

MANAGED COMPETITION IN HEALTH INSURANCE

Liran Einav
Stanford University

Jonathan Levin
Stanford University

Abstract

Rising healthcare costs have sparked debate about the best way to provide high-quality affordable health insurance. We discuss the potential for regulated insurance markets to outperform single-payer public insurance. We use as an example the private Medicare plans that now provide insurance to almost a third of seniors in the United States. The evidence suggests that private plans can limit costs and potentially appeal to enrollees, and that well-designed risk adjustment can mitigate market failures due to adverse selection. However, fostering competition between insurers, especially in smaller markets, is difficult. We discuss how future research might illuminate the relative advantages of public and private health insurance. (JEL: D12, H20, H71, L81)

1. Introduction

The high and rising costs of healthcare are a topic of considerable interest in many countries. They have been especially salient in the United States, where healthcare spending has grown to almost 20% of GDP, and is currently around twice the OECD per capita average (Figure 1).¹ Of particular concern is that the high rate of spending does not seem to translate into easily measurable benefits in terms of population health (OECD 2013). The renowned Dartmouth Atlas Report (Fisher et al. 2009) similarly has documented wide variation in US Medicare spending—states such as Texas and Louisiana spend more than twice as much per capita as Minnesota or North Dakota—without clearly measured health benefits. These findings have led to concerns that the US healthcare system is deeply inefficient.

The editor in charge of this paper was Juuso Välimäki.

Acknowledgments: This paper is based on the Alfred Marshall Lecture delivered by Einav at the EEA–ESEM meetings in Toulouse on August 28, 2014. We thank Juuso Välimäki and an anonymous referee for comments, Jay Bhattacharya, Yiwei Chen, and Vilsa Curto for related discussions, and the Stanford Institute for Economic Policy Research for research support. Einav and Levin are both Research Associates at the NBER.

E-mail: leinav@stanford.edu (Einav); jdlevin@stanford.edu (Levin)

1. Fuchs (2014) attributes the difference primarily to a more expensive mix of services and procedures in the United States. For related discussions, see Hall and Jones (2007) and Garber and Skinner (2008).




FIGURE 1. Healthcare spending in the United States and other OECD countries. Source: OECD Health Data (October 2010).

The United States is also distinct in its mix of public and private provision of health insurance. Individuals and employers purchase health plans from private insurance firms, while most seniors have publicly administered Medicare insurance.² This dichotomy has generated a long debate between advocates of private insurance markets and advocates of single-payer public insurance. The former argue that private insurers are more cost-efficient, and that regulated “managed competition” marketplaces can harness this efficiency (Enthoven 1993). The latter argue that public insurance is simpler and has lower administrative costs, and that private markets suffer from risk selection and other competitive failures.

In recent years, the United States increasingly has moved in the direction of managed competition. Around 30% of US seniors now enroll in private Medicare Advantage plans rather than public Medicare insurance. Among Americans with employer-sponsored insurance, half now have a choice of competing insurance plans (AHRQ 2012). And under the Affordable Care Act, the individual health insurance market now revolves around regulated state insurance exchanges. While each of these settings has its own details, they all share commonalities. All rely on regulations to ensure that plans offer certain baseline benefits, on consumer choice among approved plans, and on market rules to promote competition and limit incentives for risk

2. In the average OECD country, the government funds about 75%–80% of health expenditures, compared to 50% in the United States, although in countries such as the Netherlands, Switzerland, and Germany, government funding flows to competing private insurers.

selection. These features also appear in the managed competition systems used in European countries such as the Netherlands, Switzerland, and Germany.

The growth of managed competition has generated a corresponding flurry of research, and our goal in this paper is to ask what lessons this research can teach us about how well managed competition can work.³ There are really two parts to this question. The first relates to productive efficiency. Are private insurers more cost efficient than the public sector, and can they provide a given level of insurance benefits at lower cost? As we discuss in the next section, there is some evidence in the United States that they can, although this evidence remains more speculative than one would like. The second part of the question is whether the competition part of managed competition can work. That is, does competition between insurers lead to higher quality or to lower prices for consumers and the taxpayers who subsidize insurance purchases?

We focus in particular on two central problems with health insurance competition. The first is risk selection. In the settings already mentioned, healthy and sick enrollees usually pay the same premiums to obtain insurance. One may worry as a result that insurers will want to target only healthy enrollees. A central tenet of managed competition is that risk-adjustment schemes can mitigate this incentive, by compensating insurers with relatively sick enrollees. In Section 3, we argue that risk-adjustment can be fairly successful and we review some of the supporting evidence. We also observe that risk-adjustment systems tend to rely on highly imperfect predictive risk models, and we explain some of the reasons for this, such as the need to use transparent and interpretable models, and the potential for insurers to manipulate risk scores to receive higher payments. A key lesson is that risk-adjustment mechanisms should be evaluated based on the incentives they create for insurers, rather than solely on predictive power.

A second concern with managed competition is that insurers may not end up with strong incentives to compete on price and quality. Insurance markets tend to be highly concentrated. For example, in the market for Medicare private plans, the top two insurers in a local market on average combine to have around 85% of private plan enrollees. The incentive for insurers to compete on price and quality also depends on consumers being informed and price-sensitive, and neither of these conditions is guaranteed. In Section 4, we discuss some recent evidence showing how the combination of market power and limited consumer price sensitivity can lead to high margins on insurance plans, pushing up the price for consumers and taxpayers. One possibility is that competitive entry by insurers, or the provision of information to consumers might improve competitive incentives. However, the key lesson we take away is that in the absence of these effects, careful market design choices are essential to provide effective incentives for insurers.

3. We will not attempt a comprehensive list of references, but any such list would include work on employer-sponsored health plan choice, Medicare Advantage, the Medicare Part D prescription drug market, the Massachusetts state health insurance exchange, the state exchanges established under the Affordable Care Act, as well as work on European managed competition systems.

Throughout the paper, we use the US Medicare program, and its Medicare Advantage managed competition system, as our leading example. Apart from its large size and importance (US taxpayers spend over \$130 billion annually on Medicare Advantage), one reason to focus on Medicare is that Medicare's private insurance plans operate in parallel to, and compete with, Medicare's publicly administered insurance. This makes it a useful test case to evaluate managed competition's costs and benefits relative to single-payer public insurance. Focusing on Medicare also allows us to draw on some of our own recent research that tries to quantify these costs and benefits (Curto et al. 2015). Of course, as we already noted, any specific setting has its idiosyncrasies, but we also think there are enough commonalities in managed competition that lessons from one market can be useful in thinking about other related markets.

The paper proceeds as follows. The next section describes the Medicare setting, and its managed competition program. We discuss the program in some detail in order to highlight the importance of market rules in shaping competitive incentives. We also discuss some of the evidence on private plan costs relative to public insurance costs. In Section 3, we take up the problem of risk selection and risk-adjustment mechanisms. Section 4 considers market power and insurer competition. The final section is more forward looking. We identify a few questions associated with healthcare competition that we view as important but relatively unexplored, and where further research would be desirable.

2. Managed Competition and US Medicare

Health systems rely on different mixes of public and private health insurance, and there is no consensus on which is preferable. As we have already observed, advocates of single-payer public insurance tend to argue that it is simpler and involves less administrative overhead, while stressing private market failures such as adverse selection, consumer confusion, and the underprovision of quality. Advocates for market competition argue that public programs are fraught with wasteful spending and political influence, and highlight the incentives that markets can create for efficiency and innovation. The US Medicare program provides a useful setting to evaluate some of these claims, because of its dual structure whereby beneficiaries may enroll in either publicly administered insurance, or in private insurance plans that compete in a managed competition marketplace.

