

This work is distributed as a Discussion Paper by the
STANFORD INSTITUTE FOR ECONOMIC POLICY RESEARCH



SIEPR Discussion Paper No. 14-015

Can Health Insurance Competition Work? Evidence from Medicare Advantage

By

Vilsa Curto, Liran Einav, Jonathan Levin, and Jay Bhattacharya

Stanford Institute for Economic Policy Research
Stanford University
Stanford, CA 94305
(650) 725-1874

The Stanford Institute for Economic Policy Research at Stanford University supports research bearing on economic and public policy issues. The SIEPR Discussion Paper Series reports on research and policy analysis conducted by researchers affiliated with the Institute. Working papers in this series reflect the views of the authors and not necessarily those of the Stanford Institute for Economic Policy Research or Stanford University

Can Health Insurance Competition Work?

Evidence from Medicare Advantage*

Vilsa Curto, Liran Einav,
Jonathan Levin, and Jay Bhattacharya[†]

December 2014

Abstract. We estimate the economic surplus created by Medicare Advantage 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 find that private plans have costs around 12% below fee-for-service costs, and generate around \$50 in surplus on average per enrollee-month, after accounting for the disutility due to enrollees having more limited choice of providers. Taxpayers provide a large additional subsidy, and insurers capture most of the private gains. We use the model to evaluate possible program changes.

*We thank Amy Finkelstein, Eugenio Miravete and many seminar participants for useful comments. Bhattacharya would like to thank the National Institute on Aging grants R37AG036791 and R21AG041112 for funding for his work on this project.

[†]Department of Economics, Stanford University (Curto, Einav, Levin), School of Medicine, Stanford University (Bhattacharya), and NBER (Bhattacharya, Einav, Levin). Email: vcurto@stanford.edu, leinav@stanford.edu, jdlevin@stanford.edu, and jay@stanford.edu.

1 Introduction

Introducing managed competition into healthcare 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, perhaps equally important, is that there often is no clear way to draw comparisons between alternative insurance systems.

The recent experience of the US Medicare Advantage program provides a potential opportunity. 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 problems of 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 adjustment that makes payments a function of enrollee health status. The program has expanded since these reforms to cover almost thirty percent of US seniors (Figure 1).

In this paper we provide a model of insurer competition under the Medicare Advantage (MA) bidding rules, and present estimates that benchmark the program's welfare effects against traditional fee-for-service Medicare. The model we propose allows us to measure the competitive incentives facing insurers, obtain estimates of private insurer costs and the surplus that accrues to private plan enrollees, and 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. Some elements of our model — the determinants of plan bids, enrollee choice, and risk selection into private plans — have been analyzed in prior and concurrent work. In this sense, we view our main contribution as having an empirical framework that ties these pieces together in a way that facilitates econometric measurement and analysis of program design.

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 fee-for-service Medicare (MedPac, 2013). An open question is whether the extra expenditure is due to inefficiencies in the way private plans operate, or results from a failure of competition,

or instead reflects extra benefits for private plan enrollees. Any plausible answer needs to account for imperfect competition — as we discuss below, local insurance markets are highly concentrated — and provide a way to estimate insurer costs and enrollee benefits. It also needs to account for the fact that private plans tend to enroll relatively healthy beneficiaries (MedPac, 2013). We try to do this using a relatively standard imperfect competition model, which we adapt to the specific bidding rules, and to allow for differential risk selection.

We develop our analysis in several steps. The first is to estimate what it would cost to cover private plan enrollees under fee-for-service (FFS) Medicare. There are several complicating factors in this exercise. Private plan enrollment tends to be higher in urban areas with high fee-for-service costs. But within a geographic area, private plan enrollees have lower risk scores, are healthier conditional on risk score, and have risk scores that increase faster due to more thorough disease coding. We construct initial estimates of FFS costs for private plan enrollees by adjusting for geography and risk score. We then use post-enrollment mortality to adjust for residual differences in the populations. We estimate that within a typical market, private plan enrollees are about seven percent less costly than FFS enrollees. However, across the entire country, the average private plan enrollee is slightly more expensive than the average FFS enrollee because MA enrollees tend to be located in expensive urban markets.

Our second step is to provide a model of insurer competition that is suitable for welfare measurement. The model we propose captures the particular 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. We start with a baseline model in which risk adjustment successfully modifies private 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 primarily affects plan costs and not marginal bidding incentives.

The key parameters of our model are the price sensitivity of Medicare beneficiaries, the differentiation among private plans, the disutility enrollees incur from having a limited network of providers, and the costs incurred by private insurers. We estimate the sensitivity of beneficiaries to plan premiums using a difference-in-differences strategy that relies on varying changes over time in the bids of plans that share identical physician networks. We find that enrollee demand is only

moderately price-sensitive: a 10 dollar reduction in the monthly premium, or 10 dollar increase in the actuarial value of insurance benefits, increases a plan's enrollment by around 10 percent. Our demand model also generates estimates of enrollee surplus. On average, enrollees receive about 76 dollars per month of actuarial insurance benefits beyond standard Medicare, but this must be weighed against limits on provider access. Netting these out, we estimate that enrollee surplus averages around 49 dollars a month.

Our model also highlights the important competitive role played by fee-for-service Medicare, and the "benchmark" subsidy that the government provides for private plan enrollments. We estimate that, holding fixed a plan's bid, a 20 dollar decrease in the monthly Medicare subsidy has approximately the same effect on a plan's enrollment as a 20 dollar decrease in the bids of every other plan in the market. We also provide direct estimates of how changes in the benchmark rate are passed through into plan bids. Our findings coincide with those of Song, Landrum and Chernew (2012, 2013) and Cabral, Geruso and Mahoney (2014). We estimate that a 20 dollar decrease in the monthly benchmark rate leads plans to reduce their bids by around 10 dollars.

The most difficult numbers to estimate, due to data limitations, are private plan costs. We obtain estimates indirectly by calculating the optimal bid markups for insurers and making an assumption that plans set their bids optimally given competition. Using this approach, we calculate that the optimal bid mark-up is on the order of \$55-140 per enrollee-month. The implied cost estimates indicate that private plans have a cost advantage over fee-for-service Medicare in around half the country. Private plan enrollment is significantly higher in these markets. As a result, in the typical MA enrollment, we estimate that the plan can provide the same insurance benefits as fee-for-service Medicare, albeit with access to a limited set of providers, for around \$77 (or 12 percent) less per month. Our estimates are also consistent with a finding of Macurdy et al. (2013) and Landon et al. (2012) that private plan costs are not highly correlated with fee-for-service Medicare costs across local markets, suggesting that private and public insurance have different cost structures.

The final part of the paper combines our estimates to calculate welfare effects from Medicare Advantage. A useful way to think about this is in terms of the total (dollar) surplus created by a private plan enrollment, and then its division among the plan, the enrollee and taxpayers. As mentioned, MA plans achieved cost savings of 77 dollars in providing the standard FFS insurance benefit. Taxpayers captured none of this. Instead, they provided an additional subsidy of around 94 dollars per month relative to government costs under fee-for-service Medicare. The beneficiaries of

these 171 dollars (cost savings of 77 plus subsidy of 94) were insurers and consumers. We estimate that insurers captured 95 dollars per enrollee-month (not accounting for any fixed administrative costs). MA enrollees gained the rest, but also incurred a disutility of 27 dollars from having a limited network of providers, so their net surplus was lower, 50 dollars (per enrollee-month). In short, we estimate fairly substantial efficiency benefits, but an even larger net cost to taxpayers.

The empirical framework is also useful for thinking about changes in program design. The most obvious levers are the local benchmark rates against which plans compete, and the rebate formula that specifies how bid savings below the benchmark are divided between taxpayers and enrollees. Currently when a plan bids below its benchmark rate, the government retains 25% of the difference and mandates that the other 75% is passed to enrollees. One way to limit taxpayer costs would be to increase the government's share. In theory, this should be effective if demand were highly elastic. Given our estimates, it is not. Instead, benchmark reductions are a more effective instrument for reducing program costs. We estimate that a 50 dollar benchmark reduction would save taxpayers around \$8 billion a year without major reductions in plan enrollment. However, similar calculations highlight the difficulty of achieving budget neutrality for taxpayers without severely curtailing enrollment; we suggest that this trade-off could be improved with a more targeted approach to setting benchmark rates.

Our analysis contributes to an emerging literature assessing different elements of the reformed Medicare Advantage program. As mentioned above, Song, Landrum, and Chernew (2012, 2013) and Cabral, Geruso and Mahoney (2014) estimate the effect of local Medicare benchmark rates on plan bids, finding similar pass-through rates as we do, around 50%.¹ Duggan, Starc, and Vabson (2014) 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. 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 (Brown et al., 2014; Morrissey et al., 2013; Newhouse et al., 2012 and 2014), which is tangentially related to some of the results reported in the paper. Finally, the paper that is conceptually closest to ours is probably by Town and Liu (2003), who study the welfare impact of Medicare HMOs in the earlier and smaller incarnation of Medicare

¹Cabral, Geruso and Mahoney (2014) 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 provides a better explanation.

Advantage. Despite major changes in the program over the last decade, which we describe in the next section, their conclusion that the program generates net surplus but of smaller magnitude than plan profits foreshadows ours.² Miller (2014) is also a more recent paper that estimates a model of plan competition. He focuses on the potential for switching costs to lock enrollees into MA plans, leading to higher plan markups. We discuss this point in more detail in our conclusion.

The paper proceeds as follows. Section 2 provides institutional background on the Medicare Advantage program and the structure of private plans. Section 3 explains how we estimate FFS costs for private plan enrollees. Section 4 describes the theoretical model that we use to analyze plan competition. Sections 5 and 6 present evidence on plan bids and pass-through rates, and then on enrollment choices and plan demand. Section 7 examines enrollee surplus, plan costs and profits and the program’s overall welfare effects. Section 8 discusses alternative ways to promote plan competition. Section 9 concludes. We provide a variety of additional information and analyses in the Appendix.

2 The Medicare Advantage Program

Background Medicare Advantage (MA) allows Medicare beneficiaries to opt out of traditional fee-for-service Medicare and enroll in a private insurance plan. The program 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 will focus 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 a capitated monthly payment from CMS.³ Historically, these payment rates were set administratively. The result was an ongoing tension between the two program goals of expanding access and limiting costs (McGuire, Newhouse, and Sinaiko, 2011). Insurers 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.⁴

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 Town and Liu’s study 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 The Centers for Medicare and Medicaid Services (CMS) is the federal agency that manages the Medicare program.

⁴ Concerns about risk selection date to the very origins of the program (Eggers and Prihoda, 1982), and are

Several reforms in the last decade aimed to address these problems.⁵ Between 2003 and 2006, CMS phased in a risk scoring system to adjust plan payments based on enrollee health.⁶ 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.

Medicare Private Plans Medicare private plans must provide at least the same insurance benefits as standard Medicare (Parts A and B). Plans typically provide additional benefits as well, mainly in the form of more generous cost sharing or supplemental coverage (e.g. of dental, vision or prescription drugs). An important feature of the program is that these benefits must be “funded”: either through a supplemental premium paid by the insuree, or through a rebate that is paid by CMS. Plan rebates, as well as enrollee premiums, are determined through the bidding process.

