Journal of Health Economics 55 (2017) 153–167

Contents lists available at ScienceDirect

Journal of Health Economics

journal homepage: www.elsevier.com/locate/econbase

Do health insurers innovate? Evidence from the anatomy of physician payments[☆]

Jeffrey Clemens[a,b], Joshua D. Gottlieb[b,c,*], Tímea Laura Molnár[c]

[a] UCSD, USA
[b] NBER, USA
[c] UBC, Canada

ARTICLE INFO

Article history:
Received 26 December 2016
Received in revised form 3 July 2017
Accepted 4 July 2017
Available online 13 July 2017

Keywords:
Health insurance
Physician payments
Medicare
Payment reform

ABSTRACT

One of private health insurers' main roles in the United States is to negotiate physician payment rates on their beneficiaries' behalf. We show that these rates are often set in reference to a government benchmark, and ask how often private insurers customize their fee schedules away from this default. We exploit changes in Medicare's payments and dramatic bunching in markups over Medicare's rates to address this question. Although Medicare's rates are influential, 25 percent of physician services in our data, representing 45 percent of covered spending, deviate from the benchmark. Heterogeneity in the pervasiveness and direction of deviations suggests that the private market coordinates around Medicare's pricing for simplicity but abandons it when sufficient value is at stake.

1. Introduction

Health insurers have a powerful ability to shape the efficiency of health care delivery. Insurers straddle the relationship between patients and medical providers, and enter into contracts with both sides of the market. Consumers or employers purchase insurance plans whose copayments and deductibles influence subsequent demand for care. At the same time, contracts with physicians and hospitals govern how these providers will be compensated for treating insured patients, and hence the caregivers' financial incentives.

The literature on optimal consumer cost-sharing is long and well-developed (Feldstein, 1973; Besley, 1988). Only recently, however, has an empirical literature begun to explore how private insurers set copayments in practice. Einav et al. (2016) show that private insurers provide more risk-protection for drugs subject to less moral hazard. They contrast this with public insurance plans, which offer relatively uniform coverage with regards to cost-sharing. Starc and Town (2015) show that insurers responsible for patients' non-pharmaceutical spending provide more generous coverage for drugs that can keep people out of the hospital.

We investigate whether insurers apply a similar logic to the payment schedules they negotiate with providers. To what extent do private insurance carriers adopt Medicare's cost-based approach to physician payment? Looking at the flip side of this question, we analyze the extent to which private insurers customize their physician reimbursements relative to Medicare's industry standard.

Despite recent high-level changes in the U.S. health insurance market, the incentive structure through which physicians are paid remains predominantly "fee for service."[1] A physician's income depends on the quantity and intensity of the treatment she provides – even when part of a larger managed care plan or "accountable

☆ We are extremely grateful to Luisa Franzini, Cecilia Ganduglia-Cazaban, Osama Mikhail, and the UTSPH/BCBSTX Payment Systems and Policies Research Program at the University of Texas School of Public Health for data access, and for their extensive assistance in navigating the BCBSTX claims data. Clemens and Gottlieb thank the Stanford Institute for Economic Policy Research and the Federal Reserve Bank of San Francisco for their hospitality. We thank SSHRC, Jon Skinner and the Dartmouth Institute for support, and Victor Saldarriaga and Andrew Vogt for excellent research assistance. Finally, we thank Leila Agha, Liran Einav, Jeff Emerson, Amy Finkelstein, John Friedman, Matt Gentzkow, R.B. Harris, David Laibson, Neale Mahoney, Matt Panhans, Maria Polyakova, Jon Skinner, Ashley Swanson, numerous seminar participants, three referees and the editor, Luigi Siciliani, for helpful comments.

* Corresponding author at: UBC, Canada.
E-mail addresses: jeffclemens@ucsd.edu (J. Clemens), Joshua.Gottlieb@ubc.ca (J.D. Gottlieb), molnar.timea.laura@alumni.ubc.ca (T.L. Molnár).

[1] Data from the 2004–2005 Community Tracking Study (CSHSC 2006, 46) show that 52 percent of physicians earn zero revenue from capitated contracts, and 79 percent earn less than a quarter of their revenue from such contracts.

http://dx.doi.org/10.1016/j.jhealeco.2017.07.001
0167-6296/

care organization" (Zuvekas and Cohen, 2016).[2] A growing body of evidence finds that these high-powered incentives for incremental care provision help drive the level and composition of medical spending (McClellan, 2011).

The structure of physician payments is a potentially powerful tool for insurers to encourage more efficient care. Relatively little is known, however, about the extent to which private insurers customize their fee-for-service payments for this purpose. Clemens and Gottlieb (2017) find that private payments rise and decline quite strongly with Medicare's payments, which raises the question of whether private insurers' payments are meaningfully independent from Medicare's rate schedule. In this paper, we shed light on this question using detailed physician payment data from a large insurer.

Medicare compensates physicians and outpatient providers through a detailed fee-for-service pricing system. Physicians submit bills for each instance in which they provide any of 13,000 recognized services. The system assigns each service a certain number of Relative Value Units, which determine the payment. These relative values aim to measure average cost but not medical value. This procurement model thus has little capacity to steer treatment towards effective – let alone cost-effective – care. It has particular difficulty managing the use of capital-intensive diagnostic imaging services, for which average-cost payments significantly exceed providers' marginal costs – as they must in order to facilitate entry.

Private reimbursement arrangements are less transparent than Medicare's. To peer into the black box of these business-to-business contracts, we begin by developing a cross-sectional method for systematically assessing whether payments are benchmarked to Medicare's rate structure. Our first approach involves a classification algorithm motivated by the bunching literature.[3] Using the outpatient claims data of Blue Cross Blue Shield of Texas (BCBS-TX), we begin by computing the ratio of each private payment to the Medicare payment for that service. Within the payments to individual physician groups, the distributions of these ratios reveal spikes that indicate exceptionally common markups. We use these spikes to identify which payments are likely benchmarked to Medicare's relative rate structure.

We complement this cross-sectional method with an analysis of updates to Medicare's structure of relative payments. If the Medicare links we identify are accurate, then payments for Medicare-benchmarked services should update when Medicare's schedule of Relative Value Units is revised. We are able to assess this pass-through at a high frequency by applying institutional knowledge of the exact dates on which BCBS-TX implements Medicare's updates to the relative value scale. The relatively high frequency at which we can conduct our analysis allows us to limit, if not eliminate, concerns about potential confounders including active contract renegotiations and payment changes connected to substantive technological advances. We find that the payments associated with 55 percent of in-network, outpatient spending (and around three quarters of services) are linked to Medicare. These estimates are quite similar to those we obtain using our cross-sectional bunching approach.

We continue our analysis with an effort to understand the circumstances under which payments are more likely, or less likely, to be benchmarked to Medicare's relative rate structure. Deviations from benchmarking exhibit several distinctive patterns. Looking across physician groups, payments to relatively large groups are less tightly benchmarked to Medicare than payments to small groups. Payments for only ten percent of services provided by the smallest firms, representing 20 percent of their spending, deviate from Medicare's relative values. The same is true of 40 percent of services – and two-thirds of spending – from firms with total billing exceeding $1 million per year.

Looking across service categories, payments are more likely to deviate from Medicare's relative values for capital-intensive services, like diagnostic imaging, than for labor-intensive services like standard office visits. Payments for roughly 45 percent of imaging services, but only 15 percent of evaluation and management services, deviate from Medicare's menu. Within imaging, Medicare distinguishes between two types of services: a capital-intensive component for taking the image and a labor-intensive component for interpreting the image. Medicare explicitly amortizes the fixed cost of the imaging equipment into the former. We find that private insurers' payments for interpretation are far less likely to deviate from Medicare rates than payments for taking the image itself. The directions of these deviations reveal that the adjustments narrow likely gaps between marginal costs and Medicare's average-cost payments. We find that payments for labor-intensive services tend to be adjusted up while payments for capital-intensive services are adjusted down.

One plausible interpretation of these findings emphasizes the complexity of the insurer-physician contracting environment. Specifically, there is a tension between gains from fine-tuning payments and costs from making contracts complex. To manage this tension, insurers may draw on Medicare's relative value scale for the purpose of contract simplification, while strategically adapting their contracts where the value is highest. This view is consistent with the heterogeneity we observe: the benefits of fine-tuning payments will tend to be largest within contracts with large physician groups and for the capital-intensive services for which Medicare's average-cost payments deviate most from marginal cost. A complementary explanation is that large firms may use their bargaining power to obtain high service-specific markups rather than, or in addition to, high overall payments per unit of care. The information content of the relative value scale on which Medicare's payments rely can also be interpreted as a knowledge standard and, more generally, as a public good.

Our results are relevant in two broader contexts. Learning how prices are set in health care – a sector comprising 18 percent of the economy – is essential for understanding macroeconomic price-setting dynamics.[4] The service sector in general (Nakamura and Steinsson, 2008), and medical care in particular (Bils and Klenow, 2004), have especially sticky prices. We provide evidence on how this stickiness arises.[5] Consistent with Anderson et al.'s (2015) evidence from retail, the complexity of physician contracting may explain both the long duration of these prices and the public-private linkages we identify.

Public policies' residual influence on private firms is relevant in a wide range of contexts. Outside of the health care context, labor contracts sometimes benchmark wage rates to the statutory minimum.[6] Within the health sector, Medicare has been found to shape aspects of private players' behavior in the pharmaceutical,

[2] Among those same physicians referenced in footnote [1], who earn little from capitated contracts, 65 percent earn more than one quarter of their revenue from managed care (CSHSC 2006, 46). CSHSC Center for Studying Health System Change (1999) reports similar estimates from 1996 to 1997.

[3] Our setting differs from standard bunching applications in that the bunching we observe is not driven by kinks or notches in budget sets (Kleven, 2016). Instead, it results from clustering around reference points.

[4] Clemens et al. (2014, 2016) show how much Medicare price regulation can impact overall inflation.

[5] In particular, our empirical evidence supports price-setting mechanisms with the flavour of Christiano et al. (2005) or Smets and Wouters (2003, 2007).

[6] A publicly posted contract template of the United Food (Commercial Workers Union, The, 2002), for example, includes the requirement that "At no time during the life of this Agreement will any of the bagger/carry-out rates be less than twenty-five ($0.25) cents an hour above the Federal minimum wage."

hospital, and physician marketplaces (Duggan and Scott Morton, 2006; Alpert et al., 2013; White, 2013; Clemens and Gottlieb, 2017). The forces we investigate here differ conceptually from those analyzed in prior work, including Clemens and Gottlieb's (2017) analysis of physician payments. Clemens and Gottlieb (2017) analyze Medicare's influence on physicians' bargaining positions through its effects on their incentives and outside options. The current paper emphasizes the Medicare payment model's role as an industry standard. More specifically, it emphasizes the Medicare payment model's role as a benchmark or default around which contract negotiations coordinate. Our analysis provides insights into the overall pervasiveness of benchmarking against Medicare's relative cost schedule and into the types of contracts in which customization is most prevalent.

We continue in Section 2 by presenting institutional background on price setting in U.S. physician markets. Section 2 concludes with a discussion of several potential explanations for the benchmarking phenomenon we examine. Section 3 introduces the claims data we analyze. Section 4 presents the empirical strategies we implement, while Sections 5 and 6 present our results. Section 7 briefly concludes.