2.1. Medicare Insurance

The Medicare program provides access to health insurance for Americans aged 65 and older, as well as others with disabilities and specific illnesses. It is one of the largest federal government programs with around 50 million beneficiaries and an annual budget of over \$500 billion.

Medicare beneficiaries can enroll in publicly administered insurance through Medicare Parts A and B. Medicare Part A provides coverage for all inpatient spending

(i.e., expenditure associated with an overnight stay at a hospital or nursing facility). Medicare Part B provides insurance for 80% of eligible outpatient expenses, such as visits to a physician, laboratory work, outpatient surgery, or radiology. Beneficiaries pay nothing to enroll in Part A, and a subsidized monthly premium for Part B (in 2014, it was \$105).⁴ Once enrolled, they have access to a wide range of physicians and hospitals. Medicare compensates these providers based on an administered price schedule for outpatient services, or in the case of hospital admissions, by paying a fixed amount based on the patient's condition.

Medicare beneficiaries alternatively may enroll in a private insurance plan through Medicare Advantage (also known as Medicare Part C). Medicare private plans cover at least the same services as public Medicare insurance, often with lower coinsurance payments by the enrollee. The tradeoff for enrollees is that private plans also restrict coverage to a limited network of providers. Beneficiaries who enroll in a private plan continue to pay the Medicare Part B premium, and usually no more than this. At the same time, Medicare pays the plan a fixed amount per month for each enrollee.

For the last decade, Medicare also has run a separate program for prescription drug insurance, known as Medicare Part D. Beneficiaries in both public Medicare and Medicare private plans may purchase Part D coverage, and about two-thirds do. Some Medicare private plans also include prescription drug coverage as part of a bundle, so beneficiaries sometimes purchase both their standard insurance benefits and drug coverage with one-stop shopping.

2.2. Medicare Advantage: A Case Study for Competition

The Medicare Advantage program dates back to the 1980s. It was started with the goal of expanding choice for Medicare beneficiaries, and also allowing taxpayers to capture some of the cost savings that health maintenance organizations appeared to be generating in the private insurance market (McGuire, Newhouse, and Sinaiko 2011). In the original program design, plans competed to enroll beneficiaries, but plan payments were set administratively by Medicare at the county level. Partly due to the difficulties in setting these payments, the program remained limited in size. Medicare reformed the program in the mid-2000s and introduced two pillars of managed competition: a risk-adjustment mechanism to match payments to enrollee health, and competitive bidding to set payments. Since then, the program has expanded to cover around 30% of Medicare beneficiaries (Figure 2).

Competition is tightly regulated in the Medicare Advantage program. First, plans must provide at least the same insurance benefits as public Medicare. Typically, they offer additional benefits as well, such as reduced cost-sharing, or dental and vision coverage, or a discount on a bundled prescription drug plan. However, plans cannot simply add on these benefits. Instead they must "fund" them at their actuarial value.

4. Beneficiaries with low incomes may receive additional subsidies. Enrollees also have the option to purchase additional Medigap insurance, offered by private insurers, which covers insurance copayments and deductibles.

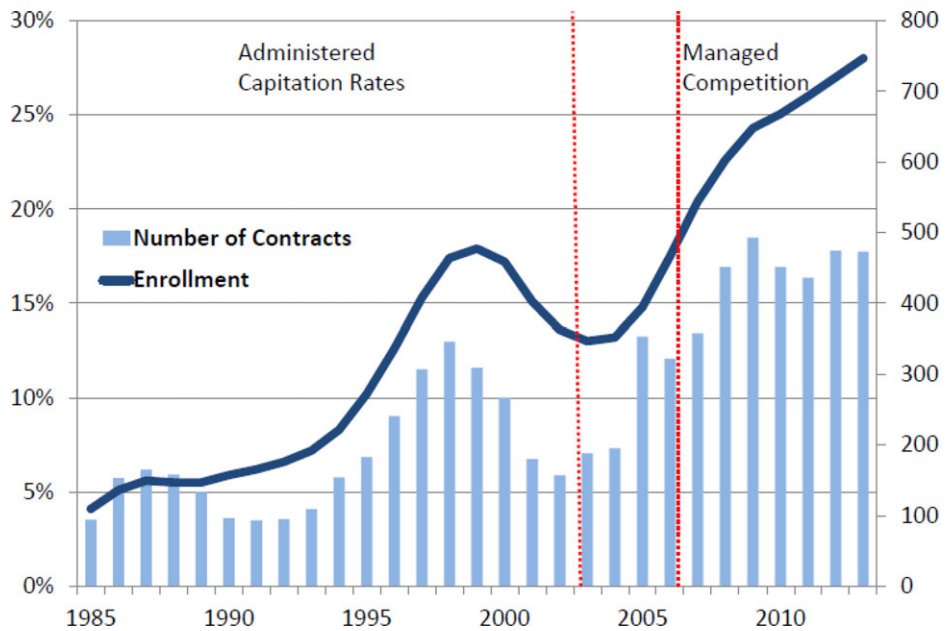


FIGURE 2. The growth of Medicare Advantage. This figure presents the growth in the number of Medicare Advantage contracts and enrollees. A contract is the primary unit by which an insurer is approved to offer Medicare Advantage plans. The two vertical lines identify the period through which the new competitive bidding and risk-scoring system has been gradually phased in.

The funding can come through a supplemental premium paid by the insuree, or through a rebate paid by Medicare. The rebate is determined in the competitive bidding process.

The bidding process plays a central role. Medicare starts by setting a benchmark reimbursement rate for each US county. Each plan then submits a bid, which is a monthly payment that the plan demands to provide standard Medicare insurance benefits to an average enrollee. Plan bids are evaluated relative to their respective local benchmark. If a plan bids below the benchmark (and almost all plans do), the discount generates a rebate that can be used to fund additional insurance benefits and make the plan coverage more attractive to enrollees.⁵

The competitive design also attempts to allow taxpayers to benefit from low plan bids. To see how it works, let B denote the benchmark rate and b a plan bid. If a beneficiary enrolls in the plan, the “savings” $B - b$ is divided. Medicare keeps 25% and rebates the other 75% to the plan. Plans must pass the rebates on to their enrollees, rather than retain them. Medicare also separately reimburses the plan for providing the standard insurance coverage, and does so according to the plan’s bid. This payment is

5. Interestingly, plans most commonly use the rebates to provide more generous insurance coverage, rather than trying to return the rebate money directly to their enrollees.

risk-adjusted. So if a plan's bid is b and an enrollee has risk score r —that is, expected costs r times that of an average beneficiary—then Medicare pays the plan rb .⁶

We will return to some of the details of the risk-scoring system and the bidding mechanism in the next two sections. For the moment, we note that there have been two frequent criticisms of Medicare Advantage. The first is that risk selection results in plans being overcompensated. Risk selection—the propensity of private plans to primarily enroll healthy beneficiaries, while sicker ones remain in public insurance—has been documented since the beginning of the Medicare Advantage program (McGuire, Newhouse, and Sinaiko 2011). The second criticism is that competition has not resulted in taxpayer savings. Studies over a long period have concluded that it typically costs taxpayers more when a beneficiary opts out of public Medicare and enrolls in a private plan (MedPAC 2013). In this sense, the program provides an ideal setting to discuss both risk selection and imperfect competition, which we see more generally as central issues for managed competition.

