NBER WORKING PAPER SERIES

THE ANATOMY OF PHYSICIAN PAYMENTS: CONTRACTING SUBJECT TO COMPLEXITY

Jeffrey Clemens
Joshua D. Gottlieb
Tímea Laura Molnár

Working Paper 21642
http://www.nber.org/papers/w21642

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2015

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 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Anatomy of Physician Payments: Contracting Subject to Complexity
Jeffrey Clemens, Joshua D. Gottlieb, and Tímea Laura Molnár
NBER Working Paper No. 21642
October 2015
JEL No. H44,H51,H57,I11,I13,L98

ABSTRACT

The reimbursement rates that private insurers pay to physicians are closely linked to those set by Medicare, despite the well-known limitations of Medicare's fee schedule. 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 an otherwise complex pricing environment. We analyze 71 million claims from a large private insurer, which represent $6.3 billion in spending over three years. Using two empirical approaches, we estimate that 75 percent of services, accounting for 65 percent of dollars, 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 pricing. When the insurer deviates from Medicare's relative prices, these deviations are consistent with adjusting 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.

Jeffrey Clemens
Department of Economics
University of California, San Diego
9500 Gilman Drive #0508
La Jolla, CA 92093
and NBER
jeffclemens@ucsd.edu

Joshua D. Gottlieb
Vancouver School of Economics
University of British Columbia
#997 - 1873 East Mall
Vancouver, BC V6T 1Z1
Canada
and NBER
joshua.gottlieb@ubc.ca

Tímea Laura Molnár
Vancouver School of Economics
University of British Columbia
#997 - 1873 East Mall
Vancouver, BC V6T 1Z1
Canada
timea@nber.org

In an exclusively public health care system, payment rates for 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 upon payments through private negotiations. This paper looks into the black box of the payments embedded in these 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 likely 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 ideal balance between public and private care financing rests 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).

In the U.S. public sector, the federal Medicare program compensates providers for out-

[1]This represents almost half of 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.

patient care 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 determines 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 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 parts. First, we provide two forms of evidence on the pervasiveness of links between Medicare's RBRVS and the payments from a single large insurer to the physicians who treat its beneficiaries. 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 RBRVS. Together, the analysis yields insights into both the extent of Medicare's influence and the economic factors underlying the insurer's approach to contracting.

Our analysis uses insurance claims data from Blue Cross Blue Shield of Texas (BCBSTX). These data have two core features that are important 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 spending on physician and clinical services from private health insurers (CMS, 2011).

We develop two methods to estimate the pervasiveness of payments linked directly to Medicare's RBRVS in the BCBSTX 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). If many services share the same ICF, this can plausibly reflect a contract specifying this particular markup. As a first pass, we infer that every claim whose ICF accounts for at least 10 percent of the provider's BCBSTX payments is contractually linked to Medicare. Under this assumption, around three quarters of BCBSTX's claims, accounting for two thirds of spending, follow the Medicare benchmark.

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

Using both the cross-sectional and update-based frameworks, we examine heterogeneity in the strength of RBRVS benchmarking across physician groups and service categories. Looking across physician groups, we find that payments to relatively large firms are less tightly linked to the RBRVS than payments to small firms. Our estimates suggest that payments for just over 80 percent of services provided by the smallest firms (representing 60 percent of their spending) are linked to Medicare's relative values. The same is true of 65 percent of services from firms with total BCBSTX 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 80 percent of evaluation and management services, but only 50 percent of imaging services, are directly linked to Medicare's menu.

Within diagnostic imaging, the RBRVS distinguishes between a capital-intensive "technical component" for taking the image, and a labor-intensive "professional component" for interpreting the image. The RBRVS explicitly amortizes the fixed cost of the imaging equipment into the technical component. We find that BCBSTX payments for the professional components are far more tightly linked to the RBRVS than are its payments for the technical components.

Finally, we show that BCBSTX's adjustments work to narrow likely gaps between marginal costs and RBRVS'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.

These results suggest that physician contracts are written to manage the tension between gains from fine-tuning payments and costs from making contracts complex. The benefits of fine-tuning payments will tend to be greatest for contracts with large physician groups, providing a rationale for why such contracts deviate more often from RBRVS. Since the RBRVS's average cost approach will have greater difficulty managing the payments for capital- than for labor-intensive services, the benefits of fine-tuning payments will tend to be greater for the former than the latter. BCBSTX appears to draw heavily on RBRVS for the purpose of contract simplification, while strategically adapting its contracts where the value of adaptations is likely to be highest.

Our analysis connects to existing work in health care and two broader literatures. First, 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 surprising responses to public sector payment idiosyncracies in pharmaceutical markets. White (2013) found a sizable positive relationship between Medicare and private hospital pricing, as did Clemens and Gottlieb (forthcoming) in the outpatient context. We make three primary contributions here. In the outpatient setting we study, we show that Medicare exerts influence over nominally independent private insurers through those insurers' adoption of Medicare's payment structure. We find evidence that private insurers deviate from this basic structure when the benefits of doing so are likely to be highest, and we show how they deviate from this benchmark.

Second, we contribute to the literature documenting how boundedly rational agents navigate complex environments. Work on 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 other applications 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. In our setting, this simplification implies that the public sector's payment model influences substantially more care than that which it finances directly.

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 pro-

[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.

vides insight into why this is the case. With some exceptions, Medicare's payment updates occur annually. We find that many private payments reflect Medicare's changes by updating with a similar frequency. Consistent with Anderson, Jaimovich and Simester's (forthcoming) evidence from retail, the complexity of setting physician payments helps to 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 have non-trivial implications for short-run measures of overall price 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 RBRVS. In section 5, we derive the empirical specifications that can estimate the Medicare-benchmarked share of payments using updates to Medicare's relative prices. 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 BCBSTX adjusts its payments when they deviate from the benchmark, and section 8 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. This section presents key features of Medicare's administrative pricing mechanisms and discusses some of the institutional detail surrounding 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 a number of Relative Value Units (RVUs) to each of 10,000 distinct health care services. The RVUs associated with service j are legislatively bound to measure the resources required to provide that service. RBRVS 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 in payment area i, the provider's fee is approximately:

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

The Reimbursement Rate, a term we use interchangeably with "price," is the amount Medicare pays for this service. The Conversion Factor (CF) is a national scaling factor, usually updated annually. The analyses in Clemens and Gottlieb (2014, forthcoming) relied on administrative changes in the CF and GAFs.

Payments across services vary primarily according to their assigned number of Relative Value Units (RVUs). These 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 will exploit these changes in one of our empirical strategies, which we introduce in section 5.

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

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] As discussed below, 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 BCBSTX-specific sources provide institutional detail that illuminates the Medicare fee schedule's role. Newsletters 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 RBRVS unless the insurer's contract specifies otherwise. Our empirical work will examine when and why this occurs. We will measure how often contracts specify exceptions, and whether BCBSTX'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 set a benchmark for an absolute price level. As in Medicare, realized private reimbursements involve the relative value scale multiplied by a constant conversion factor to convert RVUs into dollars. This conversion factor is a key subject of negotiation.

[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).

2 Conceptual Framework

Practitioners describe two modes of negotiation between providers and private insurers. Insurance carriers typically offer small provider groups payment contracts based on a fixed fee schedule. Whether this 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. Our model examines when each bargaining approach would be efficient and what each means for the welfare consequences of Medicare payment reforms.

We sketch a model of physician-insurer reimbursement rate determination that allows relative prices to 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 the substantial negotiation and coordination costs in our setting (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 payment rate from Medicare. For

[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.

[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.

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 $\alpha, \beta > 0$.

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, $u(q_1, q_2) = aq_1 + bq_2$ where q_j is the quantity of care 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 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, namely their reservation value $\overline{u}$. As $\overline{u}$ increases, insurers must increase physician reimbursements in order to induce the increase in care required to satisfy higher-$\overline{u}$ patients. The result also shows that reimbursements are decreasing in our proxies for the responsiveness of health care supply, α and β.

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* $E^* \equiv \dfrac{\overline{u}^2}{\alpha a^2 + \beta a^2}$*. 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 its optimal reimbursement ratio differs from the ratio implied by Medicare benchmarking. 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,* $\left|\dfrac{r_2^M}{r_1^M} - \dfrac{b}{a}\right|$*. 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, our assumptions about patient preferences imply that 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 prices. 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 negotiating service-specific reimbursement rates.

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 (BCBSTX). The claims database we analyze covers the universe of payments made by BCBSTX for outpatient care in 2009, 2010, and 2011. For each claim, the database provides information on the service provided, location, physician, physician group, and BCBSTX's payment to that group. Our analysis sample restricts this universe along several dimensions. The full 2009 dataset contains 54,724,994 claim lines and \$4.01 billion in spending. We clean the data as described in Appendix B.1, which initially leaves us with 41,182,992 service lines and \$2.44 billion of spending.

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 RBRVS 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 BCBSTX frequently base payments on codes outside of RBRVS. We retain over 97 percent of claims for evaluation and management, diagnostic imaging, and surgical services. 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. Other years' figures are slightly larger.

The claims data further allow us to describe the provider groups serving BCBSTX beneficiaries, at least in terms of the care they provide in that context. To enable our subsequent investigation of heterogeneity in Medicare benchmarking, we measure the total value of the care each group provides to BCBSTX patients in a given year. Our final dataset includes care provided by nearly 80,000 physician groups. Table 1 presents summary statistics on the physician groups in our final sample, and Table 2 shows the geographic distribution of the claims.

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 the simplicity of Medicare's pricing scheme to back out Blue Cross's reliance on RBRVS. The idea is straightforward: if many services from a particular physician group appear to follow the same pricing formula—Medicare's Relative Value Units ($RVUs_{j,t}$) times a constant markup—then these payments are likely benchmarked directly to Medicare's relative rates.[9] To see this formally, 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:

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

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

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

This equation defines an "implied conversion factor" (ICF)—the conversion factor that would rationalize a payment of $P_{c,j,t}$ in a Medicare-benchmarked contract. Inspection of equation (2) reveals that this pricing scheme implies a 1 for 1 relationship between log RVUs and the log of $P_{c,j,t}$. Equipped with $P_{c,j,t}$ from the claims data and CMS's published RVU assignments, we can directly investigate the prevalence of common ICFs.