2. Medical pricing institutions

Public and private payments for health care services are set through very different mechanisms. Medicare reimbursements are set based on administrative estimates of the resource costs of providing care, which we describe in Section 2.1. For patients with private health insurance, providers' reimbursements are determined through negotiations between the insurers and providers, which we describe in Section 2.2. Section 2.3 discusses several potentially complementary economic rationales for a link between reimbursement rates across these two segments of the market.

2.1. Medicare price determination[7]

In 1992, Congress established a system of centrally administered prices to reimburse physicians and other outpatient providers. This Resource-Based Relative Value Scale (RBRVS) is a national fee schedule that assigns a fixed number of Relative Value Units (RVUs) to each of 13,000 distinct health care services. Legislation specifies that the RVUs for service j are supposed to measure the resources required to provide that service. Since the costs of intermediate inputs differ across the country, RBRVS incorporates local price adjustments, called the Geographic Adjustment Factor (GAF), to compensate providers for these differences. The payment for service j to a provider in geographic region i is approximately:

$$\begin{aligned}\text{Reimbursement rate}_{i,j,t} &= \text{Conversion Factor}_t \\ \times\text{Geographic Adjustment Factor}_{i,t} & \\ &\times\text{Relative Value Units}_{j,t}.\end{aligned} \tag{1}$$

The "reimbursement rate," a term we use interchangeably with "price," is the amount Medicare pays for this service. The Conversion Factor (CF) is a national scaling factor, usually updated annually.

Variation in Medicare's payment rates is driven primarily by the number of RVUs assigned to a service. This assignment is constant across areas while varying across services. Medicare regularly updates the RVUs assigned to each service, primarily based on input from the American Medical Association, using the formal federal rule-making process. These updates are intended to account for technological and regulatory changes that alter a service's resource intensity. We exploit these changes in the empirical strategy described in Section 4.

2.2. Private sector price setting

U.S. private sector health care prices are set through negotiations between providers and private insurers.[8] The details of these negotiations are not transparent, and our limited knowledge about private sector prices comes from claims data that reveal the reimbursements paid once care is provided.[9] A common feature of physician contracts, central to both our theoretical and empirical analyses, is a form of benchmarking to Medicare.

Practitioners emphasize that Medicare's administrative pricing menu features prominently in private insurers' contracts. Newsletters that insurers distribute to participating providers frequently draw explicit links between Medicare's fee schedule and the insurer's maximum allowable charges. For example, reimbursement rates might be linked to Medicare by default unless the contract specifies otherwise. But the relative value scale does not determine an absolute price level. As in Medicare, computing private reimbursements requires multiplying RVUs by a dollar scaling factor. Practitioners describe physician contracts as involving negotiations over markups relative to Medicare, combined with payments for particular services or service bundles (Nandedkar, 2013; Gesme and Wiseman, 2010; Mertz, 2004). Our empirical work examines the pervasiveness of this benchmarking phenomenon and the circumstances under which customization occurs.

2.3. Potential rationales

Why might contracts between physicians and private insurers use Medicare's relative rate structure as a benchmark? We consider several explanations, which are broadly complementary. A first explanation is that benchmarking to Medicare's relative rate structure enables insurers and physicians to greatly simplify their contracts. A fully benchmarked contract requires negotiating over a single parameter – the markup or scaling factor. Alternative contract structures could require negotiating payments for hundreds or thousands of distinct billing codes. Medicare's payment model may serve as an industry-standard benchmark with which all parties are familiar. In Appendix A, we present a formal model of this idea, which generates predictions that our empirical analysis supports.

A second, strongly complementary explanation involves the Medicare schedule's informational content. By design, Medicare's payment model contains substantial information about the relative costs of providing physicians' services. If average-cost reimbursement is more or less what insurers desire to implement, Medicare's payment model provides useful information for private insurers to adopt. Put differently, Medicare's relative cost estimates can be interpreted as a public good. Although they may fail to reflect variations in local cost structures, the expense to insurers of independently calculating these costs may be high.

A third possibility is that providing care for Medicare beneficiaries represents physicians' primary outside option when they negotiate with private insurers (Clemens and Gottlieb, 2017). Because Medicare accounts for a large share of the market,

[7] This section draws from Clemens and Gottlieb (2014).

[8] In rare exceptions, such as in Maryland, the state government determines all hospital payment rates.

[9] A growing literature finds that physician concentration significantly affects this bargaining process. Payments are higher in markets where physicians are more concentrated (Dunn and Shapiro, 2014; Baker et al., 2014; Kleiner et al., 2015; Clemens and Gottlieb, 2017).

its payments may loom large in insurer–physician negotiations. Benchmarking private payments to Medicare's payments may be a straightforward way for contracts to acknowledge and readily adjust to changes in the value of that alternative.

A fourth possibility emphasizes insurance regulations. Regulations require insurers to ensure access to "medically necessary" services. Benchmarking payments to Medicare's rate structure may be the easiest approach to satisfying this requirement.[10] Private payments are almost universally marked up rather than marked down relative to Medicare's rates. Such a payment structure ensures that private insurers are paying sufficiently high rates to generate at least as much care access as Medicare beneficiaries enjoy.

3. Medical pricing data

Our main analysis considers firm-to-firm pricing in the context of medical claims processed by one large insurer, Blue Cross Blue Shield of Texas (BCBS). BCBS is by far the largest carrier in Texas, commanding over 40 percent of each insurance market segment (Kaiser Family Foundation, 2017). Our main database covers the universe of BCBS's payments for outpatient care in 2010; we expand our sample to cover 2008–2011 for one analysis.[11] For each claim, the database details the treatment provided, location, physician, physician group, and BCBS's payment to that group. We restrict this universe along several dimensions. The full 2010 dataset contains 57,613,494 claim lines and \$4.29 billion in spending, which we clean as described in Appendix B.1. This initial cleaning, which eliminates payments made to out-of-network physicians (who have not reached a negotiated agreement with BCBS on reimbursement rates) leaves us with 44,055,829 service lines and \$2.63 billion of spending. We will subsequently examine the other segment of the data separately.

We next merge the remaining claims with Medicare billing codes. In order for private insurers to benchmark prices to Medicare, those services would need to be billed using Medicare's billing codes. The services we cannot merge are thus clearly not benchmarked to Medicare's relative value scale. The merge retains over 97 percent of claims for evaluation and management, diagnostic imaging, and surgical services. We lose notable portions of one broad spending category, namely laboratory tests, for which both Medicare and BCBS frequently base payments on alternative codes. The remaining analysis sample includes 3,681 unique codes from the Healthcare Common Procedure Coding System (HCPCS), which comprise 23,933,577 service lines and \$2.05 billion of spending.[12]

The claims data also allow us to describe the provider groups serving BCBS beneficiaries, at least in terms of the care they provide to that sample. To enable our subsequent investigation of heterogeneity in Medicare benchmarking, we measure the total value of the care each group provides to BCBS patients in a given year. Our final dataset, which is summarized in Table 1, includes care provided by over 80,000 physician groups as identified by their billing identification number.[13] 15,000 of these groups bill more than \$10,000 annually and account for 97 percent of BCBS spending. These 15,000 groups filed an average of just under 1000 claims, had an average of 4 physicians, and saw an average of just over 400 patients.

4. Empirical approach

Our primary empirical goal is to estimate the frequency with which private reimbursement rates are benchmarked directly to Medicare's relative rate structure. We begin by presenting visually striking evidence of bunching in the ratios of physician groups' payments relative to Medicare's payments. We then develop an approach to formalizing this visual evidence in Section 4.1. Next, we present an empirical approach for exploiting policy-driven changes to Medicare's Relative Value Units (RVUs) in Section 4.2. Finally, Section 4.3 relates these approaches.

4.1. Measuring Medicare benchmarking with bunching

We begin our empirical analysis by examining the relationship between private and Medicare pricing in the cross-section. To do so, we first divide BCBS's payment to group g for service j at time t ($P_{g,j,t}$) by Medicare's allocation of RVUs to that service at that time. This defines an "Implied Conversion Factor" (ICF) as:

$$ICF_{g,j,t} = \frac{P_{g,j,t}}{RVU_{j,t}}. \tag{2}$$

While an ICF is defined for every claim, simply computing the ICF does not tell us whether the claim was, in fact, contractually benchmarked as a markup relative to Medicare's RVUs. To gauge the prevalence of contractually specified benchmarking, we analyze the regularity with which a particular group's payments reflect *the same* ICF. Specifically, we investigate the prevalence of sharp bunching in the ratio of a group's payments relative to Medicare.

Graphical illustrations, as presented in Fig. 1, can help to build intuition regarding our bunching approach's strengths and weaknesses. Panel A shows payment rates for the services provided regularly by a single physician group in the 2010 BCBS claims data.[14] Each circle on the graph is a unique payment amount for a unique service code. That is, if the group sometimes received \$45 for a standard office visit (HCPCS code 99213), and other times received \$51, those two amounts would show up as separate circles. The Blue Cross payment amount is on the y-axis and the Medicare payment for the service is on the x-axis, both shown on log scales. Taking logs of equation (2) reveals that Medicare-linked pricing implies a one-for-one relationship between the log Medicare payment and the log private reimbursement.[15]

$$\ln(P_{g,j,t}) = \ln(ICF_{g,j,t}) + \ln(RVU_{j,t}), \tag{3}$$

which has an implied coefficient of 1 on $\ln(RVU_{j,t})$. The y-intercept (in logs) is simply the log of the ICF.

Panel A shows the data from a mid-sized group (billing BCBS between \$200,000 and \$1 million in 2010) for which a single ICF dominates the payment picture. The most natural interpretation of this graph is that those services on the solid line are priced according to Medicare's fee schedule with a common ICF, while

[10] Beneficiaries may lack access to care if payment rates are too low to induce physicians to treat them. The Medicaid program's low payment rates, for example, are often linked to the possibility that beneficiaries will lack access to essential services.

[11] Our empirical results for other years are very similar to those for 2010. We focus on this one-year for brevity and show other years' results in the appendix.

[12] The HCPCS coding system is used by Medicare and many private insurers. The set of codes includes those developed for the Current Procedural Terminology (CPT) system by the American Medical Association. Appendix Table B.1 shows the exact data loss resulting from each step of cleaning. The key conclusion from this table is that, once we restrict ourselves to the relevant universe of data, additional losses from merging in Medicare codes and eliminating infrequent codes are not substantial.

[13] This is analogous to the commonly used tax ID number in Medicare claims data, but our version is anonymized.

[14] The figures exclude any code-by-payment combination that appears ten times or fewer in the data for the relevant physician group. The more systematic analysis presented below has no such exclusion. Throughout this analysis, we restrict to data from the period before BCBS implemented the RVU updates (January 1–June 30, 2010). This way our calculations are not confounded by RVU changes.

[15] Rearranging (2) and then taking logs yields:

Table 1
Summary statistics by physician group.

