Can Health Insurance Competition Work? Evidence from Medicare Advantage

Vilsa Curto

Harvard University

Liran Einav

Stanford University and National Bureau of Economic Research

Jonathan Levin

Stanford University and National Bureau of Economic Research

Jay Bhattacharya

Stanford University and National Bureau of Economic Research

We estimate the economic surplus created by the Medicare Advantage program under its reformed competitive bidding rules. We use data on the universe of Medicare beneficiaries and develop a model of plan bidding that accounts for both market power and risk selection. We estimate that the Medicare Advantage program generates substantial surplus to participants (of \$217 per enrollee-month), but that approximately two-thirds of this surplus is captured by insurers. We use the model to evaluate the impact of possible program changes, including changes that could increase competition and lead to lower profits and higher consumer surplus without raising taxpayer costs.

We thank the editor (Ali Hortaçsu), three anonymous referees, Amy Finkelstein, Eugenio Miravete, and many seminar participants for many useful comments. We are grateful to Diego Jiménez Hernández and Evan Mast for excellent research assistance. Financial support from the National Science Foundation (grant SES-1527942), the Sloan Foundation, and (to Bhattacharya) the National Institute on Aging (grants R37AG036791 and R21AG041112) is gratefully acknowledged. Replication files are provided as supplementary material online.

Electronically published January 11, 2021
[*Journal of Political Economy*, 2021, vol. 129, no. 2]

I. Introduction

Introducing managed competition into health care has been an alluring idea to economists and policy makers. Proponents argue that effectively designed market mechanisms can avoid the inefficiencies of an administrative price system. Yet there is little consensus on this claim. One reason for this is that many competitive systems do not look much like proposed ideals. Another reason, perhaps equally important, is that there is often no clear way to draw comparisons among alternative health insurance systems.

The recent experience of the US Medicare Advantage program provides an opportunity to evaluate how private health insurance competition can work. The program allows seniors to opt out of public Medicare insurance and enroll in a private insurance plan. The federal government pays plans a monthly amount for each enrollee. Historically, the payments were set administratively and the program suffered from limited uptake and cream skimming (McGuire, Newhouse, and Sinaiko 2011). In the last decade, however, Medicare introduced two new ingredients touted by advocates of managed competition (Enthoven 1993): competitive bidding to encourage plans to accept payments below a maximum benchmark rate, and risk adjustments that make payments a function of enrollee health status. Since these reforms, the program has expanded to cover more than 30% of US seniors (fig. 1).

In this paper, we study insurer competition under the Medicare Advantage (MA) bidding rules, and estimate the program's welfare effects relative to traditional Medicare (TM). The competition model we propose allows us to measure insurers' bidding incentives and the benefits that accrue to private plan enrollees, and to analyze how these might change under alternative program designs. We also adjust for non-risk-adjusted health differences between private insurance and traditional Medicare enrollees. Elements of the problem—the determinants of plan bids, enrollee choice, and risk selection into private plans—have been analyzed in prior and concurrent work. Here we contribute an empirical framework that ties the pieces together in a way that facilitates econometric measurement and analysis of program design, and apply it using comprehensive program data.

A practical motivation for our analysis is the ongoing debate over the taxpayer costs of Medicare Advantage. In 2012, the federal government spent $136 billion on payments to private insurers. The MedPAC advisory group has pointed out repeatedly that taxpayers pay less per beneficiary under traditional Medicare (MedPAC 2013). An open question is whether the extra spending is due to inefficiencies in the ways private plans operate, results from a failure of competition, or instead reflects extra benefits for private plan enrollees. Any plausible answer needs to account for

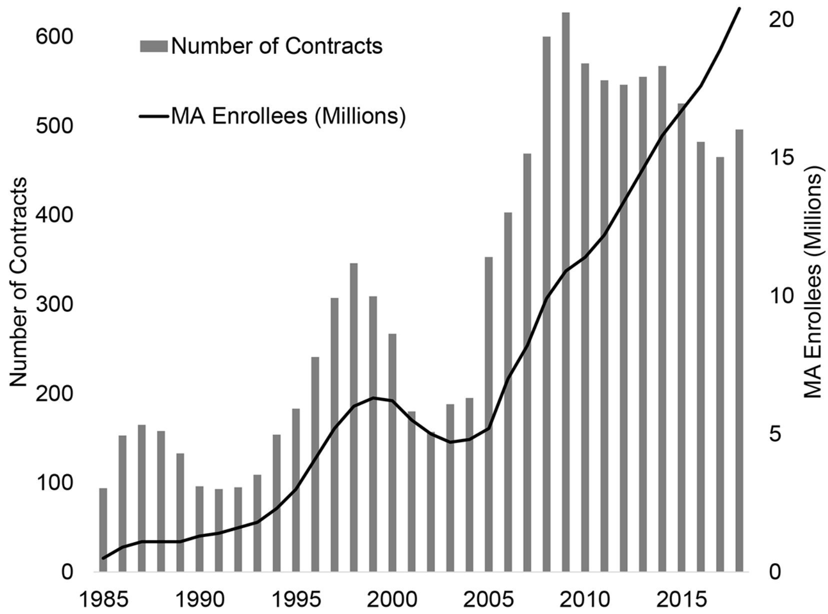

FIG. 1.—The growth of Medicare Advantage. The figure shows the number of MA contracts and the number of Medicare beneficiaries who enrolled in MA over the last three decades. As described in the text, an MA contract is the unit of observation that represents the most important decision by an insurer to enter the MA market. A contract is typically mapped into multiple distinct plans that share common features, such as provider networks, with each being offered in multiple counties. The period described experienced several important regulatory and legislative milestones, including the beginning of the Part C program in 1985, authorized under the 1982 Tax Equity and Fiscal Responsibility Act; the 1997 Balanced Budget Act, which authorized PPOs and private fee-for-service plans and raised payment rates; the 2003 Medicare Modernization Act, which instituted a competitive bidding system and a risk adjustment system based on past health diagnoses; and the 2010 Patient Protection and Affordable Care Act, which reduced payment rates and introduced bonus payments for high-quality plans. This figure is the authors' adaptation of figs. 2 and 4 from McGuire, Newhouse, and Sinaiko (2011). Contracts include HMOs, local PPOs, private FFS plans, and regional PPOs. Employer-sponsored plans are excluded. The data source is CMS's Medicare Managed Care Contract Plans Monthly Summary Reports. All data are from December of the year indicated, except for 2018, which uses data from September.

imperfect competition—as we discuss below, local insurance markets are highly concentrated—and provide a way to estimate insurer profits and enrollee benefits. It also needs to account for the fact that private plans tend to enroll relatively healthy beneficiaries (MedPAC 2013; Brown et al. 2014; Curto et al. 2019). We do this using an imperfect competition model that is adapted to MA's specific bidding rules and that allows for differential risk selection.

We begin in section II by providing a brief description of the MA program, the rules governing plan bidding, and our data. In section III, we describe an empirical model of insurer competition. The model we propose follows the structure of the MA program, in which plans submit bids to cover representative beneficiaries in a local market, and the bids translate into plan payments and enrollee premiums or benefits. In our baseline model, risk adjustment successfully scales plan payments to reflect enrollee health. We then explain how to account for imperfect risk adjustment (e.g., for risk selection on unobserved characteristics), which can distort the incentives for plan pricing. We argue that in practice the problem of risk selection can be simplified so long as plan bids, on the margin, do not have a large effect on enrollee risk composition. Then risk selection shifts plan costs but not marginal bidding incentives.

Demand for MA plans is modeled and estimated in section IV. The price sensitivity of Medicare beneficiaries and the differentiation among private plans are critical inputs into insurers' pricing incentives. We follow much of the existing literature and model demand for private plans using a standard nested logit specification. In our preferred specification, we estimate the sensitivity of beneficiaries to plan premiums using generalized method of moments (GMM) that combines both demand-side and supply-side moments. On the demand side, the key identifying variation relies on year-to-year changes in plan bids as well as variation in bids across distinct plans that share identical provider networks. On the supply side, we use our assumption about optimal bidding, which is based on our model presented in section III, to impose the fact that our estimates of price sensitivity imply plan markups that coincide, on average, with those obtained from external data on average costs for a subset of MA plans (reported in Curto et al. 2019).

In section V, we combine our demand estimates with our model of optimal bidding to assess the gains from the MA program, and how these gains may change in response to changes in market design parameters. We estimate that there are substantial gains from the MA program, although these gains are mainly captured by insurers. Specifically, we estimate that the MA program generates \$88 per enrollee-month in consumer surplus, and an additional \$129 per enrollee-month in insurer profits. We find that only a small part of these gains can be attributed to increased federal spending: we estimate that covering all MA enrollees under traditional Medicare would reduce taxpayer costs by only \$46 (6%) per enrollee-month. The current MA rules virtually ensure that taxpayer costs are not terribly different because Medicare aims to keep its benchmark reimbursement for MA close to TM costs. MA plans that are able to realize significant cost savings and submit low coverage bids are required to pass on these savings to enrollees in the form of health benefits, rather than to taxpayers.

The final section of our analysis reports on counterfactual exercises that highlight two market design parameters set by the federal government: the county-by-county benchmark reimbursement rates against which plan bids are assessed, and the rebate pass-through rate, which determines how plans must distribute benefits between taxpayers and enrollees when they submit bids below the benchmark. These exercises provide two lessons. The first is the relationship between taxpayer spending on MA, and MA enrollment: on the margin, a $5 per enrollee increase in taxpayer spending on MA increases MA enrollment by 1 percentage point. As a result, we find that directly reducing taxpayer spending by cutting benchmark plan reimbursement would shift more beneficiaries into TM. The second takeaway is more nuanced, highlighting the different effects of the two market design parameters. While increasing (decreasing) benchmark rates leads to higher (lower) consumer surplus and insurer profits, increasing the rebate pass-through rate leads to higher consumer surplus, but *lower* insurer profits. This latter effect arises because a higher rebate pass-through rate increases the sensitivity of demand to insurer bids, and therefore leads to lower equilibrium markups. Consequently, these distinct effects allow us to illustrate how combining changes to both market design parameters could shift some of the surplus from insurers to consumers, without increasing government expenditure.[1]

Our analysis contributes to a recent literature assessing different elements of the reformed Medicare Advantage program. Song, Landrum, and Chernew (2012, 2013) and Cabral, Geruso, and Mahoney (2018) estimate the reduced form effect of benchmark rates on plan bids, finding pass-through rates of around 50%. Cabral, Geruso, and Mahoney (2018) also investigate whether the incomplete pass-through they find can be explained by competitive pricing under imperfect risk adjustment, but argue that market power on the part of plans (which is the focus of our work) provides a better explanation. Duggan, Starc, and Vabson (2016) also estimate pass-through rates, using a different identification strategy, and find that plan bids are almost dollar-for-dollar responsive to the benchmark rate, so that higher benchmarks lead to little consumer benefit. Stockley et al. (2014) and Kluender and Mast (2017) investigate how benchmark changes affect different components of pricing, and Song, Cutler, and Chernew (2012) observe that the lowest plan bids are well below fee-for-service costs, and discuss whether the program bidding rules are responsible for high taxpayer costs. There is also a debate about whether Medicare's risk adjustment policy has managed to mitigate risk selection (Newhouse et al. 2012, 2015; Morrissey et al. 2013; Brown et al. 2014), which is tangentially related to some of the results reported in the paper.

