2171.12207. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1756-2171.12207.

Thomas J. Holmes. The diffusion of wal-mart and economies of density. *Econometrica*, 79(1): 253–302, 2011. ISSN 1468-0262. doi: 10.3982/ECTA7699. URL http://dx.doi.org/10.3982/ECTA7699.

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.

Bar Ifrach and Gabriel Y. Weintraub. A Framework for Dynamic Oligopoly in Concentrated Industries. *The Review of Economic Studies*, 84(3):1106–1150, 09 2016. ISSN 0034-6527. doi: 10.1093/restud/rdw047. URL https://doi.org/10.1093/restud/rdw047.

Jonathan D. Ketcham, Claudio Lucarelli, Eugenio J. Miravete, and M. Christopher Roebuck. Sinking, swimming, or learning to swim in medicare part d. *The American Economic Review*, 102 (6):2639–2673, 2012. ISSN 00028282. URL `http://www.jstor.org/stable/41724667`.

Matthew S Lewis and Kevin E Pflum. Diagnosing hospital system bargaining power in managed care networks. *American Economic Journal: Economic Policy*, 7(1):243–74, 2015.

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

Alexander MacKay and Nathan H Miller. Estimating models of supply and demand: Instruments and covariance restrictions. Technical report, 2019. URL `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3025845`.

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

Thomas G McGuire, Joseph P Newhouse, and Anna D Sinaiko. An economic history of medicare part c. *The Milbank Quarterly*, 89(2):289–332, 2011.

Keaton S Miller, Amil Petrin, Robert Town, and Michael Chernew. Optimal managed competition subsidies. Working Paper 25616, National Bureau of Economic Research, February 2019. URL `http://www.nber.org/papers/w25616`.

James A Morone. How to Think about Medicare for All. *New England Journal of Medicine*, 377(23):2209–2211, 2017. doi: 10.1056/NEJMp1713510. URL `https://doi.org/10.1056/NEJMp1713510`.

Whitney K Newey and Daniel McFadden. Large sample estimation and hypothesis testing. *Handbook of econometrics*, 4:2111–2245, 1994.

Joseph P. Newhouse, Mary Price, Jie Huang, J. Michael McWilliams, and John Hsu. Steps to reduce favorable risk selection in medicare advantage largely succeeded, boding well for health insurance exchanges. *Health Affairs*, 31(12):2618–2628, 2012. doi: 10.1377/hlthaff.2012.0345. URL `https://www.healthaffairs.org/doi/abs/10.1377/hlthaff.2012.0345`. PMID: 23213145.

Joseph P Newhouse, Mary Price, J Michael McWilliams, John Hsu, and Thomas G McGuire. How much favorable selection is left in Medicare Advantage? *American Journal of Health Economics*, 1(1):1–26, 2015. ISSN 2332-3493 (Print). doi: 10.1162/AJHE_a_00001.

Joseph P Newhouse, Mary Beth Landrum, Mary Price, J Michael McWilliams, John Hsu, and Thomas G McGuire. The comparative advantage of medicare advantage. *American journal of health economics*, 5(2):281–301, 2019.

Kathleen Nosal. Estimating switching costs for medicare advantage plans. Working paper, University of Arizona, 2012.

Barack Obama. United states health care reform: progress to date and next steps. *Jama*, 316(5): 525–532, 2016.

Ariel Pakes and Paul McGuire. Computing markov-perfect nash equilibria: Numerical implications of a dynamic differentiated product model. *The RAND Journal of Economics*, 25(4):555–589, 1994. ISSN 07416261. URL `http://www.jstor.org/stable/2555975`.

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, 2007. 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. 83(1):315–334, 2015. ISSN 00129682. doi: 10.3982/ECTA6865. URL `http://scholar.google.com/scholar?hl=en{&}btnG=Search{&}q=intitle:Moment+Inequalities+and+Their+Application{#}0{%}5Cnhttp://doi.wiley.com/10.3982/ECTA6865`.

