This work is distributed as a Discussion Paper by the

STANFORD INSTITUTE FOR ECONOMIC POLICY RESEARCH

SIEPR Discussion Paper No. 16-004

The Anatomy of Physician Payments: Contracting Subject to Complexity

By

Jeffrey Clemens , Joshua D.Gottlieb, Tımea Laura Molnar

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

The Anatomy of Physician Payments: Contracting Subject to Complexity†

Jeffrey Clemens
UCSD and NBER

Joshua D. Gottlieb
UBC and NBER

Tímea Laura Molnár*
UBC

February 9, 2016

Abstract

Why do private insurers closely link their physician payment rates to the Medicare fee schedule despite its well-known limitations? We ask to what extent this relationship reflects the use of Medicare's relative price menu as a benchmark, in order to reduce transaction costs in a complex pricing environment. We analyze 91 million claims from a large private insurer, which represent $7.8 billion in spending over four years. We estimate that 75 percent of services, accounting for 55 percent of spending, are benchmarked to Medicare's relative prices. The Medicare-benchmarked share is higher for services provided by small physician groups. It is lower for capital-intensive treatment categories, for which Medicare's average-cost reimbursements deviate most from marginal cost. When the insurer deviates from Medicare's relative prices, it adjusts towards the marginal costs of treatment. Our results suggest that providers and private insurers coordinate around Medicare's menu of relative payments for simplicity, but innovate when the value of doing so is likely highest.

---

*Clemens: jeffclemens@ucsd.edu, Gottlieb: Joshua.Gottlieb@ubc.ca, Molnár: molnar.timea.laura@alumni.ubc.ca. 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 while working on this paper. We thank SSHRC, Jon Skinner and the Dartmouth Institute for support, and Victor Saldarriaga for excellent research assistance. Finally, we thank Leila Agha, Jeff Emerson, R.B. Harris, David Laibson, Neale Mahoney, and workshop participants at Stanford Health Policy, University of Pennsylvania—Leonard Davis Institute, UBC, UC San Diego, USC, iHEA, the Junior Health Economics Summit, and the UBC Public Finance Reading Group for helpful comments.

In an exclusively public health care system, payment rates for medical providers are typically set through an administrative mechanism that applies to the entire market (Laugesen and Glied, 2011). In a multi-payer system, physicians and private insurers must agree on payments through private negotiations. This paper looks into the black box of the prices embedded in physician-insurer contracts. We analyze how these private physician payments are shaped by payment rates set by Medicare, the public health insurer for the elderly and disabled. Our results suggest that providers and private insurers coordinate around Medicare's menu of relative payments for simplicity, but innovate when the value of doing so is highest.

One of the U.S. health system's most distinctive features is the prominent position of private insurers. Despite the public sector's substantial role, private insurers directly finance roughly $1 trillion of medical spending, or one third of the total (OECD, 2015).[1] High system-wide spending, coupled with middling health outcomes, raises questions about the costs and benefits of this multi-payer approach.

The consequences of public and private care financing depend on many factors, one of which is the design of the payment systems that intermediate between patients and their health care providers. Payment systems can shape the health system's efficiency by affecting the composition of care offered (Gruber, Kim and Mayzlina, 1999; Jacobson, Earle, Price and Newhouse, 2010; Clemens and Gottlieb, 2014). Because services may differ substantially in their cost-benefit ratios (Chandra and Skinner, 2012), changes in these incentives can have first order welfare importance. If the presence of private payers generates innovation in payment system design, this innovation could be an important benefit of the multi-payer system. On the other hand, the multi-payer approach's fragmentation drives considerable administrative expense (Cutler and Ly, 2011).

[1]This represents almost half of health spending via traditional insurance plans, since it excludes out of pocket costs (12 percent of total health expenditures), research and capital investments (6 percent), public health (3 percent), as well as workers' compensation and other specified health programs.

In the U.S. public sector, the federal Medicare program compensates physicians and outpatient providers through a system known as the Resource-Based Relative Value Scale (RBRVS). The RBRVS has two key features. First, it is a remarkably detailed, fee-for-service payment model, with 13,000 distinct service codes defined. Physicians submit bills for each instance in which they provide one of these services. The RBRVS assigns each service a certain number of "relative value units" (RVUs), which determine the payment for that service. Second, these relative values are legislatively required to reflect variations in average cost, without reference to medical value. This procurement model thus has little capacity to steer care provision towards effective—let alone cost-effective—services. 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. Nevertheless, practitioners and policy makers regularly observe that private insurers' payment models lean heavily on Medicare's approach to paying for care (Borges, 2003; Gesme and Wiseman, 2010).

Our analysis has three major goals. First, we estimate the pervasiveness of links between Medicare's fee schedule and physician payments from a single large insurer. Second, we analyze how the strength of these links varies across categories of health care services and types of physician groups. Third, we measure the direction and magnitude of the private insurer's deviations from Medicare's reimbursement rates. Our results yield insights into both the extent of Medicare's influence and the economic factors underlying the insurer's approach to contracting.

We use insurance claims data from Blue Cross Blue Shield of Texas (BCBS). These data have two key features for our purposes. First, they allow us to examine the service-level payments associated with unique insurer-physician group pairings. Second, they allow us to longitudinally track these payments at high frequency.[2]

[2]Our data represent around $2 billion in annual spending, which is approximately 1 percent of national

We develop two methods to estimate the pervasiveness of payments linked directly to Medicare's RVUs in the BCBS data. We first make a straightforward observation about payments in the cross section. The payment for any service can be described as the product of its Medicare-allotted number of RVUs and a scaling of dollars per RVU. We term this scaling the "implied conversion factor" (ICF). When an ICF is shared across many services within an insurer-physician pair, we infer that the common mark-up is specified in the contract. As a baseline, we infer that every claim whose ICF accounts for at least 10 percent of the provider's BCBS payments is contractually linked to Medicare. Under this assumption, around three quarters of BCBS's claims, accounting for two thirds of spending, follow the Medicare benchmark.

Second, we measure the extent to which updates to Medicare's fees pass through to BCBS's payments. The analysis exploits institutional detail about the precise dates on which BCBS implements Medicare's annual updates. This fine-grained timing allows us to infer the share of BCBS's payments linked to Medicare without having our estimates confounded by long-run technological changes or active contract renegotiations. This method again implies that around three quarters of BCBS's payments are linked to Medicare.

Why do private insurers rely on Medicare for most, but not all, relative payment rates? We propose that physician contracts are written to manage the tension between gains from fine-tuning payments and costs from making contracts complex. This proposition—that BCBS draws on Medicare for the purpose of contract simplification, while strategically adapting its contracts where the value of adaptations is highest—predicts heterogeneity in the frequency of benchmarking across physician groups and service categories. First, the benefits of fine-tuning payments will tend to be higher for contracts with relatively large physician practices. Second, since Medicare's average cost approach has greater difficulty managing the payments for capital- than for labor-intensive services, the benefits of fine-

spending on physician and clinical services from private health insurers (CMS, 2011).

tuning payments will tend to be greater for the former than the latter.

Using both the cross-sectional and update-based frameworks, we test these predictions. Looking across physician groups, we find that payments to relatively large firms are less tightly linked to Medicare than payments to small firms. Our estimates suggest that payments for nearly 90 percent of services provided by the smallest firms (representing 80 percent of their spending) are linked to Medicare's relative values. The same is true of 60 percent of services from firms with total BCBS billing exceeding $1 million per year.

Looking across service categories, we find that payments are more closely linked to Medicare's relative values for labor-intensive services, like standard office visits, than for capital-intensive services, like diagnostic imaging. Payments for roughly 85 percent of evaluation and management services, but only 55 percent of imaging services, are directly linked to Medicare's menu.

Within diagnostic 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 BCBS payments for interpretation are far more tightly linked to Medicare rates than are its payments for the image itself.

We also show that BCBS's adjustments work to narrow likely gaps between marginal costs and Medicare's average-cost payments. Specifically, we find that payments for labor-intensive services tend to be adjusted up while payments for capital-intensive services tend to be adjusted down. This supports the view that BCBS aims to improve on Medicare's average-cost reimbursements while managing the complexity of its payment system.

When we conduct a comparable analysis on payments to out-of-network physicians we find much weaker links to Medicare's payments. The out-of-network payments, by definition, are for providers who have not reached an agreement with the insurer on reimbursement rates. This suggests that the stronger links for *in-network* prices reflect active efforts to negotiate around a simplified payment schedule.

Finally, we use our estimates to draw inferences about the cost required to shift from simple to complex contracts. Depending on the level of inefficiency that one assumes is embedded in the Medicare fee schedule, we calibrate the negotiating costs that would rationalize the share of Medicare benchmarking we see. We find, for example, that if potential efficiency gains for BCBS are 1 percent of a physician group's billings, then contracting costs of $3,000 per group would explain our results. A modest reduction in contracting costs could generate $1 billion of efficiency gains nationally.

Our findings connect to research on health care payment systems and to two more general literatures. A growing literature demonstrates significant spillovers from Medicare payment policies into the private sector. Duggan and Scott Morton (2006) and Alpert, Duggan and Hellerstein (2013) show that pharmaceutical markets respond to public sector payment idiosyncrasies. White (2013) finds a sizable positive relationship between Medicare and private hospital pricing, as do Clemens and Gottlieb (forthcoming) in the outpatient context. A key limitation of the existing literature is that prior work has not been able to link negotiated payment rates to specific physician-insurer pairs. By incoroprating such data, we make two novel contributions here. First, we show that Medicare exerts influence over nominally independent private insurers directly through those insurers' adoption of Medicare's rate structure. Second, we are able to investigate exactly when the parties deviate from Medicare's basic structure and exmaine their efforts to innovate towards efficiency.

Second, we contribute to the literature documenting how boundedly rational agents navigate complex environments. Work in behavioral economics (DellaVigna, 2009; Gabaix, 2014), macroeconomics (Sims, 2003), public finance (Chetty, Looney and Kroft, 2009; Abeler and Jäger, 2015), persuasion (Mullainathan, Schwartzstein and Shleifer, 2008), and beyond has considered how bounded rationality and computational costs shape agents' decision-making. When firms interact with each other—in our case, insurers and physician groups—little is known about how they reduce the dimensionality of the complex environments they face.[3] Benchmarking payments to Medicare's relative rates is an intriguing way to simplify the physician contracting problem. By calibrating the negotiation costs implied by our results, we provide insight into the likely magnitude of the contracting frictions relevant to our context.

Third, nominal price rigidities are central to much analysis of business cycles and monetary policy (Clarida, Galí and Gertler, 1999), and the specific form these rigidities take has significant influence on resulting dynamics (Mankiw and Reis, 2002). Detailed studies of price microdata have found that prices for services adjust less frequently than in other sectors (Nakamura and Steinsson, 2008). This is particularly true in medical care, where Bils and Klenow (2004) find that the average price persists for eleven months. Our analysis provides insight into why this is the case. With some exceptions, Medicare's payment updates occur annually. We find that private contracts incorporate Medicare's changes by updating with a similar frequency. Consistent with Anderson, Jaimovich and Simester's (forthcoming) evidence from retail, the complexity of physician contracting may explain both the long duration of these prices and the public-private linkages we estimate. Given the health sector's size, Medicare's direct and indirect influences can meaningfully affect overall inflation (Clemens, Gottlieb and Shapiro, 2014).

This paper proceeds as follows. In section 1, we describe Medicare's pricing institutions. Section 2 presents an institutionally-informed model of physician-insurer contracting. Section 3 introduces our claims data. Section 4 presents our first analysis, which investigates the cross-sectional relationship between private reimbursements and Medicare's fee schedule. In section 5, we derive the empirical specifications through which we estimate the Medicare-benchmarked share of payments using updates to Medicare's relative prices.

[3] In a different health care context, Grennan and Swanson (2015) find that hospitals are more likely to conduct active negotiations for the supplies on which they spend the most.

Section 6 presents our results from this analysis, including heterogeneity across physician groups and service categories. In section 7 we examine the direction in which BCBS adjusts its payments when they deviate from the benchmark. Section 8 examines supply responses and calibrates the magnitude of contracting frictions. Section 9 concludes.

1 Medical Pricing Institutions

Public and private payments for health care services are set through very different mechanisms. Medicare reimbursements are set to administratively determined measures of the resource costs of providing care. For patients with private health insurance, providers' reimbursements are determined through negotiations between the insurers and providers. Section 1.1 discusses key features of Medicare's administrative pricing mechanisms. Section 1.2 presents institutional details on contracting between providers and private insurers.

1.1 Medicare Price Determination[4]

Since 1992, Medicare has paid physicians and other outpatient providers through a system of centrally administered prices, based on a national fee schedule. This fee schedule, known as the Resource-Based Relative Value Scale (RBRVS), assigns an allocation of Relative Value Units (RVUs) to each of 13,000 distinct health care services. The RVUs associated with service j are legislatively bound to measure the resources required to provide that service. Medicare recognizes that goods and services have different production costs in different parts of the country; Congress mandates price adjustments, called the Geographic Adjustment Factor (GAF), to offset these differences in input costs. For service j, supplied by a provider

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

in payment area i, the provider's fee is approximately:

$$\text{Reimbursement Rate}_{i,j,t} = \text{Conversion Factor}_t \times \text{Geographic Adjustment Factor}_{i,t} \times \text{Relative Value Units}_{j,t}. \quad (1)$$

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.

Payments across services vary primarily according to their assigned number of Relative Value Units (RVUs). RVUs are constant across areas while varying across services. The RVUs associated with each service are updated on a rolling basis to account for technological and regulatory changes that alter their resource intensity. We exploit these changes in one of our empirical strategies, which we introduce in section 5.

1.2 Private Sector Price Setting

U.S. private sector health care prices are set through negotiations between providers and private insurers.[5] 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.[6] A common feature of physician contracts, central to both our theoretical and empirical analyses, is a form of benchmarking to Medicare.

Practitioners regularly emphasize that Medicare's administrative pricing menu features prominently in private insurers' contracts. Both industry-wide and BCBS-specific sources provide institutional detail that illuminates the Medicare fee schedule's role. Newsletters

[5]Some exceptions apply to this statement. For instance, private insurers' hospital payment rates in Maryland are set by a state government board.

[6]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, Bundorf, Royalty and Levin, 2014; Kleiner, White and Lyons, 2015; Clemens and Gottlieb, forthcoming).

that insurers distribute to participating providers, both in Texas and elsewhere, frequently draw explicit links between Medicare's maximum allowable charges and the insurer's fee schedule. Policies often take the form that reimbursement rates are linked to Medicare unless the insurer's contract specifies otherwise. Our empirical work examines when and why this occurs. We measure how often exceptions apply, and whether BCBS's exceptions occur systematically in cases when we would expect the cost of the Medicare menu's inefficiencies to be particularly large.

Importantly, the relative value scale itself does not determine an absolute price level. As in Medicare, realized private reimbursements involve RVUs scaled by "conversion factors," which converts RVUs into dollars. These conversion factors are key subjects of negotiation.

Practitioners describe two modes of negotiation between providers and private insurers. Insurance carriers typically offer small provider groups contracts based on fixed fee schedules. Whether the schedule is copied directly from Medicare or modified by the insurer, the parties then negotiate a constant markup over these rates (Nandedkar, 2011; Gesme and Wiseman, 2010; Mertz, 2004). In contrast, insurers are said to negotiate in more detail with hospitals and large provider groups. The model below examines when each bargaining approach would be efficient and what each means for the welfare consequences of Medicare payment reforms.

2 Conceptual Framework

We sketch a model of physician reimbursement rates that can be benchmarked to Medicare or unconstrained. Physicians and insurers can use Medicare's payments as a default relative price schedule, so that reimbursements are simply a markup over Medicare's rates.[7] Adopting this default has costs if Medicare's relative payments are suboptimal, in a sense developed below. It may nonetheless be efficient to rely on this default due to negotiation

[7]Medicare's position as the single-largest payer for health care services further reinforces its relevance as a setter of default prices. Practitioners describe the offers made by insurers to sole practitioners, for example, as being take-it-or-leave it, scalar mark-ups (or occasionally slight mark-downs) of Part B prices.

and coordination costs (Cutler and Ly, 2011).[8]

Consider an insurer that purchases two types of medical services, indexed by $j \in \{1, 2\}$, for treating its enrollees. We abstract from the physician-insurer bargaining process and assume that the insurer sets prices with full knowledge of the aggregate supply curve for each type of care. Let r_j denote the reimbursement rate that the insurer pays to physicians for providing service j, and let r_j^M be the corresponding Medicare rate. For extreme analytical simplicity, assume that the physician market supplies care to the insurer's patients according to the aggregate supply functions $s_1(r_1) = \alpha r_1$ and $s_2(r_2) = \beta r_2$, where r_j is the reimbursement rate for service j and $\alpha, \beta > 0$. If the true price-setting process is not so simple—say, if physicians are not price-takers—the model's main ideas still hold. In that case, they strive to reach a pricing agreement that maximizes joint surplus. We would simply view prices as jointly determined and negotiating costs as those incurred by both parties.

We assume that the insurer aims to minimize its medical expenses while keeping patients, or their employers, satisfied with the insurance product. This latter constraint requires that the insurer provide enough care to achieve the patient's reservation value $\overline{u}$. We assume the patients have extremely simple preferences over medical care, captured by $u(q_1, q_2) = aq_1 + bq_2$ where q_j is the quantity of service j supplied to a representative patient.

We will consider two methods of reimbursement rate determination, and then allow the insurer to choose between them. In the first case, the insurer is constrained to set reimbursements as scalar markups over Medicare rates. Let φ represent this markup, so the benchmarked payment for service j would be φr_j^M. We then obtain the following result, whose proof is in Appendix A.

Result 1 (Reimbursements Benchmarked to Medicare)**.** *When the insurer is constrained*

[8] Providers themselves may find deviating from Medicare's menu costly due to increases in the non-trivial administrative expenses associated with billing (Cutler and Ly, 2011). Regulations requiring insurers to pay sufficiently to ensure access to "medically necessary" services may also contribute to such a role for public players in these markets.

to follow Medicare's relative prices, the markup will be given by $\varphi = \dfrac{\overline{u}}{\alpha a r_1^M + \beta b r_2^M}$*. Total medical expenditures will be* $\hat{E} \equiv \varphi^2[\alpha(r_1^M)^2 + \beta(r_2^M)^2]$.

In this case, the insurer only chooses one pricing parameter: the markup φ over Medicare. Result 1 shows that this markup is increasing in our proxy for patients' demand, their reservation value $\overline{u}$. As $\overline{u}$ increases, insurers must increase physician reimbursements in order to induce the supply responses required to satisfy higher-$\overline{u}$ patients.

Next consider the insurer's behavior when relative prices are unconstrained. In this situation, the insurer sets physician reimbursements separately for each service, again aiming to minimize medical expenditures subject to the constraint that $u(q_1, q_2) \geq \overline{u}$.

Result 2 (Reimbursements When Unconstrained)**.** *When the insurer is unconstrained, reimbursement rates satisfy* $\dfrac{r_2^*}{r_1^*} = \dfrac{b}{a}$*. Medical expenditures are* $E^* \equiv \dfrac{\overline{u}^2}{\alpha a^2 + \beta b^2}$*. These expenses are weakly lower than* $\hat{E}$ *from Result 1, with equality occurring when* $\dfrac{r_2^M}{r_1^M} = \dfrac{b}{a}$*. The discrepancy between* E^* *and* $\hat{E}$ *is increasing in* $\left|\dfrac{r_2^M}{r_1^M} - \dfrac{b}{a}\right|$.

This result shows that the insurer can reduce expenditures, while maintaining patient satisfaction, whenever Medicare's reimbursement ratio differs from the ratio the insurer would prefer. Since the insurer's optimal pricing accounts for patients' relative preferences over the two services, while Medicare's reimbursements may not, relying on Medicare's payment ratio can push the insurer inefficiently far up the supply curve for one of the services. By remedying this inefficiency, the unconstrained payments can save money while maintaining patient satisfaction. The more Medicare's payment ratio deviates from the efficient one, the costlier this inefficiency is for the insurer.

We now allow the insurer to choose between the two pricing regimes. Let $\theta = \frac{r_2^M}{r_1^M}$ be the ratio of Medicare payments for the two services. If the insurer adopts this ratio, as we assumed in Result 1, it incurs no additional cost. If it chooses a different ratio, $\frac{r_2}{r_1} \neq \theta$, it incurs a fixed cost c due to the added complexity or additional negotiations required.

Result 3 (Choice of Benchmarking)**.** *Let ξ denote the insurer's savings from abandoning Medicare's payment ratio. The insurer will deviate from this ratio when $\xi > c$.*

These savings ξ are proportional to $\overline{u}^2$, and are increasing in the difference between the efficient reimbursement ratio and that implied by Medicare's payment rates, $\left|\frac{r_2^M}{r_1^M} - \frac{b}{a}\right|$. Conditional on the ratio $\frac{\beta}{\alpha}$, ξ is decreasing in the sensitivity of supply to reimbursement rates (α or β). Conditional on the ratio $\frac{b}{a}$, ξ is increasing in the amount of care required to achieve utility level $\overline{u}$ (decreasing in a or b).

This result shows that it is more worthwhile for the insurer to abandon Medicare's relative pricing, and pay the costs necessary to set prices independently, in two sets of scenarios. First, the insurer is more prone to abandon benchmarking when Medicare's default reimbursements deviate more substantially from the insurer's preferred relative prices. When the Medicare relative prices are farther from the insurer's unconditional optimum, the insurer has to spend ever more to achieve the same patient satisfaction.

Second, the insurer is more prone to abandon benchmarking when there is more money at stake. This shows up in Result 3 in three ways. First, the insurer has to spend more—both through higher prices and procuring more services—in order to provide a higher utility level $\overline{u}$. Second, when supply is less sensitive to reimbursement rates, higher payments are needed to achieve $\overline{u}$—and more so when Medicare-benchmarked prices increase the distortions. Third, when the parameters a and b in the utility function are lower, holding constant $\overline{u}$, it takes more care to achieve the requisite patient utility. Again, this implies higher costs when the insurer's preferred relative payments differ from Medicare's.

In practice, this model implies that there may be welfare gains available if the insurer and physician negotiate service- or bundle-specific reimbursement rates. Medicare's fee schedule may have its own inefficiencies, in terms of the care it encourages or division of resources it induces. Consequently, the overall quality of the health insurance product, relative to its costs, can potentially be increased by abandoning Medicare's reimbursement ratio.

3 Medical Pricing Data

We analyze health care price setting in the context of claims processed by a single large insurer, Blue Cross Blue Shield of Texas (BCBS). The claims database we analyze covers the universe of BCBS's payments for outpatient care in 2008–2011. For each claim, the database provides information on the service provided, location, physician, physician group, and BCBS's payment to that group. Our analysis sample restricts this universe along several dimensions. For example, the full 2010 dataset contains 57,613,494 claim lines and $4.29 billion in spending. We clean the data as described in Appendix B.1, which initially leaves us with 44,055,829 service lines and $2.63 billion of spending. This initial cut eliminates payments made to out-of-network physicians, who have not reached a negotiated agreement with BCBS on reimbursement rates. We will subsequently examine this segment of the data separately.