Several papers (Town and Liu 2003; Lustig 2009; Dunn 2010, 2011; Hall 2011) are conceptually closer to ours in studying the welfare impact

[1] It is important to emphasize that our entire analysis throughout keeps the availability of traditional Medicare as a public option. Assessing the impact of the Medicare Advantage program in the presence of a different public option or nonexistent public option would require us to go much further out of sample and is beyond the scope of this paper.

of Medicare private plans. While these studies estimate nested logit demand systems similar to the one we estimate, these papers all focus on an earlier period of Medicare Advantage prior to the introduction of competitive bidding and more sophisticated risk adjustment.[2] Relative to this existing work, our key contribution is that we propose and estimate a competitive bidding model that incorporates the current structure of the Medicare Advantage program, allowing us to analyze the equilibrium effects of changes in current market design parameters.[3]

II. Setting and Data

A. *The Medicare Advantage Program*

Medicare Advantage (MA) allows Medicare beneficiaries to opt out of traditional fee-for-service Medicare and enroll in a private insurance plan. The program (under its original name, Medicare Part C) was established in the early 1980s with two goals: to expand the choices available to beneficiaries and to capture cost savings from managed care. Our analysis focuses on the portion of the program targeted at individual Medicare beneficiaries. There is also a portion of the program that allows employers to sponsor plans for Medicare-eligible employees or retirees.

Private plans receive capitated monthly payments from the Centers for Medicare and Medicaid Services (CMS), which were set administratively until 2006.[4] There has historically been a tension between the two goals of expanding access and limiting costs (McGuire, Newhouse, and Sinaiko 2011). Insurers have tended to participate more in periods with higher payments, and to offer plans selectively in areas with higher payment rates. Plans also enrolled relatively healthy beneficiaries, complicating the problem of setting appropriate capitation rates.[5]

[2] Aside from the introduction of competitive bidding and risk adjustment, another major change was the introduction of Medicare Part D prescription drug coverage. In the study by Town and Liu (2003), prescription drug coverage by Medicare HMOs accounts for around half of the estimated consumer surplus. There are a number of other related studies of the earlier MA program. For instance, Dowd, Feldman, and Coulam (2003) estimate enrollment price sensitivities using data from 1999, and Pizer and Frakt (2002) is a pass-through study that examines the effect of CMS payment rates on plan benefits.

[3] Our model is static and abstracts from possible dynamic implications. Nosal (2012) provides evidence of consumer switching costs in the MA context, and Miller (2019) and Miller et al. (2019) are two more recent papers that estimate a dynamic model of plan competition, focusing on how switching costs may lock enrollees into MA plans. Like earlier work, these papers abstract from the specific details of the MA bidding rules, and are thus unable to engage in the impact of market design parameters.

[4] The Centers for Medicare and Medicaid Services (CMS) is the federal agency that manages the Medicare program.

[5] Concerns about risk selection date to the very origins of the program (Eggers and Prihoda 1982), and are discussed in McGuire, Newhouse, and Sinaiko (2011). An illustration of selective participation is that in 2005, only around 67% of Medicare beneficiaries had access to an HMO or local PPO plan (MedPAC 2009).

Several reforms have aimed to address these problems.[6] Between 2003 and 2006, CMS phased in a risk scoring system to adjust plan payments based on enrollee health.[7] In 2006, competitive bidding replaced the fixed reimbursement rate. These changes, combined with an increase in maximum capitation rates set by CMS, have coincided with the expansion of plan offerings and enrollment seen in figure 1. Interestingly, MA enrollment has continued to rise despite the fact that the Affordable Care Act, which was signed into law by President Obama in 2010, gradually reduced payments to MA insurers.

Medicare Advantage plans must provide at least the same insurance benefits as traditional Medicare (Parts A and B). They typically provide additional benefits as well, in the form of more generous cost sharing or supplemental coverage of dental, vision, or drug benefits. An important feature of the program is that MA insurers cannot simply make their plans more attractive; that is, these additional benefits must be funded, either through a supplemental premium paid by the insuree, or more frequently through a rebate paid by CMS. As we describe in more detail below, plan rebates, as well as enrollee premiums, are determined through the competitive bidding process.

From a Medicare beneficiary's perspective, MA presents a clear tradeoff. On the one hand, private plans typically restrict access to health care providers. Around 85% of MA enrollees are in HMO or PPO plans with limited provider networks and various restrictions on utilization. On the other hand, the extra benefits offered by MA plans, especially reduced cost sharing, can make private plans attractive relative to traditional Medicare, where enrollees can face large out-of-pocket costs. Traditional Medicare enrollees can insure against out-of-pocket costs by purchasing supplemental Medigap policies, but these policies often cost a few thousand dollars per year.[8] MA plans provide a "one-stop shop" to cover these costs as well as additional benefits such as dental, vision, or prescription drug coverage.[9]

[6] The reforms originate in the Balanced Budget Act of 1997 (risk scoring) and the Medicare Modernization Act of 2003 (competitive bidding and more detailed risk scoring).

[7] The risk scores are based on a formula that gives weights to chronic disease diagnoses. At the same time, CMS also reformed the enrollment process, so that beneficiaries must enroll in MA plans during a fixed annual open enrollment period, rather than being able to switch in and out of private plans on a monthly basis. There is some debate about whether and how much risk adjustment has altered plan incentives (Newhouse et al. 2012, 2015; Brown et al. 2014).

[8] In 2010, monthly premiums for Medigap insurance purchasers averaged $183 per month (Huang et al. 2013).

[9] A significant share of MA plans also include Medicare Part D prescription drug coverage, charging a supplemental fee for this coverage. An analysis of prescription drug insurance is outside the scope of this paper, but we account for this later by allowing consumers to value the convenience of the bundled package in our empirical specification of private plan demand.

Finally, there is also a class of plans known as "private fee-for-service" (PFFS). Roughly speaking, these plans mimic traditional Medicare in terms of provider access and how they reimburse nonnetwork providers. This type of plan proliferated in the mid-2000s when benchmark rates were very favorable, and in 2008 reached a 23% share of MA enrollees. However, after regulatory changes that made PFFS participation more difficult, their share dropped to 7% by 2011, and they are now relatively unimportant. Therefore, our empirical analysis below abstracts from such plans. Table E1 (tables A1, B1–B5, C1, and E1–E5 are available online) provides more detail on insurance options and how they compare to traditional Medicare.

B. Data

Our analysis is focused on the 2006–11 period.[10] We use Medicare administrative records, which contain several pieces of information. For each Medicare beneficiary, the data include her detailed demographic information, the MA plan (or traditional Medicare) she was enrolled in every month, and an annually computed risk score, which is, as we explain below, a (normalized) prediction of the beneficiary's annual health care cost based on her health care utilization over the previous 12 months. For each MA enrollee, we also observe the (capitated) payment from CMS to the insurer. For traditional Medicare enrollees, we observe every medical claim during 2006–10.[11] We merge this data set with information on benefit details and bids associated with all MA plans available to Medicare beneficiaries.

We restrict our attention to aged (nondisabled, non-ESRD [end-stage renal disease]) Medicare beneficiaries who are not enrolled in employer-sponsored plans, special needs plans, or PFFS plans. We also drop a small number of individuals for whom we are missing data fields such as risk score or private plan payment information. The resulting data set contains 171,539,631 enrollee-year observations on 42,338,991 unique enrollees. The plan data set contains 8,624 observations at the plan-year level on 2,976 unique MA plans. Exact details of the data construction are reported in appendix A and table A1 (apps. A–E are available online). We note that

[10] The year 2006 was when the modified risk adjustment system and competitive bidding were introduced. By 2012, additional tweaks to the MA program had been implemented, such as the use of plan quality metrics and offering favorable treatment to high-quality plans.

[11] CMS has very limited information on health care utilization within private plans, as CMS began to collect detailed information on encounters only in 2013. Curto et al. (2019) use claim-level data from three large MA insurers in 2010 to document health care utilization differences between MA and traditional Medicare.

the annual data on plan bids and plan payments at the plan level are unique to our paper, and are critical in facilitating the empirical analysis.

C. *Market Structure of Private Plans*

Individual markets tend to feature a large number of plans, but they are mostly small PFFS plans. Between 2006 and 2011, Medicare beneficiaries had access, on average, to three HMO plans and one PPO plan, in addition to 14 PFFS plans. The result is that local markets are highly concentrated despite the large total number of plans (see table E2). In three-quarters of US counties, three insurers serve more than 90% of MA enrollees. In nearly half of US counties, two insurers serve more than 90% of MA enrollees. Concentration is somewhat lower in urban markets and in markets with high benchmark rates (Pizer, Frakt, and Feldman 2009). For example, the mean insurer Herfindahl-Hirschman index in our study period is 47.7% for urban areas and 54.7% for rural areas.

There is less concentration at the national level. The two largest insurers, UnitedHealthcare and Humana, serve 19% and 16% of national MA enrollees, respectively (see table E3). These insurers operate in over 80% of local markets, but most insurers (there were around 130 in total during the 2006–11 period) operate in far fewer. The market structure also looks different if one accounts for traditional Medicare. Around 49% of new enrollees in a given plan were enrolled the prior year in traditional Medicare and around 12% of MA enrollees who exit a plan move into traditional Medicare.[12] The asymmetry reflects the expansion of private plans. In 2006–11, around 2.4% of traditional Medicare enrollees switched into MA each year, while only 2.0% switched in the other direction. Table E4 provides more detail on beneficiary transitions.

D. *Differences between TM and MA Enrollees*

Table 1 presents summary statistics on TM and MA enrollees for the 2006 to 2011 period. MA enrollment over this period averaged around 20%.[13] Private plan enrollees are more concentrated in urban areas. They also have lower risk scores: the average risk score of a private plan enrollee is 0.998, compared to 1.041 for traditional Medicare. Risk scores are constructed so that an individual with a risk score of 2 has twice the expected

[12] We use our analysis sample for these calculations, which excludes MA enrollees who are in PFFS or employer-sponsored plans.

[13] In our final analysis sample, MA penetration is 18%. However, as part of our sample construction we drop some MA enrollees for whom it is not possible to drop the corresponding TM enrollees, as we explain in app. A. If we were to include those MA enrollees, then MA penetration would be approximately 20%.

TABLE 1
SUMMARY STATISTICS

	Traditional Medicare				Medicare Advantage			
	Mean (1)	Standard Deviation (2)	10th Percentile (3)	90th Percentile (4)	Mean (5)	Standard Deviation (6)	10th Percentile (7)	90th Percentile (8)
Age[a]	76.4	7.8	67.0	87.4	76.5	7.3	67.6	86.8
Male (%)	42.1				41.6			
Urban[b] (%)	64.9				90.3			
New Medicare enrollee (%)	4.7				2.1			
Dual eligible (%)	13.0				8.9			
Supplemental insurance[c] (%)	72.7				10.1			
Part D coverage (%)	48.4				94.7			
Died during year (%)	4.1				3.6			
Risk score	1.041	.912	.357	2.119	.998	.844	.342	2.019
log(risk score change)[d]	.054	.525	−.533	.696	.049	.492	−.532	.659
TM monthly claims costs[e]	812	2,520	7	2,079				
MA monthly total CMS payment					806	626	323	1,542
MA monthly rebate payment					87	86	21	160
Number of enrollee-years		145,321,124				26,218,507		
Number of unique enrollees		34,872,700				7,466,291		

NOTE.—Statistics are based on full sample of beneficiaries from 2006 to 2011. The fee-for-service (FFS) Medicare sample excludes ESRD and disabled beneficiaries. The Medicare Advantage sample excludes ESRD and disabled beneficiaries as well as those enrolled in employer-sponsored plans, special needs plans, and private fee-for-service (PFFS) plans. See the text for exact sample restrictions.