The bidding process begins with CMS setting a benchmark capitation rate for each county. Each plan’s bid is assessed against its local benchmark.⁷ 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.⁸ If a plan’s bid b is below its benchmark rate B , there is no extra plan premium, and CMS makes a rebate payment to the plan of $0.75(B - b)$. The plan must use the rebate to fund additional benefits. The plan itself is paid its bid b (times the enrollee risk score) to fund standard insurance coverage.

The vast majority of plans submit bids below their benchmark rate, and rebates often are 50 dollars a month or more (see Section 5 below). The most common use of rebates is to provide extra cost-sharing or premium reductions for prescription drug coverage.⁹ The result is that plans

discussed in McGuire, Newhouse, and Sinaiko (2011). An illustration of selective participation is that in 2005, only around 67 percent of Medicare beneficiaries had access to an HMO or local PPO plan (MedPAC, 2009).

⁵The 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).

⁶The 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 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, 2014; Brown et al., 2014).

⁷Each 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 individual level by multiplying them by the individual county benchmark divided by the plan benchmark.

⁸All 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.

⁹Appendix Table A9 provides a breakdown of how plans use their rebates. In principle, rebates can be passed

generally offer more insurance than traditional Medicare. Traditional Medicare enrollees can purchase Medigap policies that cover some or all of their out-of-pocket health costs, but these policies often cost a few thousand dollars a year.¹⁰ In contrast, Medicare private plans bundle additional insurance, and features such as dental, vision, or prescription drug coverage, at relatively favorable premiums.¹¹

The trade-off is that private plans typically restrict access to healthcare providers. Around 85 percent of MA enrollees are in HMO or PPO plans with limited provider networks. 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 non-network providers. This type of plans proliferated in the mid-2000s when benchmark rates were very favorable, and in 2008 reached a 23 percent share of MA enrollees. However, their share dropped to 7 percent by 2011, and they are now relatively unimportant. Appendix Table A1 provides more detail on insurance options and how they compare to FFS Medicare.

Market Structure of Private Plans Private insurers offer Medicare Advantage plans by entering into contracts with CMS. Insurers may offer several plans as part of the same contract, with the same provider network but different benefits and bids. Larger insurers enter into multiple contracts in order to offer plans with different provider networks, for instance in different geographic areas. Individual markets tend to feature a large number of plans, but they are mostly small PFFS plans. Between 2006 and 2011, Medicare beneficiaries on average had access to three HMO plans and one PPO plan, in addition to fourteen PFFS plans.

The result is that local markets are highly concentrated despite the large total number of plans (see Table A2). In the majority of US counties, the three largest local insurers have more than 90 percent of MA enrollees. In three-quarters of US counties, the two largest insurers have more than a 75 percent market share. Concentration is somewhat lower in urban markets and in markets with high benchmark rates (Pizer, Frakt, and Feldman, 2012). For example, the mean insurer HHI in our study period was 0.477 for urban areas and 0.547 for rural areas.

directly to enrollees in cash through reductions in their Medicare Part B premium, but this is relatively uncommon. See also Stockley et al. (2014).

¹⁰In 2010, approximately 23 percent of beneficiaries were enrolled in Medigap insurance, paying monthly premiums that averaged \$183 per month (Huang et al., 2013).

¹¹One particular form of bundling is that 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.

There is less concentration at the national level. The two largest insurers, United Health Group and Humana, have 19 and 16 percent of national MA enrollees, respectively (see Table A3). These insurers operate in over 80 percent of local markets, but most insurers (there were around 130 in total in the 2006-2011 period) operate in far fewer. The market structure also looks different if one accounts for traditional Medicare. Around 35 percent of new enrollees in a given plan were enrolled the prior year in traditional Medicare (we exclude new Medicare beneficiaries from this calculation), and around 13 percent of beneficiaries who exit a plan move into traditional Medicare. The asymmetry reflects the expansion of private plans. In 2006-2011, around 2 percent of Medicare beneficiaries switched into MA each year, while only 0.6 percent switched in the other direction. Appendix Table A4 provides more detail on beneficiary transitions.

3 Constructing Cost Benchmarks under Risk Selection

We begin our analysis by constructing fee-for-service (FFS) cost estimates for the population of MA enrollees. As we have noted, this requires some care. Private plan enrollment varies by market and tends to be higher in areas with high fee-for-service costs. Private plan enrollees are also healthier in terms of measured risk scores, and potentially conditional on risk score. Finally, risk scores for continuing MA enrollees are not necessarily equivalent to those in traditional Medicare because of differences in disease coding. Our estimate of FFS costs attempts to adjust for each of these factors: geographic differences, risk score differences, unscored differences in health status, and variation in disease coding.

All of the analysis in this section, and the following sections, relies on Medicare administrative records. The data include risk score information, demographics and plan enrollment for all Medicare beneficiaries from 2006-2011, all MA plan payments from 2006-2011, and all Medicare fee-for-service claims from 2006-2010.¹² We restrict attention to aged Medicare beneficiaries who are not disabled, not dually eligible for Medicare and Medicaid, and not enrolled in employer-sponsored 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 dataset contains 153,858,811 individual-year observations on 35,504,869 unique individuals. We construct a parallel dataset on private plans: including all the private plans open to these beneficiaries, while excluding plans for dual eligibles and employer-sponsored plans. The plan dataset contains 12,311 observations at the plan-year level

¹²We have very limited information on healthcare utilization within private plans, as CMS began to collect detailed information on encounters only in 2013.

on 4,930 unique MA plans. Exact details of the data construction are reported in Appendix A and Appendix Table A5.

Differences between MA and FFS Enrollees Table 1 presents summary statistics on MA and FFS enrollees for the 2006 to 2011 period. MA enrollment over this period averaged around 20 percent.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 was 0.93, compared to 0.97 for FFS. Risk scores are constructed so that an individual with a score of 2 has twice the expected health cost of an individual with a score of 1. So based on risk score, a private plan enrollee had expected costs that were about 96 percent (that is, 0.93 divided by 0.97) those of a FFS beneficiary.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.¹⁵

The 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. An individual's codes may be incomplete and in any event are imperfect proxies for health status. In the case of MA enrollees, a variety of evidence suggests that they are somewhat healthier than FFS enrollees conditional on risk score. For instance, if one compares new MA enrollees to FFS beneficiaries with the same risk score, the "switchers" into MA have lower medical claims in the prior year than FFS beneficiaries with the same risk score. Among the FFS "stayers," prior year claims are significant predictors of current year claims conditional on risk score, suggesting that new MA enrollees are healthier than FFS 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, Chapter 7; GAO, 2013; Geruso and Layton, 2014). 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

¹³In our final analysis sample, MA penetration is 18.2 percent. However, as part of our sample construction we drop some MA enrollees for whom we do not drop the corresponding TM enrollees due to data limitations, as we explain in Appendix A. If we were to include those MA enrollees, then MA penetration would be approximately 19.9 percent.

¹⁴The risk score is calibrated to have a mean of 1 on a subsample of several million TM enrollees. We have dropped two high-risk groups – Medicaid recipients and individuals with Disability insurance – so risk scores in our sample are slightly lower than in the entire population.

¹⁵There 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.

¹⁶In our sample, the lagged claims of "switchers" into MA were about 2.3 percent 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. (2013).

more persistent.¹⁷ In 2010, CMS attempted to correct for differential coding, and deflated all MA risk scores by 3.41 percent. It applied the same adjustment in 2011, but there was no such adjustment in the years prior to 2010.

Our strategy to correct for unscored health differences is to use an ex post health measure. We focus on differences in population mortality, which (absent data on encounters) is the most reliable health outcome we observe for MA enrollees. Figure 3 plots the one-year mortality rate for MA and FFS enrollees as a function of risk score. New private plan enrollees have lower mortality conditional on risk score, consistent with self-selection. Mortality rates among all MA enrollees also are lower than for FFS enrollees, with the difference being more pronounced in the 2006-2009 period. Of course some of this difference might be explained by private plans offering superior care. However, we think this is a relatively unlikely explanation for the differential mortality. MA enrollment increased during our study period, but there is no evidence that county mortality rates fell in areas where MA enrollment increased most rapidly, which one would expect if private plan coverage reduced mortality. We present the full details of this analysis in Appendix Table A6.¹⁸

Measuring Local FFS Costs We build our cost benchmarks from measures of local FFS costs for each US county in each year. 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 mid-year enrollment, and the fact that the average risk score of FFS enrollees in our sample is not identically equal to 1. Appendix B provides more detail on the differences between our measures and the ones reported by MedPac.

Let k index county-years, and let FFS_k denote the set of individuals enrolled in FFS in county-year k . For a given individual i , let m_i denote the number of months that i was enrolled in FFS during the year, and x_i denote i 's average monthly FFS claims during those months. We define

¹⁷To 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 traditional Medicare. It is 58% vs. 45% for Polyneuropathy, and 49% vs. 33% for Angina Pectoris/Old Myocardial Infarction.

¹⁸Gowrisankaran, Town and Barrette (2011) do find some evidence that enrollment in private plans without prescription drug coverage increased mortality relative to traditional Medicare. Their study looks at data from 1993 to 2000, when Medicare Advantage was smaller and operated under different years, and Medicare Part D did not exist, so the differences between their result and those in Table A6 may not be very surprising. Note that other *ex post* measures of health might be more sensitive to differences in insurance status. For example, Table 1 shows that inpatient days are (much) lower among MA enrollees. This could be due to differences in population health, but also to differences in the way private plans handle admissions and hospital stays. In the case of inpatient stays, there are also reporting issues that bias down the private plan average in Table 1 (see Landon et al., 2012).

two measures of local FFS spending in county-year k :

$$\alpha_k = \frac{\sum_{i \in FFS_k} x_i m_i}{\sum_{i \in FFS_k} m_i} \quad \text{and} \quad a_k = \frac{\sum_{i \in FFS_k} x_i m_i}{\sum_{i \in FFS_k} r_i m_i}. \quad (1)$$

The first measure is the local FFS cost per enrollee-month. The second is the local FFS cost per enrollee risk-month. Note that $\alpha_k = a_k r_k$, where $r_k = \sum_{i \in FFS_k} r_i m_i / \sum_{i \in FFS_k} m_i$ is the (month-weighted) average enrollee risk in county-year k . In our sample, r_k is generally lower than 1, so that costs per enrollee-month are somewhat lower than costs per enrollee risk-month.

Estimating FFS Costs for MA Enrollees Next we construct estimates of FFS costs for MA enrollees. We do this in two steps, first imputing FFS costs based on each enrollee’s location and risk score, and then adjusting for health differences conditional on risk score. For the first step, we assign each individual i in county-year k an expected monthly FFS cost of $a_k r_i$, where r_i is the individual’s risk score and a_k is the local FFS cost per enrollee risk-month.¹⁹ For the second step, we rescale MA risk scores in a way that aligns conditional FFS and MA mortality rates.

Our risk score rescaling can be motivated as follows. Let $\mu_{FFS}(r)$ and $\mu_{MA}(r)$ denote the one-year mortality rates of FFS and MA enrollees. Assuming both rates are strictly increasing in r , we can define $\Lambda(r)$ to be an increasing function such that $\mu_{FFS}(\Lambda(r)) = \mu_{MA}(r)$. The idea is that an MA enrollee with observed risk score r should be compared to a FFS enrollee with risk score $\Lambda(r)$. For the purposes of cost benchmarking, we can assign FFS costs to each MA enrollee with risk score r by adjusting his or her risk score to $\Lambda(r)$.