Irene Papanicolas, Liana R. Woskie, and Ashish K. Jha. Health Care Spending in the United States and Other High-Income CountriesHealth Care Spending in the United States and Other High-Income CountriesHealth Care Spending in the United States and Other High-Income Countries. *JAMA*, 319(10):1024–1039, 03 2018. ISSN 0098-7484. doi: 10.1001/jama.2018.1150. URL `https://doi.org/10.1001/jama.2018.1150`.

Stephen T. Parente. Factors Contributing to Higher Health Care Spending in the United States Compared With Other High-Income CountriesFactors Contributing to Higher US Health Care SpendingEditorial. *JAMA*, 319(10):988–990, 03 2018. ISSN 0098-7484. doi: 10.1001/jama.2018.1149. URL `https://doi.org/10.1001/jama.2018.1149`.

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

G C Pope, R P Ellis, A S Ash, C F Liu, J Z Ayanian, D W Bates, H Burstin, L I Iezzoni, and M J Ingber. Principal inpatient diagnostic cost group model for Medicare risk adjustment. *Health care financing review*, 21(3):93–118, 2000. ISSN 0195-8631 (Print).

Matthew J Press, Rahul Rajkumar, and Patrick H Conway. Medicares new bundled payments: design, strategy, and evolution. *Jama*, 315(2):131–132, 2016.

Robert Town and Su Liu. The welfare impact of medicare hmos. *The RAND Journal of Economics*, 34(4):719–736, 2003. ISSN 07416261. URL `http://www.jstor.org/stable/1593785`.

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.

Ed Weisbart. A single-payer system would reduce us health care costs. *AMA Journal of Ethics*, 14 (11):897–903, 2012.

Steffie Woolhandler, David U Himmelstein, Marcia Angell, and Quentin D Young. Proposal of the physicians' working group for single-payer national health insurance. *JAMA*, 290(6):798–805, 2003.

Stephen Zuckerman, W Pete Welch, and Gregory C Pope. A geographic index of physician practice costs. *Journal of Health Economics*, 9(1):39–69, 1990.

8 Tables and figures

Table 1: Medicare beneficiary micro-data summary statistics

	All observations		**By MA enrollment**		**TM → MA switch**	
Variable	Mean	Std. dev.	MA	TM	Yes	No
MA enrollment indicator	.226	.418	1	0	1	0
Income ($)	46,535	68,736	38,168	48,982	39,707	44,387
Age	73.4	10.0	73.7	73.3	73.5	75.1
Outside good price ($)	2,390	590	2,487	2,361	2,398	2,382
Demographic indicators						
Female	.536	.499	.549	.532	.540	.539
Black	.079	.270	.093	.075	.084	.068
Hispanic	.010	.100	.016	.008	.013	.007
Education indicators						
Bachelor's degree or higher	.236	.425	.178	.253	.211	.232
Attended college	.304	.460	.310	.302	.304	.291
Graduated high school	.291	.454	.310	.286	.298	.301
Health status indicators						
Healthy	.784	.411	.791	.782	.806	.778
Unhealthy	.215	.411	.209	.218	.194	.222
Administrative indicators						
Employer-provided insurance	.330	.470	.011	.423	.000	.434
Non-aged eligibility	.159	.365	.164	.157	.188	.160
ESRD	.007	.081	.004	.007	.002	.007
Full-year Part A/B enrollee	.898	.302	.975	.876	.970	.898
Observations	58,444		13,225	45,219	994	22,284

Notes: An observation is a person-year. Statistics reported here are weighted according to sampling weights provided by the Medicare Current Beneficiary Survey (MCBS). Income and prices are in 2015 dollars. The outside good price is the United Healthcare premium for Medigap Plan C (see text for details). Demographic categories are defined by CMS administrative data. The first set of two columns reports means and standard deviations for all observations in the microdata. The third and fourth columns split the observations into those enrolled in MA and those enrolled in TM. The last two columns split the observations by switching behavior. Only those person-year observations for which I observe the individual enrolled in TM in last year's MCBS are included in the these last columns.