In order for private insurers to benchmark prices to Medicare, at a minimum they would need to use Medicare's billing codes. We thus merge the remaining claims with Medicare billing codes, which provides an upper bound on the potential benchmarking. This merge only loses notable portions of one broad spending category, namely laboratory tests, for which both Medicare and BCBS frequently base payments on non-standard codes. We retain over 97 percent of claims for evaluation and management, diagnostic imaging, and surgical services. The final analysis sample in 2009 includes 3,821 unique HCPCS codes, which comprise 23,933,577 service lines and $2.05 billion of spending.[9]

The claims data further 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

[9]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.

care each group provides to BCBS patients in a given year. Our final dataset includes care provided by over 80,000 physician groups. Table 1 presents summary statistics on the physician groups in our final sample.

4 Private Benchmarking to Medicare in the Cross-Section

4.1 Measuring Implied Conversion Factors in Claims Data

Our first look at the relationship between private and Medicare pricing exploits a straightforward insight: when many payments to a given physician group share a common mark-up, their payments are likely linked contractually to Medicare's relative rates. As made clear below, this claim's strength depends on the precision with which markups are rounded. Markups rounded to the nearest 2 cents per RVU, our baseline threshold, are unlikely to coincide by chance.

To flesh out our approach, we start by simplifying the Medicare payment formula from equation (1). For any one physician group, the geographic adjustment is a constant and can thus be thought of as part of the Conversion Factor.[10] Letting $P_{c,j,t}$ denote the reimbursement rate for claim c for service j in year t, equation (1) simplifies to:

$$P_{c,j,t} = \text{Conversion Factor}_t \times RVU_{j,t}. \tag{2}$$

Dividing the payment $P_{c,j,t}$ by Medicare's RVU allotment for service j, we obtain:

$$ICF_{c,j,t} = \frac{P_{c,j,t}}{RVU_{j,t}}. \tag{3}$$

This equation defines an "implied conversion factor" (ICF)—the conversion factor that would

[10]Medicare's geographic adjustments are actually slightly more complicated, but this is a close approximation. See Clemens and Gottlieb (2014) for more details.

rationalize a payment of $P_{c,j,t}$ in a Medicare-benchmarked contract. Taking logs of equation (2) reveals that this pricing scheme implies a 1 for 1 relationship between log RVUs and the log of $P_{c,j,t}$. Since we observe $P_{c,j,t}$ in the claims data and CMS publishes its RVU allocations, investigating the prevalence of common ICFs is straightforward.

Simply computing an ICF does not tell us whether claim c was actually priced according to equation (2). To gauge the relevance of this pricing scheme, we ask how often a particular group's payments reflect *the same* ICF. Figure 1 provides concrete illustrations. Each panel shows payment rates for the services provided regularly by a single physician group in the 2010 BCBS claims data.[11] Each circle on the graph is a unique payment amount for a unique service code. That is, if the group received two unique payment values for a standard office visit (HCPCS code 99213), say $45 and $51, those two amounts would show up as separate circles. The log Blue Cross payment amount is on the y-axis and the log of Medicare RVUs for the service are on the x-axis. The solid lines in Panels A and B have slopes of 1 and are drawn to coincide with each group's most common ICFs.

Panel A shows the data from a mid-sized group for which the relevance of a single ICF is readily apparent. Nearly all of this group's services share a single ICF, with a few deviations. The most natural interpretation of this graph is that those services on the solid line are priced according to Medicare RVUs with a common ICF, while 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 of the Medicare-linked share would view these and other circles off the solid line as deviations from Medicare-linked pricing.

Panel B presents an equivalently constructed graph for a larger group that provides more

[11] The figures exclude any code-by-payment combination that appears less than 10 times in the data associated with 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 each year's RVU updates (*e.g.* January 1—June 30, 2010). This way our calculations are not confounded by RVU changes.

unique services at more distinct prices. This group again has one particularly common ICF, though there is stronger evidence for the presence of a second, and possibly a third, contractually specified ICF. Finally, Panel C presents payment data for a large group that provides a substantial number of unique services. This large group has a range of ICFs, none of which visually dominate the payment picture. The scatterplot indicates the use of a remarkably complicated contract with BCBS.

To develop a summary measure of a group's links to Medicare, we make two approximations. First, we round the value of each $ICF_{c,j,t}$ to the nearest 20 cents, 10 cents, or 2 cents to explore sensitivity to allowances for rounding error. Second, we define "common ICFs" (cICFs) as those that rationalize a sufficiently large share of the BCBS's payments to a single physician group. In Figure 1, for example, the red lines in Panels A and B should undoubtedly qualify as cICFs. 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 services, then calculate the share of BCBS's payments associated with *any* of a group's cICFs.

4.2 Frequency of Common Implied Conversion Factors

Table 2 presents the share of services linked to Medicare in each year according to the methodology of section 4.1. The shares are substantial in each year and are moderately larger in 2010–11 than in 2008–09. The estimates range from 30 to 80 percent in 2008–09 and from 65 to 90 percent in 2010–11. The values increase marginally with the flexibility of our rounding threshold and decrease substantially with the stringency of the definition for a common ICF. Appendix Table B.2 shows that the results are qualitatively similar under a variety of alternative definitions.[12]

[12] 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

4.3 Heterogeneity in Share with Common ICFs

Our model of physician-insurer contracting emphasizes that we should expect to see deviations from Medicare's pricing schedule when the value of such deviations is high relative to negotiation and adjustment costs. To test this framework, this section considers heterogeneity along dimensions likely to proxy for the value of deviations.

The value of improving on Medicare's menu is driven primarily by two factors. First, the cost of maintaining inefficiencies embedded in Medicare's menu will be high when contracts cover large quantities of care. We thus anticipate relatively strong links when private insurers contract with small physician groups, and less benchmarking when considering contracts with large physician groups. Second, 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 more salient problems. Medicare rates are designed 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 better 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 heterogeneity according to services' capital and labor intensity by comparing the frequency of benchmarking across broad categories of care, such as labor-intensive evaluation and management services versus diagnostic imaging.

To adapt our ICF method for this heterogeneity analysis, we compute the share of services priced according to common Implied Conversion Factors (cICFs) at the physician group-by-service code ($j \times g$) level. We define fixed effects $\mathbb{1}_{b(j)}$ at the level of the 1-digit "Betos" classification of Berenson and Holahan (1990). To measure the relationship between group

the linked share. To overcome this problem, section 5 introduces a separate estimation strategy that is not sensitive to choices of this sort.

size and the Medicare-linked share, we categorize physician groups g according to vigintiles of their aggregate BCBS billing in a year, using $\mathbb{1}_{s(g)}$ to denote vigintile fixed effects.

Figure 2 shows the relationship between the share linked to Medicare and vigintiles of group size. Looking first at the measure shown with hollow red squares, we observe a stark negative relationship. Large groups' services have more deviations from Medicare benchmarking than small groups' services. With this measure, the variation across the vigintiles is around 20 percentage points. Appendix B presents regressions that summarize this fact. It further explores the relationship between payments and physician market structure. Among other things, we find that larger physician groups obtain higher ICFs.

To check whether the relationship between benchmarking and group size is affected by the composition of large and small groups' services, we run the following regression at the group-code level, separately by year:

$$\text{Medicare-Linked Share}_{j,g} = \eta_b \mathbb{1}_{b(j)} + \zeta_s \mathbb{1}_{s(g)} + \upsilon_{j,g}. \tag{4}$$

The orange diamonds in Figure 2 show the estimates of ζ_s, which can be interpreted as the relationship between Medicare links and group size, adjusted for service composition. The composition-adjusted relationship between group size and the Medicare-linked share remains strongly negative.

The remaining measures in Figure 2 show the Medicare-linked share in terms of dollars spent, rather than number of services. The results are quite similar to the services-based results. A stark negative relationship between firm size and the Medicare linked share is apparent in all four measures.

We next turn to differences across Betos categories, which are captured by estimates of η_b from equation (4). Column 1 of Table 3 reports estimates of equation (4), from the same regression using 2010 data that generated the group size coefficients shown in Figure 2B.

Column 2 drops the group size controls, and thus reports raw differences in means across Betos categories. The constant represents the mean benchmarking share for the omitted category, namely Evaluation & Management services. We find that benchmarking is 30–50 percent less frequent for Imaging, Procedures, and Tests than for Evaluation & Management services. Columns 3 and 4 report similar estimates when services are weighted according to the spending they represent. Table 3 and Figure 2 show that firm size and service categories independently predict variation in the prevalence of Medicare-benchmarked payments. These are precisely the types of variation that theory predicts.

5 Empirical Model Using Medicare Payment Changes

While our estimates of heterogeneity across groups and service types in the previous section are quite robust, our overall estimate of the Medicare-linked share is sensitive to the rounding and commonality thresholds chosen. Because theory provides no direct guidance as to what thresholds are most appropriate, we develop a second approach for estimating the pervasiveness of Medicare benchmarking. This empirical approach exploits updates to Medicare's allocation of RVUs.

5.1 Changes in Medicare's Relative Values

A committee of the American Medical Association, composed of representatives of various physician specialties, recommends RVU updates to CMS (Government Accountability Office, 2015). Medicare 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 incrementally 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 our sample, the newsletter announced updates taking place on July 1, 2008, on August 15, 2009, on July 1, 2010, and on September 1, 2011 (BCBS 2008; 2009; 2010; 2011). In all four years, the standard deviations of RVU changes are around 7 percent, so there is substantial pricing variation for us to exploit.

Figure 3A shows one example of how these changes can impact pricing 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 log payment rose by just under 0.05. Appendix Figure B.1 shows further examples. We next develop a method for using high frequency payment changes of this sort to infer the share of payments linked to Medicare.

5.2 Analytical Foundation

Our method exploits the institutional details we documented in section 1 about how Medicare benchmarking works in practice. When a payment $P_{g,j,t}$ is linked to Medicare's relative values, it takes the form of a scalar markup over Medicare RVUs, or

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

where g indexes physician groups, j indexes services, and t is a time period. Equation (5) implies that the scalar markup $\theta_{g,t}$ on Medicare-linked payments is additive in logs, so

$$\ln(P_{g,j,t}) = \ln(\theta_{g,t}) + \ln(RVU_{j,t}). \tag{6}$$

Equation (6) describes a linear relationship between log private insurance payments and log RVUs for a service, and in particular it predicts a regression coefficient of 1 on log RVUs. If the markup θ is a constant, it will be reflected in the constant term. If it varies across physician groups, then group fixed effects capture $\ln(\theta_g)$. If it changes over groups and across time, then group-by-time fixed effects serve the same role.

The institutional details, reflected in our model, suggested that payments may alternatively be negotiated without reference to RVUs. In this case, we denote the payment by $P_{g,j,t} = \rho_{g,j,t}$, or $\ln(P_{g,j,t}) = \ln(\rho_{g,j,t})$—with no necessary role for $\theta_{g,t}$ or $RVU_{j,t}$.

When RVUs change, these equations provide stark guidance about how private reimbursements will adjust. Consider two time periods, across which Medicare may update log RVUs 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. We can now write both types of prices in terms of service fixed effects and changes as follows. For Medicare-linked services, we have:

$$\ln(P_{g,j,t}) = \phi_j \mathbb{1}_j + \phi_g \mathbb{1}_g + \phi_{g,j} \mathbb{1}_g \cdot \mathbb{1}_j + \Delta\ln(RVU_{j,t}) \cdot \mathbb{1}_{\{t=\text{post}\}}. \tag{7}$$

For services not linked to Medicare, we have:

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

In these equations, $\mathbb{1}_{\{t=\text{post}\}}$ is an indicator for the second time period. In both types of price setting, 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 equations (7) and (8) implies a simple way to measure how many services are linked to Medicare. Equation (7) 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 appropriate fixed effects. Equation (8) 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.) The average of log reimbursements is then given by a weighted average of equations (7) and (8), and the coefficient on log RVU updates can reveal the linked share σ:

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

where we define $\epsilon_{g,j,t} = (1-\sigma) \cdot \varepsilon_{g,j,t} \cdot \mathbb{1}_{\{t=\text{post}\}}$. Equation (9) 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:

$$\ln(P_{c,g,j,t}) = \beta \Delta \ln(RVU_j) \cdot \mathbb{1}_{\{t=\text{post}\}} + \phi_t \mathbb{1}_{\{t=\text{post}\}} + \phi_j \mathbb{1}_j + \phi_g \mathbb{1}_g + \phi_{g,j} \mathbb{1}_g \cdot \mathbb{1}_j + \epsilon_{c,g,j,t}. \tag{10}$$

This is just a claims-level version of equation (9) that 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 the coefficient $\hat{\beta}$, our estimate of the share of services linked to Medicare, 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 will also estimate equation (10) weighted by the average pre-update price of each service.

For the estimate of $\hat{\beta}$ in specification (10) to equal the true Medicare-linked share σ, we must make several assumptions about active renegotiations of reimbursement rates. Since group and group-by-service fixed effects are intended to capture the level of markup θ, any changes in this markup over time may show up in the error term. In Appendix C.2, we discuss the situations in which this challenges our ability to identify the parameter σ. We emphasize there that the relatively high frequency at which we are able to estimate payment changes makes our assumptions quite plausible.

5.3 Parametric Event Study

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

$$\ln(P_{c,g,j,t}) = \sum_{t \neq 0} \beta_t \Delta \ln(RVU_j) \cdot \mathbb{1}_t + \phi_t \mathbb{1}_t + \phi_j \mathbb{1}_j + \phi_g \mathbb{1}_g + \phi_{g,j} \mathbb{1}_g \cdot \mathbb{1}_j + \epsilon_{c,g,j,t}. \tag{11}$$

When estimating equation (11), we normalize t such that $t = 1$ is the month in which BCBS has announced that it will implement RVU 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 RVU 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 RVU updates. We discuss this concern in detail in

Appendix C.

6 Results from RVU Update Analysis

6.1 Baseline Results

Figure 3B presents event study estimates of the link between Medicare's relative value scale and BCBS reimbursements. It shows estimates of equation (11) for the RVU changes implemented in 2010. BCBS's provider newsletters say that updates to Medicare's RVUs took effect that year on July 1, 2010.

The estimates reveal substantial links between RVU updates and the payments providers receive from BCBS. The coefficients imply that $\hat{\sigma} = 75$ percent of services are linked to Medicare's relative values. The dynamics in the figure are consistent with the view that this link involves the manner in which Medicare's relative values are embedded in BCBS's contracts. As in the raw data for standard office visits presented in Panel A, we see that payment changes occur at precisely the time when BCBS implements these updates.[13] 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 6.2 and 6.3 will investigate these deviations in detail. The extremely tight standard errors prior to the update in each year suggest that our fixed effects effectively capture the pre-update payments.

Figure 3C shows the variation across services that drives these results in the form of a binned scatterplot. This graph relates private reimbursement changes to Medicare fee schedule changes across different services. It shows that our results reflect pricing changes throughout the full distribution of Medicare changes.[14]

[13]The estimate for August 2009 is half of that in September and subsequent months, likely because of the mid-month RVU update date Blue Cross announced in that year (BCBS 2009).

[14]Appendix Figures C.1 and C.2 shows similar results for data from the other years.

Table 4 presents our baseline estimates of equation (10), which summarize our estimates for 2010 updates in a single coefficient. It further probes the robustness of these 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 and controls 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 RVU update. 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 three-quarters of BCBS's physician claims are linked to Medicare's relative value scale.

Panel B reports an equivalent set of specifications in which each service code is weighted according to the average BCBS payment prior to the updates. On average, the estimates imply that roughly 55 percent of BCBS's physician spending is linked to 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.[15]

The estimates presented in Figure 3 and Table 4 may differ from the true Medicare benchmarking parameter σ if changes in other terms of providers' contracts covary with the changes in RVUs. Indeed, payment changes that significantly alter physician groups' average Medicare payment can move private payments in subsequent years, due in part to the resulting changes to their bargaining positions (Clemens and Gottlieb, forthcoming). In Appendix C.4, we thus draw on institutional detail and theoretically motivated specification

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

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.

6.2 Which Servicse Are Benchmarked to Medicare?

We next investigate heterogeneity in our RVU-update estimates to explore the economic forces underlying the decision to benchmark to Medicare's payment menu. We consider heterogeneity along the same dimensions as in section 4.3, namely type of service and group size. The consistency of our results across methodologies, which differ in their strengths and weaknesses, strengthens the case for viewing the heterogeneity we uncover as reflecting systematic features of BCBS's physician contracts.

Table 5 estimates equation (10)—the relationship between private prices and changes in Medicare's relative values—separately across broad categories of services. Just as in the ICF-based results from Table 3, we observe a stronger relationship between private payments and Medicare updates for Evaluation & Management services than for Imaging. 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.[16]

Second, we divide Imaging codes into subcomponents with high capital and high labor content. Providers often bill separately for taking an image (the "Technical Component") and interpreting it (the "Professional Component"). The Professional Component is labor-intensive while the Technical Component, into which the billing codes amortize the imaging equipment's fixed cost, is capital-intensive. When the same group supplies both the Professional and Technical Components, it submits the bill as a "Global" service. The re-

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

sults in columns 5 through 7 show that payments for the Professional Component are more tightly linked to Medicare's relative values than are the payments for the 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 services.

6.3 Deviations from Benchmarking Across Physician Groups

We next consider how the strength of the link between private payments and Medicare's relative values vary across physician groups. In Figure 3D we allow our estimates of the strength of public-private payment benchmarking to vary with group size. The figure shows a binned scatterplot, analogous to Panel C, but with observations split into those coming from the largest and the smallest firms. We can see that the slope is steeper for the smaller physician groups, indicating that benchmarking is more common for their payments.

Table 6 quantifies this difference. The first column reports the baseline, service-weighted regression from Table 4. The second column introduces interactions between the RVU updates and indicators for services provided by firms of various sizes. We define mid-sized firms as those with \$200,000 to \$1,000,000 in annual billings with BCBS, and large firms as those with more than \$1,000,000 in annual billings. The estimates imply that nearly 90 percent of services provided by firms with less than \$200,000 in billings are benchmarked to Medicare, while roughly 60 percent of services provided by firms with more than \$1,000,000 in billings 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 with billings less than \$200,000 are benchmarked to Medicare, while one-third of payments to firms with more than \$1,000,000 in billings are benchmarked.[17] As with the estimates of heterogeneity across services, the heterogeneity by firm size is thus quite consistent between the ICF and

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

RVU-update methods.

6.4 Out-of-Network Payments

Our analysis thus far only includes in-network payments—those made to physician groups that have agreed with BCBS on mutually acceptable payment rates. In Appendix D we show comparable results for out-of-network payments, which arise when providers have not reached any such agreement. When a BCBS-insured patient sees an out-of-network provider, the ultimate payment reflects a complex interaction of the provider's charge, after-the-fact negotiations (as in Mahoney, 2015), and the insurance plan's coverage. So out-of-network payments are less likely to depend on a convenient benchmark such as the Medicare fee schedule.

Appendix Tables D.1 through D.3 show much weaker—if any—Medicare benchmarking in out-of-network payments. Estimates based on RVU changes, in Tables D.1 and D.2, find zero Medicare links in 2009 and 2011, and a small positive estimate in 2010. The estimates based on cICFs are higher, though still below the in-network estimates from section 4. The difference between these results and our in-network estimates suggests that the in-network prices reflect active efforts to negotiate around a simplified payment schedule.

7 How Do Private Payments Deviate from Medicare?

Thus far we have explored the frequency of deviations from strictly Medicare-linked contracts. In both the RVU-update and Implied Conversion Factor analyses, we presented evidence on how the frequency of such deviations varies across services and groups. In this section, we analyze the direction of BCBS's adjustments when it deviates from strictly Medicare-linked contracts. That is, we investigate what services BCBS rewards through upward adjustments and discourages through downward adjustments.

To measure these adjustments, we begin by estimating the following equation on claims from the pre-RVU-update period of each year—*i.e.* the initial months over which Medicare's relative values are constant:

$$\ln(P_{g,j}) = \psi \ln(RVU_j) + \delta_g + e_{g,j}. \tag{12}$$

If all payments were mechanically linked to Medicare's relative values, with a uniform contract for each group and no payment reporting error, $\psi \ln(RVU_j) + \delta_g$ would perfectly predict private payments. The group-specific δ_g estimates would account for heterogeneity in groups' markups over Medicare, and we would expect to find $\hat{\psi} = 1$. Conditional on a service's RVU allocation and group-specific markups, the prediction errors $e_{g,j}$ thus contain information about the direction of deviations from Medicare's relative values.

To understand which service categories tend to receive higher or lower payments than Medicare-benchmarking predicts we average $e_{g,j}$ by Betos category. Table 7 presents the resulting means, namely $\overline{\hat{e}_{g,j}} = \frac{1}{N_b} \sum_{j \in b} \hat{e}_{g,j}$ for each Betos group b, comprising N_b claims for all services $j \in b$ in that Betos group. The table shows that payments for Evaluation & Management and Testing services generally have positive residuals while payments for services in Imaging and Procedures have negative residuals.

Figure 4A plots the cumulative distributions of these residuals by Betos 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.[18] 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.

[18] 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.

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 substantially positive, at around 7 log points. Payment adjustments for the capital-intensive Technical Component of these services are substantially negative, averaging -12 log points. Figure 4B 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 payments structure, these results provide evidence that BCBS adjusts its contracts to increase the generosity of payments for labor-intensive services and decrease its payments for capital-intensive services.

8 Benchmarking and Payment System Efficiency

Thus far we have documented when, and how often, private contracts rely on Medicare's fee schedule and when they deviate. We now use these results to shed light on the negotiation costs and Medicare inefficiencies that can explain the benchmarking we see. For the RVU changes we study to have efficiency implications, they must generate meaningful changes in treatment. Section 8.1 provides evidence that service supply responds to price changes of the sorts we consider. In section 8.2 we then quantify the negotiation costs that would rationalize the level of benchmarking we observe.

8.1 Supply Responses

To determine whether the Medicare benchmarking we document alters the way physicians practice, we estimate how supply responds to relative price changes across services. The estimates use the same RVU changes as the foregoing analysis, which means that our estimates have three economically salient features. First, they involve short-run responses within a calendar year. Second, they involve responses to changes in the profitability of some services relative to others rather than to across-the-board changes in reimbursement rates.

Finally, they involve private payment changes that result from contractual links to changes in Medicare's relative rates.[19]

To measure supply responses, we estimate an analogue of the changes regression shown in Figure 3C in which the dependent variable is now the change in log quantity of care. We again split each calendar year into two time periods: the period before BCBS implemented the year's RVU updates, and period after it did so. The change in the log number of instances that a given physician group provided a particular service across these two time periods is our dependent variable.