To make this operational, we assume that there is a single scaling factor λ that applies in each year, so that in year t , $\Lambda_t(r) = \lambda_t r$. To estimate each λ_t , we construct a geographically balanced sample of MA and FFS enrollees for year t by randomly dropping FFS enrollees in each county until their number equals the number of MA enrollees (or the reverse if MA enrollees are the majority in a county, which is not common). We use the balanced sample to obtain nonparametric estimates of the one-year mortality rates: $\hat{\mu}_{FFS}(r)$ and $\hat{\mu}_{MA}(r)$. Then, for each MA enrollee i , we compute λ_i to satisfy $\hat{\mu}_{FFS}(\lambda_i r_i) = \hat{\mu}_{MA}(r_i)$. Finally, we average over MA enrollees in that year to obtain an estimate of λ_t .²⁰

¹⁹This construction is motivated by an underlying model of monthly FFS 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 .

²⁰In these calculations we trim a small number of observations (about 3,000 out of 55 million) with risk scores below 0.25 or above 10 because our estimates of mortality are very noisy at the extremes of the risk score distribution.

For the period 2006-2011, we estimate annual scaling factors of 1.032, 0.985, 0.976, 0.942, 0.995, and 0.990. The average for the overall period is 0.984. We adjust the predicted FFS costs of each MA enrollee in our sample by the relevant scaling factor. So an MA enrollee in year t and county-year k , with risk score r_i , has predicted FFS costs of $x_i = \lambda_{ta_k} r_i$. Note that in principle, the rescaled prediction should account both for self-selection of healthier individuals into MA, and for differences in disease coding. Of course, how well it does this depends on how well mortality differences proxy for differences in health status (and in particular expected medical claims). The adjustment is almost surely imperfect, but our view is that it is a reasonable approach and certainly preferable to making no adjustment.

Comparing Taxpayer Costs for FFS and MA Table 2 summarizes our estimates of FFS monthly costs for the traditional Medicare and private plan enrollees. The first column reports expected monthly FFS costs per enrollee. The second column reports expected monthly FFS costs per risk unit. The third and fourth columns split the sample into urban and non-urban regions, and the final columns separate the data by time period. Costs are increasing over time, and they are significantly higher in urban areas. As noted earlier, MA enrollment is higher in urban areas, and in later years. These compositional differences make the MA population about eleven percent more expensive in the first column. At the same time, MA enrollees within a given county-year have expected costs that are about five percent lower. These compositional differences are reflected in the first two columns of the table.

Our benchmark cost numbers allow a first pass at cost accounting for the MA program. The second panel of Table 2 shows that during our sample period, private plan enrollees had average expected monthly FFS costs of \$675. Plan bids were slightly below this (\$670), indicating that on average plans were willing to provide coverage for around what it would have cost to cover enrollees under traditional Medicare. Rebate payments, however, averaged \$76 per enrollee, so that on average CMS paid out \$746 per enrollee-month, about 11 percent more than we estimate CMS would have paid in FFS claims. The difference is larger in urban areas (13 percent). Total MA payments actually were almost identical to predicted FFS costs in rural areas. The taxpayer “subsidy” also was shrinking over time, from over 17 percent in 2006-07 to just under 3 percent in 2010.

These taxpayer accounting numbers, however, provide very partial information about the efficacy of the program. In particular, they provide no information as to whether rebate payments

should be viewed as a beneficial transfer to Medicare beneficiaries, or as necessary compensation to offset having a limited provider network. They also provide no information as to whether plan bids closely reflect plan costs, or whether insurers are making substantial profits. Put another way, the numbers are consistent with an optimistic view that the program generates considerable surplus for plans and enrollees, and a skeptical view that the program could not exist without a large taxpayer subsidy.

To say more, we require several pieces of information. We need an estimate of enrollment choices and the surplus realized by private plan enrollees. We also need an estimate of private plan costs and how they compare to bids. Finally, we would like to know how program changes might affect bids and enrollment, so that we can understand how policy decisions affect consumer, insurer and taxpayer surplus. The remainder of the paper pursues these goals, starting with a theoretical model of beneficiary and plan decisions.

4 Bidding Competition

This section presents our model of insurer competition under the Medicare Advantage bidding rules. We begin by describing the rules. We then present a baseline model, and explain how we will connect the theory to the data.

The Bidding Rules As described earlier, plans are reimbursed based on their bids and the local benchmark rates set by CMS. The benchmark rates are set each year at the county level. We therefore will think of each county-year as a separate market.²¹ 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.

First consider a plan that submits a bid b greater than the benchmark rate B . In this case, each enrollee must pay the plan a premium $b - B$ in addition to paying the standard Medicare Part B premium. CMS pays the plan $rb - (b - B)$, where r is the enrollee's risk score. In total, the plan receives rb . Now consider a plan that submits a bid b below the benchmark rate. In this case, the enrollee pays no plan premium.²² For an enrollee with risk r , CMS pays the plan rb . It also gives

²¹A complication that we noted earlier is that insurers often make a given plan available in several counties, and submit a single plan bid that is adjusted to the county level based on how the county benchmark compares to the average benchmark rate in the plan's service area. Despite this, we will 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.

²²This is not always right. As mentioned earlier, plans sometimes charge an additional supplemental premium to

the plan a rebate of $0.75(B - b)$, which the plan must use to provide extra benefits.

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 $b < B$. 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. The next section expands on this to model bidding incentives.

A Baseline Model We consider a single market with beneficiaries who vary in their risk type r . There are J plans, indexed by $j = 1, \dots, J$. 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 . We are interested in how plan bids translate into enrollment decisions, and how plans optimally set their bids in market equilibrium.

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 $D_{jr}(p_1, \dots, p_J)$. A natural assumption is that D_{jr} will be 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

$$Q_j(p_1, \dots, p_J) = \int r D_{jr}(p_1, \dots, p_J) dG(r). \quad (2)$$

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

Plan profits. Next we consider plan profits. A plan receives rb_j for enrolling an individual with risk r , and incurs cost rc_j . Therefore plan j 's profit given a set of plan bids p_1, \dots, p_J is

$$\pi_j(p_1, \dots, p_J) = Q_j(p_1, \dots, p_J) \cdot (p_j + B - c_j). \quad (3)$$

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 fund additional benefits. This is not common, and when it happens the amounts of these supplemental premiums are relatively low.

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.²³ For an insurer that offers a set of plans \mathcal{J} , the optimal bid for plan $j \in \mathcal{J}$ satisfies

$$0 = \sum_{l \in \mathcal{J}} (p_l + B - c_l) \frac{-\partial Q_l}{\partial p_j} + Q_j(p_1, \dots, p_J), \quad (4)$$

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

$$p_j = c_j - B + \left(\frac{-\partial \ln Q_j}{\partial p_j} \right)^{-1}. \quad (5)$$

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

Connecting the Model to Evidence We now consider how the model can be connected to evidence on bidding and enrollment decisions.

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 re-write condition (5) as $b_j = c_j + (\partial \ln Q_j / \partial b_j)^{-1}$. So an optimal mark-up will be high when plan demand is not very price sensitive, and low if it is sensitive. We provide estimates of enrollment sensitivity in Section 6.

Plan enrollment and bidding incentives also depend on both competing plan bids and the benchmark rate, which determines the relative price for traditional Medicare. A key empirical question is whether competition comes primarily from other private plans — so that for instance, additional entry may reduce bids — or from FFS traditional Medicare. A rough way to quantify this is to imagine that all competing plans bid $b_k = \beta$ and ask how changes in β and B , respectively, affect $\ln Q_j$ and $\partial \ln Q_j / \partial b_j$. We address this empirically below.

²³Complete 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 will exist are laid out in Caplin and Nalebuff (1991).

Effect of the benchmark rates. The benchmark rate plays a key role in that it largely determines the program cost to taxpayers. In turn, the “pass-through” rate of benchmark increases into plan bids determines whether marginal taxpayer dollars go to insurers or enrollees. This pass-through rate can be related to standard cost pass-through. To see this, suppose plan costs are given by $c_j = c + \gamma_j$, where c is common to all plans in a market. Then the equilibrium excess bids p_1, \dots, p_J will be a function of $B - c$, and if $p = b - B$ is a representative equilibrium bid, we will have $dp/dB = dp/d(-c)$. Therefore,

$$\frac{db}{dB} = 1 + \frac{dp}{dB} = 1 - \frac{dp}{dc} = 1 - \frac{db}{dc}. \quad (6)$$

We highlight two points. First, if the equilibrium theory is correct, and we have an accurate measure of common costs, the pass-through rate of benchmark increases should be one minus the pass-through rate of uniform cost increases. We return to this in the next section. Second, in a perfectly competitive market where $b = c$, changes in the benchmark would have no effect on plan bids. This is why Song et al. (2012) interpret their finding of positive benchmark pass-through 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 we will move from estimating bid regressions in the next section to estimating enrollment demand in Section 6.

Incentives for risk selection. The baseline model assumes that variation in a plan’s risk composition is fully compensated. However, we have seen that 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 is straightforward. It is “as if” the cost of MA plan j per enrollee risk unit is λc_j . So equation (5) becomes

$$p_j = \lambda c_j - B + \left(\frac{-\partial \ln Q_j}{\partial p_j} \right)^{-1}. \quad (7)$$

Things get more complicated if $\Lambda_j(\cdot)$ depends on p_j — either because p_j affects observed risk composition and $\Lambda_j(\cdot)$ is not proportional to r so for instance unobserved selection is more important for high or low risk enrollees, or because p_j affects risk composition conditional on r . In this

case, equation (5) becomes

$$p_j = \lambda_j c_j - B + \left(\frac{-\partial \ln Q_j}{\partial p_j} \right)^{-1} \left[1 - \frac{\partial \lambda_j}{\partial p_j} c_j \right], \quad (8)$$

where $\lambda_j = \int \Lambda_j(r, p_j, p_{-j}) D_{jr} dG(r) / \int r D_{jr} dG(r)$ is the (average) risk scaling factor for plan j . The last term appears because changes in p_j may affect the cost advantage plan j has relative to measured risk scores. For example, if an increase in p_j shifts the plan's enrollment toward lower r 's and plan j is relatively over-compensated for low-risk enrollees, there will be an extra incentive to bid high. Below, we will argue that empirically, changes in plan bids do not have much effect on a plan's risk composition, so that the final term in equation (8) is close to zero, and the former case is an appropriate description of plan incentives.

5 Determinants of Plan Bids

In this section we study the empirical determinants of plan bids. To provide some context, Figure 5 plots the distribution of plan bids relative to plan benchmarks. Over 90 percent of plan bids are below the relevant benchmark, and result in rebates and additional enrollee benefits. Here we consider how closely plan bids track benchmark changes, and how they relate to local FFS costs. As we just noted, these types of pass-through estimates can provide evidence for imperfect competition, although some care is needed in interpreting them. We turn in the following sections to our estimates of plan demand, consumer surplus and plan costs.