	Mean	Median	Std. Dev.	Min.	Max.
Panel A: all groups ($N = 80,675$)					
Number of unique services	9.70	3	27.23	1	~1,700
Number of patients	87.59	2	698.85	1	~61,930
Number of doctors	1.73	1	7.93	1	~1100
Number of claims	201	3	1763	1	~163,360
Mean allowed amount	108.91	84.43	125.16	0.64	~7680
Total BCBS revenue	25,457	383	274,700.3	0.64	~43,000,000
Panel B: groups with billings >$10,000 ($N = 15,235$)					
Number of unique services	35.99	24	53.12	1	~1700
Number of patients	424.35	151	1523	1	~61,930
Number of doctors	4.14	2	17.56	1	~1100
Number of claims	981.13	386	3860	1	~163,360
Mean allowed amount	105.52	84.65	136.3	10.75	~7680
Total BCBS revenue	124,687	44,392	606,644	10,000	~43,000,000

Note: Table shows summary statistics for data by physician group. ~ indicates rounding. *Source: Authors' calculations using claims data from BCBS.*

the remaining services are priced separately. Several of the circles below the solid line plausibly involve instances of a less common, but still contractually specified, ICF for this group. A conservative estimate would view these and other circles off the solid line as deviations from Medicare-linked pricing.

Panel B shows the full distribution of this group's markups relative to Medicare's rates. To calculate these markups, we simply divide the y-value of each dot in Panel A (the private payment) by its x-value (the Medicare rate).[16] Panel B shows a clear spike in the distribution of these ratios at around 1.4, indicating that most claims were paid based on a 40 percent markup over Medicare. This spike includes all of the services along the red line in Panel A. Other scattered values in the histogram reflect the deviations away from that line.

Panels C through F show graphs constructed analogously, but for two larger groups that provide more unique services at more distinct prices. The group shown in Panels C and D exhibits two clear spikes in the ICF frequency distribution, with a smattering of other values. The group shown in Panels E and F has a range of ICFs, none of which visually dominates the payment picture. These plots indicate a remarkably complicated contract with BCBS.

Estimating the pervasiveness of "common" ICFs requires a definition of "common." When presenting our results, we will explore sensitivity to the threshold we impose for the frequency with which an ICF must appear in a group's payments. This also requires an assumption on our rounding of the ratio of private to public payments. We explore sensitivity to the choice of rounding as well.

In addition to estimating the share of ICFs that are benchmarked, we run descriptive regressions to investigate correlates of average mark-ups. That is, we estimate

$$\ln(ICF_{g,j}) = \boldsymbol{X_{g,j}}\gamma + e_{g,j} \tag{4}$$

where $\boldsymbol{X_{g,j}}$ contains characteristics of the physician group or local market, such as firm size or concentration. We measure firm size as log total billings to the insurer. We compute firm market share within a local health care market (hospital service area) and specific service, and we measure the degree of concentration across all physician practices within that market (using the HHI at the service-by-area level). We estimate this equation at the claim level, and compute clustered standard errors that allow $e_{g,j}$ to have arbitrary correlation within clusters at the physician group level, since that is the level at which the ICFs are negotiated.

4.2. Framework for analyzing benchmarking using RVU updates

We next develop an estimation framework based on changes in Medicare's relative value scale. A committee of the American Medical Association, composed of representatives of various physician specialties, recommends RVU updates to Medicare (Government Accountability Office, 2015). These updates come in two main forms: reassessments of the resources required to provide a single service, and revisions to part of the underlying methodology. For example, a revision to the method for computing physician effort can change the weights assigned to many service codes. At least one broad update of this sort appears to occur annually over the period we study, as do hundreds of larger service-specific reassessments.

The vast majority of updates to Medicare payments go into effect on January 1 each year. But even when relying on these rates, private insurers have a choice about whether and when to shift from one year's relative value scale to the next year's (Borges, 2003). BCBS informs its providers of the date on which such updates go into effect through its provider newsletter, the *Blue Review*. During the period of our primary analysis sample, the newsletter announced updates taking place on July 1, 2010 (BCBS, 2010).[17]

Panel A of Fig. 2 presents an example of how these changes impact physician payments in our BCBS data. This graph shows average log payments by day for the most commonly billed service code, a standard office visit with an established patient (code 99213). The average log payment jumps distinctively on July 1, 2010, the day on which BCBS implemented the 2010 relative values. Medicare's log RVUs for this service rose by 0.068 between the 2009 and 2010 fee schedules. BCBS's average payment rose by approximately 4 percent. To study BCBS's payment updates systematically, we next develop a method for using high frequency payment changes to infer the share of private reimbursements linked to Medicare.

When a payment $P_{g,j,t}$ is linked to Medicare's relative values, we can write

$$P_{g,j,t} = \varphi_{g,t} \cdot RVU_{j,t} \tag{5}$$

[16] This distribution has the same sample restrictions as in Panel A; see footnote [14] for details. Note that each observation from Panel A has equal weight in the distribution in Panel B, so the distributions in Panels B, D, and F are not weighted to reflect the frequency with which we observe each markup. A weighted version would increase the relative heights of the highest bars, since the common ICFs are, by definition, more common than other markups.

[17] For one empirical extension, we incorporate additional payment changes implemented on July 1, 2008, on August 15, 2009, and on September 1, 2011 (BCBS, 2008, 2009, 2011). In all four years, the standard deviations of changes in log RVUs are around 0.07, or approximately 7 percent. This means that the Medicare changes contain sufficient variation for us to exploit and generate reasonably precise estimates.

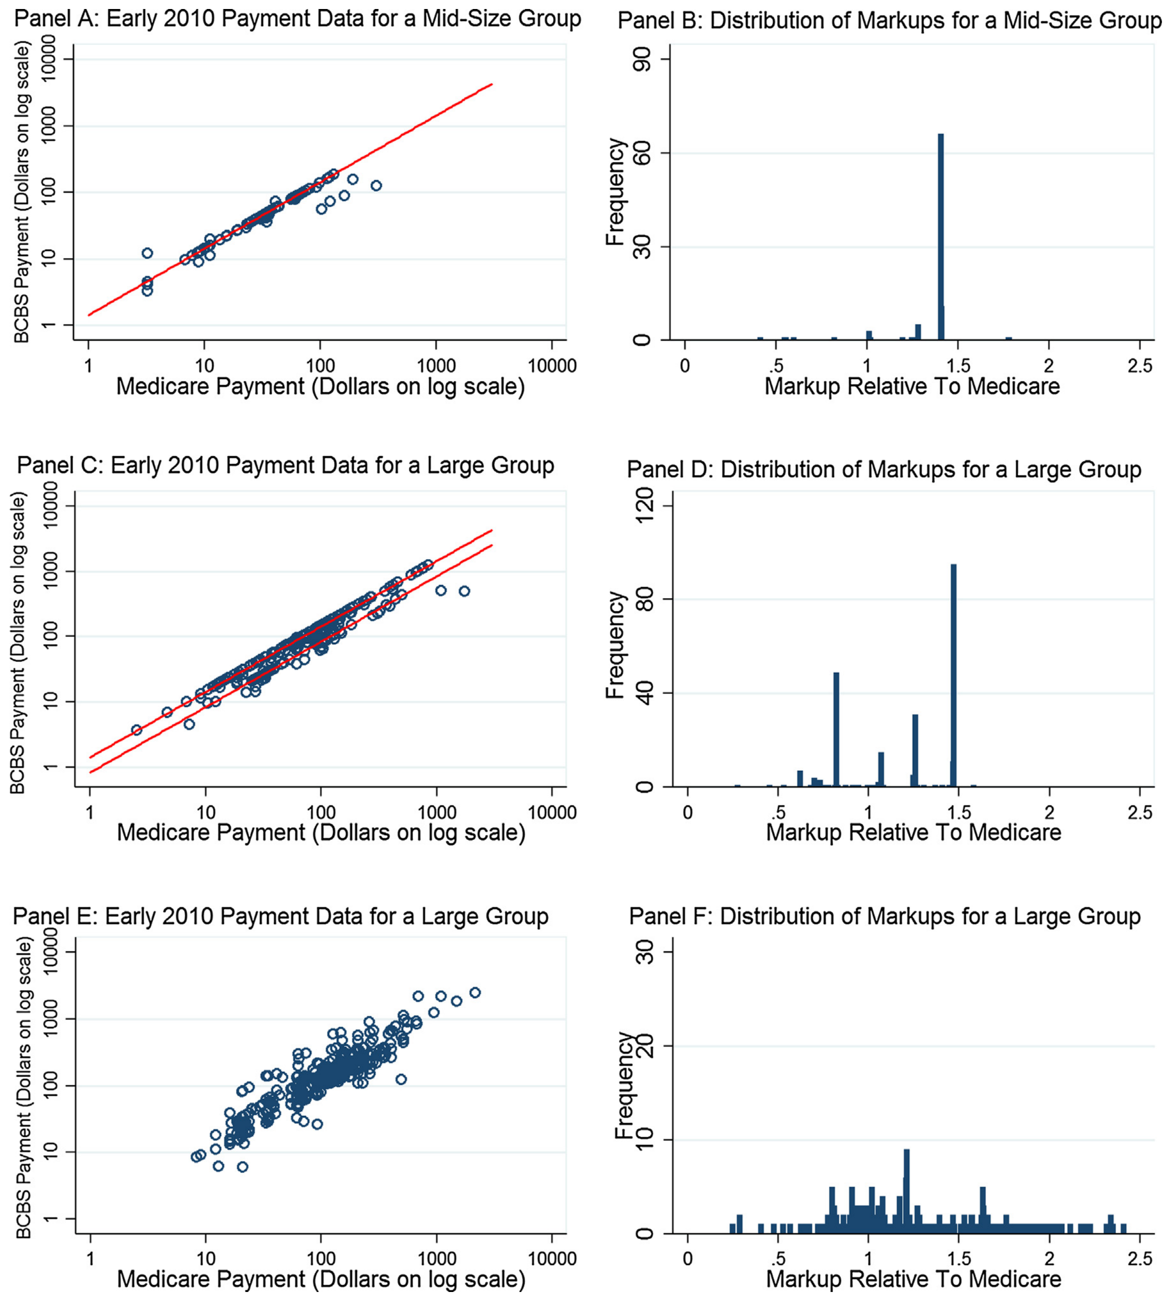

Fig. 1. Raw payments for illustrative physician groups. *Note*: Panels A–F present the raw data on BCBS reimbursement rates, and associated Medicare reimbursement, for 3 different physician groups in 2010. In Panels A, C, and E, each observation is a unique reimbursement paid for a particular service to the group. The lines have a slope of 1 (in logs) and represent the groups' most common Implied Conversion Factors. Panels B, D and F plot the distribution of markups relative to the Medicare rates for all payments each group received. They show clear spikes at the values that we identify as common Implied Conversion Factors in Panels A, C, and E. To comply with confidentiality rules, we omit from these graphs a small share of each group's claims. The share of claims whose observations are suppressed is 14.2% in Panels A and B, 1.94% in Panels C and D, and 2.95% in Panels E and F. *Source: Authors' calculations using RVUs from the Federal Register and claims data from BCBS.*

or, taking logs, $\ln(P_{g,j,t}) = \ln(\varphi_{g,t}) + 1 \cdot \ln(RVU_{j,t}),$ (6)

where $\varphi_{g,t}$ is the Implied Conversion Factor (ICF) from Section 4.1. Eq. (6) describes a linear relationship between log private insurance payments and log RVUs for a service. It describes the one-for-one relationship between log RVUs and log private payments that obtains when contracts are specified in this manner. If the markup is a constant, it will be reflected in the constant term of a regression version of (6). If the mark-up varies across physician groups, then it will be captured by group fixed effects. If it varies both across groups and across time, then it will be captured by group-by-time fixed effects.