2.3. *Private versus Public Provision of Insurance*

The Medicare program in principle also provides an opportunity to compare the relative efficiency of private and publicly administered insurance. Medicare private plans run in parallel to the public insurance program; they have enrollees of the same age; and they provide the same baseline insurance benefits. In practice, however, the comparison is not so easy because of data limitations. While thousands of studies have analyzed the services provided to public Medicare enrollees, and the payments Medicare makes for these services, there is far less direct evidence on the healthcare received by Medicare private plan enrollees and the costs that private insurers incur. Nevertheless, the available evidence does suggest that private insurers may well have lower costs than public insurance, and also a somewhat different cost structure.

One indication of this comes from the bids submitted by Medicare Advantage plans. The most important private plans in the Medicare program, in terms of enrollment, are health maintenance organization (HMO) plans. According to the Medicare Payment Advisory Commission (MedPAC), these plans submitted bids in 2014 that were on average 95% of the cost of covering the same enrollees under public Medicare (MedPAC 2014). The MedPAC numbers potentially overstate the public insurance costs by not fully accounting for the good health of private plan enrollees, but Curto et al. (2015) find that making a plausible corrective adjustment adds only a few percentage points.

At the same time, the healthcare costs incurred by the plans are most likely well below their bids. The information available on large insurers suggests that they have tended to pay out in claims around 80% to 85% of the revenue they receive.⁷ In the

6. There is also a less common case where a plan bids above the benchmark. Then the enrollee does not receive extra rebate benefits and instead must pay the excess bid $b - B$ directly to the plan. Medicare also pays the plan $rb - (b - B)$ so that in total the plan receives rb to provide the standard insurance. Between 2006 and 2011, only 7% of plan bids were above their respective benchmarks.

7. For example, United Healthcare Group is the largest Medicare Advantage insurer. In their 2014 Annual Report, United reported that across all of their health insurance products, they paid out in claims 80.9%

Medicare context, a plan's bid determines its revenue, so applying the 80%–85% claim rate would indicate that private plans pay out in medical claims around 15% to 25% less than public Medicare. This calculation matches up fairly well with cost calculations in Curto et al. (2015), which are based on subtracting an estimate of “optimal” plan bid markups from observed plan bids.

Why might private insurance costs be lower than that of public insurance? A key difference in the US context is that public Medicare insurance does relatively little to manage the utilization of healthcare services. For instance, in Medicare Part B, providers receive a fixed rate per service. This gives providers little economic incentive to choose cost-effective treatments, or to forego expensive treatments with low but positive benefits. Similarly, patients have a wide choice of providers and pay the same amount whomever they choose. They have an incentive to choose the highest-quality provider rather than the most cost-effective. In contrast, private insurers have a strong financial incentive to limit marginally valuable care, and frequently restrict the set of providers that their enrollees can access. The limited evidence that is available indicates that private plans do in fact reduce utilization of many services (Landon et al. 2012).

Of course, public Medicare enjoys a strong bargaining position when it comes to setting reimbursement levels, whereas private insurers must negotiate payment rates with providers who can have considerable market power. So it is possible that private insurers might have lower utilization, but pay more for specific services. There is some indirect evidence that private plans do in fact have different cost structures. For instance, MaCurdy et al. (2013) documented wide variation across the United States in both private insurance costs and public Medicare costs, but a relatively low correlation. Public Medicare costs tend to be high in regions where, all else equal, utilization of services is high. One hypothesis is that private cost variation results at least in part from different bargaining outcomes in setting prices (Casey 1998).

The studies we have mentioned all focus on the cost of providing insurance benefits rather than the quality of care or health outcomes. The latter are even more difficult to measure. The growth of Medicare private plan enrollment suggests that private plans do not offer noticeably worse care, at least from the perspective of Medicare beneficiaries. However, better data and more research will be needed to establish if there are measurable differences in health outcomes.⁸

of their revenue from premiums. This compares to 81.5% and 80.4% in the prior two years, and to 82.6% reported by Aetna (the second largest Medicare Advantage insurer) in their 2014 annual report. Going forward, these so-called “medical loss ratios” or MLRs may increase because of federal regulations that require insurers to have MLRs of at least 80%–85% depending on the exact market.

8. Medicare does collect disease codes on beneficiaries in private plans, but as we discuss in what follows, the coding varies significantly between public insurance and private providers. One can compare mortality rates (Gowrisankaran, Town, and Barrette 2011; Curto et al. 2015), but these are probably best studied over a relatively long time period.

3. Risk Selection and Risk Adjustment

Risk selection is a key concern associated with insurance markets. On the demand side, adverse selection can arise when individuals who are sicker or expect to utilize more services self-select into more generous insurance. This leads companies selling generous policies to price the insurance for these risky individuals, potentially leaving healthier individuals with inefficient coverage (Akerlof 1970; Einav, Finkelstein, and Cullen 2010). In extreme cases, adverse selection can cause a market to unravel entirely (Cutler and Reber 1998; Hendren 2013). On the supply side, unpriced risk heterogeneity motivates insurance companies to cater offerings to healthier, lower cost individuals, or to deny coverage to high-cost, sick individuals.

A textbook solution is to allow insurance companies to charge higher premiums for higher-risk individuals. However, this can be problematic in the context of health insurance. First, it may seem unjust for the sickest to be charged the most. Second, it is regressive. Income and health are often negatively correlated, and pricing risk is likely to mean higher premiums for lower-income households. Finally, risk-based premiums expose individuals to “reclassification risk” : variability of premiums from year to year due to changes in one’s health. Indeed, Handel, Hendel, and Whinston (2015) provide evidence that reclassification risk may lead to greater inefficiencies than adverse selection (when risk cannot be priced).⁹

Managed competition markets attempt to avoid these problems by distinguishing the premiums paid by enrollees from the fees collected by insurers. Either insurers are required to make transfers that compensate adversely selected plans, or if there is a market sponsor that pays a substantial fraction of coverage costs, enrollees can pay uniform or near-uniform premiums, while the sponsor risk adjusts its payments according to the health of a plan’s enrollees. Medicare Advantage and employer-sponsored health plan choice use the sponsorship model, whereas state insurance exchanges implement balanced budget transfers. While risk adjustment may not provide individuals with an incentive to self-select efficiently across plans (Bundorf, Levin, and Mahoney 2012), it can reduce or eliminate the motivation for plans to engage in risk selection.

3.1. Risk Adjustment in Medicare Advantage

Risk adjustment requires the measurement of enrollee costs, which can be done in a variety of ways. Costs can be measured ex post, either based on a health classification of insureds during the coverage period (as in the US state insurance exchanges) or based on realized claims (as in Medicare Part D). More frequently, payments are adjusted using ex-ante information and a predictive model of expected healthcare costs.

9. There is also an empirical question of just how important are market failures due to asymmetric information in insurance markets. The answer almost surely varies substantially depending on the market, but a number of recent studies of health insurance have found relatively small welfare losses from adverse selection (Einav, Finkelstein, and Levin 2010).