Table 2: Mean plan characteristics by price quartile

		Price quartile, if non-zero			
Variable	**Zero price**	**1st**	**2st**	**3rd**	**4th**
Annual premium ($)	0	302	600	955	1,733
Annual Part B premium reduction ($)	69	3.80	0.00	0.95	0.00
Out-of-pocket cost estimate ($)	2,405	2,456	2,308	2,254	2,013
Deductible ($)	74	69	140	105	123
Out-of-pocket cost limit ($)	3,592	3,598	4,254	4,286	3,993
Drug coverage indicator	.656	.758	.826	.858	.938
Primary doctor copay ($)	9.27	13.15	13.59	13.34	11.15
Specialist copay ($)	26.32	30.65	30.28	28.12	26.55
HMO indicator	.591	.362	.308	.336	.377
PPO indicator	.191	.251	.378	.351	.383
Generosity index	5.25	6.65	7.15	7.04	7.12
Firm “Star” rating	2.41	2.10	2.28	2.50	2.69
Enrollment	929	423	408	512	584
Observations	21,604	7,252	7,244	7,571	6,922

Notes: An observation is a plan-county-year; all reported figures are unweighted means across observations in the relevant quartile. All dollar amounts are in 2015 dollars. The out-of-pocket cost estimate is generated by CMS based on a basket of medical services consumed by an average TM beneficiary between the ages of 70 and 74 in ‘good’ health. The star rating is calculated at the firm level by CMS and ranges from zero to five.

Table 3: The association between last year's share and this year's prices and generosities

	Change in price				Change in generosity			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Last year's share	.143	.117	.254	.264	-1.36	-1.49	-.472	-.656
	(.051)	(.053)	(.052)	(.054)	(.168)	(.246)	(.245)	(.253)
County FX	No	Yes	No	Yes	No	Yes	No	Yes
Firm FX	No	No	Yes	Yes	No	No	Yes	Yes
R^2	.000	.055	.152	.177	.001	.027	.118	.132
Observations	24,250	24,250	24,250	24,250	24,250	24,250	24,250	24,250

Notes: This table reports the results of estimating Equation 2 via OLS. Observations are plan-county-years, and are restricted to cases where I observe the same plan in the same county the previous year. The left-hand side variable is the year-over-year change in the plan's price in thousands of 2015 dollars or generosity. The independent variable is the market share of the plan last year. Robust standard errors are in parentheses.

Table 4: Demand parameter estimates

Variable	Coeff.	Std. Err.
Price (per $1000)	-2.023	0.201
Price ×		
Medium income indicator	0.149	0.039
High income indicator	0.086	0.042
Generosity index	0.770	0.036
Generosity index × Healthy ind.	0.030	0.010
MA ×		
Age	1.713	0.152
Age^2	-0.118	0.010
Female indicator	-0.035	0.028
Black indicator	0.318	0.055
Hispanic indicator	0.123	0.132
Graduated high school	-0.075	0.043
Attended college	-0.160	0.043
College degree or higher	-0.504	0.047
Healthy indicator	-0.169	0.100
TM-to-MA switch ×		
Constant	-2.616	0.076
Healthy indicator	0.122	0.087
Inter-Insurer switch ×		
Constant	-1.450	0.067
Healthy indicator	0.082	0.075
Medigap price (per $1000)	0.219	0.046
Administrative indicators		
Has employer-provided insurance	-3.644	0.089
Non-aged eligibility	0.296	0.047
ESRD diagnosis	-0.951	0.196
Full year Medicare enrollment	1.779	0.061
Diagnostics		
Weighted log-likelihood	-58,247	
Individual-year obs.	58,444	
First stage F-statistic	103.156	
Plan-county-year obs.	50,593	
Derived statistics	**Mean**	**Std. Dev.**
Elasticity (if < 0)	-1.71	1.14
ds_j/dp_j	-0.021	0.039