Plan Bids, Benchmarks and FFS Costs Our model implies that the bid of a given plan j in county-year k will depend on the plan's cost, its benchmark rate, and the characteristics of the plan's residual demand. Letting x_{jk} denote a vector of residual demand characteristics, we can write:

$$b_{jk} = f(c_{jk}, B_k, x_{jk}). \quad (9)$$

We do not have a direct measure of private plan costs, but we have described a way to construct predicted FFS costs for each private plan enrollee. Let c_{jk}^{FFS} denote the predicted FFS costs of plan j 's enrollees in county-year k . We also observe the benchmark rate B_k for each county-year k . However, this leads to the following institutional wrinkle. While benchmarks are set at the county-year level, plans are defined to cover a service area that may encompass several counties. A

plan submits a single bid for its service area. If the service area encompasses several counties, CMS adjusts the plan payment to reflect not just enrollee risk scores, but the FFS costs of the enrollee’s county relative to the service area average.

We consider two ways of dealing with the geographic mismatch. The first is to define a benchmark for each plan, where we weight each county benchmark by its share of plan enrollees (this is the approach in Figure 5). We then compute predicted FFS costs across all plan enrollees, and estimate equation (9) at the plan-year level. The second is to define a separate plan bid for each county in the plan’s service area. We do this by scaling the plan bid to the county level using the CMS Intra-Service Area Rate (ISAR) adjustment (CMS, 2013). We then compute predicted FFS costs and estimate equation (9) at the plan-county-year level. As it turns out, the empirical results are similar across the two approaches.

We consider the following econometric specification, written here at the plan-county-year level:

$$b_{jk} = \beta_0 + \beta_1 c_{jk}^{FFS} + \beta_2 B_k + \eta_{year(k)} + \eta_{county(k)} + \varepsilon_{jk}. \quad (10)$$

The specification is similar to the one reported in Song, Landrum, and Chernew (2012, 2013). Including county characteristics rather than fixed effects, as they do, or adding measures of private plan market structure, has little effect on our estimates of β_1 and β_2 , which define how predicted FFS costs and benchmark rates affect plan bids. The plan-year specification is the same, but includes fixed effects for each year and each insurer contract.

Identifying the Pass-Through Rate of Benchmark Changes The benchmark rate is a key parameter because it defines the subsidy paid by taxpayers when a beneficiary enrolls in a private plan.²⁴ If the benchmark increases but bids remain constant, enrollee benefits expand. If bids increase one-for-one with the benchmark, insurers absorb the higher subsidy level. This makes the pass-through of the benchmark rate quite important. Here we describe how benchmarks are set, and the resulting variation we use to identify the effect on plan bids.

The current benchmark rates have evolved from the historical capitation rates paid to plans. Originally these were targeted to be 95 percent of the amount it would cost to cover a standard beneficiary in a given county. In the late 1990s, Congress introduced a payment floor and subsequently a separate and higher floor for urban counties (those that belong to an MSA with a population of 250,000 or more). Benchmarks were raised further in the early 2000s to encourage more plan

²⁴More specifically, taxpayer costs are rb for bids above B , and $rb + 0.75(B - b)$ for bids below B .

offerings.

During our study period of 2006-2011 a statutory formula governed annual benchmark updates. Each year, all county benchmarks were adjusted upward by the greater of 2% and the average national increase in FFS costs. In addition, CMS “rebases” each county benchmark at least once every three years. In a rebasing year, CMS calculates per-capita FFS spending in the county based on a five-year moving average. If the FFS average exceeds the benchmark rate that would be set otherwise, it becomes the new benchmark. The years 2005, 2007, and 2009 are the relevant rebasing years for our sample. In 2006, 2008, and 2010, the benchmark rates were updated based on the national growth in FFS costs. In 2011, after the passage of the Affordable Care Act, benchmarks were frozen at their 2010 levels. Figure 4 presents the distribution of county benchmarks in the first and last years of our data (2006 and 2011). A large share of the benchmark rates are at one of the two floors: 46.6 percent and 44.5 percent of the rural counties have rates at the rural floor in 2006 and 2011, while 19.8 percent and 20.7 percent of the urban counties have rates that are at the urban floor.

Our empirical specification (10) includes county and year fixed effects, so we identify the effect of benchmark rates on bids using three sources of residual variation. The first is that over time, some county benchmark rates move in and out of their respective floors. A second source is the variance of historical county-specific FFS costs from current county-specific FFS costs. The former are used in rebasing years to set county benchmarks, whereas the latter enter our specification in our construction of plan FFS costs. Finally, a third source of variation arises from the fact that benchmark rates are updated in percentage terms so that high-benchmark counties have higher dollar increases in their benchmarks. Our year fixed effects only control for average dollar increases. Our underlying assumption is that these three types of variation in benchmarks are not associated with residual variation in individual plan costs or demand. We view this as quite plausible due to the formulaic nature of the benchmark setting. The last source of variation is somewhat mechanical, so to check that it is not crucial for our results, we also report log-log specifications in Table A8.

Estimates of Plan Bid Regressions We report our main regressions in Table 3. For each specification, we also report results with a more limited set of fixed effects, so that it is possible to compare our primary specification with estimates that also make use of cross-sectional variation in benchmarks and FFS costs. The estimates are very similar across all specifications, and also very similar in the log-log specification of Table A8.

We first consider the relationship between benchmark rates and plan bids. Our estimates imply a pass-through rate in the range of 40-60 percent. This finding is consistent not just across our specifications, but also with the estimates in Song, Landrum and Chernew (2012, 2013) and Cabral, Geruso and Mahoney (2014), who consider similar although not quite identical specifications. The results do differ notably from those of Duggan, Starc, and Vabson (2014), who estimate extremely high pass-through rates, almost 100 percent. This last paper uses a different identification strategy, focusing on a small set of counties with populations that are right around the urban floor threshold.

What is the implication of this sort of pass-through rate? First it implies that changes in the benchmark rates, such as those mandated under the Affordable Care Act, will have roughly equal effects on plan profits and enrollee surplus, holding enrollment patterns fixed. Of course, enrollment patterns are not fixed, and we will use our demand estimates in the next section to provide further evidence on the consequences of benchmark changes. The pass-through estimates also provide fairly clear evidence that plans do not base bids solely on their costs, as emphasized by Song, Landrum and Chernew (2012, 2013). As noted above, however, one needs stronger assumptions about demand — which we will add in the next section — to connect pass-through rates to market power under the assumption of imperfect competition.

Table 3 also contains a second and perhaps more surprising result. The estimated coefficient on fee-for-service costs is very small, in fact nearly zero in most specifications. One way to think about this is in terms of the model of Section 4. If FFS costs are a good proxy for MA costs that are common across plans, a zero pass-through rate is somewhat surprising, and clearly at odds with the theoretical prediction that (common) cost pass-through should be equal to one minus the benchmark pass-through.

What can explain this? One possibility is that realized fee-for-service costs in a plan's service area are an accurate but very noisy measure of the expected costs that plans use in forming their bids. To explore this hypothesis, we consider additional specifications that attempt to reduce measurement error — focusing on larger counties where idiosyncratic cost shocks are more likely to average out, or averaging fee-for-service costs over multiple years. However, our basic finding remains stable: fee-for-service costs bear little relationship to plan bids conditional on benchmark rates.

This leads us toward an alternative hypothesis that the cost structure of MA plans is in fact quite different from that of fee-for-service Medicare. Private plans that are HMOs or PPOs must negotiate prices with providers, so their unit costs need not be identical to those of CMS, and they

may apply more rigorous utilization management than fee-for-service Medicare. Indeed, several recent papers (Aizcorbe et al., 2012; Baker, Bundorf, and Kessler, 2010; MaCurdy et al., 2013) have noted that Medicare FFS costs and the costs of employer-provided private health insurance are not strongly correlated across regions. The structure of Medicare private plans is probably more similar to that of employer-sponsored health plans for non-seniors than it is to that of FFS Medicare.

To explore this hypothesis, Table 4 replicates the last column of Table 3, only with separate estimates for the three types of MA plans: HMO plans, local PPO plans, and private FFS plans. As noted earlier, the private FFS plans provide similar insurance to traditional Medicare, and pay the FFS rates to providers. The results for these plans are strikingly different. While the coefficient on FFS costs is essentially zero for HMO plans, the estimates for private FFS plans are closer to 0.5 and match up with the prediction of the theoretical model that db/dB should be equal to $1 - db/dc$. The estimates for local PPO plans are in between the HMO and PFFS estimates.

In light of this, we conclude that Medicare FFS costs are not an ideal proxy for the costs of most private plans, with the exception of private FFS plans. For this reason, we will rely on estimates of plan demand, coupled with an assumption of equilibrium bidding, to back out alternative estimates of MA plan costs.

6 Estimates of Plan Demand

We next propose an empirical specification of the demand for private plans. We then use the demand specification to estimate the price sensitivity of private plan enrollees, which in turn allows us to obtain estimates of enrollee surplus and plan costs and profitability. We start this section with the demand model, then discuss the variation in plan bids that we use to identify price sensitivity and finally present the estimates.

Demand Specification We adopt a nested logit specification for plan demand, adapted slightly to capture the fact that plan revenue and costs depend on risk-weighted demand. 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 Medicare) is the only member of group $g = 0$, and all MA plans belong to the other group $g = 1$.²⁵

²⁵Our model somewhat simplifies the actual choices available because a beneficiary who chooses traditional Medicare also can choose to purchase a supplemental Medigap policy (which is not allowed for MA enrollees). In addition,

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

$$u_{ijk} = \delta_{jk} + \zeta_{ig} + (1 - \sigma)\epsilon_{ijk} \quad (11)$$

where

$$\delta_{jk} = x'_{jk}\beta - \alpha(p_{jk}^+ + 0.75p_{jk}^-) + \eta_k + \xi_{jk}. \quad (12)$$

We use x_{jk} to denote an M -dimensional vector of observable plan characteristics. These include 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. We use p_{jk} to denote the excess bid. We also make the assumption (which we test and do not reject; see Appendix Table A10) that beneficiaries respond to dollar changes in MA premiums in the same way that they respond to dollar changes in the value of the plan's benefits package. Specifically, let $p_{jk}^+ = b_{jk} - B_k$ when $b_{jk} \geq B_k$ and zero otherwise and let $p_{jk}^- = b_{jk} - B_k$ when $b_{jk} < B_k$ and zero otherwise. That is, we assume that once one accounts for CMS claiming 25 percent of under-bids, the effect of a bid change on demand is the same, whether bids are above or below the benchmark.

To complete our specification, let ξ_j denote the mean value of the unobserved (by the econometrician) plan characteristics, and let $\Delta\xi_{jk}$ denote a market-specific deviation from this mean (so $\xi_{jk} = \xi_j + \Delta\xi_{jk}$). We use η_k to denote the mean utility for MA plans in market k relative to the outside good. 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 $\zeta_{ig} + (1 - \sigma)\epsilon_{ijk}$ follows a generalized extreme value distribution. As shown in Berry (1994), this yields the nested logit specification

$$\ln(s_{jk}) - \ln(s_{0k}) = \delta_{jk} + \sigma \ln(\bar{s}_{jk}) \quad (13)$$

where \bar{s}_{jk} is the market share of plan j as a fraction of the overall share of MA 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 one to each plan

FFS beneficiaries can choose to enroll in a stand-alone Part D plan. We lump these possibilities together as the "outside good." Note that in our specifications with county and year fixed effects, the mean utility of MA relative to the outside good may vary across markets.

enrollee, we assign a weight that is proportional to the product of the enrollee’s risk score times the number of months (within the year) over which he 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_{jk} is the risk-month-weighted share of each plan.²⁶ We let δ_{jk} be a function of year and contract (or contract-county) fixed effects, as well as the full set of plan quality rating dummy variables, and the excess bid.