[a] Age as of December 31 of observation year.

[b] Urban means county with population greater than 250,000 in 2004 (when the urban floor was last set before the beginning of the sample period).

[c] Supplemental insurance refers to Medigap or retiree supplemental insurance (RSI).

[d] The natural logarithm of the risk score change—log(risk score in following year) minus log(risk score in observation year)—is defined over the sample of beneficiaries who survive until the following year.

[e] Data are for 2006–10 because 2011 FFS claims were not available.

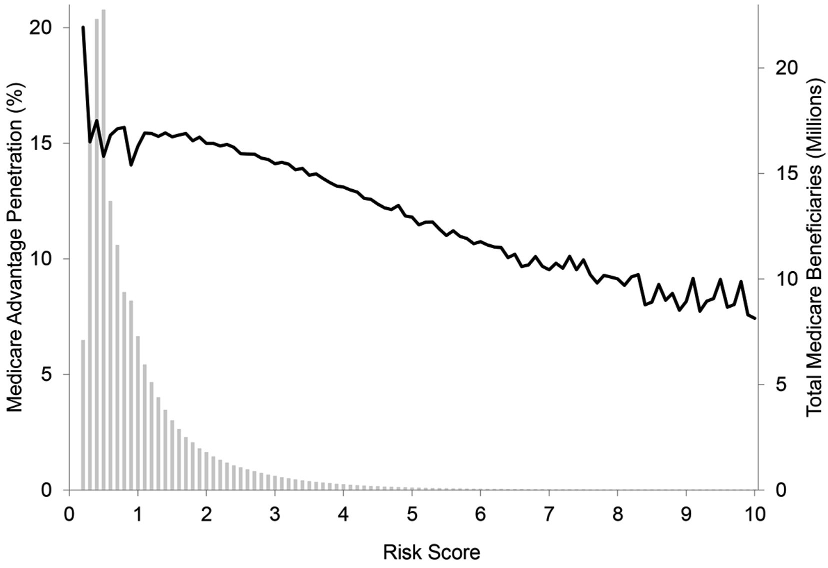

FIG. 2.—MA share by risk score. The figure shows the share of Medicare beneficiaries who select an MA plan by the beneficiary's risk score bin (pooled across years). Each bin is of 0.1 width, so that the first bin covers beneficiaries with risk score between 0 and 0.1, the second covers 0.1–0.2, and so on. The figure illustrates that the MA share is significantly lower for beneficiaries with high risk scores. The gray bars represent the underlying distribution of risk scores among Medicare beneficiaries (again, pooled across years) in order to emphasize that the range over which MA share starts declining is associated with only a small fraction of Medicare beneficiaries.

health care costs of an individual with a risk score of 1. Based on risk scores, a private plan enrollee has expected health care costs that are about 4% lower than those of a TM enrollee.[14] Figure 2 provides more detail by plotting private plan enrollment as a function of risk score. Private plan enrollment declines with risk score, and is especially low at very high scores. As is shown in the figure, MA penetration is 20% for those with risk scores ranging from 0.2 to 0.3, and is only 7% for those with risk scores above 10.[15]

Medicare risk scores are based on a statistical model that translates disease codes into predicted fee-for-service costs. It is well understood that the scores may not fully correct for health differences (Newhouse et al.

[14] The risk score is calibrated to have a mean of 1 on a subsample of several million TM enrollees. We have dropped individuals who qualify for Medicare due to a disability or due to ESRD, enrollees with primary health insurance via a current employer, enrollees in PFFS, and enrollees in employer-sponsored plans, making risk scores in our sample slightly higher than in the entire population.

[15] There are probably multiple factors that help to explain this pattern. One interpretation is that chronically ill individuals are less likely to search for a suitable MA plan, or prefer to have a wide choice of providers. Another interpretation is that MA plans try to avoid high-cost enrollees through plan design.

2012, 2015; Brown et al. 2014; Curto et al. 2019; Geruso and Layton 2020). An individual's codes may be incomplete and, in any event, are imperfect proxies for health status. A wide variety of evidence suggests that MA enrollees are somewhat healthier than TM enrollees conditional on risk score. For instance, if one compares new MA enrollees to TM enrollees with the same risk score, the "switchers" into MA have lower medical claims in the prior year than TM enrollees with the same risk score, a finding from Brown et al. (2014) that we confirm in our data (results not shown). Among the TM "stayers," prior year claims are significant predictors of current year claims conditional on risk score, suggesting that new MA enrollees are healthier than TM enrollees with identical risk scores.[16]

Disease coding is another reason that MA enrollees may be healthier conditional on risk score. Private plans, which are compensated based on risk scores, tend to code more intensively (CMS 2013, chap. 7; GAO 2013; Geruso and Layton 2020). An indication of this is that for many chronic conditions, which are unlikely to go away from one year to the next, MA coding is noticeably more persistent.[17] In 2010, CMS attempted to correct for differential coding, and deflated all MA risk scores by 3.41%. It applied the same adjustment in 2011, but there was no such adjustment in the years prior to 2010.

III. A Model of Bidding Competition

A. *The Bidding Rules*

The bidding system in MA is organized around the price at which private plans are willing to provide standardized Part A and Part B coverage to a representative beneficiary (a beneficiary with a risk score equal to 1). If the price is high, a portion of the payment to private plans is covered by an additional premium paid by the beneficiary. If the price is low, MA plans can use the bidding system to offer a lower premium or (more often) more generous coverage, with CMS netting a share (25%) of the savings. This system forces MA plans to fund reduced premia or more generous coverage through lower bids, and therefore to share some of the potential benefits with CMS.

[16] In our sample, the lagged claims of "switchers" into MA were about 2.3% lower than those of TM "stayers" conditional on risk score. The correlation of lagged TM costs and current year TM costs conditional on risk score is 0.19. The lagged claims relationship underlies the analysis of risk selection in, e.g., Brown et al. (2014).

[17] To illustrate, we compute the probability of a given disease coding in a given year conditional on having this disease coded for the same beneficiary in the previous year. This probability is generally higher for MA enrollees. For example, for chronic obstructive pulmonary disease the probability is 71% in MA vs. 66% in TM. It is 58% vs. 45% for polyneuropathy, and 49% vs. 33% for angina pectoris/old myocardial infarction.

Specifically, the bidding process in MA begins with CMS setting a benchmark capitation rate. The benchmark rates are set each year at the county level. The distribution of county benchmarks is shown in figure E1 (figs. E1 and E2 are available online). We therefore think of each county-year as a separate market. Each plan's bid is assessed against its local benchmark B.[18] Insurers submit their plan bids after the benchmarks are published. The payment rules depend on whether a plan bid is above or below the benchmark rate, and work as follows. If a plan's bid b is above its benchmark rate B, an enrollee must pay the plan a premium of $b - B$. This premium is in addition to the standard Medicare Part B premium and any supplemental premium the plan charges for additional benefits.[19] CMS pays the plan $rb - (b - B)$, where r is the enrollee's risk score. In total, the plan receives rb. If a plan's bid b is below its benchmark rate B, the enrollee pays no plan premium.[20] For an enrollee with risk r, CMS pays the plan rb. It also gives the plan a rebate of $0.75(B - b)$, which the plan must use to "fund" additional benefits it provides to enrollees.

The vast majority of plans submit bids below their benchmark rates, and rebates are often \$50 per month or more (see fig. 3). Plans typically use the rebates to provide extra cost-sharing benefits or premium reductions for prescription drug coverage.[21] Combining the two cases, we see that the plan receives rb to provide basic coverage, and the enrollee either pays a premium $b - B$ if $b > B$, or receives benefits of actuarial value $0.75(B - b)$ if $b \leq B$. So despite the rather complicated details, plans face a familiar trade-off: a lower bid reduces revenue per enrollee, but makes a plan more attractive to beneficiaries.

[18] Each plan has a "service area" that may span multiple counties. In this case, the plan's benchmark is the average of the service area benchmarks, weighted by projected plan enrollment. Subsequently, CMS payments are adjusted to the county level by multiplying them by the individual county benchmark divided by the plan benchmark. Despite this, we think of competition as occurring at the county level because insurers have the option of defining more granular plans and fine-tuning their bids to the county level.

[19] All Medicare beneficiaries pay CMS a standard premium regardless of whether they are enrolled in traditional Medicare or MA. In 2014, the Part B monthly premium was \$104.90 per month for married couples with incomes up to \$170,000; wealthier couples or individuals pay slightly more. Plans may charge a supplemental premium for additional services, but not to reduce cost sharing, which they can do only with the rebate.

[20] This is not always right. As mentioned earlier, plans sometimes charge an additional supplemental premium to fund additional benefits. This is not common, and when it happens the amounts of these supplemental premiums are relatively low.

[21] Table E5 provides a breakdown of how plans use their rebates. For instance, for plans with bids that are below the benchmark and within \$100 of the benchmark, 80% of rebate dollars are used for cost-sharing benefits and 11% of rebate dollars are used for Part D benefits. In principle, rebates can be passed directly to enrollees in cash through reductions in their Medicare Part B premiums, but this is relatively uncommon. For plans with bids that are below the benchmark and within \$200 of the benchmark, less than 1% of rebate dollars are used to reduce the Part B premium. See also Stockley et al. (2014).

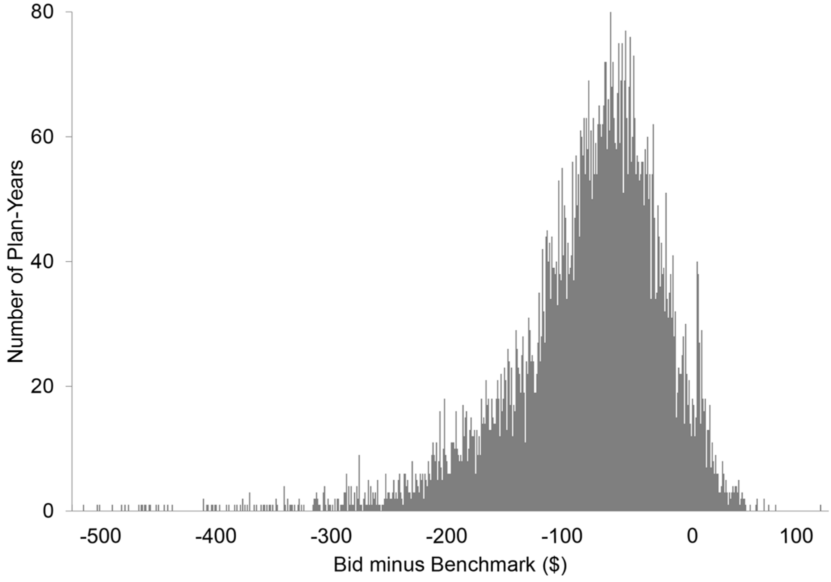

FIG. 3.—Bid distribution (relative to benchmark). The figure shows a histogram of plan bids in the data (relative to the benchmark). Bids are for covering a standardized beneficiary over a month. Each observation is a plan-year, and the figure shows the distribution of $b - B$, where b is the plan's bid and B is a plan-specific benchmark. Benchmarks are the same for all plans offering MA coverage in a given county, but plans are offered in service areas that cover multiple counties, making the relevant benchmark a plan-specific weighted average of county benchmarks. The bid is below the benchmark for 96% of observations. The mean of $b - B$ is −96 (standard deviation 69), with the 5th, 25th, 50th, 75th, and 95th percentiles being −222, −128, −84, −51, and −4, respectively. The difference $b - B$ has risen slightly in magnitude over time, with mean values changing from −84 in 2006 to −89 in 2011.