Figure 5 shows the results of this estimation for 2010, along with binned scatterplots of the underlying data.[20] Note that this is a reduced-form estimate; it relates the Medicare price change to the supply responses for privately insured patients. In Appendix F, we use an IV setup to estimate the BCBS own-price supply elasticities. We estimate elasticities of 0.05, 0.15, 0.66, and 0.37 for the individual years, three of which are significant at the $p < 0.05$ level. These positive supply elasticities imply that the pricing decisions we examine have meaningful implications for how physicians provide treatment. If Medicare sets prices inefficiently, then copying Medicare's relative prices leads to inefficient care. When BCBS deviates from Medicare rates, these positive supply responses suggest that physicians respond to payment innovation as BCBS presumably intends.

8.2 The Costs of Complex Contracting

The benchmarking that we have documented implies that Medicare's pricing decisions spill over into private insurers' payments. At the same time, private insurers limit these

[19]This final characteristic distinguishes the private payment changes we analyze from changes driven by active contract renegotiations. Price changes due to active renegotiations may be better characterized as a product of the cross-price response of private care provision to Medicare payment changes. That is, a Medicare payment increase may lead physicians to shift supply from private to public patients absent increases in the payments negotiated with private insurers.

[20]Appendix Figure F.1 shows the analogues for other years.

spillovers in contracts with large provider groups and for services with particularly large deviations between average and marginal costs. We now consider what levels of negotiating costs, and inefficiencies in Medicare payment rates, can explain these decisions.

Recall from section 6 that 55 percent of spending is linked to Medicare rates, and that this share is larger for small physician groups. The potential efficiency gains from deviating away from the Medicare benchmark are increasing in the scale of the group's business with BCBS; in particular, we assume that these potential gains are a fraction x of the group's billings, b_i. Thus it makes sense to deviate from the Medicare fee schedule whenever the costs of doing so are less than xb_i. We now consider what combinations of x and negotiation costs rationalize the decision to deviate 55 percent of the time.

To do so, we rank physician groups according to the scale of their billings b_i to BCBS, from smallest to largest. We then consider a range of potential efficiency gains, from $x = 0.1\%$ to $x = 10\%$. We also consider a range of contract complexity costs c, from $c = \$500$ to $c = \$8{,}000$ per group. For each x and c, we ask what share of spending comes from groups whose potential efficiency gains are below these contracting costs, or in other words have $xb_i < c$. We aggregate the BCBS billings for all such groups and report their share of overall BCBS billings in Table 8. The entries in this table indicate the share of spending that we would expect to be linked to Medicare, under various assumptions for x and c.

To understand the calculations in Table 8, we have to be careful about what the potential efficiency gains x mean. These gains are the potential values of improvements in the care provided by a physician group in response to a potential deviation from the Medicare fee schedule—or potential savings while providing the same value. Just as in the model from section 2, and our subsequent empirical analysis, these deviations involve changes in relative prices for the services this group provides. These are not efficiency gains from a switch from fee-for-service medicine to a value-based payment system. They are also not the gains from reimbursement changes that shift the aggregate composition of specialists or affect technology diffusion. Just as in the supply responses from section 8.1, these should be interpreted as changes in treatments from relative price changes over a modest time horizon.

Second, our calculations only make sense if x describes efficiency gains captured by the same parties that incur the negotiation costs c, namely BCBS or the relevant physician group. For a given payment system improvement's overall welfare gain, the health insurance market's opacity may limit the insurer and physician's ability to capture its incidence.[21] We term the gains split between BCBS and the physician group the "captured efficiency gain."

Table 8 highlights in bold the entries that correspond most closely to the overall Medicare-linked share we estimate, namely 55 percent. That share can be rationalized by a complexity cost on the order of \$4,000 per physician group combined with captured efficiency gains of 1 percent, or with a cost of \$2,000 and a captured efficiency gain of 0.5 percent.

The table also allows us to infer how much improvements in Medicare's fee schedule can change administrative costs in medicine. Suppose that a Medicare payment reform reduced inefficiency in half, thereby also halving the potential gains from fine tuning payments (x). Table 8 implies that the fraction of spending covered by complex contracts falls by around 14 percentage points.[22] Specifically, they would simplify the contracts with 1,057 firms whose average billings of nearly \$300,000 place them in the middle of the spending distribution. This simplification could save these physician groups and BCBS around \$2.1 million in administrative costs. An extrapolation to the national market suggests that these hypothetical improvements in Medicare's payments would, in addition to increasing the efficiency of care provision, save over \$200 million in contracting and payment processing costs.[23]

[21]The fact that physicians' contracts are considered proprietary, for example, makes it difficult for consumers to recognize differences in the efficiency of different insurers' payments systems.

[22]Looking at column 3, we are supposing a move from $x = 1\%$ to $x = 0.5\%$. This means that firms between the 40th and the 54th percentiles of the spending distribution (when ranked by firm size) stop deviating from the Medicare benchmark. There are 1,057 such firms.

[23]This reflects that fact that BCBS accounts for roughly 1/100th of the private market for outpatient services.

9 Conclusion

This paper uses physician payments from a large private insurer, Blue Cross Blue Shield of Texas (BCBS), as a window into how private firms contract for services in complex environments. Using two empirical strategies, we show that BCBS benchmarks to Medicare's schedule of relative prices to significantly simplify this problem. We estimate that roughly three quarters of services and 55 percent of BCBS's payments are directly linked to Medicare.

We find evidence that the one quarter of services and nearly half of payments that deviate from Medicare's relative rate structure involve an effort to improve the payment structure. BCBS tends to deviate when the value of doing so appears to be highest. Deviations occur disproportionately in contracts with large physician groups, where significant mutual gains can be on the line. BCBS significantly reduces its 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). BCBS hews closely to Medicare in payments for services where average-cost reimbursements will be most aligned with marginal costs, such as labor-intensive primary care services. When it deviates, the direction of BCBS's payment adjustments would tend to to encourage the provision of primary care and discourage care for which over-utilization is a more widespread concern.

The use of Medicare as a pricing backstop 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, we find that BCBS adjusts its payments 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 innovation and reform.

References

Abeler, Johannes and Simon Jäger, "Complex Tax Incentives," *American Economic Journal: Economic Policy*, August 2015, *7* (3), 1–28.

Alpert, Abby, Mark Duggan, and Judith K. Hellerstein, "Perverse reverse price competition: Average wholesale prices and Medicaid pharmaceutical spending," *Journal of Public Economics*, 2013, *103*, 44–62.

Anderson, Eric, Nir Jaimovich, and Duncan Simester, "Price Stickiness: Empirical Evidence of the Menu Cost Channel," *Review of Economics and Statistics*, forthcoming.

Baker, Laurence C., M. Kate Bundorf, Anne B. Royalty, and Zachary Levin, "Physician Practice Competition and Prices Paid by Private Insurers for Office Visits," *Journal of the American Medical Association*, 2014, *312* (16), 1653–1662.

Berenson, Robert and John Holahan, "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 December 1990.

Bils, Mark and Peter J. Klenow, "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, October 2004, *112* (5), 947–985.

Blue Cross Blue Shield of Texas, "Fee schedule update," *Blue Review*, 2008, (1), 8.

_ , "Fee schedule updates effective in August," *Blue Review*, 2009, (2), 5.

_ , "Fee schedule updates," *Blue Review*, 2010, (3), 15.

_ , "Fee schedule update," *Blue Review*, 2011, (6), 10.

Borges, Walt, "Mo' Better Blues," *Texas Medicine*, April 2003, *99* (4).

Brekke, Kurt R., Tor Helge Holmås, Karin Monstad, and Odd Rune Straume, "Do Treatment Decisions Depend on Physicians' Financial Incentives?," Working Ppaer 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).

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Robust Inference With Multiway Clustering," *Journal of Business & Economic Statistics*, 2011, *29* (2), 238–249.

Centers for Medicare and Medicaid Services, "National Health Expenditure Accounts," 2011. Available at http://www.cms.hhs.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/NationalHealthExpendData/Downloads/tables.pdf.

Chandra, Amitabh and Jonathan S. Skinner, "Technology Growth and Expenditure Growth in Health Care," Technical Report 3 September 2012.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, September 2009, *99* (4), 1145–1177.

Clarida, Richard, Jordi Galí, and Mark Gertler, "The science of monetary policy: Evidence and some theoryj," *Journal of Economic Literature*, December 1999, *37* (4), 1661–1707.

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.

_ and _ , "Do Physicians' Financial Incentives Affect Medical Treatment and Patient Health?," *American Economic Review*, April 2014, *104* (4), 1320–1349.

_ and _ , "In the Shadow of a Giant: Medicare's Influence on Private Payment Systems," *Journal of Political Economy*, forthcoming.

_ , _ , and Adam Hale Shapiro, "How Much Do Medicare Payment Cuts Reduce Inflation?," *FRBSF Economic Letter*, September 22 2014, *2014-28.*

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.

DellaVigna, Stefano, "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, June 2009, *47* (2), 315–372.

Duggan, Mark and Fiona Scott Morton, "The Distortionary Effects of Government Procurement: Evidence for Medicaid Prescription Drug Purchasing," *Quarterly Journal of Economics*, February 2006, *121*, 1–30.

Dunn, Abe and Adam Hale Shapiro, "Do Physicians Possess Market Power?," Technical Report 1 February 2014.

_ and _ , "Physician Payments Under Health Care Reform," *Journal of Health Economics*, 2015, *39*, 89–105.

Ellis, Randall P. and Thomas G. McGuire, "Provider Behavior Under Prospective Reimbursement: Cost Sharing and Supply," *Journal of Health Economics*, June 1986, *5* (2), 129–152.

Gabaix, Xavier, "A Sparsity-Based Model of Bounded Rationality," *Quarterly Journal of Economics*, November 2014, *129* (4), 1661–1710.

Gesme, Dean H. and Marian Wiseman, "How to Negotiate With Health Care Plans," *Journal of Oncology Practice*, July 2010, *6* (4), 220–222.

Government Accountability Office, "Medicare Physician Payment Rates: Better Data and Greater Transparency Could Improve Accuracy," Report GAO-15-434, GAO May 2015.

Grennan, Matthew and Ashley Swanson, "Transparency and Negotiated Prices: The Value of Information in Hospital-Supplier Bargaining," February 2015. University of Pennsylvania, 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.

Kleiner, Samuel A., William D. White, and Sean Lyons, "Market power and provider consolidation in physician markets," *International Journal of Health Economics and Management*, March 2015, *15* (1), 99–126.

Laugesen, Miriam J. and Sherry A. Glied, "Higher Fees Paid To US Physicians Drive Higher Spending For Physician Services Compared To Other Countries," *Health Affairs*, 2011, *30* (9), 1647–1656.

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

Mankiw, N. Gregory and Ricardo Reis, "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 2002, *117* (4), 1295–1328.

MedPAC, "Report to the Congress: Medicare and the Health Care Delivery System," Technical Report, Medicare Payment Advisory Commission June 2011. Available at http://www.medpac.gov/documents/reports/Jun11_EntireReport.pdf (accessed October 7, 2015).

Mertz, Gregory J., "Can You Negotiate Better Reimbursement?," *Family Practice Management*, October 2004, *11* (9), 31–34.

Mullainathan, Sendhil, Joshua Schwartzstein, and Andrei Shleifer, "Coarse Thinking and Persuasion," *Quarterly Journal of Economics*, May 2008, *123* (2), 577–619.

Nakamura, Emi and Jón Steinsson, "Five Facts About Prices: A Reevaluation of Menu Cost Models," *The Quarterly Journal of Economics*, November 2008, *123* (4), 1415–1464.

Nandedkar, Maithily, "Key principles for negotiating with insurers, deciding whether to opt in or out of Medicare," September 2011. Available at http://www.aad.org/dw/monthly/2011/september/key-principles-for-negotiating-with-insurers-deciding-whether-to-opt-in-or-out-of-medicare (accessed August 21, 2013).

OECD, "OECD.Stat Web Browser," 2015. Available at http://stats.oecd.org/ (accessed March 21, 2015).

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

Sims, Christopher A., *Journal of Monetary Economics*, April 2003, *50* (3), 665–690.

White, Chapin, "Contrary To Cost-Shift Theory, Lower Medicare Hospital Payment Rates For Inpatient Care Lead to Lower Private Payment Rates," *Health Affairs*, May 2013, *32* (5), 935–943.

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

Figure 1: Raw Payments For Illustrative Physician Groups, 2009

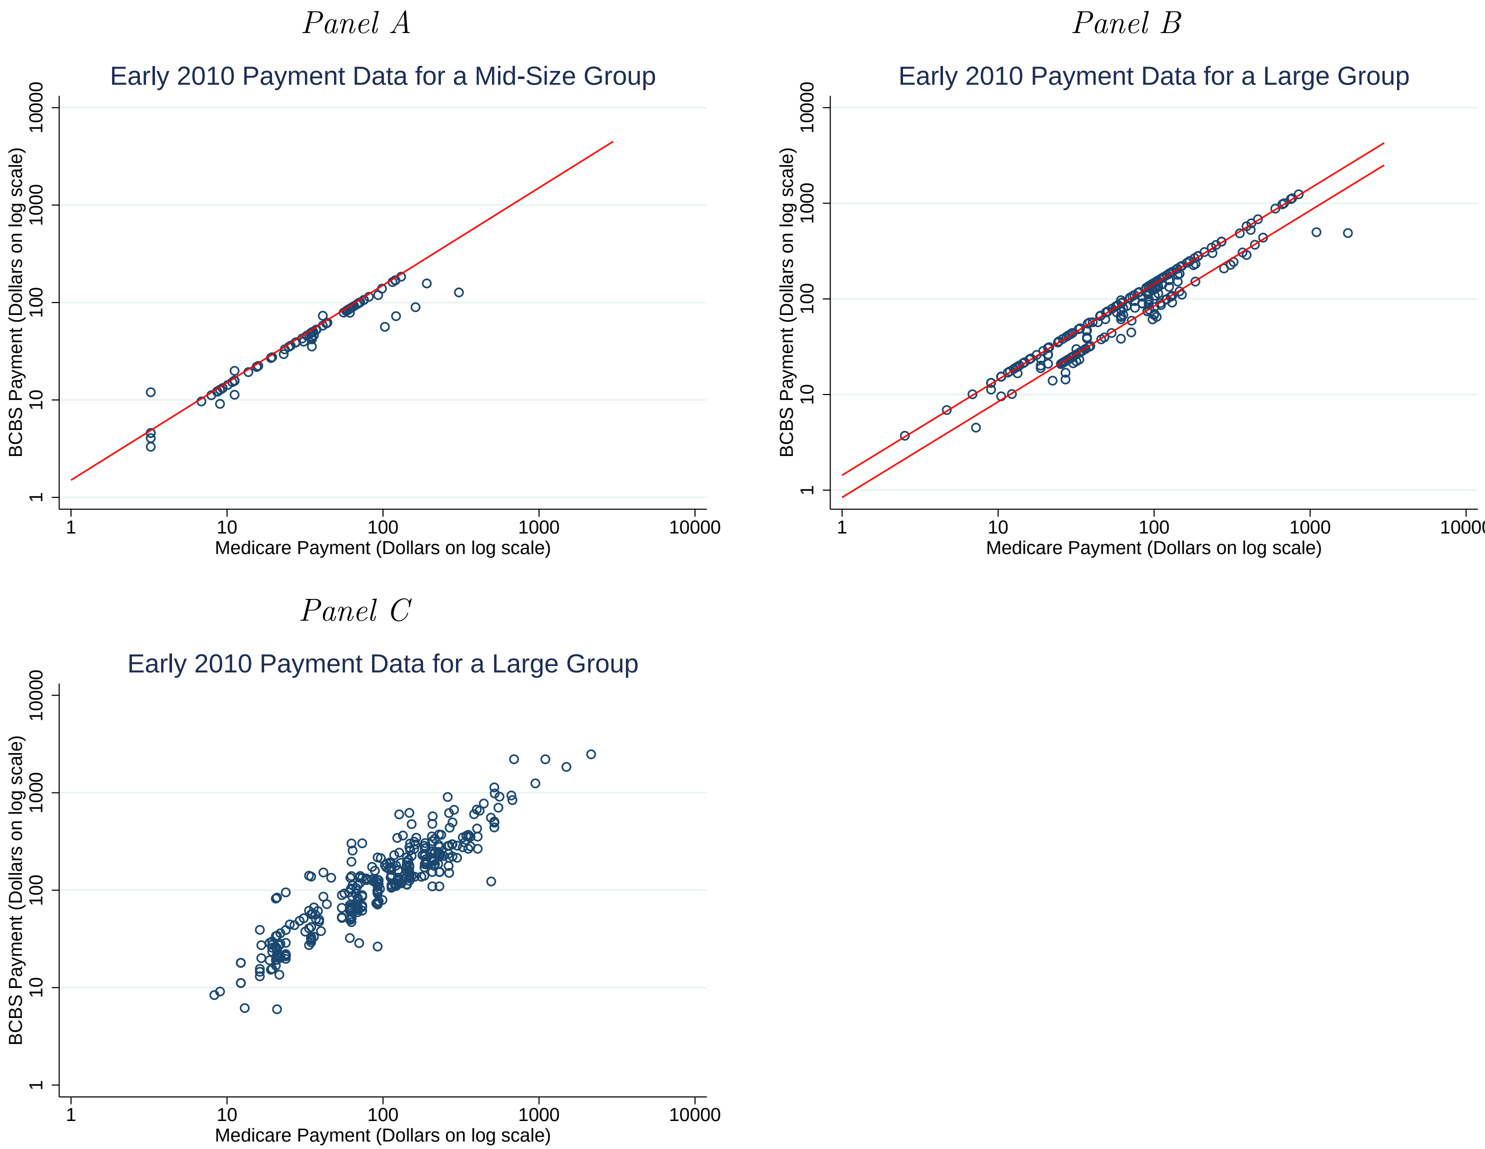

Note: The figure presents the raw data on BCBS reimbursement rates, and associated Medicare reimbursement, for 3 different physician groups in 2010. Each observation is a unique reimbursement paid for a particular service to the group. The lines have a slope of 1 (in logs) and represent the groups' most common Implied Conversion Factors. Sources: Authors' calculations using claims data from BCBS.

Figure 2: Frequency of Benchmarking and Physician Group Size

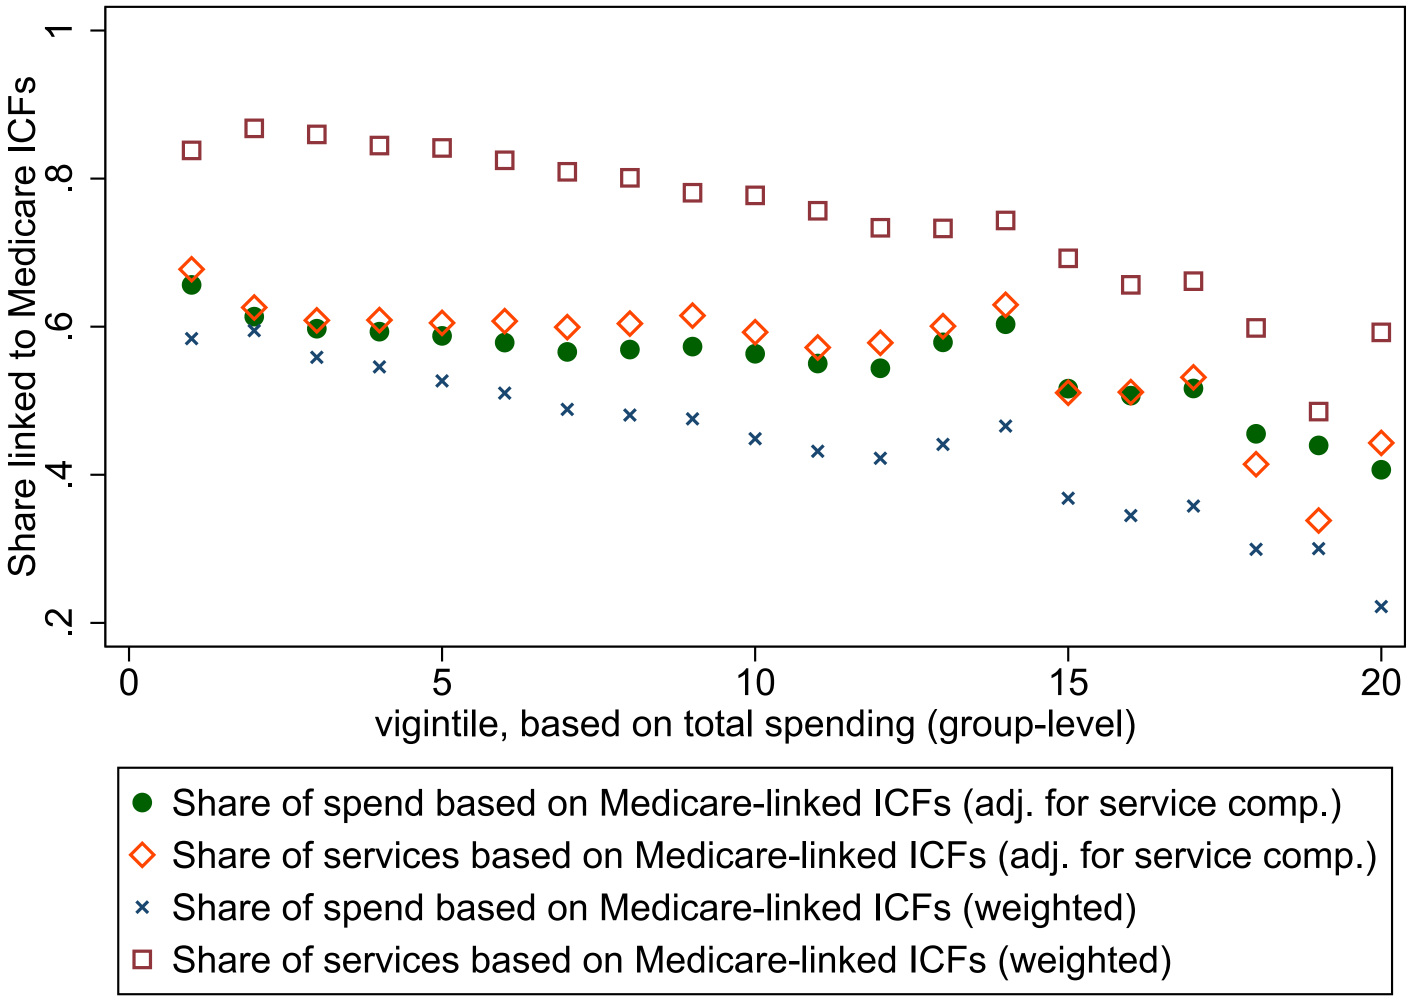

This graph shows the share of services priced according to common Implied Conversion Factors (cICFs), as defined in section 4.1, against the amount of BCBS spending on care provided by the physician group (grouped into 20 vigintiles). We interpret this as measuring the relationship between a group's Medicare-linked service share and group size. The green dots and orange diamonds show estimates of ζ_b from equation (4), which adjust for the composition of each group's services. The blue $\times$'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. Sources: Authors' calculations using claims data from BCBS.