The procedure used in Medicare Advantage is a fairly typical one. It is based on disease diagnoses over the prior twelve months. These diagnoses are categorized into 167 hierarchical classification codes, or HCCs. The term “hierarchical” is used because some codes refer to distinct conditions (e.g., diabetes, pulmonary disease, and so on), whereas others are refinements (e.g., distinguishing “diabetes” from “diabetes with complications”). Medicare uses a predictive model that maps these codes (in practice, a subset of 78 codes) into expected Medicare claims. The model is estimated on claims data from a sample of beneficiaries in public Medicare. Based on the estimates, every Medicare beneficiary, including those in private plans, is assigned a predictive risk score. The scores are normalized so that an individual with risk score $r = 1$ is expected to have average spending, while an individual with risk score $r = 2$ is expected to have double the average.

The evidence from Medicare Advantage suggests that risk adjustment can help significantly with risk-selection incentives. Prior to the development of the predictive scoring system, Medicare Advantage payments within a geographic area included small adjustments for age and low income. These adjustments accounted for only around 1% of the variation in individual cost levels (McWilliams, Hsu, and Newhouse 2012). Beneficiaries were also free to switch into and out of private plans at any time. Not surprisingly, there was a high degree of risk selection. A study by the Physician Payment Review Commission (PPRC) in 1996 showed that individuals newly enrolling in private plans had Medicare claims over the prior six months that were just 63% of average claims (PPRC 1996), and the individuals exiting private plans had claims over the subsequent six months that were 160% of average claims.

The more sophisticated scoring system as previously described was implemented between 2003 and 2006, along with rules that created an annual enrollment period to enter or exit a private plan. Since this time, the differences between private plan enrollees and the general Medicare population have decreased. McWilliams, Hsu, and Newhouse (2012) use various health measures to document that before these reforms, private plans enrollees were approximately 15%–20% healthier, while after the reform the difference declined to less than 5%. Interestingly, there is some debate about whether residual (or unscored) risk selection increased after the reforms. Brown et al. (2014) have argued that after the introduction of risk scoring, plans had an increased incentive to enroll individuals who were especially healthy for a given risk score. To assess this, they looked at whether new private plan enrollees had low lagged Medicare claims conditional on their risk scores, and found that this measure of selection increased after 2003, although the results of Newhouse et al. (2015) suggest that this result may be somewhat sensitive to the choice of sample period and set of beneficiaries being studied.

3.2. The Quality of Risk Adjustment

One of the challenges in risk adjustment is that predictive models are imperfect. Medicare’s current model predicts only around 11%–12% of the variation in health spending (measured in terms of R^2). Figure 3 illustrates its forecast accuracy by plotting




FIGURE 3. The predictive power of risk scores in Medicare Advantage. This figure is based on administrative data on aged Medicare beneficiaries in 2010 ($N = 20,491,392$), and plots the realized healthcare expenditure against the beneficiaries' risk score. To construct the figure, we split beneficiaries by their risk score bin (of width 0.1), and compute the average expenditure and the percentiles within each bin.

realized spending against risk scores for all aged beneficiaries in public Medicare in 2010. On average, the risk scores are successful—the average beneficiary with risk score of 2 indeed had about 1.85 times more realized spending than an average beneficiary. Yet there is very large variation in realized costs for a given risk score. Some of this residual variation is unavoidable because it represents unpredictable health shocks. Yet, richer modeling probably could double predictive power, to an R^2 of around 25% (Newhouse et al. 1989). The fact that all available information is not incorporated potentially opens the door to residual risk selection.

Why is information not fully used? One explanation is that risk-scoring systems face practical limitations. The Medicare system only utilizes the previous twelve months of claims. This allows beneficiaries to be incorporated quickly into the scoring system (in their second year), but throws out useful information because disease conditions may not be captured every year. As an illustration, there were around five million Medicare beneficiaries classified as having diabetes in 2006. Of these, approximately 13% were no longer classified as having diabetes in 2007 because a physician had not recorded the condition during the prior twelve months. A longer claims history would limit the problem of chronic conditions disappearing because

they are not recorded each year. However, it might require an alternative model for individuals in their first few years of the program.

A second practical limitation is that risk-scoring systems may need to be explained and justified to policy makers, creating a constraint that models be interpretable. For instance, Medicare risk scoring relies on disease conditions rather than histories of healthcare utilization, and the model includes HCCs in an additive way without interactions. This way, model coefficients can be assigned a “sensible” interpretation—the contribution of the health condition to an individual’s expected cost. The model already is fairly complicated, but arguably richer interactions or nonlinearities, or the use of complex claims histories, would lead to model that was (even more) opaque to policy makers and market participants.

A third issue is the degree to which providers and insurers can affect or manipulate risk scores. In the Medicare system, it is well understood that enrollees in private plans tend to have their health conditions coded more thoroughly than public Medicare beneficiaries (GAO 2013; Geruso and Layton 2014). Risk scores tend to increase faster for private plan enrollees, and chronic conditions are less likely to disappear. This is hardly surprising. Private plans are compensated based on the scores of their enrollees, so if a diabetes diagnosis is not reported, the insurer leaves money on the table. Some analysts have argued that this leads to plans being overcompensated, because for any given risk score, private plan enrollees are healthier than the public Medicare enrollees whose costs are used to set benchmark reimbursement rates. Since 2010, Medicare has addressed this in a somewhat ad hoc way by reducing the risk scores of all private plan enrollees (originally by 3.1%) to make their scores more comparable with public Medicare.

Finally, a fourth issue with current risk-scoring systems is that they do not account for individual healthcare costs potentially varying across plans. For instance, some insurers may have an advantage in controlling costs for certain individuals or certain diseases, but not others (Bundorf, Levin, and Mahoney 2012). In this case, the assumption that individual costs can be summarized by a single risk multiplier r is wrong. Einav et al. (2015) provide a stark example in the context of Medicare Part D. In that setting, individuals differ both in their health and in the way they respond to copayments, and the plans have different copayment rules. However, the risk scores by design predict healthcare spending for a particular assignment of individuals to plans. It follows that even if the scores are perfect predictor in sample, they may be poor predictors if people switch plans. Einav et al. (2015) show how this creates an incentive for plans to selectively enroll people who will be relatively cheap in their plan.

3.3. Risk Adjustment and Incentives

An important lesson is that risk scoring is not simply a matter of finding an accurate predictive model. Rather, the crucial issue is to design a system that beneficially affects plan incentives. The goal of risk scoring is to mitigate the natural incentive that arises for plans to target healthy enrollees, but without creating new distortions where plans

try to manipulate the scores, or select only enrollees who are healthier than their score indicates. Statistical prediction is part of this because it may limit the potential for residual risk selection, but far from all of it.

We have focused so far mainly on incentive problems with ex-ante risk scoring. If anything, the incentive issues are even more clear with ex-post risk adjustment. If, on the margin, an insurer is paid an extra dollar for every dollar it spends on an insuree, the insurer has little incentive to worry about cost. If it is the government paying the extra dollar, it then becomes unclear why the insurer is involved at all, rather than the government paying the provider directly. Of course, ex-post risk adjustment also can be based on health conditions rather than costs. However, this too can generate problematic incentives, because insurance companies may have an incentive to produce claims that increase their population risk score and payments. This same issue can occur when risk adjustment is based on past claims, but the effect is less immediate.

Incentive issues are also important in thinking about what type of information is used in risk scoring and how it is collected. The Medicare system, because it relies on disease diagnoses reported by physicians, is a process that involves considerable effort and discretion. In other industries where predictive scoring is prevalent, firms often try to base predictions on information that is collected in a more automated way. For instance, Google's advertising system scores advertisers based on how often users click on their ads. This is information that Google collects directly. Consumer credit scores generally rely on financial histories, where again the individual reporting may involve less discretion, and perhaps more automation, than in healthcare.