B. A Baseline Model

We consider a single market with beneficiaries who vary in their risk type r. There are J plans, indexed by $j = 1, \ldots, J$. Plan j has cost rc_j of covering an individual with risk r. There is also traditional Medicare. Bidding competition follows the rules described above. Let B denote the benchmark rate, and b_j the bid of plan j.

Plan demand.—Beneficiaries choose among plans taking into account a plan's fixed characteristics (its provider network, brand name, etc.) and the premium or extra rebate benefits that result from the plan's bid, or more precisely the plan's *excess bid* $p_j = b_j - B$. We write the demand for plan j among beneficiaries with risk type r as $D_{jr}(p_1, \ldots, p_J)$. A natural assumption is that D_{jr} is decreasing in a plan's own excess bid, and increasing in the excess bids of rival plans.

We define a plan's *risk-weighted enrollment* as

$$Q_j(p_1, \dots, p_J) = \int rD_{jr}(p_1, \dots, p_J)\, dG(r). \tag{1}$$

The total number of enrollees in plan j is $D_j = \int D_{jr}\, dG(r)$. In practice, the risk-weighted enrollment tends to be slightly smaller than the total number of enrollees since the average enrollee risk in most plans, $\bar{r}_j = Q_j/D_j$, is less than 1.

Plan profits.—Next we consider plan profits. A plan receives rb_j for enrolling an individual with risk r, and incurs cost rc_j. That is, we make an important assumption that a private plan's cost of enrolling an individual scales proportionally with the individual's risk score (but without restricting the plan cost of covering a "standard" beneficiary with $r = 1$). Because plan revenues are linear in r, this cost proportionality assumption significantly simplifies the model; it gets around complications that would arise from risk selection, and the plan optimal bidding problem resembles a fairly standard optimal pricing situation, except that demand is measured in risk units rather than in the number of enrollees. In particular, plan j's profit given a set of plan bids $p_1, \dots, p_J$ is

$$\pi_j(p_1, \dots, p_J) = Q_j(p_1, \dots, p_J)(p_j + B - c_j). \tag{2}$$

The plan's risk composition matters only insofar as it affects risk-weighted enrollment. This is because the effects of risk composition on costs are perfectly compensated by the risk adjustment formula. We discuss below how plan incentives might be skewed if risk adjustment is imperfect.

Equilibrium bids.—We assume that bids are generated in a complete-information Nash equilibrium. A bidding equilibrium is given by a vector of bids $p_1, \dots, p_J$, such that each insurer is maximizing profits given rival bids.[22] For an insurer that offers a set of plans $\mathcal{J}$, the optimal bid for plan $j \in \mathcal{J}$ satisfies

$$0 = \sum_{l \in \mathcal{J}} (p_l + B - c_l) \frac{-\partial Q_l}{\partial p_j} + Q_j(p_1, \dots, p_J), \tag{3}$$

where the insurer accounts for the fact that raising the bid for plan j may affect enrollment in the insurer's other plans.

In the case where an insurer offers a single plan j, we can write the formula for the optimal excess bid as

[22] Complete information is, of course, a simplification, as actual bids are made without complete knowledge of rivals' bids, but it avoids the complication of modeling an incomplete information bidding game. Conditions under which such a perfect information bidding equilibrium exists are laid out in Caplin and Nalebuff (1991).

$$p_j = c_j - B + \left(\frac{-\partial \ln Q_j}{\partial p_j}\right)^{-1}. \quad (4)$$

Adding B to both sides reveals that an optimal bid b_j equals the plan's marginal cost c_j plus a markup term that depends on the semielasticity of the risk-weighted plan demand.

C. Discussion

We have made several assumptions that allow us to analyze the MA market in a fairly standard way, and thus to explore the ways in which market design decisions may affect equilibrium outcomes. Before we take the model to the data, we discuss the model implications with respect to specific features of the MA context.

Price sensitivity and competition.—Bidding incentives depend on the price sensitivity of plan demand. Substituting $p_j = b_j - B$, and observing that $-\partial Q_j/\partial p_j = -\partial Q_j/\partial b_j$, we can rewrite condition (4) as $b_j = c_j + (-\partial \ln Q_j/\partial b_j)^{-1}$. So the model leads to the standard result: the optimal markup will be high when plan demand is not very price sensitive and low if it is sensitive.

Effect of the benchmark rates.—The benchmark rate B plays a key role in that it largely determines the program cost to taxpayers. One can observe that under the complete information assumption, from an equilibrium perspective one can view the benchmark rate as if it is a competing bid placed (not necessarily optimally) by traditional Medicare. On the one hand, traditional Medicare enrolls the majority of beneficiaries in all markets, so a lower benchmark would imply more competitive pressure on MA bids. On the other hand, it is natural to view competing MA plans as closer substitutes to each other than to traditional Medicare, suggesting that changes in benchmarks might exert weaker competitive pressure than changes in competing MA bids. One way to summarize these forces is to investigate the "pass-through" rate of benchmark increases into plan bids. If MA plans are perfect substitutes for each other (and have similar costs), benchmark changes will not affect bidding incentives. This is why Song, Landrum, and Chernew (2012) and Cabral, Geruso, and Mahoney (2018) interpret their findings of positive benchmark pass-through rates as evidence of market power. Unfortunately, pass-through rates alone are not necessarily very informative about the degree of competition (Bulow and Pfleiderer 1983), which is one reason why a more complete equilibrium analysis is required in order to analyze the effects of changes in market design parameters.

Incentives for risk selection.—The baseline model assumes that variation in a plan's risk composition is fully compensated. However, as shown in table 1 and in figure E2, MA plans enroll individuals who are relatively healthy, even conditional on risk score. We can incorporate this by assuming that individuals with risk score r who enroll in plan j have "actual" risk $\Lambda_j(r, p_j, p_{-j}) \leq r$.

The simplest case arises if $\Lambda_j(r, p_j, p_{-j}) = \lambda r < r$. Then MA plans have favorable selection conditional on observed risk score, but the effect on bidding incentives remains the same. It is "as if" the cost of MA plan j per enrollee risk unit is λc_j rather than c_j. The equilibrium analysis can then proceed in the standard way, except that this risk selection should get accounted for when MA enrollees get covered by traditional Medicare in counterfactual simulations.

The analysis is more complicated if $\Lambda_j(\cdot)$ depends on p_j. This may be because p_j affects observed risk composition and $\Lambda_j(\cdot)$ is not proportional to r so that, for instance, unobserved selection is more important for high or low risk enrollees. This may otherwise be because p_j affects risk composition conditional on r. For example, if an increase in p_j shifts the plan's enrollment toward low-risk enrollees and plan j is relatively overcompensated for low-risk enrollees, there will be an extra incentive to bid high. In table B5, we show that changes in plan bids have little effect on a plan's risk composition so that abstracting from this complication is unlikely to affect results much. Appendix B.5 provides a detailed discussion.

IV. Demand for MA Plans

A. Demand Specification

We adopt a fairly standard nested logit specification for plan demand, a specification that has been used in the past in this context (Town and Liu 2003; Dunn 2010, 2011; Hall 2011). We adapt it slightly to capture the fact that plan revenues and costs depend on risk-weighted demand rather than the number of enrollees. We consider each county-year as a separate market, indexed by k. We divide the plans into two exhaustive and mutually exclusive categories indexed by g: the outside good (traditional fee-for-service Medicare or a PFFS plan) is the only member of group $g = 0$, and all non-PFFS MA plans belong to the other group $g = 1$.[23]

The utility of beneficiary i from plan j in market k is given by

$$u_{ijk} = \delta_{h(i)jk} + \zeta_{ig} + (1 - \sigma)\,\varepsilon_{ijk}, \tag{5}$$

where

[23] Our model somewhat simplifies the range of choices because a beneficiary who chooses traditional Medicare may or may not opt to enroll in a supplemental Medigap policy, and both TM enrollees and MA enrollees whose prescription drug coverage is not bundled into their plan may enroll in a separate Medicare Part D plan.

$$\delta_{hjk} = x'_{jk}\beta - \alpha_h^- p_{jk}^- - \alpha_h^+ p_{jk}^+ + \xi_{jk} \tag{6}$$

and h indexes four consumer groups: individuals dually eligible for Medicare and Medicaid, or "dual eligibles"; and non–dual eligibles in low, medium, and high terciles of the health risk score. In this specification, p_{jk}^- and p_{jk}^+ are, respectively, plan j's excess bid below and above the benchmark,

$$\begin{aligned} p_{jk}^- &= (b_{jk} - B_k) \cdot \mathbf{1}\{b_{jk} \le B_k\}, \\ p_{jk}^+ &= (b_{jk} - B_k) \cdot \mathbf{1}\{b_{jk} > B_k\}, \end{aligned} \tag{7}$$

and x_{jk} denotes a vector of observable plan characteristics. This includes indicators for each possible plan quality rating (e.g., 3.5 stars or 5 stars), an indicator for whether the plan is bundled with supplemental benefits (such as vision or dental coverage), and an indicator for whether the plan is bundled with Part D benefits. In our preferred specification, it also includes a fixed effect for the insurer's contract with CMS underlying plan j, which specifies the network of medical providers. Importantly, we allow the demand response to be asymmetric and to vary with individual characteristics: marginal changes in the bid above the benchmark affect the premium, whereas marginal changes in the bid below the benchmark typically affect the plan's coverage generosity. Specifically, α_h^- and α_h^+ are allowed to be different, and each can vary flexibly across the four consumer groups. To complete our specification, let ξ_j denote the mean value of the unobserved plan characteristics.

In order to derive an expression for the implied market shares, we add nested logit distributional assumptions on the stochastic terms ζ_{ig} and ε_{ijk}. We assume that the ε_{ijk} are distributed i.i.d. with a type I extreme-value distribution and that ζ_{ig} is drawn from a distribution (with parameter σ) so that $\zeta_{ig} + (1 - \sigma)\,\varepsilon_{ijk}$ follows a generalized extreme-value distribution. As shown in Berry (1994), this yields the nested logit specification

$$\ln(s_{hjk}) - \ln(s_{h0k}) = \delta_{hjk} + \sigma \ln(\bar{s}_{hjk}), \tag{8}$$

where $\bar{s}_{hjk}$ is the within-MA market share of plan j for consumer group h in market k.

To estimate this model, we measure enrollment by risk-months rather than the number of beneficiaries who choose a given plan. That is, instead of assigning a weight of 1 to each plan enrollee, we assign a weight that is proportional to the enrollee's risk score times the number of months (during the observation year) she was covered by the plan. We do this because, as described earlier, plan revenues and profits are proportional to the risk-weighted enrollment, so this measure of enrollment is the one that should enter the firm's bidding decision. We convert this measure of enrollment to market shares so that s_{hjk} is plan j's risk-month-weighted share.[24]

B. *Identifying Moments*

In estimating how plan enrollment changes with plan bids, we face a common identification challenge. Differences in plan bids may reflect differences in plan structure or quality that in turn directly affect demand for the plan. In particular, although we control for CMS quality scores, there may be considerable variation in actual or perceived plan quality. Arguably the most important component of plan quality is the plan's network of providers. Even if we could perfectly observe the network, encoding it in a usable set of control variables would be challenging and imperfect. Our preferred solution is to estimate the demand model using GMM, where we rely on a set of demand-side moments and combine it with a single supply-side moment. The latter guarantees that our demand estimates, combined with the insurers' condition for optimal bidding, match external data on the observed costs, which we obtained for a subset of the plans in our data.