Payments may alternatively be negotiated without reference to RVUs. In this case, we have

$$P_{g,j,t} = \rho_{g,j,t} \text{ or } \ln(P_{g,j,t}) = \ln(\rho_{g,j,t}), \quad (7)$$

with no role for $\varphi_{g,t}$ or $RVU_{j,t}$.

When RVU allocations change, Eqs. (6) and (7) contain predictions for how private reimbursements will adjust. Consider two time periods, across which Medicare shifts payments by $\Delta \ln(RVU_{j,t})$. Let $\varepsilon_{g,j,t} = \Delta \ln(\rho_{g,j,t})$ be any change in the alternative non-benchmarked payment (as in Eq. (7)). We can now write both

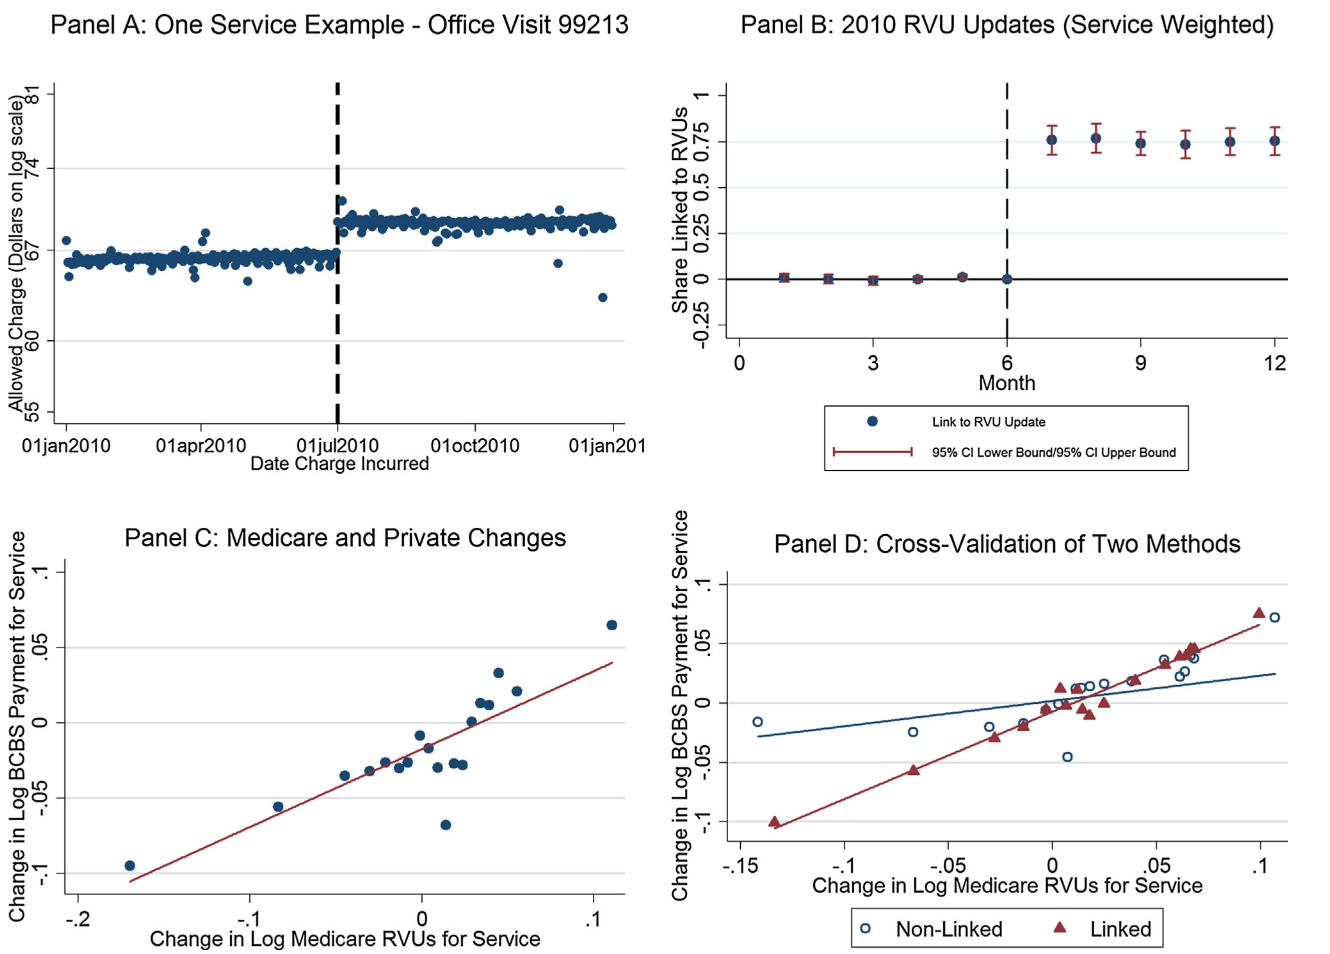

Fig. 2. Benchmarking estimates based on price changes. *Note*: All panels use data from calendar year 2010. BCBS implemented its update from the 2009 to 2010 relative value scales on July 1, 2010, as indicated by the vertical dashed line in Panels A and B. Panel A presents daily averages of BCBS's log payment for a standard office visit. Panel B shows estimates of β_t from Eq. (12), weighting observations equally. Standard errors for the estimates in Panel B are calculated allowing for arbitrary correlation among the errors associated with each HCPCS service code (including modifiers for the professional and technical components of diagnostic imaging services). Panel C presents a binned scatterplot of the relationship between Medicare payment updates (into 20 vigintiles) and changes in private payments. Private price changes are computed as the difference between service-level average payments after and before July 1, 2010. Panel D is similar, but with separate data and estimation for services that we identify as being linked to Medicare on the basis of their Implied Conversion Factors and those we identify as being non-linked. For presentation in the binned scatterplot, observations within each class of services (i.e., linked or non-linked) are grouped into twenty vigintiles on the basis of the log change in the service code's Medicare RVU allocation. The regression lines shown in Panels C and D are estimated at the underlying service-code level. *Source: Authors' calculations using RBRVS updates from the Federal Register and claims data from BCBS.*

types of prices in terms of service fixed effects and changes as follows. For Medicare-linked services, we have:

$$\ln(P_{g,j,t}) = \phi_j \mathbb{1}_j + \phi_g \mathbb{1}_g + \phi_{g,j} \mathbb{1}_g \cdot \mathbb{1}_j + \Delta \ln(RVU_{j,t}) \cdot \mathbb{1}_{\{t=\text{post}\}}. \tag{8}$$

For services not linked to Medicare, we have:

$$\ln(P_{g,j,t}) = \phi_j \mathbb{1}_j + \phi_g \mathbb{1}_g + \phi_{g,j} \mathbb{1}_g \cdot \mathbb{1}_j + \varepsilon_{g,j,t} \cdot \mathbb{1}_{\{t=\text{post}\}}. \tag{9}$$

In these equations, $\mathbb{1}_{\{t=\text{post}\}}$ is an indicator for the second time period. Under both pricing schemes, the fixed effects capture baseline payments to group g for service j in the first period, while the interaction with $\mathbb{1}_{\{t=\text{post}\}}$ captures the change between the two periods.

The linearity of Eqs. (8) and (9) implies a straightforward way to estimate the fraction of services with payments benchmared to Medicare's relative values. Eq. (8) says that a linear regression of log private payments on changes in log Medicare RVUs, for services with prices linked to Medicare, should yield a coefficient of 1 after controlling for the relevant sets of fixed effects. Eq. (9) shows that the same regression should yield a coefficient of 0 for services not priced based on Medicare, as long as the non-Medicare payment changes ($\varepsilon_{g,j,t}$) are uncorrelated with RVU updates.

More generally, suppose that both types of payments exist, and specifically that a constant share σ of payments are benchmarked to Medicare prices, while $1-\sigma$ are set independently. (We will subsequently allow for heterogeneity in σ.) The average of log reimbursements is then given by a weighted average of Eqs. (8) and (9), and the coefficient on log RVU updates can reveal the linked share σ:

$$\ln(P_{g,j,t}) = \phi_j \mathbb{1}_j + \phi_g \mathbb{1}_g + \phi_{g,j} \mathbb{1}_g \cdot \mathbb{1}_j + \sigma \cdot \Delta \ln(RVU_{j,t}) \cdot \mathbb{1}_{\{t=\text{post}\}} + \eta_{g,j,t}, \tag{10}$$

where we define $\eta_{g,j,t} = (1-\sigma) \cdot \varepsilon_{g,j,t} \cdot \mathbb{1}_{\{t=\text{post}\}}$. Eq. (10) suggests that, in a linear regression with appropriate fixed effects, we can infer the Medicare-linked share from the coefficient on log RVU changes. This motivates our baseline specification for estimating σ. We use data at the level of individual claims, indexed by c, to estimate:

$$\ln(P_{c,g,j,t}) = \beta \Delta \ln(RVU_j) \cdot \mathbb{1}_{\{t=\text{post}\}} + \phi_t \mathbb{1}_{\{t=\text{post}\}} + \phi_j \mathbb{1}_j + \phi_g \mathbb{1}_g + \phi_{g,j} \mathbb{1}_g \cdot \mathbb{1}_j + \eta_{c,g,j,t}. \tag{11}$$

Eq. (11) is a claims-level version of Eq. (10) where $\hat{\beta}$ estimates the share of payments based on Medicare rates. It adds a time period fixed effect $\mathbb{1}_{\{t=\text{post}\}}$ in case private payments shift broadly across the two time periods. This parametric difference-in-differences specification also incorporates full sets of group ($\mathbb{1}_g$), service ($\mathbb{1}_j$), and group-by-service ($\mathbb{1}_g \cdot \mathbb{1}_j$) effects to account for all time-invariant group- and service-specific terms. Thus our estimate of $\hat{\beta}$ is identified only using changes in RVUs across the two time periods. The time effect further limits the identifying variation exclusively to relative changes in RVUs across services. To obtain the share of spending linked to Medicare, we can estimate Eq. (11) weighted by the average pre-update price of each service.[18] We compute standard errors allowing for clustering in the errors $\eta_{c,g,j,t}$ at the service-code level, which is the level at which variation in RVU updates occurs.

To describe the timing with which BCBS incorporates Medicare updates into its reimbursements, we also present dynamic estimates from the following parametric event study:

$$\ln(P_{c,g,j,t}) = \sum_{t \neq 0} \beta_t \Delta \ln(RVU_j) \cdot \mathbb{1}_t + \phi_t \mathbb{1}_t + \phi_j \mathbb{1}_j + \phi_g \mathbb{1}_g + \phi_{g,j} \mathbb{1}_g \cdot \mathbb{1}_j + \eta_{c,g,j,t}. \quad (12)$$

When estimating Eq. (12), we again cluster standard errors at the service code level. We normalize t such that $t = 1$ is the month in which BCBS has announced that it will implement Medicare's fee updates. We thus expect to see $\hat{\beta}_t = 0$ for periods preceding the updates' incorporation, $t < 0$, while the $\hat{\beta}_t$ for $t > 0$ are our estimates of how often Medicare updates are incorporated into private payments. A flat profile of the post-update $\hat{\beta}_t$ estimates would suggest that all price changes correlated with Medicare changes are implemented instantaneously. An upward trend in these coefficients might suggest that our baseline estimates are affected by ongoing renegotiations between BCBS and firms whose bargaining positions are affected by Medicare updates. We discuss this concern in detail in Appendix C.