Another conceptual incentive issue is that risk-scoring inputs often reflect health conditions that are affected by health behavior. Suppose that smoking leads to pulmonary disease, and pulmonary disease is predictive of high medical spending. Then there is a sense in which risk scoring reduces the incentive for health plans to discourage smoking. That is, it removes from health plans their financial exposure to the long-term health outcomes of their enrollees. To our knowledge, this type of incentive issue has attracted very little attention, although we will return in Section 5 to some of the ways in which improved data on health conditions and health behavior might be used in creating market incentives.

As predictive modeling improves and risk scoring becomes more sophisticated, it will be important to think of these models not simply as statistical predictive tools, but rather as market design instruments. It may turn out that maximizing predictive power is the approach that leads to optimal market outcomes, but it is probably more likely—because of heterogeneous contracts, or the potential for manipulation, or the desire to provide better long-run incentives—that statistical and economic objectives will not be perfectly aligned.

4. Market Power and Imperfect Competition

We now turn our attention to the issue of market power and a lack of competitive incentives in health insurance markets. Market power concerns arise because health

insurance markets tend to be highly concentrated, and because individuals choosing between plans may not be very price-sensitive. As a result, competitive forces may not pressure plans to bring down their prices or make investments that would improve cost-efficiency. The Medicare Advantage market again provides a useful case study to illustrate these concerns.

4.1. The Concentrated Market Structure of Private Plans

To understand why health insurance markets are concentrated, we can consider the costs involved in setting up and offering a health insurance plan. The main variable costs arise in processing claims and paying for healthcare services. In addition, there are fixed costs involved in establishing and managing a network of physicians and hospitals, negotiating prices and payments with these providers, and complying with various regulatory standards and requirements.

The market structure of Medicare private plans, and the way it has evolved over time, suggests that these fixed costs may be substantial. In Medicare Advantage, HMO and preferred provider organization (PPO) plans with limited provider networks attract the vast majority of enrollees. However, insurers also offer private fee-for-service (PFFS) plans. These plans have a fee-for-service structure that is closer to Medicare public insurance. Historically, enrollees in PFFS plans could see any Medicare provider, and the plans could reimburse providers according to Medicare service prices. The PFFS plans therefore could avoid many of the fixed costs associated with establishing and maintaining a provider networks.

During the last decade, PFFS plans became very prevalent. Between 2006 and 2011, the average US county had 18 different plans from which Medicare beneficiaries could choose. Of these, 14 were PFFS plans while only 4 were HMO or PPO plans. The PFFS plans were not just prevalent but often very small, suggesting that their fixed costs were indeed low. In 2010, Medicare began to impose a network requirement on PFFS plans, at which point the number of these plans began to drop precipitously (MedPAC 2014). By 2014, the number of Medicare plans per county had dropped to 10, and a PFFS plan was only available in around half of US counties.

This evolution is consistent with the argument that building a provider network, and negotiating with these providers, is relatively costly. These same costs may also discourage the entry of insurers offering new HMO and PPO plans. Not only do local markets tend to have a limited number of these plans, they often are operated by the same insurers. This combination results in a small number of insurers tending to dominate each local market. We noted earlier that the top two insurers in each county average a 85% combined market share. The average for the top three insurers is about 95%. Concentration is somewhat lower in urban markets, but it is fairly high everywhere (Frakt, Pizer, and Feldman 2012).

Interestingly, there is less concentration at the national level, suggesting that many fixed costs are at the regional rather than national level. The national market leaders United Health Group and Humana have, respectively, 19% and 16% of national enrollees (Curto et al. 2015). Moreover, while these insurers operate nationally, many

others operate regionally. For example, Kaiser Permanente has a 40% market share in California, but operates in only six states. Its experience is consistent with regional scale being quite important. Kaiser's provider network and insurance company are integrated, a model that would seem to rely heavily on scale, but despite being very successful in its core regions, Kaiser has struggled in attempts to grow nationally.

4.2. Competitive Incentives for Plans

To what extent does market power curtail price competition? In Curto et al. (2015), we studied the competitive incentives of Medicare private plans by looking at how plan enrollment, and plan profits, respond to plan bids. To think about this, it is useful to start with a simple model.

Consider a plan that must decide on its bid b , knowing that its local benchmark rate is B . If the plan bids $b < B$, it can offer additional benefits to consumers with actuarial value $B - b$. Assuming that consumers value these benefits, the plan's enrollment will be greater if the plan submits a lower bid. However, the plan's revenue per enrollee will be lower, because a plan that bids b receives rb to cover an individual with risk r . Let us assume that the plan's costs also scale with r , so that its net profit on an enrollee with risk r is $r(b - c)$, where c is the cost (net of the additional benefits) of an individual with risk score $r = 1$.

Now, if there are several plans $1, \dots, J$, a given plan's demand also will depend on the bids of rival plans, or more precisely on the *excess bid* of each rival, $p_k = b_k - B$, which determines the extra benefits the rival plan offers. Suppose we can write plan j 's enrollment—it is convenient to measure this in risk units rather than headcount—as $Q_j(p_j, p_{-j})$. Then we can write plan j 's bidding problem as

$$\max_{p_j} (p_j + B - c_j) Q_j(p_j, p_{-j}). \quad (1)$$

This is an entirely standard pricing problem. We think of the excess bid p_j as the "price", c_j as marginal cost, and B as a subsidy paid by the government for each enrollee. The first-order condition that characterizes optimal bidding is

$$p_j = c_j - B + \left(\frac{d \ln Q_j}{dp_j} \right)^{-1} \quad \text{or equivalently} \quad b_j = c_j + \left(\frac{d \ln Q_j}{db_j} \right)^{-1}, \quad (2)$$

where the final "markup" term is the same in the two expressions. The formula can also be modified to account for insurers that operate multiple plans, but the basic idea remains the same.

It follows that a key determinant in pricing is the sensitivity of enrollment to plan benefits. The extent to which consumers are responsive to plans offering more generous benefits (through lower bids and larger rebates) will matter, as will the pricing of rival plans. If competition is very intense, small changes in a plan's benefits will be associated with large (percentage) changes in enrollment. Margins will be narrow

and bids close to costs.¹⁰ If competition is less intense, there may be a large margin between costs and bids. However, one needs empirical evidence to distinguish these cases.

One possible empirical approach is to look at how bids respond to changes in the Medicare benchmark rates (Song, Landrum, and Chernew 2012, 2013; Cabral, Geruso, and Mahoney 2014; Duggan, Starc, and Vabson 2014; Curto et al. 2015). If markets are competitive, then (exogenous) increases in B will have little or no effect on bids. Recent papers, however, have tended to find pass-through rates around 0.5. That is, for every additional dollar subsidy, plans raise their bids and profit margins by around 50 cents. This is suggestive of imperfect competition, although one difficulty with looking at pass-through rates is that they do not per se identify competitive incentives and markups.

A second approach is to look directly at how plan enrollments respond to plan bids. We do this in Curto et al. (2015), where we estimate that enrollments do not in fact appear to be very sensitive to changes in plan bids. In particular, a \$10 reduction in the monthly bid increases enrollment (measured in risk units) only by around 10%. Using these demand estimates in the context of equation (2) suggests that if firms bid optimally, they should be asking for Medicare payments that are around 10% to 25% over their costs. Operating margins of that magnitude seem roughly consistent with the publicly reported numbers of large insurers that we discussed earlier.