A standard issue in estimating equation (13) is that \bar{s}_{jk} is endogenous because unobserved changes in plan demand affect both its market share and its market share among MA plans. We use three alternative instruments for \bar{s}_{jk} : 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 dummies for which other contracts are offered in the same market. None of these is perfect, but the results across choices of instruments are very stable, and the price effects are quite consistent over an even wider set of specifications, so we use the contract dummy IV regression as our preferred specification.

Identifying Beneficiary Price Sensitivity 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 plan quality. Arguably the most important component of plan quality is the network of providers that the plan covers. Even if we could perfectly observe the network, encoding it in a usable set of control variables would be challenging and imperfect. Our solution therefore is to rely on the panel structure of the data. We use 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, one benefit is that it utilizes the structure of the program under which plans within a contract share the same provider network, and also lets us 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 market

²⁶Note 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.

share of plan j in county-year k as:

$$\ln(s_{jk}) - \ln(s_{0k}) = x'_{jk}\beta - \alpha(p_{jk}^+ + 0.75p_{jk}^-) + \sigma \ln(\bar{s}_{jk}) + \eta_{year(k)} + \eta_{contract(j)} + \xi_{jk}, \quad (14)$$

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 therefore 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 the exact benefits, i.e. the monthly premium, whether the plan includes Part D benefits. Controlling for benefits that are funded with supplemental fees (such as a bundled Part D offering), we use variation across plans within a contract in the amount of benefits generated by different rebates to identify price sensitivity. Second, plan bids change from year-to-year, and we use the idiosyncratic variation in bid changes relative to national changes to identify price sensitivity. We also report results for a variant of the demand model with contract-county fixed effects, rather than contract fixed effects. The idea is that while plans within the contract have the same provider network, the value of the network might vary across counties within a service area. The estimates are not very different, however.

Demand Estimates We first report results from standard logit specifications in Panel A of Table 5, and then from the nested logit specifications in Panel B. We use the estimates in Panel B, column 3 in subsequent calculations.

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 most tightly specified logit specification (column (5)), which uses bid variation within a contract-county for identification, implies that a bid increase of \$10 decreases the plan's risk-month-weighted market share by around 9 percent. Our preferred specification is in column (4), which uses contract fixed effects to control for variation in provider networks, but captures variation in the coverage across plans within contracts using variables that describe the benefits associated with each plan. Using this variation, the effect is about 50% greater. A (below the benchmark) bid increase of \$10 decreases the plan's risk-month-weighted market share by around 14 percent.

Using this specification, Panel B of Table 5 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, around 0.32, 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 plan, the other competing MA plans in its market (which have around 20% market share) exert just about the same competitive effect as FFS 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 FFS 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 fixed all plan bids.

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 3 percent more enrollees than a three-star plan, equivalent to a bid difference of about 3 dollars per month.

Incentives for Risk Selection Prior to the introduction of risk scoring, insurers had a clear incentive to enroll healthier beneficiaries and there is considerable evidence that they did.²⁷ We observed in Section 3 that MA enrollees today continue to have lower risk scores and lower mortality conditional on risk score. Of course if bidding is sufficiently competitive, cost savings from favorable risk selection will translate to lower bids and be competed away. From the perspective of bidding competition, a crucial issue is whether bidding incentives on the margin are affected by risk selection, in particular whether firms might avoid low bids to maintain more favorable enrollee characteristics. In the context of our empirical approach, we also propose to infer each plan’s costs from observed bids. If a higher bid leads to significantly healthier enrollees, we would need to account for this in our estimation.

We therefore want to assess whether plans have an incentive to adjust their bids in order to change the risk composition of their enrollees. To do this, we consider the following regression:

$$\bar{r}_{jm} = x'_{jk}\beta + \alpha(p_{jk}^+ + 0.75p_{jk}^-) + \mu_j + \nu_k + \varepsilon_{jk} \quad (15)$$

where \bar{r}_{jm} is the average enrollee risk for plan j in county-year k , and p_{jk} is the plan’s excess bid, as before. In Panel A of Appendix Table A11 we use the same set of specifications as in Panel A of Table 5, and Panel B of Appendix Table A11 repeats the exercise, but replacing plan average risk

²⁷There is evidence on this dating back at least to Eggers (1980) and Eggers and Prihoda (1982).

with plan mortality rate as the dependent variable, and adding plan average risk as a regressor. This allows us to look separately at selection on observables and selection on unobservables.

The results suggest that higher plan bids have a very small effect on the average risk score of a plan’s enrollees. A \$10 decrease in bid is associated, in most specifications, with a decrease in the average enrollee risk score that generally is less than 0.5%.²⁸ The second panel of Appendix Table A11 provides analogous results for the relationship between plan bids and plan mortality conditional on the plan’s average risk score. Again, a lower plan bid is associated with somewhat healthier enrollees in the sense of lower mortality. The effect is slightly bigger; a \$10 bid decrease is associated with decreased mortality rate of 0.0005 per risk unit. We will not account for this in our subsequent estimates of plan costs; taking it into account would lead to somewhat higher cost estimates following the logic of Section 3.

7 Welfare Effects of Medicare Advantage

In this section we use our demand estimates to estimate enrollee surplus from private plans, and to back out implied costs of each MA plan. We then compare these costs with the predicted FFS costs associated with the same beneficiaries, and provide estimates of the overall “gains from trade” from having MA plans, and 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.

Estimates of Enrollee Surplus 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

$$CS_k = \frac{1}{\alpha} \log \left[1 + \exp \left((1 - \sigma) \log \left[\sum_{j=1}^{J_k} \exp \left(\frac{\delta_{jk}}{1 - \sigma} \right) \right] \right) \right]. \quad (16)$$

The parameter α 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 15 to 80 dollars per Medicare enrollee-month. Of course, this surplus is concentrated among beneficiaries who actually

²⁸Note that the direction is contrary to the usual adverse selection effect. Plans that offer more generous insurance have more favorable risk selection.

enroll in MA plans. Attributing it to this narrower group implies an enrollee surplus of 103 dollars per MA enrollee-month.

This calculation suffers from some well-known limitations of the logit model, because under its assumptions there are always consumers who place very high value on any alternative (due to the idiosyncratic ϵ_{ijk}), even alternatives with very small market shares. To avoid attributing gains simply from giving consumers more options, it seems appropriate to focus attention only on the surplus that is generated through the rebate and difference out surplus that the model says would be created if private plans offered identical financial benefits to FFS Medicare.

To do this, we recompute the consumer surplus for each market, setting the component of utility that comes from the rebate equal to zero by replacing the variable $(p_{jm}^+ + 0.75p_{jm}^-)$ with p_{jm}^+ . We then subtract this non-rebate surplus from our first estimate to obtain the surplus benefit due to the rebate payment. When we do this, we estimate that enrollees receive on average 49 dollars per month in rebate surplus from their private plans. This is about two-thirds of the average actuarial value of the rebate, which is 76 dollars per enrollee-month, implying a disutility of 27 dollars a month from other features of the private plan such as the limited provider network.

Estimates of Plan Costs and Profits We next combine our demand estimates with an assumption of equilibrium bidding to infer implied plan costs, a common approach in industry studies. We start with the condition for optimal bidding in equation (5), and modify it to reflect the fact that insurers may offer multiple plans in the same market. Specifically, 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 lm th entry is equal to one if plans l and m are owned by the same insurer, and zero otherwise. With this (standard) notation, the plan costs are given by the solution to the first-order conditions for optimal bidding:

$$c = b + (\Omega \cdot D_bQ)^{-1} Q. \tag{17}$$

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. We compute costs for each plan in each county-year. In each case, the estimate is the cost of providing standard Medicare benefits to a standard Medicare Advantage enrollee, i.e. a beneficiary with risk score λ_t . To compute plan costs on a per-enrollee basis, we simply multiply by the average risk score of the plan's enrollees.

We are interested in how the estimated private plan costs compare to equivalent FFS costs. We can make this comparison at the county-year level. For each county-year, we compute the enrollment-weighted average of private plan costs. In Figure 6, we plot these costs against our estimate of the FFS costs of covering the same set of MA enrollees in that particular county-year. We compute the latter using the procedure described in Section 3. The result is that each point in Figure 6 represents a county-year, with the horizontal axis representing the predicted average FFS cost of MA enrollees in the market, and the vertical axis representing the average implied private plan costs of the same MA enrollees. The dotted line in Figure 6 is the 45-degree line, so that points above it represent markets in which MA cost is estimated to be more expensive than what it would cost to cover the same individuals in FFS, while points below the line are where MA appears to have the potential to generate cost savings.

In Figure 6, MA costs are cheaper than FFS in 53% of the markets. These markets cover about 70 percent of the population of beneficiaries who are enrolled in MA (and 61 percent of the overall Medicare population). The difference is natural; in markets where MA plans have a cost advantage, their bids are lower and hence enrollment is higher. Note, however, that MA plans do not have an obvious advantage in urban markets, where we already noted that enrollment is high. We return to this point later. Overall, we estimate the average cost of MA enrollees to be 587 dollars per enrollee-month, about 77 dollars (12 percent) less than our estimate of what it would cost to cover the same enrollees by FFS. The difference is substantial: with the current level of MA enrollment, it implies about \$13 billion a year in potential cost savings.

We can also estimate a monthly dollar markup for each plan in each county-year, by subtracting the estimated cost per standard enrollee from the plan's bid and multiplying by the plan's average risk score to obtain an enrollee-weighted average profit. We then take the enrollment-weighted average across plans for each county-year to obtain an average monthly plan profit for each market. Figure 7 plots these market-by-market estimates of private plan profits against risk-adjusted MA benchmarks for the same markets. The estimated mark-ups are substantial. We estimate the average profits accrued by MA plans for every month of enrollment to be 95 dollars. On a percentage basis, this suggests that plan margins are on the order of 16% above their (variable) costs of coverage, on average over our sample.

Summarizing the Welfare Effects We summarize our estimates of program surplus and how it is divided in the first row of Table 6. Over our study period, CMS paid an average of 756.2 dollars

for every enrollment month of an MA beneficiary. We have estimated that the private plan cost of providing baseline coverage to these beneficiaries was only 585.6 dollars. The difference, equal to 171 dollars per enrollee-month, is split between the beneficiary and the insurer. We estimate average monthly plan profits to be 95 dollars; the average rebate that went to enrollees was 76 dollars. However, we also estimate that the enrollee disutility from limited provider access was 27 dollars per month, so that the average enrollee surplus was 49 dollars per month, or about half of plan profits. Put another way, about two-thirds of the joint surplus to plans and enrollees accrued to the plans.