Figure 3: Benchmarking Estimates Based on Price Changes

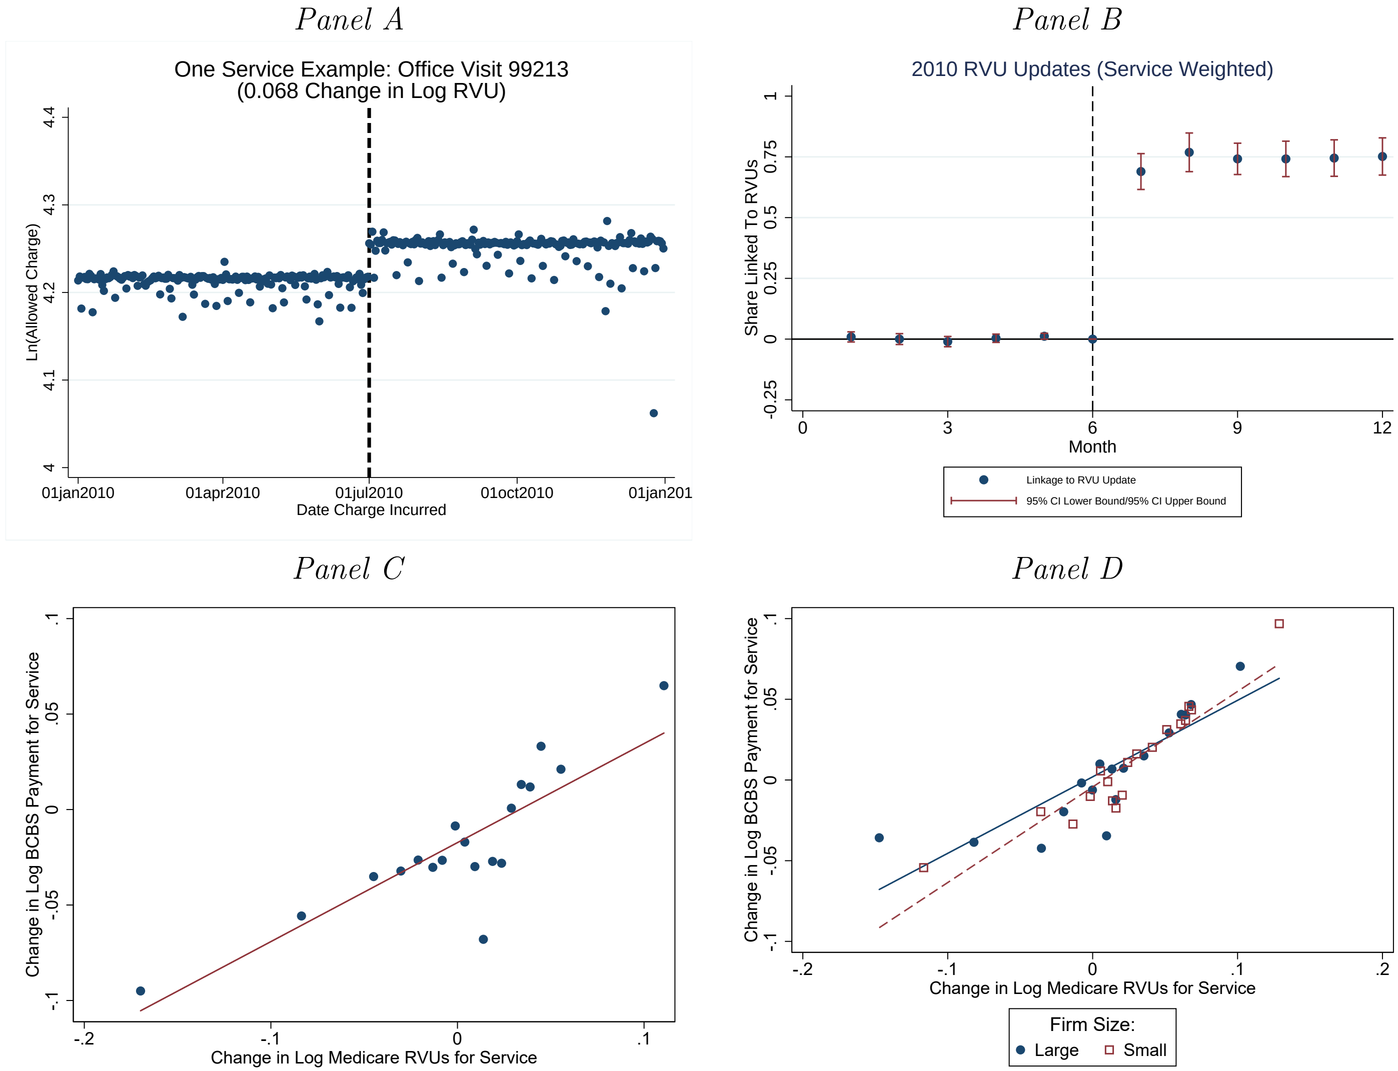

Note: All data are from calendar year 2010. BCBS implemented its update from the 2009 to 2010 relative value scales on July 1, 2010, as indicated by the vertical dashed line in Panels A and B. Panel A presents daily averages of BCBS's log payment for a standard office visit. Panel B reports estimates of the β_p from estimates of equation (11). In Panels C and D, price changes are computed between observations before and after July 1, 2010. The regressions are run at the underlying service level, but observations are grouped into twenty bins for each year, based on vigintiles of the Medicare log RVU change. Panel C reports the relationships described by equation (C.1) for RVU updates in each year, and estimates of that equation. Panel D is similar, but splits the data into services provided by the largest physician groups (those with at least $1,000,000 in annual billings) and the smallest groups (under $200,000). Sources: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.

Figure 4: **Deviations from Medicare Benchmark by Service Category**

Panel A

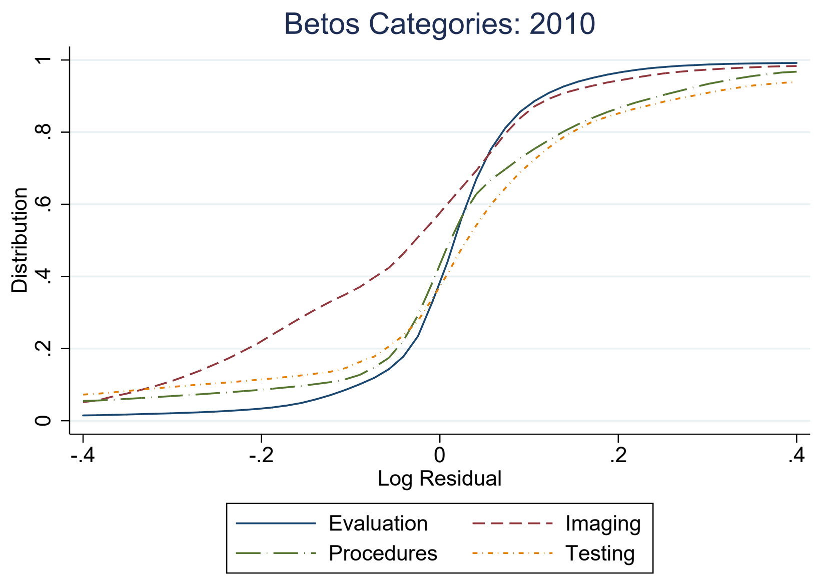

Panel B

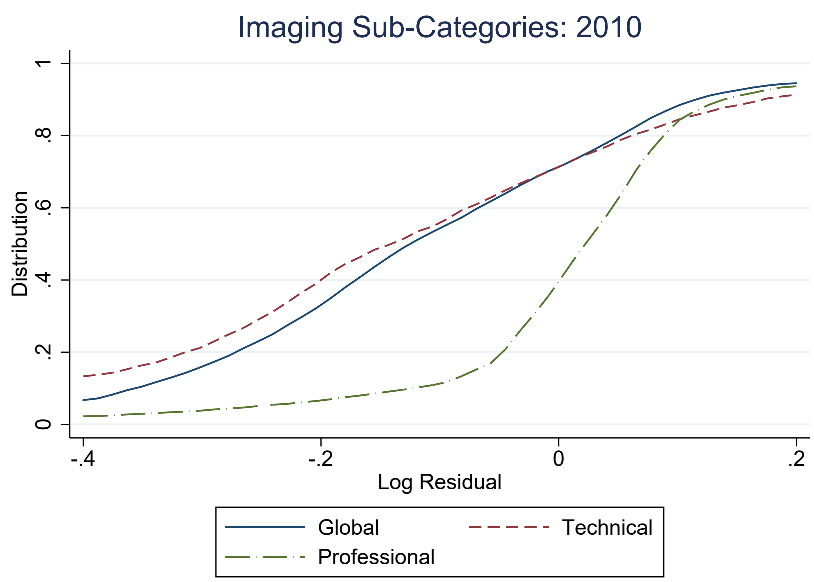

Note: The figure presents residuals $\epsilon_{g,j}$ from estimates of equation (12). The distribution of residuals is shown within either broad Betos categories (Panel A), or within the subcategories of Imaging (Panel B). The distributions are smoothed using a local linear regression, with an Epanechnikov kernel and a bandwidth of 0.01. Sources: Authors' calculations using claims data from BCBS.

Figure 5: Short-Run Supply Responses to Medicare Price Changes

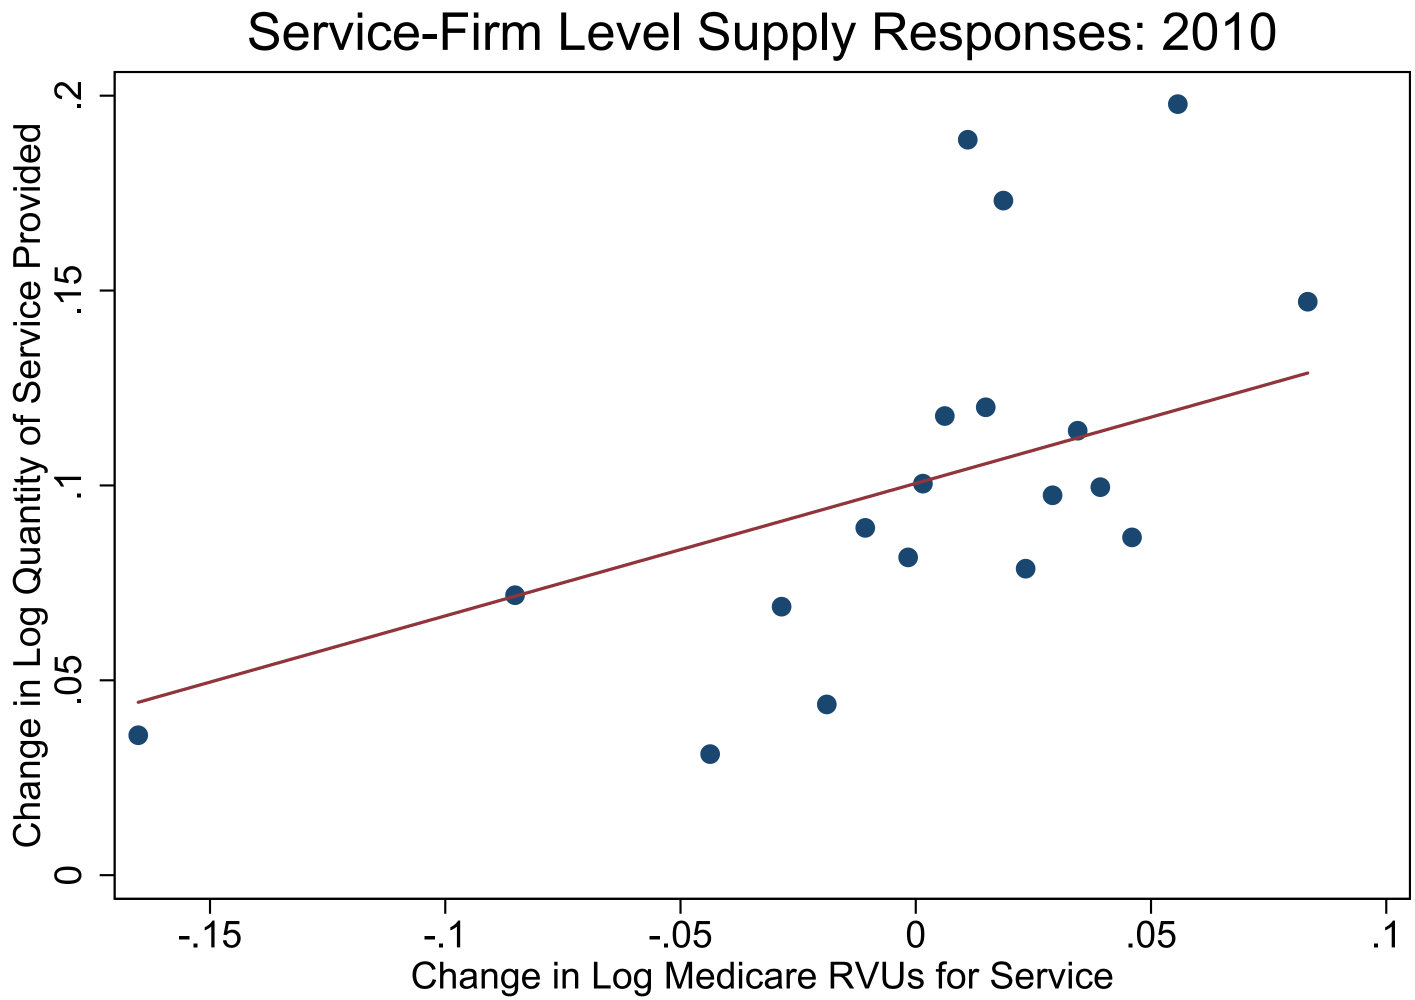

Note: The figure reports estimates of physicians' supply responses to Medicare price changes that BCBS implemented in 2010. Quantities, the dependent variable, are computed at the service-by-firm level. The figure shows estimates of changes in these quantities, measured as log differences between the period before BCBS implemented the Medicare RVU updates (on July 1, 2010) and the period after this update. Sources: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.

Table 1: Summary Statistics by Physician Group

Panel A: All Groups (N=81,741)					
	Mean	**Median**	**Std. Dev.**	**Min.**	**Max.**
Number of unique services	12.01	3	36.19	1	~2,100
Number of patients	102.2	2	2,247.25	1	~394,000
Number of doctors	1.71	1	7.84	1	~1,100
Number of claims	231	3	4,390	1	~740,000
Mean allowed amount	102.78	71.70	211.87	~5	~34,800
Total BCBS revenues	24,044	350	336,259	~5	~55,000,000
Panel B: Groups with Billings $>$ \$10,000 (N=15,235)					
	Mean	**Median**	**Std. Dev.**	**Min.**	**Max.**
Number of unique services	47.08	30	73.25	1	~2,100
Number of patients	527.26	157	5,184.05	1	~394,000
Number of doctors	4.19	2	17.90	1	~1,000
Number of claims	1,195	391	10,112	1	~740,000
Mean allowed amount	103.67	69	371.67	~6	~34,800
Total BCBS revenues	125,096	39,400	770,798	10,000	~55,000,000

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

Table 2: **Services Priced According to Common Implied Conversion Factors**

Panel A: 2008			
	Frequency Threshold:		
	5%	*10%*	*20%*
Rounding for ICFs:			
\$0.02	67%	53%	34%
\$0.10	72%	59%	40%
\$0.20	77%	65%	48%
Panel A: 2009			
	Frequency Threshold:		
	5%	*10%*	*20%*
Rounding for ICFs:			
\$0.02	67%	52%	33%
\$0.10	73%	59%	39%
\$0.20	77%	65%	47%
Panel B: 2010			
	Frequency Threshold:		
	5%	*10%*	*20%*
Rounding for ICFs:			
\$0.02	87%	81%	70%
\$0.10	89%	84%	75%
\$0.20	89%	85%	75%
Panel C: 2011			
	Frequency Threshold:		
	5%	*10%*	*20%*
Rounding for ICFs:			
\$0.02	86%	78%	66%
\$0.10	88%	82%	72%
\$0.20	88%	82%	72%

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. We restrict to data from the period before BCBS implemented each year's RVU updates (*e.g.* January 1—June 30, 2010). This way our calculations are not confounded by RVU changes that occur later in the calendar year. The cells within each panel show how this 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.

Table 3: Medicare Benchmarking by Betos Category

	(1)	(2)	(3)	(4)
Dependent variable:	Payments with Common Conversion Factors			
	Service Share		Spending Share	
Imaging	-0.380**	-0.419**	-0.458**	-0.488**
	(0.033)	(0.026)	(0.050)	(0.044)
Procedures	-0.382**	-0.416**	-0.324**	-0.351**
	(0.060)	(0.055)	(0.033)	(0.029)
Tests	-0.297**	-0.323**	-0.389**	-0.410**
	(0.064)	(0.062)	(0.053)	(0.051)
Constant	0.838**	0.788**	0.830**	0.783**
	(0.023)	(0.019)	(0.016)	(0.017)
N	542,207	542,207	542,207	542,207
Omitted Category	Evaluation & Management			
Additional Controls	Group Size	None	Group Size	None

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. This table shows estimates of the η_b coefficients in equation (4), namely the relationship between Betos category and the Medicare-linked share of services (columns 1 and 2) or spending (columns 3 and 4) at the group-service code level. 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. 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, Gelbach and Miller, 2011) by Betos category and physician group. Sources: Authors' calculations using claims data from BCBS.

Table 4: Estimating Medicare Benchmarking Using RVU Changes

	(1)	(2)	(3)	(4)
Dependent variable:	*Log private reimbursement rate*			
	Panel A: 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	3,681	3,681	3,681	3,681
	Panel B: 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	3,681	3,681	3,681	3,681
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: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. The table shows the results of OLS specifications of the forms described in section 5.2. Each column in each panel reports an estimate of $\hat{\beta}$ from equation (10). Observations are at the claim-line level and are equally weighted (Panel A), or weighted according to each service's average payment during the baseline period (Panel B). 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. Sources: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.

Table 5: **Public-Private Payment Links Across Service Categories**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent variable:				*Log private reimbursement rate*			
	Evaluation	Imaging	Procedures	Tests	Imaging Sub-Categories:		
					Global	Technical	Professional
Log RVU Change	0.841**	0.564**	0.720**	1.066**	0.545**	0.387*	0.982**
× Post-Update	(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	1,085	1,936	408	408	244	433

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. The table shows the results of OLS specifications of the forms described in section 5.2. The cells in each panel report estimates of $\hat{\beta}$ from equation (10), with samples selected to contain the HCPCS codes falling into individual 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). Additional features of each specification are described within the table. The construction of all variables is further described in the main text. Sources: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.

Table 6: **Medicare Benchmarking by Firm Size**

	(1)	(2)	(3)	(4)
Dependent variable:		*Log private reimbursement rate*		
Log RVU Change	0.750**	0.882**	0.539**	0.775**
× Post-Update	(0.038)	(0.073)	(0.061)	(0.094)
Log RVU Change		-0.074		-0.140*
× Post-Update × Midsize		(0.098)		(0.069)
Log RVU Change		-0.293*		-0.448**
× Post-Update × Large		(0.117)		(0.102)
N	23,933,577	23,933,577	23,933,577	23,933,577
Weighting:	Service	Service	Dollar	Dollar

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. 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 indicators 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). Sources: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.

Table 7: **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	0.0211	-0.0398	-0.0237	0.0759	-0.124	-0.125	0.0698
Residaul SD	(0.200)	(0.274)	(0.251)	(0.349)	(0.282)	(0.295)	(0.216)
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 equation (12) 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. 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.

Table 8: **How Much Benchmarking is Expected Depending on Assumed Inefficiency and Contracting Costs?**

	(1)	(2)	(3)	(4)	(5)
	Cost of adding complexity (c)				
Potential captured efficiency gains as share of group billings (x)	\$1,000	\$2,000	\$3,000	\$5,000	\$10,000
0.1%	**58%**	70%	78%	84%	89%
0.25	40	**54**	66	76	83
0.5	26	40	**54**	66	76
1	14	26	40	**54**	66
2	7	14	26	40	**54**
5	3	6	11	22	35
10	2	3	6	11	22

Note: Each cell reports the estimated percent of BCBS spending for which the value of shifting from a simple (Medicare-linked) contract to a complex contract would be exceeded by the assumed cost of adding complexity. The different assumptions for additional costs (c) are listed at the top of each column. For any one physician group, the value of deviating from the simple Medicare-linked contract is the magnitude of that firm's BCBS billings (b_i) times the efficiency gains (x) from switching to a more complex contract. These efficiency gains x are the overall efficiency gain times the fraction of that value that the insurer and physician can capture. For each cell defined by x and c, we aggregate the BCBS billings for all firms small enough that $xb_i < c$; that is, for all firms where it would be inefficient for the parties to deviate from the Medicare benchmark. The numbers reported in each cell are the share of total BCBS spending accounted for by these firms, or $\frac{\sum_{i;b_i<c/x} b_i}{\sum_i b_i}$. The bolded entries along the diagonal indicate the combinations of assumed complexity costs c and efficiency gains x for which the implied fraction of spending covered by Medicare-linked contracts comes closest to matching the fraction estimated in the claims data.

Appendix For Online Publication Only

A Proofs

Proof of Result 1. When relative prices are fixed, the insurer can only adjust the overall markup over Medicare, φ. Hence reimbursements are $r_1 = \varphi r_1^M$ and $r_2 = \varphi r_2^M$. Patient utility is

$$u(q_1, q_2) = u\left(s_1(\varphi r_1^M), s_2(\varphi r_2^M)\right) = \varphi\left(\alpha a r_1^M + \beta b r_2^M\right). \tag{A.1}$$

The insurer must achieve utility level $\overline{u}$ for the patients, and $\varphi = \dfrac{\overline{u}}{\alpha a r_1^M + \beta b r_2^M}$ is the minimum markup that can do so.