4.3. Competition and Market Design

The fact that insurance plans appear to have market power, and this market power encourages bids well above cost, does not necessarily imply that managed competition is a failure. To assess its overall effectiveness, one needs to compare managed competition outcomes to those of a plausible alternative. In Curto et al. (2015), we tried to do this by comparing Medicare Advantage outcomes to public Medicare. In particular, we estimated the consumer surplus associated with private plan enrollments (relative to public Medicare, private plans offer better financial coverage but restrict provider access), the effect of the private plan enrollments on taxpayers, and the insurer profits. To compute profits, we ideally would have direct measures of plan costs. As these were not available, we used a standard trick often employed in studies of imperfect competition, and assumed that insurers set bids optimally so that we could back out the costs consistent with the observed bids.

Our estimates are relatively encouraging about the prospects for private insurance provision but relatively discouraging about the efficacy of competition. Looking at the period from 2006 to 2011, we estimated that for the average private plan enrollment, it would have cost the government around \$660 a month to provide the individual with public Medicare insurance. The private plan's cost is estimated at only \$585, for a

10. For instance, in 2010 and 2011, both Aetna and United reported that their ratios of medical claims payments to plan premiums were in the range of 80%–83%, which loosely speaking would correspond to a 17%–20% margin in our bidding model (see Aetna and United Healthcare Group annual reports 2012).

saving of \$75 per enrollee-month, or around 11%. However, the average plan bid is about \$680, and factoring in the rebate, taxpayers ended up paying \$755. That is almost 15% more than they would have paid to provide government administered insurance. The additional spending goes partly to private plan enrollees. We calculated that they obtained consumer surplus of \$50 per enrollee-month (i.e., \$75 in benefits from the rebate minus approximately \$25 of disutility due to the limited network). However, insurers appear to have been the big winners during this time period, obtaining almost \$100 a month in (variable) profits from an average enrollment.

One can interpret these results in different ways. On one hand, a total surplus calculation is encouraging. It suggests that the presence of private plans generates overall positive gains from trade. The gains result from private plans having lower costs, around \$75 per enrollee-month for the population that enrolls. While our cost estimates are indirect, these savings seem to be easily enough to compensate the average enrollee for the estimated \$25 disutility of having restricted provider access. On the other hand, taxpayers captured none of the overall surplus creation. Indeed, taxpayers appear to have provided an extra \$95 subsidy for an average private plan enrollment, with a large fraction of this number flowing to insurers rather than Medicare beneficiaries.

A further finding in Curto et al. (2015) is that under the Medicare program rules, increased competition per se would not have helped taxpayers very much. The reason is that when plans reduce their bids, most of the “savings”—in fact 75%—are channeled to beneficiaries. In fact, we calculated that if during 2006–2011, plans had set their bids exactly equal to cost with no markup, taxpayer spending would have decreased only \$24 per enrollee-month. Of course, Medicare beneficiaries would have enjoyed a great deal more surplus. One conclusion we draw from this is that while market structure and competitive conditions are important, market design also has a crucial role in managed competition.

This raises a natural question, both in the Medicare context and more broadly. What market design choices will benefit consumers and limit the costs of market sponsorship, namely taxpayer costs? Within the Medicare program design, the government has two immediate levers: the benchmark reimbursement rates, and the rebate formula. Changes in the rebate formula could involve taxpayers retaining a larger share of bid savings, or moving toward a premium support model in which the government pays a fixed subsidy for each enrollment. To some extent, different combinations of benchmark and rebate rates will lead to a tradeoff—all else equal, reducing taxpayer cost is likely to reduce private plan enrollment because subsidies for enrollment are lower. However, not all changes are equal. For example, shrinking the rebate seems like a natural way to reduce taxpayer costs, but it also makes plan enrollment less sensitive to plan bids, softening competition (Curto et al. 2015). In contrast, reducing benchmark reimbursement rates is a very direct way to reduce taxpayer cost. Indeed, Medicare has begun a process of reducing these rates, and as of 2014, the taxpayer costs of a private plan enrollment had fallen to around 106% of the cost of providing public Medicare (MedPAC 2014).

More broadly, one might ask what other market design options are available. The structure of state insurance exchanges is in certain ways quite different from Medicare Advantage. In the exchanges, insurers can offer plans with different levels of generosity (gold, silver, etc.) and set the plan premiums. Lower-income consumers receive a subsidy, but the subsidy depends on the overall set of plan premiums in the market, and not on which plan the consumer chooses.¹¹ An interesting question is how this type of market design, which gives insurers more flexibility in separately setting plan generosity and premiums, compares to Medicare Advantage, and how both compare to the mechanisms used by large employers to provide their employees with health insurance options.

If one were to tackle the larger market design question, one observation that would seem to be important is that local health insurance markets can be highly heterogeneous. For example, in the Medicare market, there can be a wide range in local market structure and in the rate at which beneficiaries enroll in private plans. As of 2011, a tenth of Medicare beneficiaries were in markets where private plan enrollment was under 3%, and another tenth were in markets where enrollment exceeded 40%. Moreover, the market share of the largest insurer in 2011 ranged from 30% to 80% across the 10th and 90th percentiles of the distribution, with the Hirschman–Herfindahl Index varying from 20% to 65%. Markets also vary dramatically in the apparent cost of private insurance relative to public insurance (Figure 4). It seems likely that employer-sponsored and individual health insurance markets have similarly wide variation. This suggests that in assessing designs for managed competition, one criterion should be that the design works, or can be tailored to work, in a wide range of competitive and cost environments.

5. Some Open Questions

There are still many unknowns when it comes to healthcare competition. We have focused our discussion on the problems of risk selection and market power, which we see as particularly central issues. We would like to end, however, by considering an expanded set of issues that we think merit attention, and hopefully will receive it in coming years.

5.1. *Can We Rely on Consumers?*

Managed competition relies on consumers making informed and price-sensitive choices across competing insurance plans. One conclusion we draw from our work on Medicare Advantage, however, is that consumer demand is not very sensitive to plan bids. In the Medicare context, the issue is complicated by the fact when an insurer reduces its bid, enrollees do not see a dollar reduction in their premium, but a

11. More specifically, subsidies in a given area are tied to the second-lowest price offered by a “silver” plan, but eligible individuals receive their subsidy whichever plan they choose.