Since an ICF exists for every service with Medicare-defined RVUs, simply computing the ICF does not tell us whether service 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 BCBSTX data

[9]The strength of this claim is driven in part by the precision with which we round the markups. Specifically, we explore markups rounded to the nearest 20 cents, 10 cents, and 2 cents. Markups rounded to the nearest 2 cents per unit are relatively unlikely to coincide by accident.

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

for 2009.[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 ICF.

Panel A shows the data from a mid-sized group for which the relevance of one 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. The remaining services are priced separately. The three dots immediately below the line, and the three immediately above, may be instances of less common ICFs for this group, but a conservative interpretation would view them as deviations from Medicare-benchmarked pricing.

Panel B presents an equivalently constructed graph for a larger group that provides more unique services at more distinct prices. This group again has one particularly common ICF, though there is stronger evidence for the presence of additional ICFs below the line we have drawn in the figure. Finally, Panel C presents payment data for a large group that provides a substantial number of services. This large group has a range of different ICFs, none of which visually dominate the payment picture. The scatterplot indicates the use of a remarkably complicated contract with BCBSTX.

To develop a summary measure of a group's links to RBRVS, we make two approximations. First, we round each value of $ICF_{c,j,t}$ to the nearest dime. We explore sensitivity to altering this allowance for rounding error, which may result from both the rounding of reported payments and from the division in equation (3).

[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 incorporates these infreqeuntly used codes.

Second, we define "common ICFs" (cICFs) for a particular physician group as those that rationalize a substantial share of the group's services. In Figure 1, the red lines in Panels A and B should undoubtedly qualify as cICFs, while other values may also qualify depending on the how strictly we define a cICF. We consider a variety of thresholds for this requirement, ranging from 5 percent of a group's services to 20 percent. We then ask what share of overall BCBSTX payments belong to *any* cICF for the group providing the care.

4.2 Results From Common Implied Conversion Factors

Table 3 presents the share of services linked to Medicare in each year according to the methodology of section 4.1. The results are very similar between 2010 and 2011, but are notably lower in 2009. Depending on the thresholds imposed, we estimate Medicare links from 40 to 80 percent in 2009, and 70 to 90 percent in subsequent years. 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 C.1 shows that the results are stable to a variety of alternative definitions. If we only count the single most common ICF for each group, the annual estimates fall to 35, 64, and 59 percent—suggesting that it is important to allow for multiple ICFs per physician group. Nevertheless, theory does not provide guidance as to which thresholds are most appropriate, and the choice of threshold substantially affects our estimate of the linked share. To overcome this problem, section 5 introduces a separate estimation strategy that is not sensitive to choices of this sort.

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 RBRVS model's more salient problems. RBRVS is 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 size and link share, we categorize physician groups g according to vigintiles of their aggregate Medicare billing in a year, using $\mathbb{1}_{s(g)}$ to denote vigintile fixed effects.

Figure 2A shows the relationship between the share linked to Medicare and vigintiles of group size. In each year, we observe a stark negative relationship, meaning that large groups' services have more deviations from Medicare benchmarking than small groups'. The differences range from 30 to 40 percentage points depending on the year and measure.

Large groups could provide different services than small groups, and these services might have different propensities to be benchmarked to RBRVS. To check whether the group size relationship is affected by each group's composition of services, we run the following regression at the group-code level, separately by year:

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

This regression allows us to to purge our estimates of the group size-Medicare link relationship of bias from the different composition of services offered by different groups. Figure 2B shows the estimates of η_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.

In Panels C and D, we compute the Medicare-linked share in each vigintile in terms of dollars spent, rather than number of services. These graphs are otherwise identical to Panels A and B, respectively. In all four panels, the stark negative relationship remains apparent in all years.

We next turn to differences across Betos categories, which are captured by estimates of ζ_b from equation (4). Column 1 of Table 4 reports estimates of equation (4), from the same regression 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 15–20 percent (in 2009) and 30–50 percent (in 2010 and 2011) 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 4 and Figure 2 thus provide evidence that firm size and service categories independently predict variation in the prevalence of Medicare-benchmarked payments. This initial evidence suggests that Medicare benchmarking is less common when contracts cover substantial quantities of care, and when services are more capital-intensive.

5 Empirical Model

Despite the clarity of the relationships presented in Figure 1, it remains possible that we have misidentified the extent of BCBSTX's explicit reliance on Medicare payment rates. The ICF-based estimates are sensitive to the threshold we set for sufficient commonality, and theory does not provide guidance as to what threshold is most appropriate. To overcome this concern, we develop a new method for identifying Medicare benchmarking, which relies on changes in the relative values that Medicare assigns to individual services. This approach overcomes the ambiguity about what should count as a cICF, but at the expense of obtaining identification off of a limited number of services.

5.1 Changes in Medicare's Relative Values

Our second empirical approach exploits a series of updates to Medicare's RVU weights. Updates are recommended to CMS by a committee of the American Medical Association composed of representatives of various physician specialties. This committee is known as the Relative Value Scale Update Committee (RUC), and CMS implements the vast majority of its recommendations (Government Accountability Office, 2015).

Updates to the RBRVS come in two main forms: reassessments of the resources required to provide a single service, and revisions to part of the underlying methodology. A service's total number of RVUs comprises three parts, which are meant to account for the physician's own Work, the associated Practice Expense, and the associated Malpractice Expense.[12] A

[12] The Practice Expense RVUs are intended to capture the non-physician labor, such as nursing time, and capital inputs required to provide a service. The Malpractice Expense RVUs amortize the annual malpractice insurance premium across the services a physician provides.

revision to the method for computing the Work component, for example, 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.

For Medicare payment purposes, the vast majority of updates to the RBRVS go into effect on January 1 each year. But when relying on RBRVS, 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). BCBSTX 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 August 15, 2009, on July 1, 2010, and on September 1, 2011 (BCBSTX 2009; 2010; 2011). Table 5 summarizes each year's updates. In all three years, the standard deviations of RVU changes are around 7 percent, implying the presence of substantial pricing variation for us to exploit.

Figure 3 shows three examples of how these changes can impact pricing in our BCBSTX data. Panel A shows average log payments by day for the most commonly billed HCPCS 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 BCBSTX implemented the 2010 relative values. Medicare's log RVUs for this service rose by 0.068 between the 2009 and 2010 fee schedules. BCBSTX's average log payment rose by just under 0.05.

Panels B and C show average payments for the two parts of chest x-ray reimbursement, namely the technical component (HCPCS code 71020-TC) and professional component (71020-26), respectively. Because these services occur far less frequently than the office visit, we average their log payments at a monthly rather than daily frequency. On July 1, 2010, the log RVUs assigned to the technical component declined by 0.072, while those for the professional component did not appreciably change. BCBSTX's log payments for the former correspondingly declined by roughly the full 0.072. We now show how we estimate the share of payments linked to Medicare using this type of payment variation.

5.2 Analytical Foundation

Our empirical framework takes advantage of the institutional details we documented in section 1 regarding how Medicare benchmarking works in practice. When payment is linked to Medicare's relative values, it takes the form of a scalar markup over Medicare RVUs. We write this as

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

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 $\ln(\theta_g)$ can be captured by group fixed effects. If it changes over groups and across time, then group-by-time fixed effects serve the same role.

The institutional details, and model from section 2, suggested that payments may alternatively be negotiated without reference to RVUs. In this case, we denote the payment by $P_{g,j,t} = \rho_{g,j,t}$, which implies $\ln(P_{g,j,t}) = \ln(\rho_{g,j,t})$—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 RVUs. Let $\Delta \ln(RVU_{j,t})$ be the difference in log RVUs across the two time periods, and let $\varepsilon_{g,j,t} = \Delta \ln(\rho_{g,j,t})$ be the difference in the possible non-Medicare 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. This is consistent with the updates we saw in Figure 3B. 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 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:

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

which 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 to exclusively 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 section 5.4 below, we discuss the situations in which this challenges our ability to identify the parameter σ. We emphasize there that the plausibility of the relevant assumptions relies on the relatively high frequency at which we are able to estimate payment changes.

5.3 Parametric Event Study

To describe the timing with which BCBSTX incorporates RBRVS 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 BCBSTX 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 RBRVS 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 BCBSTX and firms whose bargaining positions are affected by RBRVS updates.

5.4 Estimation in Changes

In order to see these relationships as transparently as possible, we return to the two-period context and consider estimation in simple differences. 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:

$$\Delta \overline{\ln(P_{g,j})} = \alpha + \beta \Delta \ln(RVU_j) + (1 - \sigma)\overline{\varepsilon_{g,j}}. \tag{12}$$

In the context of price changes for one service, this equation shows how we can directly interpret the evidence from Figure 3A. This graph showed BCBSTX 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.

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.[13] 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 the relationship between private payments and changes in Medicare's relative values 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 average generosity 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.[14] The threat

[13] Specifically, the estimate of equation (12) in this case will yield:

$$\begin{aligned}\hat{\beta} &= \frac{\text{Cov}[\Delta\overline{\ln(P_{g,j})}, \Delta\ln(RVU_j)]}{\text{Var}[\Delta\ln(RVU_j)]} \\ &= \sigma + \sigma\frac{\text{Cov}[\Delta\overline{\ln(\theta_g)}, \Delta\ln(RVU_j)]}{\text{Var}[\Delta\ln(RVU_j)]} + (1-\sigma)\frac{\text{Cov}[\Delta\overline{\ln(\rho_{g,j})}, \Delta\ln(RVU_j)]}{\text{Var}[\Delta\ln(RVU_j)]}. \end{aligned} \tag{13}$$

The full derivation of equation (13) is in Appendix B.2.

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

to our estimation takes the following form: BCBSTX 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 BCBSTX, 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.

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 (13) 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 RBRVS are uncorrelated with the $\sigma_{j,g}$. In section 6.2, we will allow for heterogeneity across various dimensions in the linked shares.

6 Results from RVU Update Analysis

We begin with a cross-service analysis of RVU changes based on equation (12). Figure 4 shows a binned scatterplot of changes in average BCBSTX payment rates across services against Medicare RVU changes, together with the regression estimate of equation (12). This graph uses BCBSTX payment data from 2009, a year when BCBSTX announced it would incorporate Medicare's RVU updates on August 15. We compute changes between the average reimbursements paid for each service preceding and following that date. We estimate

$\hat{\beta} = 0.77$, which implies that 77 percent of services experiencing updates in 2009 have Blue Cross payments linked to RBRVS.