Expenditures are simply

$$\hat{E} = s_1\left(\varphi r_1^M\right)\varphi r_1^M + s_2\left(\varphi r_2^M\right)\varphi r_2^M = \alpha\varphi^2(r_1^M)^2 + \beta\varphi^2(r_2^M)^2. \tag{A.2}$$

□

Proof of Result 2. The insurer's problem is to choose reimbursement rates r_1 and r_2 to solve:

$$\min s_1(r_1)r_1 + s_2(r_2) \quad \text{subject to} \quad u(s_1(r_1), s_2(r_2)) \geq \overline{u}. \tag{A.3}$$

Given the functional form assumptions, we can write the minimization problem as:

$$\mathscr{L}(r_1, r_2) = \alpha r_1^2 + \beta r_2^2 - \lambda(\alpha a r_1 + \beta b r_2 - \overline{u}) \tag{A.4}$$

where λ is the multiplier on the patient utility constraint. The first-order conditions are:

$$r_1^* = \frac{\lambda a}{2} \tag{A.5}$$

$$r_2^* = \frac{\lambda b}{2} \tag{A.6}$$

$$\overline{u} = \alpha a r_1^* + \beta b r_2^* \tag{A.7}$$

Thus $\dfrac{r_2^*}{r_1^*} = \dfrac{b}{a}$. We can then solve for $r_1^* = \dfrac{a\overline{u}}{\alpha a^2 + \beta b^2}$. Hence medical expenditures are

$$E^* = \frac{\overline{u}^2}{\alpha a^2 + \beta b^2}. \tag{A.8}$$

To compare these expenses with those from Result 1, first define $\omega = \frac{r_2^M}{r_1^M}$ as the ratio of Medicare payments for the two services. We can then write the insurer's markup over

Medicare in the benchmarking case as

$$\varphi = \frac{\overline{u}}{(\alpha a + \beta b\omega) r_1^M} \tag{A.9}$$

and the expenditures in that case as

$$\begin{aligned} \hat{E} &= \left(\alpha + \beta\omega^2\right) \varphi^2 (r_1^M)^2 \\ &= \frac{\overline{u}^2 \left(\alpha + \beta\omega^2\right)}{(\alpha a + \beta b\omega)^2} \end{aligned} \tag{A.10}$$

It is convenient to work with the ratio of constrained to unconstrained expenditures:

$$\psi = \frac{\hat{E}}{E^*} = \frac{\left(\alpha + \beta\omega^2\right)\left(\alpha a^2 + \beta b^2\right)}{(\alpha a + \beta b\omega)^2}. \tag{A.11}$$

Note first that if $\omega = \frac{b}{a}$, then this simplifies to $\psi = 1$, as asserted in the Result. To determine what happens as ω varies, we compute the derivative:

$$\begin{aligned} \frac{d\psi}{d\omega} &= \frac{2\beta\omega\left(\alpha a + \beta b\omega\right)^2\left(\alpha a^2 + \beta b^2\right) - 2\beta b\left(\alpha + \beta\omega^2\right)\left(\alpha a^2 + \beta b^2\right)\left(\alpha a + \beta b\omega\right)}{(\alpha a + \beta b\omega)^4} \\ &= (\omega a - b)\frac{2\alpha\beta\left(\alpha a^2 + \beta b^2\right)}{(\alpha a + \beta b\omega)^3}. \end{aligned} \tag{A.12}$$

All of the terms in the fraction at the end of equation (A.12) are positive. The term in front, $\omega a - b$, is positive whenever $\omega > \frac{b}{a}$ and negative whenever $\omega < \frac{b}{a}$. Thus the ratio of expenses is increasing in ω when ω is above the privately efficient reimbursement ratio, and decreasing in ω whenever ω is below the efficient ratio. This proves that any ratio $\omega \neq \frac{b}{a}$ leads to higher medical expenditures than $\omega = \frac{b}{a}$, as the Result asserts. □

Proof of Result 3. The insurer's expenses when benchmarking to Medicare are given by equation (A.10), and those when unconstrained are given by equation (A.8). The difference between these values is

$$\begin{aligned} \xi &= \frac{\overline{u}^2\left(\alpha + \beta\omega^2\right)}{(\alpha a + \beta b\omega)^2} - \frac{\overline{u}^2}{\alpha a^2 + \beta b^2} \\ &= \overline{u}^2 \frac{\left(\alpha + \beta\omega^2\right)\left(\alpha a^2 + \beta b^2\right) - (\alpha a + \beta b\omega)^2}{\left(\alpha a^2 + \beta b^2\right)(\alpha a + \beta b\omega)^2} \\ &= \overline{u}^2 \alpha\beta \frac{a^2\omega^2 + b^2 - 2ab\omega}{\left(\alpha a^2 + \beta b^2\right)(\alpha a + \beta b\omega)^2}. \end{aligned} \tag{A.13}$$

Note that equation (A.13) is equal to zero when $\omega = \frac{b}{a}$. Otherwise it is positive, since it has a minimum at $\omega = \frac{b}{a}$.

The remainder of the Result simply requires taking derivatives of ξ:

$$\frac{d\xi}{d\overline{u}} = 2\overline{u}\alpha\beta\frac{a^2\omega^2 + b^2 - 2ab\omega}{(\alpha a^2 + \beta b^2)\,(\alpha a + \beta b\omega)^2} > 0 \tag{A.14}$$

$$\begin{aligned}
\frac{d\xi}{d\omega} &= \overline{u}^2\alpha\beta\frac{(2a^2\omega - 2ab)(\alpha a + \beta b\omega)^2 - 2\beta b(a^2\omega^2 + b^2 - 2ab\omega)(\alpha a + \beta b\omega)}{(\alpha a^2 + \beta b^2)\,(\alpha a + \beta b\omega)^4} \\
&= 2\overline{u}^2\alpha\beta\frac{a(a\omega - b)(\alpha a + \beta b\omega) - \beta b(a^2\omega^2 + b^2 - 2ab\omega)}{(\alpha a^2 + \beta b^2)\,(\alpha a + \beta b\omega)^3} \\
&= 2\overline{u}^2\alpha\beta\frac{(\alpha a^2 + \beta b^2)(\omega a - b)}{(\alpha a^2 + \beta b^2)\,(\alpha a + \beta b\omega)^3} \\
&= (\omega a - b)\frac{2\overline{u}^2\alpha\beta}{(\alpha a + \beta b\omega)^3}.
\end{aligned} \tag{A.15}$$

Inequality (A.14) shows that ξ is increasing in $\overline{u}$, which measures the generosity of insurance, or the quantity of services provided (since utility is assumed to be increasing in quantity). Equation (A.15) shows that ξ is increasing in ω when $\omega > \frac{b}{a}$, and decreasing in ω when $\omega < \frac{b}{a}$. Thus ξ is increasing in the magnitude of Medicare's deviations from the insurer's efficient pricing. $\square$

B Additional Detail on Implied Conversion Factors

B.1 Data Cleaning

This section describes our process for cleaning and merging the BCBS claims data. Table B.1 shows the data lost as we progress from the raw claims data to the final analysis sample.

For concreteness, consider the 2009 claims data. The data for this year start with 54,724,994 claim lines and \$4.01 billion in spending (row A). To reduce heterogeneity along several administrative margins, we analyze claim lines for which the payment is non-missing, the service quantity is 1, and the observation is an "original" claim line rather than an adjustment to a past payment.[24] This eliminates 5,090,024 claim lines and leaves us with \$3.24 billion in spending (row B). Next, we want to ensure that our analysis focuses on reimbursements for services that are administratively equivalent from a payments perspective, and whose payments have been agreed upon through *ex ante* negotiations. We thus retain only observations that are explicitly coded as being "outpatient" and "in network." These criteria eliminate a total of 8,302,709 claim lines and leave us with \$2.45 billion in spending (row C). Next we drop relatively rare service codes for which we have fewer than 10 observations prior to the RVU updates in a given year. In the 2009 data, this eliminates 149,269 claims and leaves us with \$2.44 billion in spending (row D). The resulting sample of 41,182,992 service lines and \$2.44 billion in spending constitutes the administratively comparable and sufficiently common billing codes we aim to understand.

In order for private insurers to benchmark prices to Medicare, at a minimum they would need to use Medicare's billing codes. On row (E), we thus merge the remaining claims with Medicare billing codes, which provides an upper bound on the potential benchmarking. The final analysis sample in 2009 includes 3,807 unique HCPCS codes, which comprise 21,941,227 service lines and \$1.89 billion of spending. The key conclusion from row (E) 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. More specifically, this merge only loses notable portions of one broad spending category, namely laboratory tests, for which both Medicare and BCBS frequently base payments on non-standard codes. We retain over 97 percent of claims for evaluation and management, diagnostic imaging, and surgical services.

B.2 Heterogeneity by Market Structure

We now consider the distinction between a group's own size and the market structure in which it operates. To begin, we estimate a variant of equation (4) that replaces vigintile fixed effects with a continuous measure of firm size:

$$\text{Medicare-Linked Share}_{j,g} = \eta_b \mathbb{1}_{b(j)} + \varsigma \text{Log Group Billings} + \upsilon_{j,g}. \tag{B.1}$$

[24] Both Medicare and private sector payment policies generate nonlinear payments in certain circumstances when multiple instances of the same service are provided per claim.

This regression summarizes the evidence from Figure 2 in the main text. Column 1 of Appendix Table B.4 shows the estimates of ς. A 10 percent increase in firm size is associated with a 2.5 to 7 percentage point decline in the share of payments benchmarked to Medicare rates.

We next consider heterogeneity in market structure by adding area characteristics to equation (B.1). In column 2 of Table B.4, we first replace the Betos category fixed effects with geographic fixed effects. Specifically, we include indicators for each hospital referral region (HRR), of which Texas has 22. Changing the fixed effects has little impact on the relationship between Medicare benchmarking and individual firm size, suggesting that this relationship was not driven by omitted geographic differences. If anything, the size-benchmarking relationship strengthens slightly in column 2.

We next consider the level of competition among local physician groups. Specifically, we estimate the local Herfinadhl-Hirschman Index (HHI) for each specialty, in each HRR, based on the level of BCBS revenue each group receives in our data. We then add this HHI to the regression from column 2, or in other words we estimate:

$$\text{Medicare-Linked Share}_{j,g} = \eta \text{Log Group Billings}_g + \zeta_b \mathbb{1}_{b(j)} + \gamma \text{HHI}_{j,g} + \upsilon_{j,g}. \tag{B.2}$$

Column 3 shows a small, insignificant negative relationship between HHI and benchmarking.

Column 4 adds an interaction between individual firm size and the specialty-area HHI measure from column 3. That is, it asks whether firm size is more or less important for benchmarking in more concentrated markets. This column shows a much stronger negative estimate on the direct effect of HHI than we observed in column 3; more concentrated markets now seem to have less Medicare benchmarking—at least for the smaller physician groups. The positive interaction term implies that the relationship with group size diminishes as HHI increases, or the relationship with HHI diminshes as group size increases. Although Medicare benchmarking is smaller for larger physician groups, and for those in concentrated markets, each of these effects diminishes as the counterpart increases.

B.3 Levels of ICFs and Group Characteristics

We next examine the levels of the common ICFs (cICFs) that we identify. Figure B.2 shows the distributions of cICFs by year. Table B.5 shows how these values relate to firm size. To avoid a mechanical relationship between the ICF levels and our firm size measure, we measure physician group size as the log number of services provided, rather than the value of billings for those services. Specifically, we estimate:

$$\ln ICF_g = \lambda \, \text{Log Group Services}_g + \upsilon_g.. \tag{B.3}$$

The six columns show two regressions in each year, one that includes all of a firm's cICFs and one that limits the sample to the most common ICF for each firm. The former regression includes standard errors clustered by physician group. All columns show a consistent positive relationship between group size and ICF. A group providing ten percent more services obtains

3 to 5 percent higher ICFs. Table B.6 runs similar regressions, but changes the dependent variable to the level of the ICF rather than its log.

In Table B.7 we consider the relationship between the level of the ICF and the frequency with which it is used. We run regressions of the form:

$$\ln ICF_{g,i} = \gamma_{5\%}\text{Linked}_{g,i}^{5\%} + \gamma_{10\%}\text{Linked}_{g,i}^{10\%} + \gamma_{20\%}\text{Linked}_{g,i}^{20\%} + v_{g,i} \tag{B.4}$$

at the level of group $g\times$ unique ICF i. In equation (B.4), $\text{Linked}_{g,i}^{x\%}$ is an indicator for whether ICF i from group g represents at least x percent of group g's billings. In some regressions, we also control for the group size. Table B.7 does not reveal any particularly clear pattern to this relationship.

Appendix Figure B.1

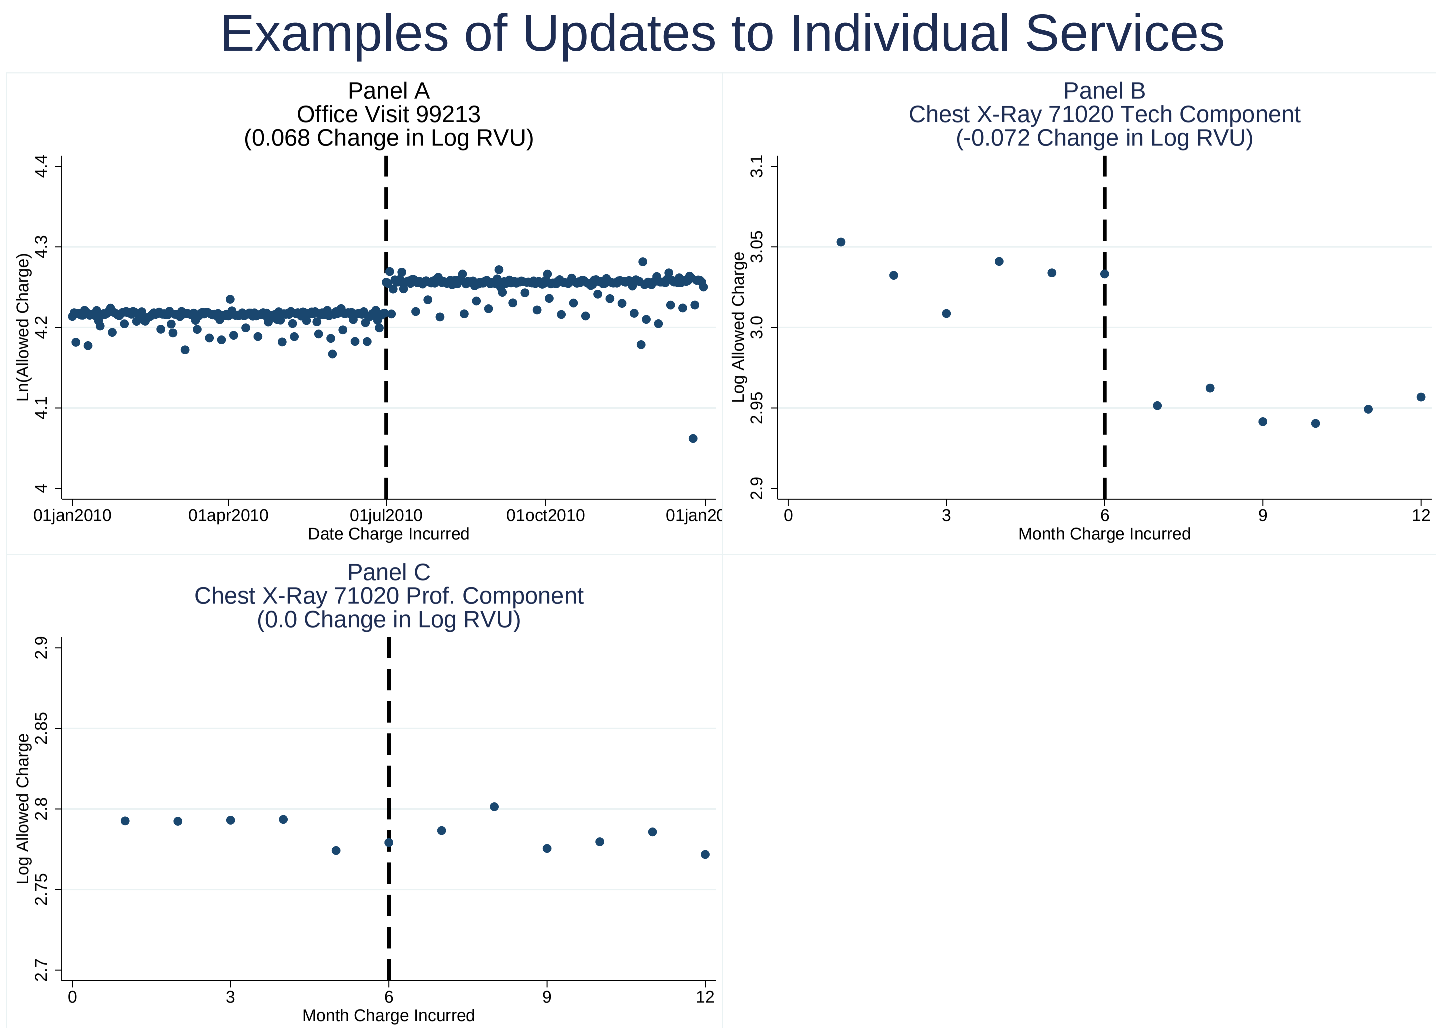

Note: The figure presents monthly averages of BCBS's log payment for the service named in each panel's title. All data are from calendar year 2010. BCBS implemented its update from the 2009 to 2010 relative value scales on July 1, 2010, as indicated by the vertical dashed line each panel.

Appendix Figure B.2: Distribution of ICFs by Firm Size

Panel A: 2008

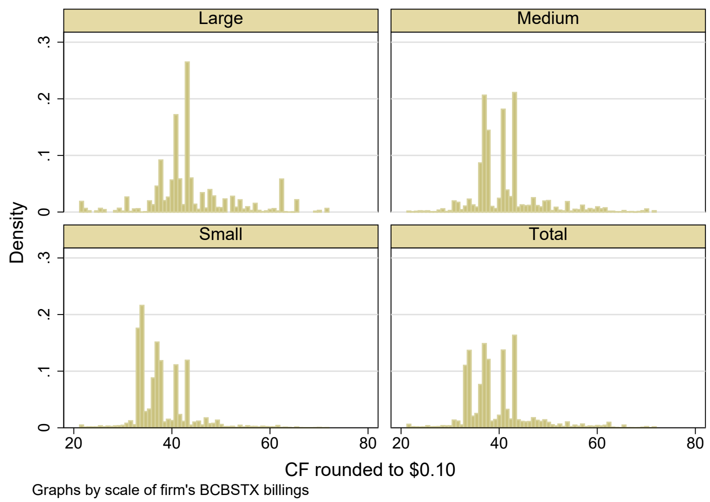

Panel B: 2009

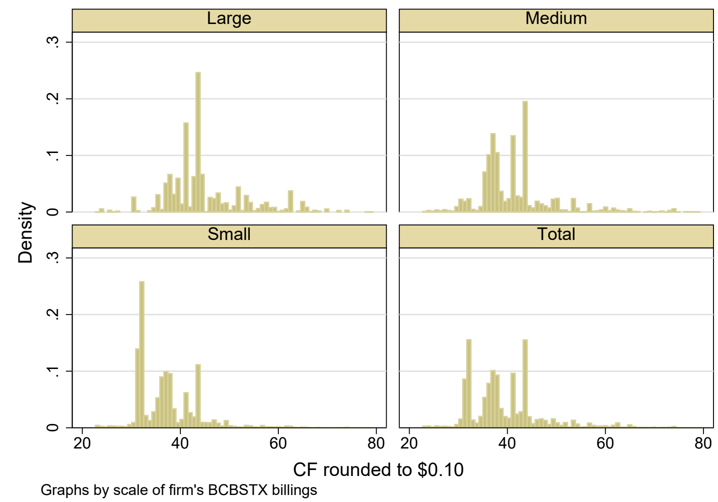

Panel C: 2010

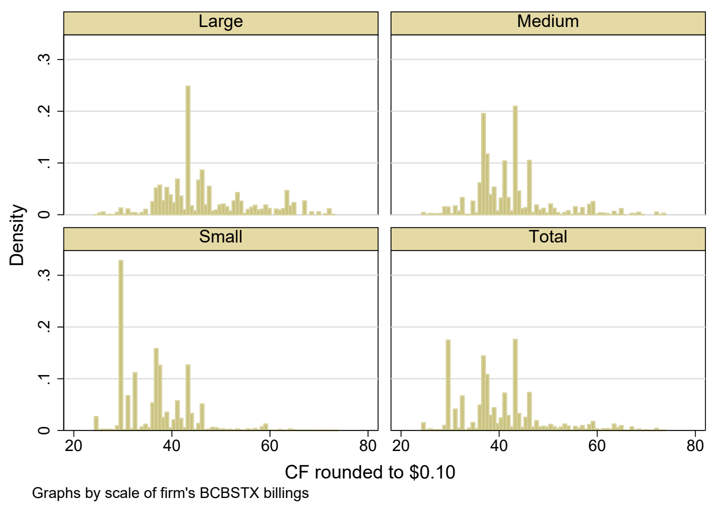

Panel D: 2011

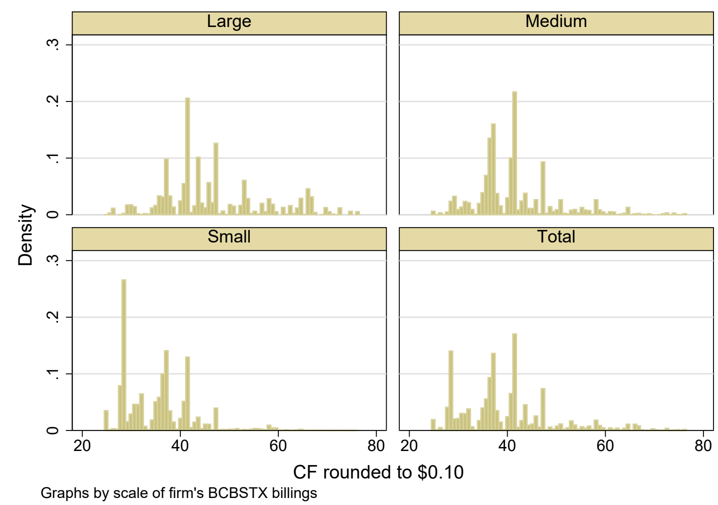

Note: The figure reports the distributions of common Implied Conversion Factors that we compute in each year. Each year's distributions are split according to the sizes of the physician groups, measured as the dollar value of the group's BCBS billings.

Appendix Table B.1: Data Cleaning

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Year:	2008		2009		2010		2011	
Measure:	Claims	Spending	Claims	Spending	Claims	Spending	Claims	Spending
(A) Initial dataset	45.5m	$3.49b	54.7m	$4.09b	57.6m	$4.29b	61.7m	$4.64b
(B) Basic cleaning	90.0%	80.2%	90.7%	80.8%	90.0%	80.0%	90.3%	80.4%
(C) In-network outpatient	74.0%	59.6%	75.5%	61.1%	76.5%	61.5%	77.3%	62.3%
(D) Exclude rare codes	73.9%	59.3%	75.3%	60.8%	76.5%	61.3%	77.3%	62.1%
(E) Medicare code merge	41.3%	47.3%	40.3%	47.1%	41.7%	47.8%	41.3%	47.8%

Note: This table quantifies the data lost at each step of our data cleaning and merge process. We show calculations for each of the four years of BCBS claims data. For each year, row (A) shows the raw number of claims (odd-numbered columns) and money spent (even-numbered columns) in that year's claims data. All subsequent rows show the share of claims on row (A) that remain after each set of cleaning steps. Row (B) shows the share of data remaining when we keep only claim lines for which the payment is non-missing, the service quantity is 1, and the observation is an "original" claim line rather than an adjustment to a past payment. These basic cleaning steps eliminate about ten percent of claims and twenty percent of spending. Row (C) further restricts our sample to the universe we consider, namely outpatient in-network claims. This eliminates approximately 15 percent more claims, and twenty percent more spending per year. Row (D) drops those relatively rare service codes for which we have fewer than 10 observations prior to the RVU updates in a given year; this has minimal effect on the sample sizes. Finally, row (E) drops claims that don't merge with Medicare's RBRVS codes. This loses 12–15 percent of observations per year. Source: Authors' calculations using claims data from BCBS.