Notes: This table reports estimates of the parameters of Equation 3. The price and generosity index parameters are estimated via 2SLS; all other parameters are estimated via maximum likelihood. Income groups are defined by terciles of the income distribution for the entire MCBS sample; low income is the omitted group. The omitted group for the switching cost interactions is 'Unhealthy' health. I weight the likelihood function using the MCBS sample weights to obtain nationally representative estimates. In the second stage, I include firm fixed effects and instrument for price using BLP instruments and non-contiguous-county Hausman instruments.

Table 5: Marginal cost parameter estimates

Variable	Coeff.	Std. Err.
FFS spending on healthy × healthy indicator	0.985	(0.144)
FFS spending on unhealthy × unhealthy indicator	1.112	(0.197)
Generosity	-0.249	(0.144)
Generosity × unhealthy indicator	0.058	(0.110)
Generosity × generosity	0.064	(0.015)
Age	-0.017	(0.119)
Log median household income	-0.055	(0.004)
Log market population	-0.009	(0.033)
Constant	0.057	(0.038)
Constant × unhealthy indicator	0.118	(0.266)
Plan-county-year obs.	8,984	

Notes: This table reports estimates of the parameters of Equation 8 using the two-step method-of-moments approach described in Section 5.2. The generosity index here differs slightly from the index reported in Table 2 due to the weighting approach discussed in Section 5.3. All units are in thousands of 2015 dollars.

Table 6: Comparing Medicare Advantage Costs to Traditional Medicare Costs

	Healthy	Unhealthy
Zero-premium plans		
Mean generosity	2.55	
Marginal cost ($)	7,784	19,280
Positive-premium plans		
Mean premium ($)	513	
Mean generosity	3.69	
Marginal cost ($)	8,220	19,780
Average TM costs ($)	8,080	17,250
Plan-county-year obs.	8,984	

Notes: The generosity index here differs slightly from the index reported in Table 2 due to the weighting approach discussed in Section 5.3. All dollar amounts here are 2015 dollars.

Table 7: Comparing Estimated Costs from Static and Dynamic Approaches

	Mean	Std. dev.
Zero-premium plans		
Statically-estimated marginal cost ($)	9,335	1,137
Dynamically-estimated marginal cost ($)	10,728	2,011
Fraction with dynamic > static cost	0.840	
Positive-premium plans		
Statically-estimated marginal cost ($)	10,132	1,127
Dynamically-estimated marginal cost ($)	11,073	2,037
Fraction with dynamic > static cost	0.674	
Average TM costs ($)	10,111	1,372
Plan-county-year obs.	8,984	

Notes: Dynamic marginal cost estimates come from the parameters in Table 5. Static marginal cost estimates are calculated using the first-order conditions for static profit maximization. Reported costs are for a risk 1.0 enrollee.

Table 8: The impact of the prior-year share on cost estimates

	Estimator input		
	(1)	(2)	(3)
	Original	Incr. share	Duopoly
Average prior-year market share			
Healthy	0.022	0.091	0.502
Unhealthy	0.020	0.087	0.502
Marginal cost for zero-premium plans ($)			
Healthy	7,784	8,087	8,135
Unhealthy	19,280	16,364	12,912
Marginal cost for positive-premium plans ($)			
Healthy	8,220	8,522	8,613
Unhealthy	19,780	16,806	13,851
Plan-county-year obs.	8,984	8,984	1,635

Notes: The "Original" column repeats the results of Table 6. To create the data used in the "Incr. share" column, I increase market shares in each market proportionally until the total inside share equals one. To create the "Duopoly" column, I kept only the top two firms in each market by share and increased their shares proportionately; some markets have only one firm and thus the average share is slightly greater than 0.5. See text for details. All dollar amounts here are 2015 dollars.