Demand-side variation.—The identifying variation on the demand side uses a difference-in-differences strategy that relies on variation within an insurer contract, either over time or across counties, or both. This approach has some limitations, but one benefit is that it utilizes the structure of the program under which plans within a contract share the same provider network. It also allows us to take advantage of the large number of plan-years for which we have data.

Consider the empirical specification described above, writing the risk-month-weighted consumer group h market share of plan j in county-year k as

$$\begin{aligned} \ln(s_{hjk}) - \ln(s_{h0k}) = {} & x'_{jk}\beta - \alpha_h^- p_{jk}^- - \alpha_h^+ p_{jk}^+ \\ & + \sigma \ln(\bar{s}_{hjk}) + \eta_{\text{year}(k)} + \eta_{\text{contract}(j)} + \xi_{jk}, \end{aligned} \tag{9}$$

where x_{jk} includes the plan's CMS quality score, as well as characteristics such as whether the plan includes a prescription drug offering. Our specification makes use of two sources of variation in plan bids. First, within a given year, a contract between an insurer and CMS may include several plans that have the same provider network but vary in their exact benefits, i.e., the monthly premium or whether the plan includes Part D benefits. Controlling for benefits that are funded with supplemental fees (such as

[24] Note that this is a bit of a shortcut and does not follow immediately from the utility specification. If we modeled the enrollment share of each risk type r as nested logit, then added up over risk types, the left-hand side would be the risk-month-weighted average of the log shares, rather than the log of the risk-month-weighted share.

a bundled Part D offering), we identify price sensitivity using variation across plans within a contract in the number of benefits generated by different rebates. Second, plan bids change from year to year, and we identify price sensitivity using the idiosyncratic variation in bid changes relative to national changes. In appendix B, we explore a host of alternative strategies to estimate demand (without using the supply-side moment described below), including different sample restrictions, more granular sets of fixed effects, or various instrumental variable strategies that have been used in the literature, all of which motivate our choice of the above specification as our preferred empirical strategy.

Finally, a standard issue in estimating equation (8) is that $\bar{s}_{hjk}$ is endogenous because unobserved changes in plan demand affect both its market share and its market share among MA plans. We use several alternative instruments for $\bar{s}_{hjk}$: the first is the number of MA plans offered in the market, the second is the number of MA contracts offered in the market, the third is a set of dummy variables for which other contracts are offered in the same market, and the fourth is the number of insurers competing in the same service area. The results across choices of instruments are very stable, and the price effects are quite consistent over an even wider set of specifications, which we explore in table B3. In our preferred specification, we use the number of MA contracts offered in the market and the number of insurers competing in the same service area.

Supply-side moment.—If we rely only on the demand-side moments, we obtain demand elasticities that are "too low," in the sense that when combined with our model of optimal bidding (which is used in our counterfactual exercises below) they give rise to insurer markups that are much greater than those observed in external data. Therefore, in order to obtain our preferred demand estimates, we add a supply-side moment that relies on these external data on costs.

Specifically, as is standard in many industry studies, the condition for optimal bidding in equation (4)—modified to reflect the fact that insurers may offer multiple plans in the same market—allows us to translate demand to implied cost. Consider a given county-year with J private plans. Let D_bQ be the estimated matrix of own- and cross-bid derivatives, and let Ω denote an ownership matrix, such that its lmth entry is equal to 1 if plans l and m are owned by the same insurer, and zero otherwise. With this (standard) notation, the implied plan costs are given by the solution to the first-order conditions for optimal bidding:

$$c = b + (\Omega \cdot D_bQ)^{-1}Q. \tag{10}$$

Here the multiplication operator refers to an element-by-element product, and c, b, and Q are J-dimensional vectors of the implied costs, observed bids, and observed shares of each plan in the market.

In Curto et al. (2019) we use data on claims paid by three large Medicare Advantage insurers (Aetna, Humana, and UnitedHealthcare) and report (table 1, col. 8) insurer costs to be, on average, \$590 per enrollee-month. This average is based on a subset of the Medicare Advantage plans in 2010, all of which are part of the plans covered in the sample used in this paper. Given an estimated set of demand parameters, equation (10) allows us to compute implied costs for each plan in each market,[25] and in particular for each of the plans that contribute to the \$590 average reported in Curto et al. (2019).

Estimation.—We estimate the model using GMM and compute bootstrapped standard errors. The complete details are described in appendix C.2 and are mostly standard, and here we highlight two specific, less standard details.

The first nonstandard aspect in our estimation is that we attach a higher weight to the single supply-side moment and have adjusted the weight so that this moment is perfectly matched, so one can view our estimation procedure as a constrained optimization, where demand is estimated using demand moments alone subject to the constraint that the single supply-side moment is matched. This decision is subjective in nature, and we made it for two reasons. Most importantly, the single supply-side moment relies on markups that are directly observed in data (for a subset of the plans), while the traditional demand-side moments rely on assumed orthogonality conditions that, while reasonable in our view, are clearly imperfect, so (subjectively) we felt that it is much more important to match the observed supply-side markup, and less critical to perfectly match the assumed demand-side conditions. In addition, because there is only a single supply-side moment and many demand parameters (and associated demand-side moments), forcing the supply-side moment to match exactly still leaves many degrees of freedom and thus does not constrain any specific demand parameter in any particular way.

The second nonstandard aspect in our estimation is computational. Optimizing nonlinearly over all our demand parameters is difficult given the large number of contract fixed effects (and other controls). Instead, we partition the model parameters into two sets. The first includes the price coefficients (we have eight of them, α^- and α^+ for each one of the four consumer groups) and the σ parameter, and the second includes everything else. This is useful because it allows us to take advantage of the property of the logit (and nested logit) specification, which implies that optimal markups are only a function of observed market

[25] Because bids are for providing standard Medicare benefits to a standard Medicare Advantage enrollee, the implied costs are also standardized. To compute plan costs on a per-enrollee basis, we simply multiply the implied (standardized) cost by the average risk score of the plan's enrollees, which is observed.

shares and the first set of coefficients (the α's and σ). We can then nonlinearly optimize over these parameters alone, and then (in each iteration of the optimization process) obtain a closed-form solution for the other parameters (as a function of the values of the α's and σ), which makes computation feasible.

C. Demand Estimates

Table 2 reports results from several demand specifications, leading to our preferred specification in column 4, which is the one we use in our counterfactual exercises in the next section. As mentioned earlier, results from a host of other specifications are reported in appendix B and tables B1–B4.

The main coefficients to emphasize capture the effect of the changes in the plan bid on plan enrollment. Recall that the vast majority of plans bid below the benchmark. In this range, the results imply that a bid increase of $10 decreases the plan's risk-month-weighted market share by around 9%. Recall that when the bid is below the benchmark, marginal changes in the bid primarily translate to changes in plan coverage generosity. Consequently, in this range of bids, dual eligibles (who are exempted from large components of the plan's cost sharing) are the least sensitive to changes in the bid, and high-risk-score individuals are the most sensitive given that they are expected to utilize the plan coverage the most.

When the plan bid is above the benchmark (which is not common), changes in bid translate to changes in the consumer premium. The demand elasticity in this range is approximately doubled relative to changes below the benchmark. This difference arises for two different reasons. First, some of this increased sensitivity is "mechanical" given that only 75% of changes in the bid below the benchmark are passed through to enrollees (in the form of benefits), while pass-through is 100% for bids above the benchmark. Second, individuals appear to value a $1 reduction in premium more than they value a $1 increase in expected coverage. For bids above the benchmark, there is limited heterogeneity in demand elasticity across consumer groups, with the exception of nondual high-risk-score enrollees, who tend to be less elastic than others.

Using this specification, table B3 reports nested logit estimates using several choices of instruments (for the within-MA market share). The results are extremely similar across columns, giving rise to similar demand elasticities with respect to the excess price. We also estimate the nested logit parameter σ to be significant, at around 0.23, consistent with our expectation that MA private plans are closer substitutes to one another than they are to traditional Medicare. The implication is that for a given

TABLE 2

DEMAND ESTIMATES

	Dependent Variable: ln(plan share) − ln(TM share)			
	(1)	(2)	(3)	(4)
$(b - B) \times I\{b \leq B\}$	−.002***			
	(<.001)			
$(b - B) \times I\{b \leq B\} \times$ nondual $\times$ lowest risk tercile		−.005***	−.006***	−.007***
		(<.001)	(.001)	(.001)
$(b - B) \times I\{b \leq B\} \times$ nondual $\times$ middle risk tercile		−.007***	−.007***	−.006***
		(<.001)	(.001)	(<.001)
$(b - B) \times I\{b \leq B\} \times$ nondual $\times$ highest risk tercile		−.016***	−.015***	−.011***
		(.001)	(.002)	(.001)
$(b - B) \times I\{b \leq B\} \times$ dual		−.001*	−.001*	−.004***
		(.0005)	(.001)	(.0005)
$(b - B) \times I\{b > B\}$	−.027***			
	(<.001)			
$(b - B) \times I\{b > B\} \times$ nondual $\times$ lowest risk tercile		−.010***	−.010***	−.010***
		(<.001)	(<.001)	(<.001)
$(b - B) \times I\{b > B\} \times$ nondual $\times$ highest risk tercile		−.004***	−.004***	−.004***
		(<.001)	(<.001)	(<.001)
$(b - B) \times I\{b > B\} \times$ dual		−.010***	−.010***	−.010***
		(<.001)	(<.001)	(<.001)
ln(plan MA share)	.226***	.227***	.227***	.226***
	(<.001)	(<.001)	(<.001)	(<.001)

Supplemental benefits	.015	.224***	.220***	.068***
	(.019)	(.031)	(.029)	(.016)
Part D benefits	.693***	.649***	.689***	1.166***
	(.023)	(.023)	(.020)	(.020)
Mean of dependent variable	−4.515	−4.515	−4.515	−4.515
Mean elasticity (risk-month weighted)	−2.137	−7.686	−7.472	−6.591
Number of observations	106,400	106,400	106,400	106,400
Number of clusters (markets)	8,886	8,886	8,886	8,886
Supply-side moment	.635	1.157	.608	.044
Bootstrap mean of supply-side moment	.659	.725	1.246	.084
Number of bootstrap samples	50	50	50	50
Instruments for ln(plan MA share)	Number of contracts in county, number of insurers in service area	Number of contracts in county, number of insurers in service area	Number of contracts in county, number of insurers in service area	Number of contracts in county, number of insurers in service area
Plan rating fixed effects (for each half-star)	Yes	Yes	Yes	Yes
Year fixed effects	No	No	Yes	Yes
Contract fixed effects	No	No	No	Yes

Note.—The table presents demand estimates. The estimates are based on the GMM results at the market-plan-group level, where a demographic group is dual eligibles or non–dual eligibles in the lowest, middle, and highest risk terciles. The unit of observation is a market-plan group (a market is a county-year). Bootstrapped standard errors are shown, where bootstrap samples are drawn at the market level. The plan's standardized bid is denoted by b and the plan's benchmark is denoted by B. The table reports the mean of the own-price demand elasticities with respect to b (weighted by enrollee risk-months; further details are provided in appendix C).

* $p < .1$.

*** $p < .01$.

plan, the other competing MA plans in its market (which have around 20% market share) exert almost the same competitive effect as traditional Medicare (with 80% market share). That is, an increase in the plan bid causes enrollees to substitute away in roughly equal shares to other MA plans and to traditional Medicare. Similarly, a given plan attracts about the same number of enrollees from a uniform increase in competing MA plan bids as it does from the same increase in the benchmark rate, holding all plan bids fixed.