Appendix Table B.2: **Alternative Measures of Pricing According to Common Implicit Conversion Factors**

	Panel A: 2008			
Benchmarking Measure:	Services	Dollars	Services Q1	Dollars Q1
Rounding for ICFs:				
$0.02	67%	60%	68%	62%
$0.10	73%	66%	74%	67%
$0.20	77%	71%	78%	72%
	Panel B: 2009			
Benchmarking Measure:	Services	Dollars	Services Q1	Dollars Q1
Rounding for ICFs:				
$0.02	67%	60%	68%	62%
$0.10	73%	66%	74%	67%
$0.20	77%	70%	78%	71%
	Panel C: 2010			
Benchmarking Measure:	Services	Dollars	Services Q1	Dollars Q1
Rounding for ICFs:				
$0.02	87%	83%	88%	84%
$0.10	89%	86%	89%	86%
$0.20	89%	87%	90%	87%
	Panel D: 2011			
Benchmarking Measure:	Services	Dollars	Services Q1	Dollars Q1
Rounding for ICFs:				
$0.02	86%	81%	86%	82%
$0.10	87%	85%	88%	85%
$0.20	88%	85%	88%	85%

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. The different cells within a panel show this statistic according to slightly different measures and using different rounding thresholds to define cICFs. The column labeled "Rounding" indicates the rounding applied to each estimated ICF. We then declare an ICF to be "common" for the payments to a physician group if it accounts for at least 5 percent of the group's services in a given year. The first column shows the share of services priced using cICFs, just as in Table 2. The column labeled "Dollars" shows a dollar-weighted measure. The dollar-weighted estimates are lower than the service-weighted measure because lower-value services are more likely to be priced using common ICFs. The remaining columns report equivalent measures for which the claims data are restricted to the first quarter of a given year. Source: Authors' calculations using claims data from BCBS.

Appendix Table B.3: Medicare Benchmarking by Betos Category

	(1)	(2)	(3)	(4)
Dependent variable:	Payments with Common Conversion Factors			
	Service Share		Spending Share	
	Panel A: 2008 (N=593,779)			
Imaging	-0.155**	-0.243**	-0.174**	-0.258**
	(0.052)	(0.052)	(0.048)	(0.047)
Procedures	-0.183**	-0.282**	-0.191**	-0.287**
	(0.054)	(0.055)	(0.043)	(0.042)
Tests	-0.150**	-0.218**	-0.200**	-0.266**
	(0.054)	(0.057)	(0.044)	(0.045)
Constant	0.603**	0.355**	0.605**	0.365**
	(0.037)	(0.051)	(0.032)	(0.040)
	Panel B: 2009 (N=593,779)			
Imaging	-0.155**	-0.243**	-0.174**	-0.258**
	(0.052)	(0.052)	(0.048)	(0.047)
Procedures	-0.183**	-0.282**	-0.191**	-0.287**
	(0.054)	(0.055)	(0.043)	(0.042)
Tests	-0.150**	-0.218**	-0.200**	-0.266**
	(0.054)	(0.057)	(0.044)	(0.045)
Constant	0.603**	0.355**	0.605**	0.365**
	(0.037)	(0.051)	(0.032)	(0.040)
	Panel C: 2011 (N=651,901)			
Imaging	-0.317**	-0.357**	-0.420**	-0.454**
	(0.032)	(0.026)	(0.053)	(0.046)
Procedures	-0.431**	-0.470**	-0.361**	-0.395**
	(0.059)	(0.052)	(0.030)	(0.026)
Tests	-0.334**	-0.362**	-0.422**	-0.446**
	(0.046)	(0.042)	(0.037)	(0.033)
Constant	0.808**	0.764**	0.799**	0.760**
	(0.023)	(0.020)	(0.014)	(0.016)
Omitted Category	Evaluation & Management			
Additional Controls	Group Size	None	Group Size	None

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. This table shows estimates of the η_b coefficients in equation (4), namely the relationship between Betos category and the Medicare-linked share of services (columns 1 and 2) or spending (columns 3 and 4) at the group-service code level. Medicare links are measured using the common Implied Conversion Factors (cICFs) defined in section 4.1. 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, Gelbach and Miller, 2011) by Betos category and physician group. Sources: Authors' calculations using claims data from BCBS.

Appendix Table B.4: Medicare Benchmarking by Firm Size and HHI

	(1)	(2)	(3)	(4)
Dependent variable:	Share of Payments with Common Conversion Factors			
	Panel A: 2009 (N=438,673)			
Log firm size	-0.071***	-0.052***	-0.049***	-0.060***
	(0.007)	(0.010)	(0.010)	(0.012)
Specialty HHI			-0.072	-0.715**
			(0.069)	(0.231)
Log firm size × HHI				0.053**
				(0.016)
	Panel B: 2010 (N=430,509)			
Log firm size	-0.024***	-0.058***	-0.052***	-0.064***
	(0.004)	(0.007)	(0.005)	(0.005)
Specialty HHI			-0.171	-0.927***
			(0.141)	(0.207)
Log firm size × HHI				0.062***
				(0.010)
	Panel C: 2011 (N=513,590)			
Log firm size	-0.025***	-0.061***	-0.056***	-0.067***
	(0.005)	(0.007)	(0.006)	(0.006)
Specialty HHI			-0.154	-0.850***
			(0.136)	(0.199)
Log firm size × HHI				0.056***
				(0.010)
Fixed Effects	Betos Cat.	HRR	HRR	HRR
Sample	In-network	In-network	In-network	In-network

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. This table shows estimates of equation (B.1), namely the relationship between the Medicare-linked share of services and firm size (columns 1 and 2) and/or market structure (columns 3 and 4) at the group-service code level. Medicare links are measured using the common Implied Conversion Factors (cICFs) defined in section 4.1. 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, Gelbach and Miller, 2011) by Betos category and physician group. Sources: Authors' calculations using claims data from BCBS.

Appendix Table B.5: Log Implicit Conversion Factors

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable:	Log Implicit Conversion Factor							
Log total services by firm	0.022**	0.017**	0.028**	0.022**	0.044**	0.033**	0.048**	0.038**
	(0.003)	(0.004)	(0.003)	(0.004)	(0.004)	(0.003)	(0.005)	(0.004)
Constant	3.511**	3.543**	3.454**	3.498**	3.347**	3.412**	3.281**	3.345**
	(0.021)	(0.026)	(0.019)	(0.023)	(0.027)	(0.020)	(0.034)	(0.027)
N	151,965	44,432	173,356	52,390	317,409	50,963	386,220	65,390
Year	2008	2008	2009	2009	2010	2010	2011	2011
ICFs Included	All ICFs	Top ICF	All ICFs	Top ICF	All ICFs	Top ICF	All ICFs	Top ICF
Standard Errors	Clustered	Robust	Clustered	Robust	Clustered	Robust	Clustered	Robust
Number of Clusters	49,591		58,253		53,848		69,489	

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. Each column shows an estimate of equation (B.3), for different years and different samples of ICFs. In columns 1, 3, and 5, standard errors are clustered by physician group.

Appendix Table B.6: **Levels of Implicit Conversion Factors**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable:	Level of Implicit Conversion Factor							
Log total services by firm	0.927**	0.703**	1.168**	0.889**	1.880**	1.297**	2.040**	1.528**
	(0.162)	(0.197)	(0.146)	(0.173)	(0.215)	(0.160)	(0.256)	(0.211)
Constant	33.103**	34.424**	30.922**	32.795**	26.424**	29.805**	23.794**	27.048**
	(0.989)	(1.195)	(0.935)	(1.104)	(1.394)	(0.995)	(1.744)	(1.377)
N	151,965	44,432	173,356	52,390	317,409	50,963	386,220	65,390
Year	2008	2008	2009	2009	2010	2010	2011	2011
ICFs Included	All ICFs	Top ICF	All ICFs	Top ICF	All ICFs	Top ICF	All ICFs	Top ICF
Standard Errors	Clustered	Robust	Clustered	Robust	Clustered	Robust	Clustered	Robust
Number of Clusters	49,591		58,253		53,848		69,489	

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. Each column shows an estimate of equation (B.3) but with the dependent variable changed to the level of the Implicit Conversion Factor, for different years and different samples of ICFs. In columns 1, 3, and 5, standard errors are clustered by physician group.

Appendix Table B.7: Implicit Conversion Factors and their Frequency

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable:	Log Implicit Conversion Factor							
ICF represents at least 5%	-0.014*	0.004	-0.017*	0.000	0.026	0.051**	-0.021	0.001
of group's claims	(0.007)	(0.005)	(0.007)	(0.006)	(0.017)	(0.013)	(0.020)	(0.016)
ICF represents at least 10%	-0.035**	-0.027**	-0.018**	-0.013*	0.060**	0.066**	0.043*	0.043**
of group's claims	(0.006)	(0.005)	(0.006)	(0.005)	(0.014)	(0.011)	(0.020)	(0.016)
ICF represents at least 20%	0.006	0.008+	-0.019**	-0.009+	-0.062**	-0.055**	-0.049*	-0.024
of group's claims	(0.005)	(0.005)	(0.006)	(0.005)	(0.012)	(0.009)	(0.024)	(0.017)
Log total services by firm		0.029**		0.033**		0.042**		0.044**
		(0.003)		(0.003)		(0.004)		(0.004)
Constant	3.709**	3.483**	3.706**	3.436**	3.645**	3.296**	3.673**	3.288**
	(0.008)	(0.023)	(0.008)	(0.023)	(0.017)	(0.029)	(0.017)	(0.035)
N	626,974	626,974	730,196	730,196	613,586	613,586	790,056	790,056
Year	2008	2008	2009	2009	2010	2010	2011	2011
ICFs Included	All ICFs	All ICFs	All ICFs	All ICFs	All ICFs	All ICFs	All ICFs	All ICFs
Standard Errors	Clustered	Clustered	Clustered	Clustered	Clustered	Clustered	Clustered	Clustered
Number of Clusters	50,367	50,367	59,137	59,137	54,258	54,258	70,090	70,090

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. Each column shows an estimate of equation (B.4), for one of the years in our sample. Standard errors are clustered by physician group.

C Estimation in Changes and Threats to Identification

This appendix justifies our measure of Medicare benchmarking based on estimation in simple differences, in section C.1. Appendix C.2 then discusses potential bias from active renegotiations of physician-insurer contracts contemporaneously with the implementation of Medicar RVU updates. Finally, Appendix C.3 computes the bias that would result in such a case.

C.1 Estimation in Changes

We simplify our main estimating equations to two time periods in order to see the Medicare-private price relationships as transparently as possible. This approach will also clearly highlight the assumptions necessary for our estimate of $\hat{\beta}$ to equal the true Medicare-linked share σ. Averaging equation (10) within each time period, and then taking the difference across the two, yields:

$$\Delta\overline{\ln(P_{g,j})} = \alpha + \beta\Delta\ln(RVU_j) + (1-\sigma)\overline{\varepsilon_{g,j}}. \tag{C.1}$$

In the context of price changes for one service, this equation shows how we can directly interpret the evidence from Figure 3C. This graph showed BCBS average log payments for a standard office visit increasing by 70 percent of the Medicare log RVU change. Hence the implied estimate of σ, in the absence of contemporaneous active negotiations, is also 70 percent.

C.2 Threats to Identification From Active Renegotiations

This interpretation is threatened by the possibility of actively negotiated changes in $\ln(\theta_g)$ and $\ln(\rho_{g,j,p})$, which would show up in the error term. If they also covary with the updates to Medicare's relative values, then our estimate of $\hat{\beta}$ would be biased relative to the true parameter σ. (We compute the bias in Appendix C.3 below.) This might arise endogenously because changes in Medicare's relative values could alter groups' bargaining positions, and perhaps do so differentially across services. We quantify the potential influence of these changes on our estimates of Medicare's bencmarking in two ways.

First, note that when we estimate β on the full sample of physician groups, it could be biased away from σ by active renegotiations of both $\ln(\rho_{g,j,t})$ and $\ln(\theta_{g,t})$. If we estimate β on the data for a single firm, however, $\Delta\ln(\theta_g)$ is a constant. In the levels specification of equation (10), we can similarly account for changes in each group's average log payment by allowing for a full set of group-by-period effects. If estimates of β change little as a result of adding firm-by-period effects to such a specification, we can rule out the possibility that changes in the overall level of each firm's payments are biasing our attempt to recover σ.

Second, the channel through which active renegotiations might bias our attempt to recover σ involves changes in bargaining power *induced* by the RVU changes.[25] The threat to

[25] Actively negotiated payment changes that are driven by the RVU updates themselves may plausibly

our estimation takes the following form: BCBS may pursue renegotiations with firms whose average Medicare payment has fallen, with these negotiations resulting in declines in their payments. Similarly, physician groups whose average Medicare payment has increased may pursue renegotiations with BCBS, with these negotiations resulting in increases in their payments. This pattern would imply a positive bias to our estimates of σ. To investigate the potential relevance of this source of bias, we first construct the average change in the RVUs for the specific services provided by each firm. This allows us to gauge the extent to which each firm is affected. We then investigate whether we obtain larger estimates $\hat{\beta}$ on a sample of firms that were significantly affected compared with firms that experienced little change in their average RVUs.

C.3 Deriving the Bias in our Medicare Link Estimate

The biased coefficient $\hat{\beta}$ we would estimate from equation (C.1) in the presence of simultaneous updates to non-benchmarked prices or group-specific markups is:

$$\begin{aligned}
\hat{\beta} &= \frac{\text{Cov}[\Delta\overline{\ln(P_{g,j})}, \Delta\ln(RVU_j)]}{\text{Var}[\Delta\ln(RVU_j)]} \\
&= \frac{\text{Cov}[\sigma\Delta\overline{\ln(\phi_g)} + \sigma\Delta\ln(RVU_j) + (1-\sigma)\Delta\overline{\ln(\rho_{g,j})} + \Delta\epsilon_{g,j,p}, \Delta\ln(RVU_j)]}{\text{Var}[\Delta\ln(RVU_j)]} \\
&= \sigma\frac{\text{Cov}[\Delta\ln(RVU_j), \Delta\ln(RVU_j)]}{\text{Var}[\Delta\ln(RVU_j)]} + \sigma\frac{\text{Cov}[\Delta\overline{\ln(\phi_g)}, \Delta\ln(RVU_j)]}{\text{Var}[\Delta\ln(RVU_j)]} \\
&+ (1-\sigma)\frac{\text{Cov}[\Delta\overline{\ln(\rho_{g,j})}]}{\text{Var}[\Delta\ln(RVU_j)]} + \frac{\text{Cov}[\Delta\epsilon_{g,j,p}, \Delta\ln(RVU_j)]}{\text{Var}[\Delta\ln(RVU_j)]} \\
&= \sigma + \sigma\frac{\text{Cov}[\Delta\overline{\ln(\phi_g)}, \Delta\ln(RVU_j)]}{\text{Var}[\Delta\ln(RVU_j)]} + (1-\sigma)\frac{\text{Cov}[\Delta\overline{\ln(\rho_{g,j})}, \Delta\ln(RVU_j)]}{\text{Var}[\Delta\ln(RVU_j)]},
\end{aligned} \tag{C.2}$$

where the third equality follows from the properties of covariances and the fourth from the fact that $\frac{\text{Cov}[\Delta\ln(RVU_{j,t}), \Delta\ln(RVU_j)]}{\text{Var}[\Delta\ln(RVU_j)]} = 1$ and $\frac{\text{Cov}[\Delta\epsilon_{g,j,t}, \Delta\ln(RVU_j^M)]}{\text{Var}[\Delta\ln(RVU_j)]} = 0$.

One separate source of bias in the estimate of $\hat{\beta}$ could arise if the linked share σ varies across firms and services. This would imply additional terms in equation (C.2) describing our regression estimates, involving covariances between the RVU updates used for identification and the service-by-group linked shares $\sigma_{j,g}$. Recovering σ also requires us to assume that these covariance terms are 0, which will be true if updates to Medicare's rates are uncorrelated with the $\sigma_{j,g}$. In section 6.2, we will allow for heterogeneity across various dimensions in the linked shares.

covary with these changes. There is no *a priori* reason to suspect that changes renegotiated for other reasons would covary with the RVU updates and bias our estimates.

C.4 Checks for the Relevance of Active Contract Renegotiation

The estimates presented in Figure 3 and Table 4 may differ from the true Medicare benchmarking parameter σ if changes in other terms of providers' contracts covary with the changes in RVUs. Indeed, payment changes that significantly alter physician groups' average Medicare payment can move private payments in subsequent years, due in part to the resulting changes to their bargaining positions (Clemens and Gottlieb, forthcoming). 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.

The most relevant institutional detail is the relatively short time horizon of our event studies. Dunn and Shapiro (2015) report that physician contracts tend to remain in force for around 3 years. Within each of our single-year event studies, we thus anticipate that roughly one-third of the groups in our sample engage in active contract re-negotiations, which could affect our estimates. Unlike the payment changes analyzed by Clemens and Gottlieb (forthcoming), which significantly shifted certain specialties' average Medicare payments, those we consider here are relatively diffused across specialties, so unlikely to affect groups' overall outside options.

Nevertheless, we investigate the potential relevance of active contract renegotiation with two analyses. First, we consider the potential effect of scheduled RVU changes on a firm's bargaining position. We construct a variable that, for each firm, reports the average change in RVUs for the services it provides. Firms experiencing a negative average change have seen their bargaining positions deterioriate. Firms experiencing an average RVU increase have seen their bargaining positions improve. Using the average RVU change to which each firm was exposed, we construct an indicator for groups whose bargaining positions were significantly affected.

Second, we investigate the potential relevance of changes in groups' average log reimbursement by adding full sets of group-by-period fixed effects to our specification. For this regression, we restrict our sample to the 100 largest firms in each year, primarily for computational ease. Note, however, that large firms are precisely those for which we would expect active renegotiations to be most frequent.

Table C.3 presents these results. Column 1 reports our baseline specification, unchanged from Table 4. Column 2 allows our coefficient of interest to vary with an indicator for whether a firm's average Medicare reimbursement rate was significantly affected by a year's RVU updates. The point estimate on this interaction varies across years, but is negative in each case. This is the opposite of what we would expect if significant RVU updates were driving active contract renegotiations. Column 3 limits the baseline specification to the services provided by the 100 largest physician groups. A comparison of column 3 with column 1 reveals that, on average across the years we analyze, the largest firms have contracts that are less linked to Medicare than are contracts in the full sample, a result that we explore further in section 6.3. Most relevant for our current purposes, however, column 4 reveals that adding group-by-period effects to the previous specification has essentially no impact on our coefficient of interest. These results provide evidence against the concern that that active contract renegotiations confound the relationship between BCBS's and Medicare's payments over the intervals we analyze. Thus they bolster the case for interpreting our estimates of $\hat{\beta}$ as unbiased estimates of the fraction of services tied directly to Medicare.

Appendix Figure C.1: Strength of Public Private Payment Relationships

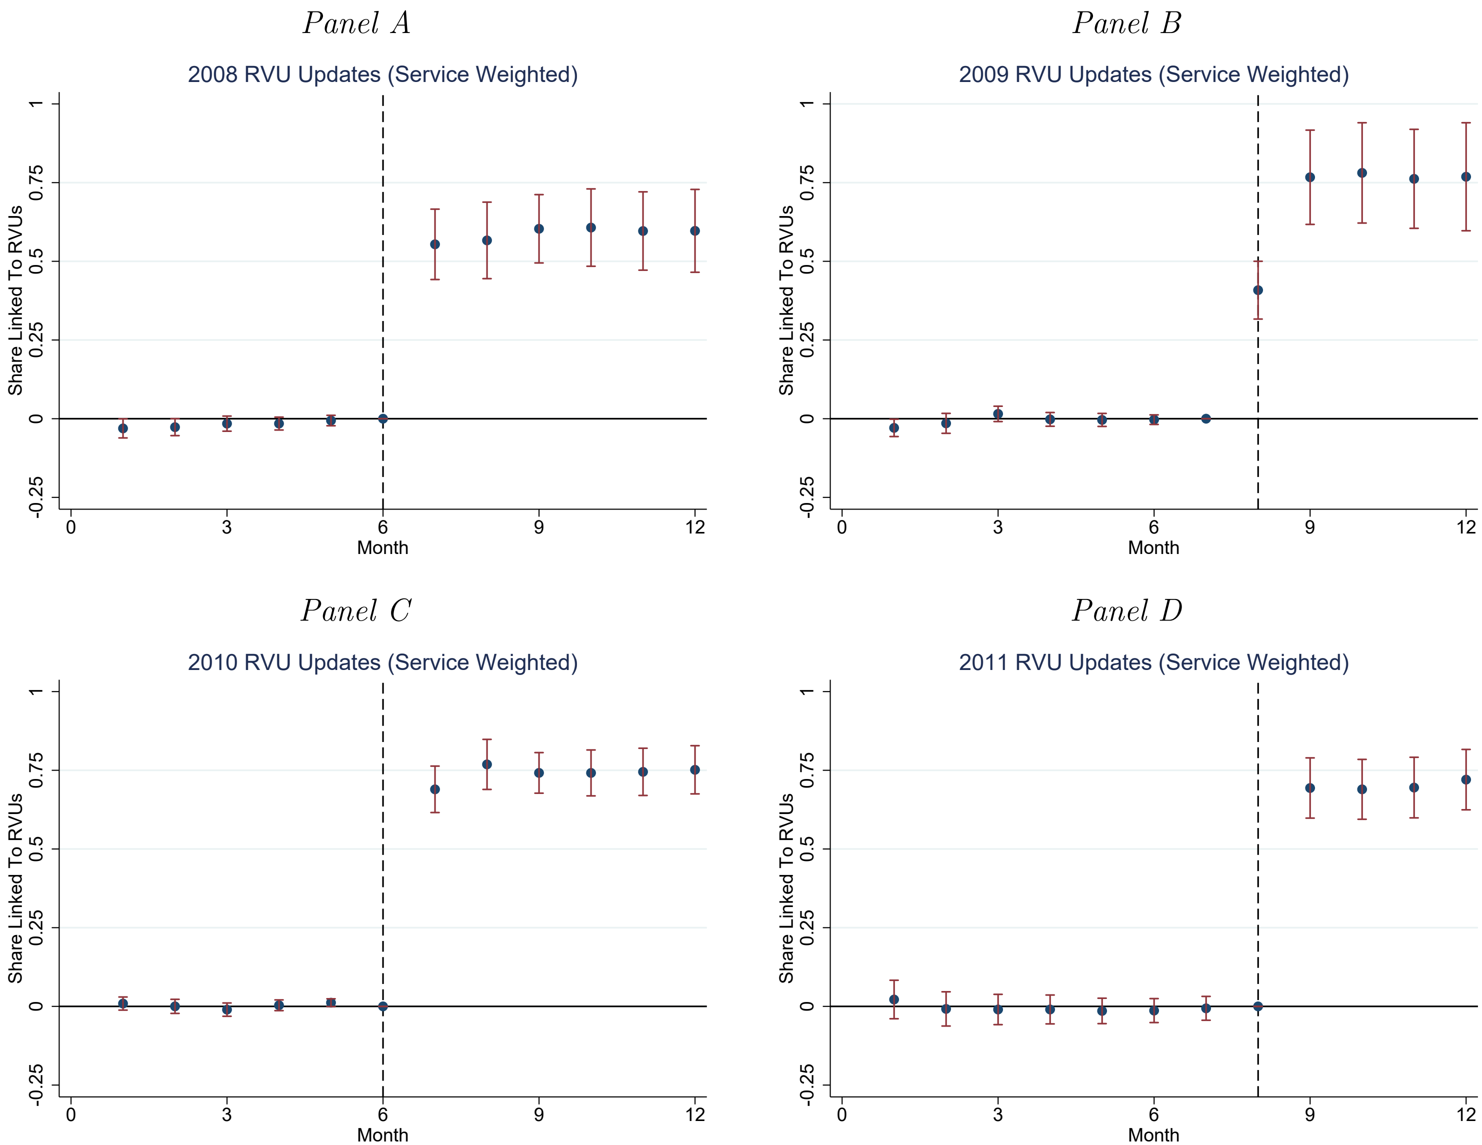

Note: The figure reports estimates of the β_p from estimates of equation (11). The vertical dashed line in each panel corresponds with the month during each year in which BCBS implemented its update from the prior year's relative value scale. These updates occurred on July 1, 2008, August 15, 2009, July 1, 2010, and September 1, 2011.