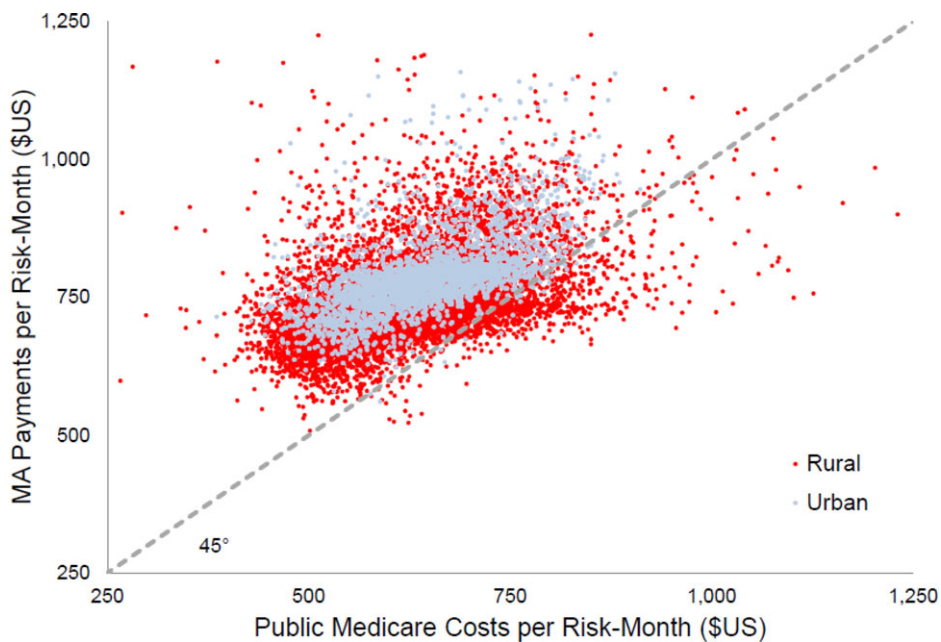


FIGURE 4. Heterogeneity across markets in Medicare Advantage. This figure presents average, risk-adjusted Medicare Advantage payments (from the government to private plans) against average, risk-adjusted healthcare expenditure for beneficiaries covered by public Medicare. Each point represents a market (county–year pair). Data cover all US counties during 2006–2010.

75 cent improvement in plan generosity. However, a fairly wide set of studies have found that consumers may not be highly price-sensitive in choosing health plans, whether because they view plans as differentiated and value particular plan characteristics (Bundorf, Levin, and Mahoney 2012), or because behavioral factors such as choice inertia (Handel 2013), inattention (Ho, Hogan, and Scott Morton 2015), or optimization mistakes (Abaluck and Gruber 2011) lead to low price elasticity.

In thinking about consumer choice, it is important to recognize two roles it can play. One is to generate price sensitivity, so that insurers have an incentive to set lower prices. The other is to sort consumers to higher-quality, or more appropriate plans. The latter improves allocative efficiency and rewards insurers for valuable plan design and administration. Creative market design choices can help with the former. For example, assigning specific default choices to inattentive or inertial consumers, or providing extra rewards for firms that set low enough prices (as in the health insurance exchange in Massachusetts) can generate higher demand elasticity even if consumers on their own would not respond sharply to price reductions. The latter problem of getting consumers to sort efficiently across plans is more challenging because it may require active choices, and informed consumers.

One reason the consumer information is a difficult issue is that health plans provide coverage that bundles many contingencies. At the time of plan choice, most consumers

do not yet know what type of healthcare they will need, and what aspects of healthcare they will value most. Is it the quality of the physician? The proximity of a provider to one's residence? The ability to obtain second opinions? Or the out-of-pocket cost of services? It is only after health problems occur that consumers become informed about what they need and what they value. Yet, at that point consumers have little choice, at least when it comes to their insurance plan. Perhaps one aspect of market design should be to allow people who find themselves in a mismatched plan some flexibility to make changes, provided this could be done in a way that does not create adverse selection problems. Alternatively, it may be useful to think about how to focus consumers' attention on information about plans that is specifically targeted to their needs, and their likely needs.

5.2. The Scope of Insurance Contracts

A somewhat remarkable aspect of healthcare markets is that the transactable unit is almost always the person-year. That is, in all of the markets we identified at the outset, consumers enroll with an insurer on an annual basis, and the insurer receives a capitated payment for covering the person for that year. It is interesting to consider whether this is the "right" level of transaction, and one can imagine arguments in a variety of directions.

One argument is that this transactable unit is too narrow. Health, in many cases, is a slow-moving process, and the tradeoffs associated with certain aspects of care are often related to a much longer time horizon. For example, preventive services and lifestyle advice focus on preventing longer-run risks, and certain procedures, such as hip or knee replacement, provide quality-of-life benefits that extend much longer than a single year. Focusing on a single year may therefore put too much weight on short-run incentives.¹² Sometimes the justification of shorter contracts is practicality—some insurers may leave the market over longer horizons, or in employer-provided markets workers come and go—but especially when the government sponsors the market, these types of objections to longer-term contracts seem less relevant.

However, one also can argue that the transactable unit is too broad. For example, it is not obvious that all healthcare services should be bundled into a single insurance contract. While externalities across related treatments make some bundling a natural solution, there are many procedures and services that are reasonably separable from overall health. Just as many insurance products separate dental, vision, and pharmaceutical drug insurance from basic healthcare, one could imagine separating other aspects of healthcare, such as hip and knee replacements, or labor and delivery. Bundling has obvious benefits (Porter and Kaplan 2014), but also has costs. For example, it may limit the ability of providers to enjoy economies of scale at the level

12. Time tradeoffs can be relevant for consumers as well as insurers. Even if consumers are able to accurately weigh the costs and benefits of different services, the simple fact that their insurance deductible renews each year can affect the services they obtain (Cabral 2013; Einav, Finkelstein, and Schrimpf 2015).

of individual services, or with limited network insurance plans, to allow consumers to mix-and-match among providers depending on their needs.

Some version of unbundling already occurs at the upstream level, where physicians and providers may contract with insurers to provide a certain scope of covered services at capitated rates. An interesting question is why such unbundling does not make its way through the vertical chain all the way to consumers.

5.3. Health versus Healthcare

Throughout this paper, and in most healthcare discussions, the focus is on healthcare services rather than health outcomes. This is primarily due to data limitations. Current measurements do not capture individual's health status very well. Insurance contracts cover healthcare services rather than health outcomes; most data sets come in the form of healthcare claims rather than direct health measures; and market incentives are not generally focused on health per se.

Going forward, it seems plausible that better, more objective, and more frequent measures of health will gradually become the norm. Once this is so, it may enable contractual relationships that are based on health conditions and health outcomes rather than healthcare services. It is easy to imagine rewarding insurance plans or providers for achieving good health outcomes, but one can also envision insurers rewarding individuals to take better care of themselves, or to engage in healthy activities. There is also a concern that better health data, such as genetic test results, might lead to new forms of risk selection. It is somewhat early to gauge how these possibilities might develop, but at least some of them are likely to become more salient in the next decade or two.

5.4. The Goals of Competition

When asked what healthcare competition should aim to achieve, policy makers and health economists often respond that the goal is better healthcare at a cheaper price. Of course, this is easy to say and harder to achieve, partly because of the difficulty in assessing what better even means, and partly because one cannot always receive more while paying less.

As a result, many managed competition programs aim for somewhat less lofty goals. One is providing the "same" health insurance at lower cost, while the other is providing "better" health insurance at the same cost. The Medicare Part D program is an example of the former. Private plans procure the same pre-packaged pharmaceutical drugs that could have otherwise been procured by the government. The hope is that private insurers can do this in a more cost-effective way. In contrast, Medicare Advantage has aimed (for a long time, and so far unsuccessfully) to keep taxpayer costs the same as under public Medicare, while private plans distinguish themselves by offering more generous insurance benefits.

A third potential objective, which perhaps should become part of the discussion, is an objective of achieving cost savings for healthcare that is perhaps worse, but is still

“good enough”. Such an objective could be politically difficult to deliver. Yet, precisely for this reason it could be an objective for which private healthcare markets have a comparative advantage. If certain individuals prefer more basic forms of healthcare in return to more residual money for consumption, there is no reason that such options would not be offered. Of course, this requires some regulation to make sure that “basic” healthcare is sufficient and that consumers are aware of what services they may be giving up.