The remaining results are generally as expected. We find that consumers place considerable value on plans that offer bundled Part D benefits. The table does not display the effects of the plan quality ratings, but they are clearly predictive of enrollment. Most ratings are from three to four stars, and, all else equal, we estimate that a four-star plan attracts around 31% more enrollees than a three-star plan, equivalent to a bid difference of about $33 per month. The coefficient on supplemental benefits is insignificant, presumably reflecting the fact that these increased benefits are "funded" by an additional premium, so the two effects offset each other.

V. Welfare Effects of Medicare Advantage

In this section, we use our demand estimates to quantify enrollee surplus from private plans and to back out implied costs of each MA plan.[26] We then compare these costs with predicted TM costs associated with the same beneficiaries, which we construct. Taken together, this exercise provides estimates of the overall "gains from trade" from having MA plans, and estimates of how this surplus is divided among taxpayers, plans, and beneficiaries. The following section considers how these calculations might change under various reforms of the bidding rules. We focus our exposition on the main findings, leaving some details of the calculations to appendix C.

A. *Estimates of Enrollee Surplus and Plan Costs*

We first consider the surplus that accrues to private plan enrollees. The nested logit formula for the consumer surplus associated with the MA plans offered in county-year k is

$$CS_k = \sum_i \frac{1}{\alpha^{+}_{h(i)}} \ln\left(1 + \exp\left\{(1-\sigma)\ln\left[\sum_{j=1}^{J_k} \exp\left(\frac{\delta_{h(i)jk}}{1-\sigma}\right)\right]\right\}\right). \qquad (11)$$

[26] Recall that our supply-side moment only matched the average costs for a subset of plans for which we observe costs in 2010. The resultant estimates allow us to back out implied costs for all plan-year combinations in our sample.

The parameter $\alpha^{+}_{h(i)}$ is the marginal utility of income, so that consumer surplus is measured in dollar terms. We compute this for every county-year. The estimates range from 30 cents to $98 per Medicare risk-month. Of course, this surplus is concentrated among beneficiaries who actually enroll in MA plans. Attributing it to this narrower group implies an enrollee surplus of $88 per MA risk-month.

We then combine our demand estimates with our model of equilibrium bidding (as described in sec. IV.C) to obtain the implied plan costs and markup. We take the risk-weighted average across plans for each county-year to obtain an average monthly plan profit for each market. Figure 4*A* plots these market-by-market estimates of private plan profits (i.e., markups) against MA benchmarks for the same markets. The estimated markups are substantial. We estimate the average profits accrued by MA plans for every month of enrollment to be $129. On a percentage basis, this translates to plan margins of approximately 21% above (variable) costs on average, which is (by design) similar to the results in Curto et al. (2019).[27]

B. Constructing TM Cost Benchmarks

We are interested in how the estimated private plan costs compare to equivalent TM costs. We can make this comparison at the county-year level. For each county-year, we have an estimate of the enrollment-weighted average of private plan costs, and we describe below how we use our data to compute the counterfactual costs had the same set of MA enrollees been covered instead by traditional Medicare.

We build our cost benchmarks from measures of local TM costs for each US county in each year.[28] Let k index county-years, and let TM_k denote the set of individuals enrolled in TM in county-year k. For a given individual i, let m_i denote the number of months that i was enrolled in TM during the year, and x_i denote i's average monthly TM costs during those months. We define two measures of local TM spending in county-year k:

$$\gamma_k = \frac{\sum_{i \in TM_k} x_i m_i}{\sum_{i \in TM_k} m_i} \qquad \text{and} \qquad a_k = \frac{\sum_{i \in TM_k} x_i m_i}{\sum_{i \in TM_k} r_i m_i}. \tag{12}$$

[27] The reported numbers in Curto et al. (2019) imply a 19% markup on average; it is not exactly the same because the estimates reported here use the Curto et al. (2019) statistics to estimate the demand coefficients, but then apply these demand estimates to a much larger set of plans and years.

[28] Our construction of these measures is somewhat different from the one used by CMS for public reporting (MedPAC 2012) in order to account for mortality and midyear enrollment, and the fact that the average risk score of TM enrollees in our sample is not identically equal to 1. Appendix D provides more detail on the differences between our measures and the ones reported by MedPAC.

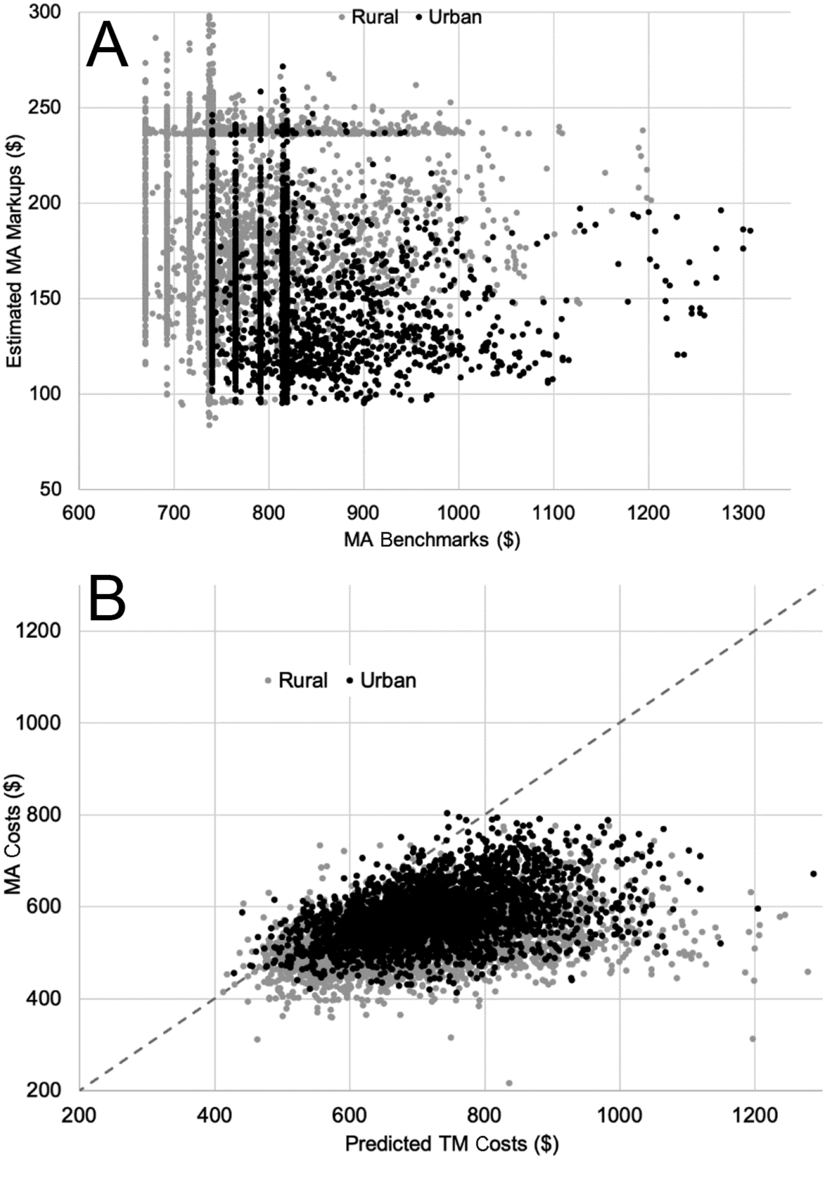

Fig. 4.—Estimated markups and costs of MA plans. Observations are at the market (county-year) level; there are 8,884 markets with MA enrollees in the analysis sample as well as predicted FFS costs (which are not available for 2011). Counties that qualified as "urban" according to the definition used to set the urban floor in 2004 are shown in black, while all other counties are shown in gray. Estimates are derived from the insurer's first-order condition for optimal bidding (see text for details), using the preferred demand specification (table 2, col. 4). *A* presents estimated markups against county benchmarks. Market-level markups are obtained by weighting market-plan-level markups by MA enrollee-months. Statistics

The first measure is the local TM cost per enrollee-month. The second is the local TM cost per enrollee risk-month. Note that $\gamma_k = a_k r_k$, where $r_k = \sum_{i \in TM_k} r_i m_i / \sum_{i \in TM_k} m_i$ is the (month-weighted) average enrollee risk in county-year k.

We then construct estimates of TM costs for MA enrollees in two steps, first imputing TM costs based on each enrollee's location and risk score, and then adjusting for health differences conditional on risk score. The procedure is similar to the one used in Curto et al. (2019). For the first step, we assign each individual i in county-year k an expected monthly TM cost of $a_k r_i$, where r_i is the individual's risk score and a_k is the local TM cost per enrollee risk-month.[29] For the second step, we rescale MA risk scores using a single, year-specific scaling factor λ_t in a way that aligns conditional (on risk score) TM and MA mortality rates. We describe the precise procedure in appendix D. We then adjust the predicted TM costs of each MA enrollee in our sample by the relevant scaling factor. Thus, an MA enrollee in year t and county-year k, with risk score r_i, has predicted TM costs of $\lambda_t a_k r_i$. We note that, in principle, the rescaled prediction should account for both self-selection of healthier individuals into MA and differences in disease coding. The key assumption is that differences in mortality rates can proxy for differences in health status (and in particular in expected medical claims).

Figure 4*B* shows a comparison of MA and TM costs for covering the same set of MA enrollees. Each point in the figure represents a county-year, with the horizontal axis representing the predicted average TM cost (per risk-month) of MA enrollees in the market, and the vertical axis representing the average implied private plan costs (per risk-month) of the same MA enrollees. The dotted line is the 45° line, so that points above it represent markets in which MA is estimated to be more expensive than TM, while points below the line are where MA appears to have the potential to generate cost savings.

are weighted by market-level MA enrollee-months. The mean markup is 128 (standard deviation of 23). The mean benchmark is 842 (standard deviation of 95). *B* presents estimated MA cost against estimated FFS cost for the same set of Medicare beneficiaries. For clarity, predicted Medicare FFS costs and implied MA costs above 1,300 are not shown. Implied MA costs are derived from the insurer's first-order condition for optimal bidding (see text for details), using the preferred demand specification (table 2, col. 4). The market-level implied cost is computed as an enrollment-weighted mean of the implied costs of all plans bidding in the market. The predicted FFS costs for MA enrollees is obtained using the method described in the text. The mean implied MA cost is 605 (standard deviation of 66), and the mean difference between implied MA cost and the predicted FFS cost is −160.

[29] This construction is motivated by an underlying model of monthly TM costs in which individual i in county-year k has stochastic costs equal to $A_k r_i \varepsilon_i$, where A_k is the expected monthly cost of a risk score 1 individual in county-year k, and ε_i is a stochastic term with $E[\varepsilon_i | r_i, k] = 1$. In this case, our sample statistic a_k will be a consistent estimator of A_k.

As one can see, MA is cheaper than TM in almost all (97%) markets. Overall, we estimate the average cost of MA enrollees to be $607 per enrollee risk-month, about $152 (20%) less than our estimate of the cost to cover the same enrollees in traditional Medicare. The difference is substantial: with the current level of MA enrollment (20.4 million enrollees in 2018), it implies about $38 billion per year in potential cost savings.