Figure 1: A potential example of dynamic firm pricing behavior

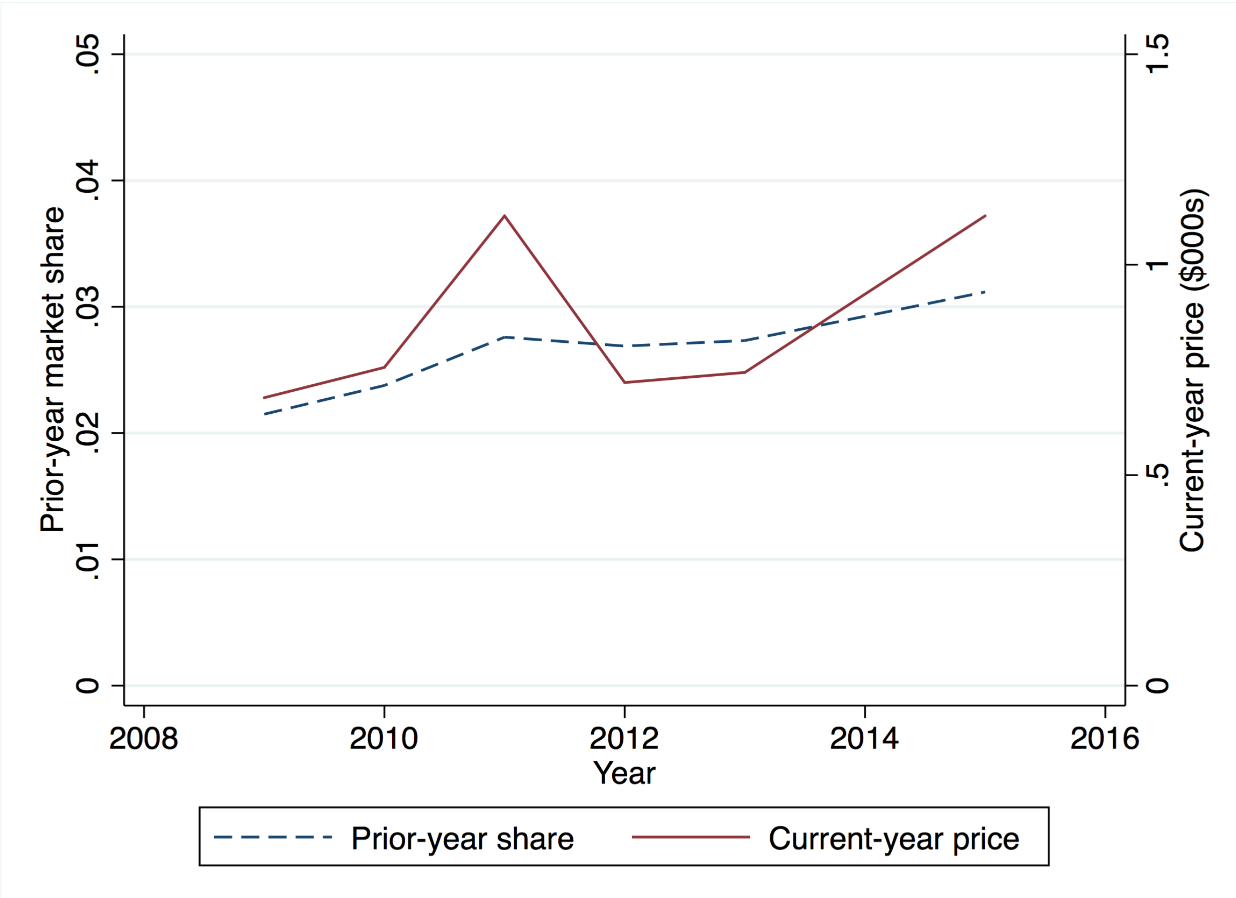

Notes: This graph plots the prior-year share and the current-year price for an example plan. In other words, the points for 2012 indicate the share the plan obtained in 2011 (and thus those "locked in" to the plan's network) and the price the firm chose for the plan in 2012. The plan illustrated is Excellus Health Plan's Medicare Blue PPO Plan Two (CMS plan ID H3335-14-0) offered in Cayuga County, New York.