Appendix Figure C.2: Benchmarking Estimates Based on Price Changes Across Services

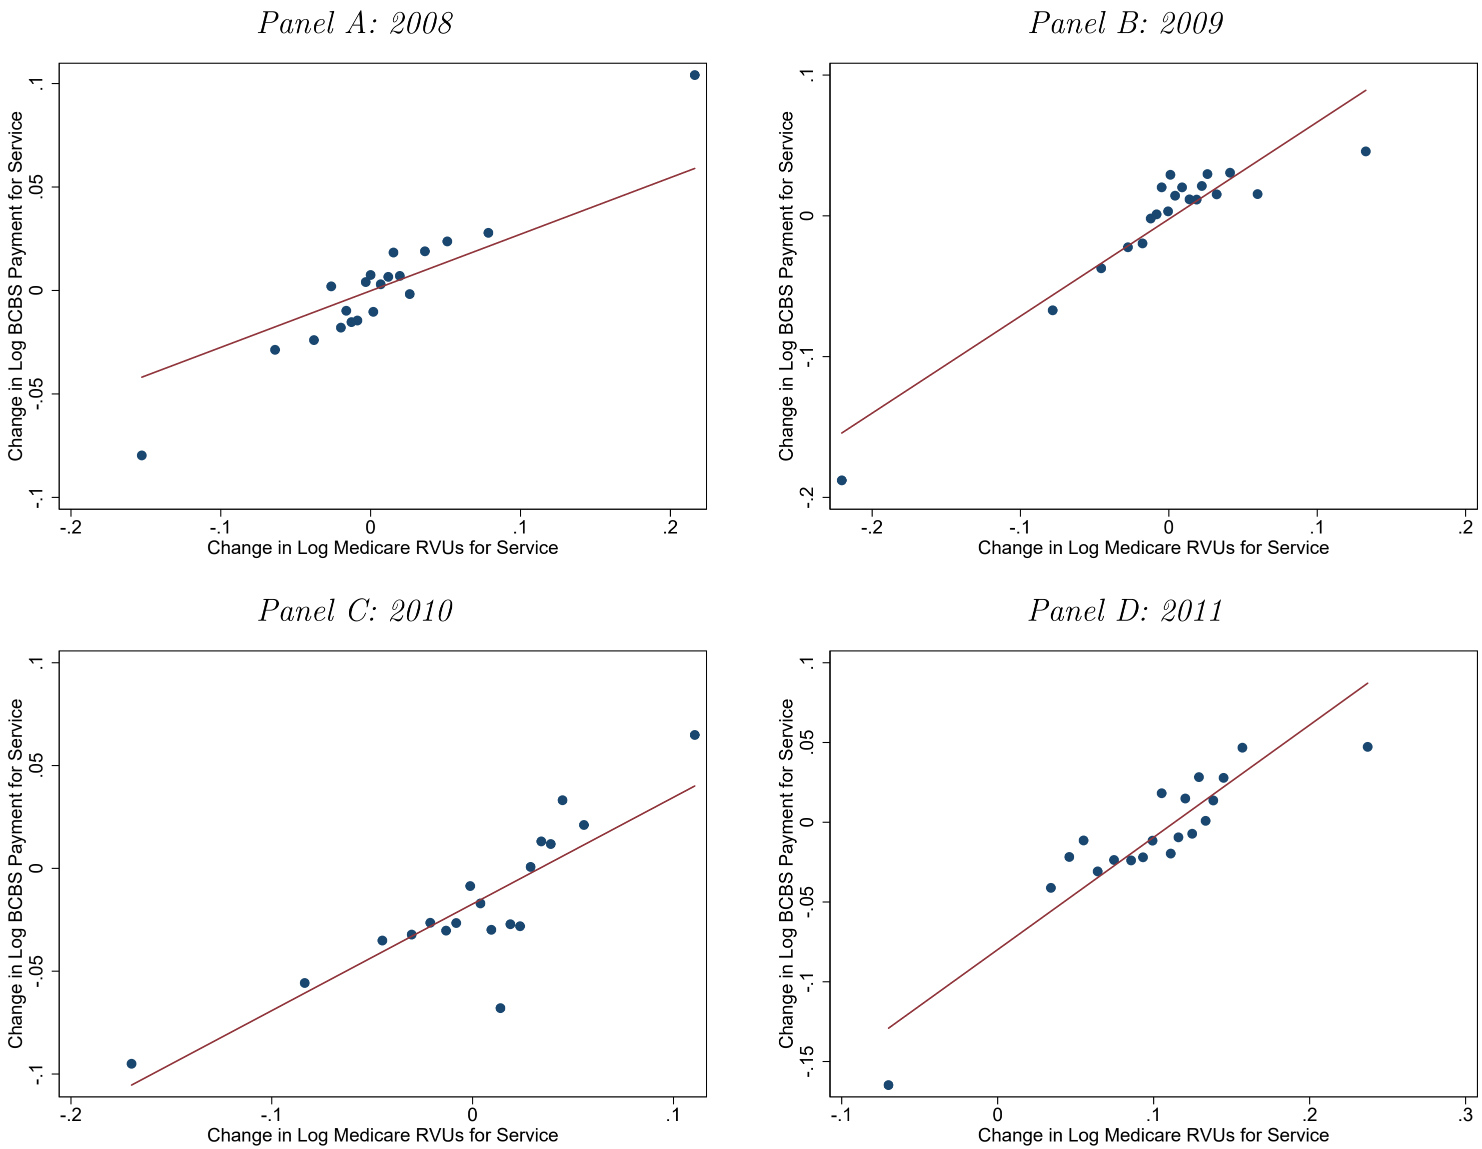

Note: The figure reports the relationships described by equation (C.1) for RVU updates in each year, and estimates of that equation. Each panel shows a separate year's estimates, measured as log differences between the period before BCBS implemented the Medicare RVU updates and the period after this update. The years are split at July 1, 2008, August 15, 2009; July 1, 2010; and September 1, 2011. The regressions are run at the underlying service level, but observations are grouped into twenty bins for each year, based on vigintiles of the Medicare log RVU change.

Appendix Table C.1: **Other Years' Estimates of Benchmarking Using RVU Changes**

	(1)	(2)	(3)	(4)
Dependent variable:	*Log private reimbursement rate*			
	Panel A: All Services: 2008 RVU Updates			
Log RVU Change × Post	0.602**	0.597**	0.539**	0.602**
	(0.061)	(0.061)	(0.060)	(0.061)
N	19,552,096	19,552,096	19,552,096	19,552,096
No. of Clusters	3,505	3,505	3,505	3,505
	Panel B: All Services: 2009 RVU Updates			
Log RVU Change × Post	0.778**	0.778**	0.792**	0.778**
	(0.081)	(0.078)	(0.070)	(0.081)
N	21,941,227	21,941,227	21,941,227	21,941,227
No. of Clusters	3,807	3,807	3,807	3,807
	Panel C: All Services: 2011 RVU Updates			
Log RVU Change × Post	0.704**	0.689**	0.679**	0.704**
	(0.046)	(0.052)	(0.048)	(0.046)
N	25,404,007	25,404,007	25,404,007	25,404,007
No. of Clusters	4,091	4,091	4,091	4,091
Group-by-Code Effects	Yes	No	Yes	Yes
Code Effects	No	Yes	No	No
Cubic Time x RVU Change	No	No	Yes	No
Cubic Time x Post	No	No	No	Yes
Weighting	Service	Service	Service	Service

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. The table shows the results of OLS specifications of the forms described in section 5.2. Each column in each panel reports an estimate of $\hat{\beta}$ from equation (10). Panel A shows estimates using RBRVS updates and BCBS claims data for 2008, Panel B for 2009, and Panel C for 2011. Observations are at the claim-line level and are equally weighted. 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. Sources: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.

Appendix Table C.2: **Dollar-Weighted Estimates of Benchmarking Using RVU Changes**

	(1)	(2)	(3)	(4)
Dependent variable:	*Log private reimbursement rate*			
	Panel A: All Services: 2008 RVU Updates			
Log RVU Change × Post	0.421**	0.413**	0.359**	0.420**
	(0.075)	(0.075)	(0.071)	(0.075)
N	19,552,096	19,552,096	19,552,096	19,552,096
No. of Clusters	3,505	3,505	3,505	3,505
	Panel B: All Services: 2009 RVU Updates			
Log RVU Change × Post	0.618**	0.627**	0.669**	0.618**
	(0.046)	(0.045)	(0.052)	(0.046)
N	21,941,227	21,941,227	21,941,227	21,941,227
No. of Clusters	3,807	3,807	3,807	3,807
	Panel C: All Services: 2011 RVU Updates			
Log RVU Change × Post	0.749**	0.739**	0.738**	0.749**
	(0.044)	(0.043)	(0.047)	(0.044)
N	25,404,007	25,404,007	25,404,007	25,404,007
No. of Clusters	4,091	4,091	4,091	4,091
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
Weighting	Dollars	Dollars	Dollars	Dollars

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. The table shows the results of OLS specifications of the forms described in section 5.2. Each column in each panel reports an estimate of $\hat{\beta}$ from equation (10). Panel A shows estimates using RBRVS updates and BCBS claims data for 2008, Panel B for 2009, and Panel C for 2011. Observations are at the claim-line level and are weighted according to each service's average payment during the baseline period. 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. Sources: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.

Appendix Table C.3: Checks for the Relevance of Active Contract Negotiations

	(1)	(2)	(3)	(4)
Dependent variable:	*Log private reimbursement rate*			
	Panel A: All Services: 2009 RVU Updates			
Log RVU Change × Post	0.778**	0.847**	0.696**	0.666**
	(0.081)	(0.085)	(0.093)	(0.081)
Log RVU Change × Post × Update Impact		-0.077		
		(0.114)		
N	21,941,227	21,941,227	4,097,283	4,097,283
No. of Clusters	3,807	3,807	3,496	3,496
	Panel B: All Services: 2010 RVU Updates			
Log RVU Change × Post	0.750**	0.992**	0.740**	0.747**
	(0.038)	(0.076)	(0.048)	(0.052)
Log RVU Change × Post × Update Impact		-0.393**		
		(0.099)		
N	23,933,577	23,933,577	4,708,213	4,708,213
No. of Clusters	3,681	3,681	3,450	3,450
	Panel C: All Services: 2011 RVU Updates			
Log RVU Change × Post	0.704**	0.804**	0.544**	0.523**
	(0.046)	(0.084)	(0.051)	(0.067)
Log RVU Change × Post × Update Impact		-0.162		
		(0.106)		
N	25,404,007	25,404,007	5,069,260	5,069,260
No. of Clusters	4,091	4,091	3,825	3,825
Group × Post-Update Effects	No	No	No	Yes
Sample	Full	Full	Largest Firms	Largest Firms

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. The table shows the results of OLS specifications of the forms described in section 5.2. Column 1 replicates the baseline specification from column 1 of Table 4. Column 2 augments the baseline specification with interaction terms allowing the effect of RVU updates to vary with the extent of the average impact of each year's RVU updates on a physician group's average Medicare reimbursement rate. In columns 3 and 4 the sample is restricted to each year's 100 largest physician groups, as sorted by total bills submitted. The specification in column 3 is the baseline specification, while the specification in column 4 includes a full set of post-by-group interactions. Panel A shows estimates using RBRVS updates and BCBS claims data for 2009, Panel B for 2010, and Panel C for 2011. Observations are at the claim-line level and are equally weighted. 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. Sources: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.

Appendix Table C.4: Public-Private Payment Links Across Service Categories

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent variable:	*Log private reimbursement rate*						
	Panel A: 2008 RVU Updates by Betos Categories						
	Evaluation	Imaging	Procedures	Tests	Imaging Sub-Categories:		
					Global	Technical	Professional
Log RVU Change	0.541***	0.644***	0.495***	0.786***	0.665***	0.494***	0.945***
× Post-Update	(0.115)	(0.092)	(0.116)	(0.055)	(0.103)	(0.112)	(0.228)
N	9,851,995	3,221,634	3,851,609	1,292,912	1,688,102	192,569	1,340,963
No. of Clusters	207	1,069	1,817	385	400	235	434
	Panel B: 2009 RVU Updates by Betos Categories						
	Evaluation	Imaging	Procedures	Tests	Imaging Sub-Categories:		
					Global	Technical	Professional
Log RVU Change	0.857**	0.775**	0.399**	0.933**	0.702**	0.769**	0.680**
× Post-Update	(0.209)	(0.066)	(0.064)	(0.052)	(0.072)	(0.068)	(0.184)
N	11,498,770	3,524,642	3,861,539	1,449,803	1,769,522	222,026	1,533,094
No. of Clusters	219	1,133	2,036	388	422	262	449
	Panel C: 2011 RVU Updates by Betos Categories						
	Evaluation	Imaging	Procedures	Tests	Imaging Sub-Categories:		
					Global	Technical	Professional
Log RVU Change	0.794**	0.616**	0.900**	0.439*	0.816**	0.692**	0.709**
× Post-Update	(0.065)	(0.100)	(0.075)	(0.221)	(0.048)	(0.067)	(0.058)
N	13,116,657	3,696,733	5,233,336	1,659,485	1,929,095	193,577	1,574,061
No. of Clusters	238	1,143	2,246	436	424	264	455

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. The table shows the results of OLS specifications of the forms described in section 5.2. The cells in each panel report estimates of $\hat{\beta}$ from equation (10), with samples selected to contain the HCPCS codes falling into individual broad service categories. The name of the relevant service category accompanies each point estimate. Panel A shows estimates using RBRVS updates and BCBS claims data for 2008, Panel B for 2009, and Panel C for 2011. 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. Sources: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.

Appendix Table C.5: Medicare Benchmarking by Firm Size

	(1)	(2)	(3)	(4)
Dependent variable:	*Log private reimbursement rate*			
Panel A: 2008 RVU Updates (N = 19,552,096)				
Log RVU Change	0.602**	0.560**	0.421**	0.418**
× Post-Update	(0.061)	(0.074)	(0.075)	(0.089)
Log RVU Change		0.130*		-0.059
× Post-Update × Midsize		(0.065)		(0.072)
Log RVU Change		-0.000		0.064
× Post-Update × Large		(0.101)		(0.085)
Panel B: 2009 RVU Updates (N = 21,941,227)				
Log RVU Change	0.778**	0.755**	0.618**	0.756**
× Post-Update	(0.081)	(0.090)	(0.046)	(0.070)
Log RVU Change		0.078		-0.110
× Post-Update × Midsize		(0.059)		(0.071)
Log RVU Change		-0.035		-0.271*
× Post-Update × Large		(0.094)		(0.109)
Panel C: 2011 RVU Updates (N = 25,404,007)				
Log RVU Change	0.704**	0.812**	0.749**	0.774**
× Post-Update	(0.046)	(0.063)	(0.044)	(0.052)
Log RVU Change		-0.140+		-0.036
× Post-Update × Midsize		(0.075)		(0.100)
Log RVU Change		-0.183*		-0.023
× Post-Update × Large		(0.075)		(0.116)
Firm Size × Post-Update Controls	No	Yes	No	Yes
Weighting	Services	Services	Dollars	Dollars

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. 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 indicators 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. 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). Sources: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.

D Results for Out-of-Network Payments

This appendix presents analogues of our baseline estimates, but for payments to out-of-network physicians. This analysis allows us to determine whether the benchmarking that we document reflects active decisions as opposed to a purely mechanical force. Table D.1 replicates Table 4 in the main text, but for out-of-network payments. Table D.2 is a dollar-weighted version of the same regressions. In both cases, we obtain small and precisely estimated coefficients. This means that out-of-network payments—which don't represent the outcome of the *ex ante* negotiations we described in section 1.2—are not priced in the same way.

Table D.3 complicates the analysis somewhat. It reveals that around half of out-of-network services appear to be priced according to cICFs. This share is much larger than the results from Tables D.1 and D.2 would suggest, though still far below the in-network results from Table 2 in the main text. The difference with the in-network results is especially pronounced in 2010 and 2011, and when using a more stringent cICF threshold (20 percent). In these cases, only 30 percent of out-of-network prices appear to be benchmarked to Medicare, compared with 70 percent of in-network payments. Nevertheless, the ambiguity over the correct definition again demonstrates the advantage of the update-based benchmarking measure in Tables D.1 and D.2.

Appendix Table D.1: Estimating Medicare Benchmarking for Out-of-Network Payments Using RVU Changes

	(1)	(2)	(3)	(4)
Dependent variable:	*Log private reimbursement rate*			
	Panel A: All Services: 2009 RVU Updates			
Log RVU Change × Post	0.018	0.007	0.084*	0.018
	(0.043)	(0.047)	(0.037)	(0.043)
N	2,585,681	2,585,681	2,585,681	2,585,681
No. of Clusters	2,456	2,456	2,456	2,456
	Panel B: All Services: 2010 RVU Updates			
Log RVU Change × Post	0.302**	0.351**	0.170**	0.302**
	(0.073)	(0.074)	(0.044)	(0.073)
N	2,386,575	2,386,575	2,386,575	2,386,575
No. of Clusters	2,051	2,051	2,051	2,051
	Panel C: All Services: 2011 RVU Updates			
Log RVU Change × Post	0.106*	0.094+	0.047	0.105*
	(0.047)	(0.054)	(0.037)	(0.047)
N	2,626,264	2,626,264	2,626,264	2,626,264
No. of Clusters	2,473	2,473	2,473	2,473
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
Weighting	Service	Service	Service	Service

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. The table shows the results of OLS specifications of the forms described in section 5.2, except using data from out-of-network payments. Each column in each panel reports an estimate of $\hat{\beta}$ from equation (10). Panel A shows estimates using RBRVS updates and BCBS claims data for 2009, Panel B for 2010, and Panel C for 2011. Observations are at the claim-line level and are equally weighted. 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. Sources: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.

Appendix Table D.2: Dollar-Weighted Estimates of Medicare Benchmarking for Out-of-Network Payments Using RVU Changes

	(1)	(2)	(3)	(4)
Dependent variable:	*Log private reimbursement rate*			
	Panel A: All Services: 2009 RVU Updates			
Log RVU Change × Post	0.036	0.004	-0.043	0.036
	(0.048)	(0.055)	(0.079)	(0.048)
N	2,585,681	2,585,681	2,585,681	2,585,681
No. of Clusters	2,456	2,456	2,456	2,456
	Panel B: All Services: 2010 RVU Updates			
Log RVU Change × Post	0.244**	0.315**	0.203*	0.242**
	(0.063)	(0.066)	(0.082)	(0.063)
N	2,386,575	2,386,575	2,386,575	2,386,575
No. of Clusters	2,051	2,051	2,051	2,051
	Panel C: All Services: 2011 RVU Updates			
Log RVU Change × Post	-0.016	-0.045	0.053	-0.016
	(0.068)	(0.075)	(0.066)	(0.067)
N	2,626,264	2,626,264	2,626,264	2,626,264
No. of Clusters	2,473	2,473	2,473	2,473
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
Weighting	Dollars	Dollars	Dollars	Dollars

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. The table shows the results of OLS specifications of the forms described in section 5.2, except using data from out-of-network payments. Each column in each panel reports an estimate of $\hat{\beta}$ from equation (10). Panel A shows estimates using RBRVS updates and BCBS claims data for 2009, Panel B for 2010, and Panel C for 2011. Observations are at the claim-line level and are equally weighted. 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. Sources: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.

Appendix Table D.3: **Out-of-Network Services Priced According to Common Implied Conversion Factors**

	Panel A: 2009		
	Frequency Threshold:		
	5%	*10%*	*20%*
Rounding for ICFs:			
$0.02	54%	42%	26%
$0.10	60%	46%	30%
$0.20	64%	52%	35%
	Panel B: 2010		
	Frequency Threshold:		
	5%	*10%*	*20%*
Rounding for ICFs:			
$0.02	57%	45%	32%
$0.10	61%	48%	34%
$0.20	65%	52%	37%
	Panel C: 2011		
	Frequency Threshold:		
	5%	*10%*	*20%*
Rounding for ICFs:			
$0.02	57%	43%	29%
$0.10	61%	47%	32%
$0.20	66%	51%	35%

Note: Each cell shows the share of out-of-network services for which payments are associated with a common Implied Conversion Factor (cICF), as defined in the main text. The cells within each panel show how this 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.

E Cross-Sectional vs. RVU-Update Approaches

This appendix motivates and presents the results of an analysis that allows us to compare the Medicare price links we estimate using our cross-sectional and update-based approaches. We begin by developing a cross-sectional metric for deviations from Medicare's pricing menu. We then combine this metric with our changes-based approach to examine whether the services that appear to receive Medicare-benchmarked payments in the cross-section also follow Medicare's RVU updates.

E.1 Testing Consistency of Medicare Links

Section 7 presented an estimate of cross-sectional relationships between Medicare and private payments, and focused on the directions of the residuals from equation (12). Aside from the directions, these prediction errors across services and groups also contain information about the frequency and magnitude of deviations from Medicare's relative values.

Figure E.1 illustrates these errors. The three colors of dots illustrate the different magnitudes of this regression's prediction errors, allowing us to investigate how services in these different categories respond to RVU updates.