C. *Who Benefits from the MA Program*

Summarizing the welfare effects from the observed MA program.—We summarize our estimates of program surplus and how it is divided in the first row of table 3. Over our study period, taxpayer costs were $805 per MA enrollee risk-month. We estimate that the private plan cost of providing baseline coverage to these beneficiaries was only $607. The difference, equal to $197 per enrollee risk-month, is split between the beneficiary and the insurer. We estimate average monthly plan profits to be $129, the average rebate that went to enrollees was $81, and the average premium paid by enrollees was less than $1. We estimate that the average consumer rebate surplus is $63 per month (not shown in the table), or nearly half of plan profits. This implies that the enrollee disutility from limited provider access is $18 per month. Put another way, about two-thirds of the joint surplus to plans and enrollees accrues to the plans.

The comparison to traditional Medicare is also informative. We estimate that the average Medicare Advantage enrollee over this period would have cost taxpayers $759 per enrollee risk-month under traditional Medicare. Therefore, of the $197 of net surplus that accrued to plans and enrollees, only $134 per enrollee risk-month were true “gains from trade.” This amount represents the cost savings ($152) minus the enrollee disutility from private plan restrictions ($18). The remaining surplus to plans and beneficiaries comes from taxpayers who, relative to traditional Medicare, provide a subsidy of $63 per enrollee risk-month. The subsidy reflects the fact that plan payments have somewhat exceeded TM costs for the same enrollees.

Translating these monthly per-enrollee amounts into annual dollars, the potential surplus gains from private plans are not trivial. At enrollment levels of around 20.4 million Medicare beneficiaries in private plans in 2018, the program surplus amounts to about $48 billion per year, with $33 billion per year in true gains from trade. Taxpayers provide a $15 billion subsidy. The result is that private plans receive around $32 billion in annual surplus, while enrollee gains are about half this amount, or $15 billion per year. Of course, our calculations ignore all of the fixed costs of operating private plans, and may overstate the potential gains to the extent that these are larger than the administrative costs associated with traditional Medicare, although these seem to be relatively small (Curto et al. 2019) and are therefore unlikely to affect the overall conclusions.

Exploring the impact of changes in market design.—Our estimates identify two factors that contribute to high taxpayer spending on private plan enrollments. The first is that MA plan costs may sometimes exceed TM costs, despite being lower on average. The second is that plans have limited incentives to bid aggressively to increase enrollment. So cost savings translate largely into insurer surplus. As a result, it is interesting to explore how market design changes might alter this situation. Here, we report on a number of alternative policy simulations using the estimated model. These calculations all hold fixed the set of MA plans and their plan designs, and in practice some changes in market design could prompt insurers to add or drop plans, or to adjust plan details (see, e.g., Frakt, Pizer, and Feldman 2012). The results we report should be interpreted with this caveat in mind.

We start by demonstrating how, under the current rebate rules, plan bids largely determine the split of surplus between insurers and enrollees, without having much effect on taxpayer costs per enrollee-month. To see this, we first consider an extreme case where MA plans bid their costs. As shown in table 3 (row 2), this results in only a $38 per enrollee risk-month (4.7%) reduction in government spending (the difference calculated from the table differs slightly due to rounding). Rebates jump up, however, leading to more generous coverage and an MA enrollment rate of 29%. Although the marginal enrollees derive less surplus than the inframarginal ones, average enrollee surplus increases to $209 per month. At the other extreme, we consider a case where plans bid exactly the benchmark (row 3 of table 3), with zero rebates. In this case, MA enrollment drops almost in half, showing that more generous coverage is crucial to attract enrollees. Government spending is a bit higher, by $39 (4.8%), while the division of surplus between insurers and enrollees shifts sharply toward insurers.

It follows that changes in the intensity of competition, holding fixed the program rules, only have a moderate effect on per-enrollment taxpayer costs. This leads us to focus on two main program design parameters, the benchmark rates and rebate formula. In each exercise, we change the program design parameters and use our model and estimates of demand and costs to recompute the bidding equilibrium.

Rows 4–6 of table 3 report results from changes in the benchmark rates, holding the rebate formula fixed. Uniform changes in the benchmark rates primarily translate to changes in MA enrollment rates. A uniform increase in benchmarks by $100, shown in row 6, implies a large increase in MA share of nearly 10 percentage points (64%) and a $67 (8.3%) increase in taxpayer spending per enrollee. Consumer surplus is naturally higher, rising by $83 (94%), and insurer profits per enrollee remain approximately

TABLE 3
SUMMARY OF RESULTS AND POLICY EXPERIMENTS

A. Results for Entire Sample of Medicare Advantage Enrollees

Description of Policy Experiment	Taxpayer Cost[a] (1)	MA Cost[b] (2)	Predicted TM Cost[c] (3)	Insurer Profits (4)	MA Share (5)	Consumer Surplus (6)
1. No change	805	607	759	129	15.4	88
2. Plans bid cost	767	607	759	0	29.2	209
3. Plans bid benchmark	843	607	759	236	8.5	44
4. Benchmarks lowered by $200	656	597	764	132	4.3	26
5. Benchmarks lowered by $100	734	603	761	129	8.4	45
6. Benchmarks raised by $100	872	609	757	133	25.3	172
7. Rebate passed through at 25%	826	615	742	217	8.2	42
8. Rebate passed through at 50%	810	612	751	182	10.0	52
9. Rebate passed through at 100%	826	604	762	101	22.7	150
10. Benchmarks lowered by $75, rebate passed through at 100%	764	600	769	106	13.0	75
11. Benchmarks lowered by $50, rebate passed through at 100%	784	601	766	103	15.9	95
12. Benchmarks lowered by $25, rebate passed through at 100%	805	603	764	101	19.2	120

B. Results by Demographic Groups of Medicare Advantage Enrollees

	MA Share: Low Risk (1)	MA Share: Medium Risk (2)	MA Share: High Risk (3)	MA Share: Dual Eligible (4)	Consumer Surplus: Low Risk (5)	Consumer Surplus: Medium Risk (6)	Consumer Surplus: High Risk (7)	Consumer Surplus: Dual Eligible (8)
1. No change	19.3	18.5	14.3	13.6	51	48	131	33
2. Plans bid cost	32.2	30.0	31.1	20.3	93	85	336	51
3. Plans bid benchmark	12.2	11.7	6.1	9.8	32	30	56	24

4. Benchmarks lowered by $200	5.1	4.4	3.9	4.7	15	12	39	13
5. Benchmarks lowered by $100	11.2	10.5	6.8	9.2	29	26	62	23
6. Benchmarks raised by $100	28.6	27.0	26.6	18.2	81	74	272	45
7. Rebate passed through at 25%	11.6	11.1	5.9	9.3	30	28	55	23
8. Rebate passed through at 50%	13.8	13.4	7.9	10.7	36	33	71	26
9. Rebate passed through at 100%	26.1	24.7	23.4	17.2	73	67	237	43
10. Benchmarks lowered by $75, rebate passed through at 100%	16.3	15.4	12.0	12.1	43	40	111	30
11. Benchmarks lowered by $50, rebate passed through at 100%	19.4	18.4	15.2	13.8	52	48	144	34
12. Benchmarks lowered by $25, rebate passed through at 100%	22.7	21.6	19.1	15.5	62	57	186	38

Note.—All reported statistics are per MA enrollee-month, and are calculated using the preferred demand specification (table 2, col. 4). The taxpayer cost is total MA payments, including rebate payments. Implied MA cost is computed using the insurer's first-order condition, and incorporates co-ownership of multiple plans by the same MA parent organization. Insurer profits are computed by adding payments made directly to plans (excluding rebate payments) and enrollee premiums, and subtracting implied MA cost. The consumer surplus is computed using the nested logit formula given in the text. Consumer surplus is divided by MA share to obtain the consumer surplus per MA enrollee. For the counterfactuals, the reported means are weighted using enrollment weights that are recalculated according to the market shares implied by the counterfactual bids.

[a] The table reports taxpayer costs that are calculated using realized enrollment weights and actual bids to impute MA payments, which differ slightly from observed taxpayer costs reported in table 1.

[b] The table reports MA costs for 2006–11, which differ slightly from those reported in fig. 4, which uses data from 2006–10.

[c] Data are for 2006–10 because 2011 TM claims were not available.

the same, rising by only $4 (3%) (but, recall, at a much higher rate of enrollment). Similarly, a uniform reduction in the benchmark rates implies lower taxpayer spending per enrollee but also significantly lower MA enrollment rates, eliminating nearly all of the consumer surplus associated with Medicare Advantage. The key observation is that raising (lowering) the benchmark rates has a large and positive (negative) effect on all outcomes of interest: MA enrollment, insurer profits, taxpayer spending, and consumer surplus.

Rows 7–9 of table 3 report results from changes in the other market design instrument, the rebate formula, holding benchmarks fixed. To examine this, we apply different rebate pass-through rates (50% or 100%, instead of 75%) to the difference between the bid and benchmark; recall that these rebate dollars go toward more generous plan benefits. These changes have two competing effects. Holding bids fixed, a larger rebate increases government spending and increases enrollee benefits. At the same time, plans have an incentive to reduce their bids because a larger share of any reduction is passed through to enrollees, making demand more responsive to bids. This leads to lower equilibrium markups. The results in table 3 suggest that the effects are fairly strong, with increased rebate pass-through rates leading to a rise in consumer surplus and a decline in insurer profits.

These results suggest that combining the two policy instruments may be fruitful. By combining an increase in the rebate pass-through rate with a reduction in benchmark rates, one can keep taxpayer costs at similar levels but increase MA enrollment and consumer surplus. At the same time, this change can reduce insurer profits, "shifting" surplus from insurers to consumers, which (arguably) may be a desirable objective. Rows 10–12 of table 3 show several policy experiments that combine an increase in the rebate pass-through rate with a reduction in benchmark rates. Row 12 of table 3 reports results from lowering benchmark rates by $25 and raising the rebate pass-through rate from 75% to 100%: we estimate 0 change in taxpayer cost, a $28 (22%) decrease in insurer profits, a 4 percentage point (24%) increase in the MA share, and a $32 (36%) increase in consumer surplus.

Figure 5 summarizes these and several additional market design counterfactuals by plotting, for each policy experiment, the estimated taxpayer savings against MA enrollment share. One can see that many counterfactuals span the "frontier," where higher taxpayer spending on the MA program leads to higher MA enrollment (and greater consumer surplus). Yet, figure 5 also highlights several policy experiments that are inside the frontier, and are dominated by alternative market design policies.[30]

[30] The counterfactual exercises reported in this section focus on illustrating the conceptual trade-offs associated with the two key market design instruments. As a referee pointed

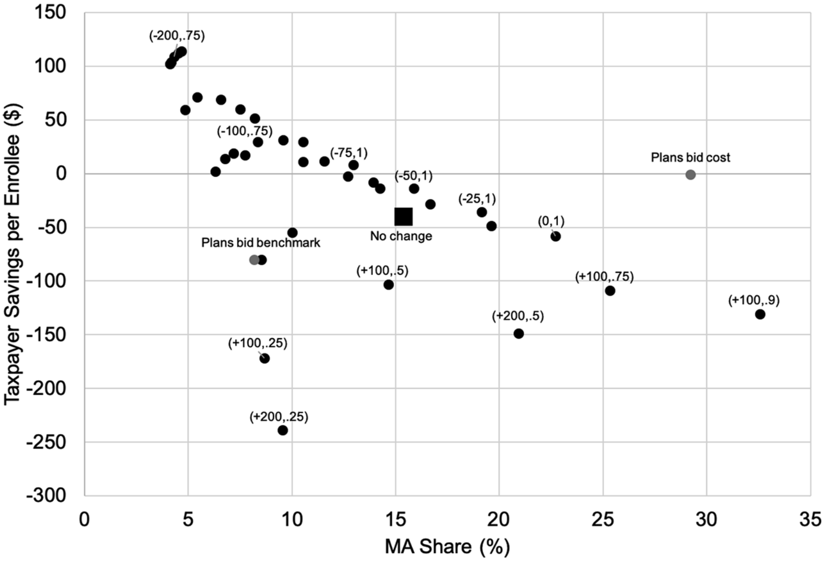

Fig. 5.—Predicted outcomes under different market design parameters. The figure summarizes results from equilibria under a host of counterfactual market designs. The large square represents our estimates under the current MA system, and each circle represents a different counterfactual. The figure shows the resulting taxpayer savings per enrollee-month and MA share from various counterfactual values of the benchmark and rebate pass-through rate. Some counterfactuals are labeled, where the first entry is the benchmark change and the second entry is the rebate pass-through rate. The figure includes all counterfactuals listed in tables 3 and C1. The "no change" counterfactual is represented by a black square. The "plans bid cost" and "plans bid benchmark" counterfactuals are represented by gray circles, as it is not feasible to implement these outcomes, but they are shown as a benchmark for illustrative purposes. All other counterfactuals are represented by black circles. "Taxpayer savings" is defined as taxpayer costs minus predicted TM costs. Appendix C provides details on how counterfactual outcomes are calculated.