4.3. Relating our approaches

The analyses we implement have complementary strengths and weaknesses. A shortcoming of the cross-sectional analysis of bunching is that it requires us to observe a constant markup across many services. Thus it may fail to detect genuine Medicare linkages involving markups that are common across relatively small numbers of services. These linkages would be detected, however, by our analysis of Medicare payment updates. The latter approach does not depend on the commonality of the markup. If the service is benchmarked, the payments will change when the underlying relative values change.

A shortcoming of the payment update approach, on the other hand, is that it could be biased if Medicare updates occur contemporaneously with changes driven by new contract negotiations. Our robustness analysis and our investigation of the precise timing of Medicare-linked changes provide evidence that new contract negotiations are unlikely to underlie our results. Nonetheless, these analyses cannot rule out active contract renegotiations altogether.

The bunching and changes approaches are thus complementary in that they have distinctive strengths and weaknesses. We relate these two analyses, and demonstrate consistency between them, by showing that the service-firm pairs we identify as benchmarked are strongly correlated across our approaches. We do this by dividing the data into subsamples according to the benchmarking results we obtain in our bunching analysis. We then estimate Eq. (11) separately on these subsamples.

[18] Since the unweighted regression treats each claim equally, it effectively weights service codes by the frequency with which they are used.

Table 2
Services priced according to common implied conversion factors.

	Frequency threshold		
	5%	10%	20%
Panel A: dollar-weighted			
Rounding for ICFs:			
$0.02	83%	76%	66%
$0.10	86%	80%	71%
$0.20	87%	80%	71%
Panel B: service-weighted			
$0.02	87%	81%	70%
$0.10	89%	84%	75%
$0.20	89%	85%	75%

Note: Each cell shows the share of services for which payments are associated with a common Implied Conversion Factor (cICF), as defined in the main text. Data are from January 1 – June 30, 2010, over which time BCBS used the 2009 version of Medicare's Resource Based Relative Value Scale. The cells within each panel show how the linked share varies as we apply different thresholds for the frequency required to quality as a cICF. The column labeled "Rounding" indicates the rounding applied to each estimated ICF. An ICF is defined as "common" for the payments to a physician group if it accounts for at least the fraction of services associated with the specified Frequency Threshold. *Source: Authors' calculations using claims data from BCBS.*

5. Baseline benchmarking results

5.1. Bunching estimates

Table 2 presents estimates of the share of services linked to Medicare according to the bunching method from Section 4.1. The results show that at least two-thirds of spending, and three-quarters of services, have prices linked to Medicare. The full table explores sensitivity to two key assumptions. First, we round the value of each $ICF_{c,g,j,t}$ to the nearest 20 cents, 10 cents, or 2 cents to explore sensitivity to rounding error. Second, we define "common ICFs" as those that rationalize a sufficiently large share of the insurer's payments to a single physician group. In Fig. 1, for example, the red line in Panel A should undoubtedly qualify as common. Other values may also qualify depending on the strictness of the threshold we apply. We consider thresholds ranging from 5 to 20 percent of a group's claims, then calculate the share of the insurer's payments associated with *any* of a group's common ICFs.

The Medicare-benchmarked shares range from 65 to 90 percent depending on the rounding and frequency thresholds; they decrease substantially with the stringency of the definition for a common ICF, but are not sensitive to the choice of rounding threshold. Appendix Table B.2 shows that alternative measures generate qualitatively similar results.[19]

Going forward, we require as our baseline that common ICFs account for 10 percent of a group's claims, when rounded to the nearest $0.02. The motivation for adopting a stringent rounding threshold is to be conservative in the extent to which our method detects false positives. At the same time, the 10 percent threshold ensures that multiple ICFs can be readily detected. Using this definition, over half of firms have just one common ICF. Fewer than 5 percent have more than 2 common ICFs.

Having identified these ICFs, we use them to describe how the generosity of BCBS reimbursements relates to firm and market characteristics. Table 3 presents estimates of Eq. (4), which regresses the ICF values themselves (in logs) against physician

[19] If we only count the single most common ICF for each group, the estimates are very similar to those reported in Table 2 when imposing a 20 percent threshold. Unfortunately, theory does not provide guidance as to which threshold is most appropriate, and the choice of threshold substantially affects our estimate of the linked share. Our changes-based estimation strategy is not sensitive to choices of this sort.

Table 3
Firm size and implied conversion factors.

Dependent variable:	(1)	(2)	(3)	(4)	(5)
	Log implied conversion factor (ICF)				
Firm SIZE (Log Spending)	0.058**			0.058**	0.040**
	(0.004)			(0.005)	(0.006)
Firm market share		0.241**		−0.158**	−0.092**
		(0.015)		(0.037)	(0.029)
Market concentration			0.238**	0.318**	0.159**
			(0.020)	(0.036)	(0.028)
N	20,736,449	20,736,449	20,736,449	20,736,449	20,736,449
No. of clusters	23,098	23,098	23,098	23,098	23,098
Code EFFECTS	No	No	No	No	Yes
HSA fixed effects	No	No	No	No	Yes

Note: This table shows estimates of relationship between the level of physicians' reimbursements, measured using Implied Conversion Factors (ICFs), and measures of firm size and/or market concentration. Standard errors are calculated allowing for arbitrary correlation among the errors associated with each physician group. The construction of all variables is discussed in the main text. *Source: Authors' calculations using claims data from BCBS.*

** Statistical significance at 0.01.

* Statistical significance at 0.05.

+ Statistical significance at 0.10.

group and market characteristics.[20] Columns 1–3 reveal that each of firm size, market share, and market concentration is, by itself, positively correlated with the generosity of the firms' payments. Consistent with other work on health care pricing (Dunn and Shapiro, 2014; Baker et al., 2014; Kleiner et al., 2015; Cooper et al., 2015; Clemens and Gottlieb, 2017), payments to large firms in markets with high levels of concentration are more generous than payments to small firms in markets with low levels of concentration. Columns 4 and 5 include all three characteristics together, with column 5 also adding fixed effects for service codes and geographic areas. Firm size remains a strong predictor of the average generosity of a firm's payments, as does overall market concentration. Market share switches signs, likely because of collinearity with log firm size.

5.2. Results from Medicare fee change analysis

We next move on from estimating ICFs to exploiting Medicare's RVU changes. Using the method from Section 4.2, Panel B of Fig. 2 presents event study estimates of the link between Medicare's relative value scale and BCBS reimbursements. It shows estimates of Eq. (12) for the Medicare payment changes implemented on July 1, 2010. The regression underlying the figure weights each claim by the dollar value of the service.[21]

The estimates reveal substantial – but not universal – links between Medicare updates and the payments providers receive from BCBS. The coefficients imply that $\hat{\sigma} = 75$ percent of services have payments linked to Medicare's relative values. The dramatic dynamics in the figure suggest that this reflects a contractual link between Medicare's relative values and BCBS payments. As in the raw data for standard office visits presented in Panel A, we see that payment changes occur when we expect. Importantly, the estimates of σ are both economically and statistically larger than 0 and smaller than 1, implying that payments for a substantial share of services deviate from strict benchmarking to Medicare's relative values. Sections 6.1 and 6.2 will investigate these deviations in detail.

Column 1 of Table 4 presents our baseline estimates of Eq. (11), which summarizes this result in a single coefficient. The estimate in column 1 of Panel A shows that roughly 55 percent of BCBS's spending is linked to Medicare's relative values. This estimate corresponds with the analysis reported if panel B of figure 2. In Panel B, we weight service codes equally rather than according to baseline payments. The unweighted estimate implies that roughly three quarters of BCBS's physician claims are paid based on Medicare's relative value scale. The difference in coefficients between Panels A and B implies that payments for relatively expensive services are less likely to be benchmarked to Medicare than are payments for low-cost services.[22]

5.3. Robustness and cross-validation of the two approaches

Table 4 probes the robustness of our changes-based estimates to a variety of specification checks. Column 1 of each panel reports our baseline specification, which includes a full set of group-by-HCPCS code fixed effects. In this baseline, we control for time effects with a simple post-update indicator. Column 2 drops the group-by-HCPCS code fixed effects in favor of a more parsimonious set of HCPCS code fixed effects. Column 3 augments the baseline specification by controlling for a cubic trend in the day of the year, which we interact with the size of each service's Medicare fee change. Column 4 allows the cubic trend in day to differ between the periods preceding and following the fee schedule update, as in a standard regression discontinuity design. The table shows that these specification changes have essentially no effect on the estimated coefficient $\hat{\beta}$. This reinforces the interpretation that, among services billed using standard HCPCS codes, roughly 55 percent of BCBS's spending is linked to Medicare's relative value scale.

Fig. 3 shows dynamic estimates that pool together data from 2008 to 11, and use RVU changes from 2008 to 10 simultaneously. This estimation allows us to check whether any given change is offset by other changes in subsequent years. The figure shows no evidence of such an offset. The short-run responses to RVU changes persist for many quarters thereafter. The graph also shows flat pre-trends over long time periods, such as 6 quarters before the 2009 RVU changes and 10 quarters before the 2010 changes.

The estimates presented thus far may differ from the true Medicare benchmarking parameter σ if changes in other terms of providers' contracts covary with the Medicare changes. Indeed, payment changes that significantly alter physician groups' average

[20] Appendix Table B.3 also shows how these same characteristics relate to the frequency of deviations from Medicare benchmarking, and the value of the deviations when they occur.

[21] Appendix Fig. C.1 shows an unweighted version of this graph, for each year.

[22] Appendix Tables C.1 and C.2 replicate Panels A and B, respectively, in other years' data.

Table 4
Estimating Medicare benchmarking using RVU changes.

Dependent variable:	(1)	(2)	(3)	(4)
	Log private reimbursement rate			
Panel A: weighted by price				
Log RVU change × post	0.539**	0.544**	0.568**	0.538**
	(0.061)	(0.061)	(0.060)	(0.061)
N	23,933,577	23,933,577	23,933,577	23,933,577
No. of clusters	3681	3681	3681	3681
Panel B: unweighted				
Log RVU change × post	0.750**	0.748**	0.765**	0.749**
	(0.038)	(0.038)	(0.043)	(0.038)
N	23,933,577	23,933,577	23,933,577	23,933,577
No. of clusters	3681	3681	3681	3681
Group-by-code effects	Yes	No	Yes	Yes
Code effects	No	Yes	No	No
Cubic time × RVU change	No	No	Yes	No
Cubic time × post	No	No	No	Yes

Note: The table shows the results of OLS specifications of the form described in Section 4.2. Each column in each panel reports an estimate of $\hat{\beta}$ from Eq. (11). Observations are at the claim-line level and are equally weighted (Panel B), or weighted according to each service's average payment during the baseline period (Panel A). Data are from 2010. Standard errors are calculated allowing for arbitrary correlation among the errors associated with each HCPCS service code (including modifiers for the professional and technical components of diagnostic imaging services). Additional features of each specification are described within the table. The construction of all variables is further described in the main text. *Source: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.*

** Statistical significance at 0.01.

* Statistical significance at 0.05.

\+ Statistical significance at 0.10.