One way to think about this is to imagine that we could distinguish between “pure” forms of healthcare and healthcare amenities, such as the ability to secure a private hospital room, or to schedule a physical with little to no wait time. One can make a reasonable argument that the government should do its best to provide the former, but that it is less well-suited to mandate, subsidize, or guarantee amenities. In contrast, starting from a basic floor, markets can be a natural way to let people make tradeoffs in deciding whether they want the “economy” product, or instead to “fly business”, which, judging by the comparative costs of US healthcare, seems to be quite common in our current system.

6. Conclusions

Our starting point in this paper was the possibility for structural changes that might contain the growth of healthcare costs and improve efficiency. Although there are many views of what these changes might be, promoting private sector competition is clearly among the leading ideas. We have tried in this paper to provide some sense of the challenges, and the possibilities, for competitive health insurance markets. In doing so, we have argued that market design choices are critical in addressing two of the main challenges, mitigating risk selection and promoting insurer competition, and we have also pointed out that the premise on which arguments for managed competition rest—namely that private insurers can improve on efficiency relative to the public sector—still calls out for more supporting evidence.

Over a longer horizon, there is also an important dimension that we did not consider, which is the argument that competitive markets will spur new technological or organizational innovations. For example, one commonly hears the view that information technology—perhaps in the form of electronic medical records or vastly improved population data—could transform the healthcare sector. Is such a transformation more likely to happen, or to happen quickly, with market competition? This would be yet another interesting issue for future work to consider.

References

- Abaluck, Jason and Jonathan Gruber (2011). “Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program.” *American Economic Review*, 101(4), 1180–1210.

- Agency for Healthcare Research and Quality (AHRQ), Center for Financing, Access and Cost Trends (2012). *Medical Expenditure Panel Survey* (Insurance Component).
- Akerlof, George (1970). "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*, 84, 488–500.
- Brown, Jason, and 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.
- Bundorf, M. Kate, Jonathan Levin, and Neale Mahoney (2012). "Pricing and Welfare in Health Plan Choice." *American Economic Review*, 102(7), 3214–3248.
- Cabral, Marika (2013). "Claim Timing and Ex Post Adverse Selection." Working paper, UT Austin.
- 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.
- Casey, Michelle (1998). "Serving Rural Medicare Risk Enrollees: HMOs' Decisions, Experiences, and Future Plans." *Health Care Financing Review*, 20, 73–81.
- Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya (2015). "Can Health Insurance Competition Work? Evidence from Medicare Advantage." NBER Working Paper No. 20818.
- Cutler, David M. and Sarah J. Reber (1998). "Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection." *Quarterly Journal of Economics*, 113, 433–466.
- 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.
- Einav, Liran, Amy Finkelstein, and Mark R. Cullen (2010). "Estimating Welfare in Insurance Markets using Variation in Prices." *Quarterly Journal of Economics*, 125, 877–921.
- Einav, Liran, Amy Finkelstein, and Jonathan Levin (2010). "Beyond Testing: Empirical Models of Insurance Markets." *Annual Review of Economics*, 2, 311–336.
- Einav, Liran, Amy Finkelstein, and Paul Schrimpf (2015). "The Response of Drug Expenditure to Nonlinear Contract Design: Evidence from Medicare Part D." *Quarterly Journal of Economics*, 130, 841–899.
- Einav, Liran, Amy Finkelstein, Ray Kluender, and Paul Schrimpf (2015). "Beyond Statistics: The Economic Content of Risk Scores." Working paper, Stanford University.
- Enthoven, Alain C. (1993). "The History and Principles of Managed Competition." *Health Affairs*, 12, 24–48.
- Fisher, Elliott, David Goodman, Jonathan Skinner, and Kristen Bronner (2009). "Health Care Spending, Quality, and Outcomes." The Dartmouth Institute for Health Policy and Clinical Practice, <http://www.dartmouthatlas.org>.
- Frakt, Austin, Steven D. Pizer, and Roger Feldman (2012). "The Effects of Market Structure and Payment Rate on the Entry of Private Health Plans into the Medicare Market." *Inquiry*, 49, 15–36.
- Fuchs, Victor R. (2014). "Why Do Other Rich Nations Spend So Much Less on Healthcare?" *The Atlantic*, July 23, 2014.
- General Accounting Office (GAO) (2013). "Medicare Advantage: Substantial Excess Payments Underscore Need for CMS to Improve Accuracy of Risk Score Adjustments." Report to Congress, GAO-13-206.
- Garber, Alan M. and Jonathan Skinner (2008). "Is American Health Care Uniquely Inefficient?" *Journal of Economic Perspectives*, 22, 27–50.
- Geruso, Michael and Timothy Layton (2014). "Risk Selection, Risk Adjustment, and Manipulable Medical Coding: Evidence from Medicare." Working paper, UT Austin.
- 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.
- Hall, Robert E. and Charles I. Jones (2007). "The Value of Life and the Rise in Health Spending." *Quarterly Journal of Economics*, 122, 39–72.
- Handel, Benjamin R. (2013). "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review*, 103(7), 2643–2682.
- Handel, Ben, Igal Hendel, and Michael D. Whinston (2015). "Equilibria in Health Exchanges: Adverse Selection vs. Reclassification Risk." *Econometrica*, forthcoming.

- Hendren, Nathaniel (2013). "Private Information and Insurance Rejections." *Econometrica*, 81, 1713–1762.
- Ho, Kate, Joe Hogan, and Fiona Scott Morton (2015). "The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program." Working paper, Columbia University
- 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, 2609–2617.
- MaCurdy, Thomas, Jay Bhattacharya, Daniella Perlroth, Jason Shafrin, Anita Au-Yeung, Hani Bashour, Camille Chicklis, Kennan Cronen, Brandy Lipton, Shahin Saneinejad, Elen Shrestha, and 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, 289–332.
- McWilliams, J Michael, John Hsu, and Joseph P. Newhouse (2012). "New Risk-Adjustment System was Associated with Reduced Favorable Selection in Medicare Advantage." *Health Affairs*, 31, 2630–2640.
- Medicare Payment Advisory Commission (MedPAC) (2007–2014). "A Data Book: Health Care Spending and the Medicare Program." Chapter 9 on Medicare Advantage.
- Newhouse, Joseph P., Willard G. Manning, Emmett B. Keeler, and Elizabeth M. Sloss (1989). "Adjusting Capitation Rates Using Objective Health Measures and Prior Utilization." *Health Care Financing Review*, 10, 41–54.
- Newhouse, Joseph P., Mary Price, J. Michael McWilliams, John Hsu, and Thomas McGuire (2015). "How Much Favorable Selection is Left in Medicare Advantage?" *American Journal of Health Economics*, 1(1), 1–26.
- OECD (2013). *Health at a Glance 2013: OECD Indicators*. OECD Publishing.
- Physician Payment Review Commission (PPRC) (1996). Annual Report to Congress. Washington, DC.
- Porter, Michael E. and Robert S. Kaplan (2014). "How Should We Pay for Health Care?" Harvard Business School Working Paper No. 15-041.
- Song, Zirui, Mary Beth Landrum, and Michael E. Chernew (2012). "Competitive Bidding in Medicare: Who Benefits From Competition?" *American Journal of Managed Care*, 18, 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, 1301–1312.