Figure 2: Example policy functions generated by the model

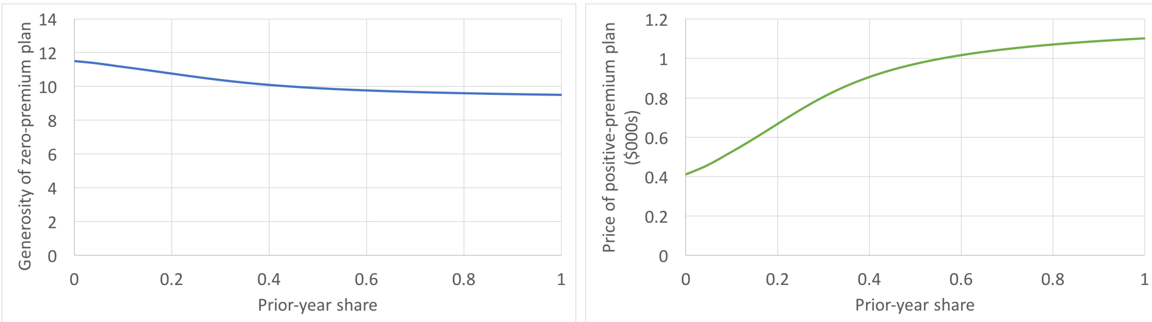

Notes: These figures illustrate policy functions for a firm in Colbert County, Alabama, at the parameters reported in Table 5. The left graph illustrates the generosity set by the firm for its zero-premium plan, and the right graph illustrates the premium of the positive-premium plan.

Figure 3: The distribution of prior-year shares and current-year prices

(a) The distribution of prior-year shares for zero-premium plans

(b) The distribution of prior-year shares and current-year premia for positive-premium plans

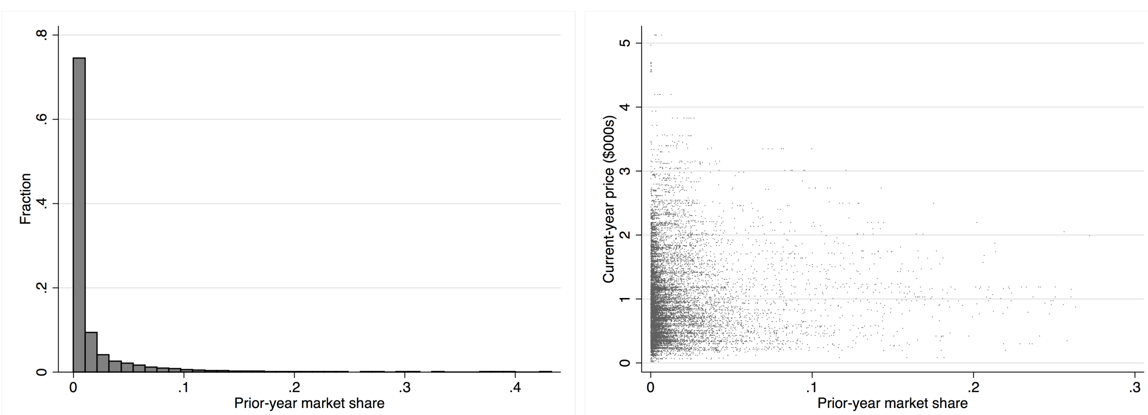

Notes: These figures illustrate the joint distribution of prior-year shares and current-year prices in the raw data. Panel (a) focuses on zero-premium plans and is thus a histogram while Panel (b) is a scatter plot of positive-premium plans.