We next examine the dynamics of Blue Cross's payments to determine if the private payments incorporate RVU updates when we would expect. To this end, Figure 5 presents event study estimates of the strength of the link between Medicare's relative value scale and BCBSTX reimbursements. Panel A presents estimates of equation (11) for the RVU changes implemented in 2009, Panel B for those implemented in 2010 and Panel C for 2011. BCBSTX's provider newsletters say that updates to Medicare's RVUs went into effect on August 15, 2009, July 1, 2010, and September 1, 2011.

The estimates again reveal substantial links between RVU updates and the payments providers receive from BCBSTX. Interpreted as estimates of σ, the results for 2009, 2010, and 2011 imply that 78, 75, and 70 percent of services were linked to Medicare's relative values. The dynamics displayed in the figure are consistent with the view that this link involves the manner in which Medicare's relative values are embedded in BCBSTX's contracts. As in the raw data for standard office visits presented in Figure 3, we see that payment changes occur in each year in precisely the period during which BCBSTX implemented these updates.[15] 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 the predictors and determinants of 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 predicted payment.

Table 6 presents our baseline estimates of equation (10), which summarizes our estimates of $\hat{\beta}$ for each year in a single coefficient. It further probes the robustness of these estimates

[15] 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 (BCBSTX 2009).

to a variety of specification checks. Panel A presents the estimates for 2009, Panel B for 2010, and Panel C for 2011. 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 for allows for the cubic trend in day to differ between the periods preceding and following the RBRVS 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 BCBS's physician claims are linked to Medicare's relative value scale.

Table 7 reports an equivalent set of specifications in which each service code is weighted according to the average BCBSTX payment prior to the updates. On average, the estimates imply that roughly two-thirds of BCBS's physician spending is linked to Medicare's relative value scale. The (modest) difference in coefficients between Tables 6 and 7 implies that payments for relatively expensive services are less likely to be benchmarked to Medicare than are payments for low-cost services.

6.1 Checks for the Relevance of Active Contract Renegotiation

The estimates presented in Figure 5 and Table 6 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 deteriorate. 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 8 presents these results. Column 1 reports our baseline specification, unchanged from Table 6. 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 the RBRVS 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 BSBCTX's payments and Medicare's RBRVS 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's RBRVS.

6.2 Deviations from Benchmarking Across Service Categories

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 BCBSTX's physician contracts.

Table 9 estimates equation (10)—the relationship between private prices and changes in Medicare's relative values—separately across broad categories of services. The estimates for each category exhibit non-trivial variation from year to year. This likely reflects the fact that different sets of services experience significant RVU changes in any given year. Once we narrow our focus to particular categories, the variation across years in which services undergo RVU changes—and hence drive our identification—becomes more significant. While this recommends cautious interpretation, we nonetheless observe some stable patterns in Table 9. First, we consistently observe a stronger relationship between private payments and RBRVS updates for Evaluation & Management services than for Imaging. While the estimated difference is modest in the 2011 data, it is substantial in both 2009 and 2010. Averaging across all three years, the estimates imply that 20 percent more of the payments for Evaluation & Management services are linked directly to Medicare's relative values than for Imaging services. The results across Table 9 are consistent with the ICF-based measures from Table 4.

Second, we divide Imaging codes into subcomponents with high capital and high labor content. Providers generally 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 Medicare amortizes 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 results in the bottom portions of each panel reveal that payments for the Professional Component are consistently more tightly linked to Medicare's relative values than are the payments for the Technical Component. This difference is quite large in the 2010 data and relatively modest in the 2009 and 2011 data. Both of these patterns are consistent with 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. Inefficiencies in Medicare's relative values matter more when an insurer contracts with a large provider group than with a smaller one; the payment model for contracts with large groups govern the incentives guiding large quantities of care provision. In Table 10 we thus allow our estimates of the strength of public-private payment benchmarking to vary with size of the group's business.

The first column of Table 10 reports the baseline, service-weighted regression from Table 6. 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 BCBSTX, and large firms as those with more than $1,000,000 in annual billings. Although the coefficients on the firm size interactions vary across years, on average the estimates imply that just over 80 percent of services provided by firms with less than $200,000 in billings are benchmarked to Medicare, while roughly two-thirds 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 payments to firms with billings less than $200,000 are benchmarked to Medicare, while closer to half of payments to firms with more than $1,000,000 in billings are benchmarked. As with the estimates of heterogeneity across services, the heterogeneity by firm size is thus quite consistent between the ICF and RVU-update methods.

7 How Do Private Payments Deviate from RBRVS?

Thus far we have explored the *pervasiveness* of deviations from strictly RBRVS-linked contracts. In both the RVU-update and Implicit 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 BCBSTX's adjustments when it deviates from strictly RBRVS-linked contracts. That is, we investigate what services BCBSTX rewards through upward adjustments and discourages through downward adjustments.

To measure these adjustments, we begin by estimating the following equation on samples restricted to the pre-RVU-update period of each year—*i.e.* the initial months over which

Medicare's relative values are constant:

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

If all payments were mechanically linked to Medicare's relative values, with a uniform contract for each group and no payment reporting error, the above equation would perfectly predict private payments. The group-specific δ_g would account for heterogeneity in groups' markups over Medicare, and we would expect to estimate $\hat{\psi} = 1$. 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.

Figure 6 illustrates this relationship, though disregarding the group-specific markups. Using 2009 data, it shows the average BCBSTX payment for service j against the corresponding Medicare reimbursement. The regression line is an estimate of (14), and the coefficient $\hat{\psi}$ is quite close to 1. Appendix C shows that Medicare links based on these deviations are consistent with our earlier estimates based on RVU updates.

To understand which services tend to receive higher or lower payments than Medicare-benchmarking predicts we examine the cross-sectional residuals from equation (14). We average these residuals by Betos category. Table 11 presents the resulting means, which we construct as $\overline{\hat{e}_{g,j}} = \frac{1}{N_b} \sum_{j \in b} \hat{e}_{g,j}$ for each Betos group b, comprising N_b claims. 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 7 plots the cumulative distributions of these residuals by Betos category. The distributions for Imaging and Procedures show far more density of negative residuals than those for Evaluation & Management services. Testing has more positive residuals, most notably in 2009, although that is largely driven by one outlier.[16] Compared to the relative payments implied by Medicare's relative values, BCBSTX thus adjusts its contracts to favor Evaluation & Management services. This coincides with the conventional wisdom that Medicare's relative values "underpay" for these labor intensive services relative to other services.

Differences in BCBSTX's adjustments for labor- and capital-intensive services are particularly sharp across the sub-categories of diagnostic imaging. Payment adjustments for the labor-intensive Professional Component of these services are substantially positive, averaging 7 log points across 2009, 2010, and 2011. Payment adjustments for the capital-intensive Technical Component of these services are substantially negative, averaging -10 log points across 2009, 2010, and 2011. The bottom panels of Figure 7 show that this pattern holds throughout the distribution. While it is clear that BCBSTX reimbursements lean heavily on Medicare's relative values for their basic payments structure, these results provide evidence that BCBSTX adjusts its contracts to increase the generosity of payments for labor-intensive services and decrease its payments for capital-intensive services.

8 Conclusion

This paper uses the setting of physician payments from a large private insurer, Blue Cross Blue Shield of Texas (BCBSTX), as a window into how private firms contract for services in complex environments. Using two empirical strategies, we show that BCBSTX uses Medicare's Relative Value Resource Based Scale (RBRVS) to significantly simplify this problem. We estimate that roughly three quarters of services and two thirds of BCBSTX's payments are directly linked to the RBRVS.

Despite Medicare's prominent role, the one quarter of services and one third of payments that deviate from RBRVS appear to represent an effort to improve the structure of payments.

[16]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.

BCBSTX disproportionately deviates from RBRVS when the value of doing so appears likely to be highest. Deviations occur disproportionately in contracts with large physician groups, where significant mutual gains can be on the line. By extension, BCBSTX significantly reduces its payments for diagnostic imaging services, a category of care for which many academics and policy makers believe Medicare pays well above marginal benefit (Winter and Ray, 2008; MedPAC, 2011). BCBSTX hews closely to the RBRVS in payments for services where average-cost reimbursements will be most closely aligned with marginal costs, namely labor-intensive primary care services. Further, the direction of BCBSTX's payment adjustments would tend to to encourage the provision of primary care and discourage categories of care for which over-utilization is a more widespread concern.

The use of RBRVS as a pricing backstop implies that many inefficiencies in Medicare's reimbursements spill over into private payment models. By extension, the value of improvements to the RBRVS may ripple through private contracts in addition to improving the performance of Medicare itself. At the same time, we find that BCBSTX adjusts its payments to curb what policy analysts regard as the RBRVS's greatest inefficiencies. Both public and private players thus appear to have important roles in the process of payment system 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 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).

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, "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.

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).