The comparison to FFS Medicare is also informative. We estimate that the average Medicare Advantage enrollee over this period would have cost taxpayers 662 dollars per month under FFS Medicare. Therefore, of the $95 + 49 = 144$ dollars of net surplus that accrued to plans and enrollees, only 50 dollars per month were true “gains from trade”. This amount represents the costs savings minus the enrollee disutility from private plan restrictions. The remaining surplus to plans and beneficiaries comes from taxpayers who, relative to FFS Medicare, provided a subsidy of 94 dollars per enrollee-month. The subsidy is higher than the average rebate, reflecting the fact that plan payments have somewhat exceeded FFS 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 current enrollment levels of around 14 million Medicare beneficiaries in private plans, the surplus gains amount to about \$8-9 billion a year. Taxpayers provide an additional \$15-16 billion subsidy. The result is that private plans receive around twice the incremental surplus created, while enrollee gains are about half this amount, or \$8 billion a 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 FFS Medicare.

8 Promoting Private Plan Competition

Overall, our estimates suggest that there are two factors contributing to higher taxpayer spending when Medicare beneficiaries enroll in private plans. One factor is that MA plan costs appear to exceed FFS costs in some markets, despite being lower on average. The second is that even when MA plan costs are cheaper, the plans often enjoy considerable market power. This allows them to mark up their bids and capture much of the cost savings as profits rather than reduced taxpayer

spending.

So while our general finding is that the MA program has been beneficial on average, it is interesting to explore how further changes in the MA market design could affect the program's costs and benefits, as well as the distribution of gains. We use our model and estimates to explore such market design changes; Table 6 summarizes our results. It is important to keep in mind that all our exercises in this section hold fixed the set of MA plans and their plan design, while in practice one would expect that insurers would also respond to changes in market design by offering new plans, dropping plans, or adjusting plan details (see, e.g., Frakt et al., 2012). We therefore interpret our main findings in this section as suggestive; the exact numbers should be interpreted with some caution given the factors we hold fixed.

We start with an attempt to provide an upper bound on the potential benefits of the current MA program by assuming that plans have no market power at all, i.e. that MA plans bid their costs. This is of course an extreme case, and may serve as a benchmark for the potential gains to consumers within the current design of rebates and benchmarks. Because much of the reduction in bids translates into increased rebates, this exercise results in only about \$24 per enrollee-month reduction in government spending. Mostly, the reduction in bids translates into greater rebates and more generous coverage. In this scenario, the MA enrollment rate jumps more than 50 percent. The marginal enrollees in private plans derive less surplus than the inframarginal enrollees, but on average across enrollees, consumer surplus from the rebate still increases to around 88 dollars per month.

Therefore, one way to understand the exercise is that under the current bidding rules, and fixing the benchmark rates, plan bids largely determine the surplus split between enrollees and insurers, with taxpayer costs per enrollee near-constant. Our second exercise provides an alternative type of bound. Here we assume that plans bid exactly the benchmark, so there are zero rebates. In this case, MA enrollment drops sharply, almost in half, showing that more generous coverage through the rebates is a key factor attracting beneficiaries into private plans. At the same time in this scenario, insurer profits per enrollee increase by 65 percent and consumer surplus declines. Government spending remains similar on a per-enrollee basis, at around \$14 more per enrollee-month.

We then report two exercises that reduce the benchmark rates. One applies a uniform reduction in benchmark rates by 50 dollars per enrollee-month, while the other sets the benchmark rates, which are often well above the FFS costs, to be equal to the FFS costs in each county. These exercises, by design, significantly reduce government spending per MA beneficiary, but due to

lower rebates, also reduce the share of Medicare beneficiaries who select Medicare Advantage. For example, at benchmark rates that equal the county-level FFS costs, the MA share is about a third lower than the observed levels, while a uniform \$50 reduction of benchmark rates results in a lower loss of MA share. As mentioned earlier, this calculation does not account for the potential exit of MA plans, which might amplify decreased enrollment effects.

Our last set of exercises explores the effect of changes in the rebate formula, by either allocating a smaller fraction (50% instead of 75%) of the difference between the bid and benchmark toward more generous plan benefits, or by allocating the full difference toward additional benefits. Changes in the rebate have two offsetting effects. On one hand, a larger rebate increases government spending and creates more generous coverage, leading to higher enrollee surplus. On the other hand, it makes beneficiaries more sensitive to changes in bids because on the margin they are affected more by bid changes. This leads to lower equilibrium markups. The results, reported at the bottom of Table 6, suggest that the effects are fairly strong. We estimate that a complete pass-through of bid reductions to consumers would substantially increase MA enrollment, albeit at the cost of significantly higher government spending per enrollee, around \$32 per enrollee-month. In contrast, reducing the rebate rate to 50% leads to a sharp reduction in MA enrollment and lower consumer surplus, without very noticeable taxpayer savings. The issue here is that with a smaller rebate share, demand becomes less sensitive and plans have an incentive to raise markups. In terms of strategies to reduce taxpayer costs, it seems that cutting the rebate rate is clearly dominated by reducing benchmark rates (again, abstracting from potential exit or changes in plan design).

To summarize, our general finding is that changes in market design can have a large effect on MA enrollment, and on the distribution of surplus. Our findings here should be viewed as suggestive, partly because we ignore the potential changes in the number and type of plans being offered, and partly because we have explored just a few of the possible reforms. For instance, one might wonder if a combination of a higher rebate rate and lower benchmarks might reduce taxpayer costs without greatly reducing enrollment. Our calculations suggest that this is in fact rather difficult to achieve if one focuses on uniform changes to benchmark rates, or changes that are tied to FFS costs (see Appendix Table A12). Benchmark changes more closely tied to plan costs might go farther in this regard, and it would be worth exploring mechanisms that could achieve this.

9 Conclusions

The reform and expansion of Medicare Advantage provides an opportunity to evaluate how private health insurance competition can work, compared to a public insurance benchmark. We propose a model of imperfect competition to think about this problem, while accounting for health differences between private plan enrollees and those enrolled in traditional Medicare. Our estimates suggest that in many areas of the country private plans are able to obtain costs savings over FFS 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 about half as much surplus as plans and taxpayers providing a large additional subsidy under the generous benchmark rates that prevailed in our data.

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 (as in Miller, 2014). Incorporating these complications would give a more nuanced view of competition.

Our treatment of potential market design interventions was also relatively brief. It focused only on simple changes in the benchmark rates and the size of plan rebates, and did not address the important issue of how program rules affect entry and exit. 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 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. Finally, a recent change of Medicare Advantage makes payments partly contingent on plan quality. It would be useful to understand whether this is effective, if enrollment decisions already are sensitive to plan quality.

References

- Aizcorbe, Ana, Eli Liebman, Sarah Pack, David M. Cutler, Michael E. Chernew, and Allison B. Rosen. 2012. "Measuring health care costs of individuals with employer-sponsored health insurance in the U.S.: A comparison of survey and claims data." *Journal of the International Association for Official Statistics* 28(1): 43-51.
- Baker, Laurence C., M. Kate Bundorf, and Daniel P. Kessler. 2010. "HMO Coverage Reduces Variations in the Use of Health Care among Patients Under Age Sixty-Five." *Health Affairs* 29(11): 2068-2074.
- Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston. 2014. "How Does Risk Selection Respond to Risk adjustment? Evidence from the Medicare Advantage Program." *American Economic Review* 104(10): 3335-3364.
- Bulow, Jeremy and Paul Pfleiderer. 1983. "A Note on the Effect of Cost Changes on Prices." *Journal of Political Economy* 91(1): 182-185.
- Cabral, Marika, Michael Geruso and Neale Mahoney. 2014. "Does Privatized Medicare Benefit Patients or Producers? Evidence from the Benefits Improvement and Protection Act." NBER Working Paper No. 20470.
- 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). 2013. *Medicare Managed Care Manual*. Publication 100-16. Available at <http://www.cms.gov>.
- Dowd, Bryan E., Roger Feldman, and Robert Coulam. 2003. "The Effect of Plan Characteristics on Medicare + Choice Enrollment." *Health Services Research* 38(1): 113-135.
- Duggan, Mark, Amanda Starc, and Boris Vabson. 2014. "Who Benefits When the Government Pays More? Pass-Through in the Medicare Advantage Program." NBER Working Paper No. 19989.
- Eggers, P. 1980. "Risk differential between Medicare beneficiaries enrolled and not enrolled in an HMO." *Health Care Financing Review* 1(3): 91-9.
- Eggers, P., and R. Prihoda. 1982. "Pre-enrollment reimbursement patterns of Medicare beneficiaries enrolled in 'at-risk' HMOs." *Health Care Financing Review* 4(1): 55-73.
- Enthoven, Alain. 1993. "The History and Principles of Managed Competition." *Health Affairs* 12(supplement): 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.
- Geruso, Michael, and Timothy Layton, 2014. "Risk Selection, Risk adjustment, and Manipulable Medical Coding: Evidence from Medicare." Mimeo, UT Austin.

- GAO (General Accounting Office). 2013. "Medicare Advantage: Substantial Excess Payments Underscore Need for CMS to Improve Accuracy of Risk Score adjustments." Report to Congress, GAO-13-206.
- Gowrisankaran, Gautam, Robert J. Town and Eric Barrette. 2011. "Managed Care, Drug Benefits and Mortality: An Analysis of the Elderly." *B.E. Journal of Economic Analysis and Policy* 11(2).
- Huang, Jennifer T., Gretchen A. Jacobson, Tricia Neuman, Katherine A. Desmond, and Thomas Rice. 2013. "Medigap: Spotlight on Enrollment, Premiums, and Recent Trends." *Kaiser Family Foundation Report*.
- Landon, Bruce E., Alan M. Zaslavsky, Robert C. Saunders, L. Gregory Pawlson, Joseph P. Newhouse, and John Z. Ayanian. 2012. "Analysis of Medicare Advantage HMOs Compared with Traditional Medicare Shows Lower Use of Many Services During 2003-09." *Health Affairs* 31(12): 2609-2617.
- MaCurdy, Thomas, Jay Bhattacharya, Daniella Perloth, Jason Shafrin, Anita Au-Yeung, Hani Bashour, Camille Chicklis, Kennan Cronen, Brandy Lipton, Shahin Saneinejad, Elen Shrestha, Sajid Zaidi. 2013. "Geographic Variation in Spending, Utilization and Quality: Medicare and Medicaid Beneficiaries." Report prepared by Acumen, LLC, commissioned by the IOM Committee on Geographic Variation in Health Care Spending and Promotion of High-Value Care, Institute of Medicine, Washington, DC.
- McGuire, Thomas G., Joseph P. Newhouse, and Anna D. Sinaiko. 2011. "An Economic History of Medicare Part C." *The Millbank Quarterly* 89(2): 289-332.
- MedPac (Medicare Payment Advisory Commission). 2007-2013. "A Data Book: Health Care Spending and the Medicare Program." Chapter 9 on Medicare Advantage.
- MedPac (Medicare Payment Advisory Commission). 2012. "Issues for Risk adjustment in Medicare Advantage." Report to the Congress: Medicare and the Health Care Delivery System.
- Miller, Keaton. 2014. "Do Private Medicare Firms have Lower Costs?" Mimeo, University of Minnesota.
- 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 Research* 48(3): 1039-1056.
- 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): 2618-2628.
- Newhouse, Joseph P., Mary Price, J. Michael McWilliams, John Hsu and Thomas McGuire. 2014. "How Much Favorable Selection is Left in Medicare Advantage?" NBER Working Paper No. 20021.