We use these categories to test whether the cross-sectional errors $\hat{e}_{g,j}$ are consistently related to BCBS's benchmarking to Medicare payments. We construct a variable that, for each service j, contains the average of the absolute value of the prediction errors $\hat{e}_{g,j}$. That is, for each service we estimate $\overline{|\hat{e}_j|} = \sum_g |\hat{e}_{g,j}| / N_j$ where N_j is the number of times service j occurs in the sample. We then estimate our baseline specification on sub-samples split based on these average prediction errors. We also estimate a full-sample specification in which we allow for an interaction between $\overline{|\hat{e}_j|}$ and changes in Medicare's relative values. That is, we estimate

$$\begin{aligned}\overline{\ln(P_{c,g,j,t})} &= \psi \Delta \ln(RVU_j) \cdot \mathbb{1}_{\{t=\text{post}\}} + \xi \Delta \ln(RVU_j) \cdot \mathbb{1}_{\{t=\text{post}\}} \cdot \overline{|\hat{e}_j|} + \gamma \mathbb{1}_{\{t=\text{post}\}} \cdot \overline{|\hat{e}_j|} \\ &\quad + \phi_t \mathbb{1}_t + \phi_j \mathbb{1}_j + \phi_g \mathbb{1}_g + \phi_{g,j} \mathbb{1}_g \cdot \mathbb{1}_j + \epsilon_{c,g,j,t}. \end{aligned} \tag{E.1}$$

If services that are farther from the Medicare prediction line in the cross section are unlinked from RVU updates, then we would expect to estimate $\hat{\xi} < 0$. If the apparent cross-sectional links are unrelated to whether a service follows Medicare updates, we would estimate $\hat{\xi} = 0$.

E.2 Consistency With Cross-Sectional Links to Medicare Payments

Table E.1 presents estimates generated using the approach discussed above. In column 1, we restrict the sample to services with below-median (absolute value of) average cross-sectional prediction errors. That is, we restrict the samples to the services for which relative payments appear to hew closely to Medicare's relative values in the cross-section. Column 2 restricts the sample to services falling between the 50th and 90th percentiles of the distribution of prediction errors, while column 3 contains services in the top decile. Figure E.2 illustrates this difference graphically, with a binned scatterplot that splits the sample at the median absolute prediction error. Column 4 presents the full sample specification, equation (E.1), with the interaction term.

The results generally reveal a strong relationship between the average magnitude of the cross-sectional prediction errors and the private payment response to changes in Medicare's relative values. This relationship is particularly strong in the data for 2009 and 2010. In these years, the results in column 1 suggest that nearly all of the payments made for services with small cross-sectional residuals were linked to Medicare's relative values. The share is substantially smaller for the services analyzed in column 2, and smaller still for those analyzed in column 3. In the 2010 sample, the estimates suggest that around half of payments are linked directly to Medicare's relative values. The relationship between the cross-sectional residuals and the strength of the links between private payments and changes in Medicare's relative values appears much weaker in the 2011 sample. The cross-sectional prediction errors have fairly strong power for predicting heterogeneity in our estimates of the link between private payments and changes in Medicare's relative values.

Appendix Figure E.1

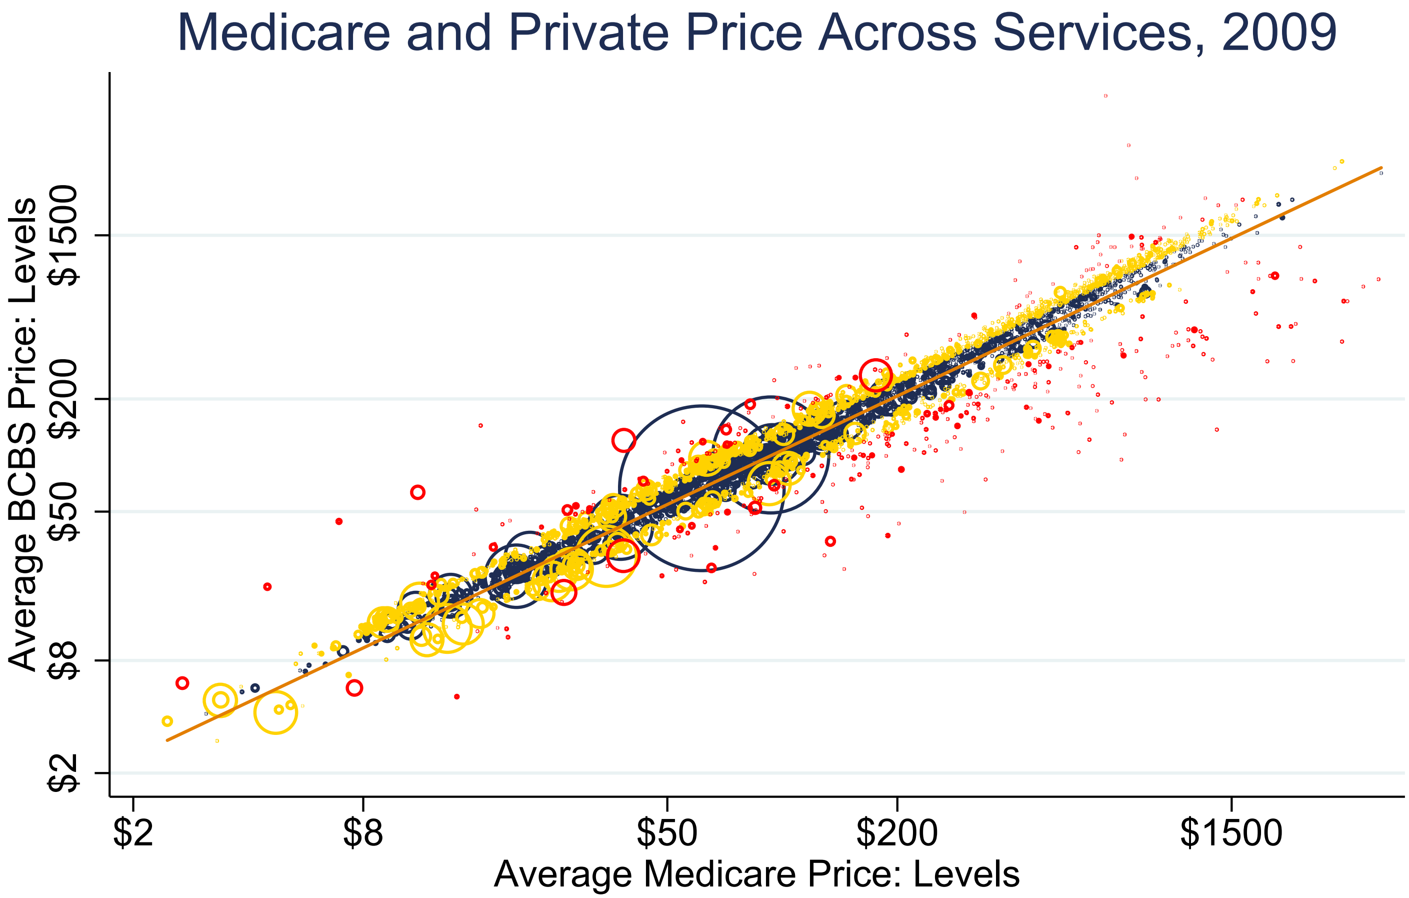

Note: Circle sizes are proportional to the number of claims a service appeared in the data.
Colors indicate magnitudes of residuals
Regression Line: Log BCBS Price = 0.963 (0.003) x Log Medicare Price + 0.23. R^2: 0.96.

Note: The figure presents the cross-sectional correlation between Medicare and BCBS reimbursement rates in 2009. Medicare reimbursement rates are calculated using each HCPCS code's 2009 allocation of relative value units, multiplied by the 2009 national conversion factor. BCBS payments are calculated as HCPCS code average across all service lines in our analysis sample.

Appendix Figure E.2

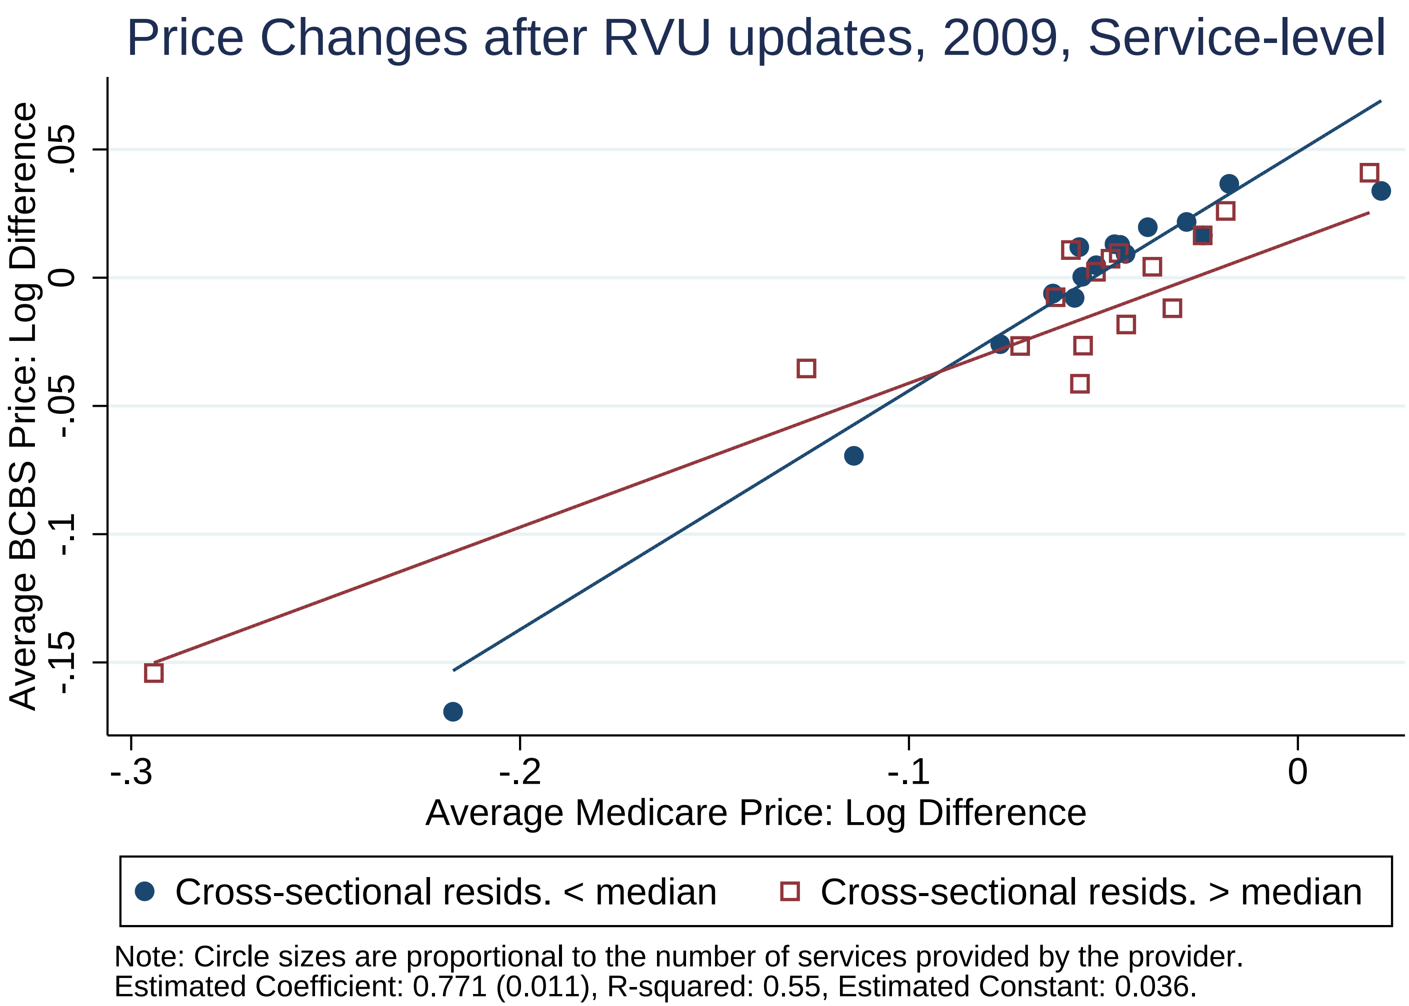

Note: The figure reports the relationship described by equation (C.1) for RVU updates in 2009, split into two sample based on the median prediction error from Figure E.1. (The blue dots in Figure E.1 correspond to the blue circles in this graph, while the yellow and red observations from Figure E.1 correspond to the red squares in this graph.) The regressions are run at the underlying service level, but observations are grouped into twenty bins for this graph, based on vigintiles of the Medicare log RVU change.

Appendix Table E.1: Relationship between the Medicare Benchmarking Estimated in Changes and Observed in the Cross Section

	(1)	(2)	(3)	(4)
	Panel A: 2009 RVU Updates			
Sample (Residual Size):	Small	Medium	Large	All
Log RVU Change	1.173***	0.870***	0.546***	1.085***
× Post-Update	(0.070)	(0.052)	(0.034)	(0.076)
Log RVU Change				-1.192***
× Post-Update × Residual				(0.215)
N	11,444,161	8,319,559	2,177,507	21,941,227
No. of Clusters	268	1,598	1,941	3,807
	Panel B: 2010 RVU Updates			
Sample (Residual Size):	Small	Medium	Large	All
Log RVU Change	0.876***	0.580***	0.464***	0.956***
× Post-Update	(0.020)	(0.058)	(0.107)	(0.061)
Log RVU Change				-1.536***
× Post-Update × Residual				(0.422)
N	11,993,795	9,567,049	2,372,733	23,933,577
No. of Clusters	398	1,347	1,936	3,681
	Panel C: 2011 RVU Updates			
Sample (Residual Size):	Small	Medium	Large	All
Log RVU Change	0.712***	0.657***	0.719***	0.755***
× Post-Update	(0.097)	(0.085)	(0.067)	(0.102)
Log RVU Change				-0.299
× Post-Update × Residual				(0.423)
N	13,059,796	9,817,494	2,526,717	25,404,007
No. of Clusters	385	1,390	2,316	4,091

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. Columns 1 through 3 of the table show the results of OLS specifications of the parameter $\hat{\beta}$ from equation (10) in section 5.2. In column 1, we restrict the sample to the HCPCS codes in the bottom half of the distribution of the average cross-sectional prediction errors generated by estimating equation (C). Column 2 restricts the sample to services falling between the 50th and 90th percentiles of the distribution of prediction errors, while column 3 contains services in the top decile. Column 4 presents estimates of $\hat{\beta}$ and $\hat{\gamma}$ from equation (E.1) in section E.1. Panel A shows estimates using RBRVS updates and BCBS claims data for 2009, Panel B for 2010, and Panel C for 2011. Observations are at the claim-line level and are equally weighted. 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 note to Table 1 and in the main text. Sources: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBS.

F Supply Elasticities Implied from Private Prices that Follow Medicare's

We use the following IV framework to estimate the own-price supply responses for physicians treating BCBS patients, in responses to reimbursement changes that follow from Medicare's RVU updates:

$$\Delta\overline{\ln(P_{g,j})} = \alpha + \beta\Delta\ln(RVU_j) + \varepsilon_{g,j} \tag{F.1}$$

$$\Delta\ln(Q_{g,j}) = \gamma + \delta\widehat{\Delta\overline{\ln(P_{g,j})}} + \epsilon_{g,j}. \tag{F.2}$$

The first stage, equation (F.1), is taken from equation (C.1) in the text. This estimates the share of private prices that respond to the Medicare RVU updates. This generates a predicted price change, which we use in the second stage equation (F.2).

The coefficient δ that we estimate in equation (F.2) is close to providing an estimate of the physicians' supply elasticity for BCBS patients, in response to BCBS prices. It is somewhat confounded, however, by the fact that the BCBS prices are changing at the same time as the prices of physicains' outside option—treating Medicare patients.[26] This would tend to bias the estimates down relative to a pure own-price supply estimate.

Table F.1 shows the results. The IV estimates scale up the reduced form estimates substantially, and range from 0.15 to 0.66. The median estimate of 0.37 occurs in 2011. For comparison, the conceptually most similar estimates in the literature are those of Brekke, Holmås, Monstad and Straume (2015). Brekke et al. (2015) estimate physicians' supply responses to a reimbursement change for one particuar service, which is also the type of price change we consider here. These are different types of elasticities than those of Clemens and Gottlieb (2014), who consider market-wide changes, or the relative price changes of Gruber et al. (1999) and Jacobson et al. (2010).

[26]Clemens and Gottlieb (2013, Appendix B) model these forces.

Appendix Figure F.1: Short-Run Supply Responses to Medicare Price Changes

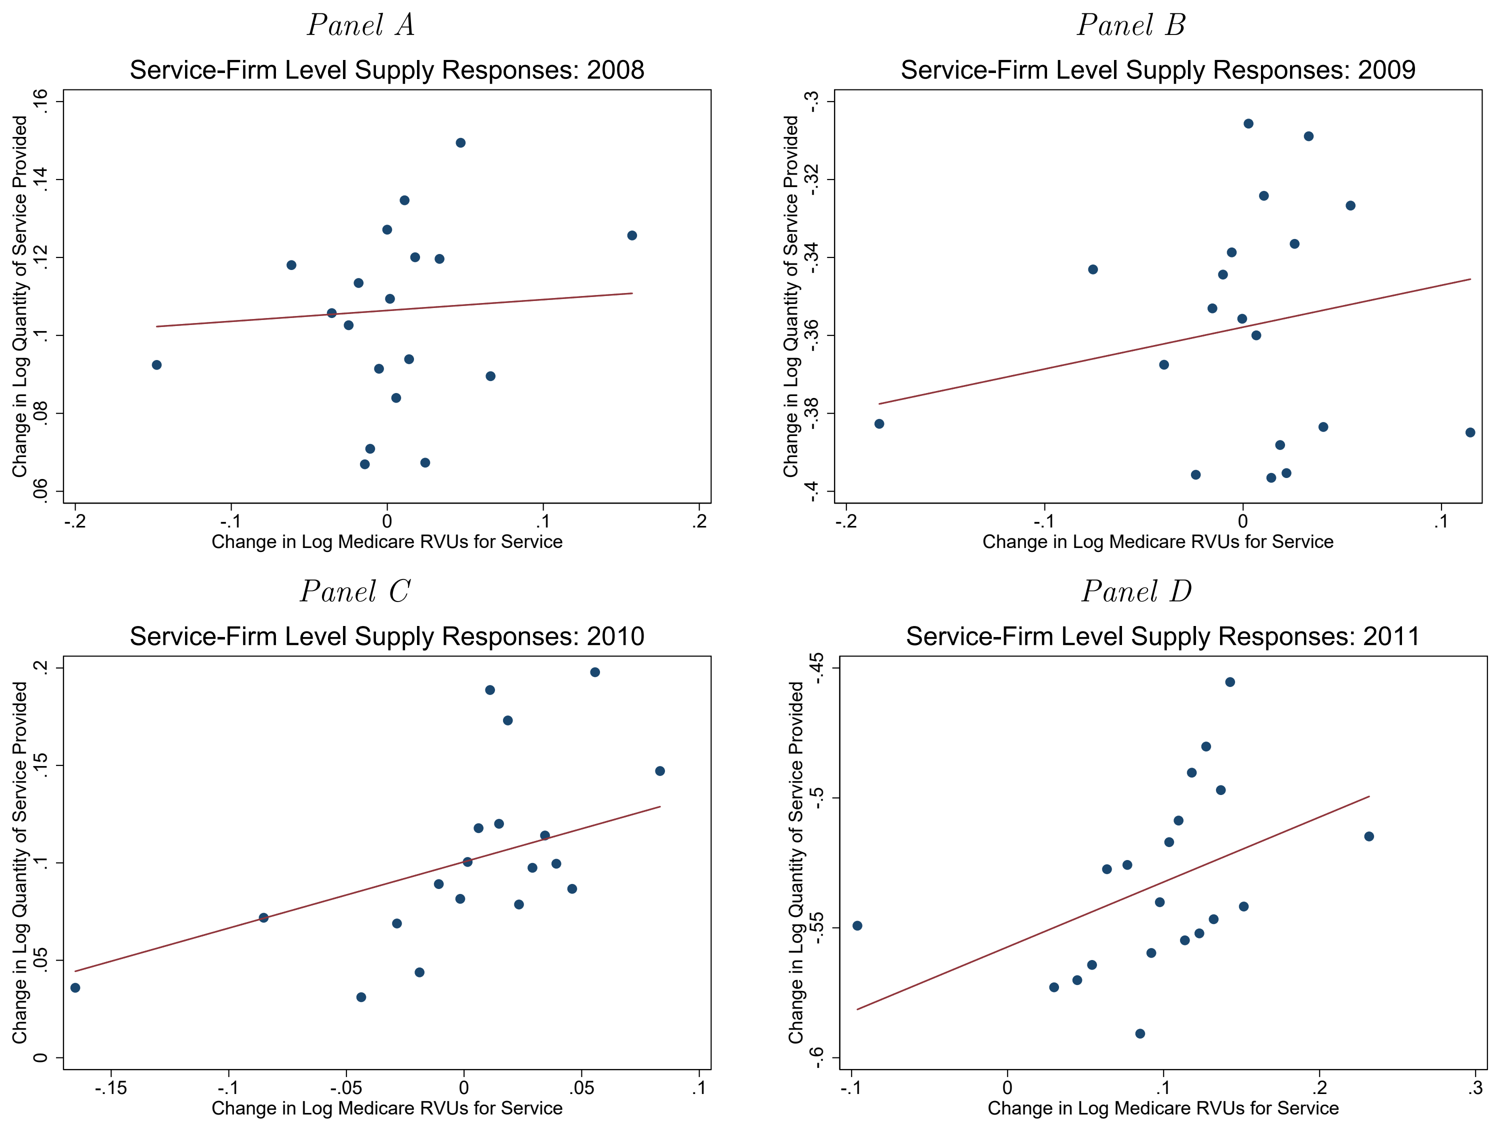

Note: The figure reports estimates of physicians' supply responses to Medicare price changes that BCBS implemented in a given year. Quantities, the dependent variable, are computed at the service-by-firm level. Each panel shows a separate year's estimates, measured as log differences between the period before BCBS implemented the Medicare RVU updates and the period after this update. The years are split at July 1, 2008, August 15, 2009; July 1, 2010; and September 1, 2011. The estimates have very different intercepts across the three panels because of the differences in the share of the year's data that are included in the periods before *vs.* after each year's update.

Appendix Table F.1: **Supply Elasticity Estimates**

	(1)	(2)	(3)	(4)
Year:	2008	2009	2010	2011
Dependent variable:	Change in log service quantity			
	Panel A: Reduced Form			
Log RVU change for service	0.027	0.095*	0.339***	0.252***
	(0.047)	(0.047)	(0.050)	(0.038)
	Panel B: IV Estimates			
Log BCBS payment change for service	0.052	0.152*	0.658***	0.365***
	(0.090)	(0.076)	(0.102)	(0.055)
N	63,526	71,354	81,294	89,936
First Stage F-Statistic	358.9	776.2	483.9	1843.0

Note: **, *, and + indicate statistical significance at the 0.01, 0.05, and 0.10 levels respectively. Panel A estimates the reduced form relationships shown in Figure 5 in the main text. Panel B shows the second stage estimates from the IV framework in equation (F.2). The robust first-stage F-statistics all easily satisfy the weak instruments test of Olea and Pflueger (2013). Source: Authors' calculations using claims data from BCBS.