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

Panel A

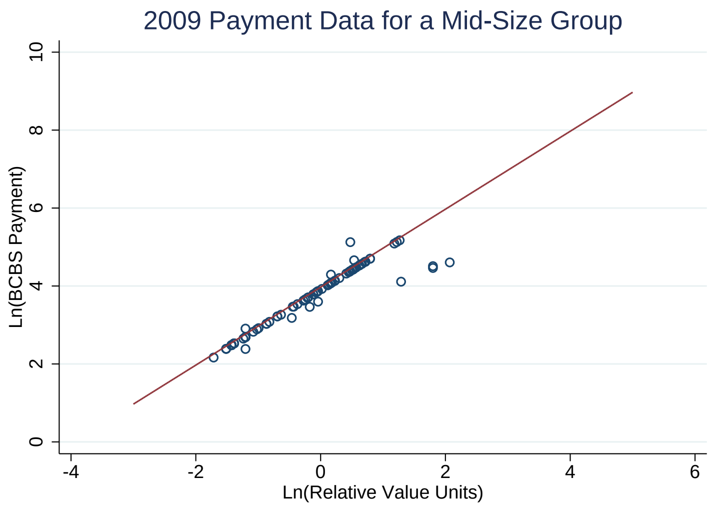

Panel B

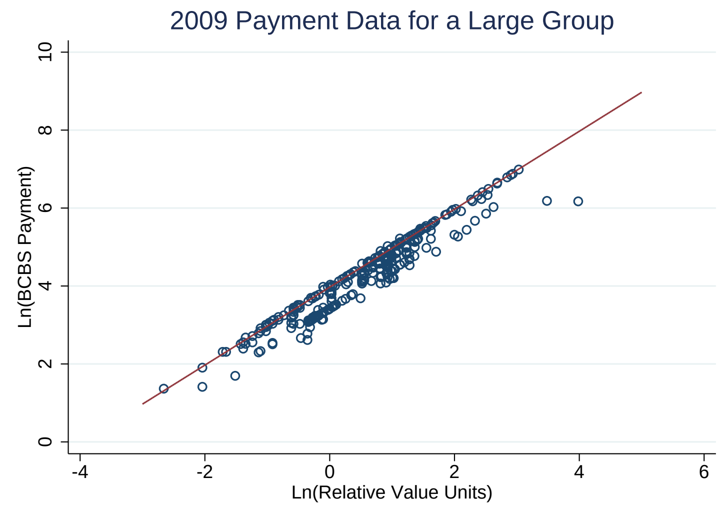

Panel C

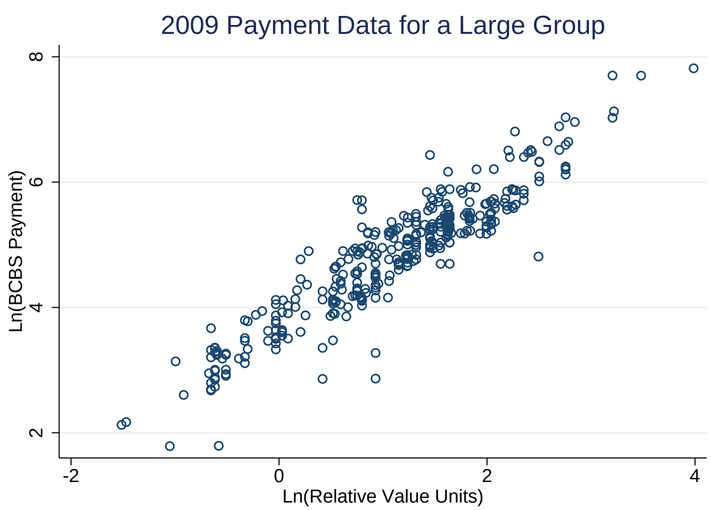

Note: The figure presents the raw data on BSBCTX reimbursement rates, and associated Medicare RVUs, for a sample of different physician groups in 2009. Each observation is a unique reimbursement paid for a particular service to the group.

Figure 2

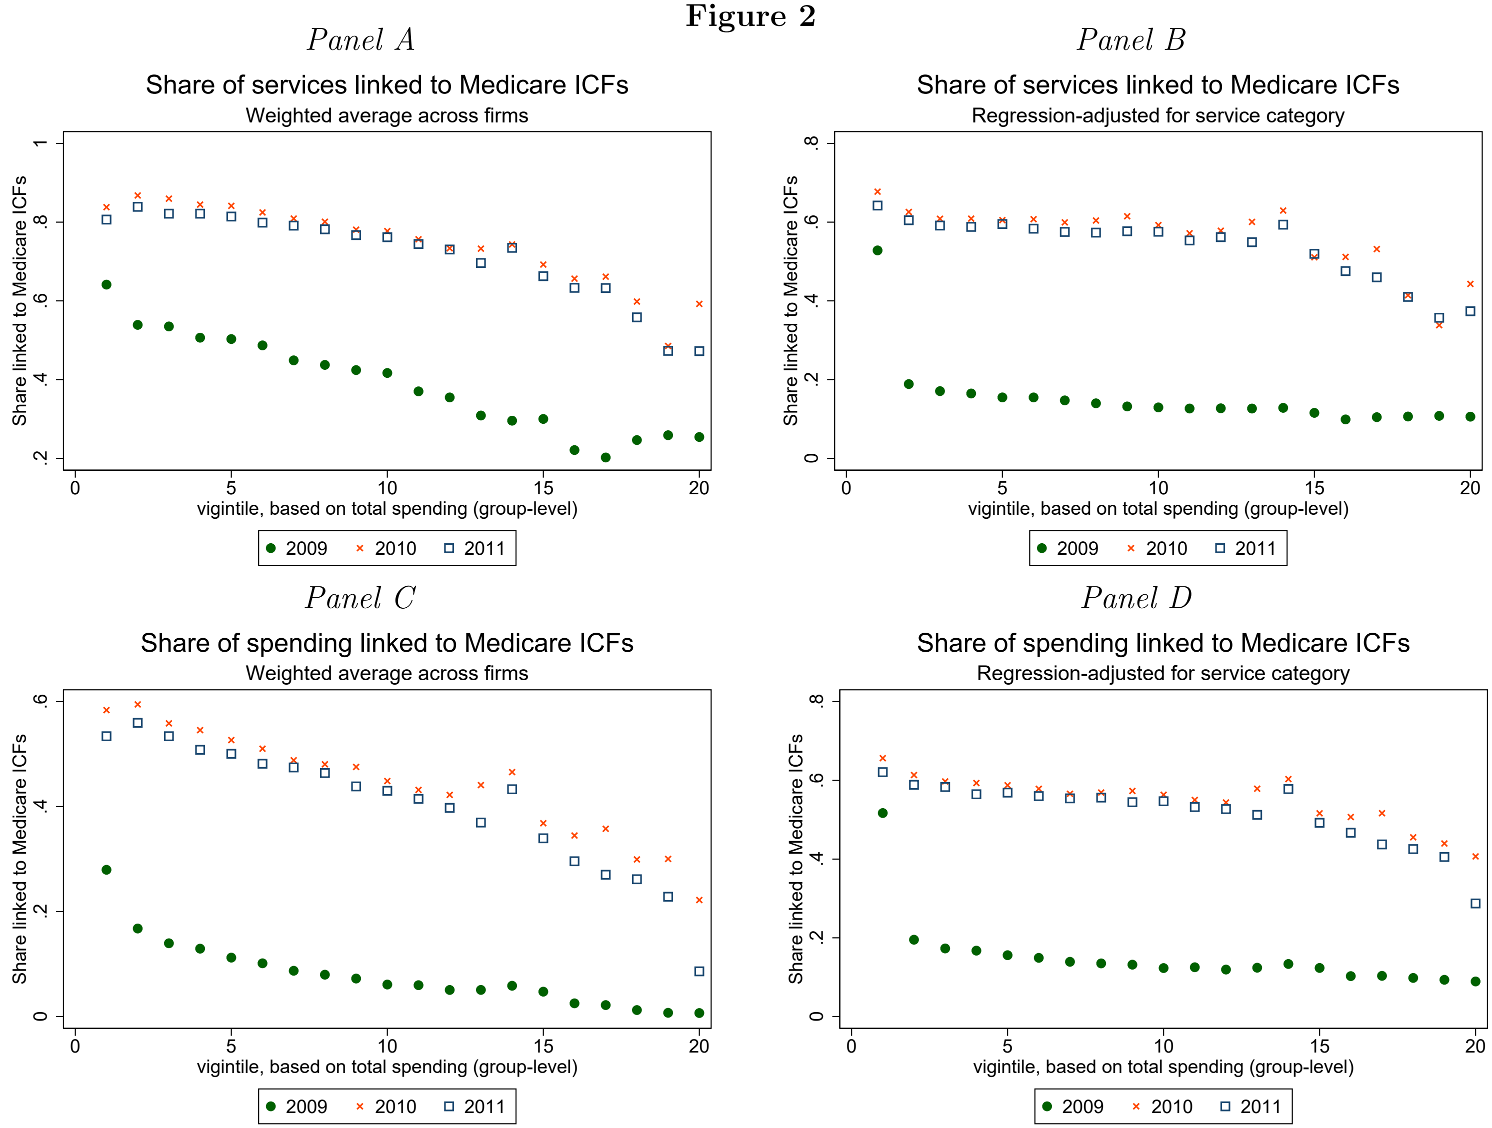

Note: Panel A shows the share of services priced according to common Implicit Conversion Factors (cICFs), as defined in section 4.1, against the amount of BCBSTX spending on care provided by the physician group (grouped into 20 vigintiles). We interpret this as measuring the relationship between a group's Medicare-linked service share and group size. Panel B shows estimates of ζ_b from equation (4), which portray the same relationship, but adjusted for the composition of each group's services. Panels C and D are analogous to Panels A and B, respectively, but measure the Medicare-linked share of spending in dollar terms as opposed to the share of services.