- Pizer, Steve D., and Austin Frakt. 2002. "Payment Policy and Competition in the Medicare + Choice Program." *Health Care Financing Review* 24(1): 83-94.
- Pizer, Steven D., Austin Frakt and Roger Feldman. 2009. "Nothing for Something? Estimating Cost and Value for Beneficiaries from Recent Increases on HMO Payments and Drug Benefits." *International Journal of Health Care Finance and Economics* 9(1): 59-81.
- Song, Zirui, David M. Cutler, Michael E. Chernew. 2012. "Potential Consequences of Reforming Medicare Into a Competitive Bidding System." *The Journal of the American Medical Association* 308(5): 459-460.
- Song, Zirui, Mary Beth Landrum, and Michael E. Chernew. 2012. "Competitive Bidding in Medicare: Who Benefits From Competition?" *The American Journal of Managed Care* 18(9): 546-552.
- Song, Zirui, Mary Beth Landrum, and Michael E. Chernew. 2013. "Competitive Bidding in Medicare Advantage: Effect of Benchmark Changes on Plan Bids." *Journal of Health Economics* 32(6): 1301-1312.
- Stockley, Karen, Thomas McGuire, Christopher Afendulis, Michael E. Chernew. 2014. "Premium Transparency in the Medicare Advantage Market: Implications for Premiums, Benefits, and Efficiency." NBER Working Paper No. 20208.
- Town, Robert and Su Liu. 2003. "The Welfare Impact of Medicare HMOs." *RAND Journal of Economics* 34(4): 719-736.
- Weyl, E. Glen and Michal Fabinger. 2013. "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition." *Journal of Political Economy* 121(3): 528-583.

Figure 1: The Growth of Medicare Advantage




Figure shows the number of MA contracts and the share 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/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 FFS 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 Figures 2 and 4 from McGuire, Newhouse, and Sinaiko (2011). Contracts include HMOs, local PPOs, Private FFS plans, and regional PPOs. 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 2014, which uses data from January.

Figure 2: MA share by risk score




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.

Figure 3: Mortality rates in MA and FFS Medicare, by risk score




Figure shows the mortality rate (over the subsequent calendar year), by risk score bin (0-0.1, 0.1-0.2, and so on), for beneficiaries who are in FFS Medicare, and in MA plans. Mortality for MA enrollees is plotted separately before and after 2010, when CMS deflated all the MA risk scores to adjust for differential coding. The Figure also shows mortality specifically for new enrollees in MA, i.e. beneficiaries who were in FFS the previous year. The gray bars show the underlying distribution of risk scores among Medicare beneficiaries.

Figure 4: The distribution of county benchmarks




Figure shows the distribution of county-specific benchmarks in 2006 (gray) and 2011 (black), the first and last years of our data. The figure illustrates that a large number of the benchmarks (68% of the counties in 2006, 66% in 2011) are within \$10 of the urban and (lower) rural floors, which were established by the Benefits Improvement and Protection Act of 2000. The rest of the counties are associated with other benchmarks, which are distributed over a range of about \$200 (per a standardized enrollee-month) above the floor.

Figure 5: Bid distribution (relative to benchmark)


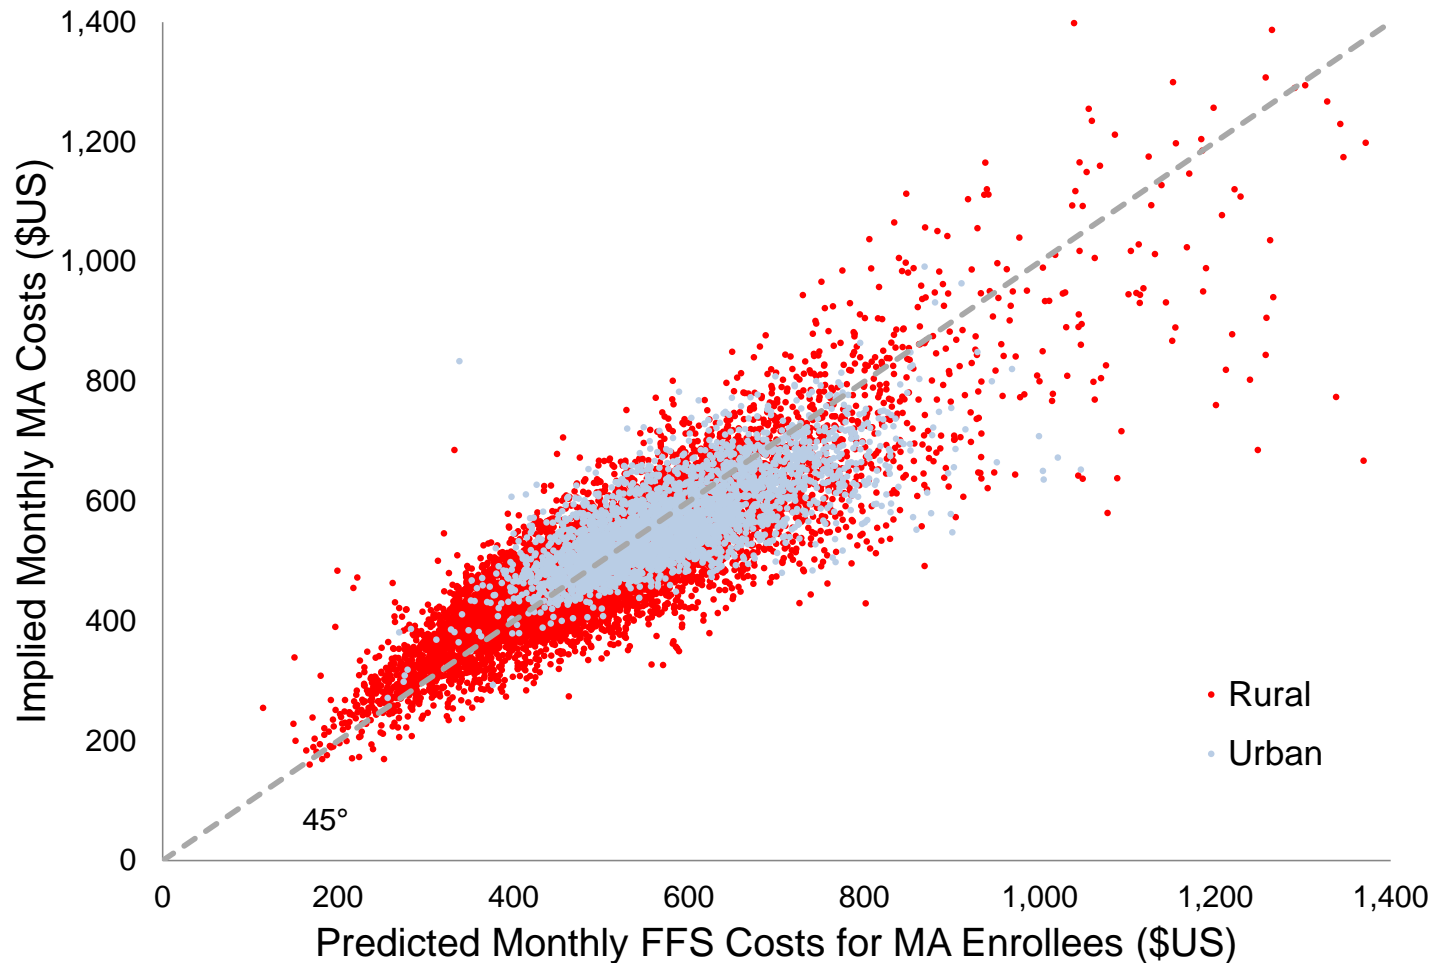



Figure shows a histogram of plan bids in the data (relative to the benchmark). Bids are for covering a standardized beneficiary over a month. Each observation is a year-plan, and we plot the distribution of the difference between the plan-specific bid in a given year and the plan-specific benchmark in that year (benchmarks are the same for all plans offering MA coverage in a given county, but plans are offered in service areas that cover multiple, often adjacent counties, and these service areas do not perfectly overlap across plans, making the relevant benchmark a plan-specific weighted average of county benchmarks). 93% of the observations are negative, which correspond to cases of bids that are below the benchmark. The average difference is -\$83 (standard deviation \$67), with the 5th, 25th, 50th, 75th, and 95th percentiles being -\$209, -\$114, -\$73, -\$40, and +\$5, respectively. The difference has risen slightly in magnitude over time, with average values changing from -\$77 in 2006 to -\$82 in 2011.

Figure 6: Implied costs of MA plans



Observations are at the market (county-year) level; there are 3,118 counties with benchmarks or 15,590 potential county-years from 2006-2010 but only the 15,211 of these with at least one MA enrollee are included in the figure. For clarity, we lose a small fraction of the observations that are associated with predicted FFS costs or implied MA costs that are greater than \$1,400. Counties that qualified as "urban" according to the definition used to set the urban floor in 2004 are shown in light blue, while all other counties are shown in darker red. Implied MA costs are derived from the insurer's first-order condition for optimal bidding (see text for details), using our preferred demand specification (Table 5(b), column 3). The market-level implied cost is computed as an enrollment-weighted average over the implied cost of all plans bidding in the market. The predicted FFS costs for MA enrollees is obtained using the method described in Section 3. Average implied MA cost is \$587 (standard deviation of \$81), and the average difference between implied MA cost and the predicted FFS cost is -\$47 (standard deviation of \$93). 47% of the markets are above the 45-degree line, with an average difference of \$40 and \$84 above and below the line, respectively.

Figure 7: Estimated MA markups



Observations are at the market (county-year) level; there are 3,118 counties with benchmarks or 15,590 potential county-years from 2006-2010 but only the 15,211 of these with at least one MA enrollee are included in the figure. For clarity, we lose a small fraction of the observations that have benchmarks greater than \$1,400 or estimated markups greater than \$200. Counties that qualified as "urban" according to the definition used to set the urban floor in 2004 are shown in light blue, while all other counties are shown in darker red. Estimates are derived from the insurer's first-order condition for optimal bidding (see text for details), using our preferred demand specification (Table 5(b), column 3). Market-level markups are enrollment-weighted and risk adjusted; they are computed by multiplying each plan's estimated mark-up by total plan risk units, adding these up for all plans in the market, and dividing by total MA enrollees in the market. The (weighted) average markup is \$96 (standard deviation of \$21). To make the benchmark comparable to the markups, we also risk adjust it; it is computed by multiplying the administrative market-level benchmark by total MA risk-months in the market and dividing it by total MA enrollee-months in the market. Statistics are weighted by market-level MA enrollee-months. The (risk-adjusted) average benchmark is \$779 (standard deviation of \$147).