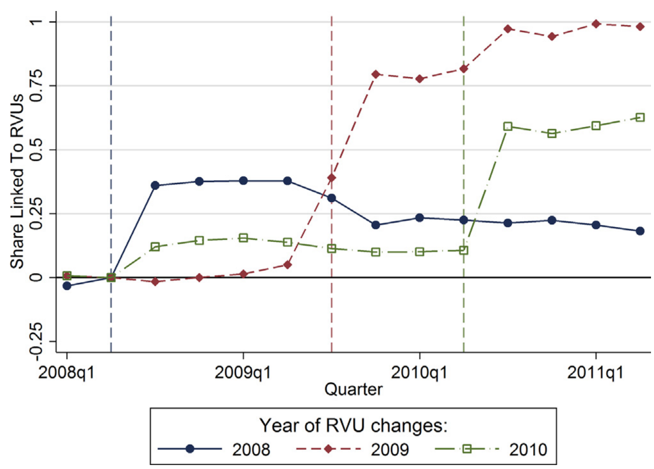

Fig. 3. Estimating multiple years' RVU updates simultaneously. *Note*: The figure reports estimates of β_t^{08}, β_t^{09} and β_t^{10} from the following modification of Eq. (12):

$$(15)\ln(P_{c,g,j,t}) = \sum_{t\neq 0}\beta_t^{08}\Delta\ln(RVU_{j,08})\cdot\mathbb{1}_t + \sum_{t\neq 0}\beta_t^{09}\Delta\ln(RVU_{j,09})\cdot\mathbb{1}_t^{10} + \sum_{t\neq 0}\beta_t\Delta\ln(RVU_{j,10})\cdot\mathbb{1}_t + \phi_t\mathbb{1}_t + \phi_j\mathbb{1}_j + \phi_g\mathbb{1}_g + \phi_{g,j}\mathbb{1}_g\cdot\mathbb{1}_j + \eta_{c,g,j,t}.$$

In this specification, $\Delta\ln(RVU_{j,T})$ refers to the log of Medicare's RVU updates from calendar year $T-1$ to calendar year T. The corresponding coefficients β_t^T indicate what share of the year-T RVU updates were incorporated into BCBS payments during calendar quarter t. Standard errors are calculated allowing for arbitrary correlation among the errors associated with each HCPCS service code (including modifiers for the professional and technical components of diagnostic imaging services). BCBS implemented its RVU updates on July 1, 2008, August 15, 2009, and July 1, 2010. The omitted interaction ($t=0$) is 2008Q2 for all of the RVU update variables. The regression line is estimated at the underlying service-code level and is dollar-weighted. *Source: Authors' calculations using claims data from BCBS.*

Medicare payment can move private payments in subsequent years, due in part to the resulting changes to their bargaining positions (Clemens and Gottlieb, 2017). In Appendix C.4, we thus draw on institutional detail and theoretically motivated specification checks to explore how much our estimates might deviate from the true share of payments benchmarked to Medicare's relative values. We find no evidence that renegotiations confound the relationship between BCBS's and Medicare's payments over the time horizons we analyze. Appendix C.4 thus bolsters the case for interpreting our estimates of $\hat{\beta}$ as measuring the fraction of services tied directly to Medicare.

To validate that our classification algorithm correctly captures services whose prices are actually linked to Medicare rates, we estimate our baseline changes-based regression separately for services identified as being more and less likely to be benchmarked according our bunching methodology. We classify each group-service pair (g, j) as Medicare-linked if all of group g's claims for service j in the pre-update period appear to be linked to Medicare rates, and as non-linked otherwise. We estimate Eq. (11) separately for these two samples.

Panel D of Fig. 2 shows two binned scatterplots analogous to Panel C, relating log BCBS price changes to log Medicare fee changes separately for the two samples. The linked sample is shown with red triangles and has a slope of 0.9, indicating that BCBS prices for 90 percent of linked services update in response to Medicare changes. The non-linked sample is shown with blue circles, and has a much smaller slope of 0.3.

In Appendix D.1 we consider the external validity of our baseline results using data from Colorado. Using data from one insurer in the Colorado All-Payer Claims Database, we obtain results generally in line with those from BCBS. This highlights that the phenomenon we investigate is not unique to our setting. Deviations from Medicare's payment structure are somewhat more common in the Colorado insurer, but the basic fact – substantial benchmarking, but far from universal – appears broadly relevant.

6. How do private payments deviate from Medicare?

In order to illuminate the economic determinants of benchmarking, we next consider variation in the strength of the link between private payments and Medicare's relative values. We consider the two primary dimensions along which payments vary: differences across physician groups and categories of services.

6.1. Deviations from benchmarking across physician groups

One key difference across groups is the scale of their business with BCBS. Size, which is likely related to market power, could influence physician-insurer negotiations in multiple ways. Table 3 found that larger firms obtain higher ICFs. Larger group size could also lead to more deviations away from this benchmarking, in particular positive deviations (higher reimbursements).

To determine how size relates to benchmarking, we measure the quantity of care each group provides in our data. We then add interactions with practice size to our baseline changes regression, Eq. (11). Table 5 shows the results. The first column reports the baseline, equally weighted regression from Table 4. The second column introduces interactions between the Medicare updates and indicators for the size of the physician group providing the care. We define mid-sized firms as those with $200,000–$1,000,000 in annual billing with BCBS, and large firms as those with more than $1,000,000 in annual billing. Each of these categories comprises one-quarter of the sample, with the remaining half of claims coming from smaller firms. The estimates imply that nearly 90 percent of services provided by firms billing less than $200,000 are benchmarked to Medicare, while roughly 60 percent of services provided by firms billing more than $1,000,000 are benchmarked. Columns 3 and 4 present similar, but dollar-weighted, estimates. The results in column 4 suggest that 77 percent of payments to firms billing less than $200,000 are benchmarked to Medicare, while one-third of payments to firms billing more than $1,000,000 are benchmarked.[23]

Fig. 4 shows that we find a similar relationship between the share linked to Medicare and physician group size using our cross-sectional bunching approach. The series in the figure reveal that this is true in both the equally-weighted and payment-weighted series. It is also true whether or not we adjust for the underlying composition of each group's services, to which we now turn.[24]

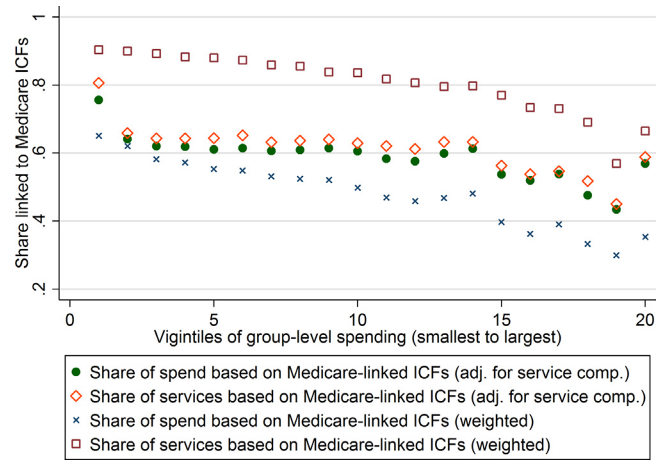

Fig. 4. Frequency of benchmarking and physician group size. *Note*: This figure shows the relationship between a group's Medicare-linked service share and group size. Specifically, it plots variation in the share of services priced according to common Implied Conversion Factors (cICFs), as defined in Section 4.1, according to physician group size. We measure group size by forming 20 vigintiles based on the group's BCBS billing. We require that cICFs account for 10 percent of a group's claims, when rounded to the nearest $0.02. The green dots and orange diamonds show estimates of ζ_b from Eq. (13), which adjust for the composition of each group's services. The blue ×'s and red squares are unadjusted, but weighted to measure the Medicare-linked share of spending in dollar terms as opposed to the share of services. All data are from 2010. *Source: Authors' calculations using claims data from BCBS.*

6.2. Which services deviate from the Medicare benchmark?

The value of improving on Medicare's menu depends on the severity of that menu's inefficiencies. Because it is difficult to systematically quantify Medicare's inefficiencies across a large range of individual services, we focus on one of the Medicare fee schedule's most salient problems. Medicare rates are computed based on average-cost reimbursement, so its reimbursements will hew closer to marginal costs for labor-intensive services than for capital-intensive services. Standard optimal payment models suggest that the latter would be more appropriately reimbursed through combinations of up-front financing of fixed costs and incremental reimbursements closer to marginal cost (Ellis and McGuire, 1986). We can proxy for services' capital and labor intensity by comparing the frequency of benchmarking across categories of care produced with different inputs, such as labor-intensive Evaluation & Management services versus capital-intensive Imaging.[25]

Table 6 estimates Eq. (11) – the relationship between private prices and changes in Medicare's relative values – separately across broad categories of services. The estimates imply that nearly 30 percent more of the payments for Evaluation & Management services are linked directly to Medicare's relative values than for Imaging services.[26]

Second, we divide Imaging codes into subcomponents with high capital and high labor content. Providers often bill separately for taking an image (the capital-intensive part, since it requires an imaging machine) and interpreting it (the labor-intensive part). When the same group supplies both components, it submits the bill as a "Global" service. The results in columns 5–7 show that

[23] Appendix Table C.5 shows similar results in data from other years.

[24] To check whether the relationship between benchmarking and group size is affected by the composition of large and small groups' services, we run a regression that allows group size and service composition to enter simultaneously. We define fixed effects $\mathbb{1}_{b(j)}$ using the "Betos" classification defined by Berenson and Holahan (1990). This hierarchical classification system goes from the broad categories we use here (such as Evaluation & Management and Imaging) to 2-digit (e.g. Advanced Imaging [MRIs and CAT scans]) and 3-digit classifications (e.g. CAT Scan: Head). We categorize all of the medical services in our data at the level of the 1-digit Betos categories.

[25] These categories are defined using the Betos categories described in footnote 24.

[26] Appendix Table C.4 replicates this analysis in other years' data.

Table 5
Medicare benchmarking by firm size.

Dependent variable:	(1)	(2)	(3)	(4)
	Log private reimbursement rate			
Log RVU change × post-update	0.750**	0.882**	0.539**	0.775**
	(0.038)	(0.073)	(0.061)	(0.094)
Log RVU change × post-update × midsize		−0.074		−0.140*
		(0.098)		(0.069)
Log RVU change × post-update × large		−0.293*		−0.448**
		(0.117)		(0.102)
N	23,933,577	23,933,577	23,933,577	23,933,577
Weighting:	Service	Service	Dollar	Dollar

Note: Columns 1 and 3 report the baseline estimates from Table 4 Panels A and B respectively. In columns 2 and 4 we augment these specifications to include interactions between firm size indicator variables and both the "Post" indicator and the interaction between the "Log RVU Change" and "Post" indicator. The omitted category is small firms, defined as those with less than $200,000 in billings. Mid-sized firms are those with billings between $200,000 and $1 million, and large firms are those with billings exceeding $1 million. Data are from 2010. Standard errors are calculated allowing for arbitrary correlation among the errors associated with each HCPCS service code (including modifiers for the professional and technical components of diagnostic imaging services). *Source: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.*

* Statistical significance at 0.05.
** Statistical significance at 0.01.
\+ Statistical significance at 0.10.

Table 6
Public–private payment links across service categories.