Figure 3

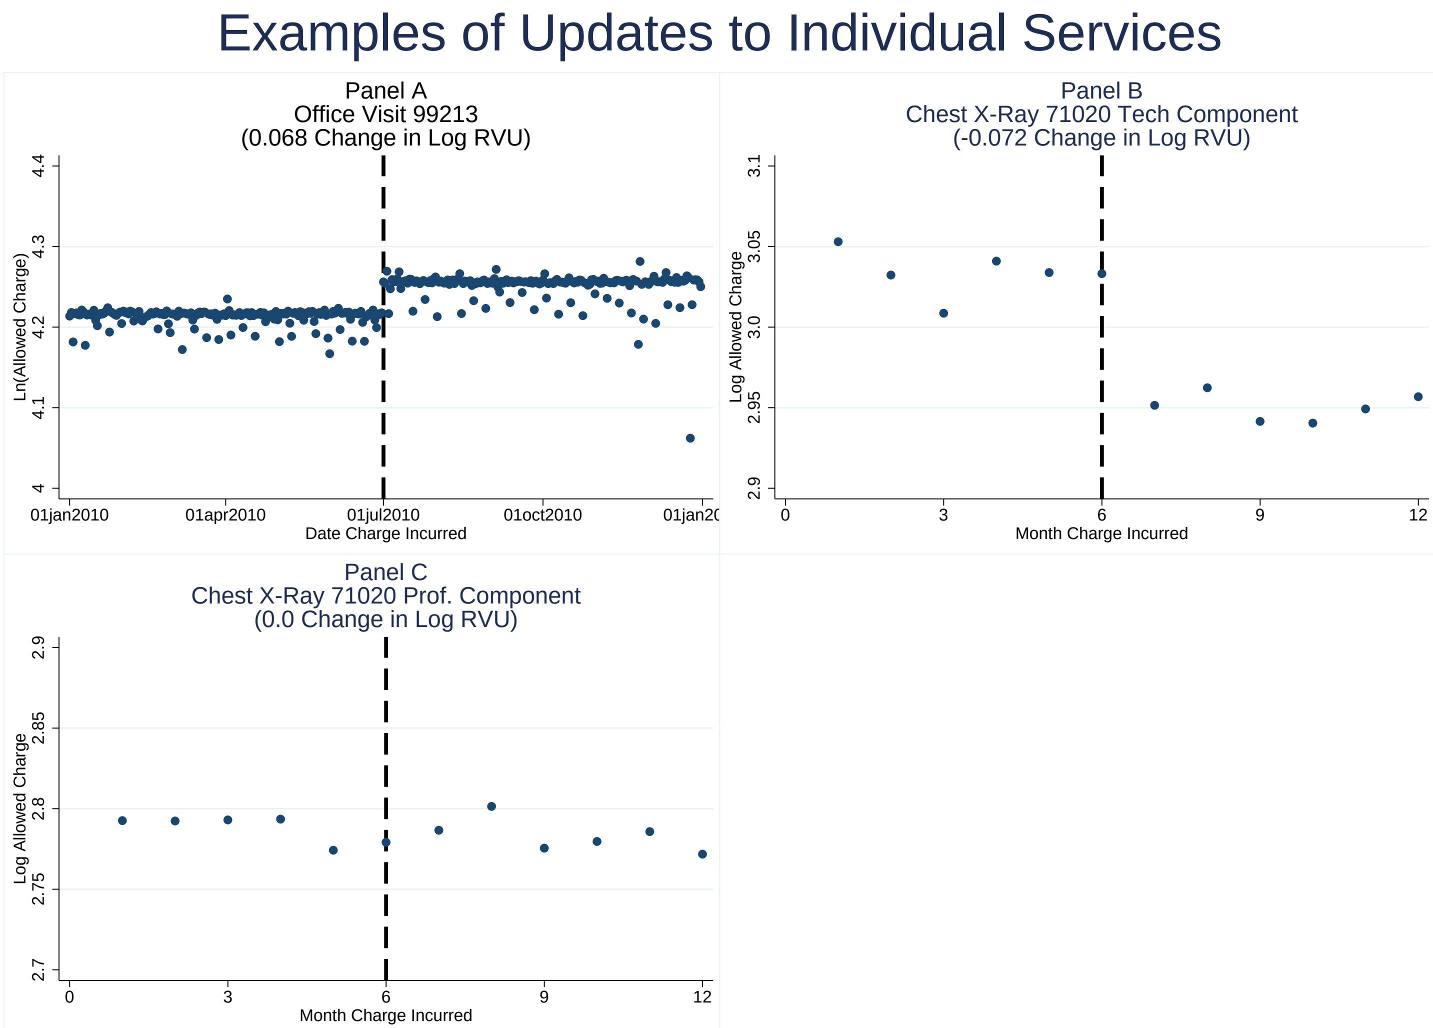

Note: The figure presents daily (Panel A) or monthly (Panels B and C) averages of BCBSTX's log payment for the service named in each panel's title. All data are from calendar year 2010. BCBSTX implemented its update from the 2009 to 2010 relative value scales on July 1, 2010, as indicated by the vertical dashed line each panel.

Figure 4

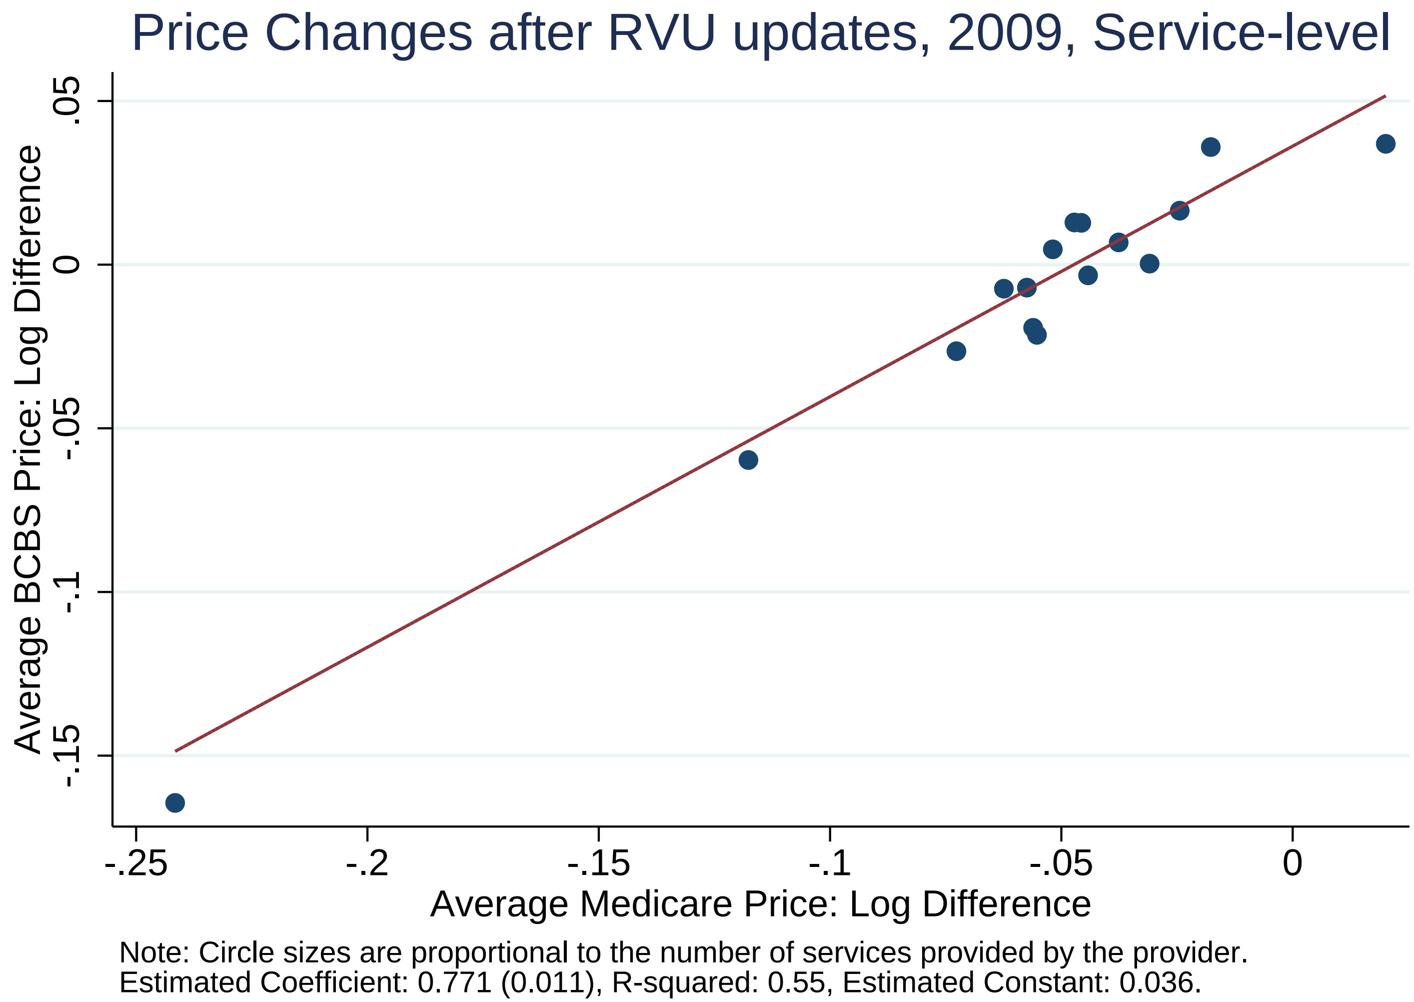

Note: The figure reports the relationship described by equation (12) for RVU updates in 2009, and an estimate of that equation. The regression is run at the underlying service level, but observations are grouped into twenty bins for this graph, based on vigintiles of the Medicare log RVU change.