Table 1: Summary statistics

Variable	Traditional Medicare (TM)				Medicare Advantage (MA)			
	Mean	Std. Dev.	10th Pctile	90th Pctile	Mean	Std. Dev.	10th Pctile	90th Pctile
No. of Obs. (beneficiary-years)		125,669,915				28,188,896		
Unique beneficiaries		30,403,528				7,842,799		
Age ^a	76.5	25.6	67.4	87.2	76.1	24.2	67.6	86.2
Male	0.435				0.432			
Urban ^b	0.649				0.839			
New Medicare enrollee	0.026				0.017			
Supplemental insurance ^c	0.727				0.061			
Part D coverage	0.410				0.937			
Inpatient days	1.49	19.32	0	4	0.83	13.01	0	2
Died during year	0.022				0.019			
Risk score	0.97	2.81	0.35	1.95	0.93	2.69	0.34	1.87
Log(Risk score change) ^d	0.05	1.75	-0.53	0.69	0.12	1.50	-0.29	0.67
FFS monthly claims costs ^e	637.3	5,316.9	9.1	1,705.2				
MA monthly total CMS payment ^f					746.8	1,988.8	303.0	1,422.6
MA monthly rebate payment ^f					75.6	54.5	18.6	146.4

Statistics are based on full sample of beneficiaries from 2006-2011. The Traditional Medicare (TM) sample excludes beneficiaries dually eligible for Medicare and Medicaid as well as Disabled beneficiaries. The Medicare Advantage sample excludes Disabled beneficiaries as well as those enrolled in employer-sponsored plans or special needs plans. See the text for exact sample restrictions. Statistics are weighted by total months of enrollment during the observation year. ^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 – $\log(\text{risk score in following year}) - \log(\text{risk score in observation year})$ -- is defined over the sample of beneficiaries who survive until the following year. ^e FFS monthly claims costs are computed by adding up total FFS claims for all TM enrollees, adding up total months of enrollment for all TM enrollees, and dividing the former by the latter. Data is for 2006-2010. ^f Mean MA monthly plan payment and rebate payment are defined analogously to FFS cost. Data is for 2006-2011. The 2006-2010 averages are almost the same (746.0 and 75.6 for total and rebate payments, respectively).

Table 2: Predicted TM cost for MA enrollees

	All Enrollees	Risk Adjusted	Urban	Rural	2006-07	2008-09	2010
<u>TM Enrollees</u>							
FFS Cost (actual)	637.3	656.6	674.1	569.6	597.9	653.7	686.2
FFS Cost (predicted)	637.3	656.6	674.1	569.6	597.9	653.7	686.2
Risk Score	0.971	1.000	0.996	0.924	0.991	0.957	0.955
<u>MA Enrollees</u>							
FFS Costs (predicted)	675.0	721.1	684.7	623.7	619.1	691.7	721.4
MA Plan Payments	670.4	716.2	691.3	560.0	657.0	681.5	668.3
MA Rebate Payments	75.6	80.8	80.4	50.5	68.8	84.1	69.5
MA Total CMS Payments	746.0	796.9	771.7	610.5	725.8	765.6	737.8
Risk Score	0.936	1.000	0.956	0.831	0.954	0.943	0.898 ^a

All numbers in the table are in current dollars per enrollee-month. Statistics are based on the sample of beneficiaries from 2006-2010, the sample period for which FFS claims are observed. Urban means county with population > 250,000 in 2004, when the urban floor was initially set. Per-enrollee-month FFS costs are computed by adding up all FFS costs for TM enrollees, adding up all months enrolled for TM enrollees, and dividing the former by the latter. Risk-adjusted FFS costs are computed by adding up all FFS costs for TM enrollees, adding up all risk-months (risk score multiplied by months enrolled) for TM enrollees, and dividing the former by the latter. All other columns report per-enrollee (not risk adjusted) amounts. Predicted FFS cost for MA enrollees is computed using a market-by-market risk-adjusted measure. For each market, we add up FFS costs for TM enrollees in the market, we add up risk-months enrolled for TM enrollees in the market, and we divide the former by the latter to obtain a market-level cost factor. For a given market, we multiply each MA enrollee's risk score by the market-level cost factor to obtain his predicted FFS costs. See text (Section 3) for further details.

^a In 2010, CMS deflated MA risk scores by 3.41% in order to calculate payments; the risk scores used in the table incorporate the deflation adjustment.

Table 3: The pass-through rate of TM cost and benchmarks to plan bids

Unit of Observation: Plan-year, Dependent Variable: Plan bid			
Sample	All plans (1)	All plans (2)	All plans (3)
Plan benchmark	0.384 (0.048)***	0.370 (0.056)***	0.466 (0.039)***
Predicted plan FFS cost	0.055 (0.073)	0.083 (0.099)	0.161 (0.073)**
Year FEs	N	Y	Y
Contract FEs	N	N	Y
Mean of dependent variable	730.1	730.1	730.1
R-squared	0.279	0.304	0.750
Observations	10,305	10,305	10,305

Unit of Observation: Plan-county-year, Dependent Variable: Plan bid			
Sample	All plans (1)	All plans (2)	All plans (3)
Plan benchmark	0.384 (0.083)***	0.370 (0.095)***	0.642 (0.044)***
Predicted plan FFS cost	0.055 (0.058)	0.083 (0.084)	-0.009 (0.047)
Year FEs	N	Y	Y
County FEs	N	N	Y
Mean of dependent variable	730.1	730.1	730.1
R-squared	0.279	0.304	0.675
Observations	181,868	181,868	181,868

Table presents regression results of plan bid on plan benchmark and predicted plan cost. The unit of observation is a plan-year in the top panel, and a plan-county-year in the bottom panel. The sample in the top (bottom) panel consists of all plan-years (all plan-county-years) in our 2006-2010 data with at least one enrollee. In the top panel, the plan benchmark is computed by multiplying the county-year-level administrative benchmark by plan risk-month enrollment in that county-year, adding this up across all counties in the plan's service area for the given year, and dividing by total plan risk-month enrollment for the given year. The predicted plan FFS cost is computed in each panel using the method described in Section 3. In the bottom panel, we weight observations by total risk-month enrollment for the given plan-county-year. Standard errors are clustered at the contract level in the top panel, and at the county level in the bottom panel.. **p < 0.05, ***p < 0.01.

Table 4: The pass-through rate, by plan type

Unit of Observation: Plan-year, Dependent Variable: Plan bid			
Sample	HMO plans (4)	Local PPO plans (5)	Private FFS plans (6)
Plan benchmark	0.452 (0.044)***	0.710 (0.048)***	0.645 (0.079)***
Plan cost	0.110 (0.079)	0.287 (0.053)***	0.516 (0.072)***
Year FEs	Y	Y	Y
County FEs	Y	Y	Y
Mean of dependent variable	731.6	745.3	715.1
R-squared	0.760	0.867	0.791
Observations	5,962	2,022	2,321

Unit of Observation: Plan-county-year, Dependent Variable: Plan bid			
Sample	HMO plans (4)	Local PPO plans (5)	Private FFS plans (6)
Plan benchmark	0.673 (0.063)***	0.678 (0.055)***	0.623 (0.024)***
Plan cost	-0.034 (0.056)	0.310 (0.059)***	0.433 (0.020)***
Year FEs	Y	Y	Y
County FEs	Y	Y	Y
Mean of dependent variable	731.6	745.3	715.1
R-squared	0.723	0.843	0.724
Observations	32,645	17,011	132,212

Table presents regression results of plan bid on the plan benchmark and predicted plan cost. The regressions are the same as those reported in Table 3, column (3), with the only difference being that they are now estimated for each plan type -- HMO, PPO, and Private FFS plans -- separately. See the notes to Table 3 for additional details about the units of observation, variable definitions, and the units on which we cluster the standard errors. **p < 0.05, ***p < 0.01.

Table 5: Demand estimates (all beneficiaries)**Panel A: Logit demand**

Dependent Variable: ln(plan risk-months market share) - ln(TM risk-months market share)					
Mean of Dep. Variable = -6.251; No. of Obs. = 206,110					
	(1)	(2)	(3)	(4)	(5)
Plan price	-0.005 (0.002) ^{***}	-0.007 (0.002) ^{***}	-0.004 (0.002) ^{**}	-0.014 (0.002) ^{***}	-0.009 (0.002) ^{***}
Supplemental benefits			-0.343 (0.176) [*]	-0.183 (0.146)	0.045 (0.135)
Part D benefits			0.585 (0.211) ^{***}	0.450 (0.257) [*]	0.487 (0.288) [*]
Plan quality rating FEs	N	N	Y	Y	Y
Year FEs	N	Y	Y	Y	Y
Contract FEs	N	N	N	Y	Y
Contract x county FEs	N	N	N	N	Y
Mean own-price elasticity	-0.368	-0.518	-0.301	-1.003	-0.638

Panel B: Nested logit demand

Dependent Variable: ln(plan risk-months market share) - ln(TM risk-months market share)			
Mean of Dep. Variable = -6.251; No. of Obs. = 206,110			
	(1)	(2)	(3)
Plan price	-0.012 (0.002) ^{***}	-0.012 (0.002) ^{***}	-0.012 (0.002) ^{***}
ln(plan MA share)	0.307 (0.029) ^{***}	0.299 (0.031) ^{***}	0.320 (0.023) ^{***}
Supplemental benefits	-0.143 (0.112)	-0.144 (0.113)	-0.142 (0.110)
Part D benefits	0.317 (0.178) [*]	0.320 (0.183) [*]	0.311 (0.177) [*]
Plan quality rating FEs	Y	Y	Y
Instrument for ln(plan MA share)	Number of plans	Number of contracts	Contract dummies
Year FEs	Y	Y	Y
Contract FEs	Y	Y	Y
Mean own-price elasticity	-1.147	-1.142	-1.158

Table presents demand regression results at the market-plan level. The unit of observation is a market-plan (a market is a county-year). Standard errors are clustered at the contract level. Letting b denote the plan's standardized bid and B denote the plan's benchmark, "plan price" is defined as $p = 0.75 \times (b - B) \times I\{b - B \leq 0\} + (b - B) \times I\{b - B > 0\}$. The mean of the own-price demand elasticities reported in the final row is weighted by enrollee risk-months (further details are provided in the appendix). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Summary of results and policy experiments

	Taxpayer Cost	Implied MA Cost	Insurer Profits	Consumer Total Surplus	Consumer Rebate Surplus	Rebate Amount	Beneficiary Premium	MA Share	Predicted FFS Cost for MA Enrollees
Observed equilibrium	756.2 ^a	585.6 ^b	95.4	102.6	48.8	75.6	0.4	18.2%	662.5 ^c
Plans bid their cost	732.6	581.8	0.0	119.1	87.9	150.8	0.0	31.4%	660.3
Plans bid the benchmark	746.5	590.5	156.0	94.8	0.0	0.0	0.0	10.2%	654.7
Benchmark lowered by \$50	715.8	586.5	73.7	99.8	37.9	57.7	2.2	15.4%	663.0
Benchmark lowered to FFS cost	730.7	598.5	75.6	103.1	39.9	66.7	10.0	13.7%	677.7
Rebate passed through at 50%	743.7	590.0	132.2	96.1	16.7	22.3	0.8	12.2%	659.9
Rebate passed through at 100%	788.6	581.5	72.8	114.0	76.9	134.4	0.2	26.4%	663.7

All reported statistics are per MA enrollee-month, and use our preferred demand specification (Table 5, Panel B, Column 3). 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. We divide this number by the MA share in the overall sample to obtain the consumer surplus per MA enrollment. We then repeat the same exercise except that we set the rebate to zero when we compute the value of each plan. We subtract the result from the previously computed consumer surplus to obtain the consumer rebate surplus. 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 In Table 1, the mean taxpayer cost is reported as \$746.8; however, here we report it as \$756.2. The \$756.2 number is calculated using the realized enrollment weights and actual bids to impute MA payments and is reported here since it is more directly comparable to the costs reported for the counterfactuals. ^b In Figure 6, we report the mean implied MA cost as \$587, which differs from the \$585.6 that we report here; this is because in Figure 6 we exclude the year 2011. ^c Based on 2006-2010, since we do not have FFS cost data for 2011.