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Log private reimbursement rate						
	Evaluation	Imaging	Procedures	Tests	Imaging sub-categories		
					Global	Technical	Professional
Log RVU change × post-update	0.841**	0.564**	0.720**	1.066**	0.545**	0.387*	0.982**
	(0.036)	(0.084)	(0.081)	(0.066)	(0.109)	(0.152)	(0.066)
N	12,259,186	3,630,019	4,750,313	1,542,254	1,826,666	209,178	1,594,175
No. of clusters	221	1085	1936	408	408	244	433

Note: The table shows the results of OLS specifications of the form described in Section 4.2. The cells in each panel report estimates of $\hat{\beta}$ from Eq. (11), with samples selected to contain the HCPCS codes falling into broad service categories. The name of the relevant service category accompanies each point estimate. Data are from 2010. Standard errors are calculated allowing for arbitrary correlation among the errors associated with each HCPCS service code (including modifiers for the professional and technical components of diagnostic imaging services). The construction of all variables is further described in the main text. *Source: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.*

* Statistical significance at 0.05.
** Statistical significance at 0.01.
\+ Statistical significance at 0.10.

payments for the labor-intensive Professional Component are more tightly linked to Medicare's relative values than are the payments for the capital-intensive Technical Component. These patterns support the hypothesis that physicians and insurers are more likely to contract away from Medicare's menu for capital-intensive services than for labor-intensive ones.

Table 7 shows that we find a similar relationship between the share linked to Medicare and service categories using our cross-sectional approach. Benchmarking is 30–50 percent less frequent for Imaging, Procedures, and Tests than for Evaluation & Management services. The results across columns reveal that we find similarly substantial differentials whether or not we control for firm size and whether services are weighted according to the spending they represent.

These results suggest that private contracts deviate when Medicare's rates are most problematic from an efficiency perspective. One-way to interpret this is in light of negotiation and adjustment costs. Private bargaining can overcome these frictions more easily when Medicare's rates are farther from the efficient or equilibrium level that would obtain under unconstrained negotiations.

6.3. How do deviations change incentives relative to Medicare?

What are physicians and insurers aiming to achieve when they negotiate reimbursements that deviate from Medicare's relative

Table 7
Medicare benchmarking by Betos category.

Dependent variable:	(1)	(2)	(3)	(4)
	Payments with common conversion factors			
	Spending share		Service share	
Imaging	−0.427**	−0.471**	−0.300**	−0.355**
	(0.053)	(0.047)	(0.030)	(0.024)
Procedures	−0.309**	−0.352**	−0.336**	−0.388**
	(0.030)	(0.028)	(0.054)	(0.052)
Tests	−0.383**	−0.415**	−0.258**	−0.297**
	(0.051)	(0.047)	(0.055)	(0.054)
Constant	0.921**	0.828**	0.941**	0.829**
	(0.015)	(0.015)	(0.020)	(0.017)
N	542,207	542,207	542,207	542,207
Omitted category		Evaluation & Management		
Additional controls	Group size	None	Group size	None

Note: This table shows estimates of the ν_b coefficients in Eq. (13), namely the relationship between Betos category and the Medicare-linked share of claim lines (columns 1 and 2) or spending (columns 3 and 4). Medicare links are measured using the common Implied Conversion Factors (cICFs) defined in Section 4.1, using data from January 1 through June 30, 2010. We require that cICFs account for 10 percent of a group's claims, when rounded to the nearest $0.02. Columns 1 and 3 show estimates after controlling for vigintile of group size, as measured with BCBS spending, and columns 2 and 4 show estimates without group size controls. Standard errors are two-way clustered (Cameron et al., 2011) by Betos category and physician group. *Source: Authors' calculations using claims data from BCBS.*

** Statistical significance at 0.01.
* Statistical significance at 0.05.
\+ Statistical significance at 0.10.

Table 8
In what direction does BCBS adjust its payments for the various service categories?

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Distributions of payment residuals by Betos categories						
	Evaluation &management	Imaging	Procedures	Tests	Imaging sub-categories		
					Global	Technical	Professional
Residual mean	0.0112	−0.0624	0.0107	0.0301	−0.122	−0.124	0.0150
Residual SD	(0.169)	(0.246)	(0.279)	(0.319)	(0.272)	(0.281)	(0.177)
N	6,010,826	1,743,011	2,312,734	751,726	883,419	102,465	757,127

Note: The table presents means and standard deviations of residuals from estimates of Eq. (14) in data from 2010. That is, we regress the log of BCBS's payments on a set of physician-group fixed effects and the log of each HCPCS code's number of Relative Value Units. This table describes the residuals from that regression. We restrict the sample to the pre-update period (January 1 through June 30, 2010) so that the relative value units are constant for each service throughout the sample. *Source: Authors' calculations using claims data from BCBS.*

prices? In this section, we present evidence on the direction of deviations from strictly Medicare-benchmarked rates to investigate what services BCBS rewards through upward adjustments and discourages through downward adjustments. We do so by describing residuals from the following regression:

$$\ln(P_{g,j}) = \psi \ln(RVU_j) + \mu_g + e_{g,j}. \qquad (14)$$

In a world of perfect benchmarking, we would find $\hat{\psi} = 1$ and $e_{g,j}$ uniformly equal to 0. So the empirical prediction errors $\hat{e}_{g,j}$ contain information about the direction of deviations from strict Medicare benchmarking. We examine heterogeneity in this prediction error across categories of services.

A subtle but important point is that this approach captures deviations from Medicare's relative prices that come through the introduction of multiple Medicare-benchmarked conversion factors. If an insurer thinks the Medicare menu's primary inefficiency is that it uniformly overpays for diagnostic imaging services relative to other services, for example, its preferred contract may simply set a low conversion factor for imaging services and a high conversion factor for other services. Our previous analyses would describe such a contract as being fully linked to Medicare. The analysis in the current section will capture the fact that this is structured to discourage the use of imaging services relative to other services.

Table 8 presents means of the residuals $\hat{e}_{g,j}$ from Eq. (14) across Betos categories.[27] The table shows that payments for Evaluation & Management and Testing services generally have positive residuals while payments for Imaging services have negative residuals. Fig. 5 Panel A plots the distributions of these residuals by service category. The distribution for Imaging shows far more density of negative residuals than those for other services. Testing has more positive residuals, although that is largely driven by one outlier code.[28] Compared to the relative payments implied by Medicare's relative values, BCBS systematically adjusts its contracts to discourage imaging services. This coincides with the conventional wisdom that Medicare's relative values underpay for labor-intensive services relative to other services, and suggests that BCBS aims to partly rectify that mispricing.

Differences in BCBS's adjustments for labor- and capital-intensive services are particularly sharp across the subcategories of diagnostic imaging. Payment adjustments for the labor-intensive Professional Component of these services are positive, at around 0.015 in logs (approximately 1.5 percent). Payment adjustments for the capital-intensive Technical Component of these services are substantially negative, averaging −0.12 in logs. Fig. 5 Panel B shows that this pattern holds throughout the distribution. While it is clear that BCBS reimbursements lean heavily on Medicare's relative values for their basic payment structure, these results suggest that BCBS adjusts its contracts to increase the generosity of payments for labor-intensive services and decrease its payments

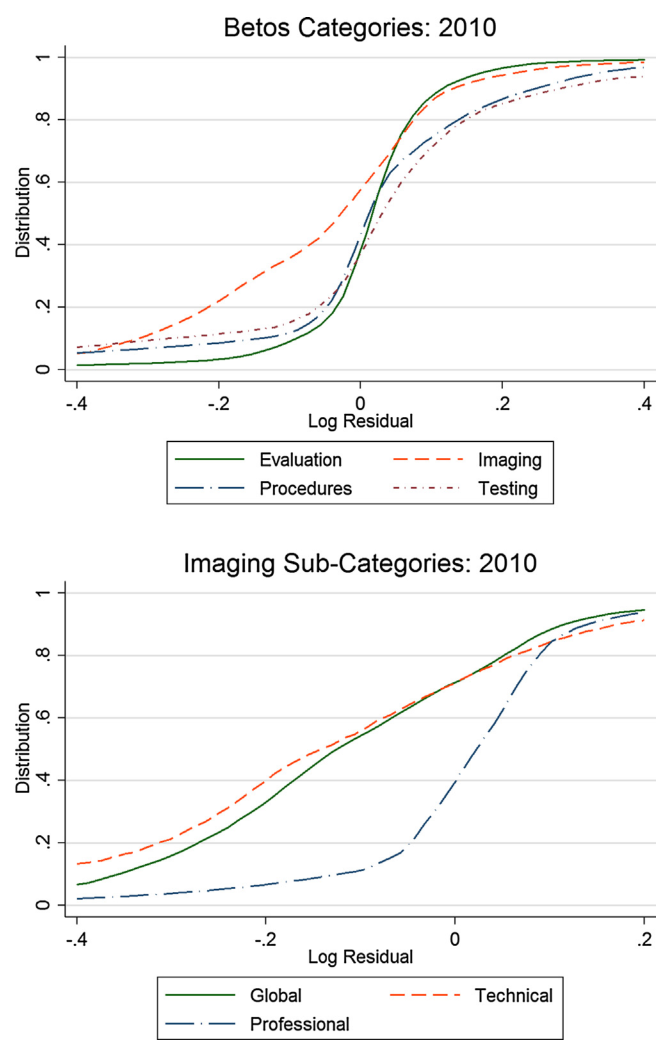

Fig. 5. Deviations from Medicare benchmark by service category. *Note*: The figure presents the distributions of empirical residuals $\hat{e}_{g,j}$ from estimates of Eq. (14). The distribution of residuals is shown within either broad Betos categories (Panel A), or within the subcategories of Imaging (Panel B). The distributions are smoothed using a local linear regression, with an Epanechnikov kernel and a bandwidth of 0.01. *Source: Authors' calculations using claims data from BCBS.*

[27] Betos categories are aggregates of related services, defined in footnote [24]. To be precise, these means are $\bar{\hat{e}}_{g,j} = \frac{1}{N_b}\sum_{j \in b} \hat{e}_{g,j}$, where each Betos group b comprises N_b claims for all services $j \in b$ in that group.

[28] In the Testing category the vast majority of residuals are negative, with the exception of one of the more common tests, which has a large and positive average residual. Recall from Section 3, however, that Testing is the one category with significant missing data problems.

for capital-intensive services. This is consistent with deviating from Medicare with an eye towards more closely targeting either marginal costs or medical value.[29]

7. Conclusion

This paper uses physician payments from a large private insurer as a window into the structure of private insurers' contracts with physicians. Using two empirical strategies, we show that payments exhibit pervasive linkages to the relative payment model of the federal Medicare program. We find that three quarters of the services and 55 percent of the spending we analyze are benchmarked to Medicare. The analysis thus reveals that benchmarking to Medicare's relative rate structure is complemented by substantial customization.

In the contracts we analyze, customization is most prevalent when its value appears likely to be highest. Deviations from benchmarking occur disproportionately in contracts with large physician groups, where significant value may be at stake. Deviations also involve reductions in payments for diagnostic imaging services, a category of care for which many academics and policy makers believe marginal benefits are low relative to costs (Winter and Ray, 2008; MedPAC, 2011). The benchmarking phenomenon is strongest in payments for services where average-cost reimbursements will be most aligned with marginal costs, such as labor-intensive primary care services. When contracts deviate from Medicare, the direction of payment adjustments would tend to encourage the provision of primary care and discourage care for which overutilization is a more widespread concern. The results are thus suggestive of effort to improve the payment structure through customization.