Appendices

A Additional tables and figures

Table A.1: The generosity index and its components

Variable	Coeff.	Mean	25th %-ile	Median	75th-%ile
Part B prem. reduction ($)	-.008	30.30	0	0	0
Out-of-pocket cost est. ($000s)	-.038	2.32	1.65	2.27	2.83
Drug coverage indicator	2.59	.763	1	1	1
Deductible ($000s)	-1.50	.094	0	0	.5
Out-of-pocket limit ($000s)	-.094	3.85	2.45	3.40	5.80
Primary doctor copay ($)	.109	11.31	5	10	15
Specialist copay ($)	.070	27.81	20	30	35
HMO indicator	1.72	.450	0	0	1
PPO indicator	4.22	.276	0	0	1
Generosity index		6.25	4.69	6.36	8.23

Notes: Coefficients are taken from Miller et al. (2019). All reported figures are from the unweighted distribution of the 50,593 plan-county-year observations reported in Table 2. The out-of-pocket cost estimate is generated by CMS based on a basket of medical services consumed by an average TM beneficiary between the ages of 70 and 74 in self-reported 'Good' health.

Table A.2: Mean plan characteristics by generosity quartile

	Generosity quartile			
Variable	**1st**	**2st**	**3rd**	**4th**
Generosity index	2.45	5.68	7.20	9.65
Annual premium ($)	291	511	536	698
Fraction with zero premium	.623	.438	.421	.227
Annual Part B premium reduction ($)	118	1.52	1.01	0.19
Out-of-pocket cost estimate ($)	2,432	2,253	2,258	2,347
Deductible ($)	59	67	105	145
Out-of-pocket limit ($)	2,603	3,301	3,684	5,799
Drug coverage indicator	.373	.779	.909	.994
Primary doctor copay ($)	6.48	9.68	12.30	16.80
Specialist copay ($)	17.97	26.30	31.22	35.75
HMO indicator	.483	.463	.636	.218
PPO indicator	.022	.117	.199	.768
Firm "Star" rating	1.80	2.13	2.79	2.89
Enrollment	507	664	900	618
Observations	12,652	12,648	12,652	12,641

Notes: An observation is a county-year-plan; all reported figures are unweighted means across observations in the relevant quartile. All dollar amounts are in 2015 dollars. The out-of-pocket cost estimate is generated by CMS based on a basket of medical services consumed by an average TM beneficiary between the ages of 70 and 74 in 'good' health. The star rating is calculated at the firm level by CMS and ranges from zero to five.

Table A.3: Marginal cost parameter estimates with altered data

	Increased share		Duopoly	
Variable	Coeff.	Std. Err.	Coeff.	Std. Err.
FFS spending on healthy × healthy indicator	1.004	(0.537)	1.022	(0.330)
FFS spending on unhealthy × unhealthy indicator	0.950	(0.352)	0.757	(0.244)
Generosity	-0.382	(0.527)	-0.278	(0.367)
Generosity × unhealthy indicator	0.005	(0.810)	0.235	(0.366)
Generosity × generosity	0.077	(0.068)	0.073	(0.070)
Age	0.012	(0.534)	-0.117	(0.360)
Log median household income	0.022	(0.734)	-0.017	(0.297)
Log market population	-0.050	(0.427)	0.078	(0.052)
Constant	0.065	(0.864)	0.0920	(0.379)
Constant × unhealthy indicator	-0.003	(1.350)	-0.182	(0.432)
Plan-county-year obs.	8,984		1,635	

Notes: This table reports estimates of the parameters of Equation 8 using the two-step method-of-moments approach described in Section 5.2. The generosity index here differs slightly from the index reported in Table 2 due to the weighting approach discussed in Section 5.3. The input to the estimation routine has been altered by increasing the market share for each firm proportionately until the total share of MA equals one. All units are in thousands of 2015 dollars.