Figure 5: Strength of Public Private Payment Relationships

Panel A

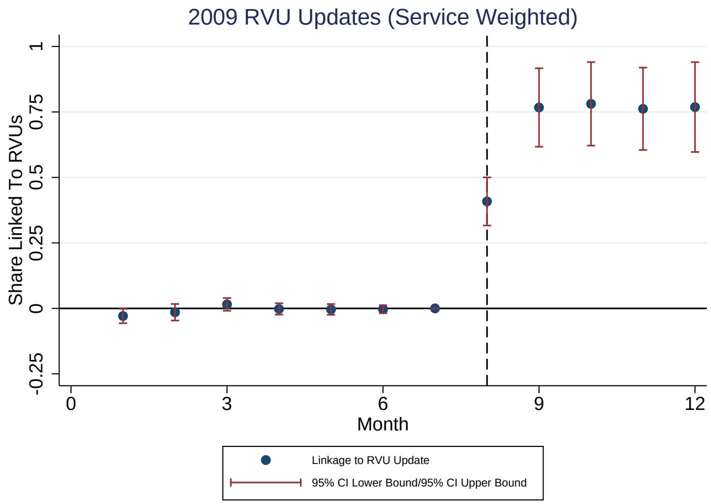

Panel B

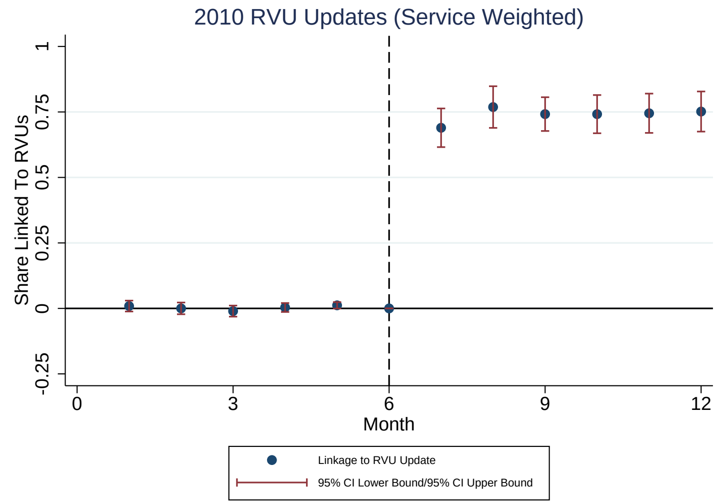

Panel C

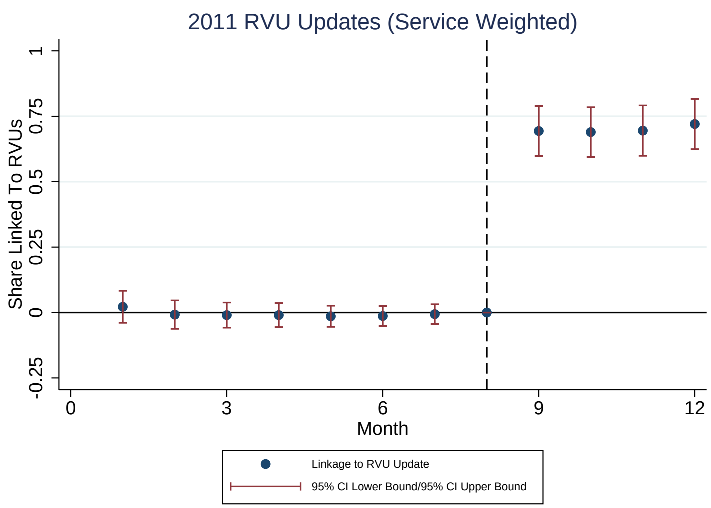

Note: The figure reports estimates of the β_p from estimates of equation (11). The vertical dashed line in each panel corresponds with the month during each year in which BCBSTX implemented its update from the prior year's relative value scale. These updates occurred on August 15, 2009, July 1, 2010,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

Figure 6

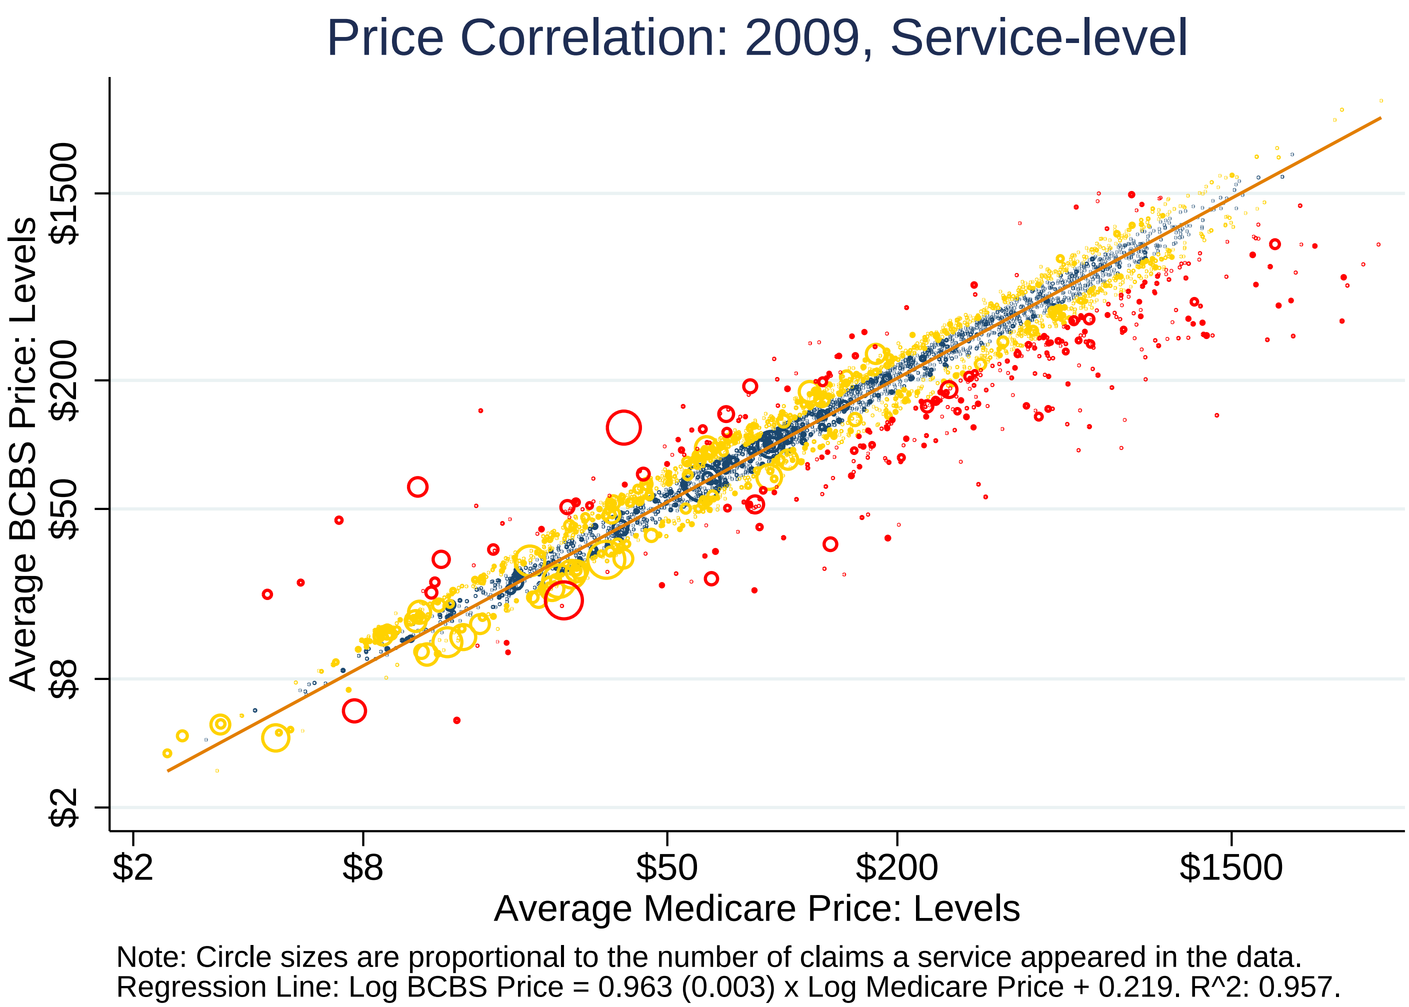

Note: The figure presents the cross-sectional correlation between Medicare and BSBCTX reimbursement rates in 2009. Medicare reimbursement rates are calculated using each HCPCS code's 2009 allocation of relative value units, multiplied by the 2009 national conversion factor. BCBSTX payments are calculated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

Figure 7

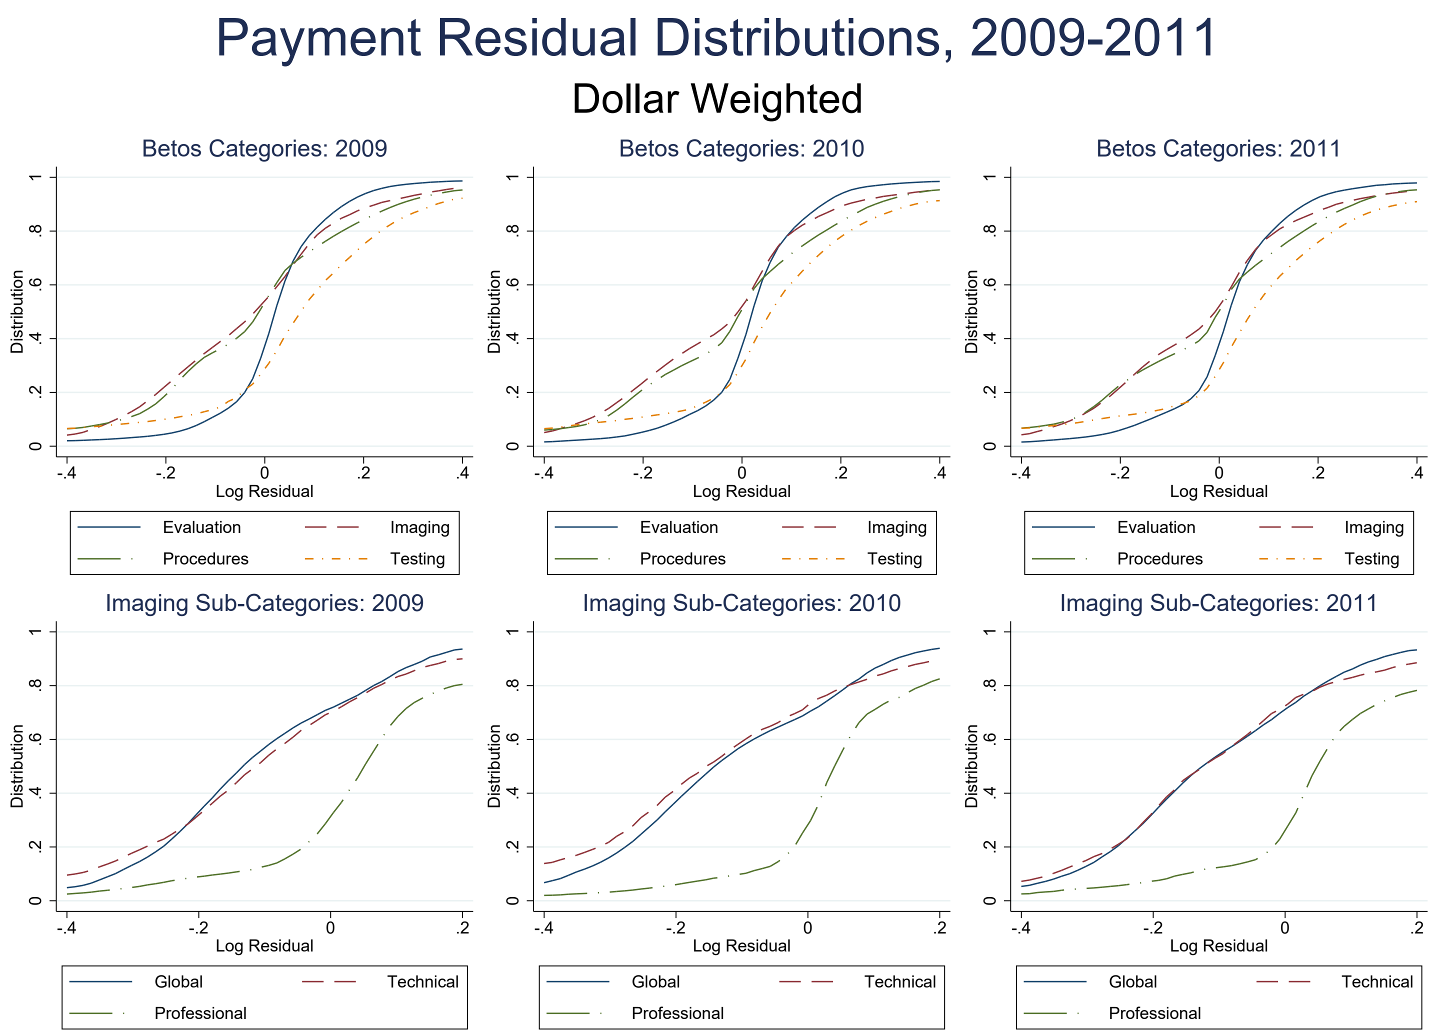

Note: The figure presents residuals $\epsilon_{g,j}$ from estimates of equation (14). For each year, the distribution of residuals is shown within either broad Betos categories (Panels A, B, and C), or within the subcategories of Imaging (Panels D, E, and F).

Table 1: Summary Statistics by Physician Group

	Panel A: All Groups (N=78,296)				
	Mean	**Median**	**Std. Dev.**	**Min.**	**Max.**
Number of unique services	12.84	3	37.37	1	~1,600
Number of patients	100.9	2	2,142.24	1	~375,000
Number of doctors	1.67	1	6.53	1	~1,000
Number of claims	231.17	3	4,240.60	1	~700,000
Mean allowed amount	89.77	69.63	77.33	~5	~620
Total BCBS revenues	23,838	351	294,971	~5	~40,000,000
	Panel B: Groups with Billings $> \$10,000$ (N=14,790)				
	Mean	**Median**	**Std. Dev.**	**Min.**	**Max.**
Number of unique services	49.56	32	73.71	1	~1,600
Number of patients	503.09	156	4,858.81	1	~375,000
Number of doctors	3.86	2	14.61	1	~1,000
Number of claims	1158	392	9,603.73	4	~700,000
Mean allowed amount	89.12	68.51	75.60	~5	~620
Total BCBSTX revenues	119,900	39,500	663,000	10,000	~40,000,000

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

Table 2: Geographic Distribution of Claims

	Claim Shares by Geographic Region, 2009 (%)	
Region	Share of Providers	Share of BCBS Spending
Dallas	27.776	21.694
Houston	12.909	15.308
Austin	9.275	8.050
San Antonio	5.371	6.426
Forth Worth	4.477	5.486
El Paso	0.768	0.892
Rest of Texas	39.424	42.144

Note: Table shows the geographic distribution of our BCBSTX claims data from 2009. Source: Authors' calculations using claims data from BCBSTX.

Table 3: Services Priced According to Common Implied Conversion Factors

	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. 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 BCBSTX.

Table 4: Medicare Benchmarking by Betos Category

	(1)	(2)	(3)	(4)
Dependent variable:	Payments with Common Conversion Factors			
	Service Share		Spending Share	
	Panel A: 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 B: 2010 (N=542,207)			
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)
	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 ζ_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 Implicit Conversion Factors (cICFs) defined in section 4.1. Columns 1 and 3 show estimates after controlling for vigintile of group size, as measured with BCBSTX 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 BCBSTX.

Table 5: Summary Statistics for RVU Changes

Variable	Mean	Std. Dev.	Min.	Max.	N
Change in Log RVUs, 08-09	-0.001	0.074	-0.691	0.738	8,512
Change in Log RVUs, 09-10	0.008	0.073	-0.691	2.45	8,624
Change in Log RVUs, 10-11	0.102	0.067	-1.208	1.211	8,659

Note: This table shows summary statistics on the distribution of Medicare RVU changes across the indicated pairs of years. Source: Authors' calculations using RVU files from CMS.

Table 6: Estimating Medicare Benchmarking Using RVU Changes

	(1)	(2)	(3)	(4)
	Panel A: 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 B: All Services: 2010 RVU Updates			
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 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 BCBSTX 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 BCBSTX.

Table 7: **Dollar-Weighted Estimates of Benchmarking Using RVU Changes**

	(1)	(2)	(3)	(4)
	Panel A: 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 B: All Services: 2010 RVU Updates			
Log RVU Change × Post	0.539**	0.544**	0.568	0.538**
	(0.061)	(0.061)	(.)	(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
	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 x RVU Change	No	No	Yes	No
Cubic Time x 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 BCBSTX claims data for 2009, Panel B for 2010, 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 BCBSTX.

Table 8: Checks for the Relevance of Active Contract Negotiations

	(1)	(2)	(3)	(4)
	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 6. 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 BCBSTX 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 BCBSTX.

Table 9: Public-Private Payment Links Across Service Categories

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Panel A: 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 B: 2010 RVU Updates by Betos Categories						
	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
	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 BCBSTX claims data for 2009, Panel B for 2010, 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 BCBSTX.

Table 10: Medicare Benchmarking by Firm Size

	(1)	(2)	(3)	(4)
Panel A: 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 B: 2010 RVU Updates (N = 23,933,577)				
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)
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 Tables 7 and 6 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. Sources: Authors' calculations using updates to Medicare's RBRVS as reported in the Federal Register and claims data from BCBSTX.

Table 11: In What Direction Does BCBSTX Adjust Its Payments for the Various Service Categories?

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Panel A: 2009 Payment Residuals by Betos Categories						
	Evaluation & Management	Imaging	Procedures	Tests	Imaging Sub-Categories:		
					Global	Technical	Professional
Residual Mean	0.0213	-0.0374	-0.0668	0.0835	-0.115	-0.0930	0.0609
Residual SD	(0.194)	(0.272)	(0.253)	(0.345)	(0.279)	(0.284)	(0.226)
N	6,992,653	2,179,969	2,371,468	897,896	1,099,857	135,847	944,265
	Panel B: 2010 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
	Panel C: 2011 Payment Residuals by Betos Categories						
	Evaluation & Management	Imaging	Procedures	Tests	Imaging Sub-Categories:		
					Global	Technical	Professional
Residual Mean	0.0233	-0.0269	-0.0296	0.0789	-0.102	-0.0924	0.0740
Residual SD	(0.204)	(0.254)	(0.319)	(0.370)	(0.245)	(0.281)	(0.225)
N	8,779,036	2,465,292	3,455,837	1,118,369	1,292,210	129,739	1,043,343

Note: The table presents means and standard deviations of residuals from estimates of equation (14). That is, in each year we regress the log of BCBSTX'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 of each year so that the relative value units are constant for each service throughout the sample.

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

$$u(q_1, q_2) = u\left(s_1(\varphi r_1^M), s_2(\varphi r_2^M)\right) = \varphi\left(\alpha a r_1^M + \beta b r_2^M\right). \tag{15}$$

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

$$\hat{E} = s_1\left(\varphi r_1^M\right)\varphi r_1^M + s_2\left(\varphi r_2^M\right)\varphi r_2^M = \alpha\varphi^2(r_1^M)^2 + \beta\varphi^2(r_2^M)^2. \tag{16}$$

□

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

$$\min s_1(r_1)r_1 + s_2(r_2) \quad \text{subject to} \quad u(s_1(r_1), s_2(r_2)) \geq \overline{u}. \tag{17}$$

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

$$\mathscr{L}(r_1, r_2) = \alpha r_1^2 + \beta r_2^2 - \lambda(\alpha a r_1 + \beta b r_2 - \overline{u}) \tag{18}$$

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

$$r_1^* = \frac{\lambda a}{2} \tag{19}$$

$$r_2^* = \frac{\lambda b}{2} \tag{20}$$

$$\overline{u} = \alpha a r_1^* + \beta b r_2^* \tag{21}$$

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

$$E^* = \frac{\overline{u}^2}{\alpha a^2 + \beta b^2}. \tag{22}$$

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

$$\varphi = \frac{\overline{u}}{(\alpha a + \beta b\omega) r_1^M} \tag{23}$$

and the expenditures in that case as

$$\begin{aligned} \hat{E} &= \left(\alpha + \beta\omega^2\right) \varphi^2 (r_1^M)^2 \\ &= \frac{\overline{u}^2 \left(\alpha + \beta\omega^2\right)}{(\alpha a + \beta b\omega)^2} \end{aligned} \tag{24}$$

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

$$\psi = \frac{\hat{E}}{E^*} = \frac{\left(\alpha + \beta\omega^2\right)\left(\alpha a^2 + \beta b^2\right)}{(\alpha a + \beta b\omega)^2}. \tag{25}$$

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:

$$\begin{aligned} \frac{d\psi}{d\omega} &= \frac{2\beta\omega\left(\alpha a + \beta b\omega\right)^2\left(\alpha a^2 + \beta b^2\right) - 2\beta b\left(\alpha + \beta\omega^2\right)\left(\alpha a^2 + \beta b^2\right)\left(\alpha a + \beta b\omega\right)}{(\alpha a + \beta b\omega)^4} \\ &= (\omega a - b)\frac{2\alpha\beta\left(\alpha a^2 + \beta b^2\right)}{(\alpha a + \beta b\omega)^3}. \end{aligned} \tag{26}$$

All of the terms in the fraction at the end of equation (26) 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 (24), and those when unconstrained are given by equation (22). The difference between these values is

$$\begin{aligned} \xi &= \frac{\overline{u}^2\left(\alpha + \beta\omega^2\right)}{(\alpha a + \beta b\omega)^2} - \frac{\overline{u}^2}{\alpha a^2 + \beta b^2} \\ &= \overline{u}^2\frac{\left(\alpha + \beta\omega^2\right)\left(\alpha a^2 + \beta b^2\right) - (\alpha a + \beta b\omega)^2}{\left(\alpha a^2 + \beta b^2\right)(\alpha a + \beta b\omega)^2} \\ &= \overline{u}^2\alpha\beta\frac{a^2\omega^2 + b^2 - 2ab\omega}{\left(\alpha a^2 + \beta b^2\right)(\alpha a + \beta b\omega)^2}. \end{aligned} \tag{27}$$

Note that equation (27) 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:

$$\frac{d\xi}{d\overline{u}} = 2\overline{u}\alpha\beta\frac{a^2\omega^2 + b^2 - 2ab\omega}{(\alpha a^2 + \beta b^2)(\alpha a + \beta b\omega)^2} > 0 \tag{28}$$

$$\begin{aligned}\frac{d\xi}{d\omega} &= \overline{u}^2\alpha\beta\frac{(2a^2\omega - 2ab)(\alpha a + \beta b\omega)^2 - 2\beta b(a^2\omega^2 + b^2 - 2ab\omega)(\alpha a + \beta b\omega)}{(\alpha a^2 + \beta b^2)(\alpha a + \beta b\omega)^4} \\ &= 2\overline{u}^2\alpha\beta\frac{a(a\omega - b)(\alpha a + \beta b\omega) - \beta b(a^2\omega^2 + b^2 - 2ab\omega)}{(\alpha a^2 + \beta b^2)(\alpha a + \beta b\omega)^3} \\ &= 2\overline{u}^2\alpha\beta\frac{(\alpha a^2 + \beta b^2)(\omega a - b)}{(\alpha a^2 + \beta b^2)(\alpha a + \beta b\omega)^3} \\ &= (\omega a - b)\frac{2\overline{u}^2\alpha\beta}{(\alpha a + \beta b\omega)^3}.\end{aligned} \tag{29}$$

Inequality (28) 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 (29) 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. □

B Data Appendix

B.1 Data Cleaning

The full 2009 dataset contains 54,724,994 claim lines and \$4.01 billion in spending. 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.[17] This eliminates 5,090,024 claim lines and leaves us with \$3.24 billion in spending. 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. 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. 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.

B.2 Deriving the Bias in our Medicare Link Estimate

The estimate of $\hat{\beta}$ in equation (13) will be:

$$\begin{aligned}
\hat{\beta} &= \frac{\mathrm{Cov}[\Delta\overline{\ln(P_{g,j})}, \Delta\ln(RVU_j)]}{\mathrm{Var}[\Delta\ln(RVU_j)]} \\
&= \frac{\mathrm{Cov}[\sigma\Delta\overline{\ln(\phi_g)} + \sigma\Delta\ln(RVU_j) + (1-\sigma)\Delta\overline{\ln(\rho_{g,j})} + \Delta\epsilon_{g,j,p}, \Delta\ln(RVU_j)]}{\mathrm{Var}[\Delta\ln(RVU_j)]} \\
&= \sigma\frac{\mathrm{Cov}[\Delta\ln(RVU_j), \Delta\ln(RVU_j)]}{\mathrm{Var}[\Delta\ln(RVU_j)]} + \sigma\frac{\mathrm{Cov}[\Delta\overline{\ln(\phi_g)}, \Delta\ln(RVU_j)]}{\mathrm{Var}[\Delta\ln(RVU_j)]} \\
&+ (1-\sigma)\frac{\mathrm{Cov}[\Delta\overline{\ln(\rho_{g,j})}]}{\mathrm{Var}[\Delta\ln(RVU_j)]} + \frac{\mathrm{Cov}[\Delta\epsilon_{g,j,p}, \Delta\ln(RVU_j)]}{\mathrm{Var}[\Delta\ln(RVU_j)]} \\
&= \sigma + \sigma\frac{\mathrm{Cov}[\Delta\overline{\ln(\phi_g)}, \Delta\ln(RVU_j)]}{\mathrm{Var}[\Delta\ln(RVU_j)]} + (1-\sigma)\frac{\mathrm{Cov}[\Delta\overline{\ln(\rho_{g,j})}, \Delta\ln(RVU_j)]}{\mathrm{Var}[\Delta\ln(RVU_j)]},
\end{aligned} \tag{30}$$

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

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

C Comparing Our Cross-Sectional and 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.

C.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 (14). 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 6 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 BCBSTX'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

$$\begin{aligned}\overline{\ln(P_{c,g,j,t})} = {}& \psi \Delta \ln(RVU_j) \cdot \mathbb{1}_{\{t=\text{post}\}} + \xi \Delta \ln(RVU_j) \cdot \mathbb{1}_{\{t=\text{post}\}} \cdot \overline{|\hat{e}_j|} + \gamma \mathbb{1}_{\{t=\text{post}\}} \cdot \overline{|\hat{e}_j|} \\ &+ \phi_t \mathbb{1}_t + \phi_j \mathbb{1}_j + \phi_g \mathbb{1}_g + \phi_{g,j} \mathbb{1}_g \cdot \mathbb{1}_j + \epsilon_{c,g,j,t}. \end{aligned} \tag{31}$$

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 $\hat{\xi} = 0$.

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

Table C.2 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 C.1 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 (31), 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 C.1

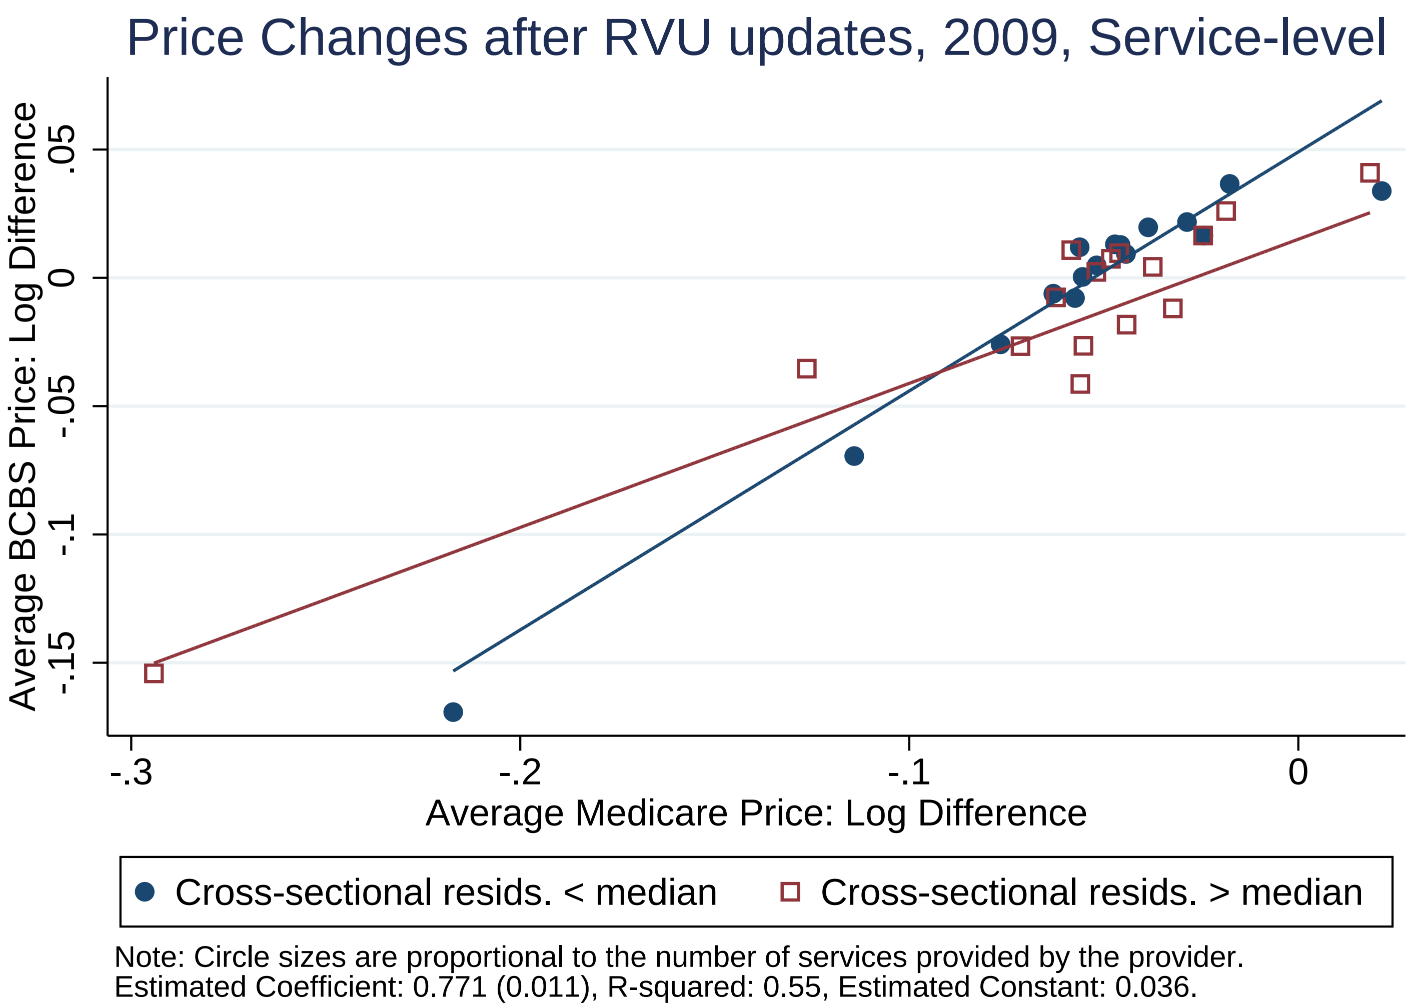

Note: The figure reports the relationship described by equation (12) for RVU updates in 2009, split into two sample based on the median prediction error from Figure 6. (The blue dots in Figure 6 correspond to the blue circles in this graph, while the yellow and red observations from Figure 6 correspond 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 C.1: **Alternative Measures of Pricing According to Common Implied Conversion Factors**

		Panel A: 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 B: 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 C: 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 3. 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 BCBSTX.

Appendix Table C.2: **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 (31) in section C.1. Panel A shows estimates using RBRVS updates and BCBSTX 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 BCBSTX.