A number of factors, including contracting frictions, market power, the information content of Medicare's fee schedule, and the burden of regulatory compliance, may contribute to the contracting patterns we observe. Disentangling these forces is an attractive goal for future research. Further insights may be gained by analyzing the relationship between hospital contracting arrangements and Medicare's inpatient payment model. Variations in the degree of contract complexity, regulatory burdens, and both the size and market power of the parties involved may generate economically interesting variations in the connection between Medicare's inpatient payment model and private insurers' hospital payments. Early work along these lines includes an analysis by Baker et al. (2016), who study the degree to which hospitals are paid retrospectively (rather than prospectively) by private insurers. Their results echo some of our findings here, in that Medicare-style prospective payments are more pervasive in similar economic conditions to those where we find more benchmarking.[30]

Regardless of the explanation, the Medicare-benchmarking phenomenon implies that many inefficiencies in Medicare's reimbursements spill over into private fee schedules. By extension, the value of improvements to public payment systems may ripple through private contracts in addition to improving the performance of Medicare itself. At the same time, the customization we observe would tend to curb what policy analysts regard as Medicare's greatest inefficiencies. Both public and private players thus appear to have important roles in the process of fee schedule improvement and payment system reform.

[29] Appendix D.2 shows that the changes in BCBS prices due to Medicare benchmarking matter in practice for the care that physicians supply to BCBS patients.

[30] Baker et al. (2016) analyze the extent to which private hospital payments reflect the Medicare program's primarily "prospective" approach. They find that private payments are more likely to be prospective when there is significant competition across the hospitals within a region, when a hospital's patients are covered primarily through managed-care arrangements, and when a large fraction of the hospital's patients are covered by Medicare.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.jhealeco.2017.07.001.

References

Alpert, A., Duggan, M., Hellerstein, J.K., 2013. Perverse reverse price competition: average wholesale prices and Medicaid pharmaceutical spending. J. Public Econ. 103, 44–62.
Anderson, E., Jaimovich, N., Simester, D., 2015. Price stickiness: empirical evidence of the menu cost channel. Rev. Econ. Stat. 97 (October (4)), 813–826.
Baker, L.C., Bundorf, M.K, Royalty, A.B., Levin, Z., 2014. Physician practice competition and prices paid by private insurers for office visits. J. Am. Med. Assoc. 312 (16), 1653–1662.
Baker, L., Bundorf, M.K, Devlin, A., Kessler, D.P., 2016. "Why Dont Commercial Health Plans Use Prospective Payment?" Working Paper No. 22709. National Bureau of Economic Research.
Berenson, R., Holahan, J., 1990, December. Using A New Type-of-Service Classification System to Examine the Growth in Medicare Physician Expenditures, 1985–1988. NTIS Publication PB91-188599, Urban Institute Health Policy Center.
Besley, T.J., 1988. Optimal reimbursement health insurance and the theory of Ramsey taxation. J. Health Econ. 7 (4), 321–336.
Bils, M., Klenow, P.J., 2004. Some evidence on the importance of sticky prices. J. Polit. Econ. 112 (October (5)), 947–985.
Blue Cross Blue Shield of Texas, 2008. Fee schedule update. Blue Rev. (1), 8.
Blue Cross Blue Shield of Texas, 2009. Fee schedule updates effective in August. Blue Rev. (2), 5.
Blue Cross Blue Shield of Texas, 2010. Fee schedule updates. Blue Rev. (3), 15.
Blue Cross Blue Shield of Texas, 2011. Fee schedule update. Blue Rev. (6), 10.
Borges, W., 2003. Mo' Better Blues. Texas Med. 99 (April (4)).
Cameron, A.C., Gelbach, J.B., Miller, D.L., 2011. Robust inference with multiway clustering. J. Bus. Econ. Stat. 29 (2), 238–249.
Center for Studying Health System Change, 1999. "Physician Survey Summary File: User's Guide and Codebook," 1999. Community Tracking Study Physician Survey, 1996-1997 (ICPSR 2597)., http://dx.doi.org/10.3886/ICPSR02597.v3.
Center for Studying Health System Change, 2006. "2004-05 Physician Survey Summary File: User's Guide and Codebook," 2006. Community Tracking Study Physician Survey, 2004-2005 (ICPSR 4584)., http://dx.doi.org/10.3886/ICPSR04584.v2.
Christiano, L.J., Eichenbaum, M., Evans, C.L., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. J. Polit. Econ. 113 (1), 1–45.
Clemens, J., Gottlieb, J.D., 2014. Do physicians' financial incentives affect medical treatment and patient health? Am. Econ. Rev. 104 (April (4)), 1320–1349.
Clemens, J., Gottlieb, J.D., 2017. In the shadow of a giant: Medicare's influence on private payment systems. J. Polit. Econ. 125 (February (1)), 1–39.
Clemens, J., Gottlieb, J.D., Shapiro, A.H., 2014, September. How much do Medicare payment cuts reduce inflation? FRBSF Econ. Lett., 2014-28.
Clemens, J., Gottlieb, J.D., Shapiro, A.H., 2016, May. Medicare payment cuts continue to restrain inflation. FRBSF Econ. Lett., 2016-15.
Cooper, Z., Craig, S.V., Gaynor, M., Van Reenen, J., 2015. The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured," Working Paper No. 21815. National Bureau of Economic Research December.
Duggan, M., Scott Morton, F., 2006, February. The distortionary effects of government procurement: evidence for medicaid prescription drug purchasing. Q. J. Econ. 121, 1–30.
Dunn, A., Shapiro, A.H., 2014. Do physicians possess market power? J. Law Econ. 57 (February (1)), 159–193.
Einav, L., Finkelstein, A., Polyakova, M., 2016, May. Private Provision of Social Insurance: Drug-Specific Price Elasticities and Cost Sharing in Medicare Part D," Working Paper No. 22277. National Bureau of Economic Research.
Ellis, R.P., McGuire, T.G., 1986. Provider behavior under prospective reimbursement: cost sharing and supply. J. Health Econ. 5 (June (2)), 129–152.
Feldstein, M.S., 1973. The welfare loss of excess health insurance. J. Polit. Econ. 81 (March–April (2)), 251–280.
Gesme, D.H., Wiseman, M., 2010. How to Negotiate With Health Care Plans. J. Oncol. Pract. 6 (July (4)), 220–222.
Government Accountability Office, 2015, May. Medicare Physician Payment Rates: Better Data and Greater Transparency Could Improve Accuracy," Report GAO-15-434. GAO.
Kaiser Family Foundation, 2015. State Health Facts, Available from: http://kff.org/state-category/health-insurance-managed-care/ (accessed 17.04.17).
Kleiner, S.A., White, W.D., Lyons, S., 2015. Market power and provider consolidation in physician markets. Int. J. Health Econ. Manag. 15 (March (1)), 99–126.
Kleven, H.J., 2016. Bunching. Annu. Rev. Econ. 8 (forthcoming).
McClellan, M., 2011. Reforming payments to healthcare providers: the key to slowing healthcare cost growth while improving quality? J. Econ. Perspect. 25 (Spring (2)), 69–92.

MedPAC, 2011. Report to the Congress: Medicare and the Health Care Delivery System," Technical Report, Medicare Payment Advisory Commission June 2011, Available from: http://www.medpac.gov/documents/reports/Jun11_EntireReport.pdf (accessed 07.11.15).

Mertz, G.J., 2004. Can you negotiate better reimbursement? Fam. Pract. Manag. 11 (October (9)), 31–34.

Nakamura, E., Steinsson, J., 2008. Five facts about prices: a reevaluation of menu cost models. Q. J. Econ. 123 (November (4)), 1415–1464.

Nandedkar, M., 2011, September. Key principles for negotiating with insurers, deciding whether to opt in or out of Medicare, Available from: http://www.aad.org/dw/monthly/2011/september/key-principles-for-negotiating-with-insurers-deciding-whether-to-opt-in-or-out-of-medicare (accessed 21.08.13).

Smets, F., Wouters, R., 2003. An estimated dynamic stochastic general equilibrium model of the euro area. J. Eur. Econ. Assoc. 1 (September (5)), 1123–1175.

Smets, F., Wouters, R., 2007. Shocks and frictions in us business cycles: a Bayesian DSGE approach. Am. Econ. Rev. 97 (June (3)), 586–606.

Starc, A., Town, R.J., 2015. Internalizing Behavioral Externalities: Benefit Integration in Health Insurance," Working Paper No. 21783. National Bureau of Economic Research December.

The United Food, Commercial Workers Union, 2002. UFCW 789 Grocery Contract, Available from: http://www.irle.berkeley.edu/library/pdf/0147.pdf.

White, C., 2013. Contrary to cost-shift theory, lower Medicare hospital payment rates for inpatient care lead to lower private payment rates. Health Affairs 32 (May (5)), 935–943.

Winter, Ariel, Ray, Nancy, 2008. "Paying Accurately For Imaging Services In Medicare,". Health Affairs 27 (6), 1479–1490.

Zuvekas, S.H., Cohen, J.W., 2016. Fee-for-service, while much maligned, remains the dominant payment method for physician visits. Health Affairs 35 (3), 411–414.

Further reading

Brekke, Kurt R., Tor Helge Holmas, Karin Monstad, and Odd Rune Straume, "Do Treatment Decisions Depend on Physicians' Financial Incentives?," Working Paper No. 07/2015, NIPE July 2015. Available online at http://www.nipe.eeg.uminho.pt/Uploads/WP_2015/NIPE_WP_07_2015.pdf (accessed December 9, 2015).

Carroll, John, "How Doctors Are Paid Now, And Why It Has to Change," Managed Care, December 2007, (12).

Clemens, Jeffrey and Joshua D. Gottlieb, "In the Shadow of a Giant: Medicare's Influence on Private Payment Systems," Working Paper No. 19503, National Bureau of Economic Research October 2013.

Cutler, David M. and Dan P. Ly, "The (Paper)Work of Medicine: Understanding International Medical Costs," Journal of Economic Perspectives, 2011, 25 (2), 3-25.

Dunn, A. Shapiro, A.H. "Physician Payments Under Health Care Reform," Journal of Health Economics, 2015, 39, 89-105.

Geruso, Michael, Timothy Layton, and Jacob Wallace, "Are All Managed Care Plans Created Equal? Evidence from Random Assignment to Medicaid HMOs," 2016. UT Austin, mimeo.

Gruber, Jon, John Kim, and Dina Mayzlina, "Physician Fees and Procedure Intensity: the Case of Cesarean Delivery," Journal of Health Economics, August 1999, 18 (4), 473-490.

Jacobson, Mireille, Craig C. Earle, Mary Price, and Joseph P. Newhouse, "How Medicare's Payment Cuts For Cancer Chemotherapy Drugs Changed Patterns Of Treatment," Health Affairs, July 2010, 29, 1391-1399.

Mahoney, Neale, "Bankruptcy as Implicit Health Insurance," American Economic Review, February 2015, 105 (2), 710-746.

Olea, Jose Luis Montiel and Carolin Pflueger, "A Robust Test for Weak Instruments," Journal of Business and Economic Statistics, July 2013, 31 (3), 358-369.