Finally, panel B of table 3 reports how the various counterfactual exercises affect different subgroups of consumers. While, naturally, all consumers are affected in the same direction, the quantitative effects are by far the largest on the high-risk, nondual consumers, who are estimated to experience large shifts in MA enrollment rates and consumer surplus as a result of the changes in the market design instruments.

---

out, our model and estimates also allow us to investigate the impact of specific changes that were made in the MA bidding incentives after our observation period. For example, we use our model to simulate the impact of varying the rebate pass-through rate by each plan's quality rating (see https://innovation.cms.gov/Files/reports/maqbpdemonstration-finalevalrpt.pdf), and we find (not reported) that these specific changes would, in fact, lead to higher markups, higher insurer profits, and lower consumer surplus.

VI. Conclusions

The reform and expansion of Medicare Advantage provides an opportunity to evaluate how private health insurance competition can work, compared to a public health insurance alternative. We propose a model of equilibrium bidding under imperfect competition, which accommodates current program rules. Using this model and estimates of demand, we estimate the welfare gains from the program and assess the equilibrium impact of changes to the two key market design parameters in the Medicare Advantage marketplace.

Our estimates suggest that private plans are able to obtain cost savings over traditional Medicare, sufficiently large to generate net surplus even after accounting for consumer disutility from having limited provider access. The division of the surplus largely benefits the plans, with enrollees capturing only about one-third of the generated surplus, and taxpayers sharing some of the costs. Yet, we identify a set of changes to the market design parameters that could enhance competition and shift some of the surplus from insurers to consumers without increasing overall taxpayer costs.

In modeling any large and complex program, there is a trade-off between capturing all of the potentially interesting dynamics, and keeping the analysis transparent and tractable. It would be interesting to expand on some of the aspects of plan competition that we simplified. First, we adopted a simple approach to rescale MA risk scores and assumed that this scaling was common across plans and invariant to plan bids. Second, we focused on the response of enrollee demand to plan bids, ignoring the fact that plans can differentiate by allocating the rebate dollars toward premium reductions or increased cost sharing. Third, our model is static, whereas enrollees may view their current plans (or traditional Medicare) as a sticky default choice, so that plans face dynamic bidding incentives (Miller 2019; Miller et al. 2019). Incorporating these complications would give a more nuanced view of competition.

Our treatment of potential market design interventions could also be extended. We focus on simple changes to benchmark rates and rebate pass-through rates, and we do not address the important issue of how program rules affect other aspects of plan design, plan quality, or plan entry and exit (Miller et al. 2019). In addition to looking at market structure, one question is whether there are ways to improve plan incentives by making enrollee demand more price sensitive. Another question is whether it might be possible to induce some competition across, as well as within, local markets, given that the national market has lower insurer concentration. Another market design proposal, which is outside the scope of this paper, is to allow benchmark rates to be a function of insurer bids. Finally, it is important to emphasize that our entire analysis keeps the availability of traditional Medicare as a public option. Assessing the impact of the Medicare Advantage program in the presence of a different public option or nonexistent public option would be an interesting avenue for future work.

References

Berry, Steven. 1994. "Estimating Discrete Choice Models of Product Differentiation." *RAND J. Econ.* 25 (2): 242–62.

Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston. 2014. "How Does Risk Selection Respond to Risk Adjustment? New Evidence from the Medicare Advantage Program." *A.E.R.* 104 (10): 3335–64.

Bulow, Jeremy, and Paul Peiderer. 1983. "A Note on the Effect of Cost Changes on Prices." *J.P.E.* 91 (1): 182–85.

Cabral, Marika, Michael Geruso, and Neale Mahoney. 2018. "Do Larger Health Insurance Subsidies Benefit Patients or Producers? Evidence from Medicare Advantage." *A.E.R.* 108 (8): 2048–87.

Caplin, Andrew, and Barry Nalebuff. 1991. "Aggregation and Imperfect Competition: On the Existence of Equilibrium." *Econometrica* 59 (1): 25–59.

CMS (Centers for Medicare and Medicaid Services). 2013. *Medicare Managed Care Manual*, publication 100–16. http://www.cms.gov.

Curto, Vilsa, Liran Einav, Amy Finkelstein, Jonathan Levin, and Jay Bhattacharya. 2019. "Healthcare Spending and Utilization in Public and Private Medicare." *American Econ. J.: Appl. Econ.* 11 (2): 302–32.

Dowd, Bryan E., Roger Feldman, and Robert Coulam. 2003. "The Effect of Plan Characteristics on Medicare+Choice Enrollment." *Health Services Res.* 38 (1): 113–35.

Duggan, Mark, Amanda Starc, and Boris Vabson. 2016. "Who Benefits When the Government Pays More? Pass-Through in the Medicare Advantage Program." *J. Public Econ.* 141 (9): 50–67.

Dunn, Abe C. 2010. "The Value of Coverage in the Medicare Advantage Insurance Market." *J. Health Econ.* 29 (6): 839–55.

———. 2011. "The Effect of Health Insurance Competition When Private Insurers Compete with a Public Option." Working paper, Bur. Econ. Analysis, US Dept. Commerce.

Eggers, Paul W., and Ronald Prihoda. 1982. "Pre-Enrollment Reimbursement Patterns of Medicare Beneficiaries Enrolled in 'At-Risk' HMOs." *Health Care Financing Rev.* 4 (1): 55–73.

Enthoven, Alain. 1993. "The History and Principles of Managed Competition." *Health Affairs* 12 (Suppl.): 24–48.

Frakt, Austin B., Steven D. Pizer, and Roger Feldman. 2012. "The Effects of Market Structure and Payment Rate on the Entry of Private Health Plans in the Medicare Market." *Inquiry* 49 (15): 15–36.

GAO (United States Government Accountability Office). 2013. "Medicare Advantage: Substantial Excess Payments Underscore Need for CMS to Improve Accuracy of Risk Score Adjustments." Report to Congress GAO-13-206.

Geruso, Michael, and Timothy Layton. 2020. "Upcoding: Evidence from Medicare on Squishy Risk Adjustment." *J.P.E.* 128 (3): 984–1026.

Hall, Anne E. 2011. "Measuring the Return on Government Spending on the Medicare Managed Care Program." *B.E. J. Econ. Analysis and Policy* 11 (2): 7.

Huang, Jennifer T., Gretchen A. Jacobson, Tricia Neuman, Katherine A. Desmond, and Thomas Rice. 2013. "Medigap: Spotlight on Enrollment, Premiums, and Recent Trends." Report, Kaiser Family Found.

Kluender, Raymond, and Evan Mast. 2017. "Information Frictions and Insurer Plan Design: Evidence from Medicare Advantage." Manuscript, Bus. School, Harvard Univ.

Lustig, Joshua. 2009. "The Welfare Effects of Adverse Selection in Privatized Medicare." Manuscript, Dept. Econ., Boston Univ.

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

MedPAC (Medicare Payment Advisory Commission). 2009. "Medicare Advantage." In *A Data Book: Health Care Spending and the Medicare Program*, sec. 9. Washington, DC: MedPAC.

———. 2013. "Medicare Advantage." In *A Data Book: Health Care Spending and the Medicare Program*, sec. 9. Washington, DC: MedPAC.

———. 2012. "Issues for Risk Adjustment in Medicare Advantage." In *Report to the Congress: Medicare and the Health Care Delivery System*, chap. 4. Washington, DC: MedPAC.

Miller, Keaton. 2019. "Estimating Costs When Consumers Have Inertia: Are Private Medicare Firms More Efficient?" Manuscript, Dept. Econ., Univ. Oregon.

Miller, Keaton, Amil Petrin, Robert Town, and Michael Chernew. 2019. "Optimal Managed Competition Subsidies." Manuscript, Dept. Econ., Univ. Oregon.

Morrissey, Michael A., Meredith L. Kilgore, David J. Becker, Wilson Smith, and Elizabeth Delzell. 2013. "Favorable Selection, Risk Adjustment, and the Medicare Advantage Program." *Health Services Res.* 48 (3): 1039–56.

Newhouse, Joseph P., Mary Price, Jie Huang, J. Michael McWilliams, and John Hsu. 2012. "Steps to Reduce Favorable Risk Selection in Medicare Advantage Largely Succeeded, Boding Well for Health Insurance Exchanges." *Health Affairs* 31 (12): 2618–28.

Newhouse, Joseph P., Mary Price, J. Michael McWilliams, John Hsu, and Thomas McGuire. 2015. "How Much Favorable Selection Is Left in Medicare Advantage?" *American J. Health Econ.* 1 (1): 1–26.

Nosal, Kathleen. 2012. "Estimating Switching Costs for Medicare Advantage Plans." Manuscript, Dept. Econ., Univ. Mannheim.

Pizer, Steven D., and Austin Frakt. 2002. "Payment Policy and Competition in the Medicare+Choice Program." *Health Care Financing Rev.* 24 (1): 83–94.

Pizer, Steven D., Austin Frakt, and Roger Feldman. 2009. "Nothing for Something? Estimating Cost and Value for Beneficiaries from Recent Medicare Spending Increases on HMO Payments and Drug Benefits." *Internat. J. Health Care Finance and Econ.* 9 (1): 59–81.

Song, Zirui, David M. Cutler, and Michael E. Chernew. 2012. "Potential Consequences of Reforming Medicare into a Competitive Bidding System." *J. American Medical Assoc.* 308 (5): 459–60.

Song, Zirui, Mary Beth Landrum, and Michael E. Chernew. 2012. "Competitive Bidding in Medicare: Who Benefits from Competition?" *American J. Managed Care* 18 (9): 546–52.

———. 2013. "Competitive Bidding in Medicare Advantage: Effect of Benchmark Changes on Plan Bids." *J. Health Econ.* 32 (6): 1301–12.

Stockley, Karen, Thomas McGuire, Christopher Afendulis, and Michael E. Chernew. 2014. "Premium Transparency in the Medicare Advantage Market: Implications for Premiums, Benefits, and Efficiency." Working Paper no. 20208 (June), NBER, Cambridge, MA.

Town, Robert, and Su Liu. 2003. "The Welfare Impact of Medicare HMOs." *RAND J. Econ.* 34 (4